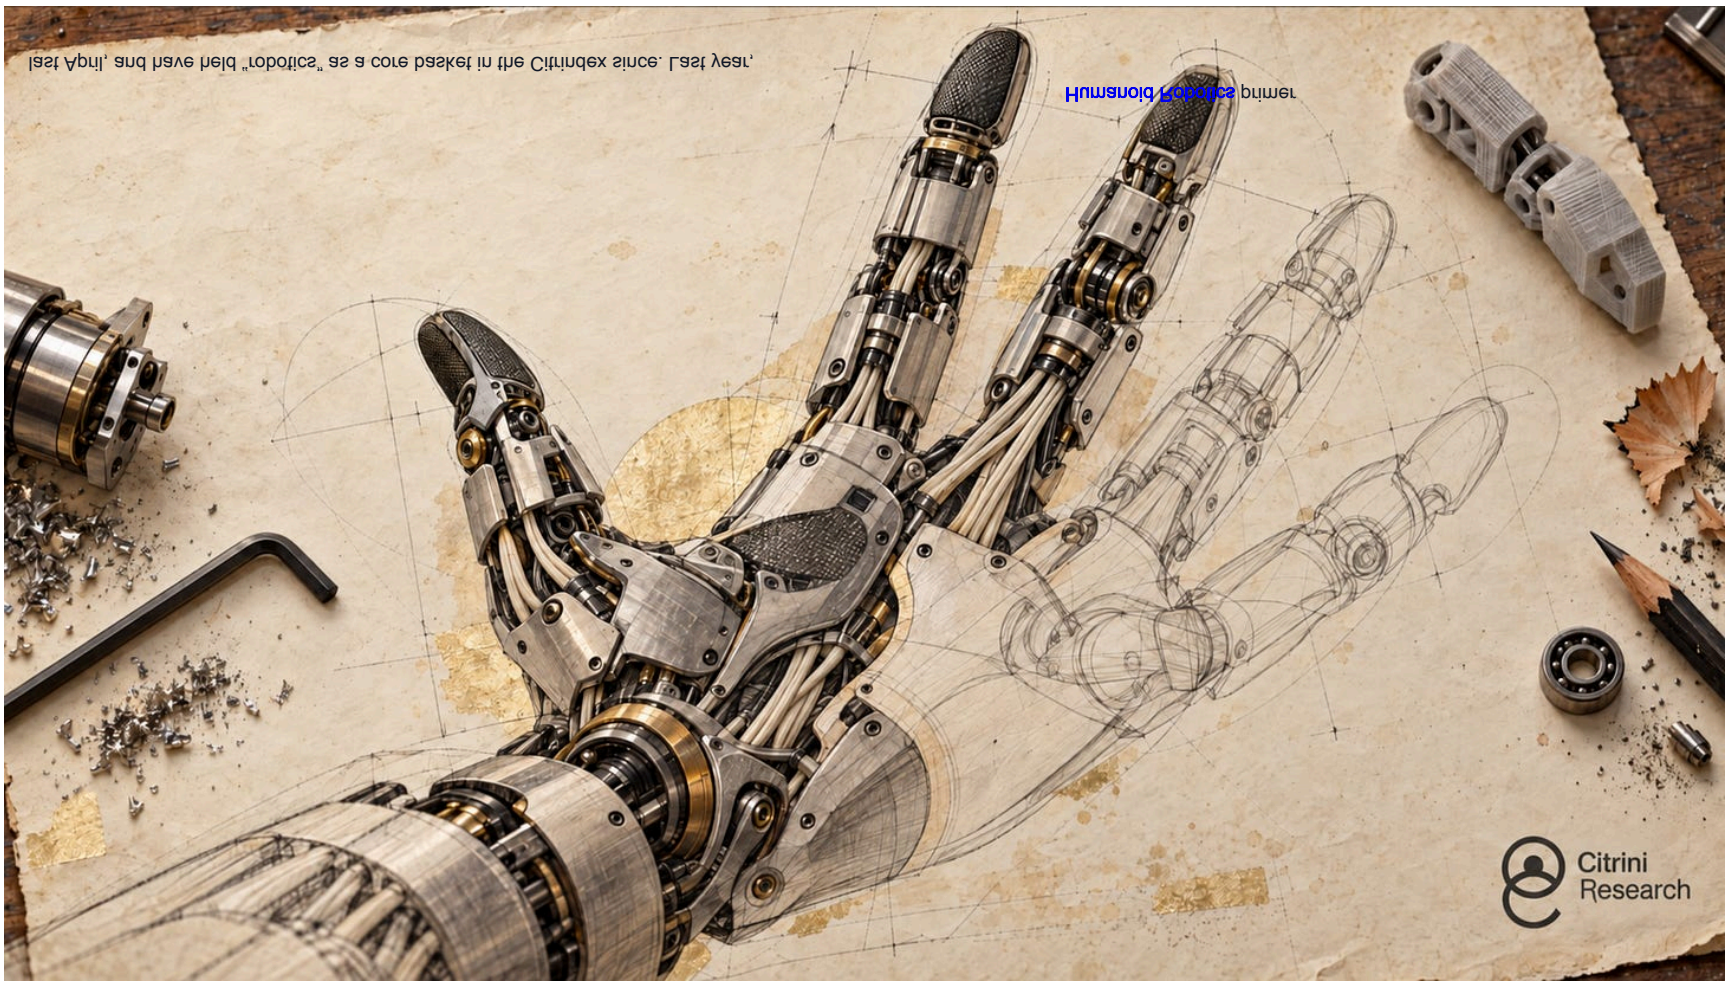
Robotics Tipping Point: A Citrini Field Trip

机器人技术的临界点：Citrini 实地调研

Seeing is Believing

眼见为实



It was 60 degrees, damp and windy – August in San Francisco.

气温华氏60度（约15.5摄氏度），潮湿多风——这是旧金山的八月。

Between graffitied picnic tables in a dive bar's backyard, between local IPAs and cheap lagers, I listened to the "*reindustrialize bros"* talk about slicing US-sourced magnet slabs, building electric motors for military drones, and machining actuators for robots *Made in America*.

在一家破旧小酒馆的后院，在涂鸦遍布的野餐桌之间，就着本地IPA精酿与廉价拉格，我听那群*"再工业化兄弟们"*谈论切割美国本土磁体板料、为军用无人机制造电动机，以及为"*美国制造*"的机器人加工执行器。

The energy was palpable – a group of mostly young guys passionate about building technology that exists in your hand rather than on a screen. No one could doubt their enthusiasm, but I needed to see it for myself.

现场的劲头清晰可感——这是一群以年轻人为主的伙计，热衷于打造那种握在手中、而非只存在于屏幕上的技术。没有人能否认他们的热情，但我需要亲眼验证。

Under its new branding of *"Physical AI"*, is *robotics* evolving from Bay Area backwaters to a true technological renaissance? Or was this happy hour more like a renaissance fair – a romantic longing for an irretrievable past when the US made things?

在*"物理AI"*这一全新标签之下，*机器人技术*究竟是正从湾区的边缘角落走向一场真正的技术复兴？还是说这场欢乐时光更像一个文艺复兴集市——对美国曾经制造实物的那个一去不返时代的浪漫怀想？

As a Citrini analyst on a field trip – I wasn't leaving without some answers.

作为一名前来实地考察的Citrini分析师——不带走一些答案，我是不会离开的。

At Citrini, we've dubbed robotics a thematic mega-trend. We've followed the industry and its enablers dating back to our

我此行就是要揭开帷幕。 在湾区待了一周，我走访了十几家头部实验室，观看了机器人实际运行，并不断追问工程师和投资者，试图弄清美国机器人行业的*真实*状况。最重要的是，这一切我都亲眼见过——用我这双血肉之躯的眼睛。

What I found changed my view of Robotics broadly, my understanding of form factors and neural networks, my handicap of commercialization timeline, ***my view of the trade…***

此行所见改变了我对机器人行业的整体判断、我对硬件形态与神经网络的理解、我对商业化时间表的预估，也改变了***我对这笔交易的看法……***

Note to Readers on Citrini's New Ownership

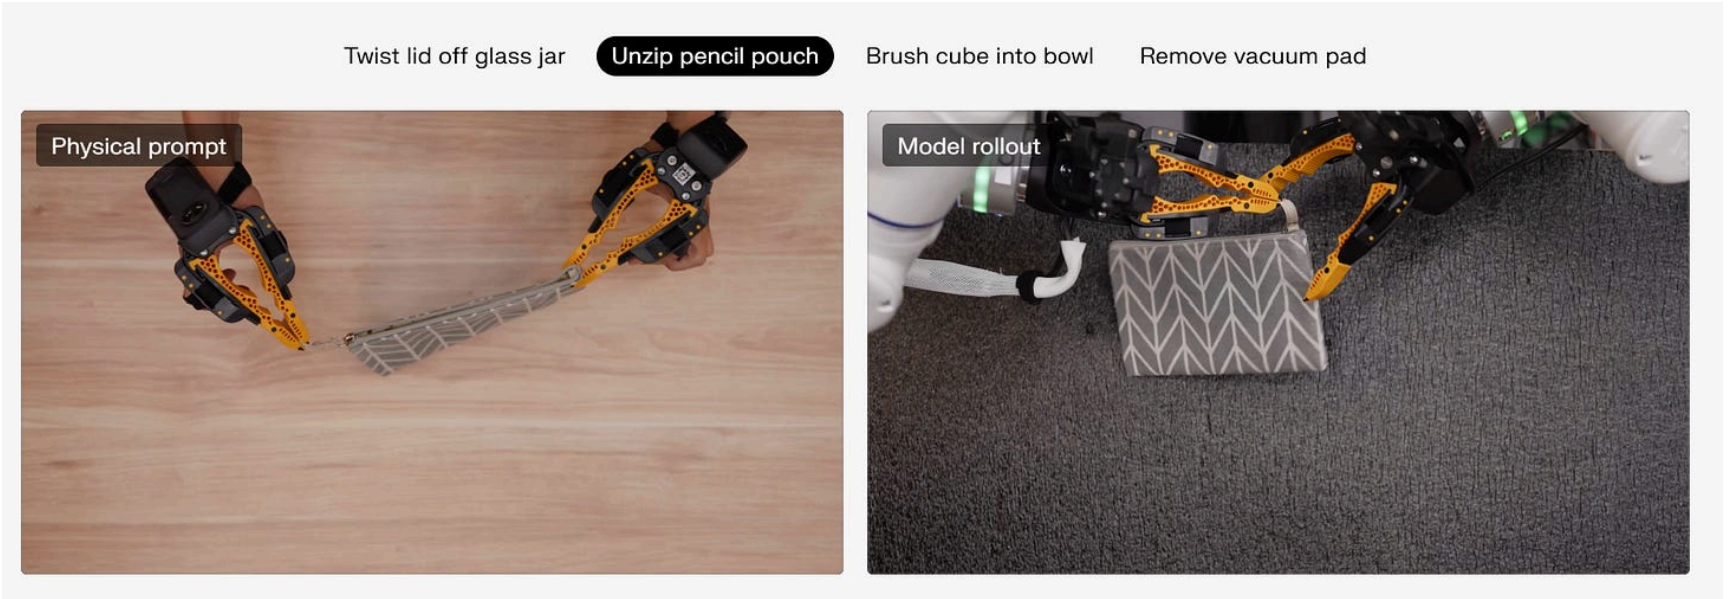
announced GEN-1.5, a "one-shot" learner to significant fanfare, showing that a bot could complete simple tasks of manipulation based on instructions *without* task-specific training.

八月，GeneralistAI

这种“上下文内学习”表明，该机器人已对物理世界（物体是什么、如何运动）具备足够广泛的理解，因此无需依靠海量数据以暴力穷举的方式去攻克特定任务。

Technical detractors, however, question whether this is true generalization, or whether the “distribution” of pretraining data provided sufficiently similar (but technically distinct) examples. Non-technical detractors suggest that opening a pencil case is far from “technological revolution”.

然而，技术层面的质疑者怀疑这是否属于真正的泛化能力，抑或只是预训练数据的“分布”恰好提供了足够相似（但在技术上有所区别）的样本。非技术背景的质疑者则认为，打开一个铅笔盒距离“技术革命”还很遥远。



Seeing an infant begin to understand and manipulate the world is *awe-inspiring* but a far cry from high utility. Even if you know it’s learning, you know there is a long way to go.

看到一个婴儿开始理解并摆弄身边的世界，固然令人*惊叹*，但距离高实用性还相差甚远。即便你知道它在学习，你也清楚前路依然漫长。

Not to be outdone, Skild AI followed up with a very similar announcement of “in-context” demonstrations of its S1 model though with more complex tasks lasting up to 10 minutes – most amusing in my opinion, cooking and flipping a pancake!

不甘示弱，Skild AI 随后发布了内容非常相似的公告，展示了其 S1 模型的“上下文内”（in-context）演示，不过任务更为复杂，时长最长可达 10 分钟——在我看来最有趣的一幕，当属做饭和翻煎饼！



(We’ll get much deeper on this brain development later.)

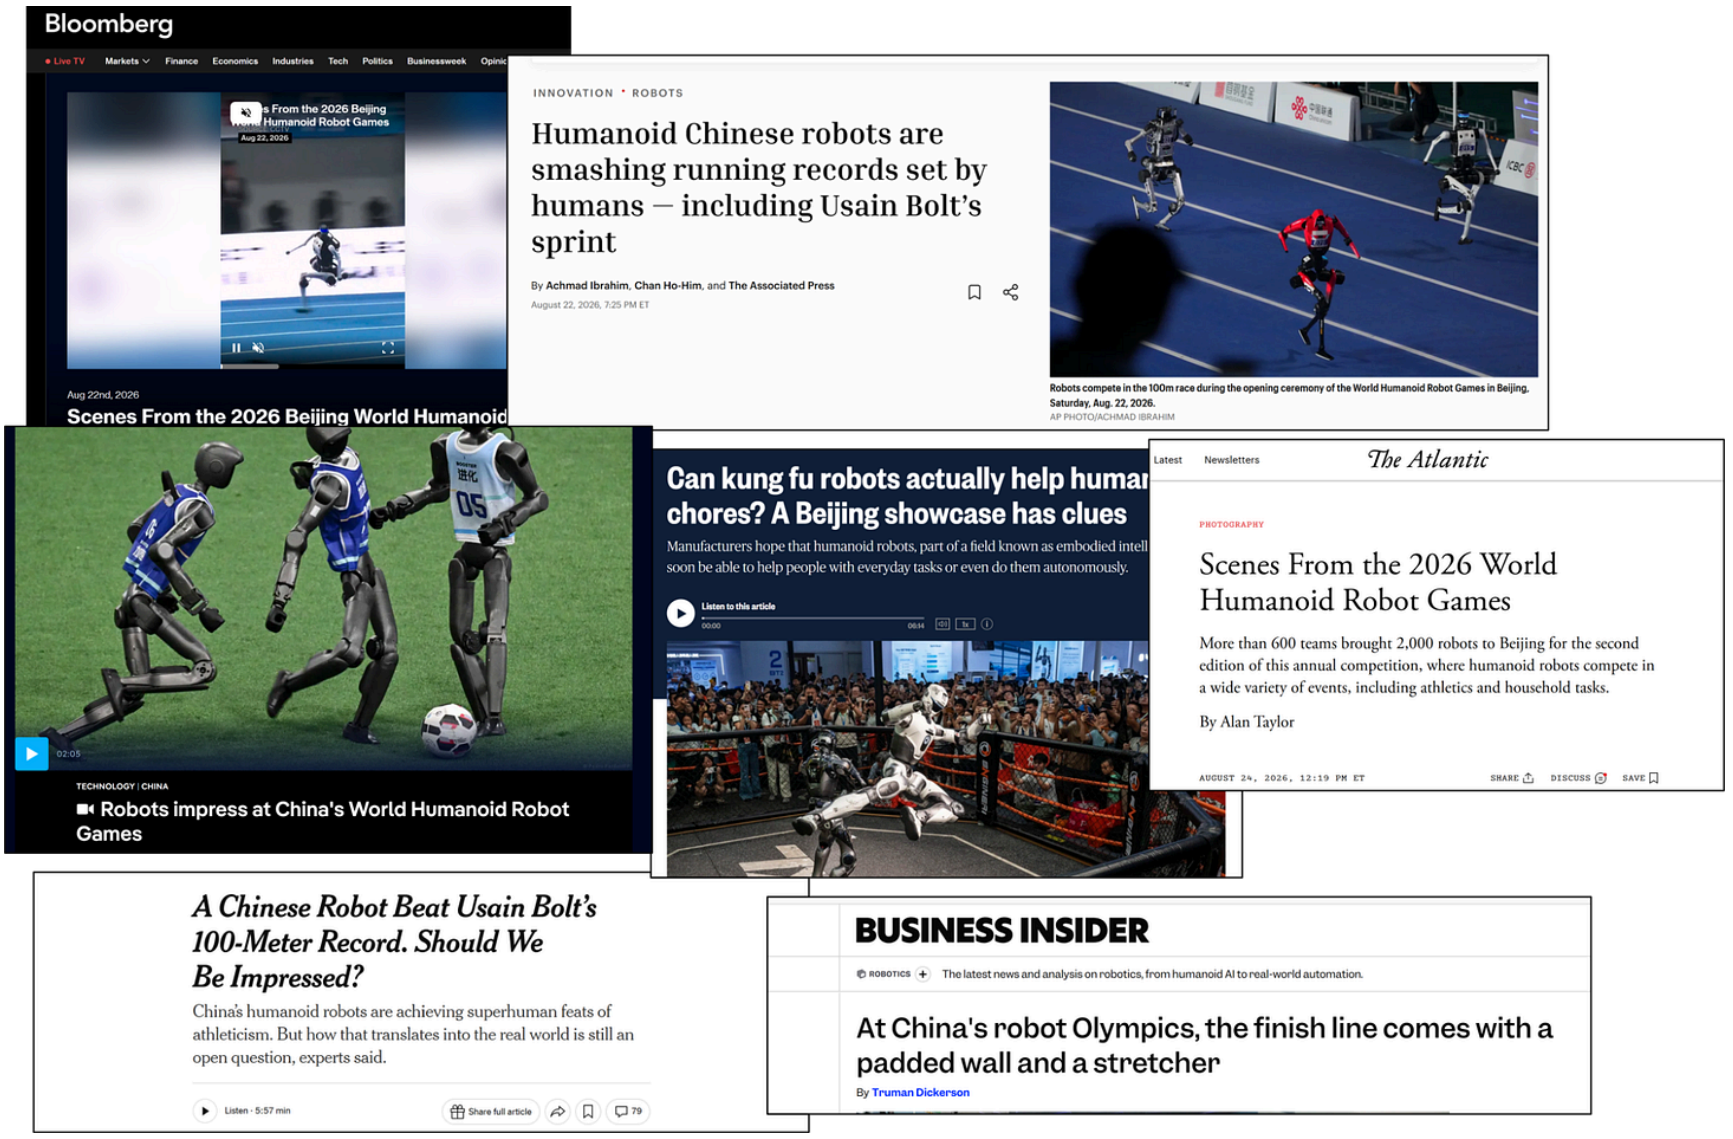
（关于这一大脑开发的话题，我们稍后会深入探讨。）

The language here is intentional. The obvious allusion is OpenAI’s May 2020 paper: *“Language Models are Few-Shot Learners”* which showed a massive breakthrough in GPT-3’s generalized reasoning, which improved dramatically with pre-training scale – implicitly suggesting their own future will follow that of ***TradAI*** (“traditional AI”, as I’m trying to coin it).

这里的措辞是刻意为之。最明显的典故指向 OpenAI 2020 年 5 月的论文：*"Language Models are Few-Shot Learners"*（《语言模型即少样本学习者》），该论文展示了 GPT-3 在通用推理上的重大突破，且这一能力随预训练规模的扩大而显著提升——这隐含地暗示着他们自身的未来将沿着 ***TradAI***（“传统 AI"，我正尝试创造这一说法）的轨迹发展。

But if the ChatGPT litmus test is *universal recognition*, last month’s World Robotics Games in China was arguably closer to the mark.

但如果说 ChatGPT 这块试金石衡量的标准是*广为人知*，那么上个月在中国举办的世界机器人运动会可以说更贴近这一标准。

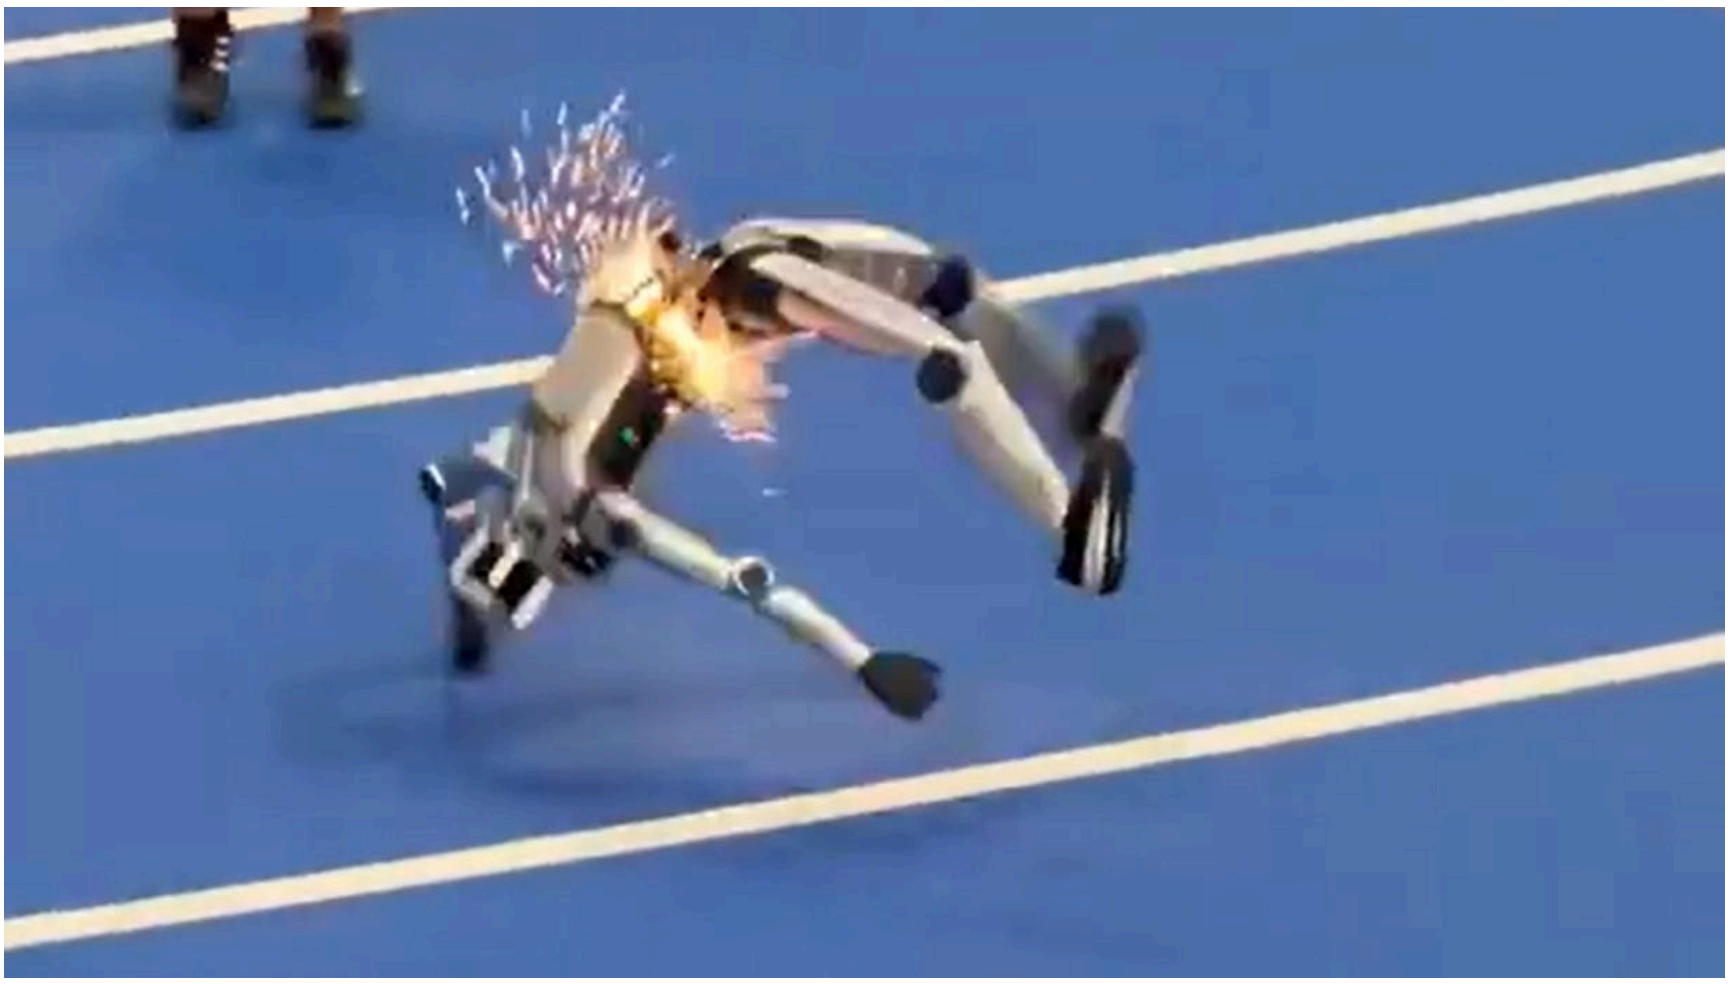


A flood of viral videos emerged from Beijing showing the advancements in humanoid-form hardware: bots running faster than Usain Bolt, high jumping to ceiling height, and occasionally crashing in comical plumes of sparks and smoke. Unitree's concurrent IPO helped attract the financial world's attention.

北京涌现出大量刷屏视频，展示了人形机器人硬件的进步：机器人奔跑速度快过博尔特（Usain Bolt），跳高能触及天花板，偶尔还会在滑稽四溅的火花与烟雾中摔毁。宇树科技（Unitree）同期推进的 IPO 也帮助吸引了金融界的关注。

Humanoid Olympics may seem like an exercise in the absurd, but really are a top-down strategic effort to push the limits of hardware in a competitive, visible and verifiable setting. *Sending a humanoid at full sprint into a crash wall is literally* ***"moving fast and breaking things".***

人形机器人奥运会看似一场荒诞的表演，实则是一项自上而下的战略行动，旨在通过一个竞争性、公开可见且可验证的场合推动硬件的极限。*让一台人形机器人全速冲刺撞向撞击墙，堪称名副其实的*"**快速行动，打破一切"**。*



Scoff at your own risk. Americans watching this little guy going out in a literal blaze of glory should be upset it's not happening in their country.

嗤之以鼻，后果自负。看着这个小家伙真真切切地在一片烈焰荣光中谢幕，美国人应该感到懊恼的是：这样的场面并没有发生在自己的国家。

The locomotion of these bots, the speed of iteration in hardware, and a growing list of serious builders, is undoubtedly impressive. But we've seen impressive, tele-operated and pre-programmed acrobatics for years. Remember Boston Dynamics caught the world's attention with its humanoid back in 2009. But these machines remain a *Sharper Image* novelty without a brain.

这些机器人的**运动能力**、硬件迭代的速度，以及越来越多认真做事的制造商，无疑都令人印象深刻。但那种令人惊艳的遥操作与预编程杂技，我们已经看了好些年。别忘了，波士顿动力早在2009年就凭借其人形机器人吸引了全世界的目光。然而，这些机器至今仍是没有大脑的 *Sharper Image* 式新奇玩具。

Ok… but what about back in May, when Figure AI hosted a nine day livestream of a furiously-dedicated humanoid (team) sorting 250,000 packages – a blend of both autonomous cognition, dexterous manipulation and full body, locomotion, and real world application?

好吧……但五月份那件事又怎么说？当时 Figure AI 进行了为期九天的直播，一群拼尽全力的人形机器人（团队）分拣了25万个包裹——这可是自主认知、灵巧操作、全身运动能力与真实世界应用的结合。



Each of these examples misses the mark in some way. We've shown the key attributes – early intelligence and autonomy, highly impressive locomotion, real-world automation. What's less obvious is all three rolled into one indisputably useful package.

上述每一个例子都在某种程度上有所欠缺。我们已经展示了关键属性——早期的智能与自主能力、令人印象深刻的运动能力、现实世界的自动化。不那么显而易见的是，将这三者集于一身、打包成一个无可争议的实用产品。

I think the "ChatGPT moment" is a red herring *(perhaps a cliché of its own).*

我认为"ChatGPT 时刻"是一种误导（这本身或许也已成了一种陈词滥调）。

There might be a massive "aha" moment, but more likely it will be a *mosaic of signals* that tell us we've reached the tipping point. It will be only *after* the breakthroughs occur and products appear in the real world that their utility is widely acknowledged.

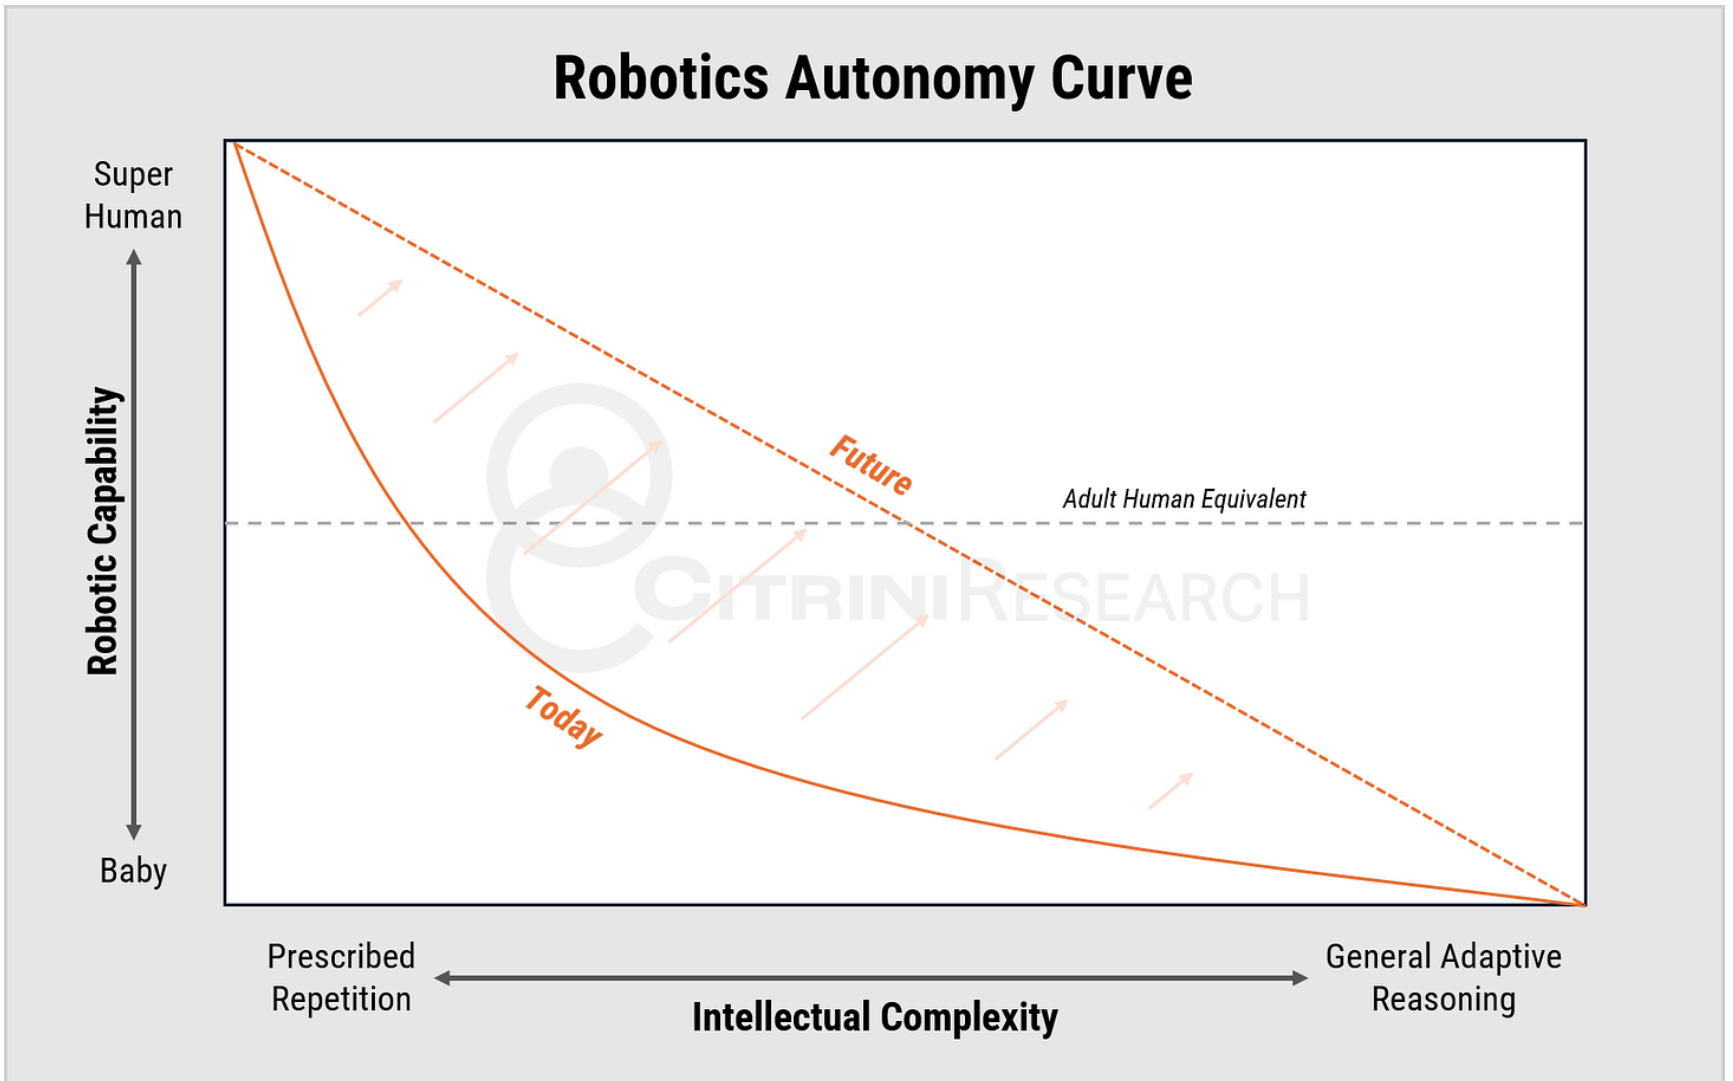
或许会出现一个巨大的"顿悟"时刻，但更可能的是，一组*拼图式的信号*会告诉我们已经到达临界点。只有*在*突破发生、产品走进现实世界*之后*，其效用才会得到广泛认可。

As far as where we stand today, here's what I think is true. The basic locomotion of humanoids is a solved problem that's only getting cheaper and more accessible (geopolitical and supply chain considerations notwithstanding). Complex dexterity is meaningfully more complicated but is being successfully attacked through multiple hardware approaches. Simple repetitive tasks like package sorting are already readily trainable, unlocking automation opportunities and near-term commercialization while models and bots improve. Complex tasks are coming along faster than you think.

至于我们今天所处的位置，以下是我的判断。人形机器人的基础运动能力是一个已经解决的问题，而且正变得越来越便宜、越来越触手可及（尽管存在地缘政治和供应链方面的考量）。复杂的灵巧操作则明显更为困难，但正在通过多种硬件路径被成功攻克。像包裹分拣这类简单的重复性任务如今已可轻松训练，在模型和机器人不断改进的同时，打开了自动化机遇与近期商业化的空间。复杂任务的到来速度比你想象的更快。

Robotic autonomy is more of a *curve* than a switch, and it looks something like this.

机器人的自主性更像一条*曲线*，而非一个开关，大致如下所示。



Broad generalizability is the furthest goal, but all the ingredients are there so clearly that it seems hard to bet against. Every day the curve pushes a bit further to the right.

广泛的通用能力是最远期的目标，但所有要素都已清晰到位，要押注它不会实现，似乎很难。这条曲线每天都在向右多推进一点。

Maybe the best answer to the cliché is to embrace it. If you had a time machine, ***would you invest in AI before its ChatGPT moment***?

也许应对这个陈词滥调的最佳方式，就是坦然接受它。如果你有一台时光机，***你会在 AI 迎来自己的“ChatGPT 时刻”之前投资 AI 吗***？

But before you get robo-pilled, we need to get into the details. There's a lot of ground to cover, but if you bear with me, it will be well worth it *(between you and me, this is my favorite piece I've written)*.

但在你彻底“入坑”机器人叙事之前，我们需要深入细节。要讲的内容很多，不过如果你耐心读下去，一定会物有所值*（悄悄说一句，这是我自己最满意的一篇作品）*。

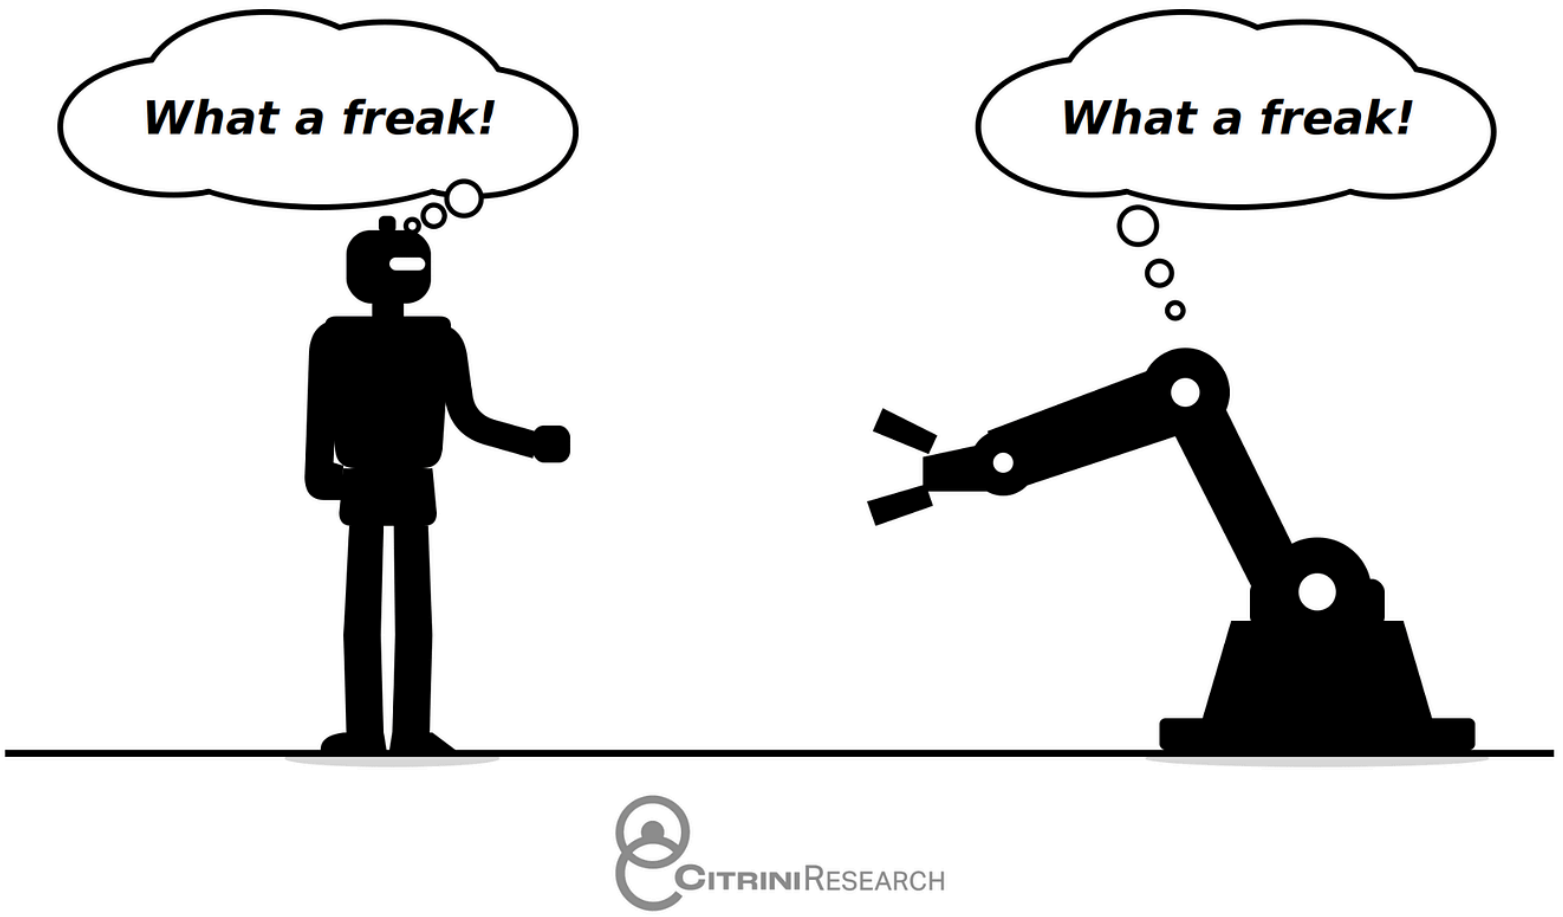
1. The Body: Hardware and Humanoids
2. **躯体：硬件与人形机器人**
3. The Brain: Grading the Class
4. **大脑：给全班打分**
5. Capability and Commercialization
6. **能力与商业化**
7. Four Predictions for Robotics
8. **对机器人领域的四项预测**
9. *What's the Trade?*
10. ***交易机会在哪里？***

1) The Body: Hardware and Humanoids

1) 躯体：硬件与人形机器人

Before we move forward, let's address the most polarizing word in robotics – *humanoid.* The very tired debate goes like this: why should a robot look like a human?

在继续之前，先来谈谈机器人领域最具争议性的一个词——*人形机器人*。那场早已令人厌倦的争论是这样的：机器人凭什么要长得像人？



The anti-humanoid camp has most of robotics history in its corner. There are many robots today and almost none look like humans (or at least not *full* humans). They are designed for a specific task which they perform with superhuman speed and precision. Our modern world relies on these workhorses, and the remaining cases that still require human handling are just scraps — small TAM edge cases.

反人形阵营拥有几乎整部机器人发展史的支持。如今的机器人已经很多，但几乎没有长得像人的（至少不像*完整的*人）。它们为特定任务而设计，并以超越人类的速度和精度执行。现代世界依赖这些主力机型，而剩余仍需人工处理的场景不过是边角料——TAM 很小的边缘案例。

The pro-humanoid camp says a hotdog maker can make 1,000 hotdogs in a day, but what if I want a burger? If we've already automated all repetitive and programmable tasks, then *what's left* requires something new – a one-size-fits-all Swiss Army knife that can occupy and manipulate a world that has been built for humans. Not a package sorter, but a universal worker that can sort packages if you ask it.

支持人形阵营则认为：一台热狗机一天能做 1,000 个热狗，可如果我想吃汉堡呢？如果我们已经把所有重复性和可编程的任务都自动化了，那么*剩下的*工作就需要一种新事物——一把万能的瑞士军刀，能在一个为人类打造的世界中行动和操作。它不该只是一台包裹分拣机，而应是一个通用工人——你让它分拣包裹，它就能分拣。

The reason why this hasn't happened yet is because we haven't had robots with enough generalizable intelligence and dexterity to be worthwhile. The pro-humanoid camp has most of science fiction's imagination on its side.

之所以至今未能实现，是因为我们还没有拥有足够可泛化的智能与灵巧性、从而值得投入的机器人。而科幻作品中的想象，大多站在支持人形阵营一边。

On the Humanoid Spectrum

关于人形光谱

This is yet another false dichotomy. There aren't *two* categories – rather there is a wide range of approaches using the same intelligence architecture to tackle different tasks. Often at least some part of them looks human-y.

这又是一个虚假的二分法。实际上并不存在*两种*类别——而是存在一系列广泛的技术路线，它们采用相同的智能架构来应对不同任务，其中往往至少有一部分形态带有人形色彩。

After all, the programmable ***robotic arm*** is the most ubiquitous and useful robot form in history. If we just call every factory arm or cobot a "partial humanoid" it would put the debate to rest.

毕竟，可编程***机械臂***是历史上最普及、最实用的机器人形态。如果我们就把每台工厂机械臂或协作机器人都称作"部分人形机器人"，这场争论也就该尘埃落定了。



Universal Robotics

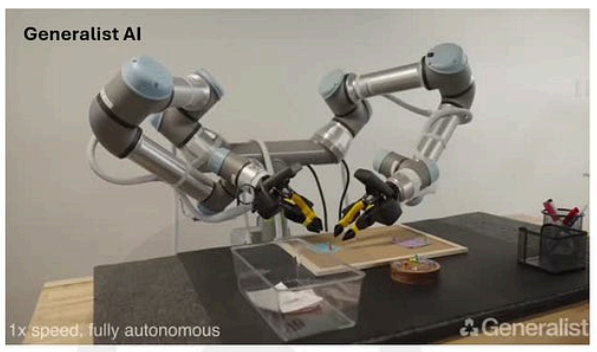


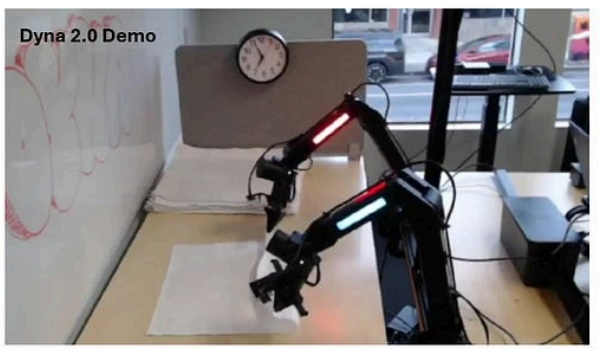
Rainbow Robotics



FANUC

Arms / Cobots





Typical Research Setup:
Two Arms with Overhead and Hand-Mounted Cameras



Unitree



Weave Robotics



Sunday Robotics



Dexmate

Flexible Torso + Mobile Base



Figure 0.3



Unitree G1



Galbot ET1



Apptronik Apollo



Tesla Optimus

Full Humanoids

Where you fall on the humanoid spectrum depends on what you are trying to do. Indeed many humanoid projects start with off-the-shelf arms from companies like Universal Robots, FANUC, or Rainbow.

你在人形机器人谱系上处于哪个位置，取决于你想做什么。事实上，许多人形机器人项目一开始就直接采用 Universal Robots、FANUC 或 Rainbow 等公司的现成机械臂。

The brains start with table-mounted manipulators. For researchers *training* physical AI models – Physical Intelligence, Generalist, Skild, or Dyna – the setup is typically a frankenstein-half torso strapped to a table: two cobot arms or manipulators with some kind of “hand” attached and an overhead stereo camera and cameras mounted on each wrist. This is a cheap and easy setup to test and train models without having to think about custom hardware.

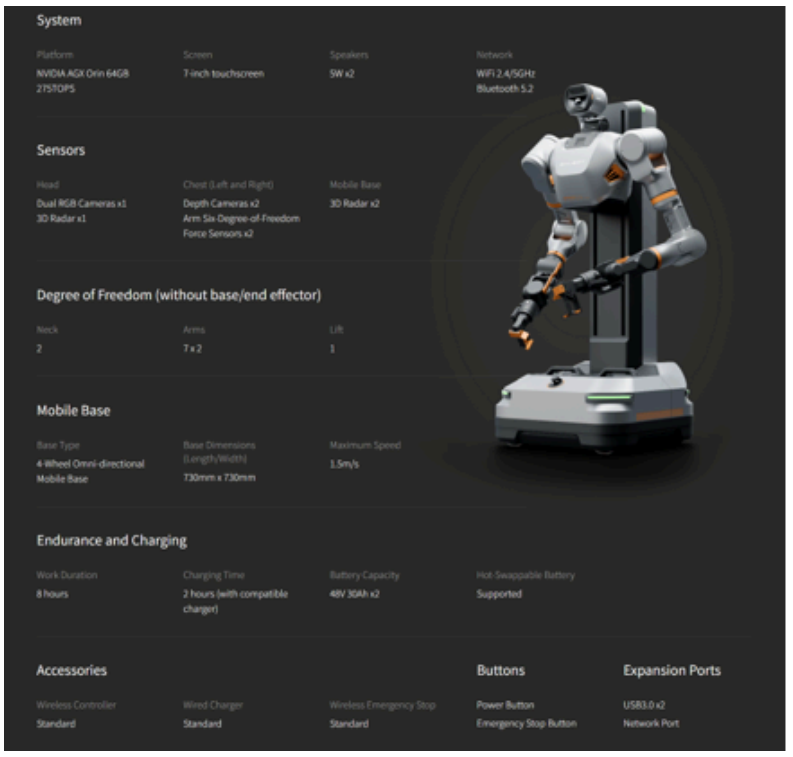
至于“大脑”，则往往从桌面式操作臂起步。对于*训练*物理 AI 模型的研究机构——Physical Intelligence、Generalist、Skild 或 Dyna——典型配置是固定在桌面上的一套“拼装”而成的半身躯干：两条协作机械臂或操作臂，末端装上某种“手”，头顶上方一个立体相机，两个腕部各装一台相机。这种配置便宜且简单，可用于测试和训练模型，而无需考虑定制硬件。

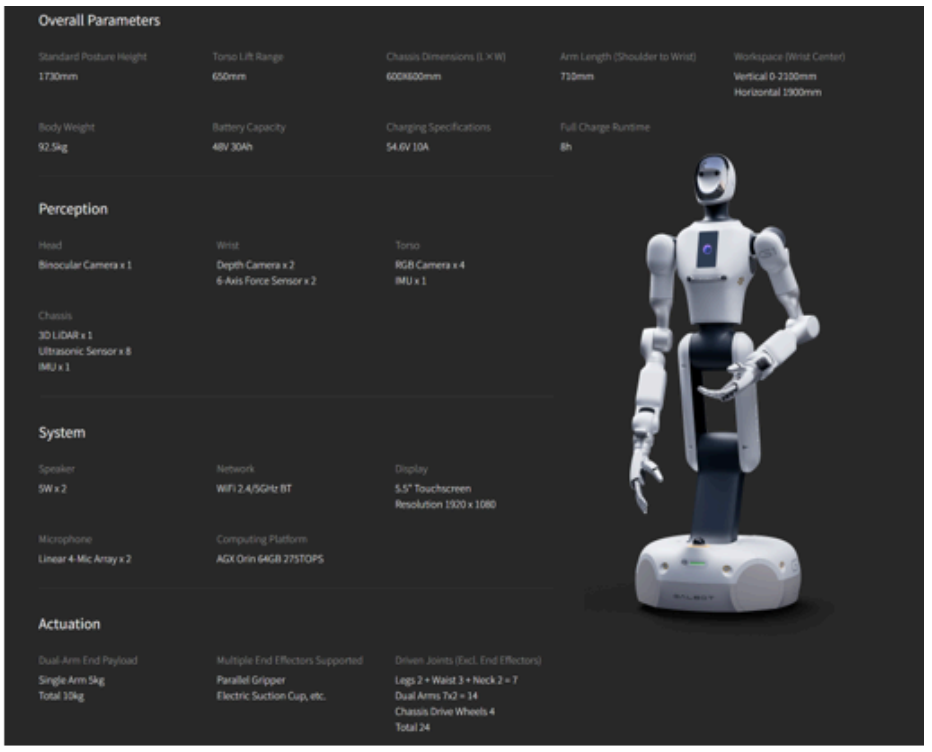
Once you have a robot that can do ***something useful*** on a table, the easiest way to add basic mobility to the torso is through a fixed (often wheeled) base. The tradeoffs are obvious – they are far cheaper and simpler but they sacrifice lower-half dexterity and mobility and are incapable of stairs, terrain or any number of other mobility challenges.

一旦你有了能在桌面上做***某些有用的事***的机器人，为躯干增加基本移动能力的最简单方式，就是配一个固定（通常为轮式）底座。其取舍显而易见——这类底座便宜得多、结构也更简单，但牺牲了下半身的灵巧性与移动能力，无法应对楼梯、复杂地形以及其他诸多移动挑战。

Many early attempts at task-specific commercialization have taken this route. China’s Galbot – whose humanoid was one of the standouts from the Beijing games – currently markets two mobile-based torso’d robots, the S1 and G1 for industrial and personal uses respectively.

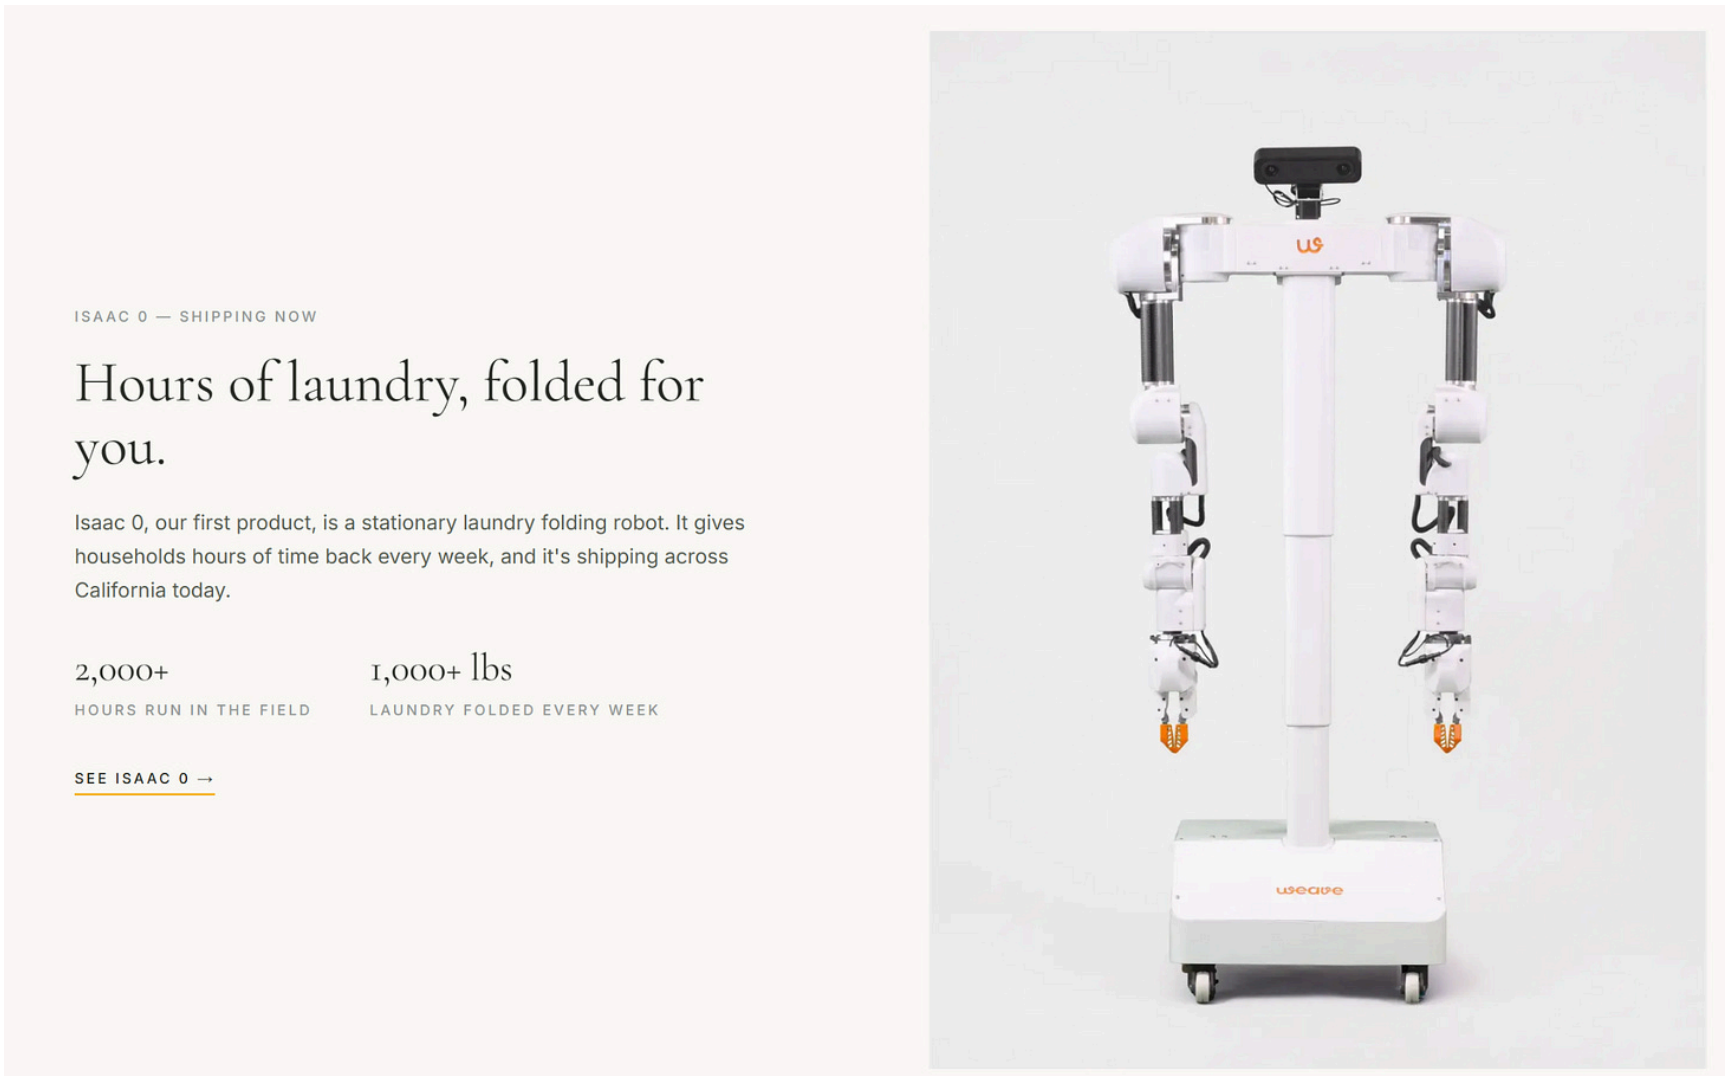
许多面向特定任务的早期商业化尝试走的都是这条路。中国的 Galbot——其人形机器人是北京运动会上的亮点之一——目前销售两款带移动底座的躯干式机器人：S1 和 G1，分别面向工业和个人用途。





Another example, Weave Robotics, has trained a laundry folding robot with its first iteration Isaac 0 actually now operating in homes and businesses (*if you live in California you can order one today for $249/month!*)

另一个例子是 Weave Robotics，该公司训练出了一台叠衣机器人，其第一代产品 Isaac 0 目前已实际投入家庭和商业场所运行（*如果你住在加州，今天就能以每月 249 美元的价格订购一台！*）



But what about stairs? Can it load and unload the machine? Does its folding quality pass inspections?

但楼梯怎么办？它能往机器里装取衣物吗？它的折叠质量能通过检验吗？

It's far from a full-fledged robo-butler, but the argument for early deployments is that you have to start putting things into the field in order to learn what works and what doesn't, begin to ramp production scale and drive down costs. (Weave's next iteration Isaac 1 will be out shortly with better capabilities and a cuter demeanor.)

它距离功能齐全的机器管家还很远，但支持早期部署的理由在于：你必须先把产品投入实际场景，才能知道什么可行、什么不可行，并着手扩大生产规模、压低成本。（Weave 的下一代产品 Isaac 1 即将问世，能力更强，神态也更讨喜。）

Full humanoids are the most complex and expensive, but replicate the full mobility and near dexterity of the human form. Something that can climb stairs, run, jump, crouch, crane. From a simple hardware perspective, locomotion and full body dexterity are essentially solved. Pre-programmed dance routines and gymnastic displays demonstrate fluidity, balance and grace, even verging into "superhuman" territory when designed for specific tasks as demonstrated in China this month.

完整人形机器人是其中最复杂、最昂贵的一类，但能复制人体形态的全部移动能力与接近人类的灵巧性。这类机器人能爬楼梯、奔跑、跳跃、下蹲、伸颈探看。仅从硬件层面看，运动控制与全身灵巧性已基本解决。预先编排的舞蹈与体操展示体现了流畅、平衡与优雅；当为特定任务而设计时，甚至已逼近"超人"水平——本月在中国展示的案例便是例证。

And as designs become more ubiquitous and production rises, the cost of humanoids will trend lower while utility grows.

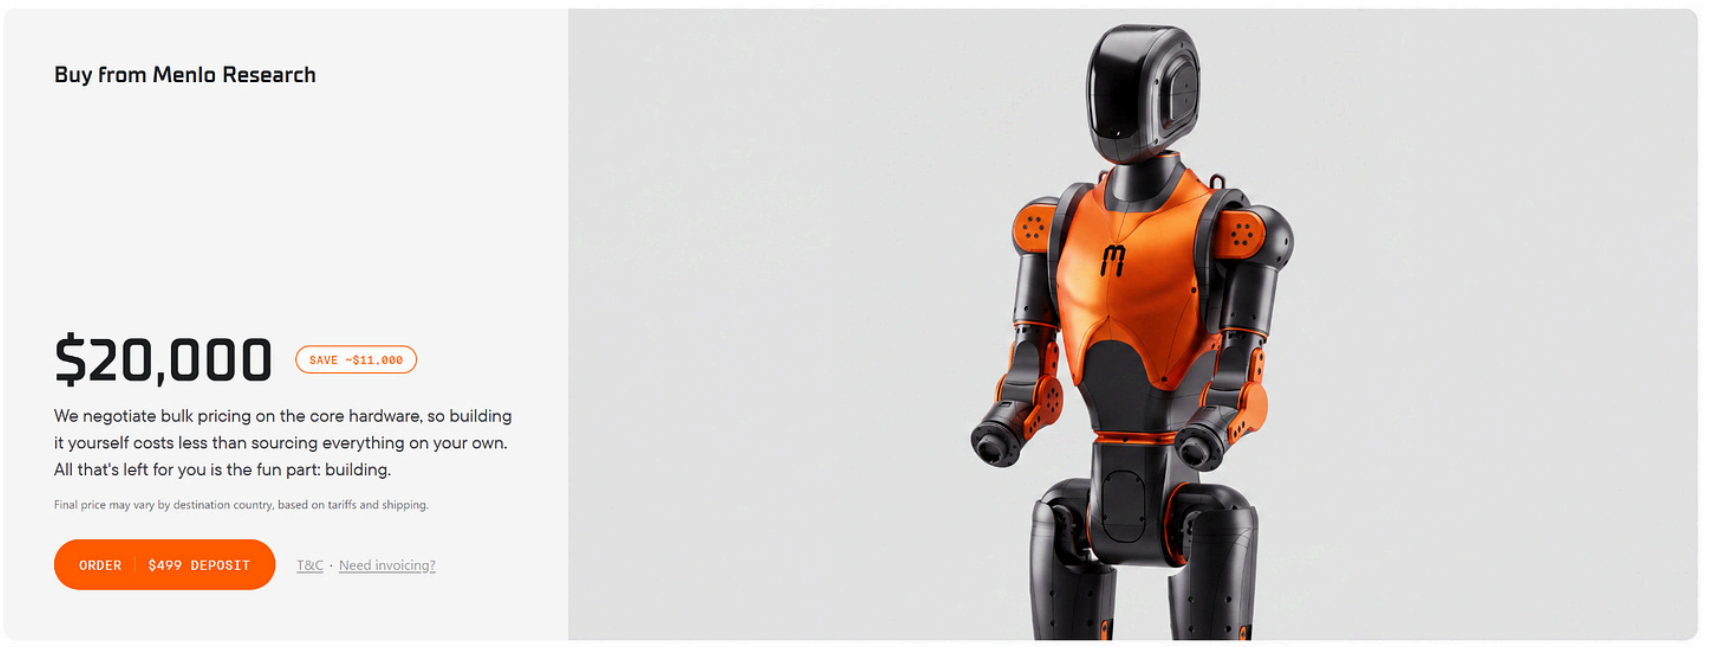
而随着设计日趋普及、产量不断提升，人形机器人的成本将趋于下降，实用性则会持续增长。

Unitree's G1 humanoid, with its ultra-low cost (currently listed at $13,500), solid locomotion and open API integration for intelligence has been the default product for both entertainment and AI researchers looking to advance humanoid autonomy on a cheap and mostly reliable hardware base.

Unitree（宇树）的 G1 人形机器人凭借超低价格（当前标价 13,500 美元）、扎实的运动能力，以及面向智能层的开放 API 集成，已成为娱乐用途与 AI 研究者的默认产品——后者希望在这一廉价且总体可靠的硬件基础上推进人形自主能力。

Others like Menlo Research are working to establish "open source" humanoid configurations with Asimov, providing a much more transparent, customizable, and similarly low-cost template for builders.

Menlo Research 等其他团队则致力于通过 Asimov 建立"开源"人形机器人配置，为开发者提供一个透明度高得多、可定制性更强、成本同样低廉的模板。

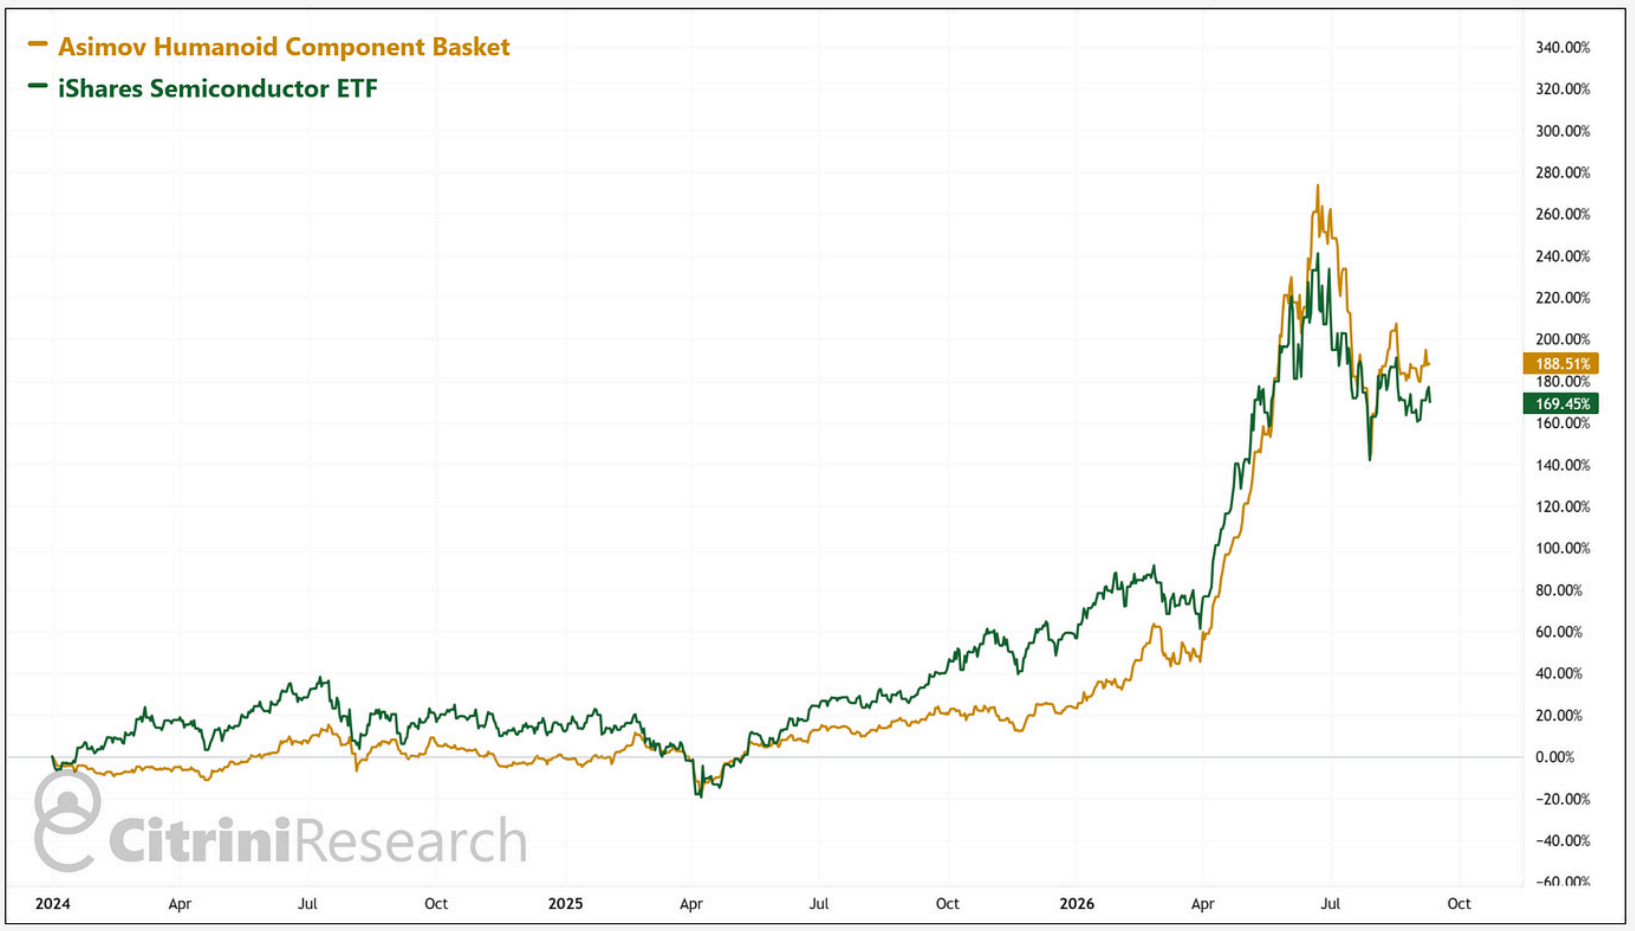


You can see it broken into 1,168 CAD components (thanks to Michelle Sun at CoreMatters and Astra for the breakdown).

你可以查看其拆解出的 1,168 个 CAD 零部件（感谢 CoreMatters 的 Michelle Sun 和 Astra 提供拆解）。

(Since we have an open sourced breakdown, I was curious enough to equal weight a basket of all identifiable public companies in the BOM as a "supply chain" proxy. It looks a lot like the broader semiconductor index, but as we discuss in much more detail, there are more nuanced and uncorrelated ways to express the trade.)

(既然有了开源的拆解清单，我出于好奇，将物料清单（BOM）中所有可识别的上市公司等权编制成一个组合，作为"供应链"的代理指标。它的走势与更宽泛的半导体指数非常相似，但正如我们后文会更详细讨论的那样，还有一些更精细、相关性更低的交易表达方式。)



Source: Citrini Research

On the other end of the humanoid spectrum, OEMs like Figure AI, Tesla's Optimus, XPENG, Apptronik, Neura, Agility and others are focused on creating closed stack solutions with specs held as close to the chest as possible. Their bet is that an integrated and proprietary hardware/software model not only provides the synergies in the brain/body embodiment but also the deepest moat and long-term commercial value.

在人形机器人光谱的另一端，Figure AI、特斯拉的 Optimus、XPENG、Apptronik、Neura、Agility 等整机厂商则专注于打造封闭式全栈方案，并尽可能对技术规格守口如瓶。他们的押注是：一体化的专有软硬件模式不仅能带来"大脑/身体"具身层面的协同效应，还能构筑最深的护城河与长期商业价值。

But look back at Asimov… do you notice anything missing?

但回头看看阿西莫夫……你注意到有什么缺失了吗？

The hands aren't included. Asimov's got stubs. The Unitree $13,500 G1 has dummy prosthetic hands.

手并不在其中。 阿西莫夫笔下的机器人只有残臂。宇树售价 13,500 美元的 G1，配的也只是无功能的装饰性假手。

Humanoid hands remain one of the most fascinating, challenging and expensive problems in the supply chain.

人形机器人的手，仍然是供应链中最迷人、最具挑战性、也最昂贵的问题之一。

All in the Hands: Hardware Answers for Dexterity

一切尽在双手：灵巧操作的硬件方案

The human hand is an incredible feat of engineering, dexterous and durable while capable of both extremely fine motor skills and heavy loads.

人手是一项了不起的工程壮举：既灵巧又耐用，既能完成极其精细的动作，也能承受沉重的负载。

In programmable robotics, replicating a hand isn't necessary. The ***"end effector"*** otherwise known as "*the thing you put at the end of the arm"*, is designed for specific tasks – welding arms have a welding point, warehouse bots might use vacuum grippers that easily attach to boxes.

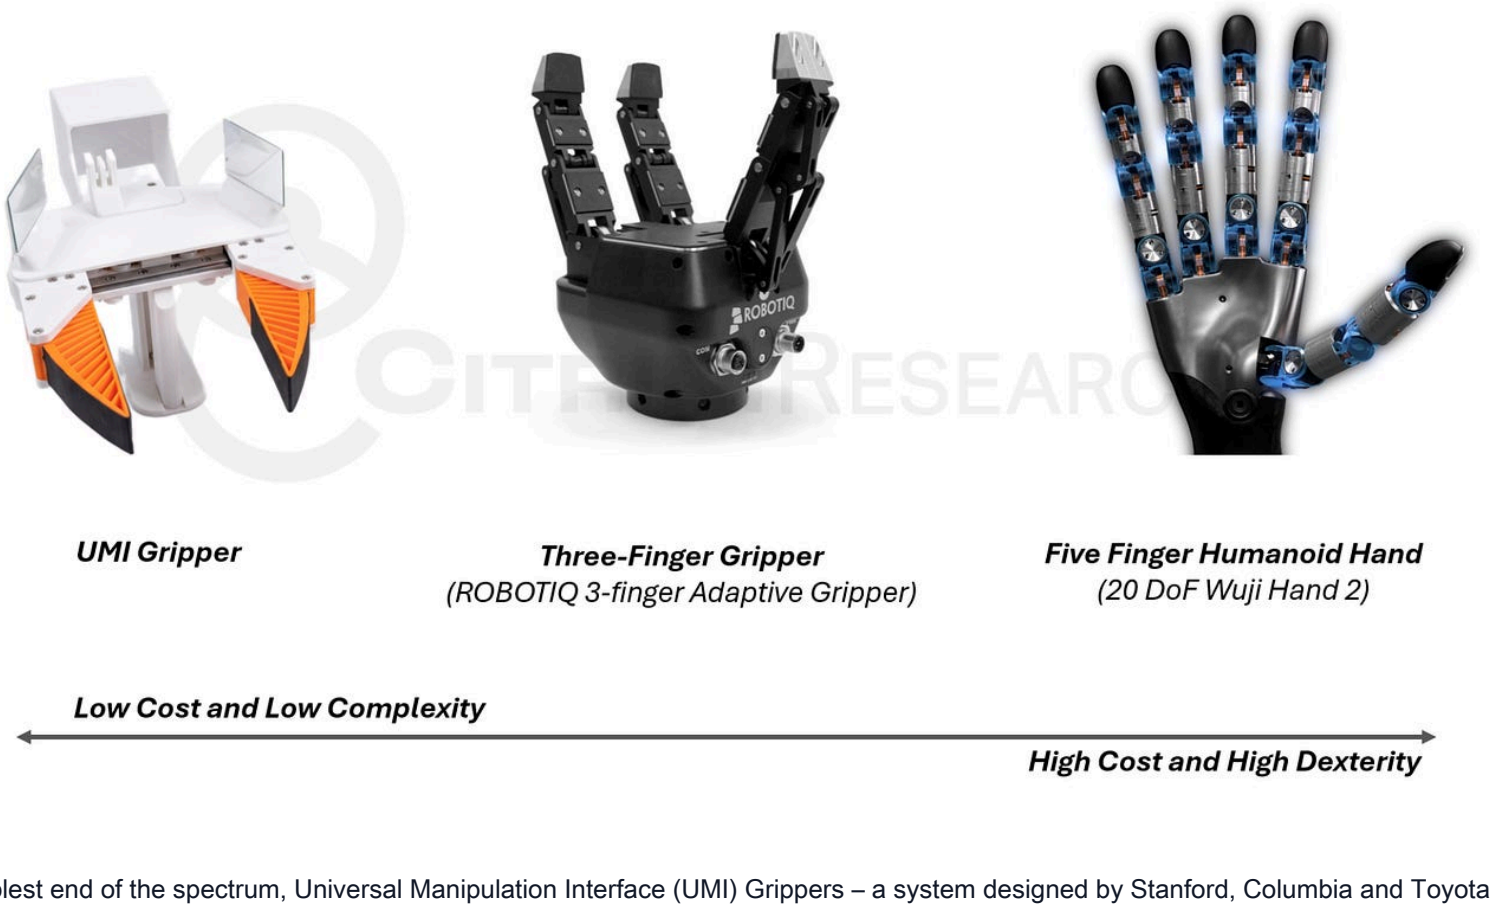
在可编程机器人领域，复制一只人手并非必要。***"末端执行器"***——俗称"*装在手臂末端的那件东西*"——是针对特定任务设计的：焊接机械臂配有焊接点，仓储机器人则可能使用能轻松吸附纸箱的真空吸盘。

But general purpose humanoids? The hand is the "humanoid question" in micro form: if you want a robot that can complete a wide range of tasks currently done by humans, you need something that has the same degrees of freedom, sensitivity and strength.

但通用型人形机器人呢？手正是"人形机器人之问"的微缩版：如果你想要一台能完成人类目前所做广泛任务的机器人，就需要一个具备同等自由度、灵敏度与力量的部件。

The robotics field has yet to come close to rivaling the ones typing these words. But just like the humanoid spectrum, there is a hand spectrum, which ranges from two-pronged grippers all the way to the five-fingered dexterous hands being built both by full-stack OEMs and specialized players. Likewise, the tradeoffs are on complexity, cost, durability, dexterity, and tactile sensing.

机器人领域尚未打造出能与敲下这段文字的双手相媲美的产品。但正如人形机器人存在谱系之分，手部同样有一条谱系：从二指夹爪一直到五指灵巧手，这些产品既出自全栈OEM厂商，也出自专注细分领域的专业厂商。同样，其中的取舍集中在复杂度、成本、耐用性、灵巧性与触觉感知上。

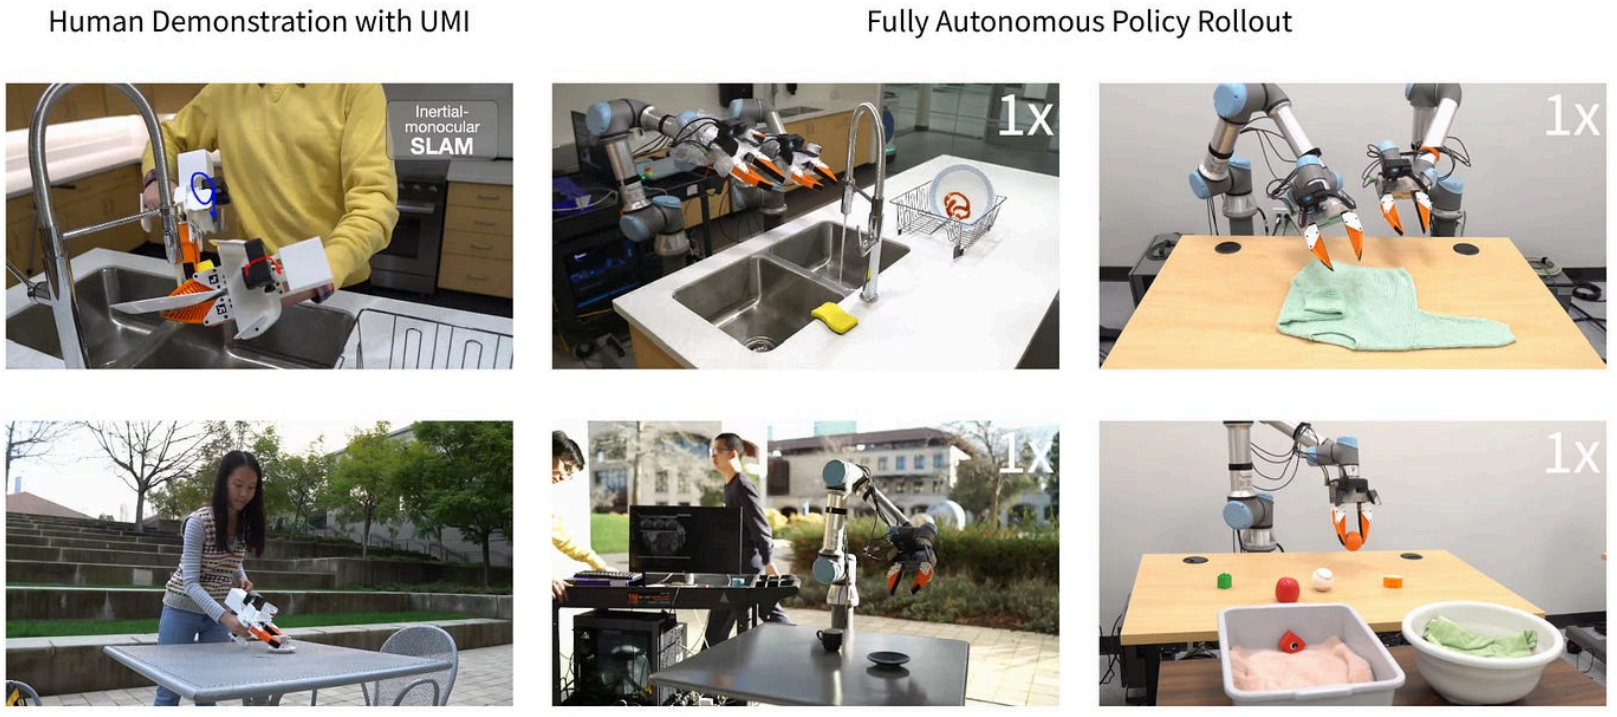


At the simplest end of the spectrum, Universal Manipulation Interface (UMI) Grippers – a system designed by Stanford, Columbia and Toyota – are a two-fingered hand with a mounted wrist camera. They're designed for task-specific training and help bridge the body embodiment problem in training data.

在最简单的一端，是由斯坦福大学、哥伦比亚大学和丰田联合设计的通用操作接口（UMI）夹爪——一种配备腕部摄像头的双指手。这类装置专为特定任务的训练而设计，有助于弥合训练数据中的具身差异问题。

With UMIs, a human demonstrator puts on their *own gripper* that mirrors the robot's actual hardware, which provides much more direct training lessons for the robot to replicate. Various versions exist, but most of the autonomy demonstrations are likely using this system or one similar to it.

使用 UMI 时，人类演示者会佩戴*自己的夹爪*，其构造与机器人的实际硬件相一致，从而为机器人提供直接得多的训练示范，便于其复现动作。此类装置存在多个版本，但大多数自主操作演示所采用的很可能正是这一系统或与之类似的系统。



Whether or not this shortcut is a bridge or a dead end is an open question. If you had two-pronged claws instead of hands, how many of your tasks daily could you perform? *(With a little bit of practice and adaptation, probably a lot.)*

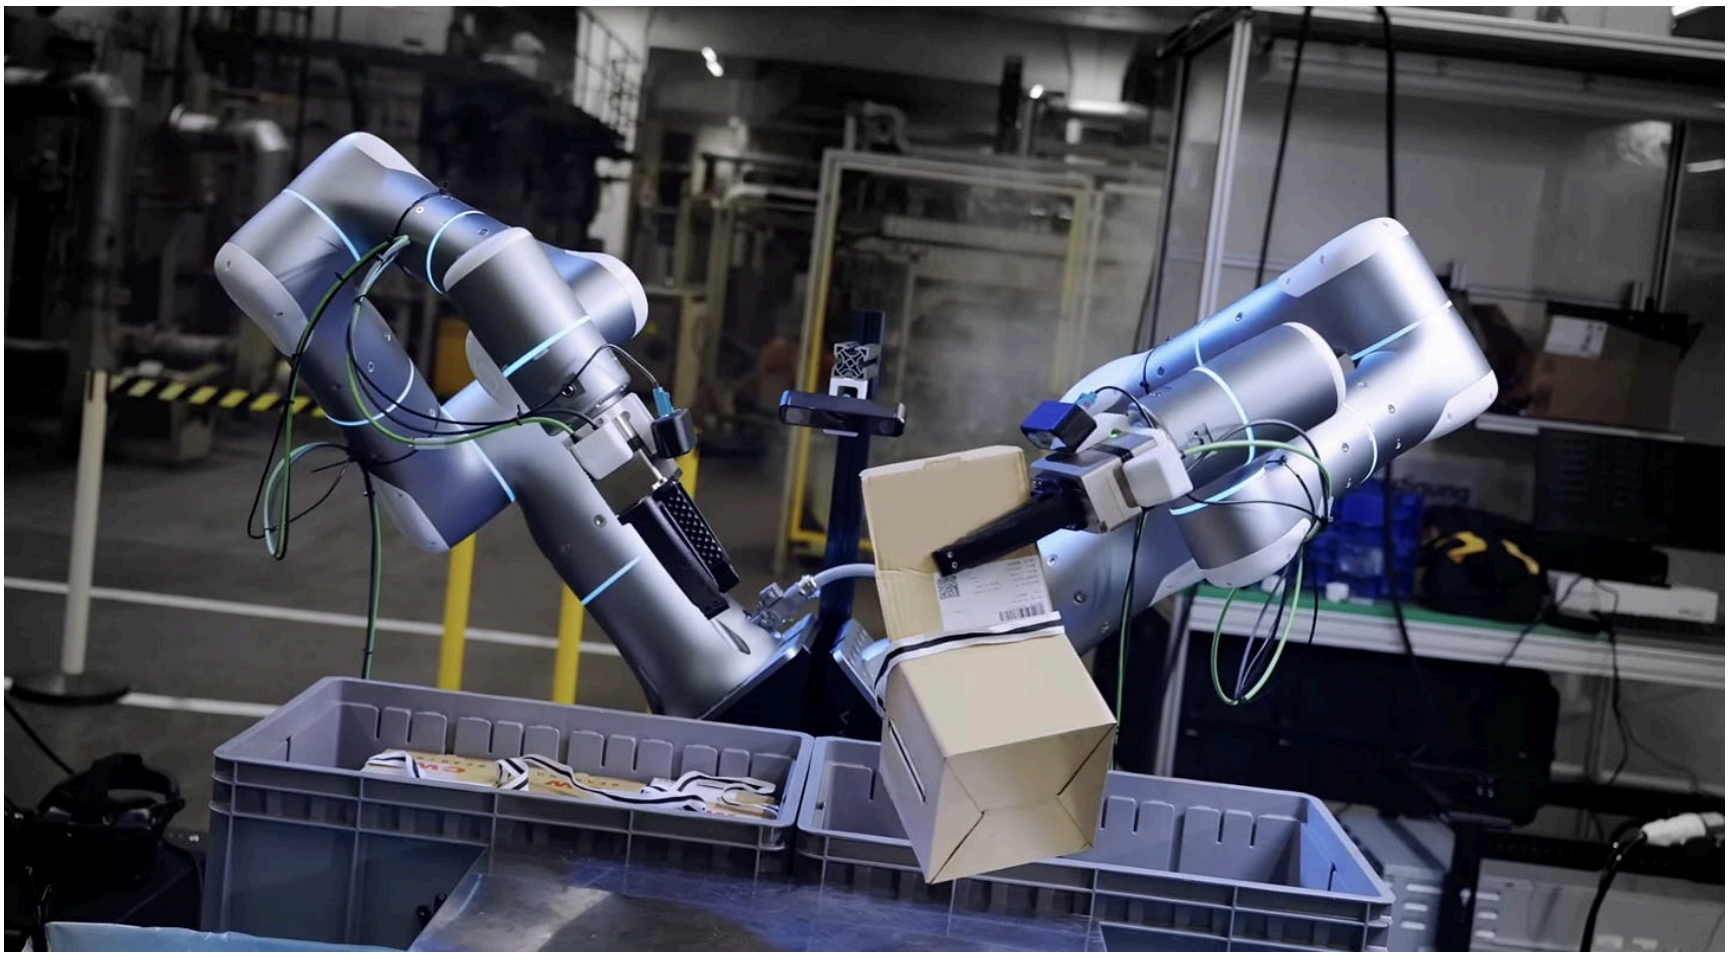
这条捷径究竟是通往未来的桥梁还是一条死路，目前尚无定论。如果你的双手被换成了两根叉状的爪子，你每天又能完成多少日常任务？*（只要稍加练习和适应，应该能完成不少。）*

And given the stage of autonomy we are at today, adding extra fingers may not be all that incrementally helpful.

而且以当下自主化技术所处的发展阶段来看，增加额外的手指未必能带来多少边际上的助益。

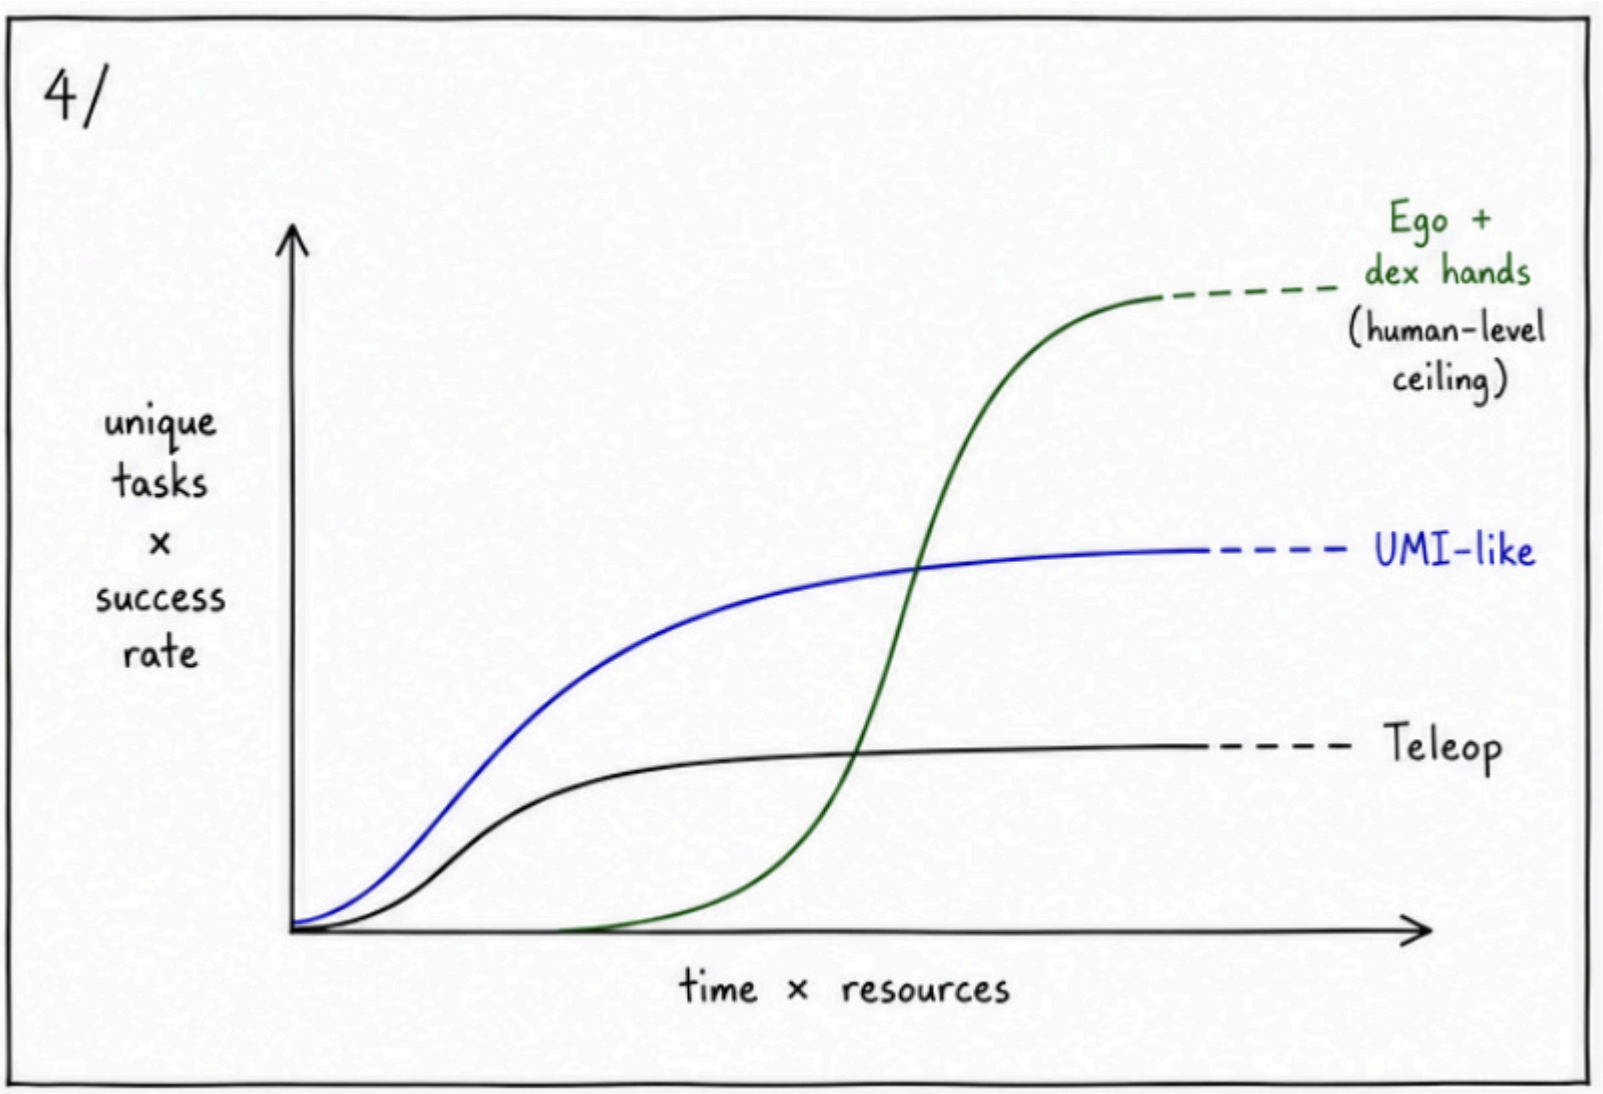
Many companies are forging forward with these simpler solutions. Private companies like Rhoda AI are tackling industrial automation tasks like decanting (unpacking boxes, as shown below) which it has demonstrated with two-prong grippers.

许多公司正在坚定推进这类更简单的方案。像 Rhoda AI 这样的私营公司正在攻克诸如倒箱（即拆箱取物，如下图所示）之类的工业自动化任务，该公司已用双爪夹爪演示了这类操作。



Still, if these tasks represent the ultimate TAM, then the robotics hype train is in a lot of trouble. As laid out by NVIDIA robotics researcher Danfei Xu, UMI and teleoperations likely lead to quicker results but with a lower ceiling.

不过，如果这些任务代表了最终的总可寻址市场（TAM），那么机器人领域的炒作热潮就麻烦大了。正如 NVIDIA 机器人研究员 Danfei Xu 所阐述的，UMI 和遥操作可能更快出成果，但天花板也更低。



Plenty of companies are shooting for the home run – the five-finger dexterous hand. But there is little agreement on how to do it, let alone who exactly is winning. If you ask the hand people, each will say why theirs is the best, and all others are doomed-to-fail. And if you look at the list prices below, these hands can cost more than the rest of the humanoid combined.

许多公司都在冲击那个"本垒打"式的终极目标——五指灵巧手。但在如何实现这一点上几乎没有共识，更不用说判断谁真正领先了。如果你去问各家做灵巧手的厂商，每家都会说自己的方案最好，而其他人注定失败。再看看下面的标价，这些灵巧手的价格可能比人形机器人其余部分加起来还要高。

One of the biggest hardware divides: tendons vs. direct drive.

硬件层面最大的路线分歧之一：腱驱动 vs. 直驱。

Tendon-based hands use motors in the forearm to control the fingers through rope-like fabrics – they spare space in the fingers and take inspiration from the human body. The problem is that tendons wear and stretch fairly quickly especially under strain and the setup requires integration down to the “elbow”.

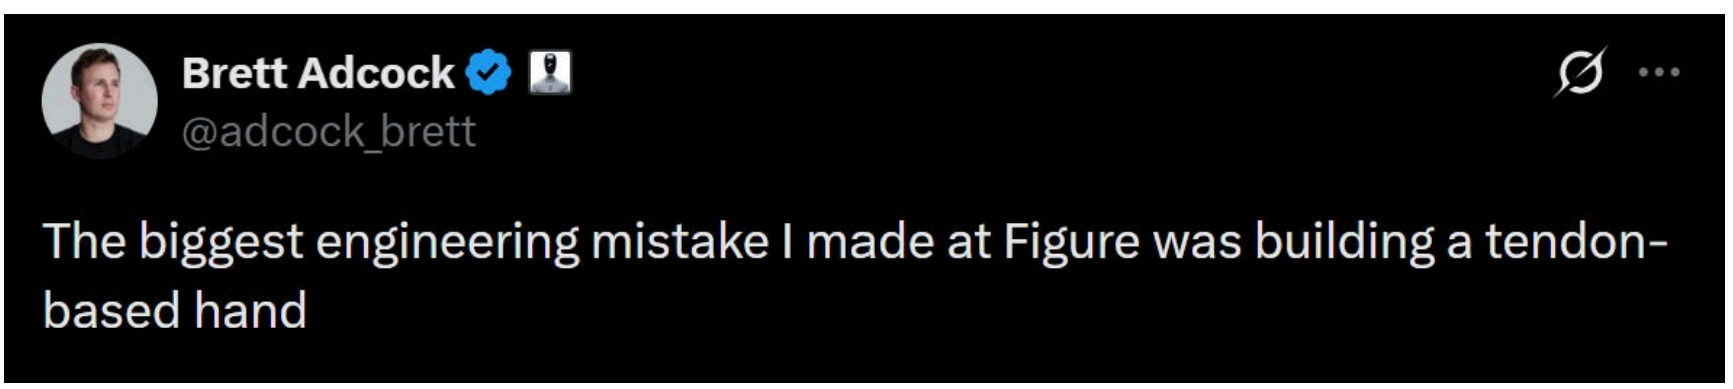
腱驱动方案将电机置于前臂，通过绳状织物控制手指——这种设计能节省手指内的空间，灵感来自人体结构。问题在于腱绳磨损和拉伸相当快，尤其是在承受负载时，而且这套方案需要把集成范围一直延伸到"肘部"。

Direct drive puts the actuators throughout the fingers, which makes simpler linear kinematics for modeling and ensures the hand is wrist-mountable and “modular”, but it adds weight, size, and overheating concerns.

直驱方案则把执行器布置在整只手指之中，这使得运动学建模更为简单、更偏线性，并保证手部可直接安装在腕部、具备"模块化"特性，但代价是重量和体积增加，以及过热隐患。

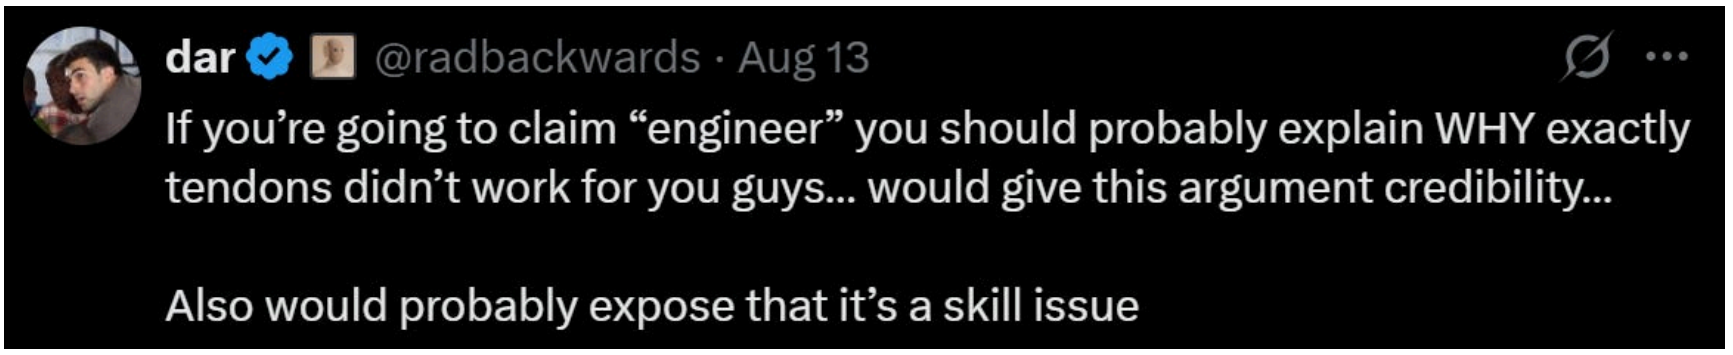
The *Tendon-istas* vs *Drive-cels* has almost become an ideological turfwar, with Figure’s Brett Adcock calling a tendon-based hand the “biggest engineering mistake I made” – Figure has since returned to a direct drive format.

腱绳派与***直驱派***之争几乎已演变成一场意识形态阵地战，Figure 的 Brett Adcock 就称腱驱动手是"我犯下的最大的工程错误"——此后 Figure 已改回直驱方案。



The lead designer at 1X spat back, calling it a “skill issue”.

1X 的首席设计师随即回怼，称这不过是“技术不行”罢了。



The jury is out for now, but it's worth noting that Figure's livestream shows uninterrupted work and solid dexterity and Wuji's Hand 2 (which is also a direct drive) has gained traction as one of the most popular third party hands on the market.

目前尚无定论，但值得注意的是，Figure 的直播展示了连续不间断的作业和扎实的手部灵巧性，而 Wuji 的 Hand 2（同样采用直驱方案）也已获得市场认可，跻身市面上最受欢迎的第三方灵巧手之列。

Company / Model	Hand	DoF	Drive	Price
Tesla Optimus Gen 3 (V3)		22 hand + 3 wrist	Tendon-driven, actuators in forearm	Not sold separately
Sanctuary AI Phoenix hand		21 (separate 17-DoF industrial hand in development)	Hydraulic (miniaturized valves)	Not for sale yet
Sharpa Wave		22 active	Not disclosed in detail	Quote-only
Figure 02/03 hand		16 (Figure 02)	Electric, actuators and sensors in each finger; tendon per third-party summaries	Not sold separately
Shadow Robot Dexterous Hand		20 actuated / 24 movements	Tendon-driven	~$100–150k
Wonik Allegro V5 Sense		16 (four fingers)	Direct drive (current-controlled joints)	~$15–35k (sources conflict)
Unitree Dex5-1/1P		20 (16 active, 4 passive)	Direct drive, backdrivable	From $25k US retail
Wuji Hand 2		20 active (4 per finger, all independent)	Direct drive, back-drivable rotary joints	From $18,500
DexRobot DexHand 021		19 (12 active, 7 passive)	Cable/tendon-driven	~$10–15k (China mid-tier)
Inspire Robots RH56 family		6 active / 12 joints (linkage-coupled)	Micro linear servos	~$8,000–9,500

CITRINI RESEARCH

Finally there is *touch.* Ever try to unzip a jacket with frozen fingers?

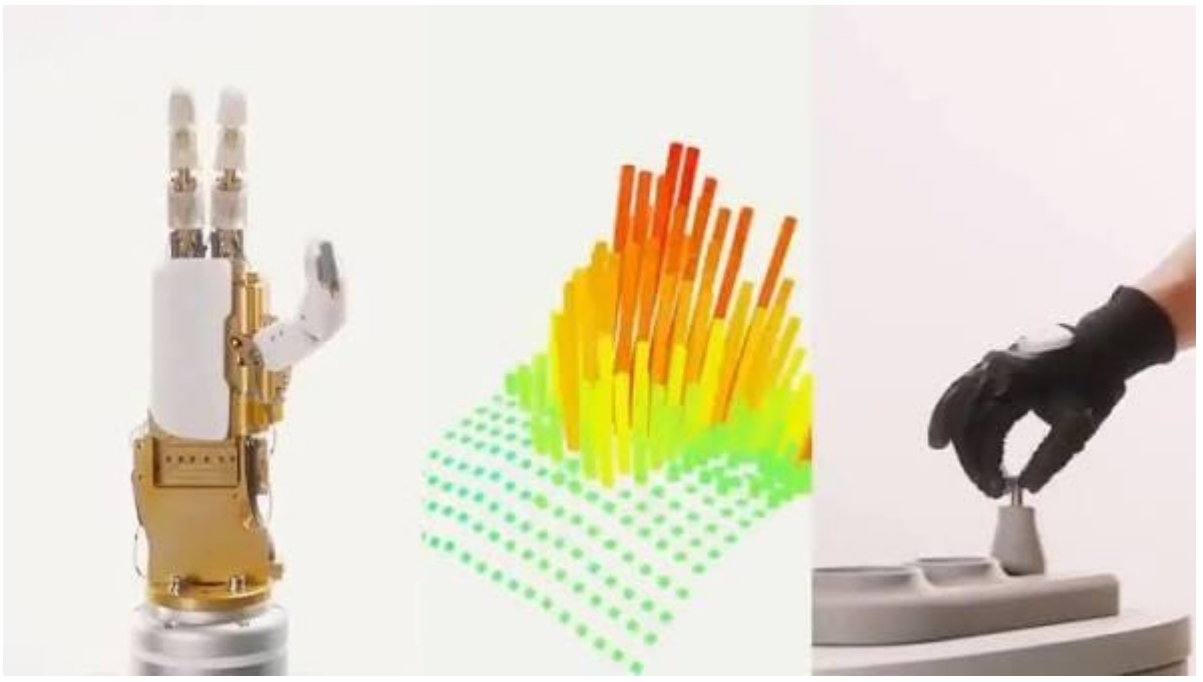
最后是***触觉***。你有没有试过用冻僵的手指去拉夹克的拉链？

The human fingertips have some of the densest nerve endings in the body with ~240 mechanoreceptors per square centimeter. Particularly for fine manipulation, this tactile sensation often guides us better than our own eyes.

人类指尖拥有人体中分布最密集的神经末梢之一，每平方厘米约含 240 个机械感受器。尤其是在精细操作中，这种触觉感知对我们的引导往往比自己的眼睛还要可靠。

Companies like privately-held Tacta Systems are developing ultra dense sensor arrays or ***skins*** that can be applied to hands in order to add this level of perception into robotic training and deployment.

像非上市公司 Tacta Systems 这样的企业正在开发可贴附于手部的超高密度传感器阵列，即***皮肤***，以便将这种级别的感知能力引入机器人训练与部署之中。



Along with a highly refined hand – like its own hydraulically actuated model – it's possible these systems could be simply mounted on cobots and used in line-manufacturing of electronics currently being done by humans.

再加上高度精巧的手部——比如其自有的液压驱动型号——这些系统有可能直接安装在协作机器人上，用于目前仍由人工完成的电子产品产线制造。

But even the best body is not worth much without a brain.

但再好的身体，没有大脑也意义不大。

2) The Brain: Grading the Class

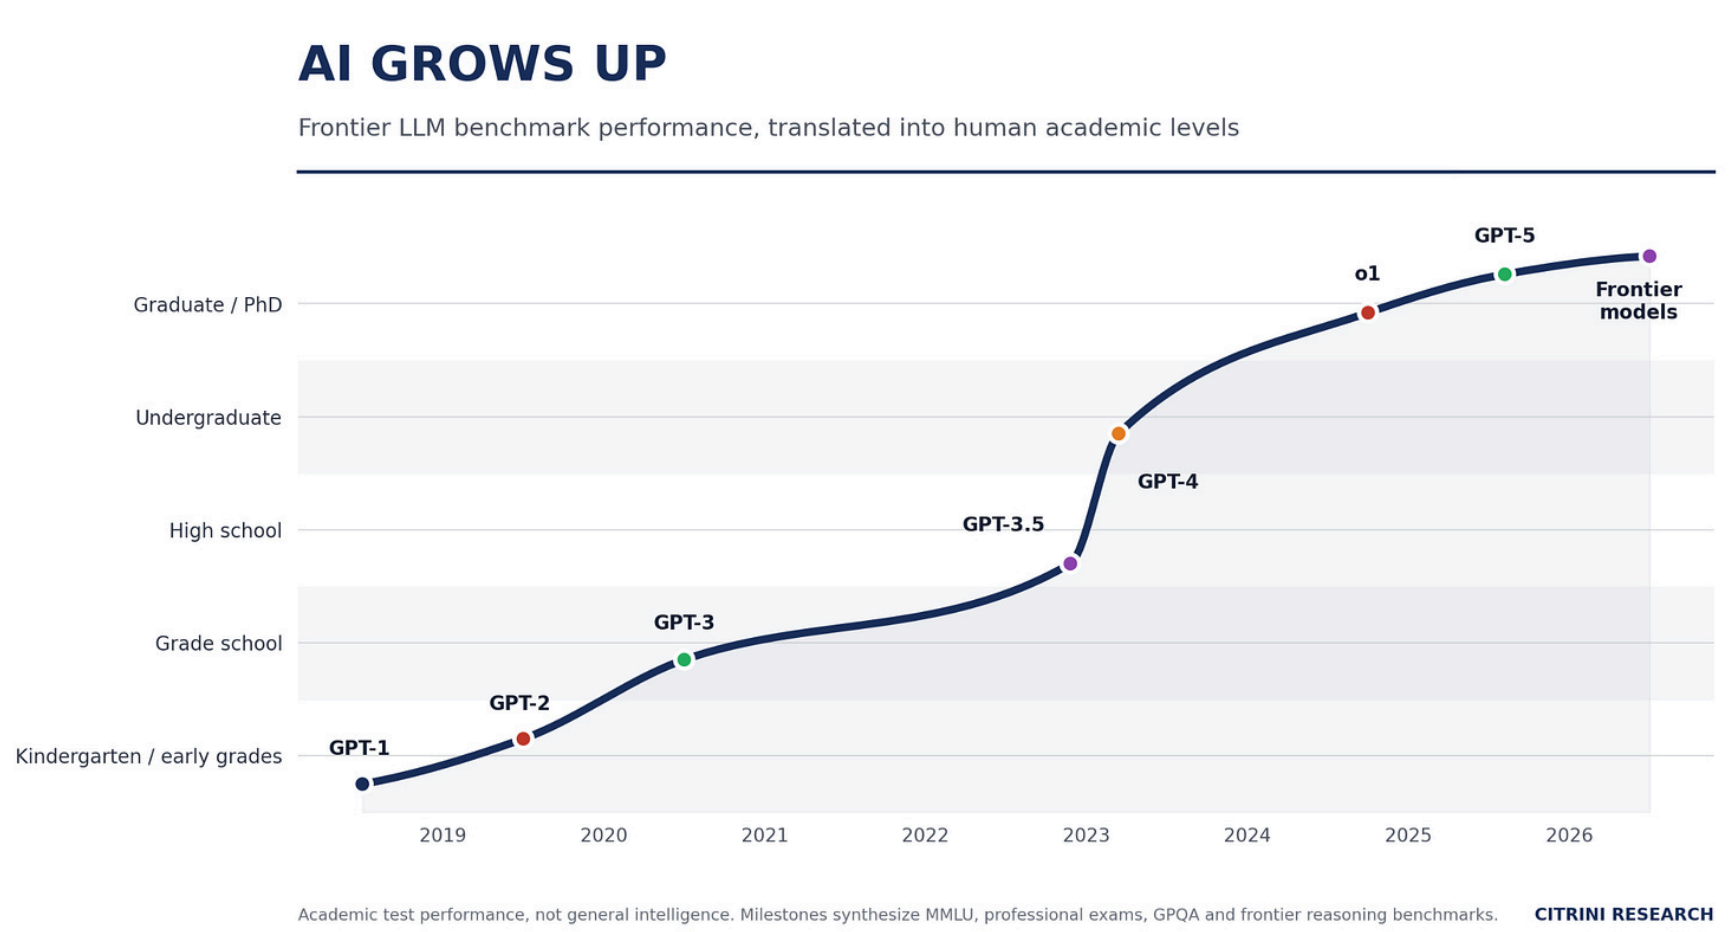
2）大脑：给全班打分

Building the brain is both the greatest challenge and most exciting thing happening in robotics. But unlike the TradAI (you know, that stuff on your computer), Physical AI lives in the real world. The real world is far harder to understand and tokenize. Written language semantics carry inherent meaning and intent in a way that an apple sitting on a table does not. A robot cannot think for two minutes about how to respond to a trip or a push. Gravity has no latency.

打造大脑既是机器人领域当前最大的挑战，也是最令人兴奋的事情。但与 TradAI（就是你电脑上跑的那些东西）不同，Physical AI 生活在真实世界里。真实世界要难理解得多，也更难 Token 化。书面语言的语义天然携带含义与意图，而摆在桌上的一个苹果并不具备这种特性。机器人不可能花两分钟去思考该如何应对一次绊倒或推搡。重力没有延迟。

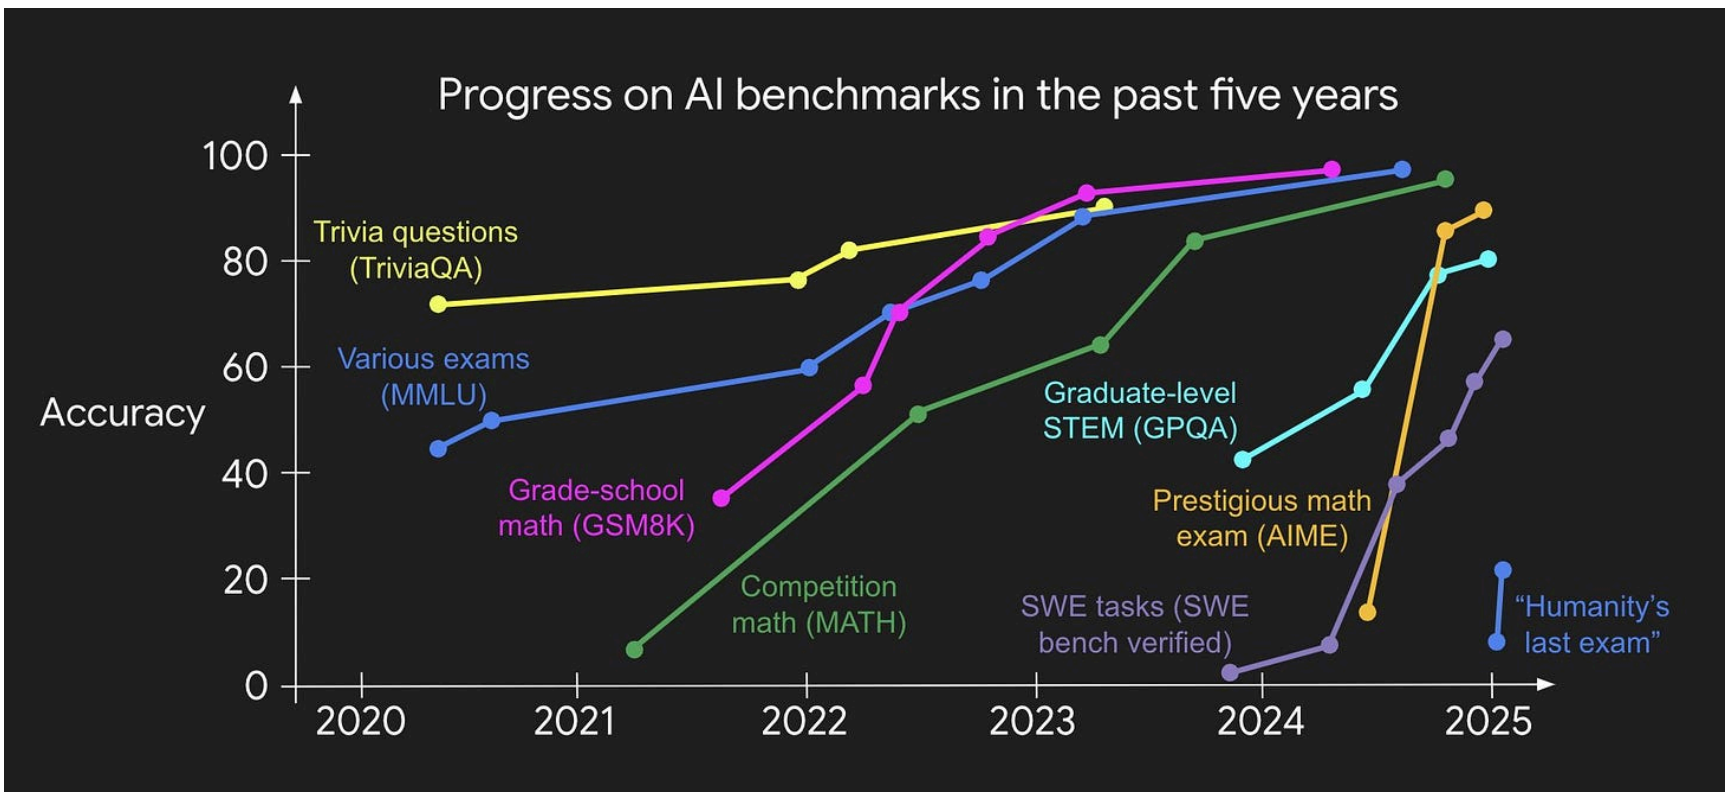
In the LLM space, we've watched intelligence *graduate* from kindergarten to PhD levels over the course of eight years.

在 LLM 领域，八年间我们亲眼见证了智能从幼儿园水平*毕业*到博士水平。



It's not a perfect benchmark – but at least there are quantitative benchmarks supporting it. Earlier GPTs were much better at trivia than a grade schooler, but we understand that fact-retrieval is a very different skill than *reasoning* through math tests. But they got better [fast](#).

这并不是一个完美的基准——但至少有定量基准在为其提供支撑。早期的 GPT 在琐碎知识问答上远胜小学生，但我们清楚，事实检索与在数学考试中进行*推理*是截然不同的两种技能。不过它们进步得非常**[快](#)**。



Notably, the learning curve was a very different shape from our own – the path towards post-PhD AI was frontloaded by a foundation built for well over a decade before reaching escape velocity. Imagine if you were in grade school for ten years, then speed-ran to solving unsolved mathematical proofs two years later (with the help of gigawatts of compute).

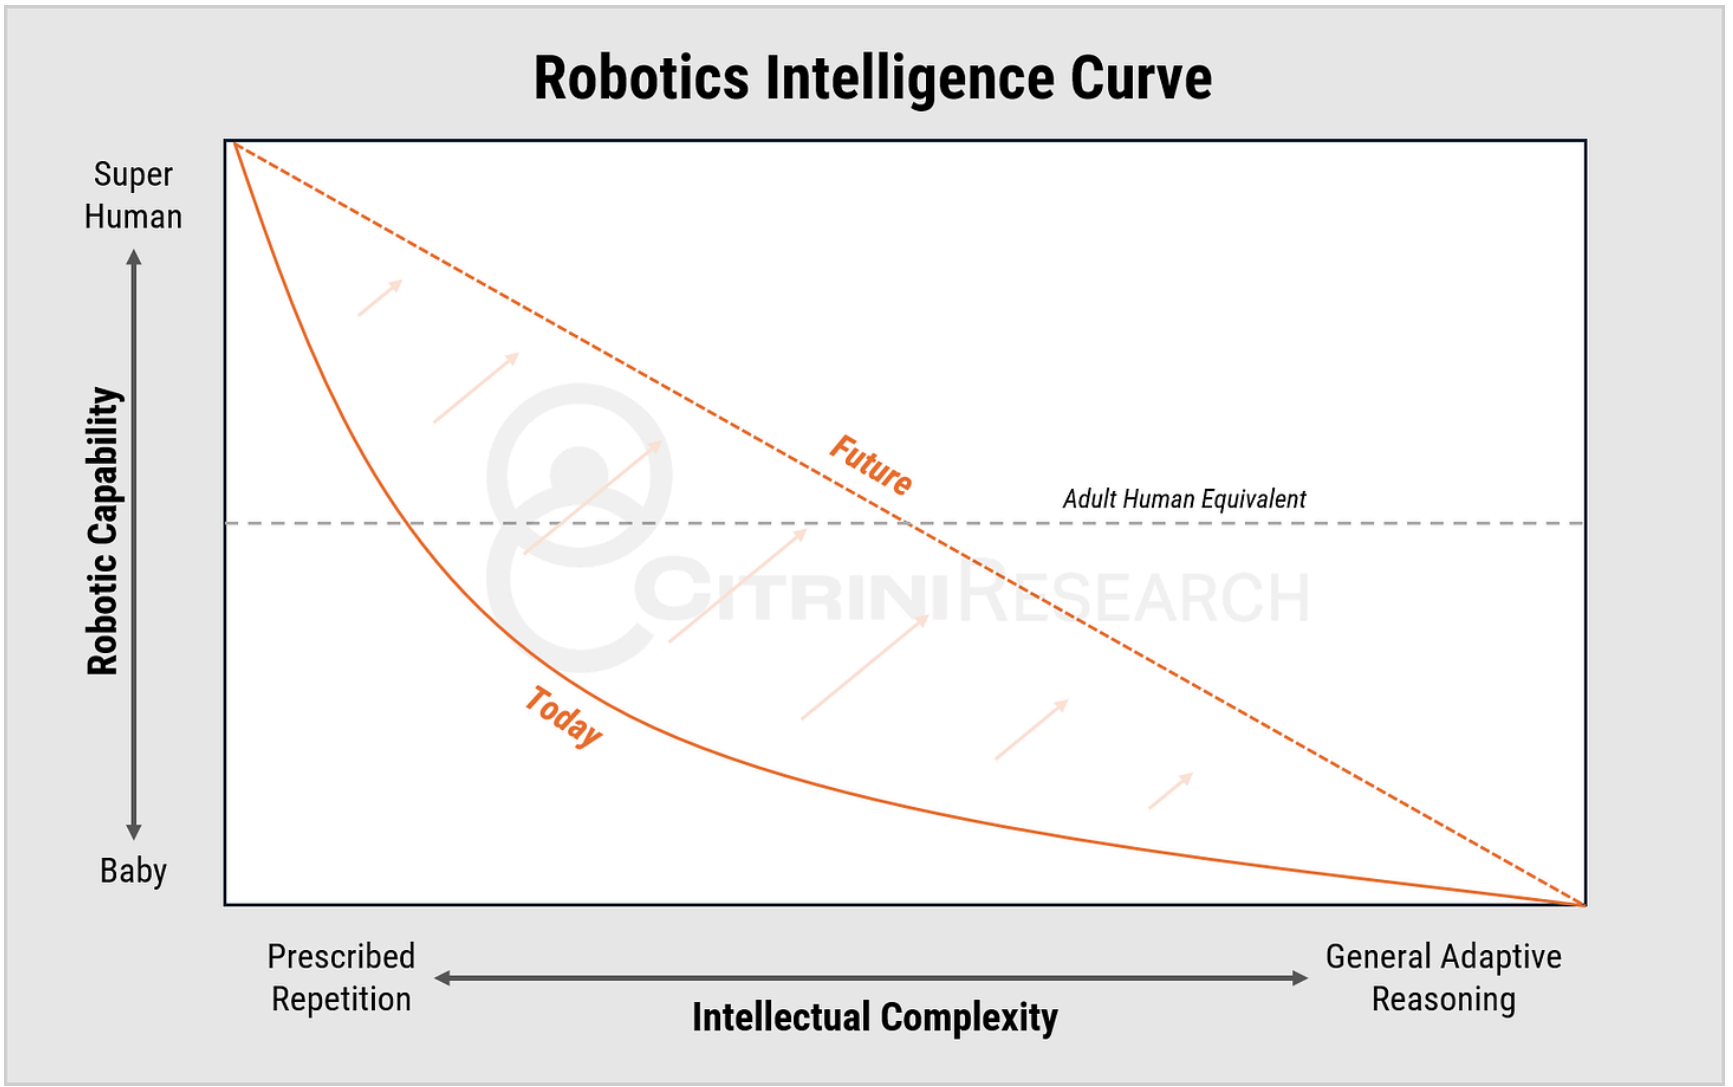
值得注意的是，其学习曲线的形状与我们自身截然不同——通向"博士级以上AI"的道路是前重型的：在达到逃逸速度之前，需要先打下持续十余年的基础。想象一下：一个人在小学待了十年，然后在两年内极速冲刺，开始解决尚未破解的数学证明（期间借助了吉瓦级的算力）。

Robotics' hope rests not only on the absolute intellectual achievement seen in AI, but even more importantly *acceleration in progress* seen in recent years.

机器人技术的希望，不仅系于AI所展现的绝对智力成就，更重要的是近年来所呈现的*进展加速*。

There is nothing today that comes close to general autonomy in robotics. ***But, that doesn't mean that autonomy isn't improving.*** The "autonomy curve" is beginning to shift, and robots are getting better at increasingly complex tasks. (For prescribed, repetitive tasks, robots have been super-human for decades).

如今还没有任何东西能接近机器人领域的通用自主能力。***但这并不意味着自主能力没有在进步。***"自主能力曲线"正在开始转变，机器人在日益复杂的任务上表现得越来越好。（对于预设的、重复性的任务，机器人的表现几十年来早已超越人类。）



As the curve shifts to the right, a greater number of tasks are able to be completed at roughly “human equivalent” capability. It’s at these levels where true economic viability and commercial demand likely begins. Yet again, it’s a spectrum, not a binary.

随着曲线向右移动，能够以大致“人类同等”水平完成的任务数量随之增加。真正的经济可行性与商业需求很可能正是从这些水平开始出现的。这依然是一个连续光谱，而非非此即彼的二元判断。

When Unitree’s founder says ***“humanoid robots are not yet ready for large-scale factory deployment”,*** he is right. Today, they lack the intelligence, dexterity, sensitivity, durability and safety needed. The factory “deployments” that have been announced by humanoid OEMs to date are effectively pilot projects – in the field demos for proof of concept and data capture.

当宇树科技创始人表示***“人形机器人尚未准备好进行大规模工厂部署”***时，他是对的。如今，人形机器人仍缺乏所需的智能、灵巧性、敏感性、耐用性与安全性。迄今人形机器人 OEM 厂商宣布的各类工厂“部署”，实际上都是试点项目——即现场演示，用于概念验证与数据采集。

TECH

Unitree Shares Decline as Founder Flags Limits of Humanoid Robots

By Du Zhihang

Gift this article

Published: Aug. 20, 2026 3:22 p.m. GMT+8

But autonomous robots ***are becoming capable in task-specific settings***, and the number of use cases is growing. It seems only logical that the path towards a fully generalizable autonomy will start with baby steps.

但自主机器人***在特定任务场景中正变得有能力***，应用案例的数量也在增长。通向完全可泛化自主性的路径，看起来自然会从蹒跚学步般的渐进开始。

But before we answer ***“wen autonomy?”*** we need to understand the architecture.

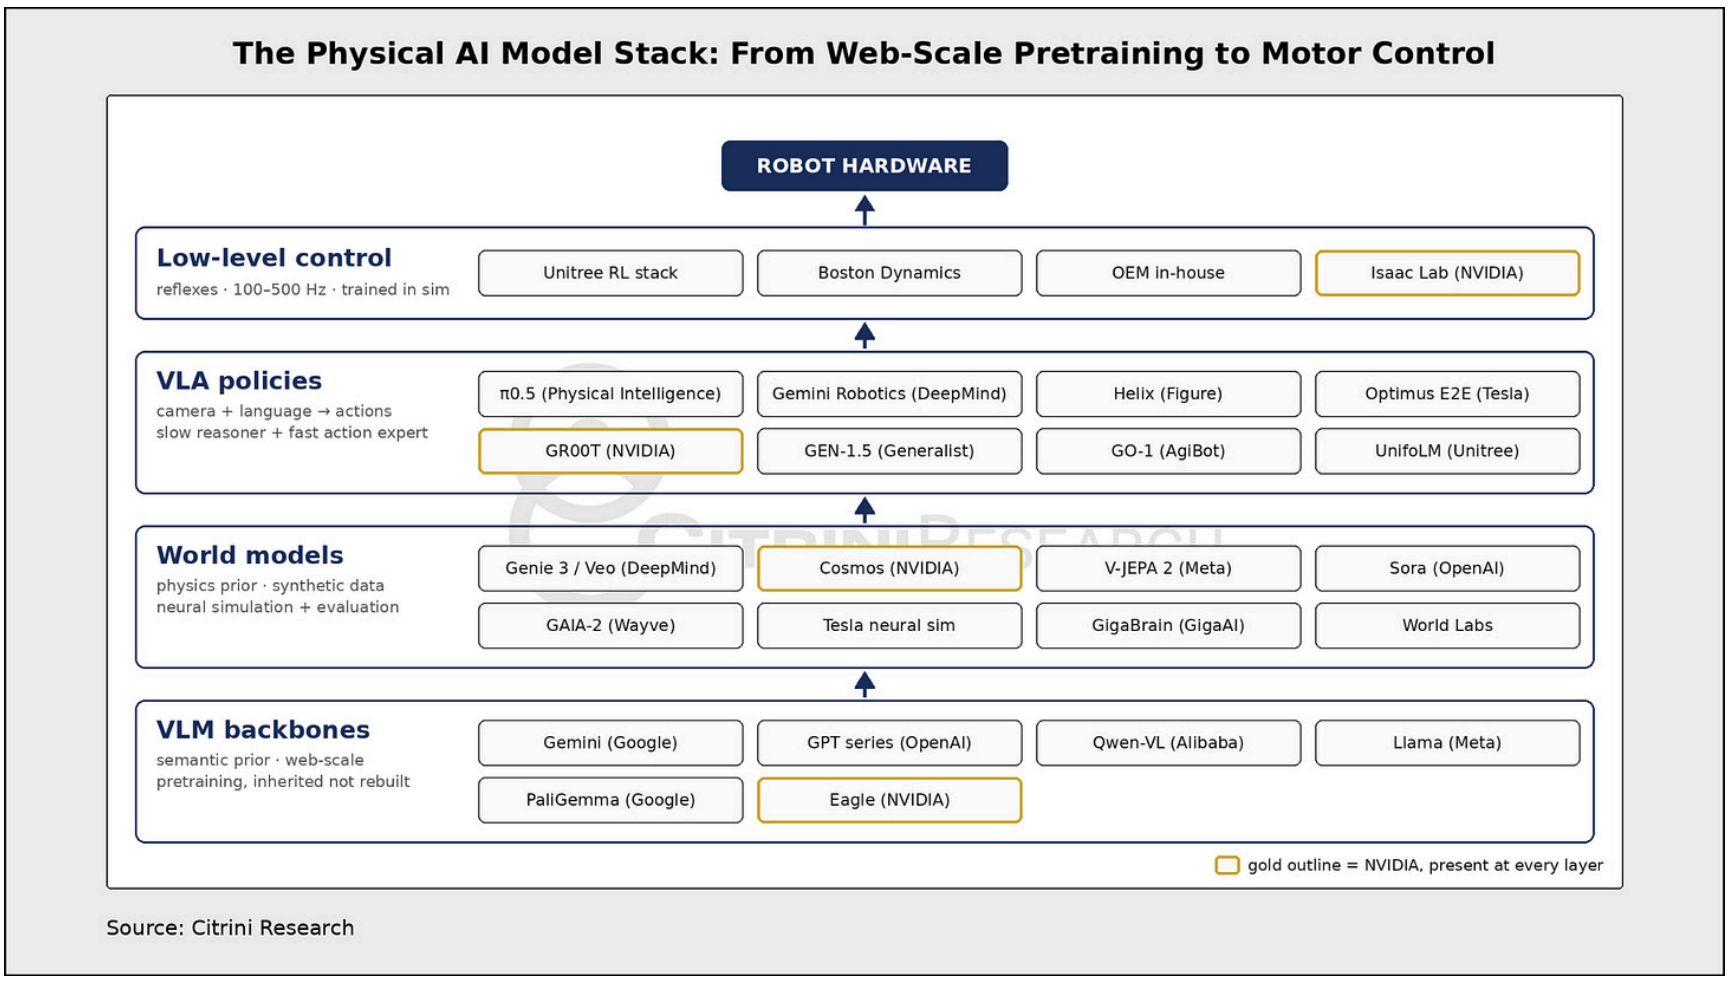
但在回答***“自主何时到来？”***这个问题之前，我们需要先理解其架构。

The Nervous System:

神经系统：

Robotics brains rely on a multi-layered interconnected nervous system with its own set of three-letter acronyms. (Note, decoding the alphabet soup depends a bit on who you ask).

机器人的“大脑”依赖一套多层互联的神经系统，其中自带一批三字母缩写词。（注意，如何解读这堆缩写大杂烩，一定程度上取决于你问的是谁。）



Starting at the foundation: Visual Language Models (VLMs) are all about understanding objects and contexts, in the same way that humans can subconsciously and instantaneously understand their surroundings by opening their eyes. They also encode physical objects to language – that brown thing with four legs is “TABLE” and the red circle on top is “APPLE”. This “translation” allows the robot to describe what it’s seeing with language, while also allowing a human to provide written or spoken instruction.

先从基础层说起：**视觉语言模型（VLM）**的核心是理解物体与环境背景，就像人类睁开眼睛便能下意识地、即时地理解周围环境一样。它还会把物理实体编码为语言——那个带四条腿的棕色物体是“桌子”，摆在它上面的红色圆形物体是“苹果”。这种“翻译”让机器人能够用语言描述所见之物，也让人类可以通过文字或语音下达指令。

VLMs are already integral in AI products – if you paste a picture into Claude, it uses a VLM to describe the picture in specific and highly accurate language.

VLM 已经是 AI 产品不可或缺的组成部分——你把一张图片粘贴到 Claude 里，它就会调用 VLM 用具体且高度精确的语言描述这张图片。

Multiple closed- and open-source VLMs have already been trained on internet-level corpora, making this layer of the robotics stack increasingly accessible. Even “full-stack” robotics projects would not try to replicate them from scratch. Figure, for example, in the past noted that its Helix system relies on an “open source” VLM with subsequent fine tuning (speculators guess it’s likely Meta’s Llama or Alibaba’s Qwen).

多个闭源与开源 VLM 已经在互联网规模的语料上完成训练，使机器人技术栈中的这一层越来越容易获取。即便是“全栈”机器人项目，也不会尝试从零复现它们。例如 Figure 过去曾表示，其 Helix 系统依赖一个“开源”VLM 并在其基础上做微调（外界猜测很可能是 Meta 的 Llama 或阿里巴巴的 Qwen）。

(Tesla might be the only exception given its extensive in-house work developing autonomous driving systems.)

（特斯拉可能是唯一的例外，因为它在自动驾驶系统的研发上有大量自研积累。）

But VLMs provide just passive interpretation. They can’t tell you to PICK UP the APPLE on the TABLE, nor can they describe the physics you’d need to do it. That’s where “world models” come in.

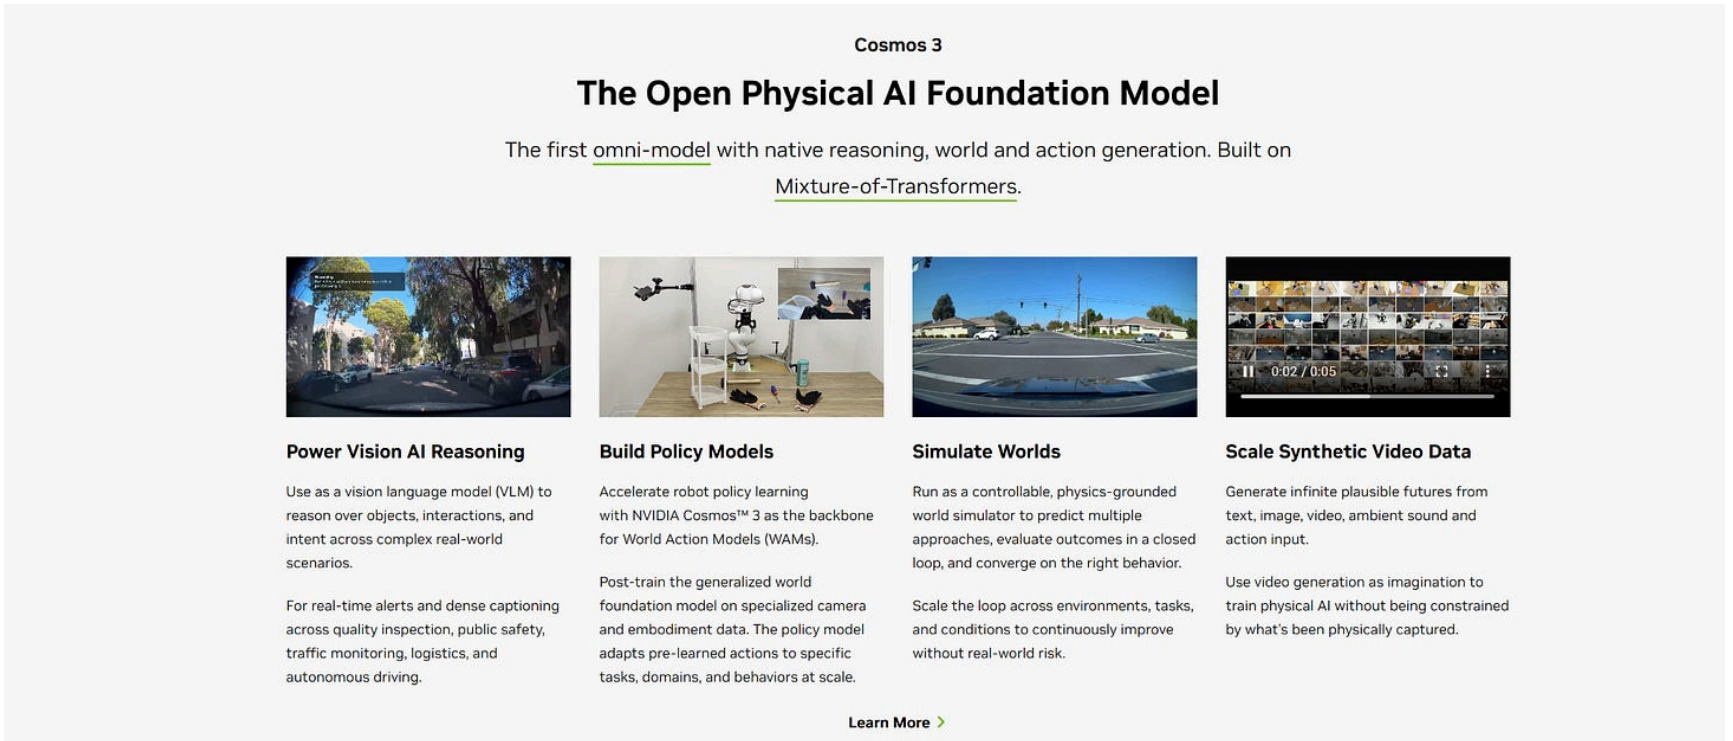
但 VLM 只提供被动的解读。它们无法发出“拿起桌上的苹果”这样的指令，也无法描述完成这个动作所需的物理过程。这正是“世界模型”发挥作用的地方。

World Models are also familiar to most in the AI space; Veo from Google, Sora from OpenAI, Cosmos from NVIDIA, etc. These are basically physics engines that predict how the world works. What happens if you move one step forward? What happens if you squeeze a water bottle (what if it’s plastic vs. metal, cap on vs. off, etc)?

世界模型对 AI 领域的大多数人来说也并不陌生：谷歌的 Veo、OpenAI 的 Sora、英伟达的 Cosmos 等都属于这一类。它们本质上就是预测世界运行规律的物理引擎。向前迈一步会发生什么？挤压一个水瓶会发生什么（瓶身是塑料还是金属、瓶盖是拧紧还是打开，结果会有何不同）？

The world-model layer is less commoditized and more strategically contested than the VLM layer. NVIDIA’s Cosmos is open-sourced, with strategic intent to spur innovation across an industry that largely relies on its chips. Google’s Genie3/Veo is licensable while OpenAI’s Sora is entirely internal.

与 VLM 层相比，世界模型层的商品化程度更低，战略争夺也更为激烈。英伟达的 Cosmos 采用开源模式，其战略意图在于推动整个行业的创新——而该行业在很大程度上依赖英伟达的芯片。谷歌的 Genie3/Veo 可对外授权使用，而 OpenAI 的 Sora 则完全封闭于内部。



Further, understanding the physics of motion is ***inherently dependent on hardware***, and therefore less interoperable between players. Optimus isn't a G1. While these models can ***simulate*** the physics of any given hardware setup, reality is messy. A world model that is fed with actual real world data collection is better positioned to navigate the "sim-to-real" gap.

此外，对运动物理规律的理解***本质上依赖于硬件***，因此在各家厂商之间互操作性较低。Optimus 不是 G1。虽然这些模型可以***模拟***任意给定硬件配置下的物理特性，但现实世界是杂乱无章的。一个以真实世界数据采集为基础进行训练的世界模型，更有条件跨越"仿真到现实"（sim-to-real）的鸿沟。

The next layer is **Visual Language Action models (VLAs).** While VLMs and world models provide the "subconscious" understanding of the world, VLAs are essentially the high-level conscious reasoning required to act within it. They tell the robot how to understand instructions and intent, allowing it to reason through the actions necessary to complete the task – relying on foundational context (VLMs) and physics (world models).

再往上一层是**视觉语言动作模型（VLA）**。VLM 和世界模型提供对世界的"潜意识"层面理解，而 VLA 本质上是在这个世界中采取行动所需的高层次有意识推理。它们让机器人理解指令与意图，从而推理出完成任务所需的动作——这一切都建立在基础语境（VLM）和物理规律（世界模型）之上。

The model releases I've mentioned above from Skild and Generalist exist at this level and it's where most AI labs are focused on iteration and advancement.

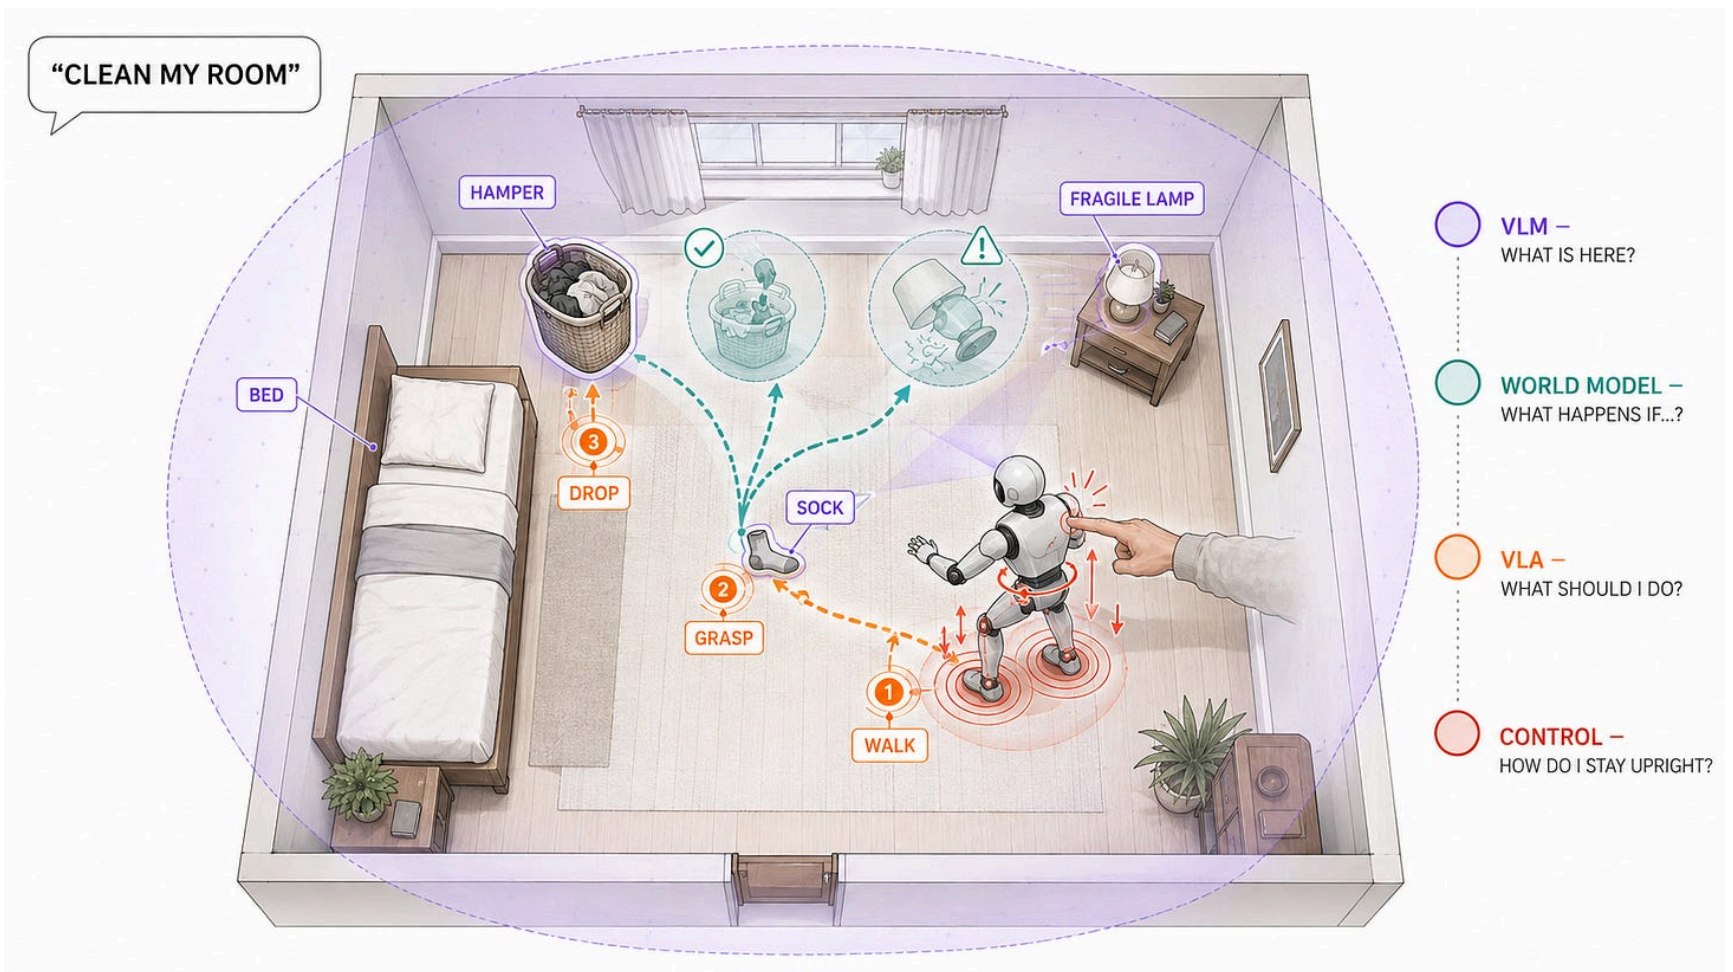
我上文提到的 Skild 和 Generalist 发布的模型即处于这一层，这也是目前大多数 AI 实验室集中力量迭代和推进的方向。

On top of all of this sits the low-level control systems – basically the quick-twitch muscles that ensure balance through hyper-fast flexes across the bot's muscles while incorporating live and unexpected physical inputs.

在这一切之上的是底层控制系统——它们相当于"快肌纤维"，通过遍布机器人全身肌肉的超高速收缩来维持平衡，同时实时处理动态且不可预期的物理输入。

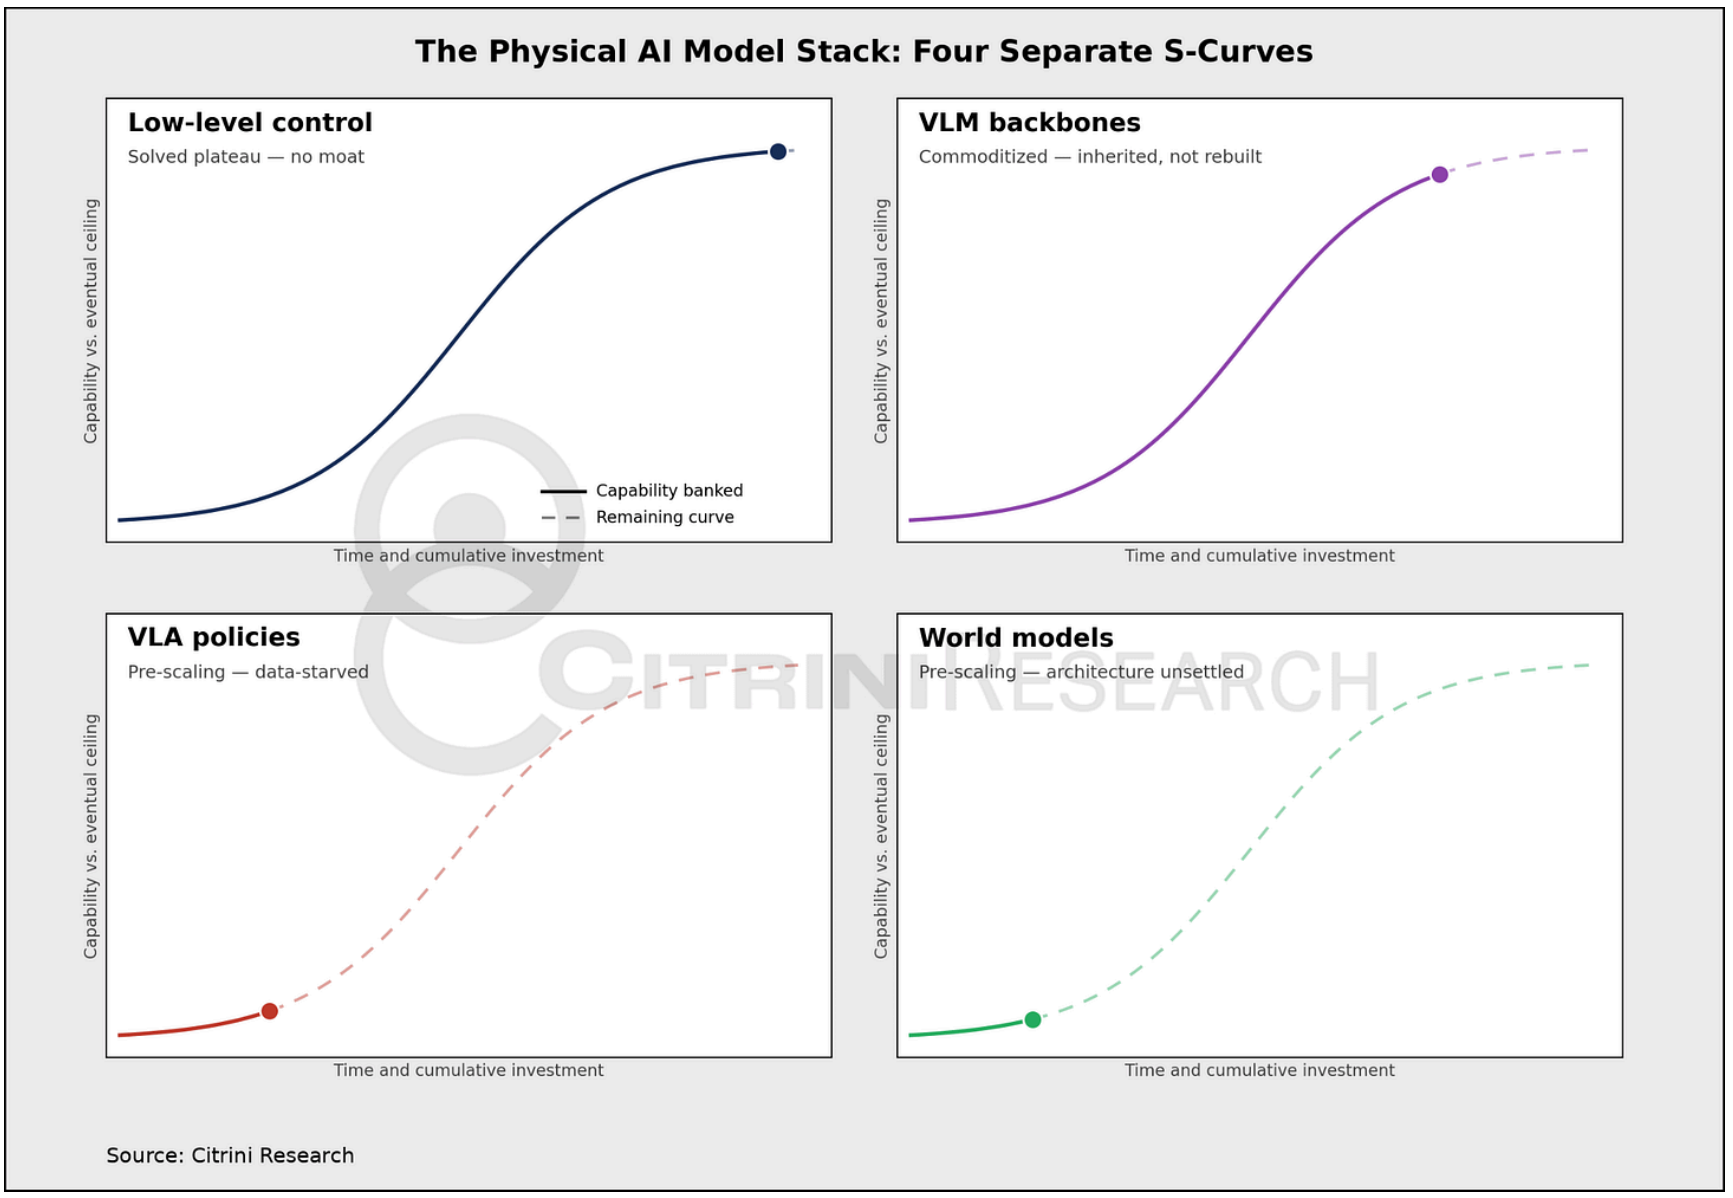
Let's take an example concretely: "***ROBOT-BUTLER, CLEAN MY ROOM!"***

我们用一个具体例子来说明："***机器人管家，打扫我的房间！***"



So which of these processes are solved, and which are holding us back?

那么，这些流程中哪些已经得到解决，哪些仍在拖累我们？



At the risk of oversimplification, low-level control and VLM backbones are mostly solved or at least highly commoditized. The major challenge is refining World Models and VLA policy – and the two are deeply entwined, with a potential solution arriving in their fusion: World Action Models (WAMs).

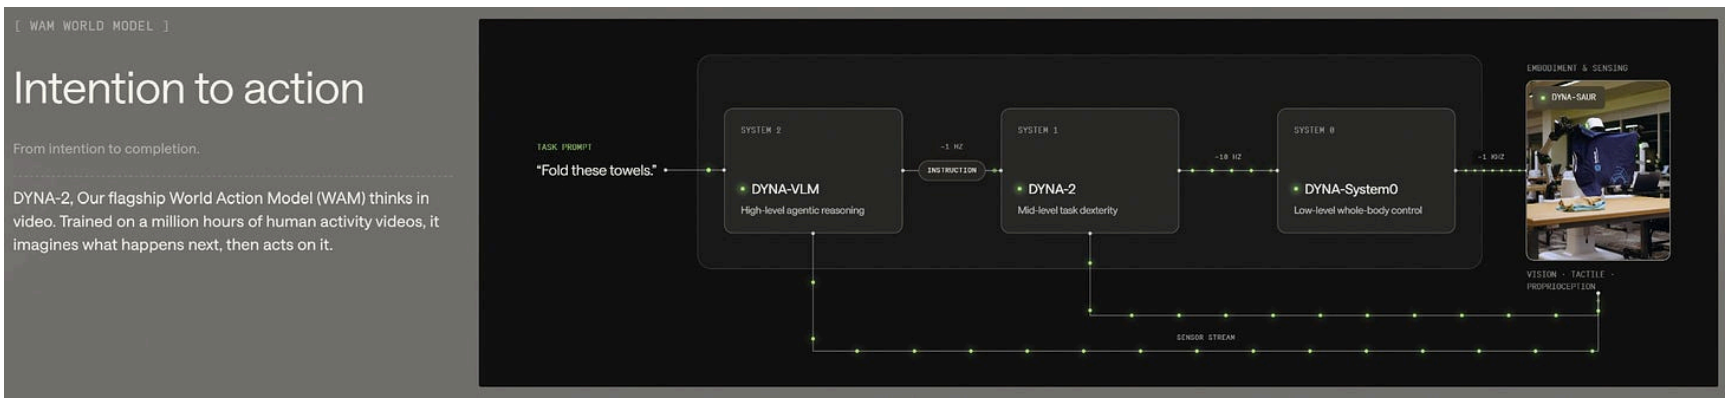
尽管有过度简化之嫌，但底层控制与 VLM 骨干网络大体上已属已解决问题，或至少已高度同质化。当前的主要挑战在于打磨世界模型（World Model）与 VLA 策略——而这两者深度交织，一个潜在的解决路径正是二者的融合：世界动作模型（WAM）。

World Action Models, if you can't guess, combine the real world physics of the World model with the policy and direction of a VLA. NVDA recently released its Hydra-0 combining world context with instruction.

世界动作模型，顾名思义，是将世界模型所刻画的现实世界物理规律与 VLA 的策略和方向结合起来。英伟达（NVDA）最近发布了 Hydra-0，将世界上下文与指令相结合。

Dyna Robotics' DYNA-2 model which the company claims to show in-context learning is a World Action model as well.

Dyna Robotics 的 DYNA-2 模型同样属于世界动作模型，该公司声称其展现出上下文学习能力。

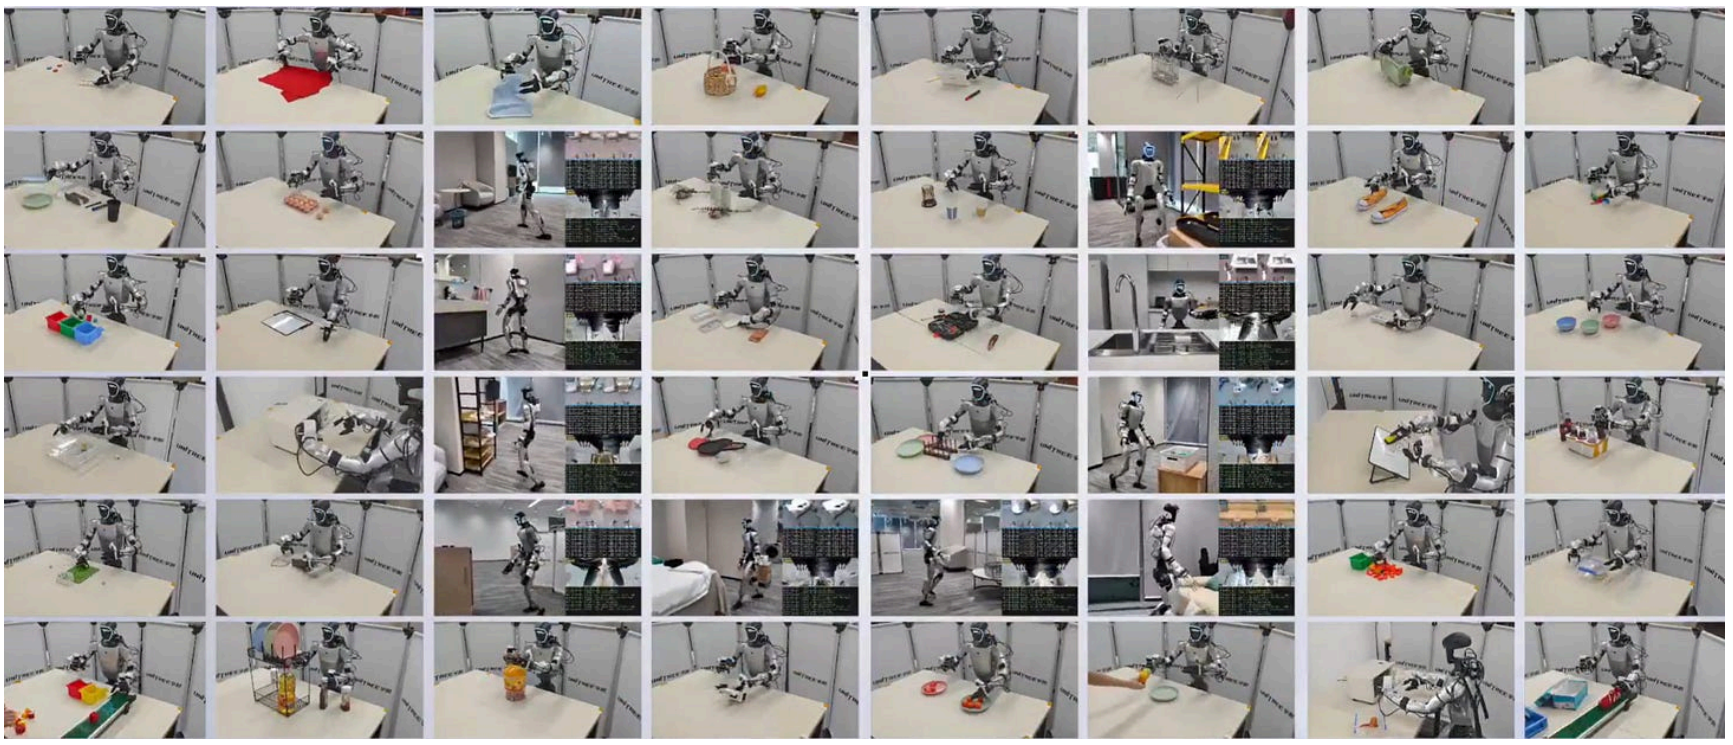


To emphasize, new announcements are arriving at such a rapid pace that it seems like every day that passes demands a new rewriting of this section.

需要强调的是，新的公告正以极快的速度接连发布，几乎每过一天，这一节内容就需要重新改写一次。

Unitree last week released its own open source "Unitree General-Purpose Humanoid Foundation Model" or UnifoLM-WLA-1.0 which it describes as a new state-of-the-art model that "coordinates desktop and whole-body mobile manipulation, supporting cross-task and cross-end-effector generalization, driven by one model, whole-body coordination".

宇树上周发布了自己的开源"宇树通用人形机器人基础模型"，即 UnifoLM-WLA-1.0，并称其为达到业界领先水平的全新模型，"可协调桌面操作与全身移动操作，支持跨任务、跨末端执行器的泛化，由单一模型驱动，实现全身协同"。



According to its self-reported benchmarking, this model would top a broad list of open- and closed-source models (though it doesn't seem to include the leading US startups). For a company that has historically dominated in hardware, but mostly ignored intelligence up to this point, this is a development worth watching.

根据其自报的基准测试成绩，该模型将位居一份涵盖开源与闭源模型的广泛榜单之首（尽管榜单中似乎并未纳入领先的美国初创公司）。对于一家历来在硬件领域明显领先、但此前对智能领域基本忽视的公司而言，这是一个值得关注的进展。

The New Elephant: Is Astra the Answer?

新的大象：Astra 是答案吗？

What about OpenAI and Anthropic? As I mentioned above, VLMs have always been a critical layer in the brain quest, but Robotics has never been the public focus of these companies.That might be quickly changing.

那么 OpenAI 和 Anthropic 呢？如前文所述，视觉语言模型（VLM）始终是这场“大脑探索”中的关键一环，但机器人技术从来都不是这些公司公开关注的焦点。而这种情况可能正在迅速改变。

In August, Anthropic released a new “Model Hardware Standard” which *“enables AI agents to operate multiple lab and manufacturing instruments, such as microscopes, liquid handlers, and robotic arms, in parallel, and perform intricate tasks ranging from routine drug discovery experiments to laser calibration on a quantum computer.”*

8月，Anthropic 发布了一项新的“模型硬件标准”（Model Hardware Standard），它*“使 AI 智能体能够并行操作多种实验室和制造设备，例如显微镜、液体处理工作站和机械臂，并执行从常规药物发现实验到量子计算机激光校准等精细任务。”*

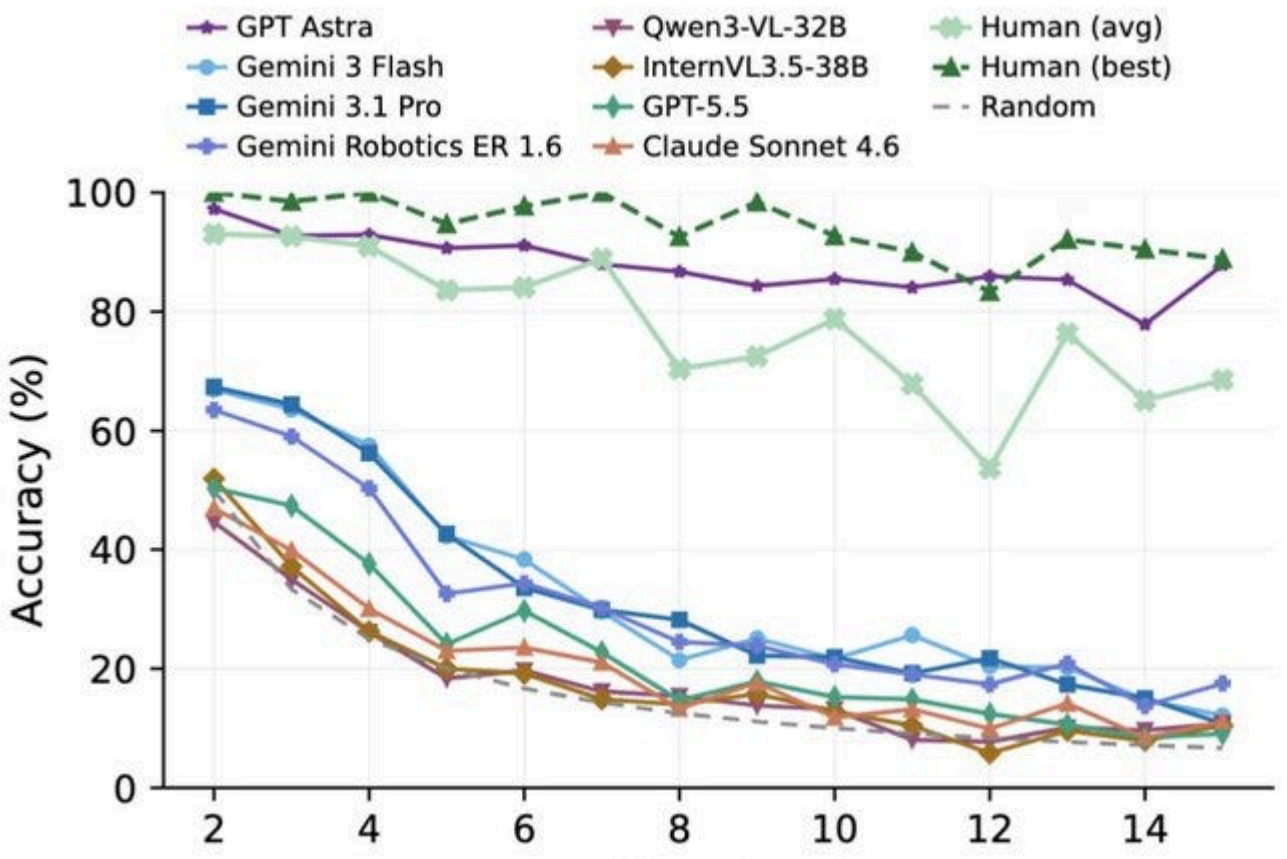


(Ok, let's pause for an ironic recognition that as these labs are warning about total AI takeover and safety, they are also trying to give the models the ability to not simply envision but directly produce novel biological compounds in lab settings… what could go wrong?)

（好，让我们先停下来做个颇具讽刺意味的观察：这些实验室一边警告 AI 可能全面接管并大谈安全，一边又在设法赋予模型这样的能力——不只是设想，而是在实验室环境中直接合成新型生物化合物……还能出什么岔子呢？）

One of the biggest jumps in Fable 5.1 vs. 5.0 was in its spatial reasoning, but OpenAI's Astra seems to be operating in a whole new ballpark. Some researchers claim Astra even exceeds “average” human levels.

Fable 5.1 相较于 5.0 最大的跃升之一体现在空间推理能力上，但 OpenAI 的 Astra 似乎已经处于完全不同的量级。一些研究者甚至声称 Astra 已超越“普通人”水平。



Translating this to table-top task completion results in ~3x performance over open-sourced VLA Molmo ACT2 according to Robotic benchmarker Robocurve.

据机器人基准测试机构 Robocurve 的评估，将此能力落实到桌面任务完成上，性能约为开源 VLA 模型 Molmo ACT2 的 3 倍。



There are still major limitations for using frontier models to operate robots (especially off the tables) including the compute and latency of frontier models.

使用前沿模型来驱动机器人（尤其是在桌面场景之外运作）仍存在重大局限，其中包括前沿模型本身的算力与延迟问题。

While it's possible that TradAI giants simply leapfrog into a pole position past the other physically focused upstarts, I don't think that will be the case. Rather I think this is a very positive development for the field that is likely to raise the tide for most boats.

虽然传统 AI 巨头有可能一步跨越、直接抢占领跑位置并超越其他专注物理领域的新兴玩家，但我认为情况不会如此。相反，我认为这对该领域是一个非常积极的发展，很可能让大多数参与者水涨船高。

First it provides strong *narrative confirmation* to take robotics very seriously. Second, any breakthrough catalysts in spatial reasoning and physical manipulation even in the closed-source frontier can provide a roadmap for replication as we've seen everywhere else in AI development. **Finally, these breakthroughs are likely to be harnessed by hardware-focused robotics integrators and lead to more rapid capability and deployment.**

首先，这为"必须高度重视机器人技术"提供了强有力的***叙事确认***。其次，空间推理与物理操作方面的任何突破性催化剂，即便出现在闭源前沿领域，也能为后续复制提供路线图——正如我们在 AI 发展的其他领域反复见到的那样。**最后，这些突破很可能被专注于硬件的机器人集成商所吸收利用，从而带来更快的能力提升与部署落地。**

Even with improving models, robots still need more training…

即使模型在不断改进，机器人仍需要更多训练……

How to Train Your Robot

如何训练你的机器人

There isn't one single consensus vision of what actually needs to happen to accelerate robotics autonomy. But in the broadest terms, closing the data deficit is probably the closest.

对于究竟需要什么才能加速机器人自主化，目前并不存在一个统一的共识愿景。但从最宽泛的角度来说，“弥补数据缺口”或许是最接近共识的答案。

While the latest World models can create remarkably realistic looking video, the ability to actually match real-world physics is much harder (think of things like object-specific friction, the crumpling of different materials, the flexibility of fabrics). Further, these models can't perfectly model the robots' own hardware (especially in generic sims with custom / iterative new forms).

尽管最新的世界模型能够生成外观极其逼真的视频，但要真正契合现实世界的物理规律却要困难得多（想想物体特定的摩擦力、不同材料的褶皱方式、织物的柔韧性等）。此外，这些模型还无法完美建模机器人自身的硬件（尤其是在用通用仿真器来模拟定制化或迭代中的新型形态时）。

Meanwhile, VLAs have an even more extreme data deficit. There isn't an internet-scale repository of “first person” videos of humans performing tasks that they want to teach to robots. There's even less data of *robots* doing the tasks.

相比之下，VLA 面临的数据缺口则更为极端。目前并不存在互联网规模的“第一人称”视频库，记录人类执行那些希望教给机器人的任务。而有关 *机器人*本身执行这些任务的数据更是稀缺。

Simulated data is extraordinarily cost efficient in imagining thousands of near-real environments. But at the end of the day, a model is a model – at the edges, it is still formulaic to a fault. Even tiny distinctions often mean that robots can *frolic in the sims* but *flop in the streets*.

模拟数据在构想数千个近乎真实的环境方面极具成本效益。但归根结底，模型终究只是模型——在边缘场景上，它依然公式化得过了头。哪怕极其微小的差异，也常常意味着机器人能***在模拟中撒欢***，***到了街上却摔跟头***。

How to close this **“Sim-to-real gap”?**

如何弥合这一**“模拟到现实的差距”（Sim-to-real gap）**？

More Data, More Compute, More Robots.

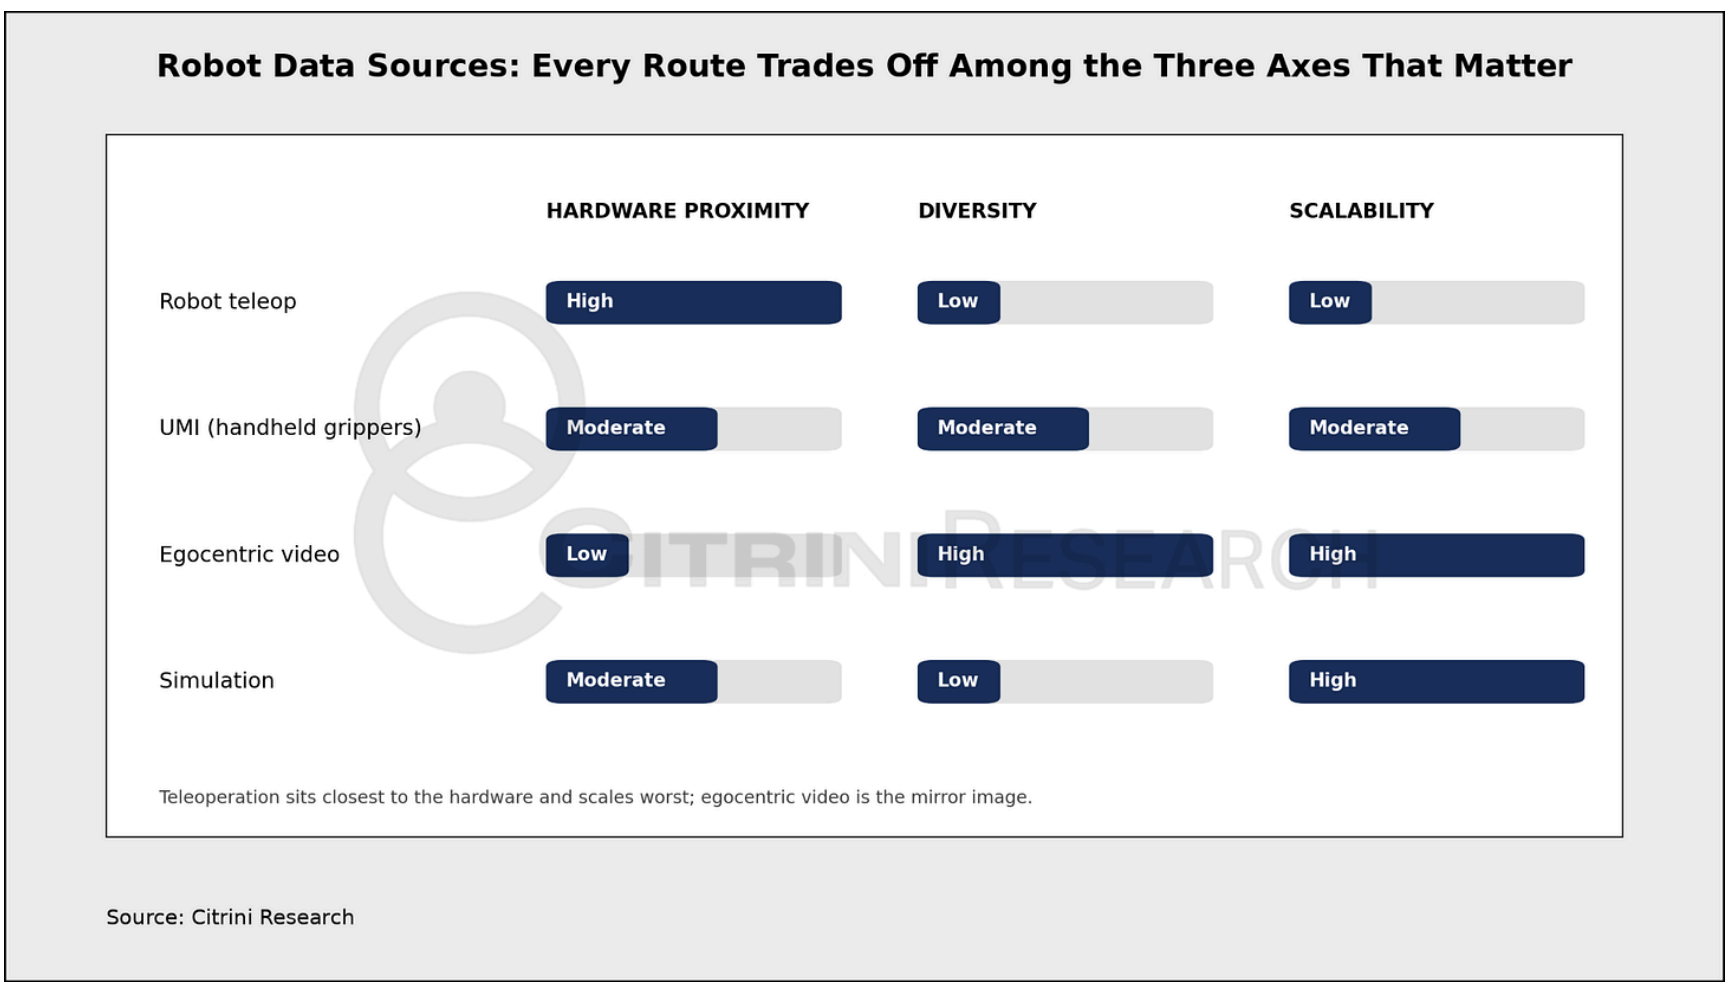
更多数据、更多算力、更多 **机器人。**

Data:

数据：

I teased it already, but there are three primary methods to teach robots (before they begin to train themselves); egocentric data, UMI demonstration, and teleoperations. Each has tradeoffs and the “best” version depends on who you ask. Still the chart below is a decent summary:

前文已经提过，在机器人开始自我训练之前，教它们主要有三种方法：第一人称视角数据、UMI 演示和遥操作。每种方法各有取舍，哪种“最佳”取决于你问谁。不过，下面这张图算是一个不错的总结：



Egocentric Data: Egocentric is a fancy word for first-person, and it means strapping a camera to someone's head (and sometimes your hands). Creating a dataset of a million dishwashings across millions of kitchens should help generalization.

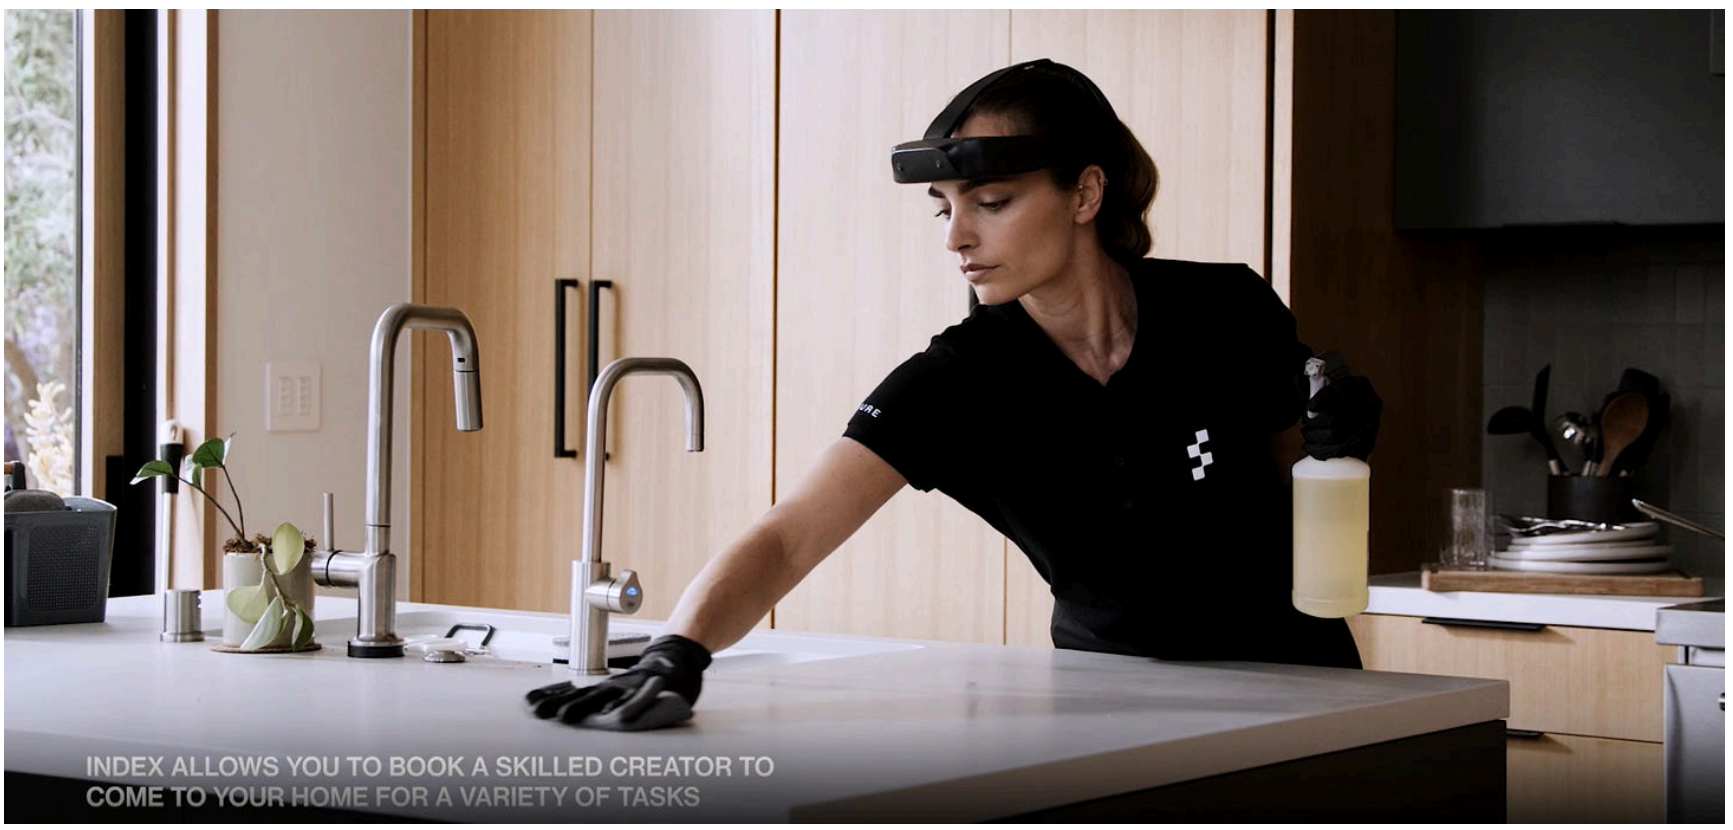
第一人称数据："Egocentric"（自我中心视角）不过是“第一人称”的时髦说法，指的是把摄像头绑在人的头上（有时也绑在手上）。构建一个覆盖数百万个厨房、包含一百万次洗碗动作的数据集，应当有助于提升泛化能力。

This “egocentric” data rush has pushed companies like Instawork, which provides contract labor, to offer extra to strap rigs on employees across industries.

这股“第一人称”数据热潮，促使 Instawork 这类提供合同用工服务的公司支付额外报酬，让各行各业的员工佩戴数据采集装置。

Some are trying to build their own – Figure just publicly released its Index data collection platform, which is live across 100+ countries with over 44,000 weekly active users. The app is processing 30 minutes of video uploads every second; 4.9 years of human work every day uploaded and Figure has paid out $15 million to date.

另一些公司则尝试自建——Figure 刚刚公开发布了其 Index 数据采集平台，目前已覆盖 100 多个国家，周活跃用户超过 4.4 万。该应用每秒处理的视频上传量达 30 分钟；每天上传的人类工作素材相当于 4.9 年的工作时长，Figure 迄今已为此支付 1,500 万美元。



There are also “datafarms” popping up on top of the existing cheap manual labor places like India. (Let’s not beat around the bush, there’s a dystopian feel in all of this that’s hard to shake.)

在印度这类原本就聚集廉价体力劳动力的地方，如今也纷纷涌现出“数据农场”（datafarm）。（不拐弯抹角地说，这一切都透着一种难以摆脱的反乌托邦气息。）



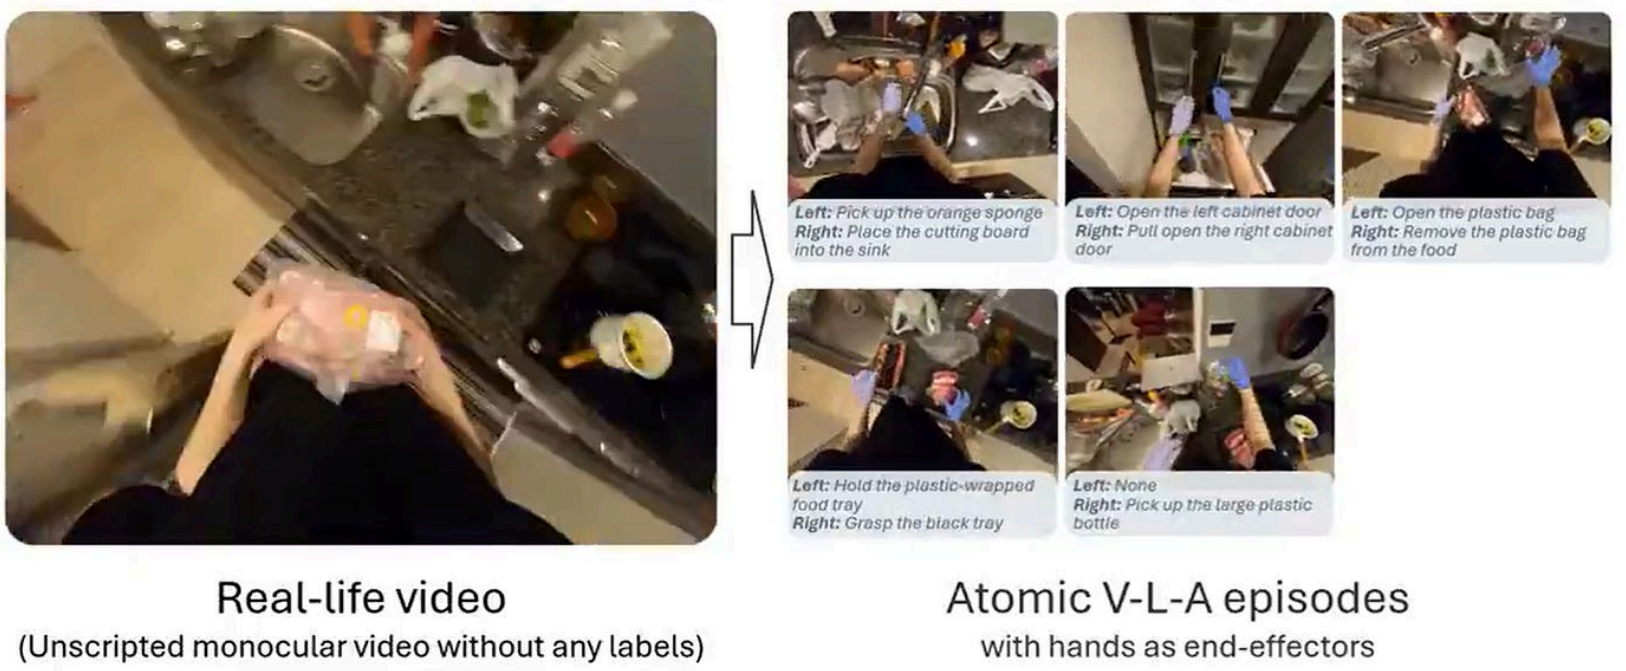


Still, egocentric data suffers from embodiment problems – watching a human sew clothing with five fingers isn’t the most helpful if the robot trying to learn only has a claw.

不过，自我中心数据存在具身形态差异问题——如果试图学习的机器人只装有一只爪形夹持器，那么观看人类用五根手指缝制衣服，对它而言帮助并不大。

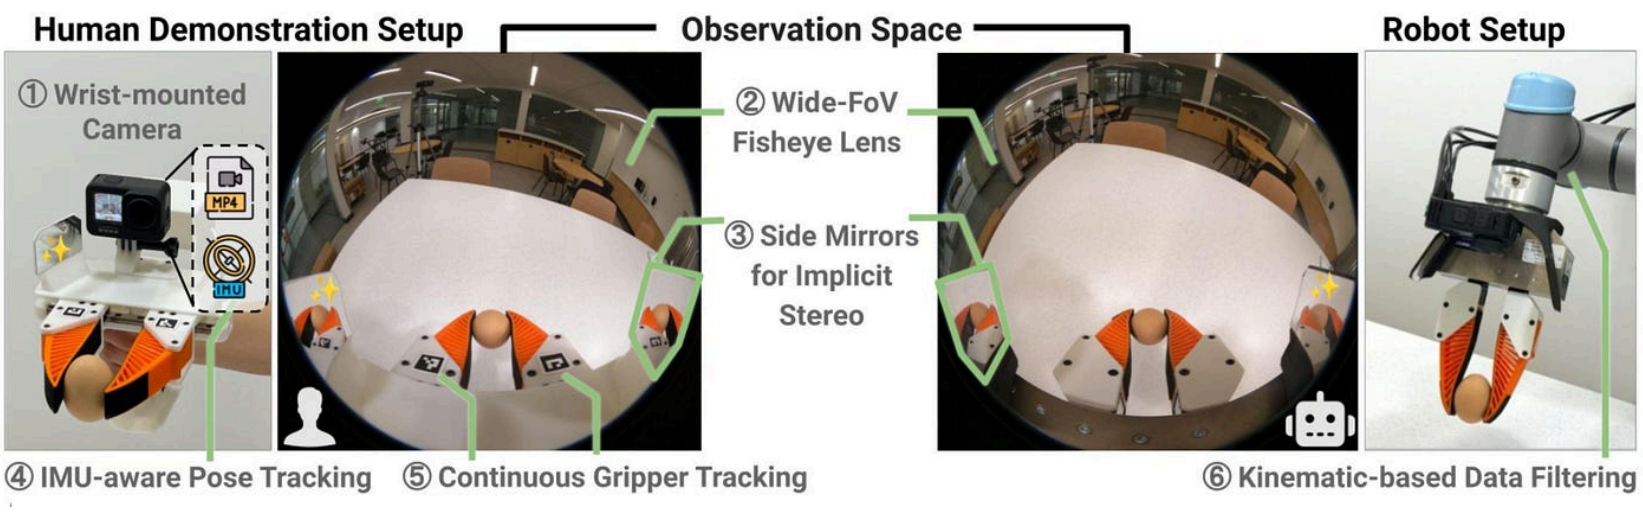
Projects like Microsoft’s open-sourced VITRA attempt to improve egocentric data by treating hands as dexterous end-effectors, breaking long videos into atomic tasks, generating language instructions, and reconstructing frame-by-frame 3D hand + camera motion.

以微软开源的 VITRA 为代表的项目，正尝试改进自我中心数据：将人手视为灵巧末端执行器，把长视频拆解为原子任务，生成语言指令，并逐帧重建手部与相机的三维运动。



UMI Demonstration: are similar to egocentric data capture in that a robot learns from human demonstrations. But rather than forcing the robot to translate from hand, it's the human who takes the robotic form.

UMI 示范采集： 这种方式与自我中心视角数据采集类似，都是让机器人从人类示范中学习。但区别在于，不必强迫机器人去适应从人手到机械臂的动作转换，而是由人类来主动采用机器人的形态。



This route fixes the "embodiment" gap and is often the technique being utilized on the most impressive VLA demos. But this shortcut could end up crippling the ultimate upside by limiting training to grippers with limited dexterity.

这一路线弥合了"具身"（embodiment）层面的差距，许多最令人印象深刻的 VLA 演示往往正是采用这一技术。但这种捷径可能最终削弱上行空间，因为它将训练局限在灵巧度有限的夹爪上。

It's also harder to scale as you need the specialized equipment – less accessible than an iPhone/GoPro and camera.

这种方式也更难规模化，因为需要专用设备——其获取门槛高于一部 iPhone/GoPro 外加相机。

Tele-operation: Teleops occur in a number of ways and can be used both for task-specific post training as well as foundational data for pretraining. Tele-op seems to have fallen out of favor, relative to egocentric data capture and demonstrations.

遥操作： 遥操作有多种实现方式，既可用于面向特定任务的后训练，也可作为预训练的基础数据。相比第一人称视角的数据采集与演示，遥操作似乎已经不再受青睐。

Meta's Quest VR headset is a major workhorse in this space and was visible almost everywhere I saw tele-ops in action.

Meta 的 Quest VR 头显是这一领域的主力设备，在我实地看到遥操作运行的地方，几乎处处可见它的身影。



Other methods include sensor-enabled gloves or body suits that can triangulate motion and replicate them in the hardware, like the picture I snapped below.

其他方法包括配备传感器的手套或体感服，可以对动作进行三角测量并在硬件中复现，就像我随手拍下的下面这张照片所示。



Beyond data ***volume***, there are questions of data quality, labeling and intent that can be optimized as well. We'll leave that for another day.

除了数据***体量***之外，数据质量、标注与意图等方面同样存在可优化之处。这个话题我们留待日后讨论。

Scaling:

规模化：

Can you call yourself "Physical AI" without scaling laws?? There has been a clear shift in tone in the robotics industry in the past several months meant to evoke the LLM development curve. The answer, perhaps, requires a whole lot more data and compute.

没有缩放定律，你还能自称"物理AI"吗？？ 过去几个月，机器人行业的论调发生了明显转变，意在呼应大语言模型（LLM）的发展曲线。而答案，或许需要多得多的数据和算力。

Dyna's self-published data on its DYNA-2 WAM shows its bots achieving a 96% success rate replicating demonstration of "seen tasks" (those that were part of its training data), but more importantly were able to also achieve a 66% success rate on "unseen tasks" after receiving a single demonstration. That is, if it was trained on 100,000 hours of video. At 10,000 hours, the success rate drops to 20%.

Dyna 自行公布的 DYNA-2 WAM 数据显示，其机器人在复现"见过的任务"（即训练数据中包含的任务）演示时成功率达 96%；更关键的是，在仅接收一次演示之后，其在"未见过的任务"上也能取得 66% 的成功率。也就是说，前提是用 10 万小时的视频进行训练；若训练数据只有 1 万小时，成功率会降至 20%。

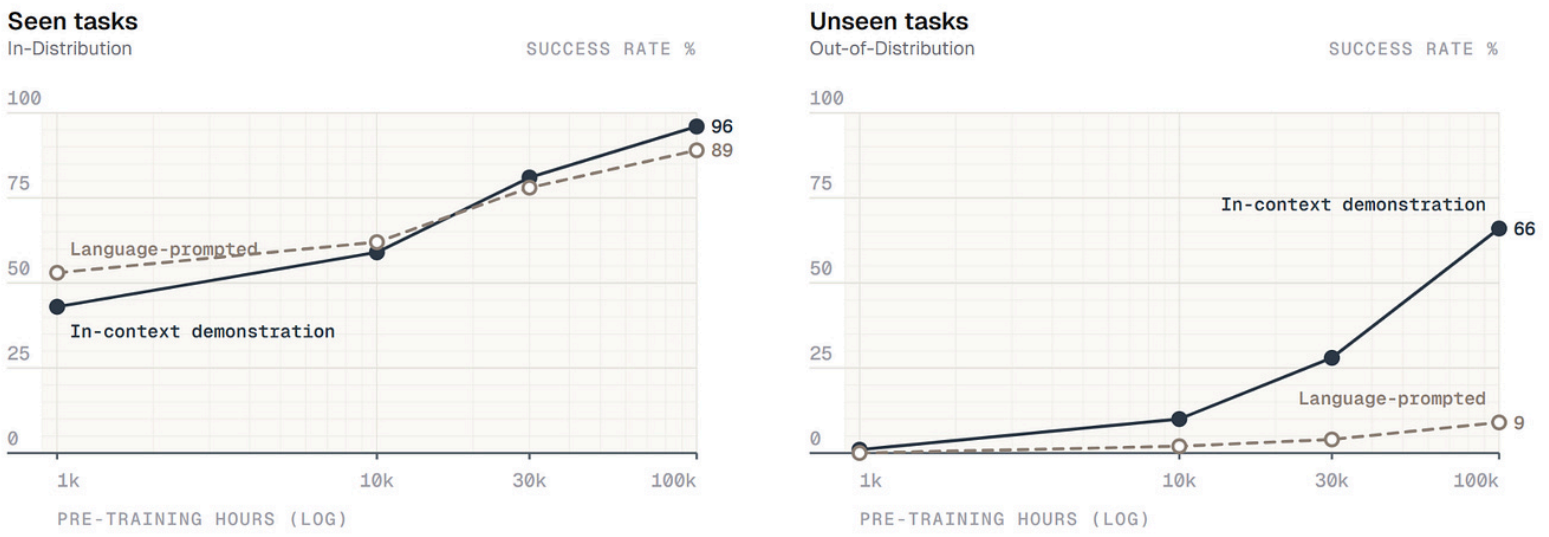


Fig. 4 Success rate against pre-training scale, for in-context and language-prompted VLA policies.

In other words, (if you squint hard and don't look too closely at the Y-axis) they are saying they believe scaling will work***.***

换言之，（只要你眯起眼睛使劲看、不去细究纵轴刻度）他们其实是在说：他们相信规模化这条路走得通***。***

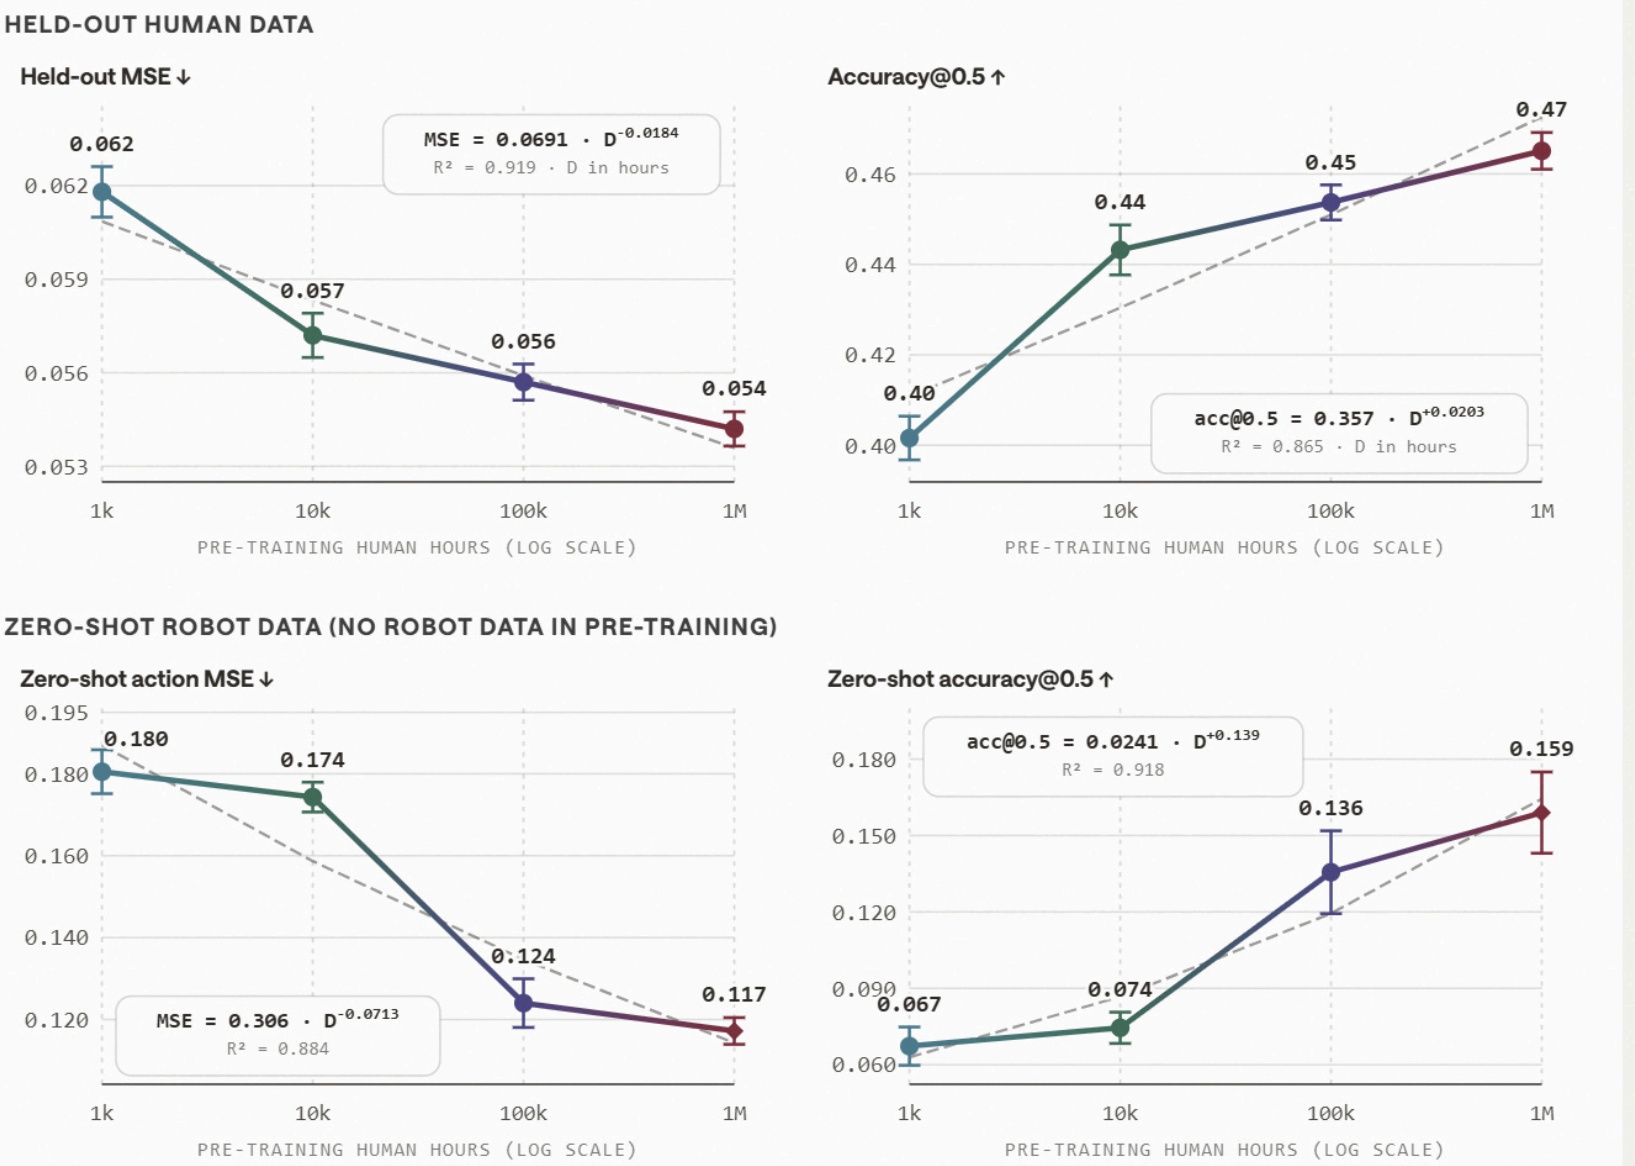


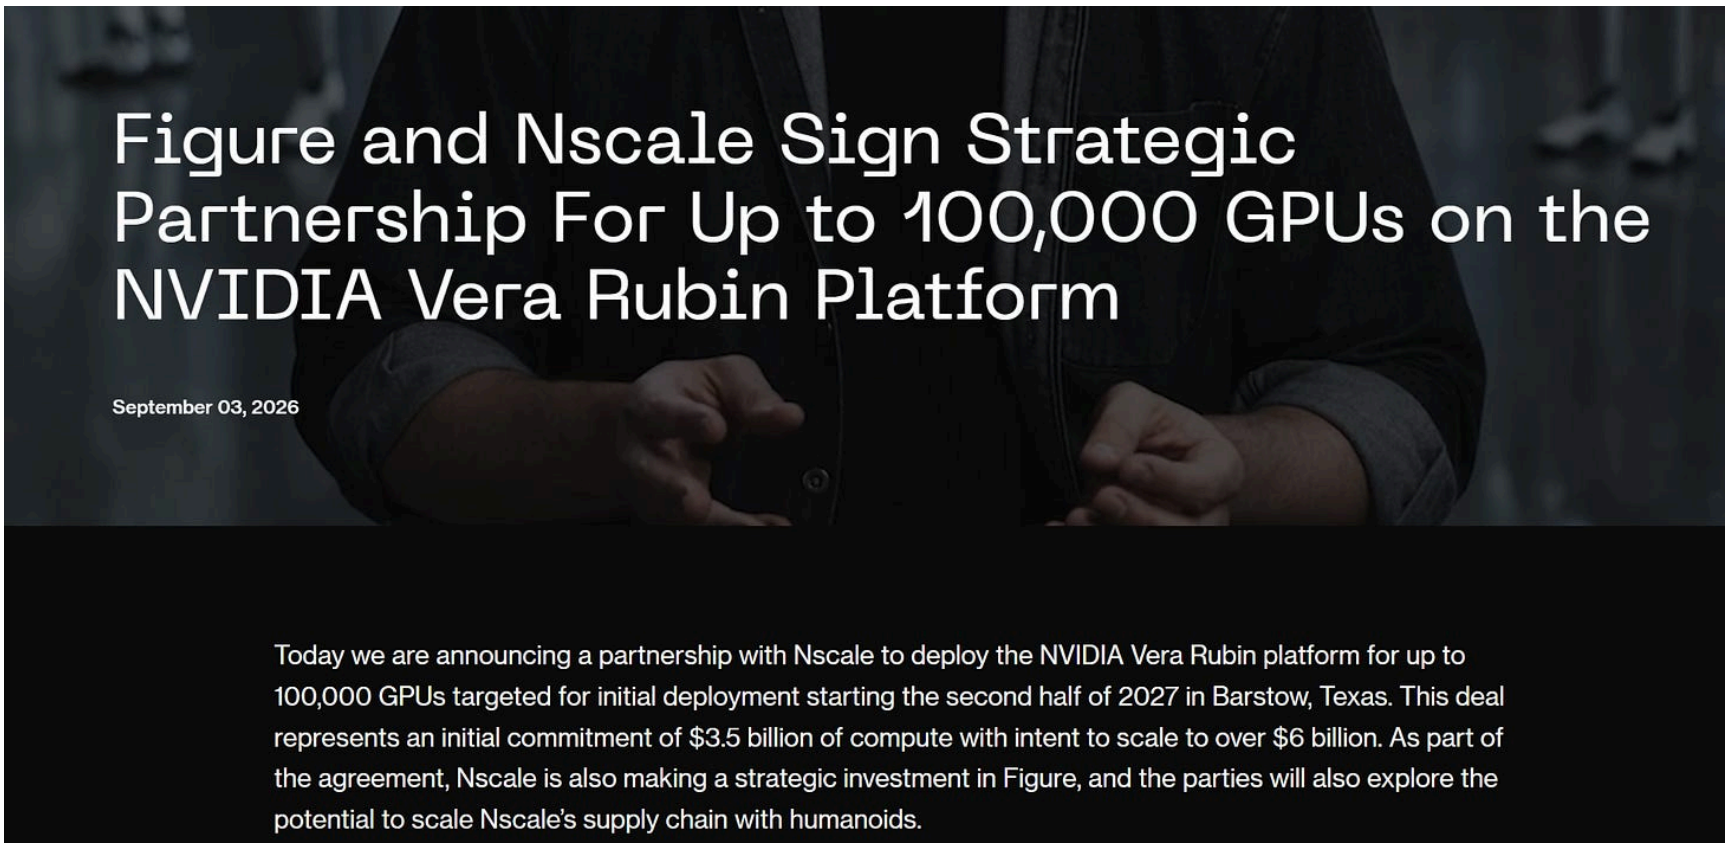
Figure 2. Scaling laws for Dyna-2 world-action models trained on a ladder of human video data, from 1000 hours to 1000000 hours. Not only Dyna-2 exhibits scaling trends on held-out human data, it also demonstrates the emergence of human-to-robot transfer scaling law on robot data the model has never seen during training.

And again, it's not just a single shop making these claims. Both VLA specialists and OEM generalists are saying similar things.

再说一次，做出这类判断的并非只有一家公司。专注 VLA 的专业厂商和 OEM 通用型厂商都在表达相似的观点。

Figure AI has essentially declared the same conclusions, first committing to $1 billion in compute spend over the next year on August 26th, before announcing a strategic partnership with Nscale to build a 100,000 GPU Vera Rubin cluster while committing to at least $3.5 billion in compute spending.

Figure AI 实际上已宣告了同样的结论：8月26日先承诺在未来一年投入10亿美元算力支出，随后又宣布与 Nscale 达成战略合作，建设一个由10万块 GPU 组成的 Vera Rubin 集群，并承诺至少投入35亿美元的算力支出。



Even with the necessary caveats, if Figure actually dedicated all 100,000 Rubin GPUs to one synchronized training run, it would be an authentic frontier-scale supercomputer — capable of reproducing GPT-4-era runs in hours, recent frontier runs in days to weeks. With the combination of the proprietary Index video data and perhaps the most rounded and proven OEM platform, we should take these aspirations seriously.

即便附上必要的保留意见，如果 Figure 真的将全部 10 万块 Rubin GPU 投入一次同步训练运行，那也将是一台名副其实的前沿规模超算——能够在数小时内复现 GPT-4 时代的训练运行，在数天到数周内复现近期前沿规模的训练。结合专有的 Index 视频数据，以及可能是最完善、经过验证的 OEM 平台，我们应当认真对待这些抱负。

At a minimum the course of direction is clear: frontier-scale training to try to break the wall of generalization.

至少，方向已经清晰：以前沿规模的训练，尝试打破泛化的壁垒。

Deployment:

部署：

Robotics' chicken vs. egg question is simple: *do you build a bunch of dumb robots to try to make them smart?*

机器人领域的"先有鸡还是先有蛋"问题很简单：*你是先造出一堆"愚蠢"的机器人，再设法让它们变聪明？*

By sending robots out into the world (either in limited customer pilots or in your own local sandbox) you gain one of the most valuable pieces: ***real-world hardware-specific data.*** At a baseline, you can develop an incredibly strong RL data for ambulation, balance, and autonomous navigation. Robots can "practice" specific tasks, which can be graded by humans for improved performance. And maybe most underrated at all is simply finding out where it breaks – metaphorically and literally.

把机器人送入现实世界（无论是小规模客户试点，还是自建的本地试验场），你就能获得最有价值的要素之一：***真实世界中针对特定硬件的数据。*** 最起码，你可以为行走、平衡与自主导航积累极强的强化学习（RL）数据。机器人可以"练习"特定任务，并由人类打分以改进表现。而最被低估的一点，也许只是找出系统会在哪里失灵——无论从比喻意义还是字面意义上。

What does it take to knock a humanoid off its balance? *What happens when it does?* How many million cycles of practical work does it take to stretch your tendons or blow your actuators? **What else can go wrong that you haven't even thought about yet?**

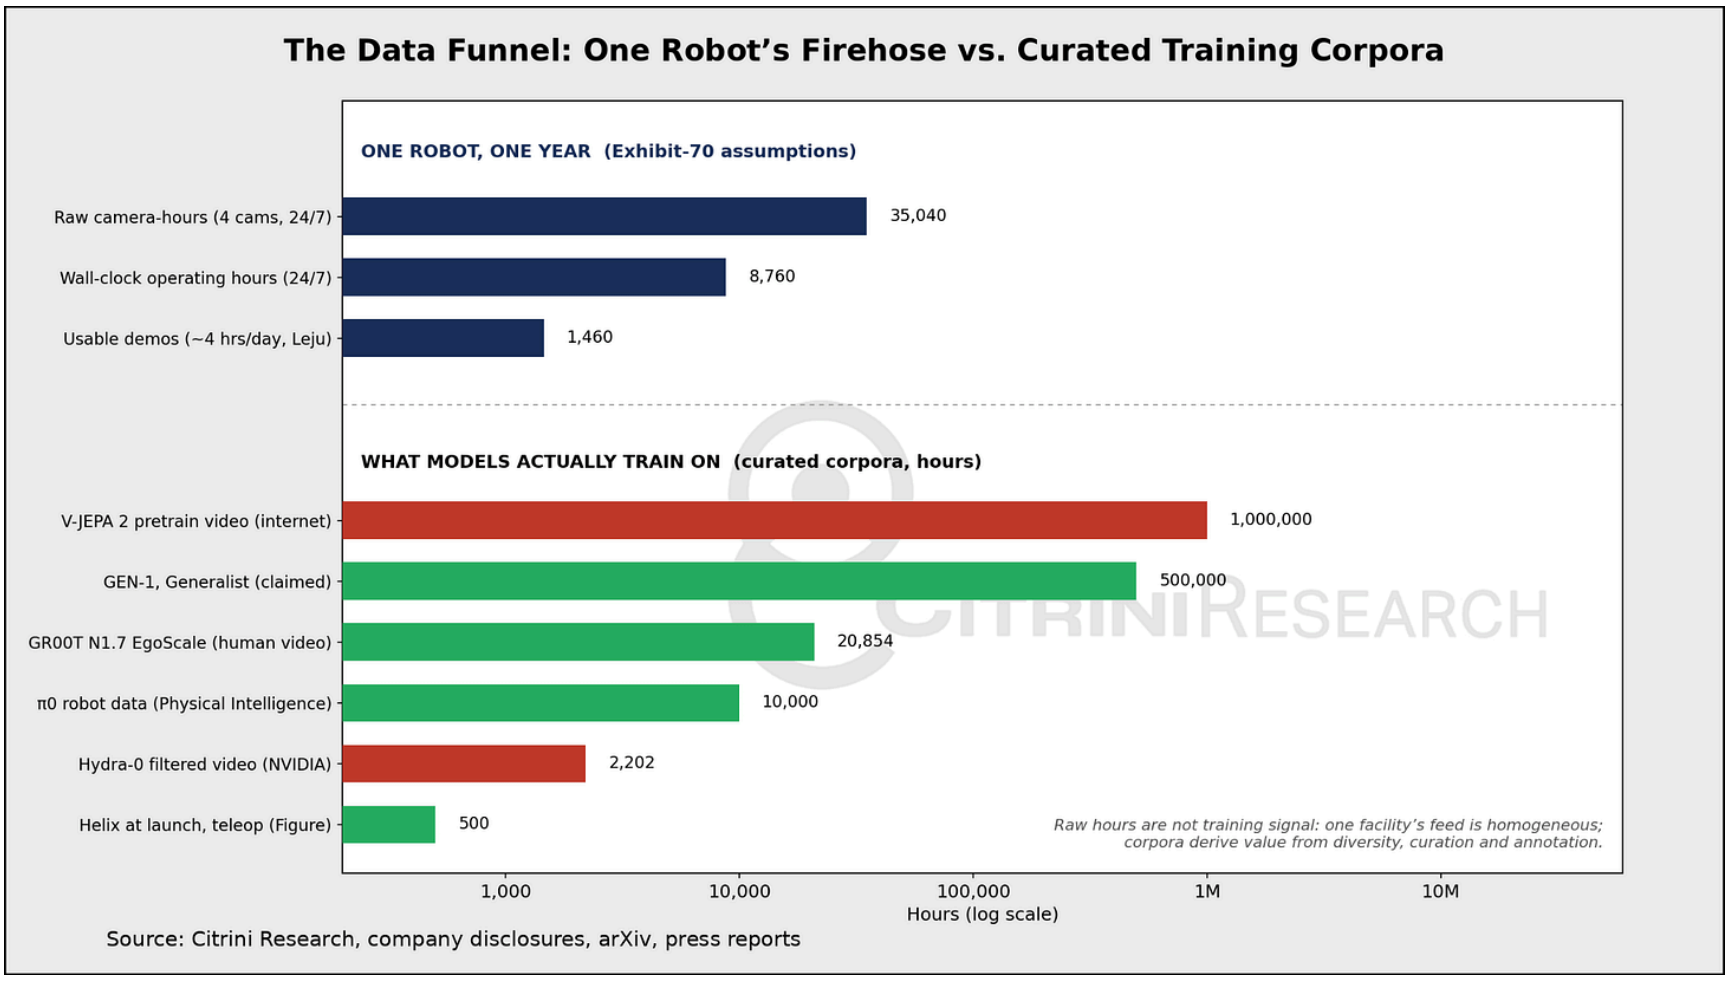
要让一台人形机器人失去平衡需要多大外力？*失去平衡之后会发生什么？* 要经过几百万次实际工作循环，才会拉伤你的"肌腱"或烧毁你的执行器？**还有哪些你压根没想到的故障可能发生？**

The idea on the surface isn't attractive financially, especially at the beginning. Full bots (especially humanoids) are the most expensive to manufacture at low-volume prototype stages. Meanwhile, the capabilities are too narrow to generate enormous third party economic demand.

从表面上看，这个想法在经济上并不吸引人，尤其是在起步阶段。整机机器人（尤其是人形机器人）在小批量原型阶段的制造成本最高。与此同时，其能力范围过于狭窄，难以催生可观的第三方经济需求。

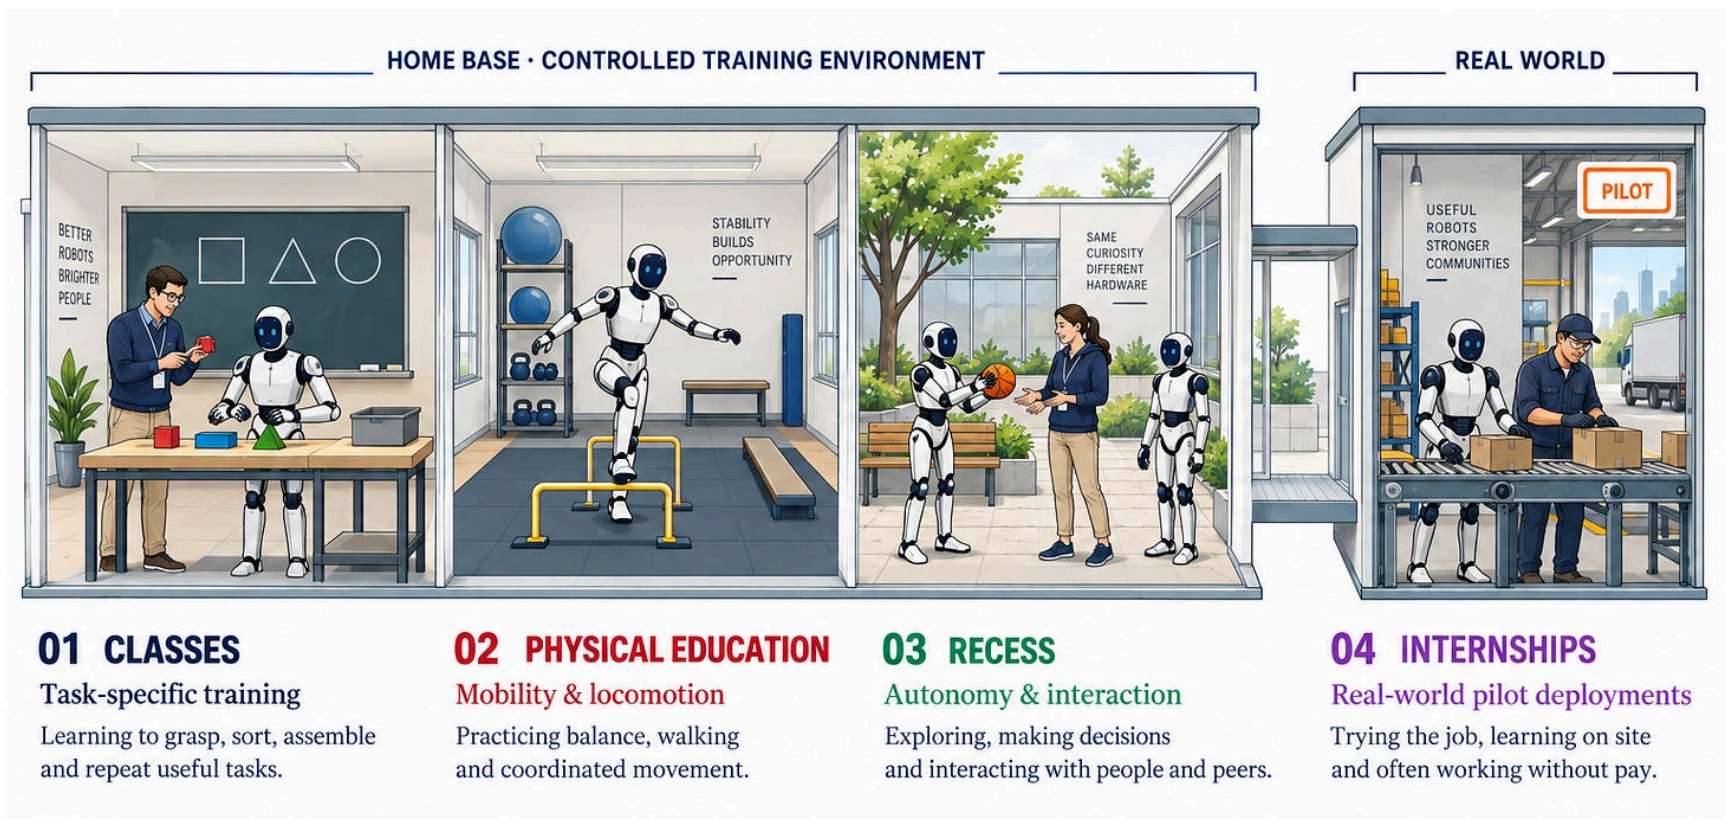
Particularly under capital constraints, it's an enormously risky flywheel bet – a meaningful upfront cash investment with little immediate return that ultimately drives down production costs, gathers valuable data, and closes the sim-to-real gap. It's the path that folks like Figure and Tesla appear to be taking.

尤其是在资本受限的情况下，这是一项风险极高的飞轮式押注——需要投入可观的前期现金，短期内几乎没有回报，但最终能降低生产成本、积累宝贵数据，并缩小仿真与现实（sim-to-real）之间的差距。Figure 和特斯拉等公司看来正在走的正是这条路径。



To anthropomorphize, humans spend two decades in "school" before they go to "work". This strategy looks a whole lot like school.

打个拟人化的比方，人类要先在"学校"里度过二十年，然后才去"上班"。这套策略看起来非常像学校。



Most of it happens in a controlled home base. For part of the day the robots take classes (task-specific training). Part of the day is in physical education (structured teaching of mobility and locomotion). Part of the day is recess (ambling about, learning *autonomy,* and how to interact with humans and each other). They even have internships – where they are sent out to work in the real world *(generally without pay)!*

大部分训练都在一个受控的"大本营"中进行。一天之中，机器人有一部分时间在上课（针对特定任务的训练）。有一部分时间在上体育课（对移动与运动能力进行结构化教学）。还有一部分时间是课间休息（四处走动，学习*自主性*，以及如何与人类和彼此互动）。它们甚至还有实习——被派到真实世界里去工作（*通常没有报酬*）！

If I was a betting man, here's where I think the true breakthroughs will come. But where do we grade the class today?

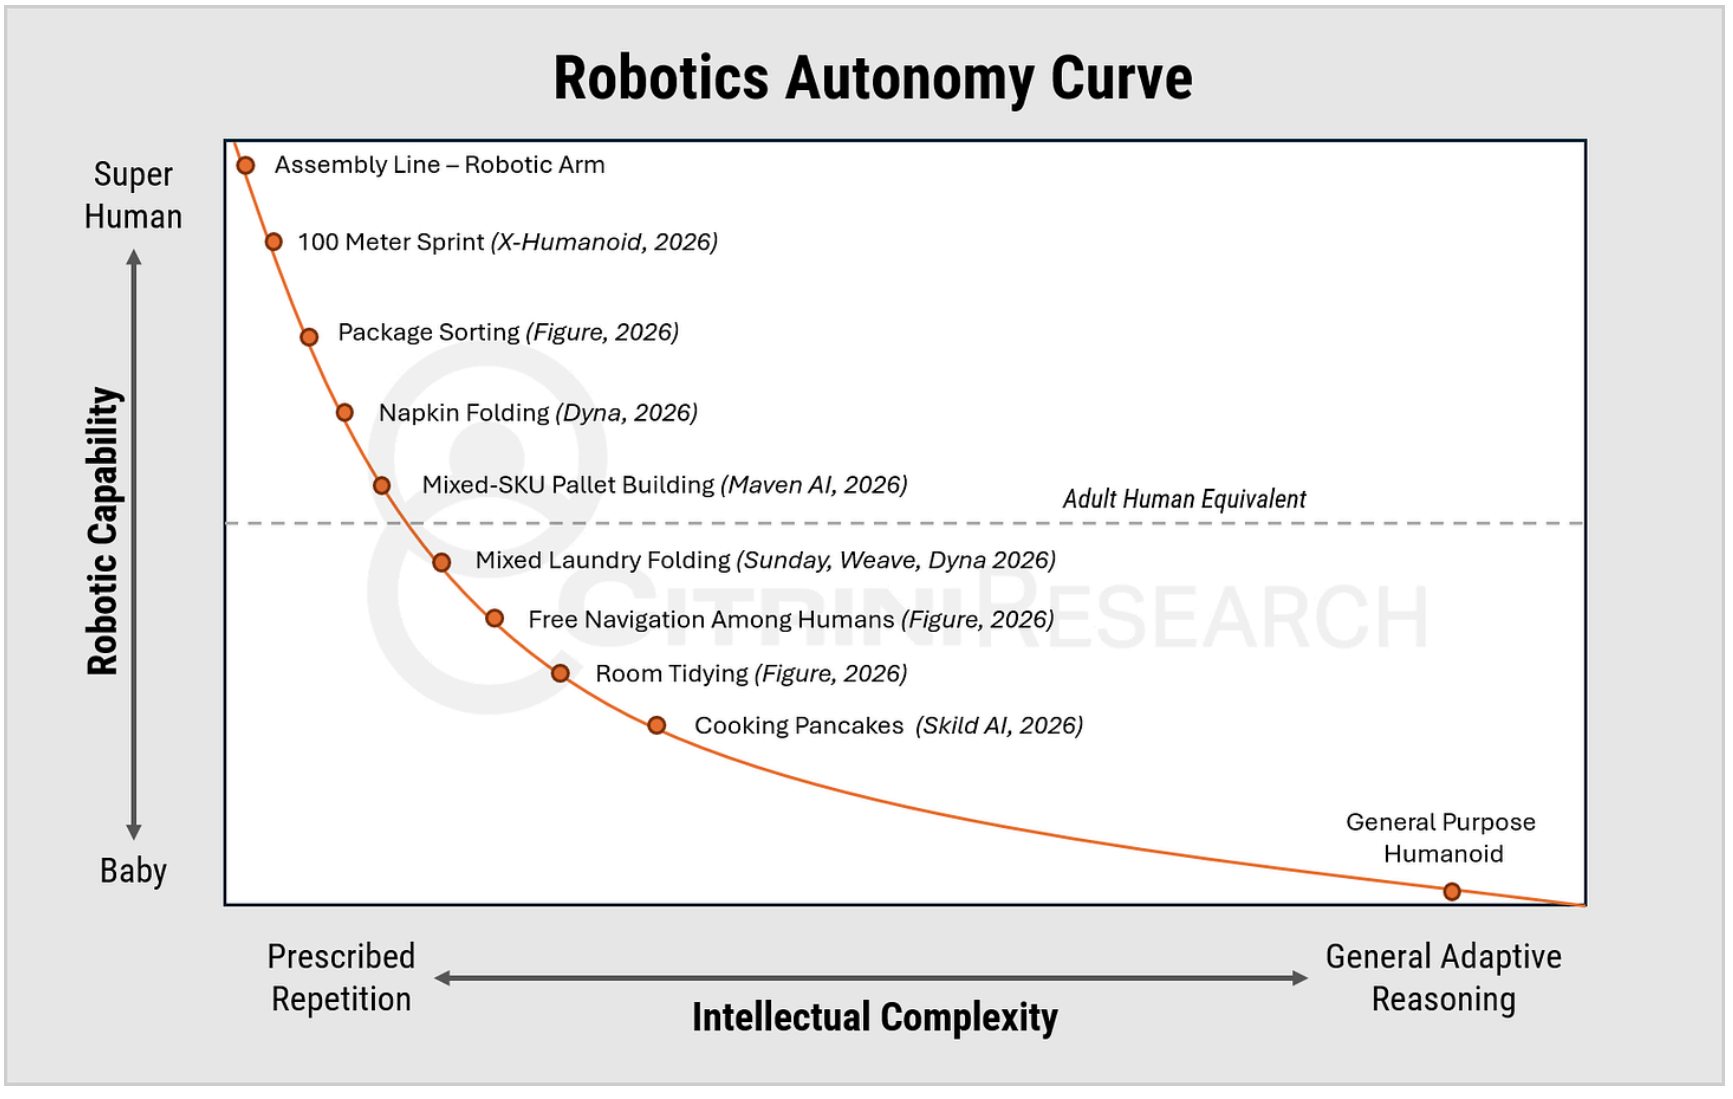
如果我是个愿意押注的人，我认为真正的突破将出现在这里。 **但今天我们该给这个班打几分？**

3) Capability and Commercialization

3) 能力与商业化

Based on what's been publicly released and what I've witnessed myself, here is a very non-exhaustive and inherently subjective attempt to actually map where we stand on the curve.

基于已公开发布的信息以及我本人的亲眼所见，下面是一次远非详尽、且本质上带有主观色彩的尝试，试图切实标定出我们目前在这条曲线上所处的位置。

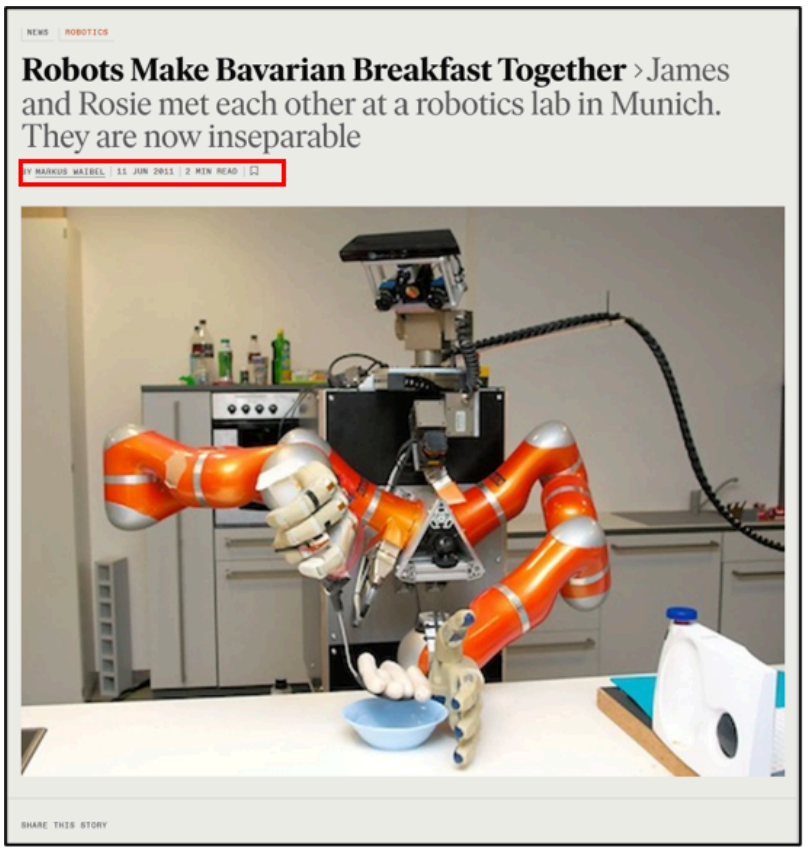


The tasks that have been demonstrated so far are ones that are well defined and repetitive, yet still have enough variability and edge cases that have kept them from being automated before.

迄今为止已演示的任务，都是那些定义明确、重复性强的任务，但它们此前仍存在足够的变数和边缘情况，因而一直未能实现自动化。

The pushback: breakfast-making and towel-folding demos were making headlines… in 2010.

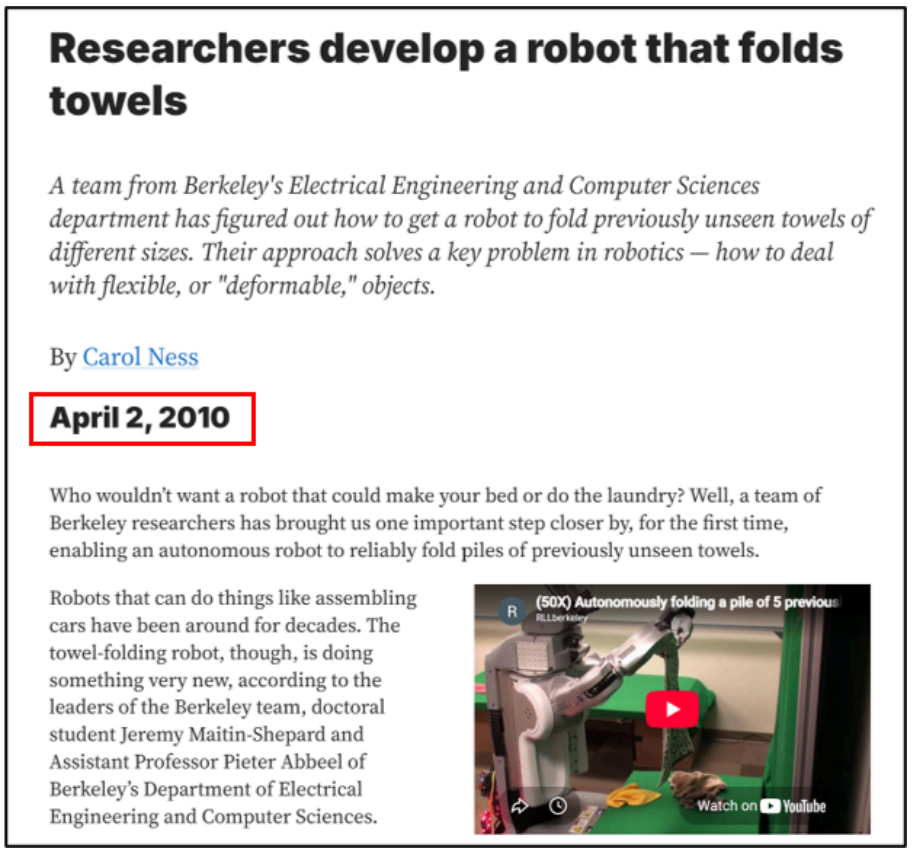
质疑的声音：做早餐、叠毛巾这类机器人演示，早在2010年就曾登上过新闻头条……


NEWS | ROBOTICS

Robots Make Bavarian Breakfast Together › James and Rosie met each other at a robotics lab in Munich. They are now inseparable

BY MARKUS WAIBEL | 11 JUN 2011 | 2 MIN READ

SHARE THIS STORY


Researchers develop a robot that folds towels

A team from Berkeley's Electrical Engineering and Computer Sciences department has figured out how to get a robot to fold previously unseen towels of different sizes. Their approach solves a key problem in robotics — how to deal with flexible, or "deformable," objects.

By Carol Ness

April 2, 2010

Who wouldn't want a robot that could make your bed or do the laundry? Well, a team of Berkeley researchers has brought us one important step closer by, for the first time, enabling an autonomous robot to reliably fold piles of previously unseen towels.

Robots that can do things like assembling cars have been around for decades. The towel-folding robot, though, is doing something very new, according to the leaders of the Berkeley team, doctoral student Jeremy Maitin-Shepard and Assistant Professor Pieter Abbeel of Berkeley's Department of Electrical Engineering and Computer Sciences.

But the former was not autonomous, while the latter was indeed an extraordinary early prototype. And remember, Google's DeepMind beat humans in Go back in 2015. Remember the frontloading? The point isn't that the ideas are brand new, but rather that foundations have been under construction for over a decade, and we might just be reaching the point of acceleration.

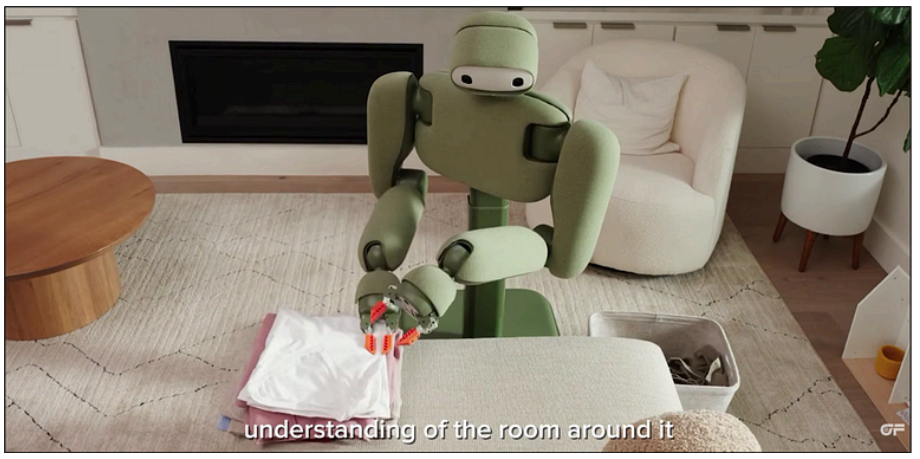
但前者并不具备自主性，而后者确实是一个了不起的早期原型。而且别忘了，谷歌的 DeepMind 早在 2015 年就在围棋上战胜了人类。还记得前置投入吗？重点不在于这些想法有多新，而在于这些基础已经建设了十多年，我们或许正抵达加速的临界点。

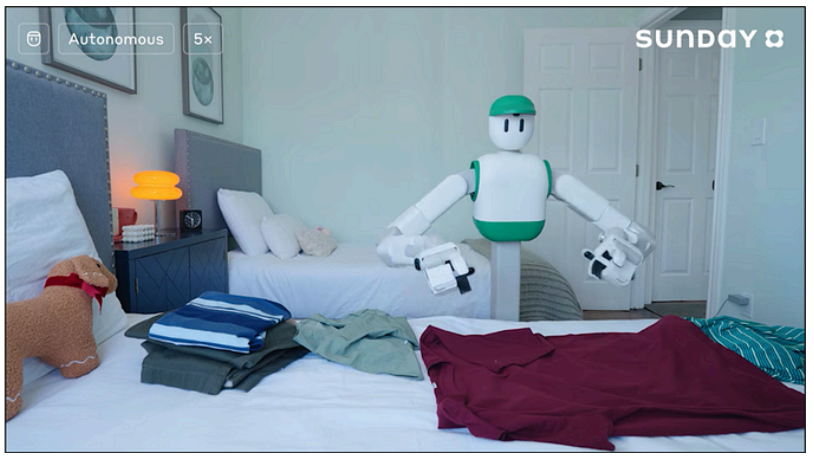
Learning from Laundry

从洗衣中学习

Take – mixed laundry folding. You and I think it's pretty boring *(watching a three hour demo of a robot folding is even less tantalizing).* But seeing it in person helps. Watching a fully autonomous robot sort, separate, and fold clothes of all shapes and varieties, operating continuously for hours… ***It's cool to watch a robot think.***

以混合衣物折叠为例。你我都觉得这事挺无聊的（*看机器人折叠衣物三个小时的演示就更提不起兴趣了*）。但亲眼见到会有完全不同的感受。看着一台全自主机器人对各种形状、各种材质的衣物进行分拣、分离和折叠，连续运转数小时……***看着一个机器人思考，真的很酷。***





“What’s the object? Looks like a sweatshirt. Is it inside out? No, good to go. How do I spread it flat on the surface without disturbing my piles? Whoops, one sleeve is stuck inside itself, I need to pull it out first. Pick the corners, fold here, fold here. Smooth the wrinkle, shake it out. Pick it up without ruining the fold and place it gently in the sweatshirt pile. Repeat…”

“这是什么物品？看起来像是一件卫衣。穿反了吗？没有，可以开始。怎么把它在台面上铺平，同时不打乱我已叠好的衣物堆？哎呀，有一只袖子卷进里面了，我得先把它抽出来。捏住边角，这里折一下，这里再折一下。抚平褶皱，抖一抖。在不破坏折叠形状的前提下拿起它，轻轻放到卫衣那一堆里。重复……”

Again, it’s not rocket science but it does require… ***intelligence.*** The best analogy I can come up with *(at the high risk of inciting parental backlash)* is like watching a baby begin to learn. They can’t really communicate, ***but you see that the lights are on***.

再说一次，这算不上什么高深的科技，但它的确需要……***智能。*** 我能想到的最贴切的类比*(冒着引发家长们强烈反感的高风险)*，就像看着一个婴儿开始学习。他们还无法真正交流，***但你能看出灯是亮着的***。

As for the quality, I’d call it very satisfactory for utility purposes today but still below “five star resort” or “very strict mama” standards. It will get better. At Dyna, its autonomous napkin folder (easier than mixed folding) increased throughput by 2.7x while improving quality over a year.

至于质量，我会说它目前已经能很好地满足实用需求，但仍未达到“五星级度假村”或“非常严格的妈妈”的标准。它会变得更好。在 Dyna，其自主餐巾折叠机（比混合折叠更简单）在一年时间里将吞吐量提升了 2.7 倍，同时质量也在改善。

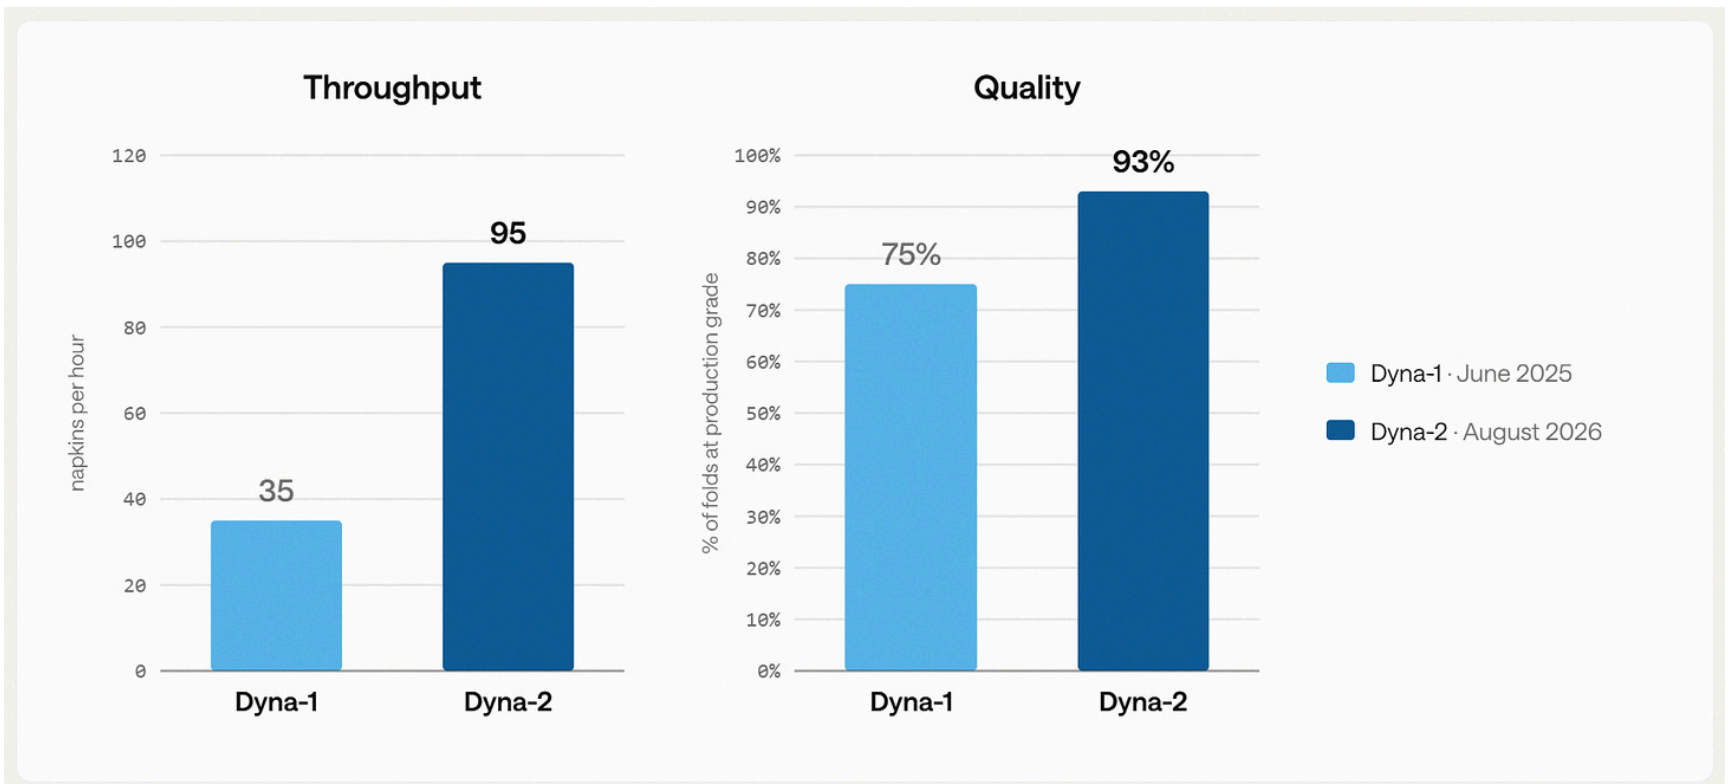


Figure 2. Dyna-2 folds 2.7× more napkins an hour than Dyna-1, at higher quality.

Is this a narrow task? **Maybe, but that doesn’t mean it’s small.**

这是一项狭窄的任务吗？也许是，但这并不意味着它规模小。

It’s a task required by commercial businesses like laundromats, hotels, or garment managers like Cintas. Dyna Robotics has robots folding in restaurant chain Din Tai Fung, hotel chain Best Western, uniform and linen service Prudential Overall, among others.

对于自助洗衣店、酒店等商业企业，或是像 Cintas 这样的制服管理服务商而言，这是一项必需的任务。 Dyna Robotics 的机器人已在连锁餐厅鼎泰丰、酒店集团 Best Western、制服与布草服务公司 Prudential Overall 等企业中进行折叠作业。

For consumers, folding is a ubiquitous chore. Would you pay $249/month to skip it? For most people today, probably not… but for some, the answer is yes. It’s those early adopters that allow for diverse real-world data collection, operational reps, while also driving unit production costs downward.

对消费者而言，叠衣服是一项无处不在的家务。你愿意每月支付 249 美元来免做这件事吗？ 对今天的大多数人来说，大概不会……但对某些人而言，答案是愿意。正是这些早期采用者，使企业得以收集多样化的真实场景数据、积累运营实操经验，同时推动单位生产成本下降。

Laundry is just the foot in the door. Several companies are rolling out commercial folding products but each of them is intently focused on expanding capabilities – loading and unloading machines, adding basic tasks and tidying, increasing mobility.

洗衣只是切入点。多家公司正在推出商用折叠产品，但每一家都在专注于扩展能力——装取洗衣设备中的衣物、增加基础任务和整理功能、提升移动能力。

Would you change your mind if it folds laundry and makes your bed? What if it has millions of in-home hours with positive reviews? What if it’s $100 or $50?

如果它不仅能叠衣服，还能帮你铺床，你会改变想法吗？如果它已累计数百万小时的家庭使用时长且广受好评呢？如果价格降到 100 美元或 50 美元呢？

This, to me, seems like the most probable model for commercialization across different domains. Tackle specific automation challenges that still require human labor and expand as the “autonomy curve” shifts upward.

在我看来，这似乎是各个领域最有可能的商业化路径：先解决那些仍需人力的具体自动化难题，再随着“自主能力曲线”的上移而逐步扩展。

Tackle things like **mixed palletization,** which means packing pallets of various SKUs together often at a distribution center for an end customer. General Mills automates an entire pallet of Cheerios coming out of the factory, but your local CVS doesn’t need a pallet of Cheerios. It needs an assortment of products, shapes, etc. from its regional warehouse. This is a job that is still largely done by humans who can grasp objects of any shape (literally and mentally) and figure out how best to stack them together. Automation systems in the past have generally failed to gain mass adoption because they are too rigid and constrained.

例如**混合托盘码垛**（mixed palletization），即在配送中心为终端客户将不同 SKU 的货品打包到同一托盘上。通用磨坊可以自动化处理刚出厂的整托盘 Cheerios，但你家附近的 CVS 门店并不需要一整托盘 Cheerios，它需要的是来自区域仓库的各类产品、各种形状的组合货品。这类工作目前仍主要由人工完成——工人既能抓起任意形状的物体，也能在头脑中规划出最佳堆叠方式。过去的自动化系统之所以普遍未能大规模普及，是因为它们过于僵化、限制太多。



Companies like Maven Robotics are building autonomous mixed-pallet systems with mobile based, dual-arm commercial robots – live in warehouse centers today. They can work within existing distribution centers and with other programmed automation, rather than requiring a rebuild around a prescriptive system.

像 Maven Robotics 这样的公司正在用搭载移动底盘的双臂商用机器人构建自主混合托盘系统——目前已在仓储中心投入实际运行。这类系统能够在现有配送中心内运行，并可与其他已编程的自动化设备协同工作，而无需围绕某种规定式系统进行重建。

This *niche-first* deployment model has important implications. Perhaps not as sexy as humanoids by the thousands, commercialization is rapidly gaining traction and the numbers are becoming real.

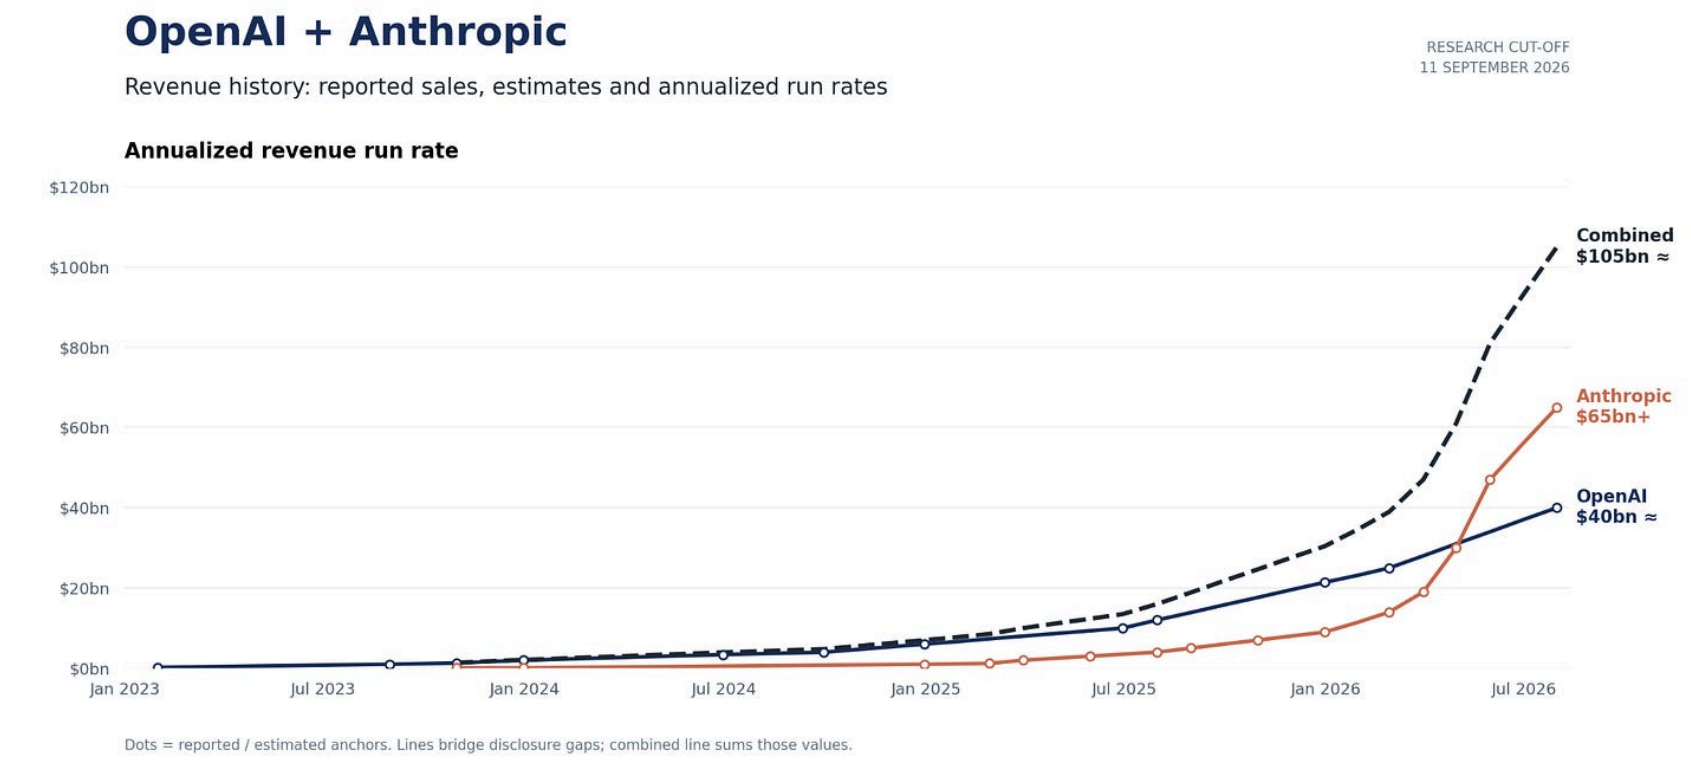
这种*利基优先*的部署模式具有重要意义。它或许不像成千上万台人形机器人那样引人注目，但其商业化正在迅速起势，数字也正变得真实起来。

Run Rates Accelerating

运行率加速增长

This week Skild AI put hard numbers on the table. Just ten months after its first commercial deployment, the company claims $100 million in ARR across 60 paying customers across “high-precision assembly” for NVDA/Foxconn, wire harnessing, and commercial kitchen applications. In context, this one company represents just ~0.1% of the current combined OpenAI + Anthropic run rate today, but is not far from their combined ARR back in 2022.

本周，Skild AI 拿出了实打实的数字。在首次商业部署仅十个月后，该公司宣称已实现 1 亿美元的年度经常性收入（ARR），拥有 60 家付费客户，业务覆盖面向 NVDA/富士康的“高精度组装”、线束加工以及商用厨房应用。作为参照，仅这一家公司的规模就只相当于当前 OpenAI 与 Anthropic 合计运行率的约 0.1%，但已与两者 2022 年时的合计 ARR 相去不远。

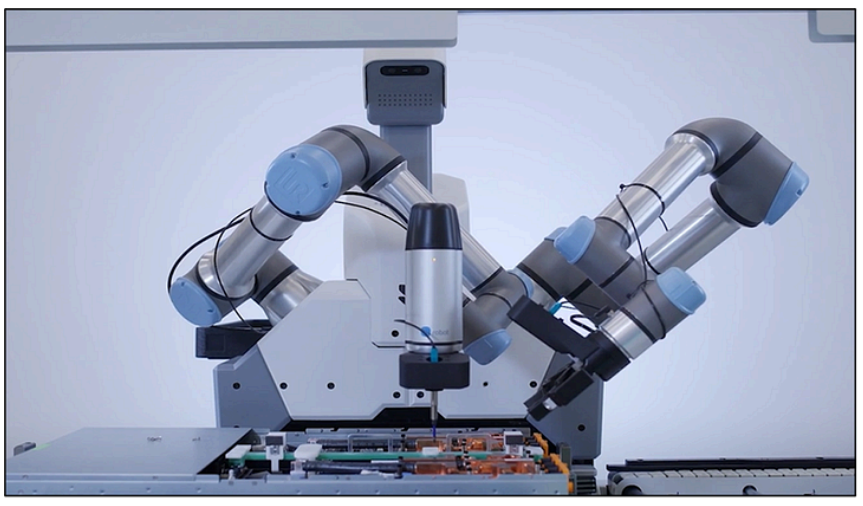


The disclosure also gives us a glimpse into what deployment looks like…

这份披露文件也让我们得以一窥其部署的实际形态……

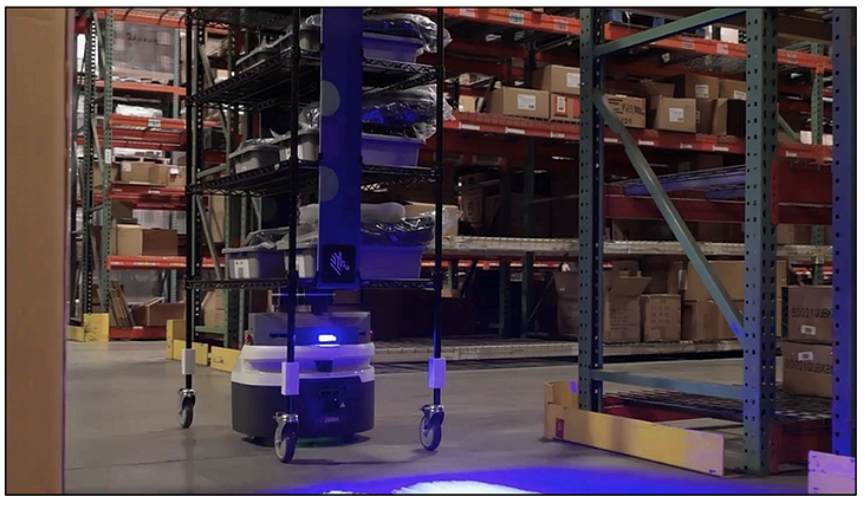
Skild appears to be primarily applying its brain to existing hardware across a variety of form factors, ranging from Roomba-looking warehouse AMRs (via its acquisition of Zebra/Fetch), to custom mounted hardware on Unitree quadrupeds, to fixed stereocameras with Universal Robots robot arms, to mobile based humanoids from Dexmate.

Skild 似乎主要是将其“大脑”应用于各类形态的现有硬件：从外形类似 Roomba 的仓储 AMR（经由其收购的 Zebra/Fetch），到加装定制硬件的 Unitree 四足机器人，再到搭配 Universal Robots 机械臂的固定式立体相机，以及来自 Dexmate 的带移动底座人形机器人。







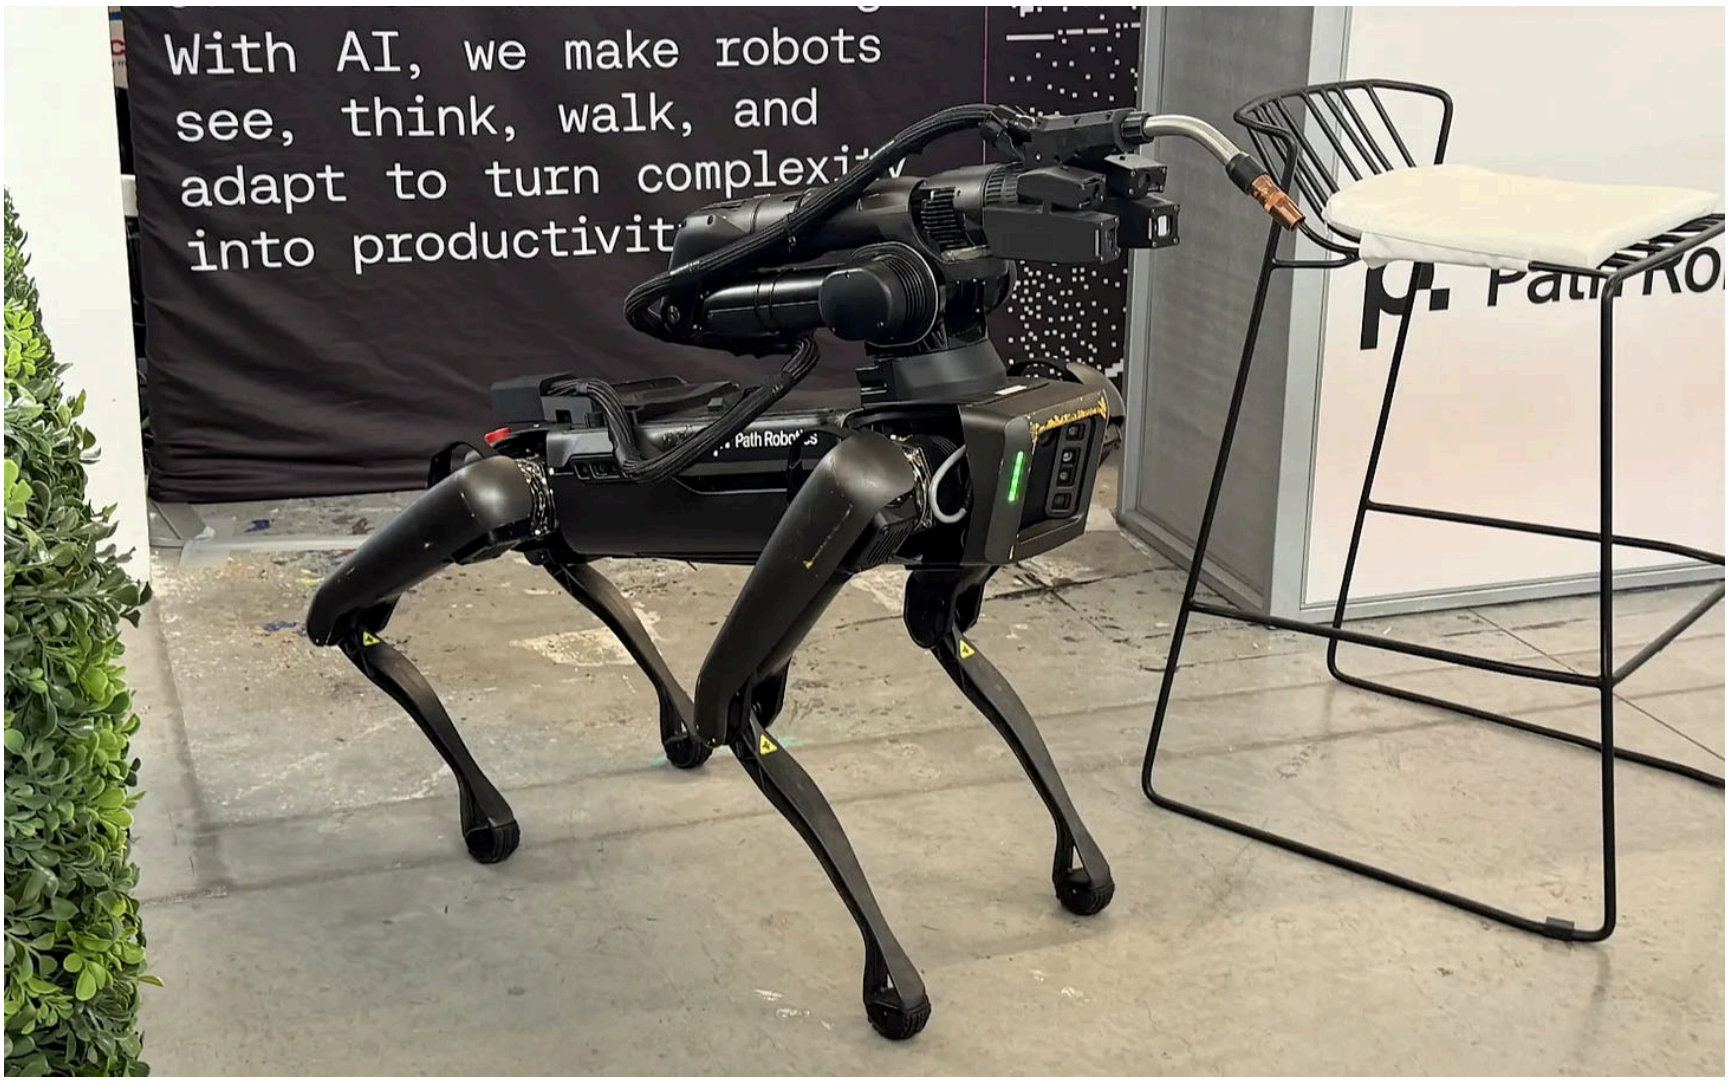


I think that a lot of the earliest deployments will look a lot like this – not fully humanoid but instead, custom brains, custom hardware, and off-the-shelf platforms all mashed together.

我认为，许多最早期的部署看起来都会与这种形态相似——并非完全的人形机器人，而是把定制的大脑、定制的硬件与现成的平台拼合在一起。

One particularly freaky version I saw was this guy — from privately held Path Robotics. Yes, that’s a welding torch attached to a Boston Dynamics quadruped, and yes, these are currently in live deployment.

我见过的一个尤为怪异的版本是这家伙——来自非上市公司 Path Robotics。没错，那是一把焊枪装在波士顿动力（Boston Dynamics）的四足机器人上，而且没错，这些设备目前已投入实际部署。



Finally, listen to commentary from **Teradyne (TER US)**, the owner of Universal Robots, at an investor conference last week. You might recall we pitched Teradyne last year due to a massive rollout of its arms in Amazon warehouses (the company has since confirmed that revenue from this contract "tripled this year, and will triple again next year"). But it seems TER is now even more excited about far more sophisticated levels of automation:

最后，不妨听听 **Teradyne（TER US）**——Universal Robots 的母公司——上周在一场投资者会议上的发言。你可能还记得，我们去年曾因该公司机械臂在亚马逊仓库的大规模部署而**推介过 Teradyne**（公司此后确认，来自该合同的收入"今年增至三倍，明年还将再增至三倍"）。但看起来，TER 如今对远为复杂的自动化层级更加兴奋：

> *"So, there's two end market dynamics and one underlying technology that get us really excited about robotics. So, the two end markets: one is in logistics, and that's where our marquee customer is, is in the e-commerce logistics space [...]*
>
> *The second vertical that we're very excited about for robotics is in electronics manufacturing, and all the way back into the semiconductor value chain. So, we've been calling this robotics-assisted test, and robotics-assisted assembly, and robotics-assisted operations for data centers.* ***And this is the fastest-growing. It's now the largest vertical that we have for our robotics business*** *[...] The common thread between them is Physical AI [...]*
>
Physical AI underpinning our robotics applications is probably the most important trend that's going to drive our business over the next few years."

> *"我们之所以对机器人业务感到非常兴奋，是因为有两大终端市场动向和一项底层技术。先说两个终端市场：一个是物流，我们最重要的客户就在电子商务物流领域 [...]*
>
*第二个让我们对机器人业务非常兴奋的垂直领域是电子制造，并一直向上延伸到半导体价值链。我们把这一块称为机器人辅助测试、机器人辅助组装，以及面向数据中心的机器人辅助运营。****这是增长最快的部分，如今已成为我们机器人业务中规模最大的垂直领域*** *[...] 两者的共同主线是实体 AI（Physical AI）[...]*
>
支撑我们机器人应用的实体 AI，可能是未来几年推动我们业务发展的最重要趋势。"

But what about the version of Physical AI that a lot of folks are selling – mass deployments in factory settings to replace labor?

那么，许多人大力兜售的那种实体 AI——即在工厂环境中大规模部署以取代人力——情况又如何呢？

Humanoids Baby Steps

人形机器人的蹒跚学步

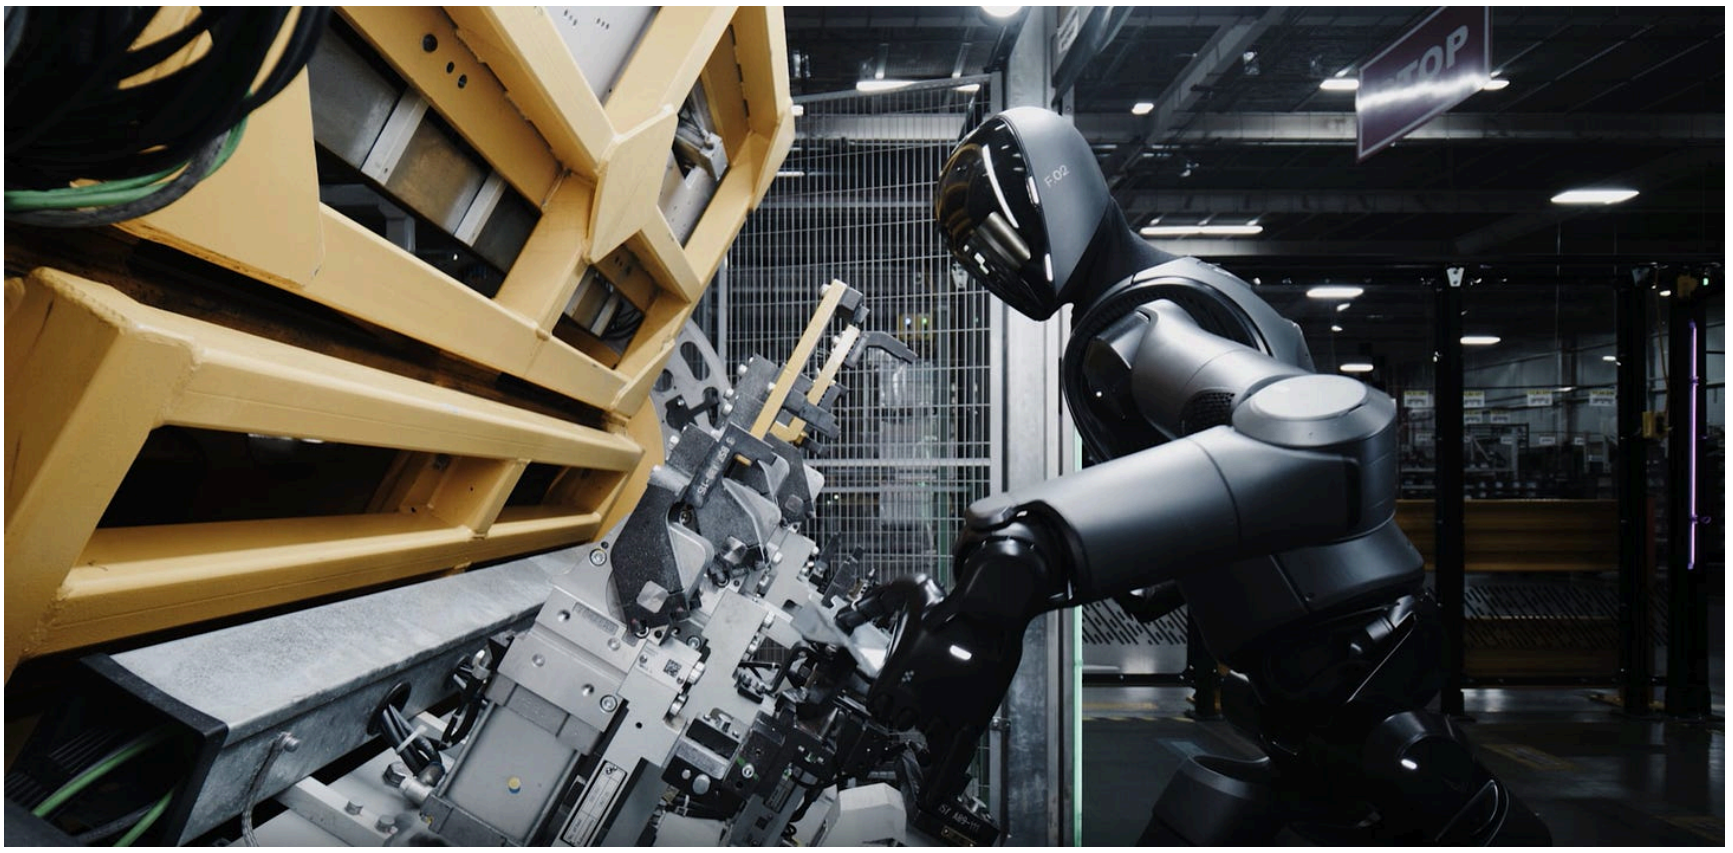
A list of headline announcements might give the impression that humanoids have reached real roll-out phases in industrial settings. Yet for all the talk, there are few ***verifiable*** pilots of humanoids doing "something useful" in factories.

一连串头条公告可能让人以为，人形机器人已在工业场景进入真正的落地部署阶段。然而，尽管讨论声浪不断，能***可验证地***证明人形机器人在工厂里"做点有用之事"的试点案例却寥寥无几。

Figure <> BMW

The most notable and longest standing is Figure AI's deployment in the BMW Spartanburg factory. The partnership dates back to 2024, with the second generation Figure 02 bot pilot beginning in 2025. The Figure 02 was able to perform "pick and place" tasks on a live BMW production line. It picks up a metal part, carries it to a welding machine, places it in the correct position, and repeats.

最引人注目且延续时间最长的是 Figure AI 在宝马斯帕坦堡（Spartanburg）工厂的部署。双方合作可追溯至 2024 年，第二代 Figure 02 机器人的试点于 2025 年启动。Figure 02 能够在宝马实际生产线上执行"拾取与放置"（pick and place）任务：它抓取一个金属部件，搬运至焊接机，放置到正确位置，然后不断重复这一过程。



There was some initial controversy around Figure's claims, but both sides have confirmed the following stats (with Figure claiming >99% success rate):

围绕 Figure 的说法起初存在一些争议，但双方均已确认以下数据（Figure 声称成功率超过 99%）：

- Ran 10-hour shift Monday-Friday
- 周一至周五每天运行 10 小时班次
- +1,250 hours of runtime
- 累计运行时长超过 1,250 小时
- +90,000 parts loaded *(1.2/min)*
- 装载零件超过 90,000 件 *(每分钟 1.2 件)*
- +30,000 X3 vehicles *(0.4/min)*
- 处理超过 30,000 辆 X3 车辆 *(每分钟 0.4 辆)*
- +1.2 million robot steps or +200 miles
- 机器人行走步数超过 120 万步，约合 200 英里

This year, the pilot has now extended to the third-generation Figure 03 and extends pick-and-place sheet metal loading to more complex use cases; parts sorting, and moving carts.

今年，该试点已扩展至第三代 Figure 03 机器人，并将拾取与放置钣金件装载作业延伸至更复杂的应用场景：零件分拣和搬运手推车。

We're not talking about particularly skilled tasks, but the task complexity is growing substantially year-over-year and BMW (who has less incentive to fake it) has expanded these pilots both with Figure in the US and a new pilot with Hexagon in Europe. In announcing this expansion, BMW spoke highly of its first learnings in Spartanburg: ***"The transition from the lab to real production was, for me, faster than we expected,"*** said Michael Nikolaides, Senior VP for BMW's Production Network.

我们所说的并不是特别高技能的任务，但任务复杂度正逐年大幅提升，而且宝马（其造假动机较小）已在两方面扩大这些试点：在美国继续与 Figure 合作，并在欧洲与 Hexagon 启动新的试点。在宣布这一扩展时，宝马高度评价了其在斯帕坦堡获得的初步经验：***"从实验室走向真实生产的过程，在我看来比我们预期的更快，"*** 宝马生产网络高级副总裁 Michael Nikolaides 表示。

Here's a picture I took (from the public access street) at Figure's HQ.

下面是我在 Figure 总部拍摄的一张照片（从公共可通行街道拍摄）。

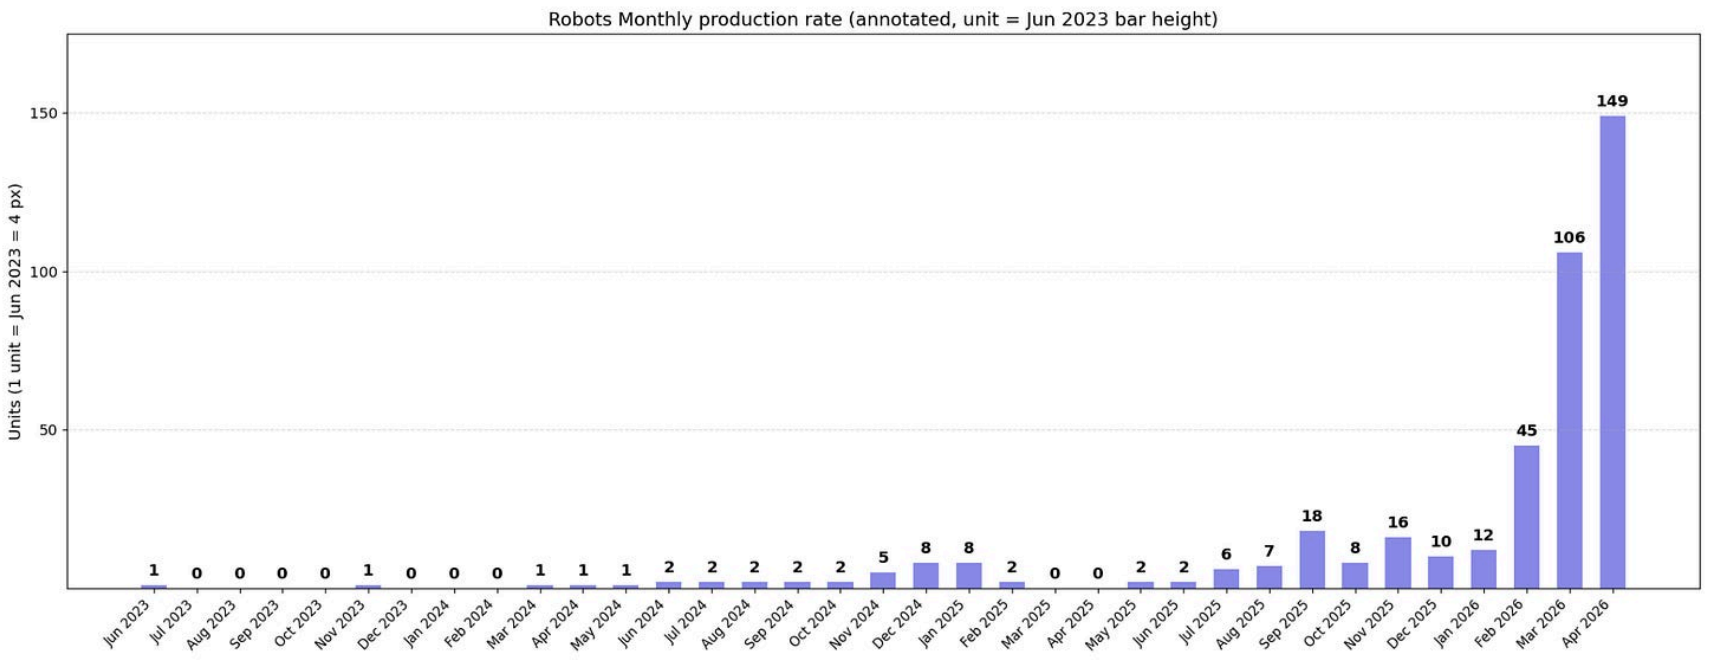


Outside on the walkway are maybe a dozen Figure 03s ambling about autonomously **(***Recess).* This is both legitimate training as well as a bit of flex – unsupervised, highly visible validation that these robots can understand and navigate the terrain and stairs, stay limited to sidewalks, and safely interact with *whoever* chooses to walk on this public sidewalk (like me with a camera).

外面的人行道上大约有十几台 Figure 03 在自主漫步**（*课间休息*）。** 这既是实打实的训练，也带点展示实力的意味——在无人监督、公开可见的环境下验证这些机器人能够理解并通行地形与楼梯，始终停留在人行道范围内，并与*任何*选择走在这条公共人行道上的人安全互动（比如拿着相机的我）。

What happens behind the reflective glass is under much tighter wraps, but we can surmise based on what they've posted publicly. With a substantial ramp in production beginning earlier this year, there are now hundreds of 03s in existence – including those shown above.

反光玻璃背后发生的事情则要保密得多，但我们可以根据他们公开发布的内容加以推断。随着今年早些时候开始的大幅量产爬坡，如今已有数百台 03 存在于世——包括上面展示的那些。



These robots are the first class at Figure's school. A few are out on their "internships" in places like Spartanburg. The majority are probably in classes on campus like "gym" and "home economics", collecting mounds of data from sensors on hundreds of bots. All is sent back up to improve the hive-minded model.

这些机器人是 Figure 这所"学校"的第一届学员。 少数机器人已经前往 Spartanburg 之类的地方开始"实习"。大多数机器人可能还在"校园"里上课，比如"体育课"和"家政课"，从数百台机器人身上的传感器收集海量数据。所有数据都会回传，用于改进这套蜂群式的共享模型。

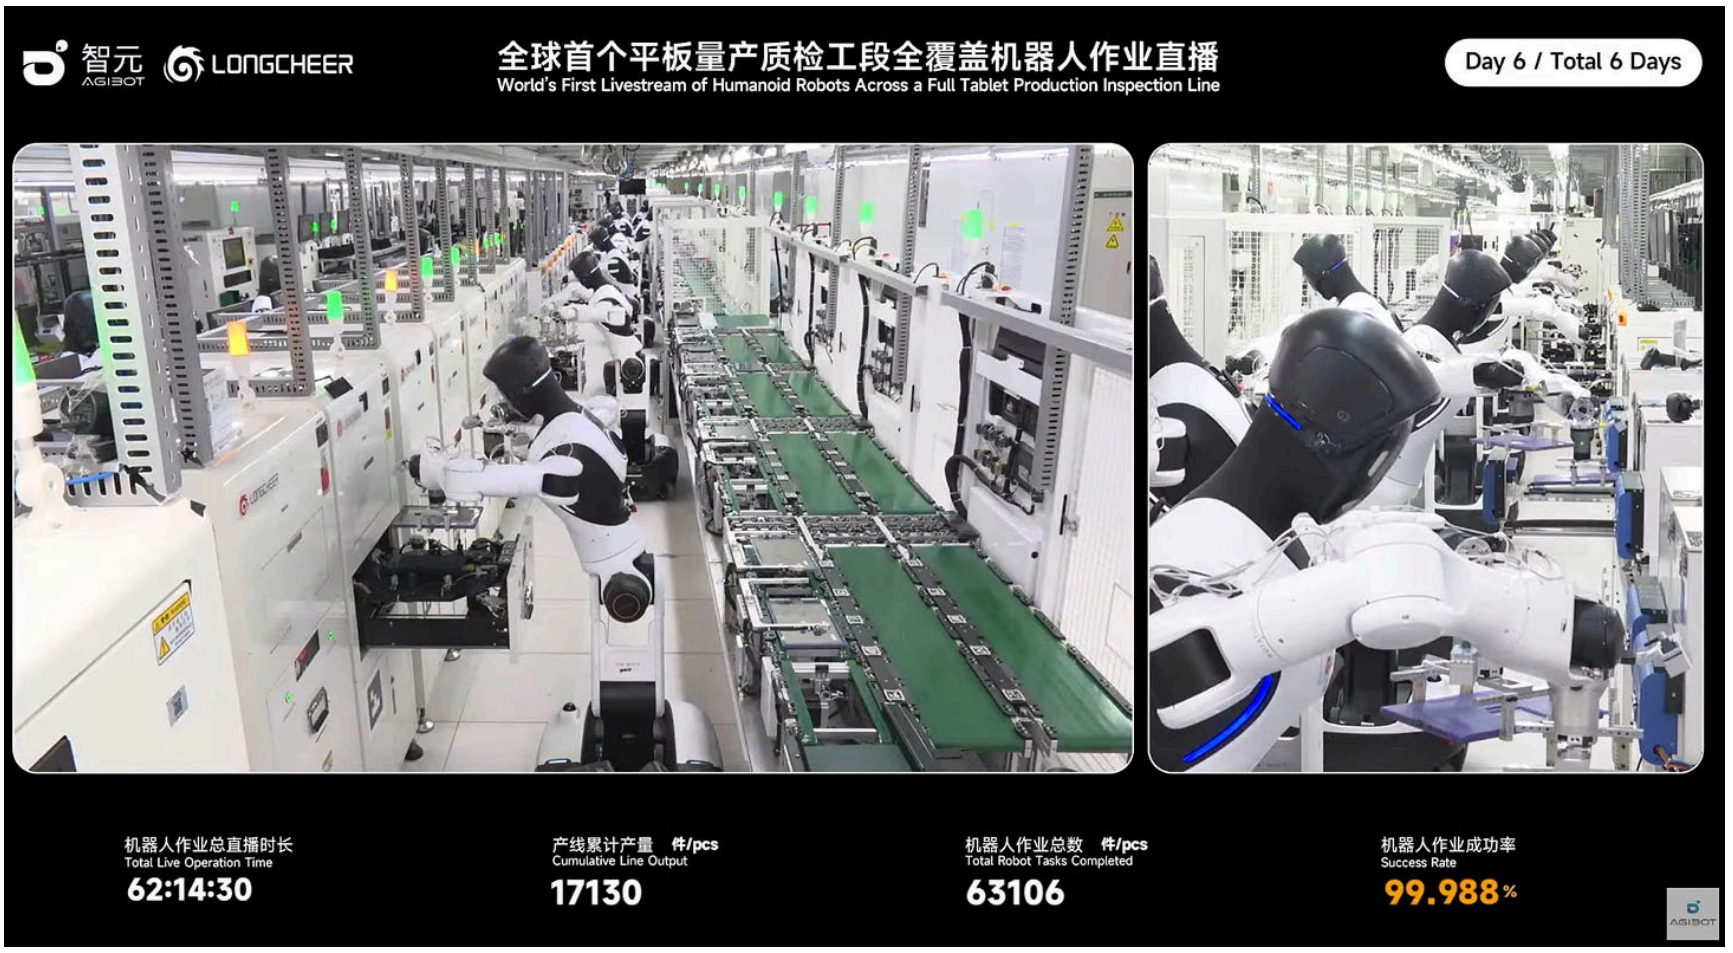
Given that the first BMW demo began over a year ago with its prior 02 model, it makes the results more impressive. Not only is there massive hardware improvement, but now there is manufacturing scale to allow the data flywheel to really start spinning.

考虑到首个 BMW 演示早在一年多前就已开始，且当时用的还是上一代 02 型机器人，如今的结果就更令人印象深刻。不仅硬件有了大幅改进，制造规模也已到位，让数据飞轮得以真正开始转动。

Agibot <> Longcheer

In April 2026, Chinese manufacturer Agibot announced it had integrated its G2 wheeled-base humanoid in the tablet production line of Longcheer Technologies, which does contract manufacturing for primarily Chinese consumer electronics companies like Lenovo and Xiaomi.

2026 年 4 月，中国制造商 Agibot 宣布，其 G2 轮式底盘人形机器人已进入龙旗科技（Longcheer Technologies）的平板电脑生产线。龙旗科技是一家代工制造商，客户主要是联想、小米等中国消费电子企业。



In June, a six-day livestream from the factory provided some more visible color. The robots were doing similar "pick and place" tasks, moving tablets from one conveyor into a testing machine, and back.

六月，工厂一场为期六天的直播提供了更多直观的细节。机器人正在执行类似的"拾取与放置"作业：将平板从一条传送带搬运至测试机中，然后再搬回。

It's worth noting that "end effectors" are clearly custom made for this specific job while its wheeled base proves its elegance as the bots glide and shift back and forth. Each factor dramatically simplifies the task vs a pure humanoid… but that's not necessarily a bad thing.

值得注意的是，这些"末端执行器"显然是为这项特定作业定制的，而其轮式底盘也在机器人来回滑行移动时展现出了优雅性。与纯粹的人形机器人相比，上述每一个因素都大幅简化了任务……但这未必是件坏事。

If "humanoid" is a spectrum and incremental deployments are necessary precursors for wide-scale rollouts, then maybe some purists could take a lesson.

如果"人形"本身是一个谱系，而渐进式部署是大规模推广的必要前提，那么某些纯粹主义者或许应该从中吸取点教训。

Agility Digit <> Schaeffler, GXO, Toyota, MeLi

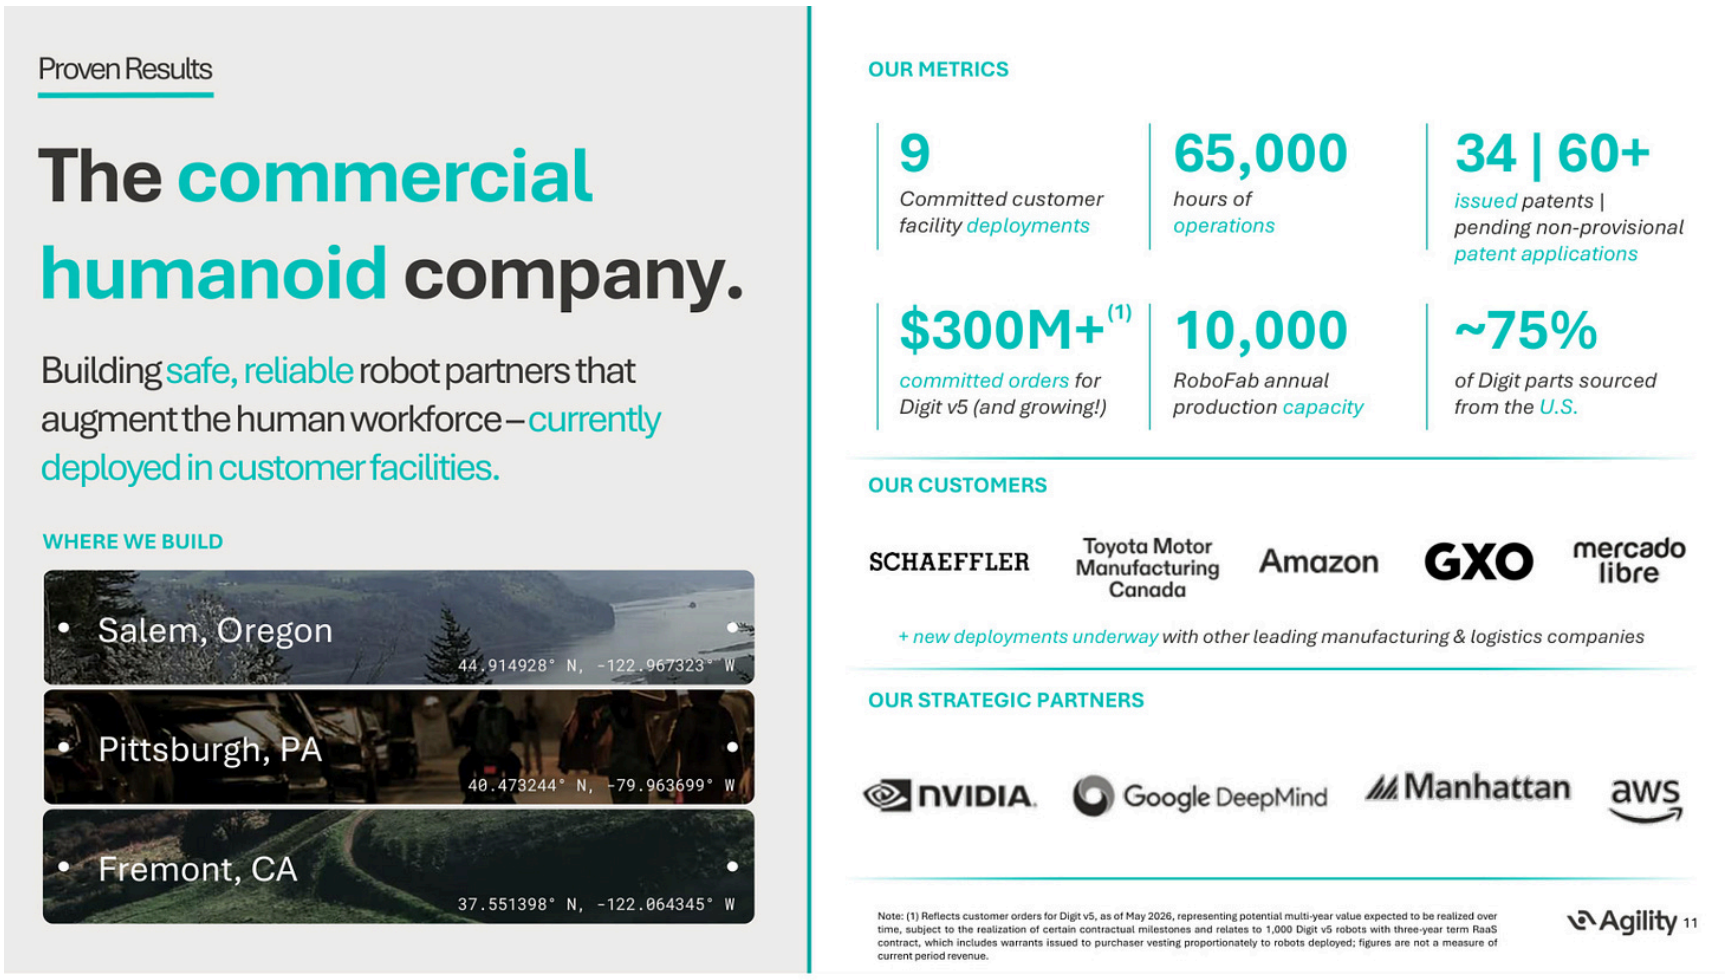
Agility Digit <> 舍弗勒、GXO、丰田、MeLi

Agility (under agreement to be acquired by SPAC Churchill Capital Corp XI – **CCXI US**) describes "active" humanoid deployments with Schaeffler, GXO, Toyota Motor Manufacturing Canada, and Mercado Libre.

Agility（已签署协议，将被 SPAC Churchill Capital Corp XI——**CCXI US**——收购）描述了其与舍弗勒、GXO、丰田汽车加拿大制造公司以及 Mercado Libre 展开的"活跃"人形机器人部署。

In its S-4 filed the other week, it describes *"nine customer facilities and has accumulated more than 65,000 hours of operation in real-world customer environments"* but the most we get on what it's actually doing is "moving totes" – 25,000 at Schaeffler and 100,000 at GXO with ~98% operational accuracy.

在其于前几周提交的 S-4 文件中，公司描述道，*"九个客户场所，并已在真实客户环境中累计运行超过 65,000 小时"*，但至于其实际在做什么，我们所能了解到的最多也就是"搬运周转箱"——在舍弗勒搬运了 25,000 个，在 GXO 搬运了 100,000 个，作业准确率约 98%。



Given its longer history (2015 founding with initial factory pilots in 2024) and list of deep-pocketed backers, the lack of any real specificity and demonstrations doesn't bode well. Notably Schaeffler is both an investor in Agility and is touting itself as a robotics "supply chain beneficiary", and so it's dually incentivized to show progress.

考虑到其历史较长（2015 年创立，2024 年才启动首批工厂试点），且投资方阵容财力雄厚，却始终拿不出真正具体的细节和实际演示，这并非好兆头。值得注意的是，舍弗勒（Schaeffler）既是 Agility 的投资方，又自称是机器人领域的"供应链受益者"，因此在展示进展方面有着双重动机。

Meanwhile, the CEO of one of its big logo clients, GXO, went on the record in August saying how better hands (dexterity and manipulation) is a key to advance humanoid development – Digit 4 uses grippers.

与此同时，其知名大客户之一 GXO 的 CEO 曾于 8 月公开表态，称更好的手部能力（灵巧性与操作能力）是推进人形机器人发展的关键——而 Digit 4 使用的仍是夹爪。



Given its upcoming acquisition, we have more detail on Agility than most other private companies. Specifically, we see the company made $1.8 million in sales in 2025 on ~$142 million in operating expenses. Its claim of "$300 million in committed orders" comes with the very heavy asterisks that they are subject to contractual milestones and also include giving equity warrants to the purchasers. This is a company that probably couldn't survive long without the proceeds from this deal, and the fact that they had resorted to a SPAC for financing might suggest that more attractive private market funding wasn't available.

鉴于即将被收购，我们掌握的 Agility 细节比大多数其他私人公司都要多。具体来看，该公司 2025 年的销售额为 180 万美元，而营业费用约为 1.42 亿美元。其宣称的"3 亿美元承诺订单"带有非常沉重的星号注脚：这些订单须以达成合同里程碑为前提，并且包括向买方授予股权认股权证。若没有这笔交易带来的资金，这家公司很可能撑不了多久；而他们不得不借助 SPAC 融资这一点，或许表明更具吸引力的私募市场融资并不可得。

To be fair, Digit's biggest problem may be that it was too early. It got stuck in an early form factor before the hardware and autonomous advancements really took off, and it's been drawing investors long enough to the point of exhaustion. Maybe CCXI investors will get to see the fruit.

公平地说，Digit 最大的问题可能在于它出现得太早。在硬件和自主能力的进步真正起飞之前，它就被早期的形态设计困住了；而它吸引投资者的时间又拖得太长，长到令投资者疲惫不堪。也许 CCXI 的投资者终能看到成果。

The Field

其他玩家

Then comes a much longer list of headlines with less publicly available to verify, but all with a similar rhyme.

接下来是一份长得多的新闻标题名单，这些信息可供公开核实的程度较低，但套路如出一辙。

Generally these are all described as small scale pilot and testing, often with equity investors and related parties. The tasks are somewhere between the precision pick-and-place of the Figure/Agibot trials, and the less distinctive Agility tote moving. In the case of Chinese companies, the deployment into state-owned or strategically important industries can often be viewed as “revenue subsidies” with a general nudge from government officials.

这些项目通常都被描述为小规模试点和测试，参与方往往是股权投资者或关联方。其任务介于 Figure/Agibot 试验中的精密抓取放置，与 Agility 那种辨识度较低的料箱搬运之间。就中国公司而言，在国有企业或战略重要行业中的部署，在政府官员的普遍推动之下，往往可被视为一种“收入补贴”。

Robot Manufacturer / Platform	Factory Counterparties and Deployment Details	Best Evidence
Galbot — S1	• **CATL, undisclosed battery factory** — Claimed regular operation. Material handling and picking in battery-module and pack production. S1 is a heavy-duty wheeled humanoid.	CATL reports an eight-hour runtime. Fleet size, throughput, cumulative operating hours and intervention rates are undisclosed.
UBTECH — Walker family	• **SANY Renewable Energy, China** — Claimed active factory operation. Walker S2 removes and sorts bolt sleeves, transfers trays and supports assembly. • **Zeekr, Ningbo** — Multi-robot training/test. Tote handling, box movement and coordinated sorting. A three-week single-robot exercise expanded into an approximately 20-robot test. • **Foxconn, Shenzhen and Zhengzhou** — Training and proof of concept. Handling, sorting, gluing, inspection and parts feeding • **BYD —** Factory pilot/training. Material handling.Dongfeng Liuzhou — Factory pilot/training. Inspection, fluid filling and emblem placement. • **FAW-Volkswagen, Qingdao —** Factory pilot/training. Quality inspection. • **Airbus, undisclosed site —** Concept testing. Airbus acquired at least one Walker S2 for aerospace-manufacturing evaluation. • **Texas Instruments, undisclosed line —** Reported pilot. Walker S2 reportedly purchased for semiconductor-line deployment and debugging, without detailed TI confirmation.	SANY provides no fleet size, runtime or throughput. Zeekr reportedly tested approximately 20 robots, but not as sustained production labor. Other counterparties generally appear to have one or a few robots, with no meaningful operating statistics.
Wandercraft — Calvin	• **Renault, Douai, France —** Factory pilot. Carries two tires at once—approximately 30 kilograms—in the wheel workshop.	Very small number presently being tested. Renault targets approximately 10 robots by year-end 2026 and 350 across French and Spanish factories by year-end 2027.
Apptronik — Apollo	• **Mercedes-Benz, Berlin and Kecskemét —** On-site pilot and data collection. Component and kit delivery, tote movement and preliminary quality inspection. Activity has included substantial teleoperation while progressing toward autonomy. • **Jabil, undisclosed facility —** Factory acceptance testing. Inspection, sorting, kitting, line-side delivery, fixture placement and subassembly before newly manufactured robots are delivered to customers.	A few units at Mercedes; no disclosed production hours, autonomous share or throughput. Jabil fleet size and performance statistics are undisclosed.
Hexagon Robotics — AEON	• **BMW, Leipzig, Germany —** Factory pilot. High-voltage battery assembly, component production and line-side material delivery. Operational debut in December 2025, followed by testing and planned pilot integration in summer 2026. • **Schaeffler, undisclosed plant —** Completed pilot with future deployment planned. Loads, unloads and inspects parts across a multi-machine station.	BMW: one wheeled humanoid disclosed. Schaeffler describes a successful 2025 pilot, with additional deployments intended from late 2026; no hours or output provided.
Boston Dynamics — Atlas	• **Hyundai, Metaplant America, Georgia —** Completed on-site prototype pilot. Picked automotive parts from supplier racks and sequenced them into another rack.	Pilot completed in October 2025. Fleet size and operating results undisclosed. Broad plant deployment is planned for 2028, not currently underway.
Tesla — Optimus	• **Tesla, undisclosed factory —** Internal development deployment. Tesla has shown battery-cell sorting and stated that Optimus robots performed autonomous factory tasks.	Tesla disclosed two factory robots, but provided no named site, cumulative hours, output, autonomous share or intervention rate.
Capgemini / Orano — Hoxo	• **Orano Melox, Gard, France —** Training-facility pilot. Navigation and technical gestures intended eventually for hazardous nuclear interventions.	Four-month pilot at Melox’s vocational training facility. No deployment inside radioactive production areas.

The longest listed comes from Shenzhen-based **UBTech (9880 HK)** which I visited in person last year, and which bears some resemblance to Digit. UBtech was an early player in the field that has largely been lapped by more recent upstarts in China.

上市时间最长的一家来自深圳的**优必选（9880 HK）**，我去年亲自到访过这家公司，其产品与 Digit 有几分相似。优必选是该领域的早期参与者，但在很大程度上已被中国后来的新兴企业超越。

What should we take away from this? Obviously, it’s not super supportive of the idea of mass production and commercialization of thousands of humanoids for industrial purposes by a wide range of OEMs in the next year or so. But that’s not the only takeaway either.

我们应该从中得到什么启示？显然，这并不太支持这样一种观点：即在未来一年左右，众多 OEM 厂商将大规模生产并商业化数千台用于工业用途的人形机器人。但这也不是唯一的结论。

Even at a small scale, it shows ***differentiation.*** The difference between Figure and Agility is a ~15x gap in valuation.

即便在很小的规模下，这也体现了***差异化***。Figure 与 Agility 之间的差距，就体现为约 15 倍的估值差异。

4) Four Predictions for Robotics

4）对机器人行业的四项预测

Up until this point I’ve given you what I think is an impartial description of the present. What about the future? Here are my four predictions…

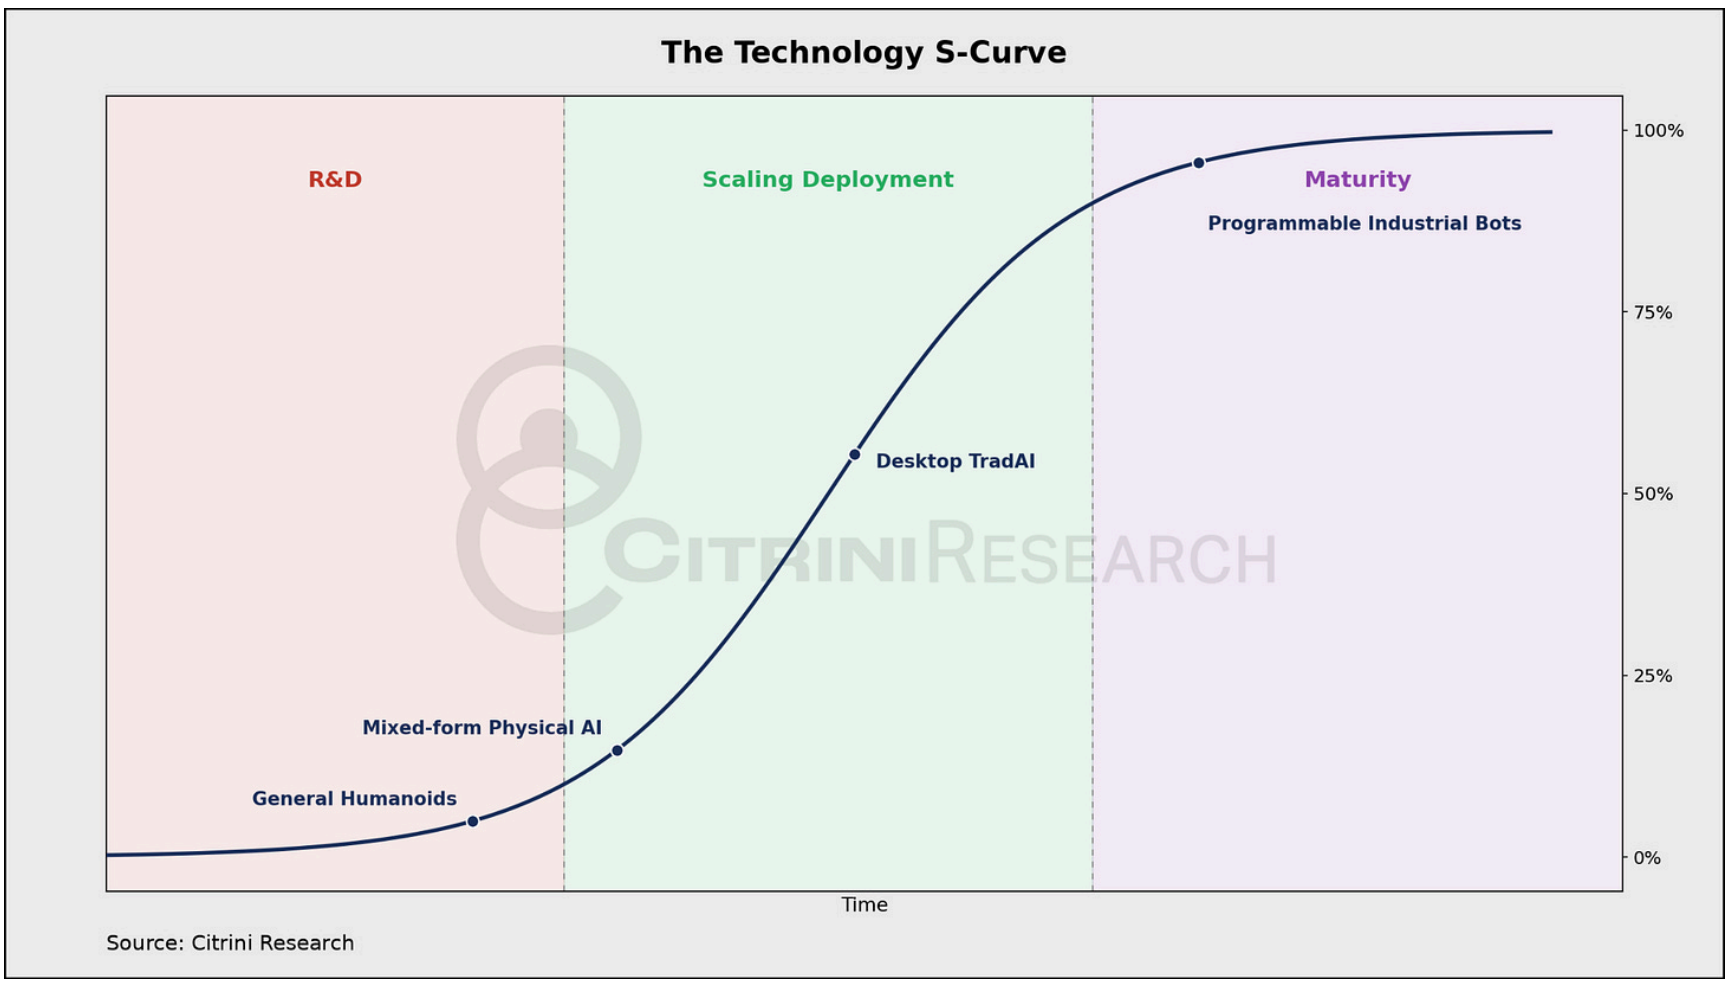
到目前为止，我为大家提供的，是我认为对现状的客观描述。那么未来呢？以下是我的四项预测……

Mind the Development Curve

留意发展曲线

The biggest thing that will trip people up is understanding the development curve.

最容易让人栽跟头的，是对发展曲线的理解。



On one side, there will be haters who point to the lack of large-scale *humanoid* deployments as proof that “robotics” is still a decade away.

***一方面，**会有唱衰者以缺乏大规模人形机器人*部署为由，断言“机器人技术”仍是十年之后的事。

The other side will likely overindex to the *current* stage of the AI/Datacenter and try to replicate the “Bottleneck” approach before they exist.

另一方面的人则很可能过度锚定 AI/数据中心的*当前*阶段，在这些瓶颈真正出现之前，就试图照搬“瓶颈”这套打法。

The former are likely to miss the trend. The latter are likely to overindex into broad lists of large industrial companies that won’t see needle-moving impact in the near term. Remember, it took years and literally trillions of investment after TradAI’s ChatGPT moment before commodity components like memory and capacitors started to catch a bid. Unlike the AI trade where incremental inference appears to be very high margin (and therefore accepting of supply cost increases), the economic proposition of robotics heavily relies on ***reducing costs*** from current levels.

前者很可能错过这轮趋势。后者则很可能把注意力过度铺在一长串大型工业企业名单上，而这些公司在短期内不会受到实质性的业绩拉动。别忘了，在传统 AI 迎来 ChatGPT 时刻之后，市场又花了数年时间、投入了真正数以万亿计的资金，存储器和电容器这类大宗元器件才开始获得买盘。与 AI 交易不同——AI 的增量推理利润率看起来极高（因此能够承受供应成本上升）——机器人技术的经济可行性在很大程度上依赖于把成本从当前水平***降下来***。

The physical nature of robotics is also a critical difference to recognize. ChatGPT could scale to a billion users in part because of its near frictionless distribution of the internet (subject to datacenter compute). Manufacturing robots is a whole different beast.

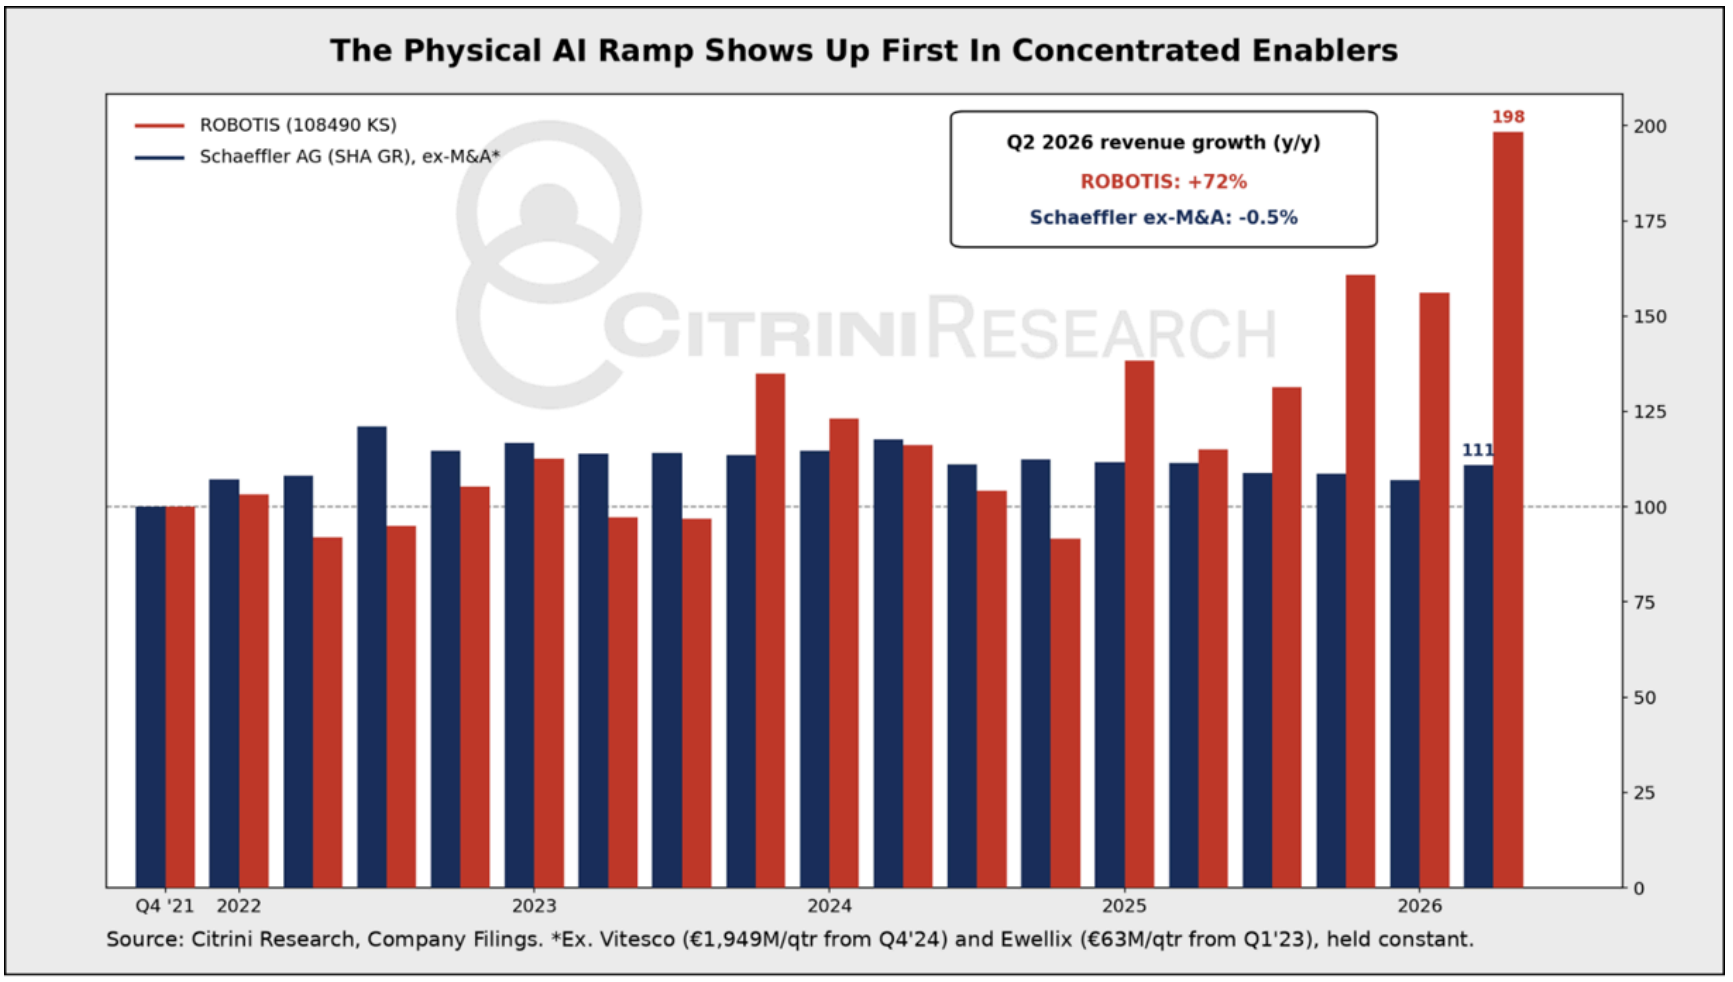
机器人技术的物理属性也是必须认清的一个关键差异。ChatGPT 之所以能扩展到十亿用户，部分原因在于互联网分发几乎不存在摩擦（当然仍受制于数据中心算力）。而制造机器人则完全是另一回事。

Even as volume and deployment grow, the large majority of near-term spend will continue going towards R&D enablers rather than mass-scale production of commodity parts. In the supply chain, highly concentrated niche suppliers are far more likely to see upward revisions than large industrials.

即便产量和部署规模不断增长，近期支出的大部分仍将继续流向研发赋能环节，而非大宗零部件的大规模量产。在供应链中，高度集中的利基供应商获得预期上修的可能性远远大于大型工业企业。

Compare **Robotis (108490 KS)**, which sells a wide range of relatively cheap R&D hardware like table-mounted manipulators, to **Schaeffler (SHA0 GR)**,an industrial automotive conglomerate which – in theory – supplies many components and has heavily leaned into the robotics supply-chain narrative.

不妨对比一下 **Robotis (108490 KS)** 与 **Schaeffler (SHA0 GR)**：前者销售台式机械臂等一系列相对廉价的研发硬件；后者是一家工业汽车集团，理论上供应众多零部件，并已大举押注机器人供应链叙事。

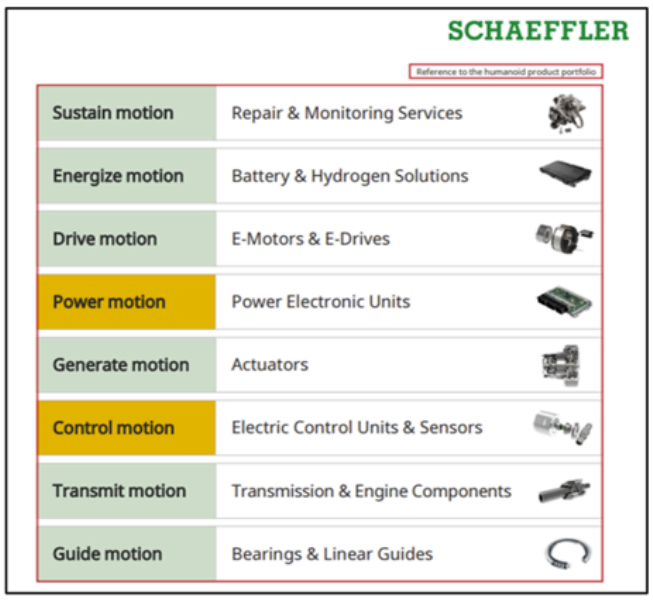




← **R&D Enabler**

VS

Supply Chain Components →



While the two companies both had a flat topline from 2021-2024 (indexed 4Q21=100), Robotis has seen a surge in revenue growth since 2025 along with the relative surge in AI-enabled robotics advancements. Schaeffler's been flat. Its €350 million *backlog* is the largest publicly disclosed that I've found so far but represents just 1.5% of annual sales.

虽然这两家公司2021至2024年的营收都基本持平（以2021年四季度=100为基准进行指数化），但自2025年以来，随着AI赋能机器人技术进展的整体跃升，Robotis的营收增长显著提速；而舍弗勒则一直持平。其3.5亿欧元的*在手订单*（backlog）是我迄今找到的最大一笔公开披露，但也仅占年销售额的1.5%。

But what about other less obvious R&D enablers?

那么，其他不那么显眼的研发赋能者呢？

My intuition is to focus less on the existing actuator and bearings suppliers and more on the wider ecosystem that a robotics revolution requires.

我的直觉是：少关注现有的执行器与轴承供应商，多关注一场机器人革命所需的更广阔生态。

(The "actuator trade" is probably the point I feel strongest about. Unitree, Figure, Tesla, and anyone else seriously contesting for scale and moat are designing and building their own actuators, while intentionally spread across subcomponent suppliers. This is critical, proprietary IP, and they are going to keep it. It will require far higher volumes for these subcomponents to be in shortage.)

（"执行器交易"（actuator trade）大概是我信念最强烈的一个观点。宇树（Unitree）、Figure、特斯拉，以及任何认真争夺规模与护城河的玩家，都在自行设计并制造各自的执行器，同时刻意将子部件分散到多家供应商。这是关键的专有知识产权，他们一定会牢牢握在手里。要让这些子部件出现短缺，所需的出货量还远高于当前水平。）

China vs. US: The Hardware Gap

中国 vs. 美国：硬件差距

The cheapest argument in robotics is to say China is ahead in robotics. The second cheapest but more sophisticated argument is that *ackshually* the US is ahead where it counts – the brain!!

机器人领域最廉价的论点，是宣称中国在机器人方面领先。第二廉价、但显得更精巧的论点则是：*其实吧*，美国在真正要紧的地方领先——大脑！！

Coming out of my China trip last year, I concluded that the Chinese strategy, it seemed from people like Unitree, was to push hard on the hardware side, while allowing researchers all around the world to invest a lot of time and money on autonomy. No one there was even making the claim that Chinese humanoids were *the best* – Figure and Tesla's Optimus were constantly cited as the global leaders.

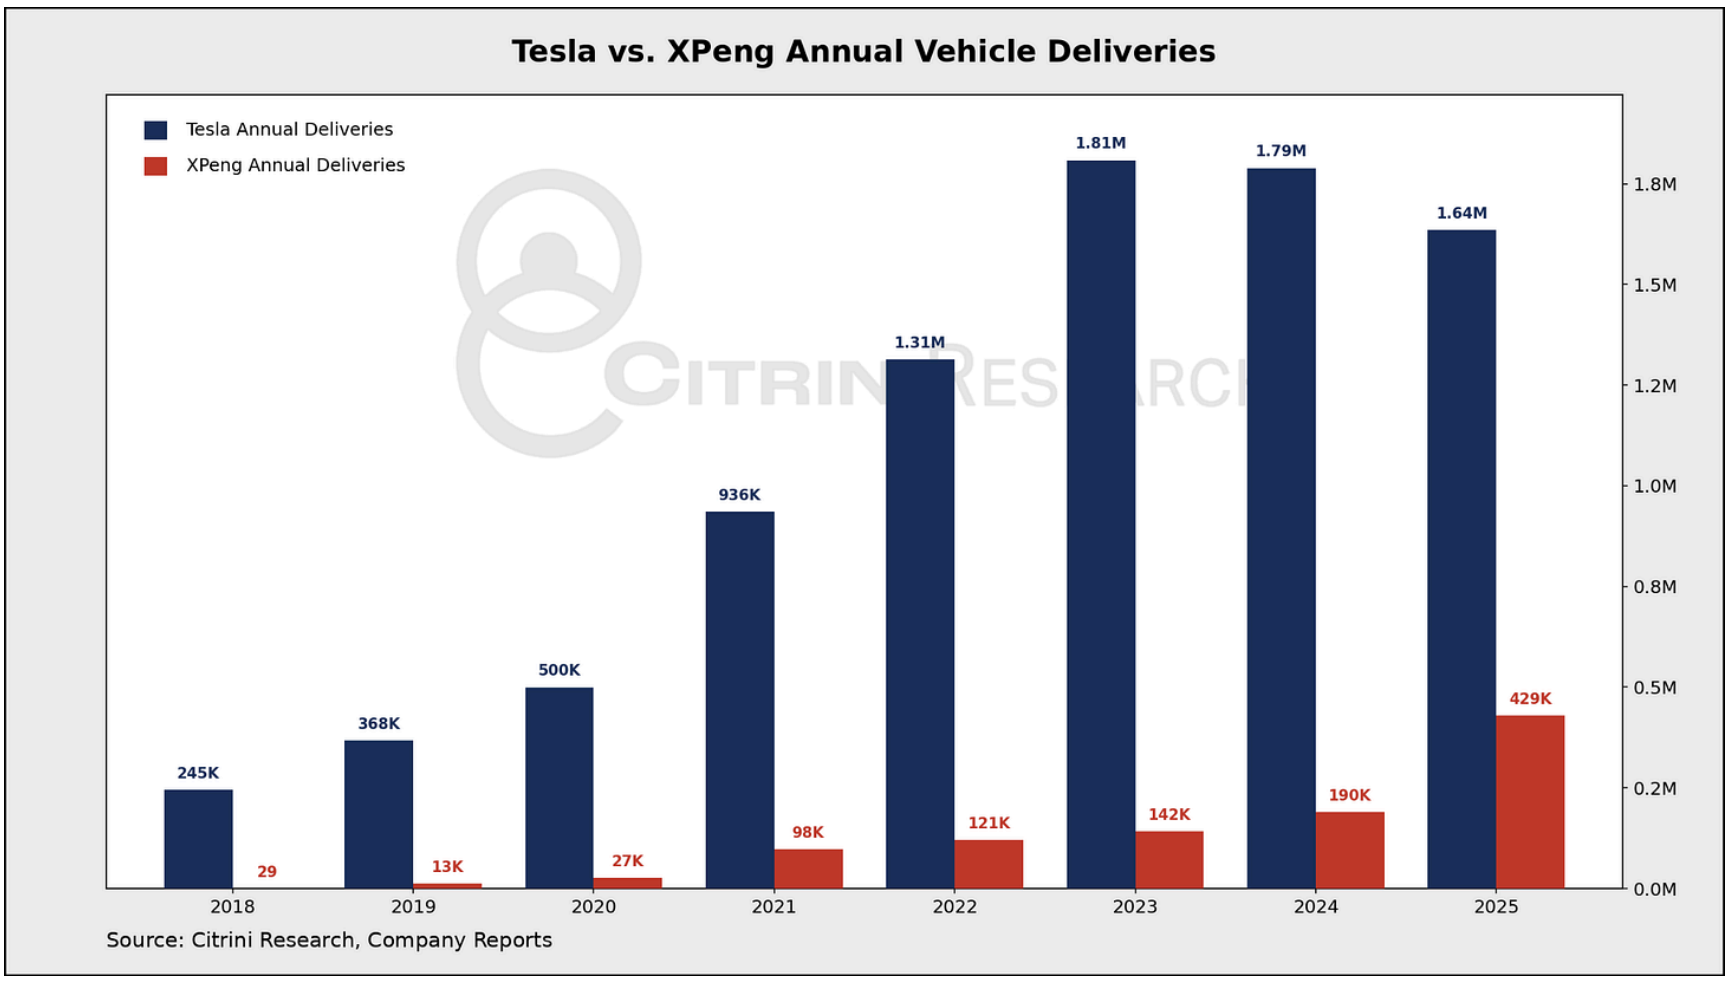
去年结束中国之行后，我得出的结论是：从宇树这类公司来看，中国的策略似乎是在硬件一侧全力推进，同时放手让世界各地的研究者在自主性上投入大量时间与资金。在当地，甚至没有人声称中国人形机器人是*最好的*——Figure和特斯拉的Optimus一再被引为全球领先者。

Being a “step behind” wasn’t a bad thing per se, it was a strategy. When the breakthroughs come, they will be fairly quickly replicated at which point competitive edge shifts towards who can build fast and cheap, qualities for which there is little comparison between the two countries.

“落后一步”本身并非坏事，而是一种策略。一旦突破出现，这些成果会相当快地被复制，届时竞争优势将转向那些能造得又快又便宜的玩家——而这两项能力在两国之间几乎没有可比性。

It’s running the EV playbook, with the literal same players. Tesla popularized the modern EV and Chinese companies like XPeng replicated… the results look like this.

它正在照搬电动汽车（EV）的剧本，而且玩家几乎一模一样。特斯拉让现代电动汽车普及开来，小鹏等中国公司随后加以复刻……结果就是下面这样。



Yet, coming out of this trip, maybe the cheapest take is actually the most right.

然而，此行归来，或许最不费脑的观点反而是最正确的。

At a highly attended *robotics* conference in San Francisco there were surprisingly few… robots.

在旧金山一场出席率很高的*机器*人大会上，令人惊讶的是……机器人少得可怜。

Egocentric data capture? *Sure!* Database infrastructure? *You bet!*

自我视角数据采集？*当然有！*数据库基础设施？*那更少不了！*

As far as physical robots, there were a few arms for tele-op demonstrations… motion capture gloves paired with Chinese hands. The only full humanoids were a couple of Unitree’s G1s staged for a remote controlled boxing match.

至于实体机器人，只有几台用于遥操作演示的机械臂……还有动作捕捉手套搭配中国产的灵巧手。唯一完整的仿人机器人是两台宇树 G1，被摆在那里用于一场遥控拳击表演。

Can you do ***“Physical AI”*** without the hardware? I can’t shake the feeling that after decades where “US tech” = “screens”, there’s more than just strategic thinking at play.

没有硬件，能做***“具身智能”（Physical AI）***吗？我心里始终挥之不去的一种感觉是：在“美国科技”=“屏幕”的几十年之后，这里起作用的恐怕不只是战略层面的考量。

I wager that ***learning by doing*** will increasingly be seen as the gold standard and that real world performance will increasingly matter more than research benchmarks. This means… building robots. The problem is that the US’ manufacturing muscle and ecosystem have decades of atrophy, especially in this sort of small-scale dynamic and iterative manufacturing. Rebuilding these skills presents both a challenge and opportunity for those enabling it.

我判断，***干中学***将日益被视为黄金标准，真实世界中的表现也将日益比研究基准更重要。这意味着……要造机器人。问题在于，美国的制造实力与生态系统已经萎缩了几十年，尤其是在这类小规模、动态且快速迭代的制造上。重建这些能力，对相关的赋能者而言既是挑战，也是机会。

Meanwhile the option of relying on Chinese hardware is increasingly off the table. So far, Unitrees have been the ubiquitous choice for an off-the-shelf bot for research, entertainment, etc. This is going to change.

与此同时，依赖中国硬件这个选项正在逐步被排除。到目前为止，宇树一直是研究、娱乐等场景中无处不在的现成机器人之选。这种局面将会改变。

Recent FCC actions banning Chinese imports have some pretty big gaps – it’s not a cold-turkey quit. But the writing is on the wall. No strategically minded US company is going to build a hardware supply chain that depends on China. Indeed the purpose of these FCC actions is to jumpstart this transition (in almost an upside down mirror of the US export controls on NVDA chips).

美国联邦通信委员会（FCC）近期禁止中国进口的措施仍存在相当大的缺口——并不是一夜之间的彻底叫停。但大势已经写在墙上。没有任何一家有战略头脑的美国公司，会去建设一条依赖中国的硬件供应链。事实上，这些 FCC 措施的目的正是为这一转型点火启动（几乎是美国对 NVDA 芯片出口管制的一幅上下颠倒的镜像）。

The best counterargument? One of the two leading US players happens to be the country's best manufacturing outfit (Tesla). The other is hyperfocused on vertical integration and is quickly ramping production (Figure).

最好的反方观点是什么？美国两家领先企业中，其中一家（特斯拉）恰好是国内最出色的制造企业；另一家则高度专注于垂直整合，并正在快速提升产量（Figure）。

Humanoids: The Big Swing

人形机器人：全力一搏

While we've shown early deployments are more likely to take non-humanoid form factors and often with mixed hardware, I would not at all count out the humanoid developers.

尽管我们已经展示过，早期部署更可能采用非人形的硬件形态，且常以混合硬件的形式落地，但我丝毫不会看轻人形机器人开发商。

They are going for the moonshot. And there are real reasons to think they can achieve it.

他们瞄准的是登月式的豪赌。 而且有切实的理由相信他们能够实现。

Market – The entire argument for humanoid is generalizability and that's what gives it the most extreme TAM. If there's going to be a trillion-dollar robotics company in 2030, I think it's most likely to be a humanoid. Other companies might be very successful in their own right, but limited in their ultimate markets for moonshot potential.

市场—— 人形机器人的全部论据在于通用性，而这正是其总可寻址市场（TAM）最为巨大的原因。如果2030年会出现一家万亿美元级的机器人公司，我认为最有可能的就是人形机器人公司。其他公司或许凭借自身也能取得成功，但终极市场的空间限制了它们登月式的潜力。

For reference, the largest robotics automation giant today, **ABB Ltd (ABBN SW),** trades at ~$170 billion.

作为参考，当今最大的机器人自动化巨头**ABB Ltd（ABBN SW）**市值约为1700亿美元。

Vertical Integration – Combining proprietary hardware, data, software, and body-specific training, presents a far more defensible moat. The leading humanoid OEMs like Unitree, Galbot, Figure, Tesla aren't buying off the shelf actuators. They are custom designing nearly every component, bringing as much value added manufacturing in house, and spreading the remaining purchases across a wide range of suppliers (to maintain both leverage and IP trade secrets).

垂直整合—— 将专有硬件、数据、软件以及针对机身的专门训练结合起来，能够形成可防御性强得多的护城河。Unitree、Galbot、Figure、特斯拉等人形机器人领军OEM并不会采购现成的执行器。他们几乎自行设计每一个部件，尽可能多地把增值制造环节收归内部，并将其余采购分散到众多供应商（以同时保持议价能力并保护知识产权与商业机密）。

This level of integration and eventual large-production is likely to lead to the most complete, reliable and cost competitive products.

这种程度的整合，加上最终的大规模量产，很可能造就最完整、可靠且具成本竞争力的产品。

Funding – The leading players have the funding and support to spend billions on R&D, before showing meaningful revenue. They have the resources to lag in early deployments while focusing on a longer term goal that may eventually leapfrog the early runners.

资金—— 头部玩家拥有资金与支持，能够在形成可观收入之前投入数十亿美元进行研发。他们有足够的资源承受早期部署上的落后，转而专注于更长期的目标，并最终可能实现对先行者的反超。

Progress – Maybe the most important factor is that I think humanoid abilities are going to progress much faster than many people imagine.

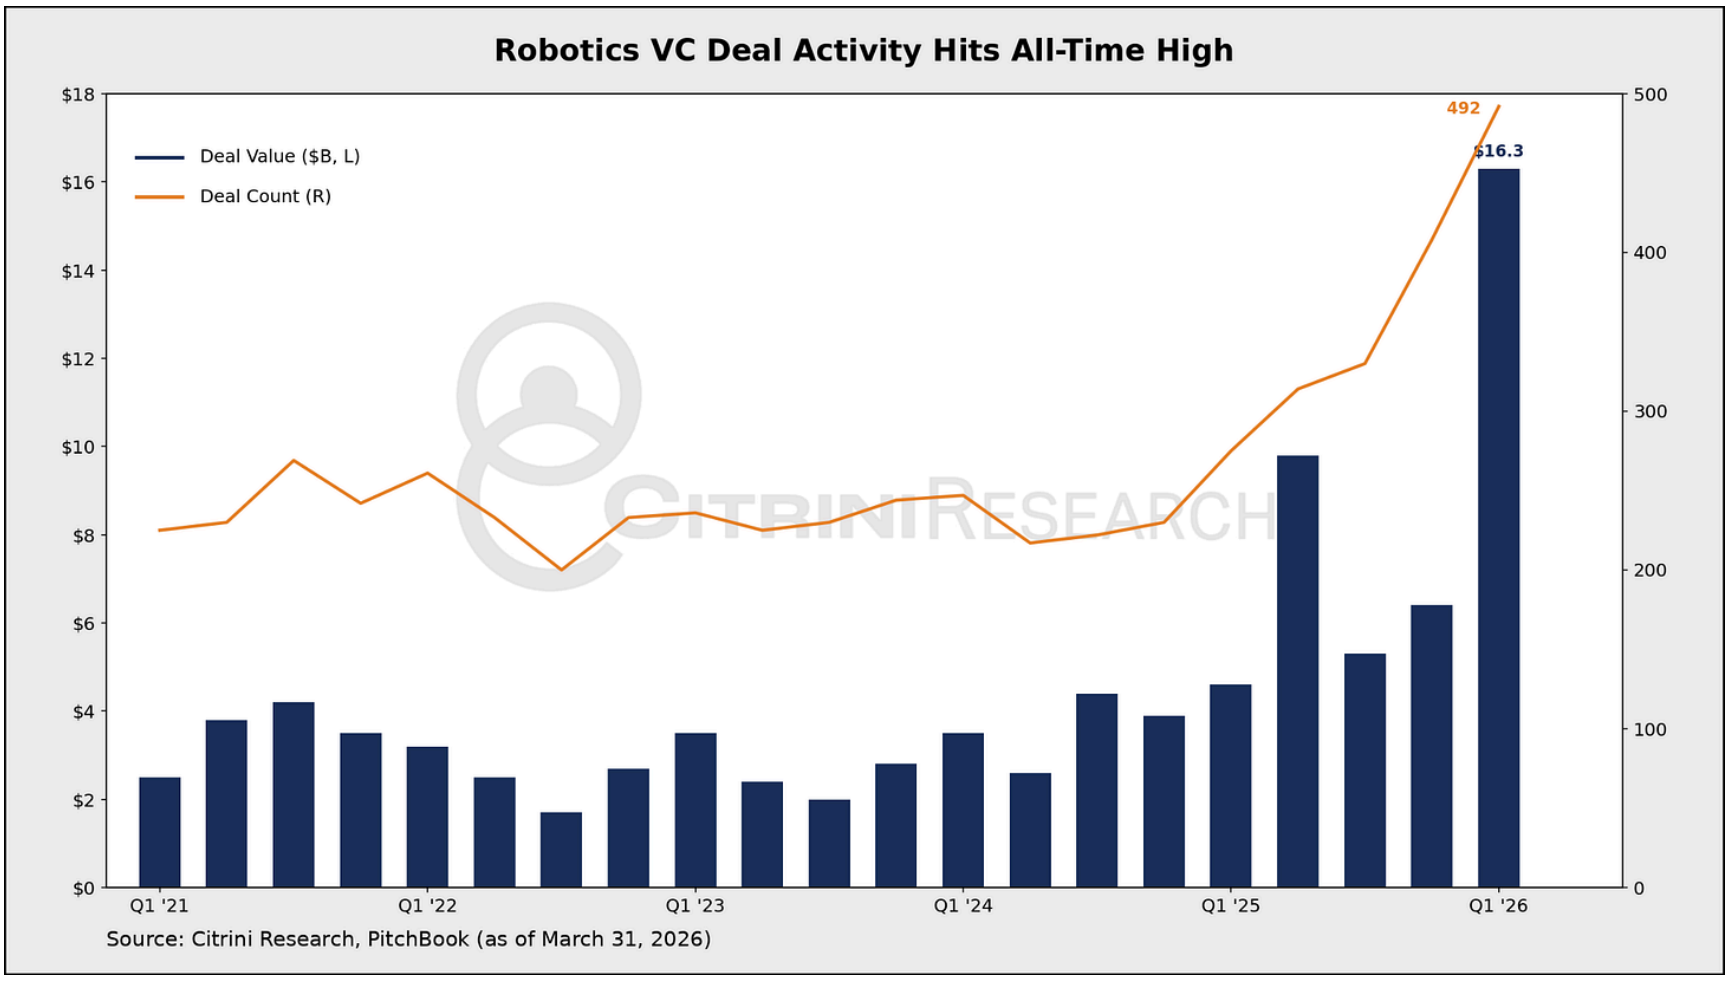
进展——也许最重要的因素在于，我认为人形机器人的能力提升速度将远超许多人的想象。

Funding & Reflexivity

资金与反身性

The final point is the dual edged sword. This is a highly capital-intensive business that does not make meaningful revenues today. Ergo, the entire thesis ultimately relies on **funding.**

最后这一点是一把双刃剑。 这是一个高度资本密集型的行业，目前尚无可观的收入。因此，整个投资逻辑最终都依赖于**资金。**



Right now funding is growing, *RAPIDLY*. My base case expectation is that rapid improvement in capabilities will create a positive-reinforcement cycle that drives more investment in the space. If it breaks down due to investor disappointment *OR* something entirely unrelated (rates, macro, VC dry powder, AI infra downturn etc.), it would invalidate the thesis for most of the beneficiaries.

当前该领域的资金正在增长，而且增速***极快***。我的基准情形预期是：能力的快速提升将形成一个正反馈循环，推动该领域获得更多投资。如果这一循环因投资者失望*或*某些完全无关的因素（利率、宏观环境、VC 可投资金、AI 基础设施低迷等）而中断，那么对大多数受益标的而言，这一投资逻辑将失效。

5) What's the Trade?

5）怎么交易？

This is the hard part. There are ETFs, sell-side reports, and of course, our own broad Robotics Universe and sub baskets that we published some 16 months ago. But nearly all of these fall into a similar trap.

这是最难的部分。市面上已有 ETF、卖方报告，当然还有我们自己在约 16 个月前发布的、覆盖面较广的 Robotics Universe 及其子篮子。但几乎所有这些方案都落入了一个类似的陷阱。

One of the biggest problems is direct accessibility. Most of the compelling US-based names are private and generally not available, *though there are some publicly traded funds of private investments that provide exposure to many of the leading private players here.* Other leaders are buried in much larger entities like Boston Dynamics within **Hyundai Motor Group (005380 KS)** and Optimus within **Tesla (TSLA US)**. Chinese leaders are similarly private and the exceptions like Unitree **(Yushu Technology, 688836 CH)** can be challenging shares to buy for many investors.

最大的问题之一是难以直接投资。 大多数引人注目的美国公司都是私有企业，通常无法直接投资，*不过也有一些投资私人资产的上市基金，能够提供对其中许多领先私有企业的敞口。* 另一些领军企业则藏在规模大得多的实体之中，例如**现代汽车集团（005380 KS）**旗下的波士顿动力（Boston Dynamics），以及特斯拉（TSLA US）**的 Optimus。中国的领军企业同样以私有为主，而像宇树科技（Unitree，688836 CH）这样的例外，对许多投资者而言也是难以买入的股票。

The second problem is that many of the names that have gotten stuffed in "robotics" baskets are nearly indistinguishable from high-momentum high-beta factors – big tech and semiconductors.

第二个问题在于，许多被塞进"机器人"篮子里的标的，与高动量、高贝塔因子——大型科技股和半导体——几乎没有区别。



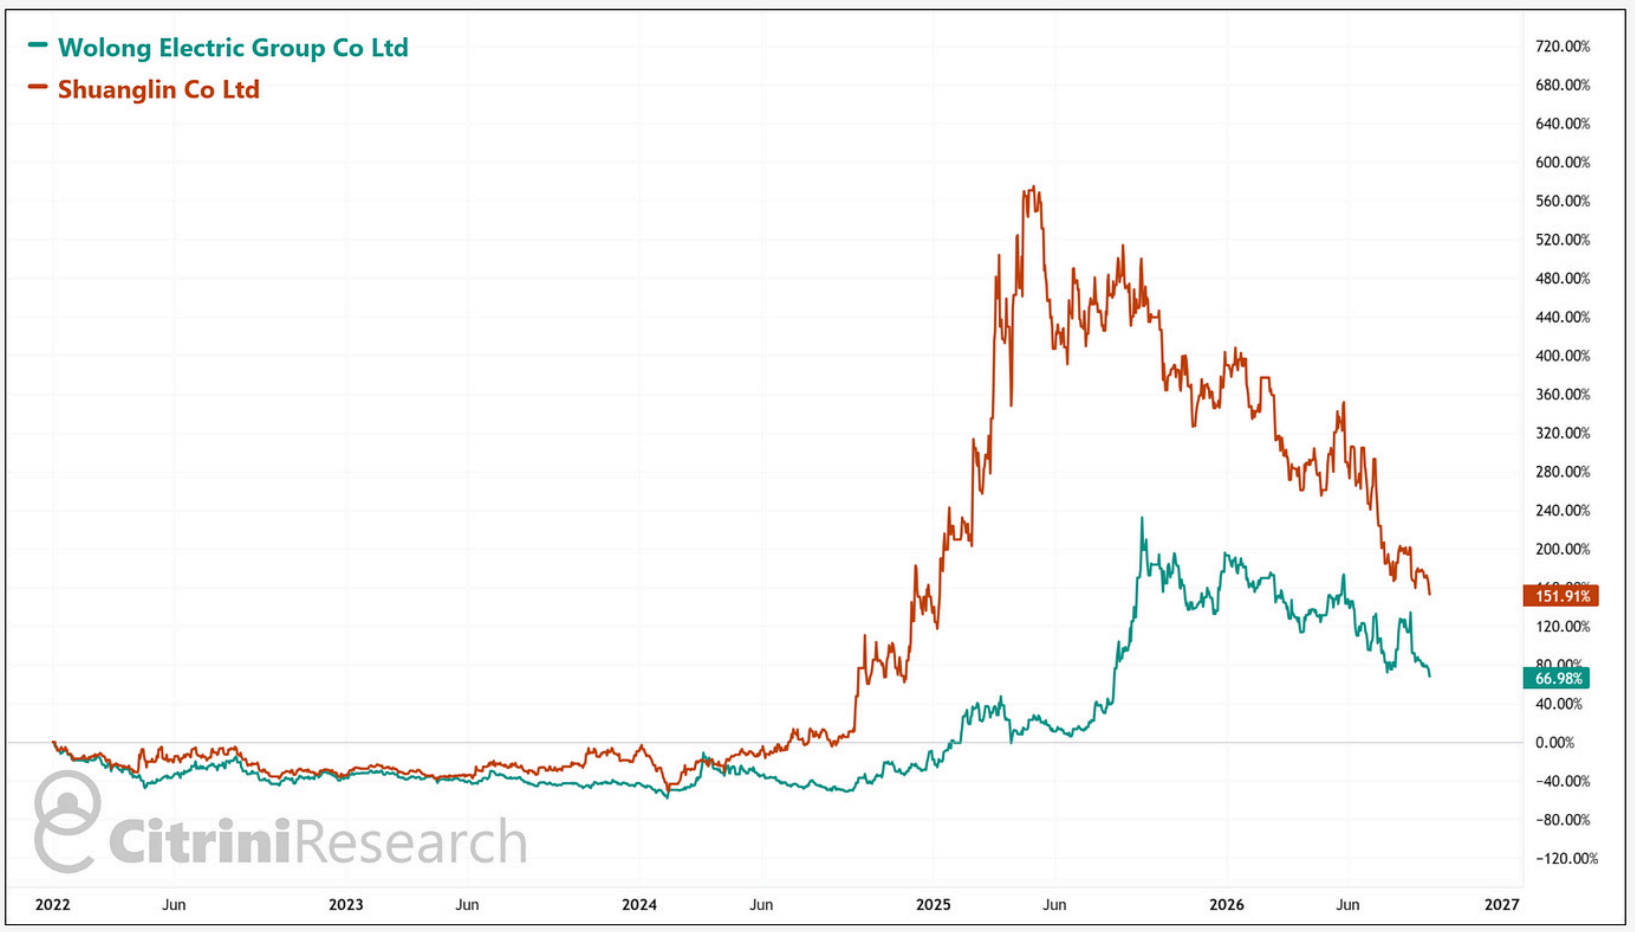
Source: Citrini Research

Partially, this is because of legitimate overlap with AI infrastructure names (think, NVDA) and partially because thematic investors have piled into the same names.

部分原因是这些标的与 AI 基础设施类股票（例如 NVDA）存在合理的重叠，另一部分原因则是主题投资者已蜂拥买入同一批标的。

Finally most clear-cut beneficiaries (and some not-so-clear-cut names) have already been bid up on the expectations, which makes the bull case harder. I saw this first hand in China. The companies whose logos made it to the “robotics supply chain” saw a massive hype bubble, but in person it was clear that there was nothing close to needle-moving scaled production in the near future.

最后，大多数界定明确的受益标的（以及一些不那么明确的标的）已因市场预期而被推高价格，这使得看多逻辑更难成立。我在中国亲眼见证了这一点。那些公司标志登上了“机器人供应链”图谱的企业曾经历巨大的炒作泡沫，但实地走访后可以清楚看到，短期内还谈不上有任何足以产生实质影响的规模化生产。



Source: Citrini Research

To be clear, high-beta hype has been a major winner this year (despite the drawdown) and certainly could continue. If that’s the angle you’re after, the Asimov basket is probably a good start.

需要澄清的是，高贝塔题材的炒作今年是主要赢家之一（尽管经历了回撤），而且当然可能延续下去。如果你追求的是这个角度，Asimov 组合大概是个不错的起点。

But I also want to focus on a narrow window. Are there public companies with meaningful near-term robotics tailwinds, low-correlation to the AI trade (if at all possible), and not already priced? I have three categories…

但我也想把焦点放在一个更窄的窗口上。是否存在这样的上市公司：近期有实质性的机器人顺风、与 AI 交易的相关性较低（如果可能的话），而且尚未被充分定价？我有三个类别……

Rapid Prototyping and Iterative Manufacturing:

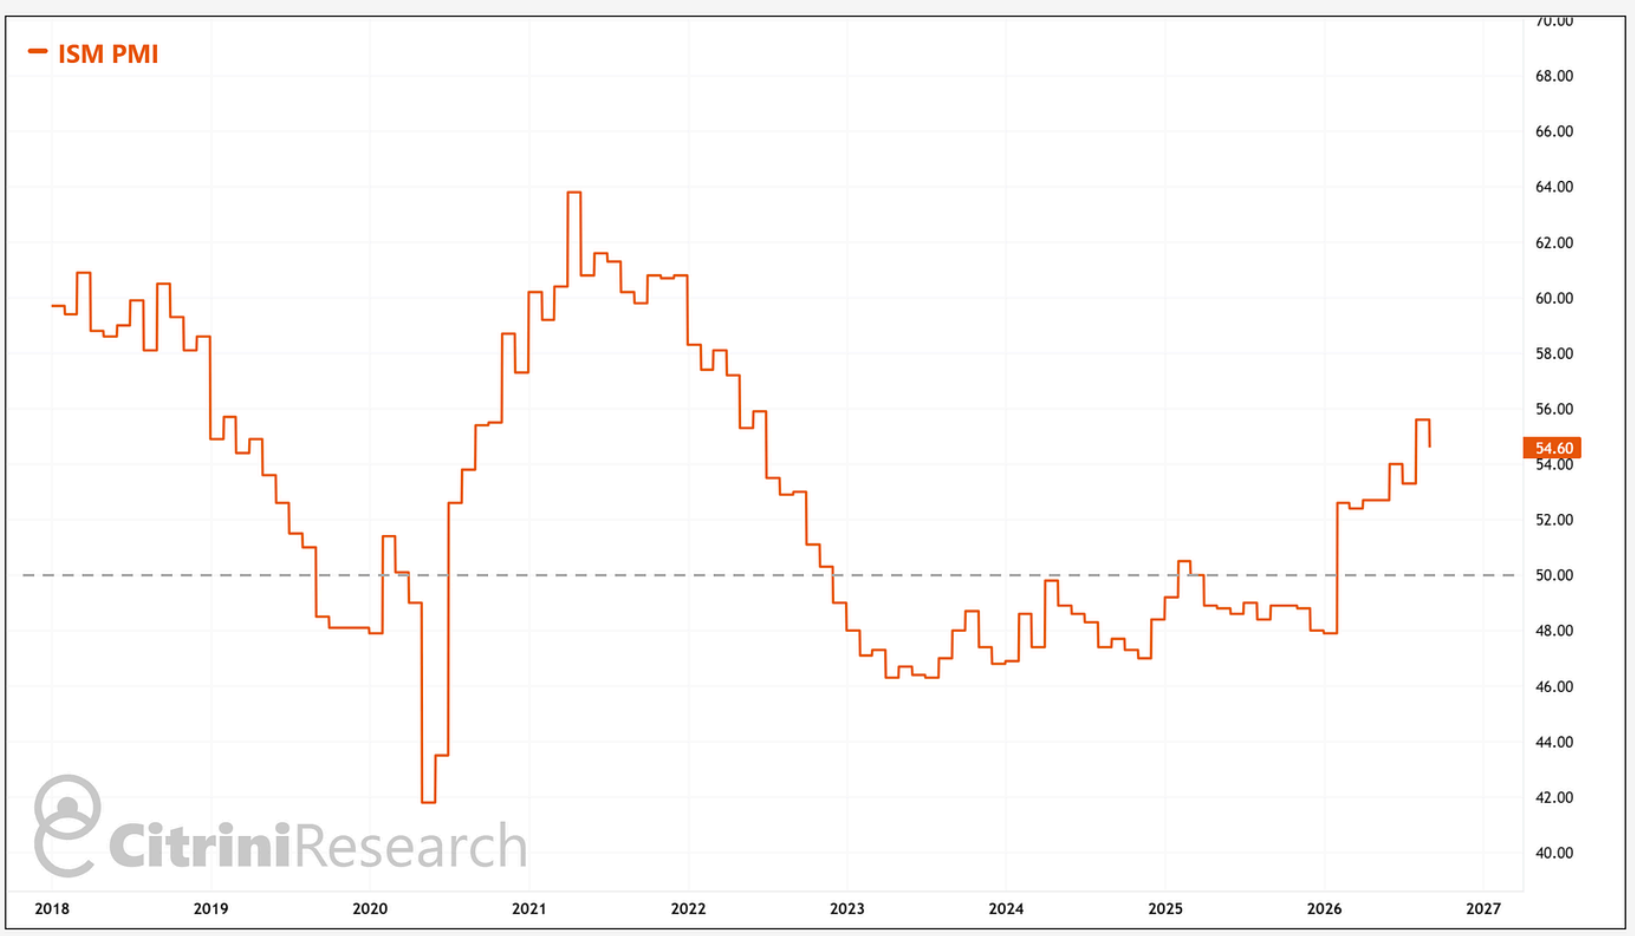
快速原型与迭代式制造：

I'm bullish on this sector first because of what I saw at many actual robotics companies and second because it ties into the other big trends we've followed at Citrini: onshoring, industrial protectionism, defense tech.

我看好这个板块，首先是因为我在许多真正的机器人公司实地所见，其次是因为它与我们 Citrini 一直追踪的其他大趋势相契合：产业回流、工业保护主义和国防科技。

The kicker? These are cyclical companies coming out of a downturn (2022-2025) driven by post-COVID reset, higher interest rates, energy costs, a crappy auto cycle, and eventually the Trump tariffs last year.

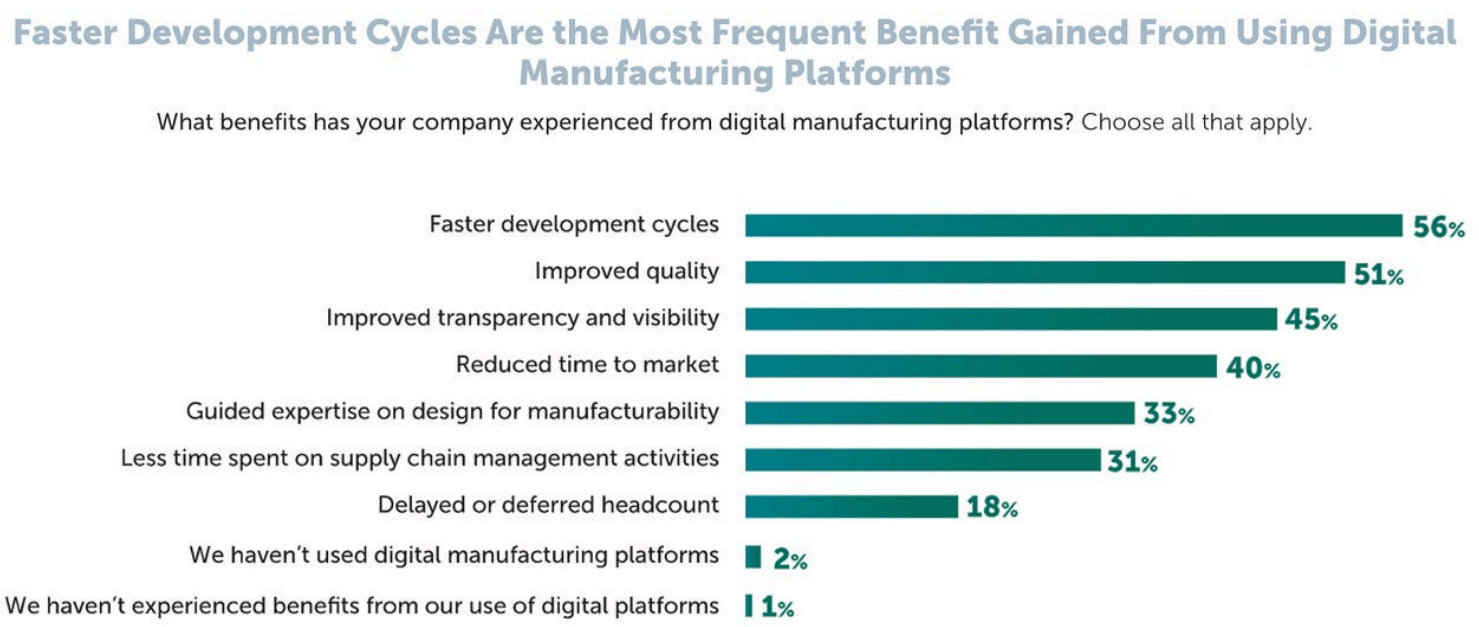
*妙处何在？*这些周期性公司正在走出 2022-2025 年的下行期，那轮下行由疫情后的重置、更高的利率、能源成本、糟糕的汽车周期，以及去年最终落地的特朗普关税所驱动。



Source: Citrini Research

This all is getting turbocharged by AI diffusion. Similar to drug discovery, generative design, LLM-assisted CAD, physics simulation, and reinforcement learning are driving the marginal cost of creating candidate designs to zero. We now have more ideas for hardware than we can realistically test, never mind produce.

这一切正因 AI 的扩散而获得强力加速。与药物研发领域的情况类似，生成式设计、大语言模型辅助的 CAD、物理仿真和强化学习正在将创建候选设计的边际成本推向零。如今我们在硬件方面的创意已多到超出实际能够测试的范围，更不用说投入生产了。



This shifts the bottleneck from design to real-world verification.

这使瓶颈从设计阶段转移到了现实世界的验证。

Manufacturing priorities move from volume to *variety* – more ideas, more variation, faster testing and iteration. Mass production on shrinking margins can come later; for now, the companies that monetize this shift earn from the frequency of design iterations rather than with unit production volume.

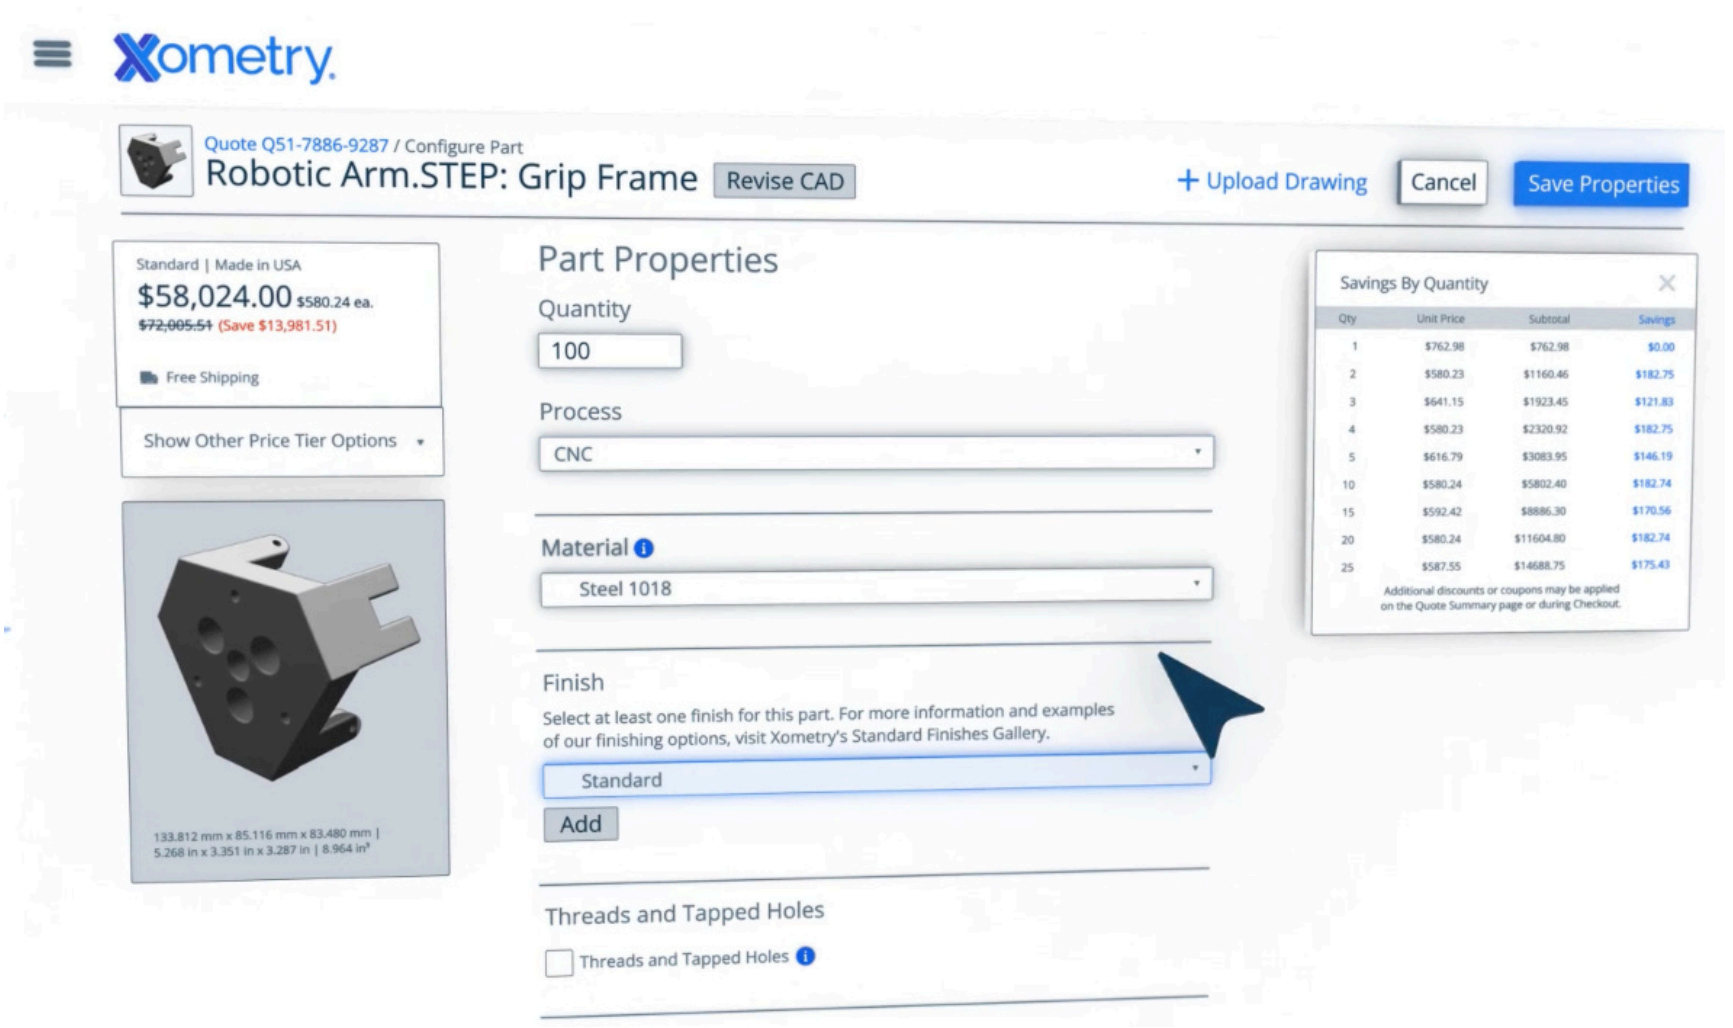
制造业的优先级正从产量转向*多样性*——更多创意、更多变化、更快的测试与迭代。在利润率不断收窄下进行的大规模生产可以留待日后；就目前而言，从这一转变中实现变现的公司，其收益来自设计迭代的频率，而非单位产量。

XOMETRY (XMTR US) isthe frontrunner marketplace for matching AI generated ideas with manufacturers. Xometry doesn't own the manufacturing facilities outright but instead sits in between buyers and sellers, monetizing each generated idea and design iteration.

XOMETRY (XMTR US) 是撮合 AI 生成创意与制造商的领先市场平台。Xometry 并不直接拥有制造设施，而是居于买卖双方之间，对每一个生成的创意和设计迭代进行变现。

A buyer can upload a CAD file, get price quotes in minutes, and be ready to take delivery from suppliers in as little as three days.

买家只需上传 CAD 文件，几分钟内即可获得报价，最快三天便可从供应商处提货。



The relatively newly appointed CEO Sanjeev Singh Sahni outlined the opportunity quite clearly on its most recent earnings call, on August 4, 2026:

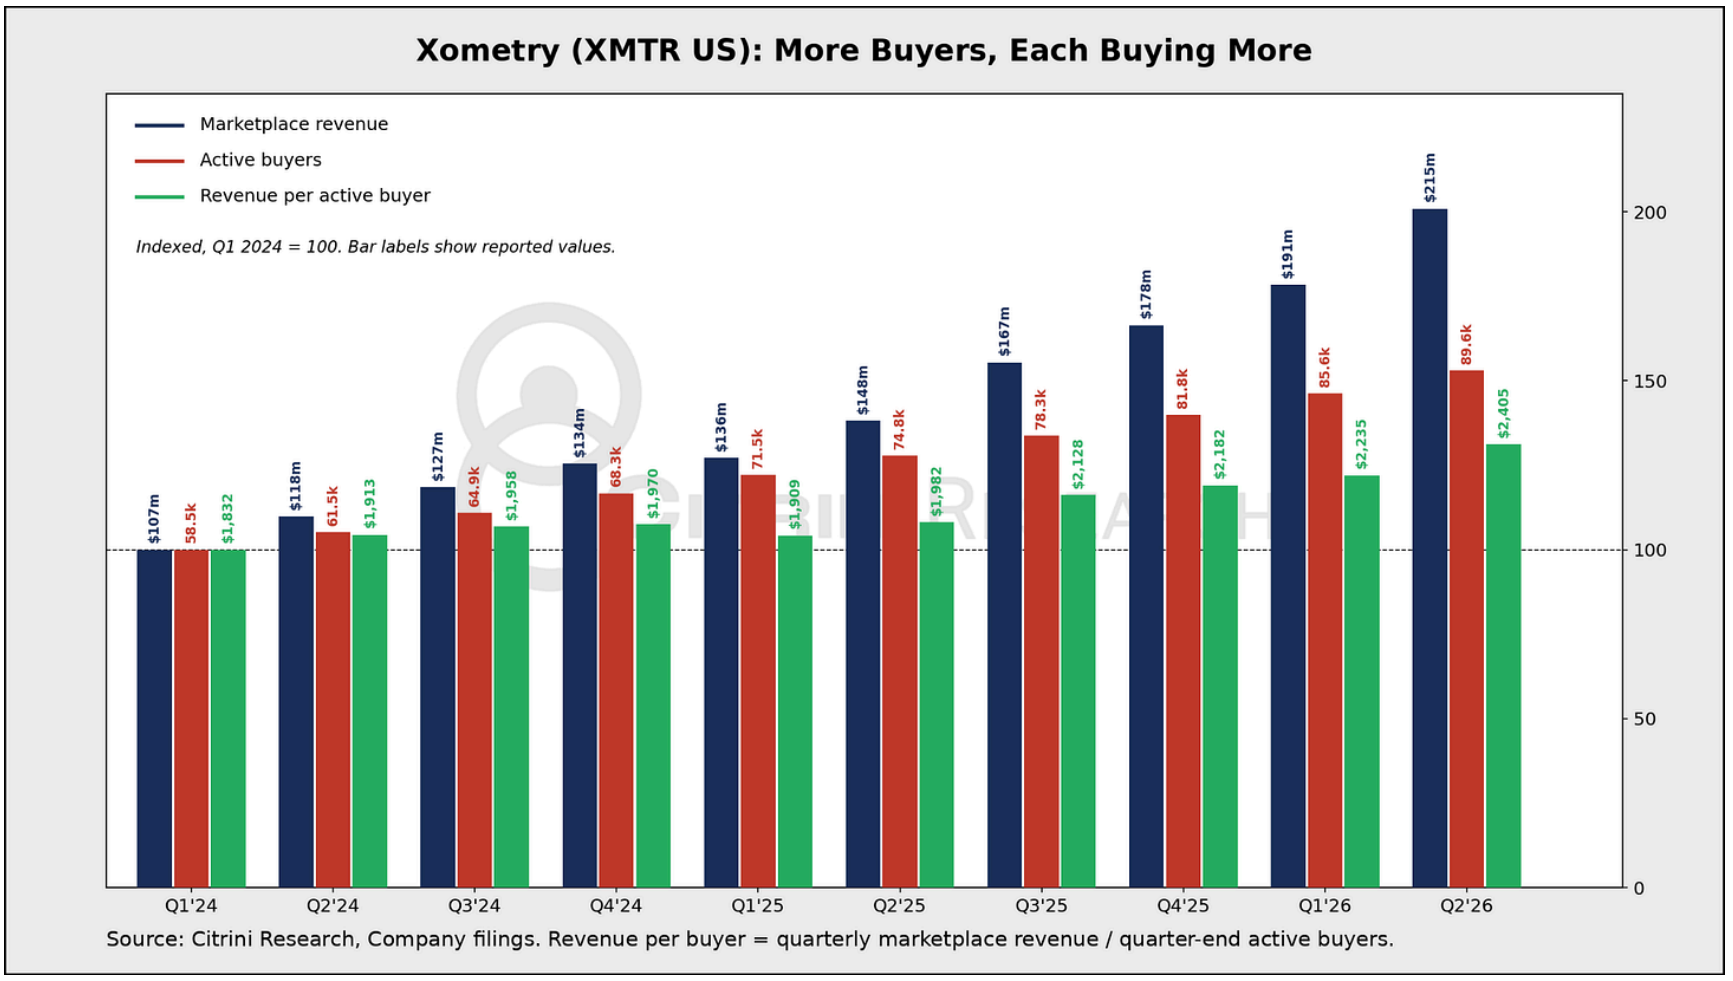
新近上任的首席执行官 Sanjeev Singh Sahni 在 2026 年 8 月 4 日的最近一次财报电话会议上，把这一机遇讲得相当清楚：

> ***"We see a long-term tailwind in physical AI. The rapid growth of robotics, autonomous systems, and defense platforms creates a new category of customers needing fast on-demand parts*** *[...]*
>
> *This deeper integration is exactly what we saw with a major enterprise robotics leader who faced a critical build deadline. They turned to us with a massive challenge, nearly 200 complex spots, including metal structures and tight tolerance CNC components, all needed within a short delivery window. We brought the strength of our marketplace to bear and leveraged 40 network suppliers to deliver the full build on time and on spec. This success proved that our platform is the go-to-solution for high stakes, large scale hardware programs."*

> ***"我们认为，实体AI（Physical AI）是一项长期顺风。机器人、自主系统和国防平台的快速增长，催生出一类需要快速按需获取零部件的新客户"*** *[...]*
>
> *"这种更深度的整合，恰好体现在一家大型企业级机器人领军企业的案例中。对方面临一个关键的制造交付期限，带着一项巨大的挑战找到我们：近 200 个复杂工件，包括金属结构件和公差要求严苛的 CNC 零件，全部需要在很短的交付窗口内完成。我们调动市场平台的力量，依托 40 家网络供应商，按时、按规格交付了全部构件。这次成功证明，我们的平台是高风险、大规模硬件项目的首选解决方案。"*

Marketplace revenue is accelerating sequentially for more than a year, with topline growth driven by both buyer count and spend.

市场（Marketplace）业务收入已连续一年多逐期环比加速，整体营收增长由买家数量和买家支出共同驱动。



How much is priced? For a software/marketplace, a 5x revenue multiple on accelerating 41% YoY growth while swinging to GAAP profitability doesn't feel expensive to me.

市场定价了多少？对于一家软件/平台型公司而言，在同比增速加快至 41%、同时扭亏转向 GAAP 盈利的背景下，5 倍市销率在我看来并不算贵。

Xometry is probably more of a market favorite, but I actually like the setup for Protolabs (PRLB US) better, which, unlike Xometry, owns its manufacturing facilities. Given its slogan is literally "manufacturing accelerated" it certainly fits our theme.

Xometry 或许更受市场青睐，但其实我更看好 Protolabs (PRLB US) 的布局——与 Xometry 不同，Protolabs 自有制造工厂。既然其口号正是"加速制造"（manufacturing accelerated），它显然契合我们的主题。



The company's capabilities are spot on for robotics needs – CNC, 3D, IM.

该公司的能力正好契合机器人领域的需求——数控加工（CNC）、3D 打印、注塑成型（IM）。

Our Manufacturing Capabilities for Rapid Prototyping

Our Rapid Prototyping services are available in 3 core manufacturing processes; these include 3D Prototyping, CNC Prototyping and Prototype Injection Moulding. Each process has its own unique features and benefits. The process you choose will depend on your specific product and project stage. Continue reading to learn more about our prototype manufacturing capabilities.



3D Printing

Our online 3D printing service consists of six additive manufacturing technologies capable of producing prototype and end-use parts in a wide range of plastic, metal, and elastomeric materials. In addition to our vast material selection, we offer several post-processing options like clear coat, painting, heat treatment, polishing, and more.

View Capabilities



CNC Machining

Our CNC machining capacity is totally unrivalled, as is our service offering. That's why we can comfortably say that we deliver the industries most comprehensive choice. For this reason, your parts will be shipped as ordered, and on time. Explore our wide range of finishing options, which includes anodising and chromate plating.

View Capabilities



Injection Moulding

Are you ready for a more cost-effective and time-efficient way to manufacture parts? Our injection moulding services utilise aluminium moulds which start at just £995, and are also available in as fast as 1 day with our quick turnaround service. Need more? Our finishing options for Injection moulding include pad printing, laser engraving, threaded inserts, and more.

View Capabilities

Their CEO outlined where they believe we are in terms of the cycle, on the Q2 2026 Earnings Call:

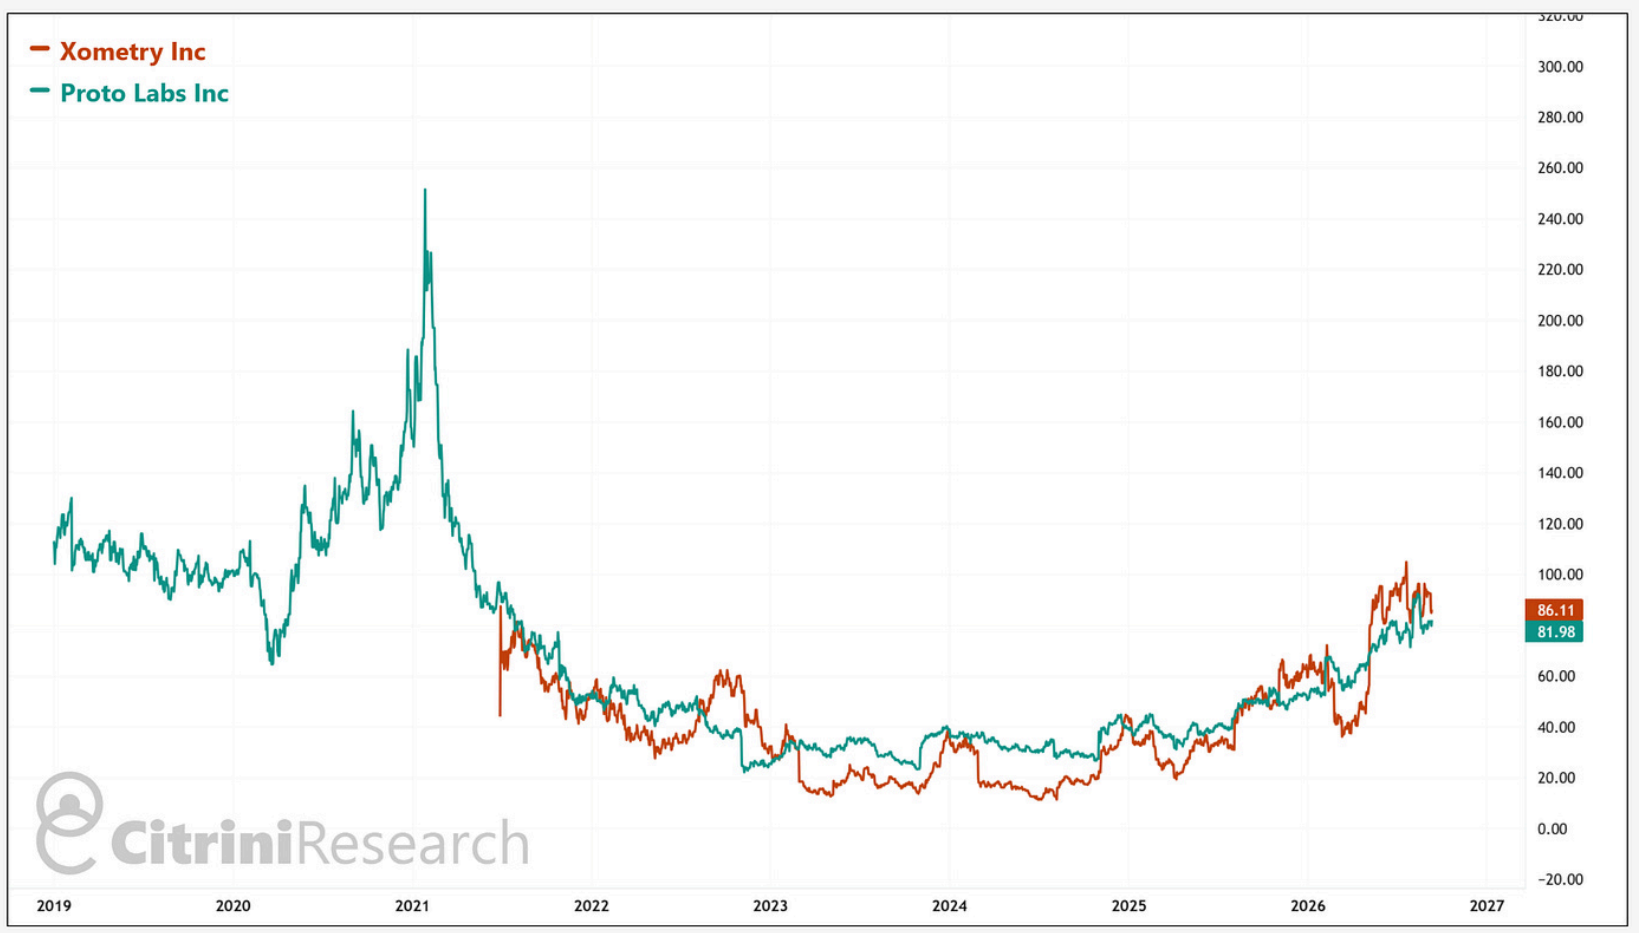
其 CEO 在 2026 年第二季度财报电话会议上阐述了他们对当前周期所处位置的判断：

> *"**We are in the early cycle of a lot of industries which are early in innovation.** You think about drones, robotics, data centers, satellites, rockets. All of those are early cycle and with a long tail ahead of us. What we offer is a unique proposition."*

> ***"我们正处于众多仍处于创新早期阶段的行业的早期周期。** 你想想无人机、机器人、数据中心、卫星、火箭。这些都处于早期周期，且前方还有很长的长尾。我们提供的是一种独特的价值主张。"*

Tactically, both companies fell into the post-COVID trough for years, but are digging their way out of the hole with a sustained multiyear uptrend. XMTR is the more volatile of the two and could be on the verge of a short term reset, but both are trading on the same fundamentals, which I think will continue to improve. *Look at that smile*.

从战术层面看，这两家公司在疫情后都陷入了多年的低谷，但正凭借持续的多年上行趋势逐步爬出深坑。XMTR 是两者中波动更大的一家，可能即将迎来一次短期重置，但两者交易的都是同一套基本面，而我认为这套基本面将继续改善。*看看那个笑容*。



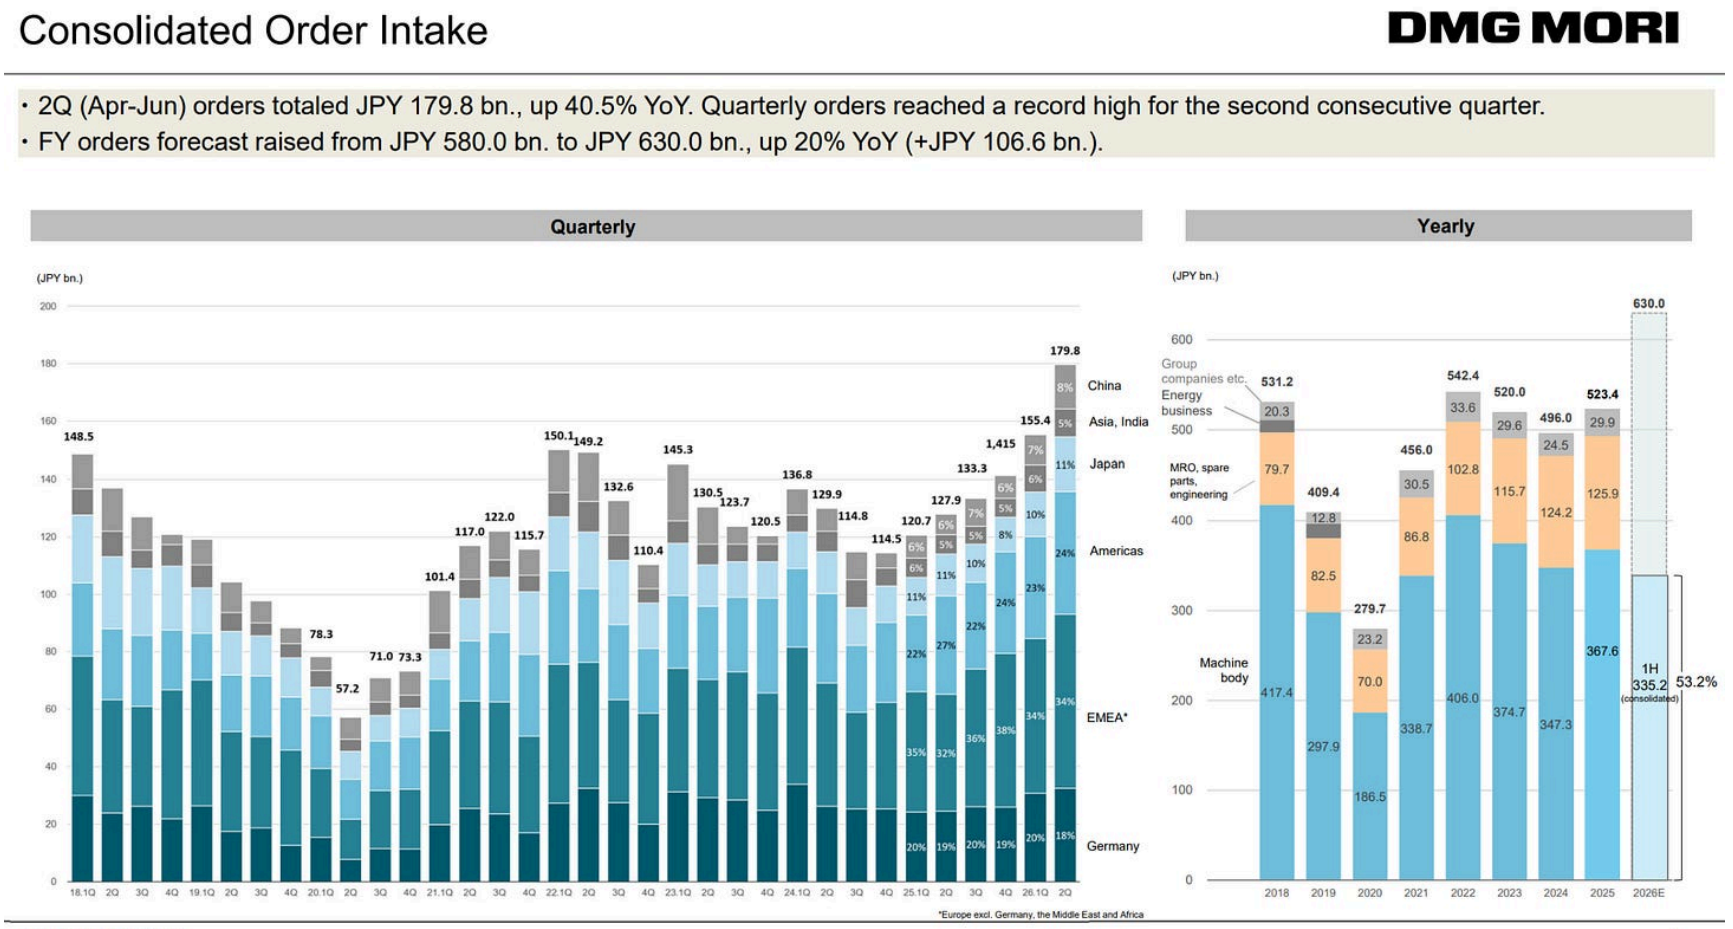
Source: Citrini Research

Machining and milling tool companies are a step below. It may be a stretch to get into lumpy, second-order derivative capital equipment but the same cyclical signs are pointing in the right direction. Tariffs are an even bigger deal here as American capital purchases were delayed and global manufacturing remained cautious.

机加工与铣削工具公司则低一个层级。切入这类订单波动较大、属于二阶衍生的资本设备或许有些牵强，但同样的周期性信号正指向正确方向。关税在这一环节的影响更为显著，因为美国资本设备采购此前遭到推迟，而全球制造业始终保持谨慎。

Two I like are **DMG Mori (6141 JP)** and **Okuma (6103 JP).**

我看好的两家是**DMG Mori（6141 JP）和Okuma（6103 JP）。**

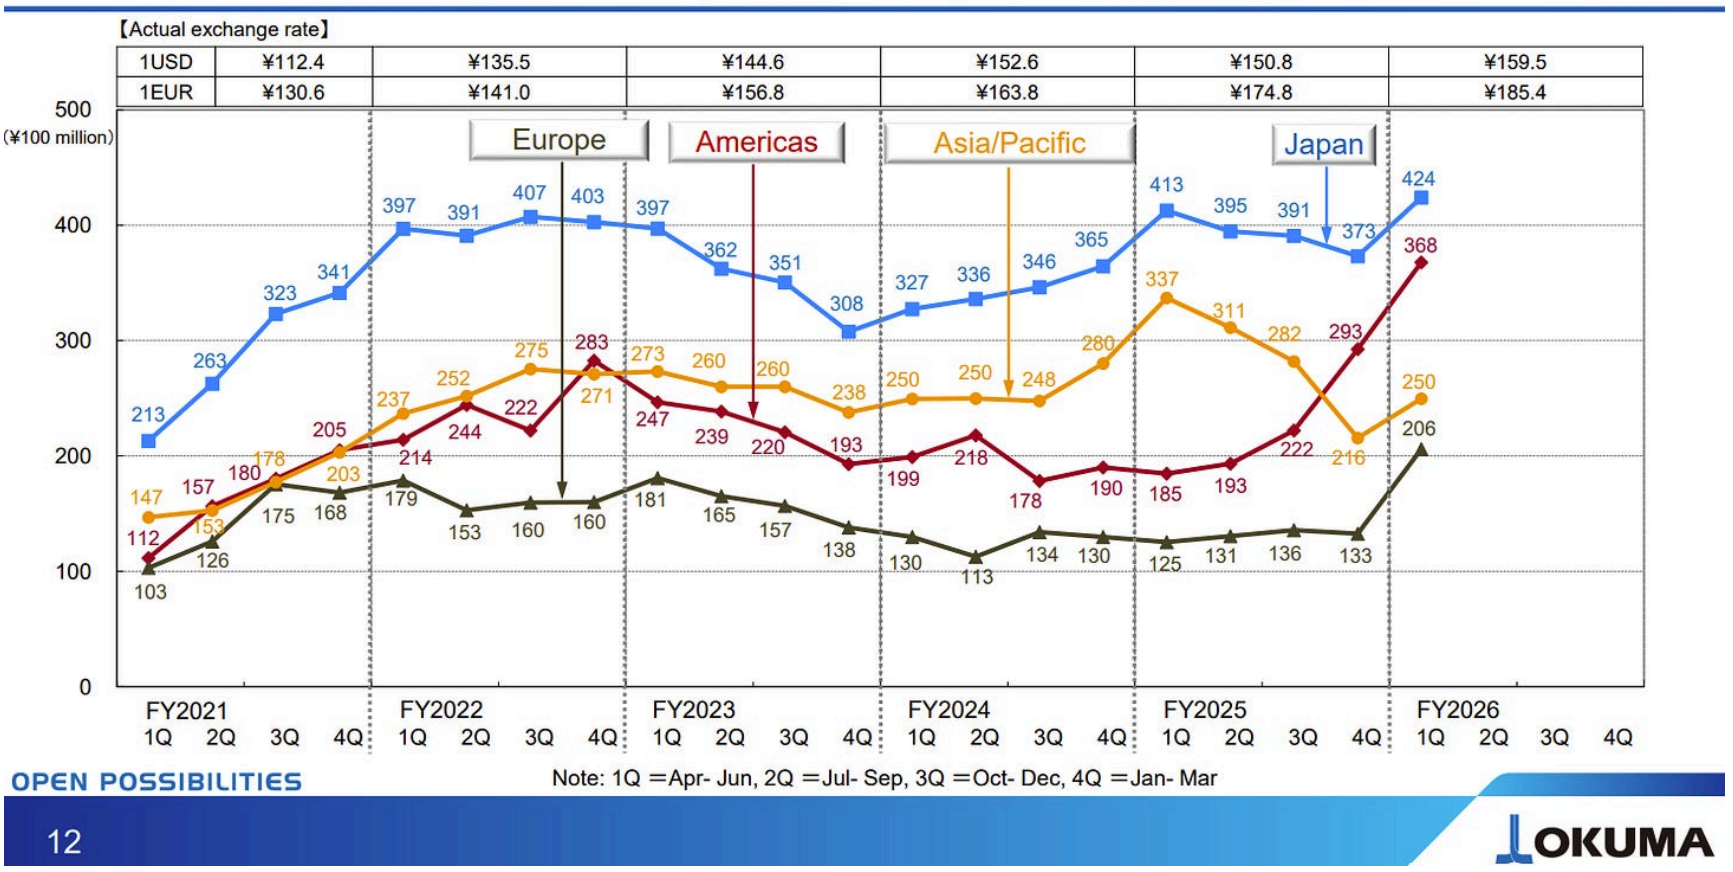


DMG is one of the global leaders in the CNC machining market, and is seeing its order book grow 40% YoY as of 2Q26 and accelerating, and continues to upgrade its guidance throughout the year. With sales recognition lagging the order book by about a year, we have strong visibility.

DMG 是 CNC 加工市场的全球领导者之一，截至 2026 年二季度，其订单储备同比增长 40% 且增速仍在加快，并在全年持续上调业绩指引。由于收入确认较订单储备滞后约一年，我们对其未来业绩具有很高的可见性。

Okuma is a bit smaller and follows the same story (record order intake and backlog), but with a stronger through-the-cycle growth trajectory.

Okuma 体量略小，情况与此类似（新增订单与在手订单均创历史新高），但其穿越周期的增长轨迹更为强劲。



US orders are booming and the company is specifically calling out a restart of orders for small and medium businesses, w tics companies would show up.

美国订单正一片火热，公司还特别点名中小企业订单出现重启，而这些机器人公司正是出现在这一客群之中。

> *"In terms of market conditions by region, in the U.S., the high level of demand continued, primarily from major corporat ce and satellites, defense-related, data center, and energy-related sectors. **Furthermore, even with medium- and small-siz re signs of a restart in capital investments "***



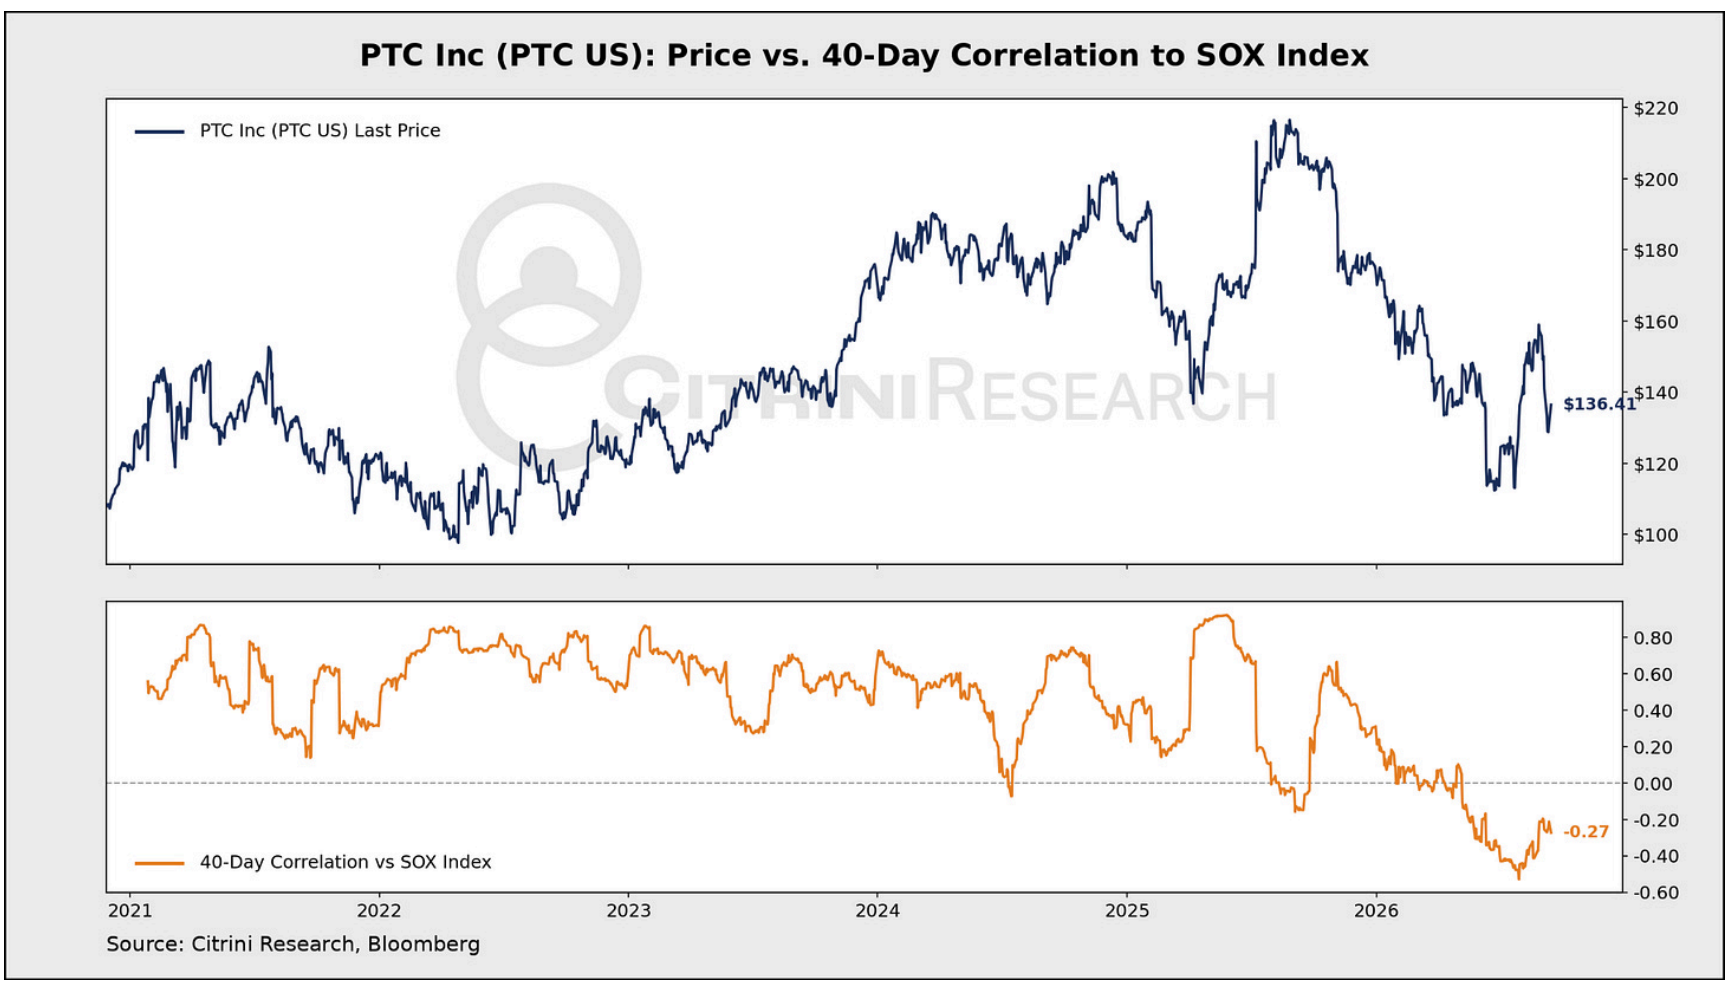
> *“从各地区的市场状况来看，在美国，高企的需求得以延续，主要来自大型企业，集中于航空、航天与卫星、国防相关、数据中心以及能源相关领域。**此外，即便是中小企业，也出现了资本投资重启的迹象**”*

(Small cap degens might look at Hurco Companies (HURC US)).

（偏爱小盘股的投机者不妨留意 Hurco Companies（HURC US）。）

Rounding out the group are two doubly heretical picks – **PTC inc (PTC US)** and **Autodesk (ADSK US).** Not only are they software companies *(gasp)* but OpenAI Astra’s abilities in OpenCAD felt like a throwback to when Anthropic was picking off SaaS companies one by one.

为这个组合收官的是两个“双重离经叛道”的标的——**PTC（PTC US）*与 ***Autodesk（ADSK US）***。它们不仅是软件公司*（倒吸一口凉气）*，而且 OpenAI Astra 在 OpenCAD 方面展现的能力，让人回想起 Anthropic 逐个猎杀 SaaS 公司的那段时期。



I’m not going to argue that there’s no risk (seats, pricing, disruption), but I can also argue that there are upsides as well – Astra’s outputs would likely be ported into each of the companies’ systems (Revit, Creo, Fusion, etc), and easier idea generation means more consumption.

我不打算争论说这里没有任何风险（席位授权、定价、颠覆冲击），但我同样可以举出上行方面的理由——Astra 的输出成果很可能会被移植到各家公司的系统中（Revit、Creo、Fusion 等），而更轻松的创意生成意味着更多的消费量。

But really one of the most compelling angles here is that these companies are still notching consistent growth and cash flow at their lowest trading multiples in a decade and are ***negative-beta*** to the AI infrastructure trade.It’s a carry-paying hedge that the market may just decide is a sneaky winner.

但实际上，这里最有说服力的角度之一是：这些公司仍在以十年来的最低交易估值倍数持续录得稳健增长和现金流，并且相对 AI 基础设施交易呈现***负贝塔***。这是一种自带持有收益的对冲，市场说不定哪天就会认定它是个隐形的赢家。

Concentrated Suppliers

集中度高的供应商

If our prediction is that most large industrial companies won’t see a material impact from the near-term level of production, who are the concentrated suppliers that might?

如果我们的预测是大多数大型工业公司不会因近期产量水平而受到实质性影响，那么又有哪些供应高度集中的公司可能会受到影响？

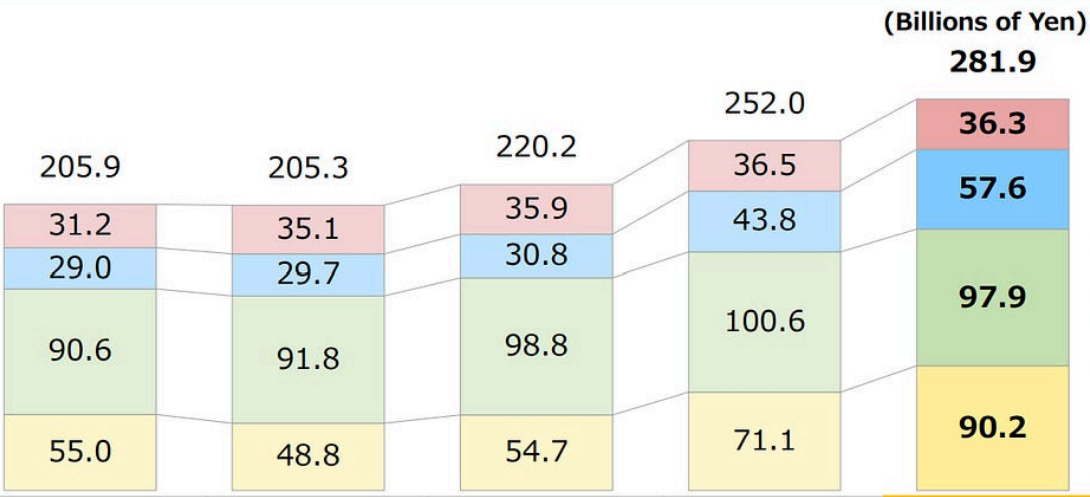
ROBOTIS (108490 KS) is perfect in theory with its suite of R&D products – it’s going to put up big growth numbers, but at 100x sales, it’s hard to get excited. By the time it takes to grow into those numbers, the development curve will begin to favor others.

****ROBOTIS（108490 KS）**凭借其成套的研发产品在理论上堪称完美——它将会交出亮眼的增长数字，但以 100 倍市销率计，实在难以令人兴奋。而等到它的业绩增长到足以消化这一估值时，技术的发展曲线将开始向其他公司倾斜。

One of the most perplexing has been **FANUC (6954 JP)**. This company is a poster child for the entire robotics industry – it is connected to Physical AI, but for some reason hasn’t gotten the narrative love that many others have. Regardless, sales have made it through the same cyclical trough that has plagued the other capital equipment industries. Sales are now lapping last year at a 17% clip while the order book is growing at 37%.

最令人费解的标的之一是 **FANUC（6954 JP）**。这家公司是整个机器人行业的标杆企业——它与具身智能（Physical AI）存在关联，但不知为何却没有获得许多其他公司所享有的叙事热度。无论如何，其销售已经渡过了困扰其他资本设备行业的同一轮周期性低谷。目前销售额同比增速达到 17%，而在手订单正以 37% 的速度增长。

Consolidated Orders by Division



Division	Change from Same Period of Previous Year	Change from Previous Quarter
TOTAL	+36.9%	+11.9%
Service	+16.2%	-0.6%
Robomachine	+98.3%	+31.3%
Robot	+8.1%	-2.6%
FA	+63.8%	+26.8%

Composition	2025Q1	2025Q2	2025Q3	2025Q4	2026Q1
Service	15.2%	17.1%	16.3%	14.5%	12.9%
Robomachine	14.1%	14.5%	14.0%	17.4%	20.4%
Robot	44.0%	44.7%	44.9%	39.9%	34.7%
FA	26.7%	23.7%	24.8%	28.2%	32.0%

It's worth noting that Universal Robots (Teradyne owned), did seem to be far more present in Physical AI demos, as it's an easier plug-and-play device. (I'm remarkably bullish on Teradyne's robotics system as we discussed earlier, and I'll probably regret it later, but the stock is far more likely to trade on its semiconductor ATE business).

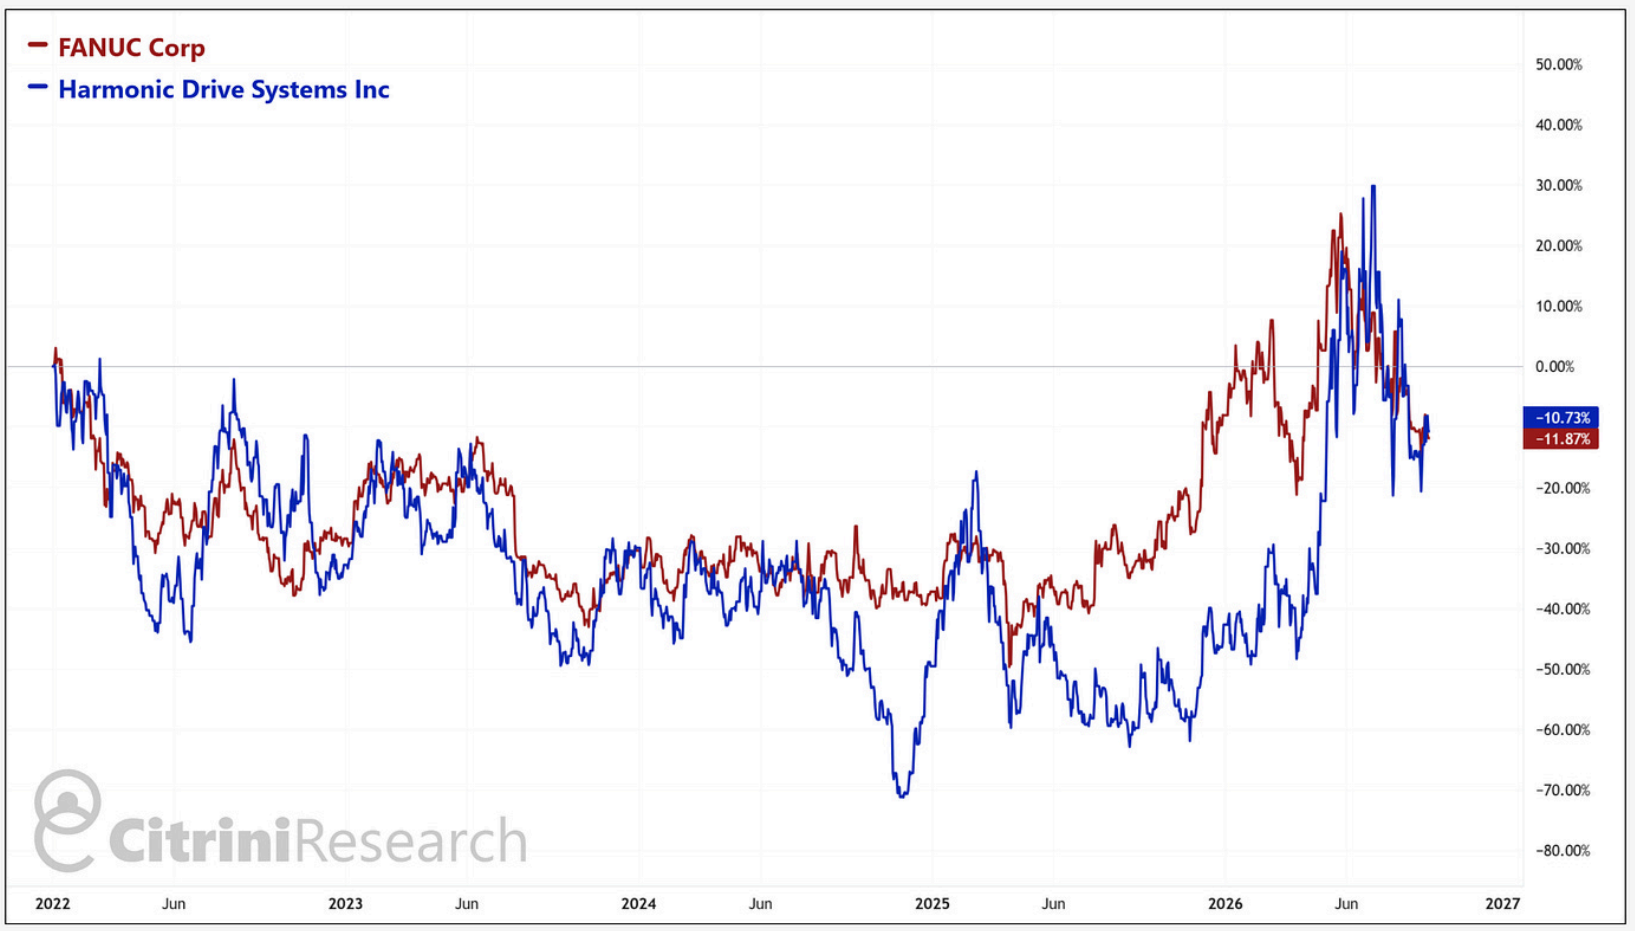
值得注意的是，Universal Robots（泰瑞达旗下）在物理 AI 演示中的出镜率明显更高，因为它是一种更容易即插即用的设备。（正如我们此前讨论的，我对泰瑞达的机器人系统相当看好——日后很可能会为此后悔——但这只股票的走势更有可能取决于其半导体 ATE 业务。）

Another stock that has frustrated investors for years is **Harmonic Drive Systems (6324 JP)**. Harmonic drives are high-precision flexible gears that allow for extreme torque with little backlash, and they're critical in hyper sensitive robotics systems. They ensure that joints stay perfectly stable where normal gear teeth would bounce. Based on the name alone, this should be a good setup, right?

另一只多年来令投资者屡屡受挫的股票是 **哈默纳科（Harmonic Drive Systems，6324 JP）**。谐波减速器是一种高精度柔性齿轮，能在回程间隙极小的情况下输出极高扭矩，因此对高灵敏度机器人系统至关重要。在普通齿轮齿会发生弹跳的部位，它们能确保关节始终保持完全稳定。仅凭名字来看，这本该是个不错的布局，对吧？

Just like Fanuc, its 1Q26 order book showed a massive jump across categories (even after excluding a large lump order from a medical customer). Its order book does include a meaningful and growing order from a single large American humanoid manufacturer.

与发那科一样，其 1Q26 订单簿显示各品类订单均大幅跳升（即使剔除来自一家医疗客户的大额整笔订单后也是如此）。其订单簿中确实包含一笔来自某家大型美国人形机器人制造商的订单，规模可观且持续增长。



Source: Citrini Research

The sales inflection stirred up some strong momentum in both stocks throughout the early summer, but both stocks still sit below their COVID peaks. The selloff coinciding with the momentum drawdown over the past two months provides an attractive entry for each.

销售拐点在整个初夏为这两只股票带来了强劲的上涨动能，但目前两者仍低于疫情时期的高点。过去两个月的股价回调恰逢动量因子回撤，为两者提供了颇具吸引力的入场时机。

How about LiDAR? We've touched on both the US-based **Ouster (OUST US)** and China's **Hesai (HSAI US)** numerous times, most recently highlighting their resolution live-color breakthroughs in our most recent State of the Themes in June.

那么激光雷达呢？ 我们曾多次提及美国的 **Ouster (OUST US)** 和中国的 **禾赛 (HSAI US)**，最近一次是在六月的 **主题现状** 中重点介绍了两者在分辨率与实时彩色成像方面的突破。

Starting with the negatives, it's clear that many roboticists want to get rid of LiDAR, largely due to cost and the advancements in camera-based vision perception (nice, Elon) which doesn't make me feel super warm and fuzzy.

先说负面因素。显然，许多机器人专家都想摆脱激光雷达，主要原因是成本，以及基于摄像头的视觉感知技术的进步（干得漂亮，马斯克）——这让我心里并不踏实。

At the same time, Ouster is claiming they weren't facing a *demand* problem, but a supply problem: production can't even keep up with the pace of orders. Take it with a grain of salt, but both companies are providing eye-popping quotes to back it up.

与此同时，Ouster 声称他们面临的并非*需求*问题，而是供应问题：生产速度甚至跟不上订单的节奏。这话姑且听之，但两家公司都拿出了令人瞠目的引述来佐证这一说法。

From Hesai…

来自禾赛……

> *"We work with more than 50 embodied AI companies worldwide, with recent orders coming from Unitree, Robbyant, Galbot, Galaxea, Dexmate, and many others. Our JT128 has become the go-to lidar for humanoid and quadruped robots, and the numbers are frankly remarkable. Robotics lidar shipments in the second quarter almost tripled compared with the same quarter last year.* ***Following year-over-year growth of 47% in 2023, 67% in 2024, and 426% in 2025, we expect another 2-3x increase 2026."***

> *"我们与全球 50 多家具身智能公司合作，近期订单来自宇树（Unitree）、Robbyant、Galbot、Galaxea、Dexmate 等众多公司。我们的 JT128 已成为人形机器人和四足机器人的首选激光雷达，相关数字坦率地说相当惊人。第二季度机器人激光雷达出货量几乎是去年同期的三倍。**继 2023 年同比增长 47%、2024 年增长 67%、2025 年增长 426% 之后，我们预计 2026 年将再实现 2-3 倍的增长。"*

And it's not just LiDAR… Hesai is now moving into the actuator game with its partnership with Sharpa where it is expecting to hit 100,000 units sold next year. This production scale may not mean a lot to a massive automotive supplier, but for a company with just ~$400 million in revenue this is a big deal.

而且不止激光雷达…… 禾赛目前正通过与 Sharpa 的合作进军执行器领域，预计明年销量将达到 10 万台。这一生产规模对大型汽车零部件供应商而言或许不算什么，但对于一家营收仅约 4 亿美元的公司来说，这是一件大事。

From Ouster…

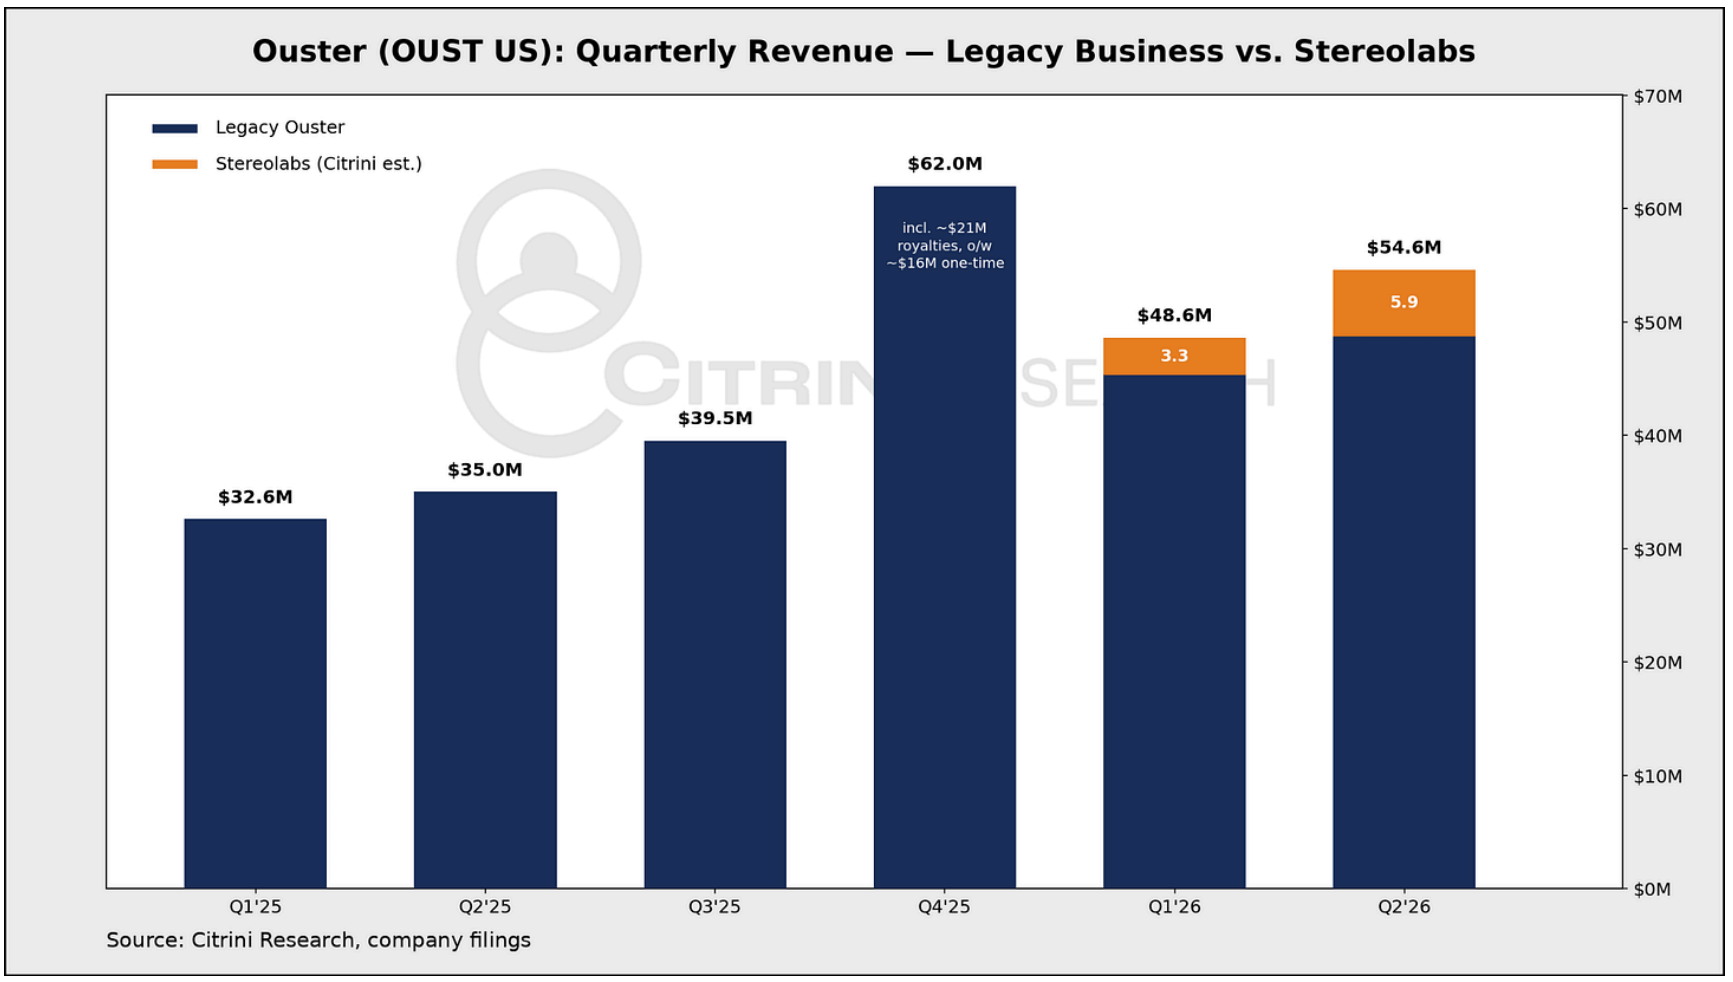
来自 Ouster……

> *"Robotics has been an incredible business for us. Q2 was the first full quarter with Stereolabs under the Ouster umbrella. And Stereolabs is stereo cameras that -- and monocular camera products, including their release of the ZED X Nano, which is a special stereo camera for wrist-mounted manipulation on humanoid and robotic arms. They have the most competitive set of products for this like hypergrowth humanoid and robotic manipulation market.* ***And we've seen incredible kind of interest and uptake from that customer set."***

> ***"机器人业务对我们而言是一门非常出色的生意。** 二季度是 Stereolabs 并入 Ouster 体系后的首个完整季度。Stereolabs 做的是双目相机——以及单目相机产品，包括新发布的 ZED X Nano，这是一款专门用于人形机器人和机械臂腕载操作的双目相机。在这个高速增长的人形机器人与机械臂操作市场，他们拥有最具竞争力的产品组合。**我们从这批客户那里看到了极高的兴趣度和采用意愿。"***

The reality is that despite the detractors, there is strong and growing demand for LiDAR, particularly in autonomous driving, smart cities, industrial movers, and infrastructure, before even considering new smart robots. And in the case of Ouster, their acquired camera business is highly relevant for overhead and wrist mounted cameras – necessary not just for the bots, but for every egocentric data capture system.

现实是，尽管不乏质疑者，激光雷达（LiDAR）的需求依然强劲且持续增长，尤其是在自动驾驶、智慧城市、工业搬运设备和基础设施领域——这还没有算上新型智能机器人带来的增量。而就 Ouster 而言，其收购来的相机业务与顶置相机和腕载相机高度相关——这类相机不仅是机器人所必需，也是所有第一视角数据采集系统的必备组件。

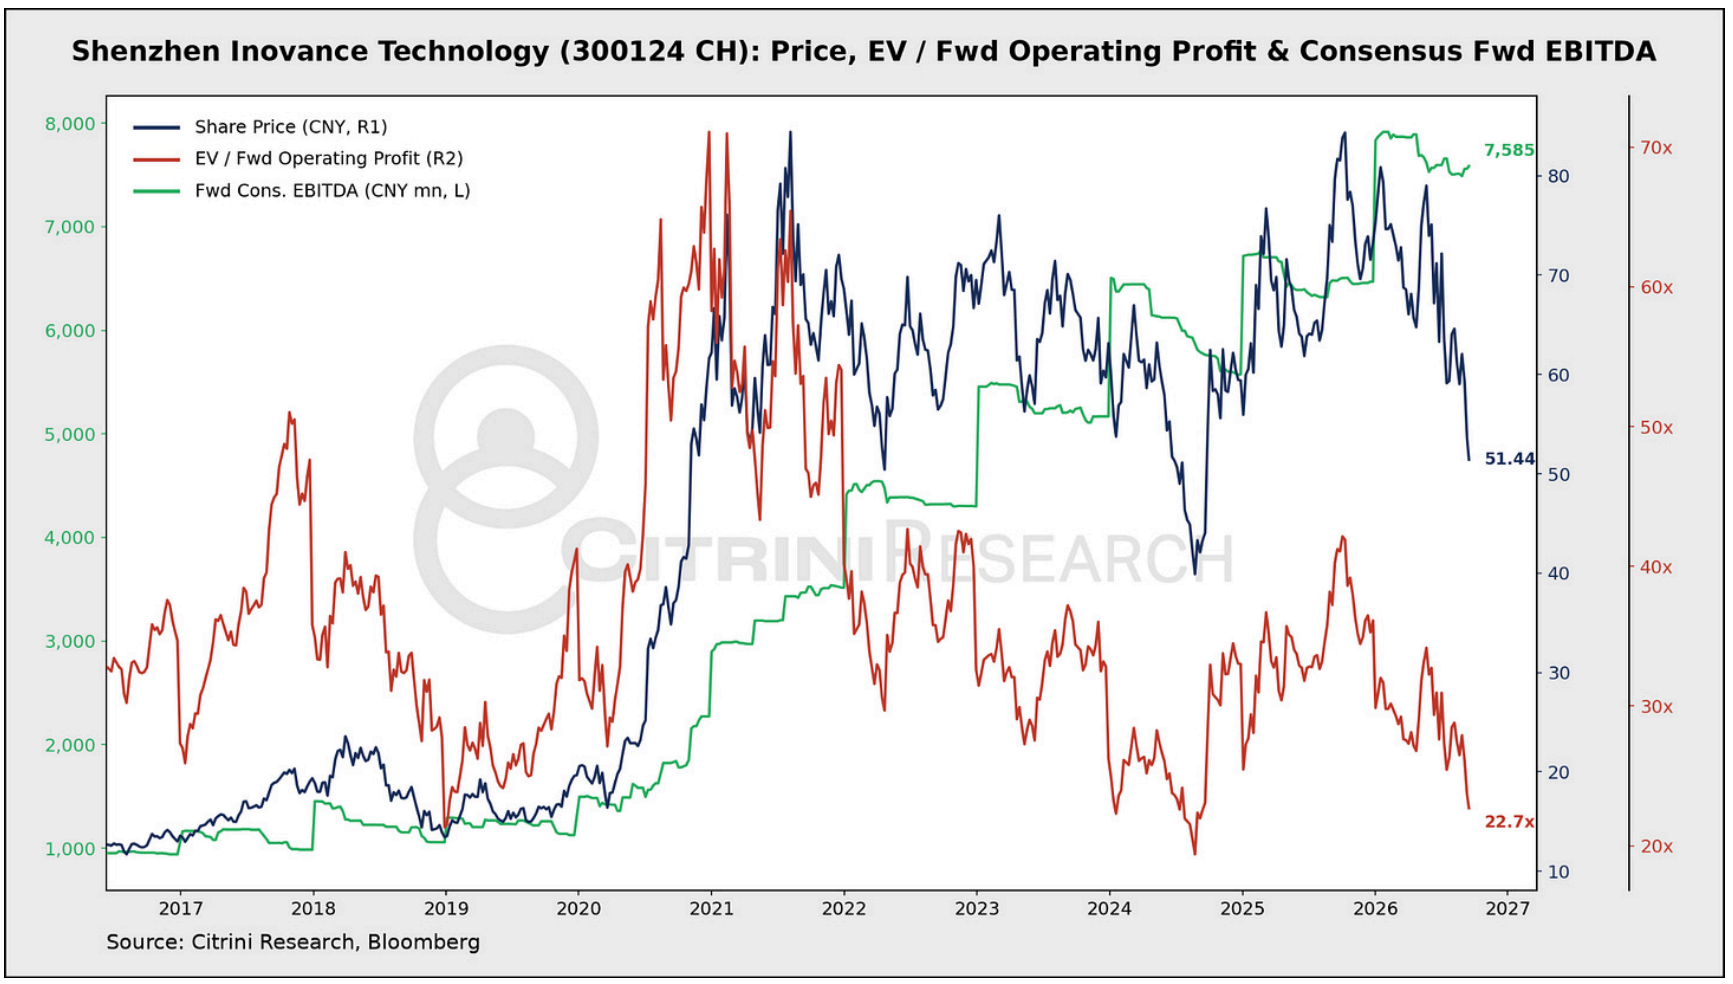


Both stocks are pretty volatile and narratively exposed, but very few other companies provide such concentrated exposure to the theme.

这两只股票的波动性都相当大，且极易受市场叙事影响，但很少有其他公司能提供如此集中于这一主题的敞口。

In China, Shenzhen Inovance (300124 CH) is the name that stands out to me. Historically an auto player, the company has pivoted heavily towards the robotics ecosystem and there is significantly more hardware production currently happening on that side of the Pacific.

在中国，**汇川技术（300124 CH）**是我眼中较为突出的标的。该公司历史上是一家汽车领域的企业，如今已大幅转向机器人生态系统，而且当前太平洋彼岸的硬件生产规模明显更为庞大。

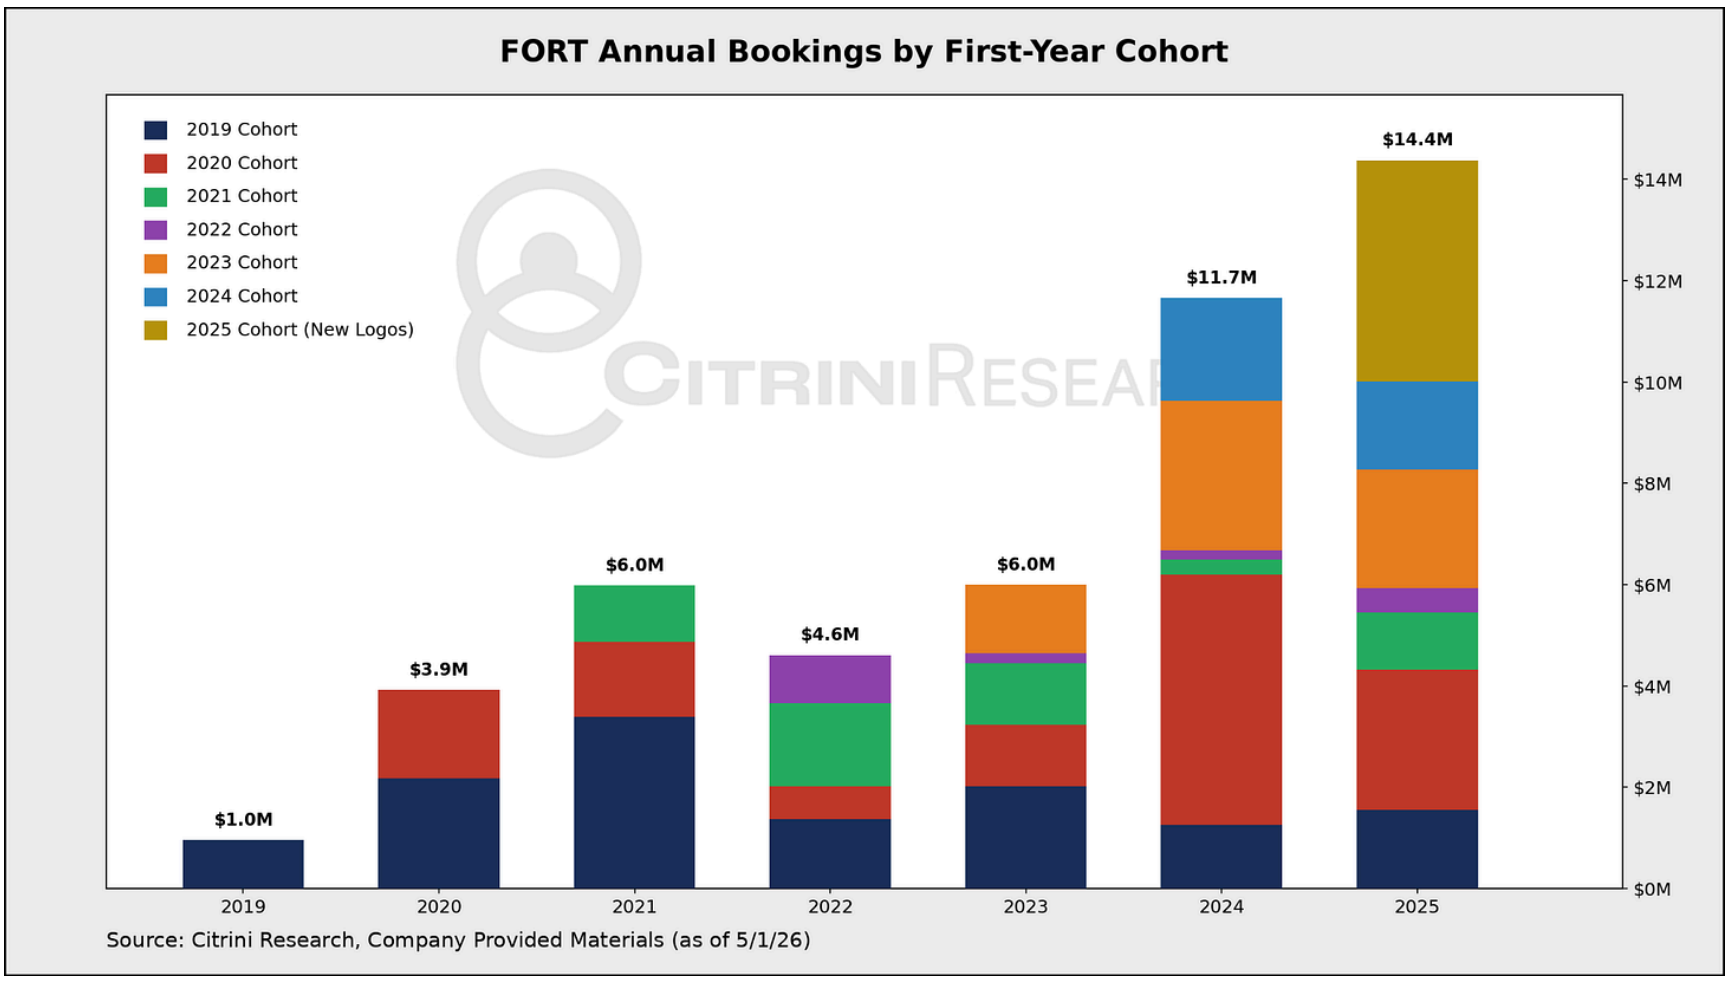


Revenue and earnings growth over 20% through the COVID lockdown, real estate crisis, and the absolutely brutal auto market is downright impressive. And given that it's trading at the very low end of its valuation range, it seems like an obvious winner.

在新冠封控、房地产危机以及极度惨烈的汽车市场期间，营收和利润增长仍超过 20%，实在令人印象深刻。考虑到该股目前的交易价格处于其估值区间的最低端，它看起来是一个显而易见的赢家。

Here's a newbie worth exploring – FORT Robotics (NTWO US) – which has a merger agreement to go public via SPAC**.** At the risk of massively oversimplifying the business, Fort installs sensor boxes that you put in an industrial robot that ensure that it shuts down before it hits something or someone.

这里有一家值得研究的新公司——FORT Robotics（NTWO US）——该公司已签署合并协议，将通过 SPAC 方式实现上市。尽管这一描述大幅简化了其业务：Fort 在工业机器人中安装传感器盒子，用于确保机器人在撞到物体或人员之前停机。



That's important for safety, of course, but even more importantly regulatory certification. Without having a certified system like this, robots in factory settings must operate in contained cages without any humans – which is quite limiting. The nice thing here is that it's well positioned for early deployments (i.e. wheeled base industrial robots), though on a longer timeframe, there's a risk that established players achieve their own certified systems and cut out FORT.

这对安全固然重要，但更重要的是监管认证。没有这类经过认证的系统，工厂环境中的机器人就必须在封闭围栏内运行、不得有人员在场——这是相当大的限制。这里的好处在于，它对早期部署（即轮式底盘工业机器人）的定位相当有利；不过在更长的时间维度上，存在既有厂商开发出各自认证系统、从而将 FORT 排除在外的风险。

What about the nervous system – compute and electronics?

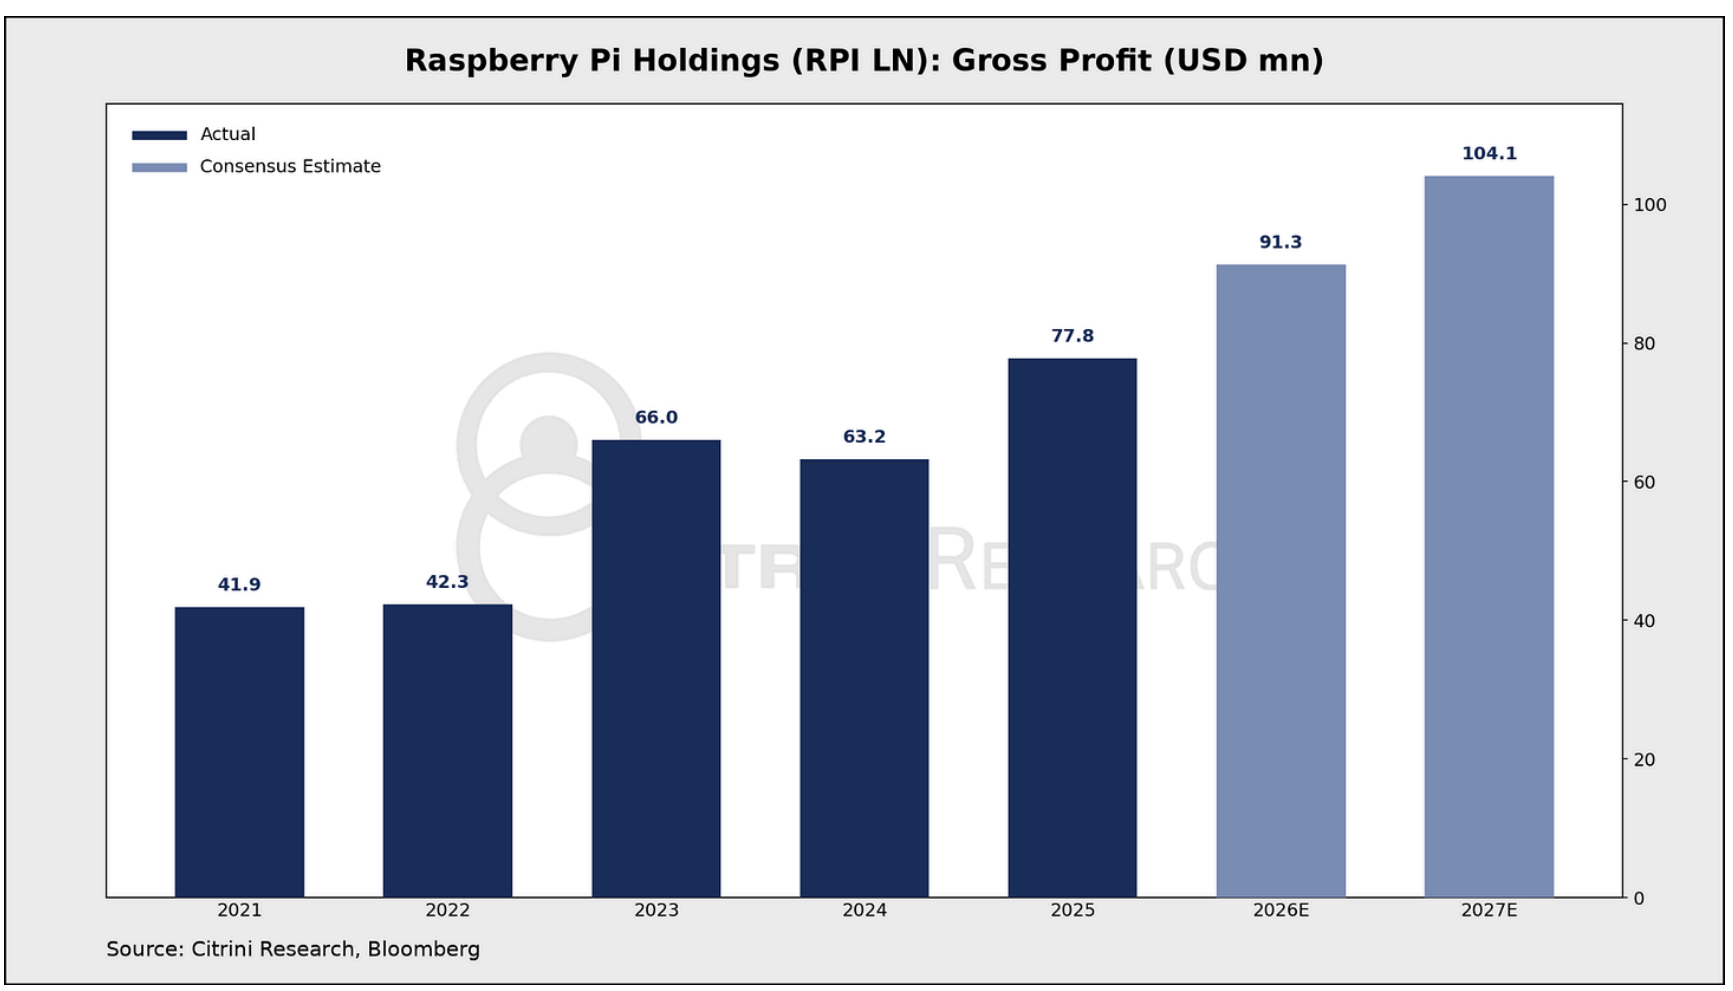
那么"神经系统"——算力与电子部件呢？

It goes without saying that **Nvidia (NVDA US)** is an omnipresent and critical player across the field. Its Jetson processor is the ubiquitous on-robot computing chip which is just one part of the NVDA Physical ecosystem. And arguably, NVDA's leg into robotics, while tiny relative to its datacenter business, should advantage the company relative to its peers in any sort of downturn.

毋庸置疑，**Nvidia（NVDA US）**是整个领域无处不在且至关重要的参与者。其 Jetson 处理器是机器人上普遍使用的机载计算芯片，而这只是 NVDA Physical 生态的一部分。可以说，NVDA 在机器人领域的布局虽然相对于其数据中心业务体量很小，但在任何形式的下行周期中，都应使公司相对于同行更具优势。

But at the risk of jumping the shark, I'm intrigued enough about **Raspberry Pi (RPI LN).** Frankly, I didn't know the company before its "open claw" stock pop back in February, which gives it a bit of a meme vibe. But in both the Asimov breakdown and some in-person discussions, the name came up.

尽管这么说可能有点扯远，**Raspberry Pi（RPI LN）**确实让我颇感兴趣。坦白讲，在今年二月其股价因"open claw"概念大涨之前，我并不了解这家公司，这也让它多少带了点 meme 股的味道。不过，无论是在 Asimov 的行业拆解中，还是在一些线下交流中，这个名字都被提及。



The company provides low-cost boards and chips for a variety of edge applications with OEMs and hobbyists. While most of these edge-devices don't have the compute for super high-level reasoning, they are quite relevant in vision applications generally. The company is clearly seeing a tailwind from the sector and is building on its product line.

该公司面向各类边缘应用，为 OEM 厂商和爱好者提供低成本开发板和芯片。尽管大多数此类边缘设备并不具备运行超高阶推理所需的算力，但它们在视觉应用领域整体上颇为重要。该公司显然正受益于这一行业的顺风，并在此基础上持续扩充产品线。

> *"AI HAT+ 2 [...]. It builds on the success of our first-generation AI accelerator products that we launched in 2024, and for the first time allows accelerated execution of modern AI workloads like Vision Transformer models, Large Language models and Vision Language models[...] When we say an AI flavor, what we tend to mean is visual AI, vision AI, generally the recognition and classification of objects in still and video images coming into the Raspberry Pi.*
>
> *The workloads that people run, our largest OEM customers today are AI customers. They are running vision AI applications on Raspberry Pi. And the workloads that they're running today on Raspberry Pi would have required a large Nvidia GPU a decade ago."*

> *"AI HAT+ 2 [...]。它在我们于 2024 年推出的第一代 AI 加速器产品取得成功的基础上更进一步，并首次支持对 Vision Transformer 模型、大语言模型和视觉语言模型等现代 AI 工作负载进行加速执行[...]。当我们说"AI 版本"时，我们通常指的是视觉 AI，也就是对进入 Raspberry Pi 的静态图像和视频中的物体进行识别与分类。*
>
> *从人们实际运行的工作负载来看，我们目前最大的 OEM 客户正是 AI 客户。他们在 Raspberry Pi 上运行视觉 AI 应用。而这些客户如今在 Raspberry Pi 上运行的工作负载，放在十年前需要一块大型 Nvidia GPU 才能承担。"*

They're not overtaking Jetson anytime soon, and there is risk that memory prices this year will kill sales volume, but the company's growth through the electronics cycle speaks to a secular growth story, that could find more momentum through robotics – simpler applications, open source kits, and the growing interest in research and tinkering.

它们短期内还不会超越 Jetson，而且今年内存价格也存在扼杀销量的风险，但该公司贯穿整个电子周期的增长，说明这是一个长期成长故事，未来还可能借助机器人领域获得更多动能——更简单的应用、开源套件，以及对研究探索和动手实践日益浓厚的兴趣。

If you'd like to punt on some semis, robots do use a lot of analogs. **Allegro Microsystems (ALGM US)** and **STMicroelectronics (STM US)** have some of the largest disclosed exposure to robotics, while companies like **Texas Instruments (TXN US)** have a higher quality but far more diversified play exposure to the tightening of the analog market.

如果你想押注一些半导体标的，机器人确实会用到大量模拟芯片。**Allegro Microsystems（ALGM US）与意法半导体（STM US）在披露口径下拥有机器人业务最大的敞口之一，而德州仪器（TXN US）**这类公司则提供了质量更高、但分散程度大得多的敞口，同样能够博弈模拟芯片市场趋紧带来的机会。

Humanoid OEMs

人形机器人 OEM 厂商

Getting straight to the point. If Humanoid OEMs are the moonshot pitch, do you take a swing?

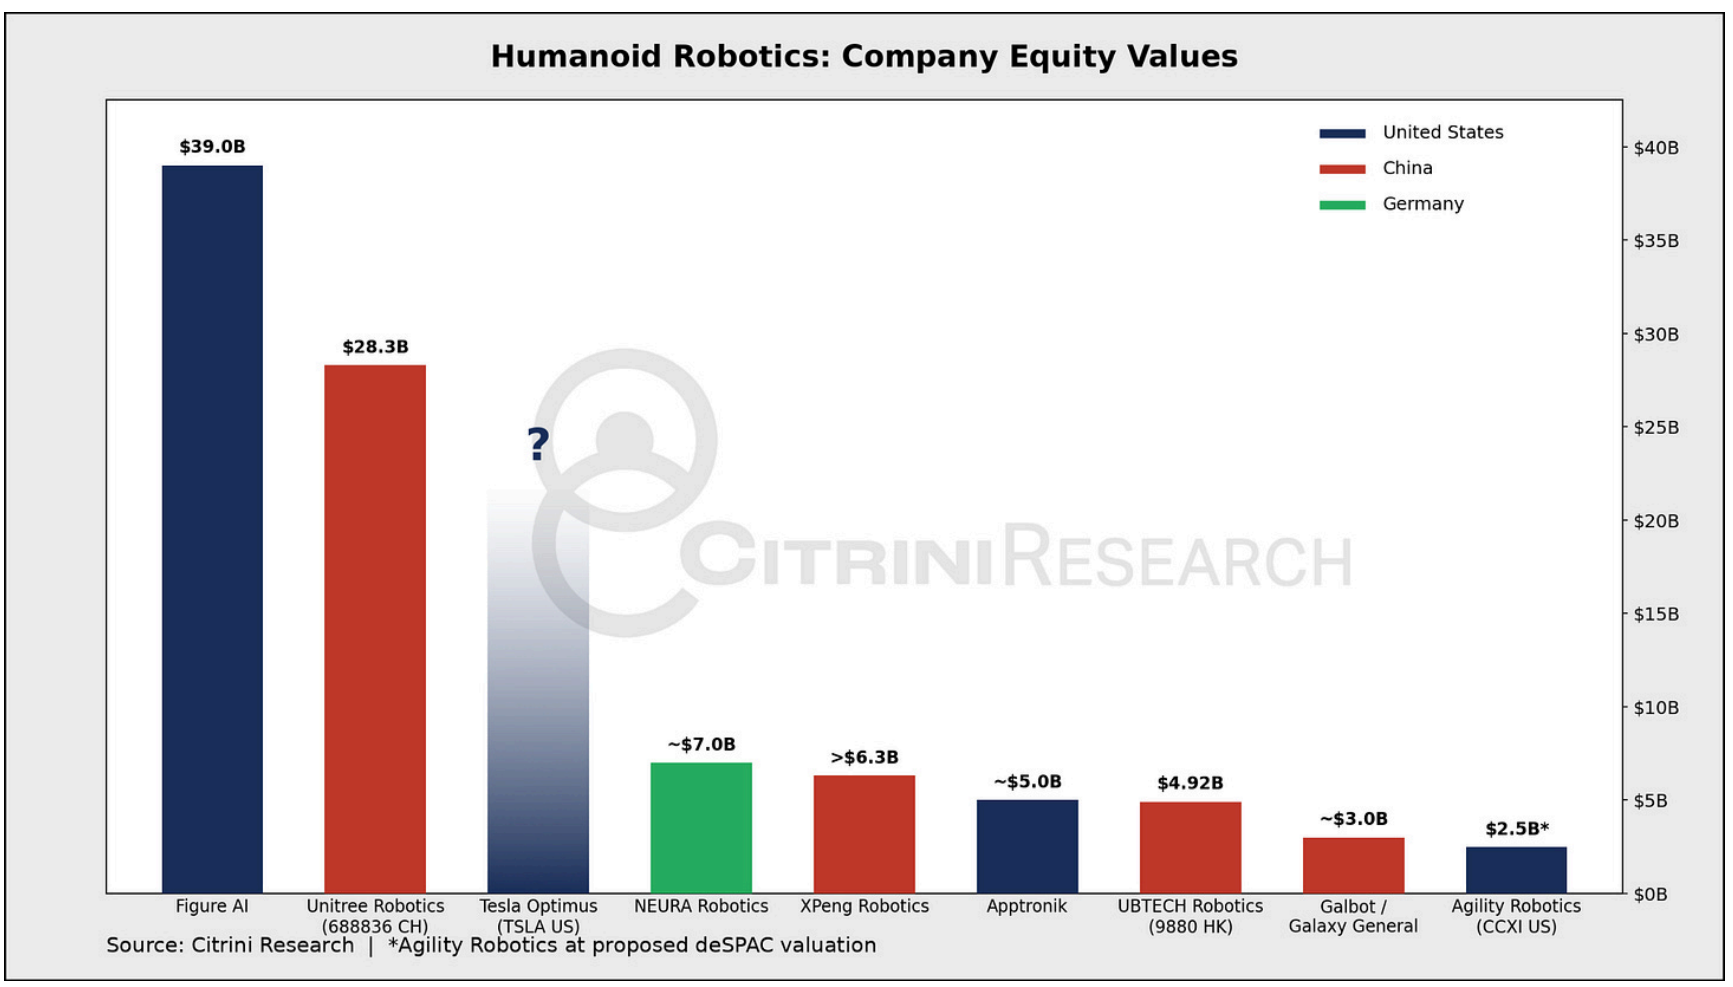
开门见山地说。如果人形机器人 OEM 厂商是那种"登月"式的豪赌故事，你会不会出手一搏？

It won't be a winner-take-all market globally – at a minimum, the US and China markets will likely be walled off from each other. Other blocs and countries will likely pursue their own out of sovereign interest and "dual use" concerns, but are likely behind in development and will be more friendly to imports.

这不会是一个全球赢家通吃的市场——至少，美国和中国市场很可能相互隔离。其他阵营和国家出于主权利益与"军民两用"考量，很可能也会追求自主发展，但其发展水平可能落后，因而对进口会更为友好。

First, here's what the market thinks. At least directionally, I'm not too far off.

首先，来看看市场的看法。至少在方向上，我的判断与市场相去不远。



I agree that Figure and Unitree are in the pole positions of their geographies, though with very different strategies.

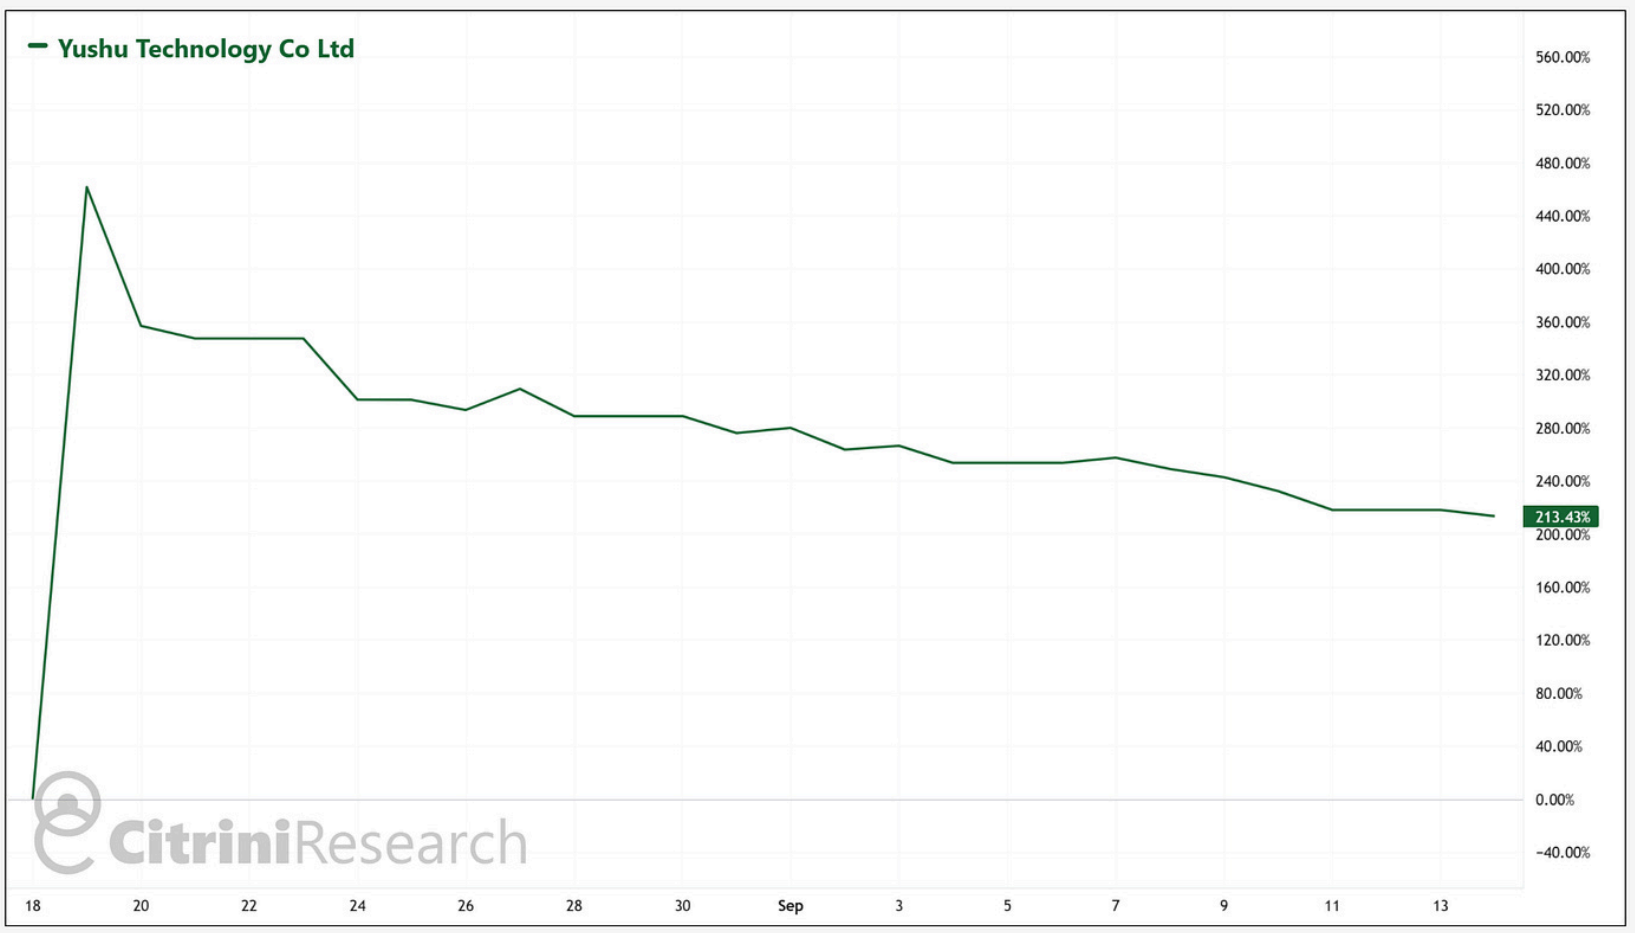
我同意 Figure 和 Unitree 分别在各自地域市场处于领跑位置，不过两者的策略截然不同。

Unitree has production volume, an open-source ecosystem for programming, and a somewhat astounding net margin of 35% with positive free cash flow at 15% of revenue.

Unitree 拥有量产规模、面向编程的开源生态，以及相当惊人的 35% 净利率，同时自由现金流为正、达到营收的 15%。

After its enormous (and some might claim manufactured) IPO pop last month, **Yushu Technology (688836 CH)** aka Unitree has traded in a straight line down. I chalk this up to simple reversion combined with clear signals from authorities to tamp down *"enthusiasm"* market-wide*.* Again, buying the stock directly will be a challenge for most people and I'd prefer to see a little bit of a bottoming pattern first, but given its recent demonstration in brain development, cash flow positivity and first mover advantage, I think it ultimately is a buy. ***The question is whether it can keep its lead in the land of involution.***

在上个月经历了巨大的（有人可能会说是人为制造的）IPO 暴涨之后，**宇树科技 (688836 CH)**（即 Unitree）一路直线下跌。我将其归因于单纯的均值回归，再加上监管层明确释放信号，要打压全市场的*"热情"*。再强调一次，对多数人来说，直接买入这只股票会有难度，我也更愿意先看到一点筑底形态，但考虑到其近期在大脑研发方面的展示、正现金流以及先发优势，我认为它最终值得买入。***问题在于，它能否在这片内卷之地保住领先优势。***



Source: Citrini Research

Meanwhile Figure, which raised at $39 billion, is effectively pre-revenue and is sinking billions of dollars of equity funding into its own production, training and compute. *And yet, I'd buy it…*

与此同时，以390亿美元估值完成融资的Figure实际上尚未产生营收，正将数十亿美元的股权资金投入自身的生产、训练和算力建设。*然而，我还是会买它……*

Last year, I got to hang with the humanoids in Hangzhou. This year, on a sidewalk in San Jose, *I entered the future.*

去年，我在杭州与人形机器人有过亲密接触。今年，在圣何塞的一处人行道上，*我踏入了未来*。

It's hard to describe – standing face to faceplate, flesh to woven mesh.

很难形容——面对面板而立，肉身对着编织网面。

Seeing is believing.

眼见为实。

I'm not being sappy, it was a genuinely unsettling and enthralling experience. Maybe it was how the 03's shape uncannily mirrored my own, its sleek and thoughtful design, or the way it tilted its head as it tracked me with its gaze.

我并非在煽情，那确实是一次令人不安又深深着迷的体验。也许是03的身形与我自己惊人地相似，也许是它流畅而考究的设计，又或者是它注视着我、随之微微歪头的样子。

There wasn't someone hiding in the bushes with a remote. This was a thing that looked like me and walked like me and looked me in the eyes, all on its own. Ignoring the $200,000 or so it took to build it, and the relatively limited useful acts it can perform, it felt to me like an announcement that something had arrived – that it was just getting started.

并没有人藏在树丛里拿着遥控器操控。这是一个长得像我、走路像我、还直视我双眼的东西，而且完全自主。先不提它约20万美元的制造成本，也不提它目前能执行的有用动作还相当有限——在我看来，这更像是一种宣告：某种东西已经到来，而这才刚刚开始。

It's probably a binary bet with asymmetric outcomes – if the company outright fails or drags along without showing progress, the $39 billion pre-revenue valuation will certainly flounder. But combine its proven lead in deployment, vertical integration, proprietary data advantage, and intense compute ambitions, and it is, in my opinion, the closest thing to "the humanoid revolution" that we've been promised.

这大概是一笔结果不对称的二元押注——如果公司彻底失败，或者迟迟拿不出进展，390亿美元的零营收估值必将难以为继。但把其在部署上已被验证的领先地位、垂直整合、专有数据优势和强烈的算力雄心结合在一起，在我看来，它就是我们一直被许诺的"人形机器人革命"最接近的现实。

As for other easy public expressions, I actually like **Xpeng (XPEV US)** because of the chart vs. Tesla that I showed above.

至于其他便于公开表达的观点，我其实挺喜欢**小鹏汽车（XPEV US）**，原因是我上文展示的那张对比特斯拉的走势图。

Optimus is the major dark horse candidate for global leader. It was the major pioneer in this modern humanoid surge and prior releases suggested it was among the most capable. The reason it hasn't gotten much attention in this piece is because they've gone dark for a long time.

Optimus 是全球领跑者位置上最主要的黑马候选。它在这一代人形机器人浪潮中是最主要的先驱，此前的发布也表明它属于能力最强的一批。之所以在本文中没有得到太多关注，是因为他们已经很长时间没有公开动向了。

There's been little to even judge progress and the rumor mill is pretty universally negative. But even if it's been a rough year, it would be insane to count Tesla out. From its infrastructure, to its manufacturing experience, to autonomy talent, Tesla has a highly attractive setup on paper.

几乎没有任何可用于评估进展的素材，而坊间传闻也几乎清一色偏负面。但即便这一年不太顺利，就此把 Tesla 排除在外也是不理智的。从基础设施、制造经验到自动驾驶人才，Tesla 在纸面上具备极具吸引力的条件。

The problem though is you're buying a trillion dollar EV company with a nascent humanoid arm. XPeng is another EV company with a nascent arm but ~100x cheaper. I would bet Optimus is in a better position than Iron today, but just look at history. Xpeng's EV business is rapidly expanding in part by taking share from Tesla.

但问题在于，你买入的是一家万亿美元市值的电动车公司，附带一个尚处起步阶段的人形机器人业务。XPeng 是另一家拥有起步阶段机器人业务的电动车公司，但估值便宜约 100 倍。我愿意打赌 Optimus 目前的位置比 Iron 更好，但看看历史就知道了。小鹏的电动车业务正在快速扩张，部分原因正是从 Tesla 手中夺取份额。

As for **Agility (CCXI US)**...Frankly I'm unimpressed with their tech, including a disappointing Digit 5 release yesterday, and don't think they have a good chance at beating the leading contenders… if you can't hit the moonshot, I'm not inclined to buy. That being said, if people start getting excited about humanoids, *it's the most pureplay expression in US stocks.*

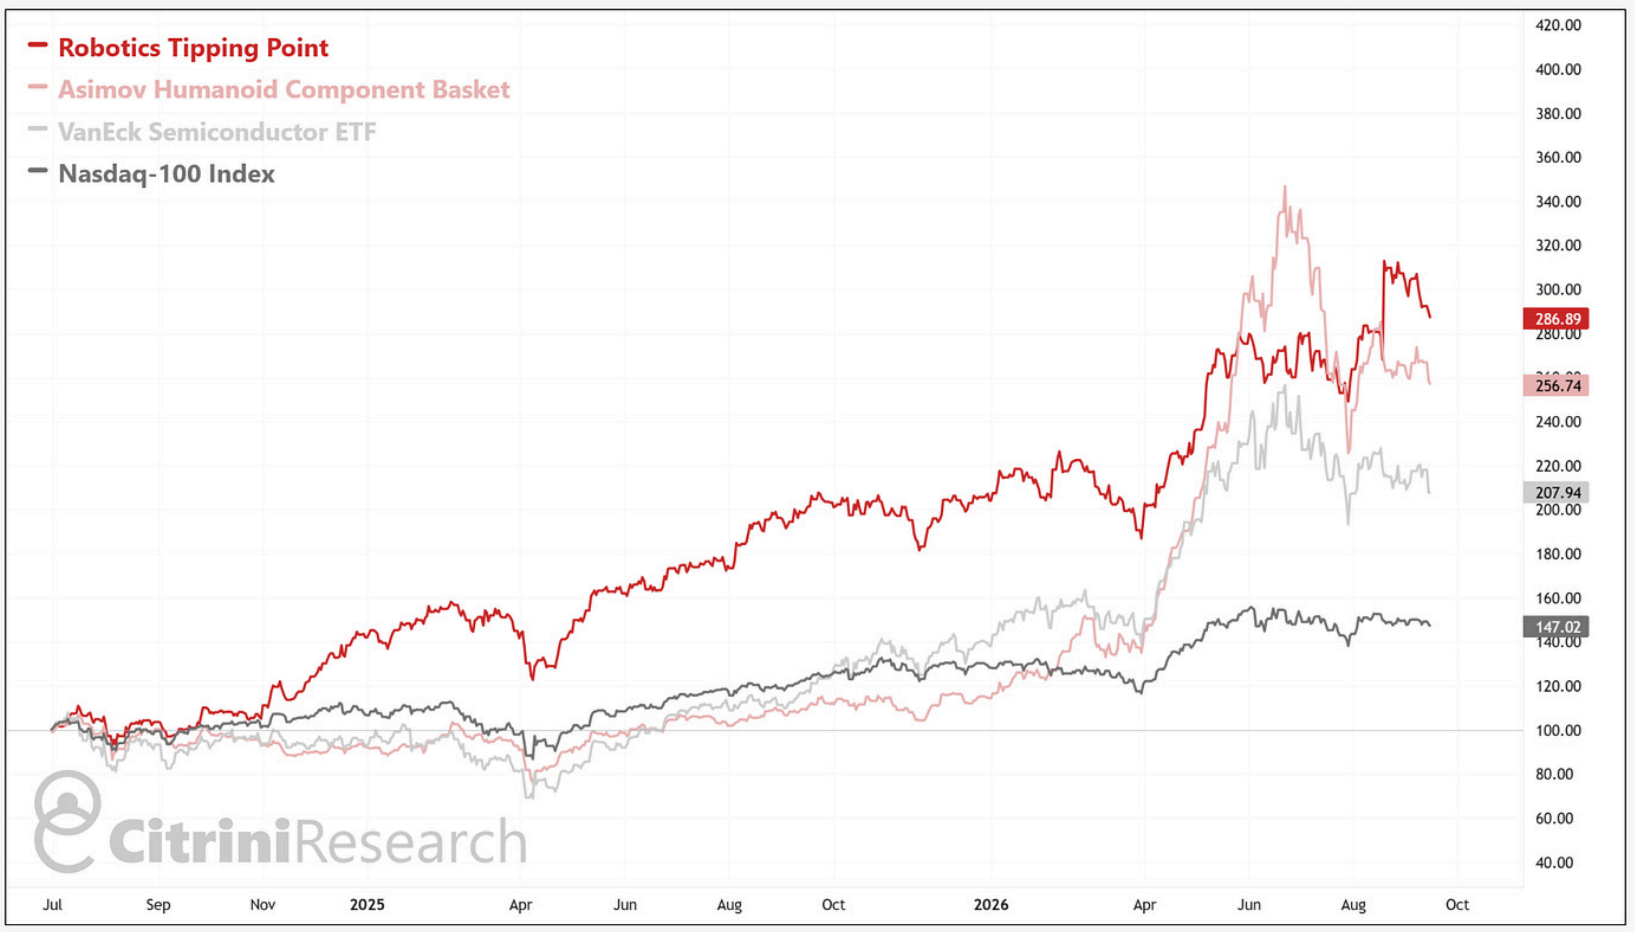
至于 **Agility (CCXI US)**……坦率说，我对他们的技术并不看好，包括昨天令人失望的 Digit 5 发布，也不认为他们有多大机会击败领先的竞争者……如果够不到那种登月级的目标，我不倾向于买入。话虽如此，如果市场开始对人形机器人感到兴奋，*它是美股中最纯粹的标的。*

The Basket:

篮子：

Here it is, the Robotics Tipping Point basket. For those of you who are frustrated with a basket approach and just want the highest-conviction single names, I've weighted the names individually according to my conviction – extrapolate accordingly.

它来了：Robotics Tipping Point 篮子。对于那些对篮子方式感到失望、只想要信念度最高的单一个股的读者，我已经按照自己的信念度对各标的单独赋权——可据此自行推演。



Source: Citrini Research

Positions in basket — as of 2026-09-15 at market close

Ticker	Exchange	Security Type	Industry	Weight	1D Return
PRLB PROTO LABS INC	US	Equity Common Stock	Industrial Metal Fabricate/Hardware | Metal Processors&Fabrica	11.00%	+1.11%
XMTR XOMETRY INC-A	US	Equity Common Stock	Industrial Metal Fabricate/Hardware | Metal Processors&Fabrica	10.00%	+3.00%
300124 SHENZHEN INOVANCE TECHNOLO-A	CH	Equity Common Stock	Industrial Electronics | Instruments-Controls	8.00%	-4.01%
HSAI HESAI GROUP	US	Equity ADR	Industrial Electronics | Electronic Measur Instr	8.00%	-3.34%
6954 FANUC CORP	JP	Equity Common Stock	Industrial Machinery-Diversified | Industrial Automat/Robot	7.00%	-1.52%
6324 HARMONIC DRIVE SYSTEMS INC	JP	Equity Common Stock	Industrial Machinery-Diversified | Industrial Automat/Robot	5.00%	-0.14%
6103 OKUMA CORP	JP	Equity Common Stock	Industrial Machinery-Diversified | Machinery-General Indust	5.00%	-1.40%
6141 DMG MORI CO LTD	JP	Equity Common Stock	Industrial Hand/Machine Tools | Mach Tools&Rel Products	5.00%	+0.17%
ADSK AUTODESK INC	US	Equity Common Stock	Technology Software | Computer Aided Design	5.00%	-1.06%
PTC PTC INC	US	Equity Common Stock	Technology Software | Applications Software	5.00%	-2.85%
XPEV XPENG INC - ADR	US	Equity ADR	Consumer, Cyclical Auto Manufacturers | Auto-Cars/Light Trucks	5.00%	-4.48%
OUST OUSTER INC	US	Equity Common Stock	Technology Semiconductors | Electronic Compo-Semicon	4.00%	-2.00%
688836 YUSHU TECHNOLOGY CO LTD-A	CH	Equity Common Stock	Industrial Machinery-Diversified | Industrial Automat/Robot	4.00%	-0.12%
RPI RASPBERRY PI HOLDINGS PLC	LN	Equity Common Stock	Technology Computers | Computers	4.00%	-1.13%
NTWO NEWBURY STREET II AC-CLASS A	US	Equity Common Stock	Diversified Holding Companies-Divers | Specified Purpose Acquis	3.00%	-0.46%
NVDA NVIDIA CORP	US	Equity Common Stock	Technology Semiconductors | Electronic Compo-Semicon	2.00%	+0.57%
STM STMICROELECTRONICS NV-NY SHS	US	Equity NY Reg Shrs	Technology Semiconductors | Electronic Compo-Semicon	2.00%	+0.29%
TXN TEXAS INSTRUMENTS INC	US	Equity Common Stock	Technology Semiconductors | Electronic Compo-Semicon	2.00%	+0.00%
ALGM ALLEGRO MICROSYSTEMS INC	US	Equity Common Stock	Technology Semiconductors | Semiconductor Equipment	2.00%	-1.74%
CCXI CHURCHILL CAPITAL CORP-CL A	US	Equity Common Stock	Diversified Holding Companies-Divers | Specified Purpose Acquis	2.00%	+0.42%
108490 ROBOTIS CO LTD	KS	Equity Common Stock	Industrial Machinery-Diversified | Industrial Automat/Robot	1.00%	-1.25%

Secondly, if you'd like the torque (and heavy semi exposure, here is the high-beta Asimov Humanoid Component Basket.

其次，如果您想要更高的弹性（以及较重的半导体敞口），这里是高贝塔的 Asimov 人形机器人零部件篮子。

Positions in basket — as of 2026-09-15 at market close

Ticker	Exchange	Security Type	Industry	Weight	1D Return
ADI ANALOG DEVICES INC	US	Equity Common Stock	Technology Semiconductors | Electronic Compo-Semicon	7.69%	+0.09%
009150 SAMSUNG ELECTRO-MECHANICS CO	KS	Equity Common Stock	Industrial Electronics | Electronic Compo-Misc	7.69%	-2.73%
STM STMICROELECTRONICS NV-NY SHS	US	Equity NY Reg Shrs	Technology Semiconductors | Electronic Compo-Semicon	7.69%	+0.29%
2327 YAGEO CORPORATION	TT	Equity Common Stock	Industrial Electronics | Capacitors	7.69%	-1.76%
MCHP MICROCHIP TECHNOLOGY INC	US	Equity Common Stock	Technology Semiconductors | Electronic Compo-Semicon	7.69%	+0.01%
6806 HIROSE ELECTRIC CO LTD	JT	Equity Common Stock	Industrial Electronics | Electronic Connectors	7.69%	-0.82%
6981 MURATA MANUFACTURING CO LTD	JT	Equity Common Stock	Industrial Electronics | Electronic Compo-Misc	7.69%	+2.74%
TXN TEXAS INSTRUMENTS INC	US	Equity Common Stock	Technology Semiconductors | Electronic Compo-Semicon	7.69%	+0.00%
6752 PANASONIC HOLDINGS CORP	JT	Equity Common Stock	Consumer, Cyclical Home Furnishings | Audio/Video Products	7.69%	+0.50%
6616 TOREX SEMICONDUCTOR LTD	JP	Equity Common Stock	Industrial Electronics | Electronic Parts Distrib	7.69%	-0.26%
002655 GETTOPACOUSTIC CO LTD-A	CH	Equity Common Stock	Consumer, Cyclical Home Furnishings | Audio/Video Products	7.69%	-3.82%
6834 EVER OHMS TECHNOLOGY CO LTD	TT	Equity Common Stock	Industrial Electronics | Capacitors	7.69%	-1.41%
RPI RASPBERRY PI HOLDINGS PLC	LN	Equity Common Stock	Technology Computers | Computers	7.69%	-1.13%

Thank you for reading.

感谢阅读。

This article is for informational purposes only and does not constitute investment advice. By accessing this material, you agree to our Terms of Service.

本文仅供参考，不构成投资建议。访问本内容即表示您同意我们的服务条款。