

Photorealistic Material Editing Through Direct Image Manipulation

KÁROLY ZSOLNAI-FEHÉR, TU Wien

PETER WONKA, KAUST

MICHAEL WIMMER, TU Wien



Fig. 1. We propose a hybrid technique to empower novice users and artists without expertise in photorealistic rendering to create sophisticated material models by applying standard image editing operations to a source image. Then, in the next step, our method proceeds to find a photorealistic BSDF that, when rendered, resembles this target image. Our method generates each of the showcased fits within 20-30 seconds of computation time and is able to offer high-quality results even in the presence of poorly-executed edits (e.g., the background of the gold target image, the gold-colored pedestal for the water material and the stitched specular highlight above it). Scene: Reynante Martinez.

Authors' addresses: Károly Zsolnai-Fehér, TU Wien, Favoritenstrasse 9-11/193-02, Vienna, Austria, 1040, zsolnai@cg.tuwien.ac.at; Peter Wonka, KAUST, Al Khwarizmi Bldg 1, Thuwal, Kingdom of Saudi Arabia, 23955-6900, pwonka@gmail.com; Michael Wimmer, TU Wien, Favoritenstrasse 9-11/193-02, Vienna, Austria, 1040, wimmer@cg.tuwien.ac.at.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2019 Copyright held by the owner/author(s).

XXXX-XXXX/2019/9-ART

<https://doi.org/10.1145/nmnnnnn.nmnnnnn>

Creating photorealistic materials for light transport algorithms requires carefully fine-tuning a set of material properties to achieve a desired artistic effect. This is typically a lengthy process that involves a trained artist with specialized knowledge. In this work, we present a technique that aims to empower novice and intermediate-level users to synthesize high-quality photorealistic materials by only requiring basic image processing knowledge. In the proposed workflow, the user starts with an input image and applies a few intuitive transforms (e.g., colorization, image inpainting) within a 2D image editor of their choice, and in the next step, our technique produces a photorealistic result that approximates this target image. Our method combines the advantages of a neural network-augmented optimizer and an encoder neural network to produce high-quality output results within 30 seconds. We also demonstrate that it is resilient against poorly-edited target

images and propose a simple extension to predict image sequences with a strict time budget of 1-2 seconds per image.

CCS Concepts: • **Computing methodologies** → **Neural networks; Rendering; Ray tracing**;

Additional Key Words and Phrases: neural networks, photorealistic rendering, material modeling, neural rendering

ACM Reference Format:

Károly Zsolnai-Fehér, Peter Wonka, and Michael Wimmer. 2019. Photorealistic Material Editing Through Direct Image Manipulation. 1, 1 (September 2019), 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The expressiveness of photorealistic rendering systems has seen great strides as more sophisticated material models became available for artists to harness. Most modern rendering systems offer a node-based shader tool where the user can connect different kinds of material models and perform arbitrary mathematical operations over them (e.g., addition and mixing), opening up the possibility of building a richer node graph that combines many of the more rudimentary materials to achieve a remarkably expressive model. These are often referred to as “principled” shaders and are commonly used within the motion picture industry [Burley and Studios 2012]. However, this expressiveness comes with the burden of complexity, i.e., the user has to understand each of the many parameters of the shader not only in isolation, but also how they influence each other, which typically requires years of expertise in photorealistic material modeling. In this work, we intend to provide a tool that can be used by a wider target audience, i.e., artists and novices that do not have any experience creating material models, but are adept at general-purpose image processing and editing. This is highly desirable as human thinking is inherently visual and is not based on physically-based material parameters [Röder et al. 2002; White 1989]. We propose a workflow in which the artist starts out with an image of a reference material and applies classic image processing operations to it. Our key observation is that even though this processed target image is often not physically achievable, in many cases, a photorealistic material model can be found that is remarkably close to it (Fig. 2). These material models can then be easily inserted into already existing scenes by the user (Fig. 3).

In summary, we present the following contributions:

- An optimizer that can rapidly match the target image when given an approximate initial guess.
- A neural network to solve the adjoint rendering problem, i.e., take the target image as an input and infer a shader that produces a material model to approximate it.
- A hybrid method that combines the advantages of these two concepts and achieves high-quality results for a variety of cases within 30 seconds.
- A simple extension of our method to enable predicting sequences of images within 1-2 seconds per image.

We provide our pre-trained neural network and the source code for the entirety of this project.

2 PREVIOUS WORK

2.1 Material Acquisition

A common workflow for photorealistic material acquisition requires placing the subject material within a studio setup and using measurement devices to obtain its reflectance properties [Marschner et al. 1999; Miyashita et al. 2016]. To import this measured data into a production renderer, it can be either used as-is, can be compressed down into a lower-dimensional representation [Papas et al. 2013; Rainer et al. 2019] or approximated through an analytic BSDF¹ model [Papas et al. 2014]. Many recent endeavors improve the cost efficiency and convenience of this acquisition step by only requiring photographs of the target material [Aittala et al. 2016, 2015; Deschaintre et al. 2018; Li et al. 2017, 2018] while still requiring physical access to these source material samples, while precomputed BSDF databases offer an enticing alternative where the user can choose from a selection of materials [Dupuy and Jakob 2018; Matusik 2003]. We aim to provide a novel way to exert direct artistic control over these material models. Our method can be related to inverse rendering approaches [Marschner and Greenberg 1998; Ramamoorthi and Hanrahan 2001], where important physical material properties are inferred from a real photograph with unknown lighting conditions. In our work, the material test scene contains a known lighting and geometry setup, but in return, enables not only the rapid discovery of new materials, but artistic control through standard and well-known image-space editing operations.

2.2 Material Editing

To be able to efficiently use the most common photorealistic rendering systems, an artist is typically required to have an understanding of physical quantities pertaining to the most commonly modeled phenomena in light transport, e.g., indices of refraction, scattering and absorption albedos and more [Burley and Studios 2012; Song et al. 2009]. This modeling time can be cut down by techniques that enable editing BRDF² models directly within the scene [Ben-Artzi et al. 2006; Cheslack-Postava et al. 2008; Sun et al. 2007], however, with many of these methods, the artist is still required to understand the physical properties of light transport, often incurring a significant amount of trial and error. Instead of editing the materials directly, other techniques enable editing secondary effects, such as caustics and indirect illumination within the output image [Ben-Artzi et al. 2008; Schmidt et al. 2013]. An efficient material editing workflow also opens up the possibility of rapid relighting previously rendered scenes [Ng et al. 2004; Wang et al. 2008, 2004]. Reducing the expertise required for material editing workflows has been a subject to a large volume of research works: an intuitive editor was proposed by pre-computing many solutions to enable rapid exploration [Hašan and Ramamoorthi 2013], carefully crafted material spaces were derived to aid the artist [Lagunas et al. 2019; Serrano et al. 2016; Soler et al. 2018], and learning algorithms have been proposed to create a latent space that adapts to the preferences of the user [Zsolnai-Fehér et al. 2018]. We endeavored to create a solution that produces the desired results *rapidly* by looking at a non-physical mockup image, requiring expertise only in

¹Bidirectional scattering distribution function.

²Bidirectional reflectance distribution function.

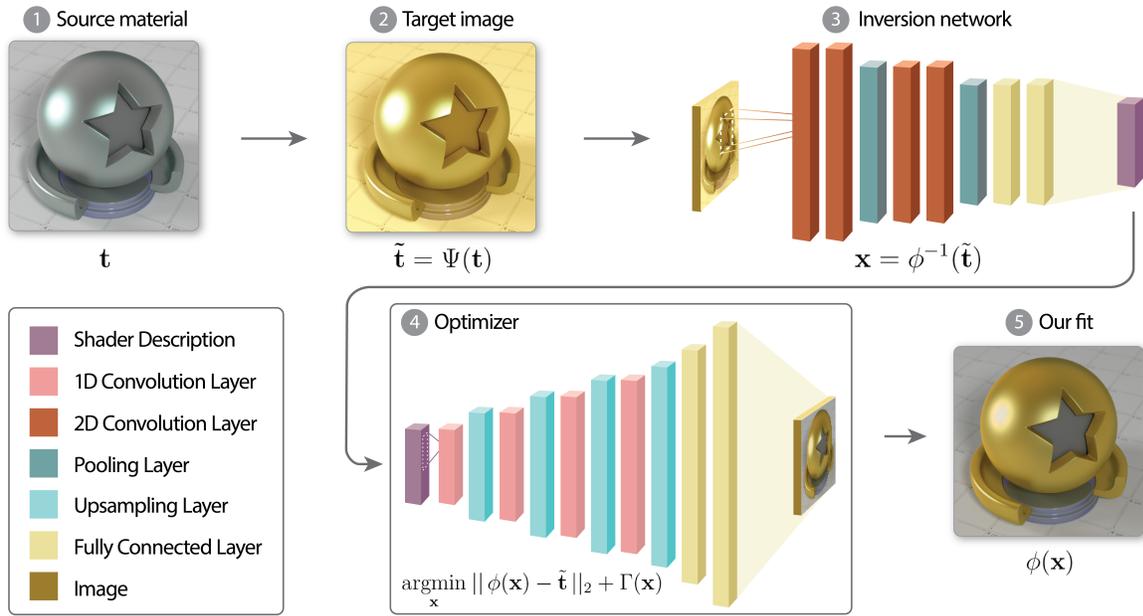


Fig. 2. Our proposed hybrid technique offers an intuitive workflow where the artist takes a source material (1) and produces the target image by applying the desired edits to it within a 2D raster image editor of their choice (2). Then, one or more encoder neural networks are used to propose a set of approximate initial guesses (3) to be used with our neural network-augmented optimizer (4), which rapidly finds a photorealistic shader setup that closely matches the target image (5). The artist then finishes the process by assigning this material to a target object and renders the final scene offline.

2D image editing, which is considered to be common knowledge by nearly all artists in the field. Generally, BRDF relighting methods are preferable when in-scene editing is a requirement, otherwise, we recommend using our proposed technique in the case of one sought material to moderate-scale problems and Gaussian Material Synthesis (GMS) for mass-scale material synthesis.

2.3 Neural Networks and Optimization

Optimization is present at the very core of every modern neural network: to be able to minimize the prescribed loss function efficiently, the weights of the networks are fine-tuned through gradient descent variants [Bottou 2010; Robbins and Monro 1951] or advanced methods that include the use of lower-order moments [Kingma and Ba 2014], while additional measures are often taken to speed up convergence and avoid poor local minima [Goh 2017; Sutskever et al. 2013]. Similar optimization techniques are also used to generate the model description and architecture of these neural networks [Elsken et al. 2018; Zoph and Le 2016], or the problem statement itself can also be turned around by using learning-based methods to discover novel optimization methods [Bello et al. 2017]. In this work, we propose two combinations of a neural network and an optimizer – first, the two can be combined *indirectly* by endowing the optimizer with a reasonable initial guess, and *directly* by using the optimizer that invokes a neural renderer at every function evaluation step to speed up the convergence by several orders of magnitude (steps 3 and 4 in Fig. 2). This results in an efficient two-stage system that is able to rapidly match a non-physical target image and does

not require the user to stay within a prescribed manifold of artistic editing operations [Zhu et al. 2016].

3 OVERVIEW

Many trained artists are adept at creating new photorealistic materials by engaging in a direct interaction with a principled shader. This workflow includes adjusting the parameters of this shader and waiting for a new image to be rendered that showcases the appropriate output material. If at most a handful of materials are sought, this is a reasonably efficient workflow, however, it also incurs a significant amount of rendering time and expertise in material modeling. Our goal is to empower novice and intermediate-level users to be able to reuse their knowledge from image processing and graphic design to create their envisioned photorealistic materials.

In this work, we set up a material test scene that contains a known lighting and geometry setup, and a fixed principled shader with a vector input of $x \in \mathbb{R}^m$ where $m = 19$. This shader is able to represent the most commonly used diffuse, glossy, specular and translucent materials with varying roughness and volumetric absorption coefficients. Each parameter setup of this shader produces a different material model when rendered. In our workflow, the user is offered a variety of images, and chooses one desired material model as a starting point. Then, the user is free to apply a variety of image processing operations on it, e.g., colorization, image inpainting, blurring a subset of the image and more. Since these image processing steps are not grounded in a physically-based framework, the resulting image is not achievable by adjusting the parameters in the vast majority of cases. However, we show that our proposed



Fig. 3. To demonstrate the utility of our system, we synthesized a new material and deployed it into an already existing scene using Blender and Cycles. In this scene, we made a material mixture to achieve a richer and foggier nebula effect inside the glass. Left: theirs, right: 50% theirs, 50% ours. Scene: Reynante Martinez.

method is often able to produce a photorealistic material that closely matches this target image.

Solution by optimization. When given an input image $\mathbf{t} \in \mathbb{R}^p$, it undergoes a series of transformations (e.g., colorization, image inpainting) as the artist produces the target image $\tilde{\mathbf{t}} = \Psi(\mathbf{t})$, where $\Psi: \mathbb{R}^p \rightarrow \mathbb{R}^p$. Then, an image is created from an initial shader configuration, i.e., $\phi: \mathbb{R}^m \rightarrow \mathbb{R}^p$, where m refers to the number of parameters within the shader and p is the number of variables that describe the output image (in our case $p = 3 \cdot 410^2$ is used with the range of 0-255 for each individual pixel). This operation is typically implemented by a global illumination renderer. Our goal is to find an appropriate parameter setup of the principled shader $\mathbf{x} \in \mathbb{R}^m$ that, when rendered, reproduces $\tilde{\mathbf{t}}$. Generally, this is not possible as a typical Ψ leads to images that cannot be perfectly matched through photorealistic rendering. However, surprisingly, we can often find a configuration \mathbf{x} that produces an image that closely resembles $\tilde{\mathbf{t}}$ through solving the minimization problem

$$\begin{aligned} \operatorname{argmin}_{\mathbf{x}} \quad & \|\phi(\mathbf{x}) - \tilde{\mathbf{t}}\|_2, \\ \text{subject to} \quad & \mathbf{x}_{\min} \leq \mathbf{x} \leq \mathbf{x}_{\max}, \end{aligned} \quad (1)$$

where the constraints stipulate that each shader parameter has to reside within the appropriate boundaries (i.e., $0 \leq x_i \leq 1$ for albedos or $x_j \geq 1$ for indices of refraction where $x_i, x_j \in \mathbf{x}$). To be able to benchmark a large selection of optimizers, we introduce

an equivalent alternative formulation of this problem where the constraints are reintroduced as a barrier function $\Gamma(\cdot)$, i.e.,

$$\begin{aligned} \operatorname{argmin}_{\mathbf{x}} \quad & (\|\phi(\mathbf{x}) - \tilde{\mathbf{t}}\|_2 + \Gamma(\mathbf{x})), \text{ where} \\ \Gamma(\mathbf{x}) = \quad & \begin{cases} 0, & \text{if } \mathbf{x} \in C, \\ +\infty, & \text{otherwise,} \end{cases} \\ C = \quad & \{\mathbf{x} \mid f_i(\mathbf{x}) \geq 0, i = 1, 2\}, \\ f_1(\mathbf{x}) = \quad & \mathbf{x}_{\max} - \mathbf{x}, \\ f_2(\mathbf{x}) = \quad & \mathbf{x} - \mathbf{x}_{\min}. \end{aligned} \quad (2)$$

where C denotes the feasible region chosen by a set of constraints described by $f_i(\cdot)$ (equivalent to the second line in (1)) and the vector comparison operator (\geq) here is considered true only when all of the vector elements exceed (or equal to) zero. In a practical implementation, the infinity can be substituted by a sufficiently large integer. This formulation enabled us to compare several optimizers (Table 3 in Appendix B), where we found Nelder and Mead's simplex-based self-adapting optimizer [1965] to be the overall best choice due to its ability to avoid many poor local minima through its contraction operator and used that for each of the reported results throughout this manuscript.

Nonetheless, solving this optimization step still takes several hours as each function evaluation invokes ϕ , i.e., a rendering step to produce an image, which clearly takes too long for day-to-day use in the industry. We introduce two solutions to remedy this limitation, followed by a hybrid method that combines their advantages.

Neural renderer. To speed up the function evaluation process, we replace the global illumination engine that implements ϕ with a neural renderer [Zsolnai-Fehér et al. 2018]. This way, instead of running a photorealistic rendering program at each step, our optimizer invokes the neural network to predict this image, thus reducing the execution time of the process by several orders of magnitude, in our case, from an average of 50 seconds to 4ms per image at the cost of restricting the material editing to a prescribed scene and lighting setup. Because of the lack of a useful initial guess, this solution still requires many function evaluations and is unable to reliably provide satisfactory solutions.

Solution by inversion. One of our key observations is that an approximate solution can also be produced *without* an optimization step by finding an appropriate inverse to ϕ : since ϕ is realized through a decoder neural network (i.e., neural renderer) that produces an image from a shader configuration, ϕ^{-1} , its inverse, can be implemented as an *encoder* network that takes an image as an input and predicts the appropriate shader parameter setup that generates this image. This adjoint problem has several advantages: first, such a neural network can be trained on the same dataset as ϕ by only swapping the inputs and outputs and retains the advantageous properties of this dataset, e.g., arbitrarily many new training samples can be generated via rendering, thereby loosening the ever-present requirement of preventing overfitting via regularization [Nowlan

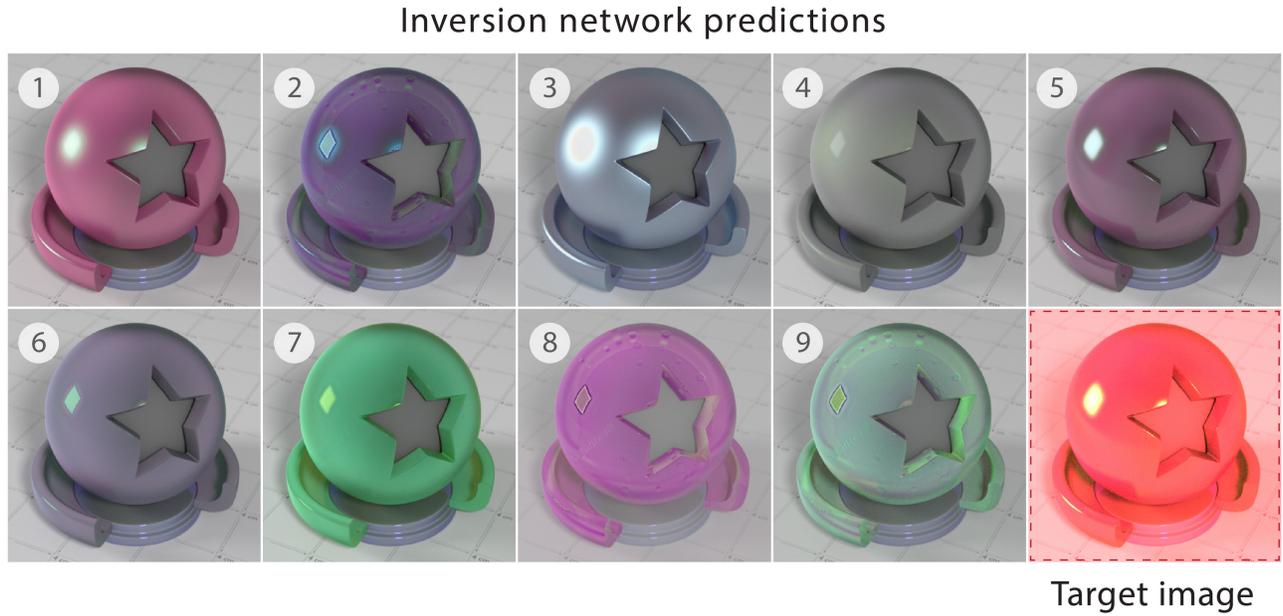


Fig. 4. Whenever the target image strays too far away from the images contained within their training set (lower right), our 9 inversion networks typically fail to provide an adequate solution and potentially predict results outside the feasible region (❷, ❸, ❹ - we discuss these cases in the supplementary materials in more detail). However, using our “best of n” scheme and our hybrid method, the best performing prediction of our neural networks can be used to equip our optimizer with an initial guess, substantially improving its results.

and Hinton 1992; Srivastava et al. 2014; Zou and Hastie 2005]. Second, we can use it to find a solution *directly* through $\mathbf{x} \approx \phi^{-1}(\tilde{\mathbf{t}})$ without performing the optimization step described in (1-2). As the output image is not produced through a lengthy optimization step, but is inferred by this encoder network, this computes in a few milliseconds. We will refer to this solution as the *inversion network* and note that our implementation of ϕ^{-1} only approximately admits the mathematical properties of a true inverse function. We also discuss the nature of the differences in more detail in Section 4. We have trained 9 different inversion network architectures and found that typically, each of them performs well on a disjoint set of inputs. Our other key observation is that because we have an atypical problem where the ground truth image ($\tilde{\mathbf{t}}$) is available and each of the candidate images can be inferred inexpensively (typically within 5 milliseconds), it is possible to compute a “best of n” solution by comparing all of these predictions to the ground truth, i.e.,

$$\mathbf{x} = \phi_{(i)}^{-1}(\tilde{\mathbf{t}}), \text{ where } i = \underset{j}{\operatorname{argmin}} \|\phi(\phi_{(j)}^{-1}(\tilde{\mathbf{t}})) - \tilde{\mathbf{t}}\|_2, \quad (3)$$

where $\phi_{(i)}^{-1}$ denotes the prediction of the i -th inversion network, $j = (1, \dots, n)$, and in our case, $n=9$ was used. This step introduces a negligible execution time increase and in return, drastically improves the quality of this inversion process for a variety of test cases. However, these solutions are only approximate in cases where the target image strays too far away from the training data (Fig. 4). In Appendix A we report the structure of the neural networks used in this figure.

Hybrid solution. Both of our previous solutions suffer from drawbacks: the optimization approach provides results that resemble $\tilde{\mathbf{t}}$ but is impracticable due to the fact that it requires too many function evaluations and gets stuck in local minima, whereas the inversion networks rapidly produce a solution, but offer no guarantees when the target image significantly differs from the ones shown in the training set. We propose a hybrid solution based on the knowledge that even though the inverse approach does not provide a perfect solution, since it can produce results instantaneously that are significantly closer to the optimum than a random input, it can be used to endow the optimizer with a reasonable initial guess. This method is introduced as a variant of (2) where $\mathbf{x}_{\text{init}} = \phi^{-1}(\tilde{\mathbf{t}})$ and a more detailed description of this hybrid solution is given below in Algorithm 1. Additionally, this technique is able to not only provide a “headstart” over the standard optimization approach but was also able to find higher quality solutions in all of our test cases.

Predicting image sequences. A typical image editing workflow takes place within a raster graphics editor program where the artist endeavors to find an optimal set of parameters, e.g., the kernel width σ in the case of a Gaussian blur operation to obtain their envisioned artistic effect. This process includes a non-trivial amount of trial and error where the artist decides whether the parameters should be increased or decreased; this is only possible in the presence of near-instant visual feedback that reflects the effect of the parameter

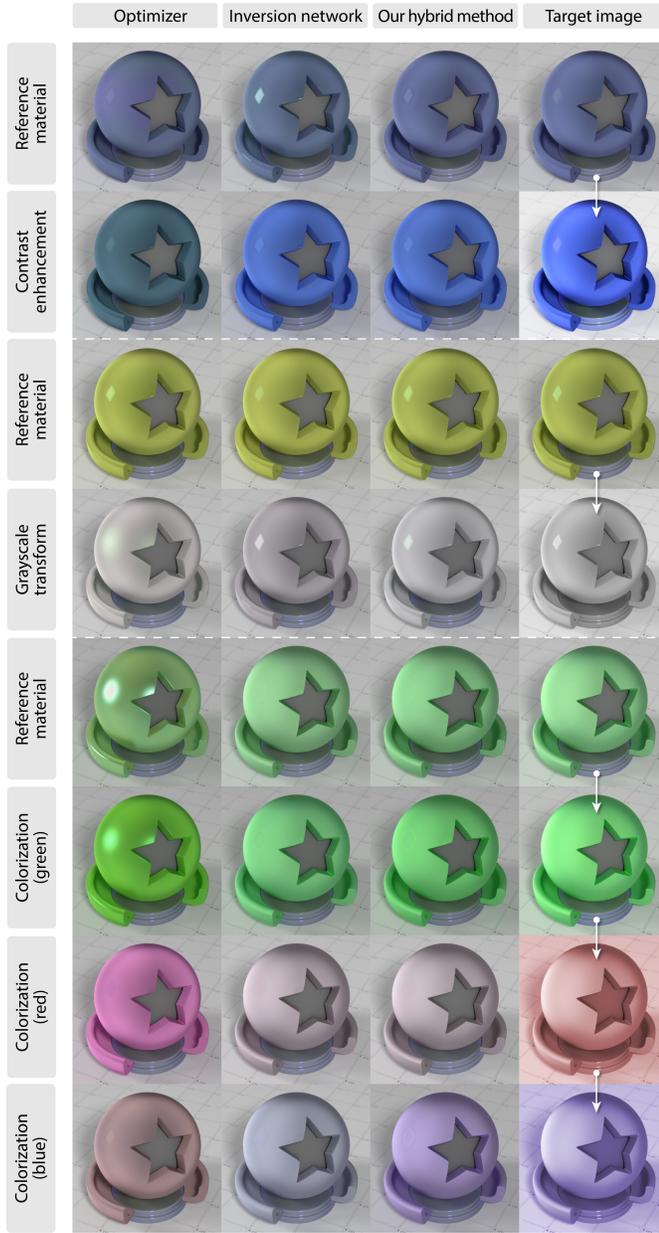


Fig. 5. Results for three techniques on common global colorization operations including saturation increase and grayscale transform. The “reference material” labels showcase materials that can be obtained using our shader and are used as source images for the materials below them, where the arrows denote the evolution of the target image.

changes on the image. We propose a simple extension to our hybrid method to accommodate these workflows: consider an example scenario where the k -th target image in a series of target images $\tilde{\mathbf{t}}_{(k)}$ are produced by subjecting a starting image \mathbf{t} to an increasingly wide blurring kernel. This operation is denoted by $\Psi_{\sigma}(\mathbf{t}) = G_{\sigma} * \mathbf{t}$, where G_{σ} is a zero-centered Gaussian, and for simplicity, the target

Algorithm 1 Photorealistic Material Editing

- 1: **Given** \mathbf{t} , $\phi(\cdot)$, $[\phi_{(1)}^{-1}(\cdot), \dots, \phi_{(n)}^{-1}(\cdot)]$, \mathbf{x}_{\min} , \mathbf{x}_{\max}
- 2: $\tilde{\mathbf{t}} \leftarrow \Psi(\mathbf{t})$ ▷ Obtain target image
- 3: **for** $i \leftarrow 1$ to n **do** ▷ Predict with n inversion networks
- 4: Compute each $\phi_{(i)}^{-1}(\tilde{\mathbf{t}})$
- 5: **Find** $i = \operatorname{argmin}_{j \in 1..n} \|\phi(\phi_{(j)}^{-1}(\tilde{\mathbf{t}})) - \tilde{\mathbf{t}}\|_2$ ▷ Find best candidate
- 6: Define $\mathbf{x}_{\text{init}} \leftarrow \phi_{(i)}^{-1}(\tilde{\mathbf{t}})$
- 7: Define $f_1(\mathbf{x}) = \mathbf{x}_{\max} - \mathbf{x}$ ▷ Set up constraints
- 8: Define $f_2(\mathbf{x}) = \mathbf{x} - \mathbf{x}_{\min}$
- 9: Define $C = \{\mathbf{x} \mid f_i(\mathbf{x}) \geq 0, i = 1, 2\}$ ▷ Construct feasible region
- 10: Define $\Gamma(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{x} \in C, \\ +\infty, & \text{otherwise} \end{cases}$ ▷ Construct barrier
- 11: **Initialize** optimizer with \mathbf{x}_{init}
- 12: **Minimize** $\operatorname{argmin}_{\mathbf{x}} (\|\phi(\mathbf{x}) - \tilde{\mathbf{t}}\|_2 + \Gamma(\mathbf{x}))$ ▷ Refine initial guess
- 13: Display $\phi(\mathbf{x})$ to user

images are produced via $\tilde{\mathbf{t}}_{(k)} = \Psi_k(\mathbf{t})$, with the initial condition of $\tilde{\mathbf{t}}_{(0)} = \mathbf{t}$. We note that many other transforms can also be substituted in the place of Ψ without loss of generality. We observe that such workflows create a series of images where each neighboring image pair shows only minute differences, i.e., for any positive non-zero k , $\|\tilde{\mathbf{t}}_{(k+1)} - \tilde{\mathbf{t}}_{(k)}\|_2$ remains small. As in these cases, we are required to propose many output images, we can take advantage of this favorable mathematical property by extending the pool of initial inversion networks with the optimized result of the previous frame by modifying Steps 3-5 of Algorithm 1 to add

$$\phi_{(n+1)}^{-1}(\tilde{\mathbf{t}}_k) = \operatorname{argmin}_{\mathbf{x}} (\|\phi(\mathbf{x}) - \tilde{\mathbf{t}}_{k-1}\|_2 + \Gamma(\mathbf{x})). \quad (4)$$

Note that this does not require any extra computation as the result of Step 12 of the previous run can be stored and reused. Intuitively, this means that *both* the inversion network predictions and the prediction of the previous image are used as candidates for the optimization (whichever is better). This way, after the optimization step is finished, the improvements can be “carried over” to the next frame. This method we refer to as *reinitialization* and in Section 4, we show that it consistently improves the quality of our output images for such image sequences, even with a strict budget of 1-2 seconds per image.

4 RESULTS

In this section, we discuss the properties of our inverse problem formulation (i.e., inferring a shader setup that produces a prescribed input image), followed by both a quantitative and qualitative evaluation of our proposed hybrid method against the optimization and inversion network solutions. We also show that our system supports a wide variety of image editing operations and can rapidly predict image sequences. To ensure clarity, we briefly revisit the three introduced methods:

Input	Initial guess		50 fun. evals		300 fun. evals		1500 fun. evals	
	Random	NN	Optimizer	Ours	Optimizer	Ours	Optimizer	Ours
Fig. 5, Row 1	41.93	5.94	33.81	4.53	9.42	2.84	5.62	2.37
Fig. 5, Row 2	78.45	32.72	68.55	32.67	40.24	32.67	40.21	32.67
Fig. 5, Row 4	35.37	18.68	30.88	16.53	17.29	14.71	16.98	14.68
Fig. 5, Row 7	41.65	22.42	38.10	22.38	26.30	22.38	26.24	22.38
Fig. 5, Row 8	29.04	19.82	26.79	18.43	22.93	15.37	22.93	15.37
Fig. 8, Row 2	23.78	12.79	20.31	11.62	8.27	7.81	8.26	7.80
Fig. 8, Row 3	21.60	9.09	16.54	8.28	6.24	5.80	6.19	5.80
Fig. 8, Row 8	29.58	9.74	22.69	7.92	6.63	5.36	6.63	5.36

Table 1. A comparison of the optimization approach (with random initialization) and our hybrid method (with “best of 9” NN initialization) on a variety of challenging global and local image editing operations in Fig. 5 and 8. The numbers indicate the RMSE of the outputs, and for reference, the first row showcases an input image that is reproducible by the shader.

F. evals	Technique	Image ID in sequence (i.e., k of $\hat{\mathbf{t}}_{(k)}$)														Σ
		0	10	20	30	40	50	60	70	80	90	100	110	120		
100	No reinitialization	1.93	1.67	2.19	2.90	3.82	4.79	5.73	6.81	7.93	9.14	10.43	11.55	12.99	81.88	
	Reinitialization	1.93	1.34	1.88	2.54	3.34	4.30	5.30	6.38	7.50	8.69	9.93	11.55	12.99	77.67	
300	No reinitialization	1.64	1.47	2.07	2.80	3.70	4.62	5.70	6.75	7.86	9.00	10.21	11.41	12.82	80.05	
	Reinitialization	1.64	1.30	1.80	2.42	3.25	4.25	5.25	6.33	7.45	8.64	9.88	11.41	12.82	76.44	
600	No reinitialization	1.57	1.44	2.06	2.77	3.66	4.60	5.69	6.74	7.83	8.96	10.12	11.41	12.80	79.65	
	Reinitialization	1.57	1.29	1.80	2.49	3.33	4.20	5.18	6.27	7.38	8.58	9.81	11.41	12.80	76.11	

Table 2. Our proposed reinitialization technique consistently outperforms per-frame computation for the image sequence shown in Fig. 6. The numbers indicate the RMSE of the outputs.

- The **optimization** approach relies on minimizing (2) with Nelder and Mead’s simplex method using a random initial guess, and implementing ϕ through a neural renderer,
- the **inversion network** refers to the “best of 9” inversion solution, i.e., $\mathbf{x} \approx \phi_{(i)}^{-1}(\hat{\mathbf{t}})$ as shown in (3),
- our **hybrid method** is obtained by combining the two above approaches as described in Algorithm 1.

Furthermore, in Appendix A, we report the structure of the neural networks used to implement each individual $\phi_{(i)}^{-1}$ shown in Fig. 4, and compare our solution to a selection of local and global minimizers in Appendix B. At the end of this section, we also compare the total time taken to synthesize 1, 10, and 100 selected materials against a recent method for mass-scale material synthesis.

Inversion accuracy. Our inversion technique leads to an approximate solution within a few milliseconds, however, because the structure of the forward and inverse networks differ, the inversion operation remains imperfect, especially when presented with a target image that includes materials that are only approximately achievable. To demonstrate this effect, we have trained 9 different inversion networks to implement ϕ^{-1} and show that none of the proposed solutions are satisfactory as a final output for the global colorization case (Fig. 4). Our goal with this experiment was to demonstrate that a solution containing only one inversion network generally produces unsatisfactory outputs, regardless of network

structure. However, these predictions can be used to equip our optimizer with an initial guess, substantially improving its results. As each neural network consumes between 300MB and 1GB of video memory, we were able to keep all of them loaded during the entirety of the work session.

Optimizer and hybrid solution accuracy. In Table 1, we compared our hybrid solution against the “best of 9” inversion network and optimization approaches and recorded the RMS error after 50, 300 and 1500 function evaluations (these roughly translate to 1, 6, and 30-second execution times) to showcase the early and late-stage performance of these methods. The table contains a selection of scenarios that we consider to be the most challenging and note that the outputs showed no meaningful change after 1500 function evaluations. Our hybrid method produced the lowest errors in each of our test cases, and surprisingly, the inversion network initialization not only provides a “headstart” for our method, but also improves the final quality of the output, thereby helping the optimizer to avoid local minima. To validate the viability of our solutions, we also ran a global minimizer [Wales and Doye 1997] with several different parameter choices and a generous allowance of 30 minutes of computation time for each; our hybrid method was often able to match (and in some cases, surpass) the quality offered by this solution (Appendix B, Table 3), further reinforcing how our inversion network initialization step helps avoid getting stuck in poor

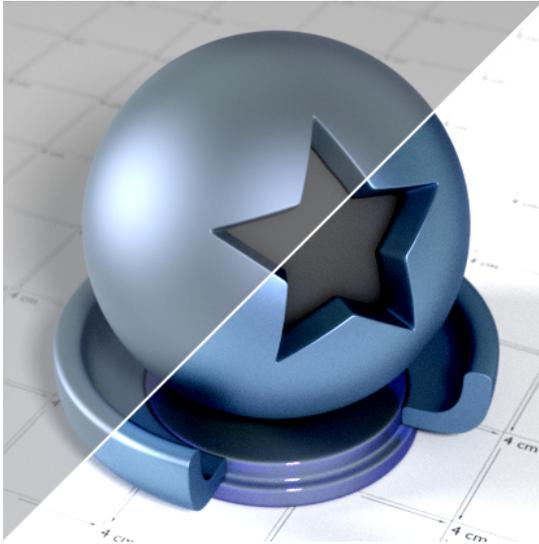


Fig. 6. Our image sequence starts with an input that is achievable using our shader (upper left), where each animation frame slightly increases its black levels. The lower right region showcases the 300th frame of the animation.

local minima. Note that the optimizer was unable to meaningfully improve the best prediction of the 9 inversion networks in Fig. 5, Row 7 – in this case, a better solution can be found by using the prediction of only the first neural network and passing it to the optimizer, improving the reported RMSE from 22.38 to 19.39 by using 300 function evaluations.

Supported image editing operations. A typical workflow using our technique includes the artist choosing a source material and applying an appropriate image editing operation (Ψ) instead of engaging in a direct interaction with the principled shader. We cluster the set of possible transforms into *global* (Fig. 5) and *local* (Fig. 8) operations: these cases include saturation increase, grayscale transform, colorization, image mixing, stitching and inpainting, and selective blurring of highlights. Both the optimizer and our hybrid method were run for 1500 function evaluations to obtain the results showcased in these two figures. As these transformations come from a 2D raster editor and are not grounded in a physically-based framework, a perfect match is often not possible, however, in each of these cases, our hybrid method proposed a solution of equivalent or better quality compared to the “best of 9” inversion network and the optimizer solutions.

Image sequence prediction. As our earlier results in Table 1 revealed that the global colorization techniques typically prove to be among the more difficult cases, we have created a challenging image sequence with an input image that is achievable with our shader, and subjected it to a slight black level increase over many frames (Fig. 6). Every image within this sequence is reproduced both with independent per-frame inference and our reinitialization technique with a strict time budget of 2, 6, and 12 seconds per image (100, 300, and 600 function evaluations). In Table 2, we show that this simple extension successfully exploits the advantageous mathematical properties of these workflows and consistently reduces

the output error for the majority of the sequence, i.e., images 1-100. We also report the RMSE of images 101-120 for reference, which we refer to as the “converged” regime in which the target images stray further and further away from the feasible domain, and the proposed solution remains the same despite these changes. Even in these cases, our reinitialization technique performs no worse than the “no reinitialization” method, and because of its negligible additional cost, we consider it to be a strictly better solution.

Modeling and execution time. In Fig. 7, we have recorded the modeling times for 1, 10, and a 100 similar materials using our method and compared them against Gaussian Material Synthesis [Zsolnai-Fehér et al. 2018] (GMS), a learning-based technique for mass-scale material synthesis. We briefly describe the most important parameters of the recorded execution times and refer the interested reader to this paper for more details – the novice and expert user timings were taken from the GMS paper and indicate the amount of time these users took to create the prescribed number of materials by hand using Disney’s “principled” shader [Burley and Studios 2012], whereas GMS and our timings contain both the modeling (i.e., scoring a material gallery in GMS and performing image processing for our technique) and execution times. If only one material is desired, our technique outperforms this previous work and nearly matches the efficiency of an expert user. When 10 materials are sought (1 base material and 9 variants), our proposed method was adapted to use the re-initialization technique and offers the best modeling times, outperforming both GMS and expert users. In the case of mass-scale material synthesis, i.e., 100 or more materials, both methods outperform experts, where GMS offers the best scaling solution. In each case, the timings for our technique include the fixed cost of loading the 9 neural networks (5.5s). Throughout this manuscript, all results were generated using a NVIDIA TITAN RTX GPU.

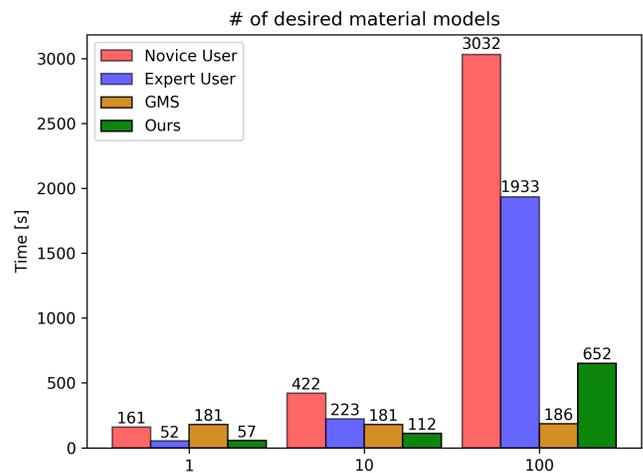


Fig. 7. The recorded modeling times reveal that if at most a handful (i.e., 1-10) of target materials are sought, our technique offers a favorable entry point for novice users into the world of photorealistic material synthesis.



Fig. 8. Results for three techniques on local image editing operations and image mixing. The “reference material” labels showcase materials that can be obtained using our shader and are used as source images for the materials below them, where the arrows denote the evolution of the target image.

5 LIMITATIONS AND FUTURE WORK

As demonstrated in Fig. 4, the results of ϕ^{-1} depend greatly on the performance of the encoder and decoder neural networks. As these methods enjoy significant research attention, we encourage further experiments in including these advances to improve them (e.g., architecture search [Real et al. 2017], capsule networks [Hinton et al. 2018; Sabour et al. 2017] and skip connections [Mao et al.

2016] among many other notable works) and adapting other neural network architectures to our problem that are more tailored to solve inverse problems [Ardizzone et al. 2018; Mataev et al. 2019]. Furthermore, strongly localized edits, e.g., blurring a small part of a specular highlight typically introduces drastic changes within only a small subset of the image and represent only a small fraction of the RMSE calculations and thus may not get proper prioritization from the optimizer. To alleviate this, the relative importance of different regions may also be controlled via weighted masks to emphasize these edits, making these edited regions “score higher” in the error metric, offering the user more granular artistic control. In specialized cases, our reinitialization technique may prove to be useful for single images by using the parameter set used to produce \mathbf{t} as an initial guess for $\hat{\mathbf{t}}$. In-scene editing still remains the key advantage of BRDF relighting techniques.

We also note that our learning technique assumes an input shader of dimensionality m and a renderer that is able to produce images of the materials that it encodes. In this work, our principled shader was meant to demonstrate the utility of this approach by showcasing intuitive workflows with the most commonly used BSDFs. However, this method needs not to be restricted to a classic principled BSDF, and is also expected to perform well on a rich selection of more specialized material models including thin-film interference [Dias 1991; Ikeda et al. 2015], fluorescence [Wilkie et al. 2001] birefringence [Weidlich and Wilkie 2008], microfacet models [Heitz et al. 2016] layered materials [Belcour 2018; Zeltner and Jakob 2018], and more.

6 CONCLUSIONS

We have presented a hybrid technique to empower novice users and artists without expertise in photorealistic rendering to create sophisticated material models by applying image editing operations to a source image. The resulting images are typically not achievable through photorealistic rendering, however, in many cases, solutions be found that are close to the desired output. Our learning-based technique is able to take such an edited image and propose a photorealistic material setup that produces a similar output, and provides high-quality results even in the presence of poorly-edited images. Our proposed method produces a reasonable initial guess and uses a neural network-augmented optimizer to fine-tune the parameters until the target image is matched as closely as possible. This hybrid method is simple, robust, and its computation time is within 30 seconds for every test case showcased throughout this paper. This low computation time is beneficial especially in the early phases of the material design process where a rapid iteration over a variety of competing ideas is an important requirement (Fig 9). Our two key insights can be summarized as follows:

- Normally, using an input image that was generated by a principled shader is not useful given that the user has to generate this image themselves with a known parameter setup. However, our main idea is that the user can subject this image to raster editing operations and “pretend” that this input is achievable, and reliably infer a shader setup to mimic it.

- Our neural networks can be combined with optimizers both *directly*, i.e., by using an optimizer that invokes a neural renderer at every function evaluation step to speed up the convergence and *indirectly* by using a set of neural networks network to endow the optimizer with a reasonable initial guess (steps ③ and ④ in Fig. 2). This combination results in a two-stage system that opens up efficient material editing workflows for artists without expertise in this area.

Furthermore, we proposed a simple extension to support predicting image sequences with a strict time budget of 1-2 seconds and believe this method will offer an appealing entry point for novices into world of photorealistic material modeling.

ACKNOWLEDGMENTS

We would like to thank Reynante Martinez for providing us the geometry and some of the materials for the Paradigm (Fig. 1) and Genesis scenes (Fig. 3), ianofshields for the Liquify scene that served as a basis for Fig. 9, Robin Marin for the material test scene, Andrew Price and Gábor Mészáros for their help with geometry modeling, Felícia Zsolnai-Fehér for her help improving our figures, Christian Freude, David Ha, Philipp Erler and Adam Celarek for their useful comments. We also thank NVIDIA for providing the hardware to train our neural networks. This work was partially funded by Austrian Science Fund (FWF), project number P27974.

A NEURAL NETWORK ARCHITECTURES

Below, we describe the neural network architectures we used to implement $\phi_{(i)}^{-1}$. The Conv2D notation represents a 2D convolutional layer with the appropriate *number of filters*, *spatial kernel sizes* and *strides*, where FC represents a dense, fully-connected layer with a prescribed number of *neurons* and *dropout probability*.

- (1) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 1x{Conv2D(64,3,1), MaxPool(2,2)} – 2x{Conv2D(128,3,1), MaxPool(2,2)} – 2x{FC(1000, 0.1)} - FC(m, 0.0)
- (2) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 2x{FC(1000, 0.1)} - FC(m, 0.0)
- (3) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 2x{FC(1000, 0.5)} - FC(m, 0.0)
- (4) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 1x{Conv2D(64,3,1), MaxPool(2,2)} – 2x{Conv2D(128,3,1), MaxPool(2,2)} – 2x{FC(3000, 0.5)} - FC(m, 0.0)
- (5) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 1x{Conv2D(64,3,1), MaxPool(2,2)} – 2x{Conv2D(128,3,1), MaxPool(2,2)} – 2x{FC(3000, 0.0)} - FC(m, 0.0)
- (6) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 2x{FC(1000, 0.0)} - FC(m, 0.0)
- (7) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 2x{FC(1000, 0.0)} - FC(m, 0.0)
- (8) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 2x{FC(100, 0.0)} - FC(m, 0.0)
- (9) 2x{Conv2D(32,3,1), MaxPool(2,2)} – 2x{FC(1000, 0.0)} - FC(m, 0.0)

Neural networks 6,7 and 9 are isomorphic and were run for a different number of epochs to test the effect of overfitting later in the training process, and therefore offer differing validation losses. The implementation of ϕ is equivalent to the one used in Zsolnai-Fehér et al.'s work [2018].

B COMPARISON OF OPTIMIZERS

In Table 3, we have benchmarked several optimizers, i.e., L-BFGS-B [Byrd et al. 1995], SLSQP [Kraft 1994], the Conjugate Gradient method [Hestenes and Stiefel 1952] and found Nelder and Mead's simplex-based self-adapting optimizer [1965] to be the overall best choice for our global and local image-editing operations. For reference, we also ran Basin-hopping [Wales and Doye 1997], a global minimizer with a variety of parameter choices and a generous allowance of 30 minutes of execution time for each test case. This method is useful for challenging non-linear optimization problems with high-dimensional search spaces. Note that when being run for long enough, this technique is less sensitive to initialization due to the fact that it performs many quick runs from different starting points, and hence, we report one result for both initialization techniques. The cells in the intersection of “Nelder-Mead” and “NN” denote our proposed hybrid method, which was often able to match, and in some cases, outperform this global minimization technique.

REFERENCES

- Miika Aittala, Timo Aila, and Jaakko Lehtinen. 2016. Reflectance modeling by neural texture synthesis. *ACM Transactions on Graphics* 35, 4 (2016), 65.
- Miika Aittala, Tim Weyrich, Jaakko Lehtinen, et al. 2015. Two-shot SVBRDF capture for stationary materials. *ACM Transactions on Graphics* 34, 4 (2015), 110–1.
- Lynton Ardizzone, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W Pellegrini, Ralf S Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. 2018. Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730* (2018).
- Laurent Belcour. 2018. Efficient Rendering of Layered Materials using an Atomic Decomposition with Statistical Operators. *ACM Transactions on Graphics* 37, 4 (2018), 1. <https://doi.org/10.1145/3197517.3201289>
- Irwan Bello, Barret Zoph, Vijay Vasudevan, and Quoc V Le. 2017. Neural optimizer search with reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 459–468.
- Aner Ben-Artzi, Kevin Egan, Frédo Durand, and Ravi Ramamoorthi. 2008. A precomputed polynomial representation for interactive BRDF editing with global illumination. *ACM Transactions on Graphics (TOG)* 27, 2 (2008), 13.
- Aner Ben-Artzi, Ryan Overbeck, and Ravi Ramamoorthi. 2006. Real-time BRDF editing in complex lighting. In *ACM Transactions on Graphics*, Vol. 25. ACM, 945–954.
- Léon Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*. Springer, 177–186.
- Brent Burley and Walt Disney Animation Studios. 2012. Physically-based Shading at Disney. In *ACM SIGGRAPH*, Vol. 2012. 1–7.
- Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16, 5 (1995), 1190–1208.
- Ewen Cheslack-Postava, Rui Wang, Oskar Akerlund, and Fabio Pellacini. 2008. Fast, realistic lighting and material design using nonlinear cut approximation. In *ACM Transactions on Graphics*, Vol. 27. ACM, 128.
- Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-image SVBRDF capture with a rendering-aware deep network. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 128.
- Maria Lurdes Dias. 1991. Ray tracing interference color. *IEEE Computer Graphics and Applications* 2 (1991), 54–60.
- Jonathan Dupuy and Wenzel Jakob. 2018. An Adaptive Parameterization for Efficient Material Acquisition and Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* (Dec. 2018).
- Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. 2018. Neural architecture search: A survey. *arXiv preprint arXiv:1808.05377* (2018).
- Gabriel Goh. 2017. Why momentum really works. *Distill* 2, 4 (2017), e6.
- Milovš Hašan and Ravi Ramamoorthi. 2013. Interactive albedo editing in path-traced volumetric materials. *ACM Transactions on Graphics (TOG)* 32, 2 (2013), 11.

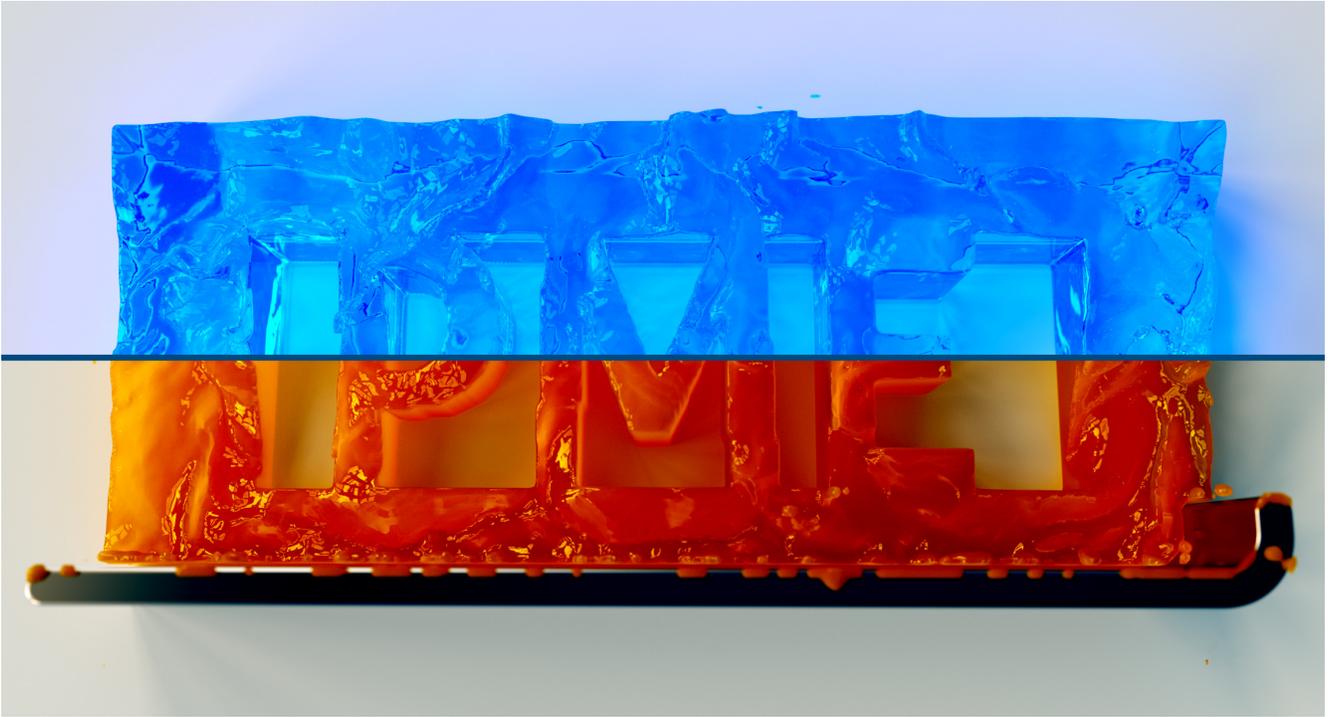


Fig. 9. Our technique is especially helpful early in the material design process where the user seeks to rapidly iterate over a variety of possible artistic effects. Both material types were synthesized using our described method. We demonstrate this workflow in our supplementary video.

Input	Init. type	Init. RMSE	Nelder-Mead	L-BFGS-B	SLSQP	CG	Basin-hopping
Fig. 5, Row 1	Rand	41.93	5.62	20.47	17.96	5.24	2.01
Fig. 5, Row 1	NN	5.94	2.37	5.84	5.94	5.94	
Fig. 5, Row 2	Rand	78.45	40.21	78.45	78.45	78.45	32.67
Fig. 5, Row 2	NN	32.72	32.67	32.72	32.72	32.72	
Fig. 5, Row 4	Rand	35.37	16.98	28.84	35.37	34.99	14.72
Fig. 5, Row 4	NN	18.68	14.68	15.33	18.18	15.90	
Fig. 5, Row 7	Rand	41.65	26.24	41.65	41.65	41.65	22.38
Fig. 5, Row 7	NN	22.42	22.38	22.42	22.42	22.42	
Fig. 5, Row 8	Rand	29.04	22.93	29.04	26.71	28.21	15.69
Fig. 5, Row 8	NN	19.82	15.37	19.82	28.87	19.82	
Fig. 8, Row 2	Rand	23.78	8.26	23.78	23.78	21.75	7.63
Fig. 8, Row 2	NN	12.79	7.80	12.79	12.79	12.79	
Fig. 8, Row 3	Rand	21.60	6.19	21.60	21.60	20.83	5.86
Fig. 8, Row 3	NN	9.09	5.80	9.09	9.09	9.09	
Fig. 8, Row 8	Rand	29.58	6.63	29.58	29.58	29.58	5.07
Fig. 8, Row 8	NN	9.74	5.36	9.61	9.61	9.68	

Table 3. A comparison of a set of classical optimization techniques revealed that when using Nelder and Mead’s simplex-based optimizer with our “best of 9” inversion network initialization, we can often match, and in some cases, outperform the results of Basin-hopping, a global minimizer. In the interest of readability, we have marked the cases where the optimizers were unable to improve upon the initial guess with red. For reference, the first two rows showcase an input image that is reproducible by the shader.

- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering microfacet BSDFs with the Smith model. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 58.
- Magnus Rudolph Hestenes and Eduard Stiefel. 1952. *Methods of conjugate gradients for solving linear systems*. Vol. 49. NBS.
- Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. 2018. Matrix capsules with EM routing. (2018).
- Sho Ikeda, Shin Watanabe, Bisser Raytchev, Toru Tamaki, and Kazufumi Kaneda. 2015. Spectral rendering of interference phenomena caused by multilayer films under global illumination environment. *ITE Transactions on Media Technology and Applications* 3, 1 (2015), 76–84.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- Dieter Kraft. 1994. Algorithm 733: TOMP–Fortran modules for optimal control calculations. *ACM Transactions on Mathematical Software (TOMS)* 20, 3 (1994), 262–281.
- Manuel Lagunas, Sandra Malpica, Ana Serrano, Elena Garces, Diego Gutierrez, and Belen Masia. 2019. A Similarity Measure for Material Appearance. *ACM Transactions on Graphics (SIGGRAPH 2019)* 38, 4 (2019).
- Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. 2017. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 45.
- Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. 2018. Materials for masses: SVBRDF acquisition with a single mobile phone image. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 72–87.
- Xiaojiao Mao, Chunhua Shen, and Yu-Bin Yang. 2016. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *Advances in neural information processing systems*. 2802–2810.
- Stephen Robert Marschner and Donald P Greenberg. 1998. *Inverse rendering for computer graphics*. Citeseer.
- Stephen R Marschner, Stephen H Westin, Eric PF LaFortune, Kenneth E Torrance, and Donald P Greenberg. 1999. Image-based BRDF measurement including human skin. In *Rendering Techniques '99*. Springer, 131–144.
- Gary Mataev, Michael Elad, and Peyman Milanfar. 2019. DeepRED: Deep Image Prior Powered by RED. (2019). [arXiv:cs.CV/1903.10176](https://arxiv.org/abs/1903.10176)
- Wojciech Matusik. 2003. *A data-driven reflectance model*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- Leo Miyashita, Kota Ishihara, Yoshihiro Watanabe, and Masatoshi Ishikawa. 2016. Zoe-Matropo: A system for physical material design. In *ACM SIGGRAPH 2016 Emerging Technologies*. ACM, 24.
- John A Nelder and Roger Mead. 1965. A simplex method for function minimization. *The computer journal* 7, 4 (1965), 308–313.
- Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. 2004. Triple product wavelet integrals for all-frequency relighting. In *ACM Transactions on Graphics (TOG)*, Vol. 23. ACM, 477–487.
- Steven J Nowlan and Geoffrey E Hinton. 1992. Simplifying neural networks by soft weight-sharing. *Neural Computation* 4, 4 (1992), 473–493.
- Marios Papas, Krystle de Mesa, and Henrik Wann Jensen. 2014. A Physically-Based BSDF for Modeling the Appearance of Paper. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 133–142.
- Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. 2013. Fabricating translucent materials using continuous pigment mixtures. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 146.
- Gilles Rainer, Wenzel Jakob, Abhijeet Ghosh, and Tim Weyrich. 2019. Neural BTF Compression and Interpolation. *Computer Graphics Forum (Proceedings of Eurographics)* 38, 2 (March 2019).
- Ravi Ramamoorthi and Pat Hanrahan. 2001. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 117–128.
- Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc Le, and Alex Kurakin. 2017. Large-scale evolution of image classifiers. *arXiv preprint arXiv:1703.01041* (2017).
- Herbert Robbins and Sutton Monro. 1951. A stochastic approximation method. *The annals of mathematical statistics* (1951), 400–407.
- Brigitte Röder, Oliver Stock, Siegfried Bien, Helen Neville, and Frank Rösler. 2002. Speech processing activates visual cortex in congenitally blind humans. *European Journal of Neuroscience* 16, 5 (2002), 930–936.
- Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*. 3856–3866.
- Thorsten-Walther Schmidt, Jan Novak, Johannes Meng, Anton S. Kaplanyan, Tim Reiner, Derek Nowrouzezahrai, and Carsten Dachsbacher. 2013. Path-Space Manipulation of Physically-Based Light Transport. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2013)* 32, 4 (Aug. 2013).
- Ana Serrano, Diego Gutierrez, Karol Myszkowski, Hans-Peter Seidel, and Belen Masia. 2016. An intuitive control space for material appearance. *ACM Transactions on Graphics* 35, 6 (2016), 186.
- Cyril Soler, Kartic Subr, and Derek Nowrouzezahrai. 2018. A Versatile Parameterization for Measured Material Manifolds. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 135–144.
- Ying Song, Xin Tong, Fabio Pellacini, and Pieter Peers. 2009. SubEdit: a representation for editing measured heterogeneous subsurface scattering. *ACM Transactions on Graphics (TOG)* 28, 3 (2009), 31.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- Xin Sun, Kun Zhou, Yanyun Chen, Stephen Lin, Jiaoying Shi, and Baining Guo. 2007. Interactive relighting with dynamic BRDFs. *ACM Transactions on Graphics* 26, 3 (2007), 27.
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. 2013. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*. 1139–1147.
- David J Wales and Jonathan PK Doye. 1997. Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A* 101, 28 (1997), 5111–5116.
- Rui Wang, Ewen Cheslack-Postava, David Luebke, Qianyong Chen, Wei Hua, Qunsheng Peng, and Hujun Bao. 2008. Real-time editing and relighting of homogeneous translucent materials. *The Visual Computer* 24, 7-9 (2008), 565–575.
- Rui Wang, John Tran, and David P Luebke. 2004. All-Frequency Relighting of Non-Diffuse Objects using Separable BRDF Approximation. In *Rendering Techniques*. 345–354.
- Andrea Weidlich and Alexander Wilkie. 2008. Realistic rendering of birefringency in uniaxial crystals. *ACM Transactions on Graphics (TOG)* 27, 1 (2008), 6.
- Randall White. 1989. Visual thinking in the ice age. *Scientific American* 261, 1 (1989), 92–99.
- Alexander Wilkie, Robert F Tobler, and Werner Purgathofer. 2001. Combined rendering of polarization and fluorescence effects. In *Rendering Techniques 2001*. Springer, 197–204.
- Tizian Zeltner and Wenzel Jakob. 2018. *Transactions on Graphics (Proceedings of SIGGRAPH)* 37, 4 (July 2018), 74:1–74:14. <https://doi.org/10.1145/3197517.3201321>
- Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. 2016. Generative visual manipulation on the natural image manifold. In *European Conference on Computer Vision*. Springer, 597–613.
- Barret Zoph and Quoc V Le. 2016. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578* (2016).
- Hui Zou and Trevor Hastie. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67, 2 (2005), 301–320.
- Károly Zsolnai-Fehér, Peter Wonka, and Michael Wimmer. 2018. Gaussian Material Synthesis. *ACM Transactions on Graphics (Proc. SIGGRAPH)* (2018).