

Steven J. Gortler; Research Statement

Computer Graphics is an important, fascinating, and very active field in Computer Science that has revolutionized many endeavors in entertainment, design, education, computer human interaction, and medicine. It is a field drawing on algorithm design and system building from the field of computer science, numerical computation from applied mathematics, simulation from applied physics and signal processing from electrical engineering. It is a growing field with one (very competitive) main conference, and a host of very healthy smaller workshops and seminars. In fact, computer graphics forms the single largest special interest subgroup of the Association for Computing Machinery (the primary computer science professional organization).

My research has focused on designing fundamental representations for computer graphics, as well as pursuing a mathematically rigorous understanding of computer graphics algorithms and techniques. This theoretical focus has been at the heart of all my research, from my early work in efficient lighting simulation, to my current work on image-based rendering and discrete geometry processing.

Image Based Rendering

Image based rendering (IBR) concerns itself with the process of capturing the appearance of objects and environments using standard video equipment and then being able to interact with the captured objects as if they were virtual objects in a computer graphics setting.

IBR offers clear benefits to computer graphics practice. In particular, one of the hardest parts of computer graphics is the creating or authoring realistic geometric models, as well as representations of their material and lighting properties. If this information can be captured as raw video data, and properly manipulated, then this authoring process can be greatly simplified.

Moreover, one can use these techniques to effectively archive the 3D appearance of actual scenes and objects. In fact, as standard camera, display, communication, and processing technology improves, we could see some IBR methods as being the foundation for a sort of 3D television system. In such a system, all of the “visual information” about some real scene can be simultaneously captured and communicated to a 3D viewing device.

One of the key tradeoffs in designing an IBR system is figuring out how much analysis must be done on the raw video data in order to use it in an interactive setting. (Many of these analytic techniques come from the related field of computer vision). Each step of the analysis creates a more abstract, and thus flexible representation, but each of these steps often introduces large errors and biases.

In our work, we have focused on fundamental representations that do as little computer-vision analysis as possible. As part of this research we developed a representation called a Lumigraph [1]. A Lumigraph is a 4-dimensional signal that represents the visual information about a real object or environment. Because this representation involves very little analysis on the data, it can represent very general kinds of scenes with very widely varying photometric properties. (The cost, understandably, is a relatively large size for these data sets). Because of this generality, it is hoped that such a representation may be the basis for future 3D television systems. In fact, a very related approach was used as part of the eye-vision instant replay system used in the 2000 Super Bowl.

Another example of this tradeoff can be found in our Image-based visual hull algorithm [7]. In this work, we began with a standard computer vision analysis approach, called shape from silhouette. Since we did not need an output model, but only an *image* of the model from a currently desired viewpoint, we were able to substantially speed up the algorithm (roughly from n^3 to n^2). This system runs in real time, and was used to help create one of the special effects in the recent movie “Minority Report.”

Our research in the IBR area has been followed up by investigations into various practical and theoretical aspects of Lumigraph rendering. Details are described in the paper listing below.

Our work offers many avenues for further investigation. In the short term, we are currently very interested in pursuing a better understanding of reconstruction filters that should be used with Lumigraph data. In the long term, we are interested in the general trend of data-driven 3D graphics - that is, the generation of new shape, material, and motion from raw real-world data. These kinds of approaches have proved to be a nice partner to the more traditional, synthetic approach to computer graphics.

Discrete Geometry Processing

Another large component of our research effort has focused on discrete-geometry-processing (DGP). This is the study of shape representations used for computer graphics. DGP research combines ideas from many areas including approximation theory, image processing, and computational geometry. Results from DGP research have allowed for more flexible geometry authoring tools, and for the ability of displaying very complex scenes using level-of-detail rendering methods.

Our research focus here has been on texture mapped triangle meshes. This representation is at the core of modern real-time graphics systems. In “texture mapping,” small images (called textures) are wrapped (or mapped) over the 3D triangle mesh. In order to apply texture mapping, one must have a way of flattening the 3D mesh onto the 2d texture image domain. This flattening is known as parameterization. There are of course many ways to do this parameterization, leading to the important and difficult question of how to do this optimally. This problem has actually been encountered for hundreds of years by people making flat maps of a curved earth. For map making, many different metrics have been used to measure the quality of the parameterization (such as angle or area preservation).

In particular, for the application of computer graphics, we have realized that a more natural metric would be one that directly quantifies how well the discrete texture samples are being used. To this end, we have constructed entirely new and relevant, yet mathematically natural metrics. We first described our metric in [11] described some useful extensions in [12].

Another difficult part of the parameterization problem is dealing with meshes of arbitrary topology. If one wishes to apply simple parameterization methods, one needs to cut the mesh open to a disk-like shape. This kind of problem is one example of a new set of algorithmic questions being explored in the field of computational topology.

In our work, we have focused on cutting methods that are specifically designed to allow for subsequent high quality parameterizations as explored in [13]. As part of that research project, we have developed the Geometry Image representation - a completely regular way of representing geometric surfaces of arbitrary topology.

Our work offers many avenues for further investigation. In the short term, we are currently very interested in investigating a host of relevant problems from computational topology. In particular we are currently investigating how discrete holomorphic forms may be a way to represent parameterizations of arbitrary-topology meshes in a way that is completely invariant to how the mesh is cut. In the long term, we are interested in advancing the field of DGP for the purpose of real-time rendering of arbitrarily complex scenes.

Annotated Bibliography of Selected Post-PhD Papers

Lightfield/Lumigraph representations

[1] The Lumigraph

Steven Gortler, Radek Grzeszczuk, Richard Szeliski and Michael Cohen.

In *Computer Graphics, Annual Conference Series 1996 (Proc. SIGGRAPH '96)*:43-54.

In this paper we describe a basic 4D representation for representing the appearance of a scene.

[2] **Time Critical Lumigraph Rendering**

Peter-Pike J. Sloan, Michael Cohen and Steven Gortler.

In *Proceedings 1997 Symposium on Interactive 3D Graphics*:17-23

Here we investigate some practical issues of Lumigraph rendering.

[3] **Polyhedral Geometry and the Two-Plane Parameterization**

Xianfeng Gu, Steven Gortler and Michael Cohen.

In *Proceedings Eighth Eurographics Workshop on Rendering, 1997*:1-12

Here we describe some of the algebraic structure of the two parallel plane parameterization used for Lumigraphs.

[4] **Dynamically Reparameterized Light Fields.**

Aaron Isaksen, Leonard McMillan and Steven Gortler.

In *Computer Graphics, Annual Conference Series 2000 (Proc. SIGGRAPH '00)*:297-306.

We investigate the role of reconstruction filters used with undersampled Lumigraphs. In particular we find that ghosting artifacts are an odd mixture of both pre and post aliasing artifacts, and show how the ghosting can be avoided.

[5] **Unstructured Lumigraph Rendering**

Chris Buehler, Mike Bosse, Leonard McMillan, Steven Gortler and Michael Cohen

In *Computer Graphics, Annual Conference Series 2001 (Proc. SIGGRAPH '01)*:425-432

We investigate the relationship between Lumigraph rendering and view-dependent texture mapping. We derive a flexible algorithm designed to utilize the strengths of both approaches. The result is a Lumigraph system that can render from unstructured input.

Other Image Based rendering approaches

[6] **Layered Depth Images**

Jonathan Shade, Steven Gortler, Li-wei He and Richard Szeliski

In *Computer Graphics, Annual Conference Series 1998 (Proc. SIGGRAPH '98)*:231-242.

We explore how efficient rendering can be done using a layered depth image representation. This representation has influenced many point and imposter based rendering techniques.

[7] **Image Based Visual Hulls**

Wojciech Matusik, Chris Buehler, Steven Gortler, Ramesh Raskar and Leonard McMillan

In *Computer Graphics, Annual Conference Series 2000 (Proc. SIGGRAPH '00)*:369-374.

We describe how visual hull computation can be done in a completely image-based framework. This algorithm allows for a real-time novel viewpoint renderer from a live video feed.

[8] **Image Based Rendering: A New Interface Between Computer Vision and Computer Graphics**

Leonard McMillan and Steven Gortler

In *Computer Graphics* 33(4), Nov 1999:57-63.

Invited survey/position paper.

Mesh Processing

[9] **Silhouette Clipping**

Pedro Sander, Steven Gortler, Hugues Hoppe and John Snyder

In *Computer Graphics, Annual Conference Series 2000 (Proc. SIGGRAPH '00)*:327-334.

Describes a very fast algorithm for real time silhouettes extraction from triangle meshes.

[10] **Discontinuity Edge Overdraw**

Pedro Sander, Hugues Hoppe, John Snyder and Steven Gortler

In *Proceedings 2001 Symposium on Interactive 3D Graphics*:167-174

We use our fast silhouette algorithm to aid in real time antialiasing.

[11] **Texture Mapping Progressive Meshes**

Pedro Sander, John Snyder, Steven Gortler and Hugues Hoppe

In *Computer Graphics, Annual Conference Series 2001 (Proc. SIGGRAPH '01)*:409-416

We introduce our geometric stretch metric to measure the quality of a parameterization.

[12] **Signal-Specialized Parametrization**

Pedro Sander, Steven Gortler, John Snyder and Hugues Hoppe.

In *Proceedings Thirteenth Eurographic Workshop on Rendering, 2002* :87-100

We introduce a parameterization metric to be used when a specific signal is known. We also derive some new properties and interpretations of the geometric stretch metric.

[13] **Geometry Images**

Xianfeng Gu, Steven Gortler and Hugues Hoppe.

In *Computer Graphics, Annual Conference Series 2002 (Proc. SIGGRAPH '02)*:355-361

We introduce a completely regular image-like representation for meshes of arbitrary topology.

Other topics

[14] **Feature-Based Cellular Texturing for Architectural Models**

Justin Legakis, Julie Dorsey and Steven Gortler

In *Computer Graphics, Annual Conference Series 2001 (Proc. SIGGRAPH '01)*:309-316

We describe a language based approach to geometric modeling with cellular type elements.

[15] **Quantum versus Classical Learnability**

Rocco Servedio and Steven Gortler

In *IEEE 2001 Conference on Computational Complexity*:138-148

We show that while quantum computation can indeed aid the computational efficiency behavior of learning algorithms, it cannot aide their information theoretic efficiency. This project was unrelated to graphics, but undertaken when I found some really interesting papers abandoned near the printer.

[16] **Rendering Techniques 2001 (Proc. EGRW)**

Editors Steven Gortler and Karol Myszkowski

Springer Verlag

The joys of being a papers chair.

[17] **Minimal Surfaces for Stereo Vision**

Chris Buehler, Steven Gortler, Michael Cohen and Leonard McMillan

In *ECCV 2002:III* 885-899

We show how the minimal discrete surface problem can be reduced to a min cut problem using a dual graph construction. This algorithm is then used to design a very accurate three-camera stereo algorithm. We are currently investing more uses for our minimal discrete surface algorithm.

Research papers to pay special attention to

In reviewing my research, I would recommend paying special attention to the papers numbered 1,5,7,12,13,17.