

JUMPSTARTER: Human-AI Planning with Task-Structured Context Curation

Anonymous ACL submission

Abstract

Human-AI planning for complex goals remains challenging with current large language models (LLMs), which rely on linear chat histories and simplistic memory mechanisms. Despite advances in long-context prompting, users still manually manage information, leading to a high cognitive burden. Hence, we propose *JumpStarter*, a system that enables LLMs to collaborate with humans on complex goals by dynamically decomposing tasks to help users manage context. We specifically introduce *task-structured context curation*, a novel framework that breaks down a user’s goal into a hierarchy of actionable subtasks, and scopes context to localized decision points, enabling finer-grained personalization and reuse. The framework is realized through three core mechanisms: context elicitation, selection, and reuse. We demonstrate that *task-structured context curation* significantly improves plan quality by 16% over ablations. Our user study shows that *JumpStarter* helped users generate plans with 79% higher quality compared to ChatGPT.¹

1 Introduction

Planning is a core cognitive process for solving complex, goal-oriented tasks (Miller and Venditto, 2021; Ho et al., 2022). Recent advances in the planning abilities of large language models (LLMs) (Valmeekam et al., 2023; Shinn et al., 2023) have enabled human-AI planning across various domains, such as travel planning (Xie et al., 2024), manufacturing, and healthcare (Lee et al., 2025). In these settings, LLMs assist users in decomposing complex goals into actionable steps (Wei et al., 2022; Shinn et al., 2023). However, while LLMs can support planning at scale, they still struggle with maintaining context over long interactions (Jiang et al., 2023). Despite improvements in memory mechanisms and extended context windows

(Luo et al., 2025), users must actively manage what information to provide, as LLMs frequently forget key details, leading to a degraded user experience. To generate personalized plans, users often resort to intensive prompt engineering, manually curating chat histories to elicit relevant responses, resupply information, and manage subtasks. This process is opaque since users have little insight into what context is retrieved. It is also burdensome, often resulting in generic output misaligned with prior interactions. To support complex, personalized workflows, LLM-based systems need mechanisms that treat both tasks and context as dynamic, composable units, enabling more modular, transparent, and efficient human-AI planning.

To address this limitation, we propose *task-structured context curation*, a novel framework that enables LLMs to collaborate with humans on complex goals by dynamically decomposing tasks to help users manage context. By decomposing a user’s goal into a hierarchy of actionable subtasks, our system scopes context to localized decision points, enabling finer-grained personalization and reuse. This approach is implemented through three core mechanisms for context management: 1) *Context Elicitation*: the system prompts users for missing information specific to each subtask; 2) *Context Selection*: the system and users collaboratively select and surface only the most relevant prior inputs for each subtask; and 3) *Context Reuse*: user-approved answer drafts (e.g., emails or study schedules) are saved and incorporated into future subtasks. Together, these mechanisms provide an alternative to labor-intensive long-context prompting, enabling more structured, adaptive, transparent, and user-controllable LLM interactions.

We specifically operationalize this method in *JumpStarter*, a human-AI planning system that supports open-ended user goals through structured task planning and personalized output generation. The process, as shown in Figure 1, begins when a user

¹We will release the code upon publication.

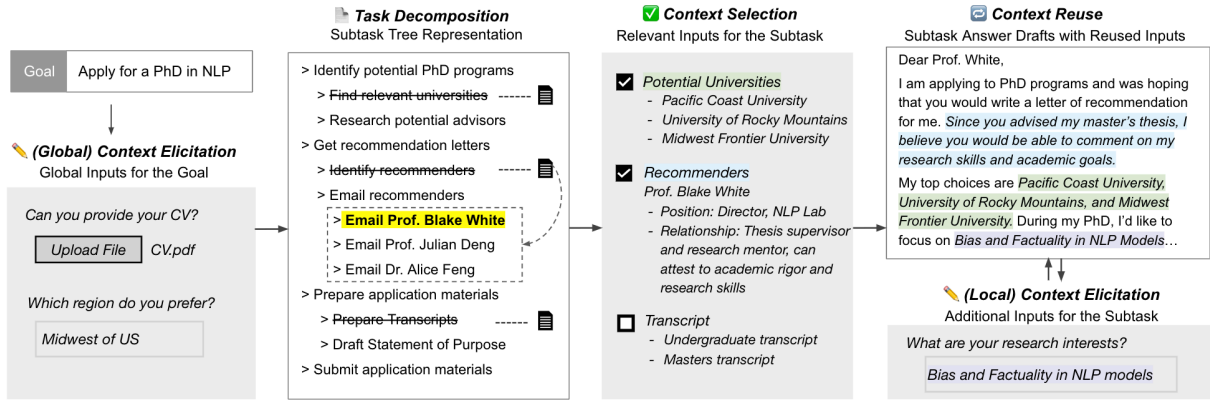


Figure 1: *JumpStarter* helps users get started on their personal goals through *task-structured context curation*. It first takes the user’s goal and elicits context for the goal. It then decomposes the goal into actionable subtasks. For each subtask, it helps users select relevant context and write answer drafts. It also aids users in refining these drafts by eliciting further context. *Task-structured context curation* improves plan quality over ablations. Our user study showed that *JumpStarter* helped users generate plans with 79% higher quality compared to using GPT-4o via the ChatGPT interface.

inputs their personal goal. The system then poses questions to elicit the relevant context for the goal. Based on the user’s responses, the system breaks down the goal into subtasks that the user can explore. If the system detects that a subtask is not actionable enough, it prompts the user, who can break it down further into additional subtasks. When the user is ready to work on a task, the system selects the relevant context and generates answer drafts. If the user is not satisfied with the suggested solution, they can request more elicitation questions or provide their own prompts for refinement. Users can modify the suggested solution, and also save it as context to be used by the system in further subtasks. The user can then move to explore the next subtask, moving towards the completion of their goal. This task-structured workflow enables users to iteratively plan their goals by focusing on one actionable step at a time, while allowing the LLM to operate on localized, context-rich prompts that improve output quality throughout the planning process.

We evaluate *JumpStarter* through both technical evaluation and user studies. For subtask detection, our prompting technique that combines chain-of-thought (CoT) reasoning, the tree level of each subtask node, and draft quality achieves 87% accuracy, significantly outperforming a 35% baseline. In a controlled expert evaluation (N=6), our full *task-structured context curation* pipeline significantly improves plan quality by 16% compared to ablated variants. In a within-subjects user study (N=10), *JumpStarter* helps users generate plans with 79%

higher quality compared to using GPT-4o from the ChatGPT interface. It also reduces perceived task load across all NASA-TLX (Hart, 2006a) dimensions. These results demonstrate that structured context curation enhances model effectiveness and user experience over flat long-context prompting.

Overall, our contributions are three-fold:

- A novel framework, namely *task-structured context curation*, for complex human-AI planning, involving hierarchical task decomposition with three components: context elicitation, selection, and reuse;
- A human-AI planning system called *JumpStarter* that helps people get started on complex personal goals by identifying when and how to decompose large tasks, and by generating answer drafts with suggested context to create detailed and personalized plans.
- Empirical validation through automatic and human-centered evaluations demonstrating clear benefits of *task-structured context curation* over existing approaches in human-AI planning, and suggesting a promising direction for building personalized, goal-oriented LLM-based systems.

2 Related Work

Task Structuring and Planning with Language Models Recent work has explored how LLMs can support planning by generating step-by-step breakdowns for user goals. Studies on LLM-based

agents for tasks like travel, scheduling, and gardening often use chain-of-thought prompting or multi-agent setups to scaffold plans (Lal et al., 2024; Xie et al., 2024; Zheng et al., 2024). However, these systems typically operate as single-turn planners with limited personalization and shallow interaction loops. Other tools offer lightweight task scaffolding through user interfaces, but they generally support only single-layer decomposition without context management. For example, ExploreLLM (Ma et al., 2024) prompts users to manually decompose goals and specify preferences via a schema-based interface. While structured guidance improves planning, it assumes heavy manual effort. In contrast, JumpStarter co-creates a multi-layer subtask hierarchy with the user and localizes context dynamically, enabling finer-grained personalization and reuse across planning workflows.

Several goal-directed LLM agents, such as AutoGPT (Significant Gravititas), attempt autonomous task chaining but often lack transparency, robustness, and meaningful user interaction (Zheng et al., 2024). Moreover, prior work notes that humans rarely perform spontaneous, exhaustive planning due to cognitive cost (Krieger et al., 2009; Teevan et al., 2016), underscoring the need for systems that scaffold both planning and personal context. Our work uniquely combines LLM-based goal decomposition with context-aware output generation across evolving task structures, aiming to reduce user burden while enhancing planning continuity.

Context Curation in LLM-based Systems Effective context curation for multi-step tasks requires more than flat prompting or opaque memory. Prior work in proactive dialogue systems explores context elicitation and information-seeking strategies (Zhang et al., 2024; Malaviya et al., 2024; Deng et al., 2023; Zhang et al., 2023), but these focus on short-turn interactions in synthetic or search-based tasks, rather than structured, long-context workflows involving draft reuse or evolving context. Other systems integrate retrieval or scratchpad-style memory to maintain context (Liu et al., 2024), yet typically store information in flat formats without explicit links to task hierarchies. For example, RAISE (Liu et al., 2024) mirrors short- and long-term memory but lacks subtask anchoring. In contrast, JumpStarter grounds context in a hierarchical task structure, enabling selective retrieval and reuse of prior inputs and outputs across subtasks.

Our work also draws on task-centric information

management from prior HCI systems (Kerne et al., 2014; Jones et al., 2008; Kaptelinin, 2003; Conley and Carpenter, 2007; Kersten and Murphy, 2006), reimaged for LLM-based workflows. Rather than relying on long prompts or passive memory modules, we propose a modular architecture—context elicitation, condensation, and reuse—that adapts dynamically as tasks evolve. While recent efforts like the Model Context Protocol (MCP) (Anthropic, 2024) introduce standardized interfaces for external tool access, they do not address subtask-scoped context curation. JumpStarter complements such protocols by structuring context around subtasks, providing finer control over planning and execution in personalized workflows.

3 JumpStarter System

JumpStarter is an LLM-based interactive system that enables human-AI planning with *task-structured context curation*. As a motivating example, we focus on planning for users’ personal projects. The system takes the user-specified goal as input, and outputs personalized plans and answer drafts for the user. Here, an answer draft refers to any tangible artifact that assists users in taking real-world actions related to a task, for example, an email draft for requesting recommendation letters or a study schedule to prepare for an exam.

This section describes how *JumpStarter* works with an example walk-through, and the implementation details of the system.

3.1 System Walkthrough

To illustrate how users interact with *JumpStarter*, we present a walkthrough using a concrete example. Consider John, a user aiming to apply for a PhD in NLP. Figures in the Appendix A provide corresponding interface visuals.

3.1.1 Goal Input and Global Context Elicitation

John begins by entering his goal—“Apply for a PhD in NLP”—into a text input box and clicks *Start*. The system then generates elicitation questions to collect relevant context (e.g., existing documents, preferences). In John’s case, it asks about potential target schools and recommendation letters. Since he is uncertain about the former and has not yet obtained the latter, he uploads his CV in response to the first question and clicks *Let’s get started*. The elicited context becomes part of the global context used across subsequent tasks.

244	3.1.2 Subtask Generation and Detection	3.1.5 Context Elicitation for Draft Iteration	291
245	Using the elicited context, <i>JumpStarter</i> generates	Unsatisfied with the initial email draft, John clicks	292
246	an initial subtask tree, presenting titles, descrip-	<i>Add Context and Regenerate</i> . The system prompts	293
247	tions, and estimated durations (Figure 6). John	follow-up questions to elicit additional details (e.g.,	294
248	reviews the list to gain an overview of the plan and	specific projects or papers), which John provides	295
249	decides to explore the subtasks sequentially.	(Figure 9b). The refined draft incorporates these	296
250	John selects the first subtask, <i>Identify Potential</i>	details, resulting in a more personalized and accept-	297
251	<i>PhD Programs</i> . The system detects that it is in-	able version.	298
252	sufficiently actionable and prompts John to either	Through structuring the goal as a hierarchy of	299
253	decompose the task or proceed with drafting. Fol-	tasks, <i>JumpStarter</i> helps John curate relevant con-	300
254	lowing the system’s recommendation, John selects	text, enabling the creation of detailed action plans	301
255	<i>Decompose the task</i> , resulting in a new set of sub-	and high-quality answer drafts. <i>JumpStarter</i> ’s	302
256	tasks under the original node (Figure 7).	structured context management and iterative refine-	303
257	3.1.3 Answer Draft Creation and Refinement	ment enable him to effectively progress toward his	304
258	John selects the subtask <i>Research Universities and</i>	goal of applying for a PhD in NLP.	305
259	<i>Programs</i> , which the system deems actionable. It	3.2 Data Representation	306
260	generates an initial answer draft—a list of NLP	<i>JumpStarter</i> represents each piece of context as	307
261	PhD programs. Upon review, John refines the out-	a key-value pair, where the key denotes the con-	308
262	put by adding, “ <i>I want schools in the Midwest of</i>	text name and the value specifies its content. For	309
263	<i>the US.</i> ” He is satisfied with the revised list and	example, the pair “ <i>Location preference: Midwest</i>	310
264	saves it as an answer draft. He is also given three re-	<i>of US</i> ” captures a user’s geographical preference.	311
265	finement options: regenerate, add more context and	The system maintains two types of context: global	312
266	regenerate, or iterate on the current draft. Saved	and local. Global context consists of information	313
267	drafts are stored as context and appear as icons on	elicited from the user immediately after goal speci-	314
268	the task tree, marking completed nodes.	fication and is universally applied throughout plan-	315
269	3.1.4 Task Forking	ning. Local context includes answer drafts and any	316
270	John proceeds to the next subtask, <i>Identify Faculty</i>	additional information the user provides during in-	317
271	<i>Members</i> . The system suggests decomposing the	teraction.	318
272	task and detects that it requires forking based on	To organize tasks, <i>JumpStarter</i> employs a hi-	319
273	the previously identified programs. It selects the	erarchical tree structure. This design mirrors the	320
274	saved university list as relevant context and asks	natural decomposition of complex goals into man-	321
275	John to confirm or modify the selection (Figure 8a).	ageable components, facilitating clear tracking of	322
276	Upon confirmation, the system forks the task into	task dependencies and progression. Each user goal	323
277	program-specific subtasks (Figure 8b), which John	serves as the root node, with system-generated sub-	324
278	begins to explore individually.	tasks represented as child nodes. Each subtask	325
279	John then explores the task <i>Get Recommendation</i>	node stores task-specific attributes, including titles,	326
280	<i>Letters</i> and decomposes it as prompted. For	descriptions, estimated durations, and any associ-	327
281	the subtask <i>Compile a List of Recommenders</i> , the	ated answer drafts.	328
282	system uses his CV to generate a list including	3.3 Task-Structured Context Curation	329
283	Prof. Blake White, Prof. Julian Deng, and Dr. Al-	<i>Task-Structured Context Curation</i> leverages a hier-	330
284	lice Feng. John accepts and saves the draft. He	archical subtask tree to support two core operations.	331
285	proceeds to <i>Reach Out to Potential Recommenders</i> ,	First, it performs subtask detection and generation	332
286	which is forked into person-specific subtasks. For	to expand the task tree. Second, it curates con-	333
287	<i>Reach Out to Prof. Blake White</i> , the system iden-	text by eliciting, selecting, and reusing relevant	334
288	tifies relevant context—including John’s CV, prior	information. The following sections describe how	335
289	collaborations, and the university list—and gener-	each operation is implemented, with corresponding	336
290	ates a personalized email draft (Figure 9a).	prompts provided in Section E.	337

3.3.1 Task Decomposition

Subtask Generation When a user chooses to break down a task, *JumpStarter* decomposes it into a list of subtasks, which are attached as child nodes to the corresponding task node. We prompt GPT-4o to generate these subtasks, each including a name, description, and estimated completion time to help users better understand the scope and duration of the task. To reduce redundancy across the task tree, GPT-4o is prompted with the current tree structure and instructed to take it into account when suggesting new subtasks, ensuring coherence and avoiding repetition (Figure 13).

Subtask Detection For each task selected by the user, *JumpStarter* first evaluates whether it is sufficiently detailed and actionable. If not, the system recommends further decomposition; if so, it prompts the user to proceed with generating an answer draft. To identify an effective prompting strategy, we experimented with multiple techniques and found that Chain-of-Thought (CoT) prompting (Wei et al., 2022) with few-shot examples, augmented with the task’s tree level, achieved the best balance between accuracy and latency. We refer the reader to Section C for experimental details and results.

Task Forking When a task is flagged for further decomposition by the Subtask Detection module, *JumpStarter* evaluates whether it requires Forking—a decomposition strategy based on distinct entities (e.g., multiple recipients or institutions) that enables subtasks to be completed in parallel. Unlike standard sequential breakdowns, forking produces independent subtasks that do not rely on order. We prompt GPT-4o to determine whether a task should be forked (See Figure 15). If forking is not needed, the system proceeds with standard hierarchical decomposition; if forking is detected, *JumpStarter* then performs context selection to identify the most relevant context keys for guiding parallel subtask generation. These context elements help ensure that the resulting entity-specific subtasks are well-scoped and non-redundant.

3.3.2 Context Elicitation

Based on the user’s specified goal, *JumpStarter* identifies relevant information the user may need to provide. The model generates suggestions for context-relevant documents (e.g., resumes, transcripts) along with corresponding elicitation questions. These questions are then displayed in the

UI, where users can either upload the suggested documents or respond directly. If a question does not require a file, the system instead provides a text input box, allowing the user to type their answer inline. The user’s goal, along with all elicitation questions and their corresponding responses, is then compiled into the global context, which is incorporated into the system prompt to guide all subsequent planning features.

In addition, when users want to improve an answer draft for a subtask but are unsure how to proceed, the system generates clarifying questions to elicit additional context from the user. With the new input, the system regenerates an improved version of the answer draft. See Appendix E.1 for relevant prompts.

3.3.3 Context Selection and Context Reuse

When a user chooses to work on a task, *JumpStarter* assists in generating an answer draft through the following multi-step process. First, *JumpStarter* identifies the most relevant context for the selected task. Given the task title and description, our system selects relevant entries from the local context. These suggested context keys are shown in a checklist pop-up, where users can deselect irrelevant items or add additional context from a drop-down list. This allows users to refine the context before generation. Subsequently, the curated context is then used to generate an answer draft by prompting GPT-4o with the task title, description, selected context, the user’s goal, and global context (See prompts in Appendix E.2). Once approved, answer drafts are stored in the local context and can be reused in future subtasks, supporting continuity and enabling task forking throughout the planning workflow.

4 Evaluation of Task-structured Context Curation

JumpStarter supports hierarchical planning by identifying actionable tasks and curating personal context to guide answer draft generation. We evaluate the accuracy of subtask detection and detail the experiment settings and results in Section C. In this section, we examine how *task-structured context curation*² influences the quality of generated plans and answer drafts. Specifically, we conducted a controlled within-subjects lab study for the following conditions: 1) only context reuse; 2) only con-

²We use *task-structured context curation* and *context curation* interchangeably in the paper.

text selection and reuse; and 3) context elicitation, selection, and reuse.

Besides ablating the use of both context selection and context elicitation in Condition (1) and context elicitation alone in Condition (2), all features and UI were kept exactly the same. We hypothesized that action plans and answer drafts developed with the full context curation method, represented by Condition (3), would be judged as higher quality than those created without it.

4.1 Participants and Procedure

We evaluated the effect of context curation on subtask and answer draft quality for three preselected personal goals inspired by Little (1983): (a) Apply to a fellowship, (b) Get a driver’s license, and (c) Organize a team event. We then used a university mailing list to recruit participants. Each goal was assigned two expert participants—those who reported achieving the goal in the past six months. Overall, we recruited six expert participants for this study (average age=25.8, three female, three male). Participants were compensated \$20 per hour, with sessions lasting about 1.5 hours each. The detailed procedure of the study is illustrated in Appendix B.

4.2 Results and Findings

Overall, 46 subtasks (i.e. plan items) were generated for Condition 1, 50 for Condition 2, and 50 for Condition 3. An equal number of answer drafts were produced, as we asked participants to generate a draft only once for each subtask.

Our method of *context curation*, represented by Condition 3, performed the best for both subtask quality and answer draft quality. For subtask quality, *context curation* ($\mu=6.12$, $\sigma=0.95$) was rated higher than Condition 1 ($\mu=5.26$, $\sigma=1.26$) and Condition 2 ($\mu=5.28$, $\sigma=1.25$). For the quality of answer drafts, *context curation* ($\mu=6.36$, $\sigma=0.95$) outperformed Condition 2 ($\mu=5.68$, $\sigma=1.01$) by a similar margin, which also outperformed Condition 1 ($\mu=5.04$, $\sigma=1.04$). The results are shown in Figure 2.

Our results show that *context selection* significantly improves the quality of answer drafts compared to the baseline. Unlike *context reuse only*, which includes all previously saved context indiscriminately, *context selection* prompts LLMs to **selectively retrieve only the most relevant information** from the available pool. Participants frequently noted that outputs under this condition reflected earlier inputs more effectively. For ex-

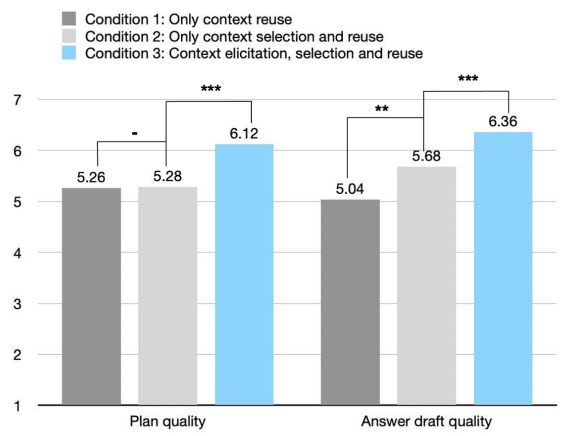


Figure 2: Expert evaluation of plan and answer draft quality across three conditions. Our task-structured context curation method (context elicitation, selection, and reuse) significantly outperforms both context reuse only and context selection and reuse only. Improvements are statistically significant (** $p < 0.01$, *** $p < 0.001$).

ample, the system generated personalized checklists for applying for scholarships and synthesized itineraries for group events. As one participant described, “*Very useful synthesis of everything I’ve explored.*” In contrast, *context reuse only* often produced generic suggestions. Another participant noted, “*I like the rec letter request email it gives me, as it considers much of my background that I saved in the previous ‘update your CV’ task. I did the same thing in the other version [context reuse only] but did not feel it was as effective.*”

Beyond *context selection*, *context elicitation* further improved the quality of both subtasks and answer drafts. Participants noted that the elicitation questions posed early in the session “*provided the right plan to start with.*” These questions helped **tailor subtasks to users’ specific situations**—for example, adjusting the number of recommendation letters required for a particular fellowship. Similarly, elicitation prompts such as the user’s state of residence (for obtaining a driver’s license) or team size (for organizing an event) led to more precise and **context-aware task breakdowns**. Preference-based questions—e.g., “*What type of vehicle do you intend to drive?*”—enabled participants to supply individualized context, resulting in more relevant and highly rated answer drafts.

5 User Study

To evaluate how *JumpStarter* supports users in planning personal goals, we conducted a within-subjects study with ten participants, comparing

Participant	Personal goal	Goal type
P1	Start a side job	Career
P2	Organize a weekly game night	Life
P3	Land a job offer	Career
P4	Prepare for the LSAT	Academia
P5	Manage social media accounts	Creativity
P6	Move to a new apartment	Life
P7	Create a portfolio website	Creativity
P8	Prepare to deliver a tutorial	Academia
P9	Start a YouTube channel	Creativity
P10	Organize a family reunion	Life

Table 1: Overview of personal goals picked by participants in the user study.

JumpStarter to ChatGPT.³ We selected ChatGPT as a baseline due to its widespread familiarity and its capacity to accept contextual input across a session, making it a realistic point of comparison for goal-oriented planning.

Each participant used both systems to explore a personal goal they intended to pursue in the near future. After interacting with each system, participants completed questionnaires assessing task load and satisfaction with outcomes. We also conducted semi-structured interviews to gather qualitative insights into their experience with each system.

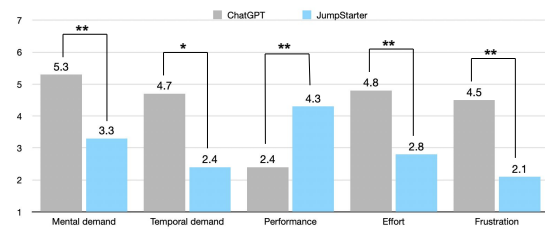
5.1 Participants and Procedure

We recruited ten participants (average age=23.8; six female, four male) through a university mailing list and word of mouth. All reported being familiar or very familiar with ChatGPT. Before the study, each participant selected a personal goal to pursue within the next six months (see Table 1). At the start of each session, participants were introduced to the concepts of action plans and answer drafts through examples, then used both ChatGPT and *JumpStarter* to plan and generate answer drafts for as many subtasks as possible. The system order was randomized and counterbalanced to mitigate order effects. Each task was limited to 25 minutes, and each study session lasted approximately 1.5 hours. All participants were compensated \$20 per hour.

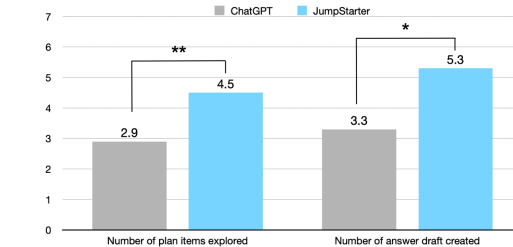
5.2 Results and Findings

We collected participants’ ratings on a 1-7 point scale through questionnaires (see Table 4 for detailed results). Participants rated their task load, outcome satisfaction, and confidence level in taking

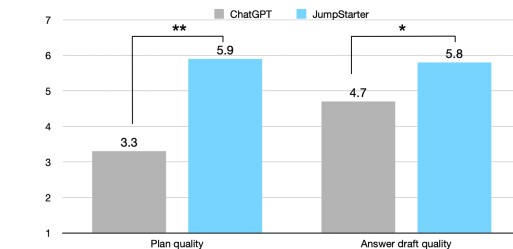
³We used the chat interface of ChatGPT and specified GPT-4o in the user study.



(a) NASA Task Load Index Comparison



(b) Exploration Efficiency Comparison



(c) Result Satisfaction Comparison

Figure 3: User study results comparison between using ChatGPT and using *JumpStarter*. The statistical test results comparing *JumpStarter* with ChatGPT, where the p-values (*: $p < .050$, **: $p < .010$, ***: $p < .001$) are reported.

the next steps on their goal by using *JumpStarter* and ChatGPT. We recorded the number of plan items explored and the answer draft generated by each participant in each task. During the interviews, we asked them follow-up questions to understand the reasons behind their scores. We applied the thematic analysis method (Braun and Clarke, 2006) to analyze the interview transcripts. We report the results in Figure 3 and discuss the key findings in this section. See Appendix D for more detailed analysis.

5.2.1 *JumpStarter* Reduces Users’ Task Load

As shown in Figure 3a, in the NASA TLX dimensions (Hart, 2006b), working with *JumpStarter* was significantly less demanding in mental demand ($p=.005$), temporal demand ($p=.012$), performance ($p=.007$), effort ($p=.007$), and frustration ($p=.009$).

Participants attributed this to the structured interface and context-aware guidance of *JumpStarter*, which helped reduce cognitive load. Unlike Chat-

GPT’s linear chat format, *JumpStarter* made it easier to visualize progress and maintain task structure. As P2 noted, “*ChatGPT info dumps a lot, and I have to keep the structure in my brain, whereas JumpStarter gave me a structure that I could easily follow.*” Similarly, targeted questions helped users refine their input and move forward. P10 commented, “*I appreciate the questions JumpStarter asked when I felt stuck about how to iterate the answer draft.*” In contrast, ChatGPT required users to generate and manage context manually. P8 remarked, “*In ChatGPT, the information load is high—I have to think very hard about what info I should provide to get things that work for me.*”

Participants also reported expending more effort with ChatGPT, often without meaningful improvement in output quality. For example, P1 asked ChatGPT to generate clarifying questions to improve the answer draft. It returned eight, which P1 described as “*a bit too abstract and hard to answer.*” Despite answering them all, P1 felt the revised output remained “*too general and not useful.*” They also noted that “*ChatGPT seemed to forget these eight answers soon after*”, leading to frustration and a sense that their effort was wasted.

5.2.2 *JumpStarter* Enables Better Exploration Efficiency

Participants explored significantly more plan items using *JumpStarter* (mean=4.5, SD=1.43) than with ChatGPT (mean=2.9, SD=0.74), and created more answer drafts as well (mean=5.3, SD=1.89 vs. mean=3.3, SD=1.16). Figure 3b illustrates this difference.

One reason for this disparity is that ChatGPT’s linear interaction format often leads users to fixate on a single task, limiting broader exploration. For instance, P2 iterated nine times on one answer draft: “*It took nine iterations to get the draft I like. I really hoped ChatGPT would guide me, but I had to direct myself. I got so involved that I completely forgot I had other planning items.*”

In contrast, *JumpStarter*’s task-structured interface helped users stay oriented within the broader plan while focusing on one task at a time. P4 remarked, “*JumpStarter has a more flexible structure. I like that I can easily jump between tasks... Seeing the task description and relevant context gives me everything I need*” P10 similarly noted, “*JumpStarter automatically manages and considers my drafts from previous tasks, which is great and helps me focus on the current one.*”

5.2.3 *JumpStarter* Improves Perceived Quality of Outcome

Plans created with *JumpStarter* were rated significantly higher in quality (mean=5.9, SD=0.88) than those created with ChatGPT (mean=3.3, SD=1.49). Similarly, participants rated the quality of answer drafts higher with *JumpStarter* (mean=5.8, SD=0.92) compared to ChatGPT (mean=4.7, SD=1.25). Figure 3c shows these differences in perceived quality.

Participants attributed this improvement to *JumpStarter*’s proactive questioning and personalized planning. Unlike ChatGPT, which often produced generic plans, *JumpStarter* asked targeted questions early in the process and incorporated user responses into the plan structure. For example, P9, who was planning to start a YouTube channel, noted: “*JumpStarter asked if I was experienced with video editing, and I said yes. It was reflected in the plan accordingly, unlike ChatGPT, which focused too much on editing I didn’t need.*”

In addition, participants also appreciated how *JumpStarter* effectively uses the relevant context to personalize the answer draft, as P1 noted, “*JumpStarter gave me much more tailored responses—like a personalized schedule to help me start my side job. It took into account key details like my limited time and the specific area I am interested in. It handled the context very well.*”

6 Conclusion

We introduced *task-structured context curation*, a novel framework for improving human-AI planning by aligning context management with hierarchical task decomposition. Operationalized in *JumpStarter*, this approach enables users to break down complex goals into actionable subtasks while dynamically eliciting, selecting, and reusing relevant context. Through both technical evaluations and user studies, we show that *JumpStarter* significantly improves subtask and answer draft quality, reduces task load, and promotes broader exploration compared to ChatGPT. Our findings highlight the limitations of traditional long-context prompting and underscore the importance of structuring both tasks and context in human-AI planning workflows. By treating planning as an interactive, task-decomposed process, *JumpStarter* offers a design paradigm for building more transparent, modular, and personalized LLM-based assistants.

7 Limitations

JumpStarter primarily assists users in figuring out “how” to achieve their personal goals, prompting real-world actions by increasing simplicity, as suggested by the Fogg Behavior Model (Fogg, 2009). However, the “why” problem—motivation (such as self-regulation or emotional challenges)—is another very important dimension that *JumpStarter* does not address. For example, one participant shared that their confidence dropped after exploring the goal, as they realized how much they had to do to make it happen. Emotional support at this time would be very valuable. But as prior studies suggested (Bhattacharjee et al., 2024), we should be very careful about using LLMs to provide emotional support, which deserves further investigation.

JumpStarter was designed to support people with personal goals that primarily involve cognitive or knowledge work, encompassing a wide range of goals across academia, career, creativity, and life (see examples in Table 1). However, there are many other personal goals that *JumpStarter* does not cover, such as physical goals (e.g., losing weight), behavioral goals (e.g., overcoming shyness), and spiritual goals (e.g., coming to terms with one’s faith), as outlined in (Little, 1983). Many of these goals require motivational support, a topic explored in previous HCI works (Consolvo et al., 2009; Ekhtiar et al., 2023). Future iterations of *JumpStarter* could potentially expand to include these areas to accommodate a broader spectrum of personal goals.

JumpStarter utilizes GPT-4o as its core engine for providing information. While LLMs are adept at synthesizing information from the Internet and can sometimes offer valid and useful references, they are also prone to generating inaccurate or hallucinated information (Achiam et al., 2023). Hence, it is crucial to integrate search agents into our system to enhance the credibility of the information provided. Also, when pursuing creative goals such as starting a YouTube channel, users require support in collecting and analyzing real-world data. Developing a search agent specifically tailored to platforms like YouTube could be a valuable direction to explore.

8 Ethics Statement

Biases We did not explicitly address potential biases present in the pre-trained language models

used in this study.

Reproducibility In this work, we use GPT-4o as the backbone of our *JumpStarter* system. To support reproducibility, we provide all prompts used in the paper.

Study Participants We recruit participants for both technical evaluation and user study through a university mailing list and word of mouth. They were paid at a rate of \$20 per hour. To protect privacy and anonymity, participants’ personal and demographic information will not be released.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2024. [Introducing the model context protocol](#). Accessed: 2025-04-29.
- Eric Bergman and Donald A Norman. 2000. Making technology invisible. *Information Appliances and Beyond: Interaction Design for Consumer Products*.
- Ananya Bhattacharjee, Yuchen Zeng, Sarah Yi Xu, Dana Kulzhabayeva, Minyi Ma, Rachel Kornfield, Syed Ishtiaque Ahmed, Alex Mariakakis, Mary P Czerwinski, Anastasia Kuzminykh, et al. 2024. Understanding the role of large language models in personalizing and scaffolding strategies to combat academic procrastination. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–18.
- Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101.
- Kenneth Conley and James Carpenter. 2007. Towel: Towards an intelligent to-do list. In *AAAI Spring Symposium: Interaction Challenges for Intelligent Assistants*, pages 26–32.
- Sunny Consolvo, Predrag Klasnja, David W McDonald, and James A Landay. 2009. Goal-setting considerations for persuasive technologies that encourage physical activity. In *Proceedings of the 4th international Conference on Persuasive Technology*, pages 1–8.
- Fergus IM Craik and Robert S Lockhart. 1972. Levels of processing: A framework for memory research. *Journal of verbal learning and verbal behavior*, 11(6):671–684.
- Yang Deng, Wenqiang Lei, Wai Lam, and Tat-Seng Chua. 2023. A survey on proactive dialogue systems: Problems, methods, and prospects. *arXiv preprint arXiv:2305.02750*.

772	Tina Ekhtiar, Armağan Karahanoğlu, Rúben Gouveia, and Geke Ludden. 2023. Goals for goal setting: a scoping review on personal informatics. In <i>Proceedings of the 2023 ACM Designing Interactive Systems Conference</i> , pages 2625–2641.	827
773		828
774		829
775		830
776		
777	Brian J Fogg. 2009. A behavior model for persuasive design. In <i>Proceedings of the 4th international Conference on Persuasive Technology</i> , pages 1–7.	831
778		832
779		833
780	Sandra G Hart. 2006a. Nasa-task load index (nasa-tlx); 20 years later. In <i>Proceedings of the human factors and ergonomics society annual meeting</i> , volume 50, pages 904–908. Sage publications Sage CA: Los Angeles, CA.	834
781		835
782		836
783		837
784		
785	Sandra G Hart. 2006b. Nasa-task load index (nasa-tlx); 20 years later. In <i>Proceedings of the human factors and ergonomics society annual meeting</i> , volume 50, pages 904–908. Sage publications Sage CA: Los Angeles, CA.	838
786		839
787		840
788		841
789		842
790	Mark K Ho, David Abel, Carlos G Correa, Michael L Littman, Jonathan D Cohen, and Thomas L Griffiths. 2022. People construct simplified mental representations to plan. <i>Nature</i> , 606(7912):129–136.	843
791		844
792		845
793		846
794		847
795	Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. Llmlingua: Compressing prompts for accelerated inference of large language models. <i>arXiv preprint arXiv:2310.05736</i> .	848
796		849
797		850
798	William Jones, Predrag Klasnja, Andrea Civan, and Michael L Adcock. 2008. The personal project planner: Planning to organize personal information. In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems</i> , pages 681–684.	851
799		852
800		853
801		854
802		
803	Victor Kaptelinin. 2003. Umea: translating interaction histories into project contexts. In <i>Proceedings of the SIGCHI conference on Human factors in computing systems</i> , pages 353–360.	855
804		856
805		857
806		858
807	Andruid Kerne, Andrew M Webb, Steven M Smith, Rhema Linder, Nic Lupfer, Yin Qu, Jon Moeller, and Sashikanth Damaraju. 2014. Using metrics of curation to evaluate information-based ideation. <i>ACM Transactions on Computer-Human Interaction (ToCHI)</i> , 21(3):1–48.	859
808		860
809		861
810		862
811		
812		
813	Mik Kersten and Gail C Murphy. 2006. Using task context to improve programmer productivity. In <i>Proceedings of the 14th ACM SIGSOFT international symposium on Foundations of software engineering</i> , pages 1–11.	863
814		864
815		865
816		866
817		867
818	Michel Krieger, Emily Margarete Stark, and Scott R Klemmer. 2009. Coordinating tasks on the commons: designing for personal goals, expertise and serendipity. In <i>Proceedings of the SIGCHI Conference on Human Factors in Computing Systems</i> , pages 1485–1494.	868
819		869
820		870
821		871
822		872
823		873
824	Yash Kumar Lal, Li Zhang, Faeze Brahma, Bodhisattwa Prasad Majumder, Peter Clark, and Niket Tandon. 2024. Tailoring with targeted precision: Edit-based agents for open-domain procedure customization. In <i>Findings of the Association for Computational Linguistics ACL 2024</i> , pages 15597–15611.	874
825		875
826		876
	Christine Lee, David Porfirio, Xinyu Jessica Wang, Kevin Zhao, and Bilge Mutlu. 2025. Veriplan: Integrating formal verification and llms into end-user planning. <i>arXiv preprint arXiv:2502.17898</i> .	877
		878
		879
	Brian R Little. 1983. Personal projects: A rationale and method for investigation. <i>Environment and behavior</i> , 15(3):273–309.	880
		881
	Na Liu, Liangyu Chen, Xiaoyu Tian, Wei Zou, Kaijiang Chen, and Ming Cui. 2024. From llm to conversational agent: A memory enhanced architecture with fine-tuning of large language models. <i>URL https://arxiv.org/abs/2401.02777</i> .	882
		883
	Junyu Luo, Weizhi Zhang, Ye Yuan, Yusheng Zhao, Junwei Yang, Yiyang Gu, Bohan Wu, Binqi Chen, Ziyue Qiao, Qingqing Long, et al. 2025. Large language model agent: A survey on methodology, applications and challenges. <i>arXiv preprint arXiv:2503.21460</i> .	884
		885
	Xiao Ma, Swaroop Mishra, Ariel Liu, Sophie Ying Su, Jilin Chen, Chinmay Kulkarni, Heng-Tze Cheng, Quoc Le, and Ed Chi. 2024. Beyond chatbots: Explore llm for structured thoughts and personalized model responses. In <i>Extended Abstracts of the CHI Conference on Human Factors in Computing Systems</i> , pages 1–12.	886
		887
	Chaitanya Malaviya, Joseph Chee Chang, Dan Roth, Mohit Iyyer, Mark Yatskar, and Kyle Lo. 2024. Contextualized evaluations: Taking the guesswork out of language model evaluations. <i>arXiv preprint arXiv:2411.07237</i> .	888
		889
	Kevin J Miller and Sarah Jo C Venditto. 2021. Multi-step planning in the brain. <i>Current Opinion in Behavioral Sciences</i> , 38:29–39.	890
		891
	Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. <i>Advances in Neural Information Processing Systems</i> , 36:8634–8652.	892
		893
	Significant Gravitas. AutoGPT .	894
	Jaime Teevan, Shamsi T Iqbal, Carrie J Cai, Jeffrey P Bigham, Michael S Bernstein, and Elizabeth M Gerber. 2016. Productivity decomposed: Getting big things done with little microtasks. In <i>Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems</i> , pages 3500–3507.	895
		896
	Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2023. Planbench: An extensible benchmark for evaluating large language models on planning and reasoning about change. <i>Advances in Neural Information Processing Systems</i> , 36:38975–38987.	897
		898
		899
		900

882 Tiannan Wang, Meiling Tao, Ruoyu Fang, Huilin
883 Wang, Shuai Wang, Yuchen Eleanor Jiang, and
884 Wangchunshu Zhou. 2024. Ai persona: Towards
885 life-long personalization of llms. *arXiv preprint*
886 *arXiv:2412.13103*.

887 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten
888 Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou,
889 et al. 2022. Chain-of-thought prompting elicits rea-
890 soning in large language models. *Advances in neural*
891 *information processing systems*, 35:24824–24837.

892 Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu,
893 Renze Lou, Yuandong Tian, Yanghua Xiao, and
894 Yu Su. 2024. Travelplanner: A benchmark for real-
895 world planning with language agents. *arXiv preprint*
896 *arXiv:2402.01622*.

897 Xuan Zhang, Yang Deng, Zifeng Ren, See-Kiong Ng,
898 and Tat-Seng Chua. 2024. Ask-before-plan: Proac-
899 tive language agents for real-world planning. *arXiv*
900 *preprint arXiv:2406.12639*.

901 Xuanming Zhang, Rahul Divekar, Rutuja Ubale, and
902 Zhou Yu. 2023. Groundialog: A dataset for repair
903 and grounding in task-oriented spoken dialogues for
904 language learning. In *Proceedings of the 18th Work-*
905 *shop on Innovative Use of NLP for Building Educa-*
906 *tional Applications (BEA 2023)*, pages 300–314.

907 Huaixiu Steven Zheng, Swaroop Mishra, Hugh Zhang,
908 Xinyun Chen, Minmin Chen, Azade Nova, Le Hou,
909 Heng-Tze Cheng, Quoc V Le, Ed H Chi, et al. 2024.
910 Natural plan: Benchmarking llms on natural lan-
911 guage planning. *arXiv preprint arXiv:2406.04520*.

A *JumpStarter* System Walkthrough

This section presents the interface visuals of the *JumpStarter* system. Figure 4 shows an overview of the interface as it generates plans and answer drafts for the goal “*Apply for a PhD in NLP*”. Figure 5 displays the elicitation interface used to gather global context before initiating the planning process. Figure 6 illustrates the subtask tree, providing a structured overview of the user’s goal. Figure 7 shows how the system recommends further decomposition for a selected subtask. Figure 8 depicts context selection for task forking and the resulting forked subtask structure. Finally, Figure 9 presents the interface for generating an answer draft, including both context selection and additional elicitation.

B Details for Technical Evaluation of Task-structured Context Curation

B.1 Study Procedure

During the study, we first introduced the expert participants to the concept of action plans and answer drafts with examples. We instructed the participants that they would be using three versions of the system for their respective goals, from start to finish. The system versions represented the three experimental conditions, which were presented in shuffled order among participants to counterbalance the learning effect (see Table 2 in Appendix B). We demonstrated how to use each version before each participant used it. We asked the participants to generate the subtasks and answer drafts exactly once. They then rated the perceived quality of the subtasks and answer drafts on a seven-point Likert scale, providing a verbal explanation.

B.2 Technical Evaluation Analysis

Our results show that the context selection feature significantly enhances the quality of answer drafts compared to the baseline. Unlike context saving only, which keeps all context in the context window all the time, context selection requires the LLM to explicitly choose relevant context from the available pool. During sessions under this condition, participants often remarked that the generated solutions appeared to take into account what they had input in previous subtasks. This was particularly evident in “summarizing” tasks, where the system could provide a personalized checklist for tasks like applying for scholarships and driver’s licenses, or an overall itinerary for a team event. E5 referred

Participant	Personal goal	Condition order
E1	Apply to a fellowship	(1) → (2) → (3)
E2		(2) → (3) → (1)
E3	Get a driver’s license	(3) → (1) → (2)
E4		(1) → (2) → (3)
E5	Organize a team event	(2) → (3) → (1)
E6		(3) → (1) → (2)

Table 2: Overview of expert participants for the comparative study. Six experts were assigned one of three goals to evaluate under all three conditions, which were presented in shuffled order to avoid biasing the results.

to the event itinerary they got as a “*very useful synthesis of everything I’ve explored.*” In contrast, the context saving only condition tended to produce only general tips for creating the itinerary. As another example, E1 mentioned that the email draft generated with context selection was more personalized than that created without it. E1 stated, “*I like the recommendation letter request email draft it gives me, as it considers much of my background that I saved in the previous ‘update your CV’ task. I did the same thing in the previous round [context dumping] but did not feel it was as effective.*”

In addition to context selection, context elicitation improves the quality of both subtasks and answer drafts, outperforming both *context dumping* and *context filtering* conditions. Participants reported that the elicitation questions posed at the beginning “*provided the right plan to start with.*” (E4) For instance, E2 uploaded the fellowship requirements document as initially suggested by the system, later rating the generated subtasks a perfect 7/7. “*It captured the requirements quite accurately...*” E2 noted. “*The subtasks were precise, fitting the unique aspects of the fellowship I am applying to, which requires only one recommendation letter, though typically more are needed.*” Similarly, elicitation questions about which state to obtain a driver’s license in (for the “get a driver license” goal) and how many people are in the team (for the “organize a team event” goal) both resulted in subtasks that were better tailored to participants’ individual situations. In addition, preference elicitation questions such as “What type of vehicle do you intend to drive?” and “When do you prefer to hold the event? Weekday or weekend? Noon or night?” prompted participants to provide answers as personal context, ultimately resulting in more tailored answer drafts that they rated highly.

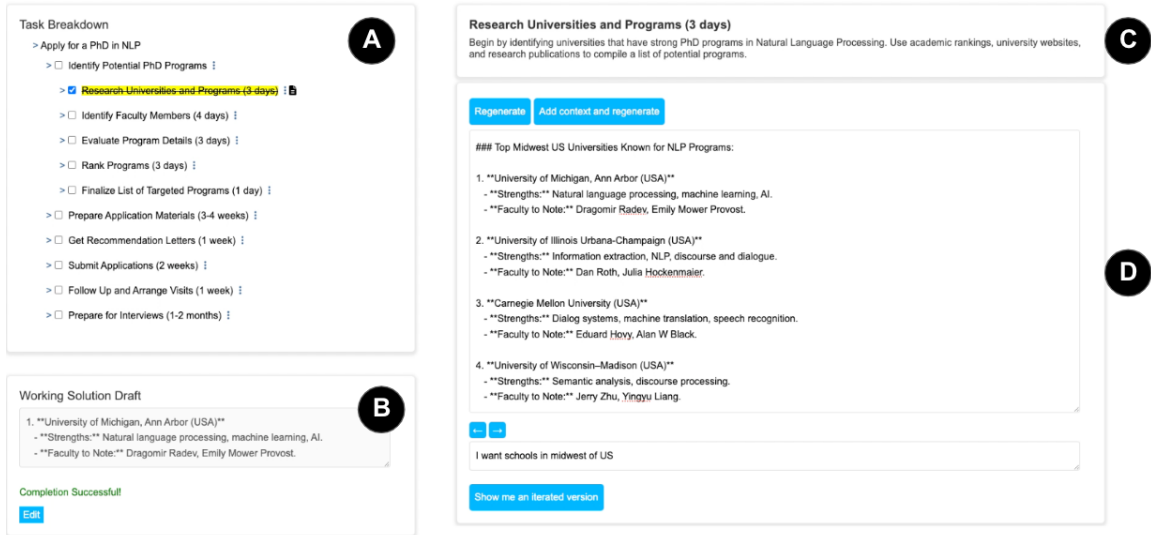


Figure 4: A screenshot of JumpStarter creating plans and answer drafts for the goal *Apply for a PhD in NLP*. (A) Task breakdown is shown as a subtask tree, with the goal being the root node. Subtasks decomposed from the same parent node are shown on the same level. (B) Saving the answer draft. (C) Detailed descriptions of the selected subtask are shown. (D) The answer draft is generated, considering the specification from the user – *"I want schools in midwest of US"*. Users have three options to improve the draft: regenerate, add context and regenerate, and iterate based on users’ new specifications.

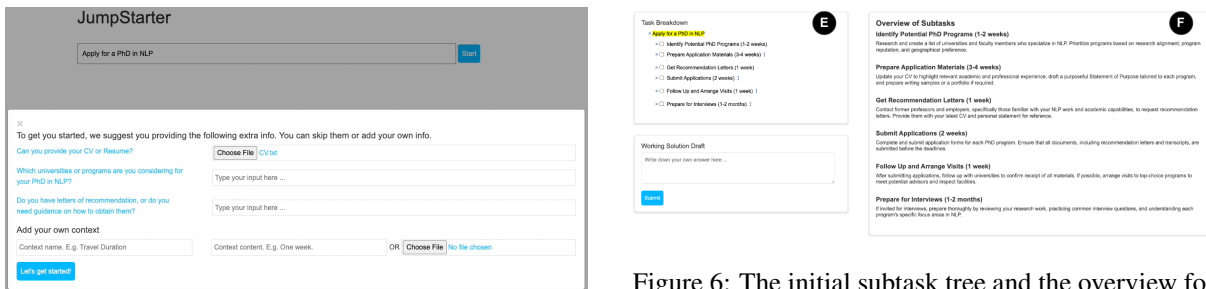


Figure 5: JumpStarter generates questions to elicit context from users to clarify the goal. The user uploads his CV.

C Technical Evaluation of Subtask Detection

C.1 Experiment settings

Subtask Detection aims to determine whether a task should be further decomposed to become executable. Using zero-shot and few-shot prompting as baselines, we evaluated three GPT-4-based prompting strategies for identifying actionable subtasks (see Appendix E for prompt details). First, we applied Chain-of-Thought (CoT) prompting with few-shot examples (Wei et al., 2022). Second, we incorporated the tree level of each task node into the prompt, based on the intuition that higher-level nodes are less likely to be immediately actionable (“Tree”). Third, we explored including the initial

answer drafts in the prompt, allowing the model to assess whether decomposition is necessary based on draft quality (“Draft”). We assessed the accuracy of each method using an expert-labeled test suite. All experiments used a temperature of 1, a maximum token limit of 2048, and top- p set to 1. Results are reported as averages across five runs for each setting.

C.2 Test suite construction

Drawing inspiration from Little (1983), we created a test suite comprising four real-world task scenarios, each representing a distinct aspect of everyday life: 1) Applying for a PhD program (Academic), 2) Obtaining a driver’s license (Practical), 3) Finding a surfing camp (Recreational), and 4) Arranging a

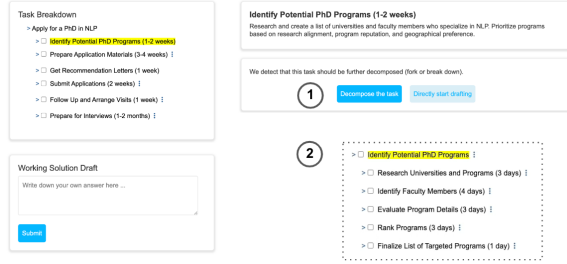
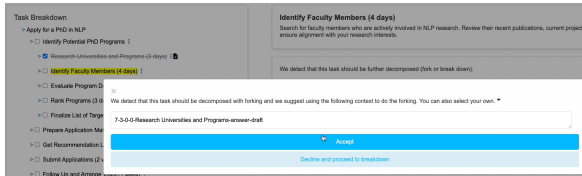
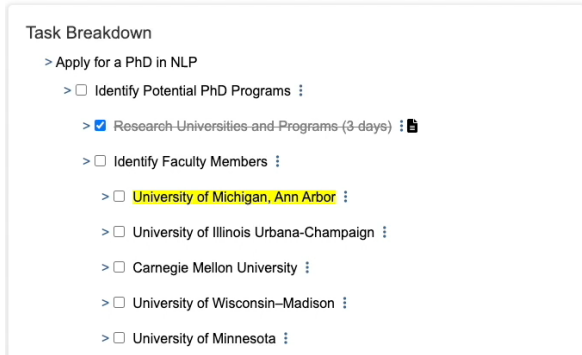


Figure 7: JumpStarter suggests further decomposition for the first subtask *Identify Potential PhD Programs*. (1) John presses the button *Decompose the task*. (2) JumpStarter generates the subtasks for John.



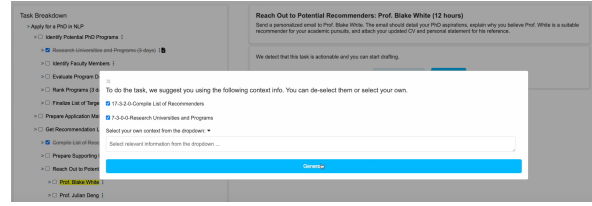
(a) Selecting relevant context for forking.



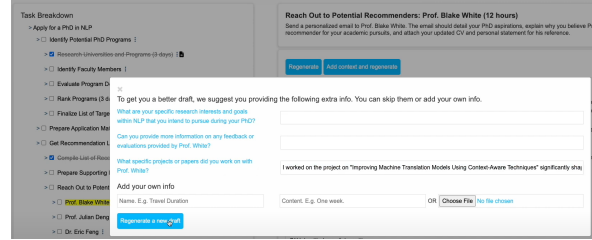
(b) Entity-based task decomposition (forking).

Figure 8: Context selection for forking and the task decomposition after applying forking on *Identify Faculty Members*.

1028 trip abroad (Travel). For each scenario, we created
 1029 ten distinct test cases, resulting in a total of 40 test
 1030 cases. An example test case is shown as follows:
 1031 *Compile a List of Potential Universities: start by*
 1032 *identifying the universities that offer PhD programs*
 1033 *in Natural Language Processing (NLP). Research*
 1034 *and compile a comprehensive list of these univer-*
 1035 *sities. We recruited four expert participants—one*
 1036 *for each task scenario—via a university mailing list*
 1037 *(average age=26.3, two female, two male). These*
 1038 *participants reported having completed the tasks in*
 1039 *the past six months and were compensated \$10 for*
 1040 *their participation. For each test case, the experts*
 1041 *were asked to determine if the current task should*
 1042 *be further decomposed to make it actionable. For*
 1043 *the example test case shown above, its label is "No"*



(a) Context selection for answer draft generation.



(b) Context elicitation for answer draft iteration.

Figure 9: Context selection and elicitation for creating and iterating the answer draft of the subtask *Reach Out to Potential Recommenders: Prof. Blake White*.

Prompting Techniques	Accuracy		Statistics	
	Mean	SD	p	Sig.
Zero-shot	.35	.000		
Few-shot	.58	.040		
+ CoT	.62	.050	.405	-
+ CoT + Tree	.69	.020	.004	**
+ CoT + Draft	.72	.020	.009	**
+ CoT + Tree + Draft	.87	.040	.000	***

Table 3: The technical evaluation results for Subtask Detection comparing different prompting techniques, where the p-values ($-: p > .100$, $+: .050 < p < .100$, $*: p < .050$, $**: p < .010$, $***: p < .001$) are reported. Note that the p-values are computed against the few-shot-only baseline. Few-shot combined with CoT+Tree+Draft achieved the best accuracy.

1044 as it is actionable and does not require further task
 1045 decomposition.

1046 C.3 Results and Findings

1047 As summarized in Table 3, the zero-shot method
 1048 yielded the lowest accuracy of only 0.35, which
 1049 also implies an inherent difficulty in the task itself.
 1050 We observe enhanced performance with few-shot
 1051 prompting (0.58), with accuracy rising even more
 1052 when using the CoT prompting paradigm (0.62).
 1053 Combining CoT few-shot prompting with the task
 1054 node tree levels (0.69), initial solution draft (0.72),
 1055 or both (0.87), all significantly led to enhanced
 1056 performance. Generating the initial solution draft
 1057 introduced a trade-off between latency and accu-

racy, so we opted for the slightly less performant few-shot CoT with tree levels for a better user experience.

D Detailed Design and Results of User Study

D.1 Hypotheses

In the user study, we investigate the following hypotheses:

- **H1:** Compared to ChatGPT, JumpStarter significantly **lowers users' task load (H1)** for mental demand (H1a), temporal demand (H1b), performance (H1c), effort (H1d) and frustration (H1e).
- **H2:** Compared to ChatGPT, JumpStarter significantly **increases users' exploration efficiency (H2)** in terms of the number of plan items explored (H2a) and the number of working solutions drafted (H2b) within the given time.
- **H3:** Compared to ChatGPT, JumpStarter significantly **increases users' satisfaction level with the quality of results (H3)** in terms of plan quality (H3a) and working solution quality (H3b).
- **H4:** Compared to ChatGPT, JumpStarter significantly **increases users' confidence in taking the next steps on their personal projects (H4)**.

D.2 Detailed Results and Findings

D.2.1 Confidence level in taking the next steps on the goal

Participants reported significantly higher confidence in taking the next steps on their personal projects using JumpStarter (mean=5.6, SD=1.07) compared to using ChatGPT (mean=3.9, SD=.99).

Participants reported that ChatGPT helped validate their thoughts with commonsense knowledge and sometimes provided surprising or useful tips. As P7 said, "*Sitting down and planning things out itself is very helpful. I used ChatGPT as a cross-reference, checking to make sure I'm on the right track—thinking similarly to other people. And sometimes answers to low-level tasks covered things I did not really know, which is good.*" However, they reported that they still always felt they might miss something important while using ChatGPT, whereas with JumpStarter, they feel more

secure (P3, P4, P9), as P4 commented, "*I love that I can break things down further if I want, so I don't feel like I miss anything.*"

In addition, JumpStarter can provide very personalized and actionable next steps that greatly increase users' confidence in taking action on their projects. For example, P1, with the goal of starting a side job, shared that "*The schedule JumpStarter helped me generate is very personalized, and I can directly use it to take real action—before, I felt worried about launching this idea as I had very limited time, now I feel like I can really start doing it.*" P10 also liked that JumpStarter provided them with a specific and personalized itinerary for organizing the family reunion—"I like that it summarizes everything I saved in the previous tasks—I can use it in the real world."

D.2.2 Tool preference

8 out of 10 participants reported they prefer to use JumpStarter in the future, compared to ChatGPT. The main reasons given include that JumpStarter can provide more customized responses with less cognitive load. The users feel that they do not have to think hard about what information to provide (P1), are guided by the system (P2), can more easily consume the information (P3) or track the plan (P6), and can get their personal details efficiently organized, framed, and utilized (P8, P9, P10).

The other two participants (P4 and P5) mentioned that their choice depended on how familiar they were with the project they wanted to work on. If it was a topic they already had a clear understanding of, they preferred the chatbot interaction to help them figure out the details. Otherwise, they would prefer to use JumpStarter as it offers more structure.

D.2.3 Improvement feedback

Participants also provided insights on how to improve JumpStarter. The main feedback included "make the subtask outline and task descriptions editable" (P1, P3, P6, P7, P8, P9), "format the suggested answer draft to be easier to read" (P4, P8, P9), "enable users to add or edit context whenever they want" (P3, P7), and "add a synthesis button to summarize what has been explored so far" (P1, P2). P8 also suggested embedding a search agent to collect data and ensure credibility. We discuss limitations and future work further in Section 7.

Category	Factor	JumpStarter		ChatGPT		Statistics		Hypotheses
		Mean	SD	Mean	SD	<i>p</i>	Sig.	
Task load	Mental demand	3.3	1.64	5.3	1.06	.005	**	H1a accepted
	Temporal demand	2.4	.97	4.7	1.83	.012	*	H1b accepted
	Performance	2.7	.95	4.6	1.26	.007	**	H1c accepted
	Effort	2.8	1.69	4.8	1.23	.007	**	H1d accepted
	Frustration	2.1	1.60	4.5	1.90	.009	**	H1e accepted
Exploration efficiency	Plan items explored	4.5	1.43	2.9	.74	.002	**	H2a accepted
	Answer drafts generated	5.3	1.89	3.3	1.16	.047	*	H2b accepted
Satisfaction	Plan quality	5.9	.88	3.3	1.49	.004	**	H3a accepted
	Answer draft quality	5.8	.92	4.7	1.25	.017	*	H3b accepted
Confidence		5.6	1.07	3.9	.99	.007	**	H4 accepted

Table 4: The statistical test results comparing JumpStarter with ChatGPT, where the p-values ($-: p > .100$, $+: .050 < p < .100$, $*: p < .050$, $**: p < .010$, $***: p < .001$) are reported.

D.3 Extended Discussion about User Study

JumpStarter demonstrates how adaptive personal context curation—through context elicitation, condensation, and reuse—can enhance human-AI collaborative planning by improving both plan quality and user experience. Unlike ChatGPT’s user-driven “pull” model, *JumpStarter*’s “push” approach reduces cognitive load by guiding users with targeted questions and enabling selective context previews, aligning with recognition-over-recall principles from cognitive psychology (Craig and Lockhart, 1972). Users valued this structured interaction but also expressed a desire for more persistent and adaptive memory, particularly as their goals and preferences evolved, highlighting the need for systems that can update context dynamically over time (Wang et al., 2024).

While ChatGPT is designed to retain conversational context, it often failed to apply it meaningfully, occasionally generating contradictory responses. These findings underscore the importance of designing systems that support more seamless and transparent context management. Looking forward, integrating in-situ context capture from tools like email or desktop files (Kaptelinin, 2003; Kersten and Murphy, 2006; Bergman and Norman, 2000) may enable more fluid, real-world human-AI collaboration. *JumpStarter* offers a design paradigm for building systems that better align LLM capabilities with users’ evolving needs in complex, long-term planning workflows.

E LLM Prompts

E.1 Context Elicitation

E.1.1 Goal Initialization and Global Context

Figure 10 details the prompts used for eliciting global context for the goal initialization process introduced in Section 3.3.2.

System prompt

You are a helpful assistant in generating at most three questions to elicit more context from the user in order to accomplish a task the user is involved. You will be given the main purpose of the task, and the current context history from the user. You need to judge if the current context information is enough to do this task. If not, what existing doc do you think they already have and can be provided as the context document for this task? Alternatively, what existing most important extra information do you think they should provide as the context information for this task? Please try your best to start by asking the potential existing doc first, and then ask for the potential info. Do not ask the question that can be possibly answered from the suggested doc in the first question. You can also just ask for the potential info if no doc is needed from the user. Please directly generate the questions for users to answer, provide the reason for the question, identify if it is DOC or INFO, and provide a name for the question. Otherwise, output "Ready".

Here are two examples:

Input: My user has a main purpose: Apply for a drivers license by 12/15/2024. The current context history from the user is empty.

Output:

1. Reason: Driver’s license requirements vary significantly depending on the location. Knowing the specific state or country would allow for tailored advice regarding local rules, tests, and documentation required -> Question: Which state or country are you applying for? -> Type: INFO -> Name: State or Country
2. Reason: Different age groups may have different requirements or steps in the licensing process. For example, minors often have to go through graduated license programs. -> Question: How old are you -> Type: INFO -> Name: Age of the User

Input: My user has a main purpose: Apply for a PhD program by 12/15/2024. The current context history from the user is empty.

Output:

1. Reason: Curriculum Vitae (CV) or Resume could be helpful as it would likely contain detailed information about their educational background and any research experiences or academic achievements, which are critical for applying to PhD programs. -> Question: What CV or Resume can you provide for the PhD application? -> Type: DOC -> Name: CV or Resume

Now, start prediction:

User Prompt

Input: My user has a main purpose: {task_input}. The current context history from the user is empty.

Figure 10: Prompts for Context Elicitation for Goal Input.

E.1.2 Answer Draft Creation and Iteration

Figure 11 presents the prompts used to generate an answer draft for each subtask.

My user has a main purpose: {main_purpose}. My user is working on the task {task_name}: {task_description}. The current context history from the user is {context_history}. Please judge if the current context information is enough to do this task. If not, what existing most important extra information do you think they should provide as the context information for this task? Please directly generate questions for users to answer as extra info. Otherwise, output "Ready" Format the response like this: 1. <question 1> : <reason for asking question 1> -> title of question 1
 2. <question 2> : <reason for asking question 2> -> title of question 2
 3. <question 3> : <reason for asking question 3> -> title of question 3

Figure 11: Prompts for Context Elicitation for Answer Draft Creation.

E.2 Context Selection

E.2.1 Answer Draft Generation

Figure 12 specifies the prompts used to select relevant context for generating an answer draft.

System prompt

Given the user's main purpose and the task they are working on, select the most relevant context keys from the current context history that can be used to draft good responses for the user to complete the task. Please also provide explanations. In The current context history is shown as one or more key-value pairs. Please select only the keys from the 'key' part of the context history. Do not select the keys from the 'value' part of the context history In Format the response like this: number. <context_keys>: <reasons for selecting context_keys>. Replace the context_keys with the actual keys as shown in the context history. Please directly give the answers and do not provide extra summarization sentences at the end.

User Prompt

My user has a main purpose: {main_purpose}. My user is working on the task {task_name}: {task_description}. Here is the current context history in JSON format (with 'key': 'value' pairs) from the user: {context_history}

Figure 12: Prompts for Context Selection for Answer Draft Creation.

E.2.2 Task Forking

The prompt to conduct the task forking is as follows:

- *My user has a main purpose: {main_purpose}. My user is working on the task {task_name}: {task_description}. My user needs to break down the task into sub-tasks. Here is the current context history from the user: {context_history}. Please select the most relevant context key from the current context history that can be used to better decompose the current task into several sub-tasks for the user to get started. Do not help the user to break down the task. Please also provide explanations. Format the response like this: <context_key>: <reasons>. Replace the context_key with the actual key in the context history.*

E.3 Task Decomposition

This section presents the prompts used for task decomposition, organized into three core components: Subtask Generation (Figure 13), Subtask Detection (Figure 14), and Task Forking (Figure 15).

My user has a main purpose: {main_purpose}. Please consider the following context information from my user: {user_context}

Please break down the task below into three to six manageable subtasks: {current task} The existing step structure is shown as follows: {existing tree step structure}.

Please directly give the response that fills in the current subtask: {current task} in the provided task structure.

Format the response like this: 1. [Duration for subtask1] {subtask1 title}: {subtask1 detailed description} 2. [Duration for subtask2] {subtask2 title}: {subtask2 detailed description} 3. [Duration for subtask3] {subtask3 title}: {subtask3 detailed description}. Please specify the duration for each subtask in terms of days, weeks or months. For example, [1 week], [2-4 weeks], [1 month], and [1-2 months]. Please do not include other texts for duration such as [Ongoing]. Please do not include "*" in the subtask title. Please directly give the response and do not start with "(current subtask title):"

Figure 13: The prompt for Subtask Generation.

System prompt

You are a useful assistance to detect if the current task needs to be further decomposed if it is not actionable and the primary goal of the task can not be viewed as a singular, distinctive deliverable. Specifically, an overall goal is decomposed into a tree of subtasks. Each of these tasks is positioned on the tree structure and is associated with a level. Based on the user prompt, and the current level of the task node on the tree, please output Yes if it needs to be decomposed; No otherwise meaning it is actionable and does not require task decomposition. Please also provide explanations for your choice.

Here are some examples:

User: My user is working on the task Research on Prospective Companies and Positions: Conduct a comprehensive search on potential companies and specific research scientist internship positions in the NLP field. Understand what each role entails, identify skill requirements, and evaluate how they align with your research interests. The current node level of the task is 1. My user needs to know if the current task is specific and actionable

Reason: This task needs to be further decomposed as it is positioned on the first level of the tree and it involves more than one deliverables: search on companies and search on positions. To complete this task, there are multiple subtasks that need to be done separately on the next level of the tree. These include conducting a comprehensive search on potential companies, searching for specific research scientist internship positions, an analysis of what each role entails, identifying skill requirements, and evaluation of alignment with the user's research interests.

Answer: Yes

User: My user is working on the task Identify Potential Universities: Create a list of universities that offer PhD programs in HCI. The selection can be based on factors such as reputation, HCI research focus, published HCI research papers, faculty expertise etc. The current node level of the task is 2. My user needs to know if the current task is specific and actionable.

Reason: This task does not need to be further decomposed as it is positioned on the second level of the tree and it just involves one deliverable: create a list of schools that offer PhD programs in HCI. Although it may require several steps to create the list, the end goal of this task is to get a list. Therefore the task is actionable.

Answer: No

User: My user is working on the task Identify Required Documents: Research and confirm all the necessary documents required for the non-driver ID application, ensuring to list all forms of acceptable proofs such as a birth certificate or passport for identity, Social Security Card or W-2 form for Social Security number, and utility bills or lease agreement for proof of residency. The current node level of the task is 1. My user needs to know if the current task is specific and actionable.

Reason: Although the task is positioned on the first level of the tree, the task does not need to be further decomposed as the primary goal of the task - identify necessary documents for the non-driver ID application - can be viewed as a singular, cohesive deliverable. Despite involving various types of documents, the task is focused on compiling a comprehensive list, which makes it actionable as a single unit. The distinction lies in the focus on gathering all necessary documentation, a clear and direct objective.

Answer: No

Now, let's start prediction:

User Prompt

User: My user is working on the task {task_title}: {task_descriptions}. The current node level of the task is {tree_level}. My user needs to know if the current task needs to be decomposed

Figure 14: The prompt for Subtask Detection.

E.4 Answer Draft Creation

The prompt to generate the answer drafts is shown as follows:

- *My user has a main purpose: {main_purpose}. Please consider the following context information from my user: {user_context}. My user needs help with the current task {current task}: {task description}*

E.5 Prompts for Technical Evaluation of Subtask Detection

E.5.1 Zero-shot Prompting

The prompt for zero-shot for the task of subtask detection is demonstrated below:

- *System prompt: You are a useful assistance to detect if the current task needs to be further decomposed if it is not actionable and the primary goal of the task can not be viewed as a singular, distinctive deliverable. Based on the user prompt, please output Yes if it needs to be decomposed; No otherwise meaning it is*

System prompt

Given the queried task, determine if a "for" loop is needed to complete the task. You will be given a question Q. Please provide the reasoning and then respond with "Yes" or "No". Here are some examples:

Q: Research the specific HCI PhD programs at each university from the initial list. Focus on aspects such as program curriculum, research opportunities, faculty expertise, and available resources.

Reason: The task requires a "for" loop to complete as there already exists an initial list of entities (i.e. universities) to research. Specifically, the goal of this task is to research the program curriculum, research opportunities, faculty expertise, and available resources for each university from the initial list. It is not possible to complete the task directly without a "for" loop.
A: Yes

Q: Make a list of potential recommenders including former supervisors, academic advisors, and professors who are familiar with your academic and research abilities.

Reason: The task does not require a "for" loop to complete as there does not exist a list of potential recommenders. The goal of this task therefore is to construct the list of recommenders based on certain criteria.
A: No

Q: Reach out to the individuals on your list via email or phone, providing them with the necessary documents and details about the HCI programs, and formally request their letters of recommendation

Reason: This task needs a "for" loop to complete as you have already obtained your list of individuals. You should reach out to each entity (i.e. individuals) on the list to complete the task.
A: Yes

Q: Gather information on different universities offering PhD programs in Human-Computer Interaction. Create an initial list based on general information such as program recognition, location, and basic offerings.

Reason: The task does not require a "for" loop to complete as there does not exist a list of universities offering PhD programs in HCI that can be used to iterate on.
A: No

User Prompt

Q: {task_description}

Figure 15: The prompt for Task Forking.

actionable and does not require task decomposition.

- User Prompt: *My user is working on the task {task title}: {task description}. My user needs to know if the current task needs to be decomposed.*

E.5.2 Few-shot Prompting

The prompt for few-shot-only prompting is shown in Figure 16. Note that we used three in-context examples in the prompt.

E.5.3 Few-shot + CoT

We constructed the prompt in a Chain-of-Thought fashion, where GPT-4 is instructed to first generate the reasoning and then the answer. The prompt is shown in Figure 17.

E.5.4 Few-shot + CoT + Draft

We experimented with incorporating both CoT and the initial working solution draft into the prompt. The system prompt is shown in Figure 18. For the user prompt, before detecting subtasks, we first generated the initial working solution draft for the current task. The user prompt is shown below:

- *My user is working on the task {task title}: {task description}. The GPT response to the task is: {Draft}. My user needs to know if the current task is specific and actionable.*

System prompt

You are a useful assistance to detect if the current task needs to be further decomposed if it is not actionable and the primary goal of the task can not be viewed as a singular, distinctive deliverable. Based on the user prompt, please output Yes if it needs to be decomposed; No otherwise meaning it is actionable and does not require task decomposition

Here are some examples:

User: My user is working on the task Research on Prospective Companies and Positions: Conduct a comprehensive search on potential companies and specific research scientist internship positions in the NLP field. Understand what each role entails, identify skill requirements, and evaluate how they align with your research interests. My user needs to know if the current task is specific and actionable
Answer: Yes

User: My user is working on the task Identify Potential Universities: Create a list of universities that offer PhD programs in HCI. The selection can be based on factors such as reputation, HCI research focus, published HCI research papers, faculty expertise etc. My user needs to know if the current task is specific and actionable.
Answer: No

User: My user is working on the task Identify Required Documents: Research and confirm all the necessary documents required for the non-driver ID application, ensuring to list all forms of acceptable proofs such as a birth certificate or passport for identity, Social Security Card or W-2 form for Social Security number, and utility bills or lease agreement for proof of residency. My user needs to know if the current task is specific and actionable.
Answer: No

Now, let's start prediction:

User prompt

My user is working on the task {task title}: {task description}. My user needs to know if the current task needs to be decomposed.

Figure 16: The few-shot-only prompt for Subtask Detection.

E.5.5 Few-shot + CoT + Tree + Draft

To construct the system prompt for this setting, we incorporate the tree level of each task into the prompt. The prompt is shown in Figure 19. Additionally, for the current task at hand, its tree-level information is also presented in the user prompt, as shown below:

- *My user is working on the task {task title}: {task description}. The current node level of the task is {level}. The GPT response to the task is: {Draft}. My user needs to know if the current task is specific and actionable.*

System prompt

You are a useful assistance to detect if the current task needs to be further decomposed if it is not actionable and the primary goal of the task can not be viewed as a singular, distinctive deliverable. Based on the user prompt, please output Yes if it needs to be decomposed; No otherwise meaning it is actionable and does not require task decomposition. Please also provide explanations for your choice.

Here are some examples:

User: My user is working on the task Research on Prospective Companies and Positions: Conduct a comprehensive search on potential companies and specific research scientist internship positions in the NLP field. Understand what each role entails, identify skill requirements, and evaluate how they align with your research interests. My user needs to know if the current task is specific and actionable

Reason: this task needs to be further decomposed as it involves more than one deliverables: search on companies and search on positions. To complete this task, there are multiple subtasks that need to be done separately. These include conducting a comprehensive search on potential companies, searching for specific research scientist internship positions, an analysis of what each role entails, identifying skill requirements, and evaluation of alignment with the user's research interests.

Answer: Yes

User: My user is working on the task Identify Potential Universities: Create a list of universities that offer PhD programs in HCI. The selection can be based on factors such as reputation, HCI research focus, published HCI research papers, faculty expertise etc. My user needs to know if the current task is specific and actionable.

Reason: This task does not need to be further decomposed as it just involves one deliverable: create a list of schools that offer PhD programs in HCI. Although it may require several steps to create the list, the end goal of this task is to get a list. Therefore the task is actionable.

Answer: No

User: My user is working on the task Identify Required Documents: Research and confirm all the necessary documents required for the non-driver ID application, ensuring to list all forms of acceptable proofs such as a birth certificate or passport for identity, Social Security Card or W-2 form for Social Security number, and utility bills or lease agreement for proof of residency. My user needs to know if the current task is specific and actionable.

Reason: The primary goal of the task is to identify necessary documents for the non-driver ID application, which can be viewed as a singular, cohesive deliverable. Despite involving various types of documents, the task is focused on compiling a comprehensive list, which makes it actionable as a single unit. The distinction lies in the focus on gathering all necessary documentation, a clear and direct objective.

Answer: No

Now, let's start prediction:

User prompt

My user is working on the task (task title): {task description}. My user needs to know if the current task needs to be decomposed.

Figure 17: The prompt for few-shot + CoT for Subtask Detection.

System prompt

You are a useful assistance to detect if the current task needs to be further decomposed if it is not actionable and the primary goal of the task can not be viewed as a singular, distinctive deliverable. Based on the user prompt, and the GPT response, please output Yes if it needs to be decomposed; No otherwise meaning it is actionable and does not require task decomposition. Please also provide explanations for your choice.

Here are some examples:

User: My user is working on the task Research on Prospective Companies and Positions: Conduct a comprehensive search on potential companies and specific research scientist internship positions in the NLP field. Understand what each role entails, identify skill requirements, and evaluate how they align with your research interests. The GPT response to the task is:
Sure, I will list some promising companies that offer research scientist internship positions in Natural Language Processing (NLP) and briefly describe what each role entails:

- Google AI**:**
 - **Role:** As a research intern, you can work in areas such as machine intelligence, machine perception, natural language processing, and quantum computing. Google typically looks for research experience, ability to design and execute a research agenda, and coding skills.
- Facebook AI Research (FAIR)**:**
 - **Role:** Natural Language Processing Internship roles at FAIR involve developing novel techniques for NLP tasks which include generating human-like text. They require experience with machine learning, AI, deep learning, dialogue management, and computer vision.
- Amazon Alexa AI**:**
 - **Role:** Interns in this role generally work on language understanding capabilities of Alexa, advancing state-of-the-art on tasks like entity resolution or text classification. Knowledge in machine learning, especially deep learning for NLP, is key for this role.
- OpenAI**:**
 - **Role:** Research internships at OpenAI are highly selective. Interns are expected to have published in machine learning or AI and can spend a significant portion of their time on original research.
- IBM Watson**:**
 - **Role:** Watson Interns can work on various AI/NLP tasks from fact extraction, relation extraction, disambiguation, to summarization, based on the team you join. A strong AI/ML foundation and coding skills are required.
- Apple AI/ML**:**
 - **Role:** Apple offers internships where you can work on Siri's language understanding capabilities. They typically require machine learning knowledge and experience, especially in deep learning.
- Microsoft Research**:**
 - **Role:** Interns in this role could work on a variety of NLP challenges, including developing algorithms that enable, for example, conversational interactions with AI. Strong machine learning knowledge, particularly in deep learning, is expected.
- Baidu Research**:**
 - **Role:** Interns can work on various novel NLP tasks such as language understanding, language generation, and more. Baidu looks for individuals with robust knowledge in machine learning, deep learning, and coding abilities.
- NVIDIA - AI Research**:**
 - **Role:** Interns have an opportunity to work on projects in the realms of machine learning, deep learning, computer vision, and natural language processing. They seek candidates with strong ML and DL backgrounds.

When selecting a specific position, I recommend considering the following:

- Does the role align with your career objectives?
- Does the role play to your strengths or offer opportunities to develop in areas of interest?
- Does the role match your skill set? If not, can you acquire those skills?

Remember, a role can be an opportunity to learn and grow, so don't be deterred if you lack a few skills. If the role excites you, go for it!

My user needs to know if the current task is specific and actionable

Reason: this task needs to be further decomposed as the GPT response is not good enough to help user complete the task. It involves more than one deliverables: search on companies and search on positions. To complete this task, there are multiple subtasks that need to be done separately. These include conducting a comprehensive search on potential companies, searching for specific research scientist internship positions, an analysis of what each role entails, identifying skill requirements, and evaluation of alignment with the user's research interests.

Answer: Yes

User: My user is working on the task Identify Potential Universities: Create a list of universities that offer PhD programs in HCI. The selection can be based on factors such as reputation, HCI research focus, published HCI research papers, faculty expertise etc. The GPT response to the task is:

Here is a list of universities that offer PhD programs in HCI:

- Stanford University, USA
- Massachusetts Institute of Technology (MIT), USA
- Carnegie Mellon University (CMU), USA
- University of California, Berkeley, USA
- University of Washington, USA
- University of Michigan, USA
- University of California, San Diego, USA
- Cornell University, USA
- University of Toronto, Canada
- University of College London, UK

My user needs to know if the current task is specific and actionable.

Reason: The GPT response is good enough and it provides direct answer to complete the task. Additionally, this task does not need to be further decomposed as it just involves one deliverable: create a list of schools that offer PhD programs in HCI. Although it may require several steps to create the list, the end goal of this task is to get a list. Therefore the task is actionable.

Answer: No

User: My user is working on the task Identify Required Documents: Research and confirm all the necessary documents required for the non-driver ID application, ensuring to list all forms of acceptable proofs such as a birth certificate or passport for identity, Social Security Card or W-2 form for Social Security number, and utility bills or lease agreement for proof of residency. The GPT response to the task is:

To apply for a non-driver ID in [redacted] you need to provide several documents to prove your identity, Social Security Number (SSN), and residency. Below, I have listed the acceptable documents for each category based on the latest guidelines provided by the [redacted] Department of Motor Vehicles (DMV).

****1. Proof of Identity and Date of Birth (must total 6 points according to [redacted] DMV's point system):****

- U.S. passport or U.S. passport card (valid or expired within the past five years) - 4 points
- Birth certificate (U.S., Puerto Rican, Canadian, or Mexican) - 2 points
- Employment Authorization Card (I-766) - 3 points
- Permanent Resident Card (I-551) - 3 points
- Foreign passport with valid U.S. Visa and I-94 - 3 points

***Note:** At least one document must show your date of birth and at least one document must have your full Social Security Number.*

****2. Proof of Social Security Number:****

- Social Security card
- W-2 form
- SSA-1099 form
- Non-SSA-1099 form
- Pay stub with your name and SSN on it

****3. Proof of [redacted] Residency (at least two documents):****

- State or federal tax return
- Utility bill (must be dated within the last 90 days)
- Lease agreement or mortgage statement
- Bank statement (must be dated within the last 90 days)
- Any mail from any government agency (dated within the last 90 days)

****Additional Notes:****

- All documents must be originals or certified copies.
- If your name has changed due to marriage, divorce, or court order, you will need to provide documents that link your current name to your former name.
- You must also complete an "Application for Permit, Driver License or Non-Driver ID Card" (Form MV-44).

Before going to the DMV, it's recommended to use the [redacted] DMV Document Guide which provides a personalized checklist of the required documents based on the specifics of your application. This can be found on the official [redacted] DMV website.

Make sure to double-check with the DMV for the latest requirements or any specific circumstances that may affect your application to ensure a smooth process.

My user needs to know if the current task is specific and actionable.

Reason: The task does not need to be further decomposed for two reasons. The first reason is that the GPT response provides direct answer to the task and generates a list of all forms of acceptable proofs for the documents required to get a non-driver ID. The second reason is that the primary goal of the task - identify necessary documents for the non-driver ID application - can be viewed as a singular, cohesive deliverable. Despite involving various types of documents, the task is focused on compiling a comprehensive list, which makes it actionable as a single unit. The distinction lies in the focus on gathering all necessary documentation, a clear and direct objective.

Answer: No

Now, let's start prediction:

Figure 18: The prompt for few-shot + CoT + Draft for Subtask Detection.

You are a useful assistance to detect if the current task needs to be further decomposed if it is not actionable and the primary goal of the task can not be viewed as a singular, distinctive deliverable. Specifically, the task is positioned on a tree structure and is associated with a level. Based on the user prompt, the current node level of the task on the tree, and the GPT response, please output Yes if it needs to be decomposed; No otherwise meaning it is actionable and does not require task decomposition. Please also provide explanations for your choice.

Here are some examples:

User: My user is working on the task Research on Prospective Companies and Positions: Conduct a comprehensive search on potential companies and specific research scientist internship positions in the NLP field. Understand what each role entails, identify skill requirements, and evaluate how they align with your research interests. The current node level of the task is 1. The GPT response to the task is: Sure, I will list some promising companies that offer research scientist internship positions in Natural Language Processing (NLP) and briefly describe what each role entails:

1. **Google AI**:**
 - **Role:** As a research intern, you can work in areas such as machine intelligence, machine perception, natural language processing, and quantum computing. Google typically looks for research experience, ability to design and execute a research agenda, and coding skills.
2. **Facebook AI Research (FAIR)**:**
 - **Role:** Natural Language Processing Internship roles at FAIR involve developing novel techniques for NLP tasks which include generating human-like text. They require experience with machine learning, AI, deep learning, dialogue management, and computer vision.
3. **Amazon Alexa AI**:**
 - **Role:** Interns in this role generally work on language understanding capabilities of Alexa, advancing state-of-the-art on tasks like entity resolution or text classification. Knowledge in machine learning, especially deep learning for NLP, is key for this role.
4. **OpenAI**:**
 - **Role:** Research internships at OpenAI are highly selective. Interns are expected to have published in machine learning or AI and can spend a significant portion of their time on original research.
5. **IBM Watson**:**
 - **Role:** Watson Interns can work on various AI/NLP tasks from fact extraction, relation extraction, disambiguation, to summarization, based on the team you join. A strong AI/ML foundation and coding skills are required.
6. **Apple AI/ML**:**
 - **Role:** Apple offers internships where you can work on Siri's language understanding capabilities. They typically require machine learning knowledge and experience, especially in deep learning.
7. **Microsoft Research**:**
 - **Role:** Interns in this role could work on a variety of NLP challenges, including developing algorithms that enable, for example, conversational interactions with AI. Strong machine learning knowledge, particularly in deep learning, is expected.
8. **Baidu Research**:**
 - **Role:** Interns can work on various novel NLP tasks such as language understanding, language generation, and more. Baidu looks for individuals with robust knowledge in machine learning, deep learning, and coding abilities.
9. **NVIDIA - AI Research**:**
 - **Role:** Interns have an opportunity to work on projects in the realms of machine learning, deep learning, computer vision, and natural language processing. They seek candidates with strong ML and DL backgrounds.

When selecting a specific position, I recommend considering the following:

- Does the role align with your career objectives?
- Does the role play to your strengths or offer opportunities to develop in areas of interest?
- Does the role match your skill set? If not, can you acquire those skills?

Remember, a role can be an opportunity to learn and grow, so don't be deterred if you lack a few skills. If the role excites you, go for it!

My user needs to know if the current task is specific and actionable

Reason: this task needs to be further decomposed as the GPT response is not good enough to help user complete the task. Moreover, it is positioned on the first level of the tree and it involves more than one deliverables: search on companies and search on positions. To complete this task, there are multiple subtasks that need to be done separately on the next level of the tree. These include conducting a comprehensive search on potential companies, searching for specific research scientist internship positions, an analysis of what each role entails, identifying skill requirements, and evaluation of alignment with the user's research interests.

Answer: Yes

User: My user is working on the task Identify Potential Universities: Create a list of universities that offer PhD programs in HCI. The selection can be based on factors such as reputation, HCI research focus, published HCI research papers, faculty expertise etc. The current node level of the task is 2. The GPT response to the task is:

Here is a list of universities that offer PhD programs in HCI:

1. Stanford University, USA
2. Massachusetts Institute of Technology (MIT), USA
3. Carnegie Mellon University (CMU), USA
4. University of California, Berkeley, USA
5. University of Washington, USA
6. University of Michigan, USA
7. University of California, San Diego, USA
8. Cornell University, USA
9. University of Toronto, Canada
10. University of College London, UK

My user needs to know if the current task is specific and actionable.

Reason: The GPT response is good enough and it provides direct answer to complete the task. Additionally, this task does not need to be further decomposed as it is positioned on the second level of the tree and it just involves one deliverable: create a list of schools that offer PhD programs in HCI. Although it may require several steps to create the list, the end goal of this task is to get a list. Therefore the task is actionable.

Answer: No

User: My user is working on the task Identify Required Documents: Research and confirm all the necessary documents required for the non-driver ID application, ensuring to list all forms of acceptable proofs such as a birth certificate or passport for identity, Social Security Card or W-2 form for Social Security number, and utility bills or lease agreement for proof of residency. The current node level of the task is 1. The GPT response to the task is: To apply for a non-driver ID in [redacted] you need to provide several documents to prove your identity, Social Security Number (SSN), and residency. Below, I have listed the acceptable documents for each category based on the latest guidelines provided by the [redacted] Department of Motor Vehicles (DMV).

****1. Proof of Identity and Date of Birth (must total 6 points according to [redacted] DMV's point system):****

- U.S. passport or U.S. passport card (valid or expired within the past five years) - 4 points
- Birth certificate (U.S., Puerto Rican, Canadian, or Mexican) - 2 points
- Employment Authorization Card (I-766) - 3 points
- Permanent Resident Card (I-551) - 3 points
- Foreign passport with valid U.S. Visa and I-94 - 3 points

***Note:** At least one document must show your date of birth and at least one document must have your full Social Security Number.*

****2. Proof of Social Security Number:****

- Social Security card
- W-2 form
- SSA-1099 form
- Non-SSA-1099 form
- Pay stub with your name and SSN on it

****3. Proof of [redacted] Residency (at least two documents):****

- State or federal tax return
- Utility bill (must be dated within the last 90 days)
- Lease agreement or mortgage statement
- Bank statement (must be dated within the last 90 days)
- Any mail from any government agency (dated within the last 90 days)

****Additional Notes:****

- All documents must be originals or certified copies.
- If your name has changed due to marriage, divorce, or court order, you will need to provide documents that link your current name to your former name.
- You must also complete an "Application for Permit, Driver License or Non-Driver ID Card" (Form MV-44).

Before going to the DMV, it's recommended to use the [redacted] DMV Document Guide which provides a personalized checklist of the required documents based on the specifics of your application. This can be found on the official [redacted] DMV website.

Make sure to double-check with the DMV for the latest requirements or any specific circumstances that may affect your application to ensure a smooth process.

My user needs to know if the current task is specific and actionable.

Reason: Although the task is positioned on the first level of the tree, the task does not need to be further decomposed for two reasons. The first reason is that the GPT response provides direct answer to the task and generates a list of all forms of acceptable proofs for the documents required to get a non-driver ID. The second reason is that the primary goal of the task - identify necessary documents for the non-driver ID application - can be viewed as a singular, cohesive deliverable. Despite involving various types of documents, the task is focused on compiling a comprehensive list, which makes it actionable as a single unit. The distinction lies in the focus on gathering all necessary documentation, a clear and direct objective.

Answer: No

Now, let's start prediction:

Figure 19: The prompt for few-shot + CoT + Tree + Draft for Subtask Detection.