

Homogeneous Codes for Energy-Efficient Illumination and Imaging

Matthew O'Toole
University of Toronto

Supreeth Achar
Carnegie Mellon University

Srinivasa G. Narasimhan
Carnegie Mellon University

Kiriakos N. Kutulakos
University of Toronto




Figure 1: *Imagine trying to acquire live video (30fps) of structured-light patterns as they are being projected onto a compact fluorescent bulb that has already been turned on (rated 1600 Lumens), or onto a face in bright sunlight (80 kLux)—with a 5-Lumen projector. One of our two prototypes, shown in (a), achieves this with an off-the-shelf laser projector and a CMOS camera with an ordinary lens. We then used it to capture the video frames in (b) and (c). We also show how to use our prototype to generate a live video feed from the projector's—rather than the camera's—point of view, shown in (d).*

Abstract

Programmable coding of light between a source and a sensor has led to several important results in computational illumination, imaging and display. Little is known, however, about how to utilize energy most effectively, especially for applications in live imaging. In this paper, we derive a novel framework to maximize energy efficiency by “homogeneous matrix factorization” that respects the physical constraints of many coding mechanisms (DMDs/LCDs, lasers, *etc.*). We demonstrate energy-efficient imaging using two prototypes based on DMD and laser illumination. For our DMD-based prototype, we use fast local optimization to derive codes that yield brighter images with fewer artifacts in many transport probing tasks. Our second prototype uses a novel combination of a low-power laser projector and a rolling shutter camera. We use this prototype to demonstrate never-seen-before capabilities such as (1) capturing live structured-light video of very bright scenes—even a light bulb that has been turned on; (2) capturing epipolar-only and indirect-only live video with optimal energy efficiency; (3) using a low-power projector to reconstruct 3D objects in challenging conditions such as strong indirect light, strong ambient light, and smoke; and (4) recording live video from a projector's—rather than the camera's—point of view.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Imaging geometry, radiometry

Keywords: energy efficiency, low-power imaging, coded illumination, coded exposure, computational photography, 3D scanning

ACM Reference Format

O'Toole, M., Achar, S., Narasimhan, S., Kutulakos, K. 2015. Homogeneous Codes for Energy-Efficient Illumination and Imaging. *ACM Trans. Graph.* 34, 4, Article 35 (August 2015), 13 pages. DOI = 10.1145/2766897 <http://doi.acm.org/10.1145/2766897>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766897>

1 Introduction

When we capture an image under controlled lighting, the power of the light source matters a lot: all things being equal, brighter sources will send more photons to the sensor during an exposure, producing a brighter and less noisy image. The brightness of the source, however, is just one way to control how much light reaches the sensor of a computational imaging system. Modern systems use an arrangement of devices to transport light from a source to the scene (or from the scene to sensor) and these devices are often programmable—galvanometers [Mertz et al. 2012], digital micro-mirror devices [Nayar et al. 2004; Hitomi et al. 2011], liquid-crystal panels [Raskar et al. 2006], phase modulators [Damberg and Heidrich 2015], *etc.* This brings up a natural question: how should we program the spatio-temporal behavior of these devices to maximize an arrangement's *energy efficiency*, *i.e.*, the energy that can be transmitted from the source to the sensor for a given imaging task, power, and exposure time?

Studies of this problem began in the 1960s for the special case of arrangements with just three “active” components: a light source that is always turned on, a light-blocking mask that is controlled by a binary code, and a sensor [Ibbett et al. 1968; Decker and Harwit 1969]. The optimal sequence of codes for this case is derived from the Hadamard matrix [Harwit and Sloane 1979] and enjoys widespread use [Schechner et al. 2007], mainly because most conventional computational imaging systems are arranged this way.

Despite the ubiquity of these arrangements, two general principles have emerged in recent years. On one hand, masks are inefficient because they waste energy whenever they block a photon [Hoskinson et al. 2010]. On the other hand, masks confer unique abilities when arranged in layers along an optical path and programmed to change repeatedly in a single exposure. This is because photons can be blocked far more selectively this way, enabling light field displays [Lanman et al. 2010; Wetzstein et al. 2012], indirect-only photography [O'Toole et al. 2012], and several other imaging functionalities [O'Toole et al. 2014].

Unfortunately, neither the original Hadamard multiplexing theory nor its recent extensions [Cossairt et al. 2012; Mitra et al. 2014a; Mitra et al. 2014b] apply to multi-layer arrangements, or to arrangements that avoid masks altogether (*e.g.*, laser-based projectors [Damberg et al. 2014; Gupta et al. 2013]). As a result, the problem of computing energy-efficient codes for multi-layer ar-

Doppler Time-of-Flight Imaging

Felix Heide^{1,2,3} Wolfgang Heidrich^{2,1} Matthias Hullin⁴ Gordon Wetzstein³
¹University of British Columbia ²KAUST ³Stanford University ⁴University of Bonn




Figure 1: We introduce a new computational imaging system that allows for metric radial velocity information to be captured instantaneously for each pixel (center row). For this purpose, we design the temporal illumination and modulation frequencies of a time-of-flight camera (left) to be orthogonal within its exposure time. The Doppler effect of objects in motion is then detected as a frequency shift of the illumination, which results in a mapping from object velocity to recorded pixel intensity. By capturing a few coded time-of-flight measurements and adding a conventional RGB camera to the setup, we demonstrate that color, velocity, and depth information of a scene can be recorded simultaneously. The results above show several frames of two video sequences. For each example, the left-most frame shows a static object (velocity map is constant), which is then moved towards (positive radial velocity) or away from (negative velocity) the camera.

Abstract

Over the last few years, depth cameras have become increasingly popular for a range of applications, including human-computer interaction and gaming, augmented reality, machine vision, and medical imaging. Many of the commercially-available devices use the time-of-flight principle, where active illumination is temporally coded and analyzed in the camera to estimate a per-pixel depth map of the scene. In this paper, we propose a fundamentally new imaging modality for all time-of-flight (ToF) cameras: per-pixel radial velocity measurement. The proposed technique exploits the Doppler effect of objects in motion, which shifts the temporal illumination frequency before it reaches the camera. Using carefully coded illumination and modulation frequencies of the ToF camera, object velocities directly map to measured pixel intensities. We show that a slight modification of our imaging system allows for color, depth, and velocity information to be captured simultaneously. Combining the optical flow computed on the RGB frames with the measured metric radial velocity allows us to further estimate the full 3D metric velocity field of the scene. The proposed technique has applications in many computer graphics and vision problems, for example motion tracking, segmentation, recognition, and motion deblurring.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning

Keywords: computational photography, time-of-flight

ACM Reference Format

Heide, F., Heidrich, W., Hullin, M., Wetzstein, G. 2015. Doppler Time-of-Flight Imaging. *ACM Trans. Graph.* 34, 4, Article 36 (August 2015), 11 pages. DOI = 10.1145/2766953 <http://doi.acm.org/10.1145/2766953>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766953>

1 Introduction

Pioneers of photography, including Eadweard Muybridge and Harold “Doc” Edgerton, advanced imaging technology to reveal otherwise invisible motions of high-speed events. Today, understanding the motion of objects in complex scenes is at the core of computer vision, with a wide range of applications in object tracking, segmentation, recognition, motion deblurring, navigation of autonomous vehicles, and defense. Usually, object motion or motion parallax are estimated via optical flow [Horn and Schunck 1981]: recognizable features are tracked across multiple video frames. The computed flow field provides the basis for many computer vision algorithms, including depth estimation. Unfortunately, optical flow is computationally expensive, fails for untextured scenes that do not contain good features to track, and it only measures 2D lateral motion perpendicular to the camera’s line of sight. Further, the unit of optical flow is pixels; metric velocities cannot be estimated unless depth information of the scene is also available. For the particular application of depth estimation, many limitations of optical flow estimation can be overcome using active illumination, as done by most structured illumination and time-of-flight (ToF) cameras. With the emergence of RGB-D imaging, for example facilitated by Microsoft’s Kinect One¹, complex and untextured 3D scenes can be tracked by analyzing both color and depth information, resulting in richer visual data that has proven useful for many applications.

In this paper, we introduce a new approach to directly imaging radial object velocity. Our approach analyzes the Doppler effect in time-of-flight cameras: object motion towards or away from the camera shifts the temporal illumination frequency before it is recorded. Conventional time-of-flight cameras encode phase information (and therefore scene depth) into intensity measurements. Instead, we propose Doppler Time-of-Flight (D-ToF) as a new imaging mode, whereby the change of illumination frequency (corresponding to radial object velocity) is directly encoded into the measured intensity. The required camera hardware is the same as for

¹microsoft.com/en-us/kinectforwindows/

Micron-scale Light Transport Decomposition Using Interferometry

Ioannis Gkioulekas
Harvard University

Anat Levin
Weizmann Institute

Frédéric Durand
Massachusetts Institute of Technology

Todd Zickler
Harvard University

Abstract

We present a computational imaging system, inspired by the optical coherence tomography (OCT) framework, that uses interferometry to produce decompositions of light transport in small scenes or volumes. The system decomposes transport according to various attributes of the paths that photons travel through the scene, including where on the source the paths originate, their pathlengths from source to camera through the scene, their wavelength, and their polarization. Since it uses interference, the system can achieve high pathlength resolutions, with the ability to distinguish paths whose lengths differ by as little as ten microns. We describe how to construct and optimize an optical assembly for this technique, and we build a prototype to measure and visualize three-dimensional shape, direct and indirect reflection components, and properties of scattering, refractive/dispersive, and birefringent materials.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Imaging geometry;

Keywords: light transport, interference, wave optics

1 Introduction

In imaging, photons leave a source, travel through a scene, and are collected by a camera. A conventional image measures the sum of all photons that arrive at each camera pixel, regardless of where on the source they originate, or which path they travel from there to the camera (Figure 1). This summation conflates information about the shapes and materials in the scene. In this paper, we show how interferometry can be used to measure decompositions of these per-pixel sums, distinguishing between photon paths that differ in their length and location of origin.

In discrete terms, where the camera sensor is partitioned into P area elements (“pixels”) and the source is partitioned into L area elements, the energy that is transferred from source to camera is often described by a $P \times L$ matrix \mathbf{T} called the *light transport matrix*. Each entry T_{pl} in this matrix represents the fraction of photons following paths originating at the l th source element and arriving at the p th camera element. With reference to Figure 1, orange versus blue paths would correspond to entries T_{pl} for different values of source location l . Conventional imaging simply measures the image \mathbf{i} (a P -vector) for some pattern \mathbf{l} (an L -vector) projected from the source according to the *light transport equation* [Ng et al. 2003]

$$\mathbf{i} = \mathbf{T}\mathbf{l}. \quad (1)$$

There are a variety of computational imaging methods for measuring elements of the transport matrix. This has applications in image-based rendering, image editing, and measuring scene shape in the presence of translucency and interreflections.

Each element of the transport matrix in Equation (1) still represents a sum of different scene paths. As depicted by the blue paths of Figure 1, sub-surface scattering and interreflections mean that there are typically many paths that originate at the same source element and arrive at the same camera element but take different routes through the scene. Computational imaging can be used to analyze these as well, by decomposing each entry T_{pl} of the transport matrix according to the contributions from photon paths that have different optical lengths τ . This leads to the notion of a *pathlength-resolved light transport matrix* \mathbf{T}^τ , $\tau \in \{\tau_{\min}, \dots, \tau_{\max}\}$, where pathlengths are discretized into a finite set of $\Delta\tau$ -sized length intervals. The full light transport matrix can then be written as

$$\mathbf{T} = \sum_{\tau} \mathbf{T}^\tau. \quad (2)$$

Each entry T_{pl}^τ is the fraction of photons originating at the l th source element that arrive at the p th camera element having traveled paths whose optical lengths are in the interval $\tau \pm \Delta\tau/2$. There is a growing number of methods for measuring either entire projections or isolated elements of the pathlength-resolved transport matrix. In addition to enhancing the applications listed above, these decompositions can be used to visualize light-in-flight through table-top scenes, and for “imaging around corners.”

We introduce a new computational imaging system that uses optical interferometry to produce high-fidelity light transport decompositions. Our system uses optical configurations that are variations of the classical Michelson interferometer, and our analysis builds on techniques that have been used for optical coherence tomography (OCT). Our system is complementary to existing computational photography methods for producing such decompositions, excelling in use cases where it is necessary to image small scenes at very high spatial and pathlength resolutions. In particular, the use of interferometry allows our system to achieve pathlength resolutions as low as $10\ \mu\text{m}$, which is necessary to analyze transport events caused by material effects like dispersion and scattering.

Our paper begins with background on the Michelson interferometer and the notions of spatial and temporal coherence length. We then introduce a mathematical model of interferometry in terms of (a continuous version of) the pathlength-resolved light transport matrix. We use it to show how sources with different coherence properties enable different kinds of light transport decompositions, differentiating light paths in terms of their endpoint locations, their optical lengths, or combinations of these two. We also characterize resolution and noise performance, and present a performance-optimized optical design that additionally allows resolving transport in terms of wavelength and polarization. Our prototype has three spectral channels, two polarization channels, a working volume of $2\text{ cm H} \times 2\text{ cm W} \times 1\text{ cm D}$, and spatial and pathlength resolutions that are both $10\ \mu\text{m}$. We use this prototype to obtain micron-scale decompositions of light transport in scenes containing reflection, refraction, dispersion, scattering, and birefringence.

2 Related Work

Light transport decomposition. Methods for decomposing light transport, both spatially and in terms of pathlength, can be described as capturing different images \mathbf{i} of a static scene under uniform illumination $\mathbf{l} = \mathbf{1}$ according to the equation

$$\mathbf{i} = \sum_{\tau} w(\tau)(\mathbf{M} \odot \mathbf{T}^\tau)\mathbf{1}. \quad (3)$$

ACM Reference Format

Gkioulekas, I., Levin, A., Durand, F., Zickler, T. 2015. Micron-Scale Light Transport Decomposition Using Interferometry. *ACM Trans. Graph.* 34, 4, Article 37 (August 2015), 14 pages. DOI = 10.1145/2766928 <http://doi.acm.org/10.1145/2766928>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766928>

Integrable PolyVector Fields

Olga Diamanti*
ETH Zurich

Amir Vaxman
Vienna University of Technology

Daniele Panozzo
ETH Zurich

Olga Sorkine-Hornung
ETH Zurich

Abstract

We present a framework for designing curl-free tangent vector fields on discrete surfaces. Such vector fields are gradients of locally-defined scalar functions, and this property is beneficial for creating surface parameterizations, since the gradients of the parameterization coordinate functions are then exactly aligned with the designed fields. We introduce a novel definition for discrete curl between unordered sets of vectors (PolyVectors), and devise a curl-eliminating continuous optimization that is independent of the matchings between them. Our algorithm naturally places the singularities required to satisfy the user-provided alignment constraints, and our fields are the gradients of an inversion-free parameterization by design.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: PolyVectors, curl-free fields, quad meshing

1 Introduction

The design of vector fields on discrete surfaces is an important problem in geometry processing, with applications in surface parameterization, texture synthesis, stylized rendering, remeshing and architectural geometry [Hertzmann and Zorin 2000; Lefebvre and Hoppe 2006; Palacios and Zhang 2007; Li et al. 2011; Liu et al. 2011; Bommes et al. 2013b]. In most of these applications, vector fields are computed to serve as a guiding basis for the construction of global parameterizations, which is also the focus of this paper.

Vector fields are typically designed by prescribing a small set of user-defined alignment constraints, which are then interpolated. The resulting field is prescribed as the desired gradient of a parameterization function, and the parameterization is computed by minimizing the difference between the tangent field and the gradient of the function in the least-squares sense, leading to a Poisson equation. In this optimization problem, perfect alignment of the gradient of the resulting parameterization function to the sparse set of constraints can be enforced using linear equality constraints [Kälberer et al. 2007; Bommes et al. 2009]. However, the computation of the function by a least-squares solution may introduce unbounded errors away from the constraints and compromise local injectivity, resulting in inverted elements. Inverted and degenerate elements make the global parameterization unsuitable for practical purposes (e.g., they lead to holes during remeshing [Ebke et al. 2013]). Previous works avoid




Figure 1: Parameterization using integrable fields. Given a set of user-provided constraints (left), we optimize for an integrable field, which exactly aligns to the corresponding parameterization (right), while staying close to the constraints.

them using heuristic weighting [Bommes et al. 2009] or relaxation by a reduction of the search space [Bommes et al. 2013a]. All these methods further increase the parameterization distortion.

The key insight that guides our approach is that the gradients of a global parameterization, with or without inversions, are curl-free vector fields. We thus directly devise conditions that guarantee that the final vector field will be curl-free, and that the corresponding parameterization will additionally be free of inverted elements. We propose an algorithm that penalizes the violation of these conditions. Using our algorithm automatically places the singularities in a way that minimizes the distortions.

We formally introduce the notion of integrability and curl for PolyVectors, which are unordered sets of tangent vectors defined on every face of the surface. We develop a formula to measure the PolyVector curl, and optimize for its removal; our formulation does not require explicit pairings of vectors between the unordered sets on adjacent elements. In this manner, the individual vector pairings (often called matchings) are free to change during the optimization, naturally inducing field singularities. Our algorithm relies on a nonlinear optimization, involving no integer variables, and supports user-provided directional constraints, which may be free to align with either the u or the v coordinate of the parameterization, without the necessity to constrain such a choice.

The conditions that ensure an inversion-free parameterization are not imposed as hard constraints by our algorithm, but as a penalization term that tends to avoid invalid configurations. This term is then minimized using a nonlinear and nonconvex optimization. Thus our algorithm is theoretically not guaranteed to find a global minimum, i.e., a fully inversion-free parameterization. However, we demonstrate its robustness and applicability by running it on a global parameterization benchmark proposed in [Myles et al. 2014], which contains 116 models. Our algorithm successfully produces integrable fields on all models, and the induced parameterizations contain no inversions.

*e-mail:olga.diamanti@inf.ethz.ch

ACM Reference Format

Diamanti, O., Vaxman, A., Panozzo, D., Sorkine-Hornung, O. 2015. Integrable PolyVector Fields. ACM Trans. Graph. 34, 4, Article 38 (August 2015), 12 pages. DOI = 10.1145/2766906
<http://doi.acm.org/10.1145/2766906>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766906>

Stripe Patterns on Surfaces

Felix Knöppel
TU Berlin

Keenan Crane
Columbia University

Ulrich Pinkall
TU Berlin

Peter Schröder
Caltech

Abstract

Stripe patterns are ubiquitous in nature, describing macroscopic phenomena such as stripes on plants and animals, down to material impurities on the atomic scale. We propose a method for synthesizing stripe patterns on triangulated surfaces, where singularities are automatically inserted in order to achieve user-specified orientation and line spacing. Patterns are characterized as global minimizers of a convex-quadratic energy which is well-defined in the smooth setting. Computation amounts to finding the principal eigenvector of a symmetric positive-definite matrix with the same sparsity as the standard graph Laplacian. The resulting patterns are globally continuous, and can be applied to a variety of tasks in design and texture synthesis.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: discrete differential geometry, digital geometry processing, texture synthesis, direction fields, singularities

1 Introduction

A diverse collection of natural phenomena exhibit the same characteristic pattern: unoriented stripes of uniform width that bifurcate at isolated points called *edge dislocations* [Kalpakjian and Schmid 2009, pp. 44–46]. Why does this bifurcation occur? Surfaces of revolution (Fig. 3) provide an instructive example: for any given stripe width there is a maximum integer number of stripes that can fit around the circumference at each height h . As the radius r grows or shrinks this number can suddenly jump—stripes must therefore branch or coalesce in order to maintain a continuous transition between neighboring regions.



Figure 1: In general, patterns without singularities exhibit uneven spacing, and patterns with even spacing must have singularities.

ACM Reference Format

Knöppel, F., Crane, K., Pinkall, U., Schröder, P. 2015. Stripe Patterns on Surfaces. ACM Trans. Graph. 34, 4, Article 39 (August 2015), 11 pages. DOI = 10.1145/2767000 <http://doi.acm.org/10.1145/2767000>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2767000>



(Photo courtesy Joshua Hill)

Figure 2: Natural phenomena like cacti (left) exhibit characteristic branching patterns in an effort to maintain evenly-spaced features. Right: Our algorithm allows one to quickly and automatically synthesize similar patterns that seamlessly cover arbitrary surfaces.

Our approach to stripe pattern synthesis is closely related to existing methods for field-aligned parameterization (Sec. 1.1). The main point of departure is that existing algorithms try to *prevent* singularities that were not specified as part of the input, whereas we intentionally *allow* new singularities, which are essential to modeling natural phenomena. In particular, we use a formulation akin to Ray *et al.* [2006], but omit both the curl correction step and the unit-norm constraint. As a result, we can formulate our problem via a simple convex quadratic energy; the practical payoff is an algorithm about an order of magnitude faster than those based on nonconvex constraints or mixed integer programming (Sec. 5.1). The final algorithm is simple to implement (App. C) requiring no global cutting or integer jumps, is robust to errors in the input (Fig. 16), and requires little work beyond factoring a single sparse matrix. The output is an angle α at each triangle corner which can be used to drive a periodic texture or shader. We also introduce a novel interpolation scheme (Sec. 4.3) that ensures the pattern is globally continuous (C^0) away from a finite collection of isolated singular points.




Figure 3: To maintain constant stripe width, each “level” of this surface of revolution must exhibit a different number of stripes. Judiciously placed branch points are therefore essential to keeping the pattern continuous. Inset: each pair of stripes from the second level perfectly branches into three.

Frame Field Generation through Metric Customization

Tengfei Jiang

Xianzhong Fang

Jin Huang*

Hujun Bao*

Yiying Tong

Mathieu Desbrun

State Key Lab of CAD&CG, Zhejiang University

Michigan State University

Caltech




Figure 1: Overview. Frame field design over surfaces is typically guided by constraints on (an)isotropy, lengths, and feature alignments (left). While current methods (e.g., MIQ [Bommes et al. 2009]) use the Euclidean-induced metric to find an appropriate frame (and further derive a quadrangulation, top), our method (bottom) first customizes a Riemannian metric from the design requirements, then generates a frame field that is a smooth cross field in this new metric—thus offering a more flexible and more intuitive design process producing high quality frame fields. Metrics are visualized using their local unit circles (middle left); the resulting frame’s sizes and directions are indicated as color isocurves (middle right); and a quadrangulation based on the frame field is then automatically extracted (right).

Abstract

This paper presents a new technique for frame field generation. As generic frame fields (with arbitrary anisotropy, orientation, and sizing) can be regarded as cross fields in a specific Riemannian metric, we tackle frame field design by first computing a discrete metric on the input surface that is compatible with a sparse or dense set of input constraints. The final frame field is then found by computing an optimal cross field in this customized metric. We propose frame field design constraints on alignment, size, and skewness at arbitrary locations on the mesh as well as along feature curves, offering much improved flexibility over previous approaches. We demonstrate the advantages of our frame field generation through the automatic quadrangulation of man-made and organic shapes with controllable anisotropy, robust handling of narrow surface strips, and precise feature alignment. We also extend our technique to the design of n -vector fields.

CR Categories: I.3.3 [Computer Graphics]: Computational Geometry & Object Modeling—Curve & surface representations.

Keywords: Frame field generation, metric customization, n -vector field design, geometry processing.

*Corresponding authors: hj@cad.zju.edu.cn, bao@cad.zju.edu.cn

ACM Reference Format

Jiang, T., Fang, X., Huang, J., Bao, H., Tong, Y., Desbrun, M. 2015. Frame Field Generation through Metric Customization. ACM Trans. Graph. 34, 4, Article 40 (August 2015), 11 pages. DOI = 10.1145/2766927 <http://doi.acm.org/10.1145/2766927>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766927>

1 Introduction

Frame fields on discrete 2-manifolds are key to a wide range of applications in computer graphics, including texture mapping and quadrangulation remeshing, as they induce a smoothly-varying assignment of a basis for all the tangent spaces of the surface. While orthonormal frame fields (also called 4-symmetry direction fields or cross fields) are most common due to the convenience afforded by unit and orthogonal bases, the need for non-uniform, anisotropic, and skewed frame fields has grown tremendously in recent years due to their versatility in geometry processing.

However, the design of generic frame fields remains a challenging task: one has to balance user-input or application-driven requirements (such as alignment to given directions and sizing constraints) with the resulting smoothness of the frame field and its inevitable singularities. Depending on the targeted application, input constraints can be either under- or over-constrained, as smoothness and direction constraints may conflict and interfere with sizing constraints, and vice-versa. This renders any user-guided design process long and tedious: automatic frame field generation can result in inaccurate alignments, requiring constraints to be interactively added or removed until the final frame field is acceptable.

In this paper, we argue that frame field design is best accomplished by first designing a non-Euclidean metric on the surface based on the input constraints, and then deriving the frame field as an orthonormal frame field in this computed metric. Abandoning the viewpoint that the surface is embedded in Euclidean space and considering instead the more general case of Riemannian manifolds significantly increases flexibility in the design of frame field, resulting in a tighter control over directional and sizing constraints.

1.1 Related work

We first review the key previous works that we will build upon or extend in our metric-based approach to frame field generation.

Computational Bodybuilding: Anatomically-based Modeling of Human Bodies

Shunsuke Saito
University of Pennsylvania, Waseda University

Zi-Ye Zhou
University of Pennsylvania

Ladislav Kavan
University of Pennsylvania




Figure 1: Given an input 3D anatomy template, we propose a system to simulate the effects of muscle, fat, and bone growth. This allows us to create a wide range of human body shapes.

Abstract

We propose a method to create a wide range of human body shapes from a single input 3D anatomy template. Our approach is inspired by biological processes responsible for human body growth. In particular, we simulate growth of skeletal muscles and subcutaneous fat using physics-based models which combine growth and elasticity. Together with a tool to edit proportions of the bones, our method allows us to achieve a desired shape of the human body by directly controlling hypertrophy (or atrophy) of every muscle and enlargement of fat tissues. We achieve near-interactive run times by utilizing a special quasi-statics solver (Projective Dynamics) and by crafting a volumetric discretization which results in accurate deformations without an excessive number of degrees of freedom. Our system is intuitive to use and the resulting human body models are ready for simulation using existing physics-based animation methods, because we deform not only the surface, but also the entire volumetric model.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics—Animation

Keywords: human body modeling, 3D anatomy, growth models, physics-based animation.

ACM Reference Format

Saito, S., Zhou, Z., Kavan, L. 2015. Computational Bodybuilding: Anatomically-based Modeling of Human Bodies. *ACM Trans. Graph.* 34, 4, Article 41 (August 2015), 12 pages. DOI = 10.1145/2766957 <http://doi.acm.org/10.1145/2766957>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766957>

1 Introduction

Human bodies exhibit large variations in size and shape due to factors such as height, muscularity, or adiposity. Although modern animations systems such as Weta Tissue can bring realistic digital humans to life by combining physics-based simulation with detailed models of 3D anatomy, creating simulation-ready anatomical models is expensive and typically involves a team of specialized digital artists experienced in 3D modeling, simulation, and anatomy. This *de facto* limits the applicability of realistic human body simulations to high-budget productions.

In this paper, we propose a method to generate a wide range of human body shapes which are ready to be simulated using existing physics-based methods. Our approach is intuitive even to inexperienced users, because it is motivated by natural biological processes. Inspired by growth laws studied in biomechanics [Taber 1995; Taber 1998; Jones and Chapman 2012; Wisdom et al. 2015], we propose mathematical models of hypertrophy and atrophy of skeletal muscles. We devise another model for growth of fat tissues. Our growth models are tightly coupled with a physics-based simulation system, which accounts for elasticity of the individual organs, e.g., as the biceps muscle grows due to hypertrophy, it exerts forces on the adjacent tissues, pushing them out of the way. We combine our models of soft tissue growth with a geometric shape deformation technique to edit the rest pose of the body, allowing us to produce human bodies with varied anthropometry, such as heights and bone lengths.

Our goal is fundamentally different from existing simulators [Lee et al. 2009] which focus on modeling the *motion* of the human body such as walking or lifting an object. Even though our method shares similar building blocks, e.g., volumetric elasticity and finite element methods, we simulate changes of the human body over the *long term*. Growth processes such as hypertrophy do not occur instantaneously and lead to quite different shapes than, e.g., muscle contractions (Figure 2). On a mathematical level, our growth mod-

Biomechanical Simulation and Control of Hands and Tendinous Systems

Prashant Sachdeva^{1*} Shinjiro Sueda^{2*} Susanne Bradley¹ Mikhail Fain¹ Dinesh K. Pai¹
¹University of British Columbia ²California Polytechnic State University




Figure 1: Our biomechanical simulation and control framework can model the human hand performing tasks such as writing (a-b), and typing on a keyboard (c). We can also simulate clinical conditions such as boutonnière deformity (d) by cutting a tendon insertion.

Abstract

The tendons of the hand and other biomechanical systems form a complex network of sheaths, pulleys, and branches. By modeling these anatomical structures, we obtain realistic simulations of coordination and dynamics that were previously not possible. First, we introduce Eulerian-on-Lagrangian discretization of tendon strands, with a new selective quasistatic formulation that eliminates unnecessary degrees of freedom in the longitudinal direction, while maintaining the dynamic behavior in transverse directions. This formulation also allows us to take larger time steps. Second, we introduce two control methods for biomechanical systems: first, a general-purpose learning-based approach requiring no previous system knowledge, and a second approach using data extracted from the simulator. We use various examples to compare the performance of these controllers.

CR Categories: I.6.8 [Simulation and Modeling]: Types of Simulation—Combined

Keywords: muscles, tendons, hands, physically based simulation, constrained strands, Lagrangian mechanics

1 Introduction

Although simulations of hands and grasp have consistently received attention in the graphics community, modeling and simulating the dynamics of the complex tendon network of the hand has remained relatively unexplored. Much of the previous simulation techniques for the hand have been based on rigid links with either joint torques [Pollard and Zordan 2005; Kry and Pai 2006; Liu 2009] or line-of-force muscles [Albrecht et al. 2003; Tsang et al. 2005; Johnson et al. 2009]. In the real hand, however, the motion of the digital

phalanges are driven by muscles originating in the forearm acting through tendons passing through a complex network of sheaths and pulleys [Valero-Cuevas et al. 2007]. This design has an important effect on the compliance and coupling of the fingers.

In our approach, we model each tendon as a physical primitive rather than using joint torques or moment arms. This is advantageous because it allows us to properly model the complexity of the tendon network, which we believe is important for obtaining realistic motions. As an added bonus, the biomechanical modeling of tendons and ligaments can give us proper joint coupling for free. As a simple example, let us consider the coordinated motion of the joints, shown in Fig. 7. The two distal joints of the finger (PIP and DIP) flex and extend in a coordinated fashion *not* because of the synchronized activation signals computed by the brain, but because of the arrangement of tendons and ligaments in the finger. In our simulator, pulling on a single tendon (flexor digitorum profundus) achieves this result, whereas with a torque-based approach, we would need to manually coordinate the torques at these two joints. Another example is in the coupling between the extensor tendons in the back of the hand. Although this lack of complete independence between the fingers is partly due to the neural control, mechanical coupling has been hypothesized as having a significant role [Lang and Schieber 2004]. We are also able to simulate hand deformities, which have applications not only in surgical planning and medicine, but also in computer graphics (Figs. 1(d) and 8). Some virtual character designs are based on deformities or injuries, and the ability to procedurally produce anatomically-based abnormal characters may prove useful, since obtaining real-world data of such characters can be a challenge.

We also introduce two methods for control of biomechanical systems. The first is a general-purpose learning method requiring no previous knowledge of the inner workings of the system and is suitable for control of any “black box” simulator. The second controller extracts controller parameters from the internals of the simulator.

Contributions We address three main issues that are especially important for biomechanical simulation. First, by applying the Eulerian-on-Lagrangian discretization of the tendon strand [Sueda et al. 2011], we greatly simplify the contact handling between tendons and bones. Second, we develop a new formulation to deal with highly stiff strands, by assuming that strain and stress propagate instantaneously through the strand, allowing us to use large time steps even for stiff tendons (§3.2). Third, we introduce and assess two methods for control of these systems (§4.2, §4.3).

*Equal contributions.

ACM Reference Format

Sachdeva, P., Sueda, S., Bradley, S., Fain, M., Pai, D. 2015. Biomechanical Simulation and Control of Hands and Tendinous Systems. ACM Trans. Graph. 34, 4, Article 42 (August 2015), 10 pages. DOI = 10.1145/2766987 http://doi.acm.org/10.1145/2766987.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: http://dx.doi.org/10.1145/2766987

GRIDiron: An interactive authoring and cognitive training foundation for reconstructive plastic surgery procedures

Nathan Mitchell
University of Wisconsin-Madison

Court Cutting
New York University

Eftychios Sifakis
University of Wisconsin-Madison




Figure 1: Simulated *Dufourmentel-Mouly* repair (see [Baker 2014]) for a large gap of excised tissue on the scalp. From top to bottom, left to right: Rendered stages of procedure, embedding lattice, real time demo on a tablet running in a web browser.

Abstract

We present an interactive simulation framework for authoring surgical procedures of soft tissue manipulation using physics-based simulation to animate the flesh. This interactive authoring tool can be used by clinical educators to craft three-dimensional illustrations of the intricate maneuvers involved in craniofacial repairs, in contrast to two-dimensional sketches and still photographs which are the medium used to describe these procedures in the traditional surgical curriculum. Our virtual environment also allows surgeons-in-training to develop cognitive skills for craniofacial surgery by experimenting with different approaches to reconstructive challenges, adapting stock techniques to flesh regions with nonstandard shape, and reach preliminary predictions about the feasibility of a given repair plan. We use a Cartesian grid-based embedded discretization of nonlinear elasticity to maximize regularity, and expose opportunities for aggressive multithreading and SIMD accelerations. Using a grid-based approach facilitates performance and scalability, but constrains our ability to capture the topology of thin surgical incisions. We circumvent this restriction by hybridizing the grid-based discretization with an explicit hexahedral mesh representation in regions where the embedding mesh necessitates overlap or nonmanifold connectivity. Finally, we detail how the front-end of our system can run on lightweight clients, while the core simulation capability can be hosted on a dedicated server and delivered as a network service.

CR Categories: I.3.5 [Computer Graphics]: Computational geometry and modeling—[Physically based modeling]

Keywords: Virtual surgery, finite elements, elasticity.

1 Introduction

The art and practice of plastic surgery is intimately tied with a lifelong learning experience both in establishing theoretical foundations, such as anatomy and pathology, as well as in acquiring and sharpening surgical skills. Although the requisite skill set overlaps with that of general surgery at large, plastic surgery has a distinct focus on the utilization of topology change as a treatment mechanism. Fortunately, a few opportunities exist for surgeons to acquire and develop these skills outside of the operating room. Some options involve operating on physical proxies, such as phantom materials, cadavers or animal tissues. Alternatively, computer technology can be leveraged to provide a non-invasive training testbed. The field of computer graphics is well positioned to contribute to this activity, given its vested interest in modeling virtual materials and digital models of anatomy. This application area, however, presents a unique set of challenges which may be uncharacteristic for traditional graphics applications. In particular, clinical measures of success (does the product improve quality of patient care?) may well differ from visual quality metrics in graphics. Additionally, given that computer-based surgical training is still an exploratory proposition, practical solutions need to be concerned with longevity, extensibility, deployment cost and ease of use.

ACM Reference Format

Mitchell, N., Cutting, C., Sifakis, E. 2015. GRIDiron: An interactive authoring and cognitive training foundation for reconstructive plastic surgery procedures. *ACM Trans. Graph.* 34, 4, Article 43 (August 2015), 12 pages. DOI = 10.1145/2766918 <http://doi.acm.org/10.1145/2766918>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766918>

Detailed Spatio-Temporal Reconstruction of Eyelids

Amit Bermanno^{1,2} Thabo Beeler¹ Yeara Kozlov^{1,2} Derek Bradley¹ Bernd Bickel^{1,3} Markus Gross^{1,2}

1) Disney Research Zurich 2) ETH Zurich 3) IST Austria




Figure 1: Our method extends high-resolution facial performance capture with a reconstruction approach that targets eyelids. We produce detailed, spatio-temporal eyelid reconstructions, even during complex deformation and folding that occur in the eye region. The result can be used to create high-fidelity digital doubles, as shown on the right.

Abstract

In recent years we have seen numerous improvements on 3D scanning and tracking of human faces, greatly advancing the creation of digital doubles for film and video games. However, despite the high-resolution quality of the reconstruction approaches available, current methods are unable to capture one of the most important regions of the face - the eye region. In this work we present the first method for detailed spatio-temporal reconstruction of eyelids. Tracking and reconstructing eyelids is extremely challenging, as this region exhibits very complex and unique skin deformation where skin is folded under while opening the eye. Furthermore, eyelids are often only partially visible and obstructed due to self-occlusion and eyelashes. Our approach is to combine a geometric deformation model with image data, leveraging multi-view stereo, optical flow, contour tracking and wrinkle detection from local skin appearance. Our deformation model serves as a prior that enables reconstruction of eyelids even under strong self-occlusions caused by rolling and folding skin as the eye opens and closes. The output is a person-specific, time-varying eyelid reconstruction with anatomically plausible deformations. Our high-resolution detailed eyelids couple naturally with current facial performance capture approaches. As a result, our method can largely increase the fidelity of facial capture and the creation of digital doubles.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning;

Keywords: Eyelid reconstruction, facial performance capture, eyelid modeling.

1 Introduction

The human face is the most important part of a person for conveying identity and emotion and therefore of central interest when creating realistic digital humans for computer games and films. Even subtle face motions can reveal information about the internal state of a person, where their attention is focused, the intention of actions, and interpretation of (non-)verbal communication. In recent years, we have witnessed tremendous progress in facial scanning, performance capture, and skin rendering. Stunning attempts at achieving a photorealistic actor such as *The Curious Case of Benjamin Button* or *The Digital Ira Project* [Alexander et al. 2013] are examples that high-resolution face scanning processes are key to cross the “uncanny valley” that divides a synthetic-looking face from a photo-realistic virtual person.

For identifying emotions, humans mainly use a consistent selective sampling of visual information from the eye region and, to a lesser extent, the mouth region [Smith et al. 2005; Peterson and Eckstein 2012]. Subtle details such as the twitch of an eyelid and the formation of small wrinkles significantly contribute to the realism of human faces and the perception of emotions. However, despite the important role of the eye region, existing capture technology is usually not able to provide an adequate level of geometric detail and motion to reproduce these subtleties. In practice, achieving realistic eyelid motions and skin deformation of the surrounding area requires significant manual modeling efforts by highly skilled artists.

Acquiring this region is extremely challenging due to several reasons. In an expressive performance, eyelids undergo extreme deformations and wrinkling. The skin rolls and folds inward when the eye opens, and stretches over the eyeball when the eye is shut. Due to concavities and eyelashes, there is significant self-shadowing, inter-reflections, and partial occlusions. Even worse, in many facial expressions a significant part of the eyelid is folded in and not visible at all. We desire an accurate performance capture that delivers consistent geometry in correspondence over time whenever visible, and deforms non-visible parts in a plausible way. Unfortunately, existing dense performance capture approaches cannot handle these extreme deformations and occlusions.

We address this problem by introducing a novel reconstruction and tracking scheme that combines a geometric deformation model with

ACM Reference Format

Bermanno, A., Beeler, T., Kozlov, Y., Bradley, D., Bickel, B., Gross, M. 2015. Detailed Spatio-Temporal Reconstruction of Eyelids. *ACM Trans. Graph.* 34, 4, Article 44 (August 2015), 11 pages. DOI = 10.1145/2766924 <http://doi.acm.org/10.1145/2766924>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766924>

Dynamic 3D Avatar Creation from Hand-held Video Input

Alexandru Eugen Ichim*
EPFL

Sofien Bouaziz†
EPFL

Mark Pauly‡
EPFL



Figure 1: Our system creates a fully rigged 3D avatar of the user from uncalibrated video input acquired with a cell-phone camera. The blendshape models of the reconstructed avatars are augmented with textures and dynamic detail maps, and can be animated in realtime.

Abstract

We present a complete pipeline for creating fully rigged, personalized 3D facial avatars from hand-held video. Our system faithfully recovers facial expression dynamics of the user by adapting a blendshape template to an image sequence of recorded expressions using an optimization that integrates feature tracking, optical flow, and shape from shading. Fine-scale details such as wrinkles are captured separately in normal maps and ambient occlusion maps. From this user- and expression-specific data, we learn a regressor for on-the-fly detail synthesis during animation to enhance the perceptual realism of the avatars. Our system demonstrates that the use of appropriate reconstruction priors yields compelling face rigs even with a minimalistic acquisition system and limited user assistance. This facilitates a range of new applications in computer animation and consumer-level online communication based on personalized avatars. We present realtime application demos to validate our method.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: 3D avatar creation, face animation, blendshapes, rigging

*alexandru.ichim@epfl.ch

†sofien.bouaziz@epfl.ch

‡mark.pauly@epfl.ch

ACM Reference Format

Ichim, A., Bouaziz, S., Pauly, M. 2015. Dynamic 3D Avatar Creation from Hand-Held Video Input. ACM Trans. Graph. 34, 4, Article 45 (August 2015), 14 pages. DOI = 10.1145/2766974
<http://doi.acm.org/10.1145/2766974>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766974>

1 Introduction

Recent advances in realtime face tracking enable fascinating new applications in performance-based facial animation for entertainment and human communication. Current realtime systems typically use the extracted tracking parameters to animate a set of pre-defined characters [Weise et al. 2009; Weise et al. 2011; Li et al. 2013; Cao et al. 2013; Bouaziz et al. 2013; Cao et al. 2014a]. While this allows the user to enact virtual avatars in realtime, personalized interaction requires a custom rig that matches the facial geometry, texture, and expression dynamics of the user. With accurate tracking solutions in place, creating compelling user-specific face rigs is currently a major challenge for new interactive applications in online communication. In this paper we propose a software pipeline for building fully rigged 3D avatars from hand-held video recordings of the user.

Avatar-based interactions offer a number of distinct advantages for online communication compared to video streaming. An important benefit for mobile applications is the significantly lower demand on bandwidth. Once the avatar has been transferred to the target device, only animation parameters need to be transmitted during live interaction. Bandwidth can thus be reduced by several orders of magnitude compared to video streaming, which is particularly relevant for multi-person interactions such as conference calls.

A second main advantage is the increased content flexibility. A 3D avatar can be more easily integrated into different scenes, such as games or virtual meeting rooms, with changing geometry, illumination, or viewpoint. This facilitates a range of new applications, in particular on mobile platforms and for VR devices such as the Oculus Rift.

Our goal is to enable users to create fully rigged and textured 3D avatars of themselves at home. These avatars should be as realistic as possible, yet lightweight, so that they can be readily integrated into realtime applications for online communication. Achieving this goal implies meeting a number of constraints: the acquisition hardware and process need to be *simple and robust*, precluding any custom-build setups that are not easily deployable. Manual assistance needs to be *minimal* and restricted to operations that can be easily performed by untrained users. The created rigs need to be *efficient* to support realtime animation, yet accurate and detailed to enable engaging virtual interactions.

Real-Time High-Fidelity Facial Performance Capture

Chen Cao^{1,2} Derek Bradley² Kun Zhou¹ Thabo Beeler²
1) State Key Lab of CAD&CG, Zhejiang University 2) Disney Research Zurich




Figure 1: We propose the first real-time facial tracking system that captures performances in online at high fidelity, including medium scale details such as wrinkles. From a monocular input (left) our system captures both the global shape as well as local details (center). The method is generic and requires no offline training or manual preprocessing steps for novel users (right).

Abstract

We present the first real-time high-fidelity facial capture method. The core idea is to enhance a global real-time face tracker, which provides a low-resolution face mesh, with local regressors that add in medium-scale details, such as expression wrinkles. Our main observation is that although wrinkles appear in different scales and at different locations on the face, they are locally very self-similar and their visual appearance is a direct consequence of their local shape. We therefore train local regressors from high-resolution capture data in order to predict the local geometry from local appearance at runtime. We propose an automatic way to detect and align the local patches required to train the regressors and run them efficiently in real-time. Our formulation is particularly designed to enhance the low-resolution global tracker with exactly the missing expression frequencies, avoiding superimposing spatial frequencies in the result. Our system is generic and can be applied to any real-time tracker that uses a global prior, e.g. blend-shapes. Once trained, our online capture approach can be applied to any new user without additional training, resulting in high-fidelity facial performance reconstruction with person-specific wrinkle details from a monocular video camera in real-time.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning;

Keywords: Real-time face reconstruction, local wrinkle model, high-fidelity performance capture.

1 Introduction

Mastering facial animation is a long-standing challenge in computer graphics. The face can describe the emotions of a character, convey their state of mind, and hint at their future actions. Audiences are particularly trained to look at faces and identify these subtle characteristics. Through decades of research, we have learned that capturing the shape and motion of real human faces can lead to realistic virtual characters.

Facial motion capture has come a long way from the original marker-based tracking approaches. Modern performance capture techniques can deliver extremely high-resolution facial geometry with very high fidelity motion information. In recent years the growing trend has been to capture faces in real-time, opening up new applications in immersive computer games, social media and real-time preview for visual effects. These methods approximate the 3D shape and motion of a face during the performance using either depth or web cameras. To make this tractable, real-time approaches use generic, low-resolution face models as a basis for the reconstruction. While these models simplify the capture problem and facilitate performance retargeting to other characters, they fail to capture the unique medium and fine scale facial details of the individual, such as wrinkles on the forehead or the so-called crows feet around the eyes. As a result, real-time performance comes at the cost of facial fidelity, and there exists a large gap in reconstruction quality between current offline and online capture methods.

In this work we aim to reduce the disparity between offline and real-time facial performance capture by presenting the first high-fidelity real-time facial capture method. Our core idea is to enhance an existing low-resolution real-time tracker by adding a local regression method that targets medium scale expression wrinkles. Our main observation is that although wrinkles appear at different locations on the face and with different orientations and scales, they exhibit very similar local appearance because of the way they are formed. The local appearance of wrinkles is a direct consequence of the shading caused by the local shape. By learning the relationship between local image appearance and local wrinkle formation, we can reconstruct plausible face wrinkles in real-time from a single RGB camera. Specifically, we train a regressor with data acquired by a high-resolution performance capture system, which can predict the shape of these features given the captured appearance in an image. The local details are stored in a displacement map that is applied to

ACM Reference Format

Cao, C., Bradley, D., Zhou, K., Beeler, T. 2015. Real-Time High-Fidelity Facial Performance Capture. ACM Trans. Graph. 34, 4, Article 46 (August 2015), 9 pages. DOI = 10.1145/2766943
<http://doi.acm.org/10.1145/2766943>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766943>

Facial Performance Sensing Head-Mounted Display

Hao Li^{†*}
Tristan Trutna[‡]

Laura Trutoiu^{†*}
Pei-Lun Hsieh[†]

Kyle Olszewski^{†*}
Aaron Nicholls[‡]

Lingyu Wei^{†*}
Chongyang Ma[‡]

[†] University of Southern California

[‡] Oculus & Facebook




Figure 1: To enable immersive face-to-face communication in virtual worlds, the facial expressions of a user have to be captured while wearing a virtual reality head-mounted display. Because the face is largely occluded by typical wearable displays, we have designed an HMD that combines ultra-thin strain sensors with a head-mounted RGB-D camera for real-time facial performance capture and animation.

Abstract

There are currently no solutions for enabling direct face-to-face interaction between virtual reality (VR) users wearing head-mounted displays (HMDs). The main challenge is that the headset obstructs a significant portion of a user's face, preventing effective facial capture with traditional techniques. To advance virtual reality as a next-generation communication platform, we develop a novel HMD that enables 3D facial performance-driven animation in real-time. Our wearable system uses ultra-thin flexible electronic materials that are mounted on the foam liner of the headset to measure surface strain signals corresponding to upper face expressions. These strain signals are combined with a head-mounted RGB-D camera to enhance the tracking in the mouth region and to account for inaccurate HMD placement. To map the input signals to a 3D face model, we perform a single-instance offline training session for each person. For reusable and accurate online operation, we propose a short calibration step to readjust the Gaussian mixture distribution of the mapping before each use. The resulting animations are visually on par with cutting-edge depth sensor-driven facial performance capture systems and hence, are suitable for social interactions in virtual worlds.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality;

Keywords: real-time facial performance capture, virtual reality, depth camera, strain gauge, head-mounted display, wearable sensors

* Authors on the first row have contributed equally.

ACM Reference Format

Li, H., Trutoiu, L., Olszewski, K., Wei, L., Trutna, T., Hsieh, P., Nicholls, A., Ma, C. 2015. Facial Performance Sensing Head-Mounted Display. *ACM Trans. Graph.* 34, 4, Article 47 (August 2015), 9 pages. DOI = 10.1145/2766939 <http://doi.acm.org/10.1145/2766939>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766939>

1 Introduction

Recent progress towards mass-market head-mounted displays (HMDs) by Oculus [Oculus VR 2014] and others, has led to a revival in virtual reality (VR). VR is drawing wide interest from consumers for gaming and online virtual worlds applications. With the help of existing motion capture and hand tracking technologies, users can navigate and perform actions in fully immersive virtual environments. However, users lack a technological solution for face-to-face communication that conveys compelling facial expressions and emotions in virtual environments. Because a user's face is significantly occluded by the HMD, established methods for facial performance tracking, such as optical sensing technologies, will fail to capture nearly the entire upper face.

To address this need, we develop a prototype HMD around an existing device. We augment the system with eight ultra-thin strain gauges (flexible metal foil sensors) placed on the foam liner for surface strain measurements and an RGB-D camera mounted on the HMD cover to capture the geometry of the visible face region. Aside from a slight increase in weight, our design unobtrusively integrates the sensors without further constraining user performance as compared to any standard virtual reality HMD.

Complex anatomical characteristics, such as individual facial tissue and muscle articulations, challenge the low dimensionality of our surface measurements across subjects. To map the input signals to a tracked 3D model in real-time, we first train a regression model by detaching the cover from the HMD to maximize visibility while the strain gauges are recording. This procedure is only performed once for each individual, and each subsequent use does not require unmounting the cover. Because of slight misplacements as well as the additional weight of the cover and the RGB-D camera, the sensitivity and measured surface locations can differ greatly between the training session and online operation (when the display is attached). For subsequent wearings by the same person, we propose a short calibration step that readjusts the Gaussian mixture distributions of the mapping [Gales 1998].

Like many real-time facial animation systems, our method uses linear blendshape models to produce output animations based on FACS expressions [Ekman and Friesen 1978]. The semantics of each blendshape mesh can be conveniently used for facial performance

The SGGX Microflake Distribution

Eric Heitz^{1,2}

Jonathan Dupuy³

Cyril Crassin²

Carsten Dachsbacher¹

¹Karlsruhe Institute of Technology

²NVIDIA

³Univ. Montréal; LIRIS, Univ. Lyon 1




Figure 1: Top-left: rendering a voxelized forest at decreasing levels of detail (left to right). Bottom-right: visualization of the voxel structure at the matching resolutions. We use the SGGX microflake distribution to represent volumetric anisotropic materials. Our representation supports downscaling and interpolation, resulting in smooth and antialiased renderings at multiple scales.

Abstract

We introduce the *Symmetric GGX* (SGGX) distribution to represent spatially-varying properties of anisotropic microflake participating media. Our key theoretical insight is to represent a microflake distribution by the projected area of the microflakes. We use the projected area to parameterize the shape of an ellipsoid, from which we recover a distribution of normals. The representation based on the projected area allows for robust linear interpolation and prefiltering, and thanks to its geometric interpretation, we derive closed form expressions for all operations used in the microflake framework.

We also incorporate microflakes with diffuse reflectance in our theoretical framework. This allows us to model the appearance of rough diffuse materials in addition to rough specular materials. Finally, we use the idea of sampling the distribution of visible normals to design a perfect importance sampling technique for our SGGX microflake phase functions. It is analytic, deterministic, simple to implement, and one order of magnitude faster than previous work.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Ray tracing;

Keywords: microflake theory, global illumination, light transport.

1 Introduction

The importance of rendering volumetric models has dramatically increased in recent years. While geometry has been used as the standard rendering primitive for decades, it is hard to represent materials like hair, fur, or fabric as surfaces. Instead, using volumetric data to describe spatially varying microstructures has proven to be a promising alternative [Zhao et al. 2011]. Furthermore, volumetric representations are also well-suited for level-of-detail (LOD) techniques that replace complex visual data (e.g. distant fine geometry) with a more compact and efficient representation of its appearance. This is an important and pressing challenge as the complexity of scenes in production rendering has reached unprecedented heights.

A physically-based definition of volumetric scattering with arbitrary microstructures, the “microflake theory”, has recently been introduced by Jakob et al. [2010]. Analogous to the microfacet theory which has been formulated for surface materials (BRDFs), this theoretical framework is used to describe volumetric scattering. It provides a phase function deduced from a given set of statistically distributed microflakes, and ensures physical correctness by constraints for energy conservation and reciprocity. While this new framework opened up new possibilities for rendering, it is currently limited by the existing microflake representations, which lack closed form operators, e.g. for evaluating the phase function and importance sampling.

In this paper we introduce a novel microflake distribution which resolves these shortcomings and thus increases the practical benefits of the microflake framework in general.

Our SGGX distribution is the first that provides lightweight storage of microstructure representation, linear parameter interpolation, and analytical evaluation and importance sampling. SGGX parameters can be generated by converting from existing data in other microflake models [Zhao et al. 2011] or created from polygonal geometry, and can be filtered to obtain volumetric level-of-detail representations.

ACM Reference Format

Heitz, E., Dupuy, J., Crassin, C., Dachsbacher, C. 2015. The SGGX Microflake Distribution. ACM Trans. Graph. 34, 4, Article 48 (August 2015), 11 pages. DOI = 10.1145/2766988 <http://doi.acm.org/10.1145/2766988>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766988>

Multi-Scale Modeling and Rendering of Granular Materials

Johannes Meng^{2,1} Marios Papas^{1,3} Ralf Habel¹
 Carsten Dachsbacher² Steve Marschner⁴ Markus Gross^{1,3} Wojciech Jarosz^{1,5*}

¹Disney Research Zürich ²Karlsruhe Institute of Technology ³ETH Zürich ⁴Cornell University ⁵Dartmouth College




Figure 1: We propose a multi-scale procedural approach for modeling granular materials. The user specifies the bounding shape for the aggregate material (top left), selects a pre-packed tile of grain bounding spheres (top middle), within which we instantiate randomly rotated copies of the selected exemplar grains (bottom left) according to the specified mixing ratios. The SANDCASTLE contains about 2 billion grains, each composed of approximately 200k triangles. We report the high-order / total render times in hours and the variance in parentheses. Our approach (top half) renders the high-order scattering over $12\times$ (50 vs. 628 hrs) faster than explicitly path tracing (EPT) the individual grains (bottom half) while providing visually indistinguishable results. The insets on the right provide equal time and equal variance comparisons.

Abstract

We address the problem of modeling and rendering granular materials—such as large structures made of sand, snow, or sugar—where an aggregate object is composed of many randomly oriented, but discernible grains. These materials pose a particular challenge as the complex scattering properties of individual grains, and their packing arrangement, can have a dramatic effect on the large-scale appearance of the aggregate object. We propose a multi-scale modeling and rendering framework that adapts to the structure of scattered light at different scales. We rely on path tracing the individual grains only at the finest scale, and—by decoupling individual grains from their arrangement—we develop a modular approach for simulating longer-scale light transport. We model light interactions within and across grains as separate processes and leverage this decomposition to derive parameters for classical radiative transport, including standard volumetric path tracing and a diffusion method that can quickly summarize the large scale transport due to many grain interactions. We require only a one-time precomputation per exemplar grain, which we can then reuse for arbitrary aggregate shapes and a continuum of different packing rates and scales of grains. We demonstrate

our method on scenes containing mixtures of tens of millions of individual, complex, specular grains that would be otherwise infeasible to render with standard techniques.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing;

Keywords: physically based rendering, granular media

1 Introduction

In this paper we consider rendering materials consisting of large assemblies of macroscopic granules. Such granular materials are common in our everyday environment: sand, gravel, and snow; sugar, salt, ground spices, laundry detergent; ocean spray or bubbles in a carbonated beverage—any large pile or aggregate object consisting of randomly oriented grains in which the individual scatterers are discernible (see Figure 2). Common among all these examples is the potential for detailed appearance at the scale of grains, but smooth large-scale appearance of the aggregate due to multiple scattering between grains. Rendering granular materials accurately and efficiently at arbitrary scales remains an open problem. Individual grains can have complex shapes and complex scattering behavior, while at the same time they can have high albedo, so that long paths with many scattering events can remain important (see Figure 3).

Treating each individual grain as explicit geometry and simulating global light transport using path tracing [Kajiya 1986] and its variants is a general solution, but is only practical for small collections of grains. At the other extreme, the aggregate object could be interpreted as a continuous medium, the smooth, large-scale appearance of which may be well expressed by participating media rendering techniques [Cerezo et al. 2005] derived from the radiative transfer equation (RTE) [Chandrasekar 1960]. Methods based on the diffu-

*The work was done while the author was employed at Disney Research.

ACM Reference Format

Meng, J., Papas, M., Habel, R., Dachsbacher, C., Marschner, S., Gross, M., Jarosz, W. 2015. Multi-Scale Modeling and Rendering of Granular Materials. ACM Trans. Graph. 34, 4, Article 49 (August 2015), 13 pages. DOI = 10.1145/2766949 <http://doi.acm.org/10.1145/2766949>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
 SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
 Copyright is held by the owner/author(s). Publication rights licensed to ACM.
 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
 DOI: <http://dx.doi.org/10.1145/2766949>

Power Particles: An incompressible fluid solver based on power diagrams

Fernando de Goes
Caltech

Corentin Wallez
École Polytechnique

Jin Huang
Zhejiang U.

Dmitry Pavlov
Imperial College London

Mathieu Desbrun
Caltech




Figure 1: *Splash de Triomphe*. A dam break near the Arc de Triomphe is simulated using 600k particles for the fluid and 80k for the solid obstacle. Bottom two rows show every sixth frame of the whole sequence, with surfacing achieved via OpenVDB [Museth 2013].

Abstract

This paper introduces a new particle-based approach to incompressible fluid simulation. We depart from previous Lagrangian methods by considering fluid particles no longer purely as material points, but also as volumetric parcels that partition the fluid domain. The fluid motion is described as a time series of well-shaped power diagrams (hence the name *power particles*), offering evenly spaced particles and accurate pressure computations. As a result, we circumvent the typical excess damping arising from kernel-based evaluations of internal forces or density without having recourse to auxiliary Eulerian grids. The versatility of our solver is demonstrated by the simulation of multiphase flows and free surfaces.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling; I.3.7 [Computer Graphics]: 3D Graphics and Realism—Animation.

Keywords: Lagrangian fluid simulation, power diagrams, incompressibility, multiphase flows, free surface.

1 Introduction

Fluid motion is a visually rich and complex phenomenon that remains, to this day, a challenge to reproduce numerically. Intricate patterns such as vortices in liquids and volutes in smoke are typ-

ical visual cues of incompressible flows that a numerical simulation needs to capture accurately. While Lagrangian methods have been shown simple and practical to efficiently generate fluid motion, they often suffer from numerical artifacts that severely impact liveliness of the flow. In particular, the manner in which (and the degree to which) incompressibility is enforced has strong implications for its dynamics. For instance, the Monte Carlo nature of kernel evaluations in Smoothed Particle Hydrodynamics (SPH) [Lucy 1977] can only control local particle density at the cost of significant motion damping [Ihmsen et al. 2014]. The addition of an auxiliary grid to provide more accurate and efficient pressure projections also causes kinetic energy dissipation and particle drifting as velocities are transferred from/to the grid [Harlow 1963; Brackbill et al. 1988]. Even moving mesh approaches struggle to enforce incompressibility due to mesh distortion and tangling [Oñate et al. 2004; Tang 2004; Clausen et al. 2013]. Consequently, none of the current Lagrangian approaches offers both strong incompressibility and low numerical dissipation.

We propose a radically different Lagrangian method to simulate incompressible fluid. Particles are considered as non-overlapping fluid parcels that partition the space occupied by the fluid through a moving power diagram. By leveraging computational tools on power diagrams, we formulate a time integrator for these “power” particles that precisely controls particle density and pressure forces, without kernel estimates or significant artificial viscosity.

1.1 Related Work

A broad variety of methods for fluid simulation have been devised in computational physics and computer graphics (see surveys in [Bridson 2008; Ihmsen et al. 2014]). Here we restrict our discussion to Lagrangian methods based on particles and moving meshes.

Early SPH work. Smoothed Particle Hydrodynamics (SPH) discretizes fluids as material points with interaction forces derived from smooth kernel functions [Monaghan 2005]. In computer graphics, this methodology was introduced for smoke and fire de-

ACM Reference Format

de Goes, F., Wallez, C., Huang, J., Pavlov, D., Desbrun, M. 2015. Power Particles: An Incompressible Fluid Solver Based on Power Diagrams. *ACM Trans. Graph.* 34, 4, Article 50 (August 2015), 11 pages.
DOI = 10.1145/2766901 <http://doi.acm.org/10.1145/2766901>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766901>

The Affine Particle-In-Cell Method

Chenfanfu Jiang[†] Craig Schroeder[†] Andrew Selle^{*} Joseph Teran^{*†} Alexey Stomakhin^{*}

[†]University of California Los Angeles ^{*}Walt Disney Animation Studios




Figure 1: APIC/PIC blends yield more energetic and more stable behavior than FLIP/PIC blends in a wine pour example. APIC/PIC blends are achieved analogously to FLIP/PIC in that it is a scaling of the particle affine matrices. ©Disney.

Abstract

Hybrid Lagrangian/Eulerian simulation is commonplace in computer graphics for fluids and other materials undergoing large deformation. In these methods, particles are used to resolve transport and topological change, while a background Eulerian grid is used for computing mechanical forces and collision responses. Particle-in-Cell (PIC) techniques, particularly the Fluid Implicit Particle (FLIP) variants have become the norm in computer graphics calculations. While these approaches have proven very powerful, they do suffer from some well known limitations. The original PIC is stable, but highly dissipative, while FLIP, designed to remove this dissipation, is more noisy and at times, unstable. We present a novel technique designed to retain the stability of the original PIC, without suffering from the noise and instability of FLIP. Our primary observation is that the dissipation in the original PIC results from a loss of information when transferring between grid and particle representations. We prevent this loss of information by augmenting each particle with a locally affine, rather than locally constant, description of the velocity. We show that this not only stably removes the dissipation of PIC, but that it also allows for exact conservation of angular momentum across the transfers between particles and grid.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation I.6.8 [Simulation and Modeling]: Types of Simulation—Animation;

Keywords: PIC, FLIP, MPM, fluids, physically-based modeling, coupling

1 Introduction

Simulating natural phenomena for virtual worlds and characters is an important application that remains extremely challenging. An artist’s need to manipulate and comprehend physical simulations imposes a significant constraint, all but requiring simulation methods to involve Lagrangian particles. In addition, the need for computational efficiency, topology change and numerical stability has led engineers toward hybrid Lagrangian/Eulerian methods. This is the cause of the ubiquity of incompressible FLIP for simulating liquids in visual effects [Zhu and Bridson 2005]. While such hybridizations solve some problems they also create numerous difficulties. Specifically, while the hybridization allows numerical algorithms to be done in the most appropriate representation, transferring between representations creates error. In this work we show how that error can be minimized with minimal effort.

While our approach will apply to a wide range of continuum phenomena, for simplicity, first consider fluid simulation. Here, pressure and viscosity updates are best done on an Eulerian grid while advection is best done with Lagrangian particles. The first and simplest method of this type is Particle-In-Cell (PIC) [Harlow 1964; Harlow and Welch 1965]. While this method is remarkably effective and simple to implement, it suffers from significant dis-

ACM Reference Format

Jiang, C., Schroeder, C., Selle, A., Teran, J., Stomakhin, A. 2015. The Affine Particle-In-Cell Method. ACM Trans. Graph. 34, 4, Article 51 (August 2015), 10 pages. DOI = 10.1145/2766996 <http://doi.acm.org/10.1145/2766996>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766996>

Restoring the Missing Vorticity in Advection-Projection Fluid Solvers

Xinxin Zhang*
UBC Computer Science

Robert Bridson†
UBC Computer Science, Autodesk Canada

Chen Greif‡
UBC Computer Science



Figure 1: Rising smoke simulations with and without IVOCK (Integrated Vorticity of Convective Kinematics). From left to right: Stable Fluids, Stable Fluids with IVOCK; BFECC, BFECC with IVOCK; MacCormack, MacCormack with IVOCK; FLIP, FLIP with IVOCK.

Abstract

Most visual effects fluid solvers use a time-splitting approach where velocity is first advected in the flow, then projected to be incompressible with pressure. Even if a highly accurate advection scheme is used, the self-advection step typically transfers some kinetic energy from divergence-free modes into divergent modes, which are then projected out by pressure, losing energy noticeably for large time steps. Instead of taking smaller time steps or using significantly more complex time integration, we propose a new scheme called IVOCK (Integrated Vorticity of Convective Kinematics) which cheaply captures much of what is lost in self-advection by identifying it as a violation of the vorticity equation. We measure vorticity on the grid before and after advection, taking into account vortex stretching, and use a cheap multigrid V-cycle approximation to a vector potential whose curl will correct the vorticity error. IVOCK works independently of the advection scheme (we present examples with various semi-Lagrangian methods and FLIP), works independently of how boundary conditions are applied (it just corrects error in advection, leaving pressure etc. to take care of boundaries and other forces), and other solver parameters (we provide smoke, fire, and water examples). For 10 ~ 25% extra computation time per step much larger steps can be used, while producing detailed vortical structures and convincing turbulence that are lost without correction.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: fluid simulation, vorticity, advection

*e-mail: zhxx@cs.ubc.ca

†e-mail: rbridson@cs.ubc.ca

‡e-mail: greif@cs.ubc.ca

ACM Reference Format

Zhang, X., Bridson, R., Greif, C. 2015. Restoring the Missing Vorticity in Advection-Projection Fluid Solvers. ACM Trans. Graph. 34, 4, Article 52 (August 2015), 8 pages. DOI = 10.1145/2766982
<http://doi.acm.org/10.1145/2766982>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766982>

Table 1: Algorithm abbreviations used through out this paper.

IVOCK	The computational routine (Alg.1) correcting vorticity for advection
SF	Classic Stable Fluids advection [Stam 1999]
SF-IVOCK	IVOCK with SF advection
SL3	Semi-Lagrangian with RK3 path tracing and clamped cubic interpolation
BFECC	Kim et al.’s scheme [2005], with extrema clamping([Selle et al. 2008])
BFECC-IVOCK	IVOCK with BFECC advection
MC	Selle et al.’s MacCormack method [2008]
MC-IVOCK	IVOCK with MacCormack
FLIP	Zhu and Bridson’s incompressible variant of FLIP [2005]
FLIP-IVOCK	FLIP advection of velocity and density, SL3 for vorticity in IVOCK.

1 Introduction

In computer graphics, incompressible fluid dynamics are often solved with fluid variables stored on a fixed Cartesian grid, known as an Eulerian approach [Stam 1999]. The advantages of pressure projection on a regular grid and the ease of treating smooth quantities undergoing large deformations help explain its success in a wide range of phenomena, such as liquids [Foster and Fedkiw 2001], smoke [Fedkiw et al. 2001] and fire [Nguyen et al. 2002]. Bridson’s text provides a good background [2008].

In Eulerian simulations, the fluid state is often solved with a time splitting method: given the incompressible Euler equation,

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{1}{\rho} \nabla p + \mathbf{f} \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0$$

fluid states are advanced by self-advection and incompressible projection. First one solves the advection equation

$$\frac{D\mathbf{u}}{Dt} = 0 \quad (2)$$

to obtain an intermediate velocity field $\tilde{\mathbf{u}}$, which is then projected to be divergence-free by subtracting ∇p , derived by a Poisson solve from $\tilde{\mathbf{u}}$ itself. We formally denote this as $\mathbf{u}^{n+1} = \text{Proj}(\tilde{\mathbf{u}})$.

Self-advection disregards the divergence-free constraint, allowing some of the kinetic energy of the flow to be transferred into diver-

A Stream Function Solver for Liquid Simulations

Ryoichi Ando*
IST Austria

Nils Thuerey†
Technische Universität München

Chris Wojtan‡
IST Austria




Figure 1: Liquid streaming out of a closed container with holes. Our solver calculates the complex formation of bubbles without explicitly solving for the gas phase. Each of the images shows FLIP particles colored by the stream function on the left, and the surface on the right. The simulation ran with a $256 \times 256 \times 128$ resolution, 54 seconds per time step, and 3.9 minutes per video frame.

Abstract

This paper presents a liquid simulation technique that enforces the incompressibility condition using a stream function solve instead of a pressure projection. Previous methods have used stream function techniques for the simulation of detailed single-phase flows, but a formulation for liquid simulation has proved elusive in part due to the free surface boundary conditions. In this paper, we introduce a stream function approach to liquid simulations with novel boundary conditions for free surfaces, solid obstacles, and solid-fluid coupling.

Although our approach increases the dimension of the linear system necessary to enforce incompressibility, it provides interesting and surprising benefits. First, the resulting flow is guaranteed to be divergence-free regardless of the accuracy of the solve. Second, our free-surface boundary conditions guarantee divergence-free motion even in the un-simulated air phase, which enables two-phase flow simulation by only computing a single phase. We implemented this method using a variant of FLIP simulation which only samples particles within a narrow band of the liquid surface, and we illustrate the effectiveness of our method for detailed two-phase flow simulations with complex boundaries, detailed bubble interactions, and two-way solid-fluid coupling.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: fluid, two-phase flow, stream function, vector potential

1 Introduction

We wish to simulate the motion of a liquid surrounded by air. The inviscid, incompressible Navier-Stokes equations describe this motion:

$$\frac{D\mathbf{u}}{Dt} = -\frac{1}{\rho}\nabla p + \mathbf{g} \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (2)$$

where \mathbf{u} is the fluid velocity field, t is time, ρ is fluid density, p is fluid pressure, and \mathbf{g} is a body acceleration. The first equation describes the evolution of fluid momentum over time, and the second equation constrains the fluid motion such that it conserves mass. The typical approach to enforcing Eq. (2) is to find a pressure p whose gradient projects the velocity field into a divergence-free state. This projection can also be interpreted as a kinetic energy minimization problem [Batty et al. 2007]:

$$\underset{p}{\text{minimize}} \quad \int_{\Omega} \frac{1}{2} \rho \|\mathbf{u}^* - \frac{\Delta t}{\rho} \nabla p\|^2 dV \quad (3)$$

ACM Reference Format

Ando, R., Thuerey, N., Wojtan, C. 2015. A Stream Function Solver for Liquid Simulations. *ACM Trans. Graph.* 34, 4, Article 53 (August 2015), 9 pages. DOI = 10.1145/2766935 <http://doi.acm.org/10.1145/2766935>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766935>

*E-mail: and@verygood.aid.design.kyushu-u.ac.jp

†E-mail: nils.thuerey@tum.de

‡E-mail: wojtan@ist.ac.at

Dihedral Angle-based Maps of Tetrahedral Meshes

Gilles-Philippe Paillé¹ Nicolas Ray² Pierre Poulin¹ Alla Sheffer³ Bruno Lévy²
¹Université de Montréal ²INRIA Nancy - Grand Est ³University of British Columbia

Abstract

We present a geometric representation of a tetrahedral mesh that is solely based on dihedral angles. We first show that the shape of a tetrahedral mesh is completely defined by its dihedral angles. This proof leads to a set of angular constraints that must be satisfied for an immersion to exist in \mathbb{R}^3 . This formulation lets us easily specify conditions to avoid inverted tetrahedra and multiply-covered vertices, thus leading to locally injective maps. We then present a constrained optimization method that modifies input angles when they do not satisfy constraints. Additionally, we develop a fast spectral reconstruction method to robustly recover positions from dihedral angles. We demonstrate the applicability of our representation with examples of volume parameterization, shape interpolation, mesh optimization, connectivity shapes, and mesh compression.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

Keywords: Volume parameterization, interpolation, compression, optimization

1 Introduction

Geometry representation plays an important role in how we think about a problem and how we solve it. A careful choice of variables can simplify the expression of a solution and lead to a straightforward implementation. Such changes of variables have significantly improved efficiency and robustness of surface maps, such as methods based on angles [Sheffer and de Sturler 2001], on metric [Springborn et al. 2008], and on curvature [Crane et al. 2011; Crane et al. 2013]. Computer animation has also benefited from representations such as barycentric coordinates [Ju et al. 2005] and modal analysis [Pentland and Williams 1989] to reduce computational time. Based on these experiences for surfaces, it is only natural to look for similar approaches for volumes. While some deformation techniques using barycentric coordinates and modal analysis trivially generalize to volumes, the body of work in this field remains limited.

In this paper, we study *dihedral angles*, a geometric quantity well-known to the volume mesh processing community. A dihedral angle is defined as the angle between two planes; each edge of a tetrahedron has a dihedral angle determined by its two incident triangles. It is commonly used to assess the quality of tetrahedral meshes [Labelle and Shewchuk 2007], and therefore, it seems natural to express an objective function that defines the quality of a mesh directly in terms of these angles, and to find a numerical solution mechanism to

optimize them. We take a step in this direction by showing that the shape of a tetrahedral mesh is completely specified by its dihedral angles, up to a global rotation and isotropic scale. Basically, this implies that manipulating these angles indeed changes the geometry of the mesh. Our goal is to study the properties and the structure of this representation, and to gain some understanding about the intricate relations between vertex positions and dihedral angles.

Our first contribution is a proof that this change of variables is equivalent to an immersion of a mesh in \mathbb{R}^3 , where an immersion is a locally injective simplicial map [Cervone 1996], which guarantees that there are no inverted tetrahedra and no multiply-covered vertices. This proof leads to the construction of a set of variables along with a complete set of constraints required for the representation to be an immersion. Our second contribution is a numerical algorithm that optimizes arbitrary input dihedral angles in such a way that they satisfy the constraints that we have exhibited. We call this process *flattening* as a parallel to surface metric flattening techniques. Our third contribution is a robust spectral reconstruction that recovers vertex positions from dihedral angles.

Related Work Our work is best interpreted as a volume equivalent of a series of angle-based surface parameterization techniques that were popularized over the last decades. The first angle-based technique was conceived by Di Battista and Vismara [1993] to draw graphs on 2D planes. This method has then been introduced to the computer graphics community for conformal parameterization of surfaces [Sheffer and de Sturler 2001]. A considerable amount of work has been done toward improving its efficiency, including by Sheffer et al. [2005] who exploit the structure of constraints to significantly decrease solving time. Convergence rate has also been improved by reformulating nonlinear constraints [Zayer et al. 2005] and linearizing the problem [Zayer et al. 2007]. More relevant related work is given along its respective applications.

The success of these methods inspired us to search for an equivalent representation for tetrahedral meshes, which lead us to dihedral angles. It is known that the dihedral angles of a single tetrahedron determine its shape [Luo 1997]. We generalize this observation to tetrahedral meshes and to build a complete framework to work with dihedral angles. While this generalization is nontrivial, we show that the same flattening-reconstruction workflow can be applied in a simple and intuitive way.

Overview Our goal is to prove that dihedral angles alone determine the shape of a tetrahedral mesh, and to show how to recover vertex positions from the sole dihedral angles (plus a global similarity transformation). We first proceed to show that the vertex positions of a tetrahedral mesh can be uniquely reconstructed from the sole dihedral angles, up to a global similarity transformation (Section 2). To prove this, we first demonstrate equivalence for a simple elementary mesh. Then, we consider a general mesh as a sequence of elementary combinatorial operations applied to a simple one, and prove the property by structural induction. In some application contexts, it may be desired that the angles match some prescribed values that do not necessarily satisfy the constraints. For these applications, we propose a constrained optimization problem that looks for the nearest set of angles that satisfy all constraints. We solve this problem using a nonlinear least-squares method (Section 3). Finally, we develop a fast spectral reconstruction method to robustly recover

ACM Reference Format

Paillé, G., Ray, N., Poulin, P., Sheffer, A., Lévy, B. 2015. Dihedral Angle-based Maps of Tetrahedral Meshes. ACM Trans. Graph. 34, 4, Article 54 (August 2015), 10 pages. DOI = 10.1145/2766900
<http://doi.acm.org/10.1145/2766900>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766900>

Conformal Mesh Deformations with Möbius Transformations

Amir Vaxman
Vienna University of Technology, Austria

Christian Müller
Vienna University of Technology, Austria

Ofir Weber
Bar Ilan University, Israel



Figure 1: A panorama of the capabilities of our framework. Deformation of a circular mesh (left). Metric-conformal deformation and interpolation in 3D (center), and an intersection-angle preserving deformation of a planar mesh (right).

Abstract

We establish a framework to design triangular and circular polygonal meshes by using face-based compatible Möbius transformations. Embracing the viewpoint of surfaces from circles, we characterize discrete conformality for such meshes, in which the invariants are circles, cross-ratios, and mutual intersection angles. Such transformations are important in practice for editing meshes without distortions or loss of details. In addition, they are of substantial theoretical interest in discrete differential geometry. Our framework allows for handle-based deformations, and interpolation between given meshes with controlled conformal error.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: Möbius transformations, discrete conformal transformations, circular meshes, handle-based editing, interpolation

1 Introduction

Editing discrete surfaces by deformations is an important branch of geometry processing. Meshes are often edited by prescribing a displacement to a subset of vertices, or by interpolating new vertex positions from existing meshes. There are common classes of meshes available for designers, each with a distinctive set of advantages. Triangle meshes are common in computer graphics, and are easy to manipulate. Polygonal meshes are more general and efficient, with less edges and vertices on average, making them popular in industrial and architectural design, where cost effectiveness plays a part. Quadrilateral meshes are viewed as discrete parametrizations of surfaces for their grid structure, and are especially comfortable for subdivision and in computer-aided design. Another important

class is *polyhedral meshes*, where every face is planar. A notable subclass of polyhedral meshes is *circular meshes*, where every face is *conyclic*, i.e., inscribed in a single circle. Considered as discrete curvature-line networks, circular meshes are useful in the field of construction and architectural geometry for their *vertex offset* properties [Liu et al. 2006; Pottmann et al. 2007], and in the field of *discrete differential geometry* for their mathematical properties [Bobenko and Suris 2008]. Circular meshes allow for proper discrete definitions of Gauss maps, shape operators, and conformal transformations. Note that triangle meshes are circular meshes by definition. Our framework is inspired by the viewpoint of *surfaces from circles* [Bobenko et al. 2006], and we show its usefulness in both design and theory. For this purpose, we employ notions from *Möbius geometry*, in which circles and spheres are fundamental objects.

A common method for mesh editing is by positional handles: the user picks and drags a small set of vertices, and the other vertices transform respecting fairness measures and constraints. Another method is the creation of meshes as averages of existing shapes by interpolation. We expect the interpolated shapes to conform to the same constraints as the originals. Moreover, if constraints are violated in the original shapes, we expect this violation to be bounded in the interpolated shapes, for a stable and intuitive result.

A common quality measure is *conformality*. Conformal transformations comprise local similarities, and consequently preserve fine details and texture (see Figure 2 for a demonstration). However, their main disadvantage is the introduction of scale variations, which might produce unintuitive results in practice. We offer a way to control and balance such variations. Conformal transformations of continuous surfaces are well defined, but there are several approaches to what discrete conformality constitutes. Moreover, some approaches are only defined in two dimensions, and cannot be generalized to 3D. We offer a unified and practical approach to several of these definitions, for both two and three dimensions.

Finally, editing circular meshes is difficult, as face concyclicity depends nonlinearly on the vertex positions. Two approaches handle this difficulty: projection of generally-deformed meshes into nearby circular meshes, and editing circular meshes in a designated subspace that preserves concyclicity. We take the latter approach.

Our contributions We offer an editing framework for triangle and circular polygonal meshes by using face-based piecewise-compatible Möbius transformations in 2D and 3D. Our approach allows for:

1. A definition and optimization of discrete conformality that unifies the following approaches:

- Preservation of circular meshes. [Bobenko et al. 2006].

ACM Reference Format

Vaxman, A., Müller, C., Weber, O. 2015. Conformal Mesh Deformations with Möbius Transformations. ACM Trans. Graph. 34, 4, Article 55 (August 2015), 11 pages. DOI = 10.1145/2766915 <http://doi.acm.org/10.1145/2766915>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766915>

Close-to-Conformal Deformations of Volumes

Albert Chern
Caltech

Ulrich Pinkall
TU Berlin

Peter Schröder
Caltech

Abstract

Conformal deformations are infinitesimal scale-rotations, which can be parameterized by quaternions. The condition that such a quaternion field gives rise to a conformal deformation is non-linear and in any case only admits Möbius transformations as solutions. We propose a particular decoupling of scaling and rotation which allows us to find near to conformal deformations as minimizers of a quadratic, convex Dirichlet energy. Applied to tetrahedral meshes we find deformations with low quasiconformal distortion as the principal eigenvector of a (quaternionic) Laplace matrix. The resulting algorithms can be implemented with highly optimized standard linear algebra libraries and yield deformations comparable in quality to far more expensive approaches.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: discrete differential geometry, digital geometry processing, conformal metric, deformations, geometric modeling

1 Introduction

Conformal mappings and deformations are in many ways the gold standard for applications ranging from parameterization [Lévy et al. 2002] and 2D morphing [Weber and Gotsman 2010] to surface deformations for geometric modeling [Crane et al. 2011], to name a few. In those settings the mappings are defined on 2D domains, where they are sufficiently flexible while maintaining mesh and texture quality throughout. The latter is due to the fact that a conformal map locally, *i.e.*, at the infinitesimal level, always looks like a scale-rotation. In particular there is no troublesome shear.

These favorable qualities have led researchers to search for ways to use such mappings for volumetric deformations as well. Alas in 3D the class of conformal deformations is not flexible at all, consisting only of Möbius transformations (rigid motions, uniform scale, and inversion in spheres) [Monge and Liouville 1850]. The challenge then becomes to pick a class of deformations which is sufficiently flexible, yet still close to conformal in a suitable sense. There have been a number of proposals along these lines in the literature (Sec. 2.1). Unfortunately they either lead to very hard non-linear optimization problems or the mappings become so flexible that they lose the properties of conformal mappings we sought to preserve in the first place.




Figure 1: While Möbius transformations are the only conformal mappings in 3D, one can get a larger variety of volumetric deformations that are close to conformal (from left to right) by blending several Möbius transformations (middle) represented as point charges (Sec. 3.4.1). (Maps courtesy FCIT)

In this paper we propose an approach which finds volumetric deformations with very low quasiconformal distortion as the minimizer of a quadratic, convex Dirichlet energy. Practically this amounts to finding the principal (belonging to the smallest eigenvalue) eigenvector of a sparse matrix. Such computational simplicity is achieved through a particular splitting of the otherwise non-linear conditions for conformality and the use of quaternions for the representation of scale-rotations. Taken together this allows us to optimize over a standard 4-dimensional vector space rather than having to deal with rotation matrices, with their implied quadratic constraints, or functions of singular values of deformation mappings. The shape of the deformations is controlled by prescribing their local scaling, *i.e.*, how much a region should grow or shrink. We will demonstrate simple ways to construct such functions based on a direct manipulation metaphor using point charges (Fig. 1), as well as a method which extends a conformal deformation of the boundary surface of the given volume into its interior (Fig. 2).

Our theory contributions include the development of the integrability condition, a proof that our energy minimization controls the weighted quasiconformal error, a justification that this weighting tends to be uniform, and an analysis of the obstruction to immersibility (torsion of the induced connection). Finally we give a novel application of the Steklov eigenproblem to extend conformal boundary deformation data to the interior. The implementation for tetrahedral meshes is described in detail. It involves, at all stages, nothing more than linear algebra for which highly optimized external libraries are readily available. Our results compare very favorably to far more expensive state of the art methods as we will demonstrate.

2 Approach

Since the quaternionic representation of scale-rotations rather than, say, scaled rotation matrices, is central to our approach we begin with a brief recall of the relevant facts.

The standard [Cayley 1845] rule for rotation of 3-vectors $v \in \mathbb{R}^3$ via a non-zero quaternion $q \in \mathbb{H}$, widely employed in computer graphics, is

$$\tilde{v} = q^{-1} v q.$$

ACM Reference Format

Chern, A., Pinkall, U., Schröder, P. 2015. Close-to-Conformal Deformation of Volumes. *ACM Trans. Graph.* 34, 4, Article 56 (August 2015), 13 pages. DOI = 10.1145/2766916 <http://doi.acm.org/10.1145/2766916>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766916>

Linear Subspace Design for Real-Time Shape Deformation

Yu Wang¹ Alec Jacobson^{2,3} Jernej Barbic⁴ Ladislav Kavan¹

¹University of Pennsylvania ²Columbia University ³ETH Zurich ⁴University of Southern California




Figure 1: Our linear subspaces are very fast to compute. This enables the users to add (or remove) control handles very quickly, allowing them to realize their creative intent in a single interactive session.

Abstract

We propose a method to design linear deformation subspaces, unifying linear blend skinning and generalized barycentric coordinates. Deformation subspaces cut down the time complexity of variational shape deformation methods and physics-based animation (reduced-order physics). Our subspaces feature many desirable properties: interpolation, smoothness, shape-awareness, locality, and both constant and linear precision. We achieve these by minimizing a quadratic deformation energy, built via a discrete Laplacian inducing linear precision on the domain boundary. Our main advantage is speed: subspace bases are solutions to a sparse linear system, computed interactively even for generously tessellated domains. Users may seamlessly switch between applying transformations at handles and *editing the subspace* by adding, removing or relocating control handles. The combination of fast computation and good properties means that designing the *right* subspace is now just as creative as manipulating handles. This paradigm shift in handle-based deformation opens new opportunities to explore the space of shape deformations.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation.

Keywords: Deformation modeling, skinning, subspace methods, reduced-order physics.

1 Introduction

Shape deformation brings life to animated characters and transforms existing geometric designs into novel variations. As shape complexities grow, so does the need for high-quality real-time shape deformation methods. Successful real-time deformation relies on building a subspace of admissible deformations, controlled by a small number of intuitive parameters. Previous attempts to design such subspaces achieve one or more desirable qualities, but no method achieves all of them.

Subspace bases derived from Euclidean-distances often boast simple, closed-form expressions, but do not respect the semantic information defined by the shape’s boundary. A basis should adapt to the given input shape: it should be *shape aware* [Joshi et al. 2007]. Recent research demonstrates the effectiveness of a variety of control structures or *handles* used to parameterized real-time deformations. Isolated point constraints, rigid bones, selected regions and exterior cages each have tasks for which they excel and those for which they fail. A good subspace should unify all handle types or combination thereof [Jacobson et al. 2011]. Finally, a deformation subspace should reproduce simple *affine* transformations (i.e. translations, rotations and scales) if applied uniformly to all handles [Hildebrandt et al. 2011]. This ensures intuitive control via simpler position constraints rather than requiring explicit full linear transformations at each handle.

Superficially, our subspace design is similar to existing variational methods: we optimize a smoothness energy within the shape with appropriate boundary conditions to support a variety of handle types. However, we make an important observation that has drastic effects on the utility of our bases. This observation begins by acknowledging the property most responsible for the success of cage-based “generalized barycentric coordinates:” linear precision or, equivalently, rest-pose reproduction. Existing variational methods inherit the *lack* of linear precision of their quadratic smoothness energies, constructed with discrete Laplacians lacking linear precision along the shape boundary. Instead, our carefully constructed volumetric Laplacian retains linear precision on the boundary of the

ACM Reference Format

Wang, Y., Jacobson, A., Barbic, J., Kavan, L. 2015. Linear Subspace Design for Real-Time Shape Deformation. ACM Trans. Graph. 34, 4, Article 57 (August 2015), 11 pages. DOI = 10.1145/2766952 <http://doi.acm.org/10.1145/2766952>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766952>

eyeSelfie: Self Directed Eye Alignment using Reciprocal Eye Box Imaging

Tristan Swedish* Karin Roesch Ik-Hyun Lee Krishna Rastogi Shoshana Bernstein Ramesh Raskar†
MIT Media Lab

Abstract

Eye alignment to the optical system is very critical in many modern devices, such as for biometrics, gaze tracking, head mounted displays, and health. We show alignment in the context of the most difficult challenge: retinal imaging. Alignment in retinal imaging, even conducted by a physician, is very challenging due to precise alignment requirements and lack of direct user eye gaze control. Self-imaging of the retina is nearly impossible.

We frame this problem as a user-interface (UI) challenge. We can create a better UI by controlling the eye box of a projected cue. Our key concept is to exploit the reciprocity, “If you see me, I see you”, to develop near eye alignment displays. Two technical aspects are critical: a) tightness of the eye box and (b) the eye box discovery comfort. We demonstrate that previous pupil forming display architectures are not adequate to address alignment in depth. We then analyze two ray-based designs to determine efficacious fixation patterns. These ray based displays and a sequence of user steps allow lateral (x, y) and depth (z) wise alignment to deal with image centering and focus. We show a highly portable prototype and demonstrate the effectiveness through a user study.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Display Algorithms H.5.2 [Information Systems]: User Interfaces—Screen Design

Keywords: retina, fundus photography, head mounted display, human computer interaction

1 Introduction

The importance of precise eye alignment to optical systems has been highlighted by the rise in popularity of consumer head mounted displays (HMDs). In these systems, we define the “eye box” as the orientation in space the eye can move while perceiving a displayed image. The eye box can be thought of as the extreme orientation (rotation and alignment) where the pupil is able to sample all desired ray angles. Many systems are designed to have an eye box that balances light efficiency and ease of alignment. However, exact alignment of HMDs cannot be guaranteed with current display methodologies. Improper alignment can lead to discomfort and nausea.

One of the most challenging alignment tasks is retinal imaging. Traditionally, acquiring retinal images involves complicated, dif-




Figure 1: Self-aligned, mobile, non-mydratic Fundus Photography. The user is presented with an alignment dependent fixation cue on a ray-based display. Once correctly aligned, a self-acquired retinal image is captured. This retinal image can be used for health, security or HMD calibration.

ficult to use and expensive equipment. The devices are designed to be used by a trained operator, need securing of head position, and non-trivial mechanical controls to ensure precise alignment. Furthermore, most retinal imaging techniques require the use of dilation drops to obtain a sufficiently large field of view (FOV). It is the combination of previous factors that make self-imaging of the retina nearly impossible.

Our key concept to drastically simplify retinal imaging is to improve the user interface by developing a framework for tighter eye box control. We develop a near eye alignment display which enables self-alignment by the user, and explore interface methodologies which ensure the precise alignment necessary to obtain self-aligned retinal images. In particular, our approach is to spread different display angles over the area of the pupil so that pupil misalignment produces a different perceived image to the user.

1.1 Contributions

- Our primary contribution is the development of an interactive “ray-cone” approach to allow for self-alignment of eyes. This method enables lateral and axial alignment of the eye to an eyepiece by producing “virtual pinholes” at the user’s pupil.
- We constructed a mathematical framework and measurement method for eye box size by evaluating the set of intersecting rays near the eye’s pupil.
- We tested this concept by establishing the first interactive system for user aligned retinal photography employing an efficient and low cost hardware design. Our self-alignment strategy for fundus photography makes retinal imaging outside a clinical setting possible.

*e-mail:tswedish@media.mit.edu

†e-mail:raskar@media.mit.edu

ACM Reference Format

Swedish, T., Roesch, K., Lee, I., Rastogi, K., Bernstein, S., Raskar, R. 2015. eyeSelfie: Self Directed Eye Alignment using Reciprocal Eye Box Imaging. ACM Trans. Graph. 34, 4, Article 58 (August 2015), 10 pages. DOI = 10.1145/2766970 <http://doi.acm.org/10.1145/2766970>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766970>

Optimal Presentation of Imagery with Focus Cues on Multi-Plane Displays

Rahul Narain¹ Rachel A. Albert² Abdullah Bulbul² Gregory J. Ward³ Martin S. Banks² James F. O'Brien²

¹University of Minnesota

²University of California, Berkeley

³Dolby Laboratories

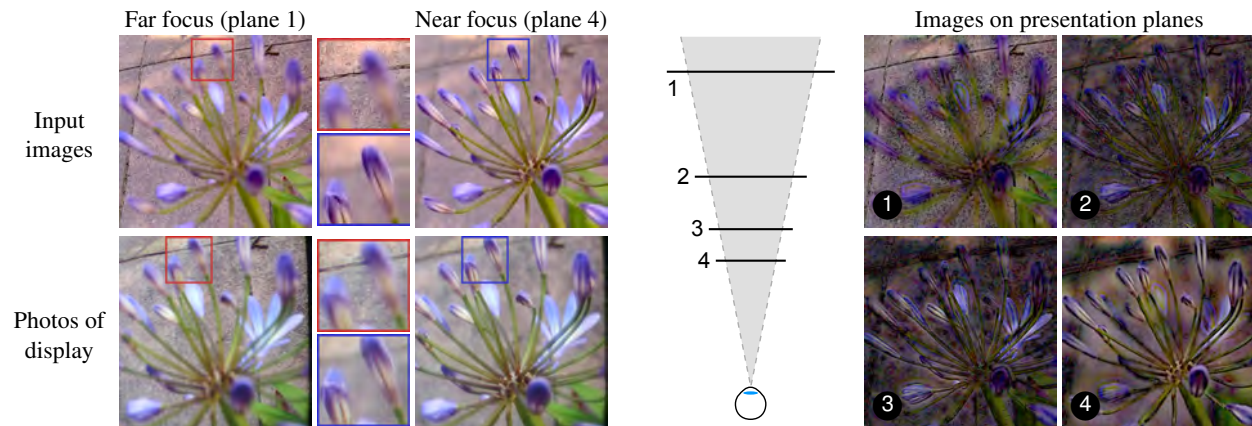


Figure 1: Reproducing a real-world scene on a multi-plane display. Given a focus stack consisting of images of a scene focused at different distances, we use optimization to determine images to show on the presentation planes of the multi-plane display so that the image seen through the display when focusing at different distances matches the corresponding image of the input scene. The presentation planes combine additively in the viewer's eye to produce an image with realistic focus cues.

Abstract

We present a technique for displaying three-dimensional imagery of general scenes with nearly correct focus cues on multi-plane displays. These displays present an additive combination of images at a discrete set of optical distances, allowing the viewer to focus at different distances in the simulated scene. Our proposed technique extends the capabilities of multi-plane displays to general scenes with occlusions and non-Lambertian effects by using a model of defocus in the eye of the viewer. Requiring no explicit knowledge of the scene geometry, our technique uses an optimization algorithm to compute the images to be displayed on the presentation planes so that the retinal images when accommodating to different distances match the corresponding retinal images of the input scene as closely as possible. We demonstrate the utility of the technique using imagery acquired from both synthetic and real-world scenes, and analyze the system's characteristics including bounds on achievable resolution.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—[Display Algorithms]

Keywords: Computational displays, multi-plane displays, eye accommodation, retinal blur, vergence-accommodation conflict

ACM Reference Format

Narain, R., Albert, R., Bulbul, M., Ward, G., Banks, M., O'Brien, J. 2015. Optimal Presentation of Imagery with Focus Cues on Multi-Plane Displays. *ACM Trans. Graph.* 34, 4, Article 59 (August 2015), 12 pages. DOI = 10.1145/2766909 <http://doi.acm.org/10.1145/2766909>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766909>

1 Introduction

The human visual system uses a number of different cues to estimate the third dimension from the 2D retinal images. Some of these—e.g. shading, perspective, and occlusion—can be reproduced in a single 2D image shown on a conventional monitor. However, other important depth cues cannot be shown on conventional displays because they arise from the geometrical relationship between scene objects and the optics of the eyes. When observing an object in a natural scene, the viewer must adjust the angle between the two eyes' lines of sight (*vergence*) to fuse the object's images in the two eyes; simultaneously, each eye adjusts the focal power of its lens (*accommodation*) to create a sharp retinal image of the object. In natural scenes, these cues are consistent, because the vergence distance, where the lines of sight intersect, and the accommodation distance, where objects are in focus, are both equal to the optical distance of the fixated object.

To display a 3D scene with vergence cues requires presenting different scenes to each eye to produce binocular disparity, while correct accommodation cues require that blur corresponds to focus changes as viewers accommodate to different distances in real time. Stereoscopic displays that provide vergence cues are now commonplace in movie theaters and are commercially available in consumer televisions. However, despite the varying depth indicated by the vergence cues in these displays, the accommodation distance to produce a sharp image remains fixed at the distance to the display surface. This *vergence-accommodation conflict* results in perceptual distortions [Watt et al. 2005], difficulty in simultaneously fusing and focusing the image [Akeley et al. 2004; Hoffman et al. 2008], and viewer discomfort and fatigue in long-term use [Emoto et al. 2005; Hoffman et al. 2008; Lambooij et al. 2009; Shibata et al. 2011]. For the displays of the future to allow effective, comfortable, and realistic viewing of stereoscopic images, they must also support correct accommodation and defocus effects.

Light field displays and volumetric displays have recently been de-

The Light Field Stereoscope

Immersive Computer Graphics via Factored Near-Eye Light Field Displays with Focus Cues

Fu-Chung Huang

Kevin Chen

Gordon Wetzstein

Stanford University




Figure 1: The light field stereoscope is a near-eye display (top left) that facilitates immersive computer graphics via stereoscopic image synthesis with correct or nearly correct focus cues. As opposed to presenting conventional 2D images, the display shows a 4D light field to each eye, allowing the observer to focus within the scene (center and right). The display comprises two stacked liquid crystal displays (LCDs) driven by nonnegative light field factorization. We implement these factorizations in real-time on the GPU; resulting patterns for front and rear LCDs, including the views for both eyes and inverse lens distortion, are shown (bottom left).

Abstract

Over the last few years, virtual reality (VR) has re-emerged as a technology that is now feasible at low cost via inexpensive cell-phone components. In particular, advances of high-resolution micro displays, low-latency orientation trackers, and modern GPUs facilitate immersive experiences at low cost. One of the remaining challenges to further improve visual comfort in VR experiences is the vergence-accommodation conflict inherent to all stereoscopic displays. Accurate reproduction of all depth cues is crucial for visual comfort. By combining well-known stereoscopic display principles with emerging factored light field technology, we present the first wearable VR display supporting high image resolution as well as focus cues. A light field is presented to each eye, which provides more natural viewing experiences than conventional near-eye displays. Since the eye box is just slightly larger than the pupil size, rank-1 light field factorizations are sufficient to produce correct or nearly-correct focus cues; *no time-multiplexed image display or gaze tracking is required*. We analyze lens distortions in 4D light field space and correct them using the afforded high-dimensional image formation. We also demonstrate significant improvements in resolution and retinal blur quality over related near-eye displays. Finally, we analyze diffraction limits of these types of displays.

ACM Reference Format

Huang, F., Chen, K., Wetzstein, G. 2015. The Light Field Stereoscope: Immersive Computer Graphics via Factored Near-Eye Light Field Display with Focus Cues. *ACM Trans. Graph.* 34, 4, Article 60 (August 2015), 12 pages. DOI = 10.1145/2766922 <http://doi.acm.org/10.1145/2766922>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766922>

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism — Virtual Reality

Keywords: computational displays, focus cues, light fields

1 Introduction

Virtual reality has gained significant traction in the last few years. Although most emerging consumer products are being advertised for gaming and entertainment applications, near-eye display technology provides benefits for society at large by providing a next-generation platform for education, collaborative work, teleconferencing, scientific visualization, remote-controlled vehicles, training and simulation, basic vision research, phobia treatment, and surgical training (e.g., [Hale and Stanney 2014]). For example, immersive VR has been demonstrated to be effective at treating post-traumatic stress disorder [Rothbaum et al. 2001] and it is an integral component of modern, minimally invasive surgery systems, such as the da Vinci surgical system¹.

To realize these applications and make VR practical for everyday and long-term use, it is crucial to create visually comfortable experiences. Current-generation VR displays support many depth cues of human vision: motion parallax, binocular disparity, binocular occlusions, and vergence. However, focus cues are usually not supported by stereoscopic displays, including head mounted displays (HMDs). Focus cues and vergence are artificially decoupled, forcing observers to maintain a fixed focal distance (on the display screen or its virtual image) while varying the vergence angle of their eyes. The resulting vergence-accommodation conflict

¹www.intuitivesurgical.com/

Decomposing Time-Lapse Paintings into Layers

Jianchao Tan*
George Mason University

Marek Dvorožňák
CTU in Prague, FEE

Daniel Sýkora
CTU in Prague, FEE

Yotam Gingold
George Mason University




Figure 1: Our approach enables rich history-based editing operations on physical paintings. From a time lapse recording of the painting process (bottom) we extract translucent paint layers into a temporal creation history. This allows artists to perform otherwise impossible spatio-temporal selections & edits which leverage temporal information to produce complex effects such as color gradients controlled by time (top left) or temporal eraser (top right). (Paint layers in this example use Porter-Duff “over” blending.) Time lapse video © Marcello Barenghi.

Abstract

The creation of a painting, in the physical world or digitally, is a process that occurs over time. Later strokes cover earlier strokes, and strokes painted at a similar time are likely to be part of the same object. In the final painting, this temporal history is lost, and a static arrangement of color is all that remains. The rich literature for interacting with image editing history cannot be used. To enable these interactions, we present a set of techniques to decompose a time lapse video of a painting (defined generally to include pencils, markers, etc.) into a sequence of translucent “stroke” images. We present transluency-maximizing solutions for recovering physical (Kubelka and Munk layering) or digital (Porter and Duff “over” blending operation) paint parameters from before/after image pairs. We also present a pipeline for processing real-world videos of paintings capable of handling long-term occlusions, such as the painter’s hand and its shadow, color shifts, and noise.

CR Categories: I.3.7 [Computer Graphics]: Picture/Image

*e-mail:tanjianchaoustc@gmail.com

ACM Reference Format
Tan, J., Dvorožňák, M., Sýkora, D., Gingold, Y. 2015. Decomposing Time-Lapse Paintings into Layers. ACM Trans. Graph. 34, 4, Article 61 (August 2015), 10 pages. DOI = 10.1145/2766960
<http://doi.acm.org/10.1145/2766960>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766960>

Generation—Bitmap and framebuffer operations I.4.6 [Image Processing and Computer Vision]; Segmentation—Pixel classification;

Keywords: images, surfaces, depth, time, video, channel, segmentation, layers, photoshop, painting

1 Introduction

A painting is an arrangement of colors on a 2D canvas. During the painting process, artists deposit color throughout the canvas via a sequence of strokes, often with a real (or simulated, in the case of digital paintings) paint brush. The final painting is, from a computational point of view, a grid of unstructured color values. Extracting structure from the final painting is extremely challenging. Yet the temporal record, which is lost in the final painting, is informative about the scene being painted. Complex drawings are drawn according to a hierarchical structure [Taylor and Tversky 1992]. Objects and parts of objects are drawn together using lower-level rules [Van Sommers 1984; Novick and Tversky 1987; Tversky 1999].

Interacting with editing history is a powerful concept in human-computer interaction. (See Nancel and Cockburn’s CAUSALITY [2014] for a recent survey and conceptual model.) This rich literature on history systems extends far beyond undo/redo. For digital image editing, this literature includes a generalization of layers for scaling, resizing, and recoloring strokes [Nancel and Cockburn 2014], revision control [Chen et al. 2011], grouping command history [Chen et al. 2012], learning from or reapplying previous commands [Grossman et al. 2010; Berthouzoz et al. 2011; Xing et al. 2014]. Wetpaint [Bonanni et al. 2009] explored a tangible “scraping” interaction to visualize layered information, such as a painting’s history. Such powerful interactions are unavailable for

Time-lapse Mining from Internet Photos

Ricardo Martin-Brualla^{1*}

David Gallup²

Steven M. Seitz^{1,2}

¹University of Washington

²Google Inc.

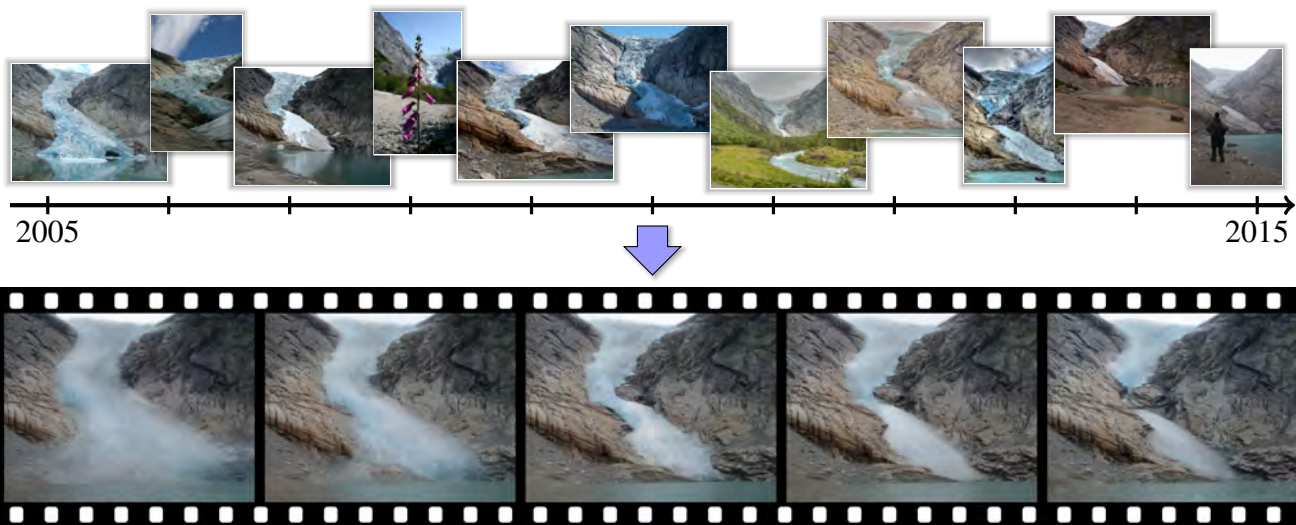


Figure 1: We mine Internet photo collections to generate time-lapse videos of locations all over the world. Our time-lapses visualize a multitude of changes, like the retreat of the Briksdalsbreen Glacier in Norway shown above. The continuous time-lapse (bottom) is computed from hundreds of Internet photos (samples on top). Photo credits: Aliento Más Allá, jirihndek, mcxurxo, elka.cz, Juan Jesús Orto, Klaus Wifßkirchen, Daikrieg, Free the image, dration and Nadav Tobias.

Abstract

We introduce an approach for synthesizing time-lapse videos of popular landmarks from large community photo collections. The approach is completely automated and leverages the vast quantity of photos available online. First, we cluster 86 million photos into landmarks and popular viewpoints. Then, we sort the photos by date and warp each photo onto a common viewpoint. Finally, we stabilize the appearance of the sequence to compensate for lighting effects and minimize flicker. Our resulting time-lapses show diverse changes in the world's most popular sites, like glaciers shrinking, skyscrapers being constructed, and waterfalls changing course.

CR Categories: I.2.10 [Artificial Intelligence]: Vision and Scene Understanding—Modeling and recovery of physical attributes; I.4.3 [Image Processing and Computer Vision]: Enhancement—Filtering and Geometric Correction; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Stereo and Time-varying imagery

Keywords: time-lapse, computational photography, image-based rendering

*This work was partially done while the first author was an intern at Google.

1 Introduction

We see the world at a fixed temporal scale, in which life advances one second at a time. Instead, suppose that you could observe an entire year in a few seconds—a 10 million times speed-up. At this scale, you could see cities expand, glaciers shrink, seasons change, and children grow continuously. Time-lapse photography provides a fascinating glimpse into these timescales. And while limited time-lapse capabilities are available on consumer cameras [Apple ; Instagram], observing these ultra-slow effects requires a camera that is locked down and focused on a single target over a period of months or years [Extreme Ice Survey].

Yet, these ultra-slow changes are documented by the billions of photos that people take over time. Indeed, an Internet image search for any popular site yields several years worth of photos. In this paper, we describe how to transform these photo collections into high quality, stabilized time-lapse videos. Figure 1 shows a few frames from one result video of a glacier receding over a decade. This capability is transformative; whereas before it took months or years to create one such time-lapse, we can now almost instantly create thousands of time-lapses covering the most popular places on earth. The challenge now is to find the interesting ones, from all of the public photos in the world. We call this problem *time-lapse mining*.

Creating high quality time-lapses from Internet photo sharing sites is challenging, due to the vast viewpoint and appearance variation in such collections. The main technical contribution of this paper is an approach for producing extremely stable videos, in which viewpoint and transient appearance changes are almost imperceptible, allowing the viewer to focus on the more salient, longer time scale scene changes. We employ structure-from-motion and stereo algorithms to compensate for viewpoint variations, and a simple but effective new temporal filtering approach to stabilize appearance. Our second significant contribution is a world-scale deployment,

ACM Reference Format

Martin-Brualla, R., Gallup, D., Seitz, S. 2015. Time-Lapse Mining from Internet Photos. ACM Trans. Graph. 34, 4, Article 62 (August 2015), 8 pages. DOI = 10.1145/2766903 <http://doi.acm.org/10.1145/2766903>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766903>

Real-Time Hyperlapse Creation via Optimal Frame Selection

Neel Joshi Wolf Kienzle Mike Toelle Matt Uyttendaele Michael F. Cohen

Microsoft Research




Figure 1: Hand-held videos often exhibit significant semi-regular high-frequency camera motion due to, for example, running (dotted blue line). This example shows how a naive 8x hyperlapse (i.e., keeping 1 out every 8 frames) results in frames with little overlap that are hard to align (black lines). By allowing small violations of the target skip rate we create hyperlapse videos that are smooth even when there is significant camera motion (pink lines). Optimizing an energy function (color-coded in Middle image) that balances matching the target rate while minimizing frame-to-frame motion results in a set of frames that are then stabilized. (Right) To illustrate the alignment we show the mean and standard deviation of three successive frames (in red box on the Left plot) after stabilization for the naive hyperlapse (Top Right) and our result (Bottom Right) – these show that our selected frames align much better than those from naive selection.

Abstract

Long videos can be played much faster than real-time by recording only one frame per second or by dropping all but one frame each second, i.e., by creating a *timelapse*. Unstable hand-held moving videos can be *stabilized* with a number of recently described methods. Unfortunately, creating a stabilized timelapse, or *hyperlapse*, cannot be achieved through a simple combination of these two methods. Two hyperlapse methods have been previously demonstrated: one with high computational complexity and one requiring special sensors. We present an algorithm for creating hyperlapse videos that can handle significant high-frequency camera motion and runs in real-time on HD video. Our approach does not require sensor data, thus can be run on videos captured on any camera. We optimally select frames from the input video that best match a desired target speed-up while also resulting in the smoothest possible camera motion. We evaluate our approach using several input videos from a range of cameras and compare these results to existing methods.

CR Categories: I.3.8 [Computer Graphics]: Applications; I.4.8 [Image Processing and Computer Vision]: Applications;

Keywords: time-lapse, hyperlapse, video stabilization

ACM Reference Format

Joshi, N., Kienzle, W., Toelle, M., Uyttendaele, M., Cohen, M. 2015. Real-Time Hyperlapse Creation via Optimal Frame Selection. ACM Trans. Graph. 34, 4, Article 63 (August 2015), 9 pages. DOI = 10.1145/2766954 <http://doi.acm.org/10.1145/2766954>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766954>

1 Introduction

The proliferation of inexpensive, high quality video cameras along with increasing support for video sharing has resulted in people taking videos more often. While increasingly plentiful storage has made it very easy to record long videos, it is still quite tedious to view and navigate such videos, as users typically do not have the time or patience to sift through minutes of unedited footage. One simple way to reduce the burden of watching long videos is to speed them up to create “timelapse” videos, where one can watch minutes of video in seconds.

When video is shot with a stationary camera, timelapse videos are quite effective; however, if there is camera motion, the speed-up process accentuates the apparent motion, resulting in a distracting and nauseating jumble. “Hyperlapse” videos are an emerging medium that addresses the difficulty of timelapse videos shot with moving cameras by performing camera motion smoothing (or “stabilization”) in addition to the speed-up process. They have a unique appealing dynamism and presence.

The two main approaches for stabilizing camera motion are hardware-based and software-based. Hardware-based methods utilizing onboard gyros can be quite successful [Karpenko 2014], but require specialized hardware at capture time, thus cannot be applied to existing videos. As they are blind to the content of the video, they also fail to stabilize large foreground objects. Software-based computer vision methods operate on the pixels themselves. They range from 2D stabilization to full 3D reconstruction and stabilization. Existing 2D approaches can work well when camera motion is slow, but breakdown when the camera has high-frequency motion. 3D approaches work well when there is sufficient camera motion and parallax in a scene [Kopf et al. 2014], but have high computational cost and are prone to tracking and reconstruction errors when there is insufficient camera translation.

In this paper, we present an algorithm for creating hyperlapse videos that runs in real-time (30 FPS on a mobile device and even faster on a desktop) and can handle significantly more camera mo-

Isotopic Approximation within a Tolerance Volume

Manish Mandad David Cohen-Steiner Pierre Alliez
INRIA Sophia Antipolis - Méditerranée

Abstract

We introduce in this paper an algorithm that generates from an input tolerance volume a surface triangle mesh guaranteed to be within the tolerance, intersection free and topologically correct. A pliant meshing algorithm is used to capture the topology and discover the anisotropy in the input tolerance volume in order to generate a concise output. We first refine a 3D Delaunay triangulation over the tolerance volume while maintaining a piecewise-linear function on this triangulation, until an isosurface of this function matches the topology sought after. We then embed the isosurface into the 3D triangulation via mutual tessellation, and simplify it while preserving the topology. Our approach extends to surfaces with boundaries and to non-manifold surfaces. We demonstrate the versatility and efficacy of our approach on a variety of data sets and tolerance volumes.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Boundary representations;

Keywords: Isotopic approximation, tolerance volume, mesh refinement, mesh simplification, intersection-free, mutual tessellation

1 Introduction

Faithful approximation of complex shapes with simplicial meshes is a multifaceted problem, involving geometry, topology and their discretization. This problem has received considerable interest due to its wide range of applications and the ever-increasing accessibility of geometric sensors. Increased availability of scanned geometric models, however, does not mean improved quality: while many practitioners have access to high-end acquisition systems, a recent trend is to replace these expensive tools with a combination of consumer-level acquisition devices. Measurement data generated by, and merged from, these heterogeneous devices are reputedly unfit for direct processing. Similarly, the growing variety of geometry processing tools often increase the net amount of defects in data: Conversion to and from various geometry representations often degrades the input, and rare are the algorithms that have stronger guarantees on their output than they have requirements on their input. As we deal with ever finer discretizations to capture intricate geometric features, this issue of offering strict geometric guarantees to be robust to the occurrence of artifacts is becoming more prevalent.

Geometric guarantees usually refer to upper bounds on the approximation error and to the absence of self-intersections. Topological guarantees refer to homotopy, homeomorphism or isotopy. In our

context *isotopy* means that there exists a smooth deformation that maps one shape to another while maintaining a homeomorphism between the two. Surface meshes with such guarantees are required for artifact-free rendering, computational engineering, reverse engineering, manufacturing and 3D printing. While *geometric* simplification can reduce the number of primitives, *topological* simplification can repair holes and degeneracies in existing discretizations. Combined, the two may also be used for reconstructing clean shapes from raw geometric data such as point sets or polygon soups.

1.1 Related Work

A vast array of methodologies has been proposed for shape approximation over the years, ranging from decimation to optimization through clustering and refinement. Fewer, however, provide error bounds. In addition, they only apply to specific types of input geometry, and often fail to satisfy geometric *and* topological guarantees as we now review.

[Agarwal and Suri 1998] proposed a polynomial-time approximation algorithm with guaranteed maximum error and minimum number of vertices, but this algorithm is too complex to be practically relevant. Approximation with bounded error has also been targeted through clustering [Kalvin and Taylor 1996], mesh decimation [Cohen et al. 1996; Klein et al. 1996; Guéziec 1996; Ciampalini et al. 1997; Cohen et al. 2003; Botsch et al. 2004; Ovreiu et al. 2012] or a combination of both [Zelinka and Garland 2002]. In general the error metric considered is the one-sided Hausdorff distance to the input mesh, but the normal deviation has also been considered [Borouchaki and Frey 2005]. In general these approaches are not generic enough to handle heterogeneous input data, and they are not designed to guarantee a valid, intersection-free output. Guarantees for intersection-free output can be obtained by preventing intersections during mesh decimation [Gumhold et al. 2003]) but this approach is not sufficient when the input itself self-intersects. Also, searching for the locus of points in space that avoids intersections when applying a decimation operator, and tunneling out of situations where every operator is forbidden is often too labor-intensive to be considered a practical solution. Another class of approaches based on Delaunay filtering and refinement, instead, provide intersection-free approximations by construction [Boissonnat and Oudot 2005]. Unfortunately, these approaches generate only isotropic meshes and thus do not target very coarse approximations and, as such, cannot be used for shape simplification.

When dealing with imperfect or heterogeneous data, methods involving repairing [Ju 2004; Bischoff et al. 2005; Attene 2010], conversion [Shen et al. 2004], or reconstruction [Hornung and Kobbelt 2006; Kazhdan et al. 2006] are designed to generate clean meshes, but do not yield low-polygon-count approximations with bounded error. Other approaches target only hole filling and do not remove the self-intersections [Dey and Goswami 2003]. A great variety of methods have been proposed for surface reconstruction. Among them, those which come with theoretical guarantees of isotopy [Amenta and Bern 1998; Amenta et al. 2000; Boissonnat and Cazals 2000; Dey 2006] assume that the sampled surface is smooth. Surface reconstruction from noisy point sets [Dey and Sun 2005; Dey 2006] has also been investigated well. A recent approach deals with boundaries but does not handle noisy data [Dey et al. 2009]. In addition, all these methods generate isotropic meshes, overly com-

ACM Reference Format

Mandad, M., Cohen-Steiner, D., Alliez, P. 2015. Isotopic Approximation within a Tolerance Volume. ACM Trans. Graph. 34, 4, Article 64 (August 2015), 12 pages. DOI = 10.1145/2766950
<http://doi.acm.org/10.1145/2766950>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766950>

Data-Driven Interactive Quadrangulation

Giorgio Marcias*
CNR of Italy, University of Pisa

Olga Sorkine-Hornung†
ETH Zurich

Kenshi Takayama‡
National Institute of Informatics

Enrico Puppo||
University of Genova

Nico Pietroni‡
CNR of Italy

Paolo Cignoni**
CNR of Italy

Daniele Panozzo§
ETH Zurich




Figure 1: Given a set of hand-made quadrangulated input models, our algorithm learns the quadrangulation patterns used to design them. This knowledge is employed in a sketch-based retopology tool to interactively quadrangulate user-sketched patches with subdivided edges. The user can sketch strokes inside a patch to suggest a specific edge-flow and the system automatically selects a quadrangulation following it.

Abstract

We propose an interactive quadrangulation method based on a large collection of patterns that are learned from models manually designed by artists. The patterns are distilled into compact quadrangulation rules and stored in a database. At run-time, the user draws strokes to define patches and desired edge flows, and the system queries the database to extract fitting patterns to tessellate the sketches' interiors. The quadrangulation patterns are general and can be applied to tessellate large regions while controlling the positions of the singularities and the edge flow. We demonstrate the effectiveness of our algorithm through a series of live retopology sessions and an informal user study with three professional artists.

CR Categories: I.3.5 [Computer Graphics]: Computational geometry and object modeling—Curve, surface, solid and object representations.

Keywords: quad meshing, polygon quadrangulation, retopology

1 Introduction

Quadrilateral meshes are ubiquitously used in the animation and CAD industry as control grids for subdivision surfaces and NURBS. Many approaches have been proposed to convert an unstructured, triangulated surface into a high-quality polygonal mesh, ranging from local decimation strategies to global field-aligned optimizations [Bommes et al. 2013b]. While dense quadrilateral meshes can be robustly created with existing methods, the creation of coarse patch layouts, or so-called *retopology*, is a challenging open problem, since the connection between the geometry of a surface and its ideal, application-dependent quad mesh is weak or nonexistent. For example, if the mesh was intended to be used for animation, its connectivity should be tailored to its articulation and optimized to reduce skinning deformation artifacts; such properties are impossible to automatically extract from static meshes.

In the industry, coarse quad layouts are manually created by professional designers, who employ their semantic knowledge and experience to adjust the layout in the context of the particular application

ACM Reference Format

Marcias, G., Takayama, K., Pietroni, N., Panozzo, D., Sorkine-Hornung, O., Puppo, E., Cignoni, P. 2015. Data-Driven Interactive Quadrangulation. *ACM Trans. Graph.* 34, 4, Article 65 (August 2015), 10 pages. DOI = 10.1145/2766964 <http://doi.acm.org/10.1145/2766964>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766964>

*e-mail:giorgio.marcias@isti.cnr.it

†e-mail:takayama@nii.ac.jp

‡e-mail:nico.pietroni@isti.cnr.it

§e-mail:panozzo@inf.ethz.ch

¶e-mail:sorkine@inf.ethz.ch

||e-mail:puppo@disi.unige.it

**e-mail:paolo.cignoni@isti.cnr.it

Convolutional Wasserstein Distances: Efficient Optimal Transportation on Geometric Domains

Justin Solomon
Stanford University

Fernando de Goes
Pixar Animation Studios

Gabriel Peyré
CNRS & Univ. Paris-Dauphine

Marco Cuturi
Kyoto University

Adrian Butscher
Autodesk, Inc.

Andy Nguyen
Stanford University

Tao Du
Stanford University

Leonidas Guibas
Stanford University



Figure 1: Shape interpolation from a cow to a duck to a torus via convolutional Wasserstein barycenters on a $100 \times 100 \times 100$ grid, using the method at the beginning of §7.

Abstract

This paper introduces a new class of algorithms for optimization problems involving optimal transportation over geometric domains. Our main contribution is to show that optimal transportation can be made tractable over large domains used in graphics, such as images and triangle meshes, improving performance by orders of magnitude compared to previous work. To this end, we approximate optimal transportation distances using entropic regularization. The resulting objective contains a geodesic distance-based kernel that can be approximated with the heat kernel. This approach leads to simple iterative numerical schemes with linear convergence, in which each iteration only requires Gaussian convolution or the solution of a sparse, pre-factored linear system. We demonstrate the versatility and efficiency of our method on tasks including reflectance interpolation, color transfer, and geometry processing.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry & Object Modeling—Geometric algorithms, languages, & systems

Keywords: Optimal transportation, Wasserstein distances, entropy, displacement interpolation.

1 Introduction

Probability distributions are ubiquitous objects in computer graphics, used to encapsulate possibly uncertain information associated with arbitrary geometric domains. Examples include image histograms, geometric features, relaxations of correspondence maps, and even physical quantities like BRDFs. To compare these objects, it is important to define an adequate notion of proximity or coverage quantifying the discrepancy or, equivalently, similarity between

distributions. These computations are commonly posed and analyzed within the theory of *optimal transportation*.

The prototypical problem in optimal transportation is the evaluation of Wasserstein (also known as Earth Mover’s) distances between distributions [Villani 2003; Rubner et al. 2000]. These distances quantify the geometric discrepancy between two distributions by measuring the minimal amount of “work” needed to move all the mass contained in one distribution onto the other. Recent developments show that incorporating these distances into optimization objectives yields powerful tools for manipulating distributions for tasks like density interpolation, barycenter computation, and correspondence estimation. As a simple example, suppose we are given two delta functions δ_x, δ_y centered at $x, y \in \mathbb{R}^2$. While the Euclidean average $(\delta_x + \delta_y)/2$ is bimodal at x and y , solving for the distribution that minimizes the sum of squared two-Wasserstein distances to δ_x and δ_y is a Dirac at the midpoint $(x+y)/2$, thus offering a geometric notion of the midpoint of two distributions.

A limiting factor in optimal transportation is the complexity of the underlying minimization problem. The usual linear program describing optimal transportation is related to minimum-cost matching, with a quadratic number of variables and time complexity scaling at least cubically in the size of the domain [Burkard and Çela 1999]. This poor complexity is largely due to the use of coupling variables representing the amount of mass transported between every pair of samples. Hence, existing large-scale methods often resort to aggressive or ad-hoc approximations that can lose connections to transportation theory or compensate with alternative formulations that apply only to restricted cases.

This paper introduces a fast, scalable numerical framework for optimal transportation over geometric domains. Our work draws insight from recent advances in machine learning approximating optimal transportation distances using entropic regularization [Cuturi 2013]. We adapt this approach to continuous domains using faithful finite elements discretizations of the corresponding optimization problems. This yields a novel approach to optimal transportation without computing or storing pairwise distances on arbitrary shapes.

After discretization, our algorithm for approximating Wasserstein distances becomes a simple iterative scheme with linear convergence, whose iterations require convolution of vectors against discrete diffusion kernels—hence the name *convolutional Wasserstein distance*. We also leverage our framework to design methods for interpolation between distributions, computation of weighted barycenters of sets of distributions, and more complex distribution-valued correspon-

ACM Reference Format

Solomon, J., de Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., Guibas, L. 2015. Convolutional Wasserstein Distances: Efficient Optimal Transportation on Geometric Domains. *ACM Trans. Graph.* 34, 4, Article 66 (August 2015), 11 pages. DOI = 10.1145/2766963 <http://doi.acm.org/10.1145/2766963>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766963>

Sampling Based Scene-Space Video Processing

Felix Klose^{12*}

Oliver Wang^{1*}

Jean-Charles Bazin¹

Marcus Magnor²

Alexander Sorkine-Hornung¹

¹ Disney Research Zurich


² TU Braunschweig



Denoising



Action shots



Virtual aperture

Figure 1: Single frames from video results created with our sampling based scene-space video processing framework. It enables fundamental video applications such as denoising (left) as well as new artistic results such as action shots (center) and virtual aperture effects (right). Our approach is robust to unavoidable inaccuracies in 3D information, and can be used on casually recorded, moving video.

Abstract

Many compelling video processing effects can be achieved if per-pixel depth information and 3D camera calibrations are known. However, the success of such methods is highly dependent on the accuracy of this “scene-space” information. We present a novel, sampling-based framework for processing video that enables high-quality scene-space video effects in the presence of inevitable errors in depth and camera pose estimation. Instead of trying to improve the explicit 3D scene representation, the key idea of our method is to exploit the high redundancy of approximate scene information that arises due to most scene points being visible multiple times across many frames of video. Based on this observation, we propose a novel pixel gathering and filtering approach. The gathering step is general and collects pixel samples in scene-space, while the filtering step is application-specific and computes a desired output video from the gathered sample sets. Our approach is easily parallelizable and has been implemented on GPU, allowing us to take full advantage of large volumes of video data and facilitating practical runtimes on HD video using a standard desktop computer. Our generic scene-space formulation is able to comprehensively describe a multitude of video processing applications such as denoising, deblurring, super resolution, object removal, computational shutter functions, and other scene-space camera effects. We present results for various casually captured, hand-held, moving, compressed, monocular videos depicting challenging scenes recorded in uncontrolled environments.

CR Categories: I.2.10 [Computer Graphics]: Picture/Image Generation—Display Algorithms I.3.7 [Artificial Intelligence]: Vision and Scene Understanding—3D/stereo scene analysis;

Keywords: Video processing, Sampling, Inpainting, Denoising, Computational Shutters

1 Introduction

Scene-space video processing, where pixels are processed according to their 3D positions, has many advantages over traditional image-space processing. For example, handling camera motion, occlusions, and temporal continuity entirely in 2D image-space can in general be very challenging, while dealing with these issues in scene-space is simple. As scene-space information, e.g., in the form of depth maps and camera pose parameters, becomes more and more widely available due to advances in tools and mass market hardware devices (such as portable light field cameras, depth-enabled smart phones [Google 2015] and RGBD cameras), techniques that leverage depth information will play an important role in future video processing approaches.

Previous work has shown that accurate scene-space information makes fundamental video processing problems simple, automatic, and robust [Bhat et al. 2007; Zhang et al. 2009a] and even enables the creation of new, compelling video effects [Kholgade et al. 2014; Kopf et al. 2014]. However, the visual output quality of such methods is highly dependent on the quality of the available scene-space information. Despite considerable advances in 3D reconstruction over the last decades, computing exact 3D information for arbitrary video recorded under uncontrolled conditions remains an elusive goal (and will likely remain so for the foreseeable future, due to inherent ambiguities in these tasks).

We propose an alternative, general-purpose framework that facilitates robust scene-space video processing in the presence of incor-

* denotes joint first authorship with equal contribution

ACM Reference Format

Klose, F., Wang, O., Bazin, J., Magnor, M., Sorkine-Hornung, A. 2015. Sampling Based Scene-space Video Processing. ACM Trans. Graph. 34, 4, Article 67 (August 2015), 11 pages. DOI = 10.1145/2766920 <http://doi.acm.org/10.1145/2766920>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766920>

audeosynth: Music-Driven Video Montage

Zicheng Liao
Zhejiang University

Yizhou Yu
The University of Hong Kong

Bingchen Gong
Zhejiang University

Lechao Cheng
Zhejiang University

Abstract

We introduce music-driven video montage, a media format that offers a pleasant way to browse or summarize video clips collected from various occasions, including gatherings and adventures. In music-driven video montage, the music drives the composition of the video content. According to musical movement and beats, video clips are organized to form a montage that visually reflects the experiential properties of the music. Nonetheless, it takes enormous manual work and artistic expertise to create it. In this paper, we develop a framework for automatically generating music-driven video montages. The input is a set of video clips and a piece of background music. By analyzing the music and video content, our system extracts carefully designed temporal features from the input, and casts the synthesis problem as an optimization and solves the parameters through Markov Chain Monte Carlo sampling. The output is a video montage whose visual activities are cut and synchronized with the rhythm of the music, rendering a symphony of audio-visual resonance.

CR Categories: I.3.0 [Computer Graphics]: General.

Keywords: audio-visual synthesis

*The word and its sound, form and its color
are vessels of a transcendental essence that we dimly surmise
As sound lends sparkling color to the spoken word
so color lends psychically resolved tone to form*

JOHANNES ITTEN, 1888 - 1967

1 Introduction

We have seen a thrilling series of work on visual media synthesis in the past decades, including Video Textures [Schödl et al. 2000], stochastic motion textures [Chuang et al. 2005], visual mosaic [Agarwala et al. 2004; Irani et al. 1996; Kopf et al. 2010], video synopsis [Pritch et al. 2008], cinemagraphs [Beck and Burg 2012] and its extensions [Bai et al. 2012; Joshi et al. 2012; Liao et al. 2013]. Perhaps at a higher level is the concept of moment images envisioned by Cohen and Szeliski [2006]. By their definition, a moment image captures a universal yet subjective state by combining visual snapshots from multiple shutter instants.

While such work successfully creates lively visual representations of the world, other dimensions of human sensation are absent, e.g.



Figure 1: Given a set of video clips and a piece of music, our system analyzes the video and music content, and produces a video montage that “dances” to the beat of the music.

hearing, touch, smell or taste. To break the status-quo, we advocate the concept of *audio-visual synthesis* by French theorist Michel Chion [1994], and introduce Music-Driven Video Montage: video montages that “dance” to a piece of background music. Music-driven video montage belongs to a type of audio-visual media called *music-driven imagery* [Goodwin 1992; Wesseling 2004], where the music drives the composition of the video content. According to musical movements and beats, video clips are reorganized to form a montage that “visually reflects the experiential properties of the music” [Wesseling 2004], e.g., body parts moving back and forth to the rhythm of music beats, footsteps on music notes, portamento when a bird glides over lake Tahoe, and on a music transition the scene cuts to the deep valley of Yosemite.

Music-driven video montage suggests a pleasant way to browse or summarize video clips collected from gatherings, adventures, sport events, or time-lapse photography. A niche of application is on the widespread mobile devices. With these devices, people nowadays shoot moments of daily life and share them online very easily. Music-driven video montage would be a nice way to organize such video clips and play them to the beat of a selected background music. The result is of higher aesthetics value, and provides a new experience for personal viewing and online sharing. However, it takes enormous manual work and artistic expertise to create such result. Our goal is to develop a computer-aided solution to creating such audio-visual composition (Figure 1).

Creating a high-quality music-driven video montage with a machine algorithm is a challenging new task. First of all, given a piece of background music and a set of video clips, it is not obvious how should the problem be formulated, because producing an audio-visual composition has a large degree of freedom. The algorithm does not have to use all the video clips or all the frames in a chosen video. It also needs to determine the order of the chosen videos on the music timeline, the cut between the videos, as well as the play speed of every chosen video. To make the problem tractable, we adopt two important thumb-of-rules when artists create music-driven imagery: *cut-to-the-beat* and *synchronization* [Wesseling 2004; Dmytryk 2010]. Cut-to-the-beat means video sequences should be cut at significant note onsets in the mu-

ACM Reference Format

Liao, Z., Yu, Y., Gong, B., Cheng, L. 2015. AudeoSynth: Music-Driven Video Montage. ACM Trans. Graph. 34, 4, Article 68 (August 2015), 10 pages. DOI = 10.1145/2766966 <http://doi.acm.org/10.1145/2766966>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766966>

High-Quality Streamable Free-Viewpoint Video

Alvaro Collet Ming Chuang Pat Sweeney Don Gillett Dennis Evseev David Calabrese
Hugues Hoppe Adam Kirk Steve Sullivan
Microsoft Corporation



Figure 1: Examples of reconstructed free-viewpoint video acquired by our system.

Abstract

We present the first end-to-end solution to create high-quality free-viewpoint video encoded as a compact data stream. Our system records performances using a dense set of RGB and IR video cameras, generates dynamic textured surfaces, and compresses these to a streamable 3D video format. Four technical advances contribute to high fidelity and robustness: multimodal multi-view stereo fusing RGB, IR, and silhouette information; adaptive meshing guided by automatic detection of perceptually salient areas; mesh tracking to create temporally coherent subsequences; and encoding of tracked textured meshes as an MPEG video stream. Quantitative experiments demonstrate geometric accuracy, texture fidelity, and encoding efficiency. We release several datasets with calibrated inputs and processed results to foster future research.

CR Categories: I.3.0 [Computer Graphics]: General

Keywords: multi-view stereo, surface reconstruction, 3D video, mesh tracking, geometry compression, MPEG

1 Introduction

Interest in free-viewpoint video (FVV) has soared with recent advances in both capture technology and consumer virtual/augmented reality hardware, e.g., Microsoft HoloLens. As real-time view tracking becomes more accurate and pervasive, a new class of immersive viewing experiences becomes possible on a broad scale, demanding similarly immersive content.

Our goal is to transform free-viewpoint video from research prototype into a rich and accessible form of media. Several system components must work together to achieve this goal: capture rigs must be easy to reconfigure and support professional production workflows; reconstruction must be automatic and scalable to high processing throughput; and results must be compressible to a data rate close to common media formats. Visual quality from any angle must be on par with traditional video, and the format must be renderable in real-time on a wide range of consumer devices.

In this paper, we discuss how we address these challenges to create an end-to-end system for realistic, streamable free-viewpoint video at significantly higher quality than the state of the art. Our approach does not require prior knowledge of the scene content. It handles challenging scenarios such as multi-figure captures, detailed hand and facial expressions, free-flowing clothing, animals, and a wide variety of props from golf clubs to ropes to furniture. The example captures in Fig. 1 and 17 illustrate a range of scenarios.

Our main contribution is an integrated capture and processing pipeline that attains fidelity, robustness, and scale suitable for professional content production. Within this pipeline, several key advancements enable the high-quality results:

- A novel multimodal reconstruction that leverages RGB, infrared (IR), and silhouette information to recover dense shape detail.
- An extension to Poisson surface reconstruction (PSR) [Kazhdan et al. 2006] that performs implicit function clipping to improve outlier removal and maintain silhouette accuracy.
- A method to automatically identify and preserve perceptually important areas (e.g., faces, hands) in performances.
- A parallelizable, non-rigid mesh tracking framework to automatically handle topology changes by segmenting the performance into a sequence of deforming keyframe meshes.
- An encoding method to encapsulate tracked textured meshes and audio information into a single low-bandwidth MPEG stream.

To support further research, we publish a dataset with thousands of FVV frames, each comprising 106 calibrated and synchronized camera images together with intermediate and final processing results.

ACM Reference Format

Collet, A., Chuang, M., Sweeney, P., Gillett, D., Evseev, D., Calabrese, D., Hoppe, H., Kirk, A., Sullivan, S. 2015. High-Quality Streamable Free-Viewpoint Video. ACM Trans. Graph. 34, 4, Article 69 (August 2015), 13 pages. DOI = 10.1145/2766945 <http://doi.acm.org/10.1145/2766945>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766945>

Bijjective Parameterization with Free Boundaries

Jason Smith*
Texas A&M University

Scott Schaefer†
Texas A&M University




Figure 1: Starting from Tutte’s parameterization (left), our optimization generates a parameterization that minimizes distortion and guarantees a bijective map (right). We show intermediate stages of the optimization where, at every step, the parameterization is bijective. As opposed to previous techniques, we do not constrain the shape of the boundary, which is free to change shape to minimize distortion.

Abstract

We present a fully automatic method for generating guaranteed bijective surface parameterizations from triangulated 3D surfaces partitioned into charts. We do so by using a distortion metric that prevents local folds of triangles in the parameterization and a barrier function that prevents intersection of the chart boundaries. In addition, we show how to modify the line search of an interior point method to directly compute the singularities of the distortion metric and barrier functions to maintain a bijective map. By using an isometric metric that is efficient to compute and a spatial hash to accelerate the evaluation and gradient of the barrier function for the boundary, we achieve fast optimization times. Unlike previous methods, we do not require the boundary be constrained by the user to a non-intersecting shape to guarantee a bijection, and the boundary of the parameterization is free to change shape during the optimization to minimize distortion.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

Keywords: parameterization, bijective mappings, free boundary

1 Introduction

Triangulated surfaces are ubiquitous in real-time graphics. However, geometry is only one aspect of how we perceive an object. Typically we annotate surfaces with other information such as color,

lighting information, or even displacements to make the surface appear more realistic and provide details beyond the resolution of the vertices of the shape. These annotations are usually performed via texture mapping, which maps two-dimensional data onto the surface of a 3D object.

Texture mapping relies on a parameterization of a surface. Given a 3D surface, we partition the shape along a connected set of edges, which we refer to as seams, into contiguous sets of triangles called charts. Parameterization is the flattening of a chart to the two-dimensional domain and the seams of the charts in 3D become the boundaries of the flattened charts in 2D. For surfaces other than developable surfaces, this flattening introduces some distortion into the shape and most parameterization methods are concerned with reducing this distortion be it in terms of deviation of angles, area, or some combination thereof.

While the quality of the parameterization is certainly important, the parameterization is of limited use for texture mapping unless it forms a bijective map between the 3D surface and the 2D texture covered by the charts. If the parameterization is not bijective, then a single point in the texture could map to multiple, disconnected regions of the surface. The result is that we cannot annotate such regions of the surface independently from one another rendering the parameterization useless for this application.

There are a few ways in which the parameterization could fail to be bijective. For example, triangles could “fold” or change orientation in the parameterization. A parameterization could also fail to be bijective if portions of the parameterization overlap. While only a handful of methods guarantee that the parameterization will be locally injective (no folds), almost none guarantee bijectivity without constraining the boundary to a non-intersecting curve. However constraining the boundary, either by user intervention or by choosing some arbitrary non-intersecting boundary curve, will produce more distortion in the parameterization than necessary since the optimization cannot modify the boundary to reduce the distortion of the parameterization.

We propose a parameterization method that guarantees the parameterization forms a bijective map. Moreover, we do not constrain the boundary to some arbitrary shape but allow the optimization to modify the boundary to reduce distortion. We discuss the class of admissible distortion metrics that guarantee local injectivity and provide a form of isometric distortion that yields an expression

*email:agjdsmith09@gmail.com

†email:schaefer@cs.tamu.edu

ACM Reference Format

Smith, J., Schaefer, S. 2015. Bijjective Parameterization with Free Boundaries. *ACM Trans. Graph.* 34, 4, Article 70 (August 2015), 9 pages. DOI = 10.1145/2766947 <http://doi.acm.org/10.1145/2766947>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766947>

Computing Locally Injective Mappings by Advanced MIPS

Xiao-Ming Fu^{*†}

Yang Liu[†]

Baining Guo^{†*}

^{*}University of Science and Technology of China

[†]Microsoft Research

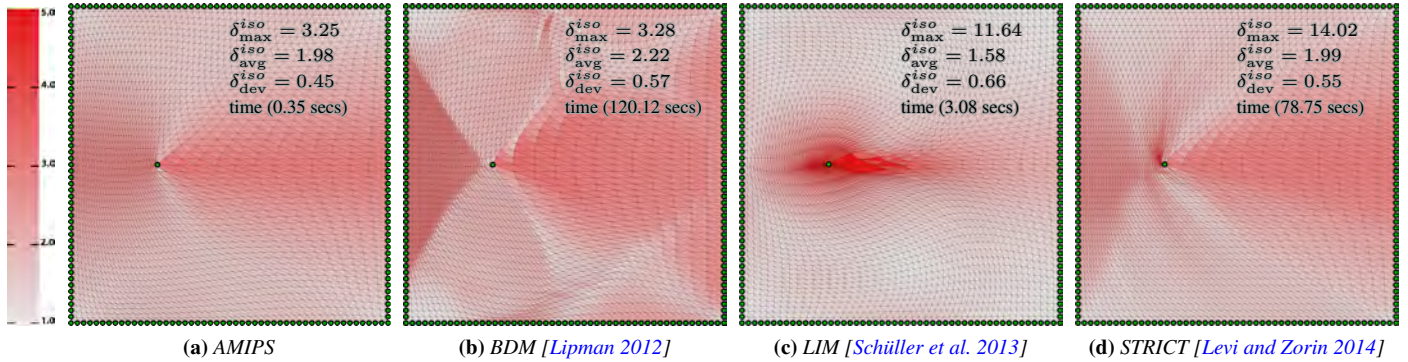


Figure 1: 2D mesh deformation: comparison with the state-of-the-art methods. The input mesh \mathcal{T} is triangulated from a 2D regular quadrilateral mesh, and an interior handle is moved from the right side to the left side while fixing all the boundary vertices. Left to right: deformation results from our AMIPS method, bounded distortion mapping [Lipman 2012], locally injective mapping [Schüller et al. 2013] and the strict minimizer [Levi and Zorin 2014]. Our method achieves the lowest maximal isometric distortion $\delta_{\max}^{\text{iso}} = \max_{\mathbf{t} \in \mathcal{T}} \max\{\sigma_{1,\mathbf{t}}^{-1}, \sigma_{2,\mathbf{t}}\}$ among all the methods with the least computation time, where $\sigma_{1,\mathbf{t}}$ and $\sigma_{2,\mathbf{t}}$ ($\sigma_{1,\mathbf{t}} \leq \sigma_{2,\mathbf{t}}$) are singular values of the Jacobian of the mapping associated with triangle \mathbf{t} . Our method also achieves more evenly and smoothly distributed distortion (see the standard deviation of the isometric distortion $\delta_{\text{dev}}^{\text{iso}}$). The color on triangles encodes the isometric distortion, with white being optimal.

Abstract

Computing locally injective mappings with low distortion in an efficient way is a fundamental task in computer graphics. By revisiting the well-known MIPS (Most-Isometric ParameterizationS) method, we introduce an advanced MIPS method that inherits the local injectivity of MIPS, achieves as low as possible distortions compared to the state-of-the-art locally injective mapping techniques, and performs one to two orders of magnitude faster in computing a mesh-based mapping. The success of our method relies on two key components. The first one is an enhanced MIPS energy function that penalizes the maximal distortion significantly and distributes the distortion evenly over the domain for both mesh-based and meshless mappings. The second is a use of the inexact block coordinate descent method in mesh-based mapping in a way that efficiently minimizes the distortion with the capability not to be trapped early by the local minimum. We demonstrate the capability and superiority of our method in various applications including mesh parameterization, mesh-based and meshless deformation, and mesh improvement.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: locally injective mapping, inexact block coordinate descent methods, parameterization, deformation, mesh improvement

1 Introduction

The task of computing a valid mapping is fundamental in many computer graphics and geometry processing applications, such as mesh parameterization, shape deformation and animation, image deformation and retargeting, remeshing and mesh improvement. A good mapping $f : \Omega \subset \mathbb{R}^d \rightarrow \mathbb{R}^d$ usually satisfies the following properties: (1) f is locally injective, i.e. $\det J(f) > 0$ where $J(f)$ is the Jacobian of f ; (2) the mapping has as low as possible distortion, for instance, the singular values of $J(f)$ should be close to 1 if an isometric mapping is desired; (3) f is smooth, which is a property especially preferred in deformation; (4) the computation of f is efficient enough to provide interactive feedback. It is also desired that the method of computing f be applicable to both 2D and 3D mappings and easily generalized to higher dimensions.

Mesh parameterization and deformation are typical applications in computer graphics where finding a low-distortion mapping is the central task. Numerous techniques have been developed in the past thirty years (cf. extensive surveys [Floater and Hormann 2005; Sheffer et al. 2006; Botsch and Sorkine 2008]). It is commonly known that linear based methods hardly satisfy all the above requirements although they are very efficient. For instance, the LSCM [Lévy et al. 2002] method cannot guarantee local injectivity and the mean value coordinate [Floater 2003a] may introduce high distortions. Instead, nonlinear approaches usually achieve lower distortion and some methods guarantee local injectivity, like MIPS [Hormann and Greiner 2000; Degener et al. 2003] and ABF++ [Sheffer et al. 2005]. However, inefficient computation and uncontrolled maximal distortions are the two main obstacles in using nonlinear approaches. Recently Schüller et al. propose a fast method to compute locally injective mappings [2013], but high distortion cannot be prevented (for instance, the region around the handle in Fig. 1c possesses high distortions). Lipman and his colleagues propose novel ways to bound the distortion explicitly [2012; 2013; 2014]. Levi and Zorin prioritize higher distortion for minimization [2014]. These methods significantly reduce distortion but with costly computation.

The compromise of low distortion and computational efficiency is

ACM Reference Format

Fu, X., Liu, Y., Guo, B. 2015. Computing Locally Injective Mappings by Advanced MIPS. ACM Trans. Graph. 34, 4, Article 71 (August 2015), 12 pages. DOI = 10.1145/2766938 <http://doi.acm.org/10.1145/2766938>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766938>

Seamless Surface Mappings

Noam Aigerman Roi Poranne Yaron Lipman
Weizmann Institute of Science




Figure 1: Two bijective seamless mappings between models of two humans are shown in (c),(d), generated by our algorithm from the two different cut-placements in (a),(b) (respectively), cuts visualized as colored curves. The two maps interpolate the same set of user-given landmarks, shown as colored spheres. The maps are visualized by texturing the male model and transferring the texture to the female model using the mappings. The algorithm is not affected by the choice of cuts: the maps do not exhibit any artifacts near the cut nor does the poor cut-correspondence (e.g. the torso in (b)) affect them, and in fact for the two different cut-placements, the produced maps are identical.

Abstract

We introduce a method for computing seamless bijective mappings between two surface-meshes that interpolates a given set of correspondences.

A common approach for computing a map between surfaces is to cut the surfaces to disks, flatten them to the plane, and extract the mapping from the flattenings by composing one flattening with the inverse of the other. So far, a significant drawback in this class of techniques is that the choice of cuts introduces a bias in the computation of the map that often causes visible artifacts and wrong correspondences.

In this paper we develop a surface mapping technique that is indifferent to the particular cut choice. This is achieved by a novel type of surface flattenings that encodes this cut-invariance, and when optimized with a suitable energy functional results in a seamless surface-to-surface map.

We show the algorithm enables producing high-quality seamless bijective maps for pairs of surfaces with a wide range of shape variability and from a small number of prescribed correspondences. We also used this framework to produce three-way, consistent and seamless mappings for triplets of surfaces.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

Keywords: surface mesh, seamless, conformal distortion, bijective simplicial mappings

ACM Reference Format

Aigerman, N., Poranne, R., Lipman, Y. 2015. Seamless Surface Mappings. *ACM Trans. Graph.* 34, 4, Article 72 (August 2015), 13 pages. DOI = 10.1145/2766921 <http://doi.acm.org/10.1145/2766921>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766921>

1 Introduction

Computation of mappings between surfaces is a core problem in computer graphics and vision, medical imaging and related fields of research. Many applications require the map to (i) be bijective, so as to establish point-to-point correspondences between the surfaces, and (ii) possess low-distortion to enable transferring of surface attributes with as little corruption as possible.

The basic task tackled in this paper is the computation of a low-distortion piecewise-linear bijective map $f : \mathbf{A} \rightarrow \mathbf{B}$ between two surfaces \mathbf{A}, \mathbf{B} , while interpolating a set of corresponding landmark points provided on the two surfaces, $p_i \in \mathbf{A}, q_i \in \mathbf{B}$.

A common approach for constructing such a map is to first cut the two surfaces consistently into topological disks (one disk, or many) and then map the disks to a common planar domain via some optimization process. The surface mapping is then induced according to the overlay of the flattened disks of one mesh over the flattened disks of the other mesh. This method often requires prescribing correspondences also along the cuts [Aigerman et al. 2014] or alternatively facing the challenging problem of optimizing also the disk-boundary correspondence and/or the cuts directly over the surfaces [Schreiner et al. 2004; Kraevoy and Sheffer 2004]. In most cases the choice of cuts affects the resulting mapping, and quite often the cut area will exhibit visible artifacts such as non-smoothness, distortion bias or wrong correspondences.

The goal of this paper is to develop an algorithm for constructing low-distortion bijective mappings of surfaces that are *seamless*, that is, indifferent to the choice of cuts, thereby relieving the user from the need to worry about cut placement. The key new insight is that there exist a set of rather simple conditions that, when enforced on the flattenings of the surfaces \mathbf{A}, \mathbf{B} , assures the resulting map $f : \mathbf{A} \rightarrow \mathbf{B}$ between the surfaces is seamless. We name flattenings which satisfy these conditions *G-flattenings*, where *G* stands for a group of transformations which relates planar copies of the mesh cut curves.

Bounded Distortion Harmonic Mappings in the Plane

Renjie Chen

University of North Carolina at Chapel Hill, USA

Ofir Weber

Bar Ilan University, Israel




Figure 1: Smooth bounded distortion harmonic mapping of the Dragon. The conformal distortion is bounded globally by 0.73 and the isometric distortion is bounded by 2.8.

Abstract

We present a framework for the computation of harmonic and conformal mappings in the plane with mathematical guarantees that the computed mappings are C^∞ , locally injective and satisfy strict bounds on the conformal and isometric distortion. Such mappings are very desirable in many computer graphics and geometry processing applications.

We establish the sufficient and necessary conditions for a harmonic planar mapping to have bounded distortion. Our key observation is that these conditions relate solely to the boundary behavior of the mapping. This leads to an efficient and accurate algorithm that supports handle-based interactive shape-and-image deformation and is demonstrated to outperform other state-of-the-art methods.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: injective mappings, conformal mappings, harmonic mappings, bounded distortion, planar deformation

1 Introduction

Computation of mappings between flat or curved surfaces is a fundamental task in computer graphics and geometry processing. Among the popular applications are shape-and-image deformation, shape interpolation-and-animation, surface parametrization, texture mapping, and finding dense correspondence between shapes for comparison and analysis. Harmonic mappings are extremely popular and were heavily used and investigated by researchers in the

last two decades [Sheffer and de Sturler 2001; Lévy et al. 2002; Zayer et al. 2005; Kharevych et al. 2006; Weber et al. 2007; Joshi et al. 2007; Springborn et al. 2008; Ben-Chen et al. 2009; Weber and Gotsman 2010].

Harmonic mappings are considered very desirable for computer graphics applications due to their smoothness and the fact that they are mathematically and computationally easy to work with. For example, fixing the boundary values of the mapping uniquely dictates its values throughout the entire domain, implying that harmonic mappings form a low dimensional subspace. This subspace is also linear in the sense that linear combinations of harmonic functions remain harmonic.

While smoothness is one of the most desirable properties when it comes to graphics, it is generally not sufficient. An artist would often like to produce deformations that are locally injective and that preserve the orientation of local shape details rather than reverse it. Moreover, it is often desirable to control and restrict the amount of angle and isometric distortion induced by the mapping. Yet, the vast majority of existing methods do not explicitly enforce constraints that impose the desirable properties but rather try to obtain these properties indirectly. The main reason is the nonconvex nature of the constraints which leads to the inability to design efficient computational methods. Recently, new techniques to efficiently handle such nonconvex problems were introduced (e.g. [Lipman 2012]), leading to an increasing interest of the scientific community in solving such challenging problems.

Methods that are based on constrained minimization (hence provide strict guarantees) are almost exclusively mesh based. They represent the mappings as piecewise linear (PWL) approximations and are nonsmooth (C^0) by construction. Refining the resolution of the meshes typically improves the result visually (at the cost of longer computation times) but cannot change the fact that mathematically PWL mappings are nonsmooth.

In contrast, the recent algorithm of [Poranne and Lipman 2014] provides the first attempt to borrow some of the advanced techniques that were developed in the mesh based setting to the meshless one. Our algorithm improves upon [Poranne and Lipman 2014] in several different aspects, allowing us to compute smooth (C^∞) harmonic and conformal deformations with strict quality guarantees.

We develop an elegant mathematical theory that provides necessary and sufficient conditions for harmonic (and as a special case conformal) mappings to be locally injective and have bounded amount

ACM Reference Format

Chen, R., Weber, O. 2015. Bounded Distortion Harmonic Mappings in the Plane. ACM Trans. Graph. 34, 4, Article 73 (August 2015), 12 pages. DOI = 10.1145/2766989 <http://doi.acm.org/10.1145/2766989>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766989>

Data-Driven Finite Elements for Geometry and Material Design

Desai Chen¹

David I.W. Levin¹²

Shinjiro Sueda¹²³

Wojciech Matusik¹

¹MIT CSAIL

²Disney Research

³California Polytechnic State University




Figure 1: Examples produced by our *data-driven* finite element method. Left: A bar with heterogeneous material arrangement is simulated 15x faster than its high-resolution counterpart. Left-Center: Our fast coarsening algorithm dramatically accelerates designing this shoe sole (up to 43x). Right-Center: A comparison to 3D printed results. Right: We repair a flimsy bridge by adding a supporting arch (8.1x) speed-up. We show a *High-Res Simulation* for comparison.

Abstract

Crafting the behavior of a deformable object is difficult—whether it is a biomechanically accurate character model or a new multimaterial 3D printable design. Getting it right requires constant iteration, performed either manually or driven by an automated system. Unfortunately, Previous algorithms for accelerating three-dimensional finite element analysis of elastic objects suffer from expensive pre-computation stages that rely on *a priori* knowledge of the object’s geometry and material composition. In this paper we introduce Data-Driven Finite Elements as a solution to this problem. Given a material palette, our method constructs a metamaterial library which is reusable for subsequent simulations, regardless of object geometry and/or material composition. At runtime, we perform fast coarsening of a simulation mesh using a simple table lookup to select the appropriate metamaterial model for the coarsened elements. When the object’s material distribution or geometry changes, we do not need to update the metamaterial library—we simply need to update the metamaterial assignments to the coarsened elements. An important advantage of our approach is that it is applicable to non-linear material models. This is important for designing objects that undergo finite deformation (such as those produced by multimaterial 3D printing). Our method yields speed gains of up to two orders of magnitude while maintaining good accuracy. We demonstrate the effectiveness of the method on both virtual and 3D printed examples in order to show its utility as a tool for deformable object design.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: Data-driven simulation, finite element methods, numerical coarsening, material design

1 Introduction

Objects with high-resolution, heterogeneous material properties are everywhere: from the output of multimaterial 3D printers to virtual characters gracing the screen in our summer blockbusters. Designing such objects is made possible by the tight coupling of design tools and numerical simulation which allows designers (or automatic algorithms) to update geometry or material parameters and subsequently estimate the physical effects of the change. Fast, accurate simulation techniques that can handle runtime changes in geometry and material composition are a necessity for such iterative design algorithms.

The gold standard technique for estimating the mechanical behavior of a deformable object under load is the finite element method (FEM). While FEM is accurate, its solution process is notoriously slow, making it a major bottleneck in the iterative design process. For this reason, there have been a large number of works on speeding up FEM simulations, and these speed improvements have enabled FEM to be used in many performance critical tasks such as computer animation, surgical training, and virtual/augmented reality. Unfortunately, even though techniques such as model reduction or numerical coarsening can achieve order-of-magnitude performance increases, they require expensive precomputation phases, typically on the order of minutes for large meshes. This precomputation requires knowledge of an object’s geometry and material composition *a priori*, something that is not known during a design task. When the user updates the model by changing the geometry or the material distribution, the preprocessing step must be run again. As shown in Fig. 2a, since this step is inside the design loop, the user cannot get rapid feedback on the changes made to the object.

We propose *Data-Driven FEM* (DDFEM), a new simulation methodology that removes these limitations and is thus extremely well-suited to the types of design problems discussed above. We divide an object into a set of deformable voxels using embedded finite elements and coarsen these voxels hierarchically. A custom metamaterial database is populated with materials that minimize the error incurred by coarsening. This database is learned *once* in a completely offline fashion and depends only on the set of materials to be used by the deformable object and not on the actual material distribution and geometry. At runtime we use the database to perform fast coarsening of an FEM mesh in a way that is agnostic to changes in geometry and material composition of the object. The key features of the algorithm are its ability to handle arbitrary, nonlinear elastic

ACM Reference Format

Chen, D., Levin, D., Sueda, S., Matusik, W. 2015. Data-Driven Finite Elements for Geometry and Material Design. ACM Trans. Graph. 34, 4, Article 74 (August 2015), 10 pages. DOI = 10.1145/2766889 <http://doi.acm.org/10.1145/2766889>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766889>

Nonlinear Material Design Using Principal Stretches

Hongyi Xu Univ. of Southern California Funshing Sin Univ. of Southern California Yufeng Zhu Univ. of Southern California Jernej Barbic Univ. of Southern California
Blizzard Entertainment Univ. of British Columbia




Figure 1: Design of isotropic nonlinear materials: The soft-body motion of the wrestler was computed using FEM, constrained to a motion capture skeletal dancing animation. Using our method, we designed a nonlinear isotropic material that performs well both during impulsive and gentle animation phases. Top row: the wrestler is performing high jumps. The soft Neo-Hookean material exhibits artifacts (belly, thighs) when the character moves abruptly. Our material and the stiff Neo-Hookean material produce good deformations. Bottom row: deformations during a gentle phase (walking while dancing) of the same motion sequence. The soft Neo-Hookean material and our method produce rich small-deformation dynamics, whereas the stiff Neo-Hookean material inhibits it. The Young's modulus of material (a) was chosen to produce good dynamics during gentle motion. We then edited it to address impulsive motion, producing (c). The stiff material in (b) is the best matching material to (c) among Neo-Hookean materials, minimizing the L_2 material curve difference to (c).

Abstract

The Finite Element Method is widely used for solid deformable object simulation in film, computer games, virtual reality and medicine. Previous applications of nonlinear solid elasticity employed materials from a few standard families such as linear corotational, nonlinear St.Venant-Kirchhoff, Neo-Hookean, Ogden or Mooney-Rivlin materials. However, the spaces of all nonlinear isotropic and anisotropic materials are infinite-dimensional and much broader than these standard materials. In this paper, we demonstrate how to intuitively explore the space of isotropic and anisotropic nonlinear materials, for design of animations in computer graphics and related fields. In order to do so, we first formulate the internal elastic forces and tangent stiffness matrices in the space of the principal stretches of the material. We then demonstrate how to design new isotropic materials by editing a single stress-strain curve, using a spline interface. Similarly, anisotropic (orthotropic) materials can be designed by editing three curves, one

for each material direction. We demonstrate that modifying these curves using our proposed interface has an intuitive, visual, effect on the simulation. Our materials accelerate simulation design and enable visual effects that are difficult or impossible to achieve with standard nonlinear materials.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling, I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: FEM, isotropic, anisotropic, material, design

1 Introduction

Three-dimensional solid Finite Element Method (FEM) simulations are widely used in computer graphics, animation and related fields. FEM simulations, however, are greatly influenced by the specific material relationship between the displacements (strains) and elastic forces (stresses). To date, applications in computer graphics use linear materials, or nonlinear materials where the strain-stress relationship is specified using a global equation, such as the linear corotational, St.Venant-Kirchhoff, Neo-Hookean, Ogden or Mooney-Rivlin materials [Bonet and Wood 1997; Bower 2011]. These materials, however, only scratch the surface of the space of all isotropic nonlinear elastic materials. The space of all anisotropic nonlinear materials is even broader. In order to more easily achieve complex visual effects, we approach the problem of how to model and design arbitrary nonlinear materials, for use in animations in computer graphics and related fields. We give a method to design isotropic

ACM Reference Format

Xu, H., Sin, F., Zhu, Y., Barbic, J. 2015. Nonlinear Material Design Using Principal Stretches. ACM Trans. Graph. 34, 4, Article 75 (August 2015), 11 pages. DOI = 10.1145/2766917 <http://doi.acm.org/10.1145/2766917>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://doi.acm.org/10.1145/2766917>

Subspace Condensation: Full Space Adaptivity for Subspace Deformations

Yun Teng^{*1,2}, Mark Meyer^{†2}, Tony DeRose^{‡2} and Theodore Kim^{§1}

¹University of California, Santa Barbara

²Pixar Animation Studios




Figure 1: (a) The simulation runs at 16 FPS, entirely within the subspace, $67\times$ faster than a full space simulation over the entire mesh. (b) Novel wall collisions begin, activating full space tets, shown in red in the inset. The simulation still runs at 2.1 FPS, a $7.7\times$ speedup. (c) Collisions produce a deformation far outside the basis, and 49% of the tets are simulated in full space. The step runs at 0.5 FPS; still a $1.9\times$ speedup. (d) The collisions are removed, and the $67\times$ speedup returns.

Abstract

Subspace deformable body simulations can be very fast, but can behave unrealistically when behaviors outside the prescribed subspace, such as novel external collisions, are encountered. We address this limitation by presenting a fast, flexible new method that allows full space computation to be activated in the neighborhood of novel events while the rest of the body still computes in a subspace. We achieve this using a method we call *subspace condensation*, a variant on the classic *static condensation* precomputation. However, instead of a precomputation, we use the speed of subspace methods to perform the condensation at every frame. This approach allows the full space regions to be specified arbitrarily at runtime, and forms a natural two-way coupling with the subspace regions. While condensation is usually only applicable to linear materials, the speed of our technique enables its application to non-linear materials as well. We show the effectiveness of our approach by applying it to a variety of articulated character scenarios.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling

Keywords: character simulation, subspace integration, static condensation, cubature, collision resolution

*yunteng.cs@cs.ucsb.edu

†mmeyer@pixar.com

‡derose@pixar.com

§kim@mat.ucsb.edu

ACM Reference Format

Teng, Y., Meyer, M., DeRose, T., Kim, T. 2015. Subspace Condensation: Full Space Adaptivity for Subspace Deformations. ACM Trans. Graph. 34, 4, Article 76 (August 2015), 9 pages. DOI = 10.1145/2766904 <http://doi.acm.org/10.1145/2766904>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766904>

1 Introduction

Subspace methods, also known as reduced-order, reduced coordinate or model reduction methods, have recently made great strides in accelerating deformable body simulations. In lieu of a full space method, also known as a full-order or full-coordinate method, which simulates every degree of freedom in a mesh, subspace methods instead constrain the motion to a subspace spanned by a compact, but expressive, set of basis vectors. Since r basis vectors are being simulated instead of N vertices, if $r \ll N$, very large speedups can be realized.

An obvious limitation arises when the expressivity of the basis vectors is insufficient, and the true, full space motion of the mesh lies outside the span of the subspace. A straightforward example of this is an external collision, such as a cannonball hitting the mesh in a novel location that was not accounted for when constructing the subspace. In these cases, subspace methods can “lock”, producing motions that are both very different from the full space solution and unrealistic in appearance (Fig. 2). A variety of strategies have been developed to account for this situation, including basis enrichment [Harmon and Zorin 2013], adaptive basis construction [Hahn et al. 2014], and falling back to brute-force, full-order computation over the entire mesh [Kim and James 2009].



Figure 2: Novel collisions cause extreme locking in a subspace-only simulation.

In this paper, we present a distinctly different approach. We observe that in many cases, particularly when dealing with external collisions, subspace methods only diverge from the full space solution in spatially localized patches. Unfortunately, there is no way to know during the precomputation stage where these patches will be,

Perceptually Based Downsampling of Images

A. Cengiz Öztireli
ETH Zurich

Markus Gross
ETH Zurich




Figure 1: It is challenging to reproduce a perceptually similar downscaled version of an image, as shown here for simple subsampling, the commonly used bicubic filter, and the state-of-the-art method of Kopf et al. [2013] (“Content-Adaptive”). Relying on a perceptual image quality measure instead of standard metrics, our method (“Perceptual”) is able to preserve perceptually important features and the overall look of the original images. Left input image courtesy of Flickr user Matteo Catanese.

Abstract

We propose a perceptually based method for downsampling images that provides a better apparent depiction of the input image. We formulate image downsampling as an optimization problem where the difference between the input and output images is measured using a widely adopted perceptual image quality metric. The downscaled images retain perceptually important features and details, resulting in an accurate and spatio-temporally consistent representation of the high resolution input. We derive the solution of the optimization problem in closed-form, which leads to a simple, efficient and parallelizable implementation with sums and convolutions. The algorithm has running times similar to linear filtering and is orders of magnitude faster than the state-of-the-art for image downsampling. We validate the effectiveness of the technique with extensive tests on many images, video, and by performing a user study, which indicates a clear preference for the results of the new algorithm.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Sampling;

Keywords: images, video, perceptual, downsampling, structural similarity, unsharp masking

1 Introduction

Image downsampling is a fundamental operation performed constantly in digital imaging. The abundance of high resolution capture devices and the variety of displays with different resolutions

make it an essential component of virtually any application involving images or video. However, this problem has so far received substantially less attention than other sampling alterations.

Classical downsampling algorithms aim at minimizing aliasing artifacts by linearly filtering the image via convolution with a kernel before subsampling and subsequent reconstruction, following the sampling theorem [Shannon 1998]. However, along with aliasing, these strategies also smooth out some of the perceptually important details and features, as shown in Figure 1, since the kernels used are agnostic to the image content. A solution to this problem is adapting the kernel shapes to local image patches [Kopf et al. 2013] in the spirit of bilateral filtering [Tomasi and Manduchi 1998], so that they are better aligned with the local image features to be preserved. This strategy can significantly increase the crispness of the features while avoiding ringing artifacts typical for post-sharpening filters. However, it still cannot capture all perceptually relevant details, and as a result, might distort some of the perceptually important features and the overall look of the input image (Figure 1, content-adaptive), or lead to artifacts such as jagged edges [Kopf et al. 2013].

Loss of some of the perceptually important features and details stems from the common shortcoming of these methods that they operate with simple error metrics that are known to correlate poorly with human perception [Wang and Bovik 2009]. Significant improvements have been obtained for many problems in image processing by replacing these classical metrics with perceptually based image quality metrics [Zhang et al. 2012; He et al. 2014].

In this paper, we propose a perceptually based method for downsampling images. We formulate the downsampling problem as an optimization where we solve for the downscaled output image given the input image. The error between the two images is measured using the widely adopted structural similarity (SSIM) index [Wang et al. 2004]. The use of SSIM in optimization problems has been hindered by the resulting non-linear non-convex error functions [Brunet et al. 2012]. However, we show that for the downsampling problem, it is possible to derive a closed-form solution to this optimization. The solution leads to a non-linear filter, which involves computing local luminance and contrast measures on the original and a smoothed version of the input image. Although the filter is seemingly different than SSIM without any covariance term, we show that it maximizes the mean SSIM between the original and

ACM Reference Format

Öztireli, A., Gross, M. 2015. Perceptually Based Downsampling of Images. ACM Trans. Graph. 34, 4, Article 77 (August 2015), 10 pages. DOI = 10.1145/2766891 <http://doi.acm.org/10.1145/2766891>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766891>

An L_1 Image Transform for Edge-Preserving Smoothing and Scene-Level Intrinsic Decomposition

Sai Bi

Xiaoguang Han

Yizhou Yu

The University of Hong Kong




Figure 1: Our algorithm for piecewise image flattening facilitates both edge-preserving smoothing and intrinsic decomposition. Two examples of edge-preserving smoothing are shown in (a)-(d), and one example of intrinsic decomposition is shown in (e)-(h). Intrinsic decomposition enables image editing, such as re-texturing (i). Original images courtesy Flickr users 47765927@N06 (a), 132341054@N03 (c) and 37213589@N08 (e).

Abstract

Identifying sparse salient structures from dense pixels is a long-standing problem in visual computing. Solutions to this problem can benefit both image manipulation and understanding. In this paper, we introduce an image transform based on the L_1 norm for piecewise image flattening. This transform can effectively preserve and sharpen salient edges and contours while eliminating insignificant details, producing a nearly piecewise constant image with sparse structures. A variant of this image transform can perform edge-preserving smoothing more effectively than existing state-of-the-art algorithms. We further present a new method for complex scene-level intrinsic image decomposition. Our method relies on the above image transform to suppress surface shading variations, and perform probabilistic reflectance clustering on the flattened image instead of the original input image to achieve higher accuracy. Extensive testing on the *Intrinsic-Images-in-the-Wild* database indicates our method can perform significantly better than existing techniques both visually and numerically. The obtained intrinsic images have been successfully used in two applications, surface re-

texturing and 3D object compositing in photographs.

CR Categories: I.4.3 [Image Processing and Computer Vision]: Enhancement—Smoothing; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Color, Shading;

Keywords: Salient Structures, Piecewise Image Flattening, Probabilistic Clustering, Intrinsic Images, Sparse Signal Recovery

1 Introduction

An image has hundreds of thousands of pixels. To effectively carry out an image manipulation or understanding task, such as image classification, segmentation, and stylization, it is vital to extract sparse salient structures, including perceptually important edges and contours, from such a large number of pixels. This is indeed the goal of contour detection and edge-preserving smoothing, which can be dated back to anisotropic smoothing [Perona and Malik 1990]. Ideally, the result of edge-preserving smoothing is a piecewise constant image with discontinuities occurring along salient edges. Spatial color variations that should be removed during such smoothing include both high-frequency and low-frequency signals. While high-frequency signals can be removed relatively easily with local filters, such as the bilateral filter [Tomasi 1998], low-frequency signals are more “stubborn” and removing them requires more global operations.

Interestingly, a similar challenge exists in intrinsic image decomposition, which attempts to separate reflectance and shading from their product, and has applications in surface re-texturing, object insertion and scene relighting. There is a recent trend on this topic to

ACM Reference Format

Bi, S., Han, X., Yu, Y. 2015. An L_1 Transform for Edge-Preserving Smoothing and Scene-Level Intrinsic Image Decomposition. *ACM Trans. Graph.* 34, 4, Article 78 (August 2015), 12 pages. DOI = 10.1145/2766946 <http://doi.acm.org/10.1145/2766946>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766946>

A Computational Approach for Obstruction-Free Photography

Tianfan Xue^{1*}

Michael Rubinstein^{2†}

Ce Liu^{2†}

William T. Freeman^{1,2}

¹MIT CSAIL

²Google Research

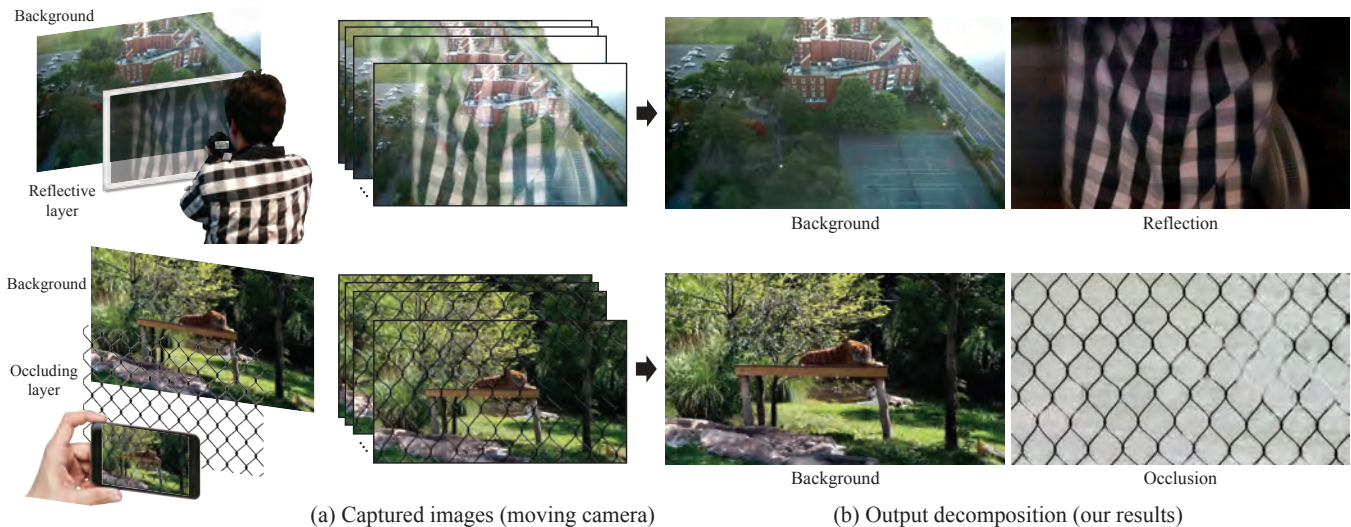


Figure 1: In this paper we present an algorithm for taking pictures through reflective or occluding elements such as windows and fences. The input to our algorithm is a set of images taken by the user while slightly scanning the scene with a camera/phone (a), and the output is two images: a clean image of the (desired) background scene, and an image of the reflected or occluding content (b). Our algorithm is fully automatic, can run on mobile devices, and allows taking pictures through common visual obstacles, producing images as if they were not there. The full image sequences and a closer comparison between the input images and our results are available in the supplementary material.

Abstract

We present a unified computational approach for taking photos through reflecting or occluding elements such as windows and fences. Rather than capturing a single image, we instruct the user to take a short image sequence while slightly moving the camera. Differences that often exist in the relative position of the background and the obstructing elements from the camera allow us to separate them based on their motions, and to recover the desired background scene as if the visual obstructions were not there. We show results on controlled experiments and many real and practical scenarios, including shooting through reflections, fences, and raindrop-covered windows.

CR Categories: I.3.7 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance; I.4.3 [Image Processing and Computer Vision]: Enhancement

Keywords: reflection removal, occlusion removal, image and video decomposition

1 Introduction

Many imaging conditions are far from optimal, forcing us to take our photos through reflecting or occluding elements. For example, when taking pictures through glass windows, reflections from indoor objects can obstruct the outdoor scene we wish to capture (Figure 1, top row). Similarly, to take pictures of animals in the zoo, we may need to shoot through an enclosure or a fence (Figure 1, bottom row). Such visual obstructions are often impossible to avoid just by changing the camera position or the plane of focus, and state-of-the-art computational approaches are still not robust enough to remove such obstructions from images with ease. More professional solutions, such as polarized lenses (for reflection removal), which may alleviate some of those limitations, are not accessible to the everyday user.

In this paper, we present a robust algorithm that allows a user to take photos through obstructing layers such as windows and fences, producing images of the desired scene as if the obstructing elements were not there. Our algorithm only requires the users to generate some camera motion during the imaging process, while the rest of the processing is fully automatic.

We exploit the fact that reflecting or obstructing planes are usually situated in-between the camera and the main scene, and as a result, have different depth than the main scene we want to capture. Thus, instead of taking a single picture, we instruct the photographer to take a short image sequence while slightly moving the camera—an interaction similar to taking a panorama (with camera motion

* Part of this work was done while the author was an intern at Microsoft Research New England.

† Part of this work was done while the authors were at Microsoft Research.

ACM Reference Format

Xue, T., Rubinstein, M., Liu, C., Freeman, W. 2015. A Computational Approach for Obstruction-Free Photography. ACM Trans. Graph. 34, 4, Article 79 (August 2015), 11 pages. DOI = 10.1145/2766940 <http://doi.acm.org/10.1145/2766940>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright held by the Owner/Author.

SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.

ACM 978-1-4503-3331-3/15/08.

DOI: <http://dx.doi.org/10.1145/2766940>

Dynamic Terrain Traversal Skills Using Reinforcement Learning

Xue Bin Peng Glen Berseth Michiel van de Panne*

University of British Columbia



Figure 1: Real-time planar simulation of a dog capable of traversing terrains with gaps, walls, and steps. The control policy for this skill is computed offline using reinforcement learning. The ground markings indicate the landing points for the front and hind legs.

Abstract

The locomotion skills developed for physics-based characters most often target flat terrain. However, much of their potential lies with the creation of dynamic, momentum-based motions across more complex terrains. In this paper, we learn controllers that allow simulated characters to traverse terrains with gaps, steps, and walls using highly dynamic gaits. This is achieved using reinforcement learning, with careful attention given to the action representation, non-parametric approximation of both the value function and the policy; epsilon-greedy exploration; and the learning of a good state distance metric. The methods enable a 21-link planar dog and a 7-link planar biped to navigate challenging sequences of terrain using bounding and running gaits. We evaluate the impact of the key features of our skill learning pipeline on the resulting performance.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: Computer Animation, Physics Simulation

1 Introduction

Physics-based simulations arguably offer some of the best prospects for creating models of human and animal motion that generalize well across a wide range of situations. The principle challenge to be solved is that of control, namely determining what actions should be applied over time in order to produce movements that solve a motion task in a natural and robust fashion. Significant progress has been made with regard to motions such as walking, running, and other specific motions, such as falling and rolling. This has been achieved using a variety of control methods including those based on hand-crafted abstractions and feedback laws; policy search to optimize suitable feedback laws; and model-predictive control methods.

In this paper we expand the capabilities of simulated articulated figures by learning control strategies for highly dynamic motions across challenging terrains that contain gaps, walls, and steps. These skills require actions, i.e., jumps and leaps, that are compatible with leaping over (or onto) upcoming terrain obstacles, while also being compatible with the current body state of the character. The interdependence of terrain-obstacle sequences, body-state sequences, and action sequences are a defining characteristic of the problem that needs to be solved. The final control policies should produce actions that are both robust and efficient.

Our solution employs a reinforcement learning (RL) method that is characterized by:

- continuous and high-dimensional states and actions;
- value-iteration based on a set of (s, a, r, s') transition tuples, using positive temporal difference (PTD) updates, followed by tuple-culling;
- non-parametric kernel-based value function and policy approximation, with outlier removal;
- an embodied state representation and an optimization process for learning a good state distance metric; and
- a progressive learning approach that alternates between ϵ -greedy exploration and value iteration.

Our contribution lies in showing that, taken together, this set of algorithmic choices opens the door to the development of new classes of skills for simulated characters. The difficulty of applying RL to continuous state and action spaces is well known, and thus our work provides insights into how RL methods can be successfully applied to such problems. We evaluate the method on a simulated 21-link planar dog and a 7-link planar biped that learn to navigate terrains with sequences of gaps, walls, and steps. We investigate the impact of the various design decisions that characterize our approach.

2 Related Work

The use of physics-based models in computer graphics has yielded tremendous dividends in rendering, geometric modeling, and the animation of phenomena such as fluids and cloth. Similar dividends are available for physics-based character animation, although progress in this area has been tempered by the difficulty of modeling the motor control skills of humans and animals that are also required to drive the simulations. This topic has a rich history and we refer the reader to a recent survey paper [Geijtenbeek and Pronost 2012] for an overview.

*e-mail: xbpeng|gberseth|van@cs.ubc.ca

ACM Reference Format

Peng, X., Berseth, G., van de Panne, M. 2015. Dynamic Terrain Traversal Skills Using Reinforcement Learning. ACM Trans. Graph. 34, 4, Article 80 (August 2015), 11 pages. DOI = 10.1145/2766910 <http://doi.acm.org/10.1145/2766910>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766910>

Online Control of Simulated Humanoids Using Particle Belief Propagation

Perttu Hämäläinen^{1*} Joose Rajamäki^{1†} C. Karen Liu^{2‡}

¹ Aalto University

² Georgia Tech



Figure 1: Our algorithm can handle complex balancing and manipulation tasks while adapting to user interactions. All our demonstrated movements emerge from simple cost functions without animation data or offline precomputation. More examples can be found in the supplemental video and on the project homepage.

Abstract

We present a novel, general-purpose Model-Predictive Control (MPC) algorithm that we call Control Particle Belief Propagation (C-PBP). C-PBP combines multimodal, gradient-free sampling and a Markov Random Field factorization to effectively perform simultaneous path finding and smoothing in high-dimensional spaces. We demonstrate the method in online synthesis of interactive and physically valid humanoid movements, including balancing, recovery from both small and extreme disturbances, reaching, balancing on a ball, juggling a ball, and fully steerable locomotion in an environment with obstacles. Such a large repertoire of movements has not been demonstrated before at interactive frame rates, especially considering that all our movement emerges from simple cost functions. Furthermore, we abstain from using any precomputation to train a control policy offline, reference data such as motion capture clips, or state machines that break the movements down into more manageable subtasks. Operating under these conditions enables rapid and convenient iteration when designing the cost functions.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: animation, motion synthesis, motion planning, optimization

*e-mail: perttu.hamalainen@aalto.fi

†e-mail: joose.rajamaki@aalto.fi

‡e-mail: karenliu@cc.gatech.edu

ACM Reference Format
Hämäläinen, P., Rajamäki, J., Liu, C. 2015. Online Control of Simulated Humanoids Using Particle Belief Propagation. *ACM Trans. Graph.* 34, 4, Article 81 (August 2015), 13 pages. DOI = 10.1145/2767002 <http://doi.acm.org/10.1145/2767002>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2767002>

1 Introduction

Research on procedural, physically based humanoid movement synthesis has attracted considerable attention, due to its promise of transforming the animator or game designer into a choreographer that directs the characters using high-level commands instead of laboriously editing animation frame-by-frame. As shown by the seminal work of Witkin and Kass [1988], physically based motion synthesis can be formulated as an optimization problem, yielding plausible movement with desirable qualities such as “squash-and-stretch” and anticipation. The field has since progressed from simple characters with a few rigid bodies towards humanoid models of increasing biomechanical detail, from offline to real-time simulation, and from purely optimization-based synthesis to utilizing libraries of motion capture and animation data. Detailed reviews of the state-of-the-art can be found in [Geijtenbeek and Pronost 2012; Guo et al. 2014; Pejsa and Pandzic 2010]. Although complex movements like bipedal locomotion can now be simulated in real-time, challenges remain in making the characters autonomously adapt to unpredictable environments, such as obstacles to foot placement and sudden changes in direction when steered interactively.

In this paper, we set us the challenge of avoiding the following approaches that most physically based real-time systems rely on: 1) using pre-scripted or recorded data such as motion capture clips, 2) employing offline precomputation to learn a control policy, or 3) designing specialized state machines that break down movement into more manageable parts, e.g., defining the recovery of balance as first lifting a foot, and then placing the foot down in a prescribed direction. All the approaches require significant offline effort and/or result in less robust real-time behavior in novel situations. Instead, we utilize the power of our online optimization algorithm to make adaptive real-time behavior emerge from simple cost functions.

Our work utilizes the fact that current personal computers can simulate rigid body dynamics hundreds of times faster than real-time. This has presented an opportunity for a new class of Model-Predictive Control (MPC) methods that use black-box dynamics simulators such as Open Dynamics Engine to simulate a controlled system forward at each control update up to a planning horizon of a few seconds. The approach has been gaining popularity in both animation and robotics research, and has been demonstrated to work with simulated humanoids and the requirements above

Intuitive and Efficient Camera Control with the Toric Space

Christophe Lino*
INRIA Rennes / IRISA, France

Marc Christie†
University of Rennes 1 / IRISA, France




Figure 1: We present (a) the Toric space, a novel and compact representation for intuitive and efficient virtual camera control. We demonstrate the potential of this representation by proposing (b) an efficient automated viewpoint computation technique, (c) a novel and intuitive screen-space manipulation tool, and (d) an effective viewpoint interpolation technique.

Abstract

A large range of computer graphics applications such as data visualization or virtual movie production require users to position and move viewpoints in 3D scenes to effectively convey visual information or tell stories. The desired viewpoints and camera paths are required to satisfy a number of visual properties (e.g. size, vantage angle, visibility, and on-screen position of targets). Yet, existing camera manipulation tools only provide limited interaction methods and automated techniques remain computationally expensive.

In this work, we introduce the *Toric space*, a novel and compact representation for intuitive and efficient virtual camera control. We first show how visual properties are expressed in this Toric space and propose an efficient interval-based search technique for automated viewpoint computation. We then derive a novel screen-space manipulation technique that provides intuitive and real-time control of visual properties. Finally, we propose an effective viewpoint interpolation technique which ensures the continuity of visual properties along the generated paths. The proposed approach (i) performs better than existing automated viewpoint computation techniques in terms of speed and precision, (ii) provides a screen-space manipulation tool that is more efficient than classical manipulators and easier to use for beginners, and (iii) enables the creation of complex camera motions such as long takes in a very short time and in a controllable way. As a result, the approach should quickly find its place in a number of applications that require interactive or automated camera control such as 3D modelers, navigation tools or 3D games.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques— [G.1.6]: Mathematics of Computing—Numerical Analysis[Optimization]

Keywords: Virtual Camera Composition, Interactive Camera

Control, Virtual Camera Interpolation

1 Introduction

Virtual camera control is an essential component of many computer graphics applications. The virtual camera – as a window on 3D contents – conveys information, sense of aesthetics, and emotion. The proper selection of viewpoints and the proper design of camera paths are therefore of prime importance to precisely convey intended effects. Furthermore, the increased availability of realistic real-time rendering workstations as well as mobile devices and their growing usage in our everyday tasks, both call for interactive and automated techniques that would simplify the creation of effective viewpoints and speed up the overall design process.

To address this task, modeling and animation software propose different camera manipulation tools. However, most tools rely on the underlying mathematical representations of cameras and camera paths. A camera is therefore manipulated through a sequence of translation and rotation operations like any other node of the scene graph, and a camera path is constructed through a sequence of manually controlled spline-interpolated key-frames. While some visual widgets may assist the users, the precise control of a viewpoint remains a complex task, especially for beginners.

In the literature, a number of techniques have been proposed to ease the control of virtual cameras through contributions such as screen-space manipulations, automated viewpoint computation from visual properties, or automated path-planning techniques. See [Christie et al. 2008] for a detailed overview. However most contributions address a single aspect at a time (viewpoint computation, camera manipulation or path-planning). Furthermore, the problem of placing and moving virtual cameras has essentially been addressed through

ACM Reference Format

Lino, C., Christie, M. 2015. Intuitive and Efficient Camera Control with the Toric Space. ACM Trans. Graph. 34, 4, Article 82 (August 2015), 12 pages. DOI = 10.1145/2766965 <http://doi.acm.org/10.1145/2766965>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766965>

*e-mail: christophe.lino@irisa.fr

†e-mail: marc.christie@irisa.fr

Interaction Context (ICON): Towards a Geometric Functionality Descriptor

Ruizhen Hu^{1,2,3*} Chenyang Zhu¹ Oliver van Kaick⁴ Ligang Liu⁵ Ariel Shamir⁶ Hao Zhang¹
¹Simon Fraser University ²SIAT ³Zhejiang University ⁴Carleton University ⁵USTC ⁶IDC

Abstract

We introduce a *contextual descriptor* which aims to provide a *geometric* description of the *functionality* of a 3D object in the context of a given scene. Differently from previous works, we do not regard functionality as an abstract label or represent it implicitly through an agent. Our descriptor, called *interaction context* or *ICON* for short, explicitly represents the geometry of object-to-object interactions. Our approach to object functionality analysis is based on the key premise that functionality should mainly be derived from interactions between objects and not objects in isolation. Specifically, ICON collects geometric and structural features to encode interactions between a central object in a 3D scene and its surrounding objects. These interactions are then grouped based on feature similarity, leading to a hierarchical structure. By focusing on interactions and their organization, ICON is insensitive to the numbers of objects that appear in a scene, the specific disposition of objects around the central object, or the objects' fine-grained geometry. With a series of experiments, we demonstrate the potential of ICON in functionality-oriented shape processing, including shape retrieval (either directly or by complementing existing shape descriptors), segmentation, and synthesis.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms.

Keywords: object functionality analysis, contextual descriptor, shape similarity, shape retrieval

"The essential definition of object classes is functional."

— Stark & Bowyer [1996]

1 Introduction

Recently in shape analysis, an increasing effort has been devoted to extracting high-level and semantic information from geometric objects and datasets [Mitra et al. 2013], especially man-made shapes. It is arguable that an important goal of some of these developments is to obtain a functional understanding of objects and object categories. The *functionality* of an object usually refers to the particular use for which the object is designed, while different interpretations of this concept are possible. For example, functionality can be defined as *the application of an object in a specific context for the accomplishment of a particular purpose* [Bogoni and Bajcsy 1995].

*ruizhen.hu@gmail.com

ACM Reference Format

Hu, R., Zhu, C., van Kaick, O., Liu, L., Shamir, A., Zhang, H. 2015. Interaction Context (ICON): Towards a Geometric Functionality Descriptor. ACM Trans. Graph. 34, 4, Article 83 (August 2015), 12 pages. DOI = 10.1145/2766914 <http://doi.acm.org/10.1145/2766914>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766914>



Figure 1: Similarity between shapes (top) vs. similarity between functionalities (bottom). A shape descriptor (LFD) considers the middle cart more similar to the desk, as shown on the left using a 2D MDS projection of the distances between objects. Our contextual descriptor, interaction context or ICON, takes into account object-to-object interactions and identifies the two carts as more similar.

Ongoing pursuits on functional shape analysis have represented functionality in different manners. Some methods take a *category-specific* approach, relying on functionality models handcrafted for specific object categories [Sutton et al. 1994; Pechuk et al. 2008]. Another line of works characterize object functionalities *implicitly* through a human agent interacting with an object [Kim et al. 2014; Zhu et al. 2014; Liu et al. 2014] to detect its *affordances*, i.e., object properties that allow a person to perform a certain action. Finally, other works represent functionalities as labels such as "to support" and "to be held" [Pechuk et al. 2008; Laga et al. 2013].

The key premise of our approach is that object functionality should mainly be derived from *interactions* between objects and not an object in isolation [Caine 1994]. Hence, to analyze the functionality of an object, the object needs to be provided in a *context*, i.e., a surrounding 3D scene, to accomplish its functional purpose [Bogoni and Bajcsy 1995]. Moreover, given that the information that is present in 3D scenes is the *geometry* of the objects, our specific focus is to represent interactions inferable from object geometry or form, reflecting an attempt to invert the well-known notion of "form follows function" and develop a *geometric functionality descriptor*.

In this work, we introduce a contextual shape descriptor we call *interaction context* or *ICON*, for short. ICON encodes pairwise and localized interaction relations between a *central* object and its surrounding objects and organizes them in a meaningful manner. In contrast to previous works, our contextual description is not category-specific, and different from affordance analyses, the interactions we consider are not confined to those involving humans. By representing the geometry of the context of interactions of the central object, which are an important cue to infer the functionality of objects, we believe that ICON can constitute the starting point for developing a geometric functionality descriptor (Figure 1).

However, there are several challenges in defining a contextual descriptor using interactions:

- The descriptor of the central object should not be sensitive to specific counts of objects or their fine-grained geometry. For example, bookcases and curio cabinets are regarded as functionally equivalent even though there is significant variation in the number and the kinds of objects displayed.

Elements of Style: Learning Perceptual Shape Style Similarity

Zhaoliang Lun¹

Evangelos Kalogerakis¹

Alla Sheffer²

¹University of Massachusetts Amherst

²University of British Columbia




Figure 1: (left) Changing the style of objects in a scene influences the sense of time and place. (right) Style similarity transcends structure: in the top row, the bed A is pronouncedly more similar, style-wise, to dresser B than C; in the bottom row, buildings A and C are stylistically more similar (insets highlight some stylistically similar elements).

Abstract

The human perception of stylistic similarity transcends structure and function: for instance, a bed and a dresser may share a common style. An algorithmically computed style similarity measure that mimics human perception can benefit a range of computer graphics applications. Previous work in style analysis focused on shapes within the same class, and leveraged structural similarity between these shapes to facilitate analysis. In contrast, we introduce the first *structure-transcending* style similarity measure and validate it to be well aligned with human perception of stylistic similarity. Our measure is inspired by observations about style similarity in art history literature, which point to the presence of similarly shaped, salient, *geometric elements* as one of the key indicators of stylistic similarity. We translate these observations into an algorithmic measure by first quantifying the geometric properties that make humans perceive geometric elements as similarly shaped and salient in the context of style, then employing this quantification to detect pairs of matching style related elements on the analyzed models, and finally collating the element-level geometric similarity measurements into an object-level style measure consistent with human perception. To achieve this consistency we employ crowdsourcing to quantify the different components of our measure; we learn the relative perceptual importance of a range of elementary shape distances and other parameters used in our measurement from 50K responses to cross-structure style similarity queries provided by over 2500 participants. We train and validate our method on this dataset, showing it to successfully predict relative style similarity with near 90% accuracy based on 10-fold cross-validation.

CR Categories: I.3.5 [Computing Methodologies]: Computer Graphics—Computational Geometry and Object Modeling;

Keywords: style similarity, crowdsourcing, machine learning

1 Introduction

Human perception of style similarity transcends structure and function; we can meaningfully discuss style similarity between a cup and a coffee pot, a bed and a dresser, or a church and a castle. Style coordination across heterogeneous object arrangements greatly contributes to their overall aesthetics, and significantly improves the believability of virtual scenes (Figure 1, left). Thus, when designing both real and virtual environments, artists and designers put significant effort into generating style coordinated object arrangements at all scales - from putting together a place-setting at a table, through room furnishing, and all the way to design of building ensembles and cityscapes. This task requires users to navigate heterogeneous object databases based on style similarity. Such style-based database navigation can be significantly accelerated by the availability of a measure that can robustly evaluate style similarity between structurally different models and detect models which share a similar style despite large functional differences (Figure 1, right). While previous work focused on evaluating style similarity between objects with similar overall structure (Section 2), we introduce the first *structure-transcending* method for style similarity evaluation between 3D shapes, and validate that it is well aligned with human perception.

Our similarity measure is motivated by observations about human perception of style hinted at by art history and appraisal literature. Art history experts often classify objects as belonging to a particular geographic or temporal style by looking at salient *geometric elements* on the objects with recurring *visual motifs* [Nutting 1928; Blumenson 1995]. For instance, classical Byzantine churches are likely to have rounded domes and arches, while Gothic structures are dominated by steep gables and flying buttresses (Figure 1, bottom row A and B). While style extends beyond the search for motif level similarity, our work focuses on the role of common motifs in style evaluation. Our style similarity metric is therefore designed around the presence of pairs of similarly shaped, or matching, salient geometric elements on the evaluated models. The relative size of such elements, their number, and the percentage of the object's surface covered by them can vary dramatically (Figure 1, right), making detection of matching elements a challenging problem that is distinctly different from partial matching or co-segmentation. We detect matching elements on the analyzed objects using a combination of bottom up clustering and top down search. We then evaluate their prevalence, their saliency, and the degree of similarity between them to generate a single style similarity measurement.

When evaluating element shape similarity and salience, we are

ACM Reference Format

Lun, Z., Kalogerakis, E., Sheffer, A. 2015. Elements of Style: Learning Perceptual Shape Style Similarity. ACM Trans. Graph. 34, 4, Article 84 (August 2015), 14 pages. DOI = 10.1145/2766929
<http://doi.acm.org/10.1145/2766929>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766929>

Style Compatibility for 3D Furniture Models

Tianqiang Liu¹

Aaron Hertzmann²

Wilmot Li²

Thomas Funkhouser¹

¹Princeton University

²Adobe Research



Figure 1: This paper proposes a method to learn a metric for stylistic compatibility between furniture in a scene. (a) The image on the left shows a plausible furniture arrangement, but with a randomly chosen mix of clashing furniture styles. The scene on the right has the same arrangement but with furniture pieces chosen to optimize stylistic compatibility according to our metric.

Abstract

This paper presents a method for learning to predict the stylistic compatibility between 3D furniture models from different object classes: e.g., how well does this chair go with that table? To do this, we collect relative assessments of style compatibility using crowdsourcing. We then compute geometric features for each 3D model and learn a mapping of them into a space where Euclidean distances represent style incompatibility. Motivated by the geometric subtleties of style, we introduce part-aware geometric feature vectors that characterize the shapes of different parts of an object separately. Motivated by the need to compute style compatibility between different object classes, we introduce a method to learn object class-specific mappings from geometric features to a shared feature space. During experiments with these methods, we find that they are effective at predicting style compatibility agreed upon by people. We find in user studies that the learned compatibility metric is useful for novel interactive tools that: 1) retrieve stylistically compatible models for a query, 2) suggest a piece of furniture for an existing scene, and 3) help guide an interactive 3D modeler towards scenes with compatible furniture.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms

Keywords: style, compatibility, crowdsourcing, scene synthesis

1 Introduction

Modeling 3D scenes is one of the most common creative tasks in computer graphics. Large online repositories of 3D models make it possible for novice users and/or automatic programs to create new scenes by assembling models of objects found online. For example, the arrangement of furniture shown in Figure 1(a) was created with a tool that allows a person to select and place objects interactively.

While many existing tools help users select the appropriate categories and placements of objects when modeling a 3D scene [Merrill et al. 2011; Yu et al. 2011], they generally ignore style compatibility — the degree to which objects “exist together in harmony” [Merriam-Webster 2004]. For example, while the scene shown in Figure 1(a) has a plausible spatial arrangement of objects appropriate for a living room, it contains a mish-mash of different styles — e.g., a casual contemporary coffee table appears in front of a formal antique sofa. The jarring juxtaposition of incompatible styles diminishes the plausibility of this scene.

The goal of this paper is to develop a mathematical representation of style compatibility between objects that can be used to guide 3D scene modeling tools. More specifically, we consider the compatibility of furniture in indoor scenes, because furniture exhibits a diverse range of distinct styles, some of which are more compatible than others, and because indoor scenes account for a large fraction of scene modeling tasks. Our work focuses on understanding how the geometry of 3D models influences their stylistic compatibility. We leave the study of compatibility for other properties (materials, colors, etc.) for future work.

There are three main challenges in developing a style compatibility metric for furniture shapes. First, a person’s notion of furniture style usually combines many subtle factors [Miller 2005] that would be hard to encode in a hand-tuned mathematical function. Instead, we learn a metric from examples. Second, furniture shapes are influenced by both function and style, with functional requirements reflected largely in the gross shapes and arrangements of parts, and styles reflected largely in the geometric details of parts (e.g., fluted legs, wing-tipped backs, etc.). Accordingly, we introduce part-aware shape features aimed at capturing the geometric details related to style. Third, style compatibility requires com-

ACM Reference Format
Liu, T., Hertzmann, A., Li, W., Funkhouser, T. 2015. Style Compatibility For 3D Furniture Models. ACM Trans. Graph. 34, 4, Article 85 (August 2015), 9 pages. DOI = 10.1145/2766898
<http://doi.acm.org/10.1145/2766898>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766898>

Semantic Shape Editing Using Deformation Handles

Mehmet Ersin Yumer*

Siddhartha Chaudhuri†
*Carnegie Mellon University

Jessica K. Hodgins*
†Cornell University

Levent Burak Kara*




Figure 1: Shapes edited using our system. Please refer to our supplementary video for real-time editing examples.

Abstract

We propose a shape editing method where the user creates geometric deformations using a set of semantic attributes, thus avoiding the need for detailed geometric manipulations. In contrast to prior work, we focus on continuous deformations instead of discrete part substitutions. Our method provides a platform for quick design explorations and allows non-experts to produce semantically guided shape variations that are otherwise difficult to attain. We crowdsource a large set of pairwise comparisons between the semantic attributes and geometry and use this data to learn a continuous mapping from the semantic attributes to geometry. The resulting map enables simple and intuitive shape manipulations based solely on the learned attributes. We demonstrate our method on large datasets using two different user interaction modes and evaluate its usability with a set of user studies.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Applications;

Keywords: semantic editing, semantic deformation, shape deformation, shape sets, crowdsourcing.

1 Introduction

The ability to edit existing digital objects is central to many modeling activities including shape design, shape exploration, and product customization. Intuitive and fast editing methods enable digital artists to build upon prior work, designers to explore shape variations, and engineers to respond to new product requirements. However, conventional shape design and editing technologies (e.g., Maya, ZBrush, AutoCAD, SketchUp) are difficult to master because they require (1) expertise in the target product domain to ensure meaningful alterations and (2) familiarity with the geometric representation and operations to realize the intended modifications. As a result, transforming high-level modeling intentions into geometric directives is often challenging. Moreover, while studies in product design demonstrate that attribute-based user ratings can assist in consumer preference modeling [Orsborn et al. 2009], it remains difficult to customize existing shapes to reflect such preferences.

In this work, we propose a shape editing method where users edit 3D shapes using a set of high-level semantic attributes (Figure 1 and 2). We focus on *deforming* an input shape rather than composing a new shape through an assembly of existing shape parts [Chaudhuri et al. 2013]. As such, our method extends to databases that were not created using a part-based framework. Additionally, it allows users to explore variations of an input shape while retaining the underlying topological structure. Our approach is particularly useful in scenarios where the user’s desires can be expressed using a set of attributes relevant to the target product, but there is no immediate means of transforming such intentions into geometric operations (e.g., “*Make this shoe more fashionable*”). Our system provides continuous deformations but does not add new components or remove existing ones from the model being edited. Hence, our approach complements existing part-based geometric modeling technologies (e.g., [Chaudhuri et al. 2013]) by enabling a semantically driven interface amenable to shape exploration and customization.

ACM Reference Format

Yumer, M., Chaudhuri, S., Hodgins, J., Kara, L. 2015. Semantic Shape Editing Using Deformation Handles. ACM Trans. Graph. 34, 4, Article 86 (August 2015), 12 pages. DOI = 10.1145/2766908
<http://doi.acm.org/10.1145/2766908>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766908>

Single-View Reconstruction via Joint Analysis of Image and Shape Collections

Qixing Huang Hai Wang
Toyota Technological Institute at Chicago

Vladlen Koltun
Intel Labs

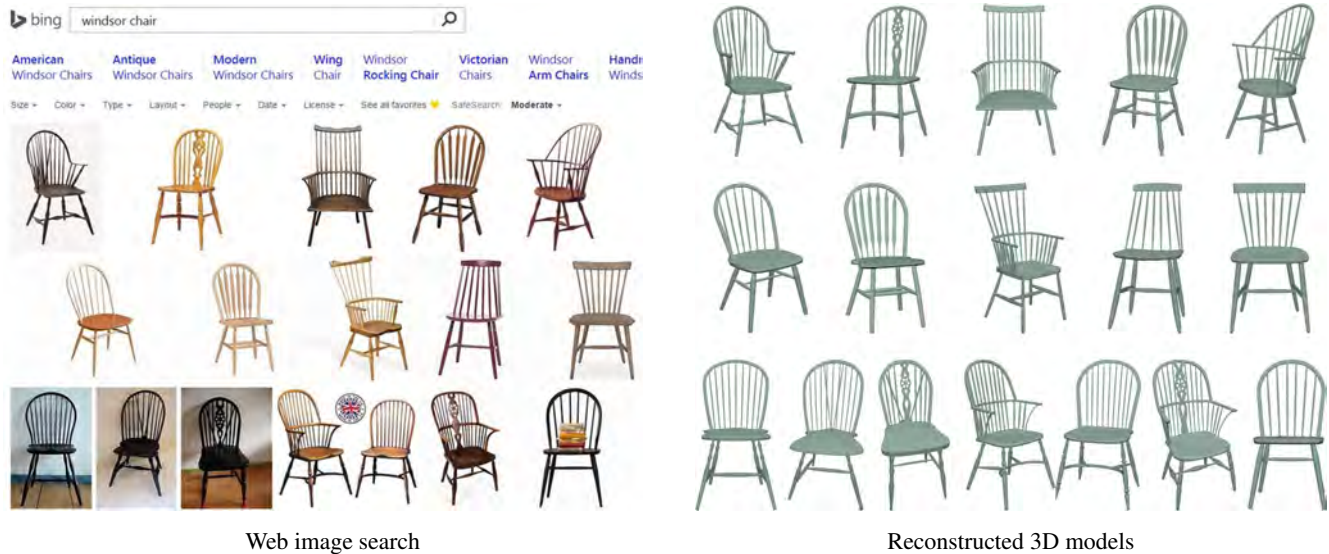


Figure 1: Our approach reconstructs objects depicted in images, even if each object is only shown in a single image. Left: Web image search for “windsor chair”, first page results. Right: 3D models automatically generated by our approach for these images.

Abstract

We present an approach to automatic 3D reconstruction of objects depicted in Web images. The approach reconstructs objects from single views. The key idea is to jointly analyze a collection of images of different objects along with a smaller collection of existing 3D models. The images are analyzed and reconstructed together. Joint analysis regularizes the formulated optimization problems, stabilizes correspondence estimation, and leads to reasonable reproduction of object appearance without traditional multi-view cues.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

Keywords: single-view reconstruction, image-based modeling

1 Introduction

Can we create 3D models of all objects in the world? High-fidelity models can be obtained from range scans and dense multi-view datasets, but acquiring such data for millions of objects demands

substantial time and labor. Can large repositories of 3D models be created using existing data that is already present on the Web?

Web images have been used to reconstruct landmark scenes, which are densely sampled by thousands of visitors [Snavely et al. 2010]. But what about the objects that populate our daily lives? Millions are already depicted on the Web. Can we exploit the regularity of object shapes [Kalogerakis et al. 2012] to reconstruct an object even if it appears in just a single image, thereby paving the way to large repositories of 3D models created by mining the Web?

In this paper, we present an approach to creating 3D models of objects depicted in images, even if each object is only shown in a single image. Our approach uses a comparatively small collection of existing 3D models to guide the reconstruction process. A key challenge is that these existing 3D models may only sparsely sample the underlying shape space. The number of high-fidelity shapes in existing 3D model repositories, such as the 3D Warehouse, is much smaller than the number of images on the Web. For many families of objects, high-quality images vastly outnumber high-quality 3D models. Thus simply retrieving the most similar existing 3D model for each input image does not yield satisfactory results: even if such retrieval is performed reliably, the closest pre-existing model is often quite different from the depicted object. We therefore develop an assembly-based approach that reconstructs objects by composing parts from pre-existing shapes.

A key idea in the presented pipeline is to jointly analyze the images and the 3D models. First, camera poses for all images are estimated by optimizing a global objective that measures the consistency of estimated poses across similar images. Then, a global network of dense pixel-level correspondences is established between natural images and rendered images. These correspondences are used to jointly segment the images and the 3D models. The computed segmentations and correspondences are used to construct new models, which are then optimized.

ACM Reference Format

Huang, Q., Wang, H., Koltun, V. 2015. Single-View Reconstruction via Joint Analysis of Image and Shape Collections. ACM Trans. Graph. 34, 4, Article 87 (August 2015), 10 pages. DOI = 10.1145/2766890 <http://doi.acm.org/10.1145/2766890>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766890>

Architecture-Scale Human-Assisted Additive Manufacturing

Hironori Yoshida^{1,*}, Takeo Igarashi¹, Yusuke Obuchi¹, Yosuke Takami¹, Jun Sato¹, Mika Araki¹, Masaaki Miki¹, Kosuke Nagata¹, Kazuhide Sakai², and Syunsuke Igarashi²

¹The University of Tokyo

²Shimizu Corporation




Figure 1: An overview of our method: (i) aggregation of printed chopsticks solidified with wood glue; (ii) a specially developed handheld printing device to consistently feed a chopstick-glue composite; (iii) print guidance system implemented by a projector-camera system; (iv) the constructed pavilion as a case study. Note that the method is still experimental, and the upper part of this pavilion was constructed separately as panels and assembled later.

Abstract

Recent digital fabrication tools have opened up accessibility to personalized rapid prototyping; however, such tools are limited to product-scale objects. The materials currently available for use in 3D printing are too fine for large-scale objects, and CNC gantry sizes limit the scope of printable objects. In this paper, we propose a new method for printing architecture-scale objects. Our proposal includes three developments: (i) a construction material consisting of chopsticks and glue, (ii) a handheld chopstick dispenser, and (iii) a printing guidance system that uses projection mapping. The proposed chopstickglue material is cost effective, environmentally sustainable, and can be printed more quickly than conventional materials. The developed handheld dispenser enables consistent feeding of the chopstickglue material composite. The printing guidance system — consisting of a depth camera and a projector — evaluates a given shape in real time and indicates where humans should deposit chopsticks by projecting a simple color code onto the form under construction. Given the mechanical specifications of the stickglue composite, an experimental pavilion was designed as a case study of the proposed method and built without scaffoldings and formworks. The case study also revealed several fundamental limitations, such as the projector does not work in daylight, which requires future investigations.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction Techniques I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling;

Keywords: 3D printing, Fabrication

1 Introduction

Digital fabrication has become more prevalent in recent years due to the increased accessibility of a range of personal digital fabrication tools. In the architectural domain, such tools have been in use for years in both professional practices and educational settings. However, unlike CAD software (which fully replaced conventional drafters), applications of these tools are limited either to custom component fabrication or scale model making. Additionally, these digital fabrication tools can never fully replace the role of humans in construction and development processes. Human labor is still essential when dealing with uncertainties in on-site construction processes because humans are able to flexibly determine case-by-case solutions. This is primarily due to the one-off nature of architectural design. Moreover, digital fabrication tools were originally designed for factory automation in controlled environments, thus not immediately applicable to construction sites [Bock 2007].

There have been some attempts to scale up 3D printer to building scale; however, a massive CNC gantry base has been an issue to be resolved and available materials are limited. The D-shape has a

ACM Reference Format

Yoshida, H., Igarashi, T., Obuchi, Y., Takami, Y., Sato, J., Araki, M., Miki, M., Nagata, K., Sakai, K., Igarashi, S. 2015. Architecture-Scale Human-Assisted Additive Manufacturing. ACM Trans. Graph. 34, 4, Article 88 (August 2015), 8 pages. DOI = 10.1145/2766951 <http://doi.acm.org/10.1145/2766951>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766951>

*e-mail: hironoriyh@gmail.com

Parametric Self-supporting Surfaces via Direct Computation of Airy Stress Functions

Masaaki Miki*
The University of Tokyo

Takeo Igarashi†
The University of Tokyo

Philippe Block‡
ETH Zurich




Figure 1: An Airy stress function is directly computed by the proposed computational approach to self-supporting surfaces.

Abstract

This paper presents a method that employs parametric surfaces as surface geometry representations at any stage of a computational process to compute self-supporting surfaces. This approach can be differentiated from existing relevant methods because such methods represent surfaces by a triangulated mesh surface or a network consisting of lines. The proposed method is based on the theory of Airy stress functions. Although some existing methods are also based on this theory, they apply its discrete version to discrete geometries. The proposed method simultaneously applies the theory to parametric surfaces directly and the discrete theory to the edges of parametric patches. The discontinuous boundary between continuous patches naturally corresponds to ribs seen in traditional vault masonry buildings. We use nonuniform rational B-spline surfaces in this study; however, the basic idea can be applied to other parametric surfaces. A variety of self-supporting surfaces obtained by the proposed computational scheme is presented.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

Keywords: NURBS surface, Airy stress function, self-supporting surface

*e-mail: masaakim@acm.org

†e-mail: takeo@acm.org

‡e-mail: pblock@ethz.ch

ACM Reference Format

Miki, M., Igarashi, T., Block, P. 2015. Parametric Self-Supporting Surfaces via Direct Computation of Airy Stress Functions. *ACM Trans. Graph.* 34, 4, Article 89 (August 2015), 12 pages. DOI = 10.1145/2766888 <http://doi.acm.org/10.1145/2766888>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766888>

1 Introduction

A self-supporting surface can support its self-weight by a purely compressive stress field. A structure of this class does not exhibit bending when design loading (typically the surface structure's self-weight plus any other dead loads) is acting on it; therefore, employing self-supporting surfaces is advantageous for construction of compressive structures, such as masonry and reinforced concrete shells. Once a self-supporting surface has been obtained, the shape will then be relayed to a structural engineering stage, typically using a series of finite element analyses, deformation, and the computation of strain and stress fields, in which both dead (design) and live loads are considered. Snow, wind, and seismic loads are typical examples of live loads.

In recent studies related to graphics, several authors have presented computational methods for self-supporting surfaces [Vouga et al. 2012; Panozzo et al. 2013; Liu et al. 2013; de Goes et al. 2013; Tang et al. 2014]. To some extent, those methods are based on Thrust Network Analysis (TNA), which was first presented by Block et al. [2007], or on discrete Airy potential polyhedral, which was first presented by Fraternali et al. [2002; 2014]. Note that Fraternali later mentioned Block's TNA [2010], implying a strong connection between TNA and discrete Airy potential polyhedral.

However, all those methods model surfaces with a triangulated mesh surface or a line network. Although applications of interesting concepts such as dual/reciprocal diagrams and a discrete Hodge star have been posed, these methods cannot handle continuous parametric surfaces without discretization. This is unfortunate because architects directly manipulate nonuniform rational B-spline (NURBS) surfaces in CAD software, such as Rhinoceros® and Blender®, and it is desirable to represent self-supporting surfaces in the same parametric representation without discretization. Another approach is to directly determine stress functions through elementary mathematical functions [Williams 1990]. However, there is only limited design freedom of stress functions in this approach.

Thus, we have developed a method that employs NURBS surfaces as surface representations at any stage of the computational process to compute parametric self-supporting surfaces. The proposed method is based on the theory of Airy stress functions, similar to some existing methods. However, while existing methods are based

Foldabilizing Furniture

Honghua Li^{1,2*} Ruizhen Hu^{1,3,4*} Ibraheem Alhashim¹ Hao Zhang¹

¹Simon Fraser University ²National University of Defense Technology ³SIAT ⁴Zhejiang University

Abstract

We introduce the *foldabilization* problem for space-saving furniture design. Namely, given a 3D object representing a piece of furniture, our goal is to apply a minimum amount of modification to the object so that it can be folded to save space — the object is thus foldabilized. We focus on one instance of the problem where folding is with respect to a prescribed folding direction and allowed object modifications include hinge insertion and part shrinking.

We develop an automatic algorithm for foldabilization by formulating and solving a nested optimization problem operating at two granularity levels of the input shape. Specifically, the input shape is first partitioned into a set of integral *folding units*. For each unit, we construct a graph which encodes conflict relations, e.g., collisions, between foldings implied by various patch foldabilizations within the unit. Finding a minimum-cost foldabilization with a conflict-free folding is an instance of the maximum-weight independent set problem. In the outer loop of the optimization, we process the folding units in an optimized ordering where the units are sorted based on estimated foldabilization costs. We show numerous foldabilization results computed at interactive speed and 3D-print physical prototypes of these results to demonstrate manufacturability.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: foldabilization, folding, furniture design, shape optimization, shape compaction

“Man, himself a collapsible being, physically and psychologically, needs and wants collapsible tools.”

– Per Mollerup (2001)

1 Introduction

Space-saving furniture designs are ubiquitous in our daily lives and workplaces. Effective space saving does not depend on down-scaling, but on smart ways of collapsing a piece of furniture or making it more collapsible. Among the many space-saving mechanisms such as stacking, implosion, and bundling [Mollerup 2001], *folding* is perhaps the most frequently observed and best practiced on furniture. Even when confined to furniture, folding can still be executed

*Honghua Li and Ruizhen Hu are joint first authors.
e-mail: howard.hhli@gmail.com, ruizhen.hu@gmail.com

ACM Reference Format
Li, H., Hu, R., Alhashim, I., Zhang, H. 2015. Foldabilizing Furniture. ACM Trans. Graph. 34, 4, Article 90 (August 2015), 12 pages. DOI = 10.1145/2766912 <http://doi.acm.org/10.1145/2766912>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766912>



Figure 1: Two automatic foldabilizations of a chair with respect to two folding directions. Shown on the right are fabrication results produced by a MakerBot Replicator II 3D printer. Folding is possible by adding hinges, shrinking parts (chair back in the top row), or allowing slanting or shearing, leading to less hinges and better structural soundness (bottom right). The foldable chair in the top row resembles the Stitch Chair by the designer Adam Goodrum.

in a rich variety of ways, offering an abundant source of appealing and challenging geometry problems.

An interesting geometry question about folding is: what makes some 3D objects more amenable to folding than others? Since rigid parts cannot be folded, hinges need to be inserted to make the parts foldable. Moreover, folding involves constrained movements of object parts and such movements often require necessary clearing space to avoid collision. Hence reducing the size or extent of furniture parts to make space is beneficial to folding, e.g., the back of the chair in Figure 1 (top) is shrunk to allow folding of the seat. Taking the two factors to the extreme, we arrive at structures formed by thin frames with many hinges; famous examples of such objects are scissoring structures such as the Hoberman spheres. However, due to structural strength and functionality considerations, furniture foldabilization can hardly go to that extreme.

In this paper, we pose the novel *foldabilization* problem for space-saving furniture design. Given a 3D object O representing a piece of furniture, our goal is to apply a *minimum* amount of modification to O to allow its parts to be folded *flatly* onto themselves or each other. We focus on one particular instance of the foldabilization problem where folding is allowed only with respect to a prescribed *folding direction* and the allowed modifications include adding (*line*) *hinges* onto furniture parts and *shrinking* the parts. Figure 1 shows two automatic foldabilization results for a chair, and the resulting folding sequences on fabricated 3D prototypes.

Foldabilization is not an easy task for humans. It resembles a 3D puzzle with a large search space and a multitude of constraints. Solving the problem requires delicate spatial reasoning and a keen foresight to adapt to the dynamic changes to the shape configuration as folding sequences proceed. While humans are highly apt at pattern recognition, they are not as skilled at precise 3D manipulation while relying solely on visual guidance. In particular, in human perception, lengths in 3D are systematically distorted due to perspective viewing [Baird 1970; Norman et al. 1996]. Thus, by

Computational Interlocking Furniture Assembly

Chi-Wing Fu*

Nanyang Technological University

Peng Song*

University of Science and Technology of China

Xiaoqi Yan Lee Wei Yang Pradeep Kumar Jayaraman
Nanyang Technological University

Daniel Cohen-Or
Tel Aviv University



Figure 1: Some snapshots showing the assembly of MULTI-FUNCTION TABLE. Our method can plan a network of joints (e.g., Figure 2) that globally interlocks the component parts in the assembly; the input component parts are just simple 3D shapes without joint geometry.

Abstract

Furniture typically consists of assemblies of elongated and planar parts that are connected together by glue, nails, hinges, screws, or other means that do not encourage disassembly and re-assembly. An alternative approach is to use an interlocking mechanism, where the component parts tightly interlock with one another. The challenge in designing such a network of interlocking joints is that local analysis is insufficient to guarantee global interlocking, and there is a huge number of joint combinations that require an enormous exploration effort to ensure global interlocking. In this paper, we present a computational solution to support the design of a network of interlocking joints that form a globally-interlocking furniture assembly. The key idea is to break the furniture complex into an overlapping set of small groups, where the parts in each group are immobilized by a local key, and adjacent groups are further locked with dependencies. The dependency among the groups saves the effort of exploring the immobilization of every subset of parts in the assembly, thus allowing the intensive interlocking computation to be localized within each small group. We demonstrate the effectiveness of our technique on many globally-interlocking furniture assemblies of various shapes and complexity.

CR Categories: I.3.5 [Computational Geometry and Object Modeling]: Curve, surface, solid, and object representations

Keywords: interlocking structure, furniture, joint, assembly

1 Introduction

Furniture generally refers to movable objects designed for supporting common human activities such as seating, storage, and office

work. Furniture is typically an assembly of elongated and planar parts that are connected together in various ways, for example, by glue, nails, hinges, and screws. However, these common connectors do not encourage furniture disassembly and re-assembly, and often harm the external appearance of the furniture and, in general, the aesthetics of the design [Postell 2012].

An alternative approach is to use an interlocking mechanism and design *interlocking furniture*, where the component parts tightly interlock with one another by a network of joints (see Figure 2). This can be achieved only by a *global* interlocking scheme rather than by an aggregate of local components [Lau et al. 2011]. Such interlocking furniture can be disassembled only through certain orders, starting from a specific *single key*, which is the only free-to-move part in the furniture assembly. This key should be taken out from the assembly in order to release the global interlocking.

Interlocking furniture has several advantages. First, the furniture can be easily assembled and disassembled repeatedly without excessive wear on its parts. Second, external fixtures such as fasteners (e.g., screws) and bearings (e.g., hinges) are not required, thus keeping the intended aesthetics of the design. In addition, since the global interlocking scheme limits the parts removal and restricts the assembly and disassembly order, component parts can tightly interlock with one another, thus facilitating the stability and strength of the furniture [Laurajax 2014; Jones 2014].

Given a 3D furniture design, the goal of this work is

To plan and construct a network of joints in the design of interlocking furniture, so that the component parts in the furniture assembly can tightly interlock with one another in a global interlocking manner.

The input of our computational method is a furniture design consisting of just a set of simple 3D parts (e.g., rectangular boxes), whereas the output is the modified 3D parts (see Figure 1) with appropriate joint geometry that interlocks the component parts.



Figure 2: Examples of four common joints: English dovetail, French dovetail, halved joints, and mortise and tenon (left to right).

*joint first authors

ACM Reference Format

Fu, C., Song, P., Yan, X., Yang, L., Jayaraman, P., Cohen-Or, D. 2015. Computational Interlocking Furniture Assembly. ACM Trans. Graph. 34, 4, Article 91 (August 2015), 11 pages. DOI = 10.1145/2766892
<http://doi.acm.org/10.1145/2766892>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766892>

LazyFluids: Appearance Transfer for Fluid Animations

Ondřej Jamriška^{1*} Jakub Fišer¹ Paul Asente² Jingwan Lu² Eli Shechtman² Daniel Sýkora¹
¹CTU in Prague, FEE ²Adobe Research

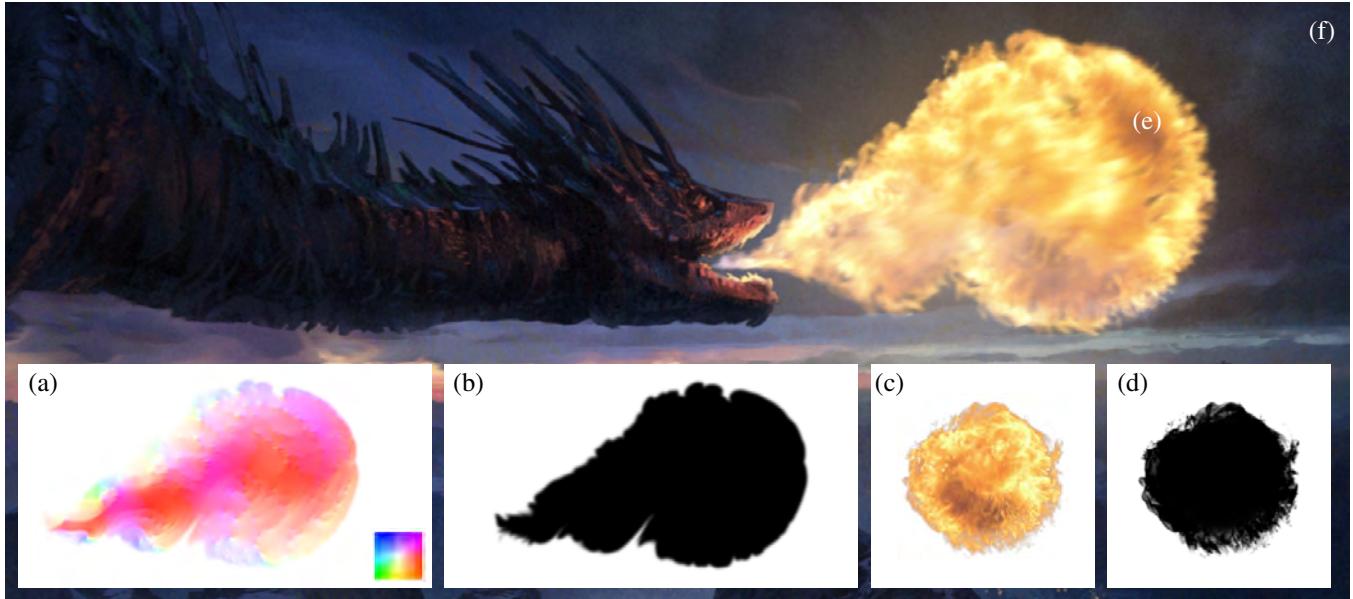


Figure 1: LazyFluids in action—an artist first designs a target fluid animation that consists of a sequence of motion fields (a) and alpha masks (b), and then selects a video exemplar of a fluid element with desired appearance (c) and alpha mask (d). Finally, our algorithm transfers appearance of the exemplar to the target animation while respecting its motion properties and boundary effects (e). The resulting animation can then be used as a part of a more complex composition (f). All alpha masks in the paper are visualised in a way that fully opaque pixels are black and fully transparent are white. Dragon painting © Jakub Javora.

Abstract

In this paper we present a novel approach to appearance transfer for fluid animations based on flow-guided texture synthesis. In contrast to common practice where pre-captured sets of fluid elements are combined in order to achieve desired motion and look, we bring the possibility of fine-tuning motion properties in advance using CG techniques, and then transferring the desired look from a selected appearance exemplar. We demonstrate that such a practical workflow cannot be simply implemented using current state-of-the-art techniques, analyze what the main obstacles are, and propose a solution to resolve them. In addition, we extend the algorithm to allow for synthesis with rich boundary effects and video exemplars. Finally, we present numerous results that demonstrate the versatility of the proposed approach.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image

*e-mail:jamriond@fel.cvut.cz

ACM Reference Format

Jamriška, O., Fišer, J., Asente, P., Lu, J., Shechtman, E., Sýkora, D. 2015. LazyFluids: Appearance Transfer for Fluid Animations. ACM Trans. Graph. 34, 4, Article 92 (August 2015), 12 pages. DOI = 10.1145/2766983 <http://doi.acm.org/10.1145/2766983>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766983>

Generation—Bitmap and framebuffer operations I.3.4 [Computer Graphics]: Graphics Utilities—Graphics editors I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

Keywords: texture synthesis, flow-guided, example-based, temporal coherency, fluid simulation

1 Introduction

Special effects based on fluid elements are ubiquitous in current digital movie production. To achieve a desired look, an artist typically makes a composition out of pre-captured videos of real fluids with a desired appearance. A key limitation here is that the motion properties of these videos remain fixed. When finer control is needed the artist has to resort to full fluid simulation followed by an advanced rendering algorithm. In this scenario, however, limited resolution and the complexity of material properties, lighting, or other parameters may hinder delivering the desired visual look.

We would like to offer artists a more practical workflow that can narrow the gap between appearance and motion controllability:

1. Quickly design the target animation using 2D CG techniques (e.g., a real-time fluid simulator [Stam 1999] or particle system [Reeves 1983]; see Fig. 1a,b).
2. Pick a photo or a video sequence containing the desired look (the source exemplar; see Fig. 1c).

Fluid Volume Modeling from Sparse Multi-view Images by Appearance Transfer

Makoto Okabe^{1,4}

Yoshinori Dobashi^{2,4}

Ken Anjyo^{3,4}

Rikio Onai¹

¹The University of Electro-Communications

²Hokkaido University

³OLM Digital, Inc.

⁴JST, CREST

Abstract

We propose a method of three-dimensional (3D) modeling of volumetric fluid phenomena from sparse multi-view images (e.g., only a single-view input or a pair of front- and side-view inputs). The volume determined from such sparse inputs using previous methods appears blurry and unnatural with novel views; however, our method preserves the appearance of novel viewing angles by transferring the appearance information from input images to novel viewing angles. For appearance information, we use histograms of image intensities and steerable coefficients. We formulate the volume modeling as an energy minimization problem with statistical hard constraints, which is solved using an expectation maximization (EM)-like iterative algorithm. Our algorithm begins with a rough estimate of the initial volume modeled from the input images, followed by an iterative process whereby we first render the images of the current volume with novel viewing angles. Then, we modify the rendered images by transferring the appearance information from the input images, and we thereafter model the improved volume based on the modified images. We iterate these operations until the volume converges. We demonstrate our method successfully provides natural-looking volume sequences of fluids (i.e., fire, smoke, explosions, and a water splash) from sparse multi-view videos. To create production-ready fluid animations, we further propose a method of rendering and editing fluids using a commercially available fluid simulator.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

Keywords: image-based modeling, natural phenomena animation, volume modeling, texture analysis/synthesis, single-view modeling

1 Introduction

Visual effects of fluid phenomena, such as fire, smoke, and explosion, are indispensable in modern movie production, and fluids have been important research topics in computer graphics for decades. However, the creation of fluid animations remains a difficult and time-consuming task, which often becomes a bottleneck in the production workflow.

There are three important methods that are commonly used to create fluid animations. The first is a physics-based fluid simulation, which allows users to design a wide variety of realistic fluids [Bridson and Müller-Fischer 2007]; however, this approach is often difficult to apply, even for experts, because the inputs to a fluid simulator (i.e., numerical parameters) are very different from the output



Figure 1: We model the natural-looking volume of fluids from sparse multi-view images (e.g., only a single-view (left) or a pair of front and side views (right)). We create the production-ready fluid animation using a volume sequence of fluids modeled from sparse multi-view videos. A fluid simulator allows users to further edit the appearance, behaviour, and shape of fluids (e.g., adding more turbulence (left)).

(appearance, behavior, and shape of fluids). The second method is to make a composite of fluid videos [Bhat et al. 2004]. Production companies typically have databases of fluid videos, whereby the artist chooses an adequate video from it, and edits and overlays it onto the scene in the post-production process. However, because fluid videos are two-dimensional (2D), it is difficult to cope with 3D effects; camera paths or lighting cannot be edited, and stereoscopic effects are not possible. The third method is image-based modeling, which takes multi-view videos of fluid phenomena and models the volume sequence [Hasinoff and Kutulakos 2007; Ihrke and Magnor 2004; Gregson et al. 2012]. Here, the input data are videos (i.e., sequences of images), which clearly correspond to the output (i.e., a volume sequence). The resulting volume sequence appears realistic, and is suitable for 3D effects; however, problems result in that there are difficulties preparing multi-view inputs, which require a studio setup with a large number of cameras.

In this paper, we propose a method of image-based modeling of fluids, which takes sparse multi-view images (e.g., only a single-view or a pair of front and side views). Figures 1 and 2 illustrate typical results obtained by our method. As will be demonstrated in this paper, while significantly reducing the number of required cameras, our method still produces high-quality and editable 3D fluids. Our method can thus provide high-quality, efficiency and flexible usability at the same time, for modeling 3D fluids. This is the main feature of our method, which has never been achieved by previous work.

When reducing the number of inputs, the visual information may become insufficient to model the volume. The top row of Figure 2 shows the volume modeled from a pair of front- and side-view inputs using the least squares method (LSM) [Kak and Slaney 1988]. This appears blurry and unnatural with novel viewing angles. Here, we ask why these images appear blurry and unnatural. To answer this question, we assume the following: *We recall the appearance of the inputs (see the leftmost and rightmost panels of Figure 2), and expect that the volume should have a similar appearance, even with novel viewing angles.*

Our method is intended to live up to this expectation. Our technical

ACM Reference Format

Okabe, M., Dobashi, Y., Anjyo, K., Onai, R. 2015. Fluid Volume Modeling from Sparse Multi-view Images by Appearance Transfer. *ACM Trans. Graph.* 34, 4, Article 93 (August 2015), 10 pages. DOI = 10.1145/2766958 <http://doi.acm.org/10.1145/2766958>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766958>

Deformation Capture and Modeling of Soft Objects

Bin Wang^{*†} Longhua Wu^{*} KangKang Yin[†] Uri Ascher[‡] Libin Liu[‡] Hui Huang^{*§}

^{*}Shenzhen VisuCA Key Lab / SIAT [†]National University of Singapore [‡]University of British Columbia



Figure 1: Our system can capture and model deformation behavior of generic soft objects from kinematic data alone. We can then synthesize new motions that satisfy user-specified constraints and respond to dynamic perturbations. Top: left - a dinosaur walking; middle - a pot holder jumping; right - a coat hanger skipping. Bottom: lotus leaves moving in an artificial wind field.

Abstract

We present a data-driven method for deformation capture and modeling of general soft objects. We adopt an iterative framework that consists of one component for physics-based deformation tracking and another for spacetime optimization of deformation parameters. Low cost depth sensors are used for the deformation capture, and we do not require any force-displacement measurements, thus making the data capture a cheap and convenient process. We augment a state-of-the-art probabilistic tracking method to robustly handle noise, occlusions, fast movements and large deformations. The spacetime optimization aims to match the simulated trajectories with the tracked ones. The optimized deformation model is then used to boost the accuracy of the tracking results, which can in turn improve the deformation parameter estimation itself in later iterations. Numerical experiments demonstrate that the tracking and parameter optimization components complement each other nicely.

Our spacetime optimization of the deformation model includes not only the material elasticity parameters and dynamic damping coefficients, but also the reference shape which can differ significantly from the static shape for soft objects. The resulting optimization problem is highly nonlinear in high dimensions, and challenging to solve with previous methods. We propose a novel splitting algorithm that alternates between reference shape optimization and deformation parameter estimation, and thus enables tailoring the optimization of each subproblem more efficiently and robustly.

^{*}email: {wangbin, wulh}@siat.ac.cn

[†]email: kkyin@comp.nus.edu.sg

[‡]email: {ascher, libinliu}@cs.ubc.ca

[§]Corresponding author: Hui Huang (hhzhiyan@gmail.com)

ACM Reference Format

Wang, B., Wu, L., Yin, K., Ascher, U., Liu, L., Huang, H. 2015. Deformation Capture and Modeling of Soft Objects. ACM Trans. Graph. 34, 4, Article 94 (August 2015), 12 pages. DOI = 10.1145/2766911 <http://doi.acm.org/10.1145/2766911>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766911>

Our system enables realistic motion reconstruction as well as synthesis of virtual soft objects in response to user stimulation. Validation experiments show that our method not only is accurate, but also compares favorably to existing techniques. We also showcase the ability of our system with high quality animations generated from optimized deformation parameters for a variety of soft objects, such as live plants and fabricated models.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: motion capture, deformation modeling, FEM

1 Introduction

Physics-based deformable models enable realistic animation of a wide range of objects and phenomena [Nealen et al. 2006]. Estimating model parameters, however, still heavily relies on either manual tuning or tedious measurements [Terzopoulos et al. 1987]. Such approaches can hardly scale to complex models with nonlinear or heterogeneous material distributions. Moreover, these methods usually employ static shapes, i.e., the static equilibrium under gravity that can be easily observed, as the original reference shapes of the deformation models. This approximation does not work for very soft objects, for instance long plant leaves, as they deform significantly due to gravity. In addition, dynamic properties such as damping coefficients have seldom been considered previously, even though they play a critical role in achieving realistic behavior, in particular for soft objects.

Data-driven methods have recently been quite successful in constructing physics-based deformable models for cloth, human organs and faces [Otaduy et al. 2012]. However, they often require measuring the dense force-displacement relationships. Such measuring processes and hardware have to be tailored to specific types of objects being modeled, and thus are hard to generalize. We wish to build a system that can learn from real-world measurements as well, and is applicable to generic objects without requiring any expensive or specialized hardware for force actuation or measurement.

In this paper, we propose a novel data-driven deformation capture and modeling framework for generic soft objects. Our system builds deformable models from pure kinematic motion trajectories,

Zoomorphic Design

Noah Duncan¹ * Lap-Fai Yu² Sai-Kit Yeung³ Demetri Terzopoulos¹

¹University of California Los Angeles

²University of Massachusetts Boston

³Singapore University of Technology and Design



(a) Input



(b) Result

Figure 1: A zoomorphic playground created by our approach.

Abstract

Zoomorphic shapes are man-made shapes that possess the form or appearance of an animal. They have desirable aesthetic properties, but are difficult to create using conventional modeling tools. We present a method for creating zoomorphic shapes by merging a man-made shape and an animal shape. To identify a pair of shapes that are suitable for merging, we use an efficient graph kernel based technique. We formulate the merging process as a continuous optimization problem where the two shapes are deformed jointly to minimize an energy function combining several design factors. The modeler can adjust the weighting between these factors to attain high-level control over the final shape produced. A novel technique ensures that the zoomorphic shape does not violate the design restrictions of the man-made shape. We demonstrate the versatility and effectiveness of our approach by generating a wide variety of zoomorphic shapes.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric Algorithms

Keywords: zoomorphic design, shape creation

* Part of the work was done when Noah was visiting SUTD.

ACM Reference Format

Duncan, N., Yu, L., Yeung, S., Terzopoulos, D. 2015. Zoomorphic Design. ACM Trans. Graph. 34, 4, Article 95 (August 2015), 13 pages. DOI = 10.1145/2766902 <http://doi.acm.org/10.1145/2766902>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766902>




Figure 2: Historic and modern zoomorphic shapes. (a) Bull-shaped vessel circa 1000 BC symbolizing fertility. (b) Piggy bank. (c) Anteater chair by artist Maximo Riera. (d) The Milwaukee Art Museum, designed in the shape of a bird in flight.

“In all things of nature there is something of the marvelous.”

— Aristotle

1 Introduction

For centuries, humanity has attempted to capture the marvels of nature in man-made objects. Such objects range from ancient pottery vessels, to modern day piggy banks, designer chairs, and even buildings (Fig. 2). Man-made shapes that have the form or appearance of an animal are called *zoomorphic*. Since the beginning of recorded history, artists have created zoomorphic shapes by “applying animalistic-inspired qualities to non-animal related objects” [Coates et al. 2009].

Zoomorphic concepts are present in architecture [Aldersey-Williams 2003], furniture [Coates et al. 2009], and product design [Bramston 2008; Lidwell and Manacsa 2011]. Research suggests that children have a natural affinity for animals, which may explain the frequent presence of zoomorphism in children’s toys [Lidwell 2014]. Fig. 1 illustrates how zoomorphic design can create a more appealing children’s playground.

We propose a novel computational approach to tackle the unique challenges involved in creating zoomorphic shapes. Some zoomor-

Shading-based Refinement on Volumetric Signed Distance Functions

Michael Zollhöfer^{1,4}
Marc Stamminger¹

Angela Dai²
Christian Theobalt⁴

Matthias Innmann¹
Matthias Nießner²

Chenglei Wu³
Matthias Nießner²

¹University of Erlangen-Nuremberg

²Stanford University

³ETH Zurich

⁴MPI Informatics




Figure 1: Our method obtains fine-scale detail through volumetric shading-based refinement (VSBR) of a distance field. We scan an object using a commodity sensor – here, a PrimeSense – to generate an implicit representation. Unfortunately, this leads to over-smoothing. Exploiting the shading cues from the RGB data allows us to obtain reconstructions at previously unseen resolutions within only a few seconds.

Abstract

We present a novel method to obtain fine-scale detail in 3D reconstructions generated with low-budget RGB-D cameras or other commodity scanning devices. As the depth data of these sensors is noisy, truncated signed distance fields are typically used to regularize out the noise, which unfortunately leads to over-smoothed results. In our approach, we leverage RGB data to refine these reconstructions through shading cues, as color input is typically of much higher resolution than the depth data. As a result, we obtain reconstructions with high geometric detail, far beyond the depth resolution of the camera itself. Our core contribution is shading-based refinement directly on the implicit surface representation, which is generated from globally-aligned RGB-D images. We formulate the inverse shading problem on the volumetric distance field, and present a novel objective function which jointly optimizes for fine-scale surface geometry and spatially-varying surface reflectance. In order to enable the efficient reconstruction of sub-millimeter detail, we store and process our surface using a sparse voxel hashing scheme which we augment by introducing a grid hierarchy. A tailored GPU-based Gauss-Newton solver enables us to refine large shape models to previously unseen resolution within only a few seconds.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Scanning;

Keywords: shading-based refinement, 3D reconstruction

1 Introduction

The advent of low-cost RGB-D cameras, such as the Microsoft Kinect, triggered the development of new algorithms that allow consumers to densely scan objects at real-time frame rates. A prominent example of this line of research is Kinect Fusion [Newcombe et al. 2011; Izadi et al. 2011] and its extensions [Roth and Vona 2012; Whelan et al. 2012; Chen et al. 2013; Nießner et al. 2013], which led to significant impact in the computer graphics, vision, and HCI communities. Their efficient reconstruction of 3D environments is beneficial to a large variety of applications, ranging from content creation to augmented reality scenarios.

Core to these approaches is the underlying surface representation of a truncated signed distance field (TSDF) [Curless and Levoy 1996]. This representation stores the distance values to the closest surface point in 3D in a voxel grid, and has several advantages compared to alternative models, such as point-based or mesh-based representations, in particular if efficiency is the goal. It enables efficient alignment and fusion of scans, while systematically considering the drastic noise and distortions in depth data of many RGB-D cameras. Further, it allows scan integration without complex connectivity and topology handling, and is the basis of many surface extraction algorithms. Unfortunately, despite these benefits, aligning and integrating depth scans on a TSDF leads to strong over-smoothing of the surfaces, as the depth data is fused projectively using a weighted average from different viewing directions. This is further compounded by drift due to alignment errors of input depth frames. Thus, while TSDFs efficiently regularize noise, resulting reconstructions lack fine-scale geometric detail. Additionally, scanning quality is limited by the limits of the depth cameras themselves. Most of them deliver images of much lower depth resolution than RGB resolution; depth is also very noisy, and may contain systematic distortions. Overall, the benefits of implicit surface representations have brought them to be prevalent in online and offline 3D reconstruction approaches, but they fail to capture fine-scale detail for noisy depth input. Note that this extends beyond TSDFs, as other implicit functions have been used as representations for 3D reconstruction; e.g., [Carr et al. 2001; Kazhdan et al. 2006; Fuhrmann and Goesele 2014].

In this paper, we address the problem of over-smoothed TSDFs, and propose a new approach to efficiently reconstruct fine-scale detail on them. To this end, we leverage RGB data to refine the implicit

ACM Reference Format

Zollhöfer, M., Dai, A., Innmann, M., Wu, C., Stamminger, M., Theobalt, C., Nießner, M. 2015. Shading-based Refinement on Volumetric Signed Distance Functions. *ACM Trans. Graph.* 34, 4, Article 96 (August 2015), 14 pages. DOI = 10.1145/2766887 <http://doi.acm.org/10.1145/2766887>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766887>

PatchTable: Efficient Patch Queries for Large Datasets and Applications

Connelly Barnes¹

¹University of Virginia

Fang-Lue Zhang²

²TNList, Tsinghua University

Liming Lou^{1,3}

Xian Wu²

³Shandong University

Shi-Min Hu²




Figure 1: Our artistic video stylization, which generalizes image analogies [Hertzmann et al. 2001] to video. An input video (a) is stylized by example in several different styles (b-d). We show a detailed crop region of a single frame of the full video, which is shown in the supplemental video. In the second row we show the exemplar pair that is used to drive the stylization, which consists of an image “before” and “after” the effect is applied. Our data structure enables this effect to be rendered on the CPU at 1024x576 resolution at 1 frame/second, which is significantly faster than previous work. Credits: Video in top row © Renars Vilnis; (b) Thomas Nast; (c) Vincent Van Gogh; (d) © Hashimoto et al. [2003].

Abstract

This paper presents a data structure that reduces approximate nearest neighbor query times for image patches in large datasets. Previous work in texture synthesis has demonstrated real-time synthesis from small exemplar textures. However, high performance has proved elusive for modern patch-based optimization techniques which frequently use many exemplar images in the tens of megapixels or above. Our new algorithm, PatchTable, offloads as much of the computation as possible to a pre-computation stage that takes modest time, so patch queries can be as efficient as possible. There are three key insights behind our algorithm: (1) a lookup table similar to locality sensitive hashing can be precomputed, and used to seed sufficiently good initial patch correspondences during querying, (2) missing entries in the table can be filled during pre-computation with our fast Voronoi transform, and (3) the initially seeded correspondences can be improved with a precomputed k-nearest neighbors mapping. We show experimentally that this accelerates the patch query operation by up to $9\times$ over k-coherence, up to $12\times$ over TreeCANN, and up to $200\times$ over PatchMatch. Our fast algorithm allows us to explore efficient and practical imaging and computational photography applications. We show results for artistic video stylization, light field super-resolution, and multi-image editing.

CR Categories: I.3.6 [Computing Methodologies]: Computer Graphics—Methodology and Techniques; I.4.9 [Computing Methodologies]: Image Processing and Computer Vision—Applications;

Keywords: Approximate nearest neighbor, patch-based synthesis

ACM Reference Format

Barnes, C., Zhang, F., Lou, L., Wu, X., Hu, S. 2015. PatchTable: Efficient Patch Queries for Large Datasets and Applications. ACM Trans. Graph. 34, 4, Article 97 (August 2015), 10 pages. DOI = 10.1145/2766934 <http://doi.acm.org/10.1145/2766934>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766934>

1 Introduction

Digital photography has recently shown clear trends towards higher resolutions and large collections of photos. Users frequently have personal collections of thousands of pictures, and billions of photographs are available from online photo-sharing websites. Resolutions range from tens of megapixels on consumer cameras up to gigapixels on specialized capture devices. There is a need for algorithms that scale to large image datasets.

Modern patch-based methods have delivered high quality results for many applications, by synthesizing or analyzing images in terms of small regions called patches (or neighborhoods). However, these methods have had a difficult time scaling to large quantities of data. For instance, the PatchMatch [Barnes et al. 2009] method is inefficient for large dataset sizes. This is because finding a correspondence in the worst case reduces to random sampling, which takes time linear in the database size. In contrast, earlier methods for texture synthesis could operate in real time [Lefebvre and Hoppe 2005], due to using low resolution inputs and the use of pre-computation in the design of their data structures.

Our aim is to accelerate modern patch-based methods so they efficiently scale to high-resolution photographs and collections of photos. In particular, we aim to accelerate the approximate nearest neighbor search problem over image patches. This problem has attracted much research interest recently [Barnes et al. 2009; Korman and Avidan 2011; Olonetsky and Avidan 2012]. To accelerate this search problem, we use an approach similar to Web search engines. We assume that modest time resources can be devoted to precomputing an *inverted index* data structure that consequently will permit very fast queries. As we demonstrate, for many applications, this index can be computed once when a user’s photographs are first loaded, and re-used for many queries. This improves interactivity and efficiency.

The main ideas behind our method are to push as much as possible of the work to a precomputation stage when the index is built, and to use array data structures to represent patch appearance space. This allows us to avoid the complex traversals that hierarchical data structures like kd-trees must make during patch queries. After the index is constructed, the time to query the patches of an image against the index is quite fast. We show improvements in query

Learning visual similarity for product design with convolutional neural networks

Sean Bell Kavita Bala
Cornell University*

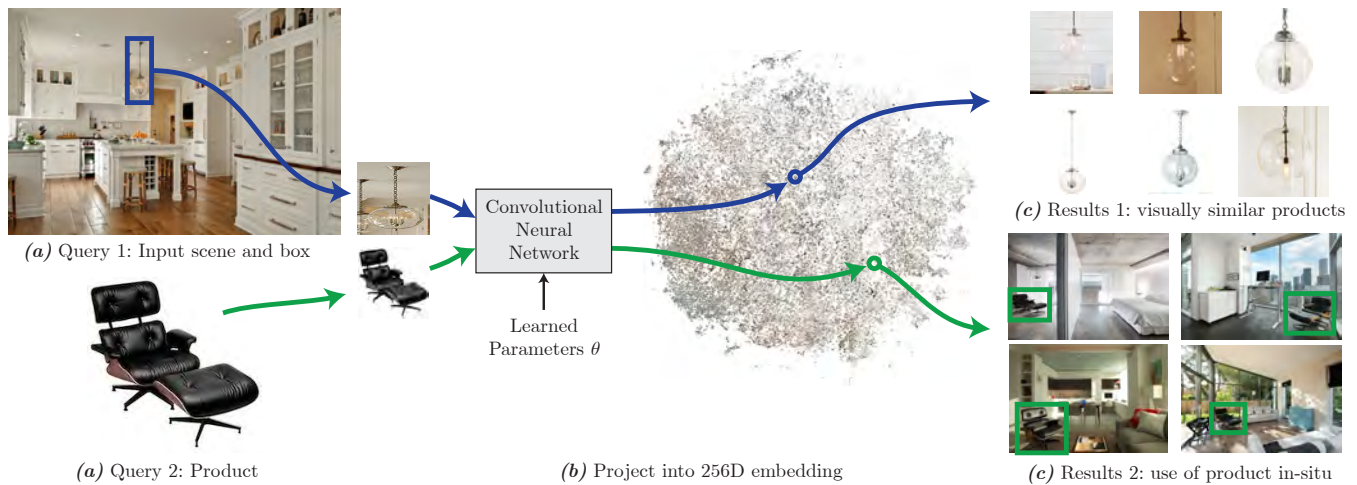


Figure 1: Visual search using a learned embedding. *Query 1:* given an input box in a photo (a), we crop and project into an embedding (b) using a trained convolutional neural network (CNN) and return the most visually similar products (c). *Query 2:* we apply the same method to search for in-situ examples of a product in designer photographs. The CNN is trained from pairs of internet images, and the boxes are collected using crowdsourcing. The 256D embedding is visualized in 2D with t-SNE. Photo credit: Crisp Architects and Rob Karosis (photographer).

Abstract

Popular sites like Houzz, Pinterest, and LikeThatDecor, have communities of users helping each other answer questions about products in images. In this paper we learn an embedding for visual search in interior design. Our embedding contains two different domains of product images: products cropped from internet scenes, and products in their *iconic* form. With such a multi-domain embedding, we demonstrate several applications of visual search including identifying products in scenes and finding stylistically similar products. To obtain the embedding, we train a convolutional neural network on pairs of images. We explore several training architectures including re-purposing object classifiers, using siamese networks, and using multitask learning. We evaluate our search quantitatively and qualitatively and demonstrate high quality results for search across multiple visual domains, enabling new applications in interior design.

CR Categories: I.3.8 [Computer Graphics]: Applications I.4.8 [Image Processing and Computer Vision]

Keywords: visual similarity, interior design, deep learning, search

*Authors' email addresses: {sbell, kb}@cs.cornell.edu

ACM Reference Format

Bell, S., Bala, K. 2015. Learning Visual Similarity for Product Design with Convolutional Neural Networks. ACM Trans. Graph. 34, 4, Article 98 (August 2015), 10 pages. DOI = 10.1145/2766959
<http://doi.acm.org/10.1145/2766959>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.
Copyright held by the Owner/Author.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
ACM 978-1-4503-3331-3/15/08.
DOI: <http://dx.doi.org/10.1145/2766959>

1 Introduction

Home owners and consumers are interested in visualizing ideas for home improvement and interior design. Popular sites like Houzz, Pinterest, and LikeThatDecor have large active communities of users that browse the sites for inspiration, design ideas and recommendations, and to pose design questions. For example, some topics and questions that come up are:

- “What is this {chair, lamp, wallpaper} in this photograph? Where can I find it?”, or, “Find me {chairs, ...} similar to this one.” This kind of query may come from a user who sees something they like in an online image on Flickr or Houzz, a magazine, or a friend’s home.
- “How has this armchair been used in designer photos?” Users can search for the usage of a product for design inspiration.
- “Find me a compatible chair matching this table.” For example, a home owner is replacing furniture in their home and wants to find a chair that matches their existing table (and bookcase).

Currently, sites like Houzz have active communities of users that answer design questions like these with (sometimes informed) guesses. Providing automated tools for design suggestions and ideas can be very useful to these users.

The common thread between these questions is the need to find *visually similar* objects in photographs. In this paper we learn a distance metric between an object *in-situ* (i.e., a sub-image of a photograph) and an *iconic* product image of that object (i.e., a clean well-lit photograph, usually with a white background). The distance is small between the in-situ object image and the iconic product image, and large otherwise. Learning such a distance metric is challenging because the in-situ object image can have many different backgrounds, sizes, orientations, or lighting when compared to the iconic product image, and, it could be significantly occluded by clutter in the scene.

LinkEdit: Interactive Linkage Editing using Symbolic Kinematics

Moritz Bächer¹

Stelian Coros^{1,2}

Bernhard Thomaszewski¹

¹Disney Research Zurich

²Carnegie Mellon University



Figure 1: Our interactive editing system allows users to adapt the shape and motion of planar linkages in an intuitive manner. As shown on the Strandbeest sculpture (1st from left), we support diverse edits including changes to the enclosure of the assembly (2nd) or the shape of individual components (3rd). Our editing tools preserve correct functioning at all times, as illustrated on a 3D-printed prototype (4th).

Abstract

We present a method for interactive editing of planar linkages. Given a working linkage as input, the user can make targeted edits to the shape or motion of selected parts while preserving other, e.g., functionally-important aspects. In order to make this process intuitive and efficient, we provide a number of editing tools at different levels of abstraction. For instance, the user can directly change the structure of a linkage by displacing joints, edit the motion of selected points on the linkage, or impose limits on the size of its enclosure. Our method safeguards against degenerate configurations during these edits, thus ensuring the correct functioning of the mechanism at all times. Linkage editing poses strict requirements on performance that standard approaches fail to provide. In order to enable interactive and robust editing, we build on a symbolic kinematics approach that uses closed-form expressions instead of numerical methods to compute the motion of a linkage and its derivatives. We demonstrate our system on a diverse set of examples, illustrating the potential to adapt and personalize the structure and motion of existing linkages. To validate the feasibility of our edited designs, we fabricated two physical prototypes.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: linkages, animation, fabrication, interactive editing

1 Introduction

As 3D printers become widely available, more and more printable content is published online. Ranging from artwork to replacement parts and mechanically functional models, there are numerous catalogs with myriad printable models to download and replicate at

home. While there is interest and value in replication, however, this usage forgoes the true potential of 3D printers: *personalization*—the opportunity to adapt content to individual needs and preferences.

A variety of existing software tools allow for intuitive editing of static digital models. Some of these tools even maintain or establish structural stability by virtue of finite element analysis. To date, however, no equivalent CAD software has been developed for mechanically-functional models. Yet, a vast array of mechanically-functional objects such as kinetic sculptures, fold-away furniture, electro-mechanical toys, and even robots stand to benefit from user-friendly editing tools and the corresponding potential for personalization that they would bring about.

We focus on the challenge of editing planar linkages, which are the beating heart of many mechanically-functional models—and notoriously difficult to design and edit: even for simple linkages, the relation between joint displacements and resulting change in motion-curves is complex and very difficult to predict. Within the group of planar linkages, our method can handle a large set of well-known and widely-used mechanisms.

Editing a linkage means changing certain aspects of its shape or motion while others are preserved. The reasons for editing linkages are diverse: for example, one may want to personalize an existing design by adapting the aesthetic appeal of its shape while maintaining its function, or one may want to alter the motion in order to reuse the linkage in a different context. In order to support such customization, we propose a set of interactive tools and high-level editing abstractions. Our method complements the recent body of work aimed at creating mechanically-functional objects [Zhu et al. 2012; Coros et al. 2013; Ceylan et al. 2013; Thomaszewski et al. 2014] and is designed to leverage the wealth of mechanical designs that can be found in online repositories such as *GrabCAD* and *Thingiverse*.

Overview & Contributions In this work, we present an interactive tool for intuitive editing of existing planar linkages, i.e., assemblies of interconnected rigid components that perform planar motion. Although mechanisms can involve other components such as gear-trains and cams, planar linkages are arguably the most challenging to design and edit. In order to accommodate the work flow that we propose, the design tool must *a)* be fast enough to provide interactive rates, *b)* offer abstractions and high-level metaphors for intuitive editing, and *c)* preserve correct functioning of the mecha-

ACM Reference Format

Bächer, M., Coros, S., Thomaszewski, B. 2015. LinkEdit: Interactive Linkage Editing using Symbolic Kinematics. ACM Trans. Graph. 34, 4, Article 99 (August 2015), 8 pages. DOI = 10.1145/2766985 <http://doi.acm.org/10.1145/2766985>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766985>

Fab Forms: Customizable Objects for Fabrication with Validity and Geometry Caching

Maria Shugrina¹

Ariel Shamir²

Wojciech Matusik¹

¹MIT CSAIL

²The Interdisciplinary Center Herzliya



Figure 1: Using offline adaptive sampling, our method converts general parametric designs into *Fab Forms*, parameterized object representations supporting interactive customization, while ensuring high-level object validity. All realized designs are functional and fabricable and can be previewed in real time. We automatically skin these representations with a Web customization UI intended for casual users.

Abstract

We address the problem of allowing casual users to customize parametric models while maintaining their valid state as 3D-printable functional objects. We define *Fab Form* as any design representation that lends itself to *interactive* customization by a novice user, while remaining *valid* and *manufacturable*. We propose a method to achieve these *Fab Form* requirements for general parametric designs tagged with a general set of automated validity tests and a small number of parameters exposed to the casual user. Our solution separates *Fab Form* evaluation into a precomputation stage and a runtime stage. Parts of the geometry and design validity (such as manufacturability) are evaluated and stored in the precomputation stage by adaptively sampling the design space. At runtime the remainder of the evaluation is performed. This allows interactive navigation in the valid regions of the design space using an automatically generated Web user interface (UI). We evaluate our approach by converting several parametric models into corresponding *Fab Forms*.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Hierarchy and geometric transformations I.3.8 [Computer Graphics]: Applications—;

Keywords: fabrication, customizable design, precomputation, CAD, geometry caching, valid region estimation

1 Introduction

Custom products offer inherent advantages over their mass-produced counterparts. Personalized objects can provide more comfort, unique aesthetic appeal, or better performance (e.g. shoes, jewelry, prosthetic limbs). Additive manufacturing promises cost-effective fabrication of such personalized products, as there is a steady progress in the 3D printing hardware and the range of available materials, as well as a reduction in cost and improved availability to consumers (e.g. through 3D printing services).

Yet, taking advantage of additive manufacturing for personalization requires developing customizable digital models and fabrication applications that can be used reliably by everyone. This task proves to be quite challenging, and fabricable end-user-customizable objects available today are very limited. Recent attempts such as one-off apps for customization of household objects [Shapeways; dreamforge], toys [Makeworld Ltd.] and jewelry [Nervous System] either restrict customization to very simple changes, or do not provide a guarantee that the customized objects are valid (e.g. stable, durable) and can truly be fabricated. Furthermore, building a dedicated software application for every customizable design is not scalable. A more general approach of putting a thin UI layer over a parametric model, as in [MakerBot], results in a frustrating user experience with no guarantee on the validity of the resulting model.

In this work, we define the requirements for a representation of an end-user-customizable object in the notion of a *Fab Form*, and provide a method for creating such a representation from a general parametric design. Formally, a *Fab Form* is a 3D design representation supporting:

1. **Customization:** intuitive controls for changing the design.
2. **Validity:** produces only functional fabricable models.
3. **Interactivity:** fast response to changing parameter values.

Our solution for creating *Fab Forms* is driven by one central assumption – the number of parameters exposed to the end-user should generally be low (typically 2-6). This goes in line with re-

ACM Reference Format

Shugrina, M., Shamir, A., Matusik, W. 2015. Fab Forms: Customizable Objects for Fabrication with Validity and Geometry Caching. ACM Trans. Graph. 34, 4, Article 100 (August 2015), 12 pages.
DOI = 10.1145/2766994 <http://doi.acm.org/10.1145/2766994>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766994>

Computational Design of Twisty Joints and Puzzles

Timothy Sun Changxi Zheng
Columbia University *



Figure 1: Twisty Armadillo. (left) A twisty puzzle in the shape of an ARMADILLO whose rotation axes are placed along a triangular prism. (right) The output of our algorithm was fabricated, assembled, and scrambled into contorted poses. The different parts of the model, such as the arms and legs, were deformed so that they do not collide with one other regardless of the configuration of the puzzle.

Abstract

We present the first computational method that allows ordinary users to create complex twisty joints and puzzles inspired by the Rubik's Cube mechanism. Given a user-supplied 3D model and a small subset of rotation axes, our method automatically adjusts those rotation axes and adds others to construct a “non-blocking” twisty joint in the shape of the 3D model. Our method outputs the shapes of pieces which can be directly 3D printed and assembled into an interlocking puzzle. We develop a group-theoretic approach to representing a wide class of twisty puzzles by establishing a connection between non-blocking twisty joints and the finite subgroups of the rotation group $SO(3)$. The theoretical foundation enables us to build an efficient system for automatically completing the set of rotation axes and fast collision detection between pieces. We also generalize the Rubik's Cube mechanism to a large family of twisty puzzles.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: Computational design, 3D fabrication, twisty puzzles, Rubik's Cube, group theory, interlocking

* e-mail: {tim, cxz}@cs.columbia.edu

ACM Reference Format
Sun, T., Zheng, C. 2015. Computational Design of Twisty Joints and Puzzles. *ACM Trans. Graph.* 34, 4, Article 101 (August 2015), 11 pages. DOI = 10.1145/2766961 <http://doi.acm.org/10.1145/2766961>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766961>

1 Introduction

Perhaps the most familiar example of a twisty joint is *Rubik's Cube* (Figure 2(a)), a 3D puzzle composed of 26 separate pieces attached to a core with six rotation axes, each of which can be rotated independently. Rubik's Cube and its variants, which are known as *twisty puzzles*, are enormously popular around the world: there are hundreds of “speedcubing” competitions for solving these puzzles every year, and new puzzle designs are being mass-produced for those seeking new challenges (Figure 2(b)). Beyond their recreational popularity, twisty joints that share similar mechanics with Rubik's Cube have found applications in many areas, including mechanical joints for robotics [Ding et al. 2011] and omnidirectional security cameras [Khoudary 2000].

The elegant and ingenious design of Rubik's Cube addresses two seemingly conflicting mechanical requirements. On one hand, it needs to be rotatable around different axes. Every rotation permutes the puzzle's pieces, which need to be aligned so that they can be rotated around other axes. On the other hand, all the pieces are interlocked such that they never fall apart in any configuration. Both goals are realized through precisely aligned rotation axes (Figure 3(a)) and the special shapes of the pieces, each with a hidden internal structure that interlocks with other pieces (Figure 3(b)). Designing new puzzles and joints is a sophisticated task, requiring expert knowledge in twisty puzzle design and often many iterations of trial and error. For example, even with CAD software, a user needs to manually check the internal mechanism for tiny undesired parts that result from misaligned cutting surfaces.




Figure 2: Twisty puzzles. *Rubik's Cube* (a) and its variants (b).

Reduced-Order Shape Optimization Using Offset Surfaces

Przemyslaw Musialski* Thomas Auzinger* Michael Birsak* Michael Wimmer* Leif Kobbelt†

*Vienna University of Technology

†RWTH Aachen University




Figure 1: We introduce a method for reduced-order shape optimization of 2-manifolds that uses offset surfaces to deform the shape. Left: a bottle model is generated using offset surfaces with constant offsets. The resulting object is unable to stand. Center: the offsets are optimized such that the bottle can stand if empty, however, if filled it is unstable. Right: the model is optimized to stand both empty and filled. In order to account for that, offset surfaces are added inside and outside of the original shape.

Abstract

Given the 2-manifold surface of a 3d object, we propose a novel method for the computation of an offset surface with varying thickness such that the solid volume between the surface and its offset satisfies a set of prescribed constraints and at the same time minimizes a given objective functional. Since the constraints as well as the objective functional can easily be adjusted to specific application requirements, our method provides a flexible and powerful tool for shape optimization. We use manifold harmonics to derive a reduced-order formulation of the optimization problem, which guarantees a smooth offset surface and speeds up the computation independently from the input mesh resolution without affecting the quality of the result. The constrained optimization problem can be solved in a numerically robust manner with commodity solvers. Furthermore, the method allows simultaneously optimizing an inner and an outer offset in order to increase the degrees of freedom. We demonstrate our method in a number of examples where we control the physical mass properties of rigid objects for the purpose of 3d printing.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

Keywords: geometry processing, geometric design optimization, shape optimization, reduced-order models, physical mass properties, digital fabrication

1 Introduction

Today’s geometric modeling software (e.g., Blender) allows for interactive design or customization of 3d geometric shapes, and many of them can now be fabricated at home using a low-cost 3d-printer. However, most such items are created in an ad-hoc fashion, i.e., their geometric and physical aspects are usually assumed intuitively or determined empirically with a series of trial-and-error iterations. While this might work well in some cases, it can also turn into a tedious and especially a costly procedure. For this reason, in the approaching age of personal digital fabrication, there is a growing demand for computational tools that not only enable ordinary users to design and 3d-print their everyday objects, but also allow them to optimize their designs for practical usability.

In this paper, we provide a novel method for shape optimization of geometric objects defined by 2-manifold surface meshes. The particular problem we are dealing with is to find a new shape that is as similar to an input shape as possible, but which at the same time satisfies various global goals. An example of such a goal is depicted in Figure 1: the bottle is intended to stand upright in a desired position when filled, however, it would fall over given its current shape. In order to prevent this, our method automatically adjusts the shape of the object, but simultaneously, it tries to preserve its volume, smoothness, and the overall detailed appearance.

Technically, we formulate this task as a continuous constrained shape optimization problem, which balances shape preservation against given design goals. We express the shape using *offset surfaces*—a simple yet powerful technique where the surface is parameterized by a vector of offset values applied to each vertex in a certain direction. This parameterization allows deforming the surface both locally and globally, depending on the chosen displacements. We describe the details in Section 4.

Since we want to preserve the characteristics of the input shape under deformations as well as possible, we favor displacements of the interior surface of the object if feasible, and penalize displacements of the outer surface explicitly. Additionally, and most importantly, we apply only low-frequency changes to the shape, which is perceptually less noticeable than high-frequency modifications, but still has a large influence on global properties, such as the object’s volume and thus mass.

ACM Reference Format

Musialski, P., Auzinger, T., Birsak, M., Wimmer, M., Kobbelt, L. 2015. Reduced-Order Shape Optimization Using Offset Surfaces. *ACM Trans. Graph.* 34, 4, Article 102 (August 2015), 9 pages. DOI: 10.1145/2766955 <http://doi.acm.org/10.1145/2766955>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766955>

RAPTER: Rebuilding Man-made Scenes with Regular Arrangements of Planes

Aron Monszpart¹

Nicolas Mellado^{1,2}

Gabriel J. Brostow¹

Niloy J. Mitra¹

¹University College London,

²Université de Toulouse; UPS; IRIT

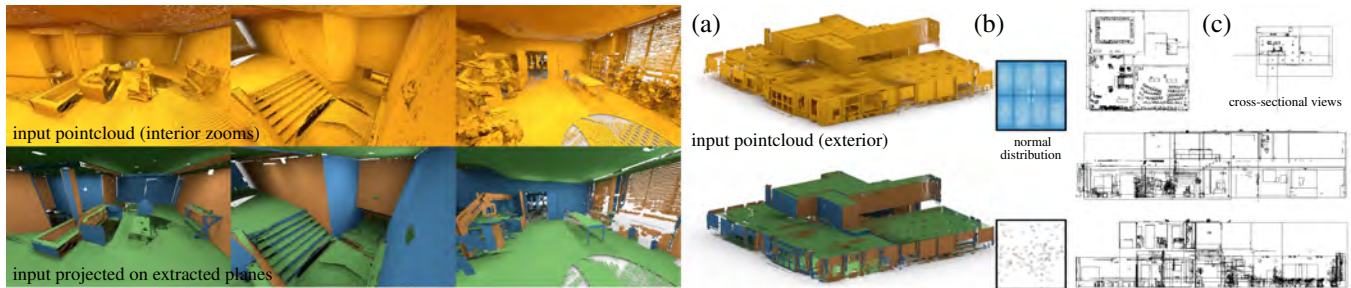


Figure 1: We present a novel approach to extract Regular Arrangements of Planes (RAP) from an unstructured and noisy raw scan (shown in gold). (a) In this example, our algorithm reconstructs a building arrangement from a raw pointcloud, pre-assembled from multiple laser scans. (b) The distribution of the initial normals is very noisy, which makes any greedy arrangement of planes error-prone. Instead, we propose a global algorithm to simultaneously select both the planes along with their sparse inter-relations. (c) Cross-sectional views reveal the discovered regularity of the extracted arrangements at multiple scales, e.g., walls, stairways, chairs, etc. Parallel planes have same color.

Abstract

With the proliferation of acquisition devices, gathering massive volumes of 3D data is now easy. Processing such large masses of pointclouds, however, remains a challenge. This is particularly a problem for raw scans with missing data, noise, and varying sampling density. In this work, we present a simple, scalable, yet powerful data reconstruction algorithm. We focus on reconstruction of man-made scenes as regular arrangements of planes (RAP), thereby selecting both local plane-based approximations along with their global inter-plane relations. We propose a novel selection formulation to directly balance between data fitting and the simplicity of the resulting arrangement of extracted planes. The main technical contribution is a formulation that allows less-dominant orientations to still retain their internal regularity, and not become overwhelmed and regularized by the dominant scene orientations. We evaluate our approach on a variety of complex 2D and 3D pointclouds, and demonstrate the advantages over existing alternative methods.

CR Categories: I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Surface fitting

Keywords: reconstruction, pointcloud, RANSAC, scene understanding, regular arrangement

1 Introduction

Pointclouds are now easy to acquire. They can be recordings of both indoor and outdoor environments, and can easily contain millions of samples. Such data volumes often retain interesting information about the captured scenes, and hence are keenly investigated in the context of scene understanding. These analyses generate valuable scene priors for various computer graphics applications. For example, they can reveal typical object arrangements in scenes [Kim et al. 2012; Mattausch et al. 2014; Chen et al. 2014], provide non-local priors for scene completion [Zheng et al. 2010], deliver workspace affordance metrics [Sharf et al. 2013; Yan et al. 2014], or aid autonomous navigation by providing scene maps [Anand et al. 2011]. A grand goal is to abstract such massive data volumes [Nießner et al. 2013; Nießner et al. 2014] to eventually produce semantic understanding of the scenes [Boulch et al. 2014].

In the context of data analysis, one interpretation of abstraction [Yumer and Kara 2012; Lafarge and Alliez 2013; Oesau et al. 2014] is to reveal interesting high-level global scene characteristics, rather than focus on local details, for example obtained via a full surface reconstruction. In the context of man-made scenes, we observe that scene characteristics are often encoded in the form of inter-part relations, inside and across objects. As many objects primarily comprise of planar faces, man-made scenes can be well abstracted as collections of planes, more importantly, planes along with their inter-relations. In this paper, we focus on the problem of reconstructing raw pointclouds as regular arrangements of planes. The resultant abstractions provide compact and simplified representations, and expose high-level structures hidden in the raw data.

In the case of noisy, incomplete, outlier-ridden data, the challenge is to balance between compliance with data, and reliance on relations to regularize the extracted planes. However, one source of difficulty is that configurations of inter-plane relations are often scene-specific due to sensitivity to the primitives' local and global environments, and hence are not known *a priori*. They can easily be missed, or even worse, wrongly identified. The same set of 3D points can be scans of two parallel walls, or a wall and an ajar door depending on the context. Specifically, it can be particularly difficult to retain non-dominant plane orientations as they can easily be masked and falsely 'regularized'. E.g., Figure 1 shows a scan

ACM Reference Format
Monszpart, A., Mellado, N., Brostow, G., Mitra, N. 2015. RAPTER: Rebuilding Man-Made Scenes with Regular Arrangements of Planes. ACM Trans. Graph. 34, 4, Article 103 (August 2015), 12 pages.
DOI = 10.1145/2766995 <http://doi.acm.org/10.1145/2766995>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766995>

Coupled Segmentation and Similarity Detection for Architectural Models

İlke Demir
Purdue University

Daniel G. Aliaga
Purdue University

Bedrich Benes
Purdue University

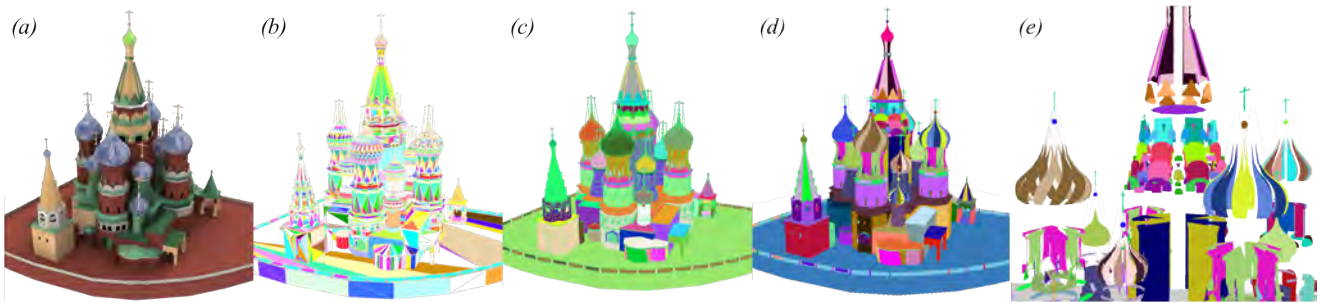


Figure 1: Overview. (a) A given input model (displayed as (b) color coded triangles), is partitioned into (c) search spaces. (d,e) Components suitable for shape analysis are generated automatically by our method.

Abstract

Recent shape retrieval and interactive modeling algorithms enable the re-use of existing models in many applications. However, most of those techniques require a pre-labeled model with some semantic information. We introduce a fully automatic approach to simultaneously segment and detect similarities within an existing 3D architectural model. Our framework approaches the segmentation problem as a weighted minimum set cover over an input triangle soup, and maximizes the repetition of similar segments to find a best set of unique component types and instances. The solution for this set-cover formulation starts with a search space reduction to eliminate unlikely combinations of triangles, and continues with a combinatorial optimization within each disjoint subspace that outputs the components and their types. We show the discovered components of a variety of architectural models obtained from public databases. We demonstrate experiments testing the robustness of our algorithm, in terms of threshold sensitivity, vertex displacement, and triangulation variations of the original model. In addition, we compare our components with those of competing approaches and evaluate our results against user-based segmentations. We have processed a database of 50 buildings, with various structures and over 200K polygons per building, with a segmentation time averaging up to 4 minutes.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric Algorithms

Keywords: segmentation, similarity detection, architectural modeling, geometry processing, shape understanding

ACM Reference Format

Demir, I., Aliaga, D., Benes, B. 2015. Coupled Segmentation and Similarity Detection for Architectural Models. ACM Trans. Graph. 34, 4, Article 104 (August 2015), 11 pages. DOI = 10.1145/2766923
<http://doi.acm.org/10.1145/2766923>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766923>

1 Introduction

Improving 3D urban modeling is beneficial to a growing number of applications in computer graphics, virtual environments, entertainment, and urban planning. One option is to efficiently re-use the existing large set of 3D polygonal models available from public databases. However these models usually lack high-level grouping or segmentation information which hampers efficient re-use and synthesis. Thus, shape-based segmentation of architectural polygon-based models is a critical research problem.

There are a variety of segmentation approaches including manual, semi-automatic, and automatic methods. Manual and semi-automatic techniques are challenging to scale to a large number of models. In this paper, we focus on fully automatic scalable methods. However, existing approaches (e.g. [Kalogerakis et al. 2010]; [Lipman et al. 2010]; [Attene et al. 2006]) hitherto first segment the model into components based on a local geometric feature analysis (e.g., creases, planar regions, primitives, etc.) and afterwards group components based on similarity. This fact does not take into account the compactness and expressiveness beneficial to editing and reuse which results in unnecessary partitioning or less-useful segmentations. For example, we want a small number of component types, each of which has many component instances, so as to permit easily changing a logically-similar geometric structure (e.g., all similar windows should be put in the same group so that all those windows can be edited equally). This is especially hard if the repeating geometric instances are not identical.

Our automatic approach couples segmentation and similarity detection of architectural 3D polygonal models while also seeking a small set of component types spanning the model with high repetition (Figure 1). We convert an arbitrary polygonal architectural model into a compact set of component types and their instances; without loss of generality, we assume the input to be a set of triangles. We define a weighted minimum exact cover problem over the input triangle set in order to reveal the implicit repetitions and similarities within the model. Seeking the minimum cover implies finding the smallest set of component types; thus minimizing the number of component types also maximizes the repetition per type. Our approach efficiently finds an approximate solution using a combinatorial optimization implemented using integer linear programming, with additional constraints. During a first phase, our method partitions the input triangles into disjoint sets – each set corresponds

Controlling Procedural Modeling Programs with Stochastically-Ordered Sequential Monte Carlo

Daniel Ritchie*
Stanford University

Ben Mildenhall*
Stanford University

Noah D. Goodman*
Stanford University

Pat Hanrahan*
Stanford University



Forward Sampling



SOSMC-Controlled Sampling



Forward Sampling



SOSMC-Controlled Sampling

Figure 1: Controlling the output of highly-variable procedural modeling programs using our Stochastically-Ordered Sequential Monte Carlo algorithm. Here, the controls encourage volumetric similarity to a target shape (shown in black).

Abstract

We present a method for controlling the output of procedural modeling programs using Sequential Monte Carlo (SMC). Previous probabilistic methods for controlling procedural models use Markov Chain Monte Carlo (MCMC), which receives control feedback only for completely-generated models. In contrast, SMC receives feedback incrementally on incomplete models, allowing it to reallocate computational resources and converge quickly. To handle the many possible sequentializations of a structured, recursive procedural modeling program, we develop and prove the correctness of a new SMC variant, Stochastically-Ordered Sequential Monte Carlo (SOSMC). We implement SOSMC for general-purpose programs using a new programming primitive: the stochastic future. Finally, we show that SOSMC reliably generates high-quality outputs for a variety of programs and control scoring functions. For small computational budgets, SOSMC's outputs often score nearly twice as high as those of MCMC or normal SMC.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems;

Keywords: Procedural Modeling, Directable Randomness, Probabilistic Programming, Sequential Monte Carlo

*e-mail: {dritchie, bmild, ngoodman, hanrahan}@stanford.edu

ACM Reference Format

Ritchie, D., Mildenhall, B., Goodman, N., Hanrahan, P. 2015. Controlling Procedural Modeling Programs with Stochastically-Ordered Sequential Monte Carlo. *ACM Trans. Graph.* 34, 4, Article 105 (August 2015), 11 pages. DOI = 10.1145/2766895 <http://doi.acm.org/10.1145/2766895>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766895>

1 Introduction

Procedural modeling has long been used in computer graphics to generate varied, detailed content with minimal human effort. Procedural models for trees, buildings, cities, and decorative patterns enrich the virtual worlds of movies and games [Měch and Prusinkiewicz 1996; Müller et al. 2006; Wong et al. 1998]. Ambitious new projects aim to produce fully-procedural, galactic-scale environments for players to explore [Procedural Reality 2014; Hello Games 2014].

This expressive power comes at a cost: procedural models often use complex, recursive control logic, resulting in emergent behavior which is difficult to direct. As a result, technical artists often must tweak parameters and massage initial conditions to achieve a desired look. This time and effort may defeat the purpose of using procedural modeling in the first place.

Fortunately, recent years have seen advances in the use of probabilistic inference techniques to control procedural models [Talton et al. 2011; Stava et al. 2014; Yeh et al. 2012]. Viewing a procedural model as sampling from a probability distribution allows for the application of Bayesian inference techniques: the prior is the procedural model itself, and the likelihood is some high-level control expressed as a scoring function.

This previous work relies on Markov Chain Monte Carlo (MCMC), but other Bayesian posterior sampling algorithms are available: another popular choice is Sequential Monte Carlo (SMC). SMC uses a set of samples, or *particles*, to represent a distribution that changes over time as new evidence is observed. As the distribution changes, SMC shifts more particles (and thus more of its computational budget) to higher-probability regions of the state space. For probabilistic models that fit this pattern of ‘evidence arriving over time,’ such as modeling the location of a mobile robot, SMC is often the method of choice: the incremental evidence it receives provides feedback early and often, allowing it to converge quickly [Doucet et al. 2001]. In contrast, MCMC receives feedback only after running through the entire model.

WorldBrush: Interactive Example-based Synthesis of Procedural Virtual Worlds

Arnaud Emilien^{1,2*} Ulysse Vimont^{1†} Marie-Paule Cani^{1‡} Pierre Poulin^{2§} Bedrich Benes^{3¶}

¹University Grenoble-Alpes, CNRS (LJK), and Inria ²LIGUM, Dept. I.R.O., Université de Montréal ³Purdue University




Figure 1: An example model (left) is used to learn parameters of distributions of trees, grass, and rocks constrained by the terrain's slope. Our model then uses these parameters to create consistent content in a larger region, using copy-paste or with user-controlled interactive brushes (center). Various adapted editing operations are introduced, inspired by painting tools, such as interpolation of parameters, that allow for creation of varying islands with trees (right).

Abstract

We present a novel approach for the interactive synthesis and editing of virtual worlds. Our method is inspired by painting operations and uses methods for statistical example-based synthesis to automate content synthesis and deformation. Our real-time approach takes a form of local inverse procedural modeling based on intermediate statistical models: selected regions of procedurally and manually constructed example scenes are analyzed, and their parameters are stored as *distributions* in a palette, similar to colors on a painter's palette. These *distributions* can then be interactively applied with brushes and combined in various ways, like in painting systems. Selected regions can also be moved or stretched while maintaining the consistency of their content. Our method captures both distributions of elements and structured objects, and models their interactions. Results range from the interactive editing of 2D artwork maps to the design of 3D virtual worlds, where constraints set by the terrain's slope are also taken into account.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Modeling packages I.3.4 [Computer Graphics]: Graphics Utilities—Graphics editors I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods—Representations (procedural and rule-based)

Keywords: computer graphics, inverse procedural modeling, interactive modeling, virtual worlds, painting systems

*e-mail:emilien@wildmagegames.com

†e-mail:ulyse.vimont@inria.fr

‡e-mail:marie-paule.cani@inria.fr

§e-mail:poulin@iro.umontreal.ca

¶e-mail:bbenes@purdue.edu

ACM Reference Format

Emilien, A., Vimont, U., Cani, M., Poulin, P., Benes, B. 2015. WorldBrush: Interactive Example-based Synthesis of Procedural Virtual Worlds. *ACM Trans. Graph.* 34, 4, Article 106 (August 2015), 11 pages. DOI = 10.1145/2766975 <http://doi.acm.org/10.1145/2766975>.

Copyright Notice

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a non-exclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766975>

1 Introduction

Geometric modeling is one of the fastest developing areas of computer graphics, yet it still presents itself with many open problems. Among them, the creation of virtual worlds is a very challenging one. Because of its diversity and specificity, the creation of virtual worlds cannot be easily automatized, while the large number of elements makes manual editing long and tedious.

Procedural methods have been for long the most popular choice for the generation of large amounts of consistent data from reduced user input [Smith 1984]. While many successful procedural methods were proposed for different domains (terrains, forests, cities, etc.) and some were combined to synthesize complex virtual worlds [Smelik et al. 2011], these methods cannot automatically respond to the designer's intent. The user then finds himself in tedious trial-and-error sessions, with no guarantee that the intended result will be finally produced. When a specific virtual world needs to be created, the user thus ends up manually manipulating a large quantity of objects, such as trees to compose a forest, or complex structured objects like houses and street segments to compose villages. In addition to taking care of scene design, the user must ensure scene consistency, which potentially breaks the artistic flow.

Just very recently, various approaches for inverse procedural modeling have been introduced [Talton et al. 2011; Vanegas et al. 2012; Št'ava et al. 2014]. Their common goal is to learn rules, or their parameters, in order to find a procedural representation of some input structure. The *key observation* of our paper is that we do not need to learn the procedural parameters of a complete scene. This task can be done instead in an interactive manner, on sub-parts of the scene or for separate categories of objects. In such a scenario, the user interactively selects parts of scenes to analyze, their parameters are learned from intermediate representations such as distributions, and they are then used to synthesize more content through interactive operations. This ability to build individual “palettes” of parameters and to reuse them during scene editing brings a simple form of inverse procedural modeling into an interactive artistic workflow.


We present a novel concept for the interactive design of consistent virtual worlds. We exploit the intuitiveness of state-of-the-art painting and editing tools that we combine with the power of procedural modeling. On the one hand, we overcome the main problem of

Advanced Procedural Modeling of Architecture

Michael Schwarz

Pascal Müller

Esri R&D Center Zurich



(a) Multi-shape coordination

(b) Boolean operations

(c) Advanced context sensitivity

(d) Best of several alternatives

Figure 1: Our novel grammar language CGA++ enables many advanced procedural modeling scenarios not possible with previous solutions (top; bottom: ours), as exemplified with a grammar for residential suburban buildings comprising a main house, a wing, and a garage, and allowing different configurations of these. (a) With CGA++, modeling decisions can be coordinated across multiple shapes, e.g., to guarantee that overall exactly one door is created. (b) CGA++ enables operations involving multiple shapes, such as Boolean operations. Hence, masses can be merged to avoid overlapping geometries, allowing, e.g., one roof covering the whole building. (c) Generic contextual information can be obtained and acted on in CGA++, whereas previous solutions at best support a narrow set of context sensitivity. While they only allow canceling windows partially occluded, CGA++ additionally enables consistently adjusting all top floor windows. (d) Traditionally, only one alternative can be pursued during one specific derivation. CGA++, however, makes it possible to investigate multiple ones and choose the best of them. On a corner lot, the building grammar may fail if it executes only one option stochastically, and the selected one causes the garage to end up on an irregular footprint. CGA++ allows all options to be explored, robustly evading such failure cases.

Abstract

We present the novel grammar language CGA++ for the procedural modeling of architecture. While existing grammar-based approaches can produce stunning results, they are limited in what modeling scenarios can be realized. In particular, many context-sensitive tasks are precluded, not least because within the rules specifying how one shape is refined, the necessary knowledge about other shapes is not available. Transcending such limitations, CGA++ significantly raises the expressiveness and offers a generic and integrated solution for many advanced procedural modeling problems. Pivotal, CGA++ grants first-class citizenship to shapes, enabling, within a grammar, directly accessing shapes and shape trees, operations on multiple shapes, rewriting shape (sub)trees, and spawning new trees (e.g., to explore multiple alternatives). The new linguistic device of events allows coordination across multiple shapes, featuring powerful dynamic grouping and synchronization. Various examples illustrate CGA++, demonstrating solutions to previously infeasible modeling challenges.

CR Categories: F.4.2 [Mathematical Logic and Formal Languages]: Grammars and Other Rewriting Systems; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

Keywords: procedural modeling, grammar language

1 Introduction

Procedural modeling techniques are successfully employed in many domains, including urban planning, computer games, and movie production, for creating numerous instances of similar but varied objects with high detail. In the important case of buildings, typically a grammar-based approach is pursued, where a set of rules describe how one shape is refined into a set of new ones. Starting from an initial shape (e.g., a lot), these rules are iteratively applied, hierarchically evolving the structure of the model and incrementally adding details.

While impressive results of high visual complexity can be created with such grammars, several advanced modeling scenarios are not feasible with current grammar languages and systems such as CGA shape [Müller et al. 2006]. In particular, they exhibit the following fundamental limitations (Fig. 1 illustrates):

1. Coordinating refinement decisions across multiple shapes is not directly supported and easily becomes impractical. Any decision that affects multiple shapes has to be made no later than at the point where these shapes' lineages diverge (and thus before these shapes exist at all) and then explicitly be passed on in rules. Notably, this implies that if this decision is influenced by properties of these (not-yet-existing) shapes or any other information only established later in the derivation process, these have to be inferred manually.
2. Operations involving multiple shapes are normally not possible. For example, neither Boolean operations, such as forming the intersection of two shapes, nor assembling a shape from multiple shapes can be expressed.
3. Contextual information is generally not available, precluding taking information from other shapes into account, as would be required for modeling objectives such as alignment. Merely for a few selected problems, ad-hoc solutions exist, such as limited occlusion queries [Müller et al. 2006].

ACM Reference Format

Schwarz, M., Müller, P. 2015. Advanced Procedural Modeling of Architecture. ACM Trans. Graph. 34, 4, Article 107 (August 2015), 12 pages. DOI = 10.1145/2766956 <http://doi.acm.org/10.1145/2766956>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766956>

Learning Shape Placements by Example

Paul Guerrero*
KAUST, Vienna University of Technology

Stefan Jeschke†
IST Austria

Michael Wimmer‡
Vienna University of Technology

Peter Wonka§
KAUST




Figure 1: New York-style skyscrapers generated by our method: Individual shape placement operations are learned during an example modeling session. Two of these operations are shown in the top row. New skyscraper variations like the ones in the bottom row can be generated from the learned model without additional user input.

Abstract

We present a method to learn and propagate *shape placements* in 2D polygonal scenes from a few examples provided by a user. The placement of a shape is modeled as an oriented bounding box. Simple geometric relationships between this bounding box and nearby scene polygons define a *feature set* for the placement. The feature sets of all example placements are then used to learn a probabilistic model over all possible placements and scenes. With this model, we can generate a new set of placements with similar geometric relationships in any given scene. We introduce extensions that enable propagation and generation of shapes in 3D scenes, as well as the application of a learned modeling *session* to large scenes without additional user interaction. These concepts allow us to generate complex scenes with thousands of objects with relatively little user interaction.

*paul@cg.tuwien.ac.at

†sjeschke@ist.ac.at

‡wimmer@cg.tuwien.ac.at

§pwonka@gmail.com

ACM Reference Format

Guerrero, P., Jeschke, S., Wimmer, M., Wonka, P. 2015. Learning Shape Placements by Example. ACM Trans. Graph. 34, 4, Article 108 (August 2015), 13 pages. DOI = 10.1145/2766933 <http://doi.acm.org/10.1145/2766933>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766933>

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems;

Keywords: modeling by example, complex model generation

1 Introduction

Recently, the amount of detail that is expected in polygonal scenes in computer graphics has increased significantly, suggesting the need for powerful tools for model creation and manipulation. Existing systems mostly employ *rule-based methods* or formulate *optimization problems*, where an expert user encodes domain-specific knowledge as a large set of rules, constraints and/or objective functions. A disadvantage of these tools is that they are very different from the commercial modeling software that most modelers are proficient with. To resolve this discrepancy, *example-based modeling methods* have recently gained interest, since they do not require a complex pre-encoding of domain knowledge, but mostly work with familiar modeling operations. One typical option is to provide one or several complete scenes as examples and to produce “similar” scenes as output. Unfortunately, the degrees of freedom grow exponentially with example scene complexity—a typical case of the “curse of dimensionality”. This means that impractically many example scenes are required to disambiguate object relations, which makes user control fairly limited.

The key problem to tackle is an efficient way to encode what valid

Skin Microstructure Deformation with Displacement Map Convolution

Koki Nagano[†] * Graham Fyffe[†] Oleg Alexander[†] Jernej Barbic[‡] Hao Li[‡] Abhijeet Ghosh^{*} Paul Debevec[†]

[†]USC Institute for Creative Technologies [‡]University of Southern California ^{*}Imperial College London




Figure 1: A rendered facial expression with (a) mesostructure only (b) static microstructure from a neutral expression (c) dynamic microstructure from convolving the neutral microstructure according to local surface strain compared to a reference photograph of the similar expression. The insets show detail from the lower-left area.

Abstract

We present a technique for synthesizing the effects of skin microstructure deformation by anisotropically convolving a high-resolution displacement map to match normal distribution changes in measured skin samples. We use a 10-micron resolution scanning technique to measure several *in vivo* skin samples as they are stretched and compressed in different directions, quantifying how stretching smooths the skin and compression makes it rougher. We tabulate the resulting surface normal distributions, and show that convolving a neutral skin microstructure displacement map with blurring and sharpening filters can mimic normal distribution changes and microstructure deformations. We implement the spatially-varying displacement map filtering on the GPU to interactively render the effects of dynamic microgeometry on animated faces obtained from high-resolution facial scans.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: skin microgeometry, deformation, measurement, displacement map, convolution

1 Introduction

Simulating the appearance of human skin is important for rendering realistic digital human characters for simulation, education, and entertainment applications. Skin exhibits great variation in color, surface roughness, and translucency over different parts of the body, between different individuals, and when it's transformed by articulation and deformation. But as variable as skin can be, human

perception is remarkably attuned to the subtleties of skin appearance, as attested to by the vast array of makeup products designed to enhance and embellish it.

Advances in measuring and simulating the scattering of light beneath the surface of the skin [Jensen et al. 2001; Weyrich et al. 2006; d'Eon et al. 2007] have made it possible to render convincingly realistic human characters whose skin appear to be fleshy and organic. Today's high-resolution facial scanning techniques (e.g. [Ma et al. 2007; Beeler et al. 2010; Ghosh et al. 2011]) record facial geometry, surface coloration, and surface *mesostructure* details at the level of skin pores and fine creases to a resolution of up to a tenth of a millimeter. By recording a sequence of such scans [Beeler et al. 2011] or performing blendshape animation using scans of different high-res expressions (e.g. [Alexander et al. 2010; Fyffe et al. 2014]), the effects of dynamic mesostructure – pore stretching and skin furrowing – can be recorded and reproduced on a digital character.

Recently, [Graham et al. 2013] recorded skin *microstructure* at a level of detail below a tenth of a millimeter for sets of skin patches on a face, and showed that texture synthesis could be used to increase the resolution of a mesostructure-resolution facial scan to one with microstructure everywhere. They demonstrated that skin microstructure makes a significant difference in the appearance of skin, as it gives rise to a face's characteristic pattern of spatially-varying surface roughness. However, they recorded skin microstructure only for static patches from neutral facial expressions, and did not record the dynamics of skin microstructure as skin stretches and compresses.

Skin microstructure, however, is remarkably dynamic as a face makes different expressions. Fig. 2 shows a person's forehead as they make surprised, neutral, and angry expressions. In the neutral expression (center), the rough surface microstructure is relatively isotropic. When the brow is raised (left), there are not only mesostructure furrows but the microstructure also develops a pattern of horizontal ridges less than 0.1 mm across. In the perplexed expression (right), the knitted brow forms vertical anisotropic structures in its microstructure. Seen face to face or filmed in closeup, such dynamic microstructure is a noticeable aspect of human expression, and the anisotropic changes in surface roughness affect the appearance of specular highlights even from a distance.

Dynamic skin microstructure results from the epidermal skin layers being stretched and compressed by motion of the tissues under-

*e-mail:nagano@ict.usc.edu

ACM Reference Format

Nagano, K., Fyffe, G., Alexander, O., Barbic, J., Li, H., Ghosh, A., Debevec, P. 2015. Skin Microstructure Deformation with Displacement Map Convolution. ACM Trans. Graph. 34, 4, Article 109 (August 2015), 10 pages. DOI = 10.1145/2766894 <http://doi.acm.org/10.1145/2766894>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766894>

Two-Shot SVBRDF Capture for Stationary Materials

Miika Aittala
Aalto University

Tim Weyrich
University College London

Jaakko Lehtinen
Aalto University, NVIDIA




Figure 1: Given an flash-no-flash image pair of a “textured” material sample, our system produces a set of spatially varying BRDF parameters (an SVBRDF, right) that can be used for relighting the surface. The capture (left) happens in-situ using a mobile phone.

Abstract

Material appearance acquisition usually makes a trade-off between acquisition effort and richness of reflectance representation. In this paper, we instead aim for both a light-weight acquisition procedure and a rich reflectance representation simultaneously, by restricting ourselves to one, but very important, class of appearance phenomena: texture-like materials. While such materials’ reflectance is generally spatially varying, they exhibit self-similarity in the sense that for any point on the texture there exist many others with similar reflectance properties. We show that the texturedness assumption allows reflectance capture using only two images of a planar sample, taken with and without a headlight flash. Our reconstruction pipeline starts with redistributing reflectance observations across the image, followed by a regularized texture statistics transfer and a non-linear optimization to fit a spatially-varying BRDF (SVBRDF) to the resulting data. The final result describes the material as spatially-varying, diffuse and specular, anisotropic reflectance over a detailed normal map. We validate the method by side-by-side and novel-view comparisons to photographs, comparing normal map resolution to sub-micron ground truth scans, as well as simulated results. Our method is robust enough to use handheld, JPEG-compressed photographs taken with a mobile phone camera and built-in flash.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance;

Keywords: appearance capture, reflectance, SVBRDF, texture synthesis

1 Introduction

Spatial variation in surface reflectance plays an immense role in the perceived realism of imagery. Traditionally, reflectance modeling has concentrated on the angular aspects of scattering (BSDF), building models that faithfully reproduce the shape and color of reflection and transmission from a single, uniform, material. However, subtle changes in reflectance, for instance those induced by wear and tear, scratches, and the like, are required to really bring a surface to life. Even simple angular reflectance models can produce results of impressive realism when their parameters are carefully varied across surface points. This observation has relatively recently led to the search for efficient and light-weight methods for capturing models that reconstruct such variation, as spatially-varying BRDF (SVBRDF), from real surfaces [Dong et al. 2011; Wang et al. 2011; Aittala et al. 2013].

In this paper, we describe a combination of acquisition setup design, simplifying material assumptions, and reconstruction algorithm that leads, we believe, to the lightest-yet setup for photographic capture of a full SVBRDF model including albedos, glossiness, and normals. This is made possible by assuming that the SVBRDF is part of a large and important class of “texture-like” materials that consist of elements or structures that (randomly) repeat over the surface, also known as *stationary materials*. This allows us to combine information across the material exemplar.

Our input consists of two approximately fronto-parallel images of a relatively flat material sample, one taken in the ambient illumination, and another taken with a flash located close to the lens, that is, under headlight illumination (see Figure 1); in practice, we use a handheld Apple iPhone 5 and perform no calibration of any kind. The flash image provides local illumination whose relative direction of incidence varies across the surface, thus providing us with lighting-dependent samples of the texture’s reflectance. The second image serves as a guide that allows us to identify similarities across the texture, under the more homogenous ambient illumination.

We present a multi-stage reconstruction pipeline that starts with gathering reflectance observations from across the image to augment information content at one representative image region, followed by structural regularization and subsequent texture statistics transfer, before we perform a non-linear optimization to fit a SVBRDF to the resulting data. From here, we propagate back this partial solution

ACM Reference Format

Aittala, M., Weyrich, T., Lehtinen, J. 2015. Two-Shot SVBRDF Capture for Stationary Materials. *ACM Trans. Graph.* 34, 4, Article 110 (August 2015), 13 pages. DOI = 10.1145/2766967 <http://doi.acm.org/10.1145/2766967>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766967>

Image Based Relighting Using Neural Networks

Peiran REN* Yue DONG Stephen LIN Xin TONG Baining GUO
Microsoft Research



Figure 1: Relighting of various scenes using light transport captured by our method from a small number of images.

Abstract

We present a neural network regression method for relighting real-world scenes from a small number of images. The relighting in this work is formulated as the product of the scene’s light transport matrix and new lighting vectors, with the light transport matrix reconstructed from the input images. Based on the observation that there should exist non-linear local coherence in the light transport matrix, our method approximates matrix segments using neural networks that model light transport as a non-linear function of light source position and pixel coordinates. Central to this approach is a proposed neural network design which incorporates various elements that facilitate modeling of light transport from a small image set. In contrast to most image based relighting techniques, this regression-based approach allows input images to be captured under arbitrary illumination conditions, including light sources moved freely by hand. We validate our method with light transport data of real scenes containing complex lighting effects, and demonstrate that fewer input images are required in comparison to related techniques.

CR Categories: I.2.10 [Artificial Intelligence]: Vision and Scene Understanding—intensity, color, photometry and thresholding; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—radiometry, reflectance, scanning

Keywords: image based relighting, light transport, neural network, clustering

*e-mail:renpeiran@gmail.com

ACM Reference Format

Ren, P., Dong, Y., Lin, S., Tong, X., Guo, B. 2015. Image Based Relighting Using Neural Networks. ACM Trans. Graph. 34, 4, Article 111 (August 2015), 12 pages. DOI = 10.1145/2766899
<http://doi.acm.org/10.1145/2766899>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766899>

1 Introduction

The appearance of a scene arises from the transport of light within it. In realistic rendering algorithms, this light transport is computed from complete scene information including geometry, reflectance properties, and lighting environment. With this information, the new appearance of the scene under different illumination can be readily determined. For a real-world scene where such data is usually unavailable, the effects of light transport can instead be inferred from images that exhibit scene appearance under different lighting conditions. Represented in a *light transport matrix* that relates image radiance to lighting condition, this light transport information can be used to relight real-world scenes through computation of a matrix-vector product [Ng et al. 2003]:

$$\mathbf{I} = \mathbf{M} \mathbf{L}, \quad (1)$$

where the outgoing radiance \mathbf{I} is expressed as a vector over N_p image pixels, the lighting condition \mathbf{L} is modeled by a vector of incident radiance from N_s light sources, and the light transport matrix \mathbf{M} is of dimension $N_p \times N_s$. Image-based relighting in this manner produces high realism without the need for scene modeling. The key challenge of this approach is in reconstructing the light transport matrix from images acquired of the scene.

Various techniques have been presented in the literature for reconstruction of the light transport matrix. A brute-force solution is to directly measure the entries of the light transport matrix [Debevec et al. 2000; Wenger et al. 2005] by capturing an image of the scene under each of the canonical light sources (corresponding to the columns of the light transport matrix). This requires acquisition of a considerable number of images, and specialized devices are needed to accurately control the lighting. Another approach is to exploit sparseness [Zongker et al. 1999; Peers et al. 2009; Sen and Darabi 2009] or coherence [Fuchs et al. 2007; Wang et al. 2009; O’Toole and Kutulakos 2010] in the light transport matrix to reduce the number of images needed for light transport reconstruction. However, these methods are either designed for specific lighting effects [Zongker et al. 1999] or rely on special hardware to capture images under particular lighting and/or viewing conditions [Fuchs et al. 2007; Wang et al. 2009; O’Toole and Kutulakos 2010]. For scenes with occlusions and high-frequency lighting effects (e.g., hard shadows, sharp specular reflections, and caustics), these techniques require numerous images to reconstruct a high-resolution light transport matrix.

Measurement-based Editing of Diffuse Albedo with Consistent Interreflections

Bo Dong¹ Yue Dong² Xin Tong² Pieter Peers¹

¹College of William & Mary ²Microsoft Research



Figure 1: Changing the diffuse albedo with consistent interreflections for a decorative table centerpiece and a bathroom scene. The difference images (inset) between the initial scene (left column) and the recolored results (middle and right) highlight the complex changes in interreflections due to a change in diffuse albedo of the planter and middle towel respectively.

Abstract

We present a novel measurement-based method for editing the albedo of diffuse surfaces with consistent interreflections in a photograph of a scene under natural lighting. Key to our method is a novel technique for decomposing a photograph of a scene in several images that encode how much of the observed radiance has interacted a specified number of times with the target diffuse surface. Altering the albedo of the target area is then simply a weighted sum of the decomposed components. We estimate the interaction components by recursively applying the light transport operator and formulate the resulting radiance in each recursion as a linear expression in terms of the relevant interaction components. Our method only requires a camera-projector pair, and the number of required measurements per scene is linearly proportional to the decomposition degree for a single target area. Our method does not impose restrictions on the lighting or on the material properties in the unaltered part of the scene. Furthermore, we extend our method to accommodate editing of the albedo in multiple target areas with consistent interreflections and we introduce a prediction model for reducing the acquisition cost. We demonstrate our method on a variety of scenes and validate the accuracy on both synthetic and real examples.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

Keywords: Diffuse Albedo Editing, Interreflection Decomposition, Computational Illumination

1 Introduction

Typical photo-editing operations directly manipulate pixel values without any notion of the underlying physics of light transport through the scene, and significant artistic skill is required to produce plausible results. One such operation is changing the color or albedo of a surface in a photograph. While changing the directly reflected appearance of the target surface is easy, correctly changing the effects of indirect lighting on other surfaces, as well as interreflections back to the target surface, is non-trivial.

ACM Reference Format

Dong, B., Dong, Y., Tong, X., Peers, P. 2015. Measurement-Based Editing of Diffuse Albedo with Consistent Interreflections. ACM Trans. Graph. 34, 4, Article 112 (August 2015), 11 pages. DOI = 10.1145/2766979 <http://doi.acm.org/10.1145/2766979>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766979>

OmniAD: Data-driven Omni-directional Aerodynamics

Tobias Martin*
ETH Zürich

Nobuyuki Umetani*
Disney Research Zürich

Bernd Bickel
IST Austria




Figure 1: (a) Example sequence of a falling heart-shaped foam object. (b) Acquired omni-directional aerodynamic property for force and torque generated by air. (c) Real-time simulation of many hearts. (d) Flyable design of a kite carrying the heart inside.

Abstract

This paper introduces “OmniAD,” a novel data-driven pipeline to model and acquire the aerodynamics of three-dimensional rigid objects. Traditionally, aerodynamics are examined through elaborate wind tunnel experiments or expensive fluid dynamics computations, and are only measured for a small number of discrete wind directions. OmniAD allows the evaluation of aerodynamic forces, such as drag and lift, for any incoming wind direction using a novel representation based on spherical harmonics. Our data-driven technique acquires the aerodynamic properties of an object simply by capturing its falling motion using a single camera. Once model parameters are estimated, OmniAD enables realistic real-time simulation of rigid bodies, such as the tumbling and gliding of leaves, without simulating the surrounding air. In addition, we propose an intuitive user interface based on OmniAD to interactively design three-dimensional kites that actually fly. Various non-traditional kites were designed to demonstrate the physical validity of our model.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: aerodynamics, real-time animation, data-driven physics, fabrication, kites

*Tobias Martin and Nobuyuki Umetani are joint first authors

ACM Reference Format
Martin, T., Umetani, N., Bickel, B. 2015. OmniAD: Data-driven Omni-directional Aerodynamics. ACM Trans. Graph. 34, 4, Article 113 (August 2015), 12 pages. DOI = 10.1145/2766919
<http://doi.acm.org/10.1145/2766919>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766919>

1 Introduction

The interaction between lightweight objects and air results in intricate motions and stunning aerodynamic effects. For example, leaves falling to the ground, while flipping, tumbling, and gliding through the air, are a fascinating natural phenomenon. Ingenious inventions, such as kites, sailboats, windmills, and airplanes, harness the power of aerodynamics and have brought tremendous benefits to mankind. The resulting motion of these objects is determined partially by the interplay between the object’s surface geometry and the dynamics of the surrounding air. There are numerous applications in science and engineering that rely on simulating this effect, from animating a simple autumn scene in a computer game to the computational design of complex aircraft.

Traditionally, wind tunnels and computational fluid dynamics (CFD) simulations have been used to measure and predict aerodynamic forces. Such elaborate machinery, however, usually requires time-consuming set-up, is computationally expensive, and is only accessible to small numbers of domain experts. This limits the potential applications in animation and in computational design tools for a more general audience.

In this paper, we present *Omni-directional AeroDynamics (OmniAD)*, a novel data-driven representation and modeling technique for simulating the aerodynamic effects of lightweight objects. Our method is tailored to allow the simple acquisition of aerodynamic properties, requiring only a single camera, and is targeted for interactive application scenarios, carefully balancing simulation speed and accuracy. Our model operates in real-time, adds only a negligible overhead to the computational cost, and can be integrated into existing rigid-body simulators. For example, the complex motion of falling objects, such as hundreds of leaves, can be animated with our model in real-time. Moreover, the quantitative accuracy of our model allows the interactive design of real-world aerodynamically functional objects.

OmniAD is data-driven and all required model data are acquired by recording videos of falling objects with a single camera. From these input videos, the three-dimensional (3D) motion trajectory is recon-

An Implicit Viscosity Formulation for SPH Fluids

Andreas Peer* Markus Ihmsen Jens Cornelis Matthias Teschner

University of Freiburg



Figure 1: Interaction of a highly viscous fluid with complex moving solids. Up to 780k particles are used. The computation time is 30s per frame.

Abstract

We present a novel implicit formulation for highly viscous fluids simulated with Smoothed Particle Hydrodynamics SPH. Compared to explicit methods, our formulation is significantly more efficient and handles a larger range of viscosities. Differing from existing implicit formulations, our approach reconstructs the velocity field from a target velocity gradient. This gradient encodes a desired shear-rate damping and preserves the velocity divergence that is introduced by the SPH pressure solver to counteract density deviations. The target gradient ensures that pressure and viscosity computation do not interfere. Therefore, only one pressure projection step is required, which is in contrast to state-of-the-art implicit Eulerian formulations. While our model differs from true viscosity in that vorticity diffusion is not encoded in the target gradient, it nevertheless captures many of the qualitative behaviors of viscous liquids. Our formulation can easily be incorporated into complex scenarios with one- and two-way coupled solids and multiple fluid phases with different densities and viscosities.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: Physically-based animation, fluid simulation, Smoothed Particle Hydrodynamics, viscosity

1 Introduction

Simulating highly viscous fluids such as honey, mud, toothpaste or dough as shown in Fig. 1 is involved. Explicit approaches require small timesteps, e.g., [Monaghan 1989; Foster and Metaxas 1996;

Morris et al. 1997], while implicit formulations need to solve a linear system, e.g., [Stam 1999; Carlson et al. 2002; Rasmussen et al. 2004; Batty and Bridson 2008]. As implicit formulations work with significantly larger timesteps, they are commonly preferred over explicit formulations.

In highly viscous fluids, viscosity and incompressibility constraints can interfere. This is particularly true for bulk viscosity that influences the divergence of the velocity field. In explicit formulations with small timesteps, competing constraints might be neglected. In implicit formulations with large timesteps, however, this issue has to be addressed. In this context, viscosity formulations for incompressible Eulerian fluids commonly assume that the input velocity field is divergence-free. As the resulting velocity field is not necessarily divergence-free, two pressure projection steps are performed, e.g. [Losasso et al. 2006; Batty and Bridson 2008].

Contribution: This paper proposes a novel implicit formulation for highly viscous SPH fluids that particularly addresses the interference of pressure and viscosity computation. Therefore, a target velocity gradient is employed that does not only encode the desired viscosity, but also preserves arbitrary velocity divergences that might have been introduced by an SPH pressure solver. In contrast to incompressible Eulerian techniques, SPH pressure solvers, e.g. implicit incompressible SPH (IISPH) [Ihmsen et al. 2014a; Ihmsen et al. 2014b], commonly introduce some divergence to the velocity field to counteract density deviations. Our formulation explicitly preserves density corrections that have been computed in the preceding pressure projection. This also means that only one pressure projection step is required, which is in contrast to state-of-the-art implicit Eulerian formulations. The proposed formulation is approximate and departs from physical models in two ways. It is parameterized with a non-physical constant, and the vorticity diffusion is not encoded in the formulation of the target velocity gradient, but considered in the reconstruction process of the final velocity field. Experiments illustrate the close relation of our approach to true viscosity, its computational efficiency and also the large range of viscosities that can be handled. Complex scenarios with one- and two-way coupled solids and multiple fluid phases are presented.

2 Related work

Lagrangian techniques: Viscosity is an important stability aspect in SPH fluid simulations. The typically required discretiza-

*e-mail:peer@cs.uni-freiburg.de

ACM Reference Format

Peer, A., Ihmsen, M., Cornelis, J., Teschner, M. 2015. An Implicit Viscosity Formulation for SPH Fluids. ACM Trans. Graph. 34, 4, Article 114 (August 2015), 10 pages. DOI = 10.1145/2766925
<http://doi.acm.org/10.1145/2766925>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766925>

Codimensional Non-Newtonian Fluids

Bo Zhu*
Stanford University

Minjae Lee*
Stanford University

Ed Quigley*
Stanford University

Ronald Fedkiw*
Stanford University
Industrial Light + Magic



Figure 1: Examples of codimensional non-Newtonian fluid phenomena. (Far Left) A brush is dragged across paint on a canvas, leaving small furrows in its wake. (Middle Left) Cheese is stretched into thin sheets and filaments as a slice of pizza is removed. (Middle Right) A knife leaves furrows in mayonnaise, then spreads a thin layer onto a slice of bread. (Far Right) Multicolored toothpaste is squeezed out of a tube in a twisting motion.

Abstract

We present a novel method to simulate codimensional non-Newtonian fluids on simplicial complexes. Our method extends previous work for codimensional incompressible flow to various types of non-Newtonian fluids including both shear thinning and thickening, Bingham plastics, and elastoplastics. We propose a novel time integration scheme for semi-implicitly treating elasticity, which when combined with a semi-implicit method for variable viscosity alleviates the need for small time steps. Furthermore, we propose an improved treatment of viscosity on the rims of thin fluid sheets that allows us to capture their elusive, visually appealing twisting motion. In order to simulate complex phenomena such as the mixing of colored paint, we adopt a multiple level set framework and propose a discretization on simplicial complexes that facilitates the tracking of material interfaces across codimensions. We demonstrate the efficacy of our approach by simulating a wide variety of non-Newtonian fluid phenomena exhibiting various codimensional features.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: non-Newtonian fluids, codimension

1 Introduction

Non-Newtonian fluids exhibit many different codimensional features that are visually interesting. The furrows made by a brush

moving through paint, the thin sheet and filaments of cheese on a pizza, and the thin filaments of toothpaste are just some of the many examples of viscoelastic phenomena to which a codimensional representation is naturally amenable. The special material properties of non-Newtonian fluids make their codimensional motions even more interesting. For example, paint (a shear thinning fluid) has low viscosity at high shear rates making it easy to apply to surfaces using a brush; however, its high viscosity at low shear rates prevents it from running or dripping after being applied. Contrast this behavior with quicksand (a shear thickening fluid) which flows at low shear rates allowing one to sink, but which becomes viscous at high shear rates preventing one from climbing out. While non-Newtonian fluids are frequently studied in both computer graphics and computational physics, most state-of-the-art methods focus on modeling them volumetrically, and few address their codimensional features such as thin sheets and filaments.

Motivated by the codimensional fluid simulation of [Zhu et al. 2014], the codimensional solid simulation of [Martin et al. 2010], and the unified framework for simulating fluids and solids in [Macklin et al. 2014], we aim to simulate a wide range of non-Newtonian fluid behaviors especially focusing on codimensional phenomena such as twisting thin films, viscoelastic filaments, and

*e-mail: {boolzhu,mjgg,equigley,rfedkiw}@stanford.edu

ACM Reference Format

Zhu, B., Lee, M., Quigley, E., Fedkiw, R. 2015. Codimensional Non-Newtonian Fluids. ACM Trans. Graph. 34, 4, Article 115 (August 2015), 9 pages. DOI = 10.1145/2766981 http://doi.acm.org/10.1145/2766981.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: http://dx.doi.org/10.1145/2766981




Figure 2: Plots of viscosity vs. shear rate for various types of fluid.

Animating Human Dressing

Alexander Clegg *

Jie Tan *

Greg Turk †

C. Karen Liu †

Georgia Institute of Technology




Figure 1: A character puts on a jacket.

Abstract

Dressing is one of the most common activities in human society. Perfecting the skill of dressing can take an average child three to four years of daily practice. The challenge is primarily due to the combined difficulty of coordinating different body parts and manipulating soft and deformable objects (clothes). We present a technique to synthesize human dressing by controlling a human character to put on an article of simulated clothing. We identify a set of *primitive actions* which account for the vast majority of motions observed in human dressing. These primitive actions can be assembled into a variety of motion sequences for dressing different garments with different styles. Exploiting both feed-forward and feedback control mechanisms, we develop a dressing controller to handle each of the primitive actions. The controller plans a path to achieve the action goal while making constant adjustments locally based on the current state of the simulated cloth when necessary. We demonstrate that our framework is versatile and able to animate dressing with different clothing types including a jacket, a pair of shorts, a robe, and a vest. Our controller is also robust to different cloth mesh resolutions which can cause the cloth simulator to generate significantly different cloth motions. In addition, we show that the same controller can be extended to assistive dressing.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation; I.6.8 [Simulation and Modeling]: Types of Simulation—Animation.

Keywords: Human figure animation, cloth simulation, path planning.

*e-mail: {aclegg3, jtan34}@gatech.edu

†e-mail: {turk, karenliu}@cc.gatech.edu

ACM Reference Format

Clegg, A., Tan, J., Turk, G., Liu, C. 2015. Animating Human Dressing. *ACM Trans. Graph.* 34, 4, Article 116 (August 2015), 9 pages. DOI = 10.1145/2766986 <http://doi.acm.org/10.1145/2766986>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766986>

1 Introduction

This paper describes a system for animating the activity of putting on clothing. Dressing is one of the most common activities that each of us carries out each day. Scenes of dressing are also common in live-action movies and television. Some of these scenes are iconic, such as the “jacket on, jacket off” drill in *The Karate Kid* (2010 version) or Spiderman pulling his mask over his head for the first time. Such dressing scenes are noticeably absent in computer animated films. Despite the importance of dressing in our lives and in film, there is as yet no systematic approach to animating a human that is putting on clothing.

Our goal is to provide a system that will allow an animator to create motion for a human character that is dressing. We want the animator to have a high degree of control over the look of the final animation. To this end, we desire a system that allows the user to describe the dressing scene as a sequence of high-level actions. Also, we would like our system to accept approximated human motion, in either the form of keyframes or motion capture, as reference for styles or aesthetics. Thus the input from the animator for a given dressing scene consists of: a character model, a garment, a sequence of dressing actions, and reference motions for the character. In order to create animation that is physically plausible, we made the choice to use physical simulation of cloth to guide the garment motions. By using cloth simulation, the human figure, made of a collection of rigid segments, can interact with the cloth in a natural manner.

The essence of animating the act of dressing is modeling the interaction between the human character and the cloth. The human’s motion must adapt to the motion of the cloth, otherwise problems occur such as the clothing slipping off or a hand getting stuck in a fold. We often take for granted the complex set of motions that are needed to put on our clothes. The seemingly simple act of putting on a jacket requires a careful coordination between the person and the jacket. Unconsciously we make constant adjustments to our hand’s position when inserting it into the jacket’s sleeve. We hold our body at an angle to keep a sleeve from sliding off our shoulder. After putting on the first sleeve, we may use any of several strategies to get our hand behind our back and within reach of the second sleeve. A system for animation of dressing must address these kinds of complexities.

A Perceptual Control Space for Garment Simulation

Leonid Sigal¹ Moshe Mahler¹ Spencer Diaz¹ Kyna McIntosh¹ Elizabeth Carter¹
Timothy Richards² Jessica Hodgins¹
¹Disney Research ²Walt Disney Animation Studios




Figure 1: Perceptual Control Space: The perceptual control model, learned from perceptual experiments conducted on a skirt, is illustrated on a skirt on the right and cape on the left. Left: Increasing levels of the silky trait. Right: Increasing the wrinkly trait.

Abstract

We present a perceptual control space for simulation of cloth that works with any physical simulator, treating it as a black box. The perceptual control space provides intuitive, art-directable control over the simulation behavior based on a learned mapping from common descriptors for cloth (e.g., flowiness, softness) to the parameters of the simulation. To learn the mapping, we perform a series of perceptual experiments in which the simulation parameters are varied and participants assess the values of the common terms of the cloth on a scale. A multi-dimensional sub-space regression is performed on the results to build a perceptual generative model over the simulator parameters. We evaluate the perceptual control space by demonstrating that the generative model does in fact create simulated clothing that is rated by participants as having the expected properties. We also show that this perceptual control space generalizes to garments and motions not in the original experiments.

CR Categories: I.3.7 [Three-Dimensional Graphics and Realism]: Animation—I.2.10: Vision and Scene Understanding—Perceptual reasoning

Keywords: perception, physical simulation, cloth simulation

1 Introduction

Cloth simulation can produce stunningly realistic examples of garments for virtual characters. The garments include detailed folding and wrinkling [Baraff and Witkin 1998; Bridson et al. 2003; Choi and Ko 2002; Grinspun et al. 2002], as well as other dynamic behaviors, for a variety of woven materials and knits [Kaldor et al. 2010]. The complex attire animated in recent feature films, such

as *Brave* and *Frozen*, provides a concrete demonstration of the versatility of modern cloth simulators. However, art-directed control of the motion of the simulated clothing is still extremely difficult. Cloth simulators, such as Maya’s nCloth and the proprietary programs used in feature production pipelines, contain tens or hundreds of parameters that modulate the dynamic behavior of the cloth and environment. These simulation parameters often loosely represent material properties and the various physical quantities of cloth, but there is no obvious mapping from the parameter labels to the common terms used to describe a physical piece of material or a garment, such as *softness*, *flowiness*, or *silkeness*.

While it is possible to measure physical parameters from samples of actual cloth [Miguel et al. 2012; Bhat et al. 2003], these parameters may not lead to the desired behavior because the parameters in the simulation model are an approximation of actual physical parameters. Moreover, dynamic parameters are particularly important to the quality of the resulting motion but are also the most difficult to measure. Further, fabric properties that do not necessarily correspond to a physical, measurable piece of fabric are often required to achieve a desired look for virtual clothing or scenery.

As a result, technical directors are tasked with the job of manually finding parameters that achieve a desired look by exploring the space of the simulator, often by sampling one or a few parameters at a time. Their job is further complicated because many parameters interact in complex ways to produce the final behavior of the garment. For example, parameters are not always independent, and there may be multiple ways to achieve the same look. Ideally, we would like to have controls for cloth simulators that are intuitive and perceptually meaningful, making the process of creating animated cloth faster and easier without reducing the ability to fine-tune the simulations as desired for a particular shot. We propose a methodology to achieve this goal by re-parameterizing cloth simulators to work with custom, perceptually discovered controls.

From conversations with technical directors, we learned that the terms used to critique their simulations were similar to the descriptions and adjectives used to describe cloth in the textile industry. For example, *china silk* is described by fabricdictionary.com as “lightweight and soft fabric”, *chiffon* as “lightweight, extremely sheer, airy, and soft fabric”, *taffeta* as “stiffened fabric with a crisp feel”. The adjectives used in these descriptions, *lightweight*, *soft*, and *stiff*, can be interpreted as traits of the fabrics (using the nomenclature of [Matusik et al. 2002]). Such traits are intuitive to understand and can effectively span the space of both the real and the

ACM Reference Format

Sigal, L., Mahler, M., Diaz, S., McIntosh, K., Carter, E., Richards, T., Hodgins, J. 2015. A Perceptual Control Space for Garment Simulation. ACM Trans. Graph. 34, 4, Article 117 (August 2015), 10 pages. DOI = 10.1145/2766971 <http://doi.acm.org/10.1145/2766971>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI = <http://dx.doi.org/10.1145/2766971>

Space-time sketching of character animation

Martin Guay*
Université de Grenoble
LJK, INRIA

Rémi Ronfard
Université de Grenoble
LJK, INRIA

Michael Gleicher
University of Wisconsin
Madison

Marie-Paule Cani
Université de Grenoble
LJK, INRIA

Abstract

We present a space-time abstraction for the sketch-based design of character animation. It allows animators to draft a full coordinated motion using a single stroke called the *space-time curve* (STC). From the STC we compute a dynamic line of action (DLOA) that drives the motion of a 3D character through projective constraints. Our dynamic models for the line's motion are entirely geometric, require no pre-existing data, and allow full artistic control. The resulting DLOA can be refined by over-sketching strokes along the space-time curve, or by composing another DLOA on top leading to control over complex motions with few strokes. Additionally, the resulting dynamic line of action can be applied to arbitrary body parts or characters. To match a 3D character to the 2D line over time, we introduce a robust matching algorithm based on closed-form solutions, yielding a tight match while allowing squash and stretch of the character's skeleton. Our experiments show that space-time sketching has the potential of bringing animation design within the reach of beginners while saving time for skilled artists.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: Sketch-based animation, space-time, stylized animation, squash-and-stretch.

1 Introduction

Creating artistic and exaggerated styles of character animation requires flexible tools that allow expressive devices such as squash and stretch, as well as animating imaginary creatures such as dragons—often precluding the use of motion capture or database look-up. Yet, creating quality movements with current free-form animation technologies is a challenge.

The main approach to free-form motion design is keyframing: character poses at specific times are interpolated to produce motion. Over the years, significant advances have been made to more naturally specify key-poses. For example, by sketching skeletons or lines of action, or by handling deformations. However, the standard keyframing approach divides spatial and temporal control, making coordination of shape over time difficult. Hence, achieving quality results with the standard approach remains beyond the ability of unskilled artists and time consuming for skilled ones.

*lemailmartin@gmail.com

ACM Reference Format
Guay, M., Ronfard, R., Gleicher, M., Cani, M. 2015. Space-time Sketching of Character Animation. ACM Trans. Graph. 34, 4, Article 118 (August 2015), 10 pages. DOI = 10.1145/2766893
<http://doi.acm.org/10.1145/2766893>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766893>




Figure 1: Current shape interpolation techniques assume point-to-point blending (first row, result shown in grey), making it hard to create path-following motions. In contrast, our space-time sketching abstraction enables animators to sketch shapes and paths with a single stroke (second row).

In this work, we introduce a novel space-time sketching concept enabling an animator to draft a full coordinated movement—that includes shape deformation over time—by sketching a single stroke. Further strokes can be used to progressively refine the animation. While strokes have been used in the past to specify both temporal and spatial iso-values of motion—with static lines of action (LOA) serving as shape abstraction at a given time as well as trajectories describing the successive positions over time of a single point—*space-time sketching* was never used to define animations. In our approach, we allow the user to control both the shape and trajectory of a character by sketching a single *space-time curve* (STC).

To illustrate our approach, consider the simple example of animating a flying dragon (Fig. 1). Animating the dragon requires the coordination of its shape over time as to follow the path's shape. With our approach, the basic animation can be created with a single sketched stroke. The stroke is used not only to provide the path of travel, but also to define how an abstraction of the character's shape (its line of action) changes over time. Additional strokes can be used to refine the movement, or add details such as the flapping of the wings. Creating such motion with existing approaches would require coordinating a large number of keyframes that specify deformations and positions along the path, or a method for puppeteering the degrees of freedom of the dragon.

The key to our approach is an interpretation of the space-time curve to define a 2D dynamic skeletal line abstraction, or *dynamic line of action* (DLOA). From the STC, we extract the DLOA's

Realtime Style Transfer for Unlabeled Heterogeneous Human Motion

Shihong Xia¹

Congyi Wang¹

Jinxiang Chai^{2*}

Jessica Hodgins³

¹Institute of Computing Technology, CAS

²Texas A&M University

³Carnegie Mellon University




Figure 1: Our realtime style generation system automatically transforms an unlabeled, heterogeneous motion sequence into a new style. (top) the input motion in the “neutral” style; (bottom) the output animation in the “proud” style. Note the more energetic arm motions and jump in the stylized motion.

Abstract

This paper presents a novel solution for realtime generation of stylistic human motion that automatically transforms unlabeled, heterogeneous motion data into new styles. The key idea of our approach is an *online learning* algorithm that automatically constructs a series of local mixtures of autoregressive models (MAR) to capture the complex relationships between styles of motion. We construct local MAR models on the fly by searching for the closest examples of each input pose in the database. Once the model parameters are estimated from the training data, the model adapts the current pose with simple linear transformations. In addition, we introduce an efficient local regression model to predict the timings of synthesized poses in the output style. We demonstrate the power of our approach by transferring stylistic human motion for a wide variety of actions, including walking, running, punching, kicking, jumping and transitions between those behaviors. Our method achieves superior performance in a comparison against alternative methods. We have also performed experiments to evaluate the generalization ability of our data-driven model as well as the key components of our system.

I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—animation;

Keywords: Character animation, realtime style transfer, online local regression, data-driven motion synthesis

*Contact author: jchai@cs.tamu.edu

ACM Reference Format

Xia, S., Wang, C., Chai, J., Hodgins, J. 2015. Realtime Style Transfer for Unlabeled Heterogeneous Human Motion. ACM Trans. Graph. 34, 4, Article 119 (August 2015), 10 pages. DOI = 10.1145/2766999 <http://doi.acm.org/10.1145/2766999>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766999>

1 Introduction

Stylistic variations in motion that suggest a particular personality, mood or role are essential for storytelling and for bringing characters to life. For everyday actions such as locomotion, the style or emotion of a motion might convey more meaning than the underlying motion itself. For example, the anger conveyed by “he stalked across the room” is likely more significant than the act of crossing the room. Thus far, one of the most successful solutions to this problem is to model style differences between example motions and use the learned model to transfer motion from one style to another one (e.g., [Amaya et al. 1996; Hsu et al. 2005]).

Hsu and colleagues [2005] introduced a linear time-invariant (LTI) model to encode style differences between example motions and achieved realtime performance for transferring the input motion from one style to another. Despite the progress made over the past decade, creating appropriate data-driven models for online style transfer remains challenging for two reasons.

First, a lifelike human character must possess a rich repertoire of activities and display a wide range of variations for a given action. This task inevitably requires data-driven models that can scale to large and heterogeneous motion data sets. However, most existing techniques are challenged when applied to input motion that contains heterogeneous behaviors such as *walking* \Rightarrow *running* \Rightarrow *jumping* because they assume that the relationship between the input and output styles can be approximated by a global linear model, such as the LTI model described by Hsu and colleagues [2005]. Global linear models are often appropriate for homogenous data sets (e.g., walking), but might not be suitable for complex heterogeneous motion data.

A second challenge is that previous work on style transfer often requires that the input motion is labeled in terms of behavior and style attributes before transferring it to a new style. However, automatically labeling the input motion in terms of behavior and style attributes remains difficult, particularly for online applications. The problem becomes even worse when the input motion cannot be classified into a single predefined behavior or style. For example, a “proud” walk performed by an old man might

Dyna: A Model of Dynamic Human Shape in Motion

Gerard Pons-Moll* Javier Romero* Naureen Mahmood* Michael J. Black*
Max Planck Institute for Intelligent Systems, Tübingen, Germany




Figure 1: Dyna: Three different animated body shapes performing different motions with soft-tissue deformations predicted by Dyna. These deformations are predicted by a learned function that depends on body shape, the angular velocity and acceleration of body parts, the velocity and acceleration of the body, and the history of previous non-rigid deformations. Dyna generalizes to new body shapes of varying body mass index and can be applied to stylized characters to automatically add realistic soft-tissue motions to animations.

Abstract

To look human, digital full-body avatars need to have soft-tissue deformations like those of real people. We learn a model of soft-tissue deformations from examples using a high-resolution 4D capture system and a method that accurately registers a template mesh to sequences of 3D scans. Using over 40,000 scans of ten subjects, we learn how soft-tissue motion causes mesh triangles to deform relative to a base 3D body model. Our *Dyna* model uses a low-dimensional linear subspace to approximate soft-tissue deformation and relates the subspace coefficients to the changing pose of the body. Dyna uses a second-order auto-regressive model that predicts soft-tissue deformations based on previous deformations, the velocity and acceleration of the body, and the angular velocities and accelerations of the limbs. Dyna also models how deformations vary with a person's body mass index (BMI), producing different deformations for people with different shapes. Dyna realistically represents the dynamics of soft tissue for previously unseen subjects and motions. We provide tools for animators to modify the deformations and apply them to new stylized characters.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: Human animation, motion capture, human shape, soft-tissue motion

*e-mail: {gerard.pons.moll,jromero,nmahmood,black}@tue.mpg.de

ACM Reference Format

Pons-Moll, G., Romero, J., Mahmood, N., Black, M. 2015. Dyna: A Model of Dynamic Human Shape in Motion. *ACM Trans. Graph.* 34, 4, Article 120 (August 2015), 14 pages. DOI = 10.1145/2766993 <http://doi.acm.org/10.1145/2766993>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright held by the Owner/Author.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
ACM 978-1-4503-3331-3/15/08.
DOI: <http://dx.doi.org/10.1145/2766993>

1 Introduction

Interest in creating and animating realistic virtual humans has led to a wide variety of models. The most realistic methods model body shape using 3D scans of many bodies and many poses [Allen et al. 2002; Anguelov et al. 2005; Hasler et al. 2009; Hirshberg et al. 2012; Chen et al. 2013]. These models can capture pose-dependent static shape deformations but cannot realistically model the effects of soft-tissue dynamics on bodies in motion. While physics-based models and finite-element modeling provide a possible solution, the complexity of such systems makes them difficult to produce and control. Instead we take a learning-based approach to modeling the deformation of visible surface geometry caused by soft-tissue dynamics. Previous learning-based methods rely on the motion capture of many markers attached to the body [Park and Hodgins 2006; Park and Hodgins 2008]. Markers provide limited spatial resolution, are time consuming to apply, may change the motion we want to observe, and large marker sets pose technical difficulties for marker identification and tracking. Consequently these approaches do not easily scale to capture detailed soft-tissue deformations on a wide variety of body shapes. What is needed is a method to capture surface deformations of the body at high spatial and temporal resolutions and a mathematical model relating these deformations to the motion and body shapes of novel characters.

To that end, we propose a new model called *Dyna* that is learned from examples and is able to produce realistic soft-tissue motions for a wide range of body shapes and motions as illustrated in Fig. 1. Dyna is an extension of the SCAPE model [Anguelov et al. 2005] to include full body deformations that are driven by the motion of the body. While SCAPE approximates static surface deformations of soft tissue, Dyna approximates dynamic surface deformations related to the motion of soft tissue. At the core, Dyna is a mathematical model that relates deformations of the body surface to changing poses of the body in time. To make this feasible, we learn a low-dimensional model of these deformations from 3D scans containing soft tissue in motion. The scans are recorded by a 4D full-body capture system that outputs high-resolution 3D meshes of the body at 60 fps, revealing detailed soft-tissue motions.

Adaptive Rendering with Linear Predictions

Bochang Moon¹, Jose A. Iglesias-Guitian¹, Sung-Eui Yoon², Kenny Mitchell¹

¹Disney Research Zürich, ²KAIST




Figure 1: Our adaptive rendering result in the Courtyard scene. Our method with 15 samples per pixel (spp) produces a high-quality reconstruction result and drastically reduces a relative mean squared error (rMSE) compared to a straight-forward method utilizing low discrepancy (LD) sampling patterns uniformly.

Abstract

We propose a new adaptive rendering algorithm that enhances the performance of Monte Carlo ray tracing by reducing the noise, i.e., variance, while preserving a variety of high-frequency edges in rendered images through a novel prediction based reconstruction. To achieve our goal, we iteratively build multiple, but sparse linear models. Each linear model has its prediction window, where the linear model predicts the unknown ground truth image that can be generated with an infinite number of samples. Our method recursively estimates prediction errors introduced by linear predictions performed with different prediction windows, and selects an optimal prediction window minimizing the error for each linear model. Since each linear model predicts multiple pixels within its optimal prediction interval, we can construct our linear models only at a sparse set of pixels in the image screen. Predicting multiple pixels with a single linear model poses technical challenges, related to deriving error analysis for regions rather than pixels, and has not been addressed in the field. We address these technical challenges, and our method with robust error analysis leads to a drastically reduced reconstruction time even with higher rendering quality, compared to state-of-the-art adaptive methods. We have demonstrated that our method outperforms previous methods numerically and visually with high performance ray tracing kernels such as OptiX and Embree.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

Keywords: Adaptive rendering, image-space reconstruction, Monte Carlo ray tracing

1 Introduction

Monte Carlo (MC) ray tracing [Kajiya 1986] has received extensive attention for synthesizing realistic rendering results, but generally requires a huge number of ray samples (e.g., more than ten thousands samples per pixel) until a converged or even visually pleasing image is generated. Unfortunately, its slow convergence directly leads to prohibitive rendering time (e.g., hours), which is often proportional to the number of ray samples generated. When a relatively small number of ray samples (e.g., less than one hundred) per pixel are allocated, images are typically corrupted by MC noise, i.e., variance.

Adaptive rendering that adjusts sampling density non-uniformly and applies smoothing locally has been actively studied recently, as the approach significantly boosts MC ray tracing by reducing the required number of ray samples drastically [Hachisuka et al. 2008; Overbeck et al. 2009]. These methods can be classified into two categories, multi-dimensional and image space adaptive rendering in terms of the dimensionality of MC samples [Moon et al. 2014]. The multi-dimensional methods [Hachisuka et al. 2008; Lehtinen et al. 2012] allocate samples and reconstruct them in a high dimensional space, where each coordinate corresponds to a random parameter in the MC integration [Kajiya 1986]. These methods can produce a high quality image even with a small number of samples (e.g., 8 samples per pixel), but managing individual samples unfortunately requires high computational and memory overhead.

On the other hand, image space methods [Rousselle et al. 2012; Li et al. 2012; Moon et al. 2014] utilize per-pixel information (e.g., colors, variances, and G-buffers) that can be easily obtained in rendering, and thus these techniques can be easily applied into existing rendering frameworks. The state-of-the-art methods (e.g., [Rousselle et al. 2012; Li et al. 2012; Moon et al. 2014]) have been shown to improve the performance of MC ray tracing by an order of magnitude. Their main target applications, however, are often limited to offline rendering frameworks, since its computational overhead is relatively large. For example, the reconstruction times of the previous methods [Rousselle et al. 2012; Li et al. 2012; Moon et al. 2014] are more than 3 s given the Courtyard scene (Fig. 1) due to their expensive reconstructions (e.g., non-local means and local regression). Especially, the recent local linear approximation [Moon et al. 2014] shows a superior reconstruction performance when a reference image has a strong linear correlation with given features (e.g., textures), but it has very expensive reconstruction time (e.g.,

ACM Reference Format

Moon, B., Iglesias-Guitian, J., Yoon, S., Mitchell, K. 2015. Adaptive Rendering with Linear Predictions. ACM Trans. Graph. 34, 4, Article 121 (August 2015), 11 pages. DOI = 10.1145/2766992
<http://doi.acm.org/10.1145/2766992>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766992>


A Machine Learning Approach for Filtering Monte Carlo Noise

Nima Khademi Kalantari

Steve Bako

Pradeep Sen

University of California, Santa Barbara



Our result with a cross-bilateral filter (4 spp)

Our result with a non-local means filter (4 spp)

Figure 1: We propose a machine learning approach to filter Monte Carlo rendering noise as a post-process. In our method, we use a set of scenes with a variety of distributed effects to train a neural network to output correct filter parameters. We then use the trained network to drive a filter to denoise a new MC rendered image. We show the result of our approach with a cross-bilateral filter for the KITCHEN scene (1200×800) on the left and with a non-local means filter for the SAN MIGUEL HALLWAY scene (800×1200) on the right. Both of these scenes are path-traced and contain severe noise at 4 samples per pixel (spp). However, our trained network is able to estimate the appropriate filter parameters and effectively reduce the noise in only a few seconds. Note, the tonemapping of the insets has been adjusted for best visibility. Scene credits: KITCHEN – Jo Ann Elliott; SAN MIGUEL HALLWAY – Guillermo M. Leal Llaguno.

Abstract

The most successful approaches for filtering Monte Carlo noise use feature-based filters (e.g., cross-bilateral and cross non-local means filters) that exploit additional scene features such as world positions and shading normals. However, their main challenge is finding the optimal weights for each feature in the filter to reduce noise but preserve scene detail. In this paper, we observe there is a complex relationship between the noisy scene data and the ideal filter parameters, and propose to *learn* this relationship using a nonlinear regression model. To do this, we use a multilayer perceptron neural network and combine it with a matching filter during both training and testing. To use our framework, we first train it in an offline process on a set of noisy images of scenes with a variety of distributed effects. Then at run-time, the trained network can be used to drive the filter parameters for new scenes to produce filtered images that approximate the ground truth. We demonstrate that our trained network can generate filtered images in only a few seconds that are superior to previous approaches on a wide range of distributed effects such as depth of field, motion blur, area lighting, glossy reflections, and global illumination.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

Keywords: Monte Carlo rendering, neural networks

1 Introduction

Producing photorealistic images from a scene model requires computing a complex multidimensional integral of the scene function at every pixel of the image. For example, generating effects like depth of field and motion blur requires integrating over domains such as lens position and time. Monte Carlo (MC) rendering systems approximate this integral by tracing light rays (samples) in the multidimensional space to evaluate the scene function. Although an approximation to this integral can be quickly evaluated with just a few samples, the inaccuracy of this estimate relative to the true value appears as unacceptable noise in the resulting image. Since the variance of the MC estimator decreases linearly with the number of samples, many samples are required to get a reliable estimate of the integral. The high cost of computing additional rays results in lengthy render times that negatively affect the applicability of MC renderers in modern film production.

One way to mitigate this problem is to quickly render a noisy image with a few samples and then filter it as a post-process to generate an acceptable, noise-free result. This approach has been the subject of extensive research in recent years [Dammertz et al. 2010; Bauszat et al. 2011; Rousselle et al. 2012; Sen and Darabi 2012; Li et al. 2012; Rousselle et al. 2013; Moon et al. 2014]. The more successful methods typically use feature-based filters (e.g., cross-bilateral or cross non-local means filters) to leverage additional scene features such as world positions that help guide the filtering process. Since these features are highly correlated with scene detail, using them in the filtering process greatly improves the quality of the results.

Some approaches have used this information to handle specific distributed effects such as global illumination [Dammertz et al. 2010; Bauszat et al. 2011] and depth of field [Chen et al. 2011]. However, the major challenge is how to exploit this additional information to denoise *general* distributed effects, which requires setting the filter weights for all features (called *filter parameters* hereafter) so that noise is removed while scene detail is preserved. To do this, Sen and Darabi [2011; 2012] proposed to use the functional dependencies between scene features and random parameters calculated us-

ACM Reference Format

Kalantari, N., Bako, S., Sen, P. 2015. A Machine Learning Approach for Filtering Monte Carlo Noise. ACM Trans. Graph. 34, 4, Article 122 (August 2015), 12 pages. DOI = 10.1145/2766977
<http://doi.acm.org/10.1145/2766977>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766977>

Gradient-Domain Path Tracing

Markus Kettunen¹ Marco Manzi² Miika Aittala¹ Jaakko Lehtinen^{1,3} Frédo Durand⁴ Matthias Zwicker²
¹Aalto University ²University of Bern ³NVIDIA ⁴MIT CSAIL



Figure 1: Comparing gradient-domain path tracing (G-PT, L_1 reconstruction) to path tracing at equal rendering time (2 hours). In this time, G-PT draws about 2,000 samples per pixel and the path tracer about 5,000. G-PT consistently outperforms path tracing, with the rare exception of some highly specular objects. Our frequency analysis explains why G-PT outperforms conventional path tracing.

Abstract

We introduce gradient-domain rendering for Monte Carlo image synthesis. While previous gradient-domain Metropolis Light Transport sought to distribute more samples in areas of high gradients, we show, in contrast, that estimating image gradients is also possible using standard (non-Metropolis) Monte Carlo algorithms, and furthermore, that even without changing the sample distribution, this often leads to significant error reduction. This broadens the applicability of gradient rendering considerably. To gain insight into the conditions under which gradient-domain sampling is beneficial, we present a frequency analysis that compares Monte Carlo sampling of gradients followed by Poisson reconstruction to traditional Monte Carlo sampling. Finally, we describe Gradient-Domain Path Tracing (G-PT), a relatively simple modification of the standard path tracing algorithm that can yield far superior results.

CR Categories: [Computer Graphics]: Rendering

Keywords: path tracing, gradient-domain, global illumination, light transport

1 Introduction

Global illumination algorithms seek to estimate the value of each image pixel defined as a complex integral over the space of light paths, relying on large numbers of Monte-Carlo samples to avoid noise. Recent advances [Lehtinen et al. 2013; Manzi et al. 2014]

take advantage of Metropolis sampling to distribute samples according not only to pixel values (or path throughput), but also finite-difference image gradients. This reduces costs by focusing computing resources onto areas of high variations. Metropolis light transport is notoriously hard to implement, however, and to the best of our knowledge, Mitsuba [Jakob 2012] is the only publicly-available implementation of Veach’s original algorithm [Veach and Guibas 1997]. Furthermore, the convergence behavior of Metropolis algorithms is often considered undesirable because they explore path space in a highly non-uniform fashion. As a consequence, error appears as entire regions being too dark or too bright, instead of as noise. This is even more challenging for gradient sampling, because the areas of high contribution are more sparse. In short, while Metropolis light transport remains unequaled for challenging light transport configurations, many practical scenarios can be more easily handled by standard approaches such as path tracing.

We show, that, maybe surprisingly, we can benefit from rendering image gradients rather than only pixel values also in standard path tracing, even without changing sample distributions (Figure 1). We extend path tracing and shoot, for each base path, additional finite-difference offset paths shifted by one pixel. This provides us not only with the image contribution for each path, but also an estimate of the finite-difference between the two pixels. In the end, we perform a screened Poisson reconstruction that combines information from both sampled pixel values (especially useful for low frequencies) and gradients (much more accurate for high frequencies).

Crucially, our gradient estimates provide much more information than finite differences of a conventionally sampled image. In both cases, a gradient is the difference of two pixels. But rather than using the difference between two sums of uncorrelated random paths, we estimate gradients by integrating the difference between pairs of paths that are carefully sampled in a highly correlated fashion. For this, we design shift mappings that generate path pairs that are “similar” so that the differences are small and result in lower variance, in a manner similar to variance reduction using correlated sampling or common random numbers in statistics.

We analytically study why gradient rendering can yield better quality than direct value rendering. A key insight is that the symmetric operations of finite differencing and Poisson image reconstruction do not cancel each other in the error of gradient rendering. This

ACM Reference Format
Kettunen, M., Manzi, M., Aittala, M., Lehtinen, J., Durand, F., Zwicker, M. 2015. Gradient-Domain Path Tracing. *ACM Trans. Graph.* 34, 4, Article 123 (August 2015), 13 pages. DOI = 10.1145/2766997
<http://doi.acm.org/10.1145/2766997>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766997>

Variance Analysis for Monte Carlo Integration

[†]Adrien Pilleboue^{1,*} [†]Gurprit Singh^{1,*} David Coeurjolly^{2,*} Michael Kazhdan^{3,*} Victor Ostromoukhov^{1,2,*}
[†]joint first authors ¹Université Lyon 1 ²CNRS/LIRIS UMR 5205 ³Johns Hopkins University

Abstract

We propose a new spectral analysis of the variance in Monte Carlo integration, expressed in terms of the power spectra of the sampling pattern and the integrand involved. We build our framework in the Euclidean space using Fourier tools and on the sphere using spherical harmonics. We further provide a theoretical background that explains how our spherical framework can be extended to the hemi-spherical domain. We use our framework to estimate the variance convergence rate of different state-of-the-art sampling patterns in both the Euclidean and spherical domains, as the number of samples increases. Furthermore, we formulate design principles for constructing sampling methods that can be tailored according to available resources. We validate our theoretical framework by performing numerical integration over several integrands sampled using different sampling patterns.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Display Algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Monte Carlo Integration;

Keywords: Stochastic Sampling, Monte Carlo Integration, Fourier Analysis, Spherical Harmonics, Global Illumination

1 Introduction

Numerical integration schemes such as Monte Carlo methods are widely used in high quality production rendering (*e.g.*, to estimate per pixel radiance). This estimation is error-prone, as Monte Carlo methods provide an approximation of the true integral. Specifically, the quality of integration depends strongly on the sampling pattern. For example, regular structures in sampling can result in strong aliasing (structured artifacts) in the final rendered image. However, by using appropriately placed stochastically-generated samples, this aliasing in rendered images can be transformed into less objectionable noise [Crow 1977; Dippé and Wold 1985; Cook 1986; Mitchell 1987; Shirley 1991], which is tightly related to variance in Monte Carlo integration.

In this work, we study variance from the sampling perspective and propose a mathematical framework that predicts variance in Monte Carlo integration. Our framework shows that the variance in integration is directly related to the power spectra of the sampling pattern and the integrand under study. We build our mathematical framework in the Euclidean space and later extend it to the

(hemi-)spherical domain, which is crucial in light transport estimation. To our knowledge, this is the first theoretical and computational analysis of variance in Monte Carlo integration on the sphere and the hemisphere. To study variance in integration we perform spectral analysis of various sampling patterns in both the Euclidean and the (hemi-)spherical domains, using Fourier analysis for Euclidean space and spherical harmonics for the (hemi-)sphere.

1.1 Contributions

Our primary contribution is a mathematical framework that establishes a direct relation between the variance in Monte Carlo integration in the Euclidean and (hemi-)spherical domains with the power spectra of the sampling pattern and the integrand. We use our mathematical framework to study:

Spectral analysis of sampling patterns We perform spectral analysis of various state-of-the-art sampling patterns on Euclidean and spherical domains. We provide theoretical insights to perform similar analysis on hemispherical domain using spherical harmonics.

Monte Carlo convergence We extend our theoretical framework to analyze the best- and worst-case variance convergence rates—for a given class of functions—in Monte Carlo integration in the Euclidean and the (hemi-)spherical domains.

The remainder of this paper is split into two main parts. In the first part, we study the relation between the error in integration and the frequency content of the integrand and the sampling pattern. We define the notion of *homogeneous sampling* (Sec. 3.1) that allows expressing error only in terms of variance in MC integration. By restricting our analysis to homogeneous sampling, we derive the closed-form expression for variance in MC integration in terms of power spectra of the integrand and the sampling pattern. We show that this relation holds in the Euclidean (Sec. 4.2) and the spherical (Sec. 5.2) domains. All mathematical notations and symbols used in our mathematical formulation can be found in Table 1. In the second part, we directly use our variance formulation to study the MC integration convergence rates of various theoretical sampling patterns with well-defined spectral profiles (Sec. 7). We then use these theoretical convergence rates to build a taxonomy to classify existing sampling methods. We show empirical results (Sec. 8) to support our theoretical study and present future directions (Sec. 9) for further exploration.

2 Related work

Sampling in Euclidean domain To improve the quality of image synthesis, researchers have studied various Monte Carlo [Crow 1977; Dippé and Wold 1985; Cook 1986; Mitchell 1987] and Quasi-Monte Carlo based [Niederreiter 1992; Lemieux 2009; Keller et al. 2012] sampling patterns. Many algorithms have been proposed to improve the quality of sampling by studying their spectral properties. Ulichney [1987] was the first to provide qualitative characterization of a good sampling pattern, which is now commonly called *Blue Noise*. Mitchell [1991] has also pointed out that energy in low-frequency part of the Fourier spectrum of the sampling pattern should be avoided. Since then, different algorithms [Balzer et al. 2009; Schlömer et al. 2011; de Goes et al. 2012] and tile-based

* firstname.lastname@liris.cnrs.fr

* misha@cs.jhu.edu

ACM Reference Format

Pilleboue, A., Singh, G., Coeurjolly, D., Kazhdan, M., Ostromoukhov, V. 2015. Variance Analysis for Monte Carlo Integration. ACM Trans. Graph. 34, 4, Article 124 (August 2015), 14 pages. DOI = 10.1145/2766930 <http://doi.acm.org/10.1145/2766930>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766930>

Single-View Hair Modeling Using A Hairstyle Database

Liwen Hu*

Chongyang Ma*

Linjie Luo†

Hao Li*

* University of Southern California

† Adobe Research




Figure 1: Our system takes as input a reference photo (a), a few user strokes (b) and a database of example hairstyles (c) to model the 3D target hairstyle (e). The retrieved examples that best match the user strokes are highlighted with the corresponding colors in (c) which are combined consistently as shown in (d). Original image courtesy of Yung-Yuan Kao.

Abstract

Human hair presents highly convoluted structures and spans an extraordinarily wide range of hairstyles, which is essential for the digitization of compelling virtual avatars but also one of the most challenging to create. Cutting-edge hair modeling techniques typically rely on expensive capture devices and significant manual labor. We introduce a novel data-driven framework that can digitize complete and highly complex 3D hairstyles from a single-view photograph. We first construct a large database of manually crafted hair models from several online repositories. Given a reference photo of the target hairstyle and a few user strokes as guidance, we automatically search for multiple best matching examples from the database and combine them consistently into a single hairstyle to form the large-scale structure of the hair model. We then synthesize the final hair strands by jointly optimizing for the projected 2D similarity to the reference photo, the physical plausibility of each strand, as well as the local orientation coherency between neighboring strands. We demonstrate the effectiveness and robustness of our method on a variety of hairstyles and challenging images, and compare our system with state-of-the-art hair modeling algorithms.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations;

ACM Reference Format

Hu, L., Ma, C., Luo, L., Li, H. 2015. Single-View Hair Modeling Using A Hairstyle Database. ACM Trans. Graph. 34, 4, Article 125 (August 2015), 9 pages. DOI = 10.1145/2766931 <http://doi.acm.org/10.1145/2766931>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766931>

Keywords: hairstyle database, data-driven modeling, sketch-based retrieval, structure-aware shape processing, piecewise helices

1 Introduction

With the dramatic performance increase of today's graphics technologies, visual details of digital humans in games, online virtual worlds, and virtual reality applications are becoming significantly more demanding, reaching nearly the quality and realism seen in film production. While compelling animations and renderings are now possible in realtime, the digitization of lifelike virtual avatars is still reserved to professional production studios, involving complex hardware equipment and talented digital artists. With the aim of scaling production and bringing 3D character creation to consumers, significant research has been dedicated to the automatic digitization of human faces [Blaiz and Vetter 1999], bodies [Zhou et al. 2010], and hands [Wang and Popović 2009] from a single input image or video. These 3D reconstruction methods typically rely on low dimensional parametric models that encode statistical shape variations of a population.

An essential but often overlooked problem is the modeling of human hair, which is challenged by the highly intricate structure of intertwined strands and the wide variation of hairstyles. Unlike anatomically compatible shapes such as faces and bodies, it is unclear how to represent or construct a unified parametric model for hair due to the vast diversity in topological structures and local details. Nevertheless, several 3D hair capture techniques have recently demonstrated the successful reconstruction of challenging hairstyles [Paris et al. 2008; Luo et al. 2013], but they typically rely on complex capture settings. Although single-view hair modeling methods [Chai et al. 2012; Chai et al. 2013] have been introduced lately for image-based rendering and editing tasks, they are only capable of handling very coarse and smooth hairstyles.

Recent advances in data-driven modeling have shown that many ill-posed single-view reconstruction problems can be addressed by leveraging geometric priors from a large collection of 3D shapes (e.g., for hand tracking [Wang and Popović 2009], pose estimation

SecondSkin: Sketch-based Construction of Layered 3D Models

Chris De Paoli Karan Singh
University of Toronto



Figure 1: 2D strokes sketched on and around 3D geometry form the input to SecondSkin (a). Layered structures are represented as solid models with volumes bounded by surface patches and curves (b). A majority (91%) of sketch strokes are perceived by viewers as one of four curve-types (c). We automatically classify these strokes based on the relationship between 2D strokes and underlying 3D geometry, producing 3D curves, surface patches, and volumes (d), resulting in layered 3D models suitable for conceptual design (e).

Abstract

SecondSkin is a sketch-based modeling system focused on the creation of structures comprised of layered, shape interdependent 3D volumes. Our approach is built on three novel insights gleaned from an analysis of representative artist sketches. First, we observe that a closed loop of strokes typically define surface patches that bound volumes in conjunction with underlying surfaces. Second, a significant majority of these strokes map to a small set of curve-types, that describe the 3D geometric relationship between the stroke and underlying layer geometry. Third, we find that a few simple geometric features allow us to consistently classify 2D strokes to our proposed set of 3D curve-types. Our algorithm thus processes strokes as they are drawn, identifies their curve-type, and interprets them as 3D curves *on* and *around* underlying 3D geometry, using other connected 3D curves for context. Curve loops are automatically surfaced and turned into volumes bound to the underlying layer, creating additional curves and surfaces as necessary. Stroke classification by 15 viewers on a suite of ground truth sketches validates our curve-types and classification algorithm. We evaluate SecondSkin via a compelling gallery of layered 3D models that would be tedious to produce using current sketch modelers.

ACM Reference Format

De Paoli, C., Singh, K. 2015. SecondSkin: Sketch-Based Construction of Layered 3D Models. ACM Trans. Graph. 34, 4, Article 126 (August 2015), 10 pages. DOI = 10.1145/2766948 <http://doi.acm.org/10.1145/2766948>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766948>

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, and systems

Keywords: sketch-based modeling, layers, shells

1 Introduction

Current 3D conceptual design tools, regardless of being based on metaphors of sketching and sculpting or traditional CAD modeling, typically focus on the creation of the *skin* or visible surface of 3D objects. Physical objects, both organic and man-made, however, are often layered assemblies: comprising parts segmented by form, function or material, built over each other (Figure 1, 2 and 13). While research in character skinning and animation has noted the importance of conceptual anatomic layers for a quarter century now [Chadwick et al. 1989], 3D conceptual design tools, to date, have largely ignored what lies beneath the skin. *SecondSkin* addresses this problem: *the fluid sketch-based creation of layered 3D structures*.

A defining aspect of layered modeling is the geometric dependence of layers on underlying layers. This is clearly evidenced in a multitude of books and tutorials on sketching and concept art [Davies and Scott 2012], where maquettes of underlying layers are used as a visual reference **on** and **around** which to draw subsequent layers (Figure 2). Prior work in sketch-based modeling [Kara and Shimada 2007; Nealen et al. 2007; Takayama et al. 2013], typically interprets sketched strokes as lying **on** template objects or the evolving 3D geometric skin of the object. In the context of layered modeling however, we expect that sketched strokes are largely drawn **around** underlying template objects, to build new layered structures.

While projecting a 2D stroke drawn from a given view **on** to 3D geometry is mathematically precise and straightforward, inferring a 3D curve from such a 2D stroke to lie **around** 3D geometry, is generally

Flow Aligned Surfacing of Curve Networks

Hao Pan * Yang Liu † Alla Sheffer ‡ Nicholas Vining ‡ Chang-Jian Li * Wenping Wang *
*The University of Hong Kong †Microsoft Research ‡University of British Columbia




Figure 1: Surfacing an artist created curve network: (a) an input curve network of the coffee maker model; (b) initial surface; (c) the computed flow field aligned with network flow lines denoted by dashed white lines on the initial surface; (d) final surface whose curvature directions (scaled by principal curvatures) (e) are well aligned with the representative flow lines in the curve network.

Abstract

We propose a new approach for automatic surfacing of 3D curve networks, a long standing computer graphics problem which has garnered new attention with the emergence of sketch based modeling systems capable of producing such networks. Our approach is motivated by recent studies suggesting that artist-designed curve networks consist of descriptive curves that convey intrinsic shape properties, and are dominated by *representative flow lines* designed to convey the principal curvature lines on the surface. Studies indicate that viewers complete the intended surface shape by envisioning a surface whose curvature lines smoothly blend these flow-line curves. Following these observations we design a surfacing framework that automatically aligns the curvature lines of the constructed surface with the representative flow lines and smoothly interpolates these representative flow, or curvature directions while minimizing undesired curvature variation. Starting with an initial triangle mesh of the network, we dynamically adapt the mesh to maximize the agreement between the principal curvature direction field on the surface and a smooth *flow field* suggested by the representative flow-line curves. Our main technical contribution is a framework for curvature-based surface modeling, that facilitates the creation of surfaces with prescribed curvature characteristics. We validate our method via visual inspection, via comparison to artist created and ground truth surfaces, as well as comparison to prior art, and confirm that our results are well aligned with the computed flow fields and with viewer perception of the input networks.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems;

Keywords: curve networks, surfacing, flow lines

1 Introduction

Surfacing of 3D curve networks, also known as *lofting* or *skinning*, is a fundamental problem in geometric modeling. Original methods focused on surfacing networks specifically designed for this task, and concentrated on achieving a desired degree of continuity across network curves and vertices [Farin and Hansford 1999]. More recent sketch-based modeling practice focuses on creating curve networks motivated by product design practices [Kara and Shimada 2007; Bordegoni and Rizzi 2011; Bae et al. 2008; Xu et al. 2014]. Recent research affirms that such artist-created 3D curve networks effectively convey uniquely imagined complex 3D shapes [Xu et al. 2014; Bessmeltsev et al. 2012]. Our work addresses the algorithmic creation of these artist-intended, imagined surfaces, given the artist designed curve network. Our framework conceptually differs from previous work in this area, e.g. [Bessmeltsev et al. 2012], in its use of a continuous geometric approach for surface fitting which, as we demonstrate, results in output surfaces better aligned with artist intent.

In our work we assume that the output surface connectivity, i.e. the set of curve cycles to be surfaced, is detected *a priori*, e.g. using the methods of [Abbasinejad et al. 2011; Zhuang et al. 2013]. Our focus instead is on computing the desired surface shape or geometry. In doing so, we are motivated by emerging research into human perception of 3D curve networks as well as observation of artistic practices in this domain [Bordegoni and Rizzi 2011; Bessmeltsev et al. 2012]. This research indicates that artist-drawn curve networks are dominated by *representative flow-line* curves [Gahan 2010; Singh et al. 2004; Xu et al. 2014], understood to largely align with lines of curvature, but which allow for artistic license at surface discontinuities, over fine details, and in umbilic regions. Perception studies suggest that the curvature lines on the viewer imagined surfaces are perceived as a blend of the flow-line segments on the designer-created curve cycles. In addition to flow lines, artists employ *trim* curves to demarcate sharp features or discontinuous transitions between surface patches. Artists frequently employ curves that serve a dual role, as trim curves for one of the attached surfaces and as a flow line on the other. Viewers leverage perceptual cues to distinguish between these curves [Xu et al. 2014].

Motivated by these observations, our framework surfaces the curve network so as to align the principal curvature directions on the surface with the representative flow lines in the artist-created network, while interpolating the trimming curves. It automatically detects the artist created feature curves, and generates smooth continuous surfaces across the rest of the network curves (Fig. 1(d)). Our method

ACM Reference Format

Pan, H., Liu, Y., Sheffer, A., Vining, N., Li, C., Wang, W. 2015. Flow Aligned Surfacing of Curve Networks. ACM Trans. Graph. 34, 4, Article 127 (August 2015), 10 pages. DOI = 10.1145/2766990
<http://doi.acm.org/10.1145/2766990>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766990>

Topology-Constrained Surface Reconstruction From Cross-sections

Ming Zou*
Washington U. in St. Louis

Michelle Holloway†
Washington U. in St. Louis

Nathan Carr‡
Adobe

Tao Ju§
Washington U. in St. Louis




Figure 1: Non-parallel cross-section curves delineating a developing chicken heart from a CT volume (a) and the reconstructed surfaces using the method of Lu et al. [2008] (b), method of Bermano et al. [2011] (c), and our method with genus-1 constraint without utilizing the CT volume (d). The first two surfaces contain numerous topological tunnels, while ours correctly captures the shell-like shape of the object.

Abstract

In this work we detail the first algorithm that provides topological control during surface reconstruction from an input set of planar cross-sections. Our work has broad application in a number of fields including surface modeling and biomedical image analysis, where surfaces of known topology must be recovered. Given curves on arbitrarily oriented cross-sections, our method produces a manifold interpolating surface that exactly matches a user-specified genus. The key insight behind our approach is to formulate the topological search as a divide-and-conquer optimization process which scores local sets of topologies and combines them to satisfy the global topology constraint. We further extend our method to allow image data to guide the topological search, achieving even better results than relying on the curves alone. By simultaneously satisfying both geometric and topological constraints, we are able to produce accurate reconstructions with fewer input cross-sections, hence reducing the manual time needed to extract the desired shape.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: Surface reconstruction, cross-section interpolation, contour stitching, topology constraint, dynamic programming

*e-mail: mingzou.cn@gmail.com

†e-mail: mavaughn@wustl.edu

‡e-mail: ncarr@adobe.com

§e-mail: taoju@cse.wustl.edu

ACM Reference Format

Zou, M., Holloway, M., Carr, N., Ju, T. 2015. Topology-Constrained Surface Reconstruction From Cross-Sections. ACM Trans. Graph. 34, 4, Article 128 (August 2015), 10 pages. DOI = 10.1145/2766976
<http://doi.acm.org/10.1145/2766976>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766976>

1 Introduction

The need to create surfaces depicting a 3D object from its cross-sections arises in many applications. For instance, in order to reconstruct the surface of an organ captured by CT or MRI, radiologists typically start by manually delineating the organ boundary on selected 2D slices of the 3D image volume. The delineated planar curves then need to be connected to form a closed surface. Figure 1 (a) shows an example stack of the delineated curves of a layer of a developing chicken heart from a CT volume (images on the right show two slices through the CT volume).

In many application domains, there exists strong prior information about the object to be depicted. One piece of such information is topology, which describes the object's connected components and their genus. For example, a biological structure usually has a known topology. Most structures have the topology of a sphere, which is a single component with zero genus. Some structures have a more complicated topology. A layer of a developing chicken heart is a cylindrical shell with genus 1. Bones in our body, such as vertebrae and the hip bone, can have a higher genus. Accurately capturing such topology not only helps with recovering the shape of the object but is also important for downstream applications of the reconstructed surface, such as shape matching, fluid dynamics simulation, and mechanical analysis.

Much research has been done for creating surfaces from cross-sections (see a more detailed review in Section 2). While the state-of-the-art algorithms are capable of creating watertight surfaces from even the most complex inputs, no method to date can guarantee that their output has a predefined topology. In fact, it is quite common for these methods to create topological errors, particularly when the cross-sections are not dense (see Figure 1 (b,c)). While it is possible for these methods to achieve the correct topology by providing them with more cross-sections, this will require additional manual labor and the data may not be available at all.

In this paper, we present a novel algorithm for reconstruction from cross-sections that allows precise control over surface topology

MultiFab: A Machine Vision Assisted Platform for Multi-material 3D Printing

Pitchaya Sitthi-Amorn^{†‡} Javier E. Ramos[†] Yuwang Wang^{†§} Joyce Kwan[†]
Justin Lan[†] Wenshou Wang[†] Wojciech Matusik[†]
MIT CSAIL[†] Chulalongkorn University[‡] Tsinghua University[§]



Figure 1: Our multi-material 3D printer (left) and a set of fabricated materials and objects (right).

Abstract

We have developed a multi-material 3D printing platform that is high-resolution, low-cost, and extensible. The key part of our platform is an integrated machine vision system. This system allows for self-calibration of printheads, 3D scanning, and a closed-feedback loop to enable print corrections. The integration of machine vision with 3D printing simplifies the overall platform design and enables new applications such as 3D printing over auxiliary parts. Furthermore, our platform dramatically expands the range of parts that can be 3D printed by simultaneously supporting up to 10 different materials that can interact optically and mechanically. The platform achieves a resolution of at least $40\mu\text{m}$ by utilizing piezoelectric inkjet printheads adapted for 3D printing. The hardware is low cost (less than \$7,000) since it is built exclusively from off-the-shelf components. The architecture is extensible and modular – adding, removing, and exchanging printing modules can be done quickly. We provide a detailed analysis of the system’s performance. We also demonstrate a variety of fabricated multi-material objects.

CR Categories: I.3.8 [Computer Graphics]: Applications;

Keywords: 3D printing, additive manufacturing, multi-material fabrication, machine vision

1 Introduction

Multi-material 3D printing holds the great promise of allowing the automated conversion of 3D models with complex structures, appearances, and properties (e.g., mechanical, electrical, chemical, optical) to physical equivalents. It has the potential to accelerate innovation – engineers and hobbyists will have the power to create objects that have been previously impossible or very difficult to fabricate. It will also give them the opportunity to iterate over their designs inexpensively and quickly. Multi-material 3D printing will also impact the general population by allowing mass customization of personalized products. This technology has a myriad of yet unexplored applications that will inspire future research and stimulate a number of industrial markets. Computer graphics researchers both in academia and in industry share this enthusiasm – they have started contributing significantly to the development of applications and tools for 3D printing [Dong et al. 2010; Bickel et al. 2010; Vidimče et al. 2013; Chen et al. 2013; Tompkin et al. 2013; Skouras et al. 2013].

Unfortunately, current multi-material additive manufacturing systems have severe shortcomings. First, access to this technology is very limited; the price of a multi-material 3D printer is prohibitive – a Stratasys Objet Connex is priced at \$250K, with materials priced at \$500 per kilogram [Stratasys]. This means that very few labs conducting research in computer graphics can afford these systems. The material library is also limited and proprietary, supporting only UV-cured photopolymers (e.g., even full-color printing is not possible). Furthermore, these printers can simultaneously use, at most, only three different materials. Finally, the hardware and software architectures for current multi-material 3D printers are proprietary and inextensible. The exposed input file format, separate STL files for each material, is limiting, e.g., the per layer raster input format is not exposed although it is internally used. Furthermore, any changes to the underlying hardware and software are virtually impossible.

Taking into account these shortcomings, we have developed a 3D printing platform with the following features:

ACM Reference Format

Sitthi-amorn, P., Ramos, J., Wang, Y., Kwan, J., Lan, J., Wang, W., Matusik, W. 2015. MultiFab: A Machine Vision Assisted Platform for Multi-Material 3D Printing. *ACM Trans. Graph.* 34, 4, Article 129 (August 2015), 11 pages. DOI = 10.1145/2766962 <http://doi.acm.org/10.1145/2766962>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766962>

Color Imaging and Pattern Hiding on a Metallic Substrate

Petar Pjanic, Roger D. Hersch

Ecole Polytechnique Fédérale de Lausanne (EPFL)

Abstract

We present a new approach for the reproduction of color images on a metallic substrate that look bright and colorful under specular reflection observation conditions and also look good under non-specular reflection observation conditions. We fit amounts of both the white ink and the classical cyan, magenta and yellow inks according to a formula optimizing the reproduction of colors simultaneously under specular and non-specular observation conditions. In addition, we can hide patterns such as text or graphical symbols in one viewing mode, specular or non-specular, and reveal them in the other viewing mode. We rely on the trade-off between amounts of white diffuse ink and amounts of cyan, magenta and yellow inks to control lightness in specular and in non-specular observation conditions. Further effects are grayscale images that alternate from a first image to a second independent image when tilting the print from specular to non-specular reflection observation conditions. Applications comprise art and entertainment, publicity, posters, as well as document security.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing and display algorithms;

Keywords: Color reproduction, document security, color prints on metal, pattern hiding, image alternations, white diffuse ink, specular versus diffuse reflection.

1 Introduction

Direct prints on a pure metallic substrate provide bright and brilliant colors when seen under specular reflection, but they look dark and dull when viewed under non-specular reflection [Pjanic and Hersch 2013]. Such prints are becoming commercially available for advertisement purposes and for the decoration of private homes.

In the present contribution, we aim at producing partly specular reflecting and partly diffusely reflecting prints that look bright and colorful under specular and also look nice under non-specular observation angles. We rely on the ability to print on top of the metallic substrate diffusely reflecting white ink halftones that reduce the specular reflection component and increase the diffuse reflection component of the print.

In addition, we offer the possibility of hiding patterns within a grayscale or color image when seen under specular viewing angles and make them appear when viewed under non-specular viewing angles or vice versa.

We rely on the fact that increasing the amount of white ink darkens the image in specular viewing mode but increases the lightness of the image in the non-specular viewing mode and that the increase of similar amounts of cyan, magenta and yellow inks darkens the image both in specular and non-specular viewing modes.

Thanks to the trade-off between the amounts of the white and of the colored inks, the lightness can be kept constant in one viewing mode and can be varied in the second viewing mode. This also enables us to produce a metallic print that shows a first grayscale image in one viewing mode, and then, by tilting it to the other viewing mode, shows a second independent grayscale image.

We had to meet a number of challenges. The first challenge consisted in creating a color reproduction workflow that enables finding surface coverages of inks yielding colors similar to the original colors under both specular and non-specular observation conditions. This is achieved by having accurate models predicting the colors of the prints as a function of the amounts of inks under either specular or non-specular observation conditions. Then, with an objective function that minimizes a difference metric between original color and specular color and between original color and non-specular color, one obtains optimal surface coverages of the diffusing white and of the classical cyan, magenta and yellow inks.

Since we would like to hide patterns under specular reflection and reveal them under non-specular reflection or vice-versa, a second challenge consisted in determining the gamut of print colors that for each color in one viewing mode provides a sufficiently large variety of lightnesses in the other viewing mode. For this purpose, a second objective function was created whose minimization provides ink surface coverages for reproducing a desired color in the first viewing mode and at the same time enables creating a modified lightness in the second viewing mode.

In contrast to prior work on appearance reproduction [e.g. Matusik et al. 2009, Lan et al. 2013], we take in respect to angular reflectances a “coarse grain” approach and consider only the two extreme cases, namely specular reflection and non-specular reflection. However in respect to color rendition, we take a “fine grain” approach which is necessary for pattern hiding. We fit surface coverages of the inks by relying on a first spectral prediction model making accurate predictions under the specular viewing conditions ($25^\circ:25^\circ$) and on a second prediction model making accurate predictions under the non-specular viewing conditions ($45^\circ:0^\circ$).

Our novel contributions are the following:

- (i) The idea of the trade-off between amounts of white diffuse ink and amounts of cyan, magenta and yellow inks to control CIELAB colors in specular and in non-specular observation conditions.
- (ii) The simultaneous consideration of two spectral prediction models, one for specular and one for non-specular viewing conditions.
- (iii) The idea that reducing the size of a gamut in one of the two viewing modes offers freedom to have lightness variations in the other viewing mode.
- (iv) Optimization formula and color separations to produce color images that can be viewed both under specular and non-specular viewing conditions.

ACM Reference Format

Pjanic, P., Hersch, R. 2015. Color Imaging and Pattern Hiding on a Metallic Substrate. *ACM Trans. Graph.* 34, 4, Article 130 (August 2015), 10 pages. DOI = 10.1145/2766944 <http://doi.acm.org/10.1145/2766944>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766944>

Computational Hydrographic Printing

Yizhong Zhang*

Chunji Yin*

Changxi Zheng†

Kun Zhou*‡

* State Key Lab of CAD&CG, Zhejiang University

† Columbia University

Abstract

Hydrographic printing is a well-known technique in industry for transferring color inks on a thin film to the surface of a manufactured 3D object. It enables high-quality coloring of object surfaces and works with a wide range of materials, but suffers from the inability to accurately register color texture to complex surface geometries. Thus, it is hardly usable by ordinary users with customized shapes and textures.

We present *computational hydrographic printing*, a new method that inherits the versatility of traditional hydrographic printing, while also enabling precise alignment of surface textures to possibly complex 3D surfaces. In particular, we propose the first computational model for simulating hydrographic printing process. This simulation enables us to compute a color image to feed into our hydrographic system for precise texture registration. We then build a physical hydrographic system upon off-the-shelf hardware, integrating virtual simulation, object calibration and controlled immersion. To overcome the difficulty of handling complex surfaces, we further extend our method to enable multiple immersions, each with a different object orientation, so the combined colors of individual immersions form a desired texture on the object surface. We validate the accuracy of our computational model through physical experiments, and demonstrate the efficacy and robustness of our system using a variety of objects with complex surface textures.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

Keywords: 3D printing, hydrographics, water transfer printing, physical simulation, viscous sheets, texture mapping

1 Introduction

Recent advances in direct digital fabrication — or “3D printing” — promise fast prototyping of objects with almost arbitrary shapes. In parallel, numerous computational methods have been devised for fabricating shapes with desired geometric (e.g., [Chen et al. 2014]), physical (e.g., [Bickel et al. 2010; Stava et al. 2012; Skouras et al. 2013]) and reflectance properties (e.g., [Weyrich et al. 2009; Lan et al. 2013]). Perhaps equally important as the shape fabrication are methods that color the surface of a fabricated piece



Figure 1: Colorful Bunny. We present a new computational approach for hydrographically printing arbitrary color textures on complex 3D surfaces and aligning the textures precisely. The input of our system includes the desired surface texture (upper left) of a physical shape (lower left). Through our hydrographic system, we color the surface of the 3D shape using the desired texture (right).

with a specified texture and work with complex surface geometry made of a wide range of materials. Such a method, unfortunately, has not been explored, and thus is the focus of this paper.

Currently, most of the 3D printed shapes at home are monochromatic, although high-end 3D printers offer the ability to fabricate colored models. These printers often support only a few colors and limited materials (e.g., plastic and sandstone), and their maintenance and printing cost can be unaffordable for many personal uses. In industrial production, there exist many ways for decorating 3D surfaces. These methods, including chemical- and electro-plating [Schlesinger and Paunovic 2011], decals [Hopper 2004], enamels [Darty 2004], and even manual painting, require heavy-duty devices, high operational cost, and skilled personnel; and they can still be limited by complex surface geometries. All these factors render them almost impossible for coloring customized 3D surfaces for personal uses.

In this paper, we propose *computational hydrographic printing*, a computational augmentation of hydrographic printing methods for physically decorating 3D surfaces with user-customized color textures. Our method works for complex surface geometries and a wide range of materials including plastic, wood and porcelain. The whole system is easy to set up for personal uses, and enjoys a low operational cost (less than 40 US cents per printing).

Traditional hydrographic printing. Hydrographic printing (also named hydroprinting or water transfer printing) is a well-known technology for transferring color images to a 3D surface [Wikipedia 2014]. It starts by printing a pixel image on a polyvinyl alcohol (PVA) film (Figure 3) using a conventional inkjet printer. This PVA color film is then put on top of water to carry color inks floating there. After that, an activator chemical is sprayed on the film, softening the color film to make it easy to stretch and also activating a binding agent. Lastly, a substrate object is slowly dipped into the water through the floating film. After

* {yizhongzhang,yinchunji}@zju.edu.cn, kunzhou@acm.org

† cxz@cs.columbia.edu

‡ Corresponding author

ACM Reference Format

Zhang, Y., Yin, C., Zheng, C., Zhou, K. 2015. Computational Hydrographic Printing. ACM Trans. Graph. 34, 4, Article 131 (August 2015), 11 pages. DOI = 10.1145/2766932 <http://doi.acm.org/10.1145/2766932>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766932>

Stable Constrained Dynamics

Maxime Tournier^{4,1,2}

Matthieu Nesme^{1,3}

Benjamin Gilles^{2,1}

François Faure^{5,3,1}

¹ INRIA

² LIRMM-CNRS

³ LJK-CNRS

⁴ RIKEN BSI-BTCC

⁵ Univ. Grenoble



Figure 1: Our method improves stability and step size for the simulation of constraint-based objects subject to high tensile forces, isolated or coupled with other types of objects. Bow: stiff 3D frame, 1D inextensible string, rigid arrow ; Trampoline: soft lateral springs, inextensible textile ; Knee: complex assembly of rigid bodies and stiff unilateral springs ; Ragdoll: rigid body assembly.

Abstract

We present a unification of the two main approaches to simulate deformable solids, namely elasticity and constraints. Elasticity accurately handles soft to moderately stiff objects, but becomes numerically hard as stiffness increases. Constraints efficiently handle high stiffness, but when integrated in time they can suffer from instabilities in the nullspace directions, generating spurious transverse vibrations when pulling hard on thin inextensible objects or articulated rigid bodies. We show that geometric stiffness, the tensor encoding the change of force directions (as opposed to intensities) in response to a change of positions, is the missing piece between the two approaches. This previously neglected stiffness term is easy to implement and dramatically improves the stability of inextensible objects and articulated chains, without adding artificial bending forces. This allows time step increases up to several orders of magnitude using standard linear solvers.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—[Physically based modeling] I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—[Animation]

Keywords: Physically based animation, Simulation, Dynamics, Constraints, Continuum mechanics, Geometric Stiffness

1 Introduction

Constraint-based simulation is very popular for implementing joints in articulated rigid bodies, and to enforce inextensibility in some directions of deformable objects such as cables or cloth. Its mathematical formulation makes it numerically robust to infinite stiffness, contrary to elasticity-based

simulation, and some compliance can be introduced in the formulation or obtained through approximate solutions. Unfortunately, when the constraint forces are large, constraint-based objects are prone to instabilities in the transverse, unconstrained directions. This occurs when pulling hard on inextensible strings and sheets, or on chains of articulated bodies. The spurious vibrations can lead to unrealistic behaviors or even simulation divergence. They can be avoided using small time steps or complex non-linear solvers, however this dramatically slows down the simulation, while many applications, especially in interactive simulation, hardly allow for one linear solution per frame. The simulation speed can only be maintained by relaxing inextensibility, or using implicit elastic bending forces, however this changes the constitutive law of the simulated objects.

In this work, we show how to perform stable and efficient simulations of both extensible and inextensible constraint-based objects subject to high tensile forces. The key to transverse stability lies in the *geometric stiffness*, a first-order approximation of the change of direction of the internal forces due to rotation or bending. Neglecting the geometric stiffness, as usually done in constraint-based simulation, is a simplification of the linearized equation system, which in turn is a simplification of the exact, non-linear implicit integration. In case of thin objects, this leaves the transverse directions unconstrained, leading to uncontrolled extensions after time integration, introducing artificial potential energy. While this is acceptable for small stiffnesses or short time steps, this may introduce instabilities in the other cases. In this paper, we show that solving the complete linear equation allows high stiffnesses and large time steps which were only achievable using much slower non-linear solvers before. We show how to handle the geometric stiffness in a numerically stable way, even for very large material stiffness. The implementation is easy to combine with existing implicit solvers, and can provide several orders of magnitude speed-ups. Moreover, it allows a unification of rigid body and continuum mechanics.

In the next section, we detail our background and motivation through an introductory example. The principle of our method is then explained in Section 3. Its application to a wide variety of cases is then presented in Section 4. We conclude and sketch future work in Section 5.

ACM Reference Format

Tournier, M., Nesme, M., Gilles, B., Faure, F. 2015. Stable Constrained Dynamics. ACM Trans. Graph. 34, 4, Article 132 (August 2015), 10 pages. DOI = 10.1145/2766969 <http://doi.acm.org/10.1145/2766969>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766969>

Air Meshes for Robust Collision Handling

Matthias Müller

Nuttapong Chentanez
NVIDIA

Tae-Yong Kim

Miles Macklin



Figure 1: A dancer with a multi-layered skirt. Our method robustly simulates the complex interaction of the layers and is able to smoothly recover from any entangled state.

Abstract

We propose a new method for both collision detection and collision response geared towards handling complex deformable objects in close contact. Our method does not miss collision events between time steps and solves the challenging problem of untangling automatically and robustly. It is conceptually simple and straight forward to parallelize due to the regularity of the algorithm.

The main idea is to tessellate the air between objects once before the simulation and by considering one unilateral constraint per element that prevents its inversion during the simulation. If large relative rotations and translations are present in the simulation, an additional dynamic mesh optimization step is needed to prevent mesh locking. This step is fast in 2D and allows the simulation of arbitrary scenes. Because mesh optimization is expensive in 3D, however, the method is best suited for the subclass of 3D scenarios in which relative motions are limited. This subclass contains two important problems, namely the simulation of multi-layered clothing and tissue on animated characters.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically Based Modeling; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation and Virtual Reality

Keywords: collision detection, mesh optimization

ACM Reference Format

Mueller-Fischer, M., Chentanez, N., Kim, T., Macklin, M. 2015. Air Meshes for Robust Collision Handling. ACM Trans. Graph. 34, 4, Article 133 (August 2015), 9 pages. DOI = 10.1145/2766907
<http://doi.acm.org/10.1145/2766907>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766907>

1 Introduction

Both collision detection and collision handling are challenging problems and active areas of research in computer graphics [Teschner et al. 2005], [Yoon et al. 2010]. Existing methods for collision detection typically use spatial acceleration structures such as uniform grids or bounding volume hierarchies to identify overlapping primitives. For thin or fast moving objects, the swept volume over a time step has to be considered, rather than the static volume of the objects not to miss any collision events. For complex objects in close proximity such as layers of clothing, this step is expensive and likely the bottleneck of the simulation.

Once pairs of overlapping primitives have been identified, appropriate response forces have to be applied in order to resolve the collisions. For thin structures such as cloth or deeply penetrating volumetric objects, finding the correct directions of the response forces is challenging. Properly recovering from an entangled state is a global problem so instead of analyzing the collision pairs individually, entire objects have to be considered [Wicke et al. 2006]. One way to solve this problem is to avoid entangled states all together via exact collision handling [Brochu et al. 2012]. These methods tend to be computationally expensive. Also, kinematic objects such as characters often force clothing into penetrating configuration so avoidance is not possible in those cases [Baraff et al. 2003].

In this project we were specifically interested in collision handling for characters with simulated multi-layered clothing, simulated tissue and the interaction of the two (see Figure 1). In this scenario it is often the case that entangled states are unavoidable at certain points in time. Examples are when a character sits or lays down or when the clothing is trapped below the armpits. Also, entangled states may result from the fact that the time for resolving collisions is limited. Typically, temporary collisions are tolerable as long as they are resolved over time. On the other hand, persistent entangled states are unacceptable.

The most important feature of our new method is therefore its ability to recover smoothly from any entangled state. This is a key piece

Aerophones in Flatland: Interactive Wave Simulation of Wind Instruments

Andrew Allen Nikunj Raghuvanshi
Microsoft Research




Figure 1: Wave fields for 2D wind instruments simulated in real-time on a graphics card. A few examples are shown, which are simplified virtual models of (a) trumpet, (b) clarinet, and (c) flute. Our interactive wave solver lets the user design and instantly perform such virtual instruments, promoting experimentation with novel designs. Dynamic changes such as opening and closing tone holes or manipulating valves automatically changes the resulting sound and radiation pattern. Synthesized musical notes can be heard in the accompanying demonstrations.

Abstract

We present the first real-time technique to synthesize full-bandwidth sounds for 2D virtual wind instruments. A novel interactive wave solver is proposed that synthesizes audio at 128,000Hz on commodity graphics cards. Simulating the wave equation captures the resonant and radiative properties of the instrument body automatically. We show that a variety of existing non-linear excitation mechanisms such as reed or lips can be successfully coupled to the instrument's 2D wave field. Virtual musical performances can be created by mapping user inputs to control geometric features of the instrument body, such as tone holes, and modifying parameters of the excitation model, such as blowing pressure. Field visualizations are also produced. Our technique promotes experimentation by providing instant audio-visual feedback from interactive virtual designs. To allow artifact-free audio despite dynamic geometric modification, we present a novel time-varying Perfectly Matched Layer formulation that yields smooth, natural-sounding transitions between notes. We find that visco-thermal wall losses are crucial for musical sound in 2D simulations and propose a practical approximation. Weak non-linearity at high amplitudes is incorporated to improve the sound quality of brass instruments.

CR Categories: H.5.5 [Information Interfaces and Presentation]: Sound and music computing—modeling; G.1.8 [Numerical Analysis]: Partial differential equations—finite difference methods; I.3.1 [Computer Graphics]: Hardware architecture—graphics processors

Keywords: wind instruments, wave equation, radiation, scattering, graphics processor (GPU), sound synthesis

1 Introduction

Making wind instruments (aerophones) is an innate human activity dating back at least 20,000 years [Baines 1967]. Wind instruments are non-linear dynamical systems with an active excitation mecha-

nism, such as a trumpet player's buzzing lips, undergoing coupled oscillation with the resonant cavity formed by the body of the instrument. This two-way coupling is essential to their operation, irreducible to a feed-forward model. For instance, the sound of oscillating lips filtered through a trumpet's resonant acoustic response results in a comb-filtered buzzing sound, not a steady musical note. The complex physics underlying their behavior has naturally invited enduring curiosity from physicists [Helmholtz 1885]. Ever since the rise of digital computers, the musical acoustics community has held sustained interest in performing direct physical simulation to create virtual instruments. In computer graphics, physically-based sound synthesis and propagation have seen rising interest over the past couple decades, with the goal of modeling sounds that are perceptually similar to reality. Our work shares this motivation: while the simulations are based closely on physical principles, the primary objective is to actively support user interaction and synthesize sounds that resemble everyday experience.

Existing real-time virtual wind instruments are based on Digital Waveguides and related techniques [Smith 2010]. The essential idea is that constant-speed wave propagation in one dimension can be expressed as a superposition of two opposing traveling waves, implementable efficiently as digital delay lines. Assuming propagating wavefronts are planar or spherical (infinite cylindrical or

ACM Reference Format

Allen, A., Raghuvanshi, N. 2015. Aerophones in Flatland: Interactive Wave Simulation of Wind Instruments. *ACM Trans. Graph.* 34, 4, Article 134 (August 2015), 11 pages. DOI = 10.1145/2767001
<http://doi.acm.org/10.1145/2767001>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2767001>

Elastic Textures for Additive Fabrication

Julian Panetta¹ *

Qingnan Zhou¹ *

Luigi Malomo²

Nico Pietroni²

Paolo Cignoni²

Denis Zorin¹

¹New York University

²Visual Computing Lab, CNR-ISTI




Figure 1: Six basic elastic textures are used to obtain a large range of homogenized isotropic material properties. A $3 \times 3 \times 1$ tiling of each pattern is shown, along with rendered (left) and fabricated (right) cell geometry below. The naming convention is explained in Section 4.

Abstract

We introduce *elastic textures*: a set of parametric, tileable, printable, cubic patterns achieving a broad range of isotropic elastic material properties: the softest pattern is over a thousand times softer than the stiffest, and the Poisson's ratios range from below zero to nearly 0.5. Using a combinatorial search over topologies followed by shape optimization, we explore a wide space of truss-like, symmetric 3D patterns to obtain a small family. This pattern family can be printed without internal support structure on a single-material 3D printer and can be used to fabricate objects with prescribed mechanical behavior. The family can be extended easily to create anisotropic patterns with target orthotropic properties. We demonstrate that our elastic textures are able to achieve a user-supplied varying material property distribution. We also present a material optimization algorithm to choose material properties at each point within an object to best fit a target deformation under a prescribed scenario. We show that, by fabricating these spatially varying materials with elastic textures, the desired behavior is achieved.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—[Geometric algorithms, languages, and systems];

Keywords: additive fabrication, microstructures, deformable objects, homogenization, shape optimization, goal-based material design

*Joint first authors

ACM Reference Format

Panetta, J., Zhou, Q., Malomo, L., Pietroni, N., Cignoni, P., Zorin, D. 2015. Elastic Textures for Additive Fabrication. *ACM Trans. Graph.* 34, 4, Article 135 (August 2015), 12 pages. DOI = 10.1145/2766937 <http://doi.acm.org/10.1145/2766937>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766937>

1 Introduction

Rapid advances in the accessibility of additive fabrication has a significant impact on how manufacturable geometric models are constructed. A key distinctive feature of common additive fabrication technologies is that the cost and time of production is practically uncorrelated with structural complexity: in fact, a complex structure using less material may be both cheaper and faster to produce.

Complex structures, aside from potentially reducing costs, open up many new possibilities, in particular for manufacturing deformable objects. By varying a small-scale structure, one can adjust a variety of material properties, from elasticity to permeability. Importantly, these properties can be varied nearly continuously over the object, something that is not commonly done in traditional processes. As it was observed in prior work, this opens up many new possibilities for object behavior.

Small-scale structures present a set of new design challenges: in all but the simplest cases, these are hard or impossible to design by hand to meet specific goals. At the same time, computational optimization of fine-scale variable structure over a whole object, even of moderate size, can easily result in numerically difficult topology and shape optimization problems with millions of variables.

In this paper, we describe *elastic volumetric textures*, a library of tileable parameterized 3D small-scale structures that can be used to control the elastic material properties of an object. Applying such textures to a hex mesh with target material properties specified per element is similar to using dithering to achieve a continuous variation of brightness or color.

In a sense, almost *all* material properties owe themselves to small-scale structures at the molecular or crystal level, and a large body of work in nanoscience aims to control material properties precisely by structure design. These works must accommodate constraints imposed by the specific properties of the elements and molecules used, the need for self-assembly, and other considerations.

Our focus is on larger-scale structures, which can be manufactured using existing 3D printing technology. With feature sizes at the scale of $10\mu\text{m}$ - $100\mu\text{m}$, these are well described by conventional elasticity theory. While this type of structure was also extensively

Microstructures to Control Elasticity in 3D Printing

Christian Schumacher^{1,2}

Bernd Bickel^{1,3}

Jan Rys²

Steve Marschner⁴

Chiara Daraio²

Markus Gross^{1,2}

¹Disney Research Zurich

²ETH Zurich

³IST Austria

⁴Cornell University



Figure 1: Given a virtual object with specified elasticity material parameters (blue=soft, red=stiff), our method computes an assemblage of small-scale structures that approximates the desired elastic behavior and requires only a single material for fabrication.

Abstract

We propose a method for fabricating deformable objects with spatially varying elasticity using 3D printing. Using a single, relatively stiff printer material, our method designs an assembly of small-scale microstructures that have the effect of a softer material at the object scale, with properties depending on the microstructure used in each part of the object. We build on work in the area of metamaterials, using numerical optimization to design tiled microstructures with desired properties, but with the key difference that our method designs families of related structures that can be interpolated to smoothly vary the material properties over a wide range. To create an object with spatially varying elastic properties, we tile the object's interior with microstructures drawn from these families, generating a different microstructure for each cell using an efficient algorithm to select compatible structures for neighboring cells. We show results computed for both 2D and 3D objects, validating several 2D and 3D printed structures using standard material tests as well as demonstrating various example applications.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling

Keywords: fabrication, topology optimization, 3D printing

1 Introduction

With the emergence of affordable 3D printing hardware and online 3D printing services, additive manufacturing technology comes with the promise to make the creation of complex functional physical artifacts as easy as providing a virtual description. Many functional objects in our everyday life consist of elastic, deformable

material, and the material properties are often inextricably linked to function. Unfortunately, elastic properties are not as easy to control as geometry, since additive manufacturing technologies can usually use only a single material, or a very small set of materials, which often do not match the desired elastic deformation behavior. However, 3D printing easily creates complex, high-resolution 3D structures, enabling the creation of *metamaterials* with properties that are otherwise unachievable with available printer materials.

Metamaterials are assemblies of small-scale structures that obtain their bulk properties from the shape and arrangement of the structures rather than from the composition of the material itself. For example, based on this principle, Lakes [1987] presented the first engineered materials that exhibit a negative Poisson's ratio. Since then, numerous designs have been proposed, usually consisting of a periodic tiling of a basic pattern, and engineering their structures is an active area of research [Lee et al. 2012].

While designing a tiled microstructure to match given *homogeneous* material properties can be achieved with modest extensions to the state of the art, designing a complex microstructural assembly to achieve *heterogeneous*, spatially varying properties is much more challenging. We face a complex inverse problem: to determine a discrete small-scale material distribution at the resolution of the 3D printer that yields the desired macroscopic elastic behavior. Inverse problems of this type have been explored for designing periodic structures that can be tiled to synthesize homogeneous volumes, but the methods are computation-intensive and do not scale to designing non-periodic structures for objects with spatially varying material properties.

Our goal is to enable users to employ metamaterials in their 3D printing workflow, generating appropriate structures specifically for their available 3D printer model and base material within seconds. Clearly, designing entire models on the fly by directly optimizing over the whole structure is impractical. Instead, we propose a data-driven approach that efficiently assembles models out of pre-computed small-scale structures so that the result both resembles the desired local elasticity and also is within the capabilities of the available output device.

First, we precompute a database of tiled structures indexed by their elastic properties. We want these structures to cover a large and ideally continuous region in the space of possible elastic behaviors. To achieve this goal, we introduce an optimization method for sampling structures that exhibit a range of desired behaviors, but are

ACM Reference Format

Schumacher, C., Bickel, B., Rys, J., Marschner, S., Daraio, C., Gross, M. 2015. Microstructures to Control Elasticity in 3D Printing. *ACM Trans. Graph.* 34, 4, Article 136 (August 2015), 13 pages. DOI = 10.1145/2766926 <http://doi.acm.org/10.1145/2766926>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766926>

By-Example Synthesis of Structurally Sound Patterns

J  r  mie Dumas*

Universit   de Lorraine, INRIA

An Lu*

INRIA, T.U. M  nchen

Sylvain Lefebvre

INRIA

Jun Wu

T.U. M  nchen

Christian Dick

T.U. M  nchen



Figure 1: Our by-example pattern synthesis algorithm produces structurally sound patterns along the surface of an object, resembling the input exemplar. The reinforcements are integrated within the pattern, by a joint optimization of appearance and structural properties. **Top:** The Stanford bunny with a variety of synthesized patterns. **Bottom:** Example patterns. These objects are printed in ABS plastic on low-cost filament printers, using a dense support. The patterns are fully connected and survived the cleaning process thanks to their reinforced structure. Yet, the reinforcements are inconspicuous as they seamlessly blend within the appearance.

Abstract

Several techniques exist to automatically synthesize a 2D image resembling an input exemplar texture. Most of the approaches optimize a new image so that the color neighborhoods in the output closely match those in the input, across all scales. In this paper we revisit by-example texture synthesis in the context of additive manufacturing. Our goal is to generate not only colors, but also structure along output surfaces: given an exemplar indicating ‘solid’ and ‘empty’ pixels, we generate a similar pattern along the output surface. The core challenge is to guarantee that the pattern is not only fully connected, but also structurally sound.

To achieve this goal we propose a novel formulation for on-surface by-example texture synthesis that directly works in a voxel shell around the surface. It enables efficient local updates to the pattern, letting our structural optimizer perform changes that improve the overall rigidity of the pattern. We use this technique in an iterative scheme that jointly optimizes for *appearance* and *structural soundness*. We consider fabricability constraints and a user-provided description of a force profile that the object has to resist.

Our results fully exploit the capabilities of additive manufacturing by letting users design intricate structures along surfaces. The structures are complex, yet they resemble input exemplars, resulting in a modeling tool accessible to casual users.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling;

Keywords: Texture Synthesis, Fabrication, By-example Modeling

1 Introduction

Additive manufacturing empowers designers and artists with an unprecedented ability to imagine and manufacture fine, intricate patterns. The patterns may be arranged in a delicate overall structure that flows in space and suggests the surface of a larger object. The sculptures *Crania Anatomica Filigre* by artist Joshua Harker [2011], the surface autoglyphs by Henry Segerman [2009] or the designs by company *Nervous System* [Rosenkrantz and Louis-Rosenberg 2007] are impressive and fascinating examples of this trend, also witnessed by the popularity of Voronoi carvings on the sharing platform *Thingiverse* (e.g. Chess Set - Voronoi Style thing:172960, Coral Candle Fixture thing:32513).

In this paper we consider the problem of automatically generating such patterns, from an input example. Our intent is to empower casual users and designers with a tool that quickly generates a compelling pattern that prints correctly and does not break in every-day use. Performing this task by hand is very difficult: the user has to be skilled in the use of CAD systems – which are often not targeted at such models – but also needs a good understanding of the limitations of additive manufacturing as well as notions of mechanical engineering to foresee when and how the object might break.

Using our approach, the user quickly obtains a solution that enforces these constraints, and resembles the input pattern. She is then able to focus on the most important task: exploiting our technique for designing interesting and intriguing objects.

We understand this problem as an instance of *by-example texture synthesis*, a long standing problem in Computer Graphics [Wei et al. 2009]. Despite the wide spectrum of available methods, only a few

ACM Reference Format

Dumas, J., Lu, A., Lefebvre, S., Wu, J., Dick, C. 2015. By-Example Synthesis of Structurally Sound Patterns. *ACM Trans. Graph.* 34, 4, Article 137 (August 2015), 12 pages. DOI = 10.1145/2766984 <http://doi.acm.org/10.1145/2766984>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766984>

*Joint first authors.

Design and Fabrication of Flexible Rod Meshes

Jesús Pérez¹

Bernhard Thomaszewski²

Stelian Coros^{2,3}

Bernd Bickel^{2,4}

José A. Canabal¹

Robert Sumner²

Miguel A. Otaduy¹

¹URJC Madrid

²Disney Research Zurich

³Carnegie Mellon University

⁴IST Austria




Figure 1: Our computational design method allows us to fabricate a deformable hat with a desired deformation behavior. The two left columns show a hat with the default rod mesh, which does not deform as desired (shown in transparent gray). The two right columns show that, by optimizing the radii and rest-shape of the rod mesh, we can fabricate in one piece a deformable hat that deforms as desired.

Abstract

We present a computational tool for fabrication-oriented design of flexible rod meshes. Given a deformable surface and a set of deformed poses as input, our method automatically computes a printable rod mesh that, once manufactured, closely matches the input poses under the same boundary conditions. The core of our method is formed by an optimization scheme that adjusts the cross-sectional profiles of the rods and their rest centerline in order to best approximate the target deformations. This approach allows us to locally control the bending and stretching resistance of the surface with a single material, yielding high design flexibility and low fabrication cost.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling

Keywords: fabrication, material design, rods

1 Introduction

Digital fabrication enables an extremely fast transition from virtual prototypes to their physical realization. Originally limited to rigid materials, 3D printers have now opened the door for fast prototyping of deformable objects. But while rigid designs merely

required the specification of geometry, deformable printing technology poses a grand challenge on computer-aided design: how to specify the desired deformation properties of the target object. The simplest approach could be to specify local material parameters at each point in the target object, and then map these parameters to the closest available material. Unfortunately, it is highly non-trivial to understand what combination of local material parameters will produce a desired global deformation behavior.

We explore the use of flexible rod meshes as an implementation of deformable surfaces. Rod meshes are ideal for building lightweight and cost-efficient physical shapes. The mesh provides a truss structure for a deformable surface and can even be filled with foam or dressed with fabric that is allowed to slide. But most importantly, the global deformation properties of a rod mesh can be adjusted simply by locally varying the cross-sectional parameters of the rods. As a result, a heterogeneous deformable object can be fabricated in one piece and from a single base material using a variety of rapid prototyping technologies. We can find instances of flexible rod meshes in furniture, apparel, architecture or accessories, for which even a standalone rod mesh might be found aesthetically pleasing. There are many other applications of rod meshes as supporting structures, which include but are not limited to toys, puppeteering, costumes, animatronics and robotics. Those use cases would potentially need volume to accommodate internal components and/or actuators, what encourages even further the use of deformable surfaces.

In this paper, we propose a computational tool for the design and fabrication of flexible rod meshes. As shown in Fig. 1, our method takes as input several poses of a deformable surface with known position and/or force constraints and automatically computes a printable representation of a rod mesh that best approximates the shapes. One of the main features of our method is the choice of design space. We use rod meshes dominated by hexagonal faces, because hexagons can stretch, shear and bend by deforming their edges. Given such a mesh, its design space consists of the rest-shape and two orthogonal radii describing the ellipsoidal cross section at each point of the rods. By adjusting two orthogonal radii conveniently

ACM Reference Format

Pérez, J., Thomaszewski, B., Coros, S., Bickel, B., Canabal, J., Sumner, R., Otaduy, M. 2015. Design and Fabrication of Flexible Rod Meshes. ACM Trans. Graph. 34, 4, Article 138 (August 2015), 12 pages. DOI = 10.1145/2766998 <http://doi.acm.org/10.1145/2766998>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766998>

Palette-based Photo Recoloring

Huiwen Chang¹

Ohad Fried¹

Yiming Liu¹

Stephen DiVerdi²

Adam Finkelstein¹

¹Princeton University

²Google




Figure 1: Palette-based photo recoloring. From left: original (computed palette below); user changes green palette entry to red (underlined), and the system recolors photo to match; user changes multiple colors to make two other styles. Photo courtesy of the MIT-Adobe FiveK Dataset [2011].

Abstract

Image editing applications offer a wide array of tools for color manipulation. Some of these tools are easy to understand but offer a limited range of expressiveness. Other more powerful tools are time consuming for experts and inscrutable to novices. Researchers have described a variety of more sophisticated methods but these are typically not interactive, which is crucial for creative exploration. This paper introduces a simple, intuitive and interactive tool that allows non-experts to recolor an image by editing a color palette. This system is comprised of several components: a GUI that is easy to learn and understand, an efficient algorithm for creating a color palette from an image, and a novel color transfer algorithm that recolors the image based on a user-modified palette. We evaluate our approach via a user study, showing that it is faster and easier to use than two alternatives, and allows untrained users to achieve results comparable to those of experts using professional software.

CR Categories: I.3.4 [Computer Graphics]: Graphics Utilities

Keywords: photo recoloring, color transformation, palette

1 Introduction

Research and commercial software offer a myriad of tools for manipulating the colors in photographs. Unfortunately these tools remain largely inscrutable to non-experts. Many features like the “levels tool” in software like Photoshop and iPhoto require the user to interpret histograms and to have a good mental model of how color spaces like RGB work, so non-experts have weak intuition about their behavior. There is a natural tradeoff between ease of use and range of expressiveness, so for example a simple hue slider, while easier to understand and manipulate than the levels tool, offers substantially less control over the resulting image. This paper introduces a tool that is easy for novices to learn while offering a broad expressive range.

Methods like that of Reinhard et al. [2001] and Yoo et al. [2013] allow a user to specify complex image modifications by simply providing an example; however an example of the kind of change the user would like to make is often unavailable. The method of Liu et al. [2014] allows users to modify the global statistics of an image by simply typing a text query like “vintage” or “new york.” However, for many desired color modifications it is hard to predict what text query would yield the desired effect. Another challenge in color manipulation is to selectively apply modifications – either locally within the image (e.g., this hat) or locally in color space (e.g., this range of blue colors) instead of globally. Selection is particularly challenging for non-experts, and a binary selection mask often leads to visual artifacts at the selection boundaries.

Our approach specifies both the colors to be manipulated and the modifications to these colors via a *color palette* – a small set of colors that digest the full range of colors in the image. Given an image, we generate a suitable palette. The user can then modify the image by modifying the colors in the palette (Figure 1). The image is changed globally such that the chosen colors are interpolated exactly with a smooth falloff in color space expressed through radial basis functions. These operations are performed in LAB color space to provide perceptual uniformity in the falloff. The naive application of this paradigm would in general lead to several kinds of artifacts. First, some pixels could go out of gamut. Simply clamping to the gamut can cause a color gradient to be lost. Therefore we formulate the radial falloff in color space so as to squeeze colors towards the gamut boundary. Second, many natural palette modifications would give rise to unpleasant visual artifacts wherein the relative brightness of different pixels is inverted. Thus, our color transfer function is tightly coupled with a subtle GUI affordance that together ensure monotonicity in the resulting changes in luminance.

This kind of color editing interface offers the best creative freedom when the user has interactive feedback while they explore various options. Therefore we show that our algorithm can easily be accelerated by a table-based approach that allows it to run at interactive frame rates, even when implemented in javascript running in a web browser. It is even fast enough to recolor video in the browser as it is being streamed over the network.

We perform a study showing that with our tool untrained users can produce similar results to those of expert Photoshop users. Finally, we show that our palette-based color transfer framework also supports other interfaces including a stroke-based interface, localized editing via a selection mask, fully-automatic palette improvement, and editing a collection of images simultaneously.

ACM Reference Format

Chang, H., Fried, O., Liu, Y., DiVerdi, S., Finkelstein, A. 2015. Palette-Based Photo Recoloring. ACM Trans. Graph. 34, 4, Article 139 (August 2015), 11 pages. DOI = 10.1145/2766978 <http://doi.acm.org/10.1145/2766978>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766978>

3DFlow: Continuous Summarization of Mesh Editing Workflows

Jonathan D. Denning*

*Taylor University

Valentina Tibaldo†

†Sapienza University of Rome

Fabio Pellacini†

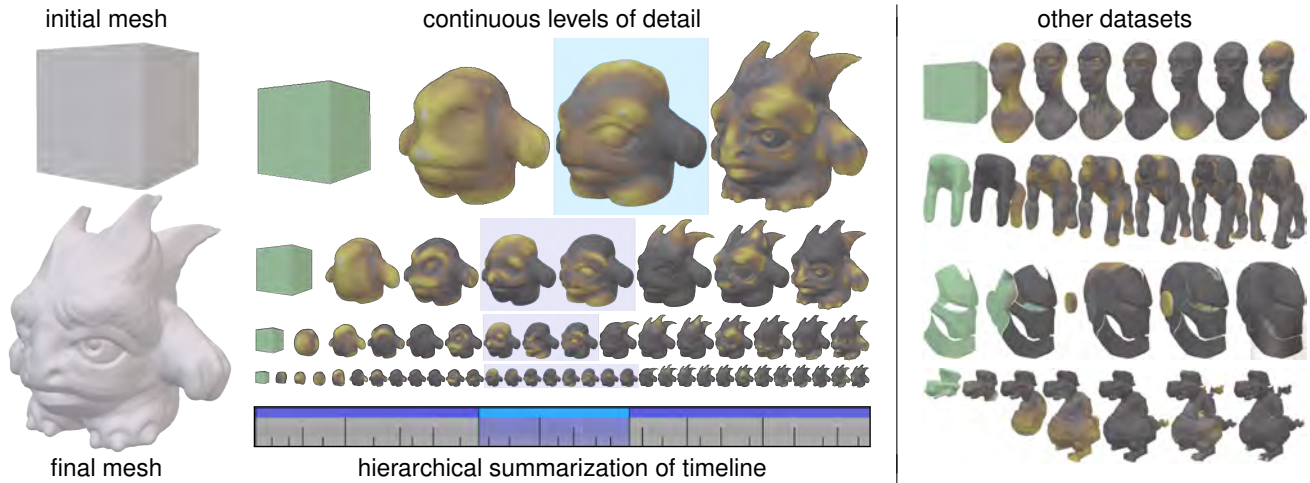


Figure 1: Continuous levels of details automatically constructed from a 30 minute digital sculpting session of a professional artist. The artist sculpted the cube (top-left) into a monster (bottom-left) in 797 strokes using dynamic remeshing techniques. The center column shows the sequence summarized in 4, 8, 16, and 32 steps (top) and the corresponding timeline (bottom). The mesh is colored green to indicate created geometry and golden to indicate the strength of change from the previous mesh. Blue highlighting and vertical black lines indicate the hierarchical summarization. Four additional sequences are shown at different levels of detail on the right.

Abstract

Mesh editing software is improving, allowing skilled artists to create detailed meshes efficiently. For a variety of reasons, artists are interested in sharing not just their final mesh but also their whole workflow, though the common media for sharing has limitations. In this paper, we present *3DFlow*, an algorithm that computes continuous summarizations of mesh editing workflows. *3DFlow* takes as input a sequence of meshes and outputs a visualization of the workflow summarized at any level of detail. The output is enhanced by highlighting edited regions and, if provided, overlaying visual annotations to indicated the artist's work, e.g. summarizing brush strokes in sculpting. We tested *3DFlow* with a large set of inputs using a variety of mesh editing techniques, from digital sculpting to low-poly modeling, and found *3DFlow* performed well for all. Furthermore, *3DFlow* is independent of the modeling software used because it requires only mesh snapshots, and uses the additional information only for optional overlays. We release *3DFlow* as open source for artists to showcase their work and release all our datasets so other researchers can improve upon our work.

1 Introduction

3D artists commonly showcase their workflows using time-lapse videos, as screen capturing software is simple to use and requires

very little interruption in the artist's work. Even for relatively simple 3D models, though, mesh editing workflows are long, ranging from tens of minutes to several hours of work, and involve thousands of operations. Time-lapse videos are not very effective at these lengths as the artist must make a trade-off between presenting the details of their workflow and keeping the presentation as short as possible. Choosing either details or brevity impacts the effectiveness of the video. Motivated by this concern, recent research has explored ways to visualize and navigate lengthy recordings of artists at work, for modeling as well as image editing. While the previous work studied ways to provide levels of interactivity beyond static images and fixed video, their systems are still limited both by the input and the output.

In this paper, we present *3DFlow*, an algorithm for producing continuous summarizations of mesh editing workflows. Our approach does not require using instrumented software, can summarize digital sculpting workflows, and is agnostic to the 3D software and editing methods used—the input needs only to be a sequence of meshes. Figure 1 shows at different levels of detail the summaries of several workflows, including low-poly modeling and sculpting sessions using dynamic or uniform remeshing. *3DFlow* is inspired by two prior works. As in *Video Tapestries* [Barnes et al. 2010], we support continuous levels of summaries to allow arbitrary temporal zooming of the editing sequence. As in *MeshFlow* [Denning et al. 2011], we add visual annotations to highlight important changes and summarize the artist's edits.

Given a sequence of meshes, *3DFlow* first detects the changes made between subsequent meshes, called *mesh deltas*, and generates a dependency graph for these deltas, called a *depgraph*, to capture the spatial and temporal dependencies of the edits. Then *3DFlow* summarizes the depgraph by repeatedly contracting the edge of least weight, computed by a cost function over the strength and distance of changes in the spatial and temporal dimensions, and merging the corresponding deltas. When only one delta remains, *3DFlow* splits the merged deltas in reverse contracting order to produce continuous levels of details.

ACM Reference Format

Denning, J., Tibaldo, V., Pellacini, F. 2015. 3DFlow: Continuous Summarization of Mesh Editing Workflows. ACM Trans. Graph. 34, 4, Article 140 (August 2015), 10 pages. DOI = 10.1145/2766936 <http://doi.acm.org/10.1145/2766936>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766936>

Practical Hex-Mesh Optimization via Edge-Cone Rectification

Marco Livesu¹ Alla Sheffer¹ Nicholas Vining¹ Marco Tarini^{2,3}
¹ University of British Columbia ² Università dell'Insubria ³ ISTI-CNR




Figure 1: We optimize poorly shaped hex meshes (a), with multiple inverted (red) elements (b), to produce high-quality, inversion-free outputs (c) using an edge-cone (d) based framework. The log-scale histogram insets highlight the improvement in worst element quality. Quality is measured using minimal Scaled Jacobian, with 1 being optimal and negative values corresponding to inverted elements.

Abstract

The usability of hexahedral meshes depends on the degree to which the shape of their elements deviates from a perfect cube; a single concave, or *inverted* element makes a mesh unusable. While a range of methods exist for discretizing 3D objects with an initial topologically suitable hex mesh, their output meshes frequently contain poorly shaped and even inverted elements, requiring a further quality optimization step. We introduce a novel framework for optimizing hex-mesh quality capable of generating inversion-free high-quality meshes from such poor initial inputs. We recast hex quality improvement as an optimization of the shape of overlapping *cones*, or unions, of tetrahedra surrounding every *directed* edge in the hex mesh, and show the two to be equivalent. We then formulate cone shape optimization as a sequence of convex quadratic optimization problems, where hex convexity is encoded via simple linear inequality constraints. As this solution space may be empty, we therefore present an alternate formulation which allows the solver to proceed even when constraints cannot be satisfied exactly. We iteratively improve mesh element quality by solving at each step a set of local, per-cone, convex constrained optimization problems, followed by a global energy minimization step which reconciles these local solutions. This latter method provides no theoretical guarantees on the solution but produces inversion-free, high quality meshes in practice. We demonstrate the robustness of our framework by optimizing numerous poor quality input meshes generated using a variety of initial meshing methods and producing high-quality inversion-free meshes in each case. We further validate our algorithm by comparing it against previous work, and demonstrate a significant improvement in both worst and average element quality.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms;

Keywords: hexahedral meshes, mesh optimization

1 Introduction

Hexahedral meshes are the finite element discretization of choice in multiple engineering domains [Blacker 2000; Shepherd and Johnson 2008; Owen 2009]. The reliability of the finite element simulation is highly dependent on the mesh element quality, or the degree to which they deviate from a perfect cube [Pébay et al. 2007]. Simulation results depend on both average and minimum element quality [Labelle and Shewchuk 2007; Pébay et al. 2007]. Even a single *inverted*, or non-convex element, makes a mesh unusable for simulation. The generation of quality hexahedral meshes is typically a two step process which first generates an initial mesh whose connectivity is designed to fit the input at hand, and then modifies the vertex positions so as to optimize the mesh element shape while keeping the connectivity fixed [Owen 2009]. The initial meshes that serve as input to the optimization step frequently contain numerous poorly shaped, and even inverted, elements. We introduce a new mesh optimization method which takes an initial low quality hex mesh as input and generates an inversion-free high quality mesh closely conforming to the input surface geometry. This conformity is important for the finite element analysis to capture the true behavior of the meshed object. Our method recasts mesh optimization as a sequence of convex, easy to optimize problems and produces high-quality meshes starting from inputs with numerous inversions and a range of diverse connectivities representative of the commonly used initial hex mesh generation approaches (Section 4).

Hex mesh optimization is an open, challenging research problem [Knupp 2001; Shepherd and Johnson 2008; Ruiz-Gironés et al. 2014b]. Existing state-of-the-art methods focus on directly optimizing a range of hex shape quality metrics such as minimal [Knupp 2001; Brewer et al. 2003] or average [Ruiz-Gironés et al. 2014b] scaled Jacobian, or condition number [Knupp 2003; Brewer et al. 2003]. Optimizing shape quality directly is difficult since both the quality metrics and the element-convexity constraints, expressed as a function of vertex positions, are non-linear; the resulting solution space is non-convex, which in turn motivates these frameworks to employ Gauss-Seidel schemes that update one vertex position at a time. Our indirect, global approach, significantly improves on the results of these methods as demonstrated in Sections 2 and 4.

The key observation behind our optimization framework is that we can formulate shape quality improvement in an easier to optimize

ACM Reference Format

Livesu, M., Sheffer, A., Vining, N., Tarini, M. 2015. Practical Hex-Mesh Optimization via Edge-Cone Rectification. ACM Trans. Graph. 34, 4, Article 141 (August 2015), 11 pages. DOI = 10.1145/2766905 <http://doi.acm.org/10.1145/2766905>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766905>

Hexahedral Mesh Re-parameterization from Aligned Base-Complex

Xifeng Gao*

Zhigang Deng[†]
University of Houston

Guoning Chen[‡]




Figure 1: (a) An input hex-mesh [Li et al. 2012]: The image on the left shows its base-complex that partitions the hexahedral mesh into different large components, illustrated with different colors on the right. Due to the misalignments between singularities, many (typically small) components arise. For instance, a strip of small components near the sharp feature is highlighted. (b) Our alignment algorithm reduces the complexity of the base-complex but leads to a hex-mesh with a large distortion. (c) Both the singularity placement and the element quality of the resulting hex-mesh are improved by our structure-aware optimization algorithm.

Abstract

Recently, generating a high quality all-hex mesh of a given volume has gained much attention. However, little, if any, effort has been put into the optimization of the hex-mesh structure, which is equally important to the local element quality of a hex-mesh that may influence the performance and accuracy of subsequent computations. In this paper, we present a first and complete pipeline to optimize the global structure of a hex-mesh. Specifically, we first extract the *base-complex* of a hex-mesh and study the misalignments among its singularities by adapting the previously introduced hexahedral sheets to the base-complex. Second, we identify the valid removal *base-complex sheets* from the base-complex that contain misaligned singularities. We then propose an effective algorithm to remove these valid removal sheets in order. Finally, we present a structure-aware optimization strategy to improve the geometric quality of the resulting hex-mesh after fixing the misalignments. Our experimental results demonstrate that our pipeline can significantly reduce the number of components of a variety of hex-meshes generated by state-of-the-art methods, while maintaining high geometric quality.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

Keywords: Hex-mesh, Singularity alignment, Optimization

*e-mail: gxf.xisha@gmail.com

[†]e-mail: zdeng4@uh.edu

[‡]e-mail: gchen16@uh.edu

ACM Reference Format

Gao, X., Deng, Z., Chen, G. 2015. Hexahedral Mesh Re-Parameterization from Aligned Base-Complex. ACM Trans. Graph. 34, 4, Article 142 (August 2015), 10 pages. DOI = 10.1145/2766941 <http://doi.acm.org/10.1145/2766941>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766941>

1 Introduction

Given a two-dimensional closed surface \mathcal{M} , producing a high quality volume parameterization $f: \mathcal{M} \rightarrow \mathbb{R}^3$, that is aligned with boundary features, is a prerequisite for a number of important applications including finite element analysis (FEA), volume sampling, 3D texture synthesis, and data compression. A hex-mesh \mathcal{H} resulting from a non-degenerate volume parameterization is suitable for partial differential equation (PDE) solving, tensor-product/trivariate spline fitting [Wang et al. 2012; Li and Qin 2012; Martin et al. 2008], Isogeometric Analysis (IGA) [Hughes et al. 2005; Bazilevs et al. 2006], and multi-grid and adaptive computations [Leonard et al. 2000; Wada et al. 2006].

Although the quality criteria of hex-meshes are application dependent, the hex-meshes that conform to boundary surfaces, preserve boundary features, have regular distribution of parameterization lines, and have low element distortion, are generally preferred [Motooka et al. 2011; Gao et al. 2014]. It is noteworthy that a hex-mesh with the above desired characteristics usually require a certain amount of well-placed singularities that are distributed either on the boundary or in the interior of the volume [Nieser et al. 2011; Huang et al. 2011; Li et al. 2012; Jiang et al. 2014]. Unlike the singularities in frame fields, the singularities of a semi-structured mesh consist of irregular mesh vertices and edges, i.e., those vertices whose valences are not 4 on a 2D quad-mesh [Bommes et al. 2011; Tarini et al. 2011], or edges whose valences are not 4 in the interior of a hex-mesh. These irregular vertices and edges correspond to those places where the parameterization is discontinuous. Intuitively, the singularities as well as the separation structures starting from them partition the domain into small regions. Figure 1(a) shows such a partition example for a Sculpture hex-mesh. Its wireframe and associated transparent surfaces visualize the partitioning structure, which is referred to as the *base-complex* in this paper. Each sub-volume of this partition, referred to as a *component* (i.e., the individual colored regions shown in the right of Figure 1(a)), can be mapped to a cube.

With the above partition, a C^2 spline basis can be fit to each component for FEA, while only C^0 continuity can be guaranteed across the boundaries of different components. For applications, such as IGA, that seek a high level of smoothness throughout the entire volumetric domain, fewer components are desired, as it would lead to more accurate and faster simulations [Hughes et al. 2005]. However, given the same set of singularities, without careful control, it

Dyadic T-mesh Subdivision

Denis Kovacs^{1,2*}

Justin Bisceglia^{1,3†}

Denis Zorin^{1‡}

¹New York University ²FiftyThree ³Blue Sky Studios

Abstract

Meshes with T-joints (T-meshes) and related high-order surfaces have many advantages in situations where flexible local refinement is needed. At the same time, designing subdivision rules and bases for T-meshes is much more difficult, and fewer options are available. For common geometric modeling tasks it is desirable to retain the simplicity and flexibility of commonly used subdivision surfaces, and extend them to handle T-meshes.

We propose a subdivision scheme extending Catmull-Clark and NURSS to a special class of quad T-meshes, dyadic T-meshes, which have no more than one T-joint per edge. Our scheme is based on a factorization with the same structure as Catmull-Clark subdivision. On regular T-meshes it is a refinement scheme for a subset of standard T-splines. While we use more variations of subdivision masks compared to Catmull-Clark and NURSS, the minimal size of the stencil is maintained, and all variations in formulas are due to simple changes in coefficients.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—[Geometric algorithms, languages, and systems];

Keywords: Subdivision surfaces, T-meshes, T-splines

1 Introduction

Subdivision surfaces are widely used to create free-form 3D shapes for computer animation and are a popular tool for conceptual design. The most obvious reason for this is that subdivision surfaces support flexible and easy-to-modify control meshes of arbitrary topology and the predominant quad mesh structure is suitable for many modeling needs.

Examination of common organic and mechanical models reveals that subdivision surface modeling often involves two tasks:

- Varying control mesh resolution between more and less feature-rich areas.
- Modifying connectivity to align features with control mesh edges.

In practice, these challenges are resolved using extraordinary vertices, most commonly in pairs of valence 3 and 5, redirecting the edge flow (Figure 1).

For commonly used schemes, the quality at the extraordinary vertices in general cannot match the surface quality of the regular parts of the surface, where it reduces to polynomial patches (Figure 5).





Figure 1: Top: T-joints (blue) allow for local refinement without extraordinary vertices (red). Bottom: successive subdivision turns each source T-face into mostly regular quads and a row of T-faces along the original T-edge.




Figure 2: T-joints can be used to keep tessellated faces roughly equal sized and thus prevent under- or over-tessellation.

T-joints significantly increase the freedom of choosing control mesh connectivity, without creating lower quality surfaces. In a restricted form produced by adaptive hierarchical refinement, these were available in subdivision surface context for a long time, both in literature [Zorin et al. 1997] and in commercial tools (e.g., Autodesk Maya). A more flexible representation based on T-splines is the primary free-form representation in Autodesk Fusion 360.

A significant body of literature has appeared on construction of smooth surfaces with T-joints, with T-splines being the most common construction. Yet, to the best of our knowledge, no subdivision schemes were proposed that operate directly on meshes with T-joints (T-meshes). The closest approach is a relatively complex T-NURCCs scheme [Sederberg et al. 2003], which has been devel-

*e-mail: kovacs@cs.nyu.edu

†e-mail: justinb@blueskystudios.com

‡e-mail: dzorin@cs.nyu.edu

ACM Reference Format

Kovacs, D., Bisceglia, J., Zorin, D. 2015. Dyadic T-Mesh Subdivision. ACM Trans. Graph. 34, 4, Article 143 (August 2015), 12 pages. DOI = 10.1145/2766972 <http://doi.acm.org/10.1145/2766972>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766972>

Lillicon: Using Transient Widgets to Create Scale Variations of Icons

Gilbert Louis Bernstein*
Stanford University

Wilmot Li†
Adobe Research




Figure 1: Icon editing with Lillicon. Starting from an input icon of a building (a), we use our system, Lillicon, to create a scale variation (b) that is more legible at small sizes (left). With Lillicon, we apply transient widgets to select and manipulate visually apparent features, like the “stroke” in (c). Creating this result involves thickening, translating, deleting strokes and adjusting contours (d). In total, these edits took less than four minutes to perform in Lillicon. The result differs from the input in several ways (e) that improve its legibility at the target size.

Abstract

Good icons are legible, and legible icons are scale-dependent. Experienced icon designers use a set of common strategies to create legible scale variations of icons, but executing those strategies with current tools can be challenging. In part, this is because many apparent objects, like hairlines formed by negative space, are not explicitly represented as objects in vector drawings. We present transient widgets as a mechanism for selecting and manipulating apparent objects that is independent of the underlying drawing representation. We implement transient widgets using a constraint-based editing framework; demonstrate their utility for performing the kinds of edits most common when producing scale variations of icons; and report qualitative feedback on the system from professional icon designers.

CR Categories: I.3.7 [Computer Graphics]—;

Keywords: downsampling, design tools, artist tools, vector graphics, constraint-based drawing, icons

1 Introduction

Icon design is a communication challenge. In part, this means choosing effective and appropriate symbols to convey an idea. But it also means making sure the depictions of those symbols are visually pleasing, consistent with other icons, and most of all *legible*. A significant component of this task is ensuring that all of these properties hold across the wide range of scales at which icons may be

viewed. For example, an icon that represents a software application typically appears at several different sizes in different contexts — as a file icon, in toolbars, and in an online marketplace on a range of different platforms. In addition, icons may be viewed on devices with different display densities that induce different physical viewing sizes.

Simple uniform scaling of icons does not suffice to produce effective, legible icons at all viewing scales. In the inline scroll example, the hairlines that delineate the top and bottom folds of the script disappear as the icon is scaled down uniformly. A similar effect arises if you hold this paper far enough away or take off glasses you may be wearing. To prevent such problems, designers adjust the proportions, spacing, and level-of-detail of icons to produce multiple scale variations that are suited for display at different target sizes.

The need for non-uniform scale variations is fundamental, despite profound changes in display technology. Recently, the proliferation of display devices in an increasing range of sizes, aspect ratios and pixel densities has been driving designers to adopt new practices to cope with this diversity while preserving the legibility and aesthetic quality of their designs. Nonetheless, neither the switch from raster to vector formats, nor the increasing saturation of users’ visual field by high-density displays will remove fundamental visual acuity limits (which vary with eyesight and viewing conditions). Nor will they eliminate the need to display icons in different parts of an application. Consequently, in carefully designed icon sets, designers will continue to hand craft scale variations to ensure legibility across all viewing situations.

While experienced icon designers are proficient at identifying the types of edits that are necessary to create effective scale variations, the process of executing such edits is often tedious. We observed two main difficulties. First, the relevant feature of the drawing that the designer wants to edit may not be conveniently exposed as a manipulable object. For example, in the script icon, the hairlines are defined as negative space outside of the filled black shape. There is no path, and thus no width parameter to edit. Second, any individual edit may require numerous secondary edits to preserve important properties of the drawing. These issues distract designers by focusing them on how to perform manipulations rather than which

*e-mail: gilbert@gilbertbernstein.com

†e-mail: wilmotli@adobe.com

ACM Reference Format
Bernstein, G., Li, W. 2015. Lillicon: Using Transient Widgets to Create Scale Variations of Icons. ACM Trans. Graph. 34, 4, Article 144 (August 2015), 11 pages. DOI = 10.1145/2766980
<http://doi.acm.org/10.1145/2766980>.

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766980>

Vector Graphics Animation with Time-Varying Topology

Boris Dalstein*
University of British Columbia

Rémi Ronfard
Inria, Univ. Grenoble Alpes, LJK, France

Michiel van de Panne
University of British Columbia




Figure 1: A space-time continuous 2D animation depicting a rotating torus, created without 3D tools. First, the animator draws key cells (in blue) using 2D vector graphics tools. Then, he specifies how to interpolate them using inbetween cells (in green). Our contribution is a novel data structure, called Vector Animation Complex (VAC), which enables such interaction paradigm.

Abstract

We introduce the Vector Animation Complex (VAC), a novel data structure for vector graphics animation, designed to support the modeling of time-continuous topological events. This allows features of a connected drawing to merge, split, appear, or disappear at desired times via keyframes that introduce the desired topological change. Because the resulting space-time complex directly captures the time-varying topological structure, features are readily edited in both space and time in a way that reflects the intent of the drawing. A formal description of the data structure is provided, along with topological and geometric invariants. We illustrate our modeling paradigm with experimental results on various examples.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations I.3.6 [Computer Graphics]: Methodology and Techniques—Graphics data structures and data types I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: vector graphics, 2D, animation, topology, space-time, cell complex, boundary-based representation

*dalboris@cs.ubc.ca

ACM Reference Format
Dalstein, B., Ronfard, R., van de Panne, M. 2015. Vector Graphics Animation with Time-Varying Topology. ACM Trans. Graph. 34, 4, Article 145 (August 2015), 12 pages. DOI = 10.1145/2766913
<http://doi.acm.org/10.1145/2766913>

Copyright Notice
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://dx.doi.org/10.1145/2766913>

1 Introduction

A fundamental difference between raster graphics and vector graphics is that the former is a *discrete* representation, while the latter is a *continuous* representation. Instead of storing individual pixels that our eyes readily interpret as curves, vector graphics stores curves that can be rendered at any resolution. As display devices spanning a wide range of resolutions proliferate, such resolution-independent representations are increasing in importance.

Similarly, a fundamental difference between traditional hand-drawn animation and 3D animation is that the former is *discrete* in time, while the latter is *continuous* in time. Instead of storing individual frames that our eyes interpret as motion, the use of animation curves allows a scene to be rendered at any frame rate.

Space-time continuous representations, i.e., representations that are resolution-independent both in the spatial domain *and* the temporal domain, are ubiquitous within computer graphics for their many advantages. They are typically based on the “model-then-animate” paradigm: a parameterized model is first developed and then animated over time using animation curves that interpolate *key values* of the parameters at *key times*. A limitation of this paradigm is the underlying assumption that the model can be parameterized by a *fixed* set of parameters that captures the desired intent. This is indeed possible for 3D animation and simple 2D animation, but it quickly becomes impractical in any 2D animation scenario where the number of strokes or how they intersect change over time. In other words, the “model-then-animate” paradigm fails when the *topology* of the model is time-dependent, which makes it challenging to represent space-time continuous animated vector graphics illustrations with time-varying topology.

In this paper, we address this problem by introducing the Vector Animation Complex (VAC). It is a cell complex immersed in space-time, specifically tailored to meet the requirements of vector graphics animation with non-fixed topology. Any time-slice of the complex is a valid Vector Graphics Complex (VGC) which make its rendering consistent with non-animated VGCs.

Accelerating Vector Graphics Rendering using the Graphics Hardware Pipeline

Vineet Batra¹, Mark J. Kilgard², Harish Kumar¹, and Tristan Lorach²

¹Adobe Systems

²NVIDIA

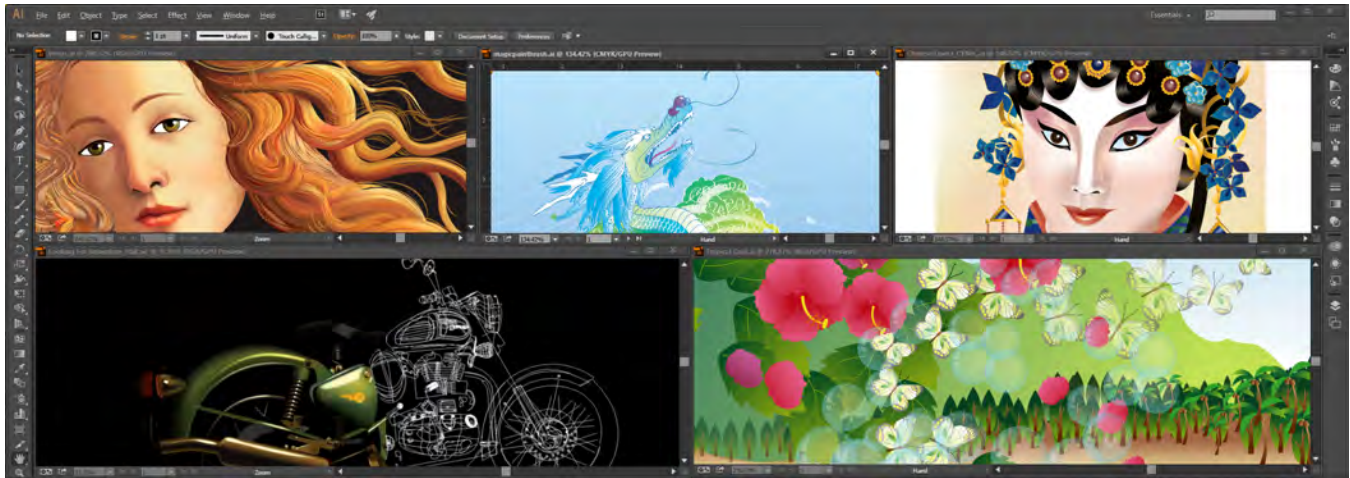


Figure 1: Five complex RGB and CMYK documents GPU-rendered by Illustrator CC; all rendered with “GPU Preview” enabled.

Abstract

We describe our successful initiative to accelerate Adobe Illustrator with the graphics hardware pipeline of modern GPUs. Relying on OpenGL 4.4 plus recent OpenGL extensions for advanced blend modes and first-class GPU-accelerated path rendering, we accelerate the Adobe Graphics Model (AGM) layer responsible for rendering sophisticated Illustrator scenes. Illustrator documents render in either an RGB or CMYK color mode. While GPUs are designed and optimized for RGB rendering, we orchestrate OpenGL rendering of vector content in the proper CMYK color space and accommodate the 5+ color components required. We support both non-isolated and isolated transparency groups, knockout, patterns, and arbitrary path clipping. We harness GPU tessellation to shade paths smoothly with gradient meshes. We do all this and render complex Illustrator scenes 2 to 6x faster than CPU rendering at Full HD resolutions; and 5 to 16x faster at Ultra HD resolutions.

CR Categories: I.3.4 [Computer Graphics]: Graphics Utilities—Graphics Editors;

Keywords: Illustrator, path rendering, vector graphics, OpenGL

1 Introduction

ACM Reference Format

Batra, V., Kilgard, M., Kumar, H., Lorach, T. 2015. Accelerating Vector Graphics Rendering using the Graphics Hardware Pipeline. ACM Trans. Graph. 34, 4, Article 146 (August 2015), 15 pages.
DOI = 10.1145/2766968 <http://doi.acm.org/10.1145/2766968>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766968>

Designers and artists worldwide rely on Adobe Illustrator to design and edit resolution-independent 2D artwork and typographic content. Illustrator was Adobe’s very first application when released over 27 years ago.

Prior to our work, no version utilized graphics hardware to accelerate Illustrator’s rendering. All rendering was performed entirely by the CPU. This situation is in stark contrast to the now ubiquitous GPU-acceleration of 3D graphics rendering in Computer-Aided Design, Animation, and Modeling applications. So while other graphical content creation applications readily benefit from the past 15 years of improvements in GPU functionality and performance, Illustrator could neither benefit from nor scale with the tremendous strides in GPU functionality and performance. Our work remedies this situation as Figure 1 shows.

The starting point for our work is OpenGL 4.4 [Khronos Group 2014] and the GPU-accelerated “stencil, then cover” path rendering functionality described in [Kilgard and Bolz 2012]. While the NV_path_rendering OpenGL extension [Kilgard 2012] provides very fast and resolution-independent rendering of first-class path objects just as we need, Illustrator’s rendering model requires much more than merely rendering paths. We had to develop our own strategies to handle features of Illustrator that, while dependent on path rendering, require considerably more sophisticated orchestration of the GPU. We focus on the GPU-based techniques we developed and productized to accomplish full GPU-acceleration of Illustrator.

Adobe originally developed Illustrator as a vector graphics editor for PostScript [Adobe Systems 1985] which implements the imaging model described by [Warnock and Wyatt 1982]. Illustrator today depends on the Portable Document Format (PDF) standard for its underlying rendering capabilities as described by the ISO 32000 standard [Adobe Systems 2008]. The evolution of PostScript to PDF introduced a number of sophisticated graphics capabilities for dealing with printer color spaces, compositing, and photorealistic artistic shading. These capabilities developed in parallel with but

Piko: A Framework for Authoring Programmable Graphics Pipelines

Anjul Patney^{1,2,*}

Stanley Tzeng^{1,2}

Kerry A. Seitz, Jr.²

John D. Owens²

¹NVIDIA

²University of California, Davis

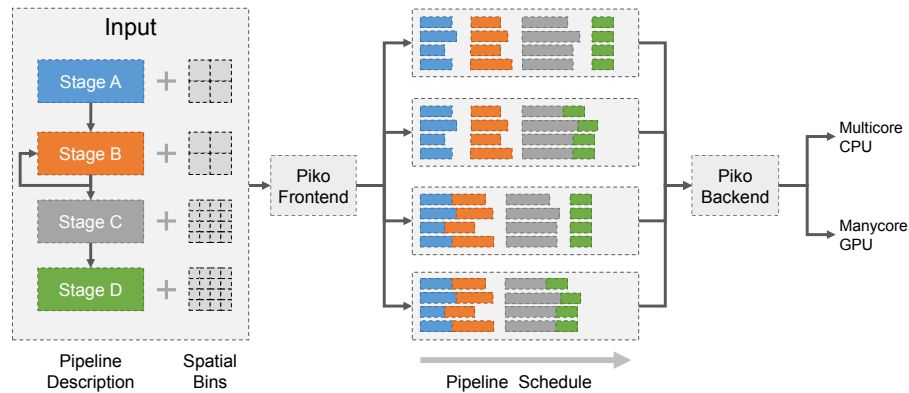


Figure 1: Piko is a framework for designing and implementing programmable graphics pipelines that can be easily retargeted to different application configurations and architectural targets. Piko’s input is a functional and structural description of the desired graphics pipeline, augmented with a per-stage grouping of computation into spatial bins (or tiles), and a scheduling preference for these bins. Our compiler generates efficient implementations of the input pipeline for multiple architectures and allows the programmer to tweak these implementations using simple changes in the bin configurations and scheduling preferences.

Abstract

We present Piko, a framework for designing, optimizing, and retargeting implementations of graphics pipelines on multiple architectures. Piko programmers express a graphics pipeline by organizing the computation within each stage into spatial bins and specifying a scheduling preference for these bins. Our compiler, *Pikoc*, compiles this input into an optimized implementation targeted to a massively-parallel GPU or a multicore CPU.

Piko manages work granularity in a programmable and flexible manner, allowing programmers to build load-balanced parallel pipeline implementations, to exploit spatial and producer-consumer locality in a pipeline implementation, and to explore tradeoffs between these considerations. We demonstrate that Piko can implement a wide range of pipelines, including rasterization, Reyes, ray tracing, rasterization/ray tracing hybrid, and deferred rendering. Piko allows us to implement efficient graphics pipelines with relative ease and to quickly explore design alternatives by modifying the spatial binning configurations and scheduling preferences for individual stages, all while delivering real-time performance that is within a factor six of state-of-the-art rendering systems.

CR Categories: I.3.1 [Computer Graphics]: Hardware Architecture—Parallel processing; I.3.2 [Computer Graphics]: Graphics Systems—Stand-alone systems

Keywords: graphics pipelines, parallel computing

ACM Reference Format

Patney, A., Tzeng, S., Seitz, K., Jr., J., Owens, J. 2015. Piko: A Framework for Authoring Programmable Graphics Pipelines. *ACM Trans. Graph.* 34, 4, Article 147 (August 2015), 13 pages. DOI = 10.1145/2766973 <http://doi.acm.org/10.1145/2766973>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH ’15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2766973>

1 Introduction

Renderers in computer graphics often build upon an underlying graphics pipeline: a series of computational stages that transform a scene description into an output image. Conceptually, graphics pipelines can be represented as a graph with stages as nodes and the flow of data along directed edges of the graph. While some renderers target the special-purpose hardware pipelines built into graphics processing units (GPUs), such as the OpenGL/Direct3D pipeline (the “OGL/D3D pipeline”), others use pipelines implemented in software, either on CPUs or, more recently, using the programmable capabilities of modern GPUs. This paper concentrates on the problem of implementing a graphics pipeline that is both highly programmable and high-performance by targeting programmable parallel processors like GPUs.

Hardware implementations of the OGL/D3D pipeline are extremely efficient and expose programmability through shaders which customize the behavior of stages within the pipeline. However, developers cannot easily customize the structure of the pipeline itself, or the function of non-programmable stages. This limited programmability makes it challenging to use hardware pipelines to implement other types of graphics pipelines like ray tracing, micropolygon-based pipelines, voxel rendering, volume rendering, and hybrids that incorporate components of multiple pipelines. Instead, developers have recently begun using programmable GPUs to implement these pipelines in software (Section 2), allowing their use in interactive applications.

Efficient implementations of graphics pipelines are complex: they must consider parallelism, load balancing, and locality within the bounds of a restrictive programming model. In general, successful pipeline implementations have been narrowly customized to a particular pipeline and often to a specific hardware target. The abstractions and techniques developed for their implementation are not easily extensible to the more general problem of creating efficient yet

*Corresponding author; apatney@nvidia.com.

Fast Grid-Free Surface Tracking

Nuttapong Chentanez^{†*}

[†]NVIDIA

Matthias Müller[†]

^{*}Chulalongkorn University

Miles Macklin[†]

Tae-Yong Kim[†]




Figure 1: Explosion of a water balloon simulated with particles and using our method for surface tracking. The liquid expands to form thin sheets and tendrils that occupy a large bounding volume.

Abstract

We present a novel explicit surface tracking method. Its main advantage over existing approaches is the fact that it is both completely grid-free and fast which makes it ideal for the use in large unbounded domains. A further advantage is that its running time is less sensitive to temporal variations of the input mesh than existing approaches. In terms of performance, the method provides a good trade-off point between speed and quality. The main idea behind our approach to handle topological changes is to delete all overlapping triangles and to fill or join the resulting holes in a robust and efficient way while guaranteeing that the output mesh is both manifold and without boundary. We demonstrate the flexibility, speed and quality of our method in various applications such as Eulerian and Lagrangian liquid simulations and the simulation of solids under large plastic deformations.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically Based Modeling; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation and Virtual Reality

Keywords: surface tracking, mesh repair, hole filling

1 Introduction

Tracking the free surface of a liquid is a challenging problem due to its complex shape and the frequent topological changes. Representing the surface as the level set of a scalar function on a grid is an

attractive approach because topological changes are handled automatically. However, the grid resolution puts a limit on the smallest feature size that can be represented by the level set. Therefore, very large grids are required to simulate thin sheets or small droplets in large domains.

The more recently introduced mesh-based surface tracking methods overcome this problem by representing the surface as a triangle mesh. There is no theoretical limit on the feature size that can be represented by explicit meshes. However, with a triangle mesh representation, topological events such as splitting and merging have to be handled explicitly which is a highly non-trivial task.

Most of the existing mesh-based surface tracking methods – while grid-free in general – still use a grid to fix the mesh in overlapping regions. As an example, Wojtan et al. used the marching cubes method [2009] and later convex hulls [2010] on a background grid to resolve overlaps. This puts a limit on the smallest size of topological events that can be handled and requires the choice of a grid resolution. Also, handling the transition from the grid-based to the grid-free representation is challenging and sensitive to numerical inaccuracies. This *stitching* process potentially introduces non-manifold meshes which might introduce difficulties at later stages.

The current state of the art grid-free single phase surface tracking method in computer graphics is El Topo by Brochu and Bridson [2009]. It produces high quality results and guarantees that the resulting mesh is overlap free but its complexity prevents the simulation of large scale scenes in reasonable time.

Our goal was to devise a grid-free method that is both robust and fast and can be used to track a surface in both grid-based and particle-based simulations. The basic idea is to delete overlapping triangles and to triangulate the resulting holes robustly and efficiently while guaranteeing manifoldness of the resulting mesh.

We achieved this with the following contributions:

- A method to remove topological noise.
- A fast way to ensure that the resulting mesh is manifold.
- A way to match holes for merging.
- A fast approximately optimal hole filling algorithm that considers the surrounding mesh. Its robustness allows the use of single precision floating point arithmetic.

ACM Reference Format

Chentanez, N., Mueller, M., Macklin, M., Kim, T. 2015. Fast Grid-Free Surface Tracking. ACM Trans. Graph. 34, 4, Article 148 (August 2015), 11 pages. DOI = 10.1145/2766991 <http://doi.acm.org/10.1145/2766991>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766991>

Double Bubbles Sans Toil and Trouble: Discrete Circulation-Preserving Vortex Sheets for Soap Films and Foams

Fang Da
Columbia University

Christopher Batty
University of Waterloo

Chris Wojtan
IST Austria

Eitan Grinspun
Columbia University



Figure 1: A variety of dynamic foam, film, and bubble scenarios captured by our method. Left: A small foam rearranges and settles to equilibrium. Center: A snapshot of an evolving catenoid soap film joining two circular wires, an instant before the film pinches apart. Right: A bubble with a wire constricting its mid-section gradually squeezes to one side.

Abstract

Simulating the delightful dynamics of soap films, bubbles, and foams has traditionally required the use of a fully three-dimensional many-phase Navier-Stokes solver, even though their visual appearance is completely dominated by the thin liquid surface. We depart from earlier work on soap bubbles and foams by noting that their dynamics are naturally described by a Lagrangian vortex sheet model in which *circulation* is the primary variable. This leads us to derive a novel circulation-preserving surface-only discretization of foam dynamics driven by surface tension on a non-manifold triangle mesh. We represent the surface using a mesh-based multimaterial surface tracker which supports complex bubble topology changes, and evolve the surface according to the ambient air flow induced by a scalar circulation field stored on the mesh. Surface tension forces give rise to a simple update rule for circulation, even at non-manifold Plateau borders, based on a discrete measure of signed scalar mean curvature. We further incorporate vertex constraints to enable the interaction of soap films with wires. The result is a method that is at once simple, robust, and efficient, yet able to capture an array of soap films behaviors including foam rearrangement, catenoid collapse, blowing bubbles, and double bubbles being pulled apart.

CR Categories: I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: fluids, vortex sheet, circulation, non-manifold mesh

ACM Reference Format

Da, F., Batty, C., Wojtan, C., Grinspun, E. 2015. Double Bubbles Sans Toil and Trouble: Discrete Circulation-Preserving Vortex Sheets for Soap Films and Foams. *ACM Trans. Graph.* 34, 4, Article 149 (August 2015), 9 pages. DOI = 10.1145/2767003 <http://doi.acm.org/10.1145/2767003>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/2767003>

1 Introduction

Although we encounter bubbles and foam on a daily basis while drinking coffee, washing the dishes, or playing with children's toys, the motion and structure of these phenomena are difficult to understand and even more difficult to simulate with a computer.

Soap bubbles are essentially layers of immiscible fluids: air on the inside, a thin film of liquid, and then air on the outside again. The surface tension of the liquid drives the bubble toward a shape with less surface area, while air pressure forces the bubble to maintain a constant volume. These behaviors are described by the Navier-Stokes equations, which can be exceptionally difficult to solve accurately, especially when confronted with the large density jumps, immiscible fluid interfaces, stiff surface tension forces, and extremely thin liquid surfaces required for bubble motion.

Prior work has successfully simulated foam dynamics using an Eulerian approach. These approaches utilize considerable computational resources for the detailed calculation of the air dynamics within each bubble and the accurate resolution of film geometry. We propose a new model which solves the problem more economically while guaranteeing several important theoretical properties by construction.

Our proposed model is the first numerical method for the dynamics of soap films, bubbles, and foams that is based on the equations of non-manifold vortex sheets.

We model soap film structures as non-manifold vortex sheets driven by surface tension forces. We propose a discrete model to evolve the sheet as it deforms and undergoes topology changes such as pinching, merging, and rearrangement.

We represent the geometry using a Lagrangian non-manifold triangle mesh decorated with material region tags. We store *circulation* variables on the vertices; by Kelvin's circulation theorem this ensures a circulation-preserving model *by construction*. Numerical results exhibit nearly dissipation-free dynamics.

The use of the vortex sheet equations ensures an exactly divergence-free velocity field via the Biot-Savart law, and integrating mesh

Simulating Rigid Body Fracture with Surface Meshes

Yufeng Zhu*
University of British Columbia

Robert Bridson†
Autodesk, University of British Columbia

Chen Greif‡
University of British Columbia

Abstract

We present a new brittle fracture simulation method based on a boundary integral formulation of elasticity and recent explicit surface mesh evolution algorithms. Unlike prior physically-based simulations in graphics, this avoids the need for volumetric sampling and calculations, which aren't reflected in the rendered output. We represent each quasi-rigid body by a closed triangle mesh of its boundary, on which we solve quasi-static linear elasticity via boundary integrals in response to boundary conditions and loads such as impact forces and gravity. A fracture condition based on maximum tensile stress is subsequently evaluated at mesh vertices, while crack initiation and propagation are formulated as an interface tracking procedure in material space. Existing explicit mesh tracking methods are modified to support evolving cracks directly in the triangle mesh representation, giving highly detailed fractures with sharp features, independent of any volumetric sampling (unlike tetrahedral mesh or level set approaches); the triangle mesh representation also allows simple integration into rigid body engines. We also give details on our well-conditioned integral equation treatment solved with a kernel-independent Fast Multipole Method for linear time summation. Various brittle fracture scenarios demonstrate the efficacy and robustness of our new method.

CR Categories: G.1.9 [Mathematics of Computing]: Numerical Analysis—Integral Equations; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: dynamics, fracture, rigid body, boundary integrals, triangle mesh, Fast Multipole Method

1 Introduction

High-quality rigid body fracture has been a popular research topic in computer graphics, with the purpose of increasing realism and plausibility in computer-generated animation. Fracturing effects such as shattering glass, destruction of a cement wall or blowing up a space battleship greatly affect the immersive user experience in computer games, virtual reality and film industry. Roughly speaking, current state-of-the-art approaches can be categorized into two groups: (1) crack propagation and (2) volume decomposition. The first category uses a volume discretization to solve for the underlying elastic dynamics, and applies some crack growth criteria, such as principal stress hypothesis (Rankine hypothesis), critical strain




Figure 1: A glass goblet broken into multiple shards (left), with corresponding fracture propagation shown in the material space (right). Only the surface mesh is used for both the physical computation and fracture propagation.

energy release rate (Griffith criterion) and minimum strain energy density (Beltrami hypothesis), to propagate fractures inside the material. These methods are generally physically realistic, but entail significant computational cost due to either the poor scaling behavior of the physical system or the complexity of the unstructured volumetric grid's remeshing operation. The second category starts with spatial decomposition of the object instantly, instead of waiting for the crack propagation process to create new fractured pieces. Such a simplified strategy gives rise to computational efficiency and flexible control of fragment generation, which is essential and attractive to graphics experts, practitioners and users. However, the potentially unpredictable physically inconsistent behavior caused by this simplification might require significant post-processing.

Our goal is to design a new rigid body fracture algorithm that is physically plausible and computationally efficient. We present a surface-only method to handle both physical computation and fracture evolution by formulating the quasi-static elasticity problem as

*e-mail:mike323@cs.ubc.ca

†e-mail:rbridson@cs.ubc.ca

‡e-mail:greif@cs.ubc.ca

ACM Reference Format

Zhu, Y., Bridson, R., Greif, C. 2015. Simulating Rigid Body Fracture with Surface Meshes. ACM Trans. Graph. 34, 4, Article 150 (August 2015), 11 pages. DOI = 10.1145/2766942 <http://doi.acm.org/10.1145/2766942>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766942>

High-Resolution Brittle Fracture Simulation with Boundary Elements

David Hahn*
IST Austria

Chris Wojtan†
IST Austria




Figure 1: *Tearing the armadillo: (a) input geometry and illustration of boundary conditions; (b) BEM mesh with maximal principal stress shown in red; (c) BEM mesh of a generated fracture; (d) high-resolution fracture surface.*

Abstract

We present a method for simulating brittle fracture under the assumptions of quasi-static linear elastic fracture mechanics (LEFM). Using the boundary element method (BEM) and Lagrangian crack-fronts, we produce highly detailed fracture surfaces. The computational cost of the BEM is alleviated by using a low-resolution mesh and interpolating the resulting stress intensity factors when propagating the high-resolution crack-front.

Our system produces physics-based fracture surfaces with high spatial and temporal resolution, taking spatial variation of material toughness and/or strength into account. It also allows for crack initiation to be handled separately from crack propagation, which is not only more reasonable from a physics perspective, but can also be used to control the simulation.

Separating the resolution of the crack-front from the resolution of the computational mesh increases the efficiency and therefore the amount of visual detail on the resulting fracture surfaces. The BEM also allows us to re-use previously computed blocks of the system matrix.

CR Categories: G.1.9 [Mathematics of Computing]: Numerical Analysis—Integral Equations; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: boundary elements, brittle fracture, crack propagation

*e-mail: david.hahn@ist.ac.at

†e-mail: wojtan@ist.ac.at

ACM Reference Format

Hahn, D., Wojtan, C. 2015. High-Resolution Brittle Fracture Simulation with Boundary Elements. *ACM Trans. Graph.* 34, 4, Article 151 (August 2015), 12 pages. DOI = 10.1145/2766896
<http://doi.acm.org/10.1145/2766896>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA.
Copyright 2015 ACM 978-1-4503-3331-3/15/08 ... \$15.00.
DOI: <http://doi.acm.org/10.1145/2766896>

1 Introduction

Computer graphics researchers have developed diverse methods for fracturing virtual objects. Purely geometric approaches use either pre-fractured models or pre-defined fracture patterns, while fracture simulations (including mass-spring systems, finite element methods, and mesh-less methods) respect the underlying physics of the fracture process. Unfortunately, these simulation methods become inefficient if all the detail of the fracture surfaces is present in the computational model, effectively limiting the visual detail that can be captured. Adaptive re-meshing (or re-sampling in mesh-less methods) is typically used to mitigate this limitation. Some methods also use heuristics to add more visual detail to fracture surfaces as a post-process, but they do not influence the simulation in any way.

The standard approach in fracture simulation for computer graphics is to cut or re-mesh one element at a time as the crack propagates through the material. This requires small time-steps and is analogous to an *Eulerian* reference frame as the crack advances through space. Our approach departs from this traditional viewpoint by adopting a *Lagrangian* reference frame for crack propagation, as illustrated in Fig. 2. This new point of view allows us to utilize techniques from front tracking, improves the achievable resolution of fracture surfaces, and still treats the underlying physics with acceptable accuracy. We combine a boundary element method (BEM) on a coarse computational mesh with a high-resolution crack propagation scheme, resulting in a fast and efficient framework for brittle fracture. Our results show that this method is capable of producing detailed, realistic fracture surfaces while avoiding complicated mesh manipulations.

Furthermore, a common artifact in computer graphics simulations is *artificial shattering*, which results from the use of stress-based fracture criteria. We introduce a new method based on linear elastic fracture mechanics (LEFM) that avoids this problem altogether by treating crack initiation and propagation separately. Our method utilizes both *strength* and *toughness* in the physically correct way, where crack *initiation* depends on stress, but crack *propagation* is governed by stress intensity factors.

The main contributions of our method are:

- an efficient symmetric Galerkin BEM implementation for quasi-static brittle fracture simulation, where each entry in the system matrix is computed only once;

Improving Light Field Camera Sample Design with Irregularity and Aberration

Li-Yi Wei^{†‡§}

Chia-Kai Liang[†]

Graham Myhre[†]

Colvin Pitts[†]

Kurt Akeley[†]

Lytro Inc.[†]

Dragoniac[‡]

Univ. Hong Kong[§]

Abstract

Conventional camera designs usually shun sample irregularities and lens aberrations. We demonstrate that such irregularities and aberrations, when properly applied, can improve the quality and usability of light field cameras. Examples include spherical aberrations for the mainlens, and misaligned sampling patterns for the microlens and photosensor elements. These observations are a natural consequence of a key difference between conventional and light field cameras: optimizing for a single captured 2D image versus a range of reprojected 2D images from a captured 4D light field. We propose designs in mainlens aberrations and microlens/photosensor sample patterns, and evaluate them through simulated measurements and captured results with our hardware prototype.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Sampling

Keywords: light field, camera, sampling, imaging, noise, misalignment, irregularity, aberration, computational photography

1 Introduction

Traditional cameras directly capture the final images. Computational cameras, in contrast, capture the original data from which the final images are computed. Such original data often contains extra information enabling flexibility and quality not possible with only the final outputs. An example is light field cameras, capturing 4D radiance data which can be processed to form 2D images with varying parameters such as view-points and focus distances [Ng 2006].

The extra powers of computational cameras also bring extra design challenges. For a light field camera, it is crucial to design its 4D sample set formed by the 2D microlens and photosensor arrays to project good 2D distributions under different focus distances. As visualized in Figure 1 top, existing light field camera designs with regular alignment of microlens and photosensor arrays can exhibit highly variable 2D distributions under different projections. This is undesirable, as images formed or captured from less uniform 2D sample distributions can have far poorer quality than those with more uniform distributions. Often, photographers have to carefully setup the scene to ensure that all objects of interests fall within ranges of good sample distributions.

The analysis and synthesis of spatial samples have been extensively researched in computer graphics [Pharr and Humphreys 2004]. However, prior works predominantly focused on optimizing a sample set which will remain fixed in the application space, such as rendering or displaying an image [Cook 1986; Heide et al. 2014], reconstructing a surface [Öztireli et al. 2010], or animating a fluid




Figure 1: Sampling design for light field cameras. Each row shows projected 2D image sample sets under different focus distances λ from a single 4D light field sample set. $\lambda = 0$ and -1 maps to the samples on the microlens and photosensor planes, respectively (see Figure 2). For better imaging quality, projected 2D sample sets should have uniform spatial coverage across different focus distances λ , blue noise distributions for microlens and photosensor arrays with improved coverage at λ away from 0, misalignment of hex-microlens/square-sensor by rotation and pitch size mismatch for similarly improved coverage, further adding spherical aberration to the mainlens for improving coverage even at λ near 0.

sequence [Schechter and Bridson 2012]. Computing a sample set to be reused under various settings in a different space with unknown input content (e.g., one 4D light field \rightarrow multiple 2D images) is a more challenging and less studied problem.

We present designs and methods to construct 4D light field sample sets that can produce more uniform 2D distributions under different projections. Our main observation is that regular sampling (Figure 1, top), as in traditional cameras, can have samples clustered at certain focus distances, causing non-uniform spatial coverage. To avoid such issues, our key idea is to introduce sample irregularities (Figure 1, middle) and lens aberrations (Figure 1, bottom) into light field camera design. Our results indicate that simple designs, such as misalignments of the microlens and sensor arrays and spherical aberrations for the mainlens, are practical to implement and can effectively improve the quality and usability of light field cameras. Moreover, we apply this principle to design simpler and less expensive lens with higher performance.

To quantify the design objectives and evaluate the design variations, we present a mathematical formulation to measure the projected spatial coverage under different focus distances of given 4D light field sample sets. We analyze various design choices, and present methods to synthesize desired misalignments and aberrations. We have built a prototype camera to evaluate our design with various captured scenes.

In sum, the contributions of this paper include:

ACM Reference Format

Wei, L., Liang, C., Myhre, G., Pitts, C., Akeley, K. 2015. Improving Light Field Camera Sample Design with Irregularity and Aberration. *ACM Trans. Graph.* 34, 4, Article 152 (August 2015), 11 pages. DOI = 10.1145/276688 <http://doi.acm.org/10.1145/276688>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. SIGGRAPH '15 Technical Paper, August 09 – 13, 2015, Los Angeles, CA. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3331-3/15/08 ... \$15.00. DOI: <http://dx.doi.org/10.1145/276688>