Robotics and Physical AI

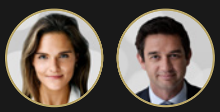

CLARE PLEYDELL-BOUVERIE AND STORM URU

ROBOTICS: THE OPPORTUNITY SET – PHYSICAL AI IS IN ITS ASCENDENCE, THIS IS THE NEXT WAVE OF AI

If 2024 was the year of the agent, 2025 is proving to be the year of the robot. Three technological breakthroughs have converged to make this possible: Firstly, in 2024 we saw a major breakthrough with vision transformers: robots can now finally see. Just as large-language models predict next word in a sentence, these predict the next pixel, enabling full-fidelity simulation – from factory layouts for new cars to software-defined robotics.

Synthetic data generation has been the next unlock: robots have to be trained on an inordinate amount of both real and synthetic data to perform a task at human-level quality. A surprise finding for us in Q4-2024 was that simulations trained solely on synthetic data (now incorporating structured and unstructured inputs) already match – and often exceed – the performance of models trained on real-world data.

Finally, is reinforcement learning – these models require extensive post-training. Whereas reinforcement learning once demanded armies of human data label-workers, it is now performed by a combination of AI and human supervisors – collapsing the iteration loop.

We believe that identifying the winners of the autonomy ecosystem today presents an opportunity analogous to spotting to identifying the winners in the AI infrastructure ecosystem two to three years ago. By 2035, there will be 16 million robots deployed each year, representing a 31% CAGR in annual deployment from today.

Collaborative robots (cobots) are the immediate opportunity – robots that work with people. They are safe to stand up, easy to redeploy, and pay-back periods are now falling below 18 months. Over a one-to-three year horizon, service robots take centre stage. These robots actually perform a task for you, often in spaces that were never designed for automation – think warehouse tugs, hospital runners, retail shelf-scanners. Beyond that come humanoid . The technology is real – legged human scale machines are capable of dancing and even driving a forklift today – but the economics need another design cycle (likely two) before mass adoption.

Pace of robotic mass adoption

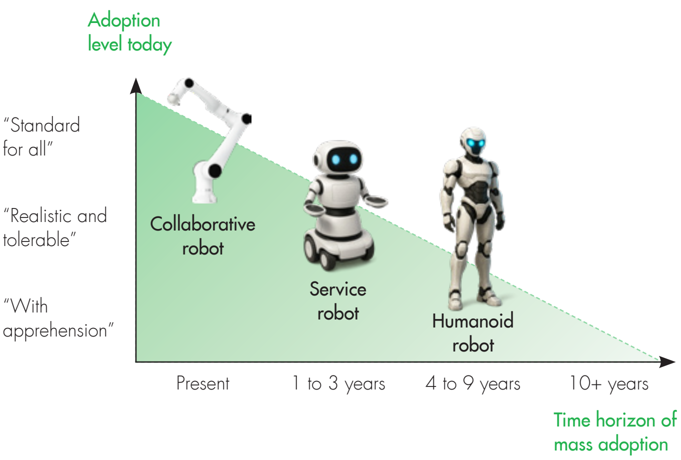

Source: Melexis 2024. All use of company logos, images or trademarks in this document are for reference purposes only.

Projected unit shipments with Optimus ramp: 1 million by 2030, 16 million by 2035

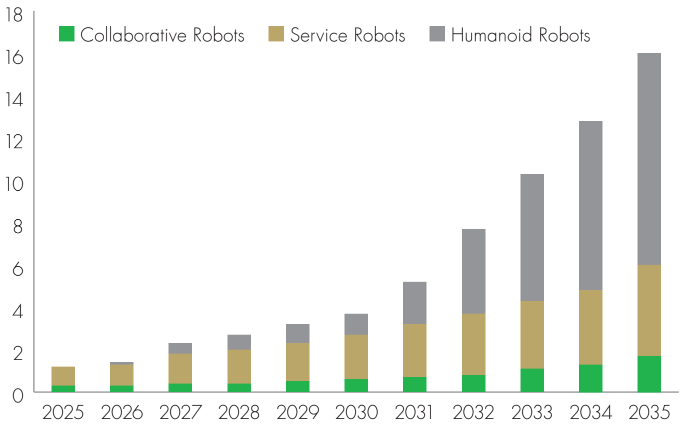

Source: Tesla Q3 25 earnings release and Liontrust estimates, November 2025.

WHY NOW? AUTOMATE OR RISK BECOMING UNCOMPETITIVE – 58% ROI FROM FACTORY AUTOMATION

The economics of automation are already shifting. For example, a BMW 'medium-automation' factory delivers an ROI of 58%, pays back within 1.8 years, and reduces cost per job of over 73%, saving $8.7mn a year. The gap stems from cost structure. A labour-intensive factory carries the full weight of wages, benefits, insurance, turnover and supervisory overhead.

Its automated counterpart replaces those with robot hardware, integration, maintenance, energy, software and a thin layer of human oversight. Integration remains the largest line item, often four to six times the cost of the robots themselves, but it is paid once. After payback, operating leverage is substantial: the automated line runs close to 24 hours a day with minimal downtime, cutting total operating costs by roughly a third at the factory level and delivering recurring annual savings of 25–40%.

A medium-automated plant typically deploys between 400 and 1,000 robots, with BMW's own sites averaging closer to 1,400 units. Each robot becomes a fixed asset sweating capital at near-continuous utilisation. At the cell level, the step-change is sharper. In early configurations, eleven robots performing the work of nine human operators achieve payback in roughly three and a half years.

As density rises, ten robots now displace the work of twenty-three humans, compressing the payback materially while throughput per cell rises and unit labour cost falls. Fanuc's own factories in Japan illustrate the end-state of this curve: robots now build robots, one unit every eighty seconds, running dark and unassisted; Xiaomi's dark factories now run lights-out, 24/7, producing on the order of one smartphone per three seconds on fully automated lines.

This is the new capital efficiency frontier of manufacturing: machines that replicate human dexterity while delivering industrial-grade ROI. Today we sit around Level 2 autonomy, where robots plan their own tasks, map dynamic environments, and traverse open floors without rails or guides. That alone shortens payback by reducing integration time and boosting utilisation. The next step, Level 3 mobile manipulation (currently in pilot stage), will compress returns again: robots will behave more like human employees – travelling to the task, adapting to variance, executing multi-step functions – and become revenue-positive within days or weeks of deployment, because the marginal cost of onboarding is near zero. Beyond that lies Level 4, where tactile feedback and fine manipulation unlock the remaining labour cost: force-dependent, high-skill work that has resisted automation.

Automated vs. labour factory – 58% ROI

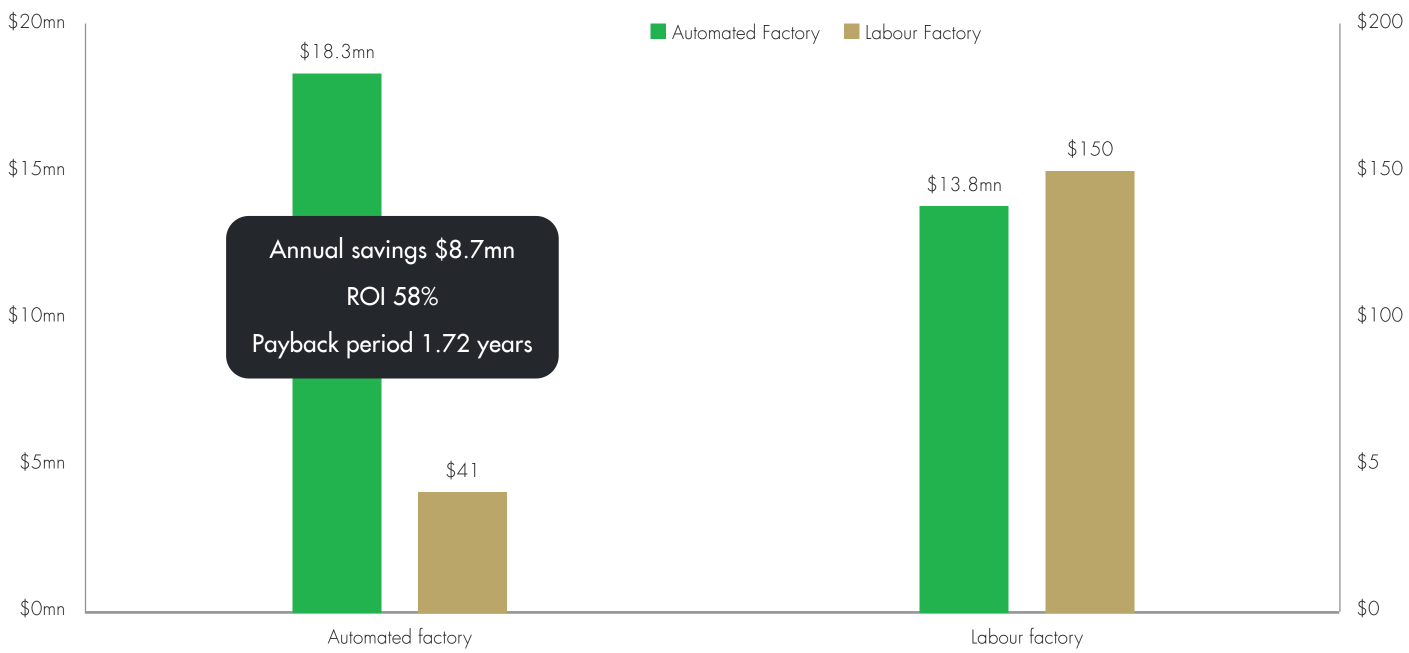

Source: Liontrust, Semianalysis (2025)

AUTONOMOUS DRIVING – THE FIRST SCALED USE CASE OF PHYSICAL AI

Autonomous driving is the first large-scale proof that physical AI creates real economic value. The stack is near-identical to robotics, stripped of dexterity – sensors for perception, high-performance compute for inference, simulation for training, and a control policy that links them – but differs in one crucial respect: the perception layer is optimised for vision rather than touch. Cars are mobile robots in a two-dimensional world; they don't grasp or manipulate, they navigate. That makes autonomy a cleaner problem and a faster commercialisation path.

The ROI is clear. The average human driver costs $0.97–1.00 per mile on long-haul routes once wages, insurance, and downtime are included. Fully autonomous systems already demonstrate pilot-level economics of $0.65–0.85 per mile, saving fleets 25–40% immediately. At scale, robotaxi and freight models push this to $0.30–0.40 per mile as utilisation climbs toward twenty hours a day. That's a 60–70% cost reduction relative to human-operated services – the same scale of margin expansion seen in the shift from manual to robotic manufacturing.

Cumulative autonomous miles driven – June 2025

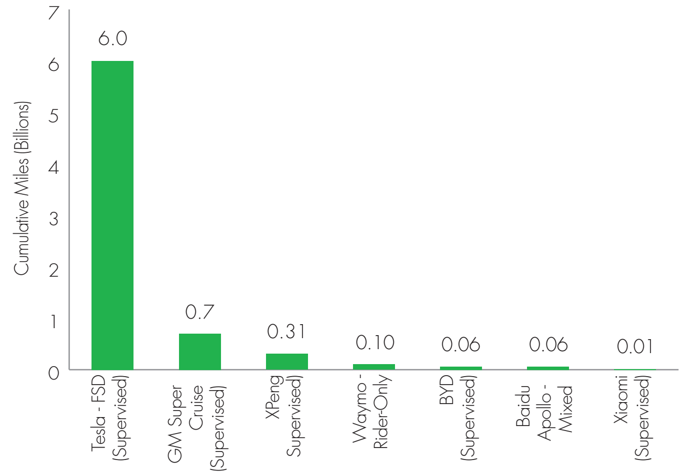

Source: Tesla company reports, Liontrust estimates, October 2025

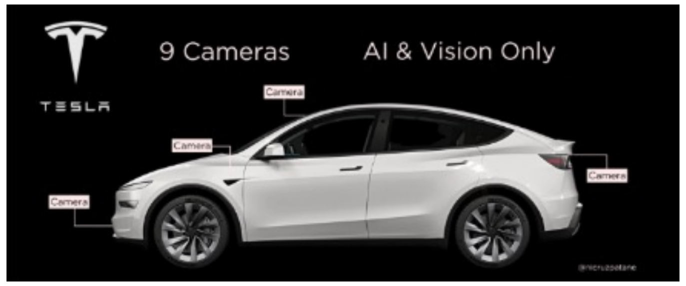

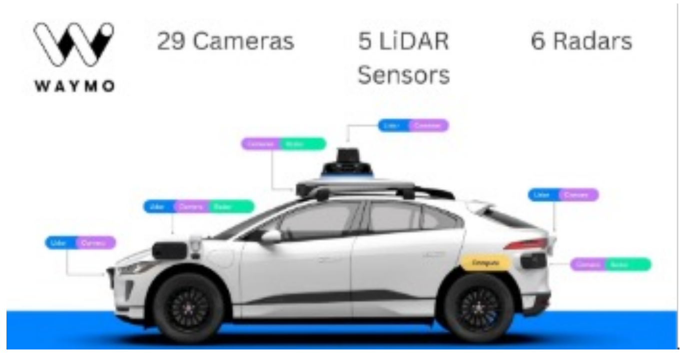

Vision-only approaches are the superior economic model. Hardware simplicity drives scale – and therefore margin. Camera-based systems such as Tesla's cost just $300–400 per vehicle versus c.$12,000 for Waymo's 5th gen LiDAR/Radar heavy stack – one reason why self-driving Tesla can be built for roughly $25,000–30,000 versus $150,000–200,000 for a Waymo-class robotaxi. Vision sensors capture richer data – colour, texture, and motion cues – at a fraction of the cost, and every additional mile driven compounds the training set. Despite Waymo's first-mover advantage, Tesla's FSD has already driven more than 45 times as many autonomous miles and is widening the gap exponentially each month. Tesla's 8.5-million-strong fleet, compared with Waymo's 2,000+ robo-taxis, highlights the importance of scale over first-mover advantage. The more data, the safer and cheaper the system becomes – a feedback loop that favours vision-first architectures.

Profit pools are consolidating at the top of the stack. Software and data flywheels dominate value capture. Tesla's end-to-end neural network for Full Self-Driving, trained on billions of miles of real-world footage, and Aurora Innovation's vertically integrated trucking platform both demonstrate how data density compounds competitive advantage. Once trained and deployed, autonomy software delivers 70–90% incremental gross margins, turning each additional vehicle into a royalty stream.

At the base of the stack, the winners are the lowest-cost component suppliers hoovering up share as autonomy scales. Sensor content per car is rising even as per-unit prices fall. Hesai, China's LiDAR leader, now ships high-performance hybrid solid-state sensors for under $400, while U.S. peers still target $1,000–1,500 per unit at volume. That cost delta ensures China dominates LiDAR volume and export markets, embedding its optics in the world's autonomy fleets.

HUMANOID ROBOTS – TIMELINES & UNIT ECONOMICS

Today, in 2025, humanoid robots are not yet cheaper than human labour. A factory or warehouse worker in developed markets costs roughly $25–30 per hour fully loaded (including benefits, overtime, and insurance), while an early-generation humanoid robot – priced around $100,000–150,000 per unit – still equates to roughly $36–37 per effective labour hour once utilisation, maintenance, and supervision are factored in.

That gap closes steadily through the decade. As production of humanoids scale toward one million units a year, hardware cost falls to $20,000–25,000 per robot, software and service fees compress, maintenance intervals extend, and human supervision fades from one-to-one to one-to-many. Energy consumption per robot-hour is already below $0.50, declining further as actuators, drivetrains, and control loops become more efficient.

The kingmaker remains utilisation. Moving from 50% uptime in 2025 to roughly 70–80% by 2028–2030 transforms the economics: by 2028, the fully loaded humanoid cost falls to around $22–23 per hour, undercutting the human equivalent in most developed markets. By 2029, improved durability and fleet management push that to roughly $14–15 per hour, and by 2030, with supervision ratios near 1:20, the effective cost approaches $10 per hour.

The steepness of that cost-decline curve ultimately depends on two variables: battery technology and dexterous manipulation. Faster gains in battery density, charge speed, and thermal efficiency drive higher utilisation and lower energy cost; delays there push the crossover right. Similarly, advances in compliant actuation, tactile sensing, and robotic hands determine how quickly supervision ratios collapse. Together, these two technologies control whether the slope of this curve falls sharply or flattens mid-decade.

On these conservative assumptions, the first cost crossover appears around 2028, initially in logistics, assembly, and handling roles where utilisation is highest and downtime lowest. Beyond that point, each incremental design cycle – lighter actuators, lower-cost compute, and cheaper sensors – compounds the margin advantage. By the end of the decade, humanoids will not just substitute for labour; they'll redefine it, delivering a structural break in productivity per unit of labour hour.

Form-factor and first adopters matter. Tesla's Optimus is c. 10kg lighter than UBTech's nearest production-bound humanoid – a small but meaningful design edge. Weight directly translates to lower energy consumption, longer battery life, and faster, safer motion; every 10 kg reduction cuts power draw by roughly 5–8% per operating hour and reduces actuator stress, extending service intervals. In robotics, that's what separates a robot that runs four hours on a charge from one that runs six.

The first scaled adoption will come from automotive manufacturing, where humanoids can work alongside existing industrial robots without retooling the floor. BMW is validating Figure 01 humanoids on its assembly lines, BYD has begun integrating UBTech units in factory workflows, and Tesla is moving Optimus from internal trials to line-side logistics and materials handling. The playbook is consistent: start on night shifts and repetitive handling tasks – moving parts, bolting, stacking – where uptime matters more than finesse, then expand into more complex functions as reliability improves.

Humanoid robot effective labour cost vs. human labour 2025–2030

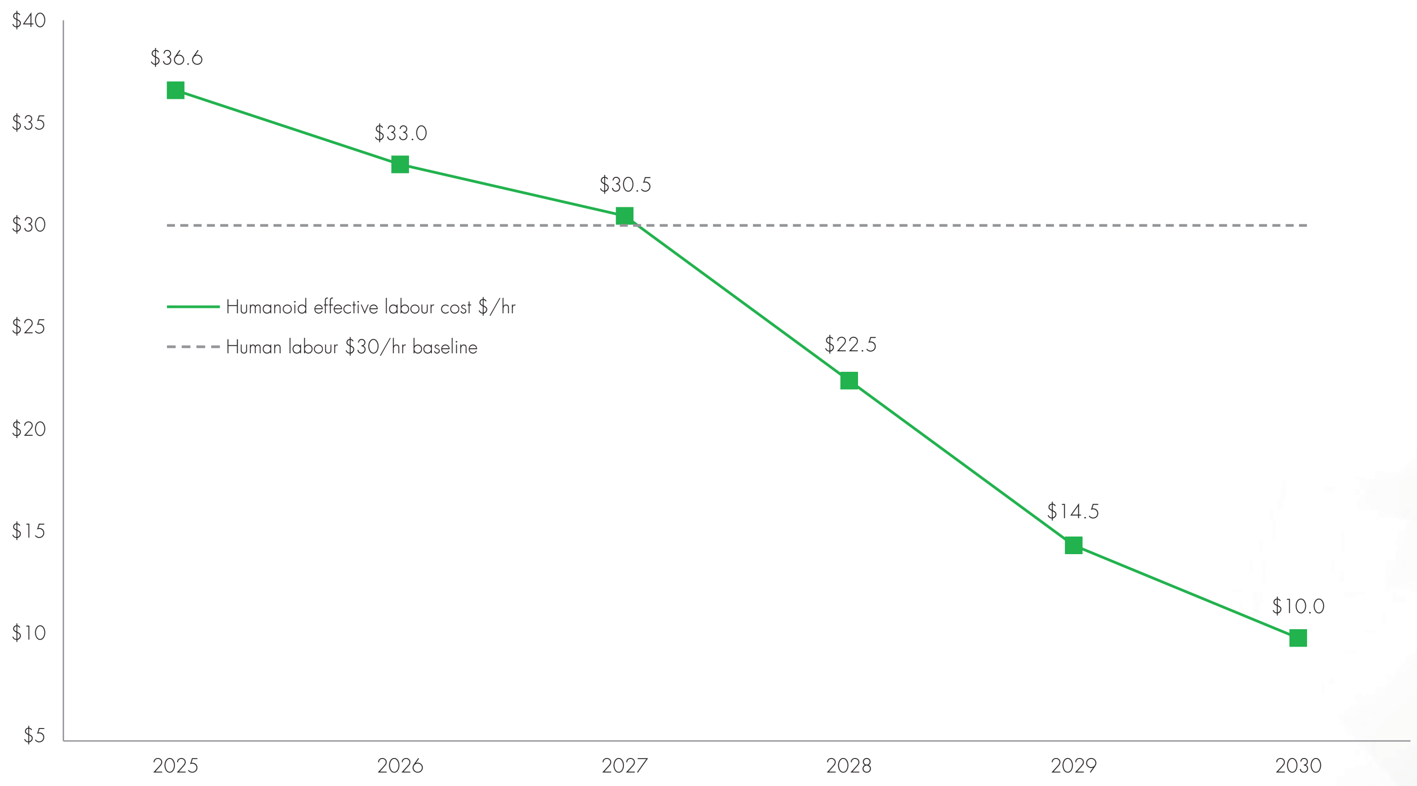

Source: Liontrust (2025)

VALUE CAPTURE – WHO HAS THE RIGHT TO WIN IN THIS NEW CYCLE OF ROBOTICS?

Having defined where value is being created, the next question is who captures it. Japanese incumbents enjoy a powerful first-mover advantage: robotic powerhouses Fanuc, Yaskawa and Kuka control roughly 36% of the world's industrial robot units, with Fanuc alone commanding 16%. However, low-cost Chinese competitors investing aggressively in innovation, and becoming increasingly visible. Incumbent management teams face the innovator's dilemma; those who acknowledge the threat are pivoting, those who do not will be disrupted.

What does pivoting the business look like? Firstly, ramping R&D spend to match Chinese counterparts. Fanuc, for example, has already accelerated its product-launch cadence 1.5x in the past 18 months alone. Management is under no illusion of the threat China poses to the dominant maker of yellow robots found on almost every factory floor. However, the resounding takeaway from our meeting with the entire ecosystem in Japan was that low end robotics is commoditising at an eye-watering pace. Value does remain in high-end, high-touch approaches – in companies like Fanuc and Keyence – the middle will be squeezed.

Robotic industry market share

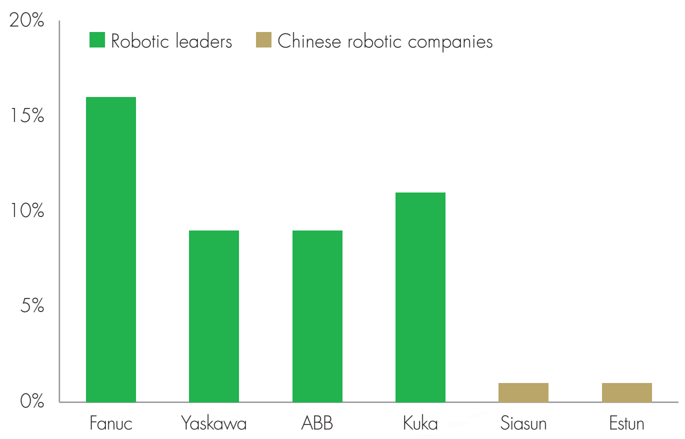

Source: Company reports 2024. Liontrust 2025

Investing in innovation (R&D/Sales)

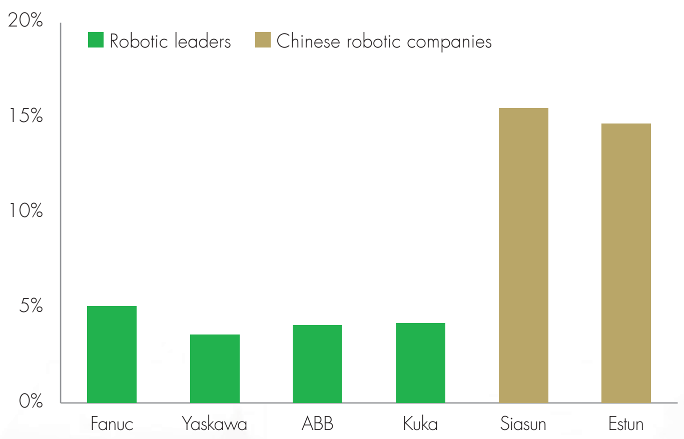

Source: Company reports 2024. Liontrust 2025

THE FULL STACK ADVANTAGE – VALUE CAPTURE ACROSS THE ROBOTICS STACK SHIFTS FROM C.75% HARDWARE TO C.75% SOFTWARE AND COMPUTE

Ultimately value will accrue to different layers of the robotics stack in this next cycle – historically, robotics has been a hardware story, dominated by sales of actuators, transmissions, grippers, sensors, power systems, mechanical structures and chassis, and cabling. Together these robotic hardware components have taken roughly 75% of the value in an average industrial robot (actuation and transmission taking the most of these), with software and silicon capturing just 10% and 15% value respectively. The robots of the future flip this playbook on its head: as automation capabilities increase, software and silicon gains in importance relative to hardware. This is where Nvidia and Tesla, two companies outside the traditional robot league, come in with their full-stack approaches.

Today's robots in deployment – which are capable of conducting repetitive, standardised tasks autonomously (most with a human in the loop) see c.25% value going to software, 17% to compute, and 58% to robotic hardware. Looking forward to the end of the decade, when we are likely to see humanoid robots in mass production and more scaled autonomous driving companies ramping, this hardware chunk shrinks by an order of magnitude to just 27%, with compute taking 28% and software 45%. So for every Optimus robot – priced at c.$20,000 once volume hits 1milion pa – chips comprise c.$5000-6000; the software c.$9000; the hardware $5000-5,500.

Value capture across robotics stack

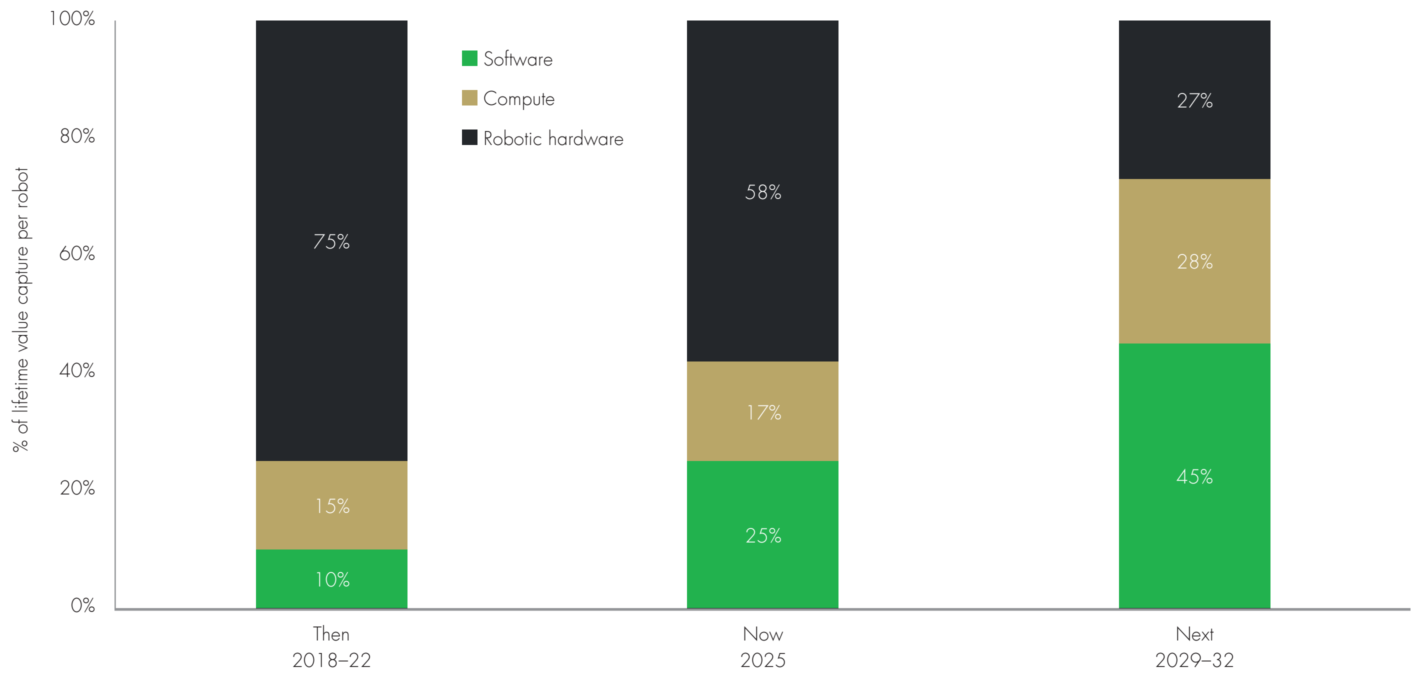

Source: Liontrust Estimates, November 2025

INVESTING BEHIND A NEW SUPPLY CHAIN THE FULL STACK ADVANTAGE – HARDWARE AND SOFTWARE

Nvidia – the compute royalty layer powering physical AI

Nvidia's ambition in physical AI mirrors what it has achieved in AI training and inference for the digital world: to become the platform on which every motion model runs. No company straddles every layer of physical AI like Nvidia. Its robotics strategy revolves around three computers – for training, simulation, and edge inference – forming a self-reinforcing loop that ties the data centre to the robot itself. The key is understanding that every robot, every factory, every data centre becomes part of the same neural network.

- **The AI data centre** – training the brain. At the core sits the Blackwell architecture, built for trillion-parameter model training and reinforcement learning. This is where robot, vehicle, and factory models are born: massive clusters of GPUs train the neural policies and reinforcement-learning agents that control motion. These GB200 NVL72 racks are the new capital equipment of the automation age: every advanced factory will need its own AI brain – and it's already happening. Foxconn is literally using Blackwell to build Blackwell, training robots on Nvidia systems to assemble Nvidia systems.
- **Digital-twin simulation** – teaching the robot. Between cloud and factory lies the simulation layer. Nvidia's RTX PRO Servers render photoreal synthetic environments in Omniverse and Cosmos, generating training data and digital twin synthetic environments. Within this world, Isaac Sim and Isaac Lab run GPU-accelerated physics loops that let robots practise 10 000x faster than real time, allowing robots to rehearse and perfect tasks millions of times before they ever touch the floor. TSMC already saves months in work by designing fabs virtually using Nvidia's digital twin infrastructure. Nvidia's flagship Project GR00T – a foundation model for humanoid motion baked into the Isaac platform – learns directly from these synthetic worlds (essentially a 'GPT-for-motion'.) Whoever ships the best silicon and the best pretrained motion model will capture the software royalty layer across every OEM.
- **Edge inference** – running in the real world, bringing the mind to the machine. Nvidia's third computer brings data-centre-class AI directly to the edge; Jetson AGX Thor powers humanoids and industrial robots; Drive AGX Thor runs autonomous vehicles. Each delivers up to 800 TFLOPS of AI performance in a 250 W envelope – data-centre intelligence shrunk into an embedded form factor, whether that be a humanoid, drone or autonomous vehicle.

These chips close the loop, executing the learned policy on the robot itself and streaming data back to the GB200 cluster for continual learning.

This three-pronged architecture reshapes the semiconductor stack. Training (GB200) fuels simulation (RTX PRO Servers, Omniverse), which perfects the hardware that inference (Jetson, Drive Thor) runs on – a self-reinforcing system where Nvidia captures value at every layer: silicon, software, data, and the digital-twin infrastructure itself.

Physical AI could become the largest accelerated compute market of all for Nvidia – an AI data centre beside every factory, running GPU clusters to train local models and coordinate fleets of robots, and silicon running on the robot or device itself. This three-fold opportunity implies tens of millions of Blackwell-class equivalents worldwide by the 2030s. Competitors – Google, AMD, Intel, Qualcomm – are chasing fragments of this vision, but none control the full stack from cloud to chassis. Nvidia already does.

Tesla – the integrated autonomy flywheel

With its vision-only autonomy stack and neural networks, Tesla commands a formidable data advantage when it comes to autonomous driving – and an equally significant compute advantage when it comes to humanoid robots, given that the same stack underpins both domains.

Tesla's competitive moat in physical AI is data density: 8.46 million Tesla vehicles on the road today and 4.5 billion miles of real-world autonomous driving data feed a single neural network that trains Optimus as well as the entire vehicle fleet. No other competitor can replicate this asset since no other AV player has adopted the same vision-only, neural networking approach. Using a pure neural-networking approach is what enables exponential growth through leaps in hardware and software progress: Tesla's AV capabilities before Elon pivoted to end-to-end networking were improving at a rate of c.2-4x a year; moving from deterministic software to neural networks with FSD 12 led to AV capabilities improving at a c.5-10% rate a month. FSD 14, Tesla's soon to be released AV software will be the biggest leap forward since FSD 12; its new chip A5 will be c.40x better than A4.

The underappreciated factor is that the same compute stack and vision transformers run across both cars and humanoids. Unlike traditional OEMs that buy autonomy as a component, Tesla's stack is end-to-end – from sensor to silicon to policy. This shows up in unit economics: once Tesla's Optimus line hits a run rate of 1 million units a year, Costs of good solds are estimated at $20–25k, implying margins akin to automotive once scaled. Tesla is effectively turning its factory into a live reinforcement-learning lab at industrial scale – robots training robots, with factory telemetry feeding back into the same model family that drives the fleet.

Aurora Innovation – first mover advantage powering commercial autonomy with a capital-light driver-as-a-service model

Aurora is the purist play on commercial autonomy as the first mover in autonomous trucking, with 2026 set to be the commercialisation inflection point as they turn highway autonomy from pilots into product. The model is software-first, capital light and, importantly, end-to-end. The technology is proven – its driverless trucks are on the road, operating commercially today – and now ready to scale.

Aurora Driver is a vertically integrated Level-4 stack – perception (FirstLight LiDAR, imaging radar, cameras), fused world model, prediction, planning and control – governed by a documented Safety Case. Unlike fragmented competitors who bolt third-party software onto OEM vehicles, Aurora controls every layer of the autonomy software chain – from its FirstLight LiDAR and imaging radar to HD maps, virtual testing, and remote-assist infrastructure – which runs on Nvidia hardware.

This end-to-end architecture is what enables industrialisation: it allows the playbook to rapidly be replicated lane by lane, with unit economics that are software-like – capital costs associated with trucks and factories are born by the OEM. Aurora charges fleets a per-mile Driver-as-a-Service fee designed to undercut the roughly $0.97 per mile human driver cost by 25-40% to $0.65–$0.85 per mile, with a further $0.15 per mile in secondary savings from lower insurance and turnover. As the network scales, hardware costs fall by around 50% with third-generation technology (in 2027), and remote-assist intensity declines sharply, driving a transition from positive gross profit in 2026 to an expected 70% gross margin by 2028.

THE DESIGN SOFTWARE BEHIND ROBOTIC AI – THE INVISIBLE LAYER OF PHYSICAL AI SITS IN SIMULATION

Cadence Design Systems – the physics layer of physical AI

Cadence sits at the quiet centre of the next AI wave; the company designs the software that designs not only the chips, but also the digital twin infrastructure that simulates and verifies robotic AI. Physical AI is a double opportunity for Cadence: every robot, car, and edge device needs new silicon for on-device inference, and every data centre still needs new silicon to train those models. Cadence designs both.

The invisible layer of physical AI sits in simulation – the software that allows robots to think before they move. Every motion policy, chip architecture, and energy system begins its life inside design tools that replicate physics, not just code. As robots grow more autonomous and physically capable, simulation becomes a profit pool in its own right: it shortens design cycles, prevents thermal and mechanical failures, and eliminates the cost of physical iteration.

In this space, Cadence Design Systems is emerging as the critical enabler. Its Reality Digital Twin and Optimality platforms simulate heat, vibration, airflow, and electromagnetic interference at chip and system level, letting engineers test physical constraints in software before building hardware. Cadence is now effectively the compiler for physical AI: its tools bridge AI model training, chip design, and system verification – the foundation of robots that think in heat, current, and motion.

As the AI market shifts from cloud to edge, systems design becomes the new frontier. Cadence's strength in automotive, robotics and life sciences, combined with its deepening partnership with ARM on edge-AI architectures, positions it as the physics backbone of physical AI – the company that makes sure the intelligence we build can survive the real world.

HARDWARE & COMPONENTS – INVEST IN THE BOTTLENECKS, SENSORS AND ACTUATORS; BEST-IN-CLASS OR LOWEST COST

Keyence – the nervous system of automation

Keyence's optical sensors, cameras, and laser profilers give robots their senses – the eyes and ears of roughly 40–45% of the world's industrial robots. Its technology forms the interface between digital control and the physical world, detecting position, distance, texture, and light with micron-level precision across factories, logistics, and inspection lines.

In humanoids, the sensor stack accounts for 4–7% of total robot cost (around 15–25% of hardware cost) and covers vision, proprioception, safety, and tactile sensing. Keyence plays across the first three: vision through high-resolution cameras and laser profilers, proprioception through distance and displacement sensors, and safety through light curtains and area scanners that create secure work zones. As humanoids reach Level 3 dexterity, sensor density jumps – more cameras, more proximity sensors, and, eventually, tactile arrays feeding force and motion data in real time.

The same trend is visible in factories. Moving from a human QC line to an AI-enabled line can 2–3x the total number of sensor points per station, as robots require constant, full-field feedback instead of sample-based inspection. Keyence is positioned at the centre of that sensor expansion: every additional actuator, camera, and vision station adds another node to its installed base.

The company's operating margins exceed 50%, higher than any software company in Japan, because every new sensor sold expands a proprietary data network used to optimise the next – effectively training data for the physical world. That feedback loop, combined with a relentless cadence of innovation – over 70% of annual sales come from products launched within the last three years, with 70% of those products ranked no. 1 in their category – keeps Keyence far ahead of peers.

The moat is execution. Keyence sells direct to more than 350,000 global clients, bypassing distributors to deliver engineering support and rapid iteration on-site. This high-touch model, paired with deep domain expertise, turns customer problems into product roadmaps at scale. In an era where factories are becoming intelligent organisms, Keyence is the sensory system feeding them awareness – indispensable, data-rich, and compounding.

Nabtesco – the muscle and the joint

Nabtesco dominates the global market for RV (cycloidal) precision reducers, the gear systems that convert a motor's high-speed spin into smooth, controllable torque – effectively the muscle and stability in a robot's heavy joints. Nabtesco holds roughly 60% global share of this category, supplying most of the world's industrial arms and factory robots where precision, stiffness, and shock resistance are non-negotiable.

As the automation cycle extends into humanoids and mobile manipulation, Nabtesco's reducers migrate naturally into the hips, knees, and waists of these new platforms – the same high-load, low-backlash environments its reducers already serve. In humanoids, lower-body joints require cycloidal or RV drives to handle dynamic gait, balance and impact forces; strain-wave gears alone can't absorb the load. That makes Nabtesco the de facto standard for the torque stack in the most demanding joints, while Harmonic Drive remains dominant in smaller, upper-body actuators.

The company's products deliver multi-kilonewton-millimetre stiffness, sub-arc-minute backlash, and 5× overload tolerance, enabling robots to move heavy payloads with repeatability measured in microns. Its gearsets appear in virtually every major industrial robot brand – Fanuc, Yaskawa, ABB, KUKA, and Kawasaki – and are now being designed into emerging humanoid platforms as torque density and durability benchmarks.

As humanoid production scales, the content opportunity multiplies: each humanoid carries 40–50 actuators, with the lower body accounting for the highest mechanical cost. Even if only half those joints use RV-class reducers, unit demand expands exponentially as humanoid volumes scale. Nabtesco's installed share and precision-manufacturing moat give it structural leverage in that transition.

Hesai – the eyes of autonomy

Hesai Technology is building a near monopoly in LiDAR – the critical sensing layer for autonomy across vehicles, robots, and infrastructure. It has achieved what most Western suppliers could not: volume and cost. Hesai's long-range hybrid solid-state sensors now ship at < \$400 per unit (vs \$1,000–1,500 for U.S. peers at volume), two years ahead of global competitors, and the company is on track for > 2 million units of annual capacity. Its in-house ASIC and simplified MEMS design deliver high-density point clouds at roughly half the cost of incumbents.

In the field, each autonomous vehicle or factory station typically carries 2–4 LiDAR units, providing 360-degree coverage and redundancy for mapping, perception, and safety. In humanoids, LiDAR is optional – many rely on stereo or structured-light cameras – but attach rates will rise in higher-speed or safety-critical tasks where spatial precision matters.

As autonomy scales, sensor content per car or robot increases even as ASPs fall. The math is simple: more devices deployed, more sensors per unit, more data to monetise. Hesai's customer footprint spans Li Auto, XPeng, and Baidu Apollo, and it has begun supplying top-five European OEMs. Beyond automotive, LiDAR demand is exploding in robotics: warehouses, ports, and logistics hubs are embedding Hesai sensors into mobile robots and mapping systems to enable precise navigation and collision avoidance.

Gross margins of roughly 25% on hardware and >70% at the group level (including software and services) prove that scale in optics can be profitable. Hesai is the optical analogue of Keyence – a company that has turned perception into a data business. In the race for physical AI, it supplies the eyes through which machines see.

EARLY ADOPTERS OF PHYSICAL AI – CREATING FIRST MOVER ADVANTAGES

Xiaomi – the dark factory operator

Xiaomi is one of the first end-to-end operators proving what a fully automated manufacturing loop looks like in practice. Its Smart Factory 2.0 in Beijing runs lights-out, 24 hours a day, producing roughly one smartphone every three seconds with no human intervention. Every production stage – SMT, assembly, testing, and packaging – is executed by robotic cells and AI-driven inspection systems built largely in-house. The plant assembles smartphones with micron-level precision, rejecting fewer than two units in a thousand, and improves yield >20% vs traditional human-run lines.

The economics mirror the BMW automated factory 58% ROI model but on consumer-electronics cadence: once integration is paid, utilisation drives compounding returns. Labour cost is near-zero; energy and maintenance dominate OpEx. Payback is measured in quarters, not years. The model scales – Xiaomi plans to replicate the architecture across multiple dark sites with robot-to-worker ratios > 500:1.

As an early adopter of physical AI, Xiaomi is creating a first mover advantage. It is already deploying computer-vision inspection, predictive-maintenance algorithms, and self-learning assembly software that can autonomously re-route production when defects appear. Owning this operational lead means Xiaomi isn't just automating output – it's codifying the playbook every other manufacturer will have to license or replicate.

Key Risks

Past performance does not predict future returns. You may get back less than you originally invested. We recommend any fund is held long term (minimum period of 5 years). We recommend that you hold funds as part of a diversified portfolio of investments.

The Funds managed by the Global Innovation Team:

- May hold overseas investments that may carry a higher currency risk. They are valued by reference to their local currency which may move up or down when compared to the currency of a Fund.
- May have a concentrated portfolio, i.e. hold a limited number of investments. If one of these investments falls in value this can have a greater impact on a Fund's value than if it held a larger number of investments.
- May encounter liquidity constraints from time to time. The spread between the price you buy and sell shares will reflect the less liquid nature of the underlying holdings.
- Outside of normal conditions, may hold higher levels of cash which may be deposited with several credit counterparties (e.g. international banks). A credit risk arises should one or more of these counterparties be unable to return the deposited cash.
- May be exposed to Counterparty Risk: any derivative contract, including FX hedging, may be at risk if the counterparty fails.
- Do not guarantee a level of income.

The risks detailed above are reflective of the full range of Funds managed by the Global Innovation Team and not all of the risks listed are applicable to each individual Fund. For the risks associated with an individual Fund, please refer to its Key Investor Information Document (KIID)/PRIIP KID.

Disclaimer

LIONTRUST
COURAGE · POWER · PRIDE