



Extended Abstract

Google AI and Turing's social definition of intelligence

Teresa Numerico^{1,*}

¹ Dept. Philosophy, Communication, Performing Arts, University of Rome III , V. Ostiense 234 – 00146 – Rome, Italy / Address

E-Mails: teresa.numerico@uniroma3.it (T. N.);

* Author to whom correspondence should be addressed; Tel.: +39-33867 10 425; Fax: +39-0657 338340

Accepted: march 2015

Introduction

Over the past few years, major Internet, hardware and software companies such as Google, Facebook, Amazon, Microsoft, and IBM have conducted considerable research in AI and robotics, and made substantial investments in firms that are active in those fields. For example, at the beginning of 2014, Google acquired DeepMind, an AI company based in London and co-founded by a former child chess prodigy, Demis Hassabis, for around £400m. In October of that year, Google also bought two other AI British spin-offs of Oxford University: Dark Blue, a company working on machine learning techniques for natural languages, and Vision Factory, who are working on computer vision. The Web search giant also launched a partnership between the Google AI structure and the AI research group at Oxford University. All these AI acquisitions, partnerships and investments followed a long list of acquisitions of robotics companies over the past few years. A major component of this list was Boston Dynamics, a company whose primary clients are the US Army, Navy and Marine Corps (Cohen 2014).

Google clearly needs this kind of technology to deal not only with their core business, but also with their more innovative developments, such as Google Glass, self-driving cars, etc. But why do AI and robotics appear to be so strategic to Google? Even their core business deals with language understanding and translation, both in processing texts to help formulate responses to query requests, and in understanding the queries themselves through speech recognition. Visual recognition is used in the retrieval of images and videos, and in the development of Google Glass and self-driving cars. All

of these areas of interest seem to be covered by the traditional fields of AI that from the 1960s onward has promised solutions to the same kinds of problems that Google is trying to solve today.

Another potential area for AI is the management of Big Data, which, as argued above (cf. 4.5.1), lies at the heart of the marketing and profiling activities of the search engine. In order to process huge amounts of data, adequate correlation algorithms are needed, based mainly on machine learning techniques developed in the field of AI during the 1970s and 1980s. It is not by chance that the director of Research at Google Inc. is now Peter Norvig, a key figure in the field, who was co-author, along with Stuart Russell, of the classical textbook on AI, *Artificial Intelligence: a modern approach* since 1995.

Some old thoughts about AI definition

According to the paper “Intelligent machinery”, written in 1948 by Alan Turing, considered one of the fathers of the modern thought about AI, the key areas of the field were: “(i) Various games ... (ii) The learning of languages, (iii) translation of languages, (iv) Cryptography, (v) mathematics” (Turing 1948/2004, 420). With the inclusion of modern image recognition, these still appear to be the main areas of interest in AI.

Any field in which the same projects remain unaccomplished for so long naturally raises suspicions. But Turing offered another interesting point of view to understand why the potential of AI deserves attention: “the extent to which we regard something as behaving in an intelligent manner is determined as much by our own state of mind and training as by the properties of the object under consideration. If we are able to explain and predict its behavior or if there seems to be a little underlying plan, we have little temptation to imagine intelligence” (Turing 1948/2004, 431).

As suggested by the title of the last paragraph of his paper, intelligence is an emotional and subjective concept, and depends just as much on the observer’s own conception of intelligence, and on his/her mental condition and training, as on the properties of the object under investigation.

According to Turing, the definition of intelligence should not be considered too crucial because: “at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machine thinking without expecting to be contradicted” (Turing 1950/2004, 449).

The suggestion was thus that, in the 50 years following the publication of his paper, the transformation of conscience and beliefs of the general public would be extensive enough to elevate “machine thinking” to the level of the generally perceived notion of “intelligence”.

These two assumptions taken together: the training and mental state of the researcher, and the change in the mentality of the general audience, form the main presupposition of AI. In future it will be agreed that what machines will be able to do will be a reasonable definition of machine intelligence.

The ‘intelligent’ results of DeepMind

The AI results obtained by DeepMind in April 2014 on the First Day of the Tomorrow Technology conference in Paris (See video presentation: <https://www.youtube.com/watch?v=EfGD2qveGdQ>), as described in (Mnih, Hassabis, *et al.* 2015), offer an interesting perspective on what may be considered intelligent, for a task achieved by a machine.

The video shows the AI computer program as it learns to play an Arcade game, Breakout, which belongs to the family of the earliest Atari 2600 games, prototyped in 1975 by Wozniak and Jobs (Twilley 2015). The aim of the game is to destroy a wall brick by brick, using a ball launched against it. The “intelligent” program behaves in a very progressive way: at the start it is not very clever, but it quickly learns. After half an hour it becomes as clever as a non-clever human, and after 300 games the program stop missing the ball, and becomes more expert that the most expert human.

According to the description – published as a letter in *Nature* on the 26th of February 2015 – behind the great success of the Deep Q-network (DQN) lies two different kinds of “intelligent” tools: the ability to create a representation of an environment using high-dimensional hierarchical sensory inputs, and the ability to reinforce learning strategies based solely on that environment, without any rearrangement of the data, or pre-programming of the learning strategies. DQN was able to develop a winning strategy by simply analyzing the pixels on the screen and the game scores: successful actions resulted from analysis of the sensory description provided by the deep neural network.

According to the description of the experiment: “To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations” (Mnih, Hassabis, *et al.* 2015, 529). The deep neural network was described by comparing its performance to human behavior. Part of the enthusiasm surrounding this achievement was attributed to the assumption, as declared by the AI experts who described the results, that game-playing expertise can be compared to that needed to solve real-world complexity problems. However this assumption could be called into question, considering the low level of complexity in 1970s arcade games. Their working environments look very different from real world problems, both in terms of the nature of the rudimentary task, and in the computer reconstruction of the limited sensory experience.

To successfully compare human strategies with the neural networks requires an interpretation and analysis all the metaphors used to describe the objectives and accomplishments of the Deep-Q network: “humans and other animals seem to solve this problem through a harmonious combination of reinforcement learning and hierarchical sensory processing systems ..., the former evidenced by a wealth of neural data revealing notable parallels between the phasic signals emitted by dopaminergic neurons and temporal difference reinforcement learning algorithms”(Mnih, Hassabis, *et al.* 2015, p. 529). This reveals an underlying assumption that there is a symmetry between the behavior of dopaminergic neurons, and the function of the reinforcement learning algorithm used by the DQN. When declaring this parallel the researchers did not provide any convincing argument to prove their hypothesis. To this may be added the doubts already mentioned about whether simple success in playing Breakout should be included in any possible definition of human intelligence. Despite these undemonstrated assumptions, many appear ready to appreciate and support the result as an outstanding progress of AI research. This fact may thus be one of the social transformations of the definition of intelligence anticipated by Turing’s 1950 paper.

The current situation thus does not differ too much from that in the 1960s and the 1970s where the game-like problem solutions of AI software prototypes were rhetorically transformed into the first step in climbing the mountain of human-like intelligent performances.

Conclusion: when the controller and the controlled are the same agent

Another noteworthy effect of this exploitation of the AI effect within Google is their creation of an “ethics board”, as required by the DeepMind people when they were taken over by Google. However, it is surely a strange practice to appoint an in-house ethics commission aimed at self-regulation. The board’s task is twofold: on one hand it aims to judge and absolve Google of any potential “sin” when managing AI software, and on the other hand it guarantees that Google is not doing any evil, when trying to emulate certain human abilities via software.

Google thus acts as if it were both the controller and the agent under control: it accepts no authority except itself, when discussing what it does that is right or wrong. This “affirmative discourse” approach (Bunz 2014) to international, social and geopolitical problems became very successful, and sufficiently convincing that other international authorities also accepted it, as suggested by the decision of the European Court of Justice on the right to be forgotten, published in May 2014⁽³⁾. The European Court attributed to Google the role of the European guarantor of the right to be forgotten. Google accepted the role, but at the same time published a report⁽⁵⁾ written by the Advisory Council, also nominated by Google, describing how it may act as an advisory committee in performing the role of guarantor. The report suggests rules that the company should follow in protecting the right of some people to have their data forgotten, while maintaining the right of the rest of the Europeans to have their data both available and preserved. It is the same situation of the ethics board for AI: Google is asking itself to be the guarantor of the ethical control over its operations of technological advancement.

References and Notes

1. Bunz, M. 2014, *The silent revolution*, Palgrave Macmillan, New York
2. Cohen R. “What’s Driving Google’s Obsession with Artificial Intelligence and robots?”, *Forbes*, 1/28/2014, <http://www.forbes.com/sites/reuvencohen/2014/01/28/whats-driving-googles-obsession-with-artificial-intelligence-and-robots/print/> .
3. Court of Justice of the European Union’s ruling in *Google Spain and Inc. vs. Agencia Española de Protección de Datos (AEPD) and Mario Costeja Gonzalez C131/12*. Sentence issued in May 2014.
4. Mnih V., Kavukcuoglu K., Hassabis D. *et al.* “Human-level control through deep reinforcement learning”, in *Nature*, 26 feb. 2015, Vol. 518, pp. 529-533.
5. Report of the Advisory Council to Google on the Right to be Forgotten <https://drive.google.com/file/d/0B1UgZshetMd4cEI3SjlvV0hNbDA/view>, published on the 6th of February 2015.
6. Turing A. M. “Intelligent Machinery” *Report, National Physics Laboratory*, 1948, in B. Meltzer D. Michie (Eds.), *Machine intelligence*, 5, Edinburgh Univ. Press, 1969: 3-23; reprint in Copeland J. (Ed.) *The Essential Turing*, Clarendon Press, Oxford, 2004, 410-432.
7. Turing A.M. “Computing Machinery and Intelligence”, *MIND*, 59 (1950), pp.433-460, reprinted in *Collected Works of A. M. Turing: mechanical intelligence*, D. C. Ince, (Ed.) North-Holland, Amsterdam 1992, pp. 133-160 and in Copeland J. (Ed.) *The Essential Turing*, Clarendon Press, Oxford, 2004, pp.441-464, URL: <http://mind.oxfordjournals.org/content/LIX/236/433.full.pdf>.
8. Twilley N. “Artificial Intelligence goes to the arcade”, *The New Yorker*, 2/25/2015, <http://www.newyorker.com/tech/elements/deepmind-artificial-intelligence-video-games>.

© 2015 by the authors; licensee MDPI and ISIS. This abstract is distributed under the terms and conditions of the Creative Commons Attribution license.