



Lessons From Automotive for Robotics

Global demand for robotics is surging. With applications ranging from industrial robots to [cobots](#), and from [autonomous mobile robots](#) to humanoids, [the market is expected to grow](#) from nearly \$50 billion in 2025 to over \$100 billion in 2030. With that rapid expansion come significant challenges: scaling efficiently, keeping costs down, integrating complex hardware and software, and sustaining performance in harsh, real-world environments.

The automotive industry has addressed each of these challenges for decades — at global scale and under stringent safety and reliability requirements. Vehicles are engineered to perform in extreme temperatures, high vibration and unpredictable conditions while delivering consistent quality across millions of units. This foundation of validated architectures, disciplined engineering and industrialized supply chains provides a proven model for robotics.

The parallels between the two industries are striking. Consider advanced driver-assistance systems (ADAS), which use intelligent perception to enable vehicles to make safe, real-time decisions about where to navigate. Or the high-performance connectors that ensure reliable power and data transmission to those vital sensors in demanding environments. Or the real-time, embedded software platforms that support the increasingly AI-powered applications that use these inputs. Combined with global manufacturing scale and supply chain strength, these capabilities, hardened in the automotive industry, can help robotics companies accelerate industrialization and deploy optimized systems with confidence.

ROBOTICS INDUSTRY CHALLENGES

Early robotics programs were typically bespoke systems in which engineers integrated compute units, sensors, power distribution systems, software modules and safety mechanisms one piece at a time. While this approach enabled innovation, elements that worked well in initial trials could show their fragility when faced with real-world performance demands. Systems were difficult and expensive to scale, challenging to maintain and not always robust enough for long service lives.

Today, several requirements are pushing the field toward a more robust, standardized architecture:

- **Modular, Scalable Platforms**
Organizations in industries including logistics, manufacturing, healthcare and mobility want to deploy fleets of robots, not isolated pilots. This requires standardized layers that can be reused across product lines.
- **Intelligent Perception Systems**
Robots rely not just on robust sensors, but also on the deep learning, multi-sensor fusion and real-time inference driven by AI and machine learning that make sense of the world around them. These workloads demand high-performance compute, efficient data transport and rigorous system optimization.
- **Robust Electrical Systems**
Robots require power and data architectures that can hold up under tough environmental conditions and be reliable over long time periods.
- **Safety, Security and Regulatory Requirements**
As robots operate closer to people, functional safety, redundancy, cybersecurity and certification frameworks become essential. These cannot be bolted on — they must be architected in.

All of these requirements have been addressed in the automotive industry, especially as it has moved to vehicles that are more software-defined. Taken together, these elements create a platform — one that provides the hardware and software foundation that OEMs can leverage to build applications that are differentiated and innovative.

SOFTWARE FOUNDATION

Software-defined vehicles (SDVs) and modern robotics are built on the same foundational technology stack. Both domains rely on tightly integrated combinations of hardware, software and communications protocols — spanning real-time operating systems (RTOS), middleware, edge intelligence and physical components such as sensors, actuators and high-reliability connectors. These shared building blocks form the backbone of physical AI.

SDVs are, in effect, highly sophisticated mobile robots: They must perceive their environment, make real-time decisions and execute precise physical actions — all under strict constraints for safety, determinism and reliability. Those same requirements define robotics applications across industries, from manufacturing to logistics to service robotics.

Because these systems interact directly with the physical world, they depend on deterministic real-time operating systems (RTOS) to guarantee predictable performance. Determinism enables coordinated motion control, accurate sensor fusion and immediate response to safety-critical events such as obstacle detection or emergency stops.

In the great ongoing example of a “road tested” RTOS, VxWorks has been used on Mars missions since the 1997 Pathfinder lander and has powered numerous rovers and spacecraft ever since. Its proven reliability has led to continued use decades later, including as [a key RTOS within critical subsystems of NASA's Artemis II](#) crewed mission. The same RTOS principles that enable

reliable operation in space also underpin safe, deterministic performance in SDVs — and, increasingly, in next-generation robotic systems.

For higher-level functions such as perception, AI processing and system integration, Linux can be very effective. For example, Yaskawa's MOTOMAN NEXT robot uses Wind River Linux to support AI-driven perception and adaptive behavior in industrial automation. Built on Yocto, [Wind River Linux](#) allows teams to tailor a minimal, hardware-specific OS while maintaining long-term support and security updates. It works well with frameworks like ROS 2 and is commonly used to host non-real-time workloads — such as planning, data handling and connectivity — alongside real-time control systems.

PERCEPTION AND DATA

The robotics industry is entering a phase that automotive has already navigated: the transition from breakthrough capability to dependable, scalable performance. Nowhere is this more critical than in perception. In both industries, perception is the foundational layer — feeding


every downstream function, from motion planning to actuation. Automotive has demonstrated that robust perception is not just about sensing the environment, but also about structuring, validating and operationalizing data so that machines can make safe, real-time decisions under uncertainty.

Decades of investment in automotive perception systems offer a clear blueprint. Advanced driver assistance systems and automated driving platforms rely on multimodal sensor fusion to create a redundant, high-confidence understanding of the environment. The critical lesson for robotics is that no single sensor modality is sufficient. Optimization requires tightly integrated sensor fusion architectures, where complementary data streams are synchronized, time-stamped and processed through unified perception pipelines to reduce uncertainty and enable deterministic behavior.

Equally important is the data lifecycle surrounding perception. In automotive, perception performance is continuously improved through large-scale data collection, annotation and model training pipelines. By implementing

Automotive Technologies on AMRs

Over decades, the automotive industry has developed, refined and hardened a wide set of advanced technologies that are also essential to robotics.



AUTONOMY & MOTION INTELLIGENCE

- Autonomy
- AI/ML-based motion planning

PERCEPTION & SENSOR FUSION

- Base perception
- Depth maps
- Occupancy grid
- Object detection

OPERATING SYSTEMS & SUPPORTING SOFTWARE

- Linux
- Middleware
- RTOS

closed-loop data architectures — where operational data is systematically fed back into development — robotics systems can evolve from reactive machines into adaptive learning systems.

Safety is where perception, data and architecture converge. Automotive has established rigorous functional safety processes (such as ISO 26262) that require not just accurate perception, but also predictable system response under failure conditions. A proper approach should integrate redundancy at both the hardware and software layers — diverse sensor modalities, independent processing paths and fallback strategies — to ensure graceful degradation rather than catastrophic failure. For robotics operating in dynamic, human-populated environments, this model is essential.

Finally, optimization in robotics will depend on platformization — the ability to deploy perception capabilities across diverse applications without reinventing the stack each time. Automotive has already moved toward centralized compute and software-defined architectures, where perception algorithms run on scalable, high-performance compute platforms with standardized interfaces. For robotics, adopting similar principles means decoupling perception software from hardware configurations, enabling rapid deployment for a wide variety of devices.

HIGH-PERFORMANCE CONNECTORS

For automotive OEMs, connectors are highly engineered yet relatively inexpensive components — manufactured at massive scale and optimized for cost without sacrificing performance. This balance of robustness and affordability has made automotive-grade connectors increasingly attractive to industries seeking reliable solutions without the expense of bespoke designs.

At their core, connectors must resolve a fundamental engineering tension: maintaining secure mating under sustained vibration and environmental stress, while still allowing for serviceability. Automotive designs address this through features such as secondary assurance mechanisms — redundant locking systems that ensure connection integrity even under mechanical strain. These performance requirements are met consistently at scale.

As vehicle electronics have evolved, connectors have become integral to high-speed data ecosystems. They now support communication protocols such as CAN and, increasingly, automotive Ethernet (including single-pair Ethernet), enabling reliable transmission of both power and high-bandwidth data within complex distributed architectures. Cost pressures have further driven material innovation, favoring widely available, manufacturable materials over specialized alternatives.

Standards such as those developed by USCAR have reinforced this evolution, emphasizing interoperability, durability and data integrity. Their growing adoption in robotics and industrial automation reflects a broader shift toward proven, cross-industry architectures.

For robotics manufacturers, the implications are significant. Leveraging automotive connector ecosystems enables faster design cycles, reduced qualification risk, and access to mature global supply chains. This, in turn, supports production scaling — moving from low-volume builds to repeatable, high-volume manufacturing with consistent quality and cost control. By adopting automotive-proven components and standards, robotics companies can accelerate industrialization while maintaining the reliability required for real-world deployment.

SCALE AND SUPPLY CHAIN

For a company to move forward, bridging the gap between prototype and production is paramount. Doing so successfully requires a partner that can provide both cost optimization and a resilient supply chain. A true scalability partner should meet current needs and be capable of supporting future expansion, so growth remains unimpeded.

Companies like Aptiv that have traditionally supplied the automotive industry have well-established capacity in property, plant and equipment that can be used to meet current and projected growth demand. That minimizes the need for incremental investment, lowers risk and delivers inherent cost advantages.

Cost Optimization

As fleets expand, cost pressures shift from secondary concern to primary constraint. Architectures that rely on expensive sensors and general-purpose compute platforms become difficult to scale economically, especially in markets where competition is high. This creates a transition point: Systems that are optimized for proof-of-concept must evolve into systems designed for production efficiency.

In high-volume industries such as automotive, this transition has already driven a shift toward sensor and compute configurations that balance performance with cost at scale — favoring approaches that minimize reliance on high-cost components while maintaining robustness.

This is where a different design philosophy becomes critical. By leveraging components and architectures developed for high-volume domains, robotics systems can benefit from economies of scale. For example, sensors and compute platforms produced at automotive scale offer substantial cost advantages compared to specialized robotics components, without sacrificing performance when paired with effective sensor fusion and AI-driven perception. In this context, AI is applied selectively —

augmenting traditional methods where needed to improve robustness and coverage, rather than as a blanket solution. This approach enables robotics platforms to move beyond early deployments and toward scalable, cost-efficient production, aligning system performance with the economic realities of large fleet deployment.

Supply Chain Resilience

Organizations at minimum should insist on Tier 3 visibility, which gives insight beyond direct suppliers into the sub-suppliers and raw-material providers where many of the most significant risks originate. By mapping the supply chain to this depth, companies can identify hidden dependencies, geographic concentrations and single points of failure that would otherwise remain intractable, untraceable problems.

This transparency enables earlier detection of disruption risk, more credible human-rights and environmental due diligence, and better-informed sourcing and continuity decisions. In practice, Tier 3 visibility strikes a critical balance: It is deep enough to surface material operational, regulatory and reputational risks, while remaining practical to maintain and act upon. For customers, it translates into greater resilience, faster response to unexpected events and higher confidence that supply chain decisions are based on a realistic view of how goods and materials actually flow.

GOING FORWARD

From the resilience of the 1970s Martian rovers to the impressiveness of modern humanoids, robots depend on technologies ranging from real-time operating systems to intelligent perception to fundamental connector hardware. Success requires a trustworthy technology partner with verified experience in developing mission-critical products and systems, along with a background in regulated markets and the ability to self-certify products to cut down on testing time and red tape.

Scaling robotics systems demands a shift from tightly coupled, prototype-driven architectures to platforms that support modularity, standardized interfaces and lifecycle management at scale. It also requires experience navigating the practical realities of deployment — thermal constraints, power management, environmental durability and system validation under real-world conditions. Organizations that succeed in scaling are those that can integrate these considerations early, rather than retrofitting them after initial deployment. By applying lessons learned from automotive-scale production— where reliability, cost and performance must be balanced across millions of units — robotics companies can move beyond proof-of-concept systems toward industrial-grade deployments. This approach not only accelerates time to market, but also establishes a foundation for continuous improvement, ensuring that robotics platforms can evolve alongside advancing AI capabilities and increasingly complex operational demands.

Taken together, Aptiv's deep expertise in automotive-grade hardware integration, safety-critical software and large-scale deployment uniquely positions the company at the intersection of SDVs and robotics. Rather than separate domains, these industries represent a continuum of physical AI — one where lessons proven on the road can be directly applied to the factory floor, the warehouse and beyond.

SENSE, THINK, ACT, OPTIMIZE

Today, most people associate artificial intelligence with tools like ChatGPT and the large cloud infrastructures that support them. But these tools exist in a purely digital world. *Physical AI* marks a pivotal shift in artificial intelligence — the fusion of advanced AI models, including generative models, with machines that operate in the physical world.

Any software-based AI can model physical laws in the abstract — calculating thrust from aerodynamic equations, for example. Physical AI systems don't just model physics; they are designed to experience it, with every action shaped by real-world friction, uncertainty, latency and consequence, continuously feeding back into perception and decision-making.

This perspective has roots in academic research on *Embodied AI*, whose proponents contended that true intelligence requires sensory experience — learning not just from data, but from interaction with the world itself.

A second branch of physical AI's lineage comes from classical robotics. Beginning in the 1980s, robots were typically organized around a simple control loop: *Sense, Plan, Act*. Sensors captured the environment, an internal model represented it and actions were executed based on that representation. Over time, "Plan" evolved into "Think," but the conceptual triad remained the same: perception → deliberation → action.

A critical turning point in this evolution came from roboticist Rodney Brooks, who challenged the dominant Sense, Plan, Act paradigm by arguing that intelligent behavior does not require explicit internal representations or detailed world models. His behavior-based approach emphasized fast, reactive sensor-actuator loops, asserting that robust intelligence emerges from direct interaction with the physical world rather than abstract reasoning about it.

Today's robotics, filtered through both Brooks' ideas and the arguments of his critics, now contain a critical fourth dimension, in which the triad has become a tetrad: *Sense, Think, Act, Optimize*.

Optimization closes the loop across time. It allows physical AI systems to learn from outcomes, adapt to changing conditions and improve performance, safety and efficiency at scale. This is no longer a static control system — it's a living one, refining itself with every mile driven, every maneuver executed and every edge case encountered.

As AI increasingly moves beyond cloud-based data centers — and the digital-in/digital-out interactions they typically power — to address physical world challenges and opportunities, the promise of physical AI is fulfilled: machines that don't just operate in the world but progressively master it.

