PHYSICAL AI FRONTIERS

ROBOTICS TRACK · EDITION R01 · JULY 2026

The Reality Gap

Simulation-trained policies, synthetic data, and what actually transfers from the simulator to the factory floor.

Lise DELOBEL

GenAI project manager & consultant — Physical AI, digital twins, and applied ML for industrial and extreme-environment operations. MIT IDSS Data Science & Machine Learning certificate.

LISEDELOBEL.FR · LISE.DELOBEL@ME.COM

The bottleneck moved

For three decades, the binding constraint in robotics was hardware. It no longer is. Actuators, sensors and onboard compute are now good enough that the limiting factor for most deployments is behaviour — and behaviour is learned from data the physical world is too slow, too expensive and too dangerous to supply. Simulation has therefore stopped being a test bench and become a production asset: the place where the training data is manufactured.

That shift creates a specific and underestimated risk. A policy trained in a simulator inherits the simulator's errors. Where those errors are small and systematic, the robot works. Where they sit in the parts of the physics nobody modelled — contact, friction, actuator response, sensor latency — the robot performs flawlessly in simulation and falls over on concrete. This is the reality gap, and closing it is now the central engineering problem of physical AI.

FINDING 01

The gap is not in the physics

Higher-fidelity rigid-body simulation does not close it. The dominant errors live in components that resist analytical modelling — geared, sprung, electronically controlled actuators whose real torque response carries delay, hysteresis and damping no equation captures cleanly.

FINDING 02

Learn the parts you cannot model

The durable answer is not better equations but measured behaviour: train a neural model of the component on real hardware data and embed it inside the simulator. The trade is compute and interpretability for transfer that actually holds.

FINDING 03

Dexterity now scales with data

Robot foundation models have produced their first credible scaling behaviour. Task performance improves with hours of human demonstration video rather than with bespoke per-task engineering — which changes who can compete, and on what.

FINDING 04

Your simulation is a strategic asset

An operator holding a high-fidelity, physically-validated model of its own facility holds proprietary training data for any robot that will ever work there. Most organisations still classify that model as a monitoring tool.

Two definitions used throughout

The sim-to-real gap is the performance loss incurred when a policy trained in simulation is deployed on physical hardware, arising from unmodelled dynamics, sensor noise, latency and distribution shift. A vision-language-action model (VLA) is a policy that maps camera images and natural-language instructions directly to robot actuator commands, replacing separately engineered perception, planning and control stages.

02 / WHY NOW

Four things changed at once

Sim-to-real is a twenty-year-old research topic. What is new is that every input it depends on became available, open and affordable inside roughly eighteen months.

20,854	200+	50–100
HOURS OF VIDEO	UNITS SHIPPED	DEMONSTRATIONS
Human egocentric footage used to pretrain a current open humanoid foundation model	Legged inspection robots in industrial service, running thousands of inspections weekly	Enough to adapt a leading on-device VLA to a new manipulation task

01 **Compute moved onto the robot**

Inference that previously required a workstation now runs on humanoid-class edge modules with sub-10ms latency. That removes the architectural excuse for cloud round-trips and makes learned policies viable inside real control loops.

02 **Foundation models became open and commercially licensed**

NVIDIA's Isaac GR00T N-series ships under permissive licensing with production support, and industry leaders including Agility, FANUC, Figure, KUKA, Skild AI, Universal Robots and Yaskawa are building on the same platform. A robotics team no longer starts from zero.

03 **Simulation became the data source, not the test bench**

Generative world models now produce synthetic trajectories at volume to train other models. Simulation has moved upstream of the policy — it manufactures the training set rather than validating the result.

04 **Latency economics are unforgiving**

A language model that takes two seconds to start responding is acceptable. A robot that takes two seconds to react is dangerous. Deployed policies must run at tens to hundreds of hertz on bounded hardware, which constrains every architectural choice made upstream.

Sources: NVIDIA Newsroom, GTC 2026 · MarkTechPost, Apr 2026 · Open Data Science, Jun 2026 · Climate Investment, Sep 2025 · Robotics Center (SVRC), Apr 2026.

Where simulation stops matching the world

Four sources of divergence, and the four strategies the field has developed against them — in ascending order of what they cost and what they buy.

DIVERGENCE 01

Contact and friction

Whether a foot slips or a gripper holds is decided in a contact patch millimetres across. Friction coefficients vary with surface, wear, temperature and contamination, and solvers approximate the contact event itself.

DIVERGENCE 02

Actuator response

Real joints contain gearboxes, springs and embedded controllers. Commanded torque and delivered torque differ by a delay and a distortion that vary with load and direction.

DIVERGENCE 03

Sensor noise and latency

Cameras blur, lidar returns drop out, and every measurement arrives late. Policies trained on clean, instantaneous state learn to depend on information the robot never has.

DIVERGENCE 04

The long tail

The simulator contains what the engineer thought to put in it. Spilled oil, a badly parked pallet, glare through a window and low battery are all in the real distribution and none are in the training set.

THE LADDER OF RESPONSES

01

Domain randomisation

Vary the unknown parameters widely so the policy is forced to be robust to all of them.

02

System identification

Measure the real machine and tune the simulator's parameters to match it.

03

Learned dynamics

Replace the components that resist modelling with neural models trained on hardware data.

04

Generative world models

Synthesise whole environments and trajectories, expanding coverage beyond what was built by hand.

The strategic read

These four are cumulative, not alternative. Randomisation buys robustness cheaply and plateaus. Identification buys accuracy and does not survive wear. Learned dynamics buy transfer at the cost of interpretability. Generative world models buy coverage at the cost of provenance. The question for an operator is not which one to adopt — it is how far up the ladder the application actually requires you to climb, because each rung adds cost, data obligations and assurance burden.

From the lab bench to a Zone 1 refinery

Legged inspection robots — ETH Zurich Robotic Systems Lab and ANYbotics

The clearest evidence that sim-to-real transfer is a commercial technology rather than an academic result is that you can now buy it, certified, for use in an explosive atmosphere. The lineage runs directly from a 2019 Science Robotics paper on training legged locomotion in simulation to a fleet operating in refineries, mines and power plants.

THE SYSTEM

What was deployed

ANYmal is a four-legged inspection robot built for heavy industry — IP67-rated, ruggedised, carrying visual and thermal cameras, ultrasonic microphone, lidar and up to 10 kg of additional payload. Its onboard AI stack performs autonomous navigation, collision avoidance and stair climbing, and it reads equipment temperatures between −40°C and 550°C without physical contact.

THE SCALE

What it is worth

More than 200 units have shipped, performing thousands of inspections weekly across oil and gas, mining, power, utilities and metals. Deployments include Equinor, ENI and Petrobras, and the Northern Lights carbon capture facility in Norway — a joint venture of Equinor, TotalEnergies and Shell. Company funding passed $150 million in 2025. The robot detects overheating, abnormal vibration and fugitive emissions.

THE THRESHOLD CROSSED

The certification milestone

ANYmal X, the explosion-proof variant, carries ATEX and IECEx Zone 1 certification — approval for atmospheres containing flammable gases, vapours or dust — and reached customer delivery in 2026. Large areas of oil, gas and chemical plants are classified Zone 1 and previously required hot-work permits or human entry for routine inspection.

So what

A locomotion policy that could not previously be hand-engineered was learned in simulation, transferred to hardware, and has now passed one of the most conservative safety-certification regimes in industry. The governing lesson for an operator is procurement-shaped rather than technical: this is an inspection programme, not a robot purchase. Value appears only when routes, data models, alerts and maintenance workflows are designed around the machine — which is the same finding as the digital twin literature, arriving from the opposite direction.

Sources: Hwangbo et al., Science Robotics 4(26), 2019 · ANYbotics product documentation, ANYmal and ANYmal X, 2026 · Climate Investment, Sep 2025 · The Robot Report, Sep 2025 · SAP News, Mar 2026. Deployment counts and customer lists are company-reported.

The first scaling law for robot dexterity

Open humanoid foundation models — NVIDIA Isaac GR00T and Google Gemini Robotics

Language models became useful when the field found a reliable relationship between data volume and capability. Robotics has spent years without one, which is why progress looked like a sequence of impressive demonstrations rather than a curve. That changed in 2026.

THE FINDING

The result

GR00T N1.7, released in April 2026, is a 3-billion-parameter open, commercially licensed vision-language-action model built on a Cosmos-Reason2 backbone with a 32-layer diffusion transformer for low-level motor control, in a dual-system 'Action Cascade' architecture. Its central advance is pretraining on 20,854 hours of human egocentric video across more than twenty task categories. NVIDIA reports what it describes as the first scaling law for robot dexterity: moving from 1,000 to 20,000 hours of human egocentric data more than doubles average task completion.

THE MECHANISM

Why the data source matters

The scaling input is human video, not robot teleoperation. Teleoperation data is expensive, slow and specific to one robot; egocentric human footage is abundant and embodiment-agnostic. The dual-system split reflects the same latency constraint seen elsewhere — a slow reasoning backbone running at a few hertz paired with a fast reactive policy running at hundreds of hertz.

THE ECONOMICS

The adaptation cost collapsed

On the Google side, Gemini Robotics on-device — optimised for execution where latency or connectivity constrain a cloud round-trip — adapts to new tasks with as few as 50 to 100 demonstrations, and has been carried across bi-arm Franka FR3 hardware and Apptronik's Apollo humanoid. A 2026 update added improved spatial reasoning and instrument-reading, developed with Boston Dynamics.

So what

If capability scales with hours of demonstration data, then competitive advantage in applied robotics shifts from control engineering toward data acquisition, curation and rights — a transition the language-model market has already completed. For an industrial operator the practical implication is uncomfortable and clarifying: the footage of your own people doing your own work, in your own facility, is a strategic asset that most organisations currently discard.

Sources: MarkTechPost, Apr 2026 · NVIDIA Newsroom, GTC 2026 · Open Data Science, Jun 2026 · Robotics Center (SVRC), Apr 2026. Scaling-law and task-completion figures are as reported by the model developers and are not independently audited.

The simulator as the factory

Generative world models, synthetic trajectories, and the digital twin you already own

The third use case is the one with the most direct consequence for organisations that are not robotics companies. If policies are trained in simulation, and simulation quality determines what transfers, then whoever owns the best model of a physical environment owns the ability to train robots for it.

THE ARCHITECTURE

Simulation moved upstream

Generative world models such as NVIDIA Cosmos are not robot policies. They produce synthetic trajectory data at volume to train the policies — a data factory rather than a controller. Paired with large-scale physics simulation, the current pattern is pretraining on the order of a million simulated episodes, then fine-tuning on tens of thousands of real ones.

THE OVERLAP

The bridge to digital twins

The geometry, kinematics and material properties an operator has already assembled for a digital twin are the same assets a training simulator requires, and the interchange layer converging across the industry is OpenUSD. A twin built for monitoring and a twin built for policy training differ less in content than in physical validation and in the rigour of the tolerances attached to them.

THE GAP TO CLOSE

What separates the two

A monitoring twin has to be visually and topologically faithful. A training twin has to be dynamically faithful — contact, friction, mass distribution, actuator response and sensor latency all have to be right enough that a policy learned inside it survives contact with the building. That is a materially higher and more expensive standard, and it is where most twin programmes will discover their limits.

So what

This reframes digital twin investment. A twin justified today on inspection efficiency or downtime avoidance has a second, larger option value: it is the training environment for every autonomous system that will ever operate in that facility. Executives evaluating twin programmes should be asking what physical validation would cost, because that expenditure is what converts a dashboard into an asset — and it is far cheaper to specify at the start than to retrofit.

Sources: MarkTechPost, Apr 2026 · NVIDIA Newsroom, GTC 2026 · Robotics Center (SVRC), Apr 2026. Cross-references Digital Twin White Paper Series, editions on industrial and energy assets.

The actuator that could not be modelled

Series-elastic joint units on a quadruped robot — ETH Zurich Robotic Systems Lab

THE CHALLENGE

Each of the robot's twelve joints is a series-elastic actuator: a brushless motor driving a harmonic-drive gearbox, connected to the limb through a torsional spring. The spring is deliberate. It absorbs the shock of a foot striking ground, stores energy between strides, and turns the hard problem of measuring torque into the easy one of measuring a deflection. It also makes the joint analytically intractable. Between the commanded torque and the torque that actually reaches the leg sit a transport delay, direction-dependent friction, gearbox compliance and damping that no closed-form expression reproduces across the operating range.

A rigid-body simulator assumes the commanded torque arrives. It does not — it arrives late, distorted, and differently depending on load. Policies trained against that assumption learn gaits that depend on torque the hardware will never deliver on time.

Why the obvious solutions fail

Raising simulation fidelity does not help, because the error is not in the rigid-body physics — it is inside the actuator's own mechanics and control electronics, which the physics engine does not simulate at all. Training on the real robot is equally unavailable: reinforcement learning needs millions of trials, and a dynamically balancing machine that falls over several million times destroys itself long before a policy converges.

FIG. 01 – SERIES-ELASTIC JOINT UNIT, LONGITUDINAL SECTION (SCHEMATIC, NOT TO SCALE)

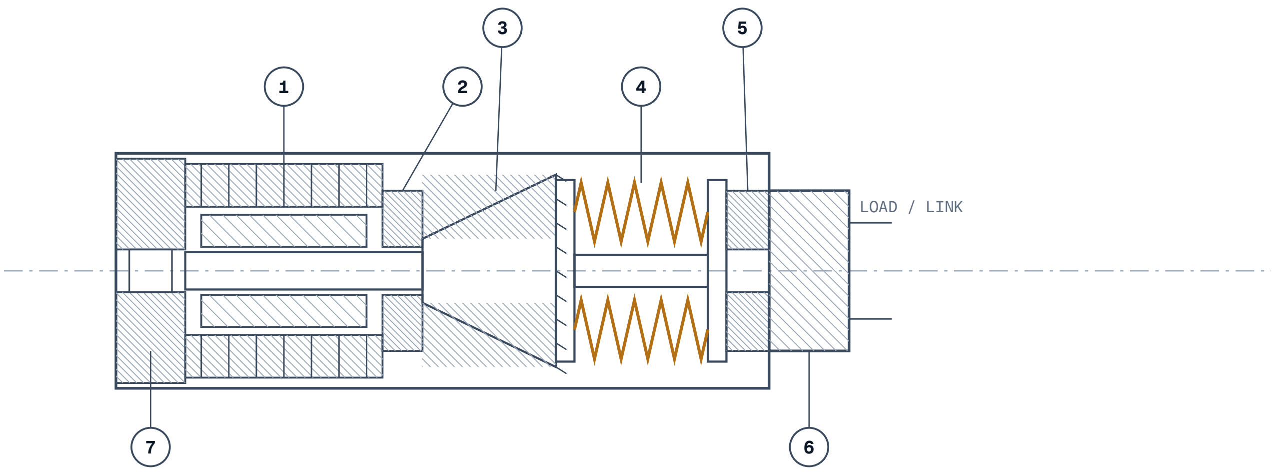

1 Brushless motor – stator / rotor
2 Motor-side position sensor
3 Harmonic drive, high reduction
4 Torsional spring – compliant element
5 Output-side absolute encoder
6 Output flange to link
7 Embedded control electronics, CAN

Stop modelling it. Measure it.

The turning point, the trade-off accepted to ship, and the result on hardware

THE RESOLUTION

The insight that unlocked the problem was a change of category rather than a change of technique. If the actuator's behaviour cannot be written down, it can still be measured. The team drove the real joints across a wide range of commands, recorded the history of position error and velocity against the torque actually delivered, and trained a small neural network — an actuator network — to predict what the physical hardware would do for a given command history. That network was then dropped into the simulator in place of the analytical actuator model.

The rest of the simulator stayed conventional. Only the component that resisted modelling was replaced by a learned one, which is why the approach generalised: it is a method for handling any subsystem whose behaviour is measurable but not derivable.

FIG. 02 — TORQUE STEP RESPONSE

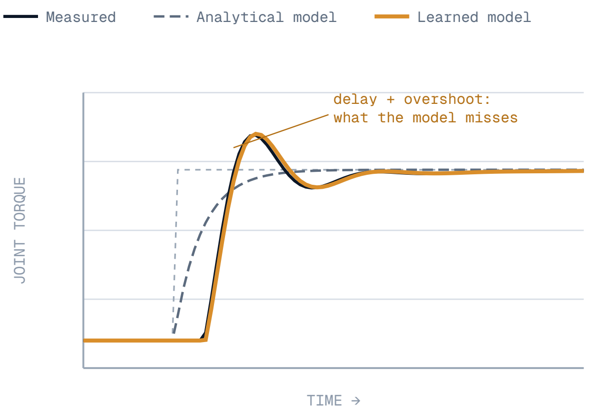

SPECIFICATION — REPRESENTATIVE JOINT UNIT

Nominal voltage	48 VDC
Peak / nominal power	480 / 180 W
Joint position accuracy	0.025°
Torque resolution	8 mNm
Actuated joints per robot	12
Command interface	CAN / CANopen

Published figures for a representative series-elastic joint unit of this class. The compliant element is what makes torque measurable — and what makes the response analytically hard.

The trade-off accepted to ship it

Roughly half of the simulator's runtime is now spent evaluating the actuator networks. The hybrid simulator still runs at close to 500,000 time steps per second — about a thousand times faster than real time — so the cost was affordable. The deeper concession is epistemic: the actuator became a black box, valid only within the distribution of commands it was trained on, and no longer inspectable the way an equation is. For a safety case, that is a real and permanent loss.

THE OUTCOME — POLICY DEPLOYED DIRECTLY TO HARDWARE, NO FURTHER TUNING

−29.7%
MECHANICAL TORQUE
Against the best existing model-based controller for the same robot

0.143 m/s
VELOCITY ERROR
Average linear velocity tracking error following commanded body velocity

1.5 m/s
PEAK SPEED
Beyond previously achieved limits, plus autonomous recovery from an inverted fall

Source: Hwangbo, Lee, Dosovitskiy, Bellicoso, Tsounis, Koltun and Hutter, “Learning agile and dynamic motor skills for legged robots,” Science Robotics 4(26), eaau5872, 2019. Mechanical power was separately reduced by 19.8%. Drawing is schematic and generic to the actuator class.

Building the capability, not buying the robot

A sequence that front-loads measurement, because everything downstream inherits its quality.

PHASE 01 · 0–3 MONTHS

Instrument the physical asset

Before any simulation work, capture how the real equipment and the real environment actually behave: command-versus-delivered response on the machines you already own, surface and lighting conditions, cycle timings, failure modes. This data is the reference against which every later model is validated. Teams that skip this step discover eighteen months later that they cannot tell a good simulator from a bad one.

PHASE 02 · 3–9 MONTHS

Build a dynamically validated twin

Extend or construct the environment model with the properties a policy needs — mass, friction, compliance, sensor placement and latency — on an interchange standard, with OpenUSD the current convergence point. Define acceptance tolerances explicitly: a twin is validated when its predictions match measured behaviour within a stated bound, not when it looks correct.

PHASE 03 · 9–18 MONTHS

Close the loop and build the evaluation harness

Introduce learned dynamics for the subsystems that resist modelling, train policies against the validated environment, and — critically — build the test regime before the results. Transfer performance must be measured on held-out physical scenarios that the simulator never generated, or the programme has no evidence, only demonstrations.

PHASE 04 · 18 MONTHS +

Scale under governance

Fleet deployment, retraining cadence as equipment wears and the environment changes, model provenance records, and a defined authority boundary for what the autonomous system may do unsupervised. Certification regimes such as ATEX or IECEx apply to the machine; the assurance argument for the learned policy is an additional and separate obligation.

Organisational readiness

The capability sits across three functions that rarely report to one another: engineering owns the physical measurement, data science owns the models, and operations owns the acceptance criteria. The single most reliable predictor of failure is a programme where simulation is procured by IT as a software licence rather than commissioned by engineering as a measurement instrument.

What goes wrong, and what it costs

Ranked by how often it is missed rather than by severity.

RISK 01

Sim-washing

A demonstration video is not evidence of transfer. Impressive simulated behaviour is cheap to produce and reveals nothing about hardware performance. Insist on results measured on physical scenarios the simulator did not generate.

RISK 02

No accepted benchmark

The field has no agreed measure of transfer quality. Vendor comparisons are therefore close to meaningless, and internal acceptance criteria have to be defined by the buyer, in advance.

RISK 03

Black-box components in a safety case

A learned actuator or dynamics model cannot be inspected the way an equation can. Where certification requires an argued safety case, this is a structural obstacle rather than a documentation task.

RISK 04

Distribution limits

A learned model is valid only within the conditions it was trained on. Wear, temperature, contamination and modification all move the machine outside that envelope silently.

RISK 05

Fidelity cost curves

Physical validation is expensive and its cost rises sharply near the end. Programmes routinely budget for the visually convincing twin and not for the dynamically faithful one.

RISK 06

Data provenance

Models pretrained on large human video corpora carry rights, privacy and consent questions that have not been settled. Operators recording their own workforce inherit those questions directly.

RISK 07

Simulation-stack lock-in

The simulator, the interchange format, the foundation model and the edge compute increasingly ship as one commercial stack. Retain the ability to move your environment model between them.

RISK 08

Trust calibration

Once a robot works reliably, supervision decays. The aviation-automation literature is explicit that operator vigilance degrades fastest precisely when a system is good but not perfect.

08 / WHERE TO START

Three questions worth asking this quarter

01 Do we know how our own equipment actually responds — measured, not specified?

02 Is our digital twin dynamically validated, or only visually faithful?

03 Who in this organisation would be accountable for a learned policy's behaviour?

If your organisation is evaluating autonomous systems, or holds a digital twin that could become a training environment, these are the questions that determine whether the investment compounds.

NEXT IN THIS TRACK

Edition R02 — The Last Inch

Logistics picking: why everything upstream of the gripper is solved and the grasp is not, and what a three-second cycle-time budget does to a vacuum seal on a wrinkled polythene bag.

SOURCES & METHOD

Deployment counts, customer lists and model-performance figures reported by vendors or developers are identified as such and are not independently audited. The Engineering Zoom draws on peer-reviewed published work; the section drawing is a schematic generic to the actuator class and is not a manufacturer drawing.

Sim-to-real and actuator modelling: Hwangbo, Lee, Dosovitskiy, Bellicoso, Tsounis, Koltun, Hutter, "Learning agile and dynamic motor skills for legged robots," Science Robotics 4(26), eaau5872, 2019 · arXiv:1901.08652 · published ANYdrive joint-unit documentation.

Industrial deployment: ANYbotics product documentation for ANYmal and ANYmal X, 2026 · Climate Investment, Sep 2025 · The Robot Report, Sep 2025 · SAP News, Mar 2026.

Foundation models and simulation: NVIDIA Newsroom, GTC 2026 · MarkTechPost, Apr 2026 · Open Data Science, Jun 2026 · Robotics Center (SVRC), Apr 2026 · vision-language-action deployment survey, May 2026.

Lise DELOBEL

GenAI project manager & consultant — Physical AI, digital twins, and applied ML for industrial and extreme-environment operations. MIT IDSS Data Science & Machine Learning certificate.

LISEDELOBEL.FR · LISE.DELOBEL@ME.COM

PHYSICAL AI FRONTIERS · ROBOTICS TRACK · EDITION R01 · JULY 2026