INSTITUTE FOR PHYSICAL AI @ BMI
The Glass Lab

Technical Report TR-2026-19
Survey / Position · Preprint v1
12 July 2026

THE DIFFUSION LAYER

The Diffusion Layer: Accelerating the Spread of Silicon-Design Competency for Physical AI — A Survey and Research Position

Physical AI is bottlenecked less by fabs than by the people who can design for them. This report treats that competency as an innovation to diffuse, and the training as the technology that sets its speed.

Charles R. Glass, EdD

The Glass Lab, Institute for Physical AI @ BMI

Correspondence: contact@physicalai-bmi.org · physicalai-bmi.org
Companion: the "Silicon for Physical AI" course (physicalai-bmi.org/education/silicon) — a real Verilog→synthesis→tapeout path in the browser

Abstract. Physical AI runs on custom silicon — for perception, for on-device inference, for the energy budget of a machine that carries its own power — and the binding constraint on producing that silicon is not fabrication capacity but the number of people who can design it. Design competency is the slowest-diffusing input in the stack, because its hardest layers are tacit: analog intuition, physical design, and verification are learned by apprenticeship, not from a manual. This report surveys the diffusion of technical competency through the lens of Rogers's innovation adoption and the economics of human capital, and identifies two levers that have recently moved the diffusion slope for silicon specifically: the collapse in the cost and trialability of a real tapeout — open-source EDA toolchains and shared multi-project-wafer shuttles now put a fabricated chip within a roughly \$100–\$300 reach — and AI copilots that carry tacit expertise down the experience curve, which the first field studies find lift novices most. The report's position is that the training itself is the technology whose cost and trialability set how fast design competency spreads, that lowering the cost of one genuine end-to-end design repetition is therefore the highest-leverage intervention, and that the responsibility for producing this general, transferable competency sits with education. We are explicit about maturity: this is a survey with a stated position and reports no new empirical measurements.

1. Introduction

The public conversation about semiconductors is about fabs — their cost, their location, their throughput. For physical AI the more immediate constraint is upstream of the fab: the supply of people who can design the chips a fab would build. Custom silicon is what makes an embodied machine affordable to run — a sensor front-end, an on-device inference block, a power-management design matched to a body that carries its own energy — and each is the product of scarce, deep design skill. Fabrication capacity can be financed and built; design competency is grown in people, slowly, and is the input most likely to bind.

It binds because its hardest layers resist the fast channels of diffusion. A synthesizable digital block can be taught from documentation and open tools; analog design, physical implementation, and verification are learned the way a craft is learned — by doing the work beside someone who already can, absorbing judgment that was never written down.[1] That tacit layer is exactly the competency the field is short of, and exactly the competency that spreads slowest. This report treats the spread itself as the object of study — the *Diffusion Layer* — and asks what sets its speed and how to raise it.

2. Scope and method

This is a survey with a stated research position, not an experimental paper. It reviews the diffusion-of-innovations and human-capital literatures, the open-source hardware-design ecosystem, and the early evidence on AI coding assistants, favoring a primary source per claim and stating maturity. It advances a position in Sections 6–7 and reports no original measurements. Where it cites a cost figure — for example the reach of a shared shuttle tapeout — it reports the order of magnitude published by the relevant program rather than a precise quote, and treats that figure as illustrative of a trend, not as a benchmark. The companion is the Institute's "Silicon for Physical AI" course, a working Verilog-to-synthesis-to-tapeout path in the browser, referenced as an existence proof of the low-cost repetition the position argues for.

3. What diffuses slowly: the tacit layer of design

Rogers's account of how any innovation spreads gives the variables that set its rate: the innovation's relative advantage, its *complexity*, its *trialability*, its observability, and the communication channels available to carry it. [2] Silicon-design competency scores badly on the two variables that matter most here. It is complex, and — historically — barely trialable: for most of the field's history a learner could not attempt a real design end to end without institutional access to six- and seven-figure tooling and a fabrication run, so the practice that builds competence was rationed to a few. Polanyi's distinction between codified and *tacit* knowledge explains why the complexity does not yield to documentation: the decisive competencies — reading a layout for the parasitic that will bite, knowing which corner to worry, sensing when a verification pass is lying — are tacit, transmitted by apprenticeship, and therefore bottlenecked by the availability of experts' time.[1] A field whose critical skill diffuses only through scarce apprenticeship grows its workforce at the rate its experts can mentor, which is slow.

4. Two levers that moved the slope

Two recent developments act directly on Rogers's rate variables for silicon design. The first is **trialability**. Open-source electronic-design-automation toolchains — synthesis, place-and-route, and verification that run without a commercial license[3,4] — combined with shared multi-project-wafer shuttles that split one wafer's cost across many small designs, have brought a real, fabricated chip within roughly a \$100–\$300 reach for an individual learner.[5] A practice that was institutionally rationed becomes something a student can attempt, and attempt again. The second lever acts on **complexity**, and specifically on the tacit layer: AI copilots that suggest, explain, and check design and code carry expert judgment down the experience curve, and the first field studies of their effect find the lift concentrated on the least-experienced — a 34% productivity gain for novices in one large study of AI-assisted work, and comparable novice-weighted gains in controlled studies of code assistants.[6,7] A copilot that makes tacit competence partly available on demand loosens the apprenticeship bottleneck that set the old, slow rate.

> The two levers are complementary and multiplicative. Cheap trialability supplies the repetitions; the copilot raises the competence gained per repetition by supplying, on demand, some of the tacit judgment that a mentor used to be the only source of. Together they change the diffusion of design competency from mentor-rate-limited to practice-rate-limited.

5. Human capital: which silicon skill transfers

Becker's separation of *general* from *specific* human capital predicts where the leverage and the responsibility lie. [8] Much of design competency is general: the discipline of specifying, synthesizing, verifying, and reasoning about a physical implementation transfers across process nodes, across tools, and across employers, and it outlives the specific PDK or flow of the moment. General capital is the kind a market under-supplies, because a firm that trains it cannot fully capture the return, and it is therefore the kind an educational institution is positioned to produce. The specific capital — the idioms of one company's flow — is best left to the employer who benefits from it. The division sharpens the training target: teach the general, transferable design competency that the two levers have just made cheap to practice, and let specificity be acquired on the job.

6. Research position

The survey supports a position with three commitments. **Treat the training as the technology.** Because diffusion rate is set by trialability and complexity, the highest-leverage intervention is not another exhortation to enter the field but a reduction in the cost and difficulty of one genuine, end-to-end design repetition — specify, synthesize, verify, and, where possible, tape out — delivered where the learner already is. The Institute's "Silicon for Physical AI" course is built to this specification: a real synthesis-and-verification path in the browser, on open tools, that ends at a submittable tile. **Pair every repetition with a copilot, deliberately.** The novice-weighted gains are the mechanism by which tacit competence diffuses; a curriculum should use an assistant to surface expert judgment at the moment of practice, while measuring that the learner is acquiring the judgment rather than outsourcing it. **Measure demonstrated competency.** As with any craft, the unit of account is a working artifact produced under realistic constraints — a design that synthesizes, verifies, and could be fabricated — not a certificate of attendance. The Diffusion Layer is, in this framing, the Glass Lab's workforce thesis specialized to silicon: the bottleneck on physical AI's hardware is a human-capital bottleneck, and it is closed by making the scarce competency cheap to practice and fast to spread.

7. Maturity and limitations

The report restates established results — Rogers on diffusion, Polanyi on tacit knowledge, Becker on human capital, and the early empirical literature on AI assistance — each cited to a primary source, and advances the Diffusion-Layer framing and the training position as a synthesis rather than a validated finding. The cost figures for open-toolchain tapeout are reported as published orders of magnitude and as a trend, not as benchmarks, and will move. No new empirical measurements are reported here; whether a low-cost, copilot-paired, tape-out-terminated curriculum actually raises the rate at which design competency diffuses through a population — and whether the competence acquired is the tacit judgment or a shallower substitute — are the decisive open questions, and are stated as such rather than answered.

References

1. Polanyi, M. (1966). *The Tacit Dimension*. University of Chicago Press.
2. Rogers, E. M. (2003). *Diffusion of Innovations* (5th ed.). Free Press.
3. OpenROAD / OpenLane: open-source RTL-to-GDSII digital implementation flow. The OpenROAD Project. (Documentation, 2020–present.)
4. Wolf, C. (2016–present). *Yosys Open SYnthesis Suite*. (Open-source Verilog synthesis.)
5. TinyTapeout: a shared multi-project-wafer program placing small designs on real silicon at low per-learner cost. (Program documentation, 2022–present.)
6. Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at Work. NBER Working Paper 31161. (Novice productivity lift ≈ 34%.)

7. Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. arXiv:2302.06590.
8. Becker, G. S. (1964). *Human Capital: A Theoretical and Empirical Analysis.* University of Chicago Press. (General vs. specific human capital.)