

DeepSurfels: Learning Online Appearance Fusion

Marko Mihajlovic¹ Silvan Weder¹ Marc Pollefeys^{1,2} Martin R. Oswald¹
¹Department of Computer Science, ETH Zurich ²Microsoft Mixed Reality and AI Zurich Lab

OnlineReconstruction.github.io/DeepSurfels

Abstract

We present DeepSurfels, a novel hybrid scene representation for geometry and appearance information. DeepSurfels combines explicit and neural building blocks to jointly encode geometry and appearance information. In contrast to established representations, DeepSurfels better represents high-frequency textures, is well-suited for online updates of appearance information, and can be easily combined with machine learning methods. We further present an end-to-end trainable online appearance fusion pipeline that fuses information from RGB images into the proposed scene representation and is trained using self-supervision imposed by the reprojection error with respect to the input images. Our method compares favorably to classical texture mapping approaches as well as recent learning-based techniques. Moreover, we demonstrate lower runtime, improved generalization capabilities, and better scalability to larger scenes compared to existing methods.

1. Introduction

Realistic 3D model reconstruction from images and depth sensors has been a central and long-studied problem in computer vision. Appearance mapping is often treated as a separate post-processing step that follows 3D surface reconstruction and is usually approached using batch-based optimization methods [18, 19, 23, 87] that are unsuitable for many applications that do not have access to the entire dataset at processing time, for instance, robot navigation [8, 9, 26], augmented reality [51, 68], and virtual reality [14, 43, 44] applications, Simultaneous Localization and Mapping (SLAM) systems [93], online scene perception methods [28, 69], and many others.

Common online fusion methods like KinectFusion [50] are well suited for online geometry fusion and can efficiently handle noise and topological changes. However, due to their high memory requirements at high voxel resolutions, they have strong limitations when it comes to encoding high-frequency appearance details on the surface. On the other hand, meshes with high-resolution tex-

ture maps [19, 23, 87] are well-suited for encoding high-frequency appearance information in an efficient manner, but they have difficulties in handling topology changes in an online reconstruction setting. Moreover, recent learning-based approaches [49, 54, 72, 73] have achieved high-quality results by learning geometry and texture mapping directly from RGB images. However, they are not well suited for local online updates, do not scale to large-scale scenes, and easily overfit to the training data.

In this paper, we approach the problem of online appearance reconstruction from RGB-D images by combining the advantages of **1**) implicit grids, which easily handle topological changes and where low resolution is often sufficient to encode the scene topology, **2**) scalable high-frequency appearance along the surface via texture maps or learned feature maps, and **3**) a learned scene representation to build a framework for learning-based appearance fusion that allows for online processing and scalability to large scenes. To this end, we propose a novel scene representation **DeepSurfels** and an efficient learning-based **online appearance fusion pipeline** which is illustrated in Figure 1.

Our DeepSurfels representation is a hybrid between an implicit surface that encodes the topology and low-frequency geometric details and a surfel representation that encodes high-frequency geometry and appearance information in form of surface-aligned patches. These patches are arranged in a sparse grid and consist of surface-aligned texels that encode appearance information *either* in the classical form of RGB color values *or*, as proposed, via learned feature vectors. The sparse grid allows for efficient volumetric rendering and enables explicit scene updates that are crucial for online fusion, while the 2D patches enable quadratic memory storage complexity like meshes or sparse grid structures. Depending on the DeepSurfel parameters it can approximate between simple colored voxels (high grid resolution, 1×1 patches) and textured meshes with high texture atlas resolutions (lower grid resolution, higher patch resolution). Our online appearance fusion pipeline iteratively fuses RGB-D frames into estimated DeepSurfels geometry and is optimized by using a differentiable renderer for self-supervision and the reprojection error as train-

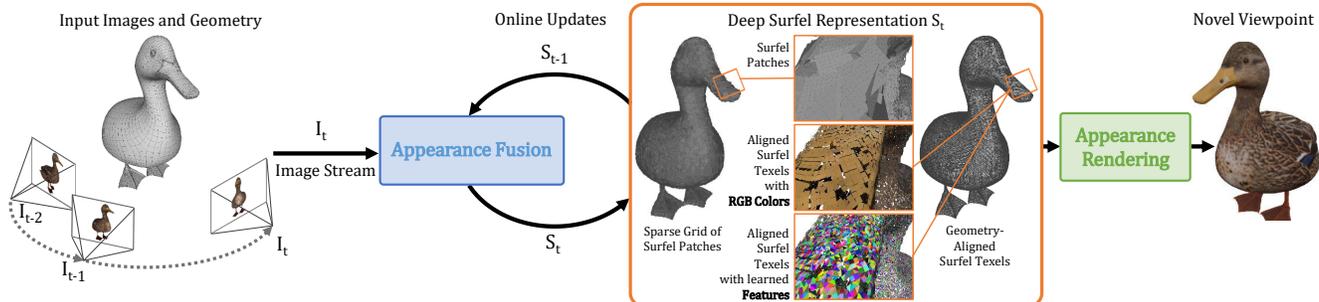


Figure 1. **Overview of our online appearance fusion pipeline and the DeepSurfel scene representation.** The **Appearance Fusion** network efficiently aggregates appearance information from a stream of camera views into the proposed DeepSurfel representation S_{t-1} that maintains high-frequency geometric and appearance information. DeepSurfels is a sparse grid of 2D patches that consist of surface-aligned texels, which encode appearance information either as RGB color values or learned feature vectors. The proposed **Appearance Rendering** network interprets aggregated and interpolated geometric and appearance information stored in DeepSurfels for rendering novel viewpoints. In this example we used DeepSurfels with a sparse 64^3 patch grid with 8×8 resolution surfel patches.

ing signal. In this way, the pipeline does not require any ground-truth texture maps and the training procedure allows for efficient transfer to new sensors and scenes without the need for acquiring costly ground-truth data.

While we eventually target full online reconstruction of both geometry and texture from monocular video, we only focus on online appearance estimation in this paper. Even in a setting with known geometry, our online approach has scalability advantages: We can fuse arbitrary numbers of input frames and the grid-aligned surfels have performance advantages during feature aggregation across local neighbors and for controlling the sampling density. Our grid-aligned surfel patches can also be seen as a spatial alignment of per-voxel sub-features being anchored along the surface. In contrast to works that only save a single feature vector per voxel, *e.g.* DeepVoxels [72], we can directly relate sub-features with particular image pixels via projective mapping and as such simplify the network learning task and improve output accuracy. As opposed to many novel view synthesis works [49, 72, 73], we do not overfit onto a single scene, but train a network that generalizes over multiple scenes without re-training. While those methods iterate many times over each input image in a slow optimization process, our method processes every image only once with a single network forward pass and is thus much faster. From the application point of view, our approach is thus closer to classical texture mapping methods like [18, 19, 23, 87].

We compare our novel scene representation and appearance fusion pipeline to existing methods on single and multi-object datasets and show that our scene representation better captures high-frequency textures. Moreover, our method generalizes well and compares favorably even to existing texture optimization methods that jointly optimize all images together. This is a crucial step towards a fully end-to-end appearance fusion method that can be deployed to real-world applications. In sum, our key **contributions** are:

- **DeepSurfels.** A novel scalable and memory-efficient 3D scene representation closing the gap between traditional interpretable and modern learned representations.
- **Online Appearance Fusion Pipeline.** An end-to-end differentiable and efficient online appearance fusion pipeline compatible with classical and learned texture mapping. The method yields competitive texturing results without heavy optimization as every input frame is processed only *once* with a single network forward pass.
- **Generalized Novel View Synthesis.** Contrary to other learning-based methods [49, 72, 73] that overfit onto a single scene, our method generalizes to new scenes without retraining.

2. Related Work

Our method relates to, and builds upon previous work on scene representations and appearance estimation which are reviewed in the following subsections.

2.1. Scene Representations

Scene representations can be broadly divided into explicit geometric and learned representations.

Explicit Geometric Representations. The major advantage of explicit geometric representations is their direct interpretability. *Point clouds* [1, 20, 56] are a lightweight and flexible 3D representation being the raw output of many 3D scanners, RGB-D cameras, and LiDARs. However, they are less suitable for the extraction of watertight surfaces due to lacking topology and connectivity information. This also impedes realistic rendering with detailed textures and complex lighting. *Mesh* representations [29, 33, 41, 89] scale well and texture mapping is convenient. However, topological changes are difficult to handle in an online process. *Voxel-grids* [10, 24, 40, 62, 77, 97] – as a natural extension of pixels to 3D space – easily handle topological changes but are difficult to use with textures and com-

plex light models. Another problem arises from the cubic memory complexity of the dense representation, which makes it expensive to capture precise shape details of complex objects. *Surface elements (Surfels)* [59, 70, 88, 94] are non-connected point primitives that reduce geometry to the essentials needed for rendering, thus being more memory efficient than meshes while still providing good rendering properties. Our *DeepSurfel* representation provides several advantages over existing explicit representations: **1)** it maintains better connectivity information than point clouds and surfels, **2)** scales better than voxel grids while still compatible with octree [75, 101, 102] and voxel hashing [32, 53] approaches that improve memory efficiency, **3)** provides better rendering quality than point clouds or voxel grids, **4)** enables fast rendering, and **5)** allows local updates.

Learned Shape Representations. Recent learned implicit representations have achieved impressive results in modeling geometry [13, 46, 47, 56]. They learn a neural network to predict the signed distance or occupancy for a given query point. These approaches struggle to scale to larger scenes and to capture high-frequency details as they tend to learn low-frequency functions which often results in over-smoothed geometry [60]. This problem has been approached by the idea of using local features [11, 31, 57, 98] or directly regressing local geometries [91, 92] for improved scalability and representation power of implicit representations. Our proposed *DeepSurfels* follows this direction using only local learned features for scalability. Unlike most learning-based methods, we do not encode the scene by network optimization. Instead, we train a network to estimate latent features following a data-driven approach, which better generalizes and facilitates online updates.

2.2. Appearance Estimation

Classical Texture Mapping. The classical way of coloring a surface from a set of input images with known camera pose is to un-project the image information onto the surface and perform a selection or blending operation to fuse the color information [3, 18, 95]. Due to errors in the camera alignment or in the surface geometry, blurry textures or patch seams affect results and additional texture alignment procedures have been proposed [5, 19, 22, 25, 37, 38, 79, 82, 87] to tackle these problems. Better texture mapping results have been achieved with an optical flow-like correction in texture space [19, 23, 87], patch-based optimization [6], or via 2D perspective warp techniques [36]. With significantly more computation effort, it is also possible to better leverage the redundancy of multiple surface observations from different views and to compute super-resolved texture maps via energy minimization [23, 27, 85, 86] or with deep learning techniques [39, 63]. All previously mentioned methods share the strategy of aggregating appearance information in patches or texture atlases with corresponding co-

ordinates onto a mesh-based surface, while other works use voxel grids [35, 45, 50, 71, 78, 104], or mesh colors [4, 99]. An overview of texture mapping methods with different representations is given in [80, 100].

Learned Appearance Representations. Recent learned appearance representations have achieved state-of-the-art results and outperformed most classical texturing methods. They encode visual information into learned features and store them in voxel grids [21, 44, 48, 58, 61, 74], point clouds [2], or meshes [64, 83, 103] which are rendered using neural networks. [54, 55] use a neural network conditioned on geometry to generate a learned texture representation. [52] combines geometry and appearance to generate a joint implicit representations. Worrall *et al.* [96] learn a disentangled representation to interpret and manipulate learned feature-based scene representations. SRNs [73] encode the scene into a neural network and render novel views using a neural ray-tracer. NeRF [49] inputs the viewing direction together with point coordinates which allows to also model illumination and complex non-Lambertian surfaces. Other works [65, 66] take advantage of local features for higher representation power, while [30, 84] uses appropriate loss terms to correct for geometric misalignment. Recent trends and applications of neural renderers are summarized in [81]. However, these methods are currently limited to fixed-size scenes, do not scale well to larger real-world scenes, or are unsuitable for online processing of appearance information. The global volumetric appearance reconstruction approach [7] additionally separates albedo, roughness, and lighting. Liu *et al.* [42] present a learned approach for shape and texture reconstruction that linearly fuses shape and color information in a voxel grid as in [17] and post-process the grid with a multi-resolution neural network. However, pure post-processing methods may not be able to revert errors of an incorrect earlier linear fusion.

Global vs Local Appearance Representations. Existing learned scene representations can be separated into global and local approaches. [49, 54] are global approaches encoding the scene into a single feature vector or the weights of a neural network, while [72] can be considered a local approach that uses a dense grid of feature vectors. For better scalability and higher representation power, we follow the local direction. Further, the local storage of appearance keeps the updates of the encoded information local. Moreover, it allows to exploit geometric relations better constraining the learning problem for improved generalization.

Online Appearance Aggregation. Most texture mapping methods process all input images in a batch-based way after the geometry estimation step and are implemented as a separate post-processing step, whereas only a minority addresses the problem of online appearance reconstruction. A popular work is KinectFusion [50] and related works [36, 45, 104], which estimate surface and appear-

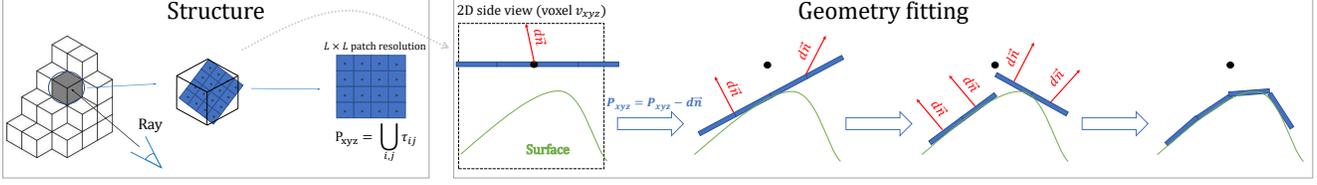


Figure 2. **DeepSurfel surface fitting.** In a recursive fitting procedure, we align the texels of each patch with the underlying SDF surface by shifting its location and adjusting its normal. d denotes queried signed distance, \vec{n} denotes SDF gradient ∇SDF in x, y, z directions.

ance information from a stream of RGB-D images. Other works fuse both geometry and appearance information directly into an oriented surfel cloud [70, 88, 94]. The major drawback of these methods is limited capacity to store high-frequency appearance along the surface and low-quality renderings. Therefore, we propose an efficient online appearance estimation pipeline mitigating these limitations.

3. DeepSurfels 3D Scene Representation

We propose *DeepSurfels* as a powerful, scalable, and easy-to-use alternative to mitigate previously mentioned problems of many scene representations.

Data Structure. DeepSurfels is a set of patches with $L \times L$ texels that can either store color information or learned feature vectors. The elementary building block is an oriented texel $\tau \in \mathbb{R}^c$ that is associated with its weight parameter ω and is stored on the objects' surface, where c denotes the number of feature channels. This number can be chosen arbitrarily for learned appearance fusion as suited for the problem setting, while we set $c = 3$ for deterministic RGB texturing. The texels τ are arranged in an $L \times L$ resolution patch $P_{xyz} : \{i, j \rightarrow \tau_{ij}; i, j \in [1, L]\}$ that is located in a sparse patch grid $\mathcal{P} = \{P_{xyz}\}_{x \leq X, y \leq Y, z \leq Z}$, where X, Y, Z represent DeepSurfels' grid resolution. Although the spatial patch size can be chosen arbitrarily, we empirically observed that texturing works best when the patch size is equal to the grid cell size such that there is no overlap between neighboring patches. For efficiency reasons, it is sufficient to store patches only for grid cells that intersect the objects' surface. However, it is also possible to allocate more layers around the iso-surface to account for noisy geometry as it is common for geometric fusion approaches [50].

Surface Fitting. We propose a recursive algorithm to align each texel τ_{ij} of the patch with the implicit surface of the geometry. We compute the patch position and orientation from a signed distance function (SDF) representing the Euclidean distance to the closest surface.

Initially, every patch P_{xyz} in the grid \mathcal{P} is positioned at the center of its grid cell. Then, the patches are shifted to the closest surface by using the pre-computed SDF, oriented according to the SDF gradient ∇SDF in all x, y, z directions, and rotated to maximize the surface coverage. These patches are subdivided into κ^2 non-overlapping patches of

$\frac{L}{\kappa} \times \frac{L}{\kappa}$ resolution, where $\kappa \geq 2$ is the smallest integer to non-trivially divide L . Each sub-patch is aligned again using the SDF field, where we trilinearly interpolate the SDF value at non-integer grid positions. This patch subdivision and alignment is repeated recursively until the resolution reaches 1×1 when patches represent texels that lie on the isosurface. This process is visually illustrated in Figure 2.

4. Online Appearance Fusion Pipeline

We also propose a pipeline for learning appearance fusion (depicted in Figure 3) that incrementally fuses RGB measurements into DeepSurfels at every time step t and yields DeepSurfel state S_t . The input to our pipeline are intrinsic K_t and extrinsic R_t camera parameters, an RGB image $I_t \in \mathbb{R}^{H \times W \times C}$, and corresponding depth map $D_t \in \mathbb{R}^{H \times W}$, where H, W and C denote image height, width, and the number of channels respectively. The pipeline consists of four main components detailed in the following.

Differentiable Projection II. The projection module renders a super-resolved feature map $\hat{F}_{t-1} \in \mathbb{R}^{kH \times kW \times c}$, where k is an upsampling factor inspired by [16] to ensure dense coverage of the geometry. There are three steps to render this feature map from already stored scene content.

First, each pixel in the incoming frame D_t is subdivided into k^2 distinct sub-pixels p_{ij}^t ($i \in [1, kH], j \in [1, kW]$), thus forming an upsampled image grid.

Second, by leveraging camera and depth information, the center of the sub-pixel p_{ij}^t is un-projected into the scene. From the un-projected scene point, the closest DeepSurfel texel and all texels within the surrounding l_∞ ball are selected. The size of this ball is chosen proportional to the size of the un-projected sub-pixel in the world space.

Third, an efficient uniform average of the selected texels determines the value of the feature entry $\hat{f}_{ij}^{t-1} \in \hat{F}_{t-1}$ (1):

$$\hat{f}_{ij}^{t-1} = \frac{1}{|T_{ij}^t|} \sum_{\tau \in T_{ij}^t} \tau, \quad (1)$$

where T_{ij}^t is the set of selected texels. This algorithm is simple, leverages the grid representation for fast rendering, and can flexibly render further optionally stored features or a surface normal map \hat{N}_{t-1} that we jointly denote as meta features \hat{M}_{t-1} . Note that all operations are differentiable

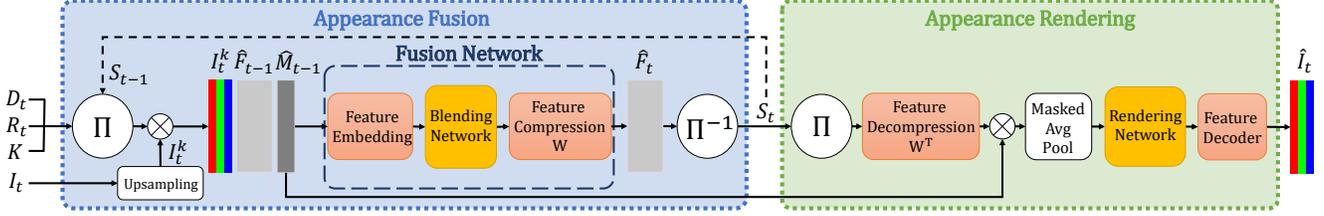


Figure 3. **Overview of our learned appearance fusion pipeline.** The pipeline consists of an **Appearance Fusion** module that integrates a new RGB measurement I_t into DeepSurfels S_{t-1} and a differentiable **Appearance Rendering** module that interprets and renders the content of representation for a given viewpoint. White blocks denote differentiable deterministic operations, rectangular blocks denote data, rounded rectangular blocks are trainable modules, and \otimes is a feature stacking operation.

and the selection can be implemented as a differentiable multiplication by an indicator function.

Fusion Network. The input image I_t is deterministically upsampled $I_t^k \in \mathbb{R}^{kH \times kW \times C}$ by factor k (nearest-neighbor interpolation) and stacked \otimes with the super-resolved features $\hat{F}_{t-1} \otimes \hat{M}_{t-1} \otimes I_t^k$. This stacked representation is embedded into a higher-dimensional feature space by a trainable linear transformation (*Feature Embedding* module Figure 3). Then, the embeddings are refined by *Blending Network* that consists of five convolutional layers (3×3 kernel size) interleaved with dropout and leaky ReLU activations. This network, based on a small receptive field, produces refined features aware of neighboring information that alleviates the problem of discretization artifacts, which can occur for low DeepSurfels resolutions. Lastly, these features are compressed by *Feature Compression* W layer to a lower dimensional feature space that is defined by DeepSurfels’ number of channels. The final output is an updated feature map \hat{F}_t that blends old information from \hat{F}_{t-1} with the new appearance information from I_t .

Inverse Projection Π^{-1} . While the fusion module and the explicit geometry representation preserve spatial coherence, this module is responsible for integrating the new appearance information in a temporally coherent way. Without temporal coherence, a new observation could overwrite old states minimizing the reprojection error for the current frame while erasing valuable prior information. The inverse projection module Π^{-1} integrates the updated feature map \hat{F}_t into the representation S_{t-1} to produce the new state S_t . For efficiency reasons, only texel values $\forall \tau_{t-1} \in \bigcup_{i,j}^{kH, kW} T_{i,j}^{t-1}$ and their weights ω_{t-1} that were intersected by at least one of the sub-pixels are updated using the following moving average scheme:

$$\tau_t = \frac{1}{\omega_{t-1} + 1} \left(\tau_{t-1} \omega_{t-1} + \frac{\sum_{i,j}^{kH, kW} \hat{f}_{ij}^t \mathbb{I}_{\tau_{t-1} \in T_{ij}^{t-1}}}{\sum_{i,j}^{kH, kW} \mathbb{I}_{\tau_{t-1} \in T_{ij}^{t-1}}} \right),$$

$$\omega_t = \omega_{t-1} + 1, \quad (2)$$

where \mathbb{I}_E is an indicator function being one, if E is true, and zero otherwise. The texel weights are initialized to $\omega_0 = 0$.

The new state S_t is optimally computed in 2D space without interrupting the gradient flow. This way, the scene is seamlessly stored in RAM or disk and can only be partially loaded and updated, which is crucial for scalability.

Appearance Rendering Module. In a first step, this module extracts compressed scene content S_t using Π and embeds these features into a higher dimensional space via a transposed linear compressor (*Feature Decompression* W^T) which acts as a regularizer. Pre-computed meta features \hat{M}_{t-1} are optionally appended and all features are downsampled by a custom masked average pooling with a stride of k and $k \times k$ kernel size, where the mask indicates which features to ignore (features that are empty or located outside the scene space). The current $H \times W$ resolution feature map is passed through the seven-layer convolutional *Rendering Network* refining features and filling potential holes that occur when the scene representation is sparsely populated. Lastly, the high-level features are decoded to RGB values by (*Feature Decoder*) three linear layers interleaved with Leaky ReLU activation functions. The final output is activated using HardTanh activation for generating valid normalized RGB values.

Loss and Optimization. The entire pipeline is trained end-to-end from scratch until convergence using the reprojection error between the rendered image \hat{I}_t and input image I_t as self-supervision. Thus, the network can learn to optimally fuse and encode appearance information from 2D training data without any ground-truth textures. Our pipeline is trained using a weighted combination of \mathcal{L}_1 and \mathcal{L}_2 loss between input image I_t and rendered image \hat{I}_t given by

$$\mathcal{L}(I_t, \hat{I}_t) = \frac{1}{C \cdot H \cdot W} \sum_{p \in I_t, \hat{p} \in \hat{I}_t} \|p - \hat{p}\|_1 + \frac{1}{2} \|p - \hat{p}\|_2 \quad (3)$$

We empirically found that a $1 : \frac{1}{2}$ weight ratio worked best in our experiments. The entire pipeline has less than 0.6M parameters and was optimized using the Adam optimizer [34] with a learning rate of 10^{-4} and batch size 1, except for the generalization experiment, where we used 2. Please see the supplementary material for more details.

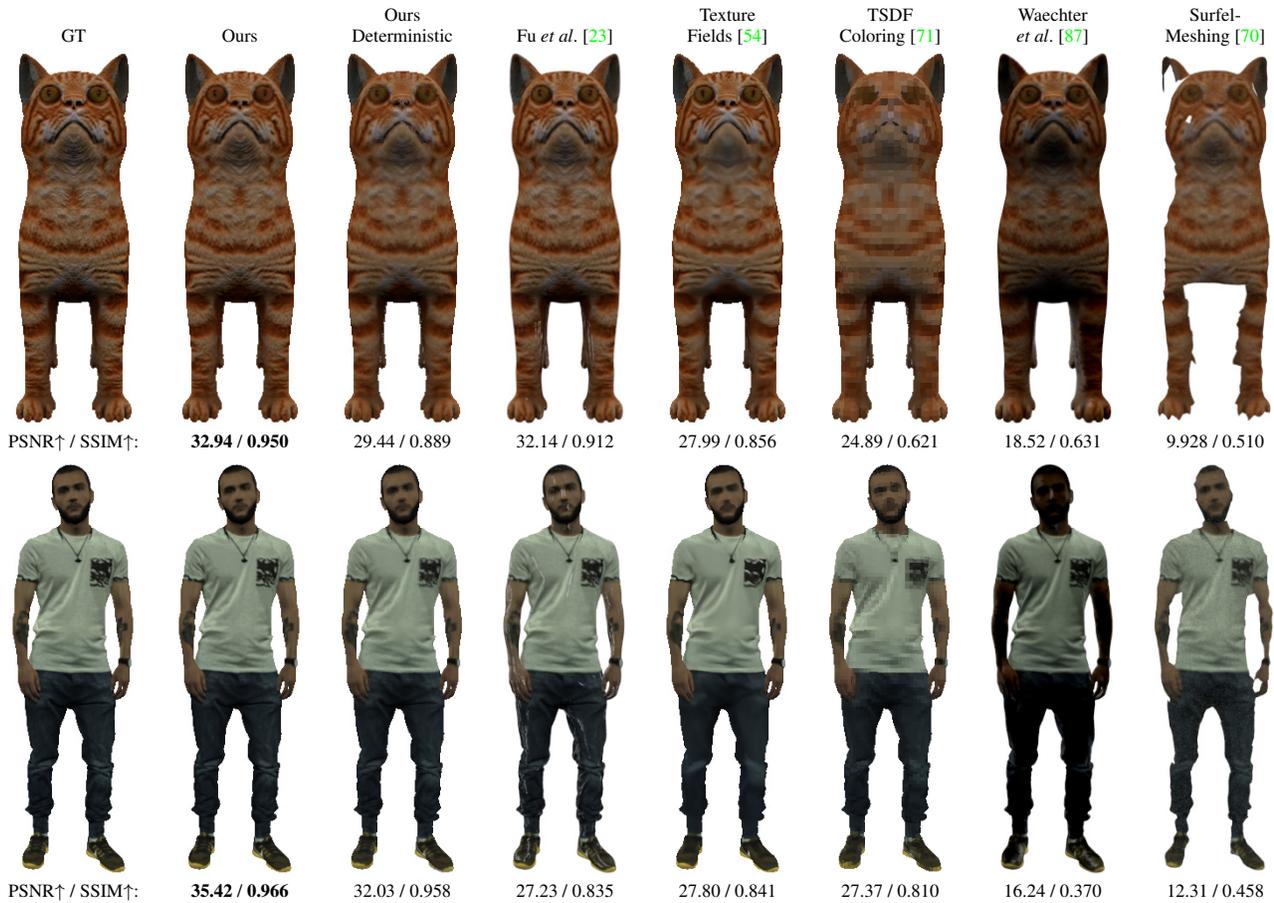


Figure 4. **Qualitative and quantitative comparison** on novel view synthesis with DeepSurfels on a 128^3 sparse grid with learned 3-channel 4×4 feature patches. The experiment demonstrates that our scene representation is able to better represent high-frequency textures compared to other state-of-the-art methods. "Ours deterministic" shows direct rendering from RGB surfel patches. Please note that SurfelMeshing [70] is the only method in this comparison which also estimates geometry while the other methods use known geometry.

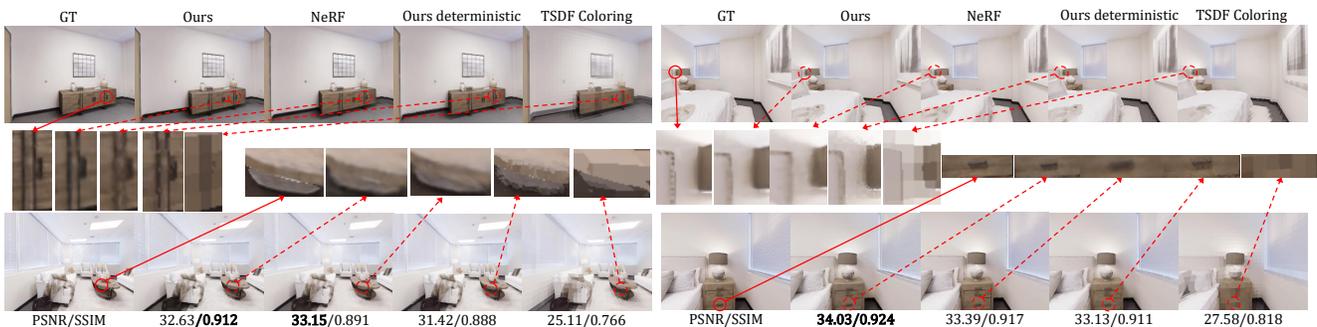


Figure 5. **Novel view synthesis for Replica [76] indoor scenes.** The figure shows different views on two scenes (left and right). Our learned approach has been trained only on the room on the left. NeRF [49] is optimized separately on both scenes.

5. Evaluation

We evaluate our method by comparing the representation power of both, our learned and our deterministic approach (direct rendering from RGB surfels) with state-of-the-art methods on novel view synthesis tasks. We further demon-

strate how our method generalizes to different scenes for a small number of distinct training samples and provide an ablation study to validate the design choices for our model. The supplementary material provides further details.

Datasets. We conduct experiments on datasets generated from Shapenet [12], publicly available human and

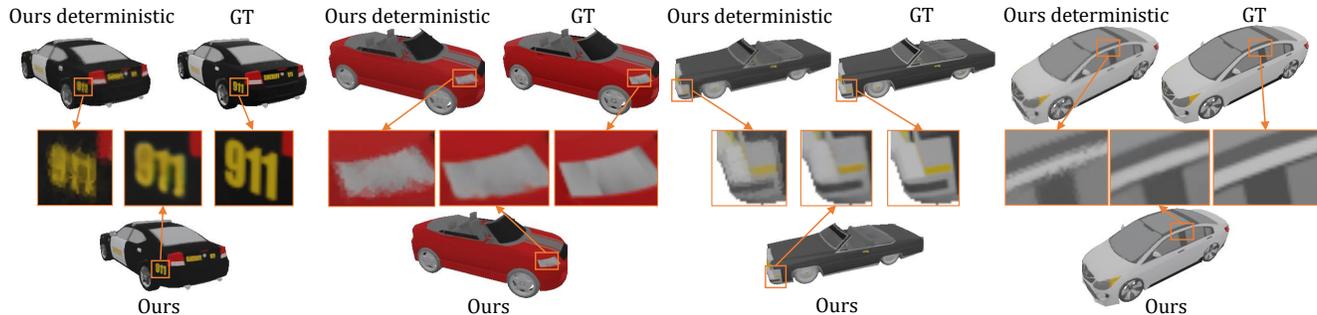


Figure 6. **Qualitative results of our model on unseen scenes from ShapeNet [12].** Compared to deterministic rendering from RGB values, DeepSurfels with learned (3+3)-channel 6×6 patches on a sparse 32^3 grid yields significantly more high-frequency details. As it is shown in the close ups, storing learned features in the texels is particularly useful to correct discretization artefacts.

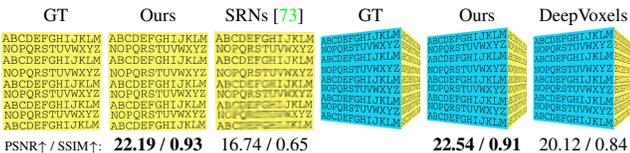


Figure 7. **Comparison of SRNs [73] and DeepVoxels [72] to our learned DeepSurfel fusion with a 64^3 grid of 8-channel 1×1 resolution feature patches on the synthetic cube dataset from [72].** Our method produces fewer blur artifacts and multi-view inconsistencies and overall yields significantly better images reconstruction results than both baselines. Note that both baselines perform global appearance fusion with unknown geometry.

cat¹ models, the indoor Replica dataset [76], and the cube scene from [72]. Replica dataset images were rendered with Habitat-Sim [67] and all other models with Blender [15].

Metrics. We quantify model performances with the following two metrics [90]. **PSNR:** The Peak Signal-to-Noise Ratio is the ratio between the maximum pixel value in the ground-truth image and the pixel-wise mean-squared error between ground-truth and rendered image. **SSIM:** The Structural Similarity Index measures similarity between patches of rendered and ground-truth images. We omit other perceptron based metrics because we are interested in recovering the true pixel value as our fusion approach can be used for more general types of data.

Novel View Synthesis. The model is optimized on 500 randomly rendered 512×512 training images for the cat and human model and the results for a single unseen frontal viewpoint² are compared with state-of-the-art batch (Fu *et al.* [23], Texture Fields [54], Waechter *et al.* [87]), and on-line methods (SurfelMeshing [70], TSDF Coloring [17]) on a 128^3 grid. The results in Figure 4 demonstrate that our approach compares favorably even to slower batch-based methods in representing high-frequency textures. Figure 7 further shows that our approach does not suffer from blurry

artifacts as the recently proposed SRNs [73], or from multi-view consistency issues like DeepVoxels [72]. Note that these approaches jointly estimate geometry and appearance while we only estimate appearance. Table 2 shows the effect of varying the number of channels on the cube dataset for DeepSurfels of 4×4 patches on a 64^3 sparse grid.

Generalization. Our pipeline scales and generalizes well on realistic room-size scenes. We trained our pipeline on 288 480×640 images of one Replica [76] room represented with DeepSurfels of 11cm voxel size with (3+3)-channel 6×6 resolution patches. We disentangled 3 color channels to improve generalization. The pipeline evaluation is performed on every 25th unseen frame in a sequence of frames generated by a moving agent in the Habitat Sim [67]. Figure 5 shows results of our trained pipeline on an optimized (left) and a non-optimized (right) scene. Our learned approach outperforms baselines in representing fine details.

We further demonstrate that our pipeline generalizes well when trained on a larger set of distinct scenes. We render 100 312×312 training images from 150 Shapenet [12] car scenes and test the pipeline on 50 unseen scenes by fusing 100 views and evaluating results on additional 60 unseen viewpoints. The whole pipeline is trained to be frame order independent by randomly shuffling scenes and frames after each optimization step. Results on test scenes (Figure 6, Table 1) indicate that our learned approach improves for discretization artifacts and overall yields sharper results which are supported by higher PSNR and SSIM scores.

Ablation Study. For the unobserved test car scenes, we quantify in Table 1 the impact of: (i) depth as a meta feature that helps our method to reason about the confidence of updates since pixels with larger depth values are less important; (ii) multi-view consistency regularization that corrects for geometric misalignments and improves interpolation among neighboring viewpoints by adding an additional error signal (3) for a viewpoint closest to the fused frame; (iii) pixel ray directions with surface orientation map to improve reasoning about light information and non-Lambertian sur-

¹3D models from free3d.com and turbosquid.com.

²For a fair comparison with the results of Texture Fields [54].

Method	PSNR \uparrow	SSIM \uparrow
SurfelMeshing [70]	13.92	0.2748
Wachter <i>et al.</i> [87]	18.27	0.4753
Fu <i>et al.</i> [23]	18.84	0.5196
TSDF Coloring [17] (32^3)	21.57	0.6375
TSDF Coloring [17] (64^3)	24.05	0.7552
TSDF Coloring [17] (128^3)	26.68	0.8526
Ours Det. (32^3 , 6×6 , 3)	27.20	0.8723
Ours Det. (64^3 , 4×4 , 3)	28.73	0.9036
<hr/>		
Learned (32^3 , 6×6 , $3+3$)	28.27	0.8777
+ depth	28.31	0.8782
+ multi-view consist.	28.36	0.8889
+ viewing direction & surface orientation	28.89	0.8907
<hr/>		
32^3 , 1×1 , $213+3$	22.95	0.7083
64^3 , 1×1 , $213+3$	25.41	0.7940
64^3 , 4×4 , $3+3$	29.92	0.9086
64^3 , 5×5 , $3+3$	30.15	0.9126
64^3 , 6×6 , $3+3$	30.27	0.9147
128^3 , 1×1 , $213+3$	26.75	0.8324
128^3 , 2×2 , $3+3$	30.23	0.9133
128^3 , 3×3 , $3+3$	30.51	0.9181
128^3 , 4×4 , $3+3$	30.60	0.9196
128^3 , 5×5 , $3+3$	30.63	0.9200
128^3 , 6×6 , $3+3$	30.64	0.9202

#Channels	PSNR \uparrow	SSIM \uparrow
2	25.95	0.9432
4	26.72	0.9506
6	27.33	0.9568
10	28.27	0.9638

Table 1. **Ablation study** on ShapeNet [12] cars. The **top part** of the table compares various baselines. Our deterministic coloring at 32^3 is still better than TSDF Coloring at 128^3 resolution. The **mid part** shows the impact of the proposed losses. The **bottom part** shows the influence of the voxel grid, surfel patch and channel resolution, demonstrating that quality improvements saturate for higher resolutions. $3+3$ denotes 3 feature and 3 color channels (disentangled) per texel. We also compare to 1×1 patches with $213+3$ channels corresponding to the same number of features as for 6×6 , $3+3$, demonstrating the benefit of a spatial sub-feature alignment in our network.

Table 2. **Varying number of feature channels** for the cube [72] dataset on 64^3 sparse grid with 4×4 patches. Additional feature channels improve the reconstruction quality.

faces; (iv) DeepSurfel parameters (grid and patch resolution). For almost all experiments use 3 feature and 3 disentangled color channels (denoted as $3+3$ in Table 1) which outperforms the baselines. We observe that every attribute (i-iii) improves generalization and that higher DeepSurfel resolution (iv) consistently benefits reconstruction quality. (v) Lastly, we demonstrate the benefit of explicitly modeling the reprojection of pixel colors via texels which can be seen as subfeatures of a large feature vector stored in every voxel. We compare the proposed 6×6 patches, $3+3$ channels to 1×1 patches, $213+3$ channels which amounts to the same number of features per voxel. We argue that the benefit of using more features per voxel quickly saturates unless subfeatures are anchored along the surface and trained with separate pixel data as supported by our results.

Real-world data. In Fig. 8 we present results on unseen real-world data from [45] for which our method yields the most detailed appearance reconstructions.

Runtime. Our method takes 57ms and 21ms for fusing and rendering a single 312×312 frame on 32^3 DeepSurfels with 6-channel patches with resolution 6×6 (Table 1). This is significantly faster compared to other deep learning methods that overfit on a single scene. For example, the state-

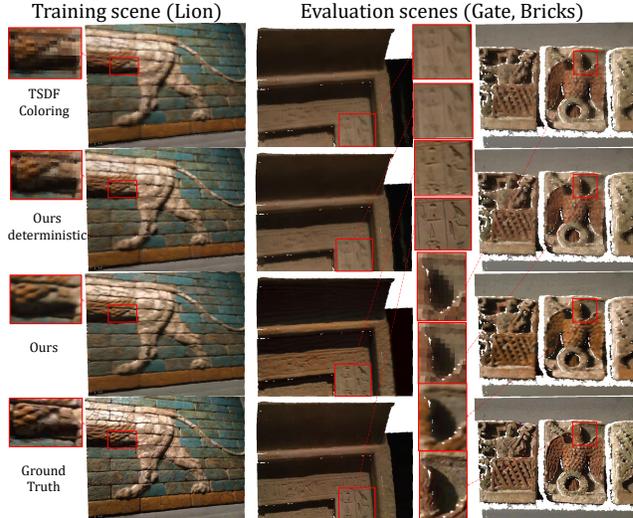


Figure 8. **Novel-view synthesis on unseen real-world data [45].** Our DeepSurfel method with 4×4 patches is trained on the *Lion* scene for 80 training iterations and then evaluated on the *Gate* and the *Bricks* scenes. The images show novel viewpoints.

of-the-art method NeRF [49] requires ~ 2 days for training on a single scene being unable to generalize to other scenes, while our method can easily be used on unseen scenes without any optimization as demonstrated in Figure 5, which is a speed up of over a thousand times on unobserved scenes for comparable or even favorable results.

Limitations. Similar to traditional methods like TSDF Coloring, our method is sensitive to camera-geometry misalignment which can lead to blurry results. Moreover, we currently do not model view-dependent effects which may result in washed-out colors due to the local feature averaging. When trained on small datasets our method may distort colors that were not seen during training.

6. Conclusion

We introduced DeepSurfels, a novel scene representation for geometry and appearance encoding, that combines explicit and implicit scene representation to improve for scalability and interpretability. It is defined on a sparse voxel grid to maintain topology relations and implements 2D geometry oriented patches to store high-frequency appearance information. We further presented a learned approach for online appearance fusion that compares favorably to existing offline and online texture mapping methods since it learns to correct for typical noise and discretization artifacts.

As future work we consider the joint online fusion of shape and appearance and address some weaknesses of our appearance fusion pipeline such as the limitation in filling large missing parts and rendering translucent surfaces.

Acknowledgments. This research was partly supported by Toshiba and Innosuisse funding (Grant No. 34475.1 IP-ICT).

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *Proc. International Conference on Machine Learning (ICML)*, 2018. 2
- [2] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 3
- [3] Cédric Allène, Jean-Philippe Pons, and Renaud Keriven. Seamless image-based texture atlases using multi-band blending. In *Proc. International Conference on Computer Vision (ICCV)*, 2008. 3
- [4] Matthieu Armando, Jean-Sébastien Franco, and Edmond Boyer. Adaptive mesh texture for multi-view appearance modeling. In *Proc. International Conference on 3D Vision (3DV)*, 2019. 3
- [5] Fausto Bernardini, Ioana M. Martin, and Holly E. Rushmeier. High-quality texture reconstruction from multiple scans. *IEEE Transactions on Visualization and Computer Graphics*, 2001. 3
- [6] Sai Bi, Nima Khademi Kalantari, and Ravi Ramamoorthi. Patch-based optimization for image-based texture mapping. *ACM Transactions on Graphics*, 2017. 3
- [7] 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 *Proc. European Conference on Computer Vision (ECCV)*, 2020. 3
- [8] Andreas Bircher, Kostas Alexis, Michael Burri, Philipp Oettershagen, Sammy Omari, Thomas Mantel, and Roland Siegwart. Structural inspection path planning via iterative viewpoint resampling with application to aerial robotics. In *IEEE International Conference on Robotics and Automation*, 2015. 1
- [9] Andreas Breitenmoser and Roland Siegwart. Surface reconstruction and path planning for industrial inspection with a climbing robot. In *Proc. International Conference on Applied Robotics for the Power Industry (CARPI)*, 2012. 1
- [10] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016. 2
- [11] Rohan Chabra, Jan Eric Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 3
- [12] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 6, 7, 8
- [13] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [14] Hang Chu, Shugao Ma, Fernando De la Torre, Sanja Fidler, and Yaser Sheikh. Expressive telepresence via modular codec avatars. In *Proc. European Conference on Computer Vision (ECCV)*, Lecture Notes in Computer Science, 2020. 1
- [15] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2020. 7
- [16] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 1984. 4
- [17] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 1996. 3, 7, 8
- [18] Paul E Debevec, Camillo J Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 1996. 1, 2, 3
- [19] M. Eisemann, B. De Decker, M. Magnor, P. Bekaert, E. de Aguiar, N. Ahmed, C. Theobalt, and A. Sellent. Floating Textures. *Computer Graphics Forum*, 2008. 1, 2, 3
- [20] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [21] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [22] Yanping Fu, Qingan Yan, Jie Liao, and Chunxia Xiao. Joint texture and geometry optimization for RGB-D reconstruction. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [23] Yanping Fu, Qingan Yan, Long Yang, Jie Liao, and Chunxia Xiao. Texture mapping for 3d reconstruction with RGB-D sensor. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 3, 6, 7, 8
- [24] Matheus Gadelha, Subhransu Maji, and Rui Wang. 3d shape induction from 2d views of multiple objects. In *Proc. International Conference on 3D Vision (3DV)*, 2017. 2
- [25] Ran Gal, Yonatan Wexler, Eyal Ofek, Hugues Hoppe, and Daniel Cohen-Or. Seamless montage for texturing models. *Computer Graphics Forum*, 2010. 3
- [26] Santiago Garrido, María Malfaz, and Dolores Blanco. Application of the fast marching method for outdoor motion planning in robotics. *Robotics and Autonomous Systems*, 2013. 1
- [27] Bastian Goldlücke, Mathieu Aubry, Kalin Kolev, and Daniel Cremers. A super-resolution framework for high-accuracy multiview reconstruction. *International Journal of Computer Vision*, 2014. 3
- [28] Christian Häne, Lionel Heng, Gim Hee Lee, Friedrich Fraundorfer, Paul Furgale, Torsten Sattler, and Marc Polle-

- feys. 3d visual perception for self-driving cars using a multi-camera system: Calibration, mapping, localization, and obstacle detection. *Image and Vision Computing*, 2017. [1](#)
- [29] Rana Hanocka, Gal Metzer, Raja Giryes, and Daniel Cohen-Or. Point2mesh: A self-prior for deformable meshes. *ACM Transactions on Graphics*, 2020. [2](#)
- [30] Jingwei Huang, Justus Thies, Angela Dai, Abhijit Kundu, Chiyu Jiang, Leonidas J Guibas, Matthias Nießner, and Thomas Funkhouser. Adversarial texture optimization from RGB-D scans. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [3](#)
- [31] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, and Thomas Funkhouser. Local implicit grid representations for 3d scenes. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [3](#)
- [32] Olaf Kähler, Victor Adrian Prisacariu, Carl Yuheng Ren, Xin Sun, Philip Torr, and David Murray. Very high frame rate volumetric integration of depth images on mobile devices. *IEEE Transactions on Visualization and Computer Graphics*, 2015. [3](#)
- [33] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. [2](#)
- [34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. International Conference on Learning Representations (ICLR)*, 2015. [5](#)
- [35] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. In *Proc. International Conference on Computer Vision (ICCV)*, 1999. [3](#)
- [36] Joo Ho Lee, Hyunho Ha, Yue Dong, Xin Tong, and Min H Kim. Texturefusion: High-quality texture acquisition for real-time rgb-d scanning. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [3](#)
- [37] Victor Lempitsky and Denis Ivanov. Seamless mosaicing of image-based texture maps. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007. [3](#)
- [38] Hendrik P. A. Lensch, Wolfgang Heidrich, and Hans-Peter Seidel. A silhouette-based algorithm for texture registration and stitching. *Graphical Models*, 2001. [3](#)
- [39] Yawei Li, Vagia Tsiminaki, Radu Timofte, Marc Pollefeys, and Luc Van Gool. 3d appearance super-resolution with deep learning. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [3](#)
- [40] Yiyi Liao, Simon Donne, and Andreas Geiger. Deep marching cubes: Learning explicit surface representations. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [41] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proc. International Conference on Computer Vision (ICCV)*, 2019. [2](#)
- [42] Z. Liu, Y. Cao, Z. Kuang, L. Kobbelt, and S. Hu. High-quality textured 3d shape reconstruction with cascaded fully convolutional networks. *IEEE Transactions on Visualization and Computer Graphics*, 2019. [3](#)
- [43] Stephen Lombardi, Jason M. Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Transactions on Graphics*, 2018. [1](#)
- [44] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics*, 2019. [1, 3](#)
- [45] Robert Maier, Kihwan Kim, Daniel Cremers, Jan Kautz, and Matthias Nießner. Intrinsic3d: High-quality 3d reconstruction by joint appearance and geometry optimization with spatially-varying lighting. In *Proc. International Conference on Computer Vision (ICCV)*, 2017. [3, 8](#)
- [46] Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [3](#)
- [47] Marko Mihajlovic, Yan Zhang, Michael J. Black, and Siyu Tang. LEAP: Learning articulated occupancy of people. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [3](#)
- [48] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics*, 2019. [3](#)
- [49] 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 *Proc. European Conference on Computer Vision (ECCV)*, 2020. [1, 2, 3, 6, 8](#)
- [50] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-time dense surface mapping and tracking. In *IEEE International Symposium on Mixed and Augmented Reality*, 2011. [1, 3, 4](#)
- [51] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. Dtam: Dense tracking and mapping in real-time. In *Proc. International Conference on Computer Vision (ICCV)*, 2011. [1](#)
- [52] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [3](#)
- [53] Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Marc Stamminger. Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics*, 2013. [3](#)
- [54] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proc. International Conference on Computer Vision (ICCV)*, 2019. [1, 3, 6, 7](#)
- [55] Michael Oechsle, Michael Niemeyer, Christian Reiser, Lars Mescheder, Thilo Strauss, and Andreas Geiger. Learning implicit surface light fields. In *Proc. International Confer-*

- ence on 3D Vision (3DV), 2020. 3
- [56] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3
- [57] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 3
- [58] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics*, 2017. 3
- [59] Hanspeter Pfister, Matthias Zwicker, Jeroen Van Baar, and Markus Gross. Surfels: Surface elements as rendering primitives. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 2000. 3
- [60] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *Proc. International Conference on Machine Learning (ICML)*, 2019. 3
- [61] Konstantinos Rematas and Vittorio Ferrari. Neural voxel renderer: Learning an accurate and controllable rendering tool. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [62] Danilo Jimenez Rezende, SM Ali Eslami, Shakir Mohamed, Peter Battaglia, Max Jaderberg, and Nicolas Heess. Unsupervised learning of 3d structure from images. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2016. 2
- [63] Audrey Richard, Ian Cherabier, Martin R. Oswald, Vagia Tsiminaki, Marc Pollefeys, and Konrad Schindler. Learned multi-view texture super-resolution. In *Proc. International Conference on 3D Vision (3DV)*, 2019. 3
- [64] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 3
- [65] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proc. International Conference on Computer Vision (ICCV)*, 2019. 3
- [66] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [67] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *Proc. International Conference on Computer Vision (ICCV)*, 2019. 7
- [68] Thomas Schöps, Martin R Oswald, Pablo Speciale, Shuoran Yang, and Marc Pollefeys. Real-time view correction for mobile devices. *IEEE Transactions on Visualization and Computer Graphics*, 2017. 1
- [69] Thomas Schöps, Torsten Sattler, Christian Häne, and Marc Pollefeys. Large-scale outdoor 3d reconstruction on a mobile device. *Computer Vision and Image Understanding*, 2017. 1
- [70] T. Schöps, T. Sattler, and M. Pollefeys. SurfelMeshing: Online surfel-based mesh reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 3, 4, 6, 7, 8
- [71] Steven M Seitz and Charles R Dyer. Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision*, 1999. 3, 6
- [72] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. DeepVoxels: Learning persistent 3d feature embeddings. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 7, 8
- [73] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2, 3, 7
- [74] Pratul P Srinivasan, Richard Tucker, Jonathan T Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. Pushing the boundaries of view extrapolation with multiplane images. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [75] Frank Steinbrucker, Christian Kerl, and Daniel Cremers. Large-scale multi-resolution surface reconstruction from rgb-d sequences. In *Proc. International Conference on Computer Vision (ICCV)*, 2013. 3
- [76] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 6, 7
- [77] David Stutz and Andreas Geiger. Learning 3d shape completion from laser scan data with weak supervision. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [78] Richard Szeliski and Polina Golland. Stereo matching with transparency and matting. In *Proc. International Conference on Computer Vision (ICCV)*, 1998. 3
- [79] Takeshi Takai, Adrian Hilton, and Takashi Mastuyama. Harmonised texture mapping. In *Proc. International Conference on 3D Vision (3DV)*, 2010. 3
- [80] Marco Tarini, Cem Yuksel, and Silvain Lefebvre. Rethinking texture mapping. In *ACM SIGGRAPH 2017 Courses*, 2017. 3
- [81] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner, R. Pandey, S. Fanello, G. Wetzstein, J.-Y. Zhu, C. Theobalt, M. Agrawala, E. Shechtman, D. B Goldman, and M. Zollhöfer. State of the art on neural rendering. *Computer Graphics Forum*, 2020. 3
- [82] Christian Theobalt, Naveed Ahmed, Hendrik P. A. Lensch, Marcus A. Magnor, and Hans-Peter Seidel. Seeing people

- in different light-joint shape, motion, and reflectance capture. *IEEE Transactions on Visualization and Computer Graphics*, 2007. 3
- [83] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics*, 2019. 3
- [84] Justus Thies, Michael Zollhöfer, Christian Theobalt, Marc Stamminger, and Matthias Nießner. Image-guided neural object rendering. In *Proc. International Conference on Learning Representations (ICLR)*, 2020. 3
- [85] Vagia Tsiminaki, Wei Dong, Martin R. Oswald, and Marc Pollefeys. Joint multi-view texture super-resolution and intrinsic decomposition. In *Proc. of the British Machine and Vision Conference (BMVC)*, page 15. BMVA Press, 2019. 3
- [86] Vagia Tsiminaki, Jean-Sébastien Franco, and Edmond Boyer. High resolution 3d shape texture from multiple videos. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 3
- [87] Michael Waechter, Nils Moehrle, and Michael Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *Proc. European Conference on Computer Vision (ECCV)*, 2014. 1, 2, 3, 6, 7, 8
- [88] Kaixuan Wang, Fei Gao, and Shaojie Shen. Real-time scalable dense surfel mapping. In *IEEE International Conference on Robotics and Automation*, 2019. 3, 4
- [89] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 2
- [90] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004. 7
- [91] Silvan Weder, Johannes L Schönberger, Marc Pollefeys, and Martin R Oswald. RoutedFusion: Learning real-time depth map fusion. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [92] Silvan Weder, Johannes L Schönberger, Marc Pollefeys, and Martin R Oswald. NeuralFusion: Online depth fusion in latent space. In *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [93] Thomas Whelan, Michael Kaess, Hordur Johannsson, Maurice Fallon, John J Leonard, and John McDonald. Real-time large-scale dense rgb-d slam with volumetric fusion. *International Journal of Robotics Research*, 2015. 1
- [94] Thomas Whelan, Renato F. Salas-Moreno, Ben Glocker, Andrew J. Davison, and Stefan Leutenegger. Elasticfusion: Real-time dense SLAM and light source estimation. *International Journal of Robotics Research*, 2016. 3, 4
- [95] Daniel N Wood, Daniel I Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 2000. 3
- [96] Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Interpretable transformations with encoder-decoder networks. In *Proc. International Conference on Computer Vision (ICCV)*, 2017. 3
- [97] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2016. 2
- [98] Qiengeng Xu, Weiyue Wang, Duygu Ceylan, Radomír Mech, and Ulrich Neumann. DISN: deep implicit surface network for high-quality single-view 3d reconstruction. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2019. 3
- [99] Cem Yuksel, John Keyser, and Donald H. House. Mesh colors. *ACM Transactions on Graphics*, 2010. 3
- [100] Cem Yuksel, Sylvain Lefebvre, and Marco Tarini. Rethinking texture mapping. In *Computer Graphics Forum*, 2019. 3
- [101] Ming Zeng, Fukai Zhao, Jiayang Zheng, and Xinguo Liu. A memory-efficient kinectfusion using octree. In *Proc. International Conference on Computational Visual Media*, 2012. 3
- [102] Ming Zeng, Fukai Zhao, Jiayang Zheng, and Xinguo Liu. Octree-based fusion for realtime 3d reconstruction. *Graphical Models*, 2013. 3
- [103] Xiuming Zhang, Sean Fanello, Yun-Ta Tsai, Tiancheng Sun, Tianfan Xue, Rohit Pandey, Sergio Orts-Escolano, Philip Davidson, Christoph Rhemann, Paul Debevec, et al. Neural light transport for relighting and view synthesis. *ACM Transactions on Graphics*, 2021. 3
- [104] Michael Zollhöfer, Angela Dai, Matthias Innmann, Chenglei Wu, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Shading-based Refinement on Volumetric Signed Distance Functions. *ACM Transactions on Graphics*, 2015. 3