

SuperRVol: Super-Resolution Shape and Reflectance Estimation in Inverse Volume Rendering

Mohammed Brahimi^{1,3} Bjoern Haefner^{1,3} Tarun Yenamandra^{1,3} Bastian Goldluecke² Daniel Cremers^{1,3}
¹ Technical University of Munich, ² University of Konstanz, ³ Munich Center for Machine Learning
 {mohammed.brahimi, bjoern.haefner, tarun.yenamandra, cremers}@tum.de
 bastian.goldluecke@uni-konstanz.de

Abstract

We propose an end-to-end inverse rendering pipeline called SuperRVol that allows us to recover 3D shape and material parameters from a set of color images in a super-resolution manner. To this end, we represent both the bidirectional reflectance distribution function's (BRDF) parameters and the signed distance function (SDF) by multi-layer perceptrons (MLPs). In order to obtain both the surface shape and its reflectance properties, we revert to a differentiable volume renderer with a physically based illumination model that allows us to decouple reflectance and lighting. This physical model takes into account the effect of the camera's point spread function thereby enabling a reconstruction of shape and material in a super-resolution quality. Experimental validation confirms that SuperRVol achieves state of the art performance in terms of inverse rendering quality. It generates reconstructions that are sharper than the individual input images, making this method ideally suited for 3D modeling from low-resolution imagery.

1. Introduction

The reconstruction of 3D shape and appearance is among the classical challenges in computer vision. While we have witnessed significant progress on this task with suitably designed neural representations, the resulting reconstructions of shape and appearance are typically limited by the resolution of the input images where high-quality models invariably require high-resolution input images. At the same time, the concept of super-resolution modeling has been studied intensively in classical variational optimization approaches. The aim of this work is to bring both of these developments together and introduce super resolution modeling into neural differentiable volume rendering approaches in order to allow high-resolution reconstructions of 3D shape and reflectance even from low-resolution input images – see Figure 1 for high quality geometry and super-resolution image reconstruction on two real world datasets.




Figure 1. Given a set of low-resolution input images (left column), SuperRVol recovers the geometry (top) and the material properties (full reconstruction at the bottom). The combination of a realistic physical camera model with inverse volumetric rendering gives rise to reconstructions that are more crisp than competing ones and even sharper than the input images.

More specifically, we consider a setting where we capture images of an opaque non-metallic object from various viewpoints. The so called *photometric* images are illuminated only from a white point light source colocated with the camera. We want to separately recover geometry and different components of an isotropic BRDF. Note that due to the nature of the setup, we can only reconstruct the slice of the BRDF for which illumination direction is equal to viewing direction. However, we will see that we generalize well to novel relighting scenarios. The images are assumed to be calibrated, i.e. camera ex- and intrinsics are given, e.g. from COLMAP [47]; this is similar to [32, 70, 71].

In detail, our contributions are as follows:

- Given a set of photometric images, we put forward an end-to-end inverse rendering approach for jointly estimating high quality shape of arbitrary topology and its corresponding reflectance properties.
- Within the image formation model, we explicitly parameterize the degradation process induced by the camera’s sensor via modeling its point spread function, allowing us to estimate super-resolved shape and material properties.
- In numerous experiments, we validate that the proposed method gives rise to state-of-the-art reconstructions of shape and reflectance. In particular, the reconstructed objects are significantly more detailed and sharper than the individual input images.

2. Related Work

In the following we recall neural inverse rendering and view synthesis as well as super-resolution approaches for 3D reconstruction.

2.1. Neural Inverse Rendering and View Synthesis

Neural approaches for inverse rendering [9,23,24,33,48,51–53,63,74] and novel view synthesis [3,11,18,26,28,32,41,44,50,61] have gained a lot of attention over the recent years. Their expressivity within a lightweight architecture such as a multilayer perceptron (MLP) form a great basis for these complex tasks [27,31,32,43,57,67,72]. However, those approaches can only recover geometry of reduced quality due to their underlying volume rendering based on the scene’s density. In contrast, surface rendering approaches [36,66] perform better when estimating geometry, but require mask supervision and can still get stuck at unsatisfactory local minima with severe geometry artefacts. This limits their usage to relatively simple objects. To get the best out of both worlds, [39,58,65] express the underlying shape implicitly via occupancy fields [39] or SDFs [58,65]. [65] even provide a sampling procedure of the volume integral to theoretically upper bound the opacity error, which is a missing desired property in [39,58]. While [39,58,65] allow to reliably reconstruct accurate geometry of complex scenes without mask supervision, these approaches lack the important ability to estimate other scene properties such as reflectance.

Oftentimes it is desired to recover the object’s BRDF properties along with the geometry, as this allows for relighting the scene under novel illumination. While some works focus on relighting based on static, unknown illumination of the scene [5,49,71,73], others enforce a change of illumination [4,29,34,70] by resorting to photometric

images. In particular, this can lead to a well-posed optimization problem of the underlying scene properties [6].

Density based volume rendering approaches such as [4,5,49,73] also estimate material parameters, but they inherit the same limitations as [32]: inaccurate geometry and the fact that the object’s material parameters are not properly defined on the surface, but everywhere in the volume. This severely limits the editing capability, and a volume rendering step has to be used during inference, making it unusable in common graphics pipelines. [29,34] are mesh-based classical inverse rendering approaches and require masks and a good initialization of the geometry, causing their convergence to be fragile, as mesh-based optimization is inherently non-differentiable at depth discontinuities and difficult to handle if topological changes or self-intersections arise. More recently, [70] proposed an inverse rendering approach which eliminates the disadvantages of mesh-based approaches to some extent due to the use of a neural SDF representation and an edge-aware surface renderer. However, the weak convergence properties of surface rendering made them resort to a two-step approach, where the first step consists of a volume rendering step [58] which is used to initialize the second step based on surface rendering.

Compared to prior works using photometric images, our approach is based solely on volume rendering using SDFs, causing reliable geometry and reflectance reconstruction without resorting to other rendering techniques such as surface rendering, thus avoiding a multi-step pipeline. Additionally, thanks to our novel problem formulation our method is still applicable for standard graphics pipelines, as a mesh and each surface point’s reflectance property can be easily recovered. Next, we discuss the state-of-the-art in super-resolution for 3D reconstruction and how we leverage that to further improve our methodology.

2.2. Super-Resolution for 3D Reconstruction

The problem of super-resolution (SR) has been extensively studied in the past [1,35,40,54,55,60,64,68]. Different problem statements of SR lead to different approaches, e.g. the case of single image SR [12,19,38], video SR [7,8,20], or depth SR [16,45,46,56,62]. Given that we are interested in a 3D reconstruction of the scene from a set of photometric images, we do not perform SR in 2D image space, but in 3D scene space [13,30,57]. [57] are able to synthesize images of higher resolution than the individual input images by resorting to supersampling, i.e. a low-resolution pixel is sampled at each super-resolution pixel’s center, allowing for a denser sampled radiance field. [13,30] are classical approaches optimizing SR geometry and textures. While [30] integrate low-resolution depth and color from an RGB-D sensor into SR keyframes and fuse these keyframes into a texture map, [13] describe their SR process using a convolution with a Gaussian kernel. This is a well

motivated image formation model of a camera sensor element and straightforward to carry out as they work in a discrete pixel grid. However, this is not easily applicable in a continuous case, i.e. when using neural networks to implicitly express shape and reflectance. While [13, 30, 57] share the benefit of increasing the input resolution of the individual input images to result in a sharper, high detailed output, they all lack the possibility of representing the scene’s intrinsic properties, i.e. shape and material. Either the full scene is represented in a neural network [57] or the estimated textures consist of lighting cues baked in to the reflectance properties, making faithful relighting impossible [13, 30].

Contrary to the existing SR works mentioned here, we mathematically formulate the camera’s image formation process in the continuous case leading to a principled sampling heuristic applicable for neural approaches which allows us to invert the camera’s image formation model resulting in reconstructions beyond the input in terms of resolution and quality. Additionally, our model can recover super-resolved geometry and BRDF parameters, free from baked in lighting cues, allowing faithful photorealistic reconstructions with full control over the scene’s properties.

3. Preliminaries: VolSDF

VolSDF [65] leverages volume rendering similarly to NeRF [32], however aims at overcoming certain limitations of NeRF by decoupling geometric representation and appearance. To this end, VolSDF models the scene geometry within a volume $\Omega \subset \mathbb{R}^3$ by means of a density $\sigma : \Omega \rightarrow \mathbb{R}_{\geq 0}$, which is related to its signed distance function (SDF) $d : \Omega \rightarrow \mathbb{R}$ by the transformation

$$\sigma(\mathbf{x}) = \alpha \Psi_{\beta}(-d(\mathbf{x})). \quad (1)$$

Here, Ψ_{β} is the Cumulative Distribution Function of the Laplace distribution with zero mean and scale β , and both $\alpha, \beta > 0$ are learnable parameters. Given this parameterization of the density in terms of the underlying SDF, we can set up the volume rendering equation to obtain radiance \mathcal{L}_p at a pixel $p \in \mathbb{R}^2$ within the image of a camera located at $\mathbf{c} \in \mathbb{R}^3$. Let $\mathbf{v} \in \mathbb{S}^2$ be the viewing direction from \mathbf{c} through p , then

$$\mathcal{L}_p = \int_0^{\infty} w(t) L(\mathbf{x}(t), \mathbf{n}(t), \mathbf{v}) dt, \quad (2)$$

where we integrate along the ray $\mathbf{x}(t) = \mathbf{c} + t\mathbf{v}, t \in \mathbb{R}$. The weights $w(t)$ form a probability distribution along the ray [65], and are given by

$$w(t) = \sigma(\mathbf{x}(t)) \exp\left(-\int_0^t \sigma(\mathbf{x}(s)) ds\right). \quad (3)$$

Finally, $L : \Omega \times \mathbb{S}^2 \times \mathbb{S}^2 \rightarrow \mathbb{R}^3$ is the radiance field, which depends on location, normal and viewing direction.




Figure 2. Geometric reconstruction using our approach based on two different rendering strategies. Surface based rendering [66] gets easily stuck in bad local minima with less detail (hair on neck) and undesirable artifacts (cheek). Our approach based on volume rendering [65] does not suffer from these issues, resulting in highly detailed reconstructions.

Since positive values of the SDF, d inside the surface are assumed, the normal vector is obtained as $\mathbf{n} = \nabla d / \|\nabla d\|$. In practice, the integral (2) is being approximated using the well-known quadrature rule at a discrete set of samples $t_1 < t_2 < \dots < t_m$ for each pixel,

$$\mathcal{L}_p \approx \sum_{i=1}^{m-1} (t_{i+1} - t_i) w(t_i) L(\mathbf{x}(t_i), \mathbf{n}(t_i), \mathbf{v}). \quad (4)$$

Note, that the integral in (3) to compute $w(t_i)$ is accumulated in a similar way while iteratively computing the sum (4). VolSDF represents the scene using two separate MLPs, one is used to describe the SDF, d_{φ} and a global geometry feature map $z_{\varphi} : \Omega \rightarrow \mathbb{R}^{256}$, while a second MLP is used to describe the radiance L_{ψ} , both with their corresponding learnable network parameters φ, ψ . Additionally to its stable convergence compared to surface rendering approaches [36, 66], cp. Figure 2, VolSDF satisfies a theoretical guarantee to upper bound the opacity error compared to similar approaches [39, 58]. Specifically, after convergence the rendered image showing the geometry of the SDF using a volume renderer is almost indistinguishable from the same rendered image based on a surface renderer using e.g. a sphere tracing algorithm [17]. This observation was shown in [39] and has the consequence that the object can be rendered using standard surface renderer frameworks [25, 37, 69]. Additionally, the appearance is then defined on the object’s surface, despite being learned as a volumetric quantity. Nevertheless, VolSDF is unable to recover the reflectance properties, since both the material and lighting are baked into the radiance network. Hence, the re-

constructed 3D model can only be rendered with the same material under the same static illumination. In the next section, we show how we extend VolSDF to enable joint estimation of shape and material, which allows material editing and view synthesis under novel lighting conditions using a traditional graphics pipeline for surface rendering.

4. Method

We will first show how to decouple appearance into reflectance and lighting, which allows to estimate a high quality material in addition to the shape. Following this, we will extend this into a novel framework for super-resolved shape and BRDF estimation to allow rendering of novel views with more detail than the individual input images.

4.1. Radiance field with explicit BRDF model

We express the radiance field $L(\mathbf{x}, \mathbf{n}, \mathbf{v})$ in (2) in terms of the BRDF and lighting using the rendering equation [21],

$$L(\mathbf{x}, \mathbf{n}, \mathbf{v}) = \int_{H_{\mathbf{n}}} L_i(\mathbf{x}, \boldsymbol{\omega}) f_r(\mathbf{x}, \mathbf{n}, \mathbf{v}, \boldsymbol{\omega})(\boldsymbol{\omega} \cdot \mathbf{n}) d\boldsymbol{\omega}, \quad (5)$$

where $H_{\mathbf{n}} \subset \mathbb{S}^2$ is the half-sphere in direction \mathbf{n} , $L_i(\mathbf{x}, \boldsymbol{\omega})$ denotes the radiance incoming at \mathbf{x} from direction $\boldsymbol{\omega}$, and f_r the spatially varying BRDF (SVBRDF). Since we assume an achromatic point light source colocated with the camera center \mathbf{c} , (5) simplifies to

$$L(\mathbf{x}, \mathbf{n}, \mathbf{v}) = \frac{L_0}{\|\mathbf{x} - \mathbf{c}\|^2} f_r(\mathbf{x}, \mathbf{n}, \mathbf{v}, \mathbf{v})(\mathbf{v} \cdot \mathbf{n}), \quad (6)$$

where L_0 corresponds to the scalar light intensity. We use the simplified Disney BRDF [22], as this provides a compact model expressive enough to represent a wide variety of materials, which was successfully used in several prior works [29, 71]. Here, the SVBRDF is parametrized with a diffuse RGB albedo $\rho : \Omega \rightarrow \mathbb{R}_{\geq 0}^3$, a roughness $\alpha_r : \Omega \rightarrow \mathbb{R}_+$ and a specular albedo $\alpha_s : \Omega \rightarrow [0, 1]$.

We implement these three components using two MLPs. The first MLP $\rho(\mathbf{x}; \gamma_1)$ is used to compute the diffuse component of the BRDF at a point \mathbf{x} , the second MLP $\alpha(\mathbf{x}; \gamma_2) = (\alpha_s(\mathbf{x}; \gamma_2), \alpha_r(\mathbf{x}; \gamma_2))$ computes the respective specular components. The combined network parameters for BRDF are denoted with $\gamma = (\gamma_1, \gamma_2)$. As mentioned earlier, we model the geometry using a third MLP $d(\mathbf{x}; \theta)$ for the SDF with its own network parameters θ . Note that we do not incorporate a global geometric feature map into our framework. The main motivation behind such a map is for the radiance field L to account for indirect lighting and self-shadows. We follow the spirit of multiple works [4, 15, 29, 34, 70] showing that satisfactory results can be achieved without modeling indirect lighting explicitly. Furthermore, the use of the robust L^1 -norm in (10)

increases robustness to those effects. On the other hand, self-shadows are not present in our captures anyway, due to our colocated camera-light setup.

In order to perform inverse rendering, we train our neural networks from the available input images. For comparing the images to the rendered image in the loss function, we model the physical camera capturing process in the next section and show how this leads to a super-resolved reconstruction of the scene.

4.2. Super-resolution image formation model

As common in many volume rendering based approaches [32, 58, 65], it is assumed that the image brightness at a given pixel corresponds exactly to the accumulated radiance of the volume rendering (2),

$$I_p = \mathcal{L}_p(\theta, \gamma). \quad (7)$$

As described in [57], even if a framework can render novel views at any resolution during inference, the performance will significantly decrease when the inference resolution becomes larger than the one of the input images. Indeed, during training, the networks are sampled exactly at the pixel locations of the training images, which means that there is no training data for points of the surface whose projections do not coincide with these locations. Thanks to the interpolation property of neural networks, we can usually still compute a reasonable value for the points which were not seen during training, in particular since they are typically in between points which have been trained. However, the sampling rate of the input images band-limits the components of the reconstructed SDF and BRDF, and higher frequency details are not magically generated if one only renders at higher resolution.

In order to restore higher frequency content, we can exploit that the camera capturing process does not only sample the exact centers of the pixels. Instead, a camera performs an integration over a subset of incoming rays which can be modeled by the so-called point spread function (PSF). In essence, it describes blur during the capturing process, for example caused by integration across a sensor element, diffraction, lens aberrations, or objects being not perfectly in focus [10]. This intrinsic blur can actually be beneficial to avoid aliasing on the captured low resolution images, as it reduces high-frequency content in the captured scene, but leads to unavoidable loss of detail.

Taking this physical process into account, we follow [2, 13] and consider the effect of the PSF in our image formation model. Thus, generalizing (7), we convolve the accumulated radiance with the PSF kernel in order to obtain image irradiance,

$$I_p = (\mathcal{L}(\theta, \gamma) * \text{PSF})(p) = \int_{\mathbb{R}^2} \mathcal{L}_{p-q}(\theta, \gamma) \text{PSF}(q) dq. \quad (8)$$




Figure 3. Visualization of standard sampling strategies [32, 65, 70, 71], compared to our sampling strategy for a scaling factor of four (in each direction) in a 3×3 LR grid. Considering the low-resolution pixel in the center, a trivial Dirac kernel (bottom-left) leads to an evaluation of the network at a single (red) point (top-left). Our principled approach with a Gaussian kernel with standard deviation of half a low-resolution pixel’s side length (bottom-right) naturally leads to a much denser sampling (top-right) at multiple (blue) points within a squared region (blue square) allowing for a much denser sampled network optimization.

For the special case that the PSF is a Dirac delta distribution, one arrives at the original model (7), so this is indeed a generalization. In practice, we assume that the PSF is a Gaussian distribution, as [42] have shown the validity of such an approximation, and [13] successfully used it to achieve texture super-resolution. In our experiments, we choose half the size of a pixel in the low-resolution input images as the standard deviation. For computational efficiency, we approximate the convolution shown in (8) by Monte Carlo integration

$$I_p \approx \frac{1}{N_s} \sum_{k=1}^{N_s} \mathcal{L}_{p-q_k}(\theta, \gamma), \quad (9)$$

where $(q_k)_{k=1 \dots N_s}$ are samples drawn from the proposed PSF. See Figure 3 for a comparison of the sampling process based on a Dirac kernel and a Gaussian kernel.

4.3. Final training objective

Our final objective consists of three terms. The first term E_{RGB} is the data term, ensuring that the rendered images fit the input images,

$$E_{\text{RGB}}(\theta, \gamma) = \sum_p \|I_p - (\mathcal{L}(\theta, \gamma) * \text{PSF})(p)\|_1. \quad (10)$$

Note that we use an L^1 -norm to improve robustness against outliers. The second term is the Eikonal term E_{eik} , which encourages $d(\mathbf{x}; \theta)$ to approximate an SDF, this is similar as in [14],

$$E_{\text{eik}}(\theta) = \sum_{\mathbf{x}} (\|\nabla_{\mathbf{x}} d(\mathbf{x}; \theta)\| - 1)^2. \quad (11)$$

Finally, we largely follow [58] and introduce an *optional* mask loss $E_{\text{mask}}(\theta)$, allowing to impose silhouette consistency,

$$E_{\text{mask}}(\theta) = \sum_p \text{BCE}(M_p, W_p(\theta)), \quad (12)$$

where M_p is the given binary mask value at the pixel p , $W_p(\theta) = \sum_{i=1}^{m-1} w(t_i; \theta)$ is the sum of the weights at the sampling locations t_i used in (4), and BCE is the binary cross entropy loss [58]. We would like to emphasize that the only aim of the mask loss is to reduce computation time as it allows us to use more samples per pixel inside the mask, and one single sample per pixel outside. Additionally, if the mask loss is used, we also truncate the volume rendering integral (2) to the unit sphere, as the background is already modeled by the mask loss, further reducing the computational requirements. If the mask loss is not used, we use an inverse sphere parameterization for the background [72]. Both use cases are directly inspired from [58].

The final loss becomes the sum of the three terms with additional weighting parameters $\lambda_1, \lambda_2 \geq 0$,

$$E(\theta, \gamma) = E_{\text{RGB}}(\theta, \gamma) + \lambda_1 E_{\text{eik}}(\theta) + \lambda_2 E_{\text{mask}}(\theta). \quad (13)$$

This loss function can in principle be used to estimate a super-resolved geometry and BRDF with one single training pass. However, since the full super-resolution model is computationally expensive, we first run an initialization pass with a Dirac kernel for the PSF and without the mask loss. Indeed, in Section 5 we show that this super-resolution free approach already retrieves state-of-the-art results even without mask supervision. After that, we have accurate silhouettes from projecting the estimated geometry into the input images, and can significantly decrease computation time by optimizing sampling as described above. Furthermore, convergence of the full super-resolution model is accelerated since we already start from a reasonable initialization of the scene.

5. Results

We evaluate our framework SuperVol on both synthetic and real world data. To this end, we create a synthetic dataset which consists of a combination of two geometries and two materials, designated in the following as *dog1*, *dog2*, *girl1* and *girl2*. Each of them is rendered from 60 different viewpoints for training. 30 other viewpoints are rendered for testing, including non-colocated illumination to evaluate our approach’s generalization capability. Our real world dataset consist of four scans: *bird* and *squirrel*, which we captured ourselves, and *pony* and *dragon* from [4]. From these, we consider between 30 and 60 images for training and around 30 images for testing.

	↑PSNR		↑SSIM [59]		↓MAE		↑PSNR			↑SSIM [59]			↓MAE		
	[70]	noSR	[70]	noSR	[70]	noSR	[70]	noSR	SuperVol	[70]	noSR	SuperVol	[70]	noSR	SuperVol
synthetic	30.4949	35.2933	0.9095	0.9529	7.8286	4.4653	28.9161	31.3557	32.4413	0.8643	0.8983	0.9170	9.3410	5.8471	5.2795
synthetic non-colocated	31.7933	36.4801	0.9026	0.9505	7.8286	4.4653	30.4858	33.5028	34.2321	0.8695	0.9100	0.9249	9.3410	5.8471	5.2795
real world	29.1411	32.2888	0.8721	0.9276	×	×	28.9149	31.5179	31.7068	0.8817	0.9143	0.9157	×	×	×

(a) high resolution training

(b) low resolution training

Table 1. Average PSNR, SSIM [59] and MAE across the datasets using high and the more challenging low resolution input for training, respectively. In both training scenarios, SuperVol and its simplified counterpart noSR (with a Dirac kernel) outperform IRON [70] quantitatively. Geometric reconstruction and image synthesis quality is in favor of our approach for the tested novel and unseen viewpoints. Additionally, our method can generalize well to a non-colocated lighting setup. Note, SuperVol and noSR trained on low resolution images can perform better than IRON [70] trained on high resolution input.

Evaluation. We consider an ablation study that we name "noSR", which consists of our framework SuperVol, but with a Dirac kernel instead of a Gaussian kernel, and we evaluate both noSR and SuperVol together with IRON [70]. IRON has demonstrated state-of-the-art performance for inverse rendering with photometric images, significantly outperforming prior work [4, 29] for which no open-source code is available online. Further, we consider two complementary scenarios. In *high resolution training*, we use the original (high resolution) images of the scans for both training and testing. In *low resolution training*, we down-sample the resolution of the training images of the scans by a factor of four to mimic the effect of a low-resolution image capture, and then test on the high resolution test images.

The first scenario is typically considered in inverse rendering works, and allows to assess the quality of our differentiable volume renderer with an explicit BRDF model. Since both, training and testing resolutions are the same in this case, we only compare noSR against IRON. For the second scenario, we include SuperVol in the comparison, allowing to measure the impact of modeling the PSF.

High resolution training. As can be seen in Table 1 (a), noSR is quantitatively superior to IRON [70] in both cases, i.e. geometric quality and image synthesis. Qualitatively, geometric estimates are shown in Figure 4. It can be seen that noSR yields more detailed shape reconstruction in convex parts of the objects, e.g. the beak, eyes, and wings of *bird* or the hair, and tail of *pony*. Additionally, thanks to the volume rendering based approach, it is also more reliable in highly concave parts, e.g. the tail of *bird* or the wings of *pony*. IRON [70], as a surface rendering based approach, fails to accurately recover these geometric details and gets stuck in local optima. This further exhibits the advantage of using volume rendering instead of surface rendering for a better convergence [58, 65], see Figure 2. Image synthesis and normal maps are visualized in Figure 5, where noSR demonstrates more crisp results compared to the blurred reconstructions of IRON [70]. Our evaluation reveals a better inverse rendering performance when using our noSR approach even without properly modeling the PSF.




Figure 4. Estimated geometry after high resolution training. IRON [70] produces significant artifacts at concave parts. Additionally, it lacks of detail in convex regions. Instead, with its inverse volume rendering, noSR produces sharp and artifact free reconstructions in both regions of the objects.

Low resolution training. Table 1 (b) demonstrates quantitatively the effectiveness of correctly modeling the PSF, where we can see that our framework SuperVol outperforms both IRON [70] and noSR in terms of geometric quality, as well as novel view synthesis under both colocated and non-colocated lighting. Regarding novel view synthesis, Figure 6 shows that SuperVol can recover some crisp details that are barely visible in the individual low resolution images and not properly reconstructed by the competitors. When it comes to geometric accuracy, as shown in Figure 7, the refinement induced from modeling the PSF is not limited to an increased quality of image synthesis, but also has a positive impact on the geometry. SuperVol recovers geometric information, which is difficult to see in the low resolution visualization, and at most partially recovered or smoothed by the other approaches, IRON [70] and noSR.




Figure 5. Image synthesis and normal map results of novel viewpoints with a co-located light source after high resolution training. Our simplified approach noSR results in a significantly sharper reconstruction, accurately reproducing fine scale details that IRON [70] fails to recover. Notice the loss of detail in the hair of *girl2*. The *dragon*’s scales are partially lost. The *squirrel* seems to be overall blurry. noSR reconstructs sharper images across all datasets.

Generalization to non-located relighting. Our approach estimates BRDF parameters from a collocated camera-light setup, which only captures a slice of the BRDF. To assess the generalization of our approach to non-located lighting setups, we provide quantified results in Table 1 under “synthetic non-located”. Qualitative results are shown in the supplementary material, which further validates the accuracy of the estimated material.

6. Conclusion

We propose to enhance suitable neural representations of shape and material parameters with a physical model of the camera that includes the blurring and down-sampling due to the point spread function (PSF). This leads to a neural approach for recovering shape and material properties at a resolution that is superior to that of baseline methods and even superior to that of the input images. We carefully mo-

tivate the choice of representation including the use of volume rendering over surface rendering and propose a generalization of standard approaches via explicitly modeling the PSF. In qualitative and quantitative studies we demonstrate that the proposed approach outperforms the state-of-the-art and offers a method for high resolution 3D modeling of shape and appearance even from low-resolution imagery.

Limitations and future work. The computational demands are substantial, primarily driven by volume rendering and super-sampling. Future endeavors could consider integrating our framework with [33], yielding a significant speed-up, while preserving quality. Moreover, this work assumes opaque objects. An intriguing avenue for future work would be an extension to (semi-)transparent objects. **Acknowledgment** This work was supported by the ERC Advanced Grant SIMULACRON. We also thank Sai Bi for providing his data.




Figure 6. Image synthesis results of novel viewpoints with a colocated light source after low resolution training. The rendering obtained with our complete SuperRVol model is much sharper, with unprecedented details that are lost in the remaining approaches.




Figure 7. Estimated geometry after low resolution training. HR and LR denote high and low resolution visualizations of the ground truth geometry, respectively. Our framework allows to obtain significantly more detailed reconstructions than IRON [70] even when trained on low resolution images (noSR). The results are further improved in the super-resolution reconstruction (SuperRVol), which contains some details which are barely visible in the low resolution visualization (LR).

References

- [1] Saeed Anwar, Salman Khan, and Nick Barnes. A deep journey into super-resolution: A survey. *ACM Computing Surveys (CSUR)*, 53(3):1–34, 2020. 2
- [2] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(9):1167–1183, 2002. 4
- [3] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2
- [4] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. In *European Conference on Computer Vision*, pages 294–311. Springer, 2020. 2, 4, 5, 6
- [5] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerf: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. 2
- [6] M Brahimi, Y Quéau, B Haefner, and D Cremers. *On the Well-Posedness of Uncalibrated Photometric Stereo Under General Lighting*, chapter Advances in Photometric 3D-Reconstruction, pages 147–176. Springer International Publishing, Cham, 2020. 2
- [7] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvnr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4947–4956, 2021. 2
- [8] Kelvin CK Chan, Shangchen Zhou, Xiangyu Xu, and Chen Change Loy. Basicvnr++: Improving video super-resolution with enhanced propagation and alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5972–5981, 2022. 2
- [9] Ziang Cheng, Junxuan Li, and Hongdong Li. Wildlight: In-the-wild inverse rendering with a flashlight. *arXiv preprint arXiv:2303.14190*, 2023. 2
- [10] Mauricio Delbracio. *Two problems of digital image formation: recovering the camera point spread function and boosting stochastic renderers by auto-similarity filtering*. PhD thesis, École normale supérieure de Cachan-ENS Cachan, 2013. 4
- [11] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5712–5721, 2021. 2
- [12] Daniel Glasner, Shai Bagon, and Michal Irani. Super-resolution from a single image. In *2009 IEEE 12th international conference on computer vision*, pages 349–356. IEEE, 2009. 2
- [13] Bastian Goldlücke, Mathieu Aubry, Kalin Kolev, and Daniel Cremers. A super-resolution framework for high-accuracy multiview reconstruction. *International journal of computer vision*, 106(2):172–191, 2014. 2, 3, 4, 5
- [14] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. 5
- [15] Bjoern Haefner, Simon Green, Alan Oursland, Daniel Andersen, Michael Goesele, Daniel Cremers, Richard Newcombe, and Thomas Whelan. Recovering real-world reflectance properties and shading from hdr imagery. In *2021 International Conference on 3D Vision (3DV)*, pages 1075–1084. IEEE, 2021. 4
- [16] Bjoern Haefner, Songyou Peng, Alok Verma, Yvain Quéau, and Daniel Cremers. Photometric depth super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(10):2453–2464, 2019. 2
- [17] John C Hart. Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer*, 12(10):527–545, 1996. 3
- [18] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021. 2
- [19] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5197–5206, 2015. 2
- [20] Takashi Isobe, Xu Jia, Xin Tao, Changlin Li, Ruihuang Li, Yongjie Shi, Jing Mu, Huchuan Lu, and Yu-Wing Tai. Look back and forth: Video super-resolution with explicit temporal difference modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17411–17420, 2022. 2
- [21] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, 1986. 4
- [22] Brian Karis and Epic Games. Real shading in unreal engine 4. *Proc. Physically Based Shading Theory Practice*, 4(3):1, 2013. 4
- [23] Berk Kaya, Suryansh Kumar, Carlos Oliveira, Vittorio Ferrari, and Luc Van Gool. Uncalibrated neural inverse rendering for photometric stereo of general surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3804–3814, 2021. 2
- [24] Petr Kellnhofer, Lars C Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4287–4297, 2021. 2
- [25] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 37(6):222:1–222:11, 2018. 3
- [26] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021. 2

- [27] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. [2](#)
- [28] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Transactions on Graphics (TOG)*, 40(6):1–16, 2021. [2](#)
- [29] Fujun Luan, Shuang Zhao, Kavita Bala, and Zhao Dong. Unified shape and svbrdf recovery using differentiable monte carlo rendering. In *Computer Graphics Forum*, volume 40, pages 101–113. Wiley Online Library, 2021. [2](#), [4](#), [6](#)
- [30] Robert Maier, Jörg Stückler, and Daniel Cremers. Super-resolution keyframe fusion for 3d modeling with high-quality textures. In *2015 International Conference on 3D Vision*, pages 536–544. IEEE, 2015. [2](#), [3](#)
- [31] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. [2](#)
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#), [3](#), [4](#), [5](#)
- [33] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. [2](#), [7](#)
- [34] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. *ACM Transactions on Graphics (TOG)*, 37(6):1–12, 2018. [2](#), [4](#)
- [35] Kamal Nasrollahi and Thomas B Moeslund. Super-resolution: a comprehensive survey. *Machine vision and applications*, 25(6):1423–1468, 2014. [2](#)
- [36] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. [2](#), [3](#)
- [37] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics (TOG)*, 38(6):1–17, 2019. [3](#)
- [38] Ben Niu, Weilei Wen, Wenqi Ren, Xiangde Zhang, Lianping Yang, Shuzhen Wang, Kaihao Zhang, Xiaochun Cao, and Haifeng Shen. Single image super-resolution via a holistic attention network. In *European conference on computer vision*, pages 191–207. Springer, 2020. [2](#)
- [39] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5569–5579. IEEE, 2021. [2](#), [3](#)
- [40] Sung Cheol Park, Min Kyu Park, and Moon Gi Kang. Super-resolution image reconstruction: a technical overview. *IEEE signal processing magazine*, 20(3):21–36, 2003. [2](#)
- [41] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. [2](#)
- [42] David peter Capel. Image mosaicing and super-resolution. *3rd Chapter-Geometric Registration Robotics Research Group Department of Engineering Science University of Oxford*, 2001. [5](#)
- [43] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. [2](#)
- [44] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12216–12225, 2021. [2](#)
- [45] Gernot Riegler, Matthias Rüther, and Horst Bischof. Atgynet: Accurate depth super-resolution. In *European conference on computer vision*, pages 268–284. Springer, 2016. [2](#)
- [46] Lu Sang, Bjoern Haefner, and Daniel Cremers. Inferring super-resolution depth from a moving light-source enhanced rgb-d sensor: a variational approach. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1–10, 2020. [2](#)
- [47] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European conference on computer vision*, pages 501–518. Springer, 2016. [1](#)
- [48] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David W Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8598–8607, 2019. [2](#)
- [49] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7495–7504, 2021. [2](#)
- [50] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. [2](#)
- [51] Tatsunori Taniai and Takanori Maehara. Neural inverse rendering for general reflectance photometric stereo. In *International Conference on Machine Learning*, pages 4857–4866. PMLR, 2018. [2](#)

- [52] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. In *Computer Graphics Forum*, volume 39, pages 701–727. Wiley Online Library, 2020. [2](#)
- [53] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, W Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *Computer Graphics Forum*, volume 41, pages 703–735. Wiley Online Library, 2022. [2](#)
- [54] Jing Tian and Kai-Kuang Ma. A survey on super-resolution imaging. *Signal, Image and Video Processing*, 5(3):329–342, 2011. [2](#)
- [55] JD Van Ouwerkerk. Image super-resolution survey. *Image and vision Computing*, 24(10):1039–1052, 2006. [2](#)
- [56] Oleg Voynov, Alexey Artemov, Vage Egiazarian, Alexander Notchenko, Gleb Bobrovskikh, Evgeny Burnaev, and Denis Zorin. Perceptual deep depth super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5653–5663, 2019. [2](#)
- [57] Chen Wang, Xian Wu, Yuan-Chen Guo, Song-Hai Zhang, Yu-Wing Tai, and Shi-Min Hu. Nerf-sr: High quality neural radiance fields using supersampling. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 6445–6454, 2022. [2](#), [3](#), [4](#)
- [58] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. [2](#), [3](#), [4](#), [5](#), [6](#)
- [59] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. [6](#)
- [60] Zhihao Wang, Jian Chen, and Steven CH Hoi. Deep learning for image super-resolution: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3365–3387, 2020. [2](#)
- [61] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8534–8543, 2021. [2](#)
- [62] Qingxiong Yang, Ruigang Yang, James Davis, and David Nistér. Spatial-depth super resolution for range images. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [2](#)
- [63] Wenqi Yang, Guanying Chen, Chaofeng Chen, Zhenfang Chen, and Kwan-Yee K Wong. Ps-nerf: Neural inverse rendering for multi-view photometric stereo. In *European Conference on Computer Vision*, pages 266–284. Springer, 2022. [2](#)
- [64] Wenming Yang, Xuechen Zhang, Yapeng Tian, Wei Wang, Jing-Hao Xue, and Qingmin Liao. Deep learning for single image super-resolution: A brief review. *IEEE Transactions on Multimedia*, 21(12):3106–3121, 2019. [2](#)
- [65] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [66] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33:2492–2502, 2020. [2](#), [3](#)
- [67] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. [2](#)
- [68] Linwei Yue, Huanfeng Shen, Jie Li, Qiangqiang Yuan, Hongyan Zhang, and Liangpei Zhang. Image super-resolution: The techniques, applications, and future. *Signal Processing*, 128:389–408, 2016. [2](#)
- [69] Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, and Shuang Zhao. Path-space differentiable rendering. *ACM Trans. Graph.*, 39(4):143:1–143:19, 2020. [3](#)
- [70] Kai Zhang, Fujun Luan, Zhengqi Li, and Noah Snavely. Iron: Inverse rendering by optimizing neural sdfs and materials from photometric images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5565–5574, 2022. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [71] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. 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*, pages 5453–5462, 2021. [1](#), [2](#), [4](#), [5](#)
- [72] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [2](#), [5](#)
- [73] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 40(6):1–18, 2021. [2](#)
- [74] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18643–18652, 2022. [2](#)