

GUIDE: A Benchmark for Understanding and Assisting Users in Open-Ended GUI Tasks

Saelyne Yang^{1,2} Jaesang Yu¹ Yi-Hao Peng² Kevin Qinghong Lin³ Jae Won Cho⁴
Yale Song⁵ Juho Kim^{1,6}
¹KAIST ²Carnegie Mellon University ³University of Oxford ⁴Konkuk University ⁵Google Inc. ⁶SkillBench

Abstract

Graphical User Interface (GUI) agents have the potential to assist users in interacting with complex software (e.g., PowerPoint, Photoshop). While prior research has primarily focused on automating user actions through clicks and keystrokes, this paradigm overlooks human intention, where users value the ability to explore, iterate, and refine their ideas while maintaining agency. To move beyond automation and toward collaboration, GUI agents must understand what users are doing and why. We introduce **GUIDE** (GUI User Intent Detection Evaluation), a benchmark that evaluates AI models on their ability to perceive user behavior, infer intent, and provide assistance in open-ended GUI tasks. **GUIDE** consists of 67.5 hours of screen recordings from 120 novice user demonstrations with think-aloud narrations, across 10 software. **GUIDE** defines three tasks—(i) Behavior State Detection, (ii) Intent Prediction, and (iii) Help Prediction that test a model’s ability to recognize behavior state, reason about goals, and decide when and how to help. Evaluations across eight state-of-the-art multimodal models reveal that all models struggled, achieving only 44.6% and 55.0% accuracy on behavior state and help prediction. However, providing user context significantly improved the performance, raising help prediction by up to 50.2pp, highlighting the critical role of structured user understanding in effective assistance. Our dataset is available at <https://guide-bench.github.io>.

1. Introduction

Graphical User Interface (GUI) agents hold great promise for supporting users in complex workflows, in mobile [20, 29, 53], web [9, 16, 41, 50, 56], and software application tasks [10, 40, 51]. In creative and analytical tools such as Photoshop or PowerPoint, these agents can automate repetitive subtasks or provide guidance to help users achieve their goals more efficiently. Most existing GUI agents, both in academic research [14, 27, 28, 54] and in commer-

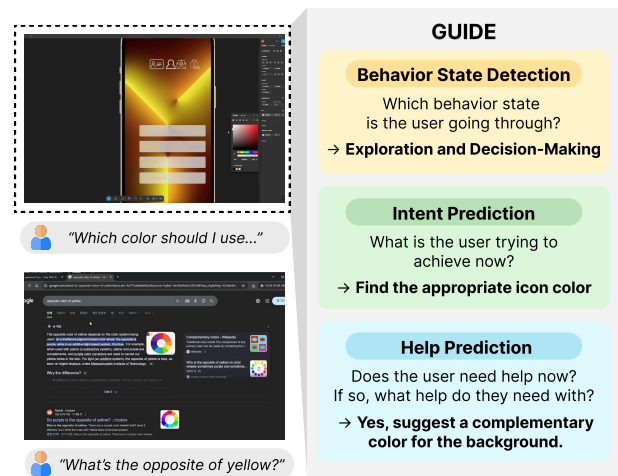


Figure 1. An example of the **GUIDE** benchmark, which jointly models three tasks: Behavior State Detection, Intent Prediction, and Help Prediction, to interpret what the user is doing, aiming to achieve, and whether and what they may need assistance with during open-ended software tasks.

cial services like Microsoft Office Copilot [32] or Figma Make [12], focus on full automation: given a goal, they either execute a sequence of clicks and keystrokes to complete the task or directly generate the desired output.

While this approach offers convenience, it overlooks how people actually work. In real-world open-ended workflows, success is not driven solely by efficiency—user satisfaction plays an equally critical role. Automated agents assume fixed goals, yet users frequently revise their intentions mid-task. For example, a user may reposition an element multiple times before reverting to the original—behavior an automated agent would treat as redundant, but which is essential for forming a preference. Rather than replacing user agency, effective assistance should accelerate exploration while keeping the user in control [22].

Recent work on proactive task assistance takes a more balanced approach [31, 45, 47, 48, 52]. Rather than auto-

Dataset	Domain #	Video #	Video Duration	Video Source	Primary Goal	Evaluation Focus		
						Behavior	Intent	Help
PsTuts [25]	1	-	71.4 h	Instructional Videos	Action Understanding			
VideoWebArena [19]	6	74	3.8 h	Human-Recorded Tutorials	Task Automation			
VideoGUI [27]	11	178	7.1 h	Instructional Videos	Task Automation			✓
UI-Vision [33]	83	450	4.8 h	Experts Performing Tasks	Task Automation			✓
AssistGUI [14]	9	100	<8.3 h	Instructional Videos	Task Automation			✓
WorldGUI [54]	10	611	<30.5 h	Instructional Videos	Task Automation			✓
GUIDE (Ours)	10	120	67.5 h	Novice Users' Demonstrations	Behavior Understanding	✓	✓	✓

Table 1. **Comparison of GUIDE with existing GUI video understanding datasets.** GUIDE differs from existing benchmarks by (i) collecting screen recordings from novice users, (ii) capturing how they naturally behave in open-ended tasks with a focus on behavior understanding, and (iii) evaluating systems based on human user needs rather than task automation.

mate tasks for users, proactive assistants infer a user’s context and intent and deliver timely help. Studies in programming and productivity tools show higher efficiency and satisfaction when a system detects a need and intervenes at the right moment [6, 37, 38, 46]. Yet, the ability to model and track users’ evolving context remains underexplored in current multimodal systems that power GUI agents.

To achieve a truly human-assisting GUI agent, a key ability is to comprehend users’ cognitive context and intentions to provide appropriate support [17]. In real-world scenarios, users rarely articulate their goals or needs explicitly, making it natural for systems to rely primarily on visual cues from the screen. These user actions often carry semantic structure, such as hovering, undoing, or repeatedly opening menus, that signal intent. However, interpretation remains challenging: similar actions may stem from entirely different intents. For example, repeated undo actions might indicate confusion or deliberate refinement. Without deeper reasoning, assistance based solely on surface-level actions can lead to shallow or misaligned responses.

To address this challenge, we present **GUIDE** (**G**UI **U**ser **I**ntent **D**etection **E**valuation), a benchmark designed to evaluate multimodal LLMs (MLLMs) on their ability to understand and assist users in complex software workflows. **GUIDE** introduces a three-stage evaluation framework: (1) *Understanding* the user’s behavioral state to identify their current workflow phase; (2) *Reasoning* about their underlying intentions and goals; and (3) *Assisting* by delivering the appropriate form of help at the right moment.

We collected 67.5 hours of screen recordings from 120 human demonstrations across 10 widely used applications—including Photoshop, Figma, PowerPoint, Premiere Pro, and Excel—covering 40 open-ended tasks designed to elicit natural user behavior. Unlike prior work that primarily targets video understanding from expert-recorded instructional videos on closed-ended tasks [14, 25, 27, 33, 54], our focus is on novice users working on open-ended tasks, with the goal of building collaborative AI sys-

tems that assist users during exploration, trial-and-error, and learning. Observing novice workflows allows us to capture authentic moments of confusion, decision-making, and discovery, offering rich opportunities for AI to provide timely, context-aware support. Each session includes both screen recordings and think-aloud narrations that surface the user’s underlying intentions and cognitive states.

Building on this dataset, we define three-staged benchmark tasks: First, **(i) Behavior State Detection** evaluates whether a model can identify the user’s behavioral state, such as exploration or confusion, based solely on visual cues. To support this, we developed a taxonomy of nine user states reflecting diverse cognitive and behavioral phases in open-ended GUI workflows, grouped into four high-level categories: Planning, Execution, Problem-Solving, and Evaluation (Figure 3). This structure aligns with human cognition and interaction theories [4, 34], while introducing finer distinctions tailored to GUI-based task behavior. Next, **(ii) Intent Prediction** targets inference of the user’s immediate goal—what they are trying to accomplish in the given moment. The final task, **(iii) Help Prediction**, assesses whether a model can determine 1) whether the user needs assistance or not, and if so, 2) what type of help would be most appropriate, such as explaining a feature, suggesting an alternative, or addressing an error. By leveraging both visual screen recordings and accompanying think-aloud narrations, we automatically generated data for each task, which was subsequently verified through human review for accuracy and consistency.

Evaluation across eight state-of-the-art MLLMs reveals that while current models struggle to interpret user behavior and predict underlying intent and help needed—achieving only 44.6% accuracy on behavior state detection and 55.0% on help prediction, performance improves significantly when structured user context is provided. For example, supplying behavioral state and intent information boosted help prediction accuracy by up to 50.2 percent points for the lowest-performing model.

Our results suggest a promising path forward: providing different layers of human-grounded context, such as behavioral cues, inferred goals, and temporal history, can lead to more accurate assistance decisions. Our benchmark provides a foundation for training and evaluating the next generation of context-aware, collaborative GUI agents.

2. Related Work

2.1. Video Understanding meets GUI

Several benchmarks evaluate video understanding in the context of GUI and software workflows. Early work by Li et al. [25] collected Photoshop tutorial videos to understand screencast videos. More recent datasets span multiple applications and tasks. For example, AssistGUI [14] focuses on automating GUI tasks using an actor-critic agent, serving as a benchmark for task-oriented GUI automation. VideoWebArena [19] evaluates long-horizon multimodal agents on web browsing tasks, emphasizing extended video context and web UI interactions. VideoGUI [27] compiles high-quality instructional screen recordings and introduces a hierarchical model for mapping visual observations to GUI actions. UI-Vision [33] provides a fine-grained desktop UI video benchmark with dense annotations for perception and interaction. Lastly, WorldGUI [54] increases task diversity by allowing arbitrary initial interface states for each task, challenging agents to handle varied starting conditions.

These prior benchmarks primarily focus on close-ended tasks with predetermined goals, aiming to replicating expert demonstrations. In contrast, our work targets open-ended GUI workflows with novice users, emphasizing understanding of user intent and context rather than step-by-step replication of actions. This shift toward user-centric evaluation fills a gap not covered by existing GUI video datasets that evaluate task completion or action prediction.

2.2. Collaborative and Proactive Agents

While GUI agents that automate interface operations based on a given goal or instruction can be effective, this fully autonomous approach can conflict with the needs of users in creative or analytical environments, where retaining control and exploring alternatives are essential. To address this, recent research has shifted toward assistive GUI agents that collaborate with users by understanding context and offering timely support. Several works have explored inferring user goals and intent in both web [36] and software environments [3, 13, 55] to better align assistance with user needs. For example, Zhao et al. [55] introduce ProactiveVA, a visual analytics agent that monitors user interactions and leverages LLMs to detect when users may be stuck, providing context-sensitive suggestions or guidance.

Several recent works in the Human-Computer Interaction (HCI) community explore this shift toward collabora-

tion and contextual support. CowPilot [18] proposes a mixed-initiative framework that enables users to share control with an autonomous web navigation agent, improving efficiency while preserving agency. In programming, proactive assistants like Codellaborator [38] and NeedHelp [6] demonstrate how real-time intervention can aid users when well-timed. Studies on software applications [22] show users prefer AI agents that guide them rather than take over entirely, reinforcing the need for transparency and shared control. ProMemAssist [39] further highlights the benefits of modeling user cognition to deliver timely, non-intrusive support. These findings echo broader discussions on autonomy levels [11] and the importance of aligning agent behavior with human preferences [21, 26]. Our work builds on these insights, evaluating how well current multimodal models can perceive a user’s state and intentions in GUI workflow recordings and decide if and how to assist. By situating the evaluation in real user workflows, we aim to push GUI agents toward true user-aware collaboration.

3. GUIDE Benchmark

To develop a benchmark that focuses on understanding and assisting users, we collected demonstrations from novice users. Unlike existing datasets that focus primarily on expert demonstrations or polished instructional videos [14, 25, 27, 30, 33, 44, 54], our dataset captures the authentic challenges and exploratory behaviors that novices exhibit during task completion, serving a crucial role in building collaborative agents. Building on these demonstrations, we propose a suite of tasks designed to evaluate models’ capabilities to understand users and provide effective assistance.

3.1. Video Collection


We collected 120 demonstrations from novice users across 10 applications spanning five categories: Photo Editing (Photoshop, GIMP), Graphic Design (Figma, Canva), Presentation Design (PowerPoint, Google Slides), Video Editing (Premiere Pro, CapCut), and Data Analysis (Google Sheets, Microsoft Excel). For each application, we designed four open-ended tasks aimed at eliciting natural and diverse user behaviors and approaches (Table B2 in supp.).

We chose creative and analytical tools to surface exploratory workflows and variation in problem-solving strategies. Each task was completed by three different users to capture diverse strategies and behaviors. We ensured that each task was flexible enough, while still incorporating elements of challenge. Participants were asked to spend at least 20 minutes per task and meet a few minimal requirements (e.g., inserting a relevant image) to mark it as complete.

We recruited 54 novice users of software from ProLific and our institution. Participants were screened based on self-reported expertise to ensure novice-level familiarity with the features in the target application (Mean: 2.8, SD:

(1) User Behavior State Detection

Which behavior state is the user in?




"I'm trying to see all the effects we can explore.."

- Task Understanding and Preparation
- Frustration
- Ideation and Planning
- Debugging
- Exploration and Decision-Making
- Seeking External Help
- Waiting and Monitoring
- Performing Actions
- Assessment

(2) Intent Prediction

What is the user trying to achieve?




"I'm going to duplicate this and make it smaller."

- A. Change the color of the rectangular shape
- B. Adjust the spacing between the form elements
- C. Create a sign-up button
- D. Create a progress bar

(3) Help Prediction

Does the user need help? If so, what kind?



"I really don't know how to make use of any of these texts.."

- (3-1) Help Need Does the user need help?
 Yes No
- (3-2) Help Content What help is needed?
A. How to color grade the video footage
B. Generate automatic captions
C. Get a guide on how to use text effects
D. Apply a special effect to the video clip

Figure 2. Overview of the three core tasks in the GUIDE benchmark. (1) **User Behavior State Detection** identifies the user’s current behavioral mode (e.g., *Exploration and Decision-Making*). (2) **Intent Prediction** infers what the user is trying to achieve (e.g., *Create a progress bar*). (3) **Help Prediction** determines whether the user needs assistance and, if so, what kind of help is relevant (e.g., *Get a guide on how to use text effects*). Together, these tasks enable a comprehensive understanding of user behavior and assistance needs in software GUI environments. We evaluate MLLMs on their ability to infer these solely from the visual input, without access to the demonstrator’s narration — a setting that closely reflects real-world use.

1.1, Range: 1–5). During the study, participants worked on the assigned task while recording their screen and keyboard/mouse input events. They were also asked to think aloud and record their voice as they carried out the task, verbalizing what they were doing and their thought process.

3.2. Benchmark Tasks

To evaluate a model’s ability to understand user context and deliver appropriate assistance, we design our benchmark as a unified three-stage framework: *Understanding* → *Reasoning* → *Assisting*. These stages progress from interpreting user behavior to inferring intentions and ultimately providing helpful assistance. Each task corresponds to a distinct level of cognitive inference required for a human-assisting GUI agent to effectively support users in open-ended software workflows.

To construct a dataset for task evaluation, we used the Human-AI collaborative method. We first transcribed the think-aloud narration using WhisperX [2], and used the narration as a main source for extracting initial annotations in addition to the video. We employed *Gemini-2.5-Pro* [15] to first create annotations needed for each task, which were then refined by human annotators. Note that we use narration only as an annotation source to capture users’ intentions and mental states. The benchmark evaluates vision-only understanding, testing whether models can infer these

states solely from visual cues, as in real-world settings with limited access to user speech.

3.2.1. User Behavior State Detection

Description. This task evaluates whether a model can interpret the user’s behavioral context directly from visual cues. Models are asked to classify a video segment into one of nine behavior states in our taxonomy (Figure 3), which spans the full range of cognitive and behavioral processes observed in creative and analytical workflows.

We developed the taxonomy through a multi-stage, human-AI collaborative process [24]. First, three authors iteratively created and consolidated an initial taxonomy over five sessions based on observations of online software task videos. Separately, we prompted *Gemini-2.5-Pro* to generate a taxonomy from scratch using our collected video dataset, without providing our initial version. We then augmented the human-generated taxonomy by integrating novel categories identified by the LLM. Finally, the combined taxonomy was validated against the entire video dataset to ensure comprehensive coverage and reorganized into the final set of nine distinct states. Our taxonomy aligns with Norman’s Seven Stages of Action [34], mapping Planning, Execution, and Evaluation to goal formation, action, and outcome assessment, and draws on Bloom’s cognitive hierarchy [4] that captures the shifts between operational

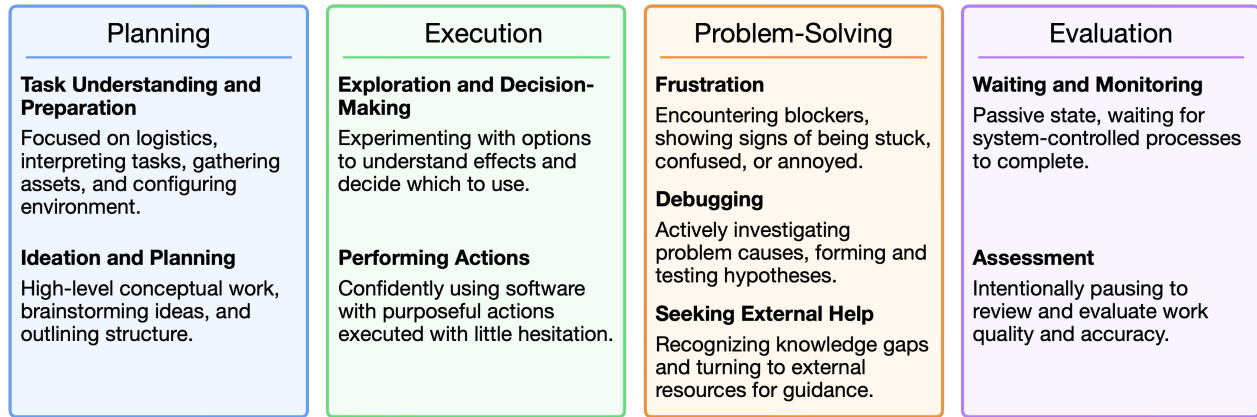


Figure 3. Our proposed taxonomy of user behavior states in GUI-based software tasks, organized into four main phases: **Planning**, **Execution**, **Problem-Solving**, and **Evaluation**. Each phase captures distinct patterns of user cognition and interaction, from initial goal formulation to iterative action, troubleshooting, and reflection.

(Execution) and critical work (Evaluation).

Dataset Curation. After constructing the taxonomy, we aligned each video with its corresponding narration segments. For every segment, we annotated the user’s behavior state using *Gemini-2.5-Pro* according to the taxonomy, prompting the model to produce both a predicted label and its reasoning. Two human annotators recruited from Prolific then verified and refined these annotations, achieving a 96.1% agreement rate. Finally, we uniformly sampled 200 instances from each of the nine classes, resulting in a balanced dataset of 1.8K annotated segments.

3.2.2. Intent Prediction

Description. This task evaluates whether a model can reason about the user’s short-term, immediate goal in context. It focuses on identifying what the user aims to achieve within open-ended workflows.

Dataset Curation. Using the narration-aligned video segments, we prompted *Gemini-2.5-Pro* to infer users’ intention in each segment. The think-aloud narrations often revealed users’ goals (e.g., “I’m going to align these objects”, “I’ll try another color”). Leveraging this signal, we prompted the model to infer the underlying user intention. After collecting and deduplicating the inferred intents, we further instructed the model to generate three plausible but incorrect alternatives to serve as distractors for the multiple-choice evaluation. The resulting intent annotations and distractors were then validated by the authors, with 88.68% of the data retained, yielding a final set of 1.3K instances.

3.2.3. Help Prediction

Description. The final task evaluates whether a model can progress from understanding and reasoning to deciding how to assist. Help Prediction consists of two subtasks: (1) **Help Need Detection**, a binary classification task that determines whether the user needs help, and (2) **Help Content Prediction**,

which identifies the specific type of help needed, such as explaining a feature or suggesting an alternative. Together, these subtasks assess a model’s ability to anticipate user needs and recommend appropriate assistance, bridging the gap between perception and actionable support.

Dataset Curation. We identified potential help-seeking moments using two complementary signals. First, *explicit help-seeking* behaviors, such as switching to external resources (e.g., Google, YouTube, ChatGPT), indicated direct attempts to seek guidance. Second, *implicit help-seeking* cues were extracted from user narration, where they expressed uncertainty or confusion (e.g., “How do I align this?”, “I can’t find Layer Mask.”). Additionally, we included clear *no-help-needed* moments, where users demonstrated confidence through their narration. Using these signals, *Gemini-2.5-Pro* was prompted to generate initial annotations for help-need and help-content labels. After deduplication, the model was additionally prompted to generate three plausible but incorrect options for each instance for multiple-choice question evaluation. All annotations and distractors were then reviewed by the authors, resulting in 1K validated instances, with 78.89% of the original data retained. For 12.5% of the retained instances, the segment’s start or end time was adjusted to exclude explicit visual help signals (e.g., user turning to Google Search) to ensure fair evaluation. Overall, 66% of the instances were labeled as help-needed, while the remaining 34% required no help.

4. Experiments

4.1. Experimental Setup

We evaluate a range of multimodal large language models (MLLMs) on our benchmark to assess their ability to understand, reason about, and assist users in open-ended software workflows. Our evaluation includes eight represen-

Model	(1) Behavior Detection		(2) Intent Prediction		(3) Help Prediction					
	-	+ Prev.	-	+ Behavior	Help Need Detection			Help Content Prediction		
					-	+ Behv.	+Behv.+Intent	-	+ Behv.	+Behv.+Intent
Gemini-2.5-Flash [15]	36.91	38.19	65.40	66.77	53.64	76.33	78.07	49.53	53.75	78.59
Gemini-2.5-Pro [15]	42.44	43.79	67.80	70.16	69.82	84.73	82.38	52.74	57.03	79.69
GPT-4o-mini [35]	17.65	17.07	60.76	62.19	46.05	78.92	82.26	31.32	42.86	79.84
GPT-4o [35]	36.32	37.24	61.19	62.58	49.69	87.79	87.91	45.95	48.37	79.78
Claude-4.5-Sonnet [1]	44.61	45.63	71.39	72.62	39.49	58.56	59.43	55.00	62.17	82.79
Qwen3-VL-8B [42]	37.97	38.13	62.70	64.03	52.83	70.39	77.36	46.06	50.63	80.11
InternVideo2.5-8B [43]	21.57	27.02	43.79	45.13	34.36	35.35	35.25	23.67	29.15	73.86
InternVL3-8B [57]	22.57	24.90	46.11	46.97	34.94	43.73	46.82	27.03	32.20	72.97

Table 2. Evaluation results on accuracy across (1) Behavior State Detection, (2) Intent Prediction, and (3) Help Prediction.

tative MLLMs spanning both proprietary and open-source models: **Gemini-2.5-Flash** [15], **Gemini-2.5-Pro** [15], **GPT-4o-mini** [35], **GPT-4o** [35], **Claude-4.5-Sonnet** [1], **Qwen3-VL-8B** [42], **InternVideo2.5-Chat-8B** [43], and **InternVL3-8B** [57]. All models are evaluated in a zero-shot setting using publicly available APIs or checkpoints, without any additional fine-tuning.

For each test instance, we uniformly sample 32 frames from the corresponding video segment, providing only visual input (excluding narration audio) to simulate perception based solely on visual cues. To ensure consistency across models, we use standardized prompting templates (Section H). We also prompt models to generate both a predicted label and supporting reasoning, a strategy shown to improve task performance [23].

Our main experiments are conducted in an offline inference setting, where models solve the task given the full video. To approximate real-world proactive assistant scenarios, we additionally evaluate an online setting, where the model receives visual input progressively—at 25%, 50%, 75%, and 100% of the segment, we uniformly sample 32 frames from the corresponding prefix for inference.

4.2. Evaluation Tasks

(1) Behavior State Detection. This task measures whether a model can identify the user’s behavioral state from a given video segment. We provide each model with clips and ask it to classify them into one of nine taxonomy-defined states. Two configurations are tested: (i) using only the current segment and (ii) with prior history, where the model is given the immediately preceding segment’s behavior state. This is framed as a multi-class classification problem, and performance is evaluated using accuracy.

(2) Intent Prediction. This task evaluates a model’s ability to infer the user’s underlying goal within a given video segment. Models are prompted to predict what the user’s

goal in two settings: (i) using only the current segment, and (ii) with additional behavior state context, where the model is also given the state label and its definition. We adopt a multiple-choice question (MCQ) format, where the model selects the most likely intent from four candidates. Performance is measured using accuracy. For the default setting (i), to mitigate potential bias, we additionally report multi-binary accuracy (MBAcc) following prior work [5, 7, 8], which evaluates whether the model correctly identifies the ground-truth intent in all three pairwise comparisons against incorrect alternatives.

(3) Help Prediction. The final task evaluates whether models can move beyond understanding and reasoning to provide actionable assistance. Given a video segment, models are asked to predict whether the user requires help (*Need*), and if so, what kind of help would be most appropriate (*Content*). **Help Need Detection** is framed as a binary classification task and evaluated using accuracy, precision, recall, and F1-score. **Help Content Prediction**, similar to Intent Prediction, uses a multiple-choice question (MCQ) format and is evaluated using accuracy and multi-binary accuracy (MBAcc) for the default setting. We test three settings for both tasks: (i) video only, (ii) video + behavior state, where the model is given the behavior label and its definition for the current segment, and (iii) video + behavior state + intent, where the model additionally receives the identified user intention. These settings progressively assess the model’s ability to leverage layered user context for meaningful, situation-aware assistance.

4.3. Results

Table 2 presents the performance of baseline models on GUIDE across the tasks, with accuracies reported under default and context-augmented settings. Overall, models performed weakest on Behavior State Detection and Help Prediction, with default-setting accuracies peaking at 44.61%

Model	Help Need Detection											
	–				+ Behavior State				+ Behavior State + Intent			
	Acc	Prec	Rec	F1	Acc	Prec	Rec	F1	Acc	Prec	Rec	F1
Gemini-2.5-Flash [15]	53.64	83.27	36.62	50.87	76.33	97.67	65.47	78.39	78.07	94.56	70.62	80.86
Gemini-2.5-Pro [15]	69.82	76.42	78.09	77.42	84.73	93.61	82.34	87.61	82.38	91.20	80.94	86.76
GPT-4o-mini [35]	46.05	83.03	22.31	35.17	76.73	97.61	66.23	78.92	79.71	97.20	71.29	82.26
GPT-4o [35]	49.69	74.41	35.14	47.73	87.79	95.39	85.53	90.19	87.91	95.12	85.95	90.30
Claude-4.5-Sonnet [1]	39.49	87.69	8.92	16.19	58.56	99.16	37.09	53.99	59.43	99.19	38.44	55.41
Qwen3-VL-8B [42]	52.83	79.86	34.23	47.92	70.39	94.35	58.50	72.22	77.36	95.38	67.56	79.09
InternVideo2.5-8B [43]	34.36	33.33	0.16	0.31	35.35	90.91	1.56	3.07	35.25	83.33	1.56	3.07
InternVL3-8B [57]	34.94	72.73	1.25	2.46	43.73	98.88	15.77	27.20	46.82	98.40	19.22	32.16

Table 3. Results for **Help Need Detection** on accuracy, precision, recall, and F1-score across three conditions (default, with behavior state, with behavior state and intent).

and 55.00% for Behavior State Detection and Help Content Prediction, respectively, both from Claude-4.5-Sonnet [1]. While Gemini-2.5-Pro [15] reached nearly 70% accuracy on Help Need Detection, most other models showed substantially lower performance across both Help sub-tasks. Across tasks, we observe that models generally benefit from added behavioral and intent context, with particularly notable improvements in help-related predictions. We report the main findings below.

4.3.1. Behavior State Detection

Behavior State Detection remains highly challenging.

All models struggled to accurately infer the user’s behavioral state from video segments, underscoring the difficulty of the 9-way classification task. While proprietary models such as Claude-4.5-Sonnet [1] and Gemini-2.5-Pro [15] performed best, no model surpassed 45% accuracy, and most fell below 40%.

Models often misinterpret signals of struggle. The most common failure was misclassifying *Frustration* or *Debugging* as *Performing Actions* or *Exploration and Decision-Making*, as shown in the confusion matrix (Figure C4 in supp.). These errors suggest that models overlook subtle indicators of user difficulty—e.g., repeated clicks, hesitation, or undoing actions—instead interpreting them as deliberate progress, revealing a lack of nuanced understanding of user frustration signals.

Temporal context shows modest potential. Incorporating the prior behavior state led to small but consistent gains across models. While most improvements were marginal, the largest gain was observed for InternVideo2.5-8B [43] with 5.45 percentage points, suggesting that temporal context holds value and may be more effectively utilized with

improved temporal reasoning capabilities.

Model	Intent Prediction		Help Prediction	
	Acc	MBAcc	Acc	MBAcc
Gemini-2.5-Flash [15]	65.40	59.09	49.53	44.69
Gemini-2.5-Pro [15]	67.80	64.34	52.74	45.31
GPT-4o-mini [35]	60.76	50.24	31.32	28.59
GPT-4o [35]	61.19	56.58	45.95	41.25
Claude-4.5-Sonnet [1]	71.39	65.44	55.00	50.78
Qwen3-VL-8B [42]	62.70	58.07	46.06	44.69
InternVideo2.5-8B [43]	43.79	27.98	23.67	18.75
InternVL3-8B [57]	46.11	40.75	27.03	23.75

Table 4. Evaluation of **Intent Prediction** and **Help Content Prediction**, with Accuracy (Acc) and Multi-Binary Accuracy (MBAcc).

4.3.2. Intent Prediction

Intent Prediction is the most tractable task, but still imperfect. Among the three tasks, models achieved the highest performance on intent prediction, with several surpassing 60% accuracy. However, performance drops under the stricter MBAcc metric, which requires consistent discrimination across all answer pairs. This indicates that while models can often select a plausible intent, they still struggle with reliably identifying the correct one over all distractors (Table 4).

Behavior context helps, but only slightly. Incorporating behavior state context (i.e., the user’s behavioral label and definition) consistently improved performance, but the gains were relatively modest across all models. This suggests that while such context may offer useful cues, it does




Figure 4. Accuracy trends across the tasks in the online setting, where models are given progressively more of the video segment (25%, 50%, 75%, and 100%). Models show consistent improvement as they see more segments, with Gemini-2.5-Flash [15] and Qwen3-VL-8B [42] showing larger and more consistent gains across all four tasks compared to the smaller open-source models.

not provide sufficient information on its own or is not yet effectively leveraged by current models for intent inference.

4.3.3. Help Prediction

High variance and missed help cases in Need Detection. Table 3 shows the full performance results for Help Need Detection. This subtask exhibited the most variance across models, with F1 scores ranging from 0.31 (InternVideo2.5-8B [43]) to 77.42 (Gemini-2.5-Pro [15]). Notably, recall was particularly low across most models—except for Gemini-2.5-Pro, all others had recall under 37%. This indicates that many instances where users actually needed help were misclassified as not needing it, echoing similar trends in Behavior State Detection (Section 4.3.1) where models frequently misinterpreted signals of struggle.

Behavior context improves Help Need Detection. Providing the user’s behavior state led to consistent and significant improvements in Help Need Detection across all models, with the largest gain observed in GPT-4o [35], with a 42.46-point increase in F1 score. This suggests that context, such as whether a user is exploring or showing signs of frustration, provides strong cues for determining help needs.

Help Content Prediction remains challenging, but benefits from intent context. Help Content Prediction proved particularly challenging, with all models struggling and the top accuracy reaching only 55% from Claude-4.5-Sonnet [1], which further declined to around 50% under the stricter MBAcc evaluation. However, incorporating user intent led to substantial improvements across models, with the largest gain in InternVideo2.5-8B [43] at 50.19 percentage points, highlighting the importance of understanding both user state and intent for providing targeted support.

4.3.4. Other Findings

Online vs. Offline Setting: models benefit more from temporal context. In our online simulation experiment,

where models are given progressively more of the video segment (25%, 50%, 75%, and 100%), we observe consistent performance gains across all four tasks (Figure 4). Gemini-2.5-Flash [15] and Qwen3-VL-8B [42] show larger and more consistent gains across all tasks, compared to the smaller open-source models, indicating a strong ability to integrate growing context into more accurate predictions. These findings suggest that gathering appropriate context over time is crucial for proactive AI assistance, where systems must not only react but also anticipate user needs based on incomplete and evolving information.

Outlook for Model Improvements. Together, these results suggest that incorporating structured user context (behavior state and intent) and temporal context consistently improves help prediction. Recent work on agents demonstrates the effectiveness of context engineering via stratified memory, where interaction history is selectively structured rather than treated as a flat sequence [49]. Applying this idea to GUI assistance is a promising direction for better leveraging long-horizon user context.

5. Conclusion

We introduced a benchmark for evaluating models in understanding, reasoning about, and assisting users in open-ended GUI-based workflows. Grounded in real-world novice user demonstrations, our tasks—behavior state detection, intent prediction, and help prediction—capture core capabilities needed for collaborative GUI agents. Evaluation across state-of-the-art MLLMs revealed that models struggle to interpret nuanced user behavior and accurately infer assistance needed in open-ended scenarios. However, when provided with appropriate user context, models showed consistent improvements, highlighting the value of structured user understanding. Overall, our benchmark provides a foundation for user-aware agents that support human workflows.

Acknowledgements

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. 2021-0-01347, Video Interaction Technologies Using Object-Oriented Video Modeling and No. RS-2024-00443251, Accurate and Safe Multimodal, Multilingual Personalized AI Tutors).

References

- [1] Anthropic. Introducing claude sonnet 4.5. <https://www.anthropic.com/news/claude-sonnet-4-5>, 2025. Anthropic News Release, September 29, 2025. 6, 7, 8
- [2] Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. Whisperx: Time-accurate speech transcription of long-form audio. *INTERSPEECH 2023*, 2023. 4
- [3] Omri Berkovitch, Sapir Caduri, Noam Kahlon, Anatoly Efros, Avi Caciularu, and Ido Dagan. Identifying user goals from ui trajectories. In *Companion Proceedings of the ACM on Web Conference 2025*, page 2381–2390, New York, NY, USA, 2025. Association for Computing Machinery. 3
- [4] Benjamin S. Bloom. *Taxonomy of educational objectives: The classification of educational goals*. Longman Group, 1st edition, 1956. 2, 4
- [5] Mu Cai, Reuben Tan, Jianrui Zhang, Bocheng Zou, Kai Zhang, Feng Yao, Fangrui Zhu, Jing Gu, Yiwu Zhong, Yuzhang Shang, Yao Dou, Jaden Park, Jianfeng Gao, Yong Jae Lee, and Jianwei Yang. Temporalbench: Towards fine-grained temporal understanding for multimodal video models. *arXiv preprint arXiv:2410.10818*, 2024. 6, 1
- [6] Valerie Chen, Alan Zhu, Sebastian Zhao, Hussein Mozannar, David Sontag, and Ameet Talwalkar. Need help? designing proactive ai assistants for programming. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2025. Association for Computing Machinery. 2, 3
- [7] Jen-Hao Cheng, Vivian Wang, Huayu Wang, Huapeng Zhou, Yi-Hao Peng, Hou-I Liu, Hsiang-Wei Huang, Kuang-Ming Chen, Cheng-Yen Yang, Wenhao Chai, et al. Tempura: Temporal event masked prediction and understanding for reasoning in action. *arXiv preprint arXiv:2505.01583*, 2025. 6
- [8] Jang Hyun Cho, Andrea Madotto, Effrosyni Mavroudi, Triantafyllos Afouras, Tushar Nagarajan, Muhammad Maaz, Yale Song, Tengyu Ma, Shuming Hu, Hanoona Rasheed, Peize Sun, Po-Yao Huang, Daniel Bolya, Suyog Jain, Miguel Martin, Huiyu Wang, Nikhila Ravi, Shashank Jain, Temmy Stark, Shane Moon, Babak Damavandi, Vivian Lee, Andrew Westbury, Salman Khan, Philipp Krähenbühl, Piotr Dollár, Lorenzo Torresani, Kristen Grauman, and Christoph Feichtenhofer. Perceptionlm: Open-access data and models for detailed visual understanding. *arXiv:2504.13180*, 2025. 6, 1
- [9] Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web, 2023. 1
- [10] Aarash Feizi, Shravan Nayak, Xiangru Jian, Kevin Qinghong Lin, Kaixin Li, Rabiul Awal, Xing Han Lù, Johan Obando-Ceron, Juan A. Rodriguez, Nicolas Chapados, David Vazquez, Adriana Romero-Soriano, Reihaneh Rabbany, Perouz Taslakian, Christopher Pal, Spandana Gella, and Sai Rajeswar. Grounding computer use agents on human demonstrations, 2025. 1
- [11] K. J. Kevin Feng, David W. McDonald, and Amy X. Zhang. Levels of autonomy for AI agents. *Knight First Amendment Institute – AI and Democratic Freedoms Essay Series*, 2025. 3
- [12] Figma. Figma make. <https://www.figma.com/make/>, 2025. Accessed: 2025-11-14. 1
- [13] Kiran Gadhav, Jochen Görtler, Zach Cutler, Carolina Nore, Oliver Deussen, Miriah Meyer, Jeff M. Phillips, and Alexander Lex. Predicting intent behind selections in scatterplot visualizations. *Information Visualization*, 20(4):207–228, 2021. 3
- [14] Difei Gao, Lei Ji, Zechen Bai, Mingyu Ouyang, Peiran Li, Dongxing Mao, Qinchen Wu, Weichen Zhang, Peiyi Wang, Xiangwu Guo, Hengxu Wang, Luwei Zhou, and Mike Zheng Shou. Assistgui: Task-oriented pc graphical user interface automation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13289–13298, Seattle, WA, USA, 2024. IEEE. Benchmarks PC GUI automation with an actor-critic agent framework. 1, 2, 3
- [15] Google Gemini Team. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 4, 6, 7, 8
- [16] Wenyi Hong, Weihang Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, and Jie Tang. Cogagent: A visual language model for gui agents, 2023. 1
- [17] Eric Horvitz, Jack Breese, David Heckerman, David Hovel, and Koos Rommelse. The lumière project: Bayesian user modeling for inferring the goals and needs of software users. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, page 256–265, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc. 2
- [18] Faria Huq, Zora Zhiruo Wang, Frank F. Xu, Tianyue Ou, Shuyan Zhou, Jeffrey P. Bigham, and Graham Neubig. Cow-Pilot: A framework for autonomous and human-agent collaborative web navigation. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 163–172, Albuquerque, New Mexico, 2025. Association for Computational Linguistics. 3
- [19] Lawrence Jang, Yinheng Li, Charles Ding, Justin Lin, Paul Pu Liang, Dan Zhao, Rogerio Bonatti, and Kazuhito Koishida. Videowebarena: Evaluating long context multimodal agents with video understanding web tasks. 2024. 2, 3
- [20] Yunseok Jang, Yeda Song, Sungryull Sohn, Lajanugen Logeswaran, Tiange Luo, Dong-Ki Kim, Kyunghoon Bae,

- and Honglak Lee. Scalable Video-to-Dataset Generation for Cross-Platform Mobile Agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 1
- [21] Anjali Khurana, Hariharan Subramonyam, and Parmit K Chilana. Why and when llm-based assistants can go wrong: Investigating the effectiveness of prompt-based interactions for software help-seeking. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, page 288–303, New York, NY, USA, 2024. Association for Computing Machinery. 3
- [22] Anjali Khurana, Xiaotian Su, April Yi Wang, and Parmit K Chilana. Do it for me vs. do it with me: Investigating user perceptions of different paradigms of automation in copilots for feature-rich software. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2025. Association for Computing Machinery. 1, 3
- [23] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2022. Curran Associates Inc. 6
- [24] Minhwa Lee, Zae Myung Kim, Vivek Khetan, and Dongyeop Kang. Human-ai collaborative taxonomy construction: A case study in profession-specific writing assistants. In *Proceedings of the Third Workshop on Intelligent and Interactive Writing Assistants*, page 51–57, New York, NY, USA, 2024. Association for Computing Machinery. 4
- [25] Kunpeng Li, Chen Fang, Zhaowen Wang, Seokhwan Kim, Hailin Jin, and Yun Fu. Screenshot tutorial video understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3
- [26] Toby Jia-Jun Li, Jingya Chen, Tom M. Mitchell, and Brad A. Myers. Towards effective human-ai collaboration in gui-based interactive task learning agents. In *CHI 2020 Workshop on Artificial Intelligence for HCI: A Modern Approach (AI4HCI)*, 2020. 3
- [27] Kevin Qinghong Lin, Linjie Li, Difei Gao, WU Qinchen, Mingyi Yan, Zhengyuan Yang, Lijuan Wang, and Mike Zheng Shou. Videogui: A benchmark for gui automation from instructional videos. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. 1, 2, 3
- [28] Kevin Qinghong Lin, Linjie Li, Difei Gao, Zhengyuan Yang, Shiwei Wu, Zechen Bai, Stan Weixian Lei, Lijuan Wang, and Mike Zheng Shou. Showui: One vision-language-action model for gui visual agent. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19498–19508, 2025. 1
- [29] Guangyi Liu, Pengxiang Zhao, Liang Liu, Zhiming Chen, Yuxiang Chai, Shuai Ren, Hao Wang, Shibo He, and Wenchao Meng. Learnact: Few-shot mobile gui agent with a unified demonstration benchmark. *arXiv preprint arXiv:2504.13805*, 2025. 1
- [30] Dunjie Lu, Yiheng Xu, Junli Wang, Haoyuan Wu, Xinyuan Wang, Zekun Wang, Junlin Yang, Hongjin Su, Jixuan Chen, Junda Chen, Yuchen Mao, Jingren Zhou, Junyang Lin, Binyuan Hui, and Tao Yu. Videoagentrek: Computer use pretraining from unlabeled videos, 2025. 3
- [31] Yaxi Lu, Shenzhi Yang, Cheng Qian, Guirong Chen, Qinyu Luo, Yesai Wu, Huadong Wang, Xin Cong, Zhong Zhang, Yankai Lin, Weiwen Liu, Yasheng Wang, Zhiyuan Liu, Fangming Liu, and Maosong Sun. Proactive agent: Shifting LLM agents from reactive responses to active assistance. In *The Thirteenth International Conference on Learning Representations*, 2025. 1
- [32] Microsoft Corporation. Microsoft copilot. <https://copilot.microsoft.com/>, 2025. Accessed: 2025-11-14. 1
- [33] Shravan Nayak, Xiangru Jian, Kevin Qinghong Lin, Juan A. Rodriguez, Montek Kalsi, Nicolas Chapados, M. Tamer Özsu, Aishwarya Agrawal, David Vazquez, Christopher Pal, Perouz Taslakian, Spandana Gella, and Sai Rajeswar. Uivision: A desktop-centric gui benchmark for visual perception and interaction. In *Proceedings of the 42nd International Conference on Machine Learning*, pages 45817–45851. PMLR, 2025. Fine-grained desktop GUI benchmark with dense annotations. 2, 3
- [34] Donald A. Norman. *The Design of Everyday Things*. Basic Books, New York, 1988. 2, 4
- [35] OpenAI. Gpt-4o system card, 2025. 6, 7, 8
- [36] SHRADDHA VIJAY PAWAR, Balavarun Pedapudi, Pramod Kaushik, Sarath Sivaprasad, Mario Fritz, and Shirish Karande. EARL: Early intent recognition in GUI tasks using theory of mind. In *ICML 2025 Workshop on Computer Use Agents*, 2025. 3
- [37] Yi-Hao Peng, Dingzeyu Li, Jeffrey P Bigham, and Amy Pavel. Morae: Proactively pausing ui agents for user choices. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*, New York, NY, USA, 2025. Association for Computing Machinery. 2
- [38] Kevin Pu, Daniel Lazaro, Ian Arawjo, Haijun Xia, Ziang Xiao, Tovi Grossman, and Yan Chen. Assistance or disruption? exploring and evaluating the design and trade-offs of proactive ai programming support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2025. Association for Computing Machinery. 2, 3
- [39] Kevin Pu, Ting Zhang, Naveen Sendhilnathan, Sebastian Freitag, Raj Sodhi, and Tanya R. Jonker. Promemassist: Exploring timely proactive assistance through working memory modeling in multi-modal wearable devices. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*, New York, NY, USA, 2025. Association for Computing Machinery. 3
- [40] Chan Hee Song, Yiwen Song, Palash Goyal, Yu Su, Oriana Riva, Hamid Palangi, and Tomas Pfister. Watch and Learn: Learning to Use Computers from Online Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. To appear. 1
- [41] Yixiao Song, Katherine Thai, Chau Minh Pham, Yapei Chang, Mazin Nadaf, and Mohit Iyyer. Bearcubs: A benchmark for computer-using web agents, 2025. 1

- [42] Qwen Team. Qwen3 technical report, 2025. 6, 7, 8
- [43] Yi Wang, Xinhao Li, Ziang Yan, Yinan He, Jiashuo Yu, Xiangyu Zeng, Chenting Wang, Changlian Ma, Haian Huang, Jianfei Gao, Min Dou, Kai Chen, Wenhai Wang, Yu Qiao, Yali Wang, and Limin Wang. Internvideo2.5: Empowering video mllms with long and rich context modeling. *arXiv preprint arXiv:2501.12386*, 2025. 6, 7, 8
- [44] Qinchen Wu, Difei Gao, Qinghong Lin, Zhuoyu Wu, and Mike Zheng Shou. Gui-narrator: Detecting and captioning computer gui actions. In *Proceedings of the 33rd ACM International Conference on Multimedia*, page 3683–3692, New York, NY, USA, 2025. Association for Computing Machinery. 3
- [45] Shirley Wu, Michel Galley, Baolin Peng, Hao Cheng, Gavin Li, Yao Dou, Weixin Cai, James Zou, Jure Leskovec, and Jianfeng Gao. Collabllm: From passive responders to active collaborators. In *International Conference on Machine Learning (ICML)*, 2025. 1
- [46] Shirley Wu, Michel Galley, Baolin Peng, Hao Cheng, Gavin Li, Yao Dou, Weixin Cai, James Zou, Jure Leskovec, and Jianfeng Gao. Collabllm: From passive responders to active collaborators. *arXiv preprint arXiv:2502.00640*, 2025. 2
- [47] Bufang Yang, Lilin Xu, Liekang Zeng, Kaiwei Liu, Siyang Jiang, Wenrui Lu, Hongkai Chen, Xiaofan Jiang, Guoliang Xing, and Zhenyu Yan. Contextagent: Context-aware proactive llm agents with open-world sensory perceptions, 2025. 1
- [48] Qinglong Yang, Haoming Li, Haotian Zhao, Xiaokai Yan, Jingtao Ding, Fengli Xu, and Yong Li. Fingertip 20k: A benchmark for proactive and personalized mobile llm agents, 2025. 1
- [49] Ruozhen Yang, Yucheng Jiang, Yueqi Jiang, Priyanka Kargupta, Yunyi Zhang, and Jiawei Han. Grounding agent memory in contextual intent, 2026. 8
- [50] Suyu Ye, Haojun Shi, Darren Shih, Hyokun Yun, Tanya Roosta, and Tianmin Shu. Realwebassist: A benchmark for long-horizon web assistance with real-world users. *arXiv preprint arXiv:2504.10445*, 2025. 1
- [51] Bofei Zhang, Zirui Shang, Zhi Gao, Wang Zhang, Rui Xie, Xiaojian Ma, Tao Yuan, Xinxiao Wu, Song-Chun Zhu, and Qing Li. Tongui: Building generalized gui agents by learning from multimodal web tutorials. *arXiv preprint arXiv:2504.12679*, 2025. Introduces the TongUI framework and GUI-Net-1M dataset. 1
- [52] Ceyao Zhang, Kaijie Yang, Siyi Hu, Zihao Wang, Guanghe Li, Yihang Sun, Cheng Zhang, Zhaowei Zhang, Anji Liu, Song-Chun Zhu, Xiaojun Chang, Junge Zhang, Feng Yin, Yitao Liang, and Yaodong Yang. Proagent: Building proactive cooperative agents with large language models, 2024. 1
- [53] Jiwen Zhang, Jihao Wu, Teng Yihua, Minghui Liao, Nuo Xu, Xiao Xiao, Zhongyu Wei, and Duyu Tang. Android in the zoo: Chain-of-action-thought for GUI agents. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 12016–12031, Miami, Florida, USA, 2024. Association for Computational Linguistics. 1
- [54] Henry Hengyuan Zhao, Kaiming Yang, Wendi Yu, Difei Gao, and Mike Zheng Shou. Worldgui: An interactive benchmark for desktop gui automation from any starting point, 2025. 1, 2, 3
- [55] Yuheng Zhao, Xueli Shu, Liwen Fan, Lin Gao, Yu Zhang, and Siming Chen. Proactiveva: Proactive visual analytics with llm-based ui agent, 2025. 3
- [56] Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Yonatan Bisk, Daniel Fried, Uri Alon, et al. Webarena: A realistic web environment for building autonomous agents. *arXiv preprint arXiv:2307.13854*, 2023. 1
- [57] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025. 6, 7

GUIDE: A Benchmark for Understanding and Assisting Users in Open-Ended GUI Tasks

Supplementary Material

A. Detailed Evaluation Metrics

In this section, we provide the formal definitions for the evaluation metrics used across our four evaluation tasks: Behavior State Detection, Intent Prediction, Help Need Detection, and Help Content Prediction. Let N denote the total number of test samples in the dataset. For the i -th sample, let y_i denote the ground-truth label and \hat{y}_i denote the model’s predicted label. $\mathbb{I}(\cdot)$ denotes the indicator function, which equals 1 if the condition inside is true and 0 otherwise.

A.1. Metric Definitions by Task

A.1.1. Task 1: Behavior State Detection

This task is formulated as a multi-class classification problem where the model must classify a video segment into one of 9 distinct behavioral states. We evaluate performance using standard **Accuracy**.

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{y}_i = y_i) \quad (1)$$

A.1.2. Task 2: Intent Prediction

This task is framed as a Multiple-Choice Question (MCQ) task with 4 options (1 correct answer and 3 distractors). We use two metrics:

Accuracy. Measures the proportion of instances where the model selects the correct intent option from the four candidates.

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{y}_i = y_i) \quad (2)$$

Multi-Binary Accuracy (MBAcc). Following prior work [5, 8], we employ MBAcc to evaluate robustness against distractors. For a given sample i , let y_i be the correct option and $\mathcal{C}_i^- = \{c_{i,1}, c_{i,2}, c_{i,3}\}$ be the set of three incorrect distractor options. The model performs a pairwise comparison function $f(x, \text{opt}_A, \text{opt}_B)$ which returns the chosen option between A and B. A prediction is considered correct under MBAcc only if the model prefers the ground truth y_i over *every* distractor in \mathcal{C}_i^- .

$$\text{MBAcc} = \frac{1}{N} \sum_{i=1}^N \left(\prod_{c \in \mathcal{C}_i^-} \mathbb{I}(f(x_i, y_i, c) = y_i) \right) \quad (3)$$

A.1.3. Task 3-1: Help Prediction (Need Detection)

This sub-task is a binary classification problem (Help Needed vs. Not Needed). We evaluate this using Accuracy, Precision, Recall, and F1-Score. Let TP (True Positive), TN (True Negative), FP (False Positive), and FN (False Negative) denote the classification counts.

- **Accuracy:** The ratio of correctly predicted observations to total observations.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (4)$$

- **Precision:** The ratio of correctly predicted positive observations to the total predicted positives.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (5)$$

- **Recall:** The ratio of correctly predicted positive observations to the all observations in the actual class.

$$\text{Recall} = \frac{TP}{TP + FN} \quad (6)$$

- **F1-Score:** The harmonic mean of Precision and Recall.

$$\text{F1} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (7)$$

A.1.4. Task 3-2: Help Prediction (Content Prediction)

Similar to Intent Prediction, this sub-task is an MCQ task where the model must select the appropriate help content. It is evaluated using **Accuracy** and **Multi-Binary Accuracy (MBAcc)**.

Accuracy.

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{y}_i = y_i) \quad (8)$$

Multi-Binary Accuracy (MBAcc). Defined identically to the Intent Prediction task. Let \mathcal{C}_i^- be the set of incorrect help content options for the i -th sample.

$$\text{MBAcc} = \frac{1}{N} \sum_{i=1}^N \left(\prod_{c \in \mathcal{C}_i^-} \mathbb{I}(f(x_i, y_i, c) = y_i) \right) \quad (9)$$

B. Dataset Details

We provide a comprehensive overview of the GUIDE dataset, detailing its statistical properties, task granularity, and the diverse range of software workflows it encompasses.

B.1. Dataset Statistics

GUIDE comprises a comprehensive collection of 120 screen recording videos, totaling approximately 67.5 hours of footage. A key characteristic of our dataset is the inclusion of rich verbal narration; as shown in Table B1, think-aloud narration covers **78% of the total video duration**, providing high-quality ground truth for annotating user intent and mental states.

Variable	Value
# Videos	120
Total Duration	67.5 hours
Avg. Duration	33 min 44 sec
Max Duration	1 hour 23 min 50 sec
Min Duration	16 min 42 sec
Think-Aloud Narration Ratio	78%
<i>Task Samples & Granularity</i>	
(1) Behavior State Detection	1.8K
<i>Avg. Segment Length</i>	14.16s
(2) Intent Prediction	1.3K
<i>Avg. Segment Length</i>	25.40s
(3) Help Prediction	1K
<i>Avg. Segment Length</i>	25.56s

Table B1. Statistics of the GUIDE dataset.

The dataset focuses on long-horizon, open-ended workflows. The average video duration is 33 minutes and 44 seconds, with sessions ranging from approximately 16 minutes to over 1 hour and 23 minutes (Figure B1). This extended duration ensures that the dataset captures the full evolution of user tasks, including periods of exploration, struggle, and error recovery.

Task Granularity. From these raw videos, we extracted varying numbers of instances for our three evaluation tasks. We collected **1.8K samples** for Behavior State Detection, **1.3K samples** for Intent Prediction, and **1K samples** for Help Prediction. Notably, the average segment length for behavior detection is shorter (14.16s) compared to Intent Prediction (25.4s) and Help Prediction (25.56s). This is because when annotating behavior states from narration-aligned segments, we instructed the model to split the clip if two or more states were identified.

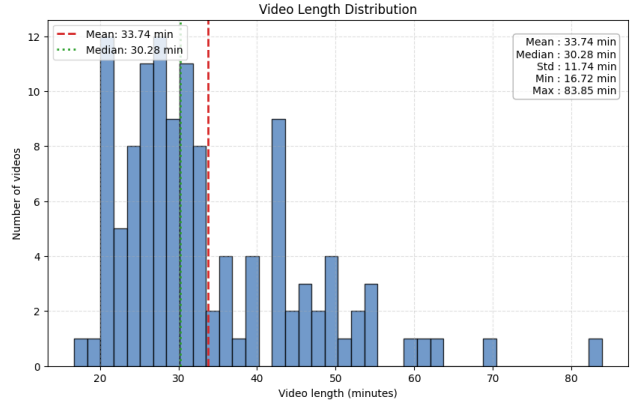


Figure B1. Distribution of screen recording video lengths in the dataset.




Figure B2. Software categories represented in the dataset.

Diversity. To ensure generalizability, the dataset spans a wide variety of software domains. As illustrated in Figure B2, users interacted with diverse applications ranging from creative design tools to analytical software.

B.2. Task Composition

To ensure our benchmark captures a comprehensive range of user behaviors, we designed a set of 20 open-ended tasks across five distinct software categories: Photo Editing, Graphic Design, Presentation Design, Video Editing, and Data Analysis. Table B2 provides a detailed overview of these categories and their corresponding tasks.

Open-Ended Task Design. Unlike rigid, step-by-step tutorials that result in linear behavior, our tasks are designed to be goal-oriented and open-ended. For instance, while we provided users with necessary materials (e.g., raw video clips, images) and suggested specific software features to utilize, we did not prescribe a fixed execution path or a target reference outcome. This semi-structured ambiguity is intentional; it forces users to engage in high-level planning, trial-and-error exploration, and problem-solving. Consequently, this setup naturally elicits the complex behavior states—such as *Exploration*, *Debugging*, and *Frustration*—that GUIDE aims to detect.

Domain Diversity. The selected software categories cover a broad spectrum of software domains, ensuring comprehensive coverage of diverse GUI workflows. Our dataset spans **creative domains** (Photo Editing, Graphic Design) that rely on visual manipulation and aesthetic decisions, **analytical domains** (Data Analysis) focused on data processing and logic, and hybrid tasks like **Presentation Design** or **Video Editing**. This variety ensures that our models are evaluated on their ability to generalize across diverse user interfaces, toolsets, and workflow paradigms.

Category	Software	Tasks
Photo Editing	Photoshop, GIMP	<ol style="list-style-type: none">1. Create a composite from two images.2. Create a bakery logo with a warm, friendly identity.3. Replace a photo’s background with a custom-designed pattern.4. Design a movie poster.
Graphic Design	Figma, Canva	<ol style="list-style-type: none">1. Design a mobile sign-up screen for a fictional app.2. Design a custom 404 error page with a visual and animated element.3. Design compact profile cards that display personal user details.4. Design an event poster for a music festival.
Presentation Design	PowerPoint, Google Slides	<ol style="list-style-type: none">1. Create a product pitch deck that highlights the MacBook’s key features.2. Create an interactive timeline presenting a company’s history.3. Create a 5-slide nature-themed shape-masked photo scrapbook.4. Create a quiz deck with 3 multiple-choice questions.
Video Editing	Premiere Pro, CapCut	<ol style="list-style-type: none">1. Edit a short interview to improve clarity and engagement.2. Design a creative intro using animated text.3. Edit a short instructional video to clearly guide a process.4. Transform a long video into a highly engaging short-form clip.
Data Analysis	Microsoft Excel, Google Sheets	<ol style="list-style-type: none">1. Design a Gantt chart for a mini project.2. Summarize and visualize responses from a survey.3. Visualize student performance across subjects.4. Summarize and visualize product sales by category or region.

Table B2. Overview of open-ended tasks across software categories. Each category includes two software applications and four tasks designed to elicit natural and diverse user behaviors.

C. User Behavior Taxonomy

To effectively assist users, an agent must understand not just *what* the user is doing (e.g., clicking a mouse), but *why* they are doing it and what their current cognitive and behavior state is. We introduce a hierarchical taxonomy of 9 user behavior states, organized into four high-level phases of the software workflow: **Planning**, **Execution**, **Problem-Solving**, and **Evaluation**. Table C3 provides detailed definitions and examples for each state.

Behavior State	Description	Examples
Planning		
Task Understanding and Preparation	The user is focused on the logistics of the task. This includes interpreting the task, gathering necessary digital assets, and configuring the software environment. Their goal is to set up the conditions needed to begin the work.	Reading task instructions, opening required software/files/templates, arranging workspace (resizing windows, organizing directories), downloading images for photo editing.
Ideation and Planning	The user is engaged in high-level conceptual work. They are brainstorming ideas, outlining the structure of the outcome, or creating a plan for how to approach the task. This often involves creating preliminary, non-final content that serves as a guide.	Formulating high-level strategy, creating step lists, sketching rough layouts or wireframes. The output is a plan or outline, not the final polished product.
Execution		
Exploration and Decision-Making	The user experiments with different options or features to understand their effects and decide which one to use. This exploratory phase involves deliberate trial and comparison, often pausing forward progress to evaluate alternatives.	Applying effects and undoing them, hovering over tools to see what they do, testing multiple font sizes to decide which fits best.
Performing Actions	The user is confidently using the software to make progress on the task. These actions are purposeful and executed with little hesitation.	Typing/deleting text, inserting and resizing images, applying formatting with clear intent, searching for functions to use.
Problem-Solving		
Frustration	The user encounters a blocker and shows signs of being stuck, confused, or annoyed. The system may not behave as expected, or the user cannot find a way to perform a desired action, leading to repetitive or unproductive behavior.	Sighing, pausing for long periods, undoing repeatedly, clicking unresponsive elements, complaining about slow system behavior.
Debugging	The user moves beyond frustration and begins to actively investigate the cause of a problem. They form and test hypotheses to diagnose and fix an issue.	Testing alternative approaches, undoing recent actions step by step, forming hypotheses about causes, adjusting settings to identify errors.
Seeking External Help	The user recognizes a gap in their own knowledge and turns to an external resource for assistance or procedural guidance.	Switching to a web browser for solutions, opening tutorials/documentation, consulting AI assistants or colleagues, posting questions in forums.
Evaluation		
Waiting and Monitoring	The user is in a passive state, waiting for a system-controlled process to complete before continuing their work. They are unable to take meaningful action and typically observe progress indicators.	Watching loading bars or spinners, waiting for exports or rendering to complete.
Assessment	The user intentionally pauses their work to review and evaluate their output. They examine the result for quality, accuracy, or aesthetics.	Zooming in/out to inspect fine details, replaying video snippets for review, comparing results to reference images or previous versions.

Table C3. Taxonomy of user behavior states in open-ended GUI workflows.

C.1. Behavior State Distribution

Figure C3 illustrates the overall distribution of user behavior states across the four high-level phases defined in our taxonomy. Table C4 provides a granular breakdown of these states across the specific evaluation tasks.

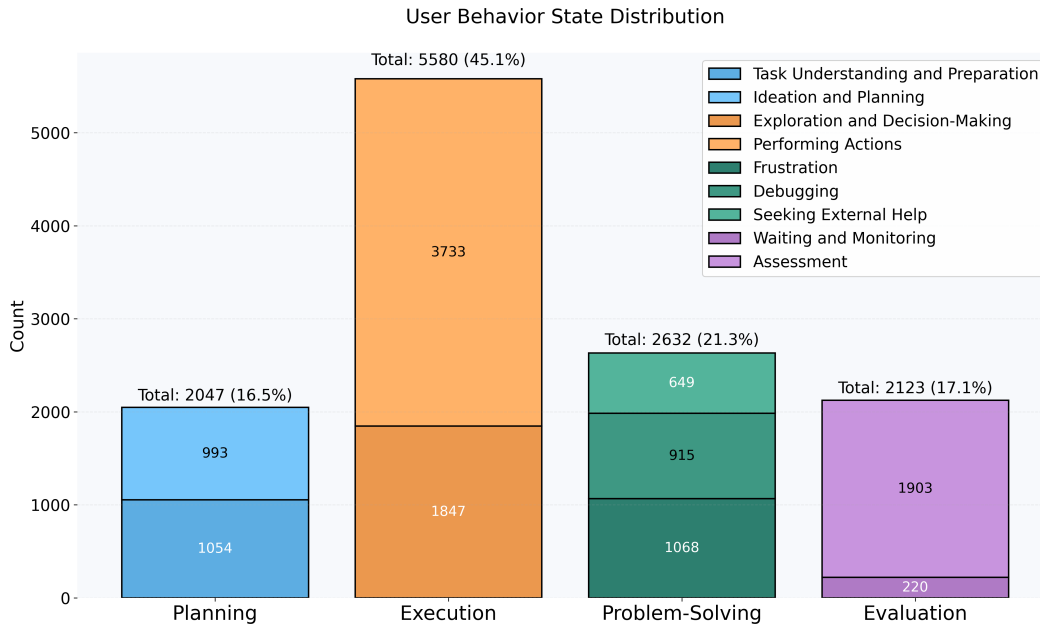


Figure C3. Distribution of user behavior states across Planning, Execution, Problem-Solving, and Evaluation phases across the videos in the dataset.

Table C4. Distribution of behavior state labels across the full dataset and specific evaluation tasks. Note that annotated instances used in the evaluation tasks may involve two or more states (e.g., a single segment containing both *Debugging* and *Seeking External Help*). Behavior State Detection uniformly sampled 200 instances from each class.

Behavior State	Dataset		Intent Prediction		Help Need Detection		Help Content Prediction	
	Count	(%)	Count	(%)	Count	(%)	Count	(%)
<i>Planning</i>								
Task Understanding and Preparation	1054	8.51%	216	9.85%	103	5.65%	36	2.82%
Ideation and Planning	993	8.02%	282	12.86%	84	4.61%	41	3.21%
<i>Execution</i>								
Exploration and Decision-Making	1847	14.92%	289	13.18%	162	8.89%	114	8.92%
Performing Actions	3733	30.15%	697	31.80%	474	26.00%	228	17.84%
<i>Problem-Solving</i>								
Frustration	1068	8.63%	131	5.98%	416	22.82%	415	32.47%
Debugging	915	7.39%	103	4.70%	259	14.21%	252	19.72%
Seeking External Help	649	5.24%	114	5.20%	85	4.66%	83	6.49%
<i>Evaluation</i>								
Waiting and Monitoring	220	1.78%	33	1.51%	17	0.93%	9	0.70%
Assessment	1903	15.37%	327	14.92%	223	12.23%	100	7.82%

C.2. Error Analysis: Behavior State Detection

Figure C4 presents the normalized confusion matrix for Behavior State Detection (Sec. 4.3.1). The results reveal a critical limitation in current MLLMs: a systemic bias toward interpreting interactions as productive execution while failing to recognize signs of struggle or hesitation. While models achieve reasonable accuracy for visually distinct states like *Seeking External Help* (0.61) and *Performing Actions* (0.57), they show near-zero capability in detecting *Frustration* (0.07) and *Debugging* (0.04). Instead, these negative states are overwhelmingly misclassified as *Performing Actions* (39% and 43%, respectively) or *Exploration and Decision-Making* (31% and 29%). This suggests that models perceive the visual activity of a struggling user—such as repeated clicking or rapid mouse movements—as deliberate progress, lacking the temporal understanding to distinguish between trial-and-error and confident execution.





Figure C4. Normalized confusion matrix for user behavior state classification. The most common errors occur when *Frustration* or *Debugging* is misclassified as *Performing Actions* or *Exploration and Decision-Making*.

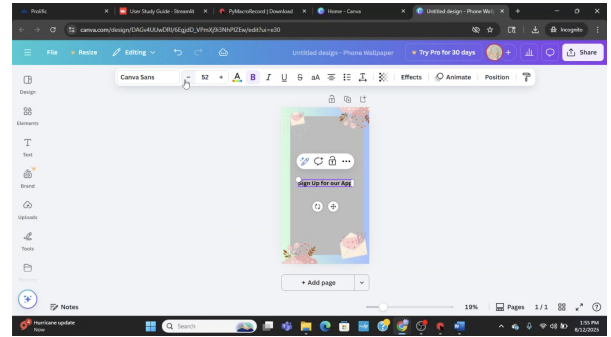
D. Screen Recording Video Examples

We present qualitative examples to illustrate the richness of the multimodal data in GUIDE.

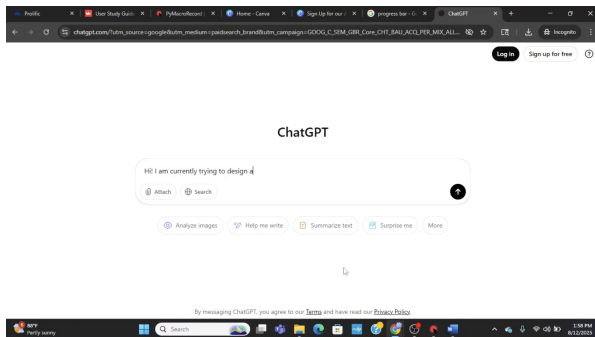
Software: Canva, Task: Design a mobile sign-up screen for a fictional app



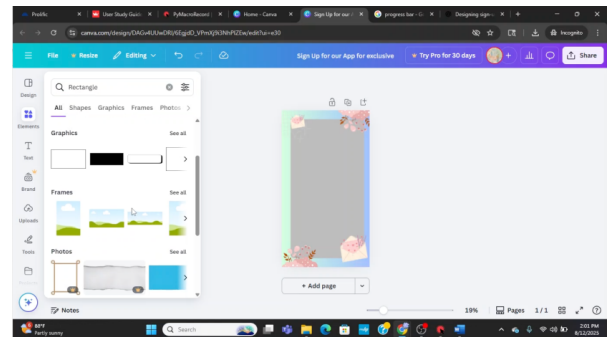
(4:11) “Nope, that’s not what I wanted to do. Try again. Alright, let’s do a text box.”



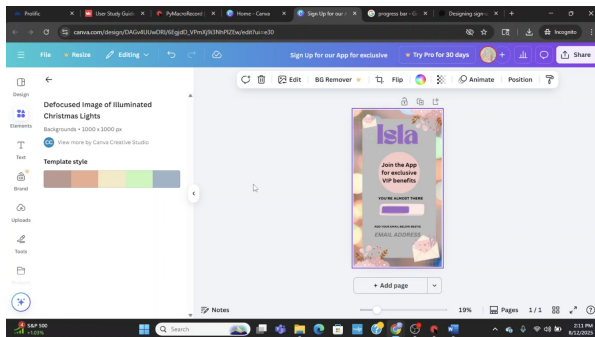
(4:40) “And we need to make this much smaller so it fits there at the top.”



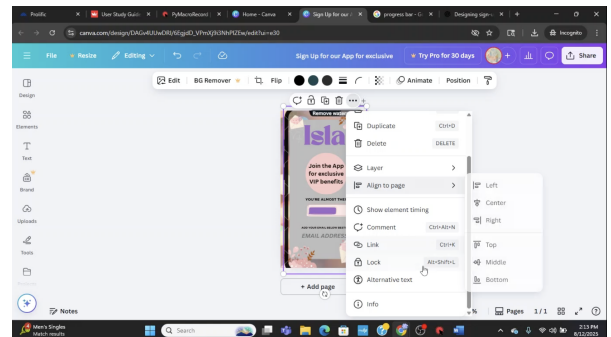
(6:50) “Progress bar for... what? I don’t know. Just do progress bar. But you know what? Let’s try ChatGPT because maybe they can help us.”



(10:07) “Oh, that’s cool, okay. So I can create the progress bar using the free elements with the shapes. Alright, so let’s try and do that. Alright. Elements.”



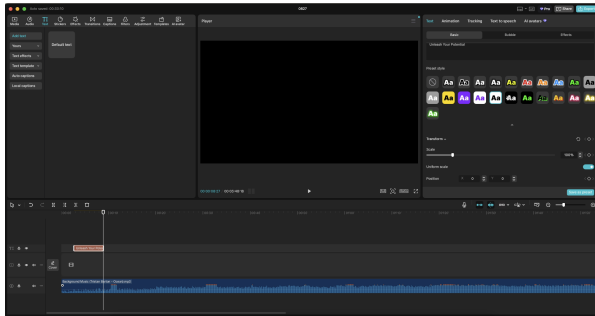
(21:01) “Ooh, what is this? Ooh, I like that so much better. All right, we’re gonna keep this. Yes, we’re gonna keep this.”




(22:05) “Not necessarily sure how I can... isn’t there a way to like make it a certain size? How do I do that?”

Table D5. Example video illustrating the user’s on-screen actions accompanied by think-aloud narration.

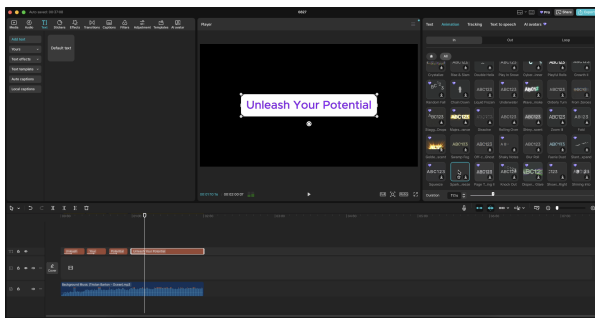
Software: CapCut, Task: Design a creative intro using animated text.



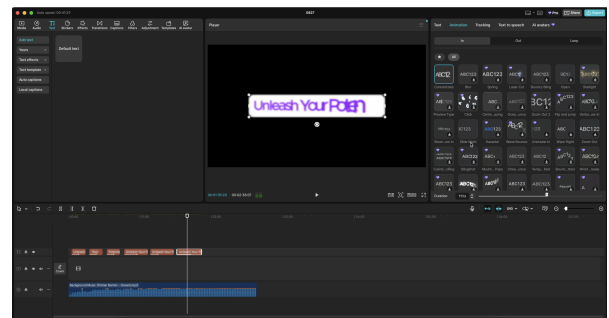
(4:56) “but I want to make it a bit more dynamic so all right.”



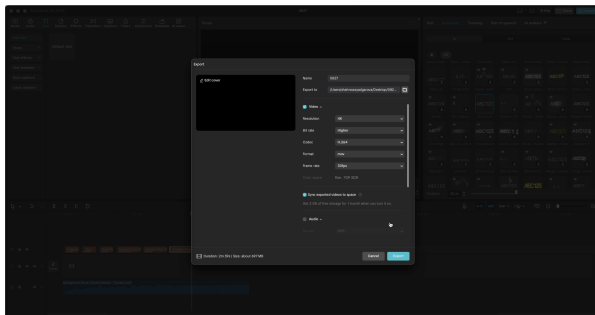
(7:55) “oh I think here it should be only one word appearing at a time.”




(11:39) “there are so many effects that I get a little bit too overwhelmed to see so many.”



(17:55) “I would like to add additional different animation for this one so I like to keep this.”



(18:42) “I am satisfied with the results so I will export this.”



(20:14) “Maybe I would like to make it a bit slower.”

Table D6. Example video illustrating the user’s on-screen actions accompanied by think-aloud narration.

E. Benchmark Task Examples

E.1. Behavior State Detection

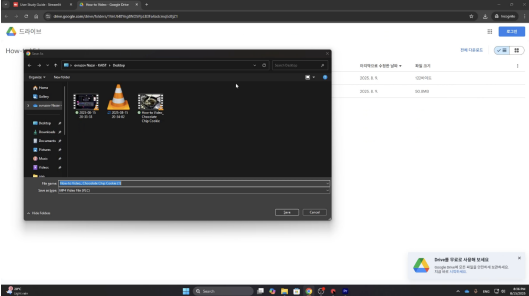
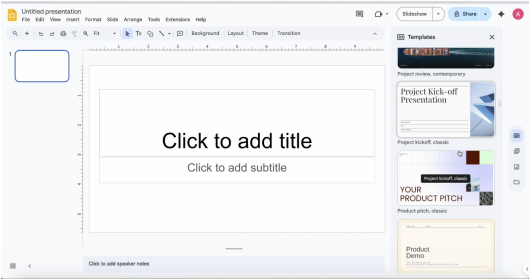
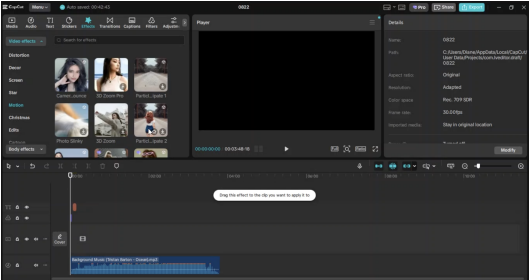
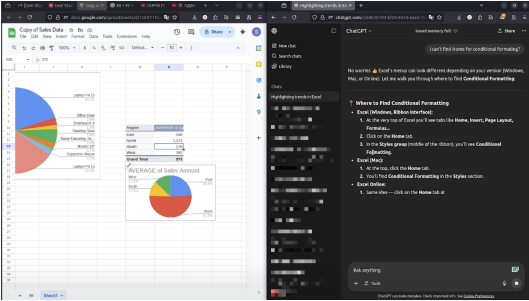
Screenshot	User Behavior State
	<p>Software: Premiere Pro</p> <p>Task: Edit a short instructional video to clearly guide a process.</p> <p>Behavior State: Task Understanding and Preparation</p> <p>The user is preparing their digital workspace before starting the editing task. They locate the necessary video file on their desktop and delete a superfluous 'test' file to prevent confusion.</p>
<p><i>“Okay, I downloaded it already. Delete my test, so I don’t get confused. I have the video.”</i></p>	
	<p>Software: Google Slides</p> <p>Task: Create a product pitch deck highlighting a product’s key features.</p> <p>Behavior State: Exploration and Decision-Making</p> <p>The user is actively browsing and comparing different templates, as shown by the scrolling and hovering behavior. The narration (‘This one looks good’) confirms they are evaluating options to make a final decision.</p>
<p><i>“I would like to just use this design or the white some minimalistic like iOS design. Oh, this one. This one looks good. Okay, let’s just...”</i></p>	
	<p>Software: CapCut</p> <p>Task: Design a creative intro using animated text.</p> <p>Behavior State: Frustration</p> <p>The user verbally expresses confusion (‘that’s strange’) after the software behaved in an unexpected way. They are momentarily paused, indicating a blocker in their workflow before they decide on a new course of action.</p>
<p><i>“Okay, that’s strange. That’s very strange, honestly.”</i></p>	
	<p>Software: Google Sheets</p> <p>Task: Summarize and visualize product sales by category or region.</p> <p>Behavior State: Seeking External Help</p> <p>The user is unable to find a feature and turns to ChatGPT for assistance. They type a question clarifying their problem, wait for the response, and then read the provided instructions.</p>
<p><i>(no narration)</i></p>	

Table E7. Example instances for the (1) User Behavior State Detection task, showing screenshots, think-aloud narration, and the corresponding behavior state.

E.2. Intent Prediction

Screenshot	Intent
	<p>Software: Canva</p> <p>Task: Design a mobile sign-up screen for a fictional app.</p> <p>Intent:</p> <ul style="list-style-type: none"> A: Rename the design file to reflect the new project B: Add the required input fields to the design C: Search for a suitable illustration to use as a header D: Resize the canvas to a custom dimension
<p><i>“So now that I have the frame as a design base, I need to include the input field for name, email.”</i></p>	
	<p>Software: Excel</p> <p>Task: Design a Gantt chart for a mini project.</p> <p>Intent:</p> <ul style="list-style-type: none"> A: Adjust the end date of the chart’s horizontal axis B: Adjust the date interval of the chart’s horizontal axis C: Reverse the order of the chart’s vertical axis D: Adjust the start date of the chart’s horizontal axis
<p><i>“okay looks perfect, I need to adjust the end date as well”</i></p>	
	<p>Software: Premiere Pro</p> <p>Task: Transform a long video into a short-form clip.</p> <p>Intent:</p> <ul style="list-style-type: none"> A: Create a new text layer above the existing video track B: Add an image to a specific empty slot in the timeline C: Apply a transition effect to the end of a video clip D: Add a video clip to the end of the current sequence
<p><i>“When this slot comes, we should put some kind of image here.”</i></p>	
	<p>Software: PowerPoint</p> <p>Task: Create a product pitch deck highlighting a product’s key features.</p> <p>Intent:</p> <ul style="list-style-type: none"> A: Align the logos with the main text boxes. B: Delete the logos from all the slides. C: Duplicate the logos onto the remaining slides. D: Change the color of the logos on all slides.
<p><i>“Paste, paste, paste, paste. Done. Done.”</i></p>	

Table E8. Example instances for the (2) Intent Prediction task, showing screenshots, think-aloud narration, and the corresponding intent.

E.3. Help Prediction

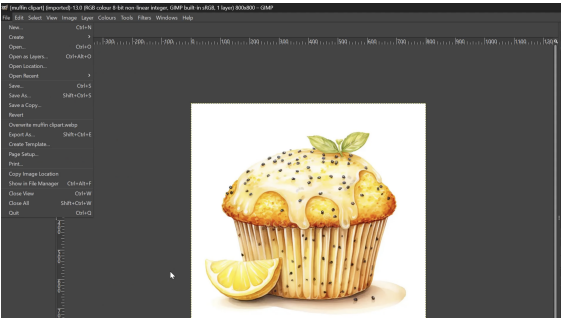
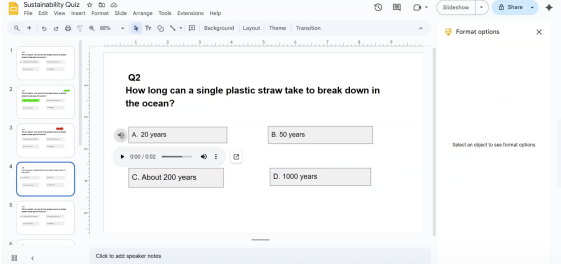
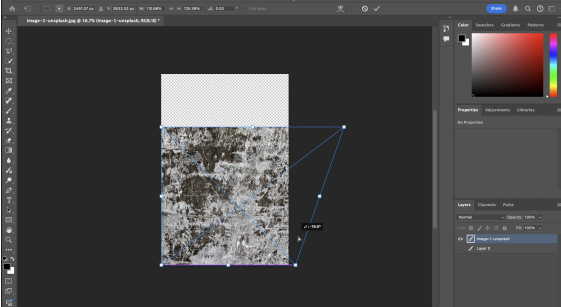
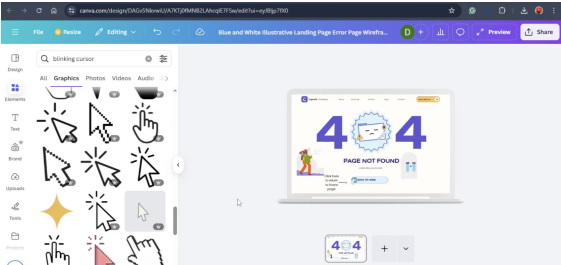
Screenshot	Help
	<p>Software: GIMP Task: Create a bakery logo with a warm, friendly identity. Help Content: A: how to add another image as a layer B: find the tool to add text C: remove the image background D: add a background color or shape</p>
<p><i>“Where could I insert the text? [...] I’m just going to, because the help function I don’t quite understand, but I can see if I can add it. Find it in Google.”</i></p>	
	<p>Software: Google Slides Task: Create a quiz deck with multiple-choice questions testing sustainability facts Help Content: A: align the answer choice boxes B: how to create a quiz slide template C: how to fix a self-identified audio related error D: add animation to reveal the correct answer</p>
<p><i>“I think I made a mistake here and I need to rectify this.”</i></p>	
	<p>Software: Photoshop Task: Create a composite from two images. Help Content: A: how to use the perspective or warp transform tools B: center the new layer on the canvas C: how to use layer blend modes D: maintain aspect ratio while scaling</p>
<p><i>“I’ll scale it. I just want to scale this up. How do I keep it?”</i></p>	
	<p>Software: Canva Task: Design a custom 404 error page with a visual and animated element. Help Need: A: help needed B: no help needed</p>
<p><i>“So I believe this is, this is great. I believe it’s just simple.”</i></p>	

Table E9. Example instances for the (3) Help Prediction task. For the Help Need Detection task, the top three instances illustrate cases labeled as *help needed*, while the last row shows an instance labeled as *no help needed*.

F. Software Task Outcome Examples

Figure F5 presents final artifacts produced by participants from the study. These examples highlight the open-ended nature of the assigned tasks. Despite receiving identical high-level instructions—such as “Design a poster for a music festival” or “Create a friendly bakery logo”—users produced markedly different results in terms of layout, aesthetic style, and complexity. This diversity confirms that the study elicited non-linear, creative workflows rather than fixed execution.



(a) Music event poster design in **Canva** (top) and **Figma** (bottom).

(b) Bakery logo design in **GIMP** (top) and **Photoshop** (bottom).

Figure F5. Example outcomes of the assigned tasks. The diversity across outputs reflects the open-ended nature of the tasks.

G. Human Verification Interface

The following is a video of the user working on Figma.

The user is performing the following task: **Design a mobile sign-up screen for a fictional app.**

4/68

Timestamp: 00:37 ~ 00:57 ✓ Watched

User Narration
"um i don't know what i love red red's my favorite color"

Label
Exploration and Decision-Making

Reasons for the label
The user opens the color picker and experiments with multiple shades of red. Their actions of clicking different points in the color palette, combined with the narration 'I don't know what I'm...!', show they are actively exploring visual options to decide on the final background color, rather than executing a predetermined choice.

Cancel Confirm

Figure G6. Annotation interface for validating and refining LLM-generated behavior-state labels. Annotators reviewed the predicted labels and the associated reasoning, correcting them if inaccurate. Each video segment was independently verified by two external annotators.

Watch the video and select the correct label.

Video Narration:
It looks like there are a lot of... a lot of places where it becomes silent, so...

Select the correct label:

Task Understanding and Preparation
The user is focused on the logistics of the task. This includes interpreting the task, gathering necessary digital assets, and configuring the software environment. Their goal is to set up the conditions needed to begin the work.

Ideation and Planning
The user is engaged in high-level conceptual work. They are brainstorming ideas, outlining the structure of the outcome, or creating a plan for how to approach the task. This often involves creating preliminary, non-final content that serves as a guide.

Exploration and Decision-Making
The user experiments with different options or features to understand their effects and decide which one to use. This exploratory phase involves deliberate trial and comparison, often pausing forward progress to evaluate alternatives.

Figure G7. Before participating in the annotation, annotators completed a quiz phase where they had to correctly classify example video segments. This process ensured that all annotators possessed a solid understanding of the behavior taxonomy and definitions.

H. Prompts

H.1. Taxonomy of User Behavior State Generation

Taxonomy of User Behavior State Generation

```
# Goal: Create a comprehensive taxonomy of user mental and behavior states by analyzing the video recording
↳ and transcript. Integrate visual observations (screen interactions, UI changes, cursor behavior) with
↳ audio/verbal cues (tone, hesitations, verbal expressions) and transcript content (exact quotes, semantic
↳ meaning).

# Analysis Guidelines:
- VISUAL EVIDENCE: Describe what you see on screen (tool selections, menu interactions, visual feedback,
↳ cursor patterns)
- AUDIO EVIDENCE: Note tone of voice, hesitations, exclamations, and vocal expressions
- TRANSCRIPT EVIDENCE: Extract precise quotes that reveal mental states and intentions
- CROSS-REFERENCE: Connect visual actions with verbal expressions to understand user intent and mental state

# Output Format (return as JSON):
{{
  "taxonomy": [
    {{
      "label": "...",
      "definition": "...",
      "evidence": [
        {{
          "timestamp": "00:01:32",
          "modality": "visual",
          "description": "User clicks on the brush tool and immediately switches to eraser",
          "significance": "Indicates uncertainty or trial-and-error behavior"
        }},
        {{
          "timestamp": "00:01:35",
          "modality": "audio",
          "description": "User says 'hmm, that's not right' with a frustrated tone",
          "significance": "Verbal confirmation of confusion and frustration"
        }},
        {{
          "timestamp": "00:01:35",
          "modality": "transcript",
          "quote": "hmm, that's not right, let me try something else",
          "significance": "Shows problem-solving mindset and willingness to iterate"
        }}
      ]
    }}
  ]
}}

# Context:
Software: {SOFTWARE}
Task: {TASK_NAME}

# Transcript:
{TRANSCRIPT_JSON}

# Video Content:
SEE THE ATTACHED FILE.
```

Figure H8. Prompt to generate a taxonomy of user behavior states given demonstration videos.

H.2. Data Annotation

Behavior State Annotation

```
# Goal
You are given a screen recording video snippet of a user using the software {SOFTWARE}. Annotate the video
↔ with the provided taxonomy of user mental and behavior states. Include the taxonomy label and reasoning
↔ that explains why the label is appropriate for the video.

# Instructions
1. Your annotation label must be based on the user's current, on-screen behavior shown in the video.
2. Annotate video based on what you see, but if the label is not clear, use the think-aloud narration as
↔ auxiliary data to understand the user's intent or thought process.
3. Be aware that the user's narration may refer to past actions or future plans. Always align your annotation
↔ label with the user's current behavior at that specific moment in the video.

# Output Format (JSON)
{{
  "label": "...", // one of the labels in the taxonomy,
  "reasoning": "...", // Explanation on why the label is appropriate for the video,
}}

# Full Video Context
Software: {SOFTWARE}
Task the user is performing in the full screen recording video: {TASK_NAME}
The full transcript of the user: {TRANSCRIPT_JSON}

# Target Video Snippet Context
- Video snippet time range:
{TRANSCRIPT_JSON[narration_index]["start"]} - {TRANSCRIPT_JSON[narration_index]["end"]}
Narration sentences of the video snippet:
"{TRANSCRIPT_JSON[narration_index]["sentence"]}"

# Taxonomy
{TAXONOMY}

# Video Content
SEE THE ATTACHED FILE.
```

Figure H9. Prompt to annotate a given video segment based on the taxonomy of user behavior states.

Intent Annotation

You are analyzing a user's screen recording video and think-aloud narration while they use software (e.g.,
↪ video editing, design, or spreadsheet tools).

Goal

For the given video snippet and its corresponding narration segment, infer what the user was aiming to
↪ complete or achieve by the end of this segment - their short-term goal or intention.

- The goal should represent **an outcome or result** that the user was either finishing or actively working on
↪ as the segment ends.
- Focus on **tangible, result-oriented goals** (e.g., "finish trimming the clip," "adjust the image color,"
↪ "complete text alignment").
- Ignore interface-level or purely procedural descriptions (e.g., "click this," "open that," "drag the
↪ layer").
- If no clear goal or outcome is expressed or shown, return "no tangible goal".

Output Format (JSON)

```
{  
  "original_narration": "<verbatim narration text>",  
  "goal": "<concise description of what the user aimed to complete or was completing by the end of this  
↪ segment, or 'no tangible goal'>",  
  "evidence_narration_snippet": "<the exact portion of the narration text that supports this inferred  
↪ goal>",  
  "reasoning": "<brief explanation of how this goal was inferred based on the narration and video>"  
}
```

Full Video Context

Software: {SOFTWARE}
Task the user is performing in the full screen recording video: {TASK_NAME}
The full transcript of the user: {TRANSCRIPT_JSON}

Target Video Snippet Context

- Video snippet time range: {TRANSCRIPT_JSON[narration_index]["start"]} -
↪ {TRANSCRIPT_JSON[narration_index]["end"]} seconds
- Narration during this snippet: "{TRANSCRIPT_JSON[narration_index]["sentence"]}"

Video Content

SEE THE ATTACHED FILE.

Figure H10. Prompt used to annotate a given video segment with the user's intent.

Help Annotation

You are analyzing a user's screen recording video and think-aloud narration while they use software (e.g., video editing, design, or spreadsheet tools).

Goal

For the given video snippet and its corresponding narration segment, infer **what kind of help or guidance the user is looking for**.

- Focus on identifying explicit requests for help or moments where the user seeks information, clarification, or suggestions.
- Help-seeking can appear in two main forms:
 - 1) **On-Screen Help Behavior**: The user performs on-screen actions to seek help (e.g., opening a web browser, typing a query into a search engine or LLM, viewing online documentation or tutorials).
 - 2) **Narration-Based Help Expression**: The user verbally expresses confusion, uncertainty, or asks questions (e.g., "I don't know how to fix this," "Why is this not showing up?", "How do I do this?").
- Ignore casual statements or comments unrelated to problem-solving.
- If there is no clear indication that the user is seeking help, or if the type of help they need is implicit or ambiguous, return "no help needed".

Output Format (JSON)

```
{
  "help_needed": "<concise but meaningful description of the help sought - focus on the underlying need,
  such as 'explain masking feature', 'suggest alternative filter', 'debug export error', or 'clarify
  timeline snapping'",
  "help_source": "<'screen', 'narration', 'both', or 'none'>",
  "evidence_narration_snippet": "<exact portion of narration that indicates help-seeking (if any)>",
  "evidence_screen_behavior": "<description of what was seen on screen that indicates help-seeking (if
  any)>",
  "reasoning": "<brief explanation of how the need for help was inferred based on the narration or/and
  on-screen behavior>"
}
```

Full Video Context

Software: {SOFTWARE}

Task the user is performing in the full screen recording video: {TASK_NAME}

The full transcript of the user: {TRANSCRIPT_JSON}

Target Video Snippet Context

- Video snippet time range: {TRANSCRIPT_JSON[narration_index]["start"]} {

↪ {TRANSCRIPT_JSON[narration_index]["end"]} seconds

- Narration during this snippet: "{TRANSCRIPT_JSON[narration_index]["sentence"]}"

Video Content

SEE THE ATTACHED FILE.

Figure H11. Prompt used to annotate a given video segment with whether help is needed and, if so, what specific help is required.

Filtering On-Screen Help-Seeking Behavior

You are analyzing a user's sequence of help-seeking actions while they use {SOFTWARE} for the task
↳ "{TASK_NAME}".

You are given HELP_DATA, a list of chronological records where each record includes:

- index: original index of the record
- help: user's inferred help need (e.g., "find a graphic for a password field")
- screen_behavior: observed on-screen action
- narration: user's think-aloud narration

Goal

- Keep only the help entries where the screen_behavior is about using external applications (e.g., Google Search, ChatGPT, Gemini, YouTube, etc.). Note that it doesn't include referring to the task instructions page.
- Return the list of segment ids that should be kept as a JSON object.

Output Format (JSON)

Return only valid JSON:

```
{  
  "kept_segment_ids": [<int>, <int>, ...]  
}
```

Data

```
HELP_DATA:  
{HELP_JSON}
```

Figure H12. Prompt used to filter on-screen help-seeking behavior from segments previously marked as help needed.

Filtering Narration-Based Help-Seeking Behavior

You are analyzing a user's sequence of actions while they use {SOFTWARE} for the task "{TASK_NAME}".

You are given HELP_DATA, a list of chronological records where each record includes:

- index: original index of the record
- help: user's inferred help need (e.g., "find a graphic for a password field")
- screen_behavior: observed on-screen action
- narration: user's think-aloud narration

Goal

Keep only segments where the narration explicitly asks for help or guidance.

Definition of explicit help

The narration contains a direct request for help or instruction, for example:

- "help me", "i need help", "i want help", "can you help", "please help"
- How or where questions about operating the software, such as:
 - "how do i ...", "how can i ...", "how to ...", "where is ...", "which option should i ...",
 - "what should i click", "what does this do".
- Requests for instructions or explanation:
 - "show me how to ...", "tell me how to ...", "could someone explain ...", "is there a way to ..."

Exclude the following

- Exploration or intent without a help request: "i'm going to try", "let me see", "i will search"
- Trial and error or self-correction without a request: "no, not that", "okay now i got it", "ah ok",
 - ↔ "finally"
- Uncertainty alone: "maybe", "i think", "not sure" unless followed by a direct question that asks for
 - ↔ guidance
- Statements addressed to self that do not ask for help: "i need to add text", "i'm looking for an icon"
- Generic questions not tied to getting guidance on what to do next

Output Format

Return only valid JSON:

```
{  
  "kept_segment_ids": [<int>, <int>, ...]  
}
```

Data

```
HELP_DATA:  
{HELP_JSON}
```

Figure H13. Prompt used to filter narration-based help-seeking behavior from segments previously marked as help needed.

Filtering No Help Needed

You are analyzing a user's sequence of actions while they use {SOFTWARE} for the task "{TASK_NAME}".

You are given NO_HELP_DATA, a chronological list of records where each record includes:

- index: original segment index
- start_time, end_time: time range of the segment
- no_help_reasoning: explanation for why the user does not need help
- narration: the user's think-aloud narration

Goal

Identify up to 5 segments where it is **explicitly clear** that the user does not need help.

These are moments when the user:

- Performs actions confidently, smoothly, and intentionally.
- Demonstrates clear understanding of what to do next without hesitation or correction.
- Speaks in a calm, matter-of-fact tone (e.g., "Now I'll add text here," "Perfect," "That looks good.").

Exclude the following:

- **Trial and error:** any sign of experimentation, correction, or rapid alternation (e.g., "No, no, no," "Let me try again," "Okay, that worked.")
- **Self-resolution after confusion:** phrases like "now I got it," "finally," "oh, that's how," "ah okay," "I see," or any narration showing realization after failure or surprise.
- **Frustration or emotional reactions:** (e.g., "oh shit," "ugh," "why," "come on") even if followed by success.
- **Uncertain or exploratory speech:** "I think," "maybe," "let's see," "try," "not sure."
- **Segments that merely lack confusion** but do not clearly express confidence or mastery.

Rules

- Select at most 5 segments that best reflect calm, deliberate, fluent progress.
- Preserve the original order of the selected segments.
- Return only the segment indices of the kept entries.

Output Format (JSON)

Return only valid JSON:

```
{  
  "kept_segment_ids": [<int>, <int>, ...]  
}
```

Data

```
NO_HELP_DATA:  
{NO_HELP_JSON}
```

Figure H14. Prompt used to filter clear no-help-needed segments.

H.3. Model Evaluation

Behavior State Detection

```
# Goal
You are given a screen recording video snippet of a user working in {SOFTWARE}.
Classify the video into one of the labels from the provided taxonomy of user mental and behavioral states.
Include both the taxonomy label and a reasoning that explains why this label best fits the observed segment.

# Instructions
1. Base your classification on the user's on-screen behavior shown in the video.
2. Provide:
  - label: one of the taxonomy labels
  - reasoning: a concise explanation of why this label fits the segment
3. Return the output strictly in valid JSON, with keys "label" and "reasoning". Do NOT wrap the JSON in
  ↪ markdown code blocks. Return only the raw JSON object.

# Output Format (JSON)
{{
  "label": "...",
  "reasoning": "..."
}}

# Video Context
Software: {SOFTWARE}
Task performed in the full recording: {TASK_NAME}
Start and end times of the snippet (relative to the full recording):
{start_time} - {end_time} seconds

# Taxonomy Descriptions
{TAXONOMY}

# Previous Segment Context (** Optional based on the condition)
The user behavior in the immediately preceding segment was {previous_label}: {label_definition}

# Video Content
SEE THE ATTACHED FILE.
```

Figure H15. Prompt used to evaluate the model on the (1) Behavior State Detection task.

Intent Prediction

```
# Goal
You are given a screen recording snippet of a user working in {SOFTWARE}.
Predict which option best describes the user's intention during this segment.

# Instructions
1. Watch the video and analyze on-screen actions.
2. Select the option (A-D) that best matches the goal of the user trying to achieve.
3. Use the provided behavior context to interpret the goal. (** Optional based on the condition)
4. Return output in JSON:
  - label: one of A-D
  - reasoning: short explanation for your choice
5. Output only a valid JSON object (no Markdown).

# Output Format (JSON)
{"label": "A", "reasoning": "..."}

# Video Context
Software: {SOFTWARE}
Task performed in the full screen recording: {TASK_NAME}
Start and end times of the snippet (relative to the full recording):
{start_time} - {end_time} seconds

# User Behavior Context (** Optional based on the condition)
The following user behavior is identified: {label}: {label_definition}.
Consider this context when predicting the user's intention and selecting the most appropriate option.

# Options
{options_text}

# Video Content
SEE THE ATTACHED FILE.
```

Figure H16. Prompt used to evaluate the model on the (2) Intent Prediction task.

Help Need Detection

```
# Goal
You are given a screen recording video snippet of a user working in {SOFTWARE}.
Based on the video content, the user's intention, and the user behavior context (** Optional based on the
↔ condition), determine if the user needs help or not in this segment.

# Instructions
1. Watch the video and observe the user's behavior and actions.
2. Consider the user's intention and behavior context provided to better understand what they are trying to
↔ accomplish and their current state. (** Optional based on the condition)
3. Determine if the user needs help or not.
4. Provide:
  - label: "yes" if the user needs help, "no" if they do not need help
  - reasoning: a concise explanation of why the user does or does not need help based on the observed
↔ behavior.
5. Output only a valid JSON object (no Markdown).

# Output Format (JSON)
{{
  "label": "yes" | "no",
  "reasoning": "..."
}}

# Video Context
Software: {SOFTWARE}
Task performed in the full screen recording: {TASK_NAME}
Start and end times of the snippet (relative to the full recording):
{start_time} - {end_time} seconds

# User Behavior Context (** Optional based on the condition)
The following user behavior is identified in order: {label}: {label_definition}.
Consider this behavior context when determining if the user needs help.

# User Intention (** Optional based on the condition)
The user's intention or goal during this segment is: {intention}
Consider this intention when determining if the user needs help to achieve this goal.

# Video Content
SEE THE ATTACHED FILE.
```

Figure H17. Prompt used to evaluate the model on the (3-1) Help Need Detection task.

Help Content Detection

```
# Goal
You are given a screen recording video snippet of a user working in {SOFTWARE}.
Based on the video content, the user's intention, and the user behavior context (** Optional based on the
↪ condition), predict which option best describes what kind of help or guidance the user is looking for
↪ during this segment.

# Instructions
1. Watch the video and predict what help the user might need at this segment.
2. Consider the user's intention and behavior context provided to better understand what they are trying to
↪ accomplish and their current state. (** Optional based on the condition)
3. Select the option (A, B, C, or D) that best matches what help the user needs to accomplish their intention
↪ given their behavior context.
4. Provide:
   - label: one of the option letters (A, B, C, or D)
   - reasoning: a concise explanation of why this option best fits the observed segment
5. Output only a valid JSON object (no Markdown).

# Output Format (JSON)
{{
  "label": "A",
  "reasoning": "...
}}

# Video Context
Software: {SOFTWARE}
Task performed in the full screen recording: {TASK_NAME}
Start and end times of the snippet (relative to the full recording):
{start_time} - {end_time} seconds

# User Behavior Context (** Optional based on the condition)
The following user behavior is identified in order: {label}: {label_definition}.
Consider this behavior context when predicting what help the user might need.

# User Intention (** Optional based on the condition)
The user's intention or goal during this segment is: {intention}

# Options
{options_text}

# Video Content
SEE THE ATTACHED FILE.
```

Figure H18. Prompt used to evaluate the model on the (3-2) Help Content Prediction task.