

Makeup Prior Models for 3D Facial Makeup Estimation and Applications

Xingchao Yang^{1,2}, Takafumi Taketomi¹, Yuki Endo², Yoshihiro Kanamori²

¹CyberAgent ²University of Tsukuba

{you.koutyo, taketomi.takafumi}@cyberagent.co.jp, {endo, kanamori}@cs.tsukuba.ac.jp

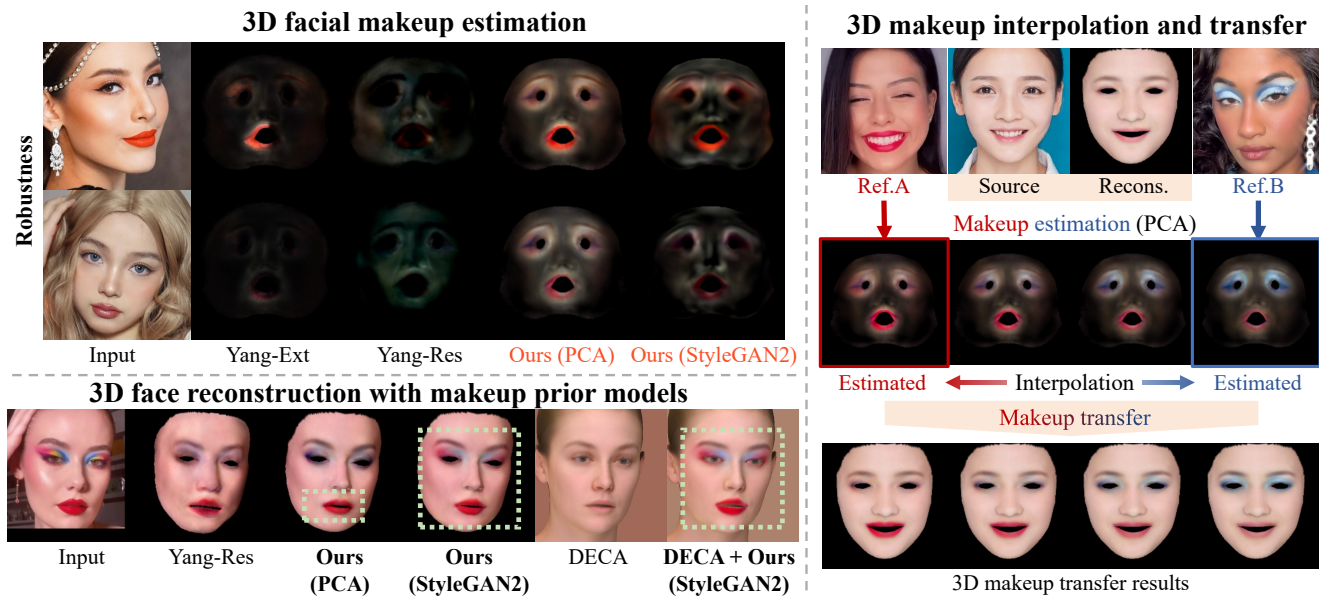


Figure 1. **Example of 3D facial makeup estimation and applications using makeup prior models.** Top left: The effectiveness of our prior models (PCA and StyleGAN2) for estimating 3D facial makeup layers. Compared to the existing method [84] (Yang-Ext and Yang-Res), our method robustly estimates makeup layers, especially in the case of self-occluded faces. Bottom left: The result of 3D face reconstruction using makeup prior models. Our method accurately recovers the makeup of 3D faces and it can be compatible with the existing 3D face reconstruction framework [24]. Right: 3D makeup interpolation and transfer applications using the PCA-based prior model. Note that the StyleGAN2-based prior model has equivalent functionality.

Abstract

In this work, we introduce two types of makeup prior models to extend existing 3D face prior models: PCA-based and StyleGAN2-based priors. The PCA-based prior model is a linear model that is easy to construct and is computationally efficient. However, it retains only low-frequency information. Conversely, the StyleGAN2-based model can represent high-frequency information with relatively higher computational cost than the PCA-based model. Although there is a trade-off between the two models, both are applicable to 3D facial makeup estimation and related applications. By leveraging makeup prior models and designing a makeup consistency module, we effectively address the challenges that previous methods faced in robustly esti-

ating makeup, particularly in the context of handling self-occluded faces. In experiments, we demonstrate that our approach reduces computational costs by several orders of magnitude, achieving speeds up to 180 times faster. In addition, by improving the accuracy of the estimated makeup, we confirm that our methods are highly advantageous for various 3D facial makeup applications such as 3D makeup face reconstruction, user-friendly makeup editing, makeup transfer, and interpolation.

1. Introduction

The first 3D face prior model, known as the 3D Morphable Model (3DMM), was published over two decades ago [5]. In recent years, a diverse array of novel 3DMMs

Table 1. **Comparative analysis of 3D makeup estimation and applications.** This table provides a functionalities comparison between our methods (PCA and StyleGAN2) and the existing techniques [84]. We evaluate various metrics including detail retention, robustness against self-occlusion, processing speed for 3D makeup estimation, and capability in applications such as makeup transfer, interpolation, and user-friendly editing. Our methods demonstrate improved performance and versatility, particularly in challenging scenarios and application potential.

Functions	Makeup Layers			3D Makeup Estimation			Applications		
	Types	Channels	Compositing	Details	Robustness	Speed	Transfer	Interpolation	User-friendly editing
Yang-Ext [84]	Textures	4	Alpha blend	✓		63.13s	✓	Alpha matte	
Yang-Res [84]	Prior	3	Residual			56.12s			
Ours (PCA)	Prior	4	Alpha blend		✓	0.31s	✓	Alpha matte / coefficient	✓
Ours (StyleGAN2)	Prior	4	Alpha blend	✓	✓	18.13s	✓	Alpha matte / coefficient	✓

have been developed [7, 8, 13, 28, 29, 33, 46, 47, 50, 57, 59, 66, 67, 70, 73, 74, 81, 82, 85, 88] and continue to play a crucial role in a wide range of face-related tasks [1, 17, 18, 20, 22, 24, 51, 52, 58, 90]. To further extend and improve the functionality of 3DMMs, several methods are being developed that incorporate additional components such as self-shadowing illumination model and specular albedo [9, 19, 64, 66]. To further enhance the realism of the generated 3D faces, the latest methods employ Generative Adversarial Networks (GANs) [35, 39] within the texture UV maps to construct a non-linear facial appearance prior model [25, 41–43, 46, 54, 62, 65, 78]. However, the extension of 3DMMs to encompass makeup, an important element of human appearance, has not been widely discussed.

Makeup is a key component in creating realistic 3D digital humans in industries such as film, games, and advertising. Previous methods for capturing 3D facial makeup relied on direct measurement techniques that required specialized equipment [34, 44]. To facilitate easier access to makeup data, Yang *et al.* [84] propose a makeup texture extraction method (hereafter referred to as “Yang-Ext”). In addition, they represent the extracted makeup texture as a residual difference from the bare skin texture to construct a PCA-based linear prior model (hereafter referred to as “Yang-Res”). However, as shown in Tab. 1, their methods fall short in computational time, makeup estimation accuracy, and the diversity of applications. More details are explained in the *supplementary material*.

In this paper, we build upon the extracted makeup UV maps from Yang-Ext [84] to construct more sophisticated makeup prior models. One is a statistical linear model that can be efficiently constructed through Principal Component Analysis (PCA) and accompanied by a ResNet [32] architecture for coefficient inference. This statistical approach is capable of capturing the global structure and color patterns of makeup, yet it may fall short in reproducing the nuanced details, such as the gradient of eyeshadows. The other is a generative nonlinear prior model using StyleGAN2 [39].

The StyleGAN2-based prior model requires more computational resources and always involves a carefully designed GAN inversion method to search for the desired coefficient. We incorporate a pSp Encoder [60] and follow it with a minor fine-tuning phase, inspired by [21, 54]. Although resource-intensive, this StyleGAN2-based approach excels at realistically reproducing intricate makeup details in images. The primary advantages of our makeup prior models lie in their reduced computational time and improved robustness in makeup estimation, as summarized in Tab. 1. The improvements are attributed to our use of a regression-based approach, as well as the design of a dedicated makeup consistency module and corresponding regularization loss functions. Furthermore, it demonstrates the versatility of various 3D facial makeup applications, as illustrated in Fig. 1.

In summary, our contributions are:

- Developing two makeup prior models - a PCA-based linear model and a StyleGAN2-based generative model - each with a tailored coefficient estimation network;
- Integrating the makeup consistency module with corresponding regularization loss functions, our specialized network architecture not only boosts robustness in makeup estimation but also minimizes computational time;
- Demonstrating applicability in advanced 3D makeup-related applications such as 3D makeup face reconstruction, user-friendly editing, makeup transfer, and interpolation.

2. Related Work

2.1. Facial Makeup Estimation

3D facial makeup significantly elevates the realism and charm of virtual digital humans. Nevertheless, obtaining 3D facial makeup remains a formidable task [34, 63]. Traditionally, this process demands considerable time and effort from artists, especially in designing shaders and materials

for makeup. Therefore, research that aims to rapidly and accurately acquire 3D facial makeup from the real world proves beneficial.

Inspired by 2D makeup transfer methods [30, 36, 48, 56], Yang *et al.* [84] introduced a pioneering method for 3D facial makeup acquisition: step-by-step texture decomposition-based extraction (Yang-Ext). Additionally, they utilize the extracted makeup textures to build a PCA-based makeup prior model and then apply an analysis-by-synthesis approach to estimate the makeup residual relative to bare skin (Yang-Res). However, the practical application of these makeup estimation methods is still challenging. It requires significant computational time and presents a challenge for self-occluded faces. Moreover, due to Yang-Res’s approach of over-fitting residual makeup, it is applicable primarily to 3D face reconstruction and not suitable for other applications like makeup transfer [10, 11, 15, 30, 36, 44, 48, 53, 55, 56, 77, 79, 80].

We address these limitations by developing effective makeup prior models and introducing corresponding makeup estimation network architectures. Our approach significantly improves the accuracy and reductions in inference time, while also expanding the potential for a wide range of applications.

2.2. Morphable Face Prior Model

The 3D Morphable Model (3DMM) is a statistical prior model that utilizes a dataset of known 3D shapes and textures [5]. Leveraging 3DMM enables a multitude of facial-related applications, such as face recognition, reconstruction, tracking, and animation [20, 90]. To capture the extensive diversity in facial shapes and textures across different ethnicities and ages, numerous enhancements have been made to the original 3DMM concept [7, 8, 13, 28, 46, 47, 50, 57, 59, 66, 74, 81, 82], including developments such as the Basel Face Model (BFM) [28] and FLAME [47].

To enhance the utility and applicability of 3DMMs, researchers have introduced various methods such as facial expression modeling [9] and advanced lighting techniques [19, 64]. The AlbedoMM [66] proposes a specular albedo prior model to better account for specular reflections. Yang *et al.* [84] introduces a residual makeup prior model (Yang-Res). Since traditional statistical linear models mainly capture broad patterns and lack fine details, modern approaches [1, 2, 4, 12, 16, 25, 26, 40–43, 46, 49, 54, 62, 65, 78, 83, 84] utilize deep generative models [35, 39] to recover detailed 3D facial appearance materials. In this paper, we use 3D facial makeup textures extracted from 2D images by Yang-Ext [84] to build two types of makeup prior models, each with its own advantages and limitations. The first is a statistical PCA model, which excels in speed, while the second is an advanced StyleGAN2 generative model, noted for its high-quality output.

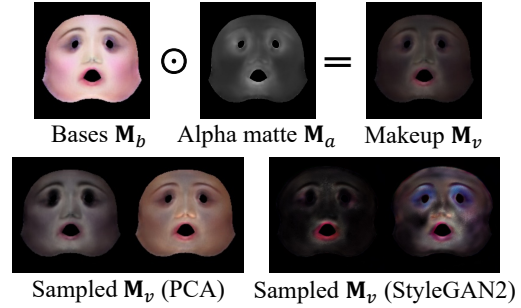


Figure 2. **Examples of makeup sampled from PCA and StyleGAN2 prior models.** The top row demonstrates visual composites M_v from the combination of makeup bases M_b and alpha matte M_a . The bottom row shows variations produced by manipulating the coefficient in the prior models.

Our models show improved performance and functionality compared to the Yang-Res model.

2.3. Model-based 3D Face Reconstruction

Model-based approaches for 3D face reconstruction focus on deriving optimal 3DMM coefficients and rendering the corresponding 3DMM to match an input image [6, 61, 71]. With the advent of deep learning, regression-based methods [3, 14, 17, 22, 24, 27, 31, 45, 68, 69, 72, 75, 87, 89] have gained prominence, offering more stable and quicker reconstruction outcomes, as detailed in various comprehensive surveys [20, 90]. In this paper, our makeup estimation network is based on a 3D face reconstruction task and introduces a makeup consistency module to enhance the quality of estimated makeup. This module facilitates the separation of makeup from facial identity and expressions, making it well-suited for applications such as makeup interpolation and transfer.

3. Approach

We first describe our makeup prior models (Sec. 3.1) and then explain the makeup estimation network (Sec. 3.2). Finally, we introduce an improvement method for the StyleGAN2 makeup prior model (Sec. 3.3).

3.1. Makeup Prior Models

Fig. 2 presents examples of makeup and an alpha matte sampled from the prior models. The makeup textures $M \in \mathbb{R}^D$ ($D = d \times d \times 4$ and d is the UV map size) generated by the prior model (PCA or StyleGAN2) consist of the 4-channel tuple (r, g, b, a) for each pixel in the UV map, where the first three channels are makeup bases $M_b \in \mathbb{R}^{d \times d \times 3}$ and the last channel is the corresponding alpha matte $M_a \in \mathbb{R}^{d \times d}$. The pixel values of M_a range from 0 to 1. We use the visually intuitive representation

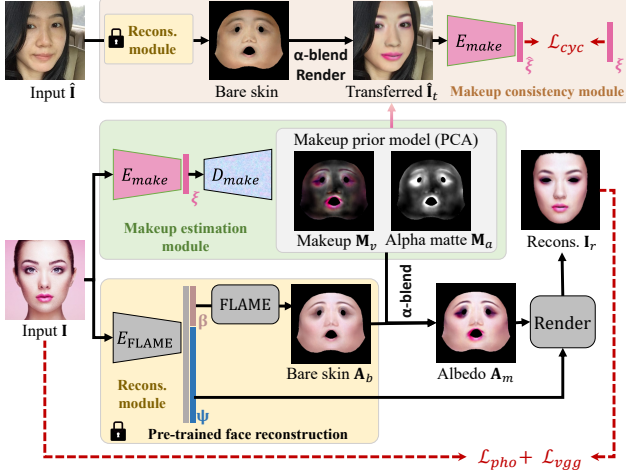


Figure 3. **Overview of the makeup estimation network architecture.** The network is composed of three modules: The Reconstruction module is pre-trained for 3D face reconstruction; The makeup estimation module employs E_{make} to infer the makeup coefficient ξ and generates associated makeup textures; The makeup consistency module enhances the effectiveness of makeup estimation.

$\mathbf{M}_v \in \mathbb{R}^{d \times d \times 3}$ to display makeup, providing a clear depiction of its color and pattern, calculated as follows:

$$\mathbf{M}_v = \mathbf{M}_b \odot \mathbf{M}_a, \quad (1)$$

where \odot denotes the element-wise multiplication.

PCA-based Prior Model. We utilize PCA to construct a linear parametric makeup prior model $\mathbf{M}(v) \in \mathbb{R}^D$ as follows:

$$\mathbf{M} = \bar{\mathbf{M}} + \mathbf{B}_m v, \quad (2)$$

where $\bar{\mathbf{M}} \in \mathbb{R}^D$ is the average facial makeup, $\mathbf{B}_m \in \mathbb{R}^{D \times 100}$ is the PCA bases, and $v \in \mathbb{R}^{100}$ is the corresponding coefficient vector for adjusting 3D facial makeup.

In contrast, Yang-Res [84] constructs its model using \mathbf{M}_v as a makeup residual, resulting in a 3-channel statistical model. Consequently, their model lacks a dedicated channel for alpha matte information.

StyleGAN2-based Prior Model. We train StyleGAN2 [39], with normally distributed noise as input, to construct a generative non-linear makeup prior model with a latent vector $\mathbf{w} \in \mathcal{W}_+$ in the latent space. The disentangled latent space for this StyleGAN2 model is defined as $\mathcal{W}_+ := \mathbb{R}^{7 \times 512}$.

3.2. Makeup Estimation Network

Our network is designed primarily for makeup estimation within the context of 3D face reconstruction. It incorporates a newly developed makeup consistency module and is

trained in a self-supervised learning manner. As depicted in Fig. 3, the network consists of three modules: the Reconstruction module, the makeup estimation module, and the makeup consistency module. The combination of the Reconstruction module and the makeup estimation module forms the core of the 3D face reconstruction functionality, while the makeup consistency module enhances the effectiveness of the makeup estimation process.

Reconstruction module. For the basic 3DMM model that does not include makeup, we use the FLAME model [47], resulting in the estimated albedo \mathbf{A}_b referred to as “bare skin”. We adopt the same pre-trained encoder E_{FLAME} and rendering methods from the existing method [84] to infer FLAME coefficients β for bare skin and ψ for rendering (i.e., geometry, camera parameters, and illumination conditions). We freeze E_{FLAME} during training our entire network.

Makeup estimation module. In the makeup estimation module, the choice of an encoder E_{make} and a decoder D_{make} depends on the employed prior model, that is, PCA or StyleGAN2. The encoder E_{make} has a role for inferring coefficients ξ (v or \mathbf{w}) of our makeup prior model. (1) Specifically, for the **PCA**-based makeup prior model, the encoder E_{make} employs the ResNet50 [32] architecture to infer the coefficients v . The decoder D_{make} follows Eq. (2) to generate makeup textures. (2) For the **StyleGAN2**-based makeup prior model, we use the pSp Encoder [60] as E_{make} to infer the coefficients \mathbf{w} . The decoder uses the pre-trained StyleGAN2 makeup prior model (Sec. 3.1) as D_{make} and removes the noise injection during generation to achieve precise control.

The final albedo \mathbf{A}_m containing makeup is calculated via alpha blending as follows:

$$\begin{aligned} \mathbf{A}_m &= \mathbf{M}_b \odot \mathbf{M}_a + (1 - \mathbf{M}_a) \odot \mathbf{A}_b \\ &= \mathbf{M}_v + (1 - \mathbf{M}_a) \odot \mathbf{A}_b, \end{aligned} \quad (3)$$

where $(1 - \mathbf{M}_a)$ refers to the inversion of \mathbf{M}_a . We use differentiable rendering to reconstruct 3D face \mathbf{I}_r from \mathbf{A}_m and ψ . Subsequently, the encoder E_{make} is trained by minimizing the error between the original image \mathbf{I} and \mathbf{I}_r to ensure the makeup similarity.

Note that the characteristics of the estimated makeup layers as summarized in Tab. 1, the difference between \mathbf{A}_m in Eq. (3) and the makeup-applied albedo in Yang-Res [84] (denoted as \mathbf{A}'_m). In the Yang-Res method, the makeup layer is represented as the residual (denoted as \mathbf{M}_Δ) between the bare skin and the input image. As shown in Fig. 8, it becomes challenging when applied to makeup transfer. \mathbf{A}'_m is calculated as follows:

$$\mathbf{A}'_m = \mathbf{A}_b + \mathbf{M}_\Delta. \quad (4)$$

Makeup consistency module. This module aims to mitigate over-fitting in makeup estimation from input image \mathbf{I} ,

accounting for variations in pose, lighting, and expression. Estimated makeup obtained from the coefficient ξ is applied to randomly selected bare faces $\hat{\mathbf{I}}$. The encoder E_{make} then re-estimates makeup on the transferred image $\hat{\mathbf{I}}_t$ to estimate makeup coefficient $\hat{\xi}$, ensuring that the new coefficient $\hat{\xi}$ consistent with the original coefficient ξ . This module is not required during the inference phase.

Loss function. The final loss function is defined as follows:

$$\begin{aligned} \mathcal{L}(\xi) = & \lambda_{pho} \mathcal{L}_{pho}(\xi) + \lambda_{vgg} \mathcal{L}_{vgg}(\xi) \\ & + \lambda_{cyc} \mathcal{L}_{cyc}(\xi) + \lambda_{reg} \mathcal{L}_{reg}(\xi), \end{aligned} \quad (5)$$

where \mathcal{L}_{pho} represents the L1 loss, and \mathcal{L}_{vgg} represents the perceptual loss [37]. \mathcal{L}_{cyc} represents the MSE loss between makeup coefficients $\hat{\xi}$ and ξ . \mathcal{L}_{reg} is a regularization term, which uses the L2 norm to prevent the inferred parameters from deviating too much from the mean value. The weight coefficients are tuned to achieve a balance between the different losses and are set to $\lambda_{pho} = 100$, $\lambda_{vgg} = 1$, $\lambda_{cyc} = 4$ for PCA, $\lambda_{cyc} = 20$ for StyleGAN2, and $\lambda_{reg} = 1 \times 10^{-4}$.

3.3. Refinement for StyleGAN2 Prior Model

We observe that the StyleGAN2-based approach occasionally yields imprecise results when inferring less common makeup styles, which we attribute to limitations in the training data leading to out-of-domain problems [76]. To overcome this issue, we follow methodologies from [21, 54] to perform minor optimization-based fine-tuning on the estimated coefficient \mathbf{w} . We remove the makeup consistency module and its associated loss function \mathcal{L}_{cyc} from \mathcal{L}_{ξ} , while other loss functions are preserved. Alternatively, we optimize the coefficient using two additional loss functions to constrain the generated makeup textures. The first is a soft symmetry loss \mathcal{L}_{sym} , ensuring bilateral symmetry in \mathbf{M} , and the second is a regulation loss \mathcal{L}'_{reg} for \mathbf{M}_a . This regulation loss acts as a preventive measure against over-fitting that may occur when the proportion of makeup is excessive. The additional loss functions can be formalized as follows:

$$\begin{aligned} \mathcal{L}'(\mathbf{w}) = & \lambda_{pho} \mathcal{L}_{pho}(\mathbf{w}) + \lambda_{vgg} \mathcal{L}_{vgg}(\mathbf{w}) \\ & + \lambda_{reg} \mathcal{L}_{reg}(\mathbf{w}) + \lambda'_{reg} \mathcal{L}'_{reg}(\mathbf{w}) \\ & + \lambda_{sym} \mathcal{L}_{sym}(\mathbf{w}), \end{aligned} \quad (6)$$

where \mathcal{L}_{sym} is the L1 loss and \mathcal{L}'_{reg} regularizes alpha matte using $\|\mathbf{M}_a\|_1$, which represents the L1 norm. The weight parameters are set to $\lambda_{sym} = 8$, and $\lambda'_{reg} = 1$.

In our experiments with the PCA-based makeup prior model, we also attempted refinement through fine-tuning. Unfortunately, we found that minor adjustments failed to yield significant improvements, while extensive perturbation of the PCA coefficients led to the generation of numerous artifacts. We hypothesize the limitations of the PCA model are reached with complex makeup, as its lower degree of freedom in the model offers limited scope for optimization. Therefore, we conduct the refinement process

to the StyleGAN2 model, which demonstrates enhanced results. The ablation comparison results confirm the necessity and effectiveness as shown in Figs. 10 and 11.

4. Implementation Details

Makeup prior models. Our prior models (PCA and StyleGAN2) are both constructed using 3,070 images of makeup bases and alpha matte, which are extracted by Yang-Ext [84] following the UV format of FLAME model [47]. These materials are sourced from makeup datasets [30, 48]. The UV map size is set to $d = 256$, following the configuration of FLAME, and only the facial area is utilized. The training of the StyleGAN2 makeup prior model requires a total duration of ten days using NVIDIA L4 GPU. To ensure stability during training, we retain the noise injection mechanism [23], while it is omitted during the inference phase.

Networks. The makeup estimation network is trained using the same dataset [30, 48] for building the makeup prior models. We set the batch size to 10 for the PCA model and 4 for the StyleGAN2 model. We use the Adam optimizer with the learning rate of 1×10^{-4} . The model is trained for 20 epochs using NVIDIA 2080 Ti GPU. The network training requires one day for the PCA model and two days for the StyleGAN2 model. For the refinement of the StyleGAN2 model, we configure the learning rate to 1×10^{-2} for 40 iterations. The computation times for inference are discussed in Sec. 5.3.

5. Experiments

The summary of functional comparisons is presented in Tab. 1. We utilize makeup datasets Wild [36] and BeautyFace [79] for experiments, which are not employed in the construction of the makeup prior model or in the network training. These datasets comprise a substantial number of images featuring challenging facial poses, expressions, and complex makeup, amounting to 380 and 2,995 images, respectively. We evaluate the makeup prior models, through a diverse set of experiments, such as makeup estimation, 3D face reconstruction, user-friendly makeup editing, makeup transfer, interpolation, and computational time. In the experiments, the red arrows denote results with deficiencies, while green arrows indicate favorable outcomes.

5.1. 3D Facial Makeup Estimation

Before discussing the various 3D facial makeup applications, we first assess the quality of makeup estimation. As shown in Figs. 1 and 4, Yang-Ext [84] faces challenges in estimating makeup on self-occluded faces. Furthermore, the step-by-step image decomposition of their method results in unstable makeup estimation, often missing details such as the color and region of the upper lip, an issue that is also present in the Yang-Res [84]. In contrast, our method




Figure 4. **Comparison with 3D facial makeup estimation methods.** The results demonstrate the robustness of our methods (PCA and StyleGAN2) in terms of stability and accuracy in estimating makeup, outperforming both Yang-Ext and Yang-Res [84], which show limitations in handling self-occluded faces.

Table 2. **Comparative analysis of 3DMM-based 3D face reconstruction on Wild [36] and BeautyFace [79] datasets.** Values in **bold** represent the best results, while those underlined denote the second-best. * denotes the integration of makeup prior models.

Method	Wild					BeautyFace				
	HM(eyes)↓	HM(lips)↓	RMSE↓	SSIM↑	LPIPS↓	HM(eyes)↓	HM(lips)↓	RMSE↓	SSIM↑	LPIPS↓
DECA [24]	0.0048	0.0088	0.0876	0.3651	0.0871	0.0043	0.0106	0.0636	0.4002	0.0840
DECA [24]* (PCA)	<u>0.0045</u>	<u>0.0081</u>	<u>0.0871</u>	0.3683	<u>0.0822</u>	<u>0.0041</u>	<u>0.0081</u>	<u>0.0595</u>	<u>0.4041</u>	<u>0.0779</u>
DECA [24]* (StyleGAN2)	0.0044	0.0076	0.0807	<u>0.3678</u>	0.0798	0.0039	0.0077	0.0568	0.4042	0.0747
Reconstruction module	0.0045	0.0093	0.0667	0.5940	0.0750	0.0037	0.0102	0.0666	0.5034	0.0745
Yang-Res [84]	<u>0.0037</u>	0.0079	<u>0.0543</u>	0.6036	<u>0.0674</u>	<u>0.0032</u>	0.0081	0.0545	<u>0.5065</u>	<u>0.0680</u>
Ours (PCA)	0.0041	<u>0.0078</u>	0.0609	<u>0.6111</u>	0.0681	0.0035	<u>0.0078</u>	0.0690	0.5013	0.0733
w/o Refine (StyleGAN2)	0.0042	0.0083	0.0618	0.6091	0.0685	0.0035	0.0084	0.0695	0.5031	0.0741
Ours (StyleGAN2)	0.0036	0.0073	0.0517	0.6240	0.0608	0.0031	0.0068	<u>0.0650</u>	0.5134	0.0673

consistently estimates the makeup regions and can handle makeup on facial areas not visible in the image. Compared to the results of the PCA prior model, the StyleGAN2 prior model captures a greater level of detail. Our effective makeup estimation facilitates subsequent 3D facial makeup applications.

5.2. Applications

3D face reconstruction. Fig. 5 displays comparative results of model-based 3D facial reconstruction with progressively increasing makeup complexity from left to right. Traditional 3DMM-based methods lack the capability to reconstruct faces with makeup, while the Yang-Res [84] often converges to local minima, resulting in 3D faces with artifacts. In contrast, our methods demonstrate effective reconstruction of facial makeup. The PCA prior model approximately restores the makeup colors, whereas the Style-

GAN2 prior model accurately recovers challenging makeup features, such as blush and gradient eyeshadow.

Our makeup prior models are compatible with other 3DMM-based reconstruction methods (require the same topology), allowing for seamless integration without additional training. Figs. 1 and 6 demonstrate the combined efficacy of our PCA and StyleGAN2 estimated makeup with the DECA [24] framework.

To measure the difference between the reconstructed faces and the original images, we conduct a quantitative evaluation using several metrics, including Histogram Matching (HM), Root Mean Square Error (RMSE), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [86] metrics. The HM metric measures the similarity of the color distributions between two images and is commonly used in makeup analysis and transfer. We perform segmentation around the eyes and lips to compare the HM metric regionally, thereby

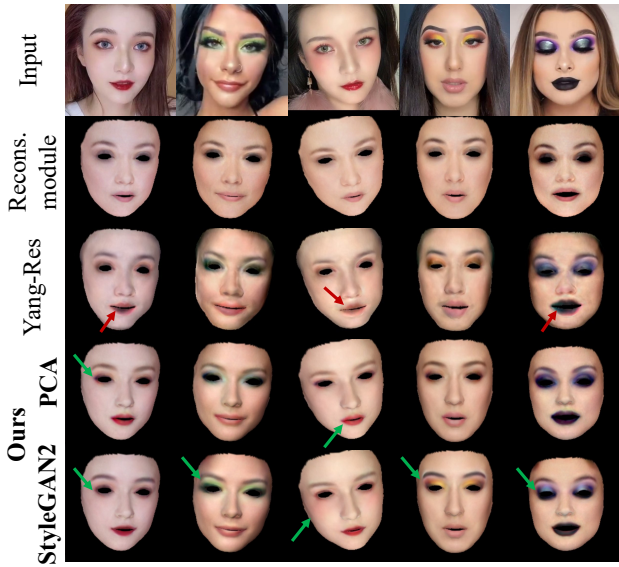


Figure 5. **Comparison with 3D face reconstruction methods using makeup prior models.** Our methods successfully reconstruct facial makeup. Specifically, the PCA model is capable of broadly recovering makeup colors, while the StyleGAN2 model achieves precise replication of complex makeup features, such as blush and gradient eyeshadow.




Figure 6. **Combining DECA and our makeup prior models.** Our makeup prior models can be incorporated into model-based 3D face reconstruction methods with the same topology, thereby endowing the capability to handle facial makeup.

assessing the reconstruction quality of eyeshadow and lipstick. The metrics presented in Tab. 2 indicate that both the PCA model and the StyleGAN2 model enhance reconstruction precision, with StyleGAN2 being superior. Detailed quantitative comparisons, including ablation results of loss functions, will be provided in the supplementary materials.

User-friendly makeup editing. Fig. 7 demonstrates the reconstructed results after editing the makeup in the original images. Our method demonstrates robustness to color variations and accommodates local patch editing, ensuring stable 3D facial applications even from less meticulous makeup edits. Yang-Res requires extra care in editing due to its over-fitting of modified images.

3D makeup transfer. The estimated makeup can be uti-

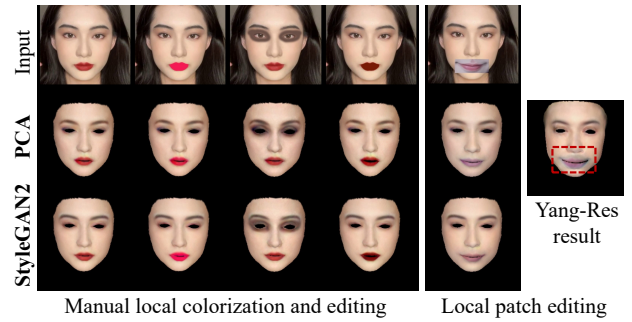


Figure 7. **Robustness and ease for makeup editing.** The results validate our method’s robustness and ease in achieving stable 3D face reconstruction with makeup edits for input images, effectively handling color changes and local patch editing. The Yang-Res approach is impacted by over-fitting issues.

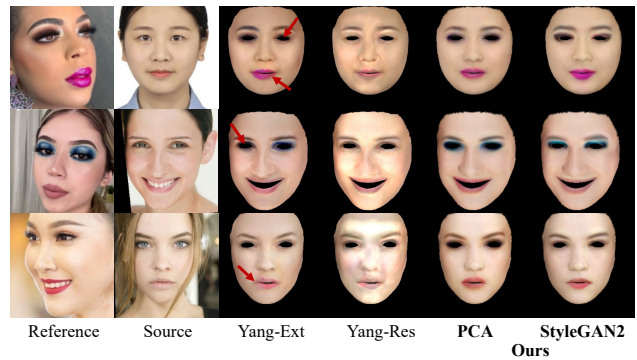


Figure 8. **Comparison of 3D makeup transfer.** The results confirm that Yang-Ext [84] struggles with large pose variations and Yang-Res [84] is constrained by its dependency on reference images, our method demonstrates accurate and reliable makeup transfer.




Figure 9. **Results of 3D makeup interpolation.** 3D makeup interpolation experiments utilizing the StyleGAN2 prior model. The top row demonstrates full interpolation between two makeup styles, blending both eyeshadow and lipstick. The bottom row highlights the capability of StyleGAN2 makeup prior model for partial makeup interpolation and transfer.

lized for a 3D makeup transfer application. As shown in Fig. 8, Yang-Ext [84] struggles with faces in large poses during estimation, leading to partial makeup loss in the makeup transfer application. The Yang-Res [84], being

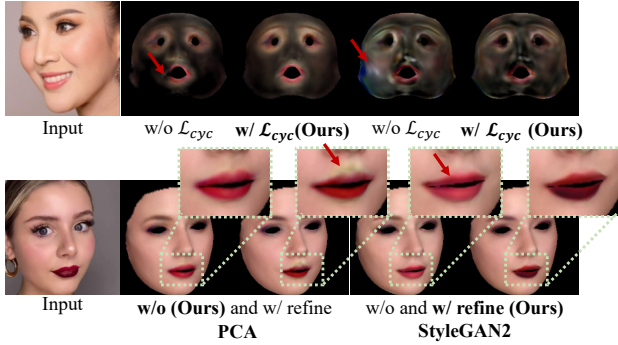


Figure 10. **Qualitative comparison for ablation study.** The top row demonstrates that incorporating the makeup consistency module contributes to the robustness of the results. The bottom row reveals that while refinement introduces artifacts in the PCA model, it significantly improves the accuracy of the StyleGAN2 model, as evidenced by the more precise makeup application on the lips.



Figure 11. **Effect of loss functions.** In the refinement process for our StyleGAN2 model, without loss functions \mathcal{L}_{sym} and \mathcal{L}'_{reg} leads to over-fitting of the input image, resulting in artifacts.

tightly coupled with the original makeup images, fails to facilitate makeup transfer. Conversely, our methods, benefiting from superior 3D makeup estimation outcomes, render the application of 3D makeup transfer both precise and effective.

3D makeup interpolation. To better visualize, we transfer the interpolated makeup to other 3D faces. Figs. 1 and 9 illustrate our experiments with makeup interpolation using the makeup prior model. Beyond the alpha matte adjustment proposed by the Yang-Ext [84], our approach leverages the inherent characteristics of the prior model, performing linear interpolation on the coefficients. Taking the StyleGAN2 prior model as an example, we demonstrate two outcomes: the top row in Fig. 9 shows the results of a full makeup interpolation between the estimated makeup coefficients of two reference images, with the noticeable blending of eyeshadow and lipstick. The bottom row illustrates our capability for partial makeup interpolation through style mixing [60]. This is achieved by retaining the selected layers of the coefficient from reference A and interpolating

the remaining layers from reference B. The lipstick color is preserved while only the eyeshadow changes. We hypothesize that this is attributable to the StyleGAN2 makeup prior model’s training, which allows for more effective disentanglement of the coefficient [38]. This represents a significant advantage over the PCA model. We intend to further investigate this feature, with the potential to apply it to applications such as makeup manipulation.

5.3. Computational Time

We measured the computational time for 3D makeup estimation. The Yang-Res and Yang-Ext methods [84] require 56.12s and 63.13s, respectively. On the other hand, our StyleGAN2-based method demonstrates a substantial speed improvement, with the inference time at 18.13s—about 3 times faster—where the makeup estimation network contributes 1.66s and the refinement process 16.47s. Our PCA-based approach achieves a remarkable 180-fold speed increase, clocking just 0.31s. We believe that while the StyleGAN2-based method offers superior quality in makeup estimation, the high-speed PCA method remains valuable for real-time situations.

5.4. Ablation Studies

As shown in Fig. 10, we conducted ablation studies on the makeup prior models (PCA and StyleGAN) to compare the outcomes with and without the use of the makeup consistency module, as well as the additional refinement process. The top row indicates that the utilization of \mathcal{L}_{cyc} leads to enhanced robustness. The bottom row shows that while refinement tends to introduce artifacts in the PCA model, it enhances the precision of the StyleGAN2 model. As a trade-off, even with the implementation of loss functions \mathcal{L}_{sym} and \mathcal{L}'_{reg} in our refinement process, the absence of the makeup consistency module may still increase the coupling between makeup and facial features. Fig. 11 shows the experiments of the loss functions \mathcal{L}_{sym} and \mathcal{L}'_{reg} . The absence of these loss functions leads to over-fitting of the input image, resulting in artifacts. Further details of ablation studies will be presented in the supplementary materials.

6. Conclusions

We have presented two makeup prior models: a PCA-based linear model for time efficiency and a StyleGAN2-based model for high-fidelity detail representation in 3D facial makeup applications. Our method addresses the challenges of self-occluded faces, enhancing robustness and reducing computation time. Demonstrated across multiple applications, our methods show improved accuracy and versatility in 3D facial makeup estimation, reconstruction, user-friendly makeup editing, transfer, and interpolation. Future work will explore the decoupling capabilities of the StyleGAN2 model for makeup manipulation.

References

- [1] Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Nießner. ClipFace: Text-guided Editing of Textured 3D Morphable Models. In *Proc. of SIGGRAPH 2023*, pages 70:1–70:11, 2023. 2, 3
- [2] Haoran Bai, Di Kang, Haoxian Zhang, Jinshan Pan, and Linchao Bao. FFHQ-UV: Normalized facial uv-texture dataset for 3D face reconstruction. In *ICCV, 2023*, pages 362–371, 2023. 3
- [3] Ziqian Bai, Zhaopeng Cui, Xiaoming Liu, and Ping Tan. Riggable 3D face reconstruction via in-network optimization. In *CVPR 2021*, pages 6216–6225, 2021. 3
- [4] Linchao Bao, Xiangkai Lin, Yajing Chen, Haoxian Zhang, Sheng Wang, Xuefei Zhe, Di Kang, Haozhi Huang, Xinwei Jiang, Jue Wang, Dong Yu, and Zhengyou Zhang. High-fidelity 3D digital human head creation from RGB-D selfies. *Transactions on Graphics*, 41(1):3:1–3:21, 2022. 3
- [5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *Proc. of SIGGRAPH 1999*, pages 187–194, 1999. 1, 3
- [6] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3D morphable model. *Transactions on Pattern Analysis and Machine Intelligence*, 25(9):1063–1074, 2003. 3
- [7] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3D morphable model learnt from 10,000 faces. In *CVPR 2016*, pages 5543–5552. 2, 3
- [8] James Booth, Anastasios Roussos, Allan Ponniah, David J. Dunaway, and Stefanos Zafeiriou. Large scale 3D morphable models. *International Journal of Computer Vision*, 126(2-4): 233–254, 2018. 2, 3
- [9] Chen Cao, Yanlin Weng, Shun Zhou, Yiyong Tong, and Kun Zhou. FaceWarehouse: A 3D facial expression database for visual computing. *Transactions on Visualization and Computer Graphics*, 20(3):413–425, 2014. 2, 3
- [10] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. PairedCycleGAN: Asymmetric style transfer for applying and removing makeup. In *CVPR 2018*, pages 40–48. Computer Vision Foundation / IEEE Computer Society, 2018. 3
- [11] Hung-Jen Chen, Ka-Ming Hui, Szu-Yu Wang, Li-Wu Tsao, Hong-Han Shuai, and Wen-Huang Cheng. BeautyGlow: On-demand makeup transfer framework with reversible generative network. In *CVPR 2019*, pages 10042–10050, 2019. 3
- [12] Jinsong Chen, Hu Han, and Shiguang Shan. Towards high-fidelity face self-occlusion recovery via multi-view residual-based GAN inversion. In *AAAI 2022*, pages 294–302, 2022. 3
- [13] Hang Dai, Nick E. Pears, William A. P. Smith, and Christian Duncan. Statistical modeling of craniofacial shape and texture. *International Journal of Computer Vision*, 128(2): 547–571, 2020. 2, 3
- [14] Radek Danecek, Michael J. Black, and Timo Bolkart. EMOCA: Emotion driven monocular face capture and animation. In *CVPR 2022*, pages 20311–20322, 2022. 3
- [15] Han Deng, Chu Han, Hongmin Cai, Guoqiang Han, and Shengfeng He. Spatially-invariant style-codes controlled makeup transfer. In *CVPR 2021*, pages 6549–6557, 2021. 3
- [16] Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, and Stefanos Zafeiriou. UV-GAN: Adversarial facial UV map completion for pose-invariant face recognition. In *CVPR 2018*, pages 7093–7102, 2018. 3
- [17] 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 *CVPR 2019 Workshops*, pages 285–295, 2019. 2, 3
- [18] Zheng Ding, Xuaner Zhang, Zhihao Xia, Lars Jebe, Zhuowen Tu, and Xiuming Zhang. DiffusionRig: Learning personalized priors for facial appearance editing. In *CVPR, 2023*, pages 12736–12746, 2023. 2
- [19] Bernhard Egger, Sandro Schönborn, Andreas Schneider, Adam Kortylewski, Andreas Morel-Forster, Clemens Blumer, and Thomas Vetter. Occlusion-aware 3D morphable models and an illumination prior for face image analysis. *IJCV 2018*, 126(12):1269–1287, 2018. 2, 3
- [20] Bernhard Egger, William A. P. Smith, Ayush Tewari, Stefanie Wuhler, Michael Zollhöfer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, Christian Theobalt, Volker Blanz, and Thomas Vetter. 3D morphable face models - past, present, and future. *Transactions on Graphics*, 39(5):157:1–157:38, 2020. 2, 3
- [21] Yuki Endo and Yoshihiro Kanamori. Controlling stylegans using rough scribbles via one-shot learning. *Computer Animation and Virtual Worlds*, 33(5), 2022. 2, 5
- [22] Haiwen Feng, Timo Bolkart, Joachim Tesch, Michael J. Black, and Victoria Fernández Abrevaya. Towards racially unbiased skin tone estimation via scene disambiguation. In *ECCV 2022*, pages 72–90, 2022. 2, 3
- [23] Ruili Feng, Deli Zhao, and Zheng-Jun Zha. Understanding noise injection in GANs. In *ICML, 2021*, pages 3284–3293, 2021. 5
- [24] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. *Transactions on Graphics*, 40(4): 88:1–88:13, 2021. 1, 2, 3, 6
- [25] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. GANFIT: Generative adversarial network fitting for high fidelity 3D face reconstruction. In *CVPR 2019*, pages 1155–1164, 2019. 2, 3
- [26] Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. OS-TeC: One-shot texture completion. In *CVPR 2021*, pages 7628–7638, 2021. 3
- [27] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T. Freeman. Unsupervised training for 3D morphable model regression. In *CVPR 2018*, pages 8377–8386, 2018. 3
- [28] Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Lüthi, Sandro Schönborn, and Thomas Vetter. Morphable face models - an open framework. In *Proceedings of International Conference on Automatic Face & Gesture Recognition*, pages 75–82, 2018. 2, 3
- [29] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner.

- Learning neural parametric head models. In *CVPR, 2023*, pages 21003–21012. IEEE, 2023. [2](#)
- [30] Qiao Gu, Guanzhi Wang, Mang Tik Chiu, Yu-Wing Tai, and Chi-Keung Tang. LADN: Local adversarial disentangling network for facial makeup and de-makeup. In *ICCV 2019*, pages 10480–10489, 2019. [3](#), [5](#)
- [31] Yudong Guo, Juyong Zhang, Jianfei Cai, Boyi Jiang, and Jianmin Zheng. CNN-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *Transactions on Pattern Analysis and Machine Intelligence*, 41(6):1294–1307, 2019. [3](#)
- [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR, 2016*, pages 770–778, 2016. [2](#), [4](#)
- [33] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. HeadNeRF: A real-time nerf-based parametric head model. In *CVPR 2022*, 2022. [2](#)
- [34] Cheng-Guo Huang, Wen-Chieh Lin, Tsung-Shian Huang, and Jung-Hong Chuang. Physically-based cosmetic rendering. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, page 190, 2013. [2](#)
- [35] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR 2017*, pages 5967–5976, 2017. [2](#), [3](#)
- [36] Wentao Jiang, Si Liu, Chen Gao, Jie Cao, Ran He, Jiashi Feng, and Shuicheng Yan. PSGAN: Pose and expression robust spatial-aware GAN for customizable makeup transfer. In *CVPR 2020*, pages 5193–5201, 2020. [3](#), [5](#), [6](#)
- [37] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV 2016*, pages 694–711, 2016. [5](#)
- [38] Tero Karras, Samuli Laine, and Timo Aila. A Style-based generator architecture for generative adversarial networks. In *CVPR, 2019*, pages 4401–4410, 2019. [8](#)
- [39] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR 2020*, pages 8107–8116, 2020. [2](#), [3](#), [4](#)
- [40] Jongyoo Kim, Jiaolong Yang, and Xin Tong. Learning high-fidelity face texture completion without complete face texture. In *CVPR 2021*, pages 13970–13979, 2021. [3](#)
- [41] Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Triantafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. AvatarMe: Realistically renderable 3D facial reconstruction “in-the-wild”. In *CVPR 2020*, pages 757–766, 2020. [2](#)
- [42] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos P Zafeiriou. AvatarMe++: Facial shape and BRDF inference with photorealistic rendering-aware GANs. *Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [43] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. FitMe: Deep photorealistic 3D morphable model avatars. In *CVPR, 2023*, pages 8629–8640, 2023. [2](#), [3](#)
- [44] Chen Li, Kun Zhou, and Stephen Lin. Simulating makeup through physics-based manipulation of intrinsic image layers. In *CVPR 2015*, pages 4621–4629, 2015. [2](#), [3](#)
- [45] Chunlu Li, Andreas Morel-Forster, Thomas Vetter, Bernhard Egger, and Adam Kortylewski. To fit or not to fit: Model-based face reconstruction and occlusion segmentation from weak supervision. In *CVPR 2023*, 2023. [3](#)
- [46] Ruilong Li, Karl Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha Prasad, Bipin Kishore, Jun Xing, and Hao Li. Learning formation of physically-based face attributes. In *CVPR 2020*, pages 3407–3416, 2020. [2](#), [3](#)
- [47] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *Transactions on Graphics*, 36(6):194:1–194:17, 2017. [2](#), [3](#), [4](#), [5](#)
- [48] Tingting Li, Ruihe Qian, Chao Dong, Si Liu, Qiong Yan, Wenwu Zhu, and Liang Lin. BeautyGAN: Instance-level facial makeup transfer with deep generative adversarial network. In *Proceedings of International Conference on Multimedia*, pages 645–653, 2018. [3](#), [5](#)
- [49] Connor Z. Lin, Koki Nagano, Jan Kautz, Eric R. Chan, Umar Iqbal, Leonidas J. Guibas, Gordon Wetzstein, and Sameh Khamis. Single-shot implicit morphable faces with consistent texture parameterization. In *Proc. of SIGGRAPH 2023*, pages 83:1–83:12, 2023. [3](#)
- [50] Feng Liu, Luan Tran, and Xiaoming Liu. 3D face modeling from diverse raw scan data. In *ICCV 2019*, pages 9407–9417, 2019. [2](#), [3](#)
- [51] Haiyang Liu, Naoya Iwamoto, Zihao Zhu, Zhengqing Li, You Zhou, Elif Bozkurt, and Bo Zheng. DisCo: Disentangled implicit content and rhythm learning for diverse co-speech gestures synthesis. In *ACMMM, 2022*, pages 3764–3773, 2022. [2](#)
- [52] Haiyang Liu, Zihao Zhu, Naoya Iwamoto, Yichen Peng, Zhengqing Li, You Zhou, Elif Bozkurt, and Bo Zheng. BEAT: A large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis. In *ECCV, 2022*, pages 612–630, 2022. [2](#)
- [53] Si Liu, Xinyu Ou, Ruihe Qian, Wei Wang, and Xiaochun Cao. Makeup like a superstar: Deep localized makeup transfer network. In *IJCAI 2016*, pages 2568–2575, 2016. [3](#)
- [54] Huiwen Luo, Koki Nagano, Han-Wei Kung, Qingguo Xu, Zejian Wang, Lingyu Wei, Liwen Hu, and Hao Li. Normalized avatar synthesis using styleGAN and perceptual refinement. In *CVPR 2021*, pages 11662–11672, 2021. [2](#), [3](#), [5](#)
- [55] Yueming Lyu, Jing Dong, Bo Peng, Wei Wang, and Tieniu Tan. SOGAN: 3D-aware shadow and occlusion robust GAN for makeup transfer. In *Proceedings of International Conference on Multimedia*, pages 3601–3609, 2021. [3](#)
- [56] Thao Nguyen, Anh Tuan Tran, and Minh Hoai. Lipstick ain’t enough: Beyond color matching for in-the-wild makeup transfer. In *CVPR 2021*, pages 13305–13314, 2021. [3](#)
- [57] Stylianos Ploumpis, Evangelos Ververas, Eimear O’ Sullivan, Stylianos Moschoglou, Haoyang Wang, Nick E. Pears, William A. P. Smith, Baris Gecer, and Stefanos Zafeiriou. Towards a complete 3D morphable model of the human head.

- Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4142–4160, 2021. 2, 3
- [58] Puntawat Ponglertnapakorn, Nontawat Tritrong, and Supasorn Suwajanakorn. DiFaReli: Diffusion face relighting. In *ICCV 2023*, pages 22589–22600, 2023. 2
- [59] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J. Black. Generating 3D faces using convolutional mesh autoencoders. In *ECCV 2018*, pages 725–741, 2018. 2, 3
- [60] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: A stylegan encoder for image-to-image translation. In *CVPR, 2021*, pages 2287–2296, 2021. 2, 4, 8
- [61] Sami Romdhani and Thomas Vetter. Estimating 3D shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *CVPR 2005*, pages 986–993, 2005. 3
- [62] Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, and Hao Li. Photorealistic facial texture inference using deep neural networks. In *CVPR 2017*, pages 2326–2335, 2017. 2, 3
- [63] Kristina Scherbaum, Tobias Ritschel, Matthias Hullin, Thorsten Thormählen, Volker Blanz, and Hans-Peter Seidel. Computer-suggested facial makeup. *Computer Graphics Forum*, 30(2), 2011. 2
- [64] Andreas Schneider, Sandro Schönborn, Bernhard Egger, Lavrenti Frobeen, and Thomas Vetter. Efficient global illumination for morphable models. In *ICCV 2017*, pages 3885–3893, 2017. 2, 3
- [65] Ron Slossberg, Ibrahim Jubran, and Ron Kimmel. Unsupervised high-fidelity facial texture generation and reconstruction. In *ECCV 2022*, pages 212–229, 2022. 2, 3
- [66] William A. P. Smith, Alassane Seck, Hannah Dee, Bernard Tiddeman, Joshua B. Tenenbaum, and Bernhard Egger. A morphable face albedo model. In *CVPR 2020*, pages 5010–5019, 2020. 2, 3
- [67] Skylar Sutherland, Bernhard Egger, and Josh Tenenbaum. Building 3D morphable models from a single scan. In *ICCVW 2021*, pages 1–11, 2021. 2
- [68] Ayush Tewari, Michael Zollhöfer, Hyeongwoo Kim, Pablo Garrido, Florian Bernard, Patrick Pérez, and Christian Theobalt. MoFA: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In *CVPR 2017*, pages 3735–3744, 2017. 3
- [69] Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeongwoo Kim, Patrick Pérez, and Christian Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 Hz. In *CVPR 2018*, pages 2549–2559, 2018. 3
- [70] Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohamed Elgharib, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. FML: Face model learning from videos. In *CVPR 2019*, pages 10812–10822, 2019. 2
- [71] Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of RGB videos. In *CVPR 2016*, pages 2387–2395, 2016. 3
- [72] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gérard G. Medioni. Regressing robust and discriminative 3D morphable models with a very deep neural network. In *CVPR 2017*, pages 1493–1502, 2017. 3
- [73] Luan Tran and Xiaoming Liu. Nonlinear 3D face morphable model. In *CVPR 2018*, pages 7346–7355, 2018. 2
- [74] Lizhen Wang, Zhiyuan Chen, Tao Yu, Chenguang Ma, Liang Li, and Yebin Liu. FaceVerse: a fine-grained and detail-controllable 3D face morphable model from a hybrid dataset. In *CVPR 2022*, pages 20301–20310, 2022. 2, 3
- [75] Yandong Wen, Weiyang Liu, Bhiksha Raj, and Rita Singh. Self-supervised 3D face reconstruction via conditional estimation. In *ICCV 2021*, pages 13269–13278, 2021. 3
- [76] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. GAN inversion: A survey. *TPAMI*, 45(3):3121–3138, 2023. 5
- [77] Jianfeng Xiang, Junliang Chen, Wenshuang Liu, Xianxu Hou, and Linlin Shen. RamGAN: Region attentive morphing GAN for region-level makeup transfer. In *ECCV 2022*, pages 719–735, 2022. 3
- [78] Shugo Yamaguchi, Shunsuke Saito, Koki Nagano, Yajie Zhao, Weikai Chen, Kyle Olszewski, Shigeo Morishima, and Hao Li. High-fidelity facial reflectance and geometry inference from an unconstrained image. *Transactions on Graphics*, 37(4):162, 2018. 2, 3
- [79] Qixin Yan, Chunle Guo, Jixin Zhao, Yuekun Dai, Chen Change Loy, and Chongyi Li. BeautyREC: Robust, efficient, and component-specific makeup transfer. In *CVPRW, 2023*, pages 1102–1110, 2023. 3, 5, 6
- [80] Chenyu Yang, Wanrong He, Yingqing Xu, and Yang Gao. EleGANt: Exquisite and locally editable GAN for makeup transfer. In *ECCV 2022*, 2022. 3
- [81] 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 *CVPR 2020*, pages 598–607, 2020. 2, 3
- [82] Kaibing Yang, Hong Shang, Tianyang Shi, Xinghan Chen, Jin Zhou, Zhongqian Sun, and Wei Yang. Asm: Adaptive skinning model for high-quality 3D face modeling. In *ICCV, 2023*, pages 20709–20717, 2023. 2, 3
- [83] Xingchao Yang and Takafumi Taketomi. BareSkinNet: Demakeup and De-lighting via 3D Face Reconstruction. *Computer Graphics Forum*, 2022. 3
- [84] Xingchao Yang, Takafumi Taketomi, and Yoshihiro Kanamori. Makeup extraction of 3D representation via illumination-aware image decomposition. *Computer Graphics Forum*, 42(2):293–307, 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [85] Tarun Yenamandra, Ayush Tewari, Florian Bernard, Hans-Peter Seidel, Mohamed Elgharib, Daniel Cremers, and Christian Theobalt. i3DMM: Deep implicit 3D morphable model of human heads. In *CVPR 2021*, pages 12803–12813, 2021. 2
- [86] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR, 2018*, pages 586–595, 2018. 6

- [87] Tianke Zhang, Xuangeng Chu, Yunfei Liu, Lijian Lin, Zhendong Yang, Zhengzhuo Xu, Chengkun Cao, Fei Yu, Changyin Zhou, Chun Yuan, et al. Accurate 3D face reconstruction with facial component tokens. In *ICCV, 2023*, pages 9033–9042, 2023. [3](#)
- [88] Yiyu Zhuang, Hao Zhu, Xusen Sun, and Xun Cao. MoFaNeRF: Morphable facial neural radiance field. In *ECCV 2022*, pages 268–285, 2022. [2](#)
- [89] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. In *ECCV 2022*, pages 250–269, 2022. [3](#)
- [90] Michael Zollhöfer, Justus Thies, Pablo Garrido, Derek Bradley, Thabo Beeler, Patrick Pérez, Marc Stamminger, Matthias Nießner, and Christian Theobalt. State of the art on monocular 3D face reconstruction, tracking, and applications. *Computer Graphics Forum*, 37(2):523–550, 2018. [2](#), [3](#)