What Do You Think About AI?

Rob Tow

Nova Lux, New Mexico, USA, Sol III

6 September 2026

A friend at a cocktail party asked what I think about AI. One might as well have asked in 1880 what one thinks about electricity: enough had happened by then to know it would matter, and almost nothing to show what for. Most forecasts assume the thing which evolves will go on being the thing we presently call an AI; I doubt it. Biology suggests otherwise. Whole evolved systems can combine rather than climb, and the individual has changed level before. Beside Edison's lamp one could predict better lamps. What one could not predict was radio.

At a recent cocktail party a friend asked me, "What do you think about AI?" I was bemused. It seemed rather like asking someone in 1880, "What do you think about electricity?"

By 1880 this would not have been an entirely foolish question. Telegraph wires crossed continents, the telephone existed, dynamos and electric motors existed, and Edison had produced a practical incandescent lamp. An intelligent observer had ample reason to conclude that electricity was going to matter, while having almost no basis for understanding what it was going to matter *for*. There was no radio, electronic amplification, television, radar, electronic computer, transistor, integrated circuit, Internet, GPS receiver, or mobile telephone. Factories had not yet been reorganized around distributed electric motors, houses had not filled with electrical appliances, and cities had not acquired the electrical nervous systems we now regard as infrastructure.

Someone looking at an incandescent lamp in 1880 could reasonably imagine a better lamp, just as someone looking at a telephone could imagine a better telephone. This would have been entirely sensible forecasting, and it would have missed most of the future.

Michael Faraday is famously supposed to have answered a question about the usefulness of electricity by asking, "What use is a newborn baby?" The provenance of the remark is untidy; Faraday himself attributed substantially the same observation to Benjamin Franklin. Scientific aphorisms, like successful genes, sometimes migrate to more favorable hosts. The metaphor remains useful because the difficulty with a newborn baby is that extrapolating from its present abilities tells you rather little about what it may become.

Artificial intelligence in 2026 is no longer a newborn. Artificial systems recognize and generate speech and images, write useful programs, predict molecular structures, discover useful patterns in enormous bodies of data, operate machinery, navigate the physical world, control complex processes, use tools, and converse about almost anything. Some outperform humans within narrow but consequential domains; others combine capabilities that until recently belonged to quite different kinds of machinery. They also hallucinate, misclassify objects a child would recognize, fail in surprising ways outside their training distributions, and occasionally exhibit the serene confidence of a drunk who has just discovered Wikipedia. There is enough there to know that something important has happened, without pretending that we yet know what sort of thing has happened.

This has naturally produced an enormous industry devoted to predicting what happens next. Models will become larger, or smaller, faster, cheaper, more reliable, and more agentic. They will replace programmers, or make programmers fantastically productive, or end civilization, or produce excellent personalized advertising immediately before ending civilization. What interests me about many of these forecasts is less whether they are right than an assumption they share: the thing which evolves will continue to be the thing we presently call an "AI".

I doubt it.

The Wrong Individual

There is another way to look at the problem, and it comes from biology. We often describe evolution as a search across a fitness landscape: variations appear, some work better than others in the environment in which they occur, successful variants reproduce, and a population consequently explores neighboring regions of possibility. It is a powerful model, but it can encourage an image of evolution as patient hill-climbing.

Biology occasionally cheats, and one of its more consequential cheats is symbiogenesis. The familiar examples are mitochondria and chloroplasts, both descended from organisms which once possessed substantially independent evolutionary histories. Those organisms entered into associations with other cells, and over immense spans of time the distinction between host and symbiont largely disappeared. What had once been one organism living inside another became an organelle.

The interesting point is not merely that the resulting cell was better at something. Each participant brought to the association an enormous package of structure already produced by previous evolution, so one lineage did not have to rediscover, mutation by mutation, all of the machinery already present in the other. Evolution had found a way to reuse the results of searches it had already performed in other lineages; repackaged and resituated. The consequence was not simply movement to a higher point on an existing fitness landscape. The combination made possible trajectories that had not previously been available to either lineage, altering the topology and even the dimensionality of the space being searched. This is a much more interesting evolutionary maneuver than climbing the neighboring hill a little faster.

Nor is it confined to events buried almost beyond recognition in the Precambrian. In 2024 researchers presented strong evidence that a cyanobacterial symbiont of the marine alga *Braarudosphaera bigelowii* has crossed the boundary into organellehood. The cyanobacterium, UCYN-A, fixes nitrogen, divides in coordination with its host, and imports proteins encoded by the host nucleus. The resulting structure has been called a *nitroplast*. The association appears to be on the order of ninety million years old, almost indecently recent compared with mitochondria, allowing us to catch organellogenesis in the act. A bacterium which once had an existence of its own has become sufficiently integrated into another cell that asking which organism it belongs to is no longer quite the useful question. Evolution has changed the individual.

This is a pattern; not a design pattern – an emergent pattern, repeated across time.

Borrowing Parts

Biology has another trick which may be even more suggestive for what is happening with artificial systems. Some organisms do not permanently incorporate another organism at all; they eat it, digest much of what they have eaten, and retain the useful machinery and use it almost like a tool.

This is called *kleptoplasty*. A number of protists and a few animals ingest photosynthetic organisms but retain their plastids, which continue performing photosynthesis inside their new owners. In some sea slugs the stolen plastids remain functional for weeks or months. The slug has not evolved a chloroplast; it has temporarily acquired one from lunch, and (partially) lives on sunlight.

Some dinoflagellates make the arrangement still more baroque. Species of *Dinophysis* acquire plastids by eating the ciliate *Mesodinium rubrum*, which in turn acquired those plastids by eating cryptophyte algae. Evolution, having discovered code reuse, appears also to have invented transitive dependency management. Very matryoshka.

These arrangements suggest a spectrum rather than a simple distinction between predator and symbiont. An organism may consume another and temporarily appropriate one of its capabilities; it may maintain a symbiont for an extended period; the association may become an obligate dependency; reproduction and metabolism may become coordinated; genes and proteins may migrate across the old boundary; and eventually the distinction between the two organisms may become largely an historical fact while a new kind of individual gets on quite contentedly with being one. There is no reason every association must travel this road, and evolution is not pursuing an architectural objective, but the examples demonstrate that the boundary around an organism is considerably less sacred than our nouns suggest.

Software has an interesting advantage in this sort of game: we deliberately build interfaces.

Evolution Discovers the API

Two biological organisms attempting an intimate association confront membranes, immune defenses, incompatible metabolisms, reproductive coordination, protein transport, and the persistent possibility that one partner will digest the other. A computer program can make an API call. And, increasingly, it can simply talk to the other system.

A language model can acquire a search engine; coupled to a database, it acquires another kind of memory; connected to a compiler, it can create an artifact, execute it, inspect what happened, and try again. A camera gives it another sensory surface, while machinery gives it a means of altering the physical world. For the duration of a particular task it can acquire symbolic mathematics, satellite imagery, financial data, another model, a robot, or a human being, and relinquish the connection when it is no longer useful. It can access the internet and the World Wide Web, noticing the world as the timeline unfolds.

This begins to look less like the development of a single increasingly capable machine and more like computational kleptoplasty. The functional boundary of the system can expand and contract according to the problem it is solving. Encapsulations occur, and associations. This matters because most of the systems being connected are not trivial components waiting to be invented. Search engines, compilers, databases, scientific instruments, communications networks, corporations, bureaucracies, and human beings are themselves the products of long periods of development and selection. Connecting them makes previously accumulated complexity available wholesale, without requiring the new assemblage to reproduce the path by which that complexity arose.

Sometimes such an association will be merely convenient and transient; sometimes it will persist because both sides benefit from it; sometimes it may become so tightly coupled that continuing to describe each component as the interesting object obscures what the larger system is doing. Biology has played this game often enough that we should at least consider the possibility that we are looking at the wrong idea of what constitutes the individual.

What Is the Organism?

Evolution has repeatedly changed the answer to the question "What is an individual?" Genes cooperate inside chromosomes, cells cooperate inside multicellular organisms, and individual insects create colonies capable of behavior no individual ant or bee possesses. At each transition, some local autonomy is constrained while capabilities become available at another level of organization. The topology changes, and selection acquires another kind of object upon which to operate.

This is one reason I find the popular image of an approaching artificial superintelligence oddly unimaginative. It generally preserves the existing ontology: there is a machine, the machine becomes progressively smarter, and eventually the machine becomes smarter than we are. Perhaps that will happen, but it strikes me as the equivalent of standing beside Edison's lamp and predicting an exceptionally bright light bulb.

The more interesting possibility is that the unit of organization changes. Consider a system containing language models, conventional programs, databases, sensors, machinery, communications networks, and people. Different components observe different portions of the world, maintain shared state, allocate work, inspect results, and act with varying degrees of local autonomy. Their actions alter the environment, changing what they subsequently observe. Successful arrangements persist, unsuccessful ones are modified or discarded, and useful relationships become progressively more difficult to disentangle.

Now connect such systems to institutions which allocate money, material, computation, and human attention. At some point asking "What is the AI?" becomes rather like asking which ant is the colony. The model may be an organ, a symbiont, or a transiently stolen chloroplast; the interesting object is increasingly the organization within which it participates.

Once that organization can perceive, act, compare results against internal criteria, retain state, and modify its future behavior, another vocabulary becomes useful: cybernetics. Organisms, institutions, military organizations, markets, and companies all close feedback loops through the world. Their actions alter the environment, which alters what they perceive, which changes their subsequent actions. Artificial cognitive systems can now enter these loops at many points, participating in observation, orientation, decision, action, and the interpretation of consequences.

We are therefore not simply adding intelligence to machinery. We are constructing new feedback systems in the world, and those systems alter the environments in which subsequent systems are selected.

Landscapes That Move

Organisms have always modified their own fitness landscapes. Beavers alter rivers, plants altered Earth's atmosphere, and predators change prey populations which in turn change predators. Evolution takes place in an environment partly constructed by the things being evolved.

Technology performs the same trick rather faster. An institution adopts an artificial cognitive system because of an existing need; the system changes what the institution can do; the institution reorganizes itself around the new capability; that makes different capabilities valuable; new systems are built for the altered institution; and people, markets, laws, and expectations change in response. The environment selects the technology while the technology changes the environment which will select what comes next.

Symbiogenesis adds the possibility that substantially complete systems can combine, making previously selected complexity available without repeating the search which produced it. Kleptoplasty demonstrates that the acquisition can be temporary. Engineered interfaces reduce the cost of forming such associations, while machine-speed experimentation allows assemblages to be created, tested, dismantled, and recombined on timescales biology could only envy.

What begins to evolve under these conditions is not merely the model but the topology of the system around it. Variation produces selection; selection retains structures; retained structures combine; some combinations persist long enough to acquire further capabilities; and eventually a combination may become a new individual upon which selection itself can operate. Evolution can then begin another search at the new level.

This is one of the ways evolution creates new scales of organization.

A Ghost at the Banquet

There is a faint ghost of Teilhard de Chardin hovering around this argument. Teilhard imagined the biosphere giving rise to another layer of organization, the *noösphere*, as human minds, institutions, communications, and technologies became progressively interconnected. He eventually carried the idea toward the Omega Point and into theology, where I do not propose to follow him, but his structural intuition was more interesting than the metaphysical destination he assigned it.

The important entity, Teilhard suspected, might be getting larger. Writing in the 1940s, he described humanity's increasingly connected machinery and communications as something approaching a "Brain of brains," a larger organization emerging from the interaction of its constituents. One need not imagine a planetary consciousness, much less an ordained destination for evolution, to notice that higher-order entities are perfectly ordinary. An ant colony exists, as does a corporation, a city, and an ecosystem. None requires a little homunculus sitting at its center to possess structure, persistence, memory, resource flows, feedback loops, and characteristic behavior.

The mistake is to confuse an entity with a person.

What has changed since Teilhard is that artificial cognitive components have entered this increasingly connected system. They can perceive, represent, communicate, plan, and act, joining humans, machines, and institutions which were already coupled. Some of those connections will be temporary, some persistent, some dependencies, and some may eventually become so integrated that describing the constituents separately tells us less than describing the larger cybernetic organization they compose.

A larger entity may indeed be in the process of coming into being. There is no particular reason to expect that it will resemble us, share our boundaries, possess a unitary consciousness, or even notice the moment at which our preferred ontology ceases to describe it very well. And I see that we may become but as mitochondria in the body of the noösphere, as we immanentize the Gaian eschaton.

Back at the Cocktail Party

So I return to the innocent question at the cocktail party: "What do you think about AI?" Asking whether artificial intelligence is good or bad seems rather like asking in 1880 whether electricity is good or bad, but the analogy now seems insufficient. Electricity eventually disappeared into almost everything else, becoming less a recognizable class of object than an enabling substrate from which innumerable other things were constructed. Artificial intelligence may similarly disappear into larger systems, with the additional complication that the boundaries of those systems may themselves be evolving.

Faraday supplies part of the answer: we are looking at something early enough in its history that extrapolation from its present uses is foolish.

Darwin supplies another: the things constructed from it will vary and be selected. Symbiogenesis shows how previously evolved complexity can combine and abruptly open otherwise inaccessible evolutionary trajectories, while kleptoplasty reminds us that such combinations need not be permanent. Cybernetics closes the loop between these evolving organizations and the world they increasingly modify. Teilhard, once relieved of the obligation to provide a destination for all this, contributes one last unsettling possibility: the interesting entity may be getting larger than anything previously existing.

The cocktail-party question assumes that "AI" identifies the thing upon which I am being asked to pass judgment. That assumption seems decidedly ephemeral, and most of the answers currently being offered strike me as banal. The interesting answer is that we should expect massive surprise, because evolution has repeatedly changed not merely the capabilities of its creatures but the boundaries, relationships, and levels of organization that define the creatures themselves.

Someone standing beside an electric lamp in 1880 could sensibly predict better lamps. What he was unlikely to predict was radio.

I suspect we are standing beside the lamp.