Preference Galleries for Material Design (sap_0537)

Eric Brochu, Abhijeet Ghosh, Nando de Freitas*
The University of British Columbia

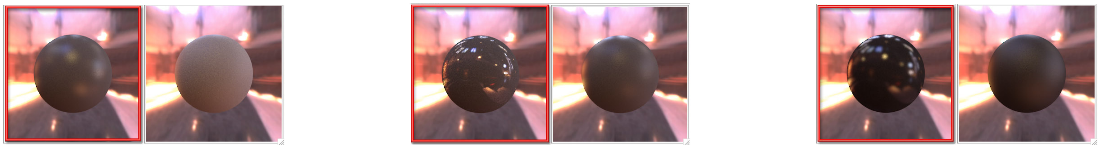

Figure 1: An example sequence of user selections (highlighted in red) of BRDFs in our preference gallery. Left-pair: Initial choices provided by the gallery. Center-pair: Intermediate choices. Right-pair: Final choices provided by the gallery before the user converges to the target high-gloss appearance.

1 Introduction

Properly modeling the appearance of a material is very important for realistic image synthesis. The appearance of a material is formalized by the notion of the Bidirectional Reflectance Distribution Function (BRDF). In computer graphics, BRDFs are most often specified using various analytical models. Analytical models that are of interest to realistic image synthesis are the ones that observe the physical laws of reciprocity and energy conservation while typically also exhibiting shadowing, masking and Fresnel reflectance phenomenon. Realistic models are hence fairly complex with many parameters that need to be adjusted by the designer for the proper material appearance. Unfortunately these parameters can interact in non-intuitive ways, and small adjustments to certain settings may result in non-uniform changes in the appearance. This can make the material design process hard for an artist or a non-expert user. To alleviate this problem, Ngan et al. [2006] recently presented an interface for navigation in a perceptually uniform BRDF space based on a metric derived from user studies. However, this is still somewhat constraining as the user has to develop an understanding of the various aspects of material appearance such as varying degrees of diffuseness, glossiness, specularity, Fresnel effects and/or anisotropy in order to navigate such an interface. An artist or a user often knows the look that he or she desires for a particular application without necessarily being interested in understanding the various subtleties of reflection! This is what we seek to address in this work with a 'preference gallery' approach to material design.

The preference gallery is to offers a number of BRDF examples (under natural illumination) to the user and he or she indicates which are most in line with what they are looking for. We propose a novel approach based on recent advances in Machine Learning to implicitly model the user's valuation function from their preferences and use the model in a principled way to generate new galleries that that present both novel BRDF examples and improvents of the settings the user has indicated preference for.

2 Active Preference Learning

While we are here interested in BRDF galleries, our preference model is very general and works with any parameter space $\chi \subseteq \mathbb{R}^d$. Assume we have shown the user M pairs of items (we talk about preferences as pairs for simplicity, but it is very easy to generalize to larger galleries). We define a set of preference relations $\mathcal{D} = \{\mathbf{r}_k \succ \mathbf{c}_k, k = 1 \ldots M, \mathbf{r}, \mathbf{c} \in \chi\}$ where $\succ$ indicates user preference. Using the techniques of Chu and Ghahramani [2005], we can then use Laplacian approximation to fit a latent Gaussian process (GP) that models the preferences.

A GP permits not only predictions, $\mu(\mathbf{x})$ for any $\mathbf{x} \in \chi$, but also estimates the prediction variance, $s(\mathbf{x})$. This allows us to define an *expected improvement* function that tells us the value of sampling at $\mathbf{x}$, $EI(\mathbf{x}) = (\mu_{\max} - \mu(\mathbf{x}))\Phi(d) + s(\mathbf{x})\phi(d)$ where $d = \frac{\mu_{\max} - \mu(\mathbf{x})}{s(\mathbf{x})}$, Φ and ϕ are the CDF and PDF of the standard Normal distribution, and $\mu_{\max}$ is the highest predicted value of $\mathcal{D}$ (that is, the example that is estimated the best one already shown to the user).

The EI function gives us a principled way of deciding which parameters to present to the user in order to maximize the user's valuation. Since sampling it is cheap and we need only approximate the maximum, it can easily be optimized using standard derivative-free techniques. Maxima may be points of high variance or high predicted valuation, or both. In practice, the first few examples will be points of high variance, since little of the space is explored (that is, the model of user valuation is very uncertain). Later samples will tend to be in regions of high valuation, as a model of the user's interest is learned. Note that we are *not* trying to learn the entire valuation function, which would take many more queries – we seek only to maximize the user's valuation, which involves accurate modelling only in the areas of high valuation. We are currently investigating different techniques to use the EI to select the image gallery, but so far asking the user to indicate preference between the images of maximum μ and maximum EI has proved very successful.

3 Example BRDF Gallery

We use our active preference learning model on an example gallery application for helping users find a BRDF. For the purposes of this example, we limit ourselves to isotropic materials and ignore wavelength dependent effects in reflection. The gallery uses the Ashikhmin-Shirley Phong model for the BRDFs and the Grace Cathedral HDR environment illumination. Our gallery demonstration presents the user with two BRDF images at a time. We start with four predetermined queries to "seed" the parameter space, and after that use the learned model to select gallery images. The GP model is updated after each preference is indicated. We use parameters of real measured materials from the MERL database for seeding the parameter space. The accompanying material shows a typical user run, where we ask the user to use the preference gallery to find a provided target image. At each step, the user need only to indicate the image they think looks most like the target. Using image pairs, it takes an average of 4 to 5 selections for the user to arrive at the target material appearance, depending on the material.

References

CHU, W., AND GHAHRAMANI, Z. 2005. Preference learning with Gaussian processes. In *Proc. ICML*2005*.

NGAN, A., DURAND, F., AND MATUSIK, W. 2006. Image driven navigation of analytical BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*, ???–???

*e-mail: {ebrochu, ghosh, nando}@cs.ubc.ca