

A Survey of Image-based Relighting Techniques

Biswarup Choudhury, Sharat Chandran

ViGIL

IIT Bombay

Mumbai

email: {biswarup, sharat}@cse.iitb.ac.in

www: www.cse.iitb.ac.in/{biswarup, sharat}

Abstract

Image-based Relighting (IBRL) has recently attracted a lot of research interest for its ability to relight real objects or scenes, from novel illuminations captured in natural/synthetic environments. Complex lighting effects such as subsurface scattering, interreflection, shadowing, mesostructural self-occlusion, refraction and other relevant phenomena can be generated using IBRL. The main advantage of image-based graphics is that the rendering time is independent of scene complexity as the rendering is actually a process of manipulating image pixels, instead of simulating light transport. The goal of this paper is to provide a complete and systematic overview of the research in Image-based Relighting. We observe that essentially all IBRL techniques can be broadly classified into three categories, based on how the scene/illumination information is captured: Reflectance function-based, Basis function-based and Plenoptic function-based. We discuss the characteristics of each of these categories and their representative methods. We also discuss about the sampling density and types of light source(s), relevant issues of IBRL.

Keywords: Image-based Relighting, Survey, Image-based Techniques, Augmented Reality.

1 Introduction

Image-based Modeling and Rendering (IBMR) synthesizes realistic images from pre-recorded images without a complex and long rendering process as in traditional geometry-based computer graphics. The major drawback of IBMR is its inherent rigidity. Most IBMR techniques assume a static illumination condition. Obviously, these assumptions cannot fully satisfy the computer graphics needs since illumination modification is a key operation in computer graphics. The ability to control illumination of the modeled scene enhances the three-dimensional illusion, which in turn improves viewers' understanding (Fig. 1). If the illumination can be modified by relighting the images, instead of rendering the geometric models, the time for image synthesis will be independent of the scene complexity. This also saves the artist/designer enormous amount of time in fine tuning the illumination conditions to achieve realistic atmospheres. Applications range from global illumination and lighting design to augmented and mixed reality, where real and virtual objects are combined with consistent illumination.

Two major motivations for IBRL are :

- Allows the user to vary illuminance of the whole (or only interesting portions of the) scene improving recognition and satisfaction.
- Brings us a step closer to realizing the use of image-based entities as the basic rendering primitives/entities.

Digital Peer Publishing Licence

Any party may pass on this Work by electronic means and make it available for download under the terms and conditions of the current version of the Digital Peer Publishing Licence (DPPL). The text of the licence may be accessed and retrieved via Internet at <http://www.dipp.nrw.de/>.

First presented at the International Conference on Computer Graphics Theory and Applications, 2006, extended and revised for JVRB



Figure 1: Specular Enhancement. Polynomial Texture Maps [MGW01].

For didactic purposes, we classify image-based relighting techniques into three categories, namely: Reflectance-based, Basis Function-based, Plenoptic Function-based. These categories should be actually viewed as a continuum rather than absolute discrete ones, since there are techniques that defy these strict categorizations.

Reflectance Function-based Relighting techniques explicitly estimate the reflectance function at each visible point of the object or scene. This is also known as the Anisotropic Reflection Model [Kaj85] or the Bidirectional Surface Scattering Reflectance Distribution Function (BSSRDF) [JMLH01]. It is defined as the ratio of the outgoing to the incoming radiance. Reflectance estimation can be achieved with calibrated light setup, which provide full control of the incident illumination. A reflectance function is modeled with the data of the scene, captured under varying illumination and view direction. Techniques then apply novel illumination and use the reflectance function calculated, for generating novel illumination effects in the scene.

Basis Function-based Relighting techniques take advantage of the linearity of the rendering operator with respect to illumination, for a fixed scene. Re-rendering is accomplished via linear combination of a set of pre-rendered “basis” images. These techniques, for the purpose of computing a solution, determine a time-independent basis – a small number of “generative” global solutions – that suffice to simulate the set of images under varying illumination and viewpoint.

Plenoptic Function-based Relighting techniques are based on the computational model, the Plenoptic Function [AB91]. The original plenoptic function aggregates all the illumination and the scene changing factors in a single “time” parameter. So most research

concentrates on view interpolation and leaves the time parameter untouched (illumination and scene static). The plenoptic function-based relighting techniques extract the illumination component from the aggregate time parameter, and facilitate relighting of scenes.

The remainder of the paper is organized as follows: Section 2, Section 3 and Section 4 respectively discuss each of the three relighting categories, along with their representative methods. In Section 5, we discuss some of the other relevant issues of relighting. We then provide some directions of future research in Section 6. Finally, we conclude with our final remarks in Section 7.

2 Reflectance Function

A reflectance function is the measurement of how materials reflect light, or more specifically, how they transform incident illumination into radiant illumination. The *Bidirectional Reflectance Distribution Function* (BRDF) [NRH⁺77] is a general form of representing surface reflectivity. A better representation is the *Bidirectional Surface Scattering Reflectance Distribution Function* (BSSRDF) [JMLH01], which model effects such as color bleeding, translucency and diffusion of light across shadow boundaries, otherwise impossible with a BRDF model. As introduced by [DHT⁺00], the reflectance function R , an 8D function, determines the light transfer between light entering a bounding volume at a direction and position $\psi_{incident}$ and leaving at $\psi_{exitant}$:

$$\mathbf{R} = \mathbf{R}(\psi_{incident}, \psi_{exitant})$$

This calculated reflectance function can be used to compute relit images of the objects, lit with novel il-

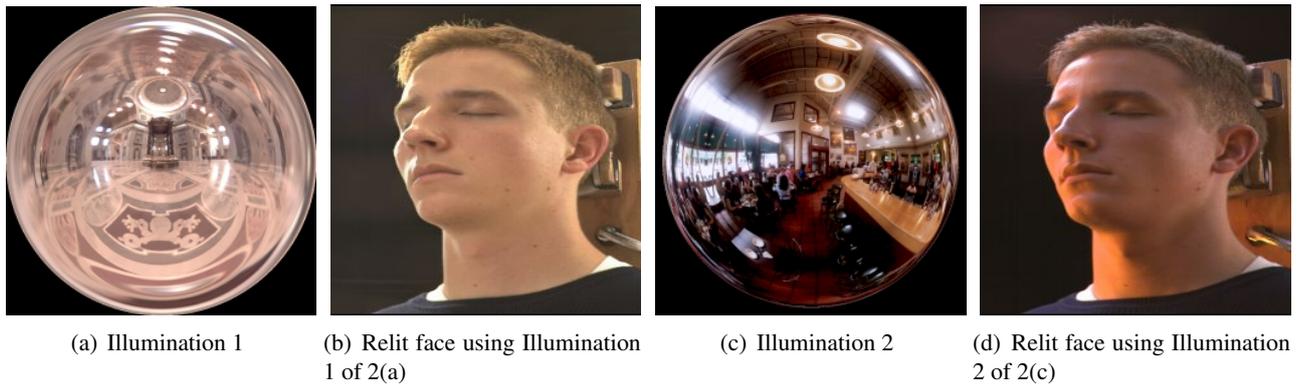


Figure 2: Mirrored ball, representing illumination of an environment, used for relighting faces [DHT⁺00].

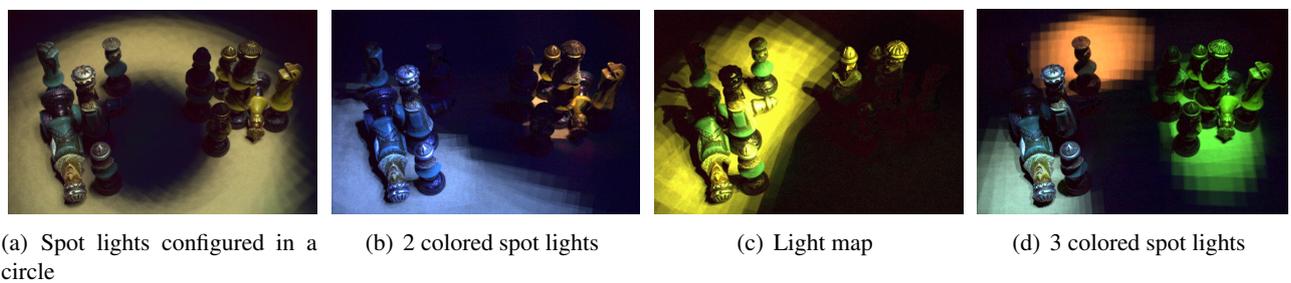


Figure 3: An arrangement of chess pieces, illuminated by different types of incident light fields [MPDW03].

illumination. The computation for each relit pixel is reduced to multiplying corresponding coefficients of the reflectance function and the incident illumination.

$$L_{exitant}(\omega) = \int_{\Omega} \mathbf{R}L_{incident}(\omega) d\omega$$

where, Ω is the space of all light directions over a hemisphere centered around the object to be illuminated ($\omega \in \Omega$). For every viewing direction, each pixel in an image stores its appearance under all illumination directions. Thus each pixel in an image is a sample of the reflectance function. The reflectance functions are sampled from real objects by illuminating the object from a set of directions while recording photographs from different viewpoints.

We classify the estimation of reflectance functions into three different categories: **Forward**, **Inverse** and **Pre-computed radiance transport**.

2.1 Forward

The forward methods of estimating reflectance functions sample these functions exhaustively and tabulate the results. For each incident illumination, they

store the reflectance function weights for a fixed observed direction. The forward method of estimating reflectance functions can further be divided into two categories, on the basis of illumination information provided, **Known** and **Unknown**.

The techniques with **known** illumination incorporate the information in their setup. The user is provided full control over the direction, position and type of incident illumination. This information is directly used for finding the reflectance properties of the scene. [DHT⁺00] use the highest resolution incident illumination with roughly 2000 directions and construct a reflectance function for each observed image pixel from its values over the space of illumination directions (Fig. 2). [MPDW04] sample the reflectance functions from real objects by illuminating the object from a set of directions while recording the photographs. They reconstruct a smooth and continuous reflectance function, from the sampled reflectance functions, using the multilevel B-spline technique. [MPDW03] exploit the richness in the angular and spatial variation of the incident illumination, and measure six-dimensional slices of the eight-dimensional reflectance field, for a fixed

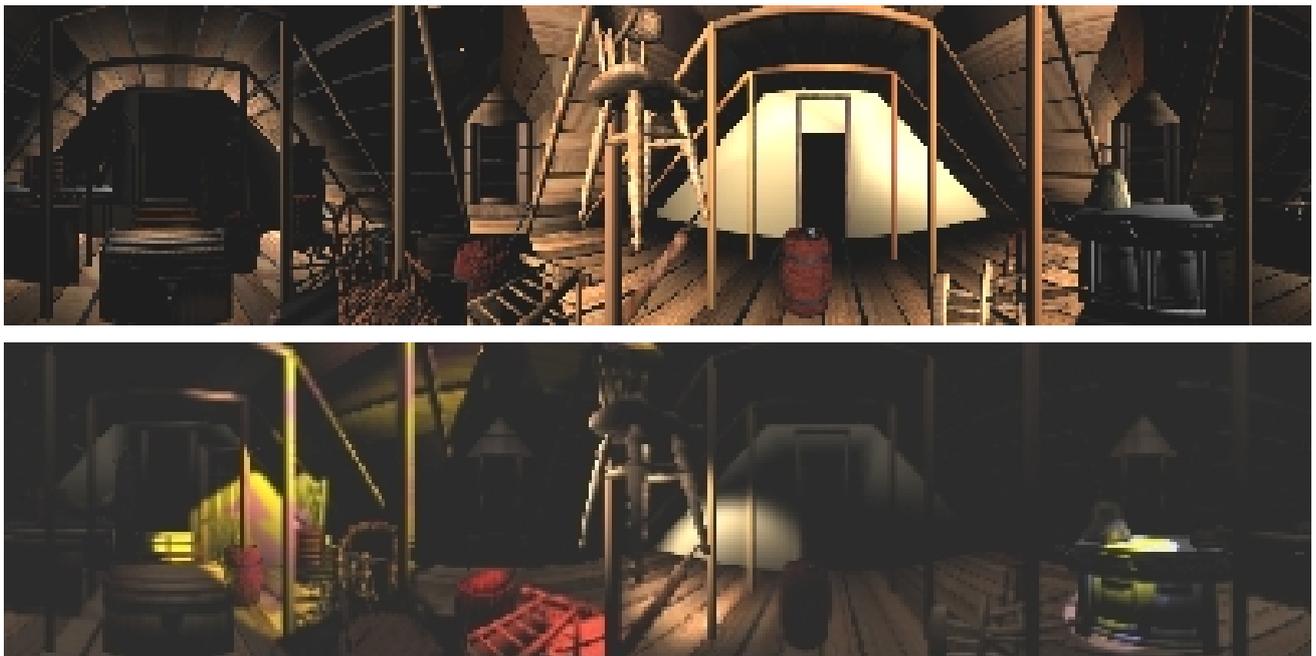


Figure 4: Attic relighting using conventional light sources(top row) and 5 spot lights(bottom row) [WHF01].

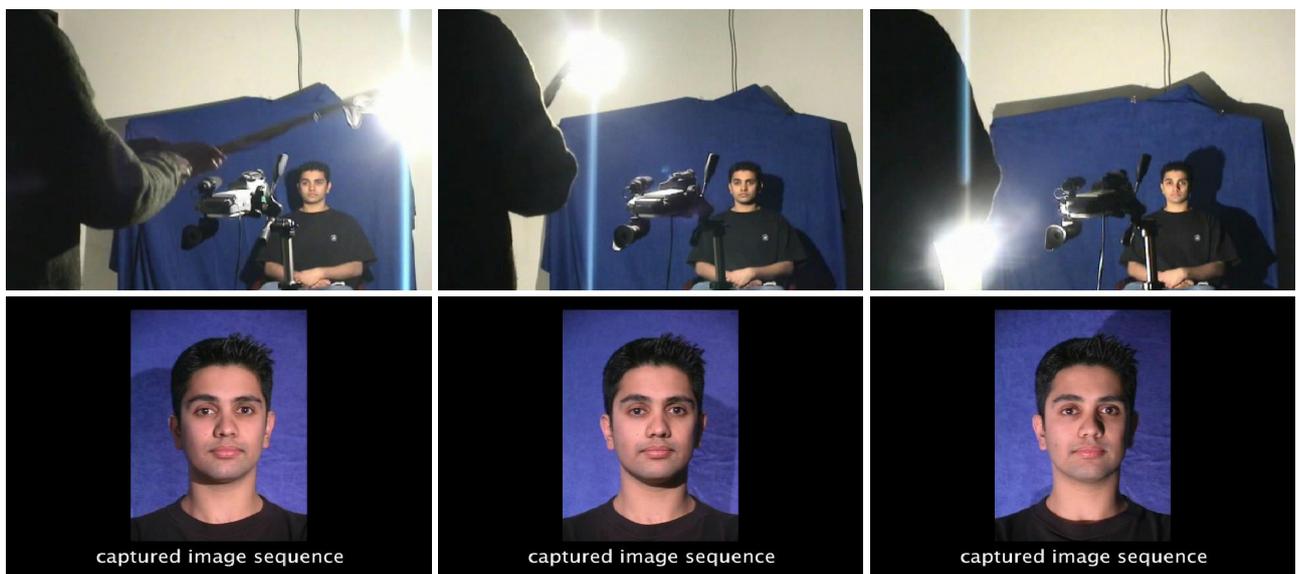


Figure 5: A light source is waived in front of the subject. Using the mathematical framework of the imaging system of the eye, one can estimate the illumination information [NN04].

viewpoint (Fig. 3). On the other hand, [MGW01] (Fig. 1) store the coefficients of a biquadratic polynomial for each texel, thereby improving upon the compactness of the representation, and uses it to reconstruct the surface color under varied illumination conditions. The reflectance functions and illumination are expressed by a set of common basis functions, en-

abling a significant speed-up in the relighting process. This method maintains all the high-frequency features such as highlights and self-shadowing.

[WHON97], [WHF01] propose a concept of *apparent-BRDF* to represent the outgoing radiance distribution passing through the pixel window on the image plane. By treating each image as an ordinary



Figure 6: Four relighting examples(top row) as linear combination of images, the coefficients being defined by novel images of a probe object(bottom, left image) which are reconstructed with the sampled probe images(right) [FBS05].

surface element, the radiance distribution of the pixel under various illumination conditions is recorded in a table (Fig. 4). [KBMK01] samples the surface's incident field to reconstruct a non-parametric apparent BRDF at each visible point on the surface. [BG01] iteratively produces an approximation of the reflectance model of diffuse, specular, isotropic or anisotropic textured objects using a single image and the 3D geometric model of the scene.

Techniques with **unknown** incident illumination information estimate the information. [NN04] use eyes of a human subject and compute a large field of view of the illumination distribution of the environment surrounding a person, using the characteristics of the imaging system formed by the cornea of an eye and a camera viewing it (Fig. 5). Their assumption of a human subject in the scene, at all times, may not be practical though. [LKG⁺03] used six steel spheres to recover the light source positions. They fit an average BRDF function to the different materials of the objects in the scene. Some other techniques [FBS05], [TSE⁺04] indirectly compute the incident illumination information by using black snooker ball/non-metallic sphere (Fig. 6). A very early work of IBRL, *Inverse Rendering* [Mar98], solves for unknown lighting and reflectance properties of a scene, for relighting purposes.

2.2 Inverse

The inverse problem of estimation of reflectance functions can be stated as follows: *Given an observation, what are the weights and parameters of the basis functions that best explain the observation?*

Inverse methods observe an output and compute the probability that it came from a particular region in the

incident illumination domain. The incident illumination is typically represented by a bounded region, such as an environment map, which is modeled as a sum of basis functions [rectangular [ZWCS99] or Gaussian kernels [CZH⁺00]] (Fig. 7). They capture an *environment matte*, which in addition to capturing the foreground object and its traditional matte, also describes how the object refracts and reflects light. This can then be placed in a new environment, where it will refract and reflect light from that scene. Techniques [MLP04], [PD03] have been proposed which progressively refine the approximation of the reflectance function with an increasing number of samples.

For a more accurate reflectance estimation, [MPZ⁺02] combine a forward method [DHT⁺00] for the low-frequency surface reflectance function and an inverse method, environment matting [CZH⁺00], for the high-frequency surface reflectance function. This is used for capturing all the complex lighting effects, like high-frequency reflections and refractions.

2.3 Pre-computed Radiance transport

A global transport simulator creates functions over the object's surface, representing transfer of arbitrary incident lighting, into transferred radiance which includes global effects like shadows, self-interreflections, occlusion and scattering effects. When the actual lighting condition is substituted at run-time, the resulting model provides global illumination effects.

The radiance transport is pre-computed using a detailed model of the scene [SKS02]. To improve upon the rendering performance, the incident illumination can be represented using spherical harmonics [KSS02], [RH01], [SKS02] or wavelets [NRH03]. The reflectance field, stored per vertex as a transfer



Figure 7: A water goblet and a pie tin digitally composed onto background images, preserving the effects of refraction [ZWCS99] [CZH⁺00].



(a) View 1: Relit Scene 1 (b) View 1: Relit Scene 2 (c) View 1: Relit Scene 3 (d) View 1: Relit Scene 4



(e) View 2: Relit Scene 1 (f) View 2: Relit Scene 2 (g) View 2: Relit Scene 3



(h) View 3: Relit Scene 1 (i) View 3: Relit Scene 2 (j) View 3: Relit Scene 3

Figure 8: Pre-computed radiance transport based IBRL, depicting all-frequency shadows, reflections and highlights along with view variation [NRH04].

matrix, can be compressed using PCA [SHHS03] or wavelets [NRH03].

[NRH04] focuses on relighting for changing illumina-

tion and viewpoint, while including all-frequency shadows, reflections and lighting (Fig. 8). They propose a novel triple product integrals based technique

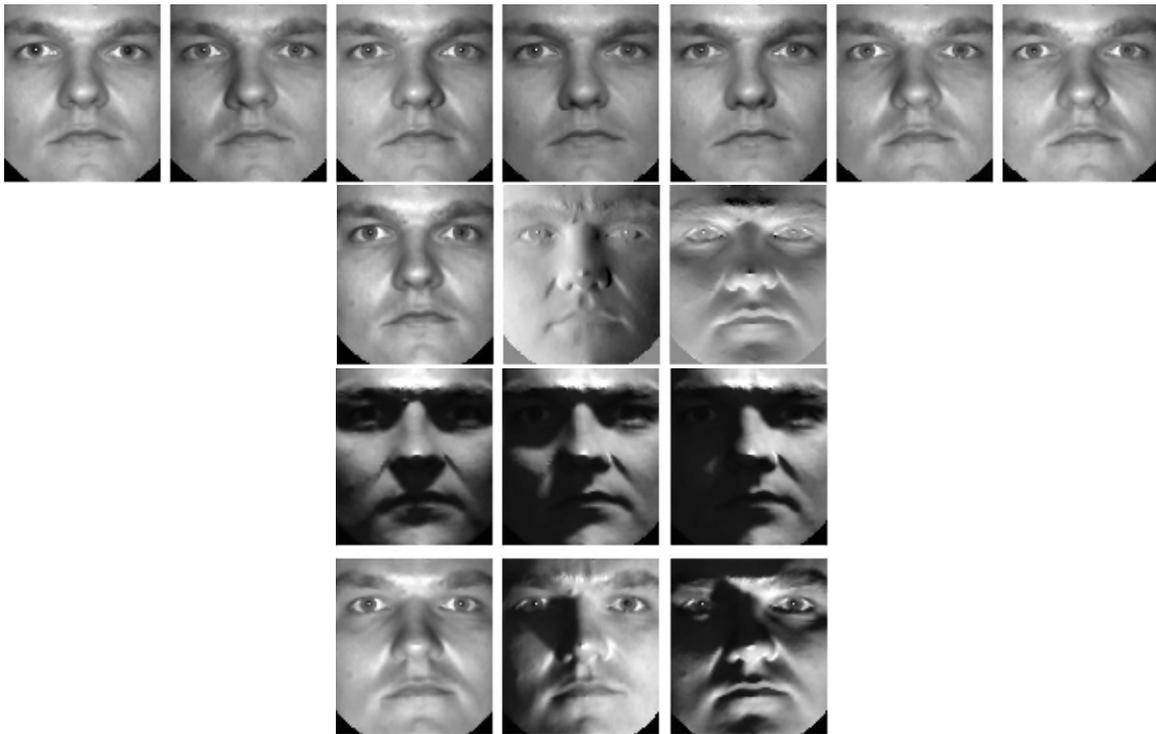


Figure 9: Illuminating faces using Illumination cones. First row: Input dataset, Second row: Basis Images, Third and Fourth row: Relit Faces [GKB98][GBK01].

of factorizing the visibility and the material properties. Recently, [WTL05] presented a method of relighting translucent objects under all-frequency lighting. They apply a two-pass hierarchical technique for computing non-linearly approximated transport vectors due to diffuse multiple scattering.

3 Basis Function

Basis Function based techniques decompose the luminous intensity distributions into a series of basis functions, and illuminances are obtained by simply summing luminance from each light source whose luminous intensity distribution obey each basis function. Assuming multiple light sources, luminance at a certain point is obtained by calculating the luminance from each light source and summing them. In general, luminance calculation obeys the two following rules of superposition:

1. The image resulting from an additive combination of two illuminants is just the sum of the images resulting from each of the illuminations independently.
2. Multiplying the intensity of the illumination

sources by a factor of α results in a rendered image that is multiplied by the same factor α .

These techniques calculate luminance in the case of alterations in the luminous distributions and the direction of light sources. The luminous intensity distribution of a point light source is expressed as the sum of a series of basis distributions. Luminance due to light source, whose luminance intensity distribution corresponds to one of the basis distributions, is calculated in advance and stored as basis luminance. Using the aforementioned property 1, the luminance due to the light source with luminous intensity distribution, is calculated by summing the pre-calculated basis luminances corresponding to each individual basis distribution. Using property 2, the luminance due to a light source, whose luminous intensity distribution can be expressed as the weighted sum of the basis distributions, is obtained by multiplying each basis luminance with corresponding weights and summing them. Thus, once the basis luminance is calculated in the pre-process, the resulting luminance can be obtained quickly by calculating the weighted sum of the basis luminances. Some desirable properties of a basis set of illumination functions are:

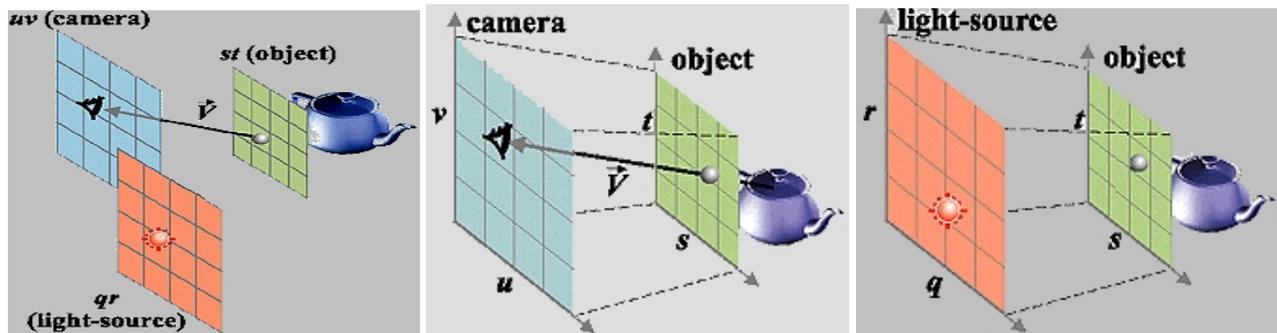


Figure 10: The tri-planar configuration and the dual light slab parametrization. Each captured value represents the radiance reflected through (s, t) and received at a certain (u, v) when the scene is illuminated by a point source positioned at (q, r) [LWS02].

1. The basis functions should be general enough to form any light source one desires.
2. The number of basis functions should be small, since this corresponds to the number of basis images we must actually store and render.

We classify the type of basis functions (used in Relighting) into five categories and provide their corresponding representative methods:

1. *Steerable Functions* [NSD94]: Under natural illuminants like the sun, whose direction varies along a path on the sphere, a desirable property of the basis is to be rotation invariant. A steerable function is one which can be written as a linear combination of rotated versions of itself.
2. *Spherical Harmonics* [DKNY95]: Spherical harmonics is an orthonormal basis defined over a unit sphere. In the context of interior lighting design, spherical harmonics can express light sources from different directions. They are effective for recovering illumination that is well localized in the frequency domain.
3. *Singular Value Decomposition/Principal Component Analysis* [GBK01] (Fig. 9), [GKB98], [OK01], [HWT⁺04]: These techniques are used for simplifying a dataset, or specifically reduce the dimensionality of a data set while retaining as much information as is possible. They compute a compact and optimal description of the data set.
4. *N-mode SVD: Multilinear algebra of higher-order Tensors* [VT04], [FKIS02], [SVLD03], [TZL⁺02]: It is an extension to tensors, of the conventional matrix singular value decomposition.

A major advantage of this model is that it is purely image-based, thereby avoiding the non-trivial problem of 3D geometric structure from the image data. Further, this formulation can disentangle and explicitly represent each of the multiple factors inherent to image formation-illumination change, viewpoint change etc.

4 Plenoptic Function

The appearance of the world can be thought of as the dense array of light rays filling the space, which can be observed by posing eyes or cameras in space. These light rays can be represented through the *plenoptic function* (from *plenus*, complete or full; and *optics*) [AB91]. As pointed by Adelson and Bergen:

The world is made of three-dimensional objects, but these objects do not communicate their properties directly to an observer. Rather, the objects fill the space around them with the pattern of light rays that constitutes the plenoptic function, and the observer takes samples from this function. The plenoptic function serves as the sole communication link between the physical objects and their retinal images. It is the intermediary between the world and the eye.

All basic visual measurements can be considered to characterize local change along one or more dimensions of a single function that describes the structure of the information in the light impinging on an observer. The plenoptic function represents the intensity of the light rays passing through the camera center at every location, at every possible viewing angle. The plenoptic function is a 7D function that models a 3D dynamic environment by recording the light rays at every space location (V_x, V_y, V_z) , from every possible direc-

tion (θ, ϕ) , over any range of wavelengths (λ) and at any time (t) , i.e.,

$$P = P^{(7)}(V_x, V_y, V_z, \theta, \phi, \lambda, t)$$

An image of a scene with a pinhole camera records the light rays passing through the camera's center-of-projection. They can also be considered as samples of the plenoptic function. Basically, the function tells us how the environment looks when our eye is positioned at $V=(V_x, V_y, V_z)$. The time parameter t actually models all the other unmentioned factors such as the change of illumination and the scene.

Plenoptic Function-based relighting techniques propose new formulations of the plenoptic function, which explicitly specify the illumination component. Using these formulations, one can generate complex lighting effects. One can simulate various lighting configurations such as multiple light sources, light sources with different colors and also arbitrary types of light sources (Section 5.1).

4.1 Representative Techniques

[WH04] discuss a new formulation of the plenoptic function, *Plenoptic Illumination Function*, which explicitly specifies the illumination component. They propose a *local illumination model*, which utilizes the rules of superposition for relighting under various lighting configurations. [LWS02] on the other hand, propose a representation of the plenoptic function, the *reflected irradiance field*. The reflected irradiance field stores the reflection of surface irradiance as an illuminating point light source moves on a plane. With the reflected irradiance field, the relit object/scene can be synthesized simply by interpolating and superimposing appropriate sample reflections (Fig. 10).

5 Discussion

In this section, we discuss some of the relevant issues involving IBRL. We also provide a brief comparison between all the categories of IBRL techniques.

5.1 Light Source type

Illumination is a complex and high-dimensional function of computer graphics. To reduce the dimensionality and to analyze their complexity and practicality, it is necessary to assume a specific type of light source. Two types of light sources most commonly used are:

1. **Directional Light Source (DLS):** A DLS emits parallel rays which do not diverge or become dimmer with distance. It is parametrized using only two variables (θ, ϕ) , which denotes the direction of the light vector. For planar surfaces lighted by a DLS, the degree of shading will be the same right across the surface. The computations required for directional lights are therefore considerably less. Using a DLS is also more meaningful, because the captured pixel value in an image tells us what the surface elements behind the pixel window *actually look like*, when all surface elements are illuminated by parallel rays in the direction of the viewing point. DLS serves well with synthetic object/scene where it is used to approximate the light coming from an extremely distant light source. But it poses practical difficulties for capturing real and large object/scene. They can be approximated with strong spotlights at a distance which greatly exceeds the size of the object/scene.
2. **Point Light Source (PLS):** A PLS shines uniformly in all directions. Its intensity decreases with the distance to the light source. A PLS is parametrized using three variables (P_x, P_y, P_z) , which denote the 3D position of the PLS in space. As a result, the angle between the light source and the normals of the various affected surfaces can change *dramatically* from one surface to the next. In the presence of multiple light sources, this means that for every vertex, one has to determine the direction of the light vector corresponding to a light source. This requires determination of the depth map of the images using computer vision algorithms, which though provide good approximations, make the lighting calculations computationally intensive. Point light source are usually close to the observer and so more practical for real and large objects/scenes.

5.2 Sampling

Sampling is one of the key issues of image-based graphics. It is a non-trivial problem because it involves the complex relationship among three elements: the depth and texture of the scene, the number of sample images, and the rendering solution. One needs to determine the *minimum sampling rate* for anti-aliased image-based rendering. Comparatively, very little research [CCST00], [SK00], [ZC04], [ZC01], [ZC03]

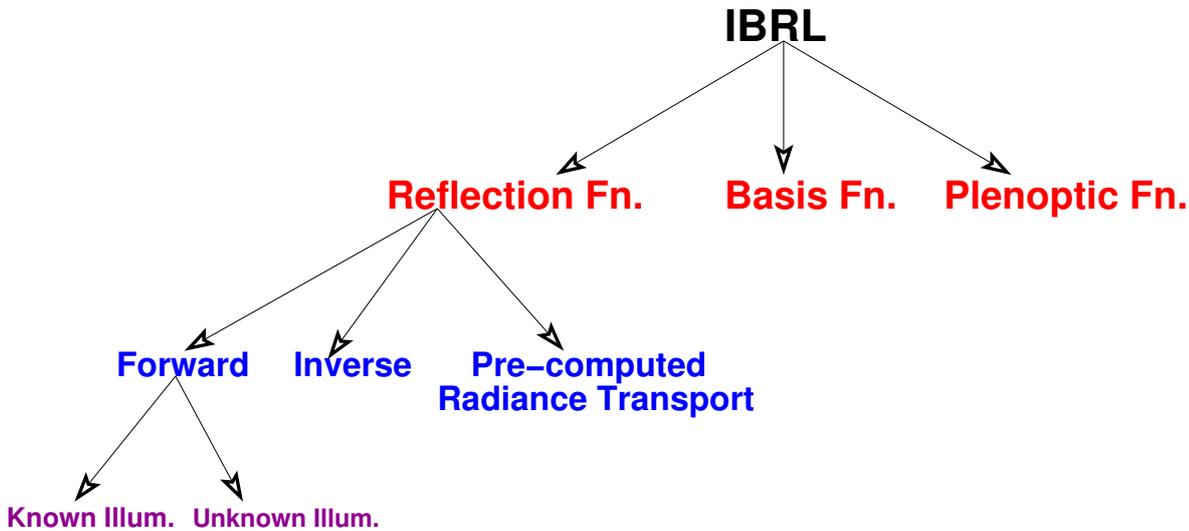


Figure 11: Pictorial representation of IBRL categorization

has gone into trying to tackle this problem.

In the context of IBRL, sampling deals with the illumination component for efficient and realistic re-lighting [WH04]. [LWS02] prove that there exists a *geometry-independent* bound of the sampling interval, which is analytically bound to the BRDF of the scene. It ensures that the intensity error in the relit image is smaller than a user-specified tolerance, thus eliminating noticeable artifacts.

5.3 Comparison

We compare all the IBRL categories (Fig. 11) on the basis of the following factors:

- *Object/Image Space*: Reflectance function-based relighting techniques are closest to the object space since they try to estimate a surface/object property, reflectance. Closest to the image space are plenoptic function-based relighting techniques. They heavily sample the scene (under different illumination conditions) and use these images for relighting. While reflection-based relighting and plenoptic function-based relighting techniques are at the extremes, basis function-based relighting techniques are in the confluence of both the object and image space.

Some reflectance function-based techniques (Pre-computed radiance transport), unlike the other two categories, require geometry of the scene

to estimate the reflectance of the objects in the scene.

- *Frequency components*: While both reflectance function and plenoptic function-based relighting techniques can capture both the high and low frequency components, basis function-based relighting techniques are usually capable of capturing only the low frequency effects. Forward reflectance function-based techniques capture the low frequency components and the inverse methods capture the high frequency components.
- *Sampling density*: For the purposes of IBRL, the sampling density of the corresponding techniques usually increases in the following order – basis function-based relighting, reflectance function-based relighting, plenoptic function-based relighting.
- *Applications*: Reflectance-based relighting techniques can be utilised to change the surface properties of an object in a scene (for example, converting a diffuse object to a specular one) in addition to doing IBRL. The techniques of the other two categories are not capable of generating such effects. They mainly illuminate the scene with a novel illumination with or without change in viewpoint.

6 Future directions

A lot of research remains to be done in IBRL. Some ideas being:

1. **Efficient Representation:** BRDF function-based IBRL techniques require huge number of samples to accurately estimate a reflectance function. Most techniques, for practical purposes, consider low-frequency components, which compromises with the visual quality of the rendered image. Almost all basis function-based techniques also require a number of images for relighting. Thus, a lot of research is required to find accurate and efficient representations of a scene, which capture all the complex phenomenas of lighting and reflectance functions. Variability in terms of viewpoint and illumination leads to huge data sets, which further incurs huge computational costs. A related area which deserves considerable investigation, is IBRL for real and large environments.
2. **Sampling:** Most techniques do not deal with the *minimum sampling density* required for anti-aliased IBRL. [LWS02] discuss about a geometry-independent sampling density based on *radiometric* tolerance. Though this serves our purpose of efficient sampling of certain scenes, what we need is a *photometric tolerance*, which takes into account the response function of human vision. [DMZ⁺05] discuss the importance of psychophysical quality scale for realistic IBRL of glossy surfaces.
3. **Compression:** No matter how much the storage and memory increase in the future, compression is always useful to keep the IBRL data at a manageable size. A high compression ratio in IBRL relies heavily on how good the images can be predicted. The sampled images for IBRL, usually have a strong inter-pixel and intra-pixel correlation, which needs to be harnessed for efficient compression. Currently, techniques such as spherical harmonics, vector quantization, direct cosine transform and spherical wavelets are used for compressing the datasets of IBRL, but all of these have their own inherent disadvantages.
4. **Dynamics:** Most IBRL techniques deal with static environments, in terms of change in geometry of the scene/object. With the development of high-end programmable graphics processors, it is

conceivable that IBRL can be applied to dynamic environments.

7 Final remarks

We have surveyed all the recent developments in the field of Image-based Relighting and, in particular, classified them into three categories based on how they capture the scene/illumination information: Reflectance Function-based, Basis Function-based and Plenoptic Function-based. We have presented each of the categories along with their corresponding representative methods. Relevant issues of IBRL like type of light source and sampling have also been discussed.

It is interesting to note the trade-off between geometry and images, needed for anti-aliased image-based rendering. Efficient representation, realistic rendering, limitations of computer vision algorithms and computational costs should motivate researchers to invent efficient image-based relighting techniques in future.

8 Acknowledgements

We would like to thank the members of ViGIL, IIT Bombay for their continued help and support.

This research was supported by the fellowship fund endowed by Infosys Technologies Limited, India.

References

- [AB91] E. H. Adelson and J. R. Bergen, *The plenoptic function and the elements of early vision*, Computational Models of Visual Processing (1991), 3–20.
- [BG01] Samuel Boivin and Andre Gagalowicz, *Image-based rendering of diffuse, specular and glossy surfaces from a single image*, SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press, 2001, pp. 107–116.
- [CCST00] Jin-Xiang Chai, Shing-Chow Chan, Heung-Yeung Shum, and Xin Tong, *Plenoptic sampling*, SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques (New York, NY,

- USA), ACM Press/Addison-Wesley Publishing Co., 2000, pp. 307–318.
- [CZH⁺00] Yung-Yu Chuang, Douglas E. Zongker, Joel Hindorff, Brian Curless, David H. Salesin, and Richard Szeliski, *Environment matting extensions: towards higher accuracy and real-time capture*, SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press/Addison-Wesley Publishing Co., 2000, pp. 121–130.
- [DHT⁺00] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar, *Acquiring the reflectance field of a human face*, SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press/Addison-Wesley Publishing Co., 2000, pp. 145–156.
- [DKNY95] Yoshinori Dobashi, Kazufumi Kaneda, Hideki Nakatani, and Hideo Yamashita, *A quick rendering method using basis functions for interactive lighting design*, Computer Graphics Forum **14** (1995), no. 3, 229–240.
- [DMZ⁺05] Olivier Dumont, Vincent Masselus, Peter Zaenen, Johan Wagemans, and Philip Dutre, *A perceptual quality scale for image-based relighting of glossy surfaces*, Tech. Rep. CW417, Katholieke Universiteit Leuven, June 2005, 2005.
- [FBS05] Martin Fuchs, Volker Blanz, and Hans-Peter Seidel, *Bayesian relighting.*, Rendering Techniques, 2005, pp. 157–164.
- [FKIS02] R. Furukawa, H. Kawasaki, K. Ikeuchi, and M. Sakauchi, *Appearance based object modeling using texture database: acquisition, compression and rendering*, EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering (Aire-la-Ville, Switzerland, Switzerland), Eurographics Association, 2002, pp. 257–266.
- [GBK01] Athinodoros S. Georghiades, Peter N. Belhumeur, and David J. Kriegman, *From few to many: Illumination cone models for face recognition under variable lighting and pose*, IEEE Trans. Pattern Anal. Mach. Intell. **23** (2001), no. 6, 643–660.
- [GKB98] A. S. Georghiades, D. J. Kriegman, and P. N. Belhumeur, *Illumination cones for recognition under variable lighting: Faces*, CVPR '98: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Washington, DC, USA), IEEE Computer Society, 1998, p. 52.
- [HWT⁺04] Tim Hawkins, Andreas Wenger, Chris Tchou, Andrew Gardner, Fredrik Göransson, and Paul E. Debevec, *Animatable facial reflectance fields.*, Rendering Techniques, 2004, pp. 309–321.
- [JMLH01] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan, *A practical model for subsurface light transport*, SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press, 2001, pp. 511–518.
- [Kaj85] James T. Kajiya, *Anisotropic reflection models*, SIGGRAPH '85: Proceedings of the 12th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press, 1985, pp. 15–21.
- [KBMK01] Melissa L. Koudelka, Peter N. Belhumeur, Sebastian Magda, and David J. Kriegman, *Image-based modeling and rendering of surfaces with arbitrary brdfs.*, CVPR (1), 2001, pp. 568–575.
- [KSS02] Jan Kautz, Peter-Pike Sloan, and John Snyder, *Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics*, EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering (Aire-la-Ville, Switzerland,

- Switzerland), Eurographics Association, 2002, pp. 291–296.
- [LKG⁺03] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel, *Image-based reconstruction of spatial appearance and geometric detail*, ACM Trans. Graph. **22** (2003), no. 2, 234–257.
- [LWS02] Zhouchen Lin, Tien-Tsin Wong, and Heung-Yeung Shum, *Relighting with the reflected irradiance field: Representation, sampling and reconstruction*, Int. J. Comput. Vision **49** (2002), no. 2-3, 229–246.
- [Mar98] Stephen Robert Marschner, *Inverse rendering for computer graphics*, Ph.D. thesis, Cornell University, 1998, Adviser-Donald P. Greenberg.
- [MGW01] Tom Malzbender, Dan Gelb, and Hans Wolters, *Polynomial texture maps*, SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press, 2001, pp. 519–528.
- [MLP04] Wojciech Matusik, Matthew Loper, and Hanspeter Pfister, *Progressively-refined reflectance functions from natural illumination.*, Rendering Techniques, 2004, pp. 299–308.
- [MPDW03] Vincent Masselus, Pieter Peers, Philip Dutre, and Yves D. Willems, *Relighting with 4d incident light fields*, ACM Trans. Graph. **22** (2003), no. 3, 613–620.
- [MPDW04] Vincent Masselus, Pieter Peers, Philip Dutre, and Yves D. Willems, *Smooth reconstruction and compact representation of reflectance functions for image-based relighting.*, Rendering Techniques, 2004, pp. 287–298.
- [MPZ⁺02] Wojciech Matusik, Hanspeter Pfister, Remo Ziegler, Addy Ngan, and Leonard McMillan, *Acquisition and rendering of transparent and refractive objects*, EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering (Aire-la-Ville, Switzerland, Switzerland), Eurographics Association, 2002, pp. 267–278.
- [NN04] Ko Nishino and Shree K. Nayar, *Eyes for relighting*, ACM Trans. Graph. **23** (2004), no. 3, 704–711.
- [NRH⁺77] F.E. Nicodemus, J.C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis, *Geometric considerations and nomenclature for reflectance*, NBS Monograph 160, National Bureau of Standards (US), 1977.
- [NRH03] Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan, *All-frequency shadows using non-linear wavelet lighting approximation*, ACM Trans. Graph. **22** (2003), no. 3, 376–381.
- [NRH04] R. Ng, Ravi Ramamoorthi, and Pat Hanrahan, *Triple product wavelet integrals for all-frequency relighting*, ACM Trans. Graph. **23** (2004), no. 3, 477–487.
- [NSD94] Jeffrey S. Nimeroff, Eero Simoncelli, and Julie Dorsey, *Efficient Re-rendering of Naturally Illuminated Environments*, Fifth Eurographics Workshop on Rendering (Darmstadt, Germany), Springer-Verlag, 1994, pp. 359–373.
- [OK01] Margarita Osadchy and Daniel Keren, *Image detection under varying illumination and pose.*, ICCV, 2001, pp. 668–673.
- [PD03] Pieter Peers and Philip Dutre, *Wavelet environment matting*, EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering (Aire-la-Ville, Switzerland, Switzerland), Eurographics Association, 2003, pp. 157–166.
- [RH01] Ravi Ramamoorthi and Pat Hanrahan, *An efficient representation for irradiance environment maps*, SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press, 2001, pp. 497–500.

- [SHHS03] Peter-Pike Sloan, Jesse Hall, John Hart, and John Snyder, *Clustered principal components for precomputed radiance transfer*, ACM Trans. Graph. **22** (2003), no. 3, 382–391.
- [SK00] Heung-Yeung Shum and Sing Bing Kang, *A review of image-based rendering techniques*, IEEE/SPIE Visual Communications and Image Processing (VCIP), 2000, pp. 2–13.
- [SKS02] Peter-Pike Sloan, Jan Kautz, and John Snyder, *Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments*, SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press, 2002, pp. 527–536.
- [SVLD03] F. Suykens, K. Vom, A. Lagae, and P. Dutre, *Interactive rendering with bidirectional texture functions*, Computer Graphics Forum **22** (2003), no. 3.
- [TSE⁺04] Chris Tchou, Jessi Stumpfel, Per Einarsson, Marcos Fajardo, and Paul Debevec, *Unlighting the parthenon*, SIGGRAPH 2004 Sketch, 2004.
- [TZL⁺02] Xin Tong, Jingdan Zhang, Ligang Liu, Xi Wang, Baining Guo, and Heung-Yeung Shum, *Synthesis of bidirectional texture functions on arbitrary surfaces*, SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press, 2002, pp. 665–672.
- [VT04] M. Alex O. Vasilescu and Demetri Terzopoulos, *Tensortextures: multilinear image-based rendering*, ACM Trans. Graph. **23** (2004), no. 3, 336–342.
- [WH04] Tien-Tsin Wong and Pheng-Ann Heng, *Image-based relighting: representation and compression*, Integrated image and graphics technologies (2004), 161–180.
- [WHF01] Tien-Tsin Wong, Pheng-Ann Heng, and Chi-Wing Fu, *Interactive relighting of panoramas*, IEEE Comput. Graph. Appl. **21** (2001), no. 2, 32–41.
- [WHON97] Tien-Tsin Wong, Pheng-Ann Heng, Siu-Hang Or, and Wai-Yin Ng, *Image-based rendering with controllable illumination*, Proceedings of the Eurographics Workshop on Rendering Techniques '97 (London, UK), Springer-Verlag, 1997, pp. 13–22.
- [WTL05] Rui Wang, John Tran, and David Luebke, *All-frequency interactive relighting of translucent objects with single and multiple scattering*, ACM Trans. Graph. **24** (2005), no. 3, 1202–1207.
- [ZC01] C. Zhang and T. Chen, *Generalized plenoptic sampling*, Tech. Rep. AMP01-06, Carnegie Mellon Technical Report, 2001., 2001.
- [ZC03] Chen. Zhang and T. Chen, *Spectral analysis for sampling image-based rendering data*, IEEE Transaction on Circuit, System on Video Technology 13, 11 (2003), 1038– 1050. 1, 2003.
- [ZC04] C. Zhang and T. Chen, *A survey on image-based rendering: Representation, sampling and compression*, SPIC 2004 **19** (2004), no. 1, 1–28.
- [ZWCS99] Douglas E. Zongker, Dawn M. Werner, Brian Curless, and David H. Salesin, *Environment matting and compositing*, SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques (New York, NY, USA), ACM Press/Addison-Wesley Publishing Co., 1999, pp. 205–214.