

# Neural Photometry-guided Visual Attribute Transfer

Carlos Rodriguez-Pardo, Elena Garces

**Abstract**—We present a deep learning-based method for propagating spatially-varying visual material attributes (e.g. texture maps or image stylizations) to larger samples of the same or similar materials. For training, we leverage images of the material taken under multiple illuminations and a dedicated data augmentation policy, making the transfer robust to novel illumination conditions and affine deformations. Our model relies on a supervised image-to-image translation framework and is agnostic to the transferred domain; we showcase a semantic segmentation, a normal map, and a stylization. Following an image analogies approach, the method only requires the training data to contain the same visual structures as the input guidance. Our approach works at interactive rates, making it suitable for material edit applications. We thoroughly evaluate our learning methodology in a controlled setup providing quantitative measures of performance. Last, we demonstrate that training the model on a single material is enough to generalize to materials of the same type without the need for massive datasets.

**Index Terms**—Artificial intelligence, Artificial neural network, Machine vision, Image texture, Graphics, Computational photography

## 1 INTRODUCTION

THE development of effective and editable material models is becoming increasingly important so that users of virtual prototyping, video-games or AR/VR applications can have compelling and realistic experiences. For instance, allowing for the creation of virtual environments that surpass the uncanny valley, or empowering artists to create breathtaking visual settings. Effective material representations require understanding the properties that uniquely define them, which we refer to as *visual material attributes*. These attributes are spatially-varying parameters that maintain spatial coherency with respect to the material structure, while remaining invariant to changes in the scene illumination or the geometry of the underlying object. For example, they may represent optical properties of a microfacet spatially-varying BRDF [1] (albedo, normals, roughness, anisotropy, etc.), but also artistic stylizations or higher-level properties, as in semantic segmentation masks.

Obtaining these attributes for *large* material samples is problematic (e.g. requiring large capturing setups or tedious manual input), and can be addressed by estimating these properties in small exemplars and, later, propagating—or *transferring*—them to the large input image, which serves as guidance. This propagation requires adapting to the local and global spatial regularities of the material, which can be challenging if the input guidance image suffers from inhomogeneous illumination, unknown scale, or affine distortions (see Figure 1, second column).

The problem has been formulated within the context of image analogies [3], and addressed via PatchMatch-based synthesis [4], look-up-tables [5], or neural networks [6], by



Fig. 1: The input exemplars on the left-bottom are transferred to the input guidances (second column) using images of the material taken under multiple illuminations (photometric input) as training data. Using this data for training, Texler *et al.* [2] fail to generalize to geometric distortions and illumination conditions not present in the training set.

looking for repetitive patterns, and match color statistics or image gradients. However, most of these methods are prohibitive due to their runtime performance or do not generalize to any kind of visual attribute. Neural networks, in particular, have proven successful and efficient for the task of video stylization [2], where the user inputs a few editing exemplars which are propagated at interactive rates to the rest of the video. However, as shown in this paper, such approach does not generalize to novel illumination conditions.

We propose a novel learning-based method to propagate any visual material attribute—estimated locally for a material—to larger samples of it. We train a neural network per material using image-to-image translation methods, making use of a policy of data augmentation that makes the transfer invariant to affine transformations (scale, rotations and shears). As opposed to other methods that use synthetic datasets for training and evaluation, our method is robustly tested using a real dataset. Further, illumination invariance is obtained by feeding the network with multiple images of the material taken under a diverse set of illuminations, which composes

- Carlos Rodriguez - Pardo is with SEDDI (28007, Madrid, Spain) and with Universidad Carlos III de Madrid (28005, Madrid, Spain).  
E-mail: carlos.rodriguezparado.jimenez@gmail.com
- Elena Garces is with SEDDI (28007, Madrid, Spain) and with Universidad Rey Juan Carlos (28933, Madrid, Spain)  
E-mail: elena.garces@seddi.com

our *Photometric dataset*. Our models can be trained in less than a minute and generalize to materials with the same microstructure.

In summary, we present the following contributions:

- The first method to use photometric data to train an image-to-image translation model capable of propagating any kind of visual material attribute to larger samples of the material regardless of the illumination conditions of the input image.
- A data augmentation policy, thoroughly evaluated with a real dataset, designed to make the transfer invariant to affine deformations.
- Exhaustive comparisons with related work demonstrating that we can achieve more predictable and higher-quality mappings with a fraction of the computational cost.
- Further, we show that our trained models generalize to materials with similar microstructure as the ones used for training.

## 2 RELATED WORK

Several computer graphics and vision problems are closely related to our method. The most similar ones are those related to any form of by-example *visual attribute transfer* (e.g color, texture, style, or geometry). Besides, we also review material estimation and capture methods.

**Visual Attribute Transfer:** Refers to the problem of transferring some visual attributes (e.g. color, style, texture, or geometry) of one or many exemplars to another exemplar while preserving its *content*.

This problem can be formulated within the context of *Image Analogies* [3], in which the goal is to stylize a target un-stylized image B, given a pair of images A (un-stylized) and A' (stylized). The most common approach to tackle this problem has been via patch-based texture synthesis [7], [8], [9]. Nevertheless, recent approaches have leveraged the capabilities of deep latent spaces within convolutional neural networks to disentangle style from content [10]. A seminal work by Gatys *et al.* [11] uses a VGG-19 convolutional neural network [12] pre-trained on ImageNet [13] as a feature descriptor for images, in which style and content are related to different layers of the network, and transferred by gradient descent optimization. Their work on *style transfer* has been extended for single images [14], [15], [16], [17] and video [18], as well as for developing image-space distance metrics that resemble human perception [19]. A limitation of many of these methods are their narrow capabilities to provide predictable edits, and considerable focus nowadays is put towards this end [20], [21]. A comprehensive review on the topic of neural style transfer is provided by Jing *et al.* [22]. Our work differs from traditional style transfer approaches in the sense that we deal with a more constrained problem that requires predictable outcomes.

Exploiting the power of deep neural networks in the image analogies problem was tackled by Liao *et al.* [23] who, by assuming a semantic prior over an exemplar input image and a target one, propose a method capable of finding a bijective mapping between both inputs, enabling two-way stylizations. Single-image generative models [24] were extended to the image analogies problem in [25], by using

convolutional neural networks to generate a new image with the style of an input *style* and the *structure* of another image. These methods, however, rely heavily on content or semantic features, making them vulnerable to lighting or geometric differences between the input images; and are computationally expensive, rendering them impractical for interactive applications. Our method is robust to both geometric distortions and illumination variations, and works at interactive rates.

Similar in spirit to our method, as it explicitly considers texture variations due to illumination, is the work of Fišer *et al.* [26], which applies patch-match to provide illumination-dependent exemplar-based stylizations to cartoon pictures. In contrast, our work is meant to be illumination-invariant. Also concerned with stylization problems, Texler *et al.* [2] present a method highly related to ours. They apply patch-based training of an encoder-decoder deep neural network using key frames stylized by a user. Resembling our approach, their algorithm also follows a few-shot learning strategy using as training data a few exemplar patches. However, as opposed to our method, they do not account for the variability of the appearance of materials under different lighting and viewpoint so, as we show in Section 7, their method does not generalize to unseen illuminations or geometric variations.

*Image colorization* is concerned with colorizing a gray-scale image given a few colorized exemplars. In this problem, it is critical to infer semantic relationships between the images so that the new scene is perceptually coherent and plausible [27], [28], [29]. Similarly, *edit propagation* methods [30], [31] work by propagating strokes provided by a user to the rest of the image, removing the need for a semantic understanding of the input image. Our work is related to the latter techniques, as we perform on the feature spaces of the CNNs and do not require a large labeled dataset to effectively solve our problem, and it can also be used to propagate segmentation masks [32].

**Textured Materials:** Many real-world materials show spatial regularities, commonly referred to as *textures*. The patterns present in textures can be parameterized, which allows for low-cost material capture or synthesis models. A way to model textured materials is through BTFs (Bidirectional Texture Functions) [33], [34], a technique that uses multiple camera views and lighting angles to capture a dense sampling of the appearance of a material. Inspired by such methods, we leverage several images of the material under different illumination conditions, however, we require less data than typical BTFs capture setups [35], [36]. Similarly, the problem of extrapolating BTFs captures to larger material samples was addressed by Steinhausen *et al.* [37], [38], who propagate measured BTFs using texture synthesis. Our method is not meant to propagate full BTFs measurements but could potentially be applied to such datasets, as we illustrate on the supplementary material.

The goal of texture synthesis is to reconstruct a larger image given a small sample leveraging structural content. This is a long-standing problem in the computer graphics field and different strategies have been proposed, for instance, using PatchMatch [39], texture transport [40], point processes [41], [42], or neural networks [43], [44], [45], [46]. Also related to our work, Li *et al.* [47] capture the appearance of materials by first estimating their BRDF and, then, synthesizing the

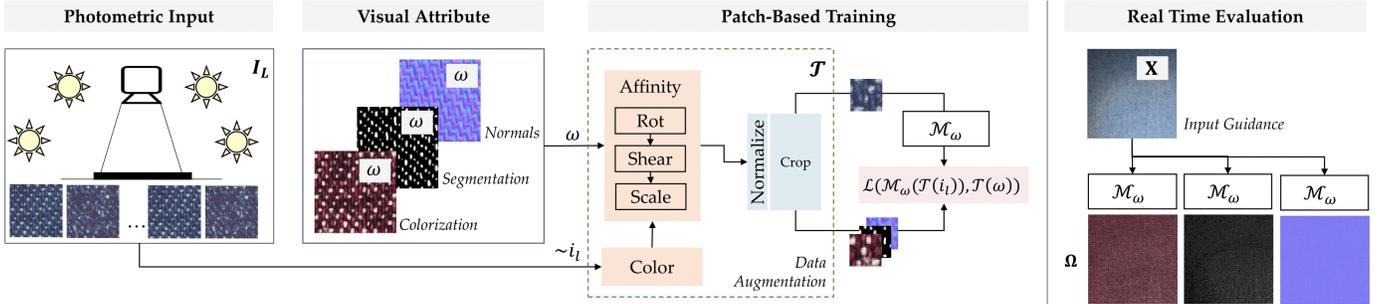


Fig. 2: Overview of the method. We learn a mapping between the photometric response  $I_L$  of the material and a visual property map  $\omega$ . We make this mapping robust to affine transformations by means of a particular policy of data augmentation used for training. We learn one model  $\mathcal{M}$  per material and visual attribute, which allows us to robustly evaluate the performance of the method under several transformations of the guidance images  $X$ . At evaluation time,  $\mathcal{M}$  can have any size. The training of  $\mathcal{M}$  per visual attribute  $\omega$  takes less than a minute.

high resolution microstructure from a dataset of measured SVBRDFs. Our problem is unlike texture synthesis, as we do not aim to create novel content but to predictably transfer visual material attributes.

**SVBRDF Estimation:** The problem of estimating a SVBRDF model from one or several images using lightweight capture setups is becoming increasingly popular in the literature. Early work [48], [49] leveraged Photometric Stereo [50] and SVBRDF manifold bootstrapping [51] for surface geometry reconstruction, while newer methods exploit the power of deep neural networks. Recent surveys by Guarnera [52] and Dong [53] contain the most relevant approaches. While our method is not meant to estimate the SVBRDF properties of a material, it can be used in combination with those techniques to create larger material assets.

There are a few methods that follow a similar paradigm to ours, transferring pre-estimated SVBRDF maps to a larger material sample. Using PatchMatch texture synthesis, Melendez *et al.* [4] transfer displacement and albedo maps from small samples of the materials. Their method is limited to daylight illumination and materials present in façades. By means of look-up-tables, and using surface normals and speculars as guidance, Riviere *et al.* [5] transfer surface reflectance captured with controlled LCD lighting to a material sample observed under natural lighting. Recently, Deschaintre *et al.* [54] fine-tune a network trained to estimate SVBRDFs [55], to work on larger material samples taking a guidance image as input. This approach is limited to transfer a pre-defined set of property maps while our method can transfer any kind. The strategy of using multiple images of the material under different illuminations as input data to is not new. Li *et al.* [56] and Ye *et al.* [57] utilize a self-augmentation strategy to make the estimation of the SVBRDF more robust to unknown environment illumination. We are inspired by these approaches to increase robustness in the model predictions.

### 3 PROBLEM FORMULATION

Our goal is to transfer a  $D$ -dimensional spatially-varying *visual attribute*  $\omega$  of a material (for example, estimated locally at high resolution) to a larger sample of it.

We formulate the problem with an image-to-image translation approach. For training, our method takes as input: a *photometric* dataset  $\mathcal{I}_L$ , and a *visual attribute*  $\omega$ . The *photometric*

dataset consisting of a number of RGB planar images of a material,  $\mathcal{I}_L = \{i_l | i_l \in \mathbb{R}^{n \times m \times 3}, |\mathcal{I}_L| \geq 1\}$ , illuminated with different light sources  $l \in L$ , of  $n \times m$  pixels. This kind of images can be either captured with specific devices [58], [59], [60], or synthetically rendered given an inverse material estimation pipeline [54], [61]. The visual attribute being a spatially-varying map  $\omega \in \mathbb{R}^{n \times m \times D}$  of any kind, and dimensions  $D$ , that maintains pixel-wise correspondence with the photometric input images  $\mathcal{I}_L$ . Figure 2 and Figure 4 show examples of these images for three different visual attributes: a stylization, a segmentation, and a normal map. Given this data for a single material, and a strategy of patch-based training, we learn a function  $\mathcal{M}$  that can be applied to a new *guidance* image of the input material (or a similar one)  $X \in \mathbb{R}^{N \times M \times 3}$  of any size  $N \times M$ , to get its corresponding visual attribute  $\Omega \in \mathbb{R}^{N \times M \times D}$ :

$$\mathcal{M} : X \rightarrow \Omega, \quad (1)$$

$$s.t. i_l \xrightarrow{\mathcal{M}} \omega, \forall i_l \in \mathcal{I}_L. \quad (2)$$

At evaluation time, the guidance image  $X$  might contain different colors, scales, illumination, or affine distortions than the images used for training. Figure 2 shows an overview of the training and evaluation processes. In Section 6, we evaluate the conditions of the input guidance image upon which the method provides robust estimations.

## 4 LEARNING FRAMEWORK

In this section, we describe the patch-based and data augmentation strategies used for training, the neural network design and loss functions, and the implementation details.

### 4.1 Patch-based Training

Each training step takes as input a pair of corresponding image patches taken from the photometric input  $\mathcal{I}_L$ , and the visual attribute  $\omega$ . Using this data alone already provides a good starting point for generalizing to unseen illumination setups. However, it is not sufficient in scenarios in which the guidance image contains variations due to image noise, a different scale, or any other affine distortion. In order to make the transfer invariant to these transformations, the network needs to be trained with the appropriate data. Data augmentation strategies are essential for reducing the amount of necessary data for training [62], [63], however,

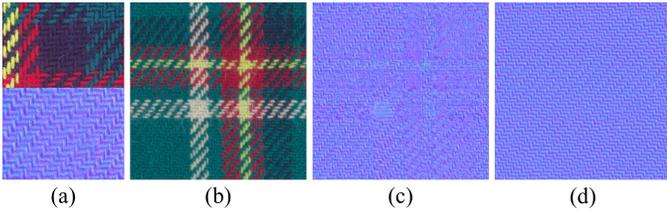


Fig. 3: The color augmentation policy takes advantage of the material structural regularities to make the transfer robust to different albedos. (a) A diffuse image of the portion of the material used for training, and its corresponding normal map. (b) Input guidance image. (c) Transferred normals without the color augmentation policy. (d) Transferred normals using color augmentation. Note that the training data does not include images containing the white yarn.

random strategies not taking into account the particular domain might degrade the quality of the prediction. We therefore follow a pre-defined data augmentation policy  $\mathcal{T}$  (illustrated in Figure 2), where random operations are performed sequentially.

**Color Augmentation:** Even if the microstructure of the material is homogeneous and can be measured using only a small patch of it, it may not be possible for the network to estimate its visual attributes for parts of the material which contain previously unseen colors. We correct this by randomly permuting the color channels of the photometric input (see Figure 3). As we show in Section 7, this data augmentation policy also helps make the model generalize the transfer to similar materials, by learning features that are more related to the structure of the material than to its color. Similar operations have been recently proposed for finding robust visual representations on self-supervised settings [64]. Only the photometric input is subject to this transformation.

**Affine Transforms:** In order to allow for material editing applications in real images, the model should generalize to images taken under camera perspectives and geometry distortions different than those present in the fronto-planar images used for training. Many of such texture irregularities can be defined as an affine transform: translation, rotation, shear, or scaling. As CNNs are shift-invariant by design [65], we propose to augment our datasets with random transformations for scalings, rotations, and shears. Those transforms can be efficiently performed to images through matrix multiplications. However, some visual property maps need to be treated specially, as the spatial transform might have a different behavior in 3D vector space, e.g. normals or tangents. In those maps, each pixel is a representation of a 3D vector. As such, we also perform the rotations and shear operations to these maps in 3D space, by multiplying each normal vector by the same affine transformation matrix applied to the 2D image. We perform the rotation around the  $Z$  axis, thus, assuming the camera sensor is parallel to the object plane.

**Cropping:** Inspired by recent work on patch-based learning [2], [66], during training, the network receives small patches of each input pair. Those patches are randomly cropped from the randomly augmented images, so the network receives a considerable amount of variations of the same material, thus making generalization possible.

## 4.2 Network Design

We follow a uni-modal image-to-image translation learning strategy, assuming there is only one correct mapping from input to output image. This approach is reasonable for the kind of transfers we test in this paper. However, it might fail for ambiguous cases where there are multiple suitable outputs for the same input, for which multi-modal approaches [67] or GANs are more advisable although harder to train. Specifically, our model  $\mathcal{M}$  is a shallow U-Net network [68] with 4 blocks of layers, containing a small number of trainable parameters, inspired on few-shot learning strategies [69], [70], [71]. This type of fully convolutional architecture is of common usage in image-space regression problems, due to its capability of efficiently learning patterns at different levels of abstraction thanks to its multi-scale design. Skip connections are added to enhance local details [68], [72]. The small number of parameters allows for faster training and inference, as well as reduced memory usage. We include further details, analysis and discussion in the supplementary material.

Note that, as opposed to previous work [54] on material transfer, that is initialized from a pre-trained network, but inspired by single-sample image synthesis methods [24], [44], we train one network per material and visual property. This strategy, although increasing training time, is key for the following reasons: First, it allows us to better understand the generalization capabilities obtained through our data augmentation policies. Second, it guarantees predictability of the trained model as every feature learned by the network is specific for each material and visual property pairs. Finally, it removes potential problems of a biased dataset as there is no cross-material or cross-domain learning. If the input dataset does not contain enough variations of the material to represent the whole material, the model will fail on areas with unseen patterns. In those cases, having a pre-trained network may help as a material prior, as in [54]. However, as we show on the supplementary material, our method obtains comparable results to pre-trained methods, with a smaller computational footprint and with the additional flexibility of not needing an expensive dataset and large models. In practice, training a single network per material and attribute is not critical as this process takes less than a minute.

**Loss function:** Choosing the appropriate loss function for a learning framework highly depends on the problem. For example, some methods [72] combine a per-pixel  $\ell_1$  metric with adversarial [73] losses, as the latter allows for better semantic mappings in multi-modal learning scenarios, whilst the former allows for improved predictability. Texler *et al.* [2] further includes a perceptual loss [74], while using a render loss is common in methods that estimate material parameters from photos [54], [55]. In our method, we show that using a  $\ell_1$  loss for training is enough for learning accurate and predictable mappings in our regression task, and binary cross-entropy loss for semantic segmentation, following standard practice in image segmentation [68]. We found that the  $\ell_2$  loss function yields overly-smooth outputs, and perceptual loss functions like *LPIPS* [19] are more prone to artifacts than pixel-wise  $\ell_{norm}$  metrics. We include an ablation study of the impact of the loss function in the supplementary material.

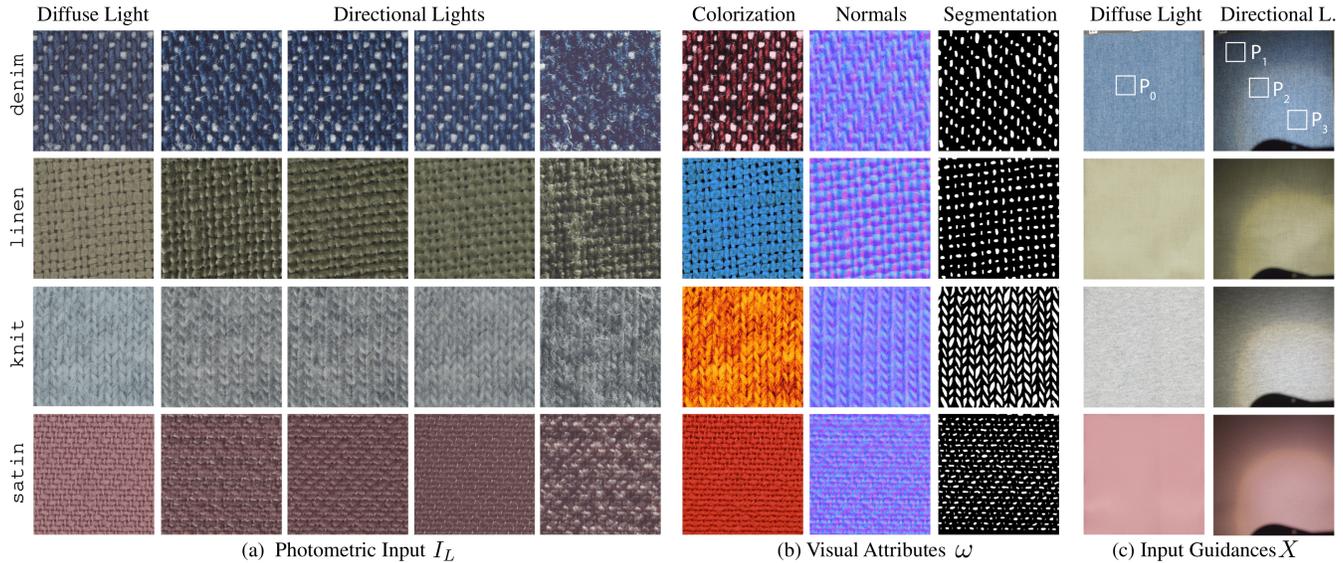


Fig. 4: An overview of our evaluation dataset. (a) Five example images of our high resolution captured data illuminated with diffuse light and four directional light sources. (b) Examples of some of our *visual property maps*: colorization, normals and yarn segmentation. (c) Evaluation images taken under diffuse and directional illumination sources. *Denim*, (c) contains labels for several patches used for the evaluation in Section 6. Those materials show different properties that may prove challenging for our visual attribute transfer task. Our *denim* and *knit* fabrics show diverse color variations (different colored dyed yarns and stochastic albedo, respectively), *linen* shows strong geometric variations and *satin* shows anisotropic optical behavior. High-resolution copies of these images are included in the supplementary material.

### 4.3 Implementation details

We use PyTorch [75] as the learning framework, Adam [76] for optimization, a learning rate of 0.002, and a batch size of 16. The training images are randomly augmented using uniform distributions by the following operations, in order: First, the photometric input is subject to the color augmentation policy. Then, both inputs and targets are randomly rotated by an angle in the  $[-90, 90]$  range, randomly sheared by an angle in the  $[-45, 45]$  range, and randomly rescaled in the  $[0.5, 2]$  range of scale factors. Then, patches of  $128 \times 128$  pixels are randomly cropped during training to generate a large data-set of images.

All inputs are always standardized using their own mean and variance. Each  $\mathcal{M}$  is trained for 1000 iterations, which takes around 1 minute on a single Nvidia 1080Ti GPU. Due to the fully convolutional nature of our models and their reduced number of trainable parameters, the guidance images  $X$  used for evaluation can be of arbitrary dimensions. We used images of up to  $5000 \times 5000$  pixels for which evaluation takes around 150ms. We refer the reader to the supplementary material for a comprehensive description and a diagram of this model, as well as further implementation details.

## 5 DATASET AND METRICS

### 5.1 Dual-Resolution Captured Data

Our method is agnostic to the capture setup [58], [59], [60], and it may work for any kind of input data (e.g. BTFs [34], [35]) as long as the photometric images are pixel-wise aligned with each other and to the visual attribute. This dataset could also be created synthetically by rendering the outputs of any SVBRDF estimation method [54]. In Section 7, we show results of our method using these acquisition pipelines.

However, for the purpose of this evaluation we chose to work with real data. The main reason is that the data

obtained with material capture devices poses extra challenges that are difficult to reproduce with render engines. For example, material irregularities, complex optical behavior, or distortions and color shifts introduced by the optical system that might cause the models to produce inaccurate estimations.

We create a dataset containing images of the same material taken with two imaging camera systems: a high-resolution camera that allows us to take pictures of  $0.7 \times 0.9$  cm, with a resolution of  $367 \times 490$  pixels and a macroscopic camera which provides images of  $11 \times 11$  cm, with a resolution of  $4800 \times 4800$  pixels. In terms of illumination, our setup has 27 different collimated light sources uniformly distributed across the hemisphere as well as diffuse illumination. We build a dataset of four different textile materials, whose complex optical behavior due to anisotropy, transmittance, directionality and microstructure [77] turns them particularly challenging for synthesis and editing operations [35], [78]. For training, we capture one image for each light source, making a total of 28 different images (Figure 4 (a)). For the evaluation set, we take one *guidance* image with diffuse illumination and another *guidance* image with a directional light source (Figure 4 (c)). The visual attributes (Figure 4 (b)) are generated automatically using photometric stereo [50] in the case of normal maps and manually by artists in the cases of colorizations and segmentation masks.

### 5.2 Attribute-specific Metrics for Evaluation

For evaluating our models, we choose domain-specific distance metrics different to those they were trained with to better understand their generalization capabilities [79]. For normal maps, we compute the cosine distance between ground truth and estimated maps, as it accounts for the geometric space in which normals lie. In the case of image

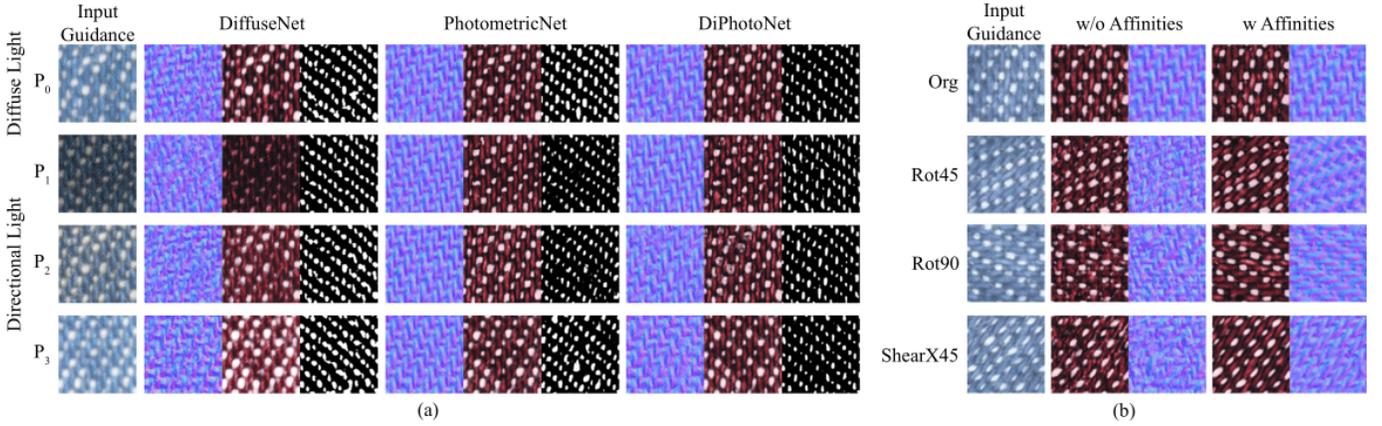


Fig. 5: Qualitative results of our models under different datasets and data augmentation configurations, for different inputs of the *denim* material. On the left (a), we show the results of our networks trained using different dataset configurations, under a guidance image taken with diffuse lighting ( $P_0$ ), and three crops of a guidance image illuminated with a directional light ( $P_1$ ,  $P_2$ ,  $P_3$ ) not present in the training set. Please refer to Figure 4 for the position of these crops on the larger guidance images. On the right (b), we show the results of our *photometricNet*, under different geometric distortions (rotations and shears) performed to its guidance images, taken under diffuse illumination.

segmentation, we evaluate the results using the Jaccard similarity coefficient (IoU) [80], which is well-suited for sparse segmentation tasks. Finally, we use the state-of-the-art metric *Learned Perceptual Image Patch Similarity* (LPIPS) [19] to evaluate the quality of the colorizations, as it has been shown to outperform  $\ell_2$  norm for visual perception tasks.

## 6 EVALUATION AND COMPARISONS

We evaluate our method in different settings. First, we assess the type and amount of photometric input data necessary for the model to generalize to different illumination conditions and image distortions. Then, we test our data augmentation strategy for affine transformations.

### 6.1 Invariance to Input Illuminations and Distortions

In this set of experiments, we aim to evaluate if the model produces the same output after changing the training data and input guidance images. To this end, we measure both the impact of different illuminations and sizes of the photometric dataset, as well as variations of the input guidance images.

#### 6.1.1 Photometric Input

In this experiment, we study the type of photometric input that makes the method invariant to different illumination conditions of the guidance image. We compare three models trained with different datasets: *diffuseNet*, which uses a single image illuminated with diffuse lighting; *photometricNet*, which takes as input 27 directional lights; and *diphotoNet*, which uses all the 28 sources. For evaluation, we use the macroscopic camera which captures a larger sample of the material at a lower resolution. We take two test guidance images with diffuse and directional lighting (Figure 4 (c)). In this experiment, all images are aligned, therefore, we follow a limited policy of data augmentation from Section 4.1, applying only color augmentation, rescaling, and crops, and leaving out rotations and shears for in-the-wild scenarios. Figure 5 shows qualitative results of our method for a selection of patches for the *denim* material. Figure 5 (a) shows that the best accuracy is in general obtained with *diphotoNet*,

i.e. training with both directional and diffuse illuminations. This is reasonable, as the model is trained using the same illuminations used during test time. Conversely, *diffuseNet*, a network only trained using diffuse lighting is less capable of generalizing under any kind of illumination source. This suggests that following a photometric approach for training material synthesis models allows for better generalization capabilities. Finally, the results for *photometricNet*, which is trained only with directional lights shows similar accuracy as *diphotoNet* while proving generalization to unseen illuminations. For all the results shown in the paper, we have chosen *photometricNet* as our model. This has an additional advantage from the usability perspective: it is relatively easy and cost-effective to build a capture setup (such as the flash light from a smartphone) or generate renders with directional lights, while it is considerably harder to recreate both types of illumination consistently.

#### 6.1.2 Dataset Size Influence

Our goal in this experiment is to understand how many images [81] taken under directional lights are needed in order to obtain the desired invariance to illumination. We train different models using reduced versions of our full 27-image dataset, and compare their results to those of the *photometricNet* trained on the full directional dataset. More precisely, we train networks using 1, 3, 9 and 18 directional lights for each material and application in our dataset following the same reduced data augmentation policy described in the previous experiment. Instead of randomly selecting light sources around the hemisphere, we perform a more sensible light source sampling. Specifically, for each reduced dataset, there is at least one light that is as close as possible to the normal of the surface in which the material lies on, thus giving more importance to frontal angles. We extend these selected lights to 3 for the reduced datasets with more than three lights. The rest of lighting sources are uniformly sampled around the hemisphere, up to a zenith angle of 70 degrees. The supplementary material contains a diagram of the distribution of lights in the hemisphere. Results are shown in Figure 6, where we see that adding more lights

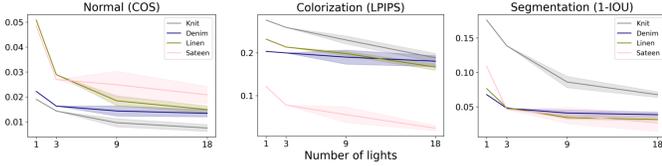


Fig. 6: Error of the reduced *photometricNets* on our ground truth guidance image (taken with diffuse illumination, not present in the training dataset) for each material in the dataset and the three visual attributes and corresponding error metrics, which are close to the normal of the surface in which the material lies. The width of the lines indicates the standard variation across 5 repeated experiments, where different light directions were randomly chosen to form the training datasets. Please refer to the supplementary material for the position of each light source.

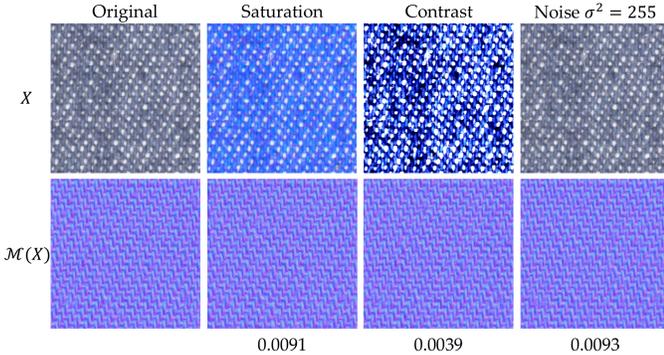


Fig. 7: Output of our model under different distorted inputs: a change of saturation, contrast, and Gaussian noise ( $\sigma^2 = 255$ ). The number at the bottom is the cosine distance with respect to the original estimation. As we can see, the output is consistent with a very small error in all cases.

to the dataset monotonically increases the generalization of every model. However, adding lights consistently shows diminishing returns, which suggests that a capture setup with around nine lights might be enough for relatively accurate estimations. Similar findings are reported in multi-image SVBRDF estimation methods [61], [82].

### 6.1.3 Image Degradations

As discussed in Section 5.1, real captured images may be subject to distortions, shifts and noise introduced by the optical capture system. To fully understand the robustness of our models with respect to these types of imperfections, we synthetically modify the saturation, contrast and noise present in the  $X$  guidance images. As shown on Figure 7,  $\mathcal{M}$  is robust to these types of degradations, even when a fair amount of details are lost in  $X$ . We refer the reader to the supplementary material for further examples of these analyses and implementation details. These images were not part of the training dataset.

## 6.2 Equivariance to Affine Transforms

Our goal in this experiment is to analyze the equivariance of our model to affine distortions of the guidance images. A function  $f$  is said to be *equivariant* with respect to a transformation  $\mathcal{T}$  if  $f(\mathcal{T}(x)) = \mathcal{T}(f(x))$ . Inspired by [83], we measure the equivariance of our model  $\mathcal{M}$  with respect to

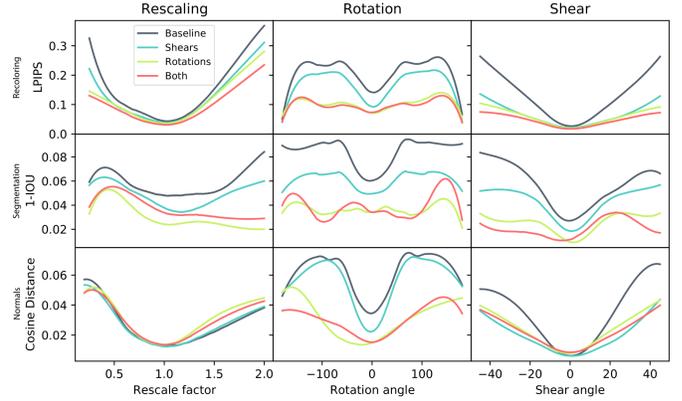


Fig. 8: Quantitative evaluation of the robustness of our networks with respect to different transformations  $\mathcal{T}$  (rescaling, rotation and shearing) of networks trained under different data augmentation policies for the *denim* material. From top to bottom, we show results on: recoloring, segmentation, and normals estimation. Lower is better for each metric.

different affine transforms  $\mathcal{T}$  performed to an input guidance image  $X$ , by computing their difference using the corresponding metric defined in Section 5.2  $d(\mathcal{M}(\mathcal{T}(X)), \mathcal{T}(\mathcal{M}(X)))$ . To achieve this, we extend the augmentation policy in the previous experiments by adding random shears and rotations to the training process. We train the *photometricNet* with and without random shears and rotations for every visual attribute in our dataset and on the *denim* material. In the case of the normals, as described in Section 4.1, we also perform the shears and rotations in the geometric space in which normals lie.

We measure their robustness with respect to three different transformations  $\mathcal{T}$ : rescalings, rotations and shears. As before, we use a guidance image taken under diffuse lighting as input for these experiments and attribute-specific distance metrics. Figure 5 (b) shows that not adding those transforms generates visual artifacts for rotations and shears. Interestingly, without affine augmentations, the models hallucinate vertical yarns, as it is the only type of data that it has seen as input. Figure 8 shows quantitative metrics for the range of transformations  $\mathcal{T}$  in which we evaluated the models. Augmenting the training dataset with both shears and rotations generally provides the best results. Furthermore, this enhanced data augmentation policy improves the robustness of the models with respect to the scale of its inputs. Notably, applying random rotations during training appears to have a larger impact than random shears on the robustness of the models. This finding suggests that applying blind policies of data augmentation [62], [64], [84], [85] may not be an optimal strategy for some applications like ours, as the networks may not learn relevant features, or overfit to noise present in the training dataset. Finally, it is worth noting that every network has the same number of parameters, and are trained for the same number of iterations, which shows that the generalization capabilities of the models can be increased at no extra training cost.

## 7 RESULTS AND COMPARISONS

In this section, we first compare our method with related approaches on image stylization and large scale SVBRDF

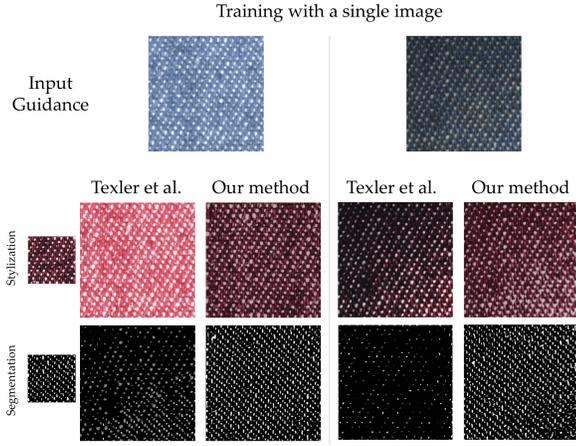


Fig. 9: Comparison of our method with Texler *et al.* [2] using a single diffuse image of the *denim* material as training data. The task is to transfer the two attributes shown (*Stylization* and *Segmentation*) to two different guidance images. Even using a single image instead of a photometric dataset, we achieve higher quality mappings at a lower cost. Further results are included in the supplementary material.

material transfer. Then, we show the capabilities of our method to generalize to similar materials to those in their training set and present its limitations<sup>1</sup>.

### 7.1 Interactive Stylizations

In the first category, the method of Texler *et al.* [74] allows artists to interactively edit a few keyframes of a video and propagate that edition to the rest of the video. As ours, they formulate this transfer problem by training an encoder-decoder network using patch-based learning. But, contrarily to us, they do not perform any data augmentation policy outside of random cropping. For a fair comparison with such method, we compare two setups: 1) using a single image as training input, and 2) using the photometric dataset. Results are shown in Figures 1 and 9. As Texler’s method is not scale invariant, in both cases the training data provided for their model has the same scale as the images used for testing. Our models have been trained with the full policy of data augmentation. In the first setup (Figure 9), we use the diffuse illumination and hence compare their output with our *diffuseNet* output. As shown, none of the methods provide high quality results but our model manages to provide closer estimations. In the second setup (Figure 1), we train their model with the photometric input. The best results are obtained with our method. Even though extending the input data using photometric cues has an impact on the quality of Texler’s results, the lack of a data augmentation policy makes the transfer fuzzier and noisier. Further, their combination of style, adversarial and pixel-wise losses fails to yield predictable mappings. These results confirm the importance of a comprehensive data-augmentation policy, such as the one we propose when using neural networks for image processing tasks of this kind. In addition, our model is trained in less time with a smaller computational footprint (1 minute vs 5 minutes).

<sup>1</sup>. In addition to the content presented in this manuscript and its supplementary material, we provide a web project which contains further results and visualizations.

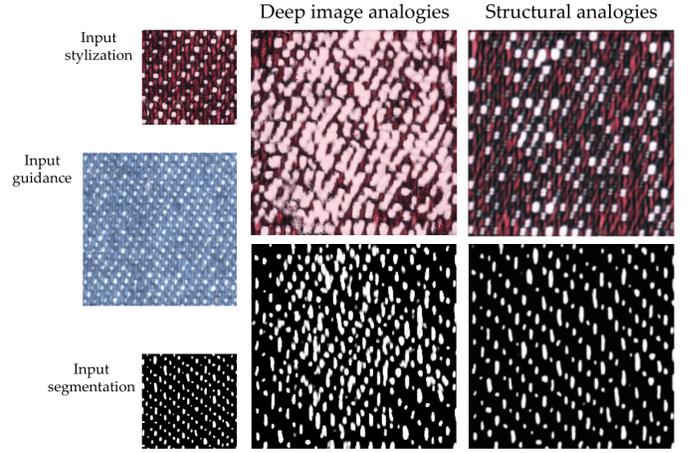


Fig. 10: Comparison of image analogies approaches for input shown on the left, of the *denim* material. From left to right: Deep Image Analogies [23] and Structural Analogies [25]. Our results shows more accurate and predictable mappings, at less computational cost than the alternatives.

Another way of formulating this visual attribute transfer problem is through *image analogies*. Using our single diffuse image for input, we compare our approach with two methods as shown in Figure 10. First, the work of Liao *et al.* [23], that uses deep latent spaces as image descriptors; and the method of Benaim *et al.* [25] that trains single-image generative models to find bijective mappings between the *structure* of one image and the *style* of another. Our method qualitatively outperforms these methods with a fraction of the computational cost: 1 minute in our case, 10 minutes in [3], 40 minutes in [23] and 10 hours in [25]. Once trained, our models can be used to evaluate any guidance image in real time for materials with similar microstructure. In contrast, image analogies methods require expensive optimizations for each guidance image. We refer the reader to the supplementary material for more comparisons with these methods.

Figure 11 shows additional results of material stylizations. In these examples, we used the method of Gatys *et al.* [11] to stylize a small patch of the material. Then, we trained a model using photometricNet and diffuseNet. As the guidance image we used a bigger image with diffuse illumination. Compared with naïve style transfer applied to the whole image, our approach provides detailed stylizations where the microstructure of the material is preserved. We further see that diffuseNet provides noisier results than photometricNet, probably due to the fact that the photometric cues help to preserve the local shading variations. We show more examples of this kind in the supplementary material.

### 7.2 Creation of Large Scale Digital Material Assets

Our method can be used to propagate SVBRDFs estimated locally in a small area of the material to larger samples. Figure 12 illustrates for a diverse set of materials that we can propagate albedo and normals estimated at high resolution in a small area of  $0.7 \times 0.9$  cm to guidance images of  $13 \times 13$  cm taken with a smartphone. Even though we train the albedo and normals models separately, which does not guarantee pixel-wise coherence between the estimated maps, the rendered images show realistic-looking materials, even with a diffuse material model. As opposed to the method

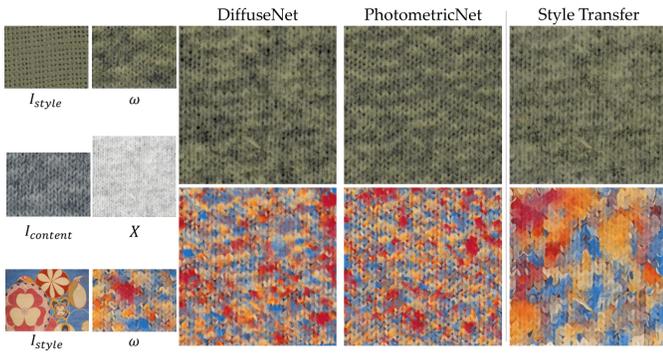


Fig. 11: Our method allows for interactive material-aware visual attribute transfers. Using an off-the-shelf style transfer algorithm [11], we can transfer the style of one image  $I_{style}$  to the content of another  $I_{content}$ , obtaining a visual attribute  $\omega$ . Training a  $\mathcal{M}_\omega$  to learn this relationship, we can find predictable style mappings, that we can transfer to guidance images  $X$ , obtaining style transfers  $\mathcal{M}_\omega(X)$ . Learning this transfer is inexpensive and allows for interactive editions. Performing this transfer directly to the guidance image generates artifacts and not-predictable mappings.

of Deschaintre *et al.* [54], which is trained to output directly Cook-Torrance [87] material layers, our method is agnostic to the parameters of the SVBRDF.

We assess the capabilities of our method on this setting using the same SVBRDF propagation scenario as proposed in [54]. Using a small crop of a synthetic SVBRDF as input, we render 27 images using the same directional light position as we used in our real dataset, and train a photometricNet to estimate the surface normals from each of those renders. We then evaluate this model using a larger area of the material, illuminated under an unknown lighting position. In Table 1, we show a quantitative comparison with [54], under different image quality metrics. As shown, our method achieves better scores on pixel-wise metrics, whilst [54] achieves better deep perceptual scores, as in the LPIPS metric [19]. This might be related to the design of our loss function: we directly minimize pixel-wise differences, while [54] is optimized using a render-aware loss. Qualitatively, as shown on Figure 13, our method obtains comparable quality mappings, with fewer artifacts. We refer the reader to the supplementary material for a larger pool of examples with a diverse set of materials. Despite capturing SVBRDF or arbitrary materials is not the goal of our method, we include in the supplementary comparisons with a direct SVBRDF acquisition method [88] using our own real dataset.

A similar setting can be used to propagate attributes leveraging BTF measurements as training data. Figure 14 shows an example using a BTF from [89]. Using Photometric Stereo [50], we compute their surface normals and train a photometricNet on a central crop of the BTF, using the captures at a camera position of  $(\phi = 0^\circ, \theta = 0^\circ)$ . We then evaluate this model using the full material surface, and a novel camera position, of  $(\phi = 15^\circ, \theta = 11^\circ)$ . As shown, our model is capable of working with captured BTF data. We provide further examples on the supplementary material.

Material ID	SSIM $\uparrow$		PSNR $\uparrow$		MSE $\downarrow$		LPIPS $\downarrow$	
	[54]	Ours	[54]	Ours	[54]	Ours	[54]	Ours
560	0,905	0,925	31,190	32,310	0,002	0,002	0,244	0,221
1581	0,645	0,675	29,740	29,920	0,006	0,005	0,398	0,446
1684	0,654	0,746	29,711	30,410	0,007	0,003	0,196	0,251
2111	0,770	0,783	32,250	32,881	0,002	0,002	0,397	0,383
Average	0,744	<b>0,782</b>	30,723	<b>31,380</b>	0,004	<b>0,003</b>	<b>0,309</b>	0,325

TABLE 1: Quantitative comparison with [54], on the studied materials and different performance metrics. As shown, our method provides better pixel-wise accuracy than [54], while their method obtains better perceptual scores.

### 7.3 Generalization to Similar Materials

In previous experiments, we have shown the performance of our models when the guidance image corresponds to the material used for training. In this experiment, we show that our models, despite being trained only on a set of captures of a single material, generalize to materials of the same category. Figure 15 shows some examples for models trained using our dataset (Figure 4), taking input guidance images of different albedos and scales. The transfer works thanks to the network design and data augmentation strategy that is designed to use microstructure details as guidance. Our method could thus be used to transfer visual attributes for a diverse set of materials by simply training one model using a single but representative material of each category.

### 7.4 Limitations

Our models are not guaranteed to provide high-quality results outside the range of input data and data augmentation policies we train them on. This limitation is common to all learning-based approaches. It is unlikely that the framework is capable of generalizing to resolutions higher to those of the training data; or down-sampled images in which the texture details are not recognizable. Furthermore, the type of transformations we apply to the training data may not represent all the possible geometric variations, or non-linear warpings that materials are subject to in the real world. For materials which exhibit a strong variation in their microstructure, which cannot be fully captured using a single photometric dataset, our patch-based approach will likely fail to generalize to the full heterogeneity. Figure 16 shows two examples of failure cases for which the test data is not included in the training data and where the affine and illumination transformations are outside the suitable range. Further examples of these limitations are included in the supplementary material.

## 8 CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed a neural visual attribute transfer framework capable of transferring, for a given material, many types of visual property maps to images of unseen patches of the same -or similar- material taken under different illumination, capture setup, and affine distortion. To our knowledge, the proposed framework is the first method capable of leveraging the optical behavior of the material to this purpose by being trained using a photometric approach. Such an approach, besides the illumination-invariance we have shown, helps the neural network learn better mappings between visual domains,

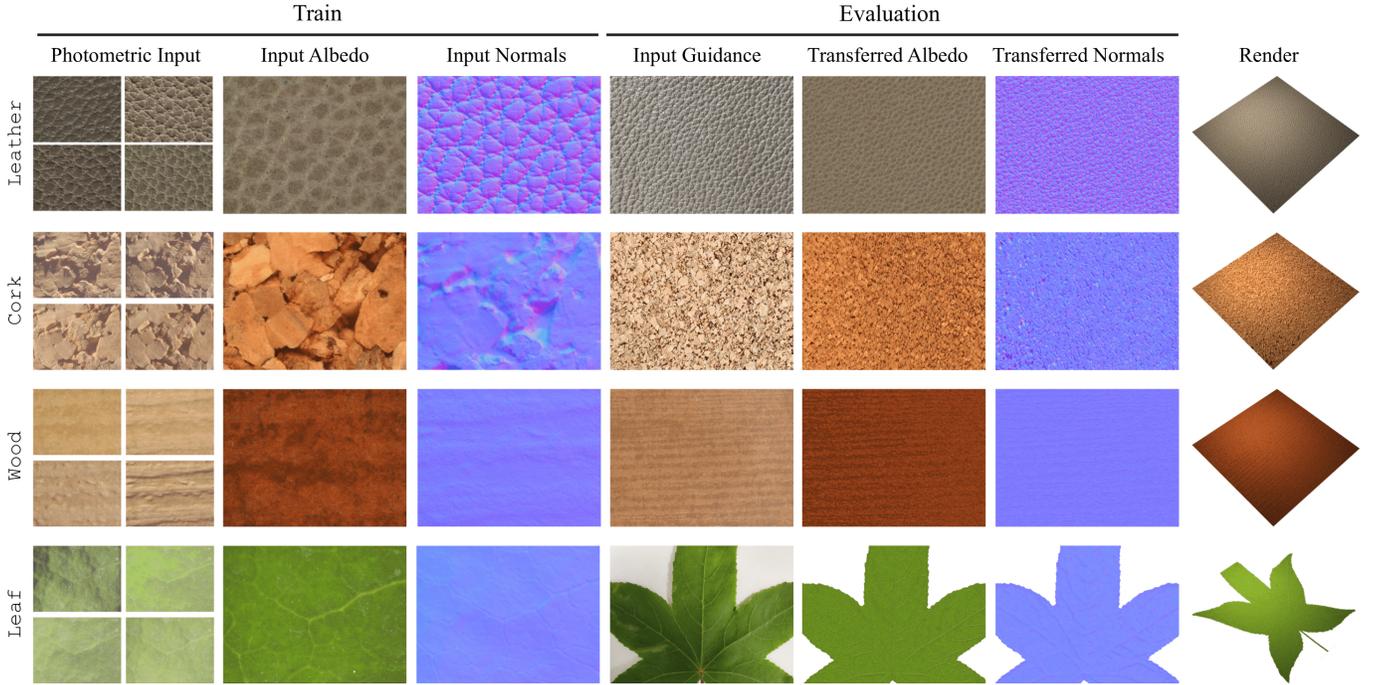


Fig. 12: Results of our framework for material capture using a smartphone. Training two models,  $\mathcal{M}_a$  and  $\mathcal{M}_n$  with a photometric dataset, we obtain respectively albedo and normals from guidance images taken under uncontrolled conditions, which can be used by render engines. Here we have used Arnold [86] and a diffuse material model. Further examples are included in the supplementary material.

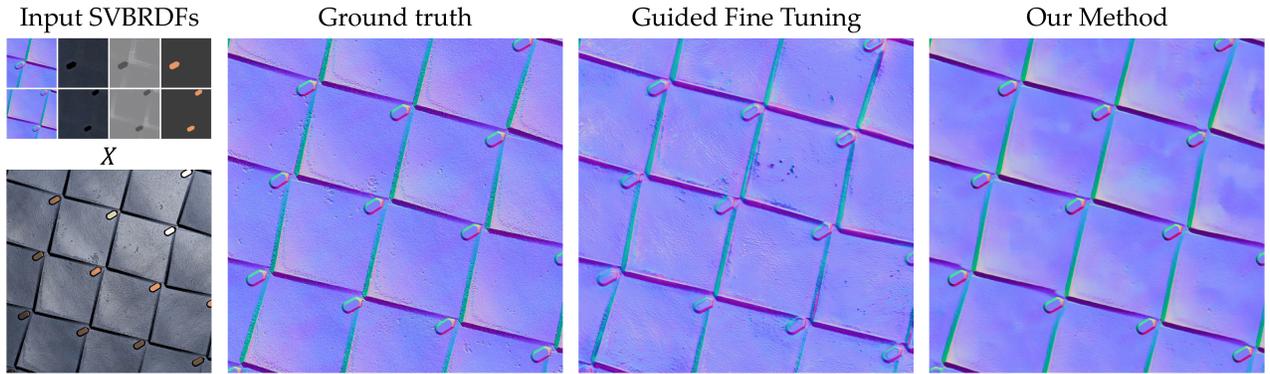


Fig. 13: Comparison of our method with the Guided Fine Tuning, by Deschaintre *et al.* [54]. Following their algorithm, we render the *Input SVBRDFs*, and train a photometricNet on those renders. As shown, our method can achieve higher quality normal maps, with fewer artifacts. Input SVBRDFs,  $X$ , ground truth and results from Guided Fine Tuning were obtained directly from [54]. Further examples are included in the supplementary material.

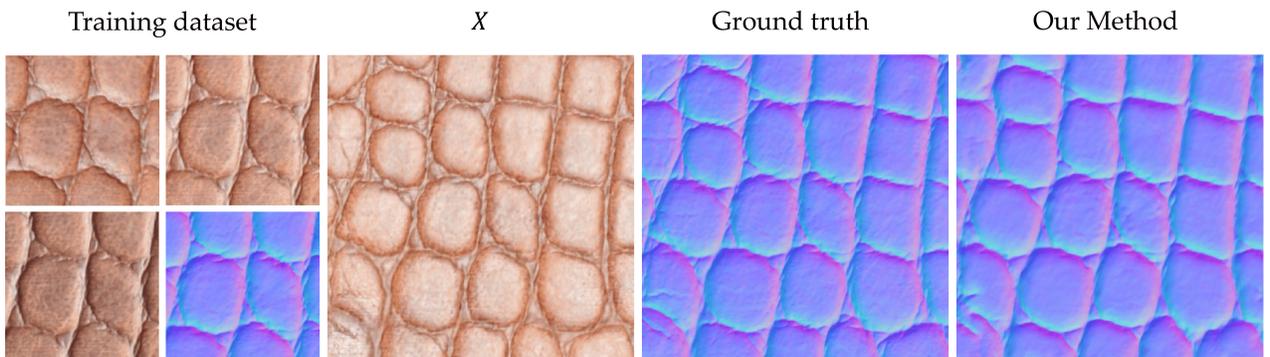


Fig. 14: Our method can work with real BTF captured data. Training a photometricNet on a crop of the material (summarized in *Training Dataset*), to output surface normals, we can estimate surface normals on larger areas of the material, even under novel viewing positions.

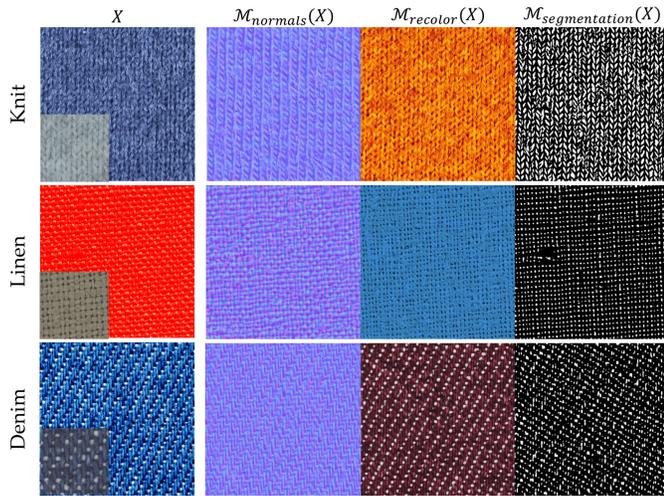


Fig. 15: Generalization capabilities of our method when evaluated on materials similar to those in their training dataset. On the top row, we show the outputs of the model trained on the *knit* in our dataset (Figure 4), and evaluated on a different guidance images. We also show examples on *linen* and *denim* fabrics, with different conditions of saturation, blur, scale and illumination. Even in very challenging cases, where the structure of the material is barely visible, as in the overly-saturated red linen or the noisy blue knit fabrics, our photometricNets can yield plausible results. We include further results and the training datasets in the supplementary material. The insets represent the training dataset by each model, as represented in Figure 4.

finding a physically-based representation of the material. Further, we have presented a comprehensive policy of data augmentation which outperforms previous work on visual attribute transfer given a single image of the material.

We have shown that our method can be used to transfer any kind of visual attribute estimated locally to larger material samples. Further, we have demonstrated that our models, although trained on a single material, generalize to materials of the same category. We think our findings will inspire future work showing that smart training strategies might alleviate the need for massive datasets.

Our method could be extended in several ways. The need for obtaining high-resolution captures taken under different illuminations may be reduced by generating rendering images through recent advances in inverse material acquisition [53]. Further training with multiple patches may help to cope with material heterogeneity [2]. Similarly, our findings suggest that extending the data augmentation policy to include 3D deformations will likely improve the accuracy. Beyond the generation of large scale digital assets for rendering, our method may have potential in other visual computing applications that require a low level understanding of the properties of the materials in real scenes. For example, the yarn segmentation application shown in the paper might be suitable as input to shape from texture applications. Specific visual attributes might be useful to identify or highlight defects for image forensics problems; or to enhance different features in real-time or AR applications.

**Acknowledgements:** Elena Garces was partially supported by a Torres Quevedo Fellowship (PTQ2018-009868). We thank

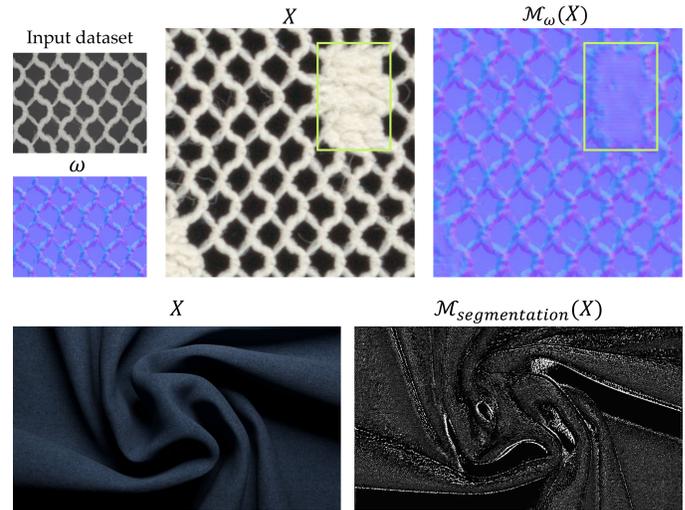


Fig. 16: Failure cases of our method. As shown on the first row, if the input dataset does not represent the heterogeneity present in the guidance  $X$ , the model fails to yield compelling results on unseen structures of the material, as shown on the green box. On the second row, we show a guidance image of the *denim* material  $X$  which exhibits strong geometric and illumination variations, outside of the range in which we train  $\mathcal{M}$  with. As such, the model shows a poor performance on the segmentation task.

Jorge López-Moreno for his feedback and David Pascual for his help with the creation of property maps.

## REFERENCES

- [1] H. C. Steinhausen, D. den Brok, M. B. Hullin, and R. Klein, "Acquiring bidirectional texture functions for large-scale material samples," 2014.
- [2] O. Texler, D. Futschik, M. Kučera, O. Jamriška, Šárka Sochorová, M. Chai, S. Tulyakov, and D. Šykora, "Interactive video stylization using few-shot patch-based training," *ACM Transactions on Graphics*, vol. 39, no. 4, p. 73, 2020.
- [3] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin, "Image analogies," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001, pp. 327–340.
- [4] F. Melendez, M. Glencross, J. Starck, and G. J. Ward, "Transfer of albedo and local depth variation to photo-textures," in *Proceedings of the 9th European Conference on Visual Media Production*, 2012, pp. 40–48.
- [5] J. Riviere, P. Peers, and A. Ghosh, "Mobile surface reflectometry," in *Computer Graphics Forum*, vol. 35, no. 1. Wiley Online Library, 2016, pp. 191–202.
- [6] I. Mazlov, S. Merzbach, E. Trunz, and R. Klein, "Neural appearance synthesis and transfer," 2019.
- [7] P. Bénard, F. Cole, M. Kass, I. Mordatch, J. Hegarty, M. S. Senn, K. Fleischer, D. Pesare, and K. Breeden, "Stylizing animation by example," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 4, pp. 1–12, 2013.
- [8] C. Barnes, F.-L. Zhang, L. Lou, X. Wu, and S.-M. Hu, "Patchtable: Efficient patch queries for large datasets and applications," *ACM Transactions on Graphics (ToG)*, vol. 34, no. 4, pp. 1–10, 2015.
- [9] O. Jamriška, Š. Sochorová, O. Texler, M. Lukáč, J. Fišer, J. Lu, E. Shechtman, and D. Šykora, "Stylizing video by example," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, pp. 1–11, 2019.
- [10] S. E. Reed, Y. Zhang, Y. Zhang, and H. Lee, "Deep visual analogy-making," in *Advances in Neural Information Processing Systems*, 2015, pp. 1252–1260.
- [11] L. A. Gatys, A. S. Ecker, and M. Bethge, "A neural algorithm of artistic style," *arXiv preprint arXiv:1508.06576*, 2015.
- [12] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *3rd International Conference on Learning Representations, ICLR*, 2015.

- [13] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [14] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1501–1510.
- [15] Y. Li, C. Fang, J. Yang, Z. Wang, X. Lu, and M.-H. Yang, "Universal style transfer via feature transforms," in *Advances in Neural Information Processing Systems*, 2017, pp. 386–396.
- [16] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proceedings of the European Conference on Computer vision (ECCV)*, 2016, pp. 694–711.
- [17] D. Chen, L. Yuan, J. Liao, N. Yu, and G. Hua, "Stylebank: An explicit representation for neural image style transfer," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1897–1906.
- [18] D. Chen, J. Liao, L. Yuan, N. Yu, and G. Hua, "Coherent online video style transfer," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1105–1114.
- [19] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [20] L. A. Gatys, A. S. Ecker, M. Bethge, A. Hertzmann, and E. Shechtman, "Controlling perceptual factors in neural style transfer," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3985–3993.
- [21] S. Gu, C. Chen, J. Liao, and L. Yuan, "Arbitrary style transfer with deep feature reshuffle," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8222–8231.
- [22] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song, "Neural style transfer: A review," *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [23] J. Liao, Y. Yao, L. Yuan, G. Hua, and S. B. Kang, "Visual attribute transfer through deep image analogy," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–15, 2017.
- [24] T. R. Shaham, T. Dekel, and T. Michaeli, "Singan: Learning a generative model from a single natural image," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 4570–4580.
- [25] S. Benaïm, R. Mokady, A. Bermano, and L. Wolf, "Structural analogy from a single image pair," in *Computer Graphics Forum*. Wiley Online Library, 2020.
- [26] J. Fišer, O. Jamriška, M. Lukáč, E. Shechtman, P. Asente, J. Lu, and D. Šykora, "Stylit: illumination-guided example-based stylization of 3d renderings," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, pp. 1–11, 2016.
- [27] M. He, D. Chen, J. Liao, P. V. Sander, and L. Yuan, "Deep exemplar-based colorization," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–16, 2018.
- [28] B. Zhang, M. He, J. Liao, P. V. Sander, L. Yuan, A. Bermak, and D. Chen, "Deep exemplar-based video colorization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8052–8061.
- [29] M. He, J. Liao, D. Chen, L. Yuan, and P. V. Sander, "Progressive color transfer with dense semantic correspondences," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 2, pp. 1–18, 2019.
- [30] X. An and F. Pellacini, "Appprop: all-pairs appearance-space edit propagation," *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3, pp. 1–9, 2008.
- [31] Y. Endo, S. Iizuka, Y. Kanamori, and J. Mitani, "Deepprop: Extracting deep features from a single image for edit propagation," in *Computer Graphics Forum*, vol. 35, no. 2. Wiley Online Library, 2016, pp. 189–201.
- [32] Y. Li, E. Adelson, and A. Agarwala, "Scribbleboost: Adding classification to edge-aware interpolation of local image and video adjustments," in *Computer Graphics Forum*, vol. 27, no. 4. Wiley Online Library, 2008, pp. 1255–1264.
- [33] K. J. Dana, B. Van Ginneken, S. K. Nayar, and J. J. Koenderink, "Reflectance and texture of real-world surfaces," *ACM Transactions On Graphics (TOG)*, vol. 18, no. 1, pp. 1–34, 1999.
- [34] T. Leung and J. Malik, "Representing and recognizing the visual appearance of materials using three-dimensional textons," *International Journal of Computer Vision*, vol. 43, no. 1, pp. 29–44, 2001.
- [35] G. Rainer, W. Jakob, A. Ghosh, and T. Weyrich, "Neural btf compression and interpolation," in *Computer Graphics Forum*, vol. 38, no. 2. Wiley Online Library, 2019, pp. 235–244.
- [36] G. Rainer, A. Ghosh, W. Jakob, and T. Weyrich, "Unified neural encoding of btfs," in *Computer Graphics Forum*, vol. 39, no. 2. Eurographics Association, 2020, pp. 1–13.
- [37] H. C. Steinhausen, D. den Brok, M. B. Hullin, and R. Klein, "Extrapolating large-scale material btfs under cross-device constraints," in *Vision, Modeling & Visualization*, D. Bommes, T. Ritschel, and T. Schultz, Eds. The Eurographics Association, 2015, pp. 143–150.
- [38] H. C. Steinhausen, R. Martin, D. den Brok, M. B. Hullin, and R. Klein, "Extrapolation of bidirectional texture functions using texture synthesis guided by photometric normals," in *Measuring, Modeling, and Reproducing Material Appearance II (SPIE 9398)*, vol. 9398, no. 14, Feb. 2015.
- [39] O. Diamanti, C. Barnes, S. Paris, E. Shechtman, and O. Sorkine-Hornung, "Synthesis of complex image appearance from limited exemplars," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 2, pp. 1–14, 2015.
- [40] M. Aittala, T. Weyrich, and J. Lehtinen, "Two-shot svbrdf capture for stationary materials," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–13, 2015.
- [41] P. Guehl, R. Allègre, J.-M. Dischler, B. Benes, and E. Galin, "Semi-procedural textures using point process texture basis functions," in *Computer Graphics Forum*, vol. 39, no. 4. Wiley Online Library, 2020, pp. 159–171.
- [42] S. Lefebvre and H. Hoppe, "Appearance-space texture synthesis," *ACM Transactions on Graphics (TOG)*, vol. 25, no. 3, pp. 541–548, 2006.
- [43] M. Elad and P. Milanfar, "Style transfer via texture synthesis," *IEEE Transactions on Image Processing*, vol. 26, no. 5, pp. 2338–2351, 2017.
- [44] Y. Zhou, Z. Zhu, X. Bai, D. Lischinski, D. Cohen-Or, and H. Huang, "Non-stationary texture synthesis by adversarial expansion," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–13, 2018.
- [45] A. Frühstück, I. Alhashim, and P. Wonka, "Tilegan: synthesis of large-scale non-homogeneous textures," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, pp. 1–11, 2019.
- [46] C. Rodriguez-Pardo, S. Suja, D. Pascual, J. Lopez-Moreno, and E. Garces, "Automatic extraction and synthesis of regular repeatable patterns," *Computers & Graphics*, vol. 83, pp. 33–41, 2019.
- [47] Y. Lin, P. Peers, and A. Ghosh, "On-site example-based material appearance acquisition," in *Computer Graphics Forum*, vol. 38, no. 4. Wiley Online Library, 2019, pp. 15–25.
- [48] A. Hertzmann and S. M. Seitz, "Example-based photometric stereo: Shape reconstruction with general, varying brdfs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 8, pp. 1254–1264, 2005.
- [49] —, "Shape and materials by example: A photometric stereo approach," in *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, vol. 1. IEEE, 2003, pp. I–I.
- [50] K. Ikeuchi, "Determining surface orientations of specular surfaces by using the photometric stereo method," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, no. 6, pp. 661–669, 1981.
- [51] Y. Dong, J. Wang, X. Tong, J. Snyder, Y. Lan, M. Ben-Ezra, and B. Guo, "Manifold bootstrapping for svbrdf capture," *ACM Transactions on Graphics (TOG)*, vol. 29, no. 4, pp. 1–10, 2010.
- [52] D. Guarnera, G. C. Guarnera, A. Ghosh, C. Denk, and M. Glencross, "Brdf representation and acquisition," in *Computer Graphics Forum*, vol. 35, no. 2. Wiley Online Library, 2016, pp. 625–650.
- [53] Y. Dong, "Deep appearance modeling: A survey," *Visual Informatics*, vol. 3, no. 2, pp. 59–68, 2019.
- [54] V. Deschaintre, G. Drettakis, and A. Bousseau, "Guided fine-tuning for large-scale material transfer," in *Computer Graphics Forum*, vol. 39, no. 4. Wiley Online Library, 2020, pp. 91–105.
- [55] V. Deschaintre, M. Aittala, F. Durand, G. Drettakis, and A. Bousseau, "Single-image svbrdf capture with a rendering-aware deep network," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–15, 2018.
- [56] 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)*, vol. 36, no. 4, pp. 1–11, 2017.
- [57] W. Ye, X. Li, Y. Dong, P. Peers, and X. Tong, "Single image surface appearance modeling with self-augmented cnns and inexact supervision," in *Computer Graphics Forum*, vol. 37, no. 7. Wiley Online Library, 2018, pp. 201–211.

- [58] G. Nam, J. H. Lee, H. Wu, D. Gutierrez, and M. H. Kim, "Simultaneous acquisition of microscale reflectance and normals," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 6, pp. 1–11, 2016.
- [59] S. Merzbach, M. Weinmann, and R. Klein, "High-quality multi-spectral reflectance acquisition with x-rite tac7," in *Proceedings of the Workshop on Material Appearance Modeling*, 2017, pp. 11–16.
- [60] R. Alcain, C. Heras, I. Salinas, J. López, and C. Aliaga, "Microscale optical capture system for digital fabric recreation," in *Proceedings of the 7th International Conference on Photonics, Optics and Laser Technology - Volume 1: PHOTOPTICS*, INSTICC. SciTePress, 2019, pp. 114–119.
- [61] Y. Guo, C. Smith, M. Hašan, K. Sunkavalli, and S. Zhao, "Materialgan: reflectance capture using a generative svbrdf model," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–13, 2020.
- [62] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, p. 60, 2019.
- [63] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila, "Training generative adversarial networks with limited data," *arXiv preprint arXiv:2006.06676*, 2020.
- [64] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [65] E. Kauderer-Abrams, "Quantifying translation-invariance in convolutional neural networks," *arXiv preprint arXiv:1801.01450*, 2017.
- [66] T. Park, A. A. Efros, R. Zhang, and J.-Y. Zhu, "Contrastive learning for unpaired image-to-image translation," *arXiv preprint arXiv:2007.15651*, 2020.
- [67] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman, "Toward multimodal image-to-image translation," in *Advances in Neural Information Processing Systems*, 2017, pp. 465–476.
- [68] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [69] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM Computing Surveys (CSUR)*, vol. 53, no. 3, pp. 1–34, 2020.
- [70] T.-C. Wang, M.-Y. Liu, A. Tao, G. Liu, B. Catanzaro, and J. Kautz, "Few-shot video-to-video synthesis," in *Advances in Neural Information Processing Systems*, 2019, pp. 5013–5024.
- [71] M.-Y. Liu, X. Huang, A. Mallya, T. Karras, T. Aila, J. Lehtinen, and J. Kautz, "Few-shot unsupervised image-to-image translation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 10 551–10 560.
- [72] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.
- [73] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, 2014, pp. 2672–2680.
- [74] O. Texler, D. Futschik, J. Fišer, M. Lukáč, J. Lu, E. Shechtman, and D. Šykora, "Arbitrary style transfer using neurally-guided patch-based synthesis," *Computers & Graphics*, vol. 87, pp. 62–71, 2020.
- [75] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [76] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR*, 2015.
- [77] C. Castillo, J. López-Moreno, and C. Aliaga, "Recent advances in fabric appearance reproduction," *Computers & Graphics*, vol. 84, pp. 103–121, 2019.
- [78] C. Kampouris, S. Zafeiriou, A. Ghosh, and S. Malassiotis, "Fine-grained material classification using micro-geometry and reflectance," in *Proceedings of the European Conference on Computer vision (ECCV)*. Springer, 2016, pp. 778–792.
- [79] L. Theis, A. v. d. Oord, and M. Bethge, "A note on the evaluation of generative models," *arXiv preprint arXiv:1511.01844*, 2015.
- [80] Y. Zhou, X. Sun, Z.-J. Zha, and W. Zeng, "Context-reinforced semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4046–4055.
- [81] J. B. Nielsen, H. W. Jensen, and R. Ramamoorthi, "On optimal, minimal brdf sampling for reflectance acquisition," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 6, pp. 1–11, 2015.
- [82] V. Deschaintre, M. Aittala, F. Durand, G. Drettakis, and A. Bousseau, "Flexible svbrdf capture with a multi-image deep network," in *Computer Graphics Forum*, vol. 38, no. 4. Wiley Online Library, 2019, pp. 1–13.
- [83] G. Benton, M. Finzi, P. Izmailov, and A. G. Wilson, "Learning invariances in neural networks," *arXiv preprint arXiv:2010.11882*, 2020.
- [84] A. Antoniou, A. Storkey, and H. Edwards, "Data augmentation generative adversarial networks," *arXiv preprint arXiv:1711.04340*, 2017.
- [85] V. Sandfort, K. Yan, P. J. Pickhardt, and R. M. Summers, "Data augmentation using generative adversarial networks (cycleGAN) to improve generalizability in ct segmentation tasks," *Scientific reports*, vol. 9, no. 1, pp. 1–9, 2019.
- [86] I. Georgiev, T. Ize, M. Farnsworth, R. Montoya-Vozmediano, A. King, B. V. Lommel, A. Jimenez, O. Anson, S. Ogaki, E. Johnston et al., "Arnold: A brute-force production path tracer," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 3, pp. 1–12, 2018.
- [87] R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," *ACM Transactions on Graphics (ToG)*, vol. 1, no. 1, pp. 7–24, 1982.
- [88] D. Gao, X. Li, Y. Dong, P. Peers, K. Xu, and X. Tong, "Deep inverse rendering for high-resolution svbrdf estimation from an arbitrary number of images," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, pp. 1–15, 2019.
- [89] M. Weinmann, J. Gall, and R. Klein, "Material classification based on training data synthesized using a btf database," in *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III*. Springer International Publishing, 2014, pp. 156–171.



**Carlos Rodriguez - Pardo** is a research engineer at SEDDI and a PhD student at the Universidad Carlos III de Madrid, Spain (UC3M). His research interests include computer vision and artificial intelligence. In 2018, he was awarded a distinction at the MSc in Artificial Intelligence at the University of Edinburgh. He completed a double BSc degree in Computer Science and Business Administration (UC3M) in 2017. He was a researcher at the Applied Artificial Intelligence Group (UC3M), working in AR applications (2013) and in data science problems (2016-2017). Carlos has served as a reviewer to conferences and journals, such as CVPR, ICCV, BMVC, ICLR, or TVCJ.



**Elena Garces** received her PhD degree in Computer Science from the University of Zaragoza in 2016. During her PhD studies, she interned twice at the Adobe (San Jose, and Seattle, USA). Her thesis dissertation focused on inverse problems of appearance capture, intrinsic decomposition from single images, video, and lightfields. She was post-doctoral researcher (2016-2018) at Technicolor R&D (Rennes, France) working on lightfields processing, and post-doctoral Juan de la Cierva Fellow (2018-2019) at the Multimodal Simulation Lab (URJC). Since 2019 she is Senior Research Scientist at SEDDI, leading the optical capture and rendering teams. She has published over 15 papers in top-tier conferences in the areas of computer graphics, vision, and machine learning, as well as authored six patents. Elena serves regularly as reviewer or PC-Member in top-tier computer vision and graphics conferences and journals such as SIGGRAPH, CVPR, ICCV, IJCV, TVCG, or EGSR.