

# THE GEOMETRY-IMAGE REPRESENTATION TRADEOFF FOR RENDERING

Sing Bing Kang, Richard Szeliski, and P. Anandan  
Vision Technology Group  
Microsoft Research  
Microsoft Corporation  
Redmond, WA 98052, USA

## Abstract

*It is generally recognized that 3-D models are compact representations for rendering. While pure image-based rendering techniques are capable of producing highly photorealistic outputs, the size of the input “model” is usually very large. The important issues in trading off geometry versus images include compactness of representation, photorealism of reconstructed views, and speed of rendering. In this paper, we describe our past work in modeling and rendering, and articulate lessons learnt. We then delineate our vision of an ideal rendering system.*

## 1 Introduction

Many commercial graphics systems and research focus on explicit 3-D modeling. This is a very traditional area, and the 3-D modeling and rendering issues have been mostly dealt with successfully. It is generally true that 3-D models are compact representations for rendering, especially when fast specialized hardware graphics accelerators are readily available. Better photorealism can be achieved at a loss in rendering speed.

Pure image-based rendering techniques (relevant surveys include [8, 12]) index images directly, and their outputs can be highly photorealistic. The important trade-offs in choosing geometric representations versus images include compactness of representation, output quality, speed, and cost.

In this paper, we review our work in extracting geometry and image-based representations from images, and summarize important lessons from our work. We also describe our concept of an ideal rendering system that is capable of using spatially varying representations to produce an optimal mix of high compression and realistic rendering.

## 2 Review of relevant work

The methods for object and scene visualization can be classified according to how geometry or image intensive the underlying representation is. On the image end of the

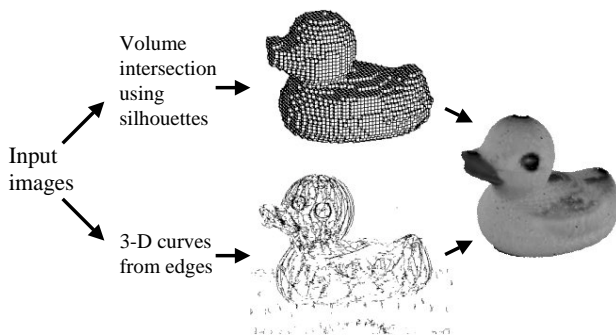


Figure 1: Two possible ways of modeling an object: using silhouettes, and using edges.

spectrum, there are systems with constrained motion (either pure rotation or translation), such as mosaics [26] or view morphing [20]. While this kind of representation is very compact, the degrees of freedom in virtual navigation are very limited. On the other hand, light-field rendering [13] and its model-assisted cousin the Lumigraph [7] provide an enhanced range of virtual motion, but at the expense of a large input database. The concentric mosaic representation [22] is more compact, but loses a degree of freedom in motion and depth effect. The important issue of sampling and artifacts in the context of light field rendering has been examined by several researchers [2, 3, 14].

Work that is primarily image-based with some implicit geometry includes interpolation of color images with single [16] or multiple [4] depth maps, multiple source rendering [9], multiple centers of projection [18], and layers [21, 1]. The compactness of these representations range from high (for flat layers) to low (with explicitly assigned depth at every pixel).

There is plenty of work on 3-D modeling, but we mention only a few as (somewhat biased) representatives in




Figure 2: 3-D modeling from multiple cylindrical panoramic images: (a) sequence of input panoramic images; (b) top view of recovered 3-D points; (c) top view of 3-D mesh; and (d) oblique view of texture-mapped model.

this area. Automatic modeling from images can take the form of reconstruction of an octree model from silhouettes [25], 3-D curves from contours [28] (Figure 1), or simply from 3-D points using panoramic images [10] (Figure 2). Examples of 3-D construction systems that require some interactivity include Façade, which uses manual placement of 3-D primitives [5], and Shum *et al.*'s system that makes use of 2-D lines drawn on panoramic mosaics [23].

### 3 Texture-mapped geometry

Texture-mapped geometric models are widely used in computer graphics. Methods for creating such models include using a CAD modeler, using a 3-D digitizer, and applying computer vision techniques to images of real objects or scenes.

Unfortunately, computer vision techniques are generally not robust enough to work under all conditions. Stereo algorithms do not work for textureless regions without prior structural knowledge (see Figure 2(d), for example). In addition, it is very difficult to capture complex visual effects such as highlights, reflections, and transparency using a texture-mapped 3-D model.

Mesh simplification is often used to reduce the complexity of a model in order to reduce its rendering time. However, simplification may lead to “equal-opportunity” smoothing, since it is difficult to automatically distinguish noise from data.

### 4 Image-based representations

A large number of image-based rendering systems have been developed in the last few years (see Section 2). Here,

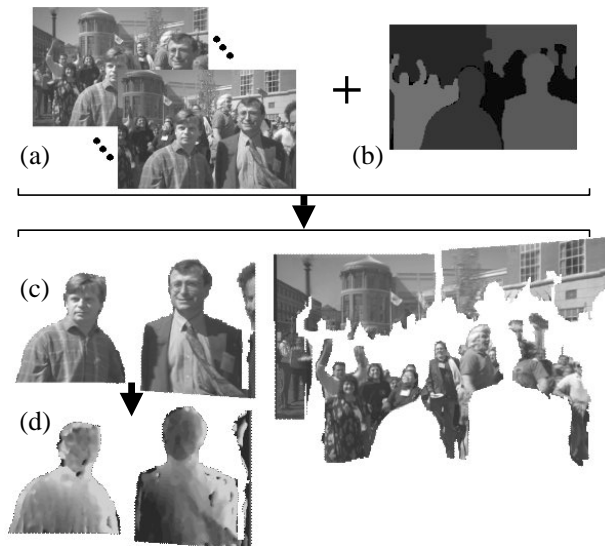


Figure 3: Layered extraction from multiple images: (a) input sequence; (b) manually segmented layers; (c) mosaicked layers; and (d) relative depth (parallax) maps.

we discuss three representative image-based representations, namely sprites with depth, layered-depth images, and the Lumigraph.

The *sprites with depth* representation [1, 21] (a.k.a. 3-D layered motion models) consists of a collection of overlapping, quasi-planar surfaces with local parallax [11]. Each sprite consists of a *matted* color image, a plane equation, and a depth map. A triangular mesh could also be used to represent the parallax.

Sprites with depth naturally capture the occlusion effects that occur in complex scenes, e.g., partial visibility, the mixing of foreground and background colors at occluding contours, and transparent mixing of light at reflections [27]. Sprites with depth can also be efficiently rendered using a combination of forward and inverse warping techniques [21]. Figure 3 shows the sprites extracted from a scene simultaneously photographed from a number of different viewpoints. Using a single depth map would be inadequate to capture all of the surfaces in this scene, while building full 3-D models would not be possible because of the limited range of viewpoints.

A variant on sprites with depth is the *Layered Depth Image* (LDI) representation [21], where multiple color/depth values can be stored at each location in a 2-D array. This representation is easier to compute than a layered sprite representation, since it does not require any segmentation of the scene into parts. On the other hand, the rendering algorithm, which uses forward mapping (*splatting*) cannot




Figure 4: Lumigraph example: (a) sampling space; (b) octree of toy lion (from image sequence); and (c) a novel view (note the detailed fur, which would be very hard to produce with a texture-mapped 3-D model).

take advantage of the contiguity between adjacent pixels in the LDI.

The most data-intensive image-based representation is the 4-D lightfield, which captures all rays passing through some viewing volume [13]. The Lumigraph [7] uses a similar representation, but also contains a rough 3-D model (see Figure 4) that enables better quality reconstructions using fewer light ray samples. Reduced (e.g., 3-D) Lumigraphs which only sample images along some viewing *path* can use far less data, while still presenting the viewer with a strong sense of parallax [24, 22].

## 5 The geometry-image continuum

Based on the complementary sets of characteristics associated with 3-D texture-mapped and image-based models, we can imagine using a continuum of representations to represent a scene. The representation would be designed to optimize compactness, speed, and visual fidelity; it could change spatially within the same scene. One such range of representations is shown in Figure 5. As can be seen, at the physically-based end of the spectrum, we have the usual single texture-mapped geometry. Along the way, we have variations of geometry models, followed by increasingly image-intensive representations. Note also that the dominant means of rendering changes from the conventional graphics pipeline to warping to interpolation, suggesting that an optimized renderer needs a different specialized hardware to handle these different modes of rendering.

A good bridge from the use of pure geometry to pure images is view-dependent geometry [17]. View-dependent geometry is useful to accommodate the fact that stereo reconstruction errors are less visible during local viewpoint perturbations, but may show dramatic effects over large view changes. In areas where stereo data is inaccurate, we may well represent these areas with view-dependent geometry, which comprises a set of geometry extracted at various positions. View-dependent geometry may also be used to capture visual effects such as highlights and transparency, which are likely to be locally coherent in image and viewpoint spaces.




Figure 5: Range of possible spatially-varying representations with their respective dominant rendering mechanisms.

The range of representations to be used for a scene could be determined on a per application basis, i.e., depending on the speed requirement and bandwidth constraint. One possible approach would be to start from the most geometric representation and progressively move to a more and more image-based representation (Figure 5) based on how well the reconstructed view match the expected view, subject to a predefined threshold. The error metric used could be a perceptually-based one as described in [19]. For example, we would expect a blank textureless wall to be represented simply by a plane with possibly view-dependent texture. On the other hand, a plant may be represented by view-dependent geometry with view-dependent texture or a Lumigraph, depending on the desired quality.

## 6 Discussion

For 2-D image viewing over the Internet, progressive JPEG is a popular format because it allows the user to immediately view the whole image, with increasingly better detail as time progresses. In a similar manner, we can imagine using a form of progressive transmission of representations from pure geometry to more image-based ones for 3-D viewing in web applications. This allows the user to virtually navigate with the normal full degrees of freedom, but with initially impoverished perceptual quality. The quality should increase with time. An interesting alternative is to download high quality images in such a manner as to allow virtual navigation, but with reduced degree of freedom in movement (say, to just rotation). In addition, within each representation, tradeoffs of rendering quality for time can still be used. An example of this is using sub-sampling or projective textures for the Lumigraph [24].

So far we have discussed the selection criteria from the end-user's or application's point of view. Important considerations exist from the content creator's perspective, such as ease of image acquisition and ease of model reconstruction (both involving expense and achievability). These factors may dictate the range of models suitable for a given application.

Another important consideration is the adaptability to the rendering platform. Platforms can vary from PDAs to PCs to supercomputers. In rendering on a PDA, for example, compactness and speed is more important than fidelity of view reconstruction due to its relatively low CPU horsepower and screen resolution. On higher end machines, compactness may be less of an issue.

How would one implement the “perfect” rendering engine? One possibility would be to utilize current hardware accelerators to produce, say, an approximate version of an LDI or a Lumigraph by replacing it with view-dependent texture-mapped sprites. The alternative is to design new hardware accelerators that can handle both conventional rendering and IBR. An example in this direction is the use of PixelFlow to render image-based models [15]. PixelFlow [6] is a high-speed image generation architecture that is based on the techniques of object-parallelism and image composition.

## 7 Conclusions

We have described various 3-D reconstruction and image-based techniques, and outlined their characteristics. Generally, while 3-D texture-mapped models are compact, their construction is not error-free, and complex visual effects such as highlights and transparency cannot be easily replicated. On the other hand, more image-based representations such as the Lumigraph are capable of producing photorealistic views, but at the expense of a large database and high cost of image acquisition. Based on these two sets of complementary (and conflicting) characteristics, we propose a set of representations that an optimal renderer may use in combination to produce acceptable levels of compactness, photorealism, and speed.

## References

- [1] S. Baker, R. Szeliski, and P. Anandan. A layered approach to stereo reconstruction. In *CVPR*, pages 434–441, Santa Barbara, CA, June 1998.
- [2] E. Camahort and D. Fussell. A geometric study of light field representations. Technical Report TR99-35, Dept. of Computer Sciences, The Univ. of Texas at Austin, Austin, TX, 1999.
- [3] J. Chai and H.-Y. Shum. Plenoptic sampling. *SIGGRAPH*, Aug. 2000.
- [4] S. E. Chen and L. Williams. View interpolation for image synthesis. *SIGGRAPH*, pages 279–288, July 1993.
- [5] P. E. Debevec, C. J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. *SIGGRAPH*, pages 11–20, Aug. 1996.
- [6] J. Eyles, S. Molnar, J. Poulton, T. Greer, A. Lastra, N. England, and L. Westover. Pixelflow: The realization. In *Siggraph/Eurographics Workshop on Graphics Hardware*, Los Angeles, CA, Aug. 1997.
- [7] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The Lumigraph. *SIGGRAPH*, pages 43–54, Aug. 1996.
- [8] S. B. Kang. A survey of image-based rendering techniques. In *Videometrics VI (SPIE Int’l Symp. on Electronic Imaging: Science and Technology)*, volume 3641, pages 2–16, San Jose, CA, Jan. 1999.
- [9] S. B. Kang and H. Q. Dinh. Multi-layered image-based rendering. In *Graphics Interface*, pages 98–106, 1999.
- [10] S. B. Kang and R. Szeliski. 3-D scene data recovery using omnidirectional multibaseline stereo. *Int’l J. of Computer Vision*, 25(2):167–183, Nov. 1996.
- [11] R. Kumar, P. Anandan, and K. Hanna. Direct recovery of shape from multiple views: A parallax based approach. In *ICPR*, pages 685–688, Jerusalem, Israel, Oct. 1994.
- [12] J. Lengyel. The convergence of graphics and vision. *IEEE Computer*, 1:46–53, 1998.
- [13] M. Levoy and P. Hanrahan. Light field rendering. *SIGGRAPH*, pages 31–42, Aug. 1996.
- [14] Z. Lin and H.-Y. Shum. On the number of samples needed in light field rendering with constant-depth assumption. In *CVPR*, Hilton Head Island, NC, June 2000.
- [15] D. K. McAllister, L. Nyland, V. Popescu, A. Lastra, and C. McCue. Real-time rendering of real world environments. In *Eurographics Workshop on Rendering*, Granada, Spain, June 1999.
- [16] L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. *SIGGRAPH*, pages 39–46, Aug. 1995.
- [17] P. Rademacher. View-dependent geometry. *SIGGRAPH*, pages 439–446, Aug. 1999.
- [18] P. Rademacher and G. Bishop. Multiple-center-of-projection images. *SIG*, pages 199–206, July 1998.
- [19] M. Ramasubramaniam, S. N. Pattanaik, and D. P. Greenberg. A perceptually based physical error metric for realistic image synthesis. *SIGGRAPH*, pages 73–82, Aug. 1999.
- [20] S. M. Seitz and C. R. Dyer. View morphing. *SIGGRAPH*, pages 21–30, Aug. 1996.
- [21] J. Shade, S. Gortler, L.-W. He, and R. Szeliski. Layered depth images. *SIGGRAPH*, pages 231–242, July 1998.
- [22] H.-Y. Shum and L.-W. He. Rendering with concentric mosaics. *SIGGRAPH*, pages 299–306, August 1999.
- [23] H. Y. Shum, M. Han, and R. Szeliski. Interactive construction of 3D models from panoramic mosaics. In *CVPR*, pages 427–433, Santa Barbara, CA, June 1998.
- [24] P.-P. Sloan, M. F. Cohen, and S. J. Gortler. Time critical Lumigraph rendering. In *Symp. on Interactive 3D Graphics*, pages 17–23, Providence, RI, Apr. 1997.
- [25] R. Szeliski. Rapid octree construction from image sequences. *CVGIP: Image Understanding*, 58(1):23–32, July 1993.
- [26] R. Szeliski. Video mosaics for virtual environments. *IEEE Computer Graphics and Applications*, pages 22–30, Mar. 1996.
- [27] R. Szeliski, S. Avidan, and P. Anandan. Layer extraction from multiple images containing reflections and transparency. In *CVPR*, Hilton Head Island, NC, June 2000.
- [28] R. Szeliski and R. Weiss. Robust shape recovery from occluding contours using a linear smoother. *IJCV*, 32(1):27–44, June 1998.