Uniformly Sampled Light Fields

Emilio Camahort *
Department of Computer Sciences
The University of Texas at Austin

Apostolos Lerios †
Align Technology, Inc.

Don Fussell *
Department of Computer Sciences
The University of Texas at Austin

Abstract

Image-based or light field rendering has received much recent attention as an alternative to traditional geometric methods for modeling and rendering complex objects. A light field represents the radiance flowing through all the points in a scene in all possible directions. We explore two new techniques for efficiently acquiring, storing, and reconstructing light fields in a (nearly) uniform fashion. Both techniques sample the light field by sampling the set of lines that intersect a sphere tightly fit around a given object. Our first approach relies on uniformly subdividing the sphere and representing this subdivision in a compact data structure which allows efficient mapping of image pixels or rays to sphere points and then to subdivision elements. We sample a light field by joining pairs of subdivision elements and store the resulting samples in a multi-resolution, highly compressed fashion that allows efficient rendering. Our second method allows a uniform sampling of all five dimensions of the light field, using hierarchical subdivision for directional space and uniform grid sampling for positional space. Light field models are acquired using parallel projections along a set of uniform directions. Depth information can also be stored for high-quality image rendering. The system can provide bounds on key sources of error in the representation and can be generalized to arbitrary scenes comprising multiple complex objects.

1 Introduction

Image-based modeling and rendering techniques have recently received much attention as an alternative to traditional geometry-based techniques for image synthesis. Synthetic images are created from a prestored set of samples of the light field of an environment instead of a geometric description of the elements in the scene. While image-based techniques have long been used to augment geometric models in the form of texture maps or environment maps, for example, newer approaches completely replace geometric information with image data. This new approach allows the construction of models directly from real-world image data without conversion to geometry. It also can be more efficient than geometric techniques for complex models which require very large amounts of geometric detail to represent.

Image-based modeling and rendering relies on the concept of a *light field*. A light field represents the radiance flowing through all the points in a scene in all possible directions. For a given wavelength, we can represent a static light field as a 5D scalar function $L(x, y, z, \theta, \phi)$, that gives radiance as a function of location and direction. Location is represented by a point (x, y, z) in 3D space. Direction is specified by a pair (θ, ϕ) of angles giving azimuth and elevation, respectively. Later we also use the alternative notation $\vec{\omega}$ to represent directions as unit vectors or, equivalently, points on a sphere's surface.

Images are discretized 2D slices of this 5D function. Given a set of viewing parameters, we can render an image by evaluating the light field function at the (x, y, z) location of the eye for a discrete set of directions within the field of view.

Current research in image-based models falls into either of two categories: (i) techniques related to computer vision, image warping, and view interpolation, and (ii) techniques related to the fundamental representation, sampling and reconstruction of the light field. With a few exceptions, most recent developments fall into the first category. Our paper, however, improves on current techniques for light field representations of single objects.

Previous light field models are based on the two plane parameterization (2PP), where the set of lines in space is parameterized by the intersection points of each line with two planes. In order to sample all the lines intersecting the convex hull of an object we can choose different arrangements of pairs of planes. Levoy and Hanrahan call each pair a *light slab* [25]; Gortler et al. [19] propose a single arrangement of six pairs of planes called the *lumigraph*. Images produced using these two approaches show noticeable artifacts when the camera crosses the boundary between two light slabs. Even arrangements of 12 light slabs do not suffice to avoid this problem, called the *disparity problem*.

In this paper we propose a solution to the disparity problem. A light field representation should allow the user to move freely around an object without noticing any resolution changes in the model. This requires the representation to be invariant under both rotations and translations. Such a representation samples the light field function by uniformly sampling the set of lines intersecting the object's convex hull. As an approximation to the convex hull we use a sphere tightly fit around the object. Given that approximation, we introduce two new uniform representations for light field models, the two-sphere parameterization and the sphere-plane parameterization. The first one was studied by the second author while a doctoral candidate at Stanford University. The second one is the work of the first and third authors at the University of Texas at Austin.

Our representations improve on previous ones for several reasons. First, they allow us to sample the light field in a (nearly) uniform fashion. Uniformity is relevant because it solves the disparity problem, but it also has some other important advantages. For example, when sampling a function whose variation is unknown a priori, uniform sampling provides a good preview of the function, that can later be refined as more information about the function is known. Also, compression theory often makes the assumption of uniformly spaced samples. For instance, the discrete Fourier transform assumes that its input is a sequence of regularly spaced function samples.

Second, our representations profit from the spherical nature of the light field function. Previous approaches avoid spherical representations due to their complexity. Instead, they use cylindrical projections or planar perspective projections for simplicity. We show in this paper that it is possible to efficiently store and render a light field using an entirely spherical representation. Furthermore, our representations implement much of the functionality of previ-

* Taylor Hall 2.124, Austin TX 78712, {ecamahor|fussell}@cs.utexas.edu, http://www.cs.utexas.edu/users/{ecamahor|fussell}

† 2991 El Camino Real, Suite 120, Redwood City CA 94061, lerios@cs.stanford.edu, http://www-graphics.stanford.edu/∼tolis/