¿Will machines ever rule the world?

¿Las máquinas dominarán el mundo?

PEDROZA, Mauricio 1; VILLAMIZAR, Gustavo 2; MENDEZ, James 3

Contents

ABSTRACT:

From ancient automatons, to the latest technologies of robotics and supercomputing, the continued progress of mankind has led man to even question his future status as a dominant species. "Will machines ever rule the world?" or more precisely: Do we want machines to rule the world? Theoretical and technological challenges involved in this choice will be considered both under the conservative approach of "machines designed as tools of man" and in the scenario of "machines thought as similar to man".
Keywords: Artificial Intelligence, Cognitive Machines, Narrow Artificial Intelligence Strong Artificial Intelligence, Artificial General Intelligence.

RESUMEN:

Desde los antiguos autómatas hasta las últimas tecnologías de la robótica, el continuo progreso de la humanidad ha llevado al hombre a cuestionar incluso su estatus futuro como especie dominante. "¿Alguna vez dominarán las máquinas el mundo?" O más precisamente: ¿Queremos que las máquinas dominen el mundo? Los retos teóricos y tecnológicos que se plantean en esta elección serán considerados tanto bajo el enfoque conservador de "máquinas que sirven al hombre" como en el escenario de "máquinas equiparables al hombre"
Palabras clave: Inteligencia Artificial, Máquinas cognitivas, Inteligencia Artificial Débil, Inteligencia Artificial Fuerte, Inteligencia Artificial General.

1. Introduction

If popular culture has taught us anything, according to Dean (2006) it is that the power of machines will increase to the point of being perceived as a threat to humanity. Bill Gates (2007) recently observed that: "the emergence of the robotics industry … is developing in much the same way that the computer business did 30 years ago" (p. 58), and according to Lin et al (2011):

If the evolution of the robotics industry is analogous to that of computers, we can expect important social and ethical challenges to rise from robotics as well. Robots are often tasked to perform the 'three Ds', that is, jobs that are dull, dirty, or dangerous … We can also expect robots to scale down as well as up: Some robots are miniature today and ever shrinking, perhaps bringing to life the idea of a "nano-bot", swarms of which might work inside our bodies or in the atmosphere or cleaning up oil spills. Even rooms or entire

buildings might be considered as robots—beyond the 'smart homes' of today— … With synthetic biology, cognitive science, and nanoelectronics, future robots could be biologically based … Again, much of this speaks to the fuzziness of the definition of robot: What we intuitively consider as robots today may change given different form-factors and materials of tomorrow (p. 944).

It is therefore natural that this set of facts and perspectives will lead to questions about the role of machines in the future of human being. This paper in particular seeks to present the differences in the conception and development of "truly" intelligent machines compared to the current state-of-the-art automatons, by approaching a controversial question: "Will machines ever rule the world?".

It is sought to show that the main lines of development in artificial intelligence (AI) are not currently focused on obtaining an adaptive, cognitive and autonomous system (being these just some of the main characteristics of an entity considered as "truly" intelligent), a system that does not require a restricted domain of operation to be successful, a system with an intelligence that we could call "general". In words of experienced AI researchers such as Peter Voss (2002): "… yet very little work is being done to specifically identify what general intelligence is, what it requires, and how to achieve it" (§ 3). Voss invite us to review our research approach, either to continue the development under the perspective of a sophisticated task automation (Narrow AI) or the construction of a human-like cognitive system (Artificial General Intelligence).

Finally, it will be argued that if we do not want machines to rule the world maybe Narrow AI is the way, also, a state-of-the-art of current technologies in robotics and automation will be presented under this conservative approach of "machines designed as tools of man". On the other hand, If we want machines to rule the world and see fulfilled many of the numerous science fiction fantasies that accompany the fact that they were able to do it, this article quotes diverse theoretical conceptions and technological developments that would be necessary in the field of Artificial General Intelligence (AGI) to pass from the generic computing device, or what Philip D. Carter (2012) called: "the new breed of thing that is language and logic mechanized" (p. 2), to a scenario of machines thought as similar to man, considering it as a parallel line of evolution of intelligent entities in the universe, or going a little further as the sociologist Lewis Mumford (1970): "it is our essential nature to transcend the limits of our biological nature <and to be ready if necessary to die in order to make such transcendence posible> "(p. 434).

2. (Narrow) Artificial Intelligence vs (General) Artificial Intelligence

In the words of AI researcher Peter Voss (2002): "Intelligence can be defined simply as an entity's ability to achieve goals – with greater intelligence coping with more complex and novel situations. Complexity ranges from the trivial – thermostats and mollusks (that in most contexts don't even justify the label 'intelligence') – to the fantastically complex; autonomous flight control systems and humans" (§ 2). This particular view reveals that what is considered "intelligent" must by definition deal with certain levels of minimal complexity to avoid being classified as a simple automatism.

In this same line of thought, according to philosopher Fred Dretske (1993) the approach to the concept of intelligence can be realized from two perspectives:

One can think of it as something like money, something almost everyone has, but some have more of than others. Or one can be thinking of it as more like wealth - something possessed by only those who have (at least) more than the average amount of money … For if intelligence is like money, questions about the possibility of artificial intelligence are questions about the possibility of machines with some mental capacity - the amount being irrelevant. If, on the other hand, intelligence is understood in a comparative way, as mental wealth, the possibility of artificial intelligence then becomes the possibility of building machines that can win games, not just machines that can play games (p. 201).

Artificial General Intelligence (AGI) aims to develop a human-level or greater intelligence

inside an artificial structure and it is also known as Strong AI. However, the vast majority of developments in artificial intelligence to date are considered as advances in the field of Narrow AI (also known as weak artificial intelligence), understood as the capacity to gather information, processing and response in a limited or narrow domain of knowledge, which results in the execution of predetermined tasks within a very specific field of operation. On the other hand, adaptability, autonomy and response to new and problematic situations is a trait that many only attribute to those we consider "intelligent beings", and it is this particular characteristic what artificial general intelligence aims to develop, and its main difference vis-à-vis the specialized task execution approach (Narrow AI).

Returning to the original question: "Will machines ever rule the world?" the academic Seth D. Baum (2014) explains:

Narrow AI is intelligent in specific domains but cannot reason outside the domain it was designed for. Narrow AI can be quite useful, and can also pose some risks. But it is not expected to take over the world, because controlling the world requires capabilities across many domains. General AI (AGI) is intelligent across a wide range of domains. Humans are also intelligent across many domains, but this does not mean that AGI would necessarily think like humans do. An AGI may not need to think like humans in order to be capable across many domains (p. 5).

3. Machines designed as tools of man

The human being as a member of the animal kingdom has within its biological characteristics the impulse to preserve its functional identity over time, since the cessation of these vital operations is what leads to the concept of death. Throughout history the needs of man have been widening, going from the ancestral requirements of food, shelter and company to a new set of parameters inherent to subsistence and coexistence within modern societies.

Herein we have cataloged this range of needs within three general groups, arguing this way the conception and design of machines as tools to serve these needs.

3.1. We seek welfare & confort

3.1.1. Food. It is well known that the ancestral need for food led historically to the emergence of agriculture, understood as the ability to provide us with food through our knowledge in the manipulation of natural elements. "The role of world agriculture will become increasingly crucial in forthcoming decades, as concerns over food, the environment, and energy increase, in the context of a world population that is predicted to reach 10 billion by the middle of the 21st century" (Murase, 2000, p. 1). With regard to the role of artificial intelligence H. Murase (2000) says:

In agricultural systems the inherent complex, dynamic and non-linear nature of its behavior has required advanced technologies, to provide better understanding, and appropriate solutions. The recent application of the technologies of artificial intelligence (i.e. computer vision, robotics and control systems, expert systems, decision support systems, natural language processing, etc.) and other advanced forms of information technology (neural networks, fuzzy logic, genetic algorithms, and photosynthesis algorithms) promises to provide solutions to problems in agricultural systems (p. 1).

According to I. Farkas (2003) regarding modeling and control in agricultural processes and intelligent control in agricultural automation, the most important recent applications can be grouped as follows:

- Greenhouse technology, e.g. climate control, hydroponics systems, nutrient supply systems and water management systems.

- Environmental and climate control of greenhouses, warehouses and animal houses.

- Post-harvest, e.g. drying systems and control, storage systems control, product quality protection.

- Animal husbandry, e.g. climate control, identification tags, feeder systems, robotic milking.

- Control issues of precision farming, e.g. site-specific operations, positioning, guidance, weed control, crop protection and management systems.

- Energy issues and alternative energy resources in agriculture (p. 1).

Additionally I. Farkas (2003) in his article titled "Artificial intelligence in agriculture" points out some topics of interest, in order to promote the consolidation of bridges between AI and its applications in agriculture and domains connected to it (in particular, environmental sciences):

Applications and practical issues:

- Successful or novel AI decision support approaches for planning, scheduling, control, monitoring, prediction or diagnosis problems in agriculture and related domains.

- Commercialized and routinely used AI-based systems, in particular expert and knowledge-based systems (in laboratories, extension services and farms).

- Critical analysis of project failures, and validation issues for AI-based decision support systems (p. 2).

As evidence of some of these developments we can cite advances in the scheme of fruit recognition for robotic harvest, for example, the performance of the fruit recognition system is influenced by factors such as the species of harvesting fruit, variable light, occlusions and many others. Therefore, the reliability of recognition methods must consider the environment in which the robot is working, and the proper selection of sensors. Regarding this, authors like Zhao et al have worked in two key elements of vision control for fruit harvesting: "fruit recognition and eye-hand coordination" (Zhao et al, 2016, p. 311). Moreover, the development of cognitive architectures for automatic gardening has even led to the operational possibility of specific treatments according to the plant individual needs. For example, the work of Agostini et al presents a cognitive system that integrates AI techniques for decision-making with robotics techniques for sensing and acting to autonomously treat plants using a real-robot platform. Artificial intelligence techniques are used to decide the amount of water and nutrients each plant needs according to the history of the plant. Robotic techniques for sensing measure plant attributes (e.g. leaves) from visual information using 3D model representations. "These attributes are used by the AI system to make decisions about the treatment to apply, and finally, acting techniques execute robot movements to supply the plants with the specified amount of water and nutrients" (Agostini et al, 2017, p. 69).

3.1.2. <u>Environment</u>. The progressive alteration of the natural balance, caused by exploitation and manipulation of resources by the hand of man has shown over time lot of adverse effects, thus promoting new initiatives to reduce the environmental impact of human activity. "However, many of these mitigation activities would imply the loss of human lives, due to the extreme conditions of the scenarios where they must be executed" (Lin et al, 2011, p. 945). Seeking to overcome these limitations today robots perform important functions in environmental remediation, such as collect trash, mop up after nuclear power plant disasters, remove asbestos, cap oil geysers, sniff out toxins, identify polluted areas, and gather data on climate warming.

3.1.3. <u>Medical and healthcare</u>. Factors such as the progressive aging of the population, a life cycle less compatible with the body's organic requirements, and the appearance of psychological alterations derived from a hectic lifestyle can lead to increased understanding of global spending on health and medical care. "This latent need has led to a shortage of human and physical resources to provide care to an increasingly large population, and this is where robotics and artificial intelligence appear as support tools" (Lin et al, 2011, p. 944). Some toy-like robots, such as PARO which looks like a baby seal, are designed for therapeutic purposes, such as reducing stress, stimulating cognitive activity, and improving socialization. Similarly, University of Southern California's socially assistive robots help coach physical-therapy and other patients. Medical robots, such as da Vinci Surgical System and ARES ingestible robots, are assisting with or conducting difficult medical procedures on their own. RIBA, IWARD, ERNIE, and other robots perform some the functions of nurses and

pharmacists, these being just a few examples of the impact of artificial intelligence in the field of health care.

3.1.4. <u>Labor and service</u>. In the search to increase our quality of life we have tried to transfer the daily and monotonous tasks to our synthetic employees. "Nearly half of the world's 7-million-plus service robots are Roomba vacuum cleaners" (Guizzo, 2010, § 1), but others exist that mow lawns, wash floors, iron clothes, and move objects from room to room. However, "the employee label is not restricted only to household chores. Robots have been employed in manufacturing for decades, particularly in auto factories, but they are also used in warehouses, movie sets, electronics manufacturing, food production, printing, fabrication, and many other industries" (Lin et al, 2011, p. 944). Moreover, a general-purpose robot, if achievable, could service many of our labor needs, as opposed to a team of robots each with its own job.

3.2. We seek power & protection

3.2.1. <u>Infrastructure and energy</u>. Artificial intelligence has an important field of application in building energy efficiency. As Wang & Srinivasan (2017) establish:

Researchers have developed various simulation tools to predict building energy use since the early 1990s. These tools can be further classified as engineering method and AI-based method … Compared with engineering methods, AI-based prediction method requires less detailed physical information of the building. There is no need for model developer to have high level of knowledge of the physical building parameters, which in return saves both time and cost for conducting the prediction. On the other hand, there is no explicit relation between the physical building parameters and model inputs, which makes it impossible to extrapolate building energy performance once the design and/or operation of the building has changed.

 (p. 807).

3.2.2. <u>Transport</u>. The need associated with the coverage of great distances has been achieved through advances in science and technology. However, the human error factor very often leads to an immense amount of lives depending on the knowledge and expertise of a single individual (subject to all imperfections of the human sensorimotor apparatus). "It is therefore optimal to transfer this responsibility to artificial mechanisms of greater robustness and precisión" (Lin et al, 2011, p. 945), as examples driverless trains today and DARPA's Grand Challenges are proof-of-concepts that robotic transportation is possible, and even commercial airplanes today are controlled autonomously for a significant portion of their flight today, never mind military UAVs.

3.2.3. <u>Military and security</u>. Despite the advance of civilization, hostility and conflict seem inherent to human nature, so ensuring the minimum conditions for peaceful coexistence often leads to the implementation of surveillance measures, repression, or ultimately the elimination of a potential threat, this explains the appearance of war robots with fierce names such as Predator, Reaper, Big Dog, Crusher, Harpy, BEAR, Global Hawk, Dragon Runner, and more. They perform a range of duties, such as spying or surveillance (air, land, underwater, space), defusing bombs, assisting the wounded, inspecting hideouts, and attacking targets. Police and security robots today perform similar functions, in addition to guarding borders and buildings, scanning for pedophiles and criminals, dispensing helpful informacion, reciting warnings, and more. There is also a growing market for homesecurity robots, which can shoot pepper spray or paintball pellets and transmit pictures of suspicious activities to their owners' mobile phones. As AI advances, we can expect robots to play more complex and a wider range of roles in society: For instance, police robots equipped with biometrics capabilities and sensors could detect weapons, drugs, and faces at a distance. "Military robots could make attack decisions on their own; in most cases today, there is a human triggerman behind those robots" (Lin et al, 2011, p. 945), these being perfect examples of pragmatic responses (though controversial) that technology can provide for deep problems.

3.3. We seek knowledge & feelings

3.3.1. <u>Research and education</u>. "It could be considered that the production and transmission of knowledge is one of the most positive impacts of robotics and applied artificial intelligence" (Lin et al, 2011, p. 948). Scientists are using robots in laboratory experiments and in the field, such as collecting ocean surface and marine-life data over extended periods (e.g., Rutgers University's Scarlet Knight) and exploring new planets (e.g., NASA's Mars Exploration Rovers). In classrooms, robots are delivering lectures, teaching subjects (e.g., foreign languages, vocabulary, and counting), checking attendance, and interacting with students.

3.3.2. <u>Entertainment</u>. Beyond the animatronics in major film and television productions, it is worth mentioning progress in the field of "edutainment" or education-entertainment robots, which include ASIMO, Nao, iCub, and others. "Though they may lack a clear use, such as for military or manufacturing, they aid researchers in the study of cognition (both human and artificial), motion, and other areas related to the advancement of robotics" (Lin et al, 2011, p. 944). Robotic toys, such as AIBO, Pleo, and RoboSapien, also serve as discovery and entertainment platforms.

3.3.3. <u>Personal care and companions</u>. The social isolation derived from the increase and improvement in telepresence technologies, as well as the limited time devoted to the family by adults in full productive stage, partly explains that the need for affection and company try to be supplied artificially. "Robots are increasingly used to care for the elderly and children, such as RI-MAN, PaPeRo, and CareBot. PALRO, QRIO, and other edutainment robots previously mentioned can also provide companionship" (Lin et al, 2011, p. 944). Moreover, due to physical limitations of its habitual users some developments propose models for a Multimodal Human Computer Interaction System (MMHCI) based on services, embedded on an assistive robot, "which is able to adapt the communication according to the type and degree of the user's disabilities" (John et al, 2016, p. 175). Surprisingly, relationships of a more intimate nature are not quite satisfied by robots yet, considering the sex industry's reputation as an early adopter of new technologies. Introduced in 2010, Roxxxy is billed as "the world's first sex robot" (Fulbright, 2010, § 1), but its lack of autonomy or capacity to "think" for itself, as opposed to merely respond to sensors, suggests that it is not in fact a robot, per the definition above.

4. Machines thought as similar to man

Moving away from the concept of machine as an executor mechanism of advanced algorithms, to a vision of machine as a thinking entity, requires delving into the concept of artificial cognition, its challenges and implications. The state of development in cognitive robotics has a great review in the work titled "Value systems for developmental cognitive robotics: A survey", below is an excerpt of that document, which outlines the technical implications to provide cognitive abilities to an artificial entity by researcher K. Merrick (2017):

Cognition by a living organism generally refers to its capacity to process perceptual information and thereby manipulate its behaviour. Human cognition encompasses a large collection of processes occurring in the human mind. "The cognitive capabilities of humans are generally considered to be manifested in self-awareness, perception, learning, knowledge, reasoning, planning and decisionmaking" (Begum & Karray, 2009). Bringing all of these capabilities together in a single artificial agent or a robot remains one of the grand challenges of our age. "However, a number of recent approaches to this challenge are emerging with the common theme of autonomous mental development" (Weng et al, 2001). "These include artificial development or developmental engineering" (Sandini et al, 1997); "epigenetic robotics" (Zlatev & Balkenius, 2001); "developmental robotics" (Cangelosi & Schlesinger, 2015) and "cognitive developmental robotics" (Asada et al, 2001). These fields of research have two main objectives: 1) to improve the capabilities of artificial systems and 2) "to progress our understanding of the foundational roots of intelligence in natural systems" (Lieto & Radicioni, 2016) (p. 38).

4.1. Reasons not to build them

According to researcher Seth D. Baum from Global Catastrophic Risk Institute (USA) the dilemma is: "Use the technology, and risk the downside of catastrophic failure, or do not use the technology, and suffer through life without it" (Baum, 2014, p. 1).

Baum points out in his paper "The great downside dilemma for risky emerging technologies" that humanity has previously faced crucial decisions regarding promising technologies with potentially catastrophic consequences, citing as particular examples the development of weapons of mass destruction and the search for contact with extraterrestrial civilizations. According to Baum, the decision about the development or implementation of such technologies should not only weigh their potential benefits but all conceivable failure scenarios, and beyond that, Baum emphasizes that the decision must be made based in the state of "real need" of the potential user of this kind of technology, the decision to implement a potentially devastating technology should only be taken in a state of absolute necessity, postponing in this way the use of technologies whose side effects are not completely understood or considered difficult to contain, until a future stage where research progress will provide a framework of operation with greater safety factor.

In relation to the potential danger of artificial general intelligence Baum (2014) expresses the following:

One line of thinking posits that an AGI, or at least certain types of AGIs, could essentially take over the world. This claim depends on two others. First, power benefits from intelligence, such that the most intelligent entities will tend to have the most power. Second, an AGI can gain vastly more intelligence than humans can, especially if the AGI can design an even more intelligent AGI, which designs a still more intelligent AGI, and so on until an "intelligence explosión" (Good, 1965) or "singularity" (Kurzweil, 2005) occurs. The resulting "superintelligent AGI" (Bostrom, 2014) "could be humanity's final invention" (Barrat, 2013) because the AGI would then be fully in control. If the AGI is "Friendly to humanity" (Yudkowsky, 2011), then it potentially could solve a great many of humanity's problems. Otherwise, the AGI will likely kill everyone inadvertently as it pursues whatever goals it happened to be programmed with Per this line of thinking, an AGI would be much like a magic genie … The genie is all-powerful but obligated to serve its master ... The AGI would do exactly what its human programmers instructed it to do, regardless of whether the programmers would, in retrospect, actually want this to happen. In attempting to program the AGI to do something desirable, the programmers could end up dead, along with everyone else on the planet (p. 6).

On the other hand, Lin et al (2011) consider the ethical and social topics given the disruptive nature of technology revolutions (AGI for example) and map the myriad issues into three broad (and interrelated) areas of ethical and social concern:

<u>Safety and errors</u>. Asbestos, DDT, and fen-phen are among the usual examples of technology gone wrong, having been introduced into the marketplace before sufficient health and safety testing… With robotics, the safety issue is with their software and design. Computer scientists, as fallible human beings, understandably struggle to create a perfect piece of complex software: somewhere in the millions of lines of code, typically written by teams of programmers, errors and vulnerabilities likely exist.

<u>Law and ethics</u>. Linked to the risk of robotic errors, it may be unclear who is responsible for any resulting harm… As robots become more autonomous, it may be plausible to assign responsibility to the robot itself, e.g., if it is able to exhibit enough of the features that typically define personhood… And if some (future) robots or cyborgs meet the necessary requirements to have rights, which ones should they have, and how does one manage such portfolios of rights, which may be unevenly distributed given a range of biological and technological capabilities?… Of course, ethical and cultural norms, and therefore law, vary around the world, so it is unclear whose ethics and law ought to be the standard in robotics; and if there is no standard, which jurisdictions would gain advantages or cause policy challenges internationally? Such challenges may require international policies, treaties, and

perhaps even international laws and enforcement bodies.

Social impact. How might society change with the Robotics Revolution? As with the Industrial and Internet Revolutions, one key concern is about a loss of Jobs … performing certain jobs better, faster, and so on—robots may displace human jobs, regardless of whether the workforce is growing or declining. The standard response is that human workers, whether replaced by other humans or machines, would then be free to focus their energies where they can make a greater impact, i.e., at jobs in which they have a greater competitive advantage; to resist this change is to support inefficiency…Yet, theory and efficiency provide little consolation for the human worker who needs a job to feed her or his family (p.945).

In conclusion, it is important to anticipate the impact that the AGI advent may generate, since unnecessary and dramatic disruptions in the natural environment, social relations, and working conditions are reflected in real human costs.

4.2. Reasons to build them

Regarding the development of AGI and its potential benefit, researcher Seth D. Baum (2014) argues several reasons to pursue AGI that could be catalogued as answers to 3 particular questions:

Would not it be better to leave things as they are? AGI is not the only threat that humanity faces. In the absence of AGI, humanity might die out anyway because of nuclear weapons, global warming, or something else. If AGI succeeds, then these other threats could go away, solved by our new computer overlords. That is a significant upside for AGI. What if an AGI has a 50% chance of killing everyone, but absent AGI, humanity has a 60% of dying out from something else? Should the AGI be launched?.

What short-term benefit would it bring? AGI development could involve basic computing resources and technologies of growing economic importance. AGI is not like nuclear weapons, which require exotic materials. AGI could be developed on any sufficiently powerful computer. Computing power is steadily growing, a trend known as Moore's Law. Meanwhile, narrow AI is of increasing technological sophistication and economic importance. At the time of this writing, driverless cars are just starting to hit the streets around the world, with promise to grow into a major industry.

What long-term benefit would it bring? Imagine having the perfect genie: unlimited wishes that are interpreted as you intended them to be, or maybe even better than you intended them to be. That could go quite well. Perhaps there would be no more poverty or pollution. Perhaps space colonization could proceed apace. Perhaps the human population could double, or triple, or grow tens, hundreds, or thousands of times larger, all with no decline in quality of life. A Friendly AGI might be able to make these things posible (p. 6).

It is not the fear of the unknown, but the critical analysis which should guide our decisions about technological innovations capable of attacking present and future problems, and citing the words of science fiction author Isaac Asimov (1978):

It is change, continuing change, inevitable change, that is the dominant factor in society today. No sensible decision can be made any longer without taking into account not only the world as it is, but the world as it will be . . . This, in turn, means that our statesmen, our businessmen, our everyman must take on a science fictional way of thinking (§ 1).

4.3. How would we build them?

The feasibility of building a cognitive and autonomous artificial system has been subject to "intense and numerous debates" (Turing, 1950, p. 433), from the earliest theoretical attempts to raise this possibility. Skeptical arguments of different nature are presented, some from a philosophical perspective: e.g., "You may be able to build systems that behave the same way as real ones, but they won't behave that way for the same reasons. The product, therefore, will not be intelligent. To get genuine intelligence you need the right kind of history, the kind of history that will establish an explanatory connection between what is represented (content) and the behavior that this content helps to explain" (Dretske, 1993,

p. 216), and others, from a design perspective: e.g., "Expecting to create an AGI without first understanding how it works is like expecting skyscrapers to fly if we build them tall enough" (Deutsch, 2012, § 1). Revealing together that the poor understanding of the concept and implications of intelligent behavior, makes unfeasible the material development of an entity with such characteristics.

In order to establish objectives and clear lines of research in the development of AGI, researchers such as Peter Voss (2002) have suggested that an artificial cognitive system in order to proactively accumulate knowledge from changing contexts demands a number of irreducible features and capabilities:

(1) Senses to obtain features from 'the world' (virtual or actual).

(2) A coherent means for storing knowledge obtained this way.

(3) Adaptive output/ actuation mechanisms (both static and dynamic).

Finally, work in AGI should focus on:

- General rather than domain-specific cognitive ability

- Acquired knowledge and skills, versus loaded databases and coded skills

- Bi-directional, real-time interaction, versus batch processing

- Adaptive attention (focus & selection), versus human pre-selected data

- Core support for dynamic patterns, versus static data

- Unsupervised and self-supervised, versus supervised learning

- Adaptive, self-organizing data structures, versus fixed neural nets or databases

- Contextual, grounded concepts, versus hard-coded, symbolic concepts

- Explicitly engineering functionality, versus evolving it

- Conceptual design, versus reverse-engineering

- General proof-of-concept, versus specific real applications development

- Animal level cognition, versus abstract thought, language, and formal logic (§ 3).

It is important to remember that Artificial General Intelligence (AGI) was the original focus of the AI field, but as express Ben Goertzel and Cassio Pennachin (2007):

…due to the demonstrated difficulty of the problem, not many AI researchers are directly concerned with it anymore. Work on AGI has gotten a bit of a bad reputation, as if creating digital general intelligence were analogous to building a perpetual motion machine. Yet, while the latter is strongly implied to be impossible by well-established physical laws, AGI appears by all known science to be quite possible. Like nanotechnology, it is 'merely an engineering problem', though certainly a very difficult one (§ 1).

According to Goertzel and Pennachin (2007) most approaches to AGI may be divided into broad categories such as:

 Symbolic AGI. The majority of ambitious AGI-oriented projects undertaken to date have been in the symbolic-AI paradigm. One famous such project was the General Problem Solver (Newell & Simon, 1961), which used heuristic search to solve problems … there is no learning involved. GPS worked by taking a general goal – like solving a puzzle – and breaking it down into subgoals. It then attempted to solve the subgoals, break-ing them down further into even smaller pieces if necessary, until the subgoals were small enough to be addressed directly by simple heuristics.

Uncertainty focused AGI. Judea Pearl's work on Bayesian networks (Pearl, 1988) introduces principles from probability theory to handle uncertainty in an AI scenario. Bayesian networks are graphical models that embody knowledge about probabilities and dependencies between events in the world. Inference on Bayesian networks is possible using probabilistic methods.

Neural net-based AGI. The neural net approach has not spawned quite so many frontal assaults on the AGI problem, but there have been some efforts along these lines. Werbos has worked on the application of recurrent networks to a number of problems (Werbos, 1977). Stephen Grossberg's work (Grossberg, 1992) has led to a host of special neural

network models carrying out specialized functions modeled on particular brain regions. Piecing all these networks together could eventually lead to a brain-like AGI system.

Evolutionary AGI. The evolutionary programming approach to AI has not spawned any ambitious AGI projects, but it has formed a part of several AGI-oriented systems, including Novamente system (Goertzel et al, 2003), de Garis's CAM-Brain machine, and John Holland's classifier systems (Holland, 1986). Classifier systems are a kind of hybridization of evolutionary algorithms and probabilistic-symbolic AI; they are AGI-oriented in the sense that they are specifically oriented toward integrating memory, perception, and cognition to allow an AI system to act in the world. Typically they have suffered from severe performance problems, but Eric Baum's recent variations on the classifier system theme seem to have partially resolved these issues (Baum & Durdanovic, 2002).

Artificial life. The artificial life approach to AGI has remained basically a dream and a vision, up till this point. Artificial life simulations have succeeded, to a point, in getting interesting mini-organisms to evolve and interact, but no one has come close to creating an Alife agent with significant general intelligence.

Program search based AGI. Program search based AGI is a newer entry into the game. It had its origins in Solomonoff, Chaitin and Kolmogorov's seminal work on algorithmic information theory in the 1960s, but it did not become a serious approach to practical AI until quite recently, with work such as Schmidhuber's OOPS system, and Kaiser's dag-based program search algorithms. This approach is different from the others in that it begins with a formal theory of general intelligence, defines impractical algorithms that are provably known to achieve general intelligence, and then seeks to approximate these impractical algorithms with related algorithms that are more practical but less universally able.

Integrative AGI. The integrative approach to AGI involves taking elements of some or all of the above approaches and creating a combined, synergistic system. This makes sense if you believe that the different AI approaches each capture some aspect of the mind uniquely well. But the integration can be done in many different ways. It is not workable to simply create a modular system with modules embodying different AI paradigms: the different approaches are too different in too many ways. Instead one must create a unified knowledge representation and dynamics framework, and figure out how to manifest the core ideas of the various AI paradigms within the universal framework (§ 1).

Finally, it is important to note that the above categorization is not intended to contain all research approaches in the field of AGI, given the intrinsic complexity of the numerous visions candidates for models of the mind. In addition, and although none of the above approaches has yet achieved tangible results regarding the construction of an artificial cognitive system, research in this branch of scientific knowledge should not conclude, since the emergence of consciousness in humans from a materialistic perspective reveals the feasibility that complex structures (the human body), formed by inert unitary elements (atoms / subatomic particles), rise to the highest levels of complexity (a conscious entity). Moreover if we think that providing this capability to artificial structures can be easily considered one of the greatest achievements in the history of mankind.

5. Conclusions

Do we want machines to rule the world? Narrow artificial intelligence answers: "No, we do not". Actually, we could not subjugate ourselves to an entity to which we have not given enough power, understood that power as the high degree of cognitive complexity that should have an entity with the intention of governing us, being this precisely the line of development of AGI but not of Narrow AI.

The development of Narrow AI has led and will probably lead to important technological revolutions capable of altering the scientific, technical and social structures of our civilization. The development of analytical and operational capacities with a higher level of complexity in computing machines leads to such systems being projected as potential substitutes for professions and specialties previously considered under the sole domain of human intellect, so the rate of technological progress will actually reveal in the future the

key point of distinction between an artificial cognitive system and the human mind (if such a difference actually exist). However, under the Narrow AI approach could be difficult to arouse intention and desires in an artificial structure that does not have autonomy in a not bounded domain of operation (the world).

Do we want machines to rule the world? Artificial general intelligence answers: "Yes, we do", or rather: "We would like to see that they try". With this purpose in mind AGI directs all its efforts to the development of a theoretical and technological framework where to consider the final product as "a conscious entity" can actually become a legitimate doubt, unlike the sophisticated automatisms (with a generally deterministic conception) which we consider our biggest breakthrough to date.

Based on what has been said above, it is clear that the original man's attempt to build an entity in his image and likeness, has represented a struggle of many generations that has finally derived in two ways, on the one hand the productive, immediate and increasingly demanded field of narrow artificial intelligence, and on the other hand, the sinuous, controversial and still yearned field of artificial general intelligence.

Will machines ever rule the world? Our final answer would be: "It is certain that without our help they will not achieve it, and it is uncertain if even with our help they would achieve it."

Acknowledgments

Bibliographic references

Agostini, A., Alenyà, G., Fischbach, A., Scharr, H., Wörgötter, F., & Torras, C. (2017). A cognitive architecture for automatic gardening. Computers and Electronics in Agriculture, 138, 69-79.

Asimov, I. (1978). My own view. In R. Holdstock, The Encyclopedia of Science Fiction. New York, NY: St. Martin's Press.

Baum, S. (2014). The great downside dilemma for risky emerging technologies. Phys. Scr, 89, 10.

Carter, P. (2012). The Emerging Story of the Machine. IGI Global.

Dean, C. (2006). SCIENTIST AT PLAY/DANIEL WILSON; If Robots Ever Get Too Smart, He'll Know How to Stop Them. Retrieved from http://query.nytimes.com/gst/fullpage.html?res=9F07E0DA123EF937A25751C0A9609C8B63&pagewanted=all

Deutsch, D. (2012). Creative blocks. Retrieved from https://aeon.co/essays/how-close-are-we-to-creating-artificial-intelligence

Dretske, F. (1993). Can Intelligence Be Artificial? Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 71(2), 201-216.

Farkas, I. (2003). Artificial intelligence in agriculture. Computers and Electronics in Agriculture, 40, 1-3.

Fulbright, Y. (2010). Meet Roxxxy, the 'woman' of your dreams. Retrieved from http://www.foxnews.com/story/0,2933,583314,00.html

Gates, B. (2007). A robot in every home. Scientific American, 58-65.

Goertzel, B., & Pennachin, C. (2007). Artificial General Intelligence. Springer.

Guizzo, E. (2010). World Robot Population Reaches 8.6 Million. Retrieved from IEEE Spectrum: http://spectrum.ieee.org/automaton/robotics/industrial-robots/041410-world-robot-population

John, E., Rigo, S., & Barbosa, J. (2016). Assistive Robotics: Adaptive Multimodal Interaction Improving People with Communication Disorders. IFAC- Papers OnLine, 49, 175-180.

Lin, P., Abney, K., & Bekey, G. (2011). Robot ethics: Mapping the issues for a mechanized world. Artificial Intelligence, 175, 942-949.

Merrick, K. (2017). Value systems for developmental cognitive robotics: A survey. Cognitive Systems Research, 41, 38-55.

Mumford, L. (1970). The Myth of the Machine. In The Pentagon of Power (Vol. II).

Murase, H. (2000). Artificial intelligence in agriculture. Computers and Electronics in Agriculture, 29, 1-2.

Nasuto, S., & Hayashi, Y. (2016). Anticipation: Beyond synthetic biology and cognitive robotics. Biosystems, 148, 22-31.

Turing, A. (1950). Computing machinery and intelligence. Mind, 49, 433-460.

Voss, P. (2002). Essentials of general intelligence: the direct path to AGI. Retrieved from http://www.kurzweilai.net/essentials-of-general-intelligence-the-direct-path-to-agi

Wang, Z., & Srinivasan, R. (2017). A review of artificial intelligence based building energy use prediction: Contrasting the capabilities of single and ensemble prediction models. Renewable and Sustainable Energy Reviews, 75, 796-808.

Zhao, Y., Gong, L., Huang, Y., & Liu, C. (2016). A review of key techniques of vision-based control for harvesting robot. Computers and Electronics in Agriculture, 127, 311-323.

1. GIDIS Research Group. SENA Regional Tolima. MSc Mechanical Engineer.mpedroza@sena.edu.co
2. Faculty of Psychology. Universidad Pontificia Bolivariana. PhD in Education. gustavo.villamizar@upb.edu.co
3. Faculty of Psychology. Universidad Pontificia Bolivariana. Undergraduate Student. jamego94@gmail.com

[Index]

[In case you find any errors on this site, please send e-mail to webmaster]