



Learning Flow-Based Feature Warping for Face Frontalization with Illumination Inconsistent Supervision

Yuxiang Wei¹, Ming Liu¹, Haolin Wang¹, Ruifeng Zhu^{2,4}, Guosheng Hu³,
and Wangmeng Zuo^{1,5}(✉)

¹ School of Computer Science and Technology, Harbin Institute of Technology,
Harbin, China

yuxiang.wei.cs@gmail.com, csmliu@outlook.com, Why_cs@outlook.com,
wmzuo@hit.edu.cn

² University of Burgundy Franche-Comté, Besançon, France

reefing.z@gmail.com

³ Anyvision, Belfast, UK

huguosheng100@gmail.com

⁴ University of the Basque Country, Eibar, Spain

⁵ Peng Cheng Lab, Shenzhen, China

Abstract. Despite recent advances in deep learning-based face frontalization methods, photo-realistic and illumination preserving frontal face synthesis is still challenging due to large pose and illumination discrepancy during training. We propose a novel Flow-based Feature Warping Model (FFWM) which can learn to synthesize photo-realistic and illumination preserving frontal images with illumination inconsistent supervision. Specifically, an Illumination Preserving Module (IPM) is proposed to learn illumination preserving image synthesis from illumination inconsistent image pairs. IPM includes two pathways which collaborate to ensure the synthesized frontal images are illumination preserving and with fine details. Moreover, a Warp Attention Module (WAM) is introduced to reduce the pose discrepancy in the feature level, and hence to synthesize frontal images more effectively and preserve more details of profile images. The attention mechanism in WAM helps reduce the artifacts caused by the displacements between the profile and the frontal images. Quantitative and qualitative experimental results show that our FFWM can synthesize photo-realistic and illumination preserving frontal images and performs favorably against the state-of-the-art results. Our code is available at <https://github.com/csyxwei/FFWM>.

Keywords: Face frontalization · Illumination preserving · Optical flow · Guided filter · Attention mechanism

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-58610-2_33) contains supplementary material, which is available to authorized users.

1 Introduction

Face frontalization aims to synthesize the frontal view face from a given profile. Frontalized faces can be directly used for general face recognition methods without elaborating additional complex modules. Apart from face recognition, generating photo-realistic frontal face is beneficial for a series of face-related tasks, including face reconstruction, face attribute analysis, facial animation, *etc.*

Traditional methods address this problem through 2D/3D local texture warping [6, 36] or statistical modeling [22]. Recently, GAN-based methods have been proposed to recover a frontal face in a data-driven manner [1, 8, 10, 28, 31, 32, 34, 35]. For instance, Yin *et al.* [32] propose DA-GAN to capture the long-displacement contextual information from illumination discrepancy images under large poses. However, it recovers inconsistent illumination on the synthesized image. Flow-based method [33] predicts a dense pixel correspondence between the profile and frontal image and uses it to deform the profile face to the frontal view. However, deforming the profile face in the image space directly leads to obvious artifacts and missing pixels should be addressed under large poses.




Fig. 1. $\pm 45^\circ$, $\pm 60^\circ$, $\pm 75^\circ$ and $\pm 90^\circ$ images of the two persons in the Multi-PIE. Each row images have the same flash in the recording environment.

The existing methods do not consider the illumination inconsistency between the profile and ground-truth frontal image. Taking the widely used benchmark Multi-PIE [4] as an example, the visual illumination conditions on several poses are significantly different from the ground-truth frontal images as shown in Fig. 1. Except $\pm 90^\circ$, the other face images are produced by the same camera type. The variation in camera types causes obvious illumination inconsistency between the $\pm 90^\circ$ images and the ground-truth frontal image. Although efforts have been made to manually color-balance those same type cameras, the illumination of resulting images within $\pm 75^\circ$ (except 0°) still look visually distinguishable with the ground-truth frontal image. Since the existing methods minimize pixel-wise loss between the synthesized image and the illumination inconsistent ground-truth, they tend to change both the pose and the illumination of the profile face image, while the latter actually is not acceptable in face editing and synthesis.

To address the above issue, this paper proposes a novel Flow-based Feature Warping Model (FFWM) which can synthesize photo-realistic and illumination

preserving frontal image from illumination inconsistent image pairs. In particular, FFWM incorporates the flow estimation with two modules: Illumination Preserving Module (IPM) and Warp Attention Module (WAM). Specifically, we estimate the optical flow fields from the given profile: the reverse and forward flow fields are predicted to warp the front face to the profile view and vice versa, respectively. The estimated flow fields are fed to IPM and WAM to conduct face frontalization.

The IPM is proposed to synthesize illumination preserving images with fine facial details from illumination inconsistent image pairs. Specifically, IPM contains two pathways: (1) Illumination Preserving Pathway and (2) Illumination Adaption Pathway. For (1), an illumination preserving loss equipped with the reverse flow field is introduced to constrain the illumination consistency between synthesized images and the profiles. For (2), guided filter [7] is introduced to further eliminate the illumination discrepancy and learns frontal view facial details from illumination inconsistent ground-truth image. The WAM is introduced to reduce the pose discrepancy in the feature level. It uses the forward flow field to align the profile features to the frontal view. This flow provides an explicit and accurate supervision signal to guide the frontalization. The attention mechanism in WAM helps to reduce the artifacts caused by the displacements between the profile and frontal images.

Quantitative and qualitative experimental results demonstrate the effectiveness of our FFWM on synthesizing photo-realistic and illumination preserving faces with large poses and the superiority over the state-of-the-art results on the testing benchmarks. Our contributions can be summarized as:

- A Flow-based Feature Warping Model (FFWM) is proposed to address the challenging problem in face frontalization, *i.e.* photo-realistic and illumination preserving image synthesis.
- Illumination Preserving Module (IPM) equipped with guided filter and flow field is proposed to achieve illumination preserving image synthesis. Warp Attention Module (WAM) uses the attention mechanism to effectively reduce the pose discrepancy in the feature level under the explicit and effective guidance from flow estimation.
- Quantitative and qualitative results demonstrate that the proposed FFWM outperforms the state-of-the-art methods.

2 Related Work

2.1 Face Frontalization

Face frontalization aims to synthesize the frontal face from a given profile. Traditional methods address this problem through 2D/3D local texture warping [6, 36] or statistical modeling [22]. Hassner *et al.* [6] employ a mean 3D model for face normalization. A statistical model [22] is used for frontalization and landmark detection by solving a constrained low-rank minimization problem.

Benefiting from deep learning, many GAN-based methods [8, 10, 27, 28, 32, 33] are proposed for face frontalization. Huang *et al.* [10] use a two-pathway GAN architecture for perceiving global structures and local details simultaneously. Domain knowledge such as symmetry and identity information of face is used to make the synthesized faces photo-realistic. Zhao *et al.* [34] propose PIM with introducing a domain adaptation strategy for pose invariant face recognition. 3D-based methods [1, 2, 31, 35] attempt to combine prior knowledge of 3D face with face frontalization. Yin *et al.* [31] incorporate 3D face model into GAN to solve the problem of large pose face frontalization in the wild. HF-PIM [1] combines the advantages of 3D and GAN-based methods and frontalizes profile images via a novel texture warping procedure. In addition to supervised learning, Qian *et al.* [17] propose a novel Face Normalization Model (FNM) for unsupervised face generation with unpaired face images in the wild. Note that FNM focuses on face normalization, without considering preserving illumination.

Instead of learning function to represent the frontalization procedure, our method gets frontal warped feature by flow field and reconstructs illumination preserving and identity preserving frontal view face.

2.2 Optical Flow

Optical flow estimation has many applications, *e.g.*, action recognition, autonomous driving and video editing. With the progress in deep learning, FlowNet [3], FlowNet2 [13] and others achieve good results by end-to-end supervised learning. While SpyNet [18], PWC-Net [26] and LiteFlowNet [11] also use coarse-to-fine strategy to refine the initial flow. It is worth mentioning that PWC-Net and LiteFlowNet have smaller size and are easier to train. Based on weight sharing and residual subnetworks, Hur and Roth [12] learn bi-directional optical flow and occlusion estimation jointly. Bilateral refinement of flow and occlusion address blurry estimation, particularly near motion boundaries. By the global and local correlation layers, GLU-Net [29] can resolve the challenges of large displacements, pixel-accuracy, and appearance changes.

In this work, we estimate bi-directional flow fields to represent dense pixel correspondence between the profile and frontal faces, which are then exploited to obtain frontal view features and preserve illumination condition, respectively.

3 Proposed Method

Let $\{I, I^{gt}\}$ be a pair of profile and frontal face image of the same person. Given a profile image I , our goal is to train a model \mathcal{R} to synthesize the corresponding frontal face image $\hat{I} = \mathcal{R}(I)$, which is expected to be photo-realistic and illumination preserving. To achieve this, we propose the Flow-based Feature Warping Model (FFWM). As shown in Fig. 2, FFWM takes U-net [20] as the backbone and incorporates with the Illumination Preserving Module (IPM) and the Warp Attention Module (WAM) to synthesize \hat{I} . In addition, FFWM uses optical flow fields which are fed to IPM and WAM to conduct frontalization. Specifically, we

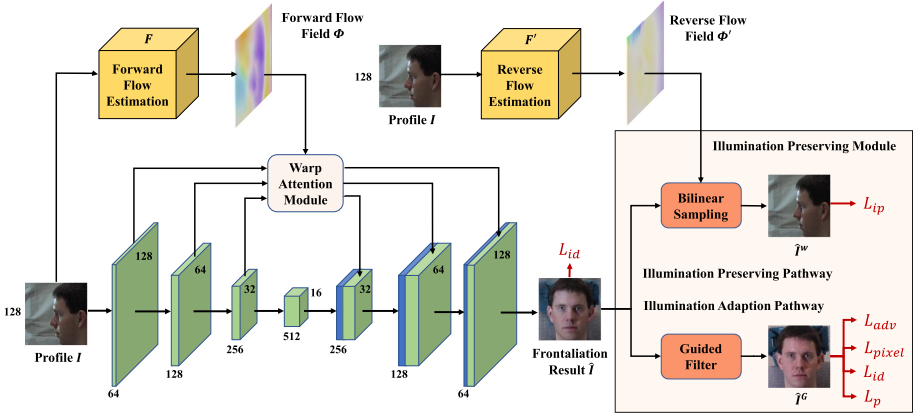


Fig. 2. The architecture of our FFWM. Illumination Preserve Module is incorporated to facilitate synthesized frontal image \hat{I} to be illumination preserving and facial details preserving in two independent pathways. Based on the skip connection, the Warp Attention Module helps synthesize frontal image effectively. Losses are shown in red color, which \hat{I}^w is the synthesized image \hat{I} warped by Φ' and the \hat{I}^G is the guided filter output. (Color figure online)

compute the forward and reverse flow fields to warp the profile to the frontal view and vice versa, respectively.

In this section, we first introduce the bi-directional flow fields estimation in Sect. 3.1. IPM and WAM are introduced in Sect. 3.2 and Sect. 3.3. Finally, the loss functions are detailed in Sect. 3.4.

3.1 Bi-directional Flow Fields Estimation

Face frontalization can be viewed as the face rotation transformation, and the flow field can model this rotation by establishing the pixel-level correspondence between the profile and frontal faces. Traditional optical flow methods [3, 13] take two frames as the input. However, we only use one profile image as the input. In this work, we adopt the FlowNetSD in FlowNet2 [13] as our flow estimation network, and change the input channel from 6 (two frames) to 3 (one image). For preserving illumination and frontalization, we estimate the reverse flow field Φ' and the forward flow field Φ from the profile image, respectively.

Reverse Flow Field. Given the profile image I , reverse flow estimation network \mathcal{F}' predicts the reverse flow field Φ' which can warp the ground-truth frontal image I^{gt} to the profile view as I .

$$\Phi' = \mathcal{F}'(I; \Theta_{\mathcal{F}'}), \quad (1)$$

$$I^{w'} = \mathcal{W}(I^{gt}, \Phi'), \quad (2)$$

Where $\Theta_{\mathcal{F}'}$ denotes the parameters of \mathcal{F}' , and $\mathcal{W}(\cdot)$ [14] is the bilinear sampling operation. To learn an accurate reverse flow field, \mathcal{F}' is pretrained with the landmark loss [15], sampling correctness loss [19] and the regularization term [19].

Forward Flow Field. Given the profile image I , forward flow estimation network \mathcal{F} predicts the forward flow field Φ which can warp I to the frontal view.

$$\Phi = \mathcal{F}(I; \Theta_{\mathcal{F}}), \quad (3)$$

$$I^w = \mathcal{W}(I, \Phi), \quad (4)$$

Where $\Theta_{\mathcal{F}}$ denotes the parameters of \mathcal{F} . To learn an accurate forward flow field, \mathcal{F} is pretrained with the same losses as \mathcal{F}' .

Then two flow fields Φ' and Φ are used for the IPM and WAM to generate illumination preserving and photo-realistic frontal images.

3.2 Illumination Preserving Module

Without considering inconsistent illumination in the face datasets, the existing frontalization methods potentially overfit to the wrong illumination. To effectively decouple the illumination and the facial details, hence to synthesize illumination preserving faces with fine details, we propose the Illumination Preserving Module (IPM). As shown in Fig. 2, IPM consists of two pathways. Illumination preserving pathway ensures that the illumination condition of the synthesized image \hat{I} is consistent with the profile I . Illumination adaption pathway ensures that the facial details of the synthesized image \hat{I} are consistent with the ground-truth I^{gt} .

Illumination Preserving Pathway. Because the illumination condition is diverse and cannot be quantified as a label, it is hard to learn reliable and independent illumination representation from face images. Instead of constraining the illumination consistency between the profile and the synthesized image in the feature space, we directly constrain it in the image space. As shown in Fig. 2, in the illumination preserving pathway, we firstly use the reverse flow field Φ' to warp the synthesized image \hat{I} to the profile view,

$$\hat{I}^w = \mathcal{W}(\hat{I}, \Phi'). \quad (5)$$

Then an illumination preserving loss is defined on the warped synthesized image \hat{I}^w to constrain the illumination consistency between the synthesized image \hat{I} and the profile I . By minimizing it, FFWM can synthesize illumination preserving frontal images.

Illumination Adaption Pathway. Illumination preserving pathway cannot ensure the consistency of facial details between the synthesized image \hat{I} and the ground-truth I^{gt} , so we constrain it in the illumination adaption pathway. Since the illumination of profile I is inconsistent with the ground-truth I^{gt} under large poses, adding constraints directly between \hat{I} and I^{gt} eliminates the illumination consistency between \hat{I} and I . So a guided filter layer [7] is firstly used to transfer the illumination of images. Specifically, the guided filter takes I^{gt} as the input image and \hat{I} as the guidance image,

$$\hat{I}^G = \mathcal{G}(\hat{I}, I^{gt}), \tag{6}$$

where $\mathcal{G}(\cdot)$ denotes the guided filter, and we set the radius of filter as the quarter of the image resolution. After filtering, the guided filter result \hat{I}^G has the same illumination with I^{gt} while keeping the same facial details with \hat{I} . Then the illumination-related losses (*e.g.*, pixel-wise loss, perceptual loss) are defined on \hat{I}^G to facilitate our model synthesize \hat{I} with much finer details. By this means, \hat{I} can become much more similar to I^{gt} in facial details without changing the illumination consistency between \hat{I} and I .

Note that the guided filter has no trainable parameters and potentially cause our model trap into local minima during training. So we apply the guided filter after several iterations, providing stable and robust initialization to our model.

3.3 Warp Attention Module

The large pose discrepancy makes it difficult to synthesize correct facial details in the synthesized images. To reduce the pose discrepancy between the profile and frontal face, Warp Attention Module (WAM) is proposed to align the profile face to the frontal one in the feature space. We achieve this by warping the profile features guided by the forward flow field Φ . The architecture of our WAM is illustrated in Fig. 3. It contains two steps: flow-guided feature warping and feature attention.

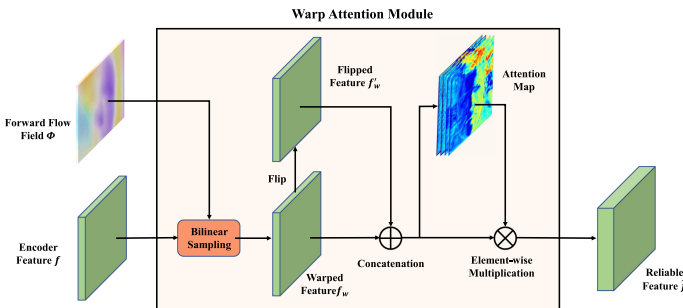


Fig. 3. The architecture of Warp Attention Module. Considering the symmetry prior of human face, WAM also contains flipped warped feature.

Flow-Guided Feature Warping. Because the profile and frontal face have different visible areas, the forward flow field Φ cannot establish a complete pixel-level correspondence between them. Hence, warping profile face directly leads to artifacts. Here we incorporate Φ with bilinear sampling operation $\mathcal{W}(\cdot)$ to warp the profile face to the frontal one in the feature space. Additionally, we use the symmetry prior of human face, and take both warped features and its horizontal flip to guide the frontal image synthesis.

$$f_w = \mathcal{W}(f, \Phi), \quad (7)$$

Where f denotes the encoder feature of the profile. Let f_w' denotes the horizontal flip of f_w , and $(f_w \oplus f_w')$ denotes the concatenation of f_w and f_w' .

Feature Attention. After warping, the warped feature encodes the backgrounds and self-occlusion artifacts, which leads to degraded frontalization performance. To eliminate above issue and extract reliable frontal feature, an attention mechanism is then used to adaptively focus on the critical parts of $(f_w \oplus f_w')$. The warped feature $(f_w \oplus f_w')$ is firstly fed into a Conv-BN-ReLU-ResidualBlock Layer to generate an attention map A , which has the same height, width and channel size with $(f_w \oplus f_w')$. Then the reliable frontalized feature \hat{f} is obtained by,

$$\hat{f} = A \otimes (f_w \oplus f_w'), \quad (8)$$

where \otimes denotes element-wise multiplication. \hat{f} is then skip connected to the decoder to help generate photo-realistic frontal face image \hat{I} .

3.4 Loss Functions

In this section, we formulate the loss functions used in our work. The background of images is masked to make the loss functions focus on the facial area.

Pixel-Wise Loss. Following [8, 10], we employ a multi-scale pixel-wise loss on the guided filter result \hat{I}^G to constrain the content consistency,

$$\mathcal{L}_{pixel} = \sum_{s=1}^S \left\| \hat{I}_s^G - I_s^{gt} \right\|_1, \quad (9)$$

Where S denotes the number of scales. In our experiments, we set $S = 3$, and the scales are 32×32 , 64×64 and 128×128 .

Perceptual Loss. Pixel-wise loss tends to generate over-smoothing results. To alleviate this, we introduce the perceptual loss defined on the VGG-19 network [25] pre-trained on ImageNet [21],

$$\mathcal{L}_p = \sum_i w_i \left\| \phi_i(\hat{I}^G) - \phi_i(I^{gt}) \right\|_1, \quad (10)$$

where $\phi_i(\cdot)$ denotes the output of the i -th VGG-19 layer. In our implementation, we use Conv1-1, Conv2-1, Conv3-1, Conv4-1 and Conv5-1 layer, and set $w = \{1, 1/2, 1/4, 1/4, 1/8\}$. To improve synthesized imagery in the particular facial regions, we also use the perceptual loss on the facial regions like eyes, nose and mouth.

Adversarial Loss. Following [24], we adopt a multi-scale discriminator and adversarial learning to help synthesize photo-realistic images.

$$\mathcal{L}_{adv} = \min_R \max_D \mathbb{E}_{I^{gt}} [\log D(I^{gt})] - \mathbb{E}_{\hat{I}^G} [\log(1 - D(\hat{I}^G))]. \quad (11)$$

Illumination Preserving Loss. To preserve the illumination of profile I on synthesized image \hat{I} , we define the illumination preserving loss on the warped synthesized image \hat{I}^w at different scales,

$$\mathcal{L}_{ip} = \sum_{s=1}^S \left\| \hat{I}_s^w - I_s \right\|_1, \quad (12)$$

Where S denotes the number of scales, and the scale setting is same as Eq. (9).

Identity Preserving Loss. Following [8, 10], we present an identity preserving loss to preserve the identity information of the synthesized image \hat{I} ,

$$\mathcal{L}_{id} = \left\| \psi_{fc2}(\hat{I}) - \psi_{fc2}(I^{gt}) \right\|_1 + \left\| \psi_{pool}(\hat{I}) - \psi_{pool}(I^{gt}) \right\|_1, \quad (13)$$

Where $\psi(\cdot)$ denotes the pretrained LightCNN-29 [30]. $\psi_{fc2}(\cdot)$ and $\psi_{pool}(\cdot)$ denote the outputs of the last pooling layer and the fully connected layer respectively. To preserve the identity information, we add the identity loss on both \hat{I} and \hat{I}^G .

Overall Losses. Finally, we combine all the above losses to give the overall model objective,

$$\mathcal{L} = \lambda_0 \mathcal{L}_{pixel} + \lambda_1 \mathcal{L}_p + \lambda_2 \mathcal{L}_{adv} + \lambda_3 \mathcal{L}_{ip} + \lambda_4 \mathcal{L}_{id}, \quad (14)$$

Where λ_* denotes the different losses tradeoff parameters.

4 Experiments

To illustrate our model can synthesize photo-realistic and illumination preserving images while preserving identity, we evaluate our model qualitatively and quantitatively under both controlled and in the wild settings. In the following subsections, we begin with an introduction of datasets and implementation details. Then we demonstrate the merits of our model on qualitative synthesis results and quantitative recognition results over the state-of-the-art methods. Lastly, we conduct an ablation study to demonstrate the benefits from each part of our model.

4.1 Experimental Settings



Fig. 4. Synthesis results of the Multi-PIE dataset by our model under large poses and illumination inconsistent conditions. Each pair presents profile (left), synthesized frontal face (middle) and ground-truth frontal face (right).

Datasets. We adopt the Multi-PIE dataset [4] as our training and testing set. Multi-PIE is widely used for evaluating face synthesis and recognition in the controlled setting. It contains 754,204 images of 337 identities from 15 poses and 20 illumination conditions. In this paper, the face images with neutral expression under 20 illuminations and 13 poses within $\pm 90^\circ$ are used. For a fair comparison, we follow the test protocols in [8] and utilize two settings to evaluate our model. The first setting (Setting 1) only contains images from Session 1. The training set is composed of all the first 150 identities images. For testing, one gallery image with frontal view and normal illumination is used for the remaining 99 identities. For the second setting (Setting 2), we use neutral expression images from all four sessions. The first 200 identities and the remaining 137 identities are used for training and testing, respectively. Each testing identity has one gallery image with frontal view and normal illumination from the first appearance.

LFW [9] contains 13,233 face images collected in unconstrained environment. It will be used to evaluate the frontalization performance in uncontrolled settings.

Implementation Details. All images in our experiments are cropped and resized to 128×128 according to facial landmarks, and image intensities are linearly scaled to the range of $[0, 1]$. The LightCNN-29 [30] is pretrained on MS-Celeb-1M [5] and fine-tuned on the training set of Multi-PIE.

In all our experiments, we empirically set $\lambda_0 = 5, \lambda_1 = 1, \lambda_2 = 0.1, \lambda_3 = 15, \lambda_4 = 1$. The learning rate is initialized by 0.0004 and the batch size is 8. The flow estimation networks \mathcal{F} and \mathcal{F}' are pre-trained and then all networks are end-to-end trained by minimizing the objective \mathcal{L} with setting $\text{lr} = 0.00005$ for \mathcal{F} and \mathcal{F}' .

4.2 Qualitative Evaluation

In this subsection, we qualitatively compare the synthesized results of our model against state-of-the-art face frontalization methods. We train our model on the training set of the Multi-PIE Setting 2, and evaluate it on the testing set of the Multi-PIE Setting 2 and the LFW [9].

Figure 4 shows the face synthesized results under large poses, and it is obvious that our model can synthesize photo-realistic images. To demonstrate the illumination preserving strength of our model, we choose the profiles with obvious inconsistent illumination. As shown in Fig. 4, the illumination of profile faces can be well preserved in the synthesized images. More synthesized results are provided in the supplementary material.

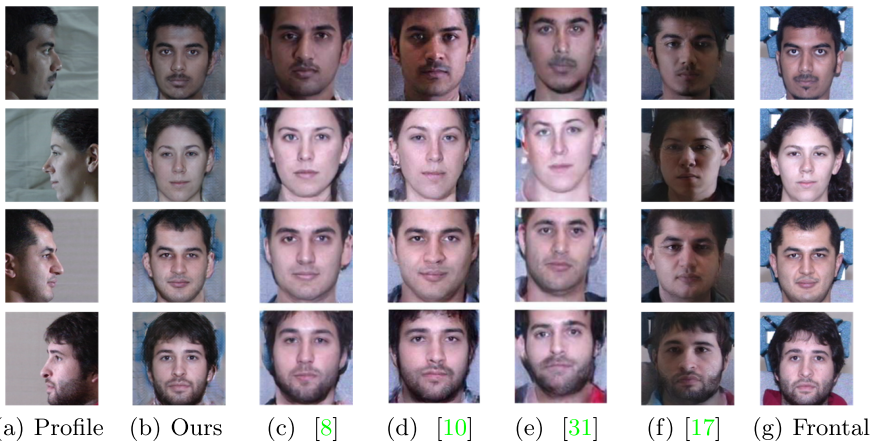


Fig. 5. Face frontalization comparison on the Multi-PIE dataset under the pose of 90° (first two rows) and 75° (last two rows).

Figure 5 illustrates the comparison with the state-of-the-art face frontalization methods [8, 10, 17, 31] on the Multi-PIE dataset. In the large pose cases, existing methods are disabled to preserve the illumination of profiles on the synthesized results. Face shape and other face components (*e.g.*, eyebrows, mustache and nose) also occur deformation. The reason is those methods are less able to preserve reliable details from the profiles. Compared with the existing methods, our method produces more identity preserving results while keeping the facial details of the profiles as much as possible. In particular, under large poses, our model can recover photo-realistic illumination conditions of the profiles, which is important when frontalized images are used for some other face-related tasks, such as face editing, face pose transfer and face-to-face synthesis.

We further qualitatively compare face frontalization results of our model on the LFW dataset with [6, 10, 17, 28, 34]. As shown in Fig. 6, the existing methods fail to recover clear global structures and fine facial details. Also they cannot




Fig. 6. Face frontalization comparison on the LFW dataset. Our method is trained on Multit-PIE and tested on LFW.

Table 1. Rank-1 recognition rates (%) across poses under Setting 2 of the Multi-PIE. The best two results are highlighted by **bold** and underline respectively.

Method	$\pm 15^\circ$	$\pm 30^\circ$	$\pm 45^\circ$	$\pm 60^\circ$	$\pm 75^\circ$	$\pm 90^\circ$	Avg
Light CNN [30]	98.59	97.38	92.13	62.09	24.18	5.51	63.31
DR-GAN [28]	94.90	91.10	87.20	84.60	–	–	89.45
FF-GAN [31]	94.60	92.50	89.70	85.20	77.20	61.20	83.40
TP-GAN [10]	98.68	98.06	95.38	87.72	77.43	64.64	86.99
CAPG-GAN [8]	99.82	99.56	97.33	90.63	83.05	66.05	89.41
PIM [34]	99.30	99.00	98.50	98.10	95.00	86.50	96.07
3D-PIM [35]	99.64	99.48	98.81	98.37	95.21	86.73	96.37
DA-GAN [32]	<u>99.98</u>	<u>99.88</u>	99.15	97.27	93.24	81.56	95.18
HF-PIM [1]	99.99	99.98	99.98	99.14	<u>96.40</u>	<u>92.32</u>	<u>97.97</u>
Ours	99.86	99.80	<u>99.37</u>	<u>98.85</u>	97.20	93.17	98.04

preserve the illumination of the profiles. Though FNM [17] generates high qualitative images, it is still disable to preserve identity. It is worth noting that our method produces more photo-realistic faces with identity and illumination well-preserved, which also demonstrates the generalizability of our model in the uncontrolled environment. More results under large poses are provided in the supplementary material.

4.3 Quantitative Evaluation

In this subsection, we quantitatively compare the proposed method with other methods in terms of recognition accuracy on Multi-PIE and LFW. The recognition accuracy is calculated by firstly extracting deep features with LightCNN-29 [30] and then measuring similarity between features with a cosine-distance metric.

Table 1 shows the Rank-1 recognition rates of different methods under Setting 2 of Multi-PIE. Our method has advantages over competitors, especially at large poses (*e.g.*, 75°, 90°), which demonstrates that our model can synthesize frontal images while preserving the identity information. The recognition rates under Setting 1 is provided in the supplementary material.

Table 2. Face verification accuracy (ACC) and area-under-curve (AUC) results on LFW.

Method	FaceNet [23]	VGG Face [16]	FF-GAN [31]	CAPG-GAN [8]	DA-GAN [32]	Ours
ACC(%)	99.63	98.95	96.42	99.37	99.56	99.65
AUC(%)	–	–	99.45	99.90	99.91	99.92

Table 2 compares the face verification performance (ACC and AUC) of our method with other state-of-the-arts [8, 16, 23, 31, 32] on the LFW. Our method achieves 99.65 on accuracy and 99.92 on AUC, which is also comparable with other state-of-the-art methods. The above quantitative results prove that our method is able to preserve the identity information effectively.

4.4 Ablation Study

In this subsection, we analyze the respective roles of the different modules and loss functions in frontal view synthesis. Both qualitative perceptual performance (Fig. 7) and face recognition rates (Table 3) are reported for comprehensive comparison under the Multi-PIE Setting 2. We can see that our FFWM exceeds all its variants in both quantitative and qualitative evaluations.

Effects of the Illumination Preserving Module (IPM). Although without IPM the recognition rates drop slightly (as shown in Table 3), the synthesized results cannot preserve illumination and are approximate to the inconsistent ground-truth illumination (as shown in Fig. 7). We also explore the contributions




Fig. 7. Model comparison: synthesis results of our model and its variants on Multi-PIE

Table 3. Incomplete variants analysis: Rank-1 recognition rates (%) across poses under Setting 2 of the Multi-PIE dataset. IAP and IPP denote the illumination adaption pathway and illumination preserving pathway in the Illumination Preserving Module (IPM). Warp, flip and att denote the three variants in Warp Attention Module (WAM).

Method	$\pm 15^\circ$	$\pm 30^\circ$	$\pm 45^\circ$	$\pm 60^\circ$	$\pm 75^\circ$	$\pm 90^\circ$	Avg
w/o IPM	99.83	99.77	99.35	98.74	97.18	93.15	98.00
IPM w/o IPP	99.84	99.74	99.36	98.47	96.73	91.56	97.62
IPM w/o IAP	99.83	99.76	99.30	98.70	97.11	92.83	97.92
w/o WAM	99.84	99.46	98.91	97.27	93.18	86.23	95.81
WAM w/o flip	99.84	99.69	99.27	98.10	96.57	92.65	97.69
WAM w/o warp	99.83	99.64	99.16	97.83	94.60	88.16	96.54
WAM w/o att	99.85	99.79	99.36	98.71	96.81	93.09	97.94
w/o \mathcal{L}_{adv}	99.83	99.72	99.28	98.57	97.09	93.11	97.93
w/o \mathcal{L}_{id}	99.85	99.62	99.12	97.42	93.93	86.05	96.00
w/o \mathcal{L}_{pixel}	99.83	99.77	99.35	98.79	97.05	92.85	97.94
w/o \mathcal{L}_p	99.81	99.75	99.33	98.62	97.13	93.10	97.96
Ours	99.86	99.80	99.37	98.85	97.20	93.17	98.04

of illumination adaption pathway (IAP) and illumination preserving pathway (IPP) in the IPM. As shown in Fig. 7, without IPP, the illumination of synthesized images tend to be inconsistent with the profiles and ground-truth images. And without IAP, the illumination of synthesized images tends to be a tradeoff between the profiles and the illumination inconsistent ground-truth images. Only integrating IPP and IAP together, our model can achieve illumination preserving image synthesis. Furthermore, our model archives a lower recognition rate when removing the IPP, which demonstrates that the IPP prompts the synthesized results to keep reliable information of the profiles.

Effects of the Warp Attention Module (WAM). We can see that without WAM, the synthesized results tend to be smooth and distorted in the self-occlusion parts (as shown in Fig. 7). As shown in Table 3, without WAM, the recognition rates drop significantly, which proves that WAM dominates in preserving identity information. Moreover, we explore the contributions of three components in the WAM, including taking flipped warped feature as additional input (w/o flip), feature warping (w/o warp) and feature attention (w/o att). As shown in Fig. 7, taking flip feature as additional input has benefits on recovering the self-occlusion parts on the synthesized images. Without the feature attention mechanism, there are artifacts on the synthesized images. Without feature warping, the synthesized results get worse visual performance. These results above suggest that each component in WAM is essential for synthesizing identity preserving and photo-realistic frontal images.

Effects of the Losses. As shown in Table 3, the recognition rates decrease if one loss function is removed. Particularly, the rates drop significantly for all poses if the \mathcal{L}_{id} loss is not adapted. We also report the qualitative visualization results in Fig. 7. Without \mathcal{L}_{adv} loss, the synthesized images tend to be blurry, suggesting the usage of adversarial learning. Without \mathcal{L}_{id} and \mathcal{L}_{pixel} , our model cannot promise the visual performance on the local textures (*e.g.*, eyes). Without \mathcal{L}_p , the synthesized faces present artifacts at the edges (*e.g.*, face and hair).

5 Conclusion

In this paper, we propose a novel Flow-based Feature Warping Model (FFWM) to effectively address the challenging problem in face frontalization, photo-realistic and illumination preserving image synthesis with illumination inconsistent supervision. Specifically, an Illumination Preserve Module is proposed to address the illumination inconsistent issue. It helps FFWM to synthesize photo-realistic frontal images while preserving the illumination of profile images. Furthermore, the proposed Warp Attention Module reduces the pose discrepancy in the feature space and helps to synthesize frontal images effectively. Experimental results demonstrate that our method not only synthesizes photo-realistic and illumination preserving results but also outperforms state-of-the-art methods on face recognition across large poses.

Acknowledgement. This work is partially supported by the National Natural Science Foundation of China (NSFC) under Grant Nos. 61671182 and U19A2073.

References

1. Cao, J., Hu, Y., Zhang, H., He, R., Sun, Z.: Learning a high fidelity pose invariant model for high-resolution face frontalization. In: Advances in Neural Information Processing Systems, pp. 2867–2877 (2018)
2. Deng, J., Cheng, S., Xue, N., Zhou, Y., Zafeiriou, S.: UV-GAN: adversarial facial UV map completion for pose-invariant face recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7093–7102 (2018)
3. Dosovitskiy, A., et al.: FlowNet: learning optical flow with convolutional networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2758–2766 (2015)
4. Gross, R., Matthews, I., Cohn, J., Kanade, T., Baker, S.: Multi-pie. Image Vis. Comput. **28**(5), 807–813 (2010)
5. Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: MS-CELEB-1M: a dataset and benchmark for large-scale face recognition. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9907, pp. 87–102. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_6
6. Hassner, T., Harel, S., Paz, E., Enbar, R.: Effective face frontalization in unconstrained images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4295–4304 (2015)

7. He, K., Sun, J., Tang, X.: Guided image filtering. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010. LNCS, vol. 6311, pp. 1–14. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15549-9_1
8. Hu, Y., Wu, X., Yu, B., He, R., Sun, Z.: Pose-guided photorealistic face rotation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8398–8406 (2018)
9. Huang, G.B., Ramesh, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: a database for studying face recognition in unconstrained environments. Technical report 07–49, University of Massachusetts, Amherst, October 2007
10. Huang, R., Zhang, S., Li, T., He, R.: Beyond face rotation: global and local perception GAN for photorealistic and identity preserving frontal view synthesis. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2439–2448 (2017)
11. Hui, T.W., Tang, X., Change Loy, C.: LiteFlowNet: a lightweight convolutional neural network for optical flow estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8981–8989 (2018)
12. Hur, J., Roth, S.: Iterative residual refinement for joint optical flow and occlusion estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5754–5763 (2019)
13. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2462–2470 (2017)
14. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in Neural Information Processing Systems, pp. 2017–2025 (2015)
15. Li, X., Liu, M., Ye, Y., Zuo, W., Lin, L., Yang, R.: Learning warped guidance for blind face restoration. In: Proceedings of the European Conference on Computer Vision, pp. 272–289 (2018)
16. Parkhi, O.M., Vedaldi, A., Zisserman, A.: Deep face recognition (2015)
17. Qian, Y., Deng, W., Hu, J.: Unsupervised face normalization with extreme pose and expression in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9851–9858 (2019)
18. Ranjan, A., Black, M.J.: Optical flow estimation using a spatial pyramid network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4161–4170 (2017)
19. Ren, Y., Yu, X., Chen, J., Li, T.H., Li, G.: Deep image spatial transformation for person image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7690–7699 (2020)
20. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
21. Russakovsky, O., et al.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**(3), 211–252 (2015)
22. Sagonas, C., Panagakis, Y., Zafeiriou, S., Pantic, M.: Robust statistical face frontalization. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 3871–3879 (2015)
23. Schroff, F., Kalenichenko, D., Philbin, J.: FaceNet: a unified embedding for face recognition and clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 815–823 (2015)

24. Shocher, A., Bagon, S., Isola, P., Irani, M.: InGAN: capturing and remapping the “DNA” of a natural image. arXiv preprint [arXiv:1812.00231](https://arxiv.org/abs/1812.00231) (2018)
25. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
26. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: PWC-NET: CNNs for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8934–8943 (2018)
27. Tian, Y., Peng, X., Zhao, L., Zhang, S., Metaxas, D.N.: CR-GAN: learning complete representations for multi-view generation. arXiv preprint [arXiv:1806.11191](https://arxiv.org/abs/1806.11191) (2018)
28. Tran, L., Yin, X., Liu, X.: Disentangled representation learning GAN for pose-invariant face recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1415–1424 (2017)
29. Truong, P., Danelljan, M., Timofte, R.: GLU-NeT: global-local universal network for dense flow and correspondences. arXiv preprint [arXiv:1912.05524](https://arxiv.org/abs/1912.05524) (2019)
30. Wu, X., He, R., Sun, Z., Tan, T.: A light CNN for deep face representation with noisy labels. *IEEE Trans. Inf. Forensics Secur.* **13**(11), 2884–2896 (2018)
31. Yin, X., Yu, X., Sohn, K., Liu, X., Chandraker, M.: Towards large-pose face frontalization in the wild. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 3990–3999 (2017)
32. Yin, Y., Jiang, S., Robinson, J.P., Fu, Y.: Dual-attention GAN for large-pose face frontalization. arXiv preprint [arXiv:2002.07227](https://arxiv.org/abs/2002.07227) (2020)
33. Zhang, Z., Chen, X., Wang, B., Hu, G., Zuo, W., Hancock, E.R.: Face frontalization using an appearance-flow-based convolutional neural network. *IEEE Trans. Image Process.* **28**(5), 2187–2199 (2018)
34. Zhao, J., et al.: Towards pose invariant face recognition in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2207–2216 (2018)
35. Zhao, J., et al.: 3D-aided deep pose-invariant face recognition. In: *IJCAI*, vol. 2, p. 11 (2018)
36. Zhu, X., Lei, Z., Yan, J., Yi, D., Li, S.Z.: High-fidelity pose and expression normalization for face recognition in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 787–796 (2015)