

# Initialization and Alignment for Adversarial Texture Optimization

Xiaoming Zhao, Zhizhen Zhao, Alexander G. Schwing

University of Illinois, Urbana-Champaign  
{xz23, zhizhenz, aschwing}@illinois.edu

[https://xiaoming-zhao.github.io/projects/advtex\\_init\\_align](https://xiaoming-zhao.github.io/projects/advtex_init_align)

**Abstract.** While recovery of geometry from image and video data has received a lot of attention in computer vision, methods to capture the texture for a given geometry are less mature. Specifically, classical methods for texture generation often assume clean geometry and reasonably well-aligned image data. While very recent methods, *e.g.*, adversarial texture optimization, better handle lower-quality data obtained from hand-held devices, we find them to still struggle frequently. To improve robustness, particularly of recent adversarial texture optimization, we develop an explicit initialization and an alignment procedure. It deals with complex geometry due to a robust mapping of the geometry to the texture map and a hard-assignment-based initialization. It deals with misalignment of geometry and images by integrating fast image-alignment into the texture refinement optimization. We demonstrate efficacy of our texture generation on a dataset of 11 scenes with a total of 2807 frames, observing 7.8% and 11.1% relative improvements regarding perceptual and sharpness measurements.

**Keywords:** scene analysis, texture reconstruction

## 1 Introduction

Accurate scene reconstruction is one of the major goals in computer vision. Decades of research have been devoted to developing robust methods like ‘Structure from Motion,’ ‘Bundle Adjustment,’ and more recently also single view reconstruction techniques. While reconstruction of geometry from image and video data has become increasingly popular and accurate in recent years, recovered 3D models remain often pale because textures aren’t considered.

Given a reconstructed 3D model of a scene consisting of triangular faces, and given a sequence of images depicting the scene, texture mapping aims to find for each triangle a suitable texture. The problem of automatic texture mapping has been studied in different areas since late 1990 and early 2000. For instance, in the graphics community [12, 29, 30], in computer vision [45, 27, 43], architecture [19] and photogrammetry [14]. Many of the proposed algorithms work very well in a controlled lab-setting where geometry is known perfectly, or in a setting where accurate 3D models are available from a 3D laser scanner.

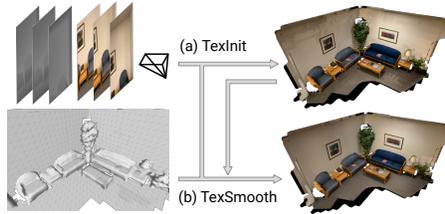


Fig. 1: We study texture generation given RGBD images with associated camera parameters as well as a reconstructed mesh. **(a) TexInit**: we initialize the texture using an assignment-based texture generation framework. **(b) TexSmooth**: a data-driven adversarial loss is utilized to optimize out artifacts incurred in the assignment step.

However, applying texture mapping techniques to noisy mesh geometry obtained *on the fly* from a recent LiDAR equipped iPad reveals missing robustness because of multiple reasons: 1) images and 3D models are often not perfectly aligned; 2) 3D models are not accurate and the obtained meshes aren’t necessarily manifold-connected. Even recent techniques for mesh flattening [41] and texturing [16, 54, 24] result in surprising artifacts due to streamed geometry and pose inaccuracies as shown later.

To address this robustness issue we find equipping of the recently-proposed adversarial texture optimization technique [24] with classical initialization and alignment techniques to be remarkably effective. Without the added initialization and alignment, we find current methods don’t produce high-quality textures. Concretely and as illustrated in Fig. 1, we aim for texture generation which operates on a sequence of images and corresponding depth maps as well as their camera parameters. Moreover, we assume the 3D model to be given and fixed. Importantly, we consider a streaming setup, with all data obtained *on the fly*, and not further processed, *e.g.*, via batch structure-from-motion. The setup is ubiquitous and the form of data can be acquired easily from consumer-grade devices these days, *e.g.*, from a recent iPad or iPhone [1]. We aim to translate this data into texture maps. For this, we first flatten the triangle mesh using recent advances [41]. We then use a Markov Random Field (MRF) to resolve overlaps in flattened meshes for non-manifold-connected data. In a next step we determine the image frame from which to extract the texture of each mesh triangle using a simple optimization. We refer to this as **TexInit**, which permits to obtain a high-quality initialization for subsequent refinement. Next we address inaccuracies in camera poses and in geometry by automatically shifting images using the result of a fast Fourier transformation (FFT) [6]. The final optimized texture is obtained by integrating this FFT-alignment component into adversarial optimization [24]. We dub this stage **TexSmooth**. The obtained texture can be used in any 3D engine for downstream applications.

To study efficacy of the proposed framework we acquire 11 complex scenes using a recent iPad. We establish accuracy of the proposed technique to generate and use the texture by showing that the quality of rendered views is superior to

prior approaches on these scenes. Quantitatively, our framework improves prior work by 7.8% and 11.1% relatively with respect to perceptual (LPIPS [53]) and sharpness ( $S_3$  [47]) measurements respectively. Besides, our framework improves over baselines on ScanNet [11], demonstrating the ability to generalize.

## 2 Related Work

We aim for accurate recovery of texture for a reconstructed 3D scene from a sequence of RGBD images. For this, a variety of techniques have been proposed, which can be roughly categorized into four groups: 1) averaging-based; 2) warping-based; 3) learning-based; and 4) assignment-based. Averaging-based methods find all views within which a point is visible and combine the color observations. Warping-based approaches either distort or synthesize source images to handle mesh misalignment or camera drift. Learning-based ones learn the texture representation. Assignment-based methods attempt to find the best view and ‘copy’ the observation into a texture. We review these groups next:

**Averaging-based:** Very early work by Marshner [29] estimates the parameters of a bidirectional reflectance distribution function (BRDF) for every point on the texture map. To compute this estimation, all observations from the recorded images where the point is visible are used. Similar techniques have been investigated in subsequent work [9].

Similarly, to compute a texture map, [30] and [12] perform a weighted blending of all recorded images. The weights take visibility and other factors into account. The developed approaches are semi-automatic as they require an initial estimate of the camera orientations which is obtained from interactively selected point correspondences or marked lines. Multi-resolution textures [33], face textures [36] and blending [31, 8] have also been studied.

**Warping-based:** Aganj *et al.* [4] morph each source image to align to the mesh. Furthermore, [54, 23] propose to optimize camera poses and image warping parameters jointly. However, this line of vertex-based optimization has stringent requirements on the mesh density and cannot be applied to a sparse mesh. More recently, Bi *et al.* [10] follow patch-synthesis [50, 7] to re-synthesize aligned target images for each source view. However, such methods require costly multiscale optimization to avoid a large number of local optima. In contrast, the proposed approach does not require those techniques.

**Learning-based:** Recently, learning-based methods have been introduced for texture optimization. Some works focus on specific object and scene categories [18, 40] while we do not make such assumptions. Moreover, learned representations, *e.g.*, neural textures, have also been developed [46, 44, 5]. Meanwhile, generative models are developed to synthesize a holistic texture from a single image or pattern [21, 32] while we focus on texture reconstruction. AtlasNet [20] and NeuTex [52] focus on learning a 3D-to-2D mapping, which can be utilized in texture editing, while we focus on reconstructing realistic textures from source images. The recently-proposed adversarial texture optimization [24] utilizes adversarial training to reconstruct the texture. However, despite advances, adver-

serial optimization still struggles with misalignments. We improve this shortcoming via an explicit high-quality initialization and an efficient alignment module. **Assignment-based:** Classical assignment-based methods operated within controlled environments [38, 39, 14, 15] or utilized special camera rigs [15, 19]. These works suggest computing for each vertex a set of ‘valid’ images, which are subsequently refined by iterating over each vertex and adjusting the assignment to obtain more consistency. Finally, texture data is ‘copied’ from the images. In contrast, we aim to create a texture in an uncontrolled setting. Consequently, 3D geometry is not accurate and very noisy. Other early work [13, 28, 22, 51, 39] focuses on closed surfaces and small-scale meshes, making them not applicable to our setting. More recently, upon finding the best texture independently for each face using cues like visibility, orientation, resolution, and distortion, refinement techniques like texture coordinate optimization, color adjustments, or scores-based optimization have been discussed [34, 48, 3].

Related to our approach are methods that formulate texture selection using a Markov Random Field (MRF) [27, 42, 16]. Shu *et al.* [42] suggest visibility as the data term and employ texture smoothness to reduce transitions. Lempitsky *et al.* [27] study color-continuity which integrates over face seams. Fu *et al.* [16] additionally use the projected 2D face area to select a texture assignment for each face. However, noisy geometry like the one shown in Fig. 2, makes it difficult for assignment-based methods to yield high quality results, which we will show later. Therefore, different from these methods, we address texture drift in a data-driven refinement procedure rather than in an assignment stage.

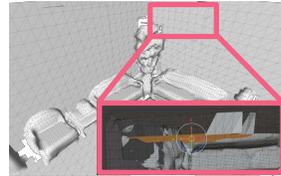


Fig. 2: Noisy geometry. The wall has two layers.

### 3 Approach

We want to automatically create the texture  $\mathcal{T}$  from a set of RGBD images  $I = \{I_1, \dots, I_T\}$ , for each of which we also know camera parameters  $\{p_t\}_{t=1}^T$ , *i.e.*, extrinsics and intrinsics. We are also given a triangular scene mesh  $M = \{\text{Tri}_i\}_{i=1}^{|M|}$ , where  $\text{Tri}_i$  denotes the  $i$ -th triangle. This form of data is easily accessible from commercially available consumer devices, *e.g.*, a recent iPhone or iPad.

We construct the texture  $\mathcal{T}$  in two steps that combine advantages of assignment-based and learning-based techniques: 1) **TexInit**: we generate a texture initialization  $\mathcal{T}_{\text{init}} \in \mathbb{R}^{H \times W \times 3}$  of height  $H$ , width  $W$  and 3 color channels in an assignment-based manner (Sec. 3.1); 2) **TexSmooth**: we then refine  $\mathcal{T}_{\text{init}}$  with an improved data-driven adversarial optimization that integrates an efficient alignment procedure (Sec. 3.2). Formally, the final texture  $\mathcal{T}$  is computed via

$$\mathcal{T} = \text{TexSmooth}(\mathcal{T}_{\text{init}}, \{I_t\}_{t=1}^T, \{p_t\}_{t=1}^T, M),$$

$$\text{where } \mathcal{T}_{\text{init}} = \text{TexInit}(\{I_t\}_{t=1}^T, \{p_t\}_{t=1}^T, M). \quad (1)$$

We detail each component next.

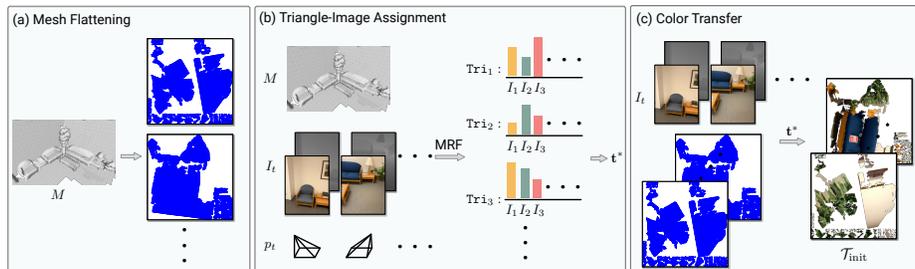


Fig. 3: **Texture initialization TexInit** (Sec. 3.1). (a) Mesh flattening: we flatten a 3D mesh into the 2D plane using overlap detection. (b) Triangle-image assignment: we develop a simple formulation to compute the triangle-image assignment  $\mathbf{t}^*$  from mesh  $M$ , frames  $I_t$  and camera parameters  $p_t$ . We assign frames to each triangle  $Tri_i$  based on  $\mathbf{t}^*$ . (c) Color transfer: based on the flattened mesh in (a) and the best assignment  $\mathbf{t}^*$  from (b), we generate the texture  $\mathcal{T}_{init}$ .

### 3.1 Texture Initialization (TexInit)

The proposed approach to obtain the texture initialization  $\mathcal{T}_{init}$  is outlined in Fig. 3 and consists of following three steps: 1) We flatten the provided mesh  $M$ . For this we detect overlaps within the flattened mesh, which may happen due to the fact that we operate with general meshes that are not guaranteed to have a manifold connectivity. Overlap detection ensures that every triangle is assigned a unique position in the texture. 2) We identify for each triangle the ‘best’ texture index  $\mathbf{t}^*$ . Hereby, ‘best’ is defined using cues like visibility and color consistency. 3) After identifying the index  $\mathbf{t}^* = (t_1^*, \dots, t_{|M|}^*)$  for each triangle, we create the texture  $\mathcal{T}_{init}$  by transferring for all  $(u, v) \in [1, \dots, W] \times [1, \dots, H]$  locations in the texture, the RGB data from the corresponding location  $(a, b)$  in image  $I_t$ .

**1) Mesh Flattening:** In a first step, as illustrated in Fig. 3 (left), we flatten the given mesh  $M$ . For this we use the recently proposed boundary first flattening (BFF) technique [41]. The flattening is fully automatic, with distortion mathematically guaranteed to be as low or lower than any other conformal mapping.

However, despite those guarantees, BFF still requires meshes to have a manifold connectivity. While we augment work by [41] using vertex duplication to circumvent this restriction, flattening may still result in overlapping regions as illustrated in Fig. 4. To fix this and uniquely assign a triangle to a position in the texture, we perform overlap detection as discussed next.

**Overlap Detection:** Overlap detection operates on flattened and possibly overlapping triangle meshes like the one illustrated in Fig. 4a. Our goal is to assign triangles to different planes. Upon re-packing the triangles assigned to different planes, we obtain the non-overlapping triangle mesh illustrated in Fig. 4b.

In order to not break the triangle mesh at a random position and end up with many individual triangles, *i.e.*, in order to maintain large triangle connectivity, we formulate this problem using a Markov Random Field (MRF). Formally, let the discrete variable  $y_i \in \mathcal{Y} = \{1, \dots, |\mathcal{Y}|\}$  denote the discrete plane index

that the  $i$ -th triangle  $\text{Tri}_i$  is assigned to. Hereby,  $|\mathcal{Y}|$  denotes the maximum number of planes which is identical to the maximum number of triangles that overlap initially at any one location. We obtain the triangle-plane assignment  $y^* = (y_1^*, \dots, y_{|M|}^*)$  for all  $|M|$  triangles by addressing

$$y^* = \arg \max_y \sum_{i=1}^{|M|} \phi_i(y_i) + \sum_{(i,j) \in \mathcal{A} \cup \mathcal{O}} \phi_{i,j}(y_i, y_j), \quad (2)$$

where  $\mathcal{A}$  and  $\mathcal{O}$  are sets of triangle index pairs which are adjacent and overlapping respectively. Here,  $\phi_i(\cdot)$  denotes triangle  $\text{Tri}_i$ 's priority over  $\mathcal{Y}$  when considering only its *local* information, while  $\phi_{i,j}(\cdot)$  refers to  $\text{Tri}_i$  and  $\text{Tri}_j$ 's joint preference on their assignments. Eq. (2) is solved with belief propagation [17].

Intuitively, by addressing the program given in Eq. (2) we want a different plane index for overlapping triangles, while encouraging mesh  $M$ 's adjacent triangles to be placed on the same plane. To achieve this we use

$$\phi_i(y_i) = \begin{cases} 1.0, & \text{if } y_i = \min \mathcal{Y}_{i, \text{non-overlap}}, \text{ and} \\ 0.0, & \text{otherwise} \end{cases}, \quad \text{and} \quad (3)$$

$$\phi_{i,j}(y_i, y_j) = \begin{cases} \mathbb{1}\{y_i = y_j\}, & \text{if } (i, j) \in \mathcal{A} \\ \mathbb{1}\{y_i \neq y_j\}, & \text{if } (i, j) \in \mathcal{O} \end{cases}. \quad (4)$$

Here,  $\mathbb{1}\{\cdot\}$  denotes the indicator function and  $\mathcal{Y}_{i, \text{non-overlap}}$  contains all plane indices where  $\text{Tri}_i$  has no overlap with others. Intuitively, Eq. (3) encourages to assign the minimum of such indices to  $\text{Tri}_i$ .

As fast MRF optimizers remove most overlaps but don't provide guarantees, we add a light post-processing to manually assign the remaining few overlapping triangles to separate planes. This guarantees overlap-free results. As mentioned before, after having identified the plane assignment  $y^*$  for each triangle we use a bin packing to uniquely assign each triangle to a position in the texture. Conversely, for every texture coordinate  $u, v$  we obtain a unique triangle index

$$i = G(u, v). \quad (5)$$

A qualitative result is illustrated in Fig. 4b. Next, we identify the image which should be used to texture each triangle.

**2) Textures from Triangle-Image Assignments:** Our goal is to identify a suitable frame  $I_{t_i}$ ,  $t_i \in \{1, \dots, T+1\}$ , for each triangle  $\text{Tri}_i$ ,  $i \in \{1, \dots, |M|\}$ . Note that the  $(T+1)$ -th option  $I_{T+1}$  refers to an empty texture. We compute the texture assignments  $\mathbf{t}^* = (t_1^*, \dots, t_{|M|}^*)$  using a purely local optimization:

$$\mathbf{t}^* = \arg \max_{\mathbf{t}} \sum_{i=1}^{|M|} \psi_i(t_i). \quad (6)$$

Here  $\psi_i$  captures *unary* cues. Note, we also studied *pairwise* cues but did not observe significant improvements. Please see the Appendix for more details. Due to better efficiency, we therefore only consider unary cues. Intuitively, we want

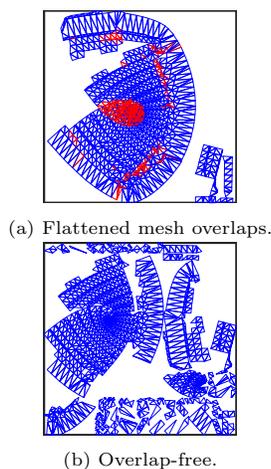


Fig. 4: **Flattening.** (a) Red triangles indicate where overlap happens. (b) The proposed method (Sec. 3.1) resolves this issue while keeping connectivity of areas.



Fig. 5: **Initialization comparison.** (a) We use PyTorch3D’s rendering pipeline [37] to project each pixel of every RGB image back to the texture. The color of each pixel in the texture is the average of all colors that project to it. This texture minimizes the  $\mathcal{L}_2$  loss of the difference between the rendered and the ground truth images. We dub it L2Avg. (b)  $\mathcal{T}_{\text{init}}$  from Sec. 3.1 permits to maintain details. The seam artifacts will be optimized out using TexSmooth (Sec. 3.2). Besides over-smoothness, without taking into account misalignments of geometry and camera poses, L2Avg produces texture that overfits to available views, *e.g.*, the sofa’s blue colors are painted onto the wall.

the program given in Eq. (6) to encourage triangle-image assignment to be ‘best’ for each triangle  $\text{Tri}_i$ . We describe the unary cues to do so next.

**Unary Potentials**  $\psi_i(t_i)$  for each pair of triangle  $\text{Tri}_i$  and frame  $I_{t_i}$  are

$$\psi_i(t_i) = \begin{cases} \psi_i^c(t_i), & \text{if } \psi_i^v(t_i) = 1 \\ -\infty, & \text{otherwise} \end{cases}, \quad (7)$$

where  $\psi_i^v(t_i)$  and  $\psi_i^c(t_i)$  represent validity check and potentials from cues respectively. Concretely, we use

$$\psi_i^v(t_i) = \mathbb{1}\{I_{t_i} \in \mathcal{S}_i^v\}, \quad (8)$$

$$\psi_i^c(t_i) = \omega_1 \cdot \psi_i^{c_1}(t_i) + \omega_2 \cdot \psi_i^{c_2}(t_i) + \omega_3 \cdot \psi_i^{c_3}(t_i), \quad (9)$$

where  $\mathcal{S}_i^v$  denotes the set of valid frames for  $\text{Tri}_i$  and  $\omega_1, \omega_2, \omega_3$  represent weights for potentials  $\psi_i^{c_1}, \psi_i^{c_2}, \psi_i^{c_3}$ . We discuss each one next:

- **Validity (V).** To assess whether frame  $I_{t_i}$  is valid for  $\text{Tri}_i$ , we check the visibility of  $\text{Tri}_i$  in  $I_{t_i}$ . We approximate this by checking visibility of  $\text{Tri}_i$ ’s three vertices as well as its centroid. Concretely, we transform the vertices and centroid from world coordinates to the normalized device coordinates of the  $t_i$ -th camera. If all vertices and centroid are visible, *i.e.*, their coordinates are in the interval  $[-1, 1]$ , we add frame  $I_{t_i}$  to the set  $\mathcal{S}_i^v$  of valid frames for triangle  $\text{Tri}_i$ .
- **Triangle area (C<sub>1</sub>).** Based on a camera’s pose  $p_{t_i}$ , a triangle’s area changes. The larger the area, the more detailed is the information for  $\text{Tri}_i$  in frame  $I_{t_i}$ .

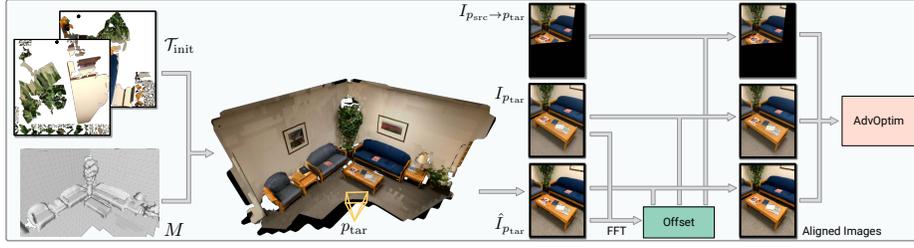


Fig. 6: **Texture smoothing TexSmooth** (Sec. 3.2). We utilize adversarial optimization (AdvOptim) [24] to refine the texture  $\mathcal{T}_{\text{init}}$  from Sec. 3.1. Differently: 1) We initialize with  $\mathcal{T}_{\text{init}}$ . 2) To resolve the issue of misalignment between rendering and ground truth (GT), we integrate an alignment module based on the fast Fourier transform (FFT).

We encourage to assign  $\text{Tri}_i$  to frames  $I_{t_i}$  with large area by defining  $\psi_i^{\mathcal{C}_1}(t_i) = \text{Area}_{t_i}(\text{Tri}_i)$  and set  $\omega_1 > 0$ .

- **Discrepancy between  $z$ -buffer and actual depth ( $\mathcal{C}_2$ ).** For a valid frame  $I_{t_i} \in \mathcal{S}_i^V$ , a triangle’s vertices and its centroid project to valid image coordinates. We compute the discrepancy between: 1) the depth from frame  $I_{t_i}$  at the image coordinates of the vertices and centroid; 2) the depth of vertices and centroid in the camera’s coordinate system. We set  $\psi_i^{\mathcal{C}_2}(t_i)$  to be the sum of absolute value differences between both depth estimates while using  $\omega_2 < 0$ .

- **Perceptual consistency ( $\mathcal{C}_3$ ).** Due to diverse illumination, triangle  $\text{Tri}_i$ ’s appearance changes across frames. Intuitively, we don’t want to assign a texture to  $\text{Tri}_i$  using a frame that contains colors that deviate drastically from other frames. Concretely, we first average all triangle’s three vertices color values across all valid frames, *i.e.*, across all  $I_{t_i} \in \mathcal{S}_i^V$ . We then compare this global average to the local average obtained independently for the three vertices of every valid frame  $I_{t_i} \in \mathcal{S}_i^V$  using an absolute value difference. We require  $\omega_3 < 0$ .

**3) Color Transfer:** Given the inferred triangle-frame assignments  $\mathbf{t}^*$  we complete the texture  $\mathcal{T}_{\text{init}}$  by transferring RGB data from image  $I_{t_i^*}$  for  $\text{Tri}_i, i \in \{1, \dots, |M|\}$ . For this we leverage the camera pose  $p_{t_i^*}$  which permits to transform the texture coordinates  $(u, v)$  of locations within  $\text{Tri}_i$  to corresponding image coordinates  $(a, b)$  in texture  $I_{t_i^*}$  via the mapping  $F: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ , *i.e.*,

$$(a, b) = F(u, v, t_i^*, p_{t_i^*}). \quad (10)$$

Intuitively, given the  $(u, v)$  coordinates on the texture in a coordinate system which is local to the triangle  $\text{Tri}_i$ , and given the camera pose  $p_{t_i^*}$  used to record image  $I_{t_i^*}$ , the mapping  $F$  retrieves the image coordinates  $(a, b)$  corresponding to texture coordinate  $(u, v)$ . Using this mapping, we obtain the texture  $\mathcal{T}_{\text{init}}$  at location  $(u, v)$ , *i.e.*,  $\mathcal{T}_{\text{init}}(u, v)$ , from the image data  $I_{t_i^*}(a, b) \in \mathbb{R}^3$  via

$$\mathcal{T}_{\text{init}}(u, v) = I_{t_i^*}(F(u, v, t_i^*, p_{t_i^*})). \quad (11)$$

Note, because of the overlap detection, we obtain a unique triangle index  $i = G(u, v)$  for every  $(u, v)$  coordinate from Eq. (5). Having transferred RGB data for all coordinates within all triangles results in the texture  $\mathcal{T}_{\text{init}} \in \mathbb{R}^{H \times W \times 3}$ , which

we compare to standard L2 averaging initialization in Fig. 5. We next refine this texture via adversarial optimization. We observe that this initialization  $\mathcal{T}_{\text{init}}$  is crucial to obtain high-quality textures, which we will show in Sec. 4.

### 3.2 Texture Smoothing (TexSmooth)

As can be seen in Fig. 5b, the texture  $\mathcal{T}_{\text{init}}$  contains seams that affect visual quality. To reconstruct a seamless texture  $\mathcal{T}$ , we extend recent adversarial optimization (AdvOptim). Different from prior work [24] which initializes with blank (paper) or averaged (code release<sup>1</sup>) textures, we initialize with  $\mathcal{T}_{\text{init}}$ . Also, we find AdvOptim doesn't handle common camera pose and geometry misalignment well. To resolve this, we develop an efficient alignment module. This is depicted in Fig. 6 and will be detailed next.

**Smoothing with Adversarial Optimization:** To optimize the texture, AdvOptim iterates over camera poses. When optimizing for a specific target camera pose  $p_{\text{tar}}$ , AdvOptim uses three images: 1) the ground truth image  $I_{p_{\text{tar}}}$  of the target camera pose  $p_{\text{tar}}$ ; 2) a rendering  $\hat{I}_{p_{\text{tar}}}$  for the target camera pose  $p_{\text{tar}}$  from the texture map  $\mathcal{T}$ ; and 3) a re-projection from another camera pose  $p_{\text{src}}$ 's ground truth image, which we refer to as  $I_{p_{\text{src}} \rightarrow p_{\text{tar}}}$ . It then optimizes by minimizing an  $\mathcal{L}_1$  plus a conditional adversarial loss. However, we find AdvOptim to struggle with alignment errors due to inaccurate geometry. Therefore, we integrate an efficient alignment operation into AdvOptim. Instead of directly using the input images, we first compute a 2D offset  $(\Delta h_{p_{\text{tar}}}, \Delta w_{p_{\text{tar}}})$  between  $I_{p_{\text{tar}}}$  and  $\hat{I}_{p_{\text{tar}}}$ , which we apply to align  $I_{p_{\text{tar}}}$  and  $\hat{I}_{p_{\text{tar}}}$  as well as  $I_{p_{\text{src}} \rightarrow p_{\text{tar}}}$  via

$$I^{\text{A}} \doteq \text{Align}(I, (\Delta h, \Delta w)), \quad (12)$$

where  $I^{\text{A}}$  marks aligned images. We then use the three aligned images as input:

$$\begin{aligned} \mathcal{L} = & \lambda \|I_{p_{\text{tar}}}^{\text{A}} - \hat{I}_{p_{\text{tar}}}^{\text{A}}\|_1 + \mathbb{E}_{I_{p_{\text{tar}}}^{\text{A}}, I_{p_{\text{src}} \rightarrow p_{\text{tar}}}^{\text{A}}} [\log D(I_{p_{\text{src}} \rightarrow p_{\text{tar}}}^{\text{A}} | I_{p_{\text{tar}}}^{\text{A}})] \\ & + \mathbb{E}_{I_{p_{\text{tar}}}^{\text{A}}, \hat{I}_{p_{\text{tar}}}^{\text{A}}} [\log(1 - D(\hat{I}_{p_{\text{tar}}}^{\text{A}} | I_{p_{\text{tar}}}^{\text{A}}))]. \end{aligned} \quad (13)$$

Here,  $D$  is a convolutional deep-net based discriminator. When using the unaligned image  $I$  instead of  $I^{\text{A}}$ , Eq. (13) reduces to the vanilla version in [24]. We now discuss a fast way to align images.

**Alignment with Fourier Transformation:** To align ground truth  $I_{p_{\text{tar}}}$  and rendering  $\hat{I}_{p_{\text{tar}}}$ , one could use naïve grid-search to find the offset which results in the minimum difference of the shifted images. However, such a grid-search is prohibitively costly during an iterative optimization, especially with high-resolution images (*e.g.*, we use a resolution up to  $1920 \times 1440$ ). Instead, we use the fast Fourier transformation (FFT) to complete the job [6]. Specifically, given a misaligned image pair of  $I \in \mathbb{R}^{h \times w \times 3}$  and  $\hat{I} \in \mathbb{R}^{h \times w \times 3}$ , we compute for every

<sup>1</sup> <https://github.com/hjwdzh/AdversarialTexture>

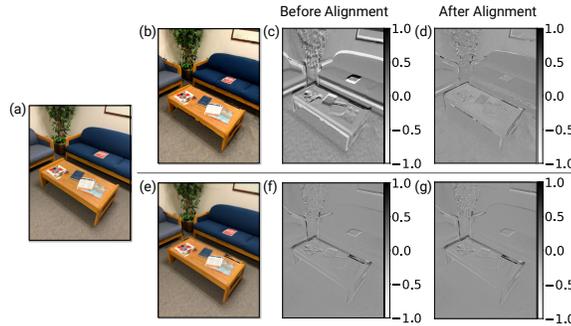


Fig. 7: **Alignment with fast Fourier transformation (FFT)** (Sec. 3.2). We show results for  $\mathcal{T}_{\text{init}}$  (Sec. 3.1) and L2Avg (Fig. 5) in top and bottom rows. (a): ground-truth (GT); (b) and (e): texture rendering with (a)’s corresponding camera; (c) and (d): difference between (a) and (b); (f) and (g): difference between (a) and (e). The top row: FFT successfully aligns GT image and rendering from  $\mathcal{T}_{\text{init}}$ . Within expectation, there is almost no misalignment for the texture L2Avg as it overfits to available views (Fig. 5).

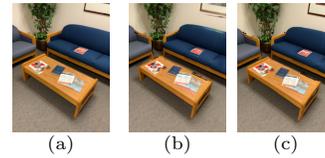


Fig. 8: **Alignment is important for evaluation (Sec. 4.1)**. Clearly, (c) is more desirable than (b). However, before alignment, LPIPS yields 0.3347 and 0.4971 for (a)-(b) and (a)-(c) pairs respectively. This is misleading as lower LPIPS indicates higher quality. After alignment, LPIPS produces 0.3343 and 0.2428 for the same pairs, which provides correct signals for evaluation.

channel the maximum correlation via

$$\operatorname{argmax}_{(i,j)} \operatorname{FFT}^{-1} \left( \operatorname{FFT}(I) \cdot \overline{\operatorname{FFT}(\hat{I})} \right) [i, j]. \quad (14)$$

Here,  $\operatorname{FFT}(\cdot)$  represents the fast Fourier transformation while  $\operatorname{FFT}^{-1}(\cdot)$  denotes its inverse and  $\overline{\operatorname{FFT}(\hat{I})}$  refers to the complex conjugate. After decoding the maximum correlation response and averaging across channels, we obtain the final offset  $(\Delta h, \Delta w)$ . As can be seen in Fig. 7(d), the offset  $(\Delta h, \Delta w)$  is very accurate. Moreover, the computation finishes in around 0.4 seconds even for  $1920 \times 1440$ -resolution images. Note, we don’t need to maintain gradients for  $(\Delta h, \Delta w)$ , since the offset is only used to shift images and not to backpropagate through it.

## 4 Experiments

### 4.1 Experimental Setup

**Data Acquisition.** We use a 2020 iPad Pro and develop an iOS app to acquire the RGBD images  $I_t$ , camera pose  $p_t$ , and scene mesh  $M$  via Apple’s ARKit [1]. **Uoffl Texture Scenes.** We collect a dataset of 11 scenes: four indoor and seven outdoor scenes. This dataset consists of a total of 2807 frames, of which 91, 2052, and 664 are of resolution  $480 \times 360$ ,  $960 \times 720$ , and  $1920 \times 1440$  respectively. For each scene, we use 90% of its views for optimization and the remainder for evaluation. In total, we have 2521 training frames and 286 test frames. This setting is more challenging than prior work where [24] “select(s) 10 views uniformly

distributed from the scanning video” for evaluation while using up to thousands of frames for texture generation. On average, the angular differences between test set view directions and their nearest neighbour in the training sets are  $2.05^\circ$  (min  $0.85^\circ$ /max  $13.8^\circ$ ). Angular distances are computed following [25]. Please see Appendix for scene-level statistics.

**Implementation.** We compare to five baselines for texture generation: L2Avg, ColorMap [54], TexMap [16], MVSTex [48], and AdvTex [24]. For ColorMap, TexMap, and MVSTex, we use their official implementations.<sup>2</sup> For AdvOptim (Sec. 3.2) used in both AdvTex and ours, we re-implement a PyTorch [35] version based on their official release in TensorFlow [2].<sup>1</sup> We evaluate AdvTex with two different initializations: 1) blank textures as stated in the paper (AdvTex-B); 2) the initialization used in the official code release (AdvTex-C). We run AdvOptim using the Adam optimizer [26]. See Appendix for more details. For our TexInit (Sec. 3.1), we use a generic set of weights across all scenes:  $\omega_1 = 1e^{-3}$  (triangle area),  $\omega_2 = -10$  (depth discrepancy), and  $\omega_3 = -1$  (perception consistency), which makes cue magnitudes roughly similar.

On a 3.10GHz Intel Xeon Gold 6254 CPU, ColorMap takes less than two minutes to complete while TexMap’s running time ranges from 40 minutes to 4 hours. MVSTex can be completed in no more than 10 minutes. Our  $\mathcal{T}_{\text{init}}$  (Sec. 3.1) completes in two minutes. Additionally, the AdvOptim takes around 20 minutes for 4000 iterations to complete with an Nvidia RTX A6000 GPU.

**Evaluation metrics.** To assess the efficacy of the method, we study the quality of the texture from two perspectives: perceptual quality and sharpness. 1) For perceptual quality, we assess the similarity between rendered and collected ground-truth views using the Structural Similarity Index Measure (SSIM) [49] and the Learned Perceptual Image Patch Similarity (LPIPS) [53]. 2) For sharpness, we consider measurement  $S_3$  [47] and the norm of image gradient (Grad) following [24]. Specifically, for each pixel, we compute its  $S_3$  value, whose difference between the rendered and ground truth (GT) is used for averaging across the whole image. A similar procedure is applied to Grad. For all four metrics, we report the mean and standard deviation across 11 scenes.

**Alignment in evaluation.** As can be seen in Fig. 8, evaluation will be misleading if we do not align images during evaluation. Therefore, we propose the following procedure: 1) for each method, we align the rendered image and the GT using an FFT (Sec. 3.2); 2) to avoid various resolutions caused by different methods, we crop out the maximum common area across methods. 3) we then compute metrics on those cropped regions. The resulting comparison is fair as all methods are evaluated on the same number of pixels and aligned content.

## 4.2 Experimental Evaluation

**Quantitative evaluation.** Tab. 1 reports aggregated results on all 11 scenes. The quality of our texture  $\mathcal{T}$  (Row 6) outperforms baselines on LPIPS,  $S_3$  and

<sup>2</sup> ColorMap: <https://github.com/intel-is1/Open3D/pull/339>; TexMap: <https://github.com/fdp0525/G2LTex>; MVSTex: <https://github.com/nmoehrle/mvs-texturing>

Table 1: **Aggregated quantitative evaluation on Uoffl Texture Scenes.** We report results in the form of mean $\pm$ std. Please see Fig. 9 for qualitative texture comparisons and Appendix ?? for scene-level quantitative results.

		SSIM $\uparrow$	LPIPS $\downarrow$	$S_3$ $\downarrow$	Grad $\downarrow$
1	L2Avg	<b>0.610</b> $\pm$ 0.191	0.386 $\pm$ 0.116	0.173 $\pm$ 0.105	7.066 $\pm$ 4.575
2	ColorMap	0.553 $\pm$ 0.193	0.581 $\pm$ 0.132	0.234 $\pm$ 0.140	7.969 $\pm$ 5.114
3	TexMap	0.376 $\pm$ 0.113	0.488 $\pm$ 0.097	0.179 $\pm$ 0.062	8.918 $\pm$ 4.174
4	MVSTex	0.476 $\pm$ 0.164	0.335 $\pm$ 0.086	0.139 $\pm$ 0.047	8.198 $\pm$ 3.936
5-1	AdvTex-B	0.495 $\pm$ 0.174	0.369 $\pm$ 0.092	0.148 $\pm$ 0.047	8.229 $\pm$ 4.586
5-2	AdvTex-C	0.563 $\pm$ 0.191	0.365 $\pm$ 0.096	0.135 $\pm$ 0.067	7.171 $\pm$ 4.272
6	Ours	0.602 $\pm$ 0.189	<b>0.309</b> $\pm$ 0.086	<b>0.120</b> $\pm$ 0.058	<b>6.871</b> $\pm$ 4.342

Table 2: **Ablation study.** We report results in the form of mean $\pm$ std.

	Adv Optim	FFT Align	$\mathcal{T}_{\text{init}}$	SSIM $\uparrow$	LPIPS $\downarrow$	$S_3$ $\downarrow$	Grad $\downarrow$
1			$\checkmark$	0.510 $\pm$ 0.175	0.342 $\pm$ 0.060	0.141 $\pm$ 0.052	8.092 $\pm$ 4.488
2	$\checkmark$	$\checkmark$		0.592 $\pm$ 0.192	0.332 $\pm$ 0.102	0.130 $\pm$ 0.066	<b>6.864</b> $\pm$ 4.211
3	$\checkmark$		$\checkmark$	0.559 $\pm$ 0.196	0.346 $\pm$ 0.082	0.125 $\pm$ 0.057	7.244 $\pm$ 4.359
4	$\checkmark$	$\checkmark$	$\checkmark$	<b>0.602</b> $\pm$ 0.189	<b>0.309</b> $\pm$ 0.086	<b>0.120</b> $\pm$ 0.058	6.871 $\pm$ 4.342

Grad, confirming the effectiveness of the proposed pipeline. Specifically, we improve LPIPS by 7.8% from 0.335 (2<sup>nd</sup>-best) to 0.309, indicating high perceptual similarity. Moreover,  $\mathcal{T}$  maintains sharpness as we improve  $S_3$  by 11.1% from 0.135 (2<sup>nd</sup>-best) to 0.120 and Grad from 7.171 (2<sup>nd</sup>-best) to 6.871. Regarding SSIM, we find it to favor L2Avg in almost all scenes (see Appendix) which aligns with the findings in [53].

**Ablation study.** We verify the design choices of TexInit and TexSmooth in Tab. 2.

**1) TexSmooth is required:** we directly evaluate  $\mathcal{T}_{\text{init}}$  and 1<sup>st</sup> vs. 4<sup>th</sup> row confirms the performance drop: -0.092 (SSIM), +0.033 (LPIPS), +0.021 ( $S_3$ ), and +1.221 (Grad). **2)  $\mathcal{T}_{\text{init}}$  is needed:** we replace  $\mathcal{T}_{\text{init}}$  with L2Avg as it performs better than ColorMap and TexMap in Tab. 1 and still incorporate FFT into AdvOptim. We observe inferior performance: -0.010 (SSIM), +0.023 (LPIPS), +0.010 ( $S_3$ ) in 2<sup>nd</sup> vs. 4<sup>th</sup> row. **3) Alignment is important:** we use the vanilla AdvOptim but initialize with  $\mathcal{T}_{\text{init}}$ . As shown in Tab. 2’s 3<sup>rd</sup> vs. 4<sup>th</sup> row, the texture quality drops by -0.043 (SSIM), +0.037 (LPIPS), +0.005 ( $S_3$ ), and +0.373 (Grad).

**Qualitative evaluation.** We present qualitative examples in Fig. 9. Fig. 9a and Fig. 9b demonstrate that L2Avg and ColorMap produce overly smooth texture. Meanwhile, due to noise in the geometry, *e.g.*, Fig. 2, TexMap fails to resolve texture seams and cannot produce a complete texture (Fig. 9c). MVSTex results in Fig. 9d are undesirable as geometries are removed. This is because MVSTex requires ray collision checking to remove occluded faces. Due to the misalignment between geometries and cameras, artifacts are introduced. We show results of AdvTex-C in Fig. 9e as it outperforms AdvTex-B from Tab. 1. Ar-

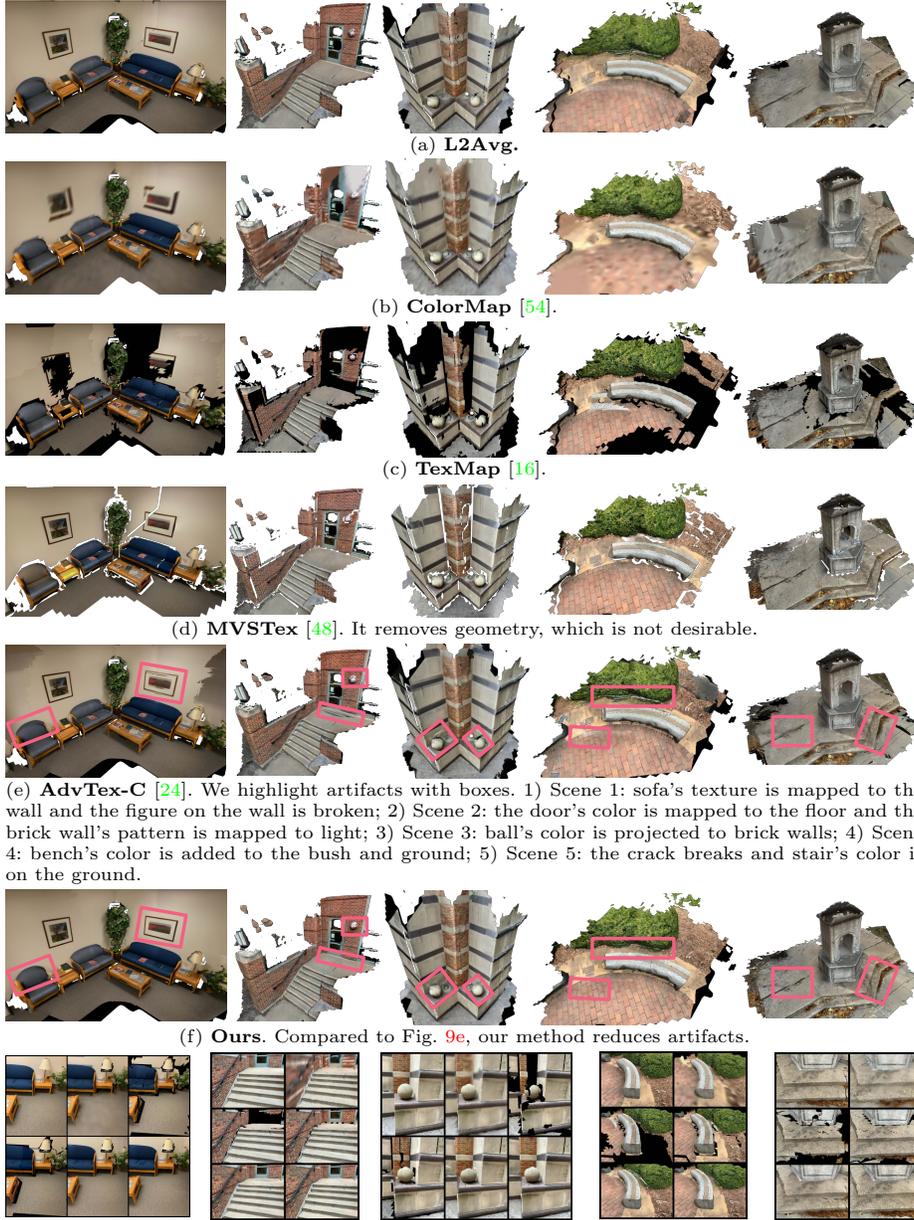


Fig. 9: **Qualitative results on UofI Texture Scenes.** For each method, we show results for Scene 1 to 5 from left to right. Best viewed in color and zoomed-in.

Table 3: **ScanNet results.** We report results in the form of mean $\pm$ std. Note, this can't be directly compared to [23]'s Tab. 2: while we reserve 10% views for evaluation, [23] reserves only 10/2011 ( $\approx 0.5\%$ ), where 2011 is the number of average views per scene.

		SSIM $\uparrow$	LPIPS $\downarrow$	$S_3$ $\downarrow$	Grad $\downarrow$
1-1	AdvTex-B	0.534 $\pm$ 0.074	0.557 $\pm$ 0.071	0.143 $\pm$ 0.028	3.753 $\pm$ 0.730
1-2	AdvTex-C	0.531 $\pm$ 0.074	0.558 $\pm$ 0.075	0.161 $\pm$ 0.044	4.565 $\pm$ 1.399
2	Ours	<b>0.571</b> $\pm$ 0.069	<b>0.503</b> $\pm$ 0.090	<b>0.127</b> $\pm$ 0.031	<b>3.324</b> $\pm$ 0.826



Fig. 10: **Remaining six scenes with our textures.** See Appendix for all methods' results on these scenes.



Fig. 11: **Results on ScanNet.** Left to right: AdvTex-B/C and ours. Ours alleviates artifacts: colors from box on the cabinet are mapped to the backpack and wall.

tifacts are highlighted. Our method can largely mitigate such seams, which can be inferred from Fig. 9f. In Fig. 9g, we show renderings, which demonstrate the effectiveness of the proposed method. Please see Appendix for complete qualitative results of scenes in Fig. 10.

**On ScanNet [11].** Following [24], we study scenes with ID  $\leq 20$  ( Fig. 11, Tab. 3). We improve upon baselines (AdvTex-B/C) by a margin on SSIM (0.534  $\rightarrow$  0.571), LPIPS (0.557  $\rightarrow$  0.503),  $S_3$  (0.143  $\rightarrow$  0.127), and Grad (3.753  $\rightarrow$  3.324).

## 5 Conclusion

We develop an initialization and an alignment method for fully-automatic texture generation from a given scene mesh, and a given sequence of RGBD images and their camera parameters. We observe the proposed method to yield appealing results, addressing robustness issues due to noisy geometry and misalignment of prior work. Quantitatively we observe improvements on both perceptual similarity (LPIPS from 0.335 to 0.309) and sharpness ( $S_3$  from 0.135 to 0.120).

**Acknowledgements:** Supported in part by NSF grants 1718221, 2008387, 2045586, 2106825, MRI #1725729, and NIFA award 2020-67021-32799.

## References

1. Augmented Reality - Apple Developer. <https://developer.apple.com/augmented-reality/> (2021), accessed: 2021-11-14 **2, 10**
2. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P.A., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zhang, X.: Tensorflow: A system for large-scale machine learning. In: OSDI (2016) **11**
3. Abdelhafiz, A., Mostafa, Y.G.: Automatic texture mapping mega-projects. *Journal of Spatial Science* (2020) **4**
4. Aganj, E., Monasse, P., Keriven, R.: Multi-view Texturing of Imprecise Mesh. In: ACCV (2009) **3**
5. Aliev, K.A., Ulyanov, D., Lempitsky, V.S.: Neural Point-Based Graphics. In: ECCV (2020) **3**
6. Anuta, P.E.: Spatial Registration of Multispectral and Multitemporal Digital Imagery Using FastFourierTransform Techniques. *Trans. Geoscience Electronics* (1970) **2, 9**
7. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: PatchMatch: a randomized correspondence algorithm for structural image editing. In: SIGGRAPH (2009) **3**
8. Baumberg, A.: Blending Images for Texturing 3D Models. *BMVC* (2002) **3**
9. Bernardini, F., Martin, I., Rushmeier, H.: High Quality Texture Reconstruction from Multiple Scans. *TVCG* (2001) **3**
10. Bi, S., Kalantari, N.K., Ramamoorthi, R.: Patch-based optimization for image-based texture mapping. *TOG* (2017) **3**
11. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T.A., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. *CVPR* (2017) **3, 14**
12. Debevec, P., Taylor, C., Malik, J.: Modeling and rendering architecture from photographs: A hybrid geometry and image-based approach. *SIGGRAPH* (1996) **1, 3**
13. Duan, Y.: Topology Adaptive Deformable Models for Visual Computing. Ph.D. thesis, State University of New York (2003) **4**
14. El-Hakim, S., Gonzo, L., Picard, M., Girardi, S., Simoni, A.: Visualization of Frescoed surfaces: Buonconsiglio Castle – Aquila Tower, “Cycle of the Months”. *IAPRS* (2003) **1, 4**
15. Früh, C., Sammon, R., Zakhor, A.: Automated texture mapping of 3D city models with oblique aerial imagery. *3DPVT* (2004) **4**
16. Fu, Y., Yan, Q., Yang, L., Liao, J., Xiao, C.: Texture Mapping for 3D Reconstruction with RGB-D Sensor. *CVPR* (2018) **2, 4, 11, 13**
17. Globerson, A., Jaakkola, T.: Fixing Max-Product: Convergent Message Passing Algorithms for MAP LP-Relaxations. In: *NIPS* (2007) **6**
18. Goel, S., Kanazawa, A., Malik, J.: Shape and Viewpoint without Keypoints. In: *ECCV* (2020) **3**
19. Grammatikopoulos, L., Kalisperakis, I., Karras, G., Petsa, E.: Automatic multi-view texture mapping of 3D surface projections. *International Workshop 3D-ARCH* (2007) **1, 4**
20. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. *CVPR* (2018) **3**

21. Henzler, P., Mitra, N.J., Ritschel, T.: Learning a Neural 3D Texture Space From 2D Exemplars. CVPR (2020) [3](#)
22. Hernández-Esteban, C.: Stereo and Silhouette Fusion for 3D Object Modeling from Uncalibrated Images Under Circular Motion. Ph.D. thesis, École Nationale Supérieure des Télécommunications (2004) [4](#)
23. Huang, J., Dai, A., Guibas, L., Nießner, M.: 3Dlite: Towards Commodity 3D Scanning for Content Creation. ACM TOG (2017) [3](#)
24. Huang, J., Thies, J., Dai, A., Kundu, A., Jiang, C.M., Guibas, L., Nießner, M., Funkhouser, T.: Adversarial Texture Optimization From RGB-D Scans. CVPR (2020) [2, 3, 8, 9, 10, 11, 13, 14](#)
25. Huynh, D.: Metrics for 3D Rotations: Comparison and Analysis. Journal of Mathematical Imaging and Vision (2009) [11](#)
26. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. ArXiv (2015) [11](#)
27. Lempitsky, V., Ivanov, D.: Seamless mosaicing of image-based texture maps. CVPR (2007) [1, 4](#)
28. Lensch, H., Heidrich, W., Seidel, H.P.: Automated texture registration and stitching for real world models. Graphical Models (2001) [4](#)
29. Marshner, S.R.: Inverse rendering for computer graphics. Ph.D. thesis, Cornell University (1998) [1, 3](#)
30. Neugebauer, P.J., Klein, K.: Texturing 3d models of real world objects from multiple unregistered photographic views. Eurographics (1999) [1, 3](#)
31. Niem, W., Wingbermühle, J.: Automatic reconstruction of 3D objects using a mobile camera. IVC (1999) [3](#)
32. Oechsle, M., Mescheder, L.M., Niemeyer, M., Strauss, T., Geiger, A.: Texture Fields: Learning Texture Representations in Function Space. ICCV (2019) [3](#)
33. Ofek, E., Shilat, E., Rappoport, A., Werman, M.: Multiresolution Textures from Image Sequences. Computer Graphics & Applications (1997) [3](#)
34. Pan, R., Taubin, G.: Color adjustment in image-based texture maps. Graphical Models (2015) [4](#)
35. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. ArXiv [abs/1912.01703](#) (2019) [11](#)
36. Pighin, F., Hecker, J., Lischinski, D., Szeliski, R., Salesin, D.H.: Synthesizing realistic facial expressions from photographs. CGIT (1998) [3](#)
37. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3D Deep Learning with PyTorch3D. arXiv:2007.08501 (2020) [7](#)
38. Rocchini, C., Cignoni, P., Montani, C., Scopigno, R.: Multiple textures stitching and blending on 3D objects. Eurographics Workshop on Rendering (1999) [4](#)
39. Rocchini, C., Cignoni, P., Montani, C., Scopigno, R.: Acquiring, stitching and blending diffuse appearance attributes on 3d models. The Visual Computer (2002) [4](#)
40. Saito, S., Wei, L., Hu, L., Nagano, K., Li, H.: Photorealistic Facial Texture Inference Using Deep Neural Networks. CVPR (2017) [3](#)
41. Sawhney, R., Crane, K.: Boundary first flattening. ACM TOG (2018) [2, 5](#)
42. Shu, J., Liu, Y., Li, J., Xu, Z., Du, S.: Rich and seamless texture mapping to 3D mesh models. Advances in Image and Graphics Technologies (2016) [4](#)
43. Sinha, S.N., Steedly, D., Szeliski, R., Agrawala, M., Pollefeys, M.: Interactive 3D architectural modeling from unordered photo collections. In: SIGGRAPH 2008 (2008) [1](#)

44. Sitzmann, V., Thies, J., Heide, F., Nießner, M., Wetzstein, G., Zollhöfer, M.: Deep-Voxels: Learning Persistent 3D Feature Embeddings. CVPR (2019) [3](#)
45. Thierry, M., David, F., Gorria, P., Salvi, J.: Automatic texture mapping on real 3D model. CVPR (2007) [1](#)
46. Thies, J., Zollhöfer, M., Nießner, M.: Deferred Neural Rendering. ACM TOG (2019) [3](#)
47. Vu, C.T., Phan, T.D., Chandler, D.M.:  $S_3$ : A spectral and spatial measure of local perceived sharpness in natural images. IEEE Transactions on Image Processing (2012) [3](#), [11](#)
48. Waechter, M., Moehrle, N., Goesele, M.: Let There Be Color! Large-Scale Texturing of 3D Reconstructions. In: ECCV (2014) [4](#), [11](#), [13](#)
49. Wang, Z., Bovik, A., Sheikh, H.R., Simoncelli, E.P.: Image Quality Assessment: from Error Visibility to Structural Similarity. IEEE Transactions on Image Processing (2004) [11](#)
50. Wexler, Y., Shechtman, E., Irani, M.: Space-time video completion. In: CVPR (2004) [3](#)
51. Wuhler, S., Atanassov, R., Shu, C.: Fully Automatic Texture Mapping for Image-Based Modeling. Tech. rep., Institute for Information Technology (2006) [4](#)
52. Xiang, F., Xu, Z., Havsan, M., Hold-Geoffroy, Y., Sunkavalli, K., Su, H.: NeuTex: Neural Texture Mapping for Volumetric Neural Rendering. In: CVPR (2021) [3](#)
53. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. CVPR (2018) [3](#), [11](#), [12](#)
54. Zhou, Q.Y., Koltun, V.: Color Map Optimization for 3D Reconstruction with Consumer Depth Cameras. ACM TOG (2014) [2](#), [3](#), [11](#), [13](#)