Author's Version to appear in Computers & Graphics

Deep scene-scale material estimation from multi-view indoor captures

Siddhant Prakash^a, Gilles Rainer^a, Adrien Bousseau^a, George Drettakis^a

^aUniversité Côte d'Azur and Inria, Sophia-Antipolis, 06902, France

ARTICLE INFO

Article history: Received 18 July 2022 Accepted 28 September 2022 Available online 10 October 2022

Keywords: Material estimation, Deep learning, Indoor scenes, Photogrammetry, Synthetic dataset, Digital 3D assets

ABSTRACT

The movie and video game industries have adopted photogrammetry as a way to create digital 3D assets from multiple photographs of a real-world scene. But photogrammetry algorithms typically output an RGB texture atlas of the scene that only serves as visual guidance for skilled artists to create material maps suitable for physically-based rendering. We present a learning-based approach that automatically produces digital assets ready for physically-based rendering, by estimating approximate material maps from multi-view captures of indoor scenes that are used with retopologized geometry. We base our approach on a material estimation Convolutional Neural Network (CNN) that we execute on each input image. We leverage the view-dependent visual cues provided by the multiple observations of the scene by gathering, for each pixel of a given image, the color of the corresponding point in other images. This image-space CNN provides us with an ensemble of predictions, which we merge in texture space as the last step of our approach. Our results demonstrate that the recovered assets can be directly used for physically-based rendering and editing of real indoor scenes from any viewpoint and novel lighting. Our method generates approximate material maps in a fraction of time compared to the closest previous solutions.

Author's Version Preprint.

1. Introduction

While physically-based rendering is now a mature technology [1], creating the digital assets to be rendered remains a major bottleneck in the creative industry. Photogrammetry has gained popularity to create digital assets from real-world scenes. A popular workflow consists in first capturing multiple photographs of the scene, then using multi-view stereo algorithms to compute an approximate 3D model from these photographs. The approximate geometry is then manually edited to create models compatible with traditional rendering pipelines – a task known as *retopology* [2]. Unfortunately, existing photogrammetry solutions typically output a simple RGB texture of the scene with baked-in lighting, which only serves as a crude initialization for artists who need create rich materials maps used by downstream physically-based renderers. Creating these maps involves significant manual work, including removing shading, shadows and highlights to form the diffuse albedo, and guessing specular strength and roughness parameters over different surfaces. We propose a learning-based approach that addresses this difficult task by augmenting the photogrammetry workflow by automatically estimating approximate material maps from multiple photographs of an indoor scene. Our goal is to provide assets that directly allow plausible renderings of the captured scene. Specifically, the output of our method are approximate Spatially-Varying Bidirectional Reflectance Distribution Function (SVBRDF) texture atlases which, combined with retopologized geometry, forms a digital asset ready for physically-based rendering of indoor scenes. We call these material maps from now on.

e-mail: siddhant.prakash@inria.fr (Siddhant Prakash), gilles.rainer@inria.fr (Gilles Rainer), adrien.bousseau@inria.fr (Adrien Bousseau), George.Drettakis@inria.fr (George Drettakis)



Figure 1: We present a method which takes multiple photographs of a scene as input (a) and predicts surface materials in the form of material maps corresponding to each input view. Merging these image-space predictions in texture space yields a texture atlas (b) that can be mapped onto retopologized geometry to produce *digital 3D assets* ready for full physically-based rendering, i.e., rendering from any viewpoint, changing/adding lights and objects, for example in (c) the golden statuette and the white mug have been added to the scene that has modified lighting and is rendered from a viewpoint not in the input.

Learning-based methods for material estimation have focused on pictures of flat surface patches [3, 4] or on single images of isolated objects [5] and scenes [6], for which the prediction can be efficiently performed in image-space using convolutional neural networks (CNNs). In contrast, approaches based on inverse rendering compute accurate material parameters in object or texture space [7, 8] to benefit from observations from multiple viewpoints. But the underlying optimization is expensive and needs to be recomputed for every new scene. We present the first method that combines ideas from these two streams of research. On the one hand, we leverage the strength of image-space CNNs to predict approximate material parameters for each photograph of the scene. On the other hand, we exploit multi-view information by gathering, for each pixel in a photograph, observations of the same scene point in other photographs. Furthermore, we aggregate the predictions given by each photograph into a common texture space to form the final texture atlas. Our method thus offers the speed of learningbased material estimation previously applied for single images, for the much harder scene-scale material estimation problem.

Our method addresses several difficulties raised by the longstanding challenge of scene-scale material estimation. First, in contrast to single-image methods [3][5][6], our multi-view setting receives a varying number of observations per pixel to be processed by the CNN. We overcome this difficulty by computing a fixed number of color statistics, which forms the initial feature maps that we feed to the CNN. Second, photographs of indoor scenes exhibit complex interactions between geometry, lighting and materials via indirect illumination. Prior work on inverse rendering model these interactions using approximate global illumination (GI) to jointly recover shape, materials and light [6]. In contrast, we consider a scenario where geometry is reconstructed with photogrammetry and retopology, such that we only need to recover material appearance. Rather than approaching this as an inverse rendering problem where global illumination must be estimated accurately to match the input observations, we train an illumination agnostic network to produce materials maps of similar appearance to the inputs. To compare the predictions to the ground truth materials, we use a rendering loss which only operates via local camera-space relighting, avoiding the underconstrained estimation and expensive computation of GI altogether. The third challenge is building a synthetic training dataset suitable for scene-scale approximate material estimation; we created a dataset from professionally modeled scenes, and provide a framework that allows the generation of new datasets for this task.

In summary, our contributions are:

- A deep neural network architecture for material estimation that exploits scene-scale multi-view input.
- A proof-of-concept solution allowing fast, scene-scale material estimation to produce digital assets suitable for physically based rendering and editing of real indoor scenes, that integrates seamlessly into the current photogrammetry workflow.
- A scene-scale synthetic dataset with ground truth SVBRDF maps and the tools to generate it, used to train our multi-view material estimation network.

We evaluate and illustrate our method on synthetic scenes that allow quantitative analysis of our algorithmic choices, and show first results on captured real scenes. We demonstrate that our automatically estimated material map atlases – albeit approximate – are of sufficient quality to allow physically-based rendering of the captured scene with novel lighting conditions and scene editing (see Fig. 1(c), 9). We will provide the source code to our system, including all the tools required to generate the training dataset from commercially-available models.

2. Related Work

We discuss the two domains that inspired our approach – optimization-based and learning-based methods for practical material estimation. We refer the interested reader to surveys on material capture [9, 10] and inverse rendering [11, 12] for more general discussions of these broad topics.

2.1. Optimization-based material capture

The recent development of differentiable rendering algorithms has reinvigorated research on inverse rendering, as pioneered by Yu et al. [7] for scene-scale material recovery. Given a 3D model of the scene, modern approaches estimate material and lighting by simulating complex global illumination effects using differentiable path-tracing [13, 8, 14]. We see our approach as complementary to such optimization-based algorithms that work on scene-scale. On the one hand optimizationbased algorithms are capable of recovering more precise information by minimizing the difference between the input images and the images re-rendered from the estimated materials. On the other hand, such a minimization typically takes 10-12 hours to converge due to complex global illumination computations and is highly sensitive to initialization. Our approach could speed-up these optimization methods by providing an initialization that is much closer to the end result compared to the random material maps that are typically used. Methods that also optimize for geometry have so far been limited to convex isolated objects often with specific lighting/capture constraints [15, 16, 17, 18, 19, 20]. In contrast, we take multiple unconstrained sparse viewpoints of the scene resulting in a variable number of observations for different scene regions and complex visibility issues due to inexact geometry.

2.2. Learning-based material capture

Recovering material appearance from a few observations is an ill-posed problem for which machine learning offers practical solutions. By leveraging large datasets of images paired with ground truth material maps, deep convolutional networks can be trained to predict per-pixel material parameters given a single picture of a flat surface patch [21, 22, 3, 23, 4, 24, 25]. Such methods were later extended to predict material, depth and normal maps of isolated objects [5, 26] and even of indoor scenes [6] from a single image. Ours is the first method to take wide baseline, scene-scale multi-view input under unknown indoor lighting for material estimation. Most related to our goal is the concurrent work by Li et al. [27], who take as input a single image of an indoor scene along with a 3D model of that scene, and build upon single view material prediction [6] to assign procedural material models to object parts. Our work is complementary, as we explore material prediction in indoor scenes under a multi-view capture scenario.

Our key insight is that the multiple images that are typically captured for photogrammetry offer complementary observations of material appearance. However, such multi-view information needs to be properly aggregated to be fed to a neural network, and to be assembled to form a valid texture atlas for later use in rendering engines. Prior methods on multi-image material prediction only partially address these challenges since they focus on small planar patches and assume that each point is visible in all input images [28, 29, 30], which is not the case when dealing with complex scenes where parts are frequently occluded or out of the field of view of many of the input photographs. While Ye et al. [31] describe how to combine image warping and max pooling to handle videos captured with a moving camera, they again focus on planar surfaces free of occlusion, and recover material parameters over a reference frame rather than over all input views of a 3D scene. We propose a solution based on image re-projection and pixel statistics to process multi-view inputs with a standard CNN architecture, and to merge multiple predictions into a single texture atlas.

Neural representations recently emerged as an effective solution to relight 3D content captured from multiple photographs [32, 33, 34, 35, 36]. However, these novel representations are not compatible with the well-established photogrammetry workflow, where artists seek to create triangular meshes and texture atlases compatible with downstream industry-standard rendering engines. While Philip et al. [37] also feeds color statistics as one of the many multi-view information to a neural renderer to perform novel-view synthesis with relighting, we predict explicit material parameters in the form of material texture maps and we assemble the predictions given by multiple views in a common texture space which is readily available to the user for further editing as desired. This post processing flexibility is missing from prior works. More recently, Munkberg et al. [38] combine neural and traditional representations within a differentiable rendering framework to recover a triangle mesh, an SVBRDF texture and an environment map, but their approach has only been demonstrated on isolated objects.

Training methods like ours require large amounts of photorealistic images with ground truth material map labels. Such a dataset is infeasible to capture so we rendered visually realistic synthetic scenes with variations in lighting, materials, geometry and viewpoints. While several datasets of indoor scenes have been described, many only provide images rendered from pre-defined viewpoints and do not allow the generation of new images, as is the case for OpenRooms [39]. Other datasets do not include the labels we are interested in, such as Hypersim [40] that provides diffuse albedo maps and a non-diffuse residual term, which is not directly compatible with existing BRDF models suitable for physically-based renderers. We built a dataset tailored to multi-view material estimation in indoor scenes, by developing an asset generation system that assembles objects from synthetic scenes modeled in Autodesk 3DS Max and then rendering with Mitsuba [41]. We hope our dataset generation tools will help foster research on scene-scale material estimation and other scene-scale learning-based tasks.

3. Overview

Our method takes as input multiple casually-captured unstructured wide-baseline images of a scene from different viewpoints using a DSLR camera. This results in multi-view observations, but in some cases capture is incomplete: In particular, lighting and some parts of the scene might be unobserved. Similar to commercial photogrammetry pipelines [2], first we obtain camera calibration and a rough multi-view stereo mesh using RealityCapture [42], followed by re-topology where an artist creates a clean version of the mesh, suitable for rendering (Fig. 2(a) and 6). This reliable 3D geometry is used as input to our method. Given the multiple input images and the corresponding geometry, our goal is to produce an atlas of mateAuthor's Version/Computers & Graphics (2022)



4



(a) Input multi-view capture and retopologized 3D mesh

(c) Per-image material predictions

(d) Material texture atlas

Figure 2: Overview of our method. We take a multi-view dataset, calibrated cameras and a retopologized mesh as input (a). We re-project information from multiple views into each input view to compute color statistics for each observed point (b). We feed this multi-view information into a convolutional neural network to predict material maps for each view (c). We finally merge this ensemble of inferred maps into texture space (d) to form a scene-scale material texture atlas.

rial parameters, i.e., spatially-varying diffuse albedo D, specular albedo S, and roughness R for a Cook-Torrance BRDF model [43]. We do not estimate normal maps since we focus on indoor scenes composed of large surfaces seen at a distance, for which our retopologized geometry provides sufficiently accurate normals (Fig. 2(a) top and 6).

We achieve this material estimation task in two main steps, illustrated in Fig. 2 (b - d). The first step relies on an image-space CNN to predict material maps for each input image separately. For each view, we use the 3D geometry to reproject image colors from other views and deduce color statistics (minimum, median and maximum color) that summarize the view-dependent appearance of each pixel. We complement these statistics with geometry buffers (surface normals and depth). In practice, we found it beneficial to split the prediction task into two tracks, one responsible for the prediction of diffuse albedo and one responsible for the prediction of specular albedo and roughness.

The second step of our method aggregates the per-view predictions into a common texture space to form the final atlas. We use simple median filtering to select a consensus from the ensemble of predictions given by all views where a surface point appears. Mapping this atlas onto the retopologized mesh gives a complete asset that is compatible with traditional physicallybased rendering (Fig. 1(c), 9), including full editing capabilities such as changing the lighting and inserting new objects.

4. Multi-View Aware Deep Material Estimation

Our problem is estimating scene-scale material properties from a multi-view dataset under unknown lighting, as opposed to the several successful deep learning methods for estimating SVBRDF maps: These start from one or a few images [21, 22, 3, 4, 24, 28, 30], typically for small planar patches of materials lit by a flash. We tackle the scene-scale material estimation problem by first processing each view with a CNN similar to the one used by Deschaintre et al. [3, 28]; We explain how we adapted this architecture to our use-case in Sec. 4.2.

The much harder scene-scale problem precludes the use of co-located flash lighting; Since we cannot benefit from the rich visual cues given by this mode of capture, our originality is to



Figure 3: Overview of our deep learning architecture and training procedure. We split the material estimation task into two tracks, one for the diffuse albedo and the other one for specular albedo and roughness (a, b). In addition, we filter the specular albedo and roughness with a bilateral solver (c). Finally, we compare the predicted maps with ground truth maps using a \mathcal{L}_1 loss as well as with a rendering loss that assesses the appearance of the materials under several lighting and viewing conditions (d).

instead leverage visual cues provided by multi-view observations. Figure 3 illustrates the main components of our architecture and its training procedure.

4.1. Network Inputs

While we run our CNN on each input view separately, we feed it with multi-view information obtained by re-projecting a set of neighboring views to the current view. For each pixel, we select at most 12 views where the corresponding point is visible, and for which the view direction is most closely aligned to the surface normal, as those views are less prone to grazing angle observations. We also add a distance term to favor cameras nearer to the surface.

We rank and pick the top-12 views by minimizing the cost composed of a view term and a distance term as shown in Fig. 4,



Figure 4: Camera selection for each surface point p. A cost function is used to select the 12 best views that consists of a visibility term and a distance term which favors normal-aligned views and cameras close to the surface.

specifically for an observation at *p*:

$$\cos t_i = \cos \alpha_i + \frac{d_{ip}}{\max(d_{jp})} \qquad \forall j \in \{1, N\}$$
(1)

where d_{ip} is the distance of point p to camera i, N is the total number of cameras, and α_i is the angle between the normal at surface point p and view direction of camera i.

Since each pixel might receive a different number of observations as not all surface may be visible in 12 or more views (this is frequently observed in the case of background pixels), we summarize this multi-view information as a fixed set of images corresponding to simple color statistics, i.e. median, maximum and minimum.

The maps that form a material map atlas have different characteristics, and can be inferred from different visual cues. On the one hand, the diffuse albedo map needs to maintain sharp texture features while being free of view-dependent highlights, shadows and indirect light. On the other hand, the specular albedo and roughness values are often nearly constant over parts of objects made of the same material, and are conveyed by highlight information from different views. These observations motivate us to predict the diffuse and specular parameters via two different tracks, and to feed each track with different visual information.

For the diffuse albedo, we complement the input image with the median image, as the median rejects highlights that only appear in a few images. We also include the maximum image to help the network locate shiny areas where highlights might need to be suppressed even if present in many images.

In addition to the images used by the diffuse track, we feed the specular track with the minimum image as it further helps locate parts where the color changes significantly across views. We also provide the specular track with normal and depth maps, which delineate different objects that often have different specular values. Our experiments revealed that providing all these extra images to the diffuse track degrades its prediction, as the network struggles to select the relevant image structures among too many visual channels. In addition, reprojection errors due to the approximations in retopologized geometry can be problematic (see Sec. 7.1).

Finally, we favor smooth specular maps by post-processing their predictions with a differentiable bilateral solver [44, 6] guided by the predicted diffuse albedo. This edge-aware smoothing attenuates discontinuities due to reprojection misalignments. Unlike Li et al. [6] we do not run our diffuse albedo maps through the solver as we observed this leads to over-smoothing of diffuse albedo maps resulting in loss of texture details.

4.2. Network Architecture

Our network architecture is based on the ones by Deschaintre et al. [3, 28] which were also designed for material estimation but works on a single flat patch of size 256x256x3. Their architectures follow the widely popular U-Net encoder-decoder [45], to which a fully-connected track responsible for processing and transmitting global information was added. We maintain the same U-Net architecture augmented with the global track but we half the number of feature layers from 8 to 4 to make the networks lighter with the feature counts in the encoder downscaling layers of 64, 128, 256, and 512. We did not observe any significant degradation in the maps due to this reduction in network capacity. We follow the same downsampling and upsampling process as described in [3] with the feature counts used in reverse order for the decoder. Instance normalization is used for stability and we also regularize by applying dropout at 50% probability on the last three layers of the decoder.

The main difference lies in the input/output; Specifically our network architecture differs in two ways. First, we separate out the network into two tracks one each for diffuse and specular maps. The two tracks serve different purposes with the diffuse track predicting only diffuse albedo maps, while the specular track outputs the specular albedo and roughness maps. The two tracks are two separate networks. While training, we can either train one of the networks without sharing any information between them or both networks by jointly computing the loss on output of both networks. We use individual training to train the networks and joint training to fine-tune the networks posttraining. Please see Sec. 6 for training details. Second, since we feed the network with multi-view information in the form of image statistics, we have more input channels compared to 3 channels for the previous networks. Concretely, since we feed the median, maximum and input image to the diffuse track which thus has 9 input channels, while for the specular track, we also provide the minimum image, depth and normal buffers resulting in 18 input channels.

For the network used to predict confidence channels to guide the bilateral solver, we follow the same CNN architecture and hyper-parameters for the solver as used in previous works [6, 44]. We include detailed breakdown of the network architecture and the parameters used in the supplemental PDF.

4.3. Loss Function

Many inverse rendering methods supervise their predictions of geometry, material and lighting by comparing re-rendered images to input images using a differentiable *rendering loss* [5, 26]. At scene-scale, such a rendering loss needs to model global illumination effects present in the input [6]. We depart from this family of methods by focusing on a scenario where geometry is given, such that our task boils down to predicting material maps only. In this context, we can supervise our method by comparing our prediction to ground truth SVBRDF maps rather than by attempting to reproduce the input. A local lighting model is sufficient for this purpose, as was originally proposed by Deschaintre et al. [3] in the context of planar surface patches. While Deschaintre et al. use point and directional lights, we improve on their approach by incorporating distant area lights modeled as spherical Gaussians.

Concretely, we use a simple differentiable renderer that takes as input the material maps along with the normals of the geometry. Our goal is to compare the local appearance of our predicted material maps to the local appearance of the ground truth SVBRDFs. To do so, we render the prediction and the ground truth under several viewing and lighting conditions and compare the resulting images under the \mathcal{L}_1 norm and E-LPIPS perceptual metric \mathcal{L}_E [46] after applying a log transform (I' = log(0.1 + I)) to compress the dynamic range of the renderings. Following Deschaintre et al., we generate random viewing conditions by sampling view vectors over the hemisphere centered around the original view direction from which the input image was rendered. We then generate lighting conditions likely to produce highlights by positioning a point light in the mirror direction of the view vector. Finally, we also create extended light sources by generating a mixture of 5 Gaussian lights with random width, color and direction distributed over the hemisphere.

We implement the shading of a point under distant area light sources by using the spherical warp introduced by Wang et al. [47]. The light, as well as the BRDF, are approximated as two Spherical Gaussians, for which a fast closed-form convolution exists. Using this approximation allows us to include extended light sources in the rendering loss without losing computational efficiency for training. Our final rendering loss averages the image differences obtained with three point-wise lighting conditions and with three extended lighting conditions.

In addition to the rendering loss, we also use \mathcal{L}_1 and \mathcal{L}_E to compare the individual predicted maps to their respective ground-truth. Denoting *I* a rendered image, *D* the diffuse albedo, *R* the roughness, and *S* the specular albedo, the total loss we use is thus:

$$\mathcal{L} = [\mathcal{L}_E(I) + \mathcal{L}_1(I)] + \mathcal{L}_E(D) + \mathcal{L}_E(R) + \mathcal{L}_E(S) + \lambda(\mathcal{L}_1(D) + \mathcal{L}_1(R) + \mathcal{L}_1(S))$$
(2)

where λ is 0.1.

4.4. Merged Renderable Scene Assets

The second step of our method merges the material maps predicted over each input view to form a single, object-space material map texture atlas suitable for rendering. We leverage the retopologized 3D mesh to identify which texel corresponds to each pixel in all input views. We select the final value of each texel as the median value of all its predictions. This median filter is especially effective at removing erroneous predictions that are not consistent across views (see Sec. 7.3.1) including the ones due to re-projection artifacts caused by approximate geometry and camera calibration in real-world scenes.



Figure 5: Renderings from our synthetic dataset used for training our network. The dataset has a variety of different lightings, materials and viewpoints.

5. Synthetic Training Dataset

We trained our method by generating a dataset of synthetic renderings with corresponding ground truth material maps. We use professional artist-modeled assets in Autodesk 3DS Max with high quality V-Ray materials, that help bridge the gap between training data and real re-topologized scenes. We purchased a set of scenes ¹ and extracted basic environments and several different objects that we recombined to create new scene configurations, on which we place objects with different materials.

We augment the initial artist-generated materials with materials from Deschaintre et al. [3], hand-picked to correspond well with the underlying geometry and to cover a wide range of everyday indoor materials such as wood, metal, plastic, rubber, leather, etc. Each choice of materials and objects provides a scene configuration, for a total of 160 scene configurations, created from 5 "base scenes" with a set of random object placements. We place area and point lights in the scene, as well as environment maps that typically illuminate the scenes through a window.

We rendered each image using a Cook-Torrance BRDF model with a Beckmann normal distribution, which we have implemented in the Mitsuba physically-based path tracer for full global illumination [41]. We subsequently denoised each rendering using the Optix denoiser [48]. For each image, we also generate the ground truth SVBRDF maps, i.e., diffuse albedo, roughness and specular albedo rendered as images. Finally, we pre-compute the per-pixel re-projected color statistics (minimum, median and maximum) for each image in our dataset.

We render each scene configuration under 40 different viewpoints, yielding a total of 6400 images at resolution 640×384 pixels (see Fig. 5 for a small selection). At training time, we extract random crops of 256×256 pixels from each image to be

¹From https://evermotion.org/shop/cat/355/all_scenes/0/0, Volume 1, 8 and 30.



Figure 6: Retopologized geometry with selected input images for REAL KITCHEN (left) and REAL HALLWAY (right) scenes.

fed to the network, which effectively augments the size of the dataset, to around 45,000 individual crops.

6. Implementation Details

We have implemented our method in python using the Py-Torch [49] framework for deep learning and C++/OpenGL shaders for all steps that require reprojection.

At inference, we run the CNN over images in 640×384 resolution. If the aspect ratio does not match, we zero-pad to fit to the nearest multiple resolution. We expect input images in linear color space, and we apply a log transformation followed by a normalization to [-1, 1] to flatten the dynamic range before processing by the CNN. For real images, we assume gamma correction of 2.2 to convert to linear space.

Our system for dataset generation includes a plug-in for 3DS Max that exports materials into MITSUBA-compatible format using our BRDF model. To handle complex material graphs, we evaluate them in 3DS Max and save texture layers subsequently used by MITSUBA.

We first train the diffuse track and then the specular track separately over 15 epochs each, using ground truth maps for the missing components when evaluating the rendering loss. We then fine-tune the two tracks jointly for 15 epochs. Overall training takes 18 hours on a 4 RTX8000 GPU cluster node. The confidence networks for roughness and specular maps are trained for 100 iterations. We use the ADAM optimizer [50] with a fixed learning rate of 2e - 5 for all training.

We will release all source code for our method, including the dataset generation system, and all training images used for our results.

7. Results and Evaluation

We show results and evaluations on two synthetic scenes (Veach Ajar, and Dining Room), and three real captured scene (Real Office, Real Kitchen, and Real Hallway). The capture details and scene lighting conditions are provided in Table 1. We also provide comparisons on two additional synthetic scenes (Living Room, and Kitchen). We provide a supplemental video showing view dependent effects over paths and image sequences. We strongly encourage the reader to view the videos to appreciate how our automatically created material maps are directly usable for physically-based rendering and scene editing.

For synthetic scenes, we render a set of views of the scene, and then use these as if they were photographs to run our entire pipeline; we use the original geometry in this case. Note that as a result the ground truth materials are encoded as a single material map texture atlas to be comparable with the results of our method. For real scenes, we take a set of photos of the scene, paying attention to capture highlights in several views, then run structure-from-motion and multi-view stereo to obtain an initial 3D reconstruction. We hired a professional artist to turn these reconstructions into a retopologized mesh (shown in Fig. 2(a) and 6). We use Blender automatic UV-unwrapper to unwrap the meshes into texture atlas. The resolution of the texture atlas is 16Kx16K pixels for the scenes we considered.

Synthetic data allows quantitative comparisons on both the material maps and the renderings; for real scenes we can only show qualitative results due to the lack of ground truth maps.

7.1. Results

In Fig. 7 and 8, we show the scene geometry textured with final *approximate* material maps obtained by our method. For synthetic scenes we also show the ground truth maps. Additionally, for both synthetic and real scenes, we show the re-rendered image, i.e., we generate the material map texture atlases using our method and then provide them to a path tracer along with the geometry to render the scene with full global illumination effects. To achieve a result as close as possible to the input image, we place the lights manually to match input conditions as much as possible. Despite these approximations, our re-renderings are plausible renditions of the input images, illustrating the efficacy of the approximate material maps we obtain.

Our method manages to capture the overall material properties of the objects even in real scenes, e.g., in the Real Office scene the desktop and the red box are shiny while the yellow box on the right and the orange are more diffuse.

In Fig. 9 we show results with modified lighting conditions and object insertion from different viewpoints on real scenes. This figure shows that we achieve our goal of creating plausible material assets for photorealistic scene editing. For each scene, we show a view in input lighting condition and the same view with modified lighting condition. We are able to remove and move shadows and highlights on most surfaces. We further augment the scene by inserting complex objects, such as a metallic statuette in the Real Office scene and a transparent water goblet and wine bottles in the Real Kitchen scene. Our material maps, together with the retopologized geometry are complete digital assets, and thus allow renderings with full GI interactions between real and virtual objects, such as color bleeding, refraction, caustics and internal reflections when rendered using a path tracer. We emphasize that such effects are not possible to reproduce using relighting based object insertion methods such as [51, 52, 53].

7.2. Evaluation

The only other methods designed to handle our input at scene-scale are differentiable rendering approaches [8, 14]; un-



Figure 7: Results on two synthetic scenes where the ground truth is available. For each scene, in the first row we show an input view and the ground truth diffuse, roughness and specular maps for that view. The second row shows the re-rendering followed by the maps obtained using our method for the input view shown. We are able to reproduce renderings which are close to the input view using the approximate material maps generated by our method.

Scenes	#images	Resolution	Pre-processing	Prediction	Texturing (×3)	Total
Veach Ajar	160	640×384	7.56	3.93	3.12	20.85
Dining Room	105	640×384	4.56	3.26	1.32	9.14
Real Office	245	640×355	11.48	5.84	4.77	31.63
Real Kitchen	296	640×411	32.05	25.25	7.91	81.03
Real Hallway	226	640×414	16.10	10.26	7.03	47.45
Mean			14.35	9.71	4.83	38.02

Table 1: Timing breakdown for each step of our method on the scenes used for our experiments. All times are reported in minutes.

fortunately neither code nor data (in the form of input images we can use for SfM/MVS) is available, precluding direct comparison. In any case, our method can be seen as complementary and could be used as an initialization for these methods, potentially accelerating their process. We report the timings for different steps of our method and compare with timings reported by these previous works to support our claims. Additionally, we present best-effort evaluation using two baselines. We also present a set of ablation studies to analyse the effect of our various design choices.

7.2.1. Speed

We show the timing breakdown of each step for our method (pre-processing (PRE-PROCESSING), single-view prediction (PREDICTION), and texture atlas generation (TEXTURING) on a

system with an Intel Xeon Gold 5218 2.30GHz Processor and Quadro RTX 5000 GPU, in Tab. 1. We observe that it is possible to create renderable assets from a multi-view dataset with approximately 30 *minutes* to an hour of computation depending on number of images and their resolution. Previous works dealing with scene-scale material estimation [8, 14] report around 10 - 12 hours for a complete scene optimization. Thus, our method is able to produce renderable material maps in *a fraction of time* as compared to previous works. Note that the timings reported do not include the time taken for reconstruction and re-topology of the mesh.

7.2.2. Comparisons

We compare to two baselines: The first mimics current practice in digital content creation and uses the retopologized geom-



Figure 8: Results on real scenes. For each scene, we show the re-rendering in closest matched capture conditions and the maps obtained for the given input view using our method.

etry to project the input images into texture space, using median filtering to create an RGB atlas (TEXTURE), while the second is based on the single image method of Li et al. [6] (LIETAL). Specifically, we run the method of Li and colleagues to generate maps for each input view, then we run the same pipeline as for our method to create a material texture atlas for the scene from the multiple predicted maps. We implement their BRDF model in MITSUBA to generate the re-renderings. We perform quantitative and qualitative comparisons for our method compared to the two baselines. Since the first baseline does not estimate maps, and LIETAL infers maps for a different BRDF model, we only compare re-renderings, i.e., we rerender the scene for a set of views. We perform quantitative comparisons by computing PSNR and DSSIM error [54]. We show the numerical results in Tab. 2, and a visual comparison in Fig. 10 and 11. For the visual comparison, we show the input Author's Version/Computers & Graphics (2022)



Figure 9: Scene editing with our method on real scenes from input and novel viewpoints. From left to right, we show the scene re-rendered with capture configurations, the same view re-rendered with modified lighting, with a virtual object inserted along with modified lights, and the modified scene from a novel viewpoint which is never captured. For example, for the real office scene we have added the statuette; notice how shadows underneath the red book and orange fruit have been (re)moved correctly on modifying lights (first row, blue and red inset) and the color bleeding of the red box and orange on the statuette (first row, green and yellow inset) due to *correct interaction between real and virtual objects* which cannot be achieved by relighting based object-insertion methods. Our method enables insertion of objects with highly complex materials such as in the real kitchen scene, we have inserted a goblet of water and two bottles of wine. Since our method produces digital assets that can be directly used for rendering, our renderings exhibit complex lighting phenomena such as caustics (second row, red inset) and the red shoe-stone in real hallway scene as illustrated by the rendered highlights (second row, red inset) with red inset; third row the specular material e.g., on the knives and the red shoe-stone in real hallway scene as illustrated by the rendered highlights (second row, red inset; third row, blue inset). Please see the supplementary video for results with moving paths and lights to better appreciate these effects.

view on which the maps are predicted. In line with our aim of generating plausible material assets for scene editing, we generate ground truth re-renderings of a modified scene by rendering the scene with modified lighting using ground truth SVBRDF maps. To show how our method compares against the two baseline for this task, we re-render the scene with modified lighting but with the material maps obtained by the three methods. We see that our approach shows much better variations of appearance, i.e., shiny materials, compared to the LIETAL that struggles to identify shiny materials and TEXTURE where everything is diffuse by construction. Our method works well for novel views which we show by re-rendering a novel view for each scene. This view was never seen by the network and no maps were predicted for this viewpoint. The quantitative results are computed on 10 views of each scene selected from a rendered path; we also show a video comparison on the rendered path along with the ground truth in the supplemental video. We see in Tab. 2 that our method is numerically best for both synthetic scenes on both PSNR and DSSIM metrics. In Veach Ajar scene our method performs significantly better than the baselines as



Figure 10: Comparisons with baselines on synthetic scenes; For each scene, we show an input view and a novel view; in the first row we show the viewpoint with original input lighting condition (ORIGINAL LIGHT GT). From second row onwards we modify the lighting in the scene and show the same view re-rendered with modified lighting condition using ground truth maps (Modified LIGHT GT), our maps (OURS), maps produced by LIETAL and the baseline static TEXTURE generated using the input images. Column 1 and 3 corresponds to an input view while column 2 and 4 corresponds to a novel view for which no maps were predicted. Notice how LIETAL fails to reconstruct highlight properly and the TEXTURE is composed of pre-baked highlights and shadows from original lighting condition.

	PSNR ↑			DSSIM↓		
Method	LIETAL	Texture	Ours	LIETAL	Texture	Ours
Veach Ajar	12.97	18.70	21.53	0.37	0.30	0.27
Dining Room	21.33	15.12	21.77	0.36	0.62	0.34
Living Room	13.67	19.97	25.61	0.44	0.44	0.16
Kitchen	9.06	16.54	21.28	0.52	0.42	0.30
Mean	14.26	17.58	22.55	0.42	0.44	0.27

Table 2: Quantitative comparison on 4 synthetic scenes. Our method performs best across the scenes which supports our qualitative observations (see also Fig. 10 and 11).

we can see how our method is able to approximately recover the gradation in specularity between the teapots while both the baselines fails to do so. In Dining Room scene, numerically our results are better yet close to LIETAL, although it's worth noticing that visually our method is able to do much better. For example, we are able to recover the spatial variation in roughness of the table top which LIETAL fails to do and ends up oversmoothing the diffuse albedo. This shows the advantage of using multi-view information which helps the network infer the spatially-varying nature of roughness locally as well as globally for each surface point.

Since we do not have ground truth for real scenes, we show qualitative results on REAL OFFICE scene in Figure 12. We show 5 re-renderings of the scene with modified lighting using our method and the baselines. In line with our observations on synthetic scenes, we observe that our method is able to predict the variations more accurately than LIETAL which fails to reconstruct highlights on glossy surfaces such as the tabletop and red box. Compared to TEXTURE we can easily see that the surfaces do not reflect the change in lighting and the baked-in shadows



Figure 11: More comparisons with baselines on synthetic scenes; The layout is the same as for Fig. 10. Similarly as in Fig. 10 we see that our approach shows much better variations of appearance, i.e., shiny materials, compared to the LIETAL that struggles with shiny material sand texture details and TEXTURE that contains pre-baked highlights and shadows from input images owing to its diffuse and static nature.

and highlights are still present due to the static and diffuse nature of the texture; please see the video to appreciate the visual importance of this effect.

7.3. Ablations

7.3.1. Multi-view Reprojections

We perform a first ablation on the two main components of our algorithm: 1) we remove the statistics reprojected from other views (No RP) and 2) we remove the merging of maps in texture space (Ours Im.). We show quantitative comparisons, where we provide error in form of Mean Squared Error (MSE) for the three maps computed for 10 randomly selected input views.

We show quantitative results in Tab. 3,4 and an example of the visual effect of the different cases of increasing multi-view information with each step in Fig. 13. Using all our components improves results in the majority of cases. The use of the reprojected image statistics makes a very significant difference in the quality of the maps. While merging in texture space may not significantly improve quality, it helps with increasing consistency between different views for the underlying surface material properties (especially for roughness and specular maps) and thus helps improve quality of the final re-rendering. Reprojection improves roughness on most of the objects in the scene, but sometimes has a negative effect on background parts that lack observations (see Fig. 13); this explains why No-RP has better MSE for roughness in Dining Room in Tab. 3.

7.3.2. Inaccurate Geometry

We study the robustness of our method for inaccurate geometry by running our pipeline on the degraded mesh obtained directly from multi-view stereo (MVS) [42] which consists of large holes and bumpy surfaces. The bumpy surfaces are an instance of extreme vertex perturbation. We show a qualitative comparison in Fig. 14. From the figure we can confirm that the maps obtained from the MVS mesh is only slightly degraded



Figure 12: Comparison with baselines in REAL OFFICE scene. We render a novel viewpoint with changing lighting conditions and show 5 frames with different lighting conditions. Notice how our method (re)move shadows (red arrow) and reconstruct highlights (green arrow) accurately while both LIETAL and TEXTURE fail to do so. Please see the supplemental video to appreciate the smooth transition of the highlights and shadows as a result of movement of the light sources.

Synthetic Veach Ajar						
	MSE↓					
Method	Diffuse	Roughness	Specular			
Ours	0.022135	0.055563	0.016522			
Ours Im.	0.030718	0.064362	0.049969			
No RP	0.141682	0.372725	0.055459			
Synthetic Dining Room						
		MSE ↓				
Method	Diffuse	Roughness	Specular			
Ours	0.026491	0.071609	0.028250			
Ours Im.	0.026642	0.077578	0.024123			
No RP	0.053104	0.048222	0.032385			

Table 3: Quantitative evaluation for the ablation on the synthetic scenes on the material maps, with 1) no reprojected statistics, 2) no multi-view merge in texture space. We see an increase in multi-view information helps improve the maps quality and/or consistency across views (see also Tab. 4 and Fig. 13).

	Veach	n Ajar	Dining Room		
Method	PSNR ↑	DSSIM \downarrow	PSNR ↑	DSSIM \downarrow	
Ours	19.520969	0.190185	19.35363	0.244253	
No RP	11.338421	0.413321	17.675568	0.343187	

Table 4: Quantitative evaluation for the ablation on the synthetic scenes on the re-renderings, with 1) no reprojected statistics, 2) no multi-view merge in texture space. Using multi-view statistics achieves better re-rendering quality than using only a single view (see also Tab. 3 and Fig. 13).

compared to the re-topologized mesh. Thus, our material estimation is robust to geometrical inaccuracies. While the maps obtained are similar, the final image obtained after re-rendering is highly degraded when rendered since the MVS geometry has bumps and holes. To obtain high quality re-renderings we need good geometry, justifying our design choice of using retopology. In future work, it may be possible to adapt previous methods (e.g., [55, 56]) to provide geometry that corrects these errors for flat surfaces, but it would be necessary to preserve the relatively well reconstructed irregular objects (such as the fruit on the table).

8. Limitations and Future Work

Despite yielding convincing rerenderings on synthetic and on real scenes, our method still has several limitations. Our final texturespace maps often suffer from limited resolution, since the texture atlas can provide only limited space for a



given object (see inset, where the texture of the red box covers only a very small part of the texture atlas). This is an inherent problem with texture-space methods, and alternative approaches (e.g., see [57]) could be a good direction for future work. While reprojection error may be a contributing factor for blurriness in results, we believe our use of median filter to merge and obtain texture-space maps alleviates this problem. We focused on estimating BRDF parameters for each texel of a texture atlas. A natural extension would be to also estimate per-texel normals expressed in a local coordinate frame, which would allow the reproduction of small geometric details not modeled in the retopologized mesh. However, such small-scale relief is often difficult to perceive when captured from far away.

In some cases, e.g., the banana and pear, our method predicts glossiness that is high; we hypothesize that this is due to the reprojection errors due to the mismatch between the re-topologized and real scene geometry. The effect of this is visible in the reprojected min image (see inset).



The dataset we created to demonstrate our approach offers limited variability, which in turns limits the ability of our method to handle diverse scenes. While we provide our toolbox to generate additional training images, rendering large datasets is costly and could benefit from strategies to reuse computation



Figure 13: Example images from SYNTHETIC DINING ROOM showing the effect of increasing multi-view information on results. Note how the quality of the maps is significantly improved by using the reprojected statistics observed in our image space predicted maps as compared to no reprojection, i.e. using only a single image. Furthermore, gathering the image space maps in texture space helps improve the consistency of the maps across views and thus improves re-rendering by assigning same material in local regions (esp. in roughness and specular maps). As a result, re-rendering is closer to ground truth.



Figure 14: Effect of inaccurate geometry. We run our pipeline on the MVS mesh obtained directly from [42]. The results from the MVS mesh is on the left column and our re-topologized mesh on the right column. We show a rerendering in the first row, followed by the albedo, roughness and specular maps obtained using our method in subsequent rows respectively.

across views [58]; we hypothesize that augmenting the variety and the number of training images seen by our network will improve results overall, possibly helping remove shadow and incorrect color residuals that are sometimes still present in our albedo maps.

We rely on fixed color statistics to aggregate the multi-view information that we feed to our per-view CNNs, and we employ a fixed median filter to merge the resulting per-view predictions into texture space. Replacing these two operations by differentiable pooling in a learned feature space could yield improved predictions, as has been done in other applications [59, 60]. However, training such an architecture end-to-end raises specific challenges, such as storing multiple CNN tracks in memory and performing differentiable re-projection in the texture atlas while doing per-image processing.

9. Conclusion

We have presented the first attempt at creating scene-scale material map textures of indoor environments using deep learning. Our solution retains the strength of image-space CNNs, which have proven successful at recovering material parameters for close-up photographs of flat surfaces and isolated objects. To apply such CNNs at scene scale, we first inject multi-view information by computing statistics about a pixel color when re-projected into neighboring views. We then move to texture space to merge per-view predictions into a single texture atlas suitable for rendering.

Our method allows automatic generation of material maps that allow plausible renderings, retaining the overall look of multiple objects in a scene, both for synthetic scenes with available ground truth and for real scenes. Our results demonstrate that by exploiting reprojected multi-view data, improving the rendering loss and exploiting state-of-the-art components it is possible to provide an operational pipeline to extract convincing materials at scene-scale from a set of images as input.

Acknowledgments

The research was funded by the ERC Advanced grant FUNGRAPH No 788065 (https://project.inria.fr/fungraph/). The authors are grateful to Inria Sophia Antipolis - Méditerranée "Nef" computation cluster and the OPAL infrastructure from Université Côte d'Azur for providing resources and support. We thank the associate editor and the anonymous reviewers for their insightful comments which helped improve the manuscript. We also thank B. Bitterli for the scenes in his rendering repository. The authors also extend their thanks to Felix Hähnlein and Emilie Yu for their helpful discussion and support for the project and especially thank the 3D artist Stefania Kousoula for mesh refinement and re-topology.

References

- Fascione, L, Hanika, J, Fajardo, M, Christensen, P, Burley, B, Green, B. Path tracing in production. In: ACM SIGGRAPH 2017 Courses. 2017, p. 1–39.
- [2] Lachambre, S, Lagarde, S, Cyril, J. Unity photogrammetry workflow. https://unity3d.com/files/solutions/photogrammetry/ Unity-Photogrammetry-Workflow_2017-07_v2.pdf; 2017.
- [3] Deschaintre, V, Aittala, M, Durand, F, Drettakis, G, Bousseau, A. Single-image svbrdf capture with a rendering-aware deep network. ACM Trans Graph 2018;37(4). URL: https://doi.org/10.1145/ 3197517.3201378. doi:10.1145/3197517.3201378.
- [4] Guo, J, Lai, S, Tao, C, Cai, Y, Wang, L, Guo, Y, et al. Highlight-aware two-stream network for single-image svbrdf acquisition. ACM Trans Graph 2021;40(4). URL: https://doi.org/10.1145/ 3450626.3459854. doi:10.1145/3450626.3459854.
- [5] Li, Z, Xu, Z, Ramamoorthi, R, Sunkavalli, K, Chandraker, M. Learning to reconstruct shape and spatially-varying reflectance from a single image. ACM Trans Graph 2018;37(6). URL: https://doi.org/10. 1145/3272127.3275055. doi:10.1145/3272127.3275055.
- [6] Li, Z, Shafiei, M, Ramamoorthi, R, Sunkavalli, K, Chandraker, M. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2020,.
- [7] Yu, Y, Debevec, P, Malik, J, Hawkins, T. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '99; USA: ACM Press/Addison-Wesley Publishing Co. ISBN 0201485605; 1999, p. 215–224. URL: https:// doi.org/10.1145/311535.311559. doi:10.1145/311535.311559.
- [8] Nimier-David, M, Dong, Z, Jakob, W, Kaplanyan, A. Material and Lighting Reconstruction for Complex Indoor Scenes with Texturespace Differentiable Rendering. In: Bousseau, A, McGuire, M, editors. Eurographics Symposium on Rendering - DL-only Track. The Eurographics Association. ISBN 978-3-03868-157-1; 2021,doi:10.2312/ sr.20211292.
- [9] Garces, E, Rodriguez-Pardo, C, Casas, D, Lopez-Moreno, J. A survey on intrinsic images: Delving deep into lambert and beyond. International Journal in Computer Vision 130, 836-868 2022;.
- [10] Guarnera, D, Guarnera, G, Ghosh, A, Denk, C, Glencross, M. Brdf representation and acquisition. Computer Graphics Forum 2016;35(2):625–650. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12867. doi:https://doi.org/10.1111/cgf.12867. [27]arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12867.
- [11] Tewari, A, Thies, J, Mildenhall, B, Srinivasan, P, Tretschk, E, Wang, Y, et al. Advances in neural rendering. 2021. URL: https://arxiv. org/abs/2111.05849. doi:10.48550/ARXIV.2111.05849.
- [12] Kato, H, Beker, D, Morariu, M, Ando, T, Matsuoka, T, Kehl, W, et al. Differentiable rendering: A survey. arXiv preprint arXiv:200612057 2020;.
- [13] Azinovic, D, Li, TM, Kaplanyan, A, Niessner, M. Inverse path tracing for joint material and lighting estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2019,.
- [14] Haefner, B, Green, S, Oursland, A, Andersen, D, Goesele, M, Cremers, D, et al. Recovering real-world reflectance properties and shading from hdr imagery. In: 2021 International Conference on 3D Vision (3DV). 2021, p. 1075–1084. doi:10.1109/3DV53792.2021.00115.
- [15] Wu, H, Wang, Z, Zhou, K. Simultaneous localization and appearance estimation with a consumer rgb-d camera. IEEE Transactions on Visualization and Computer Graphics 2016;22(8):2012–2023. doi:10.1109/TVCG.2015.2498617.
- [16] Nam, G, Lee, JH, Gutierrez, D, Kim, MH. Practical svbrdf acquisition of 3d objects with unstructured flash photography. ACM Trans Graph 2018;37(6). URL: https://doi.org/10.1145/3272127.3275017. doi:10.1145/3272127.3275017.
- [17] Goel, P, Cohen, L, Guesman, J, Thamizharasan, V, Tompkin, J, Ritchie, D. Shape from tracing: Towards reconstructing 3d object geometry and svbrdf material from images via differentiable path tracing. In: 2020 International Conference on 3D Vision (3DV). Los Alamitos, CA, USA: IEEE Computer Society; 2020, p. 1186– 1195. URL: https://doi.ieeecomputersociety.org/10.1109/ 3DV50981.2020.00129. doi:10.1109/3DV50981.2020.00129.

- [18] Ma, X, Kang, K, Zhu, R, Wu, H, Zhou, K. Free-form scanning of non-planar appearance with neural trace photography. ACM Trans Graph 2021;40(4). URL: https://doi.org/10.1145/3450626.3459679. doi:10.1145/3450626.3459679.
- [19] Luan, F, Zhao, S, Bala, K, Dong, Z. Unified shape and svbrdf recovery using differentiable monte carlo rendering. Computer Graphics Forum 2021;40(4):101-113. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/ cgf.14344. doi:https://doi.org/10.1111/cgf.14344. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14344.
- [20] Bi, S, Xu, Z, Sunkavalli, K, Kriegman, D, Ramamoorthi, R. Deep 3d capture: Geometry and reflectance from sparse multi-view images. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2020, p. 5959–5968. doi:10.1109/CVPR42600.2020. 00600.
- [21] Li, X, Dong, Y, Peers, P, Tong, X. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. ACM Trans Graph 2017;36(4). URL: https://doi.org/10. 1145/3072959.3073641. doi:10.1145/3072959.3073641.
- [22] Li, Z, Sunkavalli, K, Chandraker, M. Materials for masses: SVBRDF acquisition with a single mobile phone image. In: Ferrari, V, Hebert, M, Sminchisescu, C, Weiss, Y, editors. Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part III; vol. 11207 of *Lecture Notes in Computer Science*. Springer; 2018, p. 74–90. URL: https://doi.org/10.1007/ 978-3-030-01219-9_5. doi:10.1007/978-3-030-01219-9_5.
- [23] Gao, D, Li, X, Dong, Y, Peers, P, Xu, K, Tong, X. Deep inverse rendering for high-resolution svbrdf estimation from an arbitrary number of images. ACM Trans Graph 2019;38(4). URL: https://doi.org/ 10.1145/3306346.3323042. doi:10.1145/3306346.3323042.
- [24] Zhou, X, Kalantari, NK. Adversarial single-image svbrdf estimation with hybrid training. Computer Graphics Forum 2021;40(2):315– 325. URL: https://onlinelibrary.wiley.com/doi/abs/10. 1111/cgf.142635. doi:https://doi.org/10.1111/cgf.142635. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.142635
- [25] Henzler, P, Deschaintre, V, Mitra, NJ, Ritschel, T. Generative modelling of brdf textures from flash images. arXiv preprint arXiv:210211861 2021;.
- [26] Boss, M, Jampani, V, Kim, K, Lensch, HP, Kautz, J. Twoshot spatially-varying brdf and shape estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2020,.
- [27] Li, Z, Hold-Geoffroy, Y, Zhu, R, Xu, Z, Hasan, M, Sunkavalli, K,
- 2867. et al. Photoscene: Photorealistic material and lighting transfer for indoor scenes. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2022,.
- [28] Deschaintre, V, Aittala, M, Durand, F, Drettakis, G, Bousseau, A. Flexible SVBRDF Capture with a Multi-Image Deep Network. Computer Graphics Forum 2019;doi:10.1111/cgf.13765.
- [29] Guo, Y, Smith, C, Hašan, M, Sunkavalli, K, Zhao, S. Materialgan: Reflectance capture using a generative svbrdf model. ACM Trans Graph 2020;39(6). URL: https://doi.org/10.1145/3414685.3417779. doi:10.1145/3414685.3417779.
- [30] Asselin, L, Laurendeau, D, Lalonde, J. Deep svbrdf estimation on real materials. In: 2020 International Conference on 3D Vision (3DV). Los Alamitos, CA, USA: IEEE Computer Society; 2020, p. 1157– 1166. URL: https://doi.ieeecomputersociety.org/10.1109/ 3DV50981.2020.00126. doi:10.1109/3DV50981.2020.00126.
- [31] Ye, W, Dong, Y, Peers, P, Guo, B. Deep reflectance scanning: Recovering spatially-varying material appearance from a flash-lit video sequence. Computer Graphics Forum 2021;URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/ cgf.14387. doi:https://doi.org/10.1111/cgf.14387. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14387.
- [32] Srinivasan, PP, Deng, B, Zhang, X, Tancik, M, Mildenhall, B, Barron, JT. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In: CVPR. 2021,.
- [33] Zhang, X, Srinivasan, PP, Deng, B, Debevec, P, Freeman, WT, Barron, JT. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. arXiv preprint arXiv:210601970 2021;.
- [34] Zhang, K, Luan, F, Wang, Q, Bala, K, Snavely, N. Physg: Inverse rendering with spherical gaussians for physics-based material editing and

relighting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021, p. 5453–5462.

- [35] Boss, M, Braun, R, Jampani, V, Barron, JT, Liu, C, Lensch, HP. Nerd: Neural reflectance decomposition from image collections. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2021, p. 12684–12694.
- [36] Boss, M, Jampani, V, Braun, R, Liu, C, Barron, JT, Lensch, HP. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. In: Advances in Neural Information Processing Systems (NeurIPS). 2021,.
- [37] Philip, J, Morgenthaler, S, Gharbi, M, Drettakis, G. Free-viewpoint indoor neural relighting from multi-view stereo. ACM Transactions on Graphics 2021;URL: http://www-sop.inria.fr/reves/Basilic/ 2021/PMGD21.
- [38] Munkberg, J, Hasselgren, J, Shen, T, Gao, J, Chen, W, Evans, A, et al. Extracting Triangular 3D Models, Materials, and Lighting From Images. arXiv:211112503 2021;.
- [39] Li, Z, Yu, TW, Sang, S, Wang, S, Song, M, Liu, Y, et al. Openrooms: An open framework for photorealistic indoor scene datasets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021, p. 7190–7199.
- [40] Roberts, M, Paczan, N. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. CoRR 2020;abs/2011.02523. URL: https://arxiv.org/abs/2011.02523. arXiv:2011.02523.
- [41] Jakob, W. Mitsuba renderer. 2010.
- [42] Reality, C. Realitycapture reconstruction software. https://www. capturingreality.com/Product; 2018.
- [43] Cook, RL, Torrance, KE. A reflectance model for computer graphics. ACM Trans Graph 1982;1(1):7–24. URL: https://doi.org/10. 1145/357290.357293. doi:10.1145/357290.357293.
- [44] Barron, JT, Poole, B. The fast bilateral solver. In: Leibe, B, Matas, J, Sebe, N, Welling, M, editors. Computer Vision – ECCV 2016. Cham: Springer International Publishing. ISBN 978-3-319-46487-9; 2016, p. 617–632.
- [45] Ronneberger, O, Fischer, P, Brox, T. U-net: Convolutional networks for biomedical image segmentation. In: Navab, N, Hornegger, J, Wells, WM, Frangi, AF, editors. Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. Cham: Springer International Publishing. ISBN 978-3-319-24574-4; 2015, p. 234–241.
- [46] Kettunen, M, Härkönen, E, Lehtinen, J. E-lpips: Robust perceptual image similarity via random transformation ensembles. 2019. arXiv:1906.03973.
- [47] Wang, J, Ren, P, Gong, M, Snyder, J, Guo, B. All-frequency rendering of dynamic, spatially-varying reflectance. In: ACM SIG-GRAPH Asia 2009 Papers. SIGGRAPH Asia '09; New York, NY, USA: Association for Computing Machinery. ISBN 9781605588582; 2009,URL: https://doi.org/10.1145/1661412.1618479. doi:10. 1145/1661412.1618479.
- [48] Parker, SG, Bigler, J, Dietrich, A, Friedrich, H, Hoberock, J, Luebke, D, et al. Optix: a general purpose ray tracing engine. Acm transactions on Graphics (TOG) 2010;29(4):1–13.
- [49] Paszke, A, Gross, S, Massa, F, Lerer, A, Bradbury, J, Chanan, G, et al. Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems 32. Curran Associates, Inc.; 2019, p. 8024–8035.
- [50] Kingma, DP, Ba, J. Adam: A method for stochastic optimization. In: Bengio, Y, LeCun, Y, editors. 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings. 2015,URL: http://arxiv.org/abs/ 1412.6980.
- [51] Karsch, K, Hedau, V, Forsyth, D, Hoiem, D. Rendering synthetic objects into legacy photographs. In: Proceedings of the 2011 SIGGRAPH Asia Conference. SA '11; New York, NY, USA: Association for Computing Machinery. ISBN 9781450308076; 2011,URL: https://doi.org/10.1145/2024156.2024191. doi:10.1145/2024156.2024191.
- [52] Karsch, K, Sunkavalli, K, Hadap, S, Carr, N, Jin, H, Fonte, R, et al. Automatic scene inference for 3d object compositing. ACM Trans Graph 2014;33(3). URL: https://doi.org/10.1145/2602146. doi:10.1145/2602146.
- [53] Gardner, MA, Sunkavalli, K, Yumer, E, Shen, X, Gambaretto, E, Gagné, C, et al. Learning to predict indoor illumination from a single image. ACM Trans Graph 2017;36(6). URL: https://doi.org/10.

1145/3130800.3130891. doi:10.1145/3130800.3130891.

- [54] Loza, A, Mihaylova, L, Canagarajah, N, Bull, D. Structural similaritybased object tracking in video sequences. In: 2006 9th International Conference on Information Fusion. 2006, p. 1–6. doi:10.1109/ICIF.2006. 301574.
- [55] Yu, M, Lafarge, F. Finding Good Configurations of Planar Primitives in Unorganized Point Clouds. In: Proc. of the IEEE conference on Computer Vision and Pattern Recognition (CVPR). New Orleans, US; 2022,.
- [56] Bauchet, JP, Lafarge, F. Kinetic Shape Reconstruction. ACM Trans on Graphics 2020;39(5).
- [57] Yuksel, C, Lefebvre, S, Tarini, M. Rethinking texture mapping. In: Computer Graphics Forum; vol. 38. Wiley Online Library; 2019, p. 535– 551.
- [58] Fraboni, B, Iehl, JC, Nivoliers, V, Bouchard, G. Adaptive Multi-view Path Tracing. In: Eurographics Symposium on Rendering (EGSR). 2019,.
- [59] Su, H, Maji, S, Kalogerakis, E, Learned-Miller, EG. Multi-view convolutional neural networks for 3d shape recognition. In: Proc. ICCV. 2015,.
- [60] Kalogerakis, E, Averkiou, M, Maji, S, Chaudhuri, S. 3D shape segmentation with projective convolutional networks. In: Proc. IEEE Computer Vision and Pattern Recognition (CVPR). 2017,.