

# 다중 측면 얼굴 이미지를 이용한 개선된 얼굴 정면화 기법

최원영<sup>o</sup>                      고흥석<sup>\*</sup>

서울대학교 전기정보공학부  
{cwy9897, hsko}@graphics.snu.ac.kr

## Improved Face Frontalization Using Multiple Side-Face Images

Wonyoung Choi<sup>o</sup>                      Hyeong-Seok Ko<sup>\*</sup>

Department of Electrical and Computer Engineering, Seoul National University

### 요약

측면 얼굴 사진을 입력으로 정면 얼굴을 생성하는 연구는 얼굴 인식 분야에서 중요한 주제 중 하나이다. 이를 해결하기 위해, 본 연구에서는 여러 측면 얼굴 이미지를 입력으로 활용하여 성능을 향상시키는 방법을 제안하였다. 여러 각도의 얼굴 이미지를 사용함으로써, 특히 90도와 같은 큰 각도에서도 정확하게 정면 얼굴을 재구성할 수 있는 능력을 향상시켰다. 다중 이미지를 사용함으로써 단일 이미지 방식에서 놓칠 수 있는 세부 정보를 포착할 수 있어, 더 정확하고 신원 보존적인 정면 이미지를 생성할 수 있었다. 실험 결과, 본 연구의 방법이 기존 방식보다 성능이 크게 향상되었으며, 특히 극단적인 각도를 처리할 때 뛰어난 성능을 보였다.

### Abstract

The task of generating frontal face images is a crucial challenge in the field of facial recognition. To address this, we propose a method that improves performance by leveraging multiple side-face images as input. By using multiple views of the face, our approach enhances the ability to accurately reconstruct the frontal face, even in cases with large pose variations, such as 90-degree angles. The use of multiple images allows for the capture of details that are otherwise missed in single-image methods, leading to more precise and identity-preserving frontal images. Experimental results show that our method significantly outperforms existing approaches, especially in handling extreme pose angles, demonstrating its effectiveness in improving face frontalization performance.

**키워드:** 얼굴 정면화, 얼굴 인식, 심층신경망

**Keywords:** Face frontalize, Face recognition, Deep learning

## 1 Introduction

Face frontalization, which aims to synthesize a frontal view of a face from a non-frontal image, has been an essential task in computer vision and face recognition systems. Pose variations are one of the most significant challenges in face recognition, as non-frontal views often fail to provide complete and consistent facial information. This issue is particularly critical in real-world applications such as surveillance systems, where cameras frequently capture non-frontal faces, and in photo tagging or identity veri-

fication scenarios in uncontrolled environments. By generating a frontal view of the face, face frontalization enhances the accuracy of face recognition and facilitates tasks such as facial attribute analysis, emotion recognition, and virtual reality applications.

A well-known approach is the Face Normalization Model (FNM) [1], which uses a single side-face image to generate a frontal view. This model has been effective in many scenarios but struggles when dealing with extreme pose variations or occlusions. Another notable method is the Two-Pathway Generative Adversar-

<sup>\*</sup>corresponding author: Hyeong-Seok Ko/Seoul National University(hsko@graphics.snu.ac.kr)

ial Network (TP-GAN) [2], which similarly uses a single image to generate a frontal face. TP-GAN excels in preserving identity features, which refer to the unique characteristics of a face, such as the shape of facial structures, relative positions of key points (eyes, nose, mouth), and texture patterns that distinguish one individual from another. These features are critical for ensuring that the generated frontal image maintains the same identity as the input image. However, like FNM, TP-GAN faces challenges with large pose discrepancies and missing facial details due to occlusions in the side-view images.

These methods, while successful in frontalizing faces from a single input image, have inherent limitations. Single-image frontalization can often produce incomplete or inaccurate reconstructions, especially when the input face is at extreme angles (e.g., 90 degrees) or when parts of the face are occluded. These limitations arise because a single image cannot provide sufficient information about the hidden parts of the face, resulting in loss of critical features during frontalization.

To address these issues, the Disentangled Representation learning Generative Adversarial Network (DR-GAN) [3] was introduced, which allows the use of multiple images as input to improve the robustness of the frontalization process. DR-GAN takes advantage of multiple views of a face, learning from various poses and expressions to generate a more accurate frontal image. However, the model has a significant drawback: it requires the same number of multiple images during the training phase as well, making it impractical when only a single or fewer images are available for some identities during training. Additionally, DR-GAN does not effectively handle extreme angles, such as 90-degree side views, which can lead to degraded frontalization performance.

In this paper, we propose a new approach to face frontalization that overcomes these limitations. Our method allows the model to be trained on varying numbers of input images, removing the restriction of needing the same number of images during training and testing. Additionally, our method explicitly handles extreme facial angles, such as 90-degree side views, improving frontalization performance in these challenging cases. By using multiple side-face images as input and introducing a flexible architecture, we ensure that our model produces high-quality, identity-preserving frontal images, even under extreme conditions.

## 2 Related Works

Face frontalization has been studied extensively, and various methodologies have been explored to tackle this challenging problem. Broadly, these approaches can be categorized into traditional 3D modeling methods, deep learning techniques (especially those

based on GANs), and more recent attempts that utilize multiple images for frontalization.

### 2.1 3D Morphable Models (3DMM)

One of the earliest and most well-established methods for face frontalization is based on 3D Morphable Models (3DMM) [4]. In this approach, a 3D model of a face is constructed from a non-frontal image, which is then used to synthesize a frontal view. While 3DMM has proven effective in various applications, it has certain limitations. To achieve high-quality frontalization, the method requires precise data acquisition and significant computational resources, making it less feasible for real-time applications or cases with lower-quality input data. The optimization process of fitting the 3D model to the input image can be computationally expensive and time-consuming, hindering its scalability for large datasets or practical deployments.

### 2.2 GAN-based Face Frontalization

With the advent of deep learning, particularly Generative Adversarial Networks (GANs), face frontalization has seen substantial improvements in both speed and accuracy. GANs are well-suited for generating realistic frontal views from side-view images by learning the mapping between poses and the corresponding frontal images in a data-driven manner.

Among the notable GAN-based approaches are TP-GAN [2] and FNM [1]. TP-GAN uses a two-pathway structure to preserve both local and global features during frontalization, ensuring that identity features are maintained. FNM, on the other hand, focuses on normalizing a side-face image to a frontal view by disentangling pose and identity features, allowing for better identity-preserving frontalization. Both methods demonstrate the power of GANs in handling pose variations and generating realistic frontal images from single side-face images. However, as discussed earlier, these methods are limited in cases where extreme angles or occlusions are present, as they rely on information from a single input image.

### 2.3 Multi-Image Face Frontalization

DR-GAN, introduced by [3], aimed to improve face frontalization by utilizing multiple images as input, learning from different poses to generate a more accurate frontal image. DR-GAN’s encoder-decoder structure enabled the generation of a pose-invariant identity representation, and the model showed promising results in frontalizing faces from various angles. However, this approach also introduced some limitations. Notably, when multiple images are used as input, DR-GAN requires the same number of images

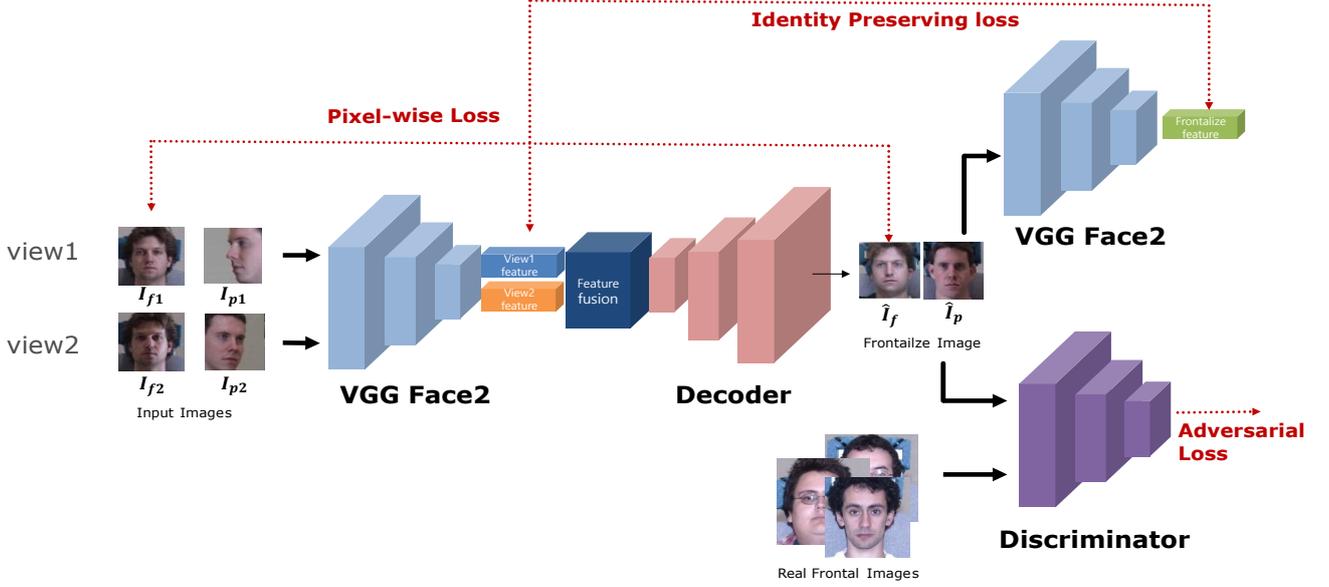


Figure 1: Overall network architecture of our proposed method, Multi-Image FNM.

during the training phase, which reduces flexibility, especially in scenarios where fewer images are available for certain identities. Moreover, DR-GAN struggles to handle extreme pose variations, such as 90-degree side profiles, leading to degraded performance in such challenging cases.

### 3 Methodology

In our method, we extend the original FNM architecture to handle multiple input images. This modified architecture, which we refer to as the Multi-Images FNM, is designed to take advantage of two side-face images, improving the robustness of the frontalization process by combining features from different perspectives.

#### 3.1 Multi-Images FNM

Our frontalization module consists of three main components: the encoder  $E$ , the decoder  $D$ , and the feature fusion module.

- The encoder  $E$  is based on VGGFace2 [5], one of the pre-trained face recognition models. This encoder transforms a side-face image of size  $224 \times 224 \times 3$  into a 2048-dimensional feature vector  $f$ . The use of a pre-trained encoder ensures that we can capture identity-related features effectively, even when the input images are captured from extreme angles. This effectiveness arises because the encoder has been trained on a large-scale dataset with diverse identities and pose variations, enabling it to learn robust and discriminative features that generalize well to unseen scenarios. Mathematically, for

an input image  $I$ , the transformation can be written as:

$$f_i = E(I_i) \quad \text{where} \quad I_i \in \mathbb{R}^{224 \times 224 \times 3}, f_i \in \mathbb{R}^{2048}$$

- The feature fusion module is responsible for integrating the multiple features extracted from different input images into a single, unified feature representation. A detailed explanation of this fusion process will be covered in Section 3.2.
- The decoder takes the fused 2048-dimensional feature vector and generates a frontal image  $I_f$  of size  $224 \times 224 \times 3$ . This process essentially reverses the encoding operation, reconstructing the frontal image from the high-level feature representation. This can be represented as:

$$I_f = D(f) \quad \text{where} \quad I_f \in \mathbb{R}^{224 \times 224 \times 3}$$

Thus, the overall process of the Multi-Images FNM architecture can be summarized by encoding multiple side-view images, fusing their features, and finally decoding them into a frontalized image, as shown in Figure 1.

#### 3.2 Feature Fusion Module

The feature fusion module is a key component of our architecture, designed to combine multiple feature vectors into a single, unified representation. This module plays a critical role in ensuring that information from different side-view images is effectively integrated to generate a high-quality frontal face.

As shown in Figure 2, we experimented with two fusion tech-

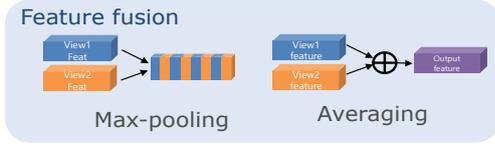


Figure 2: Two fusion methods used for feature fusion module.

niques: max pooling and average pooling. These two pooling methods are widely used in deep learning for dimensionality reduction and feature aggregation, allowing us to merge the extracted features from multiple input images efficiently.

- **Max Pooling:** For each corresponding element in the feature vectors, we select the maximum value across all input features. This method captures the most dominant feature in each dimension, ensuring that the strongest signals from each input image are retained in the final feature representation.
- **Average Pooling:** We compute the average value for each corresponding element in the feature vectors. This method helps in smoothing out the feature differences between images and results in a more generalized representation that incorporates information from all inputs equally.

In Figure 2, we show how these two fusion techniques combine two feature vectors into one. However, one of the reasons we selected these pooling methods is their flexibility in handling  $N$  input features. Whether the input consists of two, three, or more side-face images, the fusion module can seamlessly integrate any number of feature vectors, making the system highly adaptable to various input conditions.

This flexibility ensures that our model can generalize to scenarios where different numbers of side-view images are available, enhancing the robustness of the frontalization process.

## 4 Training Strategy

In this section, we describe the training strategy used for our proposed module. During training, a total of four images are input into the model. These consist of two random frontal images from the frontal image set and two random profile images from the profile image set which is a non-frontal face image set, each from a different subject.

Let  $I_{f1}$  and  $I_{f2}$  represent the two random frontal images from the same subject, and  $I_{p1}$  and  $I_{p2}$  represent the two random profile images from another same subject. These four images are processed to generate two frontalized outputs, as shown in Figure 1. Specifically, the model generates the frontalized face  $\hat{I}_f$  from  $I_{f1}$  and  $I_{f2}$ , and the frontalized face  $\hat{I}_p$  from  $I_{p1}$  and  $I_{p2}$ .

Using two random frontal images  $I_{f1}$  and  $I_{f2}$  to generate a frontalized image during training is crucial for ensuring the stability of the training process. Without this component, the training process often fails to converge. Generating a frontalized image from the same frontal image allows the model to converge quickly but results in overfitting. To address this, we included the generation of a frontalized image using two different frontal face images, which may differ in lighting conditions or minor variations. This approach improves the stability of the training process and prevents overfitting during training.

### 4.1 Loss function

The loss function used in our model consists of three components: pixel-wise loss, identity-preserving loss, and adversarial loss. Each of these losses plays a crucial role in ensuring that the generated frontal images are accurate, realistic, and identity-preserving. Below, we explain each component in detail.

- **Pixel-wise Loss:** This loss is calculated between the generated image  $\hat{I}_f$  and the input frontal images  $I_{f1}$  and  $I_{f2}$ . Since  $I_{f1}$  and  $I_{f2}$  are already frontal images, the generated image is directly compared to these input images. The pixel-wise loss is defined as:

$$\mathcal{L}_{pixel} = \|\hat{I}_f - I_{f1}\|_1 + \|\hat{I}_f - I_{f2}\|_1$$

This component of the loss function encourages the model to generate pixel-wise images similar to the ground truth, improving visual quality and ensuring stable and accurate training, as also mentioned in FNM.

- **Identity Preserving Loss:** This loss is a type of perceptual loss [6] that ensures the identity of the generated frontalized image matches the identity of the input image. In face frontalization, it is crucial that the identity in the generated frontal image remains consistent with the input profile image. This loss works by minimizing the distance between the feature space representations of the input profile images  $I_{p1}$ ,  $I_{p2}$  and the generated frontalized image  $\hat{I}_p$ , both extracted using a pre-trained face recognition model. The identity-preserving loss can be formulated as:

$$\mathcal{L}_{ip} = \|\phi(\hat{I}_p) - \phi(I_{p1})\|_2^2 + \|\phi(\hat{I}_p) - \phi(I_{p2})\|_2^2$$

where  $\phi$  represents the feature extraction function of the pre-trained face recognition model.

- **Adversarial Loss:** This loss is based on the standard GAN loss [7], which is used to ensure that the generated frontalized

image is indistinguishable from real frontal images. In a GAN framework, the generator aims to produce images that are realistic enough to fool the discriminator, while the discriminator tries to distinguish between real and generated images. The adversarial loss for the generator can be expressed as:

$$\mathcal{L}_{adv} = \mathbb{E}_{I_{gt}} [\log D_{is}(x)] + \mathbb{E}_{\hat{f}} [1 - \log D_{is}(x)]$$

where:

$\mathbb{E}_{I_{gt}}$  : Average over real frontal images.

$\mathbb{E}_{\hat{f}}$  : Average over generated frontalized images.

$D_{is}(x)$  : The discriminator function.

We employ the vanilla GAN loss function [7] to assess the performance of our model. Despite recent advancements in GANs, such as WGAN [8, 9], which show superior performance, using them for our purposes presents no issues.

The overall loss function was the weighted sum of the three aforementioned loss functions. This can be formulated as follows:

$$\mathcal{L} = \lambda_{pixel} \mathcal{L}_{pixel} + \lambda_{ip} \mathcal{L}_{ip} + \lambda_{adv} \mathcal{L}_{adv}$$

## 5 Experimental Result

### 5.1 Experimental Datasets

We utilized the CMU Multi-PIE face dataset [10] for both our training and testing sets. The Multi-PIE dataset contains over 750,000 images of 337 subjects, captured across 15 different viewpoints and 19 illumination conditions, with various facial expressions. This dataset is commonly employed for assessing face synthesis and recognition in controlled environments. Consistent with previous face frontalization research [1, 2], we adopted setting 2 to evaluate our model. Setting 2 involves using neutral expression images from all four sessions, encompassing 337 identities. We used images of the first 200 identities across 11 poses for training. For testing, a frontal view image under standard illumination was chosen as the gallery image for each of the remaining 137 identities, while the rest of the images were used as probe images.

#### 5.1.1 Implementation Details

For both training and testing, the face images were aligned using the MTCNN face detector [11], followed by cropping to a resolution of  $224 \times 224$  pixels. During the training phase, we applied the Adam optimizer with the following hyperparameters:  $lr = 10^{-4}$ ,  $\beta_1 = 0.5$ ,  $\beta_2 = 0.99$ ,  $\lambda_{pixel} = 1$ ,  $\lambda_{ip} = 1$ ,  $\lambda_{adv} = 0.1$ .

Table 1: Rank-1 recognition rate (%) performance comparison between the baseline FNM and the proposed Multi-Image FNM across different numbers of input images (n). The proposed Multi-Image FNM was trained using two input images.

number of input images	FNM [1]	Multi-Image FNM (maxpolling)	Multi-Image FNM (avgpooling)
n=1	87.39	87.39	87.39
n=2	-	94.47	93.91
n=3	-	95.63	95.81
n=4	-	95.82	96.61
n=5	-	95.89	97.15
n=6	-	95.94	97.28
n=7	-	<b>96.01</b>	97.43
n=8	-	95.94	<b>97.63</b>

## 5.2 Results

Akin to previous studies, we assessed the face frontalization performance using the rank-1 recognition rate. This metric is computed by measuring the cosine distance between the feature vectors extracted from the generated frontal faces and the gallery images of the corresponding identities, using a pre-trained face recognition network.

Our model was trained using two input images, which enables the network to learn how to effectively combine information from multiple perspectives to generate high-quality frontal images. Since our method utilizes multiple input images, we compare our results with the baseline method, FNM, which is limited to using a single input image (n=1). As illustrated in Table 1, we report recognition rates for varying numbers of input images, from n=2 to n=8, to evaluate the effectiveness of incorporating additional input images. The input images in our model are randomly selected from various angles, without restriction, ensuring that the model is exposed to diverse perspectives. Table 1 represents the results of testing a single trained model, which was trained using two input images, with varying numbers of input images (n=2 to n=8). The baseline FNM model achieves a rank-1 recognition rate of 87.39% with only one input image, whereas our proposed Multi-Image FNM shows a clear improvement in performance as the number of input images increases. For n=2, the recognition rate improves significantly, reaching 94.47% with max pooling and 93.91% with average pooling. This trend continues until n=7, where the highest recognition rates are observed: 96.01% for max pooling and 97.43% for average pooling. For n=8, max pooling remains at 95.94%, and average pooling slightly improves to 97.63%, indicating that while additional images contribute to improved frontalization, the gains may plateau beyond a certain point. The comparison between max pooling and average pooling also highlights differences in performance. While both pooling strategies show strong results, average pooling tends to perform better as the number

Table 2: Rank-1 recognition rate (%) performance comparison between the baseline FNM and the proposed Multi-Image FNM across varying pose angles.

	$-15^\circ + 15^\circ$	$-30^\circ + 30^\circ$	$-45^\circ + 45^\circ$	$-60^\circ + 60^\circ$	$-75^\circ + 75^\circ$	$-90^\circ + 90^\circ$
FNM (1-view)	97.47	96.17	94.48	89.02	78.91	68.56
Multi-Image FNM (maxpooling)	98.23	98.12	97.30	94.53	86.45	76.73
Multi-Image FNM (avgpooling)	98.12	97.77	96.97	94.03	85.78	75.62

of input images increases, particularly for  $n=3$  and beyond. This suggests that average pooling may be more effective in capturing complementary information from multiple input images, leading to higher recognition accuracy

Additionally, our method benefits from using multiple input images, allowing us to input face images from opposite directions (e.g.,  $+90$  and  $-90$  degrees), which is not possible when using a single image. In the case of single-image input, the opposite side of the face is not visible, limiting the model’s ability to fully reconstruct the frontal view. As shown in Table 2, by leveraging images from opposite angles, our multi-image input approach significantly enhances performance. Table 2 and Figure 3 illustrate the performance of the same model, trained using two input images when tested with two input images. We compare these results against the baseline FNM to highlight the performance improvements enabled by our approach.

One key advantage of our method is its ability to handle the unseen parts of the face that are missing in single-image frontalization. When only one image is used, crucial facial details on the opposite side are not visible, limiting the model’s ability to reconstruct a complete frontal view. However, by utilizing multiple side-view images from different angles, our method is able to capture and synthesize the missing details, resulting in significantly better performance, especially for extreme angles. This advantage is evident across various pose angles, as reflected in the improved recognition rates.

In addition to the quantitative results, Figure 3 provides a visual comparison between the baseline FNM and our proposed Multi-Image FNM. The figure demonstrates how our method outperforms the baseline, particularly in cases with extreme angles. The first two rows display the side-face images (Input1 and Input2) used as input, and the third row shows the results from the baseline FNM, which struggles to generate accurate frontal images from a single input. In contrast, the fourth and fifth rows show the results from our Multi-Image FNM model with max pooling and average pooling, respectively. Our method consistently generates more realistic and identity-preserving frontal images, especially when using multiple side-face inputs, which provide crucial information about

the hidden parts of the face. The final row shows the ground truth frontal images for reference, further highlighting the performance improvement of our method over the baseline FNM.

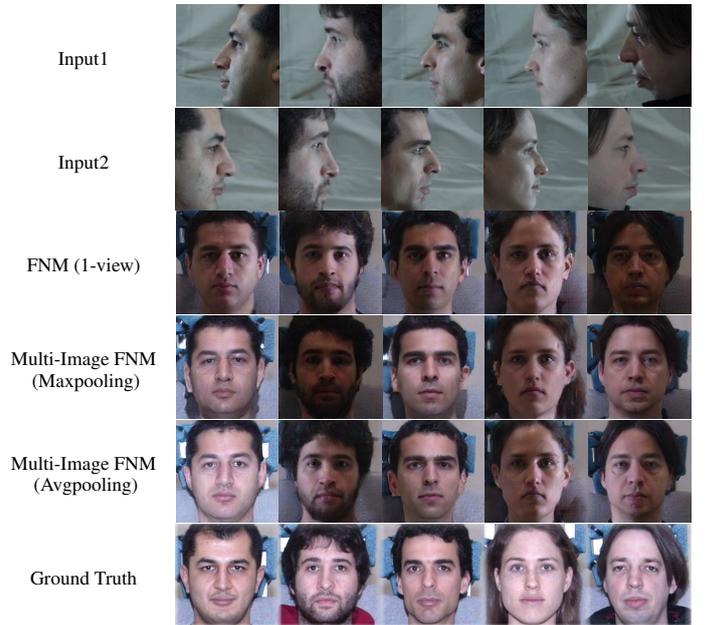


Figure 3: Visual comparison of face frontalization results between FNM and our proposed Multi-Image FNM.

### 5.3 Conclusion

In this paper, we introduced a novel face frontalization approach that leverages multiple side-face images as input to overcome the limitations of previous methods such as DR-GAN [3] and FNM [1]. Our method not only improves flexibility by allowing training with varying numbers of input images, but it also significantly enhances performance when handling extreme pose variations, such as 90-degree side profiles. By incorporating a feature fusion module and explicitly considering diverse angles, our approach generates more accurate and identity-preserving frontal images compared to single-image-based methods.

Through extensive experiments, we demonstrated that our model consistently outperforms the baseline methods, achieving higher recognition rates as the number of input images increases. The re-

sults also highlight the effectiveness of both max pooling and average pooling for feature fusion, with average pooling slightly outperforming in cases where more input images are available. Furthermore, our method shows strong robustness when reconstructing frontal faces from input images taken at extreme angles, proving its practical applicability in challenging scenarios.

Looking forward, our method holds potential for further improvements, including the incorporation of more advanced loss functions and exploring additional datasets with more diverse conditions. We believe that the flexibility and performance of our approach will contribute to the advancement of face frontalization techniques and their applications in real-world face recognition systems.

## 감사의 글

This work was supported by ASRI (Automation and Systems Research Institute at Seoul National University).

## References

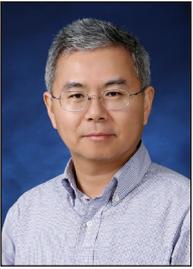
- [1] Y. Qian, W. Deng, and J. Hu, "Unsupervised face normalization with extreme pose and expression in the wild," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9851–9858.
- [2] R. Huang, S. Zhang, T. Li, and R. He, "Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2439–2448.
- [3] L. Tran, X. Yin, and X. Liu, "Disentangled representation learning gan for pose-invariant face recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1415–1424.
- [4] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3d faces," in *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '99. USA: ACM Press/Addison-Wesley Publishing Co., 1999, p. 187–194. [Online]. Available: <https://doi.org/10.1145/311535.311556>
- [5] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vggface2: A dataset for recognising faces across pose and age," in *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*. IEEE Press, 2018, p. 67–74. [Online]. Available: <https://doi.org/10.1109/FG.2018.00020>
- [6] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 694–711.
- [7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [8] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*. PMLR, 2017, pp. 214–223.
- [9] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," *Advances in neural information processing systems*, vol. 30, 2017.
- [10] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-pie," *Image and vision computing*, vol. 28, no. 5, pp. 807–813, 2010.
- [11] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, oct 2016. [Online]. Available: <https://doi.org/10.1109/2Flsp.2016.2603342>

## 〈 저자 소개 〉



### 최 원 영

- 2012년 서울대학교 전기공학부 학사 졸업
- 2012년 ~ 현재 서울대학교 전기정보공학부 석박통합과정 중
- 관심분야 : 실시간 의상 렌더링, 딥러닝, 이미지 생성
- <https://orcid.org/0009-0001-2438-6855>



### 고 형 석

- 1985년 서울대학교 계산통계학과 학부 졸업
- 1987년 서울대학교 계산통계학과 석사 졸업
- 1994년 미국 펜실베이니아 대학교 컴퓨터그래픽스전공 박사 졸업
- 1996 ~ 현재 서울대학교 전기정보공학부 교수
- 관심분야 : 디지털클로딩, 물리 기반 애니메이션
- <https://orcid.org/0009-0008-8068-5244>