

Principles of Appearance Acquisition and Representation

SIGGRAPH 2008 Class Notes

Tim Weyrich
Princeton University
USA

Jason Lawrence
University of Virginia
USA

Hendrik Lensch
Max-Planck-Institut für Informatik
Germany

Szymon Rusinkiewicz
Princeton University
USA

Todd Zickler
Harvard University
USA

Los Angeles, August 2007

Class Description

Algorithms for scene understanding and realistic image synthesis require accurate models of the way real-world materials scatter light. This class describes recent work in the graphics community to measure the spatially- and directionally-varying reflectance and subsurface scattering of complex materials, and to develop efficient representations and analysis tools for these datasets. We describe the design of acquisition devices and capture strategies for BRDFs and BSSRDFs, efficient factored representations, and a case study of capturing the appearance of human faces.

Prerequisites

Basic familiarity with the computer graphics pipeline, together with some knowledge of linear algebra and calculus.

Slides

The slides presented in class are available as separate supplemental material. They are closely aligned with the class notes and may be used as an additional reference.

Short Biographies

Jason Lawrence is an Assistant Professor in the Computer Science Department at the University of Virginia. He holds a Ph.D. in computer science from Princeton University. Jason's principal research interests focus on the acquisition and efficient representation of real-world material appearance. He recently introduced an "Inverse Shade Trees" framework for representing measured surface reflectance through a hierarchical decomposition designed for efficient rendering and editing.

Hendrik P. A. Lensch is the head of an independent research group "General Appearance Acquisition and Computational Photography" at the MPI Informatik in Saarbruecken, Germany. The group is part of the Max Planck Center for Visual Computing and Communication (Saarbruecken / Stanford). He received his diploma in computer science from the University of Erlangen in 1999. He worked as a research associate at Hans-Peter Seidel's computer graphics group at the MPI Informatik in Saarbruecken and

received his PhD from Saarland University in 2003. He spent two years (2005–2006) as a visiting assistant professor at Stanford University, USA. His research interests include 3D appearance acquisition, image-based and computational photography rendering. He has given several lectures and tutorials about this topic at various conferences including SIGGRAPH courses on realistic materials in 2002 and 2005.

Szymon Rusinkiewicz is an associate professor of Computer Science at Princeton University. His work focuses on acquisition and analysis of the 3D shape and appearance of real-world objects, including the design of capture devices and data structures for efficient representation. He also investigates algorithms for processing complex datasets of shape and reflectance, including registration, matching, completion, symmetry analysis, and sampling. In addition to data acquisition, his research interests include real-time rendering and perceptually-guided depiction. He obtained his Ph.D. from Stanford University in 2001.

Tim Weyrich is a Post-doctoral Teaching Fellow at Princeton University, working in the Computer Graphics Group Princeton. His research interests are appearance modeling, 3D reconstruction, cultural heritage acquisition, and point-based graphics. Prior to coming to Princeton in Fall 2006, he received his PhD from ETH Zurich, Switzerland, where he developed a novel face scanner to analyze human skin reflectance, allowing for photo-realistic reconstructions of human faces. He received his Diploma degree in computer science from the University of Karlsruhe (TU), Germany, in 2001.

Todd Zickler received his Ph.D. in electrical engineering from Yale University in 2004 and is currently an assistant professor in the School of Engineering and Applied Sciences at Harvard University. His research spans computer vision, computer graphics and image processing, and he is currently focused on developing representations of appearance and exploiting them for visual inference. In 2006, he was the recipient of a career award from the US NSF titled, “Foundations for Ubiquitous Image-based Appearance Capture.”

Contents

1 Radiometry and Appearance Models	1
1.1 Radiometry	2
1.2 Surface Reflectance	4
1.3 Subsurface Scattering	7
1.4 Generalizing Reflectance and Scattering	9
2 Principles of Acquisition	11
2.1 Homogeneous Reflectance: BRDF	11
2.1.1 The Gonioreflectometer	12
2.1.2 Image-based measurement of planar samples	12
2.1.3 Image-based measurement of curved samples	13
2.1.4 Image-based measurement of arbitrary shapes	13
2.2 Spatially-varying Reflectance: SVBRDF	14
2.2.1 Planar Surfaces: The Spatial Gonioreflectometer	14
2.2.2 Curved Surfaces	15
2.2.3 Separability: The Dichromatic Model	17
2.2.4 Case Study: Reflectance Sharing	17
2.3 Subsurface scattering: BSSRDF	18
2.4 Calibration	19
2.4.1 Geometric calibration	19
2.4.2 Radiometric calibration	20
2.4.3 Colorimetric calibration	20
2.4.4 Shape	21
3 Spatially-Varying Reflectance Models	23
3.1 Acquisition	24
3.2 Representation	24
3.2.1 Basis Decomposition	24
3.2.2 Parametric Models	25
3.3 The Inverse Shade Tree Framework	26
3.3.1 Alternating Constrained Least Squares	28
3.4 Conclusion and Directions of Future Research	29
3.4.1 Parametric vs. Non-Parametric	30
3.4.2 Open Problems	30
4 From BSSRDF to 8D Reflectance Fields	33
4.1 BTFs and Distant Light Reflectance Fields	33
4.2 BSSRDFs	34
4.3 Diffuse Subsurface Scattering	34
4.4 Arbitrary Light Transport	35

4.4.1	Single View – Single Projector	36
4.4.2	8D Reflectance Fields	37
4.5	Conclusions and Future Work	39
4.5.1	Open Problems	39
5	The Human Face Scanner Project	41
5.1	Previous Work	41
5.2	Skin Appearance Acquisition	42
5.2.1	Subsurface Scattering Acquisition	43
5.2.2	Reflectance Field Acquisition	43
5.3	Face Data Processing	44
5.3.1	System Calibration	44
5.3.2	Geometry Refinement	47
5.4	Reflectance Model Fit	47
5.5	Reflectance Analysis	49
5.6	Appearance Transfer	51
5.7	Conclusion	52
	List of Figures	53
	Bibliography	55

1 Radiometry and Appearance Models

Szymon Rusinkiewicz, Princeton University

Comprehending the visual world around us requires us to understand the role of *materials*. In essence, we think of the appearance of a material as being a function of how that material interacts with light. The material may simply reflect light, or it may exhibit more complex phenomena such as subsurface scattering.




Figure 1.1: *Reflectance (left) and subsurface scattering (right).*

Reflectance itself is a complex phenomenon. In general, a surface may reflect a different amount of light at each position, and for each possible direction of incident and exitant light (Figure 1.1, left). So, to completely characterize a surface's reflection we need a six-dimensional function giving the amount of light reflected for each combination of these variables (position and incident and exitant directions are two dimensions each). Note that this does not even consider such effects as time or wavelength dependence. We will consider those later, but for now let us simply ignore all time dependence, and assume that any wavelength dependence is aggregated into three color channels: red, green, and blue.

These reflectance functions embody a significant amount of information. They can tell us whether a surface is shiny or matte, metallic or dielectric, smooth or rough. Knowing the reflectance function for a surface allows us to make complete predictions of how that surface appears under any possible lighting.

For a translucent surface, the interaction with light can no longer be described as simple reflection. This is because light leaves the surface at a different point than where it entered (Figure 1.1, right). So, in order to characterize such surfaces we need a function that gives the amount of light that is scattered from each possible position (2D) to each other position

(another 2D). If we wanted to be even more correct, of course, we would need to account for the directional dependency as well.

So, now that we have some idea of *how* we can understand appearance, there remains the question of *why* we may wish to do so. In addition to the obvious application domain of image synthesis, having a complete knowledge of a materials appearance can help us interpret images. It will aid in 3D reconstruction, view interpolation, and object recognition. Furthermore, knowing how to characterize materials can help us understand how humans perceive surfaces.

This SIGGRAPH class covers the basic principles of how materials are described, how the appearance of real-world objects may be measured, and how a knowledge of appearance aids in a variety of applications. This first section of the class covers foundational topics. We will learn about *radiometry* and see the definition of the BRDF: a function describing surface reflectance at a point. We will cover generalizations of the BRDF, including spatial variation and subsurface scattering. Finally, we will consider the many different types of data that can be captured that characterize “appearance,” and how they relate to each other.

1.1 Radiometry

Let us start with the basics. Light is a form of electromagnetic energy, and so can be measured in Joules. Because it is most useful to think of continuous light flow, instead of individual pulses, we will most often be interested in the amount of energy flowing per unit time. This is known as “radiant flux” (Φ) or just “power,” and hence may be measured using the SI units of Watts.

Although having a way of characterizing the total flow of light power is useful, we will need to consider more complex quantities in order to talk about concepts such as light sources and surface reflectance.

Point Light in a Direction: Consider an ideal light source (idealized as a point in space). If the light were being emitted uniformly in all directions, describing its power (in Watts) would tell us all we wanted to know. However, it is possible that light is not being emitted equally in all directions. In this case, characterizing the power being emitted in a *particular* direction requires a different unit. In such cases, we can talk about the amount of power being emitted per unit *solid angle*.

So what exactly is a solid angle, and how is it measured? A useful analogy is to the way an angle is defined in the plane. One radian is defined as the angle subtended by an arc of a circle, with the arc length being equal to the circles radius. Equivalently, an angle in radians may be calculated by dividing the length of a circular arc by the radius.

Moving to the concept of solid angles, we will be working in three dimensions (vs. two for angles), and will be looking at a sphere (vs. a circle). The basic unit of solid angle is known as the *steradian*, and is defined as the area of some region on a sphere divided by the square of the spheres radius. A complete sphere thus has 4π steradians, and smaller solid angles define smaller regions of the space of directions.




Figure 1.2: Point light source emitting light in a direction.

So, measuring the directional power of a point light source can be done using the units of Watts per steradian. The same amount of power emitted into a smaller solid angle will result in a larger measurement (think of the case of a laser, which has a relatively small power but is concentrated into a very small solid angle).

Light Falling on a Surface: Another radiometric quantity we often wish to measure is called *irradiance*. It represents the amount of light falling onto a surface. Because the same radiant flux will be “more concentrated” when falling onto a smaller area of surface than a larger surface, we define irradiance as power per unit area.

Given this definition of irradiance, there are two immediate and easily-observed “laws” that emerge. The first is the inverse-square law: moving a point light source away from a surface reduces irradiance in proportion to the inverse square of the distance. Secondly, tilting a surface away from a point light results in a lower irradiance, in proportion to the cosine of the angle between the surface normal and the direction towards the light. This “cosine law” is often written as the dot product between the (unit-length) surface normal and light vectors.

Light Emitted from a Surface in a Direction: We now come to the final, and most complex, radiometric quantity we are going to consider, which describes the emission of light from a surface. This can be thought of as combining the two concepts we just saw: the emitted light can vary with direction (hence we must control for its directional distribution, as we did with the point-light case), and we are interested in the amount of light emitted per unit surface area. Hence, we arrive at the definition of radiance: power emitted per unit area (perpendicular to the viewing direction) per unit solid angle.

Radiance is perhaps the most fundamental unit in computer vision and graphics. It is easy to show that the irradiance on a camera sensor is proportional to the radiance of the surfaces it is imaging, so cameras “see” the radiance of surfaces. The pixel values we deal with in digital images are (ignoring nasty things like gamma) proportional to radiance!




Figure 1.3: Light emitted from a surface, in a specific direction.

1.2 Surface Reflectance

We are now ready to use what we know about radiometry to define the BRDF [NRH⁺77]. This is the Bidirectional Reflectance Distribution Function, and it describes surface reflection at a point. Formally, it is the ratio between the reflected radiance of a surface and the irradiance that caused that reflection. The radiance and irradiance are each measured at a particular angle of exitant and incident light, respectively, so the BRDF is usually written as a function of four variables: the polar angles of light coming into and out of the surface.

$$f_r(\omega_i \rightarrow \omega_o) = f_r(\theta_i, \varphi_i, \theta_o, \varphi_o) = \frac{dL_o(\omega_o)}{dE_i(\omega_i)} \quad (1.1)$$

The BRDF is often written as a differential quantity. This is to emphasize that there is no such thing as light arriving from *exactly* one direction, and being reflected into exactly one outgoing direction. Rather, we must look at non-zero ranges of incident and exitant solid angles, and consider the limit as those approach zero.

Because BRDFs are 4D functions, they are a bit tricky to visualize directly. Instead, we often visualize two-dimensional slices of this function. Figure 1.4 shows one slice of a BRDF, corresponding to one direction of incidence (the arrow) and all possible directions of reflection. The blue surface is a hemisphere stretched such that its radius in any direction is the reflected radiance in that direction, and is known as a goniometric plot.




Figure 1.4: Goniometric view of slices of a BRDF corresponding to two incident directions.

You will note that, for this particular BRDF, much of the incident light is reflected equally in all directions. This is the constant-radius (spherical) portion of the surface you see. However, there is also a bump in the surface, indicating that there is a concentrated reflection in one particular direction.

If we change the direction of incidence, we see that the constant portion of the function was unchanged, but the position of the bump moved. In fact, the bump always appears around the direction of “ideal mirror reflection” of the incident direction. This is known as a specular highlight, and it gives a surface a shiny appearance.

Properties of the BRDF: Before we look at specific BRDF models, let us discuss a few properties shared by all BRDF functions. The first is energy conservation: it is impossible for a surface to reflect more light than was incident on it! Expressing this mathematically, we see that the integral of the BRDF over all outgoing directions, scaled by a cosine term to account for foreshortening, must be less than one:

$$\int_{\Omega} f_r \cos \theta_o d\omega_o \leq 1. \quad (1.2)$$

A second, more subtle, property of BRDFs is that they must be unchanged when the angles of incidence and exitance are swapped:

$$f_r(\omega_i \rightarrow \omega_o) = f_r(\omega_o \rightarrow \omega_i). \quad (1.3)$$

This is a condition known as Helmholtz reciprocity, and is due to the symmetry of light transport. Some systems, such as Todd Zickler’s work on Helmholtz stereopsis, have relied on this property.

Some, but not all, BRDFs have a property called isotropy: they are unchanged if the incoming and outgoing vectors are rotated by the same amount around the surface normal. In this case, there is a useful simplification that may be made: the BRDF is really a 3-dimensional function in this case, and depends only on the difference between the azimuthal angles of incidence and exitance.

The inverse of isotropy is anisotropy. An anisotropic BRDF does not remain constant when the incoming and outgoing angles are rotated. In this case, a full four-dimensional function is necessary to characterize the behavior of the surface. Anisotropic materials are frequently encountered when the surface has a strongly directional structure at the small scale: brushed metals are one example (Figure 1.5).

Lambertian BRDF: We now turn to looking at specific examples of BRDFs. We will look at simple examples, such that the reflectance may be written as a mathematical formula. Real surfaces, of course, are more complex than this, and mathematical models frequently do not predict the reflectance with great accuracy.

The simplest possible BRDF is just a constant.

$$f_r = \text{const.} = \rho / \pi \quad (1.4)$$

This results in a matte or diffuse appearance, and is known as ideal Lambertian reflectance. This BRDF is frequently written as a constant ρ divided by π . In this case, ρ




Figure 1.5: *Anisotropic reflection.*

is interpreted as the diffuse albedo: it is the fraction of light that is reflected (vs. absorbed) by the surface. Plugging this BRDF into the energy conservation integral verifies that the surface conserves energy precisely when the albedo is less than or equal to one.

Blinn-Phong BRDF: Another simple analytic BRDF is the Blinn-Phong model, designed to represent glossy materials:

$$f_r = \rho / \pi + k_s (\mathbf{n} \cdot \mathbf{h})^\alpha. \quad (1.5)$$

In contrast to the Lambertian BRDF, the distribution of reflected light is not constant. In fact there is a lobe centered around the direction of ideal mirror reflection for each incident angle, containing significantly more energy than the rest of the domain. This is known as the specular lobe, and its size and fall-off are controlled by the parameters k_s and α . This lobe is what produces the specular highlights on this vase that help give it a shiny appearance.

Torrance-Sparrow BRDF: A more complex, yet more realistic, BRDF was originally developed in the physics community by Torrance and Sparrow [TS67], and was refined for computer graphics by Cook and Torrance [CT82].

$$f_r = \frac{D G F}{\pi \cos \theta_i \cos \theta_o} \quad (1.6)$$

For the purposes of calculating reflectance, this model assumes that at a very small scale the surface consists of tiny, mirror-reflective “microfacets” oriented in random directions. There are three major terms in the model that describe the angular distribution of microfacets, how many are visible from each angle, and how light reflects from each facet.

$$D = \frac{e^{-\left(\frac{\tan \beta}{m}\right)^2}}{4 m^2 \cos^4 \beta} \quad (1.7)$$

The first term D in the Torrance-Sparrow model describes what is the density of facets facing in any possible direction. Notice that part of this term resembles a Gaussian.

$$G = \min \left\{ 1, \frac{2(\mathbf{n} \cdot \mathbf{h})(\mathbf{n} \cdot \mathbf{v})}{(\mathbf{v} \cdot \mathbf{h})}, \frac{2(\mathbf{n} \cdot \mathbf{h})(\mathbf{n} \cdot \mathbf{l})}{(\mathbf{v} \cdot \mathbf{l})} \right\} \quad (1.8)$$

The next term G in the Torrance-Sparrow model accounts for the fact that not all facets are visible from all directions, because they are hidden by the facets in front of them.

Finally, the reflection from each facet is described by the Fresnel term F , which predicts that reflection increases towards grazing angles.

Other BRDF Features: Another commonly-observed characteristics of BRDFs is an increase in light reflected into all grazing angles, as is typical for “dusty” surfaces. Finally, some BRDFs include a retro-reflective component. That is, they scatter light most strongly back into the direction from which it arrived. The paint found on roads and street signs is a common example of this phenomenon. Such paint contains crystals that produce a “corner reflector” configuration

Beyond BRDFs: Although we could continue to develop mathematical BRDF formulas of increasing sophistication that explain a greater and greater variety of optical phenomena, over the past decade it has become increasingly practical to simply measure the BRDFs of real material samples [MPBM03b]. In fact, it is one of our main arguments in this class that measured data can capture a greater variety of real-world optical phenomena with greater accuracy than is possible with analytic models.

Of course, the BRDF is merely the beginning of our study of the appearance of materials. Real-world objects will exhibit more complex behaviors, such as a BRDF that changes from point to point on the surface. Adding two spatial dimensions to the four directional dimensions of the BRDF leads us to the six-dimensional Spatially-Varying BRDF. Later in this class you will hear about the challenges of capturing, representing, editing, and analyzing these complex functions [LBAD⁺06a].

1.3 Subsurface Scattering

Even the SVBRDF is not enough to characterize all materials. Many surfaces exhibit translucency: a phenomenon in which light enters the object, is reflected inside the material, and eventually re-emerges from a different point on the surface. Such sub-surface scattering can have a dramatic effect on appearance, as can be seen from these computer graphics simulations that differ in only one respect: the left image simulates surface reflection only, while the right image includes sub-surface scattering [PvBM⁺06a].

In order to cope with subsurface scattering, we will need to examine more complex appearance functions: those that can include the phenomenon of light leaving the surface at a different point than the one at which it entered.

The BSSRDF: The relevant function is known as the Bidirectional Scattering-Surface Reflection Distribution Function, or BSSRDF:

$$S(x_i, y_i, \theta_i, \varphi_i, x_o, y_o, \theta_o, \varphi_o). \quad (1.9)$$



Figure 1.6: Left: a synthesized image with surface reflection only. Right: the same model with a simulation of subsurface scattering.

You will notice that we have taken the SVBRDF and added two more variables, representing the surface location at which the light leaves the surface. We are now up to a function of 8 variables!

As we will see later in the class, the high dimensionality of this function leads to great difficulty in capturing and working with the BSSRDF directly, especially if a high sampling rate in each dimension is desired [GLL⁺04b].

Because of the enormous size of the BSSRDF, approximations to it have become quite popular. One of the most powerful approximations relies on the fact that, in many cases, the appearance is dominated by light that has reflected many times within the material. In this case, the details of each scattering event become unimportant, and the appearance is well approximated by thinking of light “diffusing” away from the location at which it enters the surface, much as heat might spread [JMLH01b].

It turns out that the pattern of diffusion is well approximated by a dipole: a combination of a point light some distance *below* the point at which light entered the surface, and a *negative* light source some (slightly larger) distance above the surface. Combining the contributions of these two light sources with Fresnel terms for light entering and leaving the surface yields a simple, yet powerful, model:

$$S = F(\theta_i) R(\|x_i - x_o\|) F'(\theta_o). \quad (1.10)$$

Because of the symmetry of diffusion, the model is effectively a function of only one variable: the distance between the points of incidence and exitance.

This dipole model, originally introduced in 2001, has become very popular for simulating subsurface scattering in many materials, and we will see applications of it later in this class.

Homogeneous and Heterogeneous Scattering: Of course, the dipole approximation assumes a uniform material: the same amount of scattering everywhere on the surface. For more realistic surfaces, you might need to add some of the complexity of the BSSRDF back in, by considering spatial variation. For example, in Figure 1.6 you can clearly see how internal structure affects the scattering.

1.4 Generalizing Reflectance and Scattering

So, does the BSSRDF cover all possible aspects of surface appearance? No!

First, we could consider all of the functions we have talked about as being dependent on the wavelength of light. Moreover, some surfaces are *fluorescent*: they emit light at different wavelengths than those present in the incident light.

Some other surfaces may have appearance that changes over time because of chemical changes, physical processes such as drying, or weathering. Other surfaces might capture light and re-emit it later, leading to phosphorescence and other such phenomena.

Thus, a complete description of light scattering at a surface needs to add at least two wavelength and two time dimensions to the BSSRDF. Scattering from a region of space would add two additional spatial dimensions.

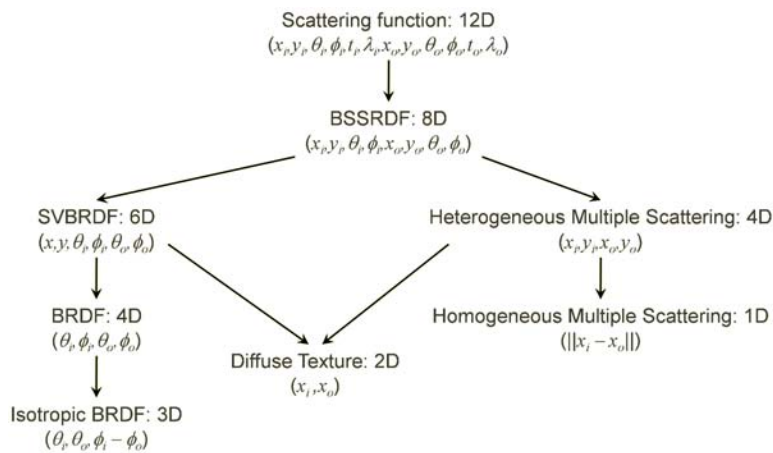


Figure 1.7: *Taxonomy of scattering and reflectance functions.*

So, we can think of all of the functions we've seen as specializations of a 12-dimensional scattering function. While nobody has really tried to capture the full function, many efforts exist to capture one or more of its low-dimensional subsets. In fact, it can be argued that over the past decade, researchers have explored most of the subsets that “make sense,” up to the limits of acquisition devices.

2 Principles of Acquisition

Todd Zickler, Harvard University

In this chapter we will consider the measurement of surface reflection properties in order of increasing complexity, from homogeneous BRDF (functions of at most five dimensions) to general subsurface scattering (a function of at most nine dimensions).

In designing an acquisition system, the four (competing) factors that need to be considered are acquisition time, accuracy and precision, cost, and generality. Here, generality refers to the breadth of materials that are to be considered. It is possible to build an efficient system for measuring the BRDF of spherical surfaces, for example, but not every material can be ‘painted on’ to a sphere.

2.1 Homogeneous Reflectance: BRDF

The BRDF is the simplest reflection model we will examine. We consider it to be a function of at most five dimensions, one of which is the spectral dimension. An isotropic BRDF has an angular domain whose dimension is reduced by one.

One generally measures a BRDF by illuminating a (locally) flat surface with a collimated beam of light with direction (ω_i) and measures the spectral radiance reflected in an output direction (ω_o). The input and output directions are assumed to be known relative to the local coordinate system of the planar surface patch.

Since the BRDF is a derivative quantity, we can only measure its average over finite spatial and angular intervals. Indeed, “truly infinitesimal elements of solid angle do not contain measurable amounts of radiant flux.”[NRH⁺77] This is not usually a problem in the angular sense. The solid angle subtended by a typical sensor is small enough that the BRDF can be considered constant within it. One must be more careful in the spatial sense. It is essential that the spatial scale of the measurements be chosen such that the assumptions underlying the BRDF (i.e., radiant flux emitted from a point is due only to flux incident at that point) are satisfied. In image-based BRDF measurement systems, where high resolution cameras are used to measure reflected flux, this generally means that images must be downsampled to obtain valid BRDF measurements.

As an example, consider the measurement geometry shown in the simple schematic on the slides. A portion of a planar sample is observed by a sensor through an optical system (e.g., by a single element of a CCD array.) The finite area of the sensor back-projects to a finite area on the surface A_o . In order for this system to provide an accurate measurement of the BRDF, we require that both the illumination and the surface scattering effects be

uniform across the surface over a larger area $A_i \supset A_o$. Also, A_i must be large enough to guarantee that flux incident outside of A_i would not contribute significantly to the radiance reflected within A_o . Detailed guidelines for BRDF measurement can be found in [NRH⁺77, Sta].

2.1.1 The Gonioreflectometer

A classic device for measuring a general, anisotropic BRDF is the four-axis gonioreflectometer. In this device, a combination of servo motors are used to position a source and detector at various locations on the hemisphere above a planar material sample. The detector is typically linked to a spectroradiometer or another optical assembly that permits recording of dense spectral measurements for each configuration of the source and detector [WSB⁺98].

This measurement process is a lengthy one, and it can require days to measure a single material. The advantage of this approach, however, is that the system can be carefully calibrated and measurements can be quite repeatable. Also, the ability to capture dense spectral information provides a tremendous advantage over the camera-based systems that are ubiquitous in the vision and graphics communities.

Acquisition time and equipment cost can be reduced if one is willing to restrict one's attention to isotropic BRDFs. In this case, the function being measured has only three angular dimensions, so one requires only three degrees of freedom in the acquisition system. In the gonioreflectometer recently built at Cornell [LFTW06], this is accomplished by having two degrees of freedom in the orientation of the planar sample and one additional degree of freedom in the angular position of the source. Using this system, one can acquire 31 spectral samples per camera/source position (roughly 10nm increments over the visible spectrum), and capturing 1000 angular samples (which is a very sparse covering of the 3D angular domain) takes approximately 10 hours.

2.1.2 Image-based measurement of planar samples

BRDF acquisition can be made less costly and time-consuming by replacing a simple photodetector by a camera. A camera's sensor contains millions of photo-sensitive elements, and by using lenses and mirrors, these elements can be used to collect a large number of reflectance samples simultaneously.

An early example of this is Ward's measurement system [War92], in which the radiance emitted by a planar sample is reflected from a half-silvered hemisphere and captured by a camera with a fish-eye lens. In this way, almost the entire output hemisphere is captured by a single image. The two degrees of freedom in the incident direction are controlled by a rotation of the source arm (about point C in the figure) and the rotation of the planar sample. A very nice property of this system is that it allows the measurement of retro-reflection directions, meaning those for which the incident and reflected directions are equal. This is not possible with either of the two gonioreflectometers described earlier. Using this system, Ward claimed that a 4D anisotropic BRDF could be measured in ten minutes.

What are we trading for this gain in efficiency? Spectral resolution, for one. If we use a white light source (uniform spectral distribution) and an RGB camera, we obtain only three spectral measurements for each angular configuration, and these measurements are weighted averages over large, overlapping intervals of the visible spectrum. This is a serious limitation if we want to be able to predict the appearance of the material under a source with a different spectral distribution. Without dense spectral information, physically-accurate color reproduction is generally unattainable. Another limitation is the complexity of the required calibration procedure. In Ward's system, we need to know the map from pixels in the camera to output directions in the coordinate system of the sample. If we want to be precise, we also need to know the exitant solid angle that is effectively subtended by each pixel. In addition to this geometric calibration information, we need radiometric information including the optical fall-off in the lens system and the radiometric camera response. The complexity of this process reduces the accuracy and repeatability of the measurements. This is an example of a design decision in which one trades precision for a decrease in acquisition time.

2.1.3 Image-based measurement of curved samples

An alternative approach is to eliminate the hemispherical reflector and to use curved material samples instead of a planar one [MWLT00, LKK98]. For isotropic materials, one can use a sphere. The surface normal varies from point to point and so does the local input and output directions. This means that each image provides a very dense (near continuous) slice of samples embedded in the 3D isotropic BRDF domain. Matusik et al. [MPBM03b] captured an extensive database of isotropic BRDFs in this way.

Using cylinders instead of spheres, one can do this for anisotropic materials as well. For example, one can cut strips of an anisotropic material at different orientations relative to the material's tangent direction and paste these strips onto a cylinder [NDM05]. The cylinder provides one degree of freedom in its surface normal, and two more degrees of freedom are obtained by rotating the cylinder and the source. The fourth and final degree of freedom comes from the number of 'strips', which is typically coarsely sampled.

2.1.4 Image-based measurement of arbitrary shapes

All methods discussed so far are limited to materials that exist as planar samples or that can be 'painted on' onto a known shape such as a sphere or a cylinder. What about materials for which this is not possible? Well, we can measure the material properties on any surface as long as the surface shape is known. So if capture the shape using a reconstruction system (laser scanner, structured-lighting, photometric stereo, etc.) and then carefully align this shape with the acquired images [MWL⁺99], we can use the very same procedure outlined above. This obviously introduces additional sources of error and bias, so here we are trading precision and accuracy for increased generality.

With all of these sources of error, BRDF measurements from captured, arbitrary shapes is often prohibitively noisy. One direction for future work is the design of systems that recover both shape and reflectance from the same image data. More on this later.

2.2 Spatially-varying Reflectance: SVBRDF

Next we allow spatial variation in the reflectance function, which increases the dimension by two. Note that despite allowing spatial variation, we maintain our assumption regarding spatial scale and sub-surface scattering. Namely, we assume that the surface area observed by each photo detector is large enough that sub-surface scattering effects are negligible and that the surface is *locally* homogeneous. This guarantees that the appearance of each small surface element can be represented by a BRDF.

Since cameras are used for the measurement of SVBRDF, essentially all acquisition systems to date have considered only sparse spectral sampling (RGB). For this reason, we can effectively ignore the spectral dimension in this section. Even with this simplification, we are left with the formidable task of measuring a function of five or six dimensions.

2.2.1 Planar Surfaces: The Spatial Gonioreflectometer

Acquisition of a spatially-varying BRDF can be thought of as the measurement of multiple, distinct BRDF—one at each point on a surface. SVBRDF acquisition systems are therefore closely related to the BRDF acquisitions systems just discussed.

A spatial gonioreflectometer is a good example and is exactly what the name suggests. It functions like a standard gonioreflectometer, except that the single photodetector is replaced by a camera. The example shown in the slides was built by David McAllister [McA02b] and is similar in spirit to an earlier version by Kristen Dana [DvGNK99]. In the example shown, the planar sample has spatially-varying reflectance, and its orientation is controlled by a two-axis pan/tilt head. The source can be rotated about the sample as well, which produces a three-axis spatial gonioreflectometer. Assuming a columnated source (and an orthographic camera for simplicity), each image yields a dense 2D spatial slice of the SVBRDF corresponding to fixed input and output directions. Of course, a three-axis device such as this one is useful for materials having an isotropic BRDF in each local region.

BRDF samples collected by this device are very non-uniformly distributed in the 5D domain. There is near continuous sampling of the spatial dimensions but only as many samples of the angular dimensions as there are positions of the sample and illuminant. This can be changed by modifying the acquisition system. As was the case for single BRDF, lenses and mirrors can be used to redirect input and output rays. Here the motive is not to decrease the acquisition time, however, but simply to alter the sampling pattern (and perhaps to improve precision).

One such system uses a beam splitter and a parabolic mirror to increase angular sampling rates at the expense of spatial sampling rates [DW04]. The parabolic mirror reflects a portion of the output hemisphere from a single small region toward the camera, thereby providing near-continuous sampling of this angular interval. The same mirror is used to direct an incident collimated beam of light toward the surface patch being observed. The direction of the incident ray is controlled by translating an aperture in front of the light source, and the surface sampling point (x, y) is changed by translating the parabolic

mirror. While the altered sampling pattern is interesting, the great strength of this system is that all required movements are pure translations. One expects this to be highly repeatable.

Before we continue, I would like to pause to clarify some terminology. In the literature, one often sees the term *Bi-directional Texture Function* (BTF), which, like the SVBRDF, is a function of two spatial and four angular dimensions. For the purpose of this session, BTF and SVBRDF are not different. Unlike a SVBRDF, a BTF $I(x, \omega_i, \omega_o)$ incorporate non-local effects such as cast shadows, occlusions, mutual illumination and sub-surface scattering that are highly dependent on the *non-local* shape of the surface. This means, for example, that the appearance at a single surface point cannot be well-represented by a parametric BRDF model. We ignore these non-local effects here and restrict our attention to SVBRDF. Arbitrary, non-local scattering will be discussed in a later section.

2.2.2 Curved Surfaces

Most interesting spatially-varying surfaces are not planar nor can they be ‘painted on’ to a planar substrate. Thus, there is often the need to measure the SVBRDF on a curved surface directly, and as for BRDF, this can be done when the 3D geometry of the surface is known. An example of a suitable acquisition system is the Stanford Spherical Gantry. This system can sample all six dimensions of the SVBRDF defined on a curved surface. In this case, one needs to sample the entire *sphere* of directions and not just a hemisphere.

To get a sense of how much data is required to densely sample the SVBRDF of a regular surface, we can perform a simple counting exercise. When the shape of the object is known and the source and view directions are given, each pixel in an image provides one sample of the SVBRDF (or the BRDF at a particular surface point). Sampling the BRDF at every surface point in 5° or 1° angular increments therefore requires millions or hundreds-of-millions of images, respectively.

Clearly, capturing millions of images per object is impractical, so we look to reduce this burden using one or more of:

- improved acquisition systems
- parametric BRDF models; and
- knowledge of general reflectance phenomena

Each is discussed separately below.

Acquisition Systems A number of acquisition systems have been developed over the past ten years. Thanks to rapid advancements in LEDs and digital cameras, it is becoming less and less expensive to build devices capable of capturing large numbers of images very quickly. In designing these systems there are a number of trade-offs. For example, it is much easier to calibrate a system like the Stanford spherical gantry or Light Probe One at USC that uses only one camera and one light source. But acquisition time for these systems will generally be much larger than systems equipped with multiple sources and cameras.

Parametric Approaches The use of parametric models generally reduces the number of required input images and hence the required acquisition time. In the general case, one must densely sample the 4D (or 3D isotropic) BRDF at each surface point. When we use a parametric model like the Phong model, however, we need only estimate a handful of parameters at each point.

Of course, this approach assumes that the appearance of the object we wish to acquire can be accurately represented by our chosen BRDF model, which may very well not be the case. For example, a Cook-Torrance model does well at representing plastics and metals, but it cannot represent retro-reflection effects. Nonetheless, if one is interesting in acquiring a particular class of surfaces that can be well-represented by a particular model, a parametric approach might be appropriate. Indeed, as we will see with the human face project in the last section of this document, when used in conjunction with other reflectance representations, parametric model-based representations can produce stunning results.

There have been a number of parametric approaches presented over the past decade (e.g., [SWI97, YDMH99, BG01, LKG⁺01, Geo03, GCHS03, MLH02]). These methods vary widely in terms of the model (or models) used, the type of input data, and the procedure used for fitting. Some of these methods are described in more detail in Sect. 4.

General Reflectance Properties Instead of relying on parametric models, or in addition to using them, we can reduce the number of required input images by exploiting what we know about reflectance in general.

We have already discussed *isotropy* and *reciprocity* as being common, and useful, properties of reflectance. These are important from an acquisition standpoint because they substantially reduce the angular domain and thus the number of required images.

We have also already discussed *compressibility*. This refers to the fact that even though the BRDF can change rapidly in some regions of its angular domain, it often changes slowly over much of it. If one represents a BRDF in a wavelet basis, for example, it is likely to be very sparse [MPBM03a]. By using appropriate sampling schemes, compressibility can be exploited for acquisition as well.

Another well-known and well-used property is *separability*. Separability refers to the fact that spatially-varying reflectance can often be written as a linear component of diffuse and specular components. This is useful because when isolated, each of these components exhibits different (and exploitable) behavior. For example, the diffuse component is typically well-represented by a Lambertian model and can often be reliably estimated over the entire surface from just a handful images.

Spatial smoothness refers to the fact that for many surfaces, reflectance is slowly varying from point to point. This is especially true for the specular reflectance component (e.g. [SWI97, ZREB06]). Thus, knowledge of the reflectance at one point on a surface can say quite a bit about the reflectance at another.

Finally, *spatial regularity* is another way of describing the correlation between the reflectance at distinct surface points on the same surface. Here, it is assumed that the reflectance at all surface points can be written as linear combinations of a single set of basis BRDF (e.g., [LKG⁺01, GCHS05b]).

Before we see an example of an approach that exploits these properties for acquisition, we will take a moment to describe separability in more detail.

2.2.3 Separability: The Dichromatic Model

The dichromatic model of reflectance is a common special case of the BRDF model, and it was introduced to the computer vision community by Shafer [Sha85] as a model for the reflectance of dielectrics. It assumes that the BRDF of the surface can be decomposed into two additive components: the interface (specular) reflectance and the body (diffuse) reflectance. Furthermore, it assumes that each of these two components can be factored into a univariate function of wavelength and a multivariate function that depends on the imaging geometry. That is,

$$f(\lambda, \omega_i, \omega_o) = g_d(\lambda)f_d(\omega_i, \omega_o) + g_s(\lambda)\tilde{f}_s(\omega_i, \omega_o).$$

If we further assume that the index of refraction on the surface is constant over the visible spectrum—a valid assumption for many materials—it follows that $g_s(\lambda)$ is a constant function [LBS90]. This leads to the common expression for the BRDF of a separable (or dichromatic) surface,

$$f(\lambda, \omega_i, \omega_o) = g_d(\lambda)f_d(\omega_i, \omega_o) + f_s(\omega_i, \omega_o), \quad (2.1)$$

where $f_s(\omega_i, \omega_o) = g_s\tilde{f}_s(\omega_i, \omega_o)$. The function $g_d(\lambda)$ is often referred to as the *spectral reflectance* of the material.

Even though it was originally used to describe the reflectance of dielectrics [Sha85], the dichromatic model has been used successfully as an approximation of the reflectance of many different materials. Empirically it has shown to be suitable for certain types of plant leaves, cloth, wood, and the skin of fruits [LBS90, TW89] in addition to a large number of dielectrics [Hea89].

Separability is useful from an acquisition standpoint because the diffuse and specular components tend to exhibit different spatio-angular characteristics that can each be exploited. If we now consider spatial variation, the diffuse component is nearly Lambertian but typically varies rapidly over the surface:

$$g_d(\lambda, \mathbf{x})f_d(\mathbf{x}, \omega_i) \approx g_d(\lambda, \mathbf{x})f_d(\mathbf{x}) \quad (2.2)$$

We often refer to such surfaces as having significant *texture*. The specular component, on the other hand, is typically non-Lambertian, but changes slowly from point to point.

2.2.4 Case Study: Reflectance Sharing

The properties described in the previous section have been exploited in various ways by both parametric and non-parametric techniques for acquisition. Rather than list them all here, we will describe one example that exploits many of them. We refer to the method as *reflectance sharing* [ZREB06], and it is a non-parametric approach that seeks to acquire the SVBRDF of a surface from a very sparse set of images.

The input is a set of images of a known three-dimensional shape, with each image being captured under a columnated illumination in a known direction. As described earlier, each pixel in one of these images provides one sample of the BRDF at the corresponding surface point. The input images are assumed to be decomposed into their specular and diffuse components. In practice this is often done using polarizing filters on the camera and light source as follows. Two exposures are captured for each view/lighting configuration, one with the polarizers aligned (to observe the sum of specular and diffuse components), and one with the source polarizer rotated by 90° (to observe the diffuse component only.) The specular component is then given by the difference between these two exposures. (See, e.g., [Mer84].)

If one assumes the diffuse component to be Lambertian it can be measured from as little as one image. We are simply required to estimate an RGB diffuse texture map $a_{\text{RGB}}(x, y)$. Estimating the specular component of the SVBRDF is more difficult, and this where isotropy, reciprocity, compressibility and spatial smoothness play an important role. Isotropy, reciprocity and compressibility are exploited by representing the angular dimensions of the SVBRDF in terms of Szymon Rusinkiewicz's halfway/difference parameterization [Rus98]. This is a natural way to 'shrink' the angular domain and to separate angular intervals that typically exhibit rapid variation from those that do not.

To exploit spatial smoothness, we view each pixel as sample lying in the (5D isotropic) SVBRDF domain, and we note that each image provides a near-continuous sampling of a 2D slice in this domain. SVBRDF estimation is formulated as a scatter-data interpolation problem, in which we simultaneously interpolate the samples in both the angular and spatial dimensions.

Using this approach, the sampling of the SVBRDF domain is highly non-uniform. We obtain only a sparse set of 2D slices (one per image), while the sampling along each 2D sheet is very dense. Fortunately, one can show that the densely-sampled regions of the SVBRDF correspond nicely with the dimensions in which we typically observe rapid angular variation, such as the half-angle dimensions near small half-angle values. For this reason, one can often recover accurate SVBRDF from a very small number of images. The example in the slide shows the specular lobes at two points on the surface in the case where only four input images are given. Plausible results are obtained even though at most four reflectance samples are available at each surface point. Once the SVBRDF is recovered, one can use it to predict the appearance of the surface in novel view and lighting conditions.

2.3 Subsurface scattering: BSSRDF

A general BSSRDF is a function of nine dimensions if we include the spectral dimension. Even if we ignore the dependency on wavelength, densely sampling an eight dimensional space is an extremely burdensome process. We will see some acquisition systems that begin to address this in the next section. Here we focus on a common simplification that relies on a factored form of the BSSRDF.

As mentioned earlier, the most common simplification of the BSSRDF comes from assuming a surface to be homogeneous and highly scattering.

$$S(\lambda, \mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) = F_t(\eta, \boldsymbol{\omega}_i)R(\lambda, \|\mathbf{x}_i - \mathbf{x}_o\|)F_t(\eta, \boldsymbol{\omega}_o). \quad (2.3)$$

In this case, provided that that index of refraction is known, the measurement process becomes one of estimating $R(\cdot)$, which is a function of one dimension [JMLH01b]. Jensen and colleagues used the setup shown in the slides to measure this function. Light is focussed at a single point on a planar, homogeneous surface, and this point is viewed by a camera. Since the surface is isotropically scattering, the radiance it emits is radially symmetric about the point of incidence. Thus, it is sufficient to examine a single surface line that contains this point, and examples of the intensities observed along one such line are shown in the graph. Following a calibration procedure, these observed intensities can be used to estimate the function R .

More complex acquisition systems can be built to measure more general BSSRDF representations that partially account for inhomogenieties. For example, Tong et al. [TWL⁺05], consider “quasi-inhomogeneous” materials that are homogeneous at a large scale but heterogeneous locally. Their representation is given by

$$S(\lambda, \mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) = f_i(\boldsymbol{\omega}_i)R_d(\mathbf{x}_i, \mathbf{x}_o)f_o(\mathbf{x}_o, \boldsymbol{\omega}_o), \quad (2.4)$$

which includes a spatially-varying “exiting function”. The acquisition system used to acquire data and fit this model incorporates a laser and a number of cameras and light sources.

2.4 Calibration

Cameras and light sources used for acquisition must be calibrated both geometrically and radiometrically. In addition, if objects of general shape are being used, their shape must be precisely known.

2.4.1 Geometric calibration

Geometric calibration involves the recovery of a camera’s extrinsic parameters (position and orientation relative to a world coordinate system) and intrinsic parameters (focal length, radial distortion parameters, etc.). Free and reliable tools for geometric camera calibration are readily available [OCV, CCT]. Likewise, geometric calibration for light sources requires the determination of their positions and orientations. The positions of point sources are typically measured by placing a number of shiny spheres at the position of the material sample. When the camera and sphere positions are known, highlights observed in the images of the illuminated spheres provide constraints on the source positions (e.g. [YNBK07]). Source orientation is typically not measured directly. Instead, one simply ‘locks down’ the sources in their known positions and measures the non-uniformity of their angular output patterns during a radiometric calibration procedure.

2.4.2 Radiometric calibration

Radiometric camera calibration involves two stages. First, one must determine the radiometric response function of the camera. This is the non-linear mapping that often exists between the irradiance incident on the image plane and the recorded intensity. Standard methods for doing this exist, and they usually involve imaging a calibration target or capturing multiple exposures of a static scene [DM97, MN99]. A second step is required to recover the optical fall-off of the lens system. An imaging system that uses an ideal thin lens exhibits a relationship between scene radiance and image irradiance that falls-off as $\cos^4 \alpha$, where α is the angle between the principle incoming ray and the optical axis. In a real camera, the optical fall-off must be measured for each zoom and aperture setting. It can be measured, for example, by capturing an image of a uniform, diffuse surface on a cloudy day.

A severe limitation that comes with the use of digital cameras for reflectometry is their limited dynamic range. In order to measure high radiance values at specular peaks while maintaining sufficient signal to noise ratios in darker regions, one must acquire several images with different exposures and merge them into a single high dynamic range (HDR) image [DM97, MN99]. If one has a large number of sources in the acquisition system, one can additionally multiplex the lighting to reduce the required dynamic range [SNB03].

Radiometric source calibration involves the measurement of output non-uniformity in a coordinate system that is rigidly associated with the source. One typically does this by imaging a planar diffuse surface (e.g., Spectralon) whose position is known relative to a camera and source.

Common sources of noise include thermal noise in the camera and temporal source variations. The former can be reduced by capturing a 'black' image for each camera (e.g., with the lens cap on) and subtracting this from all images used for measurement. Temporal source variations are more difficult to deal with, but fortunately, with advances in LED lighting, this is becoming less of an issue.

2.4.3 Colorimetric calibration

Typical camera-based acquisition systems limit spectral measurements to three wide-band observations (RGB). If this is done using a color camera equipped with Bayer filter, one must be sure to do all processing on raw data as opposed to that which is distorted by a demosaicking algorithm. Ideally, color information should be acquired using a grayscale camera equipped with a set of filters so that trichromatic or multi-spectral measurements are obtained at every pixel. Advances in electrically tunable liquid crystal spectral filters may eventually enable this approach to provide high spectral sampling rates in a reasonable amount of time.

Whenever RGB reflectance measurements are made, one must be aware that these measurements are valid only for the particular spectral power distribution of the light source. Using these triples to render synthetic images under different lighting will generally not produce physically accurate results. In addition, if multiple cameras are used, one must

compensate for variations between their filter sets. An example of system that relies on these calibration procedures will be described in the final section of this document.

2.4.4 Shape

If the materials being measured exist on arbitrarily-shaped surfaces, the shape of these surfaces must be known in the same coordinate system defined by the geometric calibration of the cameras and sources. From the perspective of measuring reflectance, three-dimensional surface reconstruction and alignment can thus be viewed as another calibration step.

In this context, one of the requirements of a 3D reconstruction system is that it recovers shape in a manner that is not biased by the material properties of the surface being measured. Indeed, we do not want the calibration of the measurement system to depend on the signal (reflectance) being measured. Common approaches to recovering shape for general materials is to use structured lighting from lasers or projectors. This can work well even for shiny and translucent surfaces [CLFS07]. One disadvantage of this approach is that the recovered shape must then be aligned with the images used for measuring reflectance. Any alignment errors are manifested in noisy reflectance samples. Another disadvantage is that it does not directly estimate surface normals, which are ultimately required for measuring reflectance.

Estimating surface normals from a range-scan requires differentiation of the discrete shape, and this is an additional source of noise for reflectance measurements. By directly estimating surface normals, this source of noise can be eliminated. Surface normals can be estimated using photometric stereo, but in its classic formulation, this violates the requirement of being independent of reflectance. Recently, we have seen the development of color-based techniques can be used to create photometric stereo systems that are independent of *dichromatic* reflectance [MZKB05], and Helmholtz stereopsis provides a means of estimating surface normals in a manner that is independent of reflectance for any surface whose reflectance can be represented by an SVBRDF [ZBK02].

Current best practice is to recover coarse geometry using structured lighting and then to combine this with surface normal estimates obtained photometrically [NRDR05]. An example of this will be discussed in the final section of this document. An alternative approach is to develop methods that simultaneously estimate shape and reflectance. To date, this approach has been restricted to methods that rely on parametric reflectance models [GCHS03, Geo03]. Doing this for more general representations of reflectance remains an open problem.

3 Spatially-Varying Reflectance Models

Jason Lawrence, University of Virginia

The appearance of opaque surfaces is characterized by the Spatially-Varying Bidirectional Reflectance Distribution Function (SVBRDF). The SVBRDF is equal to the amount of light reflected from an object's surface as a function of the position along the surface (u, v) and the direction of incidence ω_i and reflectance ω_o , parameterized with respect to the local surface normal and tangent directions:

$$S(u, v, \omega_i, \omega_o) \tag{3.1}$$

Technically, the SVBRDF also depends on wavelength λ , although we will generally ignore this by assuming it is expressed within a tristimulus color space (e.g. RGB, HSV, etc.). A useful way to think about this function is that it encodes the BRDF at every surface location, which is readily verified by fixing the spatial coordinates, resulting in a 4D function defined over incoming and reflected directions. Therefore, the SVBRDF has 6 degrees of freedom and it is precisely this high dimensionality which makes it difficult to measure and represent.

This part of the class briefly reviews modern methods for measuring the SVBRDF of real-world materials before focusing on the challenges and emerging strategies for representing these datasets.



Figure 3.1: Images of a dove greeting card taken from different viewing angles and under varying point illumination.

3.1 Acquisition

Measurements of the SVBRDF of a real-world object are captured in images taken from different viewing directions and under varying point illumination. Figure 3.1 show a few measurements of a slightly embossed greeting card that shows a metallic dove framed by a depiction of the sun. Note the change in both viewing direction and light position in these images.

This same measurement strategy also works for curved objects, although knowledge of the 3D shape is required to correctly interpret the position of each measurement within the local coordinate frame.

3.2 Representation

The datasets that result from these acquisition procedures are typically massive. For example, the dove dataset in Figure 3.1 consists of 5×400 camera x light positions for a total of 2,000 images covering a spatial area of 470×510 (totaling 5.5 GB of storage). However, many applications in computer graphics and vision require a compact representation. The clearest example from graphics is perhaps interactive rendering tasks which have strict bandwidth requirements. In the context of physically-based or global illumination rendering, it is also important that a representation support efficient sampling so that it may be used within Monte Carlo simulations of light transport. Finally, the ability to edit these datasets is beginning to receive more attention as measured materials move into production settings where designers require the same level of control that conventional parametric models have provided.

This part of the class focuses on only two of these goals which turn out to be very related: providing a representation that is both **compact** and **editable**.

3.2.1 Basis Decomposition

The general strategy is to perform some type of basis decomposition of the SVBRDF. Specifically, the input measurements are projected into a K -dimensional linear subspace spanned by the functions $\rho_k(\omega_i, \omega_o)$ (which are defined over the same domain as the BRDF) and $T_k(u, v)$, the coordinates within this basis that best fit the input:

$$S(u, v, \omega_i, \omega_o) \approx \sum_{k=1}^K T_k(u, v) \rho_k(\omega_i, \omega_o) \quad (3.2)$$

Whenever K is less than the average number of reflectance measurements at each surface location, this process compresses the data at the cost of some numerical error. In practice, K is typically several orders of magnitude smaller so the compression ratios are significant (e.g., $K = 3$ is sufficient to represent the dove dataset).

Although there are potentially many different K -dimensional bases with comparable numerical error, the goal of producing a final representation that supports editing requires

identifying those which reveal the *intuitive latent structure* in these datasets. In particular, we will often be interested in computing a decomposition that is consistent with how the input was physically manufactured. In the case of the dove greeting card we would prefer a 3-dimensional basis that clearly separates the two types of colored paper (yellow and blue) from the metallic material distributed in the shape of a dove. The main difference between the strategies we will review relate to the properties and constraints that are placed on the bases and coordinates and the specific algorithms used to execute the decomposition.

3.2.2 Parametric Models

Existing representations can be broadly classified as being either *parametric* or *non-parametric* which refers to the representation of the basis BRDFs. One of the early parametric methods was introduced by Lensch et al. [LKG⁺01, LKG⁺03] (earlier work by McAllister [McA02a] is similar, but without the critical clustering step that computes a low-dimensional basis onto which the data is then projected). Lensch et al. [LKG⁺01] acquire high-dynamic range images of an object under varying positions of a light-source and from different viewing angles. As is often the case for this type of acquisition, the position of the light is estimated from its image observed in a set of precisely aligned and highly-reflective spheres (e.g., ball bearings) placed in the camera's field of view.

A 3D model of the object's shape is obtained from either a structured light scanner (after coating the object in a removable diffuse powder) or a CT scanner. Each reflectance image is then aligned to the 3D model using an iterative optimization procedure that maximizes overlapping silhouette boundaries.

Once aligned, the set of image values at a particular surface location can be interpreted as measurements of the BRDF at that position. The parameters of a single-lobe isotropic Lafortune analytic BRDF model [LFTG97] to a subset of the entire collection of BRDF measurements collected across the surface. Because this analytic model is non-linear in its parameters, this step requires non-linear optimization for which they use the Levenberg-Marquardt (LM) algorithm.

The covariance matrix resulting from the LM algorithm provides the direction in this parameter space along which there was maximum variance in the observations. This information is used to drive a divisive clustering algorithm that repeatedly splits the cluster with the greatest inter-cluster error, generating two new cluster centers positioned on opposing sides of this principal direction of variance. This process continues until either a user specified number of clusters or an error threshold is met.

Lastly, the cluster centers are interpreted as defining a linear basis onto which the measurements at each location are projected, a process that requires solving a single linear system. This method achieves intuitive separations for several real-world datasets, although there are situations in which a clustering approach would fail to isolate a SVBRDF's component BRDFs (some examples are shown in Section 3.3).

Goldman et al. [GCHS05a] also describes a parametric approach to representing spatially-varying reflectance. However, they chose to fit measured data to the Ward BRDF model [War92] and added the surface orientation at each point as free parameters in the

optimization (in addition to the basis BRDF parameters and spatial blending weights). Finally, they specifically discuss how the final representation supports editing.

3.3 The Inverse Shade Tree Framework

Conceptually, non-parametric models are computed in an identical fashion in that some type of Expectation-Maximization algorithm is used that alternates between estimating the basis BRDFs and their blending weights. The key difference is that the basis BRDFs are represented in tabular form or within some secondary basis (e.g., Radial Basis Functions or wavelets). This approach naturally provides greater fidelity to measured data, but comes at the cost of larger datasets and more delicate interpolation and data processing.

The Inverse Shade Tree (IST) framework [LBAD⁺06b] was the first project to explore a fully editable non-parametric representation of spatially-varying reflectance. The key idea is to perform a hierarchical decomposition of measurements of the SVBRDF, shown at the root of the tree diagram in Figure 3.2, into lower-dimensional components. This achieves the combined goal of compressing the data (the nodes at each level require less storage than their parent) and providing an editable representation (the leaf nodes of this tree reveal intuitive and meaningful latent structure in the dataset that can be directly modified by a designer).

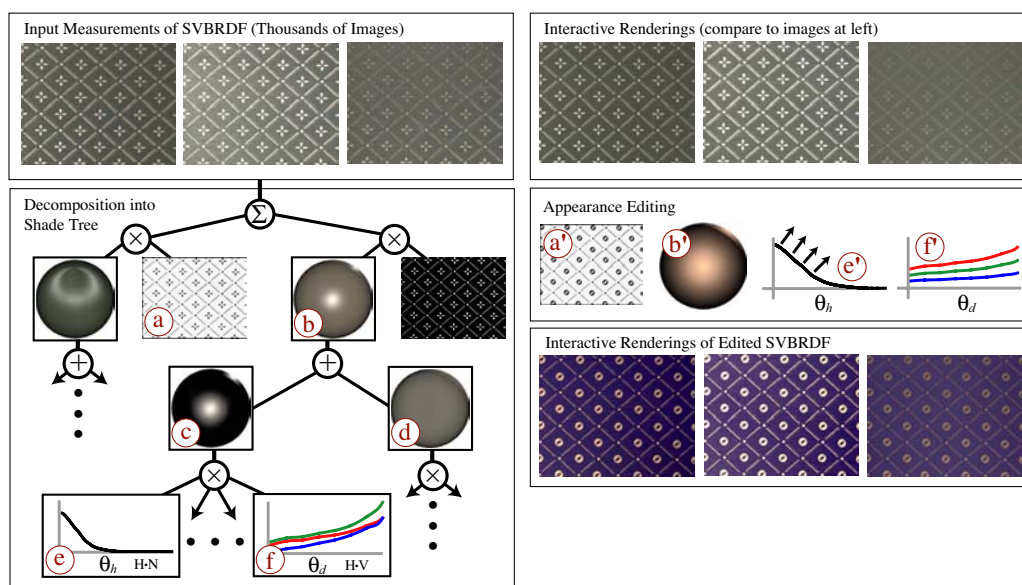


Figure 3.2: The Inverse Shade Tree framework [LBAD⁺06b] introduces techniques for decomposing measured SVBRDF data into a set of (a) spatially-varying blending weight maps and (b) basis BRDFs. The basis BRDFs are factored into sampled 2D functions corresponding to (c) specular and (d) diffuse components of reflectance (we show lit spheres rendered with these factors, not the 2D factors themselves). These 2D functions are further decomposed into (e & f) 1D curves. In addition to providing accurate interactive rendering of the original SVBRDF, this representation also supports editing either (a') the spatial distribution of the component materials or (b') individual material properties. The latter is accomplished by editing (e' & f') the sampled curves.

The concept of a *shade tree* was introduced by Rob Cook in 1984 [Coo84] as a way of assembling complex reflectance functions by combining simple parametric and sampled functions using various combination operators, working from the leaf nodes up to the root. This IST framework essentially inverts this process: decomposing a complex reflectance function into a shade tree by working from the root down toward the leaves. The other important contribution of [LBAD⁺06b] was relating the decomposition at each level to factoring an appropriately constructed matrix.

Computing a shade tree from measured SVBRDF data involves several steps. At the top level, measurements of the SVBRDF are first organized into a regularly spaced matrix by simply unrolling each image into a separate column. This preserves the spatial variation of this function along the rows of this matrix and the angular variation along its columns (color variation may be preserved along either rows or columns, although the original paper did not explore this trade-off and instead computed only colorized BRDFs). Factoring this matrix computes a sum of products of functions that depend only on spatial position and angular position (both incoming and reflected directions). These are estimates of the basis BRDFs, stored in tabular form, and the blending weights, respectively. Many standard factorization algorithms will not necessarily favor an intuitive or meaningful decomposition. Therefore, the key research challenge in this work is to design an algorithm that generates not only an accurate factorization, but one that is also intuitively editable.

A common strategy for these types of problems is to place constraints on the resulting factors that guarantee they are physically plausible. BRDFs and spatial blending weights have physical properties that can be translated into constraints. For example, reflectance functions are non-negative. Therefore, we may restrict the factorization to be non-negative to guarantee it doesn't violate this property. Another property of natural materials that becomes a constraint on the optimization is that they are typically sparse: even though a dataset might be composed of several unique BRDFs, there are typically only a few blended together at any one surface location. Finally, there are domain-specific constraints such as enforcing the basis BRDFs conserve energy and are monotonically decreasing along certain directions.

There are a variety of algorithms available for computing a matrix factorization. We briefly compare existing approaches and discuss the conditions under which they fail to provide a meaningful decomposition.

PCA/ICA: Two popular rank reduction algorithms are Principal Component Analysis (PCA) and Independent Component Analysis (ICA), along with extensions such as multilinear tensor factorization [VT04]. The main advantage of PCA is that it yields a *global* minimum in the sense of total least squares. However, these algorithms recover a basis that is orthonormal (for PCA) or statistically independent (for ICA). These restrictions are not sufficient to produce a meaningful description of the data. In particular, they allow negative values, resulting in a representation whose terms cannot be edited independently

Clustering: One popular method for clustering data is the k -means algorithm [HW79]. Like all clustering algorithms, k -means partitions the input into disjoint sets, associating each point with a representative *cluster center*. This can be interpreted as a factorization of the SVBRDF. Although clustering performs well on input with a small basis that is well-separated over the surface, it typically fails to recover a useful basis when the SVBRDF exhibits blending of its component materials.

Non-Negative Matrix Factorization: Another matrix decomposition approach is Non-Negative Matrix Factorization (NMF) [LS99]. Together with similar algorithms such as Probabilistic Latent Semantic Indexing [Hof99], NMF guarantees that both resulting factors contain only non-negative values. One motivation for this constraint is to encourage the algorithm to describe the input data as the sum of positive *parts*, thereby producing a more meaningful factorization. However, the character of the decomposition is sensitive to small changes in the data (including those due to measurement noise and misalignment), and the non-negativity constraint is not always enough to guarantee an editable separation.

3.3.1 Alternating Constrained Least Squares

Lawrence et al. [LBAD⁺06b] introduce an algorithm for computing the factorization $Z \approx WH$, subject to general linear constraints on W and H . This algorithm is built upon efficient numerical methods for solving convex quadratic programming (QP) problems of the form:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{b} - Mx\|^2 \quad \text{subject to} \quad \mathbf{l} \leq \begin{Bmatrix} x \\ Ax \end{Bmatrix} \leq \mathbf{u} \quad (3.3)$$

The n -element vector x is called the *vector of unknowns*, M is called the *least-squares matrix* and \mathbf{b} is the *vector of observations*. The vectors \mathbf{u} and \mathbf{l} provide the upper and lower *bound constraints* of both x and the linear combinations encoded in the matrix A , called the *general constraints*. There are several algorithms available for solving these types of problems. They used an inertia-controlling method that maintains a Cholesky factorization of the reduced Hessian of the objective function [GMSW84].

As with NMF, W and H are initialized to contain positive random values, and ACLS minimizes the Euclidean error of this approximation by alternately updating these two matrices. This problem is known to be convex in either W or H separately, but not simultaneously in both. As a consequence, ACLS is guaranteed to recover only a *local* minimum.

Without loss of generality, consider the case where both V and W are row vectors ($v \approx wH$). For fixed H , the current estimate of w is updated to minimize the Euclidean distance with the corresponding row in V , subject to the linear constraint $w \geq 0$. This can be accomplished by solving a QP problem in Equation 3.3, with $M = H^T$, $\mathbf{b} = v^T$, and $x = w^T$. To constrain the solution to be non-negative, set $\mathbf{l} = 0$ and $\mathbf{u} = \infty$.

The entire matrix W is determined by computing the above solution for each of its rows in turn. Similarly, H is computed one column at a time. Alternating between estimating W and H achieves a non-negative factorization of the input matrix V . Note these steps

are guaranteed to never increase the Euclidean error, thus ACLS eventually converges to a stationary point.

Sparsity is considered by modifying the objective function in a way that penalizes the sum of the non-maximum elements in w . Evaluating this measure of sparsity at several different coordinate positions shows it has the desired effect of favoring a solution that is closely aligned to one of the basis axes. This penalty is weighted against the Euclidean error in the objective function through a user-specified parameter.

We compare the blending weights obtained from ACLS to other factorization algorithms for the wood+tape datasets (please consult the talk slides for comparisons). To visualize the output of Singular Value Decomposition (SVD), we mixed positive values into the green color channel and negative values into the red color channel. Although this is the optimal result in terms of RMS error, it fails to provide an editable characterization of this dataset. Non-Negative Matrix Factorization (NMF) [LS99] shows a better separation due to its non-negativity constraint, but it still shows significant blending between the component materials. A k -means clustering algorithm provides the most sparse separation, but cannot recover a separate term for the Scotch tape in this dataset due to its inability to account for linear combinations of multiple terms. Finally, ACLS does a better job of disentangling the component materials in this dataset, but at the expected cost of larger numerical error.

Each 4D BRDF is further factored into the sum of products of 2D functions which correspond to its dominant reflectance lobes (i.e., backscattering lobe, specular lobe, diffuse lobe, etc.). Each of these are further factored into the product of 1D curves which capture salient features of the BRDF such as the shape and size of its specular highlight (as seen in a curve of the half-angle) and brightening and color shifts toward grazing angles due to Fresnel effects (as seen in a colorized curve defined over the difference-angle). It is these curves that a designer can directly manipulate to control the scattering properties of the component materials. We demonstrate three types of edits: modifications to the specular highlight of the metallic silver material in the dove dataset, replacing a subtree with curves computed from entries in the MERL/MIT database of isotropic BRDFs [MPBM03b], and interactive modifications to the spatial blending weights to alter the position of the Scotch tape.

3.4 Conclusion and Directions of Future Research

To be useful in practice, representations of measured reflectance data must be compact, accurate, and support editing. It's also important to consider how they may be integrated within existing rendering systems, although this topic was not discussed in detail (see [LRR04, CJAMJ05] for information on this topic).

Representations based on basis function decomposition typically provide the greatest fidelity to measured data. Furthermore, approximating the input in a low-dimensional basis achieves significant compression and reveal intuitive latent structure that allows editing the final result. Existing methods can be differentiated based on whether a parametric or non-parametric model is used to represent the basis BRDFs.

3.4.1 Parametric vs. Non-Parametric

It's important to understand the trade-offs involved in using a parametric or non-parametric model of the basis functions.

One real practical problem concerns scattered data interpolation. Keep in mind that these datasets often consist of measurements scattered across a high-dimensional domain. A notoriously difficult task is to reconstruct a continuous approximation from these samples, but is required to use them directly in a rendering system or, in the case of the IST framework, to generate the input matrix (see [LBAD⁺06b, Law06] for further details). On the other hand, fitting the parameters of an analytic model to scattered data is often easier and avoids any explicit interpolation stage (it can in fact be regarded as performing interpolation for a restricted set of continuous approximations). Generic data interpolation techniques have been applied in this context. Notable examples include the push-pull algorithm, first used with surface light field data [GGSC96], and techniques such as fitting Radial Basis Functions [ZERB05].

One area where non-parametric methods outperform parametric methods is in their inherent flexibility and accuracy. Obviously, a tabulated grid of numbers can represent a much wider range of functions than those captured by an analytic function of a handful of parameters. In fact, we've seen that analytic models are designed to model a specific class of materials, making a parametric approach particularly error-prone for datasets that include different types of materials (e.g., wood+tape dataset). The downside, of course, is that there is more data to consider.

Finally, please note that all of these techniques are susceptible to poor local minima. Non-parametric methods are perhaps better in this regard since they do not involve the non-linear optimization that appears with parametric methods which can be unstable for multi-lobe models or particularly sparse and noisy input. Nevertheless, they are often very sensitive to their starting position and future work should consider optimal ways of collecting input from a human user to make this automatic separation more robust.

3.4.2 Open Problems

Perhaps the clearest direction of future research is to develop techniques that bring a wider range of appearance functions into this type of inverse shade tree framework or at least focus on the goal of editing. Examples include early work in developing representations for heterogeneous translucent materials [PvBM⁺06b] and time-varying materials [GTR⁺06, WTL⁺06].

The fact that all of these basis decomposition techniques are fundamentally based on Expectation-Maximization suggests placing them in a unifying probabilistic framework. In particular, note the deep similarities between the hierarchical decomposition used in the IST framework and hierarchical probabilistic models [GaHSSR04] that are gaining traction with many problems in machine learning [Jor99, BNJ03] and computer vision [FFP05]. This type of perspective may help clarify the assumptions existing representations make and set the stage for generalizing these algorithms to work with a wider range of datasets.

Additional research in measuring these datasets is also necessary. In particular, it's worth considering ways of performing *synchronous* measurement and appearance. Objects with optical properties that are difficult to model by hand (and thus justify taking a data-driven approach) also tend to be difficult to scan. What novel optical setups support collecting reliable measurements of both shape and scattering? Also, the calibration burden of these methods remains prohibitively high. What can be done with sparse and noisy measurements and what devices might eliminate the need for such fragile calibration procedures?

4 From BSSRDF to 8D Reflectance Fields

Hendrik Lensch, Max-Planck-Institut für Informatik

In these notes we have so far concentrated on capturing surface reflectance where incident light is scattered locally at the point of incidence. The models that have been presented so far ignore global effects such as subsurface scattering, transmission, or interreflections in complicated surface geometry.

4.1 BTFs and Distant Light Reflectance Fields

A quite general representation that can represent even non-local effects are so-called bi-directional texture functions (BTFs) [DvGNK97]. In principle, a BTF captures the apparent spatially-varying BRDF at a point (x, y) for parallel incident light:

$$S(x, y, \theta_i, \varphi_i, \theta_o, \varphi_o), \quad (4.1)$$

i.e. it might contain data that cannot easily be described by a particular analytic BRDF model at each a point because the apparent reflectance might be influenced by masking, shadowing, or interreflections. To avoid the approximation errors by an analytic model, BTF data is quite often represented in tabulated form, i.e. tensors, and general compression schemes such as wavelets, higher order PCA, clustering, or spherical harmonics representations are frequently applied [MSK06].

The definition of a BTF is equivalent to the definition of a reflectance field [DHT⁺00a] for distant lighting. The only difference being that in a BTF the points (x, y) are typically defined over the surface of a small material patch while in the reflectance field approach (x, y) typically refer to locations in a camera image. For this reason, the same techniques for compressing and rendering of such reflectance fields can be applied.

For acquiring BTFs, all incident and outgoing directions need to be sampled: Various approaches have been proposed for acquiring BTFs from simulated data [TZL⁺02] or for measuring BTFs using moving robot gantries [DvGNK97, MMS⁺05, DHT⁺00a], static devices based on mirrors [HP03] for multiplexing a single camera, or multi-camera setups [MBK05]. A good survey on acquiring, processing and rendering BTFs efficiently can be found in [MMS⁺05].

4.2 BSSRDFs

BTFs and reflectance fields introduced so far assume that the light hitting the surface is parallel, originating from an infinitely far away light source. If one wants to simulate a close by light source, or, equivalently, to project a light pattern into the scene, e.g. from a spot light or as the result of two objects interacting with each other, the far-field assumption is not valid. Instead, one needs to record the so-called near-field reflectance field which couples incident to outgoing light rays. Here, the reflectance depends both on the point of incidence and the reflection point.

For surfaces, this distinction of incident and reflection point is necessary only if significant interreflections or subsurface scattering is observed. Subsurface scattering can be efficiently described by the Bi-direction Scattering Surface Reflectance Distribution Function (BSSRDF):

$$S(x_i, y_i, \theta_i, \varphi_i, x_o, y_o, \theta_o, \varphi_o) = \frac{dL(x_o, y_o, \theta_o, \varphi_o)}{d\Phi(x_i, y_i, \theta_i, \varphi_i)} \quad (4.2)$$




Figure 4.1: Acquisition setup of Jensen *et al.* [JMLH01a] for measuring homogeneous BSSRDFs.

Note that the BSSRDF is actually defined as the quotient of differential reflected radiance over the incident flux while a BRDF is the reflected radiance over the irradiance. The definition of a BSSRDF is again equivalent to the definition of a full 8D reflectance field [DHT⁺00a].

Because of their high dimensionality BSSRDFs are in general hard to represent and to acquire. A simple analytic BSSRDF model for homogeneous materials has therefore been proposed by Jensen *et al.* [JMLH01a]. It has later been updated to incorporate homogeneous multi-layer materials [DJ05a]. In the first paper, a simple measurement setup is presented in for estimating the parameters of the model: A collimated light beam hits the homogeneous slab at one point, and the resulting spatial distribution of the reflected light is measured using a

camera. A simpler device for instant measurement of a few samples has been proposed in [WMP⁺06a].

4.3 Diffuse Subsurface Scattering

For heterogeneous surfaces it is impractical to densely sample all eight dimension of the BSSRDF/reflectance fields. In order to allow for reasonable sampling effort, one strategy is to assume a less complex light transport.

Goesele et al. [GLL⁺04a] present a measurement setup for acquiring the appearance of translucent objects with a high scattering albedo. In these cases, a photon travelling some distance through the material undergoes so many scattering events that the incident light direction has actually no influence on the outgoing light direction. Since the directional dependence can be dropped from the full 8D BSSRDF, the problem can be reduced to a 4D diffuse scattering $R(x_i, y_i, x_o, y_o,)$ function that solely depends on the point where the light enters the material and the position where it leaves when being reflected. For any pair of points R indicates how much the incident irradiance at point (x_i, y_i) contributes to the outgoing radiosity at point (x_o, y_o) .

A simple 4D tensor can be used to represent this 4D function. In order to measure its entries Goesele et al. make use of a laser projector, that sweeps an individual light point over the surface of a translucent object (see Figure 4.2). A set of HDR video cameras capture




Figure 4.2: Acquisition setup (a) of Goesele et al. [GLL⁺04a] for measuring the appearance of heterogeneous translucent objects such as this alabaster figurine (b).

the reflected light at every other surface point. One of these measurements correspond to exactly one slice of the 4D tensor. Illuminating every surface point once eventually fills the entire tensor. Because of occlusions and self-shadowing it is however likely that for some parts of the object no measurements are available, in which cases texture inpainting is applied to fill the gaps. For efficient storage, a hierarchical representation of the tensor is chosen, providing a high sampling rates only close to the point of incidence, where the BSSRDF drops of quickly, while for distant points a coarser sampling is sufficient.

Recently, Peers et al. [PvBM⁺06c] presented a different setup to acquire the spatially varying 4D BSSRDF of a planar slab of material. In order to accelerate the acquisition, a grid of points is swept over the surface. A high compression on the captured tensor is achieved by aligning the main features, i.e. the point of incidence in every row, followed by a principle component analysis.

4.4 Arbitrary Light Transport

For arbitrary materials and scenes, Masselus et al. [MPDW03] presented the first acquisition system for reflectance fields that are suitable for relighting with 4D incident light

fields, i.e. where the reflected light depends on individual incident light rays. For the acquisition, a video projector swept a small block of light over the scene. In order to cope with the complexity of the acquisition problem the appearance was captured for a single viewpoint only. Additionally, the resolution of the incident light field was limited to a projector resolution of only 16×16 for a couple of projector locations. This low resolution results in clear block artifacts in the relit images.

4.4.1 Single View – Single Projector

In order to avoid those artifacts, it is necessary to measure the reflectance for every pair of rays between a camera and a projector, i.e. to acquire the reflectance for every pair of camera and projector pixels, again resulting in a fourth order tensor. While in principle this high resolution reflectance field could be acquired using scanning, it would be a too slow process.

Sen et al. [SCG⁺05] exploited the fact that for quite a number of real-world scenes the light transport matrix/reflectance field is rather sparse, i.e. that only a small fraction of the possible input rays actually contribute to the same reflected ray. In this case, it is possible to exploit the sparseness by illuminating the scene and measuring the reflected light rays for multiple illuminating light rays at once. It is possible to turn on two light rays/two projector pixels at the same time and tell their corresponding measurements apart when those two rays affect completely separated parts of the scene/the camera images. We call such two illumination rays radiometrically independent. In the same way one can also call two blocks of the projector pixels radiometrically independent if no camera pixel will be illuminated by both at the same time. Because of this property it is possible to measure the rays inside the independent block in parallel, i.e. to parallelize the exact acquisition of these two blocks. In their paper, Sen et al. propose a hierarchical approach for determining which sub-blocks are independent: Starting from a full white projector image, each block is subdivided into four children which again get subdivided. Initially, this will require one measurement per block corresponding to a sequential acquisition. At some point in time the algorithm might however detect that at some level two blocks are now radiometrically independent, allowing for parallelized subdivision of these blocks in the future. The net effect of this parallelization is significant, resulting in a complexity that is $O(\log(n))$ for n projector pixels. For quite a number of scenes the pixel-to-pixel reflectance field between a one mega-pixel projector and a camera can be acquired in only a couple of thousand images instead of a million.

Once having acquired the pixel-to-pixel light transport one can apply Helmholtz reciprocity to invert the role of projectors and cameras. Helmholtz reciprocity states that the reflectance measured for one path does not change no matter if one follows the path from the light source to the receiver or the other way around. One simply has to compute the transpose of the acquired tensor to obtain the reflectance field from the camera (the new virtual projector) to the projector (which gets the new camera). The transpose corresponds to just a resorting of rays, and therefore can be computed very efficiently (see Figure 4.3).

This dual imaging paradigm can be used to efficiently capture a 6D reflectance field from a single viewpoint, i.e. to measure the projector to camera reflectance fields for multiple

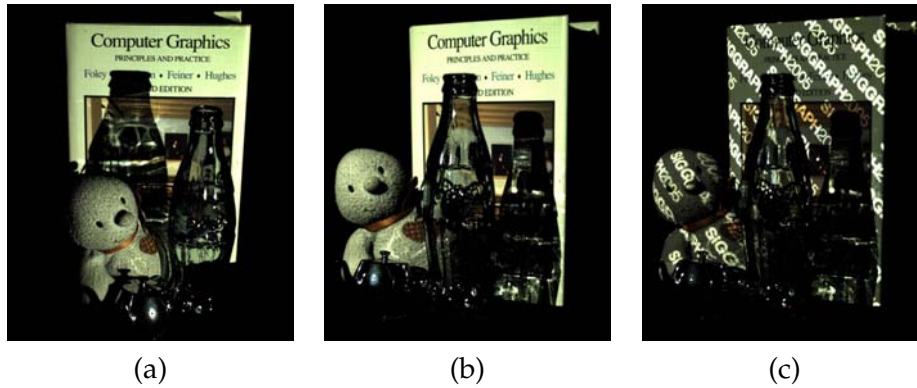


Figure 4.3: *Dual Photography: (a) Conventional photograph of a scene, illuminated by a projector with all its pixels turned on. (b) After measuring the light transport between the projector and the camera using structured illumination, dual photography is able to synthesize a photorealistic image from the point of view of the projector. This image has the resolution of the projector and is illuminated by a light source at the position of the camera. The technique can capture subtle illumination effects such as caustics and self-shadowing. Note, for example, how the glass bottle in the primal image (a) appears as the caustic in the dual image (b) and vice-versa. Because we have determined the complete light transport between the projector and camera, it is easy to relight the dual image using a synthetic light source (c).*

projectors. The problem is that during the acquisition the reflectance fields have to be captured sequentially for each projector because projectors are active devices. Their projected patterns might actually illuminate the same points in the scene causing difficulties when trying to separate their contribution. If one uses the dual setup however, where the original camera is replaced by a single projector and all projectors are replaced by cameras, one can very well acquire the projector/camera reflectance fields in parallel since cameras are passive devices which do not interfere with each other. Applying Helmholtz reciprocity, this setup can virtually be transformed into the single camera/multiple projector configuration. By swapping camera and projectors one can capture a 6D reflectance field at the same time cost as a 4D reflectance field. The resulting data now allows to relight an arbitrary complex scene with arbitrary incident light fields, i.e. with high frequency illumination patterns from various virtual projector positions.

4.4.2 8D Reflectance Fields

The previous acceleration for measuring the light transport between a projector and a camera is however limited to scenes where the light transport tensor is rather sparse. This is often the case for an individual object in a black room where few interreflections and little subsurface scattering take place. For more general cases, the light transport matrix is rather dense, i.e. every projector pixel indirectly affects every camera pixel due to multiple scattering events. The resulting light transport is however rather smooth for large blocks of the tensor. For example, illuminating one spot on a wall will have a rather similar effect to all points on an opposite wall. While this smoothness might be partially destroyed by textures on both walls, the underlying light transport still has rather low complexity or low dimensionality – it is called data sparse. In other parts in the ray space,

however, for example for direct reflections or refractions, the reflectance field might not be smooth at all.

\mathcal{H} -matrices [Hac99] are an efficient way for representing tensors which are partially data-sparse. In a \mathcal{H} -matrix the original matrix is hierarchically subdivided into smaller blocks, e.g. using a quad-tree for a 2D matrix, and for every sub-block a low-rank approximation is given, approximating the original matrix's entries. If the approximation error for one block is too large, the block is further subdivided. As \mathcal{H} -matrices have been originally developed to solve integral equations more efficiently, and since the Rendering Equation which describes the light transport in arbitrary scenes is an integral equation, reflectance fields can be very efficiently described by this data structure.

Besides resulting in a compact representation of a reflectance field, \mathcal{H} -matrices can efficiently be evaluated during relighting, where the incident light field is simply multiplied with the tensor.

\mathcal{H} -matrices further open the way for efficient acquisition of reflectance field of arbitrarily complex scenes where interreflections and scattering cannot be neglected, as well as for the acquisition of 8D reflectance fields.

Garg et al. [GTLL06] have proposed a measurement setup that forces the captured reflectance tensor to be symmetric. In the setup, every camera is paired with one projector using a beam splitter in such a way that it is possible to emit light and to measure light exactly along the same ray. In the resulting transport tensor every off-diagonal sub-block is therefore represented twice, once in its original form and once being transposed, i.e. we could capture the original and the dual image for one sub-block with just two images by fully illuminating the corresponding two projector blocks.

Since one of the images corresponds to the sum along the rows of the block and the other image to the sum along the columns of the block, it is possible to obtain a rank-1 approximation of this block with just these two images, simply as the tensor product of the two measurements obtained when first illuminating with one block of one projector, measuring the result in one block of some camera, and then measuring the transpose, i.e. measuring at the block of the first projector and emitting light from the first camera's block.

Let's look at a very simple example where the off-diagonal block B_2 has been determined to be rank-1:

$$T = \begin{pmatrix} B_1 & B_2 \\ B_2^T & B_3 \end{pmatrix} = \begin{pmatrix} 0 & B_2 \\ B_2^T & 0 \end{pmatrix} + \begin{pmatrix} B_1 & 0 \\ 0 & B_3 \end{pmatrix} \quad (4.3)$$

In this case, we can determine all entries in B_2 (and B_2^T) from just two images while B_1 and B_3 might require additional investigation.

From the intended solution one can subtract the already determined matrix

$\begin{pmatrix} 0 & B_2 \\ B_2^T & 0 \end{pmatrix}$, which leaves us with some very interesting rest which only contains the remaining, yet unknown blocks. Those two blocks are however arranged in a very interesting configuration: they are radiometrically independent, since they clearly effect completely different camera and projector regions. As a consequence, those two blocks can again be investigated further in parallel. It allows for the efficient and parallelized acquisition of even dense matrices as long as the matrices are data-sparse.

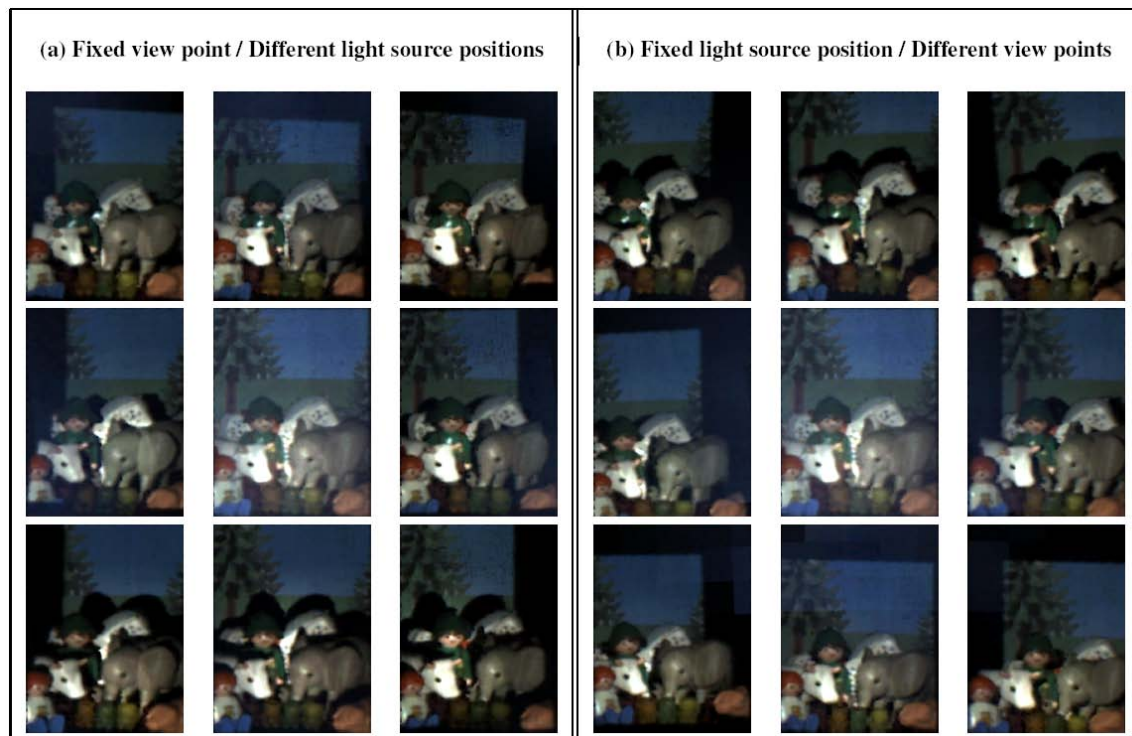


Figure 4.4: An 8D reflectance field acquired using Symmetric Photography. The scene has been recorded from 3×3 different view points and can be relit from 3×3 different projectors with full resolution.

In Figure 4.4, we show a low-resolution 8D reflectance field for 3×3 cameras and 3×3 projectors. With the symmetric photography approach one can acquire the light transport of scenes as complicated as this glass of gummy bears and faithfully reproduce the appearance of the original object (Figure 4.5).

4.5 Conclusions and Future Work

In this section we have introduced the notion of reflectance fields for relighting with spatially varying illumination patterns: from the acquisition of heterogeneous translucent objects to methods for acquiring the ray-to-ray light transport in arbitrary materials and scenes.

4.5.1 Open Problems

One big problem of sampling BSSRDFs or reflectance fields so far is the limited resolution with respect to the incident and outgoing directions. While solutions have been proposed to increase the resolution of the incident illumination [HED05, FBLS07] by using special light source arrangement, the resolution of the viewing directions is still limited to the




Figure 4.5: *Using Symmetric Photography the light transport in even very complex scenes can be efficiently captured. The synthetically relit reflectance field (a) matches the appearance of the original object under the same pattern (b).*

spacing between adjacent camera positions. A scheme for adaptively controlling the resolution in the viewing and the illumination direction still needs to be invented.

For representing reflectance fields, various bases have been proposed, wavelets, spherical harmonics, or the above mentioned \mathcal{H} -matrices. It remains to be seen how to select the optimal representation, and how to determine the dimensionality of the light transport locally.

A still outstanding goal is the acquisition of reflectance fields for relighting with 4D incident light fields for dynamic objects. While initial solutions to measuring time-varying far-field reflectance fields at interactive rates have been demonstrated [WGT⁺05, ECJ⁺06] the significantly higher complexity of near-field reflectance fields currently requires too many images for every pair of viewing and illumination directions.

5 The Human Face Scanner Project

Tim Weyrich, Princeton University

This part of the class presents a project that has leveraged principles of appearance acquisition and representation to acquire digital models of human faces. Creating digital faces that are indistinguishable from real ones is one of the biggest challenges in computer graphics. Although general rendering quality in graphics often achieves photo-realistic appearance of synthetic objects, rendering of human faces still remains a demanding task. This is not only because of the complexity of facial appearance, but also due to that fact that human observers are experts in judging whether a face is “real” or not. The process of capturing an actor’s performance and likeness has accurately been named “Digital Face Cloning.” Digital face cloning has many applications in movies, games, medicine, cosmetics, computer vision, biometrics, and virtual reality. While recent feature films already show authentic artificial renderings of character faces [BL03, Fee04], the datasets have been especially designed for an appearance in a particular scene. They are either the result of extensive manual editing to make the model appear real under hand-optimized lighting design [BL03], or they merely tabulate an actor’s lighting-dependent appearance from a single viewpoint [DHT⁺00b, HWT⁺04, Fee04]. Up to recently, there was no generic procedure that allows for the acquisition of human faces, leading to a self-contained face model that can be used for renderings in an arbitrary context. Key requirements of such a model are generality, to allow for a flexible use, and editability, that is, it should be possible to change the face’s appearance using intuitive controls. Up to date, editing of face models requires the hand of a skilled artist with a deep technological understanding.

The presented project addresses this issue, developing a face acquisition pipeline that allows for the automated acquisition of generic face models from individual subjects. This comprises the construction of a respective acquisition hardware, the development of a suited model representation, and the analysis of facial appearance across multiple subjects to derive meaningful controls for editing of the face model.

This chapter sketches selected topics of the project as presented during the class. For more in-depth information and a more comprehensive discussion of related work, refer to [WMP⁺06b, Wey06].

5.1 Previous Work

Capturing Face Appearance In accordance with the two major paradigms in appearance modeling, existing facial appearance modeling systems either use an explicit approach, explicitly modeling facial geometry and surface texture [PHL⁺98], or they em-

ploy image-based methods, such as reflectance fields over a rough impostor geometry [DHT⁺00b, HWT⁺04]. In general, explicit models are more directly accessible to editing operations, while image-based approach make it easier to achieve photo-realistic results and are less sensitive to measurement errors. In our work, however, we target an explicit modeling approach to maintain full flexibility for editing.

Skin Reflectance Models Particular attention has to be paid to the modeling of skin reflectance, as skin contributes the largest part of facial appearance. Previous work employs all major classes of reflectance models, BRDF [HK93, MWL⁺99, Sta01], BSSRDF [JMLH01b, DJ05b, DJ06], and BTF [CD02], see Chapter 2, to model skin reflectance. The presented project uses a compound model of BRDF and BSSRDF representations.




Figure 5.1: *Physiology of skin. On a micro-scale, human skin is a very heterogeneous tissue. However, at scales relevant for rendering, it is sufficient to consider the two prominent layers epidermis and dermis. The visual impact of scattering within the hypodermis is negligible. (Image from Wikipedia.)*

Appearance Editing Existing approaches to realistically alter facial appearance often focus on manual editing of facial textures, or use image-based techniques to alter skin appearance in photographs [TOS⁺03]. A higher-level editing technique is to use morphable face models to blend between shape and texture of multiple individual face models [BV99, FLS05]. In terms of expressiveness, we target similar high-level operations for our representation.

Production Environment One of the most prominent areas of application of digital face cloning are film productions. Using latest techniques in face representation, they also use a significant amount of manual labor to meet the requirements of a single shot [Wil05, Hér03, BL03, Sag04]. Increasingly, realistic face rendering gains importance in computer games.

5.2 Skin Appearance Acquisition

Covering most of face, skin is the most important aspect of facial appearance. The dominant effect in skin reflectance is due to skin's translucent layers, see Figure 5.1. In a rough approximation, skin consists of two optically active layers, the epidermis and the dermis underneath it. Light transport is mainly affected by two effects: a) Surface reflection at the air/skin interface (that is, surface reflectance that can be described by a BRDF and refraction that propagates light into the skin); b) Scattering and absorption within epidermis and dermis that can be expressed as a BSSRDF. This two-fold separation is a commonly used simplification and reflects our goal to develop a skin model of high practical value. More sophisticated, multi-layered, models exist. An excellent survey on the physiological

and anatomical properties of skin and the state of the art in skin appearance modeling has been published by Igarashi et al. [INN05].

We built two custom reflectance acquisition devices to capture surface and subsurface reflectance independently. A contact-device allows for sparse subsurface measurements across the face. Measuring pure surface reflectance, however, is difficult. Hence, our second device captures the full skin reflectance as a whole. Skin reflectance being a sum of subsurface and surface components, surface reflectance can then be derived by subtracting the measured subsurface component from the skin reflectance measurements.

5.2.1 Subsurface Scattering Acquisition

The first device, internally referred to as the “BSSRDF Gun”, is a contact-measurement device to measure purely subcutaneous light transport. The quantity we are interested in is the wavelength-dependent mean free path ℓ , or skin’s *translucency* $1/\ell$, respectively. ℓ is a measure of how far photons travel in average between two scattering events. The device feeds light into the skin using an optical fiber. A bundle of optical fibers arranged around the feeding fiber collects light exiting the skin at different distances to the feeding fiber. Digitizing the radiant exitance using an HDR camera at the end of the fiber bundle (encased by a light-proof box) allows to measure the characteristic radial fall-off, that is, the diffusion kernel due to skin’s subsurface scattering. We obtain ℓ by fitting the dipole approximation of the diffusion model BSSRDF [JMLH01b] to the measured fall-off.

As we will see in Section 5.5, translucency only varies minimally across the face. Hence, only a few measurements of the BSSRDF Gun are required to obtain representative translucency values of a face. Note that an even contact between the fiber probe and the skin is required to eliminate surface reflectance effect. We ensure this by gently evacuating the sensor using a suction pump.

5.2.2 Reflectance Field Acquisition

The second measurement device samples skin reflectance as a whole, that is, the sum of subsurface and surface effects. The device captures a 150×16 *reflectance field* of the face. The reflectance field [DHT⁺00b] is a tabulation of images from, in our case, sixteen different viewing directions under 150 different illumination directions. To that end, we built a spherical dome containing sixteen 1300×1030 firewire cameras and 150 LED light sources (each of them being a disk-shaped panel of 150 LEDs). In order to be able to associate pixels in the reflectance field with surface points on the face, we also acquire the facial geometry using a commercial single-shot 3D scanner based on stereo reconstruction of an IR random speckle projection. Our setup is completely synchronized, that is, the 150 light sources are sequentially triggered, while all cameras simultaneously acquire images at 12 frames per second. We are imaging each light condition under two different exposure times to increase the dynamic range of our measurements. Hence, a full acquisition takes about 25 seconds.

The resulting reflectance field has to be radiometrically corrected, according to the spatial location of each surface point, considering distance to the respective light source and

differences between camera sensitivities and light source characteristics. See section 5.3.1 for more details. Figure 5.2 shows a sample reflectance field before correction.

5.3 Face Data Processing

Starting from the acquired geometry and reflectance data, synthetic face models have to be constructed. Before an actual skin reflectance model can be fitted, extensive preprocessing of the input data is required. As the project aims at the construction of a face database, a large number of faces has to be processed. Hence, the processing is largely automated in a data processing pipeline that allows for an unsupervised face model construction. Figure 5.3 provides an overview over the processing pipeline. The raw geometry retrieved from the 3D scanner is cleaned, parameterized, and up-sampled to obtain a highly resolved model of the facial geometry. Based on this geometry, the acquired reflectance field is re-parameterized into a *lumitexel* representation [LKG⁺01]. This requires consideration of photometric and geometric calibrations as well as the computation of camera and light source visibilities for every vertex of the geometry. Based on the *lumitexel* representation, detailed surface normals are estimated using photometric stereo. The 3D scanner resolution itself does not suffice to capture fine-scale normal variations. As an optional path in the processing pipeline, it is possible to use the normal estimates to refine vertex positions in the source geometry, and to re-iterate the processing pipeline. As a result of this processing pipeline, *lumitexels* and a refined normal map are used in a subsequent model fit.

In the remainder, we are briefly illuminating two specific aspects of this data refinement that, although often receiving little attention, are essential for a radiometrically sound acquisition: system calibration and geometry refinement.

5.3.1 System Calibration

The acquisition hardware requires careful calibration. We generally follow the principles presented in Section 2.4. Here, we are listing the specific calibration steps necessary for each component within the face scanner hardware.

Cameras The dome cameras require geometric and photometric calibration. The geometric parameters are intrinsics (focal length, principal point, aspect ratio, distortion coefficients; determined using procedures from a standard vision library [OCV]) and extrinsics. The extrinsics of all sixteen cameras are simultaneously estimated using a Euclidean bundle adjustment optimization based on a 400-frame image sequence of an LED swept through the common viewing volume. Photometric calibration requires to model vignetting and the spectral sensitivity of the individual camera sensors. For vignetting calibration we acquire images of a large sheet of white tracing paper in front of a cloudy sky, for each camera fitting a fourth-degree bivariate polynomial to the images. For spectral (color) calibration, we take images of a color checker board under diffuse illumination. It is crucial that each camera views the color checker under the exact same conditions. We color-correct all cameras using an affine color correction model, “equalizing” all cameras to have the same characteristics as one reference camera in the acquisition



Figure 5.2: Raw reflectance images of a subject acquired with all 16 cameras and 14 (of 150) lighting conditions. Each row shows images with different camera viewpoints and the same lighting condition. Each column displays images with different lighting conditions and the same viewpoint. The images are not yet color corrected, revealing differences between camera characteristics and light source colors.

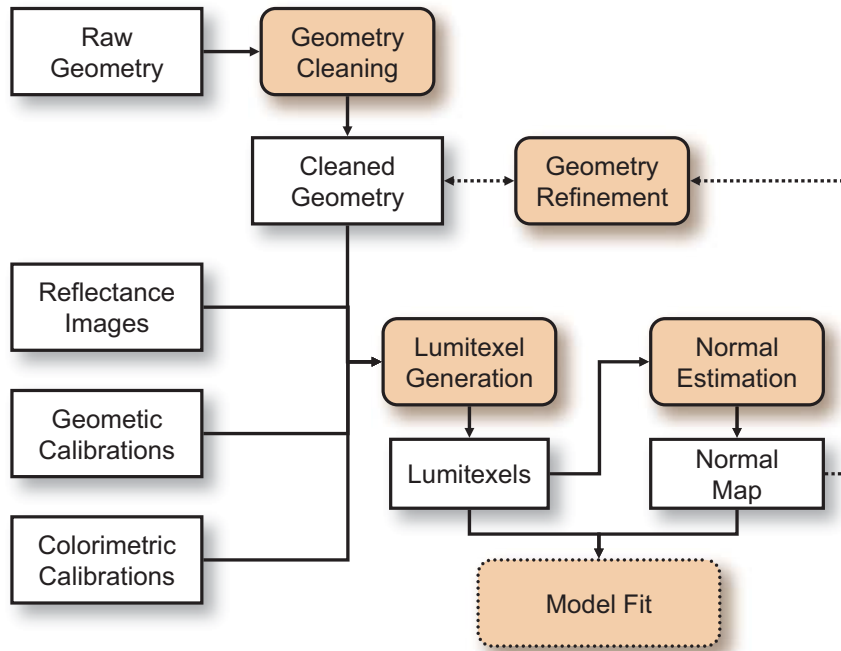


Figure 5.3: Overview over the processing pipeline. An optional feed-back loop refines the acquired geometry based on normal estimates.

dome. See [Dan92, FL00, Ami02] for a discussion of different color calibration models. A quantitative radiometric calibration is not explicitly acquired at this stage; it is implicitly performed by the light source calibration.

Light Sources The desired parameters are the light source position, color (spectral intensity) and the light-cone fall-off. Light source positions are assumed to be fixed and are taken from the CAD model of the acquisition dome. The remaining parameters are jointly measured using a Fluorilon™ reflectance target that has been calibrated to diffusely reflect 99.9% of the incident light uniformly across the spectrum. By acquiring complete reflectance fields of the target under different orientations, we collect enough data of the individual light sources to model their intensity distributions. We assume each intensity distribution to be directionally-dependent and to simply follow the $1/r^2$ law in radial direction (this holds for sufficiently large distances to the light sources). Hence, the radiometric light source calibration can be defined by a planar irradiance cross-section through its light cone. We model each light cone’s cross-section using three bivariate second-order polynomial that are independently fitted to the red, green, and blue observations of the Fluorilon™ target, respectively. Using this simple model became possible, as we are using diffuser plates in front of each light source that have experimentally been shown to produce a near-quadratic intensity fall-off.

3D Scanner The employed commercial 3D scanner comes with its own intrinsic and extrinsic calibration procedure. In addition to these proprietary calibrations, we register the 3D scanner’s coordinate system with the coordinate system of the acquisition dome using the Procrustes algorithm to match corresponding points of a 3D target that we simultaneously observe with both the 3D scanner and the dome’s cameras.

BSSRDF-Gun The subsurface scattering measurement device finally requires its own unique calibration procedure. We calibrate for the transmission from the sensing fibers to the HDR camera and for differences between the individual fibers by taking a white field measurement on a light table with an opal glass diffuser to ensure maximum uniformity of the incident light. To calibrate for spill-light within the sensor, we acquire a black image (all sensing fibers covered by a black rubber sheet) with the feeding fiber turned on. The radiometric calibration finally calibrates for the feeding fiber’s irradiance and the spectral sensitivity of the camera sensor by measuring the diffusion kernel of a sample of skim milk using our device. Skim milk has previously been measured to high accuracy [JMLH01b] and alters only minimally between different vendors. We determine the HDR camera’s color calibration so that the measured kernel meets the previously measured scattering profile of the milk, in other words, we are using skim milk as a secondary standard.

5.3.2 Geometry Refinement

The geometry from the 3D scanner suffers from imprecisions, namely from a lack of high-frequency details and from high-frequency noise. Hence, additional geometric information is obtained by estimating normals from the reflectance field using a photometric stereo implementation after [BP01]. These estimated normals, however, are not free of errors, either. Due to small calibration errors normal estimations based on different cameras tend to be biased, which can lead to an inconsistent normal field. Nehab et al. [NRDR05] analyze this general problem and propose a method to combine the reliable low-frequency information of a 3D range scan with the high-frequency content of photometric normals, thereby removing the low-frequency bias from the normal orientations.

The method by Nehab et al. is a two-step procedure. In a first stage, high-frequency details of the normal map are combined with the bias-free low-frequency normals of the smoothed geometry. Afterward, geometric detail is added to the smoothed geometry that is consistent with the new normals. We use this method to improve both normal maps and face geometry, adapting this method to our datasets. Figure 5.4 shows a schematic overview of the geometry refinement.

5.4 Reflectance Model Fit

We developed a reflectance model that is capable of describing skin reflectance based on the acquired data. We aimed at a representation that contains intuitive parameters that can be manually changed and that roughly reflects the physiology of skin. These design goals leave a wide range of potential models that are more or less strictly aligned with physical reality [INN05]. A key decision in our project, however, has been that the model should not provide more degrees of freedom than can be defined by the acquired data. A model that contains too many parameters would require ad-hoc definitions of quantities that have not been measured, which falsifies the goal of a measurement-based skin model. On the other hand, physical quantities within the model should be well-exposed to further an elementary understanding of skin reflectance.

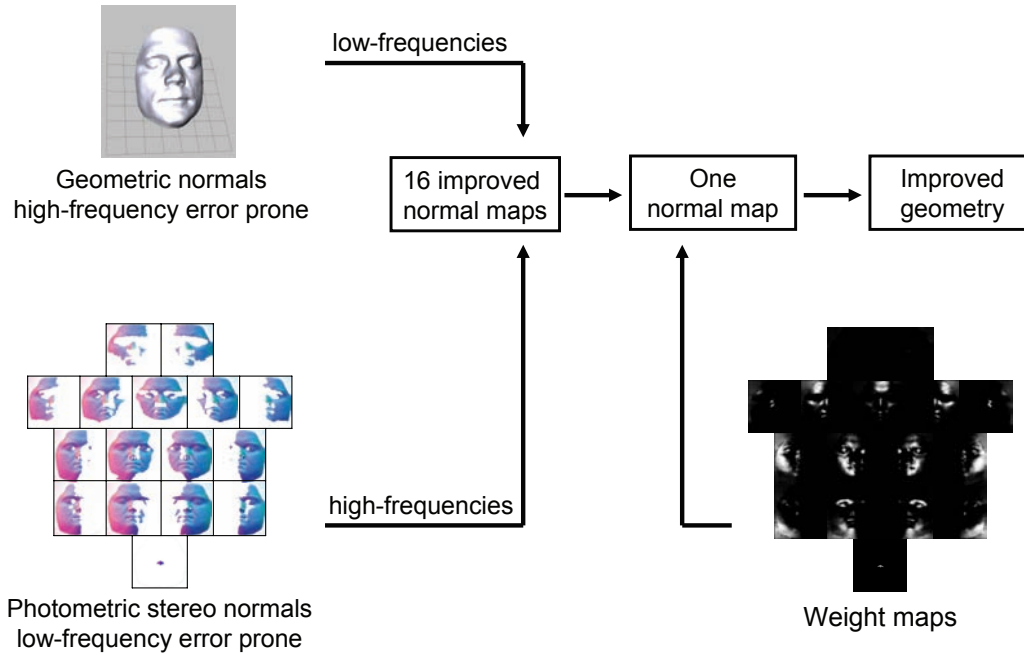


Figure 5.4: *Geometry refinement work-flow. Using only low-error frequency bands of estimated normals and 3D geometry allows us to compute an improved high-resolution geometry. We estimate for each camera (16 in total) a normal map, remove their bias to create 16 improved normal-maps and combine them to a single normal map using an appropriate weighting. Based on the improved normal map an improved geometry is computed.*

Skin Reflectance Model As previously described, we use a simple two-layer model of skin which reflects the coarse physiological classification of skin into epidermis and dermis. On top of the epidermis, we assume a thin oil layer. This decomposition is simple enough to be accessible to modeling from measured data. Our reflectance model separates reflectance into specular and diffuse reflectance. Diffuse subsurface scattering is modeled using the dipole diffusion approximation [JMLH01b], see Section 2.3. This model assumes isotropic scattering (which, as we will show, holds for facial skin) and a homogeneous material. In order to achieve spatial variation in the diffuse term, we additionally modulate the diffuse term by a modulation texture [GLL⁺04b]. For the second (specular) term, we employ the widely-used Torrance-Sparrow BRDF model, which in our experiments has proven to be best suited to model skin gloss at the oily skin/air interface.

Model Fit During the model fit, we obtain model parameters for each surface point on the face. We start by estimating the *diffuse albedo* R_d in each point by using an extension of a standard specular/diffuse separation often used in the Vision community. Our extension trades the Lambertian model of diffuse reflectance for a diffuse term that considers the transmissive Fresnel terms $F_t(\eta, \cdot)$, cf. Equation (2.3). (These terms effect an attenuation for oblique viewing and lighting angles and model the fact that due to Fresnel reflection, less light is able to pass the skin/air interface under oblique angles.) The same diffuse term occurs in the BRDF approximation of the dipole diffusion model [JMLH01b]. We parameterize our subsurface scattering BSSRDF to meet the average diffuse albedo

$\overline{R_d}$, while maintaining the translucency obtained using the “BSSRDF Gun”. The modulation texture is set to scale the underlying BSSRDF to meet R_d in every surface point. Finally, the parameters of the Torrance-Sparrow surface BRDF are obtained by fitting this BRDF to the residual reflectance samples after subtraction of the diffuse reflectance in each point. See [Wey06] for an in-depth description of the fitting procedure.

Reconstructions After the model fit, model parameters for each surface point on the face are known. We encode these parameters in floating-point textures over a common uv -parameterization of the facial geometry. Using custom shaders that implement our skin model within a Monte-Carlo raytracer, this enables us to render photo-realistic images under arbitrary illumination from arbitrary vantage points. In particular, as our reflectance field has been acquired within a fully calibrated system, it becomes possible to replay the exact illumination and viewing conditions of each reflectance field image. This allows to directly evaluate our skin model in a side-by-side comparison with photographs from within the acquisition dome. Figure 5.5 shows such a comparison. It compares single input reflectance images with synthetic images for different faces and different viewpoints. Note, that a slight blur in the fitted model reflects the fact that each surface point’s parameters are fitted to 2,400 input images simultaneously, which in the presence of noise and measurement imprecisions makes it impossible to retrieve the exact input image from the model.

5.5 Reflectance Analysis

A central goal of the project has been to analyze skin reflectance properties over a large group of subjects in order to obtain general insight in the variability of skin appearance across individuals.

The Face Database To this end, we scanned 149 subjects that were classified by skin type, gender, age, and other traits, and in each scan we manually classified facial regions, such as forehead, nose, chin, and others. This allows for a statistical analysis of characteristic variations in skin reflectance for different populations and across different facial regions. The skin type is classified according to the Fitzpatrick system [Fit88]. Table 5.1 explains the Fitzpatrick system and shows the distribution of our measurements. Figure 5.2 shows our face region classification.

Translucency Variance In an initial experiment on variation of skin translucency, we validated the isotropic assumption of our model’s subsurface scattering term. By taking subsurface scattering measurements under 16 different orientations of the sensor, we measured the degree of anisotropy of the diffusion kernel. It turns out that light diffusion is not always isotropic; abdominal skin, for instance, shows a well-expressed scattering anisotropy. All facial measurements, however, show near-isotropic diffusion kernels, which justifies an isotropic BSSRDF model for facial skin. We also analyzed spatial translucency variance measuring 52 points in two subjects’ faces. As far as accessible by our sensor, all facial regions showed a very uniform translucency. Ultimately, we decided to model skin translucency to be constant across each face and reduce the number of measurements per subject to three. Analyzing translucency variations across multiple subjects revealed a subtle difference between male and female subjects (females having a slightly



Figure 5.5: Comparison of real photographs (first and third row) to our model (second and last row). All photographs were cropped according to the 3D model to remove distracting features.

Skin Type	Skin Color	Sun Exposure Reaction	Subjects (M/F)
I	Very white	Always burn	–
II	White	Usually burn	8 / 6
III	White to olive	Sometimes burn	49 / 18
IV	Brown	Rarely burn	40 / 8
V	Dark brown	Very rarely burn	13 / 2
VI	Black	Never burn	4 / 1

Table 5.1: The Fitzpatrick skin type system and the number of subjects per skin type.



Table 5.2: 10 face regions: (1) mustache, (2) forehead, (3) eyebrows, (4) eyelids, (5) eyelashes, (6) nose, (7) cheek, (8) lips, (9) chin, and (10) beard region.

more translucent skin), while other traits did not correlate statistically significantly with translucency.

Spatial BRDF Variance A more significant variability, however, could be found in the surface reflectance. The respective Torrance-Sparrow parameters ρ_s and m vary significantly in dependence of the facial region. For each facial region, we perform principal component analysis (PCA) of these parameters, considering the BRDF fits of all subjects in the database. It turns out that the parameters do not only vary between facial regions, but depending on the region, there is also a higher variability across subjects. Exemplary observations are: the nose is quite specular, while the chin is rather non-specular; the BRDF variance on the forehead is extremely low and almost uniform across subjects, while reflectance above the lip varies highly between subjects. This shows clearly that spatial BRDF variance is an important aspect of facial appearance.

Skin Trait Variance In order to detect correlations between reflectance parameters and the traits associated with each subject, we perform canonical correlation analysis (CCA). It turns out that the surface BRDF parameters correlate the most with skin type and gender. Less surprising, albedo is highly correlated with skin type. Apart from that, there is no significant correlation of albedo with other traits.

5.6 Appearance Transfer

It is now possible to use the parameter observations within the face database to derive intuitive user controls to alter facial appearance. While the analysis performed in the previous section can generally be used as a guideline when changing skin parameters, it is desirable to have higher-level controls. In the texture synthesis procedure by Heeger and Bergen [HB95] we found a powerful tool to transfer appearance parameters between subjects and to seamlessly blend between them [MZD05]. The texture synthesis is applicable to all model parameter types and can be used to add freckles, moles, gloss variations, and

other individual effects. With the face database at hand, this provides a general appearance editing framework. Figure 5.6 shows examples where this method has been applied




Figure 5.6: *Appearance editing, altering the diffuse reflectance. From left to right: Real photograph; rendering; making the face sun-burnt; adding hair follicles in the beard area; making the skin type darker.*

to a face's diffuse reflectance, where changes are most visible. Altering other model parameters works analogously, although the effect appears to more subtle in renderings.

5.7 Conclusion

This part of the class presented a project that developed a simple and practical skin model. An important feature of this model is that all its parameters can be robustly estimated from measurements. This reduces the large amount of measured data to a manageable size, facilitates editing, and enables face appearance changes. Images from our model come close to reproducing photographs of real faces for arbitrary illumination and pose. We fit our model to data of a large and diverse group of people. The analysis of this data provides insight into the variance of face reflectance parameters based on age, gender, or skin type. The database with all statistics is available to the research community for face synthesis and analysis [Mit].

In general, there are potential extensions to our model. For example, it would be interesting to measure wavelength-dependent absorption and scattering parameters. It would also be interesting to compare the results from the diffuse dipole approximation with a full Monte Carlo subsurface scattering simulation. Other important areas that require a different modeling approach are facial hair (eyebrows, eyelashes, mustaches, and beards), hair, ears, eyes, and teeth. Very fine facial hair also leads to asperity scattering and the important "velvet" look of skin near grazing angles [KP03]. Our model does not take this into account. We captured face reflectance on static, neutral faces. Equally important are expressions and face performances. For example, it is well known that the blood flow in skin changes based on facial expressions. Our setup has the advantage that such reflectance changes could be captured in real-time.

List of Figures

1.1	Reflectance and subsurface scattering	1
1.2	Point light source emitting light in a direction	3
1.3	Light emitted from a surface, in a specific direction	4
1.4	Goniometric view of slices of a BRDF	4
1.5	Anisotropic reflection	6
1.6	Synthesized images with and without subsurface scattering	8
1.7	Taxonomy of scattering and reflectance functions	9
3.1	Reflectance measurements of a dove greeting card	23
3.2	Inverse Shade Tree framework overview diagram	26
4.1	Acquisition setup of Jensen et al. [JMLH01a]	34
4.2	Acquisition setup of Goesele et al. [GLL ⁺ 04a]	35
4.3	Dual Photography	37
4.4	8D reflectance field	39
4.5	Symmetric Photography	40
5.1	Physiology of skin	42
5.2	Raw reflectance images	45
5.3	Overview over the processing pipeline	46
5.4	Geometry refinement work-flow	48
5.5	Comparison of real photographs to our model	50
5.6	Appearance editing	52

Bibliography

- [Ami02] Isaac Amidror. Scattered data interpolation methods for electronic imaging systems: a survey. *Journal of Electronic Imaging*, 11(2):157–176, April 2002.
- [BG01] S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 107–116, 2001.
- [BL03] George Borshukov and J. P. Lewis. Realistic human face rendering for "the matrix reloaded". In *ACM SIGGRAPH 2003 Conference Abstracts and Applications (Sketch)*, page 1, New York, NY, USA, 2003. ACM Press.
- [BNJ03] D. Blei, A. Ng, and M. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, January 2003.
- [BP01] Svetlana Barsky and Maria Petrou. Colour photometric stereo: Simultaneous reconstruction of local gradient and colour of rough textured surfaces. In *ICCV*, pages 600–605, 2001.
- [BV99] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [CCT] Camera calibration toolbox for matlab. http://www.vision.caltech.edu/bouguetj/calib_doc/.
- [CD02] Oana G. Cula and Kristin J. Dana. Image-based skin analysis. In *Texture 2002, The Second International Workshop on Texture Analysis and Synthesis*, pages 35–42, Copenhagen, Denmark, June 2002.
- [CJAMJ05] Petrik Clarberg, Wojciech Jarosz, Tomas Akenine-Möller, and Henrik Wann Jensen. Wavelet importance sampling: Efficiently evaluating products of complex functions. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)*, 24(3), 2005.
- [CLFS07] T. Chen, H.P.A. Lensch, C. Fuchs, and H.P. Seidel. Polarization and Phase-Shifting for 3D Scanning of Translucent Objects. *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8, 2007.
- [Coo84] R. L. Cook. Shade trees. In *Computer Graphics*, volume 18 of *SIGGRAPH 84 Proceedings*, pages 223–231, July 1984.
- [CT82] Robert L. Cook and Kenneth E. Torrance. A reflection model for computer graphics. *ACM Transactions On Graphics*, 1(1):7–24, January 1982.

- [Dan92] James L. Dannemiller. Spectral reflectance of natural objects: how many basis functions are necessary? *Journal of the Optical Society of America A*, 9(4):507–515, April 1992.
- [DHT⁺00a] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the Reflectance Field of a Human Face. In *Proc. ACM SIGGRAPH*, pages 145–156, July 2000. ISBN 1-58113-208-5.
- [DHT⁺00b] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Computer Graphics, SIGGRAPH 2000 Proceedings*, pages 145–156, July 2000.
- [DJ05a] Craig Donner and Henrik Wann Jensen. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics*, 24(3):1032–1039, August 2005.
- [DJ05b] Craig Donner and Henrik Wann Jensen. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)*, 24(3):1032–1039, 2005.
- [DJ06] Craig Donner and Henrik W. Jensen. A spectral bssrdf for shading human skin. In *Proceedings of the 16th Eurographics Symposium on Rendering*, pages 409–417. Eurographics Association, June 2006.
- [DM97] Paul Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Computer Graphics, SIGGRAPH 97 Proceedings*, pages 369–378, Los Angeles, CA, 1997.
- [DvGNK97] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 151–157, 1997.
- [DvGNK99] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real world surfaces. *ACM Transactions on Graphics*, 1(18):1–34, 1999.
- [DW04] K.J. Dana and J. Wang. Device for convenient measurement of spatially varying bidirectional reflectance. *Journal of the Optical Society of America A*, 21(1):1–12, 2004.
- [ECJ⁺06] Per Einarsson, Charles-Felix Chabert, Andrew Jones, Wan-Chun Ma, Bruce Lamond, Tim Hawkins, Mark Bolas, Sebastian Sylwan, and Paul Debevec. Relighting human locomotion with flowed reflectance fields. In *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, pages 183–194, June 2006.
- [FBLS07] Martin Fuchs, Volker Blanz, Hendrik P.A. Lensch, and Hans-Peter Seidel. Adaptive sampling of reflectance fields. *ACM Transactions on Graphics*, 26(2), June 2007. Article 10.
- [Fee04] Catherine Feeny. Servers, spydercams and ‘Spider-Man 2’. In *VFX Pro*, http://www.uemedia.net/CPC/vfxpro/printer_9050.shtml, July 2004.

- [FFP05] L. Fei-Fei and P. Perona. A bayesian hierarchical model for learning natural scene categories. In *IEEE Computer Vision and Pattern Recognition*, 2005.
- [Fit88] T. B. Fitzpatrick. The validity and practicality of sun-reactive skin types I through VI. *Arch. Dermatology*, 124:869–871, 1988.
- [FL00] Brian V. Funt and Benjamin C. Lewis. Diagonal versus affine transformations for color correction. *Journal of the Optical Society of America A*, 17(11):2108–2112, November 2000.
- [FLS05] Martin Fuchs, Hendrik P. A. Lensch, and Hans-Peter Seidel. Reflectance from images: A model-based approach for human faces. *IEEE Transactions on Visualization and Computer Graphics*, 11(3):296–305, 2005. Member-Volker Blanz.
- [GaHSSR04] Andrew Gelman, John B. Carlin ad Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman and Hall, 2nd edition, 2004.
- [GCHS03] Dan B. Goldman, Brian Curless, Aaron Hertzmann, and Steven M. Seitz. Shape and spatially-varying brdfs from photometric stereo. Technical Report 04-05-03, University of Washington, 2003.
- [GCHS05a] Dan B Goldman, Brian Curless, Aaron Hertzmann, and Steven M. Seitz. Shape and spatially-varying BRDFs from photometric stereo. In *IEEE International Conference on Computer Vision*, 2005.
- [GCHS05b] D.B. Goldman, B. Curless, A. Hertzmann, and S.M. Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE International Conference on Computer Vision*, 2005.
- [Geo03] Athinodoros S. Georghiadis. Recovering 3-D shape and reflectance from a small number of photographs. In *Proceedings of the 14th Eurographics workshop on Rendering*, pages 230–240, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.
- [GGSC96] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. In *Computer Graphics, SIGGRAPH 96 Proceedings*, pages 43–54, New Orleans, LS, August 1996.
- [GLL⁺04a] Michael Goesele, Hendrik P. A. Lensch, Jochen Lang, Christian Fuchs, and Hans-Peter Seidel. DISCO – Acquisition of Translucent Objects. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004)*, 23(3), 2004.
- [GLL⁺04b] Michael Goesele, Hendrik P. A. Lensch, Jochen Lang, Christian Fuchs, and Hans-Peter Seidel. DISCO—Acquisition of translucent objects. *ACM Transactions on Graphics (SIGGRAPH 2004)*, 24(3):835–844, 2004.
- [GMSW84] P. Gill, W. Murray, M. Saunders, and M. Wright. Procedures for optimization problems with a mixture of bounds and general linear constraints. In *ACM Transactions on Mathematical Software*, 1984.

- [GTLL06] Gaurav Garg, Eino-Ville Talvala, Marc Levoy, and Hendrik P. A. Lensch. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, pages 251–262, June 2006.
- [GTR⁺06] Jinwei Gu, Chien-I Tu, Ravi Ramamoorthi, Peter Belhumeur, Wojciech Matusik, and Shree Nayar. Time-varying surface appearance: Acquisition, modeling and rendering. *ACM Transactions on Graphics (SIGGRAPH 2006)*, 25(3), 2006.
- [Hac99] W. Hackbusch. A Sparse Matrix Arithmetic based on \mathcal{H} -Matrices. Part I: Introduction to \mathcal{H} -matrices. *Computing*, 62(2):89–108, 1999.
- [HB95] David J. Heeger and James R. Bergen. Pyramid-based texture analysis/synthesis. In *Proceedings of SIGGRAPH 95, Computer Graphics Proceedings, Annual Conference Series*, pages 229–238, August 1995.
- [Hea89] G. Healey. Using color for geometry-insensitive segmentation. *J. Optical Society of America A*, 6(6):920–937, 1989.
- [HED05] Tim Hawkins, Per Einarsson, and Paul Debevec. A dual light stage. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, pages 91–98, June 2005.
- [Hér03] Christophe Héry. Face cloning at ILM. In *SIGGRAPH 2003 Course “Digital Face Cloning”*, 2003.
- [HK93] Pat Hanrahan and Wolfgang Krueger. Reflection from layered surfaces due to subsurface scattering. In *Computer Graphics, SIGGRAPH 93 Proceedings*, pages 165–174, Anaheim, CA, August 1993.
- [Hof99] Thomas Hofmann. Probabilistic latent semantic analysis. In *Proceedings of Uncertainty in Artificial Intelligence*, 1999.
- [HP03] Jefferson Y. Han and Ken Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics*, 22(3):741–748, July 2003.
- [HW79] J. A. Hartigan and M. A. Wong. A k-means clustering algorithm. *Applied Statistics*, 28:100–108, 1979.
- [HWT⁺04] Tim Hawkins, Andreas Wenger, Chris Tchou, Andrew Gardner, Fredrik Göransson, and Paul Debevec. Animatable facial reflectance fields. In *Rendering Techniques '04 (Proceedings of the Second Eurographics Symposium on Rendering)*, pages 309–320, Norrköping, Sweden, June 2004.
- [INN05] Takanori Igarashi, Ko Nishino, and Shree K. Nayar. The appearance of human skin. Technical report, Department of Computer Science, Columbia University CUCS-024-05, June 2005.
- [JMLH01a] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001, Computer Graphics Proceedings, Annual Conference Series*, pages 511–518, August 2001.

- [JMLH01b] Henrik Wann Jensen, Steven R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Computer Graphics, SIGGRAPH 2001 Proceedings*, pages 511–518, Los Angeles, CA, August 2001.
- [Jor99] Michael Jordan, editor. *Learning in Graphical Models*. MIT Press, 1999.
- [KP03] Jan Koenderink and Sylvia Pont. The secret of velvety skin. *Machine Vision and Application*, 14:260–268, 2003. Special Issue on Human Modeling, Analysis, and Synthesis.
- [Law06] Jason Lawrence. *Acquisition and Representation of Material Appearance for Editing and Rendering*. PhD thesis, Department of Computer Science, Princeton University, 2006.
- [LBAD⁺06a] Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (SIGGRAPH 2006)*, 25(3), 2006.
- [LBAD⁺06b] Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (SIGGRAPH 2006)*, 25(3), 2006.
- [LBS90] H. C. Lee, E. J. Breneman, and C. P. Schulte. Modeling light reflection for computer color vision. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(4):402–409, 1990.
- [LFTG97] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 117–126, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [LFTW06] H. Li, S.C. Foo, K.E. Torrance, and S.H. Westin. Automated three-axis gonioreflectometer for computer graphics applications. *Optical Engineering*, 45:043605, 2006.
- [LKG⁺01] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatially varying materials. In *Proceedings of the 12th Eurographics Workshop on Rendering*, pages 104–115, June 2001.
- [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 Transactions on Graphics*, 22(2):234–257, 2003.
- [LKK98] R. Lu, J. Koenderink, and A. Kappers. Optical properties (Bidirectional Reflectance Distribution Functions) of velvet. *Applied Optics*, 37(25):5974–5984, September 1998.

- [LRR04] Jason Lawrence, Szymon Rusinkiewicz, and Ravi Ramamoorthi. Efficient BRDF importance sampling using a factored representation. *ACM Transactions on Graphics (SIGGRAPH 2004)*, 23(3):496–505, August 2004.
- [LS99] D. Lee and H. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.
- [MBK05] Gero Müller, Gerhard H. Bendels, and Reinhard Klein. Rapid synchronous acquisition of geometry and appearance of cultural heritage artefacts. In *VAST 2005: 6th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*, pages 13–20, November 2005.
- [McA02a] David McAllister. *A Generalized Surface Appearance Representation for Computer Graphics*. Ph.d. thesis, University of North Carolina (UNC), Chapel Hill, NC, 2002.
- [McA02b] David K. McAllister. *A Generalized Surface Appearance Representation for Computer Graphics*. PhD thesis, University of North Carolina at Chapel Hill, NC, 2002.
- [Mer84] S. Mersch. Polarized lighting for machine vision applications. In *Proc. RI/SME Third Annual Applied Machine Vision Conf.*, pages 40–54. Schaumburg, IL, Feb. 1984.
- [Mit] Mitsubishi Electric Research Laboratories (MERL), ETH Zurich. The MERL/ETH skin reflectance database. <http://www.merl.com/facescanning/>.
- [MLH02] D.K. McAllister, A. Lastra, and W. Heidrich. Efficient rendering of spatial bi-directional reflectance distribution functions. *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, pages 79–88, 2002.
- [MMS⁺05] G. Müller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein. Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics Forum*, 24(1), 2005.
- [MN99] T. Mitsunaga and S.K. Nayar. Radiometric Self Calibration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 374–380, Jun 1999.
- [MPBM03a] W. Matusik, H. Pfister, M. Brand, and L. McMillan. Efficient isotropic BRDF measurement. *Proceedings of the 14th Eurographics workshop on Rendering*, pages 241–247, 2003.
- [MPBM03b] Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Transactions on Graphics (SIGGRAPH 2003)*, 22(3):759–770, July 2003.
- [MPDW03] Vincent Masselus, Pieter Peers, Philip Dutré, and Yves D. Willems. Relighting with 4d incident light fields. *ACM Transactions on Graphics*, 22(3):613–620, July 2003.

- [MSK06] Gero Müller, Ralf Sarlette, and Reinhard Klein. Data-driven local coordinate systems for image-based rendering. *Computer Graphics Forum*, 25(3):369–378, September 2006.
- [MWL⁺99] Steven R. Marschner, Steven H. Westin, Eric P. F. Lafortune, Kenneth E. Torrance, and Donald P. Greenberg. Image-based BRDF measurement including human skin. In *Proceedings of the 10th Eurographics Workshop on Rendering*, pages 139–152, Granada, Spain, June 1999.
- [MWLT00] Steven R. Marschner, Steven H. Westin, Eric P. F. Lafortune, and Kenneth E. Torrance. Image-based measurement of the Bidirectional Reflectance Distribution Function. *Applied Optics*, 39(16):2592–2600, June 2000.
- [MZD05] Wojciech Matusik, Matthias Zwicker, and Frédo Durand. Texture design using a simplicial complex of morphable textures. *ACM Transactions on Graphics*, 24(3):787–794, 2005.
- [MZKB05] Satya P. Mallick, Todd E. Zickler, David J. Kriegman, and Peter N. Belhumeur. Beyond lambert: Reconstructing specular surfaces using color. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) — Volume 2*, pages 619–626, Washington, DC, USA, 2005. IEEE Computer Society.
- [NDM05] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of brdf models. In *Proceedings of the 15th Eurographics Symposium on Rendering*, pages 117–226. Eurographics Association, 2005.
- [NRDR05] Diego Nehab, Szymon Rusinkiewicz, James Davis, and Ravi Ramamoorthi. Efficiently combining positions and normals for precise 3d geometry. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)*, 24(3):536–543, 2005.
- [NRH⁺77] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. Monograph 161, National Bureau of Standards (US), October 1977.
- [OCV] Open source computer vision library. <http://opencvlibrary.sourceforge.net/>.
- [PHL⁺98] Frédéric Pighin, Jamie Hecker, Dani Lischinski, Richard Szeliski, and David H. Salesin. Synthesizing realistic facial expressions from photographs. In *Computer Graphics*, volume 32 of *SIGGRAPH 98 Proceedings*, pages 75–84, 1998.
- [PvBM⁺06a] Pieter Peers, Karl vom Berge, Wojciech Matusik, Ravi Ramamoorthi, Jason Lawrence, Szymon Rusinkiewicz, and Philip Dutré. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics (SIGGRAPH 2006)*, 25(3), 2006.
- [PvBM⁺06b] Pieter Peers, Karl vom Berge, Wojciech Matusik, Ravi Ramamoorthi, Jason Lawrence, Szymon Rusinkiewicz, and Philip Dutré. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics (SIGGRAPH 2006)*, 25(3), 2006.

- [PvBM⁺06c] Pieter Peers, Karl vom Berge, Wojciech Matusik, Ravi Ramamoorthi, Jason Lawrence, Szymon Rusinkiewicz, and Philip Dutré. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics*, 25(3):746–753, July 2006.
- [Rus98] Szymon Rusinkiewicz. A new change of variables for efficient BRDF representation. In G. Drettakis and N. Max, editors, *Rendering Techniques '98 (Proceedings of Eurographics Rendering Workshop '98)*, pages 11–22, New York, NY, 1998. Springer Wien.
- [Sag04] Mark Sagar. Reflectance field rendering of human faces for "Spider-Man 2". In *ACM SIGGRAPH 2004 Sketches*, August 2004.
- [SCG⁺05] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. Dual photography. *ACM Transactions on Graphics*, 24(3):745–755, August 2005.
- [Sha85] S. Shafer. Using color to separate reflection components. *COLOR research and applications*, 10(4):210–218, 1985.
- [SNB03] Y.Y. Schechner, S.K. Nayar, and P.N. Belhumeur. A Theory of Multiplexed Illumination. In *IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 808–815, Oct 2003.
- [Sta] A. Standard. E1392-90, Standard practice for angle resolved optical scatter measurements on specular or diffuse surfaces. *American Society for Testing and Materials*.
- [Sta01] Jos Stam. An illumination model for a skin layer bounded by rough surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 39–52, Vienna, Austria, London, UK, June 2001. Springer-Verlag.
- [SWI97] Yoichi Sato, Mark D. Wheeler, and Katsushi Ikeuchi. Object shape and reflectance modeling from observation. In *Computer Graphics, SIGGRAPH 97 Proceedings*, pages 379–387, 1997.
- [TOS⁺03] Norimichi Tsumura, Nobutoshi Ojima, Kayoko Sato, Mitsuhiro Shiraishi, Hideto Shimizu, Hirohide Nabeshima, Syuuichi Akazaki, Kimihiko Hori, and Yoichi Miyake. Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. *ACM Transactions on Graphics (Proc. SIGGRAPH 2003)*, 22(3):770–779, 2003.
- [TS67] Kenneth E. Torrance and Ephraim M. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, September 1967.
- [TW89] S. Tominaga and B.A. Wandell. Standard surface-reflectance model and illuminant estimation. *J. Optical Society of America A*, 6(4):576–584, April 1989.
- [TWL⁺05] Xin Tong, Jiaping Wang, Stephen Lin, Baining Guo, and Heung-Yeung Shum. Modeling and rendering of quasi-homogeneous materials. *ACM Transactions on Graphics (SIGGRAPH 2005)*, 24(3):1054–1061, 2005.

- [TZL⁺02] Xin Tong, Jingdan Zhang, Ligang Liu, Xi Wang, Baining Guo, and Heung-Yeung Shum. Synthesis of bidirectional texture functions on arbitrary surfaces. In *Proceedings of SIGGRAPH*, pages 665–672. ACM Press, 2002.
- [VT04] M. Alex Vasilescu and Demetri Terzopoulos. TensorTextures: Multilinear image-based rendering. *ACM Transactions on Graphics (SIGGRAPH 2004)*, 23(3), 2004.
- [War92] G. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(Annual Conference Series):265–273, 1992.
- [Wey06] Tim Weyrich. *Acquisition of Human Faces using a Measurement-Based Skin Reflectance Model*. PhD thesis, Department of Computer Science, ETH Zurich, 2006.
- [WGT⁺05] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics*, 24(3):756–764, August 2005.
- [Wil05] Lance Williams. Case study: The gemini man. In *SIGGRAPH 2005 Course 'Digital Face Cloning'*, 2005.
- [WMP⁺06a] Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, and Markus Gross. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics*, 25(3):1013–1024, July 2006.
- [WMP⁺06b] Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, and Markus Gross. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics (Proc. SIGGRAPH 2006)*, pages 1013–1024, July 2006.
- [WSB⁺98] D.R. White, P. Saunders, S.J. Bonsey, J. van de Ven, and H. Edgar. Reflectometer for Measuring the Bidirectional Reflectance of Rough Surfaces. *Appl. Opt.*, 37:3450–3454, 1998.
- [WTL⁺06] Jiaping Wang, Xin Tong, Stephen Lin, Hujun Bao, Baining Guo, and Heung-Yeung Shum. Appearance manifolds for modeling time-variant appearance of materials. *ACM Transactions on Graphics (SIGGRAPH 2006)*, 25(3), 2006.
- [YDMH99] Yizhou Yu, Paul Debevec, Jitendra Malik, and Tim Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Computer Graphics, SIGGRAPH 99 Proceedings*, pages 215–224, Los Angeles, CA, August 1999.
- [YNBK07] Shuntaro Yamazaki, Srinivasa G. Narasimhan, Simon Baker, and Takeo Kanade. Coplanar shadowgrams for acquiring visual hulls of intricate objects. In *Proc. IEEE International Conference of Computer Vision (ICCV)*, 2007.

- [ZBK02] Todd Zickler, Peter N. Belhumeur, and David J. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *ECCV '02: Proceedings of the 7th European Conference on Computer Vision-Part III*, pages 869–884, London, UK, 2002. Springer-Verlag.
- [ZERB05] Todd Zickler, Sebastian Enrique, Ravi Ramamoorthi, and Peter Belhumeur. Reflectance sharing: Image-based rendering from a sparse set of images. In *Proceedings of the Eurographics Symposium on Rendering*, pages 253–264. Eurographics Association, 2005.
- [ZREB06] T. Zickler, R. Ramamoorthi, S. Enrique, and P.N. Belhumeur. Reflectance Sharing: Predicting Appearance from a Sparse Set of Images of a Known Shape. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 28(8):1287–1302, 2006.