We Have Never Been Intelligent: The Cognitive Constitution and the Myth of the Individual Mind

PIATRA . INSTITUTE

June 2026

Abstract

The title overstates on purpose, and the paper spends its length earning the qualification it needs: we have never been intelligent alone. Human organisms hold real and unequal cognitive capacities, but the performances usually offered as proof of intelligence, in mathematics, physics, law, engineering, administration, and now artificial intelligence, are rarely produced by an isolated brain. They are realized by assemblages of trained bodies, notations, instruments, collaborators, procedures, and inherited institutions, and then attributed, after the fact, to a single named person. This paper reads that attribution through Bruno Latour. Latour argued that the modern settlement works by purification, declaring a clean separation between nature and society while its laboratories proliferate the hybrids the separation forbids anyone to name. An analogous settlement governs intelligence. A Cognitive Constitution depends operationally on cognitive hybrids while representationally and normatively purifying their outputs into properties of bounded individuals, and the operation that performs the purification is what the paper calls attributional compression: collaborative production booked to a genius, an educational infrastructure booked to an examination score, a software ecosystem booked to a programmer's talent, a corpus and a data-center booked to a model. The argument keeps four things apart that the folk concept fuses, capacity and performance and realization and historical possibility, so that psychometrics, distributed cognition, and cultural evolution stop reading as rivals and start reading as descriptions of different levels. It rests on the defensible claims, that the realizing system for many advanced tasks exceeds the individual and that purification happens through observable attribution practices, and it declines the contested metaphysical claim that the notebook is literally part of the mind. A recent and contested line of work pushes the relocation to its limit, proposing that the capacity to learn is a generic affordance of mathematical structure that natural selection discovers rather than builds, so that random unselected networks already sit close to optimal for learning; if that holds, then crediting human biology for intelligence is the same compression run at the largest scale, credit assigned at every level and an author at none. The closing move is the sharp one: purification is not an error to be debugged but the accounting system of merit, which is why the individual mind survives every demonstration of its insufficiency. Intelligence names a settlement about who gets the credit, and we have been litigating it, under that name, for a long time.

1. The Equation on the Board

A physicist stands at a blackboard and writes a line of symbols that resolves a problem, and the room agrees it has watched intelligence happen. Look at what is in the room. There is the notation, three centuries of it, without which the line could not be written or read. There is the board itself, a surface for thinking that the hand uses to hold steps the head cannot. There is the training, two decades of it, that installed the moves the hand now makes without deliberation. There is the literature the result answers and the community that will check it, the instruments somewhere downstairs that will say whether it is true, and the institution that pays for all of it and will, if the result holds, attach the physicist's name to the achievement and no one else's. Everyone in the room can see these things. Everyone in the room credits the person.

That is the phenomenon this paper is about. Not that cognition is distributed, which is by now a settled finding, but that we keep crediting the individual anyway, in full view of the distribution, and that the crediting is not a mistake we have failed to correct but an operation we depend on. The gap between what produces a performance and what gets the credit for it is stable, structured, and productive. It has a shape worth describing.

The title says we have never been intelligent, and taken flat the title is false. Humans learn, abstract, plan, and reason, and they differ from one another in how well they do these things, reliably enough that the differences predict outcomes (Gould, 1981; Carson, 2007). The defensible claim is narrower and stranger. We have never been intelligent in the purified sense the word usually carries, the sense in which intelligence is an inner property of a bounded individual that the individual then applies to the world. Intelligence in that sense is not a thing the sciences have found inside people. It is a thing institutions assign to people, and the assignment is the subject here. The honest form of the thesis adds one word: we have never been intelligent alone.

The paper borrows its method from Bruno Latour, who made a structurally identical argument about modernity, and then it does the work Latour did not do, which is to say what the equivalent of his purification is for the case of mind, and why, unlike a mere confusion, it does not go away when you point at it.

2. Latour's Settlement

Latour's claim in *We Have Never Been Modern* is not that modernity failed to arrive but that modernity was never what it said it was (Latour, 1993). The modern settlement, what he calls the Constitution, rests on a purification: it sorts the world into two zones, an objective nature out there discovered by science and a subjective society in here negotiated by politics, and it polices the border between them. The work of the moderns is to keep the two zones clean, to insist that a fact is a fact and a value is a value and that nothing belongs to both.

The trick is that the same civilization that purifies most strenuously also manufactures, faster than any before it, the objects that belong to both zones at once. The hole in the ozone layer is a fact of atmospheric chemistry and a product of industrial policy and a problem for international law, all at the same time, and there is no scissor that cuts it into a clean fact and a clean value. Latour calls such objects hybrids, and his point is that modernity is a machine for producing them precisely because it refuses to admit they exist. The Constitution works not despite forbidding the hybrid but because it does. By guaranteeing that nature and society are separate, it frees its practitioners to mix them without limit, since officially the mixing is not happening.

The laboratory is where this is clearest, and the early laboratory studies are where Latour first saw it (Latour and Woolgar, 1986; Latour, 1987). A laboratory is a place that builds a hybrid, an entity that is at once a natural phenomenon and a social construction, and then performs the purification that lets the phenomenon walk out the door as pure nature with no fingerprints on it. What leaves the laboratory is the inscription, the graph or the number or the printed result, stripped of the apparatus and the labor and the negotiation that produced it, and it is the inscription, not the

apparatus, that travels and counts as the fact. The peers who built epistemic cultures of their own describe the same movement from the inside, the careful machinery by which a result is made and the machinery's disappearance from the result (Knorr Cetina, 1999; Galison, 1997). Purification is the name for that disappearance.

Two features of Latour's argument transfer, and one does not, and the two that transfer will do most of the work below. The first is the double structure, a public doctrine of separation running on top of a private practice of mixture, and the second is the directionality, an operation that strips the network off the product so the product can be credited to one of the two pure zones. What does not carry is the symmetry. Latour wanted nature and society treated alike; the case of intelligence is asymmetric, because the thing being purified is not a fact but a person, and a person is something institutions have strong, independent reasons to want.

3. The Cognitive Constitution

Run the same analysis on intelligence and a parallel settlement appears. Call it the Cognitive Constitution. Its public doctrine is that intelligence is an inner property of the individual mind, a quantity a person has and brings to tasks, visible most purely when the person works alone and unaided. Its private practice is that almost nothing anyone points to as intelligence is produced that way. The advanced performances are made by assemblages, couplings of trained brains with notations, diagrams, instruments, collaborators, divisions of labor, and inherited institutional procedure, and the Constitution's daily work is to take the output of an assemblage and book it to a person.

The hybrids here are cognitive. The extended-mind tradition named the smallest of them, the coupling of a brain with an external resource it uses as part of its thinking, and argued that under the right conditions the resource is not an aid to cognition but a part of it (Clark and Chalmers, 1998; Clark, 2008). The distributed-cognition tradition named the largest, the cognitive systems that span many people and many artifacts and compute things no member computes alone (Hutchins, 1995; Hollan, Hutchins and Kirsh, 2000). Between them sit the embodied and enactive accounts, for which thinking is something the whole acting body does in a structured environment rather than something a brain does to representations (Varela, Thompson and Rosch, 1991; Noë, 2004; Gallagher, 2005), and the material-engagement account, for which the clay and the wheel and the abacus are not records of thought but places where thought happens (Malafouris, 2013). The term cognitive assemblage is not new, and the debt should be marked: Hayles uses it for the couplings of human and technical cognizers through which information moves and decisions get made (Hayles, 2017).

What the traditions establish, between them, is that the hybrid is real and pervasive. What none of them quite explains is the purification. If cognition is so plainly spread across the assemblage, why does the individual mind remain the unit of every examination, every prize, every wage, every obituary? The answer is not a further fact about cognition. It is an operation performed on cognition by institutions, and the rest of the paper is about that operation.

Latour's Modern Constitution	The Cognitive Constitution
Public doctrine: nature and society are separate	Public doctrine: intelligence is an inner individual property
Private practice: laboratories manufacture nature-society hybrids	Private practice: institutions run cognition on brain-tool-group-institution hybrids
Purification strips the network from the fact	Attributional compression strips the assemblage from the performance
The inscription travels; the apparatus disappears	The score, the byline, the credential travels; the assemblage disappears
Works because forbidding the hybrid licenses making it	Works because crediting the individual licenses relying on the system
Product credited to pure nature	Product credited to the bounded person

4. Four Things "Intelligence" Names

The folk concept is a fusion, and most of the long quarrels about intelligence are quarrels between people holding different members of the fusion. Pulling them apart costs nothing and buys the rest of the argument. Four things travel under the one word.

The first is capacity, what a person can reliably learn and contribute across situations. Capacity is the proper object of psychometrics, it differs between people, and the differences are partly heritable and predict outcomes; denying this is not necessary to the thesis and would be false (Gould, 1981; Carson, 2007). The second is performance, what actually gets accomplished on a particular task on a particular day. The third is realization, the system whose operations produce that performance, which may be a brain alone but for advanced tasks is usually a brain coupled to much else. The fourth is historical possibility, the inherited infrastructure of notation, technique, and institution that makes the task thinkable at all, without which the performance is not poor but unavailable.

These are different questions about different levels, and the sciences that study them are not in conflict. Psychometrics measures capacity and finds it unequal. Distributed cognition describes realization and finds it supra-individual. Cultural evolution accounts for historical possibility and finds it cumulative, the product of a ratchet no individual turns alone (Boyd and Richerson, 1985; Henrich, 2016). A child becomes a physicist and a squirrel does not, and the reason is not that the child's brain contains physics but that the child can be enrolled into the cumulative culture that does, through the shared intentionality and high-fidelity social learning that are the human animal's actual specialism (Tomasello, 1999; Tomasello et al., 2005; Heyes, 2018; Donald, 1991). The capacity is real and located in the child. The performance, when it comes, will be realized by an assemblage, and made possible by three centuries of inherited apparatus. The folk concept assigns all four to the child and calls the sum the child's intelligence.

5. Distribution Was Never the Hard Part

The literature that established the hybrid spent two decades fighting about the wrong thing. After Clark and Chalmers proposed that Otto's notebook, the one the amnesiac patient consults the way another person consults memory, is part of Otto's cognitive system, the dispute turned metaphysical: is the notebook *literally* a part of the mind, or merely a tool the mind uses (Clark and Chalmers, 1998)? Critics answered that causal dependence is not constitution, that a thing can matter enormously to cognition without being cognition, and that the mark of the cognitive is not portable to ink (Adams and Aizawa, 2008; Rupert, 2004). Sterelny offered the useful middle, that most of what the extended-mind theorists point at is better described as scaffolding, an environment built up over a life and a history to support and shape a mind that remains, for bookkeeping purposes, inside the skin (Sterelny, 2010; Sterelny, 2012).

The metaphysical question is real, but it is not the one that matters here, and resting the thesis on it would be a mistake. Grant the critics everything. Grant that the notebook is a tool and not a mind, that constitution is a strong word the cases rarely earn, that the person at the board remains the cognitive agent and the board remains a board. The phenomenon survives untouched, because the phenomenon is not metaphysical. It is that the performance is realized by the whole coupled system and credited to one component of it. Whether the board is part of the physicist's mind is a question for philosophers. That the physics is produced by the physicist-and-board-and-training-and-community and awarded to the physicist is a question for whoever hands out the awards, and they have already answered it.

So the claim the argument leans on is the modest one, the claim even the critics grant. For many advanced tasks the realizing system exceeds the individual, demonstrably, in the sense that the same person stripped of the notation or the instruments or the team does not perform worse but fails (Hutchins, 1995; Kirsh and Maglio, 1994). And the puzzle the modest claim leaves is the one the distribution literature never took up: given that everyone can see this, what work is done by continuing to credit the individual, and by what mechanism is the crediting performed? The mechanism is attributional compression. It takes a performance distributed across an assemblage and returns a property of a person, the way the laboratory takes a hybrid and returns a fact. It is lossy by design, and the loss is the point: a score, a byline, a grade, a citation, a salary is a single number or a single name, and a single name is what the institutions downstream require. Compression is not the failure of the system. It is the product.

6. One Case, Told Six Times

The same structure appears wherever the cognition is impressive enough to be worth crediting, which is the tell. Six instances, told as one.

Quantum mechanics is understood by no unaided mind, including the minds that use it best. It is realized by a coupling of formalism, the standardized notation that does the inference, with laboratory apparatus that produces the phenomena, pedagogy that installs the moves, textbooks that freeze the consensus, and peer review that certifies it. Remove the formalism and the smartest

physicist alive cannot reconstruct the subject from intuition; the intuition is downstream of the formalism, not its source. What looks like understanding a theory is competent participation in an assemblage that holds the theory.

Mathematics is the cleanest case because its scaffold is visible on the page. The history of the field is in large part a history of notation, and the notation is not a record of mathematical thought but its machinery: the Greek diagram is a device that makes deduction possible rather than a picture of a deduction performed elsewhere (Netz, 1999), and algebraic and symbolic writing are technologies for thinking the otherwise unthinkable, so that the mathematician is in a strong sense something the symbols do (Rotman, 2000). A proof is a social and material procedure before it is a private insight. The blackboard is not where the result is recorded. It is where it is computed.

Navigation is Hutchins's case, observed directly. A naval vessel fixes its position through a process spread across a team, a set of instruments, a chart, and a fixed sequence of procedure, and the computation is a property of the bridge, not of any sailor on it; the system remembers, computes, and corrects in ways no member does alone (Hutchins, 1995). The navigator is a node with a job, and a well-run bridge can absorb a weak node and still navigate, which is the operational signature of a cognition that is not where the credit goes.

The laboratory we have already met as Latour's case, and it is the template for the others. Knowledge leaves it as inscription, the apparatus and the labor and the calibration stripped away, so that the result reads as a fact about nature rather than an achievement of a network (Latour and Woolgar, 1986; Knorr Cetina, 1999). The trained perception that makes a researcher able to see the result at all is itself an institutional product, drilled in over years until it feels like vision (Daston and Galison, 2007).

Software is the contemporary case and the most literal. A working program is realized by the programmer coupled to a codebase written mostly by others, a compiler that holds rules no programmer remembers, version control that is the project's external memory, test suites that do the checking, documentation, and a global archive of solved problems consulted continuously. The programmer is the interface to a technical-symbolic system far larger than the programmer, and the talent the industry pays for is in large part fluency at coupling. Unplug the archive and the compiler and the talent has nowhere to go.

Artificial intelligence is the case that has made the others legible, which is its real service to the argument. A model's intelligence is a corpus the size of a civilization's writing, a data-center's worth of compute, an architecture, a benchmark culture, and a great deal of underpaid human feedback, compressed into a noun and a brand. The compression is exactly the one the Cognitive Constitution performs on people, run fast enough and recently enough that we can still see the assemblage before the noun closes over it. Artificial intelligence does not introduce the problem of crediting a system's output to an agent. It exhibits the problem we had already solved, for ourselves, by forgetting we had it.

7. Why the Fiction Survives

If the individual mind were merely a mistake, it would be the most durable mistake in the human sciences, and durability of that order is usually a sign that something is being accomplished rather than missed. The Cognitive Constitution is not a confusion awaiting correction. It is an accounting system, and what it accounts for is merit.

Consider where the individuated number is actually needed. A school must decide who advances, and it cannot advance an assemblage, so it builds an instrument, the closed-book examination, whose entire design is to strip the scaffold away and read off what is left, and it calls the residue the student's intelligence. The history of that instrument is a history of deliberate purification, the construction of a unit of mind that institutions could test, rank, and compare, assembled in the same decades that built the meritocratic state and for the same purpose (Danziger, 1990; Carson, 2007; Gould, 1981). A science must decide whom to reward, and it cannot give the prize to a network, so it operates a credit economy that resolves collective production into named authorship, a convention whose machinery the social history of science lays bare (Shapin, 1994). A firm must set wages, a court must assign responsibility, a state must locate the citizen who votes and is punished and is owed. Each of these needs the output of a system to attach to a body, because bodies are what can be paid, promoted, blamed, and enfranchised. Assemblages cannot hold a credential or stand trial.

This is why the asymmetry that broke the Latour analogy matters. Latour could ask for nature and society to be treated symmetrically because neither nature nor society needed to be a person. Intelligence cannot be symmetrized that way, because the purified pole is the individual, and the individual is not an optional metaphysical posit but the indispensable unit of every institution that distributes anything. You cannot dissolve the individual mind by demonstrating that cognition is distributed, any more than you can dissolve money by demonstrating that value is social. The demonstration is correct and changes nothing, because the institution was never in the business of locating cognition. It was in the business of assigning credit, and for that it needs exactly one name.

So the folk concept of intelligence is doing a job, and the job is not description. It is settlement. It decides, in advance of any inquiry into where the thinking happened, that the thinking will be charged to a person, because the apparatus that consumes the answer can only bill a person. Attributional compression is the operation that produces the bill. The individual mind is its output, not its input, and that is why it survives every finding that should embarrass it.

8. The Gift Nobody Made

The regress seems to have a floor, and the floor is biology. Move the credit off the individual and onto the assemblage, off the assemblage and onto the cultural ratchet, and you still appear to reach solid ground at the species, since the assemblage is a human artifact, the ratchet is a human achievement, and the capacity to build and accumulate is an endowment that selection spent a very long time installing in one animal and not the others. Everything above relocates intelligence without

once leaving the human story. Even on the deflationary reading, you might suppose, intelligence bottoms out in a fact about evolved brains, which is the fact the next section will lean on. There is a recent argument that there is no bottom.

It comes from Michael Levin and his collaborators, and the provocation is that the capacity the argument has been resting on, the ability to learn and to organize behavior toward a problem, is not built by natural selection but found by it. In Levin's framing the competencies of living things ingress from a pre-existing space of mathematical structure rather than being assembled from nothing in each lineage, so that bodies are interfaces onto patterns that are discovered, not invented (Levin, 2025a; Levin, 2022). The sharp form, delivered in a recent lecture over a slide credited to Federico Pigozzi, is that the intelligence ratchet is a gift from mathematics and not from physics or biology: random, unselected networks already sit close to optimal for learning and for causal organization, so selection does not manufacture the capacity, it stumbles onto a free lunch and exploits it (Levin, 2025b).

This is more than a slide, and it is also less than settled, and both halves of that should be said. The published spine is real. Pigozzi, Goldstein, and Levin (2025) report that associative conditioning raises integrative causal emergence in gene-regulatory networks, and they propose a feedback loop in which emergence improves learning and learning in turn raises emergence, the loop they name an intelligence ratchet; Pigozzi and Levin (2026) find that the same measure of causal emergence predicts the reward a reinforcement learner eventually reaches. Beneath that sits the older formal result that a coarse-grained macro description can carry more causal power than the micro substrate it summarizes (Hoel, Albantakis, and Tononi, 2013; Hoel, 2017), extended to networks by Klein and Hoel (2020), and a convergence from several directions on the claim that learning is a generic property of dynamical matter rather than a trophy of selection (Watson and Szathmáry, 2016; Vanchurin, Wolf, Katsnelson, and Koonin, 2022), with some of that work placing emergent causal structure before the first replicators rather than after them (Walker and Davies, 2013). The less-than-settled half matters as much. The strongest claims, the optimality of random networks and the bare assertion that mathematics and not biology turns the ratchet, are talk-stage and in progress, and the metaphysics of a literal realm that forms ingress from is contested and probably unfalsifiable.

The paper does not need the metaphysics, and it should decline it for the reason it already declined the claim that the notebook is part of the mind. Grant the critics of Platonism the whole of their case. Say there is no realm forms come from, that a gift from mathematics is only a vivid way of reporting that random networks are unexpectedly learnable. The deflationary payload survives the demotion, because it was never a claim about a Platonic realm. It was a claim about where the capacity comes from, and the answer it gives is from nowhere in particular, from a generic affordance of structured matter that selection found lying around and that the human lineage happens to couple to better than most.

Take that reading and the argument is the one this paper has been making all along, run at a scale that finally makes it hard to miss. Attributional compression credits the individual for the assemblage's work. The temptation the next section names is to halt the regress at the species and

credit human biology for the capacity to learn. But if learning is a free lunch, crediting the evolved brain for intelligence is that same compression at one more remove, a generic affordance available in principle to a random network, booked to an animal the way the assemblage was booked to a genius and the corpus to a model. The names change and the operation holds. The individual was never the author of the performance, and on this telling the species was never the author of the capacity. There is credit assigned at every level and an author at none of them.

This cuts against the floor, and the cut should not be smoothed away. The next section leans on a located biological fact, that one animal is enrollable into cumulative culture at high fidelity, and the free-lunch thesis appears to dissolve precisely that. It does not, for two reasons, a cheap one and a serious one. The cheap reason is that the levels do not collide: that the capacity to learn is generic does not make every interface onto it equal, and the human animal can be the body that couples best to the ratchet even where the ratchet is no animal's invention, which is Levin's own interface picture. The serious reason is that the paper does not have to decide which floor is real, because its claim was never about the ultimate source of cognition. It was about who gets billed. Whether the source is the cortex, the culture, or the mathematics, the institution still needs one name to charge, and still charges the nearest body, and the further the real source recedes from that body the more plainly the name on the result reads as a decision rather than a finding. The free-lunch thesis, if it holds, does not rescue the individual mind. It strands it further from anything it could be a fact about.

One caution is owed, and it is owed to the paper's own standard. A gift from mathematics can be a precise claim about the learnability of random networks, or it can be a re-enchantment that moves the genius out of the skull and into the cosmos, and the second reading is the intelligence-mysticism the next section warns against, wearing a Platonist's robe. Take the precise claim and leave the robe.

9. What the Argument Leaves Standing

A thesis this large is most useful where it stops, so the closing work is to mark what relocating intelligence does not buy.

It does not deny biology, though the previous section put the standing of biology in question, and the qualification that question forces is worth stating with care. The thesis leans, at one level, on a specific, located, unequal biological fact, that the human animal is the one that can be enrolled into cumulative culture at high fidelity, through capacities for imitation, shared attention, and teaching that other animals lack or hold weakly (Tomasello, 1999; Heyes, 2018; Henrich, 2016). If the free-lunch thesis is right, that enrollability is not the origin of the capacity to learn but the quality of one interface onto a capacity older and more generic than any lineage. Either way the fact is real, it sits in individual organisms, and it differs between them, and either way the argument relocates the production of advanced performance without emptying the person who is enrolled. A brain is not the whole assemblage, but it is the part the assemblage cannot run without, and an assemblage monism that dissolved the brain into the network would be as one-sided as the individualism it replaced.

It does not abolish individual differences, and it specifies what they are differences in. The variance the tests capture is real, but it is residual by construction, what remains after the testing room has held the scaffold constant, stripped the tools, forbidden the collaborators, and equalized the notation everyone was already trained on. The test purifies on purpose, and then reports the purified residue as the person's intelligence, which is the Constitution operating in miniature (Gould, 1981; Danziger, 1990). To measure capacity that way is legitimate. To read the measurement as the location of intelligence is the error the paper names.

It does not abolish agency, genius, or contribution, and it redescribes them as coupling rather than self-sufficiency. The navigator, the mathematician, the programmer each do something indispensable, and removing them collapses the performance as surely as removing the notation does. What changes is the grammar: their contribution is a contribution to an assemblage, the skilled coupling of a trained body to a system that exceeds it, and skilled coupling is a real, unequal, creditable thing. Tacit knowledge, the part of expertise that cannot be written down and must be grown through participation, is precisely the skill of coupling, and it is no less the person's for being unstatable (Polanyi, 1966; Collins, 2010; Lave and Wenger, 1991; Dreyfus, 1992).

And it does not collapse the difference between human and artificial intelligence, though it explains why the two now look alike. They look alike because the attribution error is the same, a system's output booked to an agent, and not because the systems are. The human is scaffolded by a culture it is born into and trained by bodies; the model is scaffolded by a corpus it is fed and tuned by labor. The parity is in the bookkeeping, not in the cognition, and noticing the shared bookkeeping is the cheapest available cure for the mysticism on both sides, the genius-mysticism we apply to people and the intelligence-mysticism we are learning to apply to models.

There is a politics in this that the paper will not pretend to resolve, because it is older than the paper's terms. Every argument now under way about artificial intelligence, about who built it and who owns it and whose labor is compressed inside it, is also an argument about a natural kind, intelligence, that we have been purifying into individuals for as long as we have had examinations and prizes and wages to attach to them. The machine has not created the question of how to credit a thinking system. It has reopened a settlement, and the settlement was always about credit, never about minds.

References

Adams, F., and Aizawa, K. (2008). *The Bounds of Cognition*. Blackwell.

Boyd, R., and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press.

Carson, J. (2007). *The Measure of Merit: Talents, Intelligence, and Inequality in the French and American Republics, 1750–1940*. Princeton University Press.

Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

Clark, A., and Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.

Collins, H. (2010). *Tacit and Explicit Knowledge*. University of Chicago Press.

Danziger, K. (1990). *Constructing the Subject: Historical Origins of Psychological Research*. Cambridge University Press.

Daston, L., and Galison, P. (2007). *Objectivity*. Zone Books.

Donald, M. (1991). *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Harvard University Press.

Dreyfus, H. L. (1992). *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press.

Gallagher, S. (2005). *How the Body Shapes the Mind*. Oxford University Press.

Galison, P. (1997). *Image and Logic: A Material Culture of Microphysics*. University of Chicago Press.

Gould, S. J. (1981). *The Mismeasure of Man*. W. W. Norton.

Hayles, N. K. (2017). *Unthought: The Power of the Cognitive Nonconscious*. University of Chicago Press.

Henrich, J. (2016). *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press.

Heyes, C. (2018). *Cognitive Gadgets: The Cultural Evolution of Thinking*. Harvard University Press.

Hoel, E. P. (2017). When the map is better than the territory. *Entropy*, 19(5), 188.

Hoel, E. P., Albantakis, L., and Tononi, G. (2013). Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49), 19790–19795.

Hollan, J., Hutchins, E., and Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174–196.

Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

Kirsh, D., and Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4), 513–549.

Klein, B., and Hoel, E. (2020). The emergence of informative higher scales in complex networks. *Complexity*, 2020, 8932526.

Knorr Cetina, K. (1999). *Epistemic Cultures: How the Sciences Make Knowledge*. Harvard University Press.

Latour, B. (1987). *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press.

Latour, B. (1993). *We Have Never Been Modern* (C. Porter, Trans.). Harvard University Press.

Latour, B., and Woolgar, S. (1986). *Laboratory Life: The Construction of Scientific Facts*. Princeton University Press.

Lave, J., and Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.

Levin, M. (2022). Technological approach to mind everywhere: an experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience*, 16, 768201.

Levin, M. (2025a). Ingressing minds: causal patterns beyond genetics and environment in natural, synthetic, and hybrid embodiments. *PsyArXiv* preprint, https://doi.org/10.31234/osf.io/5g2xj.

Levin, M. (2025b). *Free Lunches: Model Systems for Studying the Agential Gifts from the Platonic Space* [lecture], https://www.youtube.com/watch?v=K8BmMU1Tm-I. The slide stating that the intelligence ratchet is a gift from mathematics is credited to Federico Pigozzi.

Malafouris, L. (2013). *How Things Shape the Mind: A Theory of Material Engagement*. MIT Press.

Netz, R. (1999). *The Shaping of Deduction in Greek Mathematics: A Study in Cognitive History*. Cambridge University Press.

Noë, A. (2004). *Action in Perception*. MIT Press.

Pigozzi, F., Goldstein, A., and Levin, M. (2025). Associative conditioning in gene regulatory network models increases integrative causal emergence. *Communications Biology*, 8, 1027.

Pigozzi, F., and Levin, M. (2026). The causally emergent alignment hypothesis: causal emergence aligns with and predicts final reward in reinforcement learning agents. *arXiv* preprint arXiv:2605.06746.

Polanyi, M. (1966). *The Tacit Dimension*. University of Chicago Press.

Rotman, B. (2000). *Mathematics as Sign: Writing, Imagining, Counting*. Stanford University Press.

Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.

Shapin, S. (1994). *A Social History of Truth: Civility and Science in Seventeenth-Century England*. University of Chicago Press.

Sterelny, K. (2010). Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465–481.

Sterelny, K. (2012). *The Evolved Apprentice: How Evolution Made Humans Unique*. MIT Press.

Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press.

Tomasello, M., Carpenter, M., Call, J., Behne, T., and Moll, H. (2005). Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences*, 28(5), 675–691.

Vanchurin, V., Wolf, Y. I., Katsnelson, M. I., and Koonin, E. V. (2022). Toward a theory of evolution as multilevel learning. *Proceedings of the National Academy of Sciences*, 119(6), e2120037119.

Varela, F. J., Thompson, E., and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.

Walker, S. I., and Davies, P. C. W. (2013). The algorithmic origins of life. *Journal of the Royal Society Interface*, 10(79), 20120869.

Watson, R. A., and Szathmáry, E. (2016). How can evolution learn? *Trends in Ecology and Evolution*, 31(2), 147–157.