

The Dynamics of Collective Creativity in Human-AI Hybrid Societies

Shota Shiiku^{1,5}, Raja Marjeh², Manuel Anglada-Tort^{3,5}, Nori Jacoby^{4,5}

¹Graduate School of Science and Technology, Shizuoka University, Hamamatsu, Japan

²Department of Psychology, Princeton University, Princeton, USA

³Department of Psychology, Goldsmiths, University of London, London, UK

⁴Department of Psychology, Cornell University, Ithaca, USA

⁵Max Planck Institute for Empirical Aesthetics, Frankfurt am Main, Germany

Abstract

Generative AI is shaping an increasingly hybrid society, where ideas and cultural artefacts are created both by humans and intelligent machines. Human creativity is influenced in complex, nonlinear ways by the actions of AI-driven agents within their social networks, but these influences are difficult to measure using traditional methods. This study examines how human-AI interactions shape the evolution of collective creation within large-scale social network experiments, where human and AI participants collectively create stories. Participants (either humans or AI) joined 5×5 grid-based networks in which stories were selected, modified, and shared over many iterations. Initially, AI-only networks showed greater creativity (rated by a separate group of human raters) and collective diversity of stories than human-only and human-AI networks. However, over time, hybrid human-AI networks became more diverse in their creations than AI-only networks. In part, this is because AI agents retained little from the original stories, while human-only networks preserved continuity. These findings highlight the value of experimental social networks in understanding human-AI hybrid societies.

Keywords: social networks; collective intelligence; large language models; creativity

Introduction

From cave paintings to symphonies, creativity has defined the human experience – our ability to imagine new possibilities and bring them to life. Today, we stand at a fascinating turning point as generative artificial intelligence (AI) is transforming the creative process by which humans formulate ideas and put them into practice (Epstein et al., 2023). AI models are not merely assisting tools but active partners in the creative process (Collins et al., 2024; Brinkmann et al., 2023), collaborating with writers, musicians, and artists in unprecedented ways. In this hybrid society, networks of interacting humans and intelligent machines constitute complex social systems for which the quality of the collective outcomes cannot be deduced from either human or AI behaviour alone (Tsvetkova, Yasseri, Pescetelli, & Werner, 2024). However, studying how collective creativity emerges in hybrid networks of humans and AI agents remains a key challenge.

While generative AI has been shown to increase the quality and efficiency of routine tasks – such as customer support, programming, and academic writing (Vaccaro, Almaatouq, & Malone, 2024), its impact on human creativity remains poorly understood. Some evidence suggests that human-AI collaboration may enhance individual’s creativity (Doshi & Hauser, 2024; Lee & Chung, 2024), increasing the efficiency and

speed of generating new ideas. In contrast, at the same time aesthetic and cultural biases embedded in generative AI models can limit global diversity, leading to homogenization effects in art and culture (Anderson, Shah, & Kreminski, 2024; Doshi & Hauser, 2024). This research demonstrates the potential impacts of generative AI on human creativity, but we still know little about how good ideas emerge and evolve within human-AI social networks, and the mechanisms that enable effective collaboration between humans and machines.

Our approach introduces a method that is both ecologically valid and open-ended, enabling the study of large-scale social interactions involving hundreds of participants within social networks. Here, we leveraged our recent capability to design experiments that incorporate real human participants within experimental social networks (Marjeh, Anglada-Tort, Griffiths, & Jacoby, 2025). By using iterative storytelling as a creative task, we examine how ideas propagate, transform, and diversify in networks that involve humans, AI, and hybrid human-AI collaborations. This framework allows us to analyze the complex interplay between human and AI agents in dynamic creative processes at a collective level.

Background

Beyond Tools

Generative AI models can now produce high-quality, human-like content across various modalities, including text (Vaswani et al., 2017; Ouyang et al., 2022), images (Ramesh et al., 2022), and speech (Kumar et al., 2023). This has led to the development of a diverse range of AI-driven tools. These include coding assistants like GitHub Copilot, applications in medicine (Jumper et al., 2021), robotics (Murphy, 2019), autonomous vehicles (Badue et al., 2021), and even AI-assisted scientific research (Romera-Paredes, 2024).

Yet, how effective are AI tools? Despite increasing investment in development, research on AI-assisted decision-making presents mixed findings. While some studies indicate that AI enhances performance in domains such as healthcare, customer service, and scientific research (Liu, Lai, & Tan, 2021; Chen, Liao, Wortman Vaughan, & Bansal, 2023), others suggest the opposite (Bansal et al., 2021; Zhang, Liao, & Bellamy, 2020). A recent systematic review and meta-analysis of 106 experimental studies (Vaccaro et al., 2024) confirmed this divide, though, on average, human-AI teams

4941

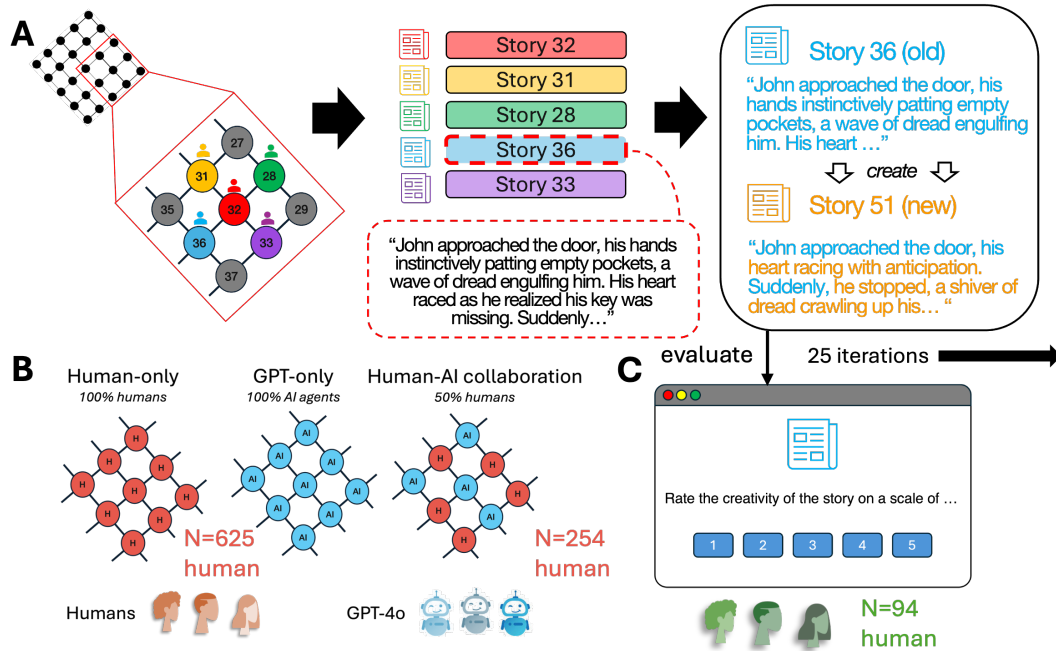


Figure 1: Experimental framework for studying collective creativity. (A) Participants join social networks and engage in a creative writing task where short stories are selected, modified, and transmitted over many iterations (B) We study three network configurations: human-only, AI-only, and human-AI. (C) The creativity of stories is assessed by a separate group of human raters.

performed better than humans working alone.

Rather than viewing AI systems merely as assisting tools, a more promising conceptual shift—drawn from the science of collaborative cognition—is to understand them as “partners in thought”: systems designed to be reasonable, insightful, knowledgeable, and trustworthy, capable of actively thinking and creating with humans (Collins et al., 2024). This perspective moves beyond the instrumental view of AI as passive aids and toward recognizing their potential as cognitive collaborators. As AI models continue to advance and proliferate, they will become increasingly integrated into human interactions and everyday activities, participating in complex, multi-modal problem-solving and contributing to collaborative creative processes.

AI and Collective Dynamics

As AI systems become embedded in everyday human activity as collaborative partners, their impact extends beyond individuals to influence collective behavior and social systems at large (Brinkmann et al., 2023; Collins et al., 2024; Tsvetkova et al., 2024). Thus, it is becoming incredibly important to understand the role of generative AI in shaping collective behavior, focusing on the emergent structures that arise from interactions between multiple individuals and AI agents simultaneously (Tsvetkova et al., 2024).

Recent advances in computational and experimental techniques now make it possible to study how collective behavior emerges in social networks (Centola, 2022; Malone

& Bernstein, 2022). Researchers have examined the effects of network size and topology in artificial social networks composed of human agents, revealing how structural factors influence consensus formation and collective intelligence (Centola & Baronchelli, 2015; Derex, Beugin, Godelle, & Raymond, 2013; Rand, Arbesman, & Christakis, 2011). For instance, Shirado and Christakis (2017) demonstrated that introducing noise into a system—such as simple bot agents—can enhance consensus by pushing the system out of local minima. However, previous studies have focused almost exclusively on simple, one-dimensional tasks, where solutions are predefined within the problem space, such as problem-solving, inference, and decision-making. Creativity, in contrast, requires engagement with open-ended challenges, vast exploration-exploitation spaces, and complex problem-solving environments.

AI and Collective Creativity

Creativity is typically assessed along two core dimensions: novelty and usefulness (Kaufman & Sternberg, 2010; Amabile, 1982). Novelty captures the extent to which an idea deviates from existing conventions or expectations, while usefulness reflects the idea’s practicality, coherence, or relevance, such as whether a short story could plausibly be published as a book. Beyond these individual-level criteria, creativity can also be understood at the collective level, where it reflects the diversity of ideas produced across a group or system (Maher, 2012; SonicRim, 2001; Parjanen, 2012). In this view, a cre-

ative ecosystem is not just one that generates isolated novel ideas, but one that fosters a wide range of distinct and meaningful outputs.

Generative AI can impact creativity in at least two ways. AI can enhance creativity by acting as a cognitive catalyst, suggesting unexpected ideas, combining concepts in novel ways, and expanding the exploratory space available to human creators (Collins et al., 2024; Muller et al., 2024). Conversely, AI can hinder creativity by promoting convergence and homogenization (Bengio et al., 2024; Anderson et al., 2024). Understanding when and how AI enhances versus constrains collective diversity is thus critical for assessing its long-term impact on human cultural production. Doshi and Hauser (2024) demonstrated that while generative AI can enhance individual creativity, it may simultaneously reduce collective diversity, the range and originality of ideas produced by a group. This study provides an important initial look at the collective consequences of AI-driven creativity, but it did not capture how these effects evolve among diverse agents interacting over time within social networks.

Creativity Research

Creativity has been studied extensively in psychology and computer science. In psychology, human creativity has been investigated using psychometric measures developed over decades (Kaufman & Sternberg, 2010). Classic “creativity tests”—such as the Structure of the Intellect divergent production tasks (Guilford, 1967) and the Torrance Tests of Creative Thinking (Torrance, 1966)—remain widely used today. These assessments typically involve language-based tasks, such as listing unusual uses for a common object or generating non-repeating words. More recent advances leverage computational models of semantic networks to address the limitations of these tests. These models allow researchers to quantitatively trace how individuals navigate conceptual spaces, revealing cognitive search strategies associated with creative thought (Beaty & Kenett, 2023). Beyond verbal tasks, some studies have explored non-linguistic forms of creativity—such as visual art production in grid-based environments—to model how people balance exploration and exploitation in open-ended tasks using computational techniques (Kumar et al., 2024; Hart et al., 2017, 2018). However, this research primarily treats creativity at the individual level and does not consider AI as an integral part of the creative process

In computer science, creativity has traditionally been approached from the perspective of computational creativity, a field that explores how machines can be programmed to perform tasks typically associated with human creative behavior (Wiggins & Forth, 2018; Colton & Wiggins, 2012). This includes the development of algorithms that can generate music, visual art, and poetry. Methods range from rule-based systems and evolutionary algorithms to neural networks and large-scale generative models. A central aim is to build systems that exhibit properties of creativity—such as novelty, value, and surprise (Boden, 2004; Lamb, Brown, &

Clarke, 2018). Recent advances in generative AI, particularly large language and image models, have expanded the scope of machine-generated content, making it possible to simulate high-level creative tasks with increasing fluency and realism (Ramesh, Dhariwal, Nichol, Chu, & Chen, 2022; Brown et al., 2020). However, this work often emphasizes output quality rather than the cognitive and social processes underlying human-AI creativity, and tends to focus on individual agents rather than collective dynamics. Integrating insights from psychology with computational methods offers a promising path forward for understanding creativity as a collective process that emerges in networks of interacting humans and machines.

Method

Participants

Participants were recruited online via Prolific¹. All participants were recruited from the UK and identified English as their native language. In total, we recruited 879 human participants and conducted 996 calls to GPT-4o (Hurst et al., 2024), using the version released on September 3, 2024. Participants provided informed consent in accordance with an approved ethics protocol (Max Planck Ethics Council #202142), and were compensated at a rate of £9 per hour.

Network Experiments on Collective Creativity

We conducted large-scale online experiments on collective creativity, where participants (either human or AI) were embedded in a 5×5 social network and engaged in a creative storytelling task. Participants selected and modified stories from their neighbours in the network, and their creations were presented to other participants over 25 iterations (Figure 1A). The same prompt was used for both human and AI agents: *Please creatively elaborate on the story, adding your own details and ideas*. Participants were randomly assigned to one location (node) in the network and only contributed once to the experiment. All networks were initialized with the same story:

“As John reached for his front door, he realized his key was missing. Panic set in as he searched his pockets, but the key was nowhere to be found. Feeling defeated, he slumped against the door, only to hear a jingle from inside—his cat had been playing with the key all along.”

We compared three experimental conditions (Figure 1B): (1) *human-only*, where all nodes were occupied by human participants (N=625), (2) *AI-only*, where all nodes were simulated using GPT-4o (625 calls to the OpenAI API), and (3) *human-AI*, consisting of an equal distribution of human participants (N=254, 50%) and GPT-4o agents (50%). To minimize bias, participants were not informed that any of the stories may have been generated by AI agents.

¹www.prolific.com

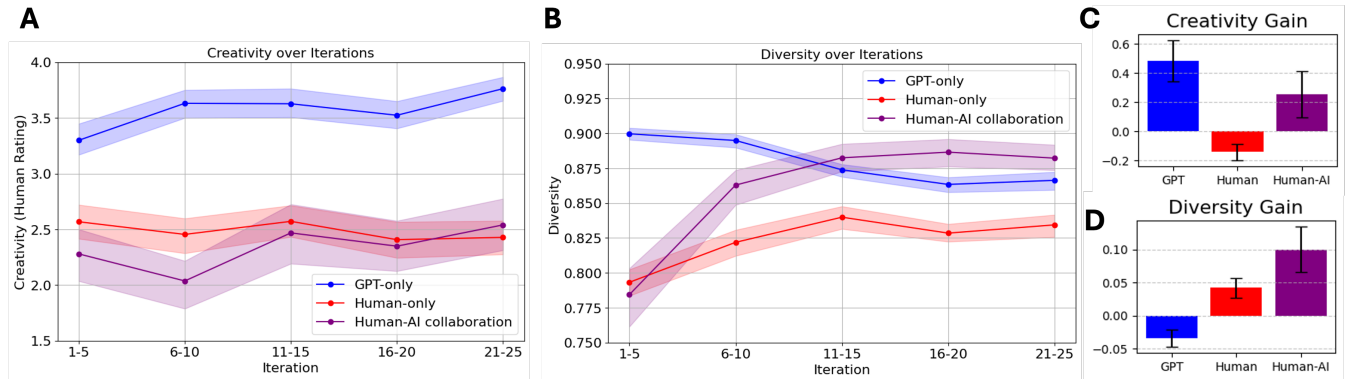


Figure 2: The dynamics of collective creativity. (A) Mean creativity ratings of stories over time, as evaluated by human participants. (B) Diversity of stories (inverse similarity) over time. The horizontal axis represents the 25 iterations, grouped into five sets of five iterations each. Error bars represent one standard deviation, computed across participants. (C-D) Creativity and diversity gain: The improvement in measured creativity and diversity from the first iteration to the last.

All experiments were conducted using PsyNet², a Python-based framework for advanced online psychological experiments (Harrison et al., 2020).

Measures of Creativity

To assess the quality of creations across conditions and iterations, we conducted a validation study with a separate group of 100 human participants. Each participant was presented with a randomized selection of stories from all experimental conditions and iterations and asked to rate their creativity on a 5-point scale, ranging from 1 (not creative at all) to 5 (extremely creative). Each participant evaluated 20 different stories. Crucially, participants were not informed whether stories were created by humans or AI.

In addition to subjective creativity, we measured the diversity of stories by computing their semantic similarity. Specifically for the diversity analysis, we embedded each story using a TF-IDF vectorization approach (Leskovec, Rajaraman, & Ullman, 2020) and calculated the pairwise cosine similarity between all stories within predefined iteration groups. Diversity was operationalized as the inverse of the mean cosine similarity, with lower similarity indicating greater diversity.

Results

Creativity and Diversity

We begin by examining the dynamics of collective creativity across the three experimental conditions. Figure 2A shows the average creativity ratings from the validation study across the three conditions over time. The *AI-only* condition exhibited the highest creativity rating ($M = 3.571$, $SD = 1.026$, 95% CI: [3.491, 3.652]), significantly higher than both *Human-only* and *Human-AI* conditions ($p < .001$). The average creativity ratings in the *Human-only* condition ($M = 2.482$, $SD = 1.026$, 95% CI: [2.389, 2.576]) were similar to the *Human-AI* condition ($M = 2.327$, $SD = 1.159$, 95% CI:

[2.184, 2.576]; $p < .001$). This result demonstrates that GPT easily surpasses humans in simple creative writing tasks, even when evaluated by human raters. Interestingly, AI-only and Human-AI networks consistently improved over time, showing a significantly positive gain by the end of the experiment (GPT: $M = 0.464$, $SD = 0.154$, $p < 0.001$, Human-AI: $M = 0.264$, $SD = 0.175$, $p < 0.001$), but Human-only networks did not improve over time ($M = -0.133$, $SD = 0.067$, $p = 1.000$) (Figure 2C).

Next, we looked at the diversity of collective creations (Figure 2B and D). Here, the *AI-only* condition exhibited the highest diversity ($M = 0.880$, $SD = 0.017$, 95% CI: [0.859, 0.900]), followed by *Human-AI* ($M = 0.860$, $SD = 0.043$, 95% CI: [0.806, 0.913]), and *Human-only*, which showed the lowest diversity ($M = 0.823$, $SD = 0.019$, 95% CI: [0.800, 0.900]).

However, the evolution of collective creativity revealed an intriguing finding. Initially (iteration 1-5), stories in the *AI-only* network were the highest in creativity and diversity, but over iterations, diversity steadily declined, with a drop of $M = -0.034$, $SD = 0.17$, 95% CI: [-0.047, -0.021] (Figure 2D). In contrast, the *Human-AI* network started with lower collective diversity (similar levels to *Human-only*) but exhibited the largest increase over time ($M = 0.098$, $SD = 0.039$, 95% CI: [0.064, 0.131]), ultimately achieving the highest overall diversity score in the final iterations.

Semantic Analysis of Stories

To visualize the difference between human and AI-generated stories, we created a semantic embedding space based on all written stories. Figure 3 shows the UMAP embedding of the stories, visualizing the dynamics between the three conditions. Each point represents a single story, and the proximity of points reflects their semantic similarity in the original high-dimensional space. To compute the high-dimensional space used for the UMAP projection, we first generated semantic embeddings for each story using a pre-trained trans-

²www.psy.net.dev

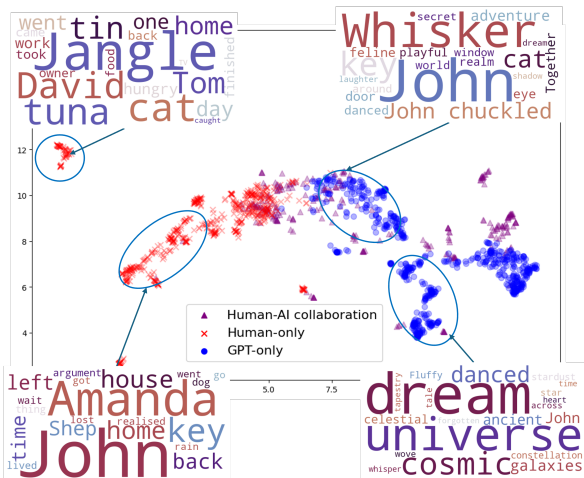


Figure 3: UMAP projection of the shared semantic embedding space, highlighting word clouds for specific clusters.

former model (all-MiniLM-L6-v2; Wang et al., 2020). We then applied UMAP to reduce these embeddings to two dimensions, allowing us to visualize the semantic similarities among the stories.

Three key insights emerge from this analysis. First, the *Human-only* and *AI-only* conditions exhibit distinct clusters, indicating a marked difference in the language and semantics employed in their stories. The word cloud on the bottom right of Figure 3 shows that space-related words such as “universe” and “cosmic” are created when AI are involved in story generation. The word clouds on the left of Figure 3 (largely composed from human data), on the other hand, suggest that story themes remain close to the original seed story (e.g., they share many of the same keywords, such as “John” and “Key”). Second, the *Human-only* condition exhibited additional human names such as “Amanda”, “David”, and “Tom” suggesting potential shifts in the protagonist or narrative content, and indicating a distinct subgroup within the generated stories. Third, in the *Human-AI* condition, humans adopted a compromise in vocabulary and semantics, resulting in narratives that interpolate between the two other conditions, suggesting that humans can be influenced by new ideas generated by AI. These clusters represent emergent “niches” in storytelling styles, characterized by unique interactions between human creativity and AI-driven semantic expansion.

Finally, to understand the shifts in narrative dynamics, we computed the frequency of dominant words in the story characterized by TF-IDF scores. We combined all stories across conditions, calculated the TF-IDF for each term, and sorted them so that the most frequent terms are identified. Figure 4 shows the dynamics of change in the narrative. Each chain shows specific term prevalence over iterations, with circle size indicating the frequency of the term in the iteration. We can see that the longevity of specific keywords vary significantly depending on the condition. For example, certain words such as “dreams”, “danced” and “celestial” only

emerge and persist under the *AI-only* condition. Conversely, words from the original story, including “John”, “cat”, and “keys” persist until the end of the experiment only in the *Human-only* condition. This suggests that humans tend to create new narratives that remain closely aligned with the original storyline, while AI outputs exhibited a unique tendency to converge on certain creative themes, such as space-related narratives, which were consistent across iterations. This convergence, while creative, can potentially indicate difference in how AI agents interpret the task compared with humans.

Discussion

This study explored how the interplay between human and AI creativity in a storytelling task influences collective creative in experimental social networks. We examined three distinct network configurations – human-only, AI-only, and human–AI – to understand whether human-AI collaboration led people towards greater creativity.

Our findings reveal that, from the outset, AI-only networks exhibited the highest levels of creativity and diversity compared to human-only and mixed networks. Over successive iterations, the creativity advantage of AI-only networks remained robust, suggesting that GPT’s ability to generate novel and divergent narratives was particularly well-suited to this task. However, while AI-generated stories initially displayed high diversity, this gradually declined over time. This suggests that although GPT can introduce novel ideas, it also exhibits a form of thematic convergence over time, leading to a reduction in overall diversity. This result complements previous evidence showing that while AI can enhance individual creativity, it may simultaneously reduce collective diversity (Doshi & Hauser, 2024).

In contrast, human-AI networks, which initially demonstrated lower diversity, ended up surpassing the diversity generated by AI-only networks. This shift was primarily driven by the distinct ways in which humans and AI approached the task. While AI frequently disregarded core narrative elements in favor of novel inventions, human participants tended to retain key story components, such as character identities (e.g., maintaining the protagonist, John) and objects (e.g., ‘keys’). The interplay between humans and AI in mixed networks enabled a dynamic balance between stability and novelty, enabling them to produce increasingly diverse outputs over time. This suggests that collaboration between humans and AI can harness the strengths of both, leading to richer and more diverse creative outcomes.

A crucial consideration in interpreting our results is the nature of the creative task itself. The storytelling task used in this study aligns with those traditionally employed in human creativity research (Kaufman & Sternberg, 2010; Beaty & Kenett, 2023). However, the task proved to be relatively easy for a large language model like GPT, which consistently outperformed humans alone both in creativity and diversity. This does not necessarily imply that AI is intrinsically more




Figure 4: Term dynamics by condition: Words are plotted along the horizontal axis, and generations are plotted along the vertical axis. A circle at position (x,y) indicates that the word x was used in a story at iteration y . A line denotes that the corresponding word was used by successive iterations. The size of the circle represents the frequency of the word’s appearance in the same iteration.

creative than humans; rather, it highlights the limitations of using tasks designed for human cognition as benchmarks for AI creativity. AI’s advantage may stem from its ability to rapidly generate highly novel yet semantically coherent text, a skill optimized by its training on vast datasets (Burton et al., 2024; Brinkmann et al., 2023). Conversely, human creativity and creative writing is influenced by memory constraints, cognitive biases, and adherence to implicit narrative conventions (Hayes, 2012).

Moving forward, creativity research should move beyond simple tasks that may be well-suited for assessing individual human creativity but fall short in evaluating AI capabilities or capturing the dynamics of human–AI synergy. For example, using multimodal creative tasks that do not only rely on language, such as music creation (Anglada-Tort, Harrison, Lee, & Jacoby, 2023) and drawings (Kumar et al., 2024; Hart et al., 2017, 2018). Moreover, future studies should focus on designing complex tasks that challenge both humans and AI in meaningful ways, allowing for a more nuanced understanding of collective creativity and underlying mechanisms.

Limitations and conclusion

Our study aimed to simulate social-naturalistic dynamics within experimental societies. However, our approach diverges from how ecologically valid social interactions unfold in the real world. Unlike our experiments, where participants engaged only once, individuals in real-world settings continuously interact with each other over time. Additionally, real-world exchanges are far more complex than the constrained task of creating simple stories used in our experiment. Another key limitation is the static structure of our simulated social network, which used a grid-based topology. Real social networks, such as groups of friends or work relationships, tend to have a modular structure with irregular patterns and

evolving connections, allowing individuals to actively shape their interactions by forming and adjusting their social ties rather than being confined to static connections. Finally, our paradigm avoids any direct communication between participants and masks agent information, such as whether your partner is a human or an AI. Prior research suggests that awareness of interacting with an AI system may alter individual behavior, which in turn could influence collective dynamics (Tsvetkova et al., 2024).

In addition, our study has several technical limitations. First, key parameters of the AI models—such as model architecture, temperature settings, and prompt design—may have influenced the results. We intend to systematically examine the impact of these variables in future work. Second, while we employed large language models (LLMs) for story generation and selection, more advanced AI agents now enable richer forms of interaction with both the environment and other agents—an aspect not addressed in the current study (Google Developers Blog, 2025). Lastly, our validation approach was relatively limited. We relied on a single measure of creativity, and participants were not provided with an explicit definition of the term. Future studies could incorporate a more comprehensive set of creativity metrics and a more robust validation framework.

Despite these limitations, none of these constraints are inherent to our methodological approach. Our framework can be extended to incorporate more ecologically valid dynamics, including persistent interactions, evolving social structures, and greater participant agency. Our approach can also be extended to non-verbal, open-ended creative task, such as music and visual art. While our study represents only an initial exploration of hybrid experimental social networks, it highlights the vast potential of this approach for future research in cognitive and computer science.

References

- Amabile, T. M. (1982). Social psychology of creativity: A consensual assessment technique. *Journal of personality and social psychology*, 43(5), 997.
- Anderson, B. R., Shah, J. H., & Kreminski, M. (2024). Homogenization effects of large language models on human creative ideation. In *Proceedings of the 16th conference on creativity & cognition* (pp. 413–425).
- Anglada-Tort, M., Harrison, P. M., Lee, H., & Jacoby, N. (2023). Large-scale iterated singing experiments reveal oral transmission mechanisms underlying music evolution. *Current Biology*, 33(8), 1472–1486.
- Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., ... Weld, D. (2021). Does the whole exceed its parts? the effect of ai explanations on complementary team performance. In *Proceedings of the 2021 chi conference on human factors in computing systems* (pp. 1–16).
- Beaty, R. E., & Kenett, Y. N. (2023). Associative thinking at the core of creativity. *Trends in cognitive sciences*, 27(7), 671–683.
- Bengio, Y., Hinton, G., Yao, A., Song, D., Abbeel, P., Darrell, T., ... others (2024). Managing extreme ai risks amid rapid progress. *Science*, 384(6698), 842–845.
- Boden, M. A. (2004). *The creative mind: Myths and mechanisms*. Routledge.
- Brinkmann, L., Baumann, F., Bonnefon, J.-F., Derex, M., Müller, T. F., Nussberger, A.-M., ... others (2023). Machine culture. *Nature Human Behaviour*, 7(11), 1855–1868.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... others (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877–1901.
- Burton, J. W., Lopez-Lopez, E., Hechtlinger, S., Rahwan, Z., Aeschbach, S., Bakker, M. A., ... others (2024). How large language models can reshape collective intelligence. *Nature Human Behaviour*, 8(9), 1643–1655.
- Centola, D. (2022). The network science of collective intelligence. *Trends in Cognitive Sciences*, 26(11), 923–941.
- Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, 112(7), 1989–1994.
- Chen, V., Liao, Q. V., Wortman Vaughan, J., & Bansal, G. (2023). Understanding the role of human intuition on reliance in human-ai decision-making with explanations. *Proceedings of the ACM on Human-computer Interaction*, 7(CSCW2), 1–32.
- Collins, K. M., Sucholutsky, I., Bhatt, U., Chandra, K., Wong, L., Lee, M., ... others (2024). Building machines that learn and think with people. *Nature Human Behaviour*, 8(10), 1851–1863.
- Colton, S., & Wiggins, G. A. (2012). Computational creativity: The final frontier? In *Ecai 2012* (pp. 21–26). IOS Press.
- Derex, M., Beugin, M.-P., Godelle, B., & Raymond, M. (2013). Experimental evidence for the influence of group size on cultural complexity. *Nature*, 503(7476), 389–391.
- Doshi, A. R., & Hauser, O. P. (2024). Generative ai enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), eadn5290.
- Epstein, Z., Hertzmann, A., of Human Creativity, I., Akten, M., Farid, H., Fjeld, J., ... others (2023). Art and the science of generative ai. *Science*, 380(6650), 1110–1111.
- Google Developers Blog. (2025). *Announcing the agent2agent protocol (a2a)*.
- Guilford, J. P. (1967). *The nature of human intelligence*. New York: Macgraw Hill.
- Harrison, P., Marjeh, R., Adolphi, F., van Rijn, P., Anglada-Tort, M., Tchernichovski, O., ... Jacoby, N. (2020). Gibbs sampling with people. *Advances in neural information processing systems*, 33, 10659–10671.
- Hart, Y., Goldberg, H., Striem-Amit, E., Mayo, A. E., Noy, L., & Alon, U. (2018). Creative exploration as a scale-invariant search on a meaning landscape. *Nature Communications*, 9(1), 5411.
- Hart, Y., Mayo, A. E., Mayo, R., Rozenkrantz, L., Tandler, A., Alon, U., & Noy, L. (2017). Creative foraging: An experimental paradigm for studying exploration and discovery. *PloS one*, 12(8), e0182133.
- Hayes, J. R. (2012). Modeling and remodeling writing. *Written Communication*, 29(3), 369–388.
- Hurst, A., Lerer, A., Goucher, A. P., Perelman, A., Ramesh, A., Clark, A., ... others (2024). GPT-4o system card. *arXiv preprint arXiv:2410.21276*.
- Kaufman, J. C., & Sternberg, R. J. (2010). *The cambridge handbook of creativity*. Cambridge University Press.
- Kumar, S., Marjeh, R., Zhang, B., Campbell, D., Hu, M. Y., Bhatt, U., ... Griffiths, T. (2024). Comparing abstraction in humans and machines using multimodal serial reproduction. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 46).
- Lamb, C., Brown, D. G., & Clarke, C. L. (2018). Evaluating computational creativity: An interdisciplinary tutorial. *ACM Computing Surveys (CSUR)*, 51(2), 1–34.
- Lee, B. C., & Chung, J. (2024). An empirical investigation of the impact of chatgpt on creativity. *Nature Human Behaviour*, 8(10), 1906–1914.
- Leskovec, J., Rajaraman, A., & Ullman, J. D. (2020). *Mining of massive data sets*. Cambridge university press.
- Liu, H., Lai, V., & Tan, C. (2021). Understanding the effect of out-of-distribution examples and interactive explanations on human-ai decision making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2), 1–45.
- Maher, M. L. (2012). Computational and collective creativity: Who’s being creative? In *Iccc* (pp. 67–71).
- Malone, T. W., & Bernstein, M. (2022). *Handbook of collective intelligence*. MIT Press.
- Marjeh, R., Anglada-Tort, M., Griffiths, T. L., & Jacoby, N. (2025). Characterizing the interaction of cultural evolu-

- tion mechanisms in experimental social networks. *arXiv preprint arXiv:2502.12847*.
- Muller, M., Houde, S., Gonzalez, G., Brimijoin, K., Ross, S. I., Moran, D. A. S., & Weisz, J. D. (2024). Group brainstorming with an ai agent: Creating and selecting ideas. In *International conference on computational creativity*.
- Parjanen, S. (2012). Experiencing creativity in the organization: From individual creativity to collective creativity. *Interdisciplinary Journal of Information, Knowledge & Management*, 7.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2), 3.
- Rand, D. G., Arbesman, S., & Christakis, N. A. (2011). Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences*, 108(48), 19193–19198.
- Shirado, H., & Christakis, N. A. (2017). Locally noisy autonomous agents improve global human coordination in network experiments. *Nature*, 545(7654), 370–374.
- SonicRim, L. S. (2001). Collective creativity. *Design*, 6(3), 1–6.
- Torrance, E. P. (1966). Torrance tests of creative thinking. *Educational and psychological measurement*.
- Tsvetkova, M., Yasserli, T., Pescetelli, N., & Werner, T. (2024). A new sociology of humans and machines. *Nature Human Behaviour*, 8(10), 1864–1876.
- Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and ai are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 1–11.
- Wang, W., Wei, F., Dong, L., Bao, H., Yang, N., & Zhou, M. (2020). Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33, 5776–5788.
- Wiggins, G. A., & Forth, J. (2018). Computational creativity and live algorithms.
- Zhang, Y., Liao, Q. V., & Bellamy, R. K. (2020). Effect of confidence and explanation on accuracy and trust calibration in ai-assisted decision making. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 295–305).