

MS-GAGA: Metric-Selective Guided Adversarial Generation Attack

Dion J. X. Ho^{†1}

dh3065@columbia.edu

Gabriel Lee Jun Rong^{†2}

2301906@sit.singaporetech.edu.sg

Niharika Shrivastava^{*2}

a103955@singaporetech.edu.sg

Harshavardhan Abichandani^{*3}

harshavardhan@u.nus.edu

Pai Chet Ng²

paichet.ng@singaporetech.edu.sg

Xiaoxiao Miao⁴

xiaoxiao.miao@dukekunshan.edu.cn

¹ Columbia University

New York, USA

² Singapore Institute of Technology

Singapore

³ Independent Researcher

⁴ Duke Kunshan University

China

Abstract

We present **MS-GAGA** (Metric-Selective Guided Adversarial Generation Attack), a two-stage framework for crafting transferable and visually imperceptible adversarial examples against deepfake detectors in black-box settings. In Stage 1, a dual-stream attack module generates adversarial candidates: MNTD-PGD applies enhanced gradient calculations optimized for small perturbation budgets, while SG-PGD focuses perturbations on visually salient regions. This complementary design expands the adversarial search space and improves transferability across unseen models. In Stage 2, a metric-aware selection module evaluates candidates based on both their success against black-box models and their structural similarity (SSIM) to the original image. By jointly optimizing transferability and imperceptibility, MS-GAGA achieves up to 27% higher misclassification rates on unseen detectors compared to state-of-the-art attacks.

1 Introduction

1.1 Background and Motivation

Rapid advances in artificial intelligence, particularly deep learning, have facilitated the creation of synthetic media known as "deepfakes". Modern deepfakes are often generated using sophisticated diffusion-based models [19, 29] which iteratively refine synthetic media until it is virtually indistinguishable from authentic content. The quality of these generated images and videos has reached a level where they are highly realistic and difficult to distinguish

from real face images, blurring the line between reality and fabrication [7, 23]. This has precipitated a host of severe societal and ethical challenges. Deepfakes have been weaponized to erode public trust, spread disinformation in political contexts, violate personal privacy through non-consensual content, and perpetrate widespread fraud [26].

In response to these growing threats, the development of accurate and reliable deepfake detection algorithms has become a critical and urgent challenge for researchers, policymakers, and industry stakeholders alike. Recent years have witnessed the emergence of a diverse range of detection approaches, broadly categorized into spatial, temporal, and adversarially robust methods. Spatial-domain methods, such as the XceptionNet-based detector [40] in the FaceForensics++ benchmark, leverage convolutional architectures to capture subtle pixel-level inconsistencies introduced during generation. Transformer-based models, including Vision Transformer (ViT) variants such as the Patch Transformer [15, 46] for Deepfake Detection, extend this capability by modeling global dependencies to identify manipulation artifacts. Temporal-domain approaches exploit frame-to-frame inconsistencies, with systems like Lip Forensics Detector (LFD) [16] analyzing lip-sync discrepancies, while multi-attentional video transformers detect temporal flicker and jitter patterns using cross-frame attention.

Despite the promising performance of deepfake detectors under normal conditions, a new and formidable challenge has emerged in the form of adversarial attacks [13]. In the context of deepfake detection, these attacks are designed to cause a detector, which would normally classify a synthetic image as "fake," to misclassify it as "real". This poses a significant threat to the reliability and effectiveness of deepfake detection systems, as it fundamentally undermines their ability to distinguish between genuine and manipulated content. The nature of these attacks represents a critical evolution in the threat landscape. Adversarial attacks do not rely on high-level, human-perceptible artifacts. Instead, they exploit the low-level vulnerabilities of the underlying neural network architecture itself [21, 31]. Recent work [5] has further shown that even pretrained face verification models, despite their strong identity representation capabilities, face challenges in reliably distinguishing genuine from manipulated identities. This signifies a far more sophisticated and difficult-to-counter threat.

This research is motivated by the critical need to understand the vulnerabilities of deepfake detection systems to adversarial attacks. Proactively developing and analyzing these attacks is a defensive strategy known as "red teaming" [12]. By intentionally exposing the weaknesses of existing detection models, researchers can devise more robust and resilient defences. The ultimate goal is to strengthen the deepfake detection algorithms' resistance to adversarial attacks and ensure the continued efficacy of media authentication systems. The study of adversarial attacks has a direct, practical application in developing countermeasures. Defense mechanisms such as adversarial training [12], which involves exposing a model to adversarial examples during its training phase, have been shown to significantly enhance a model's robustness. This work aims to advance this field by creating a more realistic and powerful form of attack, thereby pushing the boundaries of what constitutes a robust defense. A significant portion of existing research on adversarial attacks against deepfake detectors has focused on the unrealistic "white-box" threat model [23, 24, 28], where the attacker has complete knowledge of the detector's architecture and parameters. While effective at demonstrating vulnerability, this approach does not represent the real-world scenario of a malicious actor facing a deployed "black-box" model. This creates a critical gap in the literature, which this research is designed to address. The focus on black-box attacks is thus a necessary step to produce solutions that are not just academically interesting but also practically relevant to securing AI systems in the real world.

1.2 Problem Statement

The primary objective is to devise and evaluate a novel framework for generating highly transferable, imperceptible adversarial attacks on deepfake image detection algorithms, specifically targeting the more realistic black-box threat model. This study aims to create attacks that can successfully evade a wide range of unknown, state-of-the-art detectors without requiring access to their internal architecture or parameters. Our methodology prioritizes maximizing the attack’s transferability, thereby providing a robust method for stress-testing the security and reliability of deepfake detection systems in real-world scenarios.

We propose MS-GAGA 3, a dual-stream adversarial framework for attacking black-box deepfake detectors. Our key contributions are two-fold:

1. We introduce a novel dual-stream attack strategy that fuses diverse perturbation mechanisms, one driven by adaptive *MNTD-PGD* 3.1.1 optimization and the other guided by visual saliency *SG-PGD* 3.1.2, to expand the adversarial search space and improve robustness across models and content types.
2. We develop a metric-aware selection module that jointly considers black-box misclassification success (as a proxy for transferability) and perceptual similarity (via SSIM [47]), enabling dynamic, high-fidelity adversarial image selection tailored for varying input conditions.

2 Related Work

2.1 Deepfake Detection Systems

Early advances in deepfake detection relied heavily on convolutional neural networks (CNNs), exemplified by the FaceForensics++ benchmark and its Xception-based classifier [40]. Such models exploit local texture artifacts and blending inconsistencies but tend to overfit to dataset-specific characteristics, leading to sharp performance drops in cross-dataset settings.

Frequency-domain methods, such as F3-Net [30], have sought to capture GAN-specific spectral patterns, offering improved robustness over RGB-only models. However, these approaches can be undermined by compression, post-processing, or novel generative architectures like diffusion models.

Region-specific detection approaches have been explored by [8]. For instance, transfer learning has been leveraged to adapt detectors toward specific facial regions, improving sensitivity to localized manipulations. [13] empirically demonstrated that emotional expressions significantly impact facial recognition performance across diverse datasets, highlighting the importance of human-centric variability when defending against deepfake manipulations. Physiological prior-based detectors, such as DeepRhythm [35], leverage subtle cues like remote photoplethysmography (rPPG) to detect manipulation. While effective in high-resolution, controlled and region-specific settings, they degrade under occlusion, variable lighting, or low resolution, and their applicability to still images remains limited.

Transformer-based and foundation model approaches (e.g., Swin Transformer hybrids [30] or CLIP-adapted detectors [38]) have demonstrated stronger generalization to in-the-wild forgeries by modeling global context and high-level semantics. Nonetheless, these architectures are data-hungry, computationally expensive, and still vulnerable to unseen generation pipelines.

Diffusion-aware methods, such as DiffusionFake [6], reconstruct the input image using a frozen generative prior and detect inconsistencies between the original and reconstruction. While they excel at detecting manipulations from similar generative families, their reliance on prior-specific features limits performance against entirely novel architectures.

Finally, forensic approaches, including EXIF metadata consistency [22] and sensor pattern noise analysis (Noiseprint/PRNU) [8], provide interpretable detection cues but are brittle to common online transformations such as resizing, compression, and content re-hosting.

2.2 Adversarial Attacks on Deepfake Detectors

Taxonomy. Adversarial attacks are commonly classified by the attacker’s level of knowledge about the target model:

- **White-Box Attacks:** The attacker has full access to the target model’s architecture, parameters, and gradients, enabling the creation of highly effective perturbations.
- **Black-Box Attacks:** The attacker has no internal knowledge of the model and can only query it for predictions, often exploiting transferability from a surrogate model.
- **Gray-Box Attacks:** The attacker has partial knowledge, such as the architecture but not the parameters.

Formally, the attacker seeks a perturbed input x^{adv} such that a classifier f misclassifies it, with $\delta = (x^{adv} - x)$ remaining imperceptible:

$$\arg \min_{x^{adv}} \mathcal{L}[f(x^{adv}), y] \quad \text{s.t.} \quad \|x^{adv} - x\|_p \leq \varepsilon \quad (1)$$

where ε controls perceptibility under an L_p norm.

Recent Advances and Limitations. Despite improvements in accuracy, deepfake detection systems remain susceptible to adversarial perturbations. Carlini and Farid [4] showed that imperceptible pixel-level noise can cause significant misclassification rates in state-of-the-art detectors, exploiting their over-reliance on high-frequency cues—making frequency-domain methods particularly vulnerable. Recent work such as AdvForensics [31] has proposed transferable perturbations across multiple detectors, raising concerns about model-agnostic vulnerabilities. Adaptive attacks that jointly optimize against multiple architectures [22] can bypass ensemble-based systems, once believed to be more robust. Beyond pixel-level perturbations, semantic adversarial attacks alter high-level image attributes such as lighting, pose, or background while maintaining visual plausibility [5]. These exploit distributional shifts rather than low-level noise vulnerabilities, impacting both CNN- and transformer-based detectors.

Nonetheless, adversarial methods face constraints. Many require white-box access, which is unrealistic in deployed systems. Transfer-based black-box attacks often degrade under compression or social media re-encoding, which remove or distort perturbations [10]. Moreover, adversarial training can improve robustness but frequently reduces clean-sample accuracy, making it a trade-off in practical detection systems [29]. These concerns extend beyond academic benchmarks. Face recognition and generative AI pipelines are already being embedded in real-world applications such as personalized guest services in the hospitality sector [10], making adversarial robustness a pressing requirement.

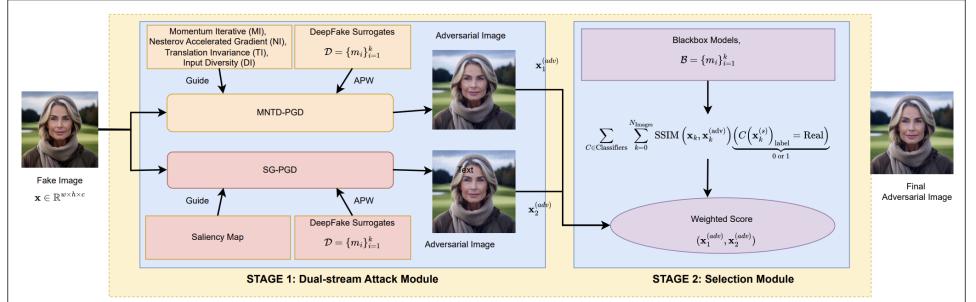


Figure 1: MS-GAGA framework with dual-stream attack module and a selection module.

3 Proposed Methodology

As shown in Figure 1, MS-GAGA is a two-stage framework for generating transferable and imperceptible adversarial examples against black-box deepfake detectors. Stage 1 creates adversarial images with MNTD-PGD and SG-PGD, while Stage 2 selects the optimal image based on black-box misclassification and SSIM.

3.1 Stage 1: Dual-stream Attack Module

To enhance both adversarial transferability and perceptual imperceptibility, MS-GAGA introduces a dual-stream adversarial generation module. Given an input AI-generated fake image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$, the objective is to craft an adversarial counterpart \mathbf{x}^{adv} that maintains high visual similarity to \mathbf{x} while successfully misleading multiple black-box deepfake detectors. This stage generates two adversarial candidates by leveraging complementary strategies aimed at (1) maximizing transferability across unseen models and (2) preserving the structural integrity of the original image.

3.1.1 Momentum-Nesterov-Translation-Diversity PGD (MNTD-PGD)

The first stream extends the standard PGD framework with a suite of robust optimization strategies: (1) Momentum Iterative (M) [8] accumulates past gradients to stabilize updates and escape local minima, (2) Nesterov Accelerated Gradient (N) [29] with a look-ahead step to refine gradient directionality, (3) Translation Invariance (T) [10] applies Gaussian smoothing to improve robustness to spatial shifts, and (4) Input Diversity (D) [28] randomizes input size and padding at each iteration to prevent overfitting to specific surrogate models and boosting attack generalization across architectures. Given an AI-generated fake image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ and a deepfake detecting surrogate ensemble $\mathcal{D} = \{m_i\}_{i=1}^k$, the adversarial example \mathbf{x}^{adv} is generated through iterative updates as follows:

$$\mathbf{x}_{t+1}^{\text{adv}} = \mathbf{x}_t^{\text{adv}} + \alpha \cdot \text{sign}(g_t) \quad (2)$$

$$g_t = \mu \cdot g_{t-1} + \frac{1}{k} \sum_{i=1}^k \nabla_{\mathbf{x}} \mathcal{L}_{\text{total}}(m_i(\mathbf{x}_t^{\text{adv}}), y) \quad (3)$$

where g_t is the momentum-accumulated gradient at iteration t , μ is the momentum decay factor, and α is the step size.

We also incorporate an SSIM-regularized [42] loss $\mathcal{L}_{\text{SSIM}}$ into our MNTD-PGD framework, directly optimizing for high visual similarity between original and adversarial images. Unlike traditional methods that only limit pixel changes, this approach explicitly shapes perturbations to be imperceptible. λ_{SSIM} is a regularization term tuned to 0.3. The total loss is defined as $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{misclassification}} + \mathcal{L}_{\text{SSIM}}$.

$$\mathcal{L}_{\text{SSIM}} = \lambda_{\text{SSIM}} \cdot (1.0 - \text{SSIM}(\mathbf{x}, \mathbf{x}^{\text{adv}})) \quad (4)$$

Additionally, we employ several novel optimizations:

- **Adaptive Per-Surrogate Weighting (APW):** Instead of treating all surrogate models equally, APW dynamically assigns higher weights to those that are harder to fool or more confident in their predictions. This strategic allocation of gradient contribution focuses optimization on the most challenging surrogates, leading to more effective and transferable adversarial examples. By learning which models provide the most informative gradients, APW introduces a meta-optimization layer that enables more intelligent, resource-aware attack strategies.
- **Epsilon Search:** The standard, fixed-epsilon perturbations are replaced by an intelligent epsilon search to find the minimal perturbation needed for a successful attack. It starts with a coarse grid search, followed by binary search refinement between the last failed and first successful values. This process minimizes perturbations and inherently improves SSIM. In addition, the epsilon grid is scaled with the pixel variance of the image - a higher variance allows for larger epsilons.
- **Preprocessing:** For each image, we adjust contrast, brightness, and add Perlin noise [54] to create more "natural" looking final images. This was key to achieving a misclassification rate of 99% on our target models.

3.1.2 Saliency-Guided PGD (SG-PGD)

To maximise the attack budget (ϵ), we made use of saliency maps [41] to identify and target the most influential pixels in an image. Given an AI-generated fake image $\mathbf{x} \in \mathbf{R}^{H \times W \times C}$ and a surrogate model $m \in \mathcal{D}$, we first compute a saliency map by taking the gradient of the loss function \mathcal{L} with respect to the input image, $M_{\text{saliency}} = |\nabla_{\mathbf{x}} \mathcal{L}(m(\mathbf{x}), y)|$. This mask isolates the most important regions, and the adversarial perturbation is subsequently multiplied by it. The update rule for creating the adversarial image, x^{adv} at each step is modified to:

$$\mathbf{x}_{t+1}^{\text{adv}} = \mathbf{x}_t^{\text{adv}} + \alpha \cdot \text{sign}(\nabla_{\mathbf{x}} \mathcal{L}(m(\mathbf{x}_t^{\text{adv}}), y)) \odot M_{\text{saliency}} \quad (5)$$

Here, α is the step size and \odot represents element-wise multiplication.

3.2 Stage 2: Selection Module

While Stage 1 generates two adversarial candidates ($x_{\text{MNTD-PGD}}^{\text{adv}}$ and $x_{\text{SG-PGD}}^{\text{adv}}$), not all perturbations are equally effective or perceptually similar. Simply choosing one arbitrarily risks sacrificing either transferability or imperceptibility. To address this, we design a metric-aware *Selection Module* that evaluates each candidate according to both (i) its success in fooling black-box detectors and (ii) its structural similarity to the original image.

Formally, given a set of classifiers $C \in \mathcal{C}$ and an input batch of N_{images} samples, we compute a joint score that multiplies attack success with perceptual quality. For each adversarial image x_k^{adv} , the Structural Similarity Index (SSIM) between the original x_k and x_k^{adv} serves

as a proxy for visual fidelity, while the classification outcome provides a binary measure of attack success. The cumulative score is defined as:

$$score = \sum_{C \in \text{Classifiers}} \sum_{k=1}^{N_{\text{images}}} SSIM(x_k, x_k^{\text{adv}}) \underbrace{\left(C(x_k^{\text{adv}})_{\text{label}} = \text{Real} \right)}_{0 \text{ or } 1} \quad (6)$$

where adversarial examples are rewarded only when they are both *successful* (fool the classifier) and *imperceptible* (high SSIM). Finally, the Selection Module compares the two streams and retains the candidate with the higher score:

$$x^{final} = \arg \max (\text{score}_{MNTD-PGD}, \text{score}_{SG-PGD}). \quad (7)$$

By coupling perceptual similarity with misclassification success, the Selection Module enforces a balance between transferability and visual fidelity.

4 Experiments

Test Set. We use the test set provided by the *AADD Challenge: ACM Multimedia 2025* [2] to validate our attack strategy. The dataset contains 693 high-quality (HQ) and 710 low-quality (LQ) deepfake images generated using the models in Table 1. A typical deepfake image is a human portrait with varied expressions and complex backgrounds. The number of images generated by GAN and Diffusion Models is numerically balanced. The LQ images are created by resizing followed by variable Quality Factor (QF) compression. This combination is designed to simulate social media compression.



Figure 2: **Top:** Low-quality deepfakes; **Bottom:** High-quality deepfakes.

Quality	Deepfake Generators
HQ	Adobe Firefly, DeepAI, Flux 1.1 Pro, HotPotAI, Nvidia SanaPAG, Stable Diffusion 3.5, StyleGAN 2, StyleGAN 3, Tencent Hunyuan
LQ	DeepAI, Flux.1, Freepik, HotPotAI, Nvidia SanaPAG, Stable Diffusion Attend and Excite, StyleGAN, StyleGAN 3, Tencent Hunyuan

Table 1: Deepfake generators used to create the test set.

Target and Surrogate Models. We used two target models, ResNet-50 [17] and DenseNet-121 [20], each configured as a binary classifier for deepfake detection in a black-box setting. In each experiment, one model served as the target while the other acted as a surrogate. To further increase architectural diversity, improve transferability, and reduce overfitting, we incorporated two additional surrogate models, EfficientNet-B0 [45] and Inception-v3 [24], as robust deepfake detectors.

Baselines. We compare MS-GAGA against the Carlini & Wagner L_∞ attack [8] (white-box) and the Square Attack [11] (black-box), both implemented via the Adversarial Robustness Toolbox (ART) [35]. While more advanced black-box attacks such as GRAPHITE [20] exist, they are computationally intensive (require days of runtime compared to the hours needed for the other methods), thereby not offering a practical or fair comparison.

All attack strategies are evaluated using the following metrics:

- **Misclassification Rate:** Number of times an adversarial deepfake image successfully fools the target deepfake detector.
- **SSIM:** Computed using `skimage.metrics.structural_similarity` with default window size=7 on each RGB channel.

5 Results

Original	MS-GAGA	Square Attack	Carlini-Wagner
			
SSIM: 1.0	SSIM: 0.83	SSIM: 0.78	SSIM: 0.86
			
SSIM: 1.0	SSIM: 0.95	SSIM: NA	SSIM: 0.93
			
SSIM: 1.0	SSIM: 0.27	SSIM: NA	SSIM: NA

Table 2: Comparison of generated adversarial deepfake images. An SSIM of “NA” indicates that the perturbed image failed to fool the classifier. In the first and third rows, MS-GAGA images have 0.8x contrast of the original.

Quantitative Analysis. Table 2 and 3 detail our experimental results. MS-GAGA achieves a 27% higher misclassification rate than the other methods in all scenarios. However, it comes at the cost of lower SSIM. To weigh the trade-offs, we calculate the score on the entire test set using (6), which shows that MS-GAGA still achieves better results overall.

Method	Misclassification Rate (%)		Avg. SSIM	Score
	ResNet-50	DenseNet-121		
No Attack	0	0	100	0
Square Attack [10]	56.3	56.6	73.4	1163
Carlini-Wagner Attack [11]	73.7	71.0	91.3	1855
MS-GAGA (Ours)	99.6	98.8	69.0	1921

Table 3: Performance comparison of various attacks

The adversarial images generated by MS-GAGA are transferable, at least between the similar architectures of ResNet-50 and DenseNet-121. In contrast, the Carlini-Wagner attack tends to overfit the surrogate model, thereby lacking transferability. Additionally, the Square attack is innately less powerful than gradient-based methods.

Ablation Study. To better understand the contribution of each stream in MS-GAGA, we conduct an ablation study where MNTD-PGD and SG-PGD are evaluated independently. Table 4 reports their quantitative performance, while Figure 3 provides visual examples of the resulting adversarial perturbations at different SSIM ranges.

Method	Misclassification Rate (%)		Avg. SSIM	Score
	ResNet-50	DenseNet-121		
MNTD-PGD 3.1.1	98.2	97.6	67.5	1854
SG-PGD 3.1.2	43.8	35.1	82.1	909

Table 4: MNTD-PGD provides transferability, SG-PGD aids higher structural similarity.

Figure 3: **Top:** MNTD-PGD ($SSIM \in [0.8 - 0.9]$); **Bottom:** SG-PGD ($SSIM \in [0.9 - 0.1]$)

6 Conclusion

We introduce MS-GAGA, a two-stage adversarial framework designed to generate highly transferable attacks against black-box deepfake detectors. The results highlight significant vulnerabilities in current deepfake detection pipelines, demonstrating the need for more resilient defences. While this study successfully moves beyond single-surrogate and white-box limitations, it is confined to CNN-based models and static image datasets. Future work will extend this framework to additional deepfake models and modalities and will evaluate its effectiveness against varied detector architectures, such as transformer-based and frequency-domain models, to provide a more comprehensive understanding of its generalizability. This research provides a critical step toward safeguarding digital media integrity by offering a rigorous method to evaluate and strengthen detection systems.

Acknowledgment

This research is supported by the Ministry of Education, Singapore, under its Academic Research Tier 1 (Grant number: GMS 956).

References

- [1] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. *CoRR*, abs/1912.00049, 2019. URL <http://arxiv.org/abs/1912.00049>.
- [2] Pietro Bongini, Sara Mandelli, Andrea Montibeller, Mirko Casu, Orazio Pontorno, Claudio Vittorio Ragaglia, Luca Zanchetta, Mattia Aquilina, Taiba Majid Wani, Luca Guarnera, Benedetta Tondi, Giulia Boato, Paolo Bestagini, Irene Amerini, Francesco De Natale, Sebastiano Battiato, and Mauro Barni. Wild: a new in-the-wild image linkage dataset for synthetic image attribution, 2025. URL <https://arxiv.org/abs/2504.19595>.
- [3] Sim Wei Xiang Calvert and Pai Chet Ng. Leveraging transfer learning for region-specific deepfake detection. In *2024 IEEE Region 10 Conference (TENCON)*, pages 310–313. IEEE, 2024.
- [4] Nicholas Carlini and Hany Farid. Evading deepfake-image detectors with white-and black-box attacks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 658–659, 2020.
- [5] Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. *CoRR*, abs/1608.04644, 2016. URL <http://arxiv.org/abs/1608.04644>.
- [6] Shen Chen, Taiping Yao, Hong Liu, Xiaoshuai Sun, Shouhong Ding, Rongrong Ji, et al. Diffusionfake: Enhancing generalization in deepfake detection via guided stable diffusion. *Advances in Neural Information Processing Systems*, 37:101474–101497, 2024.
- [7] Yunzhuo Chen, Nur Al Hasan Haldar, Naveed Akhtar, and Ajmal Mian. Text-image guided diffusion model for generating deepfake celebrity interactions. In *2023 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 348–355. IEEE, 2023.
- [8] Davide Cozzolino and Luisa Verdoliva. Noiseprint: A cnn-based camera model finger-print. *IEEE Transactions on Information Forensics and Security*, 15:144–159, 2019.
- [9] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jian-guo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [10] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

- [11] Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M. Roy. A study of the effect of jpg compression on adversarial images, 2016. URL <https://arxiv.org/abs/1608.00853>.
- [12] Ryan Feng, Neal Mangaokar, Jiefeng Chen, Earlene Fernandes, Somesh Jha, and Atul Prakash. GRAPHITE: Generating Automatic Physical Examples for Machine-Learning Attacks on Computer Vision Systems, 2022. URL <https://arxiv.org/abs/2002.07088>.
- [13] Shi Quan Marcus Foo, Vi Shean Lim, and Pai Chet Ng. Do emotions matter? an empirical exploration of facial recognition across diverse datasets. In *2025 IEEE International Conference on Service Operations and Logistics, and Informatics (SOLI)*. IEEE, 2025.
- [14] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [15] Jiazhi Guan, Hang Zhou, Zhibin Hong, Errui Ding, Jingdong Wang, Chengbin Quan, and Youjian Zhao. Delving into sequential patches for deepfake detection. *Advances in Neural Information Processing Systems*, 35:4517–4530, 2022.
- [16] Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don't lie: A generalisable and robust approach to face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5039–5049, 2021.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [18] Munir Bin Rudy Herman, Pai Chet Ng, Malcolm Yoke Hean Low, and Detlev Remy. Integrating rag with face recognition for personalized guest services for hospitality industries. In *2024 IEEE Region 10 Conference (TENCON)*, pages 1394–1397. IEEE, 2024.
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [20] Yang Hou, Qing Guo, Yihao Huang, Xiaofei Xie, Lei Ma, and Jianjun Zhao. Evading deepfake detectors via adversarial statistical consistency. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12271–12280, 2023.
- [21] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [22] Minyoung Huh, Andrew Liu, Andrew Owens, and Alexei A Efros. Fighting fake news: Image splice detection via learned self-consistency. In *Proceedings of the European conference on computer vision (ECCV)*, pages 101–117, 2018.

[23] Shehzeen Hussain, Paarth Neekhara, Malhar Jere, Farinaz Koushanfar, and Julian McAuley. Adversarial deepfakes: Evaluating vulnerability of deepfake detectors to adversarial examples. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3348–3357, 2021.

[24] Shuai Jia, Chao Ma, Taiping Yao, Bangjie Yin, Shouhong Ding, and Xiaokang Yang. Exploring frequency adversarial attacks for face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4103–4112, 2022.

[25] Aashish Chandra K, Aashutosh A V, Srijan Das, and Abhijit Das. Latent flow diffusion for deepfake video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3781–3790, June 2024.

[26] Prakash L. Kharvi. Understanding the Impact of AI-Generated Deepfakes on Public Opinion, Political Discourse, and Personal Security in Social Media . *IEEE Security & Privacy*, 22(04):115–122, July 2024. ISSN 1558-4046. doi: 10.1109/MSEC.2024.3405963. URL <https://doi.ieee.org/10.1109/MSEC.2024.3405963>.

[27] Yanli Li, Dong Yuan, Abubakar Sadiq Sani, and Wei Bao. Enhancing federated learning robustness in adversarial environment through clustering non-iid features. *Computers Security*, 132:103319, 2023. ISSN 0167-4048. doi: <https://doi.org/10.1016/j.cose.2023.103319>. URL <https://www.sciencedirect.com/science/article/pii/S0167404823002298>.

[28] Quanyu Liao, Yuezun Li, Xin Wang, Bin Kong, Bin Zhu, Siwei Lyu, Youbing Yin, Qi Song, and Xi Wu. Imperceptible adversarial examples for fake image detection. *arXiv preprint arXiv:2106.01615*, 2021.

[29] Jiadong Lin, Chuangbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. *arXiv preprint arXiv:1908.06281*, 2019.

[30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.

[31] Paarth Neekhara, Brian Dolhansky, Joanna Bitton, and Cristian Canton Ferrer. Adversarial threats to deepfake detection: A practical perspective. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 923–932, 2021.

[32] Pai Chet Ng and Konstantinos N. Plataniotis. Can pretrained face verification models distinguish true identity from deepfakes? In *BMVC Workshops*. BMVA Press, 2024.

[33] Maria-Irina Nicolae, Mathieu Sinn, Minh Ngoc Tran, Beat Bueser, Ambrish Rawat, Martin Wistuba, Valentina Zantedeschi, Nathalie Baracaldo, Bryant Chen, Heiko Ludwig, Ian Molloy, and Ben Edwards. Adversarial robustness toolbox v1.2.0. *CoRR*, 1807.01069, 2018. URL <https://arxiv.org/pdf/1807.01069>.

- [34] Ken Perlin. An image synthesizer. *ACM Siggraph Computer Graphics*, 19(3):287–296, 1985.
- [35] Hua Qi, Qing Guo, Felix Juefei-Xu, Xiaofei Xie, Lei Ma, Wei Feng, Yang Liu, and Jianjun Zhao. Deephythm: Exposing deepfakes with attentional visual heartbeat rhythms. In *Proceedings of the 28th ACM international conference on multimedia*, pages 4318–4327, 2020.
- [36] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *European conference on computer vision*, pages 86–103. Springer, 2020.
- [37] Haonan Qiu, Chaowei Xiao, Lei Yang, Xincheng Yan, Honglak Lee, and Bo Li. Semanticadv: Generating adversarial examples via attribute-conditioned image editing. In *ECCV*, 2020.
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn 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.
- [40] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1–11, 2019.
- [41] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [42] Anusha Sinha, James Lucassen, Keltin Grimes, Michael Feffer, Mary Soto, Hoda Heidari, and Nathan VanHoudnos. What can generative ai red-teaming learn from cyber red-teaming? *Carnegie Mellon University*, 2025.
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [44] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [45] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.

- [46] Junke Wang, Zuxuan Wu, Wenhao Ouyang, Xintong Han, Jingjing Chen, Yu-Gang Jiang, and Ser-Nam Li. M2tr: Multi-modal multi-scale transformers for deepfake detection. In *Proceedings of the 2022 international conference on multimedia retrieval*, pages 615–623, 2022.
- [47] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [48] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2730–2739, 2019.
- [49] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pages 7472–7482. PMLR, 2019.