Ubiquity, an ACM publication
April 2022

Ubiquity Symposium

Workings of Science

AI in 2156: The science of intelligence
by Kemal A. Delic and Jeff A. Riley

Editor's Introduction

While most people are not familiar with the details of how artificial intelligence (AI) works, the term itself is becoming more familiar to the non-scientific community, to the point that it ("AI") has almost become part of the regular vernacular of the ordinary person. AI has been with us for a long time—from first beginnings in ancient times in the form of automatons and other devices mimicking humans or other animals, through the middle of the last century when the term "artificial intelligence was actually coined, to the present times where it (the label rather than the actual technology) is entering the psyche of the general public.

This article explores the notion that the technology we call artificial intelligence is not yet ripe, but is establishing itself as a science in its own right, and that by 2156—the 200 year anniversary of the coining of the term—the technology should be in a position to deliver on its promises.

Ubiquity Symposium

Workings of Science

AI in 2156: The science of intelligence

by Kemal A. Delic and Jeff A. Riley

There is a long history of works on the imitation of human behavior, starting from the ancient stories and Greek myths via medieval devices for amusement, to modern approaches to mimic human thinking (See Appendix, "A Brief History Of AI"). The field of artificial intelligence (AI) started in a summer workshop in 1956 when this unlucky term was coined, and research was triggered. There was huge interest in AI, which was supported by generous funding, and subsequently survived (at least) two winters that were characterized by interest in AI waning and funds drying up.

These days, AI is somewhere between high hopes and strong fears, but it is certainly omnipresent in all kinds of media. AI has passed the initial period of theoretical explorations, often overpromising and later disappointing, until IT advanced to the point that many narrow applications appeared as result of experimentation, fine tuning, and engineering adaptation. We can say that AI is a very successful marketing tool if you do not ask for details, reasonably successful in deployment on mobile phones and in home gadgets, but as a general theory of intelligence is largely seen as a failure.

In this article, we submit that AI should become an established science by itself, but it must be projected on a much longer time scale, should involve several neighboring disciplines, and must have sound scientific bases in established sciences of biology, chemistry, physics, and mathematics. We project the year of 2156 as a possible horizon of this to happen, marking 200 years of the AI coinage in Dartmouth workshop.

ON INTELLIGENCE: SEVEN AI TECHNOLOGIES

The word "intelligence" derives from the Latin nouns *intelligentia* or *intellēctus***,** which in turn stem from the verb *intelligere,* to comprehend or perceive. In the Middle Ages, the word "intellectus" became the scholarly technical term for understanding, and a translation for the Greek philosophical term *nous*.

We discuss manifestation of intelligence starting from presumed, supreme human intelligence, all the way down to supposed intelligence of plants and fungus. We hypothesize about forces and factors shaping each form of intelligence, while thinking that the ultimate purpose and mission is defined by the single, well-chosen word. For example, coordination is the ultimate word describing the working mechanism of insect intelligence. Table 1 depicts five layers of intelligence observed on the Earth, starting from fungus and ending with human intelligence—all being somehow intertwined together into the mystery of life on the planet.

Table 1. Five Layers of Intelligence

INTELLIGENCE	**GOAL** to	**SHAPED** by
Human	Understand	Evolution
Animal	Survive	Competition
Insect	Reproduce	Coordination
Plants	Spread	Collaboration
Fungus	Dominate	Mimicry

"Intelligence" is a loaded word, indeed. It may have different meanings in different contexts, it casts a long shadow on philosophy, and causes significant controversies when associated with engineered machines, devices, and instruments. It becomes very challenging when confronted with the biology of life. Huge effort and literature was/is dedicated to the subject of intelligence and life on Earth, with notable examples of Schrodinger, Turing, Von Neumann. [1, 2, 3]

Over time, the AI field has developed an important number of technologies based on emulating human intelligence and mimicking human behaviors in a wide variety of activities: from game playing, via problem solving, to a highly qualified expert acting in several fields (see Table 2).

Table 2. Seven AI Technologies

Technology	**Inspiration**	**Typical Application**	**Challenges**
Neural Networks	Brain functioning	Recognition: Images, Sounds, Signals, Patterns	Self-explanation, Sensitivity on small changes in training data sets, dependency on data quality
Rule-Based Systems	Heuristics captured	Games/Gaming/Shallow Expert Systems	Hard to rescale and maintain coherence and integrity
Expert Systems	Human expertise captured	Problem Solving	Difficult to rescale and maintain
Case-Based Reasoning	Practical real- world cases used	Experience/ Knowledge-Based Systems	Working well in narrow domains and not frequently changed
Decision Trees	Logic-Based	Decision Making, Diagnostics, Analytics	Difficult to maintain for large-scale problems
Blackboard Systems	Roundtable of Collaborating Experts	Real-Time Decision- Making Systems	Few real-world deployments

Bayesian Networks	Sound mathematics combined with graph representing causes-consequences	Uncertainty Saturated Problem Solving	Creation of the structure, calculation of probabilities, establishing accurate priors

Many of these technologies produce results that approach, if not match, human performance, and while these results are typically achieved in well-defined, bounded, application domains, many of the technologies are likely to be generalizable, in time, to other domains. However, despite the much-hyped AI renaissance of the past decade or so, AI is still in its infancy, and matching the ensemble performance of the human brain is not likely to be achievable in the foreseeable future: the best we can hope for at the current stage of development is to try to mimic and match parts of the human brain that are responsible for specific activities and apply what we learn to specific domains. The hope is that over time our ensemble of technologies will converge on the ensemble performance of the human brain. Not all technologies have reached maturity or very wide applicability, but they have been the seed of computing technology augmentation—as was the most recent case with deep neural networks. Whether AI will ever match, or even get close to matching, the emergent abilities and functionality of the human brain (e.g. intuition, consciousness, etc.), whether AI could ever "think like humans", is not clear, and much debated [5, 6].

It is our belief that future advances will be based on observing intelligent behaviors in nature and trying to understand the deep principles involved—not just the external manifestations. Furthermore, aiming to mimic human intelligence is a noble goal, but we can, and should, learn more from the behaviors and intelligence of insects, animals, fish, and plants. Though not comparable to human intelligence, they have survived, and have evolved behaviors, over millions of years that we can learn from.

THIRD WAVE OR THIRD WINTER?

As we look back over the past several decades, and consider the immediate future, we clearly observe three distinct ages of AI research transforming high aims into practical systems and technologies:

The **embryonic** age is filled with enthusiasm, wildly optimistic promises, and ultimately missed steps while research is struggling to understand the basics of AI (e.g. problem solving, state space search and gaming). The focus of the field is more on mechanization and automation as first steps rather than developing a deep understanding of the biological processes that enable living systems to survive in a constantly evolving environment.

Much of the early hype emanated from uninformed commentators who equated the electronic computer with the human brain, and electronic computing and computation with human thought. Media perpetuation of the hype created quite unrealistic, sometimes bizarre, expectations, and even some fear, in the general public, while much of the AI research community went along for the ride out of self-interest: the publicity helped attract research funds.

The **embedded** age is facilitated by advances in computing technology, power, and speed. It is now possible to create embedded systems better able to emulate intelligent behavior. Embryonic search of earlier times is rescaled into an omnipresent global hyper-structure acting as an enormous knowledge base accessible to anyone, at any time, from anywhere. Mobile devices are now apparently able to understand speech and comprehend language. Machines outperform human world champions in chess, Go and natural-language knowledge quizzes. Robotics research advances from impressive demonstration devices to consumer and industrial robots doing real work in real environments. Autonomous vehicles, already prevalent in industrial and mining environments, are beginning to make appearances on our roads.

The **embodied** age is in its infancy. We expect social robots, some humanoid, to move from novelty to mainstream in the foreseeable future. Some social robots will exist purely to serve us as they interact with us—as waiters, bartenders, cooks, drivers, guides, and caregivers, etc. Others will exist simply as companions—some humanoid, but others constructed to resemble beloved pets.

We expect to see many more intelligent drones of varying capabilities and forms deployed in many more applications, all learning from their interactions with their environments. Swarms of intelligent drones, each capable of communicating with and learning from its peers, will be capable of carrying out tasks that in the past required large machines and many human operators.

We already see the occasional intrepid early adopter with microchips or other miniature electronic devices embedded in their bodies. These devices allow their hosts to monitor their own bodies and upload the data to a remote database for analysis, communicate and interact with nearby devices, and be monitored and tracked externally. We expect more integrated cyborgs—humans augmented with advanced mechanical, electronic, and AI devices—to become more prevalent in the not-too-distant future.

THE NEXT 100-PLUS YEARS: TIME TO DELIVER

There are many lessons for the AI community to learn from the past two winters, but probably the most important is that we need to set realistic expectations around the size, scope, and speed of progress to be made. After all the enthusiasm, hype, and promises of the embryonic age, what we've learned best is that AI is hard. Nature may make it look easy, but nature has had millions of years to find and refine solutions through trial and error—the AI research community has no such luxury. AI research, like all hard research, is complex, time-consuming work, characterized by many blind-alleys, missteps, and wrong turns—it will take time, progress has been slow and will continue to be slow, but progress is being made and will continue to be made (incrementally), and the field will deliver if we can ignore the uninformed hype and focus instead on what must be done to properly address the research issues.

Artificial Intelligence research is a broad, complex area of research, combining cognitive science, computer science, and robotics, but also reaching into quite diverse areas such as biomechanics and biometrics, law, ethics, and sociology to name just a few. Artificial Intelligence can no longer be considered a sub-branch of computer science—with AI producing potentially transformative applications across all areas of society including commerce, industry, health, entertainment, the arts, sport, and leisure activities, we must take a more multidisciplinary approach to AI research.

Moreover, many questions resulting from AI research cannot be answered strictly from a science or engineering perspective alone but require intimate involvement from a broad range of disciplines and participants. These questions are not limited to science. One of the goals of AI research is to develop intelligent systems that will safely interact with humans and the physical world across all of society, so questions of ethics, governance, impact, and accountability need to be addressed.

Because the product of AI research will become so pervasive, and important, across all of society world-wide, it is imperative that there be global coordination of research, and funding for research, to provide the right balance of academic, sociological, ethics, legal, scientific research, industrial innovation, and technology development skills and expertise to guide the field far into the future. The Agricultural Age lasted for some ten millennia, while the Industrial Revolution ushered in the new Industrial Age some two hundred years ago. The new age of information is now upon us, and just as the industrial revolution changed agricultural society beyond recognition, we believe the Information Age, might be driven largely by AI, will cause a profound change in society, the economy and ecology.

We will see the emergence of a hybrid society consisting of humans assisted by a wide variety of embodied, autonomic devices, working together in a smooth partnership in the new digital economy. This partnership will become essential to managing a balanced exploitation of limited resources as the Earth's population expands at an increasing high rate. Humans will likely be enhanced (augmented) and more powerful—we are destined to become a race of cyborgs—and will live very different lifestyles. Augmentation of humans will take different forms, and may, as a result, spawn different cyborg castes.

Augmentation will come in two forms:

External, wearable devices and prostheses, such as exoskeletons and artificial limbs that will enhance or replace the biomechanical functions of the human body, with embedded intelligence that will allow these "enhancers" to learn and adapt, to predict and anticipate. These devices will both enhance and protect humans—protect from the stresses and injuries caused by overtaxing the human body, as well as from external threats such as falling objects, projectiles, and predators.

Internal, or embodied, devices that will interface directly with the human brain and nervous system. These augmentations will range from enhanced sensors, intelligent bio-monitoring devices capable of regulating human functions (heart rate, chemical levels etc.), auxiliary memory and computation devices (co-processors, long-term storage arrays), devices capable of wireless interfaces to external computing and information systems (e.g. GPS, weather information systems, emergency services and alerting systems, etc.).

The science of intelligence (previously known as AI) will infuse new technologies into all aspects of social, economic, and cultural life. The composition of the world's population will change: Cyborgs and robots will join humans as functioning members of the population. Political systems will be changed beyond recognition. The role of humans vis-à-vis cyborgs vis-à-vis full robots will need to be considered. Laws to deal with the reality of different lifeforms, even to define what a lifeform is, will need to be debated and implemented.

Figure 1. The three ages of artificial intelligence.

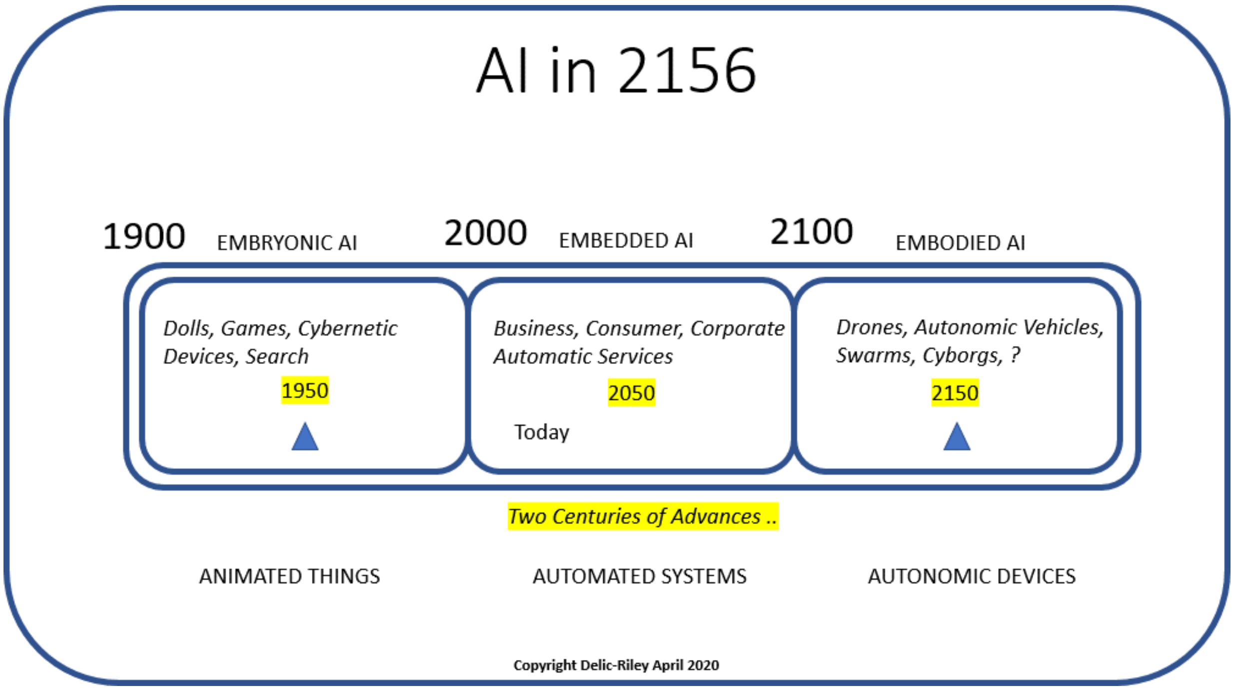

We have navigated the embryonic and embedded ages of AI, and we are just embarking upon the embodied age (Figure 1). The embryonic and embedded ages of AI have built the foundations for the embodied age, the age of augmented humans and autonomic devices. The stage is now set for AI to come into its own. AI has the potential to enhance and improve the lives of humans enormously, and so it must. The next 100 years, the embodied age of AI, is the time for AI to deliver on its promises, after surviving the third winter [6].

References

[1] Schrödinger, E. *What is Life? With Mind and Matter and Autobiographical Sketches*. Cambridge University Press, 1992.
[2] Von Neumann, J. *The Computer and the Brain*. Yale University Press, New Haven 2012.

[3] Turing, A. Computing machinery and intelligence. *Mind*, 1950.
[4] Riley, J. The elusive promise of AI: A second look. *Ubiquity* 2021, April (2021),
[5] Riley, J. Will machines ever think like humans? *Ubiquity* 2021, June (2021),
[6] *IEEE Spectrum*. The Great AI Reckoning. Special Report, October 2021.

About the Authors

Kemal A. Delic is a Senior Visiting Research Fellow working in The Center for Complexity and Design at The Open University. Previously he was a Senior Enterprise Architect and Senior Technologist with Hewlett-Packard for many years.

Jeff A. Riley is the founder of Praescientem, a company specializing in AI and IT consulting and education services. Previously he was a Master Technologist and Scientist with Hewlett-Packard for many years. Jeff holds a master's degree in IT, a Ph.D. in AI, and is a former Adjunct Principal Research Fellow of RMIT University in Melbourne, Australia.

DOI: 10.1145/3512335

Ubiquity Symposium

Workings of Science

Supplementary Material for "AI in 2156: The science of intelligence"

Kemal A. Delic and Jeff A. Riley

APPENDIX: "A BRIEF HISTORY OF AI"

The first academic conference on artificial intelligence, organized by John McCarthy, was held in 1956, and so the 1950s is generally regarded as the beginning of AI research. But there is a long, historical interest in early automata and imitated animals. But the real dawn of humankind's interest in artificial intelligence lies in our ancestors' interest in animated things, and this fascination stretches far back into antiquity.

According to Jewish legend, about 3000 years ago King Solomon's throne was constructed with mechanical animals that stretched out their feet to support him and help him rise to the next step when he stepped upon the throne, then once seated a crown was placed upon his head by a mechanical eagle.

Ktesibios, a Greek inventor born around 285 BC, is thought to have created various water powered automata. He is said to have "used water to sound a whistle and make a model owl move"—the world's first cuckoo clock.

Chinese texts record separate accounts of flying automata (artificial wooden birds) constructed by 5th century philosophers Mozi and Lu Ban.

In the late 15th and early 16th century Leonardo da Vinci created a slew of different machines, including a life-like mechanical lion automaton commissioned by Pope Leo X as a gift for French King Francois I.

In 1739 Jacques de Vaucanson unveiled his *Canard Digérateur*, or Digesting Duck, an automaton in the form of a duck that was able to move, eat, and produce faeces, while not pretending that this is a real, live duck. In another type of imitation, which was discovered as a hoax, Wolfgang von Kempelen constructed the Mechanical Turk in 1770, an automaton able to play chess. Some of these accounts are true, some probably apocryphal, but all are testament to humankind's fascination with automata stretching back thousands of years. A word first used by Homer to describe automatic door opening, "automaton" is from the Greek "acting of one's own will." Humans have a long history of interest in machines that act like humans—an early fascination that bloomed into today's artificial intelligence.

The rapid advances in computing hardware and computing methods following WWII resulted in extremely efficient technologies and a rich environment to enable ideas of artificial intelligence to advance. It is at that time that serious thinking about the possibility that machines can think is born, elaborated, and developed. Alas, hype surrounding AI research has created very high expectations and unrealistic timetables leading to (at least) two AI winters during which funding for AI research and development of technologies evaporated, interest in AI waned, and companies folded. This typically started with fascinating demos, being developed into prototypes which consistently were not able to rescale nor get fully operational in real-world circumstances. Interest waned out and prototypes were discarded and forgotten—as is the largest expert system for computer configuration tried by the major computer company in the 1970s.

Eight Decades of AI

We present a rough outline of the developments of AI during the past eight decades in Table 1 and the text that follows, where we indicate some major technology domains in a decade-span time frame and mention representative researchers and examples.

Table 1.

	DOMAIN		
Decade	**Technology**	**Names [e.g.pub]**	**E.g. App**
1940s	Perceptron/Neural Networks Problem-Solving Cybernetics	McCulloch and Pitts [12] Polya [17] Bush [7] Ashby [1, 2] Wiener [25]	Trial & Error Search Self Organization
1950s	Embryonic Theories of AI Toy Prototypes Language Systems	Turing [23]	GPS
1960s	Heuristic Search Adaptation Decision-Making Language Understanding	Minsky [14] Tsypkin [21] Simon [20]	Hearsay
1970s	Control Systems Expert Systems Q&A Systems Knowledge representation	Tsypkin [22] Nilsson [15]	Dendral-Mycin
1980s	Expert Systems Machine Learning Robotics Flocking/Swarm Intelligence	Waterman [24] Michie [13] Brooks [5] Reynolds [18]	MENACE
1990s	Autonomic Computing Games Intelligent Systems Knowledge Based Systems	Hayes-Roth [11] Russell and Norvig [19]	Watson

	Intelligent Agents		
2000s	Autonomous robots Swarm Intelligence Voice, Image, and Pattern Recognition	Stone [6] Dorigo [9]	Google Translate Siri Assistant IRobot Cleaner
2010s	Biological Systems Intelligence, Swarm Intelligence Drones - air/sea/undersea Adversarial learning Face and Signal Recognition, Video and Audio processing	Bonabeau [4] Goodfellow [10]	Swarm-Bots Scientific and Military Drones Medical Imaging Bot Assistants
2020s	Deep Learning Systems Automatic Text Generators AI Chips	Senior et al [8] Bengio et al [3] Oren Etzioni [16]	DeepMind openGPT2-XL

1940s

This period is precursor to AI research which was accelerated during the Second World War with several developments supporting military operations (radar, computers, nuclear) which later found business and civil uses.

1943. Warren McCulloch, a neuroscientist, and Walter Pitts, a logician, publish “A Logical Calculus of the Ideas Immanent in Nervous Activity” in the *Bulletin of Mathematical Biophysics* and present a simplified model of a biological neuron, leading to the development of the perceptron and laying the foundations for the development or artificial neural networks.

1945. Stanford University mathematician George Pólya publishes *How to Solve It*, describing methods of problem solving and laying out the principles of heuristic reasoning. Heuristics was explored as human experience in problem solving captured in a suitable formal framework and made popular by Herbert Simon, Nobel prize in Economics. The very idea of expert systems is based on heuristic knowledge—called also rules of thumb. It has roots in Greek's tradition of finding quick shortcuts to problem solving—as the name heuristics implies—to find and discover.

1945. Vannevar Bush, head of the U.S. Office of Scientific Research and Development (OSRD) and founder of Raytheon Corporation proposes a system that amplifies people’s own knowledge and understanding in his article “We May Think” for *The Atlantic*.

1945. Psychiatrist and cybernetics pioneer W. Ross Ashby heralds cybernetics and self-organization with his article “The Physical Origin of Adaptation by Trial and Error” in The Journal of General Psychology, in which he addresses the question of how a machine could (self-)adapt to its environment, and later lays out the principles of self-organizing systems in the 1947 article in the same journal, “Principles of the Self-organizing Dynamic System.”

1948. MIT mathematician and philosopher Norbert Wiener publishes his book *Cybernetics: Or Control and Communication in the Animal and the Machine*, pointing at the importance of the closed feedback loop for adaptive systems. As adaptation is also one of the remarkable feats of intelligent behavior, some claim that his work belongs to AI research, as well laying the theoretical foundation. His theory supported many practical systems which we see in use today, such as servomechanisms, automatic navigation, analog computing, artificial intelligence, neuroscience, and reliable communications.

1950s

1950. Alan Turing, mathematician, computer scientist, logician, cryptanalyst, philosopher, and theoretical biologist, poses the question "Can machines think?" and proposes The Imitation Game, often referred to simply as The Turing Test, as a means to test the question, in his seminal work "Computing Machinery and Intelligence," published in *Mind*, the journal of analytic philosophy.

1957. Early work of Frank Rosenblatt on Perceptron was very inspirational for later developments.

1960s

1961. In his paper "Steps Toward Artificial Intelligence" (*The Journal of the Institute of Radio Engineers*), Marvin Minsky, cognitive scientist and co-founder of MIT's AI laboratory, outlines five essential skills for AI: search, pattern-recognition, learning, planning, and induction. Minsky goes on to show how these processes could be mathematically constructed into a programmed language a machine could follow.

1966. In his paper "Adaptation, Learning and Self-learning in Automatic Systems" published in the control theory journal *Automation and Remote Control*, and in a 1971 book *Adaptation and Learning in Automatic Systems*, Yakov Tsypkin proposes the extension of classical optimization theory to include his "probabilistic iterative methods."

1968. In December 1968, Doug Englebart presented the whole slew of fascinating technologies during a singular event called "The Mother of all Demos." Several decades later, one can recognize some devices and applications in wide commercial use and having a strong similarity to things Englebart demoed. Today's web is the embodiment of the human knowledge amplifying system accurately envisioned by this demo. One may have Impression that he has demonstrated today's technologies, but he did it only more than half a century ago.

1969. Herbert Simon, Nobel Prize laureate, economist, political scientist, and cognitive psychologist, publishes his book *The Sciences of the Artificial*, an exploration of the organization of complexity, problem-solving and decision-making.

1969. Minsky and Papert have published in 1969 an entire book, *Perceptron*, arguing that one-layer neural nets cannot emulate human perception—which by itself has triggered the first AI winter. Later, Lighthill report in the UK triggered the second AI winter in 1973. Competition for the funding was the most likely cause of dispute that triggered both AI winters.

1970s

1971. Nils Nilsson, computer scientist and co-inventor of the *A** search algorithm, publishes his textbook *Problem-solving Methods in Artificial Intelligence*, outlining the theoretical ideas underlying problem-solving by heuristically guided, trial-and-error search processes.

1972. Hubert Dreyfus was one of the few sceptics challenging the AI field at that time with his book *What Computers Can'T Do: The Limits of Artificial Intelligence*. His criticism was rejected and later confirmed as being right to the point by comparing AI research to alchemy. Alchemy aimed at the impossible goal of turning dirt into gold, while inventing many useful, practical things in that process.

1980s

1985. Donald Michie, one of the pioneers of Bletchley Park and robotics (e.g. the world-first FREDDY robot project of the 1960s and '70s), publishes an article proposing robots be used as teaching and learning aids, and that experimental robotics be a branch of machine intelligence. Michie considers that what is known as artificial intelligence is effectively machine learning.

1986. Rodney Brooks, founder of iRobot, former Panasonic Professor of Robotics at the Massachusetts Institute of Technology, and former director of the MIT Computer Science and Artificial Intelligence Laboratory, publishes his layered control system for mobile robots.

1986. Donald Waterman publishes *A Guide to Expert Systems*, a basic introduction to expert systems and their use, and one of the first textbooks to appear on the subject.

1987. Craig Reynolds presented both a technical paper ("Flocks, Herds and Schools: A distributed behavioral model"), and a short, animated film ("Stanley and Stella in: Breaking the Ice") demonstrating the simulated flocking behavior of birds—a forerunner of swarm intelligence.

1987. Terry Winograd and Fernando Flores publish *Understanding Computers and Cognition*, which contributed greatly to the second AI winter.

1990s

1991. Frederick Hayes-Roth and others publish their paper, "Frameworks for Developing Intelligent Systems," describing a next-generation operating system that provides an environment for the development of intelligent systems. The environment is described by the subtitle of the paper: "The ABE Systems Engineering Environment," where "ABE" stands for "a better environment."

1995. Stuart Russel and Peter Norvig publish *Artificial Intelligence: A Modern Approach*, a university-level textbook based on the idea of Intelligent Agents. The book is now widely considered the standard text in the field of artificial intelligence, and sometimes described as “the most popular artificial intelligence textbook in the world."

2000s

2005. Rodney Brooks and other AI luminaries, discuss the progress in AI over the 25 years since 1980. Brooks is upbeat about the progress made and the future prospects of AI. He has created a couple of very successful robotic companies and advanced research in robotics, making him one of the most accomplished practitioners in the field.

2009. Marco Dorigo, AI researcher and inventor of the ant colony optimization metaheuristic, delivers an invited talk at the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WA-IAT ’09) in Milano, Italy, in which he discusses the swarm-bots and swarmanoid experiments in swarm robotics.

2010s

2010. *Natural Computing* publishes a special issue, edited by Eric Bonabeau and others, surveying the state-of-the-art in swarm intelligence—the emergent and collective intelligence of groups of simple and autonomous agents.

2014. Ian Goodfellow and others publish the first paper describing a new machine learning framework, known as generative adversarial networks, in which two neural networks are pitted against each other in contest. GANs prove to be useful for semi-supervised learning, fully supervised learning, and reinforcement learning.

2020s

2018. The Turing Award is given to Bengio, Hinton, and LeCun and ushers in major breakthroughs in artificial intelligence. They revived the neural nets (NN) research re-appearing as deep learning, which was in reality an extension of the single layer NN into network with several layers and showing that wide availability of data and computing resources can advance the AI field. Today, it is the major direction of AI research.

References

[1] Ashby, W. R. The physical origin of adaptation by trial and error. *Journal of General Psychology* 32 (1945), 13–25.
[2] Ashby, W. R. Principles of the self-organizing dynamic system. *Journal of General Psychology* 3 (1947), 125–128.
[3] Fathers of the Deep Learning Revolution Receive ACM A.M. Turing Award. ACM New, 2018.

[4] Bonabeau, E., Corne, D., Poli, R. Swarm intelligence: the state of the art special issue of natural computing. *Natural Computing* 9, 3 (2010), 655–657.
[5] Brooks, R. A robust layered control system for a mobile robot. *IEEE Journal or Robotics and Automation* 2, 1 (1986) 14–23.
[6] Stone, M. and Hirsh, H. Artificial Intelligence: The next twenty-five years. *AI Magazine* 26, 4 (2005), 85–97. https://doi.org/10.1609/aimag.v26i4.1852
[7] Bush, V. As we may think. *Atlantic Monthly*, July 1945.
[8] Senior, A.W., Evans, R., Jumper, J. et al. Improved protein structure prediction using potentials from deep learning. *Nature* 577, 706–710 (2020).
[9] Dorigo, M. Swarm-Bots and swarmanoid: Two experiments in embodied swarm intelligence. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*. September 2009.
[10] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. *Proceedings of the International Conference on Neural Information Processing Systems*. 2014, 2672–2680.
[11] Hayes-Roth, F., Davidson, J.E., Erman, L. D., and Lark, J. S. Frameworks for developing intelligent systems. *IEEE Expert* 6, 3 (1991), 30–40.
[12] McCulloch, W. and Pitts, W. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 7 (1943), 115–133.
[13] Michie, D. Steps toward robot intelligence. *Proceedings of the Royal Institute* 57 (1985), 151–165.
[14] Minsky, M. Steps toward artificial intelligence. *Proceedings of the IRE* 49, 1 (1961), 8–30. doi: 10.1109/JRPROC.1961.287775.
[15] Nilsson, N. *Problem-solving Methods in Artificial Intelligence*. McGraw-Hill, New York, 1971.
[16] Etzioni, O. et al. Green AI. *Communications of the ACM* 63, 12 (2020).
[17] Polya, G. *How to Solve It*. Princeton University Press, 1945.
[18] Reynolds, C. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 1987, 25–34.
[19] Russell, S. and Norvig, P. *Artificial Intelligence: A modern approach*. Prentice Hall, 1995.
[20] Simon, H. *The Sciences of the Artificial*. MIT Press, Cambridge, 1969.
[21] Tsypkin, Y. Z. Adaptation, learning and self-learning in automatic systems. *ARC* 27, (1966), 23–61.
[22] Tsypkin, Y. Z. *Adaptation and Learning in Automatic systems*. Academic Press, 1971.
[23] Turing, A. Computing machinery and intelligence. *Mind*. 1950.
[24] Waterman, D. *A Guide to Expert Systems*. Pearson 1986.
[25] Wiener, N. *Cybernetics, or, Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, 1948.