

Geometry Driven Progressive Warping for One-Shot Face Animation

Yatao Zhong
yazhong@microsoft.com

Applied Sciences, Microsoft
Redmond, WA, USA

Faezeh Amjadi
faamja@microsoft.com

Ilya Zharkov
zharkov@microsoft.com

Abstract

Face animation aims at creating photo-realistic portrait videos with animated poses and expressions. A common practice is to generate displacement fields that are used to warp pixels and features from source to target. However, prior attempts often produce sub-optimal displacements. In this work, we present a geometry driven model and propose two geometric patterns as guidance: 3D face rendered displacement maps and posed neural codes. The model can optionally use one of the patterns as guidance for displacement estimation. To model displacements at locations not covered by the face model (e.g., hair), we resort to source image features for contextual information and propose a progressive warping module that alternates between feature warping and displacement estimation at increasing resolutions. We show that the proposed model can synthesize portrait videos with high fidelity and achieve the new state-of-the-art results on the VoxCeleb1 and VoxCeleb2 datasets for both cross identity and same identity reconstruction. Our code is available at https://github.com/yataoz/face_reenact_GDPW.

1 Introduction

Face animation refers to the task of creating photo-realistic portrait videos with animated facial motions. The task starts with a source portrait image and a sequence of driving face poses and expressions. For each frame in the output video, a face animation model generates a new portrait image with the same pose and expression as the driving face while still preserving the source identity and appearance (see examples in Fig. 1(a)). In the early days of research, solutions often rely on purely graphics and geometry [8, 10, 31, 32]. They utilize face geometries such as facial landmarks and reconstructed 3D face models to guide face deformation. However, image synthesis in those early works is non-learning based and often resorts to handcrafted image blending heuristics, hence fails to produce photo-realistic output.

Recently with the rapid advance in image synthesis using deep neural networks, many solutions switch to a learning-based paradigm where they directly operate in the feature space of source images and warp features to achieve facial motion transfer and identity preservation. However, prior works [8, 27, 40] attempt to predict the warping displacements directly

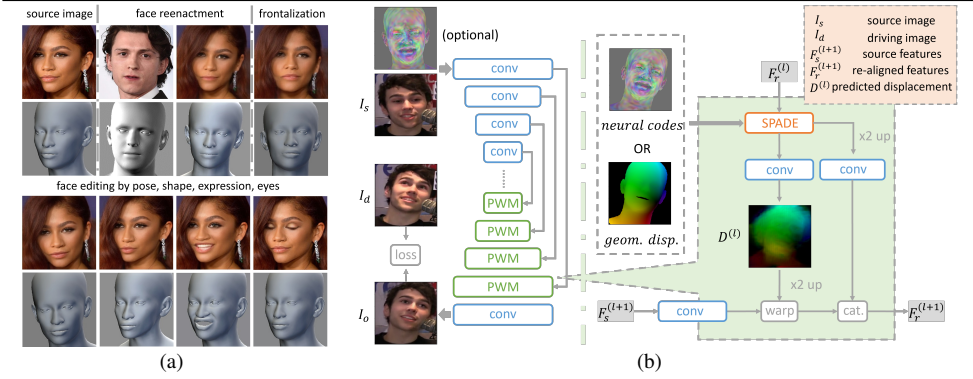


Figure 1: (a) Our model can be applied to several scenarios: face reenactment, face frontalization and fine-grained face editing by pose, shape, expression and even eyes. (b) An overview of the model. The left part shows the whole architecture. The right part shows the proposed progressive warping module (PWM) at pyramid level l and how we estimate the full displacement from the guiding geometric patterns.

from features in source image coordinates whereas the displacements are supposed to be in the target image coordinates. This misalignment issue poses a challenge to them to generate correct displacements.

With those considerations in mind, we ask how we can improve the quality of displacement fields for face animation? We present a geometry driven model and propose two geometric patterns as guidance: 3D face rendered displacement maps and posed neural codes. The model can optionally use one of the patterns as guidance for displacement estimation. To render the displacement maps, we first calculate the displacement vectors between corresponding vertices from a pair of source and driving meshes and then rasterize them to a 2-channel image. To create posed neural codes, we attach a latent embedding vector to each vertex of the 3D face and rasterize the embedded mesh to a d -channel image, where d is the embedding dimension. For choice of the 3D face model, we use FLAME [49], which is a face counterpart of the blend-skinned body model SMPL [27]. Unlike other 3DMMs such as the FaceWarehouse [8] and the Basel Face Model [26], FLAME also models articulated jaw, neck and rotating eyeballs in addition to the disentangled shape, pose and expression parameters, which makes it a more natural representation for face animation. To overcome the spatial misalignment issue present in other 3DMM based models [8, 27, 40], we warp source image features with the predicted displacement field at each pyramid scale and use the warped features to estimate the displacement field at next pyramid scale. We repeat this paradigm and progressively warp features and estimate displacements at increasing resolutions.

To summarize, our work has the following contributions: 1) We propose a geometry driven model whose displacement generation is guided by either displacement maps or posed neural codes, both rendered with the FLAME face topology. 2) We design a progressive warping module which alternates between displacement estimation and feature warping for robust facial motion transfer. 3) We demonstrate that our method sets the new-state-of-art through extensive experiments and provide our insight into the properties of the proposed guiding geometric patterns by analyzing their individual impact on the model performance.

2 Related Work

3D Face Based Warping Models. In early research of face animation, solutions often rely on purely 3D face models. For instance, [6, 31, 32] start with 3D face reconstruction and use parametric face models to control the face movement. However, image synthesis in those early works is non-learning based and often resorts to some heuristic blending techniques, hence fails to produce photo-realistic output. Later come the learning-based solutions. [17, 37] first render synthetic face images with predicted texture and lighting, which are then used as input to deep networks to create images with better realism. PIRender [27] encodes the 3DMM pose parameters into an embedding vector and uses that as modulation in AdaIN layers [14] to estimate the displacement field for face deformation. HeadGAN [8] shares a similar idea but uses a 3D face image as the source of modulation in SPADE layers [25] for displacement generation. [40] creates an approximate flow with the 3D face model for initial warping and the warped image is refined by another generator to create a more realistic image.

Keypoint Based Warping Models. Warp-Guided GANs [11] track the facial landmarks and use them to create a global displacement map for image warping. Subsequently the warped image is fed to a face refinement network to generate final image. Facial landmarks are also used in MarionETte [12] to create a global displacement map, but here warping operates in the feature space instead of image pixels. Another line of work under this category includes an ad hoc keypoint detector in the model, which is trained without direct supervision [28, 29, 35]. The keypoints are implicitly learned by the model and may not convey any interpretable meaning to humans. The displacement between each pair of keypoints controls a local motion transformation. A dense motion network is used to predict the weighting coefficient of each keypoint displacement in order to create a global warping field. MonkeyNet [28] is one of the first works that go in the direction of latent keypoint representation. Follow-up works such as FOMM [29] and 3D-FOMM [35] improve MonkeyNet [28] with first-order motion modeling and extend it from 2D to 3D respectively.

Pose-to-Face Mapping Models. Pose-to-Face refers to a class of models that transform input directly from pose encoding to face images without considering the geometry such as facial landmarks or 3D face models. They share the same spirit of pix2pix [15] and vid2vid [33, 34]. X2Face [36] is one of the first works in this direction. They have an embedding network that encodes the source identity and a driving network that uses the driving pose and embedded source face to synthesize the new face image. Bi-layer model [39] shares the similar idea of embedded face, but they have a two-stage implementation for coarse-to-fine refinement. [38] encodes an image drawn with landmarks of driving face and uses the source image features injected in AdaIN [14] to generate output. Similar to [38, 39], LSR [22] is also driven by a landmark encoded image, but LSR uses a second network trained for semantic segmentation to guide the decoder to generate better quality images. Head2Head [18] proposes a person-specific model that goes directly from driving 3D face to output image without any feature extraction from source identity. Therefore, their model does not generalize.

3 Proposed Method

We start by fitting the FLAME face model [19] to the input images (Sec. 3.1). To embed FLAME in our model and utilize its geometric information to guide facial motion transfer, we use one of the geometric patterns: the FLAME mesh rendered displacement map or the

posed neural codes. Either acts as guidance to estimate the full displacement field used for feature warping (Sec. 3.2). Afterwards we apply a novel progressive warping module and alternate between feature warping and displacement estimation to generate output images. (Sec. 3.3).

3.1 3D Face Representation

To obtain full control over the animated facial motions, we need a latent face descriptor that is compact yet expressive. In this work, we employ FLAME [14] as the underlying face representation due to its capability of modeling articulated joints (jaw, neck and eyeballs) and disentangled parameterization of shape β , pose θ and expression ψ . FLAME deforms face geometry by vertex based linear blend skinning (LBS) with corrective blendshapes. We refer readers to [14] for detailed explanation of the FLAME formulation.

Given a source portrait image I_s that we want to animate, we use DECA [8], an off-the-shelf 3D face reconstruction model, to fit the FLAME face geometry. The output of the fitting pipeline is the shape β_s , pose θ_s and expression ψ_s for the source image I_s , with which we can reconstruct a FLAME mesh M_s . In order to animate the source face with customized facial movement, we can change the values of β_s , θ_s and ψ_s to obtain the target shape β_d , pose θ_d and expression ψ_d . By substituting the updated parameters into the FLAME model, we now have the target 3D face mesh M_d that will drive the movement of source face. In the case of face reenactment where we want to transfer the facial motion from a driving portrait image I_d to the source image I_s , we do an additional FLAME fitting to I_d to generate the driving 3D face mesh M_d .

3.2 Geometric Patterns as Guidance

Our model is driven by facial geometry. We propose two geometric patterns that can guide the displacement learning process. One is a 2-channel image rendered from the displacement vectors between a pair of source and target face meshes, which we call the *geometric displacement field*. The other is the *posed neural codes*, a d -channel image rendered from a face mesh whose vertices are embedded with a set of d dimensional latent vectors. We now describe the formulation of each of the geometric patterns.

Geometric displacement field. Given the source and driving face meshes M_s and M_d , we calculate 2D displacement vectors $V_{d \rightarrow s} = \mathcal{P}(M_s) - \mathcal{P}(M_d)$, where $\mathcal{P}(\cdot)$ projects a mesh to image space. For each vertex i , we take $V_{d \rightarrow s, i}$ as its vertex attribute. Since the mesh M_d is already triangulated, we can follow the conventional graphics pipeline for rasterization. Note that we render the displacement vectors using the topology of driving mesh M_d because our goal is to warp features from the source image coordinates and place them in the target image coordinates. We also do back-face culling to ignore any triangles that are invisible from the viewing direction.

Posed neural codes. Geometric displacement field can assist with the full displacement estimation but it mostly transfers the relative motion from source to target and does not explicitly incorporate facial semantics. To enforce the full displacements conditioned on facial semantics, we propose the posed neural codes.

For a FLAME model with N vertices, we define a set of d dimensional latent vectors $E = \{e_1, e_2, e_3, \dots, e_N\}$ with $e_i \in \mathbb{R}^d$. We attach each vector e_i to vertex i and obtain a mesh embedded in $\mathbb{R}^{N \times d}$. Similar to geometric displacement field, we rasterize a d -channel image from the embedded mesh. Because the mesh topology is predefined and fixed, each vertex

comes with semantic meaningfulness, hence the assigned latent vector too. A visualization of the learned latent codes is available in Fig. 2. As can be seen, the front and back both have symmetric coloring. The left and right possess similar color distribution patterns. They suggest the facial semantics being well captured by the latent codes.

As opposed to geometric displacement field, posed neural codes does not encode relative motion directly. To recover this information, we render the latent codes with the source and target meshes respectively and provide both as input to the model. The source rendered latent codes are concatenated with the source image as input to the encoder. The target rendered latent codes are used as guidance in the decoder (Fig. 1(b)). The rendering process is fully differentiable and the latent codes are learned along with the neural network parameters.




Figure 2: Visualized latent codes on the FLAME manifold. Color encoding is obtained from t-SNE. Similar colors indicate closer distance in the embedding space.

Learning the full displacement field. Although both geometric patterns encode accurate facial geometry and motion transfer information, they do not model any surrounding pixels outside the face or head region (e.g. hair). To address this issue, we propose to use the geometric patterns as guidance to estimate the full displacement field. Imagine that we already have the spatially re-aligned features from a preceding warping module (which we will touch later in Sec. 3.3). Since the input features have been warped to align with the driving pose in target image space, we adopt them for contextual information to fill in any missing pieces that are not modeled by the 3D face. To accomplish this, we use either the rasterized geometric displacement field or the posed neural codes as the modulation input to a SPADE block [25] to guide the estimation of full displacement field (Fig. 1(b)).

3.3 Progressive Warping Module

The goal of estimating displacements is to warp the source features and rearrange them in the target image space. However, many existing works often fail to generate displacement vectors that point to the correct source locations because of the spatial misalignment issue discussed in Sec. 1. We propose a plug-in feature re-sampling block that can warp the source features to create new ones that are spatially aligned with the driving pose in target image coordinates. The re-sampling block is repeatedly applied at increasing pyramid scales to finally create a high-resolution displacement field. We term it the *progressive warping module*, or PWM for short.

Fig. 1(b) shows an overview of the model architecture. Unlike existing methods that employ a number of dedicated deep networks for feature extraction, displacement estimation, image generation and image refinement and operate in a multi-stage pipeline, our model implements a simple encoder-decoder structure. The decoder part adopts a sequence of progressive warping modules for iterative feature warping and re-alignment, which assists full displacement estimation at increasing resolutions.

Let L be the number of feature pyramid levels and l the index of the l -th level. The PWM starts from the lowest resolution at level $l = 1$ in the decoder and gradually moves to the highest resolution at $l = L$. As shown in Fig. 1(b), at each pyramid level l , the guidance map (geometric displacement field or posed neural codes) is resized to match the resolution at

level l and injected into a SPADE block [25] to modulate the re-aligned features $F_r^{(l)}$ from a preceding PWM. The SPADE output is followed by a convolutional layer to predict the full displacement field $D^{(l)}$, which is up-sampled and used to warp the shortcut features $F_s^{(l+1)}$ from encoder. The warped features are concatenated with the other route of the SPADE output to create a set of new spatially re-aligned features $F_r^{(l+1)}$, which will be used by the PWM at next level. A final convolutional layer is added on top of the last PWM to synthesize the output image I_o .

A sequence of L PWMs are employed so that we are able to generate high resolution displacement fields and features to animate portrait images. A special case is the initial PWM at the lowest resolution, where no re-aligned features from a preceding PWM is available. We go around this by skipping the SPADE and displacement estimation components. If the guidance map is the geometric displacement field, we directly adopt it (resized to match the corresponding resolution) to warp the encoder features $F_s^{(1)}$, which are concatenated with the original $F_s^{(1)}$ to produce $F_r^{(1)}$. If the guidance map is posed neural codes, we use $F_s^{(1)}$ to produce $F_r^{(1)}$ without any warping in between.

3.4 Implementation

During training, we perform self-reenactment and same identity reconstruction. For each video, we randomly sample a pair of source and driving frames and train the model with VGG-19 [80] based perceptual loss [46], Patch-GAN [15] and Hinge loss [20] based adversarial loss, and discriminator feature matching loss between the real images and the synthesized images. We use $d = 16$ as the latent code dimension and 5 pyramid levels for feature warping and re-alignment. We enable spectral normalization [23] for all convolutional layers in the model (including discriminator). We optimize the model with ADAM for 270k iterations. The training starts with an initial learning rate of 2×10^{-4} and decreases to 2×10^{-5} at 80k iterations and to 2×10^{-6} at 160k iterations. Implementation details of network architecture and training loss functions can be found in the supplementary material.

4 Experiments

4.1 Datasets and Baselines

We use the VoxCeleb1 dataset [24] for training, which contains more than 20k talking-head videos of over 1000 celebrities. We use the original train/test split provided by the authors to train and evaluate the proposed model. As a pre-processing step, we follow [27, 29] to crop the face and resize it to 256x256. We also fit the FLAME face model offline to all videos using DECA [9]. We also evaluate the VoxCeleb1 trained model on the VoxCeleb2 dataset [8], which contains 5 times more identities than VoxCeleb1. We go through the same steps for data pre-processing and use the test split of VoxCeleb2 for evaluation.

The proposed model has two main variants depending on the geometric pattern in use: geometric displacement field or posed neural codes, denoted by “PWM + geom. disp.” and “PWM + neural codes” respectively. We compare both model variants with the state-of-the-arts including HeadGAN [8], PIRender [27], FOMM [29], 3D-FOMM [65], Bi-Layer [69], LSR [27] and X2Face [66]. For all models except HeadGAN and 3D-FOMM, we adopt their officially released models for evaluation. Since no pre-trained model is available for

HeadGAN, we follow its implementation details in [8] and train the model with FLAME as the underlying 3DMM. For 3D-FOMM, we use an unofficial implementation from [10]. However, as noted by [10], they improve the original 3D-FOMM [35] by adding SPADE blocks [25] in the decoder. Some baselines [22, 27, 36, 39] are originally designed for few-shot scenarios. To match our test configuration, we tested their models in one-shot setting using a single frame without finetuning on test subjects. All models except Bi-Layer [39] are trained on VoxCeleb1, which is trained on the larger VoxCeleb2 dataset. We benchmark all models on both VoxCeleb1 and VoxCeleb2.

Table 1: Quantitative results. Best performing metric is shown in blue

VoxCeleb1										
	Same Identity Reconstruction					Cross Identity Reconstruction				
	FID ↓	CSIM ↑	AKD ↓	AED ↓	APD ↓	FID ↓	CSIM ↑	AED ↓	APD ↓	
X2Face	36.02	0.532	10.94	0.201	5.189	51.14	0.437	0.297	7.421	
Bi-Layer	79.05	0.580	3.11	0.129	0.866	87.54	0.461	0.230	1.493	
LSR	20.53	0.311	2.73	0.153	1.012	28.04	0.229	0.217	1.606	
FOMM	10.87	0.788	2.25	0.087	0.688	30.84	0.562	0.218	1.612	
3D-FOMM	5.47	0.792	2.23	0.091	0.770	22.34	0.605	0.247	2.149	
HeadGAN	10.72	0.779	3.80	0.087	1.097	25.79	0.505	0.234	1.849	
PIRender	9.47	0.745	3.41	0.119	1.148	23.82	0.530	0.222	1.916	
PWM + NMFC	5.57	0.784	2.32	0.091	0.786	15.73	0.647	0.225	1.641	
PWM + geom. disp.	3.95	0.794	2.29	0.081	0.756	14.19	0.653	0.228	1.668	
PWM + neural codes	3.93	0.793	2.15	0.067	0.670	14.32	0.584	0.174	1.214	
PWM + geom. disp. + neural codes	4.33	0.794	2.14	0.066	0.664	15.39	0.590	0.173	1.224	
PWM + geom. disp. + neural codes (BFM)	4.80	0.779	4.04	0.087	1.110	16.78	0.546	0.233	1.956	

VoxCeleb2										
	Same Identity Reconstruction					Cross Identity Reconstruction				
	FID ↓	CSIM ↑	AKD ↓	AED ↓	APD ↓	FID ↓	CSIM ↑	AED ↓	APD ↓	
X2Face	30.34	0.538	21.66	0.188	7.780	41.31	0.379	0.280	9.757	
Bi-Layer	57.26	0.541	2.84	0.138	1.153	66.22	0.436	0.221	1.780	
LSR	12.43	0.275	2.66	0.156	1.388	18.17	0.217	0.214	1.928	
FOMM	8.32	0.718	2.69	0.118	1.291	26.00	0.475	0.231	2.545	
3D-FOMM	3.49	0.723	3.11	0.118	1.498	14.52	0.567	0.240	3.189	
HeadGAN	8.59	0.704	3.87	0.112	1.665	19.47	0.473	0.227	2.647	
PIRender	7.77	0.677	3.47	0.144	1.758	19.15	0.508	0.229	2.852	
PWM + NMFC	4.17	0.725	2.64	0.118	1.162	10.68	0.601	0.230	2.142	
PWM + geom. disp.	3.14	0.732	2.69	0.111	1.203	9.83	0.612	0.234	2.564	
PWM + neural codes	3.29	0.723	2.37	0.087	0.954	10.15	0.528	0.177	1.498	
PWM + geom. disp. + neural codes	3.18	0.726	2.36	0.088	0.954	9.97	0.541	0.178	1.533	
PWM + geom. disp. + neural codes (BFM)	3.63	0.721	3.70	0.109	1.552	11.15	0.516	0.227	2.569	

4.2 Evaluation Metrics

We adopt *Frechet Inception Distance* (FID) [13], *Cosine Similarity* (CSIM) [6], *Average Keypoint Distance* (AKD), *Average Pose Distance* (APD) and *Average Expression Distance* (AED) as our evaluation metrics. FID measures the realism of the generated images and CSIM measures the capability of identity preservation. AKD, APD and AED evaluate the geometry accuracy in terms of facial landmarks, head pose (yaw, pitch & roll) and the expression parameters obtained from FLAME. More details of evaluation metrics can be found in the supplementary material.

For same identity reconstruction, we perform self-reenactment where the source portrait image and the driving portrait image come from the same video clip. During evaluation, we randomly sample pairs of source and driving frames from each video and do both forward and backward reconstruction. This is to ensure that we always cover the difficult cases. For instance, synthesizing a frontal face from a profile face is always harder than the reversed due to severe occlusions in profile face.

For cross identity reconstruction, we perform face reenactment where the source portrait

image and the driving portrait image come from different videos. One advantage of our model is that the facial motion transfer can be accomplished independent of identity due to the disentangled parameterization of shape, pose and expression in FLAME. Therefore, we replace the shape of the driving face with that of the source to preserve the facial geometry. Additionally, because facial landmarks reflect person-specific facial geometry and the driving image has a different identity and appearance from the source image, AKD can no longer be used to measure the accuracy of motion transfer on this task.



Figure 3: Qualitative results. Top two rows are same identity reconstruction and the bottom three rows are cross identity reconstruction. Zoom in for better visual comparison.

4.3 Results

Qualitative results are presented in Fig. 3. We can see that our models generate images with better realism and higher sharpness compared to other baselines. We refer readers to the supplementary material for more qualitative results. Quantitative results are shown in Tab. 1. Our method outperforms prior models on all metrics. X2Face and Bi-Layer have relatively worse numbers on all metrics. This is because they encode motion information in an embedding vector, losing spatial information that could otherwise be useful to guide their model to generate quality images. FOMM has good results on same identity reconstruction, but its performance drops significantly on cross identity reconstruction. This might be FOMM’s motion descriptor dependent on the subject’s appearance. Once the subject’s identity changes, FOMM brings not only facial motion but also the information of the new identity to image synthesis. Compared to FOMM, 3DMM based models such as HeadGAN and PIRender show better results on this task due to 3DMM’s disentangled representation of identity and facial movement. 3D-FOMM outperforms other baselines significantly, which also justifies the inclusion of 3D information for guidance. Our models achieve the best results attributed to the proposed PWM and guiding geometric patterns. We also notice that “PWM + geom. disp.” is better on image quality (FID & CSIM) while “PWM + neural codes” performs better on geometry accuracy (AKD, AED & APD), suggesting their slightly different focus in guidance, which we will analyze individually in Sec. 4.5.

4.4 Face Editing

The disentangled parameterization of FLAME enables us to edit the shape, pose and expression individually to generate customized portrait images. Fig. 4 shows some examples of face editing. Our model generates photo-realistic images with natural poses and expressions. It can even synthesize teeth in the open mouth region (see second to last column of Fig. 4).

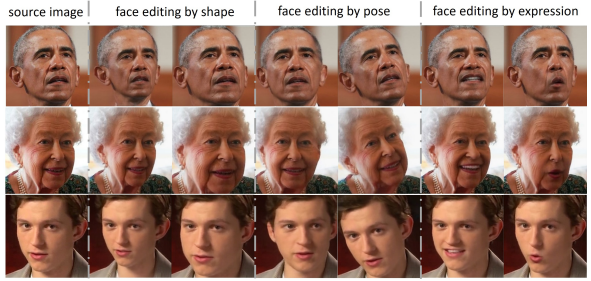


Figure 4: Face editing results.

4.5 Ablation Study

Two major components that have a significant impact on the model performance are the PWM and the guiding geometric patterns. Below we assess their individual influence.

The effectiveness of PWM. We train a model variant similar to “*PWM + neural codes*” but replace the posed neural codes with the normalized mean face coordinates (NMFC). We render a 3-channel image using the normalized coordinates of the mean face obtained from FLAME. The same representation has been used in HeadGAN [9]. We denote this model by “*PWM + NMFC*”. We compare this model variant with HeadGAN, which solely predicts the displacement field at a single pyramid scale and uses that to warp features. As shown in Tab. 1, “*PWM + NMFC*” outperforms HeadGAN by a large margin on all metrics except AED, especially on the image quality metrics FID and CSIM. This shows the importance of feature realignment at each pyramid level and the effectiveness of alternating feature warping and displacement prediction via a sequence of PWMs.

The role of geometric patterns. By comparing “*PWM + geom. disp.*” and “*PWM + neural codes*” in Tab. 1, we find that the geometric displacement field tends to generate images with higher fidelity and is better at preserving identity, but at the sacrifice of geometry accuracy. There are several reasons for this. First, unlike posed neural codes, geometric displacement field does not encode facial geometry and semantics. Instead, it only provides information about **relative** motion transfer, which makes it harder to generate images spatially aligned with the target pose, hence lower geometry accuracy. With the same architecture and complexity, weaker constraint on facial geometry and semantics makes “*PWM + geom. disp.*” focus more on realism during training, resulting in better image quality. Second, a model needs to synthesize more pixels for more aggressive motion change. For example, when the model is used for face frontalization, the more it tries to correct the face pose, the more pixels it needs to create for the originally occluded face region. By doing less motion change (less pixel synthesis), “*PWM + geom. disp.*” naturally obtains better image quality. This explains the trade-off between geometry correctness and image quality.

The above can also be justified by metrics such as L1, SSIM and LPIPS. Because those metrics are calculated pixelwise and averaged by spatial dimensions, they, to certain extent, reflect how accurate the synthesized images are spatially aligned with the ground truth. We compute the metrics for same identity reconstruction where ground truth image of the source identity is available. Results are presented in Tab. 2. The better results of “*PWM + neural codes*” can be attributed to the posed neural codes providing more information about facial geometry and semantics at the corresponding pixel locations.

We also try the combination of geometric displacements and neural codes (via channel concatenation) as guidance, shown as “PWM + geom. disp. + neural codes” in Tab. 1. Compared to models that use geometric displacements and neural codes separately, combining the two does not boost the performance further. Like neural codes guided model, the mixed model tends to favor geometric correctness by trading off realism.

The dependence on 3DMM. We are interested in the dependence of our method on the type of 3DMM in use. Therefore, we add an experiment with the “PWM + geom. disp. + neural codes” configuration but use the Bazel Face Model (BFM) [2, 2, 26] as the underlying 3DMM. As can be seen in Tab. 1, BFM underperforms FLAME but it still outperforms other baselines on most metrics, demonstrating that our method is generic. We also extend AED and APD by fitting a BFM face to compute the expression and pose parameters. Results are shown in Tab. 3. The FLAME based model has better performance even on the BFM computed AED and APD. Compared to BFM that only models the face region, we believe the better quality of FLAME benefits from its modeling of full head and neck.

Table 3: Comparison between different 3DMMs. The subscripts F and B denote a metric computed using FLAME and BFM respectively.

VoxCeleb1								
	Same Identity Reconstruction				Cross Identity Reconstruction			
	AED _F	APD _F	AED _B	APD _B	AED _F	APD _F	AED _B	APD _B
PWM + geom. disp. + neural codes (FLAME)	0.066	0.664	0.088	0.791	0.173	1.224	0.228	1.702
PWM + geom. disp. + neural codes (BFM)	0.087	1.110	0.087	1.245	0.233	1.956	0.230	2.128
VoxCeleb2								
	Same Identity Reconstruction				Cross Identity Reconstruction			
	AED _F	APD _F	AED _B	APD _B	AED _F	APD _F	AED _B	APD _B
PWM + geom. disp. + neural codes (FLAME)	0.088	0.954	0.116	1.366	0.177	1.533	0.241	2.294
PWM + geom. disp. + neural codes (BFM)	0.109	1.552	0.110	1.937	0.227	2.569	0.238	2.885

Further discussion. One limitation is the dependence on the fitting quality of 3DMM. In fact, this is a common problem because almost all models rely on either 3DMM fitting or landmark detection. Isolating and quantifying the impact of 3DMM fitting and landmark detection still remains an open question. However, we have seen [29, 35] made efforts in using self-learned landmarks (which the model deems important but are less interpretable to human) for face warping. Similarly, one future direction could be embedding a self-hosted 3DMM within the model to remove the dependence on external 3DMM fitting.

5 Conclusion

A novel geometry driven model is presented for one-shot face animation. We show that our model outperforms other baselines in both image quality and geometry accuracy. By studying each proposed component, we demonstrate the effectiveness of the progressive warping module (PWM) and find out that each geometric pattern has a different focus in guidance. The geometric displacement field achieves better image quality whereas the posed neural codes favor better geometry correctness.

Table 2: Additionally metrics for same identity reconstruction.

VoxCeleb1			
	L1 ↓	SSIM ↑	LPIPS ↓
PWM+geom. disp.	11.93	0.779	0.106
PWM+neural codes	11.38	0.782	0.103
VoxCeleb2			
	L1 ↓	SSIM ↑	LPIPS ↓
PWM+geom. disp.	13.60	0.743	0.137
PWM+neural codes	13.04	0.747	0.133

References

- [1] Unofficial implementation for one-shot free-view neural talking-head synthesis for video conferencing. URL <https://github.com/zhanglonghao1992/One-Shot-Free-View-Neural-Talking-Head-Synthesis>.
- [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] Chen Cao, Yanlin Weng, Shun Zhou, Yiyi 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.
- [4] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. Voxceleb2: Deep speaker recognition. *arXiv preprint arXiv:1806.05622*, 2018.
- [5] Kevin Dale, Kalyan Sunkavalli, Micah K Johnson, Daniel Vlasic, Wojciech Matusik, and Hanspeter Pfister. Video face replacement. In *Proceedings of the 2011 SIGGRAPH Asia conference*, pages 1–10, 2011.
- [6] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019.
- [7] 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/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [8] Michail Christos Doukas, Stefanos Zafeiriou, and Viktoriia Sharmanska. Headgan: One-shot neural head synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14398–14407, 2021.
- [9] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (TOG)*, 40(4):1–13, 2021.
- [10] Pablo Garrido, Levi Valgaerts, Ole Rehmsen, Thorsten Thormahlen, Patrick Perez, and Christian Theobalt. Automatic face reenactment. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4217–4224, 2014.
- [11] Jiahao Geng, Tianjia Shao, Youyi Zheng, Yanlin Weng, and Kun Zhou. Warp-guided gans for single-photo facial animation. *ACM Transactions on Graphics (TOG)*, 37(6): 1–12, 2018.
- [12] Sungjoo Ha, Martin Kersner, Beomsu Kim, Seokjun Seo, and Dongyoung Kim. Marionette: Few-shot face reenactment preserving identity of unseen targets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10893–10900, 2020.

- [13] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [14] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017.
- [15] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [17] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018.
- [18] Mohammad Rami Koujan, Michail Christos Doukas, Anastasios Roussos, and Stefanos Zafeiriou. Head2head: Video-based neural head synthesis. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 16–23. IEEE, 2020.
- [19] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017.
- [20] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.
- [21] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [22] Moustafa Meshry and et al. Learned spatial representations for few-shot talking-head synthesis. In *ICCV*, 2021.
- [23] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [24] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. Voxceleb: a large-scale speaker identification dataset. *arXiv preprint arXiv:1706.08612*, 2017.
- [25] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2337–2346, 2019.

- [26] 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.
- [27] Yurui Ren, Ge Li, Yuanqi Chen, Thomas H Li, and Shan Liu. Pirenderer: Controllable portrait image generation via semantic neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13759–13768, 2021.
- [28] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2377–2386, 2019.
- [29] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in Neural Information Processing Systems*, 32, 2019.
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [31] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016.
- [32] Justus Thies, Michael Zollhöfer, Christian Theobalt, Marc Stamminger, and Matthias Nießner. Headon: Real-time reenactment of human portrait videos. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018.
- [33] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. *arXiv preprint arXiv:1808.06601*, 2018.
- [34] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. Few-shot video-to-video synthesis. *arXiv preprint arXiv:1910.12713*, 2019.
- [35] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10039–10049, 2021.
- [36] Olivia Wiles, A Koepke, and Andrew Zisserman. X2face: A network for controlling face generation using images, audio, and pose codes. In *Proceedings of the European conference on computer vision (ECCV)*, pages 670–686, 2018.
- [37] Ran Yi, Zipeng Ye, Juyong Zhang, Hujun Bao, and Yong-Jin Liu. Audio-driven talking face video generation with learning-based personalized head pose. *arXiv preprint arXiv:2002.10137*, 2020.
- [38] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9459–9468, 2019.

- [39] Egor Zakharov, Aleksei Ivakhnenko, Aliaksandra Shysheya, and Victor Lempitsky. Fast bi-layer neural synthesis of one-shot realistic head avatars. In *European Conference on Computer Vision*, pages 524–540. Springer, 2020.
- [40] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3661–3670, 2021.