PHYSICAL AI FRONTIERS

ROBOTICS TRACK · EDITION R04 · JULY 2026

The Last Hundred Microns

Precision assembly when available clearance is smaller than the combined uncertainty of robot, tooling, fixturing and parts.

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

When the task is finer than the machine

The argument in one line: when total workcell uncertainty exceeds assembly clearance, precision has to be recovered through the interaction loop rather than assumed from nominal positioning.

In precision assembly the clearance between mating parts is routinely smaller than the combined uncertainty of robot pose, tool centre point, fixturing, workpiece and sensing. When that holds, conventional open-loop positioning is generally insufficient on its own — not because position control is useless, but because it can no longer guarantee the approach.

This is a different problem from the ones examined earlier in this series, and a more uncomfortable one. It is not that the model is imperfect or that a quantity is unobservable. It is that the required precision exceeds what the machine can achieve open-loop, by roughly an order of magnitude. Calibration reduces systematic error but does not eliminate fixture variation, thermal drift, part-to-part variation or contact uncertainty across production. The practical answer is to supplement positioning with compliance and contact feedback, and let the parts and the measured interaction forces guide the final alignment.

FINDING 01

The numbers do not permit open-loop positioning

Published experiments use clearances from 0.02 mm to around 0.12 mm, while the arms involved have repeatability on the order of hundredths to roughly a tenth of a millimetre. Repeatability is only one term: the governing figure also includes tool centre point, fixture, workpiece and thermal contributions.

FINDING 02

The failure mode is geometric, not a matter of force

Insertion fails through jamming: the peg contacts both walls, the friction cones overlap, and the part locks. Jamming is defined as being unable to move because of improperly applied forces and moments — so pushing harder deepens the jam rather than resolving it.

FINDING 03

In one controlled trial, wrench feedback was the enabling mechanism

Across three peg geometries at under 0.2 mm clearance, force-guided insertion succeeded 10, 8 and 7 times out of 10. With wrench sensing disabled on the same apparatus and routine, the results were 1, 0 and 0. In that setup, disabling sensing converted a viable process into one that almost always failed.

FINDING 04

Compliance can sit in the tool rather than the arm

Two architectures appear in recent work: an inherently compliant, force-sensing wrist used inside the arm's own force-control loop, and an actively controlled multi-axis end-effector that regulates contact force through its own mechanism. Both relocate the precision requirement into modular tooling.

Two definitions used throughout

Jamming is a state in which the part cannot move because the applied forces and moments fall outside the wedge of directions that would free it; it is resolved by changing the applied moment, not by increasing force. Admittance control converts a measured contact force into a commanded motion — the inverse of impedance control — and is what allows a position-controlled industrial arm to behave compliantly.

An old problem with new economics

Force-controlled assembly was demonstrated in the early 1980s. What changed is the cost of the sensing, the arrival of learned search strategies, and where the compliance now sits.

0.02 mm
TIGHTEST PUBLISHED CLEARANCE
Achieved by zero-shot sim-to-real transfer with dynamic compliance tuning

10 → 1
SUCCESS WITH AND WITHOUT FORCE
Cylindrical peg insertions per ten attempts, sensor enabled against disabled

~0.1 mm
ORDER OF ARM REPEATABILITY
In the cited studies — one term only; fixture, TCP, part and thermal error add to it

01 The sensing got small and cheap

Six-axis force-torque sensing has moved from a laboratory instrument to a component that fits inside end-of-arm tooling. Locating sensing and compliance near the contact reduces the influence of upstream joint compliance and simplifies local force regulation, though performance still depends on the whole control-loop bandwidth and mechanics.

02 Compliance became active rather than passive

The classical answer was the remote centre compliance device — a passive mechanical device whose stiffness and compliance centre are selected for a defined range of part geometries and misalignments. Active force control widens that range: the same hardware handles different holes, pose errors and unknown dimensional tolerances without mechanical re-tuning.

03 Learned policies started closing the search phase

Insertion has two phases: finding the hole, then going in. Recent work transfers search-and-insert policies from simulation to hardware with no real-world fine-tuning, reporting 90% and 100% success at 0.05 mm clearance and 70% and 90% at 0.02 mm — the sim-to-real problem of Edition R01 applied to contact.

04 The parts are getting tighter, not looser

Electrification and miniaturisation push tolerances down while volumes go up. Connectors, bearings, battery module fasteners and RF interfaces all sit in the range where an operator either solves force-controlled assembly or keeps a human doing it.

Sources: Zhang et al., "Fast skill learning for variable compliance robotic assembly," arXiv:1905.04427 · "Bridging the sim-to-real gap with dynamic compliance tuning for industrial insertion," arXiv:2311.07499 · "Delta6: a low-cost, 6-DOF force-sensing flexible end-effector," arXiv:2604.06150 · Whitney, quasi-static assembly of compliantly supported rigid parts, 1982.

Four contact states that shape the strategy

In two dimensions a peg entering a clearance-fit hole occupies one of four simplified states. The control problem is to manage the one most associated with jamming.

STATE 01

No contact

The peg is inside the hole envelope but touching nothing. This is the state an open-loop controller assumes it can reach. At very tight clearances it cannot be relied upon unless total workcell uncertainty is demonstrably below the available clearance.

STATE 02

One-point contact

The peg touches a single wall. This state is benign and informative: the contact force has a well-defined direction, so it reveals which way the peg is offset and a lateral correction resolves it. Force sensing turns this into navigation.

STATE 03

Two-point contact — potential jamming

The peg touches opposing walls at different depths. Depending on geometry, friction and the applied wrench, this can produce jamming, or wedging where geometry and friction self-lock. Increasing axial force then raises normal and friction forces rather than producing insertion. Two-point contact is a configuration, not automatically a jam.

STATE 04

Line contact

The peg lies flush against a wall along its length. Recoverable, but it signals gross misalignment rather than the fine error that force control is designed to absorb.

THE LADDER OF RESPONSES

01
Interface design
Chamfers, lead-ins and self-centring features. The cheapest precision is designed in, not controlled in.

02
Presentation & fixturing
Reduce the uncertainty that reaches the robot before any control strategy is asked to absorb it.

03
Passive compliance
A device absorbs misalignment mechanically. Fast and cheap across its designed range.

04
Active wrench control
Measure contact and command motion from it. Widens the range of geometries and tolerances.

05
Learned search
Train the strategy in simulation and transfer it, so the policy handles acquisition too.

The strategic read

Buying a higher-accuracy arm is the response most operators reach for first. It is often the most capital-intensive one, and it should be compared against interface design, fixturing, local compliance and sensing before it is specified — not because arm accuracy is worthless, but because it purchases precision across an entire workspace to solve a problem at one point. Note also that these rungs combine: most industrial assembly uses position control for the free-space approach and switches to compliant, force-limited control on contact.

Proving that the sensor is the process

Force-guided insertion against disabled sensing — a controlled comparison

Most claims about force control are directional: it helps, it improves robustness. A recent low-cost end-effector study did something more useful — it ran the same insertion strategy with wrench sensing disabled, on the same apparatus and task routine, which isolates the contribution of sensing from everything else in the stack.

THE EXPERIMENT

The setup

A six-degree-of-freedom, spring-driven, inherently compliant force-torque-sensing end-effector — used inside the arm's own force-impedance loop rather than actuating compliance itself — fitted with a gripper, inserting pegs of three geometries — cylindrical, square and hexagonal — into matching holes with clearances under 0.2 mm, which the authors note renders pure position control ineffective. The routine is three stages: initial contact, a compliant search that aligns the peg, then force-guided insertion.

THE COMPARISON

The result

Across ten trials per geometry, force-guided insertion succeeded 10/10 for the cylinder, 8/10 for the square and 7/10 for the hexagon. With wrench sensing disabled, the same apparatus and routine returned 1/10, 0/10 and 0/10. The reported traces show unwanted torques being actively suppressed, allowing continuous micro-reorientation that keeps the gripper coaxial with the bore.

THE PATTERN

Why geometry degrades the result

The cylinder is rotationally symmetric about the insertion axis, so only position and tilt must be resolved. The square and hexagon add a rotational alignment that must be found before insertion can proceed, and each flat introduces additional contact states. The lower rates for non-axisymmetric parts are consistent with that extra alignment requirement — though ten trials per geometry cannot establish a general scaling law, and finish, chamfers and controller tuning may all contribute.

So what

This is the cleanest available answer to a question that comes up in every automation business case: is wrench sensing a refinement or a requirement? In this setup, at these tolerances, it was the requirement. An evaluation that reports only aggregate success rates without the sensor-disabled control tells you very little — and asking a vendor to run precisely that comparison on your parts is cheap, fast, and unusually revealing.

Source: “Delta6: a low-cost, 6-DOF force-sensing flexible end-effector,” arXiv:2604.06150, Section on peg-in-hole tight assembly.

Training the search in simulation

Zero-shot transfer of insertion policies at industrial clearances

Edition R01 argued that policies trained in simulation fail on hardware because the simulator's unmodelled dynamics are concentrated in contact. Precision insertion is the hardest possible test of that claim, because the entire task is contact. It is therefore notable that it now transfers at all.

THE FINDING

The result

Work on dynamic compliance tuning reports zero-shot sim-to-real transfer on 3D-printed pegs and holes at clearances of 0.05 mm and 0.02 mm: 90% and 100% success for rectangular and round parts at 0.05 mm, and 70% and 90% at 0.02 mm. Comparison methods trained by imitation or reinforcement learning failed most trials under the same transfer conditions. The authors separately show that fine-tuning the Force Planner on ten real trajectories improves the harder tasks further — a result distinct from the headline zero-shot evaluation.

THE MECHANISM

Dynamic gain adaptation was central to transfer

The architecture separates a Force Planner, which plans motion and desired contact force, from a Gain Tuner that adjusts admittance gains online so real contact forces track the planned ones. It is the gain tuning rather than the trajectory that carries the transfer — a schedule for responding to contact is more portable across the reality gap than a path through space.

THE LIMIT

Where it remains hard

Rectangular parts are consistently harder than round ones for the baseline methods, and success falls as clearance tightens. These are research results on printed parts rather than production data on metal components, and they should be read as evidence that the approach is viable, not that the problem is closed.

So what

The transferable principle is that what you choose to learn determines whether it survives contact with reality. Learning a path bakes in the specific geometry and the specific error; learning how to react to force encodes a strategy that holds across both. When evaluating any simulation-trained capability, the question worth asking is not how accurate the simulator is, but what the policy was asked to output.

Source: “Bridging the sim-to-real gap with dynamic compliance tuning for industrial insertion,” arXiv:2311.07499, Table II. Results are on 3D-printed parts under laboratory conditions. Cross-references Edition R01.

Moving the precision out of the arm

Compliant end-effectors and where accuracy should be bought

The instinctive response to a precision shortfall is to buy a more precise machine. It is worth examining why the engineering literature keeps arriving somewhere else.

THE ARGUMENT

Accuracy in the arm is the expensive place to put it

A robot's positional error accumulates through every joint, link and gearbox between the base and the tool, and it drifts with temperature across a working shift. Improving it means improving the whole kinematic chain. Meanwhile the error that actually matters is a few tens of microns at one point: the contact.

THE ALTERNATIVE

A compact parallel end-effector instead

Recent work uses a lightweight multi-degree-of-freedom end-effector in a three-prismatic-revolute-spherical parallel configuration with active force control, mounted on a conventional arm. Reported force control holds an absolute mean error of 1 N against a 100 N setpoint along the insertion axis. Publisher metadata reports 0.04 N·mm mean absolute error for moments about the transverse axes; another version of the paper states the figure in different units, so it should be verified against the typeset article.

THE ECONOMICS

What this does to the capital argument

If compliance lives in a small module at the wrist, the precision capability becomes a modular tooling function rather than an intrinsic specification of the whole robot: replaceable, upgradable independently of the arm, and potentially shared across cells. The business case stops being a robot purchase and becomes a tooling decision.

So what

Before specifying a higher-accuracy arm, establish where the precision is actually required. If it is needed only at the contact — which in assembly it usually is — then buying it in the arm means paying for accuracy across an entire workspace to obtain it at a single point. That is a recurring pattern in this series: the expensive general solution is often standing in for a cheap specific one.

Source: Zhao, Hailegebrial, Wang, Wen and Zheng, “Robotic jamming-free and precise assembly for peg-in-hole operation by multi-DOF force-controlled parallel end-effector,” Robotics and Computer-Integrated Manufacturing, Volume 97, February 2026, 103083 (preprint SSRN 5186899). Versions of the paper report the moment-control error in different units; the publisher figure is quoted and flagged for verification.

The jam that pushing cannot fix

Two-point contact in clearance-fit insertion — the geometry of failure

THE CHALLENGE

Consider a peg entering a hole with a clearance of a few hundredths of a millimetre, held by an arm whose pose repeatability is around a tenth of a millimetre. The peg will not enter cleanly, because it cannot be placed accurately enough to do so. It will touch a wall — and what happens next is decided by geometry rather than by how hard the robot pushes.

If it touches at one point, the situation is recoverable and even informative: the measured contact force points in a known direction, so it can be used to command a corrective displacement. If it tilts slightly and touches both walls at different depths, the part enters two-point contact. There the friction cones at the two contacts can overlap, and no purely axial force produces motion. The assembly literature names this jamming — being unable to move because of improperly applied forces and moments — and, under self-locking combinations of geometry and friction, wedging, where the part is stuck despite high insertion force. Two-point contact does not automatically mean either; it is the configuration in which both become possible.

Why the obvious solutions fail

Pushing harder is worse than useless: in a jam the normal forces at both contacts rise with the insertion force, so friction rises in proportion and the part locks more firmly. Buying a more accurate robot does not reach the tolerance either, and thermal drift across a shift erodes whatever calibration achieved. Vision reduces initial pose error, but once the mating surfaces are occluded it cannot directly resolve the internal contact state — the same structural blindness identified in Edition R02, where the governing quantity was likewise not an optical property.

FIG. 01 – CONTACT STATES IN CLEARANCE-FIT INSERTION (SCHEMATIC, EXAGGERATED TILT)

A ONE-POINT CONTACT · B POTENTIAL JAM · C RESOLVED

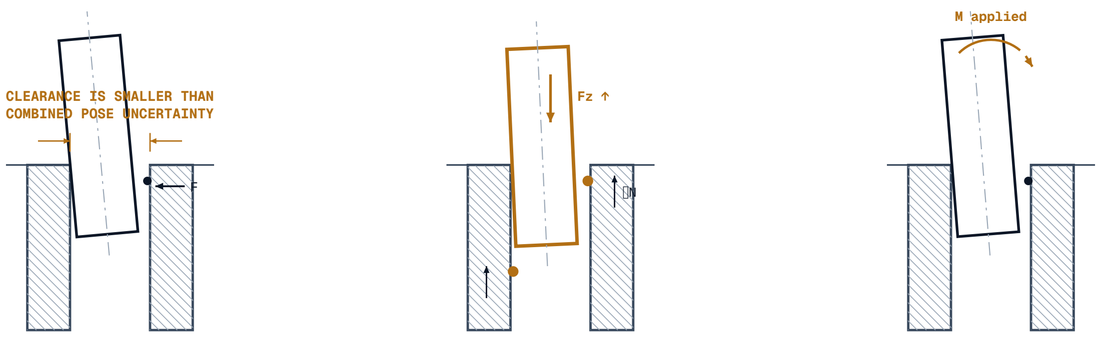

A — Peg touches one wall. The measured lateral force can command a corrective displacement.

B — Both walls loaded. Under some geometry and friction, axial force cannot advance the part.

C — Moment applied to release the second contact, then insert.

1 Clearance: tens of microns
2 Combined pose uncertainty is larger
3 Jam depends on friction and wrench
4 More force deepens the jam
5 Moment restores one-point contact
6 Insertion resumes under force control

Turn the jam back into a touch

The design move, the trade-off accepted to ship, and the published result

THE RESOLUTION

The turning point is a reframing of what the task is. Conventional strategies treat insertion as a sequence of position corrections: detect a problem, retract, adjust, try again. The alternative treats the whole operation as a reciprocal transition between a transient one-point contact and a steady two-point jam — and designs a controller whose job is simply to keep converting the second state back into the first.

The move that does the converting is a moment, not a force. Because the jam is a condition on the direction of applied load rather than its magnitude, applying a controlled moment about the transverse axis releases one of the two contacts and returns the part to one-point contact, from which insertion proceeds. Built on admittance control, this removes the need for repeated positional adjustment during insertion, and tolerates unknown dimensional tolerances rather than requiring them to be measured first.

FIG. 02 — INSERTION SUCCESS, SENSOR ON / OFF

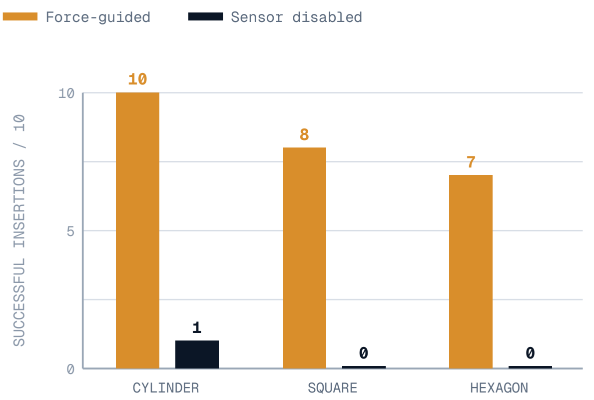

FIGURES FROM THREE INDEPENDENT STUDIES

Delta6 — clearance	< 0.2 mm
Delta6 — trials	10 per geometry
Compliance tuning — clearance	0.02 / 0.05 mm
Compliance tuning — parts	3D printed
3-PRS — force setpoint	100 N axial
3-PRS — mean error	1 N

Three separate studies on the same problem class, with different apparatus, controllers, parts and objectives. Each row is labelled with its source; the rows should not be read as one experiment.

The trade-off accepted to ship it

The part is now deliberately driven into contact and worked loose, rather than kept clear of the walls. That accepts sliding contact under load on mating surfaces — a wear and surface-finish cost on precisely the features whose tolerance justified the exercise. It also makes cycle time variable, because the number of jam-release iterations depends on the initial error, which is uncomfortable for a line balanced on fixed takt. Both are engineering inferences about production, not findings: the cited laboratory studies do not establish surface wear, cycle-time distributions or tooling life.

THE OUTCOME — PUBLISHED RESULTS ACROSS THE PROBLEM CLASS

10 / 1
SENSOR ON VS OFF
Cylindrical peg insertions per ten attempts, under 0.2 mm clearance

1 N
AT 100 N SETPOINT
Absolute mean force-control error reported for the parallel end-effector

0.02 mm
TIGHTEST TRANSFER
Clearance at which a simulation-trained policy still reached 90% zero-shot

Sources: arXiv:2604.06150 (Delta6) · SSRN 5186899 / Robotics and Computer-Integrated Manufacturing 2025 · arXiv:2311.07499 · Royal Society Open Science 6:190476 for contact-state taxonomy. Drawing is schematic with exaggerated tilt.

Buying assembly you cannot position your way into

Durations are indicative and depend on part mix, tolerance stack and existing fixturing.

PHASE 01 · 0–3 MONTHS

Measure your actual tolerance stack

Not the drawing tolerance — the realised distribution. Measure clearance across a production batch, the repeatability of the fixturing, and the drift of your existing cells across a shift. Most automation business cases in this area are built on nominal figures that the parts do not honour.

PHASE 02 · 3–9 MONTHS

Run the sensor-disabled control

Trial candidate systems on your parts, and ask for the comparison with wrench sensing disabled — run only under bounded speed, force, travel and emergency stop, since disabling feedback under rigid position control can damage parts, tooling or the robot. Where that risk is unacceptable, compare against your incumbent non-force-controlled process instead. Establish surface-finish, insertion-force and tooling-wear baselines during these trials, because there is no baseline to measure against later.

PHASE 03 · 9–18 MONTHS

Design for variable cycle time

Force-guided insertion trades a fixed duration for a distribution. Line balancing, buffering and takt assumptions all have to accommodate a tail. Programmes that specify a fixed cycle time and then discover the variance late usually blame the robot.

PHASE 04 · 18 MONTHS +

Scale on production KPIs, not demo success rates

Continue the wear monitoring begun at pilot, and manage the cell on figures a plant actually runs on: first-pass insertion yield, 95th and 99th percentile cycle time, peak force and torque, insertion-depth verification, surface-damage rate, recovery rate after failed acquisition, false-seating rate and tooling-life distribution.

Organisational readiness

This capability sits between manufacturing engineering, which owns the tolerance stack, and controls, which owns the robot. The characteristic failure is that nobody owns the interface: the tolerance is treated as a given by the controls team and the control strategy is treated as a black box by manufacturing engineering. The most useful early act is to put the realised clearance distribution and the arm's measured repeatability on the same page, in the same units, in front of both.

What goes wrong, and what it costs

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

RISK 01

Nominal tolerances in the business case

Drawing clearance and realised clearance differ, and the realised distribution has a tail. A system specified against the nominal will meet its success rate in the trial and miss it in production.

RISK 02

No sensor-disabled control

Without it, an evaluation cannot separate the contribution of force sensing from that of good fixturing. Published work shows the difference can be ten-to-one, so this is not a subtle confound.

RISK 03

Cycle time treated as fixed

Force-guided insertion produces a distribution of durations, not a constant. Line balance and buffer sizing must absorb the tail or the cell becomes the constraint.

RISK 04

Surface wear from deliberate contact

The strategy works by loading the mating surfaces. On precision features that is a slow degradation of exactly the tolerance the process depends on, and it is rarely instrumented.

RISK 05

Geometry underestimated

Rotationally symmetric parts are markedly easier than square, hexagonal or keyed ones, which add an alignment degree of freedom and more contact states. Do not extrapolate a cylinder result to a connector.

RISK 06

Thermal drift

Calibration achieved cold does not survive a shift. Any strategy that leans on absolute positional accuracy inherits the drift; force-referenced strategies largely do not, which is an underrated part of their value — though sensor bias, fixture and tool expansion and changed initial alignment all remain temperature-dependent.

RISK 07

Research results read as production data

Much of the published evidence uses printed or laboratory parts in controlled conditions. It establishes viability, not throughput, yield or tooling life on metal components at volume.

RISK 08

Where the sensing sits

Wrist-level sensing improves observability of local contact and reduces dependence on upstream dynamics, but joint-torque-controlled arms can deliver effective whole-arm compliance. The right architecture depends on bandwidth, payload, sensor noise and the workspace compliance required — it is a design decision, not a packaging detail.

08 / WHERE TO START

Three questions worth asking this quarter

01 Do we know our realised clearance distribution, or only the drawing tolerance?

02 Where in the machine are we buying precision, and where is it actually needed?

03 Has any vendor shown us the comparison with their sensor switched off?

The third question takes an afternoon and separates a genuine capability from a well-fixtured demonstration.

NEXT IN THIS TRACK

Edition R05 — The 700-Millisecond Problem

Conversational robots and the latency budget of feeling present — why humans turn take in about 200 ms, why the optimal robot response is far slower than that, and why detecting silence is the wrong mechanism entirely.

SOURCES & METHOD

Primary: “Delta6: a low-cost, 6-DOF force-sensing flexible end-effector,” arXiv:2604.06150 · Zhao, Hailegebrial, Wang, Wen and Zheng, “Robotic jamming-free and precise assembly for peg-in-hole operation by multi-DOF force-controlled parallel end-effector,” Robotics and Computer-Integrated Manufacturing, 2025 (SSRN 5186899) · “Bridging the sim-to-real gap with dynamic compliance tuning for industrial insertion,” arXiv:2311.07499.

Contact-state taxonomy and jamming definitions: Royal Society Open Science 6:190476, peg-hole disassembly using active compliance · Zhang et al., jamming analysis and force control for flexible dual peg-in-hole assembly, IEEE TIE 66(3), 2019 · Whitney, quasi-static assembly of compliantly supported rigid parts, 1982 · Zhang et al., arXiv:1905.04427 for the clearance-versus-uncertainty statement.

Method note: the figures cited come from separate published studies addressing the same problem class under different apparatus; they are not a single experiment and are attributed individually. Results are laboratory results on test parts, and establish viability rather than production throughput or tooling life. The published and preprint versions of the parallel end-effector paper state the moment-control error in different units, so that figure is reported qualitatively. Figure 01 is schematic with tilt exaggerated for legibility.

Cross-references Edition R01 on sim-to-real transfer and Edition R02 on contact-rich manipulation.

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 R04 · JULY 2026