

The Voice of the Turtle: Whatever Happened to AI?

Doug Lenat

■ *On March 27, 2006, I gave a light-hearted and occasionally bittersweet presentation on “Whatever Happened to AI?” at the Stanford Spring Symposium presentation—to a lively audience of active AI researchers and formerly active ones (whose current inaction could be variously ascribed to their having aged, reformed, given up, redefined the problem, and so on). This article is a brief chronicling of that talk, and I entreat the reader to take it in that spirit: a textual snapshot of a discussion with friends and colleagues, rather than a scholarly article. I begin by whining about the Turing test, but only for a thankfully brief bit, and then get down to my top-10 list of factors that have retarded progress in our field, that have delayed the emergence of a true strong AI.*

W e’re now well past 2001; where is HAL? When Marvin Minsky advised Arthur C. Clarke and Stanley Kubrick, 40 years ago, it seemed that achieving a full HAL-like AI by 2001 was every bit as likely as, well, commercial Pan Am flights to the moon by 2001.¹ As Bill Rawley said, the future is just not what it used to be:

When I was growing up in the 1950s we all knew what the 1990s would be like. It would be a time of great prosperity. We would live in big homes in the suburbs. There would be many labor-saving conveniences for the homemaker, and robots would do the hard chores. We would commute to work in our own helicopters. The short workweek would mean lots of leisure time that families (mom, dad and the two kids) would enjoy as “quality time” together. Space travel would be common with people living on other planets. Everyone would be happy living a fulfilling life in a peaceful world.²

Why are we still so far from having those household robots, so far from having a HAL-like AI that could pass the Turing test? The answer to that has several parts; I’ll discuss each one in turn.

For one thing, the test that Alan Turing originally proposed in *Mind* in 1950 has mutated into a much more difficult test than he ever intended, a much higher bar to clear than is needed for, for example, HAL or household-chore-performing robots.

Turing test aside, we do expect and want AIs (general ones or application-specific ones) to be capable of carrying on

back-and-forth clarification dialogues with us, their human users. But for that to be anywhere near as efficient as conversing with another human being—for the computer to not come off like an idiot savant or idiot—requires its knowing a vast panoply of facts (hot coffee is hot; jet planes fly hundreds of miles per hour), rules of thumb (if you fly into a small town airport, you’re probably not connecting there to another ongoing flight), shared experiences (what it feels like when your ears plug up as the plane descends), scripts (buckling a seatbelt), and so on.

And finally, we in AI too often allow our research vector to be deflected into some rarified area of theory or else into some short-term application. In both of those extreme cases, what we end up doing doesn’t have a large projection on what we as a field could be doing, to bring real AI into existence, so it shouldn’t come as much of a surprise that it’s still not here yet.

If you read through the literature—say all the IJCAI and AAAI proceedings and all the *AI Magazine* and *AI Journal* articles—their abstracts report a steady march of significant research strides being made, and their conclusions report great expectations, blue skies, and fair weather ahead. So much so, that when John Seely Brown and I submitted an *AI Journal* manuscript in 1983 about why some of our then-recent machine-learning programs only appeared to work—that is, on closer analysis had more the veneer of intelli-

gence rather than some seed that would germinate, some source of power that would scale up into true AI—the reviewers didn't know what to make of the article, didn't know if it even fit the criteria for publication!³

Neutering Turing's Imitation Game

Some of us treat the Turing test as a holy relic (Guccione and Tamburrini 1988), and end up so close to it—analytically pursuing some sub-sub-...-subpart of it—that we no longer see the real AI goal, a world-altering functionality that Turing's test was merely a first thought of how to test for. Others of us have turned our back on the Turing test and, more generally, on the Big AI Dream, and we build Roombas and knowledge management systems and other AI “raisins.” Part of my message here is that both of these courses are too extreme. We Eloi and Morlocks (Wells 1895) can and should come together, bound by the common goal to build a real AI—not by 2001 or (now) 2010, not as a two-year project, but at least in our lifetime. Okay, yes, in particular in my lifetime.

Turing's test was *not* (as many currently believe it to be) to see if a (human) interrogator could tell more than 50-50 whether they were talking with a human or with a computer program. In his version, “The Imitation Game,” the interrogator is told that the interrogator is talking to a man and a woman—not a person and a computer!—and that the interrogator must decide which is the woman and which is the man.

Let's suppose that men are on average able to fool the interrogator 30 percent of the time into thinking they are the woman. Now we replace the man with a computer, but the interrogator is still told that he or she is talking to a man and a woman, both of whom will claim to be the woman during a brief online conversation, and the interrogator's job is to pick the woman. If we get a computer that the interrogator picked (as being the woman) 30 percent or more of the time, then (Turing concludes) that computer, as programmed, would be intelligent. Well, at least as intelligent as men, anyway.⁴As Judy Genova puts it (Genova 1994), Turing's proposed game involves not a question of species but one of gender. In the process of making the test gender neutral, it has inadvertently been made vastly more difficult for programs to pass. In today's version, the interrogator can draw on an array of facts, experiences, visual and aural and olfactory and tactile capabilities, and so on, to ask things he or she never would have asked under Turing's original test, when the interrogator thought that he or she was trying to distinguish a human man from a human woman through a teletype.

Even worse than this, and more subtle and more interesting, is the fact that there are dozens of what I call translogical behaviors that humans regularly exhibit: illogical but predictable decisions that most people make: incorrect but predictable answers to queries. As a sort of hobby, I've collected articles distinguishing over three dozen separate phenomena of this sort.

Some of these are very obvious and heavy-handed, hence uninteresting, but still work a surprising fraction of the time—“work” meaning, here, to enable the interrogator to instantly unmask many of the programs entered into a Turing test competition as programs and not human beings: slow and errorful typing; 7 ± 2 short-term memory size; forgetting (for example, what day of the week was April 7, 1996? What day of the week was yesterday?); wrong answers to math problems (some wrong answers being more “human” than others: $93 - 25 = 78$ is more understandable than if the program pretends to get a wrong answer of 0 or -9998 for that subtraction problem (Brown and van Lehn 1980).

At the 2006 AAAI Spring Symposium, I went through a dozen more sophisticated translogical phenomena (Gleim 2001, Chapman and Malik 1995, Tversky and Kahneman 1983). While these are not yet the stumbling blocks—our programs are still stumbling over cruder ones—they may become ways of distinguishing humans from computers once those heavier-handed differences have been taken care of. For example, asked to decide which is more likely, “Fred S. just got lung cancer” or “Fred S. smokes and just got lung cancer,” most people say the latter. People worry more about dying in a hijacked flight than the drive to the airport. They see the face on Mars. They hold onto a losing stock too long because of ego. If a choice is presented in terms of rewards, they opt for a different alternative than if it's presented in terms of risks. They are swayed by ads. European countries that ask license applicants to “check this box to opt in” have a 10 percent organ donor enrollment; those whose form says “check to opt out” have 90 percent.

The basic problem is that (1) most programs are engineered to make choices as rationally as possible, whereas the early hominids were prerational decision makers, for the most part, and that (2) unfortunately, we are the early hominids. This makes the task much more daunting than Turing's original imitation game. One requirement of the gender-neutral Turing test is that the computer know more or less the totality of what humans of a certain age group, culture, time period, and so on would know. That takes us to our second point: the sweeping panoply of knowledge required for “real” natural language understanding (NLU).

Holding Up One's End of a Conversation

In the ancient cave-wall drawings by Winograd, Schank, and others, anthropologists have found numerous well-preserved examples where resolving ambiguous words, anaphoric reference, analogy, prepositional phrase attachment, scope of quantifiers, and so on, require some bit of real-world knowledge and/or reasoning capability. For example:

Which meaning of “pen” is intended in “The pen is in the box” versus “The box is in the pen”?

What is the referent of “the fool” in each of these sentences: “Joe saw his brother skiing on TV last night, but the fool didn’t recognize him” and “Joe saw his brother skiing on TV last night, but the fool didn’t have a coat on.”

Since so many people have siblings, “Mary and Sue are sisters” probably means they are each other’s sister, but if I say “Mary and Sue are mothers” it doesn’t cross your mind even for an instant that I mean that they are each other’s mother. Well, at least it doesn’t cross your conscious mind.

People who never heard of nested quantifiers can easily and correctly parse both “Every American has a mother” and “Every American has a president.”

I won’t belabor this point here; it is the foundation for the 24-year “detour” we took to do Cyc. That is, the hypothesis that progress in natural language, in speech (especially analyzing and synthesizing prosody), in robotics ... in practically all AI applications, and, more generally, almost all software applications, will be very limited unless and until there is a codification of a broad, detailed ontology and a codification of assertions involving the terms in the ontology, respectively, a skeleton and the flesh on that skeleton. What, for the examples above, it is important to know—the relative sizes of writing pens, corrals, penitentiaries, various sorts of boxes, and so on. Causes precede effects. Who can see who, when x is watching y on television. Skiing is a cold-weather outdoor activity. And so on.⁵ Today most people look to the semantic web as evolving, any year now, into that codification, that interlingua. Not believing that any such grassroots effort would converge at any level deeper than text (see Wikipedia) without some serious priming of the knowledge pump ahead of time, manually, by thoughtful design, this is why we launched the Cyc project back in 1984 and have labored at it continuously since then.

Succumbing to Temptation

Our third point was that almost everyone else in the field has given in to some “lure” or other, and is following a path at most tangentially related to

building a real HAL-like AI. At the 2006 Spring Symposium, this was the most candid portion of my talk, and the most controversial. I had quite a bit of additional material prepared, but I never got to it, because the audience began to, shall we say, discuss this energetically with me and with each other.

I started this “Succumbing to Temptation” section as a sort of “Top 10 Reasons Why Real AI Is Not Here Yet” list, although there are some extra items on it now; I guess that’s still 10 in some base. Let’s count down this list:

12. The Media and the Arts

The media and the arts aren’t always the bad guys, but they continually bombard the public with incorrect information about what computers can and can’t do, what’s hard and what’s not, and so on. On *Star Trek*, computers 200 years from now can perfectly understand speech but not do a good job synthesizing it (the same asymmetry is found in Neal Stephenson’s *The Diamond Age*); Commander Data could parse contractions but never employ them. People think that HAL-like machines already exist by now (it is after 2001 after all) and are concerned about the well-known inherent tendencies for almost all intelligent machines to kill either (1) their creators, (2) innocent bystanders, or (3) everyone.⁶ About every 20 years since world war II, premature announcements sweep the media, announcements of about-to-be-sold “home robots” that will cook the meals, mind the baby, mow the lawn, and so on—for example, Nolan Bushnell’s 1983 Androbot “product” BOB (Brains On Board).⁷ There are frankly so many examples of media hyperbole that to mention any few of them is to understate the problem. We in AI are sometimes guilty of feeding the media’s hunger for extravagant prediction and periodically pay the price of this by being the subject of media stories (often equally hyperbolic) about how project x or approach y or field z is failing, precipitating periodic AI “winters,” after which the climate thaws and the whole dysfunction cycle starts anew.

Given that sad state of affairs, how has it impeded progress in AI? Primarily in the form of suboptimal AI research and development funding decisions. Such decisions are sometimes made by individuals who believe some of the overly positive or overly negative misinformation, or individuals who do indeed “know better” but still must answer to the public—possibly answering indirectly, such as funding agencies who answer to Congress (Hulse 2003), whose members in turn answer to the public, whose state of (mis)information is shaped by arts and media.

A charming example of media miscoverage was mentioned to me the day I gave this talk (March



Courtesy M. Veloso, CMU.

Figure 1. Robocup Soccer Goalie Doing Victory Dance.

27, 2006) by Manuela Veloso of Carnegie Mellon University. Getting a team of robot dogs to play soccer is an intensely difficult, very real-time scene-understanding and interagent-coordination problem that she has spent years tackling. But what does the media inevitably show (CBS News 2005)? The little victory dance it took 20 minutes to program (the dance that the goalie robot dog does when it blocks a shot, getting up on its hind legs and raising its paws into the air; see figure 1)!

11. Smurifitis

The Smurfs is a cartoon series in which a race of tiny blue elves employ the word *smurf* here and there, to stand in for almost any word—most often a verb—mostly for cuteness. For example, “Papa, smurf me up some supper!” In AI, we often become just as enchanted with some word we’ve made up, and we keep reusing that word to mean different things. Some examples: frames, actors, scripts, slots, facets, rules, agents, explanation, semantics (this one is particularly ironic), ontology (this would be the winner if not for semantics), and AI (okay, perhaps this is actually the winner). In some cases, we commit the converse mistake: inventing new words for old concepts. And even worse, both of these errors intermix: for example, some meanings of *actors* are equivalent to some meanings of *agents*.

Most fields, scientific and otherwise, quickly learned the value of having the discipline of a more or less stable sort of terminological base. In medicine, for example, there are still discussions

and news articles appearing about the September 2000 change of the definition of the term *myocardial infarction*—a decision deliberated on by and announced by a joint committee of the American College of Cardiology and the European Heart Society. Without this sort of well-founded ontological anchoring, AI research is adrift: researchers are misled when (1) the same words mean different things (in different articles), and (2) different words mean the same (or almost the same) thing.

10. Cognitive Science

Sometimes it’s useful to have a concrete example of a small, portable, physically embodied intelligence—say, for example, me. *Homo sapiens* does indeed stand as a compelling existence proof that machine intelligence should be possible. Unless you ascribe some sort of magical, religious, or quantum-dynamical “soul” to people (and, perhaps even more strangely, you deny such a thing to computers).

Sometimes this analogy is useful: we can be inspired by biological or cognitive models and design data structures and algorithms that are more powerful than the ones we would have come up with on our own. But at other times, this analogy is harmful to progress in AI. We take human limitations too seriously, even when they are the result of hardware limitations and evolutionary kludges (as Marvin Minsky is fond of pointing out). For example, having a short-term memory size of 7 should not be a design goal for our software; only being able to “burn” one fact into long-term memory every several seconds should not be a design goal.

The various translogical phenomena that humans more or less universally fall prey to, described above, mostly fall into this category.

Consider this parable: Aliens come to Earth and observe a computer (the PC sitting on my desk, say) as it solves a set of differential equations. They want to build something that can solve the same problems (for irony, let’s suppose their alien technology is mostly biological—tailoring organisms to produce specialized new organisms that can be used as tools, weapons, vessels, calculators, whatever). They carefully measure the sounds coming out of the computer, the heat output, the strength of the air flow from its vent, the specific set of buttons on the monitor and the CPU, the fact that the computer seems to be composed of those two parts with a flexible cable interconnecting them, and so on, and they struggle to have all these properties duplicated by their new organic creations. We know it’s silly; those properties are merely artifactual. In much the same way, it is quite likely that many of the properties of the human brain and the human information-processing system are just as artifactual, just as beside the point.

And yet so many hours, so many careers, of brilliant researchers like Allen Newell and Herb Simon have been sacrificed to the altar of such verisimilitude. Generations of Carnegie Mellon University graduate students built Pure Production Systems, with the goal of duplicating the timing delays, error rates, and so on, of humans (for example, see Laird and Rosenbloom [1994]). Today we look back with disdain at ancients who pursued similar sorts of sympathetic magic rituals, but are we really that much different? I fear that some of our practices in faithfully duplicating human cognitive architecture fall under what Richard Feynman dubbed *cargo cult science*, with the same faith—part reverence and part hope—as the deluded South Seas islanders circa 1946 with their bamboo headphones, meticulously duplicating the form but totally misunderstanding the functionality, and the salience.

There is a more serious error that many researchers make, succumbing to the lure of analogizing too closely to human development. To summarize the reasoning behind that paradigm: suppose we very carefully and faithfully model a neonate; then that AI “baby” should be able to learn and grow up just as a human baby does. The Blue Fairy doesn’t need to intercede to turn it into a real boy, it will just evolve that way, inevitably, à la the robots in the movies such as *AI* and *I, Robot*.

I’m not questioning here whether machines can learn—I know firsthand that they can learn; see AM and Eurisko (Lenat and Brown 1984). The question is more one of placing a bet: how much priming of the pump should go on, versus letting the computer learn everything on its own? My point here is that I fear that too many brilliant AI researchers have placed that bet too far toward the tabula rasa end of that spectrum, and as a result AI is not as far along as it could have been by now. I’ve been working at the opposite end—priming the pump in the form of building Cyc—for 24 plus years now, and after 900 person-years of such effort, we finally, in the last couple years, have been ready to use it as an inductive bias for learning. This has led to a dramatic shift in activity at Cycorp: most of our effort these days goes into utilizing Cyc for machine learning by discovery or by language understanding and clarification dialogue, not hand-coding assertions one by one.

I’ll mention one more error of this analogy-to-humans sort that has slowed down progress in AI: an almost mystical worship of physical embodiment. Putting our programs inside a robot with a tangible, mobile body and seeing it interact with the real world is great fun. It’s rewarding, like watching your baby take its first step. There’s nothing wrong with this “situatedness,” exactly—and it will be required, for instance, when it’s time to have real home robots that cook, clean, mind the baby, mow the lawn, and so on—but it is yet one

more potential distraction until we get close to that goal.

When reporters visit Cycorp, they frequently ask to see Cyc’s metallic “body” and are not satisfied with the answer that it’s just software. They are sure we are hiding Cyc’s robot body somewhere. As we take them on a tour of the building, they often surreptitiously snap photos of our air-conditioning unit—it’s the only thing in the building that’s very big, metallic, and has several blinking lights and levers on it.

In the 1950s, the term *hacker* used to refer to model-railroading enthusiasts, forever tinkering with their track layouts and engines and such; there is a bit of that sort of “hacking” going on with robotics still today. Great fun, but like model railroading, perhaps best left for after work. Today there is less of an AI brain drain into this sort of Embodiment fetishism than there was 10–20 years ago,⁸ but integrated out, that is still a serious loss of forward progress due to this one attractive nuisance.

9. Physics Envy

When I see physics students walking around with T-shirts that have Maxwell’s Equations on them, or TOE unified equations on them, I just get so jealous. Where the hell are our “Maxwell’s Equations of Thought”? Computer scientists in general are guilty of this sort of envy, not just AI scientists. If Jerry Seinfeld ever gives a keynote talk for one of our conferences, he could lead off with something like: “Have you ever noticed how fields that are not quite self-confident about the size of their own scientific credibility often overcompensate by adding the word *science* into their name: food science, military science, political science, computer science.” All kidding aside, there is a deeply ingrained belief in almost all AI researchers that ultimately there will be some sort of T-shirt with the AI “laws” writ in large font on it. In recent years this faith in short answers has been blunted a bit,⁹ as computing power has made most mathematics proofs too long (too many cases) for humans to really do, or even read, by hand, as the Human Genome Project and Wikipedia and the ESP game (van Ahn and Dabish 2004) and other massive collaborative efforts show the power of a vast amount of elbow grease applied to some problems. But 50 years of “looking under the lamppost” is the price that our field has paid—that is, exploring the space of formal models that meet various academic criteria because that’s where the light is. That’s what got published, earned tenure, won awards. I’ll return to this later on, but it’s one reason so many of us who have worked on Cyc are either so old, like me, that they don’t care about basking in that light any more or so young that they don’t yet need it.

8. Jerks in Funding

The problem has not been the quality of the individuals making these decisions¹⁰ but rather the inevitable up-and-down nature of the funding streams themselves. As a former physics student, I remember that the first, second, and third derivatives of position with respect to time are termed velocity, acceleration, and jerk.¹¹ Not only are there changes in funding of AI research and development (velocity), and changes in the rate of change of funding (acceleration), there are often big changes in the rate of those accelerations year by year—hence what I meant by jerks in funding.

One of the unfortunate trends of the past 15 years is the gradual downward trend in many funding streams (especially U.S. ones): smaller awards, shorter awards, ones where the program itself ends earlier than it was originally slated to, and ones contingent on winning head-to-head “bake-offs.” In general it’s good to have performance milestones, but coupled with the other changes, this translates into a larger and larger fraction of an investigator’s time going to keeping the funding balls in the air. Many of the funding sources have very targeted programs, which means that if one’s goal is to build a general AI, one must tack like a sailboat does with (or against) the funding winds. That is, only a fraction of the wind power—the funding dollars—for such a project is harnessed to move the research forward toward AI; the other component moves it orthogonally to that direction, orthogonally to progress toward Big AI. Year after year.

7. Academic Inbreeding

Numerous AI paradigms formed—and persist—with concomitant impact on everything from funding decisions, to hiring and promotion decisions, to peer review. Each paradigm collectively and perhaps unconsciously decides what problems are valid to work on, what data sets to use, what methods are acceptable, and so on. These cliques often run their own workshops and conferences, hire each other’s students, accept each other’s paper submissions, and conversely don’t do those things with outsiders, with those not in that clique. I could cite very concrete examples—okay, yes some we are guilty of with Cyc, too—but my point is that this is happening throughout AI and to some extent throughout academia. Thomas Kuhn (1962) and many others have bemoaned this—it’s not a new phenomenon, but it remains an albatross around our necks. Let me give one AI-specific example.

In 1980, Bob Balzer and a few dozen of us came together in a sort of “caucus” to form AAAI, to combat the trend we were seeing at IJCAI conferences: the trend where there would be so many tracks that you never ended up hearing any pre-

sentations outside your subarea. One of our defining constraints was that there would only be one track at an AAAI conference, ever, so everyone could go to every talk and hear about the latest, best work in every subfield of AI, in every paradigm of every subfield. So what happened? Track creep. Today there are multiple tracks at almost all the large conferences, and there are a plethora of small, specialized workshops and conferences, and we as a field are back at the same degree of isolation (among the subfields of AI) or worse than we were in 1980.

Academe contributes another retarding factor on achieving real AI, in the form of publication in journals. This is a slow, “old-boy” process, but even worse it often enables researchers to produce one or two working examples, document them, and then wave their hands about how this will generalize, scale up, and so on. As I remarked earlier, this leads to the strange situation where there are 10,000 peer-reviewed AI articles that talk about how one or another technique is working and great progress has been made, and yet year by year, decade by decade, we still aren’t much closer to producing a real AI.

6. The Need for Individuation

Abraham Maslow might predict that after satisfying our survival and safety needs (like point 8) and our belongingness needs (like point 7), we would begin to work on satisfying our self-esteem needs—the need for individuation, self-respect, and respect of our peers—and then finally we would turn to work on satisfying our self-actualization need, the drive to bring real AI into existence. This point, 6, is all about how frequently we get stuck on the self-esteem needs and don’t get to the self-actualization need.

As an undergraduate or graduate student (at least here in America), we each want and need to distinguish ourselves from our classmates. As a Ph.D. student, we want our thesis to be clearly distinguished not only from other theses but from any work going on more broadly by our advisor or at our lab generally. As an assistant professor, we want to publish enough, and have those articles be clearly enough differentiated from everyone else’s work, that we will get professional respect and tenure. As a tenured professor or career researcher within a company or agency, we want to publish and run research teams and laboratories that have their own unique character and focus, to enhance our reputation, to garner more peer respect, and to enable or facilitate continued funding.

The problem is that all these self-esteem urges make it difficult to run large, coordinated AI projects: the students won’t be there too long and don’t want to become a cog in a big machine. The assistant professors can’t afford to work on projects

where they'll publish just one article every decade (even if it's a seminal Gödelesque one) or else they'll be out of a job, so they had better work on a project that will conclude or have a phase that concludes and can be written up in months, maybe one year tops. And so on up the line.

That was okay when the world was young, in the 1960s and 1970s, when AI projects could pluck low-hanging fruit (Feigenbaum and Feldman 1963, Minsky 1968). Most AI doctoral theses took a year or so of work and included the full source code as an appendix—often just a few pages long. Those small projects were exactly the right ones to do, at the time. Since then, the projects driving progress in AI have grown in size, so that they no longer fit the “isolated Ph.D. student working for a year on their thesis” model, or the “professor and two students working for nine months” model. The models still persist, rather, but the optimality of the projects that fit those models has declined. Take this test: imagine you could work on any path to AI you wanted, you could have any size staff, any duration before you had to deliver something, what would you choose? This is not just a case of bigger is better, this is a case of us no longer working on the right problems, any more, to bring about AI as expeditiously as possible. At least not real AI, in our lifetime. In my lifetime.

5. Automated Evolution

And, more generally, the problem of Scaling Down Too Much. There are several variations of “throwing the baby out with the bathwater,” scaling down so much that one may still appear to be—and believe that one is—solving problem *X* but actually is now solving problem *Y*. We'll return to this theme again in point 4 later on.

One of the approaches that some researchers have pursued, to get to an AI, is to let it emerge through a generate-and-test process akin to biological evolution: the generator is random mutation (of code, or of a feature vector that represents the organism's genome), and the test is natural selection, competition with other mutant organisms for resources in a simulated world. The mutation process itself can be guided or constrained in various ways; for example, the only offspring allowed might be those that get half of their feature vector from one existing individual and the other half from a second individual, thereby emulating sexual reproduction. Or the mutation process can involve duplicating feature elements (akin to gene duplication) and then mutating the new copy.

One of the earliest serious approaches of this sort was done in the mid-1960s by Larry Fogel, Alvin Owens, and Mike Walsh.¹² An even earlier attempt to get programs to evolve to meet an input/output (I/O) “spec”¹³ was done at IBM in the

1950s by Robert Friedberg, in which 64-step machine-language programs were mutated until the resulting program met the I/O spec. The most recent instances of this paradigm include of course A-Life, genetic algorithms (for example, Holland [1975]), simulated annealing, and all their descendants.

One problem that many of these techniques suffer from is that they tend to get stuck on local maxima. In the case of evolving a program to tell whether a number is prime or not, the program synthesis system that Cordell Green and I did in 1973 kept getting stuck on the local maximum of storing small prime numbers and then saying No for every other input. Friedberg found the same annoying stumbling block: for many I/O specs, it was faster to randomize the entire machine language program until a solution was found than it was to try to hill-climb to a solution by making one small change after another!¹⁴

So why have I listed this as an impeding factor? These programs show very promising behavior on small examples, and this leads one to believe that there will be some way for them to scale up. By a “small example” here I mean one in which the search space to evolve through is relatively small and highly parameterized even though the scope of the real-world problem might be enormous—for example, combating a global pandemic. One of the problems with many of these algorithms is that their computational complexity goes up exponentially with the length of the program, feature vector, and so on to be modified or learned. More generally, many AI researchers pursue techniques—whether they be in parsing, robotics, machine learning—that work on small problems, hence lead to great publishable papers, but just won't scale up.

As a graduate student, I saw firsthand several robotics “solutions” of this sort: Is the robot's hand too shaky to put the nut on the bolt? File the bolts to a point, so that level of shakiness or visual inaccuracy can be overcome. Is the contrast not high enough for the computer vision system to distinguish part *A* from part *B*? Paint them each a different bright color. Is parallax not working well enough? Draw a grid of known square sizes and reference marks on each part and on the table on which they sit. Are the two robotic arms colliding too much? Turn one off. Yes, these are extreme examples, but they are also real examples, from one single year, from one single famous lab. This sort of scaling down and corner cutting happens very often in AI, with the researchers (and their students, reviewers, and so on) often not even realizing that they have transformed the problem into one that superficially resembles the original task but is in fact vastly simpler to solve. That brings us to point 4.

4. More Overscaling-Down: Automated Learning

One example of that unconscious transformation of a hard problem into a simple one that appears just as hard is the work done in machine learning, and to be as harsh on myself as anyone else, I'll pick on the automated discovery work done in the 1970s and early 1980s. I had good company falling into this trap—Tom Mitchell, Pat Langley (1977), Herb Simon, and others. In Lenat and Brown (1984), which I mentioned before, we analyzed the sources of power that our programs were tapping into. To the extent that such programs appear to discover things, they are largely discharging potential energy that was stored in them (often unconsciously) by their creator. Potential energy in the form of a set of perfect data. Potential energy in the form of a set of prechosen variables to try to find a “law” interrelating. Potential energy in the form of knowing before it even begins that the program will succeed to find a “law” if it searches for a low-order polynomial as the answer, as the form that the law would take (a linear or quadratic or cubic equation connecting that set of perfect data about a perfectly selected set of variables). This changes the search space from breathtakingly large and impressive to, well, a bit small and shabby. It changes the nature of what these learning programs did from recreating a monumental conceptual leap of science to something more like running a cubic spline curve-fitting algorithm. As Michael Gorman (1987) says, this species of program “merely searches for quantitative relationships between arbitrary terms supplied by the programmers, who interpret its results to be ‘discoveries’ of Kepler’s third law. One feels they have missed the point of Kepler’s genius entirely.”

Arthur Koestler and many others have observed that learning occurs just at the fringe of what one already knows. One learns a new thing, which is an extension, combination, refinement of two or three things that one already knows. To put this in a positive way: the more one knows, the more one can learn. This has certainly been the case at a societal level, in terms of technological progress. But the negative way to think of this is: If you (a program) know more or less nothing, then it’s going to be a slow and painful process to learn almost anything. To the extent that our programs (AM, BACON, Eurisko, and others) knew—and AI programs today still know—next to nothing, the most we could do is to get them to learn or appear to learn very simple things indeed.

The reason that this is on the top-12 list is that there is a kind of chicken-and-egg problem going on here. On the one hand, we want to produce an AI as automatically as possible; we want it to discover most of what it knows on its own, not be hand-coded. On the other hand, since learning

occurs at the fringe of what the learner already knows, if we start the learner out bereft of knowledge then it isn’t going to get to human-level knowledge and reasoning capabilities very quickly, say before the recollapse or entropic heat death of the universe, which is a special case of not getting there in my lifetime.

3. Natural Language Understanding

The problem of understanding natural language (NL) is another chicken-and-egg one. On the one hand, our real AI must be able to surf the web, read books, carry on conversations, tutor and be tutable, and so on—things that human beings primarily do through natural language. On the other hand, as we saw above (with numerous examples involving resolving polysemy, anaphora, ambiguous prepositional phrase attachment, noun-noun phrases,¹⁵ interpreting metaphor and analogy, recognizing irony and sarcasm, and so on) understanding the next sentence one reads or hears might require almost any piece (out of many many millions) of prior presumed knowledge.¹⁶

One of the reasons this impediment is on my list is that AI researchers—NL researchers—have watered down the problem to the point where they “succeed” on every project, and yet natural language understanding of the sort we need is not here, nor is it, frankly, just around the corner yet. They talk about parsing an English sentence into logical form, for example, but then turn out to mean transforming it into something that structurally, syntactically fits the definition of “in logical form.” As a caricature: add parentheses at the very start and end of the sentence, and declare success. So what do I¹⁷ mean for a set of English sentences to be semantically represented in logical form, then? I mean that a theorem prover could produce the same set of answers (to queries whose answers were entailed by those sentences) that a person would, who has just read those sentences. Or, if placed in a situation where some of that knowledge could form (part of) a solution to a problem, a program could deduce or induce that relevance and correctly apply that knowledge to that (sub)problem.¹⁸

The reason this is so near the number-1 spot on my top-12 list is that the potential gain from solving this problem, from dragging this mattress out of the road, is enormous. Given the Internet and the ubiquity of cellular telephones and the propensity for people to play games such as ESP (von Ahn and Dabbish 2004), even by today’s technology standards there is a huge amount of static text and dialogue available. The long tail of the distribution is here—the distribution of what an intelligent program should either already know or, like us, should be able to quickly find (and understand) on the web. Peter Norvig and others have been point-

ing this out for years; let's start listening to them.

2. We Can Live without It

This is in many ways the most insidious roadblock to AI. Let me give a few examples.

When computer games like Adventure and Zork and the Infocom games were developing, they were a strong driving force for better and better natural language understanding, a steady wind. But then, game designers found that players would settle for clicking on a small menu of possible dialogue responses. So that motive force—fiscal and otherwise—atrophied. Game designers also found that players could be seduced into playing games where everyone around them had been killed, turned into zombies, mysteriously vanished, or for some other reason there was no one to talk with in-game. It seems most video games today are of this “everything is loot or prey or else has one small menu of things to say” sort. Or they are massively multiplayer online role-playing games (MMORPGs), where most of the characters one encounters are avatars of other living, breathing human beings—artificial intelligence, like the Mechanical Turk (both the 18th-century “device” and the recently launched Amazon.com microjob site).

When a person goes to the web with some question, that person would really like to be able to just type it in and get back an answer. Period. But by now the public has been conditioned to put up with simple keyword searching (most people don't even use Boolean operators in their searches). Some sites such as ask.com have recognized the need to provide answers, not just hits, but even they are still mostly just serving up keyword matches. For example, if you ask “Is the Eiffel Tower taller than the Space Needle?” you get a set of pages that are relevant to the query, but you must then go and read enough of them to figure out the respective heights and do the arithmetic yourself to answer the question.

Microsoft Word 2007 has lots of great new features, but when I misspell a word like *miss* as *moss*, or *like* as *lake*, it doesn't catch that error. Later on, I will talk about the history of the Boston Computer Museum. Actually I don't do that, but Word doesn't call my attention to this omission. Why doesn't Microsoft improve its spelling-checking, its grammar- and style-checking, and add some simple content-checking (for blatant omissions and inconsistencies)? Because it doesn't have to: We all buy and use Word, and we use it the way it is; we can and do live without it even partially understanding what we're writing in it, just as we all can and do live with keyword-matching search engines. Just as we can and do live with “loot or prey or small menu” game worlds.

The decades-away dreams of the future that

12. The Media and the Arts
11. Smurifitis
10. Cognitive Science
9. Physics Envy
8. Jerks in Funding
7. Academic Inbreeding
6. The Need for Individuation
5. Automated Evolution
4. More Overscaling-Down: Automated Learning
3. Natural Language Understanding
2. We Can Live without It
1. There Is One Piece Missing

Top 12 Reasons Why Real AI Is Not Here Yet.

Arthur C. Clarke sculpted and Bill Rowley reminisced about are still decades away, and to first order I lay the blame right here, with impediment 2—that is, that we all just put up with it, we put up with far less than cutting-edge machine intelligence in our applications.

1. There Is One Piece Missing

For decades, I've concluded this talk with a description of how all the various subfields of AI keep hitting the same brick wall: no general ontology of concepts, no general corpus of rules, assertions, axioms, and so on about those concepts (and interrelating or defining or constraining them, attached to that ontological skeleton). The main change today compared to the 1984 version of this rant is that these days I more rarely have to explain what an ontology is. I usually conclude by talking about how we're putting our money and our professional lives where our mouth is: leaving academe, buckling down to “prime the pump” manually, having a small army of logicians handcrafting Cyc down in Texas.

In a way, this might sound like the counterargument to points 4 and 5; namely, here we are saying that the world does have a lot of special cases in it, worth treating specially—axiomatizing or writing special-purpose reasoning modules for. Engineering is the key: the resolution of the apparent contradiction between point 1 and points 4 and 5 is that building an AI needs to be thought of as a large-scale engineering project. As Shapiro and Goker put it (in a private communication to me), the lesson is that researchers should build systems, and design approaches, that merge theory with empirical data, that merge science with large-scale engineering, that merge general methods with an

You Recommended *What?*

John Riedl

Some years ago Net Perceptions was installing its recommender software in a major catalog retailer. Our top lead consultant had been working with them for months to integrate the software into their back-end databases and front-end call center software (no mean feat: these were Windows PCs simulating IBM “green screens”!), and tuning the recommender to produce high quality recommendations that were successful against historical sales data. The recommendations were designed to be delivered in real time to the call center agents during live inbound calls. For instance, if the customer ordered the pink housecoat, the recommender might suggest the fuzzy pink slippers to go with it, based on prior sales experience.

The company was ready for a big test: our lead consultant was standing behind one of the call center agents, watching her receive calls. Then the moment came: the IT folk at the company pushed the metaphoric big red button and switched her over to the automated recommender system. The first call came in: a woman ordered the complete two-week diet package: a self-contained package with all the powdered food you needed to eat and drink to lose 10 pounds in just two weeks. Our consultant watched nervously as the agent entered the order. In seconds, the recommendations came back. The top of the list was a 5-pound tinned ham! The agent’s eyes were wide as she looked at the screen. She looked back at our consultant with eyebrows raised: should I? He smiled nervously, and gave her his best “I have no idea” shrug.

After a couple of seconds of silence, her training kicked in: “Would you like a tinned ham with that?” Without missing a beat, the woman on the other end replied, “Sure, that would be great!”

We eventually won the contract, probably in spite of recommending hams to customers who wanted to go on a diet.

John Riedl is a professor in the Department of Computer Science at the University of Minnesota.

immense battery of special-case facts, rules of thumb, and reasoners.

While Cyc is still far from complete, the good news is that it is complete enough that it is worth leveraging, and we have made it available at no cost for research and development purposes as ResearchCyc (courtesy of DARPA IPTO), and we have made its underlying ontology available at no cost even for commercial purposes as OpenCyc. If you’ve read this far, you probably already know what Cyc is, but if not, take a look at cyc.com. I’m not going to talk about it more here. We’re far enough along that we have a much smaller fraction of logicians on our staff who are manually extending Cyc and a much larger number of researchers and developers on our staff who are applying it and automatically extending it, using it for natural language understanding and generation, and using its existing ontology and content as an inductive bias in learning.

In one of our current applications, Cyc is used to help parse queries by medical researchers and to

answer those queries by integrating information that originally was in separate databases (whose schemas have been mapped to Cyc, an inherently linear rather than quadratic process). We had hoped that we could show that the number of general preapplication Cyc assertions used in doing a typical user query was at least 4 or 5; even we were stunned to find that that number averaged 1,800.

Conclusions

Building an AI of necessity means developing several components of cognition, and it’s only natural for us as a field to pursue these in parallel. But specialized AI workshops test the individual strands, not so much whether we are making progress towards AI—in effect they are one way we avoid having to take that higher-order test. Telling oneself that AI is too hard, and that the first order results and application “raisins” are what matter, is another way of avoiding taking the tests. I was heartbroken to see the whole field retreat in about

1993 (Crevier 1993) from the “strong AI” goal to weak echoes of it, to what Patrick Hayes and Pat Winston and Esther Dyson and others call the “raisin bread” notion of AI: instead of aiming high at human-level AI, aim at something we can hit, like embedding little tiny flickers of intelligence here and there in applications like Roombas, where the AI is analogous to the raisins in raisin bread.

Don’t settle for that! We have enough tools today—yes, they can be improved, but stop tinkering with them and start really using them. Let’s build pieces that fit together to form an AI (Lenat and Feigenbaum 1991) (not write articles about components that one day in the far future could in theory fit together to form an AI). A real AI could amplify humans’ intelligence, and thereby humanity’s intelligence, in much the same way—and to the same degree if not more—than electric power (and the ensuing stream of appliances using it) amplified human muscle power one century ago. We all have access to virtually limitless, fast machine cycles, virtually limitless Internet (and online volunteers), and ResearchCyc. Use those tools, recapture the dream that brought you to AI, the dream that brought me to AI, and let’s make it happen.

On the one hand, my message is: Be persistent. Twenty-four years of persistence with Cyc has brought it to the point where it is just starting to be harnessed to do somewhat intelligent things reliably and broadly. A similar positive ant vs. grasshopper lesson could be drawn from decades of work on handwriting recognition, on recommender systems, and other blossoming components of AI. The other part of the message is to take advantage of what’s going on around you—the web, semantic web, Wikipedia, SNA, OpenCyc and ResearchCyc, knowledge acquisition through games (Matuszek 2005), and so on and so on and so on. These are tools; ignoring them, staying in your safe paradigm, tinkering with yet one more tool, will delay us getting to AI. In my lifetime. Use them, and keep your eye out for the next tools, so you can make use of them as soon as possible.

Is what you’re working on going to lead to the first real AI on this planet? If not, try to recapture the dream that got you into this field in the first place, figure out what you can do to help realize that dream, and do it.

Yes, I know, this has turned out to be yet one more AI article with a hopeful, positive conclusion; don’t say I didn’t warn you.

Acknowledgements

I thank Ed Feigenbaum, Bruce Buchanan, Alan Kay, Marvin Minsky, John McCarthy, Allen Newell, Herb Simon, and George Polya for shaping my thinking about AI and instilling in me the tenacity to keep pursuing one’s hypotheses until and unless Nature—empirical results—tells you that (and how) you’re wrong. I thank Mehmet Goker, Mary Shepherd, Dan Shapiro, and Michael Witbrock for many wise and useful comments on this article, but they should not be held—what am I saying, go ahead and give them a hard time—for anything you don’t like in it.

Notes

1. As a gimmick, Pan American World Airlines (PanAm) and Trans World Airlines (TWA) started taking reservations for those 2001 spaceflights back in 1968, shortly after the *Apollo 11* landing; as an American teenager growing up in the midst of the Space Race, I signed up right away.
2. From “The Future Is Not What It Used to Be” by Admiral William Rawley, 1995 (www.maxwell.af.mil/au/awc/awcgate/awc-ofut.htm).
3. It was accepted by and appeared in the *AI Journal* in 1984 as “Why AM and EURISKO Appear to Work.”
4. Turing’s original motivation for this was probably the “game” of Allied and Axis pilots and ground stations each trying to fool the enemy into thinking they were friendlies in World War II. Somewhat creepily, many humans today in effect play this game each day: men trying to “crash” women-only chatrooms, pedophiles pretending online to be 10-year-olds, and so on.
5. Not a footnote; Etc. to the fifth power, for dramatic emphasis, given the many millions of additional bits of knowledge that are needed to understand a single text, for example this morning’s entire *New York Times*.

6. Most people who’ve read or watched *2001: A Space Odyssey*—even most AI researchers—don’t know that Marvin Minsky and Arthur C. Clarke made HAL’s actions be logical, understandable, and hence tragic in the dramatic sense. HAL had been ordered to never lie to the crew, but then (just before launch) it was given the contradictory order to lie to the Jupiter crew about their true mission. Tragically, HAL deduced a mathematically elegant solution: just kill the entire Jupiter crew, after which HAL would neither have to lie to them nor not lie to them! HAL just hadn’t been told that killing someone is worse than lying to them.

7. See the poignant brochures and the various “in the media” links at www.smallrobot.com/androbot.html.

8. See www.geocities.com/fastiland/embodiment.html for examples of several flavors of this belief. For example, from Seitz (2000): “We do not simply inhabit our bodies; we literally use them to think with.”

9. Today’s variant of this is to believe that even if the answer is not short and simple, a short, simple, elegant algorithm plus fast enough hardware will lead to AI. Then we can at least fit that algorithm on the T-shirt.

10. Though to borrow from one of my favorite Dan Quayle-isms, a cynic might point out that half of all the people involved in making AI funding decisions are below the median! But so are the researchers....

11. The fourth derivative is called snap, which has no relevance to this discussion but could go on some sort of T-shirt no doubt. The fifth and sixth are called crackle and pop, a joke because the fourth derivative is called snap.

12. As they said, “Computer technology is now entering a new phase, one in which it will no longer be necessary to specify exactly how the problem is to be solved. Instead, it will only be necessary to provide an exact statement of the problem in terms of a ‘goal’ and the ‘allowable expenditure,’ in order to allow the evolution of a best program by the available computation facility.” For a good survey of recent work in this paradigm, see *Evolutionary Computation* by Larry’s son David Fogel.

13. Given a specification of $\{(3, 13.5) (1, 0.5) (10, 500)\}$, one “very fit” program cubes its input and divides by 2.

14. Simulated annealing was an attempt to recover from this bad situation, by occasionally making large leaps from (hopefully) one “hill” to a separate and possibly higher “hill.”

15. For example, is a “computer doctor” someone who diagnoses and fixes computers or an automated physician?

16. By prior presumed knowledge, I mean something that the speaker or author can reasonably expect the listener or reader to already know because it has recently been discussed, because of the intended audience, because of personal knowledge they have about the listener or reader such as their age, job, place of birth, having a Texas driver's license, being a Gold member of the American Airlines Advantage program, standing at a bus stop, wearing a Cubs baseball cap, and so on. This also includes more or less all of what is called common sense—if you turn your glass of water upside down, the water will fall out, for example.

17. Others have said similar things far more diplomatically. See for example Rosen-schein and Shieber (1982).

18. This is one place where vast amounts of raw computational cycles, as is going on in your brain as you read this sentence, might turn out to be a necessity, not a luxury. Thus hardware might be a limiting factor to AI. But not a simple 1000x speedup, else where is the “real AI” that functions in 1000x real-time today?

1. In a private communication to me in 2007.

References

- Ahn, L. v., and Dabbish, L. 2004. Labeling Images with a Computer Game. In *Proceedings of the 2004 Conference on Human Factors in Computing Systems, CHI 2004*, Vienna, Austria. New York: Association for Computing Machinery.
- Brown, J. S., and Van Lehn, K. 1980. Repair Theory: A Generative Theory of Bugs in Procedural Skills. *Cognitive Science* 4(4): 379–426.
- CBS News. 2005. “Pentagon Goes to the Robotic Dogs,” July 21. (www.cbsnews.com/stories/2005/07/21/eveningnews/main710795.shtml).
- Chapman, G. B., and Malik, M. 1995. The Attraction Effect in Prescribing Decisions and Consumer Choice. *Medical Decision Making* 15: 414ff.
- Crevier, D. 1993. *The Tumultuous History of the Search for Artificial Intelligence*. New York: Basic Books.
- Gleim, V. 2001. New Definition May Increase Number of Heart Attack Cases. The University Record, University of Michigan, Ann Arbor, MI. (www.umich.edu/~urecord/0001/Apr02_01/13.htm).
- Feigenbaum, E., and Feldman, J., eds. 1963. *Computers and Thought*. New York: McGraw-Hill.
- Fogel, D. 2005 *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*. New York: Wiley.
- Fogel, L. J., Owens, A. J., and Walsh, M. J. 1966. *Artificial Intelligence through Simulated Evolution*. New York: Wiley.
- Friedberg, R. M. 1958. A Learning Machine. *IBM Journal of Research and Development* 2(1): 2–13.
- Genova, J. 1994. Turing's Sexual Guessing Game. *Journal of Social Epistemology* 8(4): 313–326.
- Gorman, M. E. 1987. Will the Next Kepler Be a Computer? *Science & Technology Studies* 5(2): 63–65.
- Guccione, S., and Tamburrini, G. 1988. Turing's Test Revisited. In *Proceedings of the 1998 IEEE International Conference on Systems, Man and Cybernetics*, Vol. 1. Beijing and Shenyang, China, 38–41. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Ann Arbor, MI: University of Michigan Press.
- Hulse, C. 2003. Pentagon Abandons Plan for Futures Market on Terror. *The New York Times*, Washington Section, July 29. (www.nytimes.com/2003/07/29/politics/29WIRE-PENT.html?ex=1374897600&en=61db3b09007fa0a3&ei=5007).
- Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.
- Laird, J. E., and Rosenbloom, P. 1994. The Evolution of the Soar Cognitive Architecture. In *Mind Matters: A Tribute to Allen Newell*, ed. D. Steir and T. Mitchell. Mahwah, NJ: Erlbaum and Associates.
- Langley, P. 1977. BACON: A Production System That Discovers Empirical Laws. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge. Los Altos, CA: William Kaufmann Inc.
- Lenat, D. B., and Brown, J. S. 1984. Why AM and EURISKO Appear to Work. *Artificial Intelligence* 23(3): 269–294
- Lenat, D. B., and Feigenbaum, E. A. 1991. On the Threshold of Knowledge. *Artificial Intelligence* 47(1): 185–250.
- Matuszek, C.; Witbrock, M.; Kahlert, R.; Cabral, J.; Schneider, D.; Shah, P.; and Lenat, D. B. 2005. Searching for Common Sense: Populating Cyc from the Web. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, Pittsburgh. Menlo Park, CA: AAAI Press.
- Minsky, M., ed. 1968. *Semantic Information Processing*. Cambridge, MA: The MIT Press.
- Rosenschein, S. J., and Shieber, S. M. 1982. Translating English into Logical Form. In *Proceedings of the Twentieth Annual Meeting on Association For Computational Linguistics*. Toronto: University of Toronto.
- Seitz, J. 2000. The Bodily Basis of Thought. *New Ideas in Psychology: An International Journal of Innovative Theory in Psychology* 18(1): 23–40.
- Stork, D., ed. 1996 *HAL's Legacy: 2001's Computer as Dream and Reality*. Cambridge, MA: The MIT Press.
- Turing, A. M. 1950. Computing Machinery and Intelligence. *Mind* 59(236): 433–460. (www.loebner.net/Prizef/TuringArticle.html).
- Tversky, A., and Kahneman, D. 1983. Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment. *Psychological Review* 90(4): 293–315.
- Wells, H. G. 1895. *The Time Machine: An Invention*. London: Heinemann.

Doug Lenat has worked in diverse parts of AI—natural language understanding and generation, automatic program synthesis, expert systems, machine learning, and so on—for going on 40 years now, just long enough to dare to write this article. His 1976 Stanford Ph.D. dissertation, AM, demonstrated that creative discoveries in mathematics could be produced by a computer program (a theorem proposer, rather than a theorem prover) guided by a corpus of hundreds of heuristic rules for deciding which experiments to perform and judging “interestingness” of their outcomes. That work earned him the IJCAI Computers and Thought Award and sparked a renaissance in machine-learning research. Lenat was on the computer science faculties at Carnegie Mellon University and Stanford, was one of the founders of Teknowledge, and was in the first batch of AAAI Fellows. He worked with Bill Gates and Nathan Myhrvold to launch Microsoft Research Labs, and to this day he remains the only person to have served on the technical advisory boards of both Apple and Microsoft. Since 1984, he and his team have been constructing, experimenting with, and applying a broad real-world knowledge base and reasoning engine, collectively “Cyc.” For 10 years he did this as the principal scientist of the Microelectronics and Computer Technology Corporation (MCC) and since 1994 as the chief executive officer of Cycorp. He is on the technical advisory board of TTI Vanguard, and his interest and experience in national security has led him to regularly consult for several U.S. agencies and the White House. He edits the *Journal of Applied Artificial Intelligence* and the *Journal of Applied Ontology* and has authored three books, almost 100 refereed articles, and one almost-refereed article, namely this one.