NEWS RELEASE

# The Car Is the Context: Why Automotive AI Will Be Won Beyond the Model

2026-11-18

Inspired by a presentation by Thomas Dannemann, Vice President of Cockpit Electronics Business Development at Visteon, at the 3rd European Automotive Intelligent Cockpit and Smart Driving Summit 2026, organized by ECV International.

Automotive AI will not be differentiated by the model alone. The advantage will come from controlling the vehicle context and the architecture that makes it useful.

The automotive AI race is increasingly framed around models: which is most capable, fastest or best suited to run in the vehicle. But models alone will not define the quality of automotive AI. Context will.

Ask an AI assistant, “Will we make the airport?” and a generic model can estimate travel time. An intelligent vehicle can also consider location and traffic, remaining range, charging requirements, cabin occupants and the destination itself.

The value comes from combining AI with information unique to the vehicle and the journey. This changes the question for automakers from which model to deploy to how vehicle context is collected, controlled and made available to AI.

Automotive AI needs a different architecture

Vehicles impose constraints that consumer AI applications do not. Connectivity is inconsistent, some decisions are latency-sensitive, vehicle and cabin data can be sensitive, and functions may need to remain available offline. As

assistants develop persistent memory, that memory also becomes part of the relationship between the customer, the vehicle and the brand.

This makes a purely cloud-centric architecture difficult to justify for every workload. Complex reasoning and rapidly changing world knowledge remain strong candidates for cloud inference. Tasks that depend on immediate vehicle state, predictable latency, privacy or continuous sensor input often favor local processing.

The practical answer is not edge or cloud, but an architecture capable of using both. A request can be evaluated against latency, cost, connectivity, privacy and consent requirements before being routed to the appropriate compute environment.

The economics reinforce the point. Continuously transmitting four compressed camera streams at 1–5 Mb/s for 30 driving hours per month could generate approximately 54–270 GB of data transport each month, before inference cost is considered. Local AI can instead interpret sensor data in the vehicle and determine which information is valuable enough to send upstream.

Edge AI is therefore not simply cloud AI running closer to the customer. It changes what needs to leave the vehicle in the first place.

AI decisions begin before hardware freeze

AI needs to be considered at the system level, whether compute is integrated into the vehicle platform or added through a dedicated AI accelerator. Use cases establish requirements for latency, NPU performance, memory, power and thermal management, while concurrent workloads, fallback strategies and future model requirements determine how that compute should be deployed.

As voice, vision, cabin perception, vehicle data, retrieval and agentic reasoning converge, the challenge is no longer adding individual AI features but designing an architecture that can support and evolve them together.

Models will change. Context is more durable.

The model landscape will change far faster than the lifecycle of a vehicle. An automotive AI architecture therefore needs a stable orchestration layer between the vehicle experience and individual models.

Memory, vehicle context, retrieval, permissions, tools, evaluation and workload routing should persist even when the underlying model changes. This is the principle behind **Visteon's CognitoAI™** approach: an OEM-controlled, model-agnostic AI layer designed to work across different models, speech technologies, operating environments

and compute platforms.

The benefit is not just technical flexibility. It allows automakers to adopt better AI technologies as they emerge without surrendering control of the customer experience around them.

The strategic asset is the context

A vehicle knows its destination and physical state. It can access information about the journey, cabin and surrounding environment. Over time, it can connect those signals with services, preferences and driver intent. That context is difficult for a generic AI provider to replicate.

Foundation models will improve rapidly, costs will change and today's leaders will be challenged by new models. The vehicle's context, however, belongs to the automotive experience.

For OEMs, the durable competitive advantage may therefore be less about owning the model and more about controlling the context, orchestration and architecture around it.