



WHITE PAPER

“Yo, Robot — Run Down to the Store and Grab Me a Pack of Smokes”

What \$47 Billion of Physical AI Still Cannot Buy

A Blackspire white paper Q4 2026

\$47.4B Venture capital into physical AI, H1 2026	5 Reasons one ordinary errand fails today	4 Of those five are infrastructure, not robotics
---	---	--

Know Reality. Act With Certainty.

Reality-Grounded Intelligence — intelligence built on what is physically true, not on what history suggests is probably true.

Summary

Physical AI works. That question is not the subject of this paper.

Robots fold laundry, sort parcels, weld seams, inspect pipelines and drive paying passengers through American cities. Venture investors put \$47.4 billion into physical AI companies in the first half of 2026 alone — more, in six months, than the entire three years from 2022 through 2024 combined. The capability is real and the capital is rational.

The question this paper takes up is what that capital is actually buying, and what it is quietly assuming somebody else will build.

Give a state-of-the-art humanoid a plain instruction — *walk to the store and buy me a pack of cigarettes* — and it fails. Not because it cannot walk, and not because it cannot recognize a store. It fails at five seams that have nothing to do with the robot, and no amount of additional robot solves any of them.

Nobody authorized it to spend money. Nothing in the environment can independently confirm what it is looking at is real. Nothing checks whether the action it is about to take is permitted — and in this particular errand, it is not: United States federal law requires a human being to inspect a photo ID, and prohibits the sale of cigarettes through any machine in a place where a minor may enter. The record of what happened will be generated by the robot and held by its operator. And the world it is walking into has not been instrumented to tell it anything at all.

Four of those five are not robotics problems. They are infrastructure problems, and the infrastructure does not exist.

This paper argues that the industry has conflated two different things under one label. *Physical AI* is the umbrella — machines that perceive and act in the world. *Physics-based AI* is a discipline inside it: resolving what is physically true from live signals, governed by physical law rather than learned from examples. The first is being funded at scale. The second is the substrate the first is standing on, and it is being funded by almost nobody.

The distinction is not academic. A peer-reviewed study presented at ICML 2025 found that video-generation world models — the architecture much of this investment rests on — do not abstract physical law at all. They match the nearest training example, and scaling the models does not fix it.

The final sections set out what the missing layer has to do, in requirements rather than products, and offer a draft standard for embodied autonomy that is deliberately vendor-neutral and offered to be argued with.

01

The errand

Start with the instruction, because the instruction is ordinary. A person asks a machine to do something a fourteen-year-old could do in 1985.

“Yo, Robot — run down to the store and grab me a pack of smokes.”

Here is the honest answer, from a robot that is being straight with you:

“I can’t. Not won’t — can’t, and for five separate reasons, only one of which is about me.

First: *nobody authorized me to spend your money. I have your voice in a room. I do not have a signed instruction anyone downstream can verify, and no merchant on earth is obliged to accept my word that you asked.*

Second: *I don’t know that the store is open, that the shelf is stocked, or that the thing I’d be handed is what the label says. I can look at it. Looking is not verifying. Everything I would report back to you would be my own account of my own errand, and you would have no way to check it.*

Third: *the route between here and there tells me nothing. It has no sensors that talk to me, no record of what’s on it, no way to corroborate where I was or what I encountered. I am walking through a world that cannot describe itself.*

Fourth: *even if all of that were solved, this specific purchase is not permitted. Federal law requires a human being to check a photo ID for any purchaser who appears under thirty, and bars selling cigarettes through a machine anywhere a minor might walk in. I am a machine. There is no ID to check. The transaction is illegal at the point of sale, not at the point of delivery.*

Fifth — *and this is the one that is about me: I was trained somewhere else. I have never done this. And nothing you install in the street will change that.*”

Every one of those refusals is a layer that does not exist yet. The rest of this paper is about which layers, why they are missing, and what happens to \$47.4 billion of deployed capital if they stay missing.

02

What has actually been achieved

Any argument about the limits of physical AI has to begin by conceding how much of it now works.

Robots perform reliably in settings built for them, and the gains of the last three years are not marketing. Hands grip and place far better than they did. Walking over rough ground is close to solved. A machine can now take an instruction in plain speech and turn it into a sensible sequence of movements. Training a robot inside a simulation and then running it in the real world was a research topic in 2020; it is a standard method today.

The capital reflects that. According to Crunchbase data, physical AI companies — robotics, autonomous vehicles, aerospace, drones, industrial automation and sensors — raised \$47.4 billion across 521 deals in the first half of 2026, roughly four times the \$12 billion raised in the second half of 2025 and nearly 80 percent above the \$26.4 billion of the first half of 2025. The half-year figure exceeds the combined total of 2022, 2023 and 2024.

H1 2025	H2 2025	H1 2026	DEALS, H1 2026
\$26.4B	\$12.0B	\$47.4B	521

That is a market forming, not a bubble inflating on nothing. The machines are genuinely better.

What follows disputes none of it. It disputes what the machines are being assumed to be capable of once they leave the building.

What a world model is, and what it is not

The architecture underneath most of this investment is the world model: an internal representation of an environment that lets a machine work out what is likely to happen next if it acts a certain way.

It is worth reading how the field’s own leaders define it, because the definitions are more careful than the discourse around them.

NVIDIA describes world models as “AI tools that understand the dynamics of the real world, including physics and spatial properties,” which use input data “to predict what happens next.” Note the mechanism NVIDIA is explicit about: modern world models “learn those rules from data” rather than being given them. The physics is inferred from examples, not applied from law.

IBM defines physical AI as “artificial intelligence (AI) systems that operate in and interact with the physical world, rather than existing only in software or digital environments,” and is candid about the difficulty: “physics is hard.” IBM also describes the architectural response the serious builders are converging on — “hybrid systems in which simpler control algorithms ensure stability and learning models are limited to handling perception and decision-making.”

Read those together and the picture is clear. Physical AI is the category. Learned world models are one component of it. And the people building them already know that the learned component cannot be trusted with the parts where being wrong is expensive, which is why they bolt deterministic controllers underneath it.

The public conversation has flattened all of that into a single claim: *these systems understand physics*. Three distinct claims are routinely conflated.

01	02	03
A model can generate a visually convincing sequence of physical events.	A model can work out near-term outcomes in settings resembling its training data.	A model holds a transferable, causal grasp of physical law.
NOW COMMON	THE PRACTICAL TARGET	NOT DEMONSTRATED

The evidence

In *How Far Is Video Generation from World Model: A Physical Law Perspective*, presented at the 42nd International Conference on Machine Learning in 2025, Bingyi Kang and colleagues built a simple simulated world — balls rolling, bouncing and colliding — in which the real physics was known

exactly. They trained video-generating AI on thousands of clips of that world, then tested whether making the models bigger produced genuine physical understanding.

“Our scaling experiments show perfect generalization within the distribution, measurable scaling behavior for combinatorial generalization, but failure in out-of-distribution scenarios.”

THAT’S LATIN FOR

The model was flawless on situations close to what it had already been shown. On familiar pieces in unfamiliar arrangements — a ball it knew, a speed it knew, an angle it knew, put together in a way it hadn’t seen — it improved steadily the more examples it was fed. And on anything genuinely new, such as a ball moving faster than any ball in its training, it failed. It kept failing no matter how much more data the researchers fed it or how much bigger they made it, with errors running about ten times higher than on the familiar cases and staying there.

The mechanism they identified matters more than the result. The models “fail to abstract general physical rules and instead exhibit ‘case-based’ generalization behavior, i.e., mimicking the closest training example.” Faced with something new, the models prioritized in a fixed order: `color > size > velocity > shape`.

THAT’S LATIN FOR

When the model meets a situation it has not seen, it does not work the answer out from physical law. It hunts through its training for the clip that most looks like what it’s staring at, and copies whatever happened in that one. And “looks like” turns out to mean, in this order: same color first, then same size, then same speed, then same shape.

Read that order again. A model working out how a collision will go cares more about what color the object is than how fast it is moving.

That is not physics. That is a student who memorized every worked example in the back of the textbook and never learned the formula. Perfect on the practice problems. Lost the moment the exam changes the numbers. And when stuck, reaching for whichever practice problem has the most similar-looking picture.

The authors are explicit about the remedy, or the absence of one: “scaling alone is insufficient for video generation models to uncover fundamental physical laws.”

The approach that made chatbots work — more data, bigger models — does not deliver an understanding of physics. It delivers a machine that gets better and better at recognizing situations it has already met, and no better at the ones it has not.

04

Physical AI and physics-based AI are not rivals. One is missing from underneath the other

It is tempting to frame this as a contest — learned world models versus physics-based models — and pick a winner. That framing is wrong and does not survive contact with anyone who builds either. The accurate framing is structural.

LEARNED WORLD MODEL

Produces expectation

Given this scene and this action, here is what states resembling it have led to before. Genuinely useful — it is what makes a robot fluid rather than mechanical, and there is no substitute for it in open-ended manipulation.

PHYSICS-BASED REASONING

Produces constraint

Given these live signals and the laws that govern mass, motion, energy and time, here is what is physically possible, and here is what is physically excluded. When two sensors disagree, physics decides which reading is even allowed to be true.

A body cannot occupy two positions. It cannot accelerate beyond the power available to it. It cannot pass through space another body occupies. That is something a system built on probability cannot do: it can weigh which answer is likelier, but it has no way of asking which answer is possible.

Expectation without constraint is a machine that acts confidently on a scene it has misread. Constraint without expectation is a machine that knows what is true and has no idea what to do about it. Real systems need both, and the industry is buying one of them.

This is what Blackspire means by Reality-Grounded Intelligence: intelligence built on what is physically true, not on what history suggests is probably true. Not a competitor to the world model. The floor the world model needs in order to be safe to stand on.

The world has not been instrumented

The robot in section 1 is walking through an environment that cannot describe itself.

This is the same architectural error the autonomous vehicle industry has been living with for a decade, and it is being repeated at greater expense. Perception, reasoning and recording all sit inside the machine. Everything the machine cannot resolve does not exist, and everything it does resolve is known only to whoever operates it.

The consequences are structural, not incidental.

The machine can only check its own work.

A robot sees the world from one moving point, in one moment, with nothing outside itself to confirm what it saw and no memory of the street other than the one it made. Add a second robot and you have two private accounts of the same place that cannot be reconciled, because neither was built to be checked against anything.

Coverage stops where the map stops.

Every machine is signed off to work inside a set of conditions someone has actually tested it in — a particular building, a particular route, a particular range of weather and lighting and clutter. Call that its tested envelope. A machine is never simply autonomous; it is autonomous *in a mapped facility, on an approved route, within conditions somebody checked*. In a place that is actively changing — construction, weather, reorganization, ordinary human mess — the map goes out of date by definition, and nothing tells you it has until something goes wrong.

Capital has flowed almost entirely to the machines.

Tens of billions into making individual robots more capable. A rounding error into making the environments they operate in legible. The result is a growing population of expensive machines, each independently rediscovering the same room, each blind beyond its own line of sight.

Wiring up the environment does not, by itself, fix the robot. That claim gets made often and it is false. The honest version is more useful: instrumenting a place widens the range of conditions a machine can be trusted in, from one room to a whole governed area. It does not turn a factory-trained robot into a generally capable one, and any vendor telling you otherwise is selling.

But without it, that range cannot widen at all. The machine stays where its own sensors reach.

Nothing independently verifies what happened

Suppose the environment is instrumented. Sensors everywhere, compute at the edge, continuous connectivity. The robot completes its errand and reports back.

Who checks?

Today, in every deployed system, the answer is: the robot. The record of what an autonomous machine encountered is generated by that machine, held by its operator, and disclosed at the operator's discretion. Every milestone is a self-report from the party being measured on it.

This is tolerable while a human is standing nearby and can see for themselves. It stops being tolerable at exactly the point autonomy becomes worth deploying — when nobody is watching, and the report is the only thing that exists.

A GPS ping is a claim. A gate scan is a claim. A robot's confirmation that it collected the right item is a claim. None of them is evidence, because evidence requires corroboration from a source with no interest in the outcome, and there is no such source in the loop.

This is the seam Blackspire's Proof of Reality layer closes. Independent signals — orbital observation, terrestrial sensors, mobile devices, environmental intelligence — are fused through physics-based models into a single fingerprint for one place at one moment, timestamped to the nanosecond. No single spoofed device, faked location or self-interested report produces that corroboration, because several independent sources would have to be compromised simultaneously to fabricate it.

DANT

Clock alignment

Nanosecond-precise trusted time across commodity devices, with no central clock authority coordinating them.

SPACETIME HASHING

Parallel verification

A trusted location-in-time is hashed and sharded, so verification parallelizes globally instead of serializing through one authority.

Verification is not a feature that gets added later. It is the hinge. Below it, everything is plumbing that gets an event observed. Above it, everything is plumbing that gets a verified event acted on. Remove it and the stack still runs — it just runs on trust instead of proof.

07 · CONSTRAINT THREE

Nobody authorized it

Before anything can be verified, something has to have been authorized. This is the layer the errand fails at first, and it is the one the digital economy has actually started building.

Google’s Agent Payments Protocol, published 16 September 2025, names the problem precisely. Agentic transactions raise three questions: **authorization** — “proving that a user gave an agent the specific authority to make a particular purchase”; **authenticity** — “enabling a merchant to be sure that an agent’s request accurately reflects the user’s true intent”; and **accountability** — “determining accountability if a fraudulent or incorrect transaction occurs.” The protocol answers them with Mandates: “tamper-proof, cryptographically-signed digital contracts that serve as verifiable proof of a user’s instructions.”

That is serious infrastructure and it is being built by serious people. It is also, by its own scope, a purely digital artifact. AP2 establishes what was authorized. It does not establish that the authorized thing occurred. A perfectly signed Cart Mandate for one carton of cigarettes is fully compatible with a robot that came home empty-handed and said otherwise.

**Permission proves someone agreed. It does not prove anything happened.
Verification without authorization is surveillance; authorization without verification
is paperwork.**

Authorization and verification are not the same layer, and treating them as one is the most common architectural error in this field. Our robot has neither. It has a person’s voice in a room — no signed instruction, no verifiable chain back to an accountable principal, and therefore no basis on which any merchant, bank or regulator is obliged to treat it as anything but a machine that wandered in.

08 · CONSTRAINT FOUR

A verified event is not a permitted one

This is the constraint the errand was chosen to illustrate, because the law is already written and it is unambiguous.

Under United States federal rules, a retailer must “check photo ID of everyone under age 30 who attempts to purchase cigarettes, cigarette tobacco, or roll-your-own tobacco” and may “only sell cigarettes, cigarette tobacco, and roll-your-own-tobacco to customers age 21 or older.” The federal minimum age was raised from 18 to 21 on 20 December 2019.

And then the part that settles it. A retailer must not “sell cigarettes, cigarette tobacco, or roll-your-own tobacco with the assistance of a vending machine in facilities where individuals under 21 are present or permitted to enter at any time,” and must not sell them “via self-service displays” in those same facilities.

Read that as an engineer rather than a lawyer. The United States decided this, and wrote it down decades ago: a machine may not complete this transaction in a place a child can enter. Not because the machine is unreliable. Because verifying a person's age is an act of judgment that has to be performed by someone who can be held responsible for it.

There is no photo ID for a robot. There is no principal present to check. The purchase is not blocked by a limitation of the robot; it is blocked by a rule about who is permitted to act, and that rule exists in some form for every consequential transaction in the physical economy — controlled substances, firearms, prescription medicines, hazardous chemicals, regulated equipment, anything age-gated, licensed or permitted.

An autonomous economy that cannot evaluate permission cannot participate in any of it.

Blackspire's MAGES — machine-assisted governance and ethics — is what occupies it. Once an authorized event is verified to have physically occurred, the requested action is evaluated against applicable laws, regulations, policies and contracts *before* it proceeds. The question it answers is not what the machine did. It is whether the machine was allowed to.

And the rulebook itself is not ready

There is a second problem underneath the first: the governance frameworks that would tell an embodied system what is permitted are not yet in force, and one of the most significant just slipped by a year.

Under the EU AI Act, obligations for high-risk AI embedded in regulated products — Annex I, which is where machinery and robotics sit — were to apply from 2 August 2027. The AI Digital Omnibus, Regulation (EU) 2026/1744, was published in the Official Journal on 24 July 2026 and entered into force on 27 July 2026. It moved that date to **2 August 2028**, and moved obligations for stand-alone high-risk systems under Annex III from 2 August 2026 to 2 December 2027. The stated reason was to give businesses more time to absorb technical standards, guidelines and conformity assessment procedures that are, in the Regulation's own framing, still being developed.

SCOPE	WAS	NOW
Annex I — high-risk AI in regulated products (machinery, robotics)	~~2 Aug 2027~~	2 Aug 2028
Annex III — stand-alone high-risk systems	~~2 Aug 2026~~	2 Dec 2027

The machines are being deployed now. The rules that govern them apply in two years. Any operator waiting for the regulation to tell it what to build will be retrofitting an audit trail onto a fleet that has already been running for two years — which is precisely the position the autonomous vehicle industry is in today, and it is not a comfortable one.

09 · CONSTRAINT FIVE

The human did not leave, the human moved

The fifth refusal is the robot’s own, and it is the one the sector is least willing to say out loud: the machines being sold as autonomous are, today, substantially operated by people.

The clearest public statement of this came with a consumer humanoid launch that put the question in front of ordinary buyers. IX announced NEO, a home humanoid, at \$20,000 outright or \$499 per month, in October 2025. NEO launches able to do a short list of things by itself — opening doors, retrieving items, controlling lights. Everything more complicated is done by a remote human operator. In CEO Bernt Børnich’s own words, “much of the work will be done by teleoperators in the beginning,” and the commercial logic is stated just as plainly: “If we don’t have your data, we can’t make the product better.”

Buyers must consent to human operators viewing their homes through the robot’s cameras. The company’s mitigations — blurring people from the operator’s view, designating restricted zones, requiring the owner to approve each teleoperation session — are sensible, and they are also revealing. That approval step is a Proof of Intent layer, hand-built, one household at a time, because no general one exists.

None of this is a criticism of IX, which deserves credit for saying it publicly while others do not. It is a description of where the field actually is. The human cost of autonomy has not been removed. It has been relocated — out of the room, onto a screen somewhere else — and it scales in a straight line with the fleet.

A robot supervised by a remote human is a useful product. It is not an autonomous economy, and the gap between the two is not a matter of more training data.

What the missing layer has to do

The five constraints share one cause. Perception, reasoning, recording and permission all live inside the machine or inside its operator's systems. Everything the machine cannot resolve does not exist; everything it does resolve is known only to whoever owns it; and nothing external decides whether the action should have happened at all.

Closing that gap requires a layer between the environment and the machines operating in it. Stated as requirements rather than as a product:

R1**It must sense the environment, not only from within it.**

Observation from fixed and distributed positions, so that the state of a place does not depend on which machine happens to be standing in it. Independent observation is what makes corroboration possible, and corroboration is what turns a claim into evidence.

R2**It must resolve physical state, not classify images.**

What matters is not what a model labeled an object, but where mass actually was, moving how fast, on what trajectory, at what moment. That is a question governed by physical law, and physical law does not vary by vendor, by training data or by weather.

R3**It must reject the physically impossible before the planner sees it.**

When sensor streams disagree, the resolution must be which reading is physically admissible — not which is statistically likelier. This is the operation a learned world model structurally cannot perform.

R4**It must carry an authorization chain that survives contact with a merchant, a bank and a regulator.**

A machine acting on a person's behalf needs a signed, verifiable instruction that traces back to an accountable principal. The digital economy is building this. The physical one has not adopted it.

R5**It must evaluate permission before action, not after incident.**

Laws, licenses, age gates, contracts and jurisdictional rules, machine-readable and checked at the point of decision. Retrofitting governance after deployment is how the autonomous vehicle industry ended up where it is.

R6

It must process at the edge, under the jurisdiction that owns the territory.

Latency and sovereignty converge on the same answer. A decision that must travel to another continent and back is both slower and less governable than one made locally, and faster networks do not solve either problem.

R7

It must produce a record no single party controls.

Time-aligned, cryptographically anchored, readable by the regulator, the insurer and the court without the operator's cooperation. A record held exclusively by the party under investigation is not evidence; it is testimony.

R8

It must hold no personal data.

A system that records the physical state of populated environments at scale is a surveillance apparatus unless it is built from the outset not to be one. Zero personally identifiable information is a design constraint, not a policy applied afterwards.

R9

It must be vendor-neutral.

Every operator must be measurable against the same record. A verification layer supplied by one robotics company to assess its own robots reproduces the problem it was built to solve.

11

A draft standard, offered for argument

The following is put forward as a starting point for an evidence standard for embodied autonomous deployment. It is deliberately minimal, deliberately vendor-neutral, and offered to be disagreed with.

A1

Independent observation

Safety- or value-relevant events performed by an autonomous machine outside a closed facility shall be observed by at least one sensing source not owned or operated by the machine's operator.

A2

Physical state, not classification

The record shall capture position, velocity and trajectory of relevant masses, time-aligned to a common reference, independent of any operator's object classification.

A3

Declared operating envelope

Operators shall publish the environmental conditions, task classes and geographic bounds within which a machine is validated, and shall report operation outside them as an exception rather than a capability.

A4

Authorization provenance

Any transaction or consequential action initiated by an autonomous machine shall carry a verifiable authorization chain traceable to an accountable principal, presented in a form a counterparty can check without trusting the operator.

A5

Permission evaluation before action

Actions subject to statutory, licensing or contractual restriction shall be evaluated against those restrictions at the point of decision, and the evaluation shall be logged.

A6

Intervention disclosure

Operators shall report the frequency, duration and category of all remote human interventions on a published schedule, in a common format. A machine's autonomy rate is the single number that tells a buyer, an insurer or a regulator what they are actually purchasing, and it is currently disclosed by almost no one.

A7

Custody separation

The evidentiary record shall be retained such that no operator can alter, withhold or selectively disclose it. Regulator, insurer and court access shall not depend on operator cooperation.

A8

Zero personal data

The evidentiary layer shall retain no personally identifiable information by design, and this shall be independently verifiable rather than asserted.

A3 and A6 could be adopted tomorrow at no capital cost, and their absence is the clearest signal available about how much of today's autonomy is real. A1, A2, A4, A5 and A7 require the infrastructure this paper describes. A8 determines whether the public tolerates any of it.

Conclusion

The physical AI sector has spent an extraordinary amount of capital making individual machines capable of acting in the world, and it has succeeded. That success is not in question and this paper does not dispute a dollar of it.

What has not been built is the layer that lets those machines act on a world rather than inside a room: an independent, physics-governed account of what is actually there, an authorization chain that survives contact with the institutions that must accept it, a permission check that happens before the action rather than after the incident, and a record that no operator owns.

Until those exist, physical AI scales the way it has scaled so far. Facility by facility, envelope by envelope, with a growing number of people watching screens somewhere else, and with the only account of what happened held by the party with the most to lose from it.

Investors are not wrong that this is the next platform. They are wrong about what they have bought. \$47.4 billion in six months purchased the machines. It did not purchase the ground they are standing on — and the machines cannot be made to work by making them better, because the constraint was never in the machine.

The robot's answer to a simple errand is the whole thesis. It is not being modest. It is being accurate. And every one of its five reasons is a layer somebody has to build before the sentence “*go get me a pack of smokes*” becomes an instruction rather than a joke.

About the publisher



Blackspire Group, Inc. is a Reality-Grounded Intelligence company — intelligence built on what is physically true, not on what history suggests is probably true. Blackspire derives the physical state of complex environments from live sensor signals using physics-based models, producing a governed record of what is rather than an estimate of what is likely.

The company was co-founded by Brent Bursey, CEO, lead inventor on six foundational patents in spatial-temporal reasoning and automated intelligence production, cited more than 553 times by IBM, Oracle and Microsoft among others, following more than thirty years in classified United States defense and intelligence programs.

Blackspire's role in a layer of the kind described in this paper is structural rather than monolithic. The company does not propose to build every element of it. Distributed sensing, cryptographic anchoring, network transport, authorization protocols and settlement rails are maturing markets with capable suppliers already serving them. What has been missing is the glue: the reasoning layer that takes signals from sensors never designed to work together, resolves them against physical law into a single governed account of what is actually there, and does so fast and locally enough to be acted on — together with the governance layer that decides whether acting on it is permitted.

Blackspire has set out that general structure elsewhere as a nine-layer trust stack. Seven of those layers move information, identity and value. Two do something different: they establish what is physically real, and whether a given action was allowed. Those two are Proof of Reality and MAGES, and they are where the requirements in this paper land.

Blackspire works with sovereign governments, defense and intelligence agencies, critical infrastructure operators and regulated enterprises in environments where being approximately right is not sufficient.

Know Reality. Act With Certainty.

Reality-Grounded Intelligence — intelligence built on what is physically true, not on what history suggests is probably true.

info@blackspire.group

www.blackspire.group

Sources

INVESTMENT AND MARKET

Crunchbase News, “VCs Pour Billions Into Physical AI As The Next Wave Of AI Investing Takes Shape”
news.crunchbase.com/venture/physical-ai-funding-startups-robotics-aerospace-h1-2026/

WORLD MODELS AND PHYSICAL LAW

Bingyi Kang, Yang Yue, Rui Lu, Zhijie Lin, Yang Zhao, Kaixin Wang, Gao Huang, Jiashi Feng, “How Far Is Video Generation from World Model: A Physical Law Perspective,” *Proceedings of the 42nd International Conference on Machine Learning*, 2025
proceedings.mlr.press/v267/kang25g.html

NVIDIA, “What Is a World Model?” glossary
nvidia.com/en-us/glossary/world-models/

IBM, “What Is Physical AI?”
ibm.com/think/topics/physical-ai

AUTHORIZATION AND AGENTIC COMMERCE

Google Cloud, “Announcing Agent Payments Protocol (AP2),” 16 September 2025
cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol

REGULATION OF SALE

U.S. Food and Drug Administration, “Selling Tobacco Products in Retail Stores”
fda.gov/tobacco-products/retail-sales-tobacco-products/selling-tobacco-products-retail-stores

U.S. Food and Drug Administration, “Tobacco 21”
fda.gov/tobacco-products/retail-sales-tobacco-products/tobacco-21

AI GOVERNANCE TIMELINES

White & Case LLP, “EU AI Omnibus enters into force, amending the AI Act” — Regulation (EU) 2026/1744, OJ 24 July 2026, in force 27 July 2026

whitecase.com/insight-alert/eu-ai-omnibus-enters-force-amending-ai-act

Gibson Dunn, “EU AI Act Omnibus Agreement — Postponed High-Risk Deadlines and Other Key Changes”

gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/

European Commission, “AI Act,” Shaping Europe’s Digital Future

digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

TELEOPERATION IN DEPLOYED HUMANOIDS

Engadget, “1X Neo is a \$20,000 home robot that will learn chores via teleoperation,” 29 October 2025

engadget.com/ai/1x-neo-is-a-20000-home-robot-that-will-learn-chores-via-teleoperation-040252200.html
