

A UNIFIED FRAMEWORK FOR MASKED AND MASK-FREE FACE RECOGNITION VIA FEATURE RECTIFICATION

Shaozhe Hao¹ Chaofeng Chen² Zhenfang Chen³ Kwan-Yee K. Wong¹

¹The University of Hong Kong ²Nanyang Technological University ³MIT-IBM Watson AI Lab

ABSTRACT

Face recognition under ideal conditions is now considered a well-solved problem with advances in deep learning. Recognizing faces under occlusion, however, still remains a challenge. Existing techniques often fail to recognize faces with both the mouth and nose covered by a mask, which is now very common under the COVID-19 pandemic. Common approaches to tackle this problem include 1) discarding information from the masked regions during recognition and 2) restoring the masked regions before recognition. Very few works considered the consistency between features extracted from masked faces and from their mask-free counterparts. This resulted in models trained for recognizing masked faces often showing degraded performance on mask-free faces. In this paper, we propose a unified framework, named Face Feature Rectification Network (*FFR-Net*), for recognizing both masked and mask-free faces alike. We introduce rectification blocks to rectify features extracted by a state-of-the-art recognition model, in both spatial and channel dimensions, to minimize the distance between a masked face and its mask-free counterpart in the rectified feature space. Experiments show that our unified framework can learn a rectified feature space for recognizing both masked and mask-free faces effectively, achieving state-of-the-art results. Project code: <https://github.com/haosz/FFR-Net>

Index Terms— Face Recognition, Masked Face, Feature Rectification, COVID-19

1 Introduction

Face recognition has practical significance in our daily life. Recent advances of deep neural networks have resulted in face recognition systems [1, 2, 3] that achieve over 99% recognition accuracy under ideal conditions [4, 5]. However, in real-world scenarios, face recognition systems often suffer from performance degradation due to occlusion by sunglasses, masks, *etc.* In particular, masked face recognition has recently been a hot topic under the COVID-19 pandemic.

To recognize occluded faces, some researchers [6, 7] proposed to discard features from the occluded regions. Although this approach can avoid distraction caused by the corrupted features of the occluded regions, it also results in loss of global face information which might be inferred from the discarded regions. Others [8, 9] proposed to restore the occluded face regions before recognition. They

commonly used mean squared errors in the pixel space to train generative adversarial networks [10] to restore a face. Despite the fact that the restored face may be visually appealing, they fail to preserve identity information in the feature space. In fact, very few works considered the consistency between features extracted from occluded faces and from their occlusion-free counterparts. In this case, models trained for recognizing occluded faces often show degraded performance on occlusion-free faces.

In this paper, we specifically focus on masked face recognition and propose a unified framework for masked and mask-free face recognition. In our model, the corrupted features from masked faces can be mapped to a rectified feature space with better identity separability. We adopt SENet [11] trained on mask-free faces as our feature extractor. We introduce rectification blocks (RecBlocks) on top of extracted features. RecBlocks are trained to maximize the consistency between masked faces and their mask-free counterparts in the rectified feature space. Instead of directly transforming the extracted features, RecBlocks output a linear transformation matrix \mathbf{M} applied to the extracted features. Ideally, \mathbf{M} should be an identity matrix for mask-free faces and should suppress features from the masked regions for masked faces. We carry out rectifications in both spatial and channel dimensions using two parallel RecBlocks. In addition to the commonly used classification loss, we employ an identity loss and a triplet loss to train our rectification network end-to-end. In summary, our main contributions are:

- We introduce a unified framework, named Face Feature Rectification Network (*FFR-Net*), for masked and mask-free face recognition. Experiments show that our framework achieves the best average performance on mixed datasets (mask and mask-free faces) with a single model.
- We propose rectification blocks (RecBlocks) to rectify features of masked or mask-free faces in both spatial and channel dimensions. RecBlocks can maximize the consistency between masked faces and their mask-free counterparts in the rectified feature space.

2 Methodology

In this section, we will introduce the proposed feature rectification block and describe our training objectives in detail. Fig. 1 gives an overview of Face Feature Rectification Network (*FFR-Net*).




Fig. 1: Overview of FFR-Net. We first extract features of masked/mask-free faces with a pre-trained encoder. We then apply channel and spatial rectification blocks to rectify features to a unified space. Finally, the original and rectified features are fused and sent to the classifier. \textcircled{S} represents self-similarity, \otimes represents matrix multiplication and \oplus represents concatenation and fusion.

2.1 Channel-and-Space Rectification

Consider a training set $\{x_i^0, x_i^1\}_{i=1}^N$ with N pairs of mask-free faces x_i^0 and masked faces x_i^1 . Let F denote the face encoder that extracts a feature map f_i^j from x_i^j , i.e., $f_i^j = F(x_i^j)$, where $j \in \{0, 1\}$. R denotes a rectification module that maps f_i^j to \hat{f}_i^j in the rectified feature space, i.e., $\hat{f}_i^j = R(f_i^j)$. Our goal is to train a rectification module to perform the mapping $f_i^j \rightarrow \hat{f}_i^j$ so as to maximize the recognition accuracy for both x_i^0 and x_i^1 based on the rectified features \hat{f}_i^j . For simplicity, we omit the subscripts and superscripts i, j in this section.

To fully exploit the inner structure of encoded features, we conduct the rectification process in both channel dimension and spatial dimension, i.e., *ChnRec* and *SpcRec*. The main idea is to output a linear mapping of original encoded features $f \in \mathbb{R}^{C \times (H \times W)}$, where C and $H \times W$ denote channel and spatial size respectively. The rectification can be expressed as

$$\hat{f}_c = M_c \otimes f \quad \hat{f}_s = f \otimes M_s \quad (1)$$

where \otimes is matrix multiplication, $M_c \in \mathbb{R}^{C \times C}$ and $M_s \in \mathbb{R}^{(H \times W) \times (H \times W)}$ are channel and spatial rectification matrices. M_c and M_s are generated with *ChnRec* and *SpcRec*. The structures are described in Sec. 3. The final rectified feature \hat{f} is obtained by fusing \hat{f}_c, \hat{f}_s with the original feature f .

Besides, We also discover two important properties that are helpful to predict M : self-similarity and spatial symmetry.

1. *Self-similarity*. Aligned mask-free faces usually have similar facial structure, while this pattern can be corrupted due to occurrence of masks. Inspired by this observation, we introduce the self-similarity feature as extra input for *ChnRec* and *SpcRec*. Specifically, the self-similarity features are defined as the cosine similarity matrices between feature vectors in both spatial and channel dimension which can be formulated as below

$$S_{u,v}^s = \left\langle \frac{f_u}{\|f_u\|_2}, \frac{f_v}{\|f_v\|_2} \right\rangle, u, v \in \{0, 1, \dots, H \times W - 1\} \quad (2)$$

$$S_{m,n}^c = \left\langle \frac{f_m}{\|f_m\|_2}, \frac{f_n}{\|f_n\|_2} \right\rangle, m, n \in \{0, 1, \dots, C - 1\} \quad (3)$$

where $\langle \cdot, \cdot \rangle$ is inner product, $f_u \in \mathbb{R}^C$ is the u -th feature vector in spatial dimension, and $f_m \in \mathbb{R}^{H \times W}$ is the m -th feature vector in channel dimension. $S^s \in \mathbb{R}^{(H \times W) \times (H \times W)}$ and $S^c \in \mathbb{R}^{C \times C}$ are spatial and channel self-similarity matrices respectively. We concatenate S^s and S^c with original feature f to obtain $f^s \in \mathbb{R}^{(C + (H \times W)) \times H \times W}$ and $f^c \in \mathbb{R}^{C \times ((H \times W) + C)}$ as inputs for *ChnRec* and *SpcRec* respectively.

2. *Spatial Symmetry*. Channel rectification is aimed to select those channels which extract features from mask-free regions and discard those from masked parts. However, the selected channels might only contain information from single side of the face, e.g., the left eye or the right cheek. Since f^c ignores spatial information, channel rectification may not be able to preserve the spatial symmetry property of mask-free faces. Thus, we introduce flipped fusion block as the last step of *ChnRec*. The flipped fusion block merges the original feature with the horizontally flipped feature, which ensures outputs from *ChnRec* are still spatially symmetric.

2.2 Loss Definition

The rectified feature space is supposed to have three properties: 1) rectified features of the same identity, no matter masked or not, should give similar representations to the original mask-free one. 2) masked features should give a rectification tendency to approach original mask-free encoded features. 3) each pair of rectified features should produce the best classification result for the masked face recognition task. Accordingly, we leverage three losses for identity consistency, mask robustness and accurate classification.

Identity consistency. We propose 2 components in identity loss, namely identical facial representations (\mathcal{L}_{ID1}) and identical self-similarity patterns (\mathcal{L}_{ID2}). Identity loss aims to achieve: (1) rectify the masked feature and (2) keep the mask-free feature uncorrupted. Therefore, each component has 2 parts to simultaneously penalize masked and mask-free data and is formulated as follows:

$$\mathcal{L}_{ID1} = \frac{1}{N} \sum_i \|\hat{f}_i^1 - f_i^0\|_2 + \|\hat{f}_i^0 - f_i^0\|_2 \quad (4)$$

$$\mathcal{L}_{ID2} = \frac{1}{N} \sum_i \|\hat{S}_i^1 - S_i^0\|_2 + \|\hat{S}_i^0 - S_i^0\|_2 \quad (5)$$

where S_i and \hat{S}_i are respectively the self-similarity matrices of the i -th original and rectified features f_i and \hat{f}_i . Superscript $\{0, 1\}$ represents the mask-free and masked data respectively.

Mask robustness. A variant triplet loss is adopted. Specifically, for a masked face sample, triplet loss minimizes the distance between its rectified representation (the anchor f_i^a) and the original representation of its mask-free counterpart (the positive f_i^p) while maximizing the distance between the rectified representation (the anchor f_i^a) and the original representation of itself (the negative f_i^n). Triplet loss $\mathcal{L}_{triplet}$ with

cosine distance is formulated as:

$$p_i = 1 - \cos(f_i^a, f_i^p) \quad n_i = 1 - \cos(f_i^a, f_i^n) \quad (6)$$

$$\mathcal{L}_{triplet} = \frac{1}{N} \sum_i [(p_i - n_i) + \delta]_+ \quad (7)$$

where p_i and n_i represent the i -th positive and negative items of triplet loss respectively. The margin value δ is set to 0.1.

Accurate classification. We choose CosFace [12] as our classification loss, formulated as

$$\mathcal{L}_{cls} = \frac{1}{N} \sum_{i=1}^N -\log \frac{e^{s(\cos \theta_{y_i, i} - \beta)}}{e^{s(\cos \theta_{y_i, i} - \beta)} + \sum_{j=1, j \neq y_i}^n e^{s \cos \theta_{j, i}}} \quad (8)$$

where y_i is the i th ground-truth label, and $\cos \theta_{j, i}$ is the cosine distance of the j th class weight and the i th feature. s is set to 30.0 and β is set to 0.4. Eq. 8 maximizes inter-class variance and minimizes intra-class variance by a cosine margin. We notice that ArcFace [13] demonstrates better mask-free face classification performance, but empirically CosFace is better in our framework. More experimental details are in Sec. 4.3.

3 Implementation

Network Structure. We use SOTA face recognition model SENet50 as face encoder [13]. We employ {Conv-PReLU-BN} residual convolutional structure for *SpcRec* and {Linear-PReLU-Linear} fully connected structure for *ChnRec*. Both blocks have 9 layers of the corresponding structure with Sigmoid at last. For feature fusion, we simply concatenate three features and feed into a {Conv-PReLU-BN} block.

Model Enhancement. Considering our model might sacrifice the accuracy of clean face recognition slightly, we propose a model enhancement method (*FFR-Net+*). For each pair of faces, we combine the similarity score \hat{s} from FFR-Net and the score s from raw SENet50. Note that SENet is a part of our model so no extra computational cost will be added. The new score s_n is formulated as $s_n = \max(s, \hat{s}) + \hat{s}$.

Training Details. The encoder SENet50 pre-trained on refined MS1M [14] is fixed and only FFR-Net is trainable. We train models on CASIA-WebFace [15] and the simulated masked data. We use Adam [16] with $lr = 0.01$ as optimizer.

Real-world Data. The collected real-world masked face dataset is refined from RMFD proposed by [17], consisting of 100 identities, 512 images, and 782 pairs of a masked face and a clean face. Evaluation on real data is shown in Sec. 4.2.

Mask Simulation. We choose MaskTheFace [18] repository¹ as our mask generator. Simulated samples are shown in Fig. 2.

4 Experimental Results

4.1 Evaluation Metric

Average Accuracy. Performance imbalance exists between mask-free and masked faces in many face recognition models. Therefore, datasets of all mask-free or all masked faces



Fig. 2: Example results of the mask generator.

are not sufficient for fair comparison. We introduce a series of testing datasets $\{D_t \mid t = 0.0, 0.1, \dots, 1.0\}$ in which t is randomly masked faces ratio. Accuracy is obtained on all datasets respectively. The mean accuracy is used as the key evaluation metric, denoted as Avg. in following experiments. Besides, we also consider another 2 criteria for comparison: 1) Mask-free: accuracy only on mask-free data and 2) Masked: accuracy only on masked data.

Baseline Models. We consider four baseline models. 1) *SENet50*: the face recognition model, i.e., our encoder. 2) *f-SENet50*: the same network structure as the first one but finetuned by the simulated CASIA-WebFace dataset. 3) *SENet50-convs*: the backbone model SENet50 with additional convolutional layers. 4) *PDSN*: pairwise differential siamese network (PDSN) [6], the SOTA method based on feature discarding mask to eliminate harmful interference of corrupted features.

4.2 Performance

Face Verification. We conduct evaluation on LFW [4], a popular pair-wise face verification benchmark dataset, following the standard protocol. We report the verification accuracy in Fig. 3 and Table 1. We can observe *f-SENet50* improves on masked faces by 1.07% but decreases by 0.68% on mask-free, because the finetuned model corrupts the original feature space to some extent. *SENet50-convs* shows no improvement, indicating only increasing convolutional layers is not sufficient. Our model increases the average accuracy by 0.41% compared to *SENet50*, outperforming the others by +0.16%. *FFR-Net+* further improves 0.14%. On real-world masked data, our method improves *SENet50* by 3.08%. *FFR-Net* slightly outperforms *FFR-Net+* because the enhancement method is to improve the performance on mask-free data instead of masked conditions. We also report the performance of our method compared with some latest works on masked face recognition in Table 2. Our method significantly surpasses CQA-Face [19] and Jeevan *et al.* [20] by 1.50%-5.45%.

Table 1: Face verification accuracy on LFW and real data

Method	#Image	Mask-free	Masked	Avg.	Real data
CosFace	5M	99.73%	—	—	—
ArcFace	0.49M	99.53%	—	—	—
<i>SENet50</i>	5M	99.60%	97.00%	98.23%	90.00%
<i>f-SENet50</i>	0.98M	98.92%	98.07%	98.48%	89.36%
<i>SENet50-convs</i>	0.98M	98.95%	97.07%	97.96%	86.03%
<i>PDSN</i>	0.98M	99.33%	97.07%	98.25%	90.77%
<i>FFR-Net</i>	0.98M	99.27%	98.22%	98.64%	93.08%
<i>FFR-Net+</i>	0.98M	99.43%	98.23%	98.78%	92.95%

¹<https://github.com/aqeelanwar/MaskTheFace>




Fig. 3: Acc@Ratio on LFW. Fig. 4: Acc@Ratio on MF1.

Table 2: Comparison with some latest works on masked LFW

Method	CQA-Face [19]	Jeevan <i>et al.</i> [20]	FFR-Net	FFR-Net+
Acc	92.78%	96.73%	98.22%	98.23%

Face Identification. We also test the model on MegaFace Challenge1 [5] (MF1) to evaluate the performance of face identification task. MF1 is composed of a gallery set with more than 1 million face images and a probe set of Face-Scrub [21]. The training set is viewed as small if it is less than 0.5M or viewed as large. Rank-1 accuracy with 1M distractors is used in experiments. From Fig. 4 and Table 3, the accuracy of *SENet50* decreases by 49.64% due to masks. Compared to *SENet50*, our method greatly improves on masked data by 25.06% while keeping considerable performance on mask-free one with only 6.48% accuracy dropping. Our model outperforms other methods by +1.54%. *FFR-Net+* further improves 2.11%. In addition, our model only increases FLOPS by 17.9% and the number of trainable parameters is 50.6% of the fixed backbone, which indicates low calculation cost of the proposed model.

Table 3: Face identification accuracy on MF1

Method	protocol	FLOPS	#params	Mask-free	Masked	Avg.
CosFace	small	6.31G	43.79M	77.11%	—	—
ArcFace	small	6.31G	43.79M	77.50%	—	—
CosFace	large	6.31G	43.79M	82.72%	—	—
<i>SENet50</i>	large	6.31G	43.79M	89.40%	39.76%	64.57%
<i>f-SENet50</i>	large	6.31G	43.79M	76.13%	68.80%	72.47%
<i>SENet50-convs</i>	large	7.00G	57.95M	76.01%	53.14%	64.62%
<i>PDSN</i>	large	6.45G	59.05M	81.35%	42.98%	62.21%
<i>FFR-Net</i>	large	7.44G	65.94M	82.92%	64.82%	<u>74.01%</u>
<i>FFR-Net+</i>	large	7.44G	65.94M	<u>86.91%</u>	65.06%	76.12%

4.3 Ablation Study

Component Ablation. We first ablate self-similarity features (*SS*), flipped fusion block (*FF*), as tag A and B in Table 4. The absence of *SS* and *FF* respectively decreases the average accuracy. Besides, we verify the significance of three features in fusion (\hat{f}_s , \hat{f}_c and f). C1 and C2 respectively exclude *ChnRec* and *SpcRec*, which result in 1.95% and 2.30% accuracy declining. This proves the efficiency of channel-and-space rectification block. C3 shows 3.26% accuracy dropping without fusing the original feature. Therefore, fusing f along with \hat{f}_s and \hat{f}_c can greatly boost the performance.

Loss Ablation. In Table 5, we conduct ablation study on identical facial representation constraint (\mathcal{L}_{ID1}) and identical self-similarity pattern constraint (\mathcal{L}_{ID2}) as tag D. Accuracy declines by +1.18% when removing \mathcal{L}_{ID1} or \mathcal{L}_{ID2} . Besides, we test our model with ArcFace [13] as tag E. ArcFace fails to

Table 4: Ablation study of components on MF1

Tag	<i>SS</i>	<i>FF</i>	\hat{f}_s	\hat{f}_c	f	Mask-free	Masked	Avg.
A		✓	✓	✓	✓	82.00%	63.29%	72.61%
B	✓		✓	✓	✓	78.57%	56.75%	67.83%
C1	✓	✓	✓		✓	81.69%	62.42%	72.06%
C2	✓	✓		✓	✓	82.39%	61.03%	71.71%
C3	✓	✓	✓	✓		79.06%	62.44%	70.75%
Full	✓	✓	✓	✓	✓	82.92%	64.82%	74.01%




Fig. 5: T-SNE visualization: colors represent different identities, \circ denote mask-free face and \triangle denotes masked face.

train an efficient model with a 39.71% drop, proving CosFace is more suitable for training of learnable module RecBlocks.

Table 5: Ablation study of losses on MF1

Tag	\mathcal{L}_{ID1}	\mathcal{L}_{ID2}	<i>AF</i>	<i>CF</i>	Mask-free	Masked	Avg.
D1				✓	81.77%	63.10%	72.54%
D2	✓			✓	82.03%	62.95%	72.55%
D3		✓		✓	82.76%	62.88%	72.83%
E	✓	✓	✓		41.45%	27.25%	34.30%
Full	✓	✓		✓	82.92%	64.82%	74.01%

4.4 Feature Space Visualization

The main idea of feature rectification is to find a proper mapping to project original features to rectified features. We visualize the original and rectified feature spaces to display the effective mapping process. The t-SNE result is shown in Fig. 5. In the original feature space, masked samples (triangles) and mask-free samples (circles) tend to gather separately and keep far away from each other. However, in the rectified feature space, triangles tend to align with circles and they are equally distributed. Compared to the original space, the rectified space possesses more compact intra-class geodesic distances and larger inter-class margins with greater similarity between the masked feature and the mask-free counterpart. This result matches our initial idea that the rectified feature space is more identity-separable and mask-resistant than the original one.

5 CONCLUSION

In this paper, we introduce a novel framework via feature rectification to uniformly recognize masked and mask-free faces. In this context, we propose channel-and-space rectification block to rectify facial representations at feature level in both spatial and channel dimensions and design powerful losses to enforce ideal properties of the unified rectified feature space. As a result, our model achieves state-of-the-art performance on mixed datasets of masked and mask-free faces. Finally, we verify contributions of the proposed blocks and losses to performance enhancement and rectification mapping.

6 References

- [1] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf, “Deepface: Closing the gap to human-level performance in face verification,” in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2014.
- [2] Yi Sun, Xiaogang Wang, and Xiaoou Tang, “Deep learning face representation from predicting 10,000 classes,” in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2014.
- [3] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang, “Deep learning face representation by joint identification-verification,” in *Advances in neural information processing systems (NeurIPS)*, 2014.
- [4] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller, “Labeled faces in the wild: A database for studying face recognition in unconstrained environments,” Tech. Rep., University of Massachusetts, Amherst, 2007.
- [5] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard, “The megaface benchmark: 1 million faces for recognition at scale,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [6] Lingxue Song, Dihong Gong, Zhifeng Li, Changsong Liu, and Wei Liu, “Occlusion robust face recognition based on mask learning with pairwise differential siamese network,” in *Proceedings of International Conference on Computer Vision (ICCV)*, 2019.
- [7] Weitao Wan and Jiansheng Chen, “Occlusion robust face recognition based on mask learning,” in *2017 IEEE International Conference on Image Processing (ICIP)*, 2017.
- [8] Lele Cheng, Jinjun Wang, Yihong Gong, and Qiqi Hou, “Robust deep auto-encoder for occluded face recognition,” in *Proceedings of the 23rd ACM international conference on Multimedia*, 2015.
- [9] Fang Zhao, Jiashi Feng, Jian Zhao, Wenhan Yang, and Shuicheng Yan, “Robust lstm-autoencoders for face de-occlusion in the wild,” *IEEE Transactions on Image Processing (TIP)*, 2018.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems (NeurIPS)*, 2014.
- [11] Jie Hu, Li Shen, and Gang Sun, “Squeeze-and-excitation networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018.
- [12] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu, “Cosface: Large margin cosine loss for deep face recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [13] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou, “Arcface: Additive angular margin loss for deep face recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [14] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao, “Ms-celeb-1m: A dataset and benchmark for large-scale face recognition,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- [15] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li, “Learning face representation from scratch,” *arXiv preprint arXiv:1411.7923*, 2014.
- [16] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [17] Zhongyuan Wang, Guangcheng Wang, Baojin Huang, Zhangyang Xiong, Qi Hong, Hao Wu, Peng Yi, Kui Jiang, Nanxi Wang, Yingjiao Pei, et al., “Masked face recognition dataset and application,” *arXiv preprint arXiv:2003.09093*, 2020.
- [18] Aqeel Anwar and Arijit Raychowdhury, “Masked face recognition for secure authentication,” *arXiv preprint arXiv:2008.11104*, 2020.
- [19] Qiangchang Wang and Guodong Guo, “Cqa-face: Contrastive quality-aware attentions for face recognition,” in *Association for the Advancement of Artificial Intelligence (AAAI)*, 2022.
- [20] Govind Jeevan, Geevar C. Zacharias, Madhu S. Nair, and Jeny Rajan, “An empirical study of the impact of masks on face recognition,” *Pattern Recognition*, 2022.
- [21] Hong-Wei Ng and Stefan Winkler, “A data-driven approach to cleaning large face datasets,” in *IEEE International Conference on Image Processing (ICIP)*, 2014.