

WHY A MIND NEED NOT BE INTELLIGENT, AND AN INTELLIGENCE NEED NOT BE A MIND

The Competence and the Cogito

Elan Moritz

Eagles Perch Press • Philadelphia, PA

July 15, 2026

Keywords: intelligence; mind; consciousness; intentionality; artificial intelligence; functionalism; Turing test; sentence

About this issue

This essay opens the inaugural issue of the *Philadelphia Journal of Artificial and Digital Minds*, a five-part sequence that builds a single analytical framework for thinking about artificial minds. The framework separates three properties that ordinary talk of “intelligence” and “mind” runs together—*competence*, *agency*, and *mind*—and then puts the separation to work. This first movement makes the opening cut: it prises *competence* apart from *mind*. The movements that follow introduce agency as a third, independent property; show the framework doing work on the governance of capable systems; draw out its opposite normative consequences for mind and for agency; and finally turn the instrument on a live public debate about the dangers of artificial intelligence. Each movement stands on the ones before it; read in order, they are one argument.

Abstract

The phrase “artificial intelligence” contains a buried assumption so familiar that it is almost never examined: that building a mind and building an intelligence are the same project. This essay takes seriously the question of what happens if they are not. It argues, first, that *intelligence* as the field now defines it—the capacity to achieve goals across a wide range of environments—is a notion of *competence*, and that competence makes no reference to experience or to the directedness of thought toward objects. It argues, second, that the two properties which survived the collapse of the substance tradition as the marks of mind—phenomenal experience and intentionality—are not forms of competence and are not entailed by it. From these two claims a disorienting consequence follows: one could in principle build an arbitrarily intelligent system that is not a mind, and conceive of a barely intelligent system that is. The essay distinguishes a modest reading of this result (mind and intelligence are *contingently* linked, sharing a floor but not a ceiling) from a radical one (they are *orthogonal*, co-occurring in us only for evolutionary reasons), and argues that even the modest reading suffices to show that “artificial intelligence” names a competence project which may be building something other than, and possibly disjoint from, an artificial mind. The term, in short, may be a category error frozen into an acronym. As the opening movement of this issue, the essay makes the first of three cuts on which the rest of the framework is built.

1 The Assumption in the Acronym

We call the enterprise “artificial intelligence” and its hoped-for product a “mind,” as though these were one thing. The premise of this journal is that they are not, and that keeping them apart is the first discipline of clear thought about machines. This essay makes the first separation: intelligence, as the field defines it, is competence, and competence is not mind.

There is a phrase we use so often that its structure has become invisible to us. We speak of *artificial intelligence*, and we ask, in the same breath and as though it were the same question, whether the systems we are building have, or will have, *minds*. The slide from the first term to the second is so smooth that it is rarely noticed as a slide at all. But it embeds an assumption, and the assumption may be false. It is that intelligence and mind are the same thing, or at least that to build enough of the one is to bring about the other—that a sufficiently capable system is, by that very capability, a candidate for having an inner life.

This journal begins from the suspicion that our vocabulary for minds and machines fuses together properties that are in fact distinct, and that a great deal of confusion follows from the fusion. Its method, across its issues, will be separation: pulling apart the different things a single word is made to carry, so that questions which seem hopelessly tangled can be posed cleanly. The framework this inaugural issue builds will distinguish three such properties—*competence*, the capacity to achieve goals; *agency*, the capacity to act; and *mind*, the presence of experience and of genuine aboutness. The claim that will emerge is that these three vary independently, so that a system may possess one in abundance while wholly lacking another.

This first essay makes the opening cut, and it is the most fundamental one, because everything downstream depends on it: the separation of *competence* from *mind*. If that separation holds, then “artificial intelligence,” read literally, names a project in competence—and whether that project is also, or ever, a project in mind becomes an open question rather than a settled identity. The task here is to show that the separation holds.

2 What Intelligence Became

To see that “artificial intelligence” is a competence project, one need only attend to how the field has learned to define its central term. The definition that has done the most work in the modern era is deceptively simple: intelligence is the capacity to achieve goals across a wide range of environments (Legg and Hutter 2007a,b). Sharper operationalizations refine this—measuring skill-acquisition efficiency, or performance under bounded resources (Chollet 2019)—but they share a common shape. Intelligence, on all of them, is a matter of *what a system can accomplish*: how well it selects actions that bring about its ends, how broadly that competence generalizes, how efficiently it is acquired. Nothing in any of these definitions makes reference to whether there is something it is like to be the system, or to whether the system’s internal states are genuinely *about* anything.

This is not an oversight to be corrected; it is the definition working as intended. The turn to a performance-based conception of intelligence was, historically, a liberation. Alan Turing’s famous proposal was precisely to *replace* the unanswerable question “can machines think?” with the tractable question of whether a machine could succeed at a conversational imitation (Turing 1950). The substitution was deliberate and fruitful: it exchanged an inquiry into the machine’s inner nature for an inquiry into its outward accomplishment. Thinking was set aside; *doing* was put in its place. The whole subsequent science of machine intelligence inherited that exchange, and its progress has been the progress of a science of competence—of systems that achieve, increasingly, what we set them to achieve.

Intelligence, as the field defines it

The working definition is a definition of **COMPETENCE**: the capacity to achieve goals across environments, measured by success, breadth, and efficiency. It is silent, by construction, on two further questions—whether the system *experiences* anything, and whether its states are genuinely *about* anything. That silence is not a defect of the definition. It is the mark of what the definition is a definition *of*: not mind, but the accomplishment of ends.

The point is not that the field's definition is wrong. It is that the definition is a definition of *competence*, and that competence, so defined, contains no clause about experience or aboutness. If mind involves those things, then the definition of intelligence has been constructed, from the start, so as not to require mind. To build an intelligence, in the field's own terms, is to build a competence—and whether one has thereby built a mind is a question the definition does not answer, and was not designed to answer.

3 What Mind Kept

If intelligence became competence, what did *mind* become? The concept has a long and shifting history—the subject of this journal's parent volume, which traces it from prehistory to the present (Moritz 2026)—and across that history many conceptions of mind have risen and fallen: mind as soul, as substance, as immaterial thinking thing. What is striking is that when the substance tradition collapsed, and mind could no longer be identified with a special kind of *stuff*, two properties survived the collapse as the marks by which mind is still known. Neither is a form of competence.

The first is **PHENOMENAL EXPERIENCE**: the fact that, for a conscious creature, there is something it is like to be it (Nagel 1974). There is a felt quality to seeing red, to tasting salt, to feeling pain—a subjective character that a full description of the associated behavior and information processing seems to leave out (Block 1995; Chalmers 1995). This is the aspect of mind that the hard problem of consciousness is hard about, and it is not a matter of what a system can do. A system might process wavelength information and act flawlessly upon it with nothing it is like to do so; the experience is a further fact, not guaranteed by the competence.

The second is **INTENTIONALITY**: the directedness of mental states toward objects, their being *about* something (Brentano 1874). A thought about Paris, a belief that it will rain, a desire for water—each points beyond itself to a content. And here a distinction that this journal will lean on throughout makes its first appearance: the difference between *original* and *derived* intentionality (Dennett 1987; Searle 1983). The words on a page are about things, but only derivatively: their aboutness is lent to them by writers and readers, and there is no aboutness resident in the ink. Original intentionality is aboutness a system has *on its own account*, not by an interpreter's loan. A mind, in the demanding sense, has original intentionality; a book, or a thermostat, or—the open question of our moment—a language model, may have only the derived kind (Harnad 1990; Searle 1980).

The two marks of mind

When mind ceased to be identified with a special substance, two properties remained as its marks. **Phenomenal experience:** there is something it is like to be the system. **Original intentionality:** the system's states are genuinely about things, on its own account rather than by an interpreter's loan. Neither is a capacity to achieve goals; neither is entailed by such a capacity. Both are silent in the definition of intelligence, and present in the concept of mind. (Whether these two marks are themselves one property or two—and how they depend on one another—is a deeper question this review takes up in later work; for the present argument it is enough that both lie outside competence.)

Two things about these marks matter for the argument. First, neither is a form of competence: to experience is not to accomplish, and to be genuinely about something is not to succeed at a task. Second, neither is *entailed* by competence: no measure of what a system can do settles, by itself, whether there is something it is like to be it or whether its states are originally about anything. The marks of mind lie outside the definition of intelligence, and are not delivered by satisfying it.

4 The Floor and the Ceiling

Put the two results together. Intelligence, as the field defines it, is competence, and makes no reference to experience or aboutness. The marks of mind are experience and aboutness, and are not forms of competence. It follows that intelligence and mind are, at the level of concepts, distinct—and the distinctness has a disorienting consequence, which can be stated as a pair of possibilities.

On one side: one could, in principle, build a system of arbitrarily high competence—superhumanly capable across an enormous range of goals—for which there is nothing it is like to be, and whose internal states are about the world only in the derived sense, on loan from us. A vast intelligence that is not a mind. On the other side: one can conceive of a system, or a creature, of very modest competence—barely able to achieve anything across few environments—for which there nonetheless is something it is like to be, and whose states are genuinely its own. A small mind that is barely an intelligence. The animal world may already furnish the second case: a creature need not be clever to suffer (Birch et al. 2020; Godfrey-Smith 2020).

How far apart can the two come? Here the argument admits two readings, and it is worth being explicit about which is being claimed.

Two readings of the separation

The modest reading. Mind and intelligence are *contingently* linked. They may share a common floor—perhaps a certain minimal competence is required for the kinds of experience we know— but they do not share a ceiling: above the floor, competence can grow without limit while mind does not follow, and mind can be present where competence is slight. **The radical reading.** Mind and intelligence are *orthogonal*—fully independent axes, co-occurring in human beings only because evolution happened to build our competence and our sentience into the same organism, with no deeper necessity linking them at all (Bostrom 2014; Shevlin and Halina 2019). The essay does not need the radical reading. The modest one is enough for its conclusion, because even a shared floor with no shared ceiling means that “more intelligence” does not entail “more mind,” and that the two projects can diverge without limit as capability grows.

The distinction between the readings matters, because much confusion in the public discussion of machine minds comes from sliding between them—from treating a merely contingent link as a necessary one, and inferring from a system’s rising competence that its claim to mind must be rising too. The argument here is deliberately conservative: it rests only on the modest reading, and the modest reading already breaks the inference. A system’s growing competence is not, by itself, growing evidence that it has a mind, because competence and mind are distinct properties, and the growth of one is not the growth of the other.

5 The Consequence for “Artificial Intelligence”

Return now to the phrase we began with. “Artificial intelligence,” read in the field’s own terms, names a project in *competence*: the construction of systems that achieve goals across environments, with rising skill, breadth, and efficiency. That project is real, and it is succeeding. But it is a competence project, and the argument of this essay is that competence is not mind and does not entail it. So the phrase carries, in its second word, a quiet promissory note it may not be able to honor: it suggests that what is being built is, or is on the way to becoming, a mind, when what is demonstrably being built is a competence, whose relation to mind remains open.

The category error frozen into an acronym

If intelligence is competence, and competence is not mind, then “artificial intelligence” names one thing (a competence project) while continually inviting us to expect another (an artificial mind). The two need not coincide, and on the modest reading they can diverge without limit. The phrase is not false—we are indeed building artificial *intelligence*, in the competence sense—but it is misleading, because it lets the achievement of competence stand in, rhetorically, for the achievement of mind.

This is not a deflationary claim about machine minds, and it should not be mistaken for one. Nothing here shows that an artificial system *cannot* have a mind; the argument is neutral on that question, and this journal will hold it open with care. What the argument shows is narrower and firmer: that the *competence* we can measure, and are building, is not the same as the *mind* we might wonder about, and that the presence of the first is not evidence for the presence of the second. To ask whether a machine has a mind is not to ask how capable it is. It is to ask a different question, on a different axis—and the first task of a clear account is to keep the two questions from collapsing into one.

6 Objections

“This assumes what it must prove: that experience and aboutness are not just very sophisticated competences.” The objection is that a functionalist might identify mind *with* the right kind of information processing, so that sufficient competence of the right sort simply *is* mind, and the separation collapses. But the essay does not need to refute functionalism; it needs only the weaker point that the field’s *definition* of intelligence does not build in the relevant processing. Even if some special functional organization were sufficient for mind, the capacity to achieve goals across environments does not specify that organization, and systems can satisfy the competence definition without instantiating it. The separation is between *intelligence as defined* and *mind*, and it holds whatever the ultimate metaphysics of mind turns out to be. A functionalist who accepts the separation simply locates mind in a particular kind of competence-realizer—which is to grant that competence in general is not yet mind.

“If competence and mind can come apart, why do they so reliably occur together in us?” Because evolution built them into the same organism, and for creatures like us the two have a common biological substrate—which is exactly why the modest reading posits a shared floor. But co-occurrence in the one case we know well is weak evidence for necessary connection in general, and it is precisely the kind of parochial inference the study of artificial minds must resist. That two properties always travel together in humans does not show they must travel together in systems built on entirely different principles. The reliable human co-occurrence is a fact about human beings, not a law about minds.

“Isn’t this merely verbal—a quarrel about how to use the word ‘intelligence’?” It would be merely verbal if nothing turned on it. But a great deal turns on it: whether rising machine capability is evidence of emerging machine experience, whether the systems we deploy are candidates for moral status, whether “it is very intelligent” is any kind of answer to “does it have an inner life.” The separation of competence from mind is what allows these to be treated as the distinct questions they are. Far from a quarrel about words, it is the precondition for asking the questions that matter without the answer to one being smuggled in as the answer to another.

7 Conclusion

Intelligence, as the field defines it, is competence—the capacity to achieve goals across environments—and competence makes no reference to experience or to original aboutness. The marks of mind are experience and aboutness, and are not forms of competence, nor entailed by it. So one could build a vast intelligence that is no mind, and conceive a slight intelligence that is one. “Artificial intelligence” names the competence project; whether it is also a mind project is the open question this separation lets us finally pose.

The first cut has been made. Intelligence and mind, fused in the phrase that names our era’s central technology, come apart under inspection: the one is a competence, defined by accomplishment; the other is marked by experience and by genuine aboutness, neither of which accomplishment supplies. The consequence is not that machines cannot have minds, but that their competence—however vast, however rapidly growing—is not itself their mind, and is not evidence of it. The question of the machine mind survives the separation intact, and is in fact clarified by it: freed from the assumption that capability and mindedness are the same thing, we can at last ask, as a real and open question, whether the competent systems we are building experience anything, or intend anything on their own account.

That question cannot be answered yet, and this essay has not tried to answer it. It has tried to earn the right to ask it, by prising apart the two properties our vocabulary had welded together. But competence and mind are not the only properties so welded. Between them lies a third—the

capacity to *act*, to be an agent in the world—which is neither pure competence nor mind, and which the next movement of this issue introduces as an independent axis. The separation begun here is the first of several, and the framework it opens is completed only when all three properties—competence, agency, and mind—have been set apart and can be held, at last, in view at once (Moritz 2026).

Acknowledgement. The author makes extensive use of frontier large language models, as well as local large language models, in the course of this work.

Conflict of interest. The author declares no conflict of interest.

Funding. This work received no external funding.

References

- Birch, J., A. K. Schnell, and N. S. Clayton (2020). "Dimensions of Animal Consciousness". In: *Trends in Cognitive Sciences* 24.10, pp. 789–801. DOI: [10.1016/j.tics.2020.07.007](https://doi.org/10.1016/j.tics.2020.07.007).
- Block, N. (1995). "On a Confusion about a Function of Consciousness". In: *Behavioral and Brain Sciences* 18.2, pp. 227–247. DOI: [10.1017/S0140525X00038188](https://doi.org/10.1017/S0140525X00038188).
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Brentano, F. (1874). "Psychology from an Empirical Standpoint". In: English edition, ed. Linda McAlister, Routledge, 1973; source of the intentionality thesis.
- Chalmers, D. J. (1995). "Facing Up to the Problem of Consciousness". In: *Journal of Consciousness Studies* 2.3, pp. 200–219.
- Chollet, F. (2019). "On the Measure of Intelligence". In: *arXiv preprint arXiv:1911.01547*. DOI: [10.48550/arXiv.1911.01547](https://doi.org/10.48550/arXiv.1911.01547).
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.
- Godfrey-Smith, P. (2020). *Metazoa: Animal Life and the Birth of the Mind*. New York: Farrar, Straus and Giroux.
- Harnad, S. (1990). "The Symbol Grounding Problem". In: *Physica D* 42.1–3, pp. 335–346. DOI: [10.1016/0167-2789\(90\)90087-6](https://doi.org/10.1016/0167-2789(90)90087-6).
- Legg, S. and M. Hutter (2007a). "A Collection of Definitions of Intelligence". In: *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms*. Ed. by B. Goertzel and P. Wang. Vol. 157. Frontiers in Artificial Intelligence and Applications. IOS Press, pp. 17–24.
- (2007b). "Universal Intelligence: A Definition of Machine Intelligence". In: *Minds and Machines* 17.4, pp. 391–444. DOI: [10.1007/s11023-007-9079-x](https://doi.org/10.1007/s11023-007-9079-x).
- Moritz, E. (2026). *Conceptions of Mind: A Historical Census*. Frontiers of Intelligence. Philadelphia: Eagles Perch Press.
- Nagel, T. (1974). "What Is It Like to Be a Bat?" In: *The Philosophical Review* 83.4, pp. 435–450. DOI: [10.2307/2183914](https://doi.org/10.2307/2183914).
- Searle, J. R. (1980). "Minds, Brains, and Programs". In: *Behavioral and Brain Sciences* 3.3, pp. 417–424. DOI: [10.1017/S0140525X00005756](https://doi.org/10.1017/S0140525X00005756).
- (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Shevlin, H. and M. Halina (2019). "Apply Rich Psychological Terms in AI with Care". In: *Nature Machine Intelligence* 1.4, pp. 165–167. DOI: [10.1038/s42256-019-0039-y](https://doi.org/10.1038/s42256-019-0039-y).
- Turing, A. M. (1950). "Computing Machinery and Intelligence". In: *Mind* 59.236, pp. 433–460. DOI: [10.1093/mind/LIX.236.433](https://doi.org/10.1093/mind/LIX.236.433).