Chapter 7

Conclusions and Future Work

In this dissertation, we have presented a new way of analyzing a basic building block in computer graphics rendering algorithms—the computational interaction between illumination and the reflective properties of a surface. We formalize the notion of reflection on a curved surface as a spherical convolution of the incident illumination and the bi-directional reflectance distribution function of the surface. This allows us to develop a signal-processing framework for reflection, leading to new frequency domain and combined angular and frequency-space methods for forward and inverse problems in computer graphics, taking into account the effects of complex natural illumination, and physically realistic BRDFs. We use the theoretical framework to develop new practical algorithms for interactive forward rendering using environment maps and inverse rendering to estimate lighting and material properties under complex illumination conditions.

The main thesis validated by this dissertation is that a deeper understanding of the computational nature of reflection and illumination is important and leads to new practical algorithms. Secondly, since the properties of the reflection operator impact a number of different areas, the same fundamental insights lead to novel algorithms for a number of different problems in forward and inverse rendering in graphics, as well as related areas in computer vision. Since this dissertation has both a theoretical and practical component, we discuss below our main conclusions, the bigger picture, and future directions for each of the components.

7.1 Computational Fundamentals of Reflection

Although the physical phenomena and mathematical framework for reflection in the spatial or angular domain are well known, there has not so far been a deep understanding of the computational issues. Therefore, it is important to lay the mathematical foundations for this area, which I refer to as the computational fundamentals of reflection. This dissertation has presented one such analysis in terms of signal-processing. Although the qualitative notion of reflection as convolution has a fairly long history of nearly 20 years in the graphics community and a number of people have tried to formalize this idea, this is the first time a precise mathematical description has been derived, with an associated convolution theorem. The analysis in chapters 2 and 3 has led to a number of basic questions being answered including the fundamental limits and conditioning of inverse problems, i.e. from observations of an object, what can we know about the lighting/reflectance? Our analysis also leads to a number of interesting observations, some of which may be lines for future work.

Relationship between Forward and Inverse Problems: Firstly, we have demonstrated a significant duality between forward and inverse problems. An ill-posed or ill-conditioned inverse problem is one where the input signal cannot be estimated from the output, because it is attenuated or truncated by the reflection operator (or BRDF filter). An example is illumination estimation from a diffuse or Lambertian surface. However, this also means that in the forward direction, most of the information in the input signal is not required to synthesize the output image. In chapter 4, we saw that this made it very efficient to represent and compute irradiance maps corresponding to reflections from diffuse or Lambertian surfaces. It is likely that similar considerations of duality can be profitably employed to analyze and create efficient and robust algorithms for forward and inverse problems in the future.

Dual Angular and Frequency-Space Representations: The Heisenberg uncertainty principle can effectively be interpreted as saying that quantities cannot simultaneously be localized in the spatial (angular) and frequency domains. Thus, quantities local in angular space have broad frequency spectra and vice-versa. However, we can turn this observation to our advantage by making use of both angular and frequency domains simultaneously.

By developing a frequency-space view of reflection, we ensure that we can use either the angular-space or frequency-space representation, or even a combination of the two. The diffuse BRDF component is slowly varying in angular-space, so we represent it in the frequency domain where it is a very local low-pass filter. Similarly, it is often useful to represent the specular BRDF component in the angular domain, where it is local, as opposed to the frequency domain where a large number of terms are required. This is the argument behind our representation in section 6.3. This principle of duality is likely to have many further implications, with algorithms being designed separately for high and low-frequency components, and investigation of tradeoffs and crossover points between the two.

If we seek to extend and apply our signal-processing framework in cases where there is no specific convolution formula, then it is likely that significant benefit can be obtained using wavelet or multiresolution representations, one of whose properties is to create basis functions local in both spatial (angular) and frequency domains. Such representations might provide a unified way to account for both high frequency (local in the angular domain) and low-frequency (local in the frequency domain) components.

Sampling Theory: An obvious extension of the signal-processing framework is to develop a sampling theory for various problems in graphics, determining how many observations or samples are needed for reconstruction. For instance, in the Lambertian case, our theory indicates that only 9 samples are required to reconstruct the irradiance map, if resampling is done correctly. A general framework for determining reflection map resolutions and sampling rates is described in chapter 5. More generally, we seek in future to address some important problems in graphics and vision. Within the context of image-based rendering, a long-standing question is how many images are needed to produce the effects we desire. Too many, and data acquisition and manipulation become unwieldy. Too few, and there are artifacts and we fail to get the effects we want. Some work on this topic under the assumption of Lambertian surfaces and with imprecise geometry has been presented by Chai et al. [33]. It is interesting to consider a framework based on our analysis, assuming good or perfect geometry, but taking into account the effects of complex illumination and reflection models. A number of questions need to be answered such as the number of images required for good reconstruction and the right interpolation approach taking into account the frequency variation of the various quantities.

Low-Dimensional Subspaces: One of the results of our theory is formalizing the notion of low-dimensional subspaces for images of a Lambertian object (9D) and similar results for other reflection functions. Similar ideas have been used by Nimeroff et al. [63] and Teo et al. [82] for lighting design. One interesting question is whether the theoretical analysis in this dissertation can be seen as a special case of a general framework for determining good low-dimensional subspace approximations of the reflection operator, and whether our theoretical analysis can be adapted to cases where information on probable lighting conditions, reflection functions and object geometry are known. One approach is to find the principal components or basis function images of the space of all possible images of an object. This is a common approach used for compression in graphics and vision, but so far without much understanding of the number of principal components required. A preliminary step has been taken by us [69], where we have shown how to analytically construct the PCA for the space of images of a Lambertian object under all possible distant illumination conditions. We have shown that under idealized assumptions, the principal components reduce to the spherical harmonic basis functions, but must be modified for a single viewpoint where only the front-facing normals are visible. We have also shown how to adapt the 9D Lambertian result to derive lower-dimensional subspaces when we have only a single viewpoint. A promising future direction is to develop a general theory allowing for derivation of low-dimensional results under arbitrary assumptions about lighting, reflectance, geometry, and the effects to be considered in the reflection operator.

Integral Equations: It may be possible to view our convolution result as an analytic formula for a special case of the integral equation corresponding to the reflection operator. Similarly, the convolution result and the spherical harmonic basis functions may be the optimal low-dimensional approximations under idealized assumptions. This leads to the possibility of analyzing the integral equation corresponding to the reflection and rendering operators to determine the ill-posedness and conditioning of inverse problems, and the rank of the kernel directly to determine low-dimensional subspaces.

Differential Framework: Of course, signal-processing is only one way of conducting a computational analysis of the reflection operator. Another approach is to consider a differential framework for reflection. This answers questions about how to interpolate and extrapolate neighbouring views, or how much information is available from nearby observations, such as the slightly separated two eyes. A significant amount of work has been carried out in the image domain in the computer vision community under the aegis of optical flow methods. However, we believe significant insight and practical algorithms may emerge from considering an analysis on the object surface, taking into account complex illumination and reflection functions.

Perception: This dissertation has focussed on the physical computation of light transport. However, humans do not perceive on a linear scale; rather, the perceptual response is logarithmically related to the physical intensity. This has been exploited within the context of perceptual theories like the retinex framework [46] that often perform a logarithmic transformation before applying particular algorithms. In this case, separation of high-frequency texture and low-frequency illumination is made easier after a logarithmic transform. Similarly, in computer graphics, a number of homomorphic factorization techniques [47, 57] have been applied to represent reflected light fields and BRDFs by considering the error and doing computations in a logarithmic space. All of these results lead to a number of interesting questions about how approximations should be done, whether nonlinearities are introduced and whether a perceptual-space analysis and convolution relation can be derived. Considering perceptual aspects remains an important direction of future work.

7.2 High Quality Interactive Rendering

The first practical application discussed in the dissertation is of interactive rendering with natural illumination and physically based BRDFs. High quality interactive rendering is important in many applications including visualization, simulation and training, and video games. However, current graphics hardware rendering is usually limited to very simple lighting models (usually point and directional sources only) and reflection functions (Phong and Lambertian BRDFs). This imposes a gap between photorealism and interactivity that we have seeked to bridge. Within this context, we expect a number of future improvements.

Interactive rendering with complex illumination, shading, visibility: In this dissertation, we have addressed interactive rendering with complex illumination and reflectance functions. Several further steps may be taken to improve the quality of interactive computer graphics. Future work will encompass effects like spatially varying illumination, interreflection, and cast shadows. In the long term, we would like to see all the effects currently expensive to compute with global illumination simulations, incorporated into interactive renderings.

Multidimensional rendering: Another area of future work concerns multidimensional rendering. Most current computer graphics algorithms are optimized for creating a single image. However, many current computer graphics techniques deal with multi-view, multiple-light or other dimensions of data. While inverse rendering techniques to infer parameters for hardware rendering of multidimensional image spaces have been demonstrated by Hakura [26], there has been very little work on efficient multidimensional rendering. Halle [29] has shown how to efficient scan-convert a scene for multiple views, but much future work needs to be done in the context of efficiently rendering surface light fields or other multidimensional animations. We believe the fast prefiltering algorithms in sections 4 and 5 provide an important first step in this endeavor.

Application of computational fundamentals: Many of the theoretical analyses discussed in the previous subsection can be applied to a number of problems in rendering. For instance, differential analysis can be applied to accurate interpolation, and incremental changes to programmable shading calculations. Determining appropriate sampling rates can lead to new methods for antialiasing or filtering bump maps. Low-dimensional subspace methods and angular-frequency tradeoffs could be profitably applied in a number of domains in rendering.

7.3 Inverse Rendering

Realism in computer-generated images requires accurate input models for lighting, textures and reflectances (BRDFs). One of the best ways of obtaining high-quality data is through measurements of scene attributes from real photographs by **inverse rendering**. Measuring scene attributes also introduces structure into the raw imagery, allowing an artist to independently manipulate the material properties or lighting. This dissertation has presented a coherent mathematical framework for inverse rendering under general illumination conditions. Besides the formal study of the well-posedness and conditioning of inverse problems, we have derived new practical representations and novel frequency domain and dual spatial and frequency domain algorithms for BRDF and lighting estimation. Chapter 6 demonstrates examples of all these algorithms, and shows that the rendered images appear nearly identical to real photographs.

There remain some limitations in our work that we seek to address in the future, such as handling interreflections and extending the results to more complex illumination conditions and entire scenes. Besides this, there remain a number of problems to be solved.

Structured BRDF and higher-order representations: There are a number of methods for estimating BRDFs. These lead to incomplete information; from this, a full BRDF must be reconstructed for practical use in applications. One future direction is to explore how various representations—parametric, generative, and tabulated can be fit as data of increasing precision arrives. We also wish to study ideas with respect to adaptive reconstruction. Also, based on differential analysis, a number of higher-order differential properties are worth estimating. Finally, structured representations of the high-dimensional datasets in graphics today are very important for representation, understanding, manipulation and compression.

Factored BRDF representations: There has been relatively little theoretical or practical work on BRDFs in between low-parameter models and full measured representations. This remains a subject of future work. It is interesting to consider curve-based or factored BRDF models, and how they can be estimated from a small number of photographs, under both simple and complex illumination conditions. More generally, the idea of *factoring* the reflected light field, or of constructing structured representations of the high-dimensional datasets currently prevalent in graphics (such as BRDFs, BTFs and light fields) is one of growing interest.

Combined estimation of illumination, materials and geometry: Still another interesting question is whether we can combine inverse rendering to find illumination and material properties with geometry estimation, or at least use reflection information to determine improved estimates of geometric properties. In any case, our work is clearly applicable in the context of computer vision algorithms that determine geometry; we seek to extend these methods to incorporate complex illumination and reflectance properties.

General frameworks for extrapolatory rendering: Taking a step back, our goals really extend beyond estimating material properties. In computer graphics, we often wish to actually use the difference (or "error") between the photographs and the predicted model. It is often this difference which gives the photograph its real-world, as opposed to computer-generated, appearance. A simple example is the use of texture maps, wherein the real-world texture can be used to modulate or multiply shading computations from a light transport model that assumes homogeneous objects.

This is very different from most use of empirical data elsewhere in science—where we expect the theory to fit the data, treating errors as experimental noise to be ignored. Thus, one goal of future research is to derive a general framework for combining empirical observations in the form of real photographs with a model-based approach, perhaps relying on inverse rendering, to get very realistic images that can be easily manipulated using standard graphics operations. For instance, we could manipulate the low-frequency information predicted by our rendering model, such as illumination effects, while preserving the high-frequency information in the photograph, such as object texture.

A general framework along these lines would allow us to unify inverse and image-based (interpolatory) rendering methods. Similar ideas are likely to be applicable to animation, where we want to effectively edit or manipulate motion-capture data, and modeling, where we want to manipulate range data.

7.4 Summary

In summary, this dissertation has presented a new approach toward forward and inverse problems in graphics and vision. We have analyzed the computational properties of the reflection operator, in this case in terms of signal-processing. The ideas from this analysis lead to new robust and efficient algorithms for domains that have hitherto been considered largely separate—interactive high quality rendering, and measurement of lighting and material properties from photographs by inverse rendering. Basri and Jacobs [2, 3] have shown the further application of these ideas to the problems of photometric stereo and lighting-invariant recognition in computer vision. This suggests a new approach to problems in graphics and vision. We first make progress in understanding the *computational fundamentals of reflection*. This in turn leads to broad practical impact for forward and inverse problems in rendering and computer vision—domains that have hitherto been considered largely distinct. This dissertation has presented one such investigation in terms of signal-processing, and we believe there is much progress to be made.