Technical Report TR-2026-18
Survey / Position · Preprint v1
12 July 2026

THE HUMAN LAYER

The Human Layer: Competency at the Frontier of Autonomous Systems — A Survey and Research Position

As machines take the loop, the human role does not disappear; it moves to the boundary. This report maps the competency that lives there, why automation erodes it, and how to build it deliberately.

Charles R. Glass, EdD

The Glass Lab, Institute for Physical AI @ BMI

Correspondence: contact@physicalai-bmi.org · physicalai-bmi.org
Interactive companion: physicalai-bmi.org/research/glass-lab — an illustrative model of the competency gap

Abstract. Autonomy does not remove the human from a physical-AI system; it relocates the human to the system's boundary — to supervision, exception handling, calibration, and accountability. The competencies that boundary demands are different from the ones the machine displaced, and they are precisely the competencies that routine operation no longer exercises. This report surveys the long literature on that relocation, from Fitts's division of labor and Bainbridge's *ironies of automation* to modern findings on human–AI teaming, and reads it through the economics of human capital and the diffusion of innovations. Its organizing claim is a frontier: as a task's autonomy level rises, the human's residual role climbs to higher-order competencies — situation assessment, model calibration, judgment under uncertainty, and ethical governance — while the hours available to practice them fall, opening a gap between the competency a system needs at its edge and the competency its operators retain. The report's research position follows: the training must target the frontier, not the displaced task; it must be measured as demonstrated competency rather than course completion; and its most effective form is deliberate, low-cost, high-repetition practice of the boundary skills, delivered where the work is. We are explicit about maturity. This is a survey with a stated position and reports no new empirical measurements; the interactive companion is an illustrative model, labeled as such.

1. Introduction

Every account of automation that has aged well makes the same correction to the intuition that a more automatic machine needs a less capable human. Fitts's 1951 allocation of functions asked which tasks suit people and which suit machines, and already implied that as the machine takes the routine middle, the human is left the ends — the setup, the anomaly, the judgment.[1] Bainbridge sharpened this into the *ironies of automation*: the more reliable the automatic system, the more its human supervisor's remaining job is the one thing the automation cannot do, and the less that supervisor ever practices it, so competence decays exactly where it is most needed.[2] Physical AI restates the irony at higher stakes. A system that flies, drives, welds, or handles a payload delegates the routine loop to a policy, and the person who remains is asked for situation assessment, calibration of trust, and accountable judgment at the moment the policy reaches the edge of its competence — the moment the person has had the least occasion to rehearse.

This report calls the set of those residual, boundary competencies the *Human Layer*, and treats it as a first-class object: something a program can name, teach, and measure, rather than a residue that survives whatever the machine did not take. The motivation is practical. In autonomous aviation and the uncrewed-aircraft workforce the Glass Lab studies, the competency that decides safety is not stick-and-rudder skill the automation now supplies but the higher-order skill of knowing when to trust it, when to intervene, and how to hold the system accountable — and that competency is scarce, unevenly distributed, and hard to build by operating a system that rarely fails.[3]

2. Scope and method

This is a survey with a stated research position, not an experimental paper. It reviews published work in human factors, human–automation interaction, the economics of human capital, and the diffusion of innovations, favoring a primary source per claim and stating the maturity of each. It advances a position in Sections 6–7 and reports no original measurements. The interactive companion referenced throughout is an illustrative model of the competency gap — a teaching device, labeled as such — not a data source. The author's expertise is in education and workforce development; the report's contribution is to organize known results around the frontier construct and to draw a training program from them, not to make new empirical claims about human performance.

3. The relocation, and the irony that follows it

Order tasks by *autonomy level* — the fraction of the operating loop the machine closes on its own — and track what is asked of the human at each level. At low autonomy the human is the controller; at intermediate autonomy the human is a co-pilot who shares control; at high autonomy the human is a supervisor who monitors and intervenes; at the highest levels the human is an exception handler and an accountable authority who may act only rarely. The taxonomy is familiar from driving automation and from the broader levels-of-automation literature,[4] and its shape is the point: the human's task does not shrink smoothly toward nothing. It jumps, at each level, to a qualitatively higher-order competency, while its duty cycle — the hours actually spent doing it — collapses. Bainbridge's irony is the composition of those two curves: demand for the boundary skill rises with autonomy while opportunity to practice it falls.[2]

Three failure modes recur across the literature and name what the Human Layer must defend against. *Skill decay*: a competency not exercised degrades, so the supervisor of a reliable system is least ready at the rare moment of need.[2] *Automation bias and complacency*: trust calibrated to a system's routine reliability over-generalizes to its edge cases, and monitors miss the failures they are there to catch.[5] *Loss of situation awareness*: out of the control loop, the human's model of what the system is doing and why goes stale, so intervention, when it comes, is slow and wrong.[3] None of these is a deficit of the displaced skill. Each is a deficit of a boundary competency that the design of the automation actively starves.

4. The competency frontier

Combine the two curves into one picture. Let the horizontal axis be autonomy level and the vertical axis be the order of competency the human's residual role requires — from manual control, through monitoring and trust calibration, to model calibration, judgment under uncertainty, and governance. The role traces a rising staircase: each increment of autonomy pushes the human up to a higher rung. Call that staircase the *competency frontier*. Against it, plot the competency an operator population actually retains, which is set by how often the boundary skills are practiced — and which the same increment of autonomy pushes *down*, because routine operation no longer rehearses them. The vertical distance between the two is the *competency gap*: the shortfall between what a system needs at its edge and what its operators can supply. Automation, left to itself, widens the gap from both sides at once. This is the report's central organizing claim, and the companion model animates exactly this pair of curves.

> The frontier reframes the workforce question for physical AI. The scarce input is not operators for the old task, which the machine now performs; it is people fluent in the boundary competencies the new frontier demands — and those are not produced as a by-product of running the system. They must be built on purpose.

5. Human capital at the frontier

The economics of human capital tells us which competencies are worth building and by whom. Becker's distinction between *general* capital, which is productive across employers and settings, and *specific* capital, which is productive only in one, predicts who invests in each: firms under-invest in general skills their workers can carry away, so general capital tends to be under-produced relative to its social value.[6] The boundary competencies of the Human Layer — situation assessment, calibrated trust, judgment under uncertainty, the disposition to hold a system accountable — are overwhelmingly *general*. They transfer across aircraft, vehicles, and manipulators; across employers; and across the specific automation of the day, which will be replaced before the competency is. That is precisely the capital a market will under-supply, and precisely the capital an educational institution exists to produce. The frontier construct and Becker's asymmetry meet at a single conclusion: the highest-leverage training targets the general, boundary competencies that autonomy makes scarce, and the responsibility for producing them sits with education rather than with the operator of any one system.

6. Diffusion: training as the technology to accelerate

If the Human Layer must be built deliberately, the question becomes how fast a new competency can spread through a workforce — which is the question Rogers's *diffusion of innovations* answers for any new practice: adoption follows an S-curve whose slope is set by the innovation's cost, its trialability, and the channels available to carry it.[7] Two modern findings bear directly on the slope for boundary competencies. First, generative-AI copilots, in the first field studies of their effect on work, lift the least-experienced workers the most — one study of customer-support agents found a 34% productivity gain concentrated among novices, consistent with a tool that diffuses the tacit competency of experts down the experience curve.[8] Second, deliberate practice — repeated, feedback-rich rehearsal of the specific sub-skill, rather than mere time on task — is what actually moves competence, and its yield rises when the practice is cheap enough to repeat often.[9] Both point the same way: the training itself is the technology whose cost and trialability set the diffusion slope, so lowering the cost of one more high-quality repetition of a boundary skill is the most direct lever on how fast the Human Layer can be built.

7. Research position

The survey supports a position with three commitments. **Target the frontier, not the displaced task.** A curriculum for an autonomous era should teach the boundary competencies autonomy makes scarce — supervision, trust calibration, model calibration, judgment, and governance — not re-teach the routine control the machine now supplies. **Measure competency, not completion.** Because skill decay is the failure mode, the unit of account must be demonstrated, retained competency under realistic conditions, resurfaced over time, rather than a certificate of attendance; a program should assume decay and design for maintenance. **Make one more repetition cheap.** Because diffusion slope is set by the cost of practice, the boundary skills should be rehearsed in low-cost, high-repetition, feedback-rich environments delivered where the work is — the same design principle behind the Institute's in-browser, on-device laboratories and its mastery-withspaced-review progress model. The Perception–Privacy–Policy track of the Glass Lab adds a fourth boundary competency the frontier will demand: governing the operator-sensing that physical-AI systems make possible, deciding what may be sensed, by whom, and with what consent — a competency, not only a policy.

8. Maturity and limitations

The report's claims are of two kinds, held to different standards. The survey (Sections 3–6) restates established results in human factors, human-capital economics, and diffusion, each cited to a primary source. The frontier construct (Section 4) and the training position (Section 7) are the report's contribution and are advanced as a well-motivated synthesis, not as empirically validated findings; the competency gap it describes is drawn as a qualitative relationship, and the companion model is an illustration of that relationship, not a measurement of it. No new human-performance data are reported. Quantifying the gap for a specific autonomous-aviation task, and measuring whether frontier-targeted training closes it faster than task-targeted training, are the obvious next steps and are stated as open questions rather than answered here.

References

1. Fitts, P. M. (ed.) (1951). *Human Engineering for an Effective Air-Navigation and Traffic-Control System.* National Research Council. (The "Fitts list" allocation of functions.)
2. Bainbridge, L. (1983). Ironies of Automation. *Automatica* 19(6), 775–779.
3. Endsley, M. R. (1995). Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors* 37(1), 32–64.
4. Sheridan, T. B., & Verplank, W. L. (1978). *Human and Computer Control of Undersea Teleoperators.* MIT Man–Machine Systems Lab. (Levels of automation.)
5. Parasuraman, R., & Manzey, D. H. (2010). Complacency and Bias in Human Use of Automation: An Attentional Integration. *Human Factors* 52(3), 381–410.
6. Becker, G. S. (1964). *Human Capital: A Theoretical and Empirical Analysis.* University of Chicago Press. (General vs. specific human capital.)
7. Rogers, E. M. (2003). *Diffusion of Innovations* (5th ed.). Free Press.
8. Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at Work. NBER Working Paper 31161. (Novice productivity lift ≈ 34%.)
9. Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The Role of Deliberate Practice in the Acquisition of Expert Performance. *Psychological Review* 100(3), 363–406.