

Current and Future Trends in AI

Kemal Delic
Hewlett-Packard Co.
Grenoble, France
kemal.delic@hp.com

Jeff Riley
Hewlett-Packard Co.
Melbourne, Australia
jeff.riley@hp.com

Abstract — During the past 70+ years of research and development in the domain of Artificial Intelligence (AI) we observe three principal, historical waves: embryonic, embedded and embodied AI. As the first two waves have demonstrated huge potential to seed new technologies and provide tangible business results, we describe likely developments of embodied AI in the next 25-35 years. We postulate that the famous Turing Test was a noble goal for AI scientists, making key, historical inroads – while we believe that Biological Systems Intelligence and the Insect/Swarm Intelligence analogy/mimicry, though largely disregarded, represents the key to further developments. We describe briefly the key lines of past and ongoing research, and outline likely future developments in this remarkable field.

Keywords—*Artificial Intelligence; Biological Systems Intelligence, Swarm Intelligence, Insect Intelligence*

I. Introduction

With computers of some type in almost every hand, handbag, briefcase or pocket, on desktops and in trains, planes, and automobiles everywhere, the possibility that these machines could communicate and co-operate to solve problems and/or achieve goals is more realistic now than it was at any time in the past.

Since electronic componentry continues to shrink in size, available computing power continues to grow and become more ubiquitous. With the footprint of massive computing power becoming very small, the possibility of massively parallel, distributed artificial intelligence, arises. The availability of very inexpensive, very small, very powerful processors gives rise to many possibilities for distributed, parallel processing.

As miniaturization continues, systems could be constructed consisting of millions of nanobots that form a co-operative, distributed swarm intelligence that has an innate resilience because of its massive redundancy – a large swarm could lose a few agents without significant loss of functionality.

Combining such swarms with evolutionary computation – genetic algorithms, genetic programming – artificial evolution gives rise to the possibility that populations of nanobots could be evolved for specific purposes. Although natural evolution is not teleological, artificial evolution strives to optimize populations to solutions of one or more goals – self-optimization for purpose is the hallmark of artificial evolution.

II. Elusive Goals and Real Achievements

While we still don't yet have truly intelligent machines, the field of Artificial Intelligence has moved from the hyperbolic expectations of the past to what is now an established research practice, and has, in some sense at least, become more mainstream over the past few decades. Perhaps just through greater interactions with computers over the past 30 or so years, the general population has come to accept – even expect – intelligent behaviour from computer systems and machines with embedded computers.

Nobody is surprised to find autopilots that intelligently fly aeroplanes, computers that can park a car, avoid collisions, and find the shortest route to our destinations while avoiding traffic congestion, personalized searches and recommendations on websites, or household appliances that take much of the thinking out of mundane chores – these are all embedded AI systems taken for granted by most people today. Much of modern medicine would not be possible without AI – from the development of new drugs, to computer aided imaging through to computer guided surgery, all of which have heavy reliance of artificial intelligence techniques.

Children of teenage age and younger wouldn't understand a world devoid of (semi-) intelligent machines. A clear example is the general trend of the population relying on services such as Google. The Millenniums generation knows much more than previous generations – or at least has access to much more information and knowledge than any generation has in the past - but in a drastically shallower manner. And all this in real-time in real-world circumstances – the social impact of Google 'search' is huge and profound.

Artificial Intelligence encompasses many areas – vision, language, cognition, movement and dexterity are just a few. Asking ten people to describe “artificial intelligence” would likely result in at least ten different descriptions. Certainly one of the loftier goals of AI has been to construct a machine capable of displaying intelligence that emulates – even rivals – human intelligence. We're not there yet, and while great progress has been made in several key areas – stunning advances in computer hardware have enabled the partial realization of some prescient science fiction writers and AI pioneers – we're not really even close.

In 2005 Rodney Brooks cited the following as worthy challenges for AI [1]:

- the generic visual object recognition capabilities of a two year old child
- the manual dexterity of a six year old child
- the social interaction and language capabilities of a ten year old child

These remain worthy challenges even today. We could add many items to Brooks’ list – for example, the ability to deal with unexpected or novel situations, the ability to display “common sense”, the ability to understand – or display – emotions, are all elusive goals and remain worthy challenges for the future AI research.

III. AI Creating Seeds for Major Technologies

We will not try to give exhaustive arguments, much less profound analysis, of the technological and economical impacts of AI research - one can easily discern traces of technologies either directly based on, or largely influenced by, AI research. It is instructive to notice that the majority of Chief Technology Officers in contemporary companies have deep roots in AI. Yet another anecdotal indicator lies in the fact that the recently launched, web-based courses on AI from Stanford University, CA [2] attracted more than 130,000 students from over 190 countries. This spawned yet another development of MOOCs (Massive Open Online Courses) – potentially changing the entire education industry and launching several start-ups.

TABLE I. KEY AI-TECHNOLOGY DOMAINS

	DOMAIN		
	<i>Technology</i>	<i>Names [e.g. pub]</i>	<i>E.g. App</i>
1940	Embryonic Theories of AI, Control Systems and Cybernetics	Bush [3] Polya [4] Turing [5] Ashby [6, 7] Wiener [8]	Trial & Error Search, Self-Organization
1950	Toy Prototypes – Language Systems – Knowledge representation	Minsky [9] Simon [10] Nilsson [11]	GPS
1960	Neural Networks – Heuristics - Language Understanding – Adaptation – Control Systems	McCulloch & Pitts[12] Tsytkin [13, 14]	HearSay
1970	Expert Systems – Q/A Systems	Waterman [15]	Dendral-Mycin
1980	Machine Learning – Robotics	Michie [16] Brooks [17]	MENACE iRobot
1990	Autonomic Computing – Games Knowledge Based Systems	Hayes-Roth [18]	Watson

	DOMAIN		
	<i>Technology</i>	<i>Names [e.g. pub]</i>	<i>E.g. App</i>
2000	Voice, Image and Pattern Recognition	Stuart & Norvig [19]	Google Siri
2010	Biological Systems Intelligence, Swarm Intelligence – air/sea/undersea Drones	Bonabeau [20] Dorigo [21]	Swarm-Bots

To roughly outline the developments of AI during the last 70 years or so, we indicate some major technology domains in a decade-span time-frame and mention representative researchers and examples as the reference (Table 1).

Some recent publicly described grand successes of AI are as follows:

- Google has used 1000 machines engaging 16,000 cores to create a neural-network of 1 billion nodes that learned, after being trained on 10 million pictures, to recognize a cat in previously unseen pictures [22]. The reported accuracy is greater than 75% and constantly improving. We can imagine how this experiment can be expanded into several, not yet explored directions and data collections.
- Self-driving and self-parking cars, Mars rover, and flying and sea-diving robots all have technologies inspired by AI research spreading into the commercial world and mass consumer products.
- Machines driven by AI have become world chess and go champions, while others are able to beat humans in top TV quizzes by depth and scale of knowledge and responding quicker. Entirely new types, and scale, of services are envisioned based on this.
- Machine translation has reached unprecedented quality and scale of use – and what’s more is generally available free of charge. AI-based programs are already able to write articles which only a very careful eye is able to recognize as a machine-created article.

There is another side to these enthusiastically embraced AI technologies of course. We are witnessing the destruction of millions of old-fashioned jobs, and the slow creation of new jobs in new technology areas to replace them. We note that historically, every significant technological transition has exhibited this phenomenon. For example, the move from an agricultural society to an industrial one destroyed the major proportion of land-based jobs (ploughing, planting, harvesting etc.), but created new jobs in new technological areas, and moved the population into big cities. A similar process may happen now: people will become more mobile and will change professions several times during their lifetime. Technology has the power of creating new and destroying the old – a necessary renewal process.

We must be vigilant though that the proliferation of technology, and more particularly intelligent technology brought about by AI research, does not cause the atrophying of

the knowledge and skills of the population. Too often people turn to “cookie-cutter” solutions because they are convenient (even redefining their problem to fit an available solution) instead of continuing to develop new knowledge and, just as importantly, maintaining existing knowledge.

iv. Three Phases of AI Revisited

The early practical interest of AI practitioners was in search and problem solving [11, 23], and we can assume that the importance of this can be judged only now that Google is our daily reality. In the early part of the second half of the 20th century, Knowledge Representation and Reasoning (KRR) was the principal challenge (often articulated as DARPA research). The logical consequence of the interest in KRR was the exploration of (human) learning aspects, reinforcing the role of knowledge representation, capture, mining and reuse. We have been intimately involved in commercial developments of knowledge-based systems supporting help-desk operations dating as far back as some 20 years ago [24, 25, 26]. In an abstract sense, the World Wide Web can be considered a huge knowledge-base in standard format and usage protocol. This is underlined by the fact that Google considers itself a knowledge-based company today.

The advent of the internet and World Wide Web produced a paradigm shift across many entire domains: e.g. advertising, retailing, entertainment – both in the delivery and ticketing/booking – the travel industry, education, gaming – both gambling and multi-user online fantasy type games – publishing etc. At the same time something very important happened in the world of finance – High Frequency Trading. The vast majority of all stock trades are now conducted online, and a majority of those are performed by machines run by quantitative analysts, generally using AI techniques and algorithms which are the best kept secrets of banks and hedge-funds. Over the past 15-20 years, embedded AI was the principal technology in the finance industry driving credit application evaluations, fraud detection, stock-exchange trades and predictive analytics.

Many embedded AI systems draw heavily on control systems theory. Control systems are ubiquitous and are found in many applications and areas, including robotics and mechatronics, autonomous machines and vehicles for medicine, manufacturing and military uses, etc. It’s hard to imagine an embedded intelligent system that doesn’t incorporate a control system of some sort. Some early work on control systems, adaptation and self-organizing systems was done during the 1940s [6, 7, 8], followed by important work by Tsykpin on adaptation and learning in automatic systems in the 1960s [13, 14].

The introduction of digital computers allowed for much greater complexity in control systems. The development of Fuzzy Logic and Fuzzy Systems [27, 28] provided a means of building systems that dealt with complex control problems in terms that human operators could understand. Fuzzy Control Systems can be found in many varied applications including autonomous vehicle control, face recognition, handwriting recognition, domestic appliances such as vacuum cleaners, dishwashers and washing machines.

Over the past decade or so we have observed the rise of air, land and sea drones – unmanned autonomous vehicles (UAVs) - mainly in the military arena, but also with clear, important civil and commercial uses. UAVs are a good illustration of embodied AI and control systems that was preceded by self-driving cars and automatic NASA probes exploring distant planets and moons. They all illustrate a potential which today has spin-off benefits even in the domestic world - house and window cleaning robots, automated repair, home surveillance, etc.

Robotic games are a staple testbed of the AI community and have been used over several decades to test hypotheses of emergent and cooperative behaviour, and to demonstrate behaviours such as swarming and flocking [29]. Through the experience of robotic game-playing AI practitioners have come to the realization that real intelligence is likely to emerge only when the agent or robot interacts with its environment – in other words, only when the intelligence is embodied [29, 30, 31, 32].

v. Conclusion & Outlook: Roadmap for the Future

We are aware of the benefits of AI research for the development of various technologies, but we are also well aware of potential unexpected uses for military and/or criminal activity. However, historically speaking, the development of technology initially for military purposes is usually followed by the transformation of that technology for commercial and domestic uses. It is a path well travelled: even Leonardo Da Vinci had some brilliant ideas initially embodied in military machines and projects.

We see the future of AI as intelligence embodied in swarming robots – perhaps even nanobots - with higher intelligence emerging from the cooperation, competition and coordination of the omnipresent robotic swarms. AI research has already created precedents with the research of distributed AI – well before other developments in computing enabled the creation of the real-world systems and prototypes.

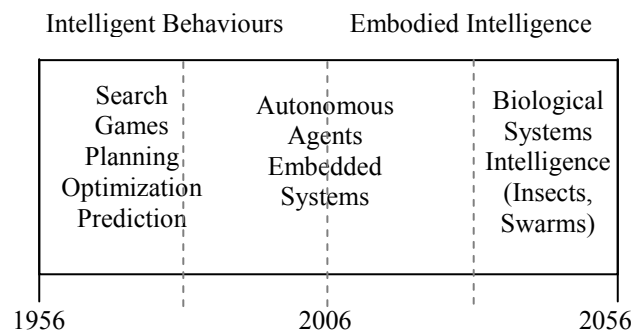


Fig. 1. 100 Years of AI (Adapted from K. Delic & U. Dayal [33])

Looking at a very long time frame – one century – we are putting our bets on only three selected domains for which we believe AI will play a positive role in preserving The Earth for

future generations (Fig. 1). Futurologists and Technologists of NASA have also made their bets public [34]. Consequently, looking into long-aim research programs of the EU and US, we see that significant investments will be made in areas we select here as *very* promising:

1. **Biological Systems Intelligence and Swarm Robotics** -- helping people during catastrophes (tsunami, hurricanes) and major industrial accidents (Deep Horizon, Fukushima) and war tragedies.

2. **Preserving Earth** – monitoring ecosystem and aiming to reduce human impact on Global Climate Change via SWARM Surveillance of Oceans, Seas and Lands.

3. **Global Bayesian Computer** -- emulating global brain functionality, sensing overall mood of population via mobile-phone swarms, reasoning about actual world problems and helping fight critical diseases - able to capture, crystallize and express emotions.

It is likely that those long-term programs will seed enough ideas to attract investment circles to put their money into it - for the benefit of life preservation on Planet Earth. This will be the ultimate and future noble goal of the century long research in AI sciences.

VI. Acknowledgement

We would like to thank anonymous reviewers for providing useful suggestions and extend our thanks to Prof Dr Lj. Vlacic for kindly pointing at Control Systems importance in robotics research.

References

- [1] R. Brooks, "AI - The next 25 years", Matthew Stone and Haym Hirsh, Eds., AI Magazine, 26(4), Winter 2005, pp.85-97.
- [2] P. Norvig, S. Thrun, "Introduction to artificial intelligence", www.ai-class.com, 2011.
- [3] V. Bush, "As we may think", Atlantic Monthly, July 1945.
- [4] G. Polya, "How to solve it", Princeton University Press, 1945.
- [5] A. Turing, "Computing machinery and intelligence", Mind 1950.
- [6] W. R. Ashby, "The physical origin of adaptation by trial and error", Journal of General Psychology (32), 1945, pp.13-25.
- [7] W. R. Ashby, "Principles of the self-organizing dynamic system", Journal of General Psychology (37), 1947, pp.125-128.
- [8] N. Wiener, "Cybernetics, or, Control and Communication in the Animal and the Machine", MIT Press, Cambridge MA 1948.
- [9] M. Minsky, "Steps toward artificial intelligence", New York, 1961.
- [10] H. Simon, "The sciences of the artificial", MIT Press, Cambridge, Mass, 1969.
- [11] N. Nilsson, "Problem-solving methods in artificial intelligence", New York: McGraw-Hill, 1971.
- [12] W. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity", Bulletin of Mathematical Biophysics (7), 1943, pp.115-133.
- [13] Y. Z. Tsytkin, "Adaptation, learning and self-learning in automatic systems", Automation and Remote Control 27(1), 1966, pp.23-61.
- [14] Y. Z. Tsytkin, "Adaptation and Learning in Automatic systems", Academic Press, 1971.
- [15] D. Waterman, "A guide to expert systems", Pearson 1986.
- [16] D. Michie, "Steps toward robot intelligence", Proceedings of the Royal Institute (57), 1985, pp.151-165.
- [17] R. Brooks, "A robust layered control system for a mobile robot", IEEE Journal of Robotics and Automation 2(1), 1986, pp.14-23.
- [18] F. Hayes-Roth, J. E. Davidson, et al. (1991). "Frameworks for developing intelligent systems." IEEE Expert 6(3): 30-40.
- [19] R. Stuart, P. Norvig, "Artificial intelligence: a modern approach", Prentice Hall 1995.
- [20] E. Bonabeau, D. Corne, R. Poli, "Swarm intelligence: the state of the art special issue of natural computing", Natural Computing 9(3), 2010, pp.655-657.
- [21] M. Dorigo, "Swarm-Bots and Swarmanoid: Two Experiments in Embodied Swarm Intelligence", WI-IAT'09: Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, September 2009.
- [22] L. Clark, "Google's artificial brain learns to find cat videos", Wired UK, www.wired.com, June 2012.
- [23] N. Nilsson, "The quest for artificial intelligence: a history of ideas and achievements", Cambridge University Press, 2010.
- [24] K. Delic, D. Lahaix, "Knowledge harvesting, articulation and delivery", HP Journal, May 1998.
- [25] J. Riley, K. Delic, "Augmenting information retrieval by knowledge infusion", IJCAI99.
- [26] K. Delic, J. Riley, "Enterprise knowledge clouds: architecture and technologies", Springer Verlag, 2010
- [27] L. Zadeh, "Fuzzy Sets", Journal of Information and Control (8), 1965.
- [28] B. Kosko, "Fuzzy Thinking: The New Science of Fuzzy Logic", Hyperion, 1993.
- [29] J. Riley, "Evolving fuzzy rules for goal-scoring behaviour in a robot soccer environment", PhD Thesis, School of Computer Science and Information Technology, RMIT University, Melbourne, Australia, 2005.
- [30] R. Brooks, "Integrated systems based on behaviors", SIGART Bulletin, 2(4), August 1991, pp.46-50.
- [31] R. Brooks, "Intelligence without representation", Artificial Intelligence Journal (47), 1991, pp.139-159.
- [32] A. Jain, P. Jain, J. Singh, "Artificial intelligence: simulated embodied agents", IJCSN International Journal of Computer Science and Network, 2(3), June 2013, pp.20-24.
- [33] K. Delic, U. Dayal, "AI-re-emerging as research in complex systems", AI-50, Monte Verita, Switzerland, 2006.
- [34] E. Vassev, R. Sterritt, C. Rouff, M. Hinchey, "Swarm technology at NASA: building resilient systems", IT Professional, 14(2), March 2012, pp.36-42.