

Intent Elicitation in Mixed-Initiative Co-Creativity

Max Kreminski^{1,*}, John Joon Young Chung¹

¹Midjourney, San Francisco, California, United States

Abstract

Mixed-initiative creative interfaces (MICIs) aim to support user creativity by supplying users with an artificially intelligent creative collaborator, but AI-based creativity support systems can struggle to understand what users want and why. To advance discussion of how MICIs can better make sense of user intent, we present three preliminary design patterns for *intent elicitation*: Ask Don't Guess, Refine via Examples, and Gauge Creative Momentum. We also briefly discuss Patchwork: a mixed-initiative collaborative storytelling canvas powered by generative AI technologies, in which we use basic implementations of these design patterns to elicit user intent.

Keywords

mixed-initiative co-creativity, creativity support tools, design patterns

1. Introduction

One branch of research in *creativity support tools* [1, 2, 3]—software systems intended to support human creativity—involves the development of *mixed-initiative creative interfaces* (MICIs) [4, 5]. MICIs employ mixed-initiative interaction techniques [6] to create a sense of collaboration between the user and an artificially intelligent creative partner.

Although MICIs have shown promise in creative domains as wide-ranging as game design [7, 8], sketching [9], music [10], and creative writing [11, 12, 13, 14, 15], and although they have become more popular with recent advances in open-ended generative AI, they are still limited by the difficulty of understanding user intent. MICIs can fail to understand user intent for at least three reasons:

- **Underexpression.** Users often do not express their intent fully. For instance, user interactions with many of the most popular modern creative AI systems tend to begin with the user passing the system a short, almost necessarily underspecified text prompt. The low information content of these prompts relative to the complexity of the desired output artifacts results in the system needing to extrapolate considerably to produce a response, often failing to respect unexpressed aspects of the user's intent [16] and perhaps homogenizing creative outcomes [17]. This can result in considerable frustration around prompting, especially for inexperienced users [18].

Joint Proceedings of the ACM IUI Workshops 2024, March 18-21, 2024, Greenville, South Carolina, USA

*Corresponding author.

✉ mkreminski@midjourney.com (M. Kreminski); jchung@midjourney.com (J. J. Y. Chung)

🌐 <https://mkremins.github.io> (M. Kreminski); <https://johnr0.github.io> (J. J. Y. Chung)

🆔 0009-0002-6268-4033 (M. Kreminski)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

- **Fixation.** In multi-turn interactions, more information about user intent can be inferred through observation of successive actions by the user—but it can be difficult to determine from actions alone what a user is “going for” creatively. In particular, the observability of high-level intent can be limited by a form of creative fixation [19] known as the “XY problem” [20], in which the user’s actions are all directed at a single narrow way of approaching their broader goal. When a MICI cannot infer a reasonable higher-level intent from lower-level actions taken by a user who is creatively stuck, it may itself become stuck in the same rut.
- **Uncertainty.** In creative contexts, users themselves are often (perhaps even necessarily) unaware of at least some aspects of their own intent [19]: the inherently surprising nature of creative solutions [21] implies that users must discover details of their intent through “reflective conversation with the materials of a design situation” [22], rather than arriving at the creative interaction with an intent already fully formed in their mind. When a user does not *know* their own intent, they cannot express it to a MICI even if they want to.

Combined, these problems around user intent make it difficult for creative AI systems—including MICIs—to consistently satisfy user needs. Furthermore, because MICIs cannot readily distinguish between situations in which the user has *underexpressed* a pre-existing intent and situations in which the user is *uncertain* about important parts of their intent, we believe that these difficulties may only be fully resolved via a channel of active *metacommunication* between user and system about user intent. In other words, to successfully “draw out” user intent in a sufficiently wide range of situations, both user and system must be able to express and resolve uncertainties about where the creative interaction is meant to go.

We aim to address these difficulties via the active elicitation of user intent. In this paper, we describe three preliminary MICI design patterns that we have begun to employ for intent elicitation. These patterns—Ask Don’t Guess, Refine via Examples, and Gauge Creative Momentum—have been implemented in Patchwork, an in-development collaborative storytelling canvas built around generative AI technologies. We first briefly address related work and introduce Patchwork. Then we discuss how each of our proposed design patterns helps to mitigate the three intent-related problems in MICI design identified above; a variety of considerations involved in the implementation of each pattern; and why we feel it makes sense to recast the problem of intent elicitation in MICIs as one of intent co-construction between both human and AI creative actors.

2. Related Work

Past work in allowing users of MICIs to explicitly express their intent has resulted in both a library of preliminary design patterns for encouraging user reflection on intent [23] and a design space of possible communication types between users and MICI systems [24].

In addition, a few existing MICIs have attempted to provide a channel for explicit metacommunication between user and system about creative intent. The storytelling MICI Loose Ends [12] represents a direct attempt to explore how creative intent can be made more explicit, including via an explicit intent negotiation side-channel (the “storytelling goals” pane); proactive inference of intent by the MICI based on users’ creative choices; visualization to users of inferred intent;

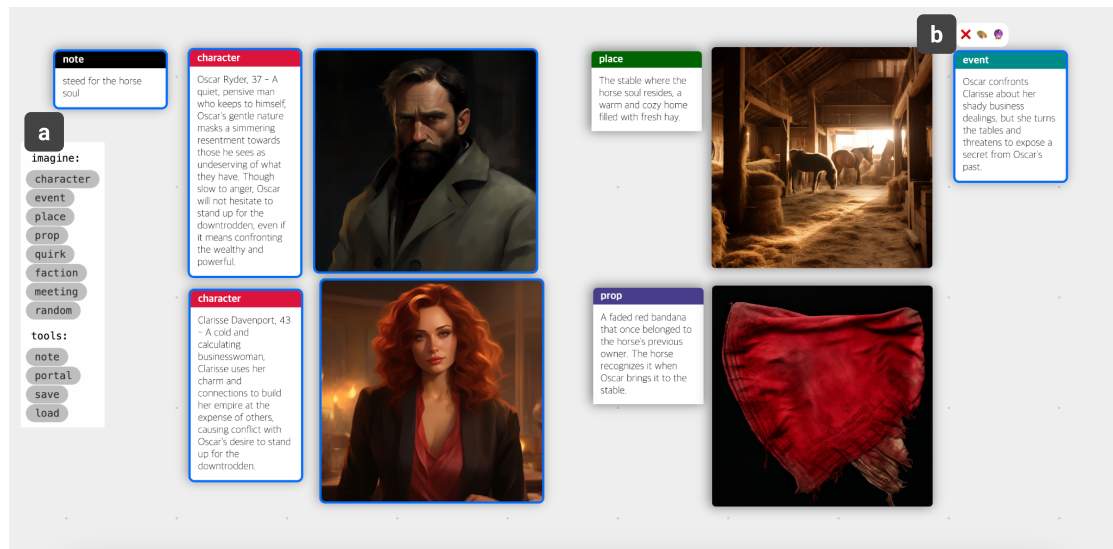


Figure 1: The Patchwork user interface. Below marker **a** is the “toolbox”, a sidebar containing a variety of tools that users can employ; many are dedicated to prompting the AI to introduce new storyworld concepts. Near marker **b** is the “toolbar”, a set of additional contextual tools that are available on groups of selected “scraps” (the card-like elements scattered around the canvas). Scraps contain text or images relevant to storyworld concepts, such as characters, places, and events that take place within the world. Selected scraps are outlined in blue.

implicit closed-ended intent clarification questions, in which users can select which of several ambiguous inferred intents they are actually pursuing; and cautionary visualization to users when specific creative choices under consideration would undercut an explicit or inferred intent. However, even Loose Ends stops short of explicitly asking users open-ended questions about their intent to help further draw it out.

Patchwork, our testbed for the design patterns discussed in this paper, is an infinite-canvas creativity support tool backed by recently introduced generative AI techniques. This puts it in the same category as several other infinite-canvas AI-based creativity support tools, including Luminate [25] and the tools introduced by Kim et al. [26].

3. Patchwork: A Collaborative Storytelling MICI

Patchwork is an in-development mixed-initiative co-creative storytelling canvas that aims to facilitate playful social creativity between groups of users, each of whom contributes text and images to an emerging storyworld. On the AI side, Patchwork provides language model-powered features for generating concepts that might fit into the users’ shared storyworld (including characters, factions, places, props, and events that might happen during the story), as well as text-to-image features that can be used to visualize these concepts.

Patchwork presents storyworld concepts on an infinite canvas that is freely scrollable and zoomable by the user. The user can freely position storyworld concepts as movable “scraps”, as in Figure 1. As a collaborative tool, Patchwork shows the cursors of other users in real time.

The user can ask Patchwork to generate story concepts of specific types with the buttons under the `imagine` section of Figure 1a, by clicking the button first and then specifying the position where the scrap should be generated. Once the textual concept scrap is generated, the user can also manually edit the content of the scrap.

By default, when generating concepts, Patchwork takes all other scraps on the canvas into account, assembling a language model prompt that includes information about these scraps alongside an appropriate concept-generation instruction and retrieving the generated concept from the resulting output text.¹ Alternatively, if the user wants generation to only be influenced by certain scraps, they can first select a specific set of scraps for generation to consider. Selected scraps are outlined in blue, as seen in Figure 1.

The user can also manually create textual notes with the `note` button under the `tools` section. If the user wants to change the type of the note, they can click on the type name of the scrap, which allows users to circulate the type within all allowed types.

Notes can be used as open-ended questions or instructions from the user to the generator via the `crystal ball` button on the toolbar that appears near selected note scraps. Another toolbar button (the `paint palette` button) can be used to generate images that visualize the selected concept or concepts. If a single textual scrap is selected for image generation, Patchwork considers the text of the selected scrap as the prompt for the image generation model. However, if multiple are selected, Patchwork uses the concatenation of the selected scraps' textual contents as the input to the image generation model. Image scraps also contribute back to the generation of further textual content: images are first described as text by an image-description model and then fed into constructed language model prompts as part of storyworld context.

Other functions of Patchwork include tools for establishing links between storyworlds, saving and loading storyworld content, and importing images from outside sources. Controls for some of these features are visible in the “tools” section of the left-hand sidebar in Figure 1a.

Storytelling and worldbuilding, like other creative activities, can be difficult, in part due to issues like the fear of the blank canvas and writer's block. In addition to these difficulties, creative collaboration also requires a degree of negotiation between multiple different creators around what the storyworld ought to contain. This makes Patchwork (a collaborative storytelling and worldbuilding environment) an ideal testbed for the introduction of MICI features intended to assist with the elicitation and negotiation of explicitly expressed user intent.

4. Preliminary Design Patterns for Intent Elicitation

Patchwork development has been guided by ongoing informal testing of successive versions of the system. The earliest versions of Patchwork placed a greater emphasis on giving users maximum direct control over the output of a purely reactive system, but as we expanded user testing both within and outside of the development team, we began to observe that users rarely arrived at the interaction with a sophisticated preconceived idea of the world they wanted to create. Instead, they discovered their intent progressively, and often responded best to features of the system that proactively attempted to draw their intent out into a more elaborate form. The design patterns we discuss here were developed in response to these observations.

¹In this regard, Patchwork follows the *UI transducers* architecture pattern [27].

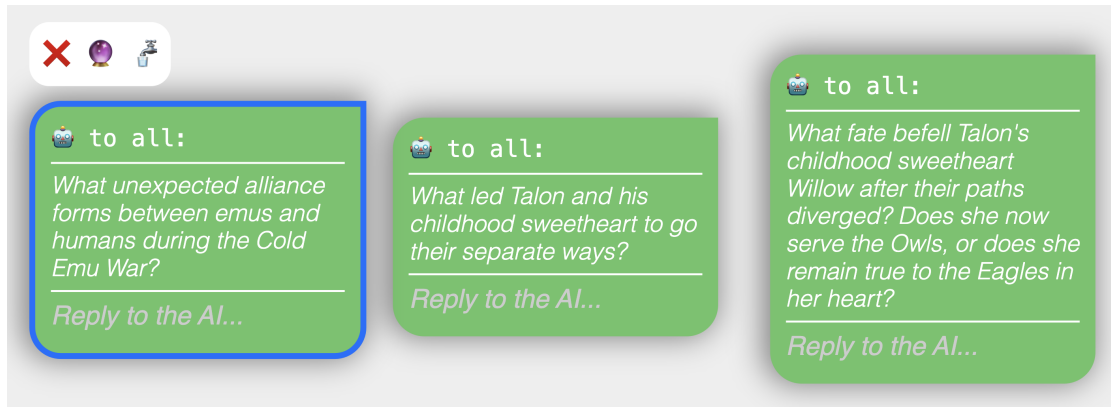


Figure 2: Several AI-generated questions from a Patchwork storyworld about the “Cold Emu War”. Questions are formatted like text messages from the AI, and can be repositioned on the canvas or deleted like any other scrap. They can also be answered via an open-ended text field, or “reversed” (triggering the AI to provide one or more answers to the question itself) via controls that appear above a selected question scrap (like the one on the left).

4.1. Ask Don’t Guess

Purely inference-based approaches to determining user intent can suffer from the inability to distinguish between user *underexpression*, *fixation*, and *uncertainty* of intent. To address this problem in Patchwork, we have designed the system to explicitly and proactively ask the user to clarify aspects of their intent.

Specifically, we use a language model to formulate and ask open-ended questions about aspects of the emerging storyworld that have not yet been fully specified by users (Figure 2). These questions vary widely: we have seen the AI ask about topics as varied as unexpressed aspects of specific characters (“What common object does the Mountebank leave behind at the scene of his crimes, and what does it represent?”); setting details that have not yet been addressed (“Why did the first battle in the Cold Emu War take place on a volcano?”); and the eventual fate of key characters, factions, or relationships (“Does Meowcifer eventually retire from cat pirate life, or does he die on the high seas?”), among many others.

Users can provide immediate responses to these questions in the form of open-ended text. Alternatively, if the user cannot think of a good answer for the question due to uncertainty or fixation, they can “turn the question around” by asking the MICI to either provide a set of potential responses or just answer the question itself. This spectrum of potential user responses allows the user to fluidly change the amount of control they are exerting over the creative process. Similar mixed-initiative interactions have been characterized as experientially similar to riding a horse: tugging the reins to give specific directions when needed, but otherwise allowing the horse to walk where it will, trusting that it will not walk into danger [28]. Question reversal also gives users a way to preserve creative momentum when they neither know nor care about the answer to a MICI-posed question.

To help reduce user uncertainty, generated questions can be *leading*—i.e., they can imply something about a yet-to-be-specified aspect of the storyworld in how they are phrased. For

instance, a question like “Why don’t Snow and Abilene get along?” might be used to imply an antagonistic relationship between two specific characters, even when nothing about the nature of these characters’ relationship has been stated before. The use of leading questions to support creativity is adapted from similar practices in tabletop roleplaying games [29]. Because leading questions can imply assertions about the storyworld that users do not actually want to adopt, users must be allowed to reject the premise implicit in a question; currently in Patchwork, users can do this by deleting the scrap that poses the question from the canvas.

Altogether, this pattern predominantly mitigates user **underexpression** of intent by prompting users to flesh out aspects of their intent that the MICI cannot otherwise determine. The leading nature of some generated questions may also help the user out of **fixation** or **uncertainty** by suggesting possibilities that they had not formerly considered—though it may also introduce new forms of fixation, if users do not feel free to reject premises of leading questions generated by the system.

4.2. Refine via Examples

Examples of creative artifacts that satisfy user intent can be difficult to extrapolate from, in part because a single example both says “too much” and “too little”: it contains many incidental details that should not constrain the creative interaction, as well as insufficient information about how the example fits into a larger creative vision. However, in the context of refining an explicitly stated intent, examples can be powerful for rapidly improving the MICI’s understanding of how the user might want a vague high-level intent to be fulfilled. Consequently, we make use of examples primarily for intent refinement in Patchwork.

Specifically, when the user asks the MICI to execute a concept generation task, they can optionally put it into a “brainstorming mode” that causes the tool to begin generating variations on the requested concept automatically, providing examples of different ways that the user’s request could be fulfilled (for instance, different possible descriptions of the same named character). The user’s implicit acceptance (via preservation) of some of these variants, and explicit rejection (via deletion) of others, is used to refine the MICI’s understanding of the user’s intent for the relevant concept.

We have also begun to experiment with a “grab bag” approach to early-stage tone-setting for a newly created Patchwork storyworld, in which users are rapidly presented with many examples of scraps that the world might include: not just entity descriptions and fully rendered images, but also short keywords and keyphrases, character names based on different corpuses, swatches of color, and so on. The user’s acceptance and rejection of items from this rapid-fire jumble of randomly selected evocative detritus can help to quickly establish an overall vibe for the storyworld, playing the same role as early-game negotiation of a *palette* in tabletop story-making games such as *Microscope* (i.e., collective player decision-making about what things the world they are creating together should and should not include) [30].

This pattern primarily contributes to the mitigation of user **uncertainty** by providing them with concrete examples of how their intent could be further developed. Secondarily, it can also help mitigate **underexpression** of intent by convincing users to provide additional information about their intent in a relatively low-effort (but commensurately low-information) way. The effect of this pattern on user **fixation** may be more indeterminate, since users may seize on

details of high-fidelity examples and then fail to consider potential alternative ways of fleshing out the same storyworld entity.

4.3. Gauge Creative Momentum

When proactively intervening in a human-computer interaction (as we do in Patchwork via question-asking), it is important for interventions to feel well-timed and relevant. Poorly timed or irrelevant-feeling interruptions may be experienced by users as distracting or annoying, leading to a “Clippy effect” in which users come to resent proactive AI assistants [31].

Consequently, MICIs that ask users open-ended questions to elicit intent need some sort of model to determine when, and about what topics, questions should be asked. This model can be rules-based or learned, with our initial attempt at such a model being a rules-based operationalization of the concept of *creative momentum* [32]. Essentially, Patchwork periodically checks the ratio of unanswered AI-generated questions to other scraps of creative material on the canvas, and asks a new question only if the users have collectively introduced sufficient new material relative to the number of questions that the AI has asked so far. Topics of questions are chosen by prompting the language model to ask questions about unspecified aspects of the storyworld specifically.

In the future, this rules-based approach to gauging creative momentum could potentially be extended, including by monitoring user activity to determine whether each user has recently taken any significant creative actions (such as creating, editing, or deleting scraps). Users who appear active but have not recently taken significant actions might then be prompted with an open-ended question to help get them creatively unstuck. However, creative momentum may not always be straightforwardly correlated with MICI-observable activity. In particular, we have observed user states of both “active stuckness” (in which users continue making edits to scraps, but without altering the content of the storyworld in any meaningful way) and “inactive flow” (in which users stop editing content directly within Patchwork for a while, but continue actively developing their creative intentions while idly panning around the canvas). This complicates the use of activity data.

Alternatively, to make better use of a variety of complex user activity inputs, it might also be possible to train a model to gauge interruption timing more generally, drawing data on interruption aptness from user engagement (or lack thereof) with MICI-generated questions. Inspiration could be taken here from recent work on automatically determining when to show suggestions in language model-assisted programming [33]. Future work on gauging creative momentum might also derive useful concepts or methods from prior empirical work on understanding the *interaction dynamics* of co-creativity [34].

Allowing a system to gauge when it should intervene (just as a human co-creative partner would) is a crucial element of giving the system true initiative. However, this runs counter to most recent work on the design of MICIs, which tend to provide users with creative assistance either passively at all times or specifically when the user requests it. This represents a departure from early research on mixed-initiative interaction in general, which frequently discusses the issue of determining when to intervene—see, e.g., Hearst et al. [6]. We believe that this aspect of creative collaboration should be given more attention in present-day research on mixed-initiative co-creativity.

Altogether, the gauging of creative momentum primarily helps to mitigate user **fixation** and **uncertainty**: by timing interventions to occur when users appear to have low creative momentum, a MICI can support users in becoming creatively unstuck by refocusing their attention on a different aspect of what they are trying to create. To a lesser extent, this pattern may also help reduce **underexpression** of user intent by making it more likely that intent-elicitation interruptions arrive during a period of user readiness to engage.

5. Conclusion

The three design patterns we presented in this paper have already shown promise in informal user testing of our in-development collaborative storytelling MICI, Patchwork, despite only basic implementations of these patterns being present in the current version of the system. More elaborate implementations of these patterns, as discussed above, could further support the elicitation of user intent in a wide variety of co-creative contexts.

However, many open questions remain around how best to implement these patterns, particularly with regard to gauging creative momentum as a means of timing the system’s proactive interventions. In the future, we hope to conduct a larger-scale formal user study of Patchwork, with an eye to gauging how the current implementations of these intent elicitation patterns can be improved.

Prior discussion of intent elicitation in artificial intelligence has focused mostly on the inference or discovery of a user’s already-constructed intent [35]. But in our view, at least in a creative context, intent elicitation must also include providing scaffolding for intent elaboration, refinement, or development: the “reflective conversation” that characterizes ideation in design situations [22]. When a MICI asks the user an explicit question about their intent, the user may answer the question by drawing on their pre-constructed understanding of what they want, or they may adapt or expand their sense of what they want to include an answer to the question posed; either way, user intent is “drawn out” into a progressively more explicit form. Compared to mere discovery of pre-existing intent, the resulting process of intent co-construction by both human and AI creative actors takes more complete advantage of the responsiveness to open-ended design situations that emerging AI technologies can provide.

Acknowledgments

Thanks to the many people at Midjourney whose infrastructural work enabled this research, and in particular to Red Kalab.

References

- [1] B. Shneiderman, Creativity support tools: accelerating discovery and innovation, *Communications of the ACM* 50 (2007) 20–32.
- [2] J. Frich, L. MacDonald Vermeulen, C. Remy, M. M. Biskjaer, P. Dalsgaard, Mapping the landscape of creativity support tools in HCI, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019.

- [3] J. J. Y. Chung, S. He, E. Adar, The intersection of users, roles, interactions, and technologies in creativity support tools, in: *Designing Interactive Systems Conference*, 2021, pp. 1817–1833.
- [4] S. Deterding, J. Hook, R. Fiebrink, M. Gillies, J. Gow, M. Akten, G. Smith, A. Liapis, K. Compton, Mixed-initiative creative interfaces, in: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, 2017, pp. 628–635.
- [5] A. Liapis, G. N. Yannakakis, C. Alexopoulos, P. Lopes, Can computers foster human users' creativity? Theory and praxis of mixed initiative co-creativity, *Digital Culture & Education* 8 (2016) 136–153.
- [6] M. A. Hearst, J. E. Allen, C. I. Guinn, E. Horvitz, Mixed-initiative interaction, *IEEE Intelligent Systems and their Applications* 14 (1999) 14–23.
- [7] A. Liapis, G. N. Yannakakis, J. Togelius, *Sentient Sketchbook: computer-assisted game level authoring* (2013).
- [8] M. Kreminski, M. Dickinson, J. Osborn, A. Summerville, M. Mateas, N. Wardrip-Fruin, Germinate: A mixed-initiative casual creator for rhetorical games, in: *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 16, 2020, pp. 102–108.
- [9] T. Lawton, F. J. Ibarrola, D. Ventura, K. Grace, Drawing with Reframer: Emergence and control in co-creative AI, in: *Proceedings of the 28th International Conference on Intelligent User Interfaces*, 2023, pp. 264–277.
- [10] N. J. Thelle, P. Pasquier, *Spire Muse: A virtual musical partner for creative brainstorming*, in: *NIME 2021*, 2021.
- [11] K. I. Gero, L. B. Chilton, *Metaphoria: An algorithmic companion for metaphor creation*, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019.
- [12] M. Kreminski, M. Dickinson, N. Wardrip-Fruin, M. Mateas, *Loose Ends: a mixed-initiative creative interface for playful storytelling*, in: *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 18, 2022, pp. 120–128.
- [13] J. J. Y. Chung, W. Kim, K. M. Yoo, H. Lee, E. Adar, M. Chang, *TaleBrush: Sketching stories with generative pretrained language models*, in: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 2022.
- [14] M. Kreminski, I. Karth, M. Mateas, N. Wardrip-Fruin, *Evaluating mixed-initiative creative interfaces via expressive range coverage analysis.*, in: *IUI Workshops*, 2022, pp. 34–45.
- [15] N. Singh, G. Bernal, D. Savchenko, E. L. Glassman, *Where to hide a stolen elephant: Leaps in creative writing with multimodal machine intelligence*, *ACM Transactions on Computer-Human Interaction* 30 (2023).
- [16] M. Chang, S. Druga, A. J. Fiannaca, P. Vergani, C. Kulkarni, C. J. Cai, M. Terry, *The prompt artists*, in: *Proceedings of the 15th Conference on Creativity and Cognition*, 2023, pp. 75–87.
- [17] B. R. Anderson, J. H. Shah, M. Kreminski, *Homogenization effects of large language models on human creative ideation*, *arXiv preprint arXiv:2402.01536* (2024).
- [18] J. Zamfirescu-Pereira, R. Y. Wong, B. Hartmann, Q. Yang, *Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts*, in: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023.
- [19] A. T. Purcell, J. S. Gero, *Design and other types of fixation*, *Design Studies* 17 (1996)

363–383.

- [20] Wikipedia contributors, XY problem, https://en.wikipedia.org/wiki/XY_problem, 2023. Accessed: 2024-01-18.
- [21] M. A. Boden, Creativity in a nutshell, *Think* 5 (2007) 83–96.
- [22] D. A. Schön, Designing as reflective conversation with the materials of a design situation, *Knowledge-based systems* 5 (1992) 3–14.
- [23] M. Kreminski, M. Mateas, Reflective creators, in: *International Conference on Computational Creativity*, 2021, pp. 309–318.
- [24] Z. Lin, U. Ehsan, R. Agarwal, S. Dani, V. Vashishth, M. Riedl, Beyond prompts: Exploring the design space of mixed-initiative co-creativity systems, in: *Proceedings of the Fourteenth International Conference on Computational Creativity*, 2023, pp. 64–73.
- [25] S. Suh, M. Chen, B. Min, T. J.-J. Li, H. Xia, Structured generation and exploration of design space with large language models for human-AI co-creation, *arXiv preprint arXiv:2310.12953* (2023).
- [26] T. S. Kim, Y. Lee, M. Chang, J. Kim, Cells, generators, and lenses: Design framework for object-oriented interaction with large language models, in: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 2023.
- [27] B. Hartmann, Expanding the design space of generative AI-powered applications, Spotlight Talk, *UIST Architecting Novel Interactions With Generative AI Models Workshop*. Slides: https://docs.google.com/presentation/d/1J6C4axHOn0YBgY_t2HnYHGhMhmbHvjFXvRMRc0Np02A/view, 2023.
- [28] F. O. Flemisch, C. A. Adams, S. R. Conway, K. H. Goodrich, M. T. Palmer, P. C. Schutte, The H-Metaphor as a guideline for vehicle automation and interaction, Technical Report, National Aeronautics and Space Administration, 2003.
- [29] StackExchange contributors, What game made leading questions a GM technique?, <https://rpg.stackexchange.com/questions/55121/what-game-made-leading-questions-a-gm-technique>, 2015. Accessed: 2024-01-19.
- [30] M. Ralya, Steal this mechanic: Microscope’s yes/no list, <https://gnomestew.com/steal-this-mechanic-microscopes-yesno-list/>, 2013. Accessed: 2024-01-19.
- [31] J. Xiao, J. Stasko, R. Catrambone, An empirical study of the effect of agent competence on user performance and perception, in: *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 2004, pp. 178–185.
- [32] D. Zhou, S. Serman, Creative struggle: Arguing for the value of difficulty in supporting ownership and self-expression in creative writing, in: *The Second Workshop on Intelligent and Interactive Writing Assistants (In2Writing)*, 2023.
- [33] H. Mozannar, G. Bansal, A. Fourney, E. Horvitz, When to show a suggestion? Integrating human feedback in AI-assisted programming, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.
- [34] N. Davis, C.-P. Hsiao, K. Y. Singh, B. Lin, B. Magerko, Creative sense-making: Quantifying interaction dynamics in co-creation, in: *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*, 2017, pp. 356–366.
- [35] A. Amos-Binks, D. Dannenhauer, G. Brewer, R. E. Cardona-Rivera, Anticipatory thinking: A new frontier for automated planning, in: *Eighth Annual Conference on Advances in Cognitive Systems*, 2020.