

Human-AI Interaction Traces as Blackout Poetry: Reframing AI-Supported Writing as Found-Text Creativity

SYEMIN PARK, Department of Industrial Design, KAIST, Republic of Korea

SOOBIN PARK, Department of Industrial Design, KAIST, Republic of Korea

YOUN-KYUNG LIM, Department of Industrial Design, KAIST, Republic of Korea




Fig. 1. A blackout poetry example [7] © Giulia Macchini.

Authors' Contact Information: Syemin Park, syeminpark@kaist.ac.kr, Department of Industrial Design, KAIST, Daejeon, Republic of Korea; Soobin Park, soobinpark@kaist.ac.kr, Department of Industrial Design, KAIST, Daejeon, Republic of Korea; Youn-kyung Lim, younlim@kaist.ac.kr, Department of Industrial Design, KAIST, Daejeon, Republic of Korea.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LLMs offer new creative possibilities for writers but also raise concerns about authenticity and reader trust, particularly when AI involvement is disclosed. Prior research has largely framed this as an issue of transparency and provenance, emphasizing the disclosure of human–AI interaction traces that account for how much the AI wrote and what the human did. Yet such audit-oriented disclosures may risk reducing creative collaboration to quantification and surveillance. In this position paper, we argue for a different lens by exploring how human–AI interaction traces might instead function as expressive artifacts that foreground the meaning-making inherent in human–AI collaboration. Drawing inspiration from blackout poetry, we frame AI-generated text as found material through which writers’ acts of curation and reinterpretation become inscribed atop the AI’s original output. In this way, we suggest that designing interaction traces as aesthetic artifacts may help readers better appreciate and trust writers’ creative contributions in AI-assisted writing.

CCS Concepts: • **Human-centered computing** → **Interaction design theory, concepts and paradigms**.

Additional Key Words and Phrases: creative activity traces, human-AI interaction, creative writing, large language models, found art

ACM Reference Format:

Syemin Park, Soobin Park, and Youn-kyung Lim. 2026. Human-AI Interaction Traces as Blackout Poetry: Reframing AI-Supported Writing as Found-Text Creativity. In *Proceedings of The ACM Association of Computing Machinery (CHI’26)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Position Statement

AI-supported writing has been rapidly transforming how creative work is produced. While these tools can offer writers meaningful benefits, such as facilitating ideation [1, 13] and providing resources for interpretation [9, 16], they also introduce new tensions. In particular, recent studies indicate that AI disclosure often elicits negative perceptions among readers [6], and readers with more negative attitudes toward AI exhibit a greater decline in trust when AI involvement is disclosed [11]. Given that writing can be seen as “*a social act that conveys an author’s intention, effort, and emotion to readers*” [8], writers now face new challenges in communicating AI involvement in ways that maintain reader trust and preserve perceptions of authentic expression [3].

In response to these challenges, HCI researchers have begun exploring systems that make AI involvement in writing more transparent by capturing traces of the AI-assisted writing process. For instance, HaLLMark [4] is a visualization system that tracks AI prompts over time and highlights the associated AI-written or AI-influenced text within a writing interface. In addition, DraftMark [12] utilizes visualization methods inspired by physical metaphors like eraser crumbs to indicate sections in the final text that received writers’ revision effort. While these works offer promising directions, writers have also shown concerns about making AI use completely transparent, as they fear such disclosure may invite harsh criticism or devalue their creative contribution [4, 5]. Consequently, recent scholarship [4, 5] has called for greater attention to the perceptual frameworks that readers bring to work created with AI assistance, encouraging more open-minded interpretations of human–AI collaboration.

Thus, instead of focusing on audit-oriented transparency tools that may invite undue scrutiny, we explore designs of activity traces that foreground writers’ creativity inherent in human–AI collaboration, helping mitigate negative biases in readers’ perceptions of AI-assisted writing. To do this, we take inspiration from blackout poetry. Blackout poetry, as a form of found art [2, 14, 15], has the “goal of using pre-existing or “found” text to create something new. Portions of the

original text are kept, while other parts are drawn over or crossed out.” [10]. As seen in Figure 1, blackout poetry often also incorporates visualization practices to construct new narratives and interpretations from pre-existing materials.

Crucially, in blackout poetry, the creative contribution lies not in writing new sentences from scratch but in the transformative act of curation: deciding which words to keep. This form of subtractive authorship resonates with how writers themselves conceptualize contribution in AI-assisted writing. Specifically, prior research [5] finds that writers often view AI-supported writing as authentic as long as they are actively deciding (e.g., selecting, adopting, removing) what ultimately enters the final text (content gatekeeping), rather than how much text they actually wrote.

Thus, we argue that AI-generated text can be framed as a form of found text, and the human writer’s role can be presented to readers as a creative sculptor of this material. The body of initially generated AI text would then become trace-bearing substrates that inscribe acts of selection (e.g., visualizing blackouts) and appropriation (e.g., abstract drawings that foreground the writer’s interpretive framing of the retained text). In this way, these trace-based representations can function as aesthetic media through which readers may better perceive and appreciate writers’ intention, taste, and meaning-making when working with AI-generated materials.

Finally, we note that certain prerequisites may be necessary to create blackout poetry-inspired traces. This visualization approach depends on AI-generated text functioning as a body of found material that writers can subsequently curate, reshape, and appropriate. When AI output is produced as polished, “perfect” prose, it may leave limited room for meaningful intervention or creative transformation. Recent work has therefore proposed generating AI output in intermediary formats [9, 16], rather than as finalized text, in order to support writers’ reflection, interpretation, and adaptation. Such generative substrates may provide especially fertile ground for blackout-style traces, enabling writers to visibly inscribe their creative and interpretive agency within the evolving text.

To summarize, drawing inspiration from blackout poetry, we propose a reframing of AI-assisted writing in which AI-generated text becomes found material transformed through writers’ subtractive authorship. We hope this work inspires further research that enables human–AI interaction traces to serve as an expressive channel through which writers can communicate their intention, taste, and interpretive agency as part of their creative practice.

Acknowledgments

We are grateful to Forking Room (Min-hyung Kang, Soo Yon Song, and Binna Choi) for inspiring this work. This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2021-NR059056).

References

- [1] Alex Calderwood, John Joon Young Chung, Yuqian Sun, Melissa Roemmele, and Max Kreminski. 2025. Phraselette: A Poet’s Procedural Palette. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS ’25)*. Association for Computing Machinery, New York, NY, USA, 2701–2717. doi:10.1145/3715336.3735832
- [2] Eric Dayton. 1999. *Art and interpretation: an anthology of readings in aesthetics and the philosophy of art*. Broadview Press.
- [3] Katy Ilonka Gero, Tao Long, and Lydia B Chilton. 2023. Social Dynamics of AI Support in Creative Writing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI ’23)*. Association for Computing Machinery, New York, NY, USA, Article 245, 15 pages. doi:10.1145/3544548.3580782
- [4] Md Naimul Hoque, Tasfia Mashiat, Bhavya Ghai, Cecilia D. Shelton, Fanny Chevalier, Kari Kraus, and Niklas Elmqvist. 2024. The HaLLMark Effect: Supporting Provenance and Transparent Use of Large Language Models in Writing with Interactive Visualization. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI ’24)*. Association for Computing Machinery, New York, NY, USA, Article 1045, 15 pages. doi:10.1145/3613904.3641895
- [5] Angel Hsing-Chi Hwang, Q. Vera Liao, Su Lin Blodgett, Alexandra Olteanu, and Adam Trischler. 2025. ‘It was 80% me, 20% AI’: Seeking Authenticity in Co-Writing with Large Language Models. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW122 (May 2025), 41 pages. doi:10.1145/3711020

- [6] Zhuoyan Li, Chen Liang, Jing Peng, and Ming Yin. 2024. How Does the Disclosure of AI Assistance Affect the Perceptions of Writing? *arXiv preprint arXiv:2410.04545* (2024).
- [7] Giulia Macchini. 2026. You have only seen the surface. <https://www.behance.net/gallery/82922035/You-have-only-seen-the-surface>. Accessed: 2026-02-08.
- [8] Hiroki Nakano, Jo Takezawa, Fabrice Matulic, Chi-Lan Yang, and Koji Yatani. 2025. Understanding Reader Perception Shifts upon Disclosure of AI Authorship. *arXiv preprint arXiv:2510.24011* (2025).
- [9] Syemin Park, Soobin Park, and Youn-kyung Lim. 2026. Constella: Supporting Storywriters' Interconnected Character Creation through LLM-based Multi-Agents. *ACM Trans. Comput.-Hum. Interact.* (Feb. 2026). doi:10.1145/3796234 Just Accepted.
- [10] Santa Clara County Library District. 2023. *What is Blackout Poetry?* <https://scclcd.org/blogs/post/what-is-blackout-poetry/> Accessed: 2026-02-08.
- [11] Oliver Schilke and Martin Reimann. 2025. The transparency dilemma: How AI disclosure erodes trust. *Organizational Behavior and Human Decision Processes* 188 (2025), 104405.
- [12] Momin N Siddiqui, Nikki Nasser, Adam Coscia, Roy Pea, and Hari Subramonyam. 2025. DraftMarks: Enhancing Transparency in Human-AI Co-Writing Through Interactive Skeuomorphic Process Traces. *arXiv preprint arXiv:2509.23505* (2025).
- [13] Nikhil Singh, Guillermo Bernal, Daria Savchenko, and Elena L. Glassman. 2023. Where to Hide a Stolen Elephant: Leaps in Creative Writing with Multimodal Machine Intelligence. *ACM Trans. Comput.-Hum. Interact.* 30, 5, Article 68 (Sept. 2023), 57 pages. doi:10.1145/3511599
- [14] Mary Lou Stribling. 1970. *Art from Found Materials: Discarded and Natural*. Crown Publishers, New York, NY.
- [15] Leeana Tankersley. 2009. *Found Art: Discovering Beauty in Foreign Places*. Zondervan, Grand Rapids, MI.
- [16] David Zhou and Sarah Serman. 2024. Ai.llude: Investigating Rewriting AI-Generated Text to Support Creative Expression. In *Proceedings of the 16th Conference on Creativity & Cognition* (Chicago, IL, USA) (C&C '24). Association for Computing Machinery, New York, NY, USA, 241–254. doi:10.1145/3635636.3656187