

Family Resemblance or Fraud? Face Morphing Attacks on Kinship Verification

Gargi S Yeole, Aarthi S, Shalvika Srivastav and Akshay Agarwal

Trustworthy BiometraVision Lab,

Department of Data Science and Engineering, IISER Bhopal, India

Abstract—Kinship verification using facial images is widely applied in forensic analysis, immigration, and child trafficking prevention. However, deep learning models for kinship verification are vulnerable to morphed images, where the facial features of two individuals are blended to create realistic but fake images. This work investigates the influence of different morphing ratios (95% child + 5% random parent to 50% child + 50% random parent) on kinship verification algorithms. Developing a morphed dataset allows us to experiment with deep learning and kinship-specific models on original and morphed child images to determine the threshold beyond which non-kin morphs are labeled kin. The experiments show a continuous rise in misclassification rates with the increasing percentage of parental features in morphed images, underscoring the difficulties encountered by current kinship verification systems. It is to be noted that the current study is the first to present significant insights into the vulnerability of existing kinship verification models against different morph ratios. It highlights the necessity for more effective verification methods to counter the risks associated with facial morphing in real-world applications.

I. INTRODUCTION

Face is one of the most widely used biometric traits, and a significant and growing application of facial biometrics is kinship verification, where the goal is to determine if two people are biologically related. This technology can help build family trees and measure how similar family members are based on facial traits. Moreover, it is crucial in several real-world applications such as forensic analysis [20], missing person identification [1], [6], and border control. By analyzing facial similarities, automated kinship verification systems aim to infer familial ties, offering a non-invasive and scalable method for identity validation. With increasing reliance on AI-based verification systems in security and legal domains, robust kinship verification becomes essential to ensure trustworthy decision-making [32]. These systems must remain reliable in various real-world scenarios while being resilient to manipulation.

In this research, we address a critical security concern in kinship verification: the susceptibility of deep learning models to face morphing attacks [2], [3], [4], [14]. Specifically, we explore how morphed images, synthetically generated by blending features of a child with those of a random parent, fool kinship verification systems into wrongly predicting a biological relationship. Such deception poses serious risks in contexts such as child trafficking and immigration fraud, where verifying family ties is vital for safety and justice. To address this, we construct a comprehensive morphed dataset

at various morphing ratios (e.g., 95-5, 90-10, 80-20, 70-30, 60-40, and 50-50, through blending of a child and a random parent images) and evaluated the robustness of different kinship verification models trained with two different loss optimizing strategies. We then assess the robustness of the model and analyze the morphing thresholds (morph ratio) at which it incorrectly classifies non-kin as kin. This study uncovers the vulnerabilities in current kinship models and highlights the need for morphing-aware verification systems.

II. RELATED WORK

Early kinship verification research often uses handcrafted features such as Local Binary Patterns (LBP) [5] and Histogram of Oriented Gradients (HOG) [10], followed by traditional classifiers such as Support Vector Machines (SVM) and Random Forests [7]. These approaches offer limited performance due to their inability to capture complex facial variations. With the emergence of deep learning, models such as DeepFace [16] and VGG-Face [26], research started to leverage convolutional neural networks (CNNs) to extract hierarchical facial features. The availability of kinship-specific datasets, including KinFaceW [23] and TSKinFace [25], further promotes data-driven kinship modeling in the research community.

Recent techniques introduce Siamese networks [22] and supervised contrastive learning [31] to improve verification accuracy by learning discriminative embeddings and contrasting kin and non-kin pairs. OR2Net [18] proposes an online re-weighting strategy that emphasizes kin-related facial regions, achieving state-of-the-art results. However, most existing research focuses exclusively on genuine kin relationships and overlooks the threat of morphed images. Morphing attacks gain attention in face recognition studies, where researchers develop GAN-based morphing detection methods [29] and Transformer-based spoofing defenses [30], but kinship verification remains largely unexplored in this context. Our work fills this gap by systematically analyzing how kinship models respond to morphing-based deception. We examine different morphing ratios to identify how models begin misclassifying non-kin pairs as kin, revealing critical weaknesses in current verification systems and emphasizing the need for morphing-aware defenses. The different morph ratios will ensure a trade-off between an attack image's naturalness and fooling rate.




Fig. 1. Example morphed images generated at different morphing ratios where a non-kin child aims to gain the kinship identity of a random parent.

III. PROPOSED FACE MORPH DATASET

This study utilizes the TSKinFace dataset [25] for kinship verification and kinship morph generation. The TSKinFace dataset [25] is a benchmark dataset for kinship verification, containing facial images of Father (F), Mother (M), Son (S), and Daughter (D). We organize these into three categories: Father-Mother-Daughter (FMD), Father-Mother-Son (FMS), and Father-Mother-Son-Daughter (FMSD), each containing 822, 855, and 912 images, respectively. To evaluate the robustness of kinship verification models under adversarial conditions, we propose a novel morphed dataset generated using two approaches: (1) traditional image morphing via OpenCV [8] and (2) a generative model-based technique using Denoising Diffusion Probabilistic Models (DDPMs) [13], [27]. Both methods aim to synthetically blend the facial features of a child and a random non-biological parent from the TSKinFace dataset to create realistic, deceptive morphed images. These morphed images are then categorized similarly to the original TSKinFace dataset, that is, FMD, FMS, and FMSD, facilitating consistent evaluation.

In the traditional OpenCV method, we adopt a simplified face-morphing approach based on pixel-level blending to generate synthetic facial images. Two aligned face images representing a child and a random parent are combined by computing a weighted average of pixel intensities. In addition to the traditional method, we introduce a second variant of the morphed dataset using a diffusion-based refinement approach [27]. Specifically, we first create intermediate morphs by linearly blending a child's and a parent's image at varying ratios. To enhance photorealism, these blended images are then passed through a pretrained stable diffusion [27], [28] image-to-image model. The visual contribution of each identity is controlled through the blending weights, effectively simulating morphing ratios similar to those used in traditional approaches.

To simulate subtle and stronger levels of identity blending, we generate morphed images at six morphing ratios: 95% Child + 5% Random Parent, 90% Child + 10% Random Parent, 80% Child + 20% Random Parent, 70% Child + 30% Random Parent, 60% Child + 40% Random Parent, and 50% Child + 50% Random Parent. These variations allow us to study the threshold (morph ratio) at which kinship verification models begin misclassifying non-kin pairs as kin, highlighting vulnerabilities in real-world scenarios such as human trafficking or identity fraud. Figure 1 illustrates

sample morphed images using OpenCV at different blending ratios, showcasing how a non-kin child can gradually adopt features of a random parent to deceive a verification model. Unlike existing face-swapping or deepfake methods [15], [19], which often aim to transfer identity or attributes without explicit control over identity proportion, our dataset introduces graded control over feature dominance. This provides a more nuanced testbed for analyzing kinship verification robustness under morphing attacks.

IV. KINSHIP VERIFICATION MODELS

We employ a Siamese network [9] architecture with three different CNN backbones: Custom CNN, ResNet50 [12], and ConvNeXt-Tiny [11] models. This network processes a pair of images (parent and child) through identical feature extractors. The extracted embeddings are then compared using distance metrics such as the L2 norm and feature concatenation to determine the probability of kinship. In the feature concatenation setup, a pair of images (parent and child) is passed through identical feature extractors to obtain embeddings. These are then concatenated and passed through a fully connected layer to determine the probability of the image pair being a kinship. The architecture for the verification of kinship based on L2 norms and concatenation-based characteristics is shown in Figure 2. Additionally, we evaluate kinship verification using GLANet [17] and MagFace [21], both leveraging feature concatenation. We further include experiments with the KFC model [24], a widely adopted framework for kinship verification tasks. The vulnerability assessment of various models ensures that the findings presented in this paper are robust and highlight the need to develop a robust kinship verification architecture. The used kinship verification models are described as follows:

- **Custom CNN (C-CNN):** A lightweight 12-layer network consisting of three convolutional layers (with 64, 128, and 256 filters), followed by ReLU activations and MaxPooling. The final fully connected layer outputs a 128-dimensional feature vector. Features extracted from the parent and child images are concatenated and passed through dense layers for kinship classification.
- **ResNet-50:** A deep residual network [12] pre-trained on ImageNet, used by removing the classification head. The resulting 2048-dimensional embeddings from both images are concatenated and fed to a classification head for kinship verification.




Fig. 2. Proposed kinship verification architecture using L2 distance and feature concatenation.

- **ConvNeXt-Tiny**: We extract 768-dimensional feature from this transformer styled architecture [11] and concatenated for classification network training.
- **GLANet (Global-Local Attention Network)**: In our implementation of GLANet [17], we incorporate two complementary feature extractors: ResNet-50 for capturing local facial details, and PVTv2-B0 (a transformer-based backbone) for modeling global contextual features. Features from both backbones are projected into a lower-dimensional space and concatenated before being embedded into a joint representation.
- **MagFace**: This model [21] produces magnitude-aware embeddings tailored for facial identity, within a Siamese network framework. Each image pair is passed through a shared MagFace backbone (iResNet-100), followed by a fully connected layer to generate compact descriptors.
- **Kinship Verification with Fair Contrastive Loss and Multi-Task Learning (KFC Model)**: It integrates Fair Contrastive Loss (FCL) to reduce bias and Multi-Task Learning (MTL) to jointly learn kinship classification and facial attribute prediction, improving robustness and fairness [24].

A. Model Learning Strategy

We train the deep-learning models for kinship verification using real parent-child pairs and test their robustness against morphed child images. The kinship verification framework leverages CNN-based and transformer-based architectures and employs two primary strategies for learning facial similarity: L2 distance with contrastive loss and feature concatenation with binary cross-entropy loss. The standard training configuration includes a batch size 32, a learning rate 0.0001, binary cross-entropy loss, and the Adam optimizer. The KFC model is trained with a batch size of 32, and fairness-aware contrastive loss is used to optimize parameters using an Adam optimizer with a learning rate of 10^{-6} . This learning strategy allows us to assess how different feature extractors perform under adversarial testing conditions (i.e., with morphed images) and how well they generalize beyond real kinship pairs. The models are trained with clean kinship pairs but evaluated on clean (100:00 ratio) and different morph ratio images.

V. KINSHIP VERIFICATION RESULTS AND ANALYSIS

This section presents the performance of various models for kinship verification on the morphed datasets, with a

TABLE I
FOOLING RATES (%) OF KINSHIP VERIFICATION MODELS ON OPENCV AND DDPM MORPHED DATASETS USING L2 NORM STRATEGY.

Ratio	OpenCV			DDPM		
	C-CNN	ResNet-50	ConvNeXt	C-CNN	ResNet-50	ConvNeXt
100:00	16.82	22.40	16.34	16.82	22.40	16.34
95:05	47.17	43.89	46.66	34.52	44.53	16.50
90:10	50.98	47.28	51.54	39.82	46.21	18.90
80:20	73.26	52.77	59.61	47.56	50.89	20.23
70:30	78.62	66.52	68.72	53.47	53.79	22.58
60:40	86.13	72.61	75.79	59.50	59.56	31.39
50:50	91.03	78.58	86.96	68.51	70.15	49.88

systematic vulnerability analysis based on morphing ratios using the fooling rate, which is computed as follows:

$$\text{Fooling Rate (\%)} = \frac{\text{No. of Non-Kin Images Classified as Kin}}{\text{Total No. of Non-Kin Images}} \times 100$$

Table I displays the fooling rates of kinship verification models on the clean (100:00) and morphed dataset generated by OpenCV and DDPM, evaluated using L2 distance. As observed, the fooling rate increases with a higher contribution of parental features in the morph, showing that models are progressively more susceptible to synthetic manipulations. Among all configurations, C-CNN at the 50:50 morphing ratio yields the highest fooling rate of 91.03% on the OpenCV dataset, demonstrating extreme susceptibility. ConvNeXt showcases the second-highest fooling rate of 86.96% for the OpenCV dataset, but on the other hand, it has the lowest fooling rate of 49.88% for the DDPM dataset. The increasing trends and significant vulnerability across all models and strategies highlight the serious risk that morphed images pose to kinship verification systems. These findings reinforce the need for more robust and morph-resilient kinship verification models, especially in contexts such as child trafficking detection, where generative morphing attacks can exploit system vulnerabilities.

Table II shows the kinship verification fooling rates of various models trained using the feature concatenation strategy on the proposed OpenCV morphed dataset. Again, as the proportion of parental features increases in the morphed child image, all models show a notable increase in the fooling rate, meaning the models are progressively more likely to classify non-biological pairs as kin incorrectly. The seriousness of this can be seen from the fact that even when the morph ratios are high (say 50-50), the generated morph images do not contain visual artifacts to be caught by any external examiner. Interestingly, sophisticated architectures, including

TABLE II
FOOLING RATES (%) AT VARIOUS MORPHING RATIOS FOR DIFFERENT MODELS ON OPENCV AND DDPM DATASETS USING *FEATURE*
CONCATENATION STRATEGY.

Ratio	ResNet50		C-CNN		ConvNeXt		GLANet		MagFace	
	OpenCV	DDPM	OpenCV	DDPM	OpenCV	DDPM	OpenCV	DDPM	OpenCV	DDPM
100:00	32.12	32.12	23.36	23.13	24.82	24.82	10.21	09.98	22.40	21.36
95:05	41.12	42.43	47.06	58.11	50.37	28.42	28.20	49.13	32.26	46.78
90:10	43.57	44.45	49.82	60.05	55.01	29.01	34.85	52.60	38.56	48.70
80:20	58.12	50.64	55.41	65.62	62.33	31.59	42.36	54.09	44.73	51.18
70:30	65.78	58.87	60.28	69.66	65.96	37.95	49.68	56.80	53.19	55.22
60:40	72.44	63.50	68.87	74.43	75.59	44.48	58.55	59.26	65.68	58.42
50:50	78.21	75.12	70.83	79.31	79.21	51.03	62.34	61.33	75.15	65.91




Fig. 3. Fooling Rates (%) of KFC Model on OpenCV Dataset.

transformer-based models such as MagFace and GLANet, demonstrate significant susceptibility, leaving a substantial security gap in the kinship verification systems. Specifically, MagFace achieves a fooling rate of 75.15% and 65.91% on OpenCV and DDPM morphed datasets, respectively. ConvNeXt-Tiny surpasses all other models, with the highest fooling rate of 79.21% at the 50:50 morphing ratio in the OpenCV morphed dataset. Moreover, the vulnerability trend is similar in the DDPM morphed dataset; we observe that C-CNN gets fooled most at a 50:50 morph ratio. As morphing approaches increasingly balance child and parent features, the boundary between genuine kin and forged similarity becomes blurred, particularly evident in the 50:50 ratio, where model predictions are most compromised.

Fig. 3 presents the fooling rates of the KFC model on the OpenCV morphed datasets across different blending ratios. It has the highest accuracy of 54.86% for the OpenCV morph dataset at a 50:50 morph ratio. The KFC model demonstrates a similar vulnerability trend to other models, showcasing that no verification model is secure.

Discussion: The results underline the importance of incorporating morph-aware defense mechanisms in kinship verification pipelines, especially in sensitive domains such as child trafficking detection, where adversarial manipulation can severely undermine system reliability. As the percentage of parent features in morphed images increases, the model becomes more confident in classifying the image as *kin*. At a 95% child + 5% parent ratio, the accuracy of kinship verification is low, indicating that the model struggles to recognize kinship due to the high dominance of child features, which is genuinely a non-kin to a random parent. Conversely,

at a 50% child + 50% parent ratio, the accuracy is high, as the increased parental features make the model easy to misclassify non-kin as kin. This analysis demonstrates that as the contribution of parental features increases in morphed images, the model shifts towards classifying them as kin with greater accuracy. Also, the L2 strategy consistently shows higher fooling rates than feature concatenation for the OpenCV morphed dataset, suggesting that simple distance metrics can also be more vulnerable to morphing attacks.

VI. CONCLUSION

This study explores vulnerabilities of deep kinship verification models by proposing multiple morphed parent-child datasets. Unlike traditional morphing methods, we systematically vary morphing ratios to assess their impact on kinship verification accuracy. We train and evaluate various deep learning models, covering pure convolutional and transformer models, to determine their effectiveness in distinguishing kin from non-kin. For similarity computation, we employ both L2 distance and feature concatenation strategies. Our results reveal a critical threshold at which non-kin morphed images, containing increasing proportions of random parental features, are misclassified as kin. This trend becomes especially evident at higher morphing ratios, such as 50:50, where fooling rates peak across all models. Being the first-ever such detailed study, our findings have critical implications for applications such as forensic investigations, missing person identification, and child trafficking detection, where precise kinship verification is essential. Therefore, the high susceptibility of models to morphed images underscores the urgent need for robust verification frameworks. While this study offers key insights into current vulnerabilities, future work could explore further refinement of kinship verification models, including adversarial training, attention-based filtering, and feature disentanglement, which could help develop models more resilient to manipulation and deception.

ACKNOWLEDGEMENT

Akshay Agarwal is partially supported through the ECRG grant of ANRF. The INSPIRE fellowship of DST, India, partially supports Aarthi S.

ETHICAL STATEMENT

Author/Reviewer Checklist:

- 1) Did you read the Ethical Impact Statement Guidelines document (provided above)? **Yes**
- 2) Is it clear that all studies and procedures described in the paper were approved (or exempted) by a valid ethical review board? Alternatively, is a valid and sufficient justification provided for why the oversight of an ethical review board was not required? **Yes**
- 3) Does the ethical impact statement provide a clear, complete, and balanced discussion of the potential risks of individual harm and negative societal impacts associated with the research? Note that this includes harm to research participants as well as harm to other individuals that may be affected by use, misuse, or misunderstanding of the research. **Yes**
- 4) Does the ethical impact statement describe reasonable, valid, and sufficient use of risk-mitigation strategies by the authors to lessen these potential risks? Alternatively, if relevant strategies were not used, is a valid and sufficient justification for this provided? **Yes**
- 5) Does the ethical impact statement provide a valid and sufficient justification for how/why the potential risks of the research are outweighed by the risk-mitigation strategies and potential benefits of the research? Note that papers with serious potential risks that are not outweighed by risk-mitigation strategies and potential benefits may be rejected. **Yes**
- 6) If the paper involves human subjects, are all of the following sub-boxes checked?
 - a) Does the main paper describe whether/how informed consent and/or assent were obtained from participants? If consent and/or assent were fully or partially obtained, were the methods used to do so valid? If not fully obtained, does the ethical impact statement provide a valid and sufficient justification for this? **Not Applicable**
 - b) Does the main paper state whether the participants explicitly consented to the use of their data in the manner described in the paper? For example, if the data was or will be shared with third parties, does it state that the participants explicitly agreed to this sharing? If some uses were not explicitly consented to, does the ethical impact statement provide a valid and sufficient justification for this? **Not Applicable**
 - c) Does the main paper explain whether/how participants were compensated? If participants were compensated, does the ethical impact statement provide a valid and sufficient justification for the form and amount of compensation provided? **Not Applicable**
 - d) If the research involves any special or vulnerable populations (e.g., minors, elderly individuals, prisoners, refugees and migrants, individuals with disabilities, individuals with mental illness, or

patients in medical settings), does the ethical impact statement provide a valid and sufficient explanation of how the rights, well-being, and autonomy of such individuals were safeguarded in the research? **Not Applicable**

REFERENCES

- [1] K. A. Afrifah, A. Badu-Boateng, S. Antwi-Akomeah, E. E. Motey, E. Boampong, P. Twumasi, P. P. Sampene, and A. Donkor. Forensic identification of missing persons using dna from surviving relatives and femur bone retrieved from salty environment. *Journal of Forensic Science and Medicine*, 6(1):40–44, 2020.
- [2] A. Agarwal and N. Ratha. Manipulating faces for identity theft via morphing and deepfake: Digital privacy. In *Handbook of Statistics*, volume 48, pages 223–241. Elsevier, 2023.
- [3] A. Agarwal and N. Ratha. Face morphing detection in social media content. In *IEEE International Conference on Image Processing*, pages 801–806, 2024.
- [4] A. Agarwal, R. Singh, M. Vatsa, and A. Noore. Magnet: Detecting digital presentation attacks on face recognition. *Frontiers in Artificial Intelligence*, 4:643424, 2021.
- [5] T. Ahonen, A. Hadid, and M. Pietikainen. Face description with local binary patterns: Application to face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 28(12):2037–2041, 2006.
- [6] G. Biosa, D. Giurghita, E. Alladio, M. Vincenti, and T. Neocleous. Evaluation of forensic data using logistic regression-based classification methods and an r shiny implementation. *Frontiers in Chemistry*, 8:738, 2020.
- [7] E. Boutellaa, M. B. López, S. Ait-Aoudia, X. Feng, and A. Hadid. Kinship verification from videos using spatio-temporal texture features and deep learning. *arXiv preprint arXiv:1708.04069*, 2017.
- [8] G. Bradski and A. Kaehler. *Learning OpenCV: Computer vision with the OpenCV library*. "O'Reilly Media, Inc.", 2008.
- [9] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 539–546, 2005.
- [10] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, volume 1, pages 886–893. Ieee, 2005.
- [11] J. Feng, H. Tan, W. Li, and M. Xie. Conv2next: Reconsidering conv next network design for image recognition. In *2022 International Conference on Computers and Artificial Intelligence Technologies (CAIT)*, pages 53–60. IEEE, 2022.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [14] M. Huber, P. C. Neto, A. F. Sequeira, and N. Damer. Fx-mad: Frequency-domain explainability and explainability-driven unsupervised detection of face morphing attacks. In *Winter Conference on Applications of Computer Vision*, pages 766–776, 2025.
- [15] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [16] N. Kohli, M. Vatsa, R. Singh, A. Noore, and A. Majumdar. Hierarchical representation learning for kinship verification. *IEEE Transactions on Image Processing*, 26(1):289–302, 2016.
- [17] D. Li and X. Jiang. Kinship verification method of face image deep feature fusion. *Acad. J. Sci. Technol.*, 5(1):57–62, 2023.
- [18] H. Li, X. Zhao, M. Wang, H. Song, and F. Sun. Or2net: Online re-weighting relation network for kinship verification. *Expert Systems with Applications*, 255:124815, 2024.
- [19] L. Li, J. Bao, H. Yang, D. Chen, and F. Wen. Faceshifter: Towards high fidelity and occlusion aware face swapping. *arXiv preprint arXiv:1912.13457*, 2019.
- [20] R. Li, N. Wang, Y. Zang, J. Liu, E. Wu, R. Wu, and H. Sun. Easykin: a flexible and user-friendly online tool for forensic kinship testing and missing person identification. *International Journal of Legal Medicine*, 137(6):1671–1681, 2023.

- [21] Q. Meng, S. Zhao, Z. Huang, and F. Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14225–14234, 2021.
- [22] A. Nandy and S. S. Mondal. Kinship verification using deep siamese convolutional neural network. In *2019 14th IEEE international conference on automatic face & gesture recognition (FG 2019)*, pages 1–5. IEEE, 2019.
- [23] H. Y. Patil and A. Chandra. Deep learning based kinship verification on kinfacew-i dataset. In *TENCON 2019-2019 IEEE Region 10 Conference (TENCON)*, pages 2529–2532. IEEE, 2019.
- [24] J. L. Peng, K. W. Chang, and S.-H. Lai. Kfc: Kinship verification with fair contrastive loss and multi-task learning. *arXiv preprint arXiv:2309.10641*, 2023.
- [25] X. Qin, X. Tan, and S. Chen. Tri-subject kinship verification: Understanding the core of a family. *IEEE Transactions on Multimedia*, 17(10):1855–1867, 2015.
- [26] R. F. Rachmadi, I. K. E. Purnama, S. M. S. Nugroho, and Y. K. Suprpto. Image-based kinship verification using dual vgg-face classifie. In *2020 IEEE international conference on internet of things and intelligence system (IoTIS)*, pages 123–128. IEEE, 2021.
- [27] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [28] Stability.AI. Stable diffusion public release. <https://stability.ai/news/stable-diffusion-public-release>, 2022.
- [29] S. Venkatesh, R. Ramachandra, K. Raja, and C. Busch. Face morphing attack generation and detection: A comprehensive survey. *IEEE transactions on technology and society*, 2(3):128–145, 2021.
- [30] N. Zhang, X. Liu, X. Li, and G.-J. Qi. Morphganformer: Transformer-based face morphing and de-morphing. *arXiv preprint arXiv:2302.09404*, 2023.
- [31] X. Zhang, X. Min, X. Zhou, and G. Guo. Supervised contrastive learning for facial kinship recognition. In *2021 16th IEEE international conference on automatic face and gesture recognition (FG 2021)*, pages 01–05. IEEE, 2021.
- [32] K. Zhang¹², Y. Huang, C. Song, H. Wu, L. Wang, and S. M. Intelligence. Kinship verification with deep convolutional neural networks. In *In British machine vision conference (BMVC)*, 2015.