

A Low-Dimensional Perceptual Space for Intuitive BRDF Editing

Weiqi Shi¹, Zeyu Wang¹, Cyril Soler² and Holly Rushmeier¹

¹Yale University, USA

²INRIA, Grenoble University, France

Abstract

Understanding and characterizing material appearance based on human perception is challenging because of the high-dimensionality and nonlinearity of reflectance data. We refer to the process of identifying specific characteristics of material appearance within the same category as material estimation, in contrast to material categorization which focuses on identifying inter-category differences [FNG15]. In this paper, we present a method to simulate the material estimation process based on human perception. We create a continuous perceptual space for measured tabulated data based on its underlying low-dimensional manifold. Unlike many previous works that only address individual perceptual attributes (such as gloss), we focus on extracting all possible dimensions that can explain the perceived differences between appearances. Additionally, we propose a new material editing interface that combines image navigation and sliders to visualize each perceptual dimension and facilitate the editing of tabulated BRDFs. We conduct a user study to evaluate the efficacy of the perceptual space and the interface in terms of appearance matching.

CCS Concepts

- Computing methodologies → Perception;

1. Introduction

Real-world materials display complex and diverse appearances and understanding how humans perceive them is not an easy task. According to [FNG15], the underlying visual processing and human perception of material appearance can be grouped into two kinds of computations: categorization and estimation. While material categorization has been widely studied in both computer graphics and vision literature [VZ08; SSKK15; BUSB15; GVPV17], not much research has been conducted to understand the process of material estimation. In other words, how to characterize the perceptual attributes of a given material appearance is still an open question.

Generally, material appearance can be represented by either analytical models or tabulated measured reflectance data. In this project, we want to focus on the latter because it captures the rich details of appearance and is not limited to a particular gamut. Previous works on studying the perception of measured data focus on either a single individual attribute such as gloss [WAKB09; PFG00] or translucency [GXZ*13; FB05], which ignores other possible dimensions that might influence perception. Other works use linear models to project the measured data to a subspace [MPBM03; SGM*16], which still requires a relatively large number of dimensions to faithfully reconstruct the input BRDF.

In this paper, we want to study the material estimation process for measured material data by building a low-dimension perceptual space. We focus on one specific category: metal materials, which is one of the four major groups ("metal", "Ceramics", "Polymers" and

"Composites") defined in material science based on chemical and mechanical structures [MJ20]. The reason to focus on one category is to minimize the possible influence of the categorization process in perception and inter-category differences. Narrowing down to one category allows us to focus on the relatively subtle intra-category differences between material samples. Also, most metal materials exhibit different levels of specularity and shininess, which makes this category unique and diversified. From the design perspective, many material design packages provide users pre-defined material "templates" (such as Blenderkit [Ble02]) to help them focus on the editing of specific material categories. Therefore, characterizing material samples within each category is important to make the design process intuitive and efficient.

We start by collecting crowdsourced data based on a psychophysical experiment, where subjects were asked to select the two images that have similar material appearances from a triplet. Given human-labeled data, we apply the Non-metric Multi-Dimensional Scaling (NMDS) algorithm on the pairwise comparisons to compute a low-dimension embedding that represents the positions of each material sample in a perception manifold. We carefully evaluate the embedding and interpolate the measured data using Gaussian Process Regression (GPR) to create a continuous perceptual space. We use GPR instead of a linear model to interpolate in the inherently non-linear perceptual manifold, while preserving linear dependency on the measured data with nice practical consequences on efficiently rendering the interpolated materials [SSN18]. Compared with other

costly non-linear predicting models, this property makes it possible to efficiently edit appearance using the proposed perceptual space.

In addition to the perceptual space, we propose a new material editing interface that combines image navigation and sliders together. The proposed interface allows users to create BRDFs in tabulated form, which maintains the subtle characteristics of measured materials and guarantees a high degree of realism. To avoid the ambiguity of perceptual terminology, we do not name each dimension of the perceptual space. Instead, we use an image navigation tool to provide the visualization of appearance variation along each dimension to help the user interpret the meanings of each dimension by themselves. To test how the perceptual space and the proposed interface could contribute to material editing, we conduct a user study of appearance matching and evaluate performance in terms of time to completion, matching errors, and self-satisfaction, etc. Interestingly, we find that (1) the interface with both image navigation and sliders outperforms the one with only sliders, and (2) users have better performance using the perceptual space than an analytical model. These two observations are in conflict with the conclusions in [KP10] which will be discussed later. In summary, our contribution includes:

- A low-dimensional embedding built on crowdsourced data to understand the material estimation and characterization process;
- a continuous perceptual space for tabulated metal material data that can be used for editing;
- an interface that combines image navigation and sliders, which outperforms the sliders-only interface according to our user study.

2. Related Works

2.1. Human Perception to Material Appearance

According to Dorsey et al. [DRS10], material appearance can be defined as “the visual impression we have for a material”. There have been many previous works focusing on understanding the general process of human perception of material appearance, and high-level reviews can be found in [And11; Fle14; FNG15; MB10]. Most of the work that has been done in the material perception area can be roughly categorized into three directions: understanding individual factors in perception (gloss, translucency, shape, and so on), defining metrics to compare BRDFs perceptually, and building perceptual spaces using low-dimensional embedding.

Understanding Individual Factors on Perception. Studying material perception in terms of a particular property can help understand a certain type of material. Throughout the years, many research projects have been done to study the perception of gloss, such as [CK15; WAKB09; PFG00]. Typically, they use crowdsourced data to fit analytical expressions and create a perceptual parameterization of analytical BRDF models. Translucency is another perceptual attribute that is widely studied. Gkioulekas et al. [GWA*15; GXZ*13] provide thorough studies to understand different the impact of phase functions and edges on the perception of translucent materials. Gigilashvili et al. [GTHP18; GUT*19] study the influence of geometry thickness on the human perception of translucency. Many other internal and external factors have also been studied to understand their impact on perceiving material appearance, such as Fresnel effects [Fau19], texture [LM01], viscosity [VBF18], shape

[HFM16; VLD07], illumination [HLM06; VBF17; KFB10] and motion [DFY*11].

Perceptual BRDF Metrics. Measuring the perceived difference between BRDFs is important for material recognition and classification. Many previous works have proposed different metrics to compare BRDFs based on human perception. Fores et al. [FFG12] and Ngan et al. [NDM05] put forward different perceptual metrics to evaluate measured BRDFs directly. Ngan et al. [NDM06], Pereira et al. [PR12] and Sun et al. [SSGM17] focus their work on evaluating and comparing BRDFs in image space. They use images with rendered materials as a medium to build the metrics and compute the difference between material samples. Another promising direction is to define the metrics based on the BRDF embedding, such as using MDS [PFG00] or PCA [SGM*16]. Lagunas et al. [LMS*19] proposed a learning-based solution for the overall material appearance to derive a similarity measure that correlates with the notion of material similarity as perceived by humans. Recently, both Serrano et al. [SCW*21] and Lagunas et al. [LSGM21] study the joint effect of geometry and illumination on material perception.

BRDF Embedding and Perceptual Space. Finding a low-dimensional representation for the high-dimensional BRDF data is the first step for building a perceptual space for material appearances. Many works have performed BRDF analysis and provided compact representations, such as spherical harmonics [WAT92] and spherical wavelets [SS95]. However, to understand human perception of materials, we need to know the relative position of each material sample in the perceptual space. Therefore, the analysis should focus on the entire space of BRDFs. Typical works include linear dimension reduction approaches such as PCA [NJR15] and nonlinear methods such as MDS. Many works build a perceptual space using linear dimension reduction approaches such as [WAKB09; MPBM03]. However, according to [MPBM03] linear approaches are unable to identify a sufficiently small subspace to facilitate practical exploration, while nonlinear methods can only generate compact embeddings without explicitly providing mappings between the measured space and the manifold [SSN18]. The study closest to our work is [SGM*16]. They create a perceptual control space by asking users to rate material appearances with 14 adjectives and then map the crowdsourced data to the principal components of the measured BRDFs using Radial Basis Functions. Different from their work that analyzes perceptual similarity based on a *a priori* categorization of perceptual traits (and possibly ambiguous English words), we collect data based on the perceived similarity of appearance and determine the most significant perceptual dimensions as a post process. Toscani [TGG*20] conduct a user study using a larger range of stimuli and build a perceptual space based on the ranking results of the similarity between the candidates and reference.

2.2. Material Editing

Many works have been devoted to the design and development of material editing applications due to the emergence of different shading models and user interaction mechanisms. We will consider material editing models from three perspectives: physical-based models, perceptual-based models, and the editing interface.

Editing Materials based on Physical Models. Currently, many

off-the-shelf software products use physical-based shading models as the foundation for material editing, such as Maya [Aut17] and Blender[Ble02]. Users can change the material appearance by tweaking the parameters of analytical BRDF models. BRDF-shop [CPK06] is a prototype that allows users to design an artist-friendly editing framework based on an extension of the Ward model. Taltton et al. [Kol09] explore the parameter space of the anisotropic Ashikhmin mode and use it to create an editing system. Shi et al. [SDR20] propose a learning-based solution to edit bi-scale materials for fabrication purposes. There are many other works that focus on efficient editing of analytical BRDF models with fast feedback for lighting changes [NKL10; CWAP08; SZC*07]. Hu et al. [HGC*20] present a deep-learning-based representation that can reduce the dimensionality of measured BRDFs for material editing. While using physical models to edit the appearance can provide accurate simulation results, the parametric space is not intuitive and consistent with human perception. The change of specific parameters may lead to subtle appearance changes that humans cannot detect.

Editing Materials based on Perceptual Models. With more and more research studying how the human visual system perceives material appearance, editing materials in perceptual space becomes possible. Pellacini et al. [PFG00], Westlund et al. [WM01] and Wills et al. [WAKB09] propose different applications to edit the gloss attribute of the material. Boyadzhiev et al. [BBPA15] develop an image-based material editing system using intuitive attributes. As mentioned previously, Serrano et al. [SGM*16] propose an intuitive control space based on the perceptual space they have developed. Mylo et al. [MGZ*17] introduce an interface to edit spatially-varying material by connecting a link between certain perceived visual properties and specific bands in its spectrum of spatial frequencies of a wavelet decomposition.

Editing Interface and User Interaction. Besides shading models, the user interface and platform also play an important role in the editing process. Most of the material editing systems mentioned so far are slider based interfaces, where users can change the parameters of material model by dragging sliders. However, the disadvantage of the slider based models is that users can only access the results in the current state represented by the sliders. Users can be easily confused by different dimensions and forget the big picture of the entire space. Many other interfaces and prototypes are also developed to improve the intuitiveness of the material editing. When the parameter space is of very low dimension, a map-based navigation is possible [SSN18]—although in this particular case no perceptual interpretation of the parameters was provided. A common technique is image navigation, which provides users with a group of materials with certain variations. Instead of modifying parameters, users can browse through the materials and pick the one that is closest to their goals. Other lines of work [MAB*97][NDM06] and software [ADO09] further explored this idea and developed interfaces based on it. Another example is the Augmented Reality material editing interface proposed by Shi et al. [SWS*17], where the real environment rather than synthetic images is provided as background to visualize appearance. Kerr et al. [KP10] introduce a user study to compare slider based systems and image navigation. Different from previous work, we propose a novel interface that combines the sliders with a variation of the image navigation sys-



Figure 1: The pipeline of our method. We collect triplet comparison data from a user study and convert them into perceptual embedding using NMDS. Next we use Gaussian Process Regression to connect the embedding and the tabulated BRDF to build a continuous perceptual space, based on which we propose a new interface for editing.

tem to help users understand the parameter space of the perceptual dimensions.

3. Overview

Before we provide technical details, we first define the goal and scope of our work. According to Fleming et al. [Fle14], the underlying visual processing of material appearance can be broadly grouped into two computations: categorization and estimation. Categorization is to assign a specific class label to a given material, while estimation is to identify the characteristics of the material. Our work focuses on the latter part. Our goal is to model a perceptual space to "locate" each material sample based on its surface reflectance properties and understand its relative position within its own class. Then we can visualize the space, interpret each dimension based on the visual clues, and edit the appearance by navigating through the perceptual dimensions. Due to the large number of material categories and their distinctive appearance, we narrow our study to the measured metal related materials collected from different datasets. We acknowledge that the conclusions and observations may be only applied to this specific material category. However, we believe that our method can be easily adapted to other categories in the future.

We first introduce the process of the user study to collect crowdsourced data using Mechanical Turk in Section 4. In Section 5, we describe learning a perceptual embedding using the NMDS model with crowdsourced data. We introduce Gaussian Process Regression to interpolate the perceptual embedding and build a perceptual space in Section 6. In Section 7, we introduce a novel material editing interface using the proposed perceptual space and conduct a user study to evaluate its performance. We conclude with a discussion in Section 8.

4. Crowdsource Data Collection

In this section, we describe the design of a crowdsourcing experiment to collect user data on material appearance perception. We adopt the Two-Alternative Forced Choice (2AFC) scheme to ask users to select two similar images from a triplet.

Stimuli. We selected measured data for 39 materials categorized as metal from multiple datasets [MPBM03; DJ18]. The selected data covers a large range of appearance within the metal category. We chose the scene Havran-2 [HFM16] to render the selected material samples since it densely samples the incoming and outgoing light directions and maximizes the information relevant for material appearance judgments. We used the Ennis environment map [FDA03] for illumination and render the stimuli with Mitsuba [Jak10]. Fig. 2 shows some examples of the stimuli and the full dataset can be found in supplemental materials. The stimuli were generated using




Figure 2: Example of stimuli in our psychophysical experiment. The Havran-2 scene [HFM16] is used to maximize the information relevant for material appearance judgments.

the same geometry under the same illumination and the only difference is the material. Therefore we believe the image similarity represents the material similarity. Since our focus is the material perception caused by surface reflectance instead of color, we generated the achromatic stimuli by averaging the RGB channels to avoid bias for the user study. We note that the color channels can be easily integrated in the future by adding the a and b channels of the CIELAB color space, as proposed in [WAKB09; SGM*16]. Also as we will introduce later, Gaussian Process Regression can be applied to interpolate colors if the stimuli are chromatic in the future.

Participants. We used Amazon Mechanical Turk to recruit participants. A total of 360 paid participants took part in our experiment. Users were not aware of the purpose of the experiment.

Procedure. We decided to use a modified 2AFC scheme instead of the ranking method used in [SGM*16] for our experiment, because of its simplicity and consistency in terms of providing numerical distances between stimuli [WAKB09]. Also, the perceptual space may be multi-dimensional and cannot be represented with a linear scale [LMS*19]. During the experiment, triplet images of the stimuli were presented to the subjects and they were asked to select any two images that share a similar appearance. Notice the method we used was slightly different from the traditional 2AFC where subjects only need to select from the two alternative candidates to match the reference. In the feedback from a pilot study using traditional 2AFC, we observed that subjects tend to choose randomly when they believe the two candidates are more similar to each other than to the reference. Therefore, we add that as a third option to cover all possible cases and reduce noise in the data.

Inspired by [TLB*11; LMS*19], we adopted an adaptive sampling scheme to reduce the total number of triplets in our experiment. Ideally, 39 different stimuli yield around $39 \times \binom{38}{2} \approx 27k$ comparisons, which would require a total of 150k responses if each comparison is evaluated by more than 5 subjects. However, this scale of the experiment is not feasible. Therefore, we turned to an adaptive sampling scheme to reduce the total comparisons to 6k. The sampling scheme selects a batch of triplets to maximize the information gain based on the previous iterations. In our case, we conducted 15 iterations and sampled 400 triplets for each iteration. The mean information gain reduced to 10^{-6} after the 15 iterations, confirming the convergence of the sampling scheme.

We conducted the experiment on Amazon Mechanical Turk. Before the subjects started the actual experiment, they were provided with a training session that included an additional 10 triplet tests with obvious answers. Subjects who failed the training session (pro-

vided more than 2 wrong answers) were not allowed to start the experiment. Each qualified subject was asked to finish 110 triplets (100 tests and 10 validation) in total without a time limit. The 10 validation triplets were randomly selected from the 100 test triplets and the images were presented in a different order. Subjects were not aware of the validation triplets. If the results from the validation triplets did not reach 80% consistency, all test triplets from the same subject would be rejected. We introduced a delay mechanism for each triplet: the users could only make actions three seconds after the triplet was displayed. This mechanism makes sure that subjects take time to examine the images instead of making fast random clicks. In the end, there were 301 out of 360 subjects who provided valid results, yielding 30,100 responses.

5. Learning A Perceptual Embedding

Given the results from the user study, our next step is to convert the triplet comparison into a low-dimensional embedding for each material sample that correctly represents the participants' answers. Specifically, the Euclidean distance between the embedding of each BRDF in the latent space should be largely consistent with the perceptual distance obtained from the triplet comparisons. Finding the embedding is an instance of the problem of multidimensional scaling (MDS). Non-metric multidimensional scaling (NMDS) is a superset of MDS that finds an embedding only based on the relative ordering of input dissimilarities and doesn't require collecting accurate magnitude information.

In this project we seek a solution using NMDS. We note that Lagunas et al. [LMS*19] propose a learning-based solution to extract 128 dimension feature vectors from triplet comparisons to predict the perceptual similarity in material appearance. However, their solution does not support intuitively interpreting and editing the high-dimensional features.

5.1. Non-metric Multidimensional Scaling

We decided to use NMDS from [AWC*07] to match the triplet comparison. We will briefly recall their algorithm and explain how it applies in our situation.

We start from the triplet (i, j, k) representing the indices of material samples answered by subjects in our user study. We denote \mathbf{x}_i as the embedding point for a material sample i , the matrix \mathbf{X} as embedding coordinates for all the material samples (one column per material), and $D_{i,j}$ as the measured dissimilarity between materials i and j . With these notations we define the set S of answers from the study to be

$$S = \{(i, j, k) \mid D_{i,j} < D_{i,k}, D_{i,j} < D_{j,k}\}. \quad (1)$$

Since each triplet is evaluated by at least 5 subjects in our experiment, the output is robust to inconsistencies and repetitions between the set S . Non-metric dimensional scaling only deals with finding the points \mathbf{x}_i to preserve the order of objects, so we can afford to use the square of the dissimilarity to replace the original one for practical reasons. After a solution is found, the Euclidean distance between the \mathbf{x}_i should match the dissimilarities above, which allows us to write the dissimilarity in terms of the Gram matrix $\mathbf{K} = \mathbf{X}^T \mathbf{X}$

as follows:

$$\begin{aligned} D_{i,j}^2 &= \|\mathbf{x}_i - \mathbf{x}_j\|_2^2 = \mathbf{x}_i^T \mathbf{x}_i - 2\mathbf{x}_i^T \mathbf{x}_j + \mathbf{x}_j^T \mathbf{x}_j \\ &= \mathbf{K}_{ii} - 2\mathbf{K}_{ij} + \mathbf{K}_{jj} \end{aligned} \quad (2)$$

\mathbf{K}_{ij} represents the element in row i and column j from the Gram matrix \mathbf{K} . Given the above conditions, the inequalities defining S become

$$\begin{aligned} D_{i,j}^2 &< D_{j,k}^2 \\ \mathbf{K}_{ii} - 2\mathbf{K}_{ij} + \mathbf{K}_{jj} &< \mathbf{K}_{jj} - 2\mathbf{K}_{jk} + \mathbf{K}_{kk}. \end{aligned} \quad (3)$$

A sufficient condition for Equation 3 is that

$$\mathbf{K}_{ii} - 2\mathbf{K}_{ij} + \mathbf{K}_{jj} + 1 \leq \mathbf{K}_{jj} - 2\mathbf{K}_{jk} + \mathbf{K}_{kk}. \quad (4)$$

This condition enforces an arbitrary but fixed difference between pairs of distances, hence preventing the solution getting arbitrarily close to $\mathbf{K} = 0$. Note that we could choose any other constant number but it would only lead to a uniform scaling of the embedding.

Besides the scaling ambiguity, Equation 4 is also oblivious to translations simultaneously applied to all points \mathbf{x}_i of the embedding. The simplest way to remove the translation ambiguity is to impose the center of the embedding to be at the origin, which is written

$$\sum_i (\sum_j \mathbf{X}_{ij})^2 = 0, \quad \text{or equivalently} \quad \sum_{jk} \sum_i \mathbf{X}_{ij} \mathbf{X}_{ik} = 0. \quad (5)$$

This turns into a constraint on the Gram matrix:

$$\sum_{jk} \mathbf{K}_{jk} = 0. \quad (6)$$

Our goal is to find a low-dimensional embedding from the triplets. The dimension of that embedding is the rank of matrix \mathbf{X} , which also equals the rank of the Gram matrix \mathbf{K} [HJ12]. To deal with the non-convexity of the objective function, we relax the rank function to its convex envelope—the trace of \mathbf{K} which is positive by construction. A detailed discussion about this relaxation can be found in [WAKB09].

Finally, to handle inconsistent measurements from the same triplet, we follow the method proposed in [AWC*07] to introduce slack variables ξ_{ijk} in every inequality constraint, which allows for violations of the inequality and augments the objective function to minimize the total violation. We also introduce a parameter λ to control the trade-off between the violation and the rank of the matrix.

Combining all these conditions, our objective is to find:

$$\mathbf{K}, \xi_{ijk} = \underset{\mathbf{K}, \xi}{\operatorname{argmin}} \quad \lambda \operatorname{Tr}(\mathbf{K}) + \sum_{(i,j,k) \in S} \xi_{ijk} \quad (7)$$

such that

$$\begin{aligned} \forall (i, j, k) \in S \quad & \mathbf{K}_{kk} - \mathbf{K}_{ii} + 2\mathbf{K}_{ij} - 2\mathbf{K}_{jk} \geq 1 - \xi_{ijk} \\ & \mathbf{K}_{kk} - \mathbf{K}_{ii} + 2\mathbf{K}_{ij} - 2\mathbf{K}_{ik} \geq 1 - \xi_{ijk} \\ & \sum_{bc} \mathbf{K}_{bc} = 0 \\ & \xi_{ijk} \geq 0 \\ & \mathbf{K} \geq 0 \end{aligned} \quad (8)$$

The last condition forces \mathbf{K} to be definite positive. We require




Figure 3: Training and testing errors (percentage of violation) for different λ in Cross validation. The best performance of the model is reached when $\lambda = 26$ with lowest testing error and relatively low training error.

that \mathbf{K} is symmetric by optimizing only one half-triangular part. These ensure that we can recover \mathbf{X} later on. Once the optimization is finished, we compute the embedding from the symmetric definite positive matrix \mathbf{K} by applying the eigen-decomposition:

$$\mathbf{K} = \mathbf{U}^T \mathbf{\Sigma} \mathbf{U} \quad \text{and then} \quad \mathbf{X} = \mathbf{U} \mathbf{\Sigma}^{\frac{1}{2}}. \quad (9)$$

To implement, we use Matlab with the SeDuMi solver [Stu99] to solve the optimization problem. The SeDuMi is a convex optimization solver that is designed to solve semi-definite problems. The average convergence time for our optimization problem is around 200 seconds on a standard PC with an AMD Ryzen 7 1700X 8-core CPU and 16 GB RAM.

5.2. Model Selection

In this section, we want to evaluate the choice of λ and the dimension of the embedding. In Equation 7 λ represents a tradeoff between the rank of the Gram matrix \mathbf{K} and the number of violations in the triplets due to the conflict responses from different subjects.

We conduct cross-validation to perform the model selection and evaluate the influence of λ on the inference accuracy. We split the subjects' responses into the training set and validation set by the ratio of 4:1. We apply the cross-validation 10 times for each value of the λ from 0 to 100, defining the error of each model as the percentage of mismatches between the inference results and the majority user response of the corresponding triplet comparison. Fig. 3 demonstrates the average training and validation errors over 10 cross-validation experiments for different λ values. We expect the training error to monotonically increase with λ because a larger value of λ means the optimization focuses on reducing the rank of matrix \mathbf{K} and allows more violations. When $\lambda = 100$, the optimized rank of \mathbf{K} is 2. However, when $\lambda = 0$, the optimization entirely focuses on reducing the number of violations, leading to smaller training errors while producing a rank of 32—which means possible overfitting. Also, this is not consistent with our goal of building a low-dimensional embedding. The validation error first decreases and then increases. As defined above, the error reaches the minimum point at $\lambda = 26$, representing the model that has the best generalization performance to deal with the violations. At this point, the first four dimensions of the embedding contain 99.2% of the total variance of all dimensions, indicating that we can safely build a 4D perceptual space.

6. Building A Perceptual Space

The perceptual embedding we have learned so far from the user responses provides dimensionality and coordinates of measured material samples on a 4D perceptual parameter space. That technique however does not provide a full mapping of the embedding space. In this section we define a continuous parameterisation space that can map any point in the perceptual space to actual reflectance data. In order to achieve this, we interpolate the material samples with a Gaussian Process Regression model and analyze each interpolated dimension of the perceptual space.

6.1. Perpetual Regression Using Gaussian Processes

We use Gaussian processes as a regression model taking as input the perceptual coordinates and as output the 4M measured BRDF data in the same format as the MERL dataset. The advantages of using a GPR model can be summarized as follows: (1) It provides an interpolating and non-linear mapping between any arbitrary low-dimensional latent space and high-dimensional data; (2) it provides a uniform linear interpolation for the data itself, regardless of the choice of latent variables and their dimension; (3) the manifold is continuous as long as the covariance/kernel function used in the Gaussian process is continuous.

We adapt the GPR model from the Gaussian process latent variable model (a.k.a. GPLVM) introduced in [SSN18] to perform the perceptual interpolation. In our case however, latent variables are not obtained through optimization of the log-likelihood of the Gaussian process, but from the NMDS described in Section 5.1. For the mathematical proofs and derivations of the properties of Gaussian process regression, please refer to [Ras03]. In our case, we denote the $N \times 4M$ matrix \mathbf{Z} to be the matrix of the N measured BRDF data used as stimuli in our experiment with $N = 39$. We denote by \mathbf{X} the $N \times 4$ matrix representing the optimized 4D perceptual embedding computed from the NMDS (which stand here for the coordinates parameterizing the manifold). Given an arbitrary point \mathbf{x}_* in the perceptual space, GP regression predicts the corresponding BRDF \mathbf{z}_* using

$$\begin{aligned} \mathbf{z}_*^\top &= \mathbf{v}_*^\top \mathbf{V}^{-1} \mathbf{Z} \\ \mathbf{v}_* &= [c(\mathbf{x}_0, \mathbf{x}_*), c(\mathbf{x}_1, \mathbf{x}_*), \dots c(\mathbf{x}_{N-1}, \mathbf{x}_*)]^\top. \end{aligned} \quad (10)$$

In the equation above, \mathbf{V} is the covariance matrix whose elements are $V_{ij} = c(\mathbf{x}_i, \mathbf{x}_j)$ for all (i, j) , and c is the covariance function that can be specified by the user and is key to modeling the non-linearity of the interpolant. Following [SSN18], we define c as a shifted Gaussian function because of its smoothness and local support, which leads to smooth transitions for the predicted BRDFs:

$$c(\mathbf{x}, \mathbf{x}') = \mu \delta(\mathbf{x}, \mathbf{x}') + e^{-\frac{||\mathbf{x} - \mathbf{x}'||^2}{2l^2}}, \quad (11)$$

where l and μ are hyperparameters that correspond to the characteristic length scale and noise-filtering parameter respectively. According to [Ras03], higher values for l lead to a smoother manifolds but make the inversion of \mathbf{V} less stable, while a non null value for μ can significantly improve numerical stability for inverting \mathbf{V} . Based on the analysis proposed in [SSN18], we set $\mu = 10^{-4}$. For l , since its value depends on the input data and our embedding are

defined up to a rotation and scale, we can set $l = 1$ and scale the embedding to the same order of magnitude. Note that although \mathbf{v}_* is non-linear with respect to the perceptual embedding, the extrapolated data \mathbf{z}_* is still linear for the measured BRDF data \mathbf{Z} . This allows the rendering of an object shaded with the interpolated BRDF by directly applying Equation 10 to precomputed shaded images for the BRDFs in the database in place of \mathbf{Z} .

6.2. Visualization and Analysis

For visualization, we continuously interpolate BRDFs along each perceptual dimension. Linearity between the BRDFs and the rendered images ensures that the correct appearance can directly be obtained by applying the same Gaussian process regression to rendered images corresponding to the input BRDFs in \mathbf{Z} . The perceptual coordinates of the 39 stimuli provide a bounding box to define the maximum and minimum value for each dimension. For better visualization we interpolate one dimension at a time and fix the coordinates of the other three dimensions at the center. We demonstrate this in Fig. 4

The images in each row depict how the BRDFs vary with the increase of a specific perceptual coordinate. Along Axis 1, the major variation is the specularity. The highlight becomes sharper and the whole appearance becomes glossier. On Axis 2, we can observe the changes in the grazing angles. The edge of the sphere changes from bright to dark along Axis 2. Axis 3 captures the specular shape and anisotropy. Notice the vertical part of the highlight starts with clear and regular dots and ends up as stretchy lines. The variation for Axis 4 depicts the increase of the diffuse part while maintaining the shape and intensity of the specular part.

Notice that the components of the perceptual coordinates are in decreasing order of the variance computed from the Gram matrix \mathbf{K} , which represents the importance of each dimension in terms of explaining the perceptual space. In our case, Axis 1 is the most important dimension that explains how humans perceive the appearances of the stimuli, and Axis 4 is the least important one. This result provides some interesting insights that align with our intuition on how humans perceive different characteristics of the metal materials. The changes of shininess and glossiness in Axis 1 demonstrates that humans are most sensitive to the intensity of highlights when recognizing metal materials. The second important thing is the reflectance caused by the Fresnel effect, which can be explained by human sensitivity towards the edge of an object. In comparison, the appearance variation displayed in Axis 3 is less significant compared with the first two, indicating that the shape of the highlight and anisotropy are subtle clues for human perception. The results in Axis 4 show that when the highlight remains the same, the appearance change of the rest part does not make a significant difference in perception, which further demonstrates the importance of the highlight. However, we note that the last observation could apply to metal materials only since all the materials used for stimuli are glossy to some extent, which outweighs the influence of the diffuse component.

We also want to demonstrate that the proposed perceptual embedding is consistent with the low-dimensional embedding of the BRDF data. We compute the PCA of our 39 stimuli using the method presented in [NJR15] and visualize the scaled first 5 principal components in 2D slice [BS12]. Fig. 5 shows the results. As we can



Figure 4: Visualization of each perceptual dimension that was automatically found by our perceptual analysis, by decreasing order of importance. Images in each row are obtained by interpolating coordinates in each dimension using a Gaussian process regression model.

see the interpreted meaning of each principal components and their importance coincide with our perceptual embedding. This result further confirms the consistency of the underlying embedding between our perceptual model and physical reflectance data. However, we cannot directly map the perceptual embedding to the first 5 principal components of the BRDFs using the GPR model and then apply reconstruction, since the first components cannot fully reconstruct the original appearance, especially for materials with high specularity and anisotropy. Fig. 6 shows such an example. Also, we do not notice a significant visual difference between our method and mapping to high dimensional principal components (such as 35D).

7. Perceptual Editing Interface

With the proposed perceptual space, our next step is to build an intuitive material editing prototype that can take advantage of the low-dimensional embedding. In contrast to previous material editing interfaces, we do not want to represent each dimension with a single word that users may interpret in different ways. Even the simple words that are used to describe physical attributes of materials (such as specularity, roughness, etc) may be abstract for novice users, not to mention the ambiguity and inconsistency caused by different implementations of different analytical BRDF models. For perceptual dimensions, it is even more difficult to find ubiquitous and intuitive words to interpret them. Traditional methods rely on fitting the perceptual dimensions with physical attributes, but this could lead to ambiguity and inconsistency.

Inspired by the image navigation interface proposed by [NDM06], we want to use the appearance variation to define each perceptual dimension. To be more specific, for each dimension, we present to users a series of images with variations corresponding to the uniform steps in our perceptual space. By observing the image variation, users are expected to understand and interpret each perceptual dimension by themselves. However, according to the results from the user study [KP10], we know that compared with the sliders, users have relatively poor performance using the interface with only image navigation in material editing. Therefore, we propose a

trade-off solution that combines the image-navigation with sliders to create an intuitive editing interface. Notice we do not intend to compare our method with [SGM*16] and [HGC*20] because 1) these methods are built based on different user studies with different material datasets as stimuli, 2) the number of controls and the interface mechanisms are significantly different, 3) our focus is to examine the contribution of image navigation in material editing instead of other factors. We believe all these factors would lead to inconclusive comparison results.

7.1. Interface Overview

Fig. 7 gives an example of our user interface, which includes three sections: image navigation (left), sliders (middle) and utilities (right).

The image navigation section is a 4×5 image grid, where each row represents the variations when increasing the value for the corresponding axis. For example, there are 5 different variations uniformly sampled with a fixed step size to demonstrate how the material appearance can vary along the first axis. Their coordinates of the first axis are increasing from left to right. The center image on each row represent the currently selected material. Once an image is selected, it will be presented in the center of each row, and the rest of the images in the grid will be updated iteratively. We provide different step sizes to control the granularity of differences between two adjacent materials on the same row. The 4 axes in our interface represent the 4-dimensional perceptual space. Given this interface, users can navigate through the entire perceptual space incrementally and iteratively. Essentially, it is a simplified 1D version of the original 2D image navigation reported in [NDM06]. However, we believe this simplification is necessary since with the help of the sliders users do not need to deal with axes selection, which is the most confusing part that makes the original image navigation less effective according to [KP10].

We also provide users with a slider to help them quickly locate the desired material and fine-tune its appearance. The slider and




Figure 5: BRDF-slices of the first 5 principal components of the 39 metal BRDF stimuli and a slice reference.




Figure 6: Comparison between our method and mapping perceptual embedding to the first 5 principle components and then perform reconstruction.




Figure 7: Example of our proposed user interface.

image navigation are codependent: changes on either one will lead to the update on the other one. The design concept is to provide users more options and information to interact with the material in a straightforward manner. For example, users can drag sliders to quickly visualize the changes along each axis and directly select the desired images from the grid.

For the utility section, users can visualize their current selected material on the right. For evaluation purposes, we also present a target material appearance for appearance matching tasks. We also provide redo and undo features to help users restore a previous or forward state of the design.

All the sample images are rendered using environment map St. Peter's Basilica [Deb08]. We use spheres as object instead of Havran-2 because it is easier for users to observe multiple images at the same time, and is consistent with existing material design systems.

7.2. User Study

Goal. We seek to evaluate the relative effectiveness of the proposed perceptual space and the interface paradigm for material design in the context of designing realistic materials. To be more specific, we want to

- (a) measure how efficiently and accurately users can perform specific material adjustments using the perceptual dimensions, and
- (b) understand how the image navigation combined with sliders could influence user behaviors for material design.

Interfaces and Implementations. For evaluation purposes, we compare three interfaces: (1) perceptual image navigation with sliders, (2) physical image navigation with sliders and (3) perceptual sliders only, all three of which are detail below. We compare user performance between interface (1) and (2) to evaluate the goal (a), and compare (1) and (3) to explore the goal (b).

Interface (1) is our proposed interface. For interface (2), we use an analytical BRDF model (Lambertian plus GGX model [WMLT07]) to control the appearance. Specifically, to match the gamut between perceptual and analytical models, we fit the measured BRDFs defined in the perceptual space to the anisotropic GGX model [CK17] using the fitting method proposed in [SJR18]. The analytical model is defined by four parameters: diffuse albedo, specular albedo, roughness and anisotropy. We use the fitting method to find the maximum and minimum values of each parameter to build a bounding box of the perceptual space. Then we can uniformly interpolate the parameter values to create a physical space as a comparison for the perceptual space. For interface (3), we use the same perceptual space but simply drop the image navigation tool and only present users with the sliders.

Tasks. In our user study, subjects are asked to finish three tasks on all the interfaces (9 trials), which focus on different reflectance behaviors of the achromatic materials. The first task requires subjects to edit the specular shape of the material. The second task focuses on the general intensity of the reflectance. The last task asks subjects to create the highly specular glass-like appearance. Before all the trials, there is a training task to help user get familiar with the interfaces. In total, subjects need to finish 10 trials and there is no time limit for each trial. Examples of each material task can be found in Fig. 8

Process. Twenty-five subjects participated in the study, 14 of them female. All subjects had normal or corrected-to-normal vision. The user study was conducted remotely and all the candidates were asked to apply first through our system. The valid candidates were selected as subjects and provided with the link to the user study.

During the experiment, each subject needed to first watch a mandatory video that introduced the basic functions of the interfaces. Subjects could not proceed to the next step without finishing




Figure 8: Starting and target materials for training and matching tasks.

the video. Then subjects would start the training trial, where they could play with our interface until they were confident to start the experiment. After they finished all the trials, they were presented with their designed materials and target materials side by side. We asked subjects to rate how satisfied with their own results from 1 to 5, with 5 representing very satisfied.

7.3. Evaluation

We evaluate the user study in two parts. First, we analyze the output of the experiment as subjects proceed through each trial, including the time to completion and matching errors. Second, we analyze user behavior during the experiment and the self-rating at the end of study. We use repeated measures analysis of variance (ANOVA) [Ste12] to compute statistical significance. This method is appropriate to calculate correlations with within-subject factors and violates the assumption of independence in standard one-way ANOVA. A p value represents the confidence of difference between two sample groups. A p value below 0.05 indicate 95% confidence of difference. In all figures, error bars represent standard error.

Time to Completion. We recorded the time that each subject used for each trial. Figure 9 shows the average time to completion for each trial over all users. Generally, time to completion on the sliders only interface is significantly higher than the image navigation with sliders interface ($p = 0.035, 0.046, 0.090$ for three tasks respectively). The time spent on the physical interface (2) is the lowest among all the interfaces. Combining with the results of matching errors introduced next, we believe this is because the influence of some physical parameters on the appearances are subtle and not linearly scaled. Subjects tend to stop editing the material when they believe changing the parameters would not make a difference in the appearance. The typical examples are the physical parameter roughness and anisotropy. According to our observations, the influence of roughness becomes less significant when its value is above a threshold. Interpolating in log space alleviates this problem to an extent, but its contribution is still less obvious compared with diffuse and specular albedo. The same applies to the anisotropy parameter.

Also we notice that the average time spent on the task 3 is lower than the other two tasks, in spite of the dramatic difference between starting and target materials. We believe this is because subjects are more sensitive to the highly specular materials and able to observe the subtle differences in appearance. This can also be demonstrated by the low perception errors and high self-rating results that will be discussed in the following sections.




Figure 9: Average time to completion for all tasks over all subjects (in seconds). Time spent on the physical image navigation and task 3 are generally lower than other scenarios.

Table 1: p values for different tasks and different error metrics

	Task 1	Task 2	Task 3
Parameter Error	0.024	0.039	0.042
MSE	0.061	0.042	0.032
Perception Error	0.048	0.058	0.095

Matching Errors. To summarize the overall performance of each interface, we compare the subject’s result with the target material using three different metrics: parameter error, image MSE error and perception error. For parameter error, we measure the MSE between the user specified parameters and the ground truth parameters of the target material. Parameters for different models are first normalized and then put into the same scale for comparison and visualization.

For image MSE error, we measure the MSE error between the rendered images. We compute the perception error using cubic root metric proposed in [NDM06].

Figure 10 shows the different error metrics for each task averaged over all subjects, and p values can be found in Table 1. It is clear that both the perceptual interfaces (1) and (3) outperform the physical interface (2) under different error metrics, especially for task 1 and 2. Notice that the standard deviation of the physical interface (2) is larger than the other interfaces, which suggests that the sensitivity of how parameters of analytical model could changes the appearances is different among the subjects. Some subjects might not notice the subtle changes of the materials when they were performing the tasks on the physical interfaces, which leads to the large errors. Comparing the interface (1) and (3), we can see that the slider only interface (3) has slightly larger errors, indicating the subjects have better performance with the help of the image navigation.

Figure 11 shows the perception error over time for two different subjects performing the three tasks with all interfaces. The general trends of the error is decreasing because the materials get closer to the target as the subject proceeded through the trial. We include error graphs for all subjects with different error metrics in supplemental material. We observe that perceptual interfaces (1) tend to converge relatively more quickly with less volatility than the the other two interfaces.

User Behaviors. To better understand user behaviors during the experiment, we compute the percentage of the total operations conducted on both the image navigation and the sliders over all subjects




Figure 10: Average errors computed over all subjects using different metrics, including parameter MSE error, image MSE error and perception error. Generally, subjects have best performance with lowest error on task 3.




Figure 11: Example graphs of perception error over time (in seconds) for two subjects. The errors on perceptual image navigation converge faster and are less volatile in general.

on interface (1) and (2). Figure 12 shows the results for all 6 trials ($p = 0.021, 0.059, 0.042, 0.047, 0.038, 0.039$ respectively). In general, we can see that for all the trials, the percentage of total operations on the image navigation and the sliders is close to half and half. This observation suggests that when editing the materials, users tend to take advantage of both the tools to achieve their goals, instead of relying on one and ignoring the other. Combining with the results of time to completion and parameter errors, we can also draw the conclusion that with the help of image navigation, users can match the appearances much faster with higher accuracy.

Self-Rating. Figure 13 demonstrates how subjects were satisfied with their matched materials compared with the targets. Overall, subjects were satisfied with their results on all the interfaces, and there is no significant difference between the ratings on the perceptual interface (1) and the physical interface (2). The average ratings on perceptual with slider only interface (3) are slightly lower than the other two for task 1 and 2. The average rating for task 3 is slightly higher than the other two tasks because of the human sensitivity for glossy materials, which makes the matching task easier. The average rating for task 2 is the lowest potentially because subjects are unfamiliar with the subtle glossy reflection with the darker appearance due to its rareness in real life. It can also be explained by the fact




Figure 12: The percentage of the total operations conducted on the image grid and sliders over all subjects for the 6 trials on image navigation interfaces (both perceptual and physical).

that task 2 suffers from the highest parameter errors and perception errors for all interfaces in Figure 10.

Discussion. Based on the results, we can conclude that the proposed perceptual space has advantages for matching material tasks under the circumstances of our user study. Compared with the physical space with analytical model, all the parameters of the proposed perceptual space have strong contributions to the appearance and




Figure 13: Average self-rated satisfaction scores over all subjects. The ratings on task 2 is the lowest while the ratings on task 3 is the highest.

the material samples are fairly spread within the space. Another conclusion is that the image navigation combined with sliders could help users for material editing in terms of efficiency and accuracy. We notice these conclusions are in conflict with Kerr et al. [KP10], who believe that there is no significant difference between the physical interface and perceptual interface, and both of them are better than the image navigation [NDM06]. However, we want to point out the two major differences between their work and ours: (1) how the perceptual space is computed, and (2) our proposed interface is different from image navigation.

In Kerr et al, the parameters of the perceptual space are directly converted from Ward and Cook-Torrance BRDF models. Theoretically, their perceptual space is a different parameterization of the physical space and should cover the exact same space but with different dimensions. Our perceptual space is computed from the embedding using the crowdsourcing data, and the physical space is generated by fitting the BRDF into the analytical model. Therefore the two spaces overlap but are not the same. We believe this might be the reason why we have different conclusions.

Our interface is inspired from the image navigation but is essentially different. Even though our interface provides the image navigation, it is a 1D visualization of each dimension of the space instead of 2D in the image navigation. According to our pilot experiment, we found that the 2D image navigation could easily confuse the users especially when they need to select two out of multiple dimensions to build a visualization and switch between other visualizations. Besides, our interface combines the image navigation with the sliders. Users can manipulate the sliders to create a rough cut and use the image navigation for fine-tuning, or vice versa, which suggests that our interface provides users recommendations and more information to assist their editing. We believe that more information is the reason that our interface outperforms the slider only interface.

8. Limitations and Conclusion

In this paper, we have proposed and analyzed a low-dimensional perceptual space for measured metal materials. Our study covers diverse material appearances collected from different datasets. We have conducted a large-scale psychophysical user study and gathered over 30K valid answers from 360 participants. We evaluated the responses using an NMDS model to extract a perceptual embedding for each material sample. We map the perceptual embedding to the reflectance data using GPR to interpolate the BRDFs, which

provides a continuous perpetual space for appearance editing. We have further designed and implemented a novel intuitive material editing interface that takes advantage of the proposed perceptual model. For evaluation, we conduct a separate user study to compare our interface with the traditional physical-based editing interface in terms of the accuracy and efficiency of appearance matching tasks.

There are limitations in our study and future work is needed. The first limitation is the material category. As we have discussed in previous sections, we only focus on understanding human perception on metal materials in this work due to their special visual appearance. In fact, a key advantage of our flexible methodology is that it can be applied to different types of materials (such as plastic and cloth, etc.) and we can build individual perceptual spaces with unique attributes for each material category, which we believe will provide more insights on material recognition and perception. However, as the first step towards studying the perception within a certain material category, more experiments and validation would definitely be encouraged in the future. Another limitation is the small size of the measured BRDF datasets. Even though we have covered a large range of metal materials, there are still many different materials that display unique appearance in real life. With the emergence of new material acquisition and capture techniques, we hope to introduce more metal BRDF and SVBRDF into our study in the future. And last, we hope our proposed editing interface could inspire additional research on intuitive material design and interaction. The paradigm of material editing is similar to color picking but with more complexity and high dimensionality. However, few efforts have been made to improve the intuitiveness and accuracy of material editing interfaces.

Acknowledgement

This work was supported in part by NSF grant IIS-2007283, and French National Research Agency project CaLiTrOp ANR-16-CE33-0026.

References

- [ADO09] ADOBE SYSTEMS INC. *Photoshop CS 4*. 2009 3.
- [And11] ANDERSON, BARTON L. “Visual perception of materials and surfaces”. *Current biology* 21.24 (2011), R978–R983 2.
- [Aut17] AUTODESK INC. *MAYA 2017*. 2017 3.
- [AWC*07] AGARWAL, SAMEER, WILLS, JOSH, CAYTON, LAWRENCE, et al. “Generalized non-metric multidimensional scaling”. *Artificial Intelligence and Statistics*. 2007, 11–18 4, 5.
- [BBPA15] BOYADZHIEV, IVAYLO, BALA, KAVITA, PARIS, SYLVAIN, and ADELSON, EDWARD. “Band-sifting decomposition for image-based material editing”. *ACM Transactions on Graphics (TOG)* 34.5 (2015), 1–16 3.
- [Ble02] BLENDER FOUNDATION. *BLENDER*. 2002 1, 3.
- [BS12] BURLEY, BRENT and STUDIOS, WALT DISNEY ANIMATION. “Physically-based shading at Disney”. *ACM SIGGRAPH*. Vol. 2012. vol. 2012. 2012, 1–7 6.
- [BUSB15] BELL, SEAN, UPCHURCH, PAUL, SNAVELY, NOAH, and BALA, KAVITA. “Material recognition in the wild with the materials in context database”. *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015, 3479–3487 1.
- [CK15] CHADWICK, ALICE C and KENTRIDGE, RW. “The perception of gloss: A review”. *Vision research* 109 (2015), 221–235 2.

[CK17] CONTY, ALEJANDRO and KULLA, CHRISTOPHER. “Revisiting physically based shading at Imageworks”. *ACM SIGGRAPH*. 2017, 7 8.

[CPK06] COLBERT, MARK, PATTANAIK, SUMANTA, and KRIVANEK, JAROSLAV. “BRDF-Shop: Creating physically correct bidirectional reflectance distribution functions”. *IEEE Computer Graphics and Applications* 26.1 (2006), 30–36 3.

[CWAP08] CHESLACK-POSTAVA, EWEN, WANG, RUI, AKERLUND, OS-KAR, and PELLACINI, FABIO. “Fast, realistic lighting and material design using nonlinear cut approximation”. *ACM Transactions on Graphics (TOG)* 27.5 (2008), 1–10 3.

[Deb08] DEBEVEC, PAUL. “Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography”. *ACM SIGGRAPH 2008 classes*. 2008, 1–10 8.

[DFY*11] DOERSCHNER, KATJA, FLEMING, ROLAND W, YILMAZ, OZGUR, et al. “Visual motion and the perception of surface material”. *Current biology* 21.23 (2011), 2010–2016 2.

[DJ18] DUPUY, JONATHAN and JAKOB, WENZEL. “An adaptive parameterization for efficient material acquisition and rendering”. *SIGGRAPH Asia 2018 Technical Papers*. ACM. 2018, 274 3.

[DRS10] DORSEY, JULIE, RUSHMEIER, HOLLY, and SILLION, FRANÇOIS. *Digital modeling of material appearance*. Elsevier, 2010 2.

[Fau19] FAUL, FRANZ. “The influence of Fresnel effects on gloss perception”. *Journal of vision* 19.13 (2019), 1–1 2.

[FB05] FLEMING, ROLAND W and BÜLTHOFF, HEINRICH H. “Low-level image cues in the perception of translucent materials”. *ACM Transactions on Applied Perception (TAP)* 2.3 (2005), 346–382 1.

[FDA03] FLEMING, ROLAND W, DROR, RON O, and ADELSON, EDWARD H. “Real-world illumination and the perception of surface reflectance properties”. *Journal of vision* 3.5 (2003), 3–3 3.

[FFG12] FORES, ADRIA, FERWERDA, JAMES, and GU, JINWEI. “Toward a perceptually based metric for BRDF modeling”. *Color and imaging conference*. Vol. 2012. 1. Society for Imaging Science and Technology. 2012, 142–148 2.

[Fle14] FLEMING, ROLAND W. “Visual perception of materials and their properties”. *Vision research* 94 (2014), 62–75 2, 3.

[FNG15] FLEMING, ROLAND W, NISHIDA, SHIN’YA, and GEGENFURTNER, KARL R. “Perception of material properties.” (2015) 1, 2.

[GTHP18] GIGILASHVILI, DAVIT, THOMAS, JEAN-BAPTISTE, HARDEBERG, JON YNGVE, and PEDERSEN, MARIUS. “Behavioral investigation of visual appearance assessment”. *Color and Imaging Conference*. 1. Society for Imaging Science and Technology. 2018, 294–299 2.

[GUT*19] GIGILASHVILI, DAVIT, URBAN, PHILIPP, THOMAS, JEAN-BAPTISTE, et al. “Impact of Shape on Apparent Translucency Differences”. *Color and Imaging Conference*. Society for Imaging Science and Technology. 2019, 132–137 2.

[GVPV17] GEORGULIS, STAMATIOS, VANWEDDINGEN, VINCENT, PROESMANS, MARC, and VAN GOOL, LUC. “Material classification under natural illumination using reflectance maps”. *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE. 2017, 244–253 1.

[GWA*15] GKIOLAKAS, IOANNIS, WALTER, BRUCE, ADELSON, EDWARD H, et al. “On the appearance of translucent edges”. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2015, 5528–5536 2.

[GXZ*13] GKIOLAKAS, IOANNIS, XIAO, BEI, ZHAO, SHUANG, et al. “Understanding the role of phase function in translucent appearance”. *ACM Transactions on graphics (TOG)* 32.5 (2013), 1–19 1, 2.

[HFM16] HAVRAN, VLASTIMIL, FILIP, JIRI, and MYSZKOWSKI, KAROL. “Perceptually motivated BRDF comparison using single image”. *Computer graphics forum*. Vol. 35. 4. Wiley Online Library. 2016, 1–12 2–4.

[HGC*20] HU, BINGYANG, GUO, JIE, CHEN, YANJUN, et al. “DeepBRDF: A Deep Representation for Manipulating Measured BRDF”. *Computer Graphics Forum*. Vol. 39. 2. Wiley Online Library. 2020, 157–166 3, 7.

[HJ12] HORN, ROGER A. and JOHNSON, CHARLES R. *Matrix Analysis*. en. Google-Books-ID: O7sgAwAAQBAJ. Cambridge University Press, Oct. 2012. ISBN: 978-1-139-78888-5 5.

[HLM06] HO, YUN-XIAN, LANDY, MICHAEL S, and MALONEY, LAURENCE T. “How direction of illumination affects visually perceived surface roughness”. *Journal of vision* 6.5 (2006), 8–8 2.

[Jak10] JAKOB, WENZEL. *Mitsuba renderer*. <http://www.mitsuba-renderer.org>. 2010 3.

[KFB10] KRIVÁNEK, JAROSLAV, FERWERDA, JAMES A, and BALA, KAVITA. “Effects of global illumination approximations on material appearance”. *ACM Transactions on Graphics (TOG)* 29.4 (2010), 1–10 2.

[Kol09] KOLTUN, VLADLEN. “Exploratory modeling with collaborative design spaces”. (2009) 3.

[KP10] KERR, WILLIAM B and PELLACINI, FABIO. “Toward evaluating material design interface paradigms for novice users”. *ACM Transactions on Graphics (TOG)*. Vol. 29. 4. ACM. 2010, 35 2, 3, 7, 11.

[LM01] LEUNG, THOMAS and MALIK, JITENDRA. “Representing and recognizing the visual appearance of materials using three-dimensional textons”. *International journal of computer vision* 43.1 (2001), 29–44 2.

[LMS*19] LAGUNAS, MANUEL, MALPICA, SANDRA, SERRANO, ANA, et al. “A similarity measure for material appearance”. *arXiv preprint arXiv:1905.01562* (2019) 2, 4.

[LSGM21] LAGUNAS, MANUEL, SERRANO, ANA, GUTIERREZ, DIEGO, and MASIA, BELEN. “The joint role of geometry and illumination on material recognition”. *arXiv preprint arXiv:2101.02496* (2021) 2.

[MAB*97] MARKS, JOE, ANDALMAN, BRAD, BEARDSLEY, PAUL A, et al. “Design galleries: A general approach to setting parameters for computer graphics and animation”. *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, 389–400 3.

[MB10] MALONEY, LAURENCE T and BRAINARD, DAVID H. “Color and material perception: Achievements and challenges”. *Journal of vision* 10.9 (2010), 19–19 2.

[MGZ*17] MYLO, MARLON, GIESEL, MARTIN, ZAIDI, QASIM, et al. “Appearance bending: A perceptual editing paradigm for data-driven material models”. *Proceedings of the conference on Vision, Modeling and Visualization*. 2017, 9–16 3.

[MJ20] MARSCHALLEK, BARBARA E and JACOBSEN, THOMAS. “Classification of material substances: Introducing a standards-based approach”. *Materials & Design* 193 (2020), 108784 1.

[MPBM03] MATUSIK, WOJCIECH, PFISTER, HANSPIETER, BRAND, MATT, and MCMILLAN, LEONARD. “A Data-Driven Reflectance Model”. *ACM Transactions on Graphics* (July 2003) 1–3.

[NDM05] NGAN, ADDY, DURAND, FRÉDO, and MATUSIK, WOJCIECH. “Experimental analysis of brdf models.” *Rendering Techniques* 2005.16th (2005), 2 2.

[NDM06] NGAN, ADDY, DURAND, FRÉDO, and MATUSIK, WOJCIECH. “Image-driven Navigation of Analytical BRDF Models.” *Rendering Techniques* 2006 (2006), 399–407 2, 3, 7, 9, 11.

[NJR15] NIELSEN, JANNIK BOLL, JENSEN, HENRIK WANN, and RAMMOORTHI, RAVI. “On optimal, minimal BRDF sampling for reflectance acquisition”. *ACM Transactions on Graphics (TOG)* 34.6 (2015), 186 2, 6.

[NKLN10] NGUYEN, CHUONG H, KYUNG, MIN-HO, LEE, JOO-HAENG, and NAM, SEUNG-WOO. “A pca decomposition for real-time brdf editing and relighting with global illumination”. *Computer Graphics Forum*. Vol. 29. 4. Wiley Online Library. 2010, 1469–1478 3.

[PFG00] PELLACINI, FABIO, FERWERDA, JAMES A, and GREENBERG, DONALD P. “Toward psychophysically-based light reflection model for image synthesis”. *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co. 2000, 55–64 1–3.

[PR12] PEREIRA, THIAGO and RUSINKIEWICZ, SZYMON. “Gamut mapping spatially varying reflectance with an improved BRDF similarity metric”. *Computer graphics forum*. Vol. 31. 4. Wiley Online Library. 2012, 1557–1566 2.

[Ras03] RASMUSSEN, CARL EDWARD. “Gaussian processes in machine learning”. *Summer School on Machine Learning*. Springer. 2003, 63–71 6.

[SCW*21] SERRANO, ANA, CHEN, BIN, WANG, CHAO, et al. “The effect of shape and illumination on material perception: model and applications”. *ACM Trans. on Graph.* (2021) 2.

[SDR20] SHI, WEIQI, DORSEY, JULIE, and RUSHMEIER, HOLLY. “Learning-Based Inverse Bi-Scale Material Fitting from Tabular BRDFs”. *IEEE Transactions on Visualization and Computer Graphics* (2020) 3.

[SGM*16] SERRANO, ANA, GUTIERREZ, DIEGO, MYSZKOWSKI, KAROL, et al. “An intuitive control space for material appearance”. *ACM Transactions on Graphics (SIGGRAPH ASIA 2016)* 35.6 (2016) 1–4, 7.

[SJR18] SUN, TIANCHENG, JENSEN, HENRIK WANN, and RAMAMOORTHI, RAVI. “Connecting Measured BRDFs to Analytic BRDFs by Data-Driven Diffuse-Specular Separation”. *ACM Transactions on Graphics (TOG)* 37.6 (2018), 273 8.

[SS95] SCHRÖDER, PETER and SWELDENS, WIM. “Spherical wavelets: Efficiently representing functions on the sphere”. *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. 1995, 161–172 2.

[SSGM17] SUN, TIANCHENG, SERRANO, ANA, GUTIERREZ, DIEGO, and MASIA, BELEN. “Attribute-preserving gamut mapping of measured BRDFs”. *Computer graphics forum*. Vol. 36. 4. Wiley Online Library. 2017, 47–54 2.

[SSKK15] SAPONARO, PHILIP, SORENSEN, SCOTT, KOLAGUNDA, ABHISHEK, and KAMBHAMETTU, CHANDRA. “Material classification with thermal imagery”. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2015, 4649–4656 1.

[SSN18] SOLER, CYRIL, SUBR, KARTIC, and NOWROUZEZHARAI, DEREK. “A versatile parameterization for measured material manifolds”. *Computer graphics forum*. Vol. 37. 2. Wiley Online Library. 2018, 135–144 1–3, 6.

[Ste12] STEVENS, JAMES P. *Applied multivariate statistics for the social sciences*. Routledge, 2012 9.

[Stu99] STURM, JOS F. “Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones”. *Optimization methods and software* 11.1-4 (1999), 625–653 5.

[SWS*17] SHI, WEIQI, WANG, ZEYU, SEZGIN, METIN, et al. “Material design in augmented reality with in-situ visual feedback”. *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. 2017, 93–103 3.

[SZC*07] SUN, XIN, ZHOU, KUN, CHEN, YANYUN, et al. “Interactive relighting with dynamic BRDFs”. *ACM SIGGRAPH 2007 papers*. 2007, 27–es 3.

[TGG*20] TOSCANI, MATTEO, GUARNERA, DAR’YA, GUARNERA, GIUSEPPE CLAUDIO, et al. “Three perceptual dimensions for specular and diffuse reflection”. *ACM Transactions on Applied Perception (TAP)* 17.2 (2020), 1–26 2.

[TLB*11] TAMUZ, OMER, LIU, CE, BELONGIE, SERGE, et al. “Adaptively learning the crowd kernel”. *arXiv preprint arXiv:1105.1033* (2011) 4.

[VBF17] VANGORP, PETER, BARLA, PASCAL, and FLEMING, ROLAND W. “The perception of hazy gloss”. *Journal of vision* 17.5 (2017), 19–19 2.

[VBF18] VAN ASSEN, JAN JAAP R, BARLA, PASCAL, and FLEMING, ROLAND W. “Visual features in the perception of liquids”. *Current biology* 28.3 (2018), 452–458 2.

[VLD07] VANGORP, PETER, LAURIJSSEN, JURGEN, and DUTRÉ, PHILIP. “The influence of shape on the perception of material reflectance”. *ACM SIGGRAPH 2007 papers*. 2007, 77–es 2.

[VZ08] VARMA, MANIK and ZISSERMAN, ANDREW. “A statistical approach to material classification using image patch exemplars”. *IEEE transactions on pattern analysis and machine intelligence* 31.11 (2008), 2032–2047 1.

[WAKB09] WILLS, JOSH, AGARWAL, SAMEER, KRIEGMAN, DAVID, and BELONGIE, SERGE. “Toward a perceptual space for gloss”. *ACM Transactions on Graphics (TOG)* 28.4 (2009), 1–15 1–5.

[WAT92] WESTIN, STEPHEN H, ARVO, JAMES R, and TORRANCE, KENNETH E. “Predicting reflectance functions from complex surfaces”. *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*. 1992, 255–264 2.

[WM01] WESTLUND, HAROLD B and MEYER, GARY W. “Applying appearance standards to light reflection models”. *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. 2001, 501–51 3.

[WMLT07] WALTER, BRUCE, MARSCHNER, STEPHEN R, LI, HONGSONG, and TORRANCE, KENNETH E. “Microfacet Models for Refraction through Rough Surfaces.” *Rendering techniques* 2007 (2007), 18th 8.