

START HERE

The Fourth Transition

A plain-language guide to why AI, like electricity, the telegraph, and the computer before it, will underdeliver until we build the missing layer around it.

June 6, 2026

The most capable software ever built is being handed to the largest companies in the world, and inside many of them it is quietly letting people down. The demo impresses. The pilot often works. Then, somewhere between the pilot and the real operation, the thing weakens and fades.

The easy explanation is that the model is not smart enough yet. I spent most of the last decade helping startups and large companies try to put new technology to work together, watching what happens after the model is already good enough, and I came to think the easy explanation is wrong. In most cases the model is not the problem. The problem is that the world around the model has not been rebuilt to fit it.

That is not a guess. It has happened three times before.

The same thing has happened three times

Every lifetime or two, a capability arrives that is so general it does not improve one corner of life, it seeps into all of them. It changes how we work, where we live, how fast news travels, and what a company even is. Historians sometimes call these civilizational transitions. There have been three in the last two centuries, and each one set something free.

The transition	What it set free	The missing layer it needed
Electrification	Energy	The power grid
Communication	Distance	Shared protocols
Digitization	Memory	The data layer

And each one arrived the same way, in three beats. The clearest place to see them is a factory.

Picture a factory in the year 1900. It is powered by one giant steam engine in the basement, turning a long iron shaft that runs the length of the building. Power reaches each machine through a tangle of belts and pulleys hanging from the ceiling, all spinning off that single shaft. The arrangement quietly dictates everything. Power fades as it travels down the line, so machines crowd close to the shaft. Buildings grow tall and narrow to stay near the spinning core. And because you cannot switch on just one part of the shaft, the whole thing turns all day even if only one machine is working. The building serves the engine, not the work.

Then electricity arrives, and the owner can finally throw out the smoky steam engine. Here is what almost everyone did first. They pulled out the steam engine and dropped one big electric motor into the same spot, turning the same shaft, spinning the same belts, running the same machines in the same places. It worked. The gains were forgettable. The factory was still shaped like a steam factory. They had swapped the engine and kept the cage. I call this first beat the bolt-on: you attach the new power to the old setup, change nothing else, and it underwhelms.

The breakthrough sounds too small to matter. Make the motors small and cheap enough to put one on every machine. Once each machine had its own motor, the ceiling full of shafts and belts came down. Power now arrived through a thin wire, so a machine could be placed anywhere, and for the first time you could lay the floor out in the order the work actually flowed. That is the ancestor of the assembly line. Factories spread into wide, single-story buildings full of natural light. With the greasy overhead belts gone, there were far fewer fires. This is the second beat, the rebuild: the real gains never came from the new power alone, they came from redesigning everything around it.

Now the part that surprises almost everyone. All of this took about forty years. Electric motors were available in the 1880s. The real jump in factory output did not arrive until the 1920s. For decades electricity was right there, ready to use, and the payoff was missing. The slow part was never building the motor. It was finding the courage and the imagination to tear up the floor plan. This is the third beat, the lag: the long, quiet gap between when a technology arrives and when it pays off, during which it looks like failure and is actually construction.

The same rhythm played twice more. The telegraph let a message cross a continent in minutes, but at first people only used it to send the same letters faster. The rebuild came when the railroads used it to track trains in real time, which let a single line carry far more traffic safely, and it went so deep that in 1883 the railroads put the country on standard time, the zones we still use today. Then computers. At first companies used them to copy the paper forms they already had, a glowing version of the same invoice, and the gains stayed small. An economist named Robert Solow caught it in a single line: you could see the computer age everywhere except in the productivity numbers. The payoff came later, and it came to the companies that rebuilt around the machine, the Walmarts and the Amazons, not the ones that simply bolted it on.

Three transitions, one rhythm: bolt-on, rebuild, lag. And in every case, the payoff waited on a missing layer that had to be built. The grid. The protocols. The data systems. None of them was free, and each one took years.

AI is the fourth time, with one difference that changes everything

AI is the fourth transition, and the rhythm is already playing. The name for this one is intelligentization, the move from machines that follow instructions to machines that decide and

act. Companies are bolting it onto their old systems and feeling let down. The rebuild has barely begun. We are sitting in the lag, which is why most of the disappointment around us is not a verdict on the technology. It is a snapshot of an unfinished transition. If you have run an AI pilot that dazzled and then stalled, you have heard the rhythm yourself.

But the fourth one breaks the pattern in a single way, and this is the whole point.

The first three set free a resource. Energy, distance, memory. In each case a human still made the decisions and simply used the resource, so the missing layer was a delivery system: reach everywhere, stay reliable, and hand the resource to a person who decides. AI begins to set free something different. It is starting to move the act of deciding itself out of human hands. For the first time the machine, not the person, is making and taking some of the choices. There is no human at the other end just using a resource.

That single fact changes what the missing layer has to be.

	The first three (deliver a resource)	The fourth (raise an actor)
What is set free	Energy, distance, memory	Deciding and acting
Who makes the decision	A person, using the resource	The machine, within set limits
What the missing layer is	A delivery system	A habitat
What the layer must do	Reach everywhere and stay reliable	Reach and stay reliable, and also let the machine learn from what it causes and earn the trust to do more

A delivery system only has to do two things well: reach everywhere and almost never fail. A habitat has to do those two things and then more, because it is not handing a resource to a person. It is raising a decision maker. You do not deliver to an actor. You raise one. And raising is a much harder job.

The missing layer is a habitat

A model is not a deployment. A brain is not an organism. A system that acts in the real world has to sense, decide, act, stay safe, learn from what it causes, and stay trusted over time. Almost none of that lives in the model. It lives in the environment we build around the model.

You can feel the difference in the physical world, because the physical world gives a small warning before it commits. A bearing vibrates differently before it fails. A vehicle drifts onto a collision path before the impact. A factory process wanders before the batch is ruined. A patient changes before they collapse. The value of an acting AI lives inside that short window, in the

moment before a consequence hardens. So does the danger. Act well inside the window and you prevent the failure. Act wrongly and you may cause one that cannot be undone. Act without trust and the operators will quietly refuse to use you when it counts.

This is why a warehouse AI that gives excellent advice is still not something a manager will trust with a real decision at three in the morning. A model can be accurate without being authorized, useful without being allowed to act. What closes that gap is not a smarter model. It is an operating layer around the model.

I call that layer the habitat: the environment that lets an intelligent system act, learn from consequence, and earn trust in the real world. It does the humble, unglamorous things that turn extraordinary the moment a machine starts acting under real consequence. It defines exactly what the job is and what must never happen. It lets the system earn authority step by step, starting narrow, proving itself, and losing authority if it fails, because institutions remember harm. It watches while the system works. And it turns every meaningful action into a record that can be inspected later. Below that line, an AI produces clever outputs. Above it, outputs become evidence, a failure becomes a reusable lesson, and one deployment begins to teach the next. That is when AI stops being a service a company rents and starts becoming an asset it builds.

Now the table from the start of this piece can take its fourth row, the one we could not fill in until the habitat had a name.

The transition	What it set free	The missing layer it needed
Electrification	Energy	The power grid
Communication	Distance	Shared protocols
Digitization	Memory	The data layer
Intelligentization	The act of deciding	The habitat

That layer is the piece too few organizations set out to build deliberately, and it is the real reason careful, well-funded AI projects stall right after the technology works. Not a model problem. A habitat problem. It is also, I am now sure, why so many of those startup and company collaborations kept spinning. Each side was building half a bridge, and the bridge itself was no one's job.

You can watch this happen right now. Robotaxis that drive into flooded streets because the world produced a situation no one had planned for. Humanoid robots that dazzle in a lab and stumble in a warehouse. Predictive maintenance that works in one depot and fails in the next. Factory AI that gives good recommendations for years and is never once allowed to act on them. The headline usually says the model failed. The deeper pattern is that the habitat was never built.

Why this matters

Three times before, the payoff was never only about owning the new power. It was also about building the layer that made the power usable, and the people who built that layer helped shape the era that followed. The models will keep getting better, and that matters. The habitat is the other half, and it is mostly still unbuilt. If the rhythm holds, and it has held three times in a row, then the payoff is less a question of whether than of when we finish building it, and who chooses to build.

The people who can hear this rhythm tend to start building the layer early, while the debate is still about the models. More often than not, they are the ones who end up building the future rather than waiting for it.

If this is the doorway, the room is next. The fuller argument, and the decade of lived experience behind it, are in the essay *The Missing Layer of Physical AI*. The long version, written first for the partners who lived it with me, is the book, *The Case for Habitat*.

The missing layer is still mostly unbuilt. Building it, alongside the models, is much of what the next decade will be about.