# An Inverse Procedural Modeling Pipeline for SVBRDF Maps

YIWEI HU, Yale University, USA CHENGAN HE, Yale University, USA VALENTIN DESCHAINTRE, Adobe Research and Imperial College London, UK JULIE DORSEY, Yale University, USA HOLLY RUSHMEIER, Yale University, USA



Textured Scene

Scene With Edited Materials

Fig. 1. On the left we show a virtual scene textured only with procedural materials generated from pixel map exemplars with our method. On the right we show examples of the editing made possible by our procedural representation on the wall and floor. We change the brick color and make it deeper and smoother. We also modify the pavement color variation, its regularity and show broken tiles. These edits are made solely through the parameters made available by our method. All of the input exemplars used to the generate procedural models are presented in supplemental material.

Procedural modeling is now the de facto standard of material modeling in industry. Procedural models can be edited and are easily extended, unlike pixel-based representations of captured materials. In this paper, we present a semi-automatic pipeline for general material proceduralization. Given Spatially-Varying Bidirectional Reflectance Distribution Functions (SVBRDFs) represented as sets of pixel maps, our pipeline decomposes them into a tree of *sub-materials* whose spatial distributions are encoded by their associated mask maps. This semi-automatic decomposition of material maps progresses hierarchically, driven by our new spectrum-aware material matting and instance-based decomposition methods. Each decomposed submaterial is proceduralized by a novel multi-layer noise model to capture local variations at different scales. Spatial distributions of these sub-materials are modeled either by a by-example inverse synthesis method recovering

Authors' addresses: Yiwei Hu, Yale University, USA, yiwei.hu@yale.edu; Chengan He, Yale University, USA, chengan.he@yale.edu; Valentin Deschaintre, Adobe Research, Imperial College London, UK, deschain@adobe.com; Julie Dorsey, Yale University, USA, julie.dorsey@yale.edu; Holly Rushmeier, Yale University, USA, holly.rushmeier@yale. edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions for permissions@acm.org.

© 2022 Association for Computing Machinery. 0730-0301/2022/1-ART18 \$15.00

https://doi.org/10.1145/3502431

Point Process Texture Basis Functions (PPTBF) [Guehl et al. 2020] or via random sampling. To reconstruct procedural material maps, we propose a differentiable rendering-based optimization that recomposes all generated procedures together to maximize the similarity between our procedural models and the input material pixel maps. We evaluate our pipeline on a variety of synthetic and real materials. We demonstrate our method's capacity to process a wide range of material types, eliminating the need for artist designed material graphs required in previous work [Hu et al. 2019; Shi et al. 2020]. As fully procedural models, our results expand to arbitrary resolution and enable high level user control of appearance.

#### CCS Concepts: • Computing methodologies $\rightarrow$ Rendering.

Additional Key Words and Phrases: procedural materials, inverse material modeling, material proceduralization

#### **ACM Reference Format:**

Yiwei Hu, Chengan He, Valentin Deschaintre, Julie Dorsey, and Holly Rushmeier. 2022. An Inverse Procedural Modeling Pipeline for SVBRDF Maps. *ACM Trans. Graph.* 41, 2, Article 18 (January 2022), 17 pages. https://doi.org/ 10.1145/3502431

#### 1 INTRODUCTION

The appearance of a scene is defined by the interaction between lighting, geometries and materials. Materials define the way the light is scattered and absorbed on and within geometries. Despite progress in material authoring tools, creating a realistic material remains time consuming, even for expert artists. To reproduce existing materials, the *Material Acquisition* field defines methods for extracting properties through measurements sampling the material reflectance at different light and camera positions [Foo 1997]. Leveraging the recent progress in lightweight material acquisition, we propose the first method aiming at generating a *procedural material representation*, allowing higher-level editability as well as arbitrary scale and resolution.Examples of the application and editing of procedural representations is shown in Fig. 1.

Procedural modeling is now common in professional material modeling with tools such as Blender or Substance Designer [Adobe 2021]. Procedural models are often represented as graphs of operations, defining sub-elements of the targeted material. Each of these operations is procedural, allowing infinite resolution or scale, and high-level control through modification of parameters.

Recent methods for simplified material acquisition [Deschaintre et al. 2018, 2019; Dong 2019; Gao et al. 2019; Li et al. 2018a] target the reconstruction of analytic material maps, which compactly represent materials, but lack the editability and resolution/scale increase possibilities offered by procedural representations.

Inverse procedural modeling of materials focuses on the creation of a procedural material representation from images, or in our case, from an existing set of analytic material maps. We postulate that this ill-posed challenge requires division of a material into meaningful components which can be represented as procedures. Recently, several methods [Hu et al. 2019; Shi et al. 2020] took a first step by proposing different inverse modeling frameworks to select a procedural graph among *existing models* and optimize the nodes parameters to best match an input texture.

In this paper, we target the challenging task of creating entirely new procedural models from a set of SVBRDF pixel maps as input, eliminating the need for pre-existing artist-designed procedural models. We propose inverting the material modeling process in a sequential way by hierarchically breaking down the material into several sub-materials, then fitting their spatial distributions and material properties with procedures.

Given a material, we hierarchically decompose it into a tree of atomic sub-materials. The different sub-materials are segmented to encourage uniform statistical variation through either a user-guided, Fourier spectrum aware KNN matting approach, or an automatic instance-based segmentation algorithm. Segmented regions provide information about local texture statistics and the properties of different sub-materials. We represent the global spatial distributions of sub-materials with mask maps. Each component of the tree is then automatically converted to a full procedural model. We present a multi-layer procedural noise model based on random phase noise [Galerne et al. 2011] to model the appearance of sub-materials to capture local variations at multiple scales. To procedurally model global spatial distributions, we propose an optimization-based inverse modeling method based on a procedural texture basis function [Guehl et al. 2020] and random sampling to convert mask maps to procedures. With the proceduralized sub-materials and their procedural mask maps organized in a tree, we build a small material graph by adding optimizable operators. The material graph is optimized in a differentiable fashion based on a rendering loss to better match the input material.

Our material representation relies solely on procedural components, enabling high level editing and arbitrary scale/resolution. We use an analytic representation as input to reduce the uncertainty inherent in single-picture material acquisition. This allows our method to seamlessly benefit from progress in lightweight material acquisition methods [Deschaintre et al. 2020; Gao et al. 2019].

We show application of our pipeline on a wide range of materials and define a new taxonomy for material decomposition complexity, highlighting the existing challenges of this ill-posed task. In summary, this paper presents a novel research direction to create a procedural representation of general analytical materials without relying on pre-existing material graphs. Specifically we present the following contributions:

- We present the first pipeline for semi-automatic *generation* of procedural representation of materials. As such, our method does *not* rely on a pre-existing library of material node graphs.
- We propose a new spectrum-aware, hierarchical segmentation method for guided material segmentation.
- We define procedures and their corresponding inverse approaches to proceduralize sub-materials and their distributions.
- We offer a differentiable rendering-based optimization routine to match our procedural material to the input material during reconstruction.
- We define a procedural representation which allows for complex multi-scale edits and arbitrary scale and/or resolution

Our implementation will be publicly available.

# 2 RELATED WORK

#### 2.1 Material Acquisition

Material acquisition seeks to recover the reflectance properties of existing surfaces or objects. Our method complements this body of work since acquired materials are used as input for our pipeline to create procedural representations. While material acquisition has challenged researchers for decades as discussed in the excellent survey by Guarnera et al. [2016], significant progress was achieved in recent years by leveraging deep learning. In particular, different methods were proposed to acquire materials from one [Deschaintre et al. 2018; Li et al. 2017, 2018a,b] or multiple [Boss et al. 2020; Deschaintre et al. 2019, 2021; Gao et al. 2019; Guo et al. 2020b] pictures of surfaces or objects. While our approach also targets the creation of a material representation, we do not seek to recovering material properties from captured sample(s), but rather to proceduralize an existing material. This difference not only allows our method to benefit from material acquisition and its progress, but also to handle any existing analytical material.

# 2.2 Texture Synthesis

In our method, we leverage texture synthesis, which creates new textures with larger scale or higher resolution from an exemplar. Several surveys provide a comprehensive overview of example-based texture synthesis [Akl et al. 2018; Raad et al. 2018; Wie et al. 2009]. Texture synthesis can be classified into three families: patch re-arrangement based, arranging patches available in the original texture to synthesize a new texture [Efros and Freeman 2001; Efros



Fig. 2. A simple example of a brick material created using node graphs in Substance Designer [Adobe 2021]. This material graph is composed by a set of procedures to generate brick patterns and add details upon them. The designed material is fully procedural thus can be further edited or expanded to any resolution.

and Leung 1999; Kaspar et al. 2015], statistics based, estimating statistics of the original texture and transferring them to a new texture [Galerne et al. 2011; Galerne et al. 2012, 2017; Gilet et al. 2014; Heeger and Bergen 1995; Heitz and Neyret 2018], and neuralnetwork-based, reproducing texture or material appearance using machine learning [Bergmann et al. 2017; Henzler et al. 2021, 2020; Niklasson et al. 2021; Shocher et al. 2019; Ulyanov et al. 2016; Zhou et al. 2018]. We use the second family of methods to reproduce the local variations of our decomposed atomic components sub-materials. Alone, these methods can reproduce micro-structures well, but fail for larger scale patterns. They are also limiting in the scope of possible editability. But by making use of progressive decomposition, extending them in a multi-layer fashion and incorporating them with an optimizable post-processing step, our method can represent complex SVBRDFs as procedural models with arbitrary scale and resolution as well as high level editability.

### 2.3 Procedural Modeling

Procedural models allow artists to retain editability and produce arbitrary scale and resolution. Specialized procedural models have been proposed but are limited to specific materials such as wood or leather and mostly stationary materials [Guo et al. 2020a]. A more generic option lies in the design software allowing artists to combine procedures to create a new material (see Fig. 2). While this procedural representation provides great editability and versatility, it remains time consuming and challenging, even for expert artists, to reproduce an existing or imagined material.

Recent methods explore inverse procedural modeling of materials, aiming at the reproduction of a given material appearance using a procedural model. Hu et al. [2019] and Shi et al. (MATch) [2020] both propose different frameworks that leverage existing Substance graph procedural models and infer their parameters to match the appearance of input textures. To do so, Hu et al. relies on neural networks trained for each individual material graph, while Shi et al. (MATch) relies on gradient descent and a differentiable version of the Substance Engine. Both of these methods, however, work on the premise of known pre-existing procedural node graphs, manually created by artists. This approach requires a large, expressive enough dataset of procedural graphs to be available and non-trivial search methods to find the closest ones among hundred of options. As opposed to these methods, our proposed pipeline, does not require predefined material graphs and simply relies on a few user scribble, which typically only require 2 minutes of interaction.

A key component of our inverse procedural modeling pipeline is the procedural representation of the structure of materials. That is, the global spatial distributions of decomposed sub-materials, which we represent using a set of binary mask maps. To preserve the procedural aspect of our results we need to represent these masks procedurally. Rosenberger et al. [2009] propose a shape synthesis method to generate layered control maps but is limited to unstructured shapes with fractal-like boundaries. Alternatively, we could use an L-system [Št'ava et al. 2010], but that approach requires predefined grammars. Rather, we make procedural generation of structured mask maps possible by proposing a by-example inverse Point Process Texture Basis Functions (PPTBF) [Guehl et al. 2020] modeling approach and leveraging random sampling.

#### 2.4 Image Segmentation and Matting

Image segmentation aims at separating an image into different regions, based on some specific criteria. In our material proceduralization framework, segmentation is used for partitioning the given material into several sub-materials. Image segmentation has been extensively studied in the last decades [Khan 2014], with recent methods leveraging deep learning [Minaee et al. 2020]. Close to our challenge are the recent work of Cimpoi et al. and Bell et al. [2015; 2016] which semantically split natural scenes into their different components (such as wood, plastic, etc.). However, these methods focus on the segmentation of complete scene images and use context [Bell et al. 2015] to better recognize materials, which is not available in our material exemplars.

Most existing segmentation methods do not perform well in our context as they result in significant error on the boundaries between two sub-materials. We therefore use image matting, originally developed to segment background and foreground with fuzzy boundaries. We found alpha matting to better represent the transition between sub-materials in a SVBRDF and to preserve good quality boundaries, even after thresholding.

Many methods for image matting suggest affinity-based solutions, solving a Laplacian clustering problem [Aksoy et al. 2017; Chen et al. 2013; Levin et al. 2008]. Recent work [Cho et al. 2016; Xu et al. 2017] attempts to learn alpha maps directly using deep neural networks, but is limited to two-layer (foreground and background) alpha matting due to available training data. Aksoy et al. [2018] proposed an improved affinity-based method by combining deep semantic features with affinity Laplacian. While it supports multiple layers, it is geared toward natural scene images and requires user to provide the number of layers. Our decomposition approach of materials is similar to the one adopted by Lawrence et al. [2006]: we aim to decompose the input spatially-varying material maps into regions of similar *sub-materials*. The approach by Lawrence et al. was further developed by AppWand [Pellacini and Lawrence 2007], AppProp [An and Pellacini 2008] and Material Matting [Lepage and Lawrence 2011], which introduce segmenting measured or analytical material images to make them easy to edit. However, in practice, these approaches seem to handle limited normal variations. More importantly, they do not target the creation of fully procedural representations. They treat the materials as SVBRDF images, limiting the editing of results to pixelwise (rather than parametric) modifications of the segmentation and/or uniform (rather than multi-scale) material editing. Beyond the editing limits, these methods do not provide for the generation of textures of arbitrary spatial extent.

We choose to build upon KNN matting [Chen et al. 2013] by proposing a spectrum-sensitive affinity-based image matting method. KNN matting supports multiple layers and enables user control to conveniently define layers of interest, which are useful features for our material segmentation purpose. The aforementioned material segmentation approaches [An and Pellacini 2008; Lepage and Lawrence 2011; Pellacini and Lawrence 2007] could produce interesting segmentation results, but we chose KNN matting because of its generality, the availability of its implementation and its efficient runtime.

# 3 OVERVIEW

Given a set of SVBRDF maps, we want to generate a procedural representation of the material. Our inverse procedural modeling pipeline starts by hierarchically decomposing input SVBRDF maps into multiple sub-materials (Sec. 4) organized as a tree structure. Each sub-material represents a statistically similar region and its local variation, while the associated segmentation masks encode the global spatial variations of these sub-materials. The tree structure provides a layered relationship between different sub-materials. We traverse the material tree to convert each component to their procedural counterparts. We represent procedural sub-materials with a multi-layer procedural noise model combining procedural noise maps (Sec. 5.1) and mask maps using Point Process Texture Basis Functions (PPTBF) [Guehl et al. 2020] and random sampling (Sec. 5.2). After proceduralization, we compose these procedural representations and use a differentiable rendering-based optimization to match the appearance of the input material (Sec. 5.3). We show an overview of our method in Fig. 3.

In this paper we demonstrate our method on common physicallybased material parameters, that can be easily acquired using recent methods [Deschaintre et al. 2020]: albedo maps, normal maps and roughness maps. As normal maps encode vectors, their channels represent 3D directions which are difficult to proceduralize. Instead of directly working in normal map space, we convert to a height map using Poisson reconstruction [Pérez et al. 2003] a produce a better proceduralization. Fig. 4 provides an example of this reconstruction. Our approach can easily handle additional gray-scale or color maps.

#### 4 SVBRDF DECOMPOSITION

The first step of our pipeline is a hierarchical decomposition of the input SVBRDF maps into multiple sub-materials in a semi-automatic way. Although it is possible to spatially decompose SVBRDF maps into multiple sub-parts in one step, we propose an iterative decomposition into a material hierarchy, allowing the encoding of a layered relationship between sub-materials for proceduralization.

Specifically, given a set of SVBRDF maps, we decompose it into multiple sub-materials using a spectrum-aware matting algorithm (Sec. 4.1), which generates mask maps for each segmented sub-material. For each masked sub-material, we let the user decide whether to further decompose it using either our matting algorithm, or a lightweight instance-based decomposition algorithm (Sec. 4.2). Using this process, users can define the elements they consider important in the texture and iterate with new sub-divisions until no salient sub-material is left, creating a tree structure of decomposed submaterials.

### 4.1 Spectrum-aware SVBRDF Decomposition

We first consider how to decompose a set of SVBRDF maps, into sub-materials. As classical segmentation [Khan 2014] does not allow for smooth boundaries, we use alpha matting [Chen et al. 2013]. Inspired by KNN matting, our algorithm allows users to draw a few strokes to conveniently indicate different regions of interest.

As an affinity-based algorithm, KNN matting [2013] relies on a feature vector X(i) for each pixel *i* of the image to compute an affinity matrix. For SVBRDF maps, traditionally used features are albedo color (r, g, b), height (h), roughness ( $\alpha$ ) and position (x,y):

$$X(i) = (r(i), g(i), b(i), h(i), \alpha(i), x(i), y(i))$$
(1)

We propose an additional feature, based on the noise spectrum, enforcing the statistical uniformity expected from decomposed submaterials. To distinguish between the differences in local noise, we take the noise Fourier spectrum into account. Notice that although position features (x, y) are taken into account, we do *not* enforce spatial continuity of the regions because we sample multiple nonlocal neighborhoods with different weights of (x, y), to explore nonlocality (as done in the original KNN matting)

We estimate the local spectrum at each pixel *i* using Welch's method [Welch 1967] and reduce its dimensionality to 3 using Principal Component Analysis (PCA). We therefore compute the affinity matrix and matting using a feature vector X(i) composed of albedo color, height, roughness, position and our spectrum estimation.

Alpha matting results in multiple alpha maps which we process to generate binary mask maps. Each pixel in the image is assigned to the binary for which it has the highest alpha map value. The thresholded alpha-matting better represent the transition between sub-materials than direct segmentation methods. Fig. 5 shows a comparison between the decomposition results with and without spectrum features on example textures. In these two images, color and position features do not provide enough hints to separate two layers because they are similar in color but differ in noise spectra. More decomposition results on material maps with user scribbles can be found in our supplemental document.

Once a material decomposed, we provide the option to further process each sub-material using the same matting algorithm with the newly generated mask map(s) as an additional constraint.

#### 4.2 Instance-based Decomposition

As an alternative to a progressive decomposition into sub-materials, we provide a lightweight solution for instance-based decomposition. This is particularly useful for repeating materials such as a tile wall composed of different types of tiles (the input shown in Fig. 6).

# An Inverse Procedural Modeling Pipeline for SVBRDF Maps • 18:5



Fig. 3. Our procedural modeling pipeline allows a user to guide the segmentation of an input material to automatically generate a matching fully procedural material model. On the left we show the procedural tree generated by our method and on the right we show the proceduralization steps taken for each sub material. We represent the spatial variation using segmentation masks and Point Process Texture Basis Functions to create a procedural version through a query-and-optimization method. We represent the sub-material parameters with a multi-scale noise matching operation after reparameterizing normals to height values. Finally we use differentiable rendering to compose all sub-materials into a complete procedural material, rendered here lit from the top. Please see our supplemental material for a graph visualisation of a few results.



Normal map

Height map

Fig. 4. Reconstruction of a height map from a normal map. Values on the height map are normalized between 0 and 1. The Normal map\* is computed from reconstructed height map by a Sobel operator. While the recomputed Normal map\* shows slightly less fine-grain detail, the general structures are well preserved.

Rather than manually segmenting each different tile, we frame this as an instance detection and clustering problem, using - in this example - the mask map of segmented tiles to extract each instances of tile. Different instances correspond to disconnected regions in the mask map segmented from the previous layer in the decomposed material tree. We then scale each instance to the same size based on its bounding box. For each instance, we estimate its color histogram and local spectrum as features and build a feature matrix for agglomerative clustering as shown in Fig. 6. The clustering result allows us to extract the different types of tiles, their frequency and their sub-material to assign them with a similar frequency in the procedural model.

#### MATERIAL PROCEDURALIZATION AND 5 RECOMPOSITION

After decomposing the input SVRBDF maps into a tree of submaterials and mask maps, we traverse the tree and convert each component into a procedural version layer by layer. Finally, we introduce optimizable parameters during the final recomposition phase to best match the input SVBRDF.



Fig. 5. Image decomposition with and without spectrum features. The first image is a synthetic image while the second image is real-world texture. User scribbles are visualized as an additional layer superimposed over the input image, where blue and green scribbles indicate different matting layers. Matting results with spectrum features (bottom row) are better than the results without spectrum features (upper row) because color features here cannot provide sufficient information to separate two layers. Mask maps generated using spectrum features are also more integrated and less fragmented.

#### Multi-layer Procedural Noise Model 5.1

In this section we describe the conversion of the segmented submaterial into procedural models. Each leaf node in our tree is a



Fig. 6. Instance-based decomposition for sub-materials. Instead of manually segmenting each sub-material, we extract instances in the input texture and build a feature matrix for each of them. Agglomerative clustering is then performed on these feature matrices, yielding multiple clusters with similar sub-materials. Their spatial distribution is visualized as a label map, where different colors indicate different clusters. Finally, we fit a different procedural model for each cluster.

sub-material while the intermediate nodes store the mask maps. Each sub-material is represented by a set of masked SVBRDF maps. We use a multi-layer procedural noise model to proceduralize their texture appearance. As our images are masked and incomplete, the spectrum of the entire image is unavailable, preventing the direct fitting of procedural noise models using noise synthesis methods e.g. [Galerne et al. 2011; Galerne et al. 2012; Gilet et al. 2014; Heitz and Neyret 2018]. While Guingo et al. [2017] propose estimating local spectra in the valid regions only, their approach cannot capture the global variation of the noise textures, e.g. when the scale of the randomness is comparable to the scale of the full image. We leverage a similar sliding window approach but propose a multi-layer procedural model -shown in Algorithm 1- to deal with incomplete images. Given an input masked image, we decompose it into several noise layers using a progressive filtering strategy. For each layer, we filter our input image with a Gaussian kernel. The kernel size of the filter becomes larger as the number of layers grows, aiming at capturing larger scale spatial variation. For each filtered layer, we estimate its local spectrum and convert it to a procedural noise model similar to [Guingo et al. 2017]. Finally, our algorithm uses the mean value of filtered image I as the base color C, and extracts the final noise layer  $N \leftarrow I - C$ .

This last layer, however, represents the lowest frequency of the input image, preventing the use of sliding window spectrum estimation. As each sub-material is masked, the full image spectrum is unavailable. We therefore inpaint the missing data [Barnes et al. 2009; Telea 2004]. To reduce artifacts introduced by this step, we apply a set of Gabor noises as a basis to approximate the power spectrum of the inpainted image [Galerne et al. 2012]. This last step and our use of multi-layer noise allows our method to produce a fully

ACM Trans. Graph., Vol. 41, No. 2, Article 18. Publication date: January 2022.

ALGORITHM 1: Multi-layer procedural noise model
<b>input</b> : image <i>I</i> , maximum number of layers $n_{max}$ , image variance
threshold $\epsilon$ , a set of filter kernel sizes $K$ , a set of local
window sizes $T$ , a set of step sizes $S$
<b>output</b> : a set of procedural models $p$ , base color $C$
$p \leftarrow \emptyset;$
$i \leftarrow 0;$
while $i < n_{max}$ do
$\sigma \leftarrow \text{Variance}(I);$
if $\sigma \leq \epsilon$ then
break;
end
$I' \leftarrow \text{Gaussian filter}(I, K_i);$
$N_i \leftarrow I - I';$
$p \leftarrow p \cup \text{Random Phase Noise matching}(N_i, T_i, S_i);$
$I \leftarrow I';$
$i \leftarrow i + 1;$
end
$C \leftarrow \operatorname{mean}(I);$
$N \leftarrow I - C;$
$p \leftarrow p \cup \text{Random Phase Noise matching}(N, T_{nmax}, S_{nmax});$



Fig. 7. Our Multi-layer procedural noise model on a colored texture image. Black regions shows masked unavailable regions. The color of noise images are more contrasted for visualization. The synthetic noise (bottom row) is procedural noise estimated from filtered noise (middle row), and the synthetic texture (middle image in the top row) is computed by adding all the synthetic procedural noise together with the base color.

procedural representation of the sub-material's properties. Fig. 7 shows an example of our multi-layer decomposition on a masked colored texture image, where each layer captures different levels of details from an input image and the last layer shows the global variation of the noise image.



Fig. 8. Comparison between multi-layer and single-layer noise on colored texture images. Compared to a single-layer approach, our multi-layer method can reconstruct a smoother texture image and capture more finegrained details, while the single-layer generates some apparent repetitive and "dirty" patterns. Textures are segmented by our spectrum-aware matting method, fitted and recomposed to generate the results.

Using the procedural noise  $p_i$  fitted to each layer i and the base color C, we reconstruct the original texture as  $\sum p_i + C$ , i = 1, 2, ..., n where n is the number of noise layers. We further improve the results of this reconstruction for SVBRDF maps by introducing a differentiable rendering-based optimization scheme, described in Sec. 5.3. Our approach provides a fully procedural representation and can reproduce fine-grained details and yield a smoother noise texture than single layer methods (see Fig. 8). Furthermore, our multi-layer approach gives users more control over level of detail.

5.1.1 *Images with Multiple Channels.* Synthesizing colored images, such as the albedo maps could result in poor color mixing during synthesis. Similar to previous approaches [Galerne et al. 2011; Guingo et al. 2017; Heeger and Bergen 1995], given a multi channel noise texture, we synthesize it in a PCA color space. We project the image from the original RGB space to the PCA space, allowing us to synthesize each channel independently before projecting it back to the original RGB space. We then match the histogram of the synthesized image to the input image as a post-processing step to ensure a matching color distribution.

## 5.2 Procedural Mask Synthesis

While we represent the sub-material local variations using our multilayer noise model, we use mask maps to represent the global spatial distributions of the sub-materials. Proceduralizing these mask maps allows us to reach a full procedural representation of the material.



Fig. 9. Examples of our procedural mask synthesis method. Given an input binary mask map (left) segmented by our spectrum-aware matting method, we first query a similar mask map (middle) from a database, and use it as an initialization of our alternating optimization algorithm to correct and match its structure to the input (right).

With it, we can easily edit, resample and extend the resolution of the spatial distribution. Similar to sub-material proceduralization, mask proceduralization is done recursively during traversal of the material tree. We model mask maps by two methods 1) Point Process Texture Basis Functions (PPTBF) [Guehl et al. 2020] to model decomposed binary mask maps by matting algorithm (Sec. 4.1); 2) Random Sampling to model decomposed instances (Sec. 4.2).

5.2.1 Inverse Mask Fitting by PPTBF. As introduced in [Guehl et al. 2020], Point Process Texture Basis Functions (PPTBF) are defined by the sparse convolution of randomly-sampled 2D points  $x_i$  with a kernel function being the product between a visual feature f and a blending window w:

$$PPTBF_{k}(\mathbf{x}) = \sum_{\mathbf{x}_{i} \in \mathcal{N}_{k}(\mathbf{x})} f(\mathbf{x} - \mathbf{x}_{i})w(\mathbf{x} - \mathbf{x}_{i}),$$
(2)

where  $N_k(\mathbf{x}) = {\mathbf{x}_1, ..., \mathbf{x}_k}$  describes the *k* closest sample points around **x**. According to Eq. (2), the behavior of PPTBF depends on 2D spatial point distribution of  $\mathbf{x}_i$ , visual feature function *f*, and blending window function *w*, which are in turn controlled by intuitive parameters such as kernel size, degree of smoothing etc. to produce a continuous scalar field. After applying a threshold, we use PPTBF to model binary mask maps segmented by our matting algorithm. In this section, We first describe the fitting of mask maps on the top level of the material tree where the mask covers the entire image, and then discuss fitting maps in the lower hierarchy where mask contains missing regions.

As PPTBF is a forward generation process, [Guehl et al. 2020] did not focus on by-example modeling. The authors suggested an approach that queries a precomputed PPTBF database and then optimizes parameters, but did not provide details. We therefore experimented with different feature representations and optimization routines and here discuss the techniques we chose. Detailed implementation of our *query-and-optimization* method can be found in our released code.

Query: As PPTBF are not fully differentiable and partly rely on discrete parameters, starting from a good initialization is particularly important for efficient optimization. We first query a database, uniformly covering the variations allowed by PPTBF, to retrieve the nearest neighbor parameters of our input mask. We use the mask image database provided in [Guehl et al. 2020]. The database contains 450K images and was generated by non-regular sampling with three different thresholds for binarization. To measure the similarity between the input mask and pre-sampled mask maps in the database, we choose Local Binary Pattern (LBP) to encode local statistics, Gram Matrix of pre-trained VGG19 [Simonyan and Zisserman 2015] deep features for global statistics, and Fourier power spectrum for regularity. We weight and concatenate these three features to build a high dimensional vector acting as a descriptor for binary mask maps. We precompute such a vector for each mask map in the database, and reduce their dimensionality by PCA to 512 dimensions for a more robust representation. The storage of the processed database is 7.81 GB in total. We build an acceleration structure for fast nearest neighbor search based on  $L_2$  distance of two feature vectors. The precomputation takes around 36 hours, but each query then requires less than 1 second.

**Optimization:** After we retrieve a pre-sampled binary mask map from the database, we use its parameters to initialize our optimization algorithm to better match the input mask. Since PPTBF contains both continuous parameters as well as discrete parameters, we apply coordinate descent to optimize continuous and discrete ones alternatively. As evaluation of PPTBF is expensive and not easily differentiable, we apply gradient-free approaches to avoid costly finite differentiation. Our implementation for PPTBF optimization is CPUbased. For continuous parameters, we apply the Powell method (SciPy [2020]) and for discrete parameters, we adopt a Bayesian optimization method with Gaussian Process (GPyOpt [2016]).

As Fig. 9 shows, our query results provide an initial guess about the general structure but it can be hard to find a perfect match from the database due to sparse sampling. Our optimization step can correct the parameters and match the structure of the input image.

5.2.2 Incomplete mask maps. Different from masks at the top layer of our tree, masks for sub-materials in the lower hierarchy are masked by higher level masks. If naively processed, this leads to poor procedural reconstruction of the sub-material distribution. In order to fit incomplete mask maps, we propose two solutions. First, we adjust our optimization by computing losses only for unmasked pixels between input mask and PPTBF output. Second, we propose proceduralizing an inpainted version – using PatchMatch [Barnes et al. 2009] – of incomplete mask maps. However, features in the mask map are not sufficient to guide inpainting. Instead, we inpaint the sub-material and compute the mask map on it as shown in



Fig. 10. Mask inpainting. We apply a hierarchical segmentation approach. We first segment the chipped area, leaving the yellow paint and spaces between the planks (Masked Input). If we directly segmented the masked input, the result (Incomplete Mask) is difficult to fit procedurally. We need to separate the influence of the parent mask (chipped area) from the sub parts to separate (planks from space in between). To do so we use in-painting. Directly inpainting the incomplete mask fails as binary masks cannot provide sufficient hints to guide the inpainting (Directly Inpainted Mask). Instead, we inpaint the input image using PatchMatch (Inpainted Input) on which we compute the segmentation, producing a better mask (Segmented Mask)

Fig. 10. Because PatchMatch inpainting is a sample-based method, this second approach works best for stochastic distributions.

5.2.3 Random Sampling. When mask maps are generated using our instance-based decomposition solution, we can model their distributions by random sampling. Suppose we have n mask maps, each representing one type of segmented instance. We count the number of instances in each mask map, and estimate their probability of occurrence. During mask synthesis, we randomly – following the estimated distribution – assign a label between 1 and n to instances in the procedural version of the mask map. The corresponding submaterials are then synthesized in the labeled regions.

#### 5.3 Recomposition

Finally, using our generated procedural noise maps and masks tree, we compose them into an output SVBRDF to reproduce the appearance of the original SVBRDF inputs. To relate this approach to a classical Substance Designer [2021] pipeline, our procedural noise maps and binary masks function as *generators nodes*, which do not rely on existing artist designed graphs.

To better match the original SVBRDF inputs, we add optimizable operators to control the appearance of procedural noise and binary maps. Given a noise map *I*, we modify its appearance by  $G(I*\alpha+\delta,\sigma)$ , where *G* is a Gaussian filter,  $\alpha$  controls the intensity,  $\delta$ biases the noise, and  $\sigma$  is the standard deviation of the Gaussian filter. For a binary mask map *M* which represents the distribution of submaterials, we model the smooth transition between the boundaries



Fig. 11. Validation of our optimization-based recomposition. In this case, the "input" material maps are generated using [Deschaintre et al. 2020] from a natural image data where lighting is not completely removed from material maps. Additionally, local variations of the different material maps cannot be well represented with simple Gaussian noise especially in the normals. Directly fitting these local textures using our multi-layer noise model leads to large fitting errors and visual artifacts (3rd row). Using our differentiable optimization approach, our method automatically refines the procedural parameters of our output during the recomposition step (2nd row). Thanks to this optimization, we improve the stability of our method for arbitrary inputs. For better visualization we directly use the KNN segmented mask rather than a procedural counterpart.

of sub-materials by  $G(M, \sigma)$ . The parameters  $\alpha$ ,  $\delta$  and  $\sigma$  are optimizable for each Gaussian filter per noise or mask map. In each leaf node of our tree, the noise maps are multiplied by their corresponding mask maps and linearly combined to reconstruct the SVBRDF maps for a procedural sub-material. Procedural Sub-materials from different layers are recursively computed and aggregated in a bottom-up fashion to build the final output SVBRDF maps.

These optimizable operators, together with our reconstructed noise maps and binary maps, build a small optimizable material graph. We optimize this material graph, using a differentiable renderingbased optimization routine. The reconstructed SVBRDF maps and the input SVBRDF maps are rendered using a Cook-Torrance, GGX shading model under randomly sampled lighting configurations. Considering that the structure of the procedural SVBRDF maps might not align perfectly with the original inputs, we define the loss as a combination of Style loss and SSIM loss. Style loss is defined as the  $L_1$  difference between Gram matrices computed over VGG [2015] features of the renderings, similar to the style transfer literature [Gatys et al. 2016]. SSIM loss is computed by the  $L_1$  difference between the structural similarity (SSIM) indices [Zhou Wang et al. 2004] of the renderings. The full loss is written as

$$L = \sum \|GM(I_i) - GM(I_i^*)\|_1 + \beta \|SSIM(I_i) - SSIM(I_i^*)\|_1$$
(3)

where GM is the Gram matrix and SSIM is an operator to compute structure similarity index;  $I_i$  and  $I_i^*$  are input/procedural albedo map, normal map (computed from height map), roughness map and

their renderings;  $\beta$  balances the weight between the style term and the SSIM term. In Fig. 11 we show an example with and without this optimization process. Often in real data local texture appearance cannot be well represented by Gaussian noise models, resulting in fitting errors and artifacts when generating procedural models (3rd row in Fig. 11). This is particularly problematic for normal/height map modeling. Our optimization helps refine the parameters of our procedural graph, significantly improving the results, even when the exemplar violates our assumption of locally uniform appearance such as in Fig. 11.

# 5.4 Final procedural representation

Our final procedural representation of material maps is the equivalent of small material graphs with (i) procedural noise maps (ii) procedural mask maps (iii) optimizable operators, similar to Fig. 2. In [Shi et al. 2020]'s terminology, each procedural noise and mask model is a generator while optimizable operators serve as filters. As a fully procedural model, editable parameters of our small material graph are (i) parameters of our generators e.g. noise models and mask models; (ii) parameters in the filter nodes. Editing can be done on individual maps or jointly on all of them. The procedural model can then generate the material at arbitrary resolution and scale.

# 6 RESULTS

We demonstrate our pipeline with a variety of materials as inputs. Each material is defined by a set of SVBRDF maps including albedo map, normal map and roughness map – these are the maps our input materials use, but our method can adapt to any additional or different gray-scale or color maps.

Our pipeline is implemented in Python and partly relies on Matlab. Differentiable optimization of our material graphs (Sec. 5.3) is implemented in PyTorch. We adopt an L-BFGS-B optimizer to optimize our material graph with a learning rate of 0.005. It takes around 200 steps – 2 minutes on a Nvidia RTX 2070 Super GPU with a CPU of Intel Core i7-9700 – to converge. A complete inverse material modeling process takes less than 5 minutes for user interaction and about 20 minutes for computing, depending on the complexity of the input material. The typical computation time is: 2 minutes for spectrum-aware matting; 13 minutes for procedural mask query (1s) and optimization (depending on initialization quality); 3 minutes for multi-layer noise modeling and 2 minutes for reference but highlight that our current implementation is not optimized.

Fig. 12 shows our inverse procedural modeling results on different materials. For each example, inputs are albedo map, normal map and roughness map only. We decompose these material maps using our hierarchical segmentation methods and visualize the computed labeled mask maps. Given material maps and computed hierarchical mask maps, we generate procedural materials and output albedo, normal and roughness maps (Sec. 5). The results show that our pipeline can reproduce well a variety of stochastic and regular materials. We see that our method recovers both large scale patterns and fine-grained details thanks to our sub-material decomposition approach. The general structure and texture appearance are not perfectly registered with the original input because our model is a

# 18:10 • Yiwei Hu, Chengan He, Valentin Deschaintre, Julie Dorsey, and Holly Rushmeier



Fig. 12. Results of our method for different materials. We show that our pipeline can proceduralize a variety of materials, and reproduce their global structures as well as local texture appearance. All of our results are entirely procedural. TL: Top lighting; SL: Side lighting. Please see our supplemental material for more results.



Fig. 13. Picture(s) as input. We apply our method on decomposed SVBRDF maps captured with a single or multiple image(s) with flash lighting [Deschaintre et al. 2019, 2020; Guo et al. 2020b], showing the diversity of material source we can handle. Our procedural results enable further creation and provide more regularity and remove the baked in lighting and irregularities due to the few image(s) capture method. The bottom two examples were segmented using the instance segmentation approach showing that it can also handle irregular patterns. CTL: Central Top light; TSL: Top Side light. Please see our supplemental material for more results. ACM Trans. Graph., Vol. 41, No. 2, Article 18. Publication date: January 2022.



Fig. 14. High-resolution material synthesis. We show the input material in the bottom right corner and render a high resolution  $(2048 \times 1024)$  version generated with our method's results with lighting from the top. By applying our pipeline, we convert these pixel-based material maps to procedural representation thus enabling arbitrary scale and resolution expansion without repetition.

procedural approximation of it. We aim at reproducing its appearance rather than a pixel-perfect match [Deschaintre et al. 2018], allowing us to preserve the material global appearance while sampling a new realization of it or editing it, as shown in Fig. 16.

#### 6.1 Natural Images as Input

Although our pipeline was designed to take material maps as input, it can also work on natural and flash images, benefiting from existing material acquisition methods [Deschaintre et al. 2018, 2019, 2020; Gao et al. 2019; Guo et al. 2020b]. Providing a set of captured images, we first apply a material acquisition method to generate material maps, and use the generated SVBRDF maps with our method. Fig. 13 shows examples of picture(s) as input. We convert the input natural image(s) to SVBRDF maps using state-of-the-art material acquisition methods [Deschaintre et al. 2019, 2020; Guo et al. 2020b].As opposed to "artist-designed" SVBRDF maps or manually post-processed SVBRDF maps, automatically generated SVBRDF maps are not perfectly clean. Different maps can be noisy and exhibit shading and color variations due to incomplete lighting removal. Normal maps in particular can represent strong height variations even within a single sub-material. This makes it challenging to recover an exact match to the input image. Furthermore, irregularities often occur in acquired SVBRDF which should be regular. We show in Fig. 13 that our pipeline is capable of matching acquired SVBRDF overall appearance, and also of enforcing better regularity.

# 6.2 Application

*6.2.1 Procedural Material Editing.* Once proceduralized using our method, materials can be further edited. Fig. 16 shows examples of operations that we can apply on our results. We can freely edit the way mask maps and noise models are combined, the large scale structures (e.g patterns and distributions) and fine-scale details (e.g. fine normals and roughness). Once generated, our procedural model allows users to get interactive feedback on each of their edits. As opposed to [Hu et al. 2019; Shi et al. 2020], we do not rely on pre-defined material graph, allowing artists to use our method to generate a small tune-able and extendable material graph to start a new design.



Fig. 15. Material super-resolution. We show our model can generate higherresolution textures but without changing its scale. This is achieved by only procedurally upscaling mask maps which controls the global spatial distributions of sub-materials.

Furthermore, Hu et al. [2019] relied on a style transfer like postprocess step to better match the input texture appearance, limiting the editability of their final results.

6.2.2 High-resolution Material Generation. As a fully procedural model, materials converted by our pipeline can be expanded to arbitrary size. Global structures can be reproduced and extended using the procedural mask map with PPTBF and random sampling, while the appearance of sub-materials can be losslessly synthesized to higher resolution using procedural noise models. Fig. 14 shows examples of high-resolution material synthesis. The resolution of the result material and rendered images is  $2048 \times 1024$ . We also show an example of material super-resolution in Fig. 15 where the resolution of the input SVBRDF is  $512 \times 512$  and we double its size in each dimension by procedurally upsampling its mask maps. Such operation will not affect its global scale, providing a super-resolved material rather than synthesize a larger scale one as demonstrated in Fig. 14. More high-resolution materials generated by our method are shown in the supplemental material.



Fig. 16. Procedural material editing. Our procedural material model provides editability to users. In the top example, we edit the optimizable filter parameters to change the (1) roughness (2) normal and (3) color of the material sequentially; for the second example, we edit parameters in the (1) mask generators (first two edits) to change spatial arrangement and distribution of sub-materials; and (2) noise generators (last edit) to change the fine-scale normals on each tile; we edit all parameters for the third example where (1) high-frequency shininess is removed; (2) global structures are changed (last two edits); and (3) normals are enhanced, leading to a completely different leathered brick material. Please zoom-in to see fine-scale controls.

#### 6.3 Comparison with Prior Work

We compare our results to previous work in inverse procedural modeling and in texture synthesis. Since we take the output of a material acquisition method as input to our pipeline, we do not compare to material acquisition methods.

6.3.1 Inverse Procedural Material Modeling. We compare our approach with state-of-the-art inverse procedural material modeling frameworks [Hu et al. 2019; Shi et al. 2020] in Figs. 18 & 19. Both are built on a collection of pre-defined material graphs and rely on model selection with parameter estimation. While Hu et al. directly applies neural network to predict parameters, Shi et al. (MATch) optimize parameters by end-to-end differential rendering. We use the model selection scheme of Hu et al. to chose input substance models for both methods. Despite not using pre-existing material graphs, we show that our pipeline well reproduces material maps with regular and stochastic patterns e.g. brick and stucco, similar to the examples shown in [Hu et al. 2019]. Hu et al. additionally propose a post-processing step (see Hu et al. (b) in Fig. 18) to enrich the details, but this step is done on the rendered images (instead of material maps) and loses its procedural aspect.

Direct comparison to (MATch) [2020] is difficult because the optimization framework of MATch strongly depends on the quality of initialization. Indeed, the MATch framework cannot optimize discrete parameters. This limitation requires the scale of the structured target to be the same as the selected procedural material. The user is then required to manually fine-tune discrete parameters in Substance Designer. Additionally, only a subset of the Substance Engine was made differentiable in MATch, limiting the pool of compatible procedural graphs. In Fig. 19 and the supplemental material we attempt to provide good initialization and discrete parameters for MATch.

6.3.2 Texture Synthesis. We also compare our approach with texture synthesis methods as they share the ability of resolution extension our method allows. We experiment with several state-ofthe-art example-based texture synthesis methods and generalize them to accept material maps as input. We stack all material maps together to build a high-dimensional texture map where each texel encodes albedo values, normal directions and roughness values. Loss functions are computed on each material map and averaged. For self-tuning texture optimization [Kaspar et al. 2015], its generalization to multi-channel material maps is not trivial and we therefore run their algorithm separately on each material map. Fig. 17 shows these comparisons in which the input resolution is  $300 \times 300$  and the output resolution is  $512 \times 512$ . We see that our method is capable of extending both structured and stochastic materials and does not suffer from artifacts seen in traditional texture synthesis approaches such as structure error. Finally, and most significantly, compared with these other texture synthesis approaches, our method enables editing and fast synthesis of larger resolution without necessarily augmenting the scale.

# 7 LIMITATIONS

# 7.1 Texture Taxonomy for Proceduralization

To characterize the limitations of our work, we describe the space of possible materials we attempt to model. We consider materials as collections of elements. The space of different element collections has four dimensions defining the complexity of the proceduralization of a material, illustrated in Fig. 20. The first axis is the number of different material element types in the texture. Both the types of elements and their spatial distribution relative to one another need to be modeled. The second axis is the nature of the spatial distribution, whether the elements on the texture are sparsely or densely positioned, with overlap for example. Dense spatial distribution are harder to segment using masks. The third axis is complexity of the contour of the spatial mask. The more complex, the more difficult it is for PPTBF or any procedural approach to represent them faithfully. The fourth axis is the texture within each element, whether it can easily be represented as noise or is more structured and semantically meaningful.

Our general pipeline that segments SVBRDF and then estimates parameters for each element type, applies to the complete space we have described. However, our current implementation is limited to the less complex end of each of the last three axes. Our segmentation method is limited in that it cannot appropriately segment dense overlapping material. As stated, PPTBF performs well with irregular simple contours but not with complex contours. Finally, for the fourth axis we have not developed a method to proceduralize semantically meaningful elements such as the bunny.

# 7.2 Specific Limitation Examples

Our results show that our pipeline is able to proceduralize different types of materials with multiple SVBRDF maps e.g. albedo, normal and roughness. In Fig. 21 we show failure cases that result from the general limitations we just described. The combination of complex



Fig. 17. Comparison of our method with example-based texture synthesis methods on SVBRDF maps. We generalize these methods to process multi-channel SVBRDF maps. We show our method; InGAN [Shocher et al. 2019]; Non-stationary Texture Synthesis by Adversarial Expansion [Zhou et al. 2018]; Self-tuning Texture Optimization [Kaspar et al. 2015]; Image Quilting [Efros and Freeman 2001]. Images shown here are rendered by GGX shading model.

contours and a semantically meaningful structure results in the failure shown in the top row of Fig. 21. The complex contours form a "flower" shape that is not preserved by the PPTBF mask. The presence of a sub-material that is itself composed of a densely packed set of elements results in the failure shown in the bottom row of Fig. 21. Neither a noise pattern nor our segmentation method can capture the arrangement of the small pebbles forming the fill material between the square elements. Finer segmentation could produce better results, but the overlapping partial shapes would still be difficult to match using PPTBF. In some cases, such as the second example of Fig. 17, the retrieved procedural mask will have small differences in the contours compared to the original. Adding a loss more sensitive to contours could improve such cases.

Finally, as our procedural material model is built upon multi-scale Gaussian noise models and Gaussian filters, it poorly reproduces highly-structured variations e.g. extremely-strong and directional normal variations seen as Fig. 22. In these cases, the segmentation of the height map is also ambiguous, leading to less faithful procedural reconstruction.

#### 8 FUTURE WORK

In this work, we propose the first complete pipeline for inverse procedural modelling of general materials and highlight the challenges of each steps. We describe here interesting future work to enable better proceduralization.

**Modelling.** New generative approaches such as deep texture or material generation [Henzler et al. 2021, 2020] could help better

reproduce complex material appearances that cannot be simply represented by noise models, while sacrificing some control.

**Segmentation.** Our pipeline relies on user input, allowing user control and specification of the artistically important elements in the material. Nevertheless, a fully automatic segmentation would provide a faster approach and allow for quickly proceduralization of a large amount of materials. Current methods [Bell et al. 2015; Chen et al. 2018; Cho et al. 2016; Xu et al. 2017] are not geared toward material segmentation and fail on the materials we consider.

**Recomposition.** Complex details and patterns might be lost during procedural modeling as our method relies on noise and mask fitting. Introducing advanced differentiable filters and generator – such as ones in Substance Designer and MATch [Shi et al. 2020] – to modify the appearance of procedural noise maps and binary maps, would allow to represent a wider range of appearances, further narrowing the gap between input and generated procedural materials.

# 9 CONCLUSION

We present the first pipeline for semi-automatic material proceduralization. Given a set of input material maps, our pipeline decomposes them into a tree of sub-materials and corresponding binary mask maps. We model the local appearance of sub-materials by procedural noise models and proceduralize binary mask map to reproduce global distribution of sub-materials.

Compared with previous work [Hu et al. 2019; Shi et al. 2020], our pipeline does not rely on any predefined material graphs. With our approach, combined with the state-of-the-art material acquisition



Fig. 18. Comparison of our method to [Hu et al. 2019]. The second column shows predicted procedural results by [Hu et al. 2019] while the third column is their style augmented results (non-procedural and no edibility). In contrast to their method, our pipeline generates fully procedural materials without a pre-existing material graph as an auxiliary input. The images are rendered using Blender with diffuse reflectance to match [Hu et al. 2019].



Fig. 19. Comparison of our method to MATch [2020]. We use Hu et al. [2019]'s framework to select a close Substance model (Default), and use MATch to optimize its appearance. As the MATch framework does not handle discrete parameters and requires good initialization, it can generate poor output if the initialization and the discrete parameters are not hand-tuned. This can give poor results for bricks as shown in the supplemental material. Materials are rendered using the GGX shading model.

methods, we enable convenient creation of high quality spatially varying procedural materials. We take a first step toward general material proceduralization and hope our work and highlighted challenges will inspire future research.



Fig. 20. Illustration of the four dimensions of variation of element collections in the space of material textures that we model.



Input Render("GT") Segmented Our Procedural Our Render

Fig. 21. Failure cases on semantically-shaped pattern and denselyoverlapped material. The input images are a captured image (first row) and an arbitrary texture (second row). The "Ground Truth" render is generated using [Deschaintre et al. 2020] results. The first row shows a case where our PPTBF-based procedural mask map fails to reproduce the structure of the segmented mask map, missing the semantic "flower" shape. In the second row, one of the sub-material contains dense, small, overlapping pebbles which cannot be well reproduced by our multi-layer noise model or easily segmented. As a result, individual pebbles cannot be distinguished in our rendered result. Please zoom in to see details.

#### ACKNOWLEDGMENT

This work was supported in part by NSF Grant No. IIS-2007283, and by Dr. Abhijeet Ghosh with his EPSRC Early Career Fellowship (EP/N006259/1)

#### REFERENCES

Adobe. 2021. Substance Designer. https://www.substance3d.com/.



Fig. 22. Failure cases on material exemplars with highly-structured variations. We show additional failure cases on extremely structured normal variations which cannot be either 1) identified by segmentation or 2) recovered by noise models and Gaussian kernels on the transition regions of adjacent sub-materials. Insets on normal map show segmentation map and our procedural map. For better visualization, the first material is rendered with a central light while the second material is rendered with a top side light.

- Adib Akl, Charles Yaacoub, Marc Donias, Jean-Pierre Da Costa, and Christian Germain. 2018. A survey of exemplar-based texture synthesis methods. *Computer Vision and Image Understanding* 172 (2018), 12–24.
- Yağiz Aksoy, Tae-Hyun Oh, Sylvain Paris, Marc Pollefeys, and Wojciech Matusik. 2018. Semantic soft segmentation. ACM Trans. Graph. 37, 4, Article 72 (2018), 13 pages.
- Yagiz Aksoy, Tunc Ozan Aydin, and Marc Pollefeys. 2017. Designing effective inter-pixel information flow for natural image matting. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 29–37.
- Xiaobo An and Fabio Pellacini. 2008. AppProp: All-Pairs Appearance-Space Edit Propagation. ACM Trans. Graph. 27, 3 (Aug. 2008), 1–9. https://doi.org/10.1145/ 1360612.1360639
- The GPyOpt authors. 2016. GPyOpt: A Bayesian Optimization framework in python. http://github.com/SheffieldML/GPyOpt.
- Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. 2009. Patch-Match: A Randomized Correspondence Algorithm for Structural Image Editing. ACM Trans. Graph. 28, 3, Article 24 (July 2009), 11 pages. https://doi.org/10.1145/ 1531326.1531330
- Sean Bell, Paul Upchurch, Noah Snavely, and Kavita Bala. 2015. Material recognition in the wild with the materials in context database. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 3479–3487.
- Urs Bergmann, Nikolay Jetchev, and Roland Vollgraf. 2017. Learning Texture Manifolds with the Periodic Spatial GAN. In Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70), Doina Precup and Yee Whye Teh (Eds.). PMLR, 469–477. http://proceedings.mlr.press/ v70/bergmann17a.html
- Mark Boss, Varun Jampani, Kihwan Kim, Hendrik P.A. Lensch, and Jan Kautz. 2020. Two-Shot Spatially-Varying BRDF and Shape Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 3982–3991.
- Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. 2018. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In Proceedings of the European Conference on Computer Vision (ECCV). 801–818.

- Qifeng Chen, Dingzeyu Li, and Chi-Keung Tang. 2013. KNN matting. IEEE Transactions on Pattern Analysis and Machine Intelligence 35, 9 (2013), 2175–2188.
- Donghyeon Cho, Yu-Wing Tai, and Inso Kweon. 2016. Natural image matting using deep convolutional neural networks. In Proceedings of European Conference on Computer Vision (ECCV). Springer, 626–643.
- Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, and Andrea Vedaldi. 2016. Deep filter banks for texture recognition, description, and segmentation. *International Journal of Computer Vision* 118, 1 (2016), 65–94.
- Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-Image SVBRDF Capture with a Rendering-Aware Deep Network. ACM Trans. Graph. 37, 4, Article 128 (Aug 2018), 15 pages.
- Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis, and Adrien Bousseau. 2019. Flexible SVBRDF Capture with a Multi-Image Deep Network. Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering) 38, 4 (July 2019), 1–13.
- Valentin Deschaintre, George Drettakis, and Adrien Bousseau. 2020. Guided Fine-Tuning for Large-Scale Material Transfer. Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering) 39, 4 (2020), 91–105. http://wwwsop.inria.fr/reves/Basilic/2020/DDB20
- Valentin Deschaintre, Yiming Lin, and Abhijeet Ghosh. 2021. Deep polarization imaging for 3D shape and SVBRDF acquisition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Yue Dong. 2019. Deep appearance modeling: A survey. Visual Informatics 3, 2 (2019), 59 - 68. https://doi.org/10.1016/j.visinf.2019.07.003
- Alexei A. Efros and William T. Freeman. 2001. Image Quilting for Texture Synthesis and Transfer. In Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01). Association for Computing Machinery, New York, NY, USA, 341–346. https://doi.org/10.1145/383259.383296
- Alexei A Efros and Thomas K Leung. 1999. Texture synthesis by non-parametric sampling. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Vol. 2. IEEE, 1033–1038.
- Sing Choong Foo. 1997. A Gonioreflectometer For Measuring The Bidirectional Reflectance Of Material For Use In Illumination Computation. Doctoral Dissertation, Cornell University.
- B. Galerne, Y. Gousseau, and J. Morel. 2011. Random Phase Textures: Theory and Synthesis. IEEE Transactions on Image Processing 20, 1 (2011), 257–267.
- Bruno Galerne, Ares Lagae, Sylvain Lefebvre, and George Drettakis. 2012. Gabor Noise by Example. ACM Trans. Graph. 31, 4, Article 73 (July 2012), 9 pages. https: //doi.org/10.1145/2185520.2185569
- B. Galerne, A. Leclaire, and L. Moisan. 2017. Texton Noise. Computer Graphics Forum 36, 8 (2017), 205–218. https://doi.org/10.1111/cgf.13073 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13073
- Duan Gao, Xiao Li, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. 2019. Deep Inverse Rendering for High-resolution SVBRDF Estimation from an Arbitrary Number of Images. ACM Trans. Graph. 38, 4, Article 134 (July 2019), 15 pages. https: //doi.org/10.1145/3306346.3323042
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2016. Image style transfer using convolutional neural networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2414–2423.
- Guillaume Gilet, Basile Sauvage, Kenneth Vanhoey, Jean-Michel Dischler, and Djamchid Ghazanfarpour. 2014. Local Random-Phase Noise for Procedural Texturing. ACM Trans. Graph. 33, 6, Article 195 (Nov. 2014), 11 pages. https://doi.org/10.1145/ 2661229.2661249
- Dar'ya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. 2016. BRDF Representation and Acquisition. Computer Graphics Forum 35, 2 (2016), 625–650.
- Pascal Guehl, Remi Allègre, Jean-Michel Dischler, Bedrich Benes, and Eric Galin. 2020. Semi-Procedural Textures Using Point Process Texture Basis Functions. Computer Graphics Forum 39, 4 (2020), 159–171. https://doi.org/10.1111/cgf.14061
- Geoffrey Guingo, Basile Sauvage, Jean-Michel Dischler, and Marie-Paule Cani. 2017. Bi-Layer Textures: a Model for Synthesis and Deformation of Composite Textures. Computer Graphics Forum (2017), 111–122. https://doi.org/10.1111/cgf.13229
- Yu Guo, Miloš Hašan, Lingqi Yan, and Shuang Zhao. 2020a. A Bayesian Inference Framework for Procedural Material Parameter Estimation. Computer Graphics Forum 39, 7 (2020), 255 – 266.
- Yu Guo, Cameron Smith, Miloš Hašan, Kalyan Sunkavalli, and Shuang Zhao. 2020b. MaterialGAN: Reflectance Capture Using a Generative SVBRDF Model. ACM Trans. Graph. 39, 6, Article 254 (Nov. 2020), 13 pages. https://doi.org/10.1145/3414685. 3417779
- David J Heeger and James R Bergen. 1995. Pyramid-based texture analysis/synthesis. In Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH). 229–238.
- Eric Heitz and Fabrice Neyret. 2018. High-Performance By-Example Noise Using a Histogram-Preserving Blending Operator. ACM Trans. Graph. 1, 2, Article 31 (Aug. 2018), 25 pages. https://doi.org/10.1145/3233304

- Philipp Henzler, Valentin Deschaintre, Niloy J Mitra, and Tobias Ritschel. 2021. Generative Modelling of BRDF Textures from Flash Images. ACM Trans Graph (Proc. SIGGRAPH Asia) 40, 6 (2021).
- Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. 2020. Learning a Neural 3D Texture Space from 2D Exemplars. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 8356 – 8364.
- Yiwei Hu, Julie Dorsey, and Holly Rushmeier. 2019. A Novel Framework for Inverse Procedural Texture Modeling. ACM Trans. Graph. 38, 6, Article 186 (Nov. 2019), 14 pages. https://doi.org/10.1145/3355089.3356516
- Alexandre Kaspar, Boris Neubert, Dani Lischinski, Mark Pauly, and Johannes Kopf. 2015. Self Tuning Texture Optimization. *Comput. Graph. Forum* 34, 2 (May 2015), 349–359. https://doi.org/10.1111/cgf.12565
- Muhammad Waseem Khan. 2014. A survey: image segmentation techniques. International Journal of Future Computer and Communication 3, 2 (2014), 89.
- Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. 2006. Inverse Shade Trees for Non-Parametric Material Representation and Editing. ACM Transactions on Graphics (Proc. SIGGRAPH) 25, 3 (July 2006).
- Daniel Lepage and Jason Lawrence. 2011. Material Matting. ACM Trans. Graph. 30, 6 (Dec. 2011), 1–10. https://doi.org/10.1145/2070781.2024178
- Anat Levin, Alex Rav-Acha, and Dani Lischinski. 2008. Spectral matting. IEEE Transactions on Pattern Analysis and Machine Intelligence 30, 10 (2008), 1699–1712.
- Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. 2017. Modeling Surface Appearance from a Single Photograph using Self-augmented Convolutional Neural Networks. ACM Trans. Graph. 36, 4, Article 45 (2017), 11 pages.
- Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. 2018a. Materials for Masses: SVBRDF Acquisition with a Single Mobile Phone Image. Proceedings of the European Conference on Computer Vision (ECCV) (2018), 72–87.
- Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. 2018b. Learning to Reconstruct Shape and Spatially-varying Reflectance from a Single Image. ACM Trans. Graph. 37, 6, Article 269 (Dec. 2018), 11 pages.
- Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. 2020. Image Segmentation Using Deep Learning: A Survey. arXiv:2001.05566 [cs.CV]
- Eyvind Niklasson, Alexander Mordvintsev, Ettore Randazzo, and Michael Levin. 2021. Self-Organising Textures. Distill (2021). https://doi.org/10.23915/distill.00027.003 https://distill.pub/selforg/2021/textures.
- Fabio Pellacini and Jason Lawrence. 2007. AppWand: Editing Measured Materials Using Appearance-Driven Optimization. ACM Trans. Graph. 26, 3 (July 2007), 54–es. https://doi.org/10.1145/1276377.1276444
- Patrick Pérez, Michel Gangnet, and Andrew Blake. 2003. Poisson Image Editing. ACM Trans. Graph. 22, 3 (July 2003), 313–318. https://doi.org/10.1145/882262.882269
- Lara Raad, Axel Davy, Agnès Desolneux, and Jean-Michel Morel. 2018. A survey of exemplar-based texture synthesis. Annals of Mathematical Sciences and Applications 3, 1 (2018), 89–148.
- Amir Rosenberger, Daniel Cohen-Or, and Dani Lischinski. 2009. Layered Shape Synthesis: Automatic Generation of Control Maps for Non-Stationary Textures. ACM Trans. Graph. 28, 5, Article 107 (Dec. 2009), 9 pages. https://doi.org/10.1145/1618452. 1618453
- Liang Shi, Beichen Li, Miloš Hašan, Kalyan Sunkavalli, Tamy Boubekeur, Radomir Mech, and Wojciech Matusik. 2020. MATch: Differentiable Material Graphs for Procedural Material Capture. ACM Trans. Graph. 39, 6, Article 196 (Dec. 2020), 15 pages.
- Assaf Shocher, Shai Bagon, Phillip Isola, and Michal Irani. 2019. InGAN: Capturing and Retargeting the "DNA" of a Natural Image. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, Yoshua Bengio and Yann LeCun (Eds.). http://arxiv.org/abs/1409.1556
- Ondrej Št'ava, Bedrich Beneš, Radomir Mčch, Daniel G Aliaga, and Peter Krištof. 2010. Inverse procedural modeling by automatic generation of l-systems. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 665–674.
- Alexandru Telea. 2004. An image inpainting technique based on the fast marching method. Journal of Graphics Tools 9, 1 (2004), 23–34.
- Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor Lempitsky. 2016. Texture Networks: Feed-Forward Synthesis of Textures and Stylized Images. In Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48 (New York, NY, USA) (ICML'16). JMLR.org, 1349–1357.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature*

Methods 17 (2020), 261-272. https://doi.org/10.1038/s41592-019-0686-2

- P. Welch. 1967. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics* 15, 2 (1967), 70–73. https://doi.org/10.1109/ TAU.1967.1161901
- Li-Yi Wie, Sylvain Lefebvre, Vivek Kwatra, and Greg Turk. 2009. State of the Art in Example-based Texture Synthesis. In *Eurographics 2009 - State of the Art Reports*, M. Pauly and G. Greiner (Eds.). The Eurographics Association. https://doi.org/10. 2312/egst.20091063
- Ning Xu, Brian Price, Scott Cohen, and Thomas Huang. 2017. Deep image matting. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2970–2979.
- Yang Zhou, Zhen Zhu, Xiang Bai, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. 2018. Non-stationary Texture Synthesis by Adversarial Expansion. ACM Transactions on Graphics (Proc. SIGGRAPH) 37, 4 (2018).
- Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. https://doi.org/10.1109/TIP.2003.819861