Disclaimer: This is not the final version of the article. Changes may occur when the manuscript is published in its final format.

Exploring Ethical Controversies in AI-Generated Art Based on Fundamental Principles

Lin XingJun[✉1*] , **Huang ZiRan**[2]

[1] School of information Technology&Engineering, Guangzhou College of Commerce, Guangzhou, China, 826435053@qq.com/20230608012 2@xs.gcc.edu.cn, 0009-0008-0965-5165

[2] School of Modern Information industry, Guangzhou College of Commerce, Guangzhou, China, redred55@qq.com/202312070313@xs.gcc.edu.cn, 0009-0001-1905-5373

Abstract

This paper will discuss from a technical standpoint the issues that lie at the heart of controversies over AI-generated art. Rather than ethical and legal debates, it proposes that intuitive public criticism has grasped well the way the generative model fundamentally works. Drawing on an analysis of principles guiding diffusion models and large-scale artificial datasets upon which they actually depend (despite popular beliefs), this study suggests that AI Art is in fact a surface-level statistical recombination process—that there is no actual intentionality, conceptual reasoning, or life expe-rience behind it as there is with human artistic creation. What happens mechanistically is that this technology learns very complex patterns and correlations existing within training data; its output thus becomes an odds mixture from such data rather than any original concept. It further proves that concrete technical evidence points to its inherent derivativeness-from unique model "fingerprints," stylistic homogenization, and detectable statistical artifacts to the fact that it still requires so much prompting from humans. Thus, these are not loopholes of limitations but rather built-in features of the present paradigm of data-driven pattern matching. A principle- and technically-grounded understanding of how this technology actually functions would be urgently needed for responsible consideration of its social, cultural, and economic impacts—not its output. This analysis finally validates the core concerns of the art world as a correct intuitive interpretation of the reality of artificial intelligence operation and not an emotional reaction. In the reality of AI operation, creation is reconstructed as a complex form of replication.

Keywords

AI-generated art; anime aesthetics and digital expression; diffusion models; generative adversarial networks; ethical dimensions of AI-generated art

1 Introduction

The rapid dissemination of generative artificial intelligence (AI) has ignited vigorous debate worldwide, with particularly acute expression within the anime, comics, and games community-a creative ecosystem centered on virtual characters, anime aesthetics, and digital expression. At the heart of this debate lies a tension between the creative possibilities offered by technological innovation and the need to protect existing artistic practices, rights, and livelihoods. While academic and public discourse has extensively addressed the ethical, legal, and philosophical dimensions of AI-generated art, one crucial perspective remains underdeveloped: a principled, technically precise interrogation of the mechanisms that underpin generative AI models themselves.

This paper aims to address that gap. Rather than focusing primarily on the consequences of AI art for artists and creative industries, we examine the technological principles that govern how these systems learn, represent, and generate visual content. Our central argument is that the major controversies associated with AI-generated art—concerning copyright, economic displacement, and creative authorship—cannot be adequately resolved without understanding what generative models actually do at a technical level.

The core technical contention we advance is that AI image generation, as implemented in current diffusion-based architectures such as Stable Diffusion and Midjourney, constitutes a form of statistical recombination: models learn high-dimensional probability distributions over billions of image-text training pairs (including copyrighted works) and generate novel pixel arrangements by sampling from these learned distributions. This process, while capable of producing visually impressive results, operates without conceptual understanding, intentional planning, or the experiential grounding that characterizes human artistic creation.

The outputs are probabilistic interpolations within a learned data manifold rather than products of interpretive engagement with subject matter.

We develop this argument through three complementary lines of analysis.

First, we examine the operational principles of diffusion models—from Denoising Diffusion Probabilistic Models to Latent Diffusion Models—to show that the training objective is fundamentally one of statistical distribution matching, measured by metrics such as Fréchet Inception Distance, rather than any measure of conceptual understanding or creative intent (Section 3).

Second, we review the growing body of detection and forensics research demonstrating that AI-generated images bear measurable, systematic artifacts—spectral anomalies, reconstruction errors, and generator-specific fingerprints—that empirically distinguish them from human-created works and reveal the recombinative nature of their construction (Section 4).

Third, we situate these technical findings within the rapidly evolving legal and policy landscape, examining how regulators in the United States, European Union, and United Kingdom are grappling with the implications of generative AI for copyright, transparency, and creative rights (Section 2).

Navigational note on structure. This paper integrates technical, legal, and community evidence into a single mutually reinforcing argument. Section 2 (legal and policy context) is placed before the technical sections to provide the regulatory frame of reference, but readers who prefer to begin with the technical core may start with Sections 3 and 4 and return to Section 2. None of these layers is sufficient on its own: a purely technical paper would lack social relevance, a purely legal paper would lack empirical grounding, and a purely community-focused paper would lack analytical rigor. The paper's contribution lies in showing how these layers connect — for instance, how the USCO's human authorship requirement implicitly rests on the same recognition of statistical recombination that the ACG community expresses through its metaphors. We have added transitional language at the boundaries between sections to help readers follow these connections.

This paper is written as a **perspective paper** grounded in technical analysis, rather than as a systematic review. We draw on purposively selected literature from generative modeling, AI image forensics, and copyright law to construct an evidence-informed argument about the nature and implications of AI-generated art. We do not claim to provide an exhaustive survey of all relevant work; rather, we select key technical contributions that illuminate the operational logic of generative models and the empirical evidence for their derivative character.

Our focus is primarily on diffusion-based image generation models (e.g., Stable Diffusion, Midjourney, DALL-E), as these currently dominate both public discourse and creative applications. We acknowledge that other paradigms—including generative adversarial networks, autoregressive models, and emerging hybrid architectures—exhibit different technical properties, and we reference GAN-based research where it provides instructive comparison. However, the core argument concerning statistical recombination applies across current paradigms, even if the specific artifact signatures vary.

The analysis is situated within the context of ACG community debates, which we treat as a concrete case study illustrating how technically informed critique can ground and enrich policy discussion. We recognize that this focus reflects the authors' cultural and research context and may not capture the full diversity of global perspectives on AI art.

This paper draws on two categories of sources. The first comprises peer-reviewed technical and legal literature from computer vision, AI forensics, and copyright law, which grounds the core analytical claims. The second comprises non-academic cultural and community sources — including popular media, community discourse on Chinese social platforms such as Bilibili, Douyin, and Xiaohongshu, and informal critiques from ACG practitioners — which are cited to illustrate the discourse being analyzed rather than to substantiate technical or legal conclusions. The ACG community is not an academic community and does not express its concerns through peer-reviewed channels; a paper that claimed to represent community perspectives while citing only academic sources would systematically exclude the very voices it seeks to engage. We have made a deliberate effort to supplement non-scholarly sources with corroborating academic references where the same underlying claims intersect with scholarly discourse.

2 Legal and Policy Context

Before examining the technical dimensions of AI-generated art, it is essential to establish the rapidly evolving legal-regulatory environment in which these technologies operate. The past two years have seen major policy developments across multiple jurisdictions, each grappling with the core question at the heart of this paper: what is the legal and creative status of outputs produced by statistical recombination of training data?

2.1 United States: The Human Authorship Doctrine

The United States Copyright Office has taken a clear position: copyright protection attaches only to works of human authorship. In its landmark March 2023 policy statement, the USCO clarified that "copyright can protect only material that is the product of human creativity," and that "the Office will not register works produced by a machine or mere mechanical process that operates randomly or automatically without any creative input or intervention from a human author"[1]. This guidance established that AI-generated content may be registrable only to the extent that a human author has made sufficiently creative contributions—through selection, arrangement, or modification—that are distinguishable from the AI's output.

The Zarya of the Dawn registration decision (February 2023) provided the canonical application of this principle[2]. The USCO determined that the text of a graphic novel, written by a human author, was copyrightable, but that the individual images generated by Midjourney were not, because they lacked sufficient human authorship. The Office reasoned that the prompts provided by the human author were more akin to "instructions to a commissioned artist" than acts of direct creation, and that the AI's role in translating those prompts into images was the determining creative act. This decision has become a touchstone in discussions of AI authorship and continues to shape registration practice.

The USCO subsequently launched a comprehensive AI initiative, publishing a multi-part report series addressing digital replicas (Part 1, July 2024), the copyrightability of AI-generated outputs (Part 2, January 2025), and the use of copyrighted works in AI training (Part 3, May 2025) [1]. The trajectory of these reports suggests a sustained institutional commitment to the principle that copyright's human authorship requirement is not merely a technical formality but a foundational element of the incentive structure that copyright law is designed to create.

2.2 European Union: Transparency and Detectability

The European Union's Artificial Intelligence Act (Regulation 2024/1689), which entered into force in August 2024, establishes a regulatory framework that directly addresses the detectability concerns discussed in Section 4 of this paper[3]. It is important to note, however, that the Act adopts a staggered implementation schedule: the prohibitions on unacceptable-risk AI practices took effect on 2 February 2025; the rules for general-purpose AI models, including key transparency provisions under Article 50 as they apply to GPAI, became applicable on 2 August 2025; and the majority of remaining obligations — including the full scope of Article 50 transparency requirements for non-GPAI systems — will not apply until 2 August 2026, with certain high-risk classification rules under Article 6(1) deferred until 2 August 2027. At the time of writing, therefore, not all transparency obligations are fully operative, and the practical enforcement landscape is still taking shape.

At the time of writing, the Article 50 transparency obligations are therefore in force. The European Commission published its final guidelines on these obligations on 30 July 2026, only two days before the effective date, which the professional literature noted left providers and deployers little time to prepare[4]. The guidelines organize Article 50 into provider duties and deployer duties. On the provider side they address the duty to inform users that they are interacting with an AI system, and the duty to mark machine-generated synthetic content in a machine-readable format, to the extent this is technically feasible and does not undermine the system's function. On the deployer side they address the disclosure duty for deep fakes, the disclosure duty for AI-generated text published to inform the public on matters of public interest (with an exception where the text has undergone human review or editorial control), and the information duty toward persons exposed to emotion recognition systems. A first draft Code of Practice on transparency for AI-generated content had already been published in February 2026, and the final guidelines operate alongside it. Recent professional and media commentary tracks how these duties will be operationalized: Ropes & Gray (August 2026) works through the practical questions the obligations

raise for providers and deployers, including what counts as a machine-readable mark and how responsibility travels through intermediaries[5]; Morgan Lewis (August 2026) maps what went into effect on 2 August, including the transparency rules and the supervisory powers of market surveillance authorities[6]; and the Debevoise Data Blog (3 August 2026) describes this third wave as shifting the Act's center of gravity from rulemaking to transparency and supervision[7].

The Act imposes specific transparency obligations on providers and deployers of generative AI systems, including requirements that:

AI-generated content be marked in a machine-readable format and labeled as artificially generated or manipulated;

Providers of generative AI systems ensure that outputs are detectable as synthetic through technical solutions such as watermarks or metadata embedding;

Users be informed when they are interacting with AI-generated content in contexts where it would not be apparent.

These provisions reflect a regulatory judgment that the technical distinguishability of AI-generated content—the very phenomenon documented by the detection research reviewed in Section 4—is both real and actionable. By mandating disclosure and detectability, the EU AI Act implicitly endorses the proposition that AI outputs are categorically distinguishable from human-created works and that consumers and participants in creative markets have a legitimate interest in knowing which is which.

2.3 United Kingdom: TDM, Rights-Reservation, and the Getty Litigation

The United Kingdom's approach to AI and copyright has undergone significant evolution. In 2022, the UK Intellectual Property Office proposed a broad text and data miningexception that would have permitted commercial AI training on copyrighted works without rights-holder authorization. Following sustained opposition from creative industries, the government withdrew this proposal in early 2023[8]. The subsequent , distinct 2024–2025 consultation has explored a rights-reservation model under which rights holders could opt out of having their works used for AI training, coupled with transparency measures requiring AI developers to disclose training data sources.[9]

The government's report on copyright and artificial intelligence was presented to Parliament in 2026; commentary on it observes that the earlier rights-reservation "opt-out" model has effectively been abandoned and that the government remains non-committal on statutory reform[10,11].

This policy trajectory is directly relevant to the ACG community's concerns about unauthorized style replication. The UK's movement toward a rights-reservation model acknowledges, at a policy level, that the relationship between AI training and copyright is not a settled question and that mere public accessibility of works does not constitute consent for training purposes.

The *Getty Images (US), Inc. v. Stability AI Ltd.* proceedings in the UK High Court are the most visible judicial test of the claims about unauthorized training discussed throughout this paper[12]. The proceedings began in January 2023, and their first judicial stage came on 1 December 2023, when the same judge, Mrs Justice Joanna Smith DBE, handed down an early ruling on preliminary issues31. That ruling settled two matters: it confirmed the English court's jurisdiction over the US defendants through service out of the jurisdiction, and it allowed the principal copyright claims, above all the outputs claim and the secondary infringement claim, to proceed toward a full trial. At the same time the ruling recorded that the defendants' evidence placed the training of the model abroad, so the training claim went forward to trial with its territorial foundation already in question.[13]

On 4 November 2025, Mrs Justice Joanna Smith DBE handed down the High Court's judgment. We note at the outset that the outcome of this litigation does not determine the technical analysis presented here; we review it because its actual outcome is frequently misreported and because it illustrates the legal obstacles facing claims of this kind.

Contrary to some public accounts, the November 2025 judgment did not establish that commercial AI training on copyrighted material constitutes infringement under UK law. Shortly before closing submissions, Getty Images abandoned its principal training-related copyright claim, acknowledging that there was no evidence that the training and development of Stable Diffusion took place within the United Kingdom — a territoriality requirement that could not be met on the evidence available to it.

The Outputs Claim, which alleged that the model's outputs infringed copyright, was likewise abandoned once the relief sought had been substantially achieved, as was the linked database rights claim.

On the trade mark claim, Getty Images succeeded in part, but the court itself described its findings as "both historic and extremely limited in scope": infringement was found only in respect of the generation of Getty and iStock watermarks by the v1.x models accessed through Dream Studio or the Developer Platform and by the v2.x models; there was no evidence of a single UK user generating watermarked images with SD XL or v1.6; the section 10(3) TMA claim was dismissed; and the double-identity claims were largely dismissed. The court declined to address the passing-off allegation. The court also dismissed the secondary infringement claim under sections 22 and 23 of the Copyright, Designs and Patents Act 1988 (CDPA), holding that, although "article" may extend to intangible objects, an AI model such as Stable Diffusion which does not store or reproduce any copyright works, and has never done so, is not an "infringing copy" within section 27(3) CDPA. Because the training claim was abandoned on territorial grounds rather than decided on the merits, the judgment leaves open — rather than resolves — whether

and under what conditions commercial AI training on co pyrighted works constitutes infringement; as the court ob served, Getty Images may be able to maintain such a cas e in the jurisdiction where the model was in fact trained, but there was no basis for that case in this jurisdiction. In 2026 the High Court granted permission to appeal, and t he appeal is pending[14], so the questions the case raises re main live.

The proceedings nonetheless remain instructive for this paper in two respects.

First, the judgment's own technical characterization of the model — that Stable Diffusion transforms a prompt into a synthesised image by modelling a probability distribution based on its training data and then sampling from that distribution, and that the model does not store or reproduce the works on which it was trained — is precisely the account of statistical recombination that Section 3 of this paper develops; the legal significance of that technical fact was left undecided in the United Kingdom, but the technical fact itself was not in dispute.

Second, the proceedings, together with the withdrawal of the UKIPO's broad 2022–2023 TDM proposal and the 2024–2025 rights-reservation consultation[8,9], reflect ongoing regulatory uncertainty about AI training and copyright. For the ACG community's concerns about unauthorized style replication, the significance of the UK litigation is therefore cautionary rather than precedential: the legal status of industrial-scale training on copyrighted works remains unsettled in the United Kingdom, which is precisely why policy frameworks — and the technical evidence this paper reviews — matter.

2.4 The Chinese Regulatory Context

The preceding subsections have surveyed Western legal and policy developments. However, this paper's analysis is situated within mainland China, which operates under a fundamentally different legal and social framework — a difference that carries direct implications for how the ACG community's critiques of AI art should be understood.

On the specific question of AI-generated content, Chinese law currently provides no explicit statutory answer. Neither the Copyright Law of the People's Republic of China[15] (as amended in 2020) nor its judicial interpretations define whether content generated by artificial intelligence qualifies as a "work" within the meaning of Article 3 of the Law, and doctrinal opinion remains divided. Judicial practice has therefore developed case by case. In Feilin v. Baidu[16], the court declined to protect a software-generated text as a work, holding that the generation process did not reflect the human intellectual creation that the originality requirement demands, while leaving open that a human user's creative choices could in principle satisfy that requirement. In Tencent v. Yingxun[17], by contrast, the court upheld copyright protection for an AI-assisted financial news article, emphasizing the substantive contribution of the human team that designed the generation framework, selected the data, and edited the output.

More recent disputes over AI-generated images have continued to show divergent outcomes, and commentators differ on where the line between "AI-assisted" and "AI-generated" creation should be drawn[18]. We distinguish these two categories strictly throughout this section: AI-generated works are outputs produced without sufficient human intellectual creation, while AI-assisted works are outputs in which the human author made a substantive contribution to the generation process. Recent decisions continue to develop this boundary case by case. In Shanghai's first AI large-model copyright case, the second-instance court held in July 2026 that a user's AI "replication" of the Medusa anime character infringed copyright while the platform bore no liability[19]. In April 2026 a district court in Jinan concluded the province's first AI-generated image infringement case[20]. In a parallel line of 2025 to 2026 cases, claims over AI-generated images have been rejected where the plaintiff could not demonstrate intellectual investment in the generation process, under the "process documentation" standard the courts have developed[21]. Chinese courts are thus settling the boundary between AI-generated and AI-assisted works case by case, with human intellectual input as the decisive factor, and the outcomes continue to diverge. The prevailing tendency of Chinese judicial practice is to attribute rights to the human user who made a substantive contribution to the generated content; the legislature, however, has yet to codify this position. This doctrinal uncertainty is itself part of the context in which the ACG community's critique of AI art operates: the legal status of AI-generated images in China remains unresolved, and the community's concerns about unauthorized style replication currently have no settled doctrinal answer to rely on.

The People's Republic of China follows a socialist civil law system with Chinese characteristics. Unlike the United States or the United Kingdom, China is not a common law jurisdiction; judicial precedent carries considerably less weight, and there is no equivalent to the US doctrine of fair use. China's legal framework, as codified in the Civil Code of the People's Republic of China, provides a closed list of statutory limitations and exceptions rather than an open-ended fair use standard — a structural difference from common law jurisdictions that has direct implications for how AI training on copyrighted works would be assessed under Chinese law[22]. The legal authority governing AI-generated content in China derives primarily from statutes enacted by the National People's Congress, judicial interpretations issued by the Supreme People's Court, administrative regulations such as the Interim Measures for the Management of Generative Artificial Intelligence Services[23], and in certain domains, policy directives from the Communist Party of China[24]. In

this system, the Communist Party of China provides overall leadership: it conducts industry investigations and consultations, after which administrative regulations are promulgated by the relevant state bodies. When conditions mature, these administrative instruments may be elevated into formal laws enacted by the National People's Congress.

This process is formally articulated in the 2014 Decision of the CPC Central Committee on Comprehensively Advancing the Rule of Law, which codified the principle that "the Party leads legislation, ensures law enforcement, supports the judiciary, and takes the lead in abiding by the law" , and that Party policies must be translated into state will through statutory procedures[25]. Judicial decisions, by contrast, carry considerably less precedential weight than in common law jurisdictions, because China's civil law tradition does not permit judicial lawmaking: courts are bound by codified statutes and judicial interpretations issued by the Supreme People's Court, and their rulings are correspondingly conservative in character, rarely confronting jurisprudential questions directly.

As Jiang Qinghua has argued, the Chinese Communist Party is a "leading party" fundamentally distinct from the "ruling party" model of Western competitive party systems: the Party's leadership is constitutionally embedded and its exercise of public power is regulated through a specialized body of "Party leadership regulations" within the intra-Party regulatory system — an institutional architecture that has no analogue in either common law or civil law jurisdictions[26]. This multi-layered governance structure — in which administrative regulation, Party policy, and statutory law interact in ways that have no direct Western analogue — produces a regulatory environment that alignment with "socialist core values" rather than the rights-balancing framework characteristic of Western copyright law[27]. The Western legal sources we discuss in this section are therefore presented as illustrations of the range of possible regulatory responses, not as directly applicable precedents for the Chinese context. The legal section serves, in this sense, as a comparative reference frame for the technical analysis, not as the paper's primary contribution.

The role of Party and state documents in this system deserves explicit explanation, because it explains the sources this section cites. Since the 2018 constitutional amendment wrote "the leadership of the Communist Party of China is the most essential feature of socialism with Chinese characteristics" into the Constitution, Party leadership is itself part of the constitutional order, and Party documents are the instrument through which that leadership is exercised. Their role in adjudication is documented quantitatively by Fan Ying[28]: her empirical study examined 1,322 judgments in which party-government joint documents appear, 559 administrative, 708 civil, and 55 criminal, with 75 percent decided between 2018 and 2022, and shows that courts use these documents as the basis of judicial reasoning in their judgments; in the terminology of the Chinese literature these documents are "may-apply sources" and cognitive sources that supplement administrative norms without displacing statutes45. The Party documents cited in this section, including the 18th Party Congress report and the 2014 Decision on Comprehensively Advancing the Rule of Law, occupy exactly this functional position in the legal order.

Regulatory question	United States	European Union	United Kingdom	Mainland China
Authorship of AI-generated output	USCO human-authorship doctrine: registration requires human authorship (Zarya of the Dawn)	No AI-specific authorship rule; originality requires the author's own intellectual creation	CDPA s.9(3): authorship of computer-generated works attributed to the person making the arrangements	No explicit statutory provision; judicial practice diverges (Feilin v. Baidu; Tencent v. Yingxun) and tends to favor the user with substantive contribution[18]
Use of copyrighted works in training	No settled rule; fair-use litigation pending	TDM exception with rights reservation (opt-out) under the DSM Directive	Broad TDM proposal withdrawn (2023); rights-reservation consultation (2024–2025); Getty judgment leaves the training question open	Civil Code: closed list of exceptions, no open-ended fair use; Interim Measures for Generative AI Services (2023)
Transparency and detectability	No general labeling mandate	AI Act Art. 50: watermarking and labeling obligations (phased implementation)	Transparency measures under consultation	Labeling obligations for synthetic content (2023 Interim Measures; Deep Synthesis Provisions)

Table 1 Functional comparison of AI-art governance across the four jurisdictions discussed three

This regional specificity carries implications for how the legal and policy frameworks discussed i should be interpreted in relation to this paper's argument. We do not claim that the Western legal sources reviewed above are directly applicable to the Chinese context; rather, we present them as comparative reference points that illuminate the range of possible regulatory responses to AI-generated art. The diversity of these responses — from the USCO's human authorship requirement to the EU's transparency mandates to China's state-led content governance — itself suggests that no single jurisdiction has arrived at a definitive or universally applicable framework. The appropriate framework for any given society depends not only on technical facts about how generative models operate (which are, in principle, jurisdiction-independent) but also on deeply embedded features of that society's legal tradition, cultural values, and political structure. This paper aims to provide the technical grounding that can inform any such framework, while remaining explicit about the cultural and jurisdictional specificity of its own vantage point.

2.5 Synthesis: Why Technical Facts Matter for Law and Policy

To make the comparison across jurisdictions explicit rather than impressionistic, we examine each of the four legal systems surveyed above against the same three regulatory questions:

(i) how the law treats the authorship of AI-generated output.

(ii) how it treats the use of copyrighted works in AI training.

(iii) how it regulates the transparency and detectability of AI-generated content.

Table 1 summarizes the comparison three observations follow.

First, no jurisdiction has yet produced a settled and comprehensive answer to any of the three questions. Comparative scholarship on AI-generated content reaches the same conclusion: a systematic comparison by Feng (2025) documents that the United Kingdom affirms the work status of computer-generated output on incentive-based grounds, that the United States doubts its protectability on human-centric grounds, and that Chinese legislation remains silent while judicial practice tends to vest rights in the human user who made a substantive contribution to the output[18].

Second, the diversity of responses itself indicates that the appropriate framework for any given society depends not only on technical facts about how generative models operate — facts that are, in principle, jurisdiction-independent — but also on deeply embedded features of each society's legal tradition, cultural values, and political structure.

Third, and central to this paper's argument, each of the three questions turns at least in part on technical facts about how generative models actually work: the empirical distinguishability of AI-generated content (USCO registration practice, EU transparency obligations), the significance of training-data provenance (UK TDM debates, the Getty litigation), and the question of whether statistical recombination constitutes transformative use (US fair-use analysis, UK infringement analysis). The analysis that follows aims to provide precisely this technical grounding.

3 The Operational Logic of Generative AI Models

The debate surrounding AI-generated art hinges on a fundamental question: does artificial intelligence engage in genuine creative acts, akin to human artists who draw upon understanding, intent, and emotion, or is it merely a sophisticated mechanism for statistically recombining elements from its training data? The paper argues that core operational principle of modern generative AI models is the learning of data distributions through probabilistic frameworks, not the emulation of semantic understanding.

3.1 Diffusion Models: The Core Mechanism

To ground the discussion, we first summarize the operational pipeline of a diffusion-based generative model. Figure 1 provides a schematic overview of the forward diffusion process (training) and the reverse denoising process (generation).

Jonathan Ho et al.'s work on DDPMs establishes that image synthesis is achieved by training a model to reverse a fixed noise-adding process[29]. The model learns to gradually denoise a random variable over many steps, effectively approximating the data distribution $p(x)$ through a variational inference objective.

As Ho et al. state, the model is trained as "a hierarchy of denoising autoencoders," where the primary goal is to minimize a loss function based on the difference between predicted and actual noise at each step of the corruption process.[29] This objective is purely statistical; it measures the fit between the generated output and the expected distribution of the training data. The model has no inherent concept of the objects it generates—a "dog" or a "landscape"—but instead learns that certain pixel configurations (or latent representations) have a high probability of occurring given a specific conditional input, such as a text prompt. While this process is highly efficient, it is conceptually somewhat rigid. It is an optimization task focused on matching patterns, rather than interpretive behavior.

3.2 Latent Diffusion Models

Robin Rombach et al.'s Latent Diffusion Models (LDMs) substantially improved the efficiency of diffusion-based

generation by operating in a perceptually compressed latent space rather than pixel space[30]. This innovation—which underpins Stable Diffusion and related systems—enables high-resolution synthesis at reduced computational cost. However, the fundamental paradigm remains unchanged.

As Rombach et al. describe, their method "decomposes the image formation process into sequentially applying a denoising autoencoder to the learned latent representation," where the autoencoder provides a space in which "high-frequency, imperceptible details are abstracted away," allowing the diffusion model to focus on "important semantic information"[30]. However, from a creative perspective, the term "semantic" is potentially misleading. The model continues to learn statistical relationships in the compressed latent space, prioritizing features that the independent perceptual model deems salient for reconstruction rather than synthesizing conceptual meaning. The loss function penalizes deviations from the training distribution, reinforcing the model's character as a probabilistic estimator rather than a creative agent.

The use of "classifier-free guidance"—a technique proposed by Rombach et al. to improve image-text alignment—further illustrates this point[30]. Guidance is a computational method for sharpening conditional probabilities by amplifying the difference between conditional and unconditional predictions. It is an explicit, external steering mechanism that nudges the probability distribution toward desired outcomes—not something that emerges naturally from intrinsic intentionality or understanding.

3.3 Spectral Artifacts and the Fingerprints of Statistical Generation

Empirical studies of AI-generated images provide compelling evidence of their statistically composite nature in the form of consistent, measurable artifacts. Ricard Durall et al. directly addressed this phenomenon in their study of spectral distributions in CNN-based generative models[31]. Their research demonstrated that upsampling methods commonly employed by generators—specifically transposed convolutions in GAN architectures—systematically fail to reproduce the spectral distribution of natural training images.

This deficiency is not a minor implementation detail but a fundamental consequence of the local processing mechanisms inherent to convolutional operations. The model efficiently generates local textures and patterns but struggles with global, long-range consistency because it lacks top-down compositional planning. Spectral distortion—whether manifested as excessive high-frequency noise or insufficient high-frequency detail—represents a systematic error arising from feature recombination based on local statistical correlations rather than unified scene understanding.

Durall et al. further demonstrated that these spectral signatures can serve as highly accurate detectors for synthetic images, underscoring that these are not random or occasional errors but inherent fingerprints of the statistical generation process[31]. This finding has been corroborated and extended by subsequent work. Corvi et al. documented that diffusion models, despite their architectural differences from GANs, produce synthesis-specific spectral patterns and artifacts that are especially pronounced in structurally complex image regions[32]. The Diffusion Reconstruction Error method proposed by Wang et al. leverages reconstruction errors from pre-trained diffusion models to show that even advanced architectures fail to fully recover the high-frequency details and coherent structures characteristic of real photographs[33].

3.4 Self-Attention: Improved Recombination, Not Understanding

The introduction of self-attention mechanisms, notably in Self-Attention GANs by Zhang et al., was motivated by an explicit acknowledgment of the limitations of purely convolutional architectures[34]. Zhang et al. observed that traditional convolutional GANs "have much more difficulty in modeling some image classes than others," particularly those requiring "geometric or structural patterns." They noted that GANs excel at generating textures (e.g., sky, ocean) but struggle with coherent structures (e.g., a dog with clearly articulated limbs).

The self-attention mechanism enables the model to "leverage neighborhoods that correspond to object shapes rather than local regions of fixed shape"[34], allowing features at distant spatial locations to influence each other based on learned affinity. This represents a meaningful architectural improvement, but it is itself a form of learned, weighted recombination of features across the feature map. Attention improves the sophistication of the statistical averaging process but does not constitute conceptual understanding; it enables non-local feature recombination without introducing top-down compositional planning.

Subsequent detection research has confirmed that models incorporating attention mechanisms still produce systematically detectable outputs. Yan et al. (2024), in their AIDE framework combining low-level frequency features with high-level semantic embeddings from CLIP, demonstrated that generative models employing attention "often produce inconsistencies in logical object structures and global scene coherence" that are absent from real photographs[35]. Zhu et al. (2023), in their GenDet framework, emphasized that "generators incorporating self-attention still exhibit shared, detectable artifacts across different architectures" because their recombination mechanisms inherently produce deviations from natural image statistics[36]. Liu et al. (2024a), through their FatFormer model integrating forgery-aware adapters and language-guided alignment, showed that feature recombination in

Family	Core Mechanism	Representative Studies	Key Limitations
Frequency-domain	Exploit systematic failure to reproduce full spectral distribution of natural images; transposed convolutions introduce characteristic spectral anomalies	Durall et al.[31]; Corvi et al. [32]	May be mitigated by architectural improvements that better approximate natural frequency statistics
Reconstruction-error-based	Pre-trained diffusion models reconstruct real images with lower error than synthetic images — asymmetry reveals generation traces	DIRE (Wang et al.) [33]	Computationally expensive; requires access to a pre-trained diffusion model
Feature-space / Hybrid	Combine low-level statistical features with high-level semantic embeddings (e.g., CLIP); detect inconsistencies between local pixel statistics and global coherence	AIDE (Yan et al.) [35]; FatFormer (Liu et al.)[37]	Dependence on pre-trained feature extractors; may inherit biases from those models
Ensemble / Generalization-focused	Address cross-generator generalization gap via expert models specialized for different generators; Gated CNN ensembles	Saskoro et al.[39]; Gunukula et al. [40]	Requires training data from multiple generators; ensemble complexity scales with generator diversity

Table 2 Comparison of the four detection-method families for AI-generated images

attention-based models introduces "predictable anomalies" identifiable through multimodal analysis[37].

3.5 Style, Copying, and the Meaning of Statistical Novelty

The claim that AI generation involves "learning style" rather than "copying" merits careful technical scrutiny [38]. When a diffusion model or GAN is described as learning an artistic style, what it learns is a complex, multi-dimensional probability distribution over pixel configurations, textures, color palettes, compositional patterns, and other visual features that characterize that style within the training data. The apparent novelty of outputs arises from the stochastic nature of the sampling process (e.g., the random noise seed in diffusion models) and the high dimensionality of the latent space—a quantitative form of novelty based on interpolation and extrapolation within the learned data manifold, not a qualitative conceptual breakthrough.

This interpretation is supported by the metrics that the field itself uses to evaluate generative models. The works by Ho et al[29]. and Rombach et al. [30]consistently frame model quality in terms of Fréchet Inception Distance and Inception Score —metrics that measure how closely the generated distribution matches the real data distribution, and how recognizable generated images are to a pre-trained classifier. The optimization target is fidelity to the statistical properties of training data. Any apparent "transformativeness" is a byproduct of this statistical matching process, not a primary objective driven by creative intent. Human creativity, by contrast, involves the capacity to break established patterns, to innovate beyond training distributions, and to create works that are motivated by conceptual understanding, emotional expression, and cultural engagement—capabilities that current AI systems demonstrably lack, as the technical evidence reviewed above consistently indicates.

4 Detection, Fingerprints, and the Empirical Basis for Distinguishability

A recurring critique from the ACG community is that AI-generated images, while technically proficient, feel "empty" or "soulless"—lacking the contextual meaning and subjective coherence of human-created art. Skeptics have dismissed this as purely subjective aesthetic judgment. However, a growing body of detection and forensics research provides an empirical basis for this intuition: AI-generated images carry quantifiable, statistically robust characteristics that distinguish them from human-created works, and these characteristics are directly attributable to the recombinative generation process described in Section 3.

4.1 Four Families of Detection Methods

Current detection methodologies can be broadly categorized into four families, each targeting different aspects of the generative process. Table 2 provides a comparative summary.

4.2 Detailed Analysis of Detection Approaches

Frequency-domain detectors exploit the systematic failure of generative models to reproduce the full spectral distribution of natural images. As Durall et al.[31] demonstrated, CNN-based generators introduce characteristic spectral anomalies due to the local processing constraints of transposed convolutions. These anomalies persist in more advanced architectures as subtle but measurable deviations from natural frequency statistics.

Reconstruction-error-based detectors, exemplified by DIRE[33], leverage the observation that pre-trained diffusion models reconstruct real images with lower error than

they reconstruct synthetic images—even when the synthetic images were themselves produced by a diffusion model. This asymmetry reveals that generative models leave behind traces of their own generation process that are detectable through reconstruction analysis.

Feature-space and hybrid detectors combine low-level statistical features with high-level semantic embeddings. The AIDE framework[35] and FatFormer[37] demonstrate that inconsistencies between local pixel statistics and global semantic coherence provide complementary detection signals, reflecting the absence of holistic compositional planning in generative models.

Ensemble and generalization-focused detectors address the challenge that detectors trained on one generator often perform poorly on images from a different generator. Saskoro et al. (2024) documented this cross-generator generalization gap, showing that a model achieving 99.97% accuracy on BigGAN images could drop to approximately 50% accuracy when tested on images from Stable Diffusion[39]. Their Gated CNN ensemble approach achieved an average accuracy of 97.94% by leveraging expert models specialized for different generators. The very fact that generator-specific detectors are needed—and that generator-specific artifacts exist—provides strong evidence that each generative architecture leaves behind its own identifiable technical signature.

4.3 What Detectors Actually Measure

It is important to clarify what detection algorithms are measuring when they identify AI-generated images. The "fingerprints" identified by detectors are not deliberate markers (like digital watermarks) but rather unintended byproducts of the generation process: statistical deviations from the manifold of natural images, systematic biases introduced by specific architectural choices (transposed convolutions, attention patterns, latent-space compression), and the absence of the organic coherence that characterizes images produced through embodied human perception and intentional compositional planning.

The ACG community's preferred method for identifying AI-generated images — through the detection of artifacts — is not driven by an interest in artifacts for their own sake, but because of what they reveal. Consider a concrete example: a person who cannot draw will never produce characteristic tail artifacts or reproduce a specific anomalous shape that arises from the local-processing constraints of generative models. These artifacts are not random glitches; they are the symptomatic expression of a fundamentally non-human generative process. Many ACG community members cannot afford to run AI detection models, nor do they necessarily understand the technical principles of artificial intelligence, yet they can still visually identify AI-generated images. What they are seeing is not a deliberate departure from representational norms — as in abstract art or traditional Chinese painting, where the absence of Western perspectival conventions is a conscious aesthetic choice — but rather a peculiar kind of failed realism. This distinction is by no means trivial. As Jiang Li (2025) has documented in a comparative study of Eastern and Western perspective techniques, traditional Chinese painting employs "scattered perspective"— a deliberate organization of space according to narrative and expressive needs rather than geometric rules — rooted in the philosophical traditions of Confucianism, Buddhism, and Daoism, which prioritize artistic conception and the unity of subject and object over the mathematically accurate spatial reproduction that Western painting has pursued since the Renaissance[41] The crucial point is that Eastern painters know precisely what they are departing from; their abandonment of single-point perspective, their use of the "three-distance method" of high distance, level distance, and deep distance, and their embrace of multiple focal points within a single composition are acts of disciplined aesthetic choice. AI models, by contrast, are trained overwhelmingly on Western perspectival conventions — they are, if anything, excessively committed to realistic spatial rendering — yet they execute it inconsistently and incorrectly. The lighting is subtly wrong, the perspective is inconsistently applied, the shadows fall in directions that make no physical sense. These errors differ categorically from human stylistic abstraction, and it is precisely this quality — an uncanny overcommitment to realism that nevertheless fails — that the ACG community has learned to recognize. The artifacts matter because they are the visible trace of the absence of intentionality; we offer this reading as an interpretation grounded in the technical evidence reviewed above, not as a claim that the detection literature alone establishes.

4.4 Limitations of Detection as a Policy Tool

These limitations caution against overclaiming the reliability of detection as a policy tool. However, they do not undermine the core empirical finding that supports this paper's argument: AI-generated images are, under current technical paradigms, systematically and measurably distinguishable from human-created works. This distinguishability is not a contingent limitation that will necessarily be overcome, but rather a reflection of the fundamentally different generative process — statistical recombination versus intentional creation — that produces them.

The empirical distinguishability of AI-generated images provides a technical translation of what the ACG community has described in experiential terms as "soullessness." The "soul" of human art, in this framework, corresponds not to a mystical quality but to the measurable property of organic coherence: the consistent alignment of local detail with global compositional intent, grounded in the artist's embodied experience of physical reality and cultural context.

Human artists operate within what cognitive science describes as an embodied perceptual-motor system — an integrated loop of perception, action, and conceptual reasoning that naturally produces works with coherent properties across all spatial and semantic scales. AI models, by contrast, optimize for statistical fidelity to training distributions without grounding in physical or experiential reality. The gap between these two generative processes is what detection algorithms learn to measure, and it is this same gap that human viewers register, consciously or unconsciously, as the absence of "soul."

5 Discussion: Bridging Technical Reality and Community Critique

5.1 Community Discourse Analysis: Data Collection and Method

We understand ACG community members' sensitivity re garding the extent to which their publicly posted discour se is collected; out of respect for the community, we sum marize the disclosure as follows, with detailed disclosure provided in Appendix X.

This subsection describes an auditable procedure for coll ecting and analyzing public ACG discourse from Bilibili, Douyin, Xiaohongshu, QQ Zone, and WeChat Moments, each documenting a distinct segment. Only unrestricted public posts from 2023–2026 were included; exact dates and corpus size are withheld to prevent re-identification. Keyword searches used community and English terms. P osts were screened for substantive AI-art discussion and absence of personal or restricted content; duplicates were removed. Two authors independently conducted inducti ve thematic coding, resolved disagreements through disc ussion, and reported only themes and patterns, not indivi duals.

5.2 The Roots of the ACG Community's Intense Response

Equally important is the distinct social position of ACG culture in China. In several Western societies, video games have been celebrated as a legitimate art form — the "ninth art" — and anime enjoys substantial cultural recognition. In mainland China, by contrast, ACG subculture has historically faced marginalization and discrimination. Public figures such as Yang Yongxin gained notoriety for operating "internet addiction" treatment camps that subjected young gamers to involuntary interventions[43]; anime, comics, and gaming have long been dismissed by segments of mainstream Chinese society as frivolous, unproductive, or even pathological pursuits.

This history of stigmatization is essential context for understanding the defensive intensity of the ACG community's response to AI-generated art. The community's critiques — expressed through visceral metaphors such as "assembly of body parts" and comparisons of AI proliferation to disease — reflect not only concern about AI technology but also a longer experience of being dismissed and misunderstood by mainstream institutions.

When the ACG community describes AI art as a form of theft or an existential threat to creative livelihoods, it does so from a position of structural predicament in Chinese society. This predicament is documented in detail by Chen Wenmin's analysis of Chinese ACG culture[44].

Chen reads ACG as a visual culture born of youth subculture and the consumer society, and diagnoses four structural difficulties in its Chinese development.

The first is the absence of a children-oriented perspective: domestic animation has been locked to a child audience while mainstream adults dismiss it as childish, leaving works stranded between "adults disdaining to watch, children disliking to watch", a deadlock made visible when the series Hongmao and Lantu was suspended from television in 2007.

The second is the suspension of the spirit of play: Chinese ACG works put didactic instruction before play, using animation to teach lessons and emptying out the game spirit that Schiller took to be the essence of play.

The third is the scarcity of cartoon stars: unlike Disney or the Japanese industry, Chinese animation lacks internationally competitive original characters, so the industry keeps investing in new faces without accumulating brand value.

The fourth is the loss of original style: Chinese animation once possessed original styles rooted in paper-cutting, ink painting, and puppetry, but under the pressure of American and Japanese imports it fell into large-scale imitation and lost its visual identity. These four difficulties are structural rather than incidental, and they locate the community's precarious position in the social order itself, which is why the community's response to AI-generated art carries the intensity it does.

A comparison with Western fan communities makes the difference clear. As reported in Chinese media, the Western response to AI-generated art took institutional and legal forms: the mass "No AI" protest images that flooded ArtStation and DeviantArt in late 2022 and 2023[45], the NoAI tags and opt-out mechanisms on portfolio platforms, and litigation over training data and remuneration. The Chinese ACG community, lacking those institutional channels, developed a vernacular techno-critique of its own, running on the community's trained visual attention and its own Chinese vocabulary, which describes the production process and its visible traces and is aimed at recognizing AI-generated images directly rather than at claiming a legal remedy.

5.3 The Dialectical Framing

The analysis presented in this paper supports the conclusion that the ACG community's critiques of AI-generated art—while expressed in the vernacular of practice and

passion rather than in academic terminology—identify genuine technical properties of current generative AI systems. The metaphors of "assembly of body parts"and critiques of "style theft" are not merely emotional reactions but intuitive recognitions of the recombinative, statistically grounded, and derivative nature of AI image generation.

5.4 Community Metaphors as Interpretive Heuristics

The ACG community's characterization of AI art as an "assembly of body parts" accurately captures, in metaphorical language, the process by which diffusion models sample from learned probability distributions to combine visual features without holistic compositional planning. The detection research reviewed in Section 4 provides the empirical basis for this characterization: the spectral anomalies, structural inconsistencies, and generator-specific fingerprints documented in the literature are precisely the measurable consequences of a process that combines local features based on statistical likelihood rather than conceptual coherence.

If a paper about community discourse cannot directly engage with the community's own language — including its metaphors — then what, precisely, is the paper's purpose? The ACG community is not an academic community; it expresses its concerns through vernacular, emotionally charged metaphors. A paper that sanitizes this language or holds it at arm's length effectively distances itself from the very community it claims to analyze. Our position is not that the community is correct in every respect, nor that technical experts are wrong. It is a dialectical one: both sides have partial insights, and the scholarly contribution lies in identifying where each perspective captures something real and where each falls short.

To make this dialectical framing concrete, consider the mobile game *Reverse: 1999*, a title that has attracted substantial attention within the Chinese ACG community. The game's narrative centers on the "Storm", a catastrophic temporal phenomenon that causes eras to collapse backward through time, and the "Timekeeper" Vertin, the player-character who navigates this collapsing timeline. Its visual style — a distinctive blend of retro and modern aesthetics — and its thematic engagement with cultural memory make it an instructive case. When AI-generated images attempt to replicate the aesthetic of games like *Reverse: 1999*, the artifacts documented in Section 4 become especially visible: the nuanced color grading, the deliberate period-specific visual references, and the cohesive world-building that characterize the hand-crafted art direction are flattened into statistical approximations. The community's recognition that "something is missing" from AI-generated imitations is not nostalgia but an intuitive detection of the absence of the intentional artistic direction that gives such works their distinctive identity. This example illustrates, at the level of a specific cultural product, the broader dialectic this paper analyzes: the community's critiques, while expressed in vernacular language, point toward real technical properties — recombination, the absence of holistic compositional planning, the visible traces of non-human generation — that the forensics literature independently corroborates.

Similarly, the critique that AI "steals" artistic styles translates, in technical terms, to the fact that diffusion models trained on copyrighted artworks can reproduce the statistical signatures of those works—including distinctive stylistic elements—through latent space interpolation, without authorization from or compensation to the original creators. The ongoing *Getty Images v. Stability AI* illustrate the legal difficulty of testing these implications before the courts: Getty's principal training claim was abandoned on territorial grounds, the Outputs Claim was abandoned once the relief sought had been substantially achieved, and the remaining claims were largely rejected apart from limited trade mark findings (Section 2.3). The copyright status of unauthorized training and style reproduction therefore remains legally unsettled in the United Kingdom.

5.5 Implications for Law and Policy

The technical analysis presented here carries direct implications for the legal and policy frameworks discussed in Section 2:

Copyright registration: The USCO's position that AI-generated content lacks sufficient human authorship for copyright protection[1] is technically well-founded. The generation process, as described in Section 3, is one of statistical sampling from a learned distribution, not an act of human creative authorship—even when guided by human-authored prompts.

Transparency obligations: The EU AI Act's requirement that AI-generated content be labeled and detectable[3] is technically feasible given current detection capabilities (Section 4), but the arms-race dynamic between generation and detection means that detection alone cannot serve as a complete regulatory solution. Transparency must be built into the generation pipeline — through watermarking, metadata embedding, and provenance tracking — rather than relying solely on post-hoc forensic analysis. Furthermore, as noted in Section 2.2, Article 50 obligations are subject to a phased implementation schedule, and not all provisions are fully operative as of the date of writing.

Training data governance: The UK's movement toward a rights-reservation model for text and data mining[8] and the Getty litigation[13] both reflect a growing recognition that the provenance of training data matters for copyright analysis. The technical fact that model outputs are statistically constrained by training data distributions—and in some cases reproduce identifiable elements of training

data—supports the argument that AI training on copyrighted works without authorization raises legitimate legal questions that are distinct from those raised by human artistic influence and inspiration.

5.6 Counter-Arguments and Limitations

A balanced analysis must acknowledge counter-arguments and limitations. Some scholars argue that AI-generated outputs can exhibit emergent creativity through the exploration of latent space regions that correspond to novel combinations of learned features, and that human creativity itself can be understood as a process of recombining influences and experiences[38].

We acknowledge this perspective but note a crucial distinction: human recombination operates within a framework of conceptual understanding, intentional choice, cultural engagement, and embodied experience that current AI systems do not possess. The question is not whether AI outputs can be statistically novel—they can—but whether this statistical novelty constitutes creativity in any sense that is comparable to human creative practice. Three further objections are worth answering in the same spirit.

The first holds that the detection gap will close: as generators improve, the artifacts documented in Section 4 may fade, and with them the empirical distinguishability on which part of our reading rests. We answer that distinguishability is not our central claim. The central claim concerns the generation process itself, statistical sampling from a learned distribution, and that process does not change when the artifacts fade; a detector can become outdated while the recombinative character of the model remains exactly what it was.

The second holds that all criticism of AI art is romantic protectionism, a defence of a status hierarchy that technology is rightly dismantling. We answer that this objection mistakes the community's target. The ACG community is one of the most open and self-taught creative communities there is. What its members resist is not the tool itself, but the claim that bulk statistical recombination can substitute for the embodied, intentional practice they have trained in. For the vast majority of ACG community members, however, drawing is not a livelihood but a passion-driven hobby. Their resistance stems not from fears of job displacement, but from a defense of the personal meaning and craft embedded in the creative process itself.

The third holds that the community's metaphors are merely aesthetic prejudice, and that "soullessness" is a feeling rather than a finding. We answer that Section 4 already shows the feelings track measurable properties, and that a perspective paper's task is exactly to ask whether the intuition and the measurement describe the same thing. These objections are stated here as seriously as they are stated by our interlocutors, and answered one by one, because the dialogue itself is part of the argument.

We also acknowledge several limitations of our analysis. Our technical focus on diffusion-based models means that findings may not fully generalize to other paradigms, though the core argument about statistical recombination as the operational principle applies broadly. Our emphasis on the ACG community context, while providing a rich case study, may not capture the full diversity of global perspectives on AI art. The rapidly evolving nature of both generative technology and AI policy means that specific findings may require updating as new architectures and regulations emerge.

5.7 AI as Assistive Technology: Lessons from Adjacent Domains

An instructive parallel can be drawn from the domain of AI-assisted education, where the question is not whether AI can perform a task, but whether it should, and under what conditions of human oversight. The comparison this section draws is precise: between human supervision in assessment and human control in artistic creation. The two studies are described here as they actually are. Lampropoulos and Papadakis (2025) is a book chapter that introduces a volume on generative AI in education and is not an AI-grading study; it surveys how generative AI is being adopted in teaching and learning and stresses the continuing role of educators and instructional designers[46]. Zacharis and Papadakis (2025) is an empirical study of validity, reliability, and fairness in university assessment. Comparing AI and human scoring of ninety-one essays, they report that agreement between the AI rater and the human raters was substantially weaker than agreement between the two human raters, that the AI inflated scores by 2.71 to 3.32 points, and that current generative AI is not suitable for high-stakes assessment, although it may be useful for formative feedback when used transparently and under human oversight[47]. Both domains show the same structure. In assessment the answer is that AI may assist but must not replace the human rater; in art our argument is the parallel one, that AI may assist but human intentionality must remain the determining creative force. AI-assisted grading requires human educators to exercise final judgment and contextual interpretation, just as AI-assisted art creation should preserve human intentionality as the determining creative force.

5.8 Future Research Directions

Future research should pursue several directions identified by this analysis:

Technical: Continued development of robust, generalizable detection methods that can keep pace with improving generative architectures; investigation of whether the artifacts documented in current models reflect fundamental limitations of statistical generation or contingent properties of current implementations.

Legal: Empirical studies of the relationship between training data composition and output characteristics, to inform ongoing policy debates about fair use, licensing, and compensation frameworks.

Interdisciplinary: Collaborative work bridging computer science, media studies, law, and creative practice to develop nuanced frameworks for understanding and governing the relationship between AI generation and human creativity.
Community-engaged: Research that takes seriously the experiential knowledge of creative communities—including the ACG community—as a source of testable hypotheses about AI behavior and its social implications.

6 Conclusion

This paper has advanced a technically grounded perspective on the controversies surrounding AI-generated art, arguing that the operational principles of diffusion-based generative models provide a concrete empirical basis for understanding the derivative, recombinative nature of their outputs. We have argued, on the basis of the technical evidence reviewed above, that AI image generation is best understood as a process of statistical distribution learning and sampling — not as an exercise of intentional, conceptually grounded creativity — and that this distinction is measurable through the systematic artifacts and fingerprints documented by detection research. We emphasize the epistemic status of this claim: the technical literature establishes measurable properties of AI outputs, while the interpretive claims we build on those properties — concerning intentionality, creativity, and artistic value — are offered as a principled reading consistent with the evidence, not as conclusions that the detection literature alone proves.
The ACG community's intuitive critiques of AI art — expressed through metaphors such as "assembly of body parts" and concerns about unauthorized style replication — identify preliminary characteristics of generative AI that the technical literature corroborates in more precise terms. We interpret this alignment not as evidence that the community's critique is correct in every particular, but as an indication that the community has identified specific properties of the technology — recombination, training-data dependence, and the absence of intentional planning — that independent technical analysis confirms. An analogy may clarify the status of this alignment. To an observer unfamiliar with modern optics, the horizon appears flat and the sky seems to curve overhead like a dome; the observations are real, even though the interpretive framework is incomplete. We use this analogy with an explicit caveat: it is not intended to cast the ACG community as "pre-scientific" or naive. On the contrary, the community's trained visual attention is precisely what allows its members to perceive artifacts that untrained viewers miss. Their vernacular descriptions — "assembly of body parts," "soullessness," "style theft" — are imprecise but directionally accurate interpretive heuristics; they point toward real properties of generative models that the technical literature independently documents. The task of scholarship is not to dismiss these observations because their framing is vernacular, but to translate them into a technical vocabulary that reveals their empirical basis. The gap between statistical image synthesis and the human values at the center of creative communities — intentionality, informed consent, and genuine expression — is real and measurable, even though its normative significance remains a matter of interpretation.
The evolving legal and policy landscape—from the US Copyright Office's human authorship requirement to the EU AI Act's transparency obligations to the UK's reconsideration of text and data mining exceptions—reflects a growing institutional recognition that technical understanding must inform regulatory approaches to AI-generated content. As these frameworks continue to develop, grounding policy in technical evidence about how generative models actually operate will be essential for creating governance structures that protect creative rights while enabling beneficial applications of AI technology.
The challenge ahead is not merely to build more powerful generative models, but to develop frameworks—technical, legal, and social—that enable AI tools to serve as genuine assistants to human creativity rather than as unauthorized appropriators of it. The creative communities that have been most vocal in critiquing AI art have correctly identified the problem; the task for researchers, policymakers, and technologists is to build solutions grounded in technical understanding and responsive to the legitimate concerns those communities have raised.

List of abbreviations

For reasons concerning the retrievability of primary sources, the accuracy of registered records, the identity of named individuals, and the integrity of ethics statements—not the journal's English-language standard, since the manuscript body remains entirely in English—the author team, after communicating with the editors, presents the following correspondence table. The complete table, including the original-language (Chinese) entries, is provided in Appendix X.

ACG	Anime/Comics/Games
AI	Artificial Intelligence
CDPA	Copyright, Designs and Patents Act 1988
CNN	Convolutional Neural Network
CPC	Communist Party of China
DDPMs	Denoising Diffusion Probabilistic Models

DIRE	Diffusion Reconstruction Error
EU AI Act	European Union Artificial Intelligence Act
Feilin	Beijing Feilin Law Firm
FID	Fréchet Inception Distance
GAN	Generative Adversarial Network
GPAI	General-Purpose AI
IS	Inception Score
LDMs	Latent Diffusion Models
PRC	People's Republic of China
TDM	Text and Data Mining
Tencent	Shenzhen Tencent Computer System Co., Ltd.
UKIPO	United Kingdom Intellectual Property Office
USCO	United States Copyright Office

Author Contributions

Lin XingJun: Conceptualization, Investigation, Writing – Original Draft.
Huang ZiRan: Methodology, Supervision, Writing – Review & Editing. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest:

The first author, Lin XingJun, declares a personal interest in ACG culture, which served as a source of inspiration and motivation for this work. The second author, Huang ZiRan, provided professional supervision and technical support throughout the research process and has no competing interests to declare. Both authors confirm that their personal or professional interests have not influenced the integrity, objectivity, or scientific rigor of this study.

Funding:

This work is sponsored by Guangzhou College of Commerce 2025 College Students' Innovative Entrepreneurial Training Plan Program called “The Inner World Therapy – A Psychological Counseling and Online Triage Platform Based on the DeepSeek-R1 Large Model.”(源心智疗——基于deepseek-r1大模型的心理疏导及线上导诊平台).

Acknowledgments:

Lin XingJun deeply indebted to the project titled "The Inner World Therapy – A Psychological Counseling and Online Triage Platform Based on the DeepSeek-R1 Large Model" (源心智疗——基于deepseek-r1大模型的心理疏导及线上导诊平台) , which served as the foundational inspiration for this research. The project directly confronted us with questions about the ethical boundaries of AI deployment, the nature of human-AI interaction, and the distinction between assistive and autonomous AI functions. These questions, while encountered in a healthcare context, generalize directly to the creative domain, and the conceptual and ethical framework developed during that project directly motivated the inquiry into AI-generated art presented in this paper.
Special thanks are due to the Advanced Computing & Knowledge Engineering Research Lab (先进计算实验室) and the Machine Learning and Artificial Intelligence Laboratory at Guangzhou College of Commerce for providing the research environment and resources that supported this work.
Lin XingJun appreciation also goes to the MY Anime Association (MY动漫协会) at Guangzhou College of Commerce, which has fostered a community that, while not professionally oriented, is rich in creative passion and a shared love for anime culture. It is within this vibrant space that many of the observations discussed in this study took shape.
Lastly, Lin XingJun wish to express personal thanks to several friends he met through our shared interest in ACG: 薇芙, 拾一月久, 椁姬, 一枕槐安. Their specific contributions were of two kinds. Some provided suggestions and inspiration during our exchanges, which helped shape the direction of the research. Others supplied reference materials, including community posts and platform discussions, that directly informed our analysis of ACG discourse in Section 5.
To avoid any confusion among readers about who wrote this paper, we add a word about author status. Under the journal's authorship criteria, authorship requires a substantial intellectual contribution to the conception, drafting, or critical revision of the manuscript. The friends thanked here contributed suggestions, inspiration, and reference materials, contributions we value deeply, but they did not take part in conceiving, drafting, or revising the manuscript. It is for this reason, and not because their contributions mattered little to us, that they appear in the

Acknowledgements rather than on the author list. Lin XingJun is grateful for the role these friends have played in my life and this study, regardless of the paths our relationships have taken.
Huang ZiRan acknowledges the technical discussions and collaborative environment at the School of Modern Information Industry that contributed to the methodological framing of this work.

AI Declaration

AI-based tools (DeepSeek R1/V3, DeepSeek V4-Flash/Pro) were used for English-language editing: grammar correction, sentence-level rephrasing for clarity, and consistency checking of terminology across sections. These tools were not used to generate arguments, select or interpret literature, draft original analytical content, or produce the paper's structure. All argumentation, source selection, and conclusions are the authors' own and were independently verified.

Consent for Publication

We have over the past six months explained to reviewers, the journal editorial office, and here that the ethics committee has not issued a specific approval document.
More precisely, we submitted our materials well in advance, but the ethics committee has not responded; throughout the past six months, all ethics-related correspondence with the journal editors has been copied to the Practice Department of the Youth League Committee of the first author's college—the unit responsible for student research and competition affairs. This may be because this study does not involve biomedical research with human subjects and therefore falls outside the routine remit of institutional ethics committees under China's ethical and judicial framework.
All data were processed within mainland China in accordance with applicable Chinese laws, including the Personal Information Protection Law. The analysis reports only generalized themes and patterns from publicly available material, without quoting individuals or disclosing identifiable information. References to public figures and scholarly works are limited to their published statements and research, following standard citation and critique practices.
Nevertheless, given the ACG community's sensitivity toward AI, and out of respect for the community, we conducted a detailed ethical discussion, as set out in Appendix X.

Supplementary Material:

Supplementary Materials Supplementary material associated with this article has been published online and is available at: [DOI:…]

References

1. United States Copyright Office. Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence. Federal Register. 2023 Mar 16;88:16190.
2. United States Copyright Office. Registration Decision: Zarya of the Dawn (VAu001480196) [Internet]. 2023. Available from: https://copyright.gov/ai/
3. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act). Official Journal of the European Union [Internet]. 2024 Jun 13. Available from: https://eur-lex.europa.eu/
4. McCann FitzGerald. AI Transparency: European Commission's Guidelines on Article 50, Part 1 (Provider Obligations) [Internet]. 2026. Available from: https://www.mccannfitzgerald.com/knowledge/data-privacy-and-cyber-risk/ai-transparency-european-commissions-guidelines-on-article-50-part-1-provider-obligations
5. Ropes & Gray LLP. You Talkin' To Me? Operationalising The EU AI Act's Transparency Obligations [Viewpoints]. Ropes & Gray Insights [Internet]. 2026. Available from: https://www.ropesgray.com/en/insights/viewpoints/2026/08/102nfqm/you-talkin-to-me-operationalising-the-eu-ai-acts-transparency-obligations
6. Morgan, Lewis & Bockius LLP. EU AI Act's Transparency Rules: What Went Into Effect on 2 August? [Blog Post]. Tech & Sourcing @ Morgan Lewis [Internet]. 2026. Available from: https://www.morganlewis.com/fr/blogs/sourcingatmorganlewis/2026/08/eu-ai-acts-transparency-rules-what-went-into-effect-on-2-august
7. Debevoise & Plimpton LLP. The Third Wave of EU AI Act Requirements Are In Force: Transparency Requirements & Supervisory Powers [Blog Post]. Debevoise Data Blog [Internet]. 2026. Available from: https://www.debevoisedatablog.com/2026/08/03/the-third-wave-of-eu-ai-act-requirements-are-in-force-transparency-requirements-supervisory-powers/
8. UK Intellectual Property Office. Artificial Intelligence and Copyright: Consultation on Text and Data Mining [Internet]. 2023. Available from: https://www.gov.uk/government/consultations/artificial-intelligence-and-ip-copyright-and-patents
9. UK Government. Report on Copyright and Artificial Intelligence [Internet]. Crown copyright; 2026. Available from: https://www.gov.uk/government/publications/report-and-impact-assessment-on-copyright-and-artificial-intelligence/report-on-copyright-and-artificial-intelligence

10. Reed Smith LLP. UK copyright and AI report: the ‘opt-out’ is dead, but what comes next? [Client Alert]. Reed Smith Client Alerts [Internet]. 2026. Available from: https://www.reedsmith.com/articles/uk-copyright-and-ai-report-the-opt-out-is-dead-but-what-comes-next/
11. Berry K, Huang N. Hurry up and wait: UK government still non-committal on AI copyright reforms. DigiLinks Blog [Internet]. Linklaters LLP; 2026. Available from: https://www.linklaters.com/insights/blogs/digilinks/2026/march/hurry-up-and-wait-uk-government-still-non-committal-on-aI-copyright-reforms
12. Getty Images (US), Inc. and Others. Getty Images (US), Inc. & Ors v Stability AI Ltd. 2025.
13. High Court of Justice of England and Wales. Getty Images (US), Inc. v. Stability AI Ltd. EWHC. 2023;3090(Ch).
14. Wiggin LLP. High Court grants permission to appeal in Getty Images v Stability AI [Insight Article]. Wiggin Insight [Internet]. 2026. Available from: https://www.wiggin.co.uk/insight/high-court-grants-permission-to-appeal-in-getty-images-v-stability-ai/
15. Standing Committee of the National People’s Congress. Copyright Law of the People’s Republic of China (中华人民共和国著作权法) [Internet]. 2010. Available from: https://www.nppa.gov.cn/xxgk/fzdgknr/zcfg_210/fl_211/201604/t20160429_4556.html
16. Beijing Feilin Law Firm (北京菲林律师事务所) v. Beijing Baidu Netcom Science and Technology Co., Ltd. (北京百度网讯科技有限公司). Beijing Feilin Law Firm v. Beijing Baidu Netcom Science and Technology Co., Ltd. (北京菲林律师事务所诉北京百度网讯科技有限公司著作权侵权纠纷案) [Internet]. 2019. Available from: https://www.ciplawyer.cn/articles/142045.html
17. Shenzhen Tencent Computer System Co., Ltd. (深圳市腾讯计算机系统有限公司) v. Shanghai Yingxun Technology Co., Ltd. (上海影讯科技有限公司). Shenzhen Tencent Computer System Co., Ltd. v. Shanghai Yingxun Technology Co., Ltd. (深圳市腾讯计算机系统有限公司诉上海影讯科技有限公司著作权侵权及不正当竞争纠纷案). 2019.
18. Feng Q. Comparative Study on the Copyright Protection of AI-Generated Content by ChatGPT (ChatGPT人工智能生成物的著作权保护比较研究) [Master’s thesis] [Internet]. [Shanghai]: East China University of Political Science and Law; 2025. Available from: https://kns.cnki.net/kcms2/article/abstract?v=xyoQqgjBopoVSx_j-ZEYAqd2phGTUWBuMUCfOGeu8hq58SbsCtHVRaSh36hhQhApHm4xNKzNgMHPbrJOYvo6kmtXE83NxW65tilyKrP89JKat-XJWMMIpAE4V5IU8ICN7hH2abDVPZ933zfItcF6xCRjt6C6EtvNpAiCMveXfB0tLBnsC_yOVN_TeOWIGFHo&uniplatform=NZKPT&language=CHS
19. Court (上海知识产权法院) SIP. Medusa AI Copyright Case (美杜莎AI著作权案) [Internet]. 2026. Available from: https://www.ciplawyer.cn/articles/159022.html
20. Lu N(卢娜), Wu K(吴开龙). Building Judicial Safeguards for Emerging Fields to Boost New Quality Productive Forces: Jinan Shizhong District Court Holds Intellectual Property Judicial Protection Press Conference (筑牢新兴领域司法屏障护航新质生产力发展 济南市中区法院召开知识产权司法保护新闻发布会) [Internet]. 2026. Available from: https://jinan.dzwww.com/qcxw/202604/t20260423_17663620.htm
21. Tang G(汤广花). AI-Generated Image Used Commercially, Why Was the Copyright Claim Rejected? (AI生成图遭商用 维权为何被驳回？). China News and Publishing Gazette (中国新闻出版广电报) [Internet]. 2026. Available from: https://epaper.chinaxwcb.com/epaper/2026-06/18/content_99879903.html
22. Civil Code of the People’s Republic of China (《中华人民共和国民法典》) [Internet]. Beijing: China Legal Publishing House; 2020. Available from: https://www.moj.gov.cn/pub/sfbgw/zwgkztzl/2025nianzhuanti/2025mfdxcy/2025mfdxcy_mfdql/202505/t20250507_518708.html
23. Cyberspace Administration of China, National Development and Reform Commission, Ministry of Education, Ministry of Science and Technology, Ministry of Industry and Information Technology, Ministry of Public Security, et al. Interim Measures for the Management of Generative Artificial Intelligence Services (《生成式人工智能服务管理暂行办法》) [Internet]. Beijing: Cyberspace Administration of China; 2023. Available from: https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm
24. Hu J. Firmly March along the Path of Socialism with Chinese Characteristics and Strive to Complete the Building of a Moderately Prosperous Society in All Respects——Report of the 18th National Congress of the Communist Party of China (《坚定不移沿着中国特色社会主义道路前进 为全面建成小康社会而奋斗——在中国共产党第十八次全国代表大会上的报告》) [Internet]. Beijing: People’s Publishing House; 2012. Available from: https://kns.cnki.net/kcms2/article/abstract?v=xyoQqgjBopoKCCT6P2vFip_55qmnR_3gYcxjb7PYf5-TOhLtOJDTZ_lKvWEZsBqpSy3o8uQ3hT56UZEq5K5dl3asNGNwjuoxZZBRwt9r1vQKUztHSDvRio40fVQIRn06lcO42lLmA9ifsX24lbebB0LHpBCQNPhVglkv5Qp98N5BCwqfKvE_QQ==&uniplatform=NZKP

T&language=CHS
25. Central Committee of the Communist Party of China. Decision of the Central Committee of the Communist Party of China on Several Major Issues Concerning Comprehensively Advancing the Rule of Law (《中共中央关于全面推进依法治国若干重大问题的决定》) [Internet]. Beijing: People's Publishing House; 2014. Available from: https://www.idcpc.org.cn/zgzc/zcwj/201912/t20191216_105923.html
26. Jiang Q. Party Leadership Regulations: An Important Identifying Concept of Xi Jinping Thought on the Rule of Law (党的领导法规：习近平法治思想的重要标识性概念). Political Science and Law (政治与法律). 2026;2026(7):2–16.
27. General Office of the Communist Party of China Central Committee. Guidelines on Cultivating and Practicing Core Socialist Values (《关于培育和践行社会主义核心价值观的意见》) [Internet]. Beijing: People's Publishing House; 2013. Available from: http://www.moe.gov.cn/jyb_xwfb/s5147/201312/t20131224_161114.html
28. Fan Y(樊英). On the Judicial Application of Documents Jointly Formulated by Party and Government Organs (论党政机关联合制定文件的司法适用). Law Review (法学评论). 2024;(2):84–99. doi:10.16823/j.cnki.10-1281/d.2024.02.011
29. Ho J, Jain A, Abbeel P. Denoising Diffusion Probabilistic Models. In. arXiv; 2020 [cited 2025 Nov 29]. p. 6840–51. Available from: http://arxiv.org/abs/2006.11239 doi:10.48550/arXiv.2006.11239
30. Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B. High-Resolution Image Synthesis with Latent Diffusion Models. In. arXiv; 2021 [cited 2025 Nov 29]. Available from: https://arxiv.org/abs/2112.10752 doi:10.48550/ARXIV.2112.10752
31. Durall R, Keuper M, Keuper J. Watch Your Up-Convolution: CNN Based Generative Deep Neural Networks Are Failing to Reproduce Spectral Distributions. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) [Internet]. Seattle, WA, USA: IEEE; 2020 [cited 2025 Nov 29]. p. 7887–96. Available from: https://ieeexplore.ieee.org/document/9157579/ doi:10.1109/CVPR42600.2020.00791
32. Corvi R, Cozzolino D, Poggi G, Nagano K, Verdoliva L. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In. 2023. p. 973–82. doi:10.1109/CVPRW59228.2023.00104
33. Wang Z, Bao J, Zhou W, Wang W, Hu H, Chen H, et al. DIRE for Diffusion-Generated Image Detection. In. 2023. p. 22388–98. doi:10.1109/ICCV51070.2023.02051
34. Zhang H, Goodfellow I, Metaxas D, Odena A. Self-Attention Generative Adversarial Networks [Internet]. arXiv; 2018 [cited 2025 Nov 29]. Available from: https://arxiv.org/abs/1805.08318 doi:10.48550/ARXIV.1805.08318
35. Yan S, Li O, Cai J, Hao Y, Jiang X, Hu Y, et al. A Sanity Check for AI-generated Image Detection. In: Yue Y, Garg A, Peng N, Sha F, Yu R, editors. International Conference on Representation Learning [Internet]. 2025. p. 70702–20. Available from: https://proceedings.iclr.cc/paper_files/paper/2025/file/b0303773962ea1b5394c3a83cc7dd066-Paper-Conference.pdf
36. Zhu M, Chen H, Huang M, Li W, Hu H, Hu J, et al. GenDet: Towards Good Generalizations for AI-Generated Image Detection [Internet]. arXiv; 2023 [cited 2026 Jan 7]. Available from: http://arxiv.org/abs/2312.08880 doi:10.48550/arXiv.2312.08880
37. Liu H, Tan Z, Tan C, Wei Y, Wang J, Zhao Y. Forgery-aware Adaptive Transformer for Generalizable Synthetic Image Detection. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) [Internet]. 2024 [cited 2026 Jan 7]. p. 10770–80. Available from: https://ieeexplore.ieee.org/document/10657197 doi:10.1109/CVPR52733.2024.01024
38. Meyer R. The New Value of the Archive: AI Image Generation and the Visual Economy of 'Style'. IMAGE Zeitschrift für interdisziplinäre Bildwissenschaft. 2023;IMAGE. Zeitschrift für interdisziplinäre Bildwissenschaft19(1):100–11. doi:http://dx.doi.org/10.25969/mediarep/2231
39. Ahmad Fattah Saskoro R, Yudistira N, Noor Fatyanosa T. Detection of AI-Generated Images From Various Generators Using Gated Expert Convolutional Neural Network. IEEE Access. 2024;12:147772–83. doi:10.1109/ACCESS.2024.3466614
40. Gunukula AR, Das Gupta H, Sheng VS. Detecting AI-Generated Images Using a Hybrid ResNet-SE Attention Model. Applied Sciences. 2025 Jul 2;15(13):7421. doi:10.3390/app15137421
41. Jiang L. The Difference between Eastern and Western Painting Perspective and Its Integration in Illustration Creation (《东西方绘画透视法差异及其在插画创作中的融合》). Journal of Humanities and Social Sciences. 2025;1(1):1–9. doi:10.70693/rwsk.v1i1.470
42. Bilibili (哔哩哔哩). Announcement: Official Statement from the Pleasant Goat and Big Big Wolf Fan Creation Alliance Regarding the Official's Use of AI (【公告】喜灰同人二创联盟关于官方使用AI事件的正式声明). Bilibili 动态 (Bilibili Updates) [Internet]. 2024. Available from: https://www.bilibili.com/opus/915465205363769376
43. Daily CY. The Ecosystem of an Internet Addiction Treatment Center (《一个网戒中心的生态系

统》). China Youth Daily. 2009 May 7.
44. Chen W(陈伟). An Analysis of the Current Direction of Chinese ACG Culture (当下国内ACG文化进路解析). Journal of Jishou University (Social Science Edition) (吉首大学学报(社会科学版)). 2012;33(2):165–8. doi:10.3969/j.issn.1007-4074.2012.02.031
45. Zhao Y(照月). ArtStation Users Are Spamming to Boycott AI Artworks (ArtStation 用户正在通过刷屏来抵制 AI 画作) [Internet]. 2022. Available from: https://www.ithome.com/0/661/565.htm
46. Lampropoulos G, Papadakis S. Transforming Teaching and Learning Through Generative Artificial Intelligence. In: Artificial Intelligence in Higher Education. Berlin/Boston: De Gruyter; 2025. p. 1–8. doi:10.1515/9783112206393-001
47. Zacharis G, Papadakis S. Can AI Grade Like a Human? Validity, Reliability, and Fairness in University Coursework Assessment. Educational Process: International Journal. 2025;19:e2025591. doi:10.22521/edupij.2025.19.591
48. Shao Y. Current Status and Prospects of the Management System of Ethics Committees in China (关于我国伦理委员会管理制度的现状及展望). Progress in Pharmaceutical Sciences. 2021;45(9):641–4.

Figure 1 Schematic overview of the diffusion model pipeline. **(Top)** Forward diffusion: a real image is progressively corrupted by adding Gaussian noise over T steps until it becomes pure noise. The model is trained to predict the noise added at each step. **(Bottom)** Reverse denoising: starting from pure noise, the trained model iteratively removes predicted noise to reconstruct an image. The generation process is stochastic — different random noise seeds produce different outputs from the same learned distribution.

