

Sparse Reconstruction of Visual Appearance for Computer Graphics and Vision

Ravi Ramamoorthi^a

^aUniversity of California, Berkeley, 525 Soda Hall, Berkeley, CA 94720, ravir@cs.berkeley.edu

ABSTRACT

A broad range of problems in computer graphics rendering, appearance acquisition for graphics and vision, and imaging, involve sampling, reconstruction, and integration of high-dimensional (4D-8D) signals. For example, precomputation-based real-time rendering of glossy materials and intricate lighting effects like caustics, can involve (pre)-computing the response of the scene to different light and viewing directions, which is often a 6D dataset. Similarly, image-based appearance acquisition of facial details, car paint, or glazed wood, requires us to take images from different light and view directions. Even offline rendering of visual effects like motion blur from a fast-moving car, or depth of field, involves high-dimensional sampling across time and lens aperture. The same problems are also common in computational imaging applications such as light field cameras.

In the past few years, computer graphics and computer vision researchers have made significant progress in subsequent analysis and compact factored or multiresolution representations for some of these problems. However, the initial full dataset must almost always still be acquired or computed by brute force. This is often prohibitively expensive, taking hours to days of computation and acquisition time, as well as being a challenge for memory usage and storage. For example, on the order of 10,000 megapixel images are needed for a 1 degree sampling of lights and views for high-frequency materials. We argue that dramatically sparser sampling and reconstruction of these signals is possible, before the full dataset is acquired or simulated. Our key idea is to exploit the structure of the data that often lies in lower-frequency, sparse, or low-dimensional spaces. Our framework will apply to a diverse set of problems such as sparse reconstruction of light transport matrices for relighting, sheared sampling and denoising for offline shadow rendering, time-coherent compressive sampling for appearance acquisition, and new approaches to computational photography and imaging.

Keywords: Sparse Representations, Multiresolution, Visual Appearance, Computer Graphics, Computer Vision, Fourier Analysis, Compressive Sensing, Wavelets, Multi-Dimensional Signal Processing

1. INTRODUCTION

Computer graphics and computer vision are undergoing an exciting revolution owing to the synergy between acquisition of real imagery, and simulation of light transport. On the simulation side, new techniques seek to rapidly compute the interplay of light in a scene and to enable real-time changes of lighting or materials, as well as complex effects such as motion blur and depth of field. The functions that need to be computed to make this possible are intricate and high-dimensional (some examples are shown in Fig. 1): 4D for incoming illumination, plus time, plus extra dimensions for materials, and the integration over the area of the camera's lens, which can result in 8D functions. This poses formidable challenges. On the acquisition side, there is also a growing need to capture the intricate interplay between light and real objects, and classical photographic challenges are revisited to beat traditional limitations using a combination of new optical capture systems and computation. This is the major thrust of the new field of computational photography, where new methods can alleviate optical problems such as motion blur or depth of field, and can in general use computation to capture more information from real scenes. The treatment of these photographic problems involves the same high-dimensional functions required for simulation or rendering, for example 5D in the combined case of depth of field and motion blur (2D for the image sensor, 2D for the lens, and 1D for time). In both the simulation and capture case, brute force approaches have sought to tackle the problem by computing or acquiring data for the full high-dimensional function, incurring extremely high precomputation or acquisition costs. We argue that one *can derive and leverage new information about the structure and sparsity of these functions, making it possible to address these problems with a dramatically reduced number of samples*. This will mean much faster rendering times, as well as capture strategies that provide high-quality visual data from sparse sampling.




Figure 1. Examples of high-dimensional appearance datasets in rendering and imaging. **Left:** Two example images of anisotropic wallpaper, showing striking differences under different light source directions and views (the full dataset is 6D, and involves several thousand images).⁵ **Middle:** Precomputed light transport of a glass vase with global illumination and caustics (a 4D dataset for all lighting directions over the image, that took days to compute).⁶ **Right:** A lenslet array enables imaging of full 4D light fields, which later enable variations in focus.⁷ In all of these cases, acquisition or precomputation is extremely expensive, requiring days of computation, thousands of images even for a sparse sampling, or drastically reduced resolutions. This makes “brute-force” sampling impractical in many cases. In this paper, we propose new frequency analysis and sparse reconstruction tools that exploit sparsity for acquisition itself (before compression), and enable orders of magnitude greater efficiency.

Most high-dimensional functions involved in rendering and acquisition exhibit much structure and sparsity. They are often locally highly anisotropic, low frequency, or low rank. While this observation has often been exploited for compression after the high-dimensional data has been simulated or acquired, *the key idea of this position paper is that sparsity must be exploited upstream to dramatically reduce the cost of rendering and acquisition, and to devise reconstruction schemes that extract more information from given data.*

For instance, consider *sparsity*, wherein only some of the wavelet coefficients of light transport are nonzero—this sparsity is an ideal fit to new developments in compressive signal-processing.¹⁻³ Another example is *space-angle and space-time bandlimitedness*. While visual appearance data can have rich all-frequency structure in space, angle or time separately, there are often simpler structures when considering the full high-dimensional space. Our recent work on a frequency analysis of motion blur⁴ has shown that the Fourier spectrum of a dynamic scene is sheared into a narrow wedge in space-time. This low-dimensional structure in space-time enables novel sheared filters and sparse sampling to render images substantially faster than previous work.

In summary, we argue that it is crucial to exploit lower-complexity structures in visual appearance datasets, to develop fundamentally more efficient algorithms and acquisition. The originality of this approach is to leverage sparsity, low-frequency and low-rank properties *before* data is simulated or acquired, thereby avoiding redundant computation or sensing. We envision theoretical and practical contributions to a broad range of topics that have historically been treated separately, and will enable intimate connections between real-time, image-based and offline rendering, and imaging applications.

In the remainder of this paper, we first briefly discuss the high-dimensional nature of visual appearance, and the key application areas that we consider. We also discuss the previous work on factored and multiresolution representations that can compress and efficiently represent these visual datasets. We then describe the types of sparsity we seek to exploit, followed by a survey of the state of the art in sparse sampling and reconstruction for these visual appearance problems.

2. BACKGROUND AND APPLICATIONS

In this section, we introduce the key concepts and application areas. We start with the reflection equation,⁸

$$L_r(\mathbf{x}, \boldsymbol{\omega}_o) = \int_{\Omega_{2\pi}} L_i(\mathbf{x}, \boldsymbol{\omega}_i) V(\mathbf{x}, \boldsymbol{\omega}_i) \rho(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \max(\boldsymbol{\omega}_i \cdot \mathbf{n}, 0) d\boldsymbol{\omega}_i, \quad (1)$$

where $L_r(\mathbf{x}, \boldsymbol{\omega}_o)$ is the reflected or outgoing radiance, corresponding to the image we want to render. The variables \mathbf{x} , $\boldsymbol{\omega}_i$ and $\boldsymbol{\omega}_o$ are respectively, the spatial location, incident and outgoing angles (since we are in 3D, they are vectors). Here, L_i is the incoming light, V is the visibility (usually binary), and ρ is the BRDF,⁹ which is in general a function of incident

and outgoing directions, and may also be a function of spatial position. As far as this paper is concerned, we will be most interested in the BRDF, $\rho(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ and generalizations thereof, that are crucial to a number of applications in rendering and imaging, noted below.

2.1 Precomputed Real-Time Rendering

For real-time rendering applications, the BRDF and visibility (as well as the cosine term) are often combined into a single transport term T , and we modify equation 1 into,

$$B(\mathbf{x}) = \int_{\Omega_{2\pi}} L_i(\boldsymbol{\omega}_i) T(\mathbf{x}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \quad (2)$$

where we make the additional simplification that we consider a fixed view (or Lambertian surfaces), so that the outgoing direction $\boldsymbol{\omega}_o(\mathbf{x})$ is fixed for a given spatial location and omitted above. We also assume distant lighting, with L_i independent of spatial location \mathbf{x} . While we have derived the above formulation from the reflection equation, it can also take into account global illumination and the rendering equation,¹⁰ since the only real assumption is the linearity of light transport. Viewpoint variation can be added by repeating this computation for a small set of views, or using factored representations of the BRDF^{11,12} or direct computation using factored wavelet triple product integrals.¹³

The advantage of this formulation is that T can be precomputed for a static scene, and then combined with dynamic lighting in real-time using a dot-product of orthonormal basis functions like spherical harmonics or wavelets. The basic idea exploits linearity of light transport^{14–16} and has recently gained much attention in the form of precomputed radiance transfer (PRT).^{6,17,18} Indeed, PRT methods are now commonplace for interactive rendering in video games. The transport is usually stored at a fixed set of spatial locations (vertices of geometry or image pixels), and projected into low-frequency spherical harmonics.^{17,19}

Light transport, as noted above is a function of 4 dimensions, 2 each for the spatial location \mathbf{x} and incident direction $\boldsymbol{\omega}_i$. For reasonable resolutions, the precomputation and storage time itself can be quite significant, with billions of samples needing to be evaluated and stored to exhaustively precompute $T(\mathbf{x}, \boldsymbol{\omega}_i)$, requiring many hours of computation and gigabytes of storage. A number of PRT papers have addressed more compact representations, such as clustered PCA²⁰ (although challenges remain, especially for higher-fidelity wavelet-based methods).

Challenge: *However, precomputation time remains a challenge, which limits the application of the method. A natural question is whether we can precompute T from a sparser set of samples, especially since we will ultimately apply a variety of compression methods.*

2.2 Appearance Acquisition

One of the challenges in computer graphics and computer vision is to acquire the appearance of real scenes, and in particular the ways in which they reflect light. (It should be noted that this challenge is not specific to visual appearance; data-driven methods have pervaded all areas of computer vision and graphics; for example, the widespread use of motion-captured and precomputed datasets for animation).

Visual appearance typically involves a number of variants of the light transport function. In general, most objects can be characterized as a black box that takes as input an incident light field, and outputs a reflected light field,^{21,22} i.e., by a transport function $T(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o)$, where \mathbf{x}_i and \mathbf{x}_o are spatial locations or light field parameterizations. Each of the spatial and directional components is 2D, so this is an 8D function. While recent efforts have provided the first measurements of the full 8D quantity,²³ a more common approach is to consider simpler slices by omitting or holding fixed some of the dimensions. For example, the spatially-varying BRDF in equation 1 can be considered a 6D slice where, for a flat surface, we fix $\mathbf{x}_i = \mathbf{x}_o$ (we assume opaque objects with no subsurface scattering). Note that these datasets can be very large. For example, using 100 samples in each dimension to capture sharp specular highlights can cause the 6D spatially varying BRDF to encompass 10^{12} samples.

With datasets of this size (4D-8D), an obvious question becomes how to represent them compactly. A number of approaches have been proposed. One general promising direction is a tree-structured factorization⁵ into lower-dimensional components such as spatial mixing weights and basis BRDFs, with the BRDFs further factored into 2D and 1D components. This is inspired by early work on factorization of measured BRDFs.²⁴

Challenge: *However, before applying factorization we must still sample and determine the full dataset by brute force. While heuristics, such as sampling more in the direction of specular highlights, or using a low sampling rate, are possible, the fundamental challenge remains. Knowing that the data will be factored into lower-dimensional or low-rank components, can we do better? Similar challenges apply to many other forms of light transport, with considerable room for sparser sampling and better reconstruction using the properties of the transport function.*

2.3 Monte Carlo High-Dimensional Rendering

Traditional offline rendering also involves high-dimensional functions, comparable in many ways to the light transport operator. We are particularly interested in efficient rendering of multidimensional effects like motion blur, depth of field, soft shadows, and antialiasing.²⁵ Monte Carlo rendering can be cast in terms of light transport operators as

$$B(x, y) = \int_x^{x+1} \int_y^{y+1} \int_{\mathbf{s}} T(u, v, \mathbf{s}) \, d\mathbf{s} \, du \, dv, \quad (3)$$

where u and v correspond to antialiasing for the given pixel (x, y) and \mathbf{s} corresponds to additional dimensions like motion blur, depth of field and so on. The traditional approach is to sample T with a large number of random (possibly stratified) Monte Carlo samples, and then evaluate the integral.

Challenge: *However, this brute force approach has a number of shortcomings. First, T can have various forms of smoothness that are not exploited. This suggests instead an adaptive signal-processing view that reconstructs T from a sparse set of samples. Second, there is considerable coherence in the Monte Carlo integrals at neighboring pixels, which suggests they should share information. Finally, the result is an integral, which gives us some flexibility in reconstruction T . The reconstruction need not be fully accurate, only the integral needs to be.*

2.4 Imaging

Many problems in imaging are related closely to the rendering challenges above, even though they have historically been studied by different groups of researchers. The simplest of challenges may involve reproducing a high-resolution video from sparse input captured with a lower framerate camera.²⁶ More broadly, many directions in the new area of computational photography are considering variants of the light transport function. For example, light field cameras⁷ capture the incident light field on the camera aperture (this camera is sought to be marketed to consumers in the near future by Lytro inc.) Other problems involve considerations of motion blur and depth of field, seeking to produce sharp images,^{27,28} which have close links to similar problems in Monte Carlo rendering.

Challenge: *The additional high-dimensional information from capturing light fields, or variations with time and depth of field necessitate camera tradeoffs. These involve either much larger numbers of measurements and file sizes, or tradeoffs of spatial resolution for the additional dimensions. A question remains if brute force capture can be replaced with sparser measurement patterns that may allow for more effective capture of a scene, as well as higher all-round resolutions. Crucially, it is important to consider the relationship with similar problems in rendering.*

3. TYPES OF SPARSITY

A solution to efficient data acquisition and sparse reconstruction for the application areas above requires understanding the types of sparsity present in light transport. In this section, we briefly discuss the canonical compression and sparsity in many situations, along with some initial recommendations for the corresponding algorithms. In the next section, we will briefly survey the state of the art in sparse reconstruction of visual appearance, and possible future directions.

	Low Frequency	Low Rank	Locally Low Rank	Sparse	Sheared
Random Projection Point Sample Parallel Meas.	Fourier Nyquist	Optical Computing ³⁹ Row-Column ^{35,36}	Sparse Precomp. ⁴³	Compressed Sensing ³⁰ Compressive ³¹	Sheared Filters ^{4,44}

Table 1. Some recent papers that represent the current state of the art in sparse reconstruction of visual appearance. This is not meant to be an exhaustive list. It is clear that improvements are possible in all of the different categories (and possibly new ones in the future), as well as many open questions. Moreover, algorithms need to be adapted for each practical application domain.

3.1 Cost Metric for Measurements

A key factor in choosing an algorithm is the cost metric for measurements. For example, most of the work in compressive sensing has assumed a **random projection** of the signal is possible at unit cost. This is indeed the case in applications like appearance acquisition, where light is obtained from a projector, and an arbitrary pattern can be used.^{29,30}

However, the application domains discussed in the previous section admit many different cost metrics. In Monte Carlo rendering and its application to precompute light transport matrices, the cost is usually per ray, or for each rendering sample. In essence, we can only **point-sample** the function of interest, not consider random projections (the cost for a random projection would be equal to the number of rays shot, i.e., the number of nonzero elements). Thus, compressive sensing ideas cannot be used or may need significant modification.³¹ On the other hand, newer approaches to matrix completion^{32,33} may be more successful.

In still other applications, it may be easy to acquire multiple **measurements in parallel**, such as taking a single image with millions of pixels using a camera. In rendering, techniques like shadow mapping often make it easy to acquire the full visibility image from a light source. Further, ideas from dual photography³⁴ can allow us to efficiently acquire full rows and columns of a light transport matrix in many applications involving both real and synthetic scenes.^{35,36}

Other measurement modalities may also be possible. In any event, a reconstruction method must consider the cost metric, whether random projection, point-sampling, parallel measurements or something else.

3.2 Light Transport Complexity

The key ingredient in a sparse reconstruction algorithm is the type of sparsity in light transport. We can consider many possibilities. First, light transport may be **low-frequency**, such as for some cases of diffuse reflection.^{37,38} In other instances, we can assume the light transport matrix is of **low rank**³⁹ but make fewer assumptions about the frequency content. Neither of these assumptions is true in general, but a much more widely held notion is that of **locally low rank**,⁴⁰ where local spatial regions behave in a low-dimensional fashion. This observation allows for precomputed light transport compression methods like Clustered Principal Component Analysis.²⁰ It is also known that light transport matrices are often **sparse** in a compressive sense, with only a few wavelet coefficients required to describe the response of a material for each pixel.³⁰ Finally, a frequency analysis in some cases shows that a low-frequency assumption is not possible in any dimension, but the frequency content lies in a **sheared wedge**.^{4,41} In general, frequency analysis of light transport⁴² is a valuable approach to analyzing the sparsity of various components.

Besides the specific properties that lower light transport complexity above, there might be other compressions to be exploited in future work, and one curious question is how best to combine knowledge of multiple types of sparsity. In any event however, a sparse reconstruction algorithm must exploit the type of sparsity for the situation at hand. It is also our belief that a consideration of the matrix of possibilities (an initial effort is shown in Table 1), given a cost metric and light transport sparsity, could lead to a good classification of existing algorithms and point towards new methods.

4. STATE OF THE ART AND FUTURE DIRECTIONS

In this section, we briefly discuss the state of the art in different areas and future directions. This is not intended to be an exhaustive survey, but rather a brief snapshot of some important recent papers in the area. A summary of methods is shown in Table 1. Here, we discuss each application area.

For **Precomputed Real-Time Rendering**, the challenge is to make precomputation fast. In this domain, if using a ray tracing-based solution for precomputation, we must point-sample, and the most salient feature of the light transport matrix

is its locally low rank nature. We have developed a sparse precomputation method,⁴³ but there is clearly lots of room for further improvements, for example using newer matrix completion methods.^{32,33}

Monte Carlo Rendering of high-dimensional effects like motion blur or soft shadows is another problem involving point-sampling, where the main sparsity feature is that the fourier spectrum lies in a narrow wedge in space-time or space-angle. We have exploited this with sheared filtering methods for motion blur⁴ and soft shadows.⁴⁴ This is by no means the end of the story, and adaptive rendering of multidimensional effects²⁵ is now a hot topic.^{45,46} Earlier work introduced the notion of row-column sampling³⁵ in cases where shadow mapping could be used to generate entire rows or columns of the light-surface interaction matrix. All of these methods share samples between nearby pixels and can introduce some artifacts, which poses some issues for their application to production rendering. An important challenge is practical adaptations of these filtering techniques for real-time and offline rendering.

For **Appearance Acquisition**, we do usually have the flexibility to use random projection patterns, which potentially leads to a host of methods for efficient and sparse acquisition. For instance, low-frequency transport can directly be sensed using Fourier basis patterns. General low-rank transport can be determined using large-matrix methods for discovering the leading eigenvalues.³⁹ An alternative is based on row-column sampling using the kernel Nystrom method.³⁶ Compressed sensing can be used to exploit sparsity in light transport.³⁰ Interesting challenges involve other assumptions on light transport, like locally low rank or sheared spectra.

Imaging applications are perhaps the biggest potential beneficiaries of sparse reconstruction methods. Initial efforts have been made using compressed sensing ideas.^{47,48} A key challenge for the future is also to exploit the analogies between rendering and imaging, that have not yet been fully explored.

5. CONCLUSION

Visual appearance, involving high-dimensional datasets, is a crucial component of research in computer vision and computer graphics. Future progress involves developing a unified framework for frequency analysis and sparse data reconstruction, without requiring the appearance datasets to be computed or captured by brute force. This research can leverage many ideas from signal processing such as compressed sensing and matrix completion, but also needs to develop novel algorithms and theories for various combinations of cost metrics and sparsity. In the long-term we hope this research will lay the theoretical foundations for signal-processing of visual appearance. In practical terms, it will lead to efficient algorithms for a variety of problem domains, including sparse reconstruction of light transport matrices for relighting, sheared sampling and denoising for offline shadow rendering, time-coherent compressive sampling for appearance acquisition, and new approaches to computational photography and imaging.

ACKNOWLEDGMENTS

This line of work is in some ways a natural progression of research on signal-processing for appearance, that goes back to my PhD dissertation and owes a debt to many students, mentors, collaborators and funding agencies (NSF, ONR, Intel, Adobe, NVIDIA, Pixar). This work and article is currently sponsored by an ONR PECASE grant N00014-09-1-0741. I also acknowledge NSF CAREER grant 09-24968 for earlier support. Finally, I am grateful to the NSF for supporting future research in this area with NSF grant 1115252, and to Fredo Durand at MIT who is a collaborator on that award and contributed to many of the ideas in this writeup.

REFERENCES

- [1] Candes, E., “Compressive sampling,” in *International Congress of Mathematics*, Number 3, 1433–1452 (2006).
- [2] Candes, E., Romberg, J., and Tao, T., “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on Pure and Applied Mathematics* **59**(8), 1207–1223 (2006).
- [3] Candes, E. and Tao, T., “Near optimal signal recovery from random projections: Universal encoding strategies?,” *IEEE Transactions on Information Theory* **52**(12), 5406–5425 (2006).
- [4] Egan, K., Tseng, Y., Holzschuch, N., Durand, F., and Ramamoorthi, R., “Frequency analysis and sheared reconstruction for rendering motion blur,” *ACM Transactions on Graphics* **28**(3) (2009).
- [5] Lawrence, J., Ben-Artzi, A., DeCoro, C., Matusik, W., Pfister, H., Ramamoorthi, R., and Rusinkiewicz, S., “Inverse Shade Trees for Non-Parametric Material Representation and Editing,” *ACM Transactions on Graphics (Proc. SIGGRAPH 06)* **25**(3), 735–745 (2006).

- [6] Ng, R., Ramamoorthi, R., and Hanrahan, P., “All-Frequency Shadows using Non-Linear Wavelet Lighting Approximation,” *ACM Transactions on Graphics (Proc. SIGGRAPH 03)* **22**(3), 376–381 (2003).
- [7] Ng, R., Levoy, M., Bredif, M., Duval, G., Horowitz, M., and Hanrahan, P., “Light Field Photography with a Hand-Held Plenoptic Camera,” Tech. Rep. CSTR 2005-02, Stanford University (2005).
- [8] Cohen, M. and Wallace, J., [*Radiosity and Realistic Image Synthesis*], Academic Press (1993).
- [9] Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., and Limperis, T., [*Geometric Considerations and Nomenclature for Reflectance (NBS Monograph 160)*], National Bureau of Standards (US) (1977).
- [10] Kajiya, J., “The Rendering Equation,” in [*SIGGRAPH 86*], 143–150 (1986).
- [11] Liu, X., Sloan, P., Shum, H., and Snyder, J., “All-Frequency Precomputed Radiance Transfer for Glossy Objects,” in [*EuroGraphics Symposium on Rendering 04*], 337–344 (2004).
- [12] Wang, R., Tran, J., and Luebke, D., “All-frequency Relighting of Non-diffuse Objects using Separable BRDF approximation,” in [*EuroGraphics Symposium on Rendering*], 345–354 (2004).
- [13] Ng, R., Ramamoorthi, R., and Hanrahan, P., “Triple Product Wavelet Integrals for All-Frequency Relighting,” *ACM Transactions on Graphics (Proc. SIGGRAPH 04)* **23**(3), 475–485 (2004).
- [14] Nimeroff, J., Simoncelli, E., and Dorsey, J., “Efficient Re-Rendering of Naturally Illuminated Environments,” in [*EuroGraphics Workshop on Rendering*], 359–373 (1994).
- [15] Dobashi, Y., Kaneda, K., Nakatani, H., Yamashita, H., and Nishita, T., “A Quick Rendering Method Using Basis Functions for Interactive Lighting Design,” in [*EuroGraphics 95*], 229–240 (1995).
- [16] Dorsey, J., Arvo, J., and Greenberg, D., “Interactive Design of Complex Time-Dependent Lighting,” *IEEE Computer Graphics and Applications* **15**(2), 26–36 (1995).
- [17] Sloan, P., Kautz, J., and Snyder, J., “Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments,” *ACM Transactions on Graphics (Proc. SIGGRAPH 02)* **21**(3), 527–536 (2002).
- [18] Ramamoorthi, R., “Precomputation-Based Rendering,” *Foundations and Trends in Computer Graphics and Vision* **3**(4), 281–369 (2009).
- [19] Ramamoorthi, R. and Hanrahan, P., “An Efficient Representation for Irradiance Environment Maps,” in [*SIGGRAPH 01*], 497–500 (2001).
- [20] Sloan, P., Hall, J., Hart, J., and Snyder, J., “Clustered Principal Components for Precomputed Radiance Transfer,” *ACM Transactions on Graphics (Proc. SIGGRAPH 03)* **22**(3), 382–391 (2003).
- [21] Gortler, S., Grzeszczuk, R., Szeliski, R., and Cohen, M., “The Lumigraph,” in [*SIGGRAPH 96*], 43–54 (1996).
- [22] Levoy, M. and Hanrahan, P., “Light Field Rendering,” in [*SIGGRAPH 96*], 31–42 (1996).
- [23] Garg, G., Talvala, E., Levoy, M., and Lensch, H., “Symmetric Photography: Exploiting Data-sparseness in Reflectance Fields,” in [*EuroGraphics Symposium on Rendering*], 251–262 (2006).
- [24] Kautz, J. and McCool, M., “Interactive Rendering with Arbitrary BRDFs using Separable Approximations,” in [*Eurographics Rendering Workshop 99*], 247–260 (1999).
- [25] Hachisuka, T., Jarosz, W., Weistroffer, R., Dale, K., Humphreys, G., Zwicker, M., and Jensen, H., “Multidimensional adaptive sampling and reconstruction for ray tracing,” *ACM Transactions on Graphics* **27**(3) (2008).
- [26] Mahajan, D., Huang, F., Matusik, W., Ramamoorthi, R., and Belhumeur, P., “Moving gradients: A path-based method for plausible image interpolation,” *ACM Transactions on Graphics (SIGGRAPH 09)* **28**(3) (2009).
- [27] Levin, A., Sand, P., Cho, T., Durand, F., and Freeman, W., “Motion-invariant photography,” *ACM Transactions on Graphics* **27**(3) (2008).
- [28] Levin, A., Hasinoff, S., Green, P., Durand, F., and Freeman, W., “4D frequency analysis of computational cameras for depth of field extension,” *ACM Transactions on Graphics* **28**(3) (2009).
- [29] Gu, J., Nayar, S., Grinspun, E., Belhumeur, P., and Ramamoorthi, R., “Compressive Structured Light for Recovering Inhomogeneous Participating Media,” in [*ECCV 08*], IV–845—IV–858 (2008).
- [30] Peers, P., Mahajan, D., Lamond, B., Ghosh, A., Matusik, W., Ramamoorthi, R., and Debevec, P., “Compressive light transport sensing,” *ACM Transactions on Graphics* **28**(1), Article 3, pages 1–18 (2009).
- [31] Sen, P. and Darabi, S., “Compressive estimation for signal integration in rendering,” *Computer Graphics Forum* **29**(4), 1355–1363 (2010).
- [32] Candes, E. and Plan, Y., “Matrix completion with noise,” in [*Proceedings of the IEEE*], (2009).
- [33] Candes, E. and Tao, T., “The power of convex relaxation: Near-optimal matrix completion.,” *IEEE Transactions on Information Theory* **56**(5), 2053–2080 (2009).

- [34] Sen, P., Chen, B., Garg, G., Marschner, S., Horowitz, M., Levoy, M., and Lensch, H., “Dual Photography,” *ACM Transactions on Graphics (SIGGRAPH 05 proceedings)* **24**(3), 745–755 (2005).
- [35] Hasan, M., Pellacini, F., and Bala, K., “Matrix row-column sampling for the many-light problem,” *ACM Transactions on Graphics (Proc. SIGGRAPH 07)* **26**(3), Article 26 (2007).
- [36] Wang, J., Dong, Y., Tong, X., Lin, Z., and Guo, B., “Kernel Nystrom method for light transport,” *ACM Transactions on Graphics* **28**(3) (2009).
- [37] Ramamoorthi, R. and Hanrahan, P., “On the relationship between Radiance and Irradiance: Determining the illumination from images of a convex Lambertian object,” *JOSA A* **18**(10), 2448–2459 (2001).
- [38] Basri, R. and Jacobs, D., “Lambertian Reflectance and Linear Subspaces,” in [*International Conference on Computer Vision*], 383–390 (2001).
- [39] O’Toole, M. and Kutulakos, K., “Optical computing for fast light transport,” *ACM Transactions on Graphics (SIGGRAPH Asia 10)* **29**(6) (2010).
- [40] Mahajan, D., Kemelmacher-Shlizerman, I., Ramamoorthi, R., and Belhumeur, P., “A theory of locally low dimensional light transport,” *ACM Transactions on Graphics (Proc. SIGGRAPH 07)* **26**(3), 62 (2007).
- [41] Chai, J., Chan, S., Shum, H., and Tong, X., “Plenoptic Sampling,” in [*SIGGRAPH 00*], 307–318 (2000).
- [42] Durand, F., Holzschuch, N., Soler, C., Chan, E., and Sillion, F., “A Frequency Analysis of Light Transport,” *ACM Transactions on Graphics (Proc. SIGGRAPH 05)* **25**(3), 1115–1126 (2005).
- [43] Huang, F. and Ramamoorthi, R., “Sparsely precomputing the light transport matrix for real-time rendering,” *Computer Graphics Forum (EGSR 10)* **29**(4), 1335–1345 (2010).
- [44] Egan, K., Hecht, F., Durand, F., and Ramamoorthi, R., “Frequency analysis and sheared filtering for shadow light fields of complex occluders,” *ACM Transactions on Graphics* **30**(2) (2011).
- [45] Overbeck, R., Donner, C., and Ramamoorthi, R., “Adaptive Wavelet Rendering,” *ACM Transactions on Graphics (SIGGRAPH ASIA 09)* **28**(5) (2009).
- [46] Lehtinen, J., Aila, T., Chen, J., Laine, S., and Durand, F., “Temporal light field reconstruction for rendering distribution effects,” *ACM Transactions on Graphics (SIGGRAPH 11)* **30**(4) (2011).
- [47] Cevher, V., Sankaranarayanan, A., Duarte, M., Reddy, D., Baraniuk, R., and Chellappa, R., “Compressive sensing for background subtraction,” in [*European Conference on Computer Vision (ECCV 08)*], II–155–II–168 (2008).
- [48] Veeraraghavan, A., Reddy, D., and Raskar, R., “Coded strobing photography: Compressive sensing of high speed periodic videos.,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* **33**(4), 671–686 (2011).