

The Effect of Wearing a Mask on Face Recognition Performance: an Exploratory Study

Naser Damer¹², Jonas Henry Grebe¹, Cong Chen¹, Fadi Boutros¹²,
Florian Kirchbuchner¹, Arjan Kuijper¹²

Abstract: Face recognition has become essential in our daily lives as a convenient and contactless method of accurate identity verification. Process such as identity verification at automatic border control gates or the secure login to electronic devices are increasingly dependant on such technologies. The recent COVID-19 pandemic have increased the value of hygienic and contactless identity verification. However, the pandemic led to the wide use of face masks, essential to keep the pandemic under control. The effect of wearing a mask on face recognition in a collaborative environment is currently sensitive yet understudied issue. We address that by presenting a specifically collected database containing three session, each with three different capture instructions, to simulate realistic use cases. We further study the effect of masked face probes on the behaviour of three top-performing face recognition systems, two academic solutions and one commercial off-the-shelf (COTS) system.

Keywords: Face recognition, COVID-19, masked face recognition.

1 Introduction

Given the current COVID-19 pandemic, it is essential to enable contactless and smooth running operations, especially in contact sensitive facilities like airports. Face recognition have been been praised as such an accurate and contactless mean of verifying identities. Wearing masks is essential to prevent the spread of contagious diseases and have been currently forced in public places in many countries. However, the performance, and thus the trust, of contactless identity verification through face recognition can be effected by wearing a mask.

Face occlusion have been repeatedly addressed in the scope of face detection solutions [Op16]. Moreover, developing occlusion invariant face recognition solutions has been a growing research challenge [So19]. However, most of these works address general occlusion that commonly appear in in-the-wild capture conditions, such as sunglasses and partial captures. Given the current COVID-19 pandemic, it is essential to study the specific effect of wearing face masks on the behaviour of face recognition system in a collaborative environment. Our work aims at studying this effect to enable the future development of solutions addressing accurate face recognition in such scenarios. To achieve that, we present a database that simulates a realistically variant collaborative face capture scenario. This database is a first version of an on-going data collection process that includes three session, each with three capture variations, per subject. We study the behaviour of three of

¹ Fraunhofer Institute for Computer Graphics Research IGD, Darmstadt, Germany

² Mathematical and Applied Visual Computing, TU Darmstadt, Darmstadt, Germany

the top performing face recognition solutions (one commercial and two academic) when encountering masked faces, in comparison to the typical no-mask baseline. We conclude with pointing out strong signs of negative effect on face recognition systems, showing the need to develop appropriate evaluation databases and recognition solutions.

2 Related work

Face recognition deployment faces a number of operational challenges. Many of these challenges, and thus the research efforts, are related to attacks on face recognition systems, such as presentation attacks (spoofing) [DD16], morphing attacks [Da19], or other unconventional attacks [Da18]. However, issues related to the biometric sample presentation, such as face occlusion, can also effect face recognition deployability. The detection of occluded faces is a well-studied issue in the computer vision domain. An example of that is the work of Optiz et al. [Op16] that proposed a novel grid loss targeting a more accurate detection of occluded faces. Focusing on masked faces, Ge et al. [Ge17] presented a solution to enhance the detection (not biometric recognition) of masked faces in in-the-wild scenarios. Their experiments did not focus on masks worn specifically for health protection reasons, but included other forms of face occlusions. However, their solution is relevant to face recognition as our experiments will show later that the investigated face recognition solutions fails in some cases to detect a face.

As stated, detecting occluded faces is a challenge that affect the operation of face biometric systems. However, the biometric recognition of these faces is a more dominant challenge. An example of the works addressing this challenge is that of Song et al. [So19] where they aim at enhancing face recognition for faces with general occlusions. Their approach tries to learn finding and discarding corrupted feature elements, linked to occlusions, from the recognition process. Focusing on masks, in a very recent work, Wang et al. [Wa20] presented, in a brief and undetailed work, crawled databases for face detection, recognition and simulated masked faces. The authors claim to enhance the recognition accuracy from 50% to 95% without providing information on their baseline, proposed algorithmic details, or clearly specifying the evaluation database. Given the current COVID-19 pandemic, a specifically collected database and evaluation of wearing real face mask on collaborative face recognition is necessary and is still missing.

3 The database

The goal of the collected database is to enable the study of face recognition performance on masked faces and drive future innovation in this domain. The database presented in this work is an initial version and further data collection efforts is on going. The data tries to simulate a collaborative, yet varying, scenario. Such as the situation in automatic border control gates or unlocking personal devices with face recognition, where the mask, illumination, and background can change.

Each of the participants was asked to collect the data on three different, not necessary consecutive days. We consider each of these days as one session. On each day, the participant will collect three videos, each of a minimum length of 5 seconds. All videos are collected from static (not hand held) webcams and the users were asked to simulate a login scenario by looking at the capture device. The images were all captured indoors, each at their residence during home-office. The capture was performed during the day (day-light) and the participants were asked to remove eyeglasses only when the frame is considered very thick. No other restrictions were imposed, such as background or mask type and its consistency over days, to simulate realistic scenarios. The three videos captured each day were as follows: 1) Face with no mask and no additional electric illumination, this will be noted as baseline (BL). 2) Face with mask on and no additional electric illumination, this will be noted as mask one (M1). 3) Face with mask on and the existing electric light in the room is turned on, this will be noted as mask two (M2). The M2 is considered to study the unknown effect of illumination variation in the case of masked face recognition, given that the mask might result in different reflection and shadow patterns.

The first session (day) is considered as the reference data (R), resulting in the baseline reference (BLR), the mask one reference (M1R), and mask two reference (M2R). The second and third sessions (days) were considered as probe data (P) and they result in the baseline probe (BLP), the mask one probe (M1P), and mask two probe (M2P), and the joint probe data from M1P and M2P referred to as M12P. From each captured video, the first second was neglected to avoid any biases related to the user interaction with the capture device. From the following three seconds, 10 frames were selected with 9 frames gap between them, as all videos are recorded at 30 frames per second. The total number of participant at this first version of the database is 24, and they all participated in all sessions. Given the number of sessions, participants and the considered frames from each video, Table 1 provide an overview on the database structure. Samples of the database are shown in Figure 1.

Session	Session 1: References			Session 2 and 3: Probes			
Data split	BLR	M1R	M2R	BLP	M1P	M2P	M12P
Illumination	No	No	Yes	No	No	Yes	Both
Number of Captures	240	240	240	480	480	480	960

Tab. 1: An overview of the database structure.

4 Face recognition

To provide a wide view on the effect of wearing a mask on face recognition performance, we analyse the performance of three face recognition algorithms. Two of these algorithms are of the top performing academic approaches, namely the ArcFace [De19] and SphereFace [Li17]. The third algorithm is a COTS algorithm from the vendor Neurotechnology [Ne]. In the following, this section provides more details on these algorithms.



Fig. 1: Samples of the collected database from the three capture types (BL, M1, and M2)

SphereFace: We chose SphereFace as it achieved competitive verification accuracy on Labeled Face in the Wild (LFW) [Hu07] 99.42% and Youtube Faces (YTF) [WHM11] 95.0% using 64-CNN layers trained on CASIA-WebFace dataset [Yi14]. SphereFace is trained using angular Softmax loss function (A-Softmax). The key idea behind A-Softmax loss is to learn discriminative features from the face image by formulating the Softmax as angular computation between the embedded features vector X and their weights W .

ArcFace: ArcFace achieved state-of-the-art performance of several face recognition benchmarks such as LFW 99.83% and YTF 98.02%. ArcFace introduced Additive Angular Margin loss (ArcFace) to enhance the discriminative power of the face recognition model. We employed ArcFace based on ReseNet-100 [He16] architecture pretrained on refined version of MS-Celeb-1M dataset [Gu16] (MS1MV2).

COTS: We used the MegaMatcher 11.2 SDK [Ne] from the vendor Neurotechnology. We chose this COTS product as Neurotechnology achieved one of the best performances in the recent NIST report addressing the performance of vendor face verification products [GP20]. The face quality threshold was set to zero to minimize neglecting masked faces. The full processes of detecting, aligning, feature extraction, and matching are part of the COTS and thus we are not able to provide their algorithmic details. Matching two faces by the COTS produces a similarity score.

For the ArcFace [De19] and SphereFace [Li17], the Multi-task Cascaded Convolutional Networks (MTCNN) [Zh16] solution is used, as recommended in [Li17], to detect (crop) and align (affine transformation) the face. Both network process the input aligned and cropped image and produce a feature vector of the size 512. To compare two faces, a distance is calculated between their respective feature vectors. This is calculated as Euclidean distance for ArcFace features, as recommended in [De19], and as Cosine distance for SphereFace features, as recommended in [Li17]. The Euclidean distance (dissimilarity) is complemented to show a similarity score and the Cosine distance shows similarity score by default.

5 Experimental setup

To baseline the performance, we evaluate the face verification performance without masks. This is done by N:N comparison of the data splits BLR and BLP (BLR-BLP). To measure the performance when wearing a mask, we perform an N:N comparison between the data splits BLR and M1P (BLR-M1P). To evaluate any induced performance change by having an additional illumination (room light) when wearing a mask, we perform an N:N comparison between the data splits BLR and M2P (BLR-M2P). To measure the overall performance including both considered illumination, we perform an N:N comparison between the data splits BLR and M12P (BLR-M12P). These four experiments are used to evaluate each of the three considered face recognition solutions.

To study the effect of wearing a mask on the recognition performance, we plot the genuine and imposter distributions of the BLR-BLP (baseline) comparisons along with the genuine and imposter score distributions of the BLR-(M1P or M2P or M12P) (mask). This allows analysing the shifts in the distributions induced by wearing a mask. We also report the mean of the genuine scores (G-mean) and mean of imposter scores (I-mean) for each experiment, to get a quantitative measure of the comparison scores shifts.

Based on the standard ISO/IEC 19795-1 [Ma06], we also enrich our performance study by a number of verification performance metrics. As the face mask induces a strong appearance change on the face, face detection might be challenging. Therefore, we report the failure to extract rate (FTX) for each experiment. FTX is proportion of failures of the feature extraction process to generate a template from the captures sample. Besides reporting the FTX, and only for the samples where a template can be created, we report algorithmic verification performance metrics. These metrics include the general Equal Error Rate (EER), which is defined as the common value of false mathc rate (FMR) and false non match rate (FNMR) at the decision threshold where they are identical. We also show the algorithmic verification performance by listing the FNMR at different operation points by presenting the achieved FMR100, FMR1000, and ZeroFMR, which are the lowest FNMR for an FMR $\leq 1.0\%$, $\leq 0.1\%$, and $\leq 0\%$, respectively. To provide an algorithmic verification performance illustration on the complete range of operation points, we plots the receiver operating characteristic (ROC) curves for all the experimental setups, for each of the investigated face recognition systems.

6 Evaluation results

Figure 2 presents the comparison between the baseline (BLR-BLP) genuine and imposter score distributions and the different masked faces experiments (BLR-M1P, BLP-M2P, BLR-M12P) on the three considered face recognition solutions. It is noticeable in all experimental setups that, when comparing masked faces probes to unmasked references, the genuine score distributions strongly shift towards the imposter distributions in comparison to the BLR-BLP setup. This indicates an expected decrease in performance and general trust in the matcher decision, as the separability between genuine and imposter samples decreases. This unwanted shift seems to be slightly stronger when the masked faces are

captured under additional artificial illumination (BLR-M2P) when compared to the natural light condition (BLR-M2P). On the other hand, the imposter score distributions do not seem to be significantly affected by the masked probes (BLR-M1P, BLP-M2P, BLR-M12P) in comparison to the unmasked baseline (BLR-BLP).

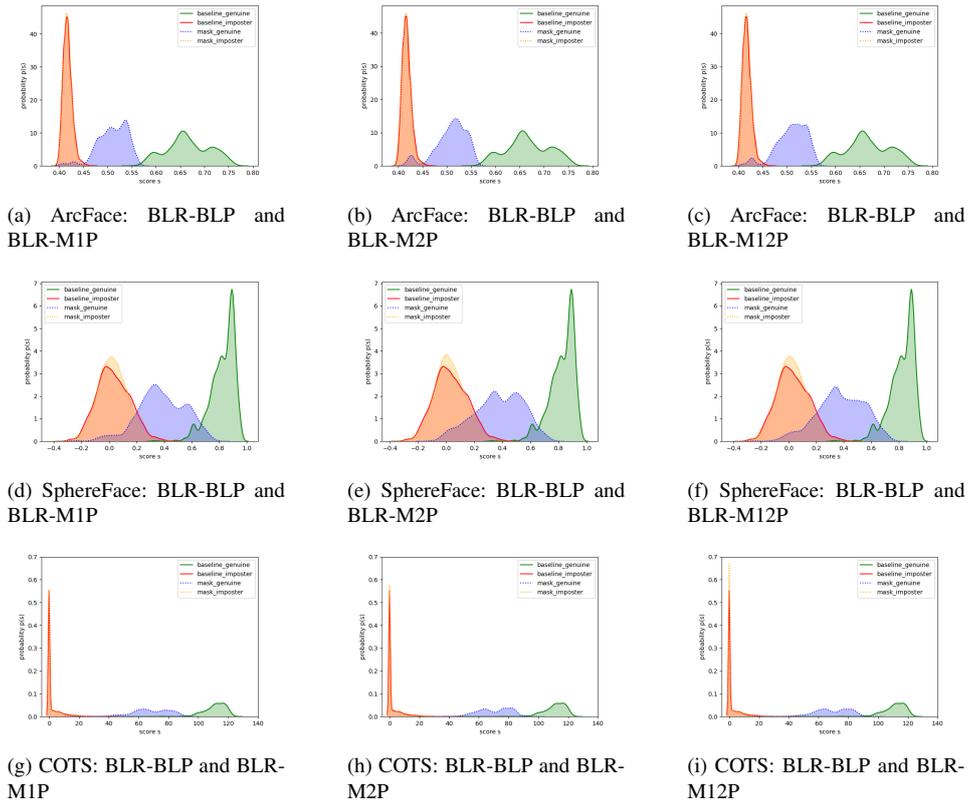


Fig. 2: The comparison score (similarity) distributions comparing the "baseline" BLR-BLP genuine and imposter distributions to those of the distributions including "masked" faces probes (BLR-M1P (a, d, g)), BLR-M2P (b, e, h), BLR-M12P (c, f, i). The shift of the genuine scores towards the imposter distribution is clear when faces are masked for all investigated system (ArcFace(a, b, c), SphereFace (d, e, f), and COTS (g, h, i)).

Tables 2, 3, and 4 present the achieved performance, given by the different evaluation metrics, on all experimental setups by the ArcFace, SphereFace, and COTS solutions, respectively. In all systems, wearing a face mask affected the ability to detect the face properly, resulting in a higher than zero (as in the baseline) FTX. Interestingly, additional illumination (typically from the top) increased the FTX in all systems (BLR-M2P compared to BLR-M1P). This is probably due to the different reflection and shadow patterns induced by the illumination, see samples in Figure 1. The FTX values for the SphereFace and ArcFace in tables 2 and 3 are identical as they both use the MTCNN network for face detection and alignment.

The verification performance (EER, FMR100, FMR1000, ZeroFMR) of the ArcFace and SphereFace is negatively affected when the probe faces are masked (BLR-M1P and BLR-M2P), see tables 2 and 3. This negative effect is stronger when the faces are captured under the effect of artificial illumination (BLR-M2P), probably due to unexpected reflections and shadowing and the fact that the BLR references were captured without such illumination. The reduction in the performance is much more dominant in the SphereFace solution in comparison to the ArcFace. For both systems, the G-mean values decreased significantly when considering the masked probes. This, despite the small size of the evaluation data, indicates a strong negative effect of the masks on the face recognition performance. On the other hand, the I-mean value when considering the masked faces, in comparison to the baseline (BLR-BLP), was not changed under the the ArcFace solution and only slightly changed under the SphereFace solution.

ArcFace	EER	FMR100	FMR1000	ZeroFMR	G-mean	I-mean	FTX
BLR-BLP	0.000%	0.000%	0.000%	0.000%	0.666	0.417	0.000%
BLR-M1P	3.163%	3.517%	3.831%	5.069%	0.511	0.417	3.750%
BLR-M2P	5.504%	6.163%	6.628%	7.616%	0.509	0.417	5.833%
BLR-M12P	4.380%	4.888%	5.229%	6.468%	0.510	0.417	4.792%

Tab. 2: The verification performance measures, the G-mean, and I-mean achieved by ArcFace on the different experimental setups. Note the performance degradation induced by the masked face probes.

SphereFace	EER	FMR100	FMR1000	ZeroFMR	G-mean	I-mean	FTX
BLR-BLP	0.216%	0.065%	0.217%	0.390%	0.825	0.033	0.000%
BLR-M1P	9.312%	27.35%	52.95%	72.91%	0.384	0.026	3.750%
BLR-M2P	12.36%	28.22%	47.66%	73.16%	0.374	0.025	5.833%
BLR-M12P	10.85%	27.86%	50.01%	73.38%	0.380	0.025	4.792%

Tab. 3: The verification performance measures, the G-mean, and I-mean achieved by SphereFace on the different experimental setups. Note the performance degradation induced by the masked face probes.

COTS	EER	FMR100	FMR1000	ZeroFMR	G-mean	I-mean	FTX
BLR-BLP	0.249%	0.000%	0.251%	0.668%	110.8	2.281	0.000%
BLR-M1P	0.443%	0.304%	0.684%	2.253%	68.09	2.221	2.500%
BLR-M2P	0.004%	0.000%	0.009%	0.050%	70.88	2.298	3.542%
BLR-M12P	0.239%	0.152%	0.341%	1.237%	69.49	2.259	3.021%

Tab. 4: The verification performance measures, the G-mean, and I-mean achieved by COTS on the different experimental setups. Although the change in the performance caused by the masked face probes is insignificant, the shift in the average genuine score towards the imposter scores is very dominant in these cases.

When it comes to verification performance metrics (EER, FMR100, FMR1000, ZeroFMR), the COTS is not significantly affected by masked faces. This is apparent in Table 4, where these performance metrics are not significantly different in all experimental setups. This might be due to the robust and high performance of the COTS solution and the limited size of the evaluation database. However, the change in the G-mean from 110.8 in the BLR-BLP to 69.46 in the BLR-M12P, while maintaining a similar I-mean, indicates a large

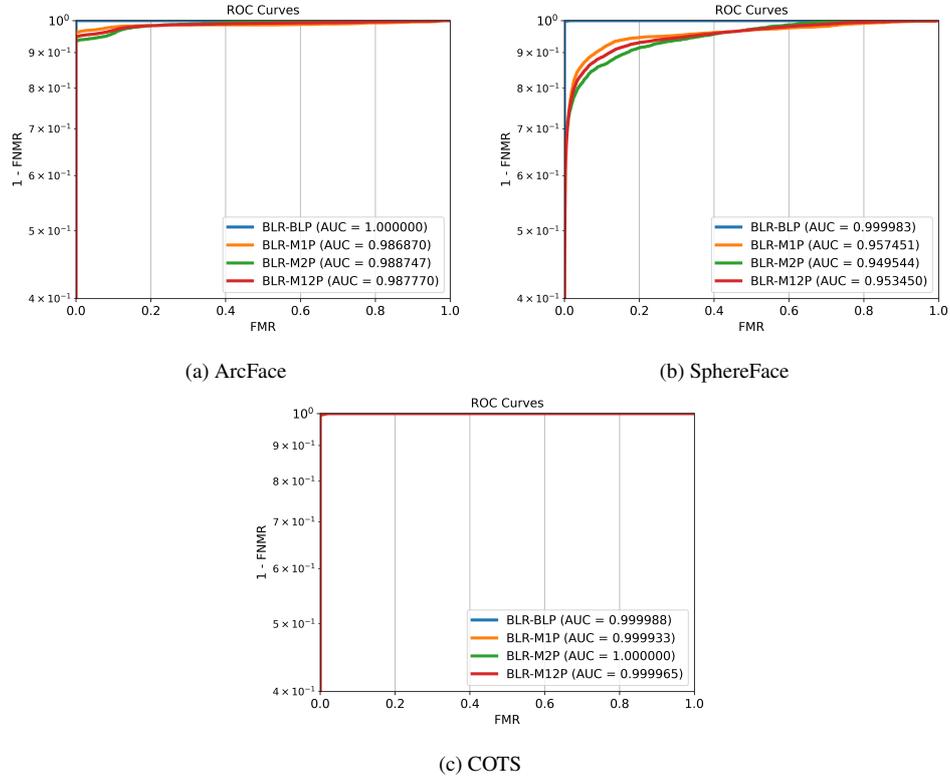


Fig. 3: The verification performance for the three investigated system (ArcFace(a), SphereFace (b), and COTS (c)) is presented as ROC curves. For each of the systems, four curves are plotted to represent the three settings that include "masked" faces probes (BLR-M1P, BLR-M2P, and BLR-M12P) and the unmasked baseline (BLR-BLP). The area under curve (AUC) is also listed for each of the ROC curves. As in Tables 2, 3, and 4, the effect of masked probes is apparent on the performance of the ArcFace and SphereFace, while the performance of the COTS is almost perfect in all experimental settings (however, with shift in genuine scores values).

change in the separability (between genuine and imposter) in the COTS decisions. This can lead to an increase in the error rate given a larger and more challenging evaluation. Such an evaluation is planned as the data presented in this paper is an initial version of a larger data being collected at the moment. To show the verification performance over a wider range of operation points, Figure 3 presents the ROC curves for the different experimental settings for each of the three investigated systems. Similar conclusions to those established from Tables 2, 3, and 4 can be made. The ArcFace and SphereFace verification performance is effected by the masked probe faces, while the COTS maintains an almost perfect verification performance. However, one must keep in mind the significant shift in the genuine score values in all three systems, as illustrated in Figure 2.

In general, the effect of wearing face masks on the face recognition behaviour is apparent on all investigated systems. The effect is most significant on the genuine scores distribution, rather than the imposter scores distribution. This renders the current face recognition solutions undependable to match masked faces with unmasked faces and, at least, requires re-evaluation.

7 Conclusion

Addressing the wide spread use of face masks as a preventive measure to the COVID-19 pandemic spread, we presented an exploratory study on the effect of wearing masks on face recognition performance in collaborative scenarios. We presented a specifically collected database captured in three different sessions, with and without wearing a mask, and is part of an ongoing effort to gather a larger scale database with realistic variations. We analysed the behaviour of two high-performing academic face recognition solutions and one of the top performing COTS solutions. Our analyses pointed out the significant effect of wearing a mask on comparison scores separability between genuine and imposter comparisons in all the investigated systems. Moreover, we point out a large drop in the verification performance of the academic face recognition solutions, even on a limited evaluation data, when considering masked face probes.

Acknowledgment

This research work has been funded by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE

References

- [Da18] Damer, Naser; Wainakh, Yaza; Boller, Viola; von den Berken, Sven; Terhöst, Philipp; Braun, Andreas; Kuijper, Arjan: CrazyFaces: Unassisted Circumvention of Watchlist Face Identification. In: 9th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2018, Redondo Beach, CA, USA, October 22-25, 2018. IEEE, pp. 1–9, 2018.
- [Da19] Damer, Naser; Saladie, Alexandra Mosegui; Zienert, Steffen; Wainakh, Yaza; Terhöst, Philipp; Kirchbuchner, Florian; Kuijper, Arjan: To Detect or not to Detect: The Right Faces to Morph. In: 2019 International Conference on Biometrics, ICB 2019, Crete, Greece, June 4-7, 2019. IEEE, pp. 1–8, 2019.
- [DD16] Damer, Naser; Dimitrov, Kristiyan: Practical View on Face Presentation Attack Detection. In (Wilson, Richard C.; Hancock, Edwin R.; Smith, William A. P., eds): Proceedings of the British Machine Vision Conference 2016, BMVC 2016, York, UK, September 19-22, 2016. BMVA Press, 2016.
- [De19] Deng, Jiankang; Guo, Jia; Xue, Niannan; Zafeiriou, Stefanos: ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 4690–4699, 2019.

- [Ge17] Ge, Shiming; Li, Jia; Ye, Qiting; Luo, Zhao: Detecting Masked Faces in the Wild with LLE-CNNs. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. IEEE Computer Society, pp. 426–434, 2017.
- [GP20] Grother Patrick, Ngan Mei, Hanaoka Kayee: Ongoing Face Recognition Vendor Test (FRVT). NIST Interagency Report, 2020.
- [Gu16] Guo, Yandong; Zhang, Lei; Hu, Yuxiao; He, Xiaodong; Gao, Jianfeng: MS-Celeb-1M: A Dataset and Benchmark for Large-Scale Face Recognition. In: Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III. pp. 87–102, 2016.
- [He16] He, Kaiming; Zhang, Xiangyu; Ren, Shaoqing; Sun, Jian: Deep Residual Learning for Image Recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. IEEE Computer Society, pp. 770–778, 2016.
- [Hu07] Huang, Gary B.; Ramesh, Manu; Berg, Tamara; Learned-Miller, Erik: Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- [Li17] Liu, Weiyang; Wen, Yandong; Yu, Zhiding; Li, Ming; Raj, Bhiksha; Song, Le: SphereFace: Deep Hypersphere Embedding for Face Recognition. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 6738–6746, 2017.
- [Ma06] Mansfield, A: Information technology–Biometric performance testing and reporting–Part 1: Principles and framework. ISO/IEC, pp. 19795–1, 2006.
- [Ne] Neurotechnology MegaMatcher 11.2 SDK. "https://www.neurotechnology.com/mm_sdk.html".
- [Op16] Opitz, Michael; Waltner, Georg; Poier, Georg; Possegger, Horst; Bischof, Horst: Grid Loss: Detecting Occluded Faces. In: Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III. volume 9907 of Lecture Notes in Computer Science. Springer, pp. 386–402, 2016.
- [So19] Song, Lingxue; Gong, Dihong; Li, Zhifeng; Liu, Changsong; Liu, Wei: Occlusion Robust Face Recognition Based on Mask Learning With Pairwise Differential Siamese Network. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 773–782, 2019.
- [Wa20] Wang, Zhongyuan; Wang, Guangcheng; Huang, Baojin; Xiong, Zhangyang; Hong, Qi; Wu, Hao; Yi, Peng; Jiang, Kui; Wang, Nanxi; Pei, Yingjiao; Chen, Heling; Miao, Yu; Huang, Zhibing; Liang, Jinbi: , Masked Face Recognition Dataset and Application, 2020.
- [WHM11] Wolf, Lior; Hassner, Tal; Maoz, Itay: Face recognition in unconstrained videos with matched background similarity. In: CVPR 2011. IEEE, pp. 529–534, 2011.
- [Yi14] Yi, Dong; Lei, Zhen; Liao, Shengcai; Li, Stan Z.: Learning Face Representation from Scratch. CoRR, abs/1411.7923, 2014.
- [Zh16] Zhang, Kaipeng; Zhang, Zhanpeng; Li, Zhifeng; Qiao, Yu: Joint face detection and alignment using multitask cascaded convolutional networks. IEEE Signal Processing Letters, 23(10):1499–1503, 2016.