

2100 AI: Reflections on the mechanisation of scientific discovery

Andrea Mannocci, Angelo A. Salatino, Francesco Osborne, Enrico Motta

Knowledge Media Institute, The Open University, Milton Keynes, UK
`name.surname@open.ac.uk`

Abstract. The pace of research is nowadays extremely intensive, with datasets and publications being published at an unprecedented rate. In this context data science, artificial intelligence, machine learning and big data analytics are providing researchers with new automatic techniques which not only help them to manage this flow of information but are also able to identify automatically interesting patterns and insights in this vast sea of information. However, the emergence of mechanised scientific discovery is likely to dramatically change the way we do science, thus introducing and amplifying serious societal implications on the role of researchers themselves, which need to be analysed thoroughly.

1 Introduction

In recent years, we have witnessed the ever-increasing availability of data produced by modern scientific instrumentation [14] (e.g. the Sloan Digital Sky Survey, the Large Hadron Collider, the Compact Muon Solenoid, etc.), the emergence of “big-data” from the Web, and an acceleration in the amount of published scientific papers¹. Digesting the information conveyed by such a “data deluge” stretches far beyond human analytical capabilities to the extent that having humans-in-the-loop could be viewed as a bottleneck and a barrier to scientific progress [9]. As a result, automated knowledge discovery has become an important new topic and leading scientists have posed “grand challenges” in semantically-enabled artificial intelligence (AI) advocating the development of new intelligent systems that, by delving into data and reasoning, can support researchers in their daily scientific work.

In particular, in 2003, Agresti [1] coined the term “Discovery Informatics” to refer to a newborn methodology “employing the full spectrum of computing and analytical science and technology to the singular pursuit of discovering new information by identifying and validating patterns in data”. Later, Gil et al. [9, 8, 15] outlined perspectives and opportunities deriving from the application of artificial intelligence to scientific discovery and founded the Discovery Informatics Initiative². More recently, Kitano proposed an even more ambitious vision by

¹ One new life science paper is published every 30 seconds; approximately 10000 papers published per day only on PubMed, according to <https://www.raconteur.net/technology/machine-learning-teaches-computers-to-tackle-big-data>

² Discovery Informatics Initiative, <http://discoveryinformaticsinitiative.org>

imagining in the near future an artificial intelligence system able to make major scientific breakthroughs in biomedical sciences, save millions of lives and maybe win a Nobel Prize [19]. This topic is also studied in research areas other than artificial intelligence, such as social and political sciences, philosophy, humanities; several authors published surveys forecasting the uptake of high-level machine intelligence (HLMI) and super (ultra) intelligence, and discussed their impact on the society [2, 22, 12, 16].

Indeed the improvements AI can bring to research endeavour are tangible, undeniable and somewhat needed if researchers want to keep up with the pace of science nowadays; nonetheless, the depicted scenario opens up to potential social issues. Science-fiction literature about artificial intelligence, alongside their derived dystopian speculations, has a longstanding tradition. In early '50s, scientific and technological progress led Isaac Asimov to anticipate and describe with fervid imagination scenarios that nowadays indeed appear forthcoming. In 1954, contemporary to the first groundbreaking findings on neural networks, in the short story “Answer”, Fredric Brown’s scientists create a super computer capable of active thinking and self-reflection³. Regarding scientific literature, in 1966 Irwin J. Good already speculated on “ultraintelligent machines” and their value and implications on the society by stating that “the man would be left far behind”, that such artefacts would “create social problems, but might also be able to solve” others and that they will be “feared and respected, and perhaps even loved” [11]. Other studies thoroughly studied the impact, sustainability and implications of this second machine revolution in the context of modern society and economic paradigms [3, 7].

In this work, we discuss and provide examples about potential scenarios stemming from the achievement of full automation of scientific discovery thanks to data-driven AI. In particular, we emphasise the implications over researchers’ community, scientific discovery and its advances.

2 The current picture

Nowadays scientists have become more and more dependant on the Web and semantic technologies in order to bolster their activity. Online services such as Google Scholar, Microsoft Academic Graph, Semantic Scholar, OpenAIRE⁴ and CORE⁵ empower researchers to find needles in a haystack across the plethora of papers published online; Scientific blogs (e.g. Science Online), Virtual Research Environments (VRE) [4], Science Gateways⁶, specially devised social networks (e.g. Academia.edu and ResearchGate) foster interdisciplinarity and render (quasi) null distances and speed of communication among research communities. A novel topic, “semantic publishing”, has emerged with the aim to create vast machine-readable data corpora describing human knowledge, and help AI agents to understand and reason on scholarly and scientific data [20].

³ Fredric Brown, “Answer”, <http://www.roma1.infn.it/anzel/answer.html>

⁴ <https://www.openaire.eu>

⁵ <https://core.ac.uk>

⁶ <http://sciencegateways.org/about/science-gateway-basics>

Furthermore, scientists learnt to leverage ontologies and the Semantic Web in order to boost performances of information retrieval and interlinking [13]. Such machine-readable information can be exploited for generating “synthesis engines” capable of digesting, reasoning and inferring new knowledge by processing (ideally) arbitrary amount of data. For example, Big Mechanism [5] is a \$45 million DARPA research program aiming at synthesising new models of cancer signalling pathways by reading automatically papers available in the literature and stitching together causal hypotheses extracted from them. Another system, Hanalyzer [21], creates a knowledge network about genes and their interactions by extracting and integrating information from the literature and other heterogeneous sources. It then reasons over such a knowledge graph and assists biologists in understanding phenomena in genomic-scale data and form new hypotheses. Wings [10] is a workflow management system (WfMS) enabling the definition of semantic constraints over workflows so as to automatically select models and algorithms and generate customised workflow instances matching the requirements provided. Nutonian⁷ and DataRobot⁸ offer solutions able to explore the hypothesis space consistently to the datasets fed in input, autonomously select the most promising ones and devise experiments in order to test them. Finally, King et al. describe Robot Scientist [18, 17], a closed-loop discovery system designed around a laboratory workstation for conducting scientific experiments in functional genomics. Without human attendance, Robot Scientist exploits AI in order to autonomously formulate hypotheses consistent with its current background knowledge, and then validate or disprove them by designing and physically running laboratory experiments and interpreting the results obtained; it then repeats the whole process over a thousands times per day.

AI-powered data-driven knowledge discovery and research offers a fair share of opportunities outside academia too; in fact, an ever-growing number of startups⁹ is currently eyeing up computational drugs discovery [6] and other bio-science and biotechnology applications. The aforementioned systems and approaches are rather specialised for narrow fields of application; nevertheless, the methodologies described can be adapted to other research areas. This is not to argue on a holistic generalisation as these methodologies might be of little or no relevance to some disciplines such as humanities and arts.

3 Research at the time of machines

With such a premise, adding on top of the “intelligent science assistant” [8], it is quite easy to imagine a forthcoming future in which an “artificial scientist” leverages artificial intelligence and Semantic Web technologies in order to combine together a number of capabilities, which in many cases are already available in state of the art systems. For instance, such an artificial scientist could be able to (i) summarise the state of the art and major findings in a given field by

⁷ <http://www.nutonian.com/products/eureqa-desktop>

⁸ <https://www.datarobot.com>

⁹ <http://www.nanalyze.com/2017/04/9-ai-computational-drug-discovery>

extracting information across a multitude of heterogeneous sources (e.g. scientific papers, patents, datasets, etc.); (ii) map current knowledge across different fields and locate “voids” (i.e. opportunities); (iii) explore interdisciplinary research opportunities; (iv) track and forecast the migration of research concepts, promising technologies and methodologies; (v) formulate hypotheses to explain observed phenomena; (vi) build models and design/run experiments validating them; (vii) find hidden data patterns and give sense to dark data; (viii) document results and findings in natural language (potentially already writing a paper draft); (ix) keep research findings up-to-date by reiterating the process whenever better tools become available.

Undoubtedly, achieving these technological advances would provide very valuable support to human researchers. Nonetheless, it is also the case that major advances in AI solutions for knowledge discovery risk to exacerbate some negative phenomena, which are already observable on a global scale. For example, a consequence of the emergence of large scale, mechanised scientific discovery could be that the discovery process could end up in the hands of a small number of organisations, able to afford the required technology for super-intelligence. This would increase inequality and strengthen a small number of dominant players, a view already discussed in [23] for knowledge producers in the Semantic Web. The companies and research institutes able to access the infrastructures for gathering, processing and reasoning over more data will essentially set trends in research and lead the scientific discovery worldwide, while smaller players will be left at the margin to face sustainability issues and eventually succumb. This perspective, superimposed on an already crippled research landscape characterised by unclear scientometrics scores and indexes, a debated peer-review process and a not always satisfactory openness and transparency, offers an explosive substrate that can seriously undermine and compromise credibility of future science. Indeed, it is important to emphasise that this trend is actually already ongoing and can be easily noted, for example, by observing authors affiliations in the most influential publications published in journals and conferences such as WWW¹⁰. Here, it is already apparent that those few large scale players with access to both big data and the infrastructure to analyse them are starting to monopolise research in key sectors such as web-scale data mining. Similarly, major discoveries in experimental physics are achievable only with access to large scale experimental infrastructures and facilities.

Furthermore, the emergence of large scale, mechanised scientific discovery could also impact on the already controversial topic concerning the attribution of groundbreaking discoveries. For example, recent major experimental physics discoveries born thanks to coordinated and extensive collaboration of physicists, scientist and technicians worldwide. The recent observation of gravitational waves for example was possible thanks to a tight collaboration between the LIGO and Virgo initiatives and the author list of the resulting publication is extenuatingly long¹¹. Similarly, behind Nobel prize attributions there is of-

¹⁰ <http://papers.www2017.com.au.s3-website-ap-southeast-2.amazonaws.com>

¹¹ <https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.116.061102>

ten the work of an entourage of many nameless individuals. For example, the discovery of W and Z bosons that awarded the Nobel prize in physics in 1984 to Carlo Rubbia and Simon van der Meer is the result of the work of over a hundred physicists, as well as the discovery of Higgs boson, result of CMS and ATLAS collaborations¹². Imagine now what would happen if an award such as the Nobel prize had to be assigned for an AI-driven discovery. Who would be rewarded? The artificial scientist? Probably not. The lead researcher of the team owning the artificial scientist? Maybe, but is it fair? What about the rest of the entourage and the technicians operating the infrastructure? And most of all, if a discovery is out of sheer brute-force hypothesis search, is it really worth a prize?

Finally, in a future dominated by AI-driven scientific discoveries, it could be the case that the main research activities would be focused on studying more deeply the technologies and the methodologies enabling more efficient and exhaustive search within the hypothesis space. Or, even worse, the researchers' daily tasks could be declassified into just checking machine arguments and the congruity of results. As a consequence, research teams composition could gradually change until the majority of the team members are technicians able to program AI, operate and fix the machines (i.e. AI as a commodity) under the supervision of one (or few) lead scientist(s) taking care of defining the line of research. As a consequence, such a vast request of highly specialised practitioners in AI may accentuate further the already present bias in both the job market revolving around research and the education system preparing new generations of AI experts for academia and industry. This would be analogous to what has already happened in many automatised shop-floors, where the robots do all the actual work of assembling components while the task of the humans is to deal with the software side and the monitoring of the operation.

4 Conclusions

Data science and semantically-enabled AI will no doubt disrupt the traditional way of doing science. New paradigms are emerging offering both new horizons to explore and potentially sustainable methodologies to keep up with today's hectic pace of research. The benefits of this synergy are indeed manifold and can relieve researchers from the heavy-lifting in their daily endeavour – and indeed the process of transitioning from research as “cottage industry” to large-scale enterprise has been going on for a while [14]. However, the emergence of mechanised scientific discovery is likely to introduce and amplify serious societal and economic implications on the role of researchers themselves and the way we do research, which need to be analysed thoroughly and promptly prevented.

For the time being we can keep calm: AI still need humans for setting up, training and operation. The data to be fed to machine learning algorithms still needs to be manually curated by researchers and the process of hypotheses generation and pruning also requires humans. Nonetheless, it is not science fiction, and is both exciting and scaring, to envisage a world where artificial intelligence

¹² <https://www.scientificamerican.com/article/expand-nobel-prize-award-teams-not-just-individuals>

could supersede humans in doing what has characterised them the most so far: science.

References

1. Agresti, W.W.: Discovery informatics. *Commun. ACM* 46(8), 25–28 (2003)
2. Baum, S.D., Goertzel, B., Goertzel: How Long Until Human-Level AI ? Results from an Expert Assessment. *Science* 78(1), 1–33 (2011)
3. Brynjolfsson, E., McAfee, A.: *The Second Machine Age - Work, progress, and prosperity in a Time of Brilliant Technologies*. WW Norton & Company (2014)
4. Candela, L., Castelli, D., Pagano, P.: Virtual Research Environments: An Overview and a Research Agenda. *Data Science Journal* 12, GRDI75–GRDI81 (2013)
5. Cohen, P.R.: DARPA’s Big Mechanism program. *Physical Biology* 12(4) (2015)
6. Duch, W., Swaminathan, K., Meller, J.: Artificial intelligence approaches for rational drug design and discovery. *Current pharmaceutical design* 13(14) (2007)
7. Ford, M.: *Rise of the robots: technology and the threat of a jobless future*. Oneworld Publications (2015)
8. Gil, Y., Greaves, M., Hendler, J., Hirsh, H.: Amplify scientific discovery with artificial intelligence. *Science* 346(6206), 171–172 (2014)
9. Gil, Y., Hirsh, H.: Discovery informatics: AI opportunities in scientific discovery. In: *AAAI Fall Symposium - Technical Report*. vol. FS-12-03, pp. 1–6 (2012)
10. Gil, Y., Ratnakar, V., Kim, J., González-Calero, P., Groth, P., Moody, J., Deelman, E.: Wings: Intelligent workflow-based design of computational experiments. *IEEE Intelligent Systems* 26(1), 62–72 (jan 2011)
11. Good, I.J.: Speculations Concerning the First Ultraintelligent Machine. *Advances in Computers* 6, 31–88 (1966)
12. Grace, K., Salvatier, J., Dafoe, A., Zhang, B., Evans, O.: When will AI exceed human performance? Evidence from AI experts. *CoRR* (2017)
13. Hendler, J.: Communication. *Science and the semantic web*. *Science (New York, N.Y.)* 299(5606), 520–521 (2003)
14. Hey, T., Tansley, S., Tolle, K.: *The fourth paradigm: data-intensive scientific discovery*. Microsoft Research (2009)
15. Honavar, V.G.: The promise and potential of big data: A case for discovery informatics. *Review of Policy Research* 31(4), 326–330 (2014)
16. Hosni, H., Vulpiani, A.: Forecasting in the Light of Big Data. *arXiv (May)* (2017)
17. King, R.D., Benway, M.: *Robot Scientist : an Autonomous Platform for Systems Biology Discovery* (2006)
18. King, R.D., Whelan, K.E., Jones, F.M., Reiser, P.G.K., Bryant, C.H., Muggleton, S.H., Kell, D.B., Oliver, S.G.: Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* 427(6971), 247–252 (2004)
19. Kitano, H.: *Artificial Intelligence to Win the Nobel Prize and Beyond: Creating the Engine for Scientific Discovery*. *AI Magazine* 37(1), 39–49 (2016)
20. Kuhn, T., Dumontier, M.: *Genuine Semantic Publishing*
21. Leach, S.M., Tipney, H., Feng, W., Baumgartner, W.A., Kasliwal, P., Schuyler, R.P., Williams, T., Spritz, R.A., Hunter, L.: Biomedical discovery acceleration, with applications to craniofacial development. *PLoS Computational Biology* 5(3) (2009)
22. Müller, V.C., Bostrom, N.: Future progress in artificial intelligence: A survey of expert opinion. *AI Matters* pp. 1–19 (2016)
23. Stutt, A., Motta, E.: Knowledge modelling an organic technology for the knowledge age. In: *The Knowledge Web*, chap. 13, pp. 211–224. Kogan Page (2000)