

Fifty Years of AI: From Symbols to Embodiment - and Back

Luc Steels

Sony Computer Science Laboratory - Paris
VUB AI Lab, Vrije Universiteit Brussel

Abstract. There are many stories to tell about the first fifty years of AI. One story is about AI as one of the big forces of innovation in information technology. It is now forgotten that initially computers were just viewed as calculating machines. AI has moved that boundary, by projecting visions on what might be possible, and by building technologies to realise them. Another story is about the applications of AI. Knowledge systems were still a rarity in the late seventies but are now everywhere, delivered through the web. Knowledge systems routinely deal with financial and legal problem solving, diagnosis and maintenance of power plants and transportation networks, symbolic mathematics, scheduling, etc. The innovative aspects of search engines like Google are almost entirely based on the information extraction, data mining, semantic networks and machine learning techniques pioneered in AI. Popular games like SimCity are straightforward applications of multi-agent systems. Sophisticated language processing capacities are now routinely embedded in text processing systems like Microsoft's Word. Tens of millions of people use AI technology every day, often without knowing it or without wondering how these information systems can do all these things. In this essay I will focus however on another story: AI as a contributor to the scientific study of mind.

Keywords: history of AI, heuristic search, knowledge systems, behavior-based robotics, semiotic dynamics.

1 Introduction

There has been a tendency to compare the performance of today's machines with human intelligence, and this can only lead to disappointment. But biology is not regarded as a failure because it has not been able to genetically engineer human-like creatures from scratch. That is not the goal of biology, and neither should the construction of human-like robots or programs that pass the Turing test be seen as the goal of AI. I argue that the scientific goal of AI is not to build a machine that somehow behaves like an intelligent system (as Turing proposed originally) but to come up with an explanation how intelligence by physically embodied autonomous systems is possible and could have originated in living systems, like us. AI takes a design stance, investigating in principle by what mechanisms intelligence is achieved, just like aircraft engineers try to figure out by what mechanisms physical systems are

able to fly. Whether the solutions we find are comparable or equal to those used by natural living systems like us is another issue about which AI itself has nothing to say. But one thing is sure, without an adequate operational theory of the mechanisms needed for intelligence, any discussion of human intelligence will always remain at the level of story telling and hand waiving, and such a theory cannot be conceived nor adequately tested only by empirical observation.

There has been a tendency to dismiss the AI achievements of the past, usually by those who are incapable to understand how they work and why they are important. For example, I think it is rather incredible that computer programs today play at grandmaster level. Note that long ago philosophers like Dreyfus were arguing that “computers would never be able to play chess” (Dreyfus, 1972). Now that everyone can see for themselves that they can (and Dreyfus lost the first game he ever played to a computer), critics argue that we can learn nothing about problem-solving or cognition from these systems. I believe these comments come from people who have never bothered to look at what is really behind game-playing programs and who have no notion of what the real issues are that problem solvers, whether natural or artificial, have to face.

I personally stumbled into AI in the early seventies through the door of (computational) linguistics, mainly inspired by the work of Terry Winograd (1972). I was interested in understanding the processes that could explain how language sentences are produced and interpreted but when building my first computer programs to test my hypotheses (using a computer which had less power than today’s mobile telephones!), it quickly became clear that my initial intuitions were totally wrong and naive, and so were the intuitions of the psychologists and linguists I encountered, simply because none of us were capable to think in computational terms except as a vague metaphor. I subsequently changed my research methodology completely and never had any regrets about it. It is indeed in the methodology that lies the uniqueness of AI. There is no unique AI theory, nor AI view on intelligence. A multitude of views have been explored, each time making contact and in turn influencing other scientific fields in a big way. I do not think that this multitude will shrink to a single view any time soon and maybe this is not necessary. The richness of AI lies in its openness towards other disciplines and the power of the experimental approach for operationalising and testing mechanistic theories of intelligence.

In this essay, I try to reconstruct, very schematically and from a very personal view of course, some of the key ideas and milestones in the non-linear AI history of the past decades. New movements seem to have come in waves, lasting about seven years before being absorbed in mainstream thinking. Then I give my opinion on what I find the most promising trends in current AI activities and which direction we might want to take in the future.

2 Prologue (1956): The Information Processing View

AI started from a single powerful idea, clearly understood and expressed at the Dartmouth founding conference (McCarthy, et.al., 1955): The brain can be viewed as an information processor. A computer is an example of an information processor but the notion of information processing is much broader. We are looking at a physical system that has a large number of states which systematically correspond to information states,

processes that transform these states systematically into others while preserving or implementing some (mathematical) function, and additional information states (programs) which regulate what processes should become active at what moments in time. From a biological point of view, neurons are cells that grow and live within larger networks, and biologists research the question how they can function and sustain themselves, similar to the way they study liver cells or the immune system. The information-processing view introduces a new perspective on the brain, neurons or networks of neurons are viewed as devices that can hold and transform information and so the question becomes: What kind of information is stored and what kind of transformations are carried out over this information. It is interesting that the switch from the biochemical, electrical point of view to the information processing point of view took place at the same time as other areas of biology, and particularly genetics and molecular and developmental biology, also started to emphasise information, viewing the genome as a program that is interpreted by the machinery of the cell. This view has become dominant today (Maynard Smith, 2000).

AI does not address the question of neural information processing directly but takes a design stance. It asks the question: What kind of information processing do living brains need in order to show the astonishing intelligent behavior that we actually observe in living organisms? For example, what information structures and processes are required to fetch from a stream of visual images the face of a person in a fraction of a second, what kind of information processing is necessary to parse a natural language sentence with all its ambiguities and complex syntactic structures and to interpret it in terms of a perceived scene, how is a plan formulated to achieve a complex collaborative action and how is this plan executed, monitored and adjusted? To investigate what kind of processes are in principle needed for intelligence makes sense for two reasons. First of all it is hard to imagine that we will be able to understand the functioning of enormously complex nervous systems with billions of elements and connections if we do not have tools to understand what they might be doing and if we can only do gross-scale measurements like fMRI or PET scanning. It is like trying to understand the structure and operation of computers that are performing millions of instructions per second and are linked with millions of other computers in giant networks like the Internet, without knowing what they might be doing and having only gross statistical measures at our disposal. Second, neural and psychological observations - brain imagining, brain disorders, effects from aging, etc. - can tell us that certain parts of the brain are involved in particular cognitive tasks, but they cannot tell us why certain types of processing are needed. To find that out, we must examine alternative mechanisms (even if they do not occur in nature) and study their impact on performance.

Information processing is a point of view and not an intrinsic property of a system. A computer can be built from anything, including TinkerToys, as Danny Hillis and Brian Silverman showed with his Tic-Tac-Toe playing computer at the MIT AI lab in the nineteen eighties, or with chemistry. So it is perfectly justified to use computers to study the behavior of natural information processing systems. It is also very straightforward to test whether a theory is valid or not. If someone has an idea what kind of information processing is needed for a task, an artificial system can be built that carries it out and anyone can observe whether it has the required performance.

The information processing view was a true revolution in the study of mind and it started to propagate rapidly after the first 1956 AI conference in Dartmouth, with

initially a major impact on psychology (particularly through Jerome Bruner, George Miller, Alan Newell, and Herbert Simon) and linguistics. Once the information processing view is adopted, the next step is to do the hard work and investigate in great detail concrete examples. Initially it was thought that there might be some very general, simple principles, like Newton's equations, that could explain all of intelligence. Newell and Simon proposed a General Problem Solver, Rosenblatt proposed Perceptrons, etc., but gradually, as more and more tasks were being tackled, it became clear that intelligence requires massive amounts of highly complex sophisticated information and very adapted information processing. Quite early, Minsky concluded that the brain must be a "huge, branchy, quick-and-dirty kludge", as opposed to a cleanly designed general purpose machine. This is of course not unusual in living systems. After all the human body is also a huge quick-and-dirty kludge with its myriad of interlocking biochemical processes and ad hoc solutions that have a high probability of breaking down. Indeed, if you come to think of it it is amazing that our bodies work so well as they do.

Most AI outsiders (including many philosophers, neuroscientists and psychologists) still cannot believe that so much complexity is required to achieve intelligence. They keep hoping that there is some golden short-cut that makes the hard work of today's vision systems, which use large numbers of algorithms and massive computation to filter, segment, and aggregate the incoming visual streams and match them against sophisticated top-down expectations stimuli, superfluous. Or they believe that language processing cannot be complicated because our brains do it effortlessly at great speeds. This echoes the trouble that biologists had in the 19th century to accept that all of life was ultimately implemented in chemistry. They initially could not believe that you needed vastly complex macromolecules (or even that such molecules could exist in nature), that you had highly complex metabolic cycles, and that a lot of these cycles were self-organised and steered by other molecules. It took hundred years for molecular biology to be accepted as the physical foundation for life and perhaps we just need a bit more time before the complexity of information processing is fully accepted and everybody is focused on what it could be rather than arguing whether the information processing view is appropriate in the study of intelligence.

3 Act I (1960s): Heuristic Search and Knowledge Representation

One reason why a mechanism for "general intelligence", similar to gravity or magnetism, may for ever remain elusive, is that almost any non-trivial problem contains a hidden combinatorial explosion. This was in fact the main lesson from the early game-playing programs. Often it is not possible to decide straight away how to explore a problem and so many avenues need to be searched. For example, almost every word in a sentence has multiple meanings and can be syntactically categorised in a multitude of ways. Hence a parse that fits with the complete sentence can only be found by elaborating different search paths, some possibly being a dead end. The growth of alternative search paths in language processing can be so high that the number of possible paths explodes. It is no longer possible to consider all of them. Such a combinatorial explosion shows up in visual perception, motor planning,

reasoning, expert problem solving, memory access, etc., in fact in every domain where intelligence comes into play. It is possible to combat combinatorial explosions with more computing power, but one of the key insights of AI in its first decade (the 60s) is that sooner or later the big numbers catch up on you and you run into a wall. So search can only be dealt with by the application of knowledge, specifically domain and task specific rules of thumb (heuristics) that quickly cut down search spaces and guide the problem solver towards reasonable solutions as fast as possible. The main goal of AI in the sixties hence became the development of techniques for organising the search process, expressing heuristics in evaluation functions or rules, and trying to learn heuristics as part of problem solving. For a while, the field of AI was even called Heuristic Search.

Towards the end of the 60s it became clear that not only the heuristics but also the way that information about the problem domain is represented plays a critical role, both to avoid search and to deal with the infinite variation and noise of the real world. So this led to the next breakthrough idea: A change in representation can mean the difference between finding a solution or getting swamped in irrelevant details. A host of new research questions was born and explored through an amazing variety of frameworks for knowledge representation and knowledge processing. Many ideas from logic could be pulled into AI and AI gave new challenges to logic, such as the problem of non-monotonic reasoning. Also the kinds of concepts that might be needed in expert problem solving or language understanding were researched and tested by building systems that used them. Initially there was some hope to find the ultimate conceptual primitives (Roger Schank's set of fourteen primitive universal concepts is one example (Schank, 1975)), but, failing that, large projects were started, such as Doug Lenat's CYC project, to do a massive analysis of human concepts and put them in a machine-usable form (Lenat, 1995).

Already in the 70s, the ideas and technologies for handling heuristic search and knowledge representation proved powerful enough to build real applications, as seen in the first wave of expert systems like MYCIN, DENDRAL, or PROSPECTOR. They were designed to mimic expert problem solving in domains like medicine or engineering, and led to a deeper analysis of problem solving in terms of knowledge level models (Steels and McDermott, 1993). Huge conferences, venture capital, spin-off companies, industrial exhibitions entered the scene and transformed AI forever. From the early 80s, applications became an integral part of AI. The emphasis on applications justified the research effort and proved that AI theory was on the right track, but it also meant that fundamental research slowed down. The initial emphasis on understanding intelligence gave way to a more pragmatic research agenda and a quest for useful applications. This trend is still going on today.

4 Act II (1980s): Neural Networks

The earliest efforts in AI (including the discussion at the Dartmouth conference) already had a kind of dual character. On the one hand, it was clearly necessary to come to grips with "symbolic" processing, and AI laboratories made major contributions to make computers and programming language sufficiently powerful to build and investigate hugely complex symbol systems. On the other hand,

sensory-motor control and pattern recognition, needed to relate symbolic structures to the real world, seem to require “subsymbolic” processing. The gap between the symbolic world and the physical world had to be bridged somewhere. Early neural networks (like Rosenblatt’s Perceptron) had already shown that some of the processing needed for these aspects of intelligence is better based on the propagation of continuous signals in networks with varying weights and thresholds instead of the transformation of symbolic expressions. However that does not mean that the information processing view is abandoned. Neural networks are just as much ‘computational’ mechanisms as heuristic search or unification (a core step in logical inference systems).

So there have always been two sides of AI, both developing in parallel. Research and early successes in symbolic processing dominated perhaps the first two decades of AI research and application, but the balance shifted in the 80s, with a renewed interest in neural networks and other forms of biologically inspired computation such as genetic algorithms. Significant advances were made throughout the 80s by studying more complex networks with “hidden” layers, or networks in which signals flow back in the network becoming recurrent, thus making it possible to deal with temporal structures (Elman, 1991). These advances lead to a steady stream of new applications in signal processing and pattern recognition, new technologies to make neural processing fast enough, and a huge impact on neuroscience, which finally began to adopt the information processing view in a serious way.

The field of computational neuroscience is now well established and it is studying what information processing is actually carried out by natural brains as well as offering new ideas to AI about what mechanisms might be needed for intelligence. The renaissance of neural network research does not make earlier work on “symbolic” AI irrelevant. Symbolic techniques remain the most adapted for studying the conceptually oriented aspects of intelligence, as in language processing or expert problem solving, and so far no adequate neural models have been proposed for language understanding, planning, or other areas in which symbolic AI excels. On the other hand, neural network techniques have proven their worth in the grounding of categories in the world or in the smooth interfacing of behavior with sensory stimuli. Intelligence is a big elephant and there is often a tendency by researchers to take the part they happen to focus on as the total. Obviously we need to avoid this trap.

5 Act III (1990s): Embodiment and Multi-agent Systems

But whatever the application successes of symbolic AI and neural networks, there was a growing feeling towards the end of the 80s that some fundamental things were missing, and two important new movements began to emerge. The first one rediscovered the body and the environment as major causal forces in the shaping of intelligent behavior (Steels and Brooks, 1994). Instead of doing very complex calculations for motor control, it was argued that it might just as well be possible to exploit the physical properties of materials and the agent-environment interaction to get smooth real-time behavior. Instead of trying to build complex ‘symbolic’ world models, which are difficult to extract reliably from real-world signals anyway, it might be possible to set up simple reactive behaviors and exploit the resulting “emergent” behavior. Ideas like this became the main dogmas of the

“behavior-based approach to AI”, which swept through leading AI laboratories in the early 90s (Pfeifer and Scheier, 2004).

Suddenly researchers started to build animal-like robots again, reminiscent of the cybernetics research that pre-dated the birth of AI in the 50s. Intense interactions started with biology, particularly ethology and evolutionary and developmental biology, and AI researchers were instrumental in helping to found the new field of Artificial Life (Langton, 1989). This wave of activity led again to the development of new software and hardware tools, new fundamental insights, and the first generation of rather astonishing animal-like robots. The amazing performance in the Robocup challenge testifies that the time was ripe to build real world robots from a behavior-based point of view. The late 90s even saw the first humanoids such as the Honda Asimo and the Sony QRIO. But all these robots are lacking any kind of “symbolic” intelligence (required for planning or natural language dialogue for example), which suggests that embodiment and neural-like dynamics in themselves are not enough to achieve cognition.

A second movement re-discovered that intelligence seldom arises in isolation. Animals and humans live in groups in which common knowledge and communication systems emerge through collective activities. Problems are solved by cooperation with others. The 90s saw the birth of the multi-agent approach to AI, which focused on how intelligence could be distributed over groups of co-operating entities and how the intelligence of individuals could be the outcome of situated interactions with others. Once again, a multitude of theoretical frameworks, programming paradigms, and formal tools sprung up to deal with the many difficult issues related to multi-agent systems and applications found a niche in software agents for the rapidly expanding Internet and for computer games. AI now established intense interactions with sociologists and anthropologists. It contributed to these fields by providing sophisticated agent-based modeling tools and got in turn inspired by integrating a social view on intelligence (Wooldridge, 2002).

6 Act IV (2000s): Semiotic Dynamics

So what is going on in AI right now? A lot of things of course and it will only become clear in retrospect what development has given the deepest long-term impact. It is obvious for example that there is at the moment a very strong trend towards statistical processing, which has shown enormous application potential in such areas as natural language processing (which is now almost entirely statistically based), web-related information retrieval, and robotics (Thrun, Burgard, and Fox, 2005). This trend is successful thanks to the availability of huge data sources and new techniques in machine learning and statistical inference. However my personal subjective choice goes to another line of current research which I believe holds great promise.

We have seen the pendulum swinging in the sixties and seventies towards knowledge representation and cognitive intelligence, counteracted by the pendulum swinging back in the other direction towards dynamics and embodiment in the eighties and nineties. Although symbolic applications of AI (as now heavily used in search engines) continue, basic AI research clearly moved away almost entirely from conceptual thinking and language. So I think this needs to be corrected again. I personally believe that the most exciting question at the moment is to see how

grounded symbolic systems can emerge in communicative interactions between embodied agents. Whereas early symbolic AI no doubt overemphasised the importance of symbols, more recent embodied AI simply ignores it, throwing away the baby with the bathwater. A more balanced view is feasible and necessary. Symbolic intelligence is the hallmark of human intelligence (some biologists have called us the “Symbolic Species” (Deacon, 1998)) and we must try to understand how intelligence that creates and builds further on symbols is possible.

In the earliest AI research (and indeed in cognitive science in general), there has always been the tacit assumption that categories and concepts for structuring the world, as required in problem solving or language communication, are static, universally shared, and hence definable a priori. This led to the search for conceptual primitives, large-scale ontologies, attempts to capture common-sense knowledge once and for all as illustrated by Lenat's CYC project or the semantic web. More recent, statistically based machine learning techniques similarly assume that there is a (static) conceptual or linguistic system out there which can be acquired, and then used without further adaptation. However observations of natural dialogue and human development show that human language and conceptualisation is constantly on the move and gets invented and aligned on the fly. Conceptualisation is often strongly shaped and reshaped by language. Conceptualisations are invented and imposed on the world based on individual histories of situated interactions, and they become shared by a dynamic negotiation process in joint collaborative tasks, including communication (Pickering and Garrod, 2004). From this point of view, ontologies and language are seen as complex adaptive systems in constant flux. Language nor the meanings expressed by language are based on a static set of conventions that can be induced statistically from language data.

Today this viewpoint is being explored under the label of semiotic dynamics (Steels, 2006). Semiotic dynamics studies how ontologies and symbol systems may emerge in a group of agents and by what mechanisms they may continue to evolve and complexify. The study of semiotic dynamics uses similar tools and techniques as other research in the social sciences concerned with opinion dynamics, collective economical decision making, etc. (Axelrod, 2006). There have already been some initial, very firm results. Groups of robots have been shown capable of self-organising a communication system with natural language like properties, including grammatical structure (Steels, 2003). These multi-agent experiments integrate insights from neural networks and embodied AI to achieve the grounding of language in sensori-motor interaction, but they also rely on sophisticated symbolic processing techniques that were developed in AI in the 60s and 70s. The current evolution of the web towards social tagging and collective knowledge development makes this research even more relevant and tangible, because when we collect data from collective human semiotic behavior that arises in these collective systems we see the same sort of semiotic dynamics as in artificial systems. Semiotic dynamics has made a deep contact with statistical physics and complex systems science. These fields have the tools to investigate the self-organising behaviors of large systems, for example by proving that a certain set of rules will lead to a globally shared lexicon or by showing how systems scale with the size of the population.

By researching how symbol systems can be self-constructed by embodied agents, AI is tackling objections raised by philosophers like Searle (in his Chinese Room

parable) or biologists like Edelman. It is true that in most AI systems of the past symbols and the conceptualisation of reality they imply were constructed and coded by designers or inferred from human use. But these scholars concluded (wrongly) that it would never be possible for physically embodied information processing systems to establish and handle symbols autonomously, whereas it is now clear that they can.

7 Epilogue: Peeking into the Future

AI is not that different from other sciences in that it is pushed forward by new technologies that enable more powerful experiments, by new challenges coming from novel applications, or by new advances in other sciences, including mathematics, which can lead to new mechanisms or new predictive tools for understanding and hence exploiting existing mechanisms. Let us see how these three aspects could influence AI in the coming decades.

+ The push from technology: In the 80s there was a high hope that new computer architectures (such as the Connection Machine (Hillis, 1986)) could lead to large advances in AI. This has not happened because standard architectures and mass-produced machines became so cheap and progressively so powerful that they overtook more exotic computational ideas. However there are strong indications that new generations of parallel computers are going to become available soon. They will make experiments possible that are too slow today or can only be done by a very limited number of people. If we want to exploit the promise of massively parallel computing, we will need to come up with novel computing paradigms, exploiting metaphors from biology or chemistry rather than from logic and mathematics. Research in membrane computing, molecular computing, amorphous computing, etc. is showing the way. Not only computer technology but also mechanical engineering and materials science are evolving rapidly and this will make it possible to build completely new kinds of artificial systems, possibly on a very small scale. These technologies are already being explored in embodied AI research and they will surely push the boundaries of autonomous systems, as well as clarify the role of information processing in achieving intelligence.

+ The push from applications: We can observe today a global trend towards collective phenomena in the information processing world. The exponential growth of websites, blogs, wikis, peer-to-peer sharing systems, folksonomies, wikipedias, etc., makes it clear that our densely interconnected world is leading to a whole new way of knowledge production and knowledge communication. Centralised knowledge production and control is swept away by powerful waves of collective tagging, collective encyclopedias, collective news media. Dense interconnectivity and the knowledge dynamics it supports can only increase, when more and more physical devices get substantial computing power and become networked so that everyone can access and change information from anywhere. Many information technologies are participating in the invention of new tools for supporting these remarkable developments, and AI can do so as well, even though some of the older ideas will no longer work, such as the reliance on the individual knowledge of an “expert” which used to be the basis for expert systems, or the idea that it is possible to define logically a universal ontology which can be imposed in a top-down fashion on the web.

A second application area which I believe is of extreme importance and to which future AI should contribute with all its force concerns the retooling of our economies so that they become sustainable. Unless we act today, ecological catastrophes are unavoidable in the not so distant future. AI is highly relevant to achieve this goal but the work has hardly started. AI can contribute with methods and techniques for measuring, interpreting, and assessing pollution and stress to our environments, for tracing and tracking the use of natural resources to organise complete healthy production and consumption cycles, for tracking energy use and managing it in a dynamical adaptive way, and for predicting and visualising the outcomes of current ecological trends so that people become more aware of the urgency of action. Biologically-inspired research results from embodied agents research are in my opinion highly relevant to orchestrate the adaptive management and optimisation of energy resources in micro-grids with local distributed production and peer-to-peer energy exchange.

+ Push from other fields: Future steps forward in theory formation are difficult to predict. Evolutionary network theory (Strogatz, 2001), which only developed in the last decade, has proven to be highly relevant for developing the theory of semiotic dynamics, and much remains to be discovered in non-linear dynamical systems theory, that may have a direct bearing on the foundations of embodied intelligence, neural networks, and agent-environment interaction, as well as on the understanding of the collective dynamics of large groups of agents which self-organise communication systems or other forms of interaction. The conceptual frameworks that have been used so far in neuroscience (basically networks of a few simplistic neurons) are much too simple to start tackling the enormous complexity of real brains so there is a growing demand from neuroscience for the kind of more sophisticated conceptual frameworks that have been common in AI. But I predict that most of the potentially useful interaction in the future will come from interactions with evolutionary biology. Biology has a great tradition of conceptual thinking and its most brilliant thinkers have now a firm grasp of the information processing paradigm (Maynard Smith, 2000). What we need to do is shift our attention away from looking at the end product of intelligence and trying to engineer that towards an understanding how intelligence dynamically arises both within a single individual and within our species. Research on the question of the origins of language is one example in this direction.

AI research, due its strong ties to engineering, has always been very pragmatic, directed by societal and industrial needs. This is in principle a good thing and keeps the field relevant. But I personally feel that we must put a much greater effort into basic research again and into distilling the lessons and principles learned from this pragmatic bottom-up AI research to translate them both in a systematic theory and into a format accessible to other sciences and the informed layman, similar to the way biologist have been able to philosophise and communicate their fundamental research results to a very broad public. A good recent example in this direction is Pfeifer and Bongard (2007). At the same time, AI is an experimental science, because that is the only way that claims and arguments from a design stance can be tested and validated. AI is at its best when doing carefully set up experiments to test the mechanisms that may play a role in intelligence. I believe that these experiments must be set up and communicated like scientific experiments (instead of “applications”), because it is only through that route that we will ever get a systematic theory.

In my view, the coming fifty years of AI look extremely bright. We will continue to make major discoveries and tackle challenges important to society. It is my personal experience that other fields are highly interested in our research results, but only if we stick to the design stance and not pretend to be modeling the human brain or human intelligence. We should get into a dialogue with the fields who have this as their major aim (cognitive science, neuroscience, etc.) but as equal partners bringing our own insights to the table.

References

1. Axelrod, R.: Agent-based modeling as a bridge between disciplines. In: Judd, K.L., Tesfatsion, L. (eds.) *Handbook of Computational Economics: Agent-Based Computational Economics*, North-Holland (2006)
2. Deacon, T.: *The Symbolic Species: The Co-evolution of Language and the Brain*. Norton Pub., New York (1998)
3. Dreyfus, H.: *What Computers Can't Do: The Limits of Artificial Intelligence*. Harpers and Row, New York (1972)
4. Elman, J.: Distributed representations: Simple recurrent networks, and grammatical structure. *Machine Learning* 7(2-3), 195–226 (1991)
5. Hillis, D.: *The Connection Machine*. MIT Press, Cambridge, MA (1986)
6. Langton, C. (ed.): *Artificial Life*. Addison-Wesley, Reading (1989)
7. Lenat, B.: Cyc: A large-scale investment in knowledge infrastructure. *Comm. ACM* 38(11), 33–38 (1995)
8. Maynard-Smith, J.: The concept of information in biology. *Philosophy of Science* 67, 177–194 (2000)
9. McCarthy, J., Minsky, M., Rochester, N., Shannon, C.: A proposal for the Dartmouth summer research project on artificial intelligence (1955), <http://www.formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
10. Pfeifer, R., Scheier, C.: *Understanding Intelligence*. MIT Press, Cambridge, MA (1999)
11. Pfeifer, R., Bongard, J.: *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, Cambridge, MA (2007)
12. Pickering, M.J., Garrod, S.: Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences* 27, 169–225 (2004)
13. Schank, R.: *Conceptual Information Processing*. Elsevier, Amsterdam (1975)
14. Steels, L., McDermott, J. (eds.): *The Knowledge Level in Expert Systems. Conversations and Commentary*. Academic Press, Boston (1993)
15. Steels, L., Brooks, R. (eds.): *The “artificial life” route to “artificial intelligence”*. Building Situated Embodied Agents. Lawrence Erlbaum Ass., New Haven (1994)
16. Steels, L.: Evolving grounded communication for robots. *Trends in Cognitive Sciences* 7(7), 308–312 (2003)
17. Steels, L.: Semiotic dynamics for embodied agents. *IEEE Intelligent Systems* 5(6), 32–38 (2006)
18. Strogatz, S.: Exploring complex networks. *Nature* 410, 268–276 (2001)
19. Thrun, S., Burgard, W., Fox, D.: *Probabilistic Robotics*. MIT Press, Cambridge, MA (2005)
20. Winograd, T.: *Understanding Natural Language*. Academic Press, London (1972)
21. Wooldridge, M.: *An Introduction to Multiagent Systems*. John Wiley & Sons, Chichester (2002)