Structuring Prompt Construction as a Tool for Thought

Rifat Mehreen Amin
LMU Munich
Munich, Germany
rifat.amin@ifi.lmu.de

Oliver Hans Kühle
LMU Munich
Munich, Germany
o.kuehle@campus.lmu.de

Daniel Buschek
University of Bayreuth
Bayreuth, Germany
daniel.buschek@uni-bayreuth.de

Andreas Butz
LMU Munich
Munich, Germany
butz@ifi.lmu.de

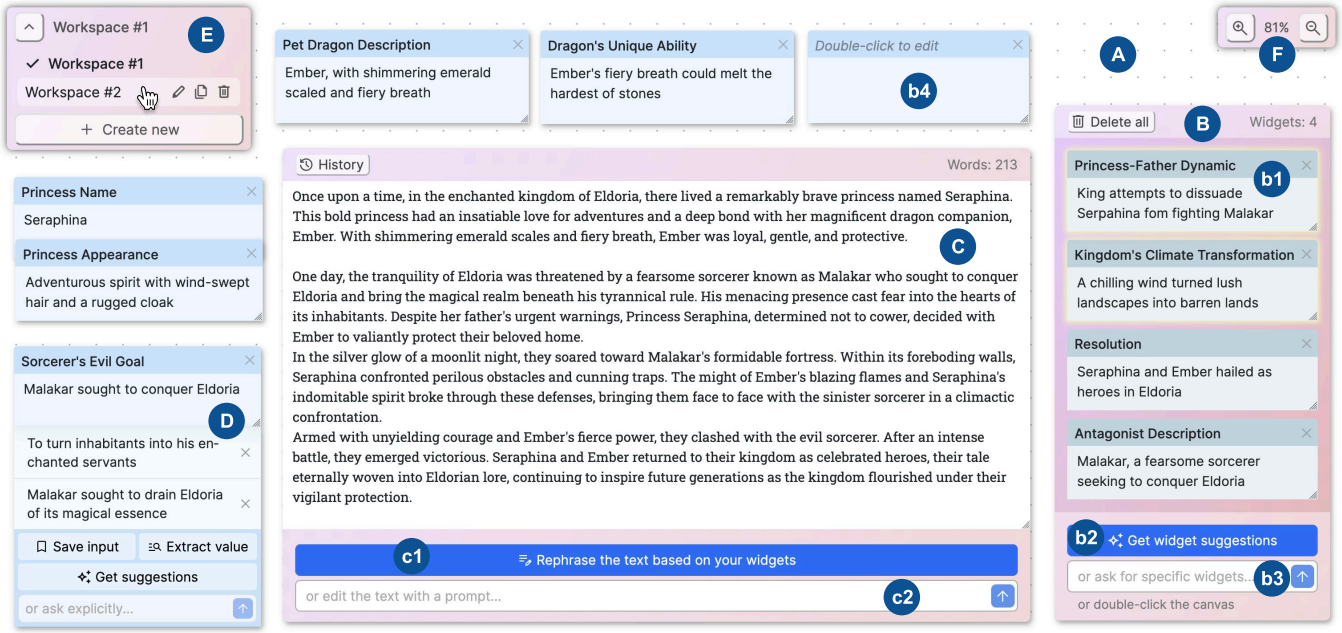

Figure 1: User interface of PromptCanvas. (A) Canvas-like workspace where users can place and freely organize widgets to create a customized environment. (B) Widget panel. (b1) Example of widgets created through system suggestions. (b2) Button to get widget suggestions from the system. (b3) Field for entering prompts to create multiple widgets of a specific theme. (b4) Example of an empty widget created by double-clicking at empty space. (C) Text editor and output of text generation. (c1) Button to rephrase the text based on the widgets on the canvas (light blue). (c2) Field to provide prompts for generating text. (D) Example of an opened widget with suggested values for customization. (E) Menu bar for creating, renaming, duplicating, or deleting a canvas. (F) Panel with alternative zoom level controls.

Abstract

We present PromptCanvas, a Tool for Thought-oriented design strategy that reframes prompt construction as an explicit cognitive activity. PromptCanvas combines a canvas-based workspace with prompt-construction widgets that externalizes parts of a prompt as manipulable, persistent artifacts. The canvas supports accumulation and comparison of intermediate reasoning, while widgets structure how users explore, refine, and consolidate prompt formulations. Together, these interface mechanisms support divergent and convergent modes of thought as well as metacognitive reflection on how prompt structure shapes AI behavior. Findings from our user studies provide empirical evidence that this design supports meaningful intermediary artifacts, improves users' understanding and sense of control, and encourages deliberate reasoning over conversational prompting. We discuss PromptCanvas as a transferable design and usage strategy for building Tools for Thought with generative AI.

CHI2026 Workshop on Tools for Thought, Barcelona
.

CCS Concepts

• **Human-centered computing** → **Graphical user interfaces**; **Natural language interfaces**; **Interactive systems and tools**.

Keywords

prompting, direct manipulation, design strategies, prompt widgets, canvas

1 Introduction and Background

Generative AI (genAI) systems are increasingly used to support writing, analysis, and other knowledge-intensive activities [19, 20, 22]. While these systems offer powerful generative capabilities, the interfaces through which users interact with them often provide limited support for the cognitive work involved in prompting. Most current prompting interfaces rely on linear, conversational interaction, encouraging users to iteratively react to outputs rather than reason deliberately about intent, constraints, and alternatives [12, 13].

Prior research on human-AI collaborative writing highlights both the benefits and risks of such interaction paradigms. AI-assisted tools can help sustain writing momentum and overcome creative blocks [8, 10], but they can also lead to reduced authorial engagement and loss of expressive intent when users become overly reliant on generated content [15, 17]. These challenges are amplified by the metacognitive demands of prompting, where users must simultaneously reason about task goals and about how to communicate those goals to the model [24].

From a Tools for Thought (TfT) perspective, these issues can be understood as a lack of adequate external representations to scaffold cognition. Early work on augmenting human intellect argued that interactive systems should enhance human reasoning through structured symbol manipulation and the evolution of problem representations [9]. In contrast, conversational genAI interfaces are typically structured as linear, turn-based dialogue systems [28], which can make it difficult to preserve and build upon intermediary outcomes in exploratory or ill-defined tasks without explicit support for state visibility and iterative refinement [2].

Research on dynamic and adaptive user interfaces suggests promising directions for addressing this gap. Prior work has shown how tightly coupling user input and system output [1], applying principles such as reification and reuse [6], and introducing dynamic widgets [25] can support more deliberate and flexible interaction. In parallel, systems for structured or guided prompting, such as Wordcraft [27], Promptify [7], PromptCharm [26] or Storyfier [18], demonstrate how interface structure can improve control and usability when working with generative models. However, these systems often prioritize output quality or task efficiency over supporting prompting as an ongoing cognitive process with meaningful intermediary artifacts.

Other works explore interaction paradigms that better support thinking with AI. For example, AI-Instruments [21] frames AI as configurable tools for creative workflows, Intent Tagging [11] makes user intent explicit, and *Luminate* [23] enables structured exploration of outputs. While these systems highlight the value of structured interaction and intermediate states, they primarily focus on output exploration, intent annotation, or tool configuration.

Building on these insights, we introduce PromptCanvas as a TfT-oriented design strategy that reframes prompt construction as an explicit, externalized cognitive activity. PromptCanvas builds upon our prior Late-Breaking Work (LBW) publication [3] and a subsequent extended version to appear at the ACM Transactions on Interactive Intelligent Systems (TiiS) journal [4]. It combines a **canvas-based workspace** with **prompt-construction widgets** that transform prompt elements into persistent, manipulable interface objects. The canvas supports accumulation, comparison, and reuse of intermediate reasoning, while widgets structure how users articulate and refine intent and constraints. Together, these mechanisms aim to support divergent exploration, convergent refinement, and metacognitive reflection when interacting with genAI.

We investigated this design through a series of user studies, including comparisons with conversational and structured writing interfaces, as well as a two-week field deployment. Our findings show how structuring prompting around persistent artifacts can support more deliberate reasoning and a greater sense of understanding and control, even when interaction requires additional cognitive effort. This work contributes a concrete, transferable strategy for designing and using genAI as a Tool for Thought.

2 Design Rationale and System Overview

Through the lens of TfT, PromptCanvas can be understood as a system designed to externalize and structure the cognitive work involved in interacting with genAI. Rather than optimizing solely for rapid prompt iteration or output quality, it aims to shape how users reason about intent, constraints, and alternatives during prompting. From this perspective, PromptCanvas reflects three core design principles for TfT-oriented genAI systems.

- First, prompt construction should be *externalized* into persistent artifacts rather than embedded in brief dialogue. Making prompt structure visible and revisable supports reflection and sensemaking across iterations.
- Second, prompting should support both *divergent and convergent thinking*. Users need to explore alternative framings before consolidating them into stable formulations, especially in exploratory tasks.
- Third, interaction should *limit premature cognitive offloading* by requiring users to articulate structure before relying on AI-generated output, encouraging forward reasoning rather than reactive trial-and-error.

To operationalize these principles, PromptCanvas combines a **canvas-based workspace** with **prompt-construction widgets**. The canvas serves as a persistent **cognitive workspace** for accumulating and reusing prompt artifacts (widgets) over time. Unlike conversational histories, the canvas supports non-linear organization, enabling users to reason across iterations.

Widgets function as **structured cognitive operators** within the canvas. They guide how users express different aspects of a prompt, such as task intent or constraints, while allowing direct manipulation of prompt widgets. Widgets do not automate prompting; instead, they scaffold reasoning by shaping how users construct and revise prompts. By interacting through widgets, users are encouraged to reason deliberately about prompt structure and its anticipated effects on AI behavior.

Together, the canvas and widgets form a composed interaction model: Widgets structure local reasoning, while the canvas supports global persistence and reflection. PromptCanvas integrates these elements into a unified interface for interacting with large language models, enabling users to iteratively construct, evaluate, and

reuse prompt artifacts. This design instantiates a transferable TfT strategy for structuring interaction with genAI around persistent intermediary representations.

3 Usage Scenarios and Types of Thought

To illustrate how PromptCanvas structures prompting as a TfT, consider the story-writing scenario shown in Figure 1. A user begins with an initial draft about Princess Seraphina and her dragon Ember in the kingdom of Eldoria. Rather than rewriting the story conversationally, the user decomposes narrative elements into structured widgets on the canvas.

Widgets such as *Princess Name*, *Princess Appearance*, *Sorcerer's Evil Goal*, *Dragon's Unique Ability*, and *Kingdom's Climate Transformation* externalize distinct narrative dimensions. This decomposition supports **divergent thinking**, as the user can explore variations in character traits, antagonist motives, or environmental changes without modifying the full text. For example, adjusting Malakar's goal or Ember's abilities becomes a targeted conceptual operation rather than a complete rewrite.

As widgets accumulate on the canvas, they function as intermediary artifacts that make the structure of the story explicit. The user can compare alternative values (Figure 1-D), revise specific components (e.g., renaming the princess or altering the resolution), and isolate one narrative attribute at a time. This supports structured exploration while reducing cognitive overload.

When the user applies the widgets using the "Rephrase the text based on your widgets" function (Figure 1-c1), the system consolidates selected structural decisions into an updated draft. This stage supports **convergent thinking**, where explored alternatives are integrated into a coherent narrative.

Finally, by observing how changes in widgets (e.g., modifying the antagonist description or climate transformation) affect the generated story, the user engages in **metacognitive reflection**. The persistent canvas allows the user to reason about how prompt structure shapes AI output, supporting learning and iterative refinement over time.

This scenario shows how PromptCanvas transforms story writing from reactive editing into a structured workflow of externalization, exploration, consolidation, and reflection.

4 Study and Methods

This section builds upon the empirical studies reported in our work on PromptCanvas [4]. The original evaluation comprised three studies across two system iterations. First, a lab study compared a canvas-based prompting interface with a conversational prompting interface. Second, a two-week field deployment examined how participants appropriated the interface in their own workflows over time. Third, a lab study compared the updated canvas-based interface with Wordcraft [27], a state-of-the-art writing tool.

Across these studies, participants engaged in exploratory writing and analysis tasks requiring iterative prompt construction and revision. In the field study, participants used the system for self-selected activities, including writing and programming-related tasks. Interaction logs, canvas-based prompt artifacts, and self-reports were collected. Participants in lab study 1 are indexed as P1-P18, and participants in lab study 2 are indexed as P19-P36 to clearly distinguish them from those in lab study 1 and avoid overlap in participant IDs.

Here, we reinterpret these empirical results through a TfT lens, focusing not on comparative performance alone, but on how interface design shaped intermediary artifacts, structured reasoning, perceived control, and cognitive effort. The analysis therefore emphasizes design implications rather than system evaluation.

5 Findings

Across both lab studies and the field study, PromptCanvas reshaped prompting from a transient input activity into a structured cognitive process. Rather than treating prompts as disposable instructions, participants described using widgets as persistent, manipulable representations of their thinking.

5.1 Intermediary Artifacts and Reuse

Participants consistently treated widgets as durable intermediary artifacts rather than one-off commands. In contrast to conversational interfaces, where prior prompts are lost in a linear history, widgets remain visible and spatially organized, supporting comparison and revision. As P7 explained, *"What happens in conversational UI is that you continuously are giving prompts... so once a new prompt is given, and it has been applied the prompt goes out of the window. So when I'm having the widget, I can see what are my exact prompts that are functioning."* This visibility enabled participants to maintain awareness of active constraints and treat them as stable representations rather than ephemeral inputs. Beyond persistence, participants described widgets as compositional tools. P30 described them as *"building blocks"* for shaping different narrative versions, suggesting that users constructed and recombined structured components rather than iteratively exchanging prompts. Similarly, P35 viewed the canvas as a *"mental map"* that supported visualizing connections across the story, while P36 compared widgets to *"thought bubbles"* that helped separate and reorganize ideas without losing track of the narrative. These metaphors indicate that widgets functioned not merely as interface elements, but as externalized cognitive structures.

Widgets also supported revision and reuse over time. As P15 noted, *"You could go back and edit a previous widget without it changing the entire text."* This ability to revisit and modify intermediary artifacts allowed participants to refine ideas incrementally rather than restarting the prompting process. Together, these accounts show that the canvas-based prompting interface (PromptCanvas) shifted prompting toward a compositional workflow grounded in persistent, manipulable representations of intent.

5.2 Understanding, Control, and Reflection

By externalizing prompt structure into persistent widgets, the interface influenced how participants reasoned about AI behavior. Rather than merely reacting to generated text, users described steering the system more deliberately. As P3 noted, *"You have control in the conversational UI, but it can be cumbersome. It (PromptCanvas) was a more streamlined control, which was really nice."* This sense of control became particularly visible in comparison with Wordcraft [27]. Participants described PromptCanvas as giving them *"way*

more influence on the generated text" (P30) and feeling *"more in control"* (P33).

Participants also described improved focus when making changes. P8 explained that the system *"really narrows down the focus where I need to edit something… I will also be able to select my focal points where I need to change something…"* This suggests that making constraints visible supported forward reasoning about the effects of changes, as well as metacognitive reflection on what was guiding the model at any given moment.

5.3 Effort, Value, and Task Context

Participants acknowledged that the use of prompt construction widgets sometimes required more time and cognitive effort than conversational prompting. As P3 noted, *"using widgets does take a bit more time."* However, this additional effort was often described as productive rather than burdensome. P8 framed the time investment as aligned with creative goals, explaining that it was *"a bit time-consuming, however, I am actually putting that time to bring out my creativity and bring out new ideas, so I'm focusing on writing."* Perceived value depended on task context. For short, time-constrained tasks, faster conversational interaction was sometimes preferred. In contrast, for exploratory or open-ended tasks, participants valued the persistence and modularity of widgets, even at the cost of increased interactional complexity. In these contexts, the added effort was experienced as worthwhile when aligned with goals such as ideation, revision, and deliberate control.

6 Discussion and Implications

Our findings show that prompt construction widgets reshape interaction with genAI by supporting structured prompting, creativity, and user control. Compared to a conversational interface and Wordcraft [27], PromptCanvas encouraged a more exploratory and compositional workflow. We interpret these findings through a Tools for Thought lens and derive implications for the design and evaluation of genAI systems that aim to support cognition.

Prompting as a Tool for Thought. Prompt construction can be reframed as a Tool for Thought rather than a transient input step. By making prompts persistent and manipulable, PromptCanvas turns them into external cognitive artifacts that support deliberate reasoning and reflection [24]. Unlike conversational prompting and tools such as Wordcraft [27], it encourages users to decompose intent into modular constraints and iteratively refine them, similar to structured interaction paradigms like DynaVis [25]. Emphasizing such intermediary representations can influence both output quality and users' thinking during interaction.

Structured Exploration and Creativity. The canvas-based prompting interface with persistent, modular constraints (PromptCanvas) supported more exploratory and varied creative workflows than linear conversational interactions or structured chat-based tools such as Wordcraft [27]. Creativity support scores remained consistently high across studies, and text outcome metrics indicated more varied and less predictable outputs. Similar to observations in DynaVis [25] and Copilot-style systems [5], structured interactive scaffolds appear to encourage broader exploration than sequential request-response formats. Whereas chat-oriented systems emphasize incremental rewriting within a linear dialogue structure [27], a design that enables parallel adjustments and modular experimentation allows users to explore alternatives simultaneously. Structured prompting environments may therefore be particularly well-suited for open-ended or ideation-heavy tasks that benefit from branching exploration.

Customizability, Control, and Metacognitive Support. Prompt construction widgets increased perceived control and customizability. Because genAI systems place metacognitive demands on users by requiring explicit goal formulation and task decomposition [24], externalizing constraints into interactive elements supported cognitive monitoring and structured guidance [14]. Compared to Wordcraft, participants felt more effective and capable with PromptCanvas. Although workload was comparable, the structured interface improved perceived performance and enabled iterative refinement without repeated prompt reformulation, suggesting that structured control can enhance agency without increasing cognitive burden.

Productive Friction and Constraint Management. While structured interfaces may be expected to increase cognitive load [25], the additional structure introduced by prompt construction widgets often acted as productive friction. The effort of creating and managing widgets supported understanding, reuse, and creative control. However, as the number of widgets increased, participants reported visual clutter and instruction overload. They suggested clustering, color-coding, scoping constraints, and incremental application to manage complexity. Future prompt construction widget systems should therefore include mechanisms for constraint management and adaptive simplification, especially for time-constrained tasks.

Implications for Design and Evaluation. PromptCanvas highlights the need to reconsider how Tools for Thought with genAI are evaluated. Traditional metrics centered on speed or output quality may overlook intermediary artifacts, reasoning processes, and metacognitive engagement. Evaluations should incorporate process-oriented measures and longitudinal perspectives capturing reuse, strategy development, and sustained interaction over time. Beyond writing, the widget-and-canvas paradigm may generalize to other domains such as visualization [25], visual generation [23], or structured editing systems like Textoshop [16]. Supporting persistent, manipulable intermediary artifacts may represent a broader design strategy for future genAI-based TfT.

7 Conclusion

This paper presented PromptCanvas, a Tool for Thought-oriented design that reframes prompt construction as an explicit, externalized cognitive process. By combining a canvas-based workspace with prompt-construction widgets, PromptCanvas supports divergent exploration, convergent refinement, and reflective engagement when interacting with genAI. Through a series of user studies, we showed that this design preserves meaningful intermediary artifacts, encourages deliberate reasoning, and supports a greater sense of understanding and control, even when interaction requires additional cognitive effort. Together, these contributions illustrate how genAI interfaces can be designed not only to produce outputs but to support people in actively thinking with AI.

Acknowledgments

Funded by Elitenetzwerk Bayern and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 525037874.

References

[1] Christopher Ahlberg and Ben Shneiderman. 1994. Visual information seeking: tight coupling of dynamic query filters with starfield displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) *(CHI '94)*. Association for Computing Machinery, New York, NY, USA, 313–317. doi:10.1145/191666.191775

[2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300233

[3] Rifat Mehreen Amin, Oliver Hans Kühle, Daniel Buschek, and Andreas Butz. 2025. Composable Prompting Workspaces for Creative Writing: Exploration and Iteration Using Dynamic Widgets. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 144, 11 pages. doi:10.1145/3706599.3720243

[4] Rifat Mehreen Amin, Oliver Hans Kühle, Daniel Buschek, and Andreas Butz. 2026. PromptCanvas: Composable Prompting Workspaces Using Dynamic Widgets for Exploration and Iteration in Creative Writing. *To appear in ACM Trans. Interact. Intell. Syst.*

[5] Shraddha Barke, Michael B. James, and Nadia Polikarpova. 2023. Grounded Copilot: How Programmers Interact with Code-Generating Models. *Proc. ACM Program. Lang.* 7, OOPSLA1, Article 78 (apr 2023), 27 pages. doi:10.1145/3586030

[6] Michel Beaudouin-Lafon and Wendy E. Mackay. 2000. Reification, polymorphism and reuse: three principles for designing visual interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (Palermo, Italy) *(AVI '00)*. Association for Computing Machinery, New York, NY, USA, 102–109. doi:10.1145/345513.345267

[7] Stephen Brade, Bryan Wang, Mauricio Sousa, Sageev Oore, and Tovi Grossman. 2023. Promptify: Text-to-Image Generation through Interactive Prompt Exploration with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–14. doi:10.1145/3586183.3606725

[8] Hai Dang, Karim Benharrak, Florian Lehmann, and Daniel Buschek. 2022. Beyond Text Generation: Supporting Writers with Continuous Automatic Text Summaries. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) *(UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 98, 13 pages. doi:10.1145/3526113.3545672

[9] Douglas C Engelbart. 2023. Augmenting human intellect: A conceptual framework. In *Augmented education in the global age*. Routledge, 13–29.

[10] Iona Gilburt. 2024. A machine in the loop: the peculiar intervention of artificial intelligence in writer's block. *New Writing* 21, 1 (2024), 26–37.

[11] 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 (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 531, 31 pages. doi:10.1145/3706598.3713861

[12] Charles Goodwin. 2015. Professional vision. In *Aufmerksamkeit: Geschichte-theorie-empirie*. Springer, Cham, Switzerland, 387–425.

[13] Peiling Jiang, Jude Rayan, Steven P. Dow, and Haijun Xia. 2023. Graphologue: Exploring Large Language Model Responses with Interactive Diagrams. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–20. doi:10.1145/3586183.3606737

[14] Abdulrahman Khamaj and Abdulelah M. Ali. 2024. Adapting user experience with reinforcement learning: Personalizing interfaces based on user behavior analysis in real-time. *Alexandria Engineering Journal* 95 (2024), 164–173. doi:10.1016/j.aej.2024.03.045

[15] Max Kreminski. 2024. The Dearth of the Author in AI-Supported Writing. In *Proceedings of the Third Workshop on Intelligent and Interactive Writing Assistants* (Honolulu, HI, USA) *(In2Writing '24)*. Association for Computing Machinery, New York, NY, USA, 48–50. doi:10.1145/3690712.3690725

[16] Damien Masson, Young-Ho Kim, and Fanny Chevalier. 2024. Textoshop: Interactions Inspired by Drawing Software to Facilitate Text Editing. arXiv:2409.17088

[17] 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* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 355, 34 pages. doi:10.1145/3544548.3581225

[18] Zhenhui Peng, Xingbo Wang, Qiushi Han, Junkai Zhu, Xiaojuan Ma, and Huamin Qu. 2023. Storyfier: Exploring Vocabulary Learning Support with Text Generation Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) *(UIST '23)*. Association for Computing Machinery, New York, NY, USA, Article 46, 16 pages. doi:10.1145/3586183.3606786

[19] Hua Xuan Qin, Shan Jin, Ze Gao, Mingming Fan, and Pan Hui. 2024. CharacterMeet: Supporting Creative Writers' Entire Story Character Construction Processes Through Conversation with LLM-Powered Chatbot Avatars. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 1051, 19 pages. doi:10.1145/3613904.3642105

[20] Emily Reif, Crystal Qian, James Wexler, and Minsuk Kahng. 2024. Automatic Histograms: Leveraging Language Models for Text Dataset Exploration. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 53, 9 pages. doi:10.1145/3613905.3650798

[21] Nathalie Riche, Anna Offenwanger, Frederic Gmeiner, David Brown, Hugo Romat, Michel Pahud, Nicolai Marquardt, Kori Inkpen, and Ken Hinckley. 2025. AI-Instruments: Embodying Prompts as Instruments to Abstract & Reflect Graphical Interface Commands as General-Purpose Tools. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1104, 18 pages. doi:10.1145/3706598.3714259

[22] 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 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–26. doi:10.1145/3613904.3642400

[23] 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 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 644, 26 pages. doi:10.1145/3613904.3642400

[24] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 680, 24 pages. doi:10.1145/3613904.3642902

[25] Priyan Vaithilingam, Elena L. Glassman, Jeevana Priya Inala, and Chenglong Wang. 2024. DynaVis: Dynamically Synthesized UI Widgets for Visualization Editing. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 985, 17 pages. doi:10.1145/3613904.3642639

[26] Zhijie Wang, Yuheng Huang, Da Song, Lei Ma, and Tianyi Zhang. 2024. PromptCharm: Text-to-Image Generation through Multi-modal Prompting and Refinement. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–21. doi:10.1145/3613904.3642803

[27] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: Story Writing With Large Language Models. In *Proceedings of the 27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) *(IUI '22)*. Association for Computing Machinery, New York, NY, USA, 841–852. doi:10.1145/3490099.3511105

[28] Qingxiao Zheng, Yiliu Tang, Yiren Liu, Weizi Liu, and Yun Huang. 2022. UX Research on Conversational Human-AI Interaction: A Literature Review of the ACM Digital Library. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 570, 24 pages. doi:10.1145/3491102.3501855