

Differential appearance editing for measured BRDFs

Apostolia Tsirikoglou* Joel Kronander Per Larsson Tanaboon Tongbuasirilai Jonas Unger
C-Research, Linköping University

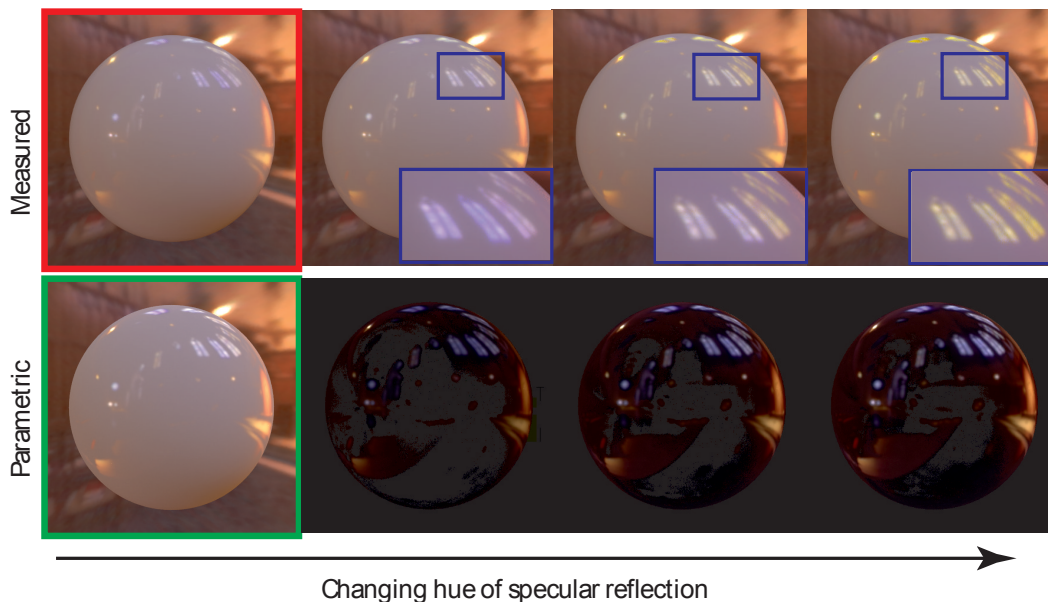


Figure 1: *Top* row shows an example edit to the measured BRDF “alimuna-oxide” from the MERL database [Matusik et al. 2003], where the hue of the specular lobe is changed. *Bottom* row shows the initial fit to an analytical, parametric, BRDF model [Ashikhmin and Shirley 2000] and the induced edits visualized as the absolute value of the residual vectors, which are used to approximate the edits made in the measured BRDF data.

1 Introduction

Data driven reflectance models using BRDF data measured from real materials, see e.g. [Matusik et al. 2003], are becoming increasingly popular in product visualization, digital design and other applications driven by the need for predictable rendering and highly realistic results. Although recent analytic, parametric, BRDFs provide good approximations for many materials some effects are still not captured well, see e.g. [Löv et al. 2012]. This makes it hard to accurately model real materials using analytic models, even if the parameters are fitted to data. In practice it is often desirable to be able to apply small edits to the measured data for artistic purposes, or to model similar materials that is not available in measured form. A drawback of data driven models is that they are often difficult to edit and do not easily lend themselves well to artistic adjustments. Existing editing techniques for measured data, for a comprehensive survey see e.g. [Schmidt et al. 2014], often use complex decompositions of the measured data that can make them difficult to use in practice.

In this work we propose and explore a simple and efficient method for intuitive editing of measured BRDF data. The artist friendliness of proposed algorithm lies in that it lets the user explore and edit the appearance of measured BRDFs by changing the parameters of analytical BRDFs familiar to the user. To maintain the physical accuracy inherent to measured data during the edits, we compute the difference between the edited analytical BRDF and the original analytical BRDF fitted to the data. We then use this analytical difference to approximate the tangent vector in the com-

plex space of measured BRDFs to induce the desired edit in the measured data. The proposed method enables physically plausible edits to the measured data, e.g. energy conservation, which can be difficult to achieve with previous techniques. We also explore how the proposed method can handle large edits to the measured data, and how the method can be used for physically plausible interpolation between measured BRDFs. Our method allows the user to rapidly create a range of artistic adjustments using familiar analytical BRDFs, for example by changing the color of the specular reflection while keeping the rest of the lobe intact, see Figures 1 and 2. More examples can be found in the supplementary video.

2 Editing

We assume that the measured BRDF data is represented in some suitable basis using coefficients $\{c_i^m\}_{i=1}^N$. In the simplest case these correspond to the canonical basis in some parameterization over the incident and outgoing directions over the hemisphere, e.g. representing the BRDF as a 3D matrix with N non-zero entries.

2.1 Parametric tangent approximation

To provide a reasonable starting-point for the user edits, we fit a parametric BRDF model, $\rho(\omega_i, \omega_o)$, to the data and expose the fitted parameter values $\{p_i\}_{i=1}^K$ to the user. We explore the use of both standard least squares fits, as well as other perceptually motivated fitting metrics [Löv et al. 2012]. The user then edits the parameters to induce a desired change in the measured data. For example, increasing the gloss parameter in an ABC-BRDF model

*apostolia.tsirikoglou@liu.se




Figure 2: We propose a simple and flexible framework for intuitive editing of measured BRDF data. our approach is based on using edits of analytical BRDF models to induce similar changes in the measured data. See the text for more details.

fit, [L ow et al. 2012], produces a set of edited parameters $\{p_i^e\}_{i=1}^K$.

The residual between the initial fit, $\{p_i\}_{i=1}^K$, and the edited parameters, $\{p_i^e\}_{i=1}^K$, is then computed using the parametric model and expressed in the same basis as the measured material, resulting in the coefficients $\{c_i^r\}_{i=1}^N$. We then update the measured data coefficients to describe the change introduced by

$$c_i^m \rightarrow \max(0, c_i^m + c_i^r) \quad (1)$$

Note that this allows us to keep the high accuracy in the measured data, and that we are only using the limited accuracy provided by the parametric model fit when representing the approximate tangent vector (induced by the user edits) in the space of measured materials. For larger edits that significantly reduce the width of the specular lobe, artifacts can appear in the measured data due to zero crossings. To handle large edits of these types we linearly interpolate the final result between the edited measured material and the parametric result. As illustrated in the supplementary video, our BRDF editing software lets the user select which analytical model to use, edit the parameters, and interactively adjust the appearance with visual feedback of the result.

2.2 Physically plausible edits

If physical plausible BRDF models are used to approximate differences and induce tangents we can take advantage of their energy conservation and reciprocity properties in the editing procedure. In fact, it is easy to see that the final, edited, measured BRDF will also be physically plausible if the tangent vectors describing the edits are computed using a physically plausible analytical model. This as the tangent induced from the physically plausible analytical model will by definition preserve energy and reciprocity for differential changes.

2.3 Large Edits

After a series of edits, the original measured material may no longer be the best representation available. However, by using a database of measured materials we can improve our method by precomputing parameters $\{p_i^m\}_{i=1}^K$ for the analytical BRDF model for each material m in the database. When a user makes a large edit we first search for the closest matching set of parameter coefficients $\{p_i^f\}_{i=1}^K$ in the database, corresponding to material f . We then compute the residual of the analytical BRDF model using the closest matching parameters $\{p_i^f\}_{i=1}^K$ and the edited $\{p_i^e\}_{i=1}^K$, and apply that change to the coefficients of the closest matching material in the database. To enable smooth transitions, the two closest measured materials can also be used and the desired edit is then interpolated between the computed results (measured data + residual). In our examples, we use the MERL-database [Matusik et al. 2003].

3 Interpolation

The proposed method can also be used for interpolating between two measured BRDF models. The simplest solution is to directly interpolate the coefficients $\{c_i^m\}_{i=1}^N$ of the measured materials. However, this approach does not preserve desirable physical invariants, such as energy and reciprocity, and can result in very complicated intermediate BRDFs, far from physically plausible. Instead we propose to use an efficient approach similar to that used to handle large edits presented in section 2.3. First we fit an analytical BRDF model to both materials, resulting in coefficients, $\{p_i^{m1}\}_{i=1}^K$ and $\{p_i^{m2}\}_{i=1}^K$. We then construct a linear map from $\{p_i^{m1}\}$ to $\{p_i^{m2}\}$ for each $i = 1 \dots k$. The user is then exposed to the parameters of this linear map of the parameter space. In practice, we expose the user to a set of sliders between 0 and 1 for each of the parameters $i = 1 \dots K$, where a setting of 0 on all sliders corresponds to the first material and a setting of 1 corresponds to the second material. Given a user setting, we then infer an intermediate analytical model $\{p_i^I\}_{i=1}^K$. By computing the difference between this intermediate analytical model and the first and second analytical fits, we get two difference vectors. The first difference vector is applied to the first measured BRDF, to compute, $\{c_i^{I1}\}_{i=1}^N$, according to (1). The second difference vector is then applied to the second measured BRDF, computing $\{c_i^{I2}\}_{i=1}^N$ in the same way. The final intermediate BRDF is then obtained as a weighted average of $\{c_i^{I1}\}_{i=1}^N$ and $\{c_i^{I2}\}_{i=1}^N$, where the weights are proportional to the distance in the parameteric space. This relatively simple technique lets us explore intuitive interpolations between measured BRDF data that preserve physically plausible invariants such as energy conservation and reciprocity, provided that the underlying analytical BRDF model preserves these.

References

- ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic phong brdf model. *Journal of graphics tools* 5, 2, 25–32.
- KERR, W. B., AND PELLACINI, F. 2010. Toward evaluating material design interface paradigms for novice users. *ACM Trans. Graph.* 29, 4.
- L OW, J., KRONANDER, J., YNNERMAN, A., AND UNGER, J. 2012. Brdf models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph.* 31, 1.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. Graph.* 22, 3.
- SCHMIDT, T.-W., PELLACINI, F., NOWROUZEZAHRAI, D., JAROSZ, W., AND DACHSBACHER, C. 2014. State of the art in artistic editing of appearance, lighting, and material. In *Eurographics STARS*, Eurographics.