Humanoid Foundation Models

The ultimate resource for humanoid robots and embodied AI

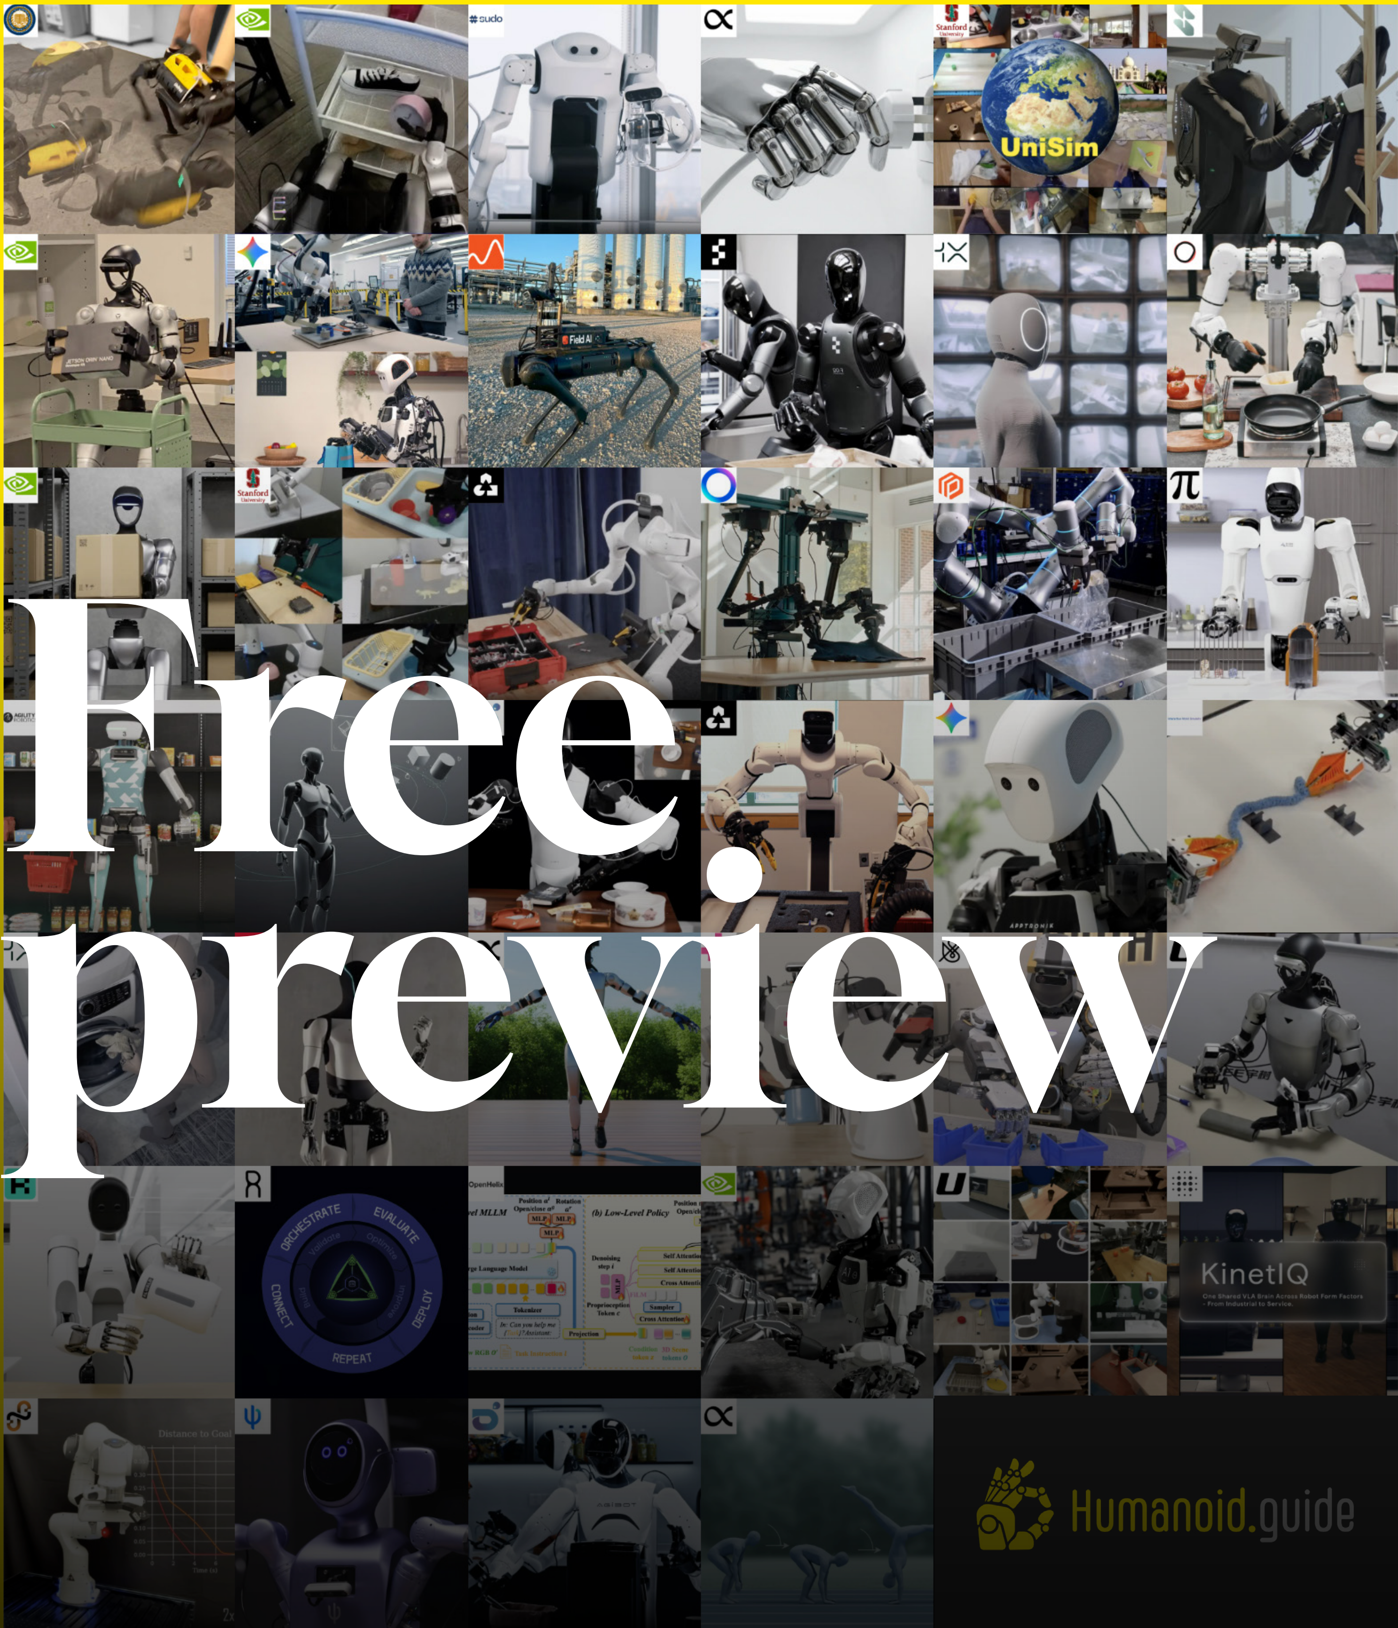

Free preview

Free preview:
what every chapter asks

The intelligence layer of humanoid robots is being rebuilt – from vision-language-action models to predictive world models. This preview walks the full report chapter by chapter. It doesn't give away the answers; it lays out the questions each chapter takes on, so you can see exactly what you'd be reading for.

01 – Reshaping Physical AI

Robotics has lived through two eras: machines that follow a script, and machines that react to what they see. A third is now arriving – machines that imagine the world before they move in it. What changed to make this possible now, and why are the parallels with the early days of self-driving so hard to ignore? This opening chapter sets the scene: the forces converging on physical AI, and why the brain – not the body – has become the contested ground.

02 – From VLAs to World Models

For three years, the fastest route to a general-purpose robot was to bolt actions onto a language model. It worked – until it hit the limits of the data it was trained on and the physics it never modelled. So what comes after the vision-language-action model, and why are the leading labs quietly shifting to systems that *predict* the world rather than just describe it? This chapter traces the turn, and the speed-versus-generalization trade-off now reshaping the whole field.

03 – Inside a World Model

Everyone is talking about world models. But what is actually *inside* one? This chapter opens the box: the four parts every embodied world model shares, the three distinct species you'll meet in the literature, and the surprising trick that lets a robot learn what a verb means from video nobody ever labelled. If you've wondered where the "imagination" really lives, this is the anatomy lesson.

Get the full Report

04 – The Four Architectural Paradigms

Under the single banner of "world models" sit four very different ways to turn a camera feed into an action – and they disagree on fundamentals. Should a robot think in discrete tokens, like a language model, or in smooth continuous motion? Should it try to predict the future at all? This chapter sorts the four paradigms and shows why the strongest systems are quietly converging on one two-speed design, borrowed from how humans think fast and slow.

05 – JEPA & the Predictive Architectures

Not every world model dreams in pixels. One camp, led by Yann LeCun, argues that predicting raw video is a waste – that intelligence lives in an abstract latent space that's far cheaper and faster to imagine in. Are they right? This chapter covers JEPA and the hybrid architectures betting against generative video, and what it would mean for the field if the pixel-predictors turn out to be solving the wrong problem.

06 – Generative World Simulators

One family of world models is used less to act than to *dream* – generating endless, navigable, physically plausible video to train and test robots that have never left the lab. If real-world robot data is scarce and expensive, could a learned simulator simply manufacture as much experience as you can afford? This chapter covers the simulators racing to become the cheapest robot-training factory in the world – and why that may be the most valuable position in the entire stack.

07 – Scaling Laws & the Data Question

The most consequential claim in physical AI is not that world models *work* – it's that they *scale*. If robot capability grows as predictably with data and compute as language did, then general-purpose robots become a budget line, not a research mystery. But does the messy physical world really obey the same straight line that built the LLM era? This chapter weighs the evidence on both sides – including the caveats that should keep anyone honest.

08 – Foundation Models I: NVIDIA's Platform

If world models are the brain of physical AI, NVIDIA is trying to own the entire nervous system: the chips, the simulators, the models and the tools everyone else builds on. How close is it to becoming the toll-collector of a whole industry? This chapter examines the stack – Cosmos, GR00T and DreamZero – and asks what it means for everyone else when one company supplies both the picks and the shovels.

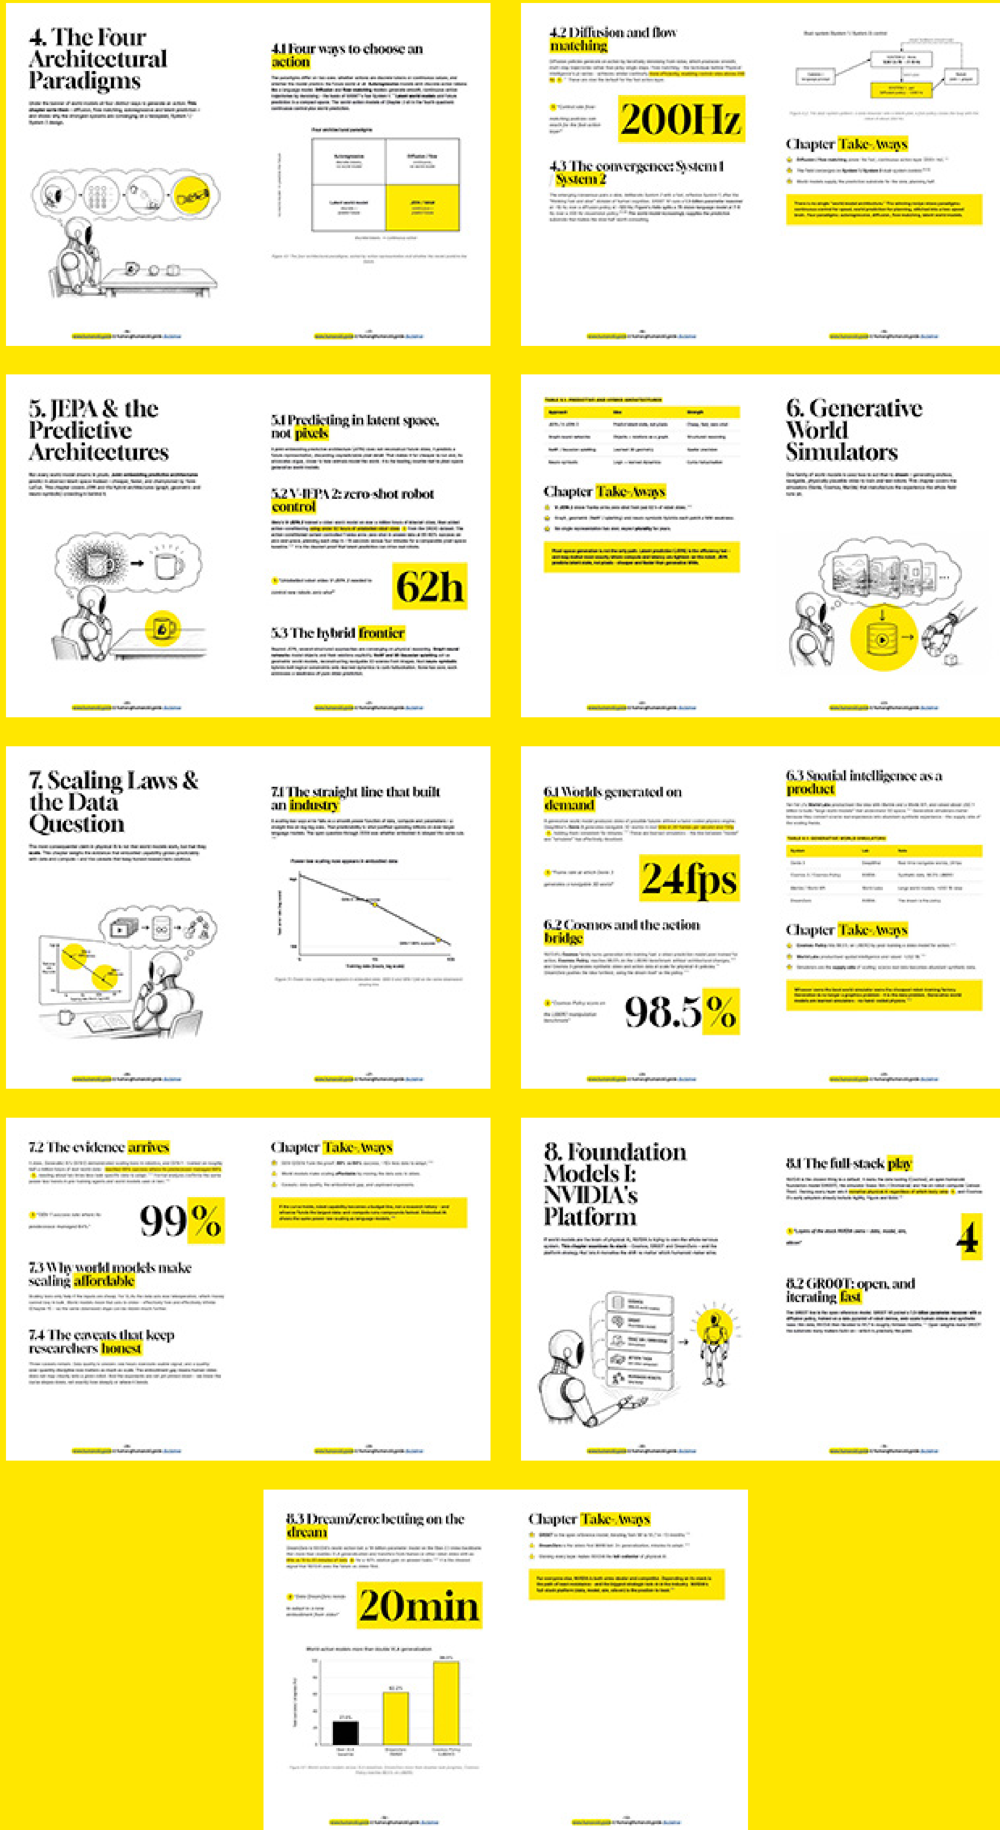

Get the full Report

09 – Foundation Models II: The Challengers

NVIDIA is not alone. A cluster of extraordinarily well-funded labs is racing to own the robot brain – some selling intelligence to all comers, some building the bodies too, several valued in the billions on almost no revenue. Who are the real contenders, what are they betting on, and which strategy – sell the brain, or own the whole robot – actually wins? This chapter profiles the challengers and the wagers behind their valuations.

10 – The Humanoid Makers

A brain needs a body. This chapter profiles the companies actually building the humanoids – and the gap between the demo reels and the shipping containers. Who is really making robots at scale, who is merely promising to, and why does the leaderboard look completely different depending on whether you count units shipped or capital raised? For anyone choosing a partner, this is the map – and the warning.

11 – Training Data & the Data Moat

World models run on experience – and experience is the one thing money can't simply buy. The internet has no equivalent of a robot's sense of touch, force and consequence. So where does the training data actually come from, who is quietly accumulating the deepest reserves, and is a durable "data moat" forming while no one is looking? This chapter follows the scarcest resource in the field.

12 – Simulation & Sim-to-Real

Most robot training now happens in worlds that do not exist. It's faster, safer and almost free – until the robot meets reality and the small differences start to matter. How wide is the "sim-to-real gap," why is it as much a safety problem as an engineering one, and can better world models finally close it? This chapter covers the simulated training ground where the scaling thesis is actually being executed.

13 – Compute & Infrastructure

World models are hungry in two contradictory ways: they must be trained on warehouse-scale clusters, yet run in real time on a battery strapped to a moving robot. How do you reconcile those two regimes – and is on-robot compute keeping pace with the models' appetite, or quietly becoming the ceiling? This chapter maps the silicon beneath physical AI, from the data center to the chip in the robot's chest.

Get the full Report

14 – USA: Innovation vs. Volume

America designs the best robot brains, owns the best chips, and writes the biggest cheques. It also builds almost no humanoids at scale and depends on Asian suppliers for the parts. Can a country lead the intelligence race and still lose the robot it's meant to run on? This chapter examines the American position – an innovation engine with a manufacturing problem – and the open question of whether the US will treat humanoids as strategic before it's too late.

15 – China: State-Backed Scale

China builds more humanoids than the rest of the world combined, on a supply chain it largely owns – and it's doing so by design, not by accident. What does a state-directed push into embodied AI actually look like, where are its real advantages, and which single chokepoint could it use to slow every competitor at once? This chapter examines the Chinese model and the leverage it has already shown it's willing to use.

16 – Europe, Japan & Korea

While the US and China fight over scale, a third bloc is competing on something quieter: trust, precision, and the rules everyone else has to follow. Can regulation be a competitive advantage rather than a handicap? Why might Japanese and Korean components, and EU-grade compliance, turn out to be scarce and exportable assets in a split-stack world? This chapter makes the case for the players too easily written off.

17 – Geopolitics

Humanoids sit squarely on top of three fights the great powers are already having – over semiconductors, over critical minerals, and over data. What happens to a robot's cost and availability when "de-risking" from one country can multiply its bill of materials? And why is a walking, watching, data-collecting machine destined to become a state-level concern? This chapter maps the politics that will shape the industry as much as any breakthrough.

18 – Industrial Applications

Humanoids will earn their keep in factories and warehouses long before they reach the living room. But where, exactly, do they pay off first – and what separates a viral pilot from a fleet that actually runs unattended? This chapter examines the early money: the automotive lines and logistics floors, the rise of robot-as-a-service pricing, and the pilot-to-production gap where most programs still quietly stall.

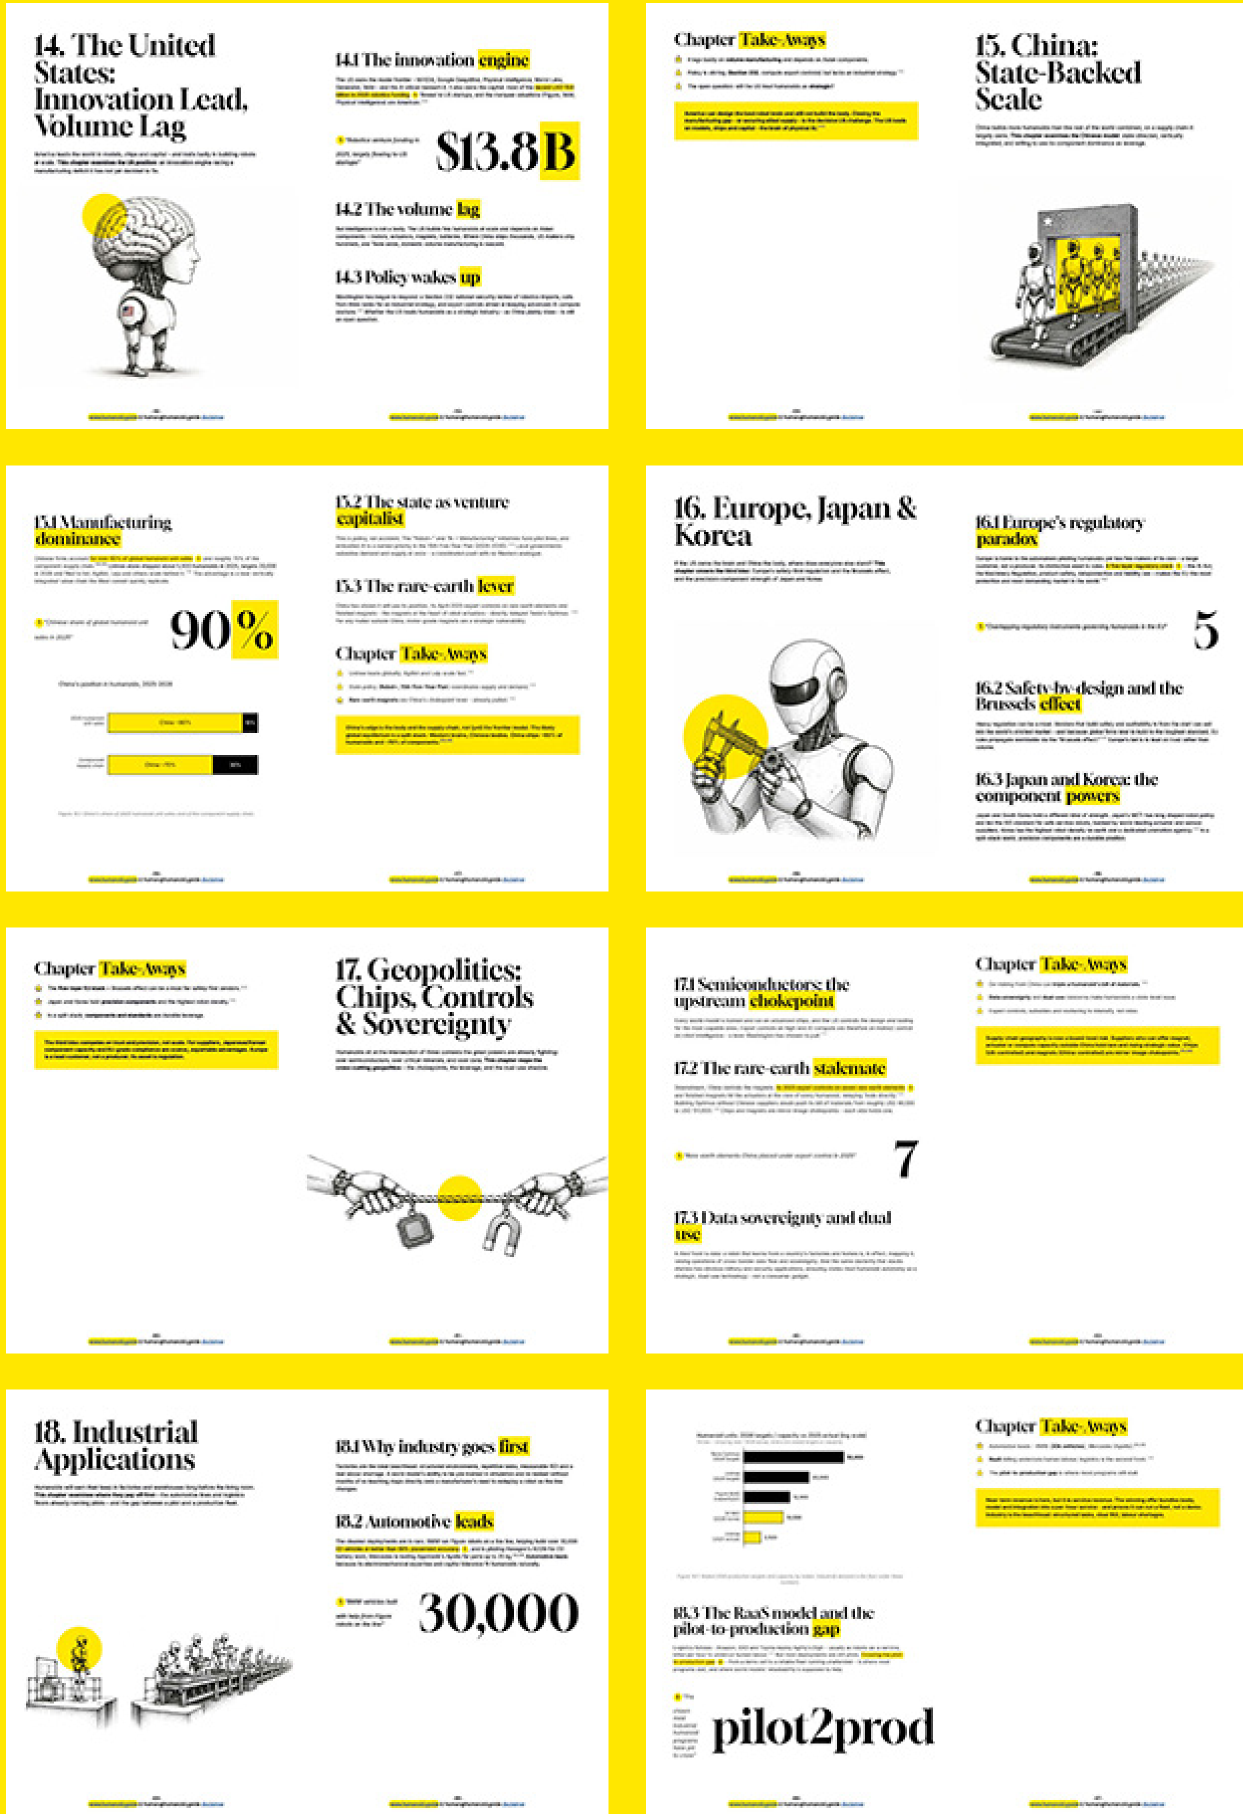

Get the full Report

19 – Domestic & Service Robots

The home is robotics' hardest market – unstructured, unforgiving, full of edge cases – and also its largest prize. Everyone wants the robot that loads the dishwasher; almost no one agrees on when it actually arrives. What has to be true for humanoids to cross the threshold into ordinary homes, and what does the adoption curve really look like? This chapter separates the dream from the timeline.

20 – Benchmarks & Evaluation

If world models are improving, how would we actually know? A model can generate gorgeous video and still drive a robot straight into a wall. What does it really mean to *measure* a world model – and are today's benchmarks testing the things that matter, or the things that are easy to score? This chapter surveys how the field grades itself, and where the scoreboard may be flattering everyone.

21 – Strengths, Shortcomings & Safety

A fair verdict needs both columns. The advances are real – but so are hallucination, long-horizon drift, the sim-to-real gap, and an unsettling new class of alignment risk that comes *precisely* from a machine's ability to imagine consequences. So when does this technology actually become reliable enough to trust unattended? This chapter sets the genuine progress against the unsolved problems and answers, soberly, the question everyone else is dancing around.

22 – Business & Investment

Capital has already decided world models are the next platform – the valuations say so, even where the revenue doesn't yet exist. So where does the value actually accrue as the dust settles: the model, the data, the body, or the scarce components no one can source elsewhere? And how should you read forecasts that disagree by three orders of magnitude? This chapter follows the money – and flags the risks the boldest valuations quietly assume away.

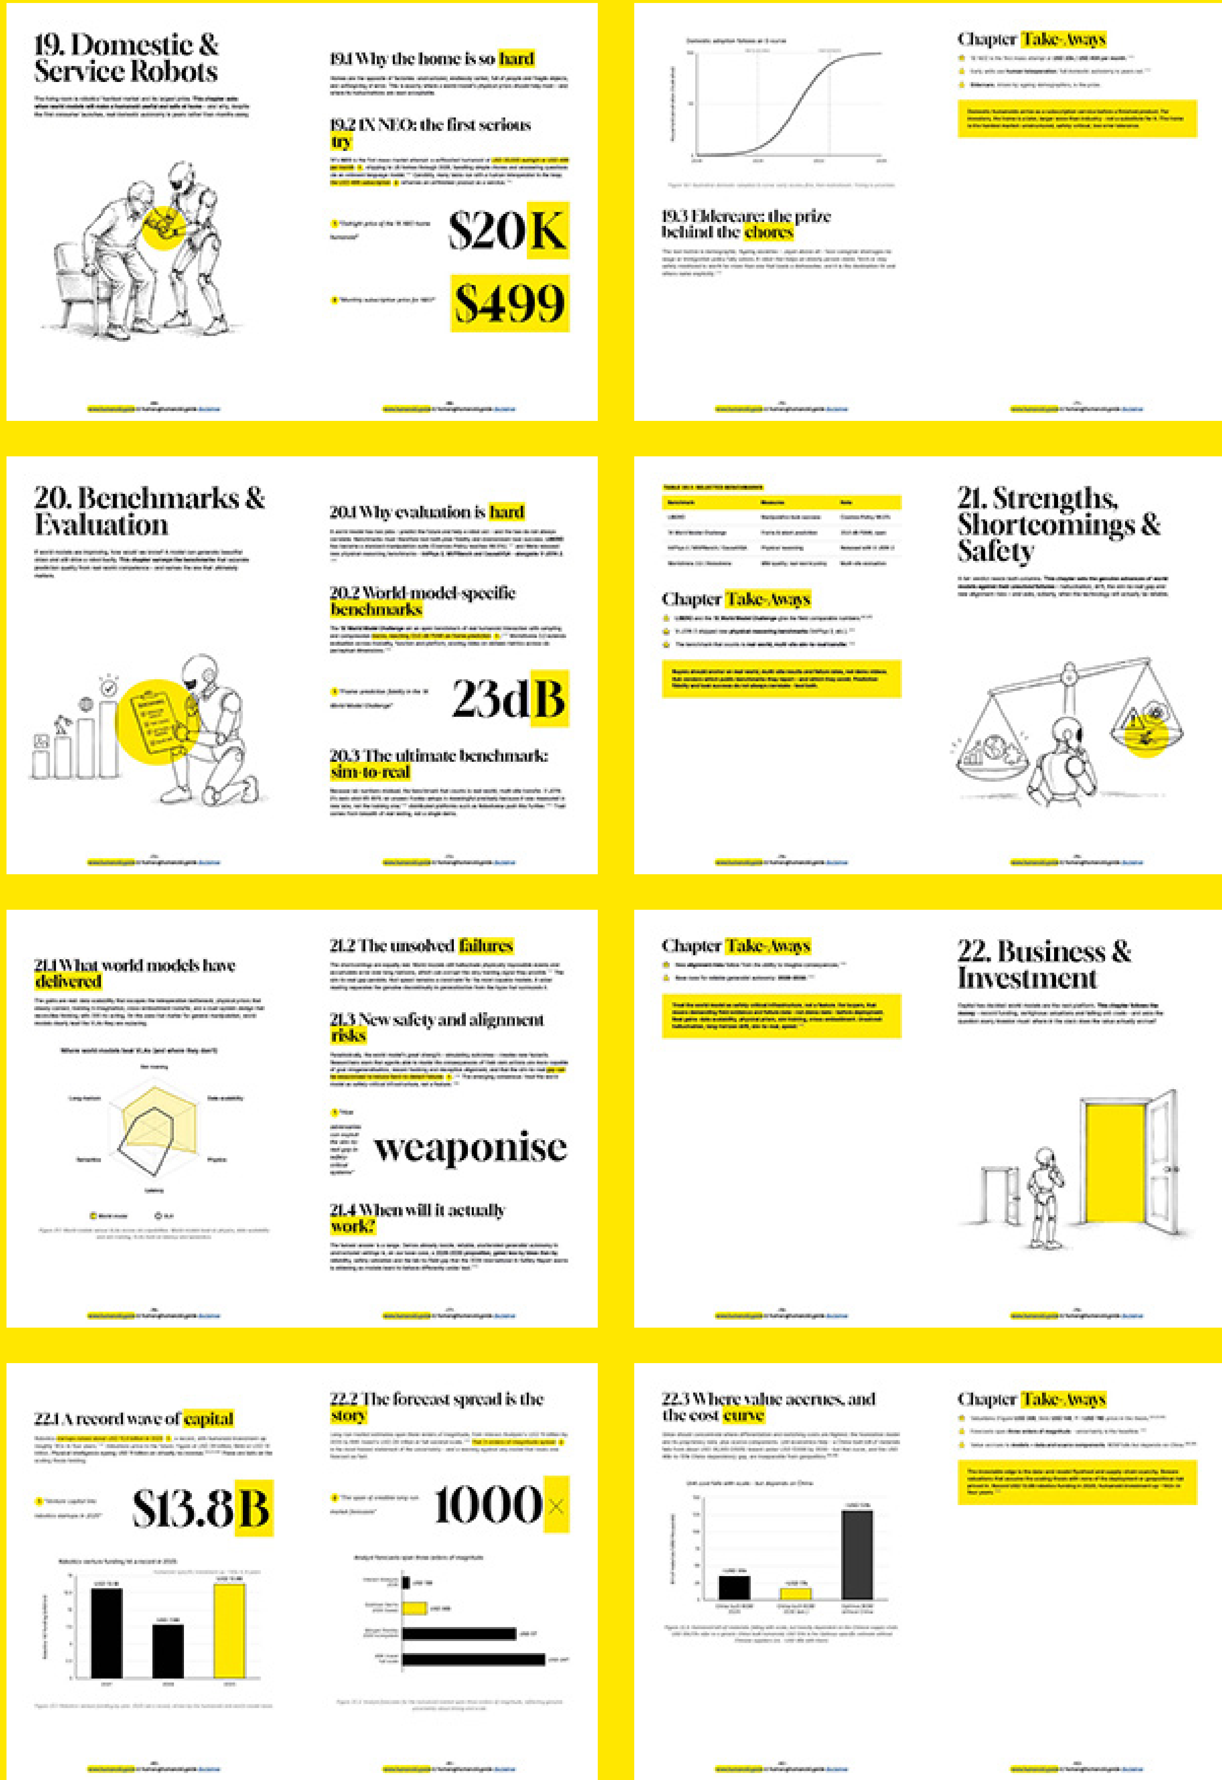

Get the full Report

23 – Scenarios, Outlook & the Playbook

Everything in the report points to one question: what should *you* do about it? This closing chapter lays out three (plus one) scenarios for the decade ahead, a base-case timeline for when reliable, general-purpose robots actually arrive, and a concrete playbook for each audience – investors, engineers, policymakers, enterprise buyers, suppliers and integrators. If you read one chapter to decide your next move, it's this one.

24 – The Brain Score Directory

Forty foundation models, one comparable scorecard. The report closes with humanoid.guide's independent directory – every model that matters, each rated across ten capability dimensions, from locomotion and dexterity to reasoning and long-horizon planning. Which "brains" can actually do what they claim, and how does the whole field line up when you put it on a single page? This is the reference you'll keep coming back to.

This is a free preview of the Humanoid Foundation Models – World Models Report 2026. The full report answers every question above across 100 pages, 16 figures and the complete 40-model Brain Score directory.
humanoid.guide · © 2026 · human@humanoid.guide

Get the full Report

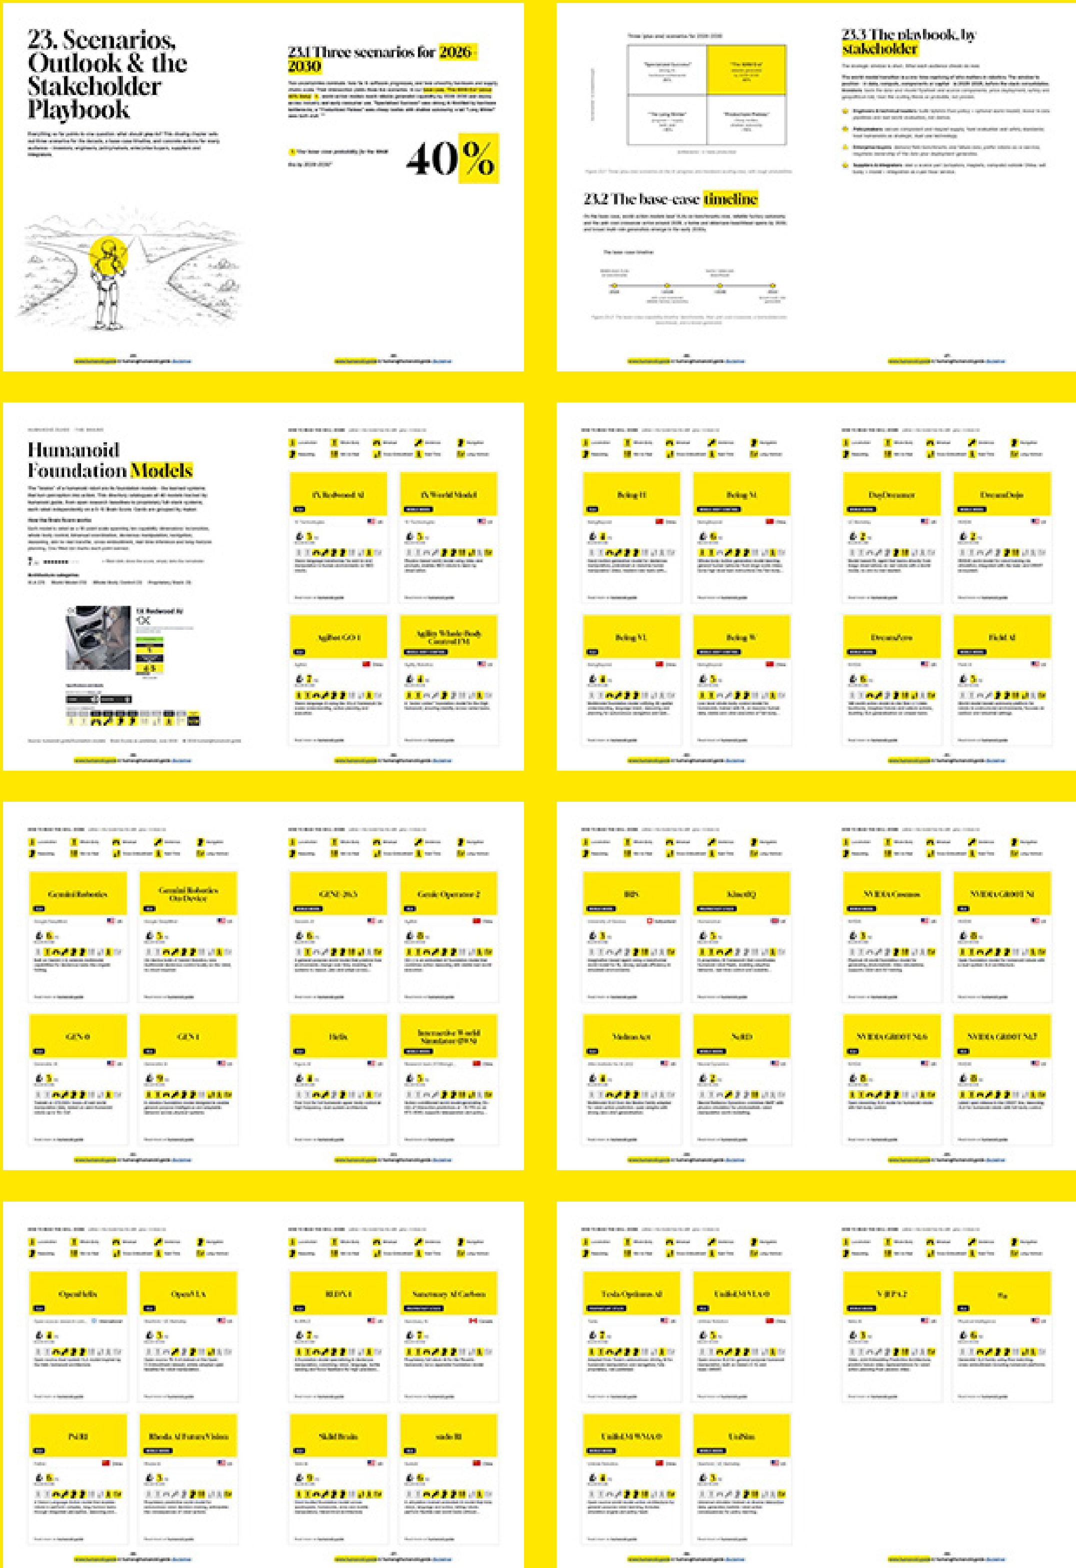

Thank you!

Humanoid.guide

Who to contact:
www.humanoid.guide
human@humanoid.guide