ISSN: 2688-8653

Substantial Similarity in Generative AI Outputs: Difficulties and Rule Reconstruction

Mingcheng Deng[1,*]

[1]Department of Logistics Management, Nanchang Jiaotong Institute, Nanchang, Jiangxi Province, 330100, China

[*]Corresponding author's e-mail: 1922757559@qq.com

Abstract

Generative AI does not make substantial similarity irrelevant. It changes the factual basis and evidentiary conditions under which the doctrine operates. A model output may resemble a prior work, but resemblance alone does not prove copying of protected expression. Training data, parametric learning, prompts, and post-processing jointly shape outputs. They therefore separate evidence of similarity from the cause of similarity. Existing scholarship has focused on training legality, copyrightability, and the allocation of responsibility. It has not fully explained how attribution should be reasoned at the output stage. This article argues that courts should reorder the inquiry within the existing infringement structure. They should filter expression first, infer source second, and apply a normative backstop last. This sequence can protect authors without turning style or genre conventions into private property.

Keywords

Generative AI; substantial similarity; copyright infringement; expression filtering; source inference.

1. Introduction

Generative AI has changed the factual basis of substantial similarity analysis. The contemporary dispute is no longer limited to whether machines can be authors. A more common issue is whether a model output resembles a prior work in character designs, composition, narrative structure, or specific language. Courts cannot stop at visual or textual comparison. They must also explain whether the overlapping elements are protected expression. They must further decide whether resemblance supports an inference of copying. Traditional copyright infringement often depends on access and substantial similarity. Only after independent creation, genre conventions, and public-domain materials are excluded can similarity carry attributional force [1].

Generative AI makes that inference unstable. Models are trained on large text, image, and multimodal datasets. After training, work-related information is often transformed into parameters, weights, and statistical associations. Outputs are then shaped by prompts, random seeds, model versions, filtering rules, and post-processing. Even where an output resembles a prior work, the right holder may struggle to identify a specific source. The resemblance may come from memorization. It may also come from a probabilistic combination of common genre materials. As a result, similarity at the level of appearance becomes separated from similarity at the level of cause.

While the doctrine of substantial similarity remains valuable in copyright adjudication, it lacks a well-defined sequential framework in generative AI disputes. Courts should not impose liability solely because two outputs appear close at the surface. At the same time, they should not leave every black-box evidentiary difficulty to right holders. A workable approach should

ISSN: 2688-8653

therefore proceed in order: protected expression must be identified first, source connection should be inferred second, and normative attribution should be evaluated last.

This issue has a distinct Chinese-law dimension. Current Chinese debates often concern three questions: whether AI-generated content can be a work, whether rights belong to users or service providers, and whether training data use is justified. Wang Qian discusses the effect of generative AI on expressive elements. Wu Handong examines machine creation from the perspectives of subject, object, and rights allocation [2][3]. Beyond these debates, a narrower question remains. When a model output resembles a prior work, how can a court write an infringement rationale that can be tested?

The scope of this article is deliberately limited. It does not address every copyright issue raised by generative AI. It examines only similarity analysis at the output stage. This limited focus avoids excessive generality. It also helps develop rules that can be used in judicial reasoning.

Accordingly, the central question is not merely whether two works are similar. The more important question is how courts should explain the source and legal significance of resemblance. Only by separating evidence, causation, and normative evaluation can substantial similarity avoid becoming a matter of visual intuition alone.

2. Main Approaches in Existing Research and Their Limits

2.1. Chinese scholarship: from copyrightability to liability

Chinese scholarship first concentrated on the copyrightability of AI-generated content. The debate often began with the requirement of human intellectual creation. One view emphasizes that works must reflect human creative contribution. Another view holds that an output may enter the copyright system if it contains original expression. Xiong Qi discussed the copyright status of AI-generated content at an early stage. His analysis suggests that protection cannot be detached from the actual relationship between humans and the generative process [4].

Research then shifted toward ownership. After the Beijing Internet Court AI text-to-image case, both the user-based view and the developer-based view gained attention. The user-based view stresses prompts, parameter choices, selection, and revision. Critics emphasize that the model plays a substantial role in visual details and expressive combinations. The U.S. Copyright Office report on copyrightability takes a comparable position. It returns protection to human creative choices and their manifestation in the final work [5]. Recent work on photography and generative AI also indicates that protection of visual works must turn on how concrete choices appear in the final expression [6]. This reasoning is relevant to Chinese doctrine as well.

A third strand concerns liability. The training stage raises issues of reproduction, licensing, and fair use or analogous limitations. The output stage concerns comparison between a specific output and a specific work. The two stages are related, but they are not identical. Treating them as one issue may obscure a key step. When an output resembles a prior work, the court must still explain why copying can be inferred.

This mixed discussion has a practical basis. Training and output are technically connected. Training data affect model capacity, and model capacity affects the output. Legal assessment, however, should not simply follow technical continuity. The training stage asks whether input use is lawful. The output stage asks whether specific expression has been reproduced. These issues should be analytically separated.

Without separation, liability may move too quickly in two opposite directions. One approach infers output legality from training legality. Another infers infringement from the mere presence of works in training data. Both conclusions are premature. Both bypass the required comparison of expression and the explanation of source.

ISSN: 2688-8653

2.2. International scholarship: from fair use to output liability

International scholarship has noticed this structural change earlier. Samuelson places generative AI copyright disputes within the broader tension between technological change and creative incentives. She shows that training, output, and authorship questions together pressure copyright law [8]. Lemley further argues that generative AI changes several assumptions in copyright law. The problem is not limited to authorship or training data. The idea-expression dichotomy and substantial similarity tests also face structural pressure [9]. Sag focuses on model memorization, data scale, and generation risks. His account shows that output risk should not be assessed only by abstract technical labels [10].

Sobel pushes the discussion toward style. He argues that the statement "style is not copyrightable" does not resolve every dispute. Some elements described as style may contain specific expressive arrangements. This is especially true in visual art. Courts should not rely on a style label alone. They should ask whether the resemblance falls within concrete expressive combinations [11]. Hacohen and Elkin-Koren also suggest that generative AI can be used to evaluate originality and copyright scope. Their argument indicates that expression-boundary analysis may obtain new technical assistance [12].

Rosati places greater emphasis on output liability. She observes that AI copyright debates have long focused on training data and text-and-data mining. Whether outputs constitute actionable copying should be treated as a separate problem [13]. Meanwhile, scholarship on the training stage shows that U.S., EU, and international approaches evaluate training data differently [14]. Comparative research also demonstrates that China, the United States, and the European Union place emphasis on different institutional concerns [15]. This direction fits the present inquiry. Training legality cannot automatically determine output legality. Output liability still requires an independent account.

2.3. Insufficiencies in existing research and this article's entry point

Existing research has reached several basic points. Generative AI should not be assessed only by final outputs. Training, prompting, selection, and dissemination may all affect legal evaluation. Copyrightability, ownership, and infringement liability should also be separated. Traditional copyright rules have not disappeared. Their assumptions require reinterpretation.

The difficulty is that these approaches remain insufficiently connected. Copyrightability research asks whether generated content can be protected. Training-data research focuses on input-side reproduction and licensing. Platform-liability research examines risk allocation. Traditional substantial similarity research compares expression between works. Each strand contributes something. Yet none fully explains how expression scope, source connection, and substitution risk should be joined in a case where an AI output is highly similar to a prior work.

This gap directly affects judicial reasoning. A court facing an AI output cannot rely only on originality rules. Nor can it decide the case only by asking whether the service provider complied with administrative duties. It must complete comparison, source explanation, and harm evaluation. Existing research has not fully connected these three steps.

The "cause of similarity" discussed here is not merely technical causation. It includes both factual causation and normative attribution. A court must explain how resemblance emerged. It must also decide whether that mode of emergence is sufficient to support responsibility.

The focus, therefore, is not whether AI can be an author. This article examines substantial similarity in output-stage infringement. The issue is not simply whether two works look alike. The key questions are why they look alike and whether the resemblance deserves legal attribution.

ISSN: 2688-8653

3. Traditional Substantial Similarity Analysis and Its Limits

3.1. The inferential structure of access and substantial similarity

Traditional copyright infringement analysis rests on an inferential structure. The right holder usually proves that the defendant had access to the prior work. The court then compares whether the two works are similar at the level of protected expression. Copying becomes plausible only when both access and similarity exist. This structure works because human creative processes are relatively visible. Drafts, sources, transactions, and dissemination paths can explain the source of similarity.

Substantial similarity is not a numerical test. Courts should not simply count identical elements. They must first filter unprotected content. Themes, facts, ideas, and public-domain materials should not be monopolized. Similarity becomes attributable only when characters, plot development, detailed arrangements, or visual combinations form concrete correspondences.

Chinese judicial practice follows a comparable method. Courts commonly identify protectable expression first. They then ask whether the accused work reproduces that expression. Liability is assessed with reference to access, use, and consequences. Relevant rules of the Supreme People's Court also require separation between original expression and public-domain elements [16].

The Shanghai Animation Film Studio v. Chongqing Yunmei case illustrates this logic. The court did not rest on an intuitive finding of resemblance. It focused on the use of core characters, plot fragments, and image-related interests. It also considered substitution and disparagement effects. Similarity analysis therefore served attributional reasoning [17][18].

3.2. The idea-expression dichotomy and expression filtering

Expression filtering is the premise of substantial similarity analysis. Copyright protects expression. It does not protect ideas, methods, concepts, styles, or facts. This principle prevents overextension of rights. Filtering is especially important for generative AI. Model outputs often resemble prior works through style. If filtering is weak, shared aesthetic language may become privatized.

Expression filtering can be divided into five levels. First, courts should exclude subject matter and themes. Second, they should exclude general stylistic descriptions. Third, they should exclude public-domain elements. Fourth, they should exclude genre conventions. Fifth, they should exclude technical generation defaults. Only the remaining concrete expression should enter further comparison.

Filtering does not mean that anything called style must be excluded. Some styles may contain stable expressive combinations. Sobel's work reminds us that a style label can conceal concrete expression [11]. Hacohen and Elkin-Koren's work on measuring originality and copyright scope also suggests that courts can examine expressive combinations more precisely, rather than stopping at the abstract label of style [12]. Courts should return to the work itself. The key question is not whether style is protectable in the abstract. It is whether the resemblance points to identifiable expressive arrangements.

3.3. The limits of the traditional mechanism in AI cases

The traditional mechanism assumes that the source of similarity can be reconstructed. In human-creation cases, access may be proved through reading, viewing, cooperation, or dissemination. AI cases are different. Users may never have accessed the prior work. Service providers may not know which work a particular output will resemble. The model also does not store the "prior work-output" relationship in a human-readable way.

If courts require right holders to prove specific extraction strictly, remedies may fail. If courts infer infringement from output similarity alone, they risk suppressing expression and

ISSN: 2688-8653

innovation. The difficulty of AI cases lies in this tension. Under-protection and over-protection are both possible. The three-step framework proposed below is designed to control both risks.

This limit is particularly visible in image generation. Many outputs do not reproduce a complete work. They may reproduce a character posture, compositional habit, or color arrangement. A single element may appear ordinary. Several elements together may form a highly similar expressive whole. Traditional rules can still respond. They need a more refined sequence.

Text generation raises a similar problem. A model may output identical sentence patterns, narrative arrangements, or distinctive errors. If these materials are highly recognizable, they should not be dismissed as common expression. If they are ordinary formulations, they should remain outside protection.

4. The Recognition Difficulties of Similarity in Generative AI Outputs

4.1. Decoupling between evidence of similarity and cause of similarity

The central difficulty in AI output infringement is causal decoupling. Resemblance between an output and a prior work proves only outward similarity. It does not directly prove copying. At least four causal paths are possible. First, the model may have memorized the work in the training set. Second, the model may have generalized from many similar works. Third, the user's prompt may have directly targeted the work. Fourth, resemblance may arise from limited subject matter or public-domain elements.

This distinction matters for evidence law. If similarity comes from model memorization, infringement risk is higher. If it is a genre result, filtering should be stronger. If the prompt directly points to a prior work, the prompt becomes an attributional fact. If resemblance comes from the public domain, liability should not follow. Appearance comparison alone cannot distinguish these situations. Courts need indirect evidence and normative evaluation.

Recent disputes already show this problem. New York Times v. OpenAI concerns reproduction of news text, summary substitution, and subscription-market effects. Getty Images v. Stability AI concerns image training, output watermarks, and cross-border liability. These cases remain in development. They do not yet provide stable doctrine. Still, they show why the output stage deserves independent analysis [19][20][21].

4.2. Confusion between style, expression, and genre elements

A second difficulty is the confusion between style and expression. Generative AI is good at reproducing visual atmosphere, brushwork, color, and composition. Users also often design prompts in the "style of" a named artist or work. Outputs can therefore resemble protected works with ease.

If courts simply repeat that style is not protected, they may miss the risk of copying expressive combinations. If they treat every stylistic resemblance as infringement, they will expand rights too far. The safer approach is to translate style questions into questions of expressive combination. Courts should ask whether the similarity reproduces non-necessary, identifiable, and original combinations of details.

4.3. Black-box proof and imbalance in evidentiary access

A third difficulty concerns evidentiary imbalance. Training data, data-cleaning rules, model versions, and output logs are usually controlled by service providers. Right holders often see only the final output. Requiring them to prove that the training set contained the work, and that the model generated the output from that work, may amount to requiring impossible proof.

Responsibility should not be reversed automatically. A better approach is dynamic allocation of proof. The right holder first identifies protected expression and distinctive correspondences. The defendant then provides materials on training sources, filtering mechanisms, prompt

ISSN: 2688-8653

records, and model versions. Courts may use confidentiality review, expert assistants, and technical investigators. This approach can protect both legal remedies and technological development.

Dynamic proof must also respect proportionality. Courts should not require firms to disclose entire training sets or model parameters. Such disclosure may harm trade secrets and technical security. The scope of disclosure should be tied to the dispute. It should be limited to materials necessary for source, prompt, and reproducibility assessment.

Technical investigation cannot replace legal judgment. Experts may explain how a model generates content. They cannot decide what counts as protected expression. Expression scope, attribution, and market substitution remain legal questions. Judges must complete normative evaluation on the basis of technical facts.

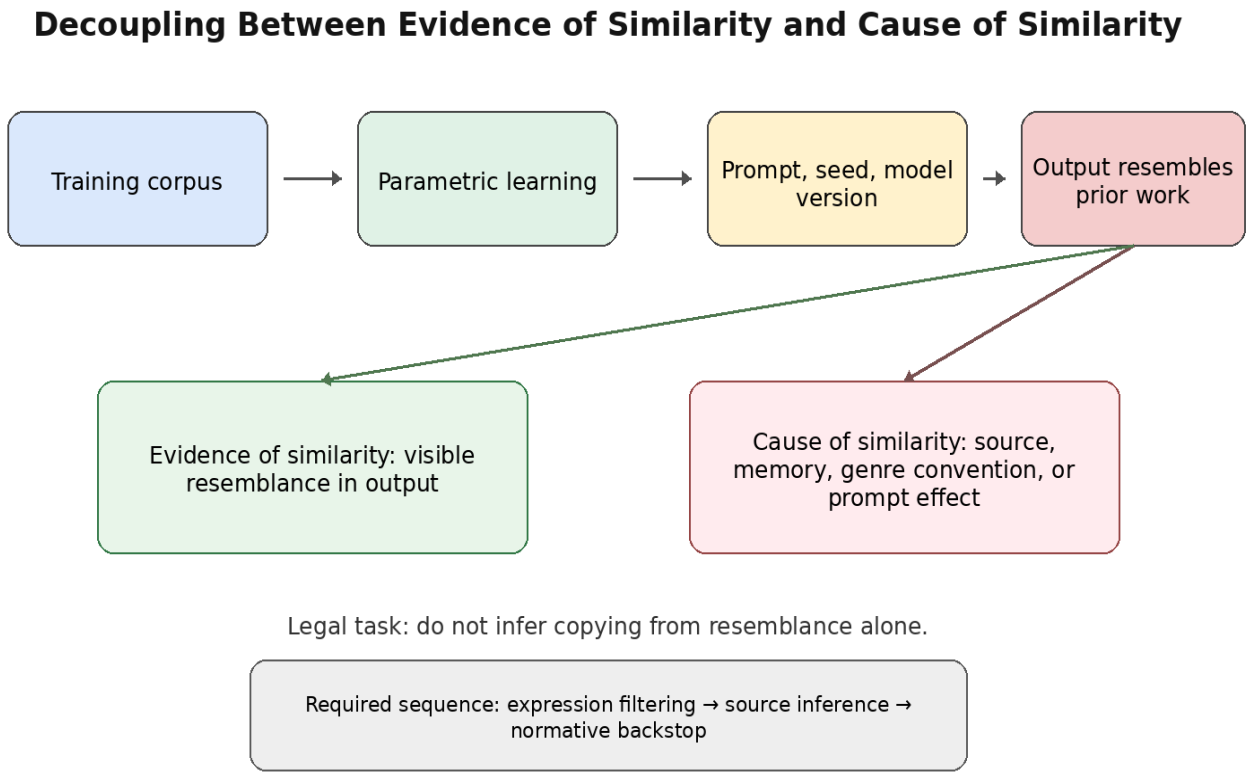


Figure 1. The evidentiary break in generative AI output similarity

5. Judicial Practice and Comparative Law Observations

5.1. Chinese cases: from copyrightability to platform responsibility

The Beijing Internet Court AI text-to-image case first confirmed the protectability of AI-generated content. The court included prompts, parameter settings, selection, and final choice in its analysis of intellectual contribution. The case mainly addressed work status and ownership. It also indicates that a claimant alleging infringement by an AI output must identify specific protected expression [7].

The significance of that case has limits. It did not directly address whether generated content infringed another work. It shows that AI outputs can be protected objects. It does not show that all AI outputs can be used freely. Protectability and infringement must remain separate.

The case therefore offers a double lesson. A claimant seeking protection must show human contribution to expression. A claimant alleging infringement must show that another output reproduced concrete expression from the claimant's work.

ISSN: 2688-8653

Public judicial reports and a reproduced judgment show that the Ultraman generative AI service case more directly touched output infringement. After users entered the name of a famous IP, the platform generated highly similar images. The available materials show that the court considered reproduction, adaptation, and the platform's duty of care. The case suggests that where a platform can foresee specific infringement triggers and has filtering capacity, its duty of reasonable care must be assessed [22][23].

That case also highlights the evidentiary value of prompts. When a prompt directly uses the name of a famous work, both the user and the platform can more easily foresee infringement risk. If the platform knows that such prompts repeatedly generate similar outputs, but still allows them, its duty of care should increase.

The Shanghai Animation Film Studio case was not an AI case, but it has methodological value. The accused short videos did not copy frame by frame. They used dialect replacement and rewriting to create a secondary-creation appearance. The court still examined core expression, image interests, and substitution risk. The same logic applies to AI outputs. Non-mechanical copying does not automatically exclude infringement [17].

5.2. U.S. and EU cases: fair use, parody, and market substitution

Campbell shows that similarity does not necessarily lead to liability. The court examined whether borrowing served criticism, comment, or a new expressive purpose. Market substitution was also central [24]. Warhol later cautioned that transformativeness should not be expanded without limit. Substitution in a commercial context still requires serious assessment [25]. Google Books, HathiTrust, and Google v. Oracle also show that technology-related copying cases usually turn on purpose, public benefit, market effect, and the nature of the copied material [26][27][28].

Deckmyn emphasizes the purpose of parody. A parody must evoke the prior work and show humor or mockery. The court must also balance the right holder's interests against freedom of expression. This standard raises challenges for AI outputs. An output may evoke a prior work without having a clear critical purpose [29].

Recent AI-related disputes further sharpen the point. The New York Times v. OpenAI and the Getty v. Stability AI proceedings raise output-oriented questions about text reproduction, image generation, watermark remnants, and market substitution [19][20][21]. Bartz v. Anthropic and Kadrey v. Meta address large-model training and fair use, while Thomson Reuters v. ROSS, though not a generative AI case, underscores the continuing importance of competitive substitution [30][31][32]. These cases are still developing and should not be read as settled doctrine. They nevertheless show that courts are now confronting training use, market effects, and output-stage attribution together. The U.S. Copyright Office report on generative AI training also treats training data, licensing markets, and market failure as independent policy issues [33].

5.3. Comparative summary: from threshold disputes to mechanism choice

Comparative law shows that the central issue is not merely the level of the similarity threshold. The U.S. path emphasizes fair use and market substitution. The EU path stresses the purpose of exceptions and interest balancing. Chinese cases focus more on work status, concrete expression, and platform duties. All three paths point to the same problem: courts must explain why similarity has attributional meaning.

China should not simply transplant any single model. A better approach is to strengthen evidence and reasoning within the existing infringement structure. Expression filtering controls the scope of protection. Source inference responds to black-box proof problems. A normative backstop prevents similarity analysis from losing contact with markets and public interests.

ISSN: 2688-8653

The value of comparative law lies in its problem awareness. U.S. law reminds us that similarity must be tested against fair use and market effect. EU law reminds us that exceptions require purpose and balancing. Chinese law should develop workable rules within the existing copyright framework.

This article therefore does not propose a separate category of AI infringement. AI disputes should return to the basic structure of copyright law. What must change is the method of proof and the order of reasoning.

Table 1. Comparison of representative domestic and foreign cases

Case type	Representative case	Main issue	Relevance to this article
AIGC copyrightability	Beijing Internet Court AI text-to-image case	Intellectual contribution, work status, ownership	Protection should return to human control and choice.
Output infringement	Ultraman generative AI service case	Famous IP prompts, reproduction/adaptation, platform duties	Output similarity may trigger responsibility.
Secondary creation	Shanghai Animation Film Studio v. Chongqing Yunmei	Character images, plot fragments, parodic or mocking use	Expression boundaries and substitution should be assessed together.
Parody/fair use	Campbell	Parody, transformativeness, market substitution	A critical purpose may justify necessary similarity.
EU parody exception	Deckmyn	Evocation of the prior work, humor/mockery, fair balance	Exceptions require visible purpose and interest balancing.
Training and licensing markets	Bartz / Kadrey / ROSS	Training purpose, data source, competitive substitution	Training conclusions do not automatically decide output liability.
Image platform liability	Getty v. Stability AI	Image training, output watermarks, cross-border liability	Output-stage claims require evidence and market analysis.

Note: The cases are selected to cover copyrightability, output infringement, secondary creation, parody or fair use, training-related disputes, and platform liability across Chinese, U.S., and EU contexts. They are used to illustrate different doctrinal functions rather than to suggest direct doctrinal equivalence.

6. Rule Reconstruction for Substantial Similarity in Generative AI Outputs

Rule reconstruction should not be understood as creating a separate AI infringement rule. A more realistic approach is to divide the existing substantial similarity inquiry into clearer reasoning steps. The analysis below asks three questions. Is the similar part protected expression? Can the source of the resemblance be inferred from evidence? Does the use deserve copyright attribution after considering purpose and market effect?

ISSN: 2688-8653

6.1. Expression filtering: preventing overprotection of style

The first step is expression filtering. Courts should identify the expression for which the plaintiff claims protection. Abstract subject matter, general style, and stock scenes should not enter the similarity comparison directly. The closer an AI output is to "style generation," the more careful filtering should be. Otherwise, right holders may use copyright to control shared aesthetic resources.

This step must avoid two extremes. One extreme excludes style, subject matter, and mood so broadly that concrete expressive combinations become unprotectable. The other treats overall impression as expression copying. Courts should provide more detailed reasoning between these poles.

Expression filtering should become a checklist-like form of reasoning. Courts should ask right holders to identify concrete corresponding elements. These may include character design, compositional arrangement, distinctive lines, non-necessary plot arrangements, or specific visual symbols. Courts should then decide whether each element is protected expression. Only elements that survive filtering should enter substantial similarity comparison.

Filtering should also account for AI-generation conventions. Default model composition, common poses, and prompt templates may cause frequent convergence. Such convergence should not automatically support attribution. Courts should distinguish model inertia, genre convention, and original expression. This distinction is essential to preventing overprotection.

6.2. Source inference: responding to black-box evidentiary difficulty

The second step is source inference. The AI black box makes direct proof difficult, but it does not force courts to choose between speculation and denial. Once the right holder identifies protected expression and distinctive correspondence, indirect evidence can play a role. Source inference does not lower the threshold of protection. It reallocates proof burdens when evidence is unequally accessible.

The object of source inference is not abstract training conduct. It concerns the connection between a specific output and a specific work. Even if the training data were generally lawful, a specific output may still reproduce protected expression. Conversely, even if the training set contained the work, a specific output may not infringe.

Source inference can examine five kinds of evidence. First, highly consistent distinctive details. Second, shared errors. Third, congruent non-necessary structural arrangements. Fourth, stable reproduction of similar outputs. Fifth, prompts, logs, or model descriptions pointing to the work. The more the evidence excludes accidental convergence, the stronger the inference becomes.

Defendants should have opportunities to rebut. Service providers may explain training sources, filtering rules, model versions, and output control measures. Users may explain prompt intent, later editing, and independent creative factors. Courts should also use expert assistance where necessary. Source inference should therefore be rebuttable rather than automatic.

The strength of inference depends on the evidence as a whole. A single similar detail usually cannot prove source. Multiple non-necessary details together increase inferential force. If shared errors and stable reproduction also appear, courts may require stronger explanations from the defendant.

Rebuttal materials should also be reviewed carefully. Service providers should not rely only on the phrase "model black box." At least reasonably available information should be disclosed, including the scope of training sources, rights-filtering measures, prompt-response mechanisms, and complaint-handling records.

ISSN: 2688-8653

Table 2. Evidence types and their evidentiary functions

Evidence type	Evidentiary function	Main risk	Stage of application
Distinctive details coincide	Shows that similarity may involve concrete expression	May confuse genre elements with original expression	Filtering / source inference
Shared errors	Supports a common-source inference	Errors may come from public sources	Source inference
Structural congruence	Shows similarity in non-necessary arrangement	May overvalue broad plot or composition	Filtering / source inference
Stable reproduction	Shows that similarity is not a one-off accident	Requires controlled prompt testing	Source inference
Prompt points to a work	Shows user intent or foreseeability	Prompt records may be incomplete	Source inference / attribution
Platform control evidence	Shows capacity to foresee, filter, or block	Over-disclosure may harm trade secrets	Platform duty
Market substitution evidence	Shows economic relevance of similarity	May overstate speculative markets	Normative backstop

Note: Some evidence types appear in more than one stage because the same fact may serve different doctrinal functions. For example, distinctive detail overlap may first help identify protected expression and then support source inference if accidental convergence can be excluded.

6.3. Normative backstop: calibrating liability through substitution risk

The third step is the normative backstop. Even when the causal chain cannot be fully reconstructed, courts should still evaluate the legal significance of resemblance. Relevant factors include use purpose, market substitution, licensing-market effects, and platform duties. The normative backstop is not an empty license for discretion. It returns similarity analysis to the purposes of copyright law.

The normative backstop should also handle licensing markets. Generative AI may not directly replace sales of the original work. It may instead replace authorization markets. Character licensing, image licensing, news-summary licensing, and style-based commercial commissions may all be affected. Courts should not look only at traditional sales markets.

Market effect should not expand without limit. Right holders cannot characterize all technological competition as market harm. Market factors should carry substantial weight only when an output appropriates the expressive value of the work or weakens a plausible licensing opportunity.

Commentary, criticism, teaching, and research uses should retain meaningful space. Campbell's transformativeness analysis is instructive. Yet an AI output that merely uses a work's attraction and enters the same market should not be justified too easily. Courts should ask whether the similarity was necessary, whether the purpose is clear, and whether the market was substituted [24].

The EU parody rule also offers guidance. Deckmyn requires a fair balance between rights protection and freedom of expression. In AI cases, courts may draw on this reasoning. If an output merely re-skins a protected character and lacks a critical purpose, its justification is

ISSN: 2688-8653

weak. If the output is used for criticism or research, normative evaluation should be more cautious [29].

Platform responsibility should also remain part of the analysis. China's Interim Measures for Generative AI Services and Provisions on Deep Synthesis already impose governance duties [34][35]. The EU General-Purpose AI Code of Practice also treats transparency and copyright compliance as important duties for model providers [36]. In copyright cases, courts should consider log retention, complaint handling, keyword filtering, and blocking of repeated infringement. These facts help determine whether the service provider exercised reasonable care.

6.4. Procedural safeguards: supporting substantive judgment through transparency

Substantive rules need procedural support. Without appropriate transparency, expression filtering and source inference may fail. Courts should consider layered disclosure. Public disclosure may cover model version, output time, and prompts. Confidential disclosure may cover training sources, filtering rules, and similarity reproduction tests. Materials involving trade secrets can be protected through confidentiality orders or closed review.

Technical assistance is also important. Expert assistants may explain model generation mechanisms. Technical investigators may help review logs and reproduction tests. Where necessary, courts may require parties to submit comparison charts. Such charts should list similar elements, reasons for protection, and reasons for exclusion. This can move judicial reasoning from intuition to testable analysis.

Procedural arrangements should also balance the interests of both sides. Right holders need necessary information. Service providers need protection against disclosure of core technology. Confidentiality orders, limited access, and expert undertakings can reduce this conflict.

The final purpose of procedure remains substantive judgment. Only with sufficient factual basis can the three-step framework operate. Otherwise, expression filtering becomes abstract, source inference becomes guesswork, and the normative backstop loses constraint.

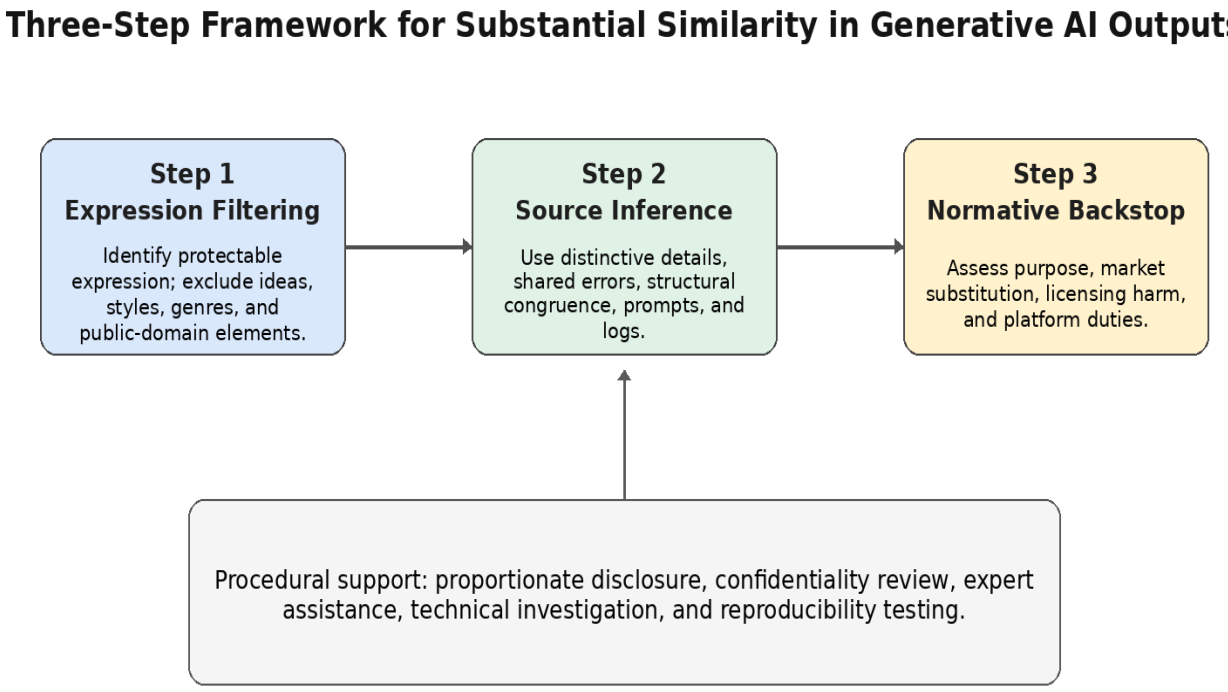


Figure 2. The three-step framework of expression filtering, source inference, and normative backstop

ISSN: 2688-8653

7. Conclusion

Generative AI does not eliminate substantial similarity. It changes the evidentiary conditions of the doctrine. Traditional cases ask mainly whether two works are alike. AI cases must also ask why they are alike and whether the resemblance should be attributed. If courts look only at output appearance, copyright boundaries may expand too far. If courts always require right holders to prove model extraction, protection may fail.

The three-step framework responds to this tension. Expression filtering first controls the boundary of protection. Source inference then addresses proof difficulties under black-box conditions. The normative backstop brings market substitution, use purpose, and public interest back into the attribution inquiry.

This framework is not a mechanical formula. It asks courts to consider protection scope first, source connection second, and normative attribution last. This order can reduce result-driven reasoning. It also allows both parties to organize evidence around specific disputes.

Further refinement will occur in individual cases. Algorithmic transparency, platform logs, expert assistance, and confidentiality review will all affect outcomes. Generative AI will continue to change content production. The substantial similarity doctrine should remain open and adaptable. Only then can copyright law maintain a workable balance between protecting creation and supporting innovation.

References

[1] Standing Committee of the National People's Congress. (2020). Copyright Law of the People's Republic of China.

[2] Wang, Q. (2024). Reconsidering the status of AI-generated content in copyright law: A third discussion. Studies in Law and Business, (3). (in Chinese)

[3] Wu, H. (2020). Questions of copyright law raised by AI-generated works. Peking University Law Journal, (3). (in Chinese)

[4] Xiong, Q. (2017). Copyright recognition of artificial intelligence generated content. Intellectual Property, (3). (in Chinese)

[5] U.S. Copyright Office. (2025). Copyright and artificial intelligence, Part 2: Copyrightability. U.S. Copyright Office.

[6] Kaiser, Z. B., Song, K., & Frankel, S. J. (2025). Reconciling copyright originality for photography and generative artificial intelligence. Columbia Journal of Law & the Arts, 49(2), 339.

[7] Beijing Internet Court. (2023). Li v. Liu, dispute over the right of authorship and the right of communication through information networks, Civil Judgment No. (2023) Jing 0491 Min Chu 11279. (in Chinese)

[8] Samuelson, P. (2023). Generative AI meets copyright. Science, 381(6654), 158–161. https://doi.org/10.1126/science.adi0656

[9] Lemley, M. A. (2024). How generative AI turns copyright law upside down. Columbia Science and Technology Law Review, 25, 21. https://doi.org/10.52214/stlr.v25i2.12761

[10] Sag, M. (2023). Copyright safety for generative AI. Houston Law Review, 61, 295.

[11] Sobel, B. L. W. (2024). Elements of style: Copyright, similarity, and generative AI. Harvard Journal of Law & Technology, 38(1), 49–107.

[12] Hacohen, U. Y., & Elkin-Koren, N. (2024). Copyright regenerated: Harnessing GenAI to measure originality and copyright scope. Harvard Journal of Law & Technology, 37(2), 555–610.

[13] Rosati, E. (2025). Infringing AI: Liability for AI-generated outputs under international, EU, and UK copyright law. European Journal of Risk Regulation, 16(2), 603–627. https://doi.org/10.1017/err.2024.72

ISSN: 2688-8653

[14] Li, K. G., Wu, H., & Dong, Y. P. (2024). Copyright protection during the training stage of generative AI: Industry-oriented U.S. law, rights-oriented EU law, and fair remuneration rights for generative AI training under the UN's international governance regime for AI. Computer Law & Security Review, 55, 106056. https://doi.org/10.1016/j.clsr.2024.106056

[15] Zhuk, A. (2024). Navigating the legal landscape of AI copyright: A comparative analysis of EU, US and Chinese approaches. AI and Ethics. https://doi.org/10.1007/s43681-023-00299-0

[16] Supreme People's Court of China. (n.d.). Interpretation on several issues concerning the application of law in the trial of copyright civil dispute cases. (in Chinese)

[17] Chongqing No. 5 Intermediate People's Court. (2019). Shanghai Animation Film Studio Co., Ltd. v. Chongqing Yunmei Information Technology Co., Ltd., Civil Judgment No. (2019) Yu 05 Min Chu 3828. (in Chinese)

[18] Supreme People's Court of China. (2023, November 3). Typical cases on film intellectual property protection: Shanghai Animation Film Studio Co., Ltd. v. Chongqing Yunmei Information Technology Co., Ltd. et al., copyright infringement dispute. https://www.court.gov.cn/zixun/xiangqing/416612.html (in Chinese)

[19] The New York Times Co. v. Microsoft Corp. and OpenAI, Inc., No. 1:23-cv-11195 (S.D.N.Y. 2023).

[20] Getty Images (US), Inc. v. Stability AI, Inc., No. 1:23-cv-00135 (D. Del. 2023).

[21] Getty Images (US), Inc. and Others v. Stability AI Ltd., [2025] EWHC 2863 (Ch).

[22] Intellectual Property Court of the Supreme People's Court of China. (2025, July 18). Intellectual property adjudication faces new challenges in the AI era. https://ipc.court.gov.cn/zh-cn/news/view-4513.html (in Chinese)

[23] China Intellectual Property Lawyers Network. (2024, February 28). Reproduced first-instance civil judgment in the first global AIGC platform infringement case: No. (2024) Yue 0192 Min Chu 113. https://www.ciplawyer.cn/articles/152937.html (reproduced from a civil judgment of Guangzhou Internet Court; in Chinese)

[24] Campbell v. Acuff-Rose Music, Inc., 510 U.S. 569 (1994).

[25] Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith, 598 U.S. 508 (2023).

[26] Authors Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015).

[27] Authors Guild, Inc. v. HathiTrust, 755 F.3d 87 (2d Cir. 2014).

[28] Google LLC v. Oracle America, Inc., 593 U.S. 1 (2021).

[29] Deckmyn and Vrijheidsfonds VZW v. Vandersteen and Others, Case C-201/13, ECLI:EU:C:2014:2132.

[30] Bartz v. Anthropic PBC, No. 3:24-cv-05417-WHA (N.D. Cal. 2025, June 23). Order on fair use.

[31] Kadrey v. Meta Platforms, Inc., No. 3:23-cv-03417-VC (N.D. Cal. 2025).

[32] Thomson Reuters Enterprise Centre GmbH v. ROSS Intelligence Inc., No. 1:20-cv-613 (D. Del. 2025, February 11).

[33] U.S. Copyright Office. (2025). Copyright and artificial intelligence, Part 3: Generative AI training. U.S. Copyright Office.

[34] Cyberspace Administration of China et al. (2023, July 10). Interim measures for generative artificial intelligence services. (in Chinese)

[35] Cyberspace Administration of China et al. (2022, November 25). Provisions on the administration of deep synthesis internet information services. (in Chinese)

[36] European Commission. (2025). The general-purpose AI code of practice.