Innoviz and the Rise of Physical AI

Bringing World Models to Life: **Part I**

Omer Keilaf, Innoviz Technologies | $INVZ

February 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 and real time. This transition, often referred to as Physical AI, represents one of the largest and longest-duration technology opportunities of the coming decades.

Unlike Digital AI, Physical AI operates under the laws of physics. 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 world model 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 Systems and NVIDIA recently announced a long-term strategic partnership to create "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.

Models became a commodity, and data is the real difference. 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 have 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 required to enable Physical AI safely, securely, and at global scale.

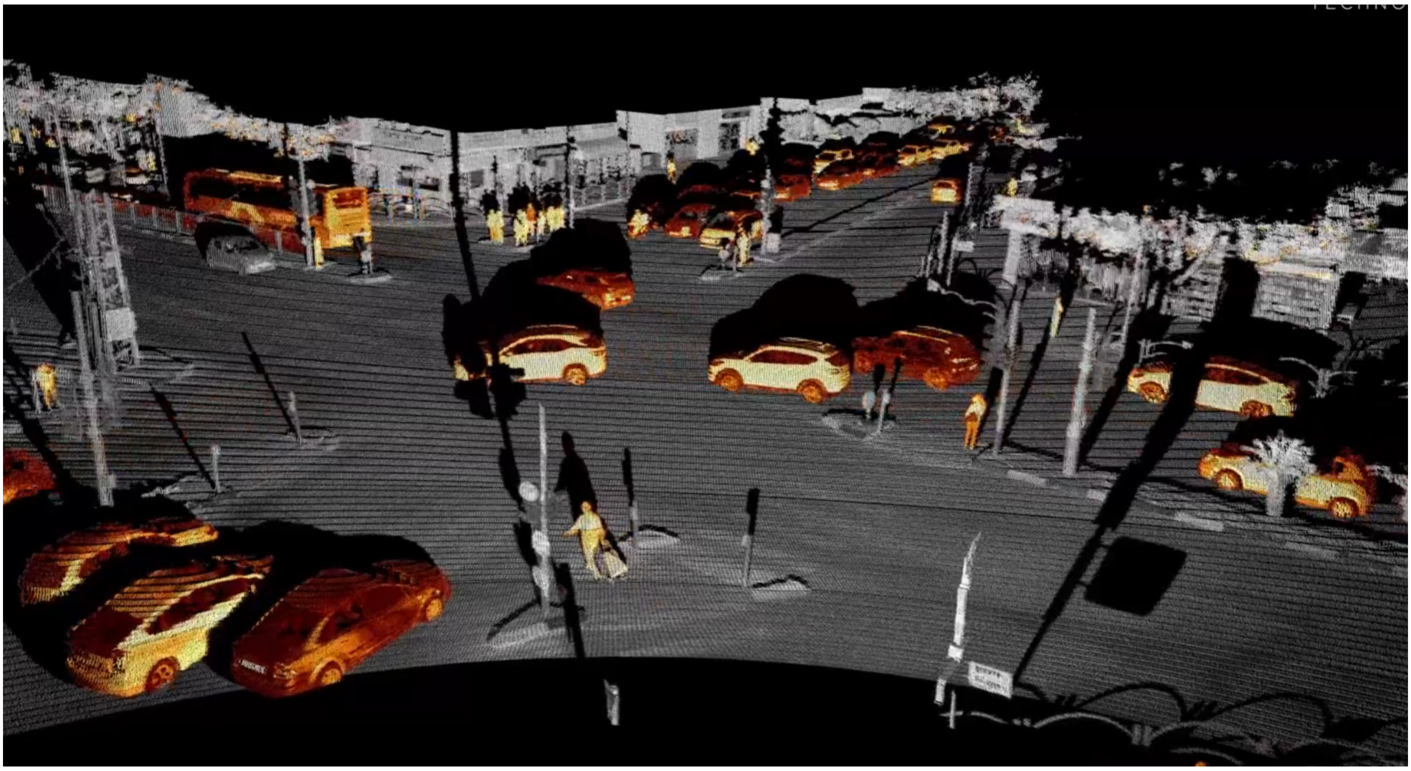

In This White Paper

01 Why Physical AI is the next major computing platform

02 Why perception and LiDAR are specifically the bottleneck

03 Hello World (Models)

04 How autonomous driving is setting the threshold

05 How Innoviz's technology is positioned as the foundational infrastructure

06 The expanding Total Addressable Market (TAM)

07 National security shaping long-term winners

08 Innoviz and the rise of live World Models

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 now becoming deployable at scale, not just demonstrable in controlled environments. 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.

Why Perception is the Bottleneck in Physical AI

Operating Under the Laws of Physics

Digital AI can tolerate ambiguity. Physical AI cannot and 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 and, 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.

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, 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, leveraging NVIDIA's accelerated computing, AI infrastructure, Omniverse technologies, and open AI models together with Dassault's decades of scientific and engineering knowledge 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. This means that 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 realtime, 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, but 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 pointcloud 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. In this way, Innoviz LiDARs act as the high-detail sensory layer that constantly feeds the industrial AI models fresh data, allowing simulated "worlds" to remain synchronized with reality and enabling Physical AI systems to make reliable, situationally-aware decisions.

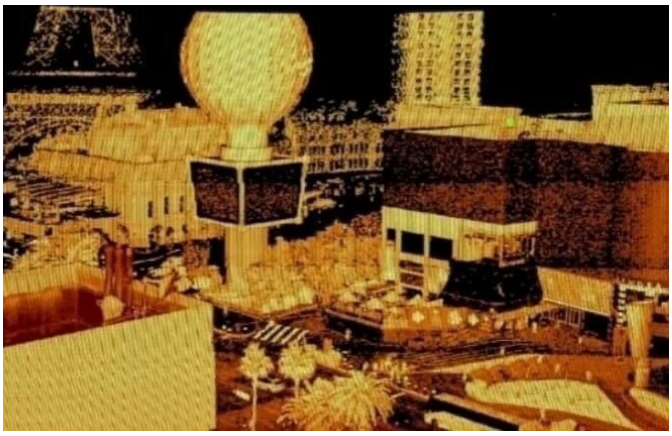

Total Addressable Market (TAM) in the Physical AI and World Model Era

Assessing the total addressable market 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, industrial automation, and drones could grow into a \$0.5–\$1.4 trillion market, with autonomous vehicles 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 comes on top of 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-City	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.

What to Expect in Part II of the White Paper

LiDARs have clearly emerged as a foundational technology for the future of Physical AI and advanced world modeling systems. The upcoming second part of the White Paper will explore how the autonomous driving industry is defining the performance and reliability thresholds for this critical sensor category. It will examine why success in automotive applications remains the ultimate proving ground, how competing sensing technologies compare, and why the LiDAR market is entering a phase of intensified consolidation. The paper will also address persistent misconceptions around LiDAR cost structures and clarify the technical and functional differences between sensors designed for L2, L3, and fully autonomous L4/L5 systems.

Building on this foundation, the paper will outline how the automotive sector is transitioning into its third major phase, where scalable autonomy and robust perception stacks become essential. It will confront the unavoidable Tesla question, analyze the strategic implications of the “last squeeze” in the LiDAR industry, and show how automotive-grade LiDAR is evolving into a cornerstone of broader Physical AI applications.

The summary will highlight Innoviz’s role in this transformation and its position within the rise of Physical AI, and next generation World Models.

INNOVIZ TECHNOLOGIES

FOR MORE UPDATES