Whitepaper

In-cabin sensing systems: from distributed sensors to context-aware, AI-defined architectures

Christoph Zorn and Yannick Saalberg, NXP Semiconductors

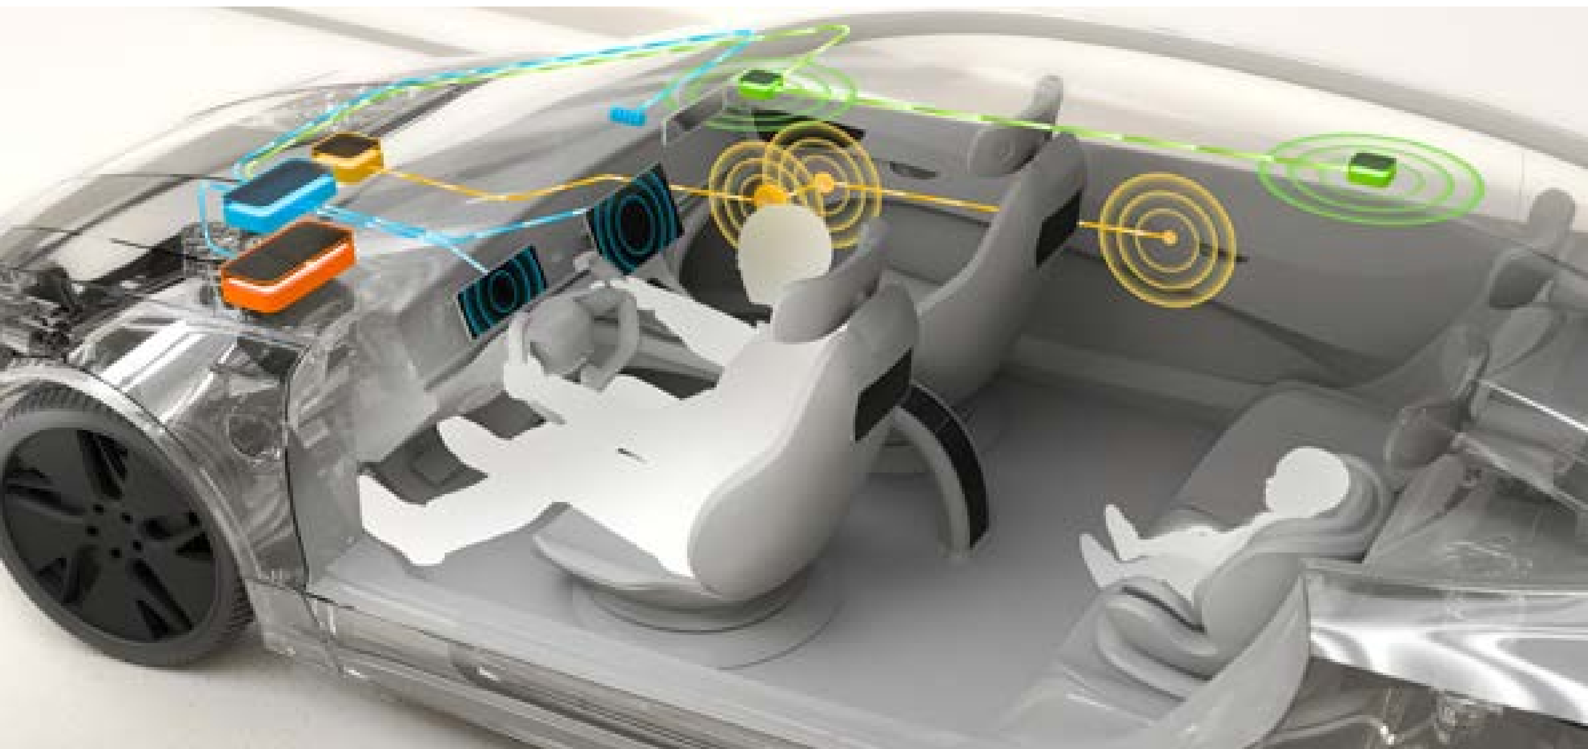

Introduction

The automotive industry is entering its next major transformation.

Over the past decade, vehicle intelligence has focused primarily on understanding the environment outside the vehicle. Advances in radar, cameras, high-performance processing, and artificial intelligence have dramatically improved a vehicle's ability to perceive its surroundings, identify hazards, and support increasingly sophisticated driver assistance systems. These technologies have laid the foundation for automated driving and represent one of the most significant innovation cycles in automotive history.

The next challenge is bringing this same level of intelligence inside the vehicle.

Future vehicles must understand not only what is happening around them, but also what is happening within the cabin. A vehicle may detect a potential hazard ahead, but it must also know whether the driver is attentive and ready to take control, whether occupants are positioned safely, whether a child or pet has been left behind, or whether a medical emergency is unfolding. The ability to perceive both the external environment and the internal state of the vehicle is becoming fundamental to the future of automated mobility.

This evolution parallels the broader development of artificial intelligence. The first generation of automotive AI focused primarily on perception, enabling vehicles to sense and interpret the world around them. More recently, advances in generative AI have introduced systems capable of reasoning and supporting increasingly sophisticated decision-making processes. The next stage is emerging now through agentic AI, where systems move beyond sensing and thinking toward acting autonomously. For vehicles to safely take action, they must possess a complete understanding of both their surroundings and their occupants.

In-cabin sensing therefore represents far more than a collection of isolated safety functions. It is becoming a critical pillar of software-defined vehicles and a key enabler of future autonomous transportation. This shift is being accelerated by growing customer expectations and an expanding regulatory landscape. Initiatives such as Euro NCAP's safe driving program, Driver Monitoring System (DMS) requirements, and Child Presence Detection (CPD) regulations are raising the bar for what vehicles must understand about their occupants.

Occupant monitoring	Driver engagement	Vehicle assistance
Seatbelt usage	Driver monitoring	Speed assistance
Occupant classification	Driver controls	ACC performance
Occupant presence – incl. CPD		Steering assistance

Euro NCAP minimum total points		
2026	60/100 pts	Soft landing ★★★★★
2027	70/100 pts	
2028+	80/100 pts	

Fig 1. Examples of occupant monitoring, driver engagement, and vehicle assistance functions increasingly reflected in safety assessment programs such as Euro NCAP.

These requirements illustrate how in-cabin sensing is evolving from a single-purpose safety function into a broad set of interconnected capabilities spanning occupant monitoring, driver engagement, and vehicle assistance. As assessment programs continue to expand, OEMs must support an increasing number of occupant-centric functions while maintaining acceptable cost, power consumption, and system complexity.

At the same time, OEMs face a significant challenge. Every new feature appears to demand additional sensors, processing power, software, validation effort, and system integration. The question is no longer how to add one more sensing function. The real challenge is how to build increasingly intelligent vehicles without continuously adding complexity and cost.

This whitepaper explores how context-aware, software-defined architectures can address that challenge. By combining existing sensing assets, centralized processing platforms, and AI-driven sensor fusion, OEMs can create scalable in-cabin sensing systems that continuously evolve through software while maintaining control of vehicle cost, power consumption, and complexity.

In-cabin sensing is becoming a system-level challenge

Historically, many in-cabin functions were developed independently to address a specific requirement. Driver engagement monitoring, seat occupancy detection, occupant classification, airbag control, and child presence detection often emerged as separate systems, each with its own sensing hardware and processing resources.

While effective for individual applications, this approach becomes increasingly difficult to sustain as vehicles move toward higher levels of automation. The number of functions that depend on an understanding of occupants continues to grow. Vehicles are expected to identify distracted drivers, determine occupant positions, recognize medical emergencies, detect forgotten children, manage seatbelt compliance, support post-crash assessment, and provide personalized comfort and wellness experiences.

Meeting these requirements through dedicated hardware for every function quickly leads to escalating cost, wiring complexity, validation effort, and power consumption. More importantly, it creates a fragmented architecture in which multiple systems attempt to understand the same environment without sharing knowledge.

As vehicle intelligence expands, this model becomes increasingly inefficient. The future requires a more holistic approach in which sensing systems contribute to a unified understanding of the cabin rather than operating as isolated subsystems.

The shift toward context-aware intelligence

The future of in-cabin sensing will not be defined by individual sensors. Instead, it will be defined by how information from multiple sensors is combined, interpreted, and transformed into actionable intelligence.

This shift is closely aligned with the broader evolution of software-defined vehicles. Modern vehicle architectures increasingly rely on centralized compute platforms capable of collecting data from many different sources and applying advanced AI algorithms to generate insights that no individual sensor could provide on its own.

Within such architectures, sensors become data sources rather than standalone systems. Cameras, microphones, pressure sensors, vehicle access systems, connectivity modules, and radar technologies contribute information to a common understanding of what is happening inside the vehicle.

This approach enables the creation of virtual sensors. Unlike traditional physical sensors, virtual sensors are software-defined capabilities that derive their outputs from multiple information sources. A child presence detection function, for example, may combine UWB radar information, seat occupancy status, microphone inputs, and vehicle state information to produce a highly reliable assessment. Similarly, driver readiness may be determined by combining gaze tracking, steering interaction, vehicle dynamics, and cabin context.

The result is a more robust, scalable, and adaptable system capable of delivering new functionality through software rather than additional hardware.

Beyond single-purpose sensor architectures

Many of today's in-cabin sensing solutions rely heavily on dedicated sensing technologies.

Camera-based systems are widely deployed for driver monitoring applications because they provide detailed information about gaze direction, eye closure, head position, facial characteristics, and gestures. These capabilities are extremely valuable for understanding driver awareness and attention. However, cameras are fundamentally limited by their field of view and their dependence on visibility conditions. Occupants in rear rows, children covered by blankets, or objects hidden behind seats may remain difficult to detect reliably.

To overcome some of these limitations, many OEMs have introduced 60 GHz radar systems. These technologies provide additional sensing capabilities by detecting movement, occupancy, and even respiration through obstructions that would challenge purely optical solutions. While highly capable, dedicated radar systems often introduce additional hardware cost, processing requirements, and system complexity.

As regulatory requirements and customer expectations continue to expand, relying on dedicated sensing hardware for every application becomes increasingly difficult to justify. The challenge is no longer finding a sensor capable of solving a problem. The challenge is finding architectures that solve many problems efficiently.

Sensor reuse as a foundation for future architectures

One of the most powerful opportunities available to OEMs is the reuse of sensing assets that are already integrated into the vehicle.

Modern vehicles contain a rich collection of sensors and computing resources that often serve only a fraction of their potential capability. Cameras, microphones, connectivity platforms, access systems, wireless technologies, and centralized processors already generate valuable information about both the vehicle and its occupants.

Rather than introducing additional dedicated hardware, software-defined architectures allow these existing resources to contribute to multiple functions simultaneously. This not only reduces BOM (Bill of Material) cost but also simplifies system development and validation. Engineers can leverage familiar sensing infrastructure, existing data pathways, and proven software frameworks to support new capabilities.

Sensor reuse therefore becomes more than a cost optimization strategy. It becomes an architectural principle that enables greater intelligence while reducing complexity.

Ultra-Wideband as an enabler of sensor fusion

Ultra-Wideband (UWB) technology illustrates this concept particularly well. Originally adopted for secure vehicle access and precise ranging applications, UWB is now emerging as an important contributor to in-cabin sensing architectures. Beyond its positioning capabilities, UWB devices provide radar functionality that can detect subtle movement and physiological signals such as respiration.

This capability makes UWB particularly attractive for applications such as Child Presence Detection, occupancy monitoring, and intrusion detection. By analyzing phase shifts and time-of-flight measurements, UWB systems can identify signs of life even when occupants are obscured from view.

Perhaps more importantly, much of the required hardware may already exist within the vehicle as part of a secure access system. This allows OEMs to introduce additional capabilities without introducing entirely new sensing platforms.

Within a sensor fusion architecture, UWB should not be viewed as a replacement for cameras or other sensing technologies. Instead, it becomes another valuable source of contextual information that contributes to a more complete understanding of the cabin environment.

The role of AI in managing complexity

As vehicles become more sophisticated, one of the greatest challenges is determining how sensing resources should be used at any given moment.

Not every sensor needs to operate continuously. Some functions require maximum sensing performance, while others require maximum energy efficiency. Future context-aware systems will increasingly rely on AI to manage these trade-offs dynamically.

Agentic AI may eventually orchestrate sensing systems in much the same way that modern operating systems manage computing resources. Low-power sensors could remain active continuously, while higher-power systems activate only when additional information is required. The vehicle would dynamically balance sensing accuracy, energy consumption, system availability, and occupant experience based on current operating conditions.

In this way, AI becomes more than a consumer of sensor data. It becomes an active participant in determining how sensing resources are allocated throughout the vehicle.

Trust as a prerequisite for safety

While regulatory compliance and detection performance are important metrics, the long-term effectiveness of in-cabin sensing systems ultimately depends on user trust.

Many safety functions are designed with the intention of protecting occupants by generating alerts when potentially unsafe situations are detected. However, systems that produce excessive false positives can quickly become a source of frustration for drivers and passengers. Warnings that are perceived as unnecessary, mistimed, or inaccurate gradually reduce confidence in the system and increase the likelihood that users will ignore alerts or disable certain functions altogether.

Paradoxically, a safety system that generates too many false alarms can ultimately reduce overall vehicle safety. Even highly capable technologies provide little benefit if drivers no longer trust the information they provide.

As the number of in-cabin sensing applications continues to grow, avoiding alert fatigue becomes increasingly important. Future vehicles will need to distinguish between situations that genuinely require intervention and those that do not. This is where context-aware sensing and AI-based sensor fusion become critical. By combining information from multiple sensing modalities and evaluating events within a broader situational context, vehicles can significantly reduce false alarms while increasing confidence in the decisions they make.

The objective is not to generate more warnings. The objective is to generate the right warning, for the right occupant at the right time. Building systems that achieve this balance will be essential for maximizing both user acceptance and real-world safety.

Sensor fusion and the future of in-cabin sensing

No single sensor technology can address every use case associated with in-cabin sensing.

The future belongs to architectures capable of combining information from multiple sensing modalities and interpreting that information within a broader context. Sensor fusion improves reliability, reduces ambiguity, and enables capabilities that would be impossible through isolated sensors alone.

More importantly, software-defined sensor fusion creates scalable platforms that can adapt to different vehicle architectures. AI algorithms can be optimized for the sensors available within a given vehicle while maintaining common software frameworks, validation methodologies, and homologation strategies across multiple programs.

This flexibility allows OEMs to support future requirements without fundamentally redesigning their sensing architectures every time a new application emerges.

Conclusion

The future of automotive intelligence depends on understanding both the world outside the vehicle and the people inside it.

As vehicles progress toward higher levels of automation, in-cabin sensing must evolve from a collection of isolated features into a coordinated, AI-driven system capable of understanding context, anticipating needs, and supporting safe autonomous operation.

Achieving this vision requires more than adding sensors. It requires a shift toward software-defined architectures built on centralized computing, sensor fusion, virtual sensors, and intelligent reuse of existing vehicle assets.

The most successful vehicles of the future will not be those with the largest number of sensors. They will be those that extract the greatest intelligence from the sensors they already have.

That transition from distributed sensing to context-aware, AI-defined architectures represents the next chapter of in-cabin intelligence and an essential step toward truly autonomous vehicles.

Authors

Christoph Zorn
Business Development, NXP Semiconductors

Yannick Saalberg
Senior Product Marketing Manager, NXP Semiconductors

Christoph Zorn is a senior business development leader with over 12 years of experience in the automotive industry. With a strong foundation in product marketing for advanced keyless entry systems, he has spent recent years driving innovation in in-cabin sensing technologies. His expertise spans Ultra-Wideband, multi-sensor fusion architectures, and safety critical occupant monitoring solutions. Christoph holds a PhD in electrical engineering and is dedicated to advancing intelligent, secure, and user-centric vehicle interiors.

Yannick Saalberg brings 10 years of experience across the automotive and aviation industries, with a focus on innovation and advanced technologies. As senior product marketing manager for automotive Ultra-Wideband products, he focuses on in-cabin UWB applications. He holds a PhD in signal processing, a master's degree in information engineering, and a bachelor's degree in electrical engineering.

Document Number: RESPROBOTSWP REV 0