

Proceedings of SIACG2011

V Ibero-American Symposium in Computer Graphics



F. Silva, D. Gutierrez, J. Rodríguez, M. Figueiredo (Eds.)

Organized by
Algarve University

In cooperation with
Grupo Português de Computação Gráfica
Eurographics Spanish Chapter
and Eurographics Association

Copyright ©2011 by Grupo Português de Computação Gráfica

All rights reserved. No part of this book may be reproduced or transmitted in any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without permission in writing from the publisher.

Published by Grupo Português de Computação Gráfica.

Proceedings of the V Ibero-American Symposium in Computer Graphics

Editors: F. Silva, D. Gutierrez, J. Rodríguez, M. Figueiredo

ISBN: 978-972-98464-6-5

Printed in Portugal

Cover design by Ideias Frescas

International Program Committee

- Abel Gomes (Univ. Beira Interior, Portugal)
Adérito Marcos (Univ. Aberta, Portugal)
Adolfo Muñoz (Univ. Zaragoza, Spain)
Adriano Lopes (Univ. Nova Lisboa, Portugal)
Aldo Reyes (Univ. Carabobo, Venezuela)
Alex Cuadros (Univ. Católica San Pablo, Peru)
Alexandra La Cruz (Univ. Simón Bolívar, Venezuela)
António Augusto Sousa (Univ. Porto, Portugal)
Carles Bosch (REVES/INRIA Sophia-Antipolis, France)
Ernesto Coto (Univ. Central de Venezuela, Venezuela)
Fernando Navarro (Microsoft, UK)
Francisco Serón (Univ. Zaragoza, Spain)
Gustavo Patow (Univ. Girona, Spain)
Hélder Araújo (Univ. Coimbra, Portugal)
Holly Rushmeier (Yale University, USA)
Javier Lluch (Univ. Politècnica de Valencia, Spain)
João Madeira Pereira (IST - Univ. Técnica Lisboa, Portugal)
Joaquim Madeira (Univ. Aveiro, Portugal)
Jorge Lopez-Moreno (Univ. Zaragoza, Spain)
Juan Carlos Torres (Univ. Granada, Spain)
Luciano Pereira Soares (TecGraf - PUC Rio, Brasil)
Luís Matey (CEIT, Spain)
Luís Paulo Santos (Univ. Minho, Portugal)
Manuel Próspero dos Santos (Univ. Nova Lisboa, Portugal)
Mateus Sbert (Univ. Girona, Spain)
Miguel Angel Otaduy (Univ. Rey Juan Carlos, Spain)
Miguel Chover (Univ. Jaume, Spain)
Miguel Sales Dias (Microsoft/IUL, Portugal)
Omaira Rodriguez (Univ. Central Venezuela, Venezuela)
Pablo Alejandro Figueiroa (UNIANDES, Colômbia)
Pedro Linares (Univ. de Carabobo, Venezuela)
Pere Brunet (Univ. Politècnica Catalunya, Spain)
Rachel McDonnell (Trinity College, Ireland)
Rafael Bidarra (TU Delft, Netherlands)
Rhadamés Carmona (Univ. Central Venezuela, Venezuela)
Robert Joan-Arinyo (Univ. Politècnica Catalunya, Spain)
Teresa Chambel (Univ. Lisboa, Portugal)
Verónica Orvalho (Univ. Porto, Portugal)

Contents

Preface	vii
Contents	xi
Tutorials	
Introduction to Programming in CUDA C Timothy Lanfear	3
Geographic Information Systems J. Rodrigues, C. Rebelo and S. Madeira	5
Invited Speakers	
3D Documents – from Affordable Creation and Indexing to Effective Usage Dieter W. Fellner	9
HCI Remixed: Interaction, Sensing and New Devices Shahram Izadi	11
Session 1:	
Data Visualization	
Digital Cartographic Generalization in Spatial Databases: Application Issues in Power Grids CAD Tools J. Rodríguez, E. Pereira and M. Canosa.....	15
3D Visualization of Sparse Geophysical Data Representing Uncertainty V. Gonçalves, P. Dias, F. Almeida, J. Madeira and B. Santos	23
Interactive Visualizations of Automatic Debugging Reports A. Ribeiro, R. Rodrigues and R. Abreu.....	31
GraphJudge: A System for Assisted Assessment of Computer Graphics Assignments L. Rocha and R. Rodrigues	37
Neural Data Exploration with Force Feedback (SP) L. Raya, M. Otaduy, and M García	43
Interacting with Outdoor Sports Performance Data (SP) J. Lopes	47
Session 2:	
Lights Fields and Image Processing	
High-quality Compression of Mipmapped Textures C. Andujar	53
Non-Uniform Spherical Light Fields J. Blasco, F. Abad, E. Camahort and R. Vivó	61
Graph-Based Reflectance Segmentation E. Garces, D. Gutierrez and J. Lopez-Moreno.....	67
Efficient Propagation of Light Field Edits A. Jarabo, B. Masia and D. Gutierrez.....	75

Non-Uniform Spherical Light Fields

J. Blasco¹, F. Abad^{1,2}, E. Camahort^{1,2}, R. Vivó^{1,2}

¹ Instituto Universitario de Automática e Informática Industrial, Universitat Politècnica de València, València, Spain
jblasco@ai2.upv.es

² Departamento de Sistemas Informáticos y Computación, Universitat Politècnica de València, València, Spain
{fjabad,camahort,rvivo}@dsic.upv.es

ABSTRACT

The study of the light field, as a function defining the radiance along the different viewing directions of a scene, started more than a decade ago and is still ongoing due to its interesting applications. Among the different parameterizations existing for the light field, we will be centered on the spherical light-field parameterization due to its quality of viewpoint-independency. In this work, we present an improved multiresolution method for rendering of spherical light fields with a non-uniform sampling, while focusing on the visualization of large storage volume light-field datasets.

Categories and Subject Descriptors (according to ACM CCS):

Computer Graphics [I.3.3]: Picture/Image Generation—Display algorithms—

Computer Graphics [I.3.5]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations—

Computer Graphics [I.3.6]: Methodology and Techniques—Graphics data structures and data types—

1. Introduction

Image-based rendering techniques are those which build three-dimensional representations of a scene using only two-dimensional information extracted from images of it. One such technique is light-field rendering [LH96]. By storing radiance information from a synthetic or real scene, light-field rendering allows the visualization of the scene from novel points of view. The information stored in a light field consists of radiance samples along a 4D set of oriented lines. Light fields can also be extended with depth information or geometric proxies as supplementary information allowing a more accurate reconstruction of the radiance function with a smaller amount of samples.

Like other image-based rendering methods, the simplified acquisition process is one of the strengths of light-field rendering. As opposed to the acquisition process of data used in traditional 3D rendering, which often requires complex scanning hardware for acquisition of real-world models, image data acquisition can be as simple as taking pictures of a model from a given set of points pointing in a given set of directions. Moreover, light fields have been studied to improve their representation [IMG00] or used in conjunction with other methods to produce novel view generation techniques [CEJ*06].

Another advantage of this class of techniques is their inherent independency of the rendered scene's geometric complexity. This fact makes them ideal for the visualization of 3D datasets of arbitrary geometric complexity while main-

taining interactive framerates. This is also true when scene complexity depends on computationally expensive shading effects. Complexity of image-based representations depends only on the image complexity, which in the specific case of light fields translates to a direct dependency on the spatial and directional resolution of the frame to be rendered.

As noted previously, light-field rendering also supports the use of hybrid information consisting of radiance and geometry. This feature allows the rendering of scenes with a great amount of detail produced by textures coupled with simple geometry supporting it. This provides light-field representations with means to integrate themselves easily with content based on more traditional rendering techniques, and can also be used to produce renderings of a higher quality with a reduced amount of samples with algorithms that use depth correction [GGSC96].

On the other hand, light-field models can present high storage and memory requirements on the rendering system if high levels of detail are to be supported. To reduce the impact of this requirements on the rendering performance, both multi-resolution and on-demand data retrieval techniques are used to enable out-of-core rendering techniques.

Moreover, recent advances in the field of 3D devices, specifically multiview displays, have further increased the relevance of light fields as a rendering technique of choice for visualization of 3D models [YCH*05, CNC*05]. Autostereoscopic displays, which produce different views of a scene for different viewpoints creating an illusion of vol-

ume, work in a manner much similar to light fields, as in both cases different images are mapped to different viewing directions. This intuitively leads to an association between both, although the specifics of this depend on the parameterization of the light-field representation used and the specifications of the devices [FNH*02, JMY*07].

This paper presents an improvement in spherical light-field rendering [IPPL97, CLF98, TRSKK08]. The spherical light-field model we use has the properties described above, including hierarchical multiresolution, support for geometric proxies and depth correction, and support for autostereoscopic rendering [EBA*09]. Spherical light-field parameterizations also grant viewpoint-independent rendering of the light field and hence uniform image quality across different points of view. In this work, we present an improvement on the light-field rendering algorithm that provides different levels of detail on arbitrary sampling directions. This allows us to reduce the directional and spatial resolution, for a pre-defined set of angles. It also allows us to reduce the dataset sizes and memory footprint of a light-field model by using less sampling on viewpoints which are less likely to be examined or are not acquirable for a given scene.

In the next pages, we review previous work and theory upon which light-field rendering finds its foundation, in order to provide a background on the specifics of the chosen rendering model. After that, we explain the features, advantages and disadvantages of this model. We then explain how these features are of relevance to our work, and present our contributions. Later, we outline some ideas for future work and discuss the conclusions derived from our current work.

2. Previous work

The concept of light field is based on the plenoptic function defined by Adelson and Bergen [AB91]. If we consider the mathematical concept of *pencil* as the set of rays passing through a point in space, we can define the light field as the radiance of every pencil in space. More formally, it can be defined as a function that captures the radiance R flowing through every point (x, y, z) in every possible direction (θ, ϕ) of 3D space. We can simplify this definition assuming radiance to be of a fixed wavelength and restricting time to represent a fixed moment, thus obtaining a 5D scalar function $R(x, y, z, \theta, \phi)$, which will give a radiance value for each position and direction in space. Note that by reducing the dimensionality of the function in this manner, the support of that function becomes the set of rays $L(x, y, z, \theta, \phi)$ of 3D space [GGSC96, LH96].

As radiance is constant along a given ray, we can further reduce the dimensionality of the function support to 4D if we assume an occlusion-free 3D region of space is being modeled. Therefore, the support of R becomes the set of oriented lines in 3D space. Once we have defined the light field in this way, all we need to do is describe this set of lines according to a given parameterization, which can be structured or unstructured. This parameterization will produce a discrete model, allowing us to quantify the radiance function by taking samples of it in a finite set of points inside the bounds of its support.




Figure 1: Planar parameterization.

2.1. Light-field Models

The two types of parameterizations, structured and unstructured, differ in the way they sample 3D space. Unstructured parameterizations do not impose restrictions on the viewpoints used to capture the images of the objects. Models based on unstructured parameterizations can be built using approximate geometric models [BBM*01], depth information [SVSG01], or measurements of the focus of the object's points [TN05]. On the other side, structured parameterizations sample the rays using regular patterns in their parameters. Most structured parameterizations can be classified into one of two types, planar and spherical.

Planar parameterizations, which incorporate concepts from holography [HH92], parameterize each line in the light field with two points P and Q . These points define the intersection of a line with two parallel planes (see Figure 1), a construction given the name of *light slab*. Therefore, a line L is defined by parameters (s, t, u, v) . According to this definition, we need a minimum of six light slabs to form a volume completely enclosing a given model [LH96]. Gortler *et al.* extended this model using geometric *proxies* to improve the reconstruction of the continuous radiance function with a depth correction algorithm, that reduced artifacts like ghosting and seams [GGSC96].

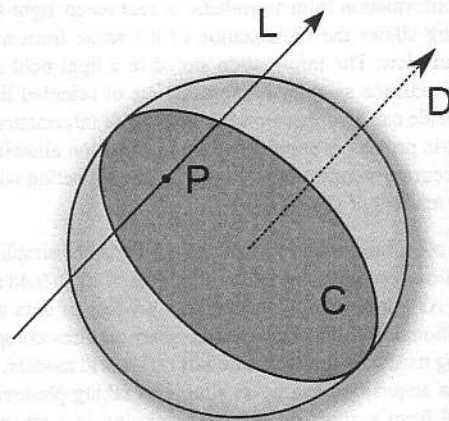


Figure 2: Spherical parameterization.

There are several types of structured spherical parameterizations. The most widely known ones are the two sphere parameterization [IPPL97] and the direction and point parameterization [CLF98]. These parameterizations have the

advantage that they produce a (quasi)uniform sampling of the line space where the light field is defined. The two-sphere parameterization defines a ray using a point on the surface of a sphere bounding the model along with a point on a sphere clinging to the previous one. On the other hand, the direction-and-point parameterization selects a point on a chosen plane and its perpendicular direction to achieve the same result. In this work we use this last parameterization, as described in detail in Section 3. Figure 2 shows the direction and point parameterization, where the ray L is given by its direction D and its intersection P with the support plane of circle C . An alternative use of this parameterization by Todt *et al.* uses parabolic images to reduce distortion during rendering [TRSKK08].

One advantage of these parameterizations with respect to planar parameterizations is that the latter need several slabs to enclose the object to model. This produces discontinuities in the radiance function and noticeable artifacts when the scene is viewed from viewpoints where more than one slab is visible at the same time. Furthermore, spherical parameterizations present interesting hierarchical multi-resolution properties.

Another alternative parameterization is surface light fields [WAA*00, CBCG02]. This parameterization needs to obtain a geometric model of the scene to be rendered. This is afterwards simplified and used as a sampling surface which closely approximates the object. A similar approach to this technique is view-dependent texture mapping [DTM96], where only a small set of sampling directions are used to capture the appearance of the object modeled.

3. Non-Uniform Spherical Light Fields

We base the model we work with on the spherical light-field parameterization, due to its multiresolution properties and the fact that it is defined uniformly on the whole scene volume. The fact this parameterization is defined in a nearly uniform manner for all possible viewpoints allows us to render light fields correctly without severe discontinuities across the possible viewpoints. This provides us with the option to selectively reduce the amount of samples for low interest viewpoints, by allowing an arbitrary degree of discontinuity or correctness as a trade-off for reduced memory requirements and enhanced performance.

3.1. Spherical Light-Field Modeling

Our model is based on the spherical light field defined by the direction and point parameterization (DPP). To parameterize a light field according to the DPP, the scene volume is firstly fit inside a sphere of unit radius. Then, we uniquely characterize each ray by its direction and its intersection point with a great circle, perpendicular to the ray direction and passing through the center of the sphere.

With this parameterization we build a discrete representation of the light field by obtaining a discretization of the set of pencils located on the surface of the sphere. For this, we pick a set of pencils, and approximate each of them by the direction passing through the center of the pencil. If we choose this set of directions so that they match the directions defined by the lines going from the sphere's center to the center of each triangle in a geodesic triangulation of the sphere, we

obtain a quasi-uniform set of directional samples. This way each directional sample passes through the center of a triangle and each triangle represents the pencil of directions approximated by the sample. We can then capture the radiance associated to each pencil by obtaining an image with an orthographic camera. The projection direction is the direction passing through the center of the triangle. This discretization allows the use of recursive subdivision of the triangles to approximate the sphere's surface with more detail, providing a simple multiresolution scheme for our representation.




Figure 3: *Triangular subdivision of the sphere: white triangles belong to the original icosahedron, gray triangles belong to the first subdivision level, and black triangles belong to the second subdivision level.*

Now, before sampling the light field following this parameterization, we must determine the resolution at which we perform the sampling. Sampling resolution will be the product of the directional resolution used, or number of directional samples, and the spatial resolution, or number of points sampled for each direction. The number of directional samples is the same as the number of triangles present in the discretization of the sphere. For our multiresolution model, we use 20 directions for the lowest level of detail (LOD), level 0. They match the number of triangles of an icosahedron. This gives us the coarsest LOD. For higher levels of detail, we subdivide the icosahedron's triangles regularly, generating a hierarchical structure where each subdivision level is made of a tessellation of the previous one with four times its number of triangles (see Figure 3). With regard to spatial resolution, it is given by the resolution of the camera used in the light-field capture process.

This light-field representation also supports depth correction. The representation stores both radiance and depth information associated to each directional sample. Depth information does not need to be very precise as it is only used to approximate the original geometry. Typically, 8-bit or 16-bit depth values provide enough precision for depth-supported rendering [EBA*09, TRSKK08].

3.2. Uniform Spherical Light-Field Rendering

Rendering of spherical light fields is much simpler than other geometry-based rendering techniques as it just requires the drawing of a relatively low amount of textured triangles with the correct projections and no lighting calculations. This algorithm is a modified version of the original Lumigraph rendering algorithm [GGSC96].

In our case, the viewing volume is enclosed inside the unit sphere representing the light field, and both are centered at the origin. Visible pencils correspond then to the sphere triangles visible in this volume. Correct radiance samples in the images representing each pencil are shown by computing the appropriate texture coordinates in each texture. We do this by projecting each vertex onto the corresponding support plane of the pencil associated to the triangle. This light-field rendering algorithm is explained in detail in the paper by Escrivà *et al.* [EBA*09].

Our sampling scheme allows generating views for any viewpoint on the unit sphere. The renderings of such models produce good results if we generate enough samples for the dataset, typically 1280 directional samples and 256x256 spatial samples. However, spherical light fields may require a larger amount of samples if depth correction is not used. When using depth correction, the storage sizes of the datasets can be reduced to a degree, but they can still remain large if high rendering detail is desired. In order to work with such large dataset sizes, the use of memory caching techniques and out-of-core rendering has been proposed, achieving acceptable results in rendering performance, but at the cost of extended usage of memory.

3.3. Non-Uniform Spherical Light-Field Modeling

In order to reduce the dataset size and memory usage of spherical light fields, we propose a simple modification of the light-field acquisition and rendering processes. As noted before, spherical light fields are characterized by the uniformity of their sampling scheme. It supports rendering a scene with the same quality regardless of the viewing parameters. This uniformity is desirable for rendering scenes, so that the user is allowed to rotate around the model freely and explore it without rendering artifacts or changes in the representation.

However, many times the observation of a scene can be narrowed to a reduced range of viewpoints, depending on the specific qualities of the model or the specific goal of the rendering. A clear example of this situation is the visualization of a light field obtained from a statue. In this case there are two circumstances leading to the use of a reduced number of viewpoints during observation of the model.

The first one is the absence of samples for parts of the viewpoint range. If captured from a real world object, chances are that a set of the light-field samples is difficult or even impossible to obtain. For instance, the samples in the bottom hemisphere of a sampled model resting on the ground, or the samples in the top part of the hemisphere if the height of the model prevents taking samples of it.

On the other hand, it is possible that a subset of the viewpoints that could be rendered is, simply, irrelevant. If we were to generate a light field from the model of a 3D-digitized statue, the model would still contain a digitized base or stand, or a flat surface in case the stand was removed from the 3D model. This surface presents little to no useful information to the viewer, and is most likely not to be observed when the dataset is rendered.

This same reasoning could be applied to arbitrary sections of a model if we wanted to focus only on certain parts of it, and is a determinant factor in the size of a dataset. If we

know in advance which viewpoints require more detail, then we can distribute samples accordingly and fine-tune both the spatial and the directional resolutions. The resulting light-field model will either use less memory or it may store more detail for selected directional samples.

3.4. Non-Uniform Spherical Light-Field Rendering

In order to implement non-uniform light fields, we need to generate and store a non-uniform dataset and implement a modified version of the rendering algorithm of Escrivà *et al.* [EBA*09]. To do so, we support selecting a different spatial resolution for each directional sample of the dataset, and storing a non-uniform triangulation hierarchy for directional sampling. Selecting different spatial resolutions allows using lower resolution images for low interest regions or, conversely, higher resolution images for regions of interest.




Figure 4: Two example views of a non-uniform sphere tessellation (regions with pencils with different subdivision levels have been separated for clarity: white: lower resolution pencils, gray: medium resolution pencils, and black: higher resolution pencils).

Independent directional resolution will allow us to control the availability of viewing points. In order to produce different directional resolutions for different views, we create a scene-dependent configuration, where samples are available only for selected pencils. Using the spherical model and its hierarchical structure, we produce different refinement levels for each of the lowest LOD pencils. Starting with the initial 20 directions, we mark the maximum refinement for each of their associated pencils. That way we avoid taking too many samples for pencils that do not need them. We only subdivide the icosahedron's triangles where higher LODs are needed.




Figure 5: Two views of the two versions of our test dataset: Top row, the back of the sculpture has less detail: Top left, the sculpture at full uniform resolution: note the absence of artifacts; and top right, the non-uniform dataset showing directional artifacts (seams) in the image, and positional artifacts (pixelation) in the inset; and bottom row, the front of the sculpture dataset has more detail and it appears at roughly the same resolution in the uniform dataset (bottom left, and the non-uniform dataset (bottom right).

Figure 4 shows two views of an example non-uniform triangulation of the sphere. Figure 5 shows a light field of a statue modeled using both a uniform and a non-uniform representation.

Note that our rendering algorithm does not show the seams present in the tessellation illustrated in Figure 3. In that Figure neighboring triangles of different resolutions show seams along their shared edges. Such seams do not appear in the light-field renderings because the projections of such triangles onto the renderer's viewing frustum completely covers the projection plane without holes or seams.

4. Results

To test our rendering algorithm for non-uniform spherical light fields, we generated a test dataset with a uniform subdivision of level 4. The dataset is made of 5120 directional samples and stores a 512×512 pixel image for each directional sample.

The dataset contains images rendered from a digitized polygonal model of 1.5 million polygons. It took us 5 days to render all the images using Blender. The complete light-field dataset occupies more than 5 gigabytes. To render the highest LOD of the uniform subdivision at interactive rates we need to store all the images in main memory (see Figure 5). That is, we need a computer outfitted with, at least, 6 gigabytes of memory. The model, however, presents a large flat area at the bottom, as well as few features in the back, leaving little information of interest in those regions. This

allows for a non-uniform sampling of the model to lower the LODs of the regions containing less information. The regions containing more information, like the front of the statue, are stored at higher LODs. A non-uniform rendering of the model can be seen in Figure 5. It requires roughly 1 gigabyte of memory, thus easily fitting in a computer's main memory.

To build the dataset, we decrease the model's LOD in a progressive manner, changing from higher LODs at the front to lower LODs at the bottom and the back of the figure. The directional resolution varies progressively from level 4 at the front of the statue to level 1 at the back of the statue. The spatial sampling is also selected in the same way as the directional samples. Specifically, higher LOD pencils use 512×512 pixel images. Likewise, lower LOD pencils use 256×256 pixel images. This variable multiresolution representation produces correct results when rendering front views. When rendering bottom and back views, artifacts appear caused by undersampling in regions that do not contain important features. Artifacts are typically seams and warping. As for the memory requirements of the non-uniform dataset, it only uses 932 MBytes instead of the 5120 MBytes of memory space used by the uniform dataset. This reduces the spatial cost of the representation by a factor of five and allows storing the entire dataset in main memory, thus reducing cache misses and disk loads. Our storage policy, which manages pencil samples hierarchically, simplified manual generation of the non-uniform representation using our knowledge of the model's features.

5. Conclusions and Future Work

In this paper we present a non-uniform subdivision scheme for storing and rendering spherical light fields. The data can be generated at different spatial and directional resolutions depending on the need of detail for the different directional samples of the light-field model. Furthermore, models with different non-uniform subsets of directional samples can be built in cases where it is impossible to acquire all the samples or there is a large sample subset with little or no interest to the viewer. In either case our approach can substantially reduce the storage requirements of the datasets while still providing good quality renderings.

We obtained the results of this paper manually generating the non-uniform light-field model. To do so we used subjective criteria based on our knowledge of the geometric model. A key goal of our future research is to create an automated or supervised acquisition process for non-uniform spherical light fields. We want to determine the relevance of each sample in a semi-automatic or automatic manner. For each directional sample, we want to be able to decide whether to capture or render an image and what resolution to use for that image. Additionally, we want to explore which criteria to use to determine whether a sample is worth generating or not.

Our idea is to use, when possible, both objective and subjective criteria. Objective criteria may for instance be based on estimations of the complexity of the sampled view. Subjective criteria can be based on test cases of scene observation by different users, expert opinions about a model's features, and remarks made by the creator of the model.

Such a methodology would allow a fine-tuned light-field

sampling where memory requirements would be minimized, allowing more rendering detail and increased cache performance.

6. Acknowledgements

The Ramesses model used in this work was provided courtesy of the INRIA by the AIM@SHAPE Shape Repository. This work was partially supported by grant TIN2009-14103-C03-03 of the Spanish Ministry of Science and Innovation.

References

- [AB91] ADELSON E. H., BERGEN J. R.: The plenoptic function and the elements of early vision. *Computational Models of Visual Processing* (1991), 3–20. 2
- [BBM*01] BUEHLER C., BOSSE M., MCMILLAN L., GORTLER S., COHEN M.: Unstructured lumigraph rendering. In *SIGGRAPH '01* (2001), ACM, pp. 425–432. 2
- [CBCG02] CHEN W.-C., BOUGUET J.-Y., CHU M. H., GRZESZCZUK R.: Light field mapping: Efficient representation and hardware rendering of surface light fields. In *SIGGRAPH 2002 Conference Proceedings* (2002), Hughes J., (Ed.), Annual Conference Series, ACM Press/ACM SIGGRAPH, pp. 447–456. 3
- [CEJ*06] CHABERT C.-F., EINARSSON P., JONES A., LAMOND B., MA W.-C., SYLWAN S., HAWKINS T., DEBEVEC P.: Relighting human locomotion with flowed reflectance fields. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Sketches* (New York, NY, USA, 2006), ACM, p. 76. 1
- [CLF98] CAMAHORT E., LERIOS A., FUSSELL D.: Uniformly sampled light fields. In *Proc. Eurographics Rendering Workshop '98* (1998), pp. 117–130. 2
- [CNC*05] CHUN W.-S., NAPOLI J., COSSAIRT O. S., DORVAL R. K., HALL D. M., II T. J. P., SCHOOLER J. F., BANKER Y., FAVALORA G. E.: Spatial 3-d infrastructure: Display-independent software framework, high-speed rendering electronics, and several new displays. In *Proceedings of SPIE-IS&T Electronic Imaging* (2005), vol. 5664, pp. 302–312. 1
- [DTM96] DEBEVEC P. E., TAYLOR C. J., MALIK J.: Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1996), ACM, pp. 11–20. 3
- [EBA*09] ESCRIVÁ M., BLASCO J., ABAD F., CAMAHORT E., VIVÓ R.: Autostereoscopic rendering of multiple light fields. *Computer Graphics Forum* 28 (December 2009), 2057–2067(11). 2, 3, 4
- [FNH*02] FAVALORA G. E., NAPOLI J., HALL D. M., DORVAL R. K., GIOVINCO M. G., RICHMOND M. J., CHUN W. S.: 100 million-voxel volumetric display. In *Proceedings of the SPIE* (2002), vol. 4712, pp. 300–312. 2
- [GGSC96] GORTLER S. J., GRZESZCZUK R., SZELISKI R., COHEN M. F.: The lumigraph. In *Proc. SIGGRAPH '96* (1996), pp. 43–54. 1, 2, 3
- [HH92] HAINES K., HAINES D.: Computer graphics for holography. *IEEE Computer Graphics and Applications* (January 1992), 37–46. 2
- [IMG00] ISAKSEN A., MCMILLAN L., GORTLER S. J.: Dynamically reparameterized light fields. In *Proc. SIGGRAPH '00* (2000), pp. 297–306. 1
- [IPPL97] IHM I., PARK S., PARK S., LEE R. K.: Rendering of spherical light fields. In *PG '97: Proceedings of the 5th Pacific Conference on Computer Graphics and Applications* (Washington, DC, USA, 1997), IEEE Computer Society, p. 59. 2
- [JMY*07] JONES A., MCDOWALL I., YAMADA H., BOLAS M., DEBEVEC P.: An interactive 360° light field display. In *SIGGRAPH '07* (2007), ACM. 2
- [LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *Proc. SIGGRAPH '96* (1996), pp. 31–42. 1, 2
- [SVSG01] SCHIRMACHER H., VOGELGSANG C., SEIDEL H., GREINER G.: Efficient free form light field rendering. In *Proceedings of Vision, Modeling, and Visualization 2001* (2001). 2
- [TN05] TAKAHASHI K., NAEMURA T.: Unstructured light field rendering using on-the-fly focus measurement. *ICME 2005. IEEE International Conference on Multimedia and Expo* (July 2005), 205–208. 2
- [TRSKK08] TODT S., REZK-SALAMA C., KOLB A., KUHNERT K.-D.: GPU-Based Spherical Light Field Rendering with Per-Fragment Depth Correction. *Computer Graphics Forum* 27, 8 (2008), 2081–2095. 2, 3
- [WAA*00] WOOD D. N., AZUMA D. I., ALDINGER K., CURLESS B., DUCHAMP T., SALESIN D. H., STUETZLE W.: Surface light fields for 3d photography. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 287–296. 3
- [YCH*05] YANG R., CHEN S., HUANG X., LI S., WANG L., JAYNES C.: Towards the light field display. *IEEE VR 2005 Workshop on Emerging Display Technologies* (March 2005). 1