Measurement I - BRDFs

Lecture #2: Wednesday, 4 October 2000
Lecturer: Szymon Rusinkiewicz, Steve Marschner
Scribe: Ben Mowery
Reviewer: Ravi Ramamoorthi, Szymon Rusinkiewicz

1 Overview

Consider the possible parameters for a photon striking a surface and the resulting photon leaving that surface:

$$(x, y, t, \theta, \phi, \lambda)_{in} \rightarrow (x, y, t, \theta, \phi, \lambda)_{out}$$

We see that in order to characterize the behavior of a surface, we need a function that gives the amount of reflected light for each possible value of these 12 parameters. Measuring such a 12-dimensional function is not really possible, so typically we make assumptions about the form of this function and only measure approximations. Figure 1 shows a taxonomy of some of these approximations.

Let us start with the original 12D space. We can eliminate a dimension by assuming that we don't have phosphorescence. That is, the light is reflected immediately after it strikes the surface, meaning that $t_{in} = t_{out}$. We can also assume that the appearance of the surface is constant over time, permitting us to eliminate both time variables completely. Next, we can assume that there is no fluorescence, meaning the incoming wavelength is equal to the outgoing wavelength. The 9D function resulting from these simplifications is called a generalized *scattering function*.

When we are dealing with typical computer graphics and computer vision applications, we can get rid of another dimension by discretizing the wavelength into bands. We typically use three bands called "red," "green," and "blue," reflecting the fact that the human eye only has three kinds of receptors.

¿From this 8-dimensional function space, we can make two different simplifications. First, we can ignoring the presence of subsurface scattering. That is, we assume that $(x, y)_{in} = (x, y)_{out}$, resulting in a six-dimensional function. We may think of this function in two ways: it is either a texture map that varies with illumination and viewing direction (thus a *Bidirectional Texture Function* or BTF) or a BRDF that varies with (x, y) position on the surface (thus a *Spatially-Varying BRDF*). During the next lecture, we will talk about some attempts to measure the BTF, as well as simplifications such as light fields, bump maps, and texture maps. Light fields and their variants (such as lumigraphs and surface light fields) result from ignoring the dependence on incident light direction. Texture maps also assume that the surface is perfectly flat and diffuse, thus resulting in a

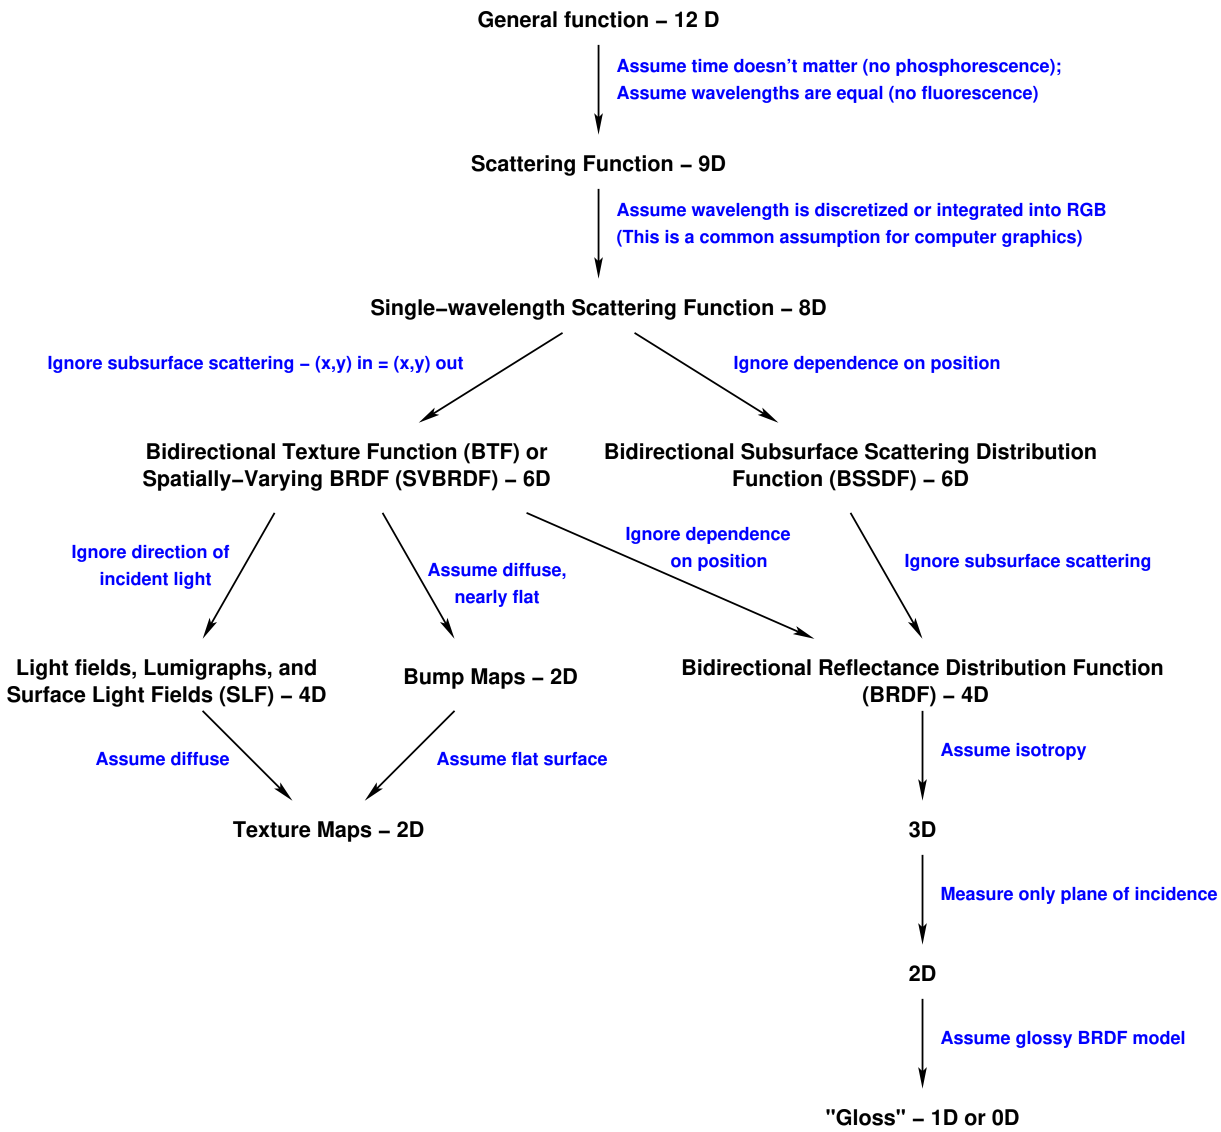

Figure 1: Taxonomy of appearance measurement.

2D function. Bump maps are slightly more flexible than simple textures, since although the surface is still assumed to be diffuse, the surface normal at each point may vary.

One may make a different simplification to the 8D scattering function, namely that the scattering behavior is constant across the surface. Thus, the spatial variation of the function depends only on $x_{out} - x_{in}$ and $y_{out} - y_{in}$. This results in a function called the *Bidirectional Subsurface Scattering Distribution Function* (BSSDF), parameterized in terms of four angles and two displacements on the surface.

For this lecture, the main function we will consider is the *Bidirectional Reflectance Distribution Function* (BRDF). Refering to the taxonomy, we see that the BRDF is a special case of either the BSSDF (assuming no subsurface scattering) or the BTF (assuming no variation across the surface). We will also consider further simplifications to the BRDF, including assuming that the BRDF is isotropic (i.e., depends only on $\phi_i - \phi_o$ rather than on both ϕ_i and ϕ_o), assuming a certain functional form for the BRDF and only measuring values in the plane of incidence, and assuming a particular BRDF model and only measuring a few "gloss" parameters.

Note that when performing any sort of appearance measurement, we must be aware of the interaction between geometry and appearance. We can, to a very large extent, trade off complexity in our geometric model for complexity in our appearance model. However, we must be careful when using the measured appearance data to use the same geometry as was assumed during measurement. For example, when representing ocean waves we may either represent them as geometry (by modeling them as a mesh or height field) or as appearance (by modeling them with a bump map, a light field, or a BTF). If we choose to model the waves as appearance, however, we must be careful when rendering not to also use the wave geometry.

2 Integrating BRDFs

The BRDF is defined as a differential quantity. In order to measure it in practice, we must integrate over a finite parameter domain, since real-world sensors cannot measure infinitesimal quantities. There are several variables over which we must integrate BRDFs:

- Incident solid angle
- Exitant solid angle
- Surface area
- Wavelength
- Time (since the BRDF is defined in terms of *power*, and we can only measure *energy*)

When doing this integration, for best results we would like to make the region of integration as small as possible, given the physical constraints of apparatus layout and detector sensitivity. The solid angle is especially important in this regard, since integrating over large regions will blur out the BRDF, especially if it contains sharp features like specular peaks. Similarly, for spatially varying BRDFs, we also want to minimize the surface area we are averaging over in order to avoid blurring of surface details. Wavelengths are typically integrated into multiple bands (RGB for most graphics applications, more and narrower bands for applications in optics, etc.). Finally, the ability to integrate over small time intervals is an important practical issue, since we would like to take many samples in rapid succession.

3 Measuring Devices

There are a couple of different ways to define the surface area over which the BRDF measurements are integrated. Three are described below:

Method #1: The light shines over a broad area and the detector integrates over a small area:

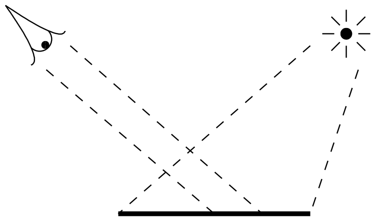

In this case the measured area is defined by the detector's field of view, and our measurements will be proportional to

$$f_r \cos\theta_i$$

Method #2: The light illuminates a small spot and the detector integrates over a wide area:

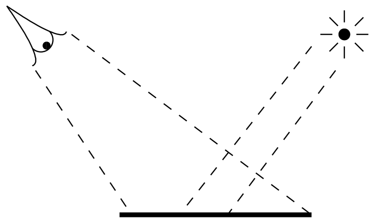

In this case the measured area is defined by the light source's beam, and our measurements will be proportional to

$$f_r \cos\theta_e$$

Method #3: The light illuminates a wide area and the detector integrates over a wide area, so that the measured area is defined by the size of the sample itself.

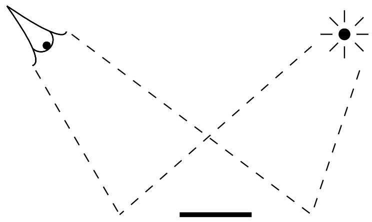

In this case the measurements are proportional to

$$f_r \cos\theta_e \ \cos\theta_i$$

Obviously we need to make sure we know which of these situations we are in when we are processing the data from a BRDF measurement.

3.1 Simple Reflectance Measurements

We now consider measuring some simple approximations to a BRDF, gradually working our way up to the full function.

One simple method for measuring color is to use a single BRDF measurement, for example with the $0/45°$ geometry:

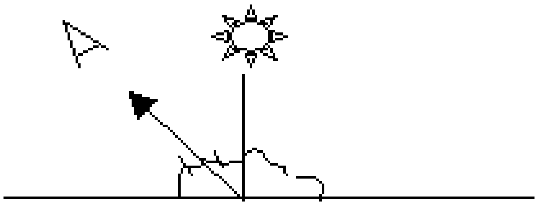

This is often used to characterize the color of paint samples, given an implicit assumption that the BRDF is Lambertian. Note that this configuration avoids specular reflection. If the solid angle of the detector changes, the amount of energy that gets into the detector changes. As long as we calibrate for this all is well.

If wish to measure specular reflectance, we can use a different setup:

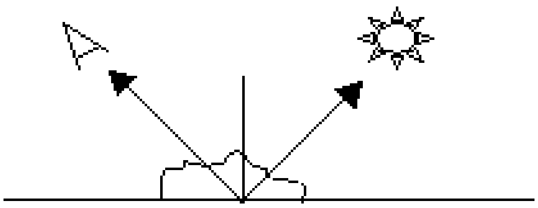

If the detector has a large enough solid angle, we will collect all the specularly reflected energy, and changing the size of this solid angle will not change the collected energy. Calibrating for a specular measurement, which is just a fraction of energy reflected, is therefore different from the diffuse case, where we are measuring BRDF values, because the diffuse measurement scales with the solid angle but the specular measurement does not. Measuring both types of reflection with the same device is therefore difficult. Also, if we are measuring the specular component of a surface that has some diffuse reflection, we want the solid angle of the detector to be the minimum size to collect all the specular energy, so as to minimize contamination from diffusely reflected light.

Measuring diffuse color with a single BRDF configuration is a little unsatisfying because, in the presence of a non-Lambertian diffuse component, color measurements are only reproducible if we carry around information about what angles were used in the measurement. A more well-defined way to measure the diffuse component is to integrate all the light that exits the sample over the whole hemisphere – this is a *directional-hemispherical* measurement. The device that is used to add up the light over the hemisphere is an integrating sphere, which is just a hollow sphere painted white on the inside. Since it's a sphere and it's highly reflective, it ends up having the same radiance everywhere inside, so we can put the detector wherever we want without changing the measurement.

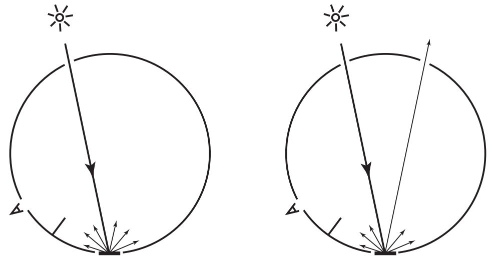

If we want to exclude the specular component from a measurement like this, we just cut a hole in the sphere where the specular beam hits and let it escape instead. Often the reciprocal setup is used: the sphere is flooded with light and the detector is positioned at a particular angle.

An even more well-defined measure of diffuse reflectance is the hemispherical-hemispherical reflectance, in which we integrate over all incident directions *and* all exitant directions. This type of reflectance is hard to measure directly, though.

None of these separate diffuse and specular measurements let us get near intuitive ideas of "shininess" and "glossiness." A method that is often used to get gloss information is to use two separate detectors: one is positioned at the correct angle to get specular reflection and the other is offset from this location to detect diffuse reflection:

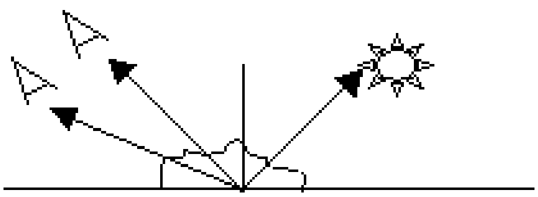

The angles for which this measurement is useful depend on the type of material being measured. For materials with broad specular peaks just the ratio of specular to diffuse reflectance is often a useful measurement. *Contrast gloss* and *luster* are measured this way. As surfaces approach a highly polished appearance, the width of the specular peak becomes the important characteristic, so the second detector is moved closer to the specular direction, and the resulting measurements have names like *haze*, *bloom*, and *distinctness of image*.

3.2 Measuring BRDFs

Now let us move to multidimensional spaces and describe how a device can be built to carry out the measurements automatically. We almost always build computer controlled devices since we must take many measurements for a single object. These are often expensive devices: many cost around $100,000.

We start with a device that can measure the BRDF only in the plane of incidence. It consists of a movable light source and sensor mounted on a circular track around the object:

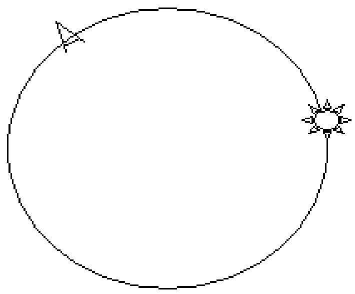

Such devices can be very accurate, since moving objects along the track can be done very precisely. Of course, since they are on the same track the sensor and source can only get so close togehter without colliding; how close they can get determines how close to retroreflection we can measure.

To obtain three degrees of freedom and measure isotropic BRDFs we must use more complicated devices. They can be aranged such that the sample remains stationary and the light and detector move around it:

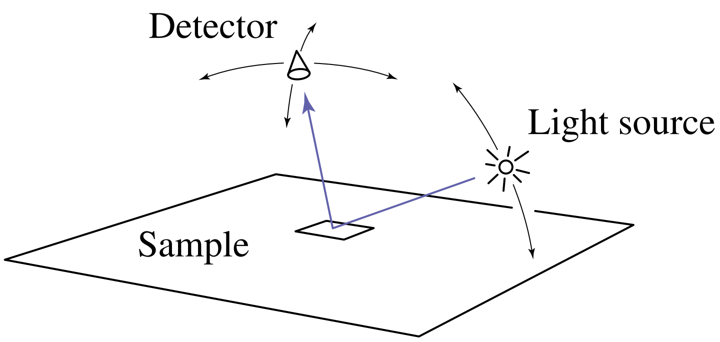

Alternatively, we may use devices that rotate the sample along two axes, with an additional degree of freedom on either the detector or light source:

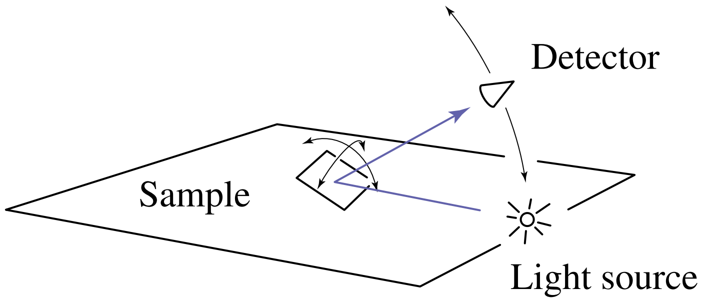

Note that either of the above configurations may be augmented to handle anisotropy by adding an additional degree of freedom on the sample, allowing it to rotate within its own plane.

4 Image-based BRDF measurement

Because the full BRDF is such a high-dimensional function and requires, in general, a large number of measurements, researchers have looked at ways of capturing a larger "slice" of the BRDF (i.e., more than just a single point sample) at a time.

An example of this kind of optimization in traditional gonioreflectometry is the use of linear CCD arrays in spectroradiometers. These devices make use of finely ruled gratings that reflect different wavelengths in different directions to spread the spectrum of light reflected from the sample across a linear CCD array. This way, the CCD elements can all make separate narrow-band measurements of the reflected light simultaneously.

Two-dimensional arrays can allow even more speedup, and for graphics it's often more useful to apply the parallelism in the geometric variables of the measurement rather than in wavelength. One design for a BRDF measurement device that uses a 2D imaging device to capture an entire hemisphere of reflected light at once was proposed by Greg Ward ("Measuring and Modeling Anisotropic Reflection," *SIGGRAPH 92*). The device uses a hemispherical mirror (actually, to be perfectly accurate, the mirror should be ellipsoidal) to reflect the light reflected by the sample to the camera, which is fitted with a fish-eye lens. The mirror is half-silvered to allow the sample to be illuminated by an external light source.

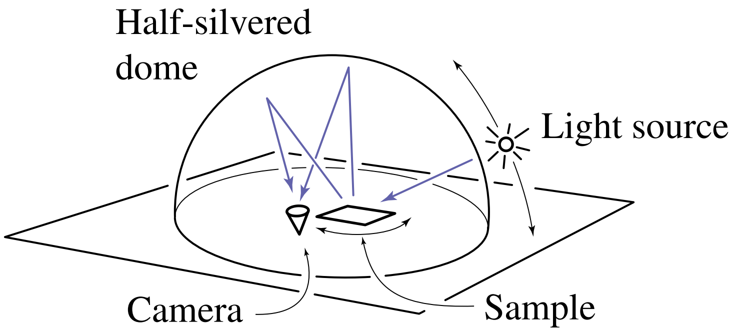

A second possibility, explored by Marschner et. al. ("Image-based bidirectional reflectance distribution measurement," *Applied Optics* 39:16, 2000) as well as in the incidence plane by Lu et. al. ("Optical properties (BRDFs) of velvet") is to use a curved object with a wide distribution of normals. The object is illuminated by a point light source, and a 2D image is taken with a camera. Each pixel in the camera's image sees a small, almost flat, bit of the object's surface. That bit of surface is illuminated from a small solid angle defined by the light source and is viewed from a small solid angle defined by the camera's aperture, so that pixel value is a BRDF measurement much like one that would be made by a normal gonioreflectometer. If the object is assumed to have known geometry and a constant BRDF, this allows us to interpret the pixels of each image as measurements that lie on two-dimensional slices of the four-dimensional BRDF space.

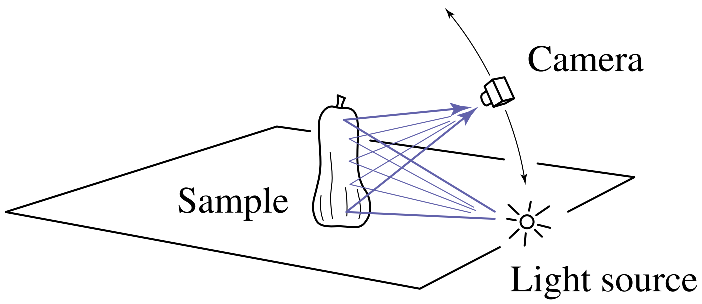

In order to acquire an entire isotropic BRDF, the angle between camera and light source is changed. This provides configurations that cover the whole domain of an isotropic BRDF. Because our ability to measure near the silhouette of the object is only limited by the resolution of the camera, this technique approach grazing angles more easily than many of the other approaches we've discussed.