

Style Agnostic 3D Reconstruction via Adversarial Style Transfer

Felix Petersen Bastian Goldluecke Oliver Deussen
University of Konstanz
{first.last}@uni-konstanz.de

Hilde Kuehne
University of Frankfurt
IBM-MIT Watson AI Lab

Abstract

Reconstructing the 3D geometry of an object from an image is a major challenge in computer vision. Recently introduced differentiable renderers can be leveraged to learn the 3D geometry of objects from 2D images, but those approaches require additional supervision to enable the renderer to produce an output that can be compared to the input image. This can be scene information or constraints such as object silhouettes, uniform backgrounds, material, texture, and lighting. In this paper, we propose an approach that enables a differentiable rendering-based learning of 3D objects from images with backgrounds without the need for silhouette supervision. Instead of trying to render an image close to the input, we propose an adversarial style-transfer and domain adaptation pipeline that allows to translate the input image domain to the rendered image domain. This allows us to directly compare between a translated image and the differentiable rendering of a 3D object reconstruction in order to train the 3D object reconstruction network. We show that the approach learns 3D geometry from images with backgrounds and provides a better performance than constrained methods for single-view 3D object reconstruction on this task.

1. Introduction

Inferring the 3D geometry of an object from an arbitrary single image is an interesting and challenging task that is easy for humans but still poses a hard computational problem. This problem is challenging because, by projecting a 3D shape onto a 2D image, depth information is discarded, and occluded regions are inherently difficult to reconstruct.

Perhaps the most interesting and promising approaches addressing this problem are those utilizing learning-based methods. These can potentially exploit previously seen examples to overcome missing information such as partial visibility, unknown lighting conditions, etc. Some of these methods rely on 3D shape supervision such as Mem3D [1], DeformNet [2], Pixel2Mesh [3], and 3D-R2N2 [4], or key-point correspondences between images [5].

Other, less supervised methods predict the geometry using a geometry reconstruction network and projecting it back onto an image space using a differentiable renderer. Such 3D representations can be based on meshes as in case of SoftRas [6], DIB-R [7], and NMR [8], voxels [9], or distance fields as used in SDFDiff [10]. The idea here is to predict a 3D model from a single image, render the silhouette of this model using a differentiable renderer, and impose a loss function that ensures that the rendered silhouette again resembles the silhouette input image. This strategy allows learning the respective shape models without the need for annotation of 3D shapes or key-point correspondences between images. An interesting problem arising in this context is that renderings need to be compared to input images for the loss function. Thus, the renderer has to produce images that look stylistically similar to the input images to allow for a meaningful comparison.

This can be challenging for differentiable renderers due to the complexity of images: If the renderer does not have enough degrees of freedom, it cannot produce a realistic, matching image. However, if it has too many degrees of freedom, trivial solutions such as a background filling the entire image with the texture of the input image become possible, and training becomes unstable. Therefore, current approaches require additional supervision and constraints, such as uniform background, known lighting conditions, materials, textures, and supervision of silhouettes [6]–[10].

To alleviate those constraints and to deal with the challenge of complex image formation, we propose a framework that combines differentiable rendering based 3D reconstruction with adversarial domain adaptation to bridge the gap between rendered and input images as shown in Fig. 1. Starting from a single input image, we predict the 3D shape of the object with a CNN that converts the image input to a 3D mesh, and apply a differentiable renderer [6] to produce a 2D rendering of the predicted mesh. To allow for comparison between the rendered representation and the input image, we propose to use an Adversarial Style-Transfer Autoencoder (ASTA) pipeline, which converts the input image to a representation that could have been generated by the renderer. The encoder and decoder are locally restricted

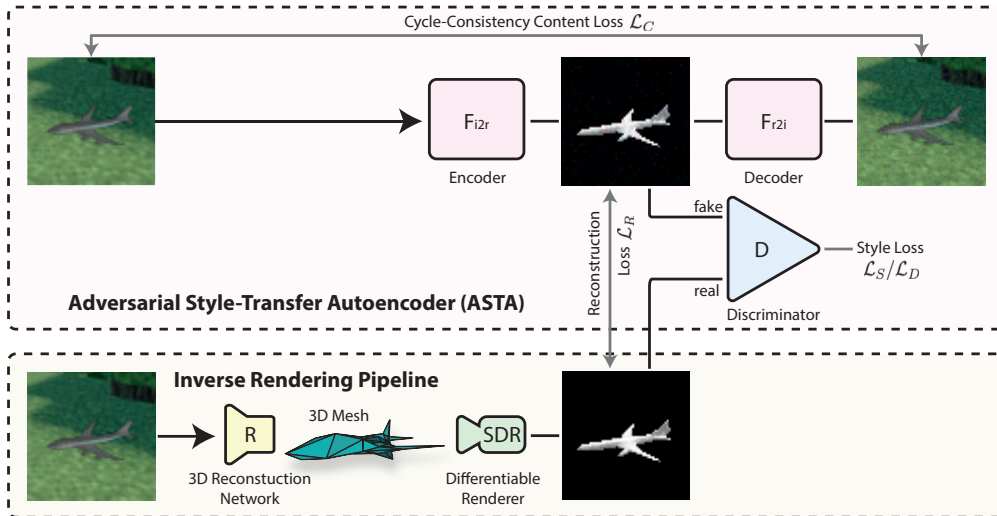


Figure 1: Overview of the framework. We extend the differentiable renderer-based 3D geometry reconstruction by an image-to-image domain adaptation network (F_{i2r}) to allow for training with arbitrary images. F_{i2r} translates input images into images that could have been rendered. We train F_{i2r} using an adversarial style loss and a cycle-consistency content loss.

U-Net architectures such that the “latent space” is a grey-scale image. Thus, the encoder first translates the input image into a representation that resembles the rendered images and the decoder translates this back into its original representation. Following the ideas of neural style transfer and adversarial domain adaptation, we use an adversarial discriminator to enforce the “latent” space to be stylistically similar to images produced by the differentiable renderer. As the autoencoder needs to reconstruct the original image, it preserves the content. With this, we can transfer the input image to the differentiable renderer’s output space, rather than forcing the renderer to reproduce the input image.

As only input images and no smoothly rendered images are initially given, we cannot directly train the style-transfer. We resolve this by adopting a co-evolutionary strategy to train the networks in an interleaved way, such that the renderings of previous reconstruction guesses are used.

We first evaluate the proposed framework on the ShapeNet data set [11]. To demonstrate the proposed framework’s abilities, we render objects of ShapeNet with and without random background scenes. We show that, for both cases, the proposed framework learns the 3D object shape. For images with backgrounds, we outperform other methods even if they are provided additional silhouette supervision, e.g., by weakly-supervised segmentation. On the ShapeNet data set with textured backgrounds, we achieve a mean intersection-over-union (IoU) of 0.3816, whereas the baseline with additional silhouette supervision achieves a mean IoU of 0.2563. Second, we apply the model trained on ShapeNet to photos of LEGO models of ShapeNet object classes, demonstrating that the method can reproduce

the performance on the synthetic images also in this case.

Contributions. We present a framework for unsupervised training of 3D reconstruction from a single image with random background scenes. We propose using an adversarial style-transfer autoencoder as image-to-image domain translator, which converts input images into images that could have been rendered by a smooth differentiable renderer. We train this domain translator to preserve the content while changing the style and—by that—bridge the gap between arbitrary images and the constraints of differentiable renderers. As the rendering domain is initially unknown, we leverage a co-evolutionary training of domain adaptation. In particular, we can relax assumptions on and do not require supervision of materials, colors, lighting, background, and silhouette—they do not need to be controlled nor constant across the data set.

2. Related Work

Since our goal is unsupervised learning of single-view 3D reconstruction, we first focus on related work based on differentiable rendering. In addition, we present related work on style-transfer, adversarial domain adaptation, and image-to-image translation.

3D Object Reconstruction via Differentiable Renderers.

Over the past years, many differentiable renderers for 3D meshes have been presented [6], [8], [12]–[21]. Those renderers are differentiable with respect to lighting, geometry [6]–[8], [10], [12]–[15], [19], [20], material [7], [16]–[18], [20], or texture [7], [16], [18], [20]. Using a differentiable renderer, a mesh optimization process can be performed

to reconstruct a single 3D object [7], [8], [12], [13], [15], [20]. For that, an initial mesh can be optimized iteratively by adapting its geometry guided by back-propagated gradients. With state-of-the-art deep learning techniques and 3D supervision, high-quality 3D object reconstruction can be done via direct prediction [2]–[4], [22]–[24]. This approach is orders of magnitude faster than optimization and allows transferring knowledge to unseen examples and occluded areas. For 3D unsupervised 3D object prediction, state-of-the-art are encoder-renderer architectures where the encoder is trained for 3D mesh prediction and the differentiable renderer plays the role of a decoder. Thus, the encoder-renderer network predicts the 3D model in the latent space [6]–[8], [14], [19]. Zhang *et al.* [25] extend this idea by integrating GANs to generate alternative views for each input image, while also relying on silhouette supervision. These encoder-render methods are restricted to reconstructing from images that look like images produced by the integrated differentiable renderer or require additional silhouette supervision, which limits their application.

The first to raise the idea of learning from 2D images without silhouette supervision were Henderson *et al.* [26]. They performed textured 3D mesh generation from images without silhouette annotation via a 2D VAE based framework. The difference is that they focus on mesh generation, while our focus lies on mesh reconstruction. To this end, they focus on making the rendering equal to the input image, while we focus on translating the input image into a rendered image.

Style-Transfer and Adversarial Domain Adaptation. Neural Style-Transfer (NST) has been first proposed by Gatys *et al.* [27], where the activations of a pretrained CNN are used for the style loss, allowing to optimize the style-transfer on a single image. Using a network to predict the style-transferred image, Isola *et al.* [28] propose the image-to-image translation architecture Pix2Pix, training it through a Generative Adversarial Network (GAN). This allows for fast inference as it shifts the process from optimization to prediction. Zhu *et al.* [29] extend this idea by introducing two GANs, translating an image as well as back-translating it, and thus including a cycle-consistency loss (CycleGAN). Based on this idea, adversarial approaches to unsupervised domain adaptation have been presented [30]–[32]. Hoffman *et al.* [30] propose extending CycleGAN by a semantic loss and a feature loss for adversarial domain adaptation. We propose a variation of these ideas, where an image-to-image domain adaptation network translates an input image to an image from the domain of the encoder-rendering pipeline’s outputs and preserves the content by a cycle-consistency through autoencoding.

3. Style-Agnostic Unsupervised Reconstruction

To perform an unsupervised style-agnostic reconstruction, we need to bridge the gap between input images and the output of the renderer. Our architecture is divided into two main parts as shown in Fig. 1. The image-to-image translation pipeline (top) translates input images into images which resemble those generated by the differentiable renderer. The reconstruction-rendering pipeline (bottom) consists of the 3D reconstruction network R and a smooth differentiable renderer SDR , forming together the encoder-renderer network. The reconstruction network translates an image input into a 3D geometry representation, which is then fed into the differentiable renderer to produce the respective rendered representation of the 3D geometry. Both parts are trained in an alternating fashion. In the remainder of this section, we discuss both pipelines in detail, introduce the necessary loss functions, and describe the overall training strategy to tie the components together.

3.1. Loss Functions for the Domain Adaptation

The network for image-to-image translation aims to translate images x drawn from the input image domain p into images that look like they have been rendered by the differentiable renderer SDR .

Training of this image-to-rendering network F_{i2r} , is controlled by two loss functions, an adversarial style loss arising from a discriminator between renderings and pseudo renderings, and a cycle-consistency loss based on the inverse mapping of the pseudo-rendered image back to the original image. Each sub-network is described in detail in the following.

Adversarial Style Loss. To convert input images into their rendered representations, we need to transfer the style from rendered images to the input images. To this end, we use the adversarial loss of a discriminator D , which discriminates between images that have actually been rendered and images created by $F_{i2r}(x)$, where x is an input image. Different from typical adversarial losses as, e.g., used in pix2pix [28] or CycleGAN [29], we can not sample from the domain of smoothly rendered images. Instead, we use those images generated by the reconstruction network and the differentiable renderer as examples of rendered objects, as we do not know the geometry of the 3D objects in advance. Note that this requires a co-evolutionary training of the reconstruction network R and the adversarial style transfer, as detailed in 3.3.

Our loss function \mathcal{L}_D for the discriminator D and style loss \mathcal{L}_S for the image-to-rendering encoder F_{i2r} are:

$$\mathcal{L}_D = \mathbb{E}_{x \sim p} \left[\log D(F_{i2r}(x)) + \log(1 - D(SDR(R(x)))) \right] \quad (1)$$

$$\mathcal{L}_S = \mathbb{E}_{x \sim p} \left[\log(1 - D(F_{i2r}(x))) \right] \quad (2)$$

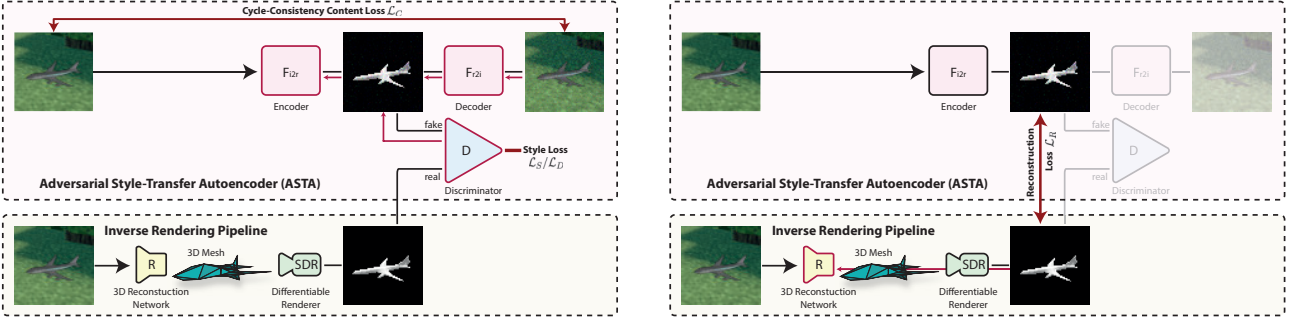


Figure 2: Training procedure of the proposed framework. Left: in the first stage, we train the domain adaptation with the cycle-consistency content loss as well as the adversarial style loss. Right: in the second stage, we train the reconstruction network with the geometry reconstruction loss. The red arrows indicate back-propagation of gradients.

To ensure that, for each domain adaptation, the predicted rendering does not differ too much from the current rendered prediction, we include the loss term

$$\mathcal{L}_J = \mathbb{E}_{x \sim p} \left[\max(0, \delta - J(\text{SDR}(R(x)), F_{i2r}(x))) \right] \quad (3)$$

to encourage that the Jaccard similarity coefficient J is larger than $\delta = 0.25$. We relax the Jaccard similarity to $J(x, y) = \frac{xy}{x+y-xy}$ as per probabilistic real-valued logic.

Cycle-Consistency Content Loss. We train the autoencoder with a cycle-consistency content loss to preserve the content during the image-to-rendering transfer. For that, we use the rendering-to-image decoder termed F_{r2i} , which is trained to compute the inverse map of F_{i2r} . We find that, as the autoencoder needs to preserve information, the easiest path (and thus the path chosen) is to keep the contents because it needs to reconstruct the input again from the stylized image. As experiments by Chu *et al.* [33] indicate, a cycle-consistency can be fulfilled even while the inner image $F_{i2r}(x)$ is almost constant because most information can be hidden in high-frequency detail. To reinforce that the content is preserved in the “latent” / smoothly rendered space, we apply random noise to the latent space before feeding it into the decoder to prevent it from storing all information in high-frequency details. Thus, we add the noise term ϵ to the input of the rendering-to-image decoder $F_{i2r}(x)$. We define the cycle-consistency loss, which is applied to both, the image-to-rendering and the rendering-to-image translation networks F_{i2r} and F_{r2i} , as

$$\mathcal{L}_C = \mathbb{E}_{x \sim p} \left[\|x - F_{r2i}(F_{i2r}(x) + \epsilon)\|_h \right], \quad (4)$$

where $\|\cdot\|_h$ denotes the (smooth) Huber norm and $\epsilon \sim \mathcal{N}(0, \sigma_{\text{noise}}^2)$ where $\sigma_{\text{noise}} = 0.15$.

In total, the style, content, and Jaccard losses for F_{i2r} and F_{r2i} accumulate to

$$\begin{aligned} \mathcal{L}_{F_{i2r}} &= \mathcal{L}_C + \mathcal{L}_S + \mathcal{L}_J \\ \mathcal{L}_{F_{r2i}} &= \mathcal{L}_C. \end{aligned} \quad (5)$$

3.2. Loss Function for the Geometry Reconstruction

To train the reconstruction network, we consider the difference between the rendered predictions for the reconstruction and the stylized input images, as shown in Fig. 2. Following Liu *et al.* [6] and Kato *et al.* [8], we sample two input images of the same object, using one of them as input to the reconstruction network, whereas the second view is only used to validate the result from an alternative view point. The reconstruction loss is therefore defined as:

$$\mathcal{L}_R = \mathbb{E}_{x, y \sim p} \left[\text{SDR}_x(R(x)) - F_{i2r}(x)\|_h + \text{SDR}_y(R(x)) - F_{i2r}(y)\|_h \right] \quad (6)$$

Here $x, y \sim p$ are pairs of images of the same object, SDR_x renders the object from the view point of x , and SDR_y renders the object from the alternative view point of y . To guide the prediction towards non-degenerate and reasonably shaped meshes, we also apply regularization losses as done by Liu *et al.* [6].

3.3. Overall Training Strategy

The challenge in training is that, to train reconstruction, we need to compare the smoothly rendered images of the predicted shape to the input, which requires domain adaptation. But, to train the domain adaptation, we need to know what a smoothly rendered image looks like. We resolve this problem by adopting a co-evolutionary strategy and by training the networks in an interleaved way. For each training cycle, we start with training the domain adaptation networks using the respective losses (see Eq. 2, Eq. 3 and Eq. 4) for a fixed number of iterations, and after that train the reconstruction network R using the reconstruction loss (see Eq. 6). Note that we start with a setting where the reconstruction is a constant sphere, a basic domain adaptation is learned, which then initializes the reconstruction network. For each new training cycle, we keep the weights of the reconstruction network, but re-initialize the domain




Figure 3: Examples of the domain adaptation in the tenth training cycle. The top row displays the input image x , the second the RGB channels of $F_{i2r}(x)$, the third the alpha channel of $F_{i2r}(x)$, and the bottom row displays $F_{r2i}(F_{i2r}(x))$.

adaptation and continue the training as described above. This re-initialization allows the domain adaptation to be trained with the current optimal guesses for the domain of rendered objects, and to recover from previous errors induced by erroneous reconstructions. In the next training cycle, training of the domain adaptation improves since the rendered reconstructions are more representative samples of the rendered image domain. Examples for the domain adaptation in the tenth training cycle for various objects are shown in Fig. 3.

4. Experiments¹

4.1. Data Sets and Evaluation Metric

ShapeNet. We evaluate the proposed approach on the ShapeNet [11] data set. ShapeNet is a 3D object data set that allows us to render sample images with uniform as well as textured background and to quantitatively compare the resulting 3D reconstruction to the 3D ground truth shape. For the comparison to state-of-the-art as well as the ablation study, we use the 13 standard benchmark classes from ShapeNet and which have been rendered by Kato *et al.* [8] at a resolution of 64×64 from 24 azimuths at an altitude of 30° using Blender and are publicly available.

To show that our method is style-agnostic, we add uniform as well as textured backgrounds to the input and training images of the data set. To this end, we render backgrounds of 135 Minecraft scenes with each 3 lighting settings and from 24 azimuths with Blender and combine them with the rendered object views. Samples of these images with backgrounds can be found in Fig. 3, 4, and 5. Note that Fig. 3 and 5 are from the training sets and Fig. 4 is from the test set.

We follow the evaluation metric proposed by Kato *et al.* [8] and used by Liu *et al.* [6] and Chen *et al.* [7] and use respective occupancy voxel grids of resolution $32 \times 32 \times 32$ and 3D intersection-over-union (IoU). We report all results based on mean 3D IoU averaged over 3 runs for 13 object classes as well as the mean over all classes.

¹Our implementation as well as our new data sets are openly available at github.com/Felix-Petersen/style-agnostic-3d-reconstruction.

ShapeNet2Lego. To evaluate the model on camera-captured photos, we built seven physical LEGO models and their CAD 3D model from ShapeNet object classes. We captured 12 photos from different angles as shown in Fig. 6 with backgrounds consisting of LEGO baseplates and walls. We use this as a second test set and apply the models trained on the synthetic ShapeNet images on those photos as well.

4.2. Implementation details

To allow comparability with previous works, we build on the source code by Liu *et al.* [6]. This setting also follows the setup by Kato *et al.* [8]. We use the same reconstruction network topology and train with a batch size of 64, a learning rate of 10^{-4} , and a uniform sphere with 642 vertices as base model. We further use—wherever applicable—the same hyperparameters, metrics, data set, and seeds to produce comparable results. During training of the reconstruction network, we follow the same strategy, i.e., we sample two input images of the same object scene, using one of them as input to the reconstruction network, whereas the second view is only used to validate the result. This is possible as the camera position is given, such that the predicted object geometry can be rendered from both perspectives. Note that we demonstrate in Section 4.5 that, even under strong perturbation of the given camera position by Gaussian noise with a standard deviation of up to 5° , the reconstruction quality does not drop significantly.

The domain adaptation is re-trained at the beginning of each training cycle. We weight the adversarial game in the style loss with a factor of 1 and the cycle-consistency loss of the transfer with a factor of 400. All reported accuracies are averaged over three runs.

4.3. Comparison to State-of-the-Art

First, we consider how our method improves existing differentiable rendering pipelines under weaker supervision. To this end, we present results of current methods with silhouette supervision and use the best performing differentiable renderer, i.e., SoftRas [6], for our method. To show the impact of our method, we consider two scenarios. In the first case we render all objects with uniform backgrounds

Table 1: Mean 3D IoU for 13 object classes from ShapeNet [11] averaged over 3 runs. We report baselines for silhouette supervision, uniform backgrounds, and textured background. In both cases (of the latter), domain adaptation significantly improves performance, and for uniform backgrounds our method outperforms older methods that require silhouettes.

| Method | Airplane | Bench | Dresser | Car | Chair | Display | Lamp | Speaker | Rifle | Sofa | Table | Phone | Vessel | Mean |
|--|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| With silhouette supervision | | | | | | | | | | | | | | |
| Yan <i>et al.</i> [9] (retrieval) | 0.5564 | 0.4875 | 0.5713 | 0.6519 | 0.3512 | 0.3958 | 0.2905 | 0.4600 | 0.5133 | 0.5314 | 0.3097 | 0.6696 | 0.4078 | 0.4766 |
| Yan <i>et al.</i> [9] (voxel) | 0.5556 | 0.4924 | 0.6823 | 0.7123 | 0.4494 | 0.5395 | 0.4223 | 0.5868 | 0.5987 | 0.6221 | 0.4938 | 0.7504 | 0.5507 | 0.5736 |
| Chen <i>et al.</i> [7] (DIB-R) | 0.570 | 0.498 | 0.763 | 0.788 | 0.527 | 0.588 | 0.403 | 0.726 | 0.561 | 0.677 | 0.508 | 0.743 | 0.609 | 0.612 |
| Kato <i>et al.</i> [8] (NMR) | 0.6172 | 0.4998 | 0.7143 | 0.7095 | 0.4990 | 0.5831 | 0.4126 | 0.6536 | 0.6322 | 0.6735 | 0.4829 | 0.7777 | 0.5645 | 0.6015 |
| Liu <i>et al.</i> [6] (SoftRas (sil.)) | 0.6419 | 0.5080 | 0.7116 | 0.7697 | 0.5270 | 0.6156 | 0.4628 | 0.6654 | 0.6811 | 0.6878 | 0.4487 | 0.7895 | 0.5953 | 0.6234 |
| Liu <i>et al.</i> [6] (SoftRas (sil.+color)) | 0.6670 | 0.5429 | 0.7382 | 0.7876 | 0.5470 | 0.6298 | 0.4580 | 0.6807 | 0.6702 | 0.7220 | 0.5325 | 0.8127 | 0.6145 | 0.6464 |
| Without silhouette supervision | | | | | | | | | | | | | | |
| Uniform backgrounds | | | | | | | | | | | | | | |
| SoftRas (no sil. supervision) | 0.1429 | 0.1112 | 0.2342 | 0.3341 | 0.1589 | 0.1713 | 0.0837 | 0.1916 | 0.0769 | 0.2125 | 0.0830 | 0.1918 | 0.1644 | 0.1659 |
| (Ours) Domain Adaptation (no sil. supervision) | 0.5502 | 0.4441 | 0.6538 | 0.7097 | 0.4995 | 0.5756 | 0.4250 | 0.6337 | 0.6082 | 0.6543 | 0.4475 | 0.7178 | 0.5431 | 0.5740 |
| Minecraft backgrounds | | | | | | | | | | | | | | |
| SoftRas (no sil. supervision) | 0.0251 | 0.0409 | 0.1618 | 0.0885 | 0.1111 | 0.0999 | 0.0437 | 0.2481 | 0.0124 | 0.1233 | 0.0836 | 0.0542 | 0.0337 | 0.0866 |
| SoftRas (sil. generated by SEAM) | 0.2905 | 0.2032 | 0.0843 | 0.5195 | 0.1626 | 0.2732 | 0.1788 | 0.1558 | 0.2573 | 0.2257 | 0.2299 | 0.3778 | 0.3732 | 0.2563 |
| (Ours) Domain Adaptation (no sil. supervision) | 0.2994 | 0.2563 | 0.5002 | 0.6078 | 0.3222 | 0.4135 | 0.2480 | 0.4208 | 0.2551 | 0.4861 | 0.2455 | 0.5485 | 0.3569 | 0.3816 |

and compare the performance with and without domain adaptation. For a second experiment we use images with textured background as described in Sec. 4.1 and shown in Fig. 3 and 4. The results of all settings are shown in Tab. 1.

As we use the same differentiable rendering pipeline as Liu *et al.* (SoftRas) [6] but with weaker supervision, we consider the performance with silhouette supervision as an upper bound for learning without silhouette supervision.

We first compare the performance in the case of uniform backgrounds. Comparing baselines with and without silhouette supervision shows that accuracy of the baselines drops significantly when using uniform backgrounds and thus no silhouette supervision, as it is not designed to operate in this setting in general. With our domain adaptation, we reach an accuracy of 57.4% compared to the 64.6% upper bound under stronger supervision and even exceeds the voxel-based perspective transformer networks [9] as well as the retrieval baseline method [9].

Second, we compare the performance in the case of textured Minecraft backgrounds to two baselines based on the SoftRas [6] algorithm: For the first baseline, we apply the algorithm as is without silhouette supervision, and for the second baseline, we provide a silhouette based on semantic segmentation. To this end, we leverage a state-of-the-art weakly-supervised semantic segmentation method (SEAM) [34] to generate silhouettes from the rendered images with textured backgrounds. We train the segmentation model on all samples of each object class with the respective class label and apply the resulting model to all samples to get the final silhouettes. Examples of the resulting segmentation are displayed in Fig. 5. The respective segmentation map is then used as silhouette supervision for the SoftRas algorithm. SoftRas (sil. generated by SEAM) is, apart from not using ground truth silhouettes, equivalent to

SoftRas (sil.) [6].

Without silhouette information, the baseline achieves a mean accuracy of 0.0866 on images with textured backgrounds, which is a significant drop compared to silhouette supervision. Compared to that, the baseline based on silhouettes generated by the weakly-supervised semantic segmentation method achieves a mean IoU of 0.2563. With the proposed domain adaptation for style-agnostic reconstruction, we achieve a mean IoU of 0.3816. Overall, the domain adaptation outperforms supervision by weakly-supervised silhouette predictions in eleven out of thirteen categories. We show qualitative results in Fig. 4.

4.4. Ablation Study

We evaluate the impact of the different elements of the network for images with textured background in an ablation study shown in Tab. 2. As expected, without the cycle-consistency content loss or the adversarial style loss, the accuracy drops significantly. Note that the Jaccard loss has the special role of preventing the domain-adapted image from differing by a too large factor from the rendered reconstruction. We found that, without the Jaccard loss, the entire architecture becomes less stable; e.g., one out of three test runs without the Jaccard loss achieves a mean IoU of only 0.0166. Without the domain adaptation, i.e., assuming that the rendered images should be the same as the input images, the mean IoU is 0.0866.

4.5. Variation of the Azimuth

As the azimuth, in ours as well as previous methods, is supervised, we perform an experiment challenging the requirement of azimuth supervision. The motivation of this experiment is to examine whether the proposed method would also be able to work in settings where the camera




Figure 4: Examples of the Geometry Reconstruction from single images with Minecraft backgrounds. For further examples, see the supplementary material.

extrinsics are computed from point correspondences and could be noisy. To simulate this setting, we add values




Figure 5: Examples of the weakly-supervised semantic segmentation with SEAM [34]

Table 2: Ablation Study. We evaluate the impact of all proposed losses as well as the entire domain adaptation.

| Method | IoU Acc. |
|-------------------------------|----------|
| Style-Agnostic Reconstruction | 0.3816 |
| Without Content-Loss (Eq. 4) | 0.2293 |
| Without Style-Loss (Eq. 2) | 0.2439 |
| Without Jaccard-Loss (Eq. 3) | 0.1852 |
| Without Domain Adaptation | 0.0866 |

Table 3: Standard deviation of angles added to the azimuth.

| Std. | 0° | 1° | 5° | 15° | 25° |
|----------|--------|--------|--------|--------|--------|
| IoU Acc. | 0.3816 | 0.3688 | 0.3726 | 0.2958 | 0.2786 |

drawn from a normal distribution to the ground truth azimuth values. We consider standard deviations of up to 25°. The results are shown in Tab. 3. We find that, up to a standard deviation of 5°, there is only a small performance penalty indicating that an approximate azimuth is sufficient.

4.6. Behavior of the Domain Adaptation

In Fig. 7, we display the training process for the first 5 training cycles. In each training cycle, after training the domain adaptation, the geometry reconstruction network learns to work in harmony with the domain adaptation. Over the course of many training cycles, the domain adaptation gets closer to the correct solution because the cycle-consistency loss always strives for a more expressive overall system. (This process is displayed in Fig. 7 on the right edge from top to bottom.)

In the initial training cycle, the rendered image is constant because there are no known 3D objects and just the initially guessed sphere is rendered. When the reconstructions improve, the data generating process for the smooth renderings also improves, improving the style-transfer in return. In each training cycle, we train the domain adaptation for 1 500 iterations and then the reconstruction network for




Figure 6: 3D reconstructions from real photos of LEGO objects. The metric is 3D IoU and the average IoU is 0.3316.

2 000 iterations. We perform a total of 20 cycles.

The input, including background, can be reconstructed by the autoencoder because it can be compressed into high-frequency details as explained in detail by Chu *et al.* [33].

In Supp. Mat. A, we investigate the impact of variations in image brightness to our method.

4.7. Reconstructions from ShapeNet2LEgo

In Fig. 6, we show single-view 3D reconstructions from photos of LEGO objects from a model trained on images with Minecraft backgrounds with our method. Here, even throughout different lighting conditions and backgrounds, 3D models can be reconstructed. We found that for the evaluated shapes, we achieve an average 3D IoU compared to the ground truth CAD models of 33.16% which is lower than the performance of the original ShapeNet test set of 38.16%, but still performs in a comparable range given the domain gap compared to the training data. Comparing the reconstructed models in detail, we find that conditions like lightning or azimuth have an impact as can be seen in the reconstruction from image two, three, and four. Here, especially, the legs appear thinner under one projection (image two) than under another angle (image three).

5. Conclusion

We proposed an adversarial style-transfer autoencoder pipeline to generalize state-of-the-art unsupervised single-view 3D geometry reconstruction methods. To this end, we employed image-to-image translation-based adversarial domain adaptation to translate from the input image domain to the rendered image domain, which relieves the need for a photo-realistic renderer, since it does not require the assumption that the renderer can generate the input images. Experiments demonstrate that our style-agnostic approach can improve state-of-the-art methods in settings with uniform backgrounds and textured backgrounds.²

²**Acknowledgements:** This work was supported by the IBM-MIT Watson AI Lab, the DFG in the SFB Transregio 161 “Quantitative Methods for




Figure 7: Training of the adversarial domain adaptation. The top row is the input (repeated for visual unity.) For each block, the top row is $F_{i2r}(x)$, the middle row is the silhouette of $F_{i2r}(x)$, and the bottom row is the image translated back to the original $F_{r2i}(F_{i2r}(x))$. Each of the five blocks represents one of the first five training cycles (from top to bottom). On the horizontal axis, training over 2 500 steps is displayed (from left to right). For further examples, see Supplementary Material D.

Visual Computing” (Project-ID 251654672), and the Cluster of Excellence ‘Centre for the Advanced Study of Collective Behaviour’.

References

- [1] S. Yang, M. Xu, H. Xie, S. Perry, and J. Xia, "Single-View 3D Object Reconstruction From Shape Priors in Memory," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [2] A. Kurenkov, J. Ji, A. Garg, V. Mehta, J. Gwak, C. Choy, and S. Savarese, "DeformNet: Free-form deformation network for 3D shape reconstruction from a single image," in *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018.
- [3] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang, "Pixel2Mesh: Generating 3D mesh models from single RGB images," in *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [4] C. B. Choy, D. Xu, J. Y. Gwak, K. Chen, and S. Savarese, "3D-R2N2: A unified approach for single and multi-view 3D object reconstruction," in *Proc. European Conference on Computer Vision (ECCV)*, 2016.
- [5] A. Kanazawa, S. Tulsiani, A. A. Efros, and J. Malik, "Learning category-specific mesh reconstruction from image collections," in *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [6] S. Liu, T. Li, W. Chen, and H. Li, "Soft Rasterizer: A Differentiable Renderer for Image-based 3D Reasoning," in *Proc. International Conference on Computer Vision (ICCV)*, 2019.
- [7] W. Chen, J. Gao, H. Ling, E. J. Smith, J. Lehtinen, A. Jacobson, and S. Fidler, "Learning to predict 3D objects with an interpolation-based differentiable renderer," in *Proc. Neural Information Processing Systems (NeurIPS)*, 2019.
- [8] H. Kato, Y. Ushiku, and T. Harada, "Neural 3D mesh renderer," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [9] X. Yan, J. Yang, E. Yumer, Y. Guo, and H. Lee, "Perspective transformer nets: Learning single-view 3D object reconstruction without 3D supervision," in *Proc. Neural Information Processing Systems (NeurIPS)*, 2016.
- [10] Y. Jiang, D. Ji, Z. Han, and M. Zwicker, "SDFDiff: Differentiable rendering of signed distance fields for 3D shape optimization," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [11] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "ShapeNet: An information-rich 3D model repository," *Computing Research Repository (CoRR) in arXiv*, 2015. [Online]. Available: <https://arxiv.org/abs/1512.03012>.
- [12] M. M. Loper and M. J. Black, "OpenDR: An approximate differentiable renderer," in *Proc. European Conference on Computer Vision (ECCV)*, 2014.
- [13] T.-M. Li, M. Aittala, F. Durand, and J. Lehtinen, "Differentiable Monte Carlo ray tracing through edge sampling," *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, vol. 37, no. 6, pp. 1–11, 2018.
- [14] P. Henderson and V. Ferrari, "Learning to Generate and Reconstruct 3D Meshes with only 2D Supervision," in *International Journal of Computer Vision*, 2019.
- [15] A. Delaunoy and E. Prados, "Gradient flows for optimizing triangular mesh-based surfaces : Applications to 3D reconstruction problems dealing with visibility," *International Journal of Computer Vision*, vol. 95, pp. 100–123, 2011.
- [16] R. Ramamoorthi and P. Hanrahan, "A signal-processing framework for inverse rendering," in *Proc. SIGGRAPH*, 2001.
- [17] A. Meka, M. Maximov, M. Zollhoefer, A. Chatterjee, H.-P. Seidel, C. Richardt, and C. Theobalt, "LIME: Live intrinsic material estimation," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [18] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok, "Synthesizing robust adversarial examples," in *International Conference on Machine Learning (ICML)*, 2018.
- [19] E. Richardson, M. Sela, R. Or-El, and R. Kimmel, "Learning detailed face reconstruction from a single image," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [20] G. Loubet, N. Holzschuch, and J. Wenzel, "Reparameterizing discontinuous integrands for differentiable rendering," *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 2019.
- [21] A. Kundu, Y. Li, and J. M. Rehg, "3D-RCNN: Instance-level 3D object reconstruction via render-and-compare," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [22] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas, "Learning Representations and Generative Models for 3D Point Clouds," in *International Conference on Machine Learning (ICML)*, 2018.
- [23] H. Fan, H. Su, and L. 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.
- [24] L. Jiang, S. Shi, X. Qi, and J. Jia, "GAL: Geometric adversarial loss for single-view 3D-object reconstruction," in *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [25] Y. Zhang, W. Chen, H. Ling, J. Gao, Y. Zhang, A. Torralba, and S. Fidler, "Image GANs meet Differentiable Rendering for Inverse Graphics and Interpretable 3D Neural Rendering," in *International Conference on Learning Representations (ICLR)*, 2021.

- [26] C. H. L. Paul Henderson Vagia Tsiminaki, "Leveraging 2d data to learn textured 3d mesh generation," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [27] L. A. Gatys, A. S. Ecker, and M. Bethge, "A neural algorithm of artistic style," *Journal of Vision*, vol. 16, no. 12, 2015.
- [28] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [29] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. International Conference on Computer Vision (ICCV)*, 2017.
- [30] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "CyCADA: Cycle-Consistent Adversarial Domain Adaptation," in *International Conference on Machine Learning (ICML)*, 2018.
- [31] G. Wilson and D. J. Cook, "A survey of unsupervised deep domain adaptation," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 11, no. 5, pp. 1–46, 2020.
- [32] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Proc. International Conference on Computer Vision (ICCV)*, 2017.
- [33] C. Chu, A. Zhmoginov, and M. Sandler, "CycleGAN, a master of steganography," in *NIPS 2017 Workshop "Machine Deception"*, 2017.
- [34] Y. Wang, J. Zhang, M. Kan, S. Shan, and X. Chen, "Self-supervised Equivariant Attention Mechanism for Weakly Supervised Semantic Segmentation," in *Proc. International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.