Generative AI–Driven Rumor Propagation Chains and the Dilemma of Legal Governance

Wenjun Wu

School of Advanced Institute of Journalism and Society, Shanghai University of Finance and Economics, Shanghai 200082, China

karry3085651896@qq.com

Abstract. Under the influence of Generative Artificial Intelligence (Generative AI), the mechanisms and chains of rumor generation and dissemination have undergone significant transformation. Generative AI not only reduces the cost of rumor fabrication but also enhances the credibility and deceptiveness of false information. Diverse motivational drivers, precision-targeted dissemination, audience psychological mechanisms, and algorithmic recommendation systems collectively accelerate the spread and amplification of AI-generated rumors. However, current legal governance faces multiple dilemmas: the rapid evolution of AI technology has outpaced legislative progress; evidentiary identification is complicated by the technical opacity of AI systems; and cross-border coordination remains fragmented. To address these challenges, this paper argues that: (1) at the technical level, regulatory mechanisms should be developed to strengthen AI content detection, traceability, and authenticity verification; (2) at the societal level, public empowerment should be enhanced through media literacy education to block rumor dissemination; and (3) at the legal level, accountability mechanisms should be refined to optimize the governance of AI-generated rumors.

Keywords: Generative Artificial Intelligence, Rumor Propagation, Legal Governance, Algorithmic Recommendation Mechanisms.

1. Introduction

In recent years, various forms of generative artificial intelligence have been increasingly applied to content creation scenarios. The emergence of DeepSeek has markedly elevated public awareness of Generative AI in China, intertwining it with everyday life. Generative AI demonstrates distinct advantages in divergent thinking, content generation, and efficiency enhancement. However, due to its low technological threshold and intrinsic characteristics—such as content mimicry, data-driven optimization, and intelligent packaging—it has progressively become a major vehicle for the creation and dissemination of online rumors.

"Rumor is the world's oldest medium of communication; it emerges wherever humans crave connection, information, and understanding." Typically originating from unofficial channels and built around specific events, rumors propagate primarily through continuous and rapid interpersonal transmission, with online media serving as the dominant carrier. On one hand, Generative AI autonomously produces false information that is subsequently disseminated by human users; on the other, some users, driven by motives such as social disruption or economic gain, deliberately employ AIGC (Artificial Intelligence Generated Content) tools to fabricate and amplify rumors across social media platforms.

The proliferation of AI-generated rumors has been especially prominent in recent years. For instance, during the February 2025 Yibin landslide incident in Sichuan Province, some online users, seeking attention and traffic revenue, used Generative AI to falsify the time, location, and even the details of the rescue process. Similar falsehoods included "A top-tier celebrity lost 1 billion in gambling at Macau" and "A boy wearing a hat was found under rubble in the 2025 Rikaze earthquake." These rumors—partly grounded in real events but augmented by fictionalized elements—exemplify the hybrid realism of AI-generated misinformation. The perpetrators were not necessarily AI experts; rather, technological advancement has significantly lowered the threshold and cost of rumor fabrication. As such, Generative AI–driven rumors circulating on social media platforms complicate public discernment between truth and falsehood and exacerbate challenges in legal governance.

To effectively regulate the generation and spread of AI-driven rumors, this paper systematically analyzes (1) the mechanisms of rumor generation and dissemination under Generative AI, (2) the corresponding dilemmas in legal governance, and (3) proposes actionable countermeasures and governance pathways to mitigate their impact.

2. Mechanism Analysis of Rumor Generation and Dissemination under Generative AI

2.1. Rumor Generation Mechanism Driven by Generative AI

2.1.1. Technological Empowerment: Efficiency Enhancement and Lowered Barriers

Technology serves as the fundamental driving force behind Generative AI, effectively reducing the threshold for fabricating false information while substantially improving the efficiency of rumor production. Traditionally, fabricating online rumors required skills such as writing, image editing, or video compositing. The advent of Generative AI has removed these technical barriers: with powerful logical reasoning and realistic simulation capabilities, users can now generate coherent text or authentic-looking visuals by merely inputting simple prompts. This democratization of content creation lowers the difficulty of producing false information and expands the pool of potential rumor creators.

Moreover, Generative AI possesses scalability and automation, enabling "assembly-line" or even "autonomous" rumor production. This mechanized process reduces the time and financial costs associated with rumor generation, making false information production faster and more frequent than ever before.

2.1.2. Content Camouflage: Multimodal Packaging and Credibility Construction

Beyond its efficiency, Generative AI also enhances the deceptive realism of rumors through multimodal integration and narrative coherence. Large language models (LLMs) transform raw data into human-like discourse, constructing narratives that seamlessly blend plausibility with persuasion. The transformation of content from "generable" to "believable" encompasses three layers: linguistic realism, visual mimicry, and evidentiary fabrication.

First, leveraging the strong linguistic modeling capacity of LLMs, AI-generated texts exhibit structural and stylistic conformity with real journalistic or conversational patterns. For example, in fabricating a false news report, Generative AI can automatically employ the inverted pyramid structure typical of journalism, embedding expressions like "according to eyewitnesses" or "official data show."

Second, with multimodal generation capabilities, Generative AI can produce highly realistic visuals, deepfakes, and composite videos. Compared to textual content, visual misinformation appeals to emotion and perception—its immediacy and vividness enhance memorability and credibility. Through techniques like face-swapping or synthetic video composition, AI exploits humans' cognitive bias of "seeing is believing," thus reinforcing the illusion of authenticity.

Third, Generative AI can fabricate evidence and falsify sources. By inventing authoritative citations, fabricating statistical data, or associating false claims with reputable organizations, it further consolidates credibility. These tactics render the boundary between authenticity and falsity increasingly obscure, making rumor detection and verification more difficult.

2.1.3. Data Bias: Model Deviation and Implicit Manipulation

The outputs of Generative AI are inherently shaped by the data on which models are trained, embedding latent biases and ideological inclinations that are difficult for users to control. Because of the black-box nature of large models, false or misleading information can emerge as a byproduct of opaque data patterns and algorithmic biases.

First, data bias in training corpora may lead to distortion. The accuracy of Generative AI depends heavily on the quality and neutrality of training data. If biased or inaccurate representations exist—

such as regional stereotypes or historical misinformation—AI systems may unconsciously reproduce or amplify them. For instance, if training data contain prejudicial portrayals of a particular region, subsequent generations about that region may unjustifiably emphasize negative aspects, contributing to rumor formation.

Second, algorithmic manipulation can occur when developers or interest groups intentionally inject subjective intent into model parameters or data curation. By tuning generation rules or input filters, these actors can steer the model toward outputs that serve specific agendas, thereby weaponizing Generative AI as a tool for misinformation dissemination.

2.1.4. Intelligent Hallucination: Fictionalization and Fact Splicing

Generative AI models generate outputs through probabilistic prediction mechanisms—they fill informational gaps based on learned patterns rather than verified facts. When prompts exceed the model's knowledge boundary or lack factual grounding, the system may produce hallucinations—syntactically correct, semantically coherent, yet factually false content.

These hallucinations often display sophisticated logical construction and detailed narrative elaboration, creating a veneer of authenticity. Beyond pure fabrication, Generative AI frequently blends falsehoods with verified facts—a process known as fact splicing. For instance, an AI-generated article might correctly cite real data but misinterpret or selectively quote it to craft misleading conclusions. Similarly, a post might build upon a true event but introduce invented subplots to amplify sensationalism.

Such hybridized misinformation blurs the line between reality and fabrication, exploiting human representativeness heuristics—our tendency to generalize credibility from partially accurate cues [10]. Consequently, audiences may unconsciously accept the overall narrative as true while overlooking its false components.

2.2. Rumor Dissemination Chain Driven by Generative AI

2.2.1. Diverse Dissemination Motives: Profit, Manipulation, and Emotion

Generative AI–driven rumor propagation is not merely a technical phenomenon but also a behavioral one, intertwined with human motives that can be categorized into three types: profit-driven, socio-political manipulation, and psychological gratification.

(1) Profit-driven motives exhibit the highest intentionality. Individuals or organizations disseminate AI-generated rumors to obtain material benefits or social influence. For example, false financial news generated by AI—such as "Company X falsified its earnings"—can manipulate stock prices or harm competitors. Similarly, during political elections, malicious actors may fabricate AI-generated scandals to damage opponents' reputations.

(2) Social manipulation motives aim to influence public sentiment. Some groups exploit AI to generate false imagery or reports—such as fabricated images of contaminated water supplies—to provoke fear, anger, or collective anxiety.

(3) Personal psychological motives are often rooted in emotional expression, curiosity, or self-display. These individualized rumors are typically fragmented and emotionally charged, circulating in small-scale networks such as WeChat friend groups. For instance, AI may be used to fabricate celebrity gossip or internal workplace rumors, fulfilling users' emotional or social needs.

2.2.2. Precision-Targeting Mechanism: Audience Profiling and Content Customization

Precision targeting bridges the gap between rumor generation and amplification. By integrating Generative AI with big data analytics, rumor dissemination can be tailored to specific audience groups with remarkable precision.

AI systems can crawl user behavior across platforms—likes, shares, comments—and synthesize multi-dimensional user profiles. For instance, investors categorized as "risk-averse" might be targeted with AI-generated rumors like "A listed company is concealing losses." Furthermore, AI adapts rumor content to platform-specific dissemination logic: the same rumor can be distributed in varied

formats (e.g., short video, microblog, long text) while real-time engagement data continually refine the narrative for maximum virality.

2.2.3. Audience Psychological Mechanisms: Cognitive Bias and Group Dynamics

Audience psychology constitutes a critical link in the rumor propagation chain under Generative AI. Through content realism and optimized delivery, AI systems strategically exploit and amplify human cognitive biases.

First, automation bias leads individuals to overtrust AI outputs. Due to AI's perceived authority and technical sophistication, users may accept generated information uncritically. The hyperrealistic style of AI-generated content reinforces this trust, causing users to share misinformation without verification.

Second, emotional contagion amplifies dissemination speed. Generative AI can mine vast datasets to detect and reproduce emotionally charged topics. Emotionally resonant rumors—particularly those evoking anxiety, outrage, or empathy—spread rapidly, as recipients prioritize emotional response over factual verification. Social validation often outweighs epistemic accuracy: users tend to trust information shared by peers rather than investigating its authenticity.

Third, group conformity effects intensify through algorithmic amplification. Generative AI can mass-produce semantically similar rumors within a short time, creating an illusion of consensus in online discourse. Algorithmic recommendation systems then reinforce this pseudo-majority effect by continuously promoting similar content, leading to deindividuation—the loss of personal judgment within crowd behavior [14]. Consequently, users follow group cues, further accelerating rumor diffusion.

2.2.4. Algorithmic Recommendation Mechanisms: Information Bubbles and Traffic Incentives

Social media algorithms play a pivotal role in scaling rumor dissemination from isolated cases to mass phenomena. Their profit-driven and engagement-optimized nature inherently aligns with the virality of sensational or emotionally charged misinformation.

On one hand, collaborative filtering algorithms personalize content feeds by analyzing user preferences and interactions. This leads to the formation of information cocoons or echo chambers, where users are repeatedly exposed to homogeneous content, reinforcing confirmation bias and emotional polarization [15]. Once rumors penetrate such clusters, algorithmic reinforcement fuels continuous circulation and internal amplification, while corrective or fact-checking information struggles to reach these insulated audiences.

On the other hand, platform algorithms have effectively replaced traditional media's "agenda-setting" and "gatekeeping" functions. Unlike editorial curation grounded in news value or public interest, algorithmic systems prioritize engagement metrics—clicks, likes, and comments—for commercial gain. AI-generated rumors, characterized by novelty and emotional appeal, naturally excel under this logic. Consequently, the algorithms promote high-performing misinformation to top rankings, increasing both visibility and dissemination velocity, and thus perpetuating the rumor cycle.

3. Legal Governance Mechanisms and Dilemmas under Generative AI

3.1. Legal Governance Mechanisms for Generative AI–Driven Rumors

Generative Artificial Intelligence (Generative AI) possesses the characteristics of high efficiency, low cost, and strong realism, which have greatly facilitated the field of information dissemination. However, these same attributes have also introduced new challenges for the generation and diffusion of rumors. To address these issues, China has gradually promulgated a series of corresponding laws and regulations.

3.1.1. Establishment of the Legal Governance Framework

The Cybersecurity Law, Data Security Law, and Personal Information Protection Law constitute the three fundamental laws of the digital era. From the perspectives of cyberspace governance, data security protection, and personal information safeguarding, these laws lay the groundwork for the governance of false information. Each of them has its own focus: Article 47 of the Cybersecurity Law explicitly prohibits individuals from producing, publishing, or disseminating false information and requires network operators to fulfill their responsibilities for content management and to take necessary measures to handle illegal information. This provision provides the upper-level legal basis for subsequent governance of AI-generated rumors. The Data Security Law focuses on data processing activities and requires data processors to ensure the authenticity, accuracy, integrity, and availability of data, which helps regulate data processing behavior and reduce the risk of rumor generation by Generative AI at the source. The Personal Information Protection Law, by restricting the misuse of personal data, indirectly curbs the behavior of fabricating personal-related rumors through AI technologies.

At present, among the normative documents promulgated in China, the Interim Measures for the Administration of Generative Artificial Intelligence Services and the Measures for the Administration of Synthetic Content Identification (2025) specifically regulate the characteristics of Generative AI, compensating for the technical adaptability deficiencies of the fundamental laws. The Interim Measures for the Administration of Generative Artificial Intelligence Services clarify the content boundaries of Generative AI and explicitly prohibit the generation of false or harmful information. The Measures for the Administration of Synthetic Content Identification require service providers to improve the accuracy and traceability of AI-generated content and to enhance transparency in the use of Generative AI. Meanwhile, the Measures also stipulate content identification obligations, requiring that AI-generated content include explicit or implicit identifiers to facilitate public recognition.

3.1.2. Administrative Punishment and Civil Remedies in Parallel

Under the current system, for rumors produced by Generative AI, two main governance paths exist: administrative punishment and civil remedies. The former relies on public authority to punish illegal conduct and protect public interests, while the latter emphasizes private relief through individual rights protection to compensate for damages.

According to the Cybersecurity Law and the Administrative Measures on Internet Information Services, national cyberspace authorities may impose penalties such as fines, business suspension for rectification, or license revocation on platforms or responsible individuals who disseminate false information. Victims of Generative AI–generated rumors, on the other hand, may file lawsuits under the Civil Code for defamation or property damage, demanding that the infringer bear civil liabilities such as cessation of infringement, apology, and compensation for losses.

3.1.3. Supplementation through Industry Self-Regulation and Ethical Norms

Industry self-regulation and ethical norms are also indispensable. By leveraging enterprises' technical capabilities and industry coordination, rumor risks can be proactively prevented and mitigated across multiple stages.

Some enterprises and platforms have established self-regulatory mechanisms through user agreements, content review systems, and "blacklist" management. For example, user agreements may require warning labels when generating content with AI models to prevent malicious dissemination; for violators, platforms may impose comprehensive restrictions, cutting off the channels for repeated rumor generation.

3.2. Dilemmas of Legal Governance under Generative AI–Driven Rumors

3.2.1. The Contradiction between Legislative Lag and Rapid Technological Development

Existing laws were primarily designed to address traditional modes of information dissemination and therefore lack clear provisions regarding the allocation of liability, evidentiary identification, and cross-border dissemination specific to AI-generated rumors. Legislation, by nature, requires rigorous procedures and values stability, with long drafting and amendment cycles. In contrast, AI technology evolves rapidly, with swiftly changing applications and potential risks. Consequently, the law struggles to respond promptly and effectively to regulate fast-evolving technologies. In practice, the absence of detailed rules often leads to interpretive uncertainty and operational difficulties.

By comparison, administrative regulations are issued more quickly and can provide preliminary guidance for industry regulation. However, due to differences between departments, the specific rules of various regulations often overlap, and the absence of an integrated system makes their application fragmented.

3.2.2. The Contradiction between Rules of Evidence and Technical Characteristics of AI

Rumors generated by Generative AI have reached a high level of realism, capable of fabricating videos, audios, and synthetic texts, making it difficult for the public to distinguish between true and false information. This not only deepens the social trust crisis but also poses challenges for evidence identification in judicial practice.

First, source tracing is difficult. Generative AI models rely on massive datasets for training, encompassing open internet resources and user-generated content, which makes it hard to trace original materials quickly. Meanwhile, the model parameters and decision mechanisms are complex and nontransparent, making it difficult to identify the origin, modification path, and dissemination nodes of AI-generated content. Furthermore, the highly realistic nature of AI-generated materials simulating real scenarios increases the difficulty of authentication. To determine the source of a rumor, forensic institutions must invest substantial resources and conduct complex operations such as model feature extraction and training data comparison.

Second, the burden of proof is heavy. Legal proceedings involving AI-generated rumors are primarily civil in nature, governed by the principle that “whoever claims must provide evidence.” In the digital age, traditional evidentiary rules no longer adapt well to technological change. Victims, often lacking technical expertise and resources, face an onerous burden of proof. They must not only demonstrate that the content is false but also that it was AI-generated and traceable to a specific tool or model. This process requires specialized knowledge such as model principles and digital authentication techniques. Most victims lack such expertise and must rely on experts, resulting in high verification costs and financial burdens.

Third, attribution of liability is difficult. The generation and dissemination of rumors driven by Generative AI involve a complex chain with multiple participants, including users who input prompts and post rumors, technology companies that provide AI services, and information platforms that facilitate dissemination. Judicial decisions struggle to clearly define each party’s fault and proportion of liability. Determining the user’s subjective intent is challenging: users may induce AI to generate false content through vague prompts and later claim ignorance. Additionally, due to rapid iterations of AI technologies and the concealment of rumor forms, traditional detection methods used by platforms have limited effectiveness. Excessive platform moderation may infringe on freedom of expression, while insufficient moderation may be manipulated by algorithmic mechanisms, indirectly fueling rumor dissemination. Current laws lack detailed standards regarding platform review responsibilities and their scope. Moreover, existing laws only regulate two types of actors—users and platforms—while insufficiently addressing model providers. It remains unclear whether and under what conditions model providers qualify as “network service providers” and what level of duty of care and review they should bear. Given that model developers can directly provide end services and that actors along the Generative AI value chain may overlap or shift roles, the dependencies among

model development, deployment, and application increase the transmissibility of risks, dispersing causality between human conduct and resulting harm.

3.2.3. The Contradiction between Cross-Border Cooperation Needs and Jurisdictional Sovereignty Barriers

Rumors generated by Generative AI can deeply embed themselves within foreign cultural and linguistic contexts, enabling transnational dissemination and global diffusion. However, the international community currently lacks a unified cooperative mechanism for rumor governance, and cross-border regulation faces multiple obstacles.

First, due to differences in cultural traditions and value systems, countries vary significantly in their legal definitions of rumors, regulatory scope, and liability attribution. Divergent legislative philosophies create a legitimacy dilemma in cross-border rumor governance, making it difficult for international platforms to formulate uniform content review standards and providing opportunities for rumor disseminators to exploit legal loopholes. Second, cross-border law enforcement cooperation is inefficient. On the one hand, judicial assistance often involves complicated diplomatic procedures that consume time; on the other, cross-border data retrieval may be blocked for violating local laws or privacy rules, breaking the evidentiary chain and hindering source tracing. Finally, the global diffusion of Generative AI may further reinforce the information power imbalance between developed and developing nations. Due to limited technical reserves and defensive capacity, developing countries are often placed in a passive position when dealing with AI-driven cross-border rumors, lacking the ability to effectively identify and intercept false information.

4. Countermeasures and Legal Pathways for Rumor Dissemination and Governance under Generative AI

4.1. Interruption of Rumor Dissemination Chains Driven by Generative AI

4.1.1. Technological Regulation

Technological development is a double-edged sword: while it can be misused by unlawful actors to inflict social harm, it can also continuously innovate to benefit society and reduce potential risks. China should employ technological means to promote the research, development, and application of AI-content recognition and traceability technologies, building a multi-tiered and three-dimensional technical-governance system.

First, the state should encourage scientific-research institutions and enterprises to develop high-precision detection tools for AI-generated content. Using watermarking, metadata labeling, and model-fingerprint techniques, platforms can automatically identify and mark AI-generated information to classify credibility levels; graded labeling of data-source reliability enhances the matching accuracy of high-quality algorithms.

Second, rumor-traceability and tracking mechanisms must be improved. Research institutes can establish cross-platform content-tracking systems that record the release, modification, and dissemination processes of AI-generated materials via blockchain or similar technologies, enabling rapid source tracing and facilitating accountability for primary offenders.

Third, platform algorithms should be continuously optimized. Enterprises may assign specialized personnel for information verification and build dynamic rumor databases to perform real-time comparison and interception. Moreover, platforms can implement traffic-control mechanisms to reduce the exposure and spread of rumor information.

4.1.2. Public Empowerment

Information governance is inherently human-centered, and the interruption of rumor dissemination driven by Generative AI relies on the active participation of the public.

High media literacy is the core weapon for the era of mixed true and false information. Government departments and public-interest organizations should conduct science-popularization campaigns on

identifying AI-generated rumors to enhance the public's capacity to discern false information; they should also organize regular legal-awareness programs to raise cognitive literacy and legal consciousness. Relevant agencies can employ posters, videos, and other multimedia forms to disseminate fundamental knowledge of algorithmic mechanisms, helping citizens understand the role and control logic of algorithms in social-relationship construction, thereby reinforcing autonomy and resisting hidden algorithmic manipulation.

In addition, to fully mobilize social participation in rumor governance, platforms should strengthen reporting mechanisms and introduce incentive systems that encourage users to engage in factual verification. For instance, platforms may create a dedicated "AI Rumor" reporting portal to motivate users to identify suspected AI-generated misinformation, rewarding effective reports with points or traffic bonuses. They can also establish a content-credibility evaluation system, adding a scoring section beneath each post to guide user assessments. Platforms should regularly aggregate these scores, adjust algorithmic biases, and prioritize the display of high-credibility content.

4.2. Legal Governance Pathways for Generative AI–Driven Rumors

4.2.1. Establishment of a Scientific Recognition and Liability Determination System

A scientific and reasonable approach to liability recognition should be established to address the difficulties of evidence identification and the ambiguity of responsibility boundaries in the governance of rumors driven by Generative AI. The law should clarify the scope of responsibility and the standards for its identification in different circumstances.

First, the liability of users should be clarified. If users enter prompts such as "fabricate," "forge," or similar instructions to induce Generative AI to produce false information, or maliciously edit and spread such content, they should be deemed to have acted with intent and bear corresponding legal responsibility. For cases where users employ obscure prompts to guide AI in fabricating rumors, relevant authorities should trace users' prompt history and analyze the correlation between prompts and generated content to determine subjective intent. When a high degree of correlation exists, users should still be held legally responsible based on the outcome, even if the prompts appear ambiguous.

Second, platforms and AI model developers must fulfill their duty of care; if content review is inadequate, they should bear negligence liability. Examples include failure to deploy AI-content detection tools, failure to manually review high-risk content, or failure to set rumor-filter thresholds within algorithmic recommendation mechanisms. Model developers should also bear negligence liability for technical safety vulnerabilities, such as failure to filter biased training data or neglect to block high-risk prompts.

Finally, competent authorities should determine penalty levels according to the extent of rumor dissemination and its social impact. Government agencies can establish dual-dimensional assessment standards that quantify harm and update data dynamically. By combining the speed of AI-generated content dissemination and the scale of derivative damage, the precision and fairness of legal application can be enhanced.

4.2.2. Optimization of Related Governance Pathways

To address problems such as legislative lag and inadequate cross-border cooperation, breakthroughs can be achieved from three perspectives: legislative improvement, multi-stakeholder collaboration, and international cooperation.

First, relevant laws and regulations should be continuously refined to systematically define the allocation of responsibility for AI-generated rumors. On the one hand, based on existing frameworks such as the Cybersecurity Law and the Interim Measures for the Administration of Generative Artificial Intelligence Services, the state should formulate specialized judicial interpretations or departmental regulations to provide detailed guidance on responsibility determination and the handling of cross-border rumor dissemination. On the other hand, the law should explicitly define the review obligations and scope of model providers within the Generative AI rumor dissemination chain to resolve the ambiguity of actor identity and dispersed responsibility. Legislation should adhere to

the principles of technological neutrality and risk prevention, encouraging innovation while preventing the misuse of AI.

Second, multi-stakeholder collaborative governance should be promoted by combining government supervision, industry self-regulation, and public participation. Government departments should strengthen supervision and law enforcement and regularly launch special rectification campaigns. Industry organizations should establish AI-ethics guidelines and content-generation standards to promote self-regulation. The public should enhance media literacy, actively report illegal activities, and play an active role in social oversight.

Third, international cooperation and rule alignment should be strengthened. In response to the cross-border dissemination of AI-generated rumors, countries should actively promote the formulation of international rules and establish rumor information–monitoring and joint law-enforcement mechanisms. They should also unify standards for rumor definition to reduce gaps arising from legal discrepancies, improve cooperation efficiency, and simplify judicial procedures. Nations should promote international exchanges and capacity building—countries with advanced AI-recognition technologies should assist those with limited capabilities in rumor detection and verification—thereby achieving collective progress, safeguarding international security, and advancing global peace and development.

5. Conclusion

Generative Artificial Intelligence can enhance the efficiency of content production but also reshape the ecology of rumor dissemination. The contradictions between legislative lag and technological evolution, evidentiary rules and AI characteristics, and cross-border coordination and sovereignty barriers have become apparent, showing that the current legal governance system requires timely adjustment. The governance of Generative AI–driven rumors can be strengthened from three aspects: technological regulation, public empowerment, and legal refinement, to effectively manage emerging risks.

In the future, as Generative AI continues to advance, China should uphold a dynamic governance mindset, consolidate the information-security defense line of the AI era, and maintain a clean and orderly cyberspace.

References

[1] Wang, P. (2024). Research on the Production, Dissemination, and Governance of Generative Artificial Intelligence Rumors. New Media and Society, (04), 90–102 + 416.

[2] Malingcat. (2009). Rumor and Rationality: Reading “Rumor: The World’s Oldest Medium.” Publishing Horizon, (04), 65.

[3] Zhao, Y., & Shi, M. X. (2024). Who Is Making the News? The Challenges of Virtual Robots to the Public Opinion Ecology and Its Governance. China Publishing, (20), 3–9.

[4] Zhang, W. X. (2025). The Public Opinion Ecological Challenges and Governance Pathways of Generative Artificial Intelligence–Produced False Information. Journal of Shandong University (Philosophy and Social Sciences), (01), 155–164. https://doi.org/10.19836/j.cnki.37-1100/c.2025.01.013

[5] Buchanan B, Lohn A, Musser M, et al. Truth, lies, and automation [J]. Center for Security and Emerging technology, 2021 (1): 2.

[6] Zhao, J. D. (2018). Research on the Forms and Characteristics of Visual Communication of Online Rumors. Communication Power Research, 2 (23), 247.

[7] Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. Business & Information Systems Engineering, 66 (1), 111–126.

[8] Ji, Z., Lee, N., & Frieske, R. (2023). Survey of Hallucination in Natural Language Generation. ACM Computing Surveys, 55 (12), 1-38.

[9] Zhou J, Zhang Y, Luo Q, et al. Synthetic lies: Understanding AI-generated misinformation and evaluating algorithmic and human solutions [C]//Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. 2023: 1-20.

[10] Kahneman, D., & Tversky, A. (1974). Judgment under uncertAInty: Heuristics and biases. Science, 185 (4157), 1124-1131.

[11] Zhao, S. G., & Zhang, Y. T. (2025). The Internal Transformation, Potential Crisis, and Countermeasures of ChatGPT's Influence on Public Opinion. New Media and Society, (01), 182–195 + 462.

[12] Yuan, G. F. (2021). Toward the "Practice" of Theoretical Pathways: Understanding Emotional Expression in Public Opinion. Journal of International Communication, 43 (06), 55–72. https://doi.org/10.13495/j.cnki.cjjc.2021.06.004

[13] Cialdini, R. B. (1984). Influence: The Psychology of Persuasion. William Morrow & Co.

[14] Le Bon, G. (1895/1960). The Crowd: A Study of the Popular Mind. Penguin Books.

[15] Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. Science, 348 (6239), 1130-1132.

[16] Peng, L. (2023). Intelligent Literacy: The Upgrading Direction of Media Literacy in the Era of Intelligent Communication. Journal of Shanxi University (Philosophy and Social Sciences), 46 (5), 101–109.