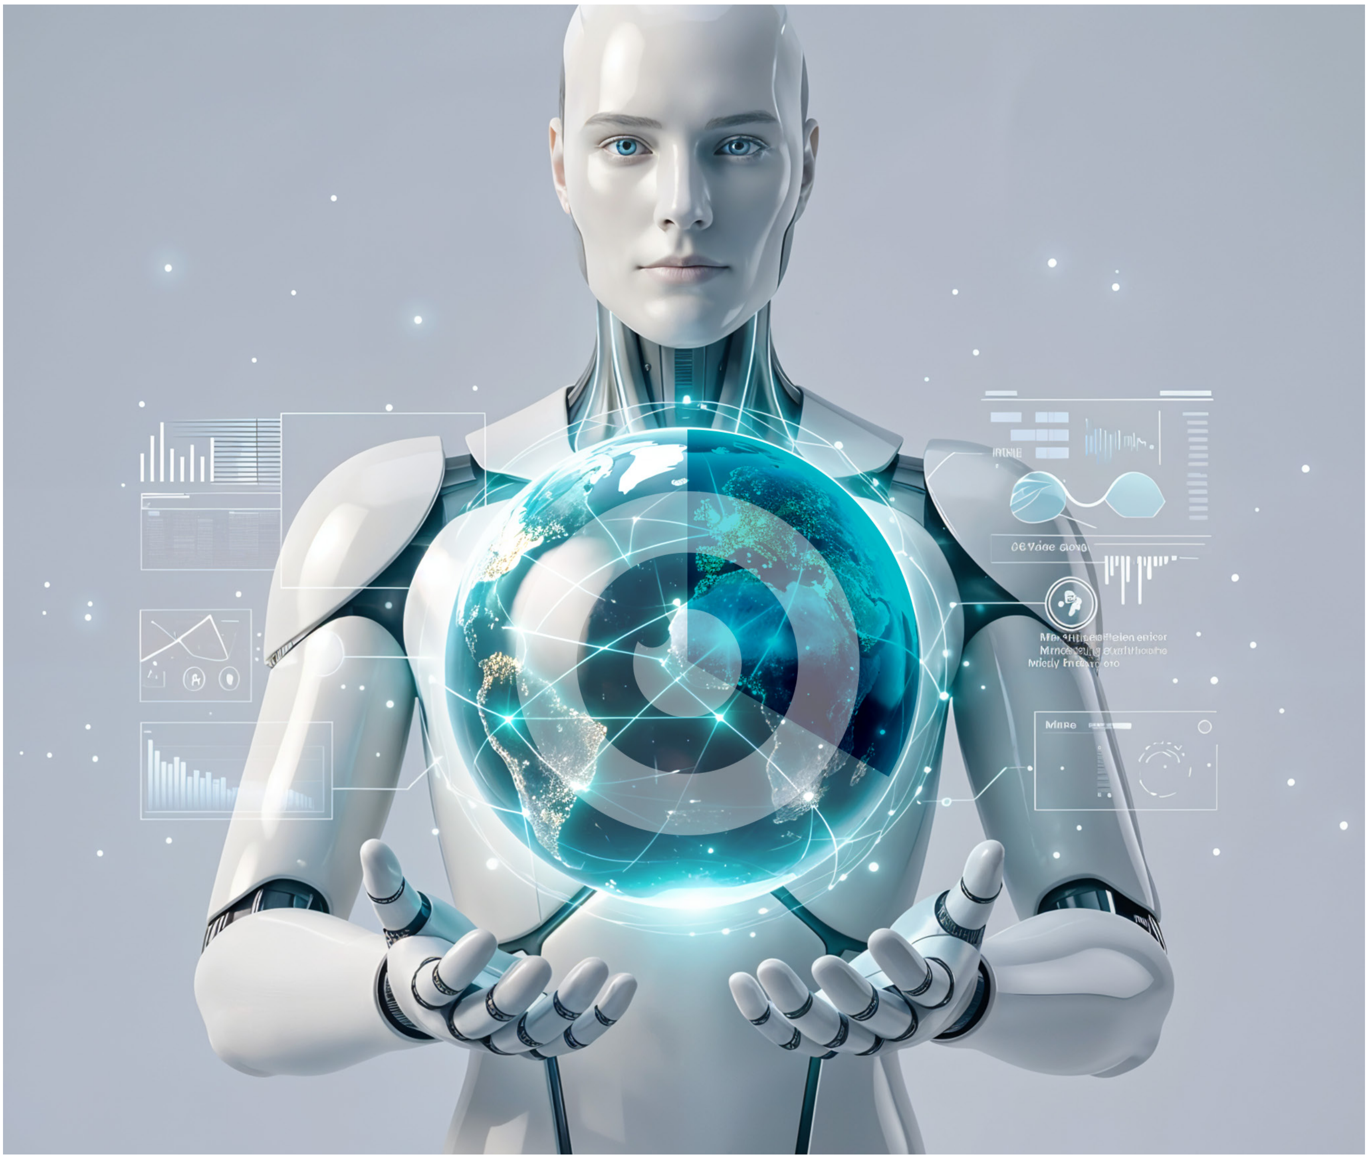

# Innoviz and the Rise of Physical AI

## Bringing World Models to Life


**Omer Keilaf, CEO and Founder**

Innoviz Technologies | $INVZ


March 2026

INNOVIZ™
TECHNOLOGIES

# Executive Summary

**The "Software is Eating the World" era is over. Now, "AI is Eating the Software". SaaS is facing massive disruption as a significant percentage of code becomes AI-generated. The moat is no longer the code; it's the data and infrastructure and the real-world integration. The biggest shift in 2026: The rise of Physical AI. AI can't operate a truck or a drone just by "guessing" the next pixel like a video generator. It needs a high-fidelity, 3D understanding of physics and geometry to train World Models capable of creating digital twins of reality. The market opportunity for Physical AI applications is significant and still developing. Physical AI applications could grow to $0.5-$1.4 trillion market by 2035, and we see LiDAR is well positioned to serve as a primary sensing layer. Innoviz is a market leader providing automotive-grade LiDAR to several global OEMs, and is well positioned for the coming Physical AI inflection point.**

Artificial Intelligence (AI) is entering a new phase. After transforming digital systems through software and Large Language Models (LLM), AI is now moving into the physical world, powering vehicles, robots, infrastructure, and machines that must perceive, reason, and act under real-world constraints in real time. This transition, often referred to as Physical AI, represents one of the largest and longest-duration technology opportunities of the coming decades.

AI can't operate a truck or a drone just by "guessing" the next pixel like a video generator. It needs high-fidelity, 3D understanding of physics and geometry to train World Models capable of creating digital twins of reality. It must function in safety-critical environments, tolerate environmental variability, and scale at infrastructure level. As a result, it requires a fundamentally different technological foundation. At the center of that foundation is perception.

World Models are a concept from machine learning and artificial intelligence where a system learns an internal representation of how the world works. You can think of a as the AI's mental model of reality, a way for it to understand environments, predict what will happen next, and plan actions. Such models require substantial processing capabilities and strong software development in a scale which has never been used before.

To achieve that, Dassault Systèmes and NVIDIA recently announced a long-term strategic partnership tocreate "Industry World Models". These are AI models grounded not in text or images, but in physics, engineering laws, materials science, and validated industrial knowledge. These Industry World Models are meant to simulate, 3D reconstruct, and emulate highly complex real-world systems, from materials to industrial robots, factories, and entire cities.

As Physical AI systems scale, they face a growing "incestuous data" problem: models are increasingly trained on synthetic or AI-generated data rather than real-world observations. This creates a feedback loop where errors, biases, and oversimplified assumptions become amplified over time, gradually distorting the model's understanding of reality.

Innoviz LiDARs will bring World Models to life. By feeding rich, real-time spatial data into the AI 'knowledge factories' running on advanced processors, our LiDAR technology transforms

World Models into living environments that continuously learn from the physical world and deliver a real-time understanding of how it operates.

The physical world is three-dimensional. Distance, occlusion, motion, and free space are geometric properties. Cameras capture two-dimensional projections of that world, requiring neural networks to infer depth based on probabilities. LiDAR, by contrast, measures distance directly by converting photons into precise geometry. In safety-critical systems, measurement is structurally more reliable than inference. As computing power continues to accelerate, the primary bottleneck in deploying Physical AI is no longer reasoning capability but access to high-fidelity, real-time 3D data.

As models have become a commodity, data has become the differentiator. World Models enable AI to build a representation of physics. To scale, they need a massive influx of ground truth tokens from the dynamic real world.

Deployment of autonomous vehicles is defining the performance threshold for perception infrastructure. The requirement for functional safety, availability, durability, cost efficiency, and high-volume manufacturability has eliminated many technologies that could not meet production-grade requirements.

After a decade of industry consolidation, only a small number of LiDAR suppliers have achieved true automotive readiness. Innoviz is among the few companies to reach series production of a Level 3 (L3) passenger vehicle and is advancing toward Level 4 (L4) deployments at scale.

Beyond automotive, Physical AI is expanding across robotics, smart infrastructure, industrial automation, defense, and mapping. When reframed through this broader lens, the Total Addressable Market for LiDARs will reach an estimated 35 to 40 billion dollars over the next decade, with multi-vertical and infrastructure-like growth characteristics.

As nations define long-term AI strategies, perception technology is becoming a strategic layer of physical intelligence systems. Trusted automotive-grade, scalable 3D perception will form a foundational component of the Physical AI era.

Innoviz is well-positioned at this inflection point, providing the deterministic, high-performance LiDAR systems that are required to enable Physical AI safely, securely, and at global scale.

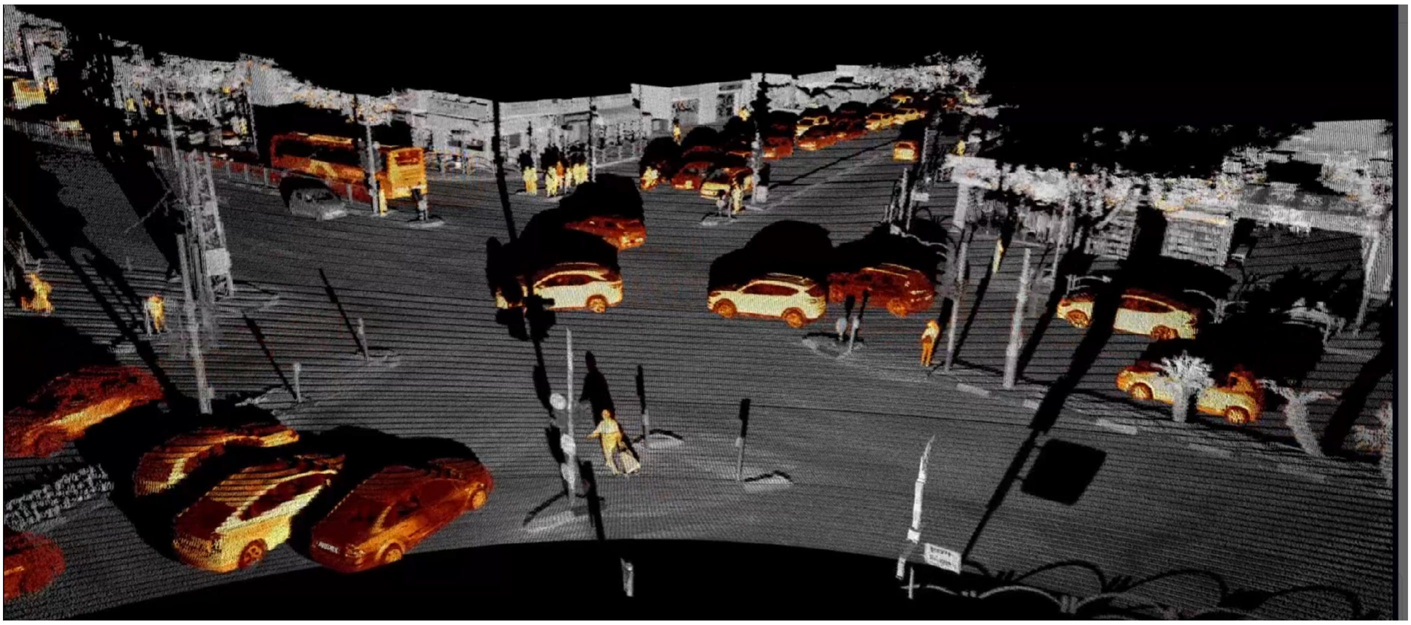

*Innoviz Point Cloud*

# In This White Paper

# PART I

Innoviz and the Rise of Physical AI

# The Shift from Digital AI to Physical AI

## From Software Intelligence to Embodied Intelligence

For decades, artificial intelligence existed primarily in the digital domain. Early systems optimized databases and classified information. More recently, Generative AI transformed how software can reason, create, and communicate.

Physical AI represents a fundamental break from this trajectory. Rather than operating on symbols and abstractions, Physical AI systems interact directly with the real world. Autonomous vehicles, robots, smart infrastructure, and industrial automation systems must perceive their surroundings, understand spatial relationships, predict outcomes governed by physics, and execute actions with real-world consequences.

This shift is not incremental; it marks a transition from intelligence that describes the world to intelligence that operates within it.

## Why This Transition is Happening Now

Several forces are converging:

- Advances in AI reasoning and planning
- Edge computing that is capable of real-time decision making
- Labor shortages driving automation
- Maturing sensor technologies capable of digitizing reality

Most critically, Physical AI is no longer just demonstrable in controlled environments; it is now becoming deployable at scale.

The remaining constraint is no longer intelligence, but high bandwidth and accurate, reliable perception.

## Economic Implications

Physical AI follows a different economic curve than Digital AI:

- Capital-intensive development
- High regulatory and safety barriers
- Constrained production and supply chain
- Infrastructure maturity and robustness versus trial-and-error

Historically, these conditions favor companies that provide foundational infrastructure that can serve as accelerators, not point solutions. It took many years for NVIDIA to establish its position as the leader in Physical AI infrastructure by developing the most advanced platform for AI training, development, and implementation. The complexity as well as the enormous cost of creating a Physical AI platform is the reason for the rather low number of infrastructure players. Nevertheless, an opportunity exists for companies that are determined to succeed in this pursuit.

Perception is one such foundational layer and only a few global LiDAR players remain. The stringent requirements are currently driven by the automotive space, which constantly pushes for lower costs, size and power, and performance that has never been seen before. The narrowing of the players has not ended yet, and we believe Innoviz is well-positioned to lead.

# Perception: The Bottleneck in Physical AI

## Operating Under the Laws of Physics

While Digital AI can tolerate ambiguity, Physical AI cannot. It must operate with certainty. Digital AI is highly dependent on the Physical AI data integrity. Data acquisition must be accurate, timely, and available at all times, regardless of environmental conditions. If the perception layer is unstable or unreliable, the Digital AI may collapse.

There are many reasons why LiDARs are employed by all vehicle manufacturers, including Tesla, to collect ground-truth data. Errors translate directly into safety risks, system failure, and/or regulatory non-compliance. Cameras infer depth indirectly and degrade under low light, glare, or poor weather. Radars lack spatial resolution. Software cannot compensate for missing or unreliable ground-truth data. Without accurate, real-time 3D perception, higher-level AI reasoning becomes fragile.

Another benefit of LiDAR technology is that it is more privacy-friendly than cameras. It captures only depth and shape information, rather than detailed images of people. Instead of recording faces, clothing, or other identifiable features, LiDAR produces anonymous point cloud data that shows where objects are located without revealing their identity. This makes it more difficult to misuse this technology for surveillance, reduces the risk of collecting personal information, and avoids many of the legal and ethical concerns that come with camera-based systems. For these reasons, we believe society will be willing to adopt LiDARs as a key element of the future Physical AI infrastructure. Consequently, LiDAR will become ubiquitous in our lives as it is safer than camera-based systems.

## Perception as Infrastructure

Perception is not a feature; it is part of the infrastructure. It must be deterministic, robust, scalable, and trustworthy. As Physical AI systems scale, value increasingly accrues to companies that provide reliable perception as part of the infrastructure.

Production of LiDARs is one of the most underappreciated challenges in the creation of the Physical AI infrastructure. As a result, only a few companies have succeeded in developing and manufacturing them. Each LiDAR technology requires a unique production process which includes accurate optical assembly, thousands of components, reliable supply chain, and a very extensive calibration and testing process that automotive industry standards demand for the LiDAR's functional safety. Because Digital AI models expect high uniformity between all sensors, this drives expectations for deterministic and highly uniform behavior across all production tolerances.

The stringent and highly complex quality standards that govern automotive manufacturing push suppliers to elevate their capabilities far beyond basic production needs. To meet these expectations, suppliers are compelled to invest in advanced infrastructures ranging from automated inspection systems to robust process control frameworks that ensure consistency, traceability, and reliability at every stage. Although these systems initially were developed to satisfy the demanding requirements of the automotive industry, they ultimately create a strong technological foundation that can later support large-scale mass production at the required quality levels. Over time, this pressure not only improves suppliers' internal processes but also prepares them for broader industrial expansion, allowing their control frameworks to scale effectively.


*Innoviz Production Line*

# Hello World (Models)

**World Models are a new class of AI systems designed to understand the real world by grounding intelligence in physics, engineering principles, materials behavior, and validated scientific knowledge, rather than in text alone.**

Unlike Large Language Models (LLMs), which predict words based on patterns, World Models simulate cause and effect, enabling AI to reason about how things work in the physical world. This allows them to model everything from molecules and biological systems to machines, factories, and large-scale industrial processes. In essence, a World Model is a computational universe, an environment where AI can safely explore, test scenarios, and make reliable predictions rooted in the laws of nature.

NVIDIA and Dassault Systèmes are partnering to create Industry World Models. They leverage NVIDIA's accelerated computing, AI infrastructure, Omniverse technologies, and open AI models together with Dassault's decades of scientific and engineering knowledge that is embedded in its Virtual Twin simulation and modeling tools. Their goal is to build AI systems that can design, simulate, optimize, and operate complex industrial systems, from manufacturing plants and supply chains to aerospace components, vehicles, and even advanced materials research. These models are "science-validated", meaning they rely on real engineering constraints and physics-based simulation rather than assumptions, making them fit for mission-critical industrial decision-making.

To power these World Models, Dassault Systèmes' cloud platform, OUTSCALE, will deploy AI factories across multiple continents, running large-scale simulations and training industrial-grade while ensuring data sovereignty and IP protection for customers. At the same time, NVIDIA is adopting Dassault's model-based systems engineering methods to design its own AI factories, starting with the Rubin platform, which uses virtual twins to simulate and optimize entire data centers before they are built. Together, these companies aim to establish a foundational architecture for Physical AI, transforming how industries innovate by enabling AI systems that deeply understand and reliably reason about the physical world.

Innoviz LiDAR sensors can help bring Industry World Models to life by supplying the real-time, high-resolution 3D data these models need to stay accurate, adaptive, and grounded in the physical world. World Models created by NVIDIA and Dassault Systèmes rely on physics-validated simulations. To remain useful in real operations, such as factories, vehicles, smart cities, or robotics, they must be continuously updated with live spatial information.

As Physical AI accelerates, it is running into a growing "incestuous data" problem, a structural risk created when models are trained predominantly on synthetic, simulated, or AI-generated data rather than on real-world sensor inputs. Synthetic data is incredibly useful for scale, but it also introduces a subtle form of model inbreeding. The outputs of one model become the training inputs of another, and over time, the entire ecosystem begins to drift away from the physical world it is meant to represent. This feedback loop amplifies hidden biases, smooths over edge cases, and reinforces the assumptions baked into the original synthetic generators. In domains like autonomous vehicles, robotics, industrial automation, and spatial intelligence where centimeter-level accuracy and precise physical reasoning are non-negotiable, this drift can be catastrophic. Without continuous grounding in high-fidelity, real-world data, Physical AI systems risk developing an increasingly distorted understanding of geometry, motion, materials, and cause-and-effect. The result is a widening gap between simulated performance and real-world reliability. Solving this problem requires a sustained investment in real sensor data, robust world modeling, and hardware-anchored perception, because Physical AI cannot afford to learn from a world that only exists inside its own imagination.

Innoviz's LiDAR technology provides precisely this: unparalleled, high-precision 3D point cloud data that captures objects, environments, and dynamic changes with exceptional fidelity, enabling accurate mapping, object validation, and environmental understanding. Additionally, Innoviz's solutions, including edge-processed real-time perception systems like InnovizSMARTer, are designed to deliver gigabits-per-second 3D sensing, compressed and processed efficiently for real-world deployment, aiming to ensure that even complex, evolving environments can be streamed back into World Models in real-time. In this way, Innoviz LiDARs have the potential to act as the high-detail sensory layer that constantly feeds the industrial AI models fresh data, which will allow simulated "worlds" to remain synchronized with reality and enabling Physical AI systems to make reliable, situationally-aware decisions.

We believe Innoviz is going to be a big part of the future of the World Models. LiDARs also have the potential to provide high-resolution 3D data for validated simulations and, when integrated with real-time perception systems, to generate precise 3D point cloud data for operational deployment. This ensures that even complex, evolving environments can be streamed back into World Models in real-time.

*Camera View of Las Vegas Scene*

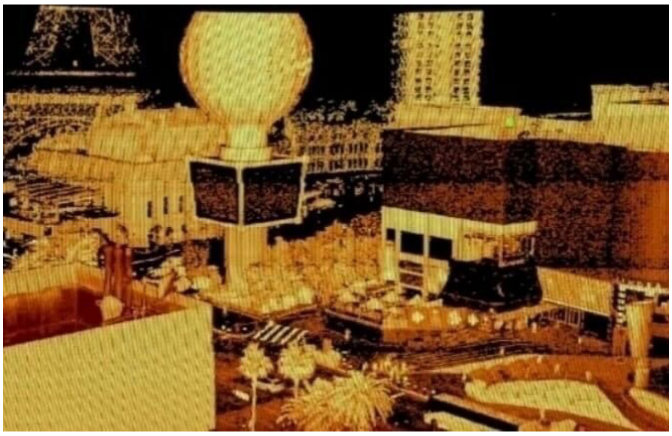

*LiDAR Point Cloud View of Las Vegas*

## Total Addressable Market in the Physical AI and World Model Era

Assessing the Total Addressable Market (TAM) for an emerging category is challenging, especially when Physical AI is only in its infancy. Still, Barclays estimates that by 2035, Physical AI applications such as robots, autonomous vehicles (AVs), industrial automation, and drones could grow into a $0.5-$1.4 trillion market, with AVs contributing nearly half of that expansion (around $550 billion). Notably, China is expected to dominate early adoption, representing 85% of new humanoid robot installations in 2025, compared with just 13% in the United States.

Traditional TAM estimates for LiDAR in autonomous driving are around $10 billion, where LiDAR typically complements the camera and radar, not replacing them. However, unlike the automotive industry, where LiDAR is added as a complementary layer, many other market segments can adopt LiDAR in place of cameras and low-resolution radar systems due to its clear performance advantages. As a result, the LiDAR market may gradually begin expanding into these adjacent addressable segments. Innoviz's third-generation sensor, InnovizThree, incorporates an onboard camera, delivering a unified 3D imaging platform that combines the strengths of LiDAR with camera-based visual capabilities.

| | Surveillance/ Security | Industrial | Smart Cities | Logistics | Robotics | UAV |
|---|---|---|---|---|---|---|
| **Radar** | $4.5B (2024) | $1.92B (2026) | $7.9B (2025) | $<1B (2026) | $0.5B (2024) | $1.4B (2024) |
| **Camera** | $16B (2025) | $2.5B (2026) | $16.4B (2025) | $2. 8B (2024) | $3.3B (2025) | $2.2B (2024) |

*Physical AI TAM for Cameras and Radars, excluding Automotive. Company estimates based on analysis of third-party market data.*

Over the next decade, both camera and radar markets are projected to continue expanding at annual growth rates of roughly 6%, indicating that these sectors remain on strong growth trajectories. At the same time, the LiDAR's unique performance advantages position it to disrupt many of these existing markets while also enabling entirely new applications.

Consequently, the LiDAR market is likely to experience robust growth as its performance advantages drive both market disruption and new application creation.

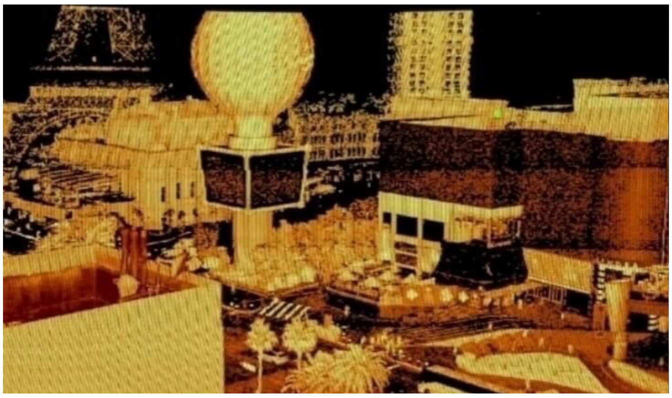



# PART II

Innoviz and the Rise of Physical AI

# Physical AI and Autonomous Driving

**Over the past decade, Physical AI has been significantly shaped by the Autonomous Driving industry, where leading technology companies, traditional automakers, and well-funded startups invested in developing highly capable autonomous driving systems. These technologies allowed AI to interpret its surroundings, anticipate events, and make real-time decisions about how a vehicle should behave. In many ways, autonomous driving has served as the "proving ground" for Physical AI, demonstrating both its potential and its challenges under some of the most demanding real-world conditions.**

Today, the AI landscape looks very different from where it began. AI development has accelerated dramatically, fueled by advances in compute, data, and model architecture. At the same time, costs associated with AI development, once a barrier for all but the largest companies, have fallen considerably. There is now broad recognition that AI can transform nearly every industry, from consumer electronics and smart cities to robotics, logistics, industrial automation, and defense. As a result, there is a growing appetite to place AI at the center of physical interaction, embedding intelligence into machines, devices, and infrastructure across society.

In parallel, LiDAR technology, a critical enabling sensor for Physical AI, has undergone a transformation of its own. Over the last decade, LiDAR systems have become significantly more advanced, more reliable, and far more economical. Improvements in semiconductor design, manufacturing processes, optical engineering, and system integration have driven substantial cost reductions, bringing sensors that once cost tens of thousands of dollars to sub-thousand-dollar price points at scale. These advancements have made LiDAR accessible not only to automotive OEMs but also to a wide range of new industries seeking to deploy Physical AI at scale. In essence, the same forces that once propelled autonomous vehicle development are now supporting the broader rise of intelligent machines across many sectors.

However, as we hand over increasingly critical decisions to AI, whether in vehicles, robots, industrial machinery, or public infrastructure, one requirement stands above all else: safety. For AI to assume responsibility in real-world environments, it must demonstrate reliability that meets or exceeds human-level performance. This demands sensing and perception systems that can consistently deliver accuracy, robustness, and redundancy under the full spectrum of environmental conditions.

Here, Physical AI can draw from the decade of rigorous development that occurred in the autonomous driving sector.

The automotive industry has been uniquely demanding, pushing LiDAR manufacturers to meet stringent requirements for quality, safety, reliability, and cost, requirements that most other industries had not previously encountered at comparable scale.

Through multiple waves of development cycles targeting Level 2, Level 3, and higher levels of driving automation, automotive OEMs continuously refined what was required for LiDAR to be viable in mass-market deployment.

This relentless pressure shaped the entire LiDAR industry. Dozens of companies entered the market in the early years, each bringing different approaches, architectures, and technologies. Over time, as OEMs raised expectations for performance, safety certifications, automotive-grade design, and long-term reliability, the field narrowed dramatically. Many LiDAR companies that once garnered significant investment were unable to meet these increasing demands. Several declared bankruptcy, others shut down operations, and others merged with competitors. The industry has since consolidated to a small number of serious contenders. The next phase of the automotive market, characterized by tighter integration, more complex safety requirements, and cost-optimized scalability is expected to further concentrate the industry around those players with proven mass-production capability, automotive-grade qualification, and scalable cost structures.

These developments create a compelling moment of opportunity. The foundation built in automotive, one of the most demanding and regulated sectors in the world, now serves as a reference point for the broader expansion of Physical AI. Industries adopting Physical AI today can leverage a decade of lessons, standards, and technological breakthroughs that were developed under rigorous conditions. Companies with demonstrated automotive-grade technology and production capabilities are well-positioned to participate in this next wave of adoption.

# The Last 10 Years in the Automotive Industry

**The trajectory of the autonomous driving market and the broader transformations in the automotive sector are best understood through the framework of the Gartner® Hype Cycle®, which illustrates how new technologies progress from early excitement to practical maturity.**

In any industry undergoing significant disruption, an initial surge of enthusiasm attracts a wave of new entrants. Startups, established companies, and investors all position themselves to capture value in what appears to be an immense and rapidly forming market. During this phase, capital is abundant, multiple technological approaches compete, and investors place numerous bets, hoping that their chosen companies will survive long enough to emerge as market leaders after the inevitable shakeout.

## 2015-2016 | Phase One – The Peak of the Hype

The autonomous driving industry reached its peak of inflated expectations around 2015-2016, a period marked by bold predictions, aggressive timelines, and widespread belief that fully autonomous vehicles were only a few years away. Billions of dollars were invested, dozens of LiDAR companies were founded, and nearly every major automotive Original Equipment Manufacturer (OEM) announced ambitious self-driving roadmaps. But as with any emerging technology, the industry soon transitioned into the more difficult phases of the Hype Cycle.

## 2017-2019 | Setbacks

Between 2017 and 2019, the narrative shifted. High-profile setbacks and missed milestones began to surface. Several programs encountered technical or operational obstacles. Some autonomous vehicle pilots were scaled back or paused. A number of LiDAR startups that had attracted early excitement were unable to meet rising technical requirements or secure long-term funding and ultimately closed their doors.

## 2020 | The Effect of COVID-19 and SPACs

Then came the COVID-19 pandemic, adding yet another layer of disruption. Factory shutdowns, delayed testing cycles, supply chain interruptions—including the global semiconductor shortage that peaked in 2021-2022 and tightened budgets—forced many companies to scale back ambitions or reprioritize. The combined impact of these challenges led some observers to conclude that the autonomous driving vision had stalled.

At the same time, the automotive industry was undergoing a second, equally transformative disruption: the accelerated global shift toward electric vehicles (EVs). This shift had a profound impact on OEM priorities and engineering resources. A number of OEMs had planned to use new EV platforms as the technological foundation for vehicles integrating advanced driver assistance and Level 3 autonomy features, as these EV architectures were designed to incorporate advanced compute, sensing, and electrical systems suitable for high-level automation.

However, as the EV market evolved, OEMs faced increasing pressure on profitability, supply chains, battery sourcing, manufacturing capacity, and regulation. Some EV programs fell behind schedule, experienced cost escalations, or required significant reengineering. As a result, certain platforms that were originally intended to launch as EVs were redesigned to support ICE (internal combustion engine) or hybrid powertrains.

This shift had two important consequences: First, timelines for next-generation vehicle platforms slipped, in some cases by several years. Second, Level 3 autonomy features that had been planned for introduction on these new platforms were delayed, since OEMs traditionally introduce major new technologies only when launching a new platform, rather than mid-cycle.

The result was a cascading delay across portions of the industry. OEMs that had planned to deploy Level 3 or advanced Level 2+ capabilities found themselves postponing releases due to EV-related complexity and resource constraints. What appeared externally as a slowdown in autonomy was, in many cases, the result of OEMs navigating the simultaneous pressures of electrification, supply chain instability, and the long development cycles inherent in safety-critical automotive systems.

This period of overlapping disruptions in autonomy, semiconductors, the pandemic, and the EV transition accelerated the consolidation already underway in the LiDAR and autonomous driving sectors. The companies that survived were those who won substantial OEM programs, able to weather long development cycles, adapt to shifting OEM priorities, and meet the rigorous standards required for mass-market automotive deployment.

The outcome is consistent with how the automotive industry historically behaves. When dealing with complex, safety-critical technologies, OEMs typically prefer a small, highly reliable supply base. They seek partners with deep expertise, long-term financial stability, and the ability to scale global production. As a result, it is entirely expected that only a small number of LiDAR manufacturers and fewer full-stack autonomous driving providers are expected to become long-term suppliers.

Rather than signaling failure, this consolidation reflects the industry's transition into maturity. The technology is advancing, the requirements are clearer, and the companies that have met these requirements are those believed to be well-positioned to lead the next stage of deployment in automotive and potentially across the broader expansion of Physical AI into robotics, infrastructure, industrial automation, and other consumer devices.

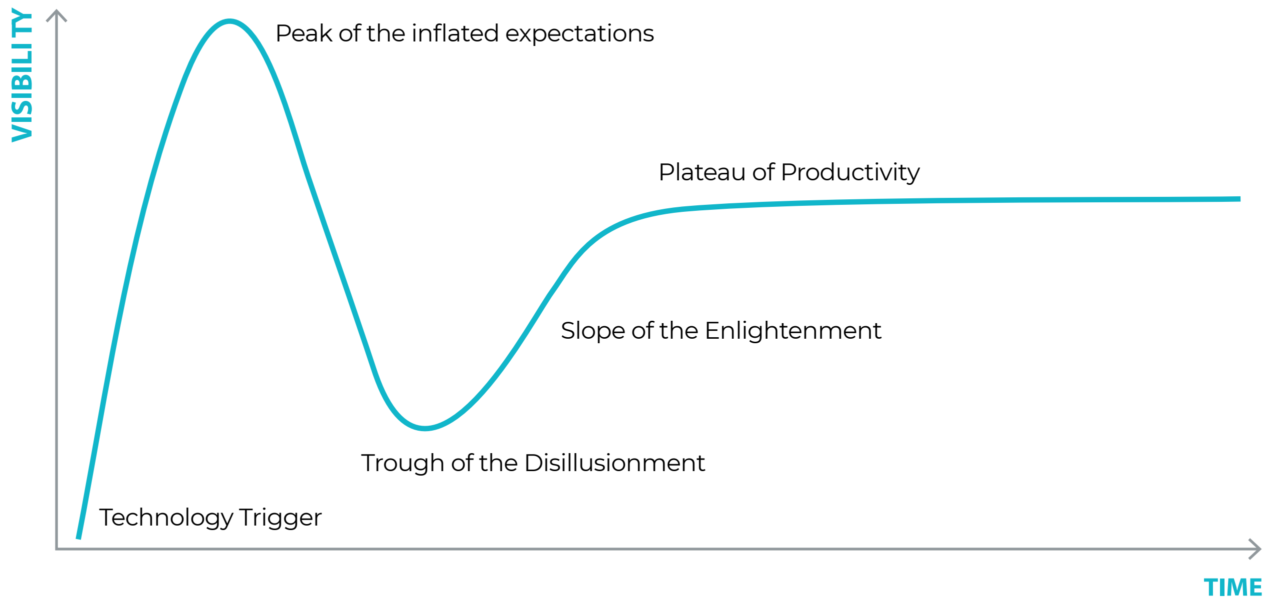


A visualization of Gartner's Hype Cycle, a methodology for describing how new technologies, and the perceptions of them, change as they emerge. Author: Jeremy Kemp. Used under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; A copy of the license is included in the section entitled GNU Free Documentation License.

"Gartner Hype Cycle", "Peak of Inflated Expectations", "Trough of Disillusionment", "Slope of Enlightenment" and "Plateau of Productivity" are registered trademarks of Gartner, Inc. and/or its affiliates.

The Special Purpose Acquisition Company (SPAC) boom of that era marked a pivotal moment for the LiDAR industry. As capital flooded into public markets, several LiDAR manufacturers seized the opportunity to secure substantial funding through de-SPAC transactions. This influx of capital allowed the remaining serious players to strengthen their balance sheets, accelerate product development, and momentarily step out of the cycle's Trough of Disillusionment®. Investors viewed LiDAR as an essential enabler for autonomy, and the public markets briefly rewarded companies promising to lead this transformation.

However, funding alone could not solve the industry's core challenge: industrialization. Many LiDAR companies that went public during the SPAC wave had promising prototypes but lacked the ability to transition from R&D units to fully automotive-grade, cost efficient, mass-manufacturable products. The gap between a lab prototype sensor and a sensor that can meet the automotive industry's lifetime reliability, safety, and cost requirements is vast. As a result, despite having access to capital, the industry underwent significant consolidation, with many SPAC-era LiDAR companies exiting, merging, or substantially scaling back operations. Over time, their once strong market positions eroded, and the industry continued to consolidate.

## 2022 | Phase Two – The Turning Point

Around 2022, the LiDAR market entered a second phase – the beginning of true maturity, characterized by the availability of more sophisticated, production-ready sensors and more disciplined technical and commercial evaluations by OEMs.

During this phase, the most consequential decision came from the Volkswagen Group, the world's second largest automotive OEM by vehicle sales, whose brands include Porsche, Audi, and more. VW selected the second-generation InnovizTwo LiDAR for one of the most significant LiDAR-based automotive programs awarded to date.

This was not merely a commercial win; it was a strategic inflection point for the entire industry. Volkswagen's decision sent a strong signal to the global automotive ecosystem that LiDAR was essential for safe deployment of Level 3 and Level 4 automation.

Volkswagen's platform, anticipated to be the first OEM-manufactured, mass-produced automotive Level 4 autonomous shuttles to deploy at scale, could serve as a reference point for OEMs evaluating the maturity of LiDAR suppliers.

## 2026 | The Third Phase in the Autonomous Vehicle Industry

Today, the market is transitioning into the third phase of the autonomous vehicle industry, a stage that can be described as the end of the Hype Cycle's Slope of Enlightenment®. Early assumptions that autonomy would be achieved quickly, cheaply, and uniformly across all vehicle platforms have given way to a more grounded understanding: true autonomy requires highly advanced, cost-efficient, automotive-grade LiDAR-technology, that only a few companies can deliver.

It is in this environment that Innoviz is introducing its third-generation LiDAR. The new generation is designed to meet the demanding requirements of mass adoption. Compared to prior generations, it is designed to significantly lower cost, provide superior performance, smaller footprint, higher reliability, and seamless integration into vehicle architectures. It represents a foundational enabler for the broader rollout of Level 3 and Level 4 systems across the industry.

In many ways, this marks the beginning of autonomy's true industrial era, with consolidation largely behind us, OEM expectations clarified, and technology reaching automotive maturity. Innoviz believes that the companies that have survived and continued to innovate are well-positioned to participate in the next decade of real-world adoption.

# Competing Technologies

**Many LiDAR solutions struggled to gain traction in automotive because the industry grappled with immature technologies, unsuitable materials, or unproven production methods that could not meet the industry's rigorous standards. Automotive-grade systems demand exceptional reliability, temperature resilience, longevity, cost efficiency, manufacturability, and functional safety. Technologies that shine in a lab often fail under real-world conditions. There was also a period when simply winning an OEM award was seen as enough to stand out. Many of those programs never reached production. In Innoviz's opinion, it's clear that the real benchmark isn't just collecting awards, it's proving you can deliver a true L3/L4 product on the road as Innoviz has done with the BMW 7 Series program. That's where Innoviz stands out; the rest is just noise.**

This challenge was compounded by a relentless push for higher performance, smaller size, lower power, and lower cost. These pressures quickly exposed weaknesses in early-generation LiDAR architectures. As a result, the LiDAR market evolved through a natural process of survival of the fittest, in which only companies with truly strong technology foundations, mature engineering, and deeply industrialized products could move forward. Automotive programs are extremely sensitive to delays or production instability, and the introduction of premature or fragile technologies has repeatedly led to program resets or cancellations.

A clear example of this can be seen in the fate of several highly promoted LiDAR technologies during the first phase of the industry:

## Optical Phased Arrays

Optical Phased Arrays (OPA) sought to leverage electronically steerable optical elements. While elegant in theory, OPA systems suffered from extremely poor optical efficiency, resulting in inadequate range and resolution. The technology never progressed beyond feasibility testing and was ultimately abandoned due to inherent physical and manufacturing constraints.

## 1550nm Fiber Laser Systems

LiDAR systems operating at 1550nm sought to exploit the higher permissible laser power at this wavelength to extend sensing range. However, this approach depended on indium gallium arsenide (InGaAs) detector arrays and high-power fiber lasers derived from the telecommunications industry. These components are intrinsically expensive, power-intensive, and physically large. They present reliability challenges under automotive environmental conditions. Despite these constraints, proponents of 1550nm architectures attracted substantial investment and customer interest, often predicated on expectations of future cost reduction and automotive qualification, rather than demonstrated feasibility at scale. In practice, these assumptions did not materialize. The reliance on non-automotive component ecosystems limited cost reduction, scalability, and robustness, preventing 1550nm systems from achieving automotive-grade maturity or economically viable industrialization.

More importantly, the core assumption behind 1550nm – that stronger laser power is required for long range performance – was ultimately disproven.

Companies like Innoviz demonstrated that 905nm Time of Flight (ToF) LiDAR can achieve and exceed required automotive range thresholds, while staying fully within Class 1 eye safety limits. This eliminated the original advantage of 1550nm and removed any justification for its higher cost, larger size, and greater power consumption.

OEM interest in 1550nm ToF architectures has declined, with announced existing projects in the industry converging around mature, cost-effective, automotive-ready 905nm solutions.

## Frequency Modulated Continuous Wave (FMCW)

Over the past decade, the automotive LiDAR industry has experienced multiple technology waves. Optical Phased Arrays and 1550nm architectures were each introduced with the promise of displacing 905nm ToF systems. More recently, Frequency Modulated Continuous Wave (FMCW) LiDAR has emerged as the latest approach positioned as a potential alternative to conventional ToF architectures.

FMCW LiDAR offers an elegant theoretical advantage: the ability to measure range and instantaneous velocity simultaneously through coherent detection. In principle, this capability enables direct Doppler measurement and enhanced object discrimination. However, translating this theoretical strength into a scalable, automotive-grade product presents significant architectural and manufacturing challenges.

## Architectural Dependencies and Silicon Photonics

Most FMCW automotive designs rely heavily on Silicon Photonics (SiPh) integration. Silicon Photonics is a highly mature and effective platform for data communications, where systems operate at relatively low optical power levels, in tightly controlled single-mode waveguide environments, and under regulated temperature conditions.

Automotive LiDAR imposes fundamentally different requirements.

Systems must:
- Operate at significantly higher optical output power
- Function across a wide automotive temperature range
- Transition from guided-wave environments into free-space propagation
- Maintain long-range performance with high signal-to-noise ratio

When adapted to free-space coherent sensing, Silicon Photonics encounters structural challenges rather than incremental optimization gaps.

## Power and Temperature Constraints

Telecom-oriented SiPh platforms are optimized for milliwatt-level laser operation and temperature-stabilized environments. In contrast, long-range automotive sensing requires substantially higher effective optical power and reliable operation across extreme temperature ranges.

To date, there is no broadly deployed, automotive-qualified, high-power, wide-temperature-range laser solution fully compatible with scalable Silicon Photonics integration. This mismatch introduces performance, thermal, and reliability challenges that directly impact range and system robustness.

## Free-Space Coherent Limitations

In fiber-based telecom systems, coherent detection benefits from well-defined single-mode propagation. Automotive LiDAR, however, operates in free space and must handle complex multipath reflections from diverse environmental surfaces.

In coherent FMCW systems, free-space multipath introduces speckle and phase noise. These effects degrade signal-to-noise ratio and can impair both range accuracy and velocity estimation, particularly in complex driving scenarios with multiple overlapping returns.

These constraints limit achievable range, signal strength, and resolution under real-world automotive conditions.

## System-Level Implications

Beyond photonic integration challenges, FMCW architecture introduces additional system-level complexity.

- Extracting point clouds from FMCW raw data requires multiple Fast Fourier Transforms (FFTs) and advanced signal processing pipelines.
- Resolving range-velocity ambiguities in multi-target environments increases real-time computational demand.
- Higher processing requirements translate into increased power consumption, larger electronic subsystems, and added thermal management burden.

These factors are particularly critical for automotive OEM installations, especially behind-the-windshield configurations, where compactness, energy efficiency, and thermal stability are mandatory.

In addition, practical implementations often require duplication of transmit and receive channels to compensate for optical and coherent detection losses, further increasing system complexity.

## Manufacturing and Cost Structure

Silicon Photonics-based FMCW systems typically depend on:

- Specialized wafer processes
- Hybrid optical-electrical packaging
- Tight alignment tolerances

Achieving the performance required for automotive deployment would likely require substantial foundry-level process investment. Such investments must be justified by high-volume demand and stable roadmaps—conditions that have not yet materialized for automotive FMCW LiDAR.

In contrast, 905nm Time-of-Flight architectures leverage decades of CMOS optimization, established silicon detector ecosystems, mature supply chains, and proven automotive qualification pathways. This results in a significantly more favorable cost structure and scalability profile.

## Industry Maturity and Deployment Readiness

FMCW remains an active area of research and development. However, to date, no FMCW automotive LiDAR platform has demonstrated large-scale, automotive-qualified deployment meeting the combined requirements of long-range, high-resolution, low-power consumption, compact integration, cost efficiency, and automotive-grade reliability.

Innoviz previously investigated Silicon Photonics integrated circuits as part of its exploration of FMCW architectures. During this process, it became evident that achieving automotive-grade performance would depend heavily on future semiconductor process advancements and substantial additional foundry investment. Parallel identification of structural technical constraints further limited feasibility. As a result, development efforts were discontinued in favor of architectures better aligned with automotive scalability and reliability requirements.

### FMCW Takeaways

FMCW LiDAR presents compelling theoretical advantages, particularly in simultaneous range and velocity measurement. However, when evaluated against the full set of automotive-grade requirements—performance, thermal robustness, manufacturability, cost structure, and scalability—Silicon Photonics-based FMCW architectures face significant structural challenges.

Although SiPh is currently the dominant high-volume technology for high-speed communications, its core building blocks do not fully overlap with those required for LiDAR. The missing components remain in their infancy, and in Innoviz's view, industry motivation to invest in their development has significantly declined.

## 905nm Time-of-Flight: The Proven Path to Automotive Success

905nm Time of Flight LiDAR remains the most mature, scalable, and cost effective technology for mass-market automotive deployment. Despite significant investment, alternative LiDAR technologies faced significant challenges scaling to automotive requirements. This was partially due to dependency on immature manufacturing ecosystems or introducing unnecessary complexity into an industry built on proven, stable, and highly optimized processes.

In contrast, 905nm ToF paired with low cost silicon detectors is the architecture that has met full automotive-grade requirements for series production at scale. These components benefit from decades of mass-production, deeply established supply chains, predictable performance, and well documented reliability. These are advantages that 1550nm and silicon photonics-based approaches have not yet demonstrated at equivalent scale in automotive production environments.

So far, 905nm ToF is the LiDAR technology that has reached series production automotive grade Level 3 on public roads. We expect that Level 4 deployment will follow. Meanwhile, a number of programs built around emerging technologies did not reach series production, and in some cases, OEM programs were revised back to Level 2. After numerous stalled efforts industry-wide, it's clear to us that the real proof isn't winning a customer. It's delivering reliably at full automotive volume series production.

Historically, some industry players dismissed 905nm LiDARs as performance-limited, assuming that higher wavelengths would unlock superior capabilities. This misconception fueled exploration into alternative architectures that ultimately proved impractical. Those early assessments overlooked the technological potential embedded in the design of 905nm ToF systems, not simply the raw component specifications. Innoviz recognized this early and invested heavily in proprietary chip design, advanced signal processing, custom detectors, and optics that extracted performance well beyond what many believed possible.

Through this approach, Innoviz demonstrated that 905nm ToF, when combined with world class system engineering and deep industrial expertise, can outperform competing architectures on every critical dimension: range, resolution, reliability, power efficiency, size, cost, and manufacturability. The maturity and scalability of 905nm components made it possible for Innoviz to push the technology far beyond industry expectations and deliver automotive-grade LiDARs that outclass systems based on more exotic and ultimately less practical approaches.

In short, the companies that attempted to disrupt automotive LiDAR with unproven technologies learned the hard way that innovation must be paired with manufacturability and reliability. The industry now recognizes that the path to mass deployment is not through experimental architectures, but through highly engineered, deeply industrialized, proven platforms, a path that aligns directly with Innoviz's technological strategy and product evolution.

## 905nm ToF Competition and China

As it becomes more evident that the technology which is ready for automotive scale is 905nm ToF, there are only a few players left, particularly when it comes to the Western market.

InnovizTwo delivers a combination of performance characteristics purpose-built for the demands of series-production automotive programs. This combination of performance, safety certification, and integration flexibility has supported Innoviz's selection by multiple global automotive OEMs for series production programs, including BMW, Volkswagen, and Daimler Truck.

Chinese LiDAR suppliers incorporated lessons derived from the development efforts of Western counterparts and converged on 905nm T0F architectures, recognizing their relative technical maturity, cost-effectiveness, and suitability for large-scale automotive deployment. The Chinese LiDAR suppliers have benefitted from strong domestic support with several new Chinese OEMs that decided to include a LiDAR in their Level 2 platform. While Level 2 driver assistance platforms do not intrinsically require LiDAR sensors, several emerging OEMs in the Chinese market initially adopted LiDAR to compensate for limited early-stage software maturity.

The inclusion of LiDAR enabled these OEMs to achieve acceptable system performance and reliability thresholds more rapidly, thereby accelerating time-to-market. As software capabilities matured, some manufacturers reassessed the cost-benefit trade off of LiDAR integration. For example, XPeng subsequently removed LiDAR from certain Level 2 platforms once its perception and decision-making stack reached sufficient maturity. This pattern is likely to persist under the intense pricing pressure and sustained losses characterizing the domestic Chinese automotive market. At the same time, as the market focus increasingly shifts toward Level 3 and higher automation, continued and potentially expanding adoption of LiDAR is expected for applications where higher system redundancy and sensing performance are required.

## Geopolitics, Regulations, and Tariffs

When expanding into Western automotive markets, Chinese LiDAR suppliers face escalating geopolitical and regulatory barriers. While several U.S. OEMs evaluated or adopted Chinese LiDARs in early R&D phases, their primary motivation was cost. Chinese suppliers initially offered significantly lower pricing compared to Western alternatives. To navigate US scrutiny, several Chinese LiDAR manufacturers have responded to US regulatory pressure by attempted workarounds, such as RoboSense establishing operations in Singapore via RoboX, a strategy aimed at reducing their exposure to geopolitical restrictions. Despite these efforts, certain Chinese LiDAR manufacturers have been designated a 'Chinese military company' on the US Department of Defense's Section 1260h list, severely limiting their ability to engage with American OEMs and Tier 1s. This regulatory trend signals a broader U.S. government policy direction of treating Chinese LiDAR suppliers as national security concerns, regardless of their stated civilian focus. This is a known risk that the US OEMs will need to manage when the time comes.

At the same time, tariffs have become an additional structural barrier for Chinese LiDAR suppliers. U.S. trade policy imposes significant import duties on Chinese automotive components, further eroding their pricing advantage and increasing total cost of ownership for American OEMs. Combined with U.S. national security concerns, blacklisting actions, and policies such as the proposed SAFE LIDAR Act, tariffs reinforce a growing divide: Chinese LiDARs are becoming increasingly impractical for Western automotive platforms, regardless of technical performance.

Regardless, the cost gap is potentially rapidly closing. As the industry shifts into large-scale Phase 3 adoption and technologies like InnovizThree deliver major reductions in size, power, and manufacturing cost, Western suppliers are now approaching the same price points Chinese vendors target.

The US government restrictions are intensifying. The proposed SAFE LIDAR Act represents a major escalation. It aims to prohibit the use of LiDAR systems developed by companies with ties to the Chinese government or military in American vehicles and critical infrastructure. The legislation is part of what seems like a broader trend in Washington that treats LiDAR technology as potentially having national security implications. Because LiDAR enables the creation of highly accurate 3D maps and can support advanced perception systems, policymakers view it as sensitive if exploited for intelligence gathering or military applications.

Taken together, the blacklisting of some Chinese suppliers, the proposed SAFE LIDAR Act, and other ongoing regulatory actions confirm a clear trajectory toward restricting LiDAR systems developed in China from US automotive platforms. This has the potential to create a structural barrier that could significantly limit Chinese vendors' ability to compete in US Level 3 programs, regardless of their technical progress or commercial scale.

With Physical AI becoming part of our day-to-day lives, LiDARs are potentially going to be used everywhere, around our roads, infrastructure, schools and even our homes. As different countries are planning their AI strategy going forward, national security, trust, and strategic alignment should be established with this critical infrastructure.

Though the market opportunity is very large, the competitive landscape has narrowed with a clear split between Western and Chinese suppliers. We believe that Innoviz is a supplier with the highest number of programs that are expected to launch in the coming years, and we have confidence in our new third-generation product which is expected to increase technological advantages and, therefore, likely to gain additional awards.

# The Last Squeeze: Cutting Cost, Power, and Size

**The LiDAR industry is now entering its third phase of evolution. In the first phase, LiDAR was primarily an R&D tool used to build early prototypes and to explore the feasibility of autonomous driving. The second phase focused on industrialization: companies worked to convert prototype systems into automotive-grade products capable of series production. Despite the large number of entrants, to our knowledge only two Western OEMs succeeded in launching a true Level 3 passenger vehicle with an automotive-grade LiDAR. Innoviz is one of the two suppliers that reached this milestone.**

We are now entering the third phase where the lessons, both technical and operational, learned in phase two are being incorporated into new programs by a broader set of OEMs. These programs are designed for large-scale production, not small-volume pilots. As a result, expectations are significantly higher and tolerance for risk is dramatically lower.

Requirements tighten with each phase. Technologies that once appeared viable can no longer meet the shrinking budgets, aggressive cost targets, and higher performance benchmarks required for mass-market Level 3 and Level 4 systems. This "last squeeze" pushes every aspect of LiDAR technology to its limits. It leaves room only for companies with mature platforms, deep experience, strong supply chains, and well-established industrial infrastructure.

From our experience, reaching this stage would require a stable and proven baseline, including:

- Lessons learned from real-world driving
- Validated performance across millions of road miles
- Fully baked automotive design considerations
- Established production lines and test infrastructure
- Quality systems built for automotive scale and consistency

In this third phase, LiDAR will be deployed at global scale, and there is zero room for design errors, manufacturing surprises, or technology immaturity, which were the same issues that plagued many second phase programs.

At the same time, Innoviz is preparing to start producing InnovizTwo, its second-generation LiDAR, initially designed for programs with Volkswagen Group, Mobileye and Daimler Truck. Building on this momentum, Innoviz is advancing forward with InnovizThree, its third-generation platform, designed from the ground-up to meet the new OEM requirements for behind-the-windshield integration, with dramatic improvements:

- 60% smaller form factor
- 40% lower power consumption
- 35% cost reduction

InnovizThree reflects the clear direction of phase three: LiDARs must become smaller, cheaper, more efficient, and more seamlessly integrated without compromising performance or automotive robustness.

We believe this phase will likely determine the long-term winners in the industry. In Innoviz's view, only technologies and suppliers capable of meeting this final squeeze would be positioned to support the automotive Level 3 and Level 4 programs in the years ahead.

# The LiDAR Industry's Dirty Little Secret

**One of the least discussed challenges in the LiDAR industry – and one that many early manufacturers chose to overlook – is the simple reality of a dirty LiDAR window. While the industry raced to close gaps in performance, cost, size, and power – gaps many believed were impossible to overcome – most companies ignored a fundamental edge case: real-world contamination.**

"Dirt" in this context includes anything that can obstruct the LiDAR's aperture: road grime, mud, rain residue, salt spray, bugs, bird droppings, and even stone chippings. These are everyday occurrences for vehicles, yet they pose one of the most critical threats to sensor availability and thus to functional safety.

LiDAR's value proposition is its role as a redundant sensor to cameras. Functional safety standards such as ISO 26262 define redundancy very clearly: two redundant sensors must have zero correlation in their failure modes. LiDAR companies often highlight the weaknesses of cameras – sun glare, low light, and dynamic range limits – but fail to acknowledge a basic truth: if a vehicle drives through a puddle of mud, it is optimistic to assume only one sensor will be obscured.

In real-world driving, both sensors could be affected simultaneously. Therefore, at least one sensing modality must be inherently resilient to window blockage, capable of maintaining minimum functional performance even when the optical surface is partially obstructed. This is particularly critical for robotaxis, trucks, and Level 4/Level 5 vehicles, which cannot afford sensor degradation below safety thresholds simply because of debris on the lens.

LiDAR companies that failed to consider this requirement may have discovered, possibly too late, that their technology is not entirely viable for autonomous driving. OEMs now have a much deeper understanding of which specifications are necessary to meet functional safety goals under harsh environmental conditions. Dirt resilience is no longer an afterthought; it is a core design requirement.

Among the most underestimated KPIs in automotive LiDAR is "availability". This does not refer to production capacity or hardware uptime; it refers to the sensor's ability to maintain its intended performance in real-world conditions. Availability accounts for:

- Weather (rain, snow, fog)
- Rapid temperature changes
- Mud, dust, and grime
- Scratches or micro abrasions
- Debris accumulation
- Any contamination affecting the optical path

During the first two phases of LiDAR adoption, the industry had not fully internalized the importance of these edge cases. Now, from what we see in the third phase of true scaled production, with driverless Level 4 programs hitting the road, availability is no longer optional. Sensors must deliver close to 100% availability, and OEMs increasingly treat this as a mandatory requirement for any supplier seeking Level 3/Level 4/Level 5 integration.

Innoviz's patented LiDAR design includes blockage resilience technology which, to the best of our knowledge, represents a significant advancement in the availability and environmental resilience required for Level 4 functional safety applications. We believe this capability is a key differentiator that positions Innoviz's LiDAR among the most capable solutions for driverless deployments. This capability allows Innoviz to deliver reliable performance ensuring that autonomous vehicles can operate safely, even under the harsh conditions that define real-world driving.

*InnovizTwo Road Resilience and Availability Demonstration*

# Are LiDARs Too Expensive?

**LiDAR is often perceived as expensive, but the reality is that the technology is simply early in its industrialization cycle. In its early days, LiDAR systems were built almost entirely by hand, using hundreds of discrete lasers, detectors, optics, and electronics. These units had to be individually assembled, calibrated, and tested, which pushed prices into the tens of thousands of dollars, yet they still delivered limited lifetime and modest performance. As the technology matured, the industry began shifting from research prototypes to products designed for real-world automotive use.**

A major factor that kept LiDAR prices high for so long was the constant shift in performance requirements. Early systems like the Velodyne HDL64 were expensive but still far from meeting automotive needs. They offered only 64 vertical lines, short range, poor resolution, and low reliability. Over the last decade, OEMs repeatedly raised expectations for range, resolution, robustness, and frame rate as they learned more about the true requirements for Level 2+, Level 3, and Level 4 systems. This continuous movement of Key Performance Indicators (KPIs) made it nearly impossible for LiDAR suppliers to stabilize designs, scale manufacturing, or invest fully in cost reduction strategies.

Now, with the industry entering its third phase of adoption, OEMs finally have a clearer understanding of what is required, and the technology has reached a point of stability. This opens the door for true industrialization and cost reduction. When you break down a modern LiDAR, particularly one designed with an efficient architecture like Innoviz's, it becomes clear that there is nothing inherently expensive in the bill of materials. A LiDAR can be built using a single low-cost 905nm laser diode, a compact silicon-based detector, and a proprietary ASIC that performs control and signal processing. In many ways, its component structure resembles that of a DVD drive, a device that once sold for under $50 in high volume.

Achieving similarly low price-points for LiDAR will require volume production, process automation, and deeper integration. None of these steps require scientific breakthroughs. They require engineering discipline and operational scale. As design requirements stabilize and production ramps across multiple OEM programs, suppliers can invest in larger ASICs, consolidated electronics, automated calibration, and streamlined manufacturing lines. This is the same cost trajectory that radar followed: thousands of dollars in the early years, then hundreds, and today well under $100 in mass production.

In short, LiDAR is not fundamentally expensive. The industry has moved past the exploratory R&D stage and is now entering a phase where cost, size, and power reduction will accelerate rapidly. The remaining work is mostly engineering and operations, not reinvention. As volumes grow, LiDAR will follow the same cost curve dynamics seen in every other maturing automotive sensing technology.

# Things You Didn't Know About the Five Levels of Autonomous Driving

**The Society of Automotive Engineers (SAE) defines five levels of driving automation. Capabilities increase with each level, but Level 3 is the first point at which a vehicle can truly be called an autonomous vehicle. The driver is permitted to exit the loop and is not responsible while the system is engaged.**

Automotive follows safety principles common to aeronautics and healthcare, where functional safety and redundancy are paramount (e.g., ISO 26262). In practice, this means no single event should be able to cause a safety incident above an extremely low probability threshold. Single points of failure may exist, but OEMs must prove the risk is negligible.

Redundancy is not merely duplicating the same component. Two identical sensors viewing the same scene can share common failure modes (e.g., glare, low light, dirt), making it difficult to prove independent behavior. Consequently, diverse redundancy, pairing different sensing modalities, is preferred to ensure uncorrelated failures.

## Levels 0-2: Driver Responsibility

Level 0 - The human drives; the vehicle may issue alerts but does not intervene.

Level 1 (Hands on, eyes on) - The human drives; the vehicle supervises and may intervene (e.g., lane keeping, AEB). Liability remains with the driver.

Level 2 (Hands off, eyes on) - The vehicle executes the driving task; the human supervises and must intervene on request. OEMs employ driver monitoring to confirm attentiveness; liability still rests with the driver. Level 2 is sometimes marketed as "supervised Full Self Driving (FSD)," although the combination of "supervised" and "FSD" is inherently contradictory. At this level, OEMs must implement driver monitoring systems to ensure the driver remains attentive and ready to intervene at all times. If the driver fails to supervise and an accident occurs, liability rests entirely with the human, not the system.

At Level 2, functional safety relies on a camera based sensing stack, with the human acting as the redundant sensing and decision making element.

When camera performance is degraded, the system issues a takeover request, and the vehicle effectively drops from Level 2 back to Level 1. In some markets, particularly China, LiDARs are added to Level 2 systems to expand availability across more weather and lighting conditions. However, because the driver remains the fallback, the LiDAR in L2 does not serve a functional safety role. In principle, even a LiDAR that resets every few minutes would not compromise L2 safety because the driver remains responsible.

The long-term future of Level 2 is uncertain. While Level 2 systems are widespread today, it is broadly recognized that passive driver supervision is unreliable and can lead to hazardous situations. Current Level 2 systems typically rely on cameras supplemented by low-cost radar sensors, but radar's limited resolution prevents it from serving as a full spatial redundancy layer for cameras. As LiDAR costs continue to decline, Level 2 systems are likely to be replaced by Level 3 systems or enhanced with LiDAR to improve availability and safety performance.

## Level 3: The First True "Eyes Off" Autonomy

Level 3 is the first stage where the human driver can fully disengage, taking both hands and eyes off the road. Because the driver is removed from the supervisory loop, the system must introduce an entirely new, non-human safety supervisor. Cameras alone cannot fulfill this role: two cameras looking at the same scene share common vulnerabilities - direct sunlight, low light, dirt, spray, and glare - which make them unsuitable as redundant sensors under functional safety standards. Instead, a different sensing modality, typically LiDAR, is required to deliver the same driving capability without failing under the same conditions. This is where the differences between LiDAR technologies truly matter.

## Why Did So Many Early Level 3 Programs Fail?

Several Level 3 programs in recent years collapsed during development because the chosen LiDAR technologies or architectures revealed problems only late in the validation process. Some LiDARs failed to meet promised performance targets; others could not withstand automotive temperature ranges, failed lifetime durability testing, or generated excessive false braking due to optical artifacts. In certain cases, lasers interfered with cameras or the LiDAR was simply not designed for high volume automotive production. These failures highlight the strict maturity and reliability requirements needed when the driver is no longer responsible that OEMs will not compromise for.

## Availability: The Key Performance Indicator That Separates L3 from L2

A critical requirement for Level 3 is availability, the LiDAR's ability to maintain its intended performance under real-world environmental conditions. Any recurring degradation, even a very brief one, compromises functional safety and triggers a driver takeover request. Many LiDARs degrade when confronted with droplets, road spray, low sun, or small blockages, making them unsuitable for Level 3. This Key Performance Indicator is the primary reason why numerous L3 programs either failed or were downgraded to L2.

This is a capability that Innoviz has achieved as a result of its superior architecture.

## From Highway L3 to "L3 Urban"

OEMs began with highways because they are controlled environments: no intersections, traffic lights, or significant numbers of vulnerable road users. Highways require extremely long range and resolution due to high speeds and stopping distances, but they are less complex from a behavioral and software standpoint. Urban environments introduce the opposite challenge. Range demands are lower, but the field of view must greatly expand to eliminate blind spots in dense, chaotic, multidirectional traffic. This has led OEMs to explore a new category: "L3 Urban" (sometimes called Private L4). In this mode, the vehicle drives unsupervised, but a driver remains available to handle rare, highly complex scenarios, such as reversing out of a tight spot or navigating a constrained parking garage. These situations represent only a tiny fraction of driving time but a disproportionate share of L4 complexity.

Unlike highway Level 3, which can be supported with a single high performance front LiDAR, Level 3 Urban requires a multi-LiDAR architecture to ensure redundancy in all directions. The number and performance specifications depend on the vehicle's feature set and how much the OEM chooses to rely on occasional driver intervention. As LiDAR technology becomes more compact, more reliable, and more cost-efficient, Level 3 Urban becomes the natural next step in expanding autonomy beyond highways.

## Level 4 (No Hands, No Eyes): Full Autonomy Within a Defined Domain

Level 4 autonomy removes the human from the equation entirely. The vehicle may not have a steering wheel, and no driver handover is available in the event of a sensor degradation. As a result, platform availability and LiDAR availability become the most critical KPIs. Unlike Level 3 where a driver can temporarily step in, Level 4 systems must maintain redundancy and situational awareness at all times, without exception.

One of the toughest challenges for Level 4 systems is ensuring that the LiDAR can genuinely act as a redundant sensing modality to cameras. This is not trivial. If a vehicle passes through mud or road spray, it is difficult to assume the camera and the LiDAR will not be obscured simultaneously. This scenario, common in real-world driving, is where most LiDARs fail. A robotaxi or autonomous truck cannot be allowed to lose both primary sensing modalities at the same moment, with no human available to take control. Placing a front LiDAR behind the windshield can help, but Level 4 also requires multiple LiDARs distributed around the vehicle to provide 360 degree coverage and continuous blockage monitoring where wipers and cleaning systems are less naturally used.

## Why Innoviz Succeeds in Level 4 Where Others Fail

We believe that Innoviz has become a leading supplier for Level 4 programs largely because its LiDAR architecture solves challenges that consistently defeat other technologies, most critically, blockage resilience, availability, and industrial maturity. Innoviz's patented blockage resilience technology enables the LiDAR to continue operating even when the window is heavily contaminated with mud, road spray, dust, cracks, bugs, or other real-world debris. This capability is not simply a performance enhancement; it is a mandatory requirement for any Level 4 vehicle that must operate with no driver available for handover. Many LiDAR architectures face significant challenges maintaining performance under contamination conditions. In our experience, blockage resilience is a capability that is required for Level 4 applications.

In addition to its technical advantages, Innoviz also benefits from a favorable geopolitical position. Innoviz's supply chain, corporate structure, and manufacturing footprint allow it to serve customers globally. In an increasingly complex regulatory environment, certain competing suppliers, particularly those with operations or ownership ties to jurisdictions subject to growing Western scrutiny, may face evolving restrictions on their ability to supply customers in the United States, Europe, and other markets. From our experience, Level 4 customers, robotaxi providers, autonomous trucking platforms, and mobility operators, regard global deployment as a core strategic requirement. We believe Innoviz's ability to supply into all major markets without geopolitical constraints is a meaningful advantage in serving global Level 4 programs, where supply chain provenance and long-term regulatory certainty are key procurement considerations.

Equally important is the maturity of Innoviz's LiDARs. Level 4 programs operate on shorter development cycles than Level 3 because their software and operations teams need to lock the sensor suite early and begin collecting large-scale datasets. These programs require a frozen LiDAR design, one that provides stable, repeatable data for training, simulation, and validation, without risk of component or architecture changes mid program. Unlike Level 3, where OEMs may still source LiDAR during the B Sample (Design Validation) phase, Level 4 programs typically require C Sample or higher (Design Freeze/Production Validation) before integration. Innoviz's readiness at this level of maturity, industrialized electronics, validated optics, frozen ASIC designs, and stable long-term supply is a key reason why Level 4 developers choose Innoviz.

Together, these advantages, technical resilience, geopolitical neutrality, and product maturity aligned with Level 4 sourcing timelines, position Innoviz as one of the very few LiDAR suppliers capable of meeting the stringent and immediate needs of the Level 4 market.

## Level 4 vs. Level 5: The Role of Domain Definition

Level 4 differs from Level 5 in that the vehicle operates within a defined operational design domain (ODD). These ODDs can include pre planned bus routes, autonomous trucks on controlled highway corridors, airport or hospital shuttles, and robotaxis deployed within specific mapped zones. The operating domain may expand over time, but the vehicle is only responsible for full autonomy within that defined scope.

# Entering Phase 3 of LiDAR Adoption

**The LiDAR industry is now transitioning into its third phase of adoption, marking the end of the Hype Cycle®'s Slope of Enlightenment and the beginning of scaled, commercially meaningful deployment. More OEMs are launching new autonomy programs, and the conditions across the global auto industry appear to have aligned to create a wave of fresh platform kickoffs.**

Several factors are driving this shift:

1. Phase 2 is complete. The early adopters of automotive LiDAR, most notably the leading German OEMs and several major Chinese manufacturers, have successfully launched their first LiDAR enabled production vehicles. The key difference between regions is that German OEMs adopted Level 3 LiDAR, while China deployed LiDAR primarily for Level 2.

2. These early adopters validated LiDAR in mass production, proved its value, and integrated it into their long term ADAS and autonomy roadmaps. Having completed their first generation of deployments, these OEMs are now shifting toward next generation platforms focused on:
   - Higher performance
   - Lower cost
   - Tighter software and hardware integration
   - New architectures supporting higher levels of automation

3. Their head start gives them a competitive advantage as the broader market enters Phase 3 and prepares for scalable LiDAR-enabled production vehicles.

4. China has officially approved L3 and L4 pilot programs. China recently authorized Level 3 and Level 4 Intelligent Connected Vehicle (ICV) pilot programs-a major milestone for the country's autonomous driving roadmap. Nine consortia, including BYD, GAC, Changan, FAW, SAIC Motor, BAIC, Yutong, SAIC Hongyan, and NIO, have been selected. The program includes five stages:

| Stage 1 | Stage 2 | Stage 3 | Stage 4 | Stage 5 |
|---|---|---|---|---|
| Pilot application | Product access pilot | Road traffic pilot | Suspension and exit | Evaluation and adjustment |

The first round of company selection is complete. Next, participating OEMs will submit product access applications to the Ministry of Industry and Information Technology (MIIT) as they prepare for broader public road deployments. This national endorsement provides massive tailwinds for scalable Level 3 and Level 4 LiDAR adoption.

5. EV disruption delayed autonomy, but that delay is ending. Many OEMs originally designed their next-generation platforms to combine EV architecture with autonomous capability. However, EV development introduced significant engineering challenges: high voltage redesign, thermal management restructuring, and major electrical architecture changes. As a result, LiDAR integration became dependent on the completion of EV related milestones, delaying autonomy roadmaps.

   Now that most OEMs have stabilized their EV platforms and finalized key architectural decisions (cooling, compute placement, and wiring topology), the backlog of delayed autonomy programs is finally moving forward. New platforms are being kicked off with autonomy once again treated as a core component.

6. The rise of Software-Defined Vehicles (SDVs). The industry is placing significantly greater emphasis on SDVs, where the vehicle's core functionality is increasingly governed by centralized compute, over-the-air (OTA) software updates, and a service-oriented end-to-end architecture. Within this paradigm, autonomous driving capabilities have become a major differentiator, driving OEMs to invest heavily in high-performance compute platforms, real-time operating systems, and scalable perception stacks. As SDV architecture matures, autonomous driving functions are no longer treated as add-on modules but as integral software workloads that shape the overall platform design, data pipeline requirements, and lifecycle update strategy.

7. Behind-the-Windshield LiDAR: the integration OEMs have been waiting for. OEMs have long waited for behind-the-windshield LiDAR integration, which they see as the ideal solution. Previous roof-mounted and grille-mounted placement options introduced significant design, performance, and durability compromises. Roof-mounted LiDAR disrupted vehicle styling, added aerodynamic drag, and was never acceptable for high-volume consumer models. Grille-mounted LiDAR solved the design issue but introduced technical limitations: the low height reduced long-range detection performance, and the sensor was highly exposed to dirt, road spray, and frequent occlusion, especially in dense traffic where other vehicles easily blocked its view.

   Behind-the-windshield placement finally resolves these issues. It preserves the vehicle's exterior design, provides an elevated and stable sensing position, and uses existing wiper and defrost systems to keep the optical path clear. This combination of protection, performance, and design compatibility is what OEMs need to confidently scale LiDAR into mainstream platforms.

Designing a LiDAR that can operate behind the windshield introduces a set of new engineering challenges that many existing LiDAR technologies will not be able to handle:

- The sensor must be significantly smaller to fit within the constrained packaging space of the cabin, often near the rearview mirror where depth and airflow are limited.
- It must be quiet in order to not create unpleasant noise close to the front passengers.
- It must maintain high performance despite optical losses from windshield attenuation, variations in glass composition, and internal back reflections that can corrupt the return signal.
- It must operate at very low power because the thermal environment behind the windshield is extremely demanding: direct solar load, limited heat sinking, and minimal airflow make heat dissipation far more difficult than in exterior mounted configurations.

These combined constraints required a fundamentally different LiDAR architecture, one that only recently became feasible. The windshield integration problem will be another screening factor for the automotive LiDAR market which only a few will manage to solve.

InnovizThree was designed for this specific purpose and is expected to allow Innoviz to continue its automotive leadership with new awards. It also has to maintain high performance despite optical losses from windshield attenuation, variations in glass composition, and internal back reflections that can corrupt the return signal.

## Vision-Only Architectures and Multi-Modal Perception

The debate around LiDAR often centers on the decision by certain manufacturers to pursue vision-only architectures. Over the past decade, camera-based systems have improved significantly through advances in neural networks, data scaling, and compute power. However, camera systems fundamentally rely on two-dimensional projections of a three-dimensional environment. Depth, distance, and occlusion must be inferred through software models rather than directly measured.

In controlled conditions, inference can approximate geometry effectively. In safety-critical environments, performance must remain robust under edge cases, including low light, glare, spray, dirt, partial occlusion, and extreme weather. These conditions can degrade visual contrast and introduce ambiguity into image-based depth estimation.

Multi-modal perception architectures reduce correlated failure modes by combining sensors with fundamentally different physical properties. LiDAR directly measures distance independent of ambient illumination, enabling deterministic three-dimensional geometry even when visual texture is limited.

Publicly available safety performance data across different autonomous systems demonstrates material variation in disengagement rates and incident statistics between architectures. While many factors contribute to safety outcomes, including software maturity, operational domain, and system design, perception modality plays a structural role in robustness.

Vision-only systems rely on layered inference. Multi-modal systems incorporate direct geometric measurement. In Physical AI, applications where failure has real-world consequences, reducing compounded uncertainty is a core design principle.

As compute continues to scale, the limiting factor in autonomous performance increasingly shifts from reasoning capacity to sensor reliability and environmental resilience. The long-term trajectory of autonomous systems is likely to favor architectures that combine high-performance compute with deterministic three-dimensional perception.

## Innoviz's Transformation from Automotive LiDAR Supplier to Physical AI Enabler

Innoviz was built for precisely this moment. On the tailwind of Automotive Driving developments and on the splash of AI demand for everything, there is a growing demand for three-dimensional (3D) sensing for many industrial applications in the short term and consumer applications in the long term.

Innoviz integrates hardware and perception software into a coherent perception platform based on its high resolution, long-range, deterministic, automotive-grade, scalable solutions. The platform turns physical reality into structured, machine-readable and reliable intelligence. Innoviz enables machines to understand reality; it is not simply selling sensors.

In machine vision, vision-based solutions are converted to 3D-like images to allow the machine vision to have a spatial understanding of its surroundings. These 3D conversions require high amounts of compute power but are inaccurate and easily mislead in various scenarios. In Innoviz's opinion, LiDARs provide the most effective path forward to allow machines to have a low latency and high bandwidth capture of the world in 3D.

Instantaneous capturing of a scene, with an accurate and detailed 3D digitalization of the real-world, enables several new ways to analyze, interact, and provide solutions in many fields.

Physical AI is now being embedded in several sectors, such as perimeter security, airports, crowd monitoring, parking lots, smart intersections, logistics, and smart infrastructure, robots and humanoids, agriculture, mining, piers, constructing sites, production sites, and basically any space where complex activities are required to be monitored or performed. Physical AI is where we digitalize the real world and create a live World Model. According to NVIDIA, those neutral networks should understand the dynamics of the real world, including physics and spatial properties. They can use input data, including text, image, video, and movement, to generate videos that simulate realistic physical environments.

Physical AI developers use World Models to generate custom synthetic data or downstream AI models for training robots and autonomous vehicles.

NVIDIA has developed the Cosmos platform, which includes a suite of world foundation models (WFMs) tailored for physical AI applications. The platform aims to accelerate the development of AI systems by providing tools that generate high-quality synthetic data, which is essential for training and evaluating AI models.

The Cosmos platform features three main types of models:

### Predict Models

These models simulate and predict future states of the world in video format, enabling realistic scenario generation.

### Transfer Models

They transform 3D simulations into photorealistic videos, enhancing the visual fidelity of generated content.

### Reason Models

These models enable AI systems to reason about scenes and make decisions based on physical understanding, facilitating complex task execution.

The Cosmos platform opens the door for World Models that are now developing future real-world applications. Platforms such as NVIDIA Cosmos, when combined with the use of Innoviz's high-performance LiDARs, have the potential to enable live digital twins of real-world environments, a capability that Innoviz's automotive-grade sensors are well-suited to support. Digitalization of the real world at the speed of light will unlock unlimited AI capabilities which were only viable to simulate or offline reconstructed modules which do not allow any real-time analysis or reaction.

3D sensors are going to be ubiquitous in the forms of cars, robots, and infrastructure. This data is going to open new capabilities with great amount of responsibility.

# Conclusion

**To our knowledge, Innoviz is a leading force in autonomous driving, holding the highest number of series production LiDAR awards in Western markets and approaching large-scale deployment of Level 3 and Level 4 systems. With its first Level 3 product already in production vehicles, the company has entered the next phase of LiDAR adoption, which is characterized by volume manufacturing and expanding production programs.**

InnovizTwo delivers a step-change in both performance and economics, achieving a 70% cost reduction and a 30x performance improvement over the first-generation InnovizOne LiDAR. The InnovizTwo platform is ramping its first series production program with Volkswagen's Level 4 ID Buzz. It is followed by multiple Mobileye Drive programs which are planned to begin in 2027 for customers, including Holon and Benteler. Each vehicle includes nine LiDARs. Additional volume production for InnovizTwo is expected from Daimler Truck shortly thereafter, alongside several other Level 4 programs on similar timelines. We believe that InnovizTwo continues to be significant in the Level 4 market due to its production readiness, superior performance, durability, and great availability.

The recently introduced InnovizThree represents a major breakthrough for Level 3 adoption, combining higher performance, lower cost, and true behind-the-windshield integration. The platform is designed to unlock the next wave of automotive awards, further strengthening Innoviz's leadership while narrowing the field of LiDAR solutions capable of meeting automotive requirements.

Over the past decade, Innoviz has built strong Tier 1 capabilities, enabling direct OEM relationships, high-volume production awards, and meaningful non-recurring engineering (NRE) revenues. In the past two years alone, the company has secured more than $110 million in NRE payment plans. While 2025 revenues were largely driven by NREs and pre-production payments, growth from 2026 onward is expected to increasingly reflect production ramps and expansion into Physical AI applications. NRE revenues are expected to remain stable, while LiDAR product sales are projected to comprise a growing share of total revenue as automotive production is planned to accelerate in 2027 and beyond.

The Physical AI market represents a substantial long-term opportunity, offering shorter sales cycles and significantly higher average selling prices, often more than 10X those of automotive applications. Innoviz expects Physical AI revenues to approach automotive levels by 2030, with the potential to exceed them over time. Automotive, however, remains the benchmark market, setting the highest technical and regulatory standards and driving industry consolidation.

With more than 2x year-over-year revenue growth in 2025, and a disciplined 20% reduction in operating expenses, improving gross margins, and accelerating penetration into the Physical AI markets, Innoviz is evolving from a development-focused company into a scalable, production-led business. As automotive programs transition from pre-production to volume manufacturing and Physical AI revenues expand through higher Average Selling Prices (ASPs) and shorter sales cycles, the company expects meaningful operating leverage to emerge. Supported by a balanced revenue mix, stable NRE contributions, and rapidly growing product revenues, Innoviz is increasingly well-positioned to achieve long-term profitability and advance its mission to become the world's premier large-scale supplier of best-in-class LiDAR solutions for autonomous driving and next-generation Physical AI applications.

INNOVIZ™
TECHNOLOGIES



**Forward-Looking Statements**
The content, information, opinions, and viewpoints contained in this white paper are those of the authors of this white paper. This white paper contains certain forward-looking statements within the meaning of the federal securities laws, including statements regarding the services offered or to be offered by Innoviz, Innoviz's competitive position, market leadership and strategic positioning, the anticipated technological capabilities, performance, reliability, availability and scalability of Innoviz's products, the markets in which Innoviz operates or is expected to operate, future demand for Innoviz's products, expected NRE payments and Innoviz's projected future operational and financial results, including revenue and NREs, market size, market definition, market growth rates, total addressable market estimates, competitive landscape estimates, industry consolidation, evolution and trends, the expected adoption, deployment, commercialization and scaling of LiDAR, autonomous driving and Physical AI technologies, management's current views and expectations and anticipated market opportunities. These forward-looking statements generally are identified by the words 'believe,' 'project,' 'expect(ed),' 'anticipate,' 'estimate(d),' 'intend(ed),' 'design(ed),' 'aiming to', 'planned to', 'strategy,' 'future,' 'opportunity,' 'plan,' 'may,' 'might,' 'should,' 'can,' 'could', 'will,' 'would,' 'will be,' 'will continue,' 'will allow,' 'will likely result,' 'likely to', 'potential' or 'potentially', 'going to be', 'to our knowledge', 'in our view' and similar expressions. Forward-looking statements are predictions, projections and other statements about future events that are based on current expectations and assumptions and, as a result, are subject to risks and uncertainties. "NRE (Non-recurring Engineering)" is booked services that may be ordered from Innoviz usually as part of a program design win and includes, among other things, application engineering, product adaptation services, testing and validation services, standards and qualification work and change requests (usually during the lifetime of a program). NREs may be paid based on milestones over the development phase of the project which may take a few years. Many factors could cause actual future events, and in the case of our forward-looking revenues and NRE bookings, actual orders or actual payments, to differ materially from the forward-looking statements in this white paper including but not limited to: the ability to implement business plans, forecasts, and other expectations; the ability to convert design wins into definitive orders and the magnitude of such orders; the possibility that NRE would be set off against liabilities and indemnities; the ability to identify and realize additional opportunities; the emerging nature of the Physical AI and World Models markets, which may not develop as anticipated or on projected timelines; the possibility that alternative sensing technologies may reduce demand for LiDAR in these markets; potential changes and developments in the highly competitive LiDAR technology and related industries; the ability to maintain and scale initial deployments into long-term commercial relationships; and Innoviz's expectations regarding the impact of the evolving conflict in Israel on its ongoing operations. The foregoing list is not exhaustive. You should carefully consider such risk and the other risks and uncertainties described in Innoviz's annual report on Form 20-F for the year ended December 31, 2025, filed with the U.S. Securities and Exchange Commission (SEC) on March 4, 2026 and in other documents filed by Innoviz from time to time with the SEC. These filings identify and address other important risks and uncertainties that could cause actual events and results to differ materially from those contained in the forward-looking statements. There can be no assurances that Innoviz will enter into definitive agreements, orders or receive payments with respect to the NRE payment plan referenced in this white paper. Forward-looking statements speak only as of the date they are made. Readers are cautioned not to put undue reliance on forward-looking statements, and Innoviz assumes no obligation and does not intend to update or revise these forward-looking statements, whether as a result of new information, future events, or otherwise. Innoviz gives no assurance that it will achieve its expectations.

**Market and Industry Data**
Market sizes referenced herein are derived from third-party research reports and/or Innoviz's own internal estimates. Third-party data involves risks and uncertainties and is subject to change based on various factors. Innoviz has not independently verified this data, and no representation is made as to the accuracy or completeness of such third-party data. Innoviz's internal estimates are based on assumptions we believe to be reasonable, but which are inherently uncertain and unpredictable. Actual market sizes and growth rates may differ materially from the estimates presented. Such information is provided "as is" and is believed to be accurate at the time publication. Innoviz makes no warranty, expressed or implied, as to the completeness or accuracy of the content contained in this white paper. The reader uses this white paper at his or her own risk, and Innoviz shall not be liable for any errors, omissions, or inaccuracies in this white paper, whether arising through negligence, oversight, or otherwise. Reliance on any information appearing in this white paper is strictly at your own risk.

**No Investment Advice / No Offer**
This white paper is provided for informational purposes only. This white paper does not take into account, nor does it provide any legal or investment advice or any opinion regarding the specific investment objectives or financial situation of any person. Nothing contained in this white paper constitutes an offer to sell or the solicitation of an offer to buy any securities. To the fullest extent permitted under applicable laws, Innoviz expressly disclaims any and all liability that may be based on this white paper or errors or omissions thereof.

**FOR MORE UPDATES**