



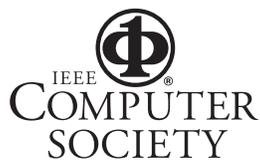
[www.computer.org/intelligent](http://www.computer.org/intelligent)

## Mindless Intelligence

*Jordan B. Pollack*

Vol. 21, No. 3  
May/June 2006

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.



© 2006 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

For more information, please see [www.ieee.org/portal/pages/about/documentation/copyright/polilink.html](http://www.ieee.org/portal/pages/about/documentation/copyright/polilink.html).

# Mindless Intelligence

Jordan B. Pollack, *Brandeis University*

**M**ost of my 25 years of professional involvement in AI have been focused on research far from its mainstream, not because of any antisocial tendencies on my part, but because of certain dilemmas inherent in the field. The first dilemma confronting AI is that both single-celled and multicelled animals survive and reproduce very

well without any nervous system at all, and “lower animals,” even insects, organize into thriving societies without any symbols, logic, or language, bee dancing and birdsong notwithstanding. These phenomena led me to delve into nonsymbolic models and ask how complex hierarchical representations and sustained state-changing procedures might naturally emerge from iterative numeric systems such as associative or connectionist neural networks.

The second dilemma is that the kind of mind we in AI seek to discover, one that “runs” on the human brain yet might be portable to another universal machine, wouldn’t even exist without having co-evolved with the brain—a chicken-and-egg problem. So, while many of my connectionist colleagues migrated with US National Institutes of Health funding into cognitive or computational neuroscience, trying to understand how the human brain works, I focused instead on what natural process could design and fabricate machinery as complex as the brain.

I ended up working closer to the field of artificial life, seeking to understand how evolution, a mindless iterative reproduction system, could eventually lead to machines whose complexity and reliability dwarfs the product of the largest teams of human engineers.

On this 50th birthday of artificial intelligence, I would like to reflect on what I feel has been its great mistake, and propose a corrective course for the next 50 years. But before analyzing this mistake, I want to say that AI is a great human endeavor with a colorful cast and many partial successes. It has provided frameworks for formally studying biological sys-

tems, animals, and humans and has spun out industries such as Lisp machines, expert systems, data mining, and even Internet search.

## Don’t promise the practically impossible

We all agree on AI’s fundamental hypothesis, that physical machines have the capacity for intelligence. Unfortunately, this hypothesis can neither be proven nor refuted scientifically, but realized only by demonstration. And until it has been convincingly demonstrated, it must remain in scientific limbo. Ordinary citizens and funding bureaucrats don’t know whether AI is *tardy*, like mechanical flight, which emerged from limbo after several hundred years of failure, or *magical*, like ESP or the alchemists’ quest to turn lead into gold. Perhaps there is even an *impossibility proof* waiting around the corner, as has put to rest quixotic notions such as time travel (Einstein) and perpetual motion (Ludwig Boltzmann). Who wants to fund a field that might be proven impossible tomorrow?

So AI, which represents one of the greatest intellectual and engineering challenges in human history—and should command the same fiscal resources as efforts to cure cancer or colonize Mars—is sometimes relegated to a laughingstock, because we can’t prevent bogus claims from cropping up in newspapers and books. We cannot seem to convince the public that humanoids and Terminators are just Hollywood special effects, as science-fictional as the little green men from Mars!

Still, some want to keep pursuing the same old AI

*AI has stalled because of its preoccupation with simulating the human mind. By studying intelligence in natural systems, outside the mind, we can reinvigorate the field.*

goals: “What are the missing pieces necessary to achieving human-level common sense?” “Let’s do a project to gain human-level performance in a (nonchess) domain.” “We will build natural language software that’s human-level in ability.” “Soon computers will be fast enough to supply human-level intelligence to humanoid robots.”

AI won’t be a gift of more CPU time. If it were, we would have already glimpsed real AI on supercomputers or large clusters, yet nothing of the kind has occurred. We don’t need faster chips to make robots smarter, since we can link a robot’s body to its supercomputer brain over wireless broadband. As the joke goes, even if AI requires an infinite loop, it should run in only five seconds on a supercomputer.

The issue isn’t the speed of running a mind-like program; it is the size and quality of the program itself. Because we routinely underestimate the complexity of evolved biological systems, and because Moore’s law doesn’t lead to a doubling of the quality of human-written software,<sup>1</sup> the same old goals are red herrings that promise the practically impossible!

### Take Mind off its pedestal

AI’s great mistake is its assumption that *human-level intelligence is the greatest intelligence that exists*, and thus, that our computational intelligences should operate “like” human cognition. Because of this mistake, most AI research has focused on “cognitive models” of intelligence, on programs that run like people think. But it turns out that we don’t think the way we think we think!

The scientific evidence coming in all around us is clear: Symbolic conscious reasoning, which is extracted through protocol analysis from serial verbal introspection, is a myth. From Michael Gazzaniga’s famous split-brain experiments, where a patient associated a snow shovel with a chicken,<sup>2</sup> through Daniel Dennett’s demolition of consciousness,<sup>3</sup> through the unconscious intelligence described most recently by Malcolm Gladwell,<sup>4</sup> it’s entirely clear that the “symbolic mind” that AI has tried for 50 years to simulate is just a story we humans tell ourselves to predict and explain the unimaginably complex processes occurring in our evolved brains.

Because of this preoccupation with mimicking human-level intelligence, as a scientific field, we’ve ignored or excluded the contributions of many alternative nonsymbolic mechanisms. Such mechanisms range from

associative and matrix models of mathematical psychology, to Markovian models, to both game and decision theories, to early neural networks (the perceptron disaster), to simulations of evolution and organic self-organization. The early success of low-hanging symbolic fruit through Lisp programming led to the pursuit of the “mythical man module,” a computer program that has the “look and feel” of human cognition yet is something more than an Eliza.

John Searle’s “Chinese Room” argument<sup>5</sup> is hateful because, in fact, he’s correct. Neither the room nor the guy in it pushing symbols “understands” Chinese. But this isn’t really a problem, because nobody actually “understands” Chinese! We only think we

The scientific evidence coming in all around us is clear: Symbolic conscious reasoning, which is extracted through protocol analysis from serial verbal introspection, is a myth.

understand it. As anyone—even a native speaker—drives further down into an explanation of his or her knowledge or behavior, instead of gaining sharper insights (as we might expect in a reductionist physical science with a better microscope), the explanations get blurrier and blurrier.

By assuming that intelligences based on human-centric cognitive architectures such as grammars or production systems are the zenith, are the most powerful intelligences in the world, our field has made the same kind of embarrassing mistake as today’s cryptocreationists, the proponents of Intelligent Design. By doubting that a mindless nonlinear iterative process such as evolution could be responsible for irreducible complexity in the designs of biological life-forms, they hold that a superhuman, super-intelligent being must have intervened.

AI also behaves as if human intelligence is next to godliness. Even the neural approach, more accepted today than ever, falls into the trap of trying to model human

cognitive structures such as verb conjugation. Why is simulating the human mind more important than simulating cellular metabolisms, insect or animal intelligence, complex pattern formation, or distributed control of complex ecologies? It must be because, as a mirror of our own intelligence, the mindless iterative and numeric computing we scientifically uncover in nature doesn’t compare to the perfectly logical indefatigable mind of Hollywood characters such as Mr. Spock and Commander Data, NP-completeness notwithstanding.

To repair this mistake and move forward as a scientific field, we must recognize that many intelligent processes in Nature perform more powerfully than human symbolic reasoning, even though they lack any of the mind-like mechanisms long believed necessary for human “competence.” Once we recognize this and start to work out these scaleable representations and algorithms without anthropomorphizing them, we should be able to produce the kind of results that will get our work funded to the level necessary for growth and deliver beneficial applications to society, without promising the intelligent English-speaking humanoid robot slaves and soldiers of science fiction.

### Defining mindless intelligence

I define “mindless intelligence” as intelligent behavior ascribed (by an observer) to any process lacking a mind-brain. Suppose some black-box process (for example, mathematical, numerical, or mechanical) exhibits behavior that appears to require intelligence. However, when we scientifically study it, we find no Lisp interpreter, no symbols, no grammars, no logic or inference engine—in fact, we realize that it works without any of the accoutrements of cognition. We can say that this process is mindlessly intelligent.

Now we can begin to seriously study intelligent performance by

- feedback-driven systems such as thermostats and steam governors;
- pattern-action systems such as Eliza programs and immune systems;
- stability and hierarchy networks such as cellular metabolisms;
- societal assemblies such as insect and colonial life-forms;
- utility-maximizing systems such as game and economic agents;
- exquisitely iterative systems such as evolution, fractals, and embryogenesis; and even

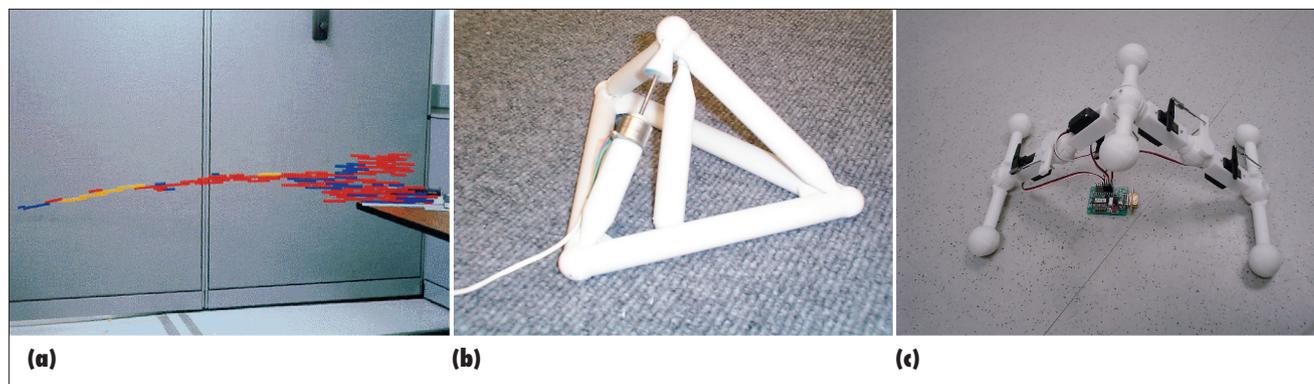


Figure 1. Three generations of evolved robots: (a) Pablo Funes' evolution of Lego discovered the cantilever,<sup>14</sup> (b) Hod Lipson's evolution of dynamic trusses invented the ratchet,<sup>15</sup> and (c) Gregory Hornby's evolution using L-systems to describe machines invented a kayaking motion.<sup>16</sup>

- mind-erasing collectives such as academic committees, crowds, and bureaucracies.

To give you a broader sense of the field, I'll briefly cover several kinds of natural processes that appear intelligent yet lack any cognitive apparatus. John Kolen and I showed how an iterated dynamical system could appear to generate a context-free or context-sensitive language, depending on the observer.<sup>6</sup> The dynamical system lacked any cognitive architecture for "generative capacity," which has been assumed by all natural language processing systems since Noam Chomsky.

Wherever we look in Nature, we see amazingly complex processes to which we can ascribe intelligence, yet we observe symbolic cognition in only one place, and only there as a result of introspection. Many of these natural processes have been studied under the aegis of complex systems or have been given the prefix "self" or "auto." Because these systems have no mind, and thus no self, I've taken the liberty of replacing those prefixes with the new term *ectomental*, which means outside (Greek) of mind (Latin).

### Ectomental organization

Evolution is the primary example of an intelligent designer who lacks a mind. There's no grammar, set of rules, library of CAD parts, or physics simulation. Simply put: a mindless reproductive system operates, transcription errors occur, and selection locks in a statistical advantage for the marginally better—or luckier—members of a population. And yet this iterative process has automatically designed machines of incredible beauty and complexity, objects that far

surpass—in complexity and reliability—anything architects, engineers, novelists, venture capitalists, or teams of software programmers can achieve.

Human teams can build systems with only 10 million to 100 million unique moving parts before the entire structure collapses, yet biological forms can have 10 billion unique moving parts.

For the past decade, my lab's goal has been to understand how evolution can produce more complex designs than a human engineering team, while lacking human-level symbolic cognition. We've focused specifically on coevolutionary machine learning systems. While we haven't yet achieved a fully open-ended design process, we have

- shown coevolutionary systems that have surpassed human performance in sorting networks and cellular-automata optimization;<sup>7</sup>
- developed theories such as Pareto coevolution,<sup>8</sup> emergent dimensionality,<sup>9</sup> and computational models of symbiogenesis;<sup>10</sup> and
- revealed the possibility of motivating a community of learners<sup>11</sup> to become their own Ideal Teachers,<sup>12</sup> resulting in novel educational software.<sup>13</sup>

Perhaps our best-known research is on the coevolution of robot bodies and brains, known as the Genetically Organized Lifelike Electro-Mechanics, or GOLEM, project. This research resulted in three generations of self-designed systems that discovered irreducibly complex components and processes such as the cantilever, ratcheting, and kayaking (see figure 1).

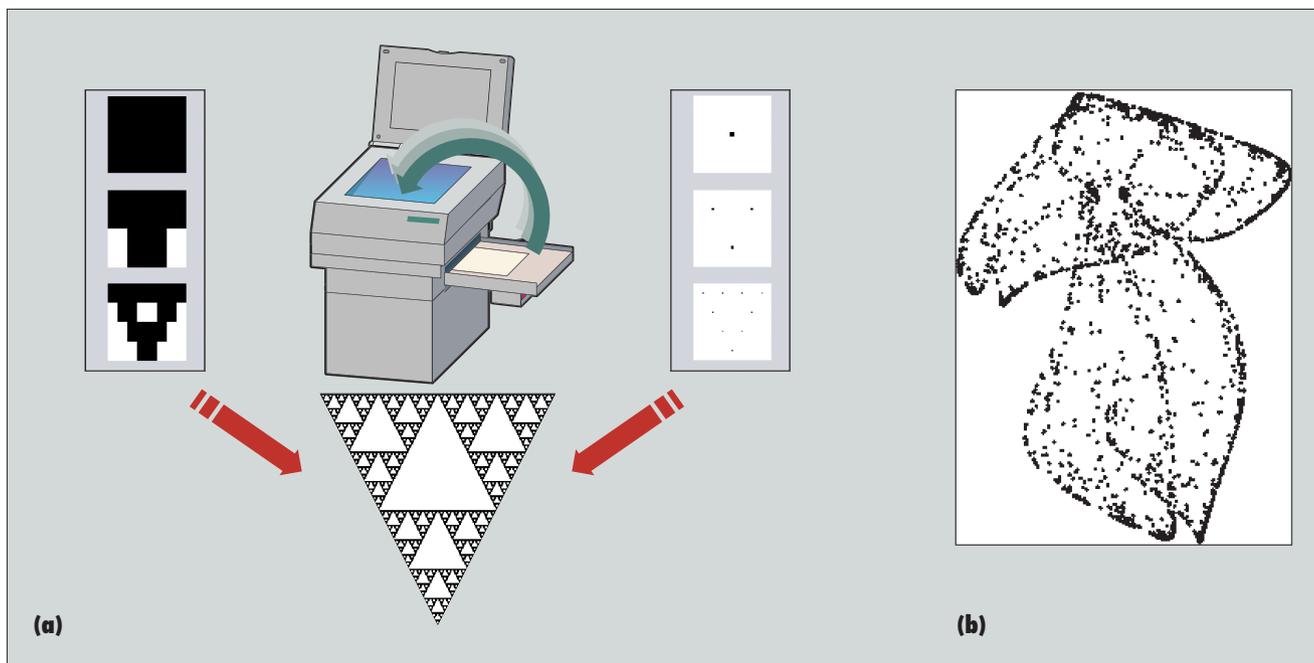
### Ectomental learning

One of the oldest AI paradigms is a self-learning or autodidactic system, a program that begins with a tabula rasa and, when dropped into an environment, gets better and better over time. Perhaps the best example of such a system is Gerald Tesauro's TD-Gammon.<sup>17</sup> He started with essentially a random neural network that could return a value for any backgammon position. Rather than training the network against an encyclopedia of human expert games, he essentially trained it against itself. After about a month of computer time on an IBM supercomputer, with the weights adjusted as a result of each game, his network, with further refinements, became one of the best players in the world.

Humans can verbalize backgammon strategies. We consider only a few plausible moves and then estimate whether one move is better or worse than another on the basis of strategic goals from models of the game (running, blocking, back-game), using all kinds of approximate and exact calculations about probability. I was a professional-level backgammon player in 1975 and felt that there were about seven different human-player "types" who, at the top of their game, achieved a rock-scissors-paper parity.

On the other hand, TD-Gammon is a mindless intelligence that dominates all human players. It uses a function to estimate values and uses a one- or two-ply look-ahead with a greedy selector to make a move. It has no logic or symbols, no strategy that looks far ahead or back in time, and no language component to discuss its strategy. Yet it's stronger than any rule-based strategy.

My lab had worked on self-learning for tic-tac-toe,<sup>18</sup> and we became interested in



**Figure 2. (a)** An iterated-function-systems fractal is like a feedback loop on a copy machine that makes more than one reduced copy of an image, resulting in the same limit for a speck of dust or a full page of ink. **(b)** The IFS theory explained the “strange automata” that emerged when recurrent neural networks were trained to recognize languages.

understanding why TD-Gammon worked. We were able to replicate the Tesauro effect using simple hill-climbing,<sup>19</sup> which led to the question of why coevolutionary self-learning worked so well for backgammon. Game theorists such as Richard Bellman recognized many years ago why a purely numeric backgammon player works better than a logical game.<sup>20</sup> He proved the existence of a value table for optimal sequential choice in Markovian games, where opponents can choose strategies yet are buffeted by random elements such as dice. Moreover, iterated approximation of the value table, through a single-ply expectimax look-ahead, leads to its convergence. So, an optimal value table combined with a one-ply greedy choice leads to the strongest-possible player.

In order to study the success of learning backgammon, I recently invented Nannon<sup>®</sup>, the smallest version of backgammon that maintains its core behaviors, yet only uses six points, three checkers, and one die per side. There are only 2,530 different board positions, and the value table converges in 15 sweeps to an error of  $10^{-7}$ .<sup>21</sup> While the full game of backgammon is much larger than Nannon, so a table can't be stored, Tesauro's choice of input representation and network size from earlier experiments led to a fortuitous convergence between TD reinforcement learning and Bell-

man's earlier mathematical work.

Perhaps many mindlessly intelligent processes in Nature are similar instances of mathematical ideals that can lead to convergence, complexity, and optimal performance in the limit.

### Ectomental repair

A marvelous characteristic of natural systems is that they can heal, or self-repair. A naïve computerized view would be to envision the algorithmic equivalent of a team of repairpersons who, under centralized supervision, consult a system model and are then deployed to a disturbance's site to apply cognition, logic, and spare parts to return the system to model behavior. However, imagining a system that contains a deployable model of itself can lead to logical conundrums.<sup>22</sup>

How might we understand self-repair in natural systems? In artificial-life research on “algorithmic chemistry,” Walter Fontana and Leo Buss described systems of simple lambda calculus programs that consume and produce each other, forming a metabolism.<sup>23</sup> When such an artificial-chemistry network had a steady-state dynamic, perturbations would return to the same attractor, like the memories in a Hopfield network.

Is the *Bauplan* of an animal a similar attractor, which the myriad of microscopic

mindless actions can't help but keep returning to? In other words, the answer to self-repair is that there's no blueprint or explicit diagram; there's just a framework and a set of parameters that mathematically define a complex attractor. Mindless and far-flung distributed operations can't help themselves; they must gravitate toward it.

Such dynamical systems with complex attractors driven by parameters are well known. One example is the Mandelbrot set, a truly exquisite iteration where the parameters define a window and each pixel computes its own color. Another example is *iterated function systems*, a union of a set of contractive maps that Michael Barnsley proved has a single fractal limit attractor akin to Cantor dust.<sup>24</sup>

Barnsley showed, much analogous to Bellman's proof, that some nonlinear iterative processes, despite having many adjustable parameters, have a single, yet complicated, limit, defined by the interaction of the parameters and rules. Simply put, an IFS fractal attractor is like repeatedly copying an image with a special copying machine that makes multiple shrunken and transformed copies of the input page (see figure 2). All nonblank starting pages, from a speck of dust to a piece of black construction paper, end up converging to the same attractor in the limit.

I came across IFSs while working to understand the relationship between recurrent neural networks and finite-state machines. As the result of trying to learn a language, a recurrent network generated an infinite-state machine with the states located on a fractal attractor.<sup>25</sup> Subsequent research used these structures for memory and hierarchical representations.<sup>26</sup>

The mindless intelligence of self-defining and self-repairing, or autopoietic,<sup>27</sup> biological forms is a big leap from Fontana's chemistries and Barnsley's fractals. Yet I am certain that biological form will one day be scientifically explained as an attractor that changes its parameters over time while it's constantly and mindlessly repaired by distributed processing at a microscopic level.

## Ectomenal assembly

Fetal development, or embryogenesis, is perhaps the perfect place to recognize the profound scale of complex behavior achievable by mindless intelligence.

Herb Simon introduced Tempus and Hora as two different kinds of watchmakers who suffer from interruptions: one uses modular construction; the other works with basic parts.<sup>28</sup> Richard Dawkins introduced the idea of the Blind Watchmaker.<sup>29</sup> Both researchers comfortably anthropomorphized what is a mindless assembly process.

Every assembly factory depends critically on human minds both as labor as well as supervision to monitor, correct, and repair ongoing processes. Yet a developing fertilized egg is also an assembly factory, without any human supervisors or any brain, that produces an exquisite, custom product with 10 billion moving parts in only nine months! Where's the mind inside the fertilized egg? Even Intelligent Design proponents might be hard pressed to defend the existence of an omniscient "Intelligent Factory Foreman" who supervises every embryo developing in the world simultaneously, deciding which creatures live or die.

Other than basic work on pattern formation, related to work by, for example, Alan Turing and Stephen Wolfram, we have a long way to go in understanding the mindless intelligence in a process that could self-assemble into a biological form with billions of parts. My lab is working on replacing the idea of a perfect robotic factory with evolutionary processes that must evolve both form and formation and overcome noise and error in physical assembly.<sup>30</sup> One of the more

interesting threads is the relationship between robotic assembly with errors and noise, and the kinds of tasks that Bellman proved could iteratively converge to optimal.<sup>31</sup> This might provide a self-construction theory involving not a blind watchmaker but a *blind chessmaster* who continuously optimizes assembly processes to maximize its own chances for successful reproduction.

## Ectomenal reproduction

Another great mystery of Nature is complex self-reproduction. Shy of a magical reverse-engineering theory (which would let us genetically engineer flying horses), we have little or no grasp on the algorithmic processes involved in the major transition

Fetal development, or embryogenesis, is perhaps the perfect place to recognize the profound scale of complex behavior achievable by mindless intelligence.

from single cells reproducing, through colonialization, to multicellular creatures with differentiated tissues and functions.

I think it's another case of dramatically underestimating the amount of intelligence in a seemingly obvious natural process. We have many simple examples of reproduction in software, from straight data copying to self-reproducing code as shown by evaluating this ditty in Common Lisp:

```
((LAMBDA (X) (LIST X (LIST 'QUOTE X))))
'(LAMBDA (X) (LIST X (LIST 'QUOTE X))))
```

Following John Von Neumann's challenge of finding self-reproduction in cellular automata, Christopher Langton helped birth the field of artificial life with his more elegant automata,<sup>32</sup> and Jason Lohn and James Reggia showed how easy it is to discover the rules for such automata.<sup>33</sup> Yet so far, all our computing reproducing systems, including Tom Ray's Tierra<sup>34</sup> and Hod Lipson's cubes,<sup>35</sup> are very simple. I'm hopeful that evolutionary

search for more complex reproductive forms holds some hope for understanding how a mindless reproductive process can become more capable over time to sustain complexity in the design of reproducible machinery.

## Ectomenal recognition, control, and regulation

Obviously, intelligence arises outside the mental sphere in so many other places in nature that I can't list them all.

The immune system is an ectomenal chemical recognition system that filters and separates millions of chemicals along the me/not-me boundary, without a central database listing which compounds are in or out.

Self-control of physical movement, of individuals and groups, is often mindless. This isn't only because time constraints push nervous-system controls to the edge but also because it's hard to find a valuable use for cognitive symbols inside mainly numeric models such as pattern generators and feedback loops.

Finally, the zenith of self-regulation is probably the planet itself. Similar to Adam Smith's "invisible hand" idea that markets are mindlessly intelligent regulators and allocators of goods and services, the Gaia hypothesis proposes that the whole biosphere operates so as to maintain the right conditions for life as we know it.<sup>36</sup> A trivial and kooky interpretation is that Gaia is a Goddess with a mind of her own, complete with symbols, logic, and language, so she might talk to us one day through a burning bush or a statue of her likeness. A deeper interpretation is to recognize that the algorithmic complexity of balancing resources, encouraging growth, and managing the network of species to maintain the "sweet spot" for life is a huge job requiring such intelligence that we better not entrust it to any elected human officials!

Under the mindless-intelligence viewpoint, both evolution itself and the global-regulation system known as Gaia are intelligent beyond and outside the mental framework based on the symbol manipulation that AI has chosen as its first 50-year focus.

I'm neither alone nor unique in wishing for a stronger scientific basis for the field. These comments certainly hearken back to many earlier calls.<sup>37</sup> Much of the world has changed in the last decade. For example, after so many years of chasing generative linguistics

tics' focus on parsing and syntax, the main thrust of both natural language processing and speech recognition has been to drive mindless statistical responses from large corpora rather than to establish carefully wrought rules and features. Intelligent-control research is also moving in a mindless way, from robotics that use shaky logical algorithms to more mathematically sophisticated nonlinear control systems.<sup>38</sup> Much cognitive-modeling research takes seriously the idea that algorithms should be not only cognitively plausible but also neurally plausible. Finally, machine learning research has progressed from its early efforts at matching human learning curves, to building strong algorithms for extracting knowledge from large statistical sources.

Yet these fields often must defend themselves from the charge that they aren't really AI. George Dyson recently visited Google and wrote that he has long considered that when "real" AI arrives on the scene, it will be surrounded by "a circle of cheerful, contented, intellectually and physically well-nourished people."<sup>39</sup> Certainly Google is based on a very large database and uses statistical machine learning techniques to choose which keywords are important in different contexts. Does Google software have any of the cognitive aspects that AI has studied for many years? The mindless market doesn't care.

As we've seen, mindless intelligence abounds in Nature, through processes that channel mathematical ideals into physical processes that can appear optimally designed yet that arise through and operate via exquisite iteration.

The hypothesis for how intelligence arises in Nature is that dynamical processes, driven by accumulated data gathered through iterated and often random-seeming processes, can become more intelligent than a smart adult human, yet continue to operate on principles that don't rely on symbols and logical reasoning. The proof lies not only in Markovian situations where a greedy sequential-choice algorithm driven by values converged under Bellman's equation, but also in the reliability, complexity, and low cost of biologically produced machines.

Because our minds aren't what they seem, symbolic explanations of our behavior that were extracted from protocol analysis and conscious introspection are misleading at best and complete fabrications at worst. Most of what our brains are doing involves mindless chemical activity not even distinguishable from

digestion of the food in the Chinese Room.

I don't mean to imply that human cognition isn't worth studying. I just want to reiterate that cognitive reporting is an always-incomplete story, a simplified verbalization of a partial insight of the working patterns of our brains. And brains aren't instruction set computers; they're complicated biological networks with all kinds of feedback at all levels, like metabolisms, gene regulatory networks, and immune systems. The software and systems that emerge from and control these networks, like evolution, embryological-development protocols, Gaian ecological regulation, or mind, will be much harder to reverse-engineer than the artifacts of human engineering culture.

Symbolic Mind is a self-aggrandized fiction told to make sense of a few pounds of mindlessly intelligent meat. It's time we wean ourselves from the fiction and start working on the science.

Emphatically then, as AI arises, it won't be organized like a good computer program, it won't speak English, and it certainly won't act like a humanoid robot from a science fiction movie. Symbolic Mind is a self-aggrandized fiction told to make sense of a few pounds of mindlessly intelligent meat. It's time we wean ourselves from the fiction and start working on the science. ■

### Acknowledgments

Thanks to my PhD students, past and present, for their collaborations, to the editors for their encouragement in finishing this manuscript, and to Carl Feynman for help coining a word.

### References

1. J. Lanier, "One Half of a Manifesto," *Wired*, Dec. 2000; [www.wired.com/wired/archive/8.12/lanier.html](http://www.wired.com/wired/archive/8.12/lanier.html).
2. M.S. Gazzinga, *The Social Brain: Discovering the Networks of the Mind*, Basic Books, 1985.

3. D. Dennett, *Consciousness Explained*, Little, Brown and Co., 1991.
4. M. Gladwell, *Blink: The Power of Thinking without Thinking*, Little, Brown and Co., 2005.
5. J.R. Searle, "Minds, Brains, and Programs," *Behavioral and Brain Sciences*, vol. 3, no. 3, 1980, pp. 417-457.
6. J.F. Kolen and J.B. Pollack, "The Observers' Paradox: Apparent Computational Complexity in Physical Systems," *J. Experimental and Theoretical Artificial Intelligence*, vol. 7, no. 3, 1995, pp. 253-269.
7. H. Juille and J.B. Pollack, "Coevolutionary Learning and the Design of Complex Systems," *Advances in Complex Systems*, vol. 2, no. 4, 1999, pp. 371-393.
8. S.G. Ficici and J.B. Pollack, "Pareto Optimality in Coevolutionary Learning," *Advances in Artificial Life: 6th European Conf. (ECAL 2001)*, LNAI 2159, Springer, 2001, pp. 316-325.
9. A. Bucci and J.B. Pollack, "Focusing versus Intransitivity: Geometrical Aspects of Coevolution," *Proc. 2003 Genetic and Evolutionary Computation Conf., LNCS 2723*, Springer, 2003, pp. 250-261.
10. R.A. Watson and J.B. Pollack, "A Computational Model of Symbiotic Composition in Evolutionary Transitions," *Biosystems*, vol. 69, nos. 2-3, 2002, pp. 187-209.
11. E. Sklar and J.B. Pollack, "A Framework for Enabling an Internet Learning Community," *J. Educational Technology & Society*, vol. 3, no. 3, 2000, pp. 393-408.
12. E.D. De Jong and J.B. Pollack, "Ideal Evaluation from Coevolution," *Evolutionary Computation*, vol. 12, no. 2, 2004, pp. 159-192.
13. A. Bader-Natal and J.B. Pollack, "Motivating Appropriate Challenges in a Reciprocal Tutoring System," *Proc. AI in Education 2005 (AIED 05)*, IOS Press, 2005, pp. 49-56.
14. P. Funes and J.B. Pollack, "Evolutionary Body Building: Adaptive Physical Designs for Robots," *Artificial Life*, vol. 4, no. 4, 1998, pp. 337-357.
15. H. Lipson and J.B. Pollack, "Automatic Design and Manufacture of Robotic Lifeforms," *Nature*, vol. 406, no. 6799, 2000, pp. 974-978.
16. G.S. Hornby and J.B. Pollack, "Creating High-Level Components with Generative Representation for Body-Brain Evolution," *Artificial Life*, vol. 8, no. 3, 2002, pp. 223-246.
17. G. Tesauro, "Temporal Difference Learning of Backgammon Strategy," *Proc. Int'l Conf. Machine Learning (ICML 92)*, Morgan Kaufmann, 1992, pp. 451-457.

The Author



**Jordan B. Pollack** is a professor of computer science and complex systems at Brandeis University and the director of the DEMO (Dynamical & Evolutionary Machine Organization) laboratory. He has broad interests and publications across AI, artificial life, neural and evolutionary computing, robotics, complex and dynamic systems, educational technology, and intellectual-property law. He has been involved in startup companies including Abuzz, Affinnova, and Thinmail. His laboratory's work on automatically designed and manufactured robots made front-page news worldwide in 2000, and he was named one of MIT's *Technology Review* "TR 10" in 2001. He received his PhD in computer science from the University of Illinois. Contact him at the Computer Science Dept., Brandeis Univ., 415 South St., Waltham, MA 02454; pollack@cs.brandeis.edu; www.jordanpollack.com.

18. P.J. Angeline and J.B. Pollack, "Competitive Environments Evolve Better Solutions to Complex Problems," *Proc. 5th Int'l Conf. Genetic Algorithms (ICGA 93)*, Morgan Kaufmann, 1993, pp. 264–270.

19. J.B. Pollack and A. Blair, "Co-Evolution in the Successful Learning of Backgammon Strategy," *Machine Learning*, vol. 32, no. 3, 1998, pp. 225–240.

20. R.E. Bellman, *Dynamic Programming*, Princeton Univ. Press, 1957.

21. J.B. Pollack, "Nannon: A Nano Backgammon for Machine Learning Research," *Proc. 2005 Int'l Conf. Computational Intelligence in Games*, 2005, pp. 277–284; <http://csapps.essex.ac.uk/cig/2005/papers/p1055.pdf>.

22. M. Minsky, "Matter, Mind, and Models," *Proc. IFIP Congress*, 1965, pp. 45–49; <http://groups.csail.mit.edu/medg/people/doyle/gallery/minsky/mmm.html>.

23. W. Fontana and L.W. Buss, "The Arrival of the Fittest": Toward a Theory of Biological

Organization," *Bull. Mathematical Biology*, vol. 56, 1994, pp. 1–64.

24. M.F. Barnsley, *Fractals Everywhere*, Academic Press, 1988.

25. J.B. Pollack, "The Induction of Dynamical Recognizers," *Machine Learning*, vol. 7, nos. 2–3, 1991, pp. 227–252.

26. S. Levy, O. Melnik, and J.B. Pollack, "Infinite RAAM: A Principled Connectionist Basis for Grammatical Competence," *Proc. 22nd Ann. Conf. Cognitive Science Soc.*, Cognitive Science Soc., 2000, pp. 298–303; [www.cogsci.rpi.edu/CSJarchive/Proceedings/2000/COGSCI00.pdf](http://www.cogsci.rpi.edu/CSJarchive/Proceedings/2000/COGSCI00.pdf).

27. F.J. Varela, H.R. Maturana, and R. Uribe, "Autopoiesis: The Organization of Living Systems, Its Characterization and a Model," *BioSystems*, vol. 5, no. 4, 1974, pp. 187–196.

28. H.A. Simon, *The Sciences of the Artificial*, MIT Press, 1969.

29. R. Dawkins, *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe without Design*, W.W. Norton and Co., 1986.

30. J. Rieffel and J.B. Pollack, "Evolving Assembly Plans for Fully Automated Design and Assembly," *Proc. NASA/DoD Conf. Evolvable Hardware (EH 05)*, IEEE Press, 2005, pp. 165–170.

31. S. Viswanathan and J.B. Pollack "On the Robustness Achievable with Stochastic Development Processes," *Proc. NASA/DoD Conf. Evolvable Hardware (EH 05)*, IEEE Press, 2005, pp. 34–39.

32. C.G. Langton, "Self-Reproduction in Cellular Automata," *Physica D*, vol. 10, nos. 1–2, 1984, pp. 135–144.

33. J.D. Lohn and J.A. Reggia, "Automatic Discovery of Self-Replicating Structures in Cellular Automata," *IEEE Trans. Evolutionary Computation*, vol. 1, no. 3, 1997, pp. 165–178.

34. T.S. Ray, "An Approach to the Synthesis of Life," *Artificial Life II*, C. Langton et al., eds., Addison-Wesley, 1991, pp. 371–408.

35. V. Zykov et al., "Self-Reproducing Machines," *Nature*, vol. 435, no. 7038, 2005, pp. 163–164.

36. J.E. Lovelock, *Gaia: A New Look at Life on Earth*, Oxford Univ. Press, 1979.

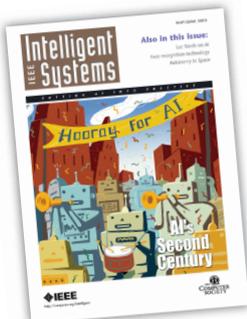
37. R.A. Brooks, "Intelligence without Representation," *Artificial Intelligence*, vol. 47, nos. 1–3, 1991, pp. 139–159.

38. F. Zhao, "Extracting and Representing Qualitative Behaviors of Complex Systems in Phase Spaces," *Artificial Intelligence*, vol. 69, nos. 1–2, 1994, pp. 51–92.

39. G. Dyson, "Turing's Cathedral," *Edge*, 24 Oct. 2005, [www.edge.org/3rd\\_culture/dyson05/dyson05\\_index.html](http://www.edge.org/3rd_culture/dyson05/dyson05_index.html).

# See the Future of Computing Now in IEEE Intelligent Systems

Tomorrow's PCs, handhelds, and Internet will use technology that exploits current research in artificial intelligence. Breakthroughs in areas such as intelligent agents, the Semantic Web, data mining, and natural language processing will revolutionize your work and leisure activities. Read about this research as it happens in **IEEE Intelligent Systems**.



**SUBSCRIBE NOW!** <http://computer.org/intelligent>

