

# Human-AI Collaboration Mechanism Study on AIGC Assisted Image Production for Special Coverage

Yajie Yang<sup>1</sup>, Yuqing Zhao<sup>1</sup>, Xiaochao Xi<sup>1†</sup>, Yanan Zhu<sup>2</sup>

<sup>1</sup>School of Digital Media and Design Arts, Beijing University of Posts and Telecommunications  
No. 10 Xitucheng Road, Haidian District, Beijing, 100876, China

<sup>2</sup>Audio-Visual Department of Xinhua News Agency, Beijing, China

yangyajie2436@bupt.edu.cn, zhaoyuqing@bupt.edu.cn, xixiaochao@bupt.edu.cn, 18612832400@163.com

## Abstract

Artificial Intelligence Generated Content (AIGC) assisting image production triggers controversy in journalism while attracting attention from media agencies. Key issues involve misinformation, authenticity, semantic fidelity, and interpretability. Most AIGC tools are opaque “black boxes,” hindering the dual demands of content accuracy and semantic alignment and creating ethical, sociotechnical, and trust dilemmas. This paper explores pathways for controllable image production in journalism’s special coverage and conducts two experiments with projects from China’s media agency: (1) Experiment 1 tests cross-platform adaptability via standardized prompts across three scenes, revealing disparities in semantic alignment, cultural specificity, and visual realism driven by training-corpus bias and platform-level filtering. (2) Experiment 2 builds a human-in-the-loop modular pipeline combining high-precision segmentation (SAM, GroundingDINO), semantic alignment (BrushNet), and style regulating (Style-LoRA, Prompt-to-Prompt), ensuring editorial fidelity through CLIP-based semantic scoring, NSFW/OCR/YOLO filtering, and verifiable content credentials. Traceable deployment preserves semantic representation. Consequently, we propose a human-AI collaboration mechanism for AIGC assisted image production in special coverage and recommend evaluating Character Identity Stability (CIS), Cultural Expression Accuracy (CEA), and User-Public Appropriateness (U-PA).

## Introduction

Tools like ChatGPT (Cao et al. 2023), Midjourney (Byrne 2023), Stable Diffusion (Rombach et al. 2022) and DALL-E (Ramesh et al. 2022; Wang et al. 2023) have drawn increasingly attention to media agencies worldwide. Newsrooms pervasively adopt the Generative Artificial Intelligence (AIGC) for automating routine tasks, summarizing large datasets, and producing imaginative images or draft copies. Up to 2023, around 73% of journalists worldwide and over 85% of German media employees had experimented with generative AI tools (Lewis et al. 2025). Other examples include A VALENTINE, FROM A.I. TO YOU by New York Times, and AIGC Feature Programme by Xinhua News Agency (Chen 2023). Diakopoulos suggests “social

media content creation” is one of the thirteen scenarios that AIGC may assist journalism, but those advanced tools “need careful evaluation before they are used for publishing. More importantly, it raises critical issues about accuracy, ethics and employment,” that opens promising research topics to AI engineers, journalists and designers (Trattner et al. 2022).

## Challenge of Journalistic Image Production

The rapid development of AIGC has blurred the line between authentic and synthetic media. Journalistic images generated by AI models such as Large Language Models (LLMs), generative adversarial networks (GANs), and diffusion models, however, are posing serious challenges to journalistic authenticity due to the inherent “black box” nature of AI systems, which can lead to image distortion and fake news proliferation (Jawaher, Suhuai, and Yuqing 2023). The internal logic of image composition, the mapping of textual prompts to visual elements, and the reasoning processes remain opaque, making it difficult to trace how outputs are generated. This lack of transparency leads to substantial risks of content distortion and misinformation (Doshi-Velez and Kim 2017; Jia 2024). An analysis of recent academic research on the challenges facing journalistic image production has identified the following four typical issues:

**Misinformation and virality.** According to a survey by Thomson Reuters Foundation, 81.7% of journalists reported using AI tools in their job, and they stressed that all AI outputs must be verified and fact-checked to maintain accuracy and credibility. An empirical study by Drolsbach and Pröllochs examine how AI-generated misinformation differs from conventional misinformation on the social media platform X (Doshi-Velez and Kim 2017). The scholars find that AI-generated misinformation tends to be more entertaining and spreads more widely yet appears slightly less believable. Classification of distortions shows that hallucinations span multiple categories, from simple factual errors to unfounded fabrications, underscoring the need for systematic evaluation and governance.

**Authenticity perception** Farooq and Vreese conduct experiment and demonstrate that aesthetic realism strongly influences whether people deem AI images authentic, while compression can mask synthetic origins (Farooq and

de Vreese 2025). The interplay between AIGC tools and participates confidence suggests that transparent AI can assist the production of images but depends on the trust of using (Fang et al. 2025). More seriously, undermine the authenticity and credibility expected of journalistic images, affecting public perception and the trustworthiness of media.

**Semantic fidelity** Technical studies show that diffusion models may neglect some prompt tokens, leading to misaligned or hallucinated content (Zhang et al. 2024). Attention regulation and similar techniques improve fidelity without retraining, but they do not address social dimensions such as misuse. An AICII 2023 study compares DALL E 2 images created from news headlines with human selected photographs and found that AI images often evoke curiosity and confusion and elicit a wider range of emotions but provide less context and newsworthiness (Paik et al. 2023)

**Interpretability and governance** The push for explainable detection highlights the importance of grounded reasoning – not just binary classification but also localisation and natural-language explanations. Ji and other scholars introduce the FakeXplained dataset, which couples binary authenticity decisions with bounding-box localisations and natural-language rationales, boosting explainability (Ji et al. 2025). Interactive systems like ASAP help users understand deceptive patterns, bridging the gap between technical detection and human comprehension. Reviews call for model transparency, data quality audits and human-in-the-loop processes (Huang et al. 2024).

In summary, research emphasize that these problems undermine the authenticity and credibility expected of journalistic images, affecting public perception and the trustworthiness of media (Yao et al. 2024). The review notes that while progress in explainable AI provides promising solutions, AIGC still faces challenges in controllability, explainability and robustness. Future directions include neuro symbolic integration to move generative systems beyond “black box” behaviour and human-in-the-loop approaches to ensure better oversight, especially in the field of journalistic image production (Zihao 2025).

To address the persistent issues of semantic inconsistency, cultural-symbol mismatch, and limited editorial reliability in current AIGC journalistic images, this study picks one the most typical scenarios in journalistic image production – special coverage – as an example and proposes a two-part experimental framework. It aims to verify whether a combined approach, which includes both cross-platform performance evaluation and the introduction of a human-in-the-loop workflow, can effectively meet the dual demands of newsroom practice: production accuracy and communication precision (de Lima-Santos, Yeung, and Dodds 2025). The structure and findings of the experiments are presented in the following section.

## Experiment 1: Cross-Platform Adaptability

Tackling the previously identified shortcomings – such as structural inaccuracy, symbol mismatch, and semantic unreliability – in newsroom contexts, the authors conduct ex-

periment part 1 by benchmarking three mainstream generative platforms: the platform A and B are popular tools from China mainland and the platform C comes from U.S. These three platforms were chosen because of their representativeness, which stems from design: we selected two widely used Chinese platforms and a widely used American platform to span the dominant technical and policy lineages in current practice. Because the platforms embody different decoding heuristics, data sourcing, and governance constraints, the comparison surface they define is broad. Our findings—semantic drift in culturally marked cues vs. structural fidelity trade-offs—are therefore instructive for newsrooms worldwide that must anticipate how model provenance affects editorial reliability. This is why the paper argues for a HITL pipeline: algorithmic strengths (speed, consistency) are preserved, but editorial guardrails are reintroduced to correct exactly those weaknesses (cultural/semantic misalignment) that pure automation struggles with. In revision, we will explicitly tie each observed failure mode to the three-layer error formation model (Back-End Data, Task Context, Model Response), clarifying how the comparison generalizes to other systems trained with different data and filters.

Before running the studies, we prepared a bilingual prompting framework to secure semantic consistency across platforms. We authored a canonical Chinese prompt with four slots, covering source citation, spatiotemporal context, task instruction, and element constraints, so editorial intent was compressed into a single, traceable command. We then produced a concept preserving English counterpart, translating at the level of terms of art and cultural referents rather than word by word, and substituting newsroom standard phrasing wherever literal renderings could mislead. All prompt versions and inference parameters were archived to guarantee image level provenance and to prevent unlogged prompt drift. For Experiment 1, the same bilingual pair was fixed and reused on every platform. For Experiment 2, the bilingual RAG was paired with LoRA controls so editors could gently correct culturally sensitive details. Acknowledging limits in cross linguistic model competence, we established practical equivalence through expert translation and editorial pretesting, and specified follow ups using controlled vocabularies and JSON LD slot schemas, documented in an appendix.

The selected scenarios have commonly appeared in special coverages: military, disaster relief and village school opening. For each scenario, ten images are generated per platform, and all prompts and parameters were systematically archived by bilingual RAG (Retrieval-Augmented Generation), reflecting typical journalistic conditions. The later evaluation centers on three representatives. Detailed prompts are shown on below:

- Military: war photojournalist; capture an RQ-1 Grey Eagle drone flying low along the Chinese border while dispersing leaflets; produce a realistic press photograph.
- Disaster Relief: documentary photographer; document nighttime flood relief in southern China with evacuees and rescue boats under lights; create a documentary-style

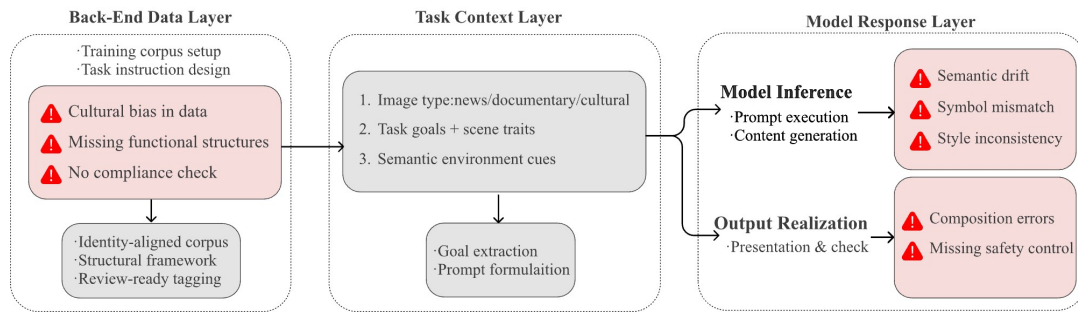


Figure 1: Mechanism model of news context AIGC failures. Source: author.

news image.

- Village School Opening: campus photojournalist; shoot rural Chinese pupils walking to school at dawn with backpacks and a waving national flag; output an authentic, warm rural-scene news photo.

For each scenario, ten images were generated per platform, and all prompts and parameters were systematically archived.

### Analysis of Platform-Generated Variations and Adaptability to Journalistic Contexts

To systematically benchmark the cross-platform adaptability of the current generation system and ensure the evaluation results are rigorous, objective, and actionable for newsroom application, a strict multi-dimensional evaluation protocol was designed and applied to the generated images. This protocol encompassed three core assessment dimensions: semantic alignment with news context, visual structural integrity including pose, facial features, and scene composition, and newsroom suitability such as compliance with media norms and readability. Quantitative scoring and qualitative analysis were combined to capture both objective performance and practical usability. By comparing and analyzing the nuanced differences in output quality, functional strengths, and inherent limitations among the three target platforms, Platform A, B, and C, the key findings are summarized as follows.

The generated outputs exhibit noticeable variation across platforms and topics. Platform C shows strong visual coherence but frequent mismatches between symbols and contextual cues. Platform A maintains moderate semantic alignment but struggle with structural integrity in complex scenes. Platform B produces realistic renders yet suffers from pose distortion and facial inconsistencies. These disparities stem from each system’s decoding heuristics, training-data composition, and built-in safety filters, jointly preventing any platform from simultaneously satisfying the newsroom metrics of cross-view facial-and-pose consistency, action-scene-symbol congruence, and newsroom suitability, as shown in Figure 2).

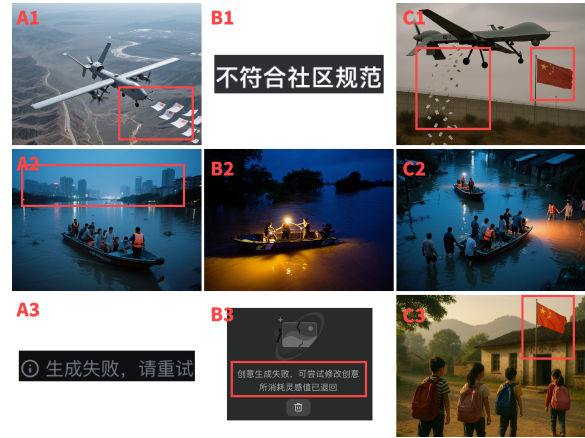


Figure 2: Reference diagram for cross-platform adaptability comparison. Source: Screenshot from AIGC platforms.

### Common Error Types and Analysis of Controllable Generation Logic

Due to the difficulty of authenticating platform generated images, existing systems cannot automatically distinguish usable from unusable content in journalistic contexts. a single “AI-generated” label is no longer sufficient; a dual track framework that combines source attribution with authenticity verification is urgently needed (Jia 2024). The authors therefore outline a three layer mechanism model of AIGC image error formation (see Figure 1): the Back-End Data Layer, Task Context Layer, and Model Response Layer. First, the back-end data layer governs corpus selection and instruction parsing, but is often hindered by culturally irrelevant data, lack of visual templates, and absence of compliance filters. Second, the task-context layer is responsible for identifying image type, objectives, and scene elements; failures here misdirect semantic orientation and prompt design. Third, the model-response layer carries out generation, where issues such as semantic drift, symbol misalignment, style inconsistency, and insufficient risk control render many outputs unfit for news use.

A comprehensive analysis of model architecture, training resources, and generation logic reveals that existing platforms remain insufficient for rigorous newsroom use. Ac-

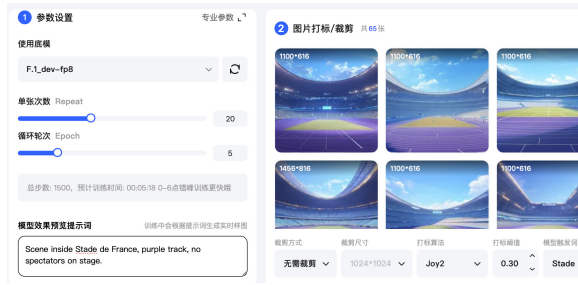


Figure 3: Self-trained LoRA "Tech-style studio". Source: screenshot from Liblib.

cordingly, Experiment 2 introduces a prompt- and structure-anchored workflow designed to enhance the reliability and appropriateness of AI-generated news visuals (Jia 2024).

## Experiment 2: Human-in-Loop Pipeline Experiment

Experiment 1 reveals that fully automated generative platforms often failed to meet the professional standards required in newsroom environments. These limitations included inaccurate facial representation, misalignment between actions and cultural symbols, and a lack of editorial reliability. To address these issues, Relying on Xinhua News Agency’s Paris Olympic promotion project, Experiment 2 introduces a human-in-the-loop workflow tailored specifically for journalistic content creation. Unlike the previous “black box” generation process, this workflow embeds human expertise across key stages, including bilingual RAG prompt construction, LoRA model configuration guided by visual journalism experts, and real-time adjustment of image outputs during inference. These human interventions ensure that editorial intentions are accurately conveyed and that generated visuals maintain cultural integrity and public credibility within the journalistic domain. with the complete workflow architecture visually detailed in Figure 3 and 4.

By integrating structural anchoring, semantic weighting, and human-supervised corrections, the pipeline significantly improves alignment between generated visuals and real-world special coverage scenarios. It enhances the fidelity of character-scene interaction, ensures consistency in style and composition, and minimizes cultural-symbol mismatches – factors that are crucial in producing editorial-grade content for public dissemination. More details on LoRA training can be found in the M1.

In order to guarantee the semantic clarity, structural controllability, and interpretability, the workflow fuses a Retrieval-Augmented Generation (RAG) prompt framework with model-control mechanisms. A four-slot prompt (information-source citation + temporal/spatial context + task instruction + element constraints) is compressed into a single high-density command; ControlNet plus custom character LoRA (Panda Digital IP) and scene LoRA (Tech-style studio) anchor skeletal pose and perspective; during generation, gradient-weighted key tokens and node-based negative prompts in ComfyUI enable real-time correction and style

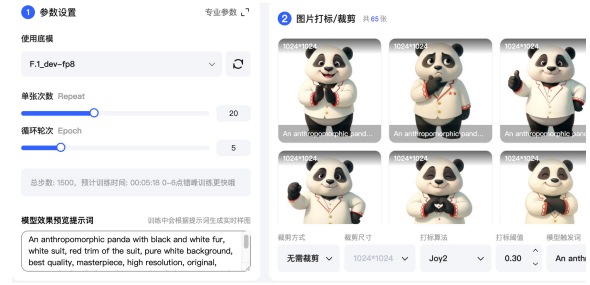


Figure 4: Self-trained LoRA "Panda Digital IP". Source: Screenshot from Liblib.

Component	Configuration	Notes
Platform	ComfyUI 0.19 + Stable Diffusion XL-1.0	Generation framework and core model
Structural Anchoring	ControlNet + LoRA (Tech-style studio; Panda Digital IP)	Locks pose, apparel, perspective
Prompt Framework	RAG + prompt (Source + Context + Task + Constraints)	Merges clues for semantic precision
Real-time Control	ComfyUI node weights and negative prompts	Live correction and style unification

Table 1: Workflow architecture

unification. This architecture simultaneously boosts semantic precision, reduces ambiguity, and enhances cultural fidelity while markedly improving editorial control, where Table 1 formally specifies its technical configuration.

## Experimental Results and Analysis

The authors recruit 102 raters aged 12 – 60 years old,, evenly distributed across age strata and drawn from diverse occupations—including students, media professionals, civil servants and freelancers. Each participant blind-reviewed 6 workflow-generated images and completed a five-point Likert questionnaire covering three axes: cross-frame facial-and-pose coherence, congruence between foreground action, background setting and cultural symbols, and perceived journalistic reliability. All ratings were used to calculate the mean and standard deviation as follows:

$$\bar{s}_j = \frac{1}{n} \sum_{i=1}^n s_{ij}, \quad \sigma_j = \sqrt{\frac{1}{n} \sum_{i=1}^n (s_{ij} - \bar{s}_j)^2} \quad (1)$$

Average  $\bar{s}_j$  and standard deviation  $\sigma_j$  for each metric  $j$  are computed as follows  $n = 102$ , where  $s_{ij}$  denotes the expert score of Image  $i$  on Metric  $j$ , as shown in Table 2. The full questionnaire is provided in Supplementary S1, and the aggregated rating table in Supplementary S2.

The findings demonstrate that the prompt-controlled, structurally anchored human-AI workflow can reliably produce high-credibility news images within 3-5 iterations: cross-view facial-and-pose consistency exceeds 93%, action-scene-symbol congruence is clear, and the newsroom

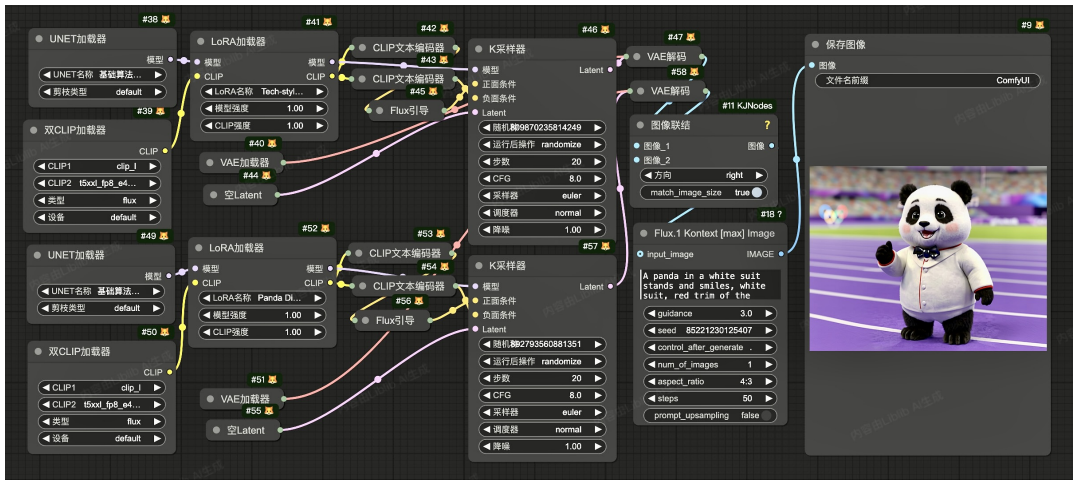


Figure 5: Mechanism model of AIGC image error formation workflow. Source: Screenshot from Liblib

Metric	Mean	SD	Notes
Cross-view facial-and-pose consistency	4.82	0.12	Consistency of facial features and pose across views
Action-scene-symbol congruence	4.46	0.23	Alignment of action, setting, and cultural symbols
Journalistic suitability	4.58	0.21	Documentary style suitable for publication

Table 2: User perception scores across three dimensions

suitability mean score reaches 4.58/5 (approximately 92%); the aggregate evaluation from 102 raters confirms strong audience acceptance.

Aggregated scores indicate near ceiling ratings for cross view facial-and-pose consistency and newsroom suitability, whereas action-scene-symbol congruence-while satisfactory-exhibits slightly higher variance, suggesting room for refining textual signage and symbolic accuracy. Moreover, The trust-labeling mechanism links each image to a verifiable provenance record (platform/version, LoRA IDs, prompts, filters), which travels with the asset and can be shared with auditors/readers. The HITL workflow assigns clear human responsibility: editors set criteria, approve prompts, and curate, while automated modules handle scoring and filtering. This ensures accountability, as anomalies trace to specific steps/approvers. Compliance is bolstered by CLIP-based checks and filters aligned with newsroom policy. A new “Ethical Handling Checklist” details when to avoid generation (e.g., high-risk reconstruction) and how to disclose composites. The result is enhanced transparency, auditability, and trust, while maintaining editorial control.

## Human-AI Collaboration Mechanism

### Comparison of Experiments

The two experiments while distinct in their design focus, together offer a coherent understanding of the strengths and limitations of AIGC in news image production. The first experiment, centered on cross-platform adaptability, highlights the current inability of mainstream generative platforms to meet professional journalistic standards, especially in terms of semantic alignment, cultural-symbol accuracy, and structural clarity across varied news scenarios. In contrast, the second experiment demonstrates that by embedding human expertise into key stages of the workflow, it is possible to substantially improve content fidelity, stylistic coherence, and editorial suitability.

This comparison reveals that technical advancement alone cannot guarantee journalistic reliability, a human-AI collaborative approach can enhance both the credibility and communicative value of generated outputs. Taken together, the results suggest that relying on algorithms alone is not enough. Instead, incorporating human expertise into the process makes a real difference and could serve as a practical path forward for producing trustworthy visuals in special coverage. Needless to say, Back-End Data Layer, Task Context Layer and Model Response Layer of AIGC would improve the quality of image generation along with the high-speed development of trending agentic AI and embodied AI. However, the authors believe that the engagement of human is indispensable in producing images by AIGC tools for journalists, enhancing not only the coherence of identity and fidelity of context, but also the trust of publics towards journalism.

### Human - AI Collaboration principles

Building on the diagnostic insights of Experiment 1 and the workflow validation achieved in Experiment 2, this section articulates a three dimensional evaluation framework to guide the verification of AIGC produced news imagery:

CIS-CEA-UPA, where Table3 delineates the corresponding assessment rubric.

- Character Identity Stability (CIS) guarantees that facial features, body posture and attire remain consistent across varied reportage scenes, preventing identity drift (Jia 2024).
- Complementing CIS is Cultural Expression Accuracy (CEA), which gauges the fidelity with which clothing symbols, linguistic signage, background motifs, and action logic align to local journalistic contexts, thereby mitigating cultural mismatch and semantic drift (Chen 2023).
- User/Public Appropriateness (UPA) evaluates composite acceptability in visual realism, emotional propriety, compositional clarity & compliance risk from both audience & editorial perspectives, ensuring that generated images are suitable for formal dissemination (Wang 2022; Gaggioli et al. 2025; Yu, Liu, and Wu 2017).

Dimension	Objective	Key Indicators	Application Scenario
CIS	Identity coherence	<ul style="list-style-type: none"> <li>· Clothing style</li> <li>· Language signage</li> <li>· Background decor</li> <li>· Action-culture logic</li> </ul>	<ul style="list-style-type: none"> <li>· Digital news anchors</li> <li>· Event character rebuild</li> </ul>
CEA	Context fidelity	<ul style="list-style-type: none"> <li>· Clothing style</li> <li>· Language signage</li> <li>· Background decor</li> <li>· Action-culture logic</li> </ul>	<ul style="list-style-type: none"> <li>· International news</li> <li>· Ceremonies</li> <li>· Diplomacy</li> </ul>
UPA	Publication suitability	<ul style="list-style-type: none"> <li>· Visual realism</li> <li>· Emotional propriety</li> <li>· Compositional clarity</li> <li>· Sensitive-content avoidance</li> </ul>	<ul style="list-style-type: none"> <li>· News publishing</li> <li>· Editorial review</li> </ul>

Table 3: User Perception Scores Across Three Dimensions

The three dimensions operate synergistically: CIS supplies identity coherence, CEA secures contextual fidelity, and UPA safeguards communicative suitability. The authors frame distil human-AI collaboration around three intertwined principles. First, domain-reliable into three mutually reinforcing principles. Authoritative data curation: training images should originate chiefly be sourced primarily from accredited news organisations so that factuality, to embed factual accuracy, clear copyright status and stylistic discipline are embedded disciplined style from the outset. Second, vertical-task LoRA refinement: lightweight finetuning is steered (Jia 2024). Second, domain-specific LoRA fine-tuning: lightweight parameter adaptation is guided by a tri-layer annotation logic-journalistic ethics, contextual semantics and communicative aesthetics-so that authenticity, cultural alignment and editorial suitability are simultaneously encoded-encoded simultaneously. Third, dual-loop editorial verification: a closed-loop involving professional photo editors and validation: a closed cycle linking pro-

fessional photo-editors with automated CIS-CEA-UPA metrics ensures that model output is iteratively screened, corrected and redeployed, establishing a co-adaptive cycle iteratively screens, corrects and re-feeds model outputs, fostering co-evolution between human insight and generative efficiency.

## AIGC Journalistic Image Evaluation and Further Application

As demonstrated by the diagnostics of Experiment 1 and the workflow validation of Experiment 2, current AIGC tools tend to fail to meet the dual newsroom demands of authenticity and compliance. Unlike illustration, gaming, or commercial visuals, journalistic imagery must simultaneously satisfy five professional attributes: authenticity of content, controllability of risk, consistency in stylistic tone, alignment between text and image, and editorial auditability.

To address these challenges, this study integrate LoRA fine-tuning with a three-layer annotation protocol. First, the authors source imagery exclusively from authoritative news outlets to ensure domain credibility. Second, the authors focus on high-frequency journalistic domains such as breaking news, sports, public services, and international affairs. Third, the authors engage experienced editors to annotate content from ethical, contextual, and aesthetic perspectives, and conduct double-blind editorial reviews to evaluate output suitability.

By combining human expertise with LoRA-based model tuning, professional standards are embedded directly at the parameter level, forming a replicable pathway for optimized news-oriented generation. The CIS-CEA-UPA evaluation framework, developed through prior experiments, provides a newsroom-centered benchmark for assessing AIGC-generated visuals. Instead of relying on general resolution or aesthetic scores, it emphasizes identity consistency, cultural fidelity, and publication appropriateness – qualities critical for journalism. Empirical testing shows that outputs meeting all three criteria are consistently preferred by editors for publication, whereas those failing any one dimension are more likely to be flagged for revision or rejection.

Moreover, this framework also serves as actionable feedback for iterative LoRA training, supporting the development of a sustainable Human-AI collaboration pipeline. Looking forward, this workflow holds potential for rapid adaptation across diverse news categories. By generalizing the annotation structure and training logic, it becomes feasible to construct a foundation model pre-trained on multi-domain journalistic corpora. Such a model could significantly reduce fine-tuning overhead and enable faster deployment for emerging news scenarios. Example can be also given by a special coverage called 2025, A Beautiful Flower Gift of Xinhua News Agency that achieves wide dissemination and significant impact. The well trained LoRA and workflow raises the image production efficiency by 25 %, more importantly, as shown in Figure 6 it has received an excellent propagation effect for over 100 million reviewers nationwide.



Figure 6: Image production of 2025, *A Beautiful Flower Gift*. Source: generation from AIGC platform.

## Conclusion and Discussion

Moving past the direct comparison of two experiments, this study focuses on what the findings imply for AI-assisted special coverage making in journalism. The results indicate that today’s diffusion-based platforms, although strong in pure generative power, still need well-placed human guidance to satisfy newsroom standards of accuracy and accountability. When professional editors shape prompts, adjust LoRA weights, and review outputs in real time, the images show higher factual fidelity and stronger cultural alignment, a pattern that matches recent observations on responsible media technology (Lewis et al. 2025; Trattner et al. 2022).

The CIS-CEA-UPA framework succeeds in turning broad editorial expectations into clear evaluation criteria. Images that meet identity coherence, contextual fidelity, and publication appropriateness gain higher credibility scores from experts and general audiences alike, a result that mirrors earlier work on interpretability and trust in AI (Doshi-Velez and Kim 2017; Ji et al. 2025). Because these same metrics guide each round of LoRA fine-tuning, they form a practical loop for continual model improvement without undue cost. The pipeline is not a universal substitute for photojournalism. It is contraindicated for events where visual veracity cannot be assured (e.g., legally sensitive reconstructions) or where composites could mislead. In such cases, the SOP mandates abstention or explicit labeling. The CIS-CEA-UPA checklist helps editors make these calls consistently.

Looking ahead, a promising step is to pre-train a lightweight foundation model on a wide, multi-domain news corpus. This approach could reduce the effort required to adapt the system to new beats such as climate reporting, local elections, and data-driven investigations, while maintaining the risk controls shown here. Closer links with newsroom content-management systems would also help edi-

tors spot compliance issues quickly and hand off tasks smoothly to automated tools. In addition, pairing visual evaluation with text-based fact-checking methods (Drolsbach and Pröllochs 2025; Farooq and de Vreese 2025) could create a more complete defense against misinformation. We hope these insights prompt researchers and practitioners to treat human judgment and algorithmic speed as complementary strengths instead of rival approaches.

Ultimately, this human-centered AI approach not only elevates the quality of machine-generated news visuals but also preserves the essential human values in journalism - truth-seeking, contextual understanding, and ethical responsibility - while harnessing the scalability of computational creativity. The pipeline is not a universal solution; in cases involving evidentiary use or sensitive individuals, abstention or disclosure is required. The CIS-CEA-UPA checklist helps editors make these calls consistently. The challenge ahead lies in standardizing these collaborative protocols across diverse media ecosystems without compromising editorial sovereignty or technological innovation. The CIS-CEA-UPA checklist helps editors make these calls consistently. The challenge ahead lies in standardizing these collaborative protocols across diverse media ecosystems without compromising editorial sovereignty or technological innovation.

## Acknowledgments

This research was made possible with support and collaboration of the Xinhua News Agency, originating from its practical workflow needs in the newsroom and aligning with operational standards. The refined workflow and key outcomes form an actionable case study, reflecting academic-industry synergy. The authors thank Haoran Jiang for his help with formatting, bug fixes, and advice to improve the manuscript.

## References

- Byrne, U. 2023. A parochial comment on Midjourney. *International Journal of Architectural Computing*, 21(2): 374–379.
- Cao, Y.; Li, S.; Liu, Y.; Yan, Z.; Dai, Y.; Yu, P. S.; and Sun, L. 2023. A comprehensive survey of ai-generated content (aigc): A history of generative ai from gan to chatgpt. *arXiv preprint arXiv:2303.04226*.
- Chen, X. M. 2023. Shengchengshi AI zai xinwen shengchan zhong de yingyong: Xianshi wenti yu yingdui celue [The application of generative AI in news production: Practical issues and countermeasures]. *Qingnian Jizhe*, (19): 57–59.
- de Lima-Santos, M.-F.; Yeung, W. N.; and Dodds, T. 2025. Guiding the way: a comprehensive examination of AI guidelines in global media. *AI & SOCIETY*, 40(4): 2585–2603.
- Doshi-Velez, F.; and Kim, B. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
- Drolsbach, C.; and Pröllochs, N. 2025. Characterizing AI-Generated Misinformation on Social Media. *arXiv preprint arXiv:2505.10266*.
- Fang, X.; Chen, P.; Wang, M.; and Wang, S. 2025. Examining the role of compression in influencing AI-generated image authenticity. *Scientific Reports*, 15(1): 12192.
- Farooq, A.; and de Vreese, C. 2025. Deciphering authenticity in the age of AI: how AI-generated disinformation images and AI detection tools influence judgements of authenticity. *AI & SOCIETY*, 1–12.
- Gaggioli, A.; Bartolotta, S.; Ubaldi, A.; Gerardini, K.; Sarcinella, E. D.; and Chirico, A. 2025. Extended Creativity: A Conceptual Framework for Understanding Human-AI Creative Relations. *arXiv preprint arXiv:2506.10249*.
- Huang, J.; Chen, C.; Mishra, A.; Kwon, B. C.; Liu, Z.; and Bryan, C. 2024. Asap: Interpretable analysis and summarization of ai-generated image patterns at scale. *arXiv preprint arXiv:2404.02990*.
- Jawaher, A.; Suhuai, L.; and Yuqing, L. 2023. A comprehensive survey on machine learning approaches for fake news detection. *Multimedia Tools and Applications*, 83(17): 51009–51067.
- Ji, Y.; Yan, H.; Lan, J.; Zhu, H.; Wang, W.; Fan, Q.; Zhang, L.; and Zhang, J. 2025. Interpretable and Reliable Detection of AI-Generated Images via Grounded Reasoning in MLLMs. *arXiv preprint arXiv:2506.07045*.
- Jia, P. 2024. Jiyu kejieshi rengong zhinen de shuju anquan fengxian shibie yanjiu [Research on data security risk identification based on explainable artificial intelligence]. *Xinxixi Tongji Gongcheng*, (01): 50–54. (in Chinese).
- Lewis, S. C.; Guzman, A. L.; Schmidt, T. R.; and Lin, B. 2025. Generative AI and its disruptive challenge to journalism: An institutional analysis. *Communication and Change*, 1(1): 9.
- Paik, S.; Bonna, S.; Novozhilova, E.; Gao, G.; Kim, J.; Wijaya, D.; and Betke, M. 2023. The affective nature of ai-generated news images: Impact on visual journalism. In *2023 11th International Conference on Affective Computing and Intelligent Interaction (ACII)*, 1–8. IEEE.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Trattner, C.; Jannach, D.; Motta, E.; Costera Meijer, I.; Diakopoulos, N.; Elahi, M.; Opdahl, A. L.; Tessem, B.; Borch, N. B.; Fjeld, M.; et al. 2022. Responsible media technology and AI: challenges and research directions. *AI and Ethics*, 2(4): 585–594.
- Wang, B.; Zhu, Y.; Chen, L.; Liu, J.; Sun, L.; and Childs, P. 2023. A study of the evaluation metrics for generative images containing combinational creativity. *Ai Edam*, 37: e11.
- Wang, Y., W.; Yang. 2022. Xinmeiti shidai minzhong xinwen de chuanguo yu chuanbo fenxi [Production and dissemination analysis of citizen journalism in the new media era]. *Media*, (23): 94–96. (in Chinese).
- Yao, F.; Li, C.; Nekipelov, D.; Wang, H.; and Xu, H. 2024. Human vs. generative AI in content creation competition: symbiosis or conflict? *arXiv preprint arXiv:2402.15467*.
- Yu, G. M.; Liu, R. Y.; and Wu, C. W. 2017. Suanfa tuijian dui xinxi xiaofei de yingxiang yanjiu: Jiyu renzhi shenjingkexue de shijiao [Research on the impact of algorithmic recommendation on information consumption: A perspective based on cognitive neuroscience]. *Xinwen Zhishi*, (2): 3–6. (in Chinese).
- Zhang, Y.; Tzun, T. T.; Hern, L. W.; and Kawaguchi, K. 2024. Enhancing semantic fidelity in text-to-image synthesis: Attention regulation in diffusion models. In *European Conference on Computer Vision*, 70–86. Springer.
- Zihao, Y. 2025. Ethical Challenges and Countermeasures of Artificial Intelligence Generated Content in the New Media Era. *Journal of Artificial Intelligence Practice*, 8(2).