Learning-Based Inverse Bi-Scale Material Fitting from Tabular BRDFs

Weiqi Shi, Julie Dorsey, Holly Rushmeier

Abstract—Relating small-scale structures to large-scale appearance is a key element in material appearance design. Bi-scale material design requires finding small-scale structures – meso-scale geometry and micro-scale BRDFs – that produce a desired large-scale appearance expressed as a macro-scale BRDF. The adjustment of small-scale geometry and reflectances to achieve a desired appearance can become a tedious trial-and-error process. We present a learning-based solution to fit a target macro-scale BRDF with a combination of a meso-scale geometry and micro-scale BRDF. We confront challenges in representation at both scales. At the large scale we need macro-scale BRDFs that are both compact and expressive. At the small scale we need diverse combinations of geometric patterns and potentially spatially varying micro-BRDFs. For large-scale macro-BRDFs, we propose a novel 2D subset of a tabular BRDF representation that well preserves important appearance features for learning. For small-scale details, we represent geometries and BRDFs in different categories with different physical parameters to define multiple independent continuous search spaces. To build the mapping between large-scale macro-BRDFs and small-scale details, we propose an end-to-end model that takes the subset BRDF as input and performs classification and parameter estimation on small-scale details to find an accurate reconstruction. Compared with other fitting methods, our learning-based solution provides higher reconstruction accuracy and covers a wider gamut of appearance.

Index Terms—reflectance, material appearance, bi-scale materials

1 Introduction

Bringing virtually designed materials into the real world is a challenging problem. Physical properties including structures and reflectance behaviors determine the appearance of a material. For manufacturing industries, small-scale geometric patterns are introduced for molding textures [1] in order to enhance material performance and increase the gamut of appearance. However, achieving a particular large-scale target appearance is challenging given the diverse choices for meso-scale geometry and micro-scale material reflectance. Without an automatic solution, the traditional design process relies on tweaking parameters to match a specific appearance, which can be a tedious (and computationally expensive) process of trial and error.

Our goal is to solve this problem by providing an automatic solution to fit a target large-scale macro-BRDF with a meso-scale geometry and a micro-scale BRDF. We use the term "large-scale appearance" to refer to the material appearance when viewed from a relatively far away distance (typically measured by meters), which can be modeled by a macro-BRDF. When we zoom in, we can take a close up view of the "small-scale details," which are made up of the meso-scale geometric structures (typically measured by millimeters). We use the micro-scale BRDF to define the basic reflectance of each facet at this level. The difference in the macro-BRDF and the micro-BRDF is caused by the shadowing and masking effects created by the meso-scale geometries. Figure 2 shows an example of a physical material at two different scales. We will refer to the combination of the meso-scale geometry and the micro-scale BRDF as the "small-scale details." We do not seek a unique solution for the small-scale details but the reconstructed large-scale appearance should be identical (or very close) to the target. At first glance, this is similar to previous BRDF fitting problems [3], [4], [7], [18], [34], [38], [46], where a single analytical model is fit to BRDF data and only a few parameters need to be optimized. However, fitting bi-scale materials is different because the search space of small-scale details is not completely continuous. The meso-scale geometries are defined by different categories, such as woven or bricks structures, and each category has unique properties and representations. There is no "universal formula" that can model the appearances of different categories, which makes the fitting problem highly nonlinear. Therefore, it is almost impossible to apply previous optimization-based solutions in a straightforward manner.

In this paper, we propose a learning-based solution using a convolutional neural network with synthetic data to learn the mapping between two scales. There are major challenges at both scales – the representation of materials at the large scale and the diverse combinations of geometric patterns and basic BRDFs for each facet at the small scale.

In computer vision and graphics the most common representation for learning materials is an image of an instance of the material in a particular context. Images are used due to the simplicity of data collection. However, material appearance in an image context depends on the scene description (including lighting, camera angle, and geometric normals) that are irrelevant to the the fitting and could produce unwanted bias in the result. For reliability and efficiency in our method, we propose a novel training pipeline to learn material properties directly from tabulated BRDFs. However, due to its high dimension, it is inefficient to conduct training on densely sampled tabulated BRDFs. We propose a new sampling method to create 2D BRDF slices for training. Compared with existing BRDF tabulations, our method well preserves the spatial features of BRDF by taking advantage of retro-reflection, which according to [14] provides an important

• *W. Shi, J. Dorsey, H. Rushmeier are with the Department of Computer Science, Yale University, New Haven, CT, 06511.*

1. http://www.stpaulengraving.com/texturing-plastic-injection-mold.html

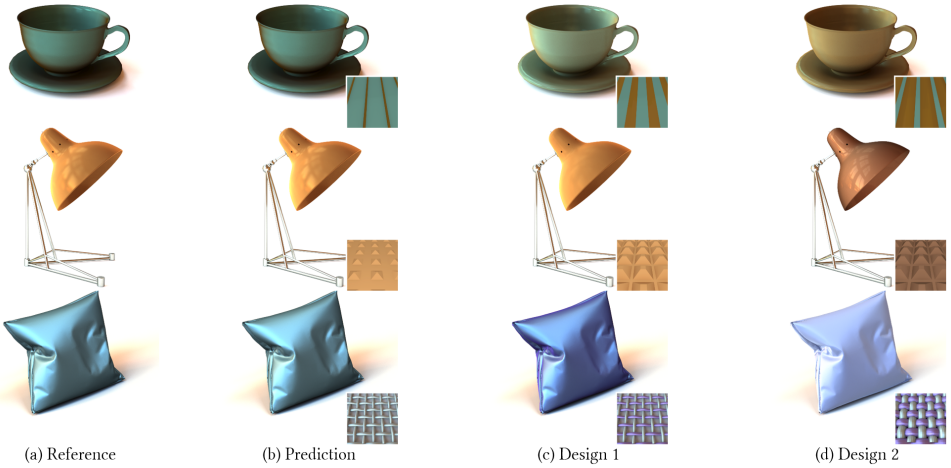

(a) Reference (b) Prediction (c) Design 1 (d) Design 2

Fig. 1: We present a new bi-scale material fitting framework. Given the large-scale appearance of materials (a), our framework predicts the meso-scale geometry and micro-scale BRDF and uses them to accurately reconstruct the materials (b). The output can be used as a starting point for material design (c) and (d). The predicted small-scale details provide a new material editing scheme to change the large-scale appearance by varying meso-scale geometry (top), switching micro-scale BRDF (middle) and editing anisotropy (bottom).

Fig. 2: Example of bi-scale material in real life. The large-scale appearance of the top of a step stool (left) that appears to have a pattern with two materials, but actually it is the same micro-BRDF with two different meso-scale geometries in the wavy pattern (right).

cue to the microfacet normal distribution function. We analyze the physical interpretations and compare the performance of other BRDF tabular formats to provide explanations for our findings.

For small-scale details, we represent geometries and BRDFs in different categories with different physical parameters to define multiple independent continuous search spaces. According to [16], [17], learning materials can be interpreted as two processes: classification and estimation. Classification identifies the boundaries separating different material categories and estimation establishes the true position within each category. We follow this framework and put forward an end-to-end model that takes as input a 2D BRDF slice and performs classification and parameter estimation on the small-scale details. The classification is based on predefined categories and the estimation is to fine-tune the corresponding parameters in a continuous search space. This framework also helps us to increase the generality of our model because the gamut of the search space is enlarged. Theoretically, there is no limit on the categories of geometric patterns and basic BRDFs that can be used in training as long as the computational power is available. This approach is not limited by the expressive power of a single BRDF model and we can fit materials with diverse properties.

For evaluation, we conduct experiments to compare the training performance of different tabular formats and explain why retro-reflection provides a strong cue for the learning process. We reconstruct large-scale appearance based on our fitting results and compare with ground truth. The rendering results evaluated with a user study demonstrate that there is no significant visual difference. Experimental results also show that our method can predict accurate small-scale details from real world materials.

In summary, this paper makes the following contributions:

- A new learning-based solution for bi-scale material fitting that takes large-scale BRDF as input and predicts small-scale details as output.
- A new method that efficiently learns materials using sparsely sampled BRDF by taking advantage of spatial features and retro-reflection.
- An end-to-end model to handle diverse representations and combinations of meso-scale geometries and micro-scale BRDFs by classification and parameter estimation.
- A prototype using our proposed method for material editing and fabrication.

2 Background

We present background on the bi-scale design problem, microfacet models that relate geometry and appearance, previous work using machine learning in material modeling, and work in BRDF fitting.

2.1 Bi-Scale Design

Bi-scale design plays an important role in manufacturing appearance. In manufacturing, creating meso-scale geometry is called texturing (in contrast with the graphics usage of the term texturing.) Texturing molds to create meso-scale geometry such as leather patterns or stipples can make a relatively inexpensive material (such as plastic) look more luxurious, or can make a brand look more distinctive. Meso-scale geometry can also have functional properties such as hiding imperfections and making products easier to grip. Maniscalchi et al. [44] provide an in-depth study of the performance of various mold texture patterns. Several basic template patterns are given by Wisconsin Engraving [50] that can be used as the small-scale details for mold texturing, and St. Paul Engraving [41] gives a step by step introduction to mold texturing. These industry examples demonstrate how bi-scale materials can enrich the gamut of manufactured material appearance.

Several research projects consider bi-scale design for manufacture. Rouiller et al. [43] propose a method to 3D print spatially varying BRDFs by optimizing the normal distribution function (NDF) of discrete facets to match an analytical NDF. Weyrich et al.

[49] optimize for a maximally-continuous and valley-minimizing height field to generate a desired reflectance. Zhao et al. [57] create volumetric models of woven cloth by matching an input weave pattern to a sample from a small database of volumetric exemplars. Lan et al. [30] present a method to 3D-print a surface with spatially-varying opaque reflectance and shading frames. These projects all consider specialized, rather than general bi-scale design. Wu et al. [53] propose a solution to prefilter high-resolution displacement maps and BRDFs jointly while preserving material appearance. However, their pre-computation for each micro-structure is time-consuming. Wu et al. [51] present a physically-based interactive system to design large-scale materials by editing small-scale geometry and BRDFs in low-rank matrix formulations. Their system can efficiently simulate material appearance from small-scale details with shadowing and masking effects. Their system provides the basis for our simulation. However, their algorithm suffers from a long pre-computation time and large memory allocation. We improve their algorithm to efficiently simulate a large number of materials for training.

2.2 Microfacet Material Modeling

Many analytical models describe large-scale material appearance with the geometry of a microfacet-based BRDF [2], [10], [40]. Walter et al. [48] give a review of microfacet theory and extend it to transmission through rough surfaces. Heitz [21] reviews different masking-shadowing functions in microfacet-based BRDFs. These papers illustrate basic microfacet theories and analytical models but are not concerned with meso-scale structures and bi-scale mapping.

The methods in [22], [23] introduce scattering and diffraction effects in microfacet material simulation. There also has been much progress for different types of microfacet materials, such as Aliaga et al. [1], Dong et al. [12], Nam et al. [37]. Belcour et al. [5] introduce an extension to microfacet theory for rendering iridescent effects caused by thin-films on a rough base layer. Dupuy et al. [13] propose a method to relate an analytical microfacet BRDF model to an input material by solving an eigenvector problem built from backscattering samples. Zhao et al. [56] produce large-scale fabric material appearance from small-scale structures measured with micro CT imaging. Heidrich et al. [20] use precomputed small-scale visibility to calculate large-scale BRDFs with indirect illumination. Kuznetsov et al. [29] propose a method for microfacet material rendering using generative models which can learn the generalized normal distribution functions (GNDF) from latent vectors. However, the geometric optics GNDF they used does not provide an explicit definition for geometry patterns for editing and the fabric GNDF only covers a small gamut of appearance. None of these methods are suitable for systematically simulating a large number of varying meso-scale structures for bi-scale training.

2.3 Inverse Rendering and BRDF Fitting

Traditional inverse rendering and BRDF fitting algorithms can also be used to specify materials from a particular appearance. Li et al. [31] derive a probabilistic formulation to jointly estimate the shape and BRDF of objects under known lighting from a synthetic dataset. Romeiro and Zickler [42] try to infer BRDF from a single image of a known shape in an unknown lighting environment, by assuming a statistical distribution of natural illumination. Many BRDF fitting algorithms [3], [4], [7], [18], [34], [38], [46] have focused on using analytical BRDF models to fit the MERL dataset [35]. Dupuy et al. [13] proposed an iterative method to extract the microfacet parameters from anisotropic materials. All of these methods try to map a material appearance to a specific BRDF model, which can only cover a limited gamut of appearance. None of these algorithms provide a bi-scale description. The study closest to our work is Wu et al. [52]. They propose a search-based inverse rendering method; they search through libraries of materials and meso-scale geometries to find the best combination for a specified appearance. However, the searching results are the "nearest neighbor" defined in the training set, which limits reconstruction accuracy. We provide a detailed comparison between our work and theirs in Section 5.2.

2.4 Learning Materials in Image Space

Many projects relate appearance to material descriptions using machine learning techniques. Most of them rely on image-based training, where the same material sample is repeatedly rendered with different scenes to cover different appearances [19], [24], [27], [32], [33], [36], [47], [55]. They generate synthetic images as training data, where the same material is repeatedly rendered using a single BRDF model with different lighting and camera angles on a fixed geometry (sphere or plane). Normally, the gamut of their model is limited due to the single parametric BRDF model being used (Ward or Phong in many cases). Furthermore, iterating over different lighting conditions and other irrelevant factors is a waste in terms of sampling, and significantly increases the computational time and storage for the training set. Another disadvantage is the possible overfitting of the rendered scenes since the objects and lighting used for rendering could bias the training results. Zsolnai-Feh'er et al. [58] present an image-based learning system to learn shader parameters from rendered images with fixed lighting and viewing angle. However, their setting is tied to the Disney Principal Shader [9] and is customized for the input of their design system, which cannot be applied to other scenarios. To avoid sampling redundancies and reduce the dimension of the training set, we propose a tabular BRDF training scheme to learn appearance efficiently from BRDFs instead of from rendered scenes.

3 Method

Our goal is to find the combination of meso-scale geometries and micro-scale BRDFs that can accurately reconstruct the input material appearance. Our method takes as input macro-scale BRDFs in 2D slice tabular format and outputs the categories and related parameters for meso-scale geometry and micro-scale BRDF. We conduct our training on synthetic data created from physically-based simulation (Section 3.1). We define the tabular BRDF as a 2D slice subset from the complete 4D BRDF (Section 3.2). We use an end-to-end model to find the small-scale representations that can accurately reconstruct the large-scale appearance (Section 3.3). The training data is generated in gray-scale by averaging the RGB of micro-scale BRDF for dimension reduction. Otherwise, the dimension of the training set will increase polynomially. To restore colors and enrich the gamut of our model, we propose a method to predict the multi-color weights for different parts of the meso-scale geometry after training (Section 3.4). In this paper, we assume the macro-scale BRDF has been specified by the designer.

3.1 Simulation

We use synthetic data for training since it can be collected from computation easily (labeled bi-scale material data relating macro BRDF to small-scale details is difficult to obtain for physical

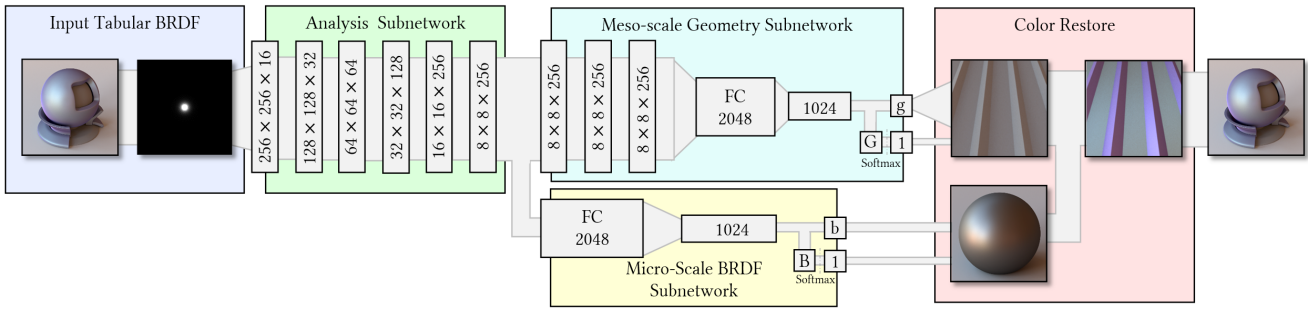

Fig. 3: The pipeline of our bi-scale learning model and network architecture. Our model takes the 2D slice tabular BRDF as input and outputs the meso-scale geometry and micro-scale BRDF. $\boldsymbol{G}$ and $\boldsymbol{B}$ represent the number of categories for classification and $\boldsymbol{g}$ and $\boldsymbol{b}$ represent the dimension of parameters for regression. Given prediction results from the network, our pipeline runs the color restore algorithm to compute color weights for different sides of the meso-scale geometry to reconstruct the large-scale appearance.

samples). However, a concern is the speed of the simulation. State-of-art microfacet modeling methods focus on accuracy. The underlying structure is resolved in pixels in the rendering to produce very high-quality material appearance. However these methods are not suitable in our case due to the lengthy rendering time and implicitly defined meso-scale structures. Therefore, based on previous work [8], [51], we improve the simulation to efficiently approximate large-scale appearance using explicitly defined meso-scale geometries. Our simulation method is physically correct in accounting for shadowing and masking effects and runs at an interactive speed.

Following [51], we define

$$\bar{f}_r(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \frac{1}{a_v(\boldsymbol{\omega}_o)} \int_{s^2} f(\boldsymbol{n}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)\gamma(\boldsymbol{n}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)d\boldsymbol{n} \quad (1)$$

where $\bar{f}_r$ is the macro-scale BRDF, f is the micro-scale BRDF, γ is the Bidirectional Visible Normal Distribution Function (BVNDF) computed from the normals of meso-scale geometry; $\boldsymbol{\omega}_i$ is the incident direction and $\boldsymbol{\omega}_o$ is the outgoing direction. $a_v(\boldsymbol{\omega}_o)$ is the visible projected area of a surface patch along $\boldsymbol{\omega}_o$, and S^2 is the surface of unit sphere. We tabulate the BVNDF N and micro-scale BRDF M as matrix format:

$$\begin{aligned} \bar{f}_r(\boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_{o,k}) &\approx \frac{\sum_j f(\boldsymbol{n}_j, \boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_{o,k})\gamma(n_j, \boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_{o,k})}{a_v(\boldsymbol{\omega}_o)} \\ &= \frac{1}{a_v(\boldsymbol{\omega}_o)} \sum_j \boldsymbol{N}_{kj} \boldsymbol{M}_{jk}{}^T \end{aligned} \quad (2)$$

$$\boldsymbol{N}_{kj} = \gamma(\boldsymbol{n}_j, \boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_{o,k}) \quad (3)$$

$$\boldsymbol{M}_{jk}{}^T = f(\boldsymbol{n}_j, \boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_{o,k}) \quad (4)$$

Each row of $\boldsymbol{N}$ represents the distribution of normals of a patch subset that are visible along $\boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_{o,k}$, and each column of $\boldsymbol{M}$ stores a tabulation of cosine-weighted BRDF, where $\boldsymbol{n}_j$ defines the local frame. We can precompute the $\boldsymbol{N}$ and M and to perform efficient simulation by querying values from the matrices. Refer to [51] for details of the tabulation and applying SVD on matrix N

For matrix $\boldsymbol{M}$, instead of uniformly mapping the rotated BRDF to a 2D grid using paraboloid maps [6] as in the original method, we use a data-driven method to importance sample the micro-scale BRDF [14], which can significantly reduce the precomputation time and artifacts in simulation. We first apply a log-relative mapping for a given BRDF, $\rho = log(\rho cosMap + \varepsilon)$ [39] to reduce the order of magnitude of the specular peak, where ε is set to 10^{-3} to avoid a singularity at zero. For each incident direction, we calculate a BRDF slice of possible outgoing directions weighted by the VNDF, which is computed from the retro-reflection. To importance sample the BRDF, we warp uniform variates on a unit grid so that their density is proportional to the luminance of the associated BRDF slice. To evaluate the BRDF values in $\boldsymbol{M}$, we first compute the half-direction vector and map it to local coordinates. Then we apply the inverse warp by the NDF and use the results to perform a lookup in the original BRDF slices to compute the BRDF value. The warp function details can be found in [14]. After filling the BRDF values in $\boldsymbol{M}$, we apply random-projection to accelerate SVD for the matrix. During the simulation, we first reconstruct $\boldsymbol{M}$ from SVD, and then apply linear interpolation between samples and undo the log-relative mapping to get the BRDF value for the specific rotated normals. Figure 4 demonstrates a comparison between our improved simulation and the original [51] using gold-metallic-paint3 from the MERL dataset. The result shows that our method provides better reconstruction accuracy due to importance sampling. Also since our method does not need to compute every single direction, the pre-computation time is decreased from the original 24 hours [51] to 40 minutes for this example.

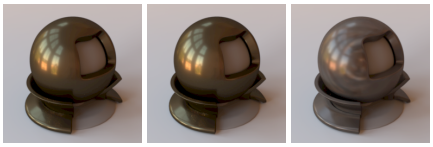

Reference Ours Wu et al.

Fig. 4: Comparison of simulations.

Our simulation is an approximation for bi-scale materials, and we only use it to build the mapping between the material scales in our training. For meso-scale structures with steep slopes coupled with high-albedo micro-scale BRDFs interreflections, which are not accounted for in our simulation, would have a noticeable impact. However, as we will demonstrate in Section 6.3, our simulation works well with our model to predict small-scale details for real world materials. Alternative simulation methods could be used if sufficient computational power is available.

3.2 Tabular BRDF Representation

The reflectance properties of a material are embedded in a 4D macro-scale BRDF, which includes millions of entries when incident and outgoing directions are densely sampled. It is almost impossible to directly learn from such a representation. We need to intelligently sample the BRDF and keep its important features for learning. We represent macro-scale BRDF as a 2D tabular format, where each entry represents the BRDF value corresponding to a specific combination of incident and outgoing directions. The benefit of the 2D format is that the directions vary continuously,

which allows us to learn BRDF gradients. We examine four tabular BRDF representations by analyzing their physical interpretations. We denote r, c as row and column index of the tabular BRDF and ω_i, ω_o as incident and outgoing directions respectively.

Sphere Table. The sphere table is constructed by fixing viewing direction and varying incident direction sampled from a hemisphere [19], [45]. Figure 5 shows the specular peak concentrates on the center of the table and the diffuse part is located around the specular peak. The edge of the sphere represents the BRDF values of grazing angles. In the following equations, resolution n is used for normalization.

$$\boldsymbol{\omega}_o = (0,0,1)$$
$$\boldsymbol{\omega}_i = (1-\frac{2r}{n}, 1-\frac{2c}{n}, \sqrt{1-(1-\frac{2r}{n})^2-(1-\frac{2c}{n})^2})$$

Retro-Reflectance Table. This table samples retro-reflectance directions along the hemisphere to stress the contribution of meso-scale facets whose normals are oriented towards the incident direction. Hence, we can see that the specular peak is more converged to the center compared with the Sphere Table.

$$\boldsymbol{\omega}_o = \boldsymbol{\omega}_i = (1-\frac{2r}{n}, 1-\frac{2c}{n}, \sqrt{1-(1-\frac{2r}{n})^2-(1-\frac{2c}{n})^2})$$

Dense-Sampled Table. Adapted from the UTIA parameterization [15], the Dense-Sampled Table uniformly samples $p = 45$ from the azimuthal angles $\theta \in [0, \frac{\pi}{2}]$ and $q = 36$ from the polar angle $\phi \in [0, 2\pi]$ for both incident (rows) and outgoing (columns) directions. We increase p and q from the original parameterization to avoid significant discontinuities and "aliasing" problems. We use fewer samples for ϕ because BRDF values vary smoothly along this direction. The left top corner of this format represents the specular peak. The longest diagonal line represents the retro-reflective measurements and its value decreases when θ increases. Other highlighted diagonal lines represent specular lobes at different angles. The right and bottom parts represent the grazing angles.

$$\boldsymbol{\omega}_{o,\theta} = floor(\frac{c}{q})\frac{\pi}{2p}, \qquad \boldsymbol{\omega}_{o,\phi} = (c \bmod q)\frac{2\pi}{q}$$
$$\boldsymbol{\omega}_{i,\theta} = floor(\frac{r}{q})\frac{\pi}{2p}, \qquad \boldsymbol{\omega}_{i,\phi} = (r \bmod q)\frac{2\pi}{q}$$

Half-Angle Table. The Half-Angle Table is constructed using the 2D slice in the Rusinkiewicz coordinate system. The slice is taken at $\phi_d = 90°$ because it maximizes the valid region of the θ_d-θ_h slice. The square root of θ_h is uniformly sampled from 0 to $\pi/2$, which provides denser sampling close to 0, in the region of the specular highlight. The physical meaning of this representation has been introduced in [9].

$$\theta_h = (\frac{r}{n})^2\frac{\pi}{2}, \qquad \theta_d = \frac{c}{n}\frac{\pi}{2}, \qquad \phi_d = \frac{\pi}{2}$$

We visualize different tabular BRDFs using two examples from the MERL dataset and demonstrate their physical interpretation in Fig. 5. We compare the training performance of these formats for learning bi-scale representations in section 5.1. We use the result of the comparison to choose the BRDF representation for our final bi-scale training.

3.3 Network Architecture

Figure 3 shows our end-to-end training pipeline. Given a target macro-scale BRDF, we first create the 2D tabular format and

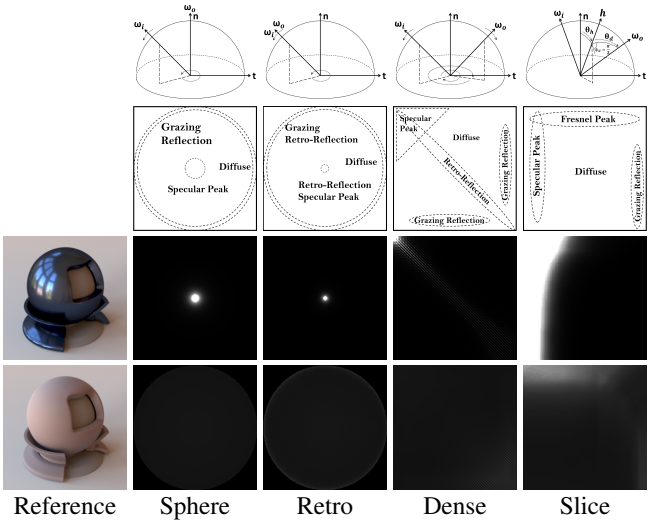

Fig. 5: Examples of different tabular BRDF formats. The first and second rows show the physical interpretations and the last two rows are examples of blue-metallic-paint2 and pink-fabric. The BRDFs are averaged over RGB-channels for the intensity comparison.

feed it into the network. Our model learns the spatial features from tabular BRDFs, then categorizes the meso-scale geometry and micro-scale BRDF into pre-defined categories and fine-tunes their corresponding parameters within each category. Based on our observations, different categories can significantly change the appearance while the physical parameters influence the details, such as the shape of the highlight. For the micro-scale BRDF, the pre-defined categories include measured BRDFs (each one is a single category) and analytical ones (a separate category). Introducing measured BRDFs can significantly increase the gamut of our fitting model due to their complex and special reflectance behaviors. The physical parameters for the the micro-scale BRDF are only defined for analytical models including roughness, diffuse and specular albedo. For the meso-scale geometry, the category is defined by different primitives that are procedurally modeled, such as grooves, woven threads and so on. Geometries generated from texture height maps are not included since they are not intuitive for editing, and they are left to future work.

Our network includes three units. The first is an analysis subnetwork of six convolution layers, each followed by a batch-normalization layer and an ReLu activation layer [25]. Each convolution has a kernel size of 7×7 and a stride of 2. The analysis subnetwork is designed to extract the spatial features from tabular BRDFs, which are fed into the following subnetworks in separate branches. The second is a meso-scale geometry unit, which includes three convolution layers and two fully connected layers. Then we split the features into two sub-branches, each of which goes through an extra fully connected output layer with $\boldsymbol{G}$ (number of categories) and $\boldsymbol{g}$ (dimension of physical parameters) hidden units respectively. The third, micro-scale BRDF, unit shares a similar branch structure ($\boldsymbol{B}$ denotes the number of category and $\boldsymbol{b}$ represent the parameters of analytical model) but with fewer convolutional layers to prevent overfitting due to its smaller search space. We find the joint learning on classification and parameter estimation have a better performance than learning separately, because the training errors on either side will be propagated to the other and back-propagated to the network. It also explains why the joint training of meso-scale geometry and micro-scale BRDF

performs better than other network structures. We will presents the results of ablation tests in Section 5.2.

The loss function for our network includes category loss ($\mathscr{L}_{GeomClass}$ and $\mathscr{L}_{BRDFClass}$) and parameters loss ($\mathscr{L}_{GeomParams}$ and $\mathscr{L}_{BRDFParams}$). We apply a softmax layer at the end of the network to calculate the Cross-Entropy for category loss, and define the L2 loss between the predicted parameters and the input labels as parameter loss. All the parameters are normalized into the same range for training. We combine those loss terms using a weighted sum. The total loss function is defined as:

$$\begin{aligned}\mathscr{L}_{total} &= \lambda_1\mathscr{L}_{GeomClass} + \lambda_2\mathscr{L}_{BRDFClass} \\ &+\lambda_3\mathscr{L}_{GeomParams} + \lambda_4\mathscr{L}_{BRDFParams}\end{aligned} \tag{5}$$

where $\lambda_1 = 1$, $\lambda_2 = 10$, $\lambda_3 = \lambda_4 = 20$.We assign more weight to the BRDF category loss to improve the network performance for distinguishing between the measured and analytical BRDFs. We found that by adding up the loss functions, the classification and regression results are penalized by each other resulting in lower training loss.

3.4 Color Restoration

Given the predicted achromatic results, we need to restore the color to match the target material. One advantage of fitting bi-scale materials is that we can can assign different colors to different facets to introduce spatially varying and anisotropic effects. We denote $\boldsymbol{C}$ as a color matrix defined in RGB and each row $\boldsymbol{c}_j = (c_{j,r}, c_{j,g}, c_{j,b})$ is a color weight for a facet. To apply the color weight, we first cluster the facets into different groups based on their normals using the BVNDF obtained from Equation 3. We can use the geometric partition to guide the clustering, where each group represents a component of geometric structure, such as the face, ridge, slope and so on. We then apply the corresponding color weights to the facets in each group by replacing the BVNDF matrix $\boldsymbol{N}$ with $(\boldsymbol{C}_r \cdot \boldsymbol{N}, \boldsymbol{C}_g \cdot \boldsymbol{N}, \boldsymbol{C}_b \cdot \boldsymbol{N})$, where

$$\boldsymbol{N}_{kj} \longleftarrow c_{j,l} \cdot \boldsymbol{N}_{kj} \quad and \quad l \in \{r,g,b\} \tag{6}$$

Assigning color weights can significantly change the appearance. For example it can introduce color variation at grazing angles. To obtain the color matrix, we use an image-based optimization with the following objective function:

$$\min_{\boldsymbol{c}_l} d(f, \bar{f}_r) \tag{7}$$

Where d is the distance metric, f is the target BRDF and $\bar{f}_r$ is defined by plugging Eq. 6 into Eq. 2.

$$\bar{f}_r(\boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_{o,k}) = \frac{1}{a_v(\boldsymbol{\omega}_o)} \sum_j Concat(||_{l\in\{r,g,b\}} c_{j,l} \cdot \boldsymbol{N}_{kj}) \boldsymbol{M}^T_{jk} \tag{8}$$

We optimize the color weights in image space, where a perfect sphere is used and lit by an environment map [11] under orthographic projection. We define the distance metric as

$$d(f_1, f_2) = \|R \cdot f_1 - R \cdot f_2\|_2 \tag{9}$$

where R is the light transport matrix transferring a BRDF into an image. Please refer to Sun et al. [46] for a detailed evaluation of different metric functions. We compute the per-pixel difference for the entire sphere to take into consideration the BRDF values at both the specular peak and the grazing angles. After optimization we can collect the color weights and reconstruct the chromatic large-scale appearance.

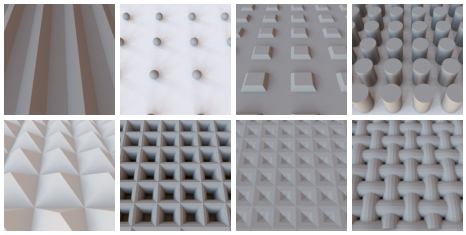

Fig. 6: Examples of different meso-scale geometry primitives used for simulation.

TABLE 1: Evaluation of Different Tabular BRDFs (8500 samples)

	Sphere	Retro-Reflectance	Dense-Sampled	Half-Angle Table
BRDF Mismatches	13	7	17	15
Geom. Mismatches	25	5	10	50
Parameter MSE	0.0064	0.0042	0.0041	0.0065

4 Implementation

We use the algorithm from Section 3.1 to create a bi-scale training set. For micro-scale BRDFs, we include 100 materials from the MERL dataset and 300 analytical BRDFs sampled from GGX Cook-Torrance models. These analytical BRDFs are created by randomly sampling roughness, diffuse and specular albedo in [0,1]. For meso-scale geometries, we procedurally generate 2000 meso-scale geometries by randomly sampling their physical parameters in their valid ranges (i.e. no collision and global scale) from 8 basic geometry categories (Fig. 6). The details of the range and definition of physical parameters can be found in supplemental materials. All the physical parameters are normalized to [0,1] during training and denormalized for reconstruction.

For comparison, we train four different networks using the same training samples but with the different tabular BRDF formats described in Section 3.2. The resolution for all tabular BRDFs is defined as $n = 256$. The Dense-Sampled Table is first created in its original resolution (45×36) and then down-sampled to 256 (directly creating Dense-Sampled Table in low resolution may lead to discontinuities and jumps). To reduce the size of the training set and time for simulation, we use achromatic tabular BRDF by averaging the color channels for measured BRDFs and sampling parameters in gray-scale for analytical BRDFs. The simulation requires roughly a week on a standard machine for each dataset. We will open source these datasets for future research projects.

We implement our network using Caffe2 [26] and the ADAM optimizer [28] with a learning rate of 10^{-3}. The total training process took around 18 hours for 250000 iterations with a batch size of 32 on a computer with an AMD Ryzen 7 1700X 8-core CPU, 16 GB RAM and an NVIDIA GeForce GTX 1070 GPU.

5 Results

We first evaluate the proposed tabular BRDF training by comparing the performance of different tabular formats and analyzing their spatial features (Section 5.1). Then we provide the experimental results from our bi-scale fitting model and conduct an ablation study (Section 5.2). We then compare our method with existing solutions (Section 5.3) and validate our method with real-world materials (Section 5.4).

TABLE 2: Ablation Study (8500 samples)

	Our Model	Model 1	Model 2
BRDF Mismatches	7	42	8
Geom. Mismatches	5	32	7
BRDF Parameter MSE	0.0037	0.0076	0.0081
Geom. Parameter MSE	0.0043	0.0056	0.0051

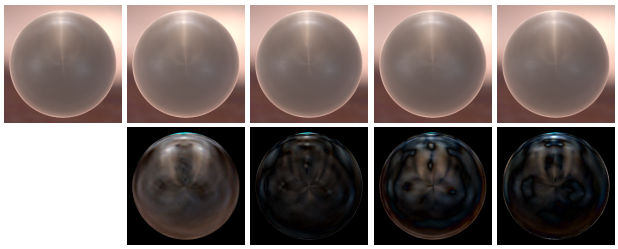

Reference Sphere Retro Dense Half-Angle

Fig. 7: Comparison of reconstructed anisotropic appearance for different tabular formats. The bottom row shows the 10 times difference between predictions and the ground truth.

5.1 Tabular BRDF Analysis

We consider the training results for different tabular formats and evaluate the embedded features. Similar to the creation of the training set, we generate 8500 test samples by randomly sampling the parameters and create a test set for each format. To avoid bias, there is no overlap between test samples and training samples. We use two metrics to evaluate network performance. For classification, we measure the number of mismatches between the predicted category and ground truth used to generate the test sample. For the parameter estimation, we compute the mean squared error for each normalized dimension between the prediction and ground truth.

We list the prediction errors in Table 1. In Fig. 7, we visualize an example of our fitting results using different tabular BRDFs. The bottom row shows the difference between the prediction and the ground truth magnified by 10 times. The results show that the Retro-Reflectance Table and the Dense-Sampled Table have better performance compared with the other two. The difference between the Sphere Table and the Retro-Reflectance Table is the shape and intensity of the specular peak. The Sphere Table has a relatively larger specular peak area and blurry boundary, which indicates that BRDF values decrease smoothly as θ_i gradually increases (small BRDF gradients). By contrast, the Retro-Reflectance Table has a smaller specular peak area and clear cut-off on the boundary, which indicates larger BRDF gradients in the retro-reflective angles. The large gradient comes from a significant jump of BRDF value, which is often caused by the change of sampling directions from the specular domain to the non-specular domain. Also note that the Retro-Reflectance Table has a higher range of BRDF values (larger intensity for specular areas and lower for diffuse areas) compared with the Sphere Table. For the Dense-Sampled Table, the sharp boundaries between highlighted diagonal lines and dark areas indicate large BRDF gradients, separating the specular and diffuse area. As for the Half-Angle Table, the specular area smoothly blends into the diffuse area with a blurry boundary due to the small gradients in both the θ_h and θ_d increasing directions. Similar to the Sphere Table, the construction of the Half-Angle Table forces the BRDF value varies slowly between neighbors. These observations, combined with the training results, suggest that the large gradient components displayed in the Retro-Reflectance Table and the Dense-Sampled Table are important features to learn and distinguish BRDF in training.

To further explain our observations, we display the feature maps from the fifth convolutional layer of the analysis sub-network in Fig. 8. As we can see, after five convolutions, basically only the shape of the specular area remains on the feature maps. We also notice that the pattern in the Retro-Reflectance Table and the Dense-Sampled Table are much clearer and the shape is more consistent with the input, while the patterns in the Sphere Table and the Half-Angle Table's feature maps look blurry. For example, the Sphere Table has a distorted circle for the specular peak with non-uniform intensity in both examples, while the Retro-Reflectance Table maintains a relatively clear and uniform circle. This indicates that the small gradients vanish and the large gradient components can remain after convolutions. Similarly, for the Half-Angle Table, most of the feature maps almost lose the pattern from the input, while the Dense-Sampled Table displays sharp diagonal line patterns in the feature maps. These results largely match with our observation: the large gradient components in the original tables lead to distinguishable patterns in the feature maps and contribute to learning material appearance.

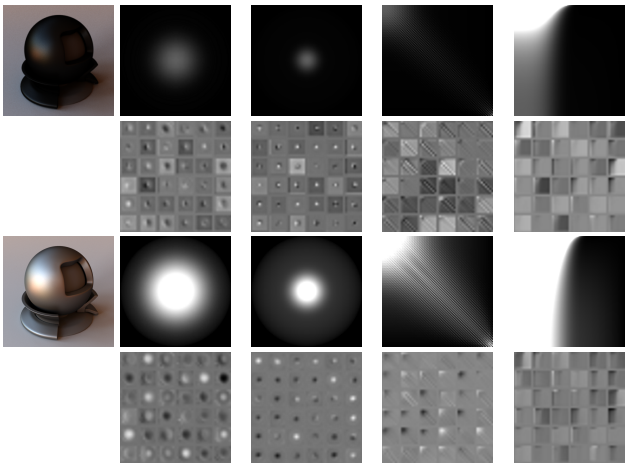

Reference Sphere Retro Dense Half

Fig. 8: Two example materials and the feature maps of tabular formats from the fifth convolutional layer. The Retro-Reflectance table and Dense-Sampled Table better maintain the shape of the specular peak compared to the other two representations.

We also experimented with 20 network settings, including different hyperparameters, weights, optimizer and number of convolutional layers, and the results can be found in the supplemental materials. Although the absolute errors vary for each experiment, the Retro-Reflectance Table in most cases maintains the lowest errors. Although our conclusion is based on experiments and observations, we still want to stress the unique property of retro-reflectance for learning materials, which has been used in recent work such as Dupuy et al. [14] to retrieve the microfacet NDF. Our experimental results can provide insights for studying retro-reflection behavior in material learning and fitting. For the rest of the paper, we use the Retro-Reflectance Table for our evaluation.

5.2 Training results

Neural network training. To evaluate bi-scale fitting results, we randomly select and render 200 samples from the previous test set. In Fig. 9, we display three of the best- and worst-cases and their four times difference from the ground truth. We use PSNR to measure the difference between our results and ground truth materials in image space. Also, we display the rendered small-scale details for comparison. In each case, the predicted images were close to indistinguishable from the ground truth for both scales.

Ablation Study. We compare our end-to-end model with two alternative structures and present results in Table 2. The first one is a separate training model, where we use three different networks to learn parameters, meso-scale geometry categories and micro-scale

TABLE 3: Comparison of 200 randomly sampled materials

	Time(sec)	Isotropic PSNR(db)	Anisotropic PSNR(db)
Ours	0.71	42.36	34.65
Wu et al.	26.78	38.43	30.38
Sun et al.	30.68	42.81	31.67

BRDF categories respectively (denoted as Model 1). We split the loss function defined in Equation 5 and use them for each network. The second one is a "one Branch" model, where no subnetwork exists and the input features go through the exact same network to output the meso-scale geometry and micro-scale BRDF (Model 2). The same loss function in Equation 5 is used. We visualize the network structures in supplemental materials. Each model is trained on the same training set and evaluated on the same test set. Compared with the results of Model 1, our network has a better performance in both classification and parameter estimation due to the existence of the shared analysis subnetwork and joint optimization. The loss of all branches back-propagates to the shared analysis subnetwork and optimize the weights together. Compared with Model 2, our network outperforms in parameter estimation with smaller MSE for micro-scale BRDF parameters. Using a single branch network to predict the meso-scale geometry and micro-scale BRDF together may cause overfitting because the search space of the micro-scale BRDF is much smaller. Our model, on the contrary, uses different subnetworks to process the meso-scale geometry and micro-scale BRDF to handle the two components, which provides more accurate fitting results.

5.3 Comparison

We compare our method with existing solutions, including the inverse bi-scale design Wu et al. [52] and Sun et al. [46] to demonstrate the fitting capability and gamut of different models. Wu et al. use a data-driven search based solution while Sun et al. fit BRDF into the GGX model using numerical optimization. We first randomly sample 200 BRDFs with uniform color-weights using our simulation method. Four randomly selected examples are presented in Fig. 10 and the statistics averaged over the 200 samples are listed in Table 3. We use the same training data for Wu et al. and our model. Based on the rendered results and PSNR value, we can observe the limited accuracy of their predictions since their results are exactly taken from the training samples. In our case, meso-scale geometries and micro-scale BRDFs are jointly optimized and defined in a continuous search space, which leads to higher reconstruction accuracy. Note that their method could be improved using their original dataset, but it is not available. We believe the comparison is still fair because the two methods are evaluated on the same training and test set. For Sun et al., the results from the two methods share similar visual appearance. However, their method does not provide small-scale details and requires more time to converge.

We also sampled 200 anisotropic materials for comparison, which are created using either anisotropic micro-scale BRDFs or multi-color weights for meso-scale geometries. Fig. 11 presents the selected results. As we can see from the zoomed images, compared with the other two methods, our method largely recovers the Fresnel effect at the grazing angle and faithfully maintains the color of the specular area from the ground truth. For Wu et al., the homogeneous color is selected from a predefined set of random colors. We can see that the discrete samples are not enough to cover a large gamut of appearances, leading to inaccurate colors for reconstruction. Sun et al. does not reconstruct BRDFs with varying color and can only yield an average color. Due to the limited gamut of the analytical model, their algorithm also fails to capture anisotropic effects. For both of these two methods, uniform colors are used to scale the achromatic BRDFs resulting in less variation in appearance. Our method, on the contrary, optimizes the spatially varying color weights on the small-scale details, leading to a wide variety of possible appearance.

We conduct a user study to further compare the three methods using Amazon Mechanical Turk. During the experiment, we gave users four images (one ground truth and three candidates) with materials rendered in the same settings, similar to the examples in Fig. 11 but without micro-structures on the corner. We asked users to compare and select the candidates in order based on the visual similarity to the ground truth (most similar to least similar). We randomly selected forty groups of materials from the previous test set with half isotropic and half anisotropic materials. We presented the candidates anonymously and in random orders. Each group of materials was evaluated by 30 users independently without a time constraint. Fig. 12 demonstrates the results. Our method gets most of the votes for most similar to ground truth in both the isotropic and anisotropic cases (76.0%), significantly outperforming the other two methods. The percentage of isotropic materials voted for our method (70.5%) is slightly less than that of anisotropic materials (81.3%). It is interesting to see that our method shares similar PSNR value with Sun et al. for isotropic materials but still gets more votes from users. We believe it is because some of the reconstructions from Sun et al. have color off on the edge of the globe (see supplemental materials). Overall, the results from the user study match with the quantitative evaluation in Table 3. See supplemental materials for all examples in the user study.

5.4 Validation on measured materials

We also validate our algorithm using 41 measured BRDFs from the material library created by Dupuy et al. [14], which are not included in the training set. Fig. 13 demonstrates the bi-scale prediction and reconstruction results. We show the rendered image of reference BRDFs, our reconstruction and predicted micro-scale BRDFs and meso-scale geometries. The average PSNR for the differences between our predictions and ground truth is 31.57dB, where the maximum is 50.56dB and the minimum is 20.75dB. From the results, we can see that our predictions capture both the intensity and color to a large extent. In Fig. 14, we qualitatively compare our results with ground truth over RGB channels separately using a slice of BRDF along the mirror direction, where we vary θ_d while fix θ_h at 0. We also show the comparison of the BRDF lopes at different angles. Our method is accurate by largely following the curve of the ground truth BRDFs in the specular part and the area of grazing angles. Refer to supplemental materials for more results.

6 Applications

6.1 Bi-scale design

The proposed bi-scale fitting provides a new material editing scheme to fine-tune a given material appearance in an intuitive way. Our algorithm supports users in designing materials with visualized small-scale details and maps them back to the large-scale appearance. Fig. 15 gives an example. Given a target material as input, our method can predict the meso-scale geometry and micro-scale BRDF, which can be used as a starting point for editing. We can let the appearance turn dimmer and duller by making the surface in the meso-scale rougher (Design 1), and increase its

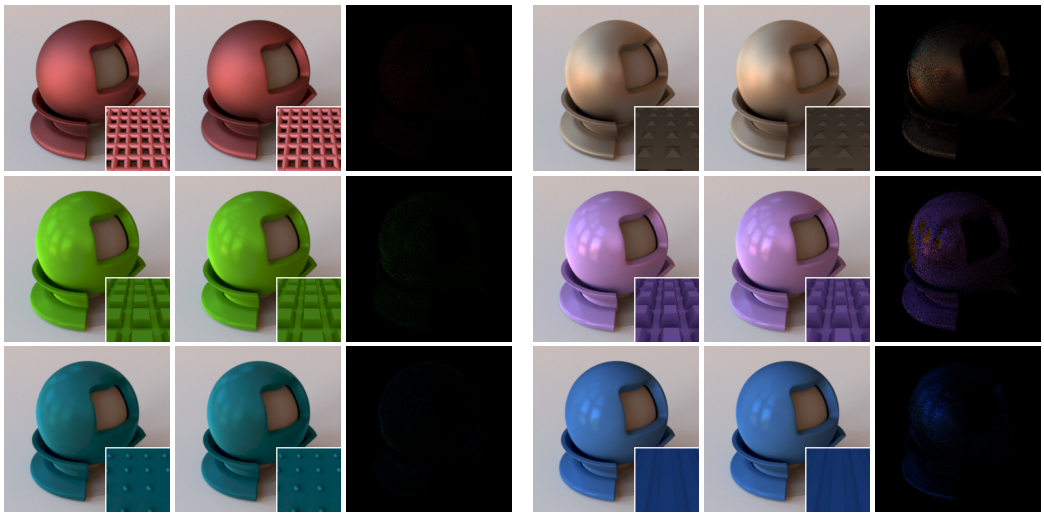

Fig. 9: The best (left side) and worst-case (right side) predictions on a set of 200 images. Averaged PSNR: 42.36dB, minimum: 22.94dB, maximum: 49.28dB

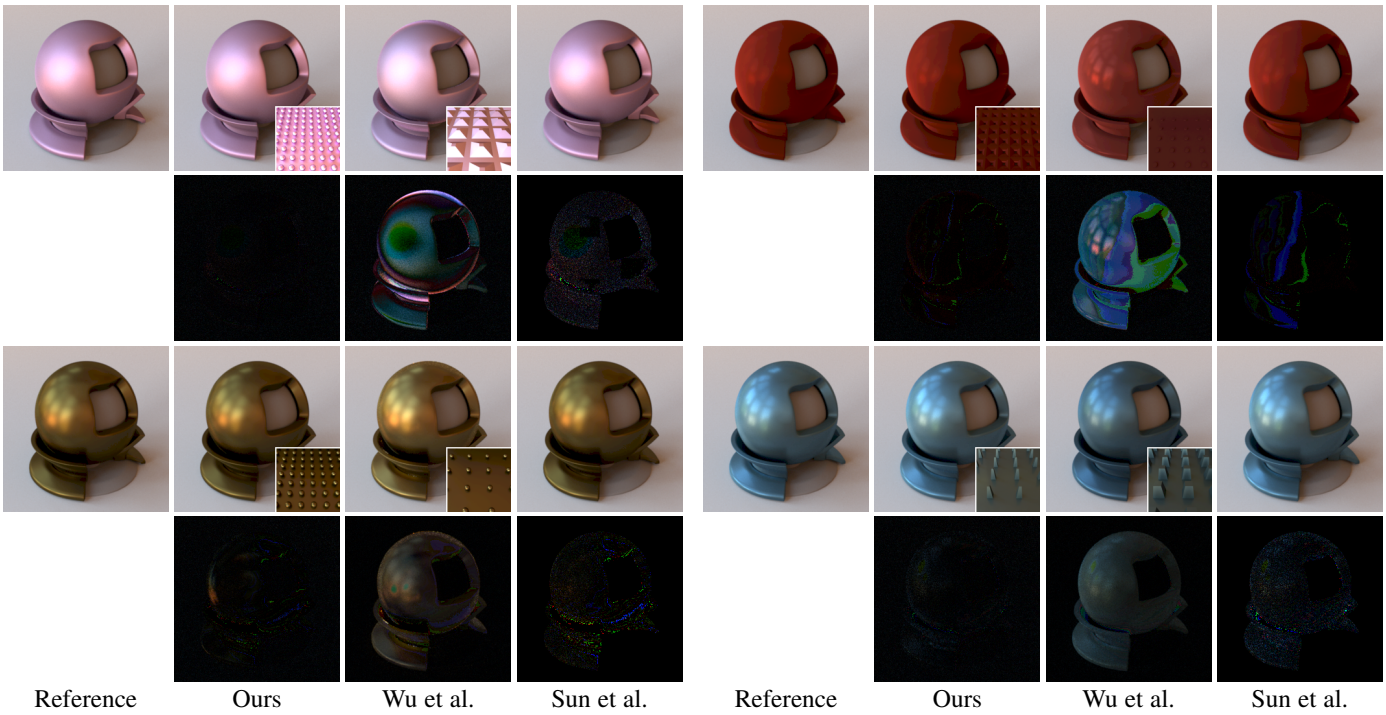

Fig. 10: Qualitative comparison (4 × difference) between our method and previous work.

gloss and brightness by smoothing the surface (Design 2). Apart from adjusting the meso-scale geometry, we can also switch the micro-scale BRDF to change the appearance (Design 3). Fig. 16 demonstrates another example to design anisotropic materials from the isotropic ones. Given fabric materials on the left, our method predicts the woven thread structures and provides reconstructions in the middle. Using the prediction results, we can adjust the colors of different woven threads to generate anisotropic silk or fabric materials displayed on the right. Overall, our method provides a bi-scale design scheme that can help users intuitively edit materials to create a large variety of appearance.

The proposed tabular BRDF training can also be extended to solve analytical fitting problems. We provide a detailed implementation and compare the results with other analytical fitting solutions in the supplemental materials.

6.2 Fabrication

The small-scale details obtained from our method could be used for material fabrication, and we show two woven examples in Fig. 17. Given a target appearance based on a fabric sample from real world, we first tabulate the BRDF and use our method to predict small-scale details. Then we render the prediction results and compare with real world sample in both scales. The same directional light and camera configurations are used to make it consistent between the rendering and photographs in large scale.

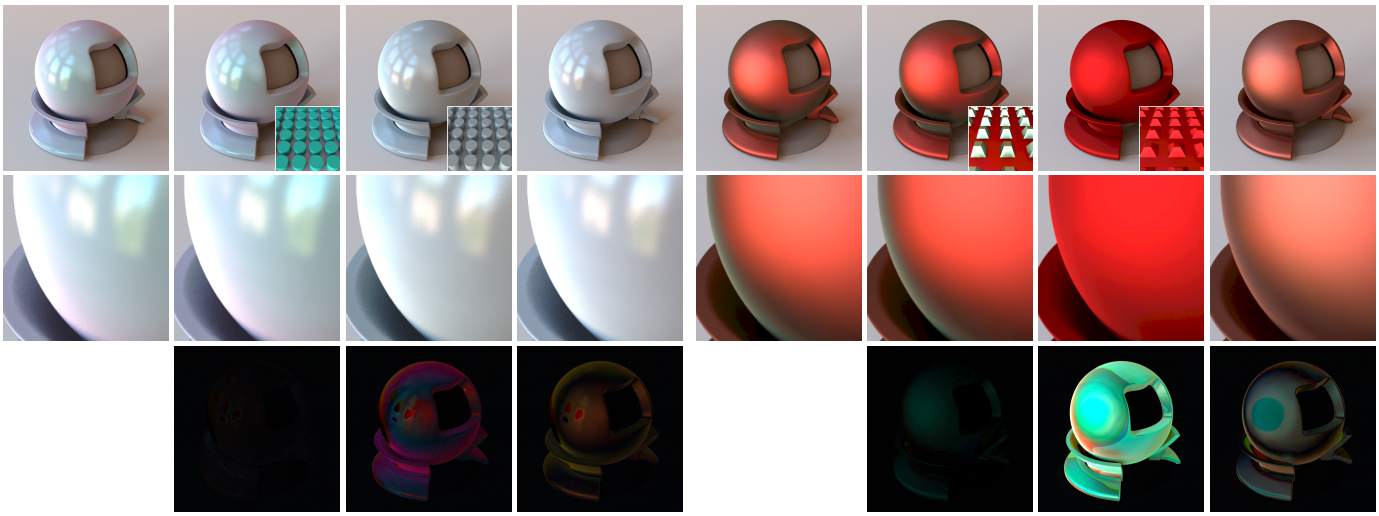

Fig. 11: Qualitative comparison (4 × difference) between our method and previous work for multi-color materials.

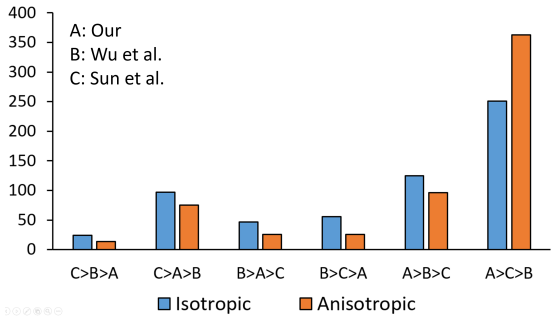

Fig. 12: Results of user study using Amazon Mechanical Turk, where we asked workers to select the images of materials from different methods in order based on the visual similarity to the ground truth. $A > B$ represents A is more similar to the ground truth than B.

Notice the inconsistency of shading for small scale is caused by the unknown lighting and camera setup of the close-up-view photographs. However, in terms of colors and structures, the predicted small-scale details largely match details of the woven fabric materials. Also considering large scale appearance, the reconstruction results are visually similar to the designed materials and the real world sample.

7 Discussion

Even with a larger fitting gamut, our bi-scale fitting model is still constrained by the diversity and complexity of training data, which could be solved by using alternative simulation methods. Fig. 18 shows a failure case. The reconstruction of material iris-purple-gem from the material library in Dupuy et al. [14] is less accurate due to the iridescent effect. Our simulation does not take into consideration the complex reflection effects like subsurface scattering and iridescent effects due to the limited computational power. We have experimented with alternative simulation solutions such as Wu et al. [53] to introduce inter-reflection, but the long precomputation time makes it impossible to simulate a large-scale training set like ours. As a trade-off, we use this relatively simple but efficient method and take different BRDF models or measured data into our training set to cover the larger fitting gamut. For future work, we could push the fitting gamut even further by introducing more diverse and complex materials or SVBRDF and find an efficient way to simulate materials with complex reflection effects.

Another limitation of our method is using a separate post-learning optimization to restore the material color. Although the optimization can produce anisotropy and multi-color at Fresnel angles, it is difficult to reconstruct the varying color in specular lobes, as demonstrated in Fig. 18. This is because the colors are assigned based on the facet groups of meso-scale geometry, which is not continuous and is low dimension. Furthermore, the local minimum of the optimization could also lead to the blending or averaging of the color of the ground truth. Ideally, given adequate computing resources, an end-to-end network that learns from chromatic data and fits colors as high dimension feature vectors would replace the optimization.

Moreover, although the focus of this paper is BRDF fitting and acquisition is out of the scope, we want to mention that the macro-scale input to our model may not be easily acquired compared with an image or photograph with target appearance. Many existing solutions [39], [54] demonstrate the possibility to reconstruct full BRDFs using sparse samples from a given image. In supplemental materials, we provide a solution to reconstruct BRDFs from the given images and then use tabular BRDFs for fitting, which enjoys higher accuracy compared with learning appearance in image space.

Finally, we want to discuss how our method fits into the real world fabrication setting. To provide a larger fitting gamut and push the boundaries, we considered many cases in this paper, such as using synthetic and measured BRDF for training and varying the color of different parts of the meso-scale geometry. Although for real life, only a limited palette of materials can be accessed, we do not consider this as a limitation. The key point of this paper is to demonstrate the possibility of using a neural representation for material fitting, which releases the traditional constraints of using a single BRDF model with a limited gamut. In other words, the fitting gamut is now defined by the training data. Therefore, simple adjustments to our method can be considered to fit in the real-world scenarios, including using the available materials as training data since measured BRDF are supported, specifying the degree of freedom for colors and varying the meso-scale geometry that is available for fabrication. Although using the limited palette of materials for fitting may lead to lower reconstruction accuracy,

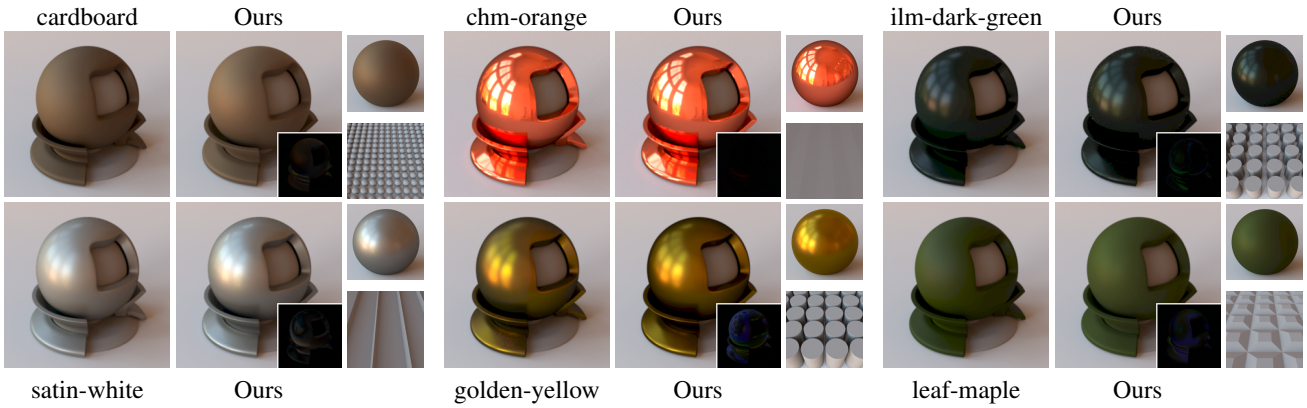

Fig. 13: Validation on measured materials from RGL material database. For each example, we show from left to right the reference, our prediction with 4 × difference, micro-scale BRDF, and meso-scale geometry.

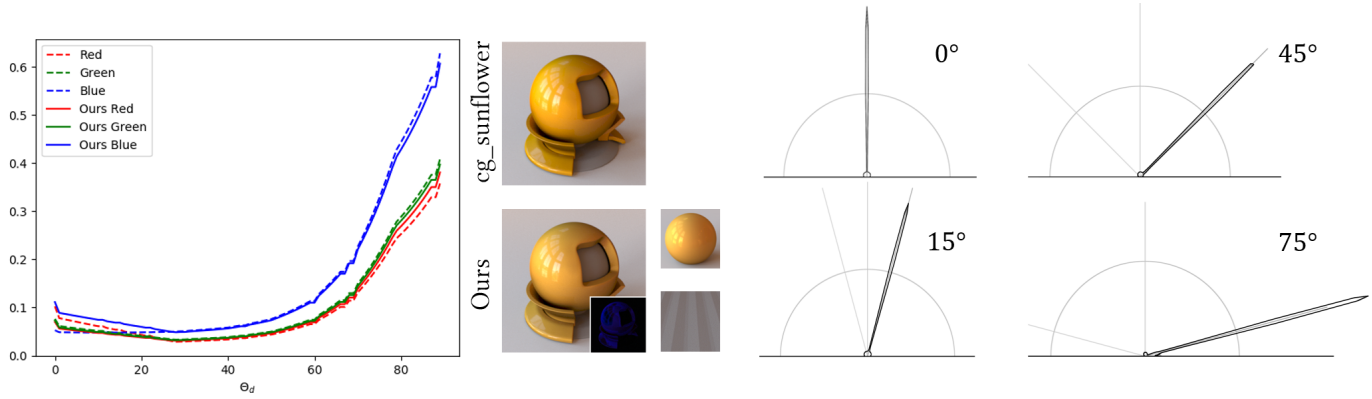

Fig. 14: On the left, we show a slice of BRDF along the mirror directions. On the right, we compare the BRDF lobes. The lobe of our results (gray) are plotted on top of and overlapped with the ground truth (black).

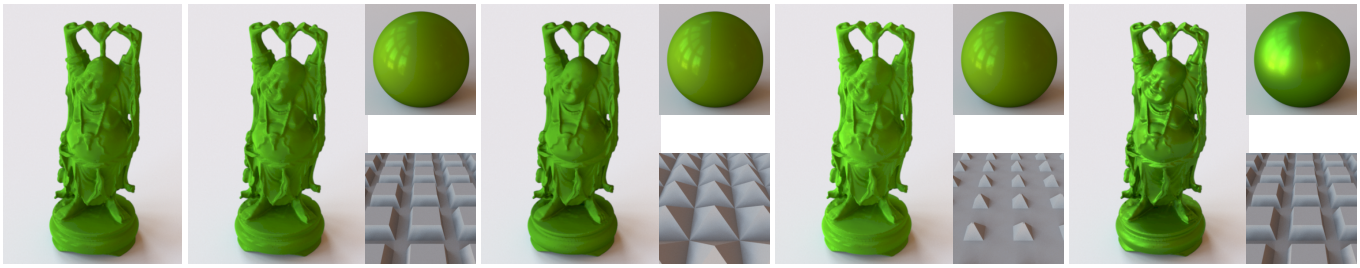

Fig. 15: Bi-scale design. Given a target material, our model can predict small-scale details, where both meso-scale geometry and micro-scale BRDF can be adjusted to create new appearances. Increasing the roughness of the surface makes the appearance more diffuse (Design 1) while smoothing the meso-scale surface makes the appearance glossy(Design 2). Users can also edit the micro-scale BRDF to change the large-scale appearance (Design 3). For each design, we present the large-scale appearance (left), micro-scale BRDF (right-top) and meso-scale geometry (right bottom).

a possibly close solution will be provided within the gamut for further iteration.

8 Conclusion

We present a learning-based solution to fit large-scale appearance using small-scale details, which provides a recipe for material design and fabrication. To accurately learn the properties from a given appearance, we introduce a new pipeline using tabular BRDF as a representation for training. Our method allows the network to directly learn the spatial features from tabular BRDFs without the bias of a particular context. Without redundant sampling for lighting and geometry, we can significantly reduce the required amount of training data and prevent overfitting to the rendering contexts. We demonstrated the effectiveness of our method by comparing with existing solutions and fitting real world materials. We believe our method can be widely used for research in efficiently and accurately learning material appearances. We also believe our bi-scale network can be used for physically realizing virtually designed materials.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 2007283

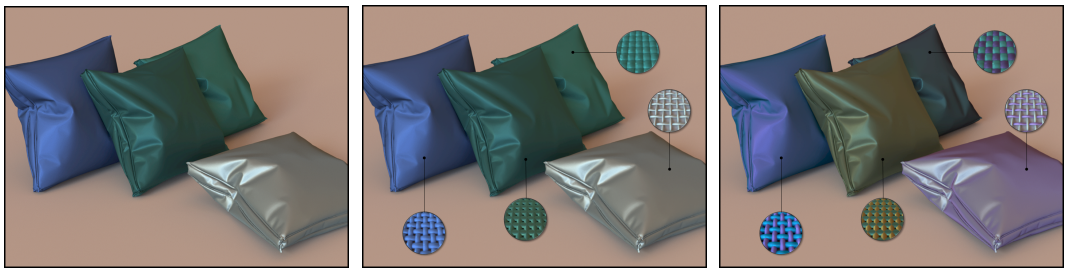

Reference Prediction and Reconstruction Editing

Fig. 16: Our algorithm can be used to design bi-scale anisotropic materials. Given target materials (left), our method can predict the small-scale woven threads (middle). Users can adjust the color of the woven threads to create various anisotropic silk or fabric materials (right).

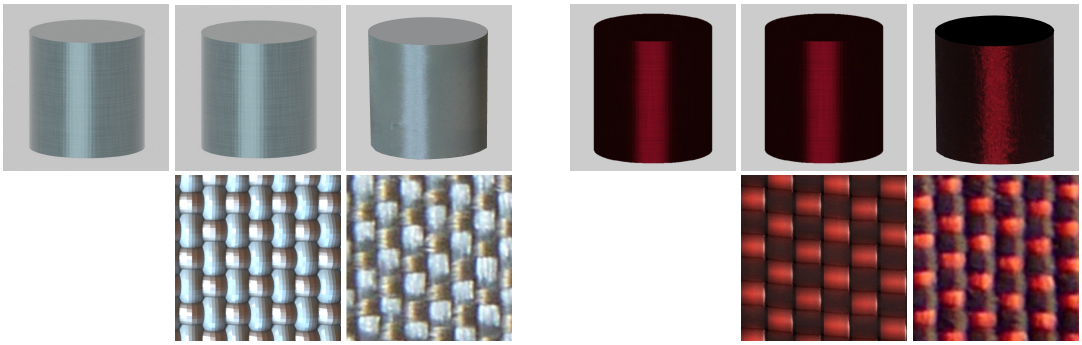

Target Prediction Fabrication Target Prediction Fabrication

Fig. 17: Physical example. We show (from left to right) the target appearance modeled as a macro-BRDF for a given sample, the reconstruction using our method and small-scale details, and an image of the real world woven sample.

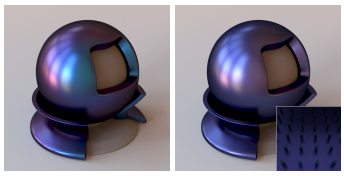

Reference Ours

Fig. 18: Failure reconstruction of material iris-purple-gem from material library [14] with iridescent effect.

REFERENCES

[1] C. Aliaga, C. Castillo, D. Gutierrez, M. A. Otaduy, J. Lopez-Moreno, and A. Jarabo. An appearance model for textile fibers. In *Computer Graphics Forum*, vol. 36, pp. 35–45. Wiley Online Library, 2017.

[2] M. Ashikmin, S. Premože, and P. Shirley. A microfacet-based brdf generator. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pp. 65–74. ACM Press/Addison-Wesley Publishing Co., 2000.

[3] M. M. Bagher, J. Snyder, and D. Nowrouzezahrai. A non-parametric factor microfacet model for isotropic brdfs. *ACM Transactions on Graphics (TOG)*, 35(5):159, 2016.

[4] M. M. Bagher, C. Soler, and N. Holzschuch. Accurate fitting of measured reflectances using a shifted gamma micro-facet distribution. In *Computer Graphics Forum*, vol. 31, pp. 1509–1518. Wiley Online Library, 2012.

[5] L. Belcour and P. Barla. A practical extension to microfacet theory for the modeling of varying iridescence. *ACM Transactions on Graphics (TOG)*, 36(4):65, 2017.

[6] S. Brabec, T. Annen, and H.-P. Seidel. Shadow mapping for hemispherical and omnidirectional light sources. In *Advances in Modelling, Animation and Rendering*, pp. 397–407. Springer, 2002.

[7] A. Brady, J. Lawrence, P. Peers, and W. Weimer. genbrdf: discovering new analytic brdfs with genetic programming. *ACM Transactions on Graphics (TOG)*, 33(4):114, 2014.

[8] E. Bruneton and F. Neyret. A survey of nonlinear prefiltering methods for efficient and accurate surface shading. *IEEE Transactions on Visualization and Computer Graphics*, 18(2):242–260, 2012.

[9] B. Burley and W. D. A. Studios. Physically-based shading at disney. In *ACM SIGGRAPH*, vol. 2012, pp. 1–7, 2012.

[10] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG)*, 1(1):7–24, 1982.

[11] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *ACM SIGGRAPH 2008 classes*, p. 32. ACM, 2008.

[12] Z. Dong, B. Walter, S. Marschner, and D. P. Greenberg. Predicting appearance from measured microgeometry of metal surfaces. *ACM Transactions on Graphics (TOG)*, 35(1):9, 2015.

[13] J. Dupuy, E. Heitz, J.-C. Iehl, P. Poulin, and V. Ostromoukhov. Extracting microfacet-based brdf parameters from arbitrary materials with power iterations. In *Computer Graphics Forum*, vol. 34, pp. 21–30. Wiley Online Library, 2015.

[14] J. Dupuy and W. Jakob. An adaptive parameterization for efficient material acquisition and rendering. In *SIGGRAPH Asia 2018 Technical Papers*, p. 274. ACM, 2018.

[15] J. Filip and R. Vávra. Template-based sampling of anisotropic brdfs. In *Computer Graphics Forum*, vol. 33, pp. 91–99. Wiley Online Library, 2014.

[16] R. W. Fleming. Visual perception of materials and their properties. *Vision research*, 94:62–75, 2014.

[17] R. W. Fleming. Material perception. *Annual review of vision science*, 3:365–388, 2017.

[18] A. Fores, J. Ferwerda, and J. Gu. Toward a perceptually based metric for brdf modeling. In *Color and Imaging Conference*, vol. 2012, pp. 142–148. Society for Imaging Science and Technology, 2012.

[19] S. Georgoulis, K. Rematas, T. Ritschel, E. Gavves, M. Fritz, L. Van Gool, and T. Tuytelaars. Reflectance and natural illumination from single-material specular objects using deep learning. *IEEE transactions on pattern analysis and machine intelligence*, 40(8):1932–1947, 2018.

[20] W. Heidrich, K. Daubert, J. Kautz, and H.-P. Seidel. Illuminating micro geometry based on precomputed visibility. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pp. 455–464. ACM Press/Addison-Wesley Publishing Co., 2000.

[21] E. Heitz. Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques*, 3(2):32–91, 2014.

[22] E. Heitz, J. Hanika, E. d'Eon, and C. Dachsbacher. Multiple-scattering microfacet bsdfs with the smith model. *ACM Transactions on Graphics (TOG)*, 35(4):58, 2016.

[23] N. Holzschuch and R. Pacanowski. A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Transactions on Graphics (TOG)*, 36(4):66, 2017.

[24] C. Innamorati, T. Ritschel, T. Weyrich, and N. J. Mitra. Decomposing single images for layered photo retouching. In *Computer Graphics Forum*, vol. 36, pp. 15–25. Wiley Online Library, 2017.

[25] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

[26] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, pp. 675–678. ACM, 2014.

[27] K. Kim, J. Gu, S. Tyree, P. Molchanov, M. Nießner, and J. Kautz. A lightweight approach for on-the-fly reflectance estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 20–28, 2017.

[28] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[29] A. Kuznetsov, M. Hašan, Z. Xu, L.-Q. Yan, B. Walter, N. K. Kalantari, S. Marschner, and R. Ramamoorthi. Learning generative models for rendering specular microgeometry. *ACM Transactions on Graphics (TOG)*, 38(6):225, 2019.

[30] Y. Lan, Y. Dong, F. Pellacini, and X. Tong. Bi-scale appearance fabrication. *ACM Trans. Graph.*, 32(4):145–1, 2013.

[31] W. Li and M. Fritz. Recognizing materials from virtual examples. In *European Conference on Computer Vision*, pp. 345–358. Springer, 2012.

[32] X. Li, Y. Dong, P. Peers, and X. Tong. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (TOG)*, 36(4):45, 2017.

[33] G. Liu, D. Ceylan, E. Yumer, J. Yang, and J.-M. Lien. Material editing using a physically based rendering network. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2261–2269, 2017.

[34] J. Löw, J. Kronander, A. Ynnerman, and J. Unger. Brdf models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG)*, 31(1):9, 2012.

[35] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, July 2003.

[36] A. Meka, M. Maximov, M. Zollhoefer, A. Chatterjee, H.-P. Seidel, C. Richardt, and C. Theobalt. Lime: Live intrinsic material estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6315–6324, 2018.

[37] G. Nam, J. H. Lee, H. Wu, D. Gutierrez, and M. H. Kim. Simultaneous acquisition of microscale reflectance and normals. *ACM Trans. Graph.*, 35(6):185–1, 2016.

[38] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of brdf models. *Rendering Techniques*, 2005(16th):2, 2005.

[39] J. B. Nielsen, H. W. Jensen, and R. Ramamoorthi. On optimal, minimal brdf sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG)*, 34(6):186, 2015.

[40] M. Oren and S. K. Nayar. Generalization of lambert's reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pp. 239–246. ACM, 1994.

[41] S. Paul. St. paul engraving. http://www.stpaulengraving.com, 2019.

[42] F. Romeiro and T. Zickler. Blind reflectometry. In *European conference on computer vision*, pp. 45–58. Springer, 2010.

[43] O. Rouiller, B. Bickel, J. Kautz, W. Matusik, and M. Alexa. 3d-printing spatially varying brdfs. *IEEE computer graphics and applications*, 33(6):48–57, 2013.

[44] S. K. Shukla and A. Kumanan. Durability of customer perceived quality of molded-in-color car bumper. Technical report, SAE Technical Paper, 2019.

[45] P.-P. J. Sloan, W. Martin, A. Gooch, and B. Gooch. The lit sphere: A model for capturing npr shading from art. In *Graphics interface*, vol. 2001, pp. 143–150, 2001.

[46] T. Sun, H. W. Jensen, and R. Ramamoorthi. Connecting measured brdfs to analytic brdfs by data-driven diffuse-specular separation. In *SIGGRAPH Asia 2018 Technical Papers*, p. 273. ACM, 2018.

[47] R. Vidaurre, D. Casas, E. Garces, and J. Lopez-Moreno. Brdf estimation of complex materials with nested learning. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1347–1356. IEEE, 2019.

[48] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pp. 195–206. Eurographics Association, 2007.

[49] T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (TOG)*, 28(3):32, 2009.

[50] Wisconsin. Wisconsin engraving. http://www.wi-engraving.com/, 2019.

[51] H. Wu, J. Dorsey, and H. Rushmeier. Physically-based interactive bi-scale material design. *ACM Transactions on Graphics (TOG)*, 30(6):145, 2011.

[52] H. Wu, J. Dorsey, and H. Rushmeier. Inverse bi-scale material design. *ACM Transactions on Graphics (TOG)*, 32(6):163, 2013.

[53] L. Wu, S. Zhao, L.-Q. Yan, and R. Ramamoorthi. Accurate appearance preserving prefiltering for rendering displacement-mapped surfaces. *ACM Transactions on Graphics (TOG)*, 38(4):137, 2019.

[54] Z. Xu, J. B. Nielsen, J. Yu, H. W. Jensen, and R. Ramamoorthi. Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM Transactions on Graphics (TOG)*, 35(6):188, 2016.

[55] Y. Yu and W. A. Smith. Pvnn: A neural network library for photometric vision. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 526–535, 2017.

[56] S. Zhao, W. Jakob, S. Marschner, and K. Bala. Building volumetric appearance models of fabric using micro ct imaging. *ACM Transactions on Graphics (TOG)*, 30(4):44, 2011.

[57] S. Zhao, W. Jakob, S. Marschner, and K. Bala. Structure-aware synthesis for predictive woven fabric appearance. *ACM Transactions on Graphics (TOG)*, 31(4):75, 2012.

[58] K. Zsolnai-Fehér, P. Wonka, and M. Wimmer. Gaussian material synthesis. *ACM Transactions on Graphics (TOG)*, 37(4):76, 2018.

Weiqi Shi is a received the B.S. degree in Electrical Engineering from Tsinghua University in 2015. He is currently a Ph.D. student in computer graphics at Yale University His research interests include material modeling, visual perception and machine learning

Julie Dorsey is the Frederick W. Beinecke Professor of Computer Science at Yale University and the founder and chief scientist of Mental Canvas, Inc. She came to Yale in 2002 from MIT, where she held tenured appointments in both the Department of Electrical Engineering and Computer Science (EECS) and the School of Architecture. She received undergraduate degrees in architecture and graduate degrees in computer science from Cornell University. Her research interests include material and texture models, sketch-based modeling, and creative applications of AI. She has received several professional awards, including MIT's Edgerton Faculty Achievement Award, a National Science Foundation CAREER Award, an Alfred P. Sloan Foundation Research Fellowship, along with fellowships from the Whitney Humanities Center at Yale and the Radcliffe Institute at Harvard.

Holly Rushmeier is the John C. Malone Professor of computer science at Yale University. Her current research interests include acquiring and modeling material appearance, applications of human perception to realistic rendering, and applications of computer graphics in cultural heritage. She received the Ph.D. in mechanical engineering from Cornell University in 1988. Between receiving the Ph.D. and arriving at Yale, she held positions with Georgia Tech, NIST, and IBM Watson Research. She is a Fellow of the Eurographics Association and the ACM, and was the recipient of the 2013 ACM SIGGRAPH Computer Graphics Achievement Award.