

Improved Intuitive Appearance Editing based on Soft PCA

Sandra Malpica, Miguel Barrio, Diego Gutierrez, Ana Serrano and Belen Masia

Universidad de Zaragoza, I3A

Abstract

*During the last few years, many different techniques for measuring material appearance have arisen. These advances have allowed the creation of large public datasets, and new methods for editing BRDFs of captured appearance have been proposed. However, these methods lack intuitiveness and are hard to use for novice users. In order to overcome these limitations, Serrano et al. [SGM*16] recently proposed an intuitive space for editing captured appearance. They make use of a representation of the BRDF based on a combination of principal components (PCA) to reduce dimensionality, and then map these components to perceptual attributes. This PCA representation is biased towards specular materials and fails to represent very diffuse BRDFs, therefore producing unpleasant artifacts when editing. In this paper, we build on top of their work and propose to use two separate PCA bases for representing specular and diffuse BRDFs, and map each of these bases to the perceptual attributes. This allows us to avoid artifacts when editing towards diffuse BRDFs. We then propose a new method for effectively navigate between both bases while editing based on a new measurement of the specularity of measured materials. Finally, we integrate our proposed method in an intuitive BRDF editing framework and show how some of the limitations of the previous model have been overcome with our representation. Moreover, our new measure of specularity can be used on any measured BRDF, as it is not limited only to MERL BRDFs [MPBM03].*

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Material perception

1. Introduction

Measurement techniques for material appearance have become more efficient and accurate over the last few years (e.g., [NJR15], [AWL15]). This has led to an increase in the number of available public databases and modeling techniques [MPBM03], [FVH14]. Yet editing material appearance has remained a challenge due to several reasons, including the disconnection between the data representation and human understanding, the complexity of some representations that require handling hundreds of parameters [ATDP11], [Bur12], and the high-dimensionality, non-linearity of real world appearance behavior that leads to unintuitive editing parameters.

However, the increase of public available data has helped in the construction of new editing techniques for the captured data [SGM*16], [WAKB09] that have arisen in the last few years overcoming the problems mentioned above. Among them, those techniques based on the use of intuitive parameters, like the work of Serrano et al. [SGM*16] stand out due to their ease of use for novice users.

In this paper, we build on the intuitive editing of material appearance method by Serrano et al. and address some of its limitations regarding the edition of certain types of materials. They make use of a BRDF representation consisting on a combination of principal

components (PCA) to reduce dimensionality, and then map these components to perceptual attributes. This PCA representation is biased towards specular materials and can fail to properly represent very diffuse BRDFs, producing artifacts when editing. The reason of this problem is the high values of the specular peaks of the measured BRDFs compared to those of the diffuse materials. Despite the log-mapping proposed by Nielsen et al. [NJR15] over the raw data before performing the Principal Component Analysis to reduce the weight of these specular peaks (that Serrano et al. also use in their work), their influence is still biasing the PCA basis in such a way that diffuse materials are not well represented. This problem could be solved with a naive approximation if we were talking only about BRDF representation and not BRDF editing, as a new PCA basis could be used to represent diffuse materials. However, it is important for us to preserve the intuitive editing process proposed by Serrano et al. and to improve it with a correct representation of diffuse materials, so that the users can carry out intuitive appearance edits on a wider range of materials.

To solve this problem, we build on top of Serrano et al.'s work and propose to use a new PCA base, together with the original one, that better supports the representation of diffuse materials. We also map this new PCA base to the perceptual attributes found by Serrano et al., thus allowing the edition technique to follow the same rules even if there are two different PCA bases.

We then propose a new method for effective navigation between both bases while editing. It relies on a new measurement of the specularity of measured materials. Using the specularity measurement to separate between diffuse and specular materials, we fix a hyperplane using support vector machines (SVM) that works as a frontier between both representations on their PCA subspaces.

Finally, we integrate our proposed method in an intuitive BRDF editing framework and show how some of the limitations of the previous model have been overcome with our representation. Moreover, our new measure of specularity can be used on any measured BRDF, as it is not only limited to BRDFs present on MERL's database [MPBM03], a public dataset of 100 measured BRDFs from a wide range of materials.

The next section shows the related work and state of the art in this particular field of computer graphics. Sections 3 to 5 show the technical details of this work, including how the new representation has been built (Section 3), how the measurement of specularity and navigation between PCA bases work (Section 4) and the integration of the representation in an interactive edition plugin (Section 5). Finally, the results are shown (Section 6) and the conclusions and future work of this paper are presented.

2. Related Work

Industrial standards. Over the years, several standards have been developed in order to characterize reflectance of real world materials [HH87]. These have been formalized in a collection of documents by the American Society for Testing and Materials (ASTM). Among them, some standards are specialized on defining a number of gloss dimensions [WAKB09] in addition to the angles from where the material should be measured [Alm87], [SGVH90]. This allows for a complete characterization of the glossiness of materials. Westlund and Meyer [WM01] find a correlation between this glossiness characterization given by the industry and parameters of some analytic reflectance models, linking in this way the industry standards of real world materials with properties of virtual BRDFs.

Editing of existing models. Some works on editing existing BRDFs include BRDF-Shop [CP06], which includes an artist-friendly interface and a model based on an extension to Ward's. Other works are based on image-driven navigation [NDM06] using as distance the difference between rendered images. Talton et al. [TGY*09] focus on a collaborative editing system over the space of Ashikhmin's model [APS00]. Other existing works are based on fast feedback of BRDF edits [SZC*07, CPWAP08, NKLN10] or specialize in a narrow variety of editions, as those needed in car paint design like the work of Ershov et al. [EKM01], focused on the physical properties of paint ingredients which affect its appearance. There are also existing works that edit measured BRDF data without fitting to any parametric models, in spite of its drawbacks (like large editing spaces [WAKB09]). Lawrence et al. [LBAD*06] decompose the BRDFs behaviour to simple 1D curves representing its physical properties. Their work is continued by Ben-ARTzi et al. [BAOR06] allowing for complex direct lighting as well as inter-reflections. Some of these methods lack a sufficiently large editing space as well as intuitive editing parameters.

Perceptual editing spaces. Perceptual strategies have been ap-

plied by many different works in computer graphics [MMG11]. On the little explored work of editing light fields, Jarabo et al. [JMB*14] use different image-based tools to show the potential of light field editing, besides favoring artistic exploration with their versatile interfaces. Applications using high dimensional spaces include crowd simulation [GKLM11], font design [OLAH14], style similarity [GAGH14], interior design taxonomy [BUSB13], translucency perception [GXZ*13], or shader design [KSI14]. After them, Boyadzhiev et al. [BBPA15] introduced intuitive attributes for image-based material editing. For garment simulation, Sigal et al. use the parameters of a custom pipeline simulator [SMD*15]. Pellacini et al. [PFG00] observed that non-linear behaviour of analytic parameters caused the editing process to be unintuitive when directly tuning these parameters. Due to this finding, they derive a perceptually uniform parameter scaling for the Ward model, used since then to study image-driven navigation spaces [NDM06] or the influence of shape in material perception [VLD07]. The concept of perceptually uniform spaces is extended by Wills et al. [WAKB09] for measured BRDFs. They propose a intuitive, low-dimensional space suited for the construction of new materials, although their space is limited to gloss editing. Kerr and Pellacini [KP10] showed that image-driven navigation seemed to be less efficient than the use of the Ward model and its perceptually linearized version, for the particular task of matching material appearance. However, they only used two simple sliders (diffuse and specular) and only studied colorless BRDFs.

Matusik et al. [MPBM03] present a data-driven reflectance model. They reduce the measured BRDFs' dimensionality either by performing a principal component analysis (PCA) or by using non-linear dimensionality reducers. Then, they define a set of *perceptual traits* and carry a single user experiment where the user decides whether a given material has or not this perceptual traits. They build trait vectors that allow the navigation in their BRDF spaces by specifying desirable changes for their traits that result in given directions for the navigation process.

Building upon Matusik's work, the work of Serrano et al. [SGM*16] is different in many ways. They also reduce the measured BRDFs' dimensionality by means of principal component analysis, and employ a set of perceptually meaningful attributes for material edition. However, their 14 attributes are found carrying out a large-scale experiment and are identified as intuitive, descriptive and discriminative when describing reflectance properties. They use the same decomposition method for dimensionality reduction and perceptual meaningful scaling used in the work of Nielsen et al. [NJR15] that makes their reduced PCA space comparable to purely perceptually derived spaces [WAKB09]. They perform a dense uniform sampling of their PCA space, synthesizing additional BRDFs from their initial MERL dataset (totaling 400), and obtain 56000 attribute ratings from 400 participants on a large scale experiment. They use these ratings to reconstruct perceptually-based complex embeddings of their attributes in the PCA space. This allows for intuitive editions of measured BRDF data, where a novice user can easily change the appearance of a material.

However, their PCA base is biased towards specular materials. This means that artifacts may appear sometimes when trying to rep-

represent diffuse materials or during the edition when lowering specular attributes. The artifacts are usually a dark halo around the specular direction. To solve this problem, in this work we follow [NJR15] suggestion and train a separate 'soft' PCA base to represent only the diffuse materials. Using the same ratings obtained by Serrano et al., we map the new PCA coefficients to the same set of 14 attributes found on their work. However, this is not enough when editing the materials since the base to be used on each moment depends on the particular BRDF characteristics that change along the edition. This is why we propose a new specularity measurement for BRDFs and a technique to navigate between both PCA representations when needed. This allows us to address the diffuse artifacts and build a editing technique that works on a wider range of materials.

3. Creation of the new representation

In this section we replicate the work of Serrano et al. First we train a PCA basis from MERL measured BRDFs and then we map the obtained PCA coefficients with user-given values for the 14 perceptual attributes.

3.1. PCA training

A principal component basis can be trained using a dataset of measured BRDFs. Once the PCA basis is trained, it can be used to represent any other BRDF as follows:

$$b = Q\alpha + \mu \quad (1)$$

Where \mathbf{b} is the BRDF represented in the basis, \mathbf{Q} is the matrix representing the PCA basis, μ is the average of the measured data and α the coefficients of each of the principal components for the particular BRDF \mathbf{b} . Prior to the PCA training, a log-relative linear mapping of the measured data is performed to reduce the effect of the specular peaks on the overall information, as Serrano et al. do on their work.

As we have discussed in previous sections, the PCA basis used in Serrano et al. [SGM*16] is biased towards specular materials, thus not allowing more diffuse BRDFs to be represented properly. This is not only a problem on the representation of existing BRDFs, but also in the process of intuitively editing materials as the specularity of BRDFs can be reduced when some of its attributes are modified (i.e. when the *matte* attribute increases) causing unwanted artifacts to appear. To address this problem, Nielsen et al. [NJR15] suggest that a separate, soft PCA basis could be used. This is the approach followed in this work, as explained in this section.

We build upon the work of Serrano et al. [SGM*16] and create a new PCA basis, using a subset of the initial 100 MERL BRDFs [MPBM03]. The range of materials that the basis can represent correctly depends to a large degree on the BRDFs used to train that PCA basis. In this case we want to represent correctly diffuse materials and avoid the bias towards specular peaks. Thus, from the 100 BRDFs that form the database, we choose a subset consisting of the less specular materials and train a PCA basis with them. Initially, the 100 BRDFs are rendered and the diffuse materials are selected manually based on their appearance and the shape

of their specular peaks. Figure 1 shows that this trained soft PCA basis can represent correctly diffuse materials. In Figure 3 we show a representation of the first five principal components of each PCA bases. According to Nielsen [NJR15] the first principal component of the 'full' PCA basis was strongly correlated with the specular peak information on the BRDFs. It is not surprising then that the soft PCA first component is widely different than the previous one, as the majority of the specular information has been removed from the training dataset. We also confirm that the soft PCA basis can't represent specular materials: Figure 2 shows an example of a specular material reconstructed from its soft PCA coefficients compared to the full PCA representation and the original BRDF from MERL.




Figure 1: The original MERL BRDF white-fabric (left), compared with the reconstruction from its soft PCA coefficients (center) and a reconstruction from the original work of Serrano et al. [SGM*16] (right). The artifacts of the diffuse material disappear with the new PCA basis.



Figure 2: The original MERL BRDF red-specular-plastic (left), compared with the reconstruction from its soft PCA coefficients (center) and a reconstruction from the original work of Serrano et al. [SGM*16] (right). The artifacts present in the soft PCA representation show that it can't correctly depict specular materials.

3.2. RBFN training

The neural network used by Serrano et al. uses a radial basis function (RBF) for each attribute that can be expressed as follows:

$$y = \varphi(\alpha) = \sum_{i=1}^{N_c} \theta_i \exp^{-\beta \|\alpha - c_i\|^2} \quad (2)$$

Where y is a value in the range [0..1] that represents the strength of said attribute, N_c is the number of neurons in the network, c_i are the centers of the neurons and θ_i the weights for each neuron. β is a parameter that controls the smoothness of Gaussian functions and α the first five coefficients of the PCA basis.

With the PCA basis trained, we need to relate the new PCA

components to the perceptual attributes found by Serrano et al. [SGM*16] listed in Table 1. To do so, we use their experimental ratings of 406 known BRDFs that give each BRDF its attribute values. With the PCA components on one hand, and the experimental ratings on the other, we train a radial basis function network (RBFN) to model the functional that, given the PCA components of a BRDF returns the value of its perceptual attributes, ranking from 0 (none or very little) to 1 (a lot). We only use the five first principal components for this as Serrano et al. show in their work that five components are enough to recover the measured BRDF correctly. This applies to all the specular or semispecular materials that are well represented on their PCA basis, and also to the diffuse materials in our soft PCA basis.

Our input data consists on 406 different BRDFs (100 from MERLs database plus 306 new synthesized BRDFs from Serrano et al. [SGM*16]) and the 56000 attribute ratings (400 BRDFs x 10 responses/BRDF x 14 questions/BRDF) gathered through Mechanical Turk by Serrano et al.. From our dataset, we separate randomly 5% of the BRDFs to use as a validation set, and do the training with the remaining 95%. We repeat this process for 10 times with different training and validation sets to find the average error of the neural network. We carry out this process with different training parameters to minimize the average validation error. At the end, we choose a neural network with 25 neurons and a sigma value of 10 (this value is given by the trainer of the neural network, and is usually achieved through experimentation; the smaller the value the tightest the fit to the data, and the bigger the smoother the results will be), which are different parameters to the ones originally used by Serrano et al.

1	<i>plastic</i>
2	<i>rubber</i>
3	<i>metallic</i>
4	<i>fabric</i>
5	<i>ceramic</i>
6	<i>soft</i>
7	<i>hard</i>
8	<i>matte</i>
9	<i>glossy</i>
10	<i>bright</i>
11	<i>rough</i>
12	<i>strength of reflections</i>
13	<i>sharpness of reflections</i>
14	<i>tint of reflections</i>

Table 1: List of the 14 attributes used in the BRDFs edition.

With the PCA basis trained and optimized, and the neural networks (one for each attribute) fitted to the data we can rebuild the editing process and see how the new soft PCA basis behaves. This test shows clearly that the soft PCA basis is not enough by itself to depict all the material variety that can be found in the MERL database. In the same way, the original full PCA representation isn't either able to correctly depict diffuse materials. To correctly represent both types of materials, we need to use each of the PCA bases when needed, according to the specularity of the material. Additionally, in order to use these representations for editing, we need

to know which representation is needed in each edition step, as the specularity of a BRDF might change when modifying some of its attributes. To solve this we propose in the next section a new specularity measure that determines when to use what PCA basis and a novel approach to navigate between both representations when needed during the editing process.

4. Navigation between representations

4.1. Measure of material specularity

As expected, it has been demonstrated that neither our soft PCA basis nor Serrano et al.'s [SGM*16] can correctly represent the full MERL's database [MPBM03] range. It is necessary that both of them are used to achieve correctness in the full range of specularity. To do so, we propose a novel measure of specularity that allows us to use one basis or the other when needed, as well as a method to navigate between representations when there isn't an original reference from which we can obtain the PCA coefficients of both representations (that is, during the editing process).

Industry standards. Many international standards define the methods and specifications for glossiness measurement of real world materials using glossmeters. This glossiness is related to the specularity of a BRDF, and the industry sets different measurement techniques and glossmeters for different types of materials. Depending on the technique used, the light and sensor are placed at different angles. According to Alman et al. [Alm87] and Saris et al. [SGVH90], 3 measurements at different angles are often enough to determine the glossiness of a metallic sample. Nevertheless, this approximation presents two drawbacks: first, each standard is specifically designed for a given material type (different materials - i.e. metal and plastic - are usually measured at different angles). And second, these measurements are meant to be obtained through a glossmeter in real life materials. To use this in virtual materials we would need to implement a virtual glossmeter that would increase the complexity and processing time of the editing process, and that is something we want to avoid as much as possible.

Parametric model fitting. Following the work on applying industry standards to virtual materials of Westlund and Meyer [WM01] we find their correlation between the glossiness of a material and the values of certain parameters of analytic models. According to Westlund and Meyer, there is a relationship between the specularity coefficient of some analytic models and the glossiness measurement of a material. Using *Alta: a library for BRDF analysis* [Alt] we fit each one of the 100 MERL BRDFs to four different analytic models: Beckmann [BS87], Blinn [Bli77], Lafor-tune [LFTG97] and Ward [War92]. From the fittings, we study the values of each model's parameters to see if there is a correlation between them and the material's specularity. We found two different parameters where this evidence can be seen (Blinn and Ward specularity coefficient). Generally the more specular materials have higher parameter values, while the more diffuse materials present lower parameter values. However this behaviour is not always consistent, which means this option is not reliable in 100% of the samples. Besides, cost of computing these parameters also increases the level of complexity during edition. The specularity measure should be needed for each of the editing process steps. Due to this rea-

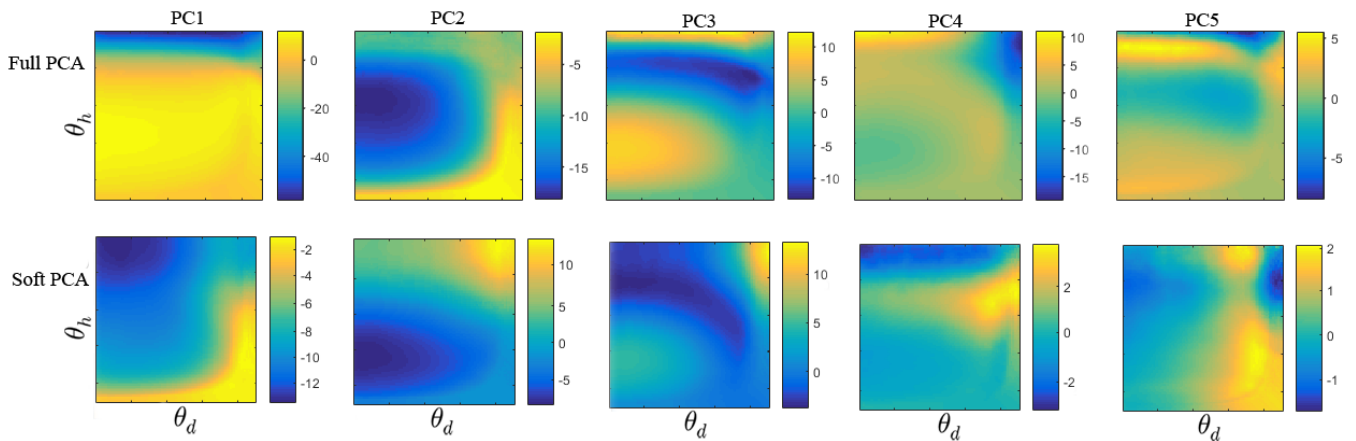


Figure 3: Comparison between the old BRDF representation (first row) with the new soft PCA basis (second row). For each row, the columns represent a 2D slice of the first five principal components eigenvectors. These slices are made using Rusinkiewicz's BRDF half-difference angle coordinates $(\phi_d, \theta_d, \theta_h)$ [Rus98] at $\phi_d = 90$.

sons, we decide to find another specularity measure for measured BRDFs.

Nielsen optimal directions. The work of Nielsen [NJR15] shows that good quality BRDFs can be acquired with few measurements in certain specific angles. This means that the characteristic behaviour of a material can be extracted from these directions. Our hypothesis is that the specularity information of the material can be obtained by just looking at the values stored on these directions, and that a difference in behaviour will appear between specular and diffuse BRDFs. To prove this, we obtain the values for the five most significant directions that Nielsen provides in his work for two very different materials and compare them, as this is the minimum number of samples required to correctly depict most BRDFs. In Figure 4 we can see a clear evidence that sustains our hypothesis. The specular material shows higher values on the specular peak as expected, while the diffuse material shows a slight descent on the same exact position. With this simple technique, we can tell whether a material is or not specular. To extend this proof of concept, we carry out the same procedure, this time with the whole MERL database. Figure 5 shows the values of the mentioned 5 most significant directions for the 100 MERL BRDFs. We can clearly distinguish two different behaviours to which most of the materials correspond, that are strongly correlated to the specularity of the BRDF. In Figure 5 the warmer colours correspond to more specular materials, while the cooler are associated with more diffuse materials.

We also show that even using the single most significant direction that Nielsen et al. provide on their work, this difference in behaviour remains. Figure 6 again shows for the 100 MERL BRDFs the value of the most significant direction given by Nielsen. Once again there is a clear difference where the diffuse BRDFs show the lowest values (below 0.4) and the specular BRDFs show higher values, in some cases several orders of magnitude higher. We propose as specularity measure an average of the RGB values for the most significant direction given by Nielsen on his work, as this

is a simple yet effective method to acquire information about the specularity of a material. To validate our specularity measure we visually explore the 100 MERL BRDFs as well as their polar plots. All the polar plots of the BRDFs whose specularity measure is below 0.4 show a gentle shape, distinctive of diffuse materials, while the specular materials show specular peaks on their polar plots that help us differentiate between these two categories of materials, thus empirically showing that the defined threshold is valid for our application.




Figure 4: Values of 2 BRDFs on the first five most significant Nielsen directions. The samples are light-brown-fabric ('*') and steel ('o'). For each one of them, the values of the three color channels of the tabulated BRDF are shown.

Figure 7 shows the specularity value for the 406 BRDFs used in the attribute training in the previous section, each one of them projected in the first three dimensions of the PCA components for both representations. Note that all values above 1 have been clamped to have a clearer vision of the diffuse region. It is interesting to see how the diffuse materials are all bound together in a small area of the full PCA space, while they occupy a much larger zone in the soft PCA representation. This shows that the soft PCA basis full-




Figure 5: Values of the 100 MERL BRDFs on the first five most significant Nielsen directions. For the sake of simplicity, only one color channel is shown for each of them. The colormap that shows the difference between diffuse and specular BRDFs is obtained with the difference from the third and fourth directions value.




Figure 6: Values of 100 BRDFs on the most significant Nielsen direction. On the right graphic we can see the values of the 100 BRDFs, while the middle and left ones are two consecutive zooms among the first one. For the sake of simplicity, only one color channel is shown for each of them. The colormap is obtained with the mean of the three color channels.

fits its purpose of maximizing the diffuse information alongside its first components, where the full PCA basis was biased towards the representation of the specular peaks of other materials.

4.2. Navigation between the representations

Once we know when do we have to use *what* PCA basis, we need a technique to represent the same BRDF on both representations and to navigate between them (i.e., given the full PCA coefficients of a BRDF, find its soft PCA coefficients when there is no original measured data during the editing process). The first idea is to find a change of basis matrix, but no linear application is found that works on any direction. The second idea is to rebuild the measured BRDF from its PCA coefficients from one basis and then from the rebuilt BRDF find the PCA coefficients of the other basis. This approximation works well when the starting point is the full PCA basis, and only if the BRDF was correctly represented on that basis before the change. This means that once the undesired artifacts have

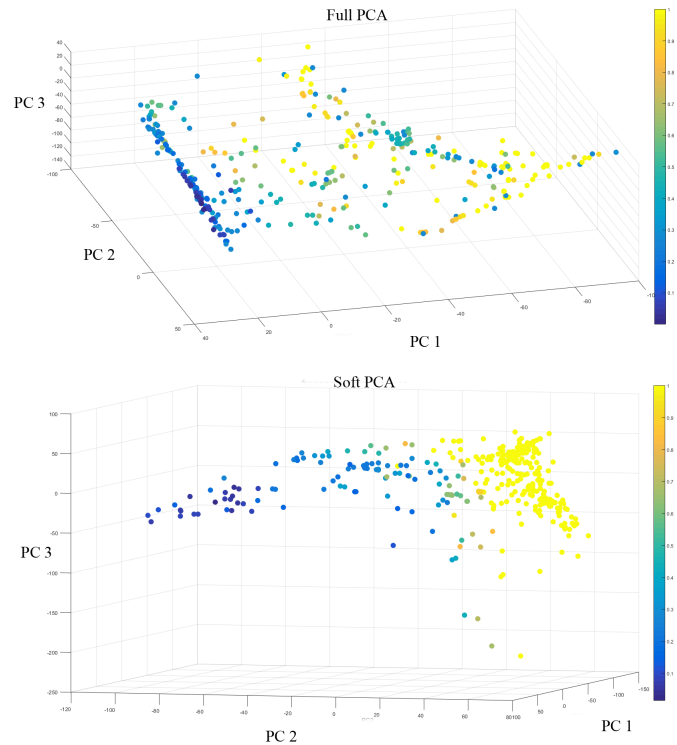


Figure 7: Shows the first three PCA coefficients of 406 BRDFs (100 of which are MERL’s database and the rest the synthesized by Serrano et al.) for full PCA (top) and soft PCA (down) representation. The color shows their specularity measure. Values under 0.4 represent diffuse materials.

appeared in a representation the missing information cannot be recovered even if the other representation is used, so the navigation between the PCA bases must occur in a point where both representations behave correctly (this happens in the frontier between specular and diffuse materials, given by the specularity threshold) to avoid this problems. Unfortunately, due to the limited dataset with wich the soft PCA basis was trained this approximation doesn’t work on the opposite direction.

To move from the soft PCA space to the full PCA representation a different method is needed. We use the same neural networks used in the attribute trainings to learn the correspondence between both bases. Once again, we perform the same cross validation process to minimize the average error and find the optimum parameters, this time with the PCA coefficients as input. This neural network has 100 neurons and a sigma value of 100. With these two methods we can finally navigate between both representations when needed.

5. Interactive plugin for material edition

One of the contributions of the work of Serrano et al. [SGM*16] is that they allow the edition of a material through simple, intuitive attributes that the user can understand. This eliminates the use of

unintuitive and non-friendly statistical parameters that drove to visually non-linear changes with which the users had to deal before. This is achieved through a gradient descent algorithm along the PCA space, in which the attribute values are used as the minimization parameter. Thanks to the functionals trained with the RBFNs, every point of the 5 dimensional space is associated with one value for each of the perceptual attributes. Therefore, a change in an attribute's value leads to a change of the 5 principal components and to a change in the material appearance. Figure 8 shows a 2D slice of the PCA space colormapped with the attribute value, where the path followed during a edition of a BRDF follows a desired change on this value.

The adaptation of this process to the new representation is simple. First, the specularity measure is obtained from the original BRDF to edit. Depending on this value and the specularity threshold, the gradient descent is performed on the corresponding PCA basis. As the material characteristics change during the editing process, a new specularity measure is obtained in each editing step. If the specularity value reaches the limit of its representation a navigation between the PCA bases occurs. To avoid swinging persistently among the representations, a certain hysteresis is used on the navigation boundary, calibrated through experimentation. This means that, for example, if the specularity threshold to jump from soft PCA to full PCA representation is set to 0.4, the specularity threshold to navigate from full PCA to soft PCA could be set to 0.35.



Figure 8: Edition of the chrome BRDF decreasing its strength of reflections. The original BRDF can be seen on the left, and the edited on the center. On the right, a 2D slice of the PCA space with the path followed for the edition (blue line). The colormap represents the attribute value, with warmer colours at its maximum.

Once the edition works with both PCA bases, we take one more step and present the two representation technique in an interactive editing plugin. This plugin uses Disney's *BRDF Explorer* to visualize BRDFs and a *Matlab* interface that gives the user a friendly tool to edit materials. As our work relies on *BRDF Explorer* to show the material that is being edited we do not consider it a standalone application but an interactive editing plugin. Figure 9 shows the interface, where 2D slices of both representations can be seen.

In order to improve the time response and the user experience of the plugin two main changes are made to our existing system:

- The boundary where both representations behave correctly is a part of the 5D space where each of them exists. As mentioned in Section 4, we empirically see that the diffuse BRDFs of our dataset show a specularity measure below 0.4 and manually fix

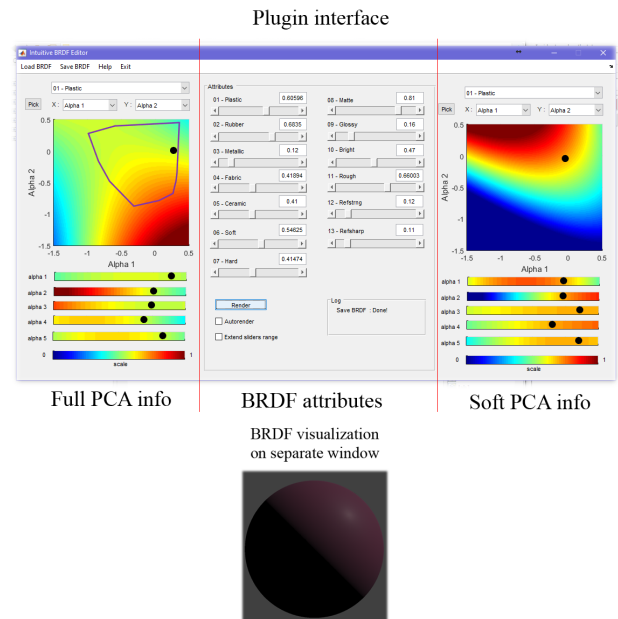


Figure 9: Our plugin interface shows a 2D slice of the full PCA space (left), the BRDF attribute values (center), and a 2D slice of the soft PCA space (right). On a separate window, a visualization of the BRDF is shown.

that as the specularity threshold. Considering such dimensionality, the fact that the manual calibration of the specularity threshold doesn't always work is not surprising. Besides, the obtaining of each specularity measurement means the reconstruction of the tabulated BRDF from its PCA coefficients. Considering the spatial location and separation of specular and diffuse BRDF on each PCA basis, as Figure 7 shows, the idea of building a spatial division for the navigation based on the specularity measurement of the available BRDFs arises. We decide to train support vector machines (SVMs) with this information, in such a way that the SVMs will tell us if a BRDF is in the space zone where the PCA basis correctly represents materials or if we have to change to the other representation. Thus we achieve an automatic frontier for the editing process. The use of the SVMs gives a more accurate criteria to distinguish between specular and diffuse BRDFs, as it learns the relation between the specularity of a BRDF and the point of the PCA space where it is located. Thus, we think that this is a better approach to classify BRDFs based on their specularity.

- The other modification is the navigation between full PCA and soft PCA representation. We find out that the navigation made with the RBFN is much faster than the reconstruction one (from some *ms* for RBFNs to around 3s for reconstruction method depending on the machine used). Removing this bottleneck will give the user a better experience and make the transition even on both directions. Thus, we train another RBFN for this purpose.

- Last but not least, the gradient descent algorithm that lead the editing process has been changed to use *Matlab's fmincon* with the *sqp* algorithm. This has made the system faster and more robust, as this method is able to avoid local minimum values.

In the next and last section we show our results and discuss our limitations, explaining their causes and setting probable future lines of work to solve them.

6. Results

As we have shown in Figure 1, the soft PCA basis is good enough to depict diffuse materials from MERL's BRDF database. This solves the BRDF representation problem, but is not enough to fix the edition drawbacks of the previous work. Thus, both representations are needed to correctly depict diffuse and specular materials. This has been achieved through a novel specular measurement, based on the information present on Nielsen most significant direction for BRDF acquisitions. Once the representation that is needed is known, navigation between the BRDF representations can be achieved through trained RBFNs. For an automation of the PCA basis frontier, we use the specular ratio to train two SVMs that designate the PCA space zone where a representation behaves correctly. We integrate this technique in a material editing plugin for a better user experience. We have overcome the spotted weakness of the previous model. Next we present our results as well as the limitations found to our model.

We have found out that the two representations technique is less robust than only using full PCA. This means that during the editing process, a one way path will work fine (changing the value of the attribute once), but constantly moving forward and backward from one representation to another will eventually lead to some kind of artifacts due to the small misalignment caused by the RBFNs that change the PCA coefficients from one basis to the other. The full PCA representation of Serrano et al. is more resistant to these type of constant changes.

During the edition, a convex hull acts as a solid boundary to limit the part of the full PCA space where the editions can take part, thus avoiding artifacts to appear. This is a conservative approach, since any BRDF inside the convex hull will keep its physical plausibility (as the space of mathematically valid BRDFs is convex [WAKB09], [MPBM03]). But there is no reason why there might not be other novel and plausible BRDFs outside the convex hull. As long as their physical properties were assured, the user should be able to explore a wider space than he is now.

Nonetheless, the improvement over the past limitations is plausible. Figure 11 shows a successful edition with a change of representation from soft to full PCA. None of the extremes would've been possible if only one representation was to be used. On the plugin interface we can see how the same BRDF behaves differently on both representations. As the material gets more glossy (more specular), the corresponding point of the 5D PCA space for the soft PCA basis moves to a zone where the BRDFs are no longer correctly represented, while the same BRDF on the full PCA basis is correctly represented and stands inside the convex hull. Figure 10 shows the opposite case, a change of representation from full PCA

to soft PCA. Figure 12 shows other examples of BRDF editions with our plugin.




Figure 10: Sample of an edition with change of PCA basis (from full to soft) of the MERL BRDF chrome. The attribute "metallic" is being modified, from original to lowest possible values. The left image corresponds to a soft PCA basis representation (low metallic value), while the right image corresponds to the original BRDF (reconstructed from its full PCA coefficients).

7. Discussion and conclusion

With our work we have shown that where a type of material is under represented, a separate PCA basis can be made to represent it better, as long as it shares a common boundary with the other (or others) PCA basis. This could be useful when no more samples can be acquired, or when the huge difference between material types biases the PCA basis towards one of them. We have also presented a novel specular measure for tabulated BRDFs, that is not restricted to a type of material, nor expensive to use.

Yet there are open questions and many opportunities for interesting future work. For example, a further analysis of the differences on the PCA bases would give us more information of their behaviour. Figure 3 shows that both bases are indeed different on their principal components, but we don't certainly know what these differences mean or how we can take advantage of them.

Another interesting line of work would be exploring other methods to navigate between the representations. We know that some misalignments exist when navigating between the PCA bases using our trained RBFNs. Studying other possible navigation methods could be useful to minimize these misalignments.

Finally, we would like to broaden the PCA space where the user can make the editions. As we've mentioned before, this is now limited by the boundaries of a convex hull. The use of the convex hull could be replaced by ensuring that the physical properties of the BRDF are met. From the three principal physical properties of a BRDF (non-negativity, reciprocity and energy conservation) the first two are assured by implementation details. Working on the third property would open the possibility of new materials for the user.

8. Acknowledgments

This research has been partially funded by an ERC Consolidator Grant (project CHAMELEON), and the Spanish Ministry of Economy and Competitiveness (projects TIN2016-78753-P and

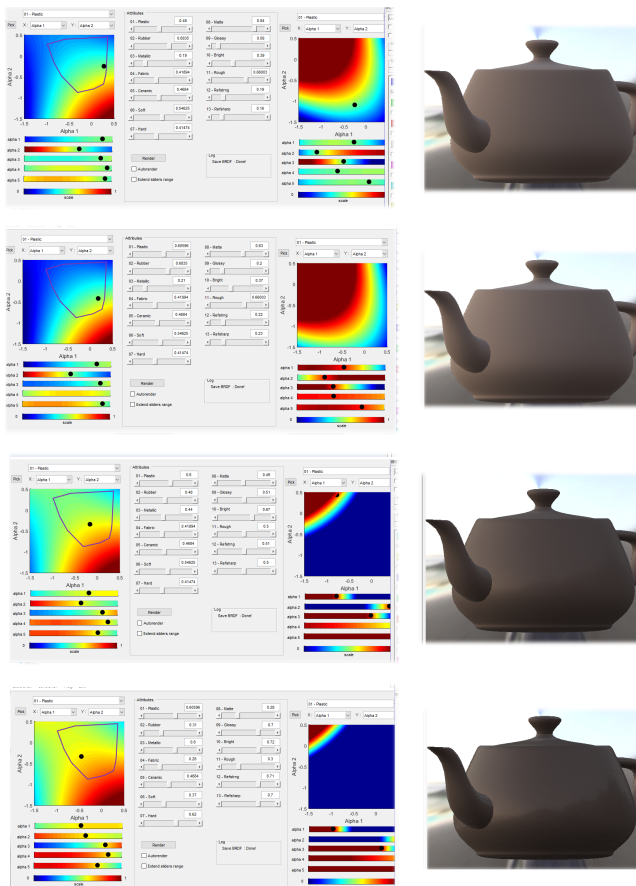


Figure 11: Sample of a edition of the MERL BRDF beige-fabric, with a change from soft to full PCA basis. The attribute "glossy" is being modified, from lowest (top) to highest (down). On the left, the editing plugin can be seen. Please note how the PCA slices are modified along the edition. On the right, the change of the BRDF appearance is shown.

TIN2016-79710-P). Ana Serrano was supported by an FPI grant from the Spanish Ministry of Economy and Competitiveness.

References

- [Alm87] ALMAN D. H.: Directional color measurement of metallic flake finishes. In *In Proceedings of the ISCC Williamsburg Conference on Appearance* (1987), pp. 53–56. 2, 4
- [Alt] Alta: a brdf analysis library. <http://alta.gforge.inria.fr/>. 4
- [APS00] ASHIKHMIN M., PREMOŽE S., SHIRLEY P.: A Microfacet-based BRDF Generator. In *Proc. of SIGGRAPH'00* (New York, NY, USA, 2000), SIGGRAPH '00, ACM Press/Addison-Wesley Publishing Co., pp. 65–74. 2
- [ATDP11] AN X., TONG X., DENNING J. D., PELLACINI F.: Appwarp: Retargeting measured materials by appearance-space warping. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 147:1–147:10. 1
- [AWL15] AITTALA M., WEYRICH T., LEHTINEN J.: Two-shot




Figure 12: Results of the editions on several BRDFs. The initial BRDF is shown on the left, and the edition result on the right. From top to bottom, these editions are: blue-fabric, increasing metallic attribute (change from soft to full PCA basis), violet rubber, increasing sharpness of reflections (change from soft to full PCA basis) and delrin, decreasing brightness.

- SVBRDF capture for stationary materials. *ACM Trans. Graph.* 34, 4 (July 2015), 110:1–13. 1
- [BAOR06] BEN-ARTZI A., OVERBECK R., RAMAMOORTHY R.: Real-time BRDF editing in complex lighting. *ACM Trans. Graph.* 25, 3 (July 2006), 945–954. 2
- [BBPA15] BOYADZHIEV I., BALA K., PARIS S., ADELSON E.: Band-sifting decomposition for image-based material editing. *ACM Trans. Graph.* 34, 5 (Oct. 2015), 163. 2
- [Bli77] BLINN J. F.: Models of light reflection for computer synthesized pictures. *SIGGRAPH Comput. Graph.* 11, 2 (July 1977), 192–198. 4
- [BS87] BECKMANN P., SPIZZICHINO A.: *The scattering of electromagnetic waves from rough surfaces*. Norwood, MA : Artech House, 1987. 4
- [Bur12] BURLEY B.: Practical Physically Based Shading in Film and Game Production - Physically Based Shading at Disney - Course Notes. In *SIGGRAPH Courses* (2012). 1
- [BUSB13] BELL S., UPCHURCH P., SNAVELY N., BALA K.: Open-surfaces: a richly annotated catalog of surface appearance. *ACM Trans. Graph.* 32, 4 (July 2013), 111:1–111:17. 2
- [CP06] COLBERT M., PATTANAİK S.: BRDF-Shop: Creating Physically Correct Bidirectional Reflectance Distribution Functions. *IEEE Computer Graphics and Applications* (2006), 30–36. 2
- [CPWAP08] CHESLACK-POSTAVA E., WANG R., AKERLUND O., PEL-

- LACINI F.: Fast, realistic lighting and material design using nonlinear cut approximation. *ACM Trans. Graph.* 27, 5 (Dec. 2008), 128:1–128:10. 2
- [EKM01] ERSHOV S., KOLCHIN K., MYSZKOWSKI K.: A realistic lighting model for computer animators. *Computer Graphics Forum* 20, 3 (2001). 2
- [FVH14] FILIP J., VAVRA R., HAVLICEK M.: Effective acquisition of dense anisotropic BRDF. In *Proc. of ICPR 2014* (August 2014), pp. 2047–2052. 1
- [GAGH14] GARCES E., AGARWALA A., GUTIERREZ D., HERTZMANN A.: A similarity measure for illustration style. *ACM Trans. Graph.* 33, 4 (July 2014). 2
- [GKLM11] GUY S. J., KIM S., LIN M. C., MANOCHA D.: Simulating heterogeneous crowd behaviors using personality trait theory. In *Proc. of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2011), SCA '11, pp. 43–52. 2
- [GXZ*13] GKIIOULEKAS I., XIAO B., ZHAO S., ADELSON E. H., ZICKLER T., BALA K.: Understanding the role of phase function in translucent appearance. *ACM Trans. Graph.* 32, 5 (Oct. 2013), 147:1–147:19. 2
- [HH87] HUNTER R. S., HAROLD R. W.: *The Measurement of Appearance (2nd Edition)*. Wiley, 1987. 2
- [JMB*14] JARABO A., MASIA B., BOUSSEAU A., PELLACINI F., GUTIERREZ D.: How do people edit light fields? *ACM Transactions on Graphics (SIGGRAPH 2014)* 33, 4 (2014). 2
- [KP10] KERR W. B., PELLACINI F.: Toward evaluating material design interface paradigms for novice users. *ACM Trans. Graph.* 29, 4 (July 2010), 35:1–35:10. 2
- [KSI14] KOYAMA Y., SAKAMOTO D., IGARASHI T.: Crowd-powered parameter analysis for visual design exploration. In *Proc. of the 27th Annual ACM Symposium on User Interface Software and Technology* (2014), UIST '14, pp. 65–74. 2
- [LBAD*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHY R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.* 25, 3 (July 2006), 735–745. 2
- [LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1997), SIGGRAPH 1997, ACM Press Addison-Wesley Publishing Co., pp. 117–126. 4
- [MMG11] MCNAMARA A., MANIA K., GUTIERREZ D.: Perception in graphics, visualization, virtual environments and animation. In *SIGGRAPH Asia 2011 Courses* (2011). 2
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph.* 22, 3 (July 2003), 759–769. 1, 2, 3, 4, 8
- [NDM06] NGAN A., DURAND F., MATUSIK W.: Image-driven Navigation of Analytical BRDF Models. In *Proc. of EGSR'06* (2006). 2
- [NJR15] NIELSEN J. B., JENSEN H. W., RAMAMOORTHY R.: On Optimal, Minimal BRDF Sampling for Reflectance Acquisition. *ACM Trans. Graph.* 34, 6 (Nov. 2015). 1, 2, 3, 5
- [NKLN10] NGUYEN C. H., KYUNG M.-H., LEE J.-H., NAM S.-W.: A PCA Decomposition for Real-time BRDF Editing and Relighting with Global Illumination. In *Eurographics Symposium on Rendering* (2010), pp. 1469–1478. 2
- [OLAH14] O'DONOVAN P., LIBEKS J., AGARWALA A., HERTZMANN A.: Exploratory font selection using crowdsourced attributes. *ACM Trans. Graph.* 33, 4 (July 2014), 92:1–92:9. 2
- [PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. In *Proc. of SIGGRAPH'00* (2000), pp. 55–64. 2
- [Rus98] RUSINKIEWICZ S.: A new change of variables for efficient BRDF representation. In *Rendering Techniques (Proc. Eurographics Workshop on Rendering)* (June 1998). 5
- [SGM*16] SERRANO A., GUTIERREZ D., MYSZKOWSKI K., SEIDEL H.-P., MASIA B.: An intuitive control space for material appearance. *ACM Transactions on Graphics (SIGGRAPH ASIA 2016)* 35, 6 (2016). 1, 2, 3, 4, 6
- [SGVH90] SARIS H. J. A., GOTTENBOS R. J. B., VAN HOUWELINGEN H.: Correlation between visual and instrumental colour differences of metallic paint films. *Color Research and Application* 15, 4 (1990), 200–205. 2, 4
- [SMD*15] SIGAL L., MAHLER M., DIAZ S., MCINTOSH K., CARTER E., RICHARDS T., HODGINS J.: A perceptual control space for garment simulation. *ACM Trans. Graph.* 34, 4 (July 2015), 117:1–117:10. 2
- [SZC*07] SUN X., ZHOU K., CHEN Y., LIN S., SHI J., GUO B.: Interactive Relighting with Dynamic BRDFs. *ACM Trans. Graph.* 26, 3 (July 2007). 2
- [TGY*09] TALTON J. O., GIBSON D., YANG L., HANRAHAN P., KOLTUN V.: Exploratory modeling with collaborative design spaces. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 167:1–167:10. 2
- [VLD07] VANGORP P., LAURIJSEN J., DUTRÉ P.: The influence of shape on the perception of material reflectance. *ACM Trans. Graph.* 26, 3 (July 2007). 2
- [WAKB09] WILLS J., AGARWAL S., KRIEGMAN D., BELONGIE S.: Toward a perceptual space for gloss. *ACM Trans. Graph.* 28, 4 (Sept. 2009), 103:1–103:15. 1, 2, 8
- [War92] WARD G. J.: Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 265–272. 4
- [WM01] WESTLUND H., MEYER G.: *Applying appearance standards to light reflection models*. Association for Computing Machinery, 2001, pp. 501–510. 2, 4