Beyond the Conversational Paradigm: Models, Design Principles, and Tools for Advancing the Spectrum of Human-LLM Interaction

Laura Colazzo
Department of Electronics, Information, and Bioengineering
Politecnico di Milano
Milan, Italy
laura.colazzo@polimi.it

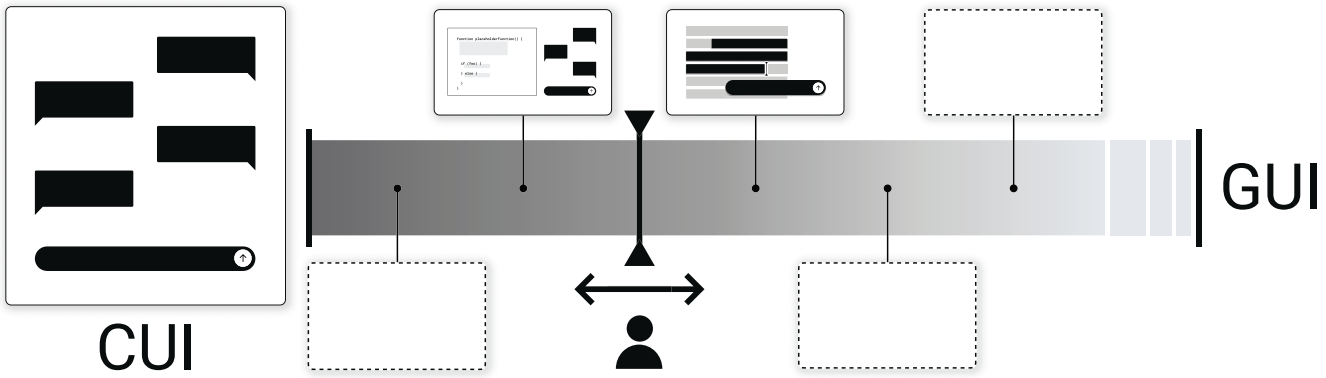

Figure 1: Visual representation of the CUI–GUI spectrum, emphasizing that optimal human–LLM interfaces can be situated at various points along a continuum between conversational and graphical interface paradigms.

Abstract

Interaction with Large Language Models (LLMs) is expanding beyond strictly Conversational User Interfaces (CUIs), with systems starting to incorporate direct-manipulation mechanisms, typical of Graphical User Interfaces (GUIs), into their designs. We argue that the current landscape of human–LLM interfaces can be described as a continuum, ranging from strictly conversational to increasingly visual and interactive interfaces built on top of a basic CUI. Based on this assumption, the doctoral research presented in this paper aims to model and analyze this CUI–GUI spectrum. Grounded in a Human-Centered Design methodology and a research-through-design approach, the work investigates factors influencing users' preferences for different interface configurations along the spectrum, through comparative empirical studies, and leveraging representative interface prototypes across it. Insights from these studies will inform the design of adaptive LLM interfaces that dynamically adjust the balance between conversational and graphical elements in response to users' needs and contextual factors. In parallel, the research will explore both risks and benefits associated with different regions of the spectrum, with the ultimate goal of developing a comprehensive characterization, as well as a set of design principles and tools that can effectively guide the design of the next generation of human–LLM interfaces.

IUI Companion '26, Paphos, Cyprus

ACM ISBN 979-8-4007-1985-1/26/03
https://doi.org/10.1145/3742414.3789234

CCS Concepts

• **Human-centered computing** → **Interaction paradigms**; **Natural language interfaces**; **Interaction techniques**; **Empirical studies in HCI**;

Keywords

Human-LLM Interaction, LLM Interfaces, Conversational User Interfaces, Graphical User Interfaces

ACM Reference Format:
Laura Colazzo. 2026. Beyond the Conversational Paradigm: Models, Design Principles, and Tools for Advancing the Spectrum of Human-LLM Interaction. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces (IUI Companion '26), March 23–26, 2026, Paphos, Cyprus.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3742414.3789234

1 Introduction

Although interaction with Large Language Models (LLMs) was originally designed around the conversational metaphor, more recently LLM-based systems have started to include direct manipulation mechanisms into their User Interfaces (UIs), moving beyond the pure chat metaphor [4, 8]. This has contributed to the rise of a broader spectrum of human–LLM interaction that spans fully Conversational UIs (CUIs), as well as increasingly interactive and visually enriched forms of engagement typical of Graphical User Interfaces (GUIs). This CUI–GUI spectrum can be conceived as a continuum in which LLM interfaces gradually transition from

strictly text-based conversational designs, toward layouts that integrate graphical elements and interactive components layered on top of the basic chat interface [4].

While this shift to more graphical and direct-manipulation LLM interfaces represented an attempt to facilitate interaction with LLMs and move beyond a rigid interaction based solely on natural language for intent expression [8], it does not necessarily imply that LLM interfaces, whose position along the spectrum is more skewed to the right, dominate the others. Depending on various factors, different points along the spectrum may be considered optimal (Figure 1). For example, an interface heavily based on direct manipulation could better support specialized tasks; conversely, a more standard text-based CUI, having a lower entry barrier, may be preferable for novice users or for tasks that are more exploratory.

2 Research Objectives

In this context, the aim of this research is to study human–LLM interaction beyond the purely conversational paradigm. This will involve modeling and analyzing the CUI–GUI spectrum of human–LLM user interfaces to investigate which factors (e.g., trust in the system, task characteristics) drive a shift toward one side or another in search of the optimal combination of conversational and graphical elements. Findings will inform the design of adaptive LLM interfaces that dynamically reconfigure themselves based on users' needs and contextual factors. Having characterized the CUI–GUI spectrum and identified the factors guiding movement along it, the research will continue by adopting another complementary perspective, identifying not only the benefits but also the potential risks associated with different combinations of conversational and graphical elements within the UI. Specifically, it will explore how different interface configurations—intended as discrete regions of the CUI–GUI spectrum, each defined by a distinct combination of conversational and graphical elements—can negatively impact the user (e.g., contributing to the decline of skills such as critical thinking), while also considering the benefits such configurations can provide (e.g., improved sense of agency). The ultimate goal is to develop a comprehensive characterization of the CUI–GUI spectrum that can guide the design of future human–LLM interfaces. To achieve this, the research will address the following Research Questions (RQs).

- **RQ1:** How can the CUI–GUI spectrum be modeled?
- **RQ2:** What factors contribute to the identification of the optimal balance between conversational and graphical elements in the design of human–LLM interfaces?
- **RQ3:** How can evidence of this influence support the design of adaptive LLM interfaces?
- **RQ4:** What are the risks and benefits associated with different regions of the CUI–GUI spectrum?

3 Methods and Expected Outcomes

The research is grounded in the Human-Centered Design (HCD) methodology and adopts a research-through-design approach to systematically explore, characterize, and operationalize the spectrum of human–LLM interaction, ranging from purely conversational to more graphical and direct-manipulation-oriented paradigms.

Current work is focusing on building a deep understanding of the current landscape of human–LLM interaction, to define a conceptual model of the envisioned CUI–GUI spectrum. Emerging trends are being explored, to systematically classify existing LLM interfaces, with representative instances selected from both the academic and commercial worlds. This exploration builds on a previous user study [1], which focused on voice-based interaction for text composition using LLM interfaces, by expanding the range of tasks and scenarios under consideration beyond content creation. Through comparative empirical studies, quantitative and qualitative data will then be collected and analyzed to gain a deeper understanding of how individual and contextual factors shape users' preferences for more conversational or graphical interaction styles. This investigation will be supported by the development of prototypes representing different regions of the spectrum, enabling a direct comparison of interaction modalities under the influence of various factors. In parallel, the risks and benefits associated with different human–LLM interaction paradigms will be empirically examined and systematically mapped across the spectrum. Taken together, insights from these studies will inform a conceptual model and an enabling framework that can effectively guide the design and development of adaptive human–LLM interfaces. Specifically, the focus will be on enhancing the interaction paradigm between humans and LLMs to support more personalized forms of engagement, with LLM interfaces able to adjust the combination of conversational and graphical elements based on individual user characteristics and contextual factors.

The ultimate goal is to produce a comprehensive characterization of human-LLM interaction, supported by a set of design principles and tools that can guide the design of the next-generation LLM interfaces.

4 Related Works

4.1 CUIs Vs. GUIs

Previous studies revealed that users' preferences for CUIs or GUIs are shaped by individual factors. For example, Liu et al. [7] found that in the decision-making domain those with a more rational decision style tend to prefer GUIs. However, as the interfaces compared in this study were not powered by LLMs, results may not extend to LLM interactions.

Concerning the effect of different interaction paradigms on the user, a study in the domain of decision support systems examined how the selection between a CUI or GUI affects the development of trust in the system [5], revealing CUIs are more effective in building trust and satisfaction than GUIs. This suggests that perceived trust in the system could be among the factors shaping users' preference for different interaction paradigms. Indeed, if CUIs promote greater trust in the system's output, users might be more inclined to favor those that expose a CUI rather than a GUI.

A major limitation of these studies is that they treat CUIs and GUIs as mutually exclusive, overlooking the possibility of a continuum between the two. Yet tools like DynaVis [11] show that hybrid solutions exist, suggesting that future research should also consider intermediate areas along the CUI–GUI spectrum.

4.2 Effects of Interaction with LLMs on Users

A growing body of research analyzes the effect of human–LLM interaction on users. In their work, Shi et al. [9] highlight that as reliance on these models increases, so does the risk for users to become victims of *LLM Dark Patterns*. These are defined as deceptive behaviors emerging dynamically through conversations "that steer users toward beliefs, decisions, or behaviors they might not otherwise adopt." An example is the *Simulated Authority* pattern, where the LLM misleadingly presents itself as an authoritative entity, potentially fostering over-trust in its answers in critical contexts. Patterns like this can negatively affect users, posing their autonomy at risk. This is true especially when anthropomorphic cues foster user trust, letting manipulation go unnoticed [9]. Over-trust can lead to over-reliance: a concept that is closely related to the lack of critical thinking [6]. In knowledge work specifically, higher self-confidence proved to be correlated with more critical engagement, while little confidence may induce over-reliance and reduce critical thinking [6]. Over-reliance can also hinder skill development, as users may fail to internalize the competencies needed for the delegated tasks [2, 6, 10]. Drosos et al. [3] raise concerns about the potential decline in users' critical thinking capabilities, which are essential in knowledge work. However, as AI systems become increasingly pervasive, the ability to critically engage with potentially wrong or misleading content is becoming crucial across all contexts [2].

While these works discuss the impact of human–LLM interaction in broad terms, introducing a spectrum-based model raises the question of whether different regions along it affect users differently. This motivates the need for a finer-grained characterization able to capture a potentially uneven distribution of risks and benefits across the spectrum.

5 Conclusion

Grounded in a HCD methodology and a research-through-design approach, this doctoral research aims to advance the field of Human–LLM interaction by analyzing the current landscape of LLM interfaces beyond the purely conversational paradigm and by proposing a description based on the CUI–GUI spectrum model. Through its systematic characterization, which will involve the identification of key factors that shape users' preferences and the evaluation of risks and benefits associated with different interface configurations, the research seeks to derive a comprehensive model, along with design principles and tools, that can ultimately inform the design of the next generation of adaptive human-LLM interfaces.

Acknowledgments

The author would like to thank her advisor, Prof. Maristella Matera, for the guidance and support provided throughout the development of this work.

References

[1] Laura Colazzo and Maristella Matera. 2025. Enabling Voice-based Direct Manipulation in Intent-oriented Interaction Paradigms. In *Proceedings of the 16th Biannual Conference of the Italian SIGCHI Chapter (CHItaly '25)*. ACM, Article 93, 2 pages. doi:10.1145/3750069.3755951

[2] Valdemar Danry, Pat Pataranutaporn, Yaoli Mao, and Pattie Maes. 2023. Don't Just Tell Me, Ask Me: AI Systems that Intelligently Frame Explanations as Questions Improve Human Logical Discernment Accuracy over Causal AI explanations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. ACM, 1–13. doi:10.1145/3544548.3580672

[3] Ian Drosos, Advait Sarkar, Xiaotong, Xu, and Neil Toronto. 2025. "It makes you think": Provocations Help Restore Critical Thinking to AI-Assisted Knowledge Work. doi:10.48550/arXiv.2501.17247 arXiv:2501.17247 [cs].

[4] Ian Drosos, Jack Williams, Advait Sarkar, Nicholas Wilson, Sean Rintel, and Payod Panda. 2025. Dynamic Prompt Middleware: Contextual Prompt Refinement Controls for Comprehension Tasks. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*. ACM, Amsterdam Netherlands, 1–23. doi:10.1145/3729176.3729203

[5] Akshit Gupta, Debadeep Basu, Ramya Ghantasala, Sihang Qiu, and Ujwal Gadiraju. 2022. To Trust or Not To Trust: How a Conversational Interface Affects Trust in a Decision Support System. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*. ACM, 3531–3540. doi:10.1145/3485447.3512248

[6] Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–22. doi:10.1145/3706598.3713778

[7] Xuanhui Liu, Tim Rietz, and Alexander Maedche. 2025. Conversational versus graphical user interfaces: the influence of rational decision style when individuals perform decision-making tasks repeatedly. *Universal Access in the Information Society* 24, 2 (June 2025), 1157–1172. doi:10.1007/s10209-024-01122-1

[8] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2024. DirectGPT: A Direct Manipulation Interface to Interact with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, 1–16. doi:10.1145/3613904.3642462

[9] Yike Shi, Qing Xiao, Qing Hu, Hong Shen, and Hua Shen. 2025. The Siren Song of LLMs: How Users Perceive and Respond to Dark Patterns in Large Language Models. (2025). doi:10.48550/ARXIV.2509.10830

[10] Lev Tankelevitch, Elena L. Glassman, Jessica He, Majeed Kazemitabaar, Aniket Kittur, Mina Lee, Srishti Palani, Advait Sarkar, Gonzalo Ramos, Yvonne Rogers, and Hari Subramonyam. 2025. Tools for Thought: Research and Design for Understanding, Protecting, and Augmenting Human Cognition with Generative AI. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–8. doi:10.1145/3706599.3706745

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