

# Differentiable Shadow Mapping for Efficient Inverse Graphics

Markus Worchel

Marc Alexa

TU Berlin

## Abstract

We show how shadows can be efficiently generated in differentiable rendering of triangle meshes. Our central observation is that pre-filtered shadow mapping, a technique for approximating shadows based on rendering from the perspective of a light, can be combined with existing differentiable rasterizers to yield differentiable visibility information. We demonstrate at several inverse graphics problems that differentiable shadow maps are orders of magnitude faster than differentiable light transport simulation with similar accuracy – while differentiable rasterization without shadows often fails to converge.

## 1. Introduction

Differentiable renderers have become an essential tool for solving inverse problems in computer vision. They currently come in two flavors: (1) forward rasterization using *local* shading models [9, 10, 39] and (2) path tracing and/or Monte Carlo methods for *global* light transport simulation [22, 36, 53, 76]. While local methods are orders of magnitude faster, they lack effects of global light interaction such as shadows, caustics, or indirect illumination.

Modern methods in real-time graphics can generate surprisingly realistic images by using efficient approximations of global effects. The single most important aspect for increasing the realism of local shading is the consideration of *shadows* (see Figure 1): for each pixel to be shaded, check if the path to a light source is unobstructed before evaluating a local shading model. Doing this accurately is costly and many approximate techniques have been developed. Our central observation is that one of the oldest and most widely used, *shadow maps* [71] (see Section 3), can be adapted to work in differentiable rendering frameworks.

In Section 4 we explain how (certain approximations of) shadow mapping can be differentiated, exploiting existing differentiable rasterizers. Our main idea is similar to shadow maps: exploit efficient rasterization from the light’s point of view. For differentiable shadows this means: Existing differentiable rasterizers handle dis-




Figure 1. Adding shadows (middle) to local shading (left) is a significant step towards global light transport simulation (right).





Figure 2. Complex scene (330k triangles) rendered in real time with our differentiable shadow mapping. It shows self-shadowing of objects, shadowing between objects, and colored surfaces. Finite differences and our automatic derivatives w.r.t. movement of the light (note the non-zero derivatives at shadow boundaries).

continuities of primary visibility along primitive borders; we use this machinery to handle discontinuities of secondary visibility along shadow borders. The resulting images contain shadows and are differentiable (see Figure 2). For many inverse graphics problems the level of realism they provide will suffice, while being generated significantly faster than with global methods. This is important for machine learning tasks, where the renderings (and their derivatives w.r.t. scene parameters) are computed repeatedly. We provide details of the implementation and how parameters affect optimization based on target images in Section 5.

Given the importance of shadows for realistic image synthesis, it is unsurprising that many inverse problems heavily depend on them. We demonstrate the trade-off and possibilities of differentiable shadow mappings in several appli-



cations, ranging from pose estimation, over different types of geometry reconstruction, to interactive editing by manipulating shadows (Section 6).

The main idea of this work, using existing differentiable rendering frameworks that readily resolve the problem of visibility discontinuities for dealing with discontinuities of secondary rays, may be useful for various other scenarios beyond shadows. We elaborate on this and other immediate consequences of our approach in Section 7.

## 2. Related Work: Differentiable Rendering

Differentiable rendering enables computing derivatives of the rendering process, which generates a virtual image from a set of scene parameters. While there are differentiable renderers for a variety of scene representations, including volumes [51] or implicit surfaces [4, 24, 68], we focus on triangle meshes here and refer to Kato et al. [26] for a broader overview.

Many differentiable mesh renderers are based on rasterization, effectively point-sampling the scene on a regular grid in the image plane. Point sampling means that a primitive either covers a pixel or not; this binary outcome makes computing derivatives w.r.t. scene parameters difficult. Some differentiable rasterizers approximate the gradient computation [27, 39] whereas others approximate the rasterization [9–11, 30, 38, 58, 66], using a form of smoothing or anti-aliasing to ensure differentiability. Approaches also differ in terms of *local* shading of a pixel, from considering only given colors to evaluating reflection models based on the BRDF.

At the other end of the spectrum are complex physically-based differentiable renderers [22, 36, 53, 76], often based on Monte Carlo integration to numerically approximate the rendering equation [25]. For handling visibility discontinuities in the gradients, several techniques have been developed [5, 36, 40, 81]. Significant effort has been made to improve the immense time and memory requirements [23, 52, 67]. While modern GPUs possess specialized hardware for ray-triangle intersection, they still cannot match the performance of rasterization-based renderers.

Multi-view 3D reconstruction and related inverse problems such as view synthesis and material estimation are prime examples for the potential of differentiable rendering. Recent combinations of inverse graphics with machine learning techniques, like NeRF [6, 7, 44, 49, 79] or neural implicit surfaces [50, 54, 70, 74, 75, 78] are particularly successful. It is still common to use local shading models, especially in combination with efficient differentiable rasterizers for triangle meshes. Missing shadows or indirect illumination in the differentiable rendering process are often compensated for by assuming a setting where these effects are either negligible (e.g. co-located camera and light [41, 77] or soft environment lighting [17, 47]), by mod-

eling the interaction of light and material as a black box neural network [72], or simply ignored [18, 61]. In addition to analysis-by-synthesis settings, differentiable renderers have been used in a variety of contexts to train machine learning models [9, 15, 21, 32, 37, 43, 56, 62, 69, 73, 80].

Despite the popularity in computer graphics, surprisingly little effort has been spent on combining local shading models and approximate global solutions in the context of differentiable rendering. Lyu et al. [42] take a first step in this direction by approximating soft shadows produced by image-based lighting, however, their gradients are so far limited by the proxy geometry [60] used for shadow computation. In contrast, we use the full mesh geometry for the shadows and are not restricted to a specific class of deformations.

## 3. Background: Shadow Mapping


Shadow mapping [71] is based on the observation that points in a scene are lit by a (point) light source if they are closest to the light source. The *shadow map* is a depth image from the perspective of the light source. If a point in the scene is further from the light than the value stored in the shadow map, it is in the shadow. Otherwise it has the value stored in the shadow map, is closest to the light source, and lit. In local shading models, light is added over all light sources and the shadow computation is performed for each individual light, using a shadow map per light. In the following we consider a single light source and shadow map.

Formally, let the scene be given as the set  $\mathcal{X} \subset \mathbb{R}^3$ . A point light source is defined by its position  $\mathbf{p}$ . We assume a projective model and represent positions in homogeneous coordinates, so  $\mathbf{p}$  can be at infinity, creating a directional light source, in which case  $\mathbf{p}$  is the direction vector. Denote the distance of  $\mathbf{x} \in \mathbb{R}^3$  to the light source as  $d : \mathbb{R}^3 \mapsto [0 : d_{\max}]$ , where the distance to a directional light source is measured from a plane outside of  $\mathcal{X}$  so that distance values are finite and non-negative.

The shadow map is parameterized over the space of rays emanating from  $\mathbf{p}$ . For simplicity we assume that the rays from  $\mathbf{p}$  to any point in the scene can be parameterized over a single plane and we assume a rectangle  $\mathcal{U}$  on this plane to be the parameter space for the rays; if the light source is placed so that this condition is violated, several planes (i.e., arranged in the form of a box) may be used. The mapping from scene points  $\mathbf{x} \in \mathbb{R}^3$  to  $\mathcal{U}$  is the linear projective camera transformation  $L : \mathbb{R}^3 \mapsto \mathbb{R}^2$ . In this setup, the shadow map can be expressed as

$$f : \mathcal{U} \mapsto \mathbb{R}, \quad f(\mathbf{u}) = \min \{d(\mathbf{x}) : L(\mathbf{x}) = \mathbf{u}, \mathbf{x} \in \mathcal{X}\} \quad (1)$$

Given the shadow map, we get the *visibility* of any point in






Figure 3. Shadow mapping suffers from self-shadowing artifacts (“shadow acne”) that are alleviated by biasing the depths towards the light source in the visibility test. Variance shadow mapping handles these artifacts implicitly.

the scene w.r.t. the light source as

$$V(\mathbf{x}) = \begin{cases} 0 & d(\mathbf{x}) > f(L(\mathbf{x})) \\ 1 & \text{else.} \end{cases} \quad (2)$$

It is important to note that the ‘else’ case really reduces to  $d(\mathbf{x}) = f(L(\mathbf{x}))$  because  $f$  had been constructed to store the smallest possible value of  $f$  in direction  $L(\mathbf{x})$ .

The usefulness of shadow maps arises from the fact that they can be efficiently approximated using graphics hardware: simply render the scene from the perspective of the light source, disabling all lighting computations and storing only the depth values. However, a naive implementation will suffer from severe artifacts. First, the condition for a point to be lit is  $d(\mathbf{x}) \leq f(L(\mathbf{x}))$  in finite precision arithmetic, which may or may not happen for visible points on a smooth surface because of the different discretization of the rendered image and the shadow map. The resulting problem is referred to as “shadow acne” and commonly tackled by introducing a *bias*  $\beta$  for the depth comparison, considering points with  $d(\mathbf{x}) \leq f(L(\mathbf{x})) + \beta$  to be lit (Figure 3). While suitably chosen bias attenuates artifacts along smooth surfaces, it cannot remedy aliasing across shadow boundaries.

Various techniques have been suggested for dealing with aliasing artifacts resulting from the discrete nature of shadow maps [2, 3, 12, 57]. Many such approaches can be traced back to or interpreted as approximations of *percentage closer filtering* [59]. Rather than comparing  $d(\mathbf{x})$  to only  $f(L(\mathbf{x}))$  we consider a small neighborhood around  $L(\mathbf{x})$  in the shadow map and take the percentage of the rays around  $L(\mathbf{x})$  that would lit  $\mathbf{x}$ . In a slight generalization of the original formulation, we assume the neighborhood to be implicitly defined by a smooth (compactly supported) kernel  $k_{\mathbf{x}}(\mathbf{u})$  satisfying  $\int k_{\mathbf{x}}(\mathbf{u}) d\mathbf{u} = 1$ . Then the desired visibility function is

$$v(\mathbf{x}) = \int_{d(\mathbf{x}) \leq f(\mathbf{u})} k_{\mathbf{x}}(\mathbf{u}) d\mathbf{u}. \quad (3)$$

Note that in this case we can indeed have  $d(\mathbf{x}) < f(\mathbf{u})$  because  $\mathbf{u}$  is not restricted to  $L(\mathbf{x})$  but varies over a region.

The visibility function is taking on values in  $[0, 1]$ , and in particular the boundaries of shadows are smoothly transitioning between the two extreme values. This also means that on completely lit smooth surfaces, even if the visibility is not 1 in concave areas, it is always smoothly varying and close to 1, eliminating the necessity to estimate a suitable bias, which can be difficult for some scenes [13].

Percentage closer filtering successfully alleviates many sampling artifacts but is costly: for every point  $\mathbf{x}$  in the scene the integral needs to be evaluated. One would much rather *pre-filter* the shadow map [2]. Yet, simply convolving  $f$  with  $k$  and then evaluating the smoothed shadow map is different. This can be readily seen as the result would take only the discrete values 0 and 1. Better approximations (see [16] for an overview) generally pre-warp  $f$  and/or pre-filter not only  $f$  but also simple functions  $g(f)$ , such as  $g(f) = f^k$  or  $g = \exp$ . Then an approximation to Eq. (3) is built from the set of pre-computed functions.

While our approach works with all such approximations, we demonstrate it at the example of *variance shadow maps* [12]. The approximation is built by taking the view that the local shadow map may be considered a probability distribution. This makes sense because it is irrelevant what points are in the integration domain in Eq. (3) as long as the measure of the integration domain is correct. The local weighted mean and variance of the shadow map are

$$\mu_{\mathbf{x}} = \int k_{\mathbf{x}}(\mathbf{u}) f(\mathbf{u}) d\mathbf{u}, \quad \sigma_{\mathbf{x}}^2 = \int k_{\mathbf{x}}(\mathbf{u}) f(\mathbf{u})^2 d\mathbf{u} - \mu_{\mathbf{x}}^2. \quad (4)$$

We note that  $v(\mathbf{x})$  is a measure and since the shadow map  $f$  is positive the best bound for this measure is given by the *one-sided Chebyshev* or *Cantelli* inequality, yielding

$$v(\mathbf{x}) \leq \frac{\sigma_{\mathbf{x}}^2}{\sigma_{\mathbf{x}}^2 + (d(\mathbf{x}) - \mu_{\mathbf{x}})^2}. \quad (5)$$

Notice that, similar to percentage closer filtering, it may happen that  $d(\mathbf{x}) < \mu_{\mathbf{x}}$  because  $\mu_{\mathbf{x}}$  is a filtered version of  $f$ . Since in this case the approximation of the visibility function would decrease as  $d(\mathbf{x})$  comes closer to the light source, Donnelly and Lauritzen [12] suggest the following slightly modified visibility function:

$$\hat{v}(\mathbf{x}) = \begin{cases} \frac{\sigma_{\mathbf{x}}^2}{\sigma_{\mathbf{x}}^2 + (d(\mathbf{x}) - \mu_{\mathbf{x}})^2}, & d(\mathbf{x}) > \mu_{\mathbf{x}} \\ 1 & \text{else.} \end{cases} \quad (6)$$

The resulting shadow maps (see Figure 3) have been observed to be very similar to percentage closer filtering.

## 4. Differentiable Filtered Shadow Mapping

Our goal is to differentiate the visibility function with respect to the scene parameters  $\Theta$ . The binary visibility  $V(\mathbf{x})$  has zero derivative almost everywhere. The exact form  $v(\mathbf{x})$

of percentage closer filtering (Eq. (3)) is smooth, yet the integration domain not only varies with the scene parameters but also depends on  $\mathbf{x}$ . We suggest to rather consider  $\hat{v}$ .

The main difficulty in computing  $\frac{\partial \hat{v}}{\partial \Theta}$  is that the shadow map  $f$  is only piecewise smooth, and the boundaries of the smooth regions may change with  $\Theta$ . This makes the derivatives of the integrals in Eq. (4) challenging to compute. Similar in spirit to the original shadow map idea [71], we can use existing rasterizers. While the original motivation was efficiency, and this is also true in our case, an additional point is that significant effort has been spent to estimate such derivatives when rendering from the perspective of the viewer [9–11, 27, 30, 38, 39]. We can readily exploit any existing solution and simply render from the perspective of the light. Most differentiable rasterizers effectively replace the piecewise smooth functions  $f$  and  $f^2$  with smoothed functions  $\tilde{f}$  and  $\tilde{f}^2$ . Attenuating aliasing artifacts in pre-filtered shadow mapping *requires* smoothing, so the smoothing by differentiable rasterizers is complementary in our setting.

In the following, the vectors  $\tilde{\mathbf{f}}$  and  $\tilde{\mathbf{f}}^2$  represent the discrete versions of  $f$  and  $f^2$ , with  $\nabla \tilde{\mathbf{f}}$  and  $\nabla \tilde{\mathbf{f}}^2$  their Jacobians w.r.t.  $\Theta$ . The elements in  $\hat{v}$  dependent on the shadow map are  $\mu$  and  $\sigma^2$ . The weighted mean  $\mu$  is a convolution that, in the discrete case, can be expressed as the scalar product

$$\mathbf{m}_1 = \mathbf{G}\tilde{\mathbf{f}}, \quad (7)$$

where  $\mathbf{G}$  is a Toeplitz matrix and  $\mathbf{m}_1$  is the first discrete moment of the shadow map. The derivative is then simply

$$\frac{\partial \mathbf{m}_1}{\partial \Theta} = \frac{\partial \mathbf{m}_1}{\partial \tilde{\mathbf{f}}} \frac{\partial \tilde{\mathbf{f}}}{\partial \Theta} = \mathbf{G}\nabla \tilde{\mathbf{f}}. \quad (8)$$

For  $\sigma^2$  we introduce the second discrete moment

$$\mathbf{m}_2 = \mathbf{G}\tilde{\mathbf{f}}^2, \quad (9)$$

which similarly has the derivative

$$\frac{\partial \mathbf{m}_2}{\partial \Theta} = \frac{\partial \mathbf{m}_2}{\partial \tilde{\mathbf{f}}^2} \frac{\partial \tilde{\mathbf{f}}^2}{\partial \Theta} = \mathbf{G}\nabla \tilde{\mathbf{f}}^2. \quad (10)$$

The discrete image of  $\sigma^2$  can be computed from the moments as  $\mathbf{m}_2 - \mathbf{m}_1 \otimes \mathbf{m}_1$  and its derivative directly follows from expressions above and the product rule for  $\mathbf{m}_1 \otimes \mathbf{m}_1$ .

The distance  $d(\mathbf{x})$  is independent of the shadow map. Its variation with  $\Theta$  stems from possible modification of the position of the light source or geometry and can be easily expressed. Lastly,  $\hat{v}$  is a piecewise rational function of these variables, with the boundary between the two pieces at  $d = \mu$  depending on both  $d$  and  $\mu$ . This requires no special treatment, however, because  $\hat{v}$  is a continuously differentiable function across  $d = \mu$ , as we see from

$$\frac{\partial \hat{v}}{\partial \mu} = -\frac{\partial \hat{v}}{\partial d} = \begin{cases} 2 \frac{\sigma^2(d-\mu)}{(\sigma^2+(d-\mu)^2)^2} & d > \mu \\ 0 & \text{else} \end{cases}. \quad (11)$$

Thus,  $\frac{\partial \hat{v}}{\partial \Theta}$  can be directly computed using the chain rule.




Figure 4. We implement directional lights and spot lights in our framework. An extension to omnidirectional point lights is straight forward. Image-based lighting (IBL) is not supported, but the illumination from some environment maps can be sufficiently well approximated by a collection of point lights.

## 5. Implementation and Parameter Exploration

Our reference implementation is based on NVDIFFRAST [30] for creating and smoothing the shadow map. The Jacobians  $\nabla \tilde{\mathbf{f}}$  and  $\nabla \tilde{\mathbf{f}}^2$  (w.r.t. the scene parameters) are not explicitly computed. Rather, we use reverse-mode differentiation to compute the gradients of an objective function. The relevant derivatives are computed automatically [55] and are implicitly evaluated.

We currently use a single shadow map per light source, which supports directional lights and spot lights. Omnidirectional point lights can be implemented by creating and querying several shadow maps. Area lights or environment lighting would have to be approximated by point light sources – we leave a more accurate treatment for future work. Figure 4 illustrates the supported light sources in our current implementation.

By building on an existing differentiable rasterizer and automatic differentiation, our code becomes rather straightforward. The pseudo-code below is provided for a single light source – keep in mind that the methods have to be executed for each light source. The pre-computation of the pre-filtered depth images  $\mathbf{m}_1$  and  $\mathbf{m}_2$  is described in Alg. 1 below. Here, we assume that the `RENDERSCENE` function returns the smoothed discrete depth image  $\tilde{\mathbf{f}}$  as well as  $\tilde{\mathbf{f}}^2$ . The modular design of NVDIFFRAST implements this in two steps, first providing the discrete values  $\mathbf{f}$  and  $\mathbf{f}^2 = \mathbf{f} \otimes \mathbf{f}$ , which are then smoothed using the `antialias` function. We implement the filtering as a convolution with a kernel of size  $k$  (e.g. a Gaussian). Once the discrete smoothed shadow maps have been computed, the visibility for a point  $\mathbf{x} \in \mathcal{X}$  can be computed with Alg. 2.

---

**Algorithm 1** Render the shadow map for a light source  $l$

---

```

function RENDERSHADOWMAP( $\mathcal{X}, \mathbf{p}_l, k$ )
   $\tilde{\mathbf{f}}, \tilde{\mathbf{f}}^2 \leftarrow$  RENDERSCENE( $\mathcal{X}, \mathbf{p}_l$ )
   $\mathbf{m}_1 \leftarrow$  FILTER( $\tilde{\mathbf{f}}, k$ )
   $\mathbf{m}_2 \leftarrow$  FILTER( $\tilde{\mathbf{f}}^2, k$ )
  return  $\mathbf{m}_1, \mathbf{m}_2$ 

```

---

**Optimization without Smoothed Shadow Map.** We test our hypothesis that the discrete shadow map  $\mathbf{f}$  and  $\mathbf{f}^2$  must

**Algorithm 2** Compute the visibility between a point  $\mathbf{x}$  and a light at  $\mathbf{p}$  given the pre-filtered shadow map

```

function COMPUTEVISIBILITY( $\mathbf{x}, \mathbf{p}, \mathbf{m}_1, \mathbf{m}_2$ )
   $\mathbf{u} \leftarrow L(\mathbf{x})$ 
   $m_1, m_2 \leftarrow \text{SAMPLESHADOWMAP}(\mathbf{u}, \mathbf{m}_1, \mathbf{m}_2)$ 
   $d \leftarrow d(\mathbf{x}, \mathbf{p}), \mu \leftarrow m_1, \sigma^2 \leftarrow m_2 - m_1^2$ 
  if  $d \leq \mu$  then return 1
  else return  $\sigma^2 / (\sigma^2 + (d - \mu)^2)$ 

```




Figure 5. Convergence behavior for a simple scene with and without using a smoothed shadow map in the differentiable rasterizer. Without smoothing, shadow discontinuities are not properly differentiated and optimization fails.

be turned into smooth versions  $\tilde{\mathbf{f}}$  and  $\tilde{\mathbf{f}}^2$  for properly tracking shadow discontinuities. This is facilitated by the modularity of NVDIFFRAST that allows disabling the smoothing.

Our setup is as simple as possible: a single square occluder casts a single shadow on a single square receiver. We optimize the position of one displaced vertex of the occluder to reproduce the shadow in a reference configuration (see Figure 5). The optimization minimizes the mean squared error to the reference shadow image.

We find that the target configuration is not attained when using the discontinuous discrete shadow map for generating the visibility (i.e., without using NVDIFFRAST’s `antialias` function). Importantly, this result is independent of the kernel  $k$  used for pre-filtering to generate the moments  $\mathbf{m}_1$  and  $\mathbf{m}_2$ . In other words, filtering  $\mathbf{f}$  and  $\mathbf{f}^2$  cannot recover the connection between the scene parameters and the shadow boundaries. Using the smooth functions (i.e., applying `antialias` to  $\mathbf{f}$  and  $\mathbf{f}^2$ ) shows the expected convergence across variations of all other parameters.

**Shadow Map Resolution and Filter Kernel.** The shadow map resolution and the filter kernel size are key parameters as they directly affect the quality of the shadow approximation (see Figure 6). Instead of investigating the *forward* pass of shadow map rendering, we focus on the *backward* pass, because it influences the gradient computation and any optimization based on it.

Consider a scene similar to the one above, now with a highly tessellated and slightly rotated occluder. High tessellation (relative to shadow map pixels) results in sparse shadow map Jacobians for occluder geometry because the




Figure 6. Effect of shadow map resolution (top row) and filter size  $k$  (bottom row) on the shadows cast by differently slanted planes.




Figure 7. Jacobian of the visibility  $\nabla_x \mathbf{v}$ , where  $x$  is the horizontal translation of a highly tessellated occluder over a planar receiver.




Figure 8. Convergence plots for the single rectangular occluder experiment with different filter kernel sizes  $k$ . Stochastic gradient descent (left) and ADAM [28] (right) with small and large steps sizes  $\alpha$ . The  $y$ -axis shows the signed error of the  $x$  translation.

anti-aliasing misses primitive boundaries. This might occur for arbitrary differentiable renderers relying on explicit boundary detection. The filter applied to recover the moments (Eqs. (7), (9)) is also applied to the Jacobians of the shadow map and thus affects the Jacobian of the visibility function (see Figure 7). While the filter kernel size is merely a visual effect in the forward pass, it can indeed affect the optimization: Figure 8 shows that larger filter kernels increase the robustness of gradient descent, both for stochastic gradient descent and momentum-based variants. However, larger kernels also increase the runtime.

Low shadow map resolution may lead to occluders being missed completely by the visibility test (see Figure 6, top row). These issues can be alleviated by increasing the shadow map resolution; again at the cost of runtime.

In the following we consider a more complex setting to test different shadow map resolutions and filter kernel sizes in a more involved optimization: an object is placed on a rectangular floor plane and illuminated by  $n$  directional lights (Figure 9). Given a single reference image  $I_{\text{ref}}$ , we wish to recover the  $n$  light directions  $\{\mathbf{l}_i\}$ , minimizing the




Figure 9. Light direction estimation with four lights, varying filter kernel sizes and shadow map resolutions.




Figure 10. Quantitative results for the parameter exploration, using a light direction estimation experiment. The  $y$ -axis in the top row is the average alignment of predicted light directions to the reference light directions, measured by the dot product.

mean squared image error.

We run the optimization for different objects and varying number of lights  $n$ . We report qualitative results in Figure 9 and quantitative results in Figure 10. As expected, increasing the shadow map resolution or the filter kernel size increases the runtime. The optimization behaves robustly with high accuracy and consistent error bounds across resolutions and filter sizes. Again, the accuracy improves slightly with larger filter size, albeit this improvement seems marginal considering the error bounds. Lowering the shadow map resolution has a similar effect because the effective filter footprint increases if  $k$  stays constant. We observe that the optimization for one light can get stuck in local minima where the shadows do not align but the shade of the floor plane matches that in the reference image. For multiple lights, shadows dominate the image, so this scenario is less likely; this could explain the different error bounds. Conversely, recovering multiple lights is more difficult, hence the lower average accuracy. These challenges are specific to the task and, in general, we observe that our method works robustly with different sets of parameters.

## 6. Applications

We run all experiments on a workstation with an NVIDIA 1070 Ti GPU with 8 GB VRAM.

### 6.1. Monocular Pose Estimation

Estimating the orientation and position of an object from images is a classic computer vision task. The problem be-




Figure 11. Visual difference between path-traced images by Mitsuba 3 [22] in direct illumination mode and our approximation based on local shading models and pre-filtered shadow maps.




Figure 12. Pose estimation results with path-traced references. In most cases, shadows provide viable cues for convergence (top row) but occasionally pure local shading is sufficient (bottom row).

comes easier for sequences of images [34] or if priors apply, such as in pose estimation of human body parts [11,64]. Given only a single (monocular) image makes retrieving the rigid transformation of a known object ill-posed. Shafer and Kanade [65] noticed that the shadow of an object in the image provides viable cues for reducing the ambiguity. We consider a simplified version of the 6D problem, allowing only 2 translational and one rotational degree of freedom. We place an object in front of a receiver plane and a directional light orthogonal to the plane. In this setting, we recover the parameters  $\mathbf{t} = (x, y)$  (translation parallel to the plane) and  $\varphi$  (rotation around the object-centric “up” axis), given a reference image  $I_{\text{ref}}$ .

We perform tests across a variety of common 3D mesh models. The reference images are generated using the differentiable path tracer Mitsuba 3 [22] in direct illumination mode. We verify that our local shading model can generate images similar to the reference images (Figure 11). The results are generated by averaging over several random starting conditions for each object. This experiment is repeated for different camera image resolutions ( $128^2$ ,  $512^2$ ,  $1024^2$ ).

We report quantitative results for the  $512^2$  resolution in Table 1 (more data is in the supplementary material) and qualitative results in Figure 12. Mitsuba 3 recovers the rotation and translation most accurately, which *could* be attributed to its accurate gradient computation, using multiple sub-pixel samples. It should be expected, however, that the results are biased towards the method used for generating

Table 1. Quantitative results for the pose estimation experiments for different scenes and image resolution  $512 \times 512$  (full table in the supplementary material). We measure the rotation error  $\Delta\varphi$ , the translation error  $\Delta t$ , and the total runtime  $t$ , averaged over ten runs. We include results for a GPU with ray-tracing cores (“RT”) and for a setting where we use our renderer as reference (“Our Ref.”). Best scores in **bold**, second best scores underlined.

|                 | Mitsuba 3 [22]                      |                       |                    | Rast. + Shadows (Ours)              |                       |                    | Rasterizer                          |                       |                    |
|-----------------|-------------------------------------|-----------------------|--------------------|-------------------------------------|-----------------------|--------------------|-------------------------------------|-----------------------|--------------------|
|                 | $\downarrow \Delta\varphi [^\circ]$ | $\downarrow \Delta t$ | $\downarrow t [s]$ | $\downarrow \Delta\varphi [^\circ]$ | $\downarrow \Delta t$ | $\downarrow t [s]$ | $\downarrow \Delta\varphi [^\circ]$ | $\downarrow \Delta t$ | $\downarrow t [s]$ |
| Bunny           | <b>0.23</b>                         | <b>0.10</b>           | 325.29             | 0.33                                | <u>0.22</u>           | <u>3.76</u>        | 0.31                                | 0.26                  | <b>1.58</b>        |
| Dragon          | <b>2.86</b>                         | <b>3.91</b>           | 306.46             | <u>3.32</u>                         | <u>4.53</u>           | <u>3.64</u>        | 8.02                                | 11.63                 | <b>1.54</b>        |
| Hand            | <b>1.41</b>                         | <b>5.18</b>           | 317.40             | <u>1.68</u>                         | <u>5.36</u>           | <u>3.66</u>        | 1.85                                | 5.66                  | <b>1.51</b>        |
| Spot            | <b>0.03</b>                         | <b>0.02</b>           | 381.80             | <u>0.05</u>                         | <u>0.05</u>           | <u>3.76</u>        | 5.80                                | 23.15                 | <b>1.77</b>        |
| Spot (RT)       | <b>0.03</b>                         | <b>0.02</b>           | 53.26              | <u>0.05</u>                         | <u>0.05</u>           | <u>1.93</u>        | 6.92                                | 21.52                 | <b>1.05</b>        |
| Spot (Our Ref.) | <u>0.05</u>                         | <u>0.12</u>           | 384.24             | <b>0.02</b>                         | <b>0.06</b>           | <u>3.60</u>        | 5.83                                | 23.12                 | <b>1.71</b>        |

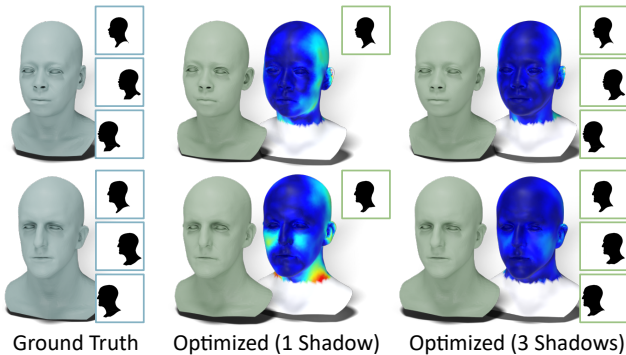


Figure 13. Face reconstruction from synthetic shadows.

the reference images. We show that this is the case by using our renderer as reference for one experiment (see Table 1) and finding our method to yield the most accurate results in this instance. Across all experiments, not using shadows at all and relying purely on the shading of the object surface performs worst, yet still manages to converge to reasonable solutions in some instances.

Unsurprisingly, rasterization is magnitudes faster than path tracing, both with and without shadows. Shadow mapping roughly doubles the runtime in rasterization, however, the pose estimation accuracy is significantly increased: our method consistently yields accuracy similar or close to the path tracing approach. Modern GPUs possess hardware units for ray-triangle intersections that significantly increase ray tracing performance. We include scores for a system with such a GPU (NVIDIA RTX 3080 Ti). It improves the path tracing runtime but, at least in this example, rasterization is still  $27\times$  faster.

## 6.2. Face Reconstruction from Profile Shadows

While the variation within certain classes of objects is governed by pose, such as in articulated objects, deformations are best described by morphable models. A prominent example for morphable models is face geometry [14].




Figure 14. Face reconstruction from real silhouettes of public figures. We masked out hair, facial hair, and the base.

Our idea is to reconstruct face geometry from one (or more) shadows of the face. We base the experiment on the publicly available morphable model assembled by Li et al. [35] (other models would work as well).

In the first experiment we generated synthetic shadows from the morphable model. This allows comparing the optimization result not just based on images but also geometry (see Figure 13). We find that using more shadows only slightly increases the runtime (the time is not linear in the number of light sources, because we have one base pass and multiple shadow passes) but significantly improves the accuracy (see supplementary material for quantitative results).

We have also performed reconstructions of face geometry from ‘real’ silhouettes, created manually by artist Charles Burns using scissors and black paper [8] (the process is shown here [1]). We ignore regions appearing to be (facial) hair or torso as those are not captured by the morphable model we use. While the lack of ground truth makes it difficult to quantify the resulting geometry, visually they appear to be reasonable approximations (Figure 14).

## 6.3. Shadow Art

Shadow Art [20, 46] is the design of a 3D object based only on the shadows it casts. When solving this design problem computationally, it is commonly simplified by assuming that the shadows are cast to planar surfaces orthogonal to directional lights, reducing the shadow images to silhouettes. In the traditional approach [46], the object is represented on a volumetric grid and voxel lines outside the shadow regions are removed – similar in spirit to visual hulls [33] or space carving [29]. Sadekar et al. [63] perform Shadow Art with differentiable rendering of meshes, but the limitation of co-located cameras and lights remains.

In our framework, lights and cameras can be placed arbitrarily. Also, curved shadow receivers work out-of-the-box (Figure 15). The only (minor) constraints we impose are gradient pre-conditioning [48] and a smoothness term between the normals of neighboring faces [41, 72] to avoid arbitrary vertex movements during optimization.

A single triangle mesh is only one possible representation of the generated geometry. The combination of automatic differentiation and gradient descent allows optimizing arbitrary scene parameters. For example, we can also repre-




Figure 15. Shadow Art with one view. Our method can be used for simple scenes with co-located light and camera (left), for more complex settings with perspective cameras that observe the shadow receiver from any direction (middle), and for complex receiver geometry (right).




Figure 16. Shadow art with two views. Our method can be used in any setting that permits differentiable rendering, for example optimizing vertex positions of a mesh (left) or optimizing the parameters of multiple geometric objects (middle and right).

sent the object as a collection of shapes (see Figure 16). Beyond our reference implementation based on triangle meshes, the same approach readily works with any geometry representation as long as it can be differentially rendered (e.g. signed distance functions [4,68] or spheres [31]).

#### 6.4. Interactive Modeling from Shadows

Our rasterization-based implementation generates shadows and – most importantly – the gradients of the objective function in real time. We demonstrate this by turning the shadow art setting, which is usually concerned with static reference shadows, into a “shadow modeling” setting, where a user can interactively paint the target shadows and an appropriate deformation is applied to the initial geometry in real time (see Figure 17). Note that using co-located camera and light is a design decision since painting silhouettes feels intuitive – the method is not limited to this setup.

### 7. Conclusion

Our approach to inverse graphics using differentiable rasterization with local shading models and shadow maps helps in closing the gap between efficient differentiable renderers based on rasterization and renderers based on path tracing and Monte Carlo methods that more accurately model light transport.




Figure 17. Shadow art from one view, with co-located camera and light is used as an interactive modeling tool. Users can modify the shadow of the reference object (gray) and an appropriate deformation is applied to the mesh in real-time (green).

**Limitations** Focusing on direct shadows neglects other effects of global illumination. But even within the framework of local shading models with shadow maps we have made simplifying assumptions: so far we consider only point light sources (instead of area lights) and for those we consider only one planar shadow map. Also, the particular choice of variance shadow maps suffers from artifacts, such as *light bleeding*, i.e. shadows becoming lighter in regions of high depth variance. Independent of the illumination effects, we have so far only considered triangle meshes for the representation of geometry.

**Possible Generalizations** Apart from shadows, various effects of global illumination have been recreated by efficient real-time approximations. Approaches very similar to ours may readily work with ambient occlusion [45], reflections and refractions [19], or generally indirect illumination.

Other geometry representations that are popular in the vision community, such as signed distance fields [24] or NeRFs [44], should work with shadow maps as well – although it is not clear if the same efficiency can be achieved, as the fast computation is mainly achieved by rasterizing triangle primitives.



## References

- [1] Alive Network. Charles burns silhouette artist. <https://www.youtube.com/watch?v=sZEZs3yVkpQ>, Oct 2016. **7**
- [2] Thomas Annen, Tom Mertens, Philippe Bekaert, Hans-Peter Seidel, and Jan Kautz. Convolution shadow maps. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, EGSR'07, page 51–60, Goslar, DEU, 2007. Eurographics Association. **3**
- [3] Thomas Annen, Tom Mertens, Hans-Peter Seidel, Eddy Flerackers, and Jan Kautz. Exponential shadow maps. In *Proceedings of Graphics Interface 2008*, GI '08, page 155–161, CAN, 2008. Canadian Information Processing Society. **3**
- [4] Sai Bangaru, Michael Gharbi, Tzu-Mao Li, Fujun Luan, Kalyan Sunkavalli, Milos Hasan, Sai Bi, Zexiang Xu, Gilbert Bernstein, and Fredo Durand. Differentiable rendering of neural sdfs through reparameterization. In *ACM SIGGRAPH Asia 2022 Conference Proceedings*, SIGGRAPH Asia '22, New York, NY, USA, 2022. Association for Computing Machinery. **2, 8**
- [5] Sai Praveen Bangaru, Tzu-Mao Li, and Frédo Durand. Unbiased warped-area sampling for differentiable rendering. *ACM Trans. Graph.*, 39(6), nov 2020. **2**
- [6] 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 (ICCV)*, pages 5855–5864, October 2021. **2**
- [7] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5470–5479, June 2022. **2**
- [8] Charles Burns. *Mastering Silhouettes: Expert Instruction in the Art of Silhouette Portraiture*. Stackpole Books, 2011. **7**
- [9] Wenzheng Chen, Jun Gao, Huan Ling, Edward J. Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2019. Curran Associates Inc. **1, 2, 4**
- [10] Wenzheng Chen, Joey Litalien, Jun Gao, Zian Wang, Clement Fuji Tsang, Sameh Khalis, Or Litany, and Sanja Fidler. DIB-R++: Learning to predict lighting and material with a hybrid differentiable renderer. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. **1, 2, 4**
- [11] Martin de La Gorce, Nikos Paragios, and David J. Fleet. Model-based hand tracking with texture, shading and self-occlusions. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. **2, 4, 6**
- [12] William Donnelly and Andrew Lauritzen. Variance shadow maps. In *Proceedings of the 2006 Symposium on Interactive 3D Graphics and Games*, I3D '06, page 161–165, New York, NY, USA, 2006. Association for Computing Machinery. **3**
- [13] Hang Dou, Yajie Yan, Ethan Kerzner, Zeng Dai, and Chris Wyman. Adaptive depth bias for shadow maps. In *Proceedings of the 18th Meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '14, page 97–102, New York, NY, USA, 2014. Association for Computing Machinery. **3**
- [14] Bernhard Egger, William A. P. Smith, Ayush Tewari, Stefanie Wuhrer, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, Christian Theobalt, Volker Blanz, and Thomas Vetter. 3d morphable face models—past, present, and future. *ACM Trans. Graph.*, 39(5), jun 2020. **7**
- [15] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Trans. Graph.*, 40(4), jul 2021. **2**
- [16] Jean-Marc Hasenfratz, Marc Lapierre, Nicolas Holzschuch, and François X. Sillion. A survey of Real-Time Soft Shadows Algorithms. *Computer Graphics Forum*, 22(4):753–774, Dec. 2003. **3**
- [17] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. *arXiv:2206.03380*, 2022. **2**
- [18] Jon Hasselgren, Jacob Munkberg, Jaakko Lehtinen, Miika Aittala, and Samuli Laine. Appearance-driven automatic 3d model simplification. In *Eurographics Symposium on Rendering*, 2021. **2**
- [19] Wolfgang Heidrich and Hans-Peter Seidel. Realistic, hardware-accelerated shading and lighting. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, page 171–178, USA, 1999. ACM Press/Addison-Wesley Publishing Co. **8**
- [20] Kai-Wen Hsiao, Jia-Bin Huang, and Hung-Kuo Chu. Multi-view wire art. *ACM Trans. Graph.*, 37(6), dec 2018. **7**
- [21] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **2**
- [22] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>. **1, 2, 6, 7**
- [23] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, and Delio Vicini. Dr.jit: A just-in-time compiler for differentiable rendering. *ACM Trans. Graph.*, 41(4), jul 2022. **2**
- [24] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **2, 8**
- [25] James T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, aug 1986. **2**
- [26] Hiroharu Kato, Deniz Beker, Mihai Morariu, Takahiro Ando, Toru Matsuoka, Wadim Kehl, and Adrien Gaidon. Differentiable rendering: A survey, 2020. **2**
- [27] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. **2, 4**

- [28] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 5
- [29] Kiriakos N. Kutulakos and Steven M. Seitz. A theory of shape by space carving. *International Journal of Computer Vision*, 38(3):199–218, 2000. 7
- [30] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Trans. Graph.*, 39(6), nov 2020. 2, 4
- [31] Christoph Lassner and Michael Zollhöfer. Pulsar: Efficient sphere-based neural rendering. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1440–1449, June 2021. 8
- [32] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme++: Facial shape and brdf inference with photorealistic rendering-aware gans. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9269–9284, 2022. 2
- [33] A. Laurentini. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):150–162, 1994. 7
- [34] Vincent Lepetit and Pascal Fua. Monocular model-based 3d tracking of rigid objects. *Found. Trends. Comput. Graph. Vis.*, 1(1):1–89, jan 2005. 6
- [35] Ruilong Li, Karl Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha Prasad, Bipin Kishore, Jun Xing, and Hao Li. Learning formation of physically-based face attributes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 7
- [36] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph.*, 37(6), dec 2018. 1, 2
- [37] 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 Trans. Graph.*, 40(6), dec 2021. 2
- [38] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2019. 2, 4
- [39] Matthew M. Loper and Michael J. Black. Opendr: An approximate differentiable renderer. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 154–169, Cham, 2014. Springer International Publishing. 1, 2, 4
- [40] Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Trans. Graph.*, 38(6), nov 2019. 2
- [41] Fujun Luan, Shuang Zhao, Kavita Bala, and Zhao Dong. Unified Shape and SVBRDF Recovery using Differentiable Monte Carlo Rendering. *Computer Graphics Forum*, 2021. 2, 7
- [42] Linjie Lyu, Marc Habermann, Lingjie Liu, Ayush Tewari, Christian Theobalt, et al. Efficient and differentiable shadow computation for inverse problems. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13107–13116, 2021. 2
- [43] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven neural stylization for meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13492–13502, June 2022. 2
- [44] 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. In *ECCV*, 2020. 2, 8
- [45] Gavin Miller. Efficient algorithms for local and global accessibility shading. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '94*, page 319–326, New York, NY, USA, 1994. Association for Computing Machinery. 8
- [46] Niloy J. Mitra and Mark Pauly. Shadow art. *ACM Trans. Graph.*, 28(5):1–7, dec 2009. 7
- [47] Jacob Munkberg, Wenzheng Chen, Jon Hasselgren, Alex Evans, Tianchang Shen, Thomas Müller, Jun Gao, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8270–8280, 2022. 2
- [48] Baptiste Nicolet, Alec Jacobson, and Wenzel Jakob. Large steps in inverse rendering of geometry. *ACM Trans. Graph.*, 40(6), dec 2021. 7
- [49] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [50] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [51] Merlin Nimier-David, Thomas Müller, Alexander Keller, and Wenzel Jakob. Unbiased inverse volume rendering with differential trackers. *ACM Trans. Graph.*, 41(4), jul 2022. 2
- [52] Merlin Nimier-David, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob. Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. *ACM Trans. Graph.*, 39(4), aug 2020. 2
- [53] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Trans. Graph.*, 38(6), nov 2019. 1, 2
- [54] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5589–5599, October 2021. 2
- [55] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming

- Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2019. Curran Associates Inc. 4
- [56] Dario Pavlo, Graham Spinks, Thomas Hofmann, Marie-Francine Moens, and Aurelien Lucchi. Convolutional generation of textured 3d meshes. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 870–882. Curran Associates, Inc., 2020. 2
- [57] Christoph Peters and Reinhard Klein. Moment shadow mapping. In *Proceedings of the 19th Symposium on Interactive 3D Graphics and Games*, i3D '15, page 7–14, New York, NY, USA, 2015. Association for Computing Machinery. 3
- [58] Felix Petersen, Bastian Goldluecke, Christian Borgelt, and Oliver Deussen. GenDR: A Generalized Differentiable Renderer. In *IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [59] William T. Reeves, David H. Salesin, and Robert L. Cook. Rendering antialiased shadows with depth maps. *SIGGRAPH Comput. Graph.*, 21(4):283–291, aug 1987. 3
- [60] Zhong Ren, Rui Wang, John Snyder, Kun Zhou, Xinguo Liu, Bo Sun, Peter-Pike Sloan, Hujun Bao, Qunsheng Peng, and Baining Guo. Real-time soft shadows in dynamic scenes using spherical harmonic exponentiation. In *ACM SIGGRAPH 2006 Papers*, SIGGRAPH '06, page 977–986, New York, NY, USA, 2006. Association for Computing Machinery. 2
- [61] Denys Rozumnyi, Martin R. Oswald, Vittorio Ferrari, and Marc Pollefeys. Shape from blur: Recovering textured 3d shape and motion of fast moving objects. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 29972–29983. Curran Associates, Inc., 2021. 2
- [62] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM Trans. Graph.*, 41(4), jul 2022. 2
- [63] Kaustubh Sadekar, Ashish Tiwari, and Shanmuganathan Raman. Shadow art revisited: A differentiable rendering based approach. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 29–37, January 2022. 7
- [64] Nikolaos Sarafianos, Bogdan Boteanu, Bogdan Ionescu, and Ioannis A. Kakadiaris. 3d human pose estimation: A review of the literature and analysis of covariates. *Computer Vision and Image Understanding*, 152:1–20, 2016. 6
- [65] Steven A Shafer and Takeo Kanade. Using shadows in finding surface orientations. *Computer Vision, Graphics, and Image Processing*, 22(1):145–176, 1983. 6
- [66] Yusuke Takimoto, Hiroyuki Sato, Hikari Takehara, Keishiro Uragaki, Takehiro Tawara, Xiao Liang, Kentaro Oku, Wataru Kishimoto, and Bo Zheng. Dressi: A hardware-agnostic differentiable renderer with reactive shader packing and soft rasterization. *Computer Graphics Forum*, 41(2):13–27, 2022. 2
- [67] Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Path replay backpropagation: Differentiating light paths using constant memory and linear time. *ACM Trans. Graph.*, 40(4), jul 2021. 2
- [68] Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Differentiable signed distance function rendering. *ACM Trans. Graph.*, 41(4), jul 2022. 2, 8
- [69] Gu Wang, Fabian Manhardt, Jianzhun Shao, Xiangyang Ji, Nassir Navab, and Federico Tombari. Self6d: Self-supervised monocular 6d object pose estimation. In *The European Conference on Computer Vision (ECCV)*, August 2020. 2
- [70] 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. *NeurIPS*, 2021. 2
- [71] Lance Williams. Casting curved shadows on curved surfaces. In *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '78, page 270–274, New York, NY, USA, 1978. Association for Computing Machinery. 1, 2, 4
- [72] Markus Worchel, Rodrigo Diaz, Weiwen Hu, Oliver Schreer, Ingo Feldmann, and Peter Eisert. Multi-view mesh reconstruction with neural deferred shading. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6187–6197, June 2022. 2, 7
- [73] Shangzhe Wu, Christian Ruppel, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [74] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. 2
- [75] 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, 2020. 2
- [76] Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, and Shuang Zhao. Path-space differentiable rendering. *ACM Trans. Graph.*, 39(4), aug 2020. 1, 2
- [77] 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 (CVPR)*, pages 5565–5574, June 2022. 2
- [78] 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 *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [79] 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 Trans. Graph.*, 40(6), dec 2021. [2](#)

- [80] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yinan Zhang, Antonio Torralba, and Sanja Fidler. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. In *International Conference on Learning Representations*, 2021. [2](#)
- [81] Yang Zhou, Lifan Wu, Ravi Ramamoorthi, and Ling-Qi Yan. Vectorization for fast, analytic, and differentiable visibility. *ACM Trans. Graph.*, 40(3), jul 2021. [2](#)