

SIGGRAPH 2015 Course Notes / Materials

Submission Sample Guidelines

Course Notes Sample

- Sample
 - The course notes sample concisely demonstrates proposed published content for course proposal review
- Simple
 - The course notes sample only needs to be long enough to evaluate anticipated quality. Recommended length: 12-18 pages. Overly long submissions will not necessarily improve reviews.
- Sufficient
 - Reviewers will evaluate content for aspects such as clarity, visual quality, topic development and detail, ease of use, etc.
- Complete
 - In addition to the course notes sample, please outline any additional resources to be published with the course notes (such as animation files, source code, bibliographies, web resources, etc.)

Submission Checklist

- Samples should be submitted in Portable Document Format (PDF).
 - Important for uniform support of reviewers across many platforms.
- The excerpt should demonstrate development of a significant topic or theme from the syllabus.
 - If the course notes originate from your presentation slides, consider providing speaker annotations where a slide might require additional explanatory details.
 - Individual presenters may each provide a brief excerpt of their work as part of the review sample.
 - SIGGRAPH 2015 presentation slide templates will be available on the SIGGRAPH 2015 website.
- Provide a summary page of any additional materials intended for publication (images, video, source code, web sources, etc.) as a supplementary reference.

Past Examples

- The attached examples demonstrate the qualities of highly regarded course note samples
 - Jaroslav Krivanek, et. al: Recent Advances in Light Transport Simulation: Some Theory and a lot of Practice
 - Ann McNamara, et. al: Perceptually-Motivated Graphics, Visualization and 3D Displays
 - Note: any one example would be sufficient for submission review

**Recent Advances in Light Transport Simulation: Some
Theory and a lot of Practice
(Submission ID: general_0486)**

SIGGRAPH 2014 Course

Course Notes Sample

Complete course materials will be available from
<http://cgg.mff.cuni.cz/~jaroslav/papers/2014-ltscourse/index.htm>

Organizers

Jaroslav Křivánek
Charles University in Prague

Alexander Keller
NVIDIA

Lecturers

Iliyan Georgiev
Solid Angle, Saarland University

Anton S. Kaplanyan
Karlsruhe Institute of Technology

Marcos Fajardo
Solid Angle

Mark Meyer
PIXAR Animation Studios

Ondřej Karlík
Charles University in Prague, Render Legion s.r.o.

Juan Cañada
Next Limit Technologies

Abstract

We are witnessing a renewed research interest in robust and efficient light transport simulation based on statistical methods. This research effort is propelled by the desire to accurately render general environments with complex materials and light sources, which is often difficult with the industry-standard ad hoc solutions. In addition, it has been recognized that advanced methods, which are able to render many effects in one pass without excessive tweaking, increase artists productivity and allow them to focus on their creative work. For this reason, the movie industry is shifting away from approximate rendering solutions towards physically-based rendering methods, which poses new challenges in terms of strict requirements on high image quality and algorithm robustness.

Many of the recent advances in light transport simulation, such as new Markov chain Monte Carlo methods, the robust combination of bidirectional path tracing with photon mapping, or path space filtering are made possible by interpreting light transport as an integral in the space of light paths. However, there is a great deal of confusion among practitioners and researchers alike regarding these path space methods.

The main contribution of the theoretical part of the course is a coherent review of the path integral formulation of light transport and its applications, including the most recent ones. We show that rendering algorithms that may seem complex at first sight, are in fact naturally derived from this general framework. We also show that the path integral framework makes the extension of the surface-based algorithm to volumetric media extremely simple. The course includes an extensive empirical comparison of the various light transport algorithms. A substantial part of the course is then devoted to the application of advanced light transport simulation and path sampling methods in practical rendering tasks in architectural visualization and VFX.

Intended audience

Industry professionals and researchers interested in recent advances in robust light transport simulation for realistic rendering with global illumination.

Prerequisites

Familiarity with rendering and with basic concepts of global illumination computation is expected.

Level of difficulty

Intermediate

Syllabus

1. Introduction & Welcome (*Křivánek*)
(5 min)
2. Path Integral Formulation of Light Transport (*Křivánek*)
(25 min)
 - Light transport simulation as an integral over the space of light paths
 - Path sampling methods and path probability density
 - Combining different path sampling techniques
 - Bidirectional path tracing
 - Joint importance sampling in participating media
3. Combining Bidirectional Path Tracing and Photon Mapping (*Georgiev*)
(20 min)
 - (Progressive) photon mapping
 - Photon mapping as a path sampling technique
 - Combining with bidirectional path tracing techniques
 - Discussion: advantages, limitations, and best practices
4. Path Space Filtering (*Keller*)
(20 min)
 - Smoothing the path contributions before image reconstruction
 - Path space similarity
 - Algorithmic aspects
 - Example applications in light transport simulation, multi-view rendering, shadows, ambient occlusion, translucency, real-time rendering
5. Comparison of Various Light Transport Methods (*Kaplanyan*)
(20 min)
 - Comparison: Monte Carlo methods
 - Video: Femtosecond light propagation
 - Comparison: Metropolis light transport with different mutation strategies
 - Comparison: Energy redistribution path tracing
 - Comparison: Photon mapping with Metropolis sampling
 - Comparison: Population Monte Carlo light transport
- Break
(15 min)
6. Efficiency = Good Importance Sampling (*Fajardo/Georgiev*)
(20 min)
 - Importance sampling area lights
 - Importance sampling single scattering in participating media
 - BSSRDF importance sampling

7. Pixar's Fast Lighting Preview (*Meyer*)
(20 min)
 - Physically based lighting in computer animation
 - Technology and integration - production pipeline, OptiX and Katana
 - Showcase the pipeline using artistic controls on full production shots and assets
8. Corona Renderer: It's all about Usability (*Karlík*)
(20 min)
 - Users requirements in architectural visualization
 - Biased vs. unbiased rendering from the users perspective
 - Examples of solutions to specific issues: energy clamping, artistic tweaking
9. Advanced Light Transport in the VFX/Archviz industry (*Cañada*)
(20 min)
 - Stage of the Industry - the reasons for accurate light transport in practice
 - Current problems, solutions, and workarounds
 - What's next?
10. Conclusions / Q & A (*all*)
(10 min)

Course presenter information

Jaroslav Krivánek *Charles University in Prague*
jaroslav.krivanek@mff.cuni.cz

Jaroslav is an associate professor in Computer Science at Charles University in Prague. Prior to this appointment, he was a Marie Curie research fellow at Cornell University, and a junior researcher and assistant professor at Czech Technical University in Prague. Jaroslav received his Ph.D. from IRISA/INRIA Rennes and the Czech Technical University (joint degree) in 2005. In 2003 and 2004 he was a research associate at the University of Central Florida. His primary research interest is realistic rendering, global illumination, radiative transport (including light transport), and Monte Carlo methods.

Alexander Keller *NVIDIA*
keller.alexander@gmail.com

Alexander Keller is a senior research manager at NVIDIA Research and leads advanced rendering research. Before, he had been the Chief Scientist of mental images, where he had been responsible for research and the conception of future products and strategies including the design of the iray renderer. Prior to industry, he worked as a full professor for computer graphics and scientific computing at Ulm University, where he co-founded the UZWR (Ulmer Zentrum für wissenschaftliches Rechnen). Alexander Keller holds a Ph.D. in computer science, authored 25 granted patents, and published more than 50 papers mainly in the area of quasi-Monte Carlo methods and photorealistic image synthesis using ray tracing.

Iliyan Georgiev *Solid Angle, Saarland University*
georgiev@cs.uni-saarland.de

Iliyan is a rendering researcher and enthusiast. He holds a B.Sc. degree in computer science from Sofia University, Bulgaria, and a M.Sc. degree from Saarland University, Germany, with a Ph.D. degree anticipated in the near future. His primary topics of interest are high performance ray tracing and Monte Carlo methods for physically-based light transport simulation. His aspiration for practical rendering solutions has given him the opportunity to work for Disney, Weta Digital and Chaos Group (V-Ray). He currently works for Solid Angle.

Anton S. Kaplanyan *Karlsruhe Institute of Technology*
anton.kaplanyan@kit.edu

Anton S. Kaplanyan is a graphics researcher at Karlsruhe Institute of Technology (KIT), Germany. Additionally he is pursuing a Ph.D. title. His primary research and recent publications are about advanced light transport methods for global illumination. Prior to joining academia Anton had been working at Crytek for three years at various positions from senior R&D graphics engineer to lead researcher. He received his

M.Sc. in Applied Mathematics at National Research University of Electronic Technology, Moscow in 2007.

Marcos Fajardo *Solid Angle*
marcos@solidangle.com

Marcos is the founder of Madrid-based Solid Angle SL, where he leads the research and development team working on the Arnold path tracing renderer. Previously he was a visiting software architect at Sony Pictures Imageworks, a visiting researcher at USC Institute for Creative Technologies under the supervision of Dr. Paul Debevec, and a software consultant at various CG studios around the world. He studied Computer Science at University of Málaga, Spain.

Mark Meyer *PIXAR Animation Studios*
mmeyer@pixar.com

Mark Meyer is a Senior Scientist at Pixar Animation Studios. He received his BS in Computer Science and Computer Engineering from Northwestern University and his Ph.D. from Caltech. Before joining Pixar in 2003, Mark worked on virtual reality and simulation at Argonne National Laboratory and instructed Computer Graphics courses in the Computer Science department at Caltech. Mark is currently working in Pixar's Research Group on projects ranging from character articulation to lighting acceleration.

Ondřej Karlík *Charles University in Prague, Render Legion s.r.o.*
ondra@corona-renderer.com

After graduating from the Czech Technical University, Ondřej became a PhD student at the Charles University in Prague, where he has cooperated with Jaroslav Krivánek on the research of usability in realistic rendering. In 2009 he started the development of his own photorealistic renderer, Corona, which have since then became one of the fastest growing renderers in the architectural visualization field. He is currently preparing the first commercial release of the Corona renderer.

Juan Cañada *Next Limit Technologies*
juan.canada@nextlimit.com

Juan joined Next Limit in 2001 to work in the Realflow development team and later he moved to the newborn Maxwell research team. Since then Juan held several positions in the team, leading it since 2007. He holds a bachelors degree in Mechanical Engineering and a degree in Environmental Sciences.

Jaroslav Křivánek:

**Path Integral Formulation of
Light Transport**

PATH SAMPLING & PATH PDF



- So how exactly do we sample the paths and how do we compute the path PDF?

Local path sampling

- Sample one path vertex at a time




1. From an a priori distribution

- lights, camera sensors

2. Sample direction from an existing vertex

3. Connect sub-paths


- test visibility between vertices



Course: Recent Advances in Light Transport Simulation
Jaroslav Krivánek - Path Integral Formulation of Light Transport

- Most of the practical algorithms rely on local path sampling, where paths are built by adding one vertex at a time until a complete path is built.
- There are three common basic operations.
 - First, we can sample a path vertex from an a priori given distribution over scene surfaces. We usually employ this technique to start a path either on a light source or on the camera sensor.
 - Second, given a sub-path that we've already sampled with a vertex at its end, we may sample a direction from that vertex, and shoot a ray in this direction to obtain the next path vertex.
 - Finally, given two sub-paths, we may connect their end-vertices to form a full light transport path. This technique actually does not add any vertex to the path. It is more or less a simple visibility check to see if the contribution function of the path is non-zero.

Example – Path tracing




Course: Recent Advances in Light Transport Simulation
Jaroslav Křivánek - Path Integral Formulation of Light Transport

3

- Let us see how these three basic operations are used in a simple path tracer.
- First, we generate a vertex on the camera lens, usually from a uniform distribution over the lens surface – so this corresponds to operation 1.
- Second, we pick a random direction from this point such that it passes through the image plane, and shoot a ray to extend the path. This is operation 2.
- Third, we may generate an independent point on the light source (operation 1) and test visibility (operation 3) to form a complete light transport path.
- We could also continue the path by sampling a random direction and shooting a ray (operation 2), and eventually hit the light source to complete the path.
- An important thing to notice is that one single primary ray from the camera actually creates a full family of light transport paths. These paths are correlated, because they share some of the path vertices, but they are distinct entities in the path space.

Use of local path sampling



Course: Recent Advances in Light Transport Simulation
Jaroslav Křivánek - Path Integral Formulation of Light Transport

4


- These same basic operations are used to construct paths in light tracing and bidirectional path tracing.

Probability density function (PDF)

path PDF

$$\underline{p(\bar{x})} = \underline{p(x_0, \dots, x_k)}$$

joint PDF of path vertices



Course: Recent Advances in Light Transport Simulation
Jaroslav Křivánek - Path Integral Formulation of Light Transport

5


- Not that we know how to construct a path, we need to evaluate its PDF so that we can plug it into the MC estimator.
- In general the PDF of a light path is simply the **joint** PDF of the path vertices.
- That is to say, the PDF that the first vertex is where it is *and* the second vertex is where it is, etc.

Probability density function (PDF)

path PDF

$$\underline{p(\bar{x})} = \underline{p(x_0, \dots, x_k)}$$

joint PDF of path vertices



Course: Recent Advances in Light Transport Simulation
Jaroslav Křivánek - Path Integral Formulation of Light Transport


6

Probability density function (PDF)

path PDF

$$\underbrace{p(\bar{x})}_{\text{joint PDF of path vertices}} = \underbrace{p(x_0, \dots, x_k)}_{\text{joint PDF of path vertices}} = \underbrace{\begin{matrix} p(x_3) \\ p(x_2 | x_3) \\ p(x_1 | x_2) \\ p(x_0) \end{matrix}}_{\text{product of (conditional) vertex PDFs}}$$

Path tracing example:



Course: Recent Advances in Light Transport Simulation
Jaroslav Krivánek - Path Integral Formulation of Light Transport

7


- The joint path PDF is given by the product of the conditional vertex PDF.
- To see what this means, let us again take the example of path tracing, where we build a path starting from the camera.
- Vertex x_3 comes from an a priori distribution $p(x_3)$ over the camera lens (usually uniform; or the delta distribution for a pinhole camera).
- Vertex x_2 is sampled by generating a random direction from x_3 and shooting a ray. This induces a PDF for x_2 , $p(x_2 | x_3)$, which is in fact conditional on vertex x_3 .
- The same thing holds for vertex x_1 , which is sampled by shooting a ray in a random direction from x_2 .
- Finally, vertex x_0 on the light source might be sampled from an uniform distribution over the light source area with pdf $p(x_0)$, independently of the other path vertices.
- The full joint PDF is given by the product of all these individual terms.

Probability density function (PDF)

path PDF

$$\underbrace{p(\bar{x})}_{\text{joint PDF of path vertices}} = \underbrace{p(x_0, \dots, x_k)}_{\text{joint PDF of path vertices}} = \underbrace{p(x_3) p(x_2) p(x_1) p(x_0)}_{\text{product of (conditional) vertex PDFs}}$$

Path tracing example:



Course: Recent Advances in Light Transport Simulation
Jaroslav Křivánek - Path Integral Formulation of Light Transport

8

- It is customary to simplify this somewhat pedantic notation and leave out the conditional signs. Nonetheless, it is important to keep in mind that the path vertex PDFs for vertices that are not sampled independently are indeed conditional PDFs.

Vertex sampling

■ Importance sampling principle

1. Sample from an a priori distrib.

emission
sampling

2. Sample direction from an existing vertex

BRDF lobe
sampling

3. Connect sub-paths


high throughput
connections

Course: Recent Advances in Light Transport Simulation
Jaroslav Krivánek - Path Integral Formulation of Light Transport

- In accordance with the principle of importance sampling, we want to generate full paths from a distribution with probability density proportional to the measurement contribution function. That is, high-contribution paths should have proportionally high probability of being sampled.
- Local path sampling takes an approximate approach, where each local sampling operation tries to importance sample the terms of the contribution function associated with the vertex being sampled.
- For example, when starting a path on the light source, we usually sample the initial vertex from a distribution proportional to the emitted power.
- When extending the path from an existing vertex, we usually sample the random direction proportionally to the BRDF at the vertex.
- Similarly, when connecting two sub-paths with an edge, we may want to prefer connections with high throughput (though this is rarely done in practice).


Vertex sampling

- Sample direction from an existing vertex



Measure conversion

- Sample direction from an existing vertex




$$p(y) = p^\perp(x \rightarrow y) G(x \leftrightarrow y)$$

w.r.t. area

w.r.t. proj.
solid angle

$$\langle I_j \rangle = \frac{f_j(\bar{x})}{p(\bar{x})}$$

$$= \frac{\dots \rho_s(x \rightarrow y) G(x \leftrightarrow y) \dots}{\dots p^\perp(x \rightarrow y) G(x \leftrightarrow y) \dots}$$



Course: Recent Advances in Light Transport Simulation
Jaroslav Krivánek - Path Integral Formulation of Light Transport

11

- There is one important technical detail associated with computing the path PDF for vertices created by direction sampling.
- The path integral is expressed with respect to the surface area measure – we are integrating over the surface of the scene – but the direction sampling usually gives the PDF with respect to the (projected) solid angle.
- The conversion factor from the projected solid angle measure to the area measure is the geometry factor.
- This means that any vertex generated by first picking a direction and then shooting a ray has the geometry factor of the generated edge importance sampled – the only geometry factor that are not importance sampled actually correspond to the connecting edges (operation 3 in local path sampling).

Perceptually-Motivated Graphics, Visualization and 3D Displays

Ann McNamara*

Department of Visualization
Texas A&M University

Katerina Mania[†]

Department of Electronic & Computer Engineering
Technical University of Crete

Marty Banks[‡]

Visual Space Perception Laboratory
University of California, Berkeley

Christopher Healey[§]

Department of Computer Science
North Carolina State University

SAMPLE NOTES #, February 18, 2010

*ann@viz.tamu.edu

[†]mania@ced.tuc.gr

[‡]Email. ramli@ramli.com

[§]healey@csc.ncsu.edu

Abstract

This course presents timely, relevant examples on how researchers have leveraged perceptual information for optimization of rendering algorithms, to better guide design and presentation in (3D stereoscopic) display media, and for improved visualization of complex or large data sets. Each presentation will provide references and short overviews of cutting-edge current research pertaining to that area. We will ensure that the most up-to-date research examples are presented by sourcing information from recent perception and graphics conferences and journals such as ACM Transactions on Perception, paying particular attention work presented at the 2010 Symposium on Applied Perception in Graphics and Visualization.

About the Lecturers

Ann McNamara

Department of Visualization
Texas A&M University
3137 TAMU
College Station, TX 77843-3137

+1-979-845-4715
ann@viz.tamu.edu
<http://www.viz.tamu.edu/people/ann>

Ann McNamara received her undergraduate and graduate degrees from the University of Bristol, UK. Anns research focuses on the advancement of computer graphics and scientific visualization through novel approaches for optimizing an individuals experience when creating, viewing and interacting with virtual spaces. She investigates new ways to exploit knowledge of human visual perception to produce high quality computer graphics and animations more efficiently. She joined the faculty of the newly formed Department of Visualization at Texas A&M University in 2008, where she is currently an assistant professor. Ann serves on several IPCs including APGV.

Katerina Mania

Department of Electronic and Computer Engineering
Technical University of Crete
University Campus Kounoupidiana Chania Crete Greece

+30 28210 37222
k.mania@ced.tuc.gr
<http://www.music.tuc.gr/kmania>

Katerina Mania completed a B.Sc. in Mathematics, University of Crete, Greece, an M.Sc./Ph.D in Computer Science, University of Bristol, UK, funded by HP Labs. She worked at HP Labs as a researcher before serving on Faculty in the Department of Informatics, University of Sussex. Katerina spent her sabbatical at NASA Ames Research Centre (Advanced Displays and Spatial Perception Laboratory) in 2003. She is currently an Assistant Professor with tenure at the Technical University of Crete, Greece. Katerina is the program co-chair for APGV 2010 and Associate Editor of ACM Transactions on Applied Perception and Presence Teleoperators and Virtual Environments.

Marty Banks

Marty Banks' Lab
University of California, Berkeley
Vision Science, 360 Minor Hall
Berkeley, CA 94720-2020

+1-510-642-7679
martybanks@berkeley.edu
<http://bankslab.berkeley.edu/>

Martin S. Banks received his Bachelors degree at Occidental College (1970). After one year in Germany teaching, he entered the graduate program in Psychology at UC San Diego. He received a Masters degree in Experimental Psychology (1973). Banks then transferred to the graduate program at the University of Minnesota where he received his PhD. in Developmental Psychology (1976). He was Assistant & Associate Professor of Psychology at the University of Texas at Austin (1976-1985). He moved to UC Berkeley School of Optometry in 1985 where he is now a Full Professor of Optometry and Vision Science.

Christopher Healey

Department of Computer Science
North Carolina State University
890 Oval Drive #8206

+1 919.513.8112
healey@csc.ncsu.edu
<http://www.csc.ncsu.edu/faculty/healey>

Christopher G. Healey received a B.Math from the University of Waterloo in Waterloo, Canada, and an M.Sc. and Ph.D. from the University of British Columbia in Vancouver, Canada. Following a postdoctoral fellowship at the University of California at Berkeley, he joined the Department of Computer Science at North Carolina State University, where he is currently an Associate Professor. His research interests include visualization, graphics, visual perception, and areas of applied mathematics, databases, artificial intelligence, and aesthetics related to visual analysis and data management.

Course Overview

5 minutes: Welcome and Introductions

Ann McNamara

Welcome, overview of course and motivation for attending.
Speaker Introductions

40 minutes: Depth Perception and 3D Displays

Martin Banks

An overview of Depth Perception and important phenomenon when presenting information on 3D Displays

40 minutes: Visualization

Chris Healey

A look at Visual Attention, Visual Memory, and its Role in Visualization.

15 minutes: Break

30 minutes: Perceptually Motivated Rendering

Ann McNamara

Overview of how knowledge from perceptual research feeds into optimized rendering algorithms.

30 minutes: Simulation and Virtual Environments

Katerina Mania

Perceptually-based Optimizations & Fidelity Metrics for Simulation Technology

30 minutes: Leading-edge research and APGV 2010

Katerina Mania

A summary of cutting edge perceptual research selected from APGV 2010 presentations.

10 minutes: A look to the future

Ann McNamara & Katerina Mania

Discussion of trends for APGV 2010

10 minutes: Conclusion, Questions & Answers

All

Wrap up, review, questions and discussion.

Contents

1	Introduction	9
1.1	Motivation	9
1.2	Course Overview	9
1.3	Focus Areas	9
1.4	Exploitation of the limitations of the HVS to reduce rendering times	9
1.5	Exploration of incorporating perceptual and cognitive aspects to Virtual Environments (VEs).	10
1.6	Visualization	10
1.7	Stereoscopic Displays	10
1.8	Summary	10
2	Perceptually Motivated 3D Displays & Depth Perception	11
2.1	Introduction	11
3	Perceptually Motivated Visualization	12
3.1	Introduction	12
3.2	Overview	12
3.3	Visual Attention and Preattentive Processing	12
3.4	Theories of Preattentive Processing	13
3.5	Feature Integration	14
3.6	Textons	14
3.7	Guided Search	14
3.8	Boolean Maps	15
3.9	Ensemble Coding	15
3.10	Feature Hierarchies	16
3.11	Visual Memory	16
3.12	Change Blindness	17
3.13	Inattentional Blindness	17
3.14	Attentional Blink	18
3.15	Conclusions	18
4	Perceptually Motivated Rendering	19
4.1	Visual Perception in Realistic Image Synthesis	19
4.2	Visual Perception	19
4.3	The Human Visual System	20
4.4	Contrast	21
4.5	Constancy	21
4.6	Human Visual Perception	22
4.7	Lightness Perception	23
4.8	Perception and Graphics	24
4.9	Perceptually driven rendering	24
4.10	Perceptually Based Image Quality Metrics	27
4.11	Dalys Visible Differences Predictor	27

4.12 Conclusion	28
5 Perceptually Motivated Virtual Environments	29
5.1 Introduction	29
5.2 Psychophysics	29
5.3 Spatial Awareness in 3D Immersive Worlds	30
5.4 Perceptually-optimized selective rendering algorithms	31
5.5 Conclusions	32
6 Leading Research	33
6.1 APGV Session Overviews	33
6.2 Session 1: Virtual Environments	33
6.3 Session 2: Image Synthesis	33
6.4 Session 3: Display Technology	33
7 Conclusions and Look to the future	33
8 Sample Slides	34
9 Sample Bibliography	35

1 Introduction

1.1 Motivation

The (re) introduction of 3D cinema, advent of affordable stereoscopic display technology, and seamless integration of real-world scenes with computer graphics fuels our continuing ability to create and display stunning realistic imagery. With the arrival of new technology, algorithms and display methods comes the realization that gains can be made by tailoring output to the intended audience; humans. Human beings have an amazingly complex perceptual systems, which have the ability to quickly capture and process vast amounts of complex data. With all its capability however, the Human Visual System (HVS) has some surprising nuances and limitations that can be exploited to the benefit of numerous graphics applications. This new tutorial will provide insight into those aspects of the HVS and other perceptual systems that can serve as both a guide and yard-stick to further the development and evaluation of computer graphics imagery and presentations. The literature on perception provides a rich source of knowledge that can be applied to the realm of computer graphics for immediate and direct benefit, generating images that not only exhibit higher quality, but use less time and resources to process. In addition, knowledge of the HVS serves as a guide on how best to present the images to fulfill the application at hand.

1.2 Course Overview

We will present timely, relevant examples on how researchers have leveraged perceptual information for optimization of rendering algorithms, to better guide design and presentation in (3D stereoscopic) display media, and for improved visualization of complex or large data sets. Each section will provide references and short overviews of cutting-edge current research pertaining to that area. We will ensure that the most up-to-date research examples are presented by sourcing information from recent perception and graphics conferences and journals such as ACM Transactions on Perception, paying particular attention work presented at the 2010 Symposium on Applied Perception in Graphics and Visualization.

1.3 Focus Areas

We will focus on four key areas in which perceptual knowledge has been successfully interleaved with computer graphics.

1.4 Exploitation of the limitations of the HVS to reduce rendering times

while improving resulting image quality. This includes real-time and non-real time graphics, image quality metrics and high dynamic range imagery.

1.5 Exploration of incorporating perceptual and cognitive aspects to Virtual Environments (VEs).

Such principles could be applied to selective real-time rendering algorithms, positive transfer of training as well as to optimizations for latency degradations and predictive tracking.

1.6 Visualization

Discussion of recent research pertaining to psychophysics and application to scientific and information visualization. A closer look at visual attention and visual memory will provide the framework for steering perceptually informed visualizations.

1.7 Stereoscopic Displays

3D stereoscopic displays are being used in a wide range of fields. To understand how better to present information on such displays, a comprehensive understanding of depth perception is necessary. This area will focus on depth perception and applications of such to image presentation.

1.8 Summary

In summary, this course represents a whirlwind tour of insights into how the eye and brain capture and process visual information through our perceptual systems, and how we can use those insights to further advance many areas in computer graphics.

PART I: Marty Banks

2 Perceptually Motivated 3D Displays & Depth Perception

2.1 Introduction

The human visual system has evolved in an environment with constrained relationships between objects and retinal images. That relationship is often altered in stereoscopic displays, so it is important to understand the situations in which the alteration is insignificant and the situations in which it causes undesirable perceptual or ergonomic effects. I will review the current literature on visual perception and human ergonomics in the context of the viewing of stereo displays. The literature shows that stereo displays can be associated with viewer fatigue/discomfort, reduced visual performance, and distorted 3D perception. This section will also discuss ways to minimize these adverse viewer effects.

SAMPLE NOTES - COMPLETE SET WILL BE PROVIDED ON ACCEPTANCE

PART II: Christopher Healey

3 Perceptually Motivated Visualization

3.1 Introduction

Human perception plays an important role in the area of visualization. An understanding of perception can significantly improve both the quality and the quantity of information being displayed. The importance of perception has been cited by numerous visualization panels and workshops.

This section summarizes some of the recent developments in research and theory regarding human psychophysics, and discusses their relevance to scientific and information visualization. We begin with an overview of the way human vision rapidly and automatically categorizes visual images into regions and properties based on simple computations that can be made in parallel across an image. This is often referred to as preattentive processing. We describe various theories of preattentive processing, and briefly discuss related work on ensemble coding and feature hierarchies. We next explain how these perceptual theories can impact visualization design and implementation.

We next examine several recent areas of research that focus on the critical role that the viewers current state of mind plays in determining what is seen, specifically, change blindness, inattention blindness, and the attentional blink. These phenomena offer a perspective on early vision that is quite different from the older view that early visual processes are reflexive and inflexible. Instead, they highlight the fact that what we see depends critically on where attention is focused and what is already in our minds prior to viewing an image. We discuss why these perceptual phenomena must be considered during visualization.

3.2 Overview

A fundamental goal of visualization is to produce images of data that support visual analysis, exploration and discovery, and identifying novel insights. An important consideration during visualization design is the role of human visual perception. How we see details in an image can directly impact a users efficiency and effectiveness. This article surveys research on attention and visual perception, with a specific focus on results that have direct relevance to visualization and visual analytics. We discuss theories of low-level visual perception, then show how these findings form a foundation for more recent work on visual memory and visual attention.

3.3 Visual Attention and Preattentive Processing

For many years vision researchers have been investigating how the human visual system analyzes images. An important initial result was the discovery of a limited set of visual properties that are detected very rapidly by low-level and fast-acting visual processes. These properties were initially called preattentive,

since their detection seemed to precede focused attention. We now know that attention plays a critical role in what we see, even at this early stage of vision. The term preattentive continues to be used, however, since it conveys an intuitive notion of the speed and ease with which these properties are identified. Typically, tasks that can be performed on large multi-element displays in less than 200 milliseconds (msec) are considered preattentive. Eye movements take at least 200 msec to initiate, and random locations of the elements in the display ensure that attention cannot be prefocused on any particular location, yet viewers report that these tasks can be completed with very little effort. This suggests that certain information in the display is seen in parallel by low-level visual processes.

A simple example of a preattentive task is the detection of a red circle in a group of blue circles. The target object has a visual property red that the blue distractor objects do not. A viewer can tell at a glance whether the target is present or absent. Here the visual system identifies the target through a difference in hue, specifically, a red target in a sea of blue distractors. Hue is not the only visual feature that is preattentive. For example, viewers can just as easily find a red circle in a background of red squares. Here, the visual system identifies the target through a difference in curvature (or form).

A unique visual property in the target—a red hue or a curved form—allows it to pop out of a display. A conjunction target made up of a combination of non-unique features normally cannot be detected preattentively. For example, consider combining the two backgrounds and searching for a red circle in a sea of blue circles and red squares. The red circle target is made up of two features: red and circular. One of these features is present in each of the distractor objects—red squares and blue circles. The visual system has no unique visual property to search for when trying to locate the target. A search for red items always returns true because there are red squares in each display. Similarly, a search for circular items always sees blue circles. Numerous studies have shown that a conjunction target cannot be detected preattentively. Viewers must perform a time-consuming serial search through the display to confirm its presence or absence.

If low-level visual processes can be harnessed during visualization, it can draw attention to areas of potential interest in a display. This cannot be accomplished in an ad-hoc fashion, however. The visual features assigned to different data attributes—the data-feature mapping—must take advantage of the strengths of our visual system, must be well-suited to the analysis needs of the viewer, and must not produce visual interference effects (e.g., conjunction search) that could mask information.

3.4 Theories of Preattentive Processing

A number of theories have been proposed to explain how preattentive processing occurs within the visual system: feature integration, textures, guided search, and boolean maps. We provide an overview of these theories, then discuss briefly feature hierarchies, which describes situations where the visual system

favors certain visual features over others, and ensemble coding, which shows that viewers can generate summaries of the distribution of visual features in a scene, even when they are unable to locate individual elements based those same features.

3.5 Feature Integration

Anne Treisman was one of the original researchers to document the area of preattentive processing. In order to explain the phenomena, Treisman proposed a model low-level human vision made up of a set of feature maps and a master map of locations. Each feature map registers activity for a specific visual feature. Treisman suggested a manageable number of feature maps, including one for each of the opponent colors, as well as separate maps for orientation, shape, and texture. When the visual system first sees an image, all the features are encoded in parallel into their respective maps. A viewer can access a particular map to check for activity, and perhaps to determine the amount of activity. The individual feature maps give no information about location, spatial arrangement, or relationships to activity in other maps, however.

3.6 Textons

Bela Julesz was also instrumental in expanding our understanding of what we see in an image. Julesz initially focused on statistical analysis of texture patterns. His goal was to determine whether variations in a particular order statistic were detected by the low-level visual system, for example contrast a first-order statistic orientation and regularity a second-order statistic and curvature a third-order statistic. Based on these findings, Julesz suggested that the early visual system detects a group of features called textons, which fall into three general categories:

Elongated blobs line segments, rectangles, or ellipses with specific properties of hue, orientation, width, and so on.

Terminators ends of line segments.

Crossings of line segments

Julesz believed that only a difference in textons or in their density could be detected preattentively. No positional information about neighboring textons is available without focused attention. Like Treisman, Julesz suggested that preattentive processing occurs in parallel and focused attention occurs in serial.

3.7 Guided Search

More recently, Jeremy Wolfe has proposed a theory that he calls guided search. He hypothesized that an activation map based on both bottom-up and top-down information is constructed during visual search. Attention is drawn to peaks in the activation map that represent areas in the image with the largest combination of bottom-up and top-down influence.

As with Treisman, Wolfe believes early vision divides an image into individual feature maps. In his theory, there is one map for each feature type: a color map, an orientation map, and so on. Within each map a feature is filtered into multiple categories. Bottom-up activation follows feature categorization. It measures how different an element is from its neighbors. Top-down activation is a user-driven attempt to find items with a specific property or set of properties. The activation map is a combination of bottom-up and top-down activity. Hills in the activation map mark regions that generate relatively large amount of bottom-up or top-down influence, but without providing information about the source of a hill. A subject's attention is drawn from hill to hill in order of decreasing activation.

3.8 Boolean Maps

A more recent model of low-level vision has been presented by Huang et al. This theory carefully divides visual search into two parts: selection and access. Selection involves choosing a set of objects from a scene. Access determines what properties of the selected objects a viewer can apprehend. Although both operations are implicitly present in previous theories, they are often described as a whole and not as separate steps.

Huang et al. suggest that the visual system can divide a scene into exactly two parts: selected elements and excluded elements. This is the boolean map that underlies their theory. The visual system can then access certain properties of the selected elements in the map. Once a boolean map is created, two properties are available to a viewer: the label for any feature in the map, and the spatial location of the selected elements. Boolean maps can be created in two ways. First, a viewer can specify a single value of an individual feature to select all objects that contain that feature. Second, union or intersection can be applied to two existing maps. In either case, only the result is retained, since evidence suggests that a viewer can only hold and access one boolean map at a time. Viewers can chain these operations together to search for targets in a fairly complex scene.

3.9 Ensemble Coding

Existing characterizations of preattentive vision have focused on how low level-visual processes can be used to guide attention to specific location or object in a larger scene. An equally important characteristic of low-level visual processes is their ability to generate a quick summary of how simple visual features are distributed across the field of view. The ability of humans to register a rapid and in-parallel summary of a scene in terms of its simple features was first reported by Ariely. He demonstrated that observers could extract the average size of a large number of dots from only a single glimpse at a display. Yet, when observers were tested on the same displays and asked to indicate whether a single dot of a given size was present, they were unable to do so. This suggests that there is a preattentive mechanism that records summary statistics of visual features

without retaining information about the constituent elements that generated the summary.

This ability to rapidly identify scene-based averages may offer important advantages in certain visualization environments. For example, given a stream of real-time data, ensemble coding would allow viewers to observe the stream at a high frame rate, yet still identify individual frames with interesting distributions of visual features (i.e. attribute values). Ensemble coding would also be critical for any situation where viewers want to estimate the amount of a particular data attribute in a display. These capabilities were hinted at in a paper by Healey et al., but without the benefit of ensemble coding as a possible explanation.

3.10 Feature Hierarchies

One promising strategy for multidimensional visualization is to assign different visual features to different data attributes. This allows multiple data values to be shown simultaneously in a single image. A key requirement of this method is a data-feature mapping that does not produce visual interference. One example of interference is a conjunction target. Another example is the presence of feature hierarchies that appears to exist in the visual system. For certain tasks one visual feature may be more salient than another. Researches in psychophysics and visualization have demonstrated a hue-shape hierarchy: the visual system favors color over shape. Background variations in hue interfere with a viewers ability to identify the presence of individual shapes and the spatial patterns they form. If hue is held constant across the display, these same shape patterns are immediately visible. The interference is asymmetric: random variations in shape have no effect on a viewers ability to see color patterns. Similar luminance-hue and hue-texture hierarchies have also been identified.

3.11 Visual Memory

Preattentive processing asks in part: What visual properties draw our eyes, and therefore our focus of attention to a particular object in a scene? An equally interesting question is: What do we remember about an object or a scene when we stop attending to it and look at something else? Many viewers assume that as we look around us we are constructing a high-resolution, fully detailed description of what we see. Researchers in psychophysics have known for some time that this is not true. In fact, in many cases our memory for detail between glances at a scene is very limited. Evidence suggests that a viewers current state of mind can play a critical role in determining what is seen and what is not.

We present three theories that demonstrate and attempt to explain this phenomena: change blindness, inattentional blindness, and attentional blink. Understanding what we remember as we focus on different parts of a visualization is critical to designing visualizations that encourage locating and retaining the information that is most important to the viewer.

3.12 Change Blindness

New research in psychophysics has shown that an interruption in what is being seen a blink, an eye saccade, or a blank screen renders us blind to significant changes that occur in the scene during the interruption. This change blindness phenomena can be illustrated using a task similar to one shown in comic strips for many years. A viewer is shown two pairs of images. A number of significant differences exists between the images. Many viewers have a difficult time seeing any difference and often have to be coached to look carefully to find it. Once they discover it, they realize that the difference was not a subtle one. Change blindness is not a failure to see because of limited visual acuity; rather, it is a failure based on inappropriate attentional guidance. Some parts of the eye and the brain are clearly responding differently to the two pictures. Yet, this does not become part of our visual experience until attention is focused directly on the objects that vary.

The presence of change blindness has important implications for visualization. The images we produce are normally novel for our viewers, so prior expectations cannot be used to guide their analyses. Instead, we strive to direct the eye, and therefore the mind, to areas of interest or importance within a visualization. This ability forms the first step towards enabling a viewer to abstract details that will persist over subsequent images.

3.13 Inattentional Blindness

A related phenomena called inattentional blindness suggests that viewers fail to perceive objects or activities that occur outside of the focus of attention. This phenomena is illustrated through an experiment conducted by Neisser. His experiment superimposed video streams of two basketball games. Players wore white shirts in one stream and black shirts in the other. Subjects attended to one team either white or black and ignored the other. Whenever the subjects team made a pass, they were told to press a key. After about 30 seconds of video, a third stream was superimposed showing a woman walking through the scene with an open umbrella. The stream was visible for about 4 seconds, after which another 25 seconds of basketball video was shown. Following the trial, only six of twenty-eight naive observers reported seeing the woman. When subjects only watched the screen and did not count passes, 100% noticed the woman.

Additional issues with relevance to visualization are also being investigated. Most et al. are studying the relationship between inattentional blindness and attentional capture, the ability of an object to draw the focus of attention without a viewers active participation. Researchers are also studying how perceptual load affects inattentional blindness. Finally, results suggest meaningful objects (e.g., a persons name or a happy face icon) may be easier to notice.

3.14 Attentional Blink

In each the previous methods for studying visual attention, the primary emphasis is on how human attention is limited in its ability to represent the details of a scene (change blindness) and in its ability to represent multiple objects at the same time (inattention blindness). But attention is also severely limited in its ability to process information that arrives in quick succession, even when that information is presented at a single location in space. The attentional blink paradigm is currently the most widely used method to study the availability of attention across time. Its name derives from the finding that when two targets are presented in rapid succession, the second of the two targets cannot be detected or identified when it appears within approximately 100-500 msec following the first target. This suggests that that attention operates over time like a window or gate, opening in response to finding a visual item that matches its current criterion or template and then closing shortly thereafter to consolidate that item as a distinct object or event from others. The attentional blink is an index of the dwell-time needed to consolidate a rapidly presented visual item into visual short term memory.

3.15 Conclusions

This presentation surveys past and current theories of low-level visual perception and visual attention. Initial work in preattentive processing identified basic visual features that can implicitly or explicitly capture a viewer's focus of attention. More recent work has extended this to study limited visual memory for change blindness and attentional blink and being blind to objects that are outside the focus of attention in inattention blindness. Each of these phenomena have significant consequences for visualization. We strive to produce images that are salient and memorable, and that guide attention to locations of importance within the data. Understanding what the visual sees and does not see is critical to designing effective visual displays.