1

On the Origins of Artificial Intelligence

1.1. The birth of artificial intelligence (AI)

1.1.1. *The 1950s–1970s in the United States*

Alan Turing's article, published in 1950 [TUR 50], which is one of the founding works in the field of AI, begins with these words: "I propose to consider the question, 'Can machines think?'"

In 1955, "Logic Theorist", considered to be the first AI program, was developed. This work was the result of cooperation between three researchers: a computer scientist (John Shaw) and two researchers from the humanities and social sciences (Herbert Simon and Allen Newell) [SIM 76]. The application was programmed using IPL language [STE 63], created within the RAND and the Carnegie Institute of Technology[1] (a project that received funding from the US Air Force). Here we have the essential elements of AI research: a multidisciplinary approach, bringing together humanities and technology, a university investment and the presence of the military. It is important to note that although the program is described today as the first AI code, these three researchers never use the expression "artificial intelligence" or present their software as falling into this category. The expression "artificial intelligence" appeared in 1956, during a series of seminars organized at Dartmouth College by John McCarthy (Dartmouth College), Claude Shannon (Bell Telephone Laboratories), Marvin Minsky (Harvard University) and Nathaniel Rochester

1 Today, Carnegie Mellon University. In 1956 the Carnegie Institute of Technology hosted its first IBM computer. However, it was not until 1986, 30 years later, that the university established the School of Computer Science.

(IBM Corporation). The aim of this scientific event was to bring together a dozen or so researchers with the ambition of giving machines the ability to perform intelligent tasks and to program them to imitate human thought.

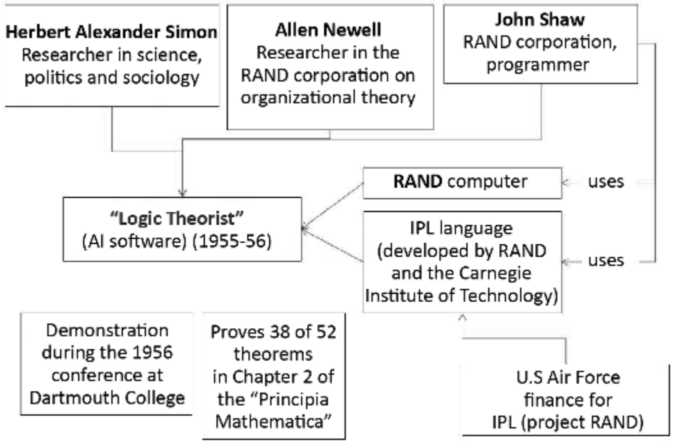

Figure 1.1. *The first artificial intelligence computer program "Logic Theorist", its designers, its results*

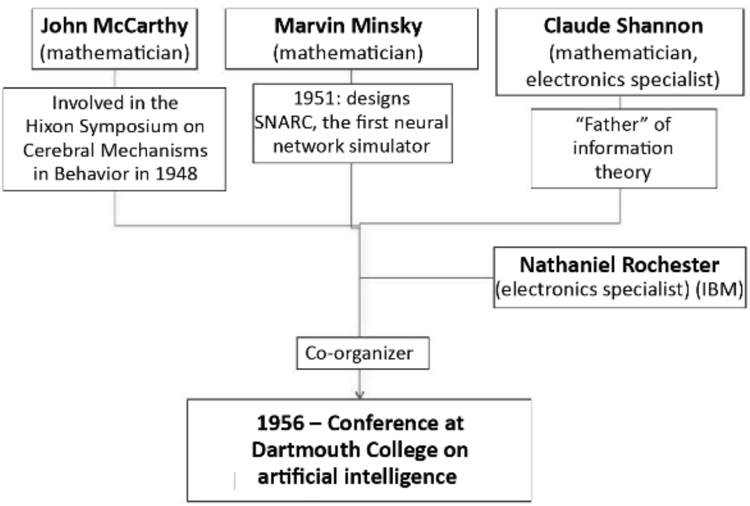

Figure 1.2. *The organizers of a "conference" (two-month program) at Dartmouth College on artificial intelligence in 1956*

While the 1956 conference was a key moment in AI history, it was itself the result of earlier reflections by key players. McCarthy had attended the 1948 *Symposium on Cerebral Mechanisms in Behavior*, attended by Claude Shannon, Alan Turing and Karl Lashley, among others. This multidisciplinary symposium (mathematicians, psychologists, etc.) introduced discussions on the comparison between the brain and the computer. The introduction of the term "artificial intelligence" in 1956 was therefore the result of reflections that had matured over several years.

The text of the proposal for the "conference" of 1956[2], dated August 31, 1955, submitted for financial support from the Rockefeller Foundation for organizing the event, defines the content of the project and the very concept of artificial intelligence:

> "The study is to proceed on the basis of the conjecture 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. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."

The project was more successful than expected because 10 people did not participate, but 43 (not including the four organizers)[3], including Herbert Simon and John Nash. This audience was composed almost entirely of North Americans (United States, Canada), and two British people. In any case, it was entirely Anglophone.

In an article titled "Steps toward artificial intelligence" [MIN 61], Marvin Minsky described, in 1961, these early days of AI research and its main objectives:

> "Our visitor[4] might remain puzzled if he set out to find, and judge these monsters for himself. For he would find only a few machines (mostly 'general-purpose' computers, programmed

2 Full document available at: http://www-formal.stanford.edu/jmc/history/dartmouth /dartmouth.html.

3 List published at http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html.

4 Marvin Minsky imagines an alien visitor arriving on Earth for the first time and discovering our computers.

> for the moment to behave according to certain specifications) doing things that might claim any real intellectual status. Some would be proving mathematical theorems of rather undistinguished character. A few machines might be playing certain games, occasionally defeating their designers. Some might be distinguishing between hand-printed letters. Is this enough to justify so much interest, let alone deep concern? I believe that it is; that we are on the threshold of an era that will be strongly influenced, and quite possibly dominated, by intelligent problem-solving machines. But our purpose is not to guess about what the future may bring; it is only to try to describe and explain what seem now to be our first steps toward the construction of 'artificial intelligence.'"

AI research is structured around new laboratories created in major universities. Stanford University created its AI laboratory in 1963. At MIT, AI was handled within the MAC project (Project on Mathematics and Computation), also created in 1963 with significant funding from ARPA.

From the very first years of its existence, the Stanford AI lab has had a defense perspective in its research. The ARPA, an agency of the US Department of Defense (DoD), subsidized the work through numerous programs. The research topics were therefore influenced by military needs, as in the case of Monte D. Callero's thesis on "An adaptive command and control system utilizing heuristic learning processes" (1967), which aimed to develop an automated decision tool for the real-time allocation of defense missiles during armed conflicts. The researcher had to model a missile defense environment and build a decision system to improve its performance based on the experiment [EAR 73]. The influence of the defense agency grew over the years. By June 1973, the AI laboratory had 128 staff, two-thirds of whom were supported by ARPA [EAR 73].

This proximity to the defense department did not, however, condition all its work. In the 1971 semi-annual report [RAP 71] on AI research and applications, Stanford University described its prospects as follows:

> "This field deals with the development of automatic systems, usually including general-purpose digital computers, that are able to carry out tasks normally considered to require human intelligence. Such systems would be capable of sensing the

> physical environment, solving problems, conceiving and executing plans, and improving their behavior with experience. Success in this research will lead to machines that could replace men in a variety of dangerous jobs or hostile environments, and therefore would have wide applicability for Government and industrial use."

Research at MIT in the 1970s, although funded by the military, also remained broad in its scope. Presenting their work to ARPA in 1973, the researchers felt that they had reached a milestone that allowed them to envisage real applications of the theoretical work carried out until then. But these applications cannot be reduced to the field of defense alone:

> "The results of a decade of work on Artificial Intelligence have brought us to the threshold of a new phase of knowledge-based programming – in which we can design computer systems that (1) react reasonably to significantly complicated situations and (2) perhaps more important for the future – interact intelligently with their operators when they encounter limitations, bugs, or insufficient information."[5]

A few leads for new lines of research are then rejected:

> "We believe that AI research can show the way to computer-based information systems far more capable than have ever been available. We plan to attack the area of Information Retrieval, both for traditional data-base and library problems and for more personal management information systems problems. The new Information Systems should help to increase the effectiveness of individuals who are responsible for complicated administrative structures, as well as for complex information problems of technical kinds. In particular, the services will be available and designed to be useable over the ARPANET, and will be designed to interact with the personal

5 MIT, Artificial Intelligence Laboratory, Proposal for research on intelligent automata and micro-automation 1974–1976, Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under Contracts N00014-70-A-0362-0003 and N00014-70-A-0362-0005, available at: http://people.csail.mit.edu/bkph/AIM/AIM-299-OCR-OPT.pdf.

systems of other individuals to recognize conflicts, and arrange communication about common concerns."

Along with university teams, the RAND Corporation also played a central role in the emergence of AI in the United States. Willis H. Ware's book, *RAND and the Information Evolution* [WAR 08], is an invaluable resource for understanding the role the organization played in the development of AI in the United States as early as the 1960s. The book covers the period 1946–2008, which is divided into two periods, one from 1946 to 1983, during which research within the agency was organized into departments, before RAND reorganized itself around programmatic actions.

In 1963, Ware recalls, of the 20 contributions that comprise the first book on AI, published by Feigenbaum and Feldman, six were written by RAND researchers. The initial work of Allen Newell and Cliff Shaw, both RAND scholars, in collaboration with Herbert Simon (Carnegie Institute of Technology) laid several foundations for AI research on learning, proof of theories and knowledge representation, among other ideas. Ware also reminds us that AI work does not develop in isolation. Their researchers built on advances in computer science, including intensive uses of new computers such as JOHNNIAC, until the mid-1960s. AI research draws on advances in computer science and many other disciplines.

The history of AI may seem America-centric. But we cannot forget that work in robotics and computer science, in an attempt to understand how the brain works, all converging on AI, mobilized researchers well beyond the North American sphere at the same time.

In the United Kingdom, Grey Walter's research, as early as the 1940s, became part of this international academic movement which was interested in the modeling of brain processes and the definition of intelligence. Grey Walter designed the "turtles" in Bristol in 1947, considered to be the first autonomous electronic robots (Luce Langevin describes the 1950s as a period when an "electronic animal menagerie" was built) [LAN 69]. These approaches were underpinned by the belief in a strong resemblance between the brain and the machine: "Physiologists admit as a working hypothesis that the functioning of the brain is that of a machine" [ASH 51].

Today's international competition between major powers is not unrelated to the history of research and early development from the 1950s to the 1960s. China and Russia, in particular, did not wait until the 2000s or 2010s to invest in the field of AI research. Their present activity in this area is based on a model which, as in the case of the United States, is several decades old.

1.1.2. *AI research in China*

Artificial intelligence research in China began in the late 1960s [XIN 83].

According to Wang Jieshu's analysis [JIE 18], the history of Chinese AI is closely linked to the history of the Soviet Union and the close relationship between the two countries. In the period 1970–1983, the main areas of research covered a broad spectrum of issues, such as:

– machine translation, a field which, in 1982, includes some of the following achievements:

- development of ECTA (English–Chinese Automatic Translation System) software,

- development of the English–Chinese Title Translation System,

- JF-111, A Universal Machine Translation System;

– natural language understanding;

– theorem proving;

– expert systems;

– robotics.

Nothing in this enumeration really distinguishes the orientation of Chinese research from its Western counterparts. Again, the approach here is multidisciplinary (mathematics, computer science, linguistics, medicine, automation, robotics, aeronautics, etc.).

In the 1980s, China already had a large number of publications, achievements, researchers and universities involved in AI research. Jieshu's article only mentions civil applications, nothing is said about the military's position on this research topic and its investment in universities.

On the basis of this article, we identify a set of universities involved (Table 1.1).

Name of University	City (Province)
Zhongshan University	Guangzhou (Guangdong)
Jilin University	Changchun (Jilin)
Zhejiang University	Hangzhou (Zhejiang)
Nanjing Technology College	Nanjing (Jiangsu)
Beijing Aeronautical Engineering Institute	Beijing
Beijing Academy of Traditional Chinese Medicine	Beijing
Institute of Automation, Academia Sinica	Beijing
Institute of Linguistics, Chinese Academy of Social Sciences	Beijing
Institute of System Science, Academia Sinica	Beijing
Mathematics Institute, Academia Sinica	Beijing
Qinghua University	Beijing
Science-Technology Information Institute and Computer Technology Institute, Academia Sinica	Beijing
Shanghai Institute of Computing Technology	Shanghai
Shenyang Institute of Automation, Academia Sinica	Shenyang (Liaoning)
Wuhan University	Wuhan (Hubei)

Table 1.1. *Universities involved in AI between 1970 and 1983 in China (classified by city). Reconstructed from [XIN 83]*

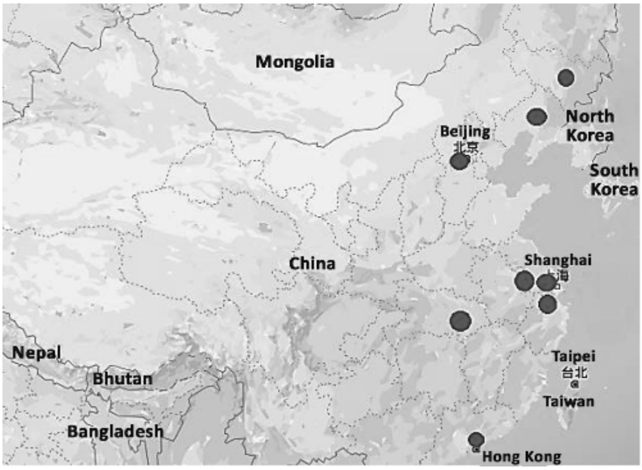

Figure 1.3. *Geographic distribution of Chinese universities investing in AI between 1970 and 1983. Reconstructed from [XIN 83]*

AI research in China thus took shape at the same time as it did in the West and has been structured around universities of excellence. This history serves as a basis for China's current ambitions, which are still expressed in research programs as well as in economic, industrial, societal and political projects.

1.1.3. *AI research in Russia*

The history of AI in Russia follows roughly the same chronology as that of the United States or the West. As early as the early 1960s, Western delegations visiting Russia noticed the presence of research teams on AI themes.

The report by E.A. Feigenbaum, who visited the USSR from June to July 1960 as a member of the American delegation to the *First Congress of the International Federation of Automatic Control* (IFAC) [FEI 60] said:

> "The program consisted of a number of welcoming speeches, and an address by the well-known scientist and Chairman of the USSR National Committee for Automatic Control, V.A. Trapeznikov."
>
> "The Soviet Deputy Premier talked on the problems which automation would bring to 'certain societies' which were not well equipped to handle this kind of technological change – change which would bring unemployment, re-education problems […]"
>
> "In general, Soviet papers could be characterized as oriented toward theory, while papers of Western delegates mixed theory and application."
>
> "In conjunction with the conference, various research institutes, educational institutions, and plants were officially opened for technical excursions by the delegates […] By far, the most popular tour was one to the Institute of Automation and Telemechanics in Moscow."

In the Soviet Union, AI was one of the components of cybernetics, in the same way as information theory, computer science or the study of military C2 [LEV 64]. Cybernetics, which appeared in the USSR in 1953, was a new and broad field, organized around various research communities which come together, in particular, at conferences dedicated to cybernetics and in numerous academic publications from the early 1960s. Military cybernetics became a sub-domain of cybernetics.

An article published in the journal *Science* on August 27, 1965 [KOP 65] introduced a new city of science, which had just been built, in Siberia: Akademgorodok, located in the suburbs of Novosibirk (Siberia).

The work of Paul R. Josephson [JOS 97] gives a whole chapter to the history of the birth of AI in the city of Akademgorodok, in the middle of the Soviet period (1960–1970). For it was there that AI in Russia was born. A Russian research community centered on AI was created there in the 1970s, with a university research center, "clubs" ("Akademia" club on Artificial Intelligence)[6] and a Council for AI (the Artificial Intelligence Council)[7], etc.

The city, now considered one of Russia's Silicon Valleys (along with Moscow and St. Petersburg), is said to be home to Russia's "cyberwar soldiers" [CLA 17]. Akademgorodok is home to a technopark, concentrating 24% of the revenue of all Russian technoparks, and 22% of the companies hosted in Russia in technoparks [LOG 16]. The Akademgorodok technocenter is currently reported to host 340 companies, 115 start-ups and nearly 10,000 employees. This ecosystem is complemented by the many university research centers that have made the city famous and unique, due to their high concentration.

In the mid-1970s, the Soviet Union envisaged the use of networked information technology as a tool for controlling, managing and planning for the Soviet economy. The project envisaged at that time was to link major production and political centers using a vast network of computers. Moscow would be the hub, but it would also pass through Leningrad (as St. Petersburg was called at the time) and Kiev. Implementation was to start around 1975 and be fully operational by 1990. Western, American (Control Data Corporation) and British (International Computers, Ltd.) companies were even involved in this project [LIE 73]. The computer and computer science, in the broad sense, was a tool of the Soviet political project, as well as posing a challenge to America which at that point faced difficulties in implementing such networks on a large scale. Soviet technological development was based

6 http://ershov.iis.nsk.su/en/archive/subgroup?nid=763577id_1=763577.

7 http://ershov.iis.nsk.su/en/archive/subgroup?nid=763551id_1=763551.

on a policy of transfer from the United States to the USSR from the end of the 1950s and accelerated from 1972 onwards under the Nixon administration [ROD 81]. The acquisition of foreign technology, especially in the field of information technology, by the USSR, was carried out through legal (sales authorized by the US government) and illegal (black market and copying) channels. AI was part of this Soviet "cybernetics" project. However, a report by the American Department of Defense in 1990 estimated that the Soviet Union had lower capabilities than America, despite research efforts and special attention to AI applications in the civil and military fields [DOD 90]:

> "The Soviet Union lags behind the United States significantly in machine intelligence and robotics. They do have a good theoretical understanding of the area and can show creativity in applying the technology to selected space and military applications. Soviet R&D on artificial intelligence (AI), under the auspices of the Academy of Sciences of the USSR, includes work on machine vision and machine learning. The value of machine intelligence to battlefield operations as well as to the domestic economy has been recognized by the Soviet government."

So, while the Soviet Union does not appear to have been truly competing with US capabilities at the time, it was nonetheless a player that added to the competitive landscape facing the United States. AI and robotics, in their various dimensions (research, development, industrialization), emerged dynamically in several countries and regions of the world: the report cites France, Europe and Japan, among others.

Research in the USSR was not isolated from the rest of the world. The USSR organized international AI conferences: for example, in Leningrad in April 1977[8] and in October 1983[9]. Its research projects and achievements were in fields relatively similar to the rest of the world: applications for automatic translation and understanding natural language (in 1982, the "Etap-1" project was created – Experimental System for Automated Translation from French into Russian)[10].

8 http://ershov.iis.nsk.su/en/archive/subgroup?nid=763546id_1=763546.

9 http://ershov.iis.nsk.su/en/archive/subgroup?nid=763550id_1=763550.

10 http://ershov.iis.nsk.su/en/archive/subgroup?nid=763438id_1=763438.

Figure 1.4. *Location of Akademgorodok (district of the city of Novosibirsk)*

1.1.4. *AI research in Japan*

The history of AI research in Japan began in the 1960s at Kyoto University, with a team formed around Professor Toshiyuki Sakai, who worked in three areas: computer vision, speech processing and natural language. In 1970, the team presented the first facial recognition system at the Osaka exhibition. During the decade, several teams took shape, at Kyushu University (around Professor Toshihiko Kurihara, with work on kanji-kana conversion systems), at the University of Osaka (with Professor Kokichi Tanaka on knowledge processing issues), at the University of Tokyo and in corporate laboratories such as NTT. However, the emergence of these various teams did not yet constitute a true Japanese AI research community. This took shape around a community of students at the University of Tokyo, the IAUEO, from which the pioneers of Japanese AI such as Hideyuki Nakashima, Koichi Hori and Hitoshi Matsubara emerged.

The government provided several hundred million dollars in funding for long-term university research programs [DON 77]. In 1971, the government launched the PIPS (Pattern Information Processing System) project, which was funded at $100 million over eight years. The computational requirements necessary to achieve the objectives of PIPS necessitated the development of new electronic chips. This phase was financed by the Japanese government as early as 1973, through a new project.

Other substantial funding was mobilized to support research in image processing, speech and natural language processing.

AI seemed to really take off in the mid-1980s [MIZ 04, NIS 12][11], after Japan launched its fifth-generation computer program, in 1982. The Institute for New Generation Computer Technology (ICOT) was the R&D arm of the national fifth-generation computer science program. Among others, ICOT produced the programming language KLIC and the legal reasoning system HELIC-II. In 1983, the ICOT had less than 50 researchers and the research themes, focused around the central project of designing the world's largest computer by 1990, directly concerned AI (e.g. automatic translation systems, automatic response systems, understanding speech, understanding images, problem-solving systems and logic programming machines, etc.).

In 1986, the Japanese Society for Artificial Intelligence (JSAI) was founded. In 1990, the association created the Pacific Rim International Conference on Artificial Intelligence (PRICAI), which aims to structure AI research in this part of the world, thus complementing or counterbalancing the Western initiatives of the IJCAI and the European Conference on Artificial Intelligence.

Alongside university research, industrial R&D activities have made a major contribution to the development of Japanese artificial intelligence. Major industrial groups set up teams dedicated to AI (NTT, Hitachi, Fujitsu, Sony, Honda, etc., are big names in the industry that invested in this field). Some of their developments received a lot of media coverage, such as AIBO (Sony), ASIMO (Honda), TAKMI – Analysis and Knowledge Mining (IBM Research Tokyo), facial recognition tools and oral translation applications for mobiles (NEC) and the humanoid robot HRP-4C (capable of singing and dancing).

Since the 1980s, the Japanese government has maintained its investment in AI, but international competition is now raging and several major powers are outperforming Japan in terms of numbers of scientific publications: while China published 41,000 articles in the field between 2011 and 2015, the United States published 25,500 and Japan 11,700. This phenomenon has

11 This is Toyoaki Nishida's interpretation of AI history. Other analyses consider that the organization of the *International Joint Conference on Artificial Intelligence* (IJCAI) in Tokyo in 1979 was the first milestone in this history (R. Mizoguchi). See [MIZ 04] and [NIS 12].

been repeated in the industrial field: the United States is said to have a thousand companies in the field, while Japan has 200–300[12].

Although Japan has been and remains one of the leaders in robotics and AI applied to this field, it is not only because it shows qualities as an integrator of the multiple aspects required (electronics, electricity, computing, automation, etc.), but also because it is a leader in many of these fields; Japan has particularly been a leader in electronics. Robotics in the 1980s required more skills in electronics, microelectronics and mechanics than in computer science.

1.1.5. *AI research in France*

In the 1950s and 1960s, interest in intelligent machines spread to many countries around the world. France was one of the players in this internationalization of research. For example, we can cite the following:

– Pierre de Latil's reflections on artificial thinking [DEL 53];

– the work of Albert Ducrocq (the son of a soldier, he studied political science and electronics and was later a journalist and essayist, and qualified as a cybernetician) who invented the "electronic fox", an autonomous robot on wheels, and who inspired the achievements in the 1970s of Bruno Lussato, inventor of zebulons, autonomous computerized handling robots. Albert Ducrocq published several works dealing with robots, weapons and AI such as *Les armes de demain* (the weapons of tomorrow) (Berger-Levrault, 1949), *Appareils et cerveaux électroniques* (electronic devices and brains) (Hachette, 1952), *L'ère des robots* (the age of robots) (Julliard, 1953) and *Découverte de la cybernétique* (discovery of cybernetics) (Julliard, 1955). He published many other works before his death in 2001;

– writings on thinking machines by Paul Chauchard [CHA 50], Paul Cossa [COS 54], Louis Couffignal [COU 52], Dominique Dubarle [DUB 48], or on the robot, with the writings of Albert Béguin [BÉG 50].

Questions are tackled from a variety of viewpoints: mathematicians, cyberneticians, philosophers, electronics engineers, etc. Are humans machines or robots? Can the brain be reproduced in a machine? Can the machine think, does it have a soul? What is a machine? Can we reduce the mechanism of

12 List of AI companies in Japan available at: https://www.data-artist.com/en/contents/ai-company-list.html.

thought or the functioning of the brain to algorithms? Are the brain and the body simple mechanics?

Louis Couffignal [COU 52] defined the machine as "an entire set of inanimate, or even, exceptionally, animate beings capable of replacing man in the execution of a set of operations proposed by man".

He listed the categories of machines: machines that can add and write, machines that can read and choose, calculating machines and thinking machines.

In 1984, Jean-Pierre Laurent [LAU 85], professor at the University of Chambéry, drew up a series of observations on the state of French AI research: identifying research potential was a difficult exercise because of the imprecise perimeter of AI, a research discipline in its own right for some, adding techniques with various applications to others; on the other hand, AI research was dispersed across France – the teams were small and spread over the whole country; AI suffered from a rather negative image in the world of academic research and in political decision-making circles, due to the overly empirical nature of the discipline, which was moving away from purely theoretical research, and the lack of convincing and spectacular results, capable of overcoming the failures of the 1960s. AI was excessively associated with "gadget" applications, but at the beginning of the 1980s, the situation seemed to be changing, and developments in the industry using expert systems seemed likely to give AI the serious image it lacked. Nevertheless, the industry still had only an imprecise vision of what AI could bring: "Some industrial organizations that are not really familiar with AI and its possibilities are inclined to believe in miracles and ask AI to solve everything (how, for example, can an expert system be built if there are no human experts?)." Researchers investing in the discipline were still poorly supported and had very limited means in terms of human resources, equipment and funding. This situation had an impact on university teaching, which attached only minor importance to AI training. The situation of academic research therefore depended on a wide range of variables: public interest, the attention of political decision-makers, industrial investment and the ability to unite isolated researchers and disciplines.

It should also be noted, when reading this study by Jean-Pierre Laurent, that the drivers of research funding depended on exogenous variables. Although the State, via its research organizations, such as the CNRS, had been funding some AI projects since 1979, the publication of the Fifth Generation MITI Project (Japan) in 1981 seemed to be a new decisive moment in political action.

1.2. Characteristics of AI research

From the outset, AI has been formed with a multidisciplinary approach, located at the crossroads of a body of research involving multiple scientific disciplines. Philosophy occupies an important place in this interdisciplinarity (H.B. Dreyfus [DRE 72] establishes links between AI and the thinking of Plato and Hobbes), alongside mathematics, computer science and cognitive sciences in particular.

> "The field draws upon theoretical constructs from a wide variety of disciplines, including mathematics, psychology, linguistics, neurophysiology, computer science, and electronic engineering." [LAW 84]

In his March 2018 report, French MP Cédric Villani recalled this interdisciplinary approach:

> "The field is so broad that it cannot be limited to a specific area of research […] AI is at the crossroad of multiple fields: computer science, mathematics (logics, optimization, analysis, probabilities, linear algebra, etc.), cognitive science, etc." [VIL 18]

Spain, in its official National Strategy Document for AI (2019), does the same: "Due to the growing complexity of its contributions, it is increasingly interdisciplinary, with synergies with biology, philosophy, the world of law, psychology, sociology and economics."[13]

AI uses a wide range of techniques or technologies:

> "each algorithm in AI is supported by a mix of techniques: semantic analysis, symbolic computing, machine learning, exploratory analysis, deep learning and neural networks […]" [VIL 18]

---

13 Spanish RDI Strategy in Artificial Intelligence, Ministry of Science, Innovation and Universities, 2019, available at: http://www.ciencia.gob.es/stfls/MICINN/Ciencia/Ficheros/Estrategia_Inteligencia_Artificial_EN.PDF.

> "These technologies include: natural language processing, smart agents, computer vision, machine learning, expert systems, autonomous cars, chatbots and voice recognition."[14]

> "Six Al technologies have significant potential for application to production and logistics processes: expert/knowledge-based systems, natural language, speech recognition, three-dimensional vision, smart robotics and neural networks. The most significant of these technologies thus far is expert systems." [MEL 89]

> "Technologies and research areas generally considered to be sub-domains of AI: • Automated Planning and Scheduling • Computer Vision • Decision Support, Predictive Analytics, and Analytic Discovery • Distributed Artificial Intelligence/Agent-based Systems • Human Language Technologies • Identity Intelligence • ML • Process Modeling • Robotics/Autonomous Systems." [DNI 19]

Finally, this new science is characterized by the extent of its field of application: in the 1970s, research tried to apply AI to the games of chess and Go, to mathematics (proving theorems), to understanding and translating written and spoken language, to controlling robots and artificial hands, to image analysis and to human–machine interaction (an intelligent dialogue)[15]. In the 1980s, there was talk of applications in:

> "bioengineering, chemistry, computer hardware, computer software, education, engineering, general purpose systems and AI utilities, law, manufacturing and industry, mathematics, medicine, the military, and resource exploration." [LAW 84]

14 Mauritius Artificial Intelligence Strategy, report, A report by the working group on artificial intelligence, November 2018, available at: http://mtci.govmu.org/English/Documents/2018/Launching%20Digital%20Transformation%20Strategy%20191218/Mauritius%20AI%20Strategy%20(7).pdf.

15 *Cryptolog* Magazine, NSA, July 1975, available at: https://www.nsa.gov/Portals/70/documents/news-features/declassified-documents/cryptologs/cryptolog_11.pdf.

Three decades later, the Villani report [VIL 18] evokes an equally broad field of applications: image and video recognition, translation, content recommendation, etc. The Villani report [VIL 18] is a report on the use of the Internet for recognizing images and videos.

The application spectrum seems infinite; the only limits are those of the imagination. Some applications that seemed quite innovative in 2020 are, already in an experimental phase or were planned in the early days of AI:

– AI in cars [TSU 79][16];

– image analysis, detection of objects in images [NAG 79];

– conflict simulators (e.g. CONSIM, *Conflict Simulator*, in 1971 (*Proceeding ACM ‘71* [CLE 71][17]).

Globally, research and development in AI have the common aim of understanding “how human cognition works by creating cognitive processes that emulate those of human beings” [VIL 18].

16 “This paper describes an automobile with artificial intelligence, which consists of a road pattern recognition unit and a problem solving unit. The vehicle is completely autonomous and can be driven without a human driver. The road pattern recognition unit involving a pair of TV cameras and a processing unit identifies obstacles in front of the vehicle and outputs data regarding the locations of the obstacles. The problem solving unit is a microcomputer system and determines control optimal to the environment around the vehicle based on the data. The algorithm employed in it is a table-look-up method, in which the location of the optimal control is addressed in the table by key words generated from the data. The table was heuristically made by means of digital simulation. The vehicle was successfully driven under various road environments at the speed below 30 Km/h.”

17 “The development and use of CONSIM (Conflict Simulator), a computer program designed to heuristically simulate decision making, is described. A typical example is used to evaluate and analyze the methodology of CONSIM. The simulation model is designed for two political opponents, each possessing any finite number of alternatives. The model is constructed utilizing techniques frequently employed in game theory. Probabilities are assigned, using a Bayesian Approach, to sets of alternatives available to the United States and China. Incorporated into the model is the capability of varying probabilities as the Vietnam Conflict evolves and the re-evaluation of risks, costs, and benefits occurs. CONSIM is easy to use and applicable wherever probabilities may be assigned to each alternative in a mutually exclusive and exhaustive set of alternatives in a dynamic situation.”

Definition of AI objectives/finalities	Source	Year
"The goal of AI research is to construct computer programs which exhibit behavior that we call 'intelligent behavior' when we observe it in human beings."	[FEI 63]	1963
"The long-term goal of general AI is to create systems that exhibit the flexibility and versatility of human intelligence in a broad range of cognitive domains, including learning, language, perception, reasoning, creativity, and planning."	National Science and Technology Council. Networking and Information Technology Research and Development Subcommittee (NITRD) (United States) [NET 16]	2016

Table 1.2. *The objectives or goals of AI*

Writing a history of AI would require taking into account not just the course of one discipline but the meeting of many. Its birth is a moment of convergence, of encounter. The destiny of AI was and will remain closely linked to this multidisciplinarity which is one of its main characteristics.

1.3. The sequences of AI history

The course of AI history would be punctuated by a succession of advances, phases of mania and periods of setbacks, sometimes called "AI winters". The path of AI is, in any case, not linear. While AI is today a worldwide success (we are in "the third era of artificial intelligence" today [MIA 18]), a new phase of retreat , which could occur rather abruptly cannot be excluded. The success of AI is linked to several conditions. Its destiny does not depend only on the progress of science, but it is linked to economic challenges and investments and to political support as well as to scientific progress.

Dividing the history of AI into several sequences makes it possible to position, within each of them, events, such as the main conflicts that have never ceased to punctuate the history of humanity, political periods and major events, in particular, events in the history of technology.

Period 1956–1970 (Spring 1): during this first period, pioneers of AI were confident in their ability to achieve the goals they had set themselves, assured that the state of science would allow for the rapid realization of AI. These pioneers were optimistic, believing that AI would have achieved its goals within the next 25 years. But difficulties quickly arose and initial ambitions had to be reconsidered. Progress was slower than expected, and many problems were still too difficult to solve for the knowledge of the time. Machine translation systems, for example, did not perform at all well. The Lighthill Report (1972) put an end to this first period for a while. Excessive negative discourse brought about the end of this dynamic period, full of hopes and utopias:

> "ARPA, the Advanced Research Projects Agency of the Defense Department, is in fact the major source of funding for AI research, and is at present under attack from several directions because its projects are not considered to be 'paying off' as rapidly as they should be for the expenditures involved. The more pragmatic of Dreyfus' criticisms, even though they may be dismissed out of hand by AI workers, are echoed by others not so easily brushed aside, for example J. Lighthill (1973) speaking for the Science Research Council in England. If enough critics succeed in discouraging ARPA research, and AI work in general, some of the potential benefits for NSA that might have been around the corner may never materialize."[18]

Period 1970–1982 (Winter 1): in the early 1970s, the so-called "AI winter" began, designating the years during which research had to persevere and obtain significant results while facing a reduction in funding. The efforts of the first period were not in vain, because, in the failures suffered, researchers understood what they had to work on.

In the 1970s, the immediate objective was no longer the creation of a universal, general intelligent machine, but of a specialized intelligence: a particular machine should be able to act intelligently in a limited domain[19]. The almost utopian hopes of AI pioneers, assured of being able to quickly transpose the mysteries of human intelligence into the machine, quickly gave way to more measured debates on the performance of the machine. It was a question of what computers could do or could not do, of their limits [DRE 72].

---

18 *Cryptolog* Journal, NSA, July 1975, available at: https://www.nsa.gov/Portals/70/documents/news-features/declassified-documents/cryptologs/cryptolog_11.pdf.

19 *Cryptolog* Journal, March–April 1986, NSA, available at: https://www.nsa.gov/Portals/70/documents/news-features/declassified-documents/cryptologs/cryptolog_101.pdf.

France appears to have been marked by strong skepticism in the 1970s and 1980s. The State invested little at the time, and the research community was still small:

> "There is a great deal of activity in research in computer vision which, as with speech understanding, is a difficult area, requiring enormous amounts of computer power. Fortunately, the indications are that the power required for many applications will be available and economically feasible in about ten years. Researchers, therefore, are currently attacking parts of the total problem with the presently available resources, leaving the integration of their efforts to the future." [HOL 76]

Period 1982–1987 (Spring 2): Japan's program for a new generation of computers revived the race for AI research, as did Ronald Reagan's so-called "'Star Wars' speech". The great nations embarked on a technological race, so as not to be totally overwhelmed and dominated by Japan. It was in the efforts for AI and the conquest of space that the struggle for power between states on the international scene was manifested. In the 1980s, expert systems proliferated. AI became commercial, and for the military, AI at last offered the prospect of applications to reinforce operational efficiency. Funding was arriving in a massive way.

Period 1987–2010 (Winter 2): the positive phase of the early 1980s was short. The AI market began to show signs of a significant slowdown as early as 1987. The specialized AI hardware industry (LISP machines) collapsed.

Period 2010-now (Spring 3): a new phase of worldwide interest in AI. The period since 2011 onwards can truly be called an "AI Spring"[20]. This period of renewal can be rooted in the conjunction of three variables[21]: Big Data, algorithms and increased computer power. This is a claim that the Finnish authorities repeat:

> "More effective and lower-cost computing capacity, vast increases in the amount of data that can be used by AI, and more advanced AI algorithms have all led to more intensive use

---

20 National artificial intelligence strategy for Qatar, January 2019, available at: https://qcai.qcri.org/wp-content/uploads/2019/02/National_AI_Strategy_for_Qatar-Blueprint_30Jan2019.pdf.

21 *Idem*.

> of artificial intelligence. In fact, we are experiencing a new spring of artificial intelligence."[22]

These three variables can also be called Big Data, Machine Learning (ML) and Cloud Computing [MIA 17].

Many hopes are once again placed on the progress and contributions of AI, there is massive state and private funding, multi-year research and development programs, and implementation in all sectors of activity where it is possible to do so. Performances well below expectations would undoubtedly be likely to slow down the current global dynamics. This would then result in a drop in funding, a lack of public interest and a decline in the commercial activities of AI products. The general or weak AI that is deployed today carries within it the seeds of its own limitations, and human-like systems, faithfully imitating intellectual capacities, are still out of reach:

> "[…] every aspect of our lives will be changed in some way. Today the progress of these intelligent machines seems limitless. Will they surpass us, or even replace us? If by definition a super-intelligence is capable of surpassing us in all areas of intelligence then our own efforts will become obsolete by then. In some ways these machines have superhuman abilities but in fact there is no machine out there that is as intelligent as a rat."[23]

Further disappointments would lead to a new downturn.

This breakdown of AI history since the 1950s is imperfect, as it reduces all aspects of AI into one. Research cycles are not necessarily the same for industry, military programs, policy or the various application areas. Research has not really gone through a period of hibernation. It has continued to be carried out throughout the world with, of course, support in terms of funding

---

22 Leading the way into the age of artificial intelligence, Final report of Finland's Artificial Intelligence Programme 2019, Publications of the Ministry of Economic Affairs and Employment, available at: https://julkaisut.valtioneuvosto.fi/bitstream/handle/10024/161688/41_19_Leading%20the%20way%20into%20the%20age%20of%20artificial%20intelligence.pdf?sequence=4.

23 Documentary broadcast on Arte.fr in September 2018, *L'intelligence artificielle va-t-elle nous dépasser?*, directed by Guillain Depardieu and Thibaut Martin. Video available at https://www.arte.tv/fr/videos/069097-000-A/l-intelligence-artificielle-va-t-elle-nous-depasser/.

and resources that vary according to the period and the country. The volume of publications, the efforts of the scientific community to structure itself and organize research around national and international associations and international symposia bear witness to this.

States' efforts in AI R&D are mainly motivated by the desire to build a technological field that provides them with economic power, offering the multiple sectors of activity tools to promote their progress, and in the military field, instruments useful for increasing their capabilities. However, during the decades from 1950 to 2000, there was no open discussion of security applications and issues, even though intelligence agencies were able to use AI to contribute to the accomplishment of their missions.

1.4. The robot and robotics

AI is closely associated with robotics. AI was established as a discipline and a scientific community in the 1950s. Specialist publications, symposiums, associations and the creation of research teams have come to participate in the construction process. AI has integrated the robot into its research questions. But it has not joined forces with robotics, which is a discipline in its own right. Robotics has its own history, its own challenges and its own research questions.

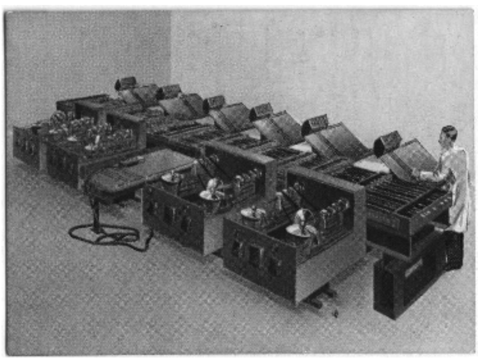

Figure 1.5. *Manchester University Robot*[24]*. For a color version of this figure, see www.iste.co.uk/ventre/artificial.zip*

24 George Arents Collection, The New York Public Library, Manchester University robot, available at: http://digitalcollections.nypl.org/items/510d47da-93ab-a3d9-e040-e00a18064a99. This illustration shows the differential analyzer designed by Vanevar Bush.

In the 18th and 19th Centuries, several achievements made by automatons were presented to the public. These were machines that took on the appearance of human beings and reproduced tasks such as writing or playing an instrument. There was, of course, no autonomous intelligence in these machines. The intelligence was that of their ingenious designers, who created precise mechanics. These machines took their place in the literature and reflected the technologies of the time: not yet called robots, they were formed of electricity, mechanics, steel and sometimes steam energy (see *The Steam Man of the Prairies* by Edward S. Ellis [ELL 68]).

Mechanization would win in the 20th Century; the industry opening the way for large-scale production (e.g. Ford factories in the 1910s). What had until then been automatons, machines and mechanics became “robots”. Karel Capek introduced the term in 1921 in a play, designating machines with a human appearance. Isaac Asimov wrote, in 1941, the short story *Liar!* [ASI 41], in which he introduced the “laws” of robotics. In the following decades, science fiction would give a substantial role to the robot and would contribute to its global success. But science fiction also contributed to forging a particular image of the robot: it is often human in appearance, it is, in any case, endowed with life and an intelligence that defies that of human.

During World War II, unmanned aircraft equipped with explosives and thrown at enemy populations and armies were called “robot bombs”. After the war, other developments were created with industrial applications. The first programmable arm was designed in 1954 by George Devol and Joe Engleberger, which was integrated into assembly lines at General Motors in 1962. Robots occupied the industrial space, the space of science and were designed to assist humanity in their experiments, and to confront hostile environments. In 1994, the Robotics Institute at Carnegie Mellon University designed the eight-legged robot DANTE II, which explored the crater of the Spurr volcano. Its shape was inspired by the shape of a spider. This achievement combined sensor technology, computer science, electronics and communication instruments.

Robotics, a scientific discipline, came about either in 1956 or 1964:

– 1956, when George Devol and Joe Engelberger met. From their collaboration, the robotics industry, a magazine (*Industrial Robot*) and an international conference (*International Symposium on Industrial Robotics*) were born. Joe Engelberger is sometimes referred to as the “father of robotics”;

– 1964, when Bernie Roth, a mechanical engineer, met AI researcher John McCarthy. This was the starting point for a fruitful collaboration, scientific publications, developments and regular scientific meetings that gradually solidified a new research community.

A third option of the date of the creation of this discipline dates back to telerobotics (developed in the 1940s in the United States to manipulate nuclear materials). Goertz’s articles in 1952 and 1954 were the starting point for the history of robotics[25].

Whatever the precise date chosen, the birth of robotics is, generally, contemporary to that of AI (1950s).

Robotics, as a discipline, has rubbed elbows with AI since its origins, but it has developed as a discipline in its own right:

> “[…] the scientific task of understanding intelligence is already well under way in artificial intelligence, a subpart of computer science. If an appropriate direct scientific field already exists, why is robotics needed in addition? Why isn’t artificial intelligence the supplier of basic scientific capital to robotics, along with other suppliers, such as computer science, mechanical engineering and electrical engineering?
>
> There is an answer and it further illuminates my basic thesis. Artificial intelligence currently shares with computer science a special view – it considers information processing divorced from energetics. This creation of an interior milieu, in which only information processes occur, is a powerful abstraction, one which helped computer science to emerge by permitting it to focus on the essential mechanisms. But the costs of the abstraction show nowhere more clearly than in the unexplored central problem of robotics – controlled perceptually coordinated motion.” [NEW 81]

25 These three versions are available in [MAS 12].

The work of the Stanford AI Research Center in the 1970s integrated R&D in robotics:

> "We are pursuing fundamental studies in several areas: problem solving, perception, automated mathematics, and learning. However, we find it productive to choose as a goal the creation of a single integrated system […] Our system consists of a mobile robot vehicle controlled by radio from a large digital computer. The principal goal is to develop software for the computer that, when used in conjunction with the hardware of the vehicle, will produce a system capable of intelligent behavior […] Before we changed computers (at the end of 1969), our robot system had achieved a primitive level of capabilities: It could analyze a simple scene in a restricted laboratory environment; plan solutions to certain problems, provided that exactly the correct data were appropriately encoded; and carry out its plans, provided nothing went wrong during execution. Therefore, when we began planning a new software system for controlling the robot from a new computer, we set more difficult short term goals: The system is to be able to operate in a larger environment, consisting of several rooms, corridors, and doorways; its planning ability must be able to select relevant data from a large store of facts about the world and the robot's capabilities; and it must be able to recover gracefully from certain unexpected failures or accumulated errors. We have not yet accomplished these goals." [RAP 71]

Like AI, robotics was built around an American center. However, robotics and robots have met with great success in many other countries, including Japan, a country that is particularly active in research and development in the field: in 1997, the first robot tournament in Japan (RoboCup) took place; in 1999, Sony presented and marketed AIBO, the first pet robot-dog; in 2000, Honda created the humanoid robot ASIMO. Japan has been particularly involved in robotics and has been a world leader in recent decades, ahead of the United States and other industrialized nations. In 2010, Japan accounted for two-thirds of the world's robot production (although not all of this robotics is AI-related). This success is largely due to the fact that Japan has, on its territory, all the technological and industrial skills and building blocks necessary for robotics (without being a leader in all of these fields: mechanics, electronics, computing, sensors, etc.) [NAR 10]. The efforts made in robotization can also be explained by specific societal needs: a lack of manpower and an aging population that no immigration can compensate for. It is therefore necessary to replace humans by machines in production, and to think about solutions to support elderly populations. Robotics is not a priority issue for national defense, especially as, since the end of World War II, the country no longer has the right to have its own army. Investments are therefore oriented in fields other than defense.

Robots are machines (AI is software):

– that are programmed;

– that perform tasks independently or semi-autonomously;

– that interact with their physical environment through sensors.

Some robots are also machines that are controlled by operators (so there is no question in their case of autonomy). The “non-intelligent” robot, as opposed to the robot that is intelligent, is only capable of repetitive tasks: taking and moving an object. For this action, it does not need human intervention, so it is, in a way, “autonomous”, but it is difficult to qualify it as “smart”. Autonomy is therefore not synonymous with intelligence. However, this robot, which only repeatedly places objects, becomes “smart” when a camera is attached to it, which allows it to distinguish objects, differentiate them and assign a different action according to the type of object. Vision and perception require AI algorithms [OWE 17]. Intelligent robots are robots that have embedded AI programs.

1.5. Example of AI integration: the case of the CIA in the 1980s

Computers entered the agency (the CIA) in the early 1960s, primarily for administrative tasks. Then, the computer emerged as a tool that could be used for the very mission of the agency, i.e. intelligence, and the collection and processing of data[26].

In the 1980s, the computer became an essential part of the agency’s activities, and it had to deal with increasing volumes of data that had to be collected, stored, analyzed and distributed within its departments[27]. The

26 https://www.cia.gov/library/readingroom/docs/CIA-RDP90G00993R000100130005-7.pdf.

27 ILO 0751-86, August 29, 1986, declassified, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP90G00993R000100130005-7.pdf.

agency was confronted with constraints imposed by data, computers and the network: large quantities of data were produced and constant and rapid flows connected machines and individuals. The speed and masses of data were two of the variables put forward by observers of the time. However, the findings did not call the technologies into question, as these seemed to bring only added value to the running of the agency and intelligence officers. The introduction of technology into the agency is portrayed as a dynamic process. In the mid-1980s, AI had not yet been adopted but the agency launched into a process of reflection which mobilized its various services and managers to envisage applications in the intelligence field.

On the other hand, the increase in the speed and density of data production by the agency (in the form of reports, notes and various documents) was not only due to technology. There was a political and governmental will within the agency to produce more knowledge on a wide range of issues every day (800 reports were produced by the CIA in 1983; but half of the staff were focused on the USSR [DDI 84]).

The treatment of AI by the American intelligence community can be analyzed using resources available online. Thus, the CIA offers a set of archival[28] documents on its site. In this corpus of data, we are looking for information that allows us to describe the relationship that the intelligence community has with a science and a set of innovative technologies, the way in which the institution organizes itself to tackle these new objects (creation of groups, designation of leaders, creation of relationships, networks, monitoring, etc.) and the applications envisaged for AI for the specific needs of intelligence, all from a chronological perspective (at what points does the agency take up the subject, does it act, what are the important milestones in this evolution, etc.). This facet of the history of the CIA has long remained unknown and in the shadows:

> "The CIA's role in the application of science and technology to the art of intelligence is far less appreciated […] However, the exploitation of science and technology has been a significant element of the CIA's activities, almost since its creation. In 1962, it resulted in the creation of the Deputy Directorate of

28 Data accessible from: https://www.cia.gov/library/readingroom/. A search on the term "artificial intelligence", using the engine available on the page, directly identifies the corpus useful for our analysis.

> Research, which was succeeded in 1963 by the Deputy Directorate for Science and Technology (renamed the Directorate of Science and Technology in 1965).” [RIC 01]

Moreover, if this story takes into account the CIA’s policies of acquisition and ownership of technology, the relationship must also be seen in the opposite direction, that of the influences of intelligence on the academic research community and industry. The CIA cannot be considered as a simple user “client”: “It is also responsible for a number of scientific advances” [RIC 01].

1.5.1. *The CIA’s instruments and methods for understanding and appropriating AI adapted to its needs*

The observations formulated by the intelligence community, indeed both the defense and the intelligence communities, were strongly marked in the early 1980s by a trend aimed at developing and experimenting with smart computer systems for data processing [ECK 84]. Defense and intelligence officials observed an imbalance between, on the one hand, an ever-growing mass of data and, on the other hand, human resources that are or will inevitably be overwhelmed by the task. The AI perspective emerged as a source of possible solutions to what was a managerial challenge and not yet essentially a security issue. The option of reversing the trend by reducing the volume of data collected, produced or provided to analysts was not considered. The only scenario is that of increasing data volumes, accelerating data flows and overwhelming human capacity with a system of machines that needs to be compensated for:

> “A central objective of Artificial Intelligence is to enable computers to emulate the intellectual functions that humans employ in analyzing and interpreting data. For these reasons the Intelligence Community is interested in assessing the potential of artificial intelligence as a technological key to computer-based smart systems for intelligence applications.” [ECK 84]

The particular interest in AI in the early 1980s was part of a special framework, a key moment in the history of the CIA, which partly recovered budgets it had lost over the previous decade. This renewal of means allowed the agency to envisage a new phase of growth, to recruit, to develop its activities, to widen its field of expertise and to increase its productivity and research activities:

> "The last four years have seen remarkable growth in the Intelligence Community's budget […] it has enabled us to restore many of the capabilities that we lost in 1970s […] In the last four years, we in the CIA for the first time have developed and implemented a comprehensive research program covering a staggering number of countries and issues […] we have rebuilt our human intelligence capability […] On the technical side, investments in new technical systems are beginning to pay off." [DDI 84]

This modernization would increase data collection capacity, which would result in new analytical needs:

> "This means a tremendous increase in the quality and precision of intelligence information we can collect. It also means that the volume of data will substantially increase. This suggests that we must begin to invest in processing systems to match our capability in collecting raw data […] We have to avoid being overwhelmed by this volume of information from all sources, and that is where developments in computers, software, and other new technologies really must help us." [DDI 84]

Although the obstacle of financial resources seems to have been removed, the implementation of new systems was no less complex. One of the main difficulties lay in defining objectives and the appropriate methods. The low level of competence of the decision-makers (managers within the agency) was also pointed out, with its consequences on the choices made:

> "A serious problem for both American industry and government is the computer illiterate senior manager […] important decisions on computers and ADP equipment are made by managers who hardly know a mainframe from a Mack truck. What do we do? We turn to the computer specialists in our own organizations who think narrowly and protect their turf and for whom larger scale planning, networking, and experimenting is anathema. And so we sometimes develop inadequate systems that can't talk to each other and meet only today's needs." [DDI 84]

Other particularities must be taken into account, such as the agency's structure and mode of operation, which are reflected in how technological systems are organized and which are sometimes referred to as "problems":

> "The major problems in the use of computers at CIA revolve around compartmentation and security. Unlike organizations of similar size in the private sector, we have to have a system that operates on a need-to-know basis, and that may involve only a handful of people." [DDI 84]

Security imperatives create particular constraints, which the systems must integrate: "We must protect ourselves at the same time against 'hackers' from the outside, and the possibility of 'moles' from the inside" [DDI 84].

The CIA, in the first half of the 1980s, structured its computer system into five independent systems: one system for analysts, one for operations, one for administration, one for personnel security and one for processing and analyzing technical intelligence data. Each system has its own software.

In 1983, the CIA drew up an action plan, the *CIA Strategic Plan* (1983–1993) [FIT 83], which sets out the areas that would specifically be the subject of an intensive development policy: AI is one of these areas, alongside human resources, information handling, space, crisis management, terrorism, secret communications, clandestine information gathering, arms control monitoring and protecting intelligence.

It is in this context of renewal that the agency was interested in AI (at the crossroads of what was then the beginning of the second wave of AI and the fifth generation of computers), in which its managers place significant hopes:

> "We will rely on AI in expert systems applications to enable us to detect indicator anomalies for warning, to synthesize combinations of data for analysis, to scan mail to pick out critical messages, or to pick out gaps in our knowledge. Applications of AI in processing huge quantities of raw data without having to translate raw data into standardized formats, as we now do should help separate the wheat from the chaff, especially in SIGINT and imagery […] More sophisticated simulation and modeling techniques will increase our ability to predict alternative outcomes of future events. AI should help analysts compare dissimilar forms of data – imagery, SIGINT, regular text […] Another application might involve accessing more data on a real time basis, especially in crises." [DDI 84]

There are, however, voices that call for the utmost caution. In the face of the hopes placed in AI's ability to profoundly transform intelligence capabilities, one should be aware of the long road ahead to achieve such achievements:

> "progress in the area of AI, in our view, is likely to be painfully slow. Promises of quick advances with practical applications should be treated with some skepticism." [DDI 84]

The Director of Intelligence (DDI), in a speech to the members of the Security Affairs Support Association, stressed the need for the institution to resist commercial pressures, marketing arguments and promise sellers:

> "I can't tell you how many contractors have tried to sell me software that will enable me flawlessly to predict the next action of the Soviet leadership. Until we can understand more about how the intuitional process works, it seems to me, it will be difficult to write 'expert system software' that can duplicate what analysts do." [DDI 84]

Skepticism is even coupled with a certain irony in the way these new technologies are viewed, at least an awareness of their limitations: "Today's Expert Systems are 'idiot savants', demonstrably useful in strictly circumscribed applications" [ECK 84].

Intelligence expectations of AI in 1983–1984 were quite similar to those of the late 2010s: detection, prediction, raw data processing, image and signal processing, etc.

We can also observe that the AI systems that have been developed since then will not have fulfilled all their promises: expert systems and other modalities for processing multisource data in heterogeneous formats will not have made it possible to predict the end of the USSR, to anticipate and prevent the attacks of September 11, 2001, and many other major actions or events that have since marked the history of both the United States and the world.

The attention paid by both the intelligence community and the military to scientific advances in the field of intelligence, however, predates the 1980s. The military, through ARPA/DARPA, supports the work of American academics, while both the CIA and the Defense Department also monitor global AI research, as evidenced by reports of missions abroad in the 1960s in the context of international colloquia or visits to organizations[29].

Type of action	Precise title	Theme/object	Date	Location
Participation in or organization of symposiums, with contributions from civil experts (industrialists, universities)	*Symposium of Artificial Intelligence Applications to Intelligence Analysis*[30]	AI applications Informing analysts and CIA officials: what is AI, how can it be useful for analysts?	November 30 to December 1, 1982 December 6–8, 1983 (500 people attended) March 19–21, 1985 March 13–14, 1986[31] 1987[32] (*Fifth Intelligence Community Artificial Intelligence Symposium. At the Defense Intelligence Agency, Bolling Air Force Base*)	CIA Headquarters (Washington D.C.); DIA (Bolling Air Force Base) (1985); NBS (Gaithersburg)/ NSA (Fort Meade); etc.

29 United States Intelligence Board, Committee on Documentation, Trip report: CODIB visit to Germany, October 11, 1962, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP80B01139A000200120018-0.pdf.

30 Science and Technology Division, Office of Scientific and Weapons Research, Directorate for Intelligence, *Symposium of Artificial Intelligence Applications to Intelligence Analysis*, November 18, 1982, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00024R000500300004-9.pdf.

31 Director of Research and Development, Announcement of AI Symposium, January 29, 1986, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP89B00297R000400870001-7.pdf.

32 Procurement Management Staff, Weekly report, Period ending on March 31, 1987, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP89-00063R000300280042-9.pdf.

Type of action	Precise title	Theme/object	Date	Location
Focus groups, coordination	Artificial Intelligence Steering Group (AISG)[33]	Reflections on the potential of AI Regular meetings (June 18, 1984 at NCARAI premises)[34]	1983 (chair: Phil Eckman)	CIA
Focus groups, coordination	Artificial Intelligence Applications Working Group (AIAWG)[35]	Coordinate R&D efforts within the agency to adapt AI to specific intelligence problems	1983	Information Systems Board (CIA)
Relations with American universities	Request for training in artificial intelligence [MAL 87]	Train a member of staff from the agency; represent the intelligence agency in the university institution (integrate students in training; or place "professors" in universities)	1987	Carnegie-Mellon University
Training	Training setup by ISGF	Two- and three-day seminars to introduce AI and its applications in intelligence	1983	Information Science Center (CIA)
	Information technology video seminar. (*first use of the headquarters cable TV system in the CIA*) [DON 86]	*Seven-part program on AI*	1986	Agency via the headquarters cable TV system

Table 1.3. *Types of actions initiated by the CIA to understand and integrate AI*

33 Created by the IRDC (Intelligence Research and Development Council).

34 NCARAI: Navy Center for Applied Research in Artificial Intelligence.

35 Information Systems Board, Artificial Intelligence applications working group, undated, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00142R000100030011-8.pdf. Charter of the Artificial Intelligence Applications Working Group of the CIA Information Systems Board, July 1983, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00142R000100050002-6.pdf.

1.5.2. *Focus groups, research, coordination*

The creation of working, reflection, research and coordination groups, in charge of organizing training initiatives and dialogue with the united industrial world, is neither an approach specific to the CIA, nor to intelligence in general, nor to the theme “AI”. In the same year, 1983, the DIA and the NSA, for example[36], also created working groups dedicated to AI. The function of each of these groups was to think about AI for the specific context of each agency (DARPA approaches AI for defense purposes, the CIA for intelligence, etc.). The AISG (Artificial Intelligence Steering Group), within the CIA, observed, in 1984, that defense invests much more in AI research than the intelligence community does, and that a rapprochement should be envisaged in order to capitalize on the advances made by defense[37].

As soon as it was implemented, the ISGF organized meetings which took the form of visits to institutions where AI work was being carried out: the National Bureau of Standards (NBS)[38] was visited on August 27, 1984, and the MITRE Corporation on May 10, 1984[39]. Through these meetings, the members of the ISAG tried to identify projects with potential applications in the field of intelligence.

Within the CIA ISB (Information Systems Board), several work streams were created, including one dedicated to AI (Artificial Intelligence Applications Working Group – AIAWG). During its 1983 meetings, one of the objectives was to identify the CIA’s AI needs or the opportunities that AI could offer the agency. The method for identifying these needs was to ask each department within the agency what AI applications they would find useful. This approach assumes that the interlocutors have a fairly precise knowledge of the possibilities of AI: “it was interesting to see common

---

36 Memo for USDRE from C/AI Steering Group, August 29, 1984, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP86M00886R000500040004-2.pdf.

37 *Idem*.

38 Artificial Intelligence Steering Group, Meeting Minutes, August 27, 1984, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP86M00886R000500050002-3.pdf.

39 Artificial Intelligence Steering Group, Trip report, May 10, 1984, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP86M00886R000500050006-9.pdf.

needs appearing in several offices, but each couched in the language and terminology of their respective application area”[40].

The draft of the summary note of the meeting of December 14, 1983 does not ask any questions about the reasons for this convergence. Perhaps this was due to a real convergence of the needs of the various services consulted. But perhaps it was also the product of a common collective representation of what AI is and what it could provide. The only real differences identified during this exercise remained in vocabulary of the formulations. Furthermore, the requirements were formulated in too general a manner: “Having identified these general requirements, it would appear that we have additional work to do in becoming more specific.”[41]

1.5.3. *The network of interlocutors outside the intelligence community*

The integration of AI into the practices and policies of the intelligence community required the involvement of a wide range of external actors.

This approach was neither specific to the new subject of AI nor specific to the intelligence community alone:

> “This country is unique in the ways in which government and private industry work together […] In the U.S., we rely on a combination of patriotism and profit motive to make our system work […] New technological devices and new analytical techniques that enable us to understand growing threats to U.S. and its people are based on the synergistic nature of the relations between private industry and government.” [DDI 84]

The CIA’s approach to AI consisted of keeping abreast of the state of research and development, identifying actors – academic, military and industrial – and engaging in dialog with them in order to understand how available AI could be transferred to the intelligence community. The CIA engaged in dialog with interlocutors outside the intelligence community on a

40 AIAWG Meeting Minutes, December 14, 1983, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00142R000100050004-4.pdf.

41 AIAWG Meeting Minutes, December 14, 1983, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00142R000100050004-4.pdf.

wide range of topics. The agency remained attentive, ready to respond and alert. For example, in 1983, some 1,200 analysts attended more than 500 conferences on a wide range of subjects [DDI 84].

The CIA was considering adapting knowledge and technology to its needs, and before even thinking about acquiring it (the documents do not specify the methods of transfer envisaged), tried to understand what was at stake:

> "The IASG of the IR&DC is in the process of preparing recommendations for Council consideration concerning where promising Artificial Intelligence technology developments are occurring and which may be of value for adaptation into the Intelligence Community."[42]

This logic guided the entire strategy of acquiring and integrating technology within the CIA, which aimed to save its financial resources and adapt existing technologies to its needs. Thus, faced with requests for acquisition whose justifications are not fundamentally questioned, the director of the CIA wrote, in 1982:

> "Before commenting on the recommendation to acquire a VHSIC/VLSI facility, the Director and I would like the Council to explore a little more exactly what would be involved. Our first blush reaction is prompted by fear that the cost may be too high for the Community to bear. Rather than fund a facility of this nature, possibly we could piggyback on existing commercial facilities to explore circuits which might be unique for intelligence programs […] We believe that the low powered, high density storage technology […] should be pursued by other DoD and commercial programs as well, thus permitting the Intelligence Community to capitalize on developmental work elsewhere."[43]

42 Proposed TDY Travel to Boston, October 25, 1985, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP87M00220R000800930013-8.pdf.

43 IR&DC report, Technology considerations for improved intelligence capabilities, December 1, 1982, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP87B00305R000801500001-4.pdf.

In 1987, the Deputy Director of Administration of the CIA expressed the desire to place a staff member at Carnegie-Mellon University for a one-year training program to represent the agency technically and politically; the second option proposed was to place a "professor" in residence for two years, with the term in quotation marks:

> "[…] the Agency should have someone at Carnegie-Mellon University who has an ADP background. […] Please […] arrange for an appropriate officer to go to Carnegie-Mellon for a year's training in the field of artificial intelligence or arrange to place an officer at Carnegie-Mellon for two years as a 'professor' in Residence […] It should be an individual who we believe can represent us both technically and politically […] Let's work at this seriously and aggressively." [MAL 87]

	Industry	**Universities**	**Policy/Government**
AI Symposium, 1982[44]	Xerox, Symbolics, DEC (Digital Equipment Corporation), VERAC Corporation, System Development Corporation, Artificial Intelligence Corporation, Texas Instruments Corporation, Hughes Research Laboratory	Yale University, University of Pittsburgh, Syracus University	Office of Research and Development (ORD)
Training		Placement of CIA agents as students or professors in American universities	

Table 1.4. *Non-intelligence actors that the CIA solicits and mobilizes and with whom it maintains various types of relationships*

44 Conference program available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00024R000500300005-8.pdf. The program was organized around 14 interventions spread over two days. Four were dedicated to a historical approach to AI research; an introduction and an institutional conclusion; the rest of the days were devoted to technical presentations and industrial demonstrations.

Background	**Dates**	**Vocabulary**
Symposium on Artificial Intelligence[45]	1982	Artificial intelligence systems Deductive reasoning Responding to *ad hoc* questions "Smart" data basis Expert systems Target tracking Ocean surveillance Information retrieval
Symposium on Artificial Intelligence	1985	Coping with increasing problems of explosive growth of information
Symposium on AI Applications in the Intelligence Community [ECK 83]	December 8, 1983 [ECK 84] (2nd edition of the symposium, the first was held in 1982, with 300 participants). Symposium sponsored by the ISGF	AI to: - merge, analyze and interpret large volumes of data - increase analyst productivity Natural Language Processing Warning indicators Automated signal processing Expert systems Automated image interpretation

45 The program of this conference is available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00024R000500300005-8.pdf.

Background	Dates	Vocabulary
AIAWG	December 14, 1983[46]	Expert systems Information retrieval, analysis and reporting Networking planning, surveillance and maintenance Consultation, training, analysis and prediction Image interpretation **Natural language** Interface to databases Text retrieval Machine translation Message recognition Message dissemination Speech understanding Speaker recognition Language identification **Image understanding** Optical character recognition Change detection Identifying, counting and classifying objects **Intelligent databases** Inferential reasoning Data resource navigation aids

46 AIAWG Meeting Minutes, December 14, 1983, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00142R000100050004-4.pdf.

Background	Dates	Vocabulary
AIAWG	December 14, 1983[47]	**Signal interpretation** Real-time TEMPEST surveillance Signal sorting and classification **AI tools and environment** Improvement to information system development process Improved user interfaces Voice input to machines AI system development aids Executive decision aids Administrative system support (budget, travel, personnel)
AI Steering Group[48]	August 29, 1984	Expert systems Natural Language Processing Intelligent database interfaces Image understanding Signal interpretations Geographic and spatial data management Intelligent workstation environments Knowledge acquisition process

47 *Idem*.

48 Memo for USDRE from C/AI Steering Group, August 29, 1984, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP86M00886R000500040004-2.pdf.

Background	Dates	Vocabulary
Mid-Level Managers' Seminar on Artificial Intelligence[49]	September 11–13, 1984	LISP Semantic networks Production systems Image processing/pattern recognition Image understanding Automating deduction and inference Natural Language Processing Expert systems

Table 1.5. *AI vocabulary within the CIA*

The document is supplemented by a copy of an article published on August 14, 1986, in *The New York Times,* presenting a three-year joint research program between Carnegie-Mellon and IBM. The project was financed by the industrialist to the tune of several million dollars.

The university did not have a candidate imposed on it by the intelligence agency. The Computer Science Department of the university, when asked about the possibilities of hosting a student from the agency, showed itself willing to engage in such an initiative, but outlined conditions: the candidate must have a Master's degree and a prior knowledge of AI. On the other hand, a financial counterpart was requested by the university department [DTE 86].

1.5.4. *What AI applications for what intelligence needs?*

The applications of AI are, for the intelligence community, one of the central points justifying the attention given to these sciences and their developments:

> "The symposium accomplished two things. First, it presented an update on artificial intelligence technologies […] and most

49 *Seminar on Artificial Intelligence*, Information Science Center, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP89B01330R000600960014-2.pdf.

> importantly, it focused on how these new technologies could be applied within the intelligence community."[50]

The intelligence community is not intended to provide *leadership* in research and focuses on the application dimension. What is in demand is an operational, functional and practical AI.

When the working groups were set up within the CIA in 1982–1983, the first findings stressed that AI was still only at the stage of developing potential applications, not producing operational systems. However, industry efforts could not succeed without a precise definition of intelligence needs. It was therefore a question, at this point, of comparing the potentialities and bringing them to the attention of the intelligence community, who then had to imagine the domains which could benefit from the contributions of AI, then formulate their needs:

> "The contracting parties have started to build teams of qualified staff; they have started to invest in the necessary IT infrastructure; and they are just beginning to understand our specific needs. These efforts will necessarily remain timid until we better understand our own needs."[51]

At the conclusion of one of its meetings in December 1983, the AIAWG produced a list of possible applications of AI in intelligence. This list[52] identified six main areas: expert systems, natural language, image understanding, smart databases, signal interpretation and AI tools and environment.

---

50 *Artificial Intelligence Symposium*, March 28, 1985, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP87M00539R000600730001-2.pdf.

51 Memo for USDRE from C/AI Steering Group, August 29, 1984, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP86M00886R000500040004-2.pdf.

52 AIAWG Meeting Minutes, December 14, 1983, available at: https://www.cia.gov/library/readingroom/docs/CIA-RDP85-00142R000100050004-4.pdf.