

This CVPR 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

3DCaricShop: A Dataset and A Baseline Method for Single-view 3D Caricature Face Reconstruction

Yuda Qiu¹ Xiaojie Xu¹ Lingteng Qiu¹ Yan Pan¹ Yushuang Wu¹ Weikai Chen² Xiaoguang Han^{1,*} ¹SRIBD, The Chinese University of Hong Kong, Shenzhen[†] ²Tencent Game AI Research Center

Abstract

Caricature is an artistic representation that deliberately exaggerates the distinctive features of a human face to convey humor or sarcasm. However, reconstructing a 3D caricature from a 2D caricature image remains a challenging task, mostly due to the lack of data. We propose to fill this gap by introducing 3DCaricShop, the first large-scale 3D caricature dataset that contains 2000 high-quality diversified 3D caricatures manually crafted by professional artists. 3DCaricShop also provides rich annotations including a paired 2D caricature image, camera parameters and 3D facial landmarks. To demonstrate the advantage of 3DCaricShop, we present a novel baseline approach for single-view 3D caricature reconstruction. To ensure a faithful reconstruction with plausible face deformations, we propose to connect the good ends of the detailrich implicit functions and the parametric mesh representations. In particular, we first register a template mesh to the output of the implicit generator and iteratively project the registration result onto a pre-trained PCA space to resolve artifacts and self-intersections. To deal with the large deformation during non-rigid registration, we propose a novel view-collaborative graph convolution network (VC-GCN) to extract key points from the implicit mesh for accurate alignment. Our method is able to generate highfidelity 3D caricature in a pre-defined mesh topology that is animation-ready. Extensive experiments have been conducted on 3DCaricShop to verify the significance of the database and the effectiveness of the proposed method. We will release 3DCaricShop upon publication.

1. Introduction

A caricature is a vivid art form of depicting persons by abstracting or exaggerating the peculiarities of the facial features. As a way to convey humor or sarcasm, caricatures



Figure 1: Left: the proposed 3DCaricShop, a large-scale repository of 3D caricatures that are manually crafted by professional artists. It's richly annotated with 2D caricature images, camera parameters, and 3D facial landmarks. **Right**: the proposed baseline method that sets the new state of the art in single-view 3D caricature reconstruction.

are widely used in entertainment, social events, electronic games and a variety of artistic creations. While 2D caricatures have gained popularity in comic graphics, there exist many scenarios, including cartoon character creation, game avatar customization, custom-made 3D printing, etc., that the 3D face caricatures remain the mainstream representations. However, creating a high-quality 3D caricature is a labor-intensive and time-consuming task even for a skilled artist. Thereby, generating expressive 3D face caricatures from a minimal input, such as a single image, is a highly-demanding but also challenging task.

Most of the prior works mainly focus on 2D caricature generation [6, 33, 16], while research on reconstructing 3D caricatures from 2D caricature images remains vary rare. Wu et al. [39] propose the first work that creates 3D caricature from 2D caricature images using an optimization based approach. They formulate caricature generation as a problem of deforming the standard 3D face. In particular, they

^{*}Corresponding email: hanxiaoguang@cuhk.edu.cn

[†]Shenzhen Research Institute of Big Data

build a intrinsic deformation space based on the exaggerated morphable models of standard faces [25]. The deformation coefficients are then optimized to reduce the landmark fitting errors. Recently, in their follow-up work [42], they employ CNN to automate the task of 2D facial landmark prediction and deformation regression. However, previous works [32, 18] have shown that the traditional 3D morphable models (3DMM) of normal faces have very limited expressiveness in modeling the intricate facial deformations in reality. Thereby, the deformation space based on a synthetically exaggerated 3DMM, as proposed in [39, 42], is far from sufficient to capture realistic 3D caricatures, which are even more diversified and complex than normal faces.

The key to tackling the above problem is a high-quality 3D caricature dataset created by artists that can provide realistic shape priors for both learning-based and optimizationbased approaches. However, there exist two challenges in constructing such a dataset. First, the 3D models crafted by artists are not topologically consistent, making it infeasible to many downstream applications, including blendshape creation, face animation, 3D landmark localization, etc. Secondly, the manually created meshes are typically not aligned with the corresponding images. While many face reconstruction techniques require an accurate registration, such misalignment makes the dataset inapplicable to projection-based applications such as landmark fitting, texture restoration and manipulation, etc.

In this work, we introduce 3DCaricShop, a large-scale 3D caricature dataset that simultaneously addresses the above issues. First of all, 3DCaricShop contains 2,000 highly diversified and high-quality 3D caricature models manually crafted by professional artists. It is constructed by requesting artists to create 3D caricatures according to 2,000 manually selected caricature images from Web-Caricature [15], that span a wide range of shape exaggerations and texturing styles. Compared to the synthetic datasets [13, 42], 3DCaricShop can provide shape priors for 3D caricatures with much higher fidelity. Secondly, all the 3D models in 3DCaricShop have been re-topologized to a consistent mesh topology that paves the way to a number of future applications, including learning a parametric shape space, batch geometry processing, etc. Thirdly, we provide accurate 3D face landmarks in 3DCaricShop, which facilitates the use of landmark fitting technique that is widely adopted in the state-of-the-art face reconstruction approaches. Last but not least, 3DCaricShop offers a paired 2D caricature image and the camera parameters that are used for mesh alignment. This enables a wide range of techniques, such as differentiable rendering, landmark fitting, etc., that rely on 2D-to-3D consistency.

To further exploit the power of 3DCaricShop, we propose a novel baseline approach to infer 3D caricatures from a single caricature image. While the methods based on deep implicit functions [14, 24] have shown promising capability of modeling objects with arbitrary topologies, it is prone to artifacts and self-intersections when applied to reconstruct 3D caricatures, which typically contain many extreme distortions. Though approaches using parametric mesh model can ensure a generation of plausible 3D face, they struggle to produce realistic faces with accurate geometry. We advocate to connect the good ends of both worlds by transferring the high-fidelity geometry learnt from the implicit reconstruction to a template mesh with a reasonable topology. To enable a faithful transfer, we propose a novel viewcollaborative graph convolution network (VC-GCN) to extract key points from the implicit mesh for accurate mesh alignment. To strike a balance between accuracy and robustness, we iteratively project the registered template mesh onto a pre-trained PCA space using 3DCaricShop to avoid overfitting to outliers. Our approach is able to generate high-quality 3D caricatures in a pre-defined mesh topology that is animation-ready.

We have conducted extensive benchmarking and ablation analysis on the proposed dataset. Experimental results show that the proposed approach trained on 3DCaricShop sets new state of the art on the task of single-view 3D caricature reconstruction from caricature images.

2. Related Work

Single-view Reconstruction Single-view reconstruction (SVR) is a classic task in computer vision. Existing methods could be classified as reconstruction for general objects [11, 36, 24] and for objects in specific categories [5, 3, 43]. It is an ill-posed problem due to the ambiguous nature. In tradition, strong priors are introduced to constrain the space of solutions. Shape-from-Shading (SfS) [26] is a kind of physical based prior on the relation between illumination and shape, which recovers the detailed shape in photos. However, it fails to analyze artists works because of the stylized shading effect. Most recently, with the success of deep learning architectures and the release of large-scale 3D shape datasets such as ShapeNet [7], learning based approaches have achieved great progress, by learning the shape priors directly from the huge datasets. According to the used 3D representations, these methods can be divided into voxel-based [20, 8, 30], point-based [27, 28], mesh-based [36, 23], and implicit-function-based frameworks [21, 31]. Among these methods, PIFu [31], an algorithm based on implicit functions, has been applied on the reconstruction of human body and achieves impressive results. In this paper, we employ PIFu to create the 3D mesh for each single caricature image.

Single-view Face Modeling A closely-related task is photo-realistic face reconstruction. Two mainstream methodologies are developed to handle this problem, i.e.

parametric based [5, 34, 10] and shape-from-shading based [29, 35] methods, and remarkable results have been achieved. However, both methods could not apply on our task directly. Parametric methods suffer from the large diversity of geometry shapes in caricature cases. For SfS algorithms, the underlying physical model could not capture various painting styles of artists.

3D Caricature Generation Following the parametric based methods of normal face reconstruction, researchers further introduce deformation to enlarge the capability of representation [19, 39, 42]. In [19], a semi-supervised manifold regularization method is proposed to learn a regressive model for mapping from 2D real faces to the enlarged training set with 3D caricatures. Wu et al. [39] formulate the 3D caricature generation as a problem of deformation from the standard 3D face. By introducing local deformation gradients, they build an intrinsic deformation representation with the capability of extrapolation. With the deformation representation, they construct an optimization framework to create caricature model guided by the landmark constraint. Following [39], Zhang et al. [42] employ CNN to learn the deformation parameters of the intrinsic deformation representations. However, due to the lack of 3D caricature data, their works are still far from satisfaction.

3D Face datasets 3D face datasets are of great value in face reconstruction tasks. In general, they could be categorized into synthetic and real captured datasets. For normal face, existing 3D datasets, including FaceWareHouse [5] and Facescape [40], are built from scanned 3D data, hence widely used in normal face tasks. They focus on the high accuracy and photo reality of the meshes. However, they could not be applied directly on caricature reconstruction. Researchers [39, 42] tried to perform deformation on real 3D face models to construct synthetic exaggerated data. Although some reasonable results are achieved, they still suffer from the lack of diversity. To tackle this problem, we propose 3DCaricShop, which is the first 3D caricature dataset built by artists, composed pairs of caricature images and meshes. Based on the dataset, 3D caricature shape could be learned in a model-free manner. We further propose a baseline method to reconstruct 3D mesh with uniform topology from single caricature image.

3. Dataset

We construct a dataset which contains 2,000 imagemodel pairs in total. All of the 3D models are annotated with 3D facial landmarks and poses w.r.t images. More details are introduced in the following aspects.

3D Model Collection WebCaricature [15] is the largestto-date dataset of 2D caricatures. It contains around 6,000 caricature images with diverse identities, geometry, and textural styles. We first selected 2,000 images from them, fur-



Figure 2: Sample caricature images and crafted 3D meshes in 3DCaricShop. Images with diverse identities, geometry, and textural styles are collected.

ther making them as diverse as possible. Then we recruited 4 paid expert Zbrush artists to create models according to images. The modeling is required to be matched with the image as much as possible, in projection manner. The contour lines for matching include edges of silhouette, lips, eyes, nose and ears. It takes around 40 minutes for each model on average, and around 40 days are cost in total. Several image-model pairs sampled from our dataset are shown in Fig. 2. Each model consists of 300k \sim 700k vertices.

Meshing Unification To support building parametric space for our 3D caricature dataset, we unify the mesh topology for all models in two steps: 1) We first manually annotate 44 3D landmarks (see details in Fig. 3) for each model; 2) The method of Non-rigid ICP [1, 9] is applied to register a pre-defined template mesh to each model, guided by the 3D landmarks. Due to the inherent difficulty to specify vertices on a 3D mesh, the landmark annotation is performed on 3 rendered views of the 3D shape. As described in [4], these 2D landmarks can be easily transformed into their corresponding 3D positions. The template mesh we use is from FaceWareHouse [5] that consists of 11,551 vertices. This procedure is illustrated by an example in Fig. 3.

Pose Annotation 3D pose estimation from a single image is the premise of our reconstruction method (see Sec. 4.1). It usually requires pose information supervision of the 3D face w.r.t the image. 3DCaricShop also provides accurate pose labels for each mesh that annotated by artists manually.

Analysis of the dataset We quantitatively analyze our dataset by comparing the shape variations with two normal face datasets (FaceWarehouse (FWH) and FaceScape), as well as one synthetic caricature dataset, FaceWarehouse with deformation (Aug. FWH). We measure the shape variation using global and part variance. In particular, the variance is computed between the models and their correspond-



Figure 3: The process of 3D landmark annotation. A raw mesh (a) is rendered from the front, left and right view. Then 2D landmarks (b) are manually annotated for each view image. 3D landmarks are obtained by projecting the mesh into a specific view and searching for the closest point on surface (d). Guiding by corresponding 3D landmarks, the template mesh (e) is deformed into the shape of the raw mesh (d), to generate mesh with correct topology (c).

ing mean shape of each dataset in terms of per-vertex displacement. The results are presented in Table 1. The shape diversity of our dataset is richer than the normal ones. For most of the face regions, 3DCaricShop has larger shape variance than Aug. FWH.

Dataset	Global	Eye	Nose	Mouth	Ear	Cheek	Face
FWH	3.41	0.71	0.61	2.60	4.41	1.43	3.40
FaceScape	2.17	0.36	0.15	2.63	5.57	1.24	2.27
Aug. FWH	5.06	1.98	6.29	2.07	9.38	5.26	5.10
Ours	8.26	4.68	3.04	10.90	9.02	8.27	6.95

Table 1: Shape variance comparisons of 3D face datasets.

4. Methodology

Overview In this section, we introduce the proposed baseline method. Given an input caricature image I, the task is to generate the corresponding 3D mesh M. With the topologically uniform 3D meshes in 3DCaricShop, a straight forward way to tackle the task is to construct a PCA basis using the 3D Morphable Model algorithm [2] to build the caricature face space. However, such a space could not handle the large variation in our data. To capture the diversity of geometry in caricature, we employ Pixel-aligned Implicit Function (PIFu) [31] to generate the 3D shape M_{I} from I. Although the implicit function models the variation in targets, it could not ensure a uniform topology for the predictions. To achieve that, we register a template mesh M_t to M_I using non-rigid registeration (NICP) [1]. Then the output of NICP is projected onto the pre-constructed PCA space, to alleviate deformation artifacts, such as selfintersections. We denote the output of NICP as $\mathbf{M}_{\mathbf{N}}$ and that of PCA as M_P . Considering the large difference between the template and target meshes, a sparse 3D landmark is needed in the stage of NICP. We propose a novel viewcollaborative graph convolution network (VC-GCN) to predict key points k from the implicit mesh, where $k \in \mathbb{R}^{44}$.

4.1. The Baseline Approach

Parametric Modeling Our parametric model space is built with standard 3D Morphable Model (3DMM) [2] algorithm. Given p caricature models with uniform topology and N vertices on each mesh, principal component analysis (PCA) is performed on the shape matrix $\mathbf{S}_M \in \mathbb{R}^{3N \times p}$, which is formed by stacking the 3D coordinates of the $N \times p$ vertices. The generated d eigen-vectors are employed as the shape basis \mathbf{S}_i , i = 1, 2, ..., d, where d is a hyper-parameter. The mean vector $\overline{\mathbf{S}}$ represents the mean shape in the mesh set. With this 3DMM, a novel caricature model \mathbf{S}_N could be represented as follows: $\mathbf{S}_N = \overline{\mathbf{S}} + \sum_{i=1}^d a_i \mathbf{S}_i$, where $\boldsymbol{a} = [a_1 \cdots a_d]^T$ is the vector of shape coefficients.

Implicit Reconstruction To capture the diversity of geometric variation, we adopt Pixel-aligned Implicit Function (PIFu) [31] to reconstruct the underlying 3D shape from images. PIFu performs 3D reconstruction by estimating the occupancy of a dense 3D shape, which determines whether a point in 3D space is inside the model or not. Given a RGB image I, its normal maps from the front view F and back view B are generated to strengthen the local details, by using a pixel2pixel-hd network [37]. Then the implicit binary function $f(\mathbf{X}, \mathbf{I}, \mathbf{F}, \mathbf{B})$ could be written as:

$$f(\mathbf{X}, \mathbf{I}, \mathbf{F}, \mathbf{B}) = \begin{cases} 1, & \text{if } \mathbf{X} \text{ is inside the mesh surface,} \\ 0, & \text{otherwise.} \end{cases}$$
(1)

where \mathbf{X} is a given 3D location in the continuous camera space. This function is modeled by a neural network. The loss function for training is formulated as:

$$\mathcal{L} = \frac{1}{n} \sum_{i=1}^{n} |f(\mathbf{X}_i, \mathbf{I}_i, \mathbf{F}_i, \mathbf{B}_i) - f^*(\mathbf{X}_i)|^2, \qquad (2)$$

where $f^*(\mathbf{X}_i)$ is the ground-truth occupancy.

3D Landmark Detection for Registration The output meshes M_I of the implicit function are not topologically uniform. In order to unify the topology, we adopt non-rigid registration[1] to deform a template M_t into the shape of M_I . As shown in [1], without landmarks the cost function of registration could run into a local minimum, where the template is collapsed onto a point on the target surface. Thus it is important to introduce the 3D landmarks of both meshes to guide the deformation. We design a novel framework to detect 3D landmarks for M_I . In short, we propose



Figure 4: Pipeline of our framework. A detail-rich mesh is generated by implicit reconstruction from 2D caricature and the corresponding normal maps. Based on this mesh, 3D landmarks of the mesh are detected to guide the non-rigid registration, which deforms a template mesh to the target one.

to perform detection on the rendered views of M_I to leverage the effectiveness of image-based CNN techniques. The process would be detailed in 4.2.

Landmark-guided Registration Since the huge difference between M_t and M_I , the deformation is likely to generate artifacts, such as self-intersection. To resolve this problem, we iteratively perform NICP and PCA projection to obtain M_N and M_P . After projection, we obtain a deformed template which is closer to M_I in shape. Fig. 4 illustrates the process of the progressive deformation.

4.2. View-collaborative 3D Landmark Detection

In this section, we discuss more details about how to detect 3D landmarks from M_I , which is the key to supporting the procedure of landmark-guided registration. A straightforward way for this detection is directly applying pointbased CNN (e.g., SparseConv [17]) to estimate landmarkware heatmap on mesh vertices. However, due to the inherent difficulty to conduct CNN on a mesh, this approach tends to produce inaccurate results. We thus propose to perform detection on the rendered views of M_I to leverage the effectiveness of image-based CNN techniques. Coarse locations of the 3D landmarks can be obtained from detected 2D landmarks on those views. More importantly, a stack of View-Collaborative GCN block (VC-GCN) is novelly designed to aggregate and enhance information from multiple views for accurate 3D landmarks locations. As illustrated in Fig. 5, single view graph features (local features) are first extracted for initialization. Then, these local features are enhanced in a progressive manner by continually fusing global information into each view. The final local features are aggregated into the global graph feature for 3D landmark prediction.

Initialization Stage In this part, more details about local feature initialization are introduced. Given 2D images rendered from 3 views {front, left, right} ({f, l, r} for simplicity), we first utilize a 2D landmark detector [12] to estimate the 2D landmarks $\hat{\mathbf{P}}^v \in \mathbb{R}^{k^v \times 2}$, where k^v denotes the key point number under the view $v \in \{f, l, r\}$. Next, 3d land-

marks on the mesh are located using the projection matrix of each view. We use the above landmarks which exist in all local views to build local graphs. After that, to enrich the information of each graph node, we extract features from the feature maps generated by the landmark detector for each node, according to their 2D coordinates. Eventually, the initial local view features $\mathbf{F}_{init}^v \in \mathbb{R}^{k^v \times (C+3)}$ for VC-GCN can be produced by concatenating the 3D landmark locations \mathbf{L}^v with related node features under each view v, where C is the image feature dimension.

View-Collaborative GCN In order to provide global information for each view, we aggregate 3 local features into a global graph feature. Then the global feature is fused into each view to enhance the local view feature. This procedure is performed by a View-Collaborative GCN block. In each VC-GCN block, local features are first sent into several GCN layers for better representations. The layer-wise operation in GCN is defined as:

$$\mathbf{F}_{\ell+1}^{v} = \sigma \left(\mathbf{D}^{v-\frac{1}{2}} \mathbf{A}^{v} \mathbf{D}^{v-\frac{1}{2}} \mathbf{F}_{\ell}^{v} \mathbf{W}_{\ell}^{v} \right)$$
(3)

where \mathbf{A}^v is the adjacency matrix with self-loops, \mathbf{D}^v is its diagonal node degree matrix to normalize \mathbf{A}^v , \mathbf{F}^v_{ℓ} represents the local feature in layer $\ell \in \{0, 1, \dots, \ell'\}$ under the view v, \mathbf{W}_{ℓ} is a trainable parameter matrix for linear projection, and $\sigma(\cdot)$ represents the non-linear activation operation. Then the obtained local features are combined into a global graph feature. For each node in the 3D landmark, its feature can be drawn from the local feature of the corresponding node under one view. Note that for the node that shared in different views, its feature is set as the average of multiple local features. The combined global graph feature is then strengthened through several GCN layers, with same operations as in Equation 3. Hence, the process of feature aggregation now can be formulated as following:

$$\begin{aligned} \mathbf{F}_{\ell'}^{v} &= \operatorname{GCN}^{v}(\mathbf{F}_{0}^{v}), v = f, l, r; \\ \mathbf{F}_{\ell'}^{g} &= \operatorname{GCN}^{g}(f_{\operatorname{comb}}(\mathbf{F}_{\ell'}^{f}, \mathbf{F}_{\ell'}^{l}, \mathbf{F}_{\ell'}^{r})). \end{aligned}$$
(4)

where $f_{\text{comb}}(\cdot)$ denotes the combination operation, $\mathbf{F}_{\ell'}^g$ is the strengthened global features. Note that the input \mathbf{F}_0^v is set



Figure 5: The pipeline of View-collaborative 3D landmark Detection. The initial 3D landmaks are obtained by exploiting the predicted 2D landmarks from 3 rendering views. A novel cascaded VC-GCN blocks is used to fuse the features form each view and sends the aggregated features into a GCN head layer for 3D displacement decoding.

as the initial local view feature \mathbf{F}_{init}^{v} in the first VC-GCN block, and is set as the prior output features in subsequent blocks.

In the second step, strengthened global features are fused into local features of each view in a non-local manner [38], so that global information can guide the model to learn more representative local features. The enhanced local features \mathbf{F}_{enh}^{v} of the view v can be obtained as following, where the non-local fusion operation is denoted as f_{n-loc} :

$$\mathbf{F}_{\text{enh}}^{v} = f_{\text{n-loc}}(\mathbf{F}_{\ell'}^{g}, \mathbf{F}_{\ell'}^{v}), v = f, l, r.$$
(5)

More details about the non-local fusion operation are described in the supplementary materials.

It usually takes numerous glimpses to adjust key points to construct a 3D face, even for an expert artist. Thus, several VC-GCN blocks are stacked to progressively enhance local features. In the connection of two blocks, the enhanced local features \mathbf{F}_{enh}^v from the former block are taken as the input \mathbf{F}_0^v of the later block.

Loss Function Given the enhanced local features from the last VC-GCN block, we combine and strengthen them using GCN layers to obtain the final global graph features. Next, it is multiplied by a GCN head layer to get the 3D landmark estimation $\hat{\mathbf{L}}^g \in \mathbb{R}^{N \times 3}$. The predicted 3D landmarks are supervised by 3D and 2D landmark ground truth simultaneously, which leads to more accurate prediction. We now formulate the loss function for the view-collaborative 3D landmark detection training as following:

$$\mathcal{L} = \sum_{i \in \Omega, v} (\mathcal{L}_{detect}(\hat{\mathbf{P}}_{i}^{v}, \mathbf{P}_{i}^{v}) + \mathcal{L}_{3D}(\hat{\mathbf{L}}_{i}^{g}, \mathbf{L}_{i}^{g}) + \mathcal{L}_{2D}(\mathbf{M}^{v}(\hat{\mathbf{L}}_{i}^{g}), \mathbf{P}_{i}^{v})),$$
(6)

where Ω is the training set, *i* is the subscript indicating each training sample, \mathcal{L}_{detect} denotes the 2D landmark detection

error in the initialization stage, \mathcal{L}_{3D} and \mathcal{L}_{2D} represents the 3D landmark prediction error in 2D and 3D space, respectively, and \mathbf{M}^{v} is the projection matrix to obtain 2D landmarks from 3D landmarks under the view v. Note that all loss terms are in smooth- l_1 form.

5. Experiments

Implementation details The proposed framework is trained on our 3DCaricShop. The dataset is separated into 1,600 for training and 400 for testing. The weights of \mathcal{L}_{detect} , \mathcal{L}_{2D} and \mathcal{L}_{3D} are set to 0.1, 0.8 and 1.0 respectively. To train the network for learning the implicit reconstruction, a RMSProp optimizer is adopted with learning rate 0.001, and the network is pre-trained with the minibatch size 2 for 80 epochs. During the training of 2D landmark detection network, an Adam optimizer is used. The learning rate is set to 24 for 30 epochs. After that, the whole framework is trained in an end-to-end manner with the same strategy as above.

Results Gallery We present some typical results of the proposed framework in Fig. 6. As illustrated, our method is robust to caricature images with diverse textures. It can also recover diversified geometric features, such as the exaggerated nose in the second sample of the first row, and the sharp long chin in the third sample of the second row Fig. 6.

5.1. Comparisons with the State-of-the-arts

We qualitatively and quantitatively compare the results of our method with a variety of state-of-the-art 3D caricature reconstruction approaches on 3DCaricShop testing set, including linear parametric model (3DMM) [35, 2], depth map (DF2Net) [41], deformation representation (AliveCaric-DL) [42], and implicit function (PIFu) [31] based methods.



Figure 6: Results gallery for our framework on 3DCaricShop. The framework has the capability to reconstruct 3D shapes from caricature images with diverse texture and geometry shapes.

Qualitative results In Fig. 7, we visualize some results of caricature reconstruction on images from 3DCaricShop. Among the parametric methods, the nonlinear deformation representation [42] based model outperforms the linear ones on fitting the exaggerated input images, but it is still not precise enough due to its limited expressiveness. Besides, other deep learning based approaches such as DF2Net and PIFu unavoidably yield artifacts like hollows and spikes. However, our method introduces the constraint of PCA parametric space into the deep model, thus can produce highly exaggerated local details upon the foundation of a plausible global shape.

Quantitative Results Considering other methods for comparison only reconstruct the face area, we adopt average point-to-surface Euclidean distance (P2S) as the evaluation metric, which measures the unidirectional distance from the source set to the target set. The average point-to-surface Euclidean distance can be computed as:

$$d_{P2S}(P,S) = \frac{1}{\|P\|} \sum_{p \in P} \min_{p' \in S} \|p - p'\|_2$$
(7)

where P is the vertex set of the reconstructed mesh and S is the corresponding ground truth surface. Besides, due to the mismatch in orientation and scale between the generated meshes and ground truth, before calculation, Procrustes alignment is performed and scaling is estimated based on least square error. As shown in Table 2, our method achieves the smallest P2S on the 3DCaricShop.

5.2. Ablation Studies

In this section, we perform ablation studies on the proposed 3D landmark detection framework and landmark guided registration process. The results show the effectiveness and robustness of our pipeline.

3D landmark detection We analyze five variants of our framework: 1) directly using the 3D landmarks selected from predicted 2D landmarks without subsequent refinement, denoted as 'w/o GCN refinement'; 2) utilizing voxel-based method [22] to estimate 3D heatmaps, denoted as

Methods	P2S	Methods	P2S
3DMM _{human}	0.295	PiFuhead	0.153
3DMM _{cari}	0.104	PiFuface	0.126
DF2Net	0.273	Ourshead	0.065
AliveCaric-DL	0.067	Oursface	0.037

Table 2: Quantitative evaluation on 3DCaricShop. Note that the meshes generated by our method and PiFu contain the entire head area which are the same with ground truth, whereas the other methods in our experiment only recover the frontal face, thus we provide two versions of results(i.e. head and face) on our method and PiFu for a more solid comparison. The results are averaged on test data.

'V2V'; 3) employing a global graph to refine the 3D landmarks from the first setting, without using VC-GCN block, which denoted as 'Global only'; 4) Only using local index to gather local features from global view, rather than using non-local operations for local feature enhancement, denoted as 'w/o G2L'; 5) The basic setting, denoted as Basic.

The metric we evaluate the results is mean per joint position error (MPJPE) which is defined as a Euclidean distance between predicted and ground truth 3D landmarks after root joint alignment. The root joint we define is the top of nose. This metric measures how accurately the root-relative 3D landmark estimation is performed. The quantitative results are listed in Table. 3. It confirm the effectiveness of each components in the design of our 3D landmark detector.

On landmark guided registration We evaluate three kinds of registration process: 1) directly perform NICP without landmarks information; 2) perform landmark-guided NICP without PCA projection; 3) the process used in our method. The visualized results are shown in Fig. 9. As [1] suggested, without landmark information, NICP could not capture the large discrepancy between M_t and M_I . Besides, the deformation without PCA projection is likely to generate meshes with self-intersection. In contrast, our method could obtain meshes with higher quality, and capture enough shape information in M_I . For example, the artifacts on ears in Fig. 9(d) are eliminated, while the shape



(a) 3DMM-Human (b) 3DMM-Cari (c) DF2Net (d) AliveCaric-DL (e) PIFu

(g) Our method

Figure 7: Qualitative results of our method compared with state-of-the-art methods, including (a) 3DMM-Human [35], (b) 3DMM-cari [2], (c) DF2Net [41], (d) AliveCaric-DL [42] and (e) PiFu [31], on 3DCaricShop. By incorporating deep models with parametric space constraint, our method (g) can reconstruct highly exaggerated geometry without distinct artifacts.



Figure 8: Ablation experiments on 3D landmark detection: (a) input; (b) ground truth; (c) pure projection; (d) global only; (e) w/o G2L; (f) ours. Our method could capture the face geometry more accurately.

Methods	MPJPE	Methods	MPJPE
w/o GCN Refinement	0.451	Global only	0.373
V2V [22]	0.407	w/o G2L	0.358
		Basic	0.291

Table 3: Ablation study for 3D landmark detection.



Figure 9: Ablation experiments on registration: (a) input; (b) GT; (c) NICP w/o landmark; (d) NICP w/o PCA projection; (e) ours. A better result is obtained with reasonable topology (e.g., the nose) by using PCA projection.

of nose is more consistent with both the groud truth mesh and the input caricature image.

6. Conclusions

We construct a new dataset and benchmark, called 3DCaricShop, for single-view 3D reconstruction from car-

Methods	P2S
Ours w/o Landmark	0.074
Ours w/o PCA projection	0.076
Ours	0.065

Table 4: Ablation study for landmark-guided registration. The similar data implies the improvement of our method concentrate on the detail structures.

icature images. 3DCaricShop is the largest collection by far of 3D caricature models crafted by professional artists. It consists of 2,000 high-quality and diversified 3D caricatures that are richly labeled with paired 2D caricature image, camera parameters, and 3D facial landmarks. A novel baseline approach is also presented to validate the usefulness of the proposed dataset. It combines the merits of flexible implicit functions and the robust parametric mesh representation. Specifically, we transfer the details from implicit reconstruction to a template mesh with the help of VC-GCN that accurately predicts 3D landmarks for the implicit mesh. Extensive benchmarking on our dataset has been performed including a variety of popular approaches. We found that reconstructing 3D caricature from a single 2D caricature image is a highly challenging task with ample opportunity for improvement. We hope 3DCaricShop and our baseline approach could shred light on future research in this field.

Acknowledge The work was supported in part by the Key Area R&D Program of Guangdong Province with grant No. 2018B030338001, by the National Key R&D Program of China with grant No. 2018YFB1800800, by Shenzhen Outstanding Talents Training Fund, and by Guangdong Research Project No. 2017ZT07X152, the National Natural Science Foundation of China 61902334, Shenzhen Fundamental Research (General Project) JCYJ20190814112007258.

References

- Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid icp algorithms for surface registration. In 2007 IEEE Conference on Computer Vision and Pattern Recognition, pages 1–8. IEEE, 2007.
- [2] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999.
- [3] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016.
- [4] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3d morphable model learnt from 10,000 faces. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 5543– 5552, 2016.
- [5] Chen Cao, Yanlin Weng, Shun Zhou, Yiying Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):413–425, 2013.
- [6] Kaidi Cao, Jing Liao, and Lu Yuan. Carigans: Unpaired photo-to-caricature translation. arXiv preprint arXiv:1811.00222, 2018.
- [7] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012, 2015.
- [8] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pages 628–644. Springer, 2016.
- [9] Hang Dai, Nick Pears, William Smith, and Christian Duncan. Statistical modeling of craniofacial shape and texture. *International Journal of Computer Vision*, 128(2):547–571, 2020.
- [10] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pages 0–0, 2019.
- [11] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 216–224, 2018.
- [12] Xiaojie Guo, Siyuan Li, Jinke Yu, Jiawan Zhang, Jiayi Ma, Lin Ma, Wei Liu, and Haibin Ling. Pfld: A practical facial landmark detector. arXiv preprint arXiv:1902.10859, 2019.
- [13] Xiaoguang Han, Chang Gao, and Yizhou Yu. Deepsketch2face: a deep learning based sketching system for 3d face and caricature modeling. ACM Transactions on graphics (TOG), 36(4):1–12, 2017.

- [14] Zeng Huang, Tianye Li, Weikai Chen, Yajie Zhao, Jun Xing, Chloe LeGendre, Linjie Luo, Chongyang Ma, and Hao Li. Deep volumetric video from very sparse multi-view performance capture. In *Proceedings of the European Conference* on Computer Vision (ECCV), pages 336–354, 2018.
- [15] Jing Huo, Wenbin Li, Yinghuan Shi, Yang Gao, and Hujun Yin. Webcaricature: a benchmark for caricature recognition. arXiv preprint arXiv:1703.03230, 2017.
- [16] Sunnie S. Y. Kim, Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Deformable style transfer. In *European Conference on Computer Vision (ECCV)*, 2020.
- [17] Baoyuan Liu, Min Wang, Hassan Foroosh, Marshall Tappen, and Marianna Pensky. Sparse convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 806–814, 2015.
- [18] Feng Liu, Luan Tran, and Xiaoming Liu. 3d face modeling from diverse raw scan data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9408– 9418, 2019.
- [19] Junfa Liu, Yiqiang Chen, Chunyan Miao, Jinjing Xie, Charles X Ling, Xingyu Gao, and Wen Gao. Semisupervised learning in reconstructed manifold space for 3d caricature generation. In *Computer Graphics Forum*, volume 28, pages 2104–2116. Wiley Online Library, 2009.
- [20] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 922–928. IEEE.
- [21] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4460–4470, 2019.
- [22] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. V2v-posenet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *Proceedings of the IEEE conference on computer* vision and pattern Recognition, pages 5079–5088, 2018.
- [23] Junyi Pan, Xiaoguang Han, Weikai Chen, Jiapeng Tang, and Kui Jia. Deep mesh reconstruction from single rgb images via topology modification networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9964–9973, 2019.
- [24] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 165–174, 2019.
- [25] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In 2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance, pages 296–301. Ieee, 2009.
- [26] Emmanuel Prados and Olivier Faugeras. Shape from shading. In *Handbook of mathematical models in computer vision*, pages 375–388. Springer, 2006.
- [27] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification

and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.

- [28] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Advances in neural information processing systems, pages 5099–5108, 2017.
- [29] Elad Richardson, Matan Sela, Roy Or-El, and Ron Kimmel. Learning detailed face reconstruction from a single image. In proceedings of the IEEE conference on computer vision and pattern recognition, pages 1259–1268, 2017.
- [30] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017.
- [31] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE International Conference* on Computer Vision, pages 2304–2314, 2019.
- [32] Matan Sela, Elad Richardson, and Ron Kimmel. Unrestricted facial geometry reconstruction using image-to-image translation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1576–1585, 2017.
- [33] Yichun Shi, Debayan Deb, and Anil K Jain. Warpgan: Automatic caricature generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10762–10771, 2019.
- [34] Luan Tran and Xiaoming Liu. Nonlinear 3d face morphable model. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 7346–7355, 2018.
- [35] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard Medioni. Extreme 3d face reconstruction: Seeing through occlusions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3935–3944, 2018.
- [36] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–67, 2018.
- [37] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [38] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
- [39] Qianyi Wu, Juyong Zhang, Yu-Kun Lai, Jianmin Zheng, and Jianfei Cai. Alive caricature from 2d to 3d. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 7336–7345, 2018.
- [40] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: a large-scale high quality 3d face dataset and detailed riggable 3d face

prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 601–610, 2020.

- [41] Xiaoxing Zeng, Xiaojiang Peng, and Yu Qiao. Df2net: A dense-fine-finer network for detailed 3d face reconstruction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2315–2324, 2019.
- [42] Juyong Zhang, Hongrui Cai, Yudong Guo, and Zhuang Peng. Landmark detection and 3d face reconstruction for caricature using a nonlinear parametric model. arXiv preprint arXiv:2004.09190, 2020.
- [43] Silvia Zuffi, Angjoo Kanazawa, and Michael J Black. Lions and tigers and bears: Capturing non-rigid, 3d, articulated shape from images. In *Proceedings of the IEEE conference* on Computer Vision and Pattern Recognition, pages 3955– 3963, 2018.