Present and Foreseeable Future of Artificial Intelligence

Chapter · January 2015

CITATIONS
2

READS
268

1 author:

Luciano Floridi
University of Oxford
461 PUBLICATIONS **14,417** CITATIONS

SEE PROFILE

Luciano Floridi

PRESENT AND FORESEEABLE FUTURE OF ARTIFICIAL INTELLIGENCE

ESTRATTO

da

NOETICA VS INFORMATICA
Le nuove strutture della comunicazione scientifica

Atti del Convegno Internazionale (Roma 19-20 novembre 2013)

A cura di Fiammetta Sabba

Leo S. Olschki Editore
Firenze

Noetica versus informatica

Le nuove strutture della comunicazione scientifica

Atti del convegno internazionale
Roma, 19-20 novembre 2013

a cura di
Fiammetta Sabba

Leo S. Olschki

NOETICA
VERSUS
INFORMATICA

LE NUOVE STRUTTURE
DELLA COMUNICAZIONE SCIENTIFICA

Atti del Convegno Internazionale
Roma, Tempio di Adriano
19-20 novembre 2013

a cura di
FIAMMETTA SABBA

LEO S. OLSCHKI EDITORE
MMXV

Casa Editrice Leo S. Olschki
Viuzzo del Pozzetto, 8
50126 Firenze
www.olschki.it

Direttore scientifico
Alfredo Serrai
Coordinatore scientifico e curatore degli atti
Fiammetta Sabba

Partner scientifici
Centro per il libro e la lettura,
Università "La Sapienza" di Roma,
Università degli Studi di Milano,
Università degli Studi di Firenze, ISTC-CNR,
University of Oxford – Oxford Internet Institute

Comitato scientifico
Domenico Bogliolo, Flavia Cristiano, Luciano Floridi,
Aldo Gangemi, Mauro Guerrini, Giorgio Montecchi,
Roberto Mondello Novelli, Alberto Petrucciani, Fiammetta Sabba, Alfredo Serrai, Giovanni Solimine

Il convegno è stato organizzato da Promoroma,
Azienda Speciale della CCIAA di Roma

La pubblicazione dei presenti atti si deve interamente
al personale sostegno del Dr. Roberto Mondello Novelli

ISBN 978 88 222 6367 4

SOMMARIO

Luciano Floridi

PRESENT AND FORESEEABLE FUTURE OF ARTIFICIAL INTELLIGENCE

Abstract

Siamo sempre più dipendenti dalle applicazioni legate all'intelligenza artificiale (tecnologie "smart") per eseguire attività che sarebbero semplicemente impossibili per un'intelligenza umana non aiutata o non aumentata. Questo è possibile perché il mondo sta diventando una "infosfera" sempre meglio adattata alle capacità limitate dell'intelligenza artificiale. Essere capaci di immaginare le esigenze di adattamento di questo processo per l'umanità può contribuire a concepire soluzioni tecnologiche che possono diminuire i loro costi antropologici.

We increasingly rely on AI-related applications (smart technologies) to perform tasks that would be simply impossible by un-aided or un-augmented human intelligence. This is possible because the world is becoming an infosphere increasingly well adapted to AI's limited capacities. Being able to imagine what adaptive demands this process will place on humanity may help to devise technological solutions that can lower their anthropological costs.

It is a well-known fact, although sometimes underestimated, that AI research seeks both to 'reproduce' the outcome of our intelligent behaviour by non-biological means, and to 'produce' the non-biological equivalent of our intelligence.

On the one hand, as a branch of engineering interested in 'intelligent behaviour reproduction', AI has been astoundingly successful, well beyond the rosiest expectations. Nowadays, we increasingly rely on AI-related applications (sometimes called smart technologies, though the expression has a wider scope) to perform tasks that would be simply impossible by un-aided or un-augmented human intelligence. Reproductive AI regularly outperforms and replaces human intelligence in an ever-larger number of contexts.

Edsger Wybe Dijkstra's famous comment that «the question of whether a computer can think is no more interesting than the question of whether a submarine can swim» is indicative of the applied approach shared by reproductive AI. Next time you experience a bumpy landing recall that that is probably because the pilot was in charge, not the computer.

On the other hand, as a 'branch of cognitive science interested in intelligence production', AI has been a dismal disappointment. Current machines have the intelligence of a toaster and we really haven't got much of a clue about how to move from there.[1]

Apparently, artefacts can be smart without being intelligent. Productive AI does not merely underperform with respect to human intelligence; it has not joined the competition yet. John McCarthy's disappointed remarks about Deep Blue's victory against Kasparov are symptomatic of the sort of productive AI that frowns upon reproductive AI. When the warning "printer not found" pops up on the screen of your computer, it may be annoying but hardly astonishing, despite the fact that the printer in question is actually right next to it.

The two souls of AI have often engaged in fratricidal feuds for intellectual predominance, academic power, and financial resources. That is partly because they both claim common ancestors and a single intellectual inheritance: a founding event, the Dartmouth Summer Research Conference on Artificial Intelligence in 1956, and a founding father, Turing, with his machine and its computational limits, and then his famous test. It hardly helps that a simulation might be used in order to check both whether the simulated source has been produced, and whether the targeted source's behaviour or performance has been reproduced or even surpassed.

The two souls of AI have been variously and not always consistently named. Sometimes the distinctions weak vs. strong AI, or Good-Old Fashioned (GOFAI) vs. New or Nouvelle AI, have been used to capture the difference. I prefer to use the less loaded distinction between light vs. strong AI.[2]

The misalignment of their goals and results has caused endless and mostly pointless diatribes. Defenders of AI point to the strong results of reproductive AI, which is really weak or light AI in terms of goals; whereas

---

[1] Luciano Floridi – Mariarosaria Taddeo – Matteo Turilli, *Turing's Imitation Game: Still a Challenge for Any Machine and Some Judges*, «Minds and Machines» 19, 2009, 1, pp. 145-150.

[2] Luciano Floridi, *Philosophy and Computing: An Introduction*. London-New York, Routledge, 1999.

detractors of AI point to the weak results of productive AI, which is really strong AI in terms of goals. Much of the current debate on the so-called singularity issue has its roots in such confusion.

In order to escape the dichotomy just outlined, one needs to realise that AI cannot be reduced to a "science of nature", or to a "science of culture",[3] because it is a "science of the artificial", to put it with.[4] As such, AI pursues neither a 'descriptive' nor a 'prescriptive' approach to the world: it investigates the constraining conditions that make possible to build and embed artefacts in the world and interact with it successfully. In other words, it 'inscribes' the world, for such artefacts are new logico-mathematical pieces of code, that is, new texts, written in Galileo's mathematical book of nature.

Until recently, the widespread impression was that such process of adding to the mathematical book of nature (inscription) required the feasibility of productive or strong AI. After all, developing even a rudimentary form of non-biological intelligence may seem to be not only the best but perhaps the only way to implement technologies sufficiently adaptive and flexible to deal effectively with a complex, ever-changing and often unpredictable when not unfriendly environment. What Descartes acknowledged to be an essential sign of intelligence – the capacity to learn from different circumstances, adapt to them and exploit them to one's own advantage – would be a priceless feature of any appliance that sought to be more than merely smart.

Such impression is not incorrect, but it is distracting. For while we were pursuing the task of inscribing AI in the world, the world itself quietly but steadily begun to change. In order to explain how, let me introduce two more concepts.[5]

Infosphere is a neologism I coined years ago on the basis of "biosphere", a term referring to that limited region on our planet that supports life. It denotes the whole informational environment constituted by all informational entities (thus including informational agents as well), their properties, interactions, processes and mutual relations. It is an environment comparable to, but different from cyberspace (which is only one of its sub-regions, as it

---

[3] Jean-Gabriel Ganascia, *Epistemology of Ai Revisited in the Light of the Philosophy of Information*, «Knowledge, Technology & Policy», 23, 2010, 1, pp. 57-73.

[4] Herbert A. Simon, *The Sciences of the Artificial*. 3rd ed., Cambridge, Mass.; London, MIT Press, 1996.

[5] Luciano Floridi, *A Look into the Future Impact of Ict on Our Lives*, «The Information Society», 23, 2007, 1, pp. 59-64.

were), since it also includes off-line and analogue spaces of information. It is an environment (and hence a concept) that is rapidly evolving.

'Re-ontologising' is another neologism that I have recently introduced in order to refer to a very radical form of re-engineering, one that not only designs, constructs or structures a system (e.g. a company, a machine or some artefact) anew, but that fundamentally transforms its intrinsic nature. In this sense, for example, nanotechnologies and biotechnologies are not merely re-engineering but actually re-ontologizing our world.

These two concepts are not indispensable – the reader is welcome to rely on any other useful shortcuts – but they are useful to clarify my previous claim, in the following way: digital ICTs are re-ontologizing the very nature of (and hence what we mean by) the infosphere, while the infosphere is progressively becoming the world in which we live. It follows that, while we were unsuccessfully pursuing the inscription of strong AI into the world, we were actually re-ontologising the world to fit light AI. Especially in recent years, the world as infosphere has been adapting to AI limited capacities increasingly well. Using a term from robotics, we have been enveloping[6] the world without fully realising it. The example of a dishwasher is elementary but still helpful to make the point. We do not build robots that wash dishes like us, we envelop micro-environments around simple robots to fit and exploit at best their limited capacities and still deliver the desired output. It is the difficulty of finding the right enveloping that makes ironing (as opposed to pressing) so time-consuming.

Enveloping used to be either a stand-along phenomenon (you buy the robot with the required envelop, like a dishwasher or a washing machine) or implemented within the walls of industrial buildings. Nowadays, enveloping the environment into an AI-friendly infosphere has started pervading any aspect of reality and is visible everywhere, on a daily basis. If driverless vehicles can move around with decreasing troubles, this is not because strong AI has finally arrived, but because the "around" they need to negotiate has become increasingly suitable to light AI applications.[7] We do not have semantically proficient technologies, but we have accumulated so much data, can rely on so many humans, and have such good statistical tools that purely syntactic technologies can bypass problems of meaning and understanding, and still deliver what we need: a translation, the right picture of a place, the preferred restaurant, the interesting book and so

[6] In robotics, an *envelope* (also known as reach envelop) is the three-dimensional space that defines the boundaries that the robot can reach.

[7] See the progressive successes of the DARPA Grand Challenge.

forth. Indeed, some of the issues we are facing today, e.g., in e-health or in financial markets, already arise within highly enveloped environments in which all relevant (and sometimes the only) data are machine-readable, and decisions as well as actions may be taken automatically, by applications and actuators that can execute commands and output the corresponding procedures, from alerting or scanning a patient, to buying or selling some bonds. Examples could easily be multiplied. It is a trend that is robust, cumulative and progressively refining: everyday sees the availability of more tags, more humans online, more documents, more statistical tools, more devices that communicate with each other, more sensors, more RFID tags, more satellites, more actuators, more data collected on all possible transitions of any system, in a word, more enveloping.

This is good news for the future of light AI and smart technologies in general, which will be exponentially more useful and successful with every step we take in the expansion of the infosphere. Enveloping is a process that has nothing to do with some sci-fi singularity, for it is not based on some unrealistic (as far as our current and foreseeable understanding of AI and computing is concerned) speculations about some super AI taking over the world in the near future. But it is a process that raises some challenges. In order to express the one I have in mind, let me use a parody.

Two people A and H are married and they really wish to make their relationship work, but A, who does increasingly more in the house, is inflexible, stubborn, intolerant of mistakes and unlikely to change, whereas H is just the opposite, but is also becoming progressively lazier and dependent on A. The result is an unbalanced situation, in which A ends up shaping the relationship and distorting H's behaviours, practically, if not purposefully. If the marriage works, that is because it is carefully tailored around A. Now, light AI and smart technologies play the role of A in the previous analogy, whereas their human users are clearly H. The risk we are running is that, by enveloping the world, our technologies might shape our physical and conceptual environments and constrain us to adjust to them because that is the best, or sometimes the only, way to make things work. After all, light AI is the stupid but laborious spouse and humanity the intelligent but lazy one, who is going to adapt to whom, given that a divorce is not an option? The reader will probably recall many episodes in real life when something could not be done, or had to be done in a very cumbersome or silly way because that was the only way to make the computerised system do what it had to do. Here is a more concrete, trivial example (philosophically, things are way more complex). The risk is that we might end up building houses with round walls and furniture with sufficiently high legs in order to

fit the capacities of a Roomba[8] much more effectively. I certainly wish our house were more Roomba-friendly. The example is useful to illustrate not only the risk but also the opportunity represented by ICT's re-ontologising power and the enveloping of the world.

There are many "roundy" places in which we live, from igloos to medieval towers, from bow windows to public buildings where corners of the rooms are rounded for sanitary reasons. If we spend most of our time inside squarish boxes that is because of another set of technologies related to the mass production of bricks and concrete infrastructures, and the ease of straight cuts of building material. It is the mechanical circular saw that, paradoxically, generates a right-angled world. In both cases, squarish and roundy places have been built following the predominant technologies, rather than through the choices of their potential inhabitants. Following this example, it is easy to see how the opportunity represented by technologies' re-ontologising power comes in three forms: rejection, critical acceptance, and proactive design. By becoming more critically aware of the re-ontologising power of light AI and smart ICT applications, we might be able to avoid the worst forms of distortion (rejection) or at least be consciously tolerant of them (acceptance), especially when it does not matter (consider the Roomba-friendly length of the legs of the furniture) or when this is a temporary solution, while waiting for a better design. In the latter case, being able to imagine what the future will be like and what adaptive demands technologies will place on their human users may help to devise technological solutions that can lower their anthropological costs. In short, intelligent design should play a major role in shaping the future of our interactions with forthcoming technological artefacts. After all, it is a sign of intelligence to make stupidity work for you.[9]

---

[8] < http://www.irobot.com/ >.

[9] I have discussed the nature of applied philosophy of information in LUCIANO FLORIDI, *On Defining Library and Information Science as Applied Philosophy of Information*, «Social Epistemology», 16, 2002, 1, pp. 37-49, and in ID., *Lis as Applied Philosophy of Information: A Reappraisal*, «Library Trends», 52, 2004, 3, pp. 658-665; and I have discussed the relevance and transparency of information in LUCIANO FLORIDI, *Understanding Epistemic Relevance*, «Erkenntnis», 69, 2008, 1, pp. 69-92, and in MATTEO TURILLI – LUCIANO FLORIDI, *The Ethics of Information Transparency*, «Ethics and Information Technology», 11, 2009, 2, pp. 105-112.

FINITO DI STAMPARE
PER CONTO DI LEO S. OLSCHKI EDITORE
PRESSO ABC TIPOGRAFIA • SESTO FIORENTINO (FI)
NEL MESE DI MARZO 2015

L'Informatica ha una capacità "mentale" adeguata per aggredire e dominare la sfera della semantica? Il dominio della linguistica elettronica, esercitatosi finora nei confini del lessico e della terminologia, potrà invadere anche i terreni del significato e sostituire il cervello umano nelle aree della gnoseologia, della indicizzazione, e della documentazione? In questo volume si offre un panorama delle visioni e delle realtà riferite alla Comunicazione, in particolare scientifica, e si mettono a fuoco i problemi relativi alla registrazione, trasmissione, e diffusione del sapere e della conoscenza nel contesto sociale ed economico modificato dai modelli e dalle tecniche dell'informatica.

L'efficienza dei mezzi elettronici lascia credere che con essi sia possibile superare, oggi, tutti i limiti che la rigidità della comunicazione a stampa imponeva nello scambio e nella trasmissione delle informazioni. In realtà si tratta di un abbaglio smentito da ciò che avviene negli effettivi processi indicali e documentari che si giovano delle soluzioni informatiche; essi si rivelano, infatti, incapaci di affrontare e risolvere gli attuali nodi della comunicazione semantica, ovverosia di quella che ha luogo per concetti. Il lettore potrà così venire a conoscenza, oltre che delle principali e diffuse teorie e tecnologie informative, anche di problematiche implicate nelle sfere della filosofia, della linguistica, della psicologia, e della neurologia.

ISBN 978 88 222 6367 4