arXiv:2404.16292v1 [cs.GR] 25 Apr 2024

One Noise to Rule Them All: Learning a Unified Model of Spatially-Varying Noise Patterns

ARMAN MAESUMI, Brown University, USA
DYLAN HU* and KRISHI SARIPALLI*, Brown University, USA
VLADIMIR G. KIM, Adobe Research, USA
MATTHEW FISHER, Adobe Research, USA
SÖREN PIRK, CAU, Germany
DANIEL RITCHIE, Brown University, USA

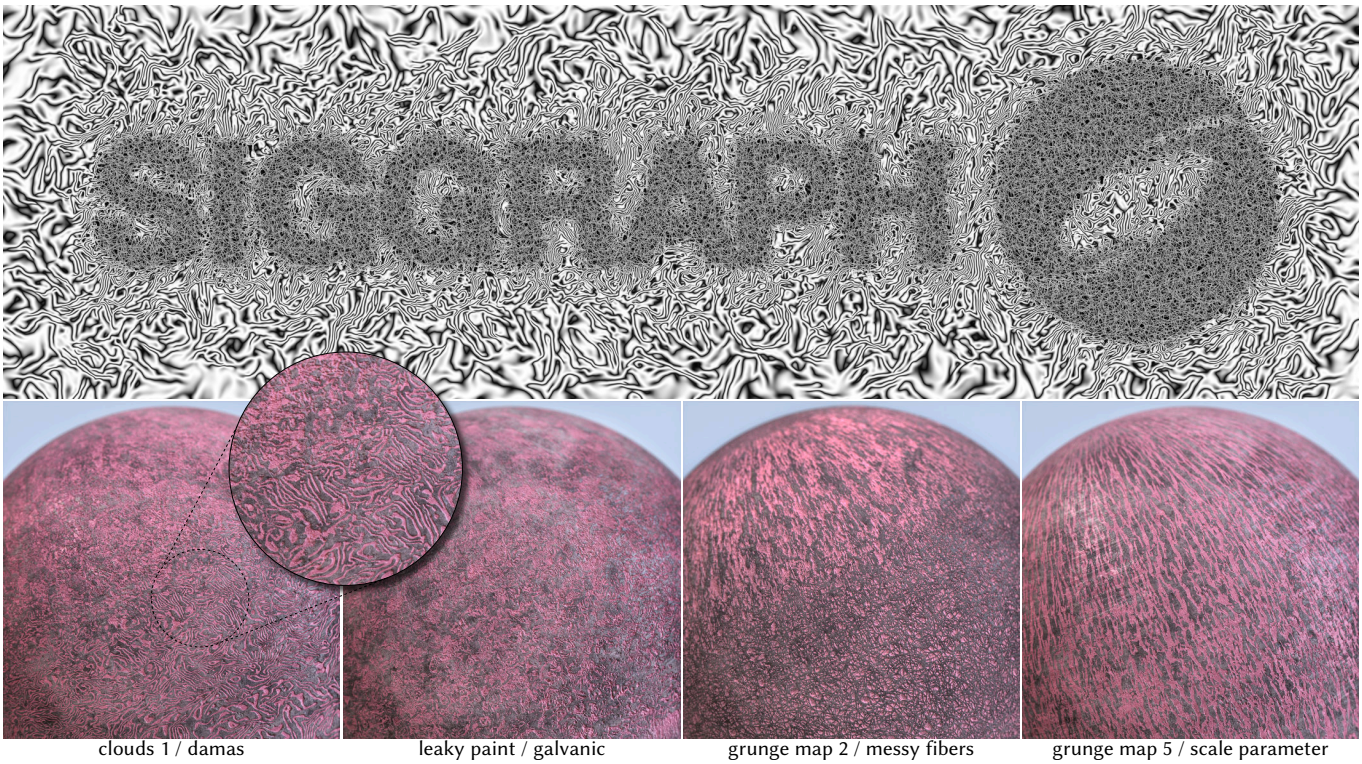

Fig. 1. Our method enables the synthesis of a wide range of noise patterns with spatially-varying characteristics. Here we show the flexibility of our unified noise model, allowing one to art-direct their noise in a granular fashion. Our model creates semantically meaningful interpolations between noise configurations; above we see the Siggraph logo written with `hay fibers` that are nested inside of `Damascus steel` striations – the scale and distortion of the steel pattern naturally interpolates into a denser pattern before transitioning into fibers. We also show renderings of a clay shader that incorporates our spatially-varying noise patterns. The first three images make use of class-interpolated noise, the final image uses parameter-interpolated noise. Please zoom into the figures for full visual detail.

Procedural noise is a fundamental component of computer graphics pipelines, offering a flexible way to generate textures that exhibit "natural" random variation. Many different types of noise exist, each produced by a separate algorithm. In this paper, we present a single generative model which can learn to generate multiple types of noise as well as blend between them. In addition, it is capable of producing spatially-varying noise blends despite not having access to such data for training. These features are enabled by training a denoising diffusion model using a novel combination of data augmentation and network conditioning techniques. Like procedural noise generators, the model's behavior is controllable via interpretable parameters plus a source of randomness. We use our model to produce a variety of visually compelling noise textures. We also present an application of our model to improving inverse procedural material design; using our model in place of fixed-type noise nodes in a procedural material graph results in higher-fidelity material reconstructions without needing to know the type of noise in advance. Open-sourced materials can be found at https://armanmaesumi.github.io/onenoise/

CCS Concepts: • **Computing methodologies** → **Texturing**; *Neural networks.*

*Both authors contributed equally to this research.

Additional Key Words and Phrases: Procedural noise, texture synthesis, texture acquisition, deep generative model

1 INTRODUCTION

Procedural noise has long been a fundamental building block in computer graphics, serving as a versatile tool for modeling fine, naturalistic details in a range of applications. It finds widespread use in creating albedo textures, bump and normal maps, terrain height fields, and density fields for volumetric phenomena such as clouds and smoke. There exists a zoo of algorithms for generating such noise patterns, for instance Perlin, Worley, and Gabor noise [Lagae et al. 2009; Perlin 1985; Worley 1996], along with systems for composing and transforming these noise patterns to synthesis more complex shaders (e.g. Blender, Adobe Substance 3D Designer, and other material editors [Adobe 2023a; Blender 2023]).

Despite the advancement of these tools, a fundamental limitation remains in the design process: the necessity to make discrete choices regarding the types of noise to employ. This constraint poses challenges in creative design, where the ideal noise type may not be immediately present among the provided "zoo", or where desired visual characteristics fall between the behaviors of existing noise types. Designers might envision patterns exhibiting different noise characteristics in various regions, transitioning smoothly without abrupt or artificial blending. The standard technique for blending between noise types—alpha blending—often yields unsatisfactory results, where features do not naturally interpolate in the image (see Figure 2).

This limitation also complicates inverse design tasks, such as recovering a procedural representation (i.e. a material node graph) of a texture from observed data. Such graphs often use multiple noise generator nodes; if the types of these nodes aren't known in advance, the search space for the inverse design problem grows combinatorially, and the need to search over discrete possible noise types limits the applicability of gradient-based optimization.

In this paper, we address these challenges by introducing a method for learning a continuous space of spatially-varying noise patterns from data that lacks any spatially-varying observations. We leverage a denoising diffusion probabilistic model (DDPM) [Ho et al. 2020] equipped with a spatially-varying conditioning module, which is trained via a novel application of CutMix data augmentation [Yun et al. 2019]. Our method enables the generation of noise that can vary smoothly across the canvas, offering a greater level of control over the resulting noise pattern. Just as with traditional procedural noise generators, the model can be controlled by changing a set of interpretable parameters or by changing the input source of randomness (i.e. random seed). Importantly, our method can generate noise images at any requested image size (beyond the size of the training data) and can also produce seamlessly tileable images.

We evaluate our model by using it to generate a variety of noise blends and spatially-varying noise patterns (including spatial variation driven by image-based masks). We also present a user interface for painting spatially-varying noise patterns, videos of this interface are included in the supplemental material. Finally, we show an application of our method to inverse procedural material design, demonstrating that using our model as a noise generator node in a material graph to be optimized can result in higher-fidelity material reconstructions without needing to know the particular noise type for that node in advance.

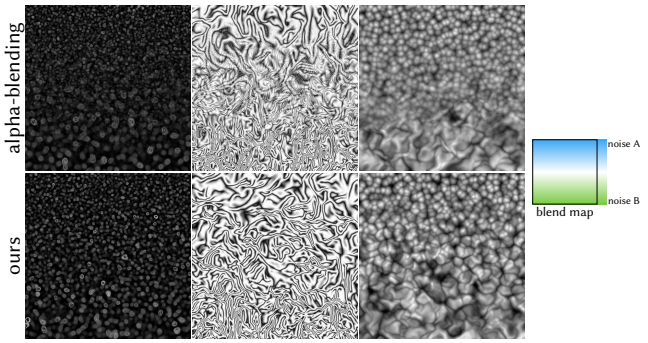

Fig. 2. Our model produces noise patterns whose characteristics (i.e. scale, ripples, distortion, etc.) interpolate naturally, creating seamless and coherent transitions. By contrast, traditional alpha-blending results in images with overlapping features, inconsistent feature opacity, and a lack of sensible transitions between the noise characteristics.

In summary, our contributions are:

(1) A generative model which can learn to generate spatially-varying noise patterns.
(2) A training scheme for the model using a novel application of CutMix augmentation, allowing the model to learn to generate spatially-varying noise patterns without having any spatially-varying training data.
(3) An application of our model to inverse procedural material design, showing improved material reconstructions without pre-specified discrete noise types.

2 RELATED WORK

Procedural noise. "Noise is the *random number generator of computer graphics*" [Lagae et al. 2010]. Procedural noise functions collectively form a body of techniques for algorithmically synthesizing patterns that mimic the randomness and irregularities in natural phenomena. These noise functions have become foundational to many applications within computer graphics, facilitating the synthesis of natural materials [Dorsey and Hanrahany 2006; Perlin 1985], terrain [Fournier et al. 1982; Galin et al. 2019; Génevaux et al. 2013; Van Der Linden et al. 2013], and motion [Hinsinger et al. 2002].

There are a wide variety of noise functions that synthesize distinct patterns, for example, Perlin [Perlin 1985, 2002], Wavelet [Cook and DeRose 2005], and Gabor noise [Lagae et al. 2009]. Such noises are often composed and transformed using various primitive operations (e.g. scalar functions and domain warping), forming patterns that are colloquially referred to as *noise* in computer graphics — while these resulting patterns may not strictly fall under the formal definition of noise [Ebert et al. 2002], we continue to use the term in this established, albeit technically imprecise sense, throughout the paper. A selection of various noise patterns is shown in Figure 5. For further background, we refer to the survey by Lagae et al. [2010].

Neural parametric texture synthesis. As noise is a central building block for textures, our work is related to methods for texture

synthesis. Texture synthesis has a deep, decades-long history, a full discussion of which is outside the scope of this paper. The most relevant methods to our work are parametric texture synthesis methods which use neural networks as their parametric model. Most of this prior work focuses on *example-based texture synthesis*, i.e. generating textures which are similar in some statistical sense to an input exemplar. Some methods do this by optimizing an image such that the features produced by passing it through a certain pre-trained neural network match those produced by the exemplar image [Gatys et al. 2015; Sendik and Cohen-Or 2017]. Others train feedforward neural networks to directly generate texture similar to an input exemplar [Li and Wand 2016; Rodriguez-Pardo and Garces 2023; Ulyanov et al. 2016; Zhou et al. 2023a, 2018]. The closest work to ours is a generative adversarial network (GAN) that does not perform example-based synthesis, but rather learns a continuous latent space of textures [Bergmann et al. 2017]. This model supports interpolation between textures, but it is not controllable by a combination of interpretable noise parameters and a source of randomness (its latent space entangles texture "style" with spatial randomness). We initially experimented with a GAN-based architecture and found that it struggled to capture a data distribution with the many different modes represented by a set of highly distinct noise types. Prior work suggests that diffusion models such as ours fare better at modeling such distributions [Xiao et al. 2022].

We also acknowledge the role of non-parametric methods in this field. These approaches synthesize textures by using the structures present within an image, without relying on a learned parametric model. For instance, Matusik et al. [2005] frame the texture interpolation problem through *morphable models* [Jones and Poggio 1998], and define a space of interpolatable textures via a simplicial complex induced by a texture similarity metric. ImageMelding [Darabi et al. 2012], on the other hand, proposes a patch-based method that smoothly "in-fills" regions of a texture via a regularized screened Poisson equation. For a more comprehensive review of this body of work, we refer to the survey by Raad et al. [2018].

Inverse procedural material design. Recovering procedural representations of materials from observations (i.e. photographs) has become an increasingly promising area of research. Material graphs—the procedural representation of choice—feature nodes that generate, transform, and filter various signals that ultimately become spatially-varying material maps (SVBRDFs). Authoring and editing these material graphs is not only time consuming, but also requires mastery of complex software [Adobe 2023a; Blender 2023]. *Inverse procedural material design* aims to alleviate this difficulty by recovering a material graph from a given image. Recent works tackle this problem in various ways, either by 1) directly predicting the parameters of a given material graph [Guo et al. 2020; Hu et al. 2019], 2) optimizing the continuous parameters of the graph via gradient-based optimization [Hu et al. 2022; Shi et al. 2020], or 3) by synthesizing a material graph from scratch [Guerrero et al. 2022; Zhou et al. 2023b].

Most relevant to our work are learned differentiable proxies, which serve as continuous relaxations for gradient-based optimization strategies to the inverse design problem. Hu et al. [2022] train several GAN-based generative models to act as differentiable proxies for *pattern* generators (i.e. a brick tiling generator node). However, these proxies are deterministic, meaning they cannot model noise functions, and critically, each pattern is represented by a separate generative model. Our unified noise DDPM is able to simultaneously capture many non-deterministic noise functions in the same continuous space, acting as a relaxation over an entire space of unique noise functions.

3 METHOD

Given a collection of procedural noise textures sampled from various parametric noise generating functions, we seek to learn a conditional generative model that captures the noise textures given their accompanying parameter configurations, enabling the synthesis of a wide range of noises from a single universal function. Our conditional generative model will be formulated in a way that facilitates synthesizing spatially-varying noise textures (i.e. spatial blends between the categorical noise type as well as noise parameters), even though the training collection does not contain such samples. That is, we assume our given collection of noise samples only exhibits globally uniform semantic properties, as shown in Figure 5.

We make use of a denoising diffusion probabilistic model (DDPM) that incorporates a spatially-varying conditioning mechanism. The model is trained using a data augmentation scheme that regulates the behavior of this conditioning module, ensuring that it behaves appropriately when given granular conditioning signals. In the following sections, we first detail the spatially-varying conditioning module, followed by the aforementioned data augmentation strategy. We defer details pertaining to our DDPM training to Section 4.

3.1 Spatially-varying conditioning

At training time, our conditioning signal is given by two vectors, $\mathbf{f}_c$, a class embedding that encodes the categorical label of a particular noise type, and $\mathbf{f}_p$, a list of parameters that determine semantic properties of the generated noise image (i.e. `scale`, `distortion`, etc.). Providing these parameters as explicit input to the model allows interpretable user control; it also disentangles control over the "style" of the generated noise (via these parameters) from the "seed" (i.e. stochastic component) of the noise functions via the initial Gaussian noise provided to the DDPM.

To facilitate synthesis of noise maps with spatially-varying properties, we employ spatially-adaptive denormalization (SPADE) [Park et al. 2019]. We first map the class and parameter vectors into a shared feature space with a small MLP, call the resulting feature vector $\mathbf{f}$. The feature vector is tiled into a spatial grid, $\mathbf{F}$, of shape $|\mathbf{f}| \times H \times W$, where H, W are the height and width respectively. This feature grid is used as input to the SPADE module. Following Dhariwal et al. [2021], we modulate the intermediate feature maps of our network's Group Normalization layers, making the SPADE conditioning function

$$S(\mathbf{h}, \mathbf{F}) = \gamma(\mathbf{F}) \odot \text{GroupNorm}(\mathbf{h}) + \beta(\mathbf{F})$$

$$\text{GroupNorm}_g(\mathbf{h}) = \frac{\mathbf{h}_g - \mu_g}{\sqrt{\sigma_g^2 + \epsilon}}$$

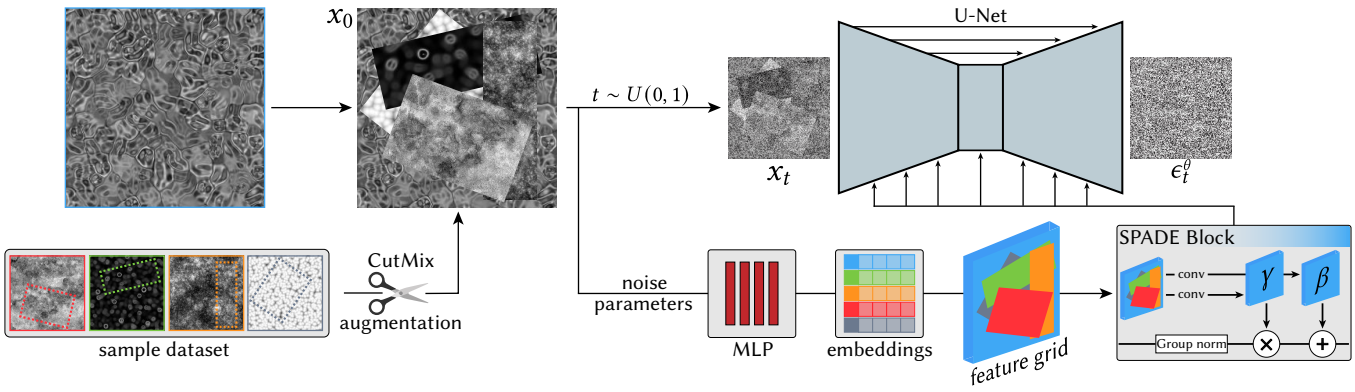

Fig. 3. Our DDPM is trained using CutMix data augmentation. We first transform the current data sample (highlighted in blue) by cutting and patching together a set of other random samples from the dataset, resulting in a training image $\mathbf{x}_0$. The noise parameters for each image patch are passed to an MLP, which projects the parameter sets into an embedding space that encodes both the noise type (class) and the noise parameters. The resulting feature vectors are tiled to form a feature grid, which is used as a conditioning signal in the U-Net's SPADE blocks, as outlined in Section 3.1.

where γ, β are convolutional layers that transform the feature grid into element-wise scales and shifts respectively, and $\mathbf{h}$ is the incoming activations of the previous layer. The scales and shifts act on the output of a Group Normalization [Wu and He 2018] block, which is computed using the mean, μ, and standard deviation, σ, of each group of channels g.

Our training data only contains noise textures that exhibit spatially *uniform* properties, meaning that all feature entries in the grid $\mathbf{F}$ are identical at training time. However, at inference time we may query the network with an artificially constructed feature grid – for instance, we can spatially blend between multiple feature vectors to produce a conditioning signal that smoothly interpolates between noise types and noise parameters, as shown in Figure 4.

3.1.1 Spherical class embeddings[1]. As previously mentioned, the network is conditioned on a class embedding, $\mathbf{f}_c$, that is learned for each noise type during training. We find that regularizing this embedding space leads to substantial improvements when interpolating between classes. In particular, we penalize the deviation of the norms of these embeddings from a target norm. For embedding vectors of dimension d that are initialized from $\mathbf{f}_c \sim \mathcal{N}^d(0, \mathbf{I}_d)$, we define the target norm, T, as the expected squared norm of a d-dimensional Gaussian vector according to the identity,

$$T_d^n := \mathbb{E}_{\mathbf{f}_c \sim \mathcal{N}^d(0,\mathbf{I}_d)} \left[||\mathbf{f}_c||_2^n \right] = 2^{n/2} \frac{\Gamma((d+n)/2)}{\Gamma(d/2)},$$

where Γ denotes Euler's gamma function. During training we then penalize the embeddings according to

$$\mathcal{L}_{\text{reg}} = \frac{1}{|C|} \sum_{c \in C} (\|\mathbf{f}_c\|_2^2 - T_d^2)^2, \tag{1}$$

this imposes a *spherical structure* on the embeddings, which means that interpolating between them can be done via spherical linear interpolation [Shoemake 1985]. We find that imposing this structure greatly improves texture blending between classes. This is likely because the MLP (which acts on these embeddings) can now learn a

[1]The following subsection introduces an enhancement to our original method, identified post-publication, which improves results in many cases.

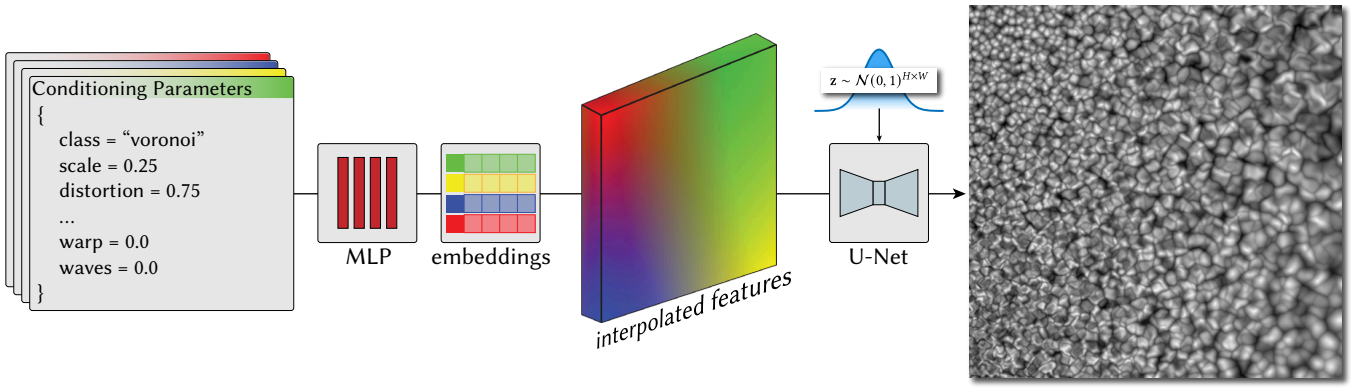

Fig. 4. At inference time, we query our network using artificially constructed feature grids, enabling a flexible way to synthesize spatially-varying noise patterns. Here we embed four sets of noise parameters, pictorially shown as one of four colors. We blend the feature vectors using bilinear interpolation, creating a smoothly-varying feature grid, which our U-Net is able to transform into a Voronoi noise pattern with non-uniform scale and distortion characteristics.

smoother mapping into the final conditioning signal. As mentioned in note 1, this regularization was identified post-publication; hence, our primary results do not make use of this change. We demonstrate its effectiveness in Figure 15 specifically.

3.2 Enhancing localized conditioning

Ideally, locally modifying the conditioning feature grid, $\mathbf{F}$, within a confined region should correspondingly alter the generated noise texture solely within that region. However, our empirical observations revealed that such localized adjustments in the conditioning signal often lead to global (or near-global) changes in the resultant noise textures. This phenomenon can be primarily attributed to the architectural design of the employed U-Net neural network – in particular, its bottleneck shape in combination with having many convolutional layers causes the receptive field of each output pixel to be relatively wide. This enables the conditioning signal to spatially propagate across the majority of the canvas, which is undesirable for the quality and controllability of our synthesized noise textures.

To address this, we incorporate a modified version of CutMix data augmentation [Yun et al. 2019] into our training procedure. CutMix, a technique that is used to enhance performance in image classification tasks, involves creating synthetic training examples by stitching together axis-aligned crops of different dataset samples. Notably, the application of CutMix has been predominantly in classification tasks, with its use in generative models being relatively unexplored and confined to augmenting the performance of GAN discriminators [Huang et al. 2021; Schonfeld et al. 2020]. In our context, we apply CutMix by combining noise textures and their corresponding conditioning feature grids, as illustrated in Figure 3. This approach not only enriches the diversity of our training dataset, but also imparts a crucial capability to our model: the ability to respond correctly to spatially localized conditioning signals.

We train our network with CutMix data augmentation applied with a probability of 0.5, i.e. for half of the training samples we train without applying CutMix. When augmenting a training image, we sample one to four (uniformly at random) auxiliary noise textures from the dataset, which are then randomly cropped (*cut*) into rectangular patches with a randomly sampled rotation $\theta \sim \mathcal{U}(0, 2\pi)$. Note that only the crop mask is rotated, not the texture itself. The resulting *mixed* sample is then a composition of a base noise texture, and one to four texture patches. It is important that all sampled patches belong to unique noise types, otherwise the resulting training image would be invalid, that is, the image would no longer belong to the distribution that we are trying to capture.

4 IMPLEMENTATION DETAILS

Noise dataset. We procure a dataset of ∼1.2 million noise images, covering 18 unique noise functions, along with dense samplings of their respective parameter sets (listed in Appendix Table 1). In particular, we sample the following noise functions from Adobe Substance 3D Designer: `cells 1`, `cells 4`, `voronoi`, `microscope view`, `grunge galvanic small`, `liquid`, `bnw spots 1`, `grunge leaky paint`, `grunge rust fine`, `grunge map 002`, `grunge map 005`, `grunge damas`, `messy fibers 3`, `perlin`, `gaussian`, `clouds 1`, `clouds 2`, `clouds 3`. A preview of our dataset is shown in Figure 5. For each noise type, we deterministically sample 16,384 parameter sets using the low-discrepancy Halton number sequence, ensuring better coverage over the space of parameters. Each parameter set is sub-sampled four times, i.e. we query the noise functions with four different seeds for each parameter set, resulting in 65,536 image samples per noise type, and 1,179,648 samples across the entire dataset. We sample at a resolution of 512 × 512px.

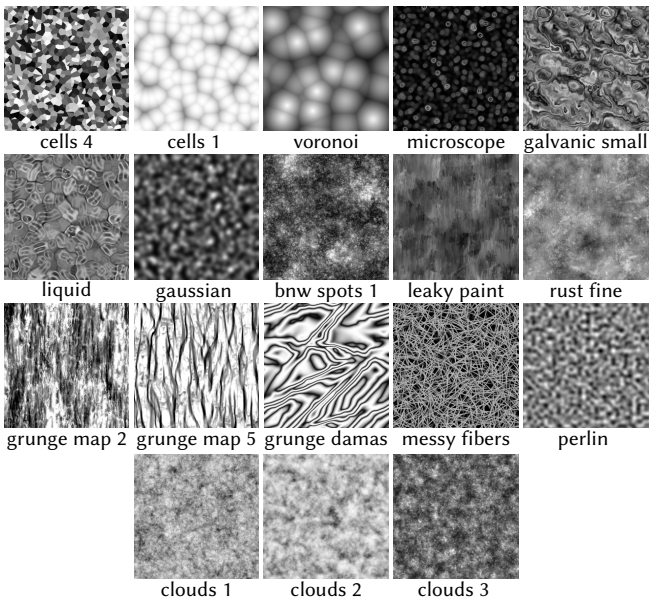

Fig. 5. Samples from our noise dataset, procured from Adobe Substance 3D Designer. We sample 18 noises along with a variety of their parameters. Note that our dataset does not contain samples with spatially-varying properties.

Network architecture. Our U-Net model is composed of three spatial levels in the encoder and decoder, each level (including the bottleneck) consists of two ResNet blocks that are conditioned on the diffusion timestep, t, and the block's respective SPADE module. The two conditioning signals are composed together as

$$\gamma_2(\mathbf{Z}) \odot (\gamma_1(t) \cdot \text{GroupNorm}(\mathbf{h}) + \beta_1(t)) + \beta_2(\mathbf{Z})$$

where γ_1, β_1 operate on t and γ_2, β_2 belong to a SPADE block. Note that the former functions produce scalar quantities, whereas the latter functions produce spatial modulation maps. Time values are first encoded using sinusoidal positional encoding. The noise parameter MLP is defined by three layers of size 128. The parameter vector $\mathbf{f}_p$ is a concatenation of all noise parameters present in the dataset, with entries zeroed out when not applicable. In total our U-Net has ∼5.1 million parameters.

Training. Following Song et al. [2020] we formulate the U-Net as a noise predictor $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{f}_c, \mathbf{f}_p)$ making the training objective

$$\mathcal{L} = \mathbb{E}_{\epsilon\sim\mathcal{N}(0,1),t\sim\mathcal{U}(0,1)} \|\epsilon - \epsilon_\theta(\mathbf{x}_t, t, \mathbf{f}_c, \mathbf{f}_p)\|_2 + \lambda\mathcal{L}_{\text{reg}} \quad (2)$$

We make use of offset noise [Guttenberg 2023], which replaces the noise sampling, $\epsilon \sim \mathcal{N}(0, 1)$, with a slightly modified distribution, $\mathcal{N}(0.1\delta, 1)$, where $\delta \sim \mathcal{N}(0, 1)$. We find that offset noise greatly helps our network resolve noise types that are extremely

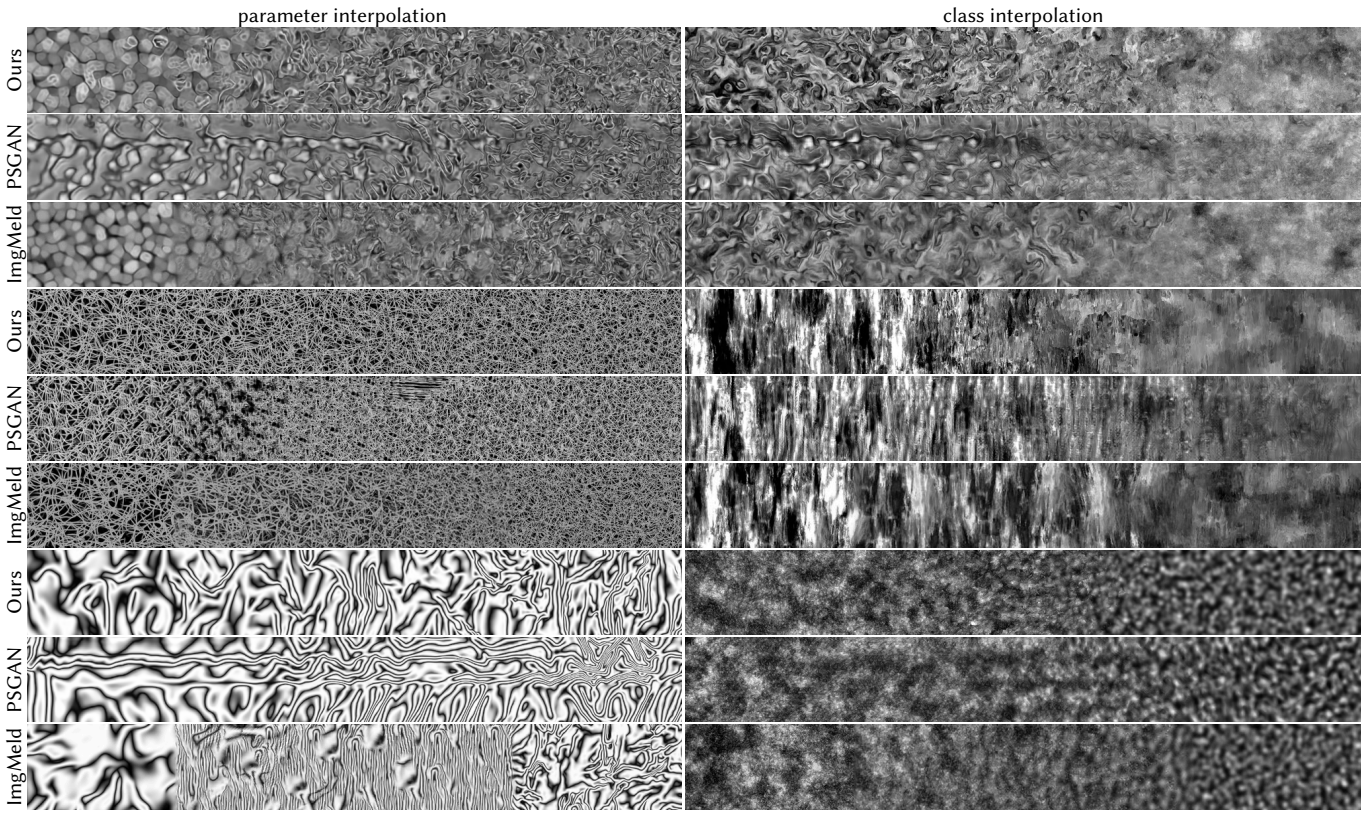

Fig. 6. We compare our method to a neural texture synthesizer, PSGAN [2017], as well as a non-parametric texture blending method, Image Melding [2012]. In the case of Image Melding, the first and last quarter of the image are given, only the remaining interior region is filled in. Both prior methods suffer from artifacts and repeated visual details, whereas our method is able to blend smoothly while synthesizing novel details throughout the canvas. We note that PSGAN produces anisotropic features that are not characteristic of the data distribution (e.g. horizontal streaks in bottom left example).

dark or bright (see Appendix Figure 16). We also note that more principled alternatives to offset noise are now available [Lin et al. 2024]. In practice the timestep sampling, $t \sim \mathcal{U}(0, 1)$, is discretized – we make use of a standard number of timesteps, i.e. 1000. We employ the cosine-beta schedule proposed by Nichol et al. [2021]. The spherical embedding regularization loss is weighted by $\lambda = 0.02$ in Equation 2. Training is done in PyTorch; we use the AdamW optimizer [Loshchilov and Hutter 2017] with a learning rate of $8 \cdot 10^{-5}$, $\beta_1 = 0.9, \beta_2 = 0.99$, and a weight decay of 0.01. We train on 8 NVIDIA RTX 3090 GPUs for ~300,000 optimization steps with batches of size 128. Noise images are downsampled to 256×256px resolution for training.

Inference details & performance. Our model can perform 80 diffusion steps per second on a single NVIDIA RTX 3090 GPU at 256×256 resolution at full precision (fp32), scaling quadratically with the resolution. We use the DDIM sampler proposed by Song et al. [2020] – unless otherwise specified, our noise figures and visualizations use 30 diffusion steps. To produce tileable noise maps, we modify all `conv2d` layers in our network to use a *circular* padding mode, making the canvas topologically toroidal.

5 RESULTS AND EVALUATION

Spatially and temporally varying noise. We demonstrate the capabilities of our noise generator by synthesizing noise patterns using an assortment of blending maps, shown in Figures 1 and 7. Our noise generator is able to be conditioned flexibly, allowing one to generate large noise patterns with diverse patterns and blends between them. We additionally show that our model is able to synthesize continuously-varying noise by interpolating the conditioning maps and/or diffusion noise (see Figure 11 and supp. video). Finally, in Figure 6 we qualitatively compare our spatially-varying noise to a GAN-based texture synthesis method, PSGAN [2017], as well as a non-parametric texture blending method, ImageMelding [2012].

Quantitative evaluation. We evaluate our noise generator's spatially-uniform noise textures against the ground truth textures from Adobe Substance 3D Designer using Frechet Inception Distance (FID), which measures the distributional similarity between synthetic samples and the data distribution [Heusel et al. 2017]. We sample 20,000 synthetic images for each noise type and evaluate the FID of each noise type separately. FID scores are reported in Appendix Table 2, alongside analogous scores for PSGAN. Our method achieves a mean FID of 20.9 and median of 13.1, while PSGAN scores 99.2 and

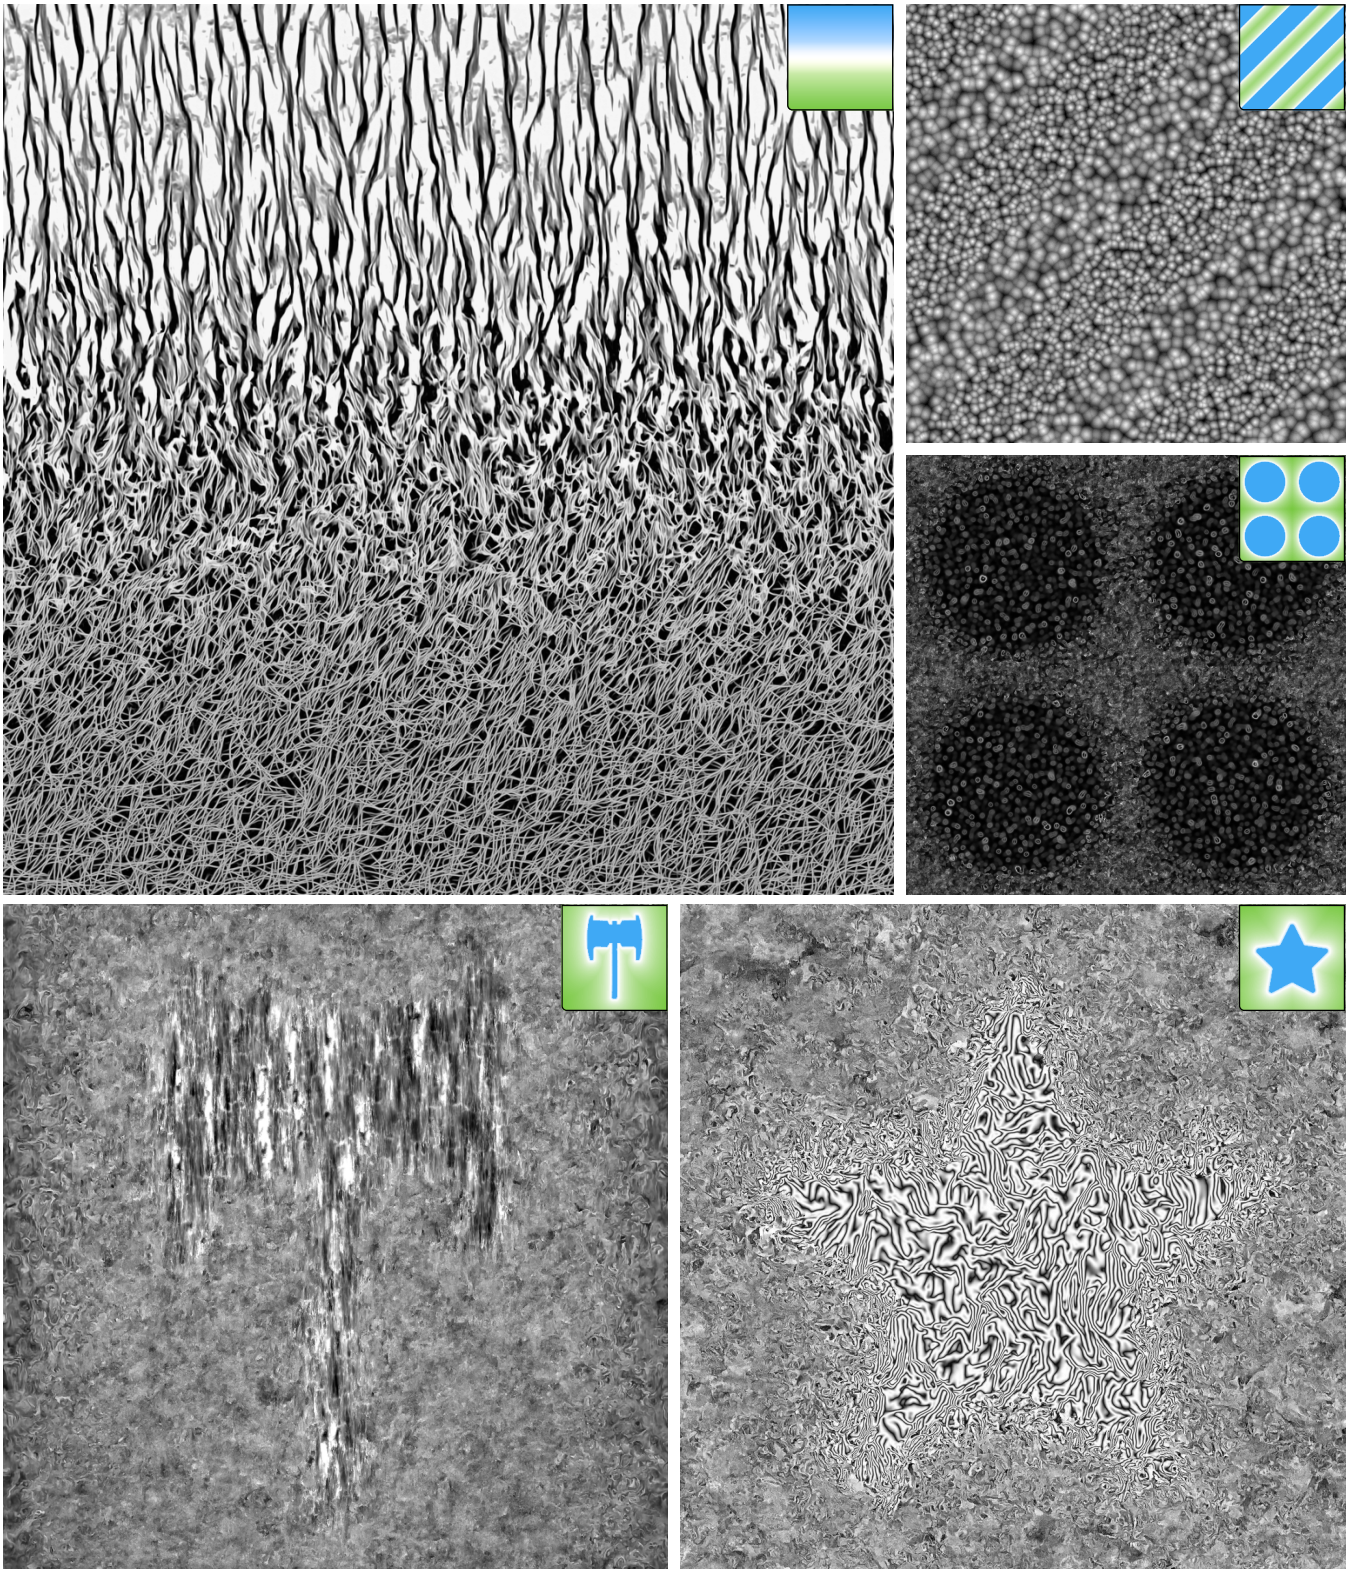

Fig. 7. Several instances of spatially-varying noise, which arise from various blended feature grids. Blend maps are shown in the top right of each panel.

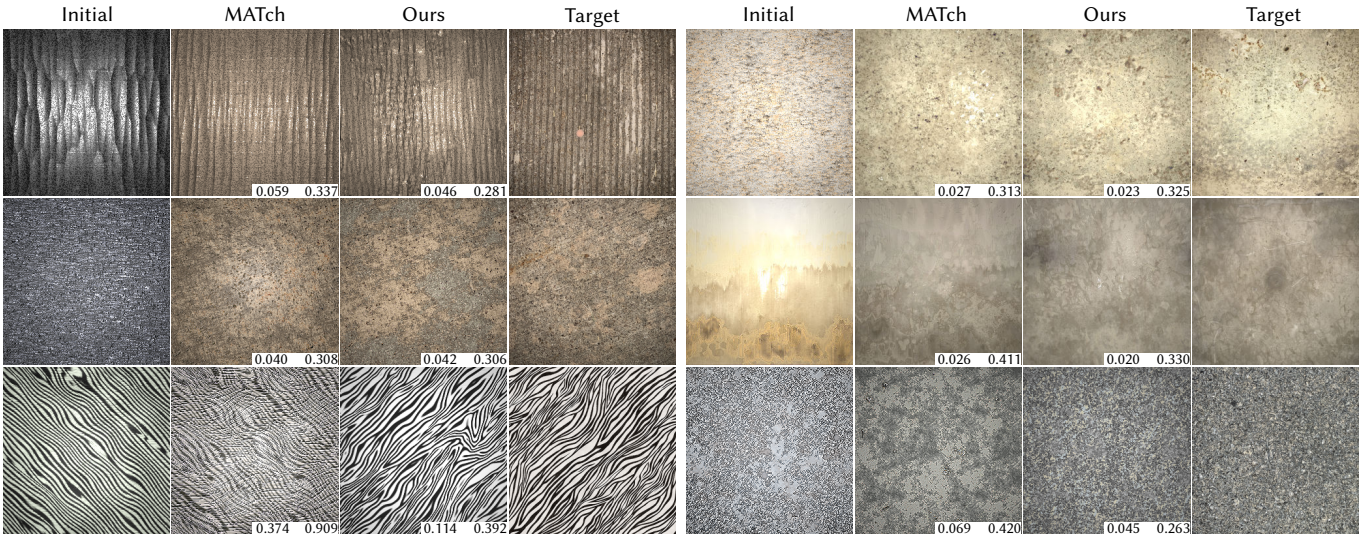

Fig. 8. We demonstrate the utility of our method in material graph optimization. Given a procedural material graph and a target photograph, our noise generative model serves as a useful prior over the space of noise functions, facilitating the recovery of non-trivial patterns that are present in the target photos, and improving the baseline result given by MATch [2020]. Two similarity scores are reported for each result with respect to the target image: MATch's feature-based texture similarity metric (left) and LPIPS [2018] (right). Lower is better for both scores.

87.5 respectively (lower is better). We note that, for many noise types, PSGAN suffers from mode collapse, which causes virtually all of the synthesized results to have similar visual features and hence its FID is severely impacted. As mentioned in Section 2, GAN-based methods struggle to faithfully capture our dataset due to the presence of many disjoint data modes [Xiao et al. 2022]. Finally, we show additional results with modified U-Net architectures in Appendix C.

Inverse material graph design. We incorporate our noise generator into the differentiable material graph library MATch [Shi et al. 2020]. Given a target photo, we select a template material graph from a selection of 88 graphs using the MATch-provided texture descriptors. In the graph, we replace one noise generator node with our model, and expose optimizable parameters for the class vector $\mathbf{f}_c$, parameter vector $\mathbf{f}_p$, and the diffusion model's latent noise $\mathbf{z}$. We include an L1 regularization on the former vectors to encourage sparsity in the optimized parameters. In Figure 8 we show that our noise generator provides a valuable prior over the space of noise functions, allowing end-to-end optimization to recover more accurate reconstructions of the target photo. We additionally show results of editing operations on an optimized material graph in Figure 9. Target photos are sampled from the Glossy dataset by Zhou et al. [2023b]. Further details for this application are in Appendix A.3.

Tileability and texture size agnosticism. Our model is able to produce noise images that are tileable, which is crucial for many downstream graphics applications. Additionally, we are able to synthesize noise at arbitrary sizes (within computational limits). In tandem, these two properties are highly desired, as it allows one to generate large tileable images that do not contain excessive repeated visual patterns, which is a common problem with tileable images. We demonstrate these capabilities in Figure 12, where we show tileable

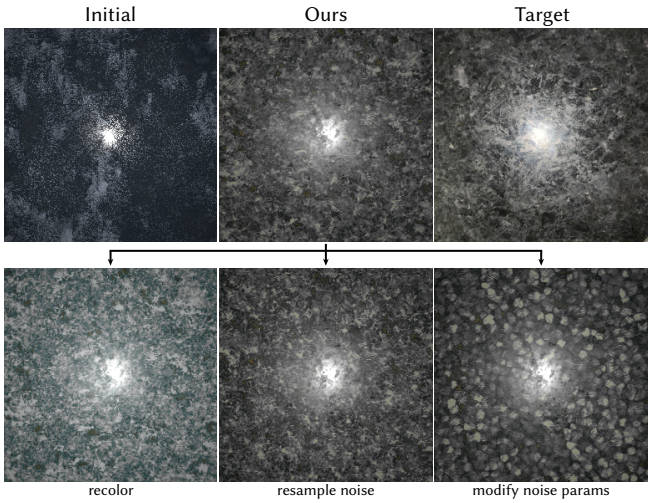

Fig. 9. After optimization, we can easily modify the resulting procedural material graphs. Here a marble material graph is optimized (first row), and edited in various ways (second row). We show the results of editing the noise colorization, resampling our diffusion noise, and modifying our model's conditioning inputs.

`grunge map 5` and `microscope` noise images, as well as the same noises synthesized on canvases that are twice as large, eliminating the repeated visual content. Additionally, in Figure 10 we show a much larger `Damascus` pattern that fills a 2048×2048 canvas – we highlight that fine-grained details are not at this size.

CutMix augmentation. We demonstrate the affect of our modified CutMix data augmentation strategy by training models with: no

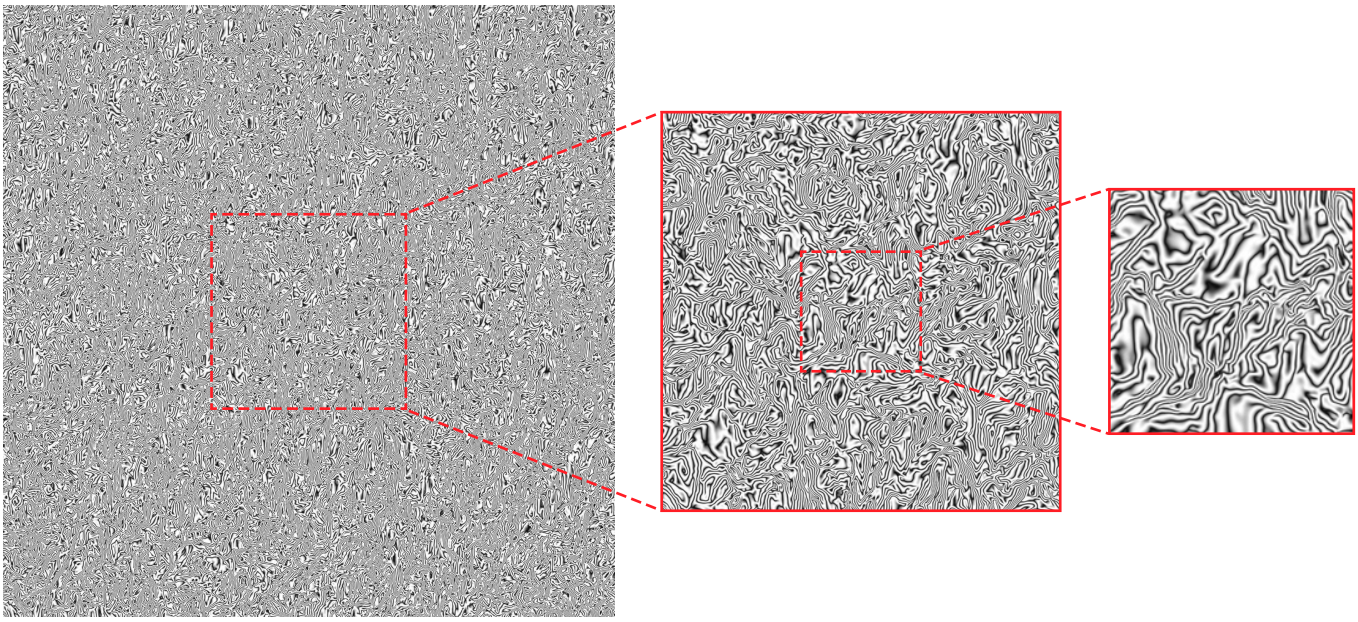

Fig. 10. Our model is capable of synthesizing noise maps at resolutions that far exceed the training resolution, thus facilitating one to paint "infinite" canvases. Here we see a 2048 × 2048px Damascus steel noise pattern with intricate fine details, which was generated via a single diffusion process. Please zoom into the figure for full details.

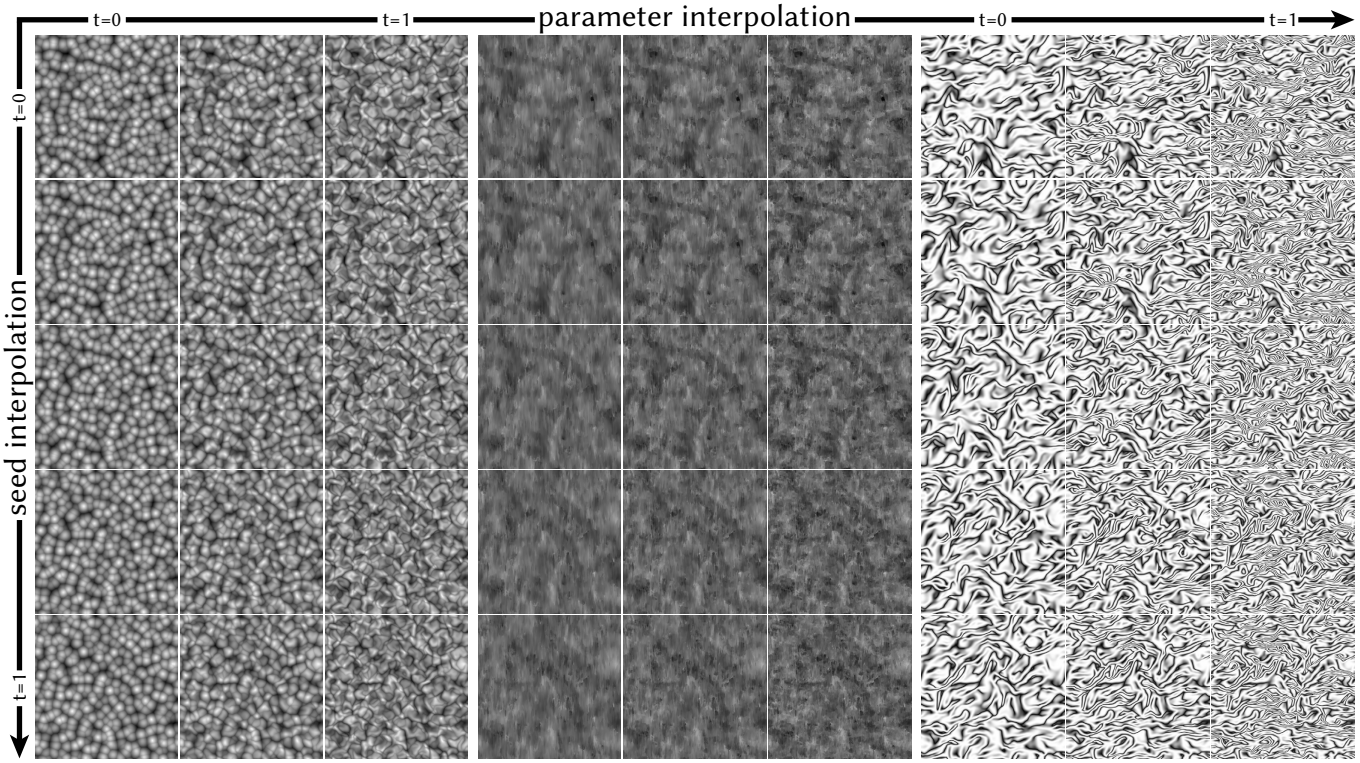

Fig. 11. Our model is able to synthesize continuously-varying textures via interpolation of the conditioning parameters and diffusion noise (analogous to the "seed" in a traditional noise function). Here we show three noise types with 1) parameters interpolating (horizontal axis), and 2) the noise seed interpolating (vertical axis). We also note that, when using the same noise seed across different noise types, we observe that similar structures appear in the output – we suggest zooming out to see this effect clearly. Noise interpolation videos are included in supplemental material.

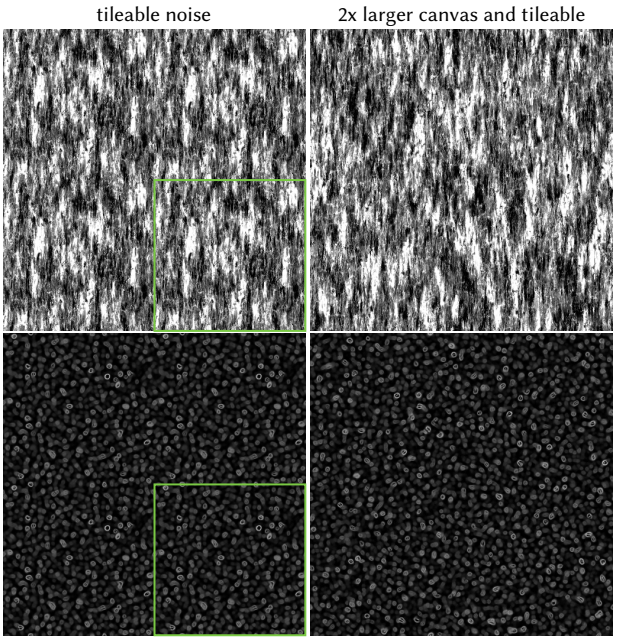

Fig. 12. Our model is able to produce seamless tileable images, as shown by the green regions, which are tiled into 2-by-2 noise maps. We also support synthesizing noise on an arbitrary canvas, avoiding noticeable repeated patterns, which is a common problem with tileable textures.

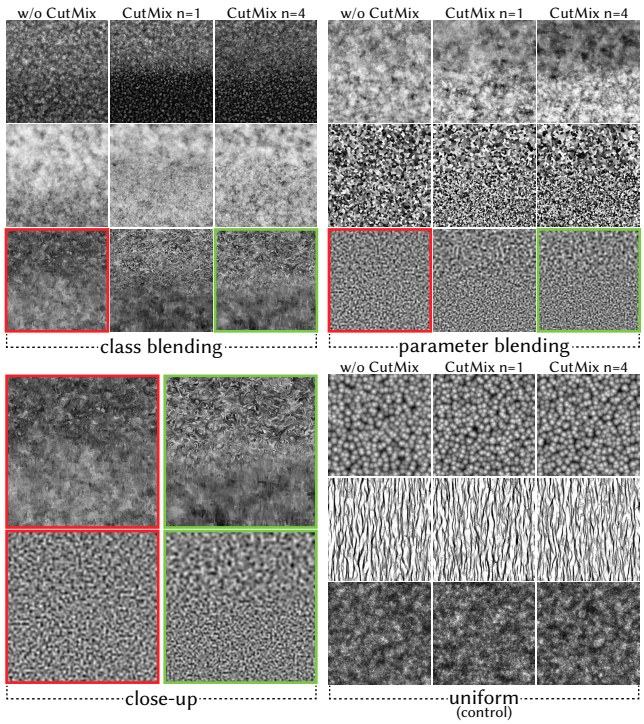

Fig. 13. Ablation study of CutMix data augmentation. Without CutMix, the U-Net fails to resolve noise maps that contain non-uniform characteristics. For instance, in row 3 of the *class blending* panel, the network cannot adequately blend between noise classes, causing the `galvanic` noise pattern to disappear entirely. We include models with one and four applications of CutMix, as detailed in Section 5. Finally, the *uniform* panel (no blending) acts as a control group – as expected, all outputs are similar.

augmentation, augmentation with only one patch (`CutMix` $n = 1$), and augmentation with one to four patches sampled uniformly at random (*CutMix* $n \in [1, 4]$). In Figure 13, we see that removing the augmentation strategy severely hinders our model's ability to respond to the conditioning signal in a local manner, producing noise patterns whose characteristics are not well-pronounced. Intuitively, this is because the conditioning signal becomes "blurred" as it propagates through the network's layers. Due to the lack of meaningful ground-truth data, we qualitatively observe that `Cutmix` $n = 1$ and `Cutmix` $n = 4$ produce similar results, with the latter being able to interpolate noise characteristics more smoothly in some cases.

6 CONCLUSION

We presented a new method for learning a single generative model for a wide range of procedural noise types. The continuous space of noise patterns learned by our model allows for smooth blending between noise types; the characteristics of the output noise are still controllable by interpretable parameters provided as input to the model. In addition, our novel approach for spatially-varying conditioning allows the model to generate spatially-varying blends between different noises despite never having seen such data at training time. We used this model to produce a variety of visually compelling textures, and we showed a proof-of-concept application of how it can be used to improve inverse procedural material design.

Our model is not without limitations. Not all pairs of noise types can be interpolated well: if the geometric features of two noise patterns are too dissimilar, the intermediate regions of a blend between them can look blurry or otherwise visually awkward (see

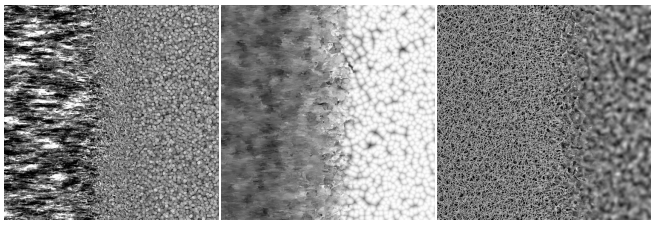

Fig. 14. Our method is unable to smoothly blend between some noise types. In the configurations above, we see artifacts in the transition regions.

Fig. 14). It may be possible to reduce these artefacts with further improvements to our CutMix-based data augmentation (e.g. non-rectangular patches; adaptive patch sampling to focus more training time on "difficult" transitions). But to some extent, such behavior is inevitable, and finding good noise pairs for interpolation is an artistic decision. It may be possible to leverage the structure of our noise embedding space to suggest good candidates for blending. Finally, diffusion models are known to model low-density modes of the data distribution less accurately [Song 2021; Um et al. 2024]; in Appendix B we demonstrate our model on additional noise types, and observe that some low-density modes of these noise functions are poorly captured. More careful data sampling techniques may be

needed when the desired noise distribution exhibits such modes; e.g. by sampling some regions of the parameter space more frequently.

Procedural material authoring tools such as Adobe Substance Designer often have deterministic pattern generators in addition to stochastic noise generators. Could our model include these in its learned space of blendable output? The deterministic nature of these patterns makes them an unnatural fit for generation via denoising diffusion model; modifications to our model and/or training procedure would be necessary to achieve this goal.

It would also be interesting to explore more techniques for interactively creating spatially-vary noise textures. For example, one might consider an interaction metaphor where the user places 'droplets' of noise on a canvas, and those droplets spread an interact with one another by e.g. solving a diffusion equation.

Finally, the application to inverse procedural material design that we presented is an early proof-of-concept; much more could be done in this direction. Distilling our model into a one-shot diffusion model [Liu et al. 2023] would make it tractable to replace *every* noise generator node in a procedural material graph with our model. Provided other graph operations are differentiable, this would open up the possibility of solving for the structure of a material graph via continuous optimization: initializing the graph to be over-complete, assigning weights to edges, and imposing a sparsity prior on edge weights. Similar approaches have been successfully employed for inferring procedural representations of 3D shapes [Kania et al. 2020; Ren et al. 2021]. In addition, allowing the optimization of spatially-varying noise generators within material graphs might allow for recovering smaller, easier-to-use graphs while still retaining the benefits of interpretable parameters.

ACKNOWLEDGMENTS

This material is based upon work that was supported by the National Science Foundation Graduate Research Fellowship under Grant No. 2040433. Part of this work was done while Arman Maesumi was an intern at Adobe Research.

REFERENCES

Adobe. 2023a. Adobe Substance 3D Designer. https://www.adobe.com/products/substance3d-designer.html.

Adobe. 2023b. Adobe Substance 3D Documentation. https://helpx.adobe.com/substance-3d-designer/substance-compositing-graphs/nodes-reference-for-substance-compositing-graphs/node-library.html.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein Generative Adversarial Networks. 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, 214–223. https://proceedings.mlr.press/v70/arjovsky17a.html

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. https://proceedings.mlr.press/v70/bergmann17a.html

Blender. 2023. Blender. https://www.blender.org.

Robert L Cook and Tony DeRose. 2005. Wavelet noise. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 803–811.

Soheil Darabi, Eli Shechtman, Connelly Barnes, Dan B Goldman, and Pradeep Sen. 2012. Image Melding: Combining Inconsistent Images using Patch-based Synthesis. *ACM Transactions on Graphics (TOG) (Proceedings of SIGGRAPH 2012)* 31, 4, Article 82 (2012), 82:1–82:10 pages.

Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* 34 (2021), 8780–8794.

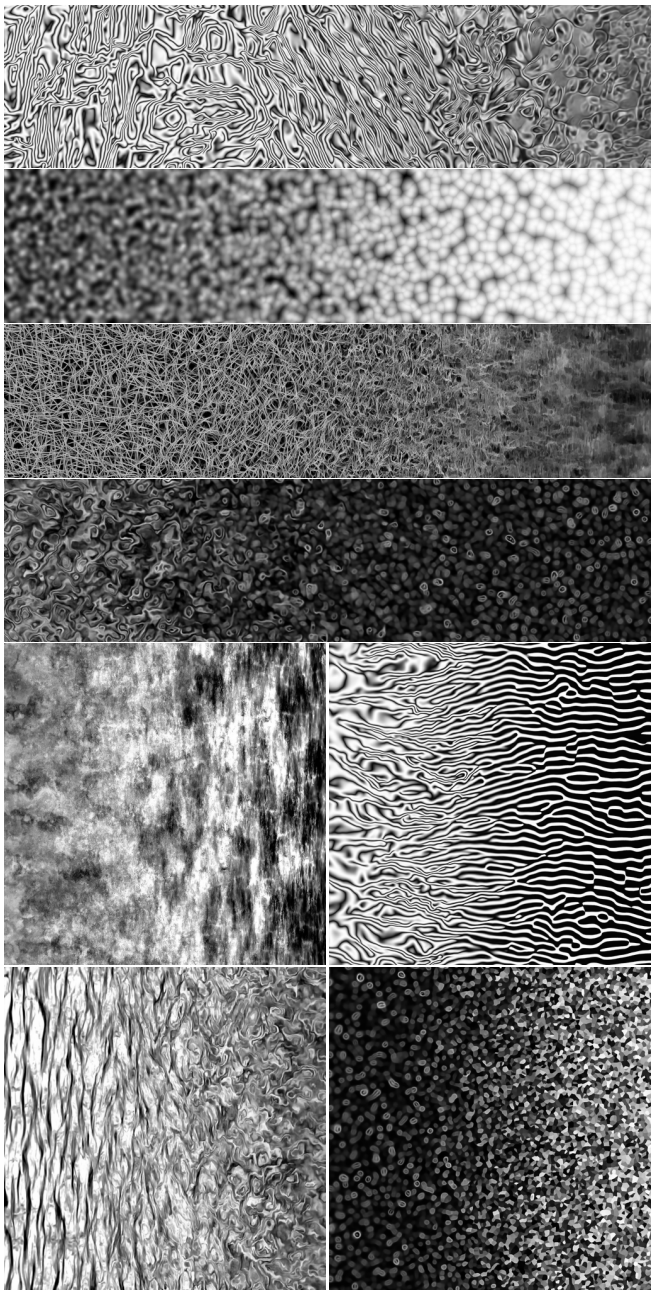

Fig. 15. Noise interpolations that arise from our spherically regularized embedding space, as described in Section 3.1.1. The blending between noise types appears organic and is largely free of the artifacts noted in Figure 14.

Julie Dorsey and Pat Hanrahany. 2006. Modeling and rendering of metallic patinas. In *ACM SIGGRAPH 2006 Courses*. 2–es.

David Ebert, Kenton Musgrave, Darwyn Peachey, Ken Perlin, and Steve Worley. 2002. *Texturing And Modeling. A Procedural Approach* (3 ed.). Morgan Kaufmann.

Alain Fournier, Don Fussell, and Loren Carpenter. 1982. Computer rendering of stochastic models. *Commun. ACM* 25, 6 (1982), 371–384.

Eric Galin, Eric Guérin, Adrien Peytavie, Guillaume Cordonnier, Marie-Paule Cani, Bedrich Benes, and James Gain. 2019. A review of digital terrain modeling. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 553–577.

Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. 2015. Texture Synthesis Using Convolutional Neural Networks. In *NeurIPS*.

Jean-David Génevaux, Éric Galin, Eric Guérin, Adrien Peytavie, and Bedrich Benes. 2013. Terrain generation using procedural models based on hydrology. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–13.

Paul Guerrero, Miloš Hašan, Kalyan Sunkavalli, Radomír Měch, Tamy Boubekeur, and Niloy J Mitra. 2022. MatFormer: A generative model for procedural materials. *arXiv preprint arXiv:2207.01044* (2022).

Yu Guo, Miloš Hašan, Lingqi Yan, and Shuang Zhao. 2020. A bayesian inference framework for procedural material parameter estimation. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 255–266.

Nicholas Guttenberg. 2023. Diffusion with Offset Noise. (2023).

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* 30 (2017).

Damien Hinsinger, Fabrice Neyret, and Marie-Paule Cani. 2002. Interactive animation of ocean waves. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*. 161–166.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.

Yiwei Hu, Julie Dorsey, and Holly Rushmeier. 2019. A novel framework for inverse procedural texture modeling. *ACM Transactions on Graphics (ToG)* 38, 6 (2019), 1–14.

Yiwei Hu, Paul Guerrero, Milos Hasan, Holly Rushmeier, and Valentin Deschaintre. 2022. Node graph optimization using differentiable proxies. In *ACM SIGGRAPH 2022 conference proceedings*. 1–9.

Zhizhong Huang, Junping Zhang, Yi Zhang, and Hongming Shan. 2021. DU-GAN: Generative adversarial networks with dual-domain U-Net-based discriminators for low-dose CT denoising. *IEEE Transactions on Instrumentation and Measurement* 71 (2021), 1–12.

Michael J Jones and Tomaso Poggio. 1998. Multidimensional morphable models. In *Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271)*. IEEE, 683–688.

Kacper Kania, Maciej Zieba, and Tomasz Kajdanowicz. 2020. UCSG-NET- Unsupervised Discovering of Constructive Solid Geometry Tree. In *Advances in Neural Information Processing Systems*.

Ares Lagae, Sylvain Lefebvre, Rob Cook, Tony DeRose, George Drettakis, David S Ebert, John P Lewis, Ken Perlin, and Matthias Zwicker. 2010. A survey of procedural noise functions. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 2579–2600.

Ares Lagae, Sylvain Lefebvre, George Drettakis, and Philip Dutré. 2009. Procedural noise using sparse Gabor convolution. *ACM Transactions on Graphics (TOG)* 28, 3 (2009), 1–10.

Chuan Li and Michael Wand. 2016. Precomputed Real-Time Texture Synthesis with Markovian Generative Adversarial Networks. In *Computer Vision – ECCV 2016*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 702–716.

Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. 2024. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 5404–5411.

Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, and Qiang Liu. 2023. InstaFlow: One Step is Enough for High-Quality Diffusion-Based Text-to-Image Generation. arXiv:2309.06380 [cs.LG]

Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017).

Wojciech Matusik, Matthias Zwicker, and Frédo Durand. 2005. Texture design using a simplicial complex of morphable textures. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 787–794.

Alexander Quinn Nichol and Prafulla Dhariwal. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*. PMLR, 8162–8171.

Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. 2019. Semantic Image Synthesis with Spatially-Adaptive Normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Ken Perlin. 1985. An Image Synthesizer. In *Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '85)*. Association for Computing Machinery, New York, NY, USA, 287–296. https://doi.org/10.1145/325334.325247

Ken Perlin. 2002. Improving noise. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. 681–682.

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.

Daxuan Ren, Jianmin Zheng, Jianfei Cai, Jiatong Li, Haiyong Jiang, Zhongang Cai, Junzhe Zhang, Liang Pan, Mingyuan Zhang, Haiyu Zhao, et al. 2021. CSG-Stump: A Learning Friendly CSG-Like Representation for Interpretable Shape Parsing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12478–12487.

Carlos Rodriguez-Pardo and Elena Garces. 2023. SeamlessGAN: Self-Supervised Synthesis of Tileable Texture Maps. *IEEE Transactions on Visualization and Computer Graphics* 29, 6 (jun 2023), 2914–2925. https://doi.org/10.1109/TVCG.2022.3143615

Edgar Schonfeld, Bernt Schiele, and Anna Khoreva. 2020. A u-net based discriminator for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8207–8216.

Omry Sendik and Daniel Cohen-Or. 2017. Deep Correlations for Texture Synthesis. *ACM Trans. Graph.* 36, 5, Article 161 (jul 2017), 15 pages. https://doi.org/10.1145/3015461

Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. 2021. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 3531–3539.

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 (Dec. 2020), 1–15.

Ken Shoemake. 1985. Animating rotation with quaternion curves. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*. 245–254.

Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising Diffusion Implicit Models. *arXiv:2010.02502* (October 2020). https://arxiv.org/abs/2010.02502

Yang Song. 2021. Generative Modeling by Estimating Gradients of the Data Distribution. https://yang-song.net/blog/2021/score/.

Thibault Tricard, Semyon Efremov, Cédric Zanni, Fabrice Neyret, Jonàs Martínez, and Sylvain Lefebvre. 2019. Procedural phasor noise. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–13.

Dmitry Ulyanov, Vadim Lebedev, Vedaldi Andrea, and Victor Lempitsky. 2016. Texture Networks: Feed-forward Synthesis of Textures and Stylized Images. In *Proceedings of The 33rd International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 48)*, Maria Florina Balcan and Kilian Q. Weinberger (Eds.). PMLR, New York, New York, USA, 1349–1357.

Soobin Um, Suhyeon Lee, and Jong Chul Ye. 2024. Don't Play Favorites: Minority Guidance for Diffusion Models. In *The Twelfth International Conference on Learning Representations*. https://openreview.net/forum?id=3NmO9lY4Jn

Roland Van Der Linden, Ricardo Lopes, and Rafael Bidarra. 2013. Procedural generation of dungeons. *IEEE Transactions on Computational Intelligence and AI in Games* 6, 1 (2013), 78–89.

Steven Worley. 1996. A cellular texture basis function. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*. Association for Computing Machinery, New York, NY, USA, 291–294. https://doi.org/10.1145/237170.237267

Yuxin Wu and Kaiming He. 2018. Group Normalization. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XIII (Lecture Notes in Computer Science, Vol. 11217)*, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). Springer, 3–19. https://doi.org/10.1007/978-3-030-01261-8_1

Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. 2022. Tackling the Generative Learning Trilemma with Denoising Diffusion GANs. In *ICLR*.

Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. 2019. CutMix: Regularization Strategy to Train Strong Classifiers with Localizable Features. In *International Conference on Computer Vision (ICCV)*.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.

Xilong Zhou, Milos Hasan, Valentin Deschaintre, Paul Guerrero, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Nima Khademi Kalantari. 2023b. PhotoMat: A Material Generator Learned from Single Flash Photos. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–11.

Yang Zhou, Kaijian Chen, Rongjun Xiao, and Hui Huang. 2023a. Neural Texture Synthesis with Guided Correspondence. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. 18095–18104.

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), 49:1–49:13.

Noise function	Sampled parameters	Range	Noise function	Sampled parameters	Range
cells 4	scale	$i[5, 50]$	grunge rust fine	base grunge contrast	$f[-0.3, 0.3]$
cells 1	scale	$i[10, 50]$		base warp intensity	$f[0.0, 256.0]$
voronoi	scale	$f[5.0, 15.0]$	grunge damas	distortion	$f[0.0, 1.0]$
	distortion intensity	$f[0.0, 1.0]$		divisions	$i[4, 16]$
	distortion scale multiplier	$f[1.0, 2.0]$		waves	$i[1, 3]$
microscope view	scale	$i[25, 64]$		details	$f[0.0, 0.75]$
	warp intensity	$f[0.0, 1.0]$		rotation random	$f[0.05, 1.0]$
bnw spots 1	scale	$i[1, 3]$	grunge map 002	N/A	N/A
liquid	scale	$i[10, 45]$	grunge map 005	N/A	N/A
	warp intensity	$f[0.1, 0.8]$	messy fibers 3	scale	$i[1, 3]$
grunge galvanic small	crispness	$f[0.0, 0.75]$	perlin	scale	$i[10, 50]$
	dirt	$f[0.0, 1.0]$	gaussian	scale	$i[10, 50]$
	micro distortion	$f[0.0, 0.6]$	clouds 1	scale	$i[1, 5]$
grunge leaky paint	leak intensity	$f[0.0, 1.0]$	clouds 2	scale	$i[1, 5]$
	leak scale	$i[1, 8]$	clouds 3	scale	$i[1, 3]$
	leak crispness	$f[0.0, 0.8]$			

Table 1. We enumerate the noise functions that are sampled from Adobe Substance 3D Designer, along with their parameters and accompanying ranges – integer and real ranges are denoted by i and f respectively.

	FID↓			FID↓	
Noise	**PSGAN**	**Ours**	**Noise**	**PSGAN**	**Ours**
cells 4	218.8	33.6	rust fine	88.5	12.3
cells 1	171.3	2.4	damas	56.1	71.0
voronoi	149.3	12.5	map 002	37.3	13.7
microscope	133.1	29.7	map 005	155.4	34.5
bnw spots 1	22.4	4.4	fibers	86.4	34.6
liquid	163.7	38.0	perlin	47.8	4.8
galvanic small	79.2	24.9	gaussian	45.2	1.7
leaky paint	155.1	44.3	clouds 1	20.4	1.5
clouds 3	38.4	2.9	clouds 2	110.8	9.0
Mean	99.2	20.9	**Median**	87.5	13.1

Table 2. We compare our FID scores for each noise type alongside PSGAN. The mean and median of all values are shown as well.

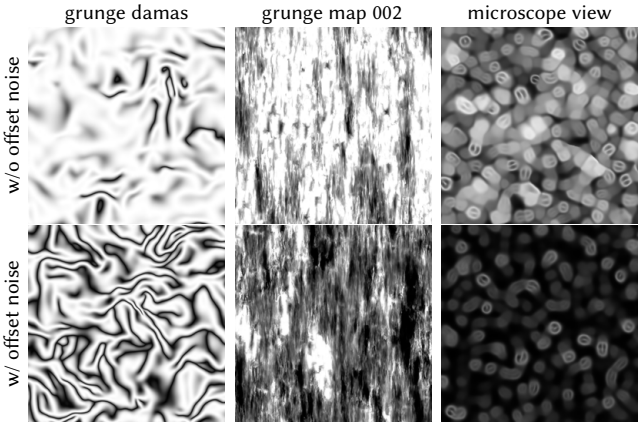

Fig. 16. Ablation of offset noise. Without offset noise, the network occasionally fails to synthesize noise images with extreme intensity distributions (i.e. intensely dark and bright images). We show representative examples above.

A EXPERIMENTAL AND IMPLEMENTATION DETAILS

A.1 Noise dataset details

We include a table enumerating all sampled noise types, the parameters that we sample, as well as the parameter ranges (see Table 1). In general, we choose to include any parameters that lead to noticeable changes in the resulting noise, with exceptions to parameters that simply act as color correction (i.e. `grunge map 005`'s *contrast* parameter). We note that, despite their names, some of these noises correspond to those that belong to graphics literature; for instance, `cells 4`, `cells 1`, and `voronoi` are variants of Worley noise [Worley 1996]. For more details about these noises, please refer to the Adobe Substance documentation [2023b].

The conditioning vector $\mathbf{f}_p$ contains an entry for each unique noise parameter – we treat identical parameter names as separate, with the exception of `scale`, which is treated as a single entry in the vector. Parameters are independently normalized to the range $[0, 1]$ before being passed to our conditioning MLP. Below we show many samples of our model's outputs with all listed noise parameters being sampled randomly (see Figs. 19 to 23).

A.2 PSGAN Baseline

We slightly modify the PSGAN [2017] architecture for the results shown in Figure 6. Since PSGAN is a purely unconditional generative model (i.e. it has no class or parameter conditioning beyond randomized latent code inputs), we endow the PSGAN model with SPADE conditioning blocks that are identical to the ones used throughout our method. We made the discriminator slightly larger to compensate for the added parameters in the generator. Finally, we use the WGAN loss [Arjovsky et al. 2017]. The remaining parts of the network and training are largely kept as-is from the open implementation. In total the generator has 30 million parameters.

A.3 Inverse Material Design Details

We expose three parameters to the MATch differentiable material graph optimizer; a soft-class vector $\mathbf{v}_c$, the parameter vector $\mathbf{f}_p$, and the diffusion Gaussian noise image $\mathbf{z}$. The soft-class vector represents a list of $[0, 1]$ values that are used to index into our class embeddings – this amounts to taking a convex combination over the class embeddings. During optimization, we apply the SoftMax function to $\mathbf{v}_c$ to ensure the values are valid. In order to encourage sparsity, we include an L1 regularization term on both $\mathbf{v}_c$ and $\mathbf{f}_p$, using a weight of 0.1 for both L1 terms in the final optimization objective. Additionally, we include a scheduled temperature parameter τ into the SoftMax operator by performing element-wise division of $\mathbf{v}_c$ by τ, $\text{softmax}(v, t) = \frac{e^{v_i/t}}{\sum_{j=1}^{K} e^{v_j/t}}$. The temperature is initialized to 0.25 and is updated at every optimization step via the update rule $\tau' \leftarrow \tau \cdot 0.97$. Finally we clamp τ to be at minimum 0.01. This scheduled temperature forces the optimization to "hone in" on a single class. We use a learning rate of 0.01 for all graphs. After $\tau = 0.01$ we perform a warm-restart of the optimizer and add noise to the exposed parameters to avoid local minima.

B TRAINING WITH ADDITIONAL NOISE TYPES

We train our model on two additional noise types that are commonly used in graphics literature: Phasor noise [2019] and Gabor noise [2009]. We utilize the same data sampling method as mentioned in Section 4. The parameters that we sampled for both noise types, as well as their ranges, are enumerated in Table 3. We use the released implementations for both noise types to accrue our training data. Examples of spatially-varying images using parameter interpolation and class interpolation for both Phasor and Gabor noise are shown in Figs. 17 and 18.

We note that our model exhibits an FID of 93.6 and 129.3 on Phasor and Gabor noise respectively, which is notably higher than the FID scores for other noise types. Upon inspection of our model's outputs, we see that some parameter configurations are not well captured by our training. For example, Gabor noise exhibits a significantly different visual appearance when its principal frequency and kernel width are very small; however, since this appearance is only captured by a narrow subset of the parameter space, the dataset thus contains much fewer of such samples. One drawback of diffusion models is that they model low-density regions of the data distribution less accurately (compared to higher-density regions) [Song 2021; Um et al. 2024], and hence our performance is worse in such situations. More careful data sampling methods may be needed when the desired noise distribution exhibits low-density modes; e.g. by sampling some regions of the parameter space more frequently.

C MODEL ARCHITECTURES

The U-Net model detailed in Section 4 has just ∼5.1 million parameters and is constructed using two downsampling blocks, a bottleneck block, and two upsampling blocks. Each block contains two ResNet sub-blocks. An additional ResNet sub-block is placed at the end of the network. We use channel dimensions of 32 and 64 for the outer

Noise function	Sampled parameters	Range
phasor noise	principal frequency, F	$f[8.0, 32.0]$
	num cells	$i[1, 12]$
	phasor density	$f[0.3, 0.5]$
	factor angle spread, θ	$f[0.0, 1.0]$
gabor noise	principal frequency, F	$f[0.02, 0.08]$
	kernel width, α	$f[0.01, 0.35]$
	kernel orientation, ω	$f[0.0, 2\pi]$

Table 3. Parameters and sampling ranges for additional noise types.

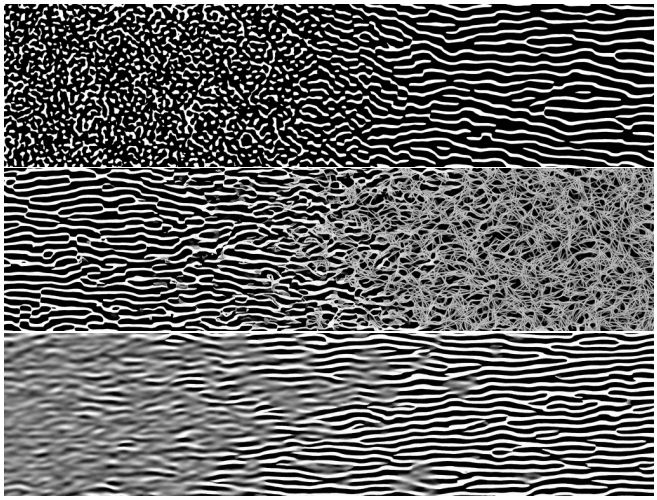

Fig. 17. Examples of our model's Phasor noise. From top to bottom: isotropic Phasor to anisotropic Phasor (parameter interpolation), anisotropic Phasor to Messy Fibers, and anisotropic Gabor to anisotropic Phasor (class interpolation).

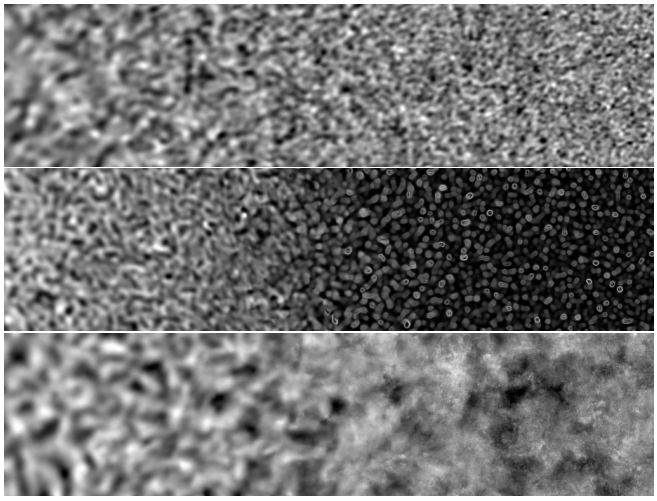

Fig. 18. Examples of our model's Gabor noise. From top to bottom: Gabor to Gabor (parameter interpolation), Gabor to Microscope View, and Gabor to Grunge Rust Fine (class interpolation).

Architecture	Block dimensions	Attention	Num. Parameters	**Mean FID↓**	**Median FID↓**	**Steps per second↑**
`Model-XS`	(32, 64, 128)	✗	5.1M	20.9	13.1	**79.5/s**
`Model-S`	(32, 64, 64, 128)	`heads` = 4, `dim` = 32	6.5M	14.0	10.8	28.6/s
`Model-M`	(64, 128, 128, 256)	`heads` = 4, `dim` = 32	22.5M	**10.3**	**10.4**	20.1/s

Table 4. Three model architectures with varying parameter counts and design choices are shown. The *Block dimensions* column shows the dimension counts for the downsampling blocks followed by the bottleneck block (the networks are symmetric; i.e. the upsampling blocks follow the same dimensions in reverse order). Our primary model (first row) is significantly faster than the other models, but compromises slightly on FID.

blocks, and 128 for the bottleneck. The ResNet sub-blocks are globally conditioned on the diffusion time t, and spatially conditioned (via SPADE) by 128-dim noise embeddings. We will refer to this model as `Model-XS` (extra small) below.

We train two additional model architectures to evaluate the potential performance of larger and more complex models. `Model-S` is identical to `Model-XS`, with the exception of added *linear attention* layers inside of each block [Shen et al. 2021], and the addition of an extra set of downsampling/upsampling blocks. Similarly, `Model-M` is identical to `Model-S`, with the exception of larger channel dimensions. In Table 4 we summarize these architectures as well as their FID scores and inference performance. The models were trained for the same number of optimization steps following the details in Section 4.

For our inverse material design application, we found that using `Model-XS` was most suitable due to the added memory cost of attention layers. For consistency, we used this model for all figures in the main text; however, in applications that do not require such light-weight networks, the larger models may be suitable.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

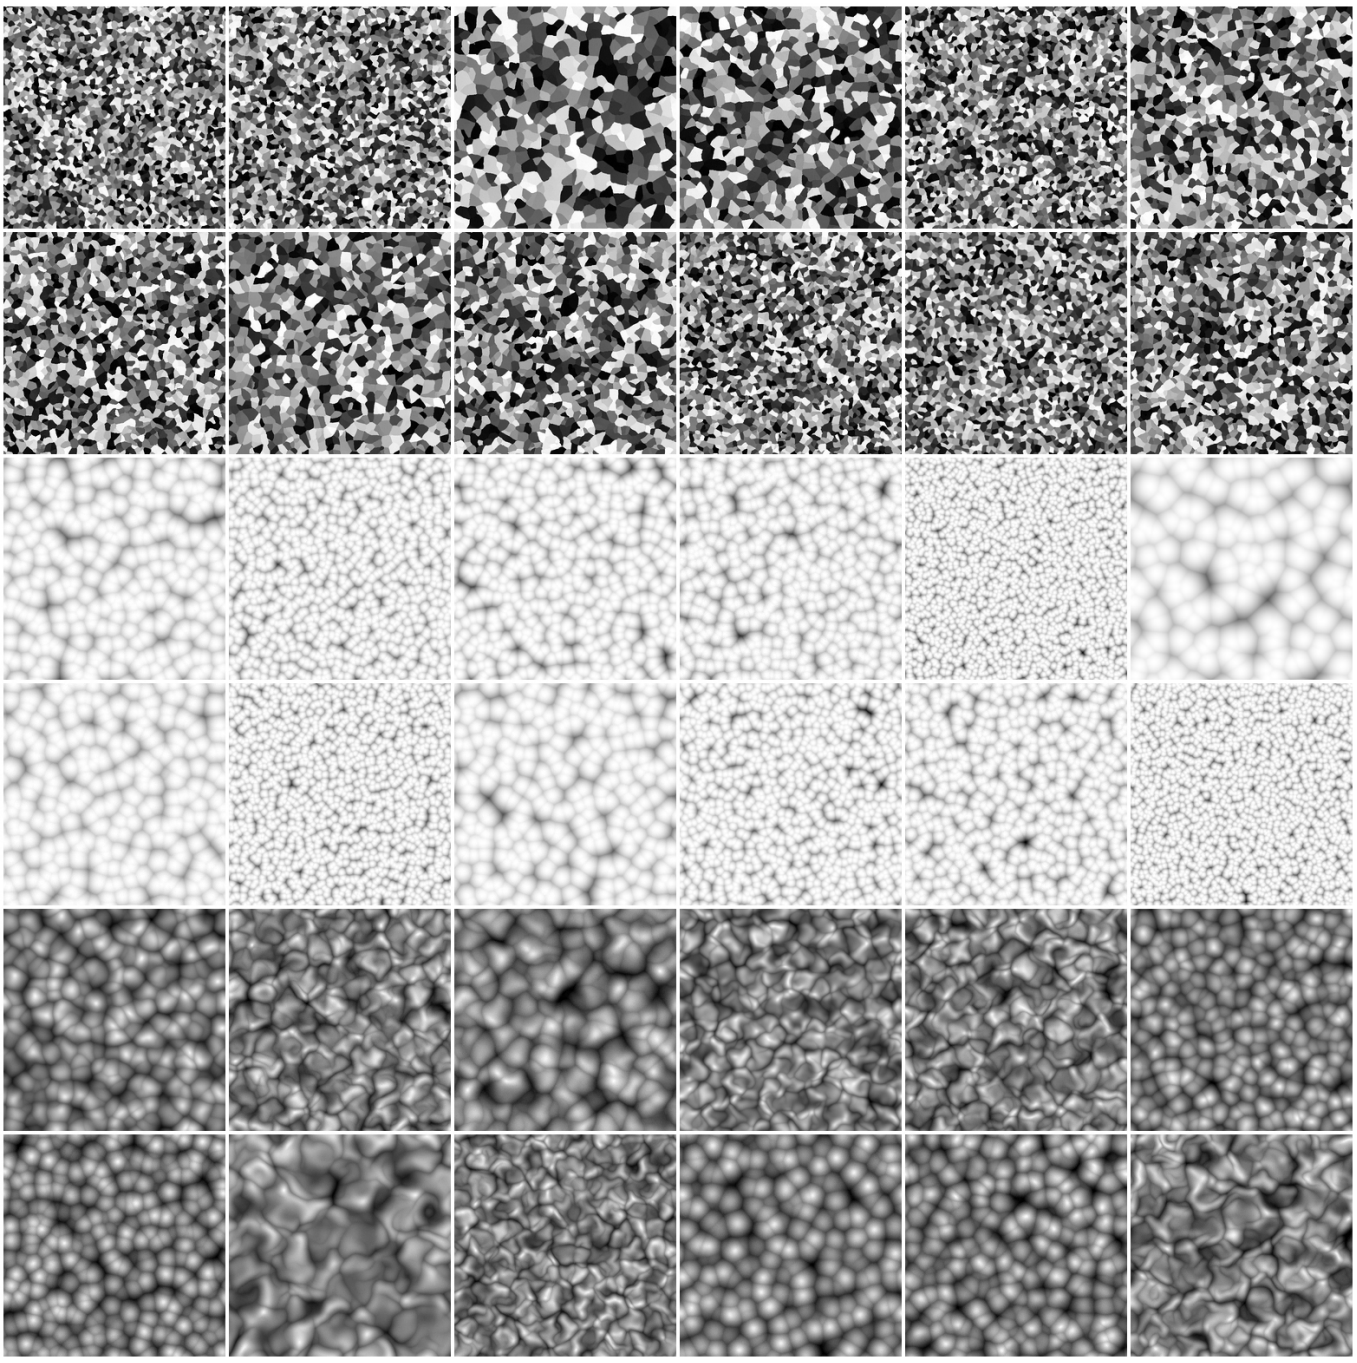

Fig. 19. Random samples of our model's `cells 4`, `cells 1`, and `voronoi` noises at 256×256 resolution.

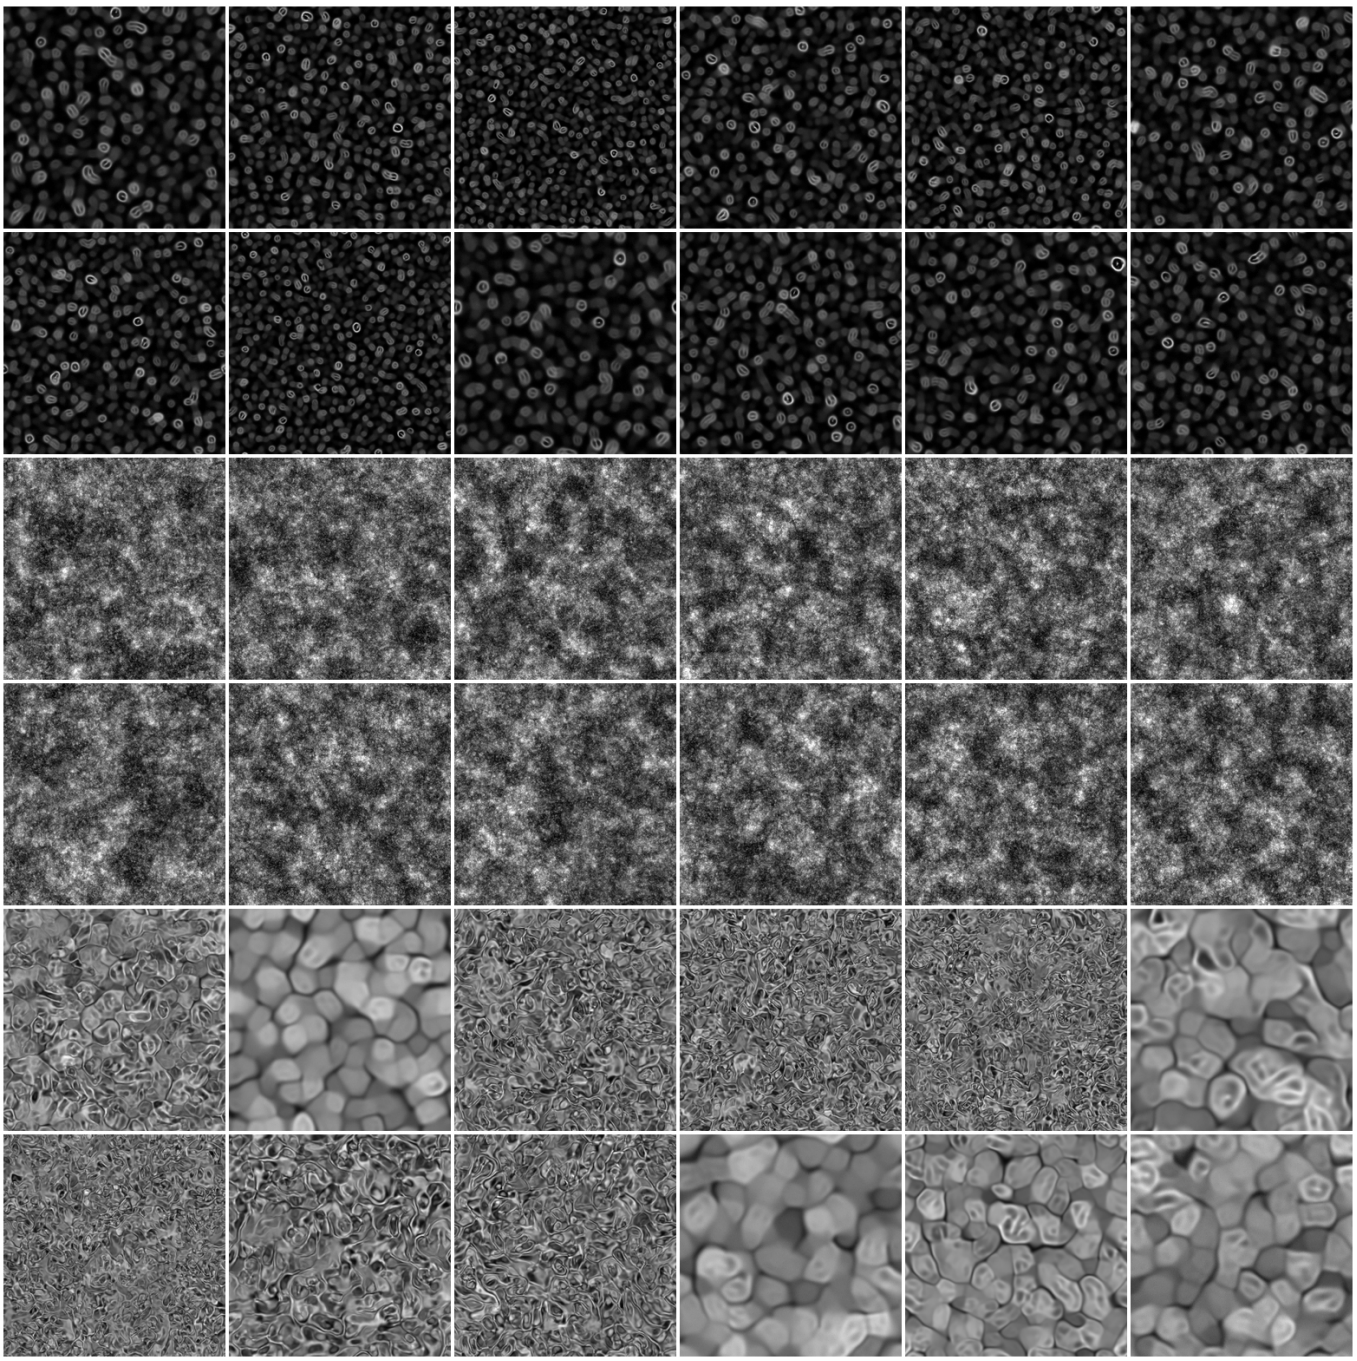

Fig. 20. Random samples of our model's `microscope view`, `bnw spots1`, and `liquid` noises at 256×256 resolution.

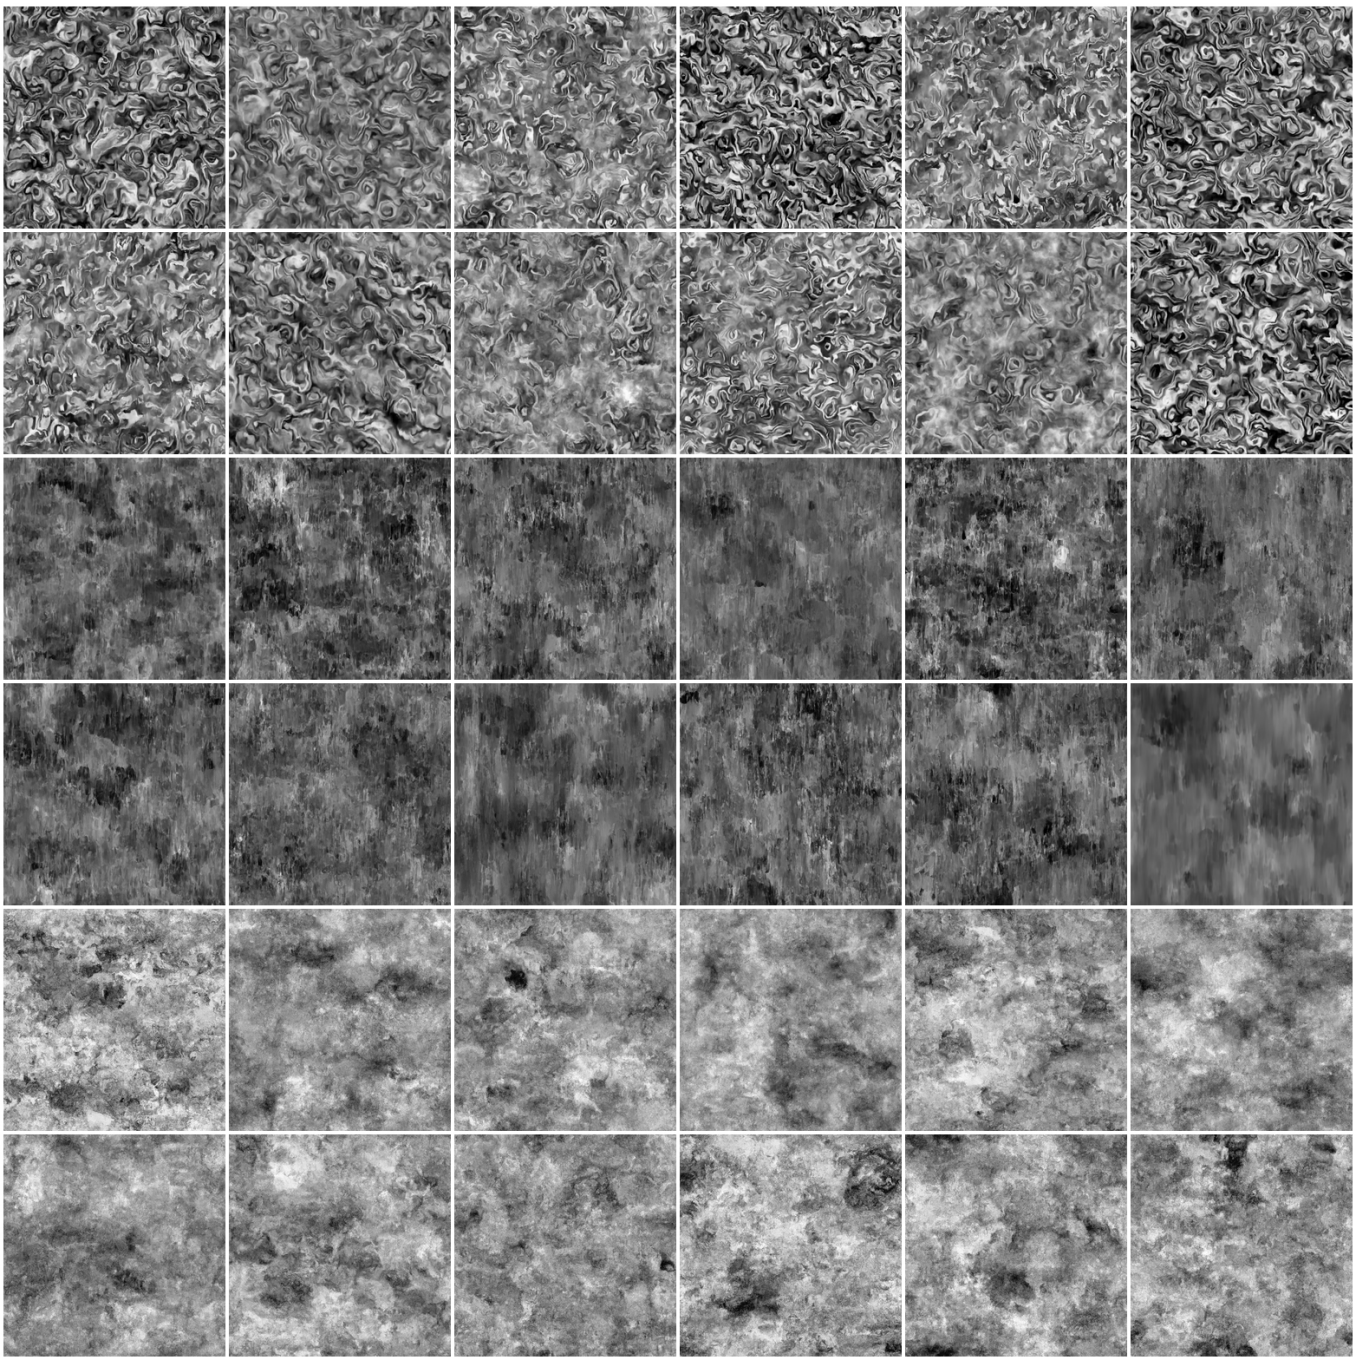

Fig. 21. Random samples of our model's grunge galvanic small, grunge leaky paint, and grunge rust fine noises at 256×256 resolution.

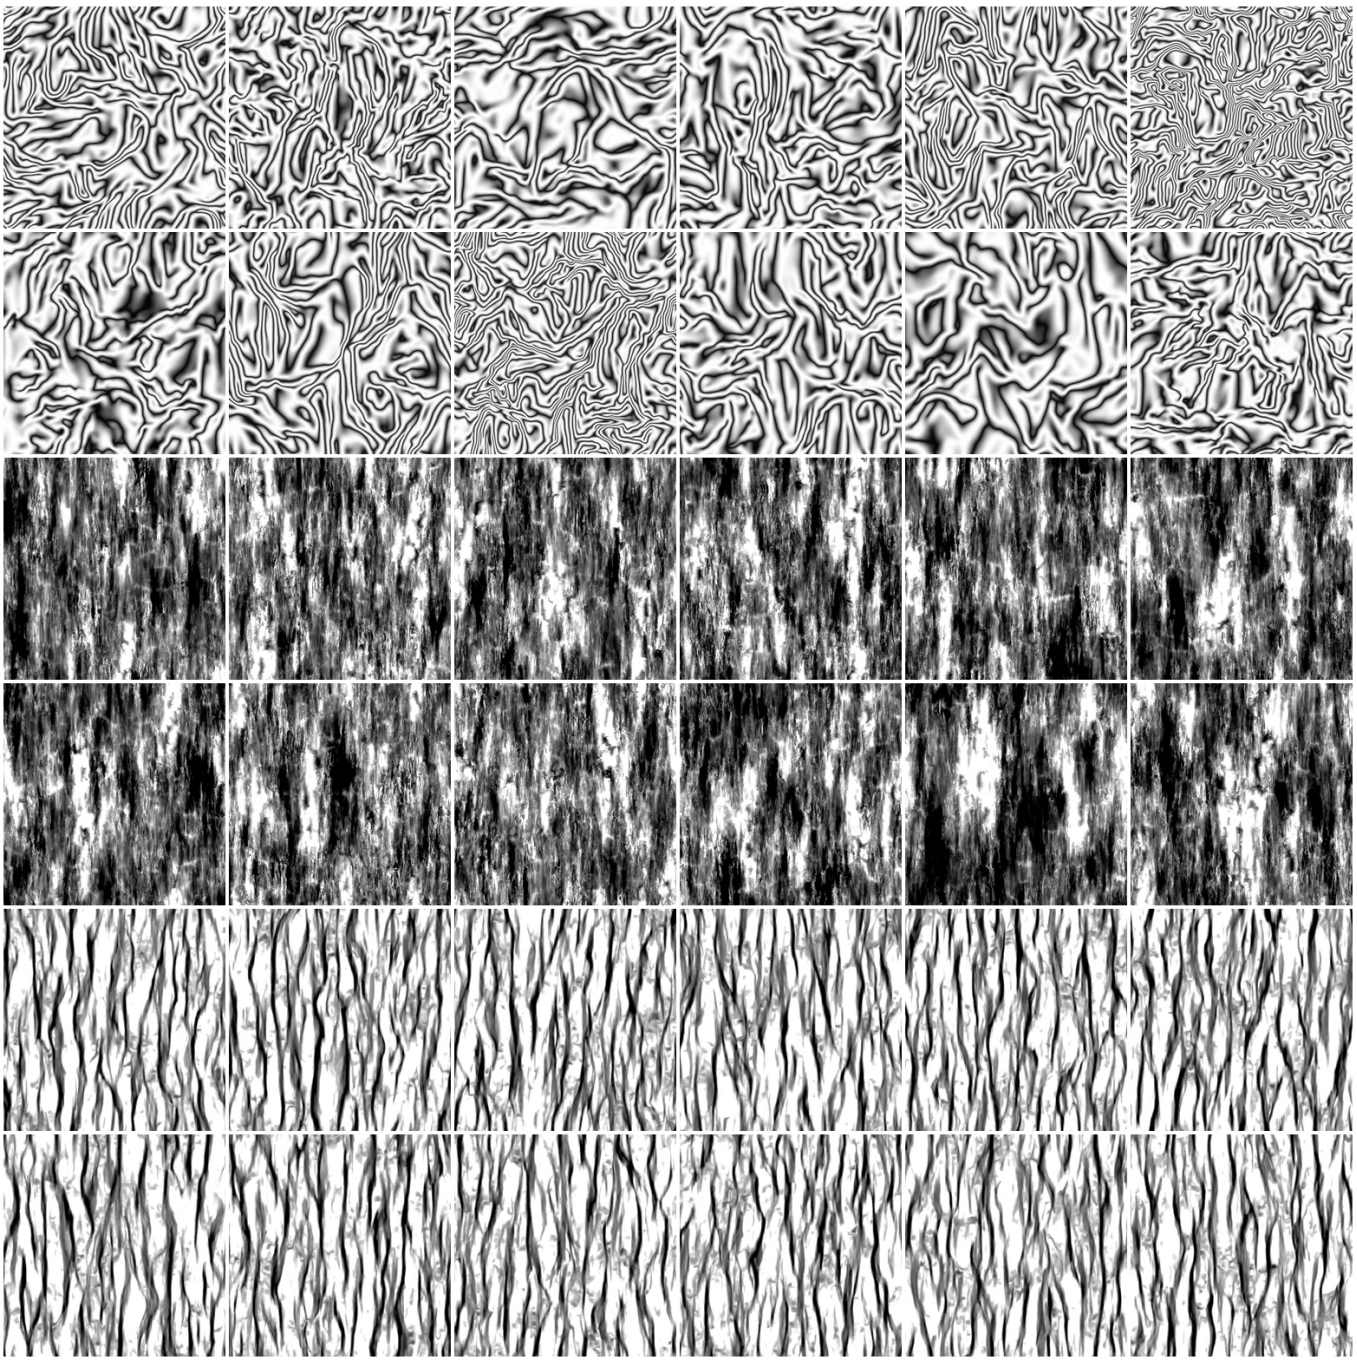

Fig. 22. Random samples of our model's `grunge damas`, `grunge map 002`, and `grunge map 005` noises at 256×256 resolution.

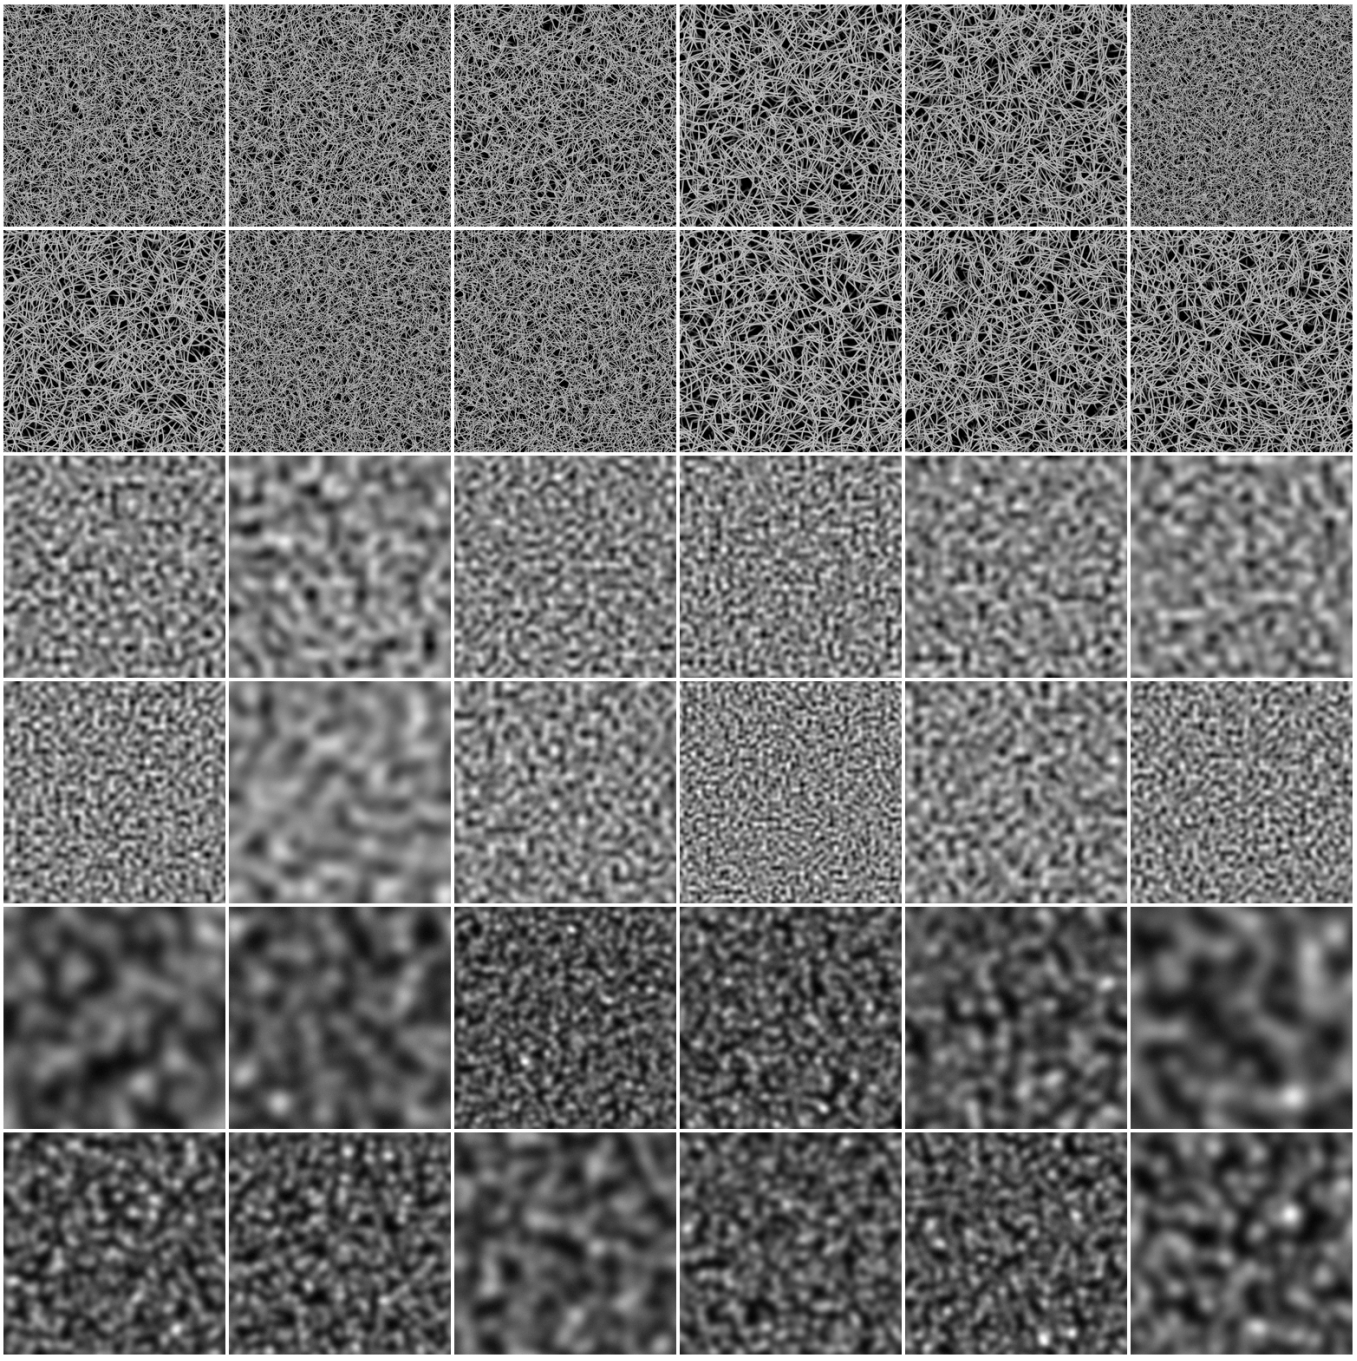

Fig. 23. Random samples of our model's `messy fibers 3`, `perlin`, and `gaussian` noises at 256×256 resolution.

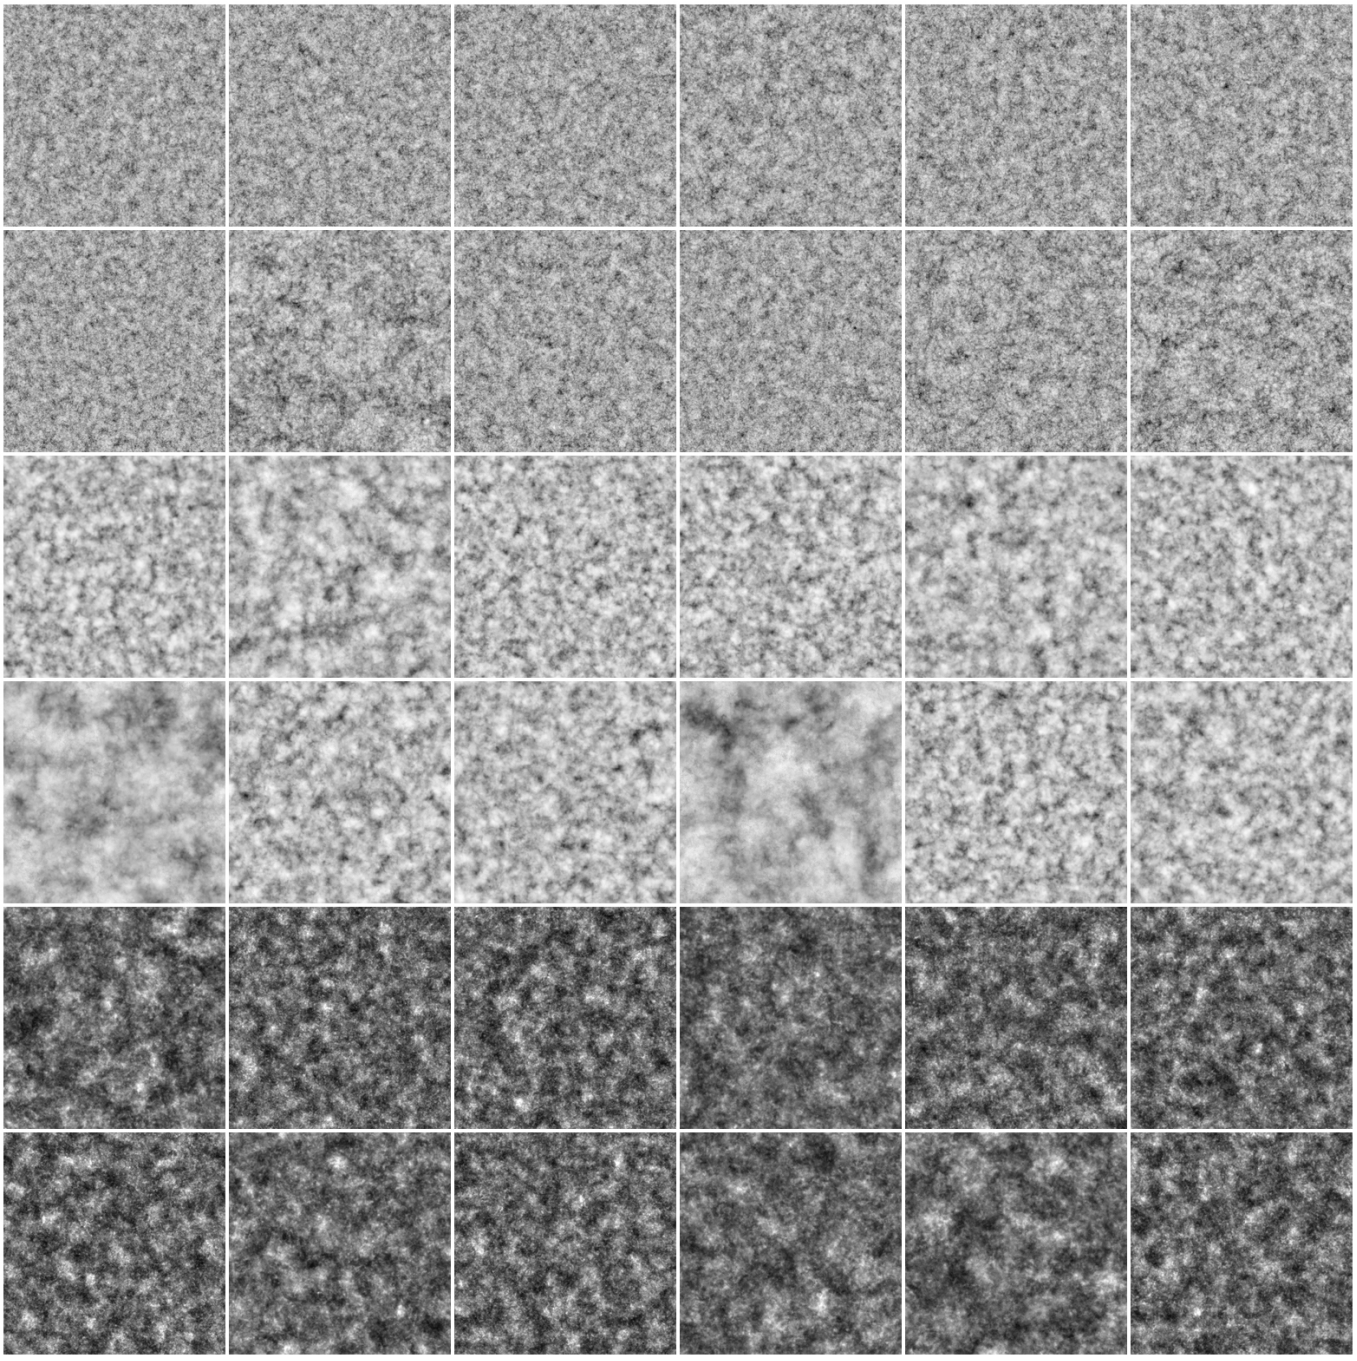

Fig. 24. Random samples of our model's `clouds 1`, `clouds 2`, and `clouds 3` noises at 256×256 resolution.