

Vulnerability of Open-source Face Recognition Systems to Blackbox Attacks: A Case Study With InsightFace

Nafiz Sadman, Kazi Amit Hasan, Elyas Rashno, Furkan Alaca, Yuan Tian, and Farhana Zulkernine

School of Computing, Queen's University

Kingston, ON, Canada

{sadman.n, kaziamit.hasan, elyas.rashno, furkan.alaca, y.tian, farhana.zulkernine}@queensu.ca

Abstract—This paper presents a comprehensive analysis of the security aspects of the InsightFace project (a popular open-source face recognition system) focusing on its susceptibility to three distinct black box attacks: *Face Swap*, *Morphing*, and *Presentation*. Open-source face recognition models are used in commercial applications, thereby motivating our security analysis. Our investigation entails a meticulous evaluation of the susceptibility of the project to false authentication when subjected to the three attacks. We observed from our experiments that InsightFace was not able to differentiate between legitimate images and manipulated images. The principal aim of this research is to draw attention to the security challenges inherent in open-source face recognition systems, often integrated into various public applications.

Index Terms—Face Recognition, Security, Attacks

I. INTRODUCTION

Face recognition systems play a crucial role in security [1]. These systems identify and authenticate individuals by analyzing facial traits and have various applications across numerous industries and other sectors [2]. Integrated into several cyber-physical systems within smart cities [3], face recognition contributes to protection and prevention measures. A key aspect of face recognition systems is they are used in authentication, either as standalone solutions or in combination with other authentication methods for two-factor or multi-factor authentication. For instance, Xin et al. [4] proposed a multi-factor authentication system for smart homes that combines face recognition with one-time passwords (OTPs) on smartphones; their system aims to overcome the limitations of traditional authentication methods, such as passwords and PINs, by offering a more secure and convenient alternative. Numerous use cases demonstrate the significance of face recognition systems in our lives [5]–[8]. However, the widespread adoption of face recognition systems also makes them attractive targets for cyber-attacks [6]–[8].

Authentication through face recognition systems involves a series of steps, including inputting facial features via a camera, face detection and localization preprocessing, feature extraction, and decision-making. Attacks can happen during the input stage or within the recognition model. Researchers have been working to develop face recognition attacks and defense protocols to strengthen security [9]–[11]. However, there is limited research on the security of open-source

face recognition systems [12]–[14]. Businesses often adopt open-source solutions to enhance the efficiency and cost-effectiveness of their system development, which necessitates a detailed security analysis.

In this research, we evaluate a popular open-source face recognition system called InsightFace against three blackbox attacks on InsightFace¹: *face swap* [15], *morphing* [16], and *presentation* [17], because of their prevalence in the literature. Face swapping, a type of deepfake attack, has around 246 research articles in Google Scholar within the last 5 years (with the search phrase: “deepfake attacks”). Morphing and presentation attack have 556 and 1900 publications respectively. Additionally, surveys on challenges of face recognition systems [18]–[21] show that these three categories of attacks are the most common forms of attacks on face recognition systems. More importantly, InsightFace uses ArcFace [22], one of the state-of-the-art loss functions used in a convolutional neural networks that outperform a few well-known face recognition models [22]. ArcFace has also been used in face authentication research [23]–[25]. At the time of writing, the InsightFace Github repository² had a significant popularity of about 17,000 stars and 5,000 forks. These statistics motivated us to choose InsightFace as the use case for this study. Our experiments demonstrate that InsightFace may fail to detect images manipulated using the aforementioned attacks, and this highlights the importance of rigorous evaluation and testing of open-source face recognition systems, especially when they play a role in safeguarding user privacy and identity. Independent of our study, InsightFace launched the Face Anti-spoofing Workshop and Challenge [26], which further highlights the interest and the need to augment open-source face recognition systems to defend against various types of attacks, including the ones that we consider in our study.

The two main contributions of this work are as follows:

- Implementation of three blackbox face recognition attacks (*Face Swap*, *Morphing*, and *Presentation*) on InsightFace using reputed public figures.
- Analyses and assessments of InsightFace’s susceptibility when subjected to the three attacks.

¹<https://insightface.ai/>

²<https://github.com/deepinsight/insightface>

The rest of the paper is organized as follows: Section II discusses related work; Section III introduces InsightFace; Section IV presents the experiments, including the dataset, threat model, and implementation; and Section V concludes the paper and discusses limitations and possible future directions.

II. RELATED WORK

Several studies [12]–[14] have analyzed the security vulnerabilities of open-source software projects. Gkortzis et al. [27] created a dataset consisting of vulnerabilities in open-source projects. To generate this dataset, the authors analyzed reports from the National Vulnerability Database (NVD) to identify vulnerable versions of open-source projects. The NVD provides detailed information about the severity of a vulnerability and its potential impact on confidentiality, integrity, and other factors. Prana et al. [5] conducted an analysis of vulnerabilities in open-source libraries used by 450 software projects developed in Java, Python, and Ruby. The authors examined various aspects of the vulnerabilities, including their types, distributions, severity, and persistence. The study found that project activity level, popularity, and developer experience do not necessarily result in better or worse handling of dependency vulnerabilities. Furthermore, the authors discovered that the types of vulnerabilities that are most common across the languages studied are “Denial of Service” and “Information Disclosure.” This finding highlights the need for developers to be particularly vigilant in addressing these types of vulnerabilities when using open-source libraries. Our study is focused on one open-source project, InsightFace, which we evaluate against several types of attacks that are specific to face recognition.

Open-source software development also faces several categories of cyber-attacks. A study by Yuki Matsuo et al. [6] investigated the vulnerability of COVID-Net, an open-source Deep Neural Network (DNN) model. The work focused on backdoor attacks, highlighting the need to consider the vulnerability of DNN models and the importance of incorporating appropriate security measures. Neslihan et al. [7] studied mask attacks on 2D and 3D face recognition systems using a MORPHO mask attacks database. They found that face recognition systems were vulnerable to spoofing attacks and their robustness varied depending on the method and modality used. The authors concluded that robust algorithms are necessary to mitigate the effects of spoofing on face recognition. The authors also showed that texture analysis may reveal more information to detect mask attacks than 3D face shape characteristics. Ulrich et. al [8] studied the vulnerability of biometric systems to morphed face attacks using two new databases of morphed images. They created databases using printed and scanned digitally morphed images, using two scanner types: a flatbed scanner and a line scanner. The paper presented a new database of morphed images, evaluated the vulnerability of different face recognition systems, and conducted a comparative study on different morphed face attack detection algorithms to assess their applicability and generalizability on the scanned morph

face database. Raghavendra et. al [28] examined the vulnerability of an extended multispectral face recognition system to presentation attacks. The system captures face images across various spectral bands and investigates each band’s susceptibility. Experiments were conducted using SpectraCam™, a commercial camera capable of capturing seven different bands. The researchers created face artifacts using laser and inkjet printers and evaluated state-of-the-art Presentation Attack Detection (PAD) algorithms. The findings indicate that the extended multispectral face recognition system is vulnerable to print attacks and suggests difficulty in detecting presentation attacks.

There are various countermeasures to defend against presentation, morphing and deepfake attacks. Presentation attack detection techniques include static texture analysis and convolutional feature extraction, liveness detection, and multi-modal fusion [29], [30]. Principle Component Analysis (PCA) can be used to detect face morphing attacks [31]. Watermark injection can be used to detect any manipulation of the image or video by deepfake algorithms [32]. These techniques are computationally efficient and can be employed on face authentication systems to enhance security. However, our preliminary testing suggests that such countermeasures are not used in the InsightFace project. Therefore, our aim is to encourage the integration of such countermeasures where applicable.

III. INSIGHTFACE

InsightFace is an open-source deep learning library for face recognition, developed at the Chinese University of Hong Kong [33]–[40]. The library implements state-of-the-art face recognition algorithms and provides pre-trained models. InsightFace uses deep convolutional neural networks (CNNs) such as ResNet [41], MobileNet [42] and DenseNet [43] to extract features from face images, and then use those features to identify individuals. The project supports a variety of face recognition tasks, including face verification, face identification, and face clustering.

The InsightFace model uses several advanced techniques to improve the accuracy of face recognition, such as ArcFace loss [22], a loss function that optimizes the angular margin between classes in the feature space. ArcFace was proposed first in 2019 as the state-of-the-art model for face recognition systems. Note that InsightFace does not claim to incorporate security against the three attacks discussed in our study.

Prior to implementing the three aforementioned black box attacks, we conducted a series of preliminary experiments on InsightFace. Our objective was to understand the recognition outcomes of the system given various inputs such as two similar faces, different faces, faces with and without masks, twins, etc. The dataset used in these experiments consisted of images of human figures with different facial features, poses, and expressions collected via web search. An example from the subset of our preliminary experiments is presented in Figure 1. InsightFace outputs a similarity score and a string of text stating whether the faces are of the same person or not. The score (0 to 1) is also an indicator of similarity. Similar



Fig. 1. Analyzing the performance of InsightFace on different well-known faces.

faces will have scores closer to 1, while dissimilar faces will have scores closer to 0.

IV. EXPERIMENT

A. Dataset

We utilized the $r100-w2m$ model provided by the InsightFace project for the purpose of face verification. $r100-w2m$ leverages the larger WebFace-2M dataset [44] for training. The dataset consists of over 2 million images and covers more than 50,000 identities. This dataset expands upon the diversity and complexity of the WebFace-600K dataset, providing an even more challenging setting for the development of advanced face recognition models.

We scrape 15 publicly available images of well-known individuals to analyze the classification outcome of InsightFace³. In order to retain the original characteristics of the images, we did not perform any pre-processing on the images. We manipulate the images using the three attacks, keeping copies of the original image. We input the original image and the manipulated image (implementing the three attacks on the actual image) to the system. The expected output should be a lower similarity score followed by InsightFace’s text output “They are NOT the same person”. The default categorization threshold used is 0.2, which we leave unchanged to ensure that our results reflect the default configuration.

B. Threat Model

The threat model of our attack is presented in Figure 2. For the scope of this research, we experimented with three attacks: *Face Swap* [15], *Morphing* [16], and *Presentation* [17]. For a presentation attack, the attacker requires an image of the legitimate user (i.e., the victim image). The attacker can manipulate that image by printing it and presenting the printed image to the camera. Manipulations can be pixel transformations or print attacks. For the other attacks, the attacker requires both the victim image and the target image. Presumably, the target image is the image that favors the attacker (e.g., an image of the attacker). The attacker’s intention is to modify the victim’s image or to replace the victim’s

³We use the web demonstration, whose source code is available at https://github.com/deepinsight/insightface/tree/f8aa2c1/web-demos/src_recognition

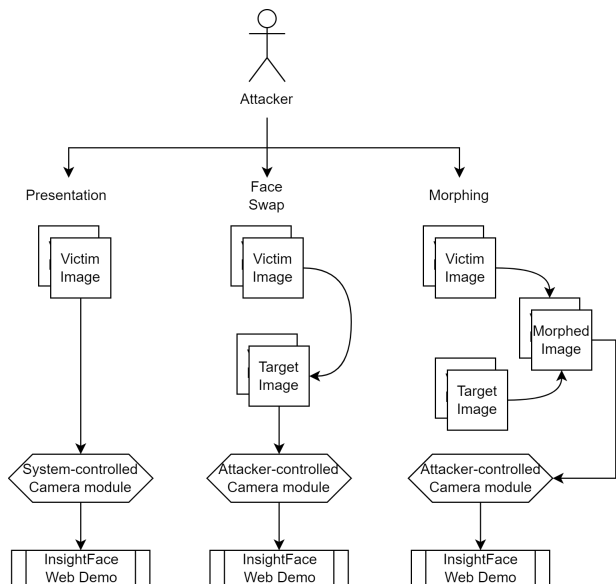


Fig. 2. Threat model considered assumed for the attacks implemented.

image with the target image to spoof the authentication system. Presentation attacks require the attacker to interact with the system’s camera that provides video input to the system. Face Swap and Morphing require a neural network model to replace the target’s facial features with that of victim’s; these attacks require the attacker to either intercept the video feed to the system and inject their own data, or to provide their own video feed to the system (e.g., if the system allows users to authenticate from any device).

In this research, we mimic the role of an ethical hacker performing the attacks presented in Fig 2. As such, we implemented the attacks and input the manipulated images to InsightFace to record the outcomes. The outcomes are presented in Table I.

C. Implementation of Attacks

The implementation of the three attack techniques considered in our threat model is described below.

Face Swap. Face swap is a type of attack in which an attacker replaces the face of a person in an image or

TABLE I

OUTCOMES OF THE TEST CASES ANALYZED BY INSIGHTFACE: THREE BLACK BOX ATTACKS (FACE SWAPPING, MORPHING, AND PRESENTATION) ARE TESTED ON THREE DIFFERENT TEST CASES (A, B, AND C). THE MANIPULATED AND ORIGINAL IMAGES ARE SUPPLIED AS INPUT TO INSIGHTFACE WITH A THRESHOLD VALUE SET TO 0.2. AN OUTCOME MORE THAN 0.2 INDICATES A MATCH.

Testcases	Face Swap	Morphing	Presentation	Expected Outcome
A	They ARE the same person; Score: 0.3053.	They ARE the same person; Score: 0.6421.	They ARE the same person; Score: 0.3399.	They are NOT the same person
B	They ARE the same person; Score: 0.3268.	They ARE the same person; Score: 0.6519.	They ARE the same person; Score: 0.3049.	They are NOT the same person
C	They ARE the same person; Score: 0.3063.	They ARE the same person; Score: 0.4536.	They are NOT the same person; Score: 0.0721.	They are NOT the same person

video with the face of another person. This can be done using various tools and techniques, such as image editing software or deep learning algorithms. For the experiment, we used Faceswapper.ai [45], an online tool that uses generative adversarial networks to swap the face in one image with a face from another image.

Figure 3 shows three face swap attacks. For test cases A, B, and C, we swap the face in ‘Subject 1’ with that in ‘Subject 2’ to create the ‘Manipulated Image’. ‘Subject 1’ as the first input image and ‘Manipulated Image’ as the second input image are then supplied as input to InsightFace. The outcome in correspondence to the inputs are shown in Table I. We observe that although InsightFace outputs lower similarity scores for the face-swapped images, they are still above the threshold to be classified as the same person.

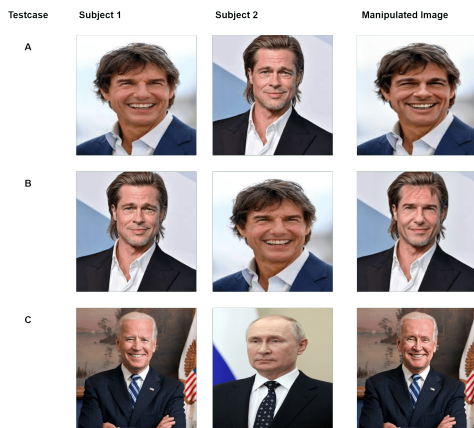


Fig. 3. Sample test cases of Face swap attack.

Morphing. Morphing is a type of attack in which an attacker combines two or more faces to create a new face that looks like a blend of the original faces. We constructed our algorithm by using the Dlib Python library’s 68 face landmarks detection model [46] to detect landmark points on the faces. Delaunay triangulation [47] is used to give better triangular meshes in a two-dimensional image plane to represent a natural face segmentation. The two landmarks are then matched and plotted on top of each other to create a morphed image.

In this study, we develop a sophisticated face-morphing algorithm by leveraging the OpenCV libraries and validated it by conducting a series of experiments on three distinct test cases, labeled A, B, and C. To perform the morphing oper-

ation, using two images named ‘Subject 1’ and ‘Subject 2’, we meticulously align and superimpose the facial landmarks of both subjects to create a composite ‘Manipulated Image’. Upon analysis, we observe a notable difference between face-swapping and face-morphing attacks. Unlike face swapping, the morphing technique preserves characteristic features from both input images, resulting in a more convincing and seamless blend. To further evaluate the effectiveness of the proposed method, we employ the InsightFace recognition system to process both ‘Subject 1’ and ‘Manipulated Image’. The outcomes of the analysis are presented in Table I. We observe that InsightFace classifies the face-morphed images as the same person; we also observe that the similarity scores are higher than for the face swap experiment discussed above. This may indicate that defending against a face morph attack may be more difficult than defending against a face swap attack.




Fig. 4. Sample test cases of face morphing attack.

Presentation Attack. Presentation attacks, also known as spoofing attacks, are attempts to deceive or bypass biometric recognition systems using artificial or manipulated samples [48]. Print attacks (a type of presentation attack) [49] specifically involve the use of printed images, photos, or other physical reproductions of a biometric sample to trick the recognition system into granting access to an unauthorized person [50]. In the context of face recognition systems, a print attack may involve presenting a high-quality photograph of the authentic face in front of the camera.

Similarly, for fingerprint recognition systems, an attacker presents a printed or molded replica of a fingerprint to

deceive the sensor. Since presentation attacks are mostly printed photo attacks, we could not find any open-source algorithm to implement the print attacks. Hence, we developed an algorithm that can generate print images from original images. The algorithm involves converting the original image to grayscale, applying a blur filter, inverting the colors, and enhancing the contrast. These transformations aid in accentuating certain features or artifacts within the image that can indicate a presentation attack.




Fig. 5. Sample test cases of presentation attack.

Figure 5 illustrates three test cases (A, B, C) demonstrating a presentation attack. The manipulated images were generated using the aforementioned custom-developed presentation attack simulation algorithm. As evident from Table 1, InsightFace struggles to differentiate between manipulated and genuine images in the first two test cases. However, in test case C, the algorithm successfully identifies the manipulated image.

V. CONCLUSION

In this research we explore the security of a popular face recognition library, InsightFace, against three black box attacks (face swap, morphing, and presentation) on face recognition systems. We evaluated the outcome of the attacks to determine if the system was able to differentiate between authentic and non-authentic faces. Our preliminary experiments show that InsightFace is susceptible to classifying manipulated images as authentic, and can thus result in false authentication of malicious users. The similarity scores also indicated that morphing attacks are harder to detect compared to face-swap and presentation attacks, given the high similarity score it achieved. False authentication can have serious consequences such as exposure of confidential information, financial loss, or identity theft. At the time of submission of this manuscript, our evaluation was still based on up-to-date code from the InsightFace repository (see Section IV).

In future work, we plan to test other face recognition models such as VGGFace2 [51]. We tested the three most common blackbox attacks. However, there are other types of attacks on face recognition systems such as whitebox attacks using adversarial eyeglasses to spoof the authentication [52]. Whereas we used a small set of public images scraped from the web to conduct our experiments, additional experiments can be performed using standard face recognition datasets. Moreover, the inclusion of faces representing different gender, race and age would help to identify biases and limitations of the systems, improve accuracy by incorporating diverse datasets, and help with building more inclusive and robust face recognition systems.

VI. ACKNOWLEDGEMENT

This research was supported by funding from the NSERC CREATE Cybersecurity Program.

REFERENCES

- [1] M. Rakhra, D. Singh, A. Singh, K. D. Garg, and D. Gupta, "Face recognition with smart security system," in *IEEE International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO)*, 2022, pp. 1–6.
- [2] V. Bruce and A. Young, "Understanding face recognition," *British Journal of Psychology*, vol. 77, no. 3, pp. 305–327, 1986.
- [3] I. Adjabi, A. Ouahabi, A. Benzaoui, and A. Taleb-Ahmed, "Past, present, and future of face recognition: A review," *Electronics*, vol. 9, no. 8, p. 1188, 2020.
- [4] T. Y. Xin, N. Katuk, and A. S. C. M. Arif, "Smart home multi-factor authentication using face recognition and one-time password on smartphone," *International Journal of Interactive Mobile Technologies*, vol. 15, no. 24, p. 33, 2021.
- [5] G. A. A. Prana, A. Sharma, L. K. Shar, D. Foo, A. E. Santosa, A. Sharma, and D. Lo, "Out of sight, out of mind? how vulnerable dependencies affect open-source projects," *Empirical Software Engineering*, vol. 26, pp. 1–34, 2021.
- [6] Y. Matsuo and K. Takemoto, "Backdoor attacks to deep neural network-based system for COVID-19 detection from chest X-ray images," *Applied Sciences*, vol. 11, no. 20, p. 9556, 2021.
- [7] N. Kose and J.-L. Dugelay, "On the vulnerability of face recognition systems to spoofing mask attacks," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 2357–2361.
- [8] U. Scherhag, R. Raghavendra, K. B. Raja, M. Gomez-Barrero, C. Rathgeb, and C. Busch, "On the vulnerability of face recognition systems towards morphed face attacks," in *IEEE International Workshop on Biometrics and Forensics*, 2017, pp. 1–6.
- [9] H. H. Nguyen, J. Yamagishi, I. Echizen, and S. Marcel, "Generating master faces for use in performing wolf attacks on face recognition systems," in *IEEE International Joint Conference on Biometrics (IJCB)*, 2020, pp. 1–10.
- [10] H. H. Nguyen, S. Marcel, J. Yamagishi, and I. Echizen, "Master face attacks on face recognition systems," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 4, no. 3, pp. 398–411, 2022.
- [11] E. Sarkar, H. Benkraouda, G. Krishnan, H. Gamil, and M. Maniatakos, "Facehack: Attacking facial recognition systems using malicious facial characteristics," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 4, no. 3, pp. 361–372, 2021.
- [12] S.-F. Wen, "Software security in open source development: A systematic literature review," in *IEEE Conference of Open Innovations Association (FRUCT)*, 2017, pp. 364–373.
- [13] A. Boulanger, "Open-source versus proprietary software: Is one more reliable and secure than the other?" *IBM Systems Journal*, vol. 44, no. 2, pp. 239–248, 2005.
- [14] A. Adewumi, S. Misra, N. Omoregbe, B. Crawford, and R. Soto, "A systematic literature review of open source software quality assessment models," *SpringerPlus*, vol. 5, no. 1, pp. 1–13, 2016.
- [15] Y. Zhang, L. Zheng, and V. L. Thing, "Automated face swapping and its detection," in *IEEE International Conference on Signal and Image Processing*, 2017, pp. 15–19.

- [16] S. Venkatesh, R. Ramachandra, K. Raja, and C. Busch, "Face morphing attack generation and detection: A comprehensive survey," *IEEE Transactions on Technology and Society*, vol. 2, no. 3, pp. 128–145, 2021.
- [17] A. Husseis, J. Liu-Jimenez, I. Goicoechea-Telleria, and R. Sanchez-Reillo, "A survey in presentation attack and presentation attack detection," in *IEEE International Carnahan Conference on Security Technology (ICCST)*, 2019, pp. 1–13.
- [18] U. Scherhag, C. Rathgeb, J. Merkle, R. Breithaupt, and C. Busch, "Face recognition systems under morphing attacks: A survey," *IEEE Access*, vol. 7, pp. 23 012–23 026, 2019.
- [19] M. Wang and W. Deng, "Deep face recognition: A survey," *Neuro-computing*, vol. 429, pp. 215–244, 2021.
- [20] L. Li, X. Mu, S. Li, and H. Peng, "A review of face recognition technology," *IEEE Access*, vol. 8, pp. 139 110–139 120, 2020.
- [21] F. Vakhshiteh, A. Nickabadi, and R. Ramachandra, "Adversarial attacks against face recognition: A comprehensive study," *IEEE Access*, vol. 9, pp. 92 735–92 756, 2021.
- [22] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699.
- [23] B. Huang, Z. Wang, G. Wang, K. Jiang, Z. He, H. Zou, and Q. Zou, "Masked face recognition datasets and validation," in *IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1487–1491.
- [24] W. Yao, V. Varkarakis, G. Costache, J. Lemley, and P. Corcoran, "Toward robust facial authentication for low-power edge-AI consumer devices," *IEEE Access*, vol. 10, pp. 123 661–123 678, 2022.
- [25] T.-V. Dang, "Smart home management system with face recognition based on ArcFace model in deep convolutional neural network," *Journal of Robotics and Control*, vol. 3, no. 6, pp. 754–761, 2022.
- [26] D. Wang, J. Guo, Q. Shao, H. He, Z. Chen, C. Xiao, A. Liu, S. Escalera, H. J. Escalante, L. Zhen, J. Wan, and J. Deng, "Wild face anti-spoofing challenge 2023: Benchmark and results," *arXiv preprint arXiv:2304.05753*, 2023.
- [27] A. Gkortzis, D. Mitropoulos, and D. Spinellis, "Vulnoss: a dataset of security vulnerabilities in open-source systems," in *International Conference on Mining Software Repositories*, 2018, pp. 18–21.
- [28] R. Raghavendra, K. B. Raja, S. Venkatesh, F. A. Cheikh, and C. Busch, "On the vulnerability of extended multispectral face recognition systems towards presentation attacks," in *IEEE International Conference on Identity, Security and Behavior Analysis*, 2017, pp. 1–8.
- [29] F. Abdullakutty, E. Elyan, and P. Johnston, "A review of state-of-the-art in face presentation attack detection: From early development to advanced deep learning and multi-modal fusion methods," *Information Fusion*, vol. 75, pp. 55–69, 2021.
- [30] S. Marcel, J. Fierrez, and N. Evans, *Handbook of Biometric Anti-Spoofing: Presentation Attack Detection and Vulnerability Assessment*. Springer Nature, 2023.
- [31] L. Dargaud, M. Ibsen, J. Tapia, and C. Busch, "A principal component analysis-based approach for single morphing attack detection," in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, January 2023, pp. 683–692.
- [32] Y. Zhao, B. Liu, M. Ding, B. Liu, T. Zhu, and X. Yu, "Proactive deepfake defence via identity watermarking," in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2023, pp. 4602–4611.
- [33] X. Ren, A. Lattas, B. Gecer, J. Deng, C. Ma, and X. Yang, "Facial geometric detail recovery via implicit representation," in *IEEE International Conference on Automatic Face and Gesture Recognition*, 2023.
- [34] J. Guo, J. Deng, A. Lattas, and S. Zafeiriou, "Sample and computation redistribution for efficient face detection," *arXiv preprint arXiv:2105.04714*, 2021.
- [35] B. Gecer, J. Deng, and S. Zafeiriou, "Ostec: One-shot texture completion," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [36] X. An, X. Zhu, Y. Xiao, L. Wu, M. Zhang, Y. Gao, B. Qin, D. Zhang, and F. Ying, "Partial FC: Training 10 million identities on a single machine," *arXiv preprint arXiv:2010.05222*, 2020.
- [37] J. Deng, J. Guo, T. Liu, M. Gong, and S. Zafeiriou, "Sub-center ArcFace: Boosting face recognition by large-scale noisy web faces," in *IEEE Conference on European Conference on Computer Vision*, 2020.
- [38] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-shot multi-level face localisation in the wild," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [39] J. Guo, J. Deng, N. Xue, and S. Zafeiriou, "Stacked dense u-nets with dual transformers for robust face alignment," in *British Machine Vision Conference*, 2018.
- [40] J. Deng, A. Roussos, G. Chrysos, E. Ververas, I. Kotsia, J. Shen, and S. Zafeiriou, "The Menpo benchmark for multi-pose 2D and 3D facial landmark localisation and tracking," *International Journal of Computer Vision*, 2018.
- [41] S. Targ, D. Almeida, and K. Lyman, "Resnet in resnet: Generalizing residual architectures," *arXiv preprint arXiv:1603.08029*, 2016.
- [42] S. Chen, Y. Liu, X. Gao, and Z. Han, "Mobilefacenet: Efficient CNNs for accurate real-time face verification on mobile devices," in *Chinese Conference on Biometric Recognition*. Springer, Aug 2018, pp. 428–438.
- [43] Y. Zhu and S. Newsam, "Densenet for dense flow," in *IEEE International Conference on Image Processing (ICIP)*, 2017, pp. 790–794.
- [44] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Lu, D. Du *et al.*, "Webface260m: A benchmark unveiling the power of million-scale deep face recognition," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 492–10 502.
- [45] FaceSwapper, "Swap photo video face online free." [Online]. Available: <https://faceswapper.ai/>
- [46] Dlib, "Face landmark detection." [Online]. Available: http://dlib.net/face_landmark_detection.py.html
- [47] Y. Xiao and H. Yan, "Facial feature location with delaunay triangulation/voronoi diagram calculation," in *ACM International Conference Proceeding Series*, vol. 147. Citeseer, 2001, pp. 103–108.
- [48] R. Raghavendra, K. B. Raja, S. Venkatesh, F. A. Cheikh, and C. Busch, "On the vulnerability of extended multispectral face recognition systems towards presentation attacks," in *IEEE International Conference on Identity, Security and Behavior Analysis (ISBA)*, 2017, pp. 1–8.
- [49] A. Anjos and S. Marcel, "Counter-measures to photo attacks in face recognition: a public database and a baseline," in *IEEE International Joint Conference on Biometrics*, 2011, pp. 1–7.
- [50] R. Ramachandra and C. Busch, "Presentation attack detection methods for face recognition systems: A comprehensive survey," *ACM Computing Surveys*, vol. 50, no. 1, pp. 1–37, 2017.
- [51] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vggface2: A dataset for recognising faces across pose and age," in *IEEE International Conference on Automatic Face & Gesture Recognition*, 2018, pp. 67–74.
- [52] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter, "Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition," in *ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 1528–1540.