



FORENSIC TRACES IN GAN-GENERATED FACE IMAGES: FORMATION MECHANISMS, EVOLUTION, AND FORENSIC INTERPRETATION

Yi Wang

Investigation and Forensic Technology Research Center of Shanxi Police College

Article DOI: <https://doi.org/10.36713/epra31768>

DOI No: 10.36713/epra31768

ABSTRACT

Generative adversarial networks (GANs) can synthesize highly realistic fictional faces, creating new challenges for digital image authentication and source attribution. Forensic examination requires more than a binary real-or-fake decision: relevant traces should be reproducible, mechanistically interpretable, and evaluated under clearly defined conditions. This structured narrative review updates the literature through October 2026 and groups forensic traces in GAN-generated face images into three mechanism-based categories: differences between physical camera imaging and neural synthesis pipelines; traces introduced by generator architecture; and anomalies in physiological, geometric, and physical consistency. The review focuses on color and noise statistics, texture, upsampling and spatial-frequency structure, generator fingerprints, StyleGAN-specific artifacts, and ocular geometry, with particular attention to whether these traces persist as generator architectures evolve. Evidence shows that StyleGAN's blob-like and progressive-growing-related phase artifacts were specifically addressed in StyleGAN2, while texture sticking was substantially reduced by the alias-free design of StyleGAN3. High-frequency spectral-falloff discrepancies can also disappear after changes to the upsampling architecture and therefore are not intrinsic properties of CNN-based generators. Frequency-domain, autocorrelation, and ocular geometric differences continue to be reported in newer models, but their expression depends on the generator, training data, image quality, and post-processing history. The forensic significance of these findings is discussed in terms of trace persistence, alternative explanations, processing history, and independent corroboration. No single anomaly should be used in isolation to support a conclusion of AI generation or a specific source.

KEYWORDS: Generative Adversarial Networks; Synthetic Faces; Forensic Traces; Forensic Image Analysis; Generator Fingerprints; Digital Image Forensics

1. INTRODUCTION

Deep generative models can now synthesize highly realistic human faces. Nightingale and Farid found that participants distinguished StyleGAN2-generated faces from real faces at approximately chance level, indicating that obvious visual defects are no longer a reliable basis for human judgment [1]. AI-generated faces encountered online may also undergo resizing, compression, and platform-specific processing, further altering the traces present in the original image [2]. Accordingly, forensic examination should focus not only on whether an image appears AI-generated, but also on identifying reproducible traces and defining their mechanisms, conditions of applicability, and evidential limits.

Automated detectors can achieve high classification performance, but a classifier score is not equivalent to interpretable forensic evidence. An INTERPOL review noted that the transfer of AI-generated media detection into forensic practice is still limited by generalization, robustness, interpretability, and the lack of appropriate evidential evaluation frameworks [3]. Existing reviews of GAN-generated faces are commonly organized around detection methods or feature types [4], whereas broader surveys of AI-generated visual media also cover disruption and authentication strategies [5]. This review therefore centers on the traces themselves: how they arise, whether they persist as generators evolve, which alternative explanations must be considered, and how far they can support forensic inference.

GAN architectures have evolved from DCGAN and PGGAN to the StyleGAN family, with substantial changes in image resolution, latent control, upsampling, and signal processing [6-11]. Latent diffusion models follow a different denoising and decoding pathway [12]. Such architectural changes reshape observable forensic traces, so findings from early generators cannot be assumed to apply to later models. This review focuses on GAN-generated faces and refers to newer generative architectures only when they help assess whether a reported trace remains generalizable.

1.1 Scope, Review Method, and Terminology

This study adopts a structured narrative review rather than a meta-analysis. The literature search covered publications available through October 2026 and included PubMed, CVF Open Access, the ACM Digital Library, NeurIPS Proceedings, publisher databases, and backward citation tracking. Search terms included combinations of GAN-generated face/image, synthetic face, forensic artifact/trace, GAN fingerprint, frequency artifact, ocular artifact, color/noise statistics, StyleGAN artifact, and forensic interpretation.

Studies were included when they met at least one of the following criteria: they reported interpretable anomalies in fully synthetic faces or GAN-generated images; examined the mechanism responsible for a trace; assessed whether a trace persisted, weakened, or disappeared in later generators; or evaluated the effects of compression, resampling, online



dissemination, or camera-source analysis on trace interpretation. Studies reporting only classification performance without analyzing interpretable traces were not emphasized. Work centered on temporal video cues, audio, or local face-swap boundaries was generally outside the scope of this review.

Three terms are distinguished throughout this review. An artifact is a deviation from the structural or statistical behavior expected in reference natural images that is introduced by

image generation or processing. A forensic trace is an observable or measurable difference that can inform source interpretation without implying a unique origin. A generator fingerprint is a statistical pattern associated with a particular generator architecture, training instance, or source [13,14]; it is therefore conceptually distinct from a device-specific camera PRNU fingerprint. The review also distinguishes authenticity assessment, generator-source attribution, and forensic interpretation, with primary emphasis on the last of these.

Table 1 Scope of this review relative to representative prior surveys

Reference	Primary focus	Organizing framework	How this review differs
van Lierop et al., 2026 [3]	Forensic detection of AI-generated images and videos	Generator architectures, detection methods, generalization, and robustness	Emphasizes forensic deployment; this review focuses more narrowly on trace persistence, alternative explanations, and evidential limits.
Wang et al., 2023 [4]	Detection of GAN-generated faces	Deep-learning, physical, physiological, and perceptual cues	Primarily detector-oriented; less emphasis on how individual traces evolve or should be interpreted in forensic casework.
Deng et al., 2025 [5]	Defenses against AI-generated visual media	Detection, disruption, authentication, and trustworthiness	Broader in scope and not organized around the mechanisms of face-specific forensic traces.
This review	GAN-generated faces, with newer generators used when needed for comparison	Formation mechanism; observable trace; persistence; alternative explanations; evidential interpretation	Uses a trace-centered forensic perspective focused on interpretability and evidential limits.

Rather than assigning numerical evidence grades, this review uses four qualitative descriptors based on the scope of published support: supported across multiple studies, condition-dependent support, model-specific support, and source-attribution cue. "Supported across multiple studies" indicates repeated observation in independent studies or across several architectures; it does not imply that the trace is universal. "Model-specific support" indicates that the trace is linked to a particular architecture or version and may weaken or disappear after model updates.

2. FORENSIC TRACES AND THEIR FORMATION MECHANISMS

GAN-generated face traces can be grouped into three mechanism-based categories: color, texture, and noise differences arising from the contrast between physical imaging and neural synthesis pipelines; generator-structure traces introduced by upsampling, convolution, normalization, or training strategy; and local or cross-region inconsistencies associated with incomplete modeling of physiological,

2.1 Imaging-Pipeline Differences

A camera image is produced through optics, sensor sampling, a color filter array, demosaicing, image signal processing, and compression. A GAN image is instead synthesized from latent variables by a neural network and mapped directly to an RGB representation. Because these signal-formation pipelines differ, the resulting images may differ in color-channel relationships, noise residuals, local pixel dependencies, and texture statistics [13,15,16]. However, bypassing a physical sensor does not imply that synthetic images must lack camera-like statistics: a generator may learn some acquisition characteristics from its training data, while screenshots or re-photography can introduce new device-related traces.




Fig. 1 Forensic differences between physical camera imaging and GAN synthesis pipelines

2.1.1 Color and Cross-Channel Statistics

Li et al. reported differences between generated and real images in the local statistics of chromatic components in HSV and YCbCr representations and in corresponding high-pass residuals [15]. Mo et al. further showed that different color-space channels can provide complementary information [17]. Tang et al. reported differences in RGB cross-channel wavelet-subband relationships in images generated by models including StyleGAN2 [18], and more recent work in online social-network settings identified exploitable differences in luminance-chrominance gradient statistics [19]. The literature therefore supports model-dependent differences in color and cross-channel statistics, rather than a fixed or universal GAN color bias.

2.1.2 Texture Statistics

Texture differences can be characterized through co-occurrence relationships, local residuals, or global texture statistics. Nataraj et al. used co-occurrence matrices to analyze local pixel dependencies in GAN-generated images [20]. Liu et al. reported global texture differences across several GAN-generated face datasets and found that some of these statistics remained informative after moderate resizing, JPEG compression, blurring, and added noise [21]. Nevertheless, human-observer experiments show that StyleGAN2 faces are difficult to distinguish reliably from real faces [1]. Visual impressions such as unusually smooth skin or abnormal hair should therefore not be treated as fixed forensic features.

2.1.3 Noise Residuals

Marra et al. adapted ideas from camera PRNU analysis and estimated artificial GAN fingerprints from denoising residuals, observing different correlation patterns across generators [13]. These residual patterns do not share the same physical origin as sensor PRNU. More recent work has shown that AI-generated images and modern smartphone or software processing can also produce false positives in camera-source analysis [22]. Noise-based evidence should therefore be interpreted in the context of the acquisition device, image resolution, and processing history.

2.2 Generator-Structure Traces

Generator-structure traces are directly shaped by upsampling, convolution, normalization, and training strategy. Their mechanisms are often relatively clear, but they are also particularly sensitive to architectural change. They may appear

as visible artifacts or as statistical patterns that are detectable only in the spatial or frequency domain.

2.2.1 Upsampling and Spatial-Frequency Traces

Many CNN-based generators progressively upsample low-resolution feature maps. Odena et al. showed that transposed convolution can produce checkerboard artifacts when kernel size and stride create uneven overlap, and that resize-convolution can reduce this effect [23]. Zhang et al. linked periodic responses in some GAN spectra to the upsampling process [24], while Frank et al. showed that DCT and other frequency-domain representations can reveal differences between some generated and real images [25].

Frequency-domain differences are not limited to regular grid patterns. Durall et al. and Dzanic et al. reported systematic differences in spectral distributions and high-frequency spectral falloff between CNN-generated and real images [26,27]. However, Chandrasegaran et al. produced counterexamples with more natural spectral falloff by modifying the final upsampling stage, showing that this discrepancy is not an intrinsic property of CNN-based generation [28]. Corvi et al. nevertheless observed frequency-domain and autocorrelation anomalies across GANs, VQ-GANs, and diffusion models, while also showing that their magnitude depends on the generator and training data [16].

Beyond amplitude spectra, phase, wavelet, and local pixel relationships have also been used to characterize generative traces. SPSL exploits phase-spectrum differences in face-forgery analysis [29]. Khayatkhoei and Elgammal attributed part of the fine-detail discrepancy in convolutional GANs to a bias against high spatial frequencies [30]. Tan et al. shifted the focus from fixed spectral peaks to neighboring-pixel relationships introduced by upsampling and reported generalizable local dependencies across several GANs and some diffusion models [31]. DCT-coefficient statistics and wavelet-packet representations provide additional, complementary descriptions [32-34].

2.2.2 Generator-Specific Fingerprints

Generator fingerprints are source-specific rather than universal GAN anomalies. Marra et al. and Yu et al. showed that different architectures, and even different training instances, can leave distinguishable statistical patterns [13,14]. Their forensic value is therefore closer to candidate-source attribution than to general authenticity assessment. Stronger interpretation requires comparable generators, training instances, and processing conditions, and a fingerprint observed



for one model should not be generalized to unknown generators.

2.2.3 StyleGAN-Specific Artifacts and Aliasing

Karras et al. identified blob-like artifacts in the original StyleGAN and linked them to the normalization design. The same study also described phase artifacts associated with progressive growing. StyleGAN2 specifically addressed these issues by redesigning the normalization scheme and removing progressive growing [10]. These artifacts should therefore be interpreted as version-specific structural traces rather than as generic GAN signatures.

StyleGAN2 still exhibits texture sticking, in which fine details depend too strongly on absolute image coordinates instead of moving consistently with the represented surface. StyleGAN3 attributed this behavior primarily to aliasing and substantially reduced it through an alias-free design with improved subpixel translation and rotation equivariance [11]. This example illustrates how architectural improvements can directly remove traces that were previously useful for discrimination.

2.3 Physiological, Geometric, and Physical Consistency Anomalies

Human faces are subject to relatively stable physiological structure and physical imaging constraints. Early studies examined visible inconsistencies involving the eyes, teeth, and facial contours [35], while related work linked some structural errors to incomplete global modeling by the generator. These traces are intuitive and easy to localize, but their occurrence depends strongly on pose, illumination, image quality, and the specific generator.

Hu et al. analyzed GAN-generated faces using the consistency of corneal specular highlights, based on the expectation that the two eyes usually receive similar illumination and view similar reflecting environments under comparable viewing geometry [36]. Guo et al. reported irregular pupil shapes in comparisons between FFHQ and StyleGAN2 images [37], and Yang et al. found statistical differences in facial landmark configurations in PGGAN-generated faces [38]. More recent work has also reported measurable ocular geometric differences in StyleGAN2-, StyleGAN3-, and diffusion-generated faces, although discriminative performance decreases substantially in heterogeneous, uncontrolled image sets [39].

Table 2 Literature support and evidential status of physiological and geometric traces

Trace	Evaluated data/models	Reported observation	Main limitations	Evidential status
Inconsistent corneal specular highlights [36]	GAN-generated faces	Interocular highlight relationships may differ from those in real faces.	Requires comparable illumination across the two eyes; side pose, occlusion, and asymmetric reflections can violate this assumption.	Condition-dependent support
Irregular pupil shape [37]	StyleGAN2 / FFHQ	Some synthetic pupils deviate from regular pupil morphology.	Pathology, eyelid occlusion, low resolution, and compression can produce similar appearances.	Model- and dataset-specific support
Facial landmark configuration [38]	PGGAN	Landmark distributions differ statistically from those of real faces.	Pose, alignment procedure, and dataset composition can shift landmark distributions.	Model-specific support
Ocular geometric differences [39]	StyleGAN2/3 and diffusion-generated faces	Measurable ocular geometric differences have been reported in newer generators.	Discriminative performance decreases in heterogeneous, uncontrolled data.	Condition-dependent support

3. RELIABILITY AND FORENSIC INTERPRETATION OF SYNTHETIC-FACE TRACES

Section 2 describes how individual traces arise, but observability does not imply stable evidential value. Forensic interpretation must also consider whether a trace persists after generator updates, whether the questioned image has undergone

transmission or processing, and whether the same observation has plausible explanations in genuine imaging or editing workflows.

3.1 Trace Persistence Across Model Evolution

Generator evolution does not eliminate all traces simultaneously; instead, it changes which traces are present and

how strongly they are expressed. StyleGAN's blob-like and phase artifacts were specifically addressed in StyleGAN2 [10], and StyleGAN2's texture sticking was substantially reduced by the alias-free design of StyleGAN3 [11]. Differences in high-

frequency spectral falloff can also change markedly after modifications to the final upsampling stage [28]. A trace that appears repeatedly in early models should therefore not be assumed to be an intrinsic property of GAN synthesis.




Fig. 2 Evolution of representative generator artifacts

Table 3 Persistence and interpretation of major traces across generator evolution

Trace	Main reported models	Proposed mechanism	Evidence from later models/studies	Current interpretation
Checkerboard / periodic upsampling artifacts [23,24]	Early CNN-GANs	Uneven overlap in transposed convolution and repeated upsampling	Modified upsampling strategies can reduce visible artifacts [23,28].	Architecture-specific; absence does not exclude GAN origin.
Blob-like artifacts [10]	StyleGAN	Normalization-related design issue	Specifically addressed in StyleGAN2 through redesigned normalization.	Version-specific trace.
Phase artifacts [10]	StyleGAN	Spatial position biases associated with progressive growing	Reduced in StyleGAN2 after progressive growing was removed.	Version-specific trace.
Texture sticking [11]	StyleGAN2	Aliasing causes fine detail to depend on absolute image coordinates.	Substantially reduced by the alias-free design of StyleGAN3.	Not a reliable trace for StyleGAN3.
High-frequency spectral-falloff discrepancies [26-28]	Multiple CNN-GANs	Final upsampling and convolution affect high-frequency statistics.	Relatively small architectural changes can produce counterexamples.	Condition-dependent; not an intrinsic property.
Mid-/high-frequency and autocorrelation anomalies [16]	GANs, VQ-GANs, and diffusion models	Generation alters spatial dependencies and frequency statistics.	Reported across model families, but strength depends on the generator and training data.	Observed across models but condition-dependent.

Newer architectures are not entirely trace-free. Corvi et al. continued to report frequency-domain and autocorrelation anomalies across several generative architectures [16], while Tan et al. identified common local dependencies in neighboring-pixel relationships across GANs and some diffusion models [31]. The more defensible conclusion is

therefore that forensic traces evolve with the generation mechanism and must be revalidated for the candidate generator and the conditions of the questioned image.

3.2 Effects of Data, Acquisition, and Post-Processing

Forensic traces are influenced by generator architecture, training instance, and training data. Generator fingerprints vary



across model instances [13,14], and biases present in training data may also be transferred to synthetic images [16]. For face datasets such as FFHQ and CelebA-HQ, alignment, cropping, resolution, and JPEG history can themselves alter color, texture, and frequency statistics. Findings obtained from a particular dataset and generator should therefore not be extrapolated directly to an unknown questioned image.

Compression, resizing, cropping, blurring, screenshots, and platform transcoding preferentially alter high-frequency, noise, and local-correlation statistics. Frequency-based traces can also be targeted by adversarial perturbations [40]. Studies of online social-network images further show that platform processing can modify color statistics in synthetic faces [19]. Whether the questioned image is an original file, has circulated through an online platform, or has undergone secondary editing should therefore be incorporated into the interpretation.

Camera-source analysis can also interact with traces introduced by AI generation. Klier and Baier reported elevated false-positive rates when AI-generated images were examined

with noise-based methods associated with modern cameras or software processing [22]. Accordingly, observations such as the absence of typical PRNU or the presence of a particular noise correlation should not be interpreted in isolation; they should be evaluated under matched device, resolution, and processing conditions.

3.3 Evidential Value and Limits of Interpretation

3.3.1 From Observation to Forensic Interpretation

A forensic trace is first an observation; it is not, by itself, a source conclusion. INTERPOL reviews emphasize that forensic use requires method validation, standardization, interpretability, and evaluation frameworks that reflect realistic case conditions [3,41]. The evidential value of a trace is stronger when its proposed mechanism is consistent with the suspected generation process, plausible alternatives have been assessed, the acquisition and processing history of the questioned image has been considered, and independent traces provide corroboration.




Fig. 3 From observed forensic traces to evidential interpretation

3.3.2 Alternative Explanations and Confounding Factors

The same observation may arise from causes other than neural image generation. Forensic analysis should therefore not ask only whether a particular anomaly is present, but whether the generative explanation is more consistent with the evidence

than plausible alternatives arising from genuine imaging, editing, or dissemination. Table 4 summarizes common confounding factors and the corresponding forensic checks for the main trace categories.

Table 4 Alternative explanations and forensic checks for major trace categories

Trace category	Plausible alternative explanations	Forensic checks	Interpretation
Corneal highlights / ocular reflections [36]	Asymmetric illumination, side pose, occlusion, local reflections	Review head pose, the environment visible in both eyes, and source resolution.	Supportive only when the underlying physical assumptions are satisfied.
Pupil / ocular geometry [37,39]	Pathology, squinting, eyelid occlusion, compression, landmark error	Assess image quality and the stability of segmentation or landmark localization.	Condition-dependent supporting evidence.
Frequency / upsampling structure [23,24,26-28]	JPEG quantization, resizing, sharpening, denoising, re-encoding	Review the original file, resolution and encoding history, and compare with matched controls.	Consistency with a generator pathway does not by itself establish synthetic origin.
Noise residuals / generator fingerprints [13,14,22]	Computational photography, denoising, screenshots, re-photography, rescaling	Document device and software processing; compare with candidate-generator reference images.	More suitable for source association; camera-PRNU thresholds should not be applied directly.
Color / cross-channel statistics [15,17,19]	White balance, filters, color management, platform transcoding	Standardize color space and encoding conditions and use matched reference data.	Supporting statistical evidence with limited specificity.
Texture statistics [20,21]	Beautification, skin retouching, blur, denoising, resampling	Review editing history, sharpness, and compression conditions.	Should not be interpreted from visual impression alone.



3.3.3 Questioned-Image Conditions and Wording of Conclusions

Before examination, the analyst should document whether the questioned image is an original file; whether it was downloaded from a social platform or captured by screenshot; its resolution and compression state; whether metadata and color-profile information are preserved; and whether there are signs of re-photography or other editing.

For traces that depend heavily on high-frequency content or noise statistics, the absence of an original file should reduce the interpretive weight assigned to a single low-level statistical observation.

Forensic conclusions can be framed at three levels. The first is descriptive, for example, "statistical features consistent with a particular generation process were observed." The second is source-supporting, for example, "the observed anomaly has been reported and validated for the candidate generator." The third is an overall opinion on authenticity or source. Progression from the first level to the third requires additional independent evidence and explicit validation of applicability. When the candidate generator is unknown, processing history is incomplete, or different traces conflict, the conclusion should preserve that uncertainty rather than rely on a single anomaly.

4. CONCLUSIONS AND FUTURE DIRECTIONS

From a forensic perspective, traces in GAN-generated face images can be organized into three broad groups: differences associated with the imaging pipeline, generator-structure traces, and physiological or geometric consistency anomalies. Their stability is not uniform. Some traces are tied to specific architectures and can be deliberately reduced by later model designs, whereas other statistical differences have been reported across multiple generator families but remain sensitive to training data, image quality, and post-processing.

Forensic examination should therefore not seek a single, permanent "AI-generated signature." A more defensible approach is to select interpretable traces appropriate to the questioned image, verify their proposed mechanisms and scope of validity, evaluate alternative explanations, and seek corroboration from traces with different underlying causes. As new generator families emerge, previously reported traces should be re-examined to determine whether they remain detectable after model updates and online dissemination and whether similar observations occur in genuine images.

Future work should establish standardized validation frameworks for forensic use. Comparison data, candidate generators, compression settings, and resampling conditions should be explicitly documented, and both the reproducibility and the failure conditions of each trace should be reported. Automated detector outputs should ultimately be translated into reviewable forensic observations rather than treated as conclusions in themselves. This trace-centered approach is more consistent with the forensic requirements of interpretability, validation, and reproducibility than an exclusive focus on classification accuracy.

Funding

This work was supported by Scientific and Technological Innovation Programs of Higher Education Institutions in Shanxi (Project No. 2023L410).

Conflict of Interest

The author declares no conflict of interest.

REFERENCES

1. Nightingale S J, Farid H. AI-synthesized faces are indistinguishable from real faces and more trustworthy[J]. *Proceedings of the National Academy of Sciences of the United States of America*, 2022, 119(8): e2120481119. DOI:10.1073/pnas.2120481119.
2. Porcile G J A, Gindi J, Mundra S, Verbus J R, Farid H. Finding AI-Generated Faces in the Wild[C]//*Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 2024: 4297-4305.
3. van Lierop S, Najdenkoska I, Worring M, Geradts Z. Interpol review of detection of AI-generated image and video deepfakes, 2022-2025[J]. *Forensic Science International: Synergy*, 2026, 13: 100704. DOI:10.1016/j.fsisy.2026.100704.
4. Wang X, Guo H, Hu S, Chang M C, Lyu S. GAN-generated faces detection: A survey and new perspectives[C]//*ECAI 2023. Frontiers in Artificial Intelligence and Applications*, 2023, 372: 2533-2542. DOI:10.3233/FAIA230558.
5. Deng J, Lin C, Zhao Z, Liu S, Peng Z, Wang Q, Shen C. A Survey of Defenses Against AI-Generated Visual Media: Detection, Disruption, and Authentication[J]. *ACM Computing Surveys*, 2025, 58(5): Article 123, 1-35. DOI:10.1145/3770916.
6. Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets[C]//*Advances in Neural Information Processing Systems*. 2014, 27: 2672-2680.
7. Radford A, Metz L, Chintala S. Unsupervised representation learning with deep convolutional generative adversarial networks[C]//*International Conference on Learning Representations*. 2016.
8. Karras T, Aila T, Laine S, Lehtinen J. Progressive growing of GANs for improved quality, stability, and variation[C]//*International Conference on Learning Representations*. 2018.
9. Karras T, Laine S, Aila T. A style-based generator architecture for generative adversarial networks[C]//*Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019: 4401-4410.
10. Karras T, Laine S, Aittala M, et al. Analyzing and improving the image quality of StyleGAN[C]//*Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020: 8110-8119. DOI:10.1109/CVPR42600.2020.00813.
11. Karras T, Aittala M, Laine S, et al. Alias-free generative adversarial networks[C]//*Advances in Neural Information Processing Systems*. 2021, 34: 852-863.
12. Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B. High-resolution image synthesis with latent diffusion models[C]//*Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022: 10684-10695.
13. Marra F, Gragnaniello D, Verdoliva L, Poggi G. Do GANs leave artificial fingerprints?[C]//*2019 IEEE Conference on Multimedia Information Processing and Retrieval*. 2019: 506-511. DOI:10.1109/MIPR.2019.00103.



14. Yu N, Davis L S, Fritz M. Attributing fake images to GANs: Learning and analyzing GAN fingerprints[C]//Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019: 7556-7566.
15. Li H, Li B, Tan S, Huang J. Identification of deep network generated images using disparities in color components[J]. *Signal Processing*, 2020, 174: 107616. DOI:10.1016/j.sigpro.2020.107616.
16. Corvi R, Cozzolino D, Poggi G, Nagano K, Verdoliva L. Intriguing properties of synthetic images: From generative adversarial networks to diffusion models[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. 2023: 973-982. DOI:10.1109/CVPRW59228.2023.00104.
17. Mo S, Lu P, Liu X. AI-generated face image identification with different color space channel combinations[J]. *Sensors*, 2022, 22(21): 8228. DOI:10.3390/s22218228.
18. Tang G, Sun L, Mao X, et al. Detection of GAN-synthesized image based on discrete wavelet transform[J]. *Security and Communication Networks*, 2021: 5511435. DOI:10.1155/2021/5511435.
19. Ghosh T, Naskar R. A robust synthetic face detector in OSN context based on Gradient of Color features[J]. *Forensic Science International*, 2026, 380: 112817. DOI:10.1016/j.forsciint.2026.112817.
20. Nataraj L, Mohammed T M, Manjunath B S, et al. Detecting GAN generated fake images using co-occurrence matrices[J]. *Electronic Imaging*, 2019, 31(5): 532-1-532-7. DOI:10.2352/ISSN.2470-1173.2019.5.MWSF-532.
21. Liu Z, Qi X, Torr P H S. Global texture enhancement for fake face detection in the wild[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020: 8060-8069. DOI:10.1109/CVPR42600.2020.00808.
22. Klier S, Baier H. Boon or Bane: Source Camera Identification meets AI-generated images[J]. *Forensic Science International: Digital Investigation*, 2026, 56(Suppl): 302064. DOI:10.1016/j.fsidi.2026.302064.
23. Odena A, Dumoulin V, Olah C. Deconvolution and checkerboard artifacts[J]. *Distill*, 2016, 1(10): e3. DOI:10.23915/distill.00003.
24. Zhang X, Karaman S, Chang S F. Detecting and simulating artifacts in GAN fake images[C]//2019 IEEE International Workshop on Information Forensics and Security. 2019: 1-6. DOI:10.1109/WIFS47025.2019.9035107.
25. Frank J, Eisenhofer T, Schönherr L, et al. Leveraging frequency analysis for deep fake image recognition[C]//Proceedings of the 37th International Conference on Machine Learning. PMLR, 2020, 119: 3247-3258.
26. Durall R, Keuper M, Keuper J. Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020: 7890-7899.
27. Dzanic T, Shah K, Witherden F. Fourier spectrum discrepancies in deep network generated images[C]//Advances in Neural Information Processing Systems. 2020, 33: 3022-3032.
28. Chandrasegaran K, Tran N T, Cheung N M. A closer look at Fourier spectrum discrepancies for CNN-generated images detection[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021: 7200-7209. DOI:10.1109/CVPR46437.2021.00712.
29. Liu H, Li X, Zhou W, et al. Spatial-phase shallow learning: Rethinking face forgery detection in frequency domain[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021: 772-781.
30. Khayatkhoei M, Elgammal A. Spatial frequency bias in convolutional generative adversarial networks[C]//Proceedings of the AAAI Conference on Artificial Intelligence. 2022, 36(7): 7152-7159. DOI:10.1609/aaai.v36i7.20675.
31. Tan C, Liu H, Zhao Y, et al. Rethinking the up-sampling operations in CNN-based generative network for generalizable deepfake detection[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2024: 28130-28139. DOI:10.1109/CVPR52733.2024.02657.
32. Bonettini N, Bestagini P, Milani S, Tubaro S. On the use of Benford's law to detect GAN-generated images[C]//25th International Conference on Pattern Recognition. 2021: 5495-5502.
33. Giudice O, Guarnera L, Battiato S. Fighting deepfakes by detecting GAN DCT anomalies[J]. *Journal of Imaging*, 2021, 7(8): 128.
34. Wolter M, Blanke F, Heese R, Garcke J. Wavelet-packets for deepfake image analysis and detection[J]. *Machine Learning*, 2022, 111: 4295-4327. DOI:10.1007/s10994-022-06225-5.
35. Matern F, Riess C, Stamminger M. Exploiting visual artifacts to expose deepfakes and face manipulations[C]//IEEE Winter Applications of Computer Vision Workshops. 2019: 83-92.
36. Hu S, Li Y, Lyu S. Exposing GAN-generated faces using inconsistent corneal specular highlights[C]//IEEE International Conference on Acoustics, Speech and Signal Processing. 2021: 2500-2504. DOI:10.1109/ICASSP39728.2021.9414582.
37. Guo H, Hu S, Wang X, Chang M C, Lyu S. Eyes tell all: Irregular pupil shapes reveal GAN-generated faces[C]//IEEE International Conference on Acoustics, Speech and Signal Processing. 2022: 2904-2908. DOI:10.1109/ICASSP43922.2022.9746597.
38. Yang X, Li Y, Qi H, Lyu S. Exposing GAN-synthesized faces using landmark locations[C]//Proceedings of the ACM Workshop on Information Hiding and Multimedia Security. 2019: 113-118. DOI:10.1145/3335203.3335724.
39. Kavar K U, Nair S, Kanwal N K, Bose D. Eyes don't lie: Unmasking fake faces with ocular clues[J]. *Forensic Science International*, 2026, 384: 112921. DOI:10.1016/j.forsciint.2026.112921.
40. Jia S, Ma C, Yao T, et al. Exploring frequency adversarial attacks for face forgery detection[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2022: 4103-4112.
41. Geradts Z, van Lierop S, Kombrink M. Interpol review of forensic image and video analysis, 2022-2025[J]. *Forensic Science International: Synergy*, 2026, 13: 100697. DOI:10.1016/j.fsisyn.2026.100697.