LinkTec Labs · Publications

Digital twins, world models, physical AI

Reading notice · V1 · 15 October 2026

Document type and scope: a synthesis of sources and a proposed reference architecture. This is not a peer-reviewed publication, a report of a LinkTec experiment or a client reference.

Scope correction: manuscript wording such as “deploys”, “builds”, “audit-ready”, “sovereign” and similar terms describes architecture objectives requiring project-specific verification. It attests no client deployment, operated equipment, certification or LinkTec result. Working hypotheses are neither completed research nor a commitment to publish.

Quoted quantitative results belong to their authors and evaluation settings. General claims, recommendations and thresholds are not guarantees for every workload. The reading review and checks of principal references are not experimental replication or exhaustive scientific validation.

Execution is on client premises or a cloud selected with the client. A model’s origin, size and openness determine neither execution location nor commercial rights. Versions, licences, access and dependencies require verification. No US-based model API without the client’s request.

Clarifications: digital twins, learned world models and physical-system control are distinct objects. NASA, Genie, Cosmos and V-JEPA systems describe their authors’ research or products, not capabilities already executed by LinkTec. Simulation does not guarantee transfer to reality; assumptions, data, validation and oversight remain necessary. Real-time operation remains experimental at LinkTec. No robot, fleet, owned datacentre or hardware supply is announced. Insufficiently sourced market sizes and forecasts are not treated as verified facts.

linktec.fr/recherche · This notice clarifies the scope of the following manuscript.

LinkTec LABS



RESEARCH NOTE N°05 · APPLIED AI & DATA

Digital twins, world models, **physical AI**

From replicas to learned simulators, and the discipline of deciding before you act.

ABSTRACT

A digital twin began as a synchronized replica of a physical asset, something to observe. The frontier has moved: a *world model* is a learned, generative simulator of an environment, something you can run forward to predict, plan, and train agents "in imagination" before touching reality. NVIDIA frames the destination plainly: physical AI must be trained digitally first, on a digital twin of itself and a digital twin of the world. This note traces the ladder from asset twins to enterprise-process twins to world models to physical AI, explains what a world model actually is, and shows why simulation is becoming the instrument by which robots, fleets, and even business processes are decided before they are deployed. We close with the simulate-to-decide architecture LinkTec Labs applies, and with the thread that ties the whole 2026 series together: model the world, simulate, measure, then act.

01 · FROM REPLICA TO SIMULATOR

The quiet redefinition of the digital twin

For two decades a digital twin meant a faithful, data-synchronized replica of a physical thing, a jet engine, a turbine, a factory line, kept in step with its real counterpart so engineers could monitor it and look back at what happened. That definition is now too small. The asset twin you observe is becoming a *world model* you can interrogate: a learned, generative simulator you run *forward* to ask what would happen next, under an action you have not yet taken.

The destination is stated most plainly by the people building the infrastructure for it. NVIDIA's Cosmos platform opens with a thesis that doubles as a definition of the field: *physical AI must be trained digitally first; it needs a digital twin of itself, the policy model, and a digital twin of the world, the world model*[5]. In other words, before a robot acts in the world, it should rehearse in a simulated one. The replica becomes a rehearsal space.

This is not a niche concern, and it is not only about robots. The digital-twin market, on the order of forty to fifty billion dollars in 2026 and compounding above 30% a year toward several hundred billion by the early 2030s, is being pulled by exactly this shift, with process-level and simulation twins as the fastest-growing segments[9]. The thesis of this note follows:

> A digital twin is no longer something you watch. It is something you ask. The value is not the replica; it is the ability to simulate a decision, and measure it, before you act.

For an organisation, robotics lab or SME alike, that reframing is the whole point, and it is the discipline this series has been building toward.

KEY TERMS

A shared vocabulary for the rest of this note. Specialists may skip ahead.

Digital twin	A virtual replica of a physical asset or process, kept synchronized with it through live data.
Digital twin of an organization	A twin of a business's processes, supply chains and workflows, not just a physical asset (Gartner: DTO).
World model	A learned, generative model of an environment's dynamics that can predict future states and be acted in.
Physical AI	AI that perceives and acts in the physical world: robots, autonomous vehicles, industrial systems.
Sim-to-real	

	Training or planning in simulation, then transferring the learned behaviour to the real world.
Synthetic data	Data generated in simulation to train models where real data is scarce, costly or unsafe.
JEPA	Joint-Embedding Predictive Architecture: a world model that predicts in a latent space rather than pixels.
Model-predictive control	Planning actions by simulating their outcomes with a model and choosing the best, step by step.
Training in imagination	Letting an agent learn inside a generated/simulated world instead of the costly real one.
Sim-to-real gap	The fidelity difference between simulation and reality that limits naive transfer.

02 · THE LADDER

Asset twin, process twin, world model, physical AI

The field is best read as a ladder of increasing capability, from systems that *observe* to systems that *act*. At the bottom sits the classic **asset twin**: a sensor-fed replica of a single machine, used for monitoring and predictive maintenance. Above it, the **process or enterprise twin**, what Gartner calls a Digital Twin of an Organization, models entire workflows, supply chains and customer journeys, so that planners can simulate a change before committing to it[9]. Higher still, the **world model** generalizes from a specific replica to a learned simulator of dynamics, able to predict states it was never explicitly shown. At the top, **physical AI** closes the loop: an agent that has rehearsed in the world model now acts in the real world[5].

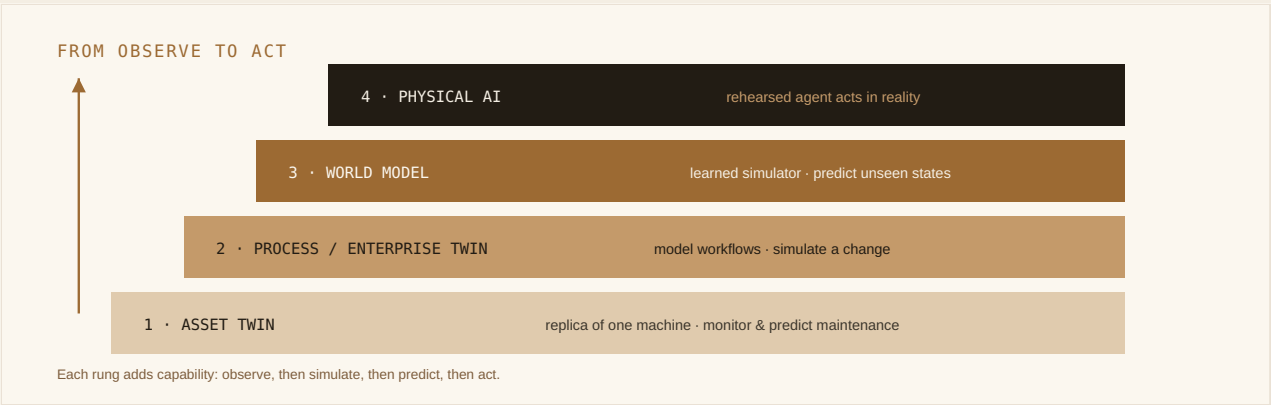


Fig. 1 · The maturity ladder. Asset twins observe; process twins simulate a change; world models predict unseen states; physical AI acts on what it rehearsed[5][9].

03 · WHAT A WORLD MODEL ACTUALLY IS

A learned simulator you can act in

A world model is, at heart, a model of *dynamics*: given the current state and a candidate action, it predicts the next state. That single capability unlocks the rest, because an agent that can predict outcomes can *plan* by trying actions in its head, and can be trained "in imagination" inside the model rather than in the costly, slow, sometimes dangerous real world. The lineage is clear. Ha and Schmidhuber showed an agent could learn a compact generative model of its environment and train largely inside it[2]. DeepMind's Genie, an 11-billion-parameter foundation world model learned from unlabelled internet video, generates action-controllable virtual worlds a user or agent can step through frame by frame[3]; Genie 3 pushes this to real-time, photorealistic, persistent 3D environments with emergent physics and object permanence[4].

Two technical families now coexist. **Generative world models** (Genie, NVIDIA Cosmos) synthesize future video or sensor streams and let agents train against them[3][5]. **Latent-predictive world models** (the JEPA line) predict in an abstract representation space rather than reconstructing pixels, which is more efficient and, its proponents argue, closer to how prediction should work; Meta's V-JEPA 2 learns from over a million hours of video to support understanding, prediction and planning[7][8]. The distinction matters in practice, but the shared idea is the engine of the field: a model you can roll forward to see what happens next.

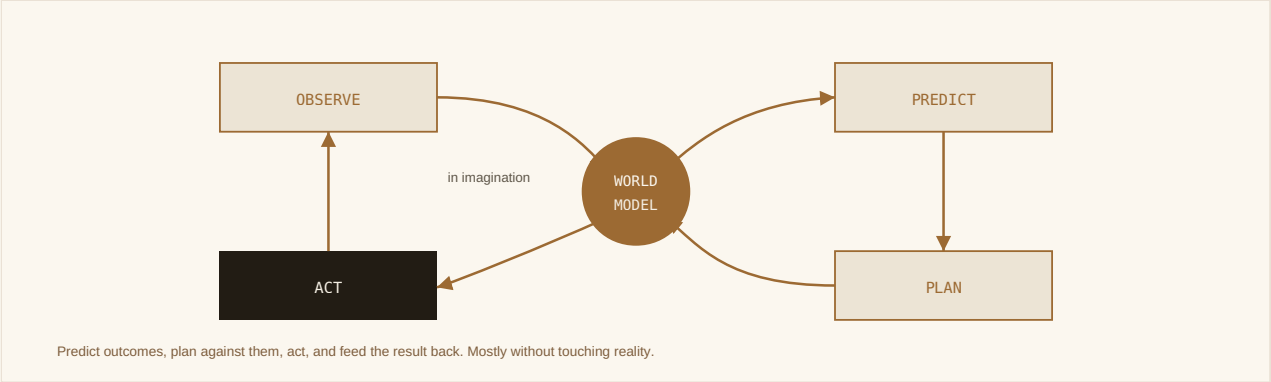


Fig. 2 · The world-model loop. From an observation, the model predicts outcomes of candidate actions and plans against them "in imagination", so only the chosen action reaches reality[2][3][7].

04 · WHY THIS CHANGES PHYSICAL AI

Rehearse in simulation, transfer to reality

The payoff lands hardest in robotics, where real-world data is slow, expensive and risky to collect. World models let a robot rehearse in simulation and then transfer, the *sim-to-real* pipeline, with results that were implausible two years ago. Meta's V-JEPA 2-AC is the sharpest illustration: after pretraining on internet video, it is post-trained on just **62 hours** of robot data, and then plans pick-and-place actions on Franka arms in labs it never saw, zero-shot, with no task-specific training and no reward engineering, using the world model for model-predictive control[7]. NVIDIA's Cosmos generates the synthetic data and simulated worlds that train robots and autonomous vehicles at scale, and Cosmos 3

extends this to omnimodal models that jointly handle language, image, video, audio and action[5][6]. Generative world models like Genie let agents be trained inside imagined environments before any hardware moves[3][4].

62 h

Of robot data sufficient for V-JEPA 2 to plan zero-shot pick-and-place on unseen arms[7].

24 fps

Real-time, photorealistic, persistent interactive worlds generated by Genie 3[4].

~$379B

Projected simulation-digital-twin software & services market by 2034, from ~$35B in 2024[9].

The principle underneath all of it is the same, and it is older than deep learning: it is cheaper, faster and safer to fail in simulation than in reality. World models industrialize that principle.

05 · THE ENTERPRISE BRIDGE

Simulate to decide, not only to demo

Most organisations are not building robots, and might conclude this frontier is not theirs. That conclusion is wrong, and the reason is the through-line of this note: the discipline is not "build a robot", it is "simulate a decision before you commit to it". A digital twin of a business process, an order pipeline, a support workflow, a logistics route, is a place to ask the same question a robot asks, what happens if I take this action, before paying the real-world cost of finding out. Model the process, calibrate the twin on real data, simulate the candidate change, measure the expected ROI and risk off-production, and only then deploy.

> WHY THIS IS THE MATURE MOVE
>
> A simulation that merely renders a process is a demo. A simulation you can *act in* to compare decisions, with measured outcomes, is an instrument. The difference between the two is the difference between a dashboard and a decision, and it is where the value of a twin actually accrues[9].

This is also where the LinkTec Labs series closes its loop. A useful process twin is built on the same substrate the earlier notes described: a context layer that retrieves the right state (Note N°01), a memory that carries it across time (Note N°02), governed and sovereign so it can run on real operational data (Note N°03), and right-sized models so the simulation is affordable to run at volume (Note N°04). World models are not a separate topic bolted on at the end. They are what the rest of the architecture lets you finally do: decide before you act.

06 · WHERE IT BREAKS

The fidelity gap and the honesty it demands

A world model is a prediction, not an oracle, and treating it as one is the characteristic failure of this technology. The *sim-to-real gap* is real: a model can hallucinate plausible-but-wrong physics, miss rare

events entirely, and grow confidently inaccurate the further it extrapolates from its training distribution. The cost of running high-fidelity world models is non-trivial, and a twin calibrated on yesterday's data drifts as the real system changes. None of this argues against the approach; it argues for the same evidence-driven discipline the rest of this series insists on. A simulated decision is a hypothesis. It is validated against reality, measured after deployment, and the twin is recalibrated when it drifts. Simulation earns trust the same way retrieval and memory do, by being measured, not by being believed.

07 · A REFERENCE APPROACH

The simulate-to-decide loop, as LinkTec Labs applies it

The pieces compose into a closed loop whose output is a decision with evidence behind it, not a rendering. Model the asset or process; calibrate the twin against real data so its predictions are grounded; simulate the candidate actions in the world model; measure expected ROI and risk off-production; decide; deploy the chosen action; and feed the realized outcome back to recalibrate. The loop is deliberately humble: the world model proposes, reality disposes, and the gap between them is the signal that keeps the twin honest.

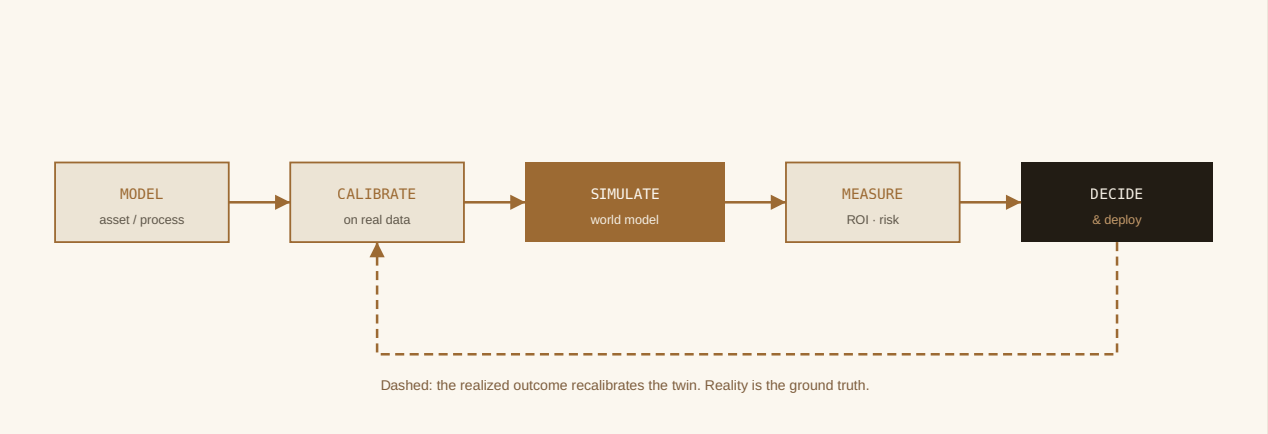


Fig. 3 · LinkTec Labs simulate-to-decide loop. Model and calibrate a twin, simulate candidate actions in a world model, measure ROI and risk, decide, deploy, and feed the realized outcome back to keep the twin grounded.

Four principles govern it. **The twin is calibrated, not assumed**, its predictions are anchored to measured reality before anyone trusts them. **Simulation is for deciding**, the output is a ranked set of actions with expected outcomes, not a pretty render. **Reality is the ground truth**, every simulated decision is checked against what actually happened, and the gap drives recalibration. And **it runs on the same governed, right-sized stack** as the rest of the series, so it is affordable, sovereign and auditable. The first three earn accuracy; the fourth makes it deployable.

08 · PRACTITIONER'S PLAYBOOK

What to do on Monday

- → **Stop treating the twin as a dashboard.** Build it to answer "what if I act", not only "what is happening"[9].
- → **Climb the ladder deliberately.** Asset twin, then process twin, then world model; do not skip calibration[5].
- → **Use simulation to fail cheaply.** Rehearse the costly or risky action in the model first[2][7].
- → **Exploit sim-to-real for data scarcity.** Synthetic data and rehearsal beat waiting for rare real events[5][6].
- → **Calibrate, then trust, then re-calibrate.** A twin grounded on stale data is a confident liar.
- → **Measure the gap.** The difference between predicted and realized outcomes is your most valuable signal.
- → **Right-size and govern the simulation.** It must be affordable to run and safe to run on real data (Notes N°03, N°04).
- → **Decide off-production.** Move the experiment into the twin so production carries only the chosen action.

A digital twin earns its keep the moment it stops being a mirror and becomes a question. Simulate the decision, measure the gap, then act. That is the whole discipline, and the whole series.

LinkTec Labs builds simulate-to-decide twins on a governed, sovereign, right-sized stack, drawing on a research culture in robotics and AI to treat simulation as a decision instrument, not a demo. This note closes our 2026 research series: context (N°01), memory (N°02), sovereignty (N°03), efficiency (N°04), and now simulation. **linktec.fr / labs**

REFERENCES

Primary & technical sources

[1] Glaessgen & Stargel. **The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles.** AIAA/NASA, 2012 (foundational digital-twin reference), https://ntrs.nasa.gov/citations/20120008178

[2] Ha & Schmidhuber. **World Models.** NeurIPS 2018. arXiv:1803.10122, https://arxiv.org/abs/1803.10122

[3] Bruce et al. (Google DeepMind). **Genie: Generative Interactive Environments.** ICML 2024. arXiv:2402.15391, https://arxiv.org/abs/2402.15391

[4] Google DeepMind. **Genie 3: a new frontier for world models** (real-time interactive environments). 2025, https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/

[5] NVIDIA. **Cosmos World Foundation Model Platform for Physical AI.** 2025. arXiv:2501.03575, https://arxiv.org/abs/2501.03575

[6] NVIDIA. **Cosmos 3: Omnimodal World Models for Physical AI** (open frontier foundation model). 2026, https://www.nvidia.com/en-us/ai/cosmos/

[7] Assran et al. (Meta AI / FAIR). **V-JEPA 2: Self-Supervised Video Models Enable Understanding, Prediction and Planning.** 2025. arXiv:2506.09985, https://arxiv.org/abs/2506.09985

[8] LeCun. **A Path Towards Autonomous Machine Intelligence** (JEPA position paper). 2022, https://openreview.net/pdf?id=BZ5a1r-kVsf

[9] Gartner / industry analyses. **Digital Twin of an Organization & digital-twin market sizing** (2025-2026; simulation-twin revenue projections). Secondary market data, attributed in-text.

© 2026 LinkTec, LinkTec Labs. Research Note N°05. Figures are original schematics by LinkTec Labs, after the cited studies. Quantitative results are attributed to their primary sources; market sizes and vendor figures are indicated as such in-text. This note is informational and does not constitute a benchmark or investment claim by LinkTec.