

# The Missing Layer of Physical AI

*Why the next AI transition will run on habitat, not models.*

There is a building in Tel Aviv that does not look like the beginning of a theory.

Eight hundred square meters. Concrete floors. High ceilings. Mediocre parking. Near a Honda motorbike dealership. Inside, for almost a decade, a handful of early-stage startups shared desks, whiteboards, and a kitchen with people from some of the world's largest corporations, companies whose combined revenue exceeds what most nations produce in a year.

Engineers from Tokyo, Paris, and Atlanta sat at the same tables as founders who had closed their first real round the month before. They ate the same local lunch. They tried to build something together that neither side could build alone.

My partners and I built that room around an idea that sounds simple and turns out to be brutally hard: help some of the largest companies on earth find the startups that could one day transform their operations, and help those agile, stubborn, often brilliant startups survive the meat grinder of deploying real technology inside organizations engineered, from the ground up, not to break.

We came to think of it as humanizing the dance between the elephant and the mouse.

Inside that dance, slowly, case by case — failure after failure, with the occasional, hard-won success — one of the most important questions in AI revealed itself to me.

At first it did not look like a question about civilization. It looked like the daily work of helping startups and large companies work together.

A startup would arrive with a powerful technology. A large industrial company would arrive with a real problem, not a fake innovation-lab problem, a real one. Trucks breaking down too often. Factory yield slipping. Fleet utilization stuck. Energy use rising. Safety margins under pressure. Customers waiting. Operators tired.

The ingredients were often better than people assume. The startup was talented. The technology was promising. The corporate champion was there and dead serious. The business case was not imaginary. The data existed. The pilot was approved.

And many times, the pilot even worked.

### **That was the strange part.**

The model performed. The demo impressed. The use case made sense. The meeting ended with momentum, and not naive momentum. People were seeing something real.

Then, somewhere between pilot and production, the thing weakened.

Sometimes it moved into another committee and never came out. Sometimes it worked at one site and failed at the next. Sometimes operators did not trust it enough to use it when it mattered. Sometimes legal, safety, compliance, IT, procurement, and the business unit each needed one more answer before letting it touch the real operation. Sometimes the startup kept improving the model while the customer kept waiting for proof that could not be produced without deeper deployment.

Each failure had an explanation. The product was too early. The customer was not ready. The data was messy. The integration was harder than expected. The champion left. The ROI was real but the organization could not absorb the change. The timing was wrong.

Each explanation was partly true.

But after you hear a different explanation for the same pattern enough times, you begin to suspect the explanations are not the thing. They are the shadow of the thing.

What I eventually saw was this: each side was building part of a bridge. The startup was building capability. The corporate was providing the world in which that capability had to work. But the bridge itself was no one's job.

More specifically, in AI deployments, the missing layer was the operating environment around the AI: the trust, permissions, governance, evidence, auditability, and learning structure that lets intelligence act safely and repeatedly in the real world.

## **I call that layer the habitat.**

This essay is the short version of the argument. The longer version lives in a book I wrote first for the partners who helped me see the pattern, *The Case for Habitat*. The book is not a product launch and it is not a victory lap. It is the story of a discovery I was lucky enough to stand close to.

*My core belief is simple: the next AI transition will not be won by the smartest model alone. It will be won by the habitat that lets intelligence act safely, economically, and repeatedly in the real world.*

## **The model is not the deployment**

The AI debate is still mostly organized around models.

Bigger models. Better models. World models. Reasoning models. Robotics foundation models. Agents. The story is powerful and, in many ways, correct: intelligence is improving quickly, and many things that looked impossible a few years ago now look inevitable.

You can see the conviction in where the money goes. Yann LeCun, one of the most important figures in the history of modern AI, has reportedly raised more than \$1 billion for AMI Labs, a new venture built around world models and AI that can reason about the physical world. At the same time, the largest labs in the world are making unprecedented capital commitments to the conviction that scaling will continue to unlock capability.

It is a real debate, between serious people, over staggering sums.

For most of the past decade, I have watched what happens after the model is already good enough, and I have come to think both sides may be seeing only half the deployment problem. Not because the models do not matter. They do, and they are becoming astonishing. But because the physical world is not waiting for intelligence alone.

The physical world does not deploy models. It admits systems that can sense, decide, act, stay safe, learn from consequence, and remain trusted over time.

A model may be the brain. But a brain is not an organism.

A robot in a lab is not yet a deployable system. A predictive model is not yet an operating capability. A fleet algorithm that predicts degradation is not yet an enterprise asset. A warehouse AI that gives excellent recommendations is not yet something a manager will trust with action at three in the morning.

**Deployment begins when the institution gives the system permission to matter.**

A model can be impressive without being trusted. Accurate without being authorized. Useful without being allowed to act. It can be technically strong while the institution around it remains too thin to absorb the consequences of its strength.

This is why so many AI pilots look better than the production systems they become. A pilot can be held together by special attention, narrow scope, friendly data, heroic people, and temporary permission. Production removes the scaffolding. The environment shifts. The operator gets busy. The second plant has different lighting. The fleet has a different duty cycle. The regulator asks harder questions.

What looked like a model problem becomes something else: a trust problem, a permission problem, a governance problem, an evidence problem, a transfer problem, a consequence problem.

**It becomes a habitat problem.**

## **Digitization gave us memory. The next transition gives us anticipation**

Every major infrastructure transition changes our relationship with something basic.

Electrification changed our relationship with energy: power no longer had to be produced where it was consumed. Telecommunications changed our relationship with distance: coordination no longer depended on proximity. Digitization changed our relationship with memory: the world became recorded, searchable, simulated, and processed, and organizations began to remember more than any human inside them could.

The transition now forming is different.

Physical AI changes our relationship with time.

Not clock time. Consequence time.

A digitized fleet knows where every vehicle is and what happened yesterday. An intelligentized fleet knows which vehicle is likely to fail tomorrow, what

intervention tonight will prevent it, and whether the system has earned enough trust to act before the breakdown becomes real.

A digitized factory records defects. An intelligentized factory sees the process drift before the batch is ruined, then adjusts within the limits it has been allowed to control.

A digitized hospital stores vital signs. An intelligentized hospital detects the pattern before deterioration becomes emergency.

**Digitization gave the world memory. Intelligentization gives it anticipation.**

That may sound like a technical upgrade. It is more than that. Once a system can see a consequence forming before it becomes irreversible, the question is no longer only whether it can predict. The question is whether it is allowed to act.

That is the heart of Physical AI.

## **The moment before consequence hardens**

In the physical world, many events give us a small gift before they become facts.

A bearing vibrates differently before it fails. A patient changes before collapse. A vehicle enters a collision path before impact. A grid begins to destabilize before blackout. A factory process drifts before the defect becomes scrap.

The consequence has not yet happened, but it has started to form.

I call this the consequence-hardening window: the interval between the earliest detectable signal that a consequence is forming and the moment that consequence becomes difficult or impossible to undo.

That window is where the value of Physical AI lives.

It is also where the danger lives.

If a system can act inside the window, it can prevent harm, reduce waste, increase uptime, protect people, and create operating capacity that human-only systems could never reach. But if it acts wrongly, the consequence may be irreversible. If it acts without trust, the institution may reject it. If it learns from bad signals, the next action may be worse. If no one can reconstruct what happened, the system may lose the right to operate even when the model itself was strong.

The physical world imposes a discipline software alone never had to carry.

The world does not wait: action must happen in time.

The world does not always forgive: some consequences cannot be reversed.

The world does not cooperate: environments change, degrade, surprise, and break assumptions.

And institutions decide whether the system may continue. Customers, regulators, operators, insurers, boards, and the public all hold veto power over machine authority.

This is why Physical AI is not simply software entering hardware. It is intelligence entering consequence.

And consequence requires an operating layer.

## **What the habitat is**

A habitat is the operating layer around an intelligent system that lets it act, learn, and earn trust in the physical world.

That can sound abstract, so let me make it simple. A habitat answers five ordinary questions that become extraordinary the moment a machine starts acting under real consequence.

## **What is the job, exactly?**

Not "detect defects" or "optimize maintenance," but the real challenge: what counts as success, what must never happen, which conditions matter, which edge cases require escalation, and what outcome the institution is actually trying to improve.

## **How does the system earn trust?**

Not by jumping from demo to autonomy, but step by step. It starts narrow. It proves itself. It earns more authority. It can lose authority. And if it loses trust, regaining it should be harder than earning it the first time, because institutions remember harm.

## **What is it allowed to touch?**

Which machines, vehicles, patients, routes, shifts, decisions, and thresholds are inside its scope? What must it escalate? What is forbidden?

## **Who watches while it works?**

Not a ceremonial review six months later, but live: cadence, oversight, thresholds, exception handling, and named accountability.

## **What record does it leave behind?**

Every meaningful action, override, near miss, failure, and authority change must become inspectable evidence.

These are humble things. They do not sound like the frontier of AI. That is exactly why they are easy to miss.

But together they change everything.

Below the habitat threshold, the AI produces outputs: alerts, predictions, classifications, recommendations. Some are useful. Many are impressive. But the institution does not necessarily become stronger.

Above the habitat threshold, outputs become evidence. A prediction becomes a documented claim. An action becomes an attributable event. A failure becomes a reusable lesson. An override becomes one of the most valuable signals in the system. One deployment begins to teach the next.

**That is when AI stops being a service the company consumes and starts becoming an institutional asset the company builds.**

## **The twin deficit**

In ordinary software, product-market fit is mostly a relationship between product and customer need.

In Physical AI, that is not enough.

The product can be real, the need can be real, the economics can be real, and the deployment can still fail, because the habitat is not ready for the product and the product is not yet mature enough for the habitat.

Both sides can be telling the truth.

The startup says, "the customer was not ready."

The customer says, "the product did not work."

Each is right, and each is incomplete.

I call this the Twin Deficit: a state in which agent capability and habitat readiness are both below the threshold the deployment requires, and each side's weakness hides the other's.

The model needs better field evidence, but the institution has not built the governance that would let it safely collect that evidence. The corporate wants proof before granting more authority, but the system needs bounded authority before it can produce the proof. The startup needs deeper access, but the

environment will not open until trust is earned. Trust cannot be earned because the evidence is not structured. The evidence is not structured because the habitat was never built.

The collaboration spins.

Everyone is active. Very little compounds.

The escape is not to blame one side. It is to mature the pairing.

The model learns the world. The habitat learns the model. The institution learns how much authority the system has earned. The evidence gets stronger. The permission expands. The next site starts warmer than the last.

That is the beginning of compounding.

## **The learning problem the physical world actually poses**

Predictive AI learns from history.

Reinforcement learning learns through exploration.

Physical AI cannot rely on either alone. History is incomplete, because the physical world keeps producing combinations that were never in the data. Exploration is dangerous, because the system cannot freely break trucks, injure workers, disrupt grids, or risk patients in order to learn.

So the real problem is this:

How does a system learn from consequence without losing the right to act?

That question changed the way I think about AI.

In real deployments, governance is not a safety wrapper bolted around learning. It shapes the learning itself. Authority determines what the system is allowed to

experience. Trust determines the richness of the consequence it sees. Evidence determines whether the next action is permitted.

The system acts within a bounded envelope. The world responds. The consequence is captured. The result becomes evidence. The evidence writes back into the habitat. The improved habitat supports broader, or narrower, authority next time.

**Trust permits action. Action elicits consequence. Consequence becomes evidence. Evidence improves the habitat. The improved habitat permits better action. That is the flywheel.**

I watched it turn once in a way I have never forgotten.

A startup had built a battery-health system for a fleet of electric delivery vehicles. They did not try to automate the depot on day one. For three months the system only advised, and the managers often ignored it. But the team logged every recommendation, every override, and every reason.

Those overrides were not noise. They were the curriculum.

By month four, the system was trusted to make small charging adjustments on its own. By month eight, it was catching faults that managers with twenty years of instinct had never been able to see. Not because the model suddenly became magical. Because the system had finally been allowed to act within bounds, and acting in the real world taught it what no dataset could.

A second depot opened later and reached in six weeks what the first had taken eight months to learn, because the hard-won lessons of the first finally had somewhere to live.

That is what compounding looks like.

It did not come from a better model alone. It came from a better habitat.

## **Where the compounding lands**

The dominant AI business model still looks too much like software-as-a-service: a vendor owns the capability, a customer rents access, and the customer pays for outputs, licenses, seats, or integration.

That worked for an earlier software era. It is a poor fit for Physical AI.

When AI acts inside a customer's institution, the most valuable thing produced is not only the immediate service. It is the operating evidence accumulated through use: what worked, what failed, which conditions mattered, which authority was earned, which actions were reversed, which near misses were prevented, which patterns transferred to the next site.

If that evidence leaves the institution and compounds only inside the vendor, the corporation may be financing someone else's moat.

The habitat changes the economic shape.

At first, the company gets efficiency. Later, if the habitat is built correctly, it gets something more valuable: an institutional memory of how intelligent action behaves under real consequence.

The first wave is service: productivity, uptime, safety, efficiency.

The second wave is asset: the institution becomes more capable because evidence, governance, and learning accumulate in a form it can reuse.

The third wave is franchise: once the habitat holds enough calibrated evidence, it can underwrite new outcome-based offerings around prevention, reliability, uptime, and capability that competitors without comparable history cannot honestly match.

The prize is not simply automating existing work. It is moving action upstream, into the moment before consequence hardens, and turning that moment into a governed, compounding economic substrate.

If Intelligentization is the transition that gives the physical world anticipation, the habitat is the layer that makes anticipation governable.

That is why I believe it may become one of the defining infrastructure layers of the next decade.

## **The category I did not know I was looking for**

I did not begin with this vocabulary.

I began with a practical question: why do so many logical, well-funded, well-intentioned collaborations fail after the technology appears to work?

Over time the answer became clearer. We had words for models, pilots, data, integration, risk, automation, digital transformation, and AI strategy. We did not have a precise enough word for the shared operating layer that lets intelligent systems earn the right to act.

So I named it.

The habitat.

Once it had a name, the rest found its place: a way to describe the constraints the physical world will not negotiate, the ladder by which a system earns and loses authority, the threshold below which outputs never become institutional evidence, and the learning engine in which consequence, under permission, becomes the real teacher.

The formal version of all of it now lives in a research series on the public archives, for anyone who wants the machinery.

But the heart of it is simple.

AI is leaving the screen and entering the world. Once it acts in the world, intelligence is no longer judged only by output quality. It is judged by governed consequence.

That is the shift.

That is the category.

And that is why the next AI transition will not be won by the smartest model alone. It will be won by the habitat that lets intelligence act safely, economically, and repeatedly in the real world.

## **An invitation**

I wrote *The Case for Habitat* first for the partners who lived this laboratory with me: the people who sat in the rooms, took the risks, opened the doors, ran the pilots, argued through the friction, and helped me see that the missing layer was not a metaphor.

It was the thing we kept needing and did not yet have a name for.

The Founder's Edition is now circulating privately among builders, operators, investors, regulators, researchers, and corporate leaders working on Physical AI.

If you are working on Physical AI deployment, governance, investment, or industrial transformation, you are welcome to request a copy.

It is not for sale. It is a gift, sent by request, and I read the requests myself.

The book is the long version.

This essay is the invitation.

---

*The Case for Habitat circulates privately among the people building, governing, funding, and studying Physical AI. If that is your work, you are welcome to one.*

**Request the Founder's Edition** → [tal@talcohen.ai](mailto:tal@talcohen.ai)

*Tell me your name, your role, and one line about what you are building or trying to understand. I read every request myself.*

*Tal*