UC Berkeley CS 294-13: Advanced Computer Graphics

Lecture 11 : Representation of Visual Appearance

Lecturer: Ravi Ramamoorthi

Scribe: Eric Biggs October, 12 2009

Introduction

The traditional production pipeline has been a tremendous success in the last few decades. We've seen industries born out of these methods, and some truly amazing things have been done with it. This production pipeline can be broken down into 3 core elements: geometry/animation, rendering algorithms, and material representations. When we reach the limits of simple parametric representation, we've had much success in animation and geometric models by incorporating data obtained from the real world i.e. 3D scanners and motion capture. These, along with rendering algorithms, have been a great success and are no longer the biggest barrier to creating truly photo-realistic imagery and animation. Our current weakest link is in the materials/reflectance element of the traditional production pipeline. To get a full richness of real world visuals we'll need to explore more complicated reflection seen in the real world and capture real world data.

Capturing Real World Data and the SVBRDF

The 4D BRDF (2 spherical incident coordinates + 2 spherical viewing coordinates) is well researched and implemented, but as will be seen later with the taxonomy of reflectance, is only a narrow slice of a much larger complete reflectance model. 2D texturing is a separate model that can be combined with the 4D BRDF to handle spatial variance over a surface, but these spaces are not disparate in the real world. To handle real world datasets they need to be combined into a 6-dimensional Spatially Varying Bidirectional Reflectance Distribution Function (SVBRDF).

Combining this SVBRDF with data acquisition methods (such as ~$150,000 gantry set ups, or light waving interns for those wishing to save cash with more post processing hacks) we can now begin to create more realistic materials that correspond to the true richness in visual perception of the real world. The 6 dimensional data, however is enormous – and this requires that we research more efficient mathematical models for this data to enable efficient offline global illumination, efficient real time rendering, and efficient real time editing. Research has begun on this and more is underway.

The SVBRDF is 6 dimensions, and thus already a bit to tackle, but many more dimensions will still need to be added should we conquer all the richness in the real world! This is an imposing problem. Each dimension creates an exponential growth in sampling and size of raw data! As many as 14 Dimensions are required to fully represent general light scattering.

Modeling – Efficiency and Intuitiveness

Not only should the ideal mathematical representation be compact enough to enable efficient rendering and editing, another worthwhile goal would be to find a representation that would allow intuitive editing by human beings. This is currently an open question. What is an intuitive representation? What sort of user interface would be suitable to artists and scientists? Some work has been done, but more is needed.

Approaches to addressing complexity are available. Low frequency spherical harmonics, sparse wavelet representations, non-linear factorizations and machine learning, clustered subspaces, etc etc.. Intuitive decomposition can be achieved with constrained matrix factorization and a tree structure combining these simple parts.

General Light Scattering and its Taxonomy

Before it was mentioned that 14 dimensions are needed to represent general light scattering, so what are they? In the realm of animation, <u>time</u> becomes another dimension (i.e. hydration changes, decay, rust). Subsurface scattering introduces the need for spatial coordinates for <u>both incident and viewing</u>. Furthermore, a <u>3rd spatial</u> dimension might need to be added for some reflectance models that are more complicated than typical 2D surfaces. What this amounts to is two 7 Dimension plenoptic spaces, one for incidence, the other for viewing. The 7 Dimensions of a plenoptic space are: 3 spatial, 2 directional, time, and light wavelength. This combines to a total of 14 dimensions that make up the General Scattering Function (GSF).

We can start to explore this GSF beast by building taxonomy around it. For a great number of materials or situations it's safe to assume away dimensions and/or ignore them. These simplified cousins of the GSF can then be studied to see where they may be applicable. In fact, many are already very familiar and well understood, such as the BRDF itself.

So at the top of this taxonomy is the root 14D General Scattering Function (GSF). Below that is an 8D Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF), which assumes away time dependence, wavelength dependence, and the 3rd spatial component (focusing on 2D surfaces). This function is perhaps currently a holy grail in current research and papers are beginning to emerge on it.

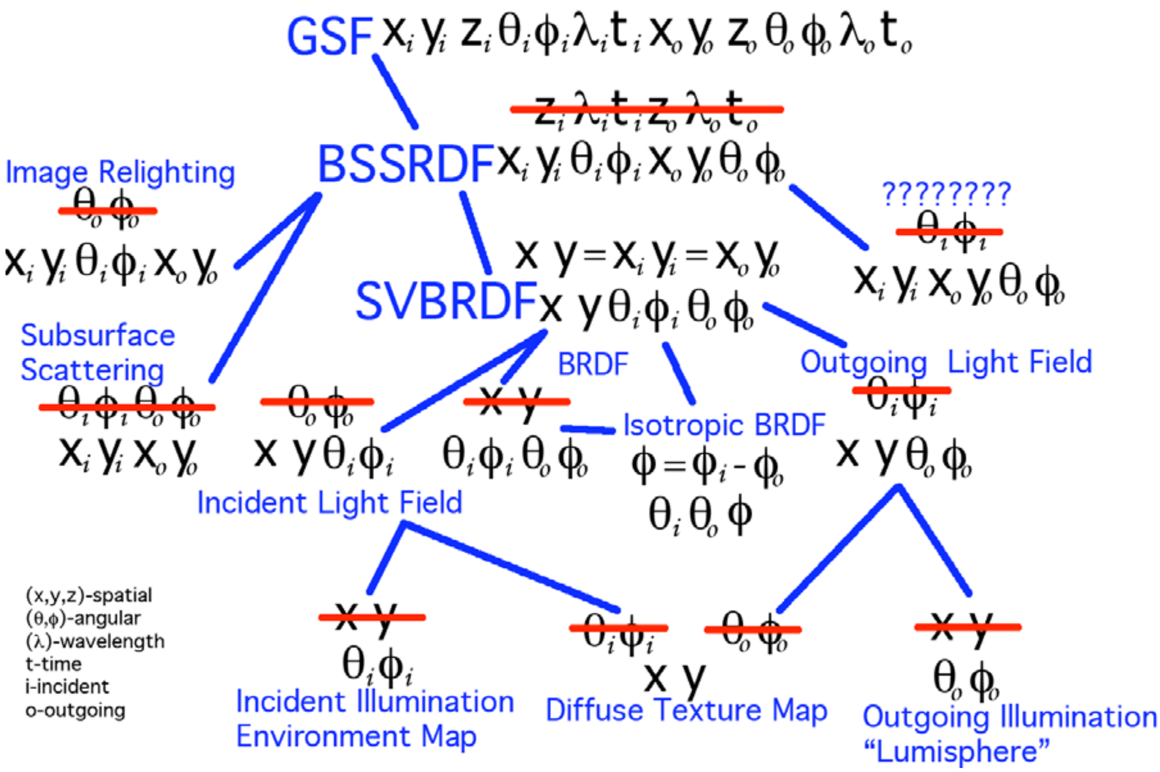

Below the BSSRDF are various familiar and unfamiliar slices of the taxonomy. Historically, slices such as the BRDF were studied first, and so have been studied for years. Every couple years or so a dimension or two would get added or studied independently under other contexts. Other slices seem to be interesting combinations but have not yet been studied for potential usefulness such as the 6 dimensional scattering function operating on the incident and outgoing 2D spatial locations and the additional 2D outgoing angles (under what conditions can we assume away incident angle? Perhaps in modeling the screen of a projector?).

<u>In Conclusion</u>

To further improve the accuracy of computer graphics, more research into the area of light reflectance is in order. The Holy Grail is making 8D BSSRDF efficient, intuitive, and as easy to manipulate as today's limiting parametric models. Research has begun but much more is needed.