Beyond Artifact Generation: Supporting the Articulation of Authors' Creative Motives in Professional Manga Creation

Nami Ogawa*
CyberAgent
Tokyo, Japan
nami_ogawa_pub@cyberagent.co.jp

Hiromu Yakura*
Max-Planck Institute for Human Development
Berlin, Germany
yakura@mpib-berlin.mpg.de

Abstract

Recent LLM-based creativity support tools primarily optimize artifact generation, such as producing plots, drafts, or story continuations to accelerate ideation. However, we argue that, for professional creators, sustainable authorship depends not only on rapid artifact production but rather on articulating a coherent creative motive: a tacit core of values, themes, and commitments that anchor long-term identity. Through formative interviews with 11 early-career manga artists, we found that highly concrete AI outputs can prematurely constrain exploration before creators' internal motives are fully formed. Based on the insight, we introduce a design shift from artifact-generation support to motive-articulation support. Grounded in Variation Theory and co-adaptive sensemaking, we built a prototype that scaffolds reflective discernment. By iteratively selecting and rejecting comparable candidates from AI models, creators externalize implicit constraints and progressively crystallize their creative motives. Here, we envision generative AI functioning not only as a content producer but also as a cognitive instrument for reflection-in-action, thereby offering a complementary direction for Tools for Thought.

CCS Concepts

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

Keywords

Creativity Support Tools, Manga, Professional Use

ACM Reference Format:
Nami Ogawa and Hiromu Yakura. . Beyond Artifact Generation: Supporting the Articulation of Authors' Creative Motives in Professional Manga Creation. In *Proceedings of ACM CHI 2026 Workshop on Tools for Thought*. ACM, New York, NY, USA, 4 pages.

1 Introduction

Generative AI has rapidly entered creative practice. In particular, recent advancements in Large Language Models (LLMs) have led to a proliferation of creativity support tools (CSTs) capable of generating complex artifacts, such as plots, characters, and entire scenes [3, 13, 21]. Yet, across writing and ideation tools, existing systems commonly focus on either accelerating brainstorming, producing drafts or partial artifacts, or optimizing immediate task performance. Such focus does not fully address how professional creators develop durable creative identities, particularly in fields where "what to express" is as central as "how to express it."

In the Japanese manga industry, early-career artists often publish self-contained one-shot stories (*yomikiri*), typically 30–100 pages long, as a proving ground. Beyond demonstrating technical skill, these one-shots help artists express consistent originality and authorial intent across independent works, signaling recurring themes and narrative sensibilities that make a creator's identity legible to editors and publishers. This developmental task is often supported through sustained dialogue with a skilled editor who helps the artist surface what they truly want to express.

Our formative work with early-career manga artists suggests that existing AI tools, especially those that produce highly concrete story continuations or ready-to-use plots, can feel generic or misaligned with an artist's internal vision. In some cases, concrete AI outputs can prematurely impose patterns before the artist's own motive has fully formed, contributing to early fixation. These observations resonate with broader concerns that early AI intervention may reduce diversity and increase homogenization in creative outputs [4, 18], and that collaboration with generative AI can introduce trade-offs in motivation and agency [14, 19].

Motivated by this gap, we propose a shift in AI design strategy: from artifact-generation support to motive-articulation support. Instead of treating ideation as the search for better artifacts, we focus on scaffolding a creator's sensemaking of their own creative motive, a tacit core of tastes, values, and "want-to-draw" commitments that anchors decision-making during creation. We introduce a prototype tool that acts as a reflective catalyst: starting from a vague core seed (e.g., "a protagonist whose efforts never pay off"), the system generates eight divergent but comparable candidate elements (e.g., "a world where effort is visualized" and "a protagonist who sees being unrewarded as natural"). The artist chooses what resonates, which becomes integrated into the core; the loop repeats. Over time, the core concept crystallizes, and the accumulated trace of selections externalizes the artist's implicit constraints, themes, and sensibilities. We offer this design strategy as a case for reorienting AI creativity support around the development of authorial identity.

2 Related Work

2.1 AI for Exploration and Reflection in Creative Domains

Recent systems use LLMs to support exploration of design spaces and narrative possibilities. For example, Luminate proposes structured generation and exploration to support human–AI co-creation [17]. More broadly, creativity support systems have explored example galleries [2, 11], design space navigation [1, 8, 20], and structured ideation workflows [9]. These systems often aim to help users

*Both authors contributed equally to this research.

consider alternatives and iterate, sometimes at multiple levels of abstraction.

Our prototype also enables exploration of design spaces, but rather than broad exploration of artifacts (e.g., story drafts), we emphasize element-level variation (e.g., emotional stance, thematic framing, or relational context) as a scaffold to articulate a creator's motive. This design choice is motivated by formative observations that highly concrete story text can "over-specify" what the user has not committed to, causing mismatch and fixation.

Kreminski and Mateas [10] proposed *process aesthetics* as a lens for examining process-focused creativity support tools, identifying reflection as a key aesthetic goal and presenting design patterns that elicit reflective experiences in autotelic creative systems. Our work extends this perspective by targeting a specific reflective outcome, i.e., the articulation of creative motives, and grounding the interaction design in Variation Theory.

2.2 Theoretical Foundations: Variation Theory

Prior HCI work has identified a key challenge in human-AI co-creation: intent elicitation and alignment, where users struggle to communicate what they want. To reduce this gap, Gmeiner et al. proposed structured representations of intent, called intent tags, as small, editable units that support granular micro-prompting and meta-intent elicitation [6]. Importantly, they report that the system supported users to figure out what they want while working on the task [6]. That is, users' needs and intents are often not completely clear upfront and emerge through iterative reflection and refinement, aligning with Schön's notion of "reflection-in-action" [6, 15]. We therefore position our work as motive-oriented alignment: rather than generating polished story drafts, we scaffold creators' reflection.

To enable this, we draw on a cognitive theory known as Variation Theory. Variation Theory originates in human concept learning: discernment arises when learners experience systematic variation in critical features against a stable background [12]. In interactive machine teaching, MOCHA uses Variation Theory to generate counterfactual examples and supports sensemaking through structured comparisons to help users refine concept definitions by comparing boundary cases, referring to the process as co-adaptive learning [5]. We draw an analogy between concept learning in machine teaching and motive articulation in creative ideation. In both cases, the user's concept boundary is initially ambiguous. By keeping an evolving "core" invariant while varying specific dimensions through multiple comparable candidates, we help creators notice which differences matter to them, surfacing critical anchors and supporting co-adaptive sensemaking.

2.3 Potential Risk of Generative AI on Creativity

Empirical work suggests that generative AI can raise individual creativity in some contexts while reducing the collective diversity of produced content [4] and that using LLM before independent ideation may lead to over-reliance on AI-generated content, suppressing divergent thinking and even undermining individuals' creative self-efficacy [14]. Another line of work reports that human–generative AI collaboration can enhance task performance but undermine intrinsic motivation, with implications for autonomy and ownership [19]. These tensions motivate designs that preserve or strengthen intrinsic drivers and support reflective agency rather than passive acceptance of AI outputs.

Our approach explicitly aims to mitigate homogenization by anchoring ideation to the artist's intrinsic creative motive. However, we treat this as an open design challenge: even systematic variation can reproduce stereotypes unless carefully designed, and reflection scaffolds can still influence direction.

3 Formative Study

We conducted semi-structured formative interviews with 11 early-career manga artists, i.e., professionals who work under assigned editors toward publication but have not yet established a serialized work, to understand how they ideate and how existing AI tools fit (or fail to fit) their workflows. Interviews included hands-on trials of screenwriting-theory-based GPT applications (e.g., plot generation following Blake Snyder's Beat Sheet [16]) and an existing creativity support tool (i.e., Luminate [17]).

Finding 1: Creative workflow as discovering the creative anchor through backwards narrative construction. Multiple artists referred to a workflow in which they first establish a "creative anchor," such as a scene they want to depict or a theme, and then construct the narrative through backwards calculation. Many creators adopt an approach of logically filling in the necessary elements to establish what they want to draw most (i.e., the creative anchor). Some creators mentioned that fragmentary scenes they want to draw appear to them visually first, implying that the core of what they want to express cannot necessarily be verbalized at the initial stage. Furthermore, a creator stated, "Even if [the theme] wavers, it feels more like I am coming to understand it myself. It is like arriving at what I really wanted to draw," suggesting that the creative process is a journey of discovering the core anchor of the work.

Finding 2: Highly concrete AI outputs can feel generic and misaligned. When artists tried tools that generate detailed story continuations or plot outlines, they reported that the results often did not match their sensibilities. They described outputs as "averaged" or "not interesting as a story." Importantly, even when outputs were presented as options, their concreteness could implicitly constrain exploration; one artist noted that presenting a result in a completed form "could nip the bud" of their own thoughts. Another artist remarked that generating a full story draft makes the output feel "too much like a final deliverable," often eliciting a detached reaction of simply "Oh, okay." This artist noted that AI's pinpoint output tends to miss their intended scope and expressed a preference for AI to intervene at a stage "one step prior to a draft [...] to provide support in approaching [the scope I want]."

Finding 3: Artists seek a reflective partner rather than an idea generator. Artists frequently emphasized the value of reflection, often realized through introspective dialogue or editorial conversations. In these interactions, editors do not simply propose ideas but ask questions to help the artist clarify their intent. This suggests that an effective AI tool in this domain should function less as an idea generator and more as a reflective partner that scaffolds sensemaking. For instance, several artists emphasized the need for AI to

prompt introspection rather than merely provide solutions. One artist specifically desired an AI that prompts reflection by asking why the creator chose a particular direction to help them delve into their deep psychology and verbalize their true intent. Similarly, another artist described a process of "response battles" where arguing against AI's shallow interpretations forces them to articulate the specific appeal of the work. This dialectical interaction, where creators refine their vision by rejecting misaligned suggestions, helps define the contours of their core concept.

Design implications. We therefore target an interaction that helps externalize and refine creative motives by supporting iterative specification: instead of generating a single best plot, the tool generates systematic variation around a core seed, enabling discernment through comparison.

4 Design Strategy: From Artifact-Generation Support to Motive-Articulation Support

We propose a design strategy shift from *artifact-generation support* to *motive-articulation support*. In artifact-generation support, AI primarily produces candidate artifacts (ideas, plots, drafts) to be selected and refined. Success is measured by immediate output quality and speed. In motive-articulation support, AI primarily scaffolds the creator's reflective articulation of their intent, motive, and constraints. We refer to this internal, often tacit core as *creative motives*, including tastes and values. The key difference is that we do not aim to generate final artifacts per se, but rather to improve the creator's self-understanding of their motive and constraints. In this framing, AI's primary role is not to contribute content as a co-author, but to generate systematic variations that help the human articulate intent. In other words, we treat AI outputs as perceptual and cognitive instruments for reflection-in-action [15], rather than as candidate drafts to be accepted. This shift is grounded in three theoretical and interactional commitments.

Variation Theory as a Scaffold for Discernment. Variation Theory posits that people learn to discern concepts by experiencing systematic variation in critical aspects while holding other aspects stable [12]. Translating this to creative ideation, the creator's current core seed serves as the invariant background. The system proposes multiple candidate additions that vary along limited dimensions (e.g., emotional stance, thematic framing, or relational context). By comparing alternatives, creators can notice resonances and mismatches with their vague inner imagery, surfacing critical anchors. While the core evolves across rounds, each round treats the current core as invariant within that comparison set, enabling contrastive discernment over the proposed variations.

Co-Adaptive Sensemaking Through Iterative Feedback. We frame the interaction as co-adaptive sensemaking [5]: the system adapts to the creator's selections, and the creator simultaneously learns and refines their own concept boundary ("what is my story really about?"). This is analogous to interactive machine teaching, where both teacher and learner co-adapt through feedback on concrete examples.

Recursive Divergence and Convergence. The interaction implements a recurring cycle: diverge by presenting diverse candidates (including conventional and deliberately contrasting options), converge through selection and rejection, and repeat. Each cycle shifts attention to what remains undecided and gradually clarifies the motive. This aligns with recursive divergence–convergence paradigms in creative ideation tools [7].

5 Prototype: Reflective Catalyst for Manga Ideation

We built a prototype tool that implements the design strategy. Given that our target users (e.g., manga artists) are not necessarily technically proficient, we prioritized ease of adoption and integration into everyday creative practice with minimal setup. Accordingly, we implemented the prototype as a Custom GPT within ChatGPT. The prompt is designed to include the following:

- Comparison set, not recommendation: Candidates are displayed as a set to compare, avoiding a ranked list that implies a correct answer.
- Limited concreteness: Candidates are kept at the *element* level rather than full story text, reducing the risk of over-specification.
- Conventional + contrasting mix: Each set intentionally includes both familiar, typical items and contrasting, nontypical items to surface boundaries.

Users experience the following interaction loop:

(1) Core seed input: The artist enters a vague seed concept (one to two sentences).
(2) Eight candidate variations: The system generates eight divergent but comparable candidate elements. Each candidate is short (one to two lines) and designed to vary one or a few dimensions while preserving the core context.
(3) Selection (or rejection): The artist selects the candidate that resonates most (or discards all and requests a re-roll). The selection is not framed as "the best idea," but as "the most resonant direction."
(4) Core concept update: The selected element is integrated into the core concept. The updated core becomes the invariant background for the next round.
(5) Trace accumulation: The system records the sequence of selections and rejected alternatives as an interpretable trace of preferences and constraints.
(6) Motive reflection (optional): At checkpoints, the system summarizes inferred motives (e.g., "You repeatedly choose constraints emphasizing unfairness and persistence") and prompts the artist to confirm, revise, or reject the inference.

As an example scenario, when an artist starts with: "A protagonist whose efforts never pay off," the system returns variations such as: "A world where effort is visualized to everyone" and "a protagonist who sees being unrewarded as natural." The artist selects one emphasizing structural injustice. In subsequent rounds, the tool continues to vary dimensions, e.g., the scale of injustice, the emotional stance (bitterness vs. quiet resilience), the genre framing (satire vs. tragedy), or the relationship to community. Over time, the artist may realize that the story is less about "effort" and more about "the pain of being unseen," or "the ethics of continuing without recognition." The trace makes visible that the artist consistently

rejects "revenge" framings and repeatedly chooses "ambiguous ending" framings, externalizing critical anchors that the artist can later use to generate alternative story premises or evaluate drafts.

6 Conclusion

We presented a design strategy shift for AI creativity support in professional manga ideation: from artifact-generation support to motive-articulation support. Grounded in Variation Theory and co-adaptive sensemaking, our prototype uses LLM-generated systematic variation as a comparison set to scaffold reflection-in-action and externalize creative motives. In the framing of Tools for Thought, this approach is designed to serve a dual role: protecting creative cognition from premature fixation imposed by overly concrete AI outputs, and augmenting reflective agency by making tacit motives visible and revisable. We hope this work contributes a complementary path for Tools for Thought: designing generative AI not only to produce content faster, but to help creators understand and commit to what they genuinely want to express. We see several open questions suited for workshop discussion: how to evaluate motive articulation quality beyond artifact-level metrics, how to extend motive-oriented support to other creative domains, and how to reconcile the reflective depth this approach affords with the short-term productivity demands that professional creators face.

References

[1] Sheng-Fen Chien and Ulrich Flemming. 2002. Design space navigation in generative design systems. *Automation in Construction* 11, 1 (2002), 1–22.

[2] Daeun Choi, Sumin Hong, Jeongeon Park, John Joon Young Chung, and Juho Kim. 2024. CreativeConnect: Supporting Reference Recombination for Graphic Design Ideation with Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 1055, 1–25 pages.

[3] John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. TaleBrush: Sketching stories with generative pretrained language models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 209, 19 pages.

[4] Anil R Doshi and Oliver P Hauser. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Sci. Adv.* 10 (2024), eadn5290. Issue 28.

[5] Simret Araya Gebreegziabher, Yukun Yang, Elena L Glassman, and Toby Jia-Jun Li. 2025. Supporting co-adaptive machine teaching through human concept learning and cognitive theories. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 533, 18 pages.

[6] Frederic Gmeiner, Nicolai Marquardt, Michael Bentley, Hugo Romat, Michel Pahud, David Brown, Asta Roseway, Nikolas Martelaro, Kenneth Holstein, Ken Hinckley, and Nathalie Riche. 2025. Intent tagging: Exploring micro-prompting interactions for supporting granular human-GenAI co-creation workflows. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 531, 31 pages.

[7] Taewook Kim, Matthew Kay, Yuqian Sun, Melissa Roemmele, Max Kreminski, and John Joon Young Chung. 2025. Scaffolding recursive divergence and convergence in story ideation. *arXiv* 2507.03307 (2025), 17 pages.

[8] Janin Koch, Andrés Lucero, Lena Hegemann, and Antti Oulasvirta. 2019. May AI?: Design Ideation with Cooperative Contextual Bandits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 633, 12 pages.

[9] Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, and Wendy E. Mackay. 2020. ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships. *Proc. ACM Hum. Comput. Interact.* 4, CSCW (2020), 045:1–045:27.

[10] M Kreminski and Michael Mateas. 2021. Reflective Creators. In *Proceedings of the International Conference on Computational Creativity*. Association for Computational Creativity, Coimbra, Portugal, 309–318.

[11] Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen I. Brafman, and Scott R. Klemmer. 2010. Designing with interactive example galleries. In *Proceedings of the 2010 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2257–2266.

[12] Ference Marton. 2014. *Necessary conditions of learning*. Routledge, London, UK.

[13] Piotr Mirowski, Kory W Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-writing screenplays and theatre scripts with language models: Evaluation by industry professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 355, 34 pages.

[14] Peinuan Qin, Chi-Lan Yang, Jingshu Li, Jing Wen, and Yi-Chieh Lee. 2025. Timing matters: How using LLMs at different timings influences writers' perceptions and ideation outcomes in AI-assisted ideation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 25, 16 pages.

[15] Donald A Schoen. 1983. *The reflective practitioner: How professionals think in action*. Routledge, London, UK.

[16] Blake Snyder. 2005. *Save the cat! The last book on screenwriting you'll ever need*. Michael Wiese Productions, Studio City, CA.

[17] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminate: Structured generation and exploration of design space with large language models for human-AI co-creation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 644, 26 pages.

[18] Samangi Wadinambiarachchi, Ryan M Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The effects of generative AI on design fixation and divergent thinking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 380, 18 pages.

[19] Suqing Wu, Yukun Liu, Mengqi Ruan, Siyu Chen, and Xiao-Yun Xie. 2025. Human-generative AI collaboration enhances task performance but undermines human's intrinsic motivation. *Sci. Rep.* 15 (2025), 15105. Issue 1.

[20] Hiromu Yakura and Masataka Goto. 2023. IteraTTA: An Interface for Exploring Both Text Prompts and Audio Priors in Generating Music With Text-to-Audio Models. In *Proceedings of the 24th International Society for Music Information Retrieval Conference*. International Society for Music Information Retrieval, British Columbia, Canada, 129–137.

[21] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: Story writing with large language models. In *Proceedings of the 27th ACM International Conference on Intelligent User Interfaces*. ACM, New York, NY, USA, 841–852.