

Deep Appearance Maps

Maxim Maximov¹, Tobias Ritschel², Mario Fritz¹

¹Max Planck Institute for Informatics, Saarland Informatics Campus

²University College London

Abstract. We propose a deep representation of appearance, i.e. the relation of color, surface orientation, viewer position, material and illumination. Previous approaches have used deep learning to extract classic appearance representations relating to reflectance model parameters (e.g. Phong) or illumination (e.g. HDR environment maps). We suggest to directly represent appearance itself as a network we call a deep appearance map (DAM). This is a 4D generalization over 2D reflectance maps, which held the view direction fixed. First, we show how a DAM can be learned from images or video frames and later be used to synthesize appearance, given new surface orientations and viewer positions. Second, we demonstrate how another network can be used to map from an image or video frames to a DAM network to reproduce this appearance, without using a lengthy optimization such as stochastic gradient descent (learning-to-learn). Finally, we generalize this to an appearance estimation-and-segmentation task, where we map from an image showing multiple materials to multiple networks reproducing their appearance, as well as per-pixel segmentation.

1 Introduction

The visual appearance of an object depends on the combination of four main factors: viewer, geometry, material and illumination. When capturing and processing appearance, we wish to change some of those factors and predict what the new appearance is. This can be achieved using methods ranging from implicit image-based representations [1] to explicit Computer Graphics-like representations [2]. Implicit methods take a couple of photos as input and allow for high quality in a limited set of conditions, but modest flexibility, e.g. , interpolating an image between two photos but not extrapolating to new views. Explicit representations allow for more flexibility when acquiring Phong parameters and HDR illumination maps, etc. [2], but incur substantial acquisition effort, e.g. , taking a large number of calibrated (HDR) photos.

We here take a step “escaping” this axis: we move away from the pixel-based nature of implicit image-based representations into a deep representation, but without reconstructing anything explicit, as we do not target a direct mapping to any explicit reflectance model or illumination either. Still, we show that such a representation can be used to solve actual tasks, such as image synthesis, appearance acquisition and estimation-and-segmentation of appearance.

This is enabled by four contributions:

First, we will introduce a generalization of reflectance maps [3], which we call an appearance map (AM). AMs are powerful enough to represent appearance for varying geometry under varying views. This allows freely changing the viewer and surface geometry, which is not possible for classic reflectance maps that fix the relation of view and illumination (cf. Fig. 1).

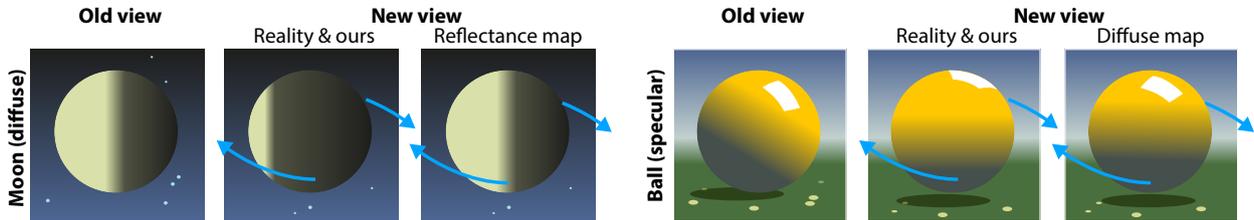


Fig. 1. Comparison of our 4D appearance maps (AM) and a classic 2D reflectance map (RM). This illustrates a view change for the diffuse Moon (**left**) and for a shiny yellow plastic ball (**right**). Rotating the view will expose the dark side of the Moon in reality. This is also possible using our AMs. A classic RM works with camera space normals that do not change for a spherical object when the viewer rotates, producing a wrong result. Even if we were to extend the 2D RM concept that depends on world-space full-sphere normals or reflection directions, this would not work for an object that is both diffuse and specular: changing the view requires a highlight and the diffuse shading to move independently, which is not possible in a 2D RM but only in a 4D AM.

Second, while classic reflectance maps (RM) can be simply represented using a single 2D image, the full appearance is a 4D phenomenon that is more challenging to represent and process. The need to move from 2D to 4D motivates using a deep network to compactly represent appearance. To this end, we suggest using a neural network to represent Deep Appearance Maps (DAM). This representation is efficient, does not require any look-ups and is differentiable. In addition, it can be learned effectively from images or video frames with a known viewer position and surface orientation. Applying this representation to new view positions and surface orientations can be done at speed competitive to classic rendering or RMs, i. e. within milliseconds for full images.

Third, as learning the deep appearance representation (DAM) from an image is not always practical (e. g. it can require long compute time), we devise another deep (convolutional) neural network that can directly map from images to the latent appearance representation directly (learning-to-learn [4–6]). This is facilitated by a neural network that predicts another neural network.

Fourth, the DAM representation can be used for joint material estimation-and-segmentation, a generalization of the previous objective. Here the input is an image with a known number of n materials, and output is n different DAMs, as well as n different appearance segmentation networks that map images to per-pixel-per-material weight maps.

We train and quantitatively test all networks on a new dataset of photo-realistically rendered images.

The dataset introduced in the next section and the network architectures from the methods section will be made publicly available upon publication.

2 Related Work

Inverse Rendering One of the main aims of inverse rendering is to recover material and illumination properties of a scene. It is a quite challenging, ill-posed and under-constrained problem that remains hard to solve for the general case. Recent works on this problem can be roughly divided into data-driven and algorithmic approaches.

Algorithmic methods are based on optimizing appearance properties for a given input [7]. These methods are usually off-line and make simplifying assumptions about the world to reduce computation time and avoid ambiguity and allow for a mathematical derivation. Most recent works [8, 9] use a set of real RGBD images to estimate appearance that are based on a specific illumination model. More refined models use data-based statistical priors to optimize for illumination and reflectance explaining an image [10, 11].

Deep-learning based approaches make a similar assumption as to how humans can recognize materials based on previous experience. Recent work [12–15] uses CNNs to estimate explicit reflectance model parameters. Similarly, illumination is reconstructed, and encoder-decoder CNNs are used to estimate the encompassing reflectance map [16] or complete illumination [17, 13]

All of these methods – data-driven or not – have in common that they rely on a specific illumination model to estimate its explicit parameters (such as Phong diffuse, specular, roughness, etc) and they represent lighting as an HDR illumination map. One of the other limitations of above mentioned CNN methods is limited feedback from a loss function: a change of estimated illumination or reflectance cannot be back-propagated through the image synthesis method. Our method does not involve a renderer, circumventing this problem.

Appearance synthesis Methods to synthesize appearance – or simply “rendering” –, can be classified as simulation-based or image-based.

Simulation-based methods require a complete explicit description of the environment that can be costly and difficult to acquire in practice [2]. A simple, yet powerful, method to represent appearance is a reflectance map [3], a 2D image that holds the color observed for a surface of a certain orientation and material under a certain illumination.

Image-based rendering (IBR) uses a set of 2D images to reconstruct a 3D model and present it in a different view or different light [1]. These methods do geometry prediction, often with manual intervention, with prediction of rendered material on top of it. Recent methods [18, 19] address this problem by using CNN models to predict completely novel views. The method of Rematas et al. [16] and StyLit [20] established a relation between surface and light transport properties and appearance given by photos, respectively, by a hand-painted exemplar, generating images “without rendering”. A simple data-driven approach

to IBR is to learn a per-pixel fully-connected neural network to reproduce per-pixel appearance [21] depending on light. A generalization of this is to shade images with per-pixel positions, normals and reflectance constraints [22]. Wang et al.[23] have used a neural network to approximate the gradient of an inverse-rendering optimization. Our method stems from the same root but neither works on pixel-based image rendering, nor does it reconstruct an explicit appearance model. We will instead use a deep representation of appearance itself.

Segmentation Classic segmentation does not take materials into account [24]. For videos of view-dependent appearance, this is particularly difficult. With adequate capture equipment, spatially varying appearance is captured routinely now [25]. Another issue is multi-materials estimation. Some work [13, 23] has used multiple materials under the same illumination, but they require pre-segmented materials. In our method we perform joint multi-material segmentation and estimation.

Learning-to-learn Learning-to-learn is motivated by the observation that a general optimizer, such as the one to find the internal parameters for a network, will never be much better than a random strategy for all problems [26]. At the same time, intelligent actors can learn very quickly, which obviously does not require a full optimization [4]. We hypothesize, after seeing a material for some time, a human, in particular a trained artist, would be able to predict its appearance in a new condition. This requires the ability to refine the learned model with new observations [5]. For convolutional networks, this was done in dynamic filter networks [6], but we are not aware of applications to appearance modeling, such as we will pursue here.

3 A Multi-view Multi-material Dataset

To work on arbitrary combinations for view, surface orientation and illumination for objects with multiple materials, we first produce a dataset. To our knowledge, no multi-material, multi-view dataset exists that allows for a controlled study. Examples from the dataset are shown in Fig. 2.

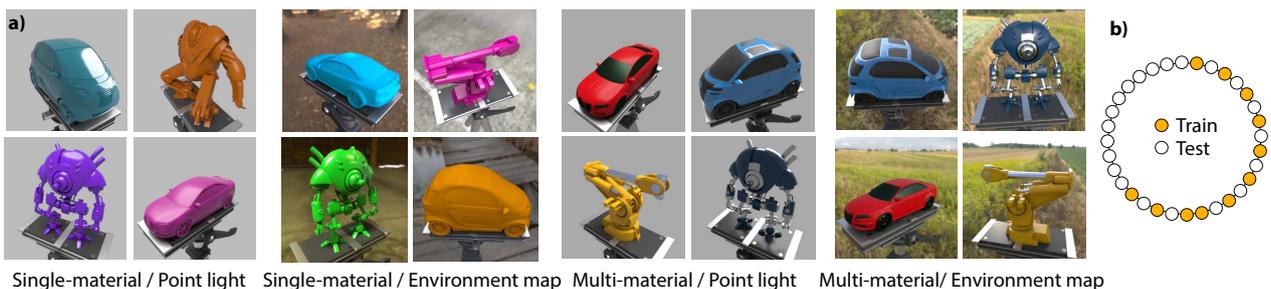


Fig. 2. Four samples from the four variants of our data set (a) and the view layout (b).

An seemingly obvious way to produce such a data set is to take many photos at many exposure settings of many geometric objects under many illumination. Regrettably, this does not scale well to a high number of samples due to the curse of dimensionality encountered with so many parameters (the product of geometry, material, illumination and view). Also it would be difficult to manually decorate them with ground-truth material segmentation. Instead, we suggest to use photo-realistically rendered imagery.

We use acquired 5 production-quality 3D objects from a model repository. Each model was assigned three physically-plausible (layered GGX shading) materials organized on the objects surface in a complex and natural spatial arrangement. Before rendering we randomize the hue of the diffuse component. For illumination, we used 20 different HDR environment maps. For each model, 32 different but equidistant view points on a circle around the object, with a random elevation were used. Overall, this results in $5 \times 20 \times 32 = 3200$ images. Note that the number of photos that would be required for such a dataset is an order of magnitude higher due to the need of capturing multiple exposures. As the views and materials are randomized, no sharing between test, train and validation exists. Geometry might occur both in test and training data. We do not split into test and training as different applications have different requirements in this respect.

For rendering, we used Blender’s [27] Cycles renderer with high-quality settings, including multiple indirect and specular bounces. Note that those light paths violate the model assumption. We add a virtual tripod to emulate the typical local reflections that validate the model assumptions of distant illumination. The resulting images were linearly tone-mapped such that the .95 percentile maps to 1 and kept linear (non-gamma corrected). For each image I in the corpus \mathcal{I} , we store many channels: appearance as RGB colors I_c , position I_p , normals I_n , and a weight map I_w with n channels, where n is the number of materials.

Additionally to the ENVMAP version, we produced a variant with POINTLIGHT illumination (technically, a single, small area light) and split the set in to flavors: MULTIMATERIAL and SINGLEMATERIAL. Using a single material, the material segmentation is sacrificed and one random material from the objects is assigned to the entire 3D objects. In the multi-material case, we proceed directly. Note, that such instrumentation would not be feasible on real photographs.

4 Deep Appearance Processing

4.1 Appearance model

We model RGB appearance L_o of a specific material f_r under a specific distant illumination L_i as a mapping from absolute world-space surface orientation \mathbf{n} and viewer direction ω_o (jointly denoted as \mathbf{x}) as in

$$L_o(\underbrace{\omega_o, \mathbf{n}}_{=\mathbf{x}}) = \int L_i(\omega_i) f_r(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle^+ d\omega_i.$$

Essentially, L_o is a six-dimensional function. In the following, we denote the two three-dimensional parameters – outgoing direction ω_o and surface orientation \mathbf{n} – as a joint parameter vector \mathbf{x} . The concept is visualized in Fig. 3: In a classic reflectance map, the normals vary (blue arrows), but the view direction is the same (orange arrows). In our generalization, both view and normals vary arbitrarily. We might even observe the same normal under different views. Classic reflectance maps [3], assume a view direction \mathbf{z} along the z axis and hold the relation of light and surface fixed, while also being limited to a single half-space:

$$L_{\text{RM}}(\mathbf{n}) \text{ where } \langle \mathbf{n}, \mathbf{z} \rangle \leq \frac{\pi}{2}.$$

Covering the 4D sphere is motivated by our applications that include freely changing the relation of view and surface orientation and observations are made covering the full world space.

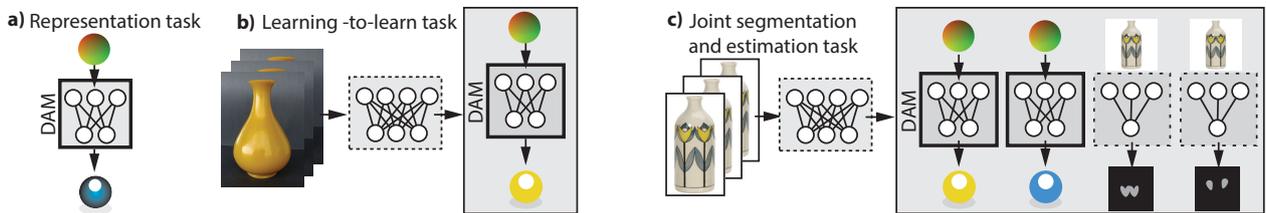


Fig. 4. Different appearance processing tasks that we address using our deep appearance maps. **a)** The first task simply reproduces a given appearance, i. e. it maps from normal and view directions to RGB values using a NN. **b)** In a learning-to-learn task a network maps an image to a DAM representation. **c)** Finally, in the segmentation-and-estimation task, a network maps an image to multiple DAMs and multiple segmentation networks.

4.2 Deep Appearance Representation

We use a deep neural network $\hat{L}_o(\mathbf{x}|\theta)$ to approximate $L_o(\mathbf{x})$ where θ denotes the networks internal parameters (Fig. 5, a). Input to such a network is the surface orientation and viewer direction parametrized as Euclidean vectors, so six numbers. This is followed by five fully-connected layers that are ultimately combined into a single RGB output value.

Here, stochastic gradient descent (SGD) is used to minimize

$$\arg \min_{\theta, \delta} c_d(\theta, W) + \alpha c_a(\theta, \delta)$$

according to the α -weighted sum of a data cost c_d that depends on the DAM model parameters and an adversarial cost c_a that also depends on the model

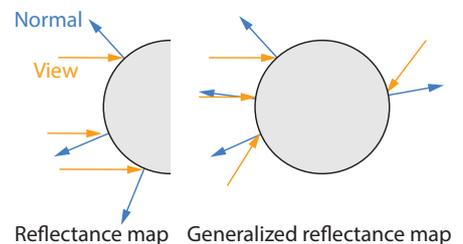


Fig. 3. Appearance map idea.

on all problems is much better than guessing [26]. However, we do not ask for “free lunch” here as we know that we do not want to solve all problems, but a specific subset of problems: learning how to map normals and view directions to appearance.

To this end, we employ a convolutional neural network $\Theta(I|\phi)$ that is executed on an image I with a second set of internal parameters ϕ . It replaces a general learning algorithm. This network can be efficiently deployed and executed on arbitrary images to produce a deep representation that can then be used for synthesis. In this sense, it is a network that predicts a new network.

Input to the network is a 256×256 RGB image I showing appearance of a single material Fig. 5, c. Using an image instead of a plain list of samples allows the network to reason from the spatial arrangement of values, e. g. detecting the shapes and relations of highlights.

Output is a 3360-dimensional vector $\Theta(I|\phi)$ that describes the internal parameters of a network producing the appearance of I . The network Θ has eight layers, reducing resolution until arriving at a fully-connected layer.

Training now minimizes for

$$\arg \min_{\phi, \delta} \sum_{I \in \mathcal{I}} c_d(\Theta(I|\phi), \mathbf{w}) + \alpha c_a(\Theta(I|\phi)|\delta), \quad (3)$$

i. e. the same cost as in the previous section, just that it is define on the parameters ϕ of a network $\Theta(I|\phi)$ to produce another network instead of the network parameters θ itself.

4.4 Multi-Material Appearance Estimation-and-Segmentation

A second application is joint appearance estimation and segmentation. Instead of holding a segmentation fix and estimating an appearance model for each segment or assuming an appearance to apply a segmentation, we jointly do both. We suggest using stochastic gradient descent (SGD) itself as an optimization method to find the segmenting and appearance-predicting networks. Here, the deep appearance representation network as well as a segmentation network are used as the latent variables to be inferred. The number of materials n is assumed to be known.

The appearance network parameters for all appearances are stacked into a matrix $\Theta(I) = (\Theta(I|\phi_1), \dots, \Theta(I|\phi_n))^T$.

Instead of inferring a per-pixel segmentation mask in the optimization, we suggest to learn a network $\Psi(\psi_i)$ with parameters ψ_i that produces the segmentation for a specific material i (Fig. 5, d). This significantly reduces the number of latent parameters per material from number of pixels to number of degrees of freedom in the network. This network again is a simple encoder-decoder with skip connections, where the encoder is shared among the materials in order to further reduce parameters. Input to this network is an image I with pixel color, normal, position, and output is a weight mask expressing how much a pixel belongs to a certain material i . There is one segmentation network parameter for each material i , and they are all stacked into a matrix $\Psi(I) = (\Psi(I|\psi_1), \dots, \Psi(I|\psi_n))^T$.

The optimization is then

$$\arg \min_{\Theta, \Psi, \delta_a, \delta_m} \sum_{i=1}^m c_d(\Theta(I|\phi_i), \Psi(I|\psi_i)) + \alpha c_a(\Theta(I|\phi_i), \delta_a) + \beta c_s(W). \quad (4)$$

Here, c_s is a sparsity term on the weight mask W that encourages a solution where most values for one pixel are zero, except one. For every channel \mathbf{w} in W it is $\sum_i \text{abs}(\mathbf{w}_i - .5)$.

5 Results

We report quantitative and qualitative results of our approach.

5.1 Protocol

Here we evaluate our deep appearance representation (Sec. 5.2), as well as their application to learning-to-learn appearance (Sec. 5.3) and joint material estimation-and-segmentation (Sec. 5.4).

Instrumentation for all tasks is performed in a similar fashion using our multi-view, multi-material data set (Sec. 3). In particular, we consider its POINTLIGHT and ENVIRONMENTMAP variants. Depending on the task, we either use a SINGLEMATERIAL or MULTIMATERIAL. The main quantity we evaluate is image similarity error (DSSIM, less is better) in respect to a path-traced reference. We consider two tasks: re-synthesizing from the SAMEVIEW as well as from a NOVELVIEW. We will use 10 of the 32 views for every sample as input and predict 22 novel views. The 20 views form a random but consecutive range of angles i. e. ca. 240 degree as detailed in Fig. 2, b.

In each application we also consider one application-specific competitor to solve the task: classic reflectance maps for appearance representation [19], SGD as a means to learn appearance and k -means clustering in RGB-Normal space to do joint material estimation-and-segmentation. We will now look into the three specific applications.

5.2 Appearance representation

Here we study how well our approach can represent appearance per-se. Most distinctly, we propose to use a 4D appearance map while other work has used 2D image representations of a reflectance map. To quantify the difference, we represent the SINGLEMATERIAL variant of our dataset as a common reflectance map, as well as using our appearance map.

To emulate common reflectance map, which is defined in view space, we take use the closest input image as a source image. Every normal in the new view is converted to camera space of the new view and the same is done for the normal in the old view. We then copy the RGB value from the old view image to the new-view image that had the most similar normal. Note, that such a multi-view

extension of RMs already is more than the state of the art that would use a single view.

Tbl. 1, top part, shows results as mean error across the data set. We see that for all data sets our method is a better representation in terms of having a lower DSSIM error. The difference in error is more pronounced for NOVELVIEW than for SAMEVIEW. A detailed plot of error distribution is seen in Fig. 7, left. This is, as classic reflectance map captures appearance for a fixed viewer location, for changing geometry, but does not generalize when the viewer moves. Arguably, classic RMs look plausible without a reference, but do not have much similarity in cases such as diffuse shading.

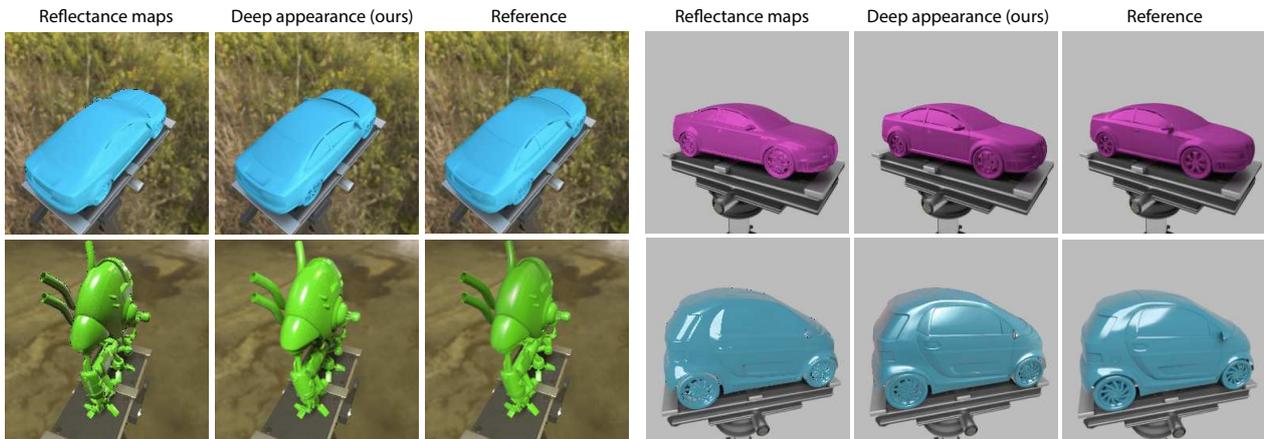


Fig. 6. Results of reflectance maps (**1st and 4th column**), our deep appearance representation (**2nd and 5th column**) and the reference (**3rd and 6th column**) for a novel view. We see that all methods look plausible, hence the popularity of RMs, but that ours is much closer to the reference, thanks to the ability to handle changes of view.

Qualitative results are seen in Fig. 6. We see that all methods produce plausible results. When looking into the actual placement of certain visual features (e.g. the highlights), we notice that while they are present in all methods, they appear in the wrong place in reflectance maps. Typical failure modes are show in Fig. 10.

5.3 Learning to learn appearance models

Here, we also follow the protocol described in Sec. 5.1. After having established the superiority of appearance maps to classic reflectance maps in the previous section, the competitor here is SGD. At best, our learning-to-learn network produces a network which is as good a running a full SGD pass.

To train the network Φ , we split the SINGLEMATERIAL variant of our dataset in 80 % train, 10 % test and 10 % validation. The middle part of Tbl. 1 summarizes the outcome when executing the resulting ϕ on the validation data set. We see that both approaches reproduce the appearance faithfully. For point lights, the mean DSSIM is .144 for SGD while it is .165 for network-based (Tbl. 1, middle

Table 1. Quantitative results. Different rows are different combination of tasks and methods (three applications, two view protocols, our two methods). Columns are different data. Error is measured as mean DSSIM across the data set (less is better).

Task	View	Method	Error (DSSIM / Percentage)			
			POINT	ENVMAP	POINT	ENVMAP
Representation (Sec. 4.2)	Same	OUR	.105	.123		
		RM	.143	.160		
	Novel	OUR	.144	.164		
		RM	.181	.193		
Learning-to-learn (Sec. 4.3)	Same	OUR	.106	.131	-0.9 %	-1.0 %
		SGD	.105	.123		
	Novel	OUR	.165	.173	-14.1 %	-5.0 %
		SGD	.144	.164		
Segmentation (Sec. 4.4)	Same	OUR	.113	.122		
		kMEANS	.132	.136		
	Novel	OUR	.161	.154		
		kMEANS	.172	.164		

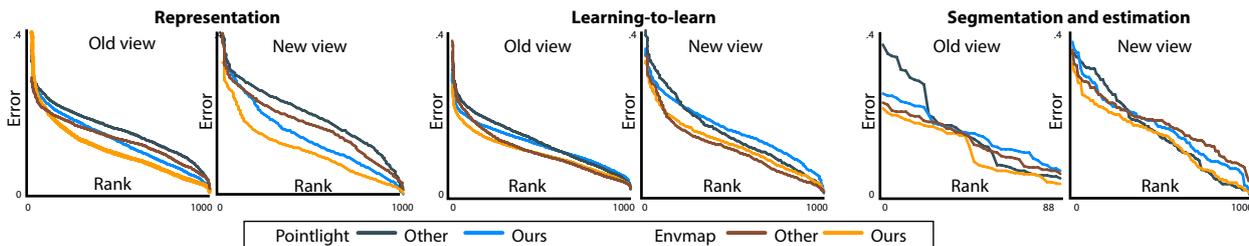


Fig. 7. Error plots. Each pair of plots is one task. In each pair, the first is the old and the second the new view. Each curve is produced by sorting the DSSIM (less is better) of all samples in the data set. Blue colors are for point light illumination, red colors for environment maps. Dark hues are the competitor (RMs for representation, SGD for Learning-to-learn and k -means for segmentation and estimation) and light hues ours.

part). Naturally, letting a network do the learning degrades quality, but only by a marginal amount, in this case 14.1 %. For environment map illumination, the mean DSSIM is increased from .164 to .173, a decrease by only 5 %. While being marginally worse, it is two orders of magnitude faster. Fig. 7 show the error distribution across the data set.

A visual comparison is found in Fig. 8. We see that replacing the SGD computation of several minutes by a network, can produce a DAM that is qualitatively similar to both the SGD’s result as well as to the reference. Overall, strength and sharpness of highlights that is already challenging for DAM per-se, seems to suffer a bit more by learning-to-learn, as also seen in Fig. 10.

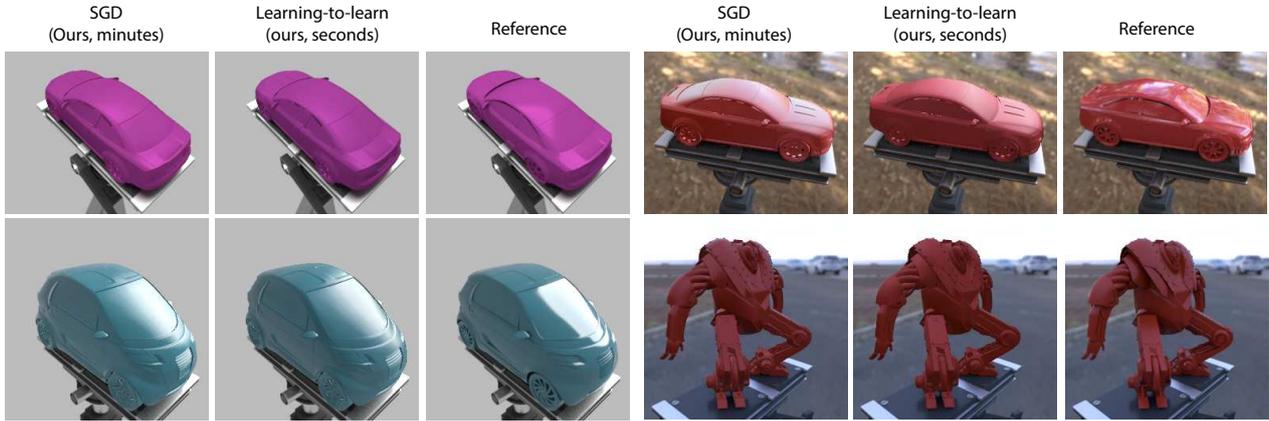


Fig. 8. Results of our DAM representation trained using stochastic gradient descent (**1st and 3rd row**), our DAMs produced by our learning-to-learn network (**2nd and 4th column**) as well as a reference (**3rd and 6th column**) in a novel-view task.

5.4 Joint Multi-Material Estimation-and-Segmentation

Finally, we quantify the joint material-and-segmentation task from Sec. 4.4. We perform the same split as in the previous section, just now on the MULTIMATERIAL variant. For the re-synthesis to new views we use the ground truth segmentation in the new view (our method only produces the segmentation in all old views).

We here compare to a competitor, where the image is first segmented using k -means clustering on normals and RGB (same weight, as both have a similar range) and material is estimated for each segment in consecution.

Tbl. 1, bottom part shows the quantitative results and Fig. 9 the qualitative outcome. On average, we achieve an DSSIM error of .133 for POINTLIGHT and .122 for ENVIRONMENTMAP. The greedy method performs worse (.161 and .154), likely as it segments highlights into individual parts. While it can understand that highlights belong “to the rest” of a material, sometimes they end up in different clusters, as seen in Fig. 10, right.

5.5 Example application: Denoising and Super-resolution

We have formally introduced and analyzed three tasks making use of DAMs above. However, we believe DAMs not not be limited to these, and would argue representing appearance itself as a network allows for interesting applications.

An example of further applications is de-noising of Monte-Carlo rendered images. This works as follows: Input to the representation task is just a number of unstructured samples of normal, view direction and RGB pairs. Therefore, we can also insert noisy path-traced pixels, either from the full image or from a (weighted) local neighborhood. The procedure will fit a NN \hat{L} to these observations that reproduces the data, yet is non-suspicious to the adversarial. This is similar to image filters that fit a linear or higher-order model to data and the re-evaluate it for smoothing, just that the model has become a neural network that is a function of normal and view direction, not pixel position. Other modern MC denoisers use

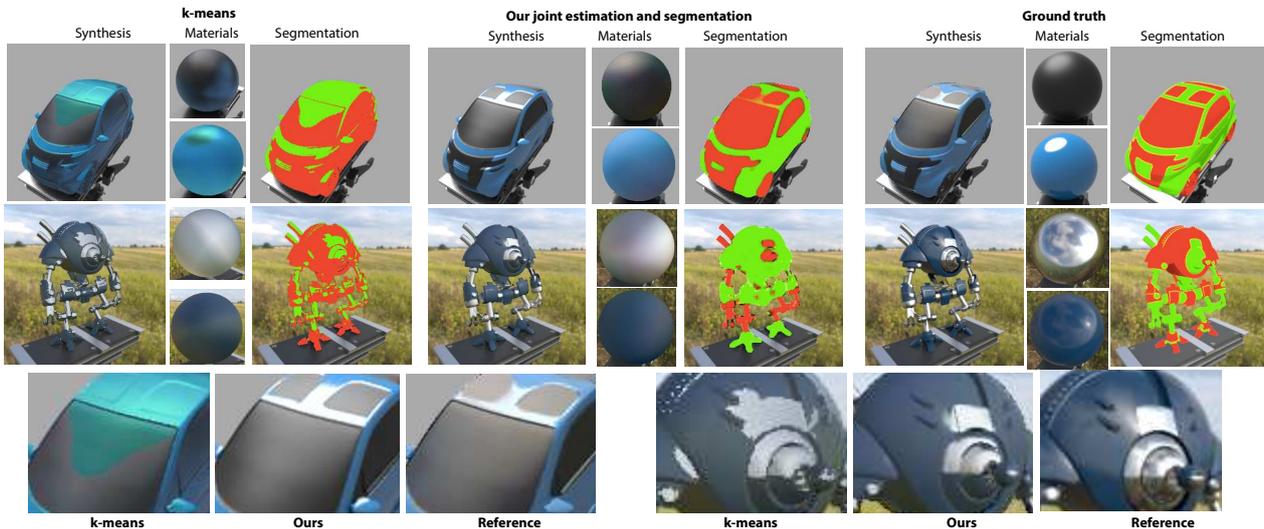


Fig. 9. Results of joint material segmentation and estimation for two samples (rows). The left part shows a greedy approach, first clustering with k means and second running our estimation. The middle part show our joint approach. The right part shows the ground truth. In every part we show a re-synthesis, as well as two estimated materials and the resulting mask. The insets in the last row show that, while not all reflection details are reproduced, ours is free of color shifts around the highlights and mostly a low-frequency approximation of the environment reflected.

deep networks, but to produce colors directly [28, 29]. Our NN \hat{L} will smooth the data, but it will preserve 4D-directional dependence. Also it is not by definition band-limited such as a linear basis like spherical harmonics: it will preserve small highlights that correctly respond to view or surface orientation changes. Consequently, it can be evaluated at arbitrary normals and view directions, including the ones it was fit to, but also for a slightly different image, or for an image in a higher resolution.

In Fig. 11, we show de-noising results: The DAM is fit to a noisy path-traced image and when re-evaluating it, the noise is disappeared, but still the highlights and shading details are preserved. In a more refined version, a DAM would be fit to local patches.

While an in-depth comparison to other de-noising alternatives is required to understand its full potential, it indicates many more applications of DAMs exist to explore. Similar applications would be possible when fitting a DAM to noisy measurements.

6 Conclusion

We have explored a novel take on appearance processing that neither works on pixel-level IBR-like representations nor by extracting classic explicit reflectance and illumination parameters. Instead, we work on a deep representation of appearance itself, defined on a generalization of reflectance maps that works in world space where observations cover all directions, such as in videos.

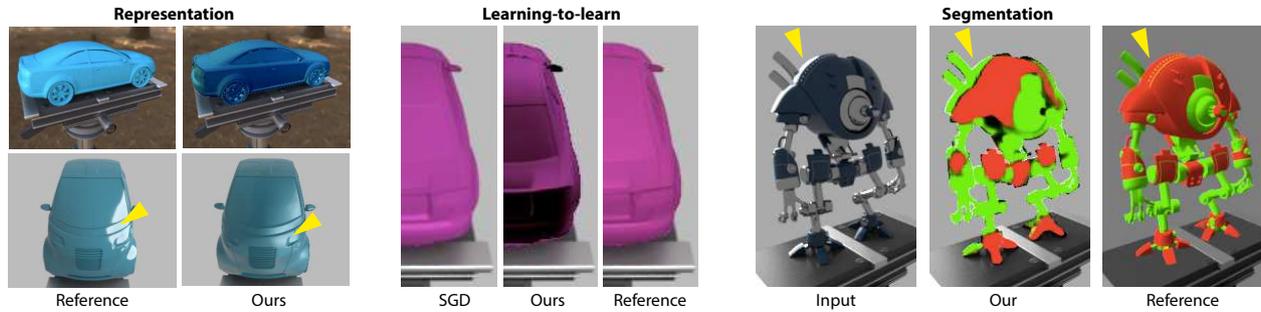


Fig. 10. Failure modes for all three tasks. For representation (**right**), the network can overshoot e. g. become darker than desired, for unobserved directions (back of the top car). More input images or a more effective GAN could suppress this. Sharp details cannot be represented by the cascade of functions of a small network. Fitting a network with more parameters might be required. For learning-to-learn (**middle**), SGD might produce the right network, but the learned network overshoots. Similarly, highlights tend to be more blurry (not shown). For segmentation (**right**) the rim highlight in the back of the character is purely white and apparently does not look enough like other highlights on blue to be understood. Consequently, it is assigned the metallic material, which is incorrect.

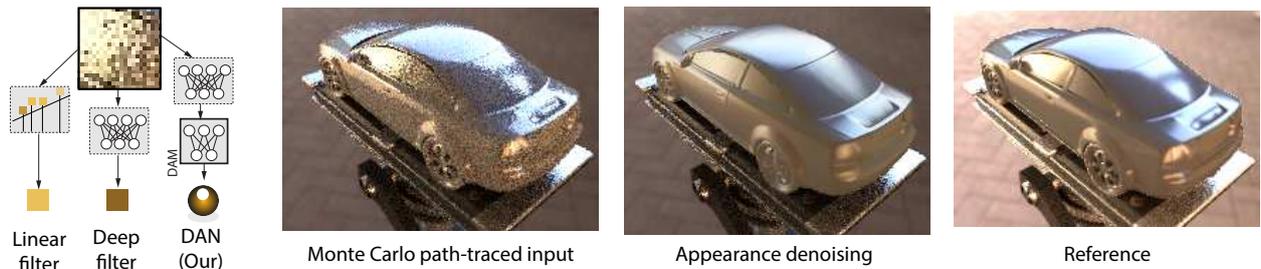


Fig. 11. Denoising idea (left) and results (right). Linear filters use a linear model to predict a new color from observed noisy Monte Carlo (MC) samples. Modern deep de-noisers [28, 29] use a network to predict a new color. Our approach uses a network to predict a DAM, that can be used to shade with any normal or view. The right part shows a noisy path-traced input image, our de-noising and a reference.

In future work, we would like to generalize this notion to form that allows independent control of illumination and reflectance, providing better material editing and relighting. We have shown this enables effective reproduction, estimation by learning-to-learn and joint material estimation-and-segmentation.

References

1. Debevec, P.E., Taylor, C.J., Malik, J.: Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In: Proc. SIGGRAPH. (1996) 11–20
2. Pharr, M., Jakob, W., Humphreys, G.: Physically based rendering: From theory to implementation. Morgan Kaufmann (2016)
3. Horn, B.K., Sjoberg, R.W.: Calculating the reflectance map. *Applied optics* **18**(11) (1979) 1770–9
4. Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S.J.: Building machines that learn and think like people. *Behavioral and Brain Sci.* **40** (2017)
5. Thrun, S., Pratt, L.: Learning to learn. Springer Science & Business Media (2012)
6. Jia, X., De Brabandere, B., Tuytelaars, T., Gool, L.V.: Dynamic filter networks. In: NIPS. (2016) 667–675
7. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. In: SIGGRAPH. (2001) 117–128
8. Wu, H., Wang, Z., Zhou, K.: Simultaneous localization and appearance estimation with a consumer RGB-d camera. *IEEE TVCG* **22**(8) (2016) 2012–2023
9. Richter-Trummer, T., Kalkofen, D., Park, J., Schmalstieg, D.: Instant mixed reality lighting from casual scanning. In: ISMAR. (2016)
10. Matusik, W., Pfister, H., Brand, M., McMillan, L.: A data-driven reflectance model. In: SIGGRAPH. (2003)
11. Lombardi, S., Nishino, K.: Reflectance and natural illumination from a single image. In: ECCV. (2012) 582–95
12. Meka, A., Maximov, M., Zollhöfer, M., Chatterjee, A., Richardt, C., Theobalt, C.: Live intrinsic material estimation. *arXiv:1801.01075* (2018)
13. Georgoulis, S., Rematas, K., Ritschel, T., Fritz, M., Tuytelaars, T., Gool, L.V.: What is around the camera? In: ICCV. (2017)
14. Liu, G., Ceylan, D., Yumer, E., Yang, J., Lien, J.M.: Material editing using a physically based rendering network. In: ICCV. (2017)
15. Kim, K., Gu, J., Tyree, S., Molchanov, P., NieBner, M., Kautz, J.: A lightweight approach for on-the-fly reflectance estimation. In: ICCV. (2017)
16. Rematas, K., Ritschel, T., Fritz, M., Gavves, E., Tuytelaars, T.: Deep reflectance maps. In: CVPR. (2016)
17. Hold-Geoffroy, Y., Sunkavalli, K., Hadap, S., Gambaretto, E., Lalonde, J.F.: Deep outdoor illumination estimation. In: CVPR. Volume 1. (2017) 6
18. Zhou, T., Tulsiani, S., Sun, W., Malik, J., Efros, A.A.: View synthesis by appearance flow. In: ECCV. (2016) 286–301
19. Rematas, K., Nguyen, C.H., Ritschel, T., Fritz, M., Tuytelaars, T.: Novel views of objects from a single image. *IEEE PAMI* **39**(8) (2017) 1576–1590
20. Fišer, J., Jamriška, O., Lukáč, M., Shechtman, E., Asente, P., Lu, J., Sýkora, D.: StyLit: Illumination-guided example-based stylization of 3d renderings. *ACM Trans. Graph.* **35**(4) (2016)
21. Ren, P., Dong, Y., Lin, S., Tong, X., Guo, B.: Image based relighting using neural networks. *ACM Trans. Graph.* **34**(4) (2015) 111
22. Nalbach O, Arabadzhyska E, M.D.S.H.R.T.: Deep shading: Convolutional neural networks for screen space shading. *Comp. Graph. Forum (Proc. EGSR)* **36**(4) (2017)
23. Wang TY, Ritschel T, M.N.: Joint material and illumination estimation from photo sets in the wild. *arXiv abs/1710.08313* (2018)

24. Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE PAMI* **22**(8) (2000) 888–905
25. Lensch, H., Kautz, J., Goesele, M., Heidrich, W., Seidel, H.P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.* **22**(2) (2003) 234–257
26. Wolpert, D.H., Macready, W.G.: No free lunch theorems for search, Santa Fe Institute (1995)
27. Blender Foundation: Blender - a 3D modelling and rendering package. (2018)
28. Bako, S., Vogels, T., McWilliams, B., Meyer, M., Novák, J., Harvill, A., Sen, P., DeRose, T., Rousselle, F.: Kernel-predicting convolutional networks for denoising monte carlo renderings. *ACM Trans. Graph. (Proc. SIGGRAPH 2017)* **36**(4) (July 2017)
29. Chaitanya, C.R.A., Kaplanyan, A.S., Schied, C., Salvi, M., Lefohn, A., Nowrouzezahrai, D., Aila, T.: Interactive reconstruction of monte carlo image sequences using a recurrent denoising autoencoder. *ACM Trans Graph. (Proc. SIGGRAPH)* **36**(4) (2017) 98

Deep Appearance Maps

— Supplemental Document —

Maxim Maximov¹, Tobias Ritschel², Mario Fritz¹

¹Max Planck Institute for Informatics, Saarland Informatics Campus

²University College London

In this supplemental document, we show and explain additional results. More specifically, we show more results on a single representation (see Fig. 1), failure cases (see Fig. 2), Learning2Learn (see Fig. 3), joint multi-material segmentation and estimation (see Fig. 5, Fig. 6), and de-noising (see Fig. 4)

Additionally, we included a short video clip ("DAMreconstruction.mp4") and local HTML-page ("DAMresults.html") with results on several appearances with different dataset options. Due to size constraints, we have compiled only a subset of views and appearance sets for this supplementary material.



Fig. 1. Results of reflectance maps (1st and 4th column), our deep appearance representation (2nd and 5th column) and the reference (3rd and 6th column) for a novel view. Our approach performs quite well with diffuse materials and with a strong point light. Note that all object are rendered with self-shadowing and global illumination turned on, which leads to some unnaturalness and artifacts in some results. Both approaches produce close to ground-truth results. However, reflectance maps are taken from the ground truth and present the best case that should be able to be estimated.



Fig. 2. Here we present most common failure types in our method. Our approach has difficulties in reconstructing high-frequency information on reflective materials (**1st row**). Small highlights are also not always stored in our network (**2nd row**). Mainly because they generate a smaller error and the network doesn't pick up on it. Similarly, the network doesn't always reconstruct highlights that have very similar hue as the diffuse albedo (**3rd row**). Unobserved sides tend to deviate from the ground truth (**4th and 5th row**).



Fig. 3. Results of our DAM representation trained using stochastic gradient descent (**2nd column**), our DAMs produced by our learning-to-learn (L2L) network (**3rd column**) from the input view (**1st column**) as well as a reference (**4th column**) in a novel-view task. L2L was trained from many sides on many appearances, it has the advantage of having more stable outputs for unobserved views than just a single DAM representation trained only on one specific appearance from few views. However, L2L struggles to reconstruct highlights.

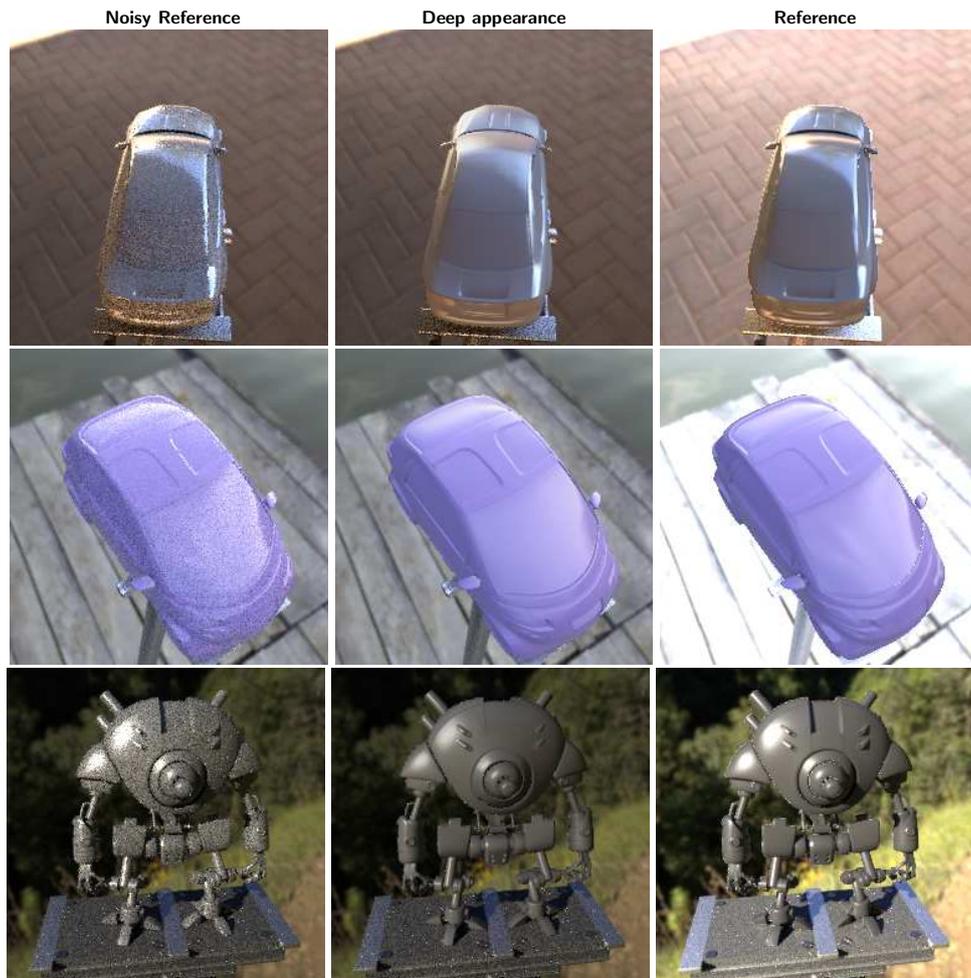


Fig. 4. Results of our DAM representation trained on noisy images (**2nd column**), noisy reference (**1st column**) and clean reference (**3rd column**) in a novel-view task. Noisy images were rendered with 10 samples and clean with 512.

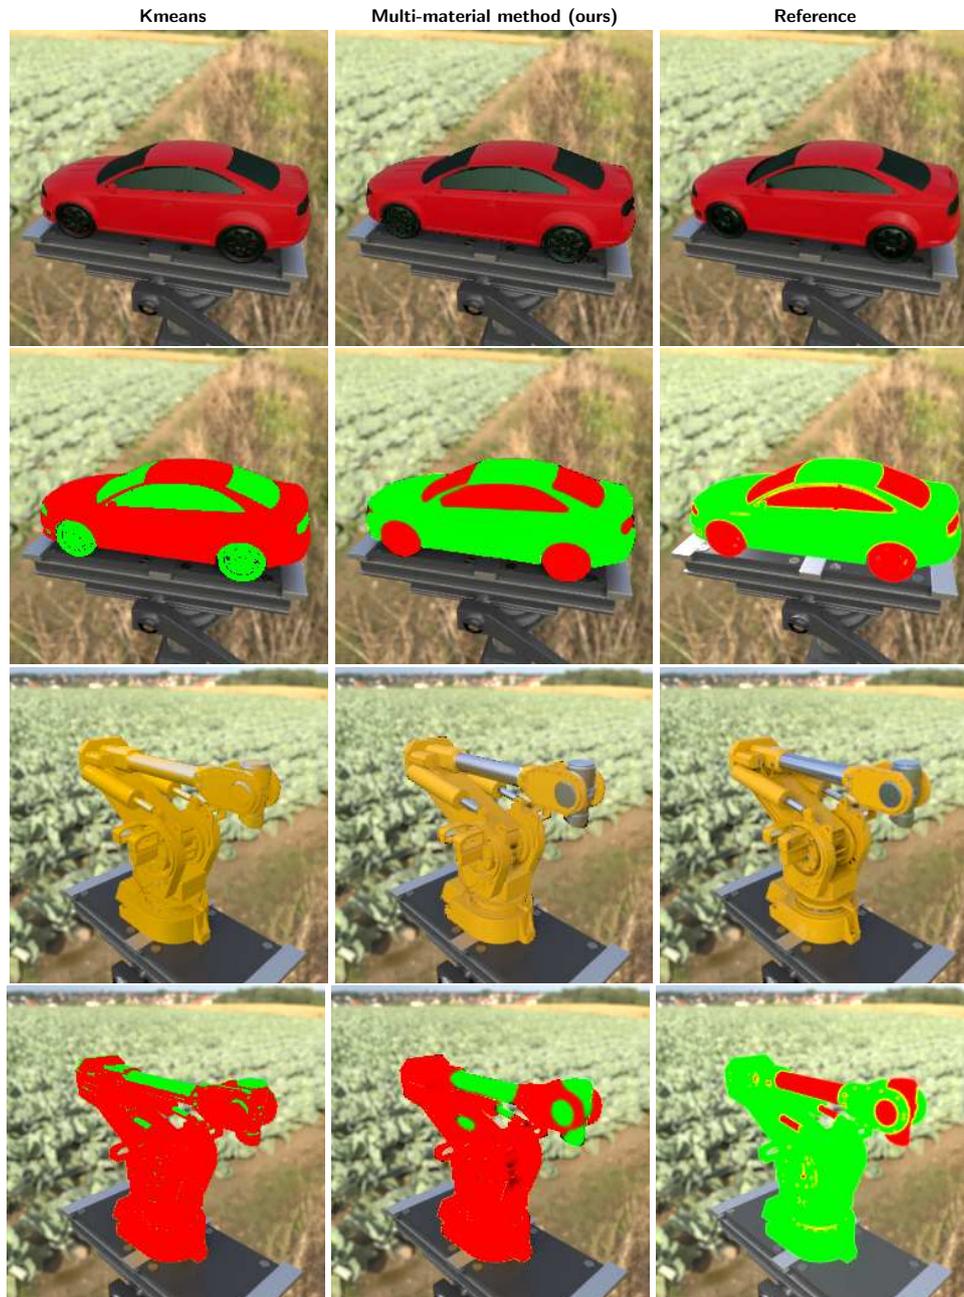


Fig. 5. Results of our joint multi-material segmentation and estimation (**2nd column**), K-means (**1st column**) and reference (**3rd column**) in a novel-view task. We can see reconstruction (**1st and 3rd row**) based on segmentation(**2st and 4rd row**).

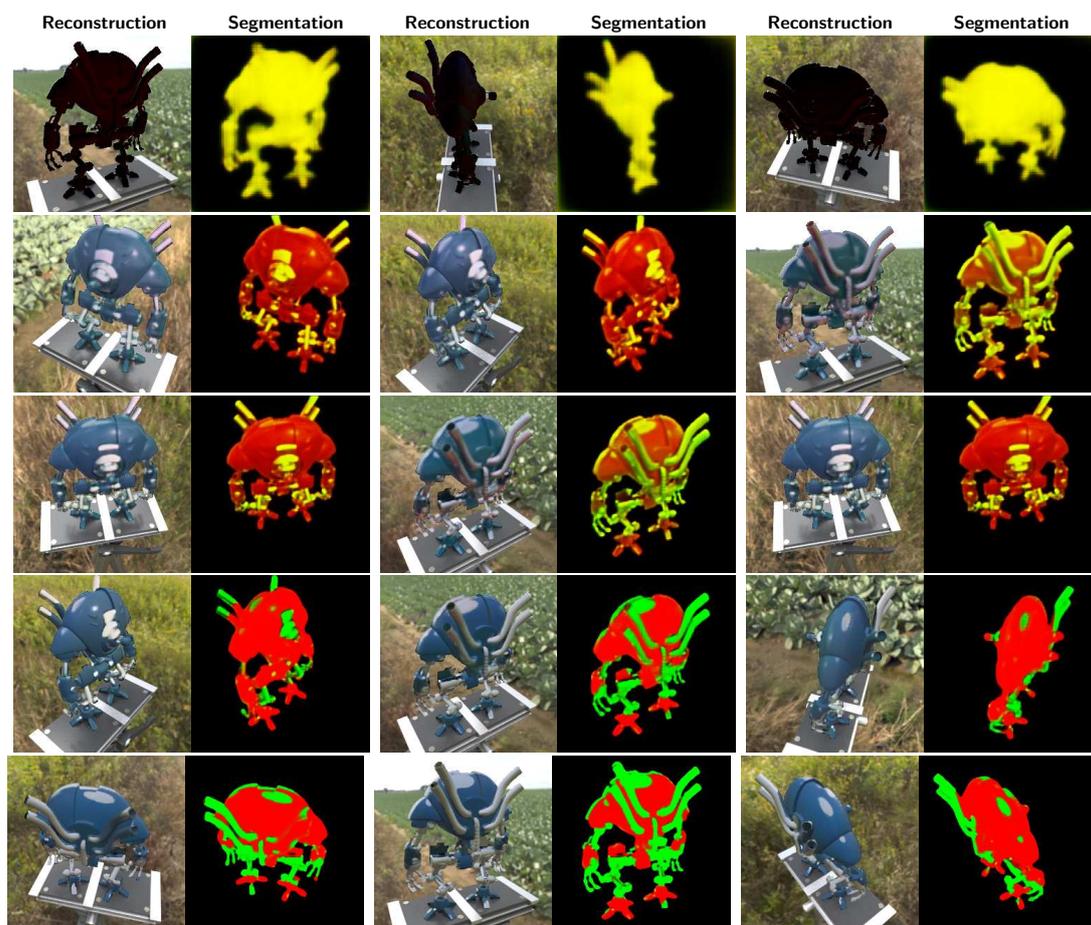


Fig. 6. Progressive multi-material segmentation and estimation process. Each row shows 3 random views of the same object presented in a pair of reconstruction and respective segmentation. All pixels of both mask channels are equal to 1. From there, the model gradually make each pixel to choose certain material based on current estimated materials.