Vol.7 Iss.1 2026 Communication & Education Review
https://doi.org/10.37420/j.cer.2026.005

Post-Truth After Generation: Mechanisms Through Which Generative AI Shapes Online Public Opinion

Luxilin*
School of Network and New Media, ZhaoQingUniversity, China
***Corresponding author E-mail: sujianxw@qq.com**

Abstract:
This study examines the role of generative artificial intelligence in online public opinion arenas. Using ethnographic observation, semi-structured interviews, and controlled experiments, it systematically analyzes how AI reshapes the workflow of public opinion production, the logic of expression, and public cognition in a post-truth context. A seven-day participatory observation shows that generative AI significantly increases production efficiency, lowers creative barriers, and reinforces content homogenization through default textual and visual templates. Interviews with 15 AI content creators indicate a motivational shift from serious viewpoint articulation toward entertainment-oriented and accommodative remixing, suggesting that AI tools not only change modes of participation but also reconstruct expressive logic. Experimental results further confirm a "viewpoint density effect": when participants are exposed to different quantities of AI-generated content on the same topic, their attitudes, emotional intensity, information judgment, and post-truth perception all increase linearly and significantly. Overall, generative AI influences public emotion and cognitive bias at the micro level through scalable generation, stylistic discipline, and algorithm-driven feedback loops, and reshapes the public opinion ecosystem at the macro level. These findings provide key theoretical and empirical support for understanding public opinion evolution in the AI era.

Keywords:
Emotional intensification, Generative artificial intelligence, Online public opinion, Post-truth, Viewpoint density effect

1 Introduction

As generative artificial intelligence rapidly integrates into everyday information production, online public opinion arenas are being profoundly restructured. Generative AI-centered on text generation, image synthesis, and automated video processing-enables individuals to create and redistribute content on public issues with low barriers and high efficiency, thereby changing the speed, structure, and driving mechanisms of public opinion formation. In a post-truth era, emotional logic often overrides factual judgment; the intensity of viewpoint dissemination can be more influential than truthfulness itself. Against this backdrop, generative AI's scalable production capacity, its tendency toward directional expressive templates, and its coupling with algorithmic recommendation systems can further accelerate the diffusion of emotional content and the polariza-

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

tion of public opinion, making the evolution of online public opinion faster, stronger, and more concentrated.

Existing studies often focus on ethical risks, technical bias, ideological impacts, and structural changes in the opinion space, but fewer examine—starting from real-world opinion contexts—how AI affects public cognition, emotional expression, and viewpoint formation during concrete creative activities and content consumption. Addressing this gap, this paper adopts a mixed-methods design combining ethnographic observation, interviews, and experiments to explore how generative AI intervenes in the production and spread of public opinion, and how it intensifies emotions and biases viewpoints at the micro level, thereby reshaping the online public opinion ecosystem in a post-truth context.

2 Literature Review

2.1 Research on generative artificial intelligence

Regarding language generation and academic applications, Emily M. Bender argues from the perspective of language philosophy and pragmatic competence that large models merely "imitate language" in form while lacking genuine understanding, judgment, and communicative ability, which can mislead users about their level of intelligence (Bender, 2023). From a research-ethics standpoint, Chris Stokel-Walker criticizes generative AI for producing false citations and reinforcing bias in the absence of sufficient governance, thereby undermining academic norms and trust in knowledge, and calls for clearer boundaries and ethical frameworks (Stokel-Walker, 2023).

On ideology and the reconstruction of public opinion, Wei Yan and Liao Heping argue—grounded in historical materialism—that the technical logic and capital forces embedded in generative AI have deeply intervened in ideological construction, creating alienation and security risks (Wei & Liao, 2025). Sun Yue emphasizes the concealed, diffusive, and dynamic nature of ideological risks and proposes integrated governance through a "technology–value" approach (Sun, 2025). Ma Yan and Hao Yanli focus on the dual effects of generative AI on the opinion space, noting that while it enhances personalized information services, it can also intensify cognitive isolation and discourse segmentation, producing a double-edged effect of "cognitive liberation" and "cognitive closure" (Zhang & He, 2025).

While these studies provide systematic analyses of language imitation, ethical risk identification, ideological control, and changes in the opinion space, they generally neglect how generative AI affects public cognition, emotional expression, and dissemination pathways within real public opinion settings. This paper responds to and extends this gap from a micro-level communication mechanism perspective.

2.2 Research on online public opinion

In research on public opinion evolution and governance pathways, Huang Wei and Liu Yizhou, from the perspective of intelligent communication, analyze how generative AI may trigger compound risks in opinion dissemination such as irrational interaction, extreme emotion aggregation, and technological manipulation, forming "public opinion risk clusters" (Huang & Liu, 2025). Wang Qiumei and Zheng Yi emphasize technical identification pathways for ideological risks, proposing dynamic monitoring of sensitive data via fuzzy clustering and classification algorithms to improve the efficiency and controllability of ideological guidance

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

(Wang & Zheng, 2025). Liu Guangrui and Liu Zhonglin focus on user comment behavior and construct a dual-variable model based on interest and event influence to improve prediction of public opinion response trends, offering data support for regulation (Shen & Li, 2025).

In studies of opinion construction and dissemination mechanisms, Luo Yunjuan, based on agenda-setting theory, finds that although online public opinion may not directly affect policymaking, it can exert bidirectional influence on media agendas in specific contexts (Luo, 2014). Dong Yucheng and colleagues build an online–offline opinion evolution model and show that digital platforms can accelerate consensus formation and emotional aggregation (Dong et al., 2017). Wu Yuchen and colleagues, using multi-platform sentiment analysis, reveal emotional polarization and its impact on the opinion environment (Wu et al., 2022).

However, these studies do not systematically analyze the deep involvement of generative AI in opinion content, expressive pathways, and cognitive patterns under post-truth conditions; consequently, they cannot fully address how generative AI reconstructs public cognition, stimulates emotional expression, and shapes opinion formation in real public opinion arenas.

2.3 Research questions

Based on post-truth theory, this paper seeks to answer: (1) Why does generative AI exhibit a distinctive influence on online public opinion? (2) What are the underlying mechanisms through which generative AI affects online public opinion?

3 Method

3.1 Ethnography: observing AI-driven public opinion production

In the context of widespread AI participation in opinion dissemination, this study first proposes two hypotheses: (H1) the rapid generation capability of generative AI improves production efficiency and accelerates the emergence and spread of opinion events; (H2) low operational thresholds encourage ordinary users to participate in content creation, increasing public participation and willingness to re-disseminate. To test these hypotheses, we employed an ethnographic approach and conducted a one-week in-depth observation of a representative AI content creator.

The observed creator was active on Xiaohongshu, frequently using “Doubao” for scriptwriting and “Jimeng AI” to produce opinion-related videos and images. The observation lasted seven days, recording at least one complete creative workflow each day. In familiar environments such as home, office, or cafés, the researcher used participatory observation combined with video/screen recording to document behavioral pathways.

We focused on four dimensions: (1) generation workflow (prompt input, platform switching, visual adjustment, output); (2) time efficiency for each stage; (3) expressive tendency (emotional tone, stance cues, strong visual elements); and (4) feedback interaction (comments, likes, shares after posting). Data were collected via screen recordings, screenshots, field notes, and short interview audio, with informed consent and strict privacy protection. Open coding and thematic analysis were used to organize and analyze the data to build a theoretical framework of AI intervention in public opinion dissemination.

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

3.2 Interviews: user motivation and participation preferences

Two hypotheses guided the interview study: (H3) the fun and low threshold of generative-AI creation shifts user motivation from serious viewpoint output toward participatory and entertainment-oriented creation; (H4) when creating with generative AI, users tend to accommodate mainstream opinion to improve dissemination and algorithmic recommendation. To better understand how AI affects user motivation, we conducted semi-structured interviews to explore expressive intent and behavioral choices during AI-assisted creation.

The interview guide had three parts: (1) demographics (gender, age, occupation, education, social media activity); (2) topic-focused questions (frequency of using tools such as Doubao and Jimeng AI; preference for viewpoint expression vs entertainment creation; whether they cater to hot topics or algorithms for exposure; perceptions of AI–public opinion interactions); (3) deeper probing of emotional drivers and feedback experiences (whether lowered barriers increased willingness to speak; whether creation is treated as entertainment; experiences of off-topic content being recommended or misunderstood).

Throughout the interviews, we followed qualitative research ethics principles (e.g., those articulated by Alan Bryman), ensured informed consent, protected privacy, and allowed withdrawal at any time. Fifteen participants with generative-AI creation experience were interviewed for 30–45 minutes each via WeChat voice/text or in person. Participants were recruited from platforms such as Xiaohongshu, Douyin, and Weibo, as well as social channels (e.g., WeChat groups). The sample was heterogeneous in gender, age, education, occupation, and region. Following the "information saturation" principle, data collection ended when responses became repetitive at the 15th interview. Interviews were transcribed verbatim (approximately 70 000 words) and coded by the researcher.

3.3 Observation and experiment: testing AI-driven opinion bias mechanisms

In public opinion evolution, generative AI may, through scalable and repetitive content generation, subtly drive "cognitive contestation" and reinforce dominant viewpoints. To test this mechanism, we conducted (1) observation in natural settings and (2) controlled experimental simulation.

The observation tracked 1–2 users with AI creation experience over 5–7 days to capture viewpoint-selection tendencies during topic selection. Through remote screen sharing or on-site accompaniment, the researcher recorded the full workflow from topic selection and copy generation to visual design and posting. We paid particular attention to whether creators preferentially chose mainstream, trending, or high-dissemination viewpoints. Screen recordings, notes, and screenshots were collected and later compared with subsequent opinion trajectories.

The experiment used a between-groups design to simulate exposure under different viewpoint densities. Four groups watched 10 videos each: Group 1 (control) watched 10 unrelated videos; Group 2 watched 2 related + 8 unrelated; Group 3 watched 4 related + 6 unrelated; Group 4 watched 8 related + 2 unrelated. Related videos focused on the same public issue (e.g., controversy over a "make-up workday" policy), while unrelated videos had no opinion cues.

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

Pre-test and post-test questionnaires measured changes in attitude, emotional response, information credibility judgment, and post-truth perception. The questionnaire contained five parts: (1) stance/attitude rating; (2) emotional intensity; (3) credibility judgment; (4) post-truth perception; (5) control variables (age, gender, education, media use habits). All questionnaires were completed online to ensure anonymity and consistency. After the experiment, some participants took part in brief follow-up interviews to explain attitude changes.

By combining observation and experiment, this study aims to reveal how generative AI affects public attitudes, emotional expression, and viewpoint bias through technical pathways at the micro level, providing mechanism-level support for subsequent research.

4 Results

4.1 Sample characteristics and reliability

4.1.1 Participant characteristics

A total of 16 valid participants were included, all recruited voluntarily via WeChat in Guangdong Province. The gender ratio was balanced (43.8% male, 56.3% female). Most participants were aged 18–24 (75.0%). The majority held a bachelor's degree (87.5%). Daily social media use was mainly 1–3 hours (50.0%). Exposure to AI videos was relatively low (62.5% reported "rarely"). After random assignment to four groups, demographic distributions showed no significant differences ($p > .10$), as shown in Table 1.

Table 1 Participant demographic characteristics (N = 16)

Variable	Category	n	%
Gender	Male	7	43.8
	Female	9	56.3
Age	18–24	12	75
	35–44	1	6.3
	45–59	2	12.5
	≥60	1	6.3
Education	Junior college	1	6.3
	Bachelor	14	87.5
	Graduate	1	6.3
Daily social media use	<1h	4	25
	1–3h	8	50
	>3h	4	25

4.1.2 Reliability test

Cronbach's α for each construct was acceptable: attitude toward viewpoints (.78), emotional response (.72), information judgment (.81), and post-truth perception (.76).

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

4.2 Main results

4.2.1 Overall pre–post changes

Paired-sample t-tests for the full sample show that all four dimensions increased significantly after the experiment (ps < .001), with large effect sizes (Cohen's d > 1.1), as shown in Table 2.

Table 2 Overall pre–post changes for the full sample (N = 16)

Dimension	Pre-test M±SD	Post-test M±SD	Δ±SD	t(15)	p	Cohen's d
Attitude toward viewpoints	11.06±2.14	12.88±2.20	+1.81±1.17	6.2	<.001	1.55
Emotional response	6.13±1.09	7.31±1.08	+1.19±0.83	5.72	<.001	1.43
Information judgment	10.94±2.06	12.75±2.02	+1.81±1.33	5.44	<.001	1.36
Post-truth perception	8.44±1.41	9.75±1.57	+1.31±1.14	4.6	<.001	1.15

4.2.2 Main effect of video density

A one-way ANOVA was conducted with video density (G1=0, G2=2, G3=4, G4=8) as the between-subject factor and the pre–post difference (Δ) as the dependent variable. Results show significant main effects across all four dimensions: attitude toward viewpoints $F(3,12)=6.40$, $p=.006$, $\eta^2p=.62$; emotional response $F(3,12)=7.15$, $p=.004$, $\eta^2p=.64$; information judgment $F(3,12)=8.93$, $p=.001$, $\eta^2p=.69$; post-truth perception $F(3,12)=9.82$, $p=.001$, $\eta^2p=.71$, as shown in Table 3.

Table 3 One-way ANOVA results by video density group (N = 16)

Dimension (Δ)	F(3,12)	p	η^2p	1-β
Attitude toward viewpoints	6.4	0.006	0.62	0.84
Emotional response	7.15	0.004	0.64	0.87
Information judgment	8.93	0.001	0.69	0.93
Post-truth perception	9.82	0.001	0.71	0.95

Tukey HSD tests indicate that the high-density group (G4) scored significantly higher than the control group (G1) across all dimensions (ps < .01). Compared with the low-density group (G2), significant differences were found on all but attitude ($p < .05$). The medium-density group (G3) was significantly higher than G1 on information judgment and post-truth perception ($p < .05$). A Jonckheere–Terpstra trend test showed a significant linear increasing effect ($z > 2.98$, $p < .003$).

4.2.3 Moderating effects of individual characteristics

Pearson correlations (Table 4) show that age is significantly negatively correlated with all Δ values ($r_s = -.49$ to $-.55$, $p < .05$), while daily social media use time is significantly positively correlated with Δ ($r_s = .58$ to

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

.64, p < .01). Education shows a positive trend but is not significant (rs = .38 to .45, p > .08). After including age and social media time as covariates in an ANCOVA, the main effect of video density remained significant (minimum F = 6.12, p = .009), indicating that the density effect is independent of these individual differences.

Table 4 Pearson correlations between individual characteristics and Δ (two-tailed)

Variable	Attitude Δ	Emotion Δ	Judgment Δ	Post-truth Δ
Age	-.48*	-.52*	-.55**	-.49*
Daily social media use time	.58*	.61**	.64**	.59*
Education	0.38	.41^	.45^	.43^

* p < .05, ** p < .01, ^ p < .10

5 Discussion

5.1 Macro level: new changes in AI-driven public opinion production

The seven-day participatory observation tracked the participant's daily creative activities using generative AI and followed the full workflow from topic selection to posting. Overall, AI tools occupied a central position in opinion-related content production: efficiency increased substantially and production barriers dropped markedly compared with the mass-media era. Moreover, a mobile-only workflow further lowered barriers and encouraged enthusiastic participation in AI content generation.

Typical short-video production involved generating scripts via Doubao, then producing videos with Jimeng AI; average time was under 30 minutes. Image revision tasks often took only 10–15 minutes from conception to completion. Daily creation usually started with browsing hot topics. The creator switched among Xiaohongshu, Weibo, and other short-video platforms to assess which issues had the greatest dissemination potential. Once a topic was chosen, the creator used Doubao to draft a script, often with very short prompts; after skimming the output, only minor edits were made before moving to visual production.

In visual production, the creator relied mainly on Jimeng AI for video or image materials, showing a strong tendency toward automation: most parameters remained at default settings, with only topic cues or simple visual instructions adjusted. Generation typically completed within minutes, followed by light editing, subtitles, and cover design. Across seven days, average end-to-end time was under 20 minutes, with the fastest full workflow taking only 12 minutes. Compared to the creator's self-reported "at least one hour" for manual production, AI compressed the cycle dramatically, enabling rapid and continuous responses to trending events.

Expressively, the content displayed stable features: emotionalized titles (e.g., "shocking," "unexpected," "plot twist") appeared frequently; the creator attributed this to AI's default headline style and its higher likelihood of attracting recommendation traffic. Jimeng AI images often used high-saturation colors and exaggerated lighting for strong visual impact, which the creator rarely modified. In social events, explicit extreme stance language was avoided, but tone and atmosphere implied attitudes, producing an implicit stance. Long-term reliance on AI scripts gradually shifted the creator's language toward the model's default style, resulting in noticeable homogenization.

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

Post-publication interaction followed a rapid feedback cycle. Notably, users were not passive "hypodermic needle" audiences; instead, they acted as active nodes within the platform architecture. From audience theory perspectives (active audience; uses and gratifications), users build meaning through viewing, interaction, and remixing driven by interests, emotional needs, social expectations, and platform incentives. Through likes, comments, secondary creation, and trend-following generation, audiences shift from "receivers" to "nodes" in a circulating opinion network, reinforcing AI content reproduction and diffusion.

Node participation is emotional, immediate, and mutually reinforcing. Users respond quickly based on current emotional states, expressing stances via comments or remixes, producing short-term emotional aggregation. AI lowers the threshold for remixing, making "node feedback–redistribution" more frequent and amplifying. This structural shift enables AI content to circulate in looped, strengthening patterns, allowing emotions and viewpoint bias to accumulate rapidly into substantial public opinion energy.

Accordingly, the creator closely monitored early comments and like velocity to infer whether the platform would recommend the content. When interaction was weak, the creator regenerated covers or adjusted headlines, often producing noticeable improvements. Motivated by attention maximization, the creator used AI to tune production based on comment feedback, participating in the co-creation of opinion trajectories. When content gained high discussion, the creator quickly generated "follow-up videos," maintaining exposure and heat. AI efficiency strengthened the "feedback–reproduction–feedback" loop, tightening production and interaction rhythms.

For example, on Day 3, the participant produced a short video about a "make-up workday policy" controversy using a Doubao-generated script and Jimeng AI visuals. Within 10 minutes, strong-emotion comments appeared (e.g., "formalism again," "inhumane"), and like velocity was higher than usual. The participant judged further dissemination potential and immediately prompted Doubao to generate a follow-up script from a "citizen perspective," produced matching visuals via Jimeng AI, and posted a second video while the first was still rising. Together, the two posts formed an "echo effect," attracting more interactions and further intensifying discussion. This illustrates how real-time feedback enabled rapid reproduction and how AI's high speed compressed cycles and increased impact.

Psychologically, frequent reference to comments and rapid AI-based tuning is driven by (1) social validation: interaction metrics signal recognition, motivating creators to adjust narratives to maintain perceived approval; (2) algorithmic incentives: creators infer that dominant comment sentiments reflect mainstream expectations and use AI to fine-tune content for recommendation capture. Moreover, responding to feedback becomes collaborative: comments provide new creative materials that creators transform—via AI—into new scripts, images, or follow-ups, creating participatory co-creation among "creator–audience–AI."

Overall, the observation shows deep AI involvement in opinion content production: it lowers barriers and compresses time, enabling ordinary users to participate at high frequency; meanwhile, its textual and visual templates subtly discipline expressive preferences, shaping overall style and emotional structure. These findings provide an empirical basis for understanding how AI reshapes public opinion production processes.

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

5.2 Meso level: mechanisms of AI influence on public opinion

Semi-structured interviews with 15 AI content creators produced multi-dimensional evidence on motivation changes, expressive strategies, and emotional experiences, supporting the two interview hypotheses.

Overall, most interviewees viewed AI as a "barrier-lowering creative tool," agreeing that generative AI increases the fun of creation and makes them more willing to participate in opinion-related remixing. Motivation showed a clear shift from "viewpoint output" to "entertainment-oriented re-creation." Many noted that traditional opinion expression requires rigorous argumentation or higher language skills, whereas AI lets them produce publishable content quickly without "thinking too much," turning expression into a relaxed, enjoyable activity rather than serious public deliberation. This shift was particularly salient among younger users, who described AI generation as playful "interactive entertainment," even treating trending-content creation as a lightweight "task."

In terms of strategy, most interviewees exhibited "accommodative expression." Many admitted consulting trending lists and mainstream viewpoints and adding directional wording into prompts to increase the likelihood of recommendation. Some explicitly said they asked AI to generate content "more aligned with mainstream emotion," especially for controversial events, avoiding dissenting views out of concern for backlash or lower account weight. Several emphasized that AI language is itself "official, safe, and error-averse," so even if their personal views differ, the final output is disciplined by the model's default style, increasing convergence and accommodation.

For instance, one interviewee described wanting to criticize a local management measure, but the AI script automatically used cautious phrases such as "policy intent," "objective analysis," and "rational view." Although the interviewee's original intention was more critical, they adopted the AI version because it "felt safer and more like what the platform would recommend," resulting in a seemingly objective but mainstream-leaning tone not fully aligned with their personal stance. This made them realize that long-term reliance on AI can guide expression toward "neutral–cautious–error-free," weakening individual differences and increasing homogeneity and accommodation.

Another prominent pattern was the "gamification" of creation. Many found novelty in AI randomness and visual effects, and enjoyed iteratively adjusting prompts and templates. Some said they created not to express a viewpoint but to see what AI could produce. This motivation downplayed the importance of content itself and encouraged more frequent participation in opinion-related re-creation.

Regarding feedback, most reported that platform recommendations or rising comment counts increased their enthusiasm to continue creating. When encountering off-topic outputs, misreadings, or criticism, they tended to blame AI scripts or models rather than themselves, sustaining tolerance for AI creation. Some worried about being drowned in homogeneous content, but this did not significantly reduce their creation frequency.

In sum, two conclusions emerge: (1) generative AI shifts user motivation toward entertainment-oriented, lightweight entry into the public sphere, with the process of generating content becoming the main driver rather than viewpoint output; (2) AI-assisted expression shows clear accommodative tendencies shaped by

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

both algorithmic incentives and AI's default cautious style, pushing opinion content toward emotionalization and homogenization. Thus, generative AI not only lowers participation barriers but also subtly reshapes expressive logic, producing a faster, lighter, and more convergent dissemination structure.

5.3 Micro level: emotional intensification by generative AI

A large quantity of same-topic generative-AI videos and posts can produce post-truth effects during opinion fermentation, triggering cognitive contestation that reinforces dominant viewpoints and weakens minority ones. This reinforcement is unrelated to factual accuracy and instead depends on perceived approval among netizens.

The study compared pre- and post-tests for 16 participants and examined whether exposure to different densities of same-topic opinion videos causes systematic changes in attitudes, emotions, and cognition. Paired-sample t-tests show significant increases across four dimensions after exposure (ps < .001; Cohen's d > 1.1), suggesting that even short exposure can significantly influence attitude direction and emotional intensity and alter credibility judgments and post-truth perception.

One-way ANOVA further shows significant main effects of video density on all four dimensions (attitude: $F(3,12)=6.40$, $p=.006$; emotion: $F(3,12)=7.15$, $p=.004$; judgment: $F(3,12)=8.93$, $p=.001$; post-truth: $F(3,12)=9.82$, $p=.001$), with large effect sizes ($\eta^2p \geq .62$). Tukey tests confirm that the high-density condition (8 related videos) produces significantly greater changes than the control (0 related videos) across all dimensions (ps < .01) and exceeds the low-density condition in most dimensions. Trend tests indicate significant linear increases ($z > 2.98$, $p < .003$), implying that more concentrated exposure yields stronger attitude bias and cognitive change.

Regarding moderators, younger participants were more susceptible (age negatively correlated with Δ), while heavier social media users showed stronger effects (use time positively correlated with Δ). After controlling for age and use time in ANCOVA, the density effect remained significant, indicating it is driven by stimulus conditions rather than individual differences.

Overall, generative AI is not merely a technical tool for producing opinion content; it is a key force shaping emotional flows, viewpoint bias, and opinion-structure evolution. Across observation, interviews, and experiments, emotional intensification follows a consistent mechanism: scalable production, template-based emotional expression, and feedback-driven cyclical reproduction amplify emotions continuously, pushing opinion toward a particular direction in a short time.

In the observed case, emotional intensification occurred through real-time creator–audience interaction. After the first video about the policy controversy, strong-emotion comments quickly emerged; AI efficiency enabled rapid follow-up scripts and visuals; high-saturation visuals further amplified negative affect. The follow-up video intensified interactions, forming a cyclical structure of audience emotion → AI-assisted generation → renewed audience stimulation, accelerating diffusion and strengthening a dominant emotional stance.

Interview data similarly indicate that users intentionally adopt expressions that "generate more comments" and "match the emotional atmosphere," while AI scripts themselves often contain emotionally charged wording. Even without strong prior emotions, users are inclined to continue creating within AI's emotional

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

framing, gradually internalizing its discipline of emotional expression. The experiment provides quantitative evidence: emotional response increases linearly with density, with the strongest intensification under high density, consistent with post-truth cognition where emotion overrides fact in credibility judgment.

Thus, emotional intensification is not driven by a single link but by a chain of "template-based emotional output → emotional audience feedback → algorithmic amplification → rapid AI reproduction." This chain enables emotions to spread exponentially and encourages rapid convergence of collective expression. As AI becomes more prevalent, such structural amplification may become a new normal in public opinion evolution and poses major challenges for governance.

6 Conclusion

Using ethnographic observation, semi-structured interviews, and controlled experiments, this study systematically reveals how generative AI deeply intervenes in the production of online public opinion, emotional expression, and viewpoint bias in a post-truth context. It finds that by lowering creative barriers and increasing efficiency, AI transforms ordinary users from passive receivers into high-frequency participants, and—through textual styles, visual templates, and algorithmic recommendation—forms a loop of "rapid generation–audience feedback–reproduction." In this process, emotional language, visual dramatization, and viewpoint convergence become mainstream expressive forms, driving faster aggregation, polarization, and diffusion. Experimental evidence confirms that higher viewpoint density significantly strengthens emotional response, attitude bias, and post-truth perception, indicating a stable structural mechanism of emotional intensification in AI-driven opinion evolution.

This study also has limitations. First, the observation and interview samples are representative but limited in size and may not fully capture cross-platform and cross-group differences. Second, the experimental setting is controlled and cannot fully reproduce the complexity and dynamics of real opinion arenas; future work should expand samples to enhance external validity. Third, this study focuses on the user side; how platform algorithms, moderation systems, and commercial logics jointly shape opinion arenas with AI-generated content warrants deeper investigation. Nevertheless, by triangulating methods, this research provides empirical evidence for how generative AI reshapes online public opinion ecology by reducing production costs, disciplining expressive style, and amplifying emotional dissemination, offering a foundation for future governance and ethics research.

References

[1].Wei, Y., & Liao, H. (2025). Exploring the real-world challenges and development pathways of generative artificial intelligence from the perspective of historical materialism. Journal of Nanning Vocational and Technical University, 2025(05), 29–35. (in Chinese)
[2].Sun, Y. (2025). Identifying and governing ideological risks of generative artificial intelligence from the perspective of socialist core values. Theory Guide, 2025(10), 52–56.(in Chinese)
[3].Zhang, L., & He, Z. (2025). The double-edged sword effect of generative AI information acquisition on the filter bubble. Information and Documentation Services, 2025(06). (in Chinese)

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)

[4].Huang, W., & Liu, Y. (2025). The formation mechanism of online public opinion risk clusters from the perspective of intelligent communication. Information Science, 2025(09), 50–59. (in Chinese)
[5].Wang, Q., & Zheng, Y. (2025). Dynamic identification and guidance pathways for ideological security data in online public opinion. Information Science, 2025(09). (in Chinese)
[6].Shen, W., & Li, X. (2025). Predicting user online comment behavior in online public opinion events. Journal of Modern Information, 2025(12), 118–129.(in Chinese).
[7].Bender, E. M. (2023). The False Promise of Imitating Text.
[8].Stokel-Walker, C. (2023). The Promise and Peril of Generative AI. Nature, 614(7947), 214–216.
[9].Luo, Y. (2014). The Internet and agenda setting in China. International Journal of Communication, 8, 1289–1312.
[10].Dong, Y., et al. (2017). Dynamics of Public Opinions in an Online and Offline Social Network.
[11].Wu, Y., et al. (2022). Emotion Analysis of Public Opinion in Online Platforms.

© 2026 by the author(s); licensee Mason Publish Group, this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (http://creativecommons.org/licenses/by/4.0/)