



## ARTIFICIAL INTELLIGENCE: THE FIRST SEVEN DECADES

Vladimir Srdanović

MISANU

e-mail: v.t.srdanovic@gmail.com

### Abstract:

The paper gives a brief overview of developments in the field of Artificial Intelligence over the past seven decades. It reflects on some issues characterising the relationship between knowledge and intelligence, as well as some phenomena related to the origin of intelligence. It also gives a short account of the beginning of AI activities in Serbia.

**Keywords:** Artificial Intelligence, knowledge, learning, AI in Serbia

### 1. Introduction

In 1955, a group of mathematicians and computer scientists, lead by John McCarthy, wrote a Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, which would explore the hypothesis that "every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it." [add reference]. The workshop took place over the summer of 1956 at Dartmouth College, where McCarthy was a young professor at the time. The legendary deliberations at the workshop have set the foundations for the field of AI.

### 2. Knowledge and Learning

Artificial Intelligence is a scientific specialty that attempts to understand the nature of intelligent action and to construct systems that are capable of intelligent action.

There is little disagreement as to whether mind and intelligence are phenomena primarily associated with humans. It is difficult, however, to come up with a single aspect of these phenomena that is exclusively human. For any particular characterization of such an aspect, it is also possible to associate it (although often at a rudimentary level) with organisms and/or systems other than humans.

Darwin, in his *The Descent of Man* (1871) had observed that "[t]here is no fundamental difference between man and the higher mammals in their mental faculties and that all the differences are differences of degree, not of kind" [add reference]. We know today that not only higher mammals but also other animals are capable of intelligent action, communicating relatively complex information, and learning. Some examples include certain birds, and even distant evolutionary "relatives" such as octopuses (which achieve these capabilities by utilizing a radically different "solution" from their neural architecture).

The amount of information necessary for survival is enormous. One way through which essential information is communicated among living organisms is their genetic code. From one generation to another, genes carry information relevant for duplicating typical characteristics a species has acquired through the process of evolution. More than that, some of the most typical responses a species exhibits in its environment, the way it deals with uncertainties it is faced with, or even in cases when it learns from its environment, all of these are guided by information that is also genetically coded.

It is obvious that effective alternate ways of acquiring this information were needed. Learning certainly is a way of doing so. This is particularly true for species that are not at the top of the food chain. The more intelligent a species is, the more it relies on its intelligence to acquire new knowledge through the learning process. The time youngsters spend with adults is typically the longest in the case of primates. It is known that a female chimpanzee would not breast-feed her offspring unless she has

spent critical years with her mother observing her breast-feeding her younger siblings. Indeed, it is striking that such a crucial function has been "dropped out" from the list of those genetically coded and shifted to those that have to be learned.

Still, another way of communicating information is via the environment. Although perhaps less frequent in the animal world, this way of communicating information is also not exclusively human. Humans have done the most in using the environment to communicate information. The earliest wall paintings in caves are dated at about fifteen thousand years ago. It took another several thousand years for picture writing to appear, developing further into the symbols and alphabets known today. All this development took place in a very short time by the standards of biological evolution.

The development of alphabets and languages allowed for practically arbitrary information to be encoded. The main point is that external media was used as a vehicle to symbolically represent knowledge. Although there are various ways to achieve that physically and different media can be used (stone, paper, magnetic, etc.), for most practical purposes knowledge represented as such was made available permanently. Moreover, this knowledge was now made available to more people, circumventing temporal and spatial limitations. Finally, it did not allow only for isolated pieces of knowledge to be accessed over and over again. The real advantage was that knowledge was made available to others for further refinement; it was now possible to accumulate it, build upon it, and increment the body of human knowledge about the universe that surrounds us [1]. The scope of the information revolution we are going through has been spelled out by Minsky: "Can you imagine that they used to have libraries where the books didn't talk to each other?" [add reference].

### 3. Symbolists vs Connectionists

The study of humans (and other intelligent systems), the types of intelligent actions they perform and the mechanisms they use to perform them resulted in the following propositions, as summarized by Newell [2]: the capability of creation and manipulation of symbol structures; use of large amounts of highly diverse knowledge; the fundamental response of a system to uncertainty is to create a space within which a resolution to that uncertainty must lie, and to search that space; the act of perception requires knowledge equal in breadth and extent to that involved in subsequent use of the resulting representation; the control of behavior towards the end can be obtained by the use of goals.

As an extension, Newell and Simon in their seminal Turing Award paper [3] proposed the Physical Symbol System Hypothesis: "A physical symbol system has the necessary and sufficient means for intelligent action."

These positions in AI research, based on high level symbolic (human readable) representations of problems, logic and search, are being referred to as **symbolic artificial intelligence**. Symbolic AI used tools such as logic programming, production rules and semantic nets and utilized them to develop applications like expert systems.

Acquiring domain knowledge in building expert systems was known, however, to be a hard and time-consuming process. Difficulties that experts encounter when asked to explicate their knowledge about the domain of their expertise has long been recognized by the researchers in AI and considered as the main contributor to the so-called "bottleneck" problem, first spelled out by Feigenbaum [4]. However, as Donald Michie [5] has pointed out, if unable to "say how", experts can "show how" they solve problems in their domains.

The medical domain, unlike many others, provides us with a relatively homogeneous "record" of how the problems are solved within it. Large amounts of medical knowledge are implicitly contained in patient case histories.

As noted by Quinlan [6], experts can articulate the framework for their knowledge and systems based on induction can produce useful knowledge from such frameworks supplied with examples. The same approach can sometimes work even when there is no human expert.

BELART, Srdanović [7] shall be noted as an example of a system utilizing these principles. The system used expert knowledge to define structure for a specific medical domain. Next, data are collected and analyzed to devise strategies used in consultation processes. BELART demonstrated robustness and high diagnostic accuracy (over 90%) in several medical domains, while using small training data sets. The system is also capable of automatically deriving diagnostic criteria, thus qualifying as an Ultra-strong learning criterion as defined by Michie [8].

So far, the most famous (ongoing) project using symbolic AI is Doug Lenat's Cyc, that began in 1984 and aimed to encode common sense in a machine. By 2017, Lenat and his team added 1.5 million terms and 24.5 million rules to Cyc, while not achieving general intelligence [9].

From the onset of AI research, another approach, referred to as **connectionism**, was also applied. The idea was that an adequate system architecture can be developed by using tools like **artificial neural networks** that can explain and model mental phenomena. Such a system can then be used to learn from and solve problems in the real-world environment.

In the 1990s, increased computer power enabled the development of the so-called deep neural networks, and the respective approach called deep learning. The last two decades provided long-awaited input: large data sets. The booming World Wide Web, smartphones, web services like YouTube and Wikipedia have flooded the Internet with enormous amounts of images, videos, and texts ready to be processed. The result was enormous success and commercial use of deep learning programs in many domains. In fact, the success was so great that today the terms *AI* and *Deep Learning* are often used as synonyms.

However, the disadvantages like the opacity of the process, inadequate mental representations, brittleness, etc., [10], stand in the way of the connectionist approach being universally applied.

Although there are efforts to combine the two approaches in order to achieve general artificial intelligence, i.e., "Strong AI", this goal still seems far away. The gap in our understanding of how our biological "software" and "hardware" operate is still wide.

#### 4. The Beginnings of AI in Serbia

Interest among individual researchers in Serbia for various aspects of AI was evident in 1980s. By mid 1980s a number of research projects on AI and its applications were under way, getting together researchers from various schools of the University of Belgrade (Electrical Engineering, Mathematics, Organizational Sciences, Center for Multidisciplinary Studies), as well as some institutes (Mathematical, Pupin, Vinča).

In 1983, the Artificial Intelligence Group (YUGAI) of the Yugoslav Society for ETAN was founded. YUGAI was a member of the European Coordination Commission for Artificial Intelligence (ECCAI) until 1991.

The University of Belgrade first introduced the Graduate Program in Artificial Intelligence in 1989. The first undergraduate course on *Application of AI – Expert Systems* was offered by the School of Science and Mathematics at the University of Belgrade in 1991.

During the 1980s, a series of six International Summer Schools on Artificial Intelligence were organized by CAS (a subsidiary of ETAN) in Dubrovnik. Lecturers included some of the world's foremost authorities in AI. The schools were attended by a total of several hundred young researchers from Yugoslavia as well as other, mainly European, countries.

#### 5. Social Implications

AI has gone through several boom and bust cycles, later called "AI Winters". These cycles were driven as much by public hype and disappointment with AI's promises and failures, as with economic downturns resulting in reduced research funds.

Present enthusiasm in the industry and research is caused by the undisputable success of "Weak AI" systems, i.e., those that are developed for particular domains. The general public, however, seems to be convinced – for the first time – in what they perceive as the undisputable/inevitable success of Strong AI. They meet it with fear rather than enthusiasm.

Out of several potential risks, two stand out:

The risk of unemployment. Although past records show that technology tends to increase unemployment, older generations witnessed that armies of typists and telephone operators have "disappeared" over a relatively short period of time without causing too much turbulence in the job market. Perhaps the next big change in the job market will be caused by the "disappearance" of truck (and car?) drivers.

The risk of weaponized AI. Though this is already real, the deaths caused to humans by other humans are, unfortunately, higher in magnitude than those caused by machines and it is likely to stay that way.

## 6. Conclusions

Over the past seven decades of its existence, Artificial Intelligence has gone through several boom and bust cycles. Although the goal of achieving general intelligent action (strong AI) has not been reached, the successes of weak AI systems has increased research funding. While it has brought enthusiasm to AI researchers, it seems to have raised fear among the general public. After so many years, the view of the future of AI, given by Newell in [11], still seems to be very fresh:

"We should, by the way, be prepared for some radical, and perhaps surprising, transformations of the disciplinary structure of science (technology included) as information processing pervades it. In particular, as we become more aware of the detailed information processes that go on in doing science, the sciences will find themselves increasingly taking a metaposition, in which doing science (observing, experimenting, theorizing, testing, archiving....) will involve understanding these information processes, and building systems that do the object-level science. Then the boundaries between the enterprise of science as a whole (the acquisition and organization of the knowledge of the world) and AI (the understanding of how knowledge is acquired and organized) will become increasingly fuzzy. Disciplines will hardly fade away. The whole population can't all read the same literature, and other span-of-control issues exist as well. Most emphatically, AI (and computer science) will not become the new Queen of the Sciences. We live in too democratic an age for that. More likely is the diffusion into all of science of the scientific quest to understand the basic nature of intelligence and information processing."

## References

- [1] Srdanović V, The Second Information Revolution. In *Philisophy of Design and Information Technologies*, Caen, 59-64.1994.
- [2] Newell A, McDermott J, Forgy C, Artificial Intelligence A Self-paced Introductory Course, CMU, Pittsburgh, 1977.
- [3] Newell A, and Simon H.A, "Computer Science as Empirical Inquiry: Symbols and Search," *Communications of the ACM*. vol. 19, No. 3, pp. 113-126, March, 1976.
- [4] Feigenbaum E.A, The art of artificial intelligence: themes and case studies in knowledge engineering. IJCAI-5, 1014-1029, 1977.
- [5] Michie D, Current development in expert systems. Proc. 2nd Australian Conf. on Applications of Expert Systems, Sydney, 1986.
- [6] Quinlan J.R, Induction of decision trees. Machine Learning 1:1, 1986.
- [7] Srdanović V, BELART – A Consultation System Integrating Knowledge and Data, MEDINFO-86, (Salamon/Blum/Jørgensen, eds.), North-Holland, 228-231, 1986.
- [8] Michie D, Machine Learning in the Next Five Years. Proceedings of the 5th European Workshop on Machine Learning, D.H. Sleeman (Ed.). London: Pitman, 1989.
- [9] Strickland E, The Turbulent Past and Uncertain Future of Artificial Intelligence, IEEE Spectrum, 30 Sep 2021. <https://spectrum.ieee.org/history-of-ai>
- [10] Choi CQ, 7 Revealing Ways AIs Fail, IEEE Spectrum, 21 Sep 2021. <https://spectrum.ieee.org/ai-failures>
- [11] Bobrow D.G, and Hayes P.J, (Eds.), Artificial Intelligence - Where Are We?, *Artificial Intelligence* 25, 375-415. 1985.