

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>

Available at www.sciencedirect.com

SciVerse ScienceDirect

journal homepage: www.elsevier.com/locate/bica

INVITED ARTICLE

The next 50 years: A personal view

Patrick Henry Winston

Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, United States

Received 5 March 2012; received in revised form 22 March 2012; accepted 23 March 2012

KEYWORDS

Biologically inspired cognitive models;
Human intelligence;
Evolution of intelligence;
Inner language;
Story understanding;
Directed perception

Abstract

I review history, starting with Turing's seminal paper, reaching back ultimately to when our species started to outperform other primates, searching for the questions that will help us develop a computational account of human intelligence. I answer that the right questions are: What's different between us and the other primates and what's the same. I answer the *what's different* question by saying that we became symbolic in a way that enabled story understanding, directed perception, and easy communication, and other species did not. I argue against Turing's reasoning-centered suggestions, offering that reasoning is just a special case of story understanding. I answer the *what's the same* question by noting that our brains are largely engineered in the same exotic way, with information flowing in all directions at once. By way of example, I illustrate how these answers can influence a research program, describing the Genesis system, a system that works with short summaries of stories, provided in English, together with low-level *common-sense rules* and higher-level *concept patterns*, likewise expressed in English. Genesis answers questions, notes abstract concepts such as *revenge*, tells stories in a listener-aware way, and fills in story gaps using precedents. I conclude by suggesting, optimistically, that a genuine computational theory of human intelligence will emerge in the next 50 years if we stick to the right, biologically inspired questions, and work toward biologically informed models.

© 2012 Elsevier B.V. All rights reserved.

The right questions

Alan Turing was born 100 years ago, in 1912, so this is a year that invites those of us interested in developing computational theories of intelligence to look back at where we have been, reflect on our successes and mistakes, and formulate goals for the future.

E-mail address: phw@mit.edu

Highlights of the past 50,000 years

In his paper, *Computing Machinery and Intelligence* (Turing, 1950), Turing introduced his "Turing Test," positing a situation in which an interrogator attempts to determine which of two texters (modern language), a computer and a person, is the computer. If the interrogator has no more than a 70% chance of making the right identification after 5 min, the computer has passed the test.

Turing's paper, in my view, had a more important contribution: he exhibited the popular arguments against the possibility of intelligent computers and demolished them. Then, 10 years later, Marvin Minsky wrote another seminal paper, *Steps Toward Artificial Intelligence* (Minsky, 1961) that organized and surveyed the field and suggested how research ought to proceed. In that same year, James Slagle completed his PhD thesis, in which he exhibited a program that impressively solved integration problems from examinations in freshman calculus.

Together, Turing, Minsky, and Slagle anchored the modern history of artificial intelligence and created enormous excitement. Turing told us there was no reason to believe computers could not be as intelligent as humans; Minsky told us what to do to make it happen; and Slagle demonstrated that a computer can do what ordinarily is considered reserved for particularly intelligent people and provided a preview of the power of expert systems.

I think it interesting to reach back further to when people first started to think about whether computers can be smarter than people. The question has a history that includes the thoughts of the Countess Lovelace, who worked with Charles Babbage on his attempt to build his analytic engine. In 1842, Lovelace wrote, as quoted in Turing's paper:

The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform.

We can excuse her naivete, because in those days it was surely inconceivable that computers could do more, and in a restricted sense it is true, but she would have conveyed a more accurate picture if she had said they can only do what we program them to do, what we teach them to do, and what they discover how to do on their own. That, however, might not have had the soothing effect she was looking for.

Lovelace contributed new thoughts about whether computers could think, but people had been thinking about thinking for a long time before that, with the Greek philosophers being conspicuous examples. Plato, in particular, wrote about thinking extensively in *The Republic*, as his description of the ideal state—consisting of the merchants, soldiers, and philosopher kings—is generally regarded as an allegorical description of the well-organized mind. In fact, *The Republic* is considered a weak translation of the untranslatable Greek title, *Politeia*, and inasmuch as scholars argue about the right translation, I suggest that one might be *Society*, in which case, *The Republic* might be better translated as *The Society of Mind*, anticipating Minsky's book with that title by more than 2000 years.

To me, however, the most important pre-Turing event was not when Lovelace thought about computer thinking, nor when Plato thought about thinking, but when we started thinking. According to Tattersall, the noted paleoanthropologist (Tattersall, 1998), the human separation from other primates occurred very recently, not much more than 50,000 years ago, when some small accident of evolution made us suddenly, in an imprecisely specified sense, symbolic. Before that we coexisted with, and were not much different from, the Neanderthals for 100,000 years or so. We made simple stone tools, so did they. We managed and perhaps knew how to light fires, and so did they. But

then, suddenly, whatever they could do, we began to do much better.

The sudden divergence from the Neanderthals and other primates raises two central questions:

- What is it about our species that makes us different from other primates?
- What is it about the biology we share with other primates that enables the human difference to emerge and matter.

I think Turing's answer to the *what's different* question would have been that we reason, and they did not; and given Turing's test, and his suggestions for how to build an intelligent machine, he did not seem to think anything other than reasoning was important.

I have a very different answer to the *what's different* question and an obvious answer to the *what's the same* question. But first, I return to the time of Turing, Minsky, and Slagle and move forward.

Highlights of the past 50 years

By the mid 1960s, pioneering thinkers had established laboratories for artificial-intelligence research, all with an emphasis on understanding reasoning in one form or another. At MIT, Minsky encouraged students to address issues he wrote about in *Steps toward Artificial Intelligence*; at Carnegie-Mellon, Allen Newell and Herbert Simon attempted to understand secrets of cognition exposed when people solve, for example, cryptarithmic problems; and at Stanford, John McCarthy attempted to bend symbolic logic to his will.

Out of these laboratories grew a field full of ambition, populated with both followers of the pioneers and dissenters to their views, together producing a cornucopia of systems, algorithms, architectures, programming languages, principles, subfields, powerful applications, and a second generation of influential thinkers.

Reflecting on the influential thinkers of the second generation, I myself felt most informed by Roger Schank and David Marr, both of whom were, in their own way, controversial figures.

Schank (1972, 1977, 1981) imprinted me with the idea that stories are important and influenced my earliest work in analogical reasoning (Winston, 1980). Although my approach to stories is different in detail, I share Schank's overall view, which persists Schank (1991), that stories have a central role to play in human thinking. When Schank turned his attention to the role of stories in human education, and retreated from the implementation of story understanding systems, it was a loss to artificial intelligence, but a gain to education.

Marr (1982) led me to believe in the importance of putting the problem before the solution. His focus on low-level vision attracted a great deal of complaint for various reasons, and his research had less than spectacular staying power, but he is much honored, especially in cognitive science, for his argument that the first step in any research program ought to be to characterize the behavior you are working to understand; then you should pick an appropriate

representation that exposes constraint; then you should start thinking about what sort of method to develop, and only then should you implement particular algorithms. Accordingly, I think Marr would not be much impressed by approaches that rise up and claim to be the answer without being particularly explicit about what the question is.

The work of Schank, Marr, and the pioneers made the late 1970s exciting times. Yet no one ever said we were ahead of schedule. There grew, by the mid 1980s, a sense that the field had succeeded somewhat as an engineering discipline, but not as a scientific inquiry into human or human-level intelligence. New approaches crystallized and others matured, with the second half of the 1980s being of special note. Minsky shifted toward an interest in how our thinking is organized and published *The Society of Mind* (Minsky, 1988), arguing that intelligence is the product of a multiplicity of collaborating ways of solving problems with a myriad of representations. In that same half decade, John Laird, Allen Newell, and Paul Rosenbloom, published "SOAR: An Architecture for General Intelligence" (Laird, Newell, & Rosenbloom, 1987) describing their thoughts on what must be common to intelligent systems, and hence, part of a general architecture. David Rumelhart and James McClelland published *Parallel Distributed Processing* (McClelland & Rumelhart, 1986), suggesting that the road to success had to go through artificial neural networks. Judea Pearl published *Probabilistic Reasoning in Intelligent Systems* (Pearl, 1988), arguing that probabilistic methods are the answer. And in 1987 Rodney Brooks submitted for publication "Intelligence without representation" (Brooks, 1991), but its anti-establishment views were evidently so controversial that the paper was not published until several years later.

As for neural nets, as described in Rumelhart and McClelland, somehow I have not been able to work up much interest in them, because they seem to be a mechanism, rather than a method that comes out of the study of a problem derived from a behavior to be understood. As a mechanism, they seem to me to be function approximators, and I do not see how models of function approximators explain much of intelligence, such as, for example, how an understanding of neural nets would help me explain either how Slagle's program solves calculus problems or the engineering found in the brain.

On the other hand, I find myself interested in Minsky's work and the work of those who develop and implement cognitive architectures such as SOAR and ACT-R, because they attempt to model that part of our intelligence that is symbolic and addresses the *what's different* question.

As for Brooks's work and that of the statistical community, both have had high engineering impact, but I think both have also created an unfortunate confusion of purpose and deserve more discussion after I express my views on the role of models.

The role of models

The word *model* has many definitions, one of which centers on the idea that a model is a surrogate that behaves in some respects like the real thing. A model's behavior may be the impressive part, or it may enable prediction, or it may aid understanding.

An ordinary thermometer has no model of the world; it just does what it does, and as such, it is extremely useful.

It is easy to model what it does mathematically, and a mathematical model certainly behaves in some respects like a real thermometer, enables prediction, and aids understanding.

Much of what we and other animals do is model free, like a thermometer, but some of what we and other animals do seems to depend on a kind of predictive simulation, enabled by the considerable reuse of perceptual and motor apparatus (Damasio, 1999; Moore & Haggard, 2008; Wolpert, 1997). As Hesslow (2002) notes in an enlightening review, we can imagine actions, and when we do, we seem to excite some of those areas in the brain that are involved in executing those actions; we can imagine perceptions, and when we do, we excite some of the areas in the brain involved in actual perceiving; and intriguingly, even in rats, simulated actions seem to invoke simulated perceptions. Such predictive power certainly suggests that a kind of model is at work.

Of course, we can increase our understanding of human and animal intelligence by making models of the thermometer-like aspects of perception and action and by making models of the simulation that seems to go on in perception and action. We also can increase our understanding of the world by making models of blood flow, image formation, and thrown rocks, and our models will be full of knowledge of Reynolds numbers, reflectance functions, and the equations of elementary physics.

The important part of the previous paragraph is often overlooked. We can, in fact, make models that not only enable prediction but also increase understanding. I think the great enabler of that kind of model making is story processing, but in any case, we do make informing models, and in particular, we can make models of our own model-making ability. Thus, there are model types that we should not confuse. We model those parts of ourselves that:

- Just do what they do.
- Perform predictive simulation.
- Make models.

And there are purposes that we should not confuse:

- To mimic behavior.
- To make predictions.
- To increase understanding.

Weizenbaum's classic Eliza program (Weizenbaum, 1965) illustrates. It is a model, in that it behaves in some respects like a Rogerian psychotherapist, perhaps it occasionally made some accurate predictions about what a Rogerian psychotherapist would say; but it sheds no light whatsoever on what is going on in the head of any psychotherapist, Rogerian or otherwise, so it is a weak, uninformative model.

Intelligence without understanding

In "Intelligence without representation," Brooks argued that his subsumption architecture is a step toward developing systems with insect-level behavior, and he described subsumption-based robots that successfully navigated halls, occasionally were made to pick up soda cans, and eventually led to the highly successful Roomba[®] vacuum cleaner. On the other hand, it is hard to believe that insects, which exhibit the level of intelligence aspired to, have separate perceptual and motor systems in each layer of a subsumption

architecture, and even harder to see any evidence that finite state machines substantially increase our understanding of how their nervous systems work. Thus, subsumption seems to me to be a model of insect behavior in the weak Eliza sense.

Brooks did not care. His goal was not to understand nature; it was to build creatures that usefully populate the physical world. Because his robots behave, in some respects, like insects, they are models, but they are models that shed little light on how insects achieve the behavior they exhibit, and so it is not surprising that no one has any clue about how to build a robot with the survival abilities of the common mosquito.

Brooks erred, however, in supposing that achieving insect-level performance his way would trivialize the development of systems with human-level intelligence and that representations would not be involved. He offered that evolution took much longer to go from first life to insects than from insects to humans and provides a seductive feathers-and-flight analogy, but because his robots exhibit only some of the behavior of nature without adding to an understanding of nature, and because his robotic work focuses on that part of animal behavior that is reactive, like a thermometer, rather than predictive or model making, I do not see it as offering substantial insight.

The retreat into statistics

Pearls's *Probabilistic Reasoning in Intelligent Systems* is a book that launched a thousand students, or maybe a lot more, and like Brooks's robots, statistical systems have enjoyed conspicuous practical success.

Nevertheless, most statistical systems do not interest me, except as engineering solutions, because, like Brooks's robots, they merely behave in some respects like the real thing, without shedding light on how the real behavior is achieved. After all, you could, equipped with statistical inference and a lot of time, make accurate predictions of how fast a ball would roll down an inclined plane, but that capability would not provide the insight provided by Newton's laws. Some might be satisfied with it, but I would not.

Of course, we humans are awash in regularity. The sun keeps coming up; rain often follows the arrival of dark clouds. So I have no quarrel with attempts to make models of how we model regularity. I only argue that statistical tools are just part of the answer to such questions, not the whole answer, nor, I think, likely to be the most important part of the whole answer.

Application success, science frustration

The proposal that led to the 1957 Dartmouth Conference included the following statement of objectives:

An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.

Today, from an applications perspective, all these objectives have been reached, at least to a minimal degree, and those who have worked in the field of artificial intelligence should be proud. Highly conspicuous systems now respond to our spoken commands and written questions, discover

exploitable regularity in data, allocate expensive resources, search the web, drive vehicles, beat champion chess players, and win Jeopardy contests. Less conspicuously, it is difficult today to build a big, important system without drawing on the products of artificial-intelligence research.

Nevertheless, given the center of gravity in artificial-intelligence research today, the science-side dream of developing a computational account of human intelligence remains elusive.

The right answers

How can we make progress toward the science-side dream? My answer is by asking better, biologically inspired questions. Some argue that intelligence is intelligence, and it does not matter if the machine exhibiting it is wet or dry inside, so why be concerned with biology. I argue that it makes no sense to wander around in a large space, full of bad ideas, largely populated with Scyllas and Charybdises.

What's different?

Noam Chomsky, who often cites Tattersall, thinks that becoming symbolic means developing the ability to combine concepts, making new concepts, without limit. To me, this sounds like the ability to construct something like a semantic net. But I would go beyond that and suggest that the most important concepts to combine are event descriptions, and that we combine event sequences into sequences, and that we move backward and forward in remembered sequences, and that with the ability to move backward and forward, we can tell stories, understand stories, and combine old stories to make new ones. And that, I think, constitutes part of the answer to the *what's different* question.

It is an especially important answer from the model perspective, because I believe that our human story manipulation ability, with all its diverse components, is what makes it possible for us to make elaborate models of both ourselves and the physical world, and therefore if we are to understand human thinking, we have to model that story-manipulation, model-enabling capability. In the end, that is what makes us different from species that have plenty of just-do-it and simulation capabilities, but whose story manipulation capability, if any, is on another, much lower level.

From this perspective, Brooks's claim that we should reject representation-centered approaches seems mysterious. If human intelligence is all about our acquisition of story understanding capabilities, then those stories are models of the world, because we use them to explain, predict, and control. And to have a model, you need a way to acquire knowledge of regularity and constraint. And to acquire knowledge of regularity and constraint, you need powerful representations. And only one animal, as far as we know, developed model-enabling representations at our human level.

There is another part to being symbolic, however, and that has to do with our ability to deploy our perceptual apparatus to solve problems using either real or imagined events. Imagine that an orangutan gives a ball to a chimpanzee. Who has the ball now? Sure, you could write a simple rule system to answer such questions, but to me, it feels like I imagine the scene and read the answer off of my imagined screen with visual perception.

From this perspective, most attempts to acquire large amounts of common-sense knowledge go forward with one boot off, because we humans develop much or most of our common sense as we need it, just in time, evidently using our visual and motor simulation apparatus. Imagine, for example, running down the street with a full bucket of water. What happens? The answer is not in the web. Watson, IBM's Jeopardy player, does not know. But you know your leg will get wet. Replace the water with nails. You will run bent over. You picture it.

Thus, being symbolic to me means that we have an inner language that supports story understanding, the acquisition of common sense from perception, and the ability to communicate with others in our social group using our myriad outer languages.

What's the same?

On the other hand, being symbolic is not all that needs to be studied. Much in what our relatives also possess by way of nonsymbolic problem solving in perceptual and motor systems is necessary too, and therefore, we need biological inspiration to help us understand nonsymbolic computing, and the ways in which we are the same as other primates, but that leads quickly to questions about how our brains are organized.

Open a standard textbook on neurobiology and you will find somewhere an illustration with an eye on one side and the primary visual cortex on the other. Signals flow from the eye's retina to the optic chiasm; then on to the lateral geniculate, explained as a kind of way station; and then to the primary visual cortex. It is a comfortable diagram. We engineer such systems, with a few modules lined up from left to right, with information flowing through them.

The trouble is, neurobiology textbooks also note that 80% of the input to the lateral geniculate comes from somewhere other than the retina. A good deal comes down from the primary visual cortex, suggesting that vision is a matter of guided hallucination. Other substantial input comes from auditory apparatus. Everything is all mixed up, with information flowing bottom to top and top to bottom and sideways too. It is a strange architecture about which we are nearly clueless.

And our brain's strange architecture exhibits other intriguing characteristics beyond information flow in all directions. Here are examples of other intriguing characteristics found in our vision system:

- Image information emerges at multiple scales.
- Information in eye coordinates becomes information in log-polar coordinates in the primary visual cortex.
- Motion aids segmentation and recognition.
- Top-down hallucination extracts features on a no-reason-not-to-believe basis.
- Processing reflects non-uniform, fovea-periphery organization.
- Multiple image processing modalities work not only together but also with other perceptual faculties.

In contrast, the typical computer vision system processes one frame at a time, uniformly, with a series of modules processing input and producing output, left to right, like a

typically engineered system, rather than an exotic one. Thus, applied vision systems are not very biologically inspired, and that, perhaps, is why, after substantially more than 50 years of research in pattern recognition and artificial intelligence, vision systems often confuse trees with people. Evidently, much remains much to be done, and much of that will require us to study the ways in which we process the kind of information that is readily processed in the same way by animals.

The work of the next 50 years

I believe our understanding of intelligence is in the state biology was in before Watson and Crick determined the structure of DNA. Back in 1950, there were antibiotics to be sure, but there was also much skepticism about whether there would ever be a deep understanding of how biology works.

Likewise, today we have important applications, but also a measure of doubt about whether there will ever be a deep understanding of how intelligence works.

Nevertheless, I am optimistic. Brain and cognitive scientists have new tools and offer many new clues. And in contrast to the early days, no matter how much computing you need, you can get it.

On the other hand, you cannot make progress unless you work on the right questions, and if you work on the right questions, progress can be fast. As J.C.R. Licklider once said to me, we tend to overestimate what we can do in a year and underestimate what we can do in 10 or 20.

What then are the right questions? At the highest level, the right question is: How are we the same and how are we different? And that answer, I think, leads to five families of critical-path, midlevel questions.

Questions that flow from the inner language: What inner-language-enabled representations constitute most of what we need when we think? To what degree can our representation repertoire expand with experience? Why, at any given time, do we have only one inner conversation going on in our heads?

Questions that flow from story understanding: How do we exploit stories so as to answer questions, find precedents, and tell stories persuasively? How do the stories in a culture influence thinking? How do names for concepts give us power over them? How can we scale up understanding systems to deal with large numbers of stories and stories that range from fairy tales to business cases?

Questions that flow from directed perception: What questions do our various faculties ask of one another? How does perceptual imagination enable common-sense problem solving in the absence of actual experience?

Questions that flow from our social nature: How do we build models of ourselves, models of others, and models of how we will react to the models others make of us. How can we combine our inner language knowledge of what can be meant with what others attempt to say in an outer language so as to produce meaning?

Questions that flow from exotic engineering: How does hallucination from inner models combine with information from the senses to produce a sense of understanding in the midst of noise? How do our diverse faculties communicate and learn to communicate with one another? How do

we learn from a fundamentally noisy, unlabeled environment? How do our brains perform so reliably and robustly.

Questions such as these invite focused study and focused study leads to sharper questions. I turn now to two representative examples.

The Genesis example

Given my personal views, it is natural that my students and I work on story understanding and our work is centered on the construction of the Genesis story understanding system. I offer it here as an example of an attempt to model the kind of human modeling that goes on when we understand, remember, retrieve, exploit, mix, and tell stories.

Shakespeare's plots, provided to Genesis in summary form, are typical of what we aspire to have Genesis understand:

Macbeth: Macbeth, Macduff, Lady Macbeth, and Duncan are persons. A thane is a kind of noble. Macbeth is a thane and Macduff is a thane. Lady Macbeth, who is Macbeth's wife, is greedy. Duncan, who is Macduff's friend, is the king, and Macbeth is Duncan's successor. Macbeth defeated a rebel. Witches had visions and talked with Macbeth. The witches made predictions. Duncan became happy because Macbeth defeated the rebel. Duncan rewarded Macbeth because Duncan became happy. Lady Macbeth, who is Macbeth's wife, wants to become the queen. Lady Macbeth persuades Macbeth to want to become the king. Macbeth murders Duncan. Then, Lady Macbeth kills herself. Dunsinane is a castle and Burnham Wood is a forest. Burnham Wood came to Dunsinane. Macduff had unusual birth. Macduff fights with Macbeth and kills him. The predictions came true.

Does Duncan end up dead? Is there revenge? The words *dead* and *revenge* do not appear, so if Genesis is to reach such conclusions, it cannot be done by looking for words. Instead, it must know not only that which appears explicitly in the story, but also have common-sense knowledge and knowledge of how to reflect on its own conclusions.

Common sense is supplied by if-then rules, supplied in English. Here is a rule for dealing with the Duncan question:

- If X kills Y, then Y becomes dead.

Higher level concept knowledge is supplied by search descriptions, also supplied in English. Here is the revenge description:

- X's harming Y leads to Y's harming X.

Genesis's common-sense rules and concept patterns are expressed in terms of a rich repertoire of representations, developed by Genesis's builders or borrowed from research variously in artificial intelligence and linguistics such as that of Borchardt (1994), Vaina and Greenblatt (1979), Jackendoff (1985), Schank (1972), and Talmy (1988). Our goal is to develop a representationally complete theory of what it is possible to know, or at least a theory that covers a large part of what we know. We already see on the horizon a ranking of representations by frequency according to the type of story involved.

Fig. 1 illustrates what Genesis does with the Macbeth summary. The story supplies the information in the white boxes; common-sense rules supply the information in the gray boxes. Boxes are connected by causes.

Interestingly, there are a lot of gray boxes, coming not from the story but rather from previous knowledge and chaining. It is as if the author supplies just enough content in the story to keep us on track, leaving us to fill in the gaps. It would seem that not only visual perception, but also story understanding, involve a substantial measure of hallucination.

Note also that there is a path from Macbeth's harming of Macduff to Macduff's harming of Macbeth. Thus, a search based on the search description for *revenge* succeeds, and Genesis reports that the Macbeth summary involves revenge.

Building on such basic capabilities, we have provided Genesis with the ability to find abstract concepts in stories (Nackoul, 2010), to read stories in a culturally biased way (Winston, 2011), to indicate if a violent situation is escalating, to identify the onset of unintended consequences, to find precedents based on conceptual understanding rather than words, to fill gaps in stories using precedents (Fay, 2012), and to tell stories in a way that is sensitive to the understanding and biases of the listener.

Work on Genesis reflects a view of what it means to be symbolic that contrasts sharply with early work in artificial intelligence, which seems to have taken the position that being symbolic meant only that we can reason logically and nothing else matters. I think that reasoning is recipe following and recipe following is a special case of story understanding.

Consider Slagle's symbolic integration program, perhaps the first reasoning program that achieved expert human-level performance. Slagle's program, in retrospect, was a rule-based system of algebraic transformation rules deployed so as to construct a goal tree in which a given problem connects via transforms to a library of integrals. Genesis, in retrospect, is a system of common-sense rules deployed so as to construct an elaboration graph in which concepts emerge from the discovery of connections between causes and downstream results. One seems a special case of another.

The propagator example

The Genesis system exploits some of the ideas in Gerald Sussman's Propagator Architecture (Radul & Sussman, 2009), which was conceived as a departure from ordinary programming, born of a desire to develop programs that are powerful, robust, flexible, agile, adaptable, and capable of dealing with the fundamentally noisy, chaotic, uncertain, confusing, ill-posed, and ambiguous world around us.

That is, Sussman developed the propagator architecture to advance programming, but his criteria were such that his goal could have been to ensure that programs could be more like people.

The first part of the propagator architecture idea is that systems should consist conceptually of boxes connected by wires. Each box watches its ports and reacts to those signal-like objects it understands, transmitting new signals, also through ports. Each box ignores inputs it does not know how to handle.

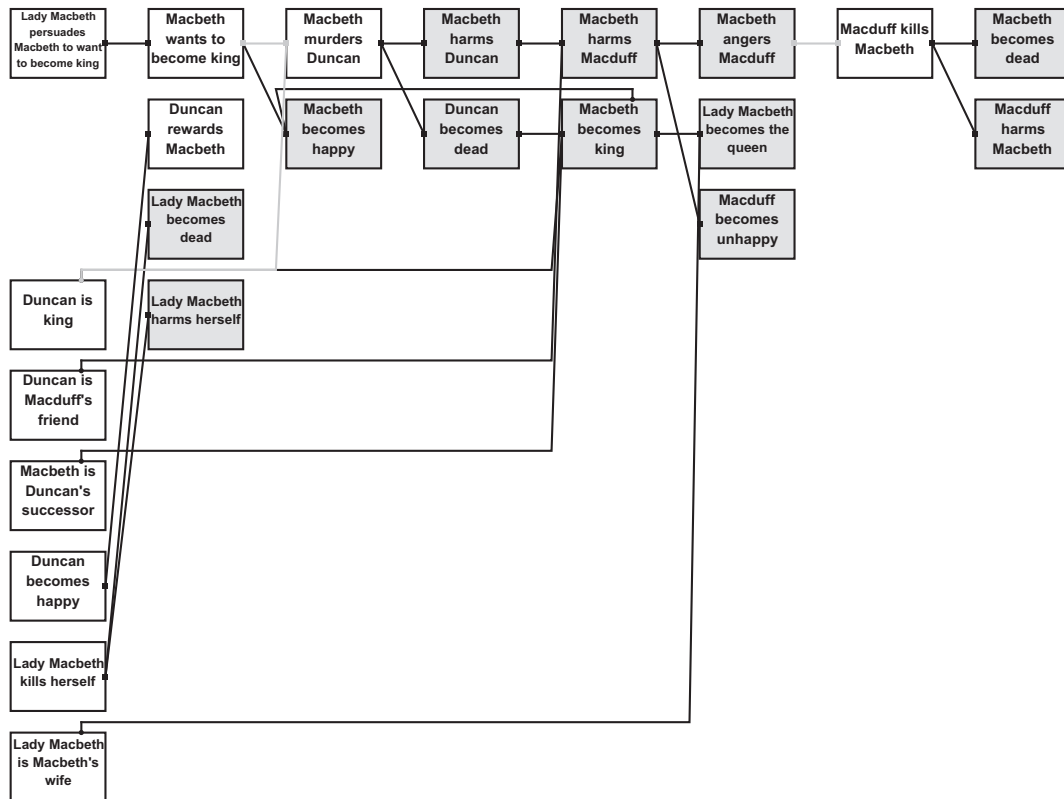


Fig. 1 Genesis's story understanding system produces a graph using common-sense rules together with a story. White boxes indicate information given explicitly in the Macbeth story. Gray boxes indicate information produced by common-sense rules. The gray connection between *Macbeth angers Macduff* and *Macduff kills Macbeth* indicates a tentative cause presumed in the absence of definite cause.

The propagator architecture distantly resembles the abstraction and interface idea, but provides even more constraint and further reduces the programmer's radius of essential comprehension. Thus, the programmer in charge of a box need not understand all the details that lie beyond the box's ports, so the newly involved programmer can get started right away and never needs to understand an entire

system, which is especially important when a large system is built by student labor.

The second part of the propagator architecture idea is that boxes do what they can with what they have, combining evidence from multiple ports to reach conclusions about the values that should be propagated to other ports. A distance estimating propagator box, for example, might have

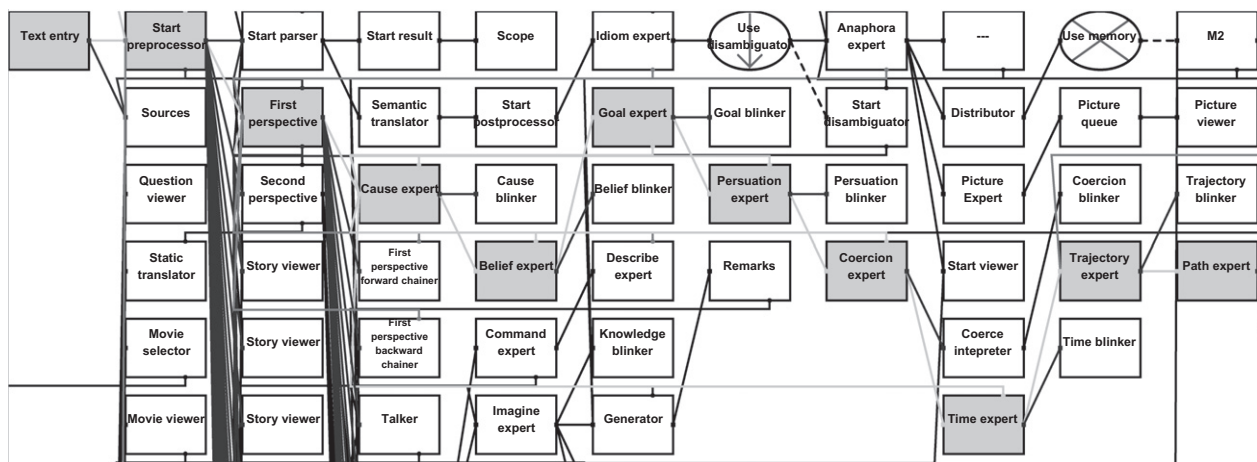


Fig. 2 Genesis exploits ideas from Sussman's propagator architecture in which modules are boxes connected conceptually by wires. In the fragment of Genesis shown, highlighted boxes are specialists that lie along an interpretation path that extracts information about cause, believe, goal, persuasion, coercion, time, trajectory, and path. Specialists not shown include those that extract information about place, transition, role, social relation, mood, property, possession, and class.

ports wired to other boxes that compute distance using, for example, stereo, or motion parallax, or focus accommodation, or knowledge of approximate object size.

The third part of the propagator architecture idea is that boxes are inherently bidirectional. Information can arrive on any port and results can propagate to any port.

So far, Genesis uses only the first part of the propagator architecture, the box-and-wire idea. Fig. 2 shows about a quarter of the Genesis system, highlighting some of the two-dozen semantic specialists that extract what they can from the parsed text flowing by.

Is such engineering reflective of the exotic engineering in the brain? Boxes and wires enforce modularity and promote extensibility and replicability. Bidirectionality ensures that information can flow in all directions. Combining information on multiple ports ensures robustness and ensures that multiple modalities can work together. All contribute to a capacity to deal with uncertainty and noise. So, yes, propagator architecture seems like a step toward an understanding of the exotic engineering in the brain, and in any event, is attractive to those attracted to biological inspiration.

Contributions

In composing this personal view, I have:

- Suggested that Turing's focus on reasoning was counter-productive and his question, can machines think, was the wrong question.
- Argued that the right high-level question is: What makes humans different from and the same as other primates.
- Articulated five families of questions that flow from thinking about the inner language, story understanding, directed perception, our social nature, and exotic engineering.
- Offered the Genesis system as an example of work motivated by a desire to understand how we humans are different from other primates, past and present.

My personal view is that attention to the five families of questions will ensure that the secrets of intelligence eventually will be revealed, just as the structure of DNA was revealed in 1953. Of course it is easier to predict the future than it is to predict exactly when the future will happen. It may take 50 years, or a hundred, but I rather think, with focused effort, 10 or 20 would be enough, and the really important consequence will not be artifacts, but rather a better understanding of ourselves and each other.

Acknowledgements

This paper is based, in part, on a talk given at the Symposium, *Computation and the Transformation of Practically Everything*, which was part of a series of symposia celebrating the sesquicentennial anniversary of MIT's founding. Research on Genesis has been supported, in part, by the National Science Foundation (IIS-0413206), the Office of Naval

Research (N00014-09-1-0597), the Air Force Office of Scientific Research (A9550-05-1-0321), and the Defense Advanced Research Projects Agency (FA8750-10-1-0076).

References

- Borchardt, G. (1994). *Thinking between the lines*. MIT Press.
- Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 47(1–3), 139–159.
- Damasio, A. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. New York: Harcourt Brace.
- Fay, M. P. (2012). *Enabling imagination through story alignment*. Master's thesis, MIT.
- Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6, 242–247.
- Jackendoff, R. (1985). *Semantics and cognition*. MIT Press.
- Laird, J. E., Newell, A., & Rosenbloom, P. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, 33(1), 1–64.
- Marr, D. (1982). *Vision*. Freeman.
- McClelland, J. L., & Rumelhart, D. E. (1986). *Explorations in parallel distributed processing*. The MIT Press.
- Minsky, M. (1961). Steps toward artificial intelligence. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. MIT Press.
- Minsky, M. (1988). *The society of mind*. Simon and Schuster.
- Moore, J., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition*, 17, 136–144.
- Nackoul, D. (2010). *Text to text: Plot unit searches generated from english*. Master's thesis, MIT.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufman.
- Radul, A., & Sussman, G. J. (2009). The art of the propagator. In *Proceedings of the 2009 international lisp conference*.
- Schank, R. C. (1972). Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 3(4), 552–631.
- Schank, R. C. (1991). *Tell me a story: A new look at real and artificial intelligence*. Northwestern University Press.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Lawrence Erlbaum Associates.
- Schank, R. C., & Riesbeck, C. K. (1981). *Inside computer understanding: Five programs plus miniatures*. Lawrence Erlbaum Associates.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12(1), 49–100.
- Tattersall, I. (1998). *Becoming human*. Harcourt.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- Vaina, L. M., & Greenblatt, R. (1979). *The use of thread memory in amnesic aphasia and concept learning*. Working paper 195, MIT Artificial Intelligence Laboratory.
- Weizenbaum, J. (1965). ELIZA—a computer program for the study of natural language communication between man and machine. *Communications of the Association for Computer Machinery*, 9.
- Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the ACM*, 23(12), 689–703.
- Winston, P. H. (2011). The strong story hypothesis and the directed perception hypothesis. In P. Langley (Ed.), *Papers from the AAAI fall symposium, Technical report FS-11-01* (pp. 345–352). AAAI Press.
- Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, 1, 209–216.