History and Culture of AI

Maria Vanina Martinez

LAIS 2024 – Day 3

CSIC
Consejo Superior de Investigaciones Científicas
IIIA-CSIC
Artificial Intelligence Research Institute

Tentative Schedule

- 2/9/2023 - Class 1: What do we mean when we say “Intelligence” in Artificial Intelligence?
- 3/9/2923 - Class 2: On thinking AI: goals, assumptions, lines of research
- 4/9/2023 – No class - Workshop
- 5/9/2023 - Class 3: On building AI and a timeline of AI
- 6/9/2023 - Class 4: AI: agents, search and knowledge representation

A timeline of significant events in AI

Source: Getty images

A timeline of AI

- We will take a look at the history of Artificial Intelligence.
- I will try to highlight some of the most interesting moments.
- Part of this presentation is based on the content of chapter 1 of the book:

  “Artificial Intelligence: A Modern Approach” Stuart Russel & Peter Norvig

The Beginning

1946-1956

The beginning: 1943-1956

- The work *“ A logical calculus of the ideas immanent in nervous activity,” WS McCulloch & W Pitts - Bulletin of Mathematical Biophysics, 5, 115-137 (1943)* is recognized today as the first AI work.
- It started from three sources:
    - The basic physiology of brain neurons
    - Russell and Whitehead's work on logic
    - And the Theory of Computation proposed by Turing several years before.

The beginning: 1943-1956

- McCulloh and Pitts showed that any computable function could be computed by some network of artificial neurons and that every logical connective could be implemented by simple network structure (the perceptron: a linear binary classifier).
- These authors suggested in their work that these networks could learn.
- In 1949, D. Hebb showed that an update rule, which modified the strength of the connection, allowed learning to occur.

The beginning: 1943-1956

- In the early 1950s, Claude Shannon and Alan Turing wrote programs that played chess.
- In 1951, two Princeton mathematics graduate students, Marvin Minsky and Dean Edmonds, built the first neural network-based computer (SNARC).
- Minsky's doctoral thesis panel was hesitant to award the degree in mathematics for this type of work, but John von Neumann, who was a member of the panel, assured:

“If this work is not mathematics today, it will be one day”

- Ironically, Minsky's later work contributed to the postponement of research in Neural Networks.

The beginning: 1943-1956

- In 1950 Alan Turing published a paper entitled:

"Computing Machinery and Intelligence" Mind , Vol LIX, No 236 (1950)

which was the beginning of a careful treatment of the concept of AI as a scientific object.

- This paper analyzed the question:

"Can Machines Think?"

- It included a test to determine whether the potential candidate was intelligent or not.
- A long list of possible objections to the possibility of a machine thinking was also discussed.

The beginning: 1943-1956

- At Princeton was another of the founding fathers of AI, John McCarthy.
- Upon graduation, he transferred to *Dartmouth College* , which would become the birthplace of AI.
- McCarthy convinced Minsky, Shannon and Rochester to hold a meeting of researchers interested in Automata Theory, Neural Networks and the study of intelligence.
- A two-month Workshop was organized at Dartmouth College during the summer of 1956.
- Two researchers who participated in it, A. Newell and H. Simon, had a predominant influence.

The beginning: 1943-1956

- Newell and Simon already had a reasoning program called *Logic Theorist* (LT).
- Simon stated: *"We have invented a computer program that is capable of non-numerical thinking, and have thus solved the Mind-Body problem."*
- This program was able to prove most of the theorems in Chapter 2 of Russell and Whitehead's *Principia Mathematica*, even finding a more elegant proof than the one contained there.
- Russell himself was delighted with the result.
- However, the *Journal of Symbolic Logic* rejected a paper authored by Newell, Simon and LT.

The beginning: 1943-1956

- The Dartmouth Workshop did not produce many significant scientific advances.
- However, he connected researchers interested in these problems and formed the community that dominated the new field for the next 20 years together with his students and colleagues at MIT, CMU, Stanford and IBM.
- The most lasting outcome of this meeting was the general agreement on John McCarthy's proposal to name the new field of study *Artificial Intelligence.*
- McCarthy himself had sporadically expressed some regret about this choice.

Much Enthusiasm …
Great Expectations

1952-1969

1952-1969

- These years were full of successes… of limited features.
- At that time there was an external consensus in the area that *“A machine can never do X”* for a large number of *Xs* that Turing took care to enumerate in his work.
- AI researchers responded (in a limited way) by showing that *X* was possible for one *X* after another.

1952-1969

- Newell and Simon's early success with the *Logic Theorist* was followed by the *General Problem Solver* (GPS).
- This program was designed to mimic human problem-solving protocols.
- This was probably the first program to embody the AI approach of “thinking humanly.”
- This work was carried out at Carnegie Mellon University (CMU) where AI and Cognitive Science continue to work together to this day.

1952-1969

- Beginning in 1952, Arthur Samuel developed a series of programs devoted to playing checkers.
- These programs began *by knowing* only the basic rules of the game and as they played they learned, eventually reaching a competitive level.
- Like many others, Samuel had trouble getting computer time and worked using IBM prototypes at night.
- This situation of not having enough computer time, which was general, led some researchers to think harder about the problem and how to solve it.

1952-1969

- In 1957, Rosenblatt built Mark I Perceptron, a hardware implementation of a perceptron, at the Cornell Aeronautical Laboratory, and was funded by the Information Systems Branch of the United States Office of Naval Research and the Rome Air Development Center.
- It was first publicly demonstrated on 23 June 1960. It was later part of a project (1963 - 1966) to develop a tool for photo-interpreters.
- The Mark I Perceptron has 3 layers:
  - An array of 400 photocells arranged in a 20x20 grid, named "sensory units" (S-units), or "input retina". Each S-unit could connect to up to 40 A-units.
  - A hidden layer of 512 perceptrons, named "association units" (A-units).
  - An output layer of 8 perceptrons, named "response units" (R-units).

1952-1969

- In 1958, McCarthy moved to MIT where he made three crucial contributions that same year.
- In *MIT AI Lab Memo No. 1* he defined the *LISP language* based on Lambda Calculus (being the second programming language defined so far).
- Motivated by the lack of computing time, he invented *time sharing*. After a while, some MIT graduate students were attracted by this novelty and later formed the Digital Equipment Corp.
- Finally, he wrote *Programs with Common Sense.*

1952-1969

- In *Programs with Common Sense*, McCarthy described *Advice Taker*, a hypothetical program that can be considered the first complete AI system.
- The program is designed to use knowledge to find solutions to problems. This knowledge was general knowledge about the world.
- Another novel feature was the ability to accept new axioms to acquire new capabilities without being reprogrammed.
- The core features of Knowledge Representation and Reasoning were already included in Advice Taker, *i.e.*, *having a formal representation of the world and manipulating it through deductive processes*.

1952-1969

- Also in 1958, Marvin Minsky moved to MIT.
- For years he worked alongside McCarthy but they slowly drifted apart as Minsky took a more practical approach as opposed to McCarthy's formalism.
- McCarthy (1963) moved to Stanford and founded the AI Lab.
- McCarthy's agenda was boosted by J.A. Robinson's publication of the resolution method which represents an algorithm for proving theorems in First Order Logic.
- The work at Stanford emphasized the development of *general-purpose methods for logical reasoning* .
- The first robotics project was developed at SRI.

1952-1969

- Minsky supervised a series of graduate students who chose problems in limited domains whose solution seemed to require intelligence.
- These restricted domains became known as *micro-worlds*.
- J. Slagle produced SAINT which could solve certain mathematical integration problems.
- T. Evans wrote ANALOGY, a program that could solve geometric analogy problems similar to those used in IQ tests.
- D. Bobrow created STUDENT, which solved algebra problems described in text form.

1952-1969

- The best-known micro-world is the *block world* which consists of solid blocks placed on a table.
- One task is to rearrange blocks in a certain way using a robot arm that can pick up one block at a time.
- This micro-world was later used by:
  - ✓ P.Winston (1970) Learning Theory.
  - ✓ D.Huffman (1971) Vision.
  - ✓ T. Winograd (1972) Natural Language.
  - ✓ S.Fahlman (1974) Planning System.
  - ✓ D.Waltz (1975) Propagation of Constraints.

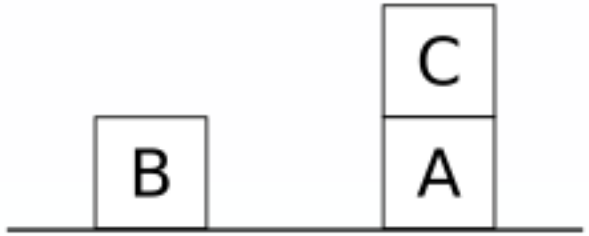

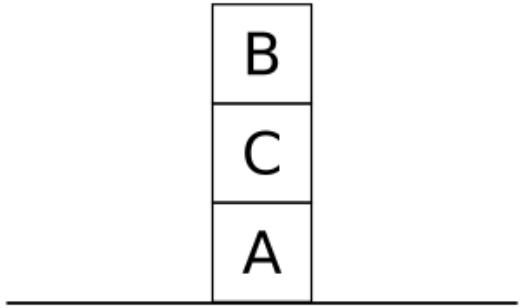

Source: wikipedia

1952-1969

- Work with neural networks had had sustained development.
- Winograd and Cowan (1963) showed that a large number of elements working collectively could represent a single concept with corresponding improvement in robustness and parallelism.
- The Hebb method for learning was also improved by Widrow and Hoff (1960-1962) and Rosenblatt (1962) for perceptrones.

Reality knocks the door…

1966-1974

1966-1974

- En 1957 Herbert Simon produced the following statement showing the euphoria of that particular moment:

“It is not my aim to surprise or shock you — but the simplest way I can summarize [this] is to say that there are now in the world machines that can think, that learn and that create. Moreover, their ability to do this things is going to increase rapidly until — in a visible future — the range of problems they can handle will be coextensive with the range to which human mind has been applied.”

- In 1958 he predicted that within 10 years a computer would be the chess champion, and that a computer would prove an important mathematical theorem.

1966-1974

- However, there were serious difficulties in using the advances achieved in the laboratory.
- In 1965 Weizenbaum produced ELIZA, a program that seemed to be able to carry on a conversation about any subject.
- This stimulated the idea that the natural language problem was soon to be solved.
- However, the first machine translators that used dictionaries as a fundamental tool and that was attempted to be used during the Cold War was very inadequate.

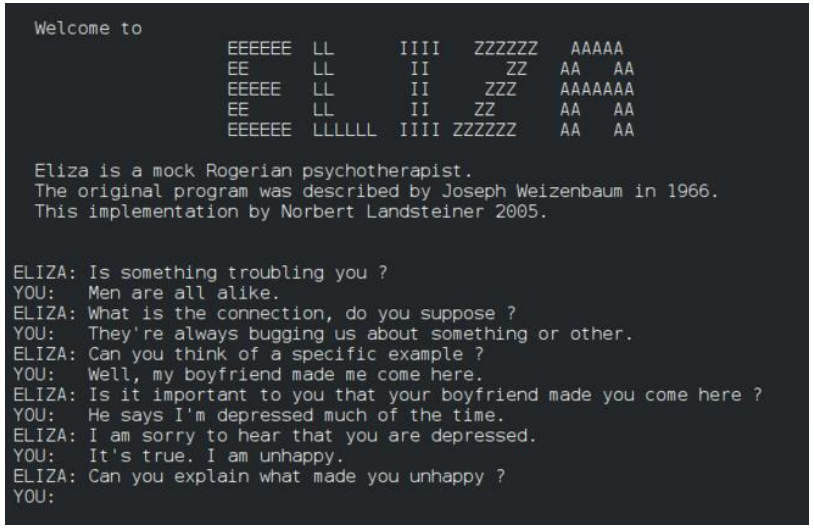

1966-1974

- An example, contained in almost any AI textbook, is the following double translation from English to Russian and back to English:

“The spirit is willing but the flesh is weak”

It was then translated into Russian and machine translated again into English:

“The vodka is good but the meat is rotten”

- In 1966 a US government committee report recommended the suspension of all funding for these programs.

1966-1974

- Another difficulty that began to be recognized was the intractability of many of the problems that AI was trying to solve.
- The programs that were being produced worked because they used micro-worlds and the hypothesis was that they could be extended to the real world simply with more powerful machines and more memory.
- With the development of *Complexity Theory,* the illusion disappeared drastically, as it was shown that such an increase in scale was impossible to achieve.
- The combinatorial explosion problem was the basis of the 1973 Lighthill report, which the British government used to suspend funding for AI.

1966-1974

- Another problem arose from the limitations of the basic structures used at the time.
- In 1969 Minsky and Papert published *Perceptrons* and in that book they showed that while perceptrons were capable of learning anything they could represent, they could actually represent very little.
- Although these results did not apply to more complex multi-layer networks, funding for neural network research dried up.
- In 1969 Bryson and Ho developed a new back-propagation learning algorithm that led to the resurgence of neural networks but 15 years later.

Knowledge-Based Systems

1969-1979

1969-1979

- Problem-solving ideas from the previous decade were based on general-purpose search methods.
- These methods have been called *"weak methods"* because they use little information about the domain.
- This leads to poor performance in some complex domains as well.
- The paradigm shift came about by developing systems that used a lot of knowledge from a narrower area.

1969-1979

- DENDRAL is an early example of this type of approach.
- It was developed at Stanford by E. Feigenbaum, a student of H. Simon, B. Buchanan, a philosopher turned computer scientist, and J. Lederberg, a Nobel Prize winner in Biology.
- DENDRAL inferred the molecular structure of a compound from information provided by a mass spectrometer and the chemical formula.
- This is a problem for which searching the space of all possible configurations had no chance of success.

1969-1979

- DENDRAL worked using rules formulated by experts in the field.
- The team that developed DENDRAL concluded that the power of this system lay in the fact that:

“All theoretical knowledge relevant to solving these problems has been brought in its general form of the predictive component of the mass spectrum (first principles of science) to efficient special forms (cooking recipes)”

1969-1979

- The importance of DENDRAL was that it was the first *Knowledge-Based System*.
- Its expert quality came from the large number of general-purpose rules it used.
- Later systems evolved into what we would today call an *architecture* in which the knowledge (*a set of rules*) and the reasoning component were separated following what McCarthy proposed for Advice Taker.
- Building on this success, Feigenbaum founded the *Heuristic Programming Project* to study the application of these methods in other areas.

1969-1979

- E. Feigenbaum, B. Buchanan and E. Shortliffe developed the MYCIN system which had the ability to expertly diagnose infectious diseases in the blood.
- An important difference with DENDRAL was that it did not have a theoretical model to support the rules it used.
- These rules came from interviews with experts in the field.
- Another important difference, which represented a significant advance, was that MYCIN used uncertainty in its rules, reflecting the uncertainty of medical knowledge itself.

1969-1979

- PROSPECTOR produced by R. Duda was a system that recommended drilling locations for geological exploration in search of molybdenum deposits.
- SHRUDLU was a system written by T. Winograd in the area of natural language understanding that used a lot of knowledge in the block world domain.
- Another important system was LUNAR, developed by W. Woods and used by geologists to ask questions, in natural language, concerning rocks brought back by the Apollo missions.

The AI Industry

1980-1988

1980-1988

- The first successful expert system was R1 .
- This system was developed by D. McDermott in 1982 for DEC to help configure orders for new computer systems.
- In 1986, DEC reported saving $40M per year.
- By 1988, DEC had 40 systems in operation.
- Dupont had 100 systems in use and 500 in development and reported savings of US$10M.

1980-1988

- In 1981, in Japan, they announced the launch of their 10-year project to develop the 5th Generation of computers.
- Machines would be built such that their machine language would be Prolog with the aim of accelerating the inference capacity of the systems.
- In this way it would be possible to manage immense knowledge bases and even achieve understanding of natural language.
- This project boosted the development of AI in USA, Europe and other countries around the world.

1980-1988

- In USA, the Microelectronics and Computer Technology Corp. (MCC) was created as a research consortium to keep up with Japan.
- In Britain, the Alvey report brought funding back to the area.
- In both cases, work was done on developing new circuits and interfaces.
- Parallel computing architectures based on Lisp and Prolog were developed.
- Specialized systems (shells) were produced for the development of expert systems.
- The industry grew from a few million in sales in 1980 to 2 billion in 1988.

Neural Networks

1986-…

1986-…

- While Minsky and Papert's book halted research into neural networks in computer science, *and funding for them*, work continued in other fields such as physics.
- Large numbers of neurons could be thought of or understood as large numbers of atoms, and physicists used methods and techniques from statistical mechanics to study these networks.
- This led to a very interesting cross-fertilization of ideas.

1986-…

- Some psychologists, D. Rumelhart and G. Hinton among others, continued with the study of neural network models for memory.
- In the mid-1980s, the back-propagation algorithm developed in 1969 by Bryson and Ho was reinvented.
- At the same time, it was concluded that creating expert systems was not as simple as it seemed.
- The impetus of research into neural networks and this decline of expert systems was seen as a contradiction between the two.

1986-…

- The late 1980s and early 1990s saw a major shift in AI research methodology.
- Some have seen this shift as a victory of the *neat (those who think that theories in AI should be supported by formal rigor)* over the *scruffy. (those who think that prototypes should be built and those that seem to work should be analyzed).*
- However, both perceptions are necessary but perhaps refocusing on more formal approaches demonstrates greater maturity of the discipline.
- The idea began to build on existing theory rather than continuing to create new theories.

1987-…

- Since the mid-90s the *Agents Revolution began.*
- The focus on designing a complete intelligent entity (an *Agent*) had led to the need to assemble different results into a single architecture.
- An agent must *sense its environment*, *must represent it*, *must reason with that representation*, and *must be able to plan its actions* .
- To plan actions you must have *your own goals*, which implies *autonomy*.

1987-…

- In multi-agent systems each agent develops its activity in an environment that may contain other agents and in this way social problems immediately appear.
- Do one agent's goals *match* those of others?
- Are one agent's goals in *conflict* with those of others?
- Do agents *collaborate* to achieve their goals?
- How is this collaboration achieved?
- How do they *communicate*?
- …

1997

- IBM’s Deep Blue beats Kasparov at chess
- In a series of six games played in a television studio, Deep Blue won two games, Kasparov won one, and three of the games ended in draws. Kasparov had defeated an earlier version of Deep Blue the year before.
- Deep Blue had plenty of computing power, and it used a “brute force” approach, evaluating 200 million possible moves a second to find the best possible one.
- Humans have the capacity to examine only about 50 moves per turn.
- The effect of Deep Blue was AI-like, but the computer was not actually thinking about strategy and learning as it played, as later systems would.
- Deep Blue’s victory over Kasparov brought AI back to the public mind.

The Future of AI in 2003

- According to John McCarthy (JACM, 50.1, January 2003)
  - ✓Achieving human-level intelligence (long term).
  - ✓Getting to the point where programs can learn from books (or other similar sources).
  - ✓Specification of programs that can interact with humans and other programs.
  - ✓Give users full control of their programming environments, that is, give users the ability to reprogram their environments without needing to understand more than necessary.

2011

- DeepQA (Watson) beat the best Jeopardy! player.
- “The goal is to build a computer that can be more effective in understanding and interacting in natural language, but not necessarily the same way humans do it.”
- The original Watson was a room-size computer consisting of 10 racks holding 90 servers, with a total of 2,880 processor cores.
- Over a period of years, Watson ingested mountains of information from Wikipedia and encyclopedias, dictionaries, religious texts, novels, plays, and books from Project Gutenberg, among other sources.
- Watson could understand questions posed in natural language and return answers that directly answer the question.

2011

- Watson’s main innovation was centered on its ability to quickly execute hundreds of algorithms to simultaneously analyze a question from many directions, find and score potential answers, gather additional supporting evidence for each answer, and evaluate everything using natural language processing.
- Watson would give a confidence level based on the coincidence among algorithms answers.
- If the confidence level was high enough, Watson was programmed to buzz in during a game of *Jeopardy!*. If not, Watson wouldn’t buzz. Watson performed all of these calculations in an average of about three seconds.
- Source: hiips://www.ibm.com/history/watson-jeopardy

2011

- An unsupervised NN classifies a cat for the first time.
- Google engineer Jeff Dean met a Stanford computer science professor named Andrew Ng and had the idea of building a large neural net, giving it massive computing power using Google's server resources, and feeding it a massive data set of images.
- The built NN ran across 16,000 server processors.
- They fed it 10 million random, unlabeled images from YouTube.
- The neural network processed the image data for three days. It then returned an output containing three blurry images depicting visual patterns it had seen over and over in the test images—a human face, a human body, and a cat. The event also marked the start of the Google Brain project.

2012

- Geoffrey Hinton and two PhD students built a computer vision neural network model called AlexNet to compete in an image recognition contest called ImageNet.
- The task is to process millions of test images and identify them with the greatest possible accuracy.
- AlexNet won the contest with an error rate of less than half that of the runner-up. In only 15.3% of cases was the correct label not in AlexNet's top five most-likely answers. The previous best score had been 26%.
- This victory was decisive to kicked off the current renaissance in deep neural networks,
- Hinton, along with fellow AI Yoshua Bengio and Yann LeCun, won the Turing Prize for 2018.

2016

- AlphaGo defeats human Go champion.
- In 2013, researchers at a British startup called DeepMind published a paper showing how they could use a neural network to play and beat 50 old Atari games. Google bought the company for $400 million.
- Several years later, DeepMind’s scientists developed a neural network model called AlphaGo that was designed to play Go, and learn by playing.
- The software played thousands of games against other AlphaGo versions, learning from both its winning and losing strategies.
- AlphaGo defeated the greatest Go player in the world, Lee Sedol, four games to one in a series of games in March 2016.

2020 - present

- Development of generative AI: DNNs with the ability for AI to generate text, images, and videos in response to text prompts. These networks learn from massive amounts of diverse content (documents, photos, and more) from across the internet.
- OpenAI built a generative pre-trained transformer (GPT) that became the architectural foundation for its two first language models GPT-1 and GPT-2, which were trained on billions of inputs.
- In 2020 the large language model GPT-3 was released; it was trained on 175 billion parameters, (GPT-2 used 1.5 billion).
- In 2021 OpenAI created and released DALL-E, a text-to-image model. The first iteration of DALL-E used a version of OpenAI's GPT-3 model and was trained on 12 billion parameters.
- In 2022, OpenAI released the AI chatbot ChatGPT based on the LLM GPT-3 .
- 2023 was a milestone year in terms of generative AI. Not only did OpenAI release GPT-4, which again built on its predecessor's power, but Microsoft integrated ChatGPT into its search engine Bing and Google released its GPT chatbot Bard.

The State of the Art besides LLMs

- Autonomous Planning Systems, *eg* in space exploration.
- Autonomous Control (vehicles)
- Medical diagnosis, knowledge-based and ML systems.
- Logistics planning, *eg* in large scale problems (hours *vs.* weeks).
- Robotics, *eg* assistance in microsurgery.