softserve

ACHIEVE ROBOTICS EXCELLENCE WITH PHYSICAL AI AND HIGH-FIDELITY VIRTUAL GYMS

How to effectively bridge simulation and reality in robotic training

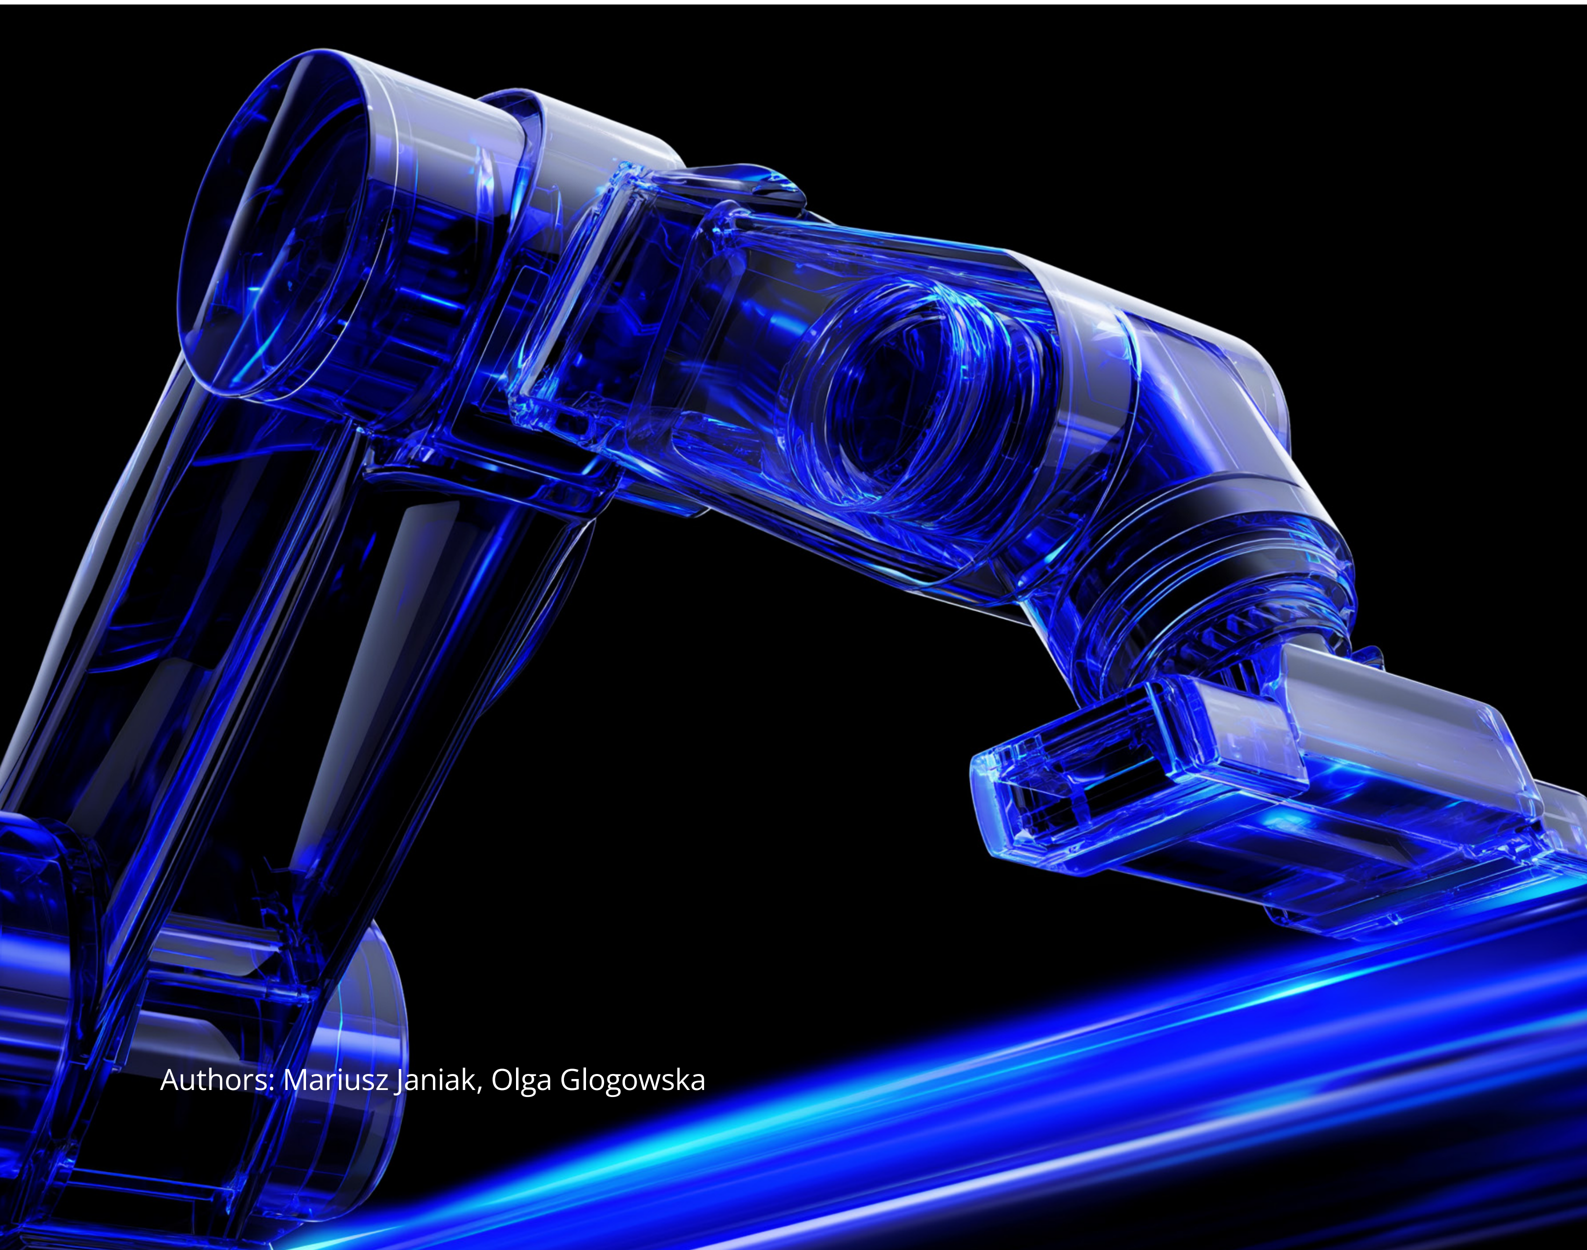

Authors: Mariusz Janiak, Olga Glogowska

Most robotics projects still depend on manual programming and extensive real-world testing. While effective in structured settings, these methods cannot cope with the variability of real industrial environments. When a part is warped, a surface is unexpectedly slippery, or a fixture is slightly misaligned, robots often fail. Such failures cause downtime, rework, and delayed scalability, reducing the return on automation investments. The core issue is that pre-programmed logic cannot handle novelty beyond predefined parameters, resulting in rigid and brittle systems.

Physical AI, combined with **high-fidelity virtual simulation**, introduces a new paradigm. Instead of relying on fixed program rules, robots learn within physics-rich environments that replicate real-world dynamics. This training builds resilience and adaptability before deployment, reducing the need for physical calibration and enabling policies that generalize across related tasks and conditions. By designing and validating robotic behavior in simulation, organizations can cut the cost and risk of prototyping, shorten integration timelines, and scale automation across sites more reliably.

Simulation-to-reality gap in robotics

The simulation-to-reality gap refers to the mismatch between how a robot behaves in a simulated environment and how it performs in the real world. Even highly detailed simulators struggle to capture the full variability of physical systems, which leads to reduced reliability and unpredictable behavior once robots leave controlled virtual settings. As one of the central challenges in robotics, bridging this gap remains the primary driving force for the constant evolution and advancement of control algorithms in robotics.

From deterministic control to adaptive systems

Early industrial robots were governed by classical control and model-based paradigms that relied on explicit mathematical descriptions of motion and force. Techniques such as **linear quadratic regulators (LQR)** and **model predictive control (MPC)** used dynamical models to generate control actions, while simpler feedback schemes such as **proportional–integral–derivative control (PID)** provided practical regulation in well-understood operating regimes. Together, these methods enabled precise motion tracking and stable trajectory execution in structured settings such as assembly lines and fixed manipulation cells, as depicted in Figure 1.

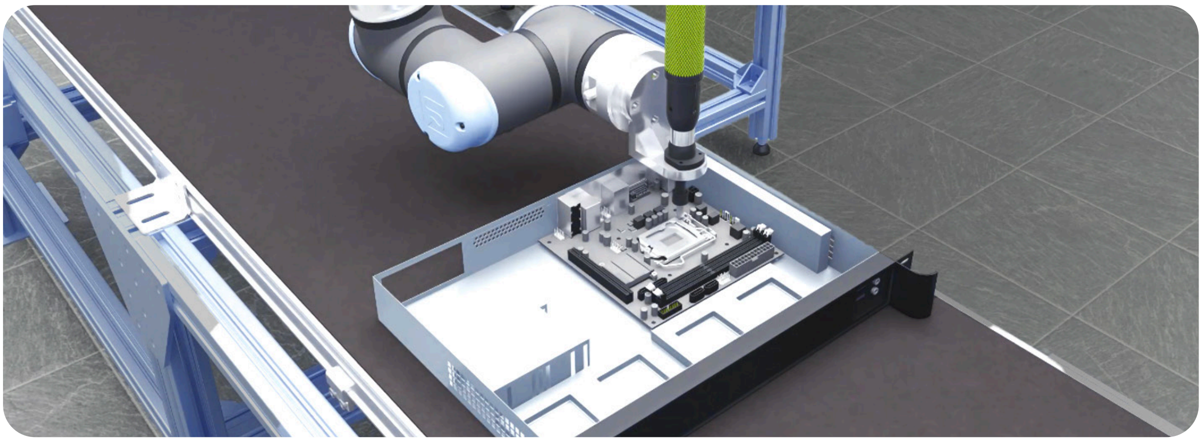

Figure 1. Comprehensive digital twin of a conveyor belt with a robotic arm performing motherboard assembly

As tasks became more complex and environments less constrained, the limitations of these methods became clearer. PID control is purely reactive and does not predict future states, while model-based controllers such as LQR and MPC depend on mathematical models that describe only the dominant and most manageable aspects of a robot's dynamics. Although classical controllers include robustness margins, their performance drops once model inaccuracies exceed those margins. Maintaining accurate models or tuning controllers for every new condition became impractical, which created the need for approaches that could learn directly from data and adapt to uncertainty.

The rise of learning-based control

The limitations of model-based systems opened the door to reinforcement learning (RL), where control policies are learned through trial and error instead of explicit modeling. In RL, an agent interacts with its environment, receives rewards or penalties based on performance, and iteratively improves its behavior.

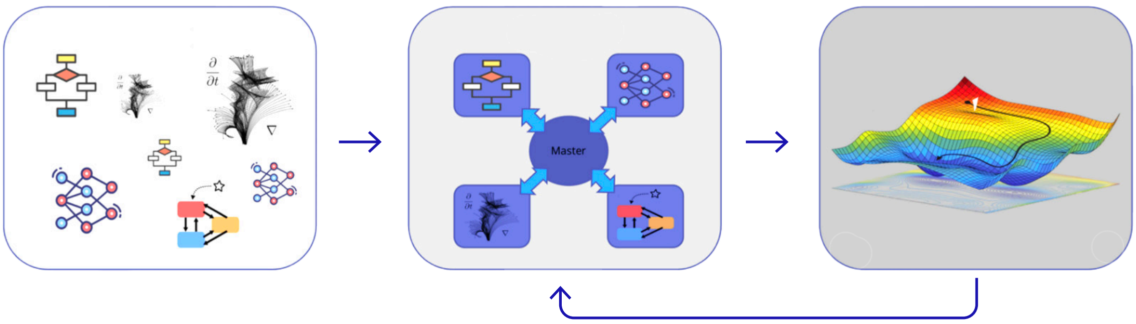

Figure 2. RL policy training as an optimization problem

RL proved capable of mastering complex, high-dimensional tasks such as autonomous navigation or dexterous manipulation (see Figure 3), which were previously intractable for model-based control. This shift was enabled by advances in neural architectures such as deep neural networks and liquid neural networks. These architectures allow RL policies to represent non-linear behaviors and adapt online to changing conditions during training and deployment.

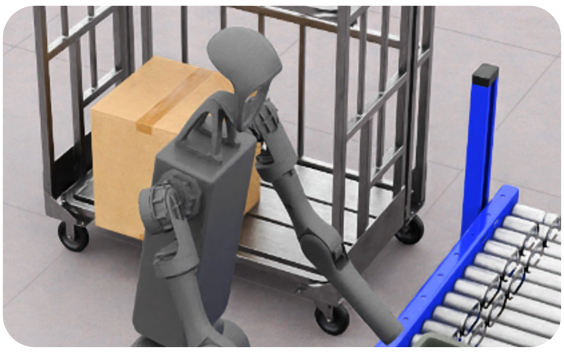

Figure 3. Humanoid robot using reinforcement learning to perform object handling in a brownfield automation environment

However, this flexibility came at a price. RL is **data-hungry** and **sample-inefficient**. Training policies robustly require millions of interactions, which are impossible to conduct safely or economically on real hardware. As a result, simulation became the only practical training ground.

New challenges in the learning era

Imitation learning and inverse reinforcement learning aim to achieve faster and more accurate policy convergence by learning directly from expert demonstrations. These approaches reduce data requirements but rarely capture the full variability of real operations. As a result, synthetic data generation in simulation remains the only scalable source of diverse scenarios.

This introduces a new dependency: the fidelity of the simulated world. Policies trained in simplified or idealized environments often fail when confronted with the noise, uncertainty, and complexity of the physical world. The mechanism and source of the simulation-to-reality gap have fundamentally changed. **The focus has drifted away from perfecting the robot's internal dynamics model to ensuring the fidelity of the external environment model for robust policy generalization.**

Although the shift to RL control policies allowed for better robustness and significantly attenuated the simulation-to-reality gap in robotics, several key issues emerged:

Data bottleneck: Even in simulation, learning from random exploration is inefficient. Collecting sufficient data for robust policies can take days or weeks of simulated time.

Generalization limits: Agents often overfit to the specific characteristics of a simulator's physics engine. Slight differences in contact modeling or sensor response can lead to real-world failure.

Domain randomization: Introducing variations in simulation, for example, in friction, lighting, or geometry, improves robustness but cannot fully replicate real-world uncertainty.

Computational inefficiency: RL training relies on stochastic sampling, making it slow to converge, which hinders scalability for industrial use.

These challenges revealed that data alone cannot close the gap. The quality and realism of simulation matter as much as the quantity of experience.

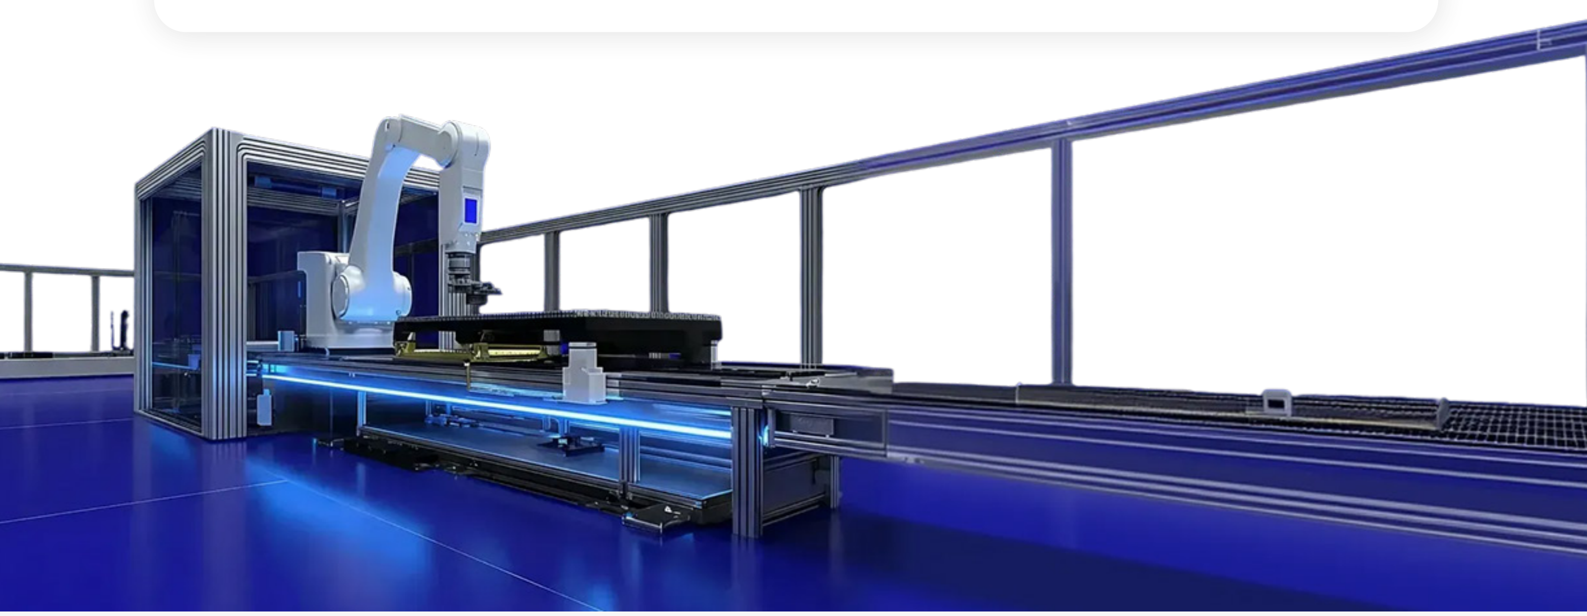

Close the Gap with high-fidelity virtual gyms

As seen in Table 1, below, the evolution of robotic control shows a clear trajectory, from brittle, model-based systems, through expensive imitation learning, toward reinforcement learning supported by simulation and domain randomization. Each step improved adaptability but left critical gaps in data efficiency and real-world robustness. To overcome these limitations, robotics is now moving toward **differentiable and physics-informed world models.**

Differentiable simulation allows gradients to propagate directly through the physics engine, which enables robots to optimize their policies using analytical feedback rather than random sampling. This significantly improves learning efficiency and stability.

At the same time, **hybrid modeling** that combines **first-principle physics with machine learning** offers a way to represent phenomena that are difficult to model explicitly, such as deformable materials, contact friction, or fluid-structure interactions. These developments promise virtual environments where learned strategies remain valid when transferred to real hardware.

High-fidelity virtual gyms provide realistic, multi-physics environments where robots learn and validate control strategies before deployment. They combine three core principles:

- Physically accurate modeling of dynamic interactions
- Scalable, parallelized training
- Differentiable simulation for efficient policy optimization

Together, these capabilities enable the development of **physical AI** — intelligent systems that can be trained virtually and transferred to physical equipment with minimal recalibration.

Paradigm	Strengths	Limitations	Impact on the Simulation-to-Reality Gap
Classical Control (PID, LQR, MPC)	Precise, deterministic, stable for known systems	Requires perfect models; fragile under real-world uncertainty	High – fails when the environment deviates from assumptions
Model-Free Reinforcement Learning	Learns directly from data; adapts to unmodeled dynamics	Extremely data-hungry; prone to overfitting simulator artifacts	Moderate – flexible but inconsistent in reality
Hybrid / Differentiable Simulation	Combines physical accuracy with efficient optimization	Computationally complex; still emerging	Low – most promising path toward closing the gap

Table 1. Evolution of robotic control paradigms

High-fidelity world models

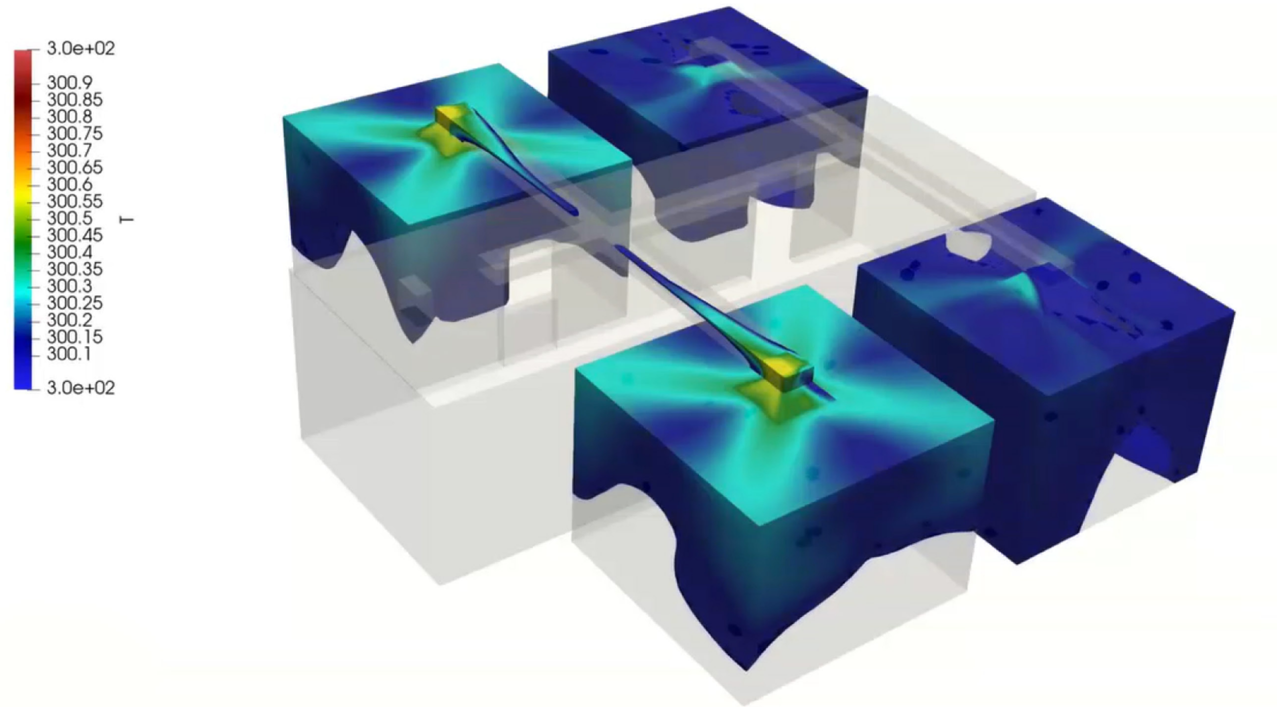

Figure 3. High-Fidelity simulation of HVAC system in a warehouse

To achieve this realism, **hybrid solvers** combine first-principle physics with data-driven residual models tuned from real-world data. This hybridization allows the simulator to represent challenging effects such as granular media behavior, non-linear friction, or soft-body deformation with greater accuracy.

For multi-physics systems, **co-simulation frameworks** integrate specialized solvers into a unified environment, as illustrated in Figure 4. The **functional mock-up interface (FMI)** standard enables interoperability through Functional Mock-up Units (FMUs). Commercial platforms such as **Ansys Twin Builder** extend FMI for strong coupling, while open-source frameworks like **preCICE** offer alternative iterative schemes for robust, tightly coupled multi-physics integration.

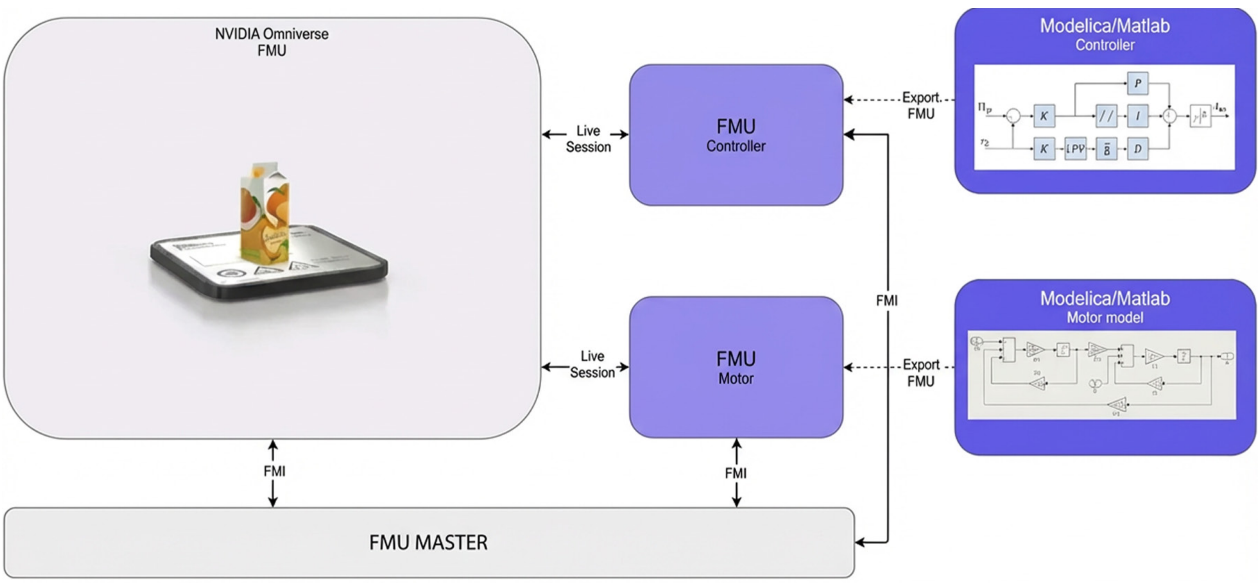

Figure 4. Co-Simulation and High-Fidelity Modelling Framework

These high-fidelity frameworks generate realistic training data, exposing control policies to the same physical subtleties they will face on the factory floor. This realism is the foundation for policies that remain stable and transferable when deployed on real hardware.

Differentiable world models

High fidelity improves realism but can be computationally expensive for large-scale RL. **Differentiable world models**, implemented in frameworks such as **PyTorch** and **JAX**, address this by allowing gradients to flow from the policy objective through the environment. Policies can then be optimized end-to-end with analytical gradients rather than relying only on stochastic sampling.

Within **actor-critic learning**, the **critic** leverages gradient information propagated through the differentiable environment to provide dense, informative feedback to the **actor**. This improves convergence, reduces the number of samples required, and lowers the amount of real-world fine-tuning after simulation.

To manage computational load while preserving essential physics, **AI/ML** surrogates replace expensive modules where appropriate. Options include **neural ordinary differential equations (Neural-ODEs)** and **physics-Informed neural networks (PINNs)**, which incorporate physical structure and conservation properties, as illustrated in Figure 5, as well as reduced-order models (ROMs) that compress CFD or FEM systems to capture dominant modes efficiently. The broader area of **physics-informed machine learning (PIML)** extends these strategies toward generalizable, fast approximations suitable for control optimization.

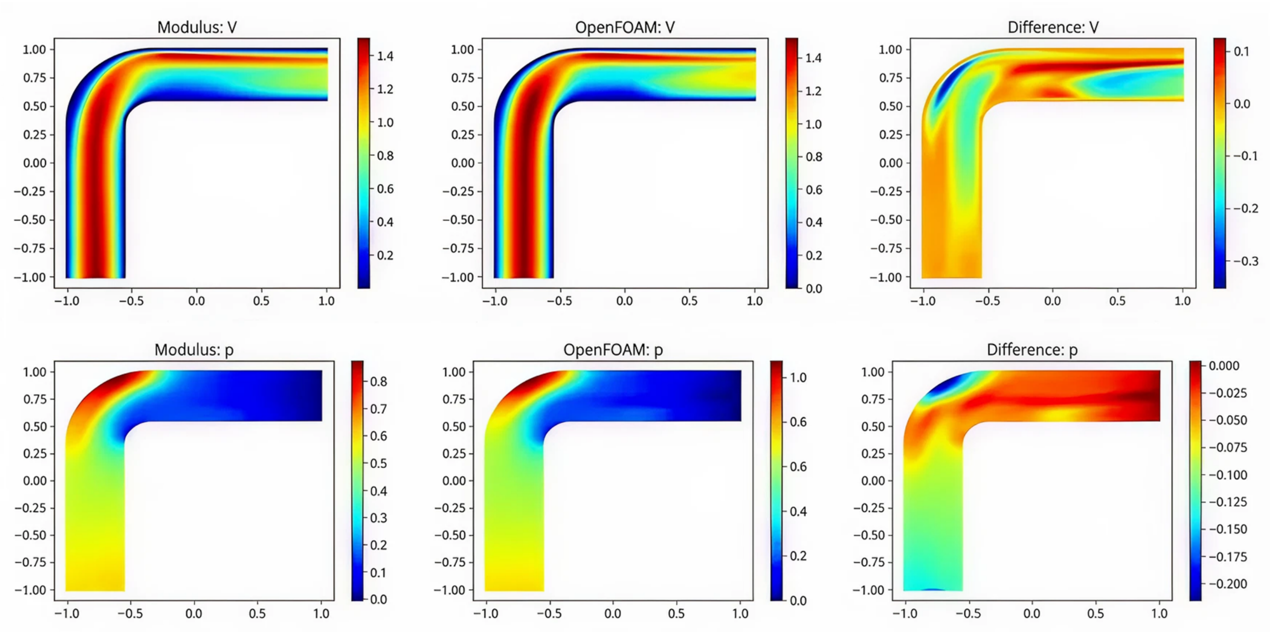

Figure 5. Development of Exhaust System Flow-Thermal Analysis Using PINNs Approach

High-fidelity virtual gyms use these components with parallelized training and domain randomization to expose policies to distributions of dynamics, sensing noise, and geometry. The result is learning that is both physically grounded and computationally efficient, positioning policies for robust transfer from simulation to real equipment.

Production-ready robotics

Deploying learned policies safely and effectively in real operations requires a structured workflow that ensures consistency between simulation and physical systems:

1. **Assess:** Identify high-variance, high-value tasks such as bin-picking of irregular parts, adaptive weld-seam tracking, or visual inspection. Define quantitative success metrics (cycle time, pick success rate, defect detection accuracy).
2. **Model:** Build a high-fidelity digital twin of the work cell using CAD geometry, sensor layouts, and material libraries reflecting real friction, mass, and compliance. The virtual gym can auto-generate a simulated cell using validated material libraries and imported CAD/assets.
3. **Train:** Within the virtual gym, apply curriculum-based reinforcement learning across parallel simulations. Enforce safety constraints such as joint limits, velocity caps, and exclusion zones
4. **Validate:** Run hardware-in-the-loop tests where the physical controller executes the learned policy while telemetry is mirrored in simulation. Confirm dynamic behavior and predictive accuracy.
5. **Deploy:** Containerize the validated policy and distribute it to edge devices. Continuous improvement follows through over-the-air updates that refine policies fleet-wide as new data is collected.

Effective transfer from simulation to the factory floor depends on maintaining alignment between digital and physical domains. The key elements of this process are summarized in Table 2:

Phase	Objective	Core Techniques	Outcome
Training in Physics-Informed Virtual Gym	Build generalizable control policies	High-fidelity simulation with differentiable surrogates; hybrid residual modeling; domain randomization across physics, lighting, and sensors	Robust policies tolerant to unseen variance
Tuning & Calibration	Align digital and real dynamics	System identification loop using telemetry to refine simulation parameters	Accurate virtual-to-physical correlation
Safe Deployment & Progressive Autonomy	Minimize operational risk during rollout	Shadow-mode operation and gradual transfer of control authority from legacy logic to learned policy	Proven reliability and safe handover

Table 2. Phases of effective simulation-to-reality transfer

This integrated workflow enables a systematic, low-risk transition from simulated learning to real-world autonomy while maintaining traceability, and it delivers tangible benefits across engineering, operations, and safety:

Faster iteration: Engineering teams can design, validate, and refine control policies entirely in simulation, eliminating costly physical prototypes and avoiding production stoppages.

Resilient performance: Physics-driven training enables robots to handle production variances, such as SKU changes, fixture misalignments, or environmental disturbances, without manual reprogramming.

Higher quality: Stress-testing policies against millions of virtual edge cases exposes and corrects failure modes before deployment, reducing scrap and rework.

Increased safety: Hazardous tasks, such as chemical handling, high-heat operations, or inspection in confined spaces, can be perfected in simulation, minimizing human exposure and extending equipment life.

Continuous improvement: Integration with CI/CD pipelines allows over-the-air updates to be delivered across fleets, ensuring robots learn collectively and remain aligned with evolving operational goals.

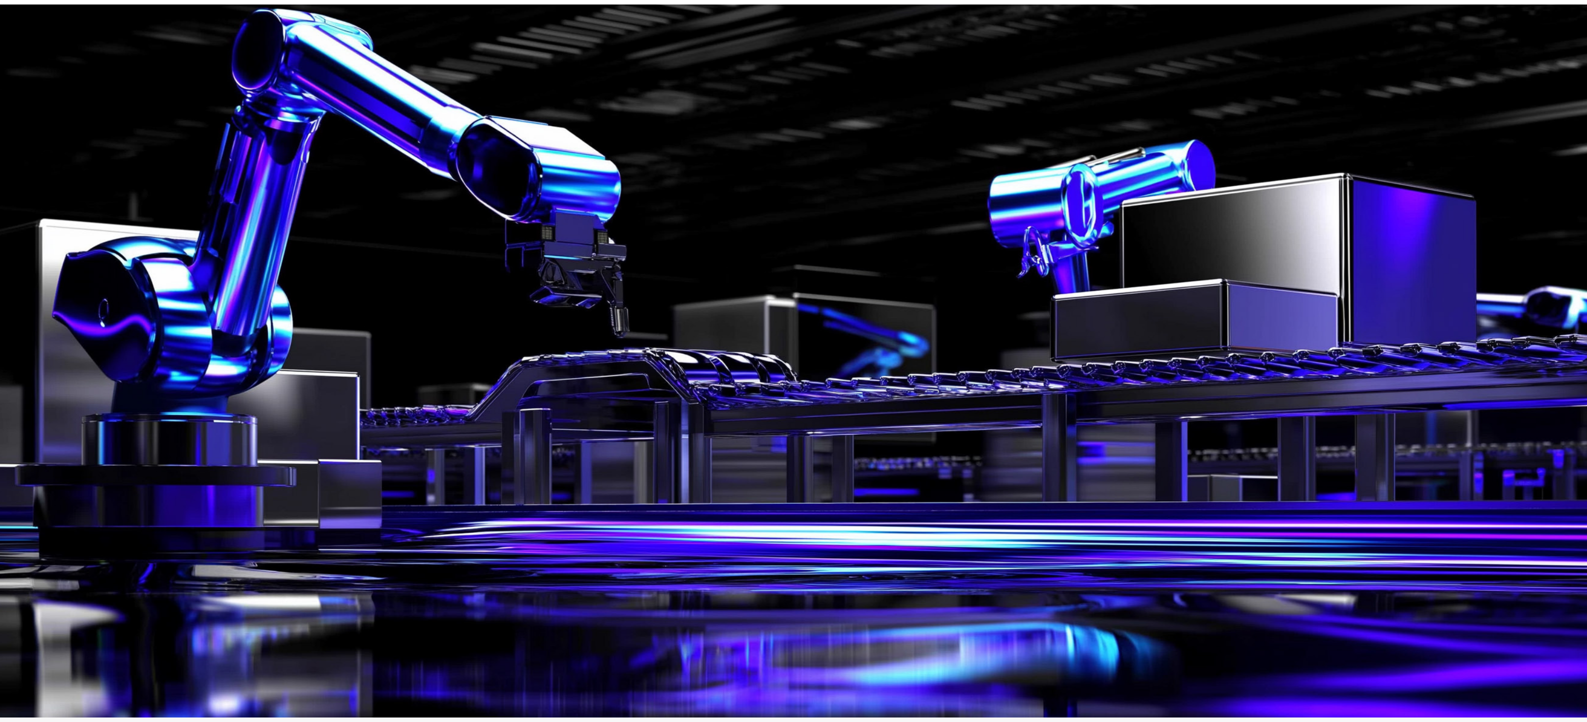

Take the leap:
From virtual gym to real-world physical AI and beyond

SoftServe delivers a **complete pathway** for organizations ready to operationalize **physical AI**. Our approach begins with the creation of **high-fidelity virtual gyms** tailored to each client's specific use case, combining **physics-based simulation** with accurate models of materials, sensors, and equipment. We select and train **AI surrogate models** that accelerate computation while preserving essential physical behavior within digital environments. **Robots are trained** inside such virtual gyms, enabling fast and reliable control policy development and validation.

Finally, SoftServe guides the transfer to physical systems through targeted **fine-tuning in real scenarios**. This **closes the simulation-to-reality gap** and ensures that deployed policies remain stable, predictable, and aligned with operational goals. With this capability in place, organizations gain scalable, adaptable automation that continues to improve throughout its lifecycle.

SoftServe stands ready to help you implement physical AI and bring robust, high-performance robotics into real production environments. Contact us to turn virtual training into lasting operational value.

Contact Us

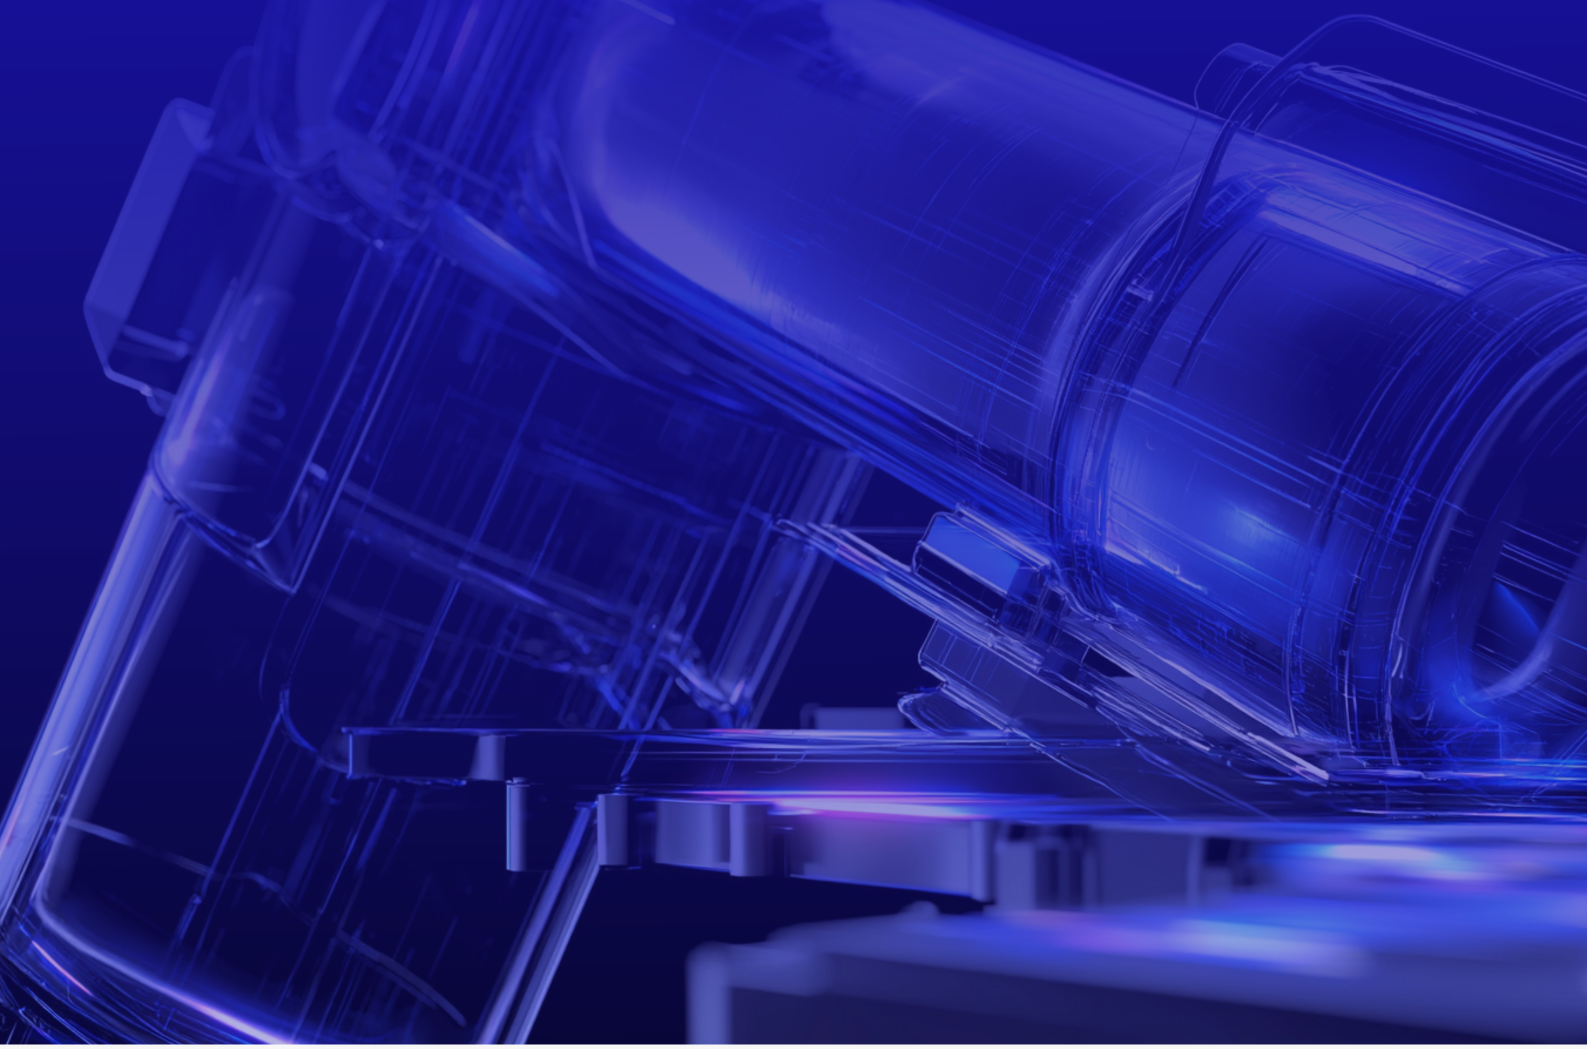

About US

SoftServe is a premier IT consulting and digital services provider.
We expand the horizon of new technologies to solve today's complex business challenges and achieve meaningful outcomes for our clients.
Our boundless curiosity drives us to explore and reimagine the art of the possible. Clients confidently rely on SoftServe to architect and execute mature and innovative capabilities, such as digital engineering, data and analytics, cloud, and AI/ML, robotics and physical AI.

Our global reputation is gained from more than 30 years of experience delivering superior digital solutions at exceptional speed by top-tier engineering talent to enterprise industries, including high tech, financial services, healthcare, life sciences, retail, energy, and manufacturing.
Visit our **website**, **blog**, **LinkedIn**, **Facebook**, and **X (Twitter)** pages for more information.

AUSTIN HQ

201 W. 5th Street, Suite 1550
Austin, TX 78701
+1 866 687 3588 (USA)
Toll Free: +1 866 687 3588

LONDON

30 Cannon Street
London EC4 6XH
United Kingdom
+44 203 807 01 41

info@softserveinc.com
www.softserveinc.com

softserve