

# Sensing What Surveys Miss: Understanding and Personalizing Proactive LLM Support by User Modeling

Ailin Liu  
LMU Munich  
Munich, Germany  
Munich Center for Machine Learning (MCML)  
Munich, Germany  
ailin.liu@lmu.de

Fiona Draxler  
University of Mannheim  
Mannheim, Germany  
fiona.draxler@uni-mannheim.de

Yesmine Karoui  
LMU Munich  
Munich, Germany  
y.karoui@campus.lmu.de

Frauke Kreuter  
LMU Munich  
Munich, Germany  
University of Maryland  
College Park, Maryland, USA  
fkreuter@umd.edu

Francesco Chiossi  
LMU Munich  
Munich, Germany  
francesco.chiossi@ifi.lmu.de

## Abstract

Difficulty spillover and suboptimal help-seeking challenge the sequential, knowledge-intensive nature of digital tasks. In online surveys, tough questions can drain mental energy and hurt performance on later questions, while users often fail to recognize when they need assistance or may satisfy, lacking motivation to seek help. We developed a proactive, adaptive system using electrodermal activity and mouse movement to predict when respondents need support. Personalized classifiers with a rule-based threshold adaptation trigger timely LLM-based clarifications and explanations. In a within-subjects study (N=32), aligned-adaptive timing was compared to misaligned-adaptive and random-adaptive controls. Aligned-adaptive assistance improved response accuracy by 21%, reduced false negative rates from 50.9% to 22.9%, and improved perceived efficiency, dependability, and benevolence. Properly timed interventions prevent cascades of degraded responses, showing that aligning support with cognitive states improves both the outcomes and the user experience. This enables more effective, personalized large language model (LLM) support in survey-based research.

## CCS Concepts

• **Human-centered computing** → *Web-based interaction; Ubiquitous and mobile computing systems and tools*; **Empirical studies in interaction design**; **Interactive systems and tools**.

## Keywords

Proactive Assistance, Physiological Sensing, Personalization, User Experience, Realtime Support

### ACM Reference Format:

Ailin Liu, Yesmine Karoui, Fiona Draxler, Frauke Kreuter, and Francesco Chiossi. 2026. Sensing What Surveys Miss: Understanding and Personalizing Proactive LLM Support by User Modeling. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3772318.3791191>

## 1 Introduction

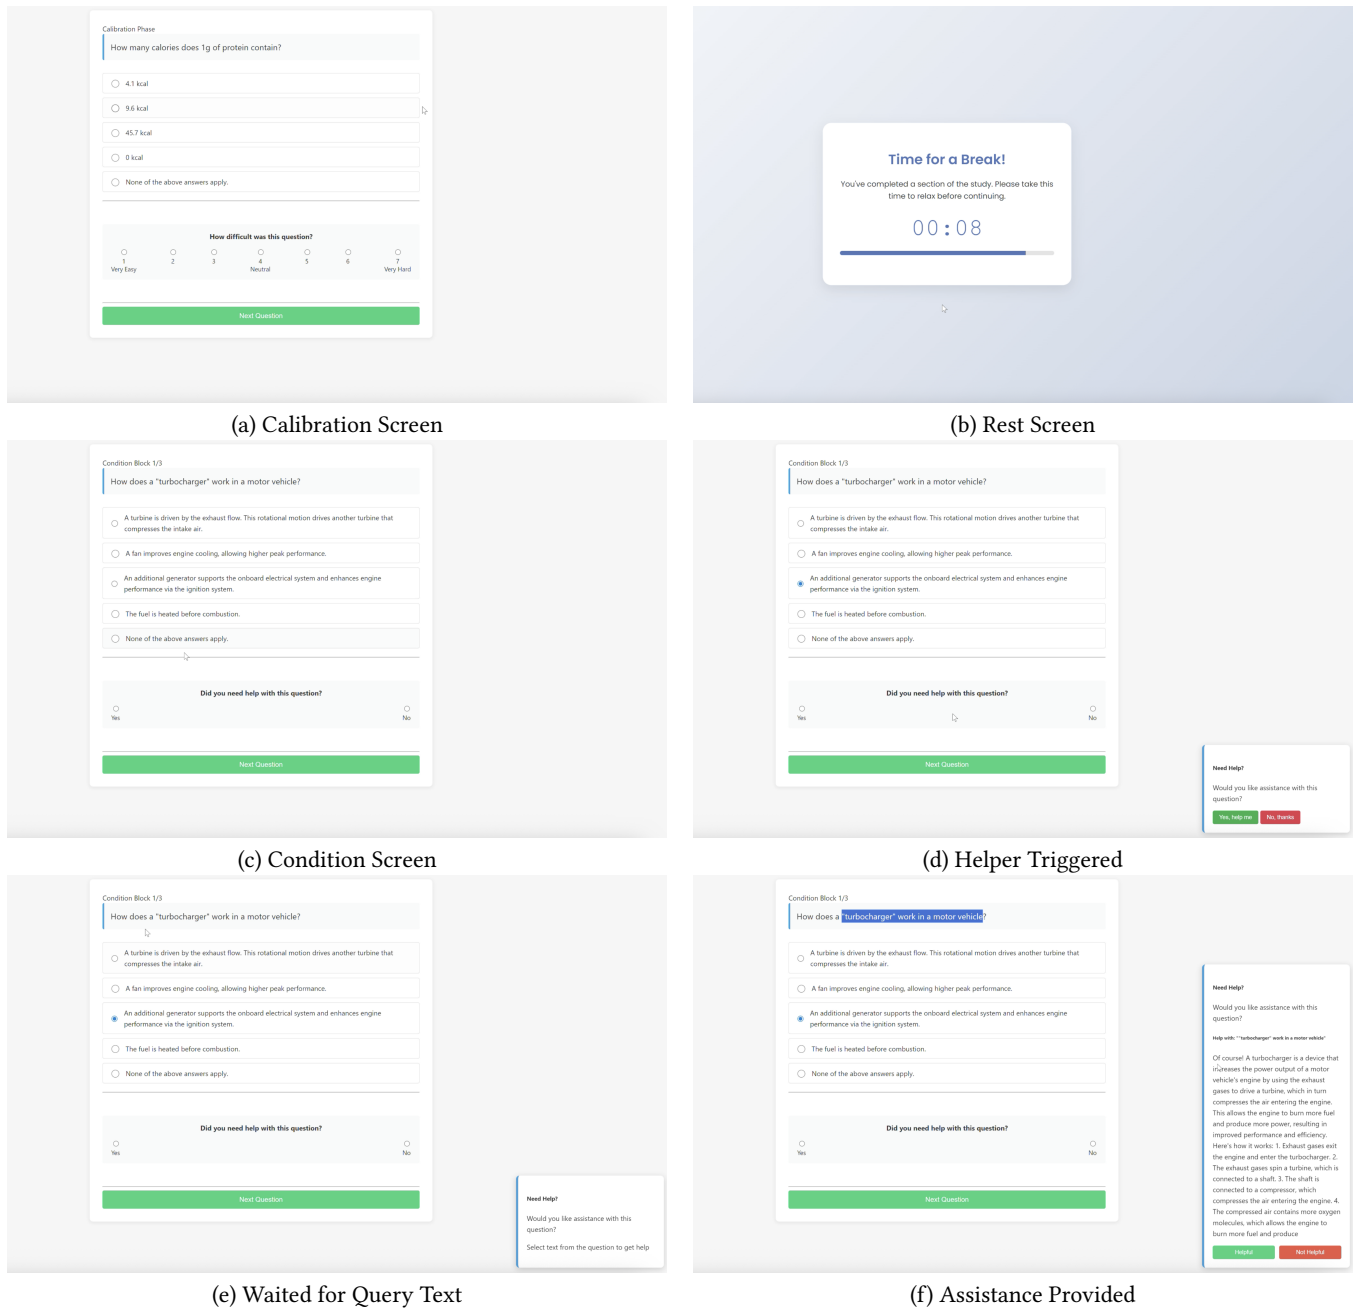
Surveys are a widely used method for collecting data across domains from product design and marketing to social research and academic feedback. Their ability to gather insights quickly, cost-effectively, and at scale makes them a core instrument for both researchers and practitioners. However, unlike interviewer-administered surveys, where trained professionals can detect hesitation, confusion, or disengagement and intervene in real time, self-administered web surveys provide no such adaptive human support. When respondents face cognitively demanding questions or feel pressured by time constraints, they may engage in satisficing behaviors, selecting suboptimal responses to reduce effort and ultimately compromising data quality [7, 70]. To mitigate these risks, survey methodologists have emphasized the importance of identifying problematic items and respondent difficulties, whether through pretesting before deployment, intervention during survey administration, or post-hoc detection of low-quality responses [21, 68].

This challenge is especially pronounced in surveys that include cognitive or general-knowledge items intended to assess participants' understanding or abilities. Such questions are widely used



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2278-3/26/04  
<https://doi.org/10.1145/3772318.3791191>



**Figure 1: Screenshots of the experimental interface: (a) calibration phase, (b) rest break, (c) condition task question, (d) adaptive helper triggered, (e) waiting for user selecting text, and (f) assistance provided by the system.**

in national social surveys, civic knowledge measures, and empirical research to measure comprehension [15, 29, 63, 79]. When respondents struggle with such questions, whether due to item complexity, ambiguous wording, or insufficient domain knowledge, the cognitive strain can have spillover effects on subsequent items, diminishing attentiveness and consistency [5]. While recent work suggests that the increase in perceived burden across items may be

modest in some contexts [48], even small increases can compound in ways that influence overall engagement and accuracy. Over time, this strain can manifest as satisficing behaviors (e.g., straightlining, skipping items) [7, 70] or early termination, ultimately reducing both the completeness and validity of collected data.

Survey methodology has investigated interactive support through requests for clarifications [18, 19, 74]. Meanwhile, LLMs make it

possible to deliver flexible, context-sensitive explanations in natural language [9, 12, 41].

However, existing approaches face fundamental limitations in their ability to provide timely, individualized support. Static personalization systems determine support strategies based on pre-task assessments (e.g., prior knowledge tests, demographic profiles) and apply the same intervention pattern to all items for a given user [18, 19, 74]. While these systems can tailor content to user characteristics, they cannot respond to within-session fluctuations in cognitive state, a critical gap given that respondent struggle varies moment-to-moment based on item difficulty, accumulated fatigue, and interference from prior items.

Reactive support systems wait for explicit user requests such as, help buttons and chat queries, before intervening [17, 20]. Previous research shows that respondents rarely request help on their own, often fewer than 10% use optional clarifications and respondents had often already formed incorrect interpretations by the time the system reacted [17, 18, 20].

Rule-based systems adjust based on fixed decision criteria (e.g., “offer help after 30 seconds of inactivity”), but these universal thresholds ignore individual differences in baseline behavior—what signals struggle for one person may be normal processing for another [40]. The core limitation across these approaches is their inability to detect and respond to the dynamic, personalized cognitive state of individual respondents in real time. When a respondent encounters an unexpectedly difficult item, existing systems either miss the moment of struggle entirely (static approaches), detect it too late (reactive approaches), or misfire by applying population-level heuristics (rule-based approaches). This timing gap is consequential: premature assistance interrupts productive problem-solving and reduces engagement, while delayed support arrives too late to prevent cascading effects on subsequent items [1, 71]. What is needed are systems that continuously monitor individual cognitive load through leading indicators—such as physiological signals that reflect arousal before behavioral degradation [8, 46], and trigger interventions precisely when each specific respondent begins to struggle with each specific item.

We present an adaptive survey assistance system that addresses this gap by combining real-time ubiquitous sensing with personalized, dynamically updating prediction models. Our system uses electrodermal activity (EDA) and mouse movement data to continuously assess respondent cognitive state, predicting moment-to-moment when respondents need help with knowledge-based survey items. Critically, we employ personalized classifiers with threshold adaptation that align with each respondent’s baseline patterns throughout the session. This enables the system to distinguish an individual’s typical processing behavior from genuine difficulty, triggering LLM-based clarifications and explanations only when cognitive load indicators suggest the respondent is struggling. Unlike static systems that apply predetermined support schedules or reactive systems that wait for explicit help requests, our approach provides proactive interventions timed to each individual’s actual cognitive state as it evolves across sequential items.

A within-subjects study (N=32) compared our aligned-adaptive (adjusted toward cognitive states) assistance timing against misaligned-adaptive (adjusted backward cognitive states) and random-adaptive control conditions across sequential survey items containing factual knowledge questions. Results demonstrate that physiological and behavioral signals can effectively predict when respondents need help. The adaptive assistance improves response accuracy from 41% to 62% and reduces missed assistance opportunities compared to the control conditions. Participants rated the aligned-adaptive system significantly higher on efficiency, dependability, and benevolence scales, while the aligned-adaptive condition achieved the highest acceptance rate, indicating positive user acceptance of real-time intervention.

Our findings highlight the critical role of timing in proactive assistance: interventions that arrive when respondents actually struggle yield greater benefits than those delivered too early or too late, or misaligned with individuals’ real-time need. By preventing small difficulties from compounding across items, aligned-adaptive timing helps preserve both response quality and user experience. This work provides a foundation for adaptive survey systems that maintain data quality through physiological and behavioral sensing, with applications that span educational assessments, healthcare questionnaires, and research surveys where the complexity of the item creates a cumulative burden on the respondent.

## 2 Related Work

We discuss related work on (1) support elements in surveys, (2) proactive assistance with LLMs, and (3) physiological and behavioral sensing for cognitive states and its adaptive applications.

### 2.1 Interactive or Intelligent Support in Survey Methodology

Survey researchers have long studied how task difficulty and motivation shape response quality. High task difficulty can lead to data issues, such as satisficing [7, 70], speeding [11] and item nonresponse [35]. Difficult questions also increase the risk of drop-out [37, 55]. Longer and more complex questions demand greater cognitive effort, which has been linked to higher breakoff rates even when other factors are controlled [64].

To address challenges in data quality and respondent experience, researchers have explored the use of interactive and pre-defined support within web surveys. Early studies focused primarily on reducing comprehension problems through system-initiated clarifications. These interventions, which included the use of hyperlinks to definitions, mouse-over explanations, and proactive messages triggered by long response times or pauses [17, 18, 20, 74], were found that respondents answered more accurately when compared to surveys without such support. However, these methods primarily addressed pre-defined clarity rather than behavioral issues.

More recent work has expanded beyond clarification to examine interactive feedback mechanisms designed to reduce satisficing, careless responding, and improve engagement. Studies have shown that providing respondents with feedback on speeding, reminders to read carefully, or commitment prompts encouraging conscientious answering can enhance response quality and improve accuracy by encouraging more thoughtful responses [16, 47]. These approaches

acknowledge that misunderstandings are common and cannot be eliminated through pretesting alone [19]. For example interactive probes can also elicit richer answers, especially from participants who are highly engaged with the survey topic [38].

Despite these successes, a limitation persists across both early and contemporary approaches: they predominantly rely on uniform, “one-size-fits all” assumptions about respondent behavior, and do not account for inter-individual or moment-by-moment variation. Most studies deploy static intervention triggers, such as a fixed response-time threshold to identify speeding, which fail to accommodate significant individual differences in reading speed, cognitive load, engagement, or interaction patterns. Help can be beneficial when offered at the right moment, but distracting or even harmful when mistimed [16]. This gap highlights the need for a more dynamic approach that tailors assistance to individual needs rather than relying on universal, predefined rules.

Recent work has introduced LLMs as a flexible tool for intelligent survey support. LLMs can generate alternative question phrasings or conversational prompts to maintain engagement [60, 85]. They can build glossaries or synonyms to lower cognitive barriers in domain-specific contexts [82]. They can also integrate contextual user data from online platforms to extend survey functionality [81]. Beyond LLMs, interactive question answering systems allow surveys to engage respondents in dialog, clarifying ambiguities, and improving accuracy [6]. These developments show the potential of intelligent, context-aware systems to enhance survey quality, but they mainly rely on textual interaction and static behavioral cues, leaving moment-to-moment cognitive variation unaddressed.

Our work complements and extends this line of research by incorporating ubiquitous and non-intrusive sensing to capture real-time changes in cognitive load, enabling personalized, proactive support that aims to be delivered precisely when respondents need it. This focus on adaptive timing and individualized state estimation fills a critical gap in existing approaches and illustrates how multimodal sensing can further advance intelligent survey interaction.

## 2.2 Proactive Assistance with LLMs

Proactive assistance systems aim to provide support before users explicitly ask for it. Early research introduced the idea of Just-In-Time Information Retrieval agents (JITIRs), which proactively present useful information based on local context in a non-intrusive way [69]. These systems reduce the cognitive effort of searching and encourage users to access information they might otherwise miss. Other work explored proactive agents that listen to ongoing conversations, detect key entities, and retrieve related information, thereby reducing the need for explicit search activity [2]. These studies established proactive support as a way to reduce mental effort and improve task performance.

Designing proactive AI assistants that deliver positive experiences remains challenging, since effective collaboration principles differ across tasks and contexts. Recent work has explored proactive assistance with LLMs across a range of domains, highlighting both the opportunities and challenges of this approach.

In care and wellbeing contexts, Liu et al. [56] developed ComPeer, a conversational agent for proactive peer support that detects significant emotional events in conversations and strategically

times interventions. Across 18 users over two weeks, they found that timing accounted for 40% of variance in intervention acceptance—identical suggestions were accepted three times more often when delivered at moments users perceived as appropriate (e.g., after expressing frustration) versus inappropriate moments (e.g., mid-task). Their finding that “timing matters more than content quality” directly motivated our focus on temporal alignment. Building on these works, our study explores proactive LLM assistance in the context of cognitive overload in survey-like interaction.

In programming contexts, Chen et al. [12] explored whether LLMs could proactively anticipate developers’ needs and offer code suggestions before being asked. Across 24 developers, they found preemptive assistance improved efficiency by 17% when timed to natural workflow pauses (e.g., after completing a function), but risked learned helplessness when mistimed—developers began waiting for AI help rather than attempting problems independently when suggestions appeared too frequently. This highlights that adaptive timing must balance support with user agency.

Pu et al. [67] explicitly evaluated “assistance versus disruption” tradeoffs in proactive programming support through 24 developer interviews. They identified “temporal appropriateness” as the primary factor distinguishing helpful from disruptive interventions—the same code suggestion was perceived as supportive when timed to natural pause points but disruptive when interrupting active coding, even if content quality was identical. They found that developers strongly preferred systems that could detect their cognitive state rather than intervening on fixed schedules.

In healthcare, Imrie et al. [41] developed proactive LLM-based explanations for medical terminology in patient-facing health forms. Their system monitored when patients hovered over unfamiliar terms and automatically generated plain-language explanations. Across 89 patients completing health history forms, proactive explanations reduced completion time by 18% and improved accuracy of self-reported symptoms compared to on-demand help buttons. However, their trigger remained behavioral—dwell time exceeding 2 seconds—requiring users to first encounter and visibly pause on problematic content before receiving help. This highlights the limitation of behavioral triggers: they detect struggle only after it has begun manifesting in observable behavior.

These studies converge on a key insight: **timing matters as much as content**. However, all rely on behavioral triggers, pause detection [41], workflow stage transitions [12, 67], or conversational cues [56], which are inherently retrospective. Behavioral indicators appear only after cognitive processes have begun to affect observable actions. Our work extends this by exploring whether *physiological* and *behavioral* triggers can enable even earlier, more prospective intervention. By detecting cognitive load through EDA before it manifests in behavior, we test whether intelligent support can be timed to the moment the struggle begins rather than after it becomes behaviorally apparent.

## 2.3 Physiological and Behavioral Sensing for Cognitive State Detection and Adaptive Applications

Physiological and behavioral sensing has become a central approach for detecting cognitive states such as workload, stress, and

fatigue, enabling adaptive systems to provide timely support. A broad range of physiological signals have been explored for this purpose. For example, eye tracking and electrocardiograms (ECG) have been shown to capture fluctuations in cognitive load and stress in virtual reality (VR) environments, allowing for real-time adaptation [31, 62]. Electroencephalography (EEG) remains a widely used modality, with recent advances demonstrating that even consumer-grade, headphone-style EEG devices can provide reliable signal quality and classification of cognitive load across diverse tasks when configured by trained experimenters [44, 45]. Similarly, multi-sensor approaches including ECG, EEG, EDA, and facial expressions have been applied to assess cognitive status, such as attention, fatigue, and workload during human–robot collaboration and human–human interview scenarios [36, 42, 53]. Planke et al. [65] further showed that online multimodal fusion of EEG, eye activity, and control inputs yields accurate and reliable inference of mental workload in driving scenarios. Among these signals, EDA is particularly prominent due to its sensitivity to sympathetic nervous system activation and its established use in workload detection and user experience evaluation [8, 32, 46, 72]. However, not all physiological measures are equally suitable for short interaction windows; for instance, cardiovascular metrics typically require sustained engagement (over one minute) to reliably reflect cognitive load [25, 75].

Complementing these physiological measures, behavioral sensing methods offer lightweight and scalable alternatives. Eye-tracking features such as fixation dynamics, combined with heart rate data, have been used to differentiate low and high cognitive load states in participant-specific models in gaming [3, 4]. More recently, mouse tracking has emerged as a promising proxy for cognitive processes in digital environments. Unlike simple response time measures, mouse trajectories reflect the continuous evolution of decision-making, capturing hesitation, uncertainty, and conflict in real time [14, 30]. This approach is especially suitable for online tasks, as it is virtually cost-free, scalable, and more robust to external distractions compared to traditional latency-based metrics [22, 39, 40]. Features such as repeated directional changes and prolonged hovering of the cursor have been associated with an increased response burden in survey tasks, providing fine-grained indicators of cognitive strain [28, 39, 51]. These findings underscore the potential of integrating physiological and behavioral sensing for real-time, fine-grained modeling of cognitive states across diverse application domains.

While these advances demonstrate the potential of multimodal sensing for adaptive support, they come almost exclusively from domains such as VR training, gaming, driving, robotics, and learning technologies, not from survey methodology. Survey research has historically relied on static interfaces and fixed rules (e.g., response-time thresholds, scripted clarification prompts), with very limited use of moment-to-moment behavioral or physiological monitoring. Thus, there is a clear gap: surveys rarely incorporate adaptive sensing to infer cognitive state, even though such techniques are well established in other domains.

Our work bridges this gap by drawing on adaptive sensing practices from HCI while grounding the application in survey methodology. We combine lightweight physiological and behavioral signals to build an LLM-based agent capable of dynamically tailoring support, offering help precisely when cognitive load is elevated, addressing challenges that traditional survey designs cannot capture.

### 3 Design of the Adaptive System for Surveys with a Proactive AI Agent

We designed an LLM-based AI agent that proactively intervenes to mitigate cognitive overload while users complete web-based multiple-choice questions.

#### 3.1 Feature Selection and Baseline Classifiers

We implemented a system that detects overload and proactively assists users. Cognitive overload detection is enabled by continuously processing physiological and behavioral data to classify users' cognitive states and dynamically trigger intervention; this classifier is further personalized and adapted throughout interaction. Further, LLM-based help for cognitive overload mitigation delivers proactive, task-related assistance when overload is detected, combining interactive clarifications and explanations that reduce cognitive strain and support task completion.

We used a dataset [54] from a web-based multiple-choice task with manipulated question difficulty, focusing on general knowledge questions. Multimodal data were collected, including mouse dynamics, eye tracking, ECG, and EDA.

To identify the most informative features to indicate cognitive load, we applied SelectKBest with `f_regression` as the scoring function. We chose `f_regression` because it evaluates linear correlations between features and the target variable, aligning with our use of linear regression models for baseline prediction. The results (cf. Appendix A) indicated that the top predictive features included three mouse movement–based measures and two EDA measures: the number of vertical direction changes in cursor movement (*ypos\_flips*), the duration of time the cursor hovered over (*hover\_time*), the total number of hovering events during a question (*hovers*), the number of peaks (*peaks\_num*) and the average tonic EDA level across the task (*tonic\_avg*). We decided to focus on mouse movement and tonic EDA measure because they can be computed robustly in real time. We excluded features unsuitable for real-time interaction. Although peak-based EDA features (*peaks\_num*) ranked highly offline, they require windowed peak detection and smoothing, introducing latency. Because `f_regression`–based feature selection does not alter the ranking of suitable features when unsuitable ones are excluded, the final feature set balances predictive power with low-latency, stable signals for live adaptation. Based on these results, we trained two baseline regression models to estimate cognitive load: one using tonic EDA and one using mouse-movement features. To reduce individual and task-related baseline effects, the EDA model used changes from task-onset baseline rather than absolute values and incorporated task difficulty (binary-coded). The mouse-based model combined cursor direction changes, hover duration, number of hover events, and task difficulty in a linear regression. We included pre-validated task difficulty [52] as a feature because easy and difficult questions were intentionally balanced and randomly interleaved. This provided relevant context for distinguishing genuine cognitive overload from responses to question difficulty alone, preventing threshold drift and over-triggering of assistance on subsequent easy questions. Importantly, task difficulty serves only as a minor contextual predictor; it does not dominate  $y_{EDA}$  or  $y_{Mouse}$  on its own. When EDA or mouse-movement patterns remain stable, the model output stays low and

does not trigger adaptation.

$$y_{EDA} = \alpha \cdot \text{tonic\_difference} + \beta \cdot \text{task\_difficulty} + \epsilon$$

$$y_{\text{Mouse}} = \alpha \cdot \text{ypos\_flips} + \beta \cdot \text{hover\_time} + \gamma \cdot \text{hovers} + \delta \cdot \text{task\_difficulty} + \epsilon$$

We used two separate unimodal models to increase robustness, allowing one modality to provide stable estimates if the other is degraded by noise, connectivity issues, or missing data.

### 3.2 Personalization and Adaptive Process

To enable personalization and adaptive support, we extended the baseline classifiers through both gradient-based fine-tuning and rule-based adaptive updates. The model of each user was personalized using gradient descent with L2 regularization, ensuring stable convergence and preventing overfitting. At the beginning of the interaction, users completed a calibration phase consisting of self-reported cognitive load questions. These calibration responses were used to fine-tune the model intercept ( $\epsilon$ ) and feature weights through a one-shot gradient-style update, thereby aligning the system's predictions with individual's subjective perception of cognitive effort. The final output was computed as the maximum of the EDA-based and mouse movement-based classifiers:

$$y_{\text{Final}} = \max(y_{\text{Mouse}}, y_{\text{EDA}})$$

This design prioritizes sensitivity, ensuring that potential overload detected by either modality is not overlooked. To trigger interventions, we defined a decision threshold  $\theta$ , which kept rule-based updating after each experimental question, i.e.,  $\theta_{\text{current}} = \theta_{\text{previous}} \pm k\delta$ ,  $k$  was determined by users' interaction. If  $y_{\text{Final}} > \theta$ , the system classified the user as experiencing cognitive overload and proactively offered help; if  $y_{\text{Final}} \leq \theta$ , no intervention was triggered. Both  $y_{\text{EDA}}$  and  $y_{\text{Mouse}}$  were implicitly aligned through the initialization and personalization procedure.

Further, the system adapted dynamically during task through rule-based threshold adjustments. These rules were grounded in two principles: (1) missed opportunities to intervene should lower the threshold to catch future overload earlier, and (2) unnecessary or unwanted interventions should raise the threshold to avoid over-assistance. These updates were guided by the outcome of each task attempt and the user's response to offered help:

- **No help offered, answer correct:** the threshold was slightly increased ( $+\delta$ ) as the user successfully managed the task independently, suggesting earlier intervention was unnecessary.
- **No help offered, answer incorrect:** the threshold was substantially decreased ( $-4\delta$ ) because the system missed an opportunity to intervene; it should step in earlier in future.
- **Help offered, accepted, and answer correct:** the threshold was slightly decreased ( $-\delta$ ), reflecting that the help was useful and should be offered somewhat earlier.
- **Help offered, accepted, and answer incorrect:** the threshold was moderately decreased ( $-2\delta$ ), as the user sought help, but the support was insufficient; earlier and possibly stronger intervention may be needed.
- **Help offered, not accepted, and answer correct:** the threshold was substantially increased ( $+4\delta$ ), as the user demonstrated that they did not need assistance, indicating the system should hold back more.

- **Help offered, not accepted, and answer incorrect:** the threshold was moderately increased ( $+2\delta$ ), since the user rejected the help but still performed poorly, suggesting reluctance to accept intervention and that the system should offer it less frequently.

Through this gradient-based personalization combined with rule-based adaptation, the system continuously refined its intervention strategy to better align with each user's behavior. We summarize the decision rules in Figure 2.

### 3.3 Development and Implementation

The system was implemented in Python and JavaScript with a Flask-based backend that coordinated real-time data collection, classification, and interaction. Mouse movement data was captured from the browser and EDA data collected via Bluetooth. Both streams were processed in real time, and the classifier dynamically evaluated the user's cognitive state at each trial.

When the classifier detected cognitive overload and crossed the intervention threshold, the system triggered a pop-up prompt asking whether the user would like assistance. If the user declined or ignored, the task continued uninterrupted. If the user accepted assistance, the interface allowed them to freely select any portion of the question text to query, ranging from a single word to a phrase to the entire question. We designed this interaction to maximize user control and flexibility. Instead of imposing predefined highlights or system-selected segments, users could indicate precisely where they experienced uncertainty. It reduces the risk of unnecessary or irrelevant explanations, ensuring that system-provided help is both context-sensitive and user-driven. The selected text was then sent to the backend for processing (cf. Figure 1 (d)-(f)).

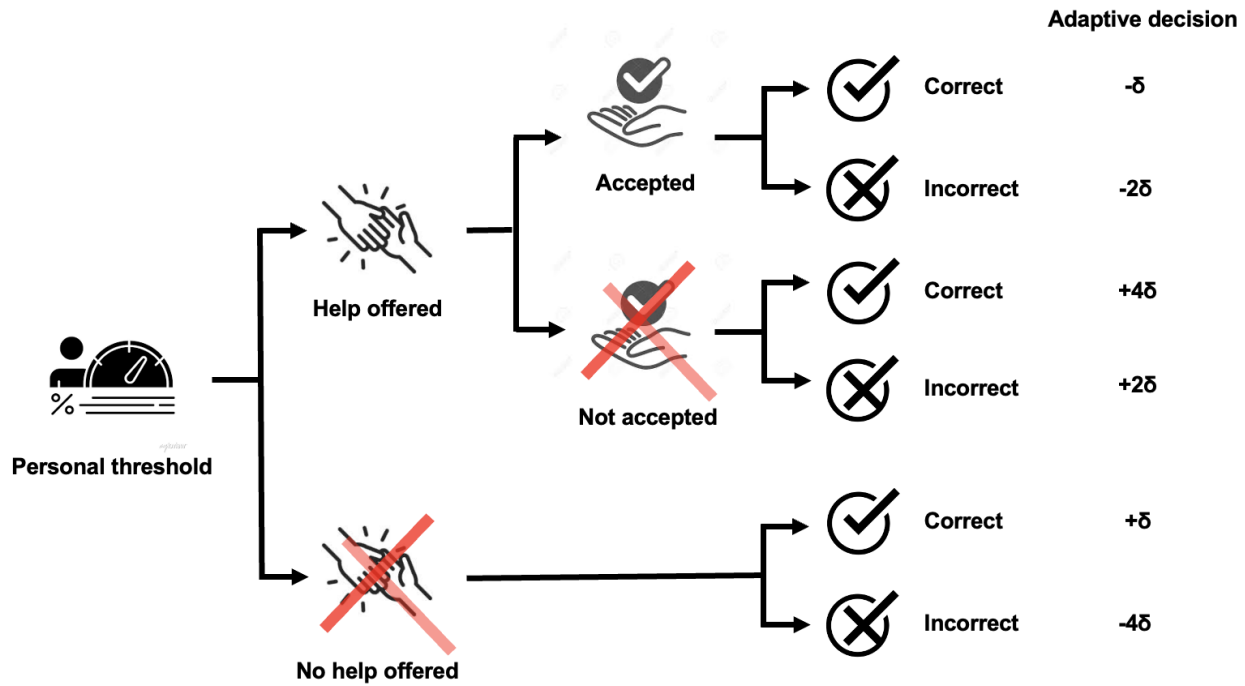
For providing explanations, we embedded a locally hosted large language model (LLaMA-2-7B [80]). Upon receiving the selected text, the model generated a natural language explanation tailored to the user's query (using a fixed prompt cf. Appendix B) and its responses were limited to a maximum of 160 tokens.

## 4 User Study

Our study investigates how adaptive LLM assistance timing affects survey response quality and user experience in sequential knowledge tasks. We engaged participants in multiple-choice questions while continuously monitoring physiological and behavioral signals to trigger real-time interventions.

We used a three-condition within-participants experimental design with the independent variable ADAPTATION STRATEGY (*Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive*). ADAPTATION STRATEGY describes how the system adjusts intervention thresholds based on user performance and help acceptance behavior. The *Aligned-Adaptive* condition aligns threshold updates with user needs, *Misaligned-Adaptive* deliberately misaligns them, and *Random-Adaptive* applies arbitrary adjustments. All conditions provide identical LLM-based assistance content, differing only in timing mechanisms. The study aimed to test central hypotheses:

- **H1 [Task Performance]:** Response accuracy will be significantly higher in the *Aligned-Adaptive* system than in the *Misaligned-Adaptive* or *Random-Adaptive* systems.



**Figure 2: Rule-based threshold adaptation logic per trial: positive changes increase the threshold (fewer future interventions), negative changes decrease it (more future interventions).**

- **H2 [User Experience]:** The *Aligned-Adaptive* system will lead to superior perceived user experience, reflected in higher ratings of efficiency, dependability, and benevolence, compared to the *Misaligned-Adaptive* or *Random-Adaptive*.
- **H3 [Workload]:** Participants will report lower subjective workload when interacting with the *Aligned-Adaptive* system, compared to the *Misaligned-Adaptive* or *Random-Adaptive*.

A total of 32 participants completed the study using a standard personal computer. The study was approved by the ethics committee of LMU Munich (EK-MIS-2025-0414-FT-d01).

#### 4.1 Task

The experimental task consisted of answering multiple-choice questions from validated German-language cognitive test batteries [52], subsequently translated into English. Each question had five options, only one of which was correct. To ensure a structured manipulation of task difficulty, questions with a reported average accuracy rate above 60% in the tests' population norm were categorized as easy, while those with a reported average accuracy rate below 60% were categorized as difficult.

The questions spanned a diverse set of general knowledge domains, including Fine Arts, Biology/Chemistry, Health, Geography, History, Politics, Mathematics, Religion, Literature, and Technology. We prepared four blocks of 20 questions each. Within each block, the distribution of easy and difficult questions was counterbalanced to ensure comparable overall difficulty across blocks. Topical coverage was balanced to prevent any single domain from disproportionately influencing performance.

#### 4.2 Study Design

The experiment employed a within-subjects design in which each participant experienced three distinct conditions, along with a calibration block presented at the very beginning. The calibration block consisted of the same number of questions as the main experimental blocks and was used to fine-tune the personalized classifier for each user based on self-reported cognitive load in a 7-point scale. Each of the three main conditions, *Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive*, manipulated the intervention threshold mechanism differently to examine its effect on user interaction, cognitive load, and task performance. To control for order effects, the sequence of the three experimental conditions and the assignment of question blocks to the conditions was randomized across participants. Furthermore, the order of questions within each block was randomized independently. The starting classifiers of all three conditions were the same, with the same feature weights and intercepts.

Based on a two-participant pilot study, we set the initial threshold to  $\theta = 12$ , requiring the estimated overload level to exceed this value before assistance was triggered. The step size was set to  $\delta = 1$  to allow gradual threshold adaptation across conditions.

**4.2.1 Aligned-Adaptive.** In this condition, the intervention threshold followed the rule-based adaptation strategy described above Figure 2. The threshold was dynamically adjusted after each question based on the user's performance and response to offered help, allowing the system to provide contextually timed interventions.

**4.2.2 Mismatched-Adaptive.** The threshold adjustments were inverted relative to the standard rule. For example, increases in the threshold in the standard rule became decreases, and vice versa. This condition tested the sensitivity of the system’s adaptation logic and its impact on user experience.

**4.2.3 Random-Adaptive.** The threshold was adjusted by a randomly selected value within the same absolute range as the adaptive rule. For example, if the maximum change in the adaptive condition was  $\pm x$ , the condition applied a random adjustment between  $-x$  and  $+x$  after each question trial. This condition served as a control to assess the effects of structured adaptation compared to unstructured, stochastic threshold changes.

### 4.3 Apparatus

The experimental setup integrated a combination of web-based task presentation, mouse movement tracking, and physiological data collection, through a Chrome browser window running our custom-built Flask web application (cf. Section 3.3).

### 4.4 EDA Recording and Processing

EDA was continuously recorded using the BITalino (r)evolution kit (PLUX Wireless Biosignals, Portugal) at a sampling rate of 100 Hz, with `revolution-python-api` provided by BITalino. To ensure high signal quality, two electrodes were attached to the index and middle fingers of the participant’s non-dominant hand, minimizing interference with the mouse-hand used for task interaction. Prior to electrode placement, the skin was cleaned and treated with a potassium chloride (KCl) solution to reduce impedance and enhance conductivity. A custom Python monitoring module was developed to manage data acquisition, segmentation, and storage in real time. It also time-aligned the physiological signals with task events.

The monitoring module maintained two levels of data handling: (1) session-level logging, where all EDA samples were continuously appended to a cumulative session file, and (2) question-level segmentation, where EDA signals were stored separately for each trial to enable fine-grained analysis. To reduce data loss in case of technical errors, periodic backups were automatically triggered at one-minute intervals. Each file was indexed by user ID, question index, and timestamp for traceability.

For each trial question, the system extracted the initial tonic level of the EDA signal at each question onset. Subsequent fluctuations were computed relative to this baseline.

During acquisition, a dedicated acquisition thread streamed raw EDA data, time-stamped each reading, and associated it with both the local question index and global session index. This ensured that EDA features could be aligned with behavioral markers such as mouse movements. At the end of each question and upon session termination, data were automatically finalized and stored.

### 4.5 Mouse Movement Recording and Processing

Mouse movement data were continuously recorded within the web-based task interface using a custom JavaScript tracking module. The tracker monitored and logged several features on a per-trial basis. One key measure was the direction flip count, defined as the number of times vertical cursor movement changed direction after exceeding a displacement threshold of 100 pixels, which served as

a proxy for disorganized or uncertain exploration. Another was the hover count, measuring how often the cursor remained stationary for longer than 500 ms, reflecting moments of hesitation or focused inspection. In addition, the system recorded the total hover time, capturing the cumulative duration of all hover events within a trial as an indicator of prolonged deliberation.

To ensure temporal alignment with task events, each mouse movement was timestamped and linked to both the local question index and the global session index. For each trial, a summary log was created containing total flips, hover events, and hover durations, which was then transmitted asynchronously to the Flask backend via REST API calls. Intermediate data (e.g., ongoing hovers) were finalized at the end of each trial, ensuring that no interaction events were lost during transitions.

### 4.6 Measures

We evaluated both user-related outcomes and system-related outcomes, combining objective task performance, self-reported perceptions, and system-level metrics. This approach provided a comprehensive understanding of how participants interacted with the system and how well the system functioned.

User-related outcomes focused on **Task Performance, Task Effort, and User Experience (Benevolence, Efficiency, and Dependability)**. System-related outcomes focused on classifier performance, quantified using confusion matrices to analyze detection accuracy, false positives, and false negatives.

**4.6.1 Task Performance and Effort.** Objective task performance was measured by the number and percentage of correctly answered multiple-choice questions. To complement behavioral outcomes, we assessed participants’ subjective perceptions of workload and performance using the unweighted NASA-TLX [34].

**4.6.2 Benevolence.** We further examined participants’ perceptions of the system’s benevolence, defined as the belief that the system acts in the user’s best interest and genuinely supports their goals [33]. We used the “benevolence” subscale of the human computer trust scale [24, 33] in 5-point Likert items:


- **Interest:** I believe that the system will act in my best interest.
- **Help:** I believe that the system will do its best to help me if I need help.
- **Preferences:** I believe that the system is interested in understanding my needs and preferences.

**4.6.3 Efficiency and Dependability.** User experience was measured using the modular UEQ+, an extension of the original User Experience Questionnaire (UEQ) [49]. As per the authors’ recommendations for partial deployment, we focused on the two scales most relevant to the system’s interaction design:

- **Efficiency:** tasks can be completed without unnecessary effort. This scale contrasts items, i.e., slow/fast, inefficient/efficient, impractical/practical, and cluttered/organized.
- **Dependability:** the user remains in control of the interaction. This scale contrasts items, i.e., unpredictable/predictable, obstructive/supportive, not secure/secure, and does not meet expectations/meets expectations.

**4.6.4 System Performance.** To evaluate the accuracy of the system’s cognitive overload detection, we collected ground-truth labels directly from participants during each question trial. After attempting a question, participants were asked whether they felt they needed help, responding with one of two options: Yes or No. These self-reported judgments served as an independent measure of perceived task difficulty and were not tied to the functionality of the AI assistance trigger.

By comparing the predicted overload states of the model with these independent self-reports, we constructed confusion matrices to assess detection accuracy, false positives, and false negatives in post-hoc validation. We further measured the number of triggered assistance and the number of accepted assistance.



**Figure 3: The task was structured in blocks. First, participants completed a personalization block to calibrate physiological and behavioral signals. This was followed by three condition blocks, *Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random*, presented in randomized order ( $\times 3$ ). After each condition block, participants provided evaluation ratings, and the session concluded with a post-study interview.**

## 4.7 Procedure

Upon arrival, participants were welcomed by the experimenter and provided with a comprehensive introduction to the study. This included an overview of the research goals, the role of the intelligent system, and the types of data being recorded (mouse movement, EDA, task responses, and self-reports). Participants received detailed instructions on the procedure and system interaction. They were informed that the AI might occasionally offer help, but were not told how assistance was triggered or about the experimental conditions. Participants could accept, decline, or ignore offers. Written informed consent was obtained, followed by a brief demographic questionnaire on age and gender.

Participants completed general knowledge multiple-choice questions using a mouse while the system continuously monitored cognitive state and triggered assistance prompts when thresholds were exceeded. When help was accepted, participants could select any portion of the question text and submit it for clarification, and were instructed on this process at the start of the study.

The session comprised four blocks (Figure 1, Figure 3): a calibration block to personalize the classifier using self-reported cognitive load, followed by three experimental blocks, *Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive*. Experimental block and question orders were randomized across participants. Each block included 20 questions, with mandatory breaks of at-least one-minute between blocks to reduce fatigue. After each experimental block, participants completed a questionnaire on workload, user experience, and system perceptions.

At the conclusion of the session, participants engaged in a semi-structured interview (cf. Appendix C). During the interview, they were asked to reflect on their overall experience with the system, describe their perceptions of the adaptive interventions, and suggest potential improvements for future iterations. The entire experiment took around one hour.

## 4.8 Participants

An a priori power analysis was conducted using *G\*Power* (version 3.1) [26], to estimate the required sample size for a study analyzed with linear mixed-effects models (LME) [26]. The analysis assumed a medium effect size ( $f^2 = 0.25$ ) based on HCI guidelines [84], an alpha level of 0.05, and a desired power of 0.80. The results indicated that a minimum total sample size of 27 participants is required to achieve a desired statistical power (actual power = 0.81). A total of 32 valid participants took part in the study, recruited from a university through message groups. The sample included 14 males and 18 females. Participants’ ages ranged from 19 to 43 years ( $M = 26$ ,  $SD = 5.77$ ). The participants received a flat payment of 12 euros.

## 5 Metadata and Validation

To contextualize the empirical findings, we report metadata related to study conditions and system performance.


### 5.1 Condition Duration

Participants spent the following amounts of completion time in each condition. For the *Aligned-Adaptive* condition, the mean duration was 17.71 minutes ( $SD = 13.08$ ); the *Random-Adaptive* condition, 14.62 minutes ( $SD = 13.52$ ); the *Misaligned-Adaptive* condition, 18.36 minutes ( $SD = 16.55$ ). A one-way ANOVA was conducted to examine differences in task durations across the three conditions. The analysis could not detect a statistically significant effect of condition on completion time ( $p = 0.56$ ).

### 5.2 System Performance

**5.2.1 Performance to know when people need help.** As shown in Figure 4, in evaluating the adaptive regression model’s ability to detect participants’ needs for assistance, clear differences emerged across conditions. The *Aligned-Adaptive* condition achieved the highest overall accuracy, averaging 61%, indicating that the system could most reliably capture participants’ actual needs. In contrast, accuracy dropped to 48% in the *Random-Adaptive* condition and further to 43% in the *Misaligned-Adaptive* condition, reflecting substantially weaker performance in both sensitivity and specificity. These results suggest that only in the *Aligned-Adaptive* condition achieved a sufficiently robust alignment with participants’ expressed help requests, whereas the *Random-Adaptive* and *Misaligned-Adaptive* conditions yielded less dependable outcomes.

To further examine errors, we compared false negative rates (FNR; proportion of missed help requests) across conditions at the participant level. This participant-level analysis provides a more fine-grained view than the aggregate confusion matrix, which collapses across all trials and participants. On average, the *Aligned-Adaptive* condition exhibited a substantially lower FNR ( $M = 21\%$ ,  $SD = 25\%$ ) compared to both the *Misaligned-Adaptive* condition ( $M = 44\%$ ,  $SD = 50\%$ ) and the *Random-Adaptive* condition ( $M = 43\%$ ,  $SD = 50\%$ ).



**Figure 4: Confusion matrices showing model performance in knowing whether users need assistance across the three experimental conditions (*Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive*). Each matrix compares the system's predicted cognitive state with the participants' self-reported need, illustrating accuracy, false positives, and false negatives per condition.**

= 47%). Paired Wilcoxon signed-rank tests with Holm-Bonferroni correction confirmed that these differences were statistically significant: *Aligned-Adaptive* vs. *Misaligned-Adaptive* ( $W = 16.0$ ,  $p_{corrected} = 0.0084$ ) and *Aligned-Adaptive* vs. *Random* ( $W = 13.0$ ,  $p_{corrected} = 0.0052$ ). These results indicate that the adaptive rule set was particularly effective at reducing false negatives, ensuring that user requests for help were less likely to go undetected compared to the *Random-Adaptive* and the *Misaligned-Adaptive* conditions.

**5.2.2 Helper Frequency and Distribution.** As shown in Figure 5, we analyzed helper presentation frequency, acceptance counts, and acceptance rates across conditions. The *Aligned-Adaptive* condition showed the highest mean number of helpers (mean = 14.03, SD = 4.30, median = 12.50), compared to *Misaligned-Adaptive* (mean = 11.50, SD = 7.76, median = 16.00) and *Random-Adaptive* (mean = 11.19, SD = 7.17, median = 14.00). A similar pattern emerged for accepted helpers, with *Aligned-Adaptive* producing more accepted helpers (mean = 9.75, SD = 3.76, median = 9.00) than *Misaligned-Adaptive* (mean = 6.03, SD = 5.50, median = 5.50) or *Random-Adaptive* (mean = 6.78, SD = 6.11, median = 6.00). Acceptance rates mirrored these trends, with *Aligned-Adaptive* achieving the highest mean rate (mean = 0.676, SD = 0.171, median = 0.700), followed by *Random-Adaptive* (mean = 0.444, SD = 0.302, median = 0.500) and *Misaligned-Adaptive* (mean = 0.388, SD = 0.256, median = 0.400). Because Shapiro-Wilk tests indicated non-normality, we applied Kruskal-Wallis tests with Dunn's post-hoc comparisons (Holm-corrected). Shown helper frequency was not detected a statistically significant effect in across conditions ( $H = 2.23$ ,  $p = 0.328$ ). Acceptance rate differed significantly ( $H = 19.86$ ,  $p < 0.001$ ), with *Aligned-Adaptive* significantly higher than both *Misaligned-Adaptive* ( $p_{corrected} < 0.001$ ) and *Random-Adaptive* ( $p_{corrected} = 0.003$ ). Accepted helper frequencies also showed a significant effect ( $H = 8.22$ ,  $p = 0.016$ ), with *Aligned-Adaptive* exceeding *Misaligned-Adaptive* ( $p_{corrected} = 0.024$ ) and marginally exceeding *Random-Adaptive* ( $p_{corrected} = 0.048$ ), while no difference emerged between *Misaligned-Adaptive* and *Random-Adaptive*.

## 6 Results

We evaluated participants' performance, workload, and perceptions across the three experimental conditions (*Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive*). Our primary interest was focused on pairwise comparisons of *Aligned-Adaptive* vs. *Misaligned-Adaptive* and *Aligned-Adaptive* vs. *Random-Adaptive*.

For the analysis of user-related outcomes (e.g., task accuracy, workload scores, user experience ratings), we employed LME, which provides a critical advantage for our study that, because our system is personalized and adapts to each individual's interaction patterns, participants may differ substantially in how thresholds shift and interventions are triggered. In all models, condition was treated as a fixed effect, with the *Aligned-Adaptive* condition set as the reference and participant ID as a random intercept. This allows direct interpretation of the fixed-effect coefficients for *Misaligned-Adaptive* and *Random-Adaptive* conditions relative to *Aligned-Adaptive*. Statistical analyses were conducted in Python using the `statsmodels` and `scipy.stats` packages.

### 6.1 Task Performance and Effort

We analyzed participants' task performance in terms of accuracy across the calibration block and the three experimental conditions. The mean accuracy for the calibration block was 41% (SD = 17%), reflecting baseline performance prior to system adaptation. Participants' mean accuracy increased in the experimental blocks: In the *Misaligned-Adaptive* condition, accuracy averaged 51% (SD = 23%), followed by 55% (SD = 20%) in the *Random-Adaptive* condition. Performance was highest in the *Aligned-Adaptive* condition, with a mean accuracy of 62% (SD = 15%). This suggests that participants performed better when interacting with the fully *Aligned-Adaptive* system compared to the other conditions.

Further, we analyzed participants' perceived task performance and effort using item-level scores from the unweighted NASA-TLX [34], as shown in Figure 6 and Table 1. For the performance dimension, where higher scores indicate better self-reported performance, the LME model revealed that participants rated their performance highest in the *Aligned-Adaptive* condition. Relative to





Figure 5: Box plots showing rates of helpers (1) shown, (2) accepted, and (3) acceptance rate distributions across the three experimental conditions (*Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive*).

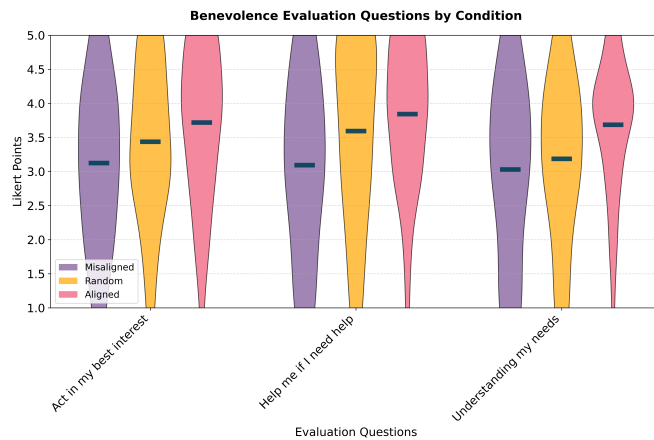
Table 1: Results of Linear Mixed Effects Models. The *Aligned-Adaptive* condition was used as the reference category.

	Coef.	Std. Error	z	p-value	CI [0.025 0.975]	Sig.
<b>Task Performance and Effort</b>						
<i>Task Accuracy</i>						
Misaligned-Adaptive	-0.103	0.026	-3.889	< 0.001	[-0.155, -0.051]	***
Random-Adaptive	-0.070	0.026	-2.637	0.008	[-0.122, -0.018]	**
Intercept	0.617	0.034	17.988	< 0.001	[0.550, 0.684]	***
Participant Variance	0.026	0.089				
<i>Perceived Performance</i>						
Misaligned-Adaptive	-0.625	0.154	-4.070	< 0.001	[-0.926, -0.324]	***
Random-Adaptive	-0.594	0.154	-3.866	< 0.001	[-0.895, -0.293]	***
Intercept	3.375	0.191	17.638	< 0.001	[3.000, 3.750]	***
Participant Variance	0.794	0.466				
<i>Perceived Effort</i>						
Misaligned-Adaptive	0.562	0.173	3.256	0.001	[0.224, 0.901]	**
Random-Adaptive	0.531	0.173	3.075	0.002	[0.193, 0.870]	**
Intercept	2.656	0.190	13.972	< 0.001	[2.284, 3.029]	***
Participant Variance	0.679	0.377				
<b>User Experience</b>						
<i>Benevolence</i>						
Misaligned-Adaptive	-0.667	0.174	-3.829	< 0.001	[-1.008, -0.325]	***
Random-Adaptive	-0.344	0.174	-1.974	0.048	[-0.685, -0.003]	*
Intercept	3.750	0.189	19.875	< 0.001	[3.380, 4.120]	***
Participant Variance	0.654	0.364				
<i>Efficiency</i>						
Misaligned-Adaptive	-0.383	0.099	-3.871	< 0.001	[-0.577, -0.189]	***
Random-Adaptive	-0.234	0.099	-2.370	0.018	[-0.428, -0.041]	*
Intercept	3.906	0.141	27.786	< 0.001	[3.631, 4.182]	***
Participant Variance	0.476	0.415				
<i>Dependability</i>						
Misaligned-Adaptive	-0.500	0.111	-4.519	< 0.001	[-0.717, -0.283]	***
Random-Adaptive	-0.281	0.111	-2.542	0.011	[-0.498, -0.064]	*
Intercept	3.602	0.116	30.923	< 0.001	[3.373, 3.830]	***
Participant Variance	0.238	0.213				

Note: Significance codes: \*\*\* p < 0.001, \*\* p < 0.01, \* p < 0.05. All models include participant as random intercept.




**Figure 6: Violin plot of participants' perceived task performance and effort of the system across conditions. The distribution of responses is shown for the *Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive* conditions, with the slash lines indicating the mean score for each question per condition. For performance, higher is better, and for effort, lower is better.**



**Figure 7: Violin plot of participants' perceived benevolence of the system across conditions. The distribution of responses is shown for the *Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive* conditions, with the slash lines indicating the mean score for each question per condition.**

*Aligned-Adaptive*, perceived performance was significantly lower in the *Misaligned-Adaptive* condition and in the *Random-Adaptive* condition. For the effort dimension, where lower scores indicate lower perceived effort, the model indicated that participants reported the least effort in the *Aligned-Adaptive* condition. Perceived effort increased significantly in the *Misaligned-Adaptive* condition and in the *Random-Adaptive* condition.



**Figure 8: Violin plot of participants' perceived user experience of the system across conditions. The distribution of responses is shown for the *Aligned-Adaptive*, *Misaligned-Adaptive*, and *Random-Adaptive* conditions, with the horizontal lines indicating the mean score for each question per condition.**

## 6.2 User Experience

Participants' perceptions of system benevolence, measured as the average of three Likert items, were generally high across all conditions, as shown in Figure 7 and Table 1. Relative to *Aligned-Adaptive*, perceived benevolence decreased significantly in the *Misaligned-Adaptive* condition and to a lesser extent in the *Random-Adaptive*.

As shown in Figure 8 and Table 1, the average efficiency score, where higher values indicate a stronger impression that tasks can be completed effectively and without unnecessary effort, showed a significant decrease in the *Misaligned-Adaptive* relative to *Aligned-Adaptive* and a smaller but still significant decrease in the *Random-Adaptive* condition. Similarly, for the average dependability score, reflecting the participants' sense of control over the interaction, ratings were significantly lower for *Misaligned-Adaptive* and *Random-Adaptive* compared to *Aligned-Adaptive*.

We analyzed four individual items of particular interest, efficient, practical, supportive, and meets expectations, to capture fine-grained perceptions of the system. For the item "efficient", participants rated the system significantly lower in *Misaligned-Adaptive* (Coef. = -0.531, SE = 0.147,  $z = -3.621$ ,  $p < 0.001$ , 95% CI [-0.819 -0.244]) and *Random-Adaptive* (Coef. = -0.344, SE = 0.147,  $z = -2.343$ ,  $p = 0.019$ , 95% CI [-0.631 -0.056]) compared to *Aligned-Adaptive*. Similar patterns were observed for "practical" (*Misaligned-Adaptive*: Coef. = -0.562, SE = 0.130,  $z = -4.315$ ,  $p < 0.001$ , 95% CI [-0.818 -0.307]; *Random-Adaptive*: Coef. = -0.469, SE = 0.130,  $z = -3.596$ ,  $p < 0.001$ , 95% CI [-0.724 -0.213]), "supportive" (*Misaligned-Adaptive*: Coef. = -0.656, SE = 0.153,  $z = -4.277$ ,  $p < 0.001$ , 95% CI [-0.957 -0.356]; *Random-Adaptive*: Coef. = -0.469, SE = 0.153,  $z = -3.055$ ,  $p = 0.002$ , 95% CI [-0.769 -0.168]), and "meets expectations" (*Misaligned-Adaptive*: Coef. = -0.625, SE = 0.186,  $z = -3.360$ ,  $p = 0.001$ , 95% CI [-0.990 -0.260]; *Random-Adaptive*: Coef. = -0.375, SE = 0.186,  $z = -2.016$ ,  $p = 0.044$ , 95% CI [-0.740 -0.010]). Random intercept variances ranged from 0.260 to 0.563 across items.

## 6.3 Qualitative Findings

To complement the quantitative results, we examined the transcriptions of the post-study interview data to gain deeper insights into participants' experiences of interacting with the intelligent support agent (cf. Appendix C). We adopted a bottom-up thematic analysis approach, allowing patterns to emerge inductively from the data. Two researchers independently coded the transcripts, generating initial codes and grouping them into broader themes. Through iterative discussion, discrepancies were resolved, and a consensus on the final thematic results was reached, which was then systematically applied back to the full set of transcripts to ensure consistent interpretation and representation of participants' perspectives. The resulting themes capture participants' perspectives on the agent's behavior, their subjective experiences with adaptive support, and broader considerations for real-world use.

**6.3.1 Motivation to Interact.** A recurring theme in the interviews was the influence of confidence and motivation on participants' willingness to engage with the system. Many participants emphasized that their decision to accept or decline assistance was directly tied to their confidence in answering a question ( $N = 27$ ). As one participant explained, "I declined only when I was really sure about the answer and otherwise whenever it popped out and I wasn't sure what's the correct answer I always went for the help" (P27).

Interestingly, participants reported different personal thresholds of certainty before accepting help, ranging from as low as 50% to as high as complete certainty. Some described interacting with the system even when fully confident, framing it as a form of reassurance or validation. For instance, one participant noted: "Because if I'm 100% sure, it cost me nothing to be more sure than the 100%. But sometimes the human brain says, I'm sure, I don't need help. But at the end, I think the AI help is much better. Or with the AI help, it's much better. No matter how sure you are, it's better to be more sure" (P5). The agent functioned not only as a compensatory tool in moments of uncertainty but also as a confidence amplifier.

Motivational factors also played a substantial role. Participants indicated that a desire to perform well or achieve higher accuracy influenced their likelihood of accepting assistance, especially when they perceived the task as evaluative or high-stakes ( $N = 18$ ). Beyond performance, curiosity emerged as an additional motivator ( $N = 7$ ). Participants explicitly described engaging with the agent out of an interest in learning new information, with one participant explaining: "That motivated me to learn more new things and get more new information" (P2).

These findings highlight that the decision to interact with the agent was shaped by a combination of situational confidence, performance goals, and intrinsic curiosity.

**6.3.2 Helpfulness and Interruptions.** Participants generally did not perceive the agent as disruptive to their workflow ( $N = 30$ ). Participants explicitly described the system as unobtrusive and likened its presence to a supportive companion rather than an interruption. As one participant reflected, "It didn't interrupt my workflow at all. It felt like when someone is thirsty and someone is like, do you need water or something, it was really helpful" (P4). This characterization highlights how the agent was largely experienced as a

contextual aid, intervening only when relevant without negatively affecting task engagement.

Nonetheless, some limitations in the presentation of assistance were noted. Some participants reported frustration when the system's responses were slightly verbose or required additional effort to extract the information most relevant to the task ( $N = 7$ ). For these participants, the need to parse lengthy or imprecise responses reduced the immediacy of the support. At the same time, preferences varied: others valued more elaborate explanations that provided additional context or opportunities for incidental learning ( $N = 3$ ). As one participant noted, "I think in a way they were still pretty helpful where it still gives you some sort of information. Maybe not the specific thing. You still learn something exactly" (P11).

These findings suggest that participants widely perceived the agent as helpful and minimally disruptive, but expectations about the granularity of responses differed. This points to the importance of personalization in balancing concise task support with opportunities for broader learning.

**6.3.3 Usage and Applications.** Across conditions, participants consistently described differences in how and when support was offered, underscoring that the timing of adaptive interventions shaped their experience as much as the content of the support itself. They explicitly noted perceivable differences between conditions, often emphasizing mismatches between when they expected assistance and when it actually appeared ( $N = 16$ ). For example, some described that help did not align with the moments they most needed it: "Across different blocks and among those questions, not always the AI gave me the help that I needed" (P3) and "For my first block, it appeared later in questions. For the second one, it appeared kind of randomly" (P14). Others contrasted two conditions based purely on temporal predictability: "In the first block, it was confused and with few help, but in the second one, I was chilling because I knew the AI would help me anyway. The first one was more stressful" (P5). Multiple participants described heightened cognitive load and negative affect when interventions arrived too late or not at all, creating unfulfilled expectations: "In the first block, the system helped me more. Then in the other block it didn't help sometimes, so I kept waiting for it, but it didn't showed up" (P13). This waiting behavior indicates that participants developed temporal expectations based on past experience, and violations of these expectations, even when help eventually arrived, created frustration: "It made me frustrated because I was waiting for help but didn't get it" (P24). These observations highlight that the perceived effectiveness of cognition-aware support hinges not only on what help is provided, but critically on when it is delivered.

Furthermore, participants frequently discussed how such a system could extend beyond the experimental setting into everyday contexts. Education, healthcare, and other professional processes were mentioned as particularly promising application areas, especially in situations where complex terminology or specialized language presents barriers to understanding. For instance, one participant reflected on healthcare contexts: "Maybe like a survey from doctors or hospitals [...] when filling those paper surveys I do feel I use Google myself. There's some complex terminology, too scientific. A system like this would help" (P32). Similarly, others highlighted bureaucratic contexts such as banking or government

forms: “I could see it actually helping me with filling forms in bureaucracy because I have to google a lot of things and having a system like this would offer some help” (P27).

Finally, participants expressed preferences for different levels of control over how and when help should be triggered. While some ( $N = 22$ ) valued the system’s proactive interventions, others ( $N = 3$ ) preferred greater autonomy, suggesting that “always-on” availability or customizable triggers would allow them to adapt the system to their own working styles, which highlights the importance of flexible design to accommodate diverse expectations.

## 7 Discussion

We present a personalized, rule-based adaptive support system to mitigate cognitive overload in knowledge-intensive online tasks and examine how the timing of AI assistance affects user experience. Our primary methodological contribution is a real-time overload detection mechanism that integrates multimodal signals, EDA and mouse movement, using individually calibrated, continuously updated thresholds. A second contribution is a systematic comparison of three adaptation strategies: *Aligned-Adaptive* (assistance timed to inferred cognitive states), *Misaligned-Adaptive* (deliberately mistimed), and *Random-Adaptive* (arbitrary timing), to assess the impact of timing in proactive support systems.

### 7.1 Summary of Results

Our adaptive system improved both performance and user experience. Task accuracy rose from 41% at baseline to 62% in the *Aligned-Adaptive* condition, compared with 51% and 55% in the other conditions. Participants also reported higher perceived performance and lower mental effort with aligned assistance.

Trust and user experience benefited as well: efficiency, dependability, and benevolence ratings were significantly higher in the *Aligned-Adaptive* condition, consistent with human–technology trust frameworks [23, 27].

Qualitative analysis showed that timely interventions were perceived as supportive, whereas misaligned timing caused stress and frustration. Help-seeking behavior was influenced by confidence, motivation, and curiosity. Participants envisioned applications in healthcare, education, and professional workflows.

### 7.2 Adaptive Support Improves Performance

We expected that response accuracy would be significantly higher in the *Aligned-Adaptive* system than in the *Misaligned-Adaptive* or *Random-Adaptive* systems. The results confirmed this expectation: participants achieved the highest accuracy in the *Aligned-Adaptive* condition, outperforming both other conditions. These findings echo evidence from survey methodology that comprehension aids can reduce errors [19, 20]. Previous research also shows that respondents rarely request help on their own: Often fewer than 10% use optional clarifications [17, 18], highlighting the value of systems that can proactively offer assistance when needed. Our results extend this literature by demonstrating that the effectiveness of proactive support depends on its alignment with the respondent’s cognitive state. Participants in the *Misaligned-Adaptive* and *Random-Adaptive* conditions did not experience the same improvements, suggesting

that when assistance is offered matters. In our study, the *Aligned-Adaptive* condition did not trigger substantially more often overall, yet it produced a significantly higher acceptance rate, indicating that support perceived as timely or appropriate is more likely to be taken up, and this higher uptake is what ultimately contributes to improved response quality. Our findings should not be interpreted as evidence about static, passively available personalization models (i.e., non-adaptive), since these factors were not experimentally manipulated. Instead, the comparison among aligned, misaligned, and random triggering suggests a design implication grounded directly in our data: adaptive assistance is most beneficial when its triggering logic produces support at moments that correspond to respondents’ actual task demands. Rather than emphasizing classifier precision abstractly, our results indicate that timing that feels meaningfully connected to respondents’ present experience, is essential for improving response quality.

### 7.3 Adaptive Systems Enhance User Experience

We expected that the *Aligned-Adaptive* system would lead to superior perceived user experience, reflected in higher ratings of efficiency, dependability, and benevolence. This hypothesis was also supported. Participants rated the *Aligned-Adaptive* system more efficient, dependable, and benevolent than the *Misaligned-Adaptive* and *Random-Adaptive* controls. These qualities are central to user trust: benevolence and dependability are established antecedents of system acceptance in HCI and organizational psychology [59, 61].

These findings are grounded in established results in adaptive assistance research, showing that help perceived as proactive but non-intrusive increases user satisfaction [2]. They also connect to recent advances in proactive LLM systems across diverse domains. In programming, adaptive LLM-based support has been shown to boost developer efficiency when aligned with task context [13, 67]; in collaborative and social settings, proactive interventions can strengthen perceptions of support and benevolence [56, 83]. However, prior work has also warned that mistimed interventions risk undermining trust by appearing disruptive [66]. Our results empirically substantiate this: when assistance was mistimed (*Misaligned-Adaptive* condition), participants rated the system as less efficient, dependable and benevolent.

We show that adaptive timing not only preserves data quality but also fosters perceptions of efficiency, dependability, and benevolence, introducing a trust-oriented lens.

### 7.4 Personalized Support Reduces Workload

We expected that participants would report lower subjective workload when interacting with the *Aligned-Adaptive* system compared to the *Misaligned-Adaptive* or *Random-Adaptive* systems. This hypothesis was confirmed: subjective workload ratings (NASA-TLX effort) were lowest and perceived performance highest in the *Aligned-Adaptive* condition. A key factor underlying this effect appears to be participants’ greater willingness to accept assistance in the aligned condition. Although the *Aligned-Adaptive* system did not deliver substantially more interventions overall, it yielded a significantly higher acceptance rate, resulting in a larger number of help episodes actually taken up. This suggests that support perceived as relevant

or well-timed is more likely to be used, and that this higher uptake, rather than sheer availability, contributes to reductions in perceived workload. Participants may have felt more supported not because the system intervened more frequently, but because its interventions were judged as meaningful enough to accept.

Task-duration metadata provide additional context for these effects. Although accepting help can introduce extra interaction steps, these differences in task duration were not statistically significant. The *Aligned-Adaptive* condition, despite a higher rate of accepted interventions, did not result in longer task times compared to the *Misaligned-Adaptive* condition. This pattern is consistent with the idea that well-timed, relevant support can reduce cognitive overload, allowing users to keep efficiency even engaging with assistance.

These results complement prior findings that proactive assistance can reduce effort by easing search, interpretation, and comprehension demands [2]. They also extend survey research showing that cognitive strain contributes to satisficing, careless responding, and breakoff [55, 64]. Our study advances this literature by demonstrating that multimodal sensing, EDA, a well-established indicator of sympathetic arousal and workload [8, 46], together with mouse dynamics [40], can be operationalized in real time to deliver support that participants are more willing to accept.

Rather than claiming that physiological and behavioral responsiveness alone “maintains” cognitive capacity, our findings more cautiously point to a design implication: systems that adapt to users’ moment-to-moment experience in ways they recognize as appropriate can meaningfully reduce perceived workload, especially in extended, sequential tasks where cumulative effort matters.

## 7.5 Over-Reliance, Reassurance-Seeking, and Shifts in User Agency

Our findings surface important ethical considerations for cognition-aware support systems, particularly the risk of over-reliance and reassurance-seeking behaviors. Several participants described anticipating or waiting for the system’s help, and some expressed frustration when support did not appear as expected. These accounts suggest that users may begin to defer judgment or delay action in hopes of being assisted, indicating a subtle shift in agency from user to system. Such patterns resonate with recent evidence that users’ self-confidence and perceived competence strongly shape their reliance on AI: increases in trust or decreases in confidence lead to greater relative AI reliance [73]. In our study, the higher acceptance rate in the aligned condition may similarly reflect a reinforcing loop in which timely support boosts trust, which in turn may heighten reliance on the system’s guidance. While this can improve task performance, it risks creating dependencies that undermine long-term autonomy or diminish users’ confidence in their own abilities.

At the same time, participants voiced divergent preferences regarding control over how and when assistance should be triggered. A substantial majority ( $N = 22$ ) explicitly favored proactive, system-initiated interventions, describing them as convenient, reassuring, or cognitively easing. While this preference speaks to the perceived value of timely support, it also signals an emergent ethical risk: when users overwhelmingly prefer to be acted upon rather than

to act, cognition-aware systems may drift toward increasingly interventionist defaults, gradually shifting normative expectations around when help should appear and who, human or system, directs the flow of interaction. Such asymmetries risk reinforcing over-reliance by making proactive intervention feel natural and self-directed effort comparatively burdensome. Others ( $N = 3$ ) explicitly preferred greater autonomy and suggested that “always-on” availability or customizable triggering options would better align with their personal working styles. These contrasting preferences underscore the ethical importance of designing adaptive systems that are not only responsive but also flexible, allowing users to calibrate the degree of automation and maintain meaningful control over their interaction. Collectively, these findings highlight a broader ethical imperative: cognition-aware systems must balance helpfulness with transparency and user agency, ensuring that adaptive support augments rather than eclipses human decision-making. Long-term deployments should examine whether and how repeated exposure amplifies reliance patterns, reshapes expectations of assistance, or alters users’ sense of competence and control.

## 7.6 Design Implications

Our study reveals core insights for designing adaptive AI support systems in cognitively demanding environments.

**7.6.1 Multimodal sensing is feasible and practical.** The technical foundation demonstrates that multimodal sensing can work in practice: combining EDA and mouse movement provides sufficient signal for real-time cognitive load detection. This approach offers a practical alternative to more invasive methods like EEG, enabling deployment in everyday digital environments where users cannot be expected to wear specialized equipment.

**7.6.2 Prioritize Temporal Precision Over Content Sophistication.** While much HCI research focuses on improving the quality of help content (e.g., more detailed explanations [57], better visualizations [58], multimodal presentations [78]), our results suggest that when help arrives may matter in addition to what is delivered. Systems should invest in accurate prediction of moments of user difficulty rather than solely optimizing content. A simple, well-timed explanation delivered at the moment of confusion may be more effective than a sophisticated explanation that arrives too early (causing interruption) or too late (after the user has already made errors or moved on). This has practical implications for proactive system design: intelligent collaborations may achieve greater impact by improving temporal detection mechanisms (e.g., physiological sensors, behavioral analytics, machine learning models) than by iterating on help content alone.

**7.6.3 Account for Individual Differences in Struggle Signals.** Our personalized classifiers with dynamic threshold adaptation proved critical. What signals struggle for one user (e.g., 15 seconds of mouse hesitation) may represent normal processing speed for another. This finding extends beyond surveys to any domain where users vary in baseline behavior, expertise, or cognitive processing styles. Adaptive systems should calibrate to individual users through brief onboarding tasks or continuous background monitoring rather than applying population-level heuristics, and implement personalized baselines that update throughout the session as the system learns

each user's typical patterns. For example, educational software should distinguish between a typically fast student pausing versus a typically slow student taking the same time.

**7.6.4 Design for Temporal Expectation Management.** Participants developed expectations about when help would arrive based on prior interaction patterns, and violations of these expectations caused frustration even when help quality remained constant. Consistency in intervention timing may be as important as accuracy. Users appear to build mental models of system behavior that influence their cognitive strategies (e.g., whether to persist through difficulty or wait for assistance). Systems should either maintain consistent temporal patterns or explicitly communicate when they are operating in different modes. If temporal patterns vary, provide subtle indicators of system state. For instance, a decision-support tool might signal “monitoring mode” versus “active assistance mode” so users know whether to expect immediate help.

**7.6.5 Mitigate Cascading Effects of Missed Interventions.** Our results demonstrated spillover effects where early struggles compounded across subsequent items. This highlights that the cost of false negatives (missing a moment of struggle) may exceed the cost of false positives (offering unnecessary help) in sequential task contexts. A single missed intervention can trigger satisficing behaviors, reduced engagement, and degraded performance on later items, creating cumulative harm that exceeds the momentary annoyance of occasionally receiving unnecessary help. Therefore, in domains with sequential dependencies (educational assessments, multi-step procedures, complex forms), calibrate interventions to favor recall over precision, better to occasionally offer help when not needed than to miss critical moments of struggle. Implement “recovery mechanisms” that detect when users have fallen into suboptimal patterns (e.g., rapid clicking, random responses) and proactively offer support to break the cascade before it progresses further.

**7.6.6 Distinguish Between Productive Struggle and Unproductive Floundering.** Not all difficulty requires intervention. Educational research distinguishes between productive struggle (moderate challenge that promotes learning) and unproductive floundering (excessive difficulty that causes frustration without benefit) [77]. The adaptive system should navigate this distinction, intervening when users flounder but avoiding premature help that short-circuits beneficial struggle. Implement multi-stage intervention strategies that begin with minimal support (e.g., subtle hints, progress indicators) and escalate only if struggle persists beyond a productive threshold. For example, an intelligent tutoring system might first highlight relevant materials, then provide a conceptual hint, and finally offer worked examples. This graduated approach respects productive struggle while preventing unproductive floundering. Physiological signals may help distinguish these states: moderate EDA increases suggest productive engagement, while sharp spikes or prolonged elevated levels indicate unproductive stress.

**7.6.7 Balance Proactivity with User Agency.** While proactive intervention improved outcomes, some participants reported wanting more control over when they received help. Overly aggressive support can undermine user autonomy and create learned helplessness where users stop attempting problems independently. Systems must balance predictive proactivity with user agency. Provide

configurable intervention modes that users can adjust based on personal preferences and task context. Options might include: (1) “Proactive” mode where the system intervenes automatically based on predicted struggle, (2) “Suggestive” mode where the system signals help availability but requires explicit acceptance, (3) “On-demand” mode where users must request help manually. Allow users to switch modes mid-task as their confidence or task difficulty changes. Track mode preferences and performance to learn optimal defaults for different user profiles.

**7.6.8 Design for Transparency in Timing Decisions.** Participants sometimes expressed confusion about why help appeared when it did (or didn't). Consider incorporating lightweight explanations of timing logic, particularly when users express confusion. For example: “I noticed you paused on this question, would you like help?” or “You've been working on similar problems successfully, so I'm giving you space to try this independently.” This transparency can help users develop more accurate mental models of system behavior and reduce frustration from unexpected timing. However, explanations must be brief to avoid adding cognitive load during already-challenging moments.

**7.6.9 Generalization Across Domains.** While our study focused on surveys with knowledge-based questions, the principle of temporal alignment applies broadly to any interactive system where users encounter intermittent cognitive demands.

Temporal precision of hint delivery may be as consequential as hint content quality. Rather than offering support on fixed schedules or after predetermined error counts, systems could leverage behavioral and physiological indicators to detect when individual students transition from productive struggle (which promotes learning) to unproductive floundering (which causes disengagement), triggering interventions precisely at this inflection point.

Medical professionals may encounter diagnostic uncertainty, yet existing systems typically present information reactively in response to explicit queries. Systems could monitor interaction patterns, dwell times, and consultation sequences to identify moments when clinicians exhibit uncertainty, proactively surfacing relevant evidence or differential diagnoses at these critical junctures rather than waiting for manual information retrieval.

Traditional approaches to software assistance rely on upfront tutorials or context-sensitive help menus, which users often ignore or find intrusive. Temporally aligned systems could instead detect when users attempt tasks inefficiently (e.g., using multiple manual steps for operations that could be automated) and deliver just-in-time feature suggestions at the moment of inefficiency, thereby reducing interruption costs while maintaining relevance.

In collaborative platforms, some participants may struggle to contribute effectively due to domain knowledge gaps, language barriers, or social anxiety. Rather than providing uniform scaffolding to all participants, temporally aligned systems could detect when specific individuals encounter barriers through participations' interaction patterns, communication hesitancy, or physiological signals, offering targeted support (e.g., background information, prompts, or facilitation) precisely when engagement falters.

These applications share in common: systems must continuously sense user state, predict moments of cognitive demand or uncertainty, and trigger interventions with temporal precision. The

specific sensing modalities will necessarily vary by domain, physiological signals and mouse dynamics for individual computer tasks, interaction logs and performance patterns for software applications, communication patterns and contribution metrics for collaborative systems, but the fundamental design imperative remains constant.

## 7.7 Limitations & Future Work

Our study demonstrates the potential of adaptive, data-driven support systems, but several limitations remain.

The controlled laboratory setting and short task duration limit the ecological validity of these findings. The demographic composition of our sample was biased toward younger and middle-aged adults, potentially limiting the generalizability of our findings to older populations, who may exhibit different cognitive and physiological patterns. In real-world surveys, interactions often span longer periods, occur across diverse devices, and take place under multitasking conditions, all of which may affect cognitive load and help-seeking behavior. Future research should extend this work to field deployments on online platforms and mobile devices to better assess generalizability. However, the modalities we analyze are increasingly accessible outside the lab: behavioral signals such as mouse movement and response timing are already available in virtually all online survey platforms [39, 40]; physiological indicators such as EDA can be captured through common consumer wearables [50, 76]. Major survey panels and commercial platforms have begun piloting integrations with smartwatch data, suggesting a realistic pathway for deploying adaptive mechanisms at scale [43]. We therefore see our findings not as limited to laboratory use but as a foundation for future real-world studies that exploit emerging, low-burden sensing technologies. Evaluating how these adaptive systems perform under naturalistic noise conditions remains an important next step.

Our evaluation also focused on knowledge-based multiple-choice questions and a single form of intervention through textual, LLM-based assistance. However, our core sensing modalities, EDA and mouse movement, capture domain-general indicators of cognitive load and behavioral uncertainty that should transfer across item types. Open-ended responses may require further adjustment to distinguish composition effort from struggle. Grid questions may produce different mouse patterns due to spatial navigation demands. Attitudinal items may generate distinct patterns reflecting evaluative rather than retrieval processes. Our personalized, continuously adapting classifiers were designed to accommodate such variations by calibrating to individual baselines, but the extent of generalization across substantially different cognitive processes remains an empirical question. Further, the adaptive rules in our system were optimized around task accuracy, a meaningful and measurable signal for knowledge-based tasks. Other relevant outcome metric will suitable to apply. In self-report settings, indicators such as satisficing behaviors (e.g., skipping or straightlining), response consistency, or attention-check performance often provide more direct measures of data quality. Future work could therefore redesign adaptive triggers around these markers to better align with other methodological goals of survey research. At the same time, our findings offer a proof of concept: even when tuned to accuracy, adaptive timing meaningfully shaped user experience and behavior,

suggesting that similar mechanisms could be repurposed to target survey-specific quality indicators once those metrics are integrated into the adaptive logic.

A central challenge lies in personalization. The system relied on a brief calibration phase to establish individual thresholds, which highlighted the difficulty in adaptive systems. The main difficulty was a cold-start problem ( $N = 6$ ): when only limited calibration data were available, the system occasionally set overly sensitive thresholds, resulting in an increased number of assistance triggers early in the task and a temporary reduction in the contrast between conditions. This effect was short-lived: thresholds stabilized quickly once the rule-based updates had incorporated several observations, after which these participants experienced no abnormal triggering behavior. While such cold-start sensitivity is a common issue in personalized adaptive systems, its practical impact in real-world survey deployments may be minimal. Many surveys already begin with warm-up items (demographic questions such as age or education), which could naturally supply the initial behavioral data needed for calibration before cognitively demanding items appear. Future work could further mitigate cold-start effects by leveraging these early low-stakes questions. Technically, one promising direction is context-aware adaptation, where additional information, such as time-on-task, behavioral variability, or short-term performance fluctuations, are used to refine early load estimates and prevent premature triggering. Another improvement involves mid-task manipulation checks, in which self-report prompts are inserted at strategic points to verify whether personalized thresholds remain aligned with participants' subjective experience. Confidence-weighted triggering could reduce early false positives by estimating predictive uncertainty (through variance measures or bootstrap ensembles) and temporarily down-weighting or gating intervention decisions when model confidence is low. These approaches would make estimation more robust, reduce sensitivity to noisy initial data, and support more stable adaptation throughout the interaction.

We used a non-task-pretrained LLM for generating clarifications. Although we constrained outputs through carefully designed prompts and a 160-token limit, some participants still perceived the explanations as overly verbose. Future systems could refine LLM outputs through domain-specific pretraining, style tuning, or adaptive summarization mechanisms that tailor the length and focus of clarifications to both the task and the user's interaction history. Beyond survey contexts, such refinements point to promising applications in formative educational settings, where the goal is not summative evaluation but scaffolding learners' reasoning processes. In these environments, an LLM operating in the background could be prompted to behave more like a reflective tutor, offering hints, asking metacognitive questions, or helping learners articulate their understanding without revealing answers outright.

Further, about participant expectations, knowing the AI might provide help could have influenced subjective ratings. In misaligned conditions, participants may have perceived the system as malfunctioning, potentially exaggerating differences in dependability ratings, consistent with expectancy violation theory [10]. However, participants were unaware of which condition they were in and only knew the system might generally use cognitive load detection. This general framing should have affected all conditions approximately equally, preserving the validity of relative comparisons

between conditions. Future work could employ between-subjects designs to further explore the effects from expectancy effects.

## 8 Conclusion

This study investigated adaptive, personalized LLM-based agents for mitigating cognitive overload in web-based tasks. In a within-subject experiment with 32 participants, we found that adaptive timing improved in acceptance rate, accuracy, workload, and perceptions of efficiency, dependability, and benevolence. These results show that well-timed, personalized support not only prevents performance loss but also fosters trust and a positive experience.

Our work advances adaptive support systems by demonstrating that closed-loop interventions driven by multimodal physiological and behavioral signals can effectively predict and respond to individual cognitive struggle in real time. This temporal design focus that intervention timing matters has implications across interactive systems where users encounter variable cognitive demands. Future research should extend this approach to longer-term deployments in other survey contexts, explore generalization across diverse item formats (open-ended responses, attitudinal scales, grid questions), and investigate alternative outcome metrics aligned with other survey methodology goals. More broadly, this work points toward a new generation of adaptive agents that continuously sense user cognitive state through physiological and behavioral signals, predict moments of struggle with personalized models, and provide proactive support aligned to individual needs. By demonstrating that such systems can maintain data quality, improve task performance, and earn user trust, we establish a foundation for physiology-aware interactive systems that offer timely, contextually appropriate assistance across domains.

## 9 Open Science & Transparency

We encourage readers to reproduce and extend our results. Therefore, our system is attached with the supplementary materials. During the preparation of this work, the authors used OpenAI's GPT-4o and Grammarly for grammar and style editing. All content was reviewed and edited by the authors, who take full responsibility for the final publication.

## Acknowledgments

This research was supported by the Federal Ministry of Research, Technology and Space (BMFTR) under grant agreement 16DKZ2019B (KODAQs). It was co-funded by the European Union – NextGenerationEU, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), project numbers 529719707; 396057129, and the Munich Center for Machine Learning (MCML).

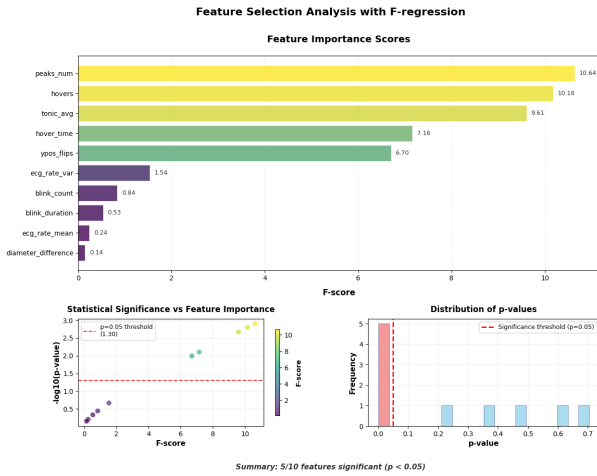
## References

- [1] Vincent Alevén, Elmar Stahl, Silke Schworm, Frank Fischer, and Raven Wallace. 2003. Help Seeking and Help Design in Interactive Learning Environments. *Review of Educational Research* 73, 3 (Sept. 2003), 277–320. doi:10.3102/00346543073003277
- [2] Salvatore Andolina, Valeria Orso, Hendrik Schneider, Khalil Klouche, Tuukka Ruotsalo, Luciano Gamberini, and Giulio Jacucci. 2018. Investigating Proactive Search Support in Conversations. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. ACM, 1295–1307. doi:10.1145/3196709.3196734
- [3] Tobias Appel, Peter Gerjets, Stefan Hoffmann, Korbinian Moeller, Manuel Ninaus, Christian Scharinger, Natalia Sevchenko, Franz Wortha, and Enkelejda Kasneci. 2023. Cross-Task and Cross-Participant Classification of Cognitive Load in an Emergency Simulation Game. *IEEE Transactions on Affective Computing* 14, 2 (April 2023), 1558–1571. doi:10.1109/taffc.2021.3098237
- [4] Tobias Appel, Natalia Sevchenko, Franz Wortha, Katerina Tsarava, Korbinian Moeller, Manuel Ninaus, Enkelejda Kasneci, and Peter Gerjets. 2019. Predicting Cognitive Load in an Emergency Simulation Based on Behavioral and Physiological Measures. In *2019 International Conference on Multimodal Interaction (ICMI '19)*. ACM, 154–163. doi:10.1145/3340555.3353735
- [5] Marilyn Ashley and Krystelle Shaughnessy. 2023. Predicting insufficient effort responding: The relation between negative thoughts, emotions, and online survey responses. *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement* 55, 3 (2023), 198. doi:10.1037/cbs0000308
- [6] Giovanni Maria Biancofiore, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, and Fedelucio Narducci. 2024. Interactive Question Answering Systems: Literature Review. *Comput. Surveys* 56, 9 (May 2024), 1–38. doi:10.1145/3657631
- [7] Danielle R. Blazek and Jason T. Siegel. 2023. Preventing satisficing: A narrative review. *International Journal of Social Research Methodology* 27, 6 (Aug. 2023), 635–648. doi:10.1080/13645579.2023.2239086
- [8] Wolfram Boucsein. 2012. *Electrodermal Activity*. Springer US. doi:10.1007/978-1-4614-1126-0
- [9] Zana Bućinca, Siddharth Swaroop, Amanda E. Paluch, Finale Doshi-Velez, and Krzysztof Z. Gajos. 2025. Contrastive Explanations That Anticipate Human Misconceptions Can Improve Human Decision-Making Skills. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, 1–25. doi:10.1145/3706598.3713229
- [10] Judee K. Burgoon. 2015. Expectancy Violations Theory. 9 pages. doi:10.1002/9781118540190.wbeic102
- [11] Charles F. Cannell, Peter V. Miller, and Lois Oksenberg. 1981. Research on Interviewing Techniques. *Sociological Methodology* 12 (1981), 389. doi:10.2307/270748
- [12] Valerie Chen, Alan Zhu, Sebastian Zhao, Hussein Mozannar, David Sontag, and Ameet Talwalkar. 2025. Need Help? Designing Proactive AI Assistants for Programming. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, 1–18. doi:10.1145/3706598.3714002
- [13] Wei-Hao Chen, Weixi Tong, Amanda Case, and Tianyi Zhang. 2025. Dango: A Mixed-Initiative Data Wrangling System using Large Language Model. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, 1–28. doi:10.1145/3706598.3714135
- [14] Paul Cisek and John F. Kalaska. 2010. Neural Mechanisms for Interacting with a World Full of Action Choices. *Annual Review of Neuroscience* 33, 1 (June 2010), 269–298. doi:10.1146/annurev.neuro.051508.135409
- [15] Edward T. Cokely, Mirta Galesic, Eric Schulz, Saima Ghazal, and Rocio Garcia-Retamero. 2012. Berlin Numeracy Test. doi:10.1037/t45862-000
- [16] Frederick Conrad, Roger Tourangeau, Mick Couper, and Chan Zhang. 2017. Reducing speeding in web surveys by providing immediate feedback. *Survey Research Methods* Vol 11 (2017), No 1 (2017). doi:10.18148/SRM/2017.V11N1.6304
- [17] Frederick G Conrad, Mick P Couper, and Roger Tourangeau. 2003. Interactive features in web surveys. In *Joint meetings of the American Statistical Association*. San Francisco, CA.
- [18] Frederick G Conrad, Mick P Couper, Roger Tourangeau, and Andrey Peytchev. 2006. Use and non-use of clarification features in web surveys. *Journal of Official Statistics* 22, 2 (2006), 245.
- [19] Frederick G. Conrad and Michael F. Schober. 2000. Clarifying Question Meaning in a Household Telephone Survey. *Public Opinion Quarterly* 64, 1 (2000), 1–28. doi:10.1086/316757
- [20] Frederick G. Conrad, Michael F. Schober, and Tania Coiner. 2007. Bringing features of human dialogue to web surveys. *Applied Cognitive Psychology* 21, 2 (Feb. 2007), 165–187. doi:10.1002/acp.1335
- [21] Jessica Daikeler, Leon Fröhling, Indira Sen, Lukas Birkenmaier, Tobias Gummer, Jan Schwalbach, Henning Silber, Bernd Weiß, Katrin Weller, and Clemens Lechner. 2025. Assessing Data Quality in the Age of Digital Social Research: A Systematic Review. *Social Science Computer Review* 43, 5 (Oct. 2025), 943–979. doi:10.1177/08944393241245395
- [22] Maria Camila Dias, Cátia Cepeda, Dina Rindlisbacher, Edouard Battagay, Marcus Cheetham, and Hugo Gamboa. 2019. Predicting Response Uncertainty in Online Surveys: A Proof of Concept. In *International Conference on Bio-Inspired Systems and Signal Processing*. <https://api.semanticscholar.org/CorpusID:88491708>
- [23] Wen Duan, Shiwen Zhou, Matthew J Scalia, Xiaoyun Yin, Nan Weng, Ruihao Zhang, Guo Freeman, Nathan McNeese, Jamie Gorman, and Michael Tolston. 2024. Understanding the Evolution of Trust Over Time within Human-AI Teams. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (Nov. 2024), 1–31. doi:10.1145/3687060
- [24] Mateusz Dubiel, Luis A. Leiva, Kerstin Bongard-Blanchy, and Anastasia Sergeeva. 2024. "Hey Genie, You Got Me Thinking about My Menu Choices!" Impact of Proactive Feedback on User Perception and Reflection in Decision-making Tasks. *ACM Transactions on Computer-Human Interaction* 31, 5 (Oct. 2024), 1–30. doi:10.1145/3685274

- [25] Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology. 1996. Heart rate variability: standards of measurement, physiological interpretation, and clinical use. *Circulation* 93, 5 (1996), 1043–1065. doi:10.1161/01.CIR.93.5.1043
- [26] Franz Faul, Edgar Erdfelder, Axel Buchner, and Albert-Georg Lang. 2009. Statistical power analyses using G\*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods* 41, 4 (Nov. 2009), 1149–1160. doi:10.3758/brm.41.4.1149
- [27] Nishani Fernando, Bahareh Nakisa, Adnan Ahmad, and Mohammad Naim Rastgoo. 2025. Adaptive XAI in High Stakes Environments: Modeling Swift Trust with Multimodal Feedback in Human AI Teams. doi:10.48550/ARXIV.2507.21158
- [28] Amanda Fernández-Fontelo, Pascal J. Kieslich, Felix Henninger, Frauke Kreuter, and Sonja Greven. 2021. Predicting Question Difficulty in Web Surveys: A Machine Learning Approach Based on Mouse Movement Features. *Social Science Computer Review* 41, 1 (July 2021), 141–162. doi:10.1177/08944393211032950
- [29] Shane Frederick. 2005. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives* 19, 4 (Nov. 2005), 25–42. doi:10.1257/089533005775196732
- [30] Jonathan B. Freeman. 2018. Doing Psychological Science by Hand. *Current Directions in Psychological Science* 27, 5 (Aug. 2018), 315–323. doi:10.1177/0963721417746793
- [31] Hong Gao and Enkelejda Kasneci. 2024. Exploring Eye Tracking as a Measure for Cognitive Load Detection in VR Locomotion. In *Proceedings of the 2024 Symposium on Eye Tracking Research and Applications (ETRA '24)*. ACM, 1–3. doi:10.1145/3649902.3655644
- [32] Vanessa Georges, François Courtemanche, Sylvain Senecal, Thierry Baccino, Marc Fredette, and Pierre-Majorique Léger. 2016. UX Heatmaps: Mapping User Experience on Visual Interfaces. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 4850–4860. doi:10.1145/2858036.2858271
- [33] Siddharth Gulati, Sonia Sousa, and David Lamas. 2019. Design, development and evaluation of a human-computer trust scale. *Behaviour & Information Technology* 38, 10 (Aug. 2019), 1004–1015. doi:10.1080/0144929x.2019.1656779
- [34] Sandra G. Hart and Lowell E. Staveland. 1988. *Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research*. Elsevier, 139–183. doi:10.1016/s0166-4115(08)62386-9
- [35] Sigrid Haunberger. 2011. To Participate or Not to Participate: Decision Processes Related to Survey Non-response. *Bulletin of Sociological Methodology/Bulletin de Méthodologie Sociologique* 109, 1 (Jan. 2011), 39–55. doi:10.1177/0759106310387721
- [36] Alexander Heimerl, Pooja Prajod, Silvan Mertes, Tobias Baur, Matthias Kraus, Ailin Liu, Helen Risack, Nicolas Rohleder, Elisabeth André, and Linda Becker. 2025. The ForDigitStress Dataset: A Multi-Modal Dataset for Automatic Stress Recognition. *IEEE Transactions on Affective Computing* 16, 2 (2025), 1219–1234. doi:10.1109/TAFFC.2024.3501400
- [37] Michael Hoerger. 2010. Participant Dropout as a Function of Survey Length in Internet-Mediated University Studies: Implications for Study Design and Voluntary Participation in Psychological Research. *Cyberpsychology, Behavior, and Social Networking* 13, 6 (Dec. 2010), 697–700. doi:10.1089/cyber.2009.0445
- [38] Jennifer L. Holland and Leah Melani Christian. 2008. The Influence of Topic Interest and Interactive Probing on Responses to Open-Ended Questions in Web Surveys. *Social Science Computer Review* 27, 2 (Dec. 2008), 196–212. doi:10.1177/0894439308327481
- [39] Rachel Horwitz, Sarah Brockhaus, Felix Henninger, Pascal J. Kieslich, Malte Schierholz, Florian Keusch, and Frauke Kreuter. 2019. Learning from Mouse Movements: Improving Questionnaires and Respondents' User Experience Through Passive Data Collection. *Advances in Questionnaire Design, Development, Evaluation and Testing* (Dec. 2019), 403–425. doi:10.1002/9781119263685.ch16
- [40] Rachel Horwitz, Frauke Kreuter, and Frederick Conrad. 2016. Using Mouse Movements to Predict Web Survey Response Difficulty. *Social Science Computer Review* 35, 3 (Jan. 2016), 388–405. doi:10.1177/0894439315626360
- [41] Fergus Imrie, Paulius Rauba, and Mihaela van der Schaar. 2023. Redefining Digital Health Interfaces with Large Language Models. doi:10.48550/arXiv.2310.03560 arXiv:2310.03560 [cs.HC]
- [42] Ashish Jaiswal, Gaurav Nale, Qiyuan An, Hamza Reza Pavel, Enamul Karim, Sneha Acharya, and Fillia Makedon. 2024. An Assistive Robotic System for Cognitive State Assessment in Individuals with Spinal Cord Injury. In *Proceedings of the 17th International Conference on Pervasive Technologies Related to Assistive Environments (PETRA '24)*. ACM, 243–251. doi:10.1145/3652037.3652068
- [43] Arie Kapteyn, Marco Angrisani, Jill Darling, and Tania Gutsche. 2024. The Understanding America Study (UAS). *BMJ Open* 14, 10 (Oct. 2024), e088183. doi:10.1136/bmjopen-2024-088183
- [44] Michael T. Knierim, Christian Zimny, Gabriel Ivucic, and Tobias Röddiger. 2025. Advancing Wearable BCI: Headphone EEG for Cognitive Load Detection in Lab and Field. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 1 (March 2025), 1–26. doi:10.1145/3712283
- [45] Jens Kohlmorgen, Guido Dornhege, Mikio L. Braun, Benjamin Blankertz, Klaus-Robert Müller, Gabriel Curio, Konrad Hagemann, Andreas Bruns, Michael Schrauf, and Wilhelm E. Kincses. 2007. *Improving Human Performance in a Real Operating Environment through Real-Time Mental Workload Detection*. The MIT Press, 409–422. doi:10.7551/mitpress/7493.003.0031
- [46] Thomas Kosch, Jakob Karolus, Johannes Zagermann, Harald Reiterer, Albrecht Schmidt, and Pawel W. Woźniak. 2023. A Survey on Measuring Cognitive Workload in Human-Computer Interaction. *Comput. Surveys* 55, 13s (July 2023), 1–39. doi:10.1145/3582272
- [47] Tanja Kunz and Marek Fuchs. 2019. Using Experiments to Assess Interactive Feedback That Improves Response Quality in Web Surveys. 247–273 pages. doi:10.1002/9781119083771.ch13
- [48] Tanja Kunz and Tobias Gummer. 2025. Effects of Objective and Perceived Burden on Response Quality in Web Surveys. *International Journal of Social Research Methodology* 28, 4 (July 2025), 385–395. doi:10.1080/13645579.2024.2393795
- [49] Bettina Laugwitz, Theo Held, and Martin Schrepp. 2008. *Construction and Evaluation of a User Experience Questionnaire*. Springer Berlin Heidelberg, 63–76. doi:10.1007/978-3-540-89350-9\_6
- [50] Evgenia Lazarou and Themis P. Exarchos. 2024. Predicting stress levels using physiological data: Real-time stress prediction models utilizing wearable devices. *AIMS Neuroscience* 11, 2 (2024), 76–102. doi:10.3934/neuroscience.2024006
- [51] Franziska M. Leipold, Pascal J. Kieslich, Felix Henninger, Amanda Fernández-Fontelo, Sonja Greven, and Frauke Kreuter. 2024. Detecting Respondent Burden in Online Surveys: How Different Sources of Question Difficulty Influence Cursor Movements. *Social Science Computer Review* 43 (04 2024). doi:10.1177/08944393241247425
- [52] Detlev Liepmann and André Beauducel. 2010. BOWIT–Bochumer Wissenstest. *Zeitschrift für Arbeits- und Organisationspsychologie A&O* (2010).
- [53] Yixiang Lim, Nichakorn Pongsakornasathien, Alessandro Gardi, Roberto Sabatini, Trevor Kistan, Neta Ezer, and Daniel J. Bursch. 2021. Adaptive Human-Robot Interactions for Multiple Unmanned Aerial Vehicles. *Robotics* 10, 1 (Jan. 2021), 12. doi:10.3390/robotics10010012
- [54] Ailin Liu, Francesco Chioffi, Felix Henninger, Lisa Bondo Andersen, Tobias Wistuba, Sonja Greven, Frauke Kreuter, and Fiona Draxler. 2026. Physiological and Behavioral Modeling of Stress and Cognitive Load in Web-Based Question Answering. arXiv:2601.17890 [cs.HC] <https://arxiv.org/abs/2601.17890>
- [55] Mingnan Liu and Laura Wronski. 2017. Examining Completion Rates in Web Surveys via Over 25, 000 Real-World Surveys. *Social Science Computer Review* 36, 1 (Feb. 2017), 116–124. doi:10.1177/0894439317695581
- [56] Tianjian Liu, Hongzheng Zhao, Yuheng Liu, Xingbo Wang, and Zhenhui Peng. 2024. ComPeer: A Generative Conversational Agent for Proactive Peer Support. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. ACM, 1–22. doi:10.1145/3654777.3676430
- [57] Jiaju Ma, Lei Shi, Kenneth Aleksander Robertsen, and Peggy Chi. 2025. AmBigChat: Interactive Hierarchical Clarification for Ambiguous Open-Domain Question Answering. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. ACM, 1–18. doi:10.1145/3746059.3747686
- [58] Callum Mackinnon and Mary Ellen Foster. 2025. Lowering the Barrier: Conversational Interfaces as a Novice-Friendly Path to Data Visualisation. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25)*. ACM, 1–7. doi:10.1145/3719160.3737623
- [59] Roger C. Mayer, James H. Davis, and F. David Schoorman. 1995. An Integrative Model of Organizational Trust. *The Academy of Management Review* 20, 3 (July 1995), 709. doi:10.2307/258792
- [60] Ted K. Mburu, Kangxuan Rong, Campbell J. McColley, and Alexandra Werth. 2025. Methodological foundations for artificial intelligence-driven survey question generation. *Journal of Engineering Education* 114, 3 (May 2025). doi:10.1002/jee.70012
- [61] D. Harrison McKnight, Vivek Choudhury, and Charles Kacmar. 2002. Developing and Validating Trust Measures for e-Commerce: An Integrative Typology. *Information Systems Research* 13, 3 (Sept. 2002), 334–359. doi:10.1287/isre.13.3.334.81
- [62] Mahsa Nasri. 2025. Towards Intelligent VR Training: A Physiological Adaptation Framework for Cognitive Load and Stress Detection. In *Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP '25)*. ACM, 419–423. doi:10.1145/3699682.3727575
- [63] NORC at the University of Chicago. 2024. General Social Survey 2024. <https://gss.norc.org>.
- [64] Andy Peytchev. 2009. Survey Breakoff. *Public Opinion Quarterly* 73, 1 (2009), 74–97. doi:10.1093/poq/nfp014
- [65] Lars J. Planke, Alessandro Gardi, Roberto Sabatini, Trevor Kistan, and Neta Ezer. 2021. Online Multimodal Inference of Mental Workload for Cognitive Human Machine Systems. *Computers* 10, 6 (June 2021), 81. doi:10.3390/computers10060081
- [66] Thanawit Prasongpongchai, Pat Pataranutaporn, Monchai Lertsutthiwong, and Pattie Maes. 2025. Talk to the Hand: an LLM-powered Chatbot with Visual Pointer as Proactive Companion for On-Screen Tasks. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, 1–16. doi:10.1145/3706598.3715579
- [67] Kevin Pu, Daniel Lazaro, Ian Arawjo, Haijun Xia, Ziang Xiao, Tovi Grossman, and Yan Chen. 2025. 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 (CHI '25)*. ACM,

- 1–21. doi:10.1145/3706598.3713357
- [68] Francesco Restuccia, Nirnay Ghosh, Shameek Bhattacharjee, Sajal K Das, and Tommaso Melodia. 2017. Quality of information in mobile crowdsensing: Survey and research challenges. *ACM Transactions on Sensor Networks (TOSN)* 13, 4 (2017), 1–43. doi:10.1145/3139256
- [69] B. J. Rhodes and P. Maes. 2000. Just-in-time information retrieval agents. *IBM Systems Journal* 39, 3.4 (2000), 685–704. doi:10.1147/sj.393.0685
- [70] Caroline Roberts, Emily Gilbert, Nick Allum, and Léila Eisner. 2019. Research Synthesis. *Public Opinion Quarterly* 83, 3 (2019), 598–626. doi:10.1093/poq/nfz035
- [71] Ido Roll, Vincent Alevan, Bruce M. McLaren, and Kenneth R. Koedinger. 2011. Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Learning and Instruction* 21, 2 (April 2011), 267–280. doi:10.1016/j.learninstruc.2010.07.004
- [72] Florian Schaule, Jan Ole Johanssen, Bernd Bruegge, and Vivian Loftness. 2018. Employing Consumer Wearables to Detect Office Workers' Cognitive Load for Interruption Management. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (March 2018), 1–20. doi:10.1145/3191764
- [73] Max Schemmer, Niklas Kuehl, Carina Benz, Andrea Bartos, and Gerhard Satzger. 2023. Appropriate Reliance on AI Advice: Conceptualization and the Effect of Explanations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces (IUI '23)*. ACM, 410–422. doi:10.1145/3581641.3584066
- [74] Michael F Schober, Frederick G Conrad, and Jonathan E Bloom. 2000. Clarifying word meanings in computer-administered survey interviews. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 22.
- [75] Fred Shaffer and Jay P Ginsberg. 2017. An overview of heart rate variability metrics and norms. *Frontiers in public health* 5 (2017), 258. doi:10.3389/fpubh.2017.00258
- [76] Pekka Siirtola. 2019. Continuous stress detection using the sensors of commercial smartwatch. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers (UbiComp '19)*. ACM, 1198–1201. doi:10.1145/3341162.3344831
- [77] Tanmay Sinha and Manu Kapur. 2021. When Problem Solving Followed by Instruction Works: Evidence for Productive Failure. *Review of Educational Research* 91, 5 (June 2021), 761–798. doi:10.3102/00346543211019105
- [78] Maura Spiegelman and Frederick Conrad. 2025. Improving Understanding of Survey Questions with Multimodal Clarification. *methods data* (2025), 1–28. doi:10.12758/MDA.2025.09
- [79] Roger Tourangeau, Lance J. Rips, and Kenneth Rasinski. 2000. *The Psychology of Survey Response*. Cambridge University Press. doi:10.1017/cbo9780511819322
- [80] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shriti Bhoale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288 [cs.CL] <https://arxiv.org/abs/2307.09288>
- [81] Lev Velykoivanenko, Kavous Salehzadeh Niksirat, Stefan Teofanovic, Bertil Chapuis, Michelle L. Mazurek, and Kévin Huguenin. 2024. Designing a Data-Driven Survey System: Leveraging Participants' Online Data to Personalize Surveys. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, 1–22. doi:10.1145/3613904.3642572
- [82] Laia Vidal Sabanés and Iria da Cunha. 2025. AI as a resource for the clarification of medical terminology: An analysis of its advantages and limitations. *Terminology* 31, 1 (May 2025), 37–71. doi:10.1075/term.00083.vid
- [83] Bufang Yang, Yunqi Guo, Lilin Xu, Zhenyu Yan, Hongkai Chen, Guoliang Xing, and Xiaofan Jiang. 2025. SocialMind: LLM-based Proactive AR Social Assistive System with Human-like Perception for In-situ Live Interactions. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 1 (March 2025), 1–30. doi:10.1145/3712286
- [84] Koji Yatani. 2016. *Effect Sizes and Power Analysis in HCI*. Springer International Publishing, 87–110. doi:10.1007/978-3-319-26633-6\_5
- [85] Hye Sun Yun, Mehdi Arjmand, Phillip Sherlock, Michael K Paasche-Orlow, James W Griffith, and Timothy Bickmore. 2024. Keeping Users Engaged During Repeated Interviews by a Virtual Agent: Using Large Language Models to Reliably Diversify Questions. In *Proceedings of the ACM International Conference on Intelligent Virtual Agents (IVA '24)*. ACM, 1–10. doi:10.1145/3652988.3673929

## A Feature Importance Analysis using F-regression



**Figure 9: Visualization of feature importance scores derived from F-regression analysis. The horizontal bar chart (top) displays the influential features ranked by their F-scores, with color intensity indicating relative importance. The scatter plot (bottom left) illustrates the relationship between feature importance and statistical significance ( $-\log_{10}(p\text{-value})$ ), with the red dashed line indicating the  $p=0.05$  significance threshold. The histogram (bottom right) shows the distribution of p-values across all features, highlighting the proportion of statistically significant predictors.**

Appendix A visualized the feature selection considerations.

## B A Prompt Used in Clarification Generation

«SYS»

You are a helpful assistant who clearly explains concepts in English. Provide ONLY the context.

«/SYS»

Explain the concept "{words}" in English.

## C The Post-study Semi-structured Interview Questions

### C.1 Overall Experience:

- You just used a system that attempted to predict how difficult a question was for you and provided help based on that prediction. Can you describe your overall experience with it, especially how it felt across different questions or sessions?

### C.2 Perceived Costs and Risks:

- Did accepting the help ever interrupt your workflow or make the task more difficult in any way? (For instance, did it feel mistimed, break your focus, or require extra effort to understand?)

- What about when you declined help? Did that ever make things harder later on? (Like missing out on useful suggestions or feeling unsure afterward?)

### C.3 Decision to Accept or Decline Help:

- During the experiment, we ran three different blocks of the task: one where the system adapted positively based on your behavior, one where it adapted in the opposite direction, and one where behavior was responded to randomly. Did you notice any differences between these three blocks? How so?
- Can you walk me through your experience in any block, like what felt helpful, what felt confusing, or annoying?
- What motivated you to accept, ignore, or reject the help when it was offered?
- Were there times when you declined help even though you weren't confident, or accepted help even when you felt you didn't need it?
- What influenced that decision?

### C.4 Real-World Use & Improvements:

- If you could improve the system, what would you change? (Anything that frustrated you, or features you wish it had?)
- Would you want to use a system like this in real life? Why or why not? (And in what kind of situations or tasks would it make sense or not?)