

3D Reconstruction of Sculptures from Single Images via Unsupervised Domain Adaptation on Implicit Models

Ziyi Chang
ziyi.chang@durham.ac.uk
Durham University
Durham, UK

George Alex Koulieris
georgios.a.koulieris@durham.ac.uk
Durham University
Durham, UK

Hubert P. H. Shum*
hubert.shum@durham.ac.uk
Durham University
Durham, UK




Figure 1: Two reconstruction examples (left and middle images) and display in virtual reality (right image). Our network generates 3D sculptures from single RGB images. Results can be used in virtual reality applications such as virtual museum.

ABSTRACT

Acquiring the virtual equivalent of exhibits, such as sculptures, in virtual reality (VR) museums, can be labour-intensive and sometimes infeasible. Deep learning based 3D reconstruction approaches allow us to recover 3D shapes from 2D observations, among which single-view-based approaches can reduce the need for human intervention and specialised equipment in acquiring 3D sculptures for VR museums. However, there exist two challenges when attempting to use the well-researched human reconstruction methods: limited data availability and domain shift. Considering sculptures are usually related to humans, we propose our unsupervised 3D domain adaptation method for adapting a single-view 3D implicit reconstruction model from the source (real-world humans) to the target (sculptures) domain. We have compared the generated shapes with other methods and conducted ablation studies as well as a user

study to demonstrate the effectiveness of our adaptation method. We also deploy our results in a VR application.

CCS CONCEPTS

• Computing methodologies → Unsupervised learning; Shape modeling.

KEYWORDS

3D Reconstruction, Unsupervised Learning, Domain Adaptation, Transfer Learning, VR

ACM Reference Format:

Ziyi Chang, George Alex Koulieris, and Hubert P. H. Shum. 2018. 3D Reconstruction of Sculptures from Single Images via Unsupervised Domain Adaptation on Implicit Models. In *Proceedings of The 28th ACM Symposium on Virtual Reality Software and Technology (VRST '22)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

VRST '22, Nov. 29–Dec. 1, 2022, Tsukuba, Japan
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

While virtual reality (VR) is an effective tool for implementing virtual museums, the acquisition of the virtual equivalent of exhibits, such as sculptures that we consider in this paper, can be labour-intensive. Many exhibits are 3D in nature, such as sculptures, historical objects and specimens. Typical solutions for acquiring 3D surfaces such as 3D scanning generally require specific hardware and potentially manual post-processing. A machine-learning-based,

software solution that requires minimal human intervention is beneficial for VR applications. This paper focuses on sculptures as they have a wide variety of shape topology and are challenging due to the general lack of texture.

Deep learning based 3D reconstruction has the potential of effectively acquiring 3D models from 2D observations. View synthesis approaches such as NeRF [23] can render new views, but may incur inconsistency across views due to the lack of explicit 3D modelling. General object reconstructions such as [26] may reconstruct a wide variety of objects. However, due to the lack of prior knowledge on the type of objects, reconstruction quality may not always be consistent. We observe that sculptures generally resemble real-world humans while having variations on body structure. This motivates us to adapt the well-researched 3D human reconstruction methods such as [33] to reconstruct sculptures.

To minimise the amount of human intervention and the use of specialised equipment, implicit 3D human reconstruction from a single RGB image is preferred. Previous single-view based human reconstruction methods can be categorized into three streams: implicit reconstruction, parametric reconstruction, and their combination. The first one is to learn an implicit function such as signed distance function [26] and occupancy probability [2]. The second one builds upon various parametric human models such as SMPL [19] and uses neural network to estimate a few parameters to deform the pre-defined parametric human models. The combination [3] blends previous two streams to add surface details to parametric human models. Compared with regular human shapes, sculptures have more irregular topology. Parametric human models cannot be deformed into another topology. To solve the irregular topology of sculptures, we consider implicit models are more suitable than parametric ones for our reconstruction problem.

There are two major challenges when adapting single-view implicit 3D reconstruction for sculptures. First, obtaining 3D data are labour-intensive and sometimes infeasible [2], which is particularly true for sculptures as they may be of larger sizes and partially combined together with the environment. Second, sculptures generally consist of irregular topology, making it difficult to apply pre-trained models for reconstruction. This is known as domain shift [40] in machine learning.

Our insight is that the aforementioned challenges can be solved by unsupervised domain adaptation of 3D human reconstruction systems. To tackle the limitation of data, unsupervised learning enables our system to acquire knowledge of 3D sculptures without large well-labelled datasets [24]. To tackle domain shift, due to the similarity of human-based shapes and sculptures, domain adaptation on an implicit model can transfer knowledge learnt from one domain into another [40] without retraining while allowing the modelling of the different distribution of topology.

In this work, we propose a new method of unsupervised 3D domain adaptation for 3D sculpture reconstruction tasks. Our pipeline utilises single RGB images as input and transfers a pre-trained implicit model from source (real-world human) into target (sculpture) domain in an unsupervised manner. We consider pre-trained PIFu [33] is more suitable than other implicit models due to limited data availability. To represent domain shift, we adapt multi-layer intermediate feature architecture and Maximum Mean Discrepancy (MMD) into our problem of 3D unsupervised domain adaptation.

Multiple features in the encoder extract different patterns from low to high levels. Our MMD is based on multiple levels of features and aligns two domains. To refine geometry and surface details, we propose neighbour aggregation and pseudo labels for our problem of 3D unsupervised domain adaptation on 3D sculpture reconstruction. Guidance is designed to transmit from the source domain to the target domain through aggregation and pseudo labels. During the aggregation, we re-weight features that are used to define neighbours by balancing the importance of pixel-aligned information with depth values. To the best of our knowledge, we present the first research on unsupervised 3D domain adaptation with implicit models on the 3D reconstruction task.

Experiments show our network exceeds state-of-the-art on our collected sculptures for evaluation. We have provided qualitative results generated by various methods. As for quantitative comparison, we choose Chamfer distance and point-to-surface distance as the two metrics - in both of which our methods achieve the best results. We also conduct extra experiments on other alternative domain adaptation methods to demonstrate the effectiveness of our design. Additional ablation studies on proposed modules and design choices have been implemented. In addition, we have also invited 20 volunteers to participate in our user study. Finally, we deploy the reconstructed sculptures in a VR application.

Our contributions are as follows, with source code available at <https://github.com/mrzzzy2021/SculptureRecon>:

- We propose an unsupervised method for domain adaptation on sculpture reconstruction, solving the problem of creating 3D shapes with limited available data.
- We adapt multi-level features into our problem of unsupervised 3D domain adaptation on a 3D reconstruction task.
- We propose to use a re-weighting and neighbourhood strategy for better structural exploration of latent feature space.

2 RELATED WORK

2.1 3D Reconstruction

3D reconstruction has been applied in recovering a variety of subjects. Some approaches mainly focus on human reconstruction [5, 36]. Others are specialized for objects [10, 12]. For example, [8] is designed for humans while [25] focuses on reconstructing 3D cars. There are also methods for arbitrary subjects such as NeRF [23]. However, the reconstruction quality may not be consistent due to the lack of prior knowledge. In our paper, we observe that many sculptures are artistic interpretations of real world human beings and thus, we focus more on human reconstruction to better leverage domain knowledge and control computational cost.

Among 3D human reconstruction methods, they can be categorized into single-view and multi-view human reconstruction [41]. Single-view methods can recover 3D shape from a single image [33, 36]. Multi-view approaches make predictions based on several images or videos that carry shape information from different angles [5, 13]. While multi-view human reconstruction can achieve better performance, it requires more information as input, which may not be accessible due to limited data acquisition. On the contrary, single-view reconstruction methods are more flexible and easy to operate when data acquisition is difficult. As it may not be easy to collect multi-view images for sculptures due to complex environmental

restrictions, we are interested in single-view reconstruction. This also makes our methodology favorable for practical deployment of AR/VR applications among others, as it is easier to obtain a single image of a sculpture when compared with multiple images.

There are three main streams of methods for reconstruction of real-world humans. The first stream uses implicit models, which is proposed in recent years and has become more and more popular. Implicit models rely on a learned function to describe surfaces [41]. There are two popular implicit representations available for implicit models. One is occupancy probability, which indicates the probability of the queried 3D coordinate inside the 3D shape and many studies have trained implicit functions to predict the probability [33, 34]. They focus on how to extract more representative image features, by aligning pixel-wise features to coordinates or iteratively extracting multi-level local features, to facilitate more accurate classification on sampled locations. The other representation is the signed distance function [18, 26]. Instead of using probability to indicate different properties of 3D coordinates, signed distance functions indicate the positive or negative distance from sampled locations to the surface [39].

The second stream of methods for the reconstruction of real-world humans uses parametric models. These models are pre-defined manually and controlled by a few parameters. A large number of parametric human models have been proposed such as SMPL [19], MANO [31], SMPL-X [28] and STAR [27]. Earlier studies usually estimate the corresponding parameters directly on the provided images. To achieve better results, recent works tend to estimate parameters in cooperation with other extra information. [1] makes use of UV transformation to generate normals and displacements and then applies them to SMPL models to get a clothed 3D human body. [37] directly fits a SMPL model to silhouette and posture and then refines the texture of the SMPL model by normal and skinning maps. [42] assigns pre-defined semantic representation to every estimated SMPL vertex and voxelises mesh to better refine shape details. [8] proposes to use UV map to indicate occlusion and leverages convolution kernels on 2D images to encode style features and size features to model clothes which are then combined with SMPL parameters for reconstruction.

The third stream is to combine parametric models with implicit representations for more accurate modelling [41]. Although in the second and third streams, parametric methods can reconstruct 3D real-world humans, they cannot handle irregular topology because of their pre-defined parametric models. Therefore, we consider implicit methods are more suitable for our reconstruction problem.

2.2 Domain Adaptation

There are three types of domain adaptation methods. The first one is supervised domain adaptation [17] and requires that both the source and target domains are well-labelled. The second category is semi-supervised domain adaptation [6] where partial target data are well-labelled. The third one is unsupervised domain adaptation [40] and does not rely on target labels to transfer the learnt knowledge from the source domain to the target domain. Considering the limited data acquisition, we believe unsupervised domain adaptation is more suitable to solve our problem.

Unsupervised domain adaptation has been widely used on 2D or 2.5D tasks. Previous researches on 2D can be categorised into two groups. One is based on learning domain-invariant features and the other is to use a network to directly learn the mapping. The former category [15] minimises domain shifts by reducing the difference between feature distributions in feature space. The latter group [4] models the mapping by training a network such as CycleGAN [43]. For 2.5D, [35] utilises contrastive learning to maintain geometries with depth images. For these different inputs, there are a number of popular methods such as pseudo-labeling [32], and batch normalization tailored for domain adaptation [22].

Though lots of effort has been made on 2D images or depth maps, few works have been proposed for 3D reconstruction. Unsupervised domain adaptation for 3D reconstruction was introduced by [29]. [16] assumes their data are from the same object category and proposes a method that performs translations between 3D and 2D representations, which in other words are two domains. It focuses more on training networks by unpaired images and shapes. [38] fills the domain gap in 3D reconstruction between synthetic images and real images by extracting domain-invariant image features. This work utilises graph neural network for mesh generation and proposes DDA to extend the network to solve the synthetic-to-real difference [21, 30]. Instead of synthetic-to-real differences, our research focuses on the topology shift and we propose our domain adaptation methods to solve the sculpture reconstruction problem as a transferring problem.

3 3D SCULPTURE RECONSTRUCTION




Figure 2: Overview of our proposed unsupervised domain adaptation on 3D sculpture reconstruction.

Reconstructing 3D sculptures from single RGB images has two main challenges. First, the difficulty on data acquisition leads to limited labels and training data. It is infeasible to obtain other auxiliary information such as surface normals, prohibiting the use of some reconstruction modules [34]. The other challenge is the domain shift of topology. The irregular topology of 3D sculptures makes domain transferring non-linear. The ideal 3D representation used in the reconstruction module should be flexible enough to cover various potential 3D shapes.

To solve the aforementioned challenges, we choose PIFu [33] that is pre-trained on real-world humans and propose our unsupervised domain adaptation (UDA) pipeline. Using a human reconstruction network aids the unsupervised domain adaptation. We observe that

sculptures are usually artistic interpretations of humans. To alleviate the difficulty of transferring, we leverage the already-learned knowledge on real-world human shapes as a solid starting point. PIFu is a single-view reconstruction method and only requires a single RGB image as input, mitigating the difficulty on data acquisition. Meanwhile, PIFu is based on implicit representation that learns a function to approximate the surface. Compared with parametric representation, PIFu is not restricted by the pre-defined canonical models and is flexible enough to learn irregular topology of 3D sculptures. We also design our domain adaptation to be unsupervised to transfer knowledge of real-world humans into sculptures without requiring labels in the target domain.

3.1 Problem Definition

Considering the challenges, our goal is to propose an unsupervised domain adaptation (UDA) method for PIFu to solve the problem of reconstructing 3D sculptures from single RGB images. We will first briefly introduce the pre-trained PIFu used in our domain adaptation process. Then, we will discuss our unsupervised domain adaptation settings. The pipeline of PIFu is shown in Figure 3. PIFu

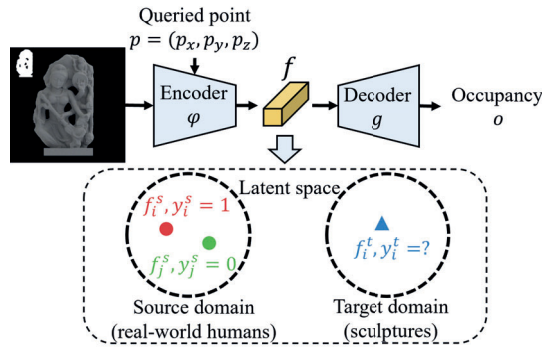


Figure 3: Domain definition. Our source and target domains are composed of pixel-aligned features for each queried point. The target domain only has features f^t from sculptures and does not have ground truth y^t .

predicts the occupancy probability for any queried point in space. It focuses on learning an implicit function $func(p, I, M)$ for a query point $p = (p_x, p_y, p_z) \in \mathbb{R}^3$, which is formulated as:

$$func(p, I, M) = g(\varphi(p_x, p_y, I, M), p_z) = o \in [0, 1], \quad (1)$$

where I and M are an image and its mask, φ and g are encoder and decoder, respectively, o is the predicted occupancy probability for the queried point p . The pixel-aligned feature f in PIFu for the queried point $p = (p_x, p_y, p_z)$ is formed by two concatenated parts: a feature based on bilinear interpolation with (p_x, p_y) and the corresponding depth value p_z . The learned shape \mathcal{S} is extracted using marching cubes:

$$\mathcal{S} = \{p \in \mathbb{R}^3 | o = 0.5\}. \quad (2)$$

Our 3D unsupervised domain adaptation (UDA) settings are based on the pixel-aligned features extracted by the PIFu encoder. For our UDA with implicit reconstruction models, we have access to the labeled source domain that consists of n_s queried points for

each real-world human mesh. The PIFu encoder then converts these queried points into their corresponding pixel-aligned features. As we use multi-level architecture (to be explained in Section 3.2), we define our source domain as $\mathcal{D}^s = \{f_i^{s,l}, y_i^s\}$ for $i = 1, \dots, n_s$ and $l = 1, \dots, L$ where the source labels y_i^s are defined as:

$$y_i^s = \begin{cases} 1 & p_i^s \text{ is inside mesh surface,} \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

and n_s is the number of pixel-aligned features. Similarly, the target domain for the l -th layer is defined as $\mathcal{D}^{t,l} = \{f_i^{t,l}\}$ for $i = 1, \dots, n_t$, which consist of n_t features and the target labels are unknown. Our aim is to estimate their corresponding occupancy probability $o^{t,l} \in \mathcal{Y}^{t,l}$ through our UDA method. Note that y in Eq.3 represents the ground truth label, which is only possibly 1 or 0. On the contrary, o in Eq.1 means the predictions of our neural network, which is within the range from 0 to 1.

3.2 Multi-level Features

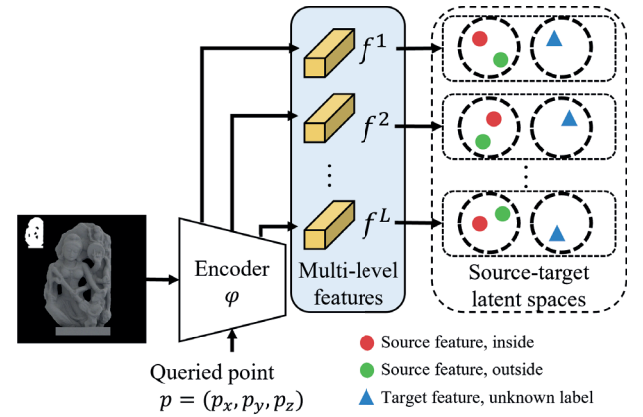


Figure 4: We propose to adapt multi-level pixel-aligned features to solve our unsupervised domain adaptation problem.

We propose to adapt the multi-level pixel-aligned features into our UDA method for the sculpture reconstruction problem, as shown in Figure 4. The latent features from different encoder layers are usually sensitive to a subset of input patterns. This facilitates our UDA to better measure domain shifts using comprehensive information from low to high network levels.

We also propose to process the multi-level features layer by layer dynamically. Instead of mixing them up, our design processes these features layer by layer. In this way, the relative information of occupancy within each layer can be maintained and perceived by our adaptation method.

The multi-level pixel-aligned features come from image features extracted by the PIFu encoder. We choose four output layers of the stacked hourglass network in the PIFu encoder to obtain the pixel-aligned features. As shown in Figure 4, we first extract the multiple image features by different layers of PIFu encoder and then we assign the pixel-aligned features to each queried point according to their pixel-aligned 2D locations. After obtaining features, we form a number of source-target latent spaces with features from

the same layers. Multi-layer features $\{f^l\}_{l=1}^L$ are extracted when given a point p and layer l .

3.3 Domain Alignment

We propose our domain alignment method to align features from the source (real-world humans) and the target (sculpture) domains for our reconstruction task. Our domain alignment has two components, minimising discrepancy and maintaining accuracy.

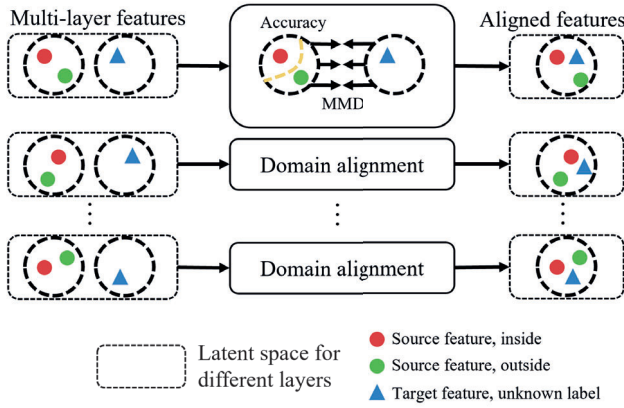


Figure 5: Our domain alignment utilises Maximum Mean Discrepancy and maintains the accuracy in the source domain.

For the first component, we propose to adapt Maximum Mean Discrepancy (MMD) [11] on pixel-aligned features from each layer to mix up the source and target domains for 3D reconstruction. MMD is a predefined kernel metric to quantify the similarity between two distributions. It leverages kernel functions to map pixel-aligned features from two domains into a latent space where the similarity can be calculated.

Compared to other learning-based methods, MMD is more suitable due to the limited data availability. As discussed in Section 3.1, data acquisition is a challenge that leads to limited training data without labels. Under this problem scenario, we choose to use this predefined measurement instead of learnt-based methods such as discriminators in adversarial training. Although learnt-based methods are more flexible, they require abundant training data to extract knowledge for obtaining appropriate representations. For example, discriminators trained on our small dataset cannot acquire enough knowledge to provide guidance for adversarial learning, which we will show results for in Section 4. By decreasing the distribution difference, the encoder is forced to pay more attention to occupancy-related features.

Following our multi-layer architecture, our MMD is imposed to pixel-aligned features for each layer and we use the average value as our similarity loss for the adaptation process. Each pixel-aligned feature is projected into the Reproducing Kernel Hilbert Space (RKHS). The empirical expectation is computed for two domains and the difference can also be computed. As there are several MMD values from multiple layers, we use the average value of the computed MMD values across layers as the similarity loss. It can

be formulated as:

$$\begin{aligned} \mathcal{L}_{sim} &= \frac{1}{L} \sum_{l=1}^L \mathcal{L}_{mmd}^l \\ &= \frac{1}{L} \sum_{l=1}^L \left\| \frac{1}{n_s} \sum_{f^{s,l} \in \mathcal{D}^{s,l}} f^{s,l} - \frac{1}{n_t} \sum_{f^{t,l} \in \mathcal{D}^{t,l}} f^{t,l} \right\|_{\mathcal{H}}^2 \end{aligned} \quad (4)$$

where \mathcal{H} corresponds to the Reproducing Kernel Hilbert Space (RKHS) with Gaussian kernel.

The second component is designed to keep accuracy on predictions in the source domain. The domain aligning operation changes pixel-aligned features to achieve domain invariance. As only the similarity is considered in MMD values, we propose to keep the accuracy of occupancy prediction in source domain where points are well-labelled. We regulate the aligned features to be meaningful by leveraging the mean square error (MSE) to maintain classification accuracy with labels $y^s \in \{0, 1\}$:

$$\mathcal{L}_{source} = \text{MSE}(g(f^{s,l}), y^s). \quad (5)$$

3.4 Latent Structural Exploration

Due to the lack of target labels, we propose our method to explore the spatial structure of latent feature space with guidance from the source domain labels. Our exploration consists of two components, neighbour aggregation and diversity maintenance.

As shown in Figure 6, the first component, neighbour aggregation, is designed to describe region density by averaging source neighbour labels and updating the target domain with pseudo labels. Generally, the higher region density a particular label has, the higher possibility that a point in this region should have the same label. Considering the noise of target predictions, we treat the top-K nearest features from only the source domain as neighbours. The averaged source neighbour label has been used to combine with the target predictions as the pseudo labels. This combination transmits the density clue provided by source labels to the target domain.

We empirically propose to increase the importance of depth value in features within neighbour aggregation. As single-view 3D reconstruction is an ill-posed problem and usually suffers from depth ambiguity, we hereby increase the importance of depth value in queried point coordinates when searching neighbours. Specifically, we re-weight pixel-aligned features by multiplying a scaling factor to the depth value which we empirically set as 256. The re-weighted features can better reveal structure and relationship when measured by Euclidean distance.




Figure 6: We propose to use the averaged labels from source neighbours as the indication of regional density.

As we are using a multi-layer architecture, we calculate the neighbourhood information for each layer. Concretely, we use the averaged labels over source domain neighbours to represent the region density for target domain points. The aggregated information for point-feature $f^{t,l}$ from the l -th layer is defined as:

$$\hat{y}^{t,l} = \frac{1}{|\mathcal{N}_k^{s,l}|} \sum_{f^{s,l} \in \mathcal{N}_k^{s,l}} y^{s,l}, \quad (6)$$

where $\hat{y}^{t,l}$ is the aggregated information for the target intermediate feature $f^{t,l}$ in l -th layer and $\mathcal{N}_k^{s,l}$ is the set of top- K nearest source features to $f^{t,l}$ in the latent space.

Our pseudo labels are designed to combine source regional density and target domain predictions. Instead of directly assigning the computed neighbourhood information as the pseudo labels for target domain points, we consider the characteristics of target structure. A momentum is used to balance the combination of target and source structural information. The aggregated result $\hat{y}^{t,l}$ is used for assigning pseudo-labels to target feature points. These labels convey information regarding the surroundings of a feature point and thus can provide guidance on target domain learning. The assigned pseudo-label can be calculated by balancing the self-momentum and its neighbourhood information:

$$\hat{y}^{t,l} \leftarrow m \times o^{t,l} + (1 - m) \times \hat{y}^{t,l}, \quad (7)$$

where $o^{t,l}$ is original prediction for target domain, m is the self-momentum and increases as $m = \frac{\text{epoch-start_epoch}}{\text{epoch_total}}$ where we set $\text{start_epoch} = 30$ and $\text{epoch_total} = 60$ due to larger noise for self-predictions at the beginning of the process. The assigned pseudo-label is then used in MSE loss on the target domain:

$$\mathcal{L}_{\text{target}} = \text{MSE}(g(f^{t,l}), \hat{y}^{t,l}). \quad (8)$$

The second component is concerned with prediction diversity of the target domain by maximizing the mutual information loss. As discussed in Section 3.2, implicit-based reconstruction requires to sample points from the entire space to represent a 3D shape. As we are moving point-features in high-dimensional latent space, the diversity is maintained to avoid the potential overly-concentration in a small area. Therefore, we compute the entropy-based loss as our diversity measurement:

$$\begin{aligned} \mathcal{L}_{\text{mi}} &= H(\mathcal{Y}^{t,l}) - H(\mathcal{Y}^{t,l} | \mathcal{D}^{t,l}), \\ &= h(\mathbb{E}_{f^{t,l} \in \mathcal{D}^t}(g(f^{t,l}))) - \mathbb{E}_{f^{t,l} \in \mathcal{D}^t}(h(g(f^{t,l}))), \end{aligned} \quad (9)$$

where \mathbb{E} is the expectation over point-features and $h(x) = -\sum_i x \log x$ is the conditional entropy.

Combining the aforementioned losses, we design the overall loss as a weighted sum:

$$\mathcal{L} = w_1 \mathcal{L}_{\text{sim}} + w_2 \mathcal{L}_{\text{source}} + w_3 \mathcal{L}_{\text{target}} + w_4 \mathcal{L}_{\text{mi}}, \quad (10)$$

where $\{w_i\}_{i=1}^4$ are weights for different loss elements and are empirically set to $w_1 = 5$, $w_2 = 2$. We use dynamical values for w_3 and w_4 as $\frac{\text{epoch-start_epoch}}{\text{epoch_total}}$ where we define $\text{start_epoch} = 30$ and $\text{epoch_total} = 60$. This is because $\mathcal{L}_{\text{target}}$ will be inaccurate and diversity will be less important at the beginning of training.

4 EXPERIMENTS

4.1 Implementation Details

As previously mentioned, obtaining 3D data of sculptures is non-trivial and we want to figure out whether our method can work with limited data size. We collected 28 sculpture meshes from the website *ScanTheWorld* and split them into a training set (20 meshes) and a testing set (8 meshes). The testing set has been hidden from all methods, including supervised and unsupervised ones, in their training process. We report quantitative performance on the testing set. As for real-world humans, we keep the same setting of what PIFu uses. We found that using only 3 real-world human meshes works for our method as we want to reduce the dependence on large amounts of labelled data required in our domain adaptation. As the collected sculptures do not have colour, we assign random colours to each vertex. Both sculptures and real-world humans are processed to obtain RGB images and corresponding masks under the same rendering settings as used in PIFu [33].

We use the point-to-surface distance (P2S) and chamfer distance (CD) for the quantitative performance evaluation. The former captures the surface details by computing the distance between points in the reconstructed surface and their nearest surface of ground truth. The latter metric captures the overall similarity of global shapes by computing the distance between points from the reconstructed shape and ground truth. The smaller both metrics are, the higher the quality is.

We run our code on one TITAN XP GPU. The training can be completed within approximately 2 days. We use the Marching Cubes algorithm [20] to extract the 0.5-level surface as the mesh shape. We filter isolated parts, rescale meshes into a unit bounding box, and place the mesh in the origin of the coordinate system.

4.2 Quantitative and Qualitative Comparisons

We firstly compare our method with untrained PIFu [33] and untrained PIFuHD [34] as baseline. Both PIFu and PIFuHD are pre-trained on real-world humans and our method is learnt from our training set in an unsupervised manner. All the three compared methods are not exposed to 3D ground truth. The quantitative performance is reported on our testing set as shown in Table 1. Our method surpasses the other two methods in both metrics. The qualitative comparison is shown in Figure 7. Our method reconstructs complete shapes and corrects surface details. When trained with the shape of real-world humans, the model only learns a canonical human shape. Therefore, the model attempts to imitate the shape of a real-world human body for 3D reconstruction. As shown in Figure 8, PIFu and PIFuHD tend to recover two legs and feet that are usually observed in the lower part of the real-world human shapes while the sculpture has a base instead of human legs and feet. Such a topology change can be handled by our method.

We then compare our method with alternative unsupervised adaptation methods to validate our unsupervised design. As no previous study is conducted on unsupervised domain adaptation for 3D reconstruction with implicit models, we reimplement two adaptation methods: contrastive learning (Contra) based on [7] and adversarial learning with gradient reversal layer (GRL) based on [9]. We train them with PIFu (denoted as PIFu + Contra and PIFu + GRL) on our training set in an unsupervised way. Table 1 and




Figure 7: Qualitative comparison. DAREC produces point clouds and we transfer them into meshes.




Figure 8: An example of topology change. Both PIFu and PIFuHD tend to mimic the real-world human shape distribution, such as trying to reconstruct two legs and feet.

Table 1: Quantitative comparison between different methods. All methods compared here do not have access to 3D ground truth of sculptures. We use † to denote a method fails on some results after filtering.

Methods	P2S	CD
PIFuHD [34]	0.084 [†]	0.099 [†]
PIFu [33]	0.044	0.048
PIFu [33] + Contra [7]	0.080	0.057
PIFu [33] + GRL [9]	0.06 [†]	0.069 [†]
DAREC [29]	0.106	0.087
Ours	0.038	0.047

Figure 7 show the quantitative and qualitative results, respectively. We also implement DAREC [29]. PIFu + GRL does not perform well because the discriminator cannot provide correct guidance due to

the limited dataset size. [29] proposes a domain adaptation method, DAREC, to solve synthetic-to-real 3D object reconstruction. Their method leverages a 3D shape autoencoder that is pretrained on a large amount of well-labelled 3D shapes. Then an image encoder is trained with two discriminators to achieve domain confusion between extracted features from images and keep extracted features lying on the shape manifold. There are two main differences between [29] and ours. The first one is that [29] manipulates the single global shape-features to learn canonical human shapes for domain confusion while ours uses a set of point-features to learn occupancy probability for domain confusion. The second one is that [29] relies on discriminators for domain confusion while ours is based on predefined MMD. The two differences enable our method to complete the domain adaptation in our research while [29] does not perform well. The manifold formed by shape-features is based on whether discriminators consider reasonable and thus, cannot handle topology changes. Our point-features form the manifold determined by whether the point is inside or not. This provides flexibility on topology changes. Similar with PIFu + GRL, the discriminators in DAREC for domain confusion require a large amount of training data and cannot provide correct guidance due to the limited dataset size. On the contrary, our predefined MMD-based method can achieve domain confusion. PIFu + Contra does not perform well because implicit models rely on samples around 0.5 to approximate where the 0.5-level surface is. Contrastive learning reduces the distance of the similar samples in the latent space. The concentration to some particular values leads to large uncertainty around 0.5-level surface as shown in Figure 9.




Figure 9: Explanation for ‘PIFu + Contra’ on its over-concentrated feature space.

We compare our unsupervised method with two supervised PIFu methods. We train the two supervised methods, denoted as retrained PIFu and fine-tuned PIFu, with 3D labels in our training set. The two supervised methods directly learn with 3D ground truth labels in our training set. The testing set is hidden during training. These unseen data are used to report the quantitative performance in this experiment. Table 2 shows quantitative results. Our method achieves the best P2S but does not in CD. However, CD is not completely consistent with visual quality [14] and thus, we provide a qualitative comparison in Figure 10. We find out that both retrained PIFu and fine-tuned PIFu recover noisy surface.

Table 2: Quantitative comparison between our unsupervised method with two supervised methods.

Methods	P2S	CD
Retrained PIFu	0.048	0.039
Fine-tuned PIFu	0.046	0.039
Ours	0.038	0.047

The two supervised methods have a severe generalization issue and, as a result, they cannot be used in real-world applications. Figure 10 shows they do not perform well on testing data although their CD is better than ours. As pointed out by [14], this is because CD is relatively insensitive to visual artifacts such as poor surface quality. This is also the reason that we use P2S to capture surface quality and visualize the reconstructed results. This generalization problem is caused by insufficiency of the training data and thus, retraining or fine-tuning is incapable of learning surface.

4.3 Ablation study

Our ablation study validates four modules (indicated by loss term) used in domain adaptation. We provide results under different modules used in domain adaptation in Table 3.

We also provide a quantitative comparison by disabling two design choices: the multi-level features and the rescaling factor. Without a multi-level architecture for unsupervised adaptation, our performance drops a lot on both metrics. This indicates that multi-level features provide more comprehensive clues for our unsupervised domain adaptation. We also find that performance improves when we use a rescaling factor to balance the importance of depth value and pixel-aligned features. With this rescaling factor, our

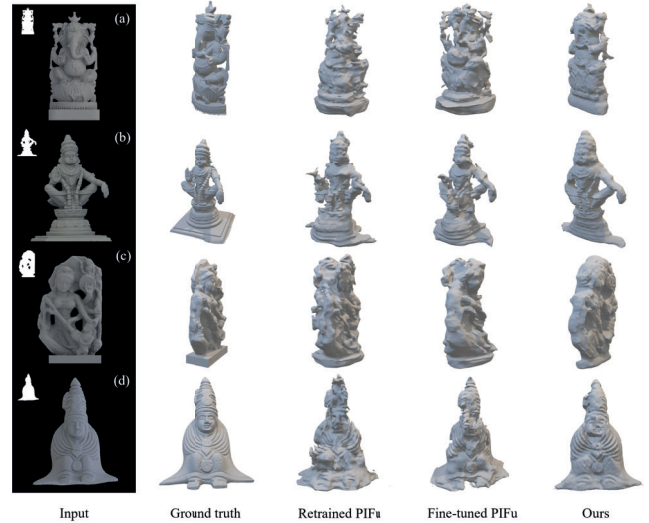


Figure 10: Qualitative comparison with supervised methods, retrained PIFu and fine-tuned PIFu. The two methods they severely struggle with surfaces.

Table 3: Ablation study. ‡ means a method fails to reconstruct meshes after filtering.

Methods	P2S	CD
Ours w/o \mathcal{L}_{mmd}	0.053	0.052
Ours w/o \mathcal{L}_{source}	‡	‡
Ours w/o \mathcal{L}_{target}	0.041	0.049
Ours w/o \mathcal{L}_{mi}	0.041	0.047
Ours w/o multi-level features	0.056	0.052
Ours w/o rescaling	0.043	0.05
Ours	0.038	0.047

adaptation leverages the latent structure in feature space better. As shown in Table 3, by applying two choices, our method can achieve the best performance.

4.4 User study

To evaluate the results from the users’ perspective, we design a user study. We invited 20 participants (10 males and 10 females, the mean age is 34.75 and the standard deviation is 16.98) coming from different educational/occupational backgrounds. Each participant was presented with one 3D shape and a corresponding RGB image at a time. To avoid bias, the results from various methods are mixed, unlabelled and shown in random order. We asked each participant to give a score from 1 to 5 (the higher, the better) based on the similarity of the 3D shape and RGB image. The results are presented in Figure 11. Our method achieves the highest score.

4.5 VR Applications

We also display our results on a commercial VR headset. With VR equipment, users can place, move and observe the generated sculptures from all directions as shown in Figure 12.




Figure 11: User study on results. We calculate the averaged score for results generated by each method.



Figure 12: VR displaying of our results. (Top image) Seeing the generated sculpture. (Bottom image) Locating the result.

5 CONCLUSION

In conclusion, we are the first to apply unsupervised domain adaptation into implicit models to solve the problem of reconstructing 3D sculptures from single-view RGB images, and our method can be deployed in creating various content for VR/AR applications. Our method addresses two critical challenges. The first one is the limited data acquisition, leading to inadequate data for supervised learning and incomplete shape information. The second one is the domain shift of irregular topology. We propose our unsupervised domain adaptation based on multi-level features with distribution alignment and neighbourhood information to transfer pre-trained PIFu to solve the problem of single-view sculpture reconstruction. Our adaptation method makes use of features from multiple layers of PIFu, combining and processing meaningful information from different levels. With multi-layer features, our adaptation explores the underlying structure of latent feature space and leverages pseudo labels to gradually obtain correct predictions on the target domain.

Experiments have shown that our domain adaptation method can be successfully applied into PIFu. Compared with other methods without 3D ground truth, our method provides better performance measured by P2S and CD. Compared with supervised methods, our method generalizes better and our training does not require 3D ground truth. We find out that supervised methods are suffering with limited datasets. They cannot handle diverse surfaces. Our ablation studies validate the necessity of our proposed modules. Finally, we invited participants for a user study.




Figure 13: One challenging case in experiment. The network is not able to obtain the correct shapes for overweight people.

In the future, we will focus on three potential directions. One potential research direction is to apply our method on artistic human paintings. Considering our adaptation network does not require labels, it is highly feasible to operate on artworks where only a 2D image is accessible and 3D ground truth is not available at all. We intend to be able to handle more highly abstract human shapes such as Picasso paintings. It potentially is difficult for our adaptation method to find an appropriate mapping relationship. The other direction is treating challenging cases as shown in Figure 13. This may be because knowledge transfer is largely influenced by the pre-trained models. We will explore deeper structures of sculpture features to minimize such influence in the future. Our domain adaptation has potential to be applied into other pretrained implicit-function models with point-features and no extra inputs. Generally, our method can be applied to a pretrained explicit reconstruction model. We will leave such attempts for future work.

REFERENCES

- [1] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. 2019. Tex2shape: Detailed full human body geometry from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2293–2303.
- [2] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. 2022. Photorealistic Monocular 3D Reconstruction of Humans Wearing Clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [3] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. 2020. Combining Implicit Function Learning and Parametric Models for 3D Human Reconstruction. In *European Conference on Computer Vision (ECCV)*. Springer.
- [4] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. 2017. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3722–3731.
- [5] Lu Chen, Sida Peng, and Xiaowei Zhou. 2021. Towards efficient and photorealistic 3d human reconstruction: a brief survey. *Visual Informatics* 5, 4 (2021), 11–19.
- [6] Shuaijun Chen, Xu Jia, Jianzhong He, Yongjie Shi, and Jianzhuang Liu. 2021. Semi-supervised domain adaptation based on dual-level domain mixing for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11018–11027.
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.

- [8] Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. 2021. SMPLicit: Topology-Aware Generative Model for Clothed People. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 11875–11885.
- [9] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research* 17, 1 (2016), 2096–2030.
- [10] Zishu Gao, En Li, Guodong Yang, Zhe Wang, Yunong Tian, Zize Liang, Rui Guo, and Shengchuan Li. 2019. Object reconstruction with deep learning: A survey. In *2019 IEEE 9th Annual International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*. IEEE, 643–648.
- [11] Eric Granger, Madhu Kiran, Jose Dolz, Louis-Antoine Blais-Morin, et al. 2020. Joint progressive knowledge distillation and unsupervised domain adaptation. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [12] Xian-Feng Han, Hamid Laga, and Mohammed Bannamoun. 2019. Image-based 3D object reconstruction: State-of-the-art and trends in the deep learning era. *IEEE transactions on pattern analysis and machine intelligence* 43, 5 (2019), 1578–1604.
- [13] Yang Hong, Juyong Zhang, Boyi Jiang, Yudong Guo, Ligang Liu, and Hujun Bao. 2021. Stereopifu: Depth aware clothed human digitization via stereo vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 535–545.
- [14] Jiongchao Jin, Akshay Gadi Patil, Zhang Xiong, and Hao Zhang. 2020. DR-KFS: A differentiable visual similarity metric for 3d shape reconstruction. In *European Conference on Computer Vision*. Springer, 295–311.
- [15] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. 2019. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4893–4902.
- [16] Berk Kaya and Radu Timofte. 2020. Self-supervised 2D image to 3D shape translation with disentangled representations. In *2020 International Conference on 3D Vision (3DV)*. IEEE, 1039–1048.
- [17] Jia Liu, Wenjie Xuan, Yuhang Gan, Juhua Liu, and Bo Du. 2022. An End-to-end Supervised Domain Adaptation Framework for Cross-Domain Change Detection. *arXiv preprint arXiv:2204.00154* (2022).
- [18] Shaohui Liu, Yinda Zhang, Songyu Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. 2020. DIST: Rendering Deep Implicit Signed Distance Function With Differentiable Sphere Tracing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [19] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* 34, 6 (2015), 1–16.
- [20] William E Lorensen and Harvey E Cline. 1987. Marching cubes: A high resolution 3D surface construction algorithm. *ACM siggraph computer graphics* 21, 4 (1987), 163–169.
- [21] Krishna Himaja Manamasa. 2020. Domain Adaptation from 3D synthetic images to real images.
- [22] Fabio Maria Carlucci, Lorenzo Porzi, Barbara Caputo, Elisa Ricci, and Samuel Rota Buló. 2017. Autodial: Automatic domain alignment layers. In *Proceedings of the IEEE international conference on computer vision*. 5067–5075.
- [23] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- [24] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. 2021. Unsupervised Domain Adaptation: A Reality Check. *arXiv preprint arXiv:2111.15672* (2021).
- [25] Naoki Nozawa, Hubert PH Shum, Qi Feng, Edmond SL Ho, and Shigeo Morishima. 2022. 3D car shape reconstruction from a contour sketch using GAN and lazy learning. *The Visual Computer* 38, 4 (2022), 1317–1330.
- [26] Hayato Onizuka, Zehra Hayirci, Diego Thomas, Akihiro Sugimoto, Hideaki Uchiyama, and Rin-ichiro Taniguchi. 2020. TetraTSDF: 3D human reconstruction from a single image with a tetrahedral outer shell. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6011–6020.
- [27] Ahmed A A Osman, Timo Bolkart, and Michael J. Black. 2020. STAR: A Sparse Trained Articulated Human Body Regressor. In *European Conference on Computer Vision (ECCV)*. 598–613.
- [28] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. 2019. Expressive Body Capture: 3D Hands, Face, and Body from a Single Image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. 10975–10985.
- [29] Pedro O Pinheiro, Negar Rostamzadeh, and Sungjin Ahn. 2019. Domain-adaptive single-view 3D reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7638–7647.
- [30] Mengshi Qi, Edoardo Remelli, Mathieu Salzmann, and Pascal Fua. 2020. Unsupervised domain adaptation with temporal-consistent self-training for 3D hand-object joint reconstruction. *arXiv preprint arXiv:2012.11260* (2020).
- [31] Javier Romero, Dimitrios Tzionas, and Michael J. Black. 2017. Embodied Hands: Modeling and Capturing Hands and Bodies Together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* 36, 6 (Nov. 2017), 245:1–245:17. <https://doi.org/10.1145/3130800.3130883>
- [32] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. 2017. Asymmetric tri-training for unsupervised domain adaptation. In *International Conference on Machine Learning*. PMLR, 2988–2997.
- [33] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. 2019. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2304–2314.
- [34] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. 2020. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 84–93.
- [35] Yuefan Shen, Yanchao Yang, Youyi Zheng, C Karen Liu, and Leonidas J Guibas. 2022. DCL: Differential Contrastive Learning for Geometry-Aware Depth Synthesis. *IEEE Robotics and Automation Letters* 7, 2 (2022), 4845–4852.
- [36] Yating Tian, Hongwen Zhang, Yebin Liu, and Limin Wang. 2022. Recovering 3d human mesh from monocular images: A survey. *arXiv preprint arXiv:2203.01923* (2022).
- [37] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. 2019. Photo wake-up: 3d character animation from a single photo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5908–5917.
- [38] Cong Yang, Housen Xie, Haihong Tian, and Yuanlong Yu. 2021. Dynamic Domain Adaptation for Single-view 3D Reconstruction. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3563–3570.
- [39] Mingyue Yang, Yuxin Wen, Weikai Chen, Yongwei Chen, and Kui Jia. 2021. Deep Optimized Priors for 3D Shape Modeling and Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 3269–3278.
- [40] Youshan Zhang. 2021. A survey of unsupervised domain adaptation for visual recognition. *arXiv preprint arXiv:2112.06745* (2021).
- [41] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. 2021. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE transactions on pattern analysis and machine intelligence* 44, 6 (2021), 3170–3184.
- [42] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. 2019. Deephuman: 3d human reconstruction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7739–7749.
- [43] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*. 2223–2232.