

Cocoa: Co-Planning and Co-Execution with AI Agents

K. J. Kevin Feng*
University of Washington
Seattle, Washington, USA
kjfeng@uw.edu

Tal August
University of Illinois
Urbana-Champaign
Urbana, Illinois, USA
talaugust6@gmail.com

Daniel S Weld
Allen Institute for AI
Seattle, Washington, USA
weld@cs.washington.edu

Kevin Pu*
University of Toronto
Toronto, Ontario, Canada
kevin.pu@mail.utoronto.ca

Pao Siangliulue
Allen Institute for AI
Seattle, Washington, USA
paos@allenai.org

Amy X. Zhang
University of Washington
Seattle, Washington, USA
axz@cs.uw.edu

Matt Latzke
Allen Institute for AI
Seattle, Washington, USA
mattl@allenai.org

Jonathan Bragg
Allen Institute for AI
Seattle, Washington, USA
jbragg@allenai.org

Joseph Chee Chang
Allen Institute for AI
Seattle, Washington, USA
josephc@allenai.org

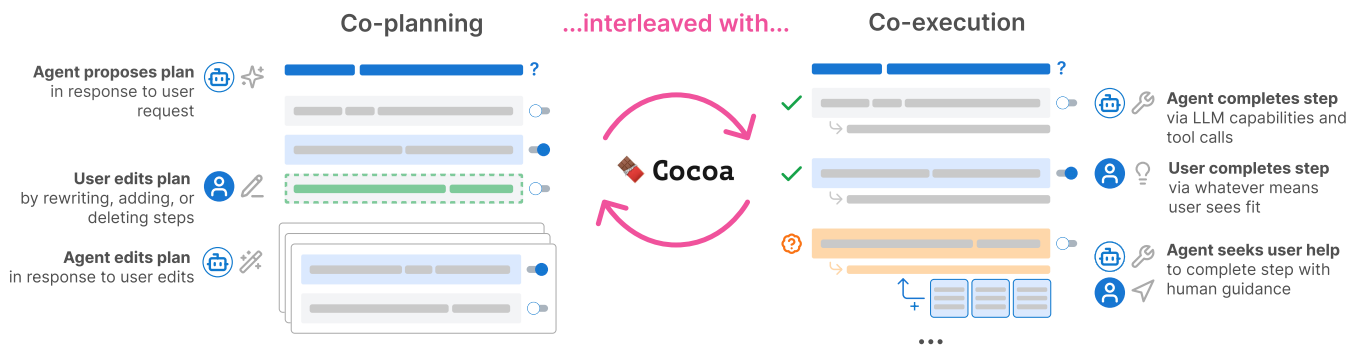


Figure 1: COCOA is an interactive system that facilitates interleaved co-planning and co-execution with AI agents in a document environment for scientific researchers. COCOA integrates AI agents into documents through an interface where a human user and an AI agent can jointly plan and execute plan steps using a shared representation of tasks, roles, and progress directly in the document.

Abstract

As AI agents take on increasingly long-running tasks involving sophisticated planning and execution, there is a corresponding need for novel interaction designs that enable deeper human-agent *collaboration*. However, most prior works leverage human interaction to fix “autonomous” workflows that have yet to become fully autonomous or rigidly treat planning and execution as separate stages. Based on a formative study with 9 researchers using AI to support their work, we propose a design that affords greater flexibility in collaboration, so that users can 1) *delegate agency* to the user or agent via a collaborative plan where individual steps can be assigned; and 2) *interleave* planning and execution so that

plans can adjust after partial execution. We introduce COCOA, a system that takes design inspiration from computational notebooks to support complex research tasks. A lab study ($n = 16$) found that COCOA enabled steerability without sacrificing ease-of-use, and a week-long field deployment ($n = 7$) showed how researchers collaborated with COCOA to accomplish real-world tasks.

ACM Reference Format:

K. J. Kevin Feng, Kevin Pu, Matt Latzke, Tal August, Pao Siangliulue, Jonathan Bragg, Daniel S Weld, Amy X. Zhang, and Joseph Chee Chang. 2026. Cocoa: Co-Planning and Co-Execution with AI Agents. 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, 23 pages. <https://doi.org/10.1145/3772318.3791673>

*Work completed during an internship at Ai2.



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).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791673>

1 Introduction

Since the advent of personal computing, researchers and practitioners in artificial intelligence (AI) and human-computer interaction (HCI) have set sights on developing AI agents that can collaborate with humans to perform complex computer-based tasks [34, 61, 66, 67]. For example, paradigms such as mixed-initiative

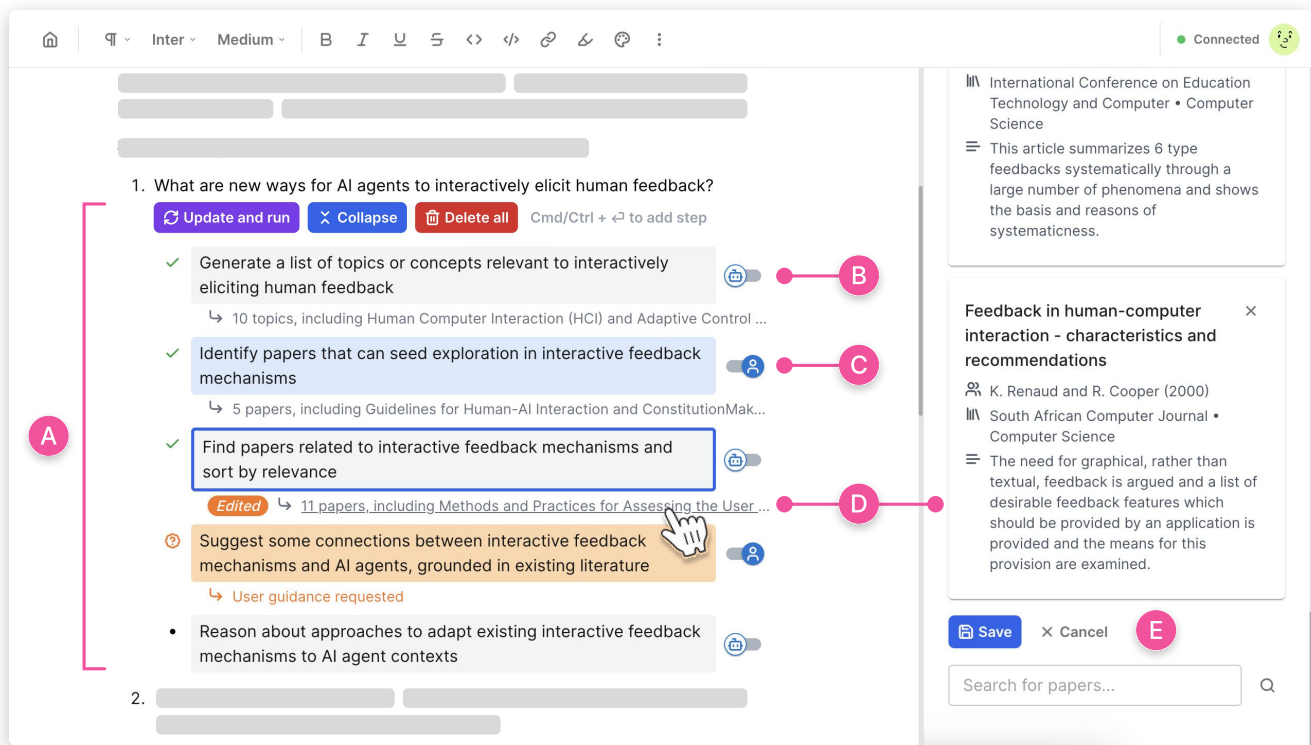


Figure 2: An overview of the COCOA user interface. An interactive plan (A) affords human-agent co-planning and co-execution: a researcher and the AI agent can collaboratively edit the plan in the document and execute the plan steps, similar to executing code cells in a computational notebook. Steps can be assigned to the AI agent (B) or the researcher (C). The researcher can freely edit the AI agent’s outputs in an interactive sidebar (D), such as adding relevant papers that the agent did not find (E) to help steer the agent with their feedback and expertise. In this example, the first three steps of a plan to summarize methods for human feedback elicitation have already been executed, and the agent is requesting guidance from the user in the next step.

interaction envisioned fluid human-AI collaboration with user-led direct data manipulation combined with AI-led autonomous task execution [34]. Recent advancements in large language models (LLMs) and agent frameworks involving reasoning, planning, memory, and tool use [50, 90, 100, 102, 110] have accelerated the development of fully autonomous AI agents to tackle complex tasks, including shopping online [108], writing software [40, 105, 107], performing “deep research” to assemble a report based on web sources [26, 76, 83], and general computer use [2, 75, 93]. Despite the push towards full autonomy, recent work has recognized the importance of deeper human-agent *collaboration* for not only bolstering agent performance on complex tasks [95], but also improving accessibility [81], user control and trust [35, 69], safety [75], and user agency more generally [21], leading to the development of more collaborative human-agent systems, such as for web browsing [35, 69, 81] and general productivity [75, 94].

However, many of these systems incorporate prescribed collaborative workflows that may limit their utility and steerability. Some works use *agent-guided workflows*: the agent dictates what input to elicit from the user, as well as when and how the elicitation

happens [35, 69, 75, 81, 94]. Users may intervene by manually interrupting or taking control of the agent [69, 75, 76], but these efforts are reactive and rely on persistent awareness of where the agent has erred. On the other hand, researchers have developed systems with *user-guided workflows*, where the user is in full control of task planning, while the agent executes scoped subtasks [1, 33, 47, 99]. Whether agent- or user-guided, these systems encode a *particular configuration* of human-agent collaboration—there are few affordances to *flexibly delegate agency* by altering who is guiding based on environmental context or user preferences [89].

Separately, researchers and industry practitioners have realized the value of decoupling human-agent planning and execution to improve control and avoid fruitless execution [33, 37, 47]. However, completely isolating the two into rigid stages can lead to poor results as well—plans often need to be iteratively adapted based on results from completed steps [69]. In contrast, fully autonomous frameworks like ReAct [110] *interleaves* planning and execution, and has shown benefits in allowing fully autonomous AI agents to readjust their plans after seeing partial execution results. We thus see a new design opportunity for human-agent interaction and

ask: **what might flexible configurations of agency in human-agent collaboration across planning and execution look like, and how will users make use of them in their own work?**

In this work, we explore a novel instantiation of mixed-initiative interaction [34] for *explicitly and flexibly delegating work* in human-agent hybrid workflows by *interleaving collaborative planning and execution*. We conduct this exploration in the domain of scientific research, a challenging domain for human-agent collaboration due to the prevalence of knowledge-intensive workflows and nuanced decisions made with domain-specific knowledge. Existing agentic research tools primarily focus on supporting human-agent collaboration either *before* or *after* task execution. One group of systems elicits user intent and preferences before execution [26, 76, 83] (e.g., asking clarification questions about users’ goals [76], allowing plan review and modification [26, 69]). Conversely, another group helps users verify and correct the agent after it has generated an output after executing its plan [31, 73]. In practice, because planning and execution are deeply interleaved, we aim to support human-agent collaboration throughout *fluid movements* between the two phases.

We introduce COCOA, an interactive system that implements our approach by supporting *hybrid human-agent workflows* for scientific researchers in a document editor—a common site for research work. This interleaving happens via the following three ways. First, COCOA supports human-agent **co-planning**: the agent, when invoked in the document, will first propose a plan of action that is fully interactive and integrated within the document. The user can then edit steps, add new ones, and assign each step to either themselves or the agent. Second, COCOA supports **human-agent co-execution**: drawing inspiration from the design of computational notebooks, COCOA allows the user and the agent to collaboratively complete plan steps, re-executing steps as desired. The user can interactively refine the agent’s intermediate outputs and also manually take over steps themselves to steer the workflow in a more desirable direction. Third, and most importantly, COCOA **interleaves co-planning and co-execution**: users can smoothly transition between the two and modify their plans based on outputs from execution, and vice versa.

We then evaluate COCOA through a lab study with 16 researchers, where we found that *interleaving co-planning and co-execution* in COCOA afforded greater agent steerability while maintaining ease of use compared to a strong chat baseline. We finally conducted a 7-day field deployment of COCOA with 7 researchers, where they found the *explicit delegation of work between themselves and the agent*, in addition to interleaved co-planning and co-execution, to be valuable in their day-to-day research.

Concretely, our work makes the following contributions:

- (1) A formative study with 9 researchers that motivates the design of interactive systems that can flexibly delegate work between the user and agent across task planning and execution.
- (2) COCOA, an interactive system that interleaves human-agent co-planning and co-execution in a document editor for scientific researchers.
- (3) A lab evaluation of COCOA with 16 researchers, where we found that participants frequently made use of interleaving co-planning and co-execution when completing research

tasks, suggesting that the two are mutually beneficial while enabling more granular steering of the agent.

- (4) A 7-day deployment study of COCOA with 7 researchers, where participants made use of self- and AI-assigned steps to flexibly delegate human-AI agency and embed their expertise into the workflow.

2 Related Work

2.1 Planning and Interactivity in LLM Agents

Prior work has identified core components in LLM agent architectures to consist of memory, reasoning, planning, and tool use [50, 63, 100, 102, 110, 114]. Central to LLM agents’ longitudinal operations is long-term planning, which demands capable reasoning abilities [32] and thus is often implemented by chain-of-thought (CoT)¹ [74, 85, 104, 110]. CoT provides a series of intermediate reasoning steps as exemplars in prompting to boost performance on complex reasoning tasks [74, 104]. CoT is crucial for facilitating task decomposition and therefore LLM agents’ planning capabilities [16, 85]. However, even with CoT, critical investigations of LLMs’ abilities to autonomously generate executable plans reveal limited success across diverse domains [101]. Indeed, although CoT can help improve model planning for tasks with well-defined, objective solutions that provide clear signals for reinforcement learning—e.g., numerical, tabular, and knowledge-based reasoning [50]—it may not be so effective in domains that require expert tacit knowledge to navigate ambiguous problem spaces with no single right answer [32]. Unlike model properties that show empirical improvement through scaling laws, limitations of planning in these domains may not resolve with scale alone as 1) tacit knowledge is not well-documented in training data and is thus difficult for a model to robustly learn [17, 22, 111], and 2) there may be no “correct” workflow for CoT to follow and verify the correctness post-hoc [41]. Scientific research is one such domain [111].

In light of this, recent work has recognized the need to interactively incorporate user feedback for improving LLM agents’ planning capabilities and beyond [35, 47, 55, 69, 75, 81, 94], particularly within scientific discovery [27, 39, 68]. For example, Lawley and MacLellan [55] architected an approach where user interaction is used to guide the model in planning for unseen tasks on-the-fly using a hierarchical network of smaller actions. OpenAI’s Operator [75] allows users to “take control” of the agent to perform tasks manually. Yet, interactive techniques in this area are still nascent, despite the new challenges identified for human-agent communication [7]. Interfaces for LLM assistants (e.g., AgentGPT [87], Devin [105]) have been primarily limited to chat interfaces targeted at *monitoring* agent activity rather than *empowering the user to proactively collaborate* with the agent.

An emerging body of work has started to experiment with more interactive techniques for agents. He et al. [33] separated agent planning from execution and found that user involvement during execution can significantly boost agent performance, but poses a high cognitive load. Kazemitabaar et al. [47] developed interfaces for data analysts to edit an execution plan of an LLM to provide more control points for steering behavior, while Google’s Deep

¹We use CoT as an umbrella term encompassing chain-of-thought, trees-of-thought [109, 115], and related methods.

Research [26] can re-generate a plan if the user provides feedback via prompting. However, these interactions are *corrective* rather than *collaborative*—that is, users would typically engage in these interactions to correct agent behavior post-hoc rather than *proactively iterate back-and-forth* with the agent at multiple points in the workflow. We do this by introducing affordances for collaborative human-AI planning and execution via a shared representation of tasks and delegation of agency. While our work shares some conceptual similarities with Kazemitabaar et al. [47], we focus on enabling human-AI co-agency *across steps* in multi-step plans rather than revealing hidden model reasoning *at each step*.

2.2 Computational Notebooks

Donald Knuth’s vision of *literate programming* interleaves computer programs with natural language documentation [53]. This vision later gave rise to the computational notebook paradigm, which organizes program imperatives, input data, and rich outputs (e.g., tables, visualizations, interactive widgets) into linearly arranged cells [15, 54]. Computational notebooks are a well-studied paradigm in HCI (e.g., [5, 15, 18, 59, 62, 72, 113]), with most prior work focusing on how data scientists leverage them in their workflows.

In our work, we take inspiration from design decisions made for computational notebooks because of the many parallels between workflows in notebooks and mixed-initiative task completion systems [28, 34, 106]. For instance, data scientists often break down a high-level data analysis task into executable cells in computational notebooks [48, 88]. As they execute each cell and inspect outputs, they can verify the quality and sensibility of outputs before moving forward in their analysis, or go back to modify existing code cells before execution to interleave between programming, execution, and output examination [49]. Finally, the notebook can serve as documentation and communication of their activities [59].

By drawing these parallels, we can reveal new opportunities to design interactive workflows and interfaces for LLM agents. Just like how a data scientist may easily add new cells in a notebook or reconfigure existing cells to adapt their analysis plan, a user may interactively edit an agent’s plan of execution to better steer the agent in productive directions. Furthermore, an agentic system may simultaneously be more usable and resource-efficient if a user could direct an agent to iterate on a subtask to improve its output before continuing onto subtasks dependent on that output. This also demands new ways of viewing and editing agent outputs on subtasks, which increases opportunities for human input over existing approaches such as simply logging agent actions [4, 77, 87]. Our work exploits these parallels to contribute new design patterns for steering agentic systems. We also apply novel extensions of these parallels for human-agent collaboration: agent-proposed plans (i.e., dynamic AI-generated notebook cells) and agent vs. user steps (i.e., AI executes some cells while the user executes others).

2.3 AI for Scientific Research

Significant efforts in recent years have advanced how AI can be used to support scientific research. Of particular interest to us are works that leverage these recent advances to help researchers with planning and execution of *literature-focused tasks*. Planning a research project is a cognitively demanding, complex process consisting of

iterative cycles of divergent and convergent thinking grounded in literature review [14, 19, 80], making it an ideal scenario to explore human-agent collaboration. One major thread of prior research has focused on leveraging AI models as tools or components in user-driven systems. These include tools for paper reading [14, 86] and skimming [25], literature review [58, 80], paper recommendation [43, 44], information retrieval and sensemaking [13, 24, 45, 103], and ideation [6, 30, 38]. On the other end of the spectrum are fully automated end-to-end systems where an AI agent has much more agency in planning and executing literature review tasks [20, 92], or even to attempt carrying out entire research projects on their own [26, 27, 39, 64, 76, 83, 91].

As AI systems become more capable, we ask: Is the generation of full research artifacts (e.g., research questions, proposals) the most desired and promising paradigm to assist researchers? Indeed, when working with other (human) collaborators, researchers often find richness in co-evolving partially completed artifacts [41, 56, 97, 98] and answering or asking questions that stimulate critical thinking and reflection [78, 79]. Further, some fully automatically generated artifacts are still found to be of lower quality than those generated with human involvement [39]. Schmidgall et al. found that even systems designed for end-to-end research automation produce significantly higher quality outputs with human involvement at each stage [91]. In this work, we build on recent advances in AI research support tools, but rather than focusing on generating pre-determined research artifacts, we contribute an interaction design pattern that allows for flexible specification of the final artifact and collaboration between a researcher and an AI agent.

3 Formative Study

Building on prior work, we aim to enhance collaboration between humans and AI agents using a shared, interactive operational representation. We focus on scientific research—a challenging domain for human-agent collaboration given its complex, multi-step workflows and the nuanced decisions researchers make upon processing information-dense scientific papers. Specifically, we investigate the potential for project documents—continuous records of ideas, updates, and tasks—to act as the agent environment. We thus conducted a formative study to answer the following research questions:

- F1:** What are the properties and opportunities of researchers’ project documents for human-agent collaboration?
- F2:** How do researchers engage in planning within project documents?
- F3:** How would researchers prefer to collaborate with an AI agent in their research workflows?

3.1 Participants, Procedure, and Data Analysis

We recruited 9 Ph.D. students (detailed demographics in Table 3 of Appendix A) through an interest form sent to a research organization’s internal Slack channel and the authors’ professional connections. We targeted Ph.D. students as they often lead the detailed planning and execution of research projects and are the primary editors of project documents.

All studies were conducted virtually over Google Meet and lasted 60 minutes. Before the study, we collected from participants an

active or past research project document and a brief description of their research interests. The study was divided into 3 parts:

- (1) An activity involving the participant’s project document to better understand their current planning behavior (25 minutes).
- (2) An activity with a Wizard-of-Oz (WoZ) design probe to explore how researchers plan and explore research ideas in a document editor with LLM support via a chatbot (25 minutes).
- (3) An exit interview where researchers reflected on their experience with the probe and using LLMs in research (10 minutes).

Details for each part can be found in Appendix B. Each participant was given a \$35 USD honorarium after the study. The study was reviewed and exempted by our organization’s internal review board. We qualitatively analyzed project documents produced from the study probe activity alongside study transcripts. Our data analysis methods can be found in Appendix C.

3.2 Findings

In this section, we redact any project-specific details for privacy and intellectual confidentiality.

3.2.1 [F1 & F3] Project documents are promising environments for human-agent collaboration. Participants’ project documents acted as a “hub” for their projects. These documents commonly included meeting notes, *planned to-do items*, *progress updates* tied to to-do items, *research questions* to discuss with collaborators, and *links to other documents* detailing a particular project component in greater depth. These documents often served as “scratch paper” for researchers to informally leave a trace of their reasoning and high-level goals, conduct short-term and long-term planning, and progress tracking of these plans. Most documents (P1, P2–4, P7–P9) contained one or more problem statements that motivated the project and stated its core contributions. From there, participants listed subgoals and ideas for planned exploration (all participants).

Participants expressed a desire for tighter integration between an AI agent and their project documents in their research workflows for two main reasons. **First**, they wanted to receive in-situ AI support as they plan and execute research tasks in their documents. P1 and P2 both envisioned an AI system “*actively engaging with the content I’m writing [...] after I write each statement, the system could retrieve relevant papers that could provide background or related work for me to read more.*” P9 appreciated “*the ability to lay out the steps when I start something. [It] was very helpful [for] visualizing the outcome.*” **Second**, participants wanted the AI agent to have access to the broader context that already existed in their project documents. When using the probe, P2 was concerned that they would need to “*prompt [the chatbot] with a lot of background,*” but realized they “*might have [the background] written down*” in their document. P3 desired closer document integration but warned that the agent may clutter their document. They suggested: “*having something on the side that does not flow into where I am writing, like a side [panel]—if I want something from it, I want to bring it back into my doc.*”

3.2.2 [F2] Literature search and understanding play a central role in research planning. In all project documents, we saw places where participants intended to initiate a multi-step plan (often via a *how?* and *can we?* question) to address different research tasks. The most common category of tasks are ones that were literature-augmented (e.g., literature search and understanding), which are often combined with a wide variety of other research actions, such as experiment/artifact design (P3), experiment/code execution (P4), data inspection and synthesis (P7), communicating and discussing results (P5), and ideation (P2). Literature was not only a significant component of *existing* plans, but was important for informing *future* plans. Some participants’ plans (e.g., P3, P4, P7) were formed from hypotheses they hoped to verify, which prior work may have already addressed; P7 pointed out that the hypotheses they had in mind can “*actually be proven from previous papers*” so they saw literature review as a planning aid. Consequently, all participants expressed a desire to use AI to help with *literature-augmented tasks*—exploring and understanding relevant literature to inform decision-making. This finding echoes recent surveys on how researchers leverage LLMs to conduct research, where the most frequent usage category was information-seeking [60].

3.2.3 [F3] Participants preferred to perform higher-level reasoning and synthesis themselves. In addition to highlighting tasks where they wanted AI assistance, participants also spoke about tasks they did not want AI to automate away. In particular, participants saw higher-level reasoning and information synthesis as critical tasks they wanted to do themselves. They were also often unsatisfied with AI’s outputs on these tasks. P3 explains: “*this tinkering process around reading and playing around with things is what gives you the ideas. I don’t know if I want those things automated because the process is as helpful as the final result.*” P5 agreed that they “*wouldn’t want it to be making the final decisions for sure. Just give me inspiration for where I can go.*” Specifically, they shared that “*the main thing I worry about is feeling enough ownership*” if AI makes more consequential decisions. An overeager AI assistant that attempts to perform tasks researchers prefer to do themselves is frustrating because it does not complement the user’s work and instead creates undesired noise for them to filter through.

3.3 Design Goals

We synthesize our formative study findings into three design goals for an interactive system that facilitates meaningful collaboration between researchers and AI agents by using plans as a shared representation.

DG1: Provide opportunities for human-agent collaboration across planning and execution. Agent developers have called for decoupling planning and execution to avoid fruitless execution [37]. Indeed, participants discussed *planning* (i.e., breaking down a problem into smaller, more actionable tasks like they did in their project documents) and *execution* (i.e., completing steps in the plan and refining outputs, with or without AI assistance) as separate but complementary actions. We thus see planning and execution as two distinct stages for fertile researcher-AI collaboration. Because these two stages are closely intertwined and synergistic, we aim to interleave the two in our system.

DG2: Allow flexible delegation of work between researcher and agent. Researchers may not want to delegate all parts of a plan to AI [F2, F3]. We also observed that this preference may also be context-dependent and constantly shifting. Our system uses a concrete implementation of a theory of flexible delegation proposed by Satyanarayan [89].

DG3: Integrate seamlessly into a document environment. Project documents contain rich externalizations of researchers' workflow, thought processes, and plans (e.g., "to-dos") [F1]. This is key information an AI agent can use to better assist the researcher. The document is also an appealing environment for researcher-agent collaboration, but requires careful information management strategies to prevent unwanted distractions [F3]. Thus, we aim to 1) allow the researcher to use familiar document editing affordances to interact with the agent, and 2) strategically manage outputs to avoid excessively cluttering the document.

4 Cocoa: System Walkthrough and Implementation

We present COCOA, a system that embeds an AI agent into a document editor and provide an affordance we call *interactive plans* that allow users to collaboratively plan (**co-plan**) with the agent—the agent proposes an initial plan of execution to tackle a user request that the user can edit to their liking. Then, users can collaboratively execute (**co-execute**) the plan with the agent—the user and the agent can build off each other's work to synergize human and AI capabilities, and provide flexible delegation of human and AI work.

To illustrate the features and functionality of COCOA, we follow the journey of Nora, a fictional researcher in human-AI interaction, as she uses her project document to further explore open questions in interactive interfaces for AI agents.

4.1 Co-Planning

Nora's project document is an informal, reverse chronological log of research progress and includes information such as meeting notes, rough ideas, links to literature, and questions for herself and her collaborators. She has a specific question in her document that she wants to explore further, but being new to the topic, she is not sure how to proceed on her own: "*What are new ways AI agents can interactively elicit human feedback?*"

4.1.1 Reviewing and selecting an initial plan. To initiate co-planning, she highlights the question to invoke the agent on it using a button that appears in a floating menu. The agent analyzes the question within the context of her entire document and proposes a few different plans via a **plan selector** UI. Noticing one of the proposed plans started with generating a list of more detailed concepts, which would potentially give her the opportunity to learn and tune the specific paper search directions, Nora selects the plan, which then becomes fully interactive within the document (Figure 3).

4.1.2 Assigning "Agent steps" and "User steps". The initial plan included steps for brainstorming relevant topics, searching for papers, and making connections between papers (Figure 3). To balance between effort and control, Nora uses the **step assignment toggle** (Figure 2 B&C, Figure 3 right) to assign tasks that are relatively low

risk but high effort to the agent, such as searching for papers and expanding on some of her preliminary ideas. In contrast, she assigned tasks that were more consequential or required higher-level thinking to herself, such as identifying seed papers for the agent to explore more broadly or making insightful connections between multiple papers.

4.1.3 Plan Editing and Interactive Replanning. Upon further review of the plan, Nora is now also curious about the prominent authors in the area and wish to learn about their papers first. She edits the **step description** to reflect this. The step description can be edited just like editing any other bulleted list in her document, making this intuitive for Nora. Similarly, Nora can also easily **add and delete steps from the plan**. For example, she notices that the agent has proposed a summary step that may be irrelevant, so she highlights the step and hits Backspace to delete it. Finally, she uses the keyboard shortcut Cmd/Ctrl+Enter while selecting the last step to add a new step below it, and edits the step description of the newly added step accordingly. These interactions mirror how one would edit items in a bulleted list in text editors, making them intuitive for Nora.

While COCOA allowed Nora full control over how she wishes to manually edit the plan, this process can be effortful for users. Therefore, the system provides support for *replanning* in the following ways: *First*, when editing an individual step, Nora can request the system to provide suggestions for alternatives, informed by the previous steps and document context. *Second*, after the user has edited some plan step, the plan may become incoherent. For example, a plan step that searches for additional relevant papers given a set of seed papers might no longer make sense if the previous step were edited to produce a list of prominent authors. In this case, COCOA **automatically detects this and replans subsequent steps** accordingly by removing those steps and "autocompleting" the plan with new ones (Figure 4).

4.2 Co-Execution

Now that Nora has made her desired edits to the plan, she is ready to co-execute it with the agent.

4.2.1 Continuous and stepwise execution. Drawing inspiration from computational notebooks, there are two modes of co-execution: continuous and stepwise. Each step has a **step completion indicator** that indicates whether the step has not yet been run (a bullet point), is in progress (a spinner), requires user input (a question mark badge), or is complete (a checkmark). This notebook-inspired design allows Nora to dispatch the agent to run autonomously on parts of a plan she trusts the agent to complete on its own. She can also iterate on specific parts of a plan by examining and potentially amending the output of each plan step before continuing.

Nora wishes to run through the entire plan once to see the output. She triggers *continuous co-execution* with the **[Run all]** button at the top of the plan. In **continuous co-execution**, the agent will automatically continue onto the next step as soon as it (or the user) has completed the previous one, until it reaches a user step or the end of the plan. By contrast, if she wishes to iterate and refine the outputs of individual steps before the agent moves onto the next one, she can opt for **stepwise co-execution**.

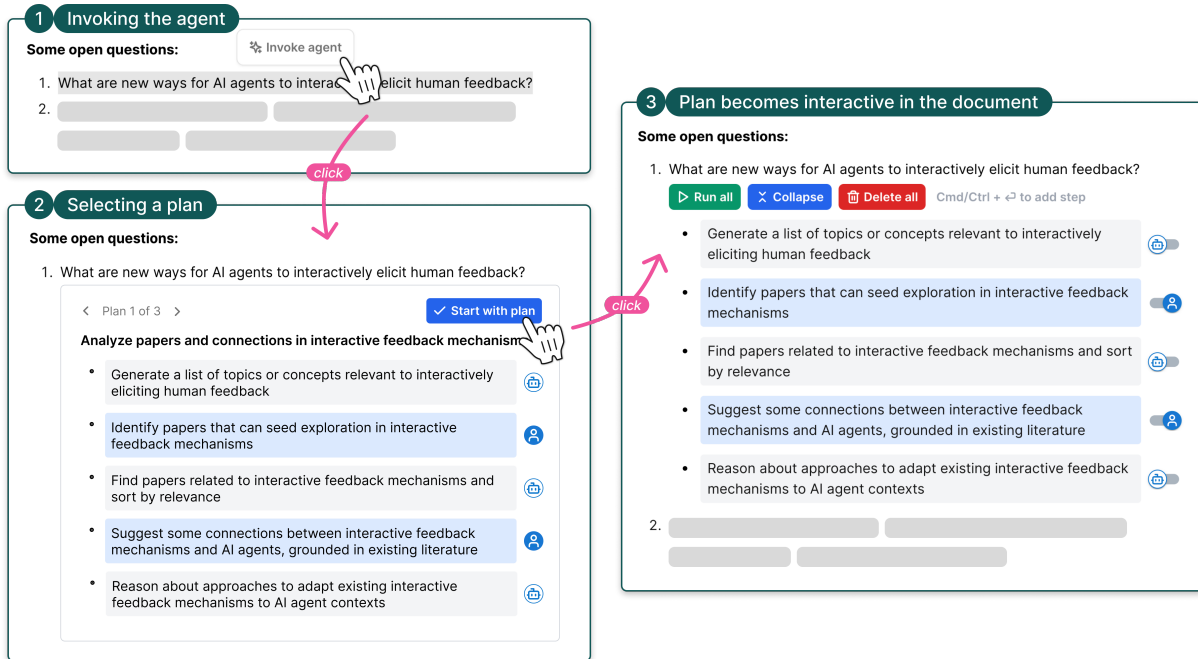


Figure 3: A user invokes the agent on a piece of text in the document by clicking on the “Invoke agent” button that appears on hover whenever text is highlighted. The agent will use the highlighted text and context from elsewhere in the document to propose a series of plans, displayed in a plan selector that appears under the highlighted text. Once the user selects a plan, it becomes fully interactive in the document.

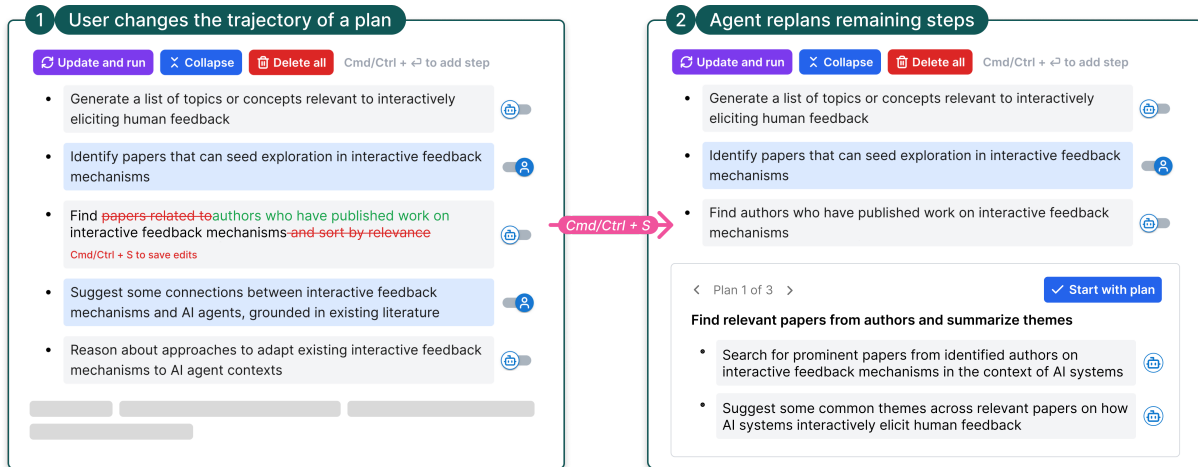


Figure 4: If major changes are made to a step that changes the rest of the plan’s trajectory, COCOA detects this and will trigger replanning. Replanning replaces subsequent steps with ones the agent suggested for “autocompleting” the plan.

Nora can flexibly switch from stepwise to continuous co-execution at any point by clicking the [Run remaining] button at the top of the plan. Conversely, Nora can switch from continuous to stepwise co-execution by clicking on the [Pause after this step] button during execution and manually running subsequent steps.²

²Cocoa does not permit out-of-order stepwise execution because some steps may rely on the outputs of previous ones. We also note that out-of-order execution in

4.2.2 *Editing outputs of agent steps.* As the agent completed a paper search step, Nora clicks into that step to examine its output. This opens a sidebar to the right that is populated with interactive paper cards (Figure 5). Nora can further guide the agent with her expertise by removing papers she considers irrelevant and/or

computational notebooks is a major user pain point identified in prior academic work [15, 54] and industry reports [29, 42].

adding papers the agent missed using a built-in paper search functionality connected to the Semantic Scholar academic database [52]. For now, she removes a couple of papers in the output that were not published recently and were published in a less relevant venue.

The interactive UI in the sidebar dynamically adapts according to the step's *output type*, which falls into one of five categories in COCOA: **papers**, **authors**, **topics**, **entities**, and **text**. Papers, authors, and topics are drawn from Semantic Scholar and are displayed as interactive cards that Nora can add, delete, or search for using the sidebar's built-in Semantic Scholar search. Entities are lists of items in natural language (search queries, research questions, etc.) and are displayed as editable pills. Text is natural language text displayed in an editable text box.

If the agent fails to complete its step (e.g., it finds no papers or runs into an error), COCOA will also proactively alert Nora so that she can complete the step instead, or retry with the agent.

4.2.3 Completing user steps. Before executing the plan, Nora had assigned step 2 to herself during co-planning. She wanted the agent to start a broad search of many relevant topics and concepts, but then wanted to contribute some seed papers she knew would be relevant to steer the agent's subsequent actions.

COCOA has now paused at step 2, highlighting the step in orange to attract Nora's attention and to seek her input. When she clicks into that step, the same interactive sidebar opens to the right (Figure 5). She could then manually review the topics generated in step 1 and uses the paper search feature built into the sidebar to browse relevant papers and add them to the output of step 2. Additionally, she recalls a few papers that a collaborator had recommended to her previously, finds them using the paper search feature, and adds them to the output.

4.2.4 Final output panel. After the final plan step has been executed, the agent adds a **plan output panel** (Figure 6) containing a modified version of the last step's output into the document just below the plan. The modification involves rewriting the output, given the context of the plan and its outputs, to more directly connect to the original user request. This brings the results of running the plan into the document so Nora can easily reference it in the context of other content. The panel is collapsed by default to save space and reduce clutter, but she can expand it to copy and paste parts of the output she finds useful and add them to her document. If she does not want the panel in her document at all, she can simply delete it; she can still access the outputs of all steps in the sidebar.

4.3 Interleaving Co-Planning and Co-Execution

In addition to supporting co-planning and co-execution independently, COCOA is designed to support **synergistic interleaving of the two**. As Nora co-executes the plan with the agent, she encounters a relevant author that the agent identified and wishes to further explore the author's papers. She edits the following step's plan description to satisfy this. She reruns that step and receives a new batch of papers, and the agent automatically reruns the rest of the plan with the updated output. She is also not satisfied with the agent's interpretation of common themes and believes that she can surface deeper insights by reading the papers herself, so she toggles that step to be a user step. After reading the papers in detail,

she returns to her document and jots down her insights. She now finds the next step repetitive, so she deletes it. Overall, her use of COCOA is highly iterative, and she smoothly transitions between co-planning and co-execution.

When designing COCOA, we were inspired by the computational notebook paradigm supported scientists via tight coupling of code and output, enabling rapid iteration and steering of complex scripts. To summarize this connection, we map COCOA's features to those in computational notebooks in Table 1.

4.4 Implementation Details

COCOA is implemented as a web application with a Next.js and TypeScript frontend communicating with a Flask backend. The frontend uses the Tiptap³ framework for the main document editor. Each document supports synchronous collaboration via Hocuspocus⁴ and all changes are auto-saved to Tiptap Cloud, which also serves as storage for all participant documents. Interactive components within the document editor are implemented as custom Tiptap extensions in TypeScript.

The Flask backend orchestrates LLM activity with calls to a custom tool-calling LLM agent⁵ powered by GPT-4o and scaffolded by LangChain.⁶ Our agent is designed to assist researchers through literature-grounded response generation. To do this, we implemented tools that give our agent access to millions of research papers via publically available Semantic Scholar APIs.⁷ This API is powered by The Semantic Scholar Open Data Platform [52] which consists of a state-of-the-art PDF extraction and knowledge graph normalization pipeline. This gave our agent access to both the full text of millions of research papers and also their rich metadata, such as authors, fields of studies, and publication venues [52]. More implementation details can be found in Appendix D.

5 Lab Study

We conducted a within-subjects task-based lab study with 16 researchers to evaluate COCOA against a custom chat baseline. Specifically, our study aimed to address the following research questions:

- L1:** How does COCOA compare to our chat baseline for ease of use, steerability, and general utility⁸ in research project documents?
- L2:** How does interleaved co-planning and co-execution contribute to ease of use, steerability, and general utility in COCOA, if there is a significant improvement over the baseline?
- L3:** What kinds of steps did researchers wish to assign to an agent and themselves in practice?

5.1 Participants

We recruited 16 Ph.D. and postdoctoral researchers (14 Ph.D.s, 2 postdocs; 10 female, 6 male) in computer science (CS) or CS-adjacent areas via university mailing lists, word of mouth, social media recruitment messages (on Twitter/X, Mastodon, Bluesky), and personal connections. We recruited on a first-come, first-serve basis

³<https://tiptap.dev/>

⁴<https://tiptap.dev/docs/hocuspocus/>

⁵<https://platform.openai.com/docs/guides/function-calling>

⁶<https://www.langchain.com/>

⁷<https://www.semanticscholar.org/product/api>

⁸General utility is a composite metric created from Q1, Q2, and Q4 from Appendix H.

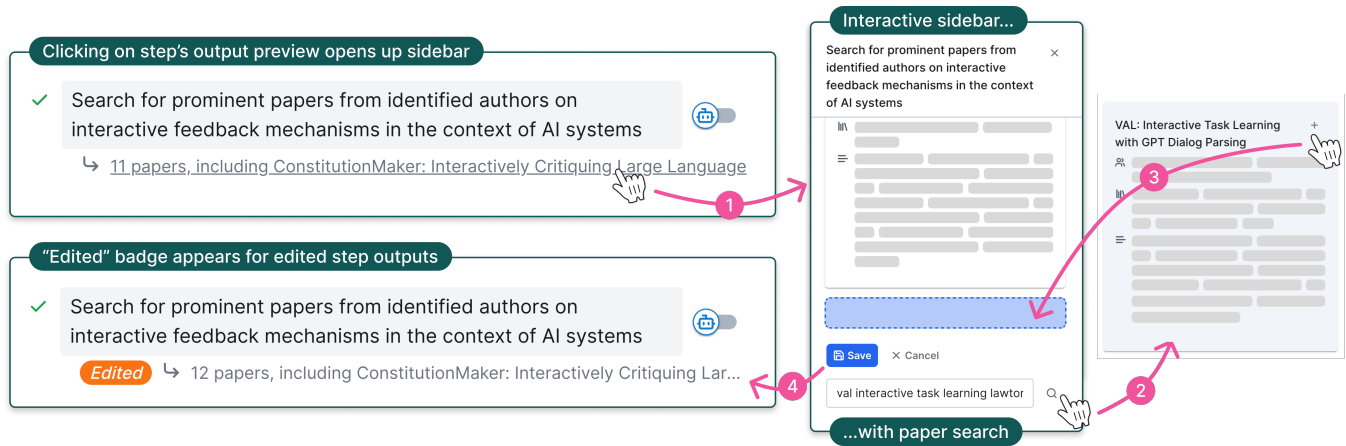


Figure 5: Users can click the output preview text beneath the plan step to access the interactive sidebar (1). By editing agent outputs in the interactive sidebar (2), users draw from their expertise to edit outputs, such as adding papers the agent might have missed (3), to guide the agent. Once edited, an indicator will appear in the output’s text preview below the step (4).

CocOA Feature	Computational Notebook Feature
Editable plan steps that can be composed into a larger, more complex task. Users can choose between stepwise and continuous execution. Users can view and improve step outputs within the plan. Users can flexibly switch between improving the plan and executing the plan steps. Plan steps assigned as user steps pause execution so users can perform tasks themselves. Completed plan can serve as documentation of agent & human work.	Code cells that can be composed into a larger, more complex program. Users can choose between running individual cells or run all. Users can view code outputs within the notebook. Users can flexibly switch between improving the code and executing code cells. Break points pause execution so users can examine current code outputs; all cells are executed programmatically. Completed notebook can serve as documentation of analyst’s work.

Table 1: Main connections between features in CocOA and those in computational notebooks.

as we conducted our studies and closed recruitment when we approached data saturation. Further details about our recruitment process, pilot studies, and participants can be found in Appendix F and Table 4 of Appendix E.

5.2 Baseline System

As our primary research question is to better understand the costs and benefits of CocOA’s co-planning and co-execution affordances, the baseline system (Figure 7) was a common chat interface powered by the same underlying LLM agent: a custom tool-calling agent powered by GPT-4o with access to the same state-of-the-art tools for searching, analyzing, summarizing, and comparing literature across millions of papers on Semantic Scholar [52], in addition to standard LLM capabilities. See D for more details about the agent. The design of the chat interface closely resembled that of existing AI-powered research systems (e.g., Assist mode in Scite [92], Google Scholar Labs, various deep research systems), and popular LLM chatbots (e.g., ChatGPT, Claude). During the study, participants positioned the chat interface beside the baseline document editor, which is the same editor that CocOA uses, except we removed the option to invoke the agent to eliminate all co-planning and co-execution interactions.

We chose this baseline because it mimicked the dominant interaction paradigm implemented in AI-powered research tools at the time of our study and allowed for strong comparisons on aspects like ease of use. We recognize that variations of this baseline, such as forcing the agent to show an editable plan before execution (e.g., [33]), are also strong baselines that can help further isolate the factors contributing to our results. We discuss ablations CocOA as alternative baselines in Section 9.

5.3 Research Task and Study Procedure

To make the study more realistic and grounded in real-world research projects, tasks for this study were open questions or unfinished items from participants’ existing project documents. Details for how we controlled for the length of the study and ensured the two tasks were valid and comparable can be found in Appendix G. Tasks were randomly assigned to CocOA and the baseline.

The first author (study facilitator) conducted 1:1 studies with the 16 participants over Google Meet between October and December 2024. Each study was 90 minutes in length. The study started off with introductions and a brief overview, followed by the two tasks, which were counterbalanced across participants. Detailed procedures for each condition can be found in Appendix G. After each condition, participants filled out a short evaluation form with

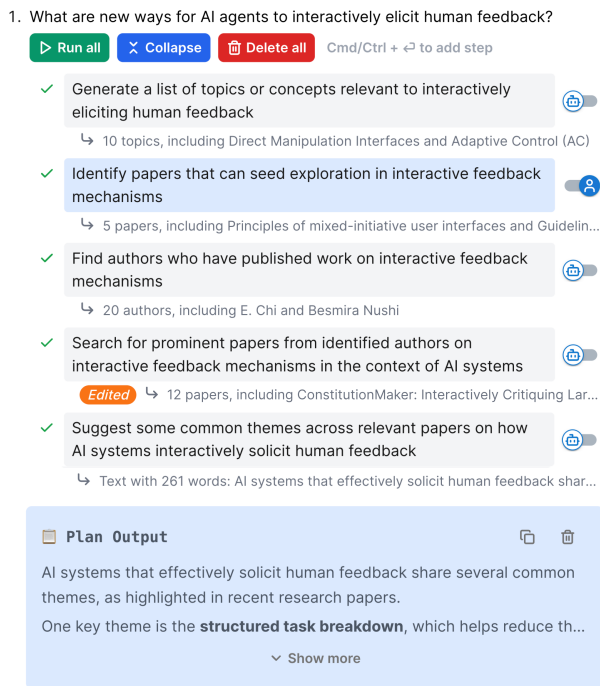


Figure 6: Once execution of the entire plan is complete, the agent inserts a plan output panel in the document with the last step’s output. This allows the user to view the plan’s results within the context of their document. The panel can be expanded or collapsed, and remains in the document even if the plan itself is collapsed for easy access.

5-point Likert scale questions for that condition (see Appendix H). These questions were inspired by the System Usability Scale [10] but adapted for the research context.

The task-based portion of the study was 75 minutes. The study concluded with a semi-structured interview that lasted around 15 minutes. Participants were asked to reflect on their experiences across the two systems and discuss the pros and cons of both. They were also asked about particular decisions the study facilitator observed when using each system.

Participants received a \$75 USD honorarium upon completion of the study. All studies were recorded and transcribed by Google Meet. This study was reviewed and approved by the Allen Institute of AI’s IRB.

5.4 Data Analysis

The first author analyzed the recording transcripts using Braun and Clarke’s reflexive thematic analysis [9]. This approach uses a hybrid inductive-deductive approach to iteratively surface codes and themes across the data. We paid close attention to codes related to participants’ comparisons of their experiences across the two systems and gradually grouped them into themes. NotebookLM⁹ was used to help iterate on themes and discover new ones. We

⁹<https://notebooklm.google.com>

performed statistical tests of participants’ ratings on the evaluation forms with the Wilcoxon signed-rank test (with the Bonferroni correction for our multi-rating analysis on significant results) given the non-parametric nature of our data. We saved all documents used in the study as well as all conversation histories in the baseline system, and referred back to them when needed.

The first author also coded all video recordings from the study to note the timestamps at which participants engaged and disengaged in a particular interaction. In COCOA, these interactions were **co-planning** (see 4.1), **co-execution** (see 4.2), and **output inspection** (when the participant passively inspected the system’s outputs). The closest equivalent interactions were used as codes in the baseline: **prompting for planning** (instructing the agent to perform an action without any attempts to modify the agent’s earlier outputs), **prompting for editing**¹⁰ (instructing the agent to modify or only consider a subset of its output in future actions), and **output inspection**.

We established these equivalences based on the functional purpose of the interactions: co-planning and prompting for planning both instruct the agent to produce new outputs, while co-execution and prompting for editing aim to improve an existing output. We note that there is not perfect correspondence between interactions across the two systems. For example, co-execution is not only about refining agent outputs, but also giving users agency to complete tasks they assigned for themselves. We acknowledge this by shading the interactions in COCOA and the baseline with slightly different colours in Fig. 8. However, because co-planning and co-execution subsume the functional roles of their baseline counterparts in the context of our study, we consider them appropriate to compare.

6 Lab Study Results

A high-level overview of participants’ interactions with COCOA and the baseline is depicted in Figure 8. In both conditions, output inspection occupied the most time out of the coded interactions, but was lengthier in the baseline than COCOA (54.1% vs. 33.6% of the approximately 25-minute session). A paired t-test found this difference to be significant ($t(15) = 3.95, p < 0.001$). Participants also engaged in co-execution in COCOA significantly more often than the closest equivalent interaction—output editing via prompting—in the baseline (15.2% vs. 3.1%, $t(15) = 4.29, p < 0.001$). They spent slightly more time on co-planning (21.5% vs. 17.4%), although this was not significant ($t(15) = 1.01, p = 0.33$). Our analysis indicates that, when provided with more affordances to interactively modify the agent’s outputs, users will take advantage of those opportunities for active engagement and shift away from passive output consumption. That is, users could **interleave co-planning and co-execution to steer the agent through tighter feedback loops**, rather than relying on slower cycles of trial-and-error prompting and output inspection.

We also computed performance metrics for both co-planning and co-execution from coded interactions and system logs. These included the number of step edits, reassignments, and output edits across all participants. Table 2 displays our results. Overall, 32.6% of steps were edited and 17.4% were reassigned. This shows that

¹⁰Since there were no user steps in the baseline, the closest equivalent interaction is attempting to edit the agent’s output via prompting.

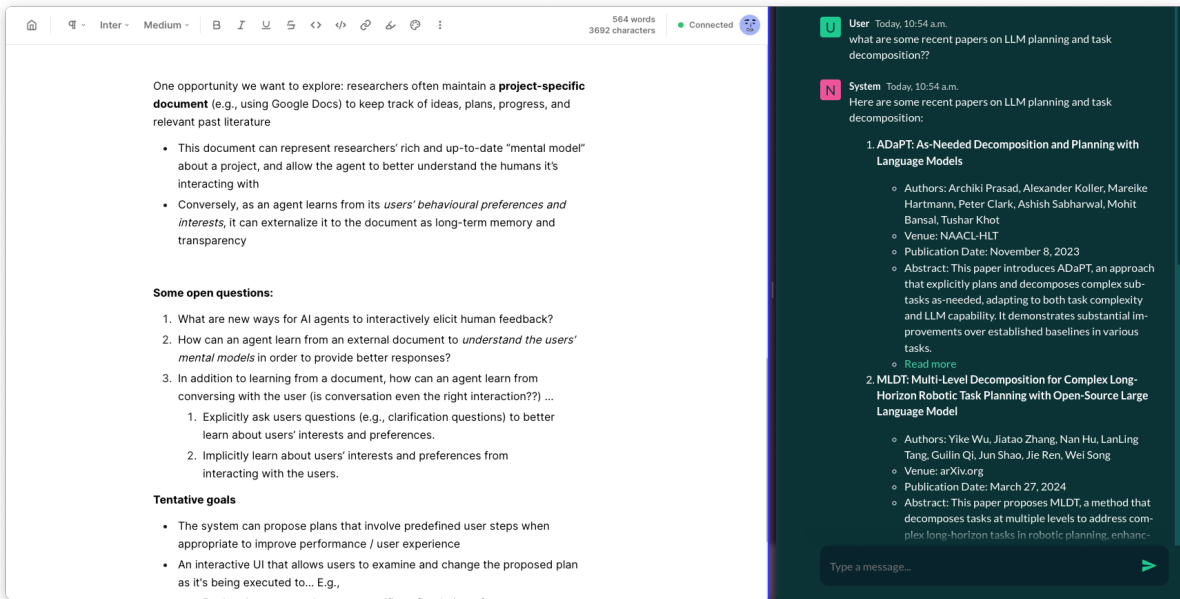


Figure 7: Participants’ typical setup for our baseline condition. On the left, the same document editor from COCOA but without an option to invoke the agent for co-planning and co-execution. On the right, the agent is situated within a chat interface.

most plan steps proposed by the agent (the content of the step as well as the assignment) were accepted as-is. During execution, the agent surfaced an average of 12.9 discrete output items (i.e., papers, authors, topics, text entities) per step to users. On average, 2.1 (16.3%) were deleted by participants, suggesting that the system had reasonable precision.

However, these metrics alone may not paint a holistic picture about COCOA’s use. To more deeply understand the context and rationale behind users’ actions, we analyzed videos from our user studies (Fig. 8) alongside qualitative transcripts from participants thinking aloud (Section 6.2).

Metric	Count (% if applicable)
Total steps completed across all plans	86
Step edits by user	28 (36.5%)
Step reassignments by user	15 (17.4%)
Step additions by user	12 (14.0%)
Step deletions by user	4 (4.7%)
Replanning triggered by system	14
Avg. output items (i.e., papers, authors, topics, text entities) per step	12.9
Avg. output items deleted per step	2.1 (16.3%)

Table 2: COCOA usage metrics from our lab study.

In the rest of this section, we present quantitative and qualitative analyses of participants’ data from our user study. We group these results by our research questions from the start of Section 5. We

denote the median rating of our baseline and COCOA as M_b and M_c , respectively, and the Wilcoxon test statistic as W .

6.1 Steerability, Ease of Use, and Utility (L1)

Introducing novel interactive systems with additional affordances can often lead to higher effort when using the systems as a trade-off for better utility [36, 96]. However, based on participants’ post-task ratings on a 5-point Likert scale, there was no significant difference in perceived *ease of use* between the baseline and COCOA with 16 participants ($M_b = 4.5$, $M_c = 4$, $p = 1.000$, $W = 7.50$). At the same time, we observed a significant difference that COCOA provided better *steerability*. In response to the question “*I could easily steer the system towards doing something helpful*,” participants rated COCOA higher than the baseline to a significant degree ($M_b = 3$, $M_c = 4$, $p = 0.005$, $W = 0$),¹¹ with $W = 0$ indicating that all participants rated COCOA’s steerability as greater than or equal to that of our baseline. In sum, these results suggest that **COCOA affords interleaving co-planning and co-execution via rich interactive affordances without sacrificing ease of use**, when compared to the simple chat interface used in the baseline condition (Figure 9).

We also tested for differences in *general utility*, which was a composite metric consisting of three measures broadly related to the usefulness and insightfulness of system outputs, but we found no significant differences ($M_b = M_c = 3.5$, $p = 0.7$, $W = 40$). This was unsurprising—the systems’ utility to the researcher may depend on a range of system-agnostic factors beyond our control (e.g., contextual interpretations of outputs by researchers, the extent to which researchers thought about the task before the study, etc.)

¹¹Here, we set the significance threshold to $\alpha = 0.05/3 = 0.015$ using Bonferroni correction. Our result is significant post-correction: $p = 0.005 < 0.015$.

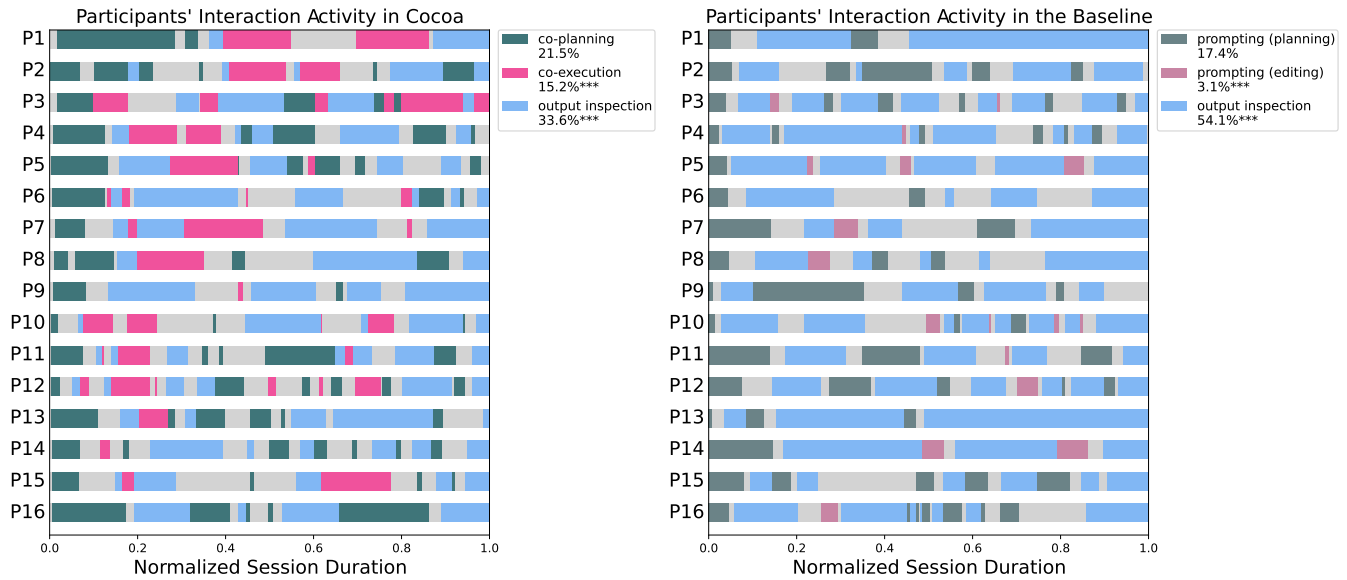


Figure 8: Participants' interaction activity in Cocoa and the baseline within the study sessions. * indicates a statistically significant difference ($p < 0.001$ via a paired t-test) with the closest equivalent interaction in the opposing condition. Gray areas represent spans of time where the participant was not engaging in any of our coded interactions because they were answering a question posed by the study facilitator and/or waiting on the system. Definitions for all coded interactions can be found in Section 5.4.**




Figure 9: Participants' Likert scale ratings of steerability, ease of use, and general utility across our baseline and Cocoa. ** indicates $p < 0.01$ via a Wilcoxon signed-rank test.

However, improved steerability may have allowed users to steer the agent away from providing downright unhelpful outputs in COCOA, as indicated by fewer lowly-rated outliers (see Figure 9).

6.2 Interleaving Co-Planning and Co-Execution Created an Effective Steering Loop (L2)

6.2.1 Outputs from co-execution drove re-planning. Qualitative results based on think alouds and post-interviews suggest that COCOA was valuable not only because it enabled human-agent co-planning or co-execution, but because *interleaving the two meant*

that co-execution outputs could be used to inform collaborative re-planning. Rather than treating the plan as a static artifact, participants used co-execution to find interesting and potentially productive directions to steer the agent towards. For example, P7 and P15 both used outputs from agent steps as a signal to “*organize thoughts and iterate on ideas*” (P7). P15 specifically iterated on an intermediate “*paper search*” step by re-executing it to cover papers from venues that they preferred. This led to subsequent plan updates that ultimately resulted in a more relevant final output. Similarly, P6 observed that “*backtracking*” to a specific step and editing it was more helpful than digging through a chat conversation: “*With chat, when it does something wrong there’s not really an easy way to fix it because there’s no concrete steps that it’s following, whereas [with*

COCOA], I can go back and be like, let me have it do something else here.”

Besides modifying intermediate steps, participants also added steps at the end of the plan to assign follow-up investigations, or deleted steps to sharpen the plan’s focus. For example, P12 added a step at the end of their plan to have the agent suggest some potential next steps to pursue after seeing its paper summaries, while P5 removed a less relevant intermediate step on searching for papers related to autonomous vehicles to avoid distractions from this tangential topic. In general, iterative output evaluation and re-planning allowed participants to better “fine-tune” the agent’s outputs and also their own research process—P1 likened their experience to using “[plan steps] as building blocks” that they could mix and match to create a custom workflow.

6.2.2 Fluid movement between co-planning and co-execution enabled granular steering. The interleaving of co-planning and co-execution allowed participants to move fluidly between the two and exert control exactly when necessary. By opting for stepwise execution of the plan, participants could pause the agent to “catch up” with agent activity and ensure outputs are up to standard before continuing. Indeed, P4 and P11 both *verified* and *manually curated* the list of papers returned by the agent from a paper search step before proceeding. P11 specifically cites resource efficiency as a practical reason: “If I update [outputs of] the first or second step [after executing all steps], it would need to rerun the following steps. I just want to reduce some API usage.” P10 similarly mentioned that they would “waste energy” running subsequent steps if they later detected errors in initial steps. Overall, many participants chose stepwise execution because it *allowed more room to interleave co-planning*; this in turn offered more granular control over the human-agent collaborative work trajectory and prevented error propagation.

There were, however, still reasons for preferring continuous execution. P8 and P9 thought continuous execution was more efficient because it parallelized human and AI efforts: “Running all [steps] is a little bit of parallelization in my head where I can get the next thing rolling, and if I end up not having to edit it, that means we’re already moving ahead” (P8). In this case, P8 does not feel a need to be “caught up” with the agent’s work because they know they can still steer the agent at their own pace by modifying the agent’s execution and triggering re-planning at any point.

Interleaving co-planning and co-execution via direct manipulation in particular also contributed to granular steering. Many participants, including P6, liked the precision with which they could edit the output of a particular step compared to a chat interface: “With a conversation agent, even if I’m really good at prompting it, I have to redo the prompt and keep changing that original prompt.” (P6). This feature also allowed P2 to quickly re-plan and repurpose the output when the agent produced unexpected (albeit helpful) paper search queries: “I was looking for more application intervention based stuff, but this is great because now I can actually use these keywords to broaden my horizon of how I was thinking about these solutions, which is I think interesting. I’m just trying to remove a couple [for future steps].” Finally, participants found that interleaved co-planning and co-execution scoped tasks in a way that naturally afforded more precise control, especially when paired with direct

manipulation, as P15 shared: “[COCOA] was easier to control because [the outputs] are so specific. My control is scoped down to these very small tasks. If I want to add more papers, that’s easy to do.”

7 Deployment Study and Results

In addition to interleaving *co-execution and co-planning*, we were also interested in studying *the delegation of human and AI work*. Thus, we tracked which steps participants assigned to the agent vs. themselves.

In our formative study (Section 3.2.3), participants preferred to assign tasks related to information retrieval and brainstorming to the agent while keeping higher-level reasoning and synthesis tasks for themselves. In our lab study, however, most participants toggled all plan steps—even higher-level ones assigned to them by default—to the agent. Some avoided user steps because they were curious about “*how the bot does it*” (P1). Understandably, testing the agent’s capabilities and limitations is a natural tendency when learning to use a new system. Many speculated that they would leverage user steps if they had more time to interact with the system outside of the lab environment. P8 explicitly said that “*I assign everything to [the agent] because I don’t think we have that much time for me to refine my [own steps].*”

To obtain a more holistic understanding of how COCOA can facilitate flexible delegation of human-AI work and more generally assist researchers day-to-day, we conducted a 7-day field deployment study with 7 participants.

7.1 Participants and Procedure

We invited all participants from our lab study to participate in this field deployment. 7 participants (P3, P5, P10, P11, P12, P15, P16) agreed to participate.

Participants used COCOA to work on a real-world, in-progress research project of their choice over a 7-day period. Before the study, the study facilitator sent each participant a user manual¹² documenting COCOA’s functionalities. The manual also contained some general guidelines for the week (see Appendix I), including aiming to spend at least 90 minutes with the system throughout the duration of the study. Each participant could log onto their personal workspace in COCOA with a set of credentials provided by the study facilitator. We used system logs to verify that there was at least semi-regular activity for each participant throughout the study.

At the end of the 7-day period, participants completed a 30-minute semi-structured exit interview where they were asked about their general experience with the system as well as specific co-planning and co-execution interactions the study facilitator identified from within their workspace. All interviews were recorded, transcribed, and analyzed with the same qualitative analysis procedure described in Section 5.4. Participants received a \$100 USD honorarium upon completing the exit interview.

7.2 Field Deployment Study Results

7.2.1 COCOA was most useful for literature synthesis and early-stage project planning. All participants agreed that COCOA was especially helpful for **literature discovery and synthesis**.

¹²See Supplementary Materials.

This illustrates the value of interleaving human-AI efforts across planning and execution in complex information foraging and sense-making activities. P15 found the structured plan for conducting literature search to be “really helpful for refining my thought process because it made me think more carefully about what I’m actually trying to look for.” P10 agreed, and also felt that the combination of plan steps involving search query generation and paper search left them with “more faith that I was getting a good coverage of the field.” P11 also appreciated COCOA’s speed at literature search and synthesis and found it **conducive to early-stage explorations**: “if I were to come up with a summary of 20 papers, it takes probably a day, but [in COCOA] it’s around a minute which is really useful to formulate early-stage research ideas.” Beyond literature search, the structure of the plan was also useful to “break down tasks and provide a sense of progress” (P16) and “identifying important actionable steps I could take with a research problem” (P12). Overall, participants found COCOA valuable in their work, with P3 in particular sharing that “I have an entire tab in my Google doc where we have stuff from COCOA.”

7.2.2 Step assignment strategy changed: more delegation of work with longitudinal use. Compared to the lab study, where steps were mostly assigned to the agent, participants in the field study assigned more plan steps to themselves. This was likely due to a combination of increased familiarity with the agent’s capabilities and system features, as well as more opportunity to reason about which steps they wanted to keep for themselves and why. These user steps often involved high-level strategic decision-making, writing and developing arguments, running experiments and evaluations, deep reading of literature, and articulating core research ideas and novelty.

The ability to assign steps to the agent and themselves helped participants **create a back-and-forth workflow with the agent**. For example, P15 first figured out “broad buckets of literature or themes or topics that I need to talk about” before passing those to the agent to “help me brainstorm ways to connect those bodies of literature.” P3, who had already conducted some interviews for their project, first had the agent look up key literature before assigning themselves a step to “pass in more of the dynamics from my interview notes” and then had the agent propose interventions based on the literature and notes. For some participants, user steps came later in the plan because of their high-level nature. P11 often assigned the last plan step to themselves because “it was about the main innovation in the project idea [...] no one has [done it] before so I don’t trust the agent.” Participants also indicated that their step assignment also **depended on the stage of research they were in**. P3 admitted that “if I was just doing initial discovery and study planning, I would have gone more agent [steps].” These findings suggest that participants benefitted from the ability to flexibly configure agency in their human-AI workflows.

7.2.3 Interleaving co-planning and co-execution supported evolving research trajectories. Just like in the lab study, participants valued interleaving of co-planning and co-execution, but the value for supporting **flexible and iterative workflow** became more apparent with longitudinal use. As participants saw the results of execution, their understanding of the problem space shifted, which led them to want to **anchor certain steps and outputs as**

“checkpoints,” and use those as a springboard for exploring new territory. For example, P12 shared that “I don’t need to run the [earlier] retrieval step again. I know that those papers are relevant” but instead could focus on “only improving the second step [...] to identify something different from the literature.” Similarly, P16 felt like they could easily “debug” suboptimal parts of the plan: “by forcing these checkpoints of brainstorming search queries and assembling lists of papers, it helps me get a sense of what parts of the process are good and which parts need work.” Outputs from co-execution also inspired new approaches for participants to **“iteratively build upon a plan or reshape a plan”** (P5). Participants even mentioned that they wanted similar interleaving strategies integrated into other LLM-powered systems; for P11, it was “one of the improvements I want to see in LLM [reasoning] models” so that users can force an early stop to the reasoning or steer it in more productive directions. Overall, our deployment of COCOA suggests that interleaving co-planning and co-execution helps support the non-linear nature of scientific inquiry.

7.2.4 Additional desiderata. Participants identified several desiderata for COCOA after spending more time with it. Some participants wanted to clearly see **what context COCOA used** to generate plans and execute plan steps. For example, P10 was unsure whether the system was “using the context from my other plans to generate the [other] plan.” The system indeed was, but did not visualize relevant context in the document. P3, P15, and P16 all wanted more transparency into the shared context pool (Section 4.2) used across all steps and the specific context drawn upon by each step. Several participants desired **tool interoperability**, such that they could swap out COCOA’s default tools for their own ones (e.g., an LLM fine-tuned on paper summaries instead of COCOA’s paper summarizer). This way, COCOA would serve as a “hub” for their custom tools. Finally, P5 took inspiration from Notion’s ability to nest pages to suggest nested plans, where plans can be organized and expanded in a hierarchical manner. These suggestions all provide valuable directions for future work.

8 Discussion

8.1 Design Considerations for Co-Agency in Human-Agent Collaboration

Researchers have called for reframing the design of human-AI systems as not merely the specification of affordances, but construction of open design spaces that support human-AI co-agency [89]. That is, rather than providing a fixed configuration of human and AI agency across all interactions, users can flexibly vary the agency they wish assert and have the AI modulate its agency in response. COCOA offers an early instantiation of this vision. We now reflect upon several key design decisions and discuss how future systems may build upon them or explore alternatives when designing for co-agency in human-agent systems.

We expect that *fluid re-planning* will continue to be important and challenging to design for in future systems. Plans are bound to shift as the user reviews agent outputs and becomes inspired with new ideas. For example, researchers in both our lab and deployment studies dynamically re-planned with the agent upon gaining more clarity on *what they actually wanted to explore* after co-executing




Figure 10: Examples from our field deployment where user and agent steps were interleaved in participants' plans. The plans steps were rewritten to redact project-specific details.

earlier steps (see Sections 6.2 and 7.2.3). In COCOA, re-planning replaces outdated steps with new ones. However, participants sometimes expressed a desire to preserve the old steps because their outputs can still be useful to reference in the future. We thus see *plan branching* as a promising design pattern for exploring counterfactual plan paths. For example, systems can allow users to swap certain steps in and out of the plan, or set checkpoints at specific steps to anchor exploration. By doing so, useful work can be better retained while supporting flexible plan redirection.

Relatedly, the *length of the plan should be calibrated depending on the task*. Some tasks, such as open-ended brainstorming, benefit from shorter, partial plans where next steps can be flexibly determined based on prior execution. Indeed, a few participants expressed a preference for a step-by-step planning approach over COCOA's default plans when conducting brainstorming, where they did not know what next step would be until they saw the results of the previous step. On the other hand, for tasks with more structured and predictable workflows, such as literature review, participants appreciated interacting with a full plan. We calibrated the lengths of plans generated by COCOA based on participants' plans in our formative study. We recommend designers of human-agent systems to consider what the appropriate plan length is for tasks they hope to support, and calibrate accordingly.

Finally, *user expertise and upskilling potential should be considered* when designing for an appropriate level of agent control. As surfaced in our formative study (Section 3.2.3), there are two primary reasons why users prefer keeping a task for themselves over delegating it to the agent. *First*, the user can perform the task more effectively than the agent. This is often the case where the user holds more expertise (e.g., brainstorming research questions in a user's primary research area). *Second*, the task is intrinsically valuable for the user to do themselves, even if an agent is capable of automating it (e.g., reading seminal papers to develop research taste). If a human-agent system is designed for expert users who want to regularly hone their skills and intuition, providing many levels of control throughout agent planning and execution (as in COCOA) may be appropriate. If the system is for navigating unfamiliar domains with agents and users are not looking to upskill, fewer controls may enhance user experience. Because interactions such as co-planning and co-execution can burden the user with additional cognitive load, determining the user's capacity and willingness to take on more cognitive load can inform the appropriate level of control to afford.

These design considerations were drawn from the specific definition of agency (from Bennett et al. [8]) and domain (scientific research) specific to our study. The design space is vast, and we are excited to witness the application, refinement, and even refutation of these considerations as future work navigates this space with alternative interpretations of agency in different domains.

8.2 The Agent Notebook Paradigm?

The design of COCOA's interface was partly inspired by the computational notebook paradigm. Just like how notebooks help break apart a script into smaller components for unit testing and quick iteration, our "agent notebook" design—featuring plan step "cells" and options for stepwise or continuous execution—help bring similar benefits to agentic workflows. We saw in both our lab study (Section 5) and field deployment (Section 7) that participants valued the ability to execute steps one at a time and interactively refine outputs before moving onto subsequent steps. We now discuss agent notebooks as a *potentially generalizable design pattern* and consider its applications to agentic systems beyond COCOA.

Consider "deep research" systems that assemble a detailed report on a particular topic upon searching for and synthesizing information from around the web [26, 76, 83]. These systems may clarify user requests in a chat conversation [76], or allow the user to prompt the agent to regenerate the plan [26] before starting the research process. However, there is no room for user involvement thereafter. Agent notebooks can be particularly valuable here because they afford iteration and control over *specific parts of planning and execution*. For example, the user may not be satisfied with the sources the agent is using to inform its answers and wants to steer the agent to use information from a narrow set of websites for a particular plan step. This kind of interaction is expensive without agent notebooks because it requires re-running the entire research process. Agent notebooks allow the user to focus on improving outcomes of a particular step by editing that step and rerunning it or directly editing the agent's outputs at that step, which can save significant time and cost.

While we envision agent notebooks to be a design pattern that is generalizable beyond COCOA, we note that the *specific affordances for co-planning and co-execution* may vary depending on the environment in which human-agent collaboration happens. COCOA's environment is a document editor for scientific research, and contains affordances designed around that environment (e.g., representing plan steps as editable bulleted lists, paper curation and search in the sidebar). Translating these affordances to coding, for

example, may need to draw upon different interaction metaphors. The plan can be represented as an executable spec file (e.g., a special Jupyter notebook with only markdown cell blocks that serve as step instructions), while co-execution can resemble live pair programming. We anticipate that making intentional design choices for these affordances by drawing from appropriate metaphors will be an important part of future work.

8.3 The Practical Significance of User Steps and Planning Ahead

In COCOA, users are presented with an interactive multi-step plan up front, which the agent adheres to and can modify while operating. Importantly, this plan also allows users to flexibly delegate a step to themselves or the agent. Both design decisions are a departure from many existing agentic AI systems, where the agent dynamically constructs a plan step-by-step as it generates outputs [110] and seeks user input only when it *automatically* determines it should do so [75, 81, 94]. We discuss two practical reasons—cost and oversight—for why such a departure may be desirable for practical agents in the real world.

AI agents can be prohibitively expensive to run [46, 107]. Because an agent’s actions often involve multiple LLM calls, costs can quickly accumulate, especially for more complex and long-running tasks. In response, academic projects have implemented cost-capping measures to alleviate this [107], and researchers have called for agent evaluations to be cost-controlled [46]. In user-facing interactive systems, cost also includes time spent waiting for agents to come back with results. These cost considerations change user behavior. When participants made the decision to assign most or all steps to the agent in our lab studies, they viewed the agent’s actions as rather inconsequential and low-cost—they could run the agent and see what happens, and edit the output post hoc (Section 6.2). However, while this was the case in the lab, it may not be true in some real-world contexts where users and/or developers need to pay monetary and time costs for each agent step. User steps become an important cost-saving measure. Users can choose to complete a step that is not particularly burdensome or is too difficult for current AI agents to complete effectively. Efficiently leveraging human effort and expertise where appropriate via user steps can be a key step towards developing cost-aware agents that are usable and affordable in practice.

The increased autonomy of AI agents also comes with increased risks, as AI harms become more difficult to anticipate, and accountability for AI actions becomes harder to trace [11, 12]. COCOA’s interface not only provides more transparency into agent actions, but also provide some assurance about and control over the agent’s execution trajectory. By generating a plan ahead of execution that the agent adheres to, harm anticipation and mitigation become more tractable. Step-by-step plan generation used by existing agent frameworks, such as ReAct [110], may also accomplish a similar goal by asking for user approval whenever a new step is generated. However, this can severely limit agent autonomy when desirable and may also disengage the user if their only form of supervision is repetitive rubber-stamping [46]. Additionally, advanced planning enables assessments of task risk pre-execution. Plan steps

determined to be of higher risk (e.g., requiring working with passwords) can be assigned to the user—even if the agent is capable of completing them—as a risk mitigation measure.

9 Limitations and Future Work

Based on our formative study, we built COCOA as a document editor because documents are natural sites of planning for researchers and present an ideal environment for agent interaction (Section 3.2.1). However, the linear nature of a document also has its limitations and does not support “forking” plans and iterating on versions of a plan in parallel. These types of interactions are better supported by node-based canvases, which have been gaining popularity as a means of sensemaking and creative exploration with LLMs (e.g., [3, 84, 99]). Future work can explore new interaction techniques for interleaved co-planning and co-execution. The default assignment of user and agent steps in COCOA is also a result of researchers’ preferences from our formative study. Past work explored automated methods for a model to decide when to defer a task to a human expert versus acting on its own [70]. Future work may investigate applications of these methods in human-agent systems and/or leverage inference-time scaling [71, 74] to improve plan quality.

Our baseline mimicked the dominant interaction paradigm in AI-powered research tools: chatbots. While this allowed us to elicit participant feedback on the usability of COCOA compared to the status quo, we could have designed additional baselines by ablating parts of our design. For example, our baseline could have been an interactive planning interface that did not support stepwise execution. In future work, we can expand our participant pool and run more studies with ablations of COCOA to pinpoint which particular affordances participants valued most. Further, the underlying agent used in COCOA also had some technical limitations. Since it was not multimodal, it was not capable of taking visual content (e.g., figures or slides) as input. The agent was also not capable of executing code—a capability which a couple of participants inquired about during the study. Future work can expand COCOA’s utility with multimodal agents and/or agents that can execute code.

Another limitation is the demographics of our participants. Our participants were researchers in CS and CS-adjacent areas. Research culture and incentives specific to CS may bias our results and limit our imagination of how COCOA can be used. P10, who used to work in wet labs in their undergraduate research, shared that a promising use case for COCOA would be helping wet lab researchers walk through the steps necessary to prepare for lab experiments. Future work may investigate novel applications of systems like COCOA in domains beyond CS.

10 Conclusion

In this paper, we presented COCOA, an interactive system for scientific researchers to flexibly collaborate with an AI agent in a document editing environment to tackle open questions and tasks within their research projects. Motivated by a formative study with 9 researchers, COCOA introduced a new approach for human-agent collaboration—*interleaved co-planning and co-execution*—by which a human and agent can iteratively create plans of action and execute the plan in a mixed-initiative manner that centers human agency.

Our lab ($n = 16$) and field deployment ($n = 7$) studies showed that participants benefited from interleaved co-planning and co-execution in the lab and in the field. Compared to a strong baseline with a more familiar chat interface, COCOA enhanced agent steerability without sacrificing ease of use. In the lab study, participants' main mode of co-execution was refining agent outputs, but when COCOA was integrated into their research workflows in the field over a 7-day period, participants took advantage of user steps to guide the agent with their expertise and exert more of their agency. Overall, our work demonstrates the promise of our approach for flexibly mediating human and AI agency. Importantly, our work offers many practical takeaways for the design of human-agent systems, such as leveraging user steps and planning ahead for agent cost management and oversight.

AI agents have exciting potential to transform our digital experiences and advance long-standing visions held by HCI and AI communities. However, these transformations and advancements also demand renewed attention to strategies for ensuring effective and flexible human-agent collaboration. Our work advances efforts in this emerging frontier.

Acknowledgments

We warmly thank all our participants from our formative, pilot, lab, and deployment studies. We thank our anonymous reviewers for feedback on our manuscript. Finally, we thank Raymond Fok and Shannon Shen for helpful discussions.

References

- [1] Tyler Angert, Miroslav Suzara, Jenny Han, Christopher Pondoc, and Hariharan Subramonyam. 2023. Spellburst: A Node-based Interface for Exploratory Creative Coding with Natural Language Prompts. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (2023). <https://api.semanticscholar.org/CorpusID:260704716>
- [2] Anthropic. 2024. Developing a computer use model. <https://www.anthropic.com/news/developing-computer-use>.
- [3] Ian Arawjo, Chelse Swoopes, Priyan Vaithilingam, Martin Wattenberg, and Elena L. Glassman. 2023. ChainForge: A Visual Toolkit for Prompt Engineering and LLM Hypothesis Testing. *ArXiv abs/2309.09128* (2023). <https://api.semanticscholar.org/CorpusID:262044762>
- [4] AutoGPT. 2024. Empower your digital tasks with AutoGPT. <https://agpt.co/>.
- [5] Amid Ayobi, Jacob Hughes, Christopher Duckworth, Jakob J Dylag, Sam James, Paul Marshall, Matthew Guy, Anitha Kumaran, Adriane Chapman, Michael J. Boniface, and Aisling Ann O'Kane. 2023. Computational Notebooks as Co-Design Tools: Engaging Young Adults Living with Diabetes, Family Carers, and Clinicians with Machine Learning Models. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:258218057>
- [6] Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. 2024. ResearchAgent: Iterative Research Idea Generation over Scientific Literature with Large Language Models. *ArXiv abs/2404.07738* (2024). <https://api.semanticscholar.org/CorpusID:269042844>
- [7] Gagan Bansal, Jennifer Wortman Vaughan, Saleema Amershi, Eric Horvitz, Adam Fournay, Hussein Mozannar, Victor Dibia, and Daniel S Weld. 2024. Challenges in Human-Agent Communication. (2024). <https://api.semanticscholar.org/CorpusID:270870360>
- [8] Dan Bennett, Oussama Metatla, Anne Roudaut, and Elisa D. Mekler. 2023. How does HCI Understand Human Agency and Autonomy? *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:256389761>
- [9] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11 (2019), 589 – 597. <https://api.semanticscholar.org/CorpusID:197748828>
- [10] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.
- [11] Alan Chan, Carson Ezell, Max Kaufmann, Kevin Wei, Lewis Hammond, Herbie Bradley, Emma Bluemke, Nitarshan Rajkumar, David Krueger, Noam Kolt, et al. 2024. Visibility into AI Agents. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 958–973.
- [12] Alan Chan, Rebecca Salganik, Alva Markelius, Chris Pang, Nitarshan Rajkumar, Dmitrii Krashennnikov, Lauro Langosco, Zhonghao He, Yawen Duan, Micah Carroll, et al. 2023. Harms from increasingly agentic algorithmic systems. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 651–666.
- [13] Joel Chan, Joseph Chee Chang, Tom Hope, Dafna Shahaf, and Aniket Kittur. 2018. Solvent: A mixed initiative system for finding analogies between research papers. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–21.
- [14] Joseph Chee Chang, Amy X. Zhang, Jonathan Bragg, Andrew Head, Kyle Lo, Doug Downey, and Daniel S. Weld. 2023. CiteSee: Augmenting Citations in Scientific Papers with Persistent and Personalized Historical Context. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:256868353>
- [15] Souti Chattopadhyay, I. V. R. K. V. Prasad, Austin Z. Henley, Anita Sarma, and Titus Barik. 2020. What's Wrong with Computational Notebooks? Pain Points, Needs, and Design Opportunities. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020). <https://api.semanticscholar.org/CorpusID:210927488>
- [16] Jiangjie Chen, Siyu Yuan, Rong Ye, Bodhisattwa Prasad Majumder, and Kyle Richardson. 2023. Put your money where your mouth is: Evaluating strategic planning and execution of llm agents in an auction arena. *arXiv preprint arXiv:2310.05746* (2023).
- [17] Inyoung Cheong, King Xia, K. J. Kevin Feng, Quan Ze Chen, and Amy X. Zhang. 2024. (A)I Am Not a Lawyer, But... Engaging Legal Experts towards Responsible LLM Policies for Legal Advice. In *Conference on Fairness, Accountability and Transparency*. <https://api.semanticscholar.org/CorpusID:267413187>
- [18] Frederick Choi, Sajjadur Rahman, Han Jun Kim, and Daz Zhang. 2023. Towards Transparent, Reusable, and Customizable Data Science in Computational Notebooks. *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:257687372>
- [19] Douglas L. Dean, Jillian M. Hender, Thomas Lee Rodgers, and Eric L. Santanen. 2006. Identifying Quality, Novel, and Creative Ideas: Constructs and Scales for Idea Evaluation. *J. Assoc. Inf. Syst.* 7 (2006), 30. <https://api.semanticscholar.org/CorpusID:15910404>
- [20] Elicit. 2024. Elicit: The AI Research Assistant. <https://elicit.com/>.
- [21] KJ Feng, David W McDonald, and Amy X Zhang. 2025. Levels of Autonomy for AI Agents. *arXiv preprint arXiv:2506.12469* (2025).
- [22] K. J. Kevin Feng, Quan Ze Chen, Inyoung Cheong, King Xia, and Amy X. Zhang. 2023. Case Repositories: Towards Case-Based Reasoning for AI Alignment. *ArXiv abs/2311.10934* (2023). <https://api.semanticscholar.org/CorpusID:265295304>
- [23] Jennifer Fereday and Eimear Muir-Cochrane. 2006. Demonstrating rigor using thematic analysis: A hybrid approach of inductive and deductive coding and theme development. *International journal of qualitative methods* 5, 1 (2006), 80–92.
- [24] Raymond Fok, Joseph Chee Chang, Tal August, Amy X. Zhang, and Daniel S. Weld. 2023. Qlarify: Recursively Expandable Abstracts for Directed Information Retrieval over Scientific Papers. <https://api.semanticscholar.org/CorpusID:263835343>
- [25] Raymond Fok, Hita Kambhamettu, Luca Soldaini, Jonathan Bragg, Kyle Lo, Andrew Head, Marti A. Hearst, and Daniel S. Weld. 2022. Scim: Intelligent Skimming Support for Scientific Papers. *Proceedings of the 28th International Conference on Intelligent User Interfaces* (2022). <https://api.semanticscholar.org/CorpusID:254591867>
- [26] Google. 2025. Gemini Deep Research. <https://gemini.google/overview/deep-research/>.
- [27] Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutarō Tanno, et al. 2025. Towards an AI co-scientist. *arXiv preprint arXiv:2502.18864* (2025).
- [28] Madeleine Grunde-McLaughlin, Michelle S. Lam, Ranjay Krishna, Daniel S. Weld, and Jeffrey Heer. 2023. Designing LLM Chains by Adapting Techniques from Crowdsourcing Workflows. *ArXiv abs/2312.11681* (2023). <https://api.semanticscholar.org/CorpusID:266362444>
- [29] Joel Grus. 2018. I don't like notebooks. https://docs.google.com/presentation/d/1n2RlMdmv1p25Xy5thJUhkKgvjV-dkAIsUXP-AL4fll/edit?slide=id.g362da58057_0_1.
- [30] Xuemei Gu and Mario Krenn. 2024. Generation and human-expert evaluation of interesting research ideas using knowledge graphs and large language models. <https://api.semanticscholar.org/CorpusID:270062620>
- [31] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025).
- [32] Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. Reasoning with Language Model is Planning with World Model. *ArXiv abs/2305.14992* (2023). <https://api.semanticscholar.org/CorpusID:>

- 258865812
- [33] Gaole He, Gianluca Demartini, and Ujwal Gadiraju. 2025. Plan-then-execute: An empirical study of user trust and team performance when using llm agents as a daily assistant. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [34] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *International Conference on Human Factors in Computing Systems*. <https://api.semanticscholar.org/CorpusID:8943607>
- [35] Faria Huq, Zora Zhiruo Wang, Frank F. Xu, Tianyue Ou, Shuyan Zhou, Jeffrey P. Bigham, and Graham Neubig. 2025. CowPilot: A Framework for Autonomous and Human-Agent Collaborative Web Navigation. *ArXiv abs/2501.16609* (2025). <https://api.semanticscholar.org/CorpusID:275931872>
- [36] Edwin L. Hutchins, James Hollan, and Donald A. Norman. 1985. Direct Manipulation Interfaces. *Hum. Comput. Interact.* 1 (1985), 311–338. <https://api.semanticscholar.org/CorpusID:16355120>
- [37] Chip Huyen. 2025. Agents. <https://huyenchip.com/2025/01/07/agents.html>.
- [38] Peter Jansen, Marc-Alexandre Côté, Tushar Khot, Erin Bransom, Bhavana Dalvi Mishra, Bodhisattwa Prasad Majumder, Oyvind Tafjord, and Peter Clark. 2024. DISCOVERYWORLD: A Virtual Environment for Developing and Evaluating Automated Scientific Discovery Agents. *arXiv preprint arXiv:2406.06769* (2024).
- [39] Peter Jansen, Oyvind Tafjord, Marissa Radensky, Pao Siangliulue, Tom Hope, Bhavana Dalvi, Bodhisattwa Prasad Majumder, Daniel S. Weld, and Peter Clark. 2025. CodeScientist: End-to-End Semi-Automated Scientific Discovery with Code-based Experimentation. <https://api.semanticscholar.org/CorpusID:277451644>
- [40] Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2023. SWE-bench: Can Language Models Resolve Real-World GitHub Issues? *ArXiv abs/2310.06770* (2023). <https://api.semanticscholar.org/CorpusID:263829697>
- [41] Marina Jirotko, Charlotte P Lee, and Gary M Olson. 2013. Supporting scientific collaboration: Methods, tools and concepts. *Computer Supported Cooperative Work (CSCW)* 22 (2013), 667–715.
- [42] Project Jupyter. 2015. Jupyter Notebook UX Survey. <https://github.com/jupyter/surveys/tree/master/surveys/2015-12-notebook-ux>.
- [43] Hyeonsu B Kang, Joseph Chee Chang, Yongsung Kim, and Aniket Kittur. 2022. Threddy: An Interactive System for Personalized Thread-based Exploration and Organization of Scientific Literature. *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (2022). <https://api.semanticscholar.org/CorpusID:251402552>
- [44] Hyeonsu B Kang, Rafal Kocielnik, Andrew Head, Jiangjiang Yang, Matt Latzke, Aniket Kittur, Daniel S. Weld, Doug Downey, and Jonathan Bragg. 2022. From Who You Know to What You Read: Augmenting Scientific Recommendations with Implicit Social Networks. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022). <https://api.semanticscholar.org/CorpusID:248299830>
- [45] Hyeonsu B Kang, Sherry Wu, Joseph Chee Chang, and Aniket Kittur. 2023. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (2023). <https://api.semanticscholar.org/CorpusID:260899915>
- [46] Sayash Kapoor, Benedikt Stroebel, Zachary S. Siegel, Nitya Nadgir, and Arvind Narayanan. 2024. AI Agents That Matter. *ArXiv abs/2407.01502* (2024). <https://api.semanticscholar.org/CorpusID:270870360>
- [47] Majeed Kazemitabaar, Jack Williams, Ian Drosos, Tovi Grossman, Austin Z. Henley, Carina Negreanu, and Advait Sarkar. 2024. Improving Steering and Verification in AI-Assisted Data Analysis with Interactive Task Decomposition. <https://api.semanticscholar.org/CorpusID:270923956>
- [48] Mary Beth Kery, Marissa Radensky, Mahima Arya, Bonnie E. John, and Brad A. Myers. 2018. The Story in the Notebook: Exploratory Data Science using a Literate Programming Tool. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018). <https://api.semanticscholar.org/CorpusID:5060661>
- [49] Mary Beth Kery, Donghao Ren, Fred Hohman, Dominik Moritz, Kanit Wongsuphasawat, and Kayur Patel. 2020. mage: Fluid Moves Between Code and Graphical Work in Computational Notebooks. *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (2020). <https://api.semanticscholar.org/CorpusID:221836345>
- [50] Joongwon Kim, Bhargavi Paranjape, Tushar Khot, and Hanna Hajishirzi. 2024. Husky: A Unified, Open-Source Language Agent for Multi-Step Reasoning. <https://api.semanticscholar.org/CorpusID:270370824>
- [51] Jeongyeon Kim, Sangho Suh, Lydia B Chilton, and Haijun Xia. 2023. Metaphorian: Leveraging Large Language Models to Support Extended Metaphor Creation for Science Writing. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 115–135.
- [52] Rodney Michael Kinney, Chloe Anastasiades, Russell Authur, Iz Beltagy, Jonathan Bragg, Alexandra Buraczynski, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Arman Cohan, Miles Crawford, Doug Downey, Jason Dunkelberger, Oren Etzioni, Rob Evans, Sergey Feldman, Joseph Forney, David W. Graham, F.Q. Hu, Regan Huff, Daniel King, Sebastian Kohlmeier, Bailey Kuehl, Michael Langan, Daniel Lin, Haokun Liu, Kyle Lo, Jaron Lochner, Kelsey MacMillan, Tyler C. Murray, Christopher Newell, Smita R Rao, Shaurya Rohatgi, Paul Sayre, Zejiang Shen, Amanpreet Singh, Luca Soldaini, Shivshankar Subramanian, A. Tanaka, Alex D Wade, Linda M. Wagner, Lucy Lu Wang, Christopher Wilhelm, Caroline Wu, Jiangjiang Yang, Angele Zamarron, Madeleine van Zuylen, and Daniel S. Weld. 2023. The Semantic Scholar Open Data Platform. *ArXiv abs/2301.10140* (2023). <https://api.semanticscholar.org/CorpusID:256194545>
- [53] Donald Ervin Knuth. 1984. Literate Programming. In *Computer/law journal*. <https://api.semanticscholar.org/CorpusID:1200693>
- [54] Sam Lau, Ian Drosos, Julia M. Markel, and Philip J. Guo. 2020. The Design Space of Computational Notebooks: An Analysis of 60 Systems in Academia and Industry. In *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, 1–11. <https://doi.org/10.1109/VL/HCC50065.2020.9127201>
- [55] Lane Lawley and Christopher J. MacLellan. 2024. VAL: Interactive Task Learning with GPT Dialog Parsing. *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2024). <https://api.semanticscholar.org/CorpusID:263609076>
- [56] Charlotte P Lee. 2007. Boundary negotiating artifacts: Unbinding the routine of boundary objects and embracing chaos in collaborative work. *Computer Supported Cooperative Work (CSCW)* 16 (2007), 307–339.
- [57] Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambsganss, David Zhou, Emad A. Alghamdi, Tal August, Avinash Bhat, Madiha Zahrah Choksi, Senjuti Dutta, Jin L.C. Guo, Md. Naimul Hoque, Yewon Kim, Seyed Parsa Neshaei, Agnia Sergeyuk, Antonette Shibani, Disha Shrivastava, Lila Shroff, Jessi Stark, S. Sterman, Sitong Wang, Antoine Bosselut, Daniel Buschek, Joseph Chee Chang, Sheral Chen, Max Kreminski, Joonuk Park, Roy Pea, Eugenia H. Rho, Shannon Zejiang Shen, and Pao Siangliulue. 2024. A Design Space for Intelligent and Interactive Writing Assistants. *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2024). <https://api.semanticscholar.org/CorpusID:268553985>
- [58] Yoonjoo Lee, Hyeonsu B Kang, Matt Latzke, Juho Kim, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2024. PaperWeaver: Enriching Topical Paper Alerts by Contextualizing Recommended Papers with User-collected Papers. In *International Conference on Human Factors in Computing Systems*. <https://api.semanticscholar.org/CorpusID:268248445>
- [59] Haotian Li, Lu Ying, Haidong Zhang, Yingcai Wu, Huamin Qu, and Yun Wang. 2023. Notable: On-the-fly assistant for data storytelling in computational notebooks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [60] Zhehui Liao, Maria Antoniak, Inyoung Cheong, Evie Yu-Yen Cheng, Ai-Heng Lee, Kyle Lo, Joseph Chee Chang, and Amy X Zhang. 2024. LLMs as Research Tools: A Large Scale Survey of Researchers' Usage and Perceptions. *arXiv preprint arXiv:2411.05025* (2024).
- [61] Henry Lieberman. 1997. Autonomous interface agents. *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems* (1997). <https://api.semanticscholar.org/CorpusID:6576547>
- [62] Yanna Lin, Haotian Li, Leni Yang, Aoyu Wu, and Huamin Qu. 2023. InkSight: Leveraging Sketch Interaction for Documenting Chart Findings in Computational Notebooks. *IEEE Transactions on Visualization and Computer Graphics* 30 (2023), 944–954. <https://api.semanticscholar.org/CorpusID:259936935>
- [63] Yue Liu, Sin Kit Lo, Qinghua Lu, Liming Zhu, Dehai Zhao, Xiwei Xu, Stefan Harrer, and Jon Whittle. 2024. Agent Design Pattern Catalogue: A Collection of Architectural Patterns for Foundation Model based Agents. *arXiv preprint arXiv:2405.10467* (2024).
- [64] Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292* (2024).
- [65] Xiao Ma, Swaroop Mishra, Ariel Liu, Sophie Ying Su, Jilin Chen, Chinmay Kulkarni, Heng-Tze Cheng, Quoc V. Le, and Ed Hui-hsin Chi. 2023. Beyond ChatBots: ExploreLLM for Structured Thoughts and Personalized Model Responses. *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:265551437>
- [66] Pattie Maes. 1994. Agents that reduce work and information overload. *Commun. ACM* 37 (1994), 30–40. <https://api.semanticscholar.org/CorpusID:59868493>
- [67] Pattie Maes and Alan Wexelblat. 1996. Interface agents. *Conference Companion on Human Factors in Computing Systems* (1996). <https://api.semanticscholar.org/CorpusID:26827189>
- [68] Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Sanchaita Hazra, Ashish Sabharwal, and Peter Clark. 2024. Data-driven Discovery with Large Generative Models. *ArXiv abs/2402.13610* (2024). <https://api.semanticscholar.org/CorpusID:267770682>
- [69] Hussein Mozannar, Gagan Bansal, Cheng Tan, Adam Fournay, Victor Dibia, Jingya Chen, Jack Gerrits, Tyler Payne, Mathews Kunzler Maldaner, Madeleine Grunden-McLaughlin, Eric Zhu, Griffin Bassman, Jacob Alber, Peter Chang, Ricky Loynd, Friederike Niedtner, Ece Kamar, Maya Murad, Rafah Hosn, and Saleema Amershi. 2025. Magentic-UI: Towards Human-in-the-loop Agentic Systems. *ArXiv abs/2507.22358* (2025). <https://api.semanticscholar.org/CorpusID:>

- 280391757
- [70] Hussein Mozannar and David Sontag. 2020. Consistent estimators for learning to defer to an expert. In *International conference on machine learning*. PMLR, 7076–7087.
- [71] Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. 2025. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393* (2025).
- [72] Jasper Tran O’Leary, Gabrielle Benabdallah, and Nadya Peek. 2023. Imprimer: Computational Notebooks for CNC Milling. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:258217042>
- [73] OpenAI. 2024. Introducing canvas. <https://openai.com/index/introducing-canvas/>.
- [74] OpenAI. 2024. Learning to Reason with LLMs. <https://openai.com/index/learning-to-reason-with-llms/>.
- [75] OpenAI. 2025. ChatGPT Agent System Card. https://cdn.openai.com/pdf/839e66fc-602c-48bf-81d3-b21eacc3459d/chatgpt_agent_system_card.pdf.
- [76] OpenAI. 2025. Introducing deep research. <https://openai.com/index/introducing-deep-research/>.
- [77] OpenAI Platform. 2024. Assistants API. <https://platform.openai.com/docs/assistants/overview>.
- [78] James C Overholser. 1993. Elements of the Socratic method: I. Systematic questioning. *Psychotherapy: Theory, Research, Practice, Training* 30, 1 (1993), 67.
- [79] Christine A Padesky. 1993. Socratic questioning: Changing minds or guiding discovery. In *A keynote address delivered at the European Congress of Behavioural and Cognitive Therapies, London*, Vol. 24.
- [80] Srishti Palani, Aakanksha Naik, Doug Downey, Amy X. Zhang, Jonathan Bragg, and Joseph Chee Chang. 2023. Relatedly: Scaffolding Literature Reviews with Existing Related Work Sections. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:256846632>
- [81] Yi-Hao Peng, Dingzeyu Li, Jeffrey P. Bigham, and Amy Pavel. 2025. Morae: Proactively Pausing UI Agents for User Choices. <https://api.semanticscholar.org/CorpusID:280985324>
- [82] Perplexity. 2024. Perplexity AI. <https://www.perplexity.ai/>.
- [83] Perplexity. 2025. Introducing Perplexity Deep Research. <https://www.perplexity.ai/hub/blog/introducing-perplexity-deep-research>.
- [84] Kevin Pu, K. J. Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2024. IdeaSynth: Iterative Research Idea Development Through Evolving and Composing Idea Facets with Literature-Grounded Feedback. *ArXiv abs/2410.04025* (2024). <https://api.semanticscholar.org/CorpusID:273186404>
- [85] Yujia Qin, Shi Liang, Yining Ye, Kunlun Zhu, Lan Yan, Ya-Ting Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Runchu Tian, Ruobing Xie, Jie Zhou, Marc H. Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023. ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs. *ArXiv abs/2307.16789* (2023). <https://api.semanticscholar.org/CorpusID:260334759>
- [86] Napol Rachatasumrit, Jonathan Bragg, Amy X. Zhang, and Daniel S. Weld. 2022. CiteRead: Integrating Localized Citation Contexts into Scientific Paper Reading. *27th International Conference on Intelligent User Interfaces* (2022). <https://api.semanticscholar.org/CorpusID:247585131>
- [87] Reworkd. 2024. AgentGPT. <https://agentgpt.reworkd.ai/>.
- [88] Adam Rule, Aurélien Tabard, and James Hollan. 2018. Exploration and Explanation in Computational Notebooks. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018). <https://api.semanticscholar.org/CorpusID:5048947>
- [89] Arvind Satyanarayan. 2024. Intelligence as Agency. In *Adjunct Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, 1–3.
- [90] Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. *ArXiv abs/2302.04761* (2023). <https://api.semanticscholar.org/CorpusID:256697342>
- [91] Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Zicheng Liu, and Emad Barsoum. 2025. Agent Laboratory: Using Llm agents as research assistants. *arXiv preprint arXiv:2501.04227* (2025).
- [92] Scite. 2025. AI for research. <https://scite.ai/>.
- [93] Omar Shaikh, Shardul Sapkota, Shan Rizvi, Eric Horvitz, Joon Sung Park, Diyi Yang, and Michael S Bernstein. 2025. Creating General User Models from Computer Use. *arXiv preprint arXiv:2505.10831* (2025).
- [94] Yijia Shao, Vinay Samuel, Yucheng Jiang, John Yang, and Diyi Yang. 2024. Collaborative Gym: A Framework for Enabling and Evaluating Human-Agent Collaboration. <https://api.semanticscholar.org/CorpusID:274964990>
- [95] Quan Shi, Michael Tang, Karthik Narasimhan, and Shunyu Yao. 2024. Can Language Models Solve Olympiad Programming? *ArXiv abs/2404.10952* (2024). <https://api.semanticscholar.org/CorpusID:269187896>
- [96] Ben Shneiderman and Pattie Maes. 1997. Direct manipulation vs. interface agents. *Interactions* 4 (1997), 42–61. <https://api.semanticscholar.org/CorpusID:27708923>
- [97] Wesley Shrum, Joel Genuth, and Ivan Chompalov. 2007. *Structures of scientific collaboration*. MIT Press.
- [98] Susan Leigh Star and James R Griesemer. 1989. Institutional ecology, translations and boundary objects: Amateurs and professionals in Berkeley’s Museum of Vertebrate Zoology, 1907–39. *Social studies of science* 19, 3 (1989), 387–420.
- [99] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (2023). <https://api.semanticscholar.org/CorpusID:258822925>
- [100] Theodore R. Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. 2023. Cognitive Architectures for Language Agents. *ArXiv abs/2309.02427* (2023). <https://api.semanticscholar.org/CorpusID:261556862>
- [101] Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. 2023. On the Planning Abilities of Large Language Models - A Critical Investigation. *ArXiv abs/2305.15771* (2023). <https://api.semanticscholar.org/CorpusID:260440590>
- [102] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi (Jim) Fan, and Anima Anandkumar. 2023. Voyager: An Open-Ended Embodied Agent with Large Language Models. *ArXiv abs/2305.16291* (2023). <https://api.semanticscholar.org/CorpusID:258887849>
- [103] Xingbo Wang, Samantha Lee Huey, Rui Sheng, Saurabh Mehta, and Fei Wang. 2024. SciDaSynth: Interactive Structured Knowledge Extraction and Synthesis from Scientific Literature with Large Language Model. *ArXiv abs/2404.13765* (2024). <https://api.semanticscholar.org/CorpusID:269293213>
- [104] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Huai hsin Chi, F. Xia, Quoc Le, and Denny Zhou. 2022. Chain of Thought Prompting Elicits Reasoning in Large Language Models. *ArXiv abs/2201.11903* (2022). <https://api.semanticscholar.org/CorpusID:246411621>
- [105] Scott Wu. 2024. Introducing Devin, the first AI software engineer. <https://www.cognition.ai/blog/introducing-devin>.
- [106] Tongshuang Sherry Wu, Haiyi Zhu, Maya Albayrak, Alexis Axon, Amanda Bertsch, Wenxing Deng, Ziqi Ding, Bill Boyuan Guo, Sireesh Gururaja, Tzu-Sheng Kuo, Jenny T Liang, Ryan Liu, Hita Mandal, Jeremiah Milbauer, Xiaolin Ni, N. Padmanabhan, Subhashini Ramkumar, Alexis Sudjianto, Jordan Taylor, Ying-Jui Tseng, Patricia Vaidos, Zhijin Wu, Wei Wu, and Chenyang Yang. 2023. LLMs as Workers in Human-Computational Algorithms? Replicating Crowdsourcing Pipelines with LLMs. *ArXiv abs/2307.10168* (2023). <https://api.semanticscholar.org/CorpusID:259982473>
- [107] John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024. SWE-agent: Agent-Computer Interfaces Enable Automated Software Engineering. *ArXiv abs/2405.15793* (2024). <https://api.semanticscholar.org/CorpusID:270063685>
- [108] Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems* 35 (2022), 20744–20757.
- [109] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *ArXiv abs/2305.10601* (2023). <https://api.semanticscholar.org/CorpusID:258762525>
- [110] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. ReAct: Synergizing Reasoning and Acting in Language Models. *ArXiv abs/2210.03629* (2022). <https://api.semanticscholar.org/CorpusID:252762395>
- [111] Xingjian Zhang, Yutong Xie, Jin Huang, Jing Ma, Zhaoying Pan, Qijia Liu, Ziyang Xiong, Tolga Ergen, Dongsu Shim, Honglak Lee, and Qiaozhu Mei. 2024. MASSW: A New Dataset and Benchmark Tasks for AI-Assisted Scientific Workflows. <https://api.semanticscholar.org/CorpusID:270370776>
- [112] Zheng Zhang, Jie Gao, Ranjodh Singh Dhaliwal, and Toby Jia-Jun Li. 2023. Visar: A human-ai argumentative writing assistant with visual programming and rapid draft prototyping. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 1–30.
- [113] Chengbo Zheng, Dakuo Wang, April Yi Wang, and Xiaojuan Ma. 2022. Telling Stories from Computational Notebooks: AI-Assisted Presentation Slides Creation for Presenting Data Science Work. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022). <https://api.semanticscholar.org/CorpusID:247594488>
- [114] Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. 2023. Language Agent Tree Search Unifies Reasoning Acting and Planning in Language Models. *ArXiv abs/2310.04406* (2023). <https://api.semanticscholar.org/CorpusID:263829963>
- [115] Yuchen Zhuang, Xiang Chen, Tong Yu, Saayan Mitra, Victor S. Bursztny, Ryan A. Rossi, Somdeb Sarkhel, and Chao Zhang. 2023. ToolChain*: Efficient Action Space Navigation in Large Language Models with * Search. *ArXiv abs/2310.13227* (2023). <https://api.semanticscholar.org/CorpusID:264405734>

A Formative Study Participants

See Table 3

B COCOA Formative Study Procedure Details

We provide more details of each portion of our formative study and its associated activities below.

B.1 Project document activity

First, we asked participants to discuss some of the ideas in their project document they shared with us, the origins of those ideas, and the ideal ways they envision AI-powered tools helping them throughout the research process. We then invited participants to a Google Doc or Slides that we converted from their project document for easy collaboration; this conversion was made by copying and pasting content from their project document if it was not already in Docs or Slides. We asked participants to brainstorm 2–3 remaining questions they have about their project that they would be interested in tackling. For each question, participants created a bulleted plan they would undertake to pursue it, which we then used to answer F2.

B.2 Probe activity

In the second activity, we presented participants with a WoZ design probe in the form of two side-by-side Google Docs. One Doc, the “planning document,” contained 3 research ideas generated by Perplexity AI [82] based on the description of research interests submitted by the participant and 1–3 of their most recent publications and/or preprints. Participants were asked to work in the doc to expand upon an idea (or combine multiple) through a self-defined plan, with the option of invoking the help of an AI assistant through requests prefaced with a “!” command. The study facilitator (the Wizard) entered this command into Perplexity AI and pasted the output to the other Doc, the “AI output document.” While waiting for the response, we asked participants to write a brief plan consisting of 3–6 steps for how they would complete the request themselves. Participants then inspected the assistant’s output and incorporated any useful text into their document. This process repeated until the participant felt like their idea was sufficiently concrete to write a short descriptive paragraph about it, typically after at least 2–3 requests to the AI. This activity helped us answer F3.

B.3 Concluding interview

The study ended with a concluding interview, where we asked participants about their experiences, perceptions, and desiderata for the AI assistant, along with preferred ways to interact with it. Each participant was given a \$35 USD honorarium after the study. The study was reviewed and approved by the Allen Institute of AI’s IRB.

C Formative Study Data Analysis

To answer F1, the first author used a hybrid inductive-deductive coding process [23] to code participants’ submitted project documents. The first 5 documents were inductively coded to surface common structural elements before the elements were deductively

applied to the remaining 4. We iterated on existing elements and added new ones as needed.

For F2, we performed inductive thematic analysis on two documents—participants’ submitted project documents and in-study planning documents. The first author sourced *intentions for planning* within submitted project documents by locating the snippet of text that initiates a plan. For example, if the plan consists of a bulleted list, the intention is often expressed in the lead-in text that immediately precedes the list. We note that an individual plan item can also signal planning intent if it contains a nested plan. The first author performed thematic analysis by inductively coding the extracted text. The first author then identified and extracted *plan steps* participants wrote in both documents before inductively coding them.

For F3, the first author performed open coding on the study transcripts before thematically analyzing the coded snippets. To enrich the data, the first author also extracted and inductively coded all requests to the assistant from participants’ planning documents. The codes were discussed and iterated upon with other project team members on a weekly basis over the course of 3 weeks.

D Further Implementation Details

D.1 Underlying LLM Agent

D.1.1 Research support tools. We created a custom tool-calling LLM agent¹³ powered by GPT-4o. While GPT-4o provides APIs for managing multi-turn conversation context, the system manages its own context using a JSON schema we developed for storing key insights from papers and other forms of output without storing the full artifact. Previous plan steps and their outputs are also represented in the schema to enrich context. To provide high-quality research support, we incorporated two best-in-class research support functionalities into our agent. The first tool allowed the agent to access The Semantic Scholar Open Data Platform [52] via its publically available search APIs.¹⁴ This tool provided the agent access to the full text and metadata (such as authors and publication venues) of millions of research papers, created by a state-of-the-art PDF extraction and knowledge graph normalization pipeline, allowing the agent to search for authors and papers based on topics, titles, and rich metadata [52]. The second tool is the *Ask this Paper*¹⁵ feature on Semantic Scholar which generates paper summaries based on the paper’s full text given an aspect of interest. This tool allowed the agent to summarize a relevant paper based on specific aspects of interest within a plan, and also synthesize specific aspects across many papers. For example, the agent can use the first tool to find papers relevant to a topic of interests from a particular author, then, used the second tool to summarize the findings of each papers, and, finally, synthesize the individual summaries into a compare and contrast statement about the set of papers for the user to review. All prompts used are provided in our Supplementary Materials.

D.2 Guidance for plan generation

We use findings from our formative study to guide the agent’s generation of initial plans. In our formative study, we had asked

¹³<https://platform.openai.com/docs/guides/function-calling>

¹⁴<https://www.semanticscholar.org/product/api>

¹⁵<https://www.semanticscholar.org/product>

P#	Gender	Age Range	Country	Research YoE	Research Area	Research Area (General)	AI Use Frequency
P1	Man	25–34	U.S.	6–10	Human-AI interaction	HCI	A couple times
P2	Man	25–34	U.S.	6–10	LLM evaluation	NLP	Occasionally
P3	Woman	18–24	U.S.	2–5	LLMs + society	NLP	Never
P4	Man	25–34	Canada	2–5	Multilingual NLP	NLP	Occasionally
P5	Woman	25–34	U.S.	6–10	Human-AI interaction	HCI	Occasionally
P6	Man	25–34	U.S.	2–5	Multimodal AI & HCI	HCI	Occasionally
P7	Woman	25–34	South Korea	2–5	LLM retrieval	NLP	Occasionally
P8	Man	25–34	U.S.	6–10	Human-AI interaction	NLP	Occasionally
P9	Woman	18–24	U.S.	0–1	Creativity support tools	HCI	A couple times

Table 3: Participants from our formative study. All participants were Ph.D. students. Country refers to the country in which the participant primarily conducts research at the time of the study. Research YoE refers to the years of experience conducting academic research. AI Use Frequency refers to how frequently the participants uses AI tools to ideate and/or iterate on research ideas. Participants selected “Occasionally” based on the description “I don’t rely on AI but sometimes tinker with it.”

participants to write down brief plans when interacting with our probe (see Appendix B). Based on participants’ preferences of which steps they prefer assigning to an AI vs. keeping for themselves (F3 from our formative study), we organized these steps into agent and user steps. We then wrote examples of how these steps are composed into plans for tackling particular questions, once again drawing from plans participants wrote in the probe activity. These examples were provided for in-context learning in COCOA’s system prompt for plan generation (see prompts in our Supplementary Materials). Finally, we also enable COCOA to learn from interactive plans the user has previously created and edited. When the agent is invoked, COCOA collects the existing plans in the document and adds them to the in-context learning examples on-the-fly. This allows the user’s co-planning efforts to be reused for similar requests elsewhere in the document.

We also enable the agent to learn from previously created plans by extracting existing plans in the document and integrating them into system context.

E Lab Study Participants

See Table 4

F Lab Study Recruitment Process and Pilot Studies

We sent out a recruitment form to collect basic demographic details, frequency of AI use in research, and a copy of a project document from an ongoing or completed project document to use during the study. One participant also participated in our formative study. 15 researchers were based in the United States, and 1 was based in Canada. Participants’ research areas ranged from large language models, to ubiquitous computing, to visualization; broadly, they

spanned HCI ($n = 10$), ML ($n = 3$), and NLP ($n = 3$). In our participant pool, many users were occasional (6) or frequent¹⁶ (6) users of AI tools in the research process.

We also recruited 4 additional participants (1 female, 2 male, 1 non-binary) from our institution for pilot studies to test out COCOA, catch usability issues, and help us refine our procedure. We learned from our pilot studies that the length of generated plans should be around 3 steps (rather than the system’s default of 5–6 steps) to fit within the study’s time constraints.

G Lab Study Procedural Details

To control for the length of the study, the first author reviewed the participants’ documents prior to the study to identify and highlight two candidate tasks. To make sure the two tasks are valid and comparable, early on in the study, the first author asked participants 1) whether the tasks have been decisively solved by the participant and whether they would still like to explore them in the study, and 2) whether the two tasks are similar in scope, such that one task does not take significantly more resources and effort to explore than the other. If these conditions were not met, the participant was prompted to rewrite one or both tasks until they were. The two tasks were then randomly assigned to COCOA and the baseline.

In the COCOA condition, participants first watched a 6-minute video demo of the system, and were then directed to a document where they would get used to using COCOA’s features on the following sample question: “How can we use AI to better society?” This practice period lasted around 10 minutes. Participants were then asked to spend 25 minutes to make as much progress as possible tackling the task assigned to this condition with COCOA, thinking aloud as they did so. They were also permitted to also use other tools (academic search engines, other AI tools, social media, etc.)

¹⁶We define “occasional” and “frequent” the same way as we do in our formative study (Table 3).

P#	Gender	Age Range	Country	Job Title	Research YoE	Research Area	Research Area (General)	AI Use Frequency
P1	Woman	25–34	U.S.	Ph.D. student	2–5	Social computing	HCI	Frequently
P2	Woman	25–34	U.S.	Ph.D. student	2–5	Social computing	HCI	Occasionally
P3	Woman	25–34	U.S.	Ph.D. student	6–10	Visualization	HCI	Occasionally
P4	Woman	25–34	U.S.	Ph.D. student	2–5	Culture + computing	HCI	A couple times
P5	Woman	25–34	U.S.	Postdoc	6–10	Software development	HCI	Occasionally
P6	Woman	25–34	U.S.	Ph.D. student	6–10	Health + computing	HCI	Occasionally
P7	Man	25–34	U.S.	Ph.D. student	2–5	LLMs	NLP	Frequently
P8	Woman	25–34	U.S.	Ph.D. student	2–5	Accessibility	HCI	Never
P9	Man	25–34	U.S.	Ph.D. student	2–5	Multimodal AI	ML	Frequently
P10	Woman	18–24	U.S.	Ph.D. student	2–5	On-device AI	ML	A couple times
P11	Man	18–24	U.S.	Ph.D. student	2–5	Ubiquitous computing	HCI	Always
P12	Man	25–34	U.S.	Ph.D. student	6–10	LLM evaluation	NLP	Frequently
P13	Man	35–44	Canada	Ph.D. student	2–5	Computational biology	ML	Frequently
P14	Woman	18–24	U.S.	Ph.D. student	2–5	LLMs	NLP	Frequently
P15	Woman	25–34	U.S.	Ph.D. student	2–5	Social computing	HCI	Occasionally
P16	Man	25–34	U.S.	Postdoc	6–10	Human-AI interaction	HCI	Occasionally

Table 4: Participants from our user study. Country refers to the country in which the participant primarily conducts research at the time of the study. Research YoE refers to the years of experience conducting academic research. AI Use Frequency refers to how frequently the participants use AI tools to ideate and/or iterate on research ideas. Participants selected “Occasionally” based on the description “I don’t rely on AI but sometimes tinker with it.”

alongside COCOA, although very few did. After the 25 minute session, participants wrote a brief set of takeaways and next steps from their exploration, and filled out a short evaluation form with 5-point Likert scale questions about their experience (see Appendix H for this form). After submitting the form, they were encouraged to further try out COCOA on other parts of their project document for another 10–15 minutes.

In the baseline condition, we did not provide participants a tutorial because the chat interface was already ubiquitous beyond interacting with AI agents. Participants were once again asked to spend 25 minutes making as much progress as possible tackling the task assigned to this condition, thinking aloud as they did so and using other tools if desired. They wrote brief takeaways and next steps after the 25 minutes was up and filled out the same evaluation form as the COCOA condition.

Because participants identified literature search and understanding as one of the areas they can benefit most from AI assistance (Section 3.2.2), we had participants tackle a literature-augmented task in both COCOA and the baseline. By *literature-augmented*, we mean tasks that are to be completed by referencing or drawing from academic literature; in the study, participants worked on not only literature review tasks, but also tasks where researchers needed to make literature-informed decisions such as study design and ideation. This stands in contrast with tasks focused on writing mechanics (e.g., rewording a sentence) or tasks that do not involve literature (e.g., booking travel) that past work has covered (e.g., [51, 57, 65, 107, 108, 112]).

H Lab Study Self-Evaluation Form Questions

All questions were answered on a 5-point Likert scale.

- The system’s outputs were useful to me for exploring the research problem at hand
- I obtained new and useful insights from using the system
- The system helped me clearly understand the steps needed to effectively tackle this particular problem
- By using the system, I’ve developed a better understanding of potential solutions to this particular problem
- I found the system easy to use
- I could easily steer the system towards doing something helpful
- I could see myself easily integrating this system into my workflow
- I’m satisfied with the artifact (summary and next steps) that I’ve created after using the system
- I feel confident about my next steps after using the system

I Deployment Study General Guidelines

- (1) Try to spend a total of at least 90 minutes with Cocoa during the 7-day use period.
- (2) It may help to jot down any notable experiences with Cocoa so we can more easily discuss them later.
- (3) Treat the document as a scratch space to tinker with and explore rough ideas.
- (4) Cocoa is just another tool in your toolbox. You can use it alongside whatever you typically use in your workflows!

- (5) If the system crashes at any point, try refreshing the page.
The document auto-saves regularly.