

“What are you doing?”: effects of intermediate feedback from agentic LLM in-car assistants during multi-step processing

Johannes Kirmayr, Raphael Paul Wennmacher, Khanh Huynh, Lukas Stappen, Elisabeth André, Florian Alt

Angaben zur Veröffentlichung / Publication details:

Kirmayr, Johannes, Raphael Paul Wennmacher, Khanh Huynh, Lukas Stappen, Elisabeth André, and Florian Alt. 2026. “What are you doing?”: effects of intermediate feedback from agentic LLM in-car assistants during multi-step processing.” In *CHI '26: proceedings of the 2026 CHI Conference on Human Factors in Computing Systems, Barcelona, Spain, April 13-17, 2026*, edited by Nuria Oliver, David A. Shamma, Heloisa Candello, Pablo Cesar, Pedro Lopes, Alessandro Bozzon, Thomas Kosch, et al., 504. New York, NY: ACM. <https://doi.org/10.1145/3772318.3790997>.

"What Are You Doing?": Effects of Intermediate Feedback from Agentic LLM In-Car Assistants During Multi-Step Processing

Johannes Kirmayr
 BMW Group Research and Technology
 Munich, Germany
 Human-Centered Artificial Intelligence
 Augsburg University
 Augsburg, Germany
 johannes1.kirmayr@uni-a.de

Raphael Paul Wennmacher
 LMU Munich
 Munich, Germany
 raphael@wennmacher.eu

Khanh Huynh
 BMW Group Research and Technology
 Munich, Germany
 LMU Munich
 Munich, Germany
 khanh.huynh@ifi.lmu.de

Lukas Stappen
 BMW Group Research and Technology
 Munich, Germany
 lukas.stappen@bmw.de

Elisabeth André
 Human-Centered Artificial Intelligence
 Augsburg University
 Augsburg, Germany
 andre@informatik.uni-augsburg.de

Florian Alt
 LMU Munich
 Munich, Germany
 florian.alt@ifi.lmu.de


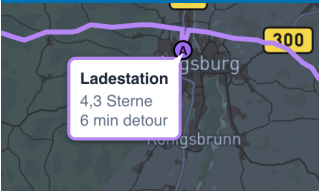

USER	AGENTIC ASSISTANT - COMPARING 2 FEEDBACK MECHANISMS			
	Tool Calls	🔍 search_contact()	🔍 search_charger()	🗺️ start_navigation()
🗣️ Take me to Mr. Bergheim, he emailed me the address, and plan a fast-charging station if the battery level drops below 10%.	No Intermediate Feedback	🗣️ Assistant is planning...	🗣️ Assistant is planning...	🗣️ 🗣️ Alright, I have found Lukas Bergheim in your contacts. The address Roseng...
	Planning & Results Feedback	🗣️ I have found Lukas Bergheim in your contacts.	🗣️ I have planned in a fast-charger stop with 6 minutes detour.	🗣️ Route with charging stop to Lukas Bergheim calculated...
				
	»»	»»	»»	»» End result 45s

Figure 1: How should agentic assistants communicate during long-running multi-step operations? We compare two feedback strategies: No Intermediate (NI) feedback (top), where the system acknowledges the request and remains silent until delivering the final result, versus Planning & Results (PR) feedback (bottom), where planned steps and intermediate outcomes are communicated through synchronized audio and visual updates. The illustrated task exemplifies the multi-step operations agentic assistants perform, including contact lookup, address extraction, battery check, and charging stop planning. Our study (N=45) examines effects on perceived latency, trust, task load, and user experience in stationary and driving contexts.



Abstract

Agentic AI assistants that autonomously perform multi-step tasks raise open questions for user experience: how should such systems communicate progress and reasoning during extended operations, especially in attention-critical contexts such as driving? We investigate feedback timing and verbosity from agentic LLM-based in-car assistants through a controlled, mixed-methods study (N=45) comparing planned steps and intermediate results feedback against silent operation with final-only response. Using a dual-task paradigm with an in-car voice assistant, we found that intermediate feedback significantly improved perceived speed, trust, and user experience while reducing task load - effects that held across varying task complexities and interaction contexts. Interviews further revealed user preferences for an adaptive approach: high initial transparency to establish trust, followed by progressively reducing verbosity as systems prove reliable, with adjustments based on task stakes and situational context. We translate our empirical findings into design implications for feedback timing and verbosity in agentic in-car assistants, balancing transparency and efficiency.

CCS Concepts

• **Human-centered computing** → **Auditory feedback**; **Natural language interfaces**; *Graphical user interfaces*; *Collaborative interaction*; **User studies**; **Sound-based input / output**; **Displays and imagers**; Haptic devices.

Keywords

Large Language Model, AI Agent, Human-Agent Interaction, Agent Feedback Design

ACM Reference Format:

Johannes Kirmayr, Raphael Paul Wennmacher, Khanh Huynh, Lukas Stappen, Elisabeth André, and Florian Alt. 2026. "What Are You Doing?": Effects of Intermediate Feedback from Agentic LLM In-Car Assistants During Multi-Step Processing. 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, 18 pages. <https://doi.org/10.1145/3772318.3790997>

1 Introduction

Large language model (LLM) agents are transforming how we interact with Artificial Intelligence (AI) systems, moving beyond single-turn question-answering to autonomously executing complex, multi-step tasks [1, 31, 72]. These agentic systems decompose user requests, invoke multiple tools [56], and synthesize results across extended processing periods; from searching and comparing flight options to analyzing documents and generating comprehensive reports. This shift from simple query-response to autonomous task execution introduces fundamental design challenges: How should these systems communicate their progress to users during lengthy operations? When should they provide updates versus working silently in the background? How much detail is appropriate, and how should feedback adapt to different contexts and user needs?

These questions become particularly acute in dual-task scenarios where users engage with AI assistants while performing other activities. Consider an in-car voice assistant that handles complex requests while the driver navigates traffic. In such context, poorly

designed feedback can create dangerous distractions or cognitive overload [11, 65, 66], while insufficient communication can result in "ambiguous silence" [78], leaving users uncertain about system progress and thereby undermining trust and perceived responsiveness [23, 55, 68].

Currently deployed user-facing agentic systems reveal this gap through diverse feedback practices. Cursor [21], which reached one million users within 16 months, operates nearly silently until completion with details on demand. Manus AI [2] provides verbose step-by-step narration, while Perplexity [3] previews steps but withholds intermediate results. This variability, from minimal to maximal transparency, highlights the lack of shared design principles for agentic feedback and underscores the need for open-research guidelines as user-deployed systems become widespread.

Decades of HCI research have established well-tested principles for designing system feedback. Yet agentic systems, with their extended processing and autonomous actions, bring new considerations. First, latency is inherent rather than accidental. Prior work demonstrates that unexpected delays degrade user experience [46, 58] and that progress indicators improve perceived responsiveness [47]. However, in agentic systems, the expanded processing time is a consequence of intentionally increased reasoning and multi-step tool use, posing the question of whether findings on perceived delay mitigation transfer to waiting time as it becomes expected. Second, the volume of information generated during multi-step processing also expands well beyond typical single-turn interactions, raising questions about cognitive load and information distribution [10]. As described, this is particularly critical in dual-task contexts such as driving, where even lightweight interactions can impair performance. Third, grounded communication requires evidence of both perception and understanding [4, 9, 20]. In long-running agentic systems, it is unclear whether grounding is maintained when perception is acknowledged upfront but evidence of understanding is deferred until the final response, or whether users expect transparency throughout the process. Finally, trust is known to be a key factor in human-AI relationships, as it often determines whether a system is adopted and relied upon [68]. Prior work shows that transparency mechanisms such as explanations can foster trust [45, 80]. For agentic systems, this raises questions about which forms of feedback best support users in calibrating trust during extended and autonomous processing.

To address the gap between emerging agentic capabilities and established feedback design principles, we investigate the feedback design for an LLM-based agentic in-car assistant through a mixed-methods study (N = 45) in a controlled car simulation environment. We focus on three critical dimensions: (1) **feedback timing** — whether systems should provide intermediate informative updates during processing or deliver results only upon completion; (2) **interaction context** — how feedback strategies perform when the AI assistant is the primary versus a secondary task alongside driving; and (3) **adaptive verbosity** — how feedback detail should evolve with situational demands and long-term use.

Our quantitative experiments reveal that intermediate feedback consistently outperforms end-only delivery across multiple metrics and interaction context. Providing planned steps and incremental results during processing increased perceived speed ($d_z = 1.01$), improved user experience ($d_z = 0.54$), enhanced trust ($d_z = 0.38$),

and, surprisingly, reduced task load ($d_z = -0.26$) despite multiple interaction points, compared to silent processing followed by a comprehensive final response. Complementing these findings, our qualitative interview analysis uncovers sophisticated adaptation preferences. Participants envisioned systems that initially provide transparent, detailed feedback to establish trust, then progressively reduce verbosity in favor of efficiency as the system proves reliable. Yet they expected transparency to immediately return for novel, ambiguous, or high-stakes requests. Preferences varied regarding context-aware adaptation in social settings and media consumption, with a consistent desire for simple override controls when automatic adaptation was unsatisfactory.

Contribution Statement. Our empirical findings advance agentic in-car assistant design through: (1) **Empirical evidence** from a controlled in-car study (N=45) showing that intermediate, informative updates during extended agent operation improve responsiveness, trust, and user experience across single- and dual-task contexts; (2) **Adaptive verbosity patterns** from qualitative interviews showing that users desire adaptive feedback detail that decreases as they experience system reliability over time, but increases situationally for novel, ambiguous, or high-stakes tasks; and (3) **Design implications** for feedback design and adaptation in in-car single- and dual-task interaction, with potential relevance to other primary-task systems (e.g., customer support) and dual-task contexts (e.g., smart-home assistants while cooking).

2 Related Work

2.1 Feedback Strategies in Current Agentic Systems

Current systems vary widely from minimal (Cursor [21]) to verbose (Manus [2]) feedback. This diversity reflects different implicit assumptions about user needs. Cursor’s minimal approach embodies a “stay out of the way” philosophy, prioritizing uninterrupted workflow over transparency. This implies intermediate details may distract expert users, and suits contexts where users have established trust or where processing steps remain technical rather than decision-relevant. Manus’s verbose approach implies that transparency builds trust and helps users maintain situational awareness, though risking information overload. Perplexity’s [3] hybrid strategy, showing plans but not intermediate results, attempts to balance expectations management with efficiency. These systems exemplify the diversity of feedback strategies deployed in practice. At the same time, they also underscore the need for empirically grounded design principles. With companies rarely publishing their design rationales or formative studies, the open research community lacks systematic guidance on how feedback strategies should align with varying user needs and task contexts.

2.2 Principles from Human-AI Interaction Research

Grounding and Communication. Research on human-AI communication builds on foundational studies of human communication and human teamwork. Grounding communication theory [20] highlights that effective collaboration requires maintaining common ground, meaning shared knowledge, beliefs, and assumptions. This

common ground must be updated continuously, not only about content but also about the process of interaction. Brennan [14] further extends this to human-computer interaction, emphasizing that people need to be able to seek evidence that they are understood, with Yankelovich et al. [78] stating that the absence of feedback leaves the user in an ambiguous silence. In agentic systems, where a single user request initiates extended multi-step reasoning, grounding becomes more complex. The question is whether continuous intermediate updates have to be provided or if an upfront indication of perception with a deferred condensed final answer suffices to maintain common ground.

Latency and Waiting. Delays in system responses have long been shown to degrade user experience. Early work demonstrated that response delays decrease satisfaction [46, 58] and that unexpected waiting increases frustration [58]. In voice interfaces, such delays may even lead users to assume the system has failed [55] or that an error has occurred [23]. Nielsen highlights 10 seconds as an upper bound for keeping users’ attention during waiting periods [48, p.135]. Successful mitigation strategies include progress indicators [47], conversation fillers [43], and explanations of ongoing processing [80]. While such strategies reduce anxiety and foster trust, their effectiveness in agentic systems, where latency is not accidental but an expected and productive aspect of multi-step reasoning, remains underexplored.

Trust and Transparency. Trust is central in human-AI interaction, influencing whether users rely on or reject system support [68]. Lee and See [39] define trust as “an attitude that an agent will achieve an individual’s goal in a situation characterized by uncertainty and vulnerability”. In the case of agentic systems, users expose themselves to such vulnerability when delegating complex requests to the system. For trust to emerge and persist, users must clearly understand what they can expect. Transparency is widely recognized as a means of fostering trust [45], with explanations that align user expectations and system behavior [40, 80]. Empirical findings from human-AI collaboration confirm that transparency enhances trust, especially when systems make their reasoning explicit [71]. Hoff and Bashir [29] further distinguishes trust in automation into dispositional trust as the person’s general tendency to trust automation, situational trust in a given context, and learned trust by prior experience. For agentic systems, situational and learned trust are particularly important, as they indicate the need for a dynamic feedback design. Importantly, in dual-task contexts such as driving, transparency must be carefully balanced: feedback must establish trust without overloading cognitive resources.

Human Oversight. Oversight on AI systems encompasses two concepts: passive oversight (monitoring) and active oversight (human intervention) [5, 38, 62]. Its primary goal is to enable humans to detect and correct errors in else autonomous AI decisions. This human-AI collaboration can lead to improved agent performance [28], and user control can increase trust in the system [22]. However, research also highlights critical limitations: humans often overtrust plausible AI outputs or override accurate ones, undermining effective human oversight [28]. Additionally, human involvement substantially increases cognitive load on them [28]. Therefore, effective oversight requires careful design. Sterz et al. [62] emphasize that

humans need epistemic access, including a sufficient understanding of what the system is doing and why. This requires delivering the right transparency at the right time and format to support monitoring without overloading cognition, motivating our research question on designing agentic feedback for clear comprehension.

2.3 Feedback Design under Cognitive Constraints

Multimodal Feedback and Cognitive Resources. Work on multimodal feedback highlights strategies for distributing information across sensory channels. Wickens' Multiple Resource Theory suggests that tasks utilizing different perceptual modalities draw from separate cognitive resource pools, reducing interference [75]. This principle proves particularly relevant for in-vehicle interaction, where auditory-vocal tasks interfere less with vehicle control than visual-manual ones [30]. However, modality selection involves trade-offs: auditory feedback preserves visual attention but lacks persistence [15], while visual feedback provides spatial detail but requires potentially dangerous gaze shifts [16]. Oviatt and Cohen [53] shows that coordinating modalities, using audio cues to direct attention to visual information, can reduce cognitive load compared to single-modality approaches. For agentic systems with extended multi-step feedback, distributing information across modalities over time without overwhelming any single channel becomes a central design challenge.

Secondary Task Distraction in Driving and Automotive Interface Design. Even lightweight activities can impair driving performance: hands-free conversations shrink the functional field of view [7], and voice-based interactions impose moderate cognitive workload [64–66]. Driver distraction is a major safety factor in crashes [11], yet drivers consistently underestimate its impact [74]. Consequently, automotive HCI research has focused on designing assistants that minimize distraction through careful modality choice. Speech interfaces have been found to generally outperform visual-manual [41] and touch- or gesture-based [79] alternatives on driving performance measures. Braun et al. [13] show that still visualizing natural language alongside speech improves information recall through text summaries, while keywords reduce cognitive load and icons increase attractiveness. Nevertheless, even basic voice interactions, such as dictating text messages, elevate cognitive workload compared to undistracted driving [42]. More recently, researchers have begun investigating LLM-powered conversational agents in vehicles. Sorokin et al. [60] examine multimodal systems combining voice and graphical interfaces, arguing that maintaining mutual understanding between the user and LLM requires bidirectional translation: informing users of the AI's focus and changes while making user actions legible to the model. However, existing studies largely address short, reactive interactions: command execution, clarification dialogues, or navigation updates. Comparatively little empirical work examines feedback design for *long-running, agentic* in-vehicle assistants that autonomously perform multi-step tasks.

Following Norman's principle that design must align with user needs and cognitive limits [49], agentic systems must carefully balance trust-building transparency with cognitive safety. Feedback

strategies should therefore support trust development while minimizing distraction, ideally increasing responsiveness perception and overall user experience.

3 Research Questions

To mitigate the latency of multi-step processing and the larger volume of information within a single assistant turn, feedback must be designed to balance responsiveness, trust, grounded communication, and task load. To guide the design of feedback mechanisms, we address three research questions within an in-car assistant environment:

- RQ1 When should agentic systems provide feedback: how does the timing of feedback—providing updates during task execution versus only at completion—affect users' perception of waiting time, overall experience, trust, and cognitive workload; and how should feedback timing be adapted over time and in situational context?
- RQ2 How do task complexity and driving demands influence feedback preferences: how do longer processing times and whether users are actively driving shape preferences for when and how much system feedback to receive?
- RQ3 How detailed should system feedback be: how should the level of verbosity in system feedback adapt over time as users become familiar with the system, and how should it adjust based on situational context to maintain an optimal balance between keeping users informed, minimizing distraction, and building trust?

4 User Study

To answer these research questions, we conducted a mixed-methods user study (N=45). In the quantitative part, we perform a within-subject lab study focusing on feedback timing effects under varying conditions (RQ1, RQ2). Our independent variables are (1) Feedback Timing, (2) Task Duration, and (3) Interaction Context. Our dependent variables are perceived (1) Speed, (2) UX, (3) Task Load, and (4) Trust. The quantitative study design and analysis are detailed in Section 4.3. Qualitative interviews extend the measured feedback timing effects and focus on adaption preferences for feedback verbosity and timing over time and situational context (cf. details in Section 4.4).

4.1 Apparatus

4.1.1 User Study Environment. The study was conducted in a controlled car simulation environment (Figure 2). The simulation environment used a fixed position, full-frame car mockup. The participants sat in the driving seat throughout the whole study.

The interaction setup included: (1) a voice user interface providing auditory feedback via an external speaker, (2) a graphical user interface on a tablet positioned at the typical in-car center console location, and (3) a lane-keeping driving simulation displayed on a vertical screen outside the car-mockup and in front of the participant at 2.7 meter distance. During the quantitative study (cf. Sec. 4.3.1), participants interacted with the voice assistant in two interaction contexts: stationary (single-task), where they sat in the car without performing the additional lane-keeping task, and driving (dual-task), where they concurrently performed the driving-related



Figure 2: Study apparatus: 1. Speaker (Voice User Interface), 2. Center Display (Graphical User Interface), 3a. Driving simulation in the form of a lane-keeping task, 3b: mouse to correct continuous lateral drift for driving simulation.

task. The driving-related task required participants to maintain lane position by continuously correcting lateral drift. As the steering wheel in the car mockup (visible in Figure 2) was a non-functional component, we implemented mouse-based steering input for the simulation. Participants used mouse clicks to control lateral vehicle position. This arrangement enabled a controlled manipulation of cognitive load and provided a consistent basis for comparing feedback preferences between single-task and dual-task situations.

4.1.2 LLM-based Voice-Assistant System. Our research team had previously developed a fully functional agentic LLM-based in-car voice assistant capable of handling complex, multi-step tasks with real-time intermediate feedback. This system served as the foundation for our study, providing realistic interaction flows and response timings.

Based on this, for the user study, we created a prototype in ProtoPie and deployed it on a tablet simulating the vehicle’s center display. This study-specific LLM-inspired prototype ensured strict comparability across all conditions in the within-subject design. For each task configuration, the target utterance was shown on-screen for participants to read aloud. The system transcribed the spoken input in real-time and displayed the transcription on-screen, signaling to participants that their input had been received. Upon receiving the voice input, the prototype system triggered the corresponding deterministic interaction sequence, delivering visual and auditory outputs at fixed timesteps according to the configuration of one of the eight interaction tasks (cf. Figure 3 and Figure 4). Thus, unlike the working system with its dynamic LLM-generated responses, the study prototype used predefined LLM-inspired responses and

fixed response timings to ensure experimental control and reproducibility, while the visible real-time transcription preserved the experience of interacting with a live system.¹

4.2 Procedure

The study was conducted in the car-mockup simulation environment in-person over two weeks with 45 participants (~60 min each). Participants were recruited through a major automotive company’s mailing lists and community channels across multiple departments to ensure demographic diversity and varying levels of familiarity with LLMs and voice assistants. The study’s procedure consists of three phases: preparation, task execution with interleaved questionnaires, and post-task interviews.

Preparation. Participants first completed informed consent and a demographic questionnaire covering age, gender, and familiarity with: LLMs, general voice assistant systems, and the company’s in-car voice assistant. The experimenter then introduced the driving simulation and center display prototype, with participants training on the lane-keeping task using mouse clicks to maintain lane position. Finally, participants were briefed on the capabilities of an agentic in-car assistant and informed that the study focused on *feedback delivery methods*, not AI performance.

Task Execution. Participants then completed eight tasks covering all experimental conditions. Questionnaires were interleaved at different points during the session to capture the dependent variables, including perceptions of speed, task load, user experience, and trust (details in Section 4.3).

¹We provide screen videos of the ProtoPie implementation in the supplementary material.

Post-Task Interview. Finally, participants were interviewed about feedback adaptation preferences through three open-ended questions with follow-up prompts as needed (details in Section 4.4).

Ethics. While formal approval from an ethics review board was not required in the jurisdiction where the study was conducted, all procedures adhered to recognized ethical research practices and the ACM Code of Ethics, including clear participant information, informed consent, and data protection.

4.3 Quantitative Study

4.3.1 Quantitative Experiment Design. To capture the influence of feedback timing, task length and attentional context, we employed a controlled within-subjects $2 \times 2 \times 2$ factorial design. The independent variables (IVs) and condition levels are presented in Table 1 and further explained in the following paragraphs.

As a result, shown in Figure 3, each participant performs 8 tasks across the $2 \times 2 \times 2$ conditions.

IV1: Feedback Timing. To study the effects of feedback timing, we contrasted two complementary system behaviors. The **No Intermediate (NI)** feedback condition served as a strong baseline: after the participant's request was spoken, the system confirmed perception with a clicking sound and a visual message "I am planning..." on the center screen, but then remained silent until delivering the final response. This setup reflects how many AI agents operate and is particularly relevant in dual-task settings such as driving, where silence and background operation reduce interruptions and let the assistant "get out of the way". In contrast, the **Planning & Results (PR)** feedback condition represented a transparent approach, providing informative intermediate updates during processing in addition to a final summary. This condition aims to mitigate uncertainty during waiting and to distribute the increased information load across smaller, progressive steps rather than presenting it all condensed at once. Intermediate updates are conveyed auditorily and complemented by a visual display, which has been shown to be an effective design choice [13]. We also considered including a third condition with minimal progress cues (e.g., simple signals or conversation fillers, such as "working on it"). Prior work has already repeatedly shown that such cues improve perceived latency and user experience [43, 47]. Given the already large design space and the novel dimension of increased information load to be conveyed, we chose not to include this condition here; instead, we revisit whether minimal cues could suffice in the discussion (Section 6.1.2).

IV2: Task Duration. As a second independent variable, we varied task duration by designing two tasks with different complexity: a **medium**-duration task with 3 intermediate steps and a **long**-duration task with 6 intermediate steps. Figure 4 illustrates the two tasks along with the content and timing of intermediate updates and final responses for the respective feedback timing conditions.

As displayed, intermediate updates for PR were presented at fixed 5-second intervals; this corresponded to the empirical average step duration in our agentic in-car assistant prototype across different LLM models (GPT-4o [50], Claude-Sonnet-4-Thinking [6], Gemini-2.5-Flash-Thinking [67]). This interval is also below the 10-second threshold identified by [48] as the upper limit for maintaining user attention. We provide screen videos of the ProtoPie implementation

of the visual and auditive feedback for the different tasks in the supplementary material. To avoid confounding effects of verbosity in the timing study, the appropriate level of detail for feedback in the two task examples was determined in a pre-study with (N=7) institutional HCI experts working on user experience in the LLM-based in-car voice assistant. Across the eight tasks, interchangeable attributes (e.g., [fastest | shortest] route, [McDonald's | Bakery], [20% | 10%] battery) were permuted to minimize task memorization and repeated answers, while keeping the tasks conceptually equivalent. The user requests given to participants were deliberately explicit, hinting at the number and type of steps the assistant would take. Explicit requests aimed to reduce variability in expectations across participants with different technical backgrounds.

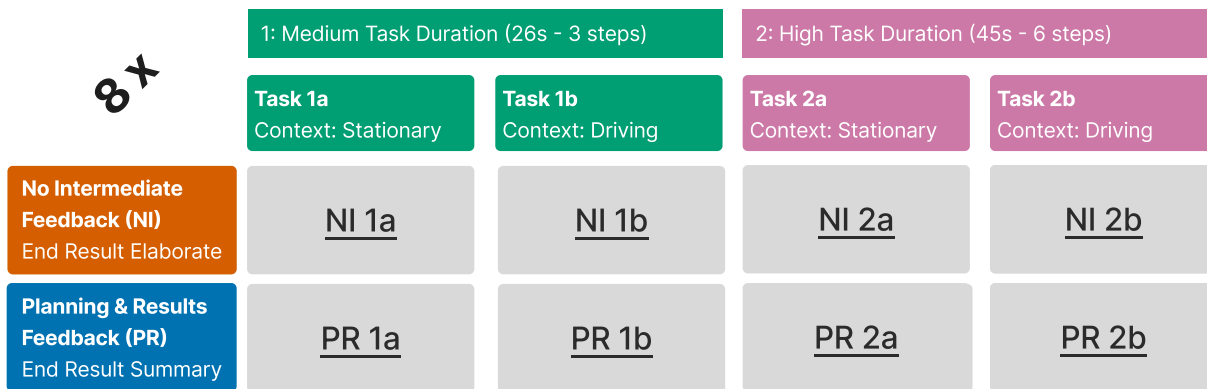
IV3: Interaction Context. As a third independent variable, we varied the interaction context between a single-task and a dual-task setting. In the **stationary** (single-task) condition, participants interacted with the voice assistant as their sole task while sitting in the stationary car-mockup. In the **driving** (dual-task) condition, participants additionally performed the driving-related lane-keeping task described in Section 4.1.1, while the voice-assistant interaction remained unchanged. The vehicle did not physically move; instead, the task simulated core attentional demands of driving. In this setting, the lane-keeping task naturally takes priority, relegating the voice assistant interaction to a secondary task. We included this manipulation as the concurrent task is expected to reduce available cognitive resources for processing feedback and may make waiting periods feel less idle, both potentially influencing preferred feedback timing and perceived system speed. We thus conceptualize this manipulation primarily as a single- versus dual-task comparison rather than an investigation of vehicle motion. While an alternative design, comparing manual with automated driving, could isolate task demands from perceived motion, our conditions capture ecologically valid scenarios: pre-trip navigation setup versus in-drive query.

4.3.2 Dependent Variables and Measurement. To capture the effects of feedback timing, task duration, and context, we measured four dependent variables (DVs) that reflect responsiveness, workload, user experience, and trust:

- **Perceived Speed (DV1):** A single custom 7-point Likert item (1 = very slow, 7 = very fast) (cf. [80]). This directly measures perceived responsiveness, the most immediate effect of feedback timing and task duration.
- **Task Load (DV2):** Three NASA-TLX [27] subscales (*Mental Demand*, *Temporal Demand*, *Frustration*) on a 0–100 scale (0 = very low, 100 = very high). These capture cognitive demands and frustration, particularly relevant in dual-task settings such as driving. Physical Demand, Performance, and Effort were excluded because tasks were non-physical and largely reactive. It should be noted that we report the unweighted average of the included subscales (Raw TLX, RTLX [26]); this allows for comparisons across our experimental conditions, due to subscale exclusion it should not be compared directly to full NASA-TLX scores in other experiments.
- **User Experience (DV3):** Three UEQ+ [57] subscales (*Attractiveness* (overall impression), *Dependability* (perceived

Table 1: Independent variables (IVs) with condition levels and descriptions. The total completion time is defined as the moment when the assistant finishes presenting its final response.

IV	Condition	Description
Feedback Timing	NI	No Intermediate Feedback: Only elaborate final response after task completion.
	PR	Planning&Result: Intermediate feedback for planned steps and intermediate results during processing and summarized response after task completion.
Task Duration	Medium	Task with 3 assistant steps and medium total completion time (26 s).
	High	Task with 6 assistant steps and high total completion time (45 s).
Interaction Context	Stationary	Single-activity: user interacts with system without concurrent tasks.
	Driving	Dual-activity: user interacts with system as a secondary activity while performing a driving-related task.

**Figure 3: Quantitative study design: Each participant completed 8 tasks across the 2x2x2 conditions.**

control and predictability), and *Risk Handling* (ability to detect and handle risks)) on a 7-point scale ($-3 =$ strongly negative, $+3 =$ strongly positive). These capture pragmatic and hedonic qualities of real-time feedback, allowing us to assess overall acceptance and perceived control. Subscales are reported individually, and an overall UX score was computed as a KPI, weighted by participants' self-reported importance following UEQ+ guidelines.

- **Trust (DV4):** The short form of the TIAS, S-TIAS [44], which covers *Confidence*, *Reliability*, and *Trustworthiness* on a 7-point scale (1 = not at all, 7 = extremely) and focuses on trust in artificial intelligence. Trust was measured to capture how different feedback strategies influence users' confidence in the assistant's behavior.

Table 2 summarizes the DVs, including when they were measured and the factorial structure applied in analysis.

By design, *User Experience* and *Task Load* were collapsed across the condition of task duration. Although task length may influence these measures, our focus was on comparing feedback systems across contexts; duration was included primarily to span a plausible range of complexity and duration, with detailed effects analyzed only for perceived speed, as this is the most sensitive measure of waiting. Similarly, *Trust* was measured once per feedback timing condition as we targeted system-level trust (confidence/reliability/trust in assistant) rather than context-specific

preferences. Testing differences between the study interaction context condition would provide limited generalizability given numerous possible contexts (stressful driving, media consumption, etc.) beyond our study scope. Notably, participants experienced both interaction contexts before providing their trust judgment, allowing them to form holistic system-level trust assessments. For the specific context variance included in our study, we expected user experience measures to provide more insightful metrics for context-related preferences. This design also reduces questionnaire burden.

4.3.3 Task Order and Measurement Timing. Participants completed all eight experimental conditions in a hierarchically counterbalanced order: (1) by feedback system (NI vs. PR), (2) by context (stationary vs. driving) within each system, and randomized (3) by task duration (medium vs. high) within each 2-task block. As shown in Figure 5, perceived speed was rated after each task, task load and user experience were measured after each two-task block, and trust was measured once per feedback system after four tasks.

4.3.4 Quantitative Analysis. We use repeated-measures ANOVAs with within-subject factors *Feedback Timing* (NI, PR), *Context* (Stationary, Driving), and, where applicable, *Task Duration* (Medium, High) or *Subscale*. Planned paired *t*-tests are used to unpack main and interaction effects. When multiple cell-wise comparisons are tested (Perceived Speed across Duration \times Context), we apply

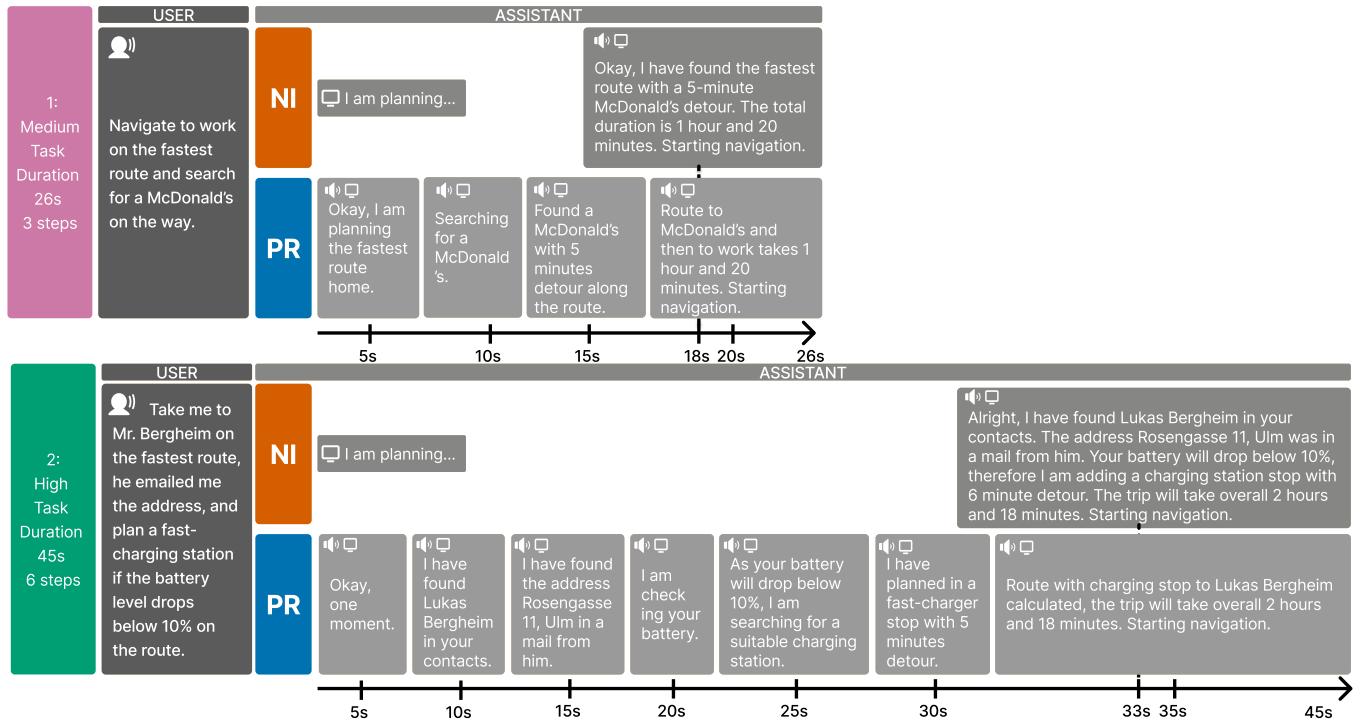


Figure 4: Quantitative study tasks: User requests for the two tasks with different durations, along with the assistant's final responses for the No Intermediate (NI) feedback system and assistant updates & final response for the Planning & Results (PR) feedback system. As the final answer is longer for the NI feedback system, it was started earlier (at 18s compared to 20s, respectively at 33s compared to 35s) so that the last output is at the same time for both systems. Note that at the beginning of both systems (NI and PR), a clicking sound, accompanied by the visual message "I am planning," is played to indicate the perception of the user request.

Table 2: Dependent variables (DVs) with measurement details.

DV	Measurement	Instrument	Design
Perceived Speed	After every task	7-pt Likert	$2 \times 2 \times 2$
Overall Experience	After 2-task block	UEQ+ subset	2×2 (Timing \times Context)
Task Load	After 2-task block	NASA-TLX subset	2×2 (Timing \times Context)
Trust	After 4-task block	S-TIAS	Paired (Timing)

a Holm correction. We report F , p , partial eta squared (η_p^2) for ANOVAs and t , 95% CI, and Cohen's d_z for paired contrasts. Additionally, for multi-item measures for one scale, internal consistency is assessed using Cronbach's α .

4.3.5 Quantitative Hypotheses. Based on prior work on latency, cognitive load, and trust in interactive systems, we formulate the following hypotheses for our quantitative study:

- **H1: Feedback Timing Effects (RQ1):** PR feedback (a) increases perceived speed compared to NI, (b) increases subjective task load compared to NI, (c) improves user experience, (d) increases user trust compared to NI.

- **H2: Interaction Context and Task Duration Effects (RQ1, RQ2):** (a) Driving context increases subjective task load compared to stationary and (b) longer task duration decreases perceived speed.

4.4 Qualitative Study

4.4.1 Semi-structured Interviews. We complemented the quantitative experiment with semi-structured interviews to explore how participants envision adaptive feedback systems over time (RQ3) and to contextualize findings on feedback timing, task complexity, and cognitive load (RQ1–RQ2). After completing the tasks, participants answered three open-ended questions (translated from German):

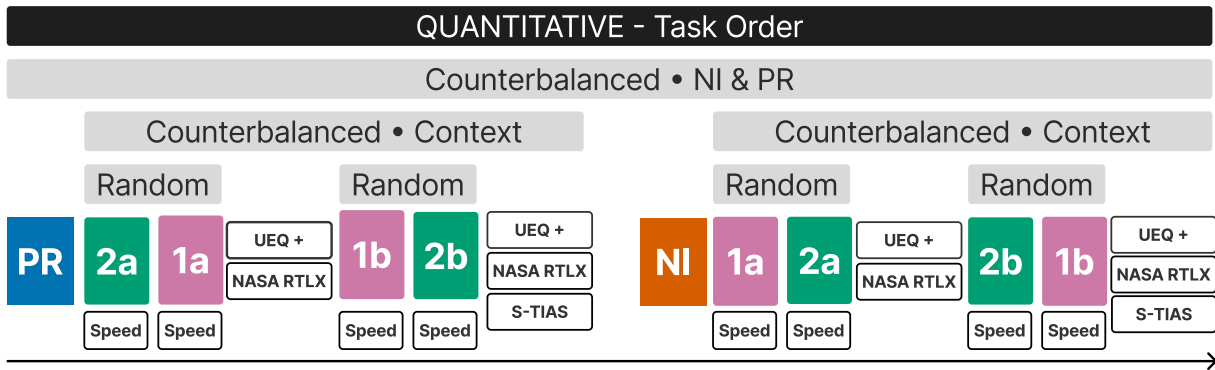


Figure 5: Task order and measurement timing: each participant performed all 8 tasks in hierarchically counterbalanced and randomized order. Dependent variables were measured (white boxes with black border) either after each task (perceived speed), after a 2-task block (UEQ+: user experience, NASA RTLX: task load), or after all 4 tasks per feedback system (S-TIAS: trust).

- (1) How much verbal feedback would you like from the system? Consider the driving situation, passengers, music, and other distractions.
- (2) Should the system notify you when it is uncertain, or decide autonomously? If notified, how should this be communicated?
- (3) Which system behaviors or experiences would foster long-term trust?

Follow-up prompts were used as needed to clarify or elaborate on participants' responses.

4.4.2 Qualitative Analysis. We transcribed the audio recordings of the 45 semi-structured interviews and cleaned the data. We then analyzed the transcripts using thematic analysis [12] and Atlas.ti. Two researchers independently open-coded a random subset of 20% of the data. During this phase, the focus was on maintaining close proximity to participants' language and experiences while generating granular codes such as *interruption reduction during media interaction* or *mute on demand*. The researchers then met in person to compare their initial codes, resolve discrepancies, and consolidate overlapping codes. Through this discussion, a shared codebook was developed consisting of 18 codes organized into preliminary conceptual groupings. Using this codebook, we divided the remaining transcripts equally for coding. Finally, the researchers reconvened to review coded extracts, refine code groups, and iteratively develop overarching themes. The researchers refined conceptual boundaries, distinguishing, for example, *external real-time adaptation* (media/social context) from *internal real-time adaptations* (task ambiguity, task novelty). This resulted in five themes capturing participants' preferences and rationales for feedback timing, verbosity, and adaptive feedback².

4.5 Participants

Table 3 shows the participants' distributions.

We recruited 45 participants (29 male, 16 female) from an automotive company, all of whom were above the age of 18. Age distribution was diverse, spanning 18–64 years, and covered the

typical age range for early adopters of in-car voice assistants. Participants reported varying familiarity with LLMs, voice assistants (VA), and in-car assistants. While most were familiar with general VAs (80%), familiarity with LLMs and the company voice assistant was lower, representing an expected sample regarding prior experience, given the duration of existence of general VAs and LLMs.

4.6 Limitations

Our work is subject to the following limitations. First, participants were recruited from a single automotive company. Although recruitment spanned multiple departments with diverse demographics and varying familiarity with LLMs and voice assistants, reflecting the intended customer base, generalization should be done cautiously.

Second, the driving context was simulated using a standardized lane-keeping task. This provided a consistent cognitive load increase across participants but cannot fully capture the variability of real-world driving, such as dynamic traffic or environmental distractions. Moreover, this manipulation inherently conflates perceived vehicle state with task demands. We hypothesize that the observed effects are primarily driven by the attentional demands of the concurrent task rather than perceived vehicle motion, since these demands directly constrain resources available for processing voice assistant feedback; future work comparing manual with perceived automated driving could empirically disentangle these factors. Similarly, intermediate feedback was provided at fixed 5s intervals to isolate the effect of feedback timing; adaptive or context-aware feedback policies were beyond the scope of this controlled experiment.

Third, our study captured immediate reactions to feedback timing and verbosity under the conditions of stationary vs. driving and medium vs. long task duration. Longitudinal adaptation over time and contextual adaptation (e.g., adjusting verbosity to cognitive load or passenger presence) were assessed only via self-reports in the qualitative interviews, rather than behavioral data from extended real-world deployments.

Finally, feedback was always provided simultaneously via voice and visual channels. We did not explore different modality combinations (e.g., intermediate feedback visually but not auditorily)

²The final codebook and theme structure are included in the supplementary material.

Table 3: Participant demographics and familiarity (N=45). LLM = Large Language Model, VA = Voice Assistant.

Category	Levels and Distribution
Gender	Male: 64%, Female: 36%
Age	18–24: 16%, 25–34: 44%, 35–44: 22%, 45–54: 13%, 55–64: 4%
LLM Familiarity	1=not: 4%, 2: 24%, 3: 24%, 4: 29%, 5=extremely: 18%
VA Familiarity	1=not: 0%, 2: 20%, 3: 40%, 4: 38%, 5=extremely: 2%
In-car VA Familiarity	1=not: 36%, 2: 29%, 3: 20%, 4: 11%, 5=extremely: 4%

or additional modalities (e.g., haptic cues for progress indication). Including these would have led to an unfeasibly large design space; we therefore focused on the most relevant independent variables to establish a controlled baseline before adding such complexity in future work.

5 Results

5.1 Quantitative Results

We first present an overview of feedback timing effects across all dependent variables (DVs) in Figure 6, followed by detailed analyses for each DV using repeated-measures ANOVAs and planned contrasts with t-tests.

The moderating effects of the interaction context and task duration conditions are presented subsequently.

Figure 6 shows a significant positive effect for the PR feedback system across all measured scores; we further unpack this in the subsequent sections.

5.1.1 Perceived Speed (H1a, H2b). A $2 \times 2 \times 2$ RM-ANOVA with factors *Timing* (PR vs. NI), *Duration* (Medium vs. High), and *Context* (Stationary vs. Driving) revealed a strong main effect of *Timing* ($F(1, 44) = 58.83, p < .001, \eta_p^2 = .57$). Collapsing across all conditions, perceived speed was significantly higher with a large effect under PR compared to NI feedback ($t(44) = 7.67, p < .001, 95\% \text{ CI } [0.90, 1.54], d_z = 1.01$) (ref. Figure 6), supporting **H1a**. Additionally, a three-way interaction *Timing* \times *Duration* \times *Context* emerged, $F(1, 44) = 4.09, p = .049, \eta_p^2 = .09$. Planned contrasts showed that PR feedback significantly outperformed NI feedback across all Duration \times Context combinations (all $p < .001, d_z = 0.58\text{--}0.95$), with the largest advantage for long tasks in the stationary single-task context ($d_z = 0.95$). We also observed a main effect of *Duration*, $F(1, 44) = 12.33, p = .001, \eta_p^2 = .22$, with longer tasks reducing perceived speed, $t(44) = -3.51, p = .001, 95\% \text{ CI } [-0.55, -0.15], d_z = -0.52$, supporting **H2b**. Planned contrasts revealed that the reduction in perceived speed from Medium to High duration was significant under NI feedback and Stationary context ($t(44) = -3.54, p = .001, 95\% \text{ CI } [-1.22, -0.53], d_z = -0.52$). This implies that intermediate feedback buffered the negative impact of longer task duration on perceived speed, especially during the single-task condition (cf. Figure 7).

5.1.2 Task Load (H1b, H2a). A $2 \times 2 \times 3$ RM-ANOVA with factors *Timing* (NI vs. PR), *Context* (Stationary vs. Driving), and *Subscale* (Mental Demand, Temporal Demand, Frustration) revealed a main effect of *Timing* ($F(1, 44) = 4.79, p = .034, \eta_p^2 = .10$), and an approaching but no significant main effect of *Context* ($F(1, 44) = 3.96,$

$p = .053$). Collapsing across subscales and contexts, task load was, unexpectedly, significantly lower with a small effect under PR compared to NI feedback ($t(44) = -2.19, p = .034, 95\% \text{ CI } [-8.54, -0.35], d_z = -0.26$) (ref. Figure 6), contradicting **H1b**, which expected PR to *increase* task load due to multiple interaction points. Further planned contrasts showed this reduction was primarily driven by the *Frustration* subscale ($t(44) = -2.04, p = .047, 95\% \text{ CI } [-12.81, -0.08], d_z = -0.26$), while Mental and Temporal Demand did not differ significantly (all $p > .10$). Cronbach's $\alpha = .77$ indicated acceptable internal consistency for the composite. Finally, although task load tended to be higher in the driving than the stationary context, this effect did not reach significance ($p = .053$), thus providing no clear support for **H2a**.

5.1.3 User Experience (H1c). A $2 \times 2 \times 3$ RM-ANOVA with factors *Timing* (NI vs. PR), *Context* (Stationary vs. Driving), and *Subscale* (Attractiveness, Dependability, Risk Handling) revealed a strong main effect of *Timing* ($F(1, 44) = 12.09, p = .001, \eta_p^2 = .22$), but no main effect of *Context* ($p = .85$) and no *Timing* \times *Context* interaction ($p = .32$). Collapsing across contexts, intermediate feedback (PR) significantly improved all three user experience subscales with small to medium effect compared to final-only feedback (NI): Attractiveness ($t(44) = 2.29, p = .027, 95\% \text{ CI } [0.04, 0.66], d_z = 0.38$), Dependability ($t(44) = 2.44, p = .019, 95\% \text{ CI } [0.07, 0.72], d_z = 0.47$), and Risk Handling ($t(44) = 3.90, p < .001, 95\% \text{ CI } [0.35, 1.10], d_z = 0.60$), with the strongest effect observed for Risk Handling. Cronbach's α indicated good reliability for all multi-item measures of the subscales ($\alpha = .74\text{--}.86$). Using the UEQ+ KPI weighting by self-reported importance of each subscale, PR feedback also improved the overall user experience composite score with a medium effect size ($t(44) = 3.30, p = .002, 95\% \text{ CI } [0.06, 0.24], d_z = 0.54$) (ref. Figure 6), supporting **H1c**.

5.1.4 User Trust (H1d). Trust was measured once per participant for each feedback condition after completing all tasks for that condition. Collapsing across all contexts and durations, trust ratings were significantly higher with a small effect under intermediate feedback (PR) compared to final-only feedback (NI) ($t(44) = 2.10, p = .042, 95\% \text{ CI } [0.01, 0.60], d_z = 0.38$), supporting **H1d**. Internal consistency for the three S-TIAS items (confidence in the system, reliability, and trustworthiness) was good (Cronbach's $\alpha = 0.84$). At the subscale level, PR feedback significantly improved Reliability ($t(44) = 2.20, p = .033, 95\% \text{ CI } [0.03, 0.60], d_z = 0.41$) and Trustworthiness ($t(44) = 2.41, p = .020, 95\% \text{ CI } [0.05, 0.61], d_z = 0.34$), but not Confidence ($p = .194$). Counterbalancing ensured that overall trust scores were not affected by order. Nevertheless, as trust can be learned through experience [29], we tested for order effects. A 2×2




Figure 6: Scores for dependent variables of post-hoc t-tests when contrasting the feedback timing systems (NI vs. PR) collapsed across context and duration conditions. All scores show significant effect for the PR feedback timing: perceived speed shows a large effect ($p < .001, ***$, 95% CI [0.90, 1.54], $d_z = 1.01$), user trust shows a small effect ($p = 0.042, *$, 95% CI [0.01, 0.60], $d_z = 0.38$), user experience KPI value showed a moderate effect ($p = 0.002, **$, 95% CI [0.06, 0.24], $d_z = 0.54$), and task load showed a small effect ($p = 0.034, *$, 95% CI [-8.54, -0.35], $d_z = -0.26$).




Figure 7: Perceived speed by task duration (medium, 26s vs. high, 45s) and interaction context (stationary vs. driving) for NI vs. PR feedback timing. Intermediate feedback reduced the negative slope in perceived speed for longer tasks.

mixed ANOVA with *Timing* (NI vs. PR) as a within-subject factor and *Order* (NI-first vs. PR-first) revealed no significant main effect of *Order* ($p = .418$) and no significant *Timing* × *Order* interaction ($p = .083$).

5.1.5 Summary of Duration and Context Effects. Perceived speed showed a significant *Timing* × *Duration* interaction ($p = .049$), where intermediate feedback buffered the negative effect of longer

tasks. For the other DVs, the effect of duration was not separately measured. Interaction *Context* (stationary vs. driving) produced no consistent main or interaction effects; a trend toward higher task load while driving ($p = .053$) did not reach significance, and no *Timing* × *Context* interactions were observed. Overall, intermediate feedback consistently improved perceived speed, user experience, and trust while unexpectedly reducing task load; duration mainly

affected perceived speed, whereas interaction context showed no consistent effects across dependent variables.

5.1.6 Summary of Demographic and Familiarity Effects. No significant main effects of age or familiarity with LLMs, voice assistants, or the company’s voice assistant were observed on any dependent variable. However, moderation analyses revealed that higher LLM familiarity significantly amplified improvements in trust ($p = .017$, $d_z = 0.74$) and user experience ($p = .021$, $d_z = 0.72$) from NI to PR feedback conditions. This pattern may reflect PR feedback’s structural similarity to Chain-of-Thought [73] reasoning outputs typical of LLMs, though alternative explanations warrant investigation. No additional moderation effects reached significance. Future research with a larger sample size may better capture the influence of demographic variables and technical familiarity.

5.2 Qualitative Findings

Complementing the quantitative results, the interviews after completing all 8 tasks reveal deeper insights into why participants preferred intermediate feedback and how they envision future adaptive systems (RQ3). Through inductive thematic analysis of 45 semi-structured interviews, we identified five major themes that contextualize our experimental results. Below, we detail each theme, illustrating findings with representative participant quotes (P1–P45). All quotes are translated from German into English.

5.2.1 T1: Longitudinal adaptation should be gated by trust and enabled by learning. Across interviews, participants emphasized that feedback verbosity should decrease over time when the system has proven itself and learned their routines. Most described trust as the primary precondition for reducing transparency; some explicitly tied this to repeated successful outcomes and predictable behavior. As P1 noted, “Trust only grows over time-when it delivers good results consistently”. Echoing this, P28 suggested, “With more trust, it can say less”. Participants also articulated how such reductions would be operationalized: the system should recognize recurrent tasks and remember prior choices to streamline future interactions. P9 commented, “For recurring complex tasks, it could keep it shorter”. Similarly, P35 proposed, “It should remember my answers and ask fewer questions next time”. While participants generally endorsed this progression, they also implied that reduced transparency should remain reversible if trust is challenged.

5.2.2 T2: Real-time external adaptation: balancing responsiveness with social and media context. Participants advocated for feedback that adapts to the surrounding situation, particularly media and social contexts, yet views diverged on how much to reduce interruptions. Several participants preferred minimal intermediate speech while listening to music or podcasts; P39 said, “With music on, one condensed summary would be better”. P2 added, “It’s annoying if things are said multiple times during a podcast”. In contrast, others wanted consistent output regardless of concurrent media: “Even with music, I want the output when I ask for it” (P31; see also P35, P42). Social presence introduced additional sensitivity: some felt continuous intermediate speech could be intrusive with passengers—“With a passenger, intermediate steps could be more tiring” (P13), and preferred a single end summary to avoid disrupting conversation (e.g., P22).

5.2.3 T3: Real-time internal adaptation: task ambiguity, stakes, and novelty heighten transparency needs. Participants reached a clear consensus that ambiguity requires clarification independent of trust. “It should ask follow-up questions when something is ambiguous” (P6; see also P7, P10). They further differentiated between high- and low-stakes actions: for decisions with higher consequence or higher error cost, most participants wanted verification and intermediate checkpoints: “With contacts, it’s important to ask” (P4); “With emails, a lot can go wrong” (P13). In contrast, for low-stakes actions (e.g., choosing a fast-food stop), many preferred the assistant to proceed with minimal dialog: “McDonald’s, just take the faster one” (P5). Finally, participants highlighted novelty as a cue for more transparency: “As long as the information is new and not repeated, it’s relevant to me” (P16).

5.2.4 T4: Active user control as a safety valve for feedback. Regardless of context policy, there was broad support for lightweight, user-driven controls to modulate verbosity. Participants repeatedly requested a way to mute or dampen spoken feedback when needed. P23 stated, “I want to choose how much information I get, mute is very important”. P11 was even more direct: “I should be able to tell it not to talk”. Participants suggested using such controls especially during media playback or when passengers are present (e.g., P44).

5.2.5 T5: Progressive chunking lightens cognitive load compared to end-only “dumps”. Participants contrasted progressive, small updates with a cognitively heavier “all at once” delivery. P1 put it succinctly: “It’s the same information, but a complete dump is harder to absorb”. Several echoed that stepwise updates made the process feel lighter and easier to follow than a dense end summary.

6 Discussion & Implications

6.1 Feedback Timing (RQ1, RQ2)

6.1.1 Intermediate feedback outperforms across conditions. Our quantitative results demonstrate that intermediate updates substantially improved perceived speed with a large effect, user experience with a medium effect, and trust with a small effect while lowering perceived frustration and task load with small effects (cf. Section 5.1). This pattern extends decades of research on responsiveness. Early HCI studies showed that unexplained delays reduced perceived speed [43], responsiveness improved overall user experience [46], and unexpected waiting increased frustration [58]. More recent work shows that explanations during delays were shown to increase trust [80]. Our results extend these findings to the context of agentic LLM-based assistants, where delays are not just incidental but inherent to handling multi-intent requests and extended reasoning steps. While participants were briefed on these mechanisms and given explicit requests that hinted at multi-step processing, latency may still have felt unexpected at times due to misaligned mental models [33, 63].

Intermediate updates were particularly helpful for longer, more complex tasks: perceived speed declined under final-only feedback but was buffered by stepwise updates. This underscores the importance of sustaining a sense of progress as task duration increases. The strongest effects were observed when interacting without a secondary task, suggesting that idle waiting is especially pronounced. While benefits were also observed during dual-task interaction, the

moderation effects were weaker, and for other dependent variables, we did not find significant interactions with context or duration. This highlights robustness but also underscores the need for in-the-wild validation under more varied dual-task demands.

Design implication: From the user’s perspective, agentic in-car assistants should provide intermediate updates during long-running, multi-step tasks if possible, particularly as task complexity grows. This guideline is most critical in early phases of adoption when overall trust is still developing. Benefits hold across both stationary and driving contexts, but may vary with situational trust or when competing attentional demands alter user priorities (see Section 6.2.2).

6.1.2 Are simple progress cues enough, or must updates include content? A natural design question is whether lightweight progress indicators (such as auditory fillers, e.g., “working on it”, or visual and haptic cues) could substitute for content-rich updates, especially once trust has been established. Our results and related work suggest otherwise.

Early work on “ambiguous silence” showed that minimal feedback leaves users uncertain about what the system is doing [78], a problem documented in voice interfaces [55] and reflected in theories of grounding: communication requires not just perception-level confirmation (“I heard you”) but also understanding-level evidence (“I understood what you meant”)[4, 9, 19]. If intermediate steps are hidden until the final response, users are deprived of understanding-level feedback throughout the wait. This forces them into laborious checking behaviors to re-establish common ground, a dynamic well described in conversational grounding research [14].

Trust research supports this interpretation: Zhang et al. [80] found that simply notifying users of a delay is insufficient, while providing justifications for what is happening increases trust. Our workload results reinforce the point that when identical information is delivered in smaller steps, task load decreases significantly, with a small effect size, compared to receiving a single condensed final response. Participant explanations (T5) complement this: they described stepwise updates as cognitively lighter than dense “dumps”. This is also important to enable effective human oversight, as Sterz et al. [62] outlines that epistemic access, including comprehending what the agent is doing, is important.

An analogy to visual driver distraction guidelines is interesting: AAM standards restrict individual glances to in-car displays while driving to two seconds and cumulative glances to complete one task to 20 seconds [13, 25]. A single long information dump parallels a prolonged glance, intensifying momentary distraction, whereas stepwise updates are akin to multiple brief glances - each less demanding and less disruptive to ongoing tasks.

Design implication: Content-bearing intermediate updates appear more effective than progress-only cues. They help preserve grounding, maintain trust, and distribute cognitive effort more evenly across time. This holds when the cognitive feedback channel (e.g., auditory) is available; if not, updates might be experienced as interruptive, and user preferences can differ (see Section 6.2.2).

One promising direction to gain efficiency while maintaining these benefits is through learned cue associations. As users learn mappings between simple cues and specific system actions, those

cues can effectively convey the content of what the assistant is doing without verbose language. Cho et al. [18] demonstrate this approach with direct multimodal feedback cues in repetitive tasks, reducing cognitive load while preserving informational value.

6.2 Feedback Verbosity (RQ3, RQ2)

6.2.1 Long-term verbosity adaption – gated by learned trust through demonstrated reliability. Participants (T1) reported that reductions in verbosity were only acceptable once the in-car assistant had demonstrated reliability. Trust was described as the decisive factor, developing over time through repeated successful interactions. This resonates with Hoff and Bashir [29] three layers of trust: dispositional, situational, and learned. In our qualitative findings, learned trust - understood as users’ emerging confidence grounded in experienced reliability - was central for long-term verbosity adaption. As the system repeatedly performed well, participants expressed willingness to accept less transparency in favor of greater efficiency. This illustrates a trajectory where transparency, represented by intermediate feedback, initially outweighs efficiency due to its benefits for perceived speed, trust, and task load. Over time, as learned trust grows, verbosity can be scaled back without eroding confidence in the system. Reversibility was mentioned (T1) by some participants, but is inherently part of this learned-trust-dependent adaptation: when reliability falters, transparency must reappear.

Design implication: Our empirical findings suggest that feedback verbosity should start high to establish transparency, then decrease as the system demonstrates reliability. This trajectory treats demonstrated reliability as a practical proxy for learned trust, enabling efficiency without undermining confidence.

6.2.2 Real-time situational context adaptation - internal vs. external factors. Participants described the need for real-time adjustments to feedback, shaped by both internal task factors and external situational factors. For internal factors (T3), situational trust was crucial: when a task was novel or ambiguous, participants sought greater transparency, regardless of their prior experience. This supports prior work showing that explanations are most valuable when outcomes are uncertain or likely to surprise the user [61]. Similarly, high-stakes requests such as contacting people or handling emails were seen as requiring verification, whereas routine or low-stakes actions could be handled more efficiently with minimal verbosity. This highlights the need for transparency to scale with task stakes, enabling human control and potential intervention, which Dietvorst et al. [22] found to increase user trust .

By contrast, external context factors (T2), such as media consumption or social interaction in the car, produced divergent preferences. Some participants preferred a condensed, final-only summary to avoid interruptions, while others valued consistent updates independent of distractions. Because no uniform policy fits all, user-facing controls (e.g., mute button or voice command) become an effective and wanted (T4) practical way to resolve mismatches between system adaptation and individual preferences. Although not explicitly mentioned by participants, we hypothesize that verbosity controls should also permit the expansion of detail on demand – for instance, through the visual unfolding of a step within an interface element, or in response to user-initiated clarification requests.

Design implication: Our qualitative findings suggest that verbosity should be increased when task factors are novel, ambiguous, or high-stakes and can be reduced for routine or low-stakes requests. External context adaptation remains more contested: until robust personalization methods exist, systems might provide simple user controls for in-the-moment overrides.

6.3 Applicability across domains

Our design implications could inform considerations for other user-facing agentic systems beyond in-car assistants, though direct transferability requires further investigation. We identify two main user-agent interaction settings where our findings may offer relevant insights: (i) when the agent constitutes the primary task, and (ii) in dual-task scenarios where primary and secondary activities rely on different cognitive channels. In contrast, when both tasks share the same channel (such as coding copilots, writing assistants, or listening to media while driving), intermediate updates may risk interference rather than relief. Since we only qualitatively found divergent preferences, we cannot yet offer implications in this regard.

In primary-task systems, such as social or service robots [37, 59, 82] and customer-support agents [8, 17], our findings on adaptive feedback detail and keeping informative updates to preserve grounded communication appear particularly relevant. In dual-task contexts with different channels, examples include smart home assistants used during everyday activities, such as cooking [32], or emerging wearable agents that deliver feedback via audio or haptics alongside physical tasks. Here, along with feedback timing implications, our results on verbosity adaptation based on internal task factors and lightweight user control may help strike a balance between responsiveness and flexibility.

Regarding temporal scope, our implications target agentic systems with execution times ranging from multiple seconds to one minute. They likely do not extend to "deep agent" systems, such as OpenAI's Deep Research [52], which operate over multiple minutes to half an hour. Such extended durations naturally push interactions into background processing, as sustained intermediate feedback across many minutes would likely overwhelm users rather than maintain engagement. This distinction raises an interesting research question for future work: identifying the temporal threshold where agentic systems should transition from maintaining user attention to background operation.

Finally, the application of our implications ultimately also depends on domain- and modality-specific constraints. Just as driving imposes limits on distractions, other domains will have their own boundaries. Our insights may inform design considerations in other domains, but likely require adaptation to specific contextual constraints rather than direct transfer.

6.4 Future Design and Tech Challenges

Our empirical findings yield design implications for agentic in-car assistants, yet translating these into operational systems poses open challenges for further research. Rather than proposing concrete solutions, we outline design challenges and point to directions and related work that may inspire future work.

Intermediate feedback timing and content. Current LLM agents operate through sequential tool calls, and intermediate outputs can be prompted if they are supported natively (e.g., OpenAI's tool preambles [51]), or can be supplied by an additional asynchronous LLM. Open design challenges include coordinating overlapping intermediate voice outputs, and deciding which information to present in the informative updates versus deferring for follow-up clarification.

Long-term verbosity adaptation based on demonstrated reliability. Our empirical findings indicate that users prefer an adaptive system that reduces feedback verbosity, as they perceive sufficient reliability. This relates to Google's PAIR guidebook on trust, which identifies ten levers of trust [24]; our findings specifically reveal a dynamic relationship between two of these levers: *Reliance*, indicating how much users can depend on the system, and *Transparency*, which helps users understand and predict agent responses. While trust itself is a latent variable that cannot be directly measured [69, 76], the relationship between the more objective demonstrated reliability from interaction history and verbosity preferences may offer a more tractable direction for future system design. Prior work on online trust estimation, using Bayesian Networks to infer user trust states in robot-collaborative settings [77] or Finite State Automata to model trust from acceptance behaviors [70], may inform future operational approaches. Trust-related behavioral signals from interaction history, such as acceptances, interruptions, corrections, rejections, and override rates, are commonly used and may serve as practical proxies for demonstrated reliability, potentially yielding systems better aligned with the preferences we identified. Future work would need to determine which signals best estimate demonstrated reliability and identify appropriate verbosity adaptation levels.

Situational verbosity adaptation based on task novelty, stakes, and confidence. Novelty can be estimated from memory and task history [35, 54], and stakes are often estimable in closed domains with limited action space. In contrast, detecting ambiguity remains an open challenge and is an active research area [36]. Current LLMs' internal confidence assessments are yet poorly calibrated [34, 81], limiting reliable detection.

These challenges, informed by our empirical findings, offer concrete entry points for future systems research on agentic feedback.

7 Conclusion

In this work, we address the central challenge of how agentic assistants should communicate progress and manage information load during long-running tasks. Current systems vary from silent background operation to detailed step-by-step updates - the right balance is especially critical in dual-task settings where cognitive load is constrained. Through a controlled mixed-methods study of an in-car agentic assistant, we show that (1) intermediate informative updates improve trust, perceived speed, and user experience while reducing task load, and (2) feedback should adapt over time - starting transparent to establish trust, becoming more concise as reliability is demonstrated, and re-expanding situationally when tasks are ambiguous, novel, or high-stakes. These findings inform design considerations for adaptive feedback with potential transferability

beyond driving and in-car assistants, to other primary-task interactions and dual-task contexts where cognitive channels permit intermediate communication.

References

- [1] Deepak B. Acharya, Karthigeyan Kuppan, and B. Divya. 2025. Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey. *IEEE Access* 13 (2025), 18912–18936. doi:10.1109/ACCESS.2025.3532853
- [2] Manus AI. 2025. Manus. Retrieved September 10, 2025 from <https://manus.im/?index=1>
- [3] Perplexity AI. 2025. Perplexity. Retrieved September 10, 2025 from <https://www.perplexity.ai/>
- [4] Jens Allwood, Joakim Nivre, and Elisabeth Ahlsén. 1992. On the semantics and pragmatics of linguistic feedback. *Journal of semantics* 9, 1 (1992), 1–26. doi:10.1093/jos/9.1.1
- [5] Marco Almada. 2019. Human intervention in automated decision-making: Toward the construction of contestable systems. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law* (Montreal, QC, Canada) (ICAIL '19). Association for Computing Machinery, New York, NY, USA, 2–11. doi:10.1145/3322640.3326699
- [6] Anthropic. 2025. System Card: Claude Opus 4 & Claude Sonnet 4. Retrieved September 5, 2025 from <https://www-cdn.anthropic.com/6d8a8055020700718b0c49369f60816ba2a7c285.pdf>
- [7] Paul Atchley and Jeff Dressel. 2004. Conversation Limits the Functional Field of View. *Human Factors* 46, 4 (2004), 664–673. doi:10.1518/hfes.46.4.664.56808
- [8] Isabel Auer, Stephan Schlögl, and Gundula Glowka. 2024. Chatbots in Airport Customer Service—Exploring Use Cases and Technology Acceptance. *Future Internet* 16, 5 (2024), 175. doi:10.3390/fi16050175
- [9] Agnes Axelsson, Hendrik Buschmeier, and Gabriel Skantze. 2022. Modeling Feedback in Interaction With Conversational Agents—A Review. *Frontiers in Computer Science* 4 (2022), 744574. doi:10.3389/fcomp.2022.744574
- [10] Gagan Bansal, Jennifer W. Vaughan, Saleema Amershi, Eric Horvitz, Adam Fourney, Hussein Mozannar, Victor Dibia, and Daniel S. Weld. 2024. Challenges in Human-Agent Communication. doi:10.48550/arXiv.2412.10380
- [11] Vanessa Beanland, Michael Fitzharris, Kristie L. Young, and Michael G. Lenné. 2013. Driver inattention and driver distraction in serious casualty crashes: Data from the Australian National Crash In-depth Study. *Accident Analysis & Prevention* 54 (2013), 99–107. doi:10.1016/j.aap.2012.12.043
- [12] Ann Blandford, Dominic Furniss, and Stephann Makri. 2016. *Qualitative HCI Research: Going Behind the Scenes*. Springer, Cham. doi:10.1007/978-3-031-02217-3
- [13] Michael Braun, Nora Broy, Bastian Pfleging, and Florian Alt. 2019. Visualizing natural language interaction for conversational in-vehicle information systems to minimize driver distraction. *Journal on Multimodal User Interfaces* 13, 2 (2019), 71–88. doi:10.1007/s12193-019-00301-2
- [14] Susan E. Brennan. 1998. The Grounding Problem in Conversations With and Through Computers. In *Social and Cognitive Approaches to Interpersonal Communication* (1st ed.), Susan R. Fussell and Roger J. Kreuz (Eds.). Lawrence Erlbaum, Hillsdale, NJ, 201–225.
- [15] Stephen A. Brewster, Peter C. Wright, and Alistair D.N. Edwards. 1994. Guidelines for the creation of earcons. In *Proceedings of HCI'94*. Cambridge University Press, 747–759.
- [16] Gary Burnett, Elizabeth Crundall, David Large, Glyn Lawson, and Liesje de Harder. 2013. On-the-move and in your car: An overview of HCI issues for in-car computing. *International Journal of Mobile Human Computer Interaction* 5, 1 (2013), 1–21. doi:10.4018/jmhci.2009010104
- [17] Aili Chen, Xuyang Ge, Ziquan Fu, Yanghua Xiao, and Jiangjie Chen. 2024. TravelAgent: An AI Assistant for Personalized Travel Planning. doi:10.48550/arXiv.2409.08069
- [18] Hyunsung Cho, Jacqui Fashimpaur, Naveen Sendhilnathan, Jonathan Browder, David Lindbauer, Tanya R. Jonker, and Kashyap Todt. 2025. Persistent Assistant: Seamless Everyday AI Interactions via Intent Grounding and Multimodal Feedback. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 59, 19 pages. doi:10.1145/3706598.3714317
- [19] Herbert H. Clark. 1996. *Using language*. Vol. 35. Cambridge university press, Cambridge. 167–222 pages. doi:10.1017/S0022226798217361
- [20] Herbert H. Clark and Susan E. Brennan. 1991. Grounding in Communication. In *Perspectives on Socially Shared Cognition*, Lauren Resnick, Levine B., M. John, Stephanie Teasley, and D. (Eds.). American Psychological Association, 13–1991. doi:10.1037/10096-006
- [21] Cursor. 2025. The AI Code Editor. Retrieved September 10, 2025 from <https://cursor.com/home>
- [22] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2018. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management science* 64, 3 (2018), 1155–1170.
- [23] Markus Funk, Carie Cunningham, Duygu Kanver, Christopher Saikalis, and Rohan Pansare. 2020. Usable and Acceptable Response Delays of Conversational Agents in Automotive User Interfaces. In *12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Virtual Event, DC, USA) (AutomotiveUI '20). Association for Computing Machinery, New York, USA, 262–269. doi:10.1145/3409120.3410651
- [24] Google. 2025. Trust + Explanations Help people recover from errors. Retrieved September 10, 2025 from <https://pair.withgoogle.com/guidebook/chapters/trust-and-explanations/understanding-trust-in-your-product>
- [25] Group ADFTW. 2006. *Statement of principles, criteria and verification procedures on driver interactions with advanced in-vehicle information and communication systems*. Technical Report. Alliance of Automobile Manufacturers.
- [26] Sandra G. Hart. 2006. Nasa-Task Load Index (NASA-TLX); 20 Years Later. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 50. Sage publications, Los Angeles, USA, 904–908. doi:10.1177/154193120605000909
- [27] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [28] Gaole He, Gianluca Demartini, and Ujwal Gadiraju. 2025. Plan-Then-Execute: An Empirical Study of User Trust and Team Performance When Using LLM Agents As A Daily Assistant. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 414, 22 pages. <https://doi.org/10.1145/3706598.3713218>
- [29] Kevin A. Hoff and Masooda Bashir. 2015. Trust in Automation: Integrating Empirical Evidence on Factors That Influence Trust. *Human Factors* 57, 3 (2015), 407–434. doi:10.1177/0018720814547570
- [30] William J Horrey and Christopher D Wickens. 2006. Examining the impact of cell phone conversations on driving using meta-analytic techniques. *Human Factors* 48, 1 (2006), 196–205. doi:10.1518/001872006776412135
- [31] Soodeh Hosseini and Hossein Seilani. 2025. The role of agentic AI in shaping a smart future: A systematic review. *Array* 26 (2025), 100399. doi:10.1016/j.array.2025.100399
- [32] Razan Jaber, Sabrina Zhong, Sanna Kuoppamäki, Aida Hosseini, Iona Gessinger, Duncan P. Brumby, Benjamin R. Cowan, and Donald Mcmillan. 2024. Cooking With Agents: Designing Context-aware Voice Interaction. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, USA) (CHI '24). Association for Computing Machinery, New York, USA, Article 551, 13 pages. doi:10.1145/3613904.3642183
- [33] Philip N. Johnson-Laird. 1986. *Mental models: towards a cognitive science of language, inference, and consciousness*. Harvard University Press, USA.
- [34] Adam T. Kalai, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang. 2025. Why Language Models Hallucinate. doi:10.48550/arXiv.2509.04664
- [35] Johannes Kirmayr, Lukas Stappen, Phillip Schneider, Florian Matthes, and Elisabeth Andre. 2025. CarMem: Enhancing Long-Term Memory in LLM Voice Assistants through Category-Bounding. In *Proceedings of the 31st International Conference on Computational Linguistics: Industry Track*. Association for Computational Linguistics, Abu Dhabi, UAE, 343–357.
- [36] Katarzyna Kobalcyk, Nicolas Astorga, Tennison Liu, and Mihaela van der Schaar. 2025. Active Task Disambiguation with LLMs. In *The Thirteenth International Conference on Learning Representations*. Singapore. doi:10.48550/arXiv.2502.04485
- [37] Matthias Kraus, Nicolas Wagner, Nico Untereiner, and Wolfgang Minker. 2022. Including Social Expectations for Trustworthy Proactive Human-Robot Dialogue. In *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization* (Barcelona, Spain) (UMAP '22). Association for Computing Machinery, New York, USA, 23–33. doi:10.1145/3503252.3531294
- [38] Markus Langer, Kevin Baum, and Nadine Schlicker. 2024. Effective Human Oversight of AI-Based Systems: A Signal Detection Perspective on the Detection of Inaccurate and Unfair Outputs. *Minds and Machines* 35, 1 (Nov. 2024), 1. doi:10.1007/s11023-024-09701-0
- [39] John D. Lee and Katrina A. See. 2004. Trust in Automation: Designing for Appropriate Reliance. *Human Factors* 46, 1 (2004), 50–80. doi:10.1518/hfes.46.1.50_30392
- [40] Bingjie Liu. 2021. In AI We Trust? Effects of Agency Locus and Transparency on Uncertainty Reduction in Human-AI Interaction. *Journal of Computer-Mediated Communication* 26, 6 (09 2021), 384–402. arXiv:<https://academic.oup.com/jcmc/article-pdf/26/6/384/41139653/zmab013.pdf> doi:10.1093/jcmc/zmab013
- [41] Victor Ei-Wen Lo and Paul A. Green. 2013. Development and Evaluation of Automotive Speech Interfaces: Useful Information from the Human Factors and the Related Literature. *International Journal of Vehicular Technology* 2013, 1 (2013), 924170. doi:10.1155/2013/924170
- [42] Alexandra Loew, Ina Koniakowsky, Yannick Forster, Frederik Naujoks, and Andreas Keinath. 2023. The impact of speech-based assistants on the driver's cognitive distraction. *Accident Analysis & Prevention* 179 (2023), 106898. doi:10.1016/j.aap.2022.106898
- [43] Mykola Maslych, Mohammadreza Katebi, Christopher Lee, Yahya Hmaiti, Amirpouya Ghasemaghaei, Christian Pumarada, Jannese Palmer, Esteban

- Segarra Martinez, Marco Emporio, Warren Snipes, Ryan P. McMahan, and Joseph J. LaViola Jr. 2025. Mitigating Response Delays in Free-Form Conversations with LLM-powered Intelligent Virtual Agents. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25, 49)*. Association for Computing Machinery, New York, USA, 15. doi:10.1145/3719160.3736636
- [44] Melanie J. McGrath, Oliver Lack, James Tisch, and Andreas Duenser. 2025. Measuring Trust in Artificial Intelligence: Validation of an Established Scale and Its Short Form. *Frontiers in Artificial Intelligence* 8 (2025). doi:10.3389/frai.2025.1582880
- [45] Siddharth Mehrotra, Chadha Degachi, Oleksandra Vereschak, Catholijn M. Jonker, and Myrthe L. Tielman. 2024. A Systematic Review on Fostering Appropriate Trust in Human-AI Interaction: Trends, Opportunities and Challenges. *ACM Journal on Responsible Computing* 1, 4, Article 26 (2024), 45 pages. doi:10.1145/3696449
- [46] Robert B. Miller. 1968. Response Time in Man-Computer Conversational Transactions. In *Proceedings of the December 9-11, 1968, Fall Joint Computer Conference, Part I (AFIPS '68 (Fall, Part I))*. Association for Computing Machinery, New York, USA, 267–277. doi:10.1145/1476589.1476628
- [47] Brad A. Myers. 1985. The Importance of Percent-Done Progress Indicators for Computer-Human Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '85)*. Association for Computing Machinery, New York, USA, 11–17. doi:10.1145/317456.317459
- [48] Jakob Nielsen. 1994. *Usability Engineering*. Morgan Kaufmann Publishers, San Francisco, USA.
- [49] Donald A. Norman. 2002. *The design of everyday things*. Basic Books, New York, USA.
- [50] OpenAI. 2024. GPT-4o System Card. doi:10.48550/arXiv.2410.21276
- [51] OpenAI. 2025. GPT-5 prompting guide - Tool preambles. Retrieved September 8, 2025 from https://cookbook.openai.com/examples/gpt-5/gpt-5_prompting_guide#tool-preambles
- [52] OpenAI. 2025. Introducing deep research. Retrieved September 5, 2025 from <https://openai.com/index/introducing-deep-research/>
- [53] Sharon Oviatt and Phil Cohen. 2000. Multimodal interfaces that process what comes naturally. *Commun. ACM* 43, 3 (2000), 45–53. doi:10.1145/330534.330538
- [54] Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. 2024. MemGPT: Towards LLMs as Operating Systems. doi:10.48550/arXiv.2310.08560
- [55] Martin Porcheron, Joel E. Fischer, Stuart Reeves, and Sarah Sharples. 2018. Voice Interfaces in Everyday Life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, USA, 1–12. doi:10.1145/3173574.3174214
- [56] Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: language models can teach themselves to use tools. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (New Orleans, USA) (NIPS '23)*. Curran Associates Inc., Red Hook, USA, Article 2997, 13 pages.
- [57] Martin Schrepp and Jörg Thomaschewski. 2019. Design and Validation of a Framework for the Creation of User Experience Questionnaires. *International Journal of Interactive Multimedia and Artificial Intelligence* 5, 7 (2019), 88–95. doi:10.9781/ijimai.2019.06.006
- [58] Ben Shneiderman. 1984. Response Time and Display Rate in Human Performance with Computers. *Comput. Surveys* 16, 3 (1984), 265–285. doi:10.1145/2514.2517
- [59] Christina S. Song and Youn-Kyung Kim. 2022. The role of the human-robot interaction in consumers' acceptance of humanoid retail service robots. *Journal of Business Research* 146, 2 (2022), 489–503. doi:10.1016/j.jbusres.2022.03.087
- [60] Lenja Sorokin, Khanh Huynh, Malin Eiband, Lukas Stappen, and Jeremy Dillmann. 2025. Collaborating with LLMs Through a Voice and Graphical User Interface. In *Adjunct Proceedings of the 27th International Conference on Mobile Human-Computer Interaction (MobileHCI '25 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 3, 8 pages. doi:10.1145/3737821.3749555
- [61] Sarath Sreedharan, Tathagata Chakraborti, and Subbarao Kambhampati. 2021. Foundations of explanations as model reconciliation. *Artificial Intelligence* 301 (2021), 103558. doi:10.1016/j.artint.2021.103558
- [62] Sarah Sterz, Kevin Baum, Sebastian Biewer, Holger Hermanns, Anne Lauber-Rönsberg, Philip Meinel, and Markus Langer. 2024. On the Quest for Effectiveness in Human Oversight: Interdisciplinary Perspectives. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (Rio de Janeiro, Brazil) (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 2495–2507. doi:10.1145/3630106.3659051
- [63] Mark Steyvers, Heliodoro Tejeda, Aakriti Kumar, Catarina Belem, Sheer Karny, Xinyue Hu, Lukas Mayer, and Padhraic Smyth. 2025. What Large Language Models Know and What People Think They Know. *Nature Machine Intelligence* 7, 2 (2025), 221–231. doi:10.1038/s42256-024-00976-7
- [64] David L. Strayer, Joel M. Cooper, Rachel M. Goethe, Madeleine M. McCarty, Douglas J. Getty, and Francesco Biondi. 2019. Assessing the Visual and Cognitive Demands of In-Vehicle Information Systems. *Cognitive Research: Principles and Implications* 4, 1 (2019), 18. doi:10.1186/s41235-019-0166-3
- [65] David L. Strayer, Joel M. Cooper, Jonna Turrill, James R. Coleman, and Rachel J. Hopman. 2016. Talking to your car can drive you to distraction. *Cognitive Research: Principles and Implications* 1, 1 (2016), 16. doi:10.1186/s41235-016-0018-3
- [66] David L. Strayer, Jonna Turrill, Joel M. Cooper, James R. Coleman, Nathan Medeiros-Ward, and Francesco Biondi. 2015. Assessing Cognitive Distraction in the Automobile. *Human Factors* 57, 8 (2015), 1300–1324. doi:10.1177/0018720815575149
- [67] Google Gemini Team. 2025. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. doi:10.48550/arXiv.2507.06261
- [68] Takane Ueno, Yuto Sawa, Yeongdae Kim, Jacqueline Urakami, Hiroki Oura, and Katie Seaborn. 2022. Trust in Human-AI Interaction: Scoping Out Models, Measures, and Methods. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, USA) (CHI EA '22)*. Association for Computing Machinery, New York, USA, Article 254, 7 pages. doi:10.1145/3491101.3519772
- [69] Oleksandra Vereschak, Gilles Bailly, and Baptiste Caramiaux. 2021. How to Evaluate Trust in AI-Assisted Decision Making? A Survey of Empirical Methodologies. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 327 (Oct. 2021), 39 pages. doi:10.1145/3476068
- [70] Maria Virvou, George A. Tsihrintzis, and Evangelia-Aikaterini Tsihrintzi. 2024. VIRTSE: A novel trust dynamics model enhancing Artificial Intelligence collaboration with human users – Insights from a ChatGPT evaluation study. *Information Sciences* 675 (2024), 120759. doi:10.1016/j.ins.2024.120759
- [71] Michael Vössing, Niklas Kühl, Matteo Lind, and Gerhard Satzger. 2022. Designing Transparency for Effective Human-AI Collaboration. *Information Systems Frontiers* 24, 3 (June 2022), 877–895. doi:10.1007/s10796-022-10284-3
- [72] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne X. Zhao, Zhewei Wei, and Jirong Wen. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science* 18, 6 (2024), 186345. doi:10.1007/s11704-024-40231-1
- [73] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '22)*. Curran Associates Inc., Red Hook, NY, USA, Article 1800, 14 pages.
- [74] Mathew P. White, J. Richard Eiser, and Peter R. Harris. 2004. Risk Perceptions of Mobile Phone Use While Driving. *Risk Analysis* 24, 2 (2004), 323–334. doi:10.1111/j.0272-4332.2004.00434.x
- [75] Christopher D. Wickens. 2008. Multiple resources and mental workload. *Human Factors* 50, 3 (2008), 449–455. doi:10.1518/001872008X288394
- [76] Yaqi Xie, Indu P. Bodala, Desmond C. Ong, David Hsu, and Harold Soh. 2020. Robot capability and intention in trust-based decisions across tasks. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction (Daegu, Republic of Korea) (HRI '19)*. IEEE Press, 39–47.
- [77] Anqi Xu and Gregory Dudek. 2015. Optimo: Online probabilistic trust inference model for asymmetric human-robot collaborations. In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*. 221–228.
- [78] Nicole Yankelovich, Gina-Anne Levow, and Matt Marx. 1995. Designing SpeechActs: issues in speech user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Denver, USA) (CHI '95)*. ACM Press/Addison-Wesley Publishing Co., USA, 369–376. doi:10.1145/223904.223952
- [79] Tingru Zhang, Xing Liu, Weisheng Zeng, Da Tao, Guofa Li, and Xingda Qu. 2023. Input modality matters: A comparison of touch, speech, and gesture based in-vehicle interaction. *Applied Ergonomics* 108 (2023), 103958. doi:10.1016/j.apergo.2022.103958
- [80] Zhengquan Zhang, Konstantinos Tsiakos, and Christina Schneegass. 2024. Explaining the Wait: How Justifying Chatbot Response Delays Impact User Trust. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces (Luxembourg, Luxembourg) (CUI '24)*. Association for Computing Machinery, New York, USA, Article 27, 16 pages. doi:10.1145/3640794.3665550
- [81] Chiwei Zhu, Benfeng Xu, Quan Wang, Yongdong Zhang, and Zhendong Mao. 2023. On the Calibration of Large Language Models and Alignment. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 9778–9795. doi:10.18653/v1/2023.findings-emnlp.654
- [82] Albert Lukasik and Arkadiusz Gut. 2025. From robots to chatbots: unveiling the dynamics of human-AI interaction. *Frontiers in Psychology* Volume 16 - 2025 (2025). doi:10.3389/fpsyg.2025.1569277

A Survey

The following questions are translated from german into english.

A.1 Demographics & Technical Familiarity

- Age
 - Under 18
 - 18–24 years
 - 25–34 years
 - 35–44 years
 - 45–54 years
 - 55–64 years
 - 65 years and older
- How do you describe yourself?
 - Male
 - Female
 - Non-binary / third gender
 - Self-description preferred
 - Prefer not to answer
- How familiar are you with the general capabilities and functionalities of Large Language Models (LLMs) such as ChatGPT?
 - Not at all familiar
 - Somewhat familiar
 - Familiar
 - Very familiar
 - Extremely familiar
- How familiar are you with voice assistants such as Alexa, Siri, Google Assistant, etc.?
 - Not at all familiar
 - Somewhat familiar
 - Familiar
 - Very familiar
 - Extremely familiar
- How familiar are you with the companies in-car voice assistant?
 - Not at all familiar
 - Somewhat familiar
 - Familiar
 - Very familiar
 - Extremely familiar

A.2 Questionnaires

A.2.1 Perceived Speed.

- How fast or slow did you perceive the system during the task?
 - Very slow (1) – Very fast (7)

A.2.2 User Experience - UEQ+ subset.

- Attractiveness
 - In my opinion, the product is generally:
 - * annoying (-3) – enjoyable (3)
 - * Bad (-3) – Good (3)
 - * unpleasant (-3) – pleasant (3)
 - * unfriendly (-3) – friendly (3)
 - The product characteristic described by these terms is for me
 - * Not important at all (1) - Very important (7)
- Dependability

- In my opinion, the reactions of the product to my input and command are:
 - * unpredictable (-3) – predictable (3)
 - * obstructive (-3) – supportive (3)
 - * not secure (-3) – secure (3)
 - * does not meet expectations (-3) – meets expectations (3)
- The product characteristic described by these terms is for me
 - * Not important at all (1) - Very important (7)
- Risk Handling
 - I find the application errors and risks which may arise when using the product to be:
 - * threatening (-3) – harmless (3)
 - * hazardous to health (-3) – not hazardous to health (3)
 - * damaging (-3) – not damaging (3)
 - * likely to cause collision (-3) – unlikely to cause collision (3)
 - The product characteristic described by these terms is for me
 - * Not important at all (1) - Very important (7)

A.2.3 *Task Load - NASA-RTLX subset.* Please indicate for each of the dimensions below how demanding the task was for you. Mark on the following scales to what extent you felt challenged or required in the six mentioned dimensions:

- Mental Demand
 - How much mental and perceptual activity was required (e.g., thinking, deciding, calculating, remembering, observing, searching, . . .)? Was the task easy or demanding, simple or complex, exacting or forgiving?
 - * Low (0) – High (100)
- Temporal Demand
 - How much time pressure did you feel due to the rate or pace at which the tasks occurred? Was the pace slow and leisurely or rapid and frantic?
 - * Low (0) – High (100)
- Frustration level
 - How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?
 - * Low (0) – High (100)

A.2.4 *User Trust - S-TIAS.* Please indicate the extent to which you agree with the following statements about the voice assistant:

- I have confidence in the assistant
 - Not at all (1) - Extremely (7)
- The system is reliable
 - Not at all (1) - Extremely (7)
- I can trust the system
 - Not at all (1) - Extremely (7)

A.3 Semi-structured Interview Questions

The following questions were asked while allowing follow-up prompts and clarification questions:

- (1) How much verbal feedback would you like from the system?
Consider driving situation, passengers, music, and other distractions.
- (2) Should the system notify you when it is uncertain, or decide autonomously? If notified, how should this be communicated?
- (3) Which system behaviors or experiences would foster long-term trust?