Value migration in embodied AI: from whole machines to sensing components

具身智能的价值迁移：从整机到感知零部件

庚辛研究院 / Glacier Institute*†

From The Musk Institute

Glacier Institute（庚辛研究院） | Working Paper No. GI-WP-2026-P3-EN | Date: 2026-09-22 | Language: English, with Chinese title and abstract

Cite as:

Glacier Institute (2026). Value migration in embodied AI: from whole machines to sensing components. Glacier Institute Working Paper GI-WP-2026-P3-EN. https://glacier.mba/research/GI-WP-2026-P3-EN.html

庚辛研究院（2026）. 《具身智能的价值迁移：从整机到感知零部件》. 庚辛研究院通讯论文 GI-WP-2026-P3. https://glacier.mba/research/GI-WP-2026-P3.html

Key Takeaways

1. **The whole machine answers "can it"; the components answer "is it stable, is it cheap enough, can it be changed".**
2. **The more transferable the upper layer's capability, the less value stays in the upper layer.**
3. **The layer that creates value and the layer that holds bargaining power are never automatically the same layer.**
4. **Meeting the mean is the prototype's job; narrowing the distribution is the production line's.**

Abstract

Public discourse on embodied AI is organised almost entirely around the whole machine: whether a robot can walk, grasp, or perform a coherent demonstration. This paper argues that such a framing cannot answer a more consequential question — at which layer of the stack durable bargaining position will settle.

We advance a single claim: value in embodied AI is migrating from the whole machine toward the perception and actuation component layer, and within perception, tactile sensing is the last to be solved and the hardest to route around. The paper predicts no prices for any firm; it argues why value should accumulate at this layer, and under what conditions the argument fails.

Three mechanisms are stated as falsifiable propositions. First, a manufacturing experience curve: for devices of this kind, unit cost declines log-linearly in cumulative output, and what actually narrows is the yield distribution, not merely the mean [6][7][8][11]. Second, a data loop: physical-interaction data cannot be synthesised at will and scales with deployed robot-hours rather than compute, so the data stock is tied to whoever ships hardware [17][19][20]. Third, full-stack traceability: whether a field failure can be traced along the signal chain back to device design determines whether the improvement cycle is measured in device generations or in model iterations [26].

Section 7 states four conditions under which the claim does not hold: vertical integration by whole-machine makers [22], a competence-destroying shift in sensing mechanism [24][25], interface standardisation commoditising the component [23], and applications that simply do not require fine tactile resolution.

Keywords: embodied AI; value migration; tactile sensing; experience curve; data flywheel; industry architecture

JEL Classification: O33 (Technological Change: Choices and Consequences; Diffusion Processes), L63 (Microelectronics; Computers; Communications Equipment),

* **Disclosure:** Glacier Institute is the research arm of Glacier Capital, and this paper is issued under the name of Glacier Institute. In the course of its business, Glacier Capital acts as financial adviser to a number of technology companies and invests its own capital in some of them; such relationships may overlap with the industries discussed here. This paper does not concern any specific mandate and uses no non-public information; all company, industry and data references are drawn from public sources and cited individually. The authors received no third-party compensation for this paper.

† **Disclaimer:** This is a methodological working paper and represents the authors' analysis at the time of writing only. It does not constitute investment advice, nor an offer or solicitation of an offer for any security, fund interest or other instrument, nor a commitment or forecast regarding the valuation, financing outcome or investment return of any company. It has not been peer reviewed and may be revised in later versions.

L22 (Firm Organization and Market Structure), G24 (Investment Banking; Venture Capital)

摘要 (Chinese abstract)

具身智能的公开叙事几乎都挂在整机上：一台机器能不能走、能不能抓、能不能演示一段连贯动作。本文认为，这条叙事线回答不了一个更要紧的问题——这个产业长期的议价位置会沉在哪一层。

本文提出并论证一个判断：具身智能的价值正在从整机 / 本体，向感知与执行零部件一侧迁移；而在感知这一侧，触觉是最迟被解决、也最难被绕开的一环。 本文不预测任何公司的价格，只论证价值为什么会向这一层沉淀，以及这个论证在什么条件下失效。

论证由三条机制构成，每一条都写成可被检验、也可被证伪的命题：量产经验曲线——第一百万颗与第一颗的差别不是数量，是良率分布的宽度；数据闭环——物理交互数据的获取成本不随算力下降，只随真机部署的小时数下降；全栈可追溯——现场失败样本能否沿信号链回到器件设计，决定了改进周期是以器件为单位还是以模型为单位。

第 7 节专门讨论这个判断不成立的四种情形：整机厂的垂直整合、使在位经验归零的感知机制更替、接口标准化对零部件的商品化，以及场景本身并不需要那么细的触觉。

本文引用庚辛资本与庚辛研究院已公开的表述时一律原样照搬，不作改写；其余判断由本文自负。分析框架，非投资建议。

Keywords (Chinese): 具身智能；价值迁移；触觉传感；经验曲线；数据闭环；产业架构

1. The Problem: The Whole-Machine Narrative Cannot Answer "Which Layer Value Settles In"

Over the past two years, public information on embodied AI has been dense but uniform in structure: almost every visible signal hangs on the whole machine. The more coherent a prototype's motion, the further it travels; the yield fluctuation of the nth batch of devices on a production line has no propagation value.

Glacier Capital's statement on this is direct:

> "Why? Because the two ask different questions. A demo answers one: can it be done. The line answers three: is it stable under harsh conditions, do the numbers work, and does the data feed back into the model — the data flywheel. A machine that cannot answer those is an exhibit."（「演示只回答一个：能不能。产线要回答三个：严苛工况下稳不稳，账算不算得过来，数据能不能反哺模型 (data flywheel)。答不上来的机器，是展品。」）
> — Glacier Capital, *Watch Fewer Launch Events, Go Look at the Line*（《少看发布会，去看产线》）[3]

That sentence is usually read as a due-diligence discipline. This paper reads it as a proposition about industry structure: **if the whole machine only answers "can it", then the question it answers is one that will be redone again and again, with a bar that falls every year.** The marginal difficulty of a demonstration is falling — public robot-learning models and datasets are turning "getting a machine to complete a sequence of motions" into a transferable capability [17][18][19][21]. The three questions the production line must answer, by contrast, each grow in time and materials, and none gets cheaper when a paper is published.

The question is therefore rewritten: on a value chain where the upper layer's capability becomes ever more transferable and the lower layer's ever more dependent on accumulation, which layer does value settle in?

Glacier Capital has already given a structurally identical answer on another industry line:

> "In other words, the entry point of AI infrastructure is moving down: from the chip down to power electronics, and down again to energy."（「也就是说，AI 基础设施的入口在下移：从芯片下移到电力电子，再下移到能源。」）
>
> — Glacier Capital, *Below Compute Is Electricity*（《算力之下是电力》）[4]

This paper argues that a downward shift in the same direction is happening on the embodied AI chain, and that it lands on perception and actuation components. It also sets out to make clear that the shift is not driven by the intuition that "lower is harder", but by three mechanisms that can each be tested separately.

> The more transferable the upper layer's capability, the less value stays in the upper layer.

2. Definition: What This Paper Means by "Value Migration", and What It Does Not

In this paper, "value migration" refers to one thing only: **on an industry chain, which layer's capability is harder to replicate, and therefore holds the stronger bargaining position in the long-run division of labour.**

Three boundaries, stated up front.

图 1 / Exhibit 1 Which layer value settles in along the embodied-AI chain. The figure is this paper's own synthesis; the three layers and the contrast between transferable capability above and cumulative capability below come from Sections 1 and 2, and the parallel downward shift on the compute chain is quoted in Section 1. No company name, price or time forecast appears, and the arrow marks the direction of migration, not chronological order.

First, this paper does not discuss any firm's price, financing arrangements or future capital-market path. It discusses the structure of the division of labour, not asset pricing.

Second, this paper does not claim the whole machine is unimportant. On the contrary, Glacier Capital's public statements hold that the robot body is the entry point for data:

> "Hardware sets the ceiling on what the model can do." (「硬件决定模型能力的上限。」)
>
> "Whoever can build the body and get it deployed holds the tap." (「谁能把本体造出来并部署下去，谁就握着数据的水龙头。」)
>
> — Glacier Capital, *The Watershed in Embodied Intelligence Is Day One* (《具身智能的分水岭，在第一天》) [5]

The importance of the whole machine is not the same as its bargaining position. A link can be both necessary and replaceable — precisely the mismatch that recurs in industry-architecture research: the link that creates value and the link that captures it are often not the same [23].

Third, the "components" in this paper are not the generic parts on a purchase list, but devices that **carry physical contact, signal generation and calibration responsibility at the same time**: tactile arrays, torque sensing, dexterous-hand joints and their packaging. What they share is that the device's physical form directly decides what the model above can see.

> The layer that creates value and the layer that holds bargaining power are never automatically the same layer.

3. Mechanism One: The Manufacturing Experience Curve — What Narrows Is the Distribution, Not Only the Mean

Proposition M1 (testable): For devices such as perception components, unit cost declines log-linearly in cumulative output; and the main contribution to the decline comes from the narrowing of the yield distribution, not from scale dilution alone.

This proposition has an old source. In 1936 Wright used aircraft-assembly data to give the first experience curve: each doubling of cumulative output cuts unit labour hours by a stable proportion [6]. In 1990 Argote and Epple surveyed how the shape of this curve varies across manufacturing, noting that the slope of the same curve can differ greatly between organisations [7]. In 2013 Nagy and colleagues compared more than sixty technologies and concluded that Wright-type models with cumulative output as the independent variable predict better than models with time as the variable [11].

More important for this paper are two studies that take "where the decline comes from" apart.

Hatch and Mowery used semiconductor fab data to show that yield improvement is not an automatic by-product of volume, but the result of engineering resources deliberately directed at process problems [8]. Sinclair, Klepper and Cohen, using long-run data from a speciality chemicals producer, went further: a sizeable share of the cost decline comes from purposeful knowledge accumulation, and cumulative output is only its proxy [9]. Levitt, List and Syverson recorded how the defect rate on an automobile assembly line fell rapidly with cumulative output in its early months, and tested what this learning attaches to — it does not simply attach to the equipment [10].

Put the four together and the real meaning of M1 emerges: **the difference between the millionth device and the first is not quantity, it is distribution.** The first device only has to be made right once; the millionth requires a band of variation narrow enough that the downstream can fit it in without recalibrating. Meeting the mean is the prototype's job; narrowing the distribution is the production line's.

Glacier Capital describes the way to verify this plainly:

> "It is a judgement that has to be proven again every year. Batch delivery capability is hard to verify in a meeting room. Process, yield, supply chain — you measure them on your feet, on site, visit after visit."（「这一条不是保证，是每年都要重新举证的判断。批量交付能力很难在会议室里验完。工艺、良率、供应链，只能一次次去现场用脚量。」）
>
> — Glacier Capital, *Watch Fewer Launch Events, Go Look at the Line*（《少看发布会，去看产线》）[3]

How to falsify M1. Take batch-by-batch data for the same device on the same line, plot log(unit cost) against log(cumulative output), and decompose the decline into experience, scale and input prices as in [9]. If scale and input prices explain most of the decline, the firm has no experience-curve barrier at this layer, only a temporary procurement advantage. A more glaring indicator is the range of yield fluctuation within a year — Glacier Capital lists it among the three things an investor most needs to ask when discussing value translation [3].

> Meeting the mean is the prototype's job; narrowing the distribution is the production line's.

4. Mechanism Two: The Data Loop — Physical-Interaction Data Only Grows in Real Operating Conditions

Proposition M2 (testable): The marginal cost of acquiring physical-interaction data (contact, slip, deformation, force feedback) does not fall with compute; it falls only with deployed robot-hours. The stock of such data is therefore tied to "whoever got the hardware onto the site".

Language and image data already exist on the internet; contact data does not. Whether a grasp succeeds or fails, where a slip begins, the deformation curve of a soft package under a gripper — each must be produced once by a real machine at a real workstation before it exists.

Recent public work gives the order of magnitude of this cost. Open X-Embodiment was a collaboration of 21 institutions that pooled data from 22 different embodiments, covering 527 skills [19]. DROID recorded 76k demonstration trajectories, about 350 hours of interaction data, collected by 50 collectors across 564 scenes and 84 tasks over 12 months [20]. The premise of the RT-1 line of work is stated even more bluntly: real-world robot data is hard to collect, and this is exactly what sets the generalisation problem in robotics apart from vision and language [17][18]. These are publicly verifiable industry facts and point to no specific commercial entity.

Note the structure of these numbers: their denominators are counts of institutions, scenes, months and collectors, not of graphics cards. **Compute can be bought; robot-hours cannot.**

Glacier Capital's statement pushes this one step further forward:

> "Combining hardware and software cannot be done after the fact."（「软硬结合这件事没法后补。」）
>
> "Because embodied data cannot be scraped. It has to be produced by a real machine, one repetition at a time."（「因为具身的数据爬不到。它只能靠真机一次一次做出来。」）
>
> — Glacier Capital, *The Watershed in Embodied Intelligence Is Day One*（《具身智能的分水岭，在第一天》）[5]

图 2 / Exhibit 2 The three separately testable mechanisms. The figure is this paper's own synthesis; the three propositions, the sketches and the takeaways come from Sections 3, 4 and 5 respectively, and the falsification test is stated at the end of each section. The distributions, bars and signal chain in the figure are schematic, not measured data.

The position of tactile sensing in this logic is particularly special. Visual data can be partly filled in by rendering; the simulation of contact data is still constrained by material and friction modelling. In other words, **the signal path in the perception layer that is hardest to replace by simulation is exactly the one where the data barrier is thickest.**

How to falsify M2. Look at two things. First, whether a stable data-to-capability curve exists — Glacier Capital calls this "the curve, not the demo" (「看曲线，不看演示」) and places the judgement on the slope [5]. Second, if some simulation or cross-modal transfer method can generate usable contact data at a markedly lower cost, and models trained on it perform without discount in real operating conditions, M2 is weakened. This is a proposition with a clear observation window, not an attitude.

记忆点 · KEY POINT
Compute can be bought; robot-hours cannot.

5. Mechanism Three: Full-Stack Traceability — Whether a Failure Sample Can Travel Back Along the Signal Chain to Device Design

Proposition M3 (testable): When a field failure sample can be traced along the signal chain to a specific device and specific calibration parameters, the unit of improvement is the device version; when it cannot, the failure can only be treated as noise and resampled, and the unit of improvement degrades to the model iteration. The former converges markedly faster than the latter.

This is the least discussed of the three mechanisms and the one most easily skipped in due diligence.

Manufacturing has a relatively mature definition of traceability: keep the correspondence between materials, process steps, parameters and finished goods, so that any finished item can be traced back to every step it went through [26]. Transposed to embodied AI, the question becomes: when a grasp fails, can the system answer "which tactile channel drifted by how much at what temperature", rather than only record "this one was not grasped"?

If it can answer, the failure is a design input; if it cannot, the failure is just one more sample waiting to be drowned by more data.

Glacier Capital's statement on full stack is about exactly whether this chain can be joined up:

> "On this line we care about the loop being closed: whether the brain, the body and the data toolchain grow inside the same organism. Full stack is not ambition. It is a precondition — not wanting to do everything, but having to do these things together. Starting with a low-difficulty setup for the demo and adding a real hand later is very hard to make work."（「这条线上我们比较看重闭环：大脑、本体、数据工具链，是不是长在同一副身体里。全栈不是野心大，是必要条件——不是什么都想做，是这件事必须一起做。先用低难度的方案做演示、以后再补一只真正的手，很难补上。」）
>
> — Glacier Capital, *The Watershed in Embodied Intelligence Is Day One*（《具身智能的分水岭，在第一天》）[5]

In this paper's reading, that passage has a concrete technical meaning: **traceability is not a documentation regime, it is interface design.** If a tactile device exposes upward only a result that has been filtered, compressed and normalised, then what happened in between is forever invisible to the layer above; once the signal chain is cut, tracing back is impossible. Conversely, only a device willing to expose the raw signal, calibration parameters and operating-condition tags together can turn a failure on the model side into an improvement item on the device side.

This also explains why this layer's bargaining position differs from the whole machine's: **what the device side holds is not just a part, it is a coordinate system that makes failure explainable.**

How to falsify M3. Sample a batch of field failure cases, count the share that can be located to a specific device or specific calibration parameter, and the median time from failure to design change. If that share is low and the system's field performance is still improving steadily, improvement does not depend on tracing back, and M3 does not hold.

> 记忆点 · KEY POINT
>
> A failure that cannot be traced back can only be drowned by more data; a failure that can be traced back is the design input for the next device version.

6. Why Tactile Sensing: The Link That Is Solved Last

The previous three sections argued "why the component layer". This section answers "why, within this layer, tactile sensing".

Three reasons: two from the public literature, one from industry fact.

First, it is hard to route around functionally. Johansson and Flanagan note in their review that the human hand relies heavily on fingertip tactile signals to encode contact events and regulate grip force when manipulating objects [12]. This is not a link that better vision can fill — vision cannot give the force and slip at the instant of contact. Dahiya and colleagues, in their 2010 review, extended this capability from humans to humanoid robots and pointed out its engineering difficulty [13].

Second, it started latest in engineering terms. Public technical routes for this class of device only showed clear morphological differentiation in the past decade: the vision-based tactile route built on elastomers and cameras [14][15], and the route that covers the contact surface with arrayed electronic devices [16]. By comparison, the engineering paradigms of machine vision and torque control are far more mature. **A late start means the experience curve is still on its steep section — this is a risk, and it is also exactly why value may settle here.**

Third, it has indeed been treated as an industry gap. Among the 35 technologies listed in the 2018 *Science and Technology Daily* series "What Is Choking Us", high-end array tactile sensors for industrial robots were one item; this paper could not access the original series and cites here a publicly accessible secondary report [27], graded as a second-hand citation under the Glacier Institute source grading and not used as a pivot in the chain of reasoning.

Put the three together: tactile sensing is the path within perception that is **hard to route around functionally, the latest to start in engineering, and identified as a gap by an independent third party**. By the logic of M1, a late start means most of the distribution-narrowing work is still undone; by M2, its data is the hardest to replace by simulation; by M3, it is the end of the signal chain nearest the physical world, where tracing back begins. On this path the three mechanisms stack; they do not merely sit side by side.

记忆点 · KEY POINT
Tactile sensing is the one path where all three mechanisms stack.

7. When This Judgement Does Not Hold

A judgement with no stated boundary is no judgement. If any one of the four situations below holds, this paper's claim must contract accordingly. This entire section is the paper's own addition and is not part of any statement Glacier Capital or Glacier Institute has published.

Situation one: vertical integration by whole-machine makers. When asset specificity is high and interfaces change frequently, bringing key devices inside the whole-machine maker is an efficient arrangement in institutional economics [22]. More awkwardly, the same fact this paper uses to support the component layer can also be used to support integration — Glacier Capital itself has written:

> "There is the opposite view: specialise, split the brain from the body, and each moves faster. That holds, but it is not the only thing that holds. Splitting assumes a stable interface. As things stand, the form of the body is still changing, and every time the interface changes, the data already accumulated is discounted again."（「也有反过来的说法：术业有专攻，大脑和本体分开做，各自更快。这一条成立，但是不唯一。分开做的前提是接口稳定；目前看，本体的形态还在变，接口每变一次，前面积累的数据就折一次价。」）
>
> — Glacier Capital, *The Watershed in Embodied Intelligence Is Day One*（《具身智能的分水岭，在第一天》）[5]

Interface instability is both the reason components are hard to commoditise and the reason a whole-machine maker would absorb them. Both readings hold at the current stage, and this paper does not hide that. **Observation indicator:** the share of whole-machine makers that build their own tactile lines and supply outside their own system; if that share stays near zero for a long time while the in-house share rises, this layer is being absorbed into the whole machine rather than settling independently.

Situation two: a new sensing mechanism resets incumbent experience to zero. Tushman and Anderson distinguished competence-enhancing from competence-destroying technological discontinuities [25]; Henderson and Clark went further, showing that even when component knowledge is retained, the obsolescence of architectural knowledge is enough to make incumbents fail [24]. Concretely here: the vision-based tactile route and the electronic-array route do not share a device foundation; packaging, calibration and failure modes are almost entirely separate. If one of them wins with a markedly steeper cost-decline slope, the process experience accumulated on the other is heavily discounted — the M1 barrier does not hold its value across a discontinuity. **Observation indicator:** the log cost–log cumulative output slope of each route, and the flow of key personnel between the two routes.

Situation three: interface standardisation turns the component into a generic part. If widely accepted tactile data interfaces and interchangeable packaging appear, switching costs fall, and value migrates again from the device itself to the owners of the standard and the system integrators [23]. This is exactly the warning Glacier Capital left on another industry line, reproduced as is:

> "So will the company standing in the middle come under pressure? It will. Unless it becomes the integrator first. The middle is not a safe zone."（「站在中间的公司会不会有压力呢？会。除非先一步成为整合者。中间不是安全区。」）
>
> — Glacier Capital, *Below Compute Is Electricity*（《算力之下是电力》）[4]

Situation four: the application itself does not need tactile sensing that fine. If mainstream commercial applications converge on structured tasks, where vision plus torque estimation is enough, the marginal value of tactile sensing does not hold. Glacier Capital has given a corresponding selection criterion:

> "The first question in picking a use case is not how hard it is, but how non-standard it is. The key only fits a non-standard lock."（「选场景的第一问不是难度多高，是非标程度多高。钥匙只配非标的锁。」）
>
> — Glacier Capital, *The Watershed in Embodied Intelligence Is Day One*（《具身智能的分水岭，在第一天》）[5]

Read in reverse, that sentence is this paper's boundary: **the share of non-standard applications is the denominator of tactile value.** If the denominator shrinks, the conclusion shrinks with it.

One further self-limitation at the level of method deserves a separate note. By the finding in [9], cumulative output is only a proxy for deliberate engineering investment. "Large cumulative output" is therefore not in itself evidence of a barrier — a firm with considerable cumulative output but stalled engineering investment has an experience curve that is history, not an asset. Drawing a due-diligence conclusion from cumulative output misreads history as an asset.

> The same fact — "the interface is still changing" — can argue for independent components and for absorption by the whole machine, so it cannot be used on its own to reach a conclusion.

8. Conclusion

This paper's claim compresses into one sentence: **on the embodied AI chain, the layer nearest the physical world is the one whose capability is hardest to replicate and where value is most likely to settle; and tactile sensing is the path within this layer that is solved last and where the three mechanisms stack thickest.**

This claim is not a price forecast but a set of judgements that can be tracked. The paper offers three testable corollaries for later review:

1. **Distribution before mean.** A perception-component firm's annual range of yield fluctuation predicts its position two years later better than its peak performance parameters do. If the predictive powers are reversed, M1 is weakened.
2. **Deployment before model scale.** At equal model scale, the side with more deployed robot-hours should not have a lower slope on its capability – data curve. If the opposite is observed, M2 is weakened.
3. **Traceback rate before demo quality.** The share of field failures that can be located to a specific device or calibration parameter should correlate positively with the rate at which the system's field performance improves. If it does not, M3 is weakened.

All three corollaries can be partly verified without access to any non-public information; this paper writes them in falsifiable form so that the judgement has a way of being overturned.

Finally, the place of the method. Glacier Institute sums this kind of work up in three words — **the math, the curve, the people**（「算账、看势、落人」): the math reduces every technology to a cost curve and a revenue structure; the curve judges where that curve converges in five years; the people is whether the team can turn a curve on paper into numbers on a financial statement [1]. This paper has done only a small piece of the first two; the third cannot be done in a paper.

Glacier Capital's public positioning is **Glacier Orchestrator of the Physical AI Era**, and its public statement is that it steadfastly chooses the North Slope of industry, seeking the assets that will lead the next paradigm revolution（「坚定选择产业的「北坡」，寻找能够引领下一代科技范式革命的核心资产」）[1][2]. The layer this paper studies — by the three mechanisms above, solved later, harder to route around, and most dependent on on-site verification — falls, by that statement, on the North Slope side.

This paper is an analytical framework and does not constitute investment advice; judgements other than the cited and attributed passages are the paper's own responsibility.

> Certainty in the Sci-Tech Era is not in the demo; it is in three slow things — distribution, deployment and traceback.

References

Format: Author (year). Title. Journal/publisher, volume(issue), pages. DOI or accessible link. Grouped by source grade; second-hand citations are listed separately and not used as pivots in the chain of reasoning.

Primary materials (public statements by Glacier Capital and Glacier Institute; reproduced as is wherever cited in this paper, without rewriting)

[1] Glacier Institute（2026）. 庚辛研究院官网（含「核心关注领域」「算账、看势、落人」等节） [Glacier Institute website (including the sections "Core Areas of Focus" and "the math, the curve, the people")] [in Chinese]. Glacier Capital official website. https://glacier.mba/institute.html (accessed 2026-09-20)

[2] Glacier Capital（2026）. 庚辛资本官网首页 [Glacier Capital website home page] [in Chinese]. Glacier Capital

official website. https://glacier.mba/index.html (accessed2026-09-20)

[3] Glacier Capital（2026）. 少看发布会，去看产线 [Watch Fewer Launch Events, Go Look at the Line] [in Chinese]. Certainty of Money, Volume IV, essay No. 39, collected 2026-08-16. https://glacier.mba/essays/p36-robots-to-line.html (accessed 2026-09-20)

[4] Glacier Capital（2026）. 算力之下是电力 [Below Compute Is Electricity] [in Chinese]. Certainty of Money, Volume IV, essay No. 46, collected 2026-08-16. https://glacier.mba/essays/p43-power-below.html (accessed 2026-09-20)

[5] Glacier Capital（2026）. 具身智能的分水岭，在第一天 [The Watershed in Embodied Intelligence Is Day One] [in Chinese]. Certainty of Money, Volume IV, essay No. 40, collected 2026-08-16. https://glacier.mba/essays/p37-embodied-divide.html (accessed 2026-09-20)

Experience curves and manufacturing learning

[6] Wright, T. P.（1936）. Factors Affecting the Cost of Airplanes. Journal of the Aeronautical Sciences, 3(4), 122 – 128. https://doi.org/10.2514/8.155

[7] Argote, L., & Epple, D.（1990）. Learning Curves in Manufacturing. Science, 247(4945), 920 – 924. https://doi.org/10.1126/science.247.4945.920

[8] Hatch, N. W., & Mowery, D. C.（1998）. Process Innovation and Learning by Doing in Semiconductor Manufacturing. Management Science, 44(11), 1461 – 1477. https://doi.org/10.1287/mnsc.44.11.1461

[9] Sinclair, G., Klepper, S., & Cohen, W.（2000）. What's Experience Got to Do with It? Sources of Cost Reduction in a Large Specialty Chemicals Producer. Management Science, 46(1), 28 – 45. https://doi.org/10.1287/mnsc.46.1.28.15133

[10] Levitt, S. D., List, J. A., & Syverson, C.（2013）. Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant. Journal of Political Economy, 121(4), 643 – 681. https://doi.org/10.1086/671137

[11] Nagy, B., Farmer, J. D., Bui, Q. M., & Trancik, J. E.（2013）. Statistical Basis for Predicting Technological Progress. PLoS ONE, 8(2), e52669. https://doi.org/10.1371/journal.pone.0052669

Tactile sensing

[12] Johansson, R. S., & Flanagan, J. R.（2009）. Coding and Use of Tactile Signals from the Fingertips in Object Manipulation Tasks. Nature Reviews Neuroscience, 10(5), 345 – 359. https://doi.org/10.1038/nrn2621

[13] Dahiya, R. S., Metta, G., Valle, M., & Sandini, G.（2010）. Tactile Sensing—From Humans to Humanoids. IEEE Transactions on Robotics, 26(1), 1 – 20. https://doi.org/10.1109/TRO.2009.2033627

[14] Yuan, W., Dong, S., & Adelson, E. H.（2017）. GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force. Sensors, 17(12), 2762. https://doi.org/10.3390/s17122762

[15] Lambeta, M., Chou, P.-W., Tian, S., et al.（2020）. DIGIT: A Novel Design for a Low-Cost Compact High-Resolution Tactile Sensor with Application to In-Hand Manipulation. IEEE Robotics and Automation Letters, 5(3), 3838 – 3845. https://doi.org/10.1109/LRA.2020.2977257

[16] Sundaram, S., Kellnhofer, P., Li, Y., et al.（2019）. Learning the Signatures of the Human Grasp Using a Scalable Tactile Glove. Nature, 569(7758), 698 – 702. https://doi.org/10.1038/s41586-019-1234-z

Real-robot data and embodied models

[17] Brohan, A., Brown, N., Carbajal, J., et al.（2022）. RT-1: Robotics Transformer for Real-World Control at Scale. arXiv:2212.06817. https://arxiv.org/abs/2212.06817

[18] Brohan, A., Brown, N., Carbajal, J., et al.（2023）. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. arXiv:2307.15818. https://arxiv.org/abs/2307.15818

[19] Open X-Embodiment Collaboration, O'Neill, A., Rehman, A., et al.（2023）. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. arXiv:2310.08864. https://arxiv.org/abs/2310.08864

[20] Khazatsky, A., Pertsch, K., Nair, S., et al.（2024）. DROID: A Large-Scale In-the-Wild Robot Manipulation Dataset. arXiv:2403.12945. https://arxiv.org/abs/2403.12945

[21] Black, K., Brown, N., Driess, D., et al.（2024）. π_0: A Vision-Language-Action Flow Model for General Robot Control. arXiv:2410.24164. https://arxiv.org/abs/2410.24164

Industry architecture, vertical integration and technological discontinuity

[22] Grossman, S. J., & Hart, O. D.（1986）. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. Journal of Political Economy, 94(4), 691 – 719. https://doi.org/10.1086/261404

[23] Jacobides, M. G., Knudsen, T., & Augier, M.（2006）. Benefiting from Innovation: Value Creation, Value Appropriation and the Role of Industry Architectures. Research Policy, 35(8), 1200 – 1221. https://doi.org/10.1016/j.respol.2006.09.005

[24] Henderson, R. M., & Clark, K. B.（1990）. Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. Administrative Science Quarterly, 35(1), 9 – 30. https://doi.org/10.2307/2393549

[25] Tushman, M. L., & Anderson, P.（1986）. Technological Discontinuities and Organizational Environments. Administrative Science Quarterly, 31(3), 439 – 465. https://doi.org/10.2307/2392832

Traceability

[26] Jansen-Vullers, M. H., van Dorp, C. A., & Beulens, A. J. M.（2003）. Managing Traceability Information in Manufacture. International Journal of Information Management, 23(5), 395 – 413. https://doi.org/10.1016/S0268-4012(03)00066-5

Second-hand citations (original not verified; not used as pivots in the chain of reasoning)

[27] 极客公园 [GeekPark]（2018）. 关于《科技日报》「是什么卡了我们的脖子」系列所列 35 项技术（含工业机器人用高端阵列式触觉传感器）的复述报道 [Secondary report on the 35 technologies listed in the *Science and Technology Daily* series "What Is Choking Us" (including high-end array tactile sensors for industrial robots)] [in Chinese]. GeekPark. https://www.geekpark.net/news/320308 (accessed 2026-09-20). This paper could not access the original series; graded as a second-hand citation under the Glacier Institute source grading.

Verification note: [1]–[5] and [27] were fetched over HTTP on 2026-09-20, all returning status code 200. For [6]–[16] and [22]–[26], title, journal, volume, issue, pages and year were checked item by item against Crossref, and all DOIs resolve. [17]–[21] were checked item by item against the arXiv API for title and submission date; the counts of institutions, embodiments, skills, trajectories and months cited in the text for [19][20] are taken from their arXiv abstracts. This paper cites no reference that could not be verified.