Light Field Rendering with Triangular Viewpoint Mesh

Contact
Dongho Kim
Department of Computer Science
The George Washington University
dkim@seas.gwu.edu

James K. Hahn
Department of Computer Science
The George Washington University

Light field rendering[1] is a representative image-based rendering system using two-plane parameterization. The dataset in this scheme is a set of rays passing through regular grids sampled on two parallel planes. And this 4D dataset in (s,t,u,v) space can be thought of as a set of 2D images that have their viewpoints located at rectangular grids on a plane. We use (s,t) coordinates to parameterize the viewpoints and (u,v) coordinates to parameterize pixels in the images.

Rendering from this dataset is a process of reconstruction of rays. Due to the nature of digital sampling with limited resolution, we cannot avoid rendering error. In this work, we propose a new triangular sampling scheme of (s,t) plane. And we compare this scheme with rectangular sampling.

Figure 1(a) shows rectangular sampling in the original light field rendering. The viewpoints are sampled regularly at rectangular grids in (s,t) space. Figure 1(b) shows new parameterization of (s,t) based on triangular sampling. Here, the viewpoints are located at the vertices of a triangular mesh. This sampling can be obtained by moving the sample points in alternate rows of the original sampling by half the distance between the samples, and we can know that the sampling density is the same in both cases. In other words, the cost for the sampling does not change in new parameterization.

Since a light field is a set of discrete samples, we have to consider the aliasing problem. Figure 2(a) shows the frequency domain spectrum after the sampling in Figure 1(a). Due to the digital sampling, the spectrum of the original signal is repeated along two dimensions of the frequency domain. The spacing between the repetitions is determined by the sampling intervals. Here, we assume that the original signal is band-limited within a circle. This assumption makes sense if the original signal does not have any directional strength. Because of the overlapping spectrums, we get aliasing error from the overlapping spectrums. Figure 2(b) shows the frequency domain spectrum after the sampling in Figure 1(b). In this case, due to the characteristics of interlaced sampling, the amount of aliasing is reduced because the distances between spectrums are increased. In other words, for the same amount of aliasing, we need a smaller number of sampling points with new parameterization.

For reconstruction in original light field rendering, quadrilinear interpolation is used in (s,t,u,v) space. This is bilinear interpolation in (s,t) viewpoint sampling space. For new triangular parameterization, we interpolate with the barycentric coordinates of the reconstructed rays. So barycentric weights determine the interpolation in (s,t) space. Since we still use bilinear interpolation in (u,v), 12 samples are involved in the reconstruction, while original light field rendering uses 16 samples. Obtaining barycentric coordinates for reconstructed rays can be done quickly by calculating the distance to the parallel lines of the triangular mesh. Or hardware-based rendering presented in Isaksen et al.[3] can be used. Figure 3 is a rendering example from the DRAGON dataset in the Stanford light field archive.[4]

In order to compare two parameterizations, the images are reconstructed using the same viewpoints. From the original 32 x 32 images in the DRAGON dataset, we use 16 x 16 images for both parameterizations so that the samplings are performed as in Figure 1. And we reconstruct at the viewpoints along the blue lines in Figure 1, where the correct images are known from the DRAGON dataset but are not used in both samplings. From a rendering of 225 viewpoints, triangular sampling showed 2.1 percent less error on average, while using fewer samples for reconstruction. Moreover, the maximum error is smaller than original light field rendering. This is due to the enhanced sampling with less aliasing.

In future work, we need more concrete verification of the suggested parameterization. One method could be rendering a lot of images from randomly sampled viewpoints and performing statistical comparison. Or it is possible that spectral analysis like that suggested by Chai et al.[2] will be desirable.

References

1. Levoy, M. & Hanrahan, P. (1996). Light field rendering. *Proceedings of SIGGRAPH 96*, 31-42.
2. Chai, J., Chan, S., Shum, H., & Tong, X. (2000). Plenoptic sampling. *Proceedings of SIGGRAPH 2000*, 307-318.
3. Isaksen, A., McMillan, L., & Gortler, S.J. (2000). Dynamically reparameterized light fields. *Proceedings of SIGGRAPH 2000*, 297-306.
4. URL: www-graphics.stanford.edu/software/lightpack/lifs.html

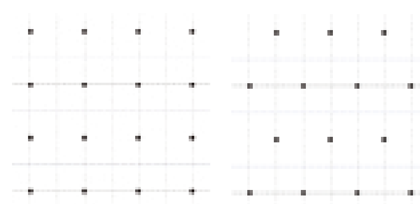

Figure 1. Sampling grids: (a) rectangular sampling, (b) triangular sampling.

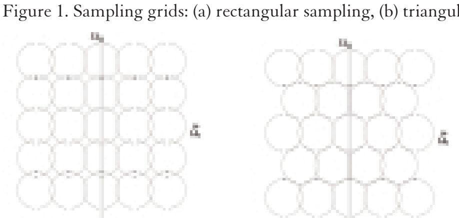

Figure 2. Frequency spectrum: (a) rectangular sampling, (b) triangular sampling.

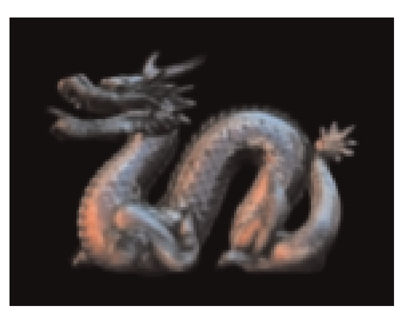

Figure 3. Reconstructed image.