User Driven Strategies

Impact of Large-Language Models on In-Vehicle UX

Chris Schreiner

June 2025

Table of Contents

Introduction

It has been over 40 years since the Nissan Silvia S110 was launched in Japan, which was the first commercially available vehicle to detect vocal sounds.[1] It has been almost 30 years since the launch of Mercedes-Benz Linguatronic in Europe.[2] And it has been 25 years since OnStar launched offboard speech recognition with their Virtual Advisor in North America. Now well into the 2020s, in-car voice assistants have slowly but steadily evolved from rule-based voice control to more fluid, dialog-driven interaction.

Today, with the rise of large language models (LLMs), the in-car assistant is evolving once again—from a functional command parser to a conversational partner capable of understanding intent, context, and nuance. Unlike legacy voice systems, which rely on rule-based parsing and limited vocabularies, LLMs offer fluid, context-aware, open-ended dialogue capabilities. With the ability to infer intent, hold memory across interactions, and respond intelligently to ambiguous input, this evolution could redefine how vehicle occupants communicate with their vehicles.

This report explores the potential of LLMs to meaningfully improve the in-vehicle experience. It investigates where these models offer incremental vs. transformative gains over current systems, identifies key future use cases, and unpacks the nuanced UX considerations that may not be immediately obvious to designers and product teams.

[1] While referred to by Nissan as "World-first: Voice recognition power-assisted windows", the system reportedly recognized only the sound of occupants who may be choking from getting their necks trapped by the power windows, which were themselves a new automotive feature at the time.

[2] Linguatronic was capable of basic digit phone dialling with a vocabulary of about 30 words. https://aclanthology.org/H01-1047.pdf

LLM Advantages

LLMs have the potential to offer a step-change in capability over today's natural language systems.

1. Conversational Understanding, Not Just Command Matching

Traditional voice UX systems are good at parsing structured input like "Call John" or "Navigate to 123 Main Street." But they break down when utterances have more complex, unstructured, or nuanced information, including the following:

- Ambiguous referents: "Take me to Union" where "Union" could refer to a street, station, or other place name
- Vague intent: "Play something relaxing" does not specify what genre or service
- Implicit goals: "I'm cold" could mean to turn up the temperature, turn on the defroster, turn off the air conditioner, or just complaining

LLMs are an improvement in these scenarios because they can interpret implied meaning and generate structured outputs from unstructured inputs. This unlocks a much broader range of interaction styles—especially for non-technical users who may not remember system limitations.

2. Flexibility in Language and Form

LLMs have fluency in:

- Colloquialisms and idioms (e.g., "Crank up the AC," "Kill the noise")
- Disfluencies (e.g., "Uh, I meant the Thai place, not sushi")
- Multilingual speech or code switching (e.g., "Quiero ir a Target, por favor")

Lack of flexibility has always been a key factor in task errors with speech systems. LLMs, by contrast, generally handle these well, potentially dramatically increasing task success rates, especially in international or multicultural markets.

3. Context Awareness and Short-Term Memory

Legacy systems typically process each voice command independently. If they have any short-term memory, they tend to be extremely limited to the last few commands. LLMs can reference entire previous interactions, making conversations feel much more fluid. Contextual memory enables the assistant to feel less like a "command interpreter" and more like a true assistant.

4. Task Composition and Chaining

LLMs can interpret multi-intent queries that require sequencing or coordination. This is particularly beneficial in the driving context, especially for advanced satnav entries. For example, "Find an open EV charger, and if it's more than 15 minutes away, show me a place I can grab lunch nearby."

These kinds of requests require logical reasoning, filtering, and data synthesis across multiple domains—beyond the scope of conventional voice systems. An well-integrated LLM could potentially handle this situation.

5. Domain Adaptability and Personalization

LLMs can be provided with domain-specific knowledge (e.g., vehicle manuals, OEM service logs, user preferences). This allows for:

- Onboarding support: "How do I turn off auto hold?"
- Diagnostics explanation: "What does the triangle light mean?"
- Personalization: "Use my usual climate settings" or "Tell me if tire pressure is low like last winter."

These interactions move voice from a functional tool to a personalized interface for the vehicle.

6. Natural Failover and Clarification

Where traditional systems often produce "Sorry, I didn't catch that" messages, LLMs can fail more gracefully by suggesting clarifications, which keeps the user in the flow, engaged, and alert, and reduces the potential for frustration and distraction.

Future LLM Use Cases

LLMs can enhance the UX for almost any complex voice query, from media to ecommerce. Future use cases for which we predict LLMs will have the greatest impact on the in-vehicle experience include:

1. Goal-Based Navigation

This is the area where LLMs can have the greatest immediate impact. More efficient navigation features are always at the top of consumers' wish lists, and here is where an LLM's context and reasoning engines shine. If LLMs are integrated with a system's navigation engine to understand previous routes, current traffic, EV chargers, and POI databases, consumers will finally be able to naturally describe their navigation goal, and the LLM can calculate the optimal path to achieve it.

For example, an LLM would be able to map out long-range EV trips, with charging stations and desired activities and boundaries along the route (e.g. "On this trip I want charging stations with a Starbucks no more than one mile off the route). An LLM could also provide suggested destinations along the route (e.g. "I have three hours to kill. Can you suggest a scenic drive, a museum, and a place that has excellent sushi?"), or provide routes beyond the traditional options that delight the user in some specific but meaningful way (e.g. "Take me on a quiet drive through forested areas.").

2. Dynamic Troubleshooting

While OEMs have increasingly implemented owners' manuals in the dash, and have long talked about providing the ability to verbally ask questions such as "What does this warning light mean?", LLMs can take this use case one step further. Drivers could theoretically ask questions like "What is that squeaking sound I'm hearing?" which could be diagnosed by the LLM with a captured audio file without any vehicle diagnostic information. This would be especially valuable in situations where vehicle issues are intermittent.

3. Deep Personalization

LLMs integrated with vehicle settings and infotainment can create a highly tailored user-specific experience. Some examples would include:

- "The adaptive cruise control is a little too aggressive, please back it off a bit."
- "When I'm on a rural road, turn off lane assistance."
- "From now on when I ask for directions avoid this road whenever possible."

4. Multi-Modal Scene Reasoning

LLMs tied into sensors or camera inputs could be able to interpret scenes and answer questions about recent events (e.g. "Did I have the right of way there?", "What kind of animal was that?").

5. Other Use Cases

There are many other automotive use cases possible with LLMs, but many of these are at a minimum duplications of use cases that can be achieved via other means, or at best provide only minimal enhancements. These include entertaining (e.g. storytelling, games), driver monitoring (e.g. detecting and responding to mood or fatigue), reminders (e.g. “Remind me to look for a present for my partner when I get home”), proactive ecommerce (e.g. “You are low on fuel. There is a Shell station 3 miles ahead that has the cheapest gas in the area”), and many more. These use cases could be beneficial add-ons, but should not be the motivation for integrating LLMs.

Key UX Considerations

When integrating LLMs into an existing speech solution, there are several key factors to consider.

1. Managing Expectations and Errors

The propensity for LLMs to hallucinate or respond with overconfidence is well documented. In-vehicle systems must include clear boundaries ("I'm not sure, let me check...") and fail-safe fallbacks to prevent misleading or distracting output. Especially when dealing with vehicle safety systems or driving-related tasks, inaccuracies and hallucinations cannot be tolerated.

2. Spoken Tone Misalignment

Written language and spoken speech can differ in how they are delivered. Unlike written commands, spoken input carries tone, mood, hesitation, urgency, sarcasm, and frustration—all of which influence meaning but are often invisible to an LLM trained primarily on text. Spoken utterances have greater spontaneity and are more likely to be unfiltered. In cognitively demanding situations in the car, users may speak more emotionally and under stress. They may make sarcastic comments such as "Nice job" when the assistant fails to reroute, or "I guess we're lost again" not as a literal query, but as a pointed comment. LLMs trained on internet-scale text may interpret these phrases too literally or respond inappropriately if tone and context are ignored.

3. Misuse of LLMs for Unsafe Commands

Users may jokingly or casually say things like "Take me off a cliff" or "Run them off the road" and LLMs might respond literally or lightly. A fully integrated LLM needs to account for these, and have a strong sense of safe driving behaviour.

4. Regional and Cultural Nuance

LLMs should adapt to local expressions, etiquette, and linguistic norms. What feels intuitive in Detroit may fall flat in Los Angeles, and even more so in Munich or Seoul. Fine-tuning by region can enhance user trust and comfort.

5. Multi-User Input

Multiple occupants speaking at the same time has been a well-understood issue with voice controls, and especially so with traditional in-vehicle speech systems. The more casual nature of interacting with LLMs due to their flexibility makes it more important to add speaker of seat identification to differentiate between multiple users who may speak at the same time.

6. Reliance on Cloud Connectivity

Many LLM-based assistants depend on constant internet connectivity, but vehicles frequently pass through areas with weak or no signal (e.g. parking garages, tunnels, rural routes). This requires delicate error handling or a hybrid solution with compressed, offline-capable reasoning engines that can degrade gracefully.

7. Data Privacy

LLMs raise new privacy challenges at a time when automotive OEMs are still struggling with the responsibilities of connected data. OEMs must define policies around voice storage, query logs, model fine-tuning using customer data, and what data is shared with the LLM partner.

8. LLM Updates

As LLMs evolve, so will risks. OEMs must prepare for regular model updates, prompt tuning, and revalidation of assistant behavior.

Conclusions

In-vehicle speech interfaces have come a long way—from the rigid, menu-driven systems of the early 2000s to more fluid, natural-language-based assistants in today's mirroring solutions and embedded in premium vehicles. However, despite the improvements, most current systems remain narrow in scope, easily confused by ambiguity, and frustratingly literal. LLMs offer a path beyond these constraints. Their ability to interpret nuanced commands, manage multi-step tasks, and infer unstated goals introduces a step-change in how drivers and passengers can interact with their vehicles, creating an experience more akin to talking to a helpful human than a dogmatic machine.

That said, LLM integration into the car is not without risks. Automakers face challenges that extend beyond technology: user trust, expectation management, privacy, and safety all become more complex in a conversational system that remembers, reasons, and sometimes improvises.

As OEMs design vehicles years in advance, they need to understand how LLMs may evolve between the time vehicle specs are completed and the end result is driven off the showroom floor. In that timeframe, LLMs are likely to become faster, smaller, and more efficient, enabling real-time, on-device reasoning with low latency and no cloud dependency. Future LLMs may be able to incorporate multi-modal inputs by default (vision, audio, location, and sensor data), allowing them to interpret not just what a user says, but also what's happening inside and outside the vehicle. Importantly, LLMs may move from being reactive assistants to proactive, collaborative agents that understand intent, learn routines, anticipate needs, and coordinate across services, all with a tone, style, and vocabulary tailored to each user. In the vehicle context, this means LLMs may replace today's assistants and manage complex tasks across HVAC, navigation, entertainment, safety, and communication — not just with smarter responses, but with true conversational intelligence that can reason through ambiguity, handle interruptions, and adapt to the shifting cognitive load of driving.

The winners in this space won't be those who simply bolt on a smarter voice model—but those who deeply integrate LLMs with automotive UX, edge AI, and backend integration to deliver in-car experiences that are beyond what can be done today.

Contact

Chris Schreiner
+44 1313 588626
www.userdrivenstrategies.com
chris@userdrivenstrategies.com