

Combining collective and machine intelligence at the knowledge frontier

E. Malliaraki, Nesta, United Kingdom

A. Berdichevskaia, Nesta, United Kingdom

Introduction

In the past decade, machine learning and deep learning have advanced significantly. They can now assist in the process of discovery, for instance, by analysing large volumes of data. On the other hand, humans have unique abilities such as creativity, intuition, contextualisation and abstraction. Moving forward, the best of both worlds must be combined. Instead of using only artificial intelligence (AI) to navigate scientific knowledge, novel AI and human collaborations could explore complexity and advance the frontiers of scientific understanding in new ways. This essay describes emerging tools and initiatives for discovering, encoding and synthesising knowledge that could help guide the way. The recommendations outline a pathway for changing scientific infrastructures, incentives and institutions to help hybrid human-AI science to flourish.

Each day, researchers publish more than 4 000 scientific papers in the field of biomedicine alone. An estimated 200 000 articles have been written to date about the COVID-19 pandemic. While the number of publications grows exponentially (Bornmann and Mutz, 2015), the number of novel scientific ideas expands only linearly (see Staša Milojević's essay in this volume). Across fields, from medicine to agriculture to computers, the effort and money required to innovate are growing.

Traditional mechanisms for getting up to speed with, and navigating, the knowledge frontier have limitations. Textbooks are updated only slowly. Meanwhile, literature reviews are often ad hoc and may overlook relevant work from other disciplines or discount disruptive ideas in favour of canonical works (Chu et al., 2021).

As fields grow, they often divide into subspecialties, each with its own literature. These cluster in discrete areas of knowledge (Foster et al., 2015). Fragmentation leads to a combinatorial explosion of “undiscovered public knowledge” (i.e. knowledge existing in the unrevealed connections between existing bodies of publicly available knowledge) (Swanson, 1989). In addition, while older or neglected findings and hypotheses from distant disciplines could become sources of new knowledge, the means are unavailable to systematically resurrect them (Swanson, 2011). This results in undiscovered correlations, undrawn conclusions and a failure to capitalise on insights from analogous problems (see the essay by Smallheiser et al. in this volume).

At the same time, science is carried out by ever-larger teams and international consortia.¹ More than ever before, science has become a collective endeavour. This is witnessed in successes ranging from mapping the human genome to pushing the frontier of neuroscience research with the Human Brain Project and proving Einstein's theory of gravitational waves. From optimising the size (Wu et al., 2019) and diversity of

teams (Nielsen et al., 2018) to leveraging emerging methods like crowdsourcing and crowd forecasting (Sell et al., 2021), an understanding of how to make the most of collective intelligence for scientific discovery is just beginning to emerge.

Reimagining current methods and tools to better harness collective and machine intelligence will help scientists better assess the state of scientific knowledge and prioritise research at the knowledge frontier (Berditchevskaia and Baeck, 2020).

Mapping the knowledge frontier

Encoding and discovering knowledge

Scientific knowledge consists of concepts and relationships represented in research papers, patents, software and other academic artefacts. However, the existing science communication infrastructure does not help researchers make the best use of the predominantly document-centric scholarly outputs. For example, the extracted text, graphics, bibliography and metadata from PDF files often require extensive cleaning before use. While words and sentences may be searched for, images, references, symbols and other semantics are currently mostly inaccessible to machines.

In recent years, it has become increasingly possible to represent scholarly knowledge in machine-actionable form. So far, the emphasis has been on representing, maintaining, and linking data and metadata about articles, people and other relevant entities. However, in metadata discovery, a combination of humans and machines allows for better detection of scientific data artefacts. For example, project RePEc uses AI to infer which datasets have been used in a publication and then requests authors to validate or reject these dataset annotations (Nathan, 2019).

Once encoded, these pieces of public knowledge must be searchable and discoverable at the right level of representation. Recent advances in natural language processing (NLP), text mining and information retrieval can now help researchers find and understand scientific information. For example, citation analyses using NLP help researchers discover intersections and emerging trends across scholarly databases, as well as lines of research that are either new or not mainstream. Moreover, improved representation learning and extreme summarisation in scholarly documents can alleviate information overload. Open Knowledge Maps is a notable example of a human-AI collaboration that clusters papers in similar subfields based on keywords, datasets and research software. It then allows users to create, edit and update their own knowledge maps (Matthews, 2021).

Several initiatives have focused on identifying and extracting more granular units of scientific information from papers. These units can be problems (Lahav et al., 2021), hypotheses (Spangler et al., 2014), methods (Fathalla et al., 2017), findings (Sebastian et al., 2017), causal relations and even automated suggestions of new hypotheses (Liekens et al., 2011). Specifically, recent advances in language models (i.e. sciBERT, bioBERT)² create much more accurate semantic representations of scientific concepts and allow for a more contextual search of scientific documents. However, these models are typically built on pretrained language representations by domain experts. They have limited generalisability outside the domains where they are developed.

This problem of limited generalisability can be addressed by harnessing the complementary pools of expertise from scientists and policy makers. A notable example is the TREC-COVID project, which extracts useful information from publications, Twitter conversations and library searches (Roberts et al., 2021). It then collects rankings and relevance judgements over submitted paper pools from medical domain experts and uses them to increase the accuracy of search algorithms. The potential for expanding these approaches across the sciences is considerable.

Finding analogies in distant fields often drives scientific discovery. However, the growing volume of academic publications makes it difficult to identify topics in a single discipline, let alone cross-domain analogies. Project Solvent (Chan et al., 2018) addresses this by collecting annotations of problems and findings in papers contributed by large groups of people. The annotated dataset is then used to train algorithms that identify semantic analogies between research papers. Similar data-driven methods for learning abstract relationships between concepts could be used to identify sub-problems and constraints and suggest novel R&D pathways based on analogy (Hope et al., 2021).

Connecting and structuring knowledge

Once relevant public knowledge is encoded and discovered, it needs to be organised and synthesised. With recent advances in knowledge representation and human-machine interaction, scholarly information can be expressed as semantically-rich knowledge graphs (Auer et al., 2018). Knowledge graphs are a way to organise the world's structured knowledge by mapping the connections between different concepts and integrating information extracted from multiple data sources (Chaudhri, Chittar and Genesereth, 10 May 2021). Current automatic approaches to create these graphs only achieve moderate accuracy and have limited coverage. Where they exist, they tend to describe scholarly knowledge in specific fields, such as mathematics, chemistry and the life sciences.

A knowledge network for science must enable parallel and synchronised encoding and augmentation of knowledge. In the near future, knowledge synthesis system(s) will need to continuously navigate the development of knowledge and integrate new submissions both from theory and empirical evidence, including qualitative and quantitative studies. Intelligent interfaces could help experts find connections between concepts and theories cloaked in archaic academic language and emerging datasets and computational models (Chen and Hitt, 2021). NLP can already automate certain elements of mapping and translation between concepts or constructs to enable a more dynamic grouping of similar or complementary ideas. In the future, the development of causal inference methods within NLP may start to estimate cause-and-effect relationships (Feder et al., 2021). This is likely to be a particularly fruitful touchpoint for human-machine collaboration, as domain experts use the resulting knowledge graphs to contextualise, test and explore emerging relationships between various knowledge entities such as methods and experimental data.

To do this complex task, ontologies would have to be developed to give a stable frame to knowledge. An ontology can be thought of as a schema for organising information. For example, an ontology might be developed to define fields and subfields of research. One such schema – developed by the Web of Science – defines approximately 250 subject areas in science, social sciences, and the arts and humanities. Many other types of ontology can also be used to formally characterise academic literature. For example, ontologies might be created to specify who is considered a research contributor or to tag parts of research papers.

The Open Research Knowledge Graph (ORKG) is a good example of combined collective and machine intelligence to give structure to academic content. It organises and connects scholarly knowledge using crowdsourced contributions from researchers who acquire, curate, publish and process this knowledge (Karras et al., 2021). Enhanced machine-interpretable semantic content would help support more systems like the ORKG, making it easier for human experts to find connections and map connections based on their domain knowledge and understanding. In the meantime, collective intelligence is being used to enrich documents and annotate articles. The Dokie.li project³ collects decentralised semantic annotations from scientists and MicroPublications, which is itself a new model to represent scientific arguments semantically.

Oversight and quality control

A knowledge synthesis infrastructure will not be complete without ongoing curation and quality assurance by domain experts, librarians and information scientists. During the COVID-19 pandemic, communities of academics came together to curate resources relevant to the crisis response. SciBeh, created in the early days of the pandemic, is one such community. SciBeh is a network of behavioural scientists aiming to create a crisis knowledge management infrastructure fit for rapid response, while maintaining the rigour of the scientific process.⁴ Collective processes of quality assurance are increasingly important, as many published findings in social sciences have been difficult to replicate (Camerer et al., 2018). The reasons are well established, ranging from poor experiment design, small sample sizes, naive data analysis practices and a publication bias towards positive findings. Investing in new tools that automatically check scientific papers for limitations,⁵ automatically categorise scientific uncertainty or predict the likelihood of replication in scholarly communication⁶ could help address part of the problem. However, such systems are unlikely to be completely foolproof. They will require augmentation by highly distributed peer review or crowdsourced intelligence from multiple experts able to detect and point to the evidence-containing parts of the publications.

Introducing these new tools into established scientific practice will take time. In the context of knowledge discovery, AI tools are more likely to be used if they are developed in response to community needs and can be adapted based on ongoing community oversight and feedback (Halfacker and Geiger, 2020). Otherwise, they risk neglect or rejection from the users they are intended to support. If such tools are not properly adapted to user needs, they might produce results of lower quality or impact than more traditional approaches.

Experiments supported through Nesta's Collective Intelligence Grants Programme over the last two years have generated early findings about some obstacles to successful human-machine collaboration.⁷ In one experiment, a serendipity-inducing recommendation algorithm, developed to improve the search for novel ideas and information, confused rather than helped groups to design new solutions to societal challenges (Gill, Peach and Steadman, 2021). These experiments highlight the importance of developing tools together with established communities of users to ensure their successful integration into existing workflows.

How to integrate combined AI-human systems into mainstream science

Scholarly communication and scientific progress have much to gain from capitalising on the emerging approaches outlined here. Some future directions on how to accelerate the integration of combined AI-human systems into mainstream science are offered below for consideration by the academic community, scientific institutions and science funders.

Invest in new infrastructure to navigate the knowledge frontier

A finer-grained and machine-actionable representation of scholarly knowledge is needed, along with the infrastructure to support knowledge curation, publishing and synthesis by humans and machines. This infrastructure should be able to support the storage of documents and datasets, as well as link this content to people and institutions.

Scholarly outputs present unique challenges for processing that necessitate the development of NLP methods optimised for this domain. This requires a more extensive research and development programme to develop processes and workflows for linking AI and human-based processing components in knowledge discovery and synthesis. Promising future directions may include new workflows that organise a pipeline of combined human-machine approaches to uploading and annotating content, followed by knowledge

graph development. Lastly, knowledge from non-traditional data sources, like tweets, drug labels, news articles and web content might be integrated into and augment literature-based discovery.

Develop AI tools that focus explicitly on co-operation and enhancing collective intelligence

Another innovation opportunity lies in creating tools to optimise the collective intelligence potential of scientific teams or wider groups through mass collaboration. This will require more research into co-operative AI systems that can enhance group problem solving and decision making for collective benefit.

Co-operative human-AI systems will be inherently complex. They will have to learn to navigate problems where the goals of different actors and organisations are in tension with one another, as well as those where actors have common agendas (Dafoe et al., 2021). Currently, this area of research has lagged behind other topics in AI when it comes to investment (Littman et al., 2021).⁸

Make use of existing social networks to experiment with human-AI collaboration

Multiple social platforms support knowledge exchange between academics and provide an infrastructure for discovering literature. Examples include ResearchGate, Academia.edu and the Loop community from the Frontiers journals group. Some of these platforms already use AI-enabled recommendation systems to tailor content for users (Matthews, 2021). In the future, such platforms should become testbeds for experimenting with new forms of combined human-AI knowledge discovery, idea generation and synthesis. Experimentation along such lines could help these platforms develop distinctive services in comparison to those of existing social networks like Twitter and LinkedIn, which researchers increasingly use to communicate about their work.

Re-think incentives for knowledge mapping and synthesis

Several institutional, educational and social conditions inhibit knowledge integration. Existing measures of publishability reward incremental advances in depth of understanding. They also motivate discoveries built on individual disciplines rather than knowledge synthesis. Editors, reviewers and academic institutions value theoretical coherence within a specific domain with their associated traditional analyses over others. Such foci are all essential to progress in science, but an exclusive focus on these approaches to knowledge creation will miss other available opportunities. To help widen perspectives, some have argued for a new market for knowledge synthesis workers (Chen and Hitt, 2021), integrative PhD programmes⁹ and/or industry research programmes to innovate based on knowledge synthesis.

Research councils and academic institutions should experiment with these proposals and support new roles and career paths. They could support what might be termed “applied metascientists” – experts in curating and maintaining information infrastructure, who can also build crucial bridges between the public, academia and industry. Other initiatives could include prize or competition mechanisms, such as Science4Cast,¹⁰ which incentivises mapping and forecasting the future of scientific research.

Strengthen collaborative institutions

Academia, government and industry must work together on a national and international scale to create the tools for better human and machine understanding of scientific knowledge. Such collaborations can also help democratise access to the latest science-of-science algorithms and maintain a common codebase for researchers and non-technical practitioners to navigate the knowledge frontier together. The United States has begun revitalising its research infrastructure. There have also been calls for the United Kingdom to establish new institutions such as the Atlas Institute¹¹ to fully map the world of scientific knowledge and identify gaps. The successful uptake of research infrastructures such as repositories for preprints (e.g. arXiv) and datasets (e.g. Zenodo) has demonstrated the research community can integrate new

infrastructures. However, new infrastructures may need to reach a critical level of uptake before being widely accepted. This takes time and endorsement by the wider research community, including funders and other institutions.

Conclusion

Some of the biggest scientific challenges today are cross-disciplinary in nature and will require interdisciplinary collaboration to solve. The persistence of isolated islands of knowledge risks slowing scientific progress. Scientific and policy networks require the tools, incentives and institutional structures to track new knowledge across fields, prioritise challenges and work across disciplines to solve them. A wide spectrum of health, economic, social and corporate challenges has already benefited from human-machine partnerships. In the current trajectory, only AI will be able to keep up with and make sense of the ever-growing scientific literature. Productive human-AI collaborations can harness the best of both worlds, the human and the machine, helping scientists to identify scientific priorities and the solutions to our most pressing problems.

References

- Auer, S. et al. (2018), "Towards a knowledge graph for science", in *Proceedings of the 8th International Conference on Web Intelligence, Mining and Semantics*, pp. 1-6, <https://doi.org/10.1145/3227609.3227689>.
- Berditchevskaia, A. and P. Baek (2020), "The future of minds and machines: How AI can scale and enhance collective intelligence", 10 February, Nesta, London, www.nesta.org.uk/mindsmachines.
- Bornmann, L. and R. Mutz (2015), "Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references", *Journal of the Association for Information Science and Technology*, Vol. 66/11, pp. 2215-2222, <https://doi.org/10.1002/asi.23329>.
- Camerer, C. et al. (2018), "Evaluating the replicability of social science experiments in nature and science between 2010 and 2015", *Nature Human Behaviour*, Vol. 2/9, pp. 637-644, <https://doi.org/10.1038/s41562-018-0399-z>.
- Chan, J. et al. (2018), "Solvent: A mixed initiative system for finding analogies between research papers", in *Proceedings of the ACM on Human-Computer Interaction*, 2 (CSCW), pp.1-21, <https://doi.org/10.1145/3274300>.
- Chaudhri, V.K., N. Chittar and M. Genesereth (10 May 2021), "An introduction to knowledge graphs", The Stanford AI Lab Blog, <http://ai.stanford.edu/blog/introduction-to-knowledge-graphs>.
- Chen, V.Z. and M.A. Hitt (2021), "Knowledge synthesis for scientific management: Practical integration for complexity versus scientific fragmentation for simplicity", *Journal of Management Inquiry*, Vol. 30/2, pp.177-192, <https://doi.org/10.1177/1056492619862051>.
- Chu, J. et al. (2021), "Slowed canonical progress in large fields of science", in *Proceedings of the National Academy of Sciences*, Vol 118/41, pp. e2021636118, <https://doi.org/10.1073/pnas.2021636118>.
- Dafoe, A. et al. (2021), "Cooperative AI: Machines must learn to find common ground", *Nature*, Vol. 593/7857, pp. 33-36, <https://doi.org/10.1038/d41586-021-01170-0>.
- Fathalla, S. et al. (2017), "Towards a knowledge graph representing research findings by semantifying survey articles" in *International Conference on Theory and Practice of Digital Libraries*, pp. 315-327, Springer, Cham, https://doi.org/10.1007/978-3-319-67008-9_25.
- Feder, A. et al. (2021), "Causal inference in natural language processing: Estimation, prediction, interpretation and beyond", *arXiv*, arXiv:2109.00725 [cs.CL], <http://arxiv.org/abs/2109.00725>.

- Foster, J.G. et al., 2015. "Tradition and innovation in scientists research strategies", *American Sociological Review*, Vol. 80/5, pp. 875-908, <https://doi.org/10.1177/0003122415601618>.
- Gill, I., K. Peach and I. Steadman (2021), "Collective intelligence grants programme: Experiments in collective intelligence design for social impact", 14 October, Nesta, London, www.nesta.org.uk/report/experiments-collective-intelligence-design-20/.
- Halfaker, A. and R.S. Geiger (2020), "ORES: Lowering barriers with participatory machine learning in Wikipedia", *Proceedings of the ACM on Human-Computer Interaction*, Vol. 4/CSCW2, Article 148, pp. 1-37, <https://doi.org/10.1145/3415219>.
- Hope, T. et al. (2021), "Scaling creative inspiration with fine-grained functional facets of product ideas", *arXiv*, arXiv:2102.09761 [cs.HC], <https://arxiv.org/abs/2102.09761>.
- Karras, O. et al. (2021), "Researcher or crowd member? Why not both! The Open Research Knowledge Graph for Applying and Communicating CrowdRE Research", *arXiv*, arXiv:2108.05085 [cs.DL], <https://arxiv.org/abs/2108.05085>.
- Lahav, D. et al. (2021), "A search engine for discovery of scientific challenges and directions", *arXiv*, arXiv:2108.13751 [cs.CL], <https://arxiv.org/abs/2108.13751>.
- Liekens, A.M. et al. (2011), "BioGraph: Unsupervised biomedical knowledge discovery via automated hypothesis generation", *Genome Biology*, Vol. 12/6, pp.1-12, <https://doi.org/10.1186/gb-2011-12-6-r57>.
- Littman, M.L. et al. (2021), *Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report*, Stanford University, Stanford, <http://ai100.stanford.edu/2021-report>.
- Matthews, D. (2021), "Drowning in the literature? These smart software tools can help", *Nature*, Vol. 597/7874, pp.141-142, <https://doi.org/10.1038/d41586-021-02346-4>.
- Nathan, P. (2019), "Human-in-the-loop AI for scholarly infrastructure", 14 September, *Derwen*, <https://medium.com/derwen/dataset-discovery-and-human-in-the-loop-ai-for-scholarly-infrastructure-e65d38cb0f8f>.
- Nielsen, M.W. et al. (2018), "Making gender diversity work for scientific discovery and innovation", *Nature Human Behaviour*, Vol. 2, pp. 726-734, <https://doi.org/10.1038/s41562-018-0433-1>.
- Roberts, K. et al. (2021), "Searching for scientific evidence in a pandemic: An overview of TREC-COVID", *arXiv*, arXiv:2104.09632 [cs.IR], <https://arxiv.org/abs/2104.09632>.
- Sebastian, Y. et al. (2017), "Emerging approaches in literature-based discovery: Techniques and performance review", *The Knowledge Engineering Review*, Vol. 32, p. e12, <https://doi.org/10.1017/S0269888917000042>.
- Sell, T.K. et al. (2021), "Using prediction polling to harness collective intelligence for disease forecasting", *BMC Public Health*, Vol. 21, pp. 2132, <https://doi.org/10.1186/s12889-021-12083-y>.
- Spangler, S. et al. (2014), "Automated hypothesis generation based on mining scientific literature", in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1877-1886, <https://doi.org/10.1145/2623330.2623667>.
- Swanson, D.R. (2011), "Literature-based resurrection of neglected medical discoveries", *DISCO: Journal of Biomedical Discovery and Collaboration*, Vol. 6, pp. 34-47, <https://doi.org/10.5210/disco.v6i0.3515>.
- Swanson, D.R. (1989), "Online search for logically-related noninteractive medical literatures: A systematic trial-and-error strategy", *Journal of the American Society for Information Science*, Vol. 40/5, pp.356-358, <https://dblp.uni-trier.de/rec/journals/jasis/Swanson89.html>.
- Wu, L. et al. (2019), "Large teams develop and small teams disrupt science and technology", *Nature*, Vol. 566, pp. 378-382, <https://doi.org/10.1038/s41586-019-0941-9>.

Notes

¹ The science policy community increasingly recognises the importance of large-scale international consortia. See, for example, the *Bold Ambition: International Large-Scale Science* report from the American Academy of Arts and Sciences initiative on Challenges for International Scientific Partnerships. Full text available at www.amacad.org/publication/international-large-scale-science.

² In particular, transformer-based language models have been leading the field. These are deep-learning models that use the method of attention to boost their training speed. The Transformer approach was first described in <https://arxiv.org/abs/1706.03762>.

³ See website for more information: <https://dokie.li/>.

⁴ The SciBeh project was established in 2020 to manage the new information, data and resources being produced by behavioural scientists during the COVID-19 pandemic, www.scibeh.org/ (accessed 24 February 2022).

⁵ SciFact, a project from the AllenAI institute, hosts a competition to develop AI models that check the veracity of scientific claims, <https://leaderboard.allenai.org/scifact/submissions/public> (accessed 24 February 2022).

⁶ The SCORE programme from DARPA aims to develop AI-enabled tools to assign confidence metrics to results and studies across behavioural and social sciences. www.darpa.mil/program/systematizing-confidence-in-open-research-and-evidence (accessed 24 February 2022).

⁷ Collective Intelligence Grants 2.0 was a GBP 500 000 fund to support experiments in collective intelligence design, focusing on the interaction between humans and machines. www.nesta.org.uk/project/collective-intelligence-grants/, (accessed 24 February 2022).

⁸ There is evidence that some funders and researchers in the AI community are taking note. The Cooperative AI Foundation was established in 2021, with an initial endowment of USD 15 million. See www.cooperativeai.com/foundation.

⁹ The Venture Science Doctorate from Day One Project is training graduates to combine research and entrepreneurship, thereby incentivising commercialisation of academic research, www.dayoneproject.org/post/forging-1-000-venture-scientists-to-transform-the-innovation-economy (accessed 24 February 2022).

¹⁰ This is an open competition that aims to develop machine-learning models capable of capturing the evolution of scientific concepts and predict which research topics will emerge. See www.iarai.ac.at/science4cast/ (accessed 24 February 2022).

¹¹ The Atlas Institute concept was proposed in a blog by the Tony Blair Institute. www.tenentrepreneurs.org/the-way-of-the-future (accessed 24 February 2022).



From:

Artificial Intelligence in Science

Challenges, Opportunities and the Future of Research

Access the complete publication at:

<https://doi.org/10.1787/a8d820bd-en>

Please cite this chapter as:

Malliaraki, Eirini and Aleks Berditchevskaia (2023), "Combining collective and machine intelligence at the knowledge frontier", in OECD, *Artificial Intelligence in Science: Challenges, Opportunities and the Future of Research*, OECD Publishing, Paris.

DOI: <https://doi.org/10.1787/dbbd48a9-en>

This document, as well as any data and map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area. Extracts from publications may be subject to additional disclaimers, which are set out in the complete version of the publication, available at the link provided.

The use of this work, whether digital or print, is governed by the Terms and Conditions to be found at <http://www.oecd.org/termsandconditions>.