Intuitive Editing of Material Appearance

Ana Serrano[1] Diego Gutierrez[1] Karol Myszkowski[2] Hans-Peter Seidel[2] Belen Masia[1]
[1] Universidad de Zaragoza [2] MPI Informatik

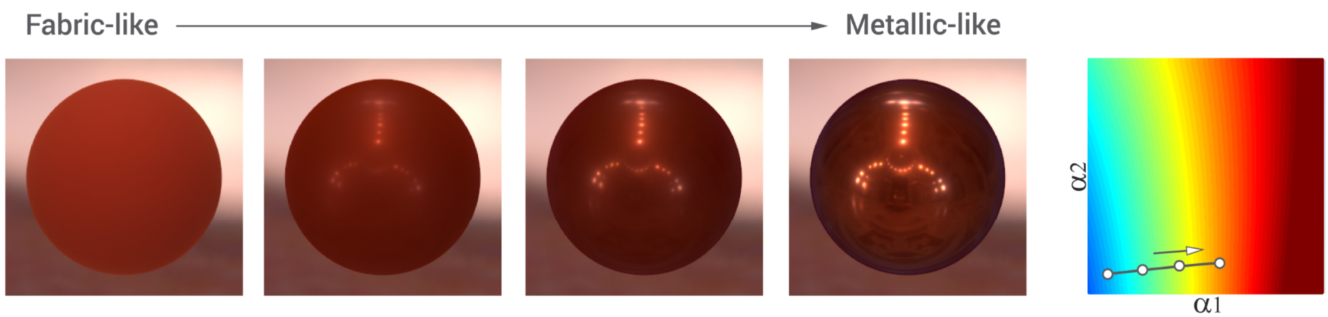

Figure 1: *Example of intuitive editing of material appearance using our framework. From a measured, fabric-like BRDF from the MERL database, we increasingly modify its metallic appearance by simply adjusting one perceptual attribute. Key to this ease of use and predictability of the results is our novel functionals, which map the coefficients of the first five principal components in a PCA space to the expected behavior of the perceptual attributes, based on a large-scale user study comprising 56000 ratings. The figure on the right shows a 2D slice of our five-dimensional PCA space illustrating this particular functional, and the path followed by the edit in such space.*

Abstract

Many different techniques for measuring material appearance have been proposed in the last few years. These have produced large public datasets, which have been used for accurate, data-driven appearance modeling. However, although these datasets have allowed us to reach an unprecedented level of realism in visual appearance, editing the captured data remains a challenge. In this work, we develop a novel methodology for intuitive and predictable editing of captured BRDF data, which allows for artistic creation of plausible material appearances, bypassing the difficulty of acquiring novel samples. We synthesize novel materials, and extend the existing MERL dataset [Matusik et al. 2003] up to 400 mathematically valid BRDFs. We design a large-scale experiment with 400 participants, gathering 56000 ratings about the perceptual attributes that best describe our extended dataset of materials. Using these ratings, we build and train networks of radial basis functions to act as functionals that map the high-level perceptual attributes to an underlying PCA-based representation of BRDFs.

We show how our approach allows for intuitive edits of a wide range of visual properties, and demonstrate through a user study that our functionals are excellent predictors of the perceived attributes of appearance, enabling predictable editing with our framework.

Keywords: BRDF editing, appearance editing, light reflection models, visual perception

1 Introduction and motivation

Measurement techniques for material appearance are gaining in accuracy, speed, efficiency, and ease of use (see for instance [Nielsen et al. 2015] for a very recent example). This has brought a paradigm shift in computer graphics towards data-driven appearance modeling techniques and databases [Matusik et al. 2003]. In this work we focus on intuitive, perceptually-based editing of captured BRDF data. Given the existence of large databases of measured BRDFs, a seemingly attractive option would be fitting them to parametric models. Unfortunately, this approach does not suit our goal of flexible material editing well, since the error introduced depends on the nature of the BRDF being represented [Ngan et al. 2005]. Moreover, the error metrics that guide such fitting do not take into account perceptual aspects, which might lead to visible artifacts for seemingly optimal approximations. Last, fitting requires a non-linear optimization which is often numerically unstable, expensive to compute, and typically involves visual inspection to judge the final outcome [Ngan et al. 2005].

Instead, we turn to a non-parametric approach, which can represent with high fidelity a wide scope of measured BRDFs, and lends itself naturally to accommodating our perceptually-based material editing framework. Recently, Nielsen and colleagues [2015] introduced a log-relative mapping that enables a convenient linear decomposition of measured BRDFs into their principal components. The first five of these components are nicely descriptive of appearance, but cannot be controlled in an intuitive manner for editing. The reason is twofold: First, their components are not able to properly isolate the different effects that characterize appearance; and second, linear variations in magnitude of the components result in highly non-linear changes in appearance.

We build upon their findings, and perform a series of experiments to build an intuitive space for editing, which shows that there is a much more intricate correlation between principal components, material appearance, and appearance perception.

2 Our approach

In our work, we make use of the PCA space derived by Nielsen and colleagues [2015], using the first five principal components of this space for BRDF representation. We first quadruple the original MERL dataset to 400 BRDFs, by synthesizing novel, mathematically valid samples from measured ones that we use in our experiments. We then find a mapping between the space of principal components and higher level perceptual attributes, and show that it can be used

SIGGRAPH '16 Posters, July 24-28, 2016, Anaheim, CA,
ACM 978-1-4503-4371-8/16/07.
http://dx.doi.org/10.1145/2945078.2945141

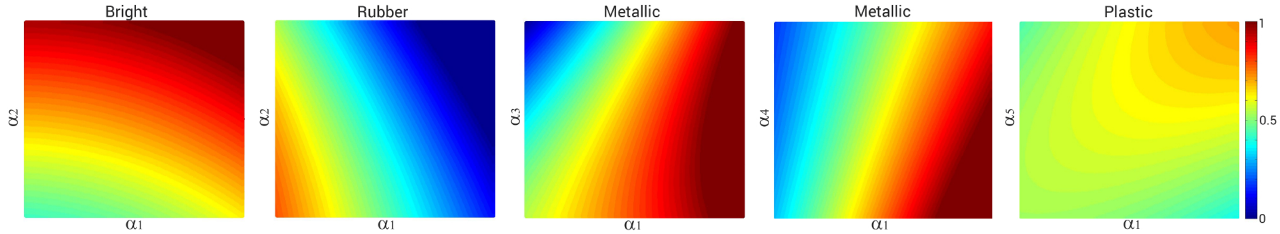

Figure 2: *Sample 2D slices of our functionals $\varphi : \mathbb{R}^5 \rightarrow \mathbb{R}$, mapping coefficients $\boldsymbol{\alpha}$ in the PCA basis to perceptual ratings for different attributes and along different dimensions.* ***From left to right:*** bright *(α_1-α_2 slice);* rubber *(α_1-α_2 slice);* metallic *(α_1-α_3 slice);* metallic *(α_1-α_4 slice); and* plastic *(α_1-α_5 slice).*

directly for intuitive material editing.

2.1 Obtaining a set of meaningful attributes

First, we performed a series of experiments to obtain a meaningful list of editing attributes. We compiled an extensive list of appearance attributes from previous works in industry and academia. Additionally, seven subjects were asked to provide, at least four attributes that described the appearance of a set of 60 BRDFs, using their own words; this yielded a second initial list of attributes. We then joined the two lists and reduced the number of entries by clustering semantically equivalent attributes; from this we obtained our initial list of 28 appearance attributes.

To further reduce the initial 28 attributes, keeping only those meaningful and understandable even by inexperienced users, we devised an experiment in which subjects had to establish, for each stimulus shown, whether each of the attributes applied to the material or not. This experiment would tell us: First, for which attributes there is a high agreement between users, and second, which attributes systematically received negative answers and thus are not representative.

The outcome of this experiment consists of a final list of eleven attributes, covering both high- and mid-level features: *plastic-like, rubber-like, metallic-like, fabric-like, ceramic-like, matte, glossy, bright, rough, strength of reflections*, and *sharpness of reflections*.

2.2 Building an editing framework

From each of the attributes of our list, a perceptual rating is obtained from a vast user study in which we gather 56000 answers, covering a wide range of attributes and BRDFs. To analyze how different materials are characterized in terms of our list of perceptual attributes, we model these attributes as Likert items. Each scale was numbered from 1 (*None, or very little*) to 5 (*A lot*). During the test, the participants were shown one rendered material at a time, plus the eleven perceptual attributes from the previous experiment; they were asked to rate each of the perceptual attributes, for each BRDF, in the Likert scale. In all tests, we also added a *control BRDF* which was used for outlier rejection.

We then trained radial basis functions networks (RBFNs) for each of the attributes, achieving a mapping between the perceptual ratings of each attribute and the underlying principal component basis coefficients.

Figure 2 shows a series of slices of the 5D space formed by the coefficients α_i of the first five principal components for different material attributes depicting our mappings obtained from RBFNs. We plot two-dimensional slices $\alpha_1 - \alpha_i$ ($i = 2..5$), since the first component (controlled by α_1) has the greatest influence on material appearance.

Observations on two-dimensional slices of our 5D PCA-space confirm that: i) analyzing each principal component of the BRDFs in isolation cannot explain how materials are perceived; and ii) our approach correlates well with human perception of materials, since we find many expected behaviors in our two-dimensional projections.

Once the RBFNs are trained, they can be used to intuitively and interactively edit a measured BRDF, yielding new, plausible BRDFs. In particular, a framework based on gradient descent is used to traverse the space defined by the coefficients of the first five principal components.

3 Conclusions and future work

Our framework enables performing intuitive, interactive modifications on measured BRDFs, as shown in Figure 1. We further validate our approach by means of a user study which confirms that our RBFNs can predict well the attribute values given by users.

Furthermore, our framework aligns changes of the attribute values with resulting visible changes in the material appearance. Some of these attributes provide a higher level description (e.g., *metallic, rubber...*), while others refer more directly to a particular property (e.g., *rough, glossy...*). This provides an appropriate trade-off between ease of use and fine control on the appearance edits.

Nevertheless, although we have shown a wide variety and range of edits, our work is in principle limited to the particular choice of stimuli BRDFs used in our experiments. Moreover, despite the fact that our extended MERL dataset provides a reasonably coverage over a very wide range of isotropic appearances, some perceptual attributes may be under-represented with respect to others. This may hinder some editing operations, since the derived functionals will be more inaccurate in those regions. On the other hand, our dataset can be easily extended, and our methodology could be applied to other databases as well. An exciting avenue of future work involves extending our methodology to increase the range of editable materials, including anisotropy or subsurface scattering effects

References

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (July), 759–769.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGSR '05, 117–126.

NIELSEN, J. B., JENSEN, H. W., AND RAMAMOORTHI, R. 2015. On optimal, minimal brdf sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG) 34*, 6 (November).