Vanishing point of attention: A platform for adaptive driver dialogue experiments

Morgan Fredriksson[1], Yanis Yaici[2], Kevin Lam[2], Jurgen Konigsmann[1], Jens Edlund[2]

[1]Nagoon AB, Stockholm, Sweden, [2]KTH Royal Institute of Technology, Stockholm, Sweden

Corresponding author: morgan@nagoon.se

Abstract

Current in-vehicle conversational agents lack awareness of the driving situation, treating all dialogue alike regardless of cognitive demands. This paper presents a modular experimental platform that integrates the CARLA driving simulator with a real-time spatial-reasoning engine to support research on situation-aware dialogue. The system enables Wizard-of-Oz studies in which human operators control conversational agents informed by live spatial-semantic analysis of the traffic environment. As initial validation, a controlled study (n = 10) tested the platform's sensitivity to conversational load effects, examining whether increasing conversational complexity produces a vanishing point of attention, a threshold where combined conversational and driving demands lead to a non-linear collapse in performance. Results revealed a sharp rise in collisions and missed hazard detections under high cognitive load, confirming the platform's sensitivity to conversational strain. The platform provides a reproducible testbed for investigating how dialogue timing, content, and environmental demands interact, offering a foundation for designing adaptive, cognitively safe in-vehicle conversational systems.

1 Introduction

Conversational agents (CAs) in vehicles promise enhanced safety and reduced workload, yet current implementations treat all dialogue equally regardless of the situation and the driving demands. This inability to adapt conversation based on situation and context is a fundamental limitation that creates a critical safety risk we term the vanishing point of attention: the threshold where conversational demands overwhelm driving resources, leading to catastrophic performance failure. This paper presents an experimental platform that integrates the CARLA driving simulator (Dosovitskiy et al., 2017) with a real-time spatial reasoning engine to support systematic investigation of situation-aware dialogue strategies. The platform addresses the gap between current static CAs and the dynamic nature of driving by providing real-time situational awareness capable of modulating conversational behaviour in response to evolving traffic and driver conditions. The system is grounded in a structured discourse model that represents dialogue as a sequence of issues under discussion, each linked to situational events such as approaching intersections, route changes, or external interruptions. This model provides the conceptual foundation for managing key mechanisms for cognitively sensitive driver interaction, such as dialogue timing, pausing, and resumption. The platform supports Wizard-of-Oz (WoZ) studies in which human operators control conversational agents informed by continuous spatial-semantic analysis of the simulated driving environment. This approach makes it possible to explore adaptive dialogue policies before full automation becomes technically feasible.

The key contributions are:

1. A reproducible technical framework integrating the CARLA high-fidelity driving simulation with real-time spatial reasoning, providing a testbed for the development and evaluation of discourse-level dialogue management and attention-aware dialogue strategies.

2. Experimental validation (n=10) demonstrating the platform's sensitivity to conversational load, revealing a pronounced, non-linear rise in collision frequency and severe degradation in hazard detection under high conversational complexity.

The platform provides researchers with a controlled yet realistic environment to address a critical question: how can in-vehicle conversational agents regulate their behaviour to prevent drivers from crossing the vanishing point of attention?

Proceedings of the 16th International Workshop on Spoken Dialogue Systems Technology, pages 231–238
February 26–March 1, 2026.

2 Background and related work

Understanding how CAs interact with the demands of real-world driving requires grounding in two complementary areas: the cognitive psychology of multitasking, and the technical evolution of in-vehicle dialogue systems. This section reviews the mechanisms by which conversation imposes cognitive load on drivers, the limitations of current system architectures, and the developments in spatial reasoning and experimental methodology that inform the present work.

2.1 Cognitive load in vehicle conversations

The integration of conversational interfaces in vehicles presents a central paradox. Speech interaction promises hands-free, eyes-free operation, apparently well suited to driving; yet empirical evidence shows that cognitive distraction from conversation can be as hazardous as visual-manual distraction. According to Multiple Resource Theory, conversational activity competes for the same central processing resources required for safe vehicle control, particularly when it demands reasoning, memory retrieval, or complex linguistic processing. Extensive experimental work has confirmed that such secondary tasks impair driving performance even when the driver's hands remain on the wheel. Strayer and colleagues (Strayer and Johnston, 2001)(Strayer et al., 2003) demonstrated that telephone conversations induce inattention blindness and delay hazard detection. Meta-analyses further show systematic degradation: as the cognitive complexity of speech tasks increases, subjective workload rises while lane-keeping and hazard-response performance deteriorate. (Engström et al., 2017) Several studies suggest that this degradation is not gradual but abrupt. (Patten et al., 2004) identified a cliff-edge pattern in which performance remains stable until cognitive demand crosses a threshold, after which collisions and reaction delays escalate sharply. This non-linear collapse aligns with what we describe as the vanishing point of attention: the moment at which the driver's available cognitive resources are exhausted. Theoretical models provide converging explanations for this phenomenon. Multiple Resource Theory (Wickens, 2008) and earlier work on resource economy (Navon and Gopher, 1979) describe how tasks compete across modalities and processing codes; when conversation and driving both engage central cognitive channels, interference becomes catastrophic. Threaded Cognition (Salvucci and Taatgen, 2008) offers a mechanistic view: concurrent tasks operate as threads sharing limited buffers, and driving performance collapses once the control thread loses priority to conversational processing. The challenge, therefore, is to design in-vehicle dialog systems that can recognise and respond to these cognitive dynamics in real time, adapting not only what they say but also when and how they say it.

2.2 Current in-vehicle conversational agents

Commercial in-vehicle conversational agents operate, presumably implicitly, under the assumption that drivers maintain consistent cognitive capacity regardless of traffic conditions or manoeuvring demands. They lack the adaptive behaviours that human passengers display naturally: pausing during complex driving, deferring non-urgent topics, or modulating conversational tempo when attention must be reclaimed. Current systems process user input and generate output in isolation from the surrounding situation, with little capacity to adjust timing or phrasing in response to risk(Du et al., 2024). This limitation stems from the absence of real-time world modelling and from the weak integration between dialogue management and situational-awareness subsystems. As a result, even modern assistants based on large language models remain reactive and largely blind to the physical and cognitive states of the driver. Research continues to address this gap, exploring mechanisms for linking dialogue flow to environmental understanding (Fernández-Rojas et al., 2019) and physiological indicators of driver distraction (Bargshady et al., 2025) Earlier generations of in-vehicle systems, including European research prototypes, demonstrated that multimodal speech interfaces could provide safe, hands-free interaction, but their architectures were tightly coupled to proprietary vehicle data and rigid simulation environments. This prevented systematic testing of critical phenomena such as interruption handling, timing control, and dialogue resumption. The resulting interactions were formally correct but behaviourally static, incapable of adapting to dynamic driving demands. To address these constraints, our work grounds dialogue management in an explicit discourse model that represents interaction as a structured sequence of issues under discussion, each linked to the evolving driving situation. This approach provides a foundation for dialogue behaviour that can be suspended, resumed, or redirected according to real-time situational demands, rather than following fixed turn-taking logic.

2.3 Spatial reasoning for situation awareness

Effective dialogue management in driving requires continuous awareness of the traffic environment, the vehicle state, and the driver's current activity. Spatial reasoning systems provide a means of transforming raw sensor or simulation data into high-level situational descriptions (e.g. an approaching intersection, a vehicle preparing to overtake, or entering high-density traffic). The spatial reasoning engine is a custom implementation built on established methods and public standards such as OpenDRIVE. It functions as a digital twin of the driving environment, and processes OpenDRIVE road-network data (Dupuis et al., 2010) and real-time telemetry to maintain a dynamic semantic representation of the surrounding scene. From these inputs it generates discrete situational events that can trigger dialogue-state transitions, inform timing decisions, or suspend speech when cognitive load is expected to peak. By linking the discourse model to this continuously updated world model, the system achieves a form of situation-aware dialogue management in which conversational behaviour is shaped by real-time driving conditions. Earlier studies recognised the importance of such adaptation. Human–human observations show that passengers naturally pause, defer topics, or change tone in demanding traffic situations (Fors and Villing, 2011). Early prototype systems began to emulate these behaviours by negotiating dialogue timing according to cognitive load (Edlund et al., 2012) or suspending conversation during critical manoeuvres (Neßelrath and Feld, 2013)(Reichel et al., 2014), typically using Wizard-of-Oz setups in simulators. However, these systems lacked a real-time spatial–semantic model capable of interpreting the driving situation at a level sufficient for adaptive dialogue management. Our framework advances this line of work by coupling high-fidelity driving simulation with continuous spatial reasoning, enabling the investigation of dialogue strategies that adjust dynamically to traffic events, driver state, and conversational progression.

2.4 Wizard-of-Oz for situated dialogue research

The Wizard-of-Oz (WoZ) paradigm offers a controlled means of investigating human–machine interaction before full automation becomes technically feasible. In driving research, it provides a safe and repeatable way to study conversational behaviour under realistic workload conditions while maintaining experimental control. Earlier automotive WoZ studies focused largely on command-and-control dialogues or information-retrieval tasks (Murao et al., 2003), with limited attention to the dynamic coupling between conversation and driving activity. Consequently, these experiments could characterise language use in vehicles but not the timing, pausing, or resumption strategies required for adaptive dialogue. Our approach extends this methodology by embedding the human operator within a real-time situation model that informs each dialogue decision. The Wizard observes live spatial-semantic data from the simulator (traffic flow, vehicle proximity, driver state) and uses this information to modulate dialogue in real time. This arrangement allows systematic exploration

of situation-aware dialogue policies: when to interrupt, when to resume, and how to adjust conversational pacing as environmental demands fluctuate. By combining WoZ control with spatial reasoning and discourse-state modelling, the platform bridges experimental flexibility and ecological realism. It supports iterative development of adaptive dialogue strategies that can later be transferred to autonomous implementations, once perception and reasoning components are sufficiently mature.

Figure 1: Simulator environment

3 System architecture

The experimental platform implements a distributed, event-driven architecture that bridges high-fidelity driving simulation with real-time spatial reasoning for situation-aware dialogue management. The design prioritises modularity, real-time performance, and experimental reproducibility. It provides a research infrastructure in which dialogue flow, situational state, and operator control are coordinated through a unified event framework.

3.1 Core components and integration

The architecture comprises three primary subsystems interconnected through a central event server. The CARLA simulator module provides the simulated driving environment, including road layouts, vehicle behaviour, and surrounding traffic. The simulator affords fine-grained control over driving scenarios through sensor simulation, vehicle telemetry, and scripted traffic patterns. A Tesla Model 3 vehicle model was selected for the initial experiments. The simulation runs in synchronous mode to ensure consistent timing across trials. The spatial reasoning engine functions as a real-time world model, an enriched digital twin that reconstructs the simulated physical environment and tracks the dynamic behaviour of all relevant actors and objects. By interpreting raw data streams from the simulator and external sensors, the engine maintains a continuously updated representation of the traffic scene, driver state, and vehicle context. It generates high-level situational insights such as relative vehicle positioning, upcoming manoeuvres, or signs of driver distraction. The event server coordinates real-time communication between all system components. It receives event streams from CARLA (vehicle position, speed), the spatial-reasoning engine (contextual state analysis), and user-monitoring sensors (gaze, driver state). The spatial-reasoning engine processes these inputs and publishes higher-level situational events through the server (e.g. approaching intersections, vehicles preparing to overtake, or prolonged gaze distraction) forming the shared basis for dialogue control.

3.2 Spatial-semantic event processing

The spatial-reasoning engine parses OpenDRIVE road-network definitions through the ODRParser component, ensuring precise alignment between authored scenarios and the simulated environment. Events can be created either through a graphical editor for prototype development or programmatically, with specific road features (for example, turns, merges, or stop signs) linked to their corresponding segments for reuse across experimental conditions. A CarlaReplayParser enables scenario replay and controlled manipulation, allowing researchers to modify parameters between experimental runs while maintaining consistent underlying traffic patterns. This capability ensures reproducible experimentation when evaluating dialogue-management strategies under matched driving conditions.

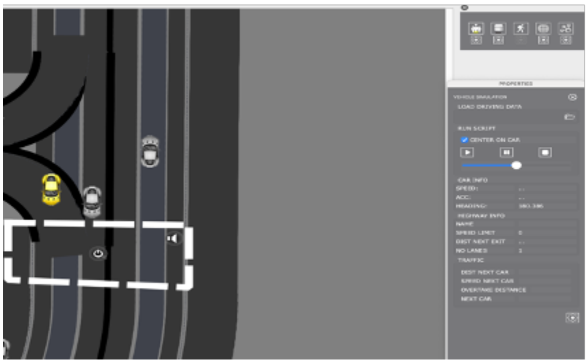

Figure 2: Screenshot from Spatial Reasoning Engine GUI running in sync with simulator.

3.3 Dialogue state management

The system maintains a finite-state representation in which each state corresponds to a specific conversational domain (navigational assistance, overtaking guidance, or idle monitoring). Each state defines entry conditions based on situational events, available utterances from pre-scripted libraries, and transition rules triggered by events or operator input. The WoZ interface enables human operators to control the conversational agent during experiments, offering graphical selection of predefined utterances, simulator-event triggering, free-text composition for real-time TTS rendering, and visualisation of dialogue-flow states. A dedicated authoring tool supports iterative scenario development through dialogue-state definition and event-to-transition mapping.

3.4 Sensor integration and data collection

Beyond CARLA's internal sensors, the setup incorporates webcams for facial-expression and pose detection, eye-trackers for detailed gaze analysis, and optional audio recording. Hardware includes a Logitech G923 force-feedback steering wheel and pedal set with a Samsung Odyssey Neo G9 ultra-wide display (49 in, 5120×1440 resolution). All experimental data (CARLA simulation recordings, user-monitoring streams, and WoZ interaction logs) are centrally recorded to enable replay, annotation, and cross-modal analysis.

4 Method

As an initial validation of the platform, we conducted a controlled study (n = 10) examining the relationship between conversational complexity and driving performance during simulated long-haul driving.

4.1 Participants

Ten participants aged 20–25 years, all students at a technical university and holding valid European driving licences, took part in a within-subjects repeated-measures experiment.

4.2 Procedure

Four levels of conversational complexity were tested in counterbalanced order:
Baseline: Silent driving
Low: Simple factual questions (name, address, preferences)
Medium: Memory recall tasks (recalling speed limits, describing passed objects)
High: Mental arithmetic and logical reasoning

Participants drove a mixed urban–highway route for approximately ten minutes per condition, with five-minute breaks between runs. Performance was assessed using

- the NASA-TLX (Hart and Staveland, 1988) for subjective workload;
- Detection Response Task (DRT) (International Organization for Standardization, 2016) for objective attention capacity, and
- driving metrics including lane invasions, speed violations, and collisions.

4.3 Ethical Considerations

The study protocol was conducted in accordance with the ethical research guidelines of the authors' institution. Participant welfare was the primary concern; although the simulation eliminated physical risk, measures were taken to minimize potential psychological stress from the demanding cognitive tasks and simulated collision events. All participants provided written informed consent, which detailed the nature of the driving simulation and the dialogue tasks. Participants were explicitly informed that they could stop the experiment at any time without penalty. All collected data was fully anonymized to protect participant privacy, and all participants were debriefed on the study's objectives following the experiment.

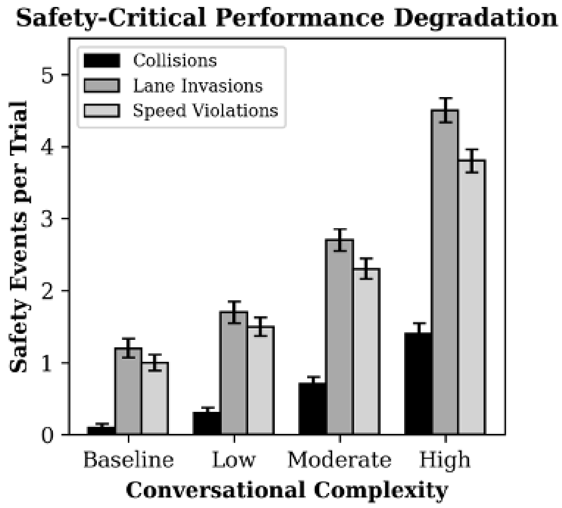

Figure 3: Safety-critical performance degradation across all metrics. Collisions (black bars) Lane Invasions (Dark Gray), and Speed Violations (Light gray) all increased with conversational complexity.

5 Results

The platform successfully captured systematic relationships between conversational complexity and driving performance, consistent with established findings on cognitive load in driving research.

5.1 Driving performance

Performance degradation was systemic, culminating in a non-linear decline in safety-critical metrics under high conversational complexity. Collision frequency increased from 0.0 events per trial in the baseline condition to 2.3 events per trial. This was accompanied by a pronounced loss of vehicle control both for lane invasions and speed violations.

5.2 Cognitive load and attention

These performance failures corresponded directly with increases in cognitive load. Perceived mental demand (NASA-TLX) rose by 136 per cent from baseline to the high-complexity condition. Objective measures showed a corresponding depletion of attentional resources: detection-response (DRT) reaction times slowed by 60 per cent, and detection hit rate fell from 93.1 per cent to 59.8 per

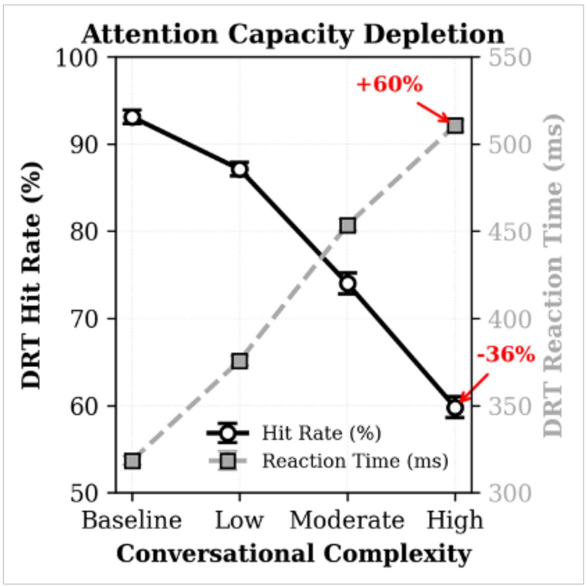

Figure 4: Objective attention capacity depletion under increasing conversational complexity. As complexity rises, DRT Hit Rate (solid line) falls from 93.1% to 59.8% (while DRT Reaction Time (dashed line) slows from 318.4 ms to 510.5 ms

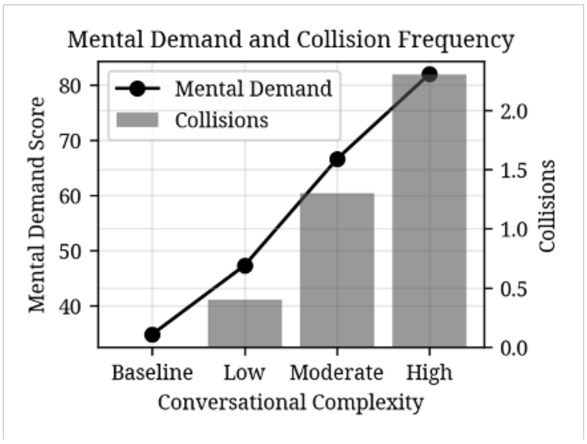

Figure 5: The relationship between conversational complexity, perceived mental demand (NASA-TLX), and collision frequency. As mental demand (line) increased from baseline () to high (), collision frequency (bars) rose from 0.0 to 2.3 events per trial.

cent. Participants thus failed to register over 40 per cent of peripheral stimuli; clearly consistent with inattention blindness.

5.3 Workload vs performance

Strong correlations emerged between subjective workload and objective performance: NASA-TLX and DRT reaction time ($r = 0.92$, $p < 0.001$) and NASA-TLX and collision frequency ($r = 0.85$, $p < 0.001$).

6 Discussion

The results confirm the experimental platform's sensitivity to cognitive-load effects across multiple measurement dimensions. The platform advances situated dialogue research by treating the conversational agent as a situation-aware co-driver rather than an isolated interface. Unlike conventional voice assistants, which operate without regard to driving conditions, the integrated system enables real-time perception of and response to the surrounding environment. This marks a shift from static interaction to dynamic adaptation based on continuous spatial–semantic understanding. The architecture facilitates reproducible experimentation that is essential for systematic progress. CARLA provides repeatable scenarios in which identical traffic sequences can be used to test different dialogue strategies, while the spatial-reasoning engine translates sensor data into structured situational events. These events drive dialogue modulation across trials, ensuring that conversational behaviour aligns with driving conditions. The modular, event-driven design with clearly defined interfaces lowers barriers for other researchers to reproduce and extend this work. Our preliminary study illustrates the vanishing point phenomenon: a threshold beyond which conversational demands appear to trigger a non-linear collapse of driving performance. This finding, while tentative, has important implications for system design. It suggests that hands-free operation alone is insufficient for safety. Future in-vehicle agents must act as active co-pilots, managing cognitive load by modulating not only the content of dialogue (what is said) but also its timing (when it is said) and manner (how it's said) in order to prevent the driver from exceeding attentional capacity. This non-linear collapse aligns with models such as Threaded Cognition. The sharp rise in collisions, mirrored by the 40 per cent DRT detection failure, indicates that once combined task demands exceed cognitive capacity, the primary driving thread becomes starved, producing inattention blindness and severe degradation of control. The platform's scientific value lies in its ability to capture this causal chain in full. Simpler setups might have recorded only subjective workload or isolated performance metrics. By integrating high-fidelity simulation (CARLA), real-time situational analysis, and objective attention measures (DRT), the system exposes the link from conversational load to attentional depletion and ultimately to safety-critical failure. The exceptionally strong correlations observed—such as that between workload and DRT reaction time—confirm this coupling. Collectively, these results demonstrate the platform's utility as a high-sensitivity testbed for developing and evaluating adaptive dialogue strategies that may help prevent drivers from reaching this vanishing point of attention.

6.1 Limitations

The modest sample (n = 10, aged 20–25, technical university students) limits generalisability. Accordingly, we interpret the results primarily as platform sensitivity validation and evidence of a threshold-like effect in this setup, not as population-level estimates. Broader validation is needed with diverse and professional driver populations. Although high-fidelity simulators are well-established tools for assessing driving performance (de Winter et al., 2009; Meule & Fraser, 2015), they cannot fully reproduce real-world risk perception; collision rates in CARLA may not directly translate to on-road conditions. The Wizard-of-Oz approach introduces operator variability and limits deployment scalability. Gaze metrics, although captured, are outside the scope of this initial report. Similarly, the linguistic content of the dialogues was not analysed; future work should examine whether performance effects stem from cognitive load, task-switching overhead, or specific linguistic factors. The validation study manipulated fixed complexity levels and did not evaluate an adaptive timing policy. Policy evaluation is left for future work. Hardware factors, including relatively stiff steering feedback, may have amplified lane-keeping deviations. Nevertheless, the ability to detect the pronounced rise in collision frequency under load demonstrates the platform's research value.

7 Future work

Future work will extend the system from Wizard-of-Oz operation towards autonomous, policy-driven control. Preliminary trials suggest that Wizard-initiated pauses during high-demand driving segments may mitigate performance degradation, though systematic evaluation of adaptive policies is ongoing. Reinforcement-learning approaches could enable an agent to interpret real-time situational cues and decide when and how to engage in dialogue. Driver-performance metrics could serve as reward signals for learning optimal attention-aware interaction policies. Subsequent research will include larger-scale studies across varied scenarios and the transfer of methods to professional simulators and, ultimately, controlled on-road testbeds. To support reproducibility, the platform builds on open-source components including CARLA and OpenDRIVE; source code and experimental materials may be shared with researchers on request.

8 Conclusion

This work presented a modular experimental platform that integrates high-fidelity driving simulation with real-time spatial reasoning to support research on situation-aware dialogue in vehicles. Rather than treating conversation as independent of driving, the system links dialogue flow to a continuously updated model of the traffic environment and driver state, enabling controlled study of adaptive timing, pausing, and resumption strategies. The pilot study demonstrated the platform's sensitivity to conversational load and revealed evidence of a vanishing point of attention—a threshold where conversational demands may precipitate a rapid collapse in driving performance. While preliminary, these results illustrate the value of the approach and its potential to inform the design of cognitively safe interaction policies. More broadly, the platform provides a reproducible infrastructure for advancing research on attention-aware dialogue systems. It enables systematic experimentation under realistic conditions and offers a bridge between conceptual models of discourse management and future autonomous implementations. As vehicle assistants evolve, such frameworks will be essential for ensuring that conversational systems help preserve, rather than erode, the attention required for safe driving.

References

Ghazal Bargshady, Hakki Gokalp Ustun, Yasaman Baradaran, Houshyar Asadi, Ravinesh C Deo, Jeroen Van Boxtel, and Raul Fernandez Rojas. 2025. Neurosafedrive: An intelligent system using fnirs for driver distraction recognition. *Sensors*, 25(10).

Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. Carla: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning (CoRL 2017), Proceedings of Machine Learning Research, Vol 78*, pages 1–16, Mountain View, CA, USA.

H. Du, X. Feng, J. Ma, M. Wang, S. Tao, Y. Zhong, and H. Wang. 2024. Towards proactive interactions for in-vehicle conversational assistants utilising large language models. *arXiv preprint arXiv:2403.09135*.

Marius Dupuis, Martin Strobl, and Hans Grezlikowski. 2010. Opendrive 2010 and beyond – status and future of the de facto standard for the description of road networks. In *Proceedings of the Driving Simulation Conference Europe (DSC Europe 2010)*, pages 231–242, Paris, France.

Johan Edlund, Clare Oertel, and Jonas Gustafson. 2012. Investigating negotiation for load-time in the gethomesafe project. In *IAST Workshop on Innovation and Applications in Speech Technology*, pages 45–48, Dublin, Ireland.

Johan Engström, Gustav Markkula, Trent Victor, and Narelle Merat. 2017. Effects of cognitive load on driving performance: The cognitive control hypothesis. *Human Factors*, 59(5):734–764.

R. Fernández-Rojas, A. Perry, H. Singh, B. Campbell, S. Elsayed, R. Hunjak, and H. A. Abbass. 2019. Contextual awareness in human-advanced-vehicle systems: a survey. *IEEE Access*, 7:33304–33328.

Kristina Lundholm Fors and Jessica Villing. 2011. Reducing cognitive load in in-vehicle dialogue system interaction. In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue, SemDial*, pages 55–62.

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, Advances in Psychology Vol 52*, pages 139–183. Elsevier, Amsterdam, NL / New York, USA.

International Organization for Standardization. 2016. Road vehicles — transport information and control systems — detection-response task (drt) for assessing attentional effects of cognitive load in driving (iso 17488:2016).

H. Murao, N. Kawaguchi, S. Matsubara, T. Yamaguchi, and Y. Inagaki. 2003. Example-based spoken dialogue system using woz system log. In *Proceedings of the Fourth SIGdial Workshop on Discourse and Dialogue*, pages 140–148.

David Navon and Daniel Gopher. 1979. On the economy of the human-processing system. *Psychological Review*, 86(3):214–255.

Robert Neßelrath and Michael Feld. 2013. Towards a cognitive load ready multimodal dialogue system for in-vehicle human-machine interaction. In *Adjunct Proceedings of the 5th International Conference on Automotive User Interfaces and Interactive Vehicular Applications. ACM, New York, NY, USA*, pages 49–52.

Christopher J.D. Patten, Albert Kircher, Joakim Östlund, and Lennart Nilsson. 2004. Using mobile telephones: Cognitive workload and attention resource allocation. *Accident Analysis & Prevention*, 36(3):341–350.

Sven Reichel, Jasmin Sohn, Ute Ehrlich, André Berton, and Michael Weber. 2014. Out-of-domain spoken dialogs in the car: a woz study. In *Proceedings of the 15th annual meeting of the special interest group on discourse and dialogue (SIGDIAL)*, pages 12–21.

Dario D. Salvucci and Niels A. Taatgen. 2008. Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review*, 115(1):101–130.

David L. Strayer, Frank A. Drews, and Dennis J. Crouch. 2003. Cell-phone-induced failures of visual attention during simulated driving. *Journal of Experimental Psychology: Applied*, 9(1):23–32.

David L. Strayer and Frank A. Johnston. 2001. Driven to distraction: Dual-task studies of simulated driving and conversing on a cellular telephone. *Psychological Science*, 12(6):462–466.

Christopher D. Wickens. 2008. Multiple resources and mental workload. *Human Factors*, 50(3):449–455.