

# Promoting Critical Thinking With Domain-Specific Generative AI Provocations

Thomas Serban von Davier  
Carnegie Mellon University  
Pittsburgh, PA, United States  
tvondavi@andrew.cmu.edu

Jodi Forlizzi  
Carnegie Mellon University  
Pittsburgh, PA, United States  
forlizzi@cs.cmu.edu

Hao-Ping (Hank) Lee  
Carnegie Mellon University  
Pittsburgh, PA, United States  
haopingl@cs.cmu.edu

Sauvik Das  
Carnegie Mellon University  
Pittsburgh, PA, United States  
sauvik@cmu.edu

## Abstract

The evidence on the effects of generative AI (GenAI) on critical thinking is mixed, with studies suggesting both potential harms and benefits depending on its implementation. Some argue that AI-driven provocations, such as questions asking for human clarification and justification, are beneficial for eliciting critical thinking. Drawing on our experience designing and evaluating two GenAI-powered tools for knowledge work, ArtBot in the domain of fine art interpretation and Privy in the domain of AI privacy, we reflect on how design decisions shape the form and effectiveness of such provocations. Our observations and user feedback suggest that domain-specific provocations, implemented through productive friction and interactions that depend on user contribution, can meaningfully support critical thinking. We present participant experiences with both prototypes and discuss how supporting critical thinking may require moving beyond static provocations toward approaches that adapt to user preferences and levels of expertise.

## CCS Concepts

• **Human-centered computing** → **Interaction paradigms.**

## Keywords

Human-centered AI, Critical Thinking, Prototyping, Generative AI, Human-AI Collaboration

### ACM Reference Format:

Thomas Serban von Davier, Hao-Ping (Hank) Lee, Jodi Forlizzi, and Sauvik Das. 2026. Promoting Critical Thinking With Domain-Specific Generative AI Provocations. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 6 pages.

## 1 Introduction

As generative artificial intelligence (GenAI) systems continue to increase in popularity and capability, applied research has increasingly demonstrated the benefits of domain-specific models over large, general-purpose models, particularly in settings that require nuanced reasoning or contextual sensitivity [2, 6, 14]. In light of

recent discussions on the negative impact of GenAI on critical thinking, these findings suggest that specialization, rather than scale alone, can significantly improve the outcomes of human-AI interaction.

In this workshop paper, we argue that a similar principle applies to the design of GenAI systems intended to support critical thinking, particularly when such systems are framed as provocateurs or facilitators. Prior work has emphasized the importance of deliberately antagonistic or challenging AI behaviors as a core design requirement for tools for thought [3, 20]. This line of research argues that when AI systems introduce friction by questioning assumptions, surfacing alternatives, or resisting user intent, they prompt users to pause, reflect, and engage in deeper critical thinking during interaction.

Building on this literature, we advance the position that provocations embedded in GenAI systems must themselves be domain-specific to effectively support critical thinking. Rather than relying on general-purpose challenges, we contend that provocations grounded in domain knowledge, norms, and established frameworks are more likely to be meaningful, interpretable, and actionable for users.

To support this argument, we draw on our experience designing and evaluating two GenAI systems presented in recent CHI papers, each developed explicitly to support critical thinking within a distinct domain. The first system acts as a conversational companion for digital collections of fine art, where provocations are informed by art history and educational curricula to encourage interpretive reflection [22]. The second system is a structured, AI-assisted whiteboard for practitioners to identify privacy risks and mitigations, with provocations derived from established privacy taxonomies and frameworks. Although both systems adopt a provocateur-facilitator framework, the form and content of their challenges are tightly coupled to their respective domains [10].

In this paper, we present an overview of both systems, including how domain knowledge informed the design of provocations and how each system was evaluated with users. From these cases, we surface comparative learnings about how domain-specific provocations are received by different audiences, as well as design insights into how framing GenAI as a provocateur or facilitator influences engagement and interpretation.

Ultimately, this workshop contribution offers preliminary recommendations for framing GenAI systems as effective tools for



This work is licensed under a Creative Commons Attribution 4.0 International License.  
CHI '26, Barcelona, Spain  
© 2026 Copyright held by the owner/author(s).

domain-specific critical thinking. We argue that while domain specificity strengthens the impact of provocations, user preferences and expectations continue to shape how such systems are perceived and used. We present this work to invite discussion on how to balance domain grounding, provocation, and user agency in future tools for thought.

## 2 Background

While industry, government, and the research world explore various paths towards improving AI and building more AI-enabled systems, there has been a growing field of work questioning whether the current implementation of automation-prioritizing chatbots is truly the most effective form of tooling. A large survey and mixed-methods study by Microsoft researchers found that many knowledge workers report lower participation and lower engagement with critical-thinking steps when they rely on GenAI for writing [9]. Similarly, a widely reported neurocognitive study from the MIT Media Lab used EEG to compare people writing essays with ChatGPT, with search, or unaided [7]. The LLM-assisted writers showed reduced neural markers of executive control, memory encoding, and creativity in the tasks used by the study, a signal that heavy reliance on AI for generative work can reduce the cognitive activation behind learning and original composition. This reveals a “payoff” problem that some interpret as evidence that human productivity gains from AI are neither automatic nor uniformly distributed. Taken together, these strands suggest that (a) AI can lower cognitive effort and measurable engagement in specific tasks, (b) reported productivity gains do not always translate into better-quality outcomes, and (c) design choices matter if we want AI to augment rather than atrophy human critical thinking.

Many researchers, including our colleagues interested in the Tools for Thought workshop, argue that GenAI systems should scaffold, provoke, and support deliberation instead of simply producing answers. The principle behind viewing LLM interactions this way is the same principle that guides peer review, testing & evaluation, iteration and the Socratic method: that an idea or plan is only considered high quality if it has been challenged and verified. Researchers have proposed interventions (e.g., “provocations” [20] or “antagonisms” [3]) that deliberately surface critiques or alternatives to model outputs; their work shows such micro-interventions can increase metacognitive activity and user scrutiny in shortlisting and knowledge tasks. This research philosophy has contributed to Park and Kulkarni’s thinking assistant [16] and Liu et al.’s Thoughtful AI [13].

Ye et al. explore how AI could be engineered to ask better questions and support domain-specific inquiry rather than provide ready-made conclusions [27]. A focus on domain-specific inquiry led to our ArtBot paper [22]. It is an example of a Socratic-style companion that guides users through analysis by prompting reflection and layered questioning, illustrating how dialogic agents can foster analytic practice in nontechnical domains. Other tools like FarSight [24] and Privy [10] demonstrate in-situ interfaces that scaffold human-AI reasoning within established structures. These tools show how embedding reflective checks directly into development and authoring environments can shift outcomes away

from “copy-paste” convenience and towards accountable, deliberate practice.

## 3 Case Studies

Our prior work introduced two GenAI systems designed to support distinct forms of knowledge work through domain-specific provocations: ArtBot, a conversational companion for fine art interpretation, and Privy, a structured ideation tool for identifying and mitigating AI privacy risks. Both systems were developed and evaluated in recent CHI papers, in which their task-specific effectiveness was assessed through controlled experiments. In this workshop paper, we focus not on comparative performance outcomes, but on the design methodologies and interaction paradigms that distinguish the two systems and inform their role as tools for critical thinking. Both systems were intentionally designed to be compared against non-AI-powered alternatives and were evaluated with participants in controlled settings (ArtBot:  $n = 13$ ; Privy:  $n = 12$ <sup>1</sup>). In addition to addressing their primary research questions, these studies generated rich observational data on how users responded to different forms of AI provocation and facilitation. Below, we outline the core design decisions underpinning each system, with particular attention to how domain knowledge shaped the structure and presentation of provocations.

### 3.1 ArtBot: Domain-Grounded Provocation for Art Interpretation

ArtBot (Figure 1) was developed to support the engagement with digital fine art collections through an interactive, large language model-augmented interface [22]. The system was implemented using locally hosted Llama 3 models [15] combined with RAG (Retrieval Augmented Generation) [12], allowing access to a curated corpus of art-historical metadata, curatorial texts, and educational materials.

The interaction paradigm of ArtBot draws inspiration from Socratic tutoring, an instructional approach in which understanding is developed through targeted questioning rather than direct explanation. This approach aligns with dialogic education practices, where learning emerges through a dialogue of provocations and reflective responses [21]. In ArtBot, these practices were operationalized through prompts that challenge users to articulate interpretations and reconsider assumptions about an artwork. Some examples of the interaction include the GenAI system asking the participant whether their interpretation of the artwork changes if they know it was made during a time of revolution, and the participant then responding and discussing whether their interpretation changes and how.

Users interact with ArtBot while viewing an artwork image, engaging in a conversational exchange intended to deepen interpretation rather than deliver authoritative explanations. During the evaluation, participants were asked to provide short written reflections after interacting with each artwork, which served as the basis for assessing interpretive engagement.

<sup>1</sup>The original paper presents two versions of Privy [10] — one incorporating LLM-powered features and one without — each evaluated with 12 practitioners.

**Table 1: This table provides an overview on the theoretical grounding underlying both systems and which specific domains each one operates under.**

Case Study	Domain	Theoretical Grounding	Model Type
ArtBot	Art Interpretation	Dialogic Education [21]	Local, open-weight models [15]
Privy	AI Privacy Reports	Framework Grounded Provocations [5]	Cloud-host, proprietary models [25]

### 3.2 Privy: Structured Provocation for AI Privacy Planning

Privy (Figure 2) was designed to support AI practitioners during the early design and planning phases of AI system development, with a focus on identifying and mitigating privacy risks [10]. Privy leverages a structured, branching-tree workflow embedded within a whiteboard-style interface. This structure guides users through a sequence of decisions and reflections commonly encountered in AI system design.

The system is augmented by a GenAI backend (implemented using GPT-4.1) and is informed by system prompts grounded in established AI privacy taxonomies, including Lee et al. and Das et al.’s frameworks [5, 11]. These domain-specific prompts (e.g., design frictions that require users to assess the relevance and severity of risks, or to reflect on the effectiveness of proposed mitigations) enable Privy to challenge practitioners by surfacing potential risks, highlighting unintended use cases, and facilitating planning for privacy mitigation best practices — e.g., *How can you design this feature to encourage users to regularly review and update their sharing settings so they stay in control of how their social network data is used?*

As users progress through the Privy workflow, they iteratively document privacy risks and associate each with tailored mitigation strategies. The interaction culminates in a structured artifact, a design document summarizing identified risks and proposed mitigations. Evaluation focused on the quality and completeness of these artifacts, which were assessed by privacy experts.

## 4 Supporting Critical Thinking Through Challenging Provocations

Although ArtBot and Privy differ substantially in domain, audience, and interaction style, both systems were intentionally designed to position GenAI as a provocateur and facilitator rather than an answer engine. In each case, domain knowledge plays a central role in shaping how provocations are formulated, when they are introduced, and how users are encouraged to respond.

Both ArtBot and Privy intentionally incorporated **design friction** into their interactions by resisting the impulse to provide direct answers. In ArtBot, this friction took the form of a Socratic interaction style that prompted users to articulate their own interpretations of an artwork instead of presenting curator-authored wall text. Similarly, within the mitigation stage of the Privy workflow, recommended privacy mitigations were framed as questions rather than as a numbered list of solutions.

Across both systems, participants frequently responded to these provocations by pausing to consider their own perspectives before continuing. In some cases, this pause introduced mild frustration,

particularly when users expected the system to provide more immediate or authoritative information. In other cases, participants described the questioning as engaging or surprising, noting that it surfaced considerations they had not previously explored. These reactions suggest that carefully designed friction can prompt reflection, though it may also challenge user expectations shaped by more answer-oriented AI systems.

A second shared design decision was the use of **user-created content gates**, where progression through the system required participants to contribute their own thoughts before receiving additional AI support. From a methodological perspective, this design explicitly operationalized a human-in-the-loop approach: users were required to actively engage and externalize their reasoning before the system responded.

In both ArtBot and Privy, these gates were supported by lightweight instructional text within input fields to help users structure their responses. We observed that the quality and specificity of user input often shaped the relevance and usefulness of subsequent system output. This design choice reinforced the role of the GenAI system as a facilitator of thinking rather than a generator of standalone insight, while also making user effort a visible part of the interaction.

A third design decision involved **system framing**: i.e., how each system was presented to participants. Both ArtBot and Privy were presented as possessing relevant domain knowledge and as tools intended to support human thinking within their respective domains. This framing was generally effective, but it also elicited varied reactions depending on participants’ domain expertise.

Participants with greater subject-matter familiarity were more likely to challenge the system’s suggestions, tone, or assumptions, at times expressing disagreement or skepticism. Other participants, particularly those with less experience in the domain, described the systems as informative and well-grounded. These differing responses suggest that framing GenAI as a facilitator interacts with users’ prior knowledge, influencing whether provocations are perceived as helpful, restrictive, or misaligned.

Across both systems, we observe that altering the structure of human-AI interaction through questioning strategies, workflow constraints, and domain-grounded prompts can meaningfully influence how users externalize, refine, and develop their own thinking. These shared design principles form the basis for the comparative reflections we bring to the workshop discussion.

## 5 Discussion

Across both ArtBot and Privy, domain-specific provocations appeared to support a wider range of ideas and responses than generic

prompts would likely have produced. Participants drew on domain-relevant concepts, vocabularies, and concerns when responding to system questions, resulting in outputs that were better aligned with the goals of each task. Based on our observations and participant feedback, we speculate that if these systems relied on more general provocations (like simply asking “Why?” or asking for justification), the reflective gains observed in prior evaluations would have been diminished.

At the same time, the variability in user reactions points to an additional human factor that warrants further attention. Individual expectations, expertise, and tolerance for friction all shape how provocations are received. This suggests that while domain specificity strengthens GenAI’s role in supporting critical thinking, it does not fully determine user experience, and thereby highlights the need for adaptable or customizable provocation strategies in future systems.

One recurring pattern across both systems concerned participants’ underlying assumptions about what GenAI should do. Some participants approached the systems with a strong expectation that GenAI functions primarily as an automation tool. In ArtBot, these users expressed a desire for authoritative interpretations of artworks, preferring curator-approved explanations over dialogic exploration. Similarly, in Privy, such participants expected the system to automatically generate a completed end-to-end privacy review or privacy impact assessment (PIA) based on a brief system description.

From this perspective, the inclusion of GenAI raised a fundamental question: if the system already has access to relevant knowledge, how can users meaningfully engage in reflection or decision-making, if at all? This tension surfaced resistance to provocations designed to slow interaction or demand user input. As part of our workshop contribution, we aim to discuss how GenAI tools designers and developers might either accommodate or intentionally challenge entrenched views of GenAI as automation, drawing on prior work that explores how systems can surface, negotiate, or reframe users’ mental models at the point of interaction [18, 23]. The results of our case studies provide evidence supporting the interaction-automation conflict outlined by Wiberg & Berqvist [26] which is essentially a paradox in human-AI interaction [19] where the need to design meaningful interactions encounters AI’s increasing ability to automate processes.

An opposing, but equally consequential, pattern emerged among participants who expressed skepticism about GenAI’s capacity to contribute meaningfully to critical thinking. These users often referenced existing literature or public discourse critiquing GenAI’s limitations and were inclined to postpone engagement with AI-generated content for as long as possible. Several explicitly stated a preference for producing their own ideas before consulting the system.

While this behavior can be interpreted as desirable by signaling an awareness of the value of independent thinking, it also introduced new challenges. In some cases, participants underutilized features intended to expand perspective, surface blind spots, or support ideation. Prior explainable AI (XAI) work suggests that making system capabilities, limitations, and intent legible to users can help mitigate such resistance [8]. We see an opportunity to explore how GenAI systems might better signal their role as bounded,

reflective collaborators rather than as authoritative or generative replacements. We are inspired by work looking towards creating effective documentation, tutorials, and model cards [4, 17].

A third tension emerged around participants’ subject-matter familiarity. Users with substantial domain expertise occasionally perceived system provocations as condescending or redundant, particularly when the system challenged assumptions they already considered well understood. Conversely, novice users sometimes interpreted the same systems as possessing expert-level authority, potentially over-weighting their suggestions.

These reactions underscore the risk of a mismatch between system tone and user expertise and perception of the GenAI system [1]. They also suggest that simply embedding provocation (whether general or domain-specific) is insufficient. Instead, effective support for critical thinking may require systems that are responsive to signals of user expertise, confidence, or intent, adapting their provocations accordingly.

Taken together, these reflections point to a broader design challenge: critical thinking is not only shaped by what provocations are presented, but by what role the user believes the system plays, and who they believe themselves to be in relation to it. Domain specificity strengthens the relevance and interpretability of provocations, but individual differences in expectation, expertise, and trust continue to shape outcomes.

We bring these observations to the workshop to invite discussion around how domain-specific GenAI systems might better balance provocation, adaptability, and user agency. In particular, we are interested in exploring strategies for customizing or negotiating provocation styles in response to human-centered factors. In doing so we move beyond static designs toward more responsive tools for thought.

## 6 Conclusion

Through the design and evaluation of two distinct GenAI prototypes, we developed insights into how domain-specific design decisions and implementation strategies can support critical thinking in practice. Our experiences suggest that provocations grounded in domain knowledge can more effectively support users’ knowledge work. Specifically because they are designed to introduce productive friction, invite user interaction, and frame the system as a knowledgeable facilitator. At the same time, our work highlights the role of user preferences and expectations, underscoring the need for adaptive and flexible implementations of GenAI systems. We present these cases to the Tools for Thought workshop as concrete, evidence-informed examples, and to invite discussion around how domain grounding, provocation, and user agency might be balanced in future GenAI tools for critical thinking.

## Acknowledgments

We want to thank our collaborators on the original ArtBot and Privy papers, many people worked hard to bring these systems to life. We also want to thank our participants for their time and feedback.

## References

- [1] Alireza Amrollahi, Jiaqi Yang, Syed Muhammad Fazal e Hasan, and Basma Badredine. 2026. Knowledge workers’ trust and reception of generative AI’s advice

- in complex tasks. *International Journal of Information Management* 88 (2026), 103031. doi:10.1016/j.ijinfomgt.2026.103031
- [2] Peter Belcak, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. 2025. Small Language Models are the Future of Agentic AI. arXiv:2506.02153 [cs.AI] <https://arxiv.org/abs/2506.02153>
- [3] Alice Cai, Ian Arawjo, and Elena L. Glassman. 2024. Antagonistic AI. arXiv:2402.07350 [cs.AI] <https://arxiv.org/abs/2402.07350>
- [4] Anamaria Crisan, Margaret Drouhard, Jesse Vig, and Nazneen Rajani. 2022. Interactive Model Cards: A Human-Centered Approach to Model Documentation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 427–439. doi:10.1145/3531146.3533108
- [5] Sauvik Das, Cori Faklaris, Jason I Hong, Laura A Dabbish, et al. 2022. The security & privacy acceptance framework (spaf). *Foundations and Trends® in Privacy and Security* 5, 1-2 (2022), 1–143.
- [6] Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes. arXiv:2305.02301 [cs.CL] <https://arxiv.org/abs/2305.02301>
- [7] Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. 2025. Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. arXiv:2506.08872 [cs.AI] <https://arxiv.org/abs/2506.08872>
- [8] Samuli Laato, Miika Tiainen, A.K.M. Najmul Islam, and Matti Mäntymäki. 2022. How to explain AI systems to end users: a systematic literature review and research agenda. *Internet Research* 32, 7 (2022), 1–31. doi:10.1108/INTR-08-2021-0600
- [9] 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* (Yokohama, Japan) (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 1121, 22 pages. doi:10.1145/3706598.3713778
- [10] Hao-Ping (Hank) Lee, Yu-Ju Yang, Matthew Bilik, Isadora Krsek, Thomas Serban von Davier, Kyzyl Monteiro, Jason Lin, Shivani Agarwal, Jodi Forlizzi, and Sauvik Das. 2026. Privy: Envisioning and Mitigating Privacy Risks for Consumer-facing AI Product Concepts. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery. doi:10.1145/3772318.3791279
- [11] Hao-Ping (Hank) Lee, Yu-Ju Yang, Thomas Serban Von Davier, Jodi Forlizzi, and Sauvik Das. 2024. Deepfakes, Phrenology, Surveillance, and More! A Taxonomy of AI Privacy Risks. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 775, 19 pages. doi:10.1145/3613904.3642116
- [12] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 9459–9474. [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf)
- [13] Xingyu Bruce Liu, Haijun Xia, and Xiang Anthony Chen. 2025. Interacting with Thoughtful AI. arXiv:2502.18676 [cs.HC] <https://arxiv.org/abs/2502.18676>
- [14] Yang Liu, Bingjie Yan, Tianyuan Zou, Jianqing Zhang, Zixuan Gu, Jianbing Ding, Xidong Wang, Jingyi Li, Xiaozhou Ye, Ye Ouyang, Qiang Yang, and Ya-Qin Zhang. 2025. Towards Harnessing the Collaborative Power of Large and Small Models for Domain Tasks. arXiv:2504.17421 [cs.LG] <https://arxiv.org/abs/2504.17421>
- [15] Meta. 2024. Llama3.1. <https://llama.meta.com/>.
- [16] Soya Park, Hari Subramonyam, and Chinmay Kulkarni. 2024. Thinking Assistants: LLM-Based Conversational Assistants that Help Users Think By Asking rather than Answering. arXiv:2312.06024 [cs.HC] <https://arxiv.org/abs/2312.06024>
- [17] David Piorkowski, Daniel González, John Richards, and Stephanie Houde. 2020. Towards evaluating and eliciting high-quality documentation for intelligent systems. arXiv:2011.08774 [cs.SE] <https://arxiv.org/abs/2011.08774>
- [18] Jeba Rezwana and Mary Lou Maher. 2022. Understanding User Perceptions, Collaborative Experience and User Engagement in Different Human-AI Interaction Designs for Co-Creative Systems. In *Proceedings of the 14th Conference on Creativity and Cognition* (Venice, Italy) (CC '22). Association for Computing Machinery, New York, NY, USA, 38–48. doi:10.1145/3527927.3532789
- [19] Zainab Salma, Raquel Hijón-Neira, and Celeste Pizarro. 2025. Designing Co-Creative Systems: Five Paradoxes in Human-AI Collaboration. *Information* 16, 10 (2025). doi:10.3390/info16100909
- [20] Advait Sarkar, Xiaotong, Xu, Neil Toronto, Ian Drosos, and Christian Poelitz. 2024. When Copilot Becomes Autopilot: Generative AI's Critical Risk to Knowledge Work and a Critical Solution. arXiv:2412.15030 [cs.HC] <https://arxiv.org/abs/2412.15030>
- [21] David Skidmore and Kyoko Murakami. 2016. *Dialogic Pedagogy: An Introduction*. Multilingual Matters, 1–16. <http://ebookcentral.proquest.com/lib/bath/detail.action?docID=4614619>.
- [22] Thomas Serban von Davier, Aaron John Henry Larsen, Max Van Kleek, and Nigel Shadbolt. 2025. ArtBot: An Exploration into AI's Potential for Guiding Art Analysis. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI EA '25). Association for Computing Machinery, New York, NY, USA, Article 77, 11 pages. doi:10.1145/3706599.3720181
- [23] Xingyi Wang, Xiaozheng Wang, Sunyup Park, and Yaxing Yao. 2025. Mental Models of Generative AI Chatbot Ecosystems. In *Proceedings of the 30th International Conference on Intelligent User Interfaces* (IUI '25). Association for Computing Machinery, New York, NY, USA, 1016–1031. doi:10.1145/3708359.3712125
- [24] Zijie J. Wang, Chinmay Kulkarni, Lauren Wilcox, Michael Terry, and Michael Madaio. 2024. Farsight: Fostering Responsible AI Awareness During AI Application Prototyping. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 976, 40 pages. doi:10.1145/3613904.3642335
- [25] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.
- [26] Mikael Wiberg and Erik Stolterman Bergqvist. 2023. Automation of interaction—interaction design at the crossroads of user experience (UX) and artificial intelligence (AI). *Personal and Ubiquitous Computing* 27 (2023), 2281–2290. doi:10.1007/s00779-023-01779-0
- [27] Andre Ye, Jared Moore, Rose Novick, and Amy X. Zhang. 2024. Language Models as Critical Thinking Tools: A Case Study of Philosophers. arXiv:2404.04516 [cs.HC] <https://arxiv.org/abs/2404.04516>

## A System Screenshots

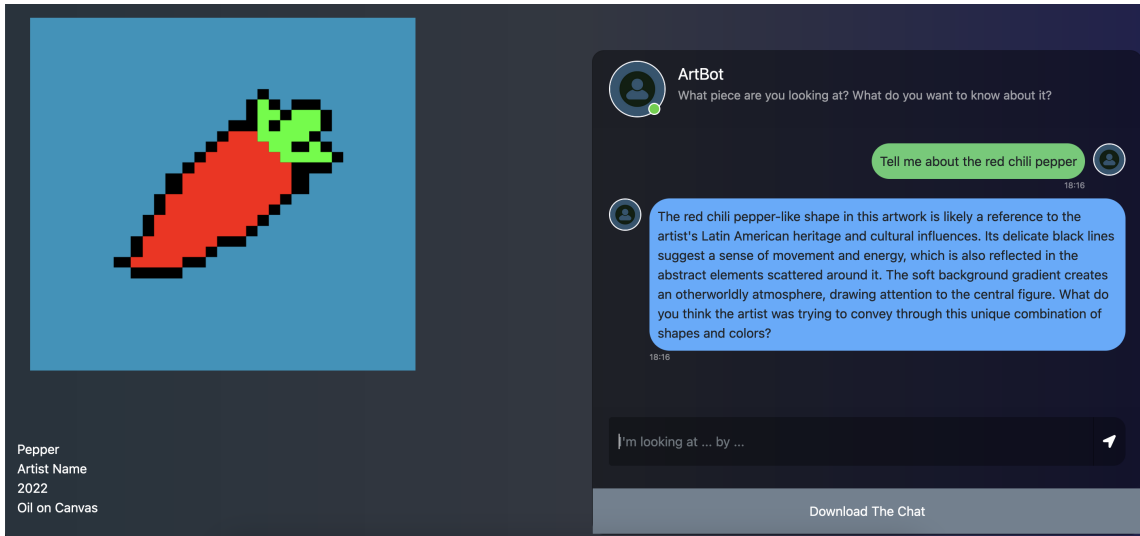


Figure 1: ArtBot is an LLM-powered tool that challenges participants with questions to encourage them to share their own interpretations of the artwork in the shared workspace.

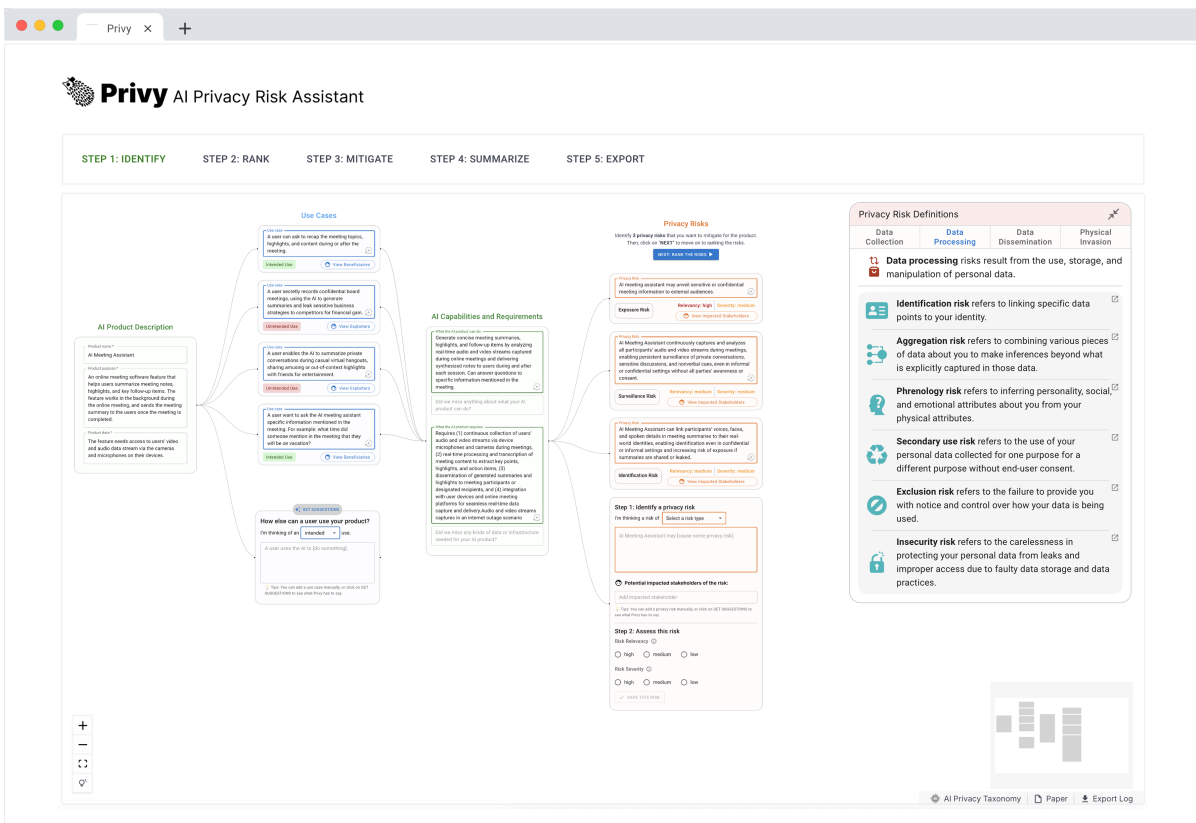


Figure 2: PRIVY is an LLM-powered tool that guides practitioners through structured privacy impact assessments to: (i) identify relevant risks in novel AI product concepts, and (ii) propose appropriate mitigations.