

Self-Supervised Robustifying Guidance for Monocular 3D Face Reconstruction

Hitika Tiwari^{1,4,*}

hitika@iitk.ac.in

Min-Hung Chen^{2,**}

vitec6@gmail.com

Yi-Min Tsai³

yi-min.tsai@mediatek.com

Hsien-Kai Kuo³

hsienkai.kuo@mediatek.com

Hung-Jen Chen³

hj.chen@mediatek.com

Kevin Jou³

kevin.jou@mediatek.com

K. S. Venkatesh¹

venkats@iitk.ac.in

Yong-Sheng Chen⁴

yschen@nycu.edu.tw

¹ Indian Institute of Technology Kanpur, India

² Microsoft AI R&D Center, Taiwan

³ MediaTek Inc., Taiwan

⁴ National Yang Ming Chiao Tung University, Taiwan

Abstract

Despite the recent developments in 3D Face Reconstruction from occluded and noisy face images, the performance is still unsatisfactory. Moreover, most existing methods rely on additional dependencies, posing numerous constraints over the training procedure. Therefore, we propose a *Self-Supervised ROBustifying GUIDanceE (ROGUE)* framework to obtain robustness against occlusions and noise in the face images. The proposed network contains 1) the *Guidance Pipeline* to obtain the 3D face coefficients for the clean faces and 2) the *Robustification Pipeline* to acquire the consistency between the estimated coefficients for occluded or noisy images and the clean counterpart. The proposed image- and feature-level loss functions aid the ROGUE learning process without posing additional dependencies. To facilitate model evaluation, we propose two challenging occlusion face datasets, *ReaChOcc* and *SynChOcc*, containing real-world and synthetic occlusion-based face images for robustness evaluation. Also, a noisy variant of the test dataset of CelebA is produced for evaluation. Our method outperforms the current state-of-the-art method by large margins (e.g., for the perceptual errors, a reduction of **23.8%** for real-world occlusions, **26.4%** for synthetic occlusions, and **22.7%** for noisy images), demonstrating the effectiveness of the proposed approach. The occlusion datasets and the corresponding evaluation code are released publicly at <https://github.com/ArcTrinity9/Datasets-ReaChOcc-and-SynChOcc>.

1 Introduction

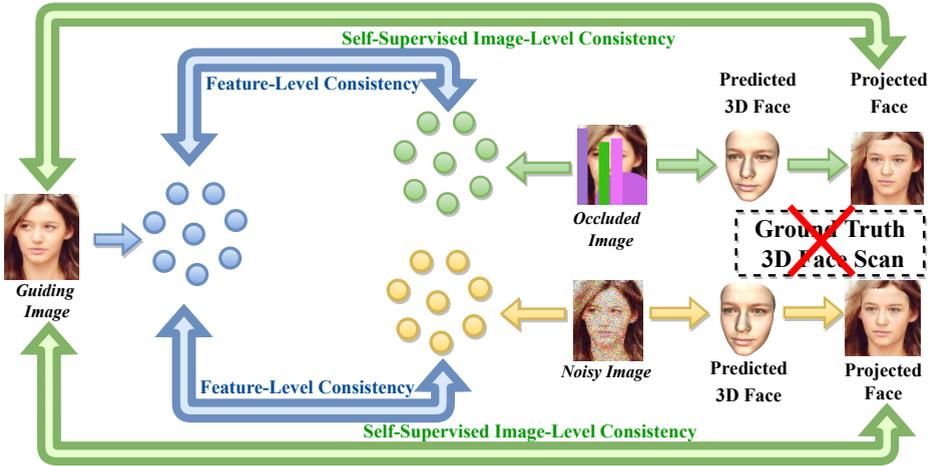


Figure 1: An overview of the proposed *Self-Supervised ROBustifying GUIDancE (ROGUE)* framework. ROGUE addresses the occlusion and noise problems in face images for 3D face reconstruction by the proposed novel image- and feature-level consistency loss functions in the self-supervised fashion, enforcing occluded and noise coefficients to be consistent with the target coefficients of the guiding image, without the requirement of 3D ground-truth face scans and any additional dependency for the training.

3D face reconstruction from monocular face images has been a longstanding problem in the field of 3D computer graphics and computer vision. Recent deep-learning-based approaches demonstrate encouraging progress with regard to perceptual accuracy and training efficiency, facilitating numerous applications such as face recognition [0, 3, 22, 33], face artifice and animation [6, 13, 35]. To address the mathematically ill-posed issue, the fitting-based method, *3D Morphable Model (3DMM)* [2], proposes a low-dimensional 3DMM search space spanning the range of human facial appearance. The coordinates from the two sub-spaces, *geometry* and *texture*, along with the illumination and pose parameters, generate a 3D face such that the corresponding face image (projection of 3D face) resembles the target image. However, most target images contain occlusions such as glasses and masks. Moreover, face images are usually not noise-free. Therefore, the fitting-based methods may drift the coordinates outside the 3DMM space or distort the 3D face geometry and texture, posing challenges to the problem of 3D face reconstruction from monocular images.

To address the above issues, several approaches have been proposed. Fitting-based optimization approaches [9] iteratively adapt the segmentation map to the target face image. 3D faces can also be obtained from occluded face images using training methodologies with different supervisions [7, 13, 28, 31, 36]. In addition, depth-based methods [18, 37] tackle noise issues for 3D face reconstruction with depth maps. However, the above methods hold several dependencies, such as skin masks, depth maps, ground-truth data, synthetic data, segmented maps, multi-images, etc., posing numerous constraints over the training procedure. Therefore, a novel training pipeline that can avoid the above-stated requisites and attain robustness against facial occlusions and image noise is desired. Moreover, there is a need for dedicated occlusion datasets to facilitate the performance evaluation of such models.

In this work, we propose two natural occlusion-based test datasets: **Real World Challenging Occlusion (ReaChOcc)**, and **Synthetic Challenging Occlusion (SynChOcc)** datasets to facilitate robust face reconstruction research, which is not well-explored in the community. Also, we propose a novel **Self-Supervised Robustifying GUIDANCE (ROGUE)** framework, which learns statistical facial coefficients for occluded, and noisy face images simultaneously in a self-supervised manner, without requiring ground truth 3D face scans. The proposed ROGUE contains two parts: 1) The *Guidance Pipeline* estimates coefficients for the clean target face using self-supervised cycle-consistent manners, and 2) the *Robustification Pipeline* enforces the estimated coefficients of occluded and noisy faces to be consistent with clean images. The training is done without additional dependencies due to our image and feature-level losses. The proposed ROGUE framework is evaluated on three datasets: ReaChOcc, SynChOcc, noise variant of the CelebA [24] dataset, and outperforms the current state-of-the-art methods by large margins. For example, for the *perceptual error*, ROGUE achieves a reduction of **23.8%** ($1.237 \rightarrow 0.943$) for real-world occlusions, **26.4%** ($1.195 \rightarrow 0.879$) for synthetic occlusions, and **22.7%** ($1.245 \rightarrow 0.963$) for noisy images. In summary, the contributions of our work are as follows:

1. **ReaChOcc and SynChOcc Testing Datasets:** To facilitate robust face reconstruction research, we propose ReaChOcc and SynChOcc datasets containing natural real-world and synthetic facial occlusions. Our datasets facilitate both *shape* and *texture* comparisons. We have publicly released the datasets and the corresponding evaluation code.
2. **Self-Supervised Robustifying Guidance Framework:** We propose a self-supervised framework with novel image- and feature-level robustification losses, dubbed **ROGUE**, to obtain accurate 3D faces by attaining robustness against the challenging facial *occlusions* and *noise* in the facial images (e.g., 25+% perceptual error reduction), without posing dependencies and the requirement of 3D ground truth.

2 Related Work

Robustness for Face Reconstruction: Egger et al. [9] aim to address the occlusion issues by segmenting the target image into face and non-face regions and iteratively adapting the face model and the segmentation to the target image. Tran et al. [31] deploy an example-based hole filling approach by utilizing the reference set of images containing a suitably similar individual as in the target image. Genova et al. [13] exploit synthetic ground truth data (with the label-free instances of real target image) to tackle the occlusions. Yuan et al. [36] exploit 3DMM to tackle the occlusions in 2D images, where the 3D ground truth data obtained by 3DDFA [38] is required. However, the above methods either only well tackle small-scale occlusions (e.g., minor beards, goggles) instead of large-scale ones (e.g., face masks, tattoos) [9, 9] or rely upon additional dependencies, such as additional images, synthetic data, 3D ground truth, etc. [9, 13, 31, 36]. Besides, our method focuses on tackling large-scale occlusions without posing additional dependencies. Moreover, the noise in the face images poses a challenge in obtaining accurate 3D faces. To our knowledge, there are no 3D face reconstruction methods [9, 9, 8, 9, 12, 13, 19, 27, 29, 30, 31, 32] aiming to reconstruct the 3D faces from the heavily noisy face images. However, there are depth-based methods [18, 37] aiming to address the issues of device-specific noise in obtaining the depth map for reconstructing 3D faces, but tackling the noises in the face images is beyond the scope of those papers. In this paper, the proposed *Self-Supervised Robustifying*

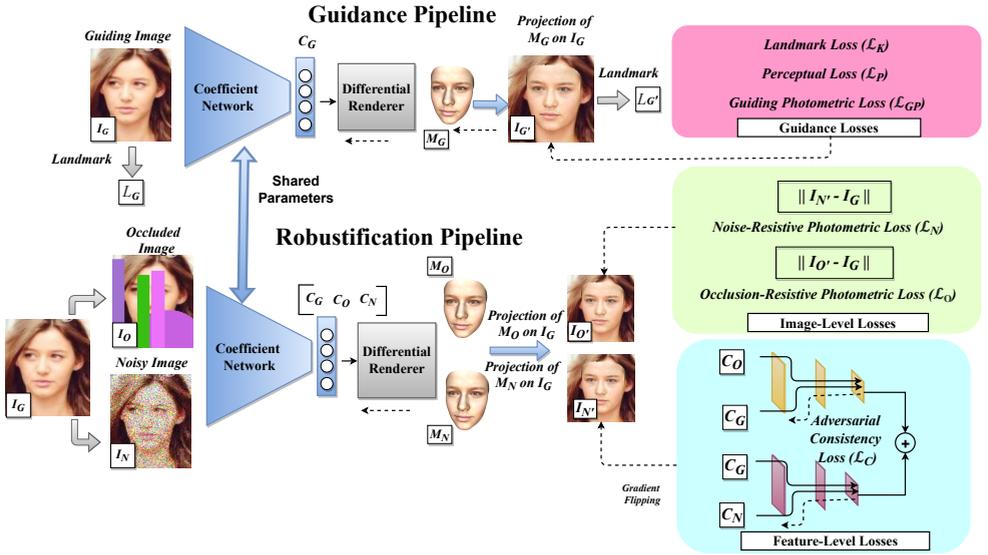


Figure 2: The overall training pipeline of the proposed *Self-Supervised ROBustifying GUIDance (ROGUE)* framework. The *Guidance Pipeline* ensures the faithful reconstruction of the 3D faces from clean *guiding* images I_G in a cycle-consistent manner, and the *Robustification Pipeline* enforces the estimated coefficients of occluded and noisy images (C_O , C_N) to be consistent with *guiding* ones (C_G). The training is done in a self-supervised fashion by the proposed self-supervised image-level losses (\mathcal{L}_O , \mathcal{L}_N) and feature-level adversarial consistency loss (\mathcal{L}_C), without the need for 3D ground truth. Here the solid lines represent the data flow, whereas the dotted lines indicate gradient flow.

Guidance framework aims to attain robustness against the image noise and facial occlusions, thus facilitating the accurate reconstruction of 3D faces from noisy and occluded images.

Occlusion-Aware Datasets: Existing real-world and synthetically occluded test datasets [17, 21, 24, 34] have the following shortcomings: lack of non-occluded ground-truth face images, restrictions on open access, and the limited number of facial occlusions. RealOccWild dataset [34] contains 270 faces with various natural occlusions. The real-world dataset in [17] consists of challenging occlusions (e.g., sunglasses, food, hats, and hands). [21] contains hand-occluded face images. NoW dataset [24] consists of 528 images with common occlusions. However, these datasets do not contain textured 2D/3D ground-truth data, and some occlusions such as bangs, beards, mustaches, turbans, and masks are absent. In addition, the unavailability of test datasets [24] in the open public domain poses constraints over the testing. Besides, our publicly released datasets are designed explicitly for reconstruction tasks and contain various occluded images and the corresponding non-occluded faces.

3 Self-Supervised ROBustifying GUIDance (ROGUE)

Despite the encouraging results obtained by the previous methods for 3D face reconstruction from occluded face images, there is still a large room for improvement with regards

to moderately to heavily occluded face images. In addition, tackling image noise is still an under-addressed issue. Moreover, these methods require several dependencies such as synthetic data, skin masks, etc., posing constraints for training (see Sec. 2 for more details). Therefore, we aim to learn 3D faces in a self-supervised manner without requiring ground truth 3D face scans and other dependencies. To achieve this goal, we propose the **Self-Supervised ROBustifying GUIDance (ROGUE)** framework, which is composed of: 1) the **Guidance Pipeline** and 2) the **Robustification Pipeline** (Fig. 2). For the preliminaries of monocular 3D face reconstruction, please refer to the Supplementary.

Guidance Pipeline: In occlusion robust monocular 3D face reconstruction, one of the main goals is learning reliable 3DMM coefficients with *the least* supervision and dependencies. Inspired by R-Net [10] which contains comparatively fewer dependencies, we propose the *Self-Supervised Guidance Pipeline* to learn the coefficients C_G by exploiting the cycle-consistency in a self-supervised manner, as shown in Fig. 2 (upper). More specifically, the *Guidance Pipeline* takes a clean (i.e., non-occluded noise-free) image I_G (named *guiding image*) as the input, renders the 3D mesh M_G , and projects back to get the 2D face image $I_{G'}$. And then C_G is learned by enforcing the consistency between I_G and $I_{G'}$, using only a single monocular face image. Moreover, C_G guides the Robustification Pipeline to attain robustness against the face occlusions and noise in the images *without relying upon external guidance* such as skin masks [1], synthetic data [13], etc. For more details on various components of the *Guidance Pipeline* please refer to the Supplementary.

Robustification Pipeline: Although the Guidance Pipeline reduces the requirement of supervision and dependencies, the two significant issues for monocular 3D face reconstruction are still not fully addressed: *occlusion* and *noise*. First of all, current methods still cannot reasonably handle the face images with the majority of facial regions occluded, where these methods drift away from their searches from the 3DMM space, resulting in the reconstruction of non-human-like 3D faces. Moreover, additional dependencies such as pre-trained face segmentation models [14], skin masks [1], etc., used by existing methods for tackling the occlusion issues constrain the efficiency of training. Furthermore, despite the progress in the 3D face reconstruction field, no approach has been proposed to tackle the issue of noise in the face image. All the above challenges motivate the need to learn 3D facial coefficients from occluded and noisy face images more accurately and efficiently. Therefore, we propose the *Self-Supervised Robustification Pipeline* to attain robustness against the occlusions and noise in the face images *with the least additional dependencies*, as shown in Fig. 2 (lower). More specifically, we exploit the guiding image I_G and the estimated coefficients C_G from Guidance Pipeline, and encourage the geometry and texture consistency between the Robustification Pipeline and the Guidance Pipeline, to make C_G consistent with the estimated coefficients C_O (from occluded face images I_O) and C_N (from noisy face images I_N). All the components of the Robustification Pipeline are presented as follows:

1) To **obtain consistency with the Guidance Pipeline** for the Robustification (occlusion and noise) coefficients, we exploit a three-layer Generative Adversarial Network (GAN) architecture and propose the *Adversarial Consistency Loss* \mathcal{L}_C as follows:

$$\mathcal{L}_C = \mathcal{L}_{CO} + \mathcal{L}_{CN}, \mathcal{L}_{CO} = \mathcal{L}_h(\mathcal{D}(C_G, C_O), [d_G, d_O]), \mathcal{L}_{CN} = \mathcal{L}_h(\mathcal{D}(C_G, C_N), [d_G, d_N]), \quad (1)$$

where \mathcal{L}_{CO} represents the *occlusion-robustification consistency loss* for tackling the occlusion issues and \mathcal{L}_{CN} denotes *noise-robustification consistency loss* for tackling the noise in the face image. In the equation, \mathcal{D} is the classifier to discriminate C_G and $C_i \in \mathbb{R}^{257}$ ($i = O/N$), and \mathcal{L}_h denotes the standard Huber loss function. In addition, $d_i \in \mathbb{R}$ ($i = G/O/N$)

represents the labels associated with the (guiding/occlusion/noise) coefficients.

2) To **ensure the guidance direction** such that the Robustification Pipeline learns through the experience of the Guidance Pipeline and not vice-versa, we directly regress the pixels of the projected 3D face obtained from the occluded face images ($I_{O'}$) and noisy face images ($I_{N'}$) over the guidance counterpart (I_G) by the proposed *Occlusion-Resistive Photometric Loss* \mathcal{L}_O and *Noise-Resistive Photometric Loss* \mathcal{L}_N , respectively, as follows:

$$\mathcal{L}_O = \|I_{O'} - I_G\|, \quad \mathcal{L}_N = \|I_{N'} - I_G\|. \quad (2)$$

The overall loss function \mathcal{L}_{robust} for the proposed method can be expressed below:

$$\mathcal{L}_{robust} = \beta_O \mathcal{L}_O + \beta_N \mathcal{L}_N - \beta_C \mathcal{L}_C, \quad (3)$$

where β_O, β_N and β_C are the weights associated with occlusion and noise-resistive photometric losses (Eq. (2)), and adversarial consistency loss (Eq. (1)), respectively. The negative sign indicates the *adversarial training*. For simplicity, the notation of the image index is ignored here. It is worth noting that the proposed *Self-Supervised Robustifying Guidance* framework leverages the novel robustification loss function \mathcal{L}_{robust} . Thus our approach bears a significant difference from R-Net [2] regarding the model, architecture, losses, and target data. Unlike R-Net, our model does not require skin masks for the training, facilitating training efficiency. Moreover, our proposed framework is the first (to the best of our knowledge) to tackle the noise in the face images for 3D face reconstruction without 3D ground truth.

4 Dataset Preparation

To obtain the training data, we exploit the training set of several standard face datasets as the clean *guiding* images and create synthetic *occluded* and *noisy* face images for our training pipeline. For testing, numerous real-world and synthetically occluded test datasets [1, 2, 24, 34] have been proposed. However, these datasets have shortcomings such as the unavailability of the dataset in the public domain, lack of non-occluded ground-truth face images, and a limited number of facial occlusions. RealOcc-Wild dataset [34] contains 270 faces with various natural occlusions. The real-world dataset in [1] consists of challenging occlusions (e.g., sunglasses, food, hats, and hands). [2] contains hand-occluded face images. For 3D face reconstruction, there is only one occlusion-based dataset, NoW test set [24], which is not publicly available and thus poses constraints on the testing. Moreover, the datasets mentioned above do not contain several types of facial occlusions such as bangs, beards, mustaches, turbans, and masks. Therefore, a dataset is required which contains numerous possible occlusions and the corresponding non-occluded facial data and should facilitate open research. For achieving the objectives, we propose two datasets: 1) *ReaChOcc* contains real-world challenging facial occlusions such as beards, food items, hands, sunglasses, and 2) *SynChOcc* consists of tough natural occlusions such as mustaches, spectacles. Furthermore, to validate the efficacy of our model against noisy cases, we construct a 3) *Noisy* variant of CelebA-test dataset [2]. Please refer to the Supplementary for more details about the training and testing datasets.

ReaChOcc Dataset: To facilitate occlusion robust 3D face reconstruction model evaluation on challenging real-world data, we introduce a new testing set **Real-World Challenging Occlusion (ReaChOcc)** consisting of 550 face images gathered from various open sources. In our dataset, we have 11 images of each subject in the set of 50 subjects such that 10 images of a subject are occluded, and 1 image is clean. The occluded and clean facial images

are unpaired (captured under different image acquisition environments). These images cover a range of tough facial occlusions, e.g., beards, hands, masks, sunglasses, mustaches, and foods (e.g., Fig. 3 (a)). Moreover, we provide 5 facial landmark coordinates to facilitate cropping and alignment, if needed. However, due to occlusions, 331 occluded face images failed to be detected by dlib [16] to produce landmark coordinates. Therefore, we manually labeled the landmark coordinates of these facial images.

SynChOcc Dataset: We also introduce a novel synthetic occlusion-based test set, **Synthetic Challenging Occlusion (SynChOcc)** dataset, to evaluate the performance of occlusion robust 3D face networks. The dataset contains 550 face images of 50 subjects such that each subject has 10 occluded facial images and 1 non-occluded face. The occluded are generated by overlaying natural occlusions (e.g., turbans, face masks, eye masks, hats, and bangs) on the clean facial images (e.g., Fig. 3 (b)); thus, we have paired data in the proposed dataset. Also, we provide 5 facial landmark coordinates to facilitate cropping and alignment of the face images. These facial landmark coordinates are derived using dlib [16].

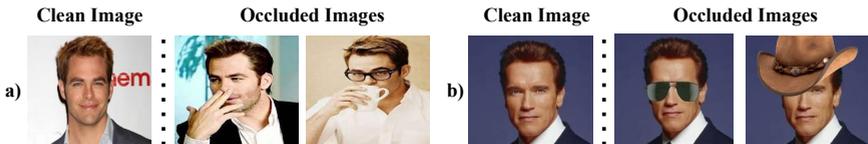


Figure 3: A demonstration of the samples from the proposed a) ReaChOcc, and b) SynChOcc datasets. Our ReaChOcc contains *unpaired clean images*, whereas SynChOcc provides *paired clean images* for comparison.

5 Experiments

In this work, unlike several recent approaches [11, 39], we aim to recover 3D face shape and texture simultaneously from occluded and noisy monocular face images without posing additional requirements. To achieve this goal, we propose two datasets designed for this problem (Sec. 4) since there is no publicly available one. For more dataset and implementation details, please refer to the Supplementary.

5.1 Evaluation Metrics

To evaluate the model performance, we use a standard evaluation metric: *perceptual error metric*, which aims at deriving the mean Euclidean L2 Distance between the feature vectors obtained from various face recognition models. The primary focus of the error metric is on obtaining the visual discrepancy between rendered 3D face and the corresponding 2D face image. We exploit a total of 2 high-performing face recognition models in the main paper: FaceNet-512 [25], and ArcFace [9]. We detail an algorithm to outline the perceptual error metric-based evaluation procedure on the proposed ReaChOcc and SynChOcc datasets in Algo. 1. Moreover, we evaluate our approach on the standard NoW [24] validation dataset to validate its effectiveness. We also present the perceptual error results from 5 other popular backbones, the details on the performance of our model on the MICC dataset in the Supplementary.

Algorithm 1 Evaluation on ReaChOcc and SynChOcc Datasets

Require: Real-World Occluded Face Dataset: $\Psi_R \in \mathbb{R}^{50 \times 10}$, Synthetically Occluded Face Dataset: $\Psi_S \in \mathbb{R}^{50 \times 10}$,
Clean Face Dataset: $\Upsilon \in \mathbb{R}^{10}$, Projection Function: ζ , Perceptual Network: ν , ROGUE: λ

Ensure: Perceptual Dissimilarities: $\mathcal{L}_R, \mathcal{L}_S$

```

while  $i \leq 50$  do
   $I_G \leftarrow \Upsilon[i]$ ;
  while  $j \leq 10$  do
     $I_{O^R} \leftarrow \Psi_R[i][j]$ ;
     $I_{O^S} \leftarrow \Psi_S[i][j]$ ;
     $C_{G_S}, C_{G_T}, C_{G_E}, C_{G_I}, C_{G_P} \leftarrow \lambda(I_G)$ ;
     $C_{O_S^R}, C_{O_T^R}, C_{O_E^R}, C_{O_I^R}, C_{O_P^R} \leftarrow \lambda(I_{O^R}^R)$ ;
     $C_{O_S^S}, C_{O_T^S}, C_{O_E^S}, C_{O_I^S}, C_{O_P^S} \leftarrow \lambda(I_{O^S}^S)$ ;
    Update  $C_{O_E} \leftarrow C_{G_E}, C_{O_I} \leftarrow C_{G_I}, C_{O_P} \leftarrow C_{G_P}$ ;
     $I_{O^{R'}} \leftarrow \zeta(C_{O_S^R}, C_{O_T^R}, C_{O_E^R}, C_{O_I^R}, C_{O_P^R}, I_G)$ ;
     $I_{O^{S'}} \leftarrow \zeta(C_{O_S^S}, C_{O_T^S}, C_{O_E^S}, C_{O_I^S}, C_{O_P^S}, I_G)$ ;
     $\mathcal{L}_R \leftarrow \nu(I_G, I_{O^{R'}})$ ;
     $\mathcal{L}_S \leftarrow \nu(I_G, I_{O^{S'}})$ ;
     $j \leftarrow j + 1$ 
  end while
   $i \leftarrow i + 1$ 
end while

```

5.2 Experimental Results

Qualitative Evaluation: We show the qualitative efficacy of our method on: 1) the *ReaChOcc* set, 2) the *SynChOcc* set, and 3) the *noisy face* set. For this purpose, we compare our results with the several latest state-of-the-art methods. 3DMM [23], Flow [14], 3DDFA [58], Sela et al. [26], Tran et al. [31], and MICA [59] are the occlusion robust 3D face shape reconstruction methods. MoFA [28], R-Net [7], and DECA [10] proposed to reconstruct 3D face shape and texture simultaneously from occluded monocular face images. Motivated by this, we break our comparisons into two categories: 1) comparison with shape recovery-focused methods (Fig. 4), and 2) comparison with texture (along with shape) recovery-based methods (Fig. 5). In Fig. 4, our method demonstrates better shape recovery from occluded images than *most* SOTA methods. These approaches focus on recovering shape, whereas texture estimation is beyond the scope of these methods. Moreover, Fig. 5 shows that our reconstructed 3D faces are visually closer to clean images compared to DECA, R-Net, and MoFA. Note that DECA aims at wrapping the input images to the recovered 3D face shapes by estimating the UV texture maps, thus reproducing occlusions on the 3D faces. Besides, unlike SOTA approaches, our method simultaneously focuses on recovering occlusion robust shape and texture to improve the visual similarity with the non-occluded facial images.

Quantitative Analysis: The state-of-the-art methods [14, 23, 26, 31, 58, 59] focus on reconstructing occlusion-aware 3D face shapes, whereas the issue of robust texture recovery is not addressed by these methods. Therefore, these methods fail to perform well on the perceptual error metric. As a result, we compare our method with MoFA, R-Net, and DECA, which reconstruct both 3D face shape and texture. Our quantitative results (Table 1) show better perceptual similarity for the reconstructed 3D face than these approaches. The proposed method reduces the perceptual error by a large margin of **23.8%** (from 1.237 to 0.943) compared to MoFA on ReaChOcc. In addition, our approach reduces **9.8%** (from 1.045 to 0.943) and **14.1%** (from 1.097 to 0.943) the perceptual errors for R-Net and DECA, respectively. On the SynChOcc dataset, our proposed method shows a large reduction of **26.4%** (from

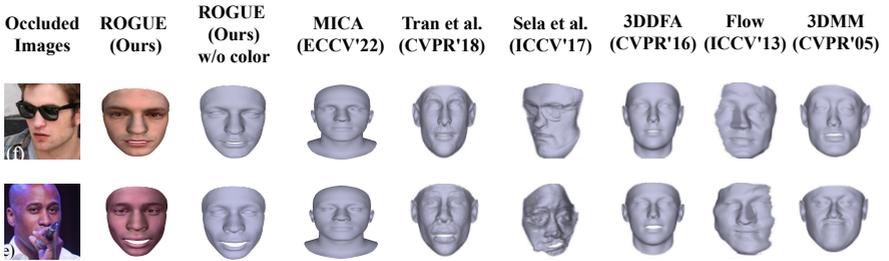


Figure 4: A qualitative comparison of our method with various methods for the case of real-world occlusions. Our results show improved reconstructed 3D faces.

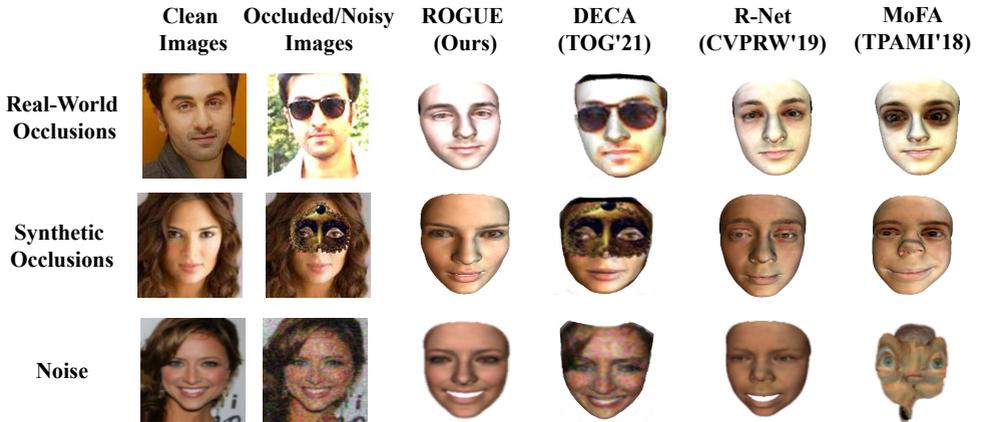


Figure 5: A qualitative comparison of our method with DECA, R-Net and MoFA on the *ReaChOcc*, *SynChOcc* and *noisy* datasets. Our results show a significant improvement in the reconstructed 3D faces. Note that DECA’s meshes are cropped for clear comparison.

1.195 to 0.879) compared to MoFA. In addition, the proposed approach reduces **8.0%** (from 0.955 to 0.879) and **7.6%** (from 0.951 to 0.879) the perceptual errors with regard to R-Net and DECA, respectively. Finally, for the noisy variant, our method reduces the perceptual errors by a large margin of **22.7%** (from 1.245 to 0.963) compared to MoFA. Moreover, our approach reduce **17.1%** (from 1.161 to 0.963) and **17.5%** (from 1.167 to 0.963) of the perceptual errors compared to R-Net and DECA, respectively. All these results demonstrate the efficacy of the proposed approach. It is worth noting that DECA estimates occlusion robust 3D face shape, whereas robust texture estimation is beyond its scope; thus, perceptual error evaluation for DECA (Table 1) is performed only to emphasize the *necessity of occlusion robust 3D texture reconstruction*. We also evaluate our model on the standard NoW [24] validation set. NoW derives the scan-to-mesh distance between the ground truth scan and the predicted meshes. It is worth noting that our approach focuses on producing robust texture and shape simultaneously, but the performance on the shape-specific (i.e., not evaluate texture accuracy) NoW dataset (Table 2) is still comparable to SOTA methods like DECA. **More Discussions:** Due to the page limit, please refer to the Supplementary for 1) the details on the testing and training datasets, 2) implementation details, 3) more comparisons with

Methods	ReaChOcc		SynChOcc		Noise	
	FaceNet-512	ArcFace	FaceNet-512	ArcFace	FaceNet-512	ArcFace
MoFA (TPAMI' 18)	1.237 ± 0.141	1.313 ± 0.114	1.195 ± 0.126	1.284 ± 0.150	1.245 ± 0.171	1.250 ± 0.274
R-Net (CVPRW' 19)	1.045 ± 0.173	1.188 ± 0.171	0.955 ± 0.187	1.131 ± 0.194	1.161 ± 0.253	1.221 ± 0.217
DECA (TOG' 21)	1.097 ± 0.176	1.196 ± 0.176	0.951 ± 0.184	1.061 ± 0.210	1.167 ± 0.295	1.170 ± 0.298
ROGUE (Ours)	0.943 ± 0.187	1.025 ± 0.168	0.879 ± 0.174	0.983 ± 0.186	0.963 ± 0.185	1.017 ± 0.146

Table 1: A quantitative comparison of the perceptual distance using the mean euclidean L2 distance metric with other approaches on the proposed ReaChOcc, SynChOcc and noisy datasets, where the error numbers are the lower, the better.

NoW Evaluation (Non-Metrical)

Methods	median	mean	std
MoFA (TPAMI' 18)	1.547	2.228	2.567
R-Net (CVPRW' 19)	1.505	2.133	2.485
ROGUE (Ours)	1.408	1.978	2.221
DECA (TOG' 21)	1.308	1.635	1.407

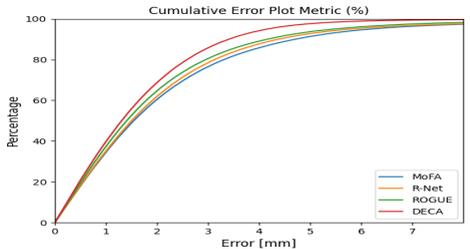


Table 2: (Left) A quantitative evaluation on the NoW validation dataset. (Right) In the plot, the x-axis shows the scan-to-mesh distance error (in mm), whereas the y-axis displays the cumulative percentage such that the higher the curve, the better the shape-based accuracy. It is worth noting that the expressions are not set to be neutral during evaluation.

other methods, 4) ablation studies for the cases of occlusions and noisy images, and 5) the discussions of potential negative societal impact and limitations.

6 Conclusions and Future Work

In this work, we presented two occluded face datasets, **ReaChOcc** and **SynChOcc**, containing various challenging real-world and synthetic occlusion-based face images for robustness tests. Moreover, we proposed a novel *Self-Supervised ROBustifying GUIDanceE (ROGUE)* framework to address the problem of occlusions and noise in the face image for monocular 3D face reconstruction in a self-supervised manner. More specifically, we trained the **Guidance Pipeline** to guide the **Robustification Pipeline** to see through occlusions (e.g., irrespective of the occlusion colors, shapes, and spatial locations) and noise in the face image. Our experiments showed that our model outperforms the current state-of-the-art methods by large margins (e.g., a reduction of **23.8%** for real-world occlusions, **26.4%** for synthetic occlusions, and **22.7%** for the noise in the face images). For future work, we aim at even fewer training dependencies. For example, we plan to waive the requirement of the Guidance Pipeline by empowering the Robustification Pipeline to self-estimate the probable non-occluded 3D faces that enable the model to gain robustness against the occlusions.

Acknowledgements

This work was supported in part by the Higher Education Sprout Project of the National Yang Ming Chiao Tung University, Ministry of Education, and Ministry of Science and Technology (MOST-110-2634-F-A49 -006), Taiwan.

References

- [1] Insaf Adjabi, Abdeldjalil Ouahabi, Amir Benzaoui, and Abdelmalik Taleb-Ahmed. Past, present, and future of face recognition: A review. *Electronics*, 9(8):1188, 2020.
- [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] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Transactions on pattern analysis and machine intelligence*, 25(9):1063–1074, 2003.
- [4] Anpei Chen, Zhang Chen, Guli Zhang, Kenny Mitchell, and Jingyi Yu. Photo-realistic facial details synthesis from single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9429–9439, 2019.
- [5] 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.
- [6] Qixin Deng, Luming Ma, Aobo Jin, Huikun Bi, Binh Huy Le, and Zhigang Deng. Plausible 3d face wrinkle generation using variational autoencoders. *IEEE Transactions on Visualization & Computer Graphics*, pages 1–1, 2021.
- [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] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5154–5163, 2020.
- [9] Bernhard Egger, Andreas Schneider, Clemens Blumer, Andreas Forster, Sandro Schönborn, and Thomas Vetter. Occlusion-aware 3d morphable face models. In *BMVC*, volume 2, page 4, 2016.
- [10] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. Joint 3d face reconstruction and dense alignment with position map regression network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 534–551, 2018.
- [11] 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.
- [12] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1155–1164, 2019.
- [13] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T Freeman. Unsupervised training for 3d morphable model regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8377–8386, 2018.

- [14] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [15] Zihao Jian and Minshan Xie. Realistic face animation generation from videos. *arXiv preprint arXiv:2103.14984*, 2021.
- [16] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [17] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5549–5558, 2020.
- [18] Peixin Li, Yuru Pei, Yicheng Zhong, Yuke Guo, and Hongbin Zha. Robust 3d face reconstruction from single noisy depth image through semantic consistency. *IET Computer Vision*, 2021.
- [19] Jiangke Lin, Yi Yuan, Tianjia Shao, and Kun Zhou. Towards high-fidelity 3d face reconstruction from in-the-wild images using graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5891–5900, 2020.
- [20] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738. IEEE Computer Society, 2015. ISBN 978-1-4673-8391-2. URL <http://dblp.uni-trier.de/db/conf/iccv/iccv2015.html#LiuLWT15>.
- [21] Behnaz Nojavanasghari, Charles E Hughes, Tadas Baltrušaitis, and Louis-Philippe Morency. Hand2face: Automatic synthesis and recognition of hand over face occlusions. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 209–215. IEEE, 2017.
- [22] Nick Pears and Ajmal Mian. 3d face recognition. In *3D Imaging, Analysis and Applications*, pages 569–630. Springer, 2020.
- [23] Sami Romdhani and Thomas Vetter. Estimating 3d shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 986–993. IEEE, 2005.
- [24] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. Learning to regress 3d face shape and expression from an image without 3d supervision. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [25] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [26] 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.
- [27] Ayush Tewari, Michael Zollhofer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Perez, and Christian Theobalt. Mofa: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1274–1283, 2017.
- [28] Ayush Tewari, Michael Zollhofer, Florian Bernard, Pablo Garrido, Hyeonwoo Kim, Patrick Perez, and Christian Theobalt. High-fidelity monocular face reconstruction based on an unsupervised model-based face autoencoder. *IEEE transactions on pattern analysis and machine intelligence*, 42(2):357–370, 2018.

- [29] Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeonwoo Kim, Patrick Pérez, and Christian Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 Hz. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2549–2559, 2018.
- [30] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6142–6151, 2020.
- [31] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard G Medioni. Extreme 3d face reconstruction: Seeing through occlusions. In *CVPR*, pages 3935–3944, 2018.
- [32] 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.
- [33] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gérard Medioni. Regressing robust and discriminative 3d morphable models with a very deep neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5163–5172, 2017.
- [34] Kenny TR Voo, Liming Jiang, and Chen Change Loy. Delving into high-quality synthetic face occlusion segmentation datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4711–4720, 2022.
- [35] Dan Ye and Chiou-Shann Fuh. 3d morphable face model for face animation. *International Journal of Image and Graphics*, 20(01):2050003, 2020.
- [36] Xiaowei Yuan and In Kyu Park. Face de-occlusion using 3d morphable model and generative adversarial network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10062–10071, 2019.
- [37] Yicheng Zhong, Yuru Pei, Peixin Li, Yuke Guo, Gengyu Ma, Meng Liu, Wei Bai, WenHai Wu, and Hongbin Zha. Face denoising and 3d reconstruction from a single depth image. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 117–124. IEEE, 2020.
- [38] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z Li. Face alignment across large poses: A 3d solution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 146–155, 2016.
- [39] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. *arXiv preprint arXiv:2204.06607*, 2022.