The Alan Turing Institute

AI Scientist Grand Challenge

Summary of discussion during workshop held in February 2020

The Alan Turing Institute

The Alan Turing Institute is the national institute for data science and artificial intelligence. We pursue research which tackles the big challenges in data science, with lasting effects for science, the economy and the world we live in. The Turing applies data science research to real-world problems, working with partners in industry, government and the third sector.

The Institute was founded in March 2015 as a registered charity (1162533) and Company Limited by Guarantee (09512457). It is governed by its Articles of Association and Charitable Object, which is the furtherance of education for the public benefit, particularly through research, knowledge exchange and public engagement in the field of data science.

Disclaimer

The following document is a summary of the discussion during an exploratory workshop on the merits of setting up an AI scientist grand challenge and how this might be achieved. Whilst the Turing hosted the event the ideas and concepts in this report reflect the discussions at the workshop and are not automatically endorsed by the Institute. References to ‘Nobel’ or ‘Turing’ are aspirational, and convey the ideas exchanged and developed by attendees during the workshop with no formalised ties established to relevant bodies.

Workshop Summary developed by Workshop Attendees

Posing an AI Scientist Grand Challenge: Artificial Intelligence Systems Capable of Nobel-Quality Discoveries

Co-Chairs:

Yolanda Gil (University of Southern California)

Ross King (The Alan Turing Institute)

Hiroaki Kitano (Okinawa Institute of Science and Technology, The Systems Biology Institute, and The Alan Turing Institute)

Sponsored by the Office of Naval Research (ONR Global)

Hosted by The Alan Turing Institute

1. Workshop Goals

The Alan Turing Institute, supported by the Office of Naval Research (ONR Global), convened a workshop to explore the merits of a grand challenge focused on AI for scientific discovery and how this might be achieved.The meeting was organized by Ross King, Yolanda Gil, and Hiroaki Kitano, bringing together three continents, in London on 26-27 February 2020.

Scientific discovery has long been of interest to AI researchers. Herbert Simon (Nobel Laureate and Turing award winner) worked on cognitive aspects of scientific discovery in the 1940s [1]. Edward Feigenbaum (Turing award winner) worked with Joshua Lederberg (Nobel Laureate) and other colleagues in Stanford in the 1960s on automating the identification of organic molecules and learning rules about their mass spectra [2].

Now decades later, driven by vastly more powerful computer hardware and software, AI is having an ever-increasing impact on science [3].

The workshop co-chairs brought forward complementary long-term visions for AI in science: robot scientists that achieve closed-loop automation in terms of hypothesis generation, experimental design, execution and verification [4], the concept of thoughtful AI systems that can be effective partners for scientific discovery [5], and the intriguing proposal for a grand challenge for AI and systems biology to develop AI systems that can make major scientific discoveries significant enough to be worthy of a Nobel Prize [6].

The workshop was attended by a diverse group of thirty invited participants. Several expected invitees had to cancel their travel in the last minute due to travel restrictions related to COVID-19. Twenty participants came from universities, five from industry research labs, and five from government. Nine were from the UK, fifteen from the US, and five from Japan, and one from Europe (outside the UK). Some participants were AI researchers and practitioners with direct experience in science domains, others were scientists with a deep interest in AI techniques. A complete list of workshop participants is included at the end of this document.

The participants of the workshop recognised the need for AI research to accelerate the speed of scientific discoveries in order to address important societal challenges of our generation, such as health and wellbeing, environmental sustainability, and technological innovation. The workshop participants converged on several key points to move this initiative forward:

1. Agreed to pursue the formulation of a Nobel Turing Challenge[1],
2. Recommended the organisation of a 'planetary' initiative to define intermediate milestones and to foster interdisciplinary collaboration and resource sharing, and
3. Proposed potential shorter-term grand challenges in biomedicine, environmental sciences, and materials science.

The rest of this document summarises the discussions on each of these items in turn.

2. The Nobel Turing Challenge: Developing AI Scientists

The Nobel Turing Challenge can be described as follows:

The Nobel Turing Challenge aims to develop 'AI Scientists': AI systems capable of making Nobel-quality scientific discoveries highly autonomously at a level comparable, and possibly superior, to the best human scientists by 2050 [7].

The challenge calls for the development of highly (if not fully) autonomous AI Scientists that can either make discoveries on their own or collaborate with other human and machine scientists in making Nobel quality discoveries. AI Scientists should be capable of independently motivating, making, understanding, and communicating discoveries. This challenge is emblematic and the central focus of the initiative.

Attaining the Nobel Turing Challenge by 2050 is an ambitious but feasible goal: given the current and planned investments in both science and AI research by governments and industry, and the potential societal benefits of accelerating research in both AI and broader areas of science [8].

The Nobel Turing Challenge is named after Alan Turing. He was the founder of AI, and the challenge is reminiscent of the imitation game to test intelligence he proposed [9], known as the Turing test; to develop machines that can carry conversations that are indistinguishable from humans. Instead of focusing on the intelligence required to realize human-level conversations, as the Turing test does, *the Nobel Turing Challenge is focused on the intelligence required to do human-level scientific research, and requires that the science is at the level of the best human scientists.*

The reference to the Nobel Prize points to the requirement that the output of an AI Scientist necessary to succeed in the challenge must be held in the highest regard by the scientific community. It is not meant to focus only on science disciplines covered by the Nobel Prize

[1] 'Nobel Turing Challenge' indicates the ambition of the initiative. There are no formal ties to the Nobel Foundation or any member of the Turing family, though engagement with relevant bodies in relation to this initiative is an aspiration.

categories. Indeed, significant awards exist to recognize exceptional discoveries in other areas of science, for example the Blue Planet Prize in environmental sciences, The Alan Turing Prize in computer science, the Fields Medal in mathematics, or the Crafoord Prize in ecology.

The initial set of science disciplines addressed at the workshop were: 1) biomedicine and biotechnology, 2) Earth, space, and environmental sciences, and 3) materials science. These three areas were recommended as the initial focus of the initiative. Other disciplines, such as chemistry, physics, as well as more applied domains, could be added in the future. For each area, participants discussed short-term and long-term challenges that would be considered as important milestones for discoveries in each field. While work can proceed in each of these areas of science independently, the Nobel Turing Challenge will be used as a central driving force to coordinate and generalize AI research efforts across scientific disciplines.

It was agreed that there is a fundamental distinction between AI Scientists, and AI tools for science. Human scientists routinely use a wide range of AI tools for specific tasks that contribute to discoveries (e.g. data analysis, text extraction, etc.). These AI tools cannot be considered AI Scientists because they only address narrow aspects of the scientific process. More importantly, they lack the autonomy of human scientists in setting their goals, interpreting results, and communicating findings. The clear and explicit goal of the challenge is to develop AI system that have highly autonomous capabilities to perform research, rather than developing AI tools.

3. Establishing a Planetary Initiative for the Nobel Turing Challenge

The Nobel Turing Challenge is an ambitious undertaking that can be best achieved through a global initiative that brings together “talent without borders”: across disciplinary boundaries, across organizational structures, and across nations and continents. In addition, fostering open sharing of resources (infrastructure, instruments, data, etc.) will enable the exploitation of every observation and insight available at a global scale. To facilitate coordination, resource development and sharing, it is desirable to setup an overarching organization that will orchestrate otherwise disconnected discipline-specific or national efforts.

The inspiration for this organization, and strong evidence for its future success, is the very successful RoboCup federation [10]. Workshop co-chair H. Kitano [11] was one of the key initial organizers of RoboCup, and he will enable the direct transfer of organizational, scientific, and practical knowledge from RoboCup to the new initiative. RoboCup was started in 1993 with the aim of using the game of football (soccer) to drive AI and robotics research. RoboCup’s original stated goal was that “by the middle of the 21st century, a team of fully autonomous humanoid robot soccer players shall win a soccer game, complying with the official rules of FIFA, against the winner of the most recent World Cup”. Several leagues were organized for different aspects of the research involved, such as a simulation league and a small robot league, and later a humanoid robot league. Currently, in addition to RoboCupSoccer, the federation runs other leagues under the rubrics of RoboCupRescue, RoboCup@Home, and RoboCup Industrial. The RoboCup president and trustees provide overall coordination and oversight. The competition rules for each league are defined by a committee elected among the participants in that league. The rules for each competition are designed to drive the research in directions proposed by researchers participating in the leagues based on their best intuition about promising directions, and anticipated topics of funded research projects. The rules for each competition are updated each year as needed to push the research. Annual research symposia enable competition participants to share ideas and advances. The best teams release software and data to other teams to provide a detailed account of their approaches immediately following the annual competition. Many teams collaborate to develop research infrastructure, such as simulators or program interfaces. Companies have often also been involved. No monetary reward is offered to the winning teams. RoboCup 2019 attracted “3,500 dedicated scientists and developers from more than 40 countries.”

Given the significance of this initiative, a network of global researchers and funding agencies shall be formed. There are ongoing and planned funding programs that explicitly support research in AI for science. In the US these include the NSF (Expeditions, National AI Research Institutes), DARPA (e.g., SIMPLEX, Big Mechanism, and AIRA programs), as well as growing efforts at the DOE, NIH, IARPA, and DOD. There are potentially relevant funding sources in other countries as well. The Japanese Moonshot program, administered by the cabinet office and run by the Japan Science and Technology Agency (JST), explicitly calls for AI scientist as one of their grand challenges. Funding agencies with the UK have also expressed interested in the AI grand-challenge, and discussions are on-going. However, the Turing Nobel challenge needs to continue and grow beyond the scope of regular scientific funding programs that typically last 5 years. A mechanism to ensure continuity of the scientific objectives, organization, community, and supporting ecosystems should be designed and implemented from the beginning. It is important to continue carrying the torch until the ultimate goal of the challenge is achieved, and its benefits realized.

To jumpstart the Nobel Turing Challenge Initiative (NTI), a steering committee could be created initially composed of scientists from the United Kingdom, the United States of America, and Japan, but will be extended beyond initial regions as soon as the structure of the initiative is established. The three main coordinating centers in the three regions could set up and manage the initial phases of the initiative: The Alan Turing Institute, the University of Southern California (NTI-USC), and the Okinawa Science and Technology Institute (NTI-OIST). These centers are committed to contributing initial infrastructure and resources for the federation. The coordinating centers may delegate specific functions to other institutions when appropriate. A dedicated international body for the initiative could be eventually created.

A series of workshops should be held in rotation in each region hosted by the coordinating centers to 1) catalyze a community, 2) design appropriate intermediate grand challenges, 3) develop the organizational and governing structures of the initiative, 4) promote infrastructure development and resource sharing, 5) create opportunities for work across disciplinary boundaries, 6) broker team formation, and 7) coordinate funding programs across continents. Workshops could also be organized outside of these regions when host organizations are identified.

4. Potential Intermediate Grand Challenges

Achieving the Nobel Turing Challenge will take several decades of hard research. Along the way, intermediate grand challenges will need to be defined to structure specific problems and demonstrate tangible progress. A key criterion for these challenges will be the synergistic aspects of discovery in scientific domains, and the development of the AI scientist technology. Thus, progress in a grand challenge will drive progress in both science and AI research.

From an AI perspective, intermediate grand challenges could be defined to drive the development of diverse intelligent capabilities:

- Demonstrating and formalizing expert human-level knowledge about some area of science
- Robots for laboratory experiments driven by AI systems for scientific discovery
- Scientific communication: generate, summarize, discuss, review, critique/compare scientific articles
- Generating compelling scientific questions and/or research plans to answer questions

These kinds of challenges should be defined for specific scientific disciplines, which will promote synergistic collaborations across challenges in different scientific areas - discussed next. It is essential that the challenge is seen as adventurous to both AI researchers and domain science experts, such that they work together to develop the new platforms to enable discovery. In the case of Materials and Chemicals the UKRI/EPSRC AI3SD Network (www.ai3sd.org) which has at its core the aim to bring these communities (Chemistry, Materials and AI) together to facilitate AI for Scientific Discovery, is very keen to work with the Nobel Turing Challenge Initiative in pursuit of these objectives. A key principle of the AI3SD Network is to enable a two-way communication between scientific domain researchers and computer science researchers. Key to this is quality data that is well described and machine readable to enable interdisciplinary interactions. Setting up the necessary community agreed ontologies is an essential first step, which in collaboration with the international standards agencies (e.g. International Union of Pure and Applied Chemistry IUPAC) is one of our current network projects.

Each of the science areas below will likely drive research in some areas of AI more heavily than others, for example knowledge synthesis and robotics in biology, machine learning in geosciences, and constraint reasoning in materials science, although all areas would benefit from any components of an AI Scientist.

A significant milestone would be when AI systems generate unexpected discoveries. In these cases, the system's ability to make a case for its novelty would be a very valuable capability.

4.1 Biology/Biomedicine

Nobel-quality discoveries in biomedicine/Biotechnology that could be targets for AI Scientists could include: discovering new principles in biology, reversing aging, curing cancer, understanding how brains generate behavior, highly efficient drug discovery and vaccine production, the creation of organisms with specific functions, etc. Possible intermediate Nobel Turing Initiative (NTI) grand challenges could include:

- *Challenges driven by a medical/societal motivation*: These could include the development of a task-specific autonomous system that could develop a vaccine for a new virus within weeks, discover a safe drug against a new disease within weeks, design drugs or treatments for escaping specific undesirable cell states caused by disease, synthesize organisms with designed functions, and understand some particular cognitive impairments.

- *Challenges driven by development of a highly autonomous and integrated research platform*: An integrated AI / robotics biology laboratory could be a standard platform to facilitate a broad range of research, protocols and data exchange, and knowledge sharing. An initial configuration could target microbial and cell culture-based research including yeast, mammalian stem cells, and synthetic biology. This challenge shall serve three purposes: (1) technology development, (2) proof-of-concept for automation of science, and (3) showcase for transforming research in microbial and cell-based systems. This challenge is closely related to the proposal of “AI-enabled self-driving laboratories” [12].

- *Challenges driven by existing biomedical repositories*: Genbank, the Cancer Genome Atlas (TCGA) [13], the Clinical Proteomic Tumor Analysis Consortium (CPTAC) [14], and the UK Biobank [15] are all very large data repositories that could be analyzed using AI to do systematic data analysis, which would be an important component of AI Scientists. While this may focus on AI and bioinformatics challenges, it enables a broader range of researchers to get involved in the challenge.

One flagship project within the initiative would be the development and operation of highly autonomous AI/Robotics experimental laboratories in selected key research centers. These would enable the separate development of the hardware and software for AI scientists, and researchers working on new AI software would not have to deal directly with the difficulties associated with running physical laboratories. Infrastructure that would need to be developed includes open and standard configurations, software and cloud platforms, APIs, and protocol sharing approaches to enable scalable and shareable AI scientist prototype components, initially targeting basic scientific research on microbials, mammalian cellular systems, and synthetic biology. An initial core center may be established in one or more of the coordinating centers, and later expanded to other sites as the proof-of-concept is validated. These laboratories would be remotely accessible.

Another project could focus more on AI and bioinformatics. Initial milestones in this area could be developed around the existing DREAM Challenges [16], which are well-known competitions for biomedical questions. Important initial progressive multi-year challenge milestones could be to win one of DREAM Challenges initially, and then wins multiple DREAM Challenges by enhancing generality of the system. Special NTI-DREAM challenges could potentially be created to address NTI requirements. Although the existing DREAM Challenges are fundamentally algorithmic, participants must have significant domain knowledge to succeed, and this makes them hard to access for AI researchers. NTI challenges could propose new formulations that would include knowledge resources to make the challenges accessible to AI researchers and/or AI systems. Many past DREAM Challenges were focused on systems biology [17], and new ones could be proposed in this area. Special criteria for NTI challenges could be formulated to address specific requirements that would lead to developing important components of AI Scientists that shall span both informatics and robotics components. For example, an intermediate NTI challenge could be created that would take advantage of sophisticated robot laboratories.

The three coordinating centers suggested above could contribute substantial resources and infrastructure to jumpstart efforts in these initial challenge areas. NTI-USC has demonstrated an AI system called DISK [18] [19] for automatically generating solutions for the NCI-CPTAC DREAM Proteogenomics Challenge [20]. NTI-OIST has spearheaded the Garuda [21] platform through an alliance of members from Japan, UK and Europe, and USA (including USC). The Alan Turing

Institute convenes academic leaders in artificial intelligence from across all UK universities, as well as industry and government partners. . Significant resources already exist in biomedical research (ontologies, databases, and software tools), although not all of them are directly usable by AI systems and would need to be extended to accommodate NTI challenges.

4.2 Geosciences

Possible intermediate NTI challenges in geosciences [22] could result from AI innovations in areas such as model-driven sensors, knowledge graphs that link data and models, and theory-driven machine learning. There has been a recent surge of interest in theory-driven machine learning [23], an area where deep learning models are constrained by physical laws and theories. Important initial progressive multi-year NTI milestones to drive this area could include:

- *Grand Challenge*: Create accurate hydrological predictions (e.g. for floods and droughts) for any river basin in the world.
    - Hydrological Prediction in Ungauged Basins (PUB), modeling sparsely or completely basins that have no historical observations from gauges or sensors (as is the case in many areas of the globe).
    - Hydrological Prediction Under Change (PUC), creating models that take into account the feedback that occurs between the cycling of water in nature and the use/management of water for human needs.
- *Grand challenge*: Seasonal to subseasonal weather prediction.
    - Breaking the chaos barrier in weather prediction, that is, increasing the time period in which weather predictions can be made (which is currently days) to several months.
    - Integrated models that use subseasonal weather predictions as drivers, e.g. for water resources, floods, droughts, transportation, epidemiology, etc.

4.3 Materials Science

Materials science is probably the area of science where the greatest interest has been shown in developing AI scientists. Possible intermediate NTI grand challenges in this area could include the following:

- *Grand challenge*: Given a target structure, decide how to synthesize it successfully, including iterating synthesis attempts until it is successful.
- *Grand challenge*: An AI system that runs without human input/decisions to produce believable ab initio calculations for many different systems and properties.
- *Grand challenge*: Given an *ab initio* calculation, an AI system provides a critique, judgement of whether the calculation is believable (this is the reverse problem of the previous one).
- *Grand challenge*: Autonomously produce large numbers of theoretical materials with scientifically interesting properties.
- *Grand challenge*: Closed system autonomous system that produces a material that is significantly better than existing ones for practical purpose.
- *Grand challenge*: Fully automate the entire electron microscopy process.

- *Grand challenge*: The inverse problem to design molecules that have specific function and then control their assembly to create materials and devices. Applications range from life sciences (e.g. drugs, anti-microbial agents, pesticides) through to Physical Sciences (e.g. molecules for efficient solar energy conversion, batteries, surface coatings for corrosion resistance or minimize friction).

5. Next steps

Following this workshop, the Alan Turing Institute, with ONRG support, has launched a study to develop a shared road map towards the achievement of NTI. Over the course of 2020 there will be a number of workshops with academic and leading thinkers from other, relevant domains. We welcome interest from those who would like to join these sessions, keep abreast of latest updates or help the development and the launch of the abovementioned initiative to share their interest via AIforScientificDiscovery@turing.ac.uk.

7. List of Workshop Participants

Mark Briers, The Alan Turing Institute, UK

Andrew Briggs, University of Oxford, UK

Alan Bundy, University of Edinburgh, UK

Ayodeji Coker, Office of Naval Research, USA

Matthew Daniels, Georgetown University, USA

Hushpreet Dhaliwal, The Alan Turing Institute, UK

Jared Dunnmon, Stanford University, USA

Saso Dzeroski, Jozef Stefan Institute, Slovenia

James Evans, University of Chicago, USA

Jeremy Frey, University of Southampton

Peter Friedland, Air Force Office of Scientific Research, USA

Yolanda Gil, University of Southern California (Co-Chair), USA

Fredrick Gregory, Army Research Office, USA

Carla Gomes, Cornell University, USA

Lawrence Hunter, University of Colorado, USA

Nick Jennings, Imperial College London, UK

Ross King, The Alan Turing Institute (Co-Chair), UK

Hiroaki Kitano, Okinawa Institute of Science and Technology, The Systems biology Institute and The Alan Turing Institute (Co-Chair), Japan

Pushmeet Kohli, Google DeepMind, UK

Vipin Kumar, University of Minnesota, USA

Patrick Riley, Google, USA

Patrick Rose, Office of Naval Research, USA

Jon Rowe, The Alan Turing Institute, UK

Jun Seita, Riken, Japan

Barl Selman, Cornell University, USA

Larissa Soldatova, University of London, UK

Michael Spranger, Sony, Japan

Koichi Takahashi, Riken, Japan

Adrian Weller, University of Cambridge and The Alan Turing Institute, UK

Chris Williams, University of Edinburgh, UK

Jiangying Zhou, DARPA, USA

---

[1] Simon, Herbert A. "Models of Discovery and Other Topics in the Methods of Science." Springer, 1977. ISBN 978-94-010-9521-1.

[2] Lindsay, Robert K., Bruce G. Buchanan, Edward A. Feigenbaum, and Joshua Lederberg. "Applications of Artificial Intelligence for Organic Chemistry: The Dendral Project." McGraw-Hill, 1980. ISBN 978-0070378957.

[3] Science. "AI Transforming Science." Special Issue, 357(6346), 7 July 2017. https://science.sciencemag.org/content/357/6346

[4] King, Ross. D. et al. "The Automation of Science." Science,324(5923), 3 April 2009. http://doi.org/10.1126/science.1165620

[5] Gil, Yolanda. "Thoughtful Artificial Intelligence: Forging A New Partnership for Data Science and Scientific Discovery." Data Science, 1(1). 2017. http://doi.org/10.3233/DS-170011

[6] Kitano, Hiroaki. "Artificial Intelligence to Win the Nobel Prize and Beyond: Creating the Engine for Scientific Discovery." AI Magazine. 37(1), 2016. https://doi.org/10.1609/aimag.v37i1.2642

[7] The original formulation of The Nobel Turing Challenge by Kitano prior to the workshop was "to develop AI system capable to autonomously perform high-impact research primarily in the area of physiology and medicine some of which may worth the Nobel Prize or beyond, and actually win it by 2050." The workshop version was modified to include broader areas beyond the scope of the Nobel Prize. In addition, any implication on winning the prize was deemphasized .

[8] Gil, Yolanda and Bart Selman (Eds). "A 20-Year Community Roadmap for Artificial Intelligence Research in the US." Computing Community Consortium (CCC) and Association for the Advancement of Artificial Intelligence (AAAI) report, published in arXiv, https://arxiv.org/abs/1908.02624v1, August 2019.

[9] Turing, A. M. "Computing Machinery and Intelligence." Mind 49: 433-460, 1950. http://cogprints.org/499/1/turing.html

[10] https://www.robocup.org/

[11] Kitano, Hiroaki, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, Eiichi Osawa, and Hitoshi Matsubara,. "RoboCup: A Challenge Problem for AI." AI Magazine, Vol. 18, No. 1, Spring 1997.. DOI: https://doi.org/10.1609/aimag.v18i1.1276 (https://www.robocup.org/) [12] Stevens, R., et al., "AI for Science," Department of Energy, February 2020. https://anl.box.com/s/f7m53y8beml6hs270h4yzh9l6cnmukph

[13] https://www.cancer.gov/about-nci/organization/ccg/research/structural-genomics/tcga

[14] https://proteomics.cancer.gov/programs/cptac

[15] https://www.ukbiobank.ac.uk/

[16] http://dreamchallenges.org/

[17] https://collections.plos.org/dream

[18] Srivastava, Arunima.Ravali Adusumilli, Hunter Boyce, Daniel Garijo, Varun Ratnakar, Rajiv Mayani, Thomas Yu, Raghu Machiraju, Yolanda Gil, and Parag Mallick. "Semantic Workflows for Benchmark Challenges: Enhancing Comparability, Reusability and Reproducibility." Proceedings of the Pacific Symposium on Biocomputing (PSB), 2019. https://psb.stanford.edu/psb-online/proceedings/psb19/srivastava.pdf

[19] Gil, Yolanda, Daniel Garijo, Varun Ratnakar, Rajiv Mayani, Ravali Adusumilli, Hunter Boyce, Arunima Srivastava, and Parag Mallick. "Towards Continuous Scientific Data Analysis and Hypothesis Evolution." Proceedings of the Thirty-First Conference of the Association for the Advancement of Artificial Intelligence (AAAI), San Francisco, CA, 2017. https://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14839/14046

[20] https://www.synapse.org/#!Synapse:syn8228304/wiki/413428

[21] Ghosh S, Matsuoka Y, Asai Y, Hsin KY, Kitano H., Software for systems biology: from tools to integrated platforms. Nature Reviews Genetics (2011) doi: 10.1038/nrg3096, http://www.garuda-alliance.org/

[22] Gil, Yolanda, Suzanne Pierce, et al. "Intelligent Systems for Geosciences: An Essential Research Agenda". Communications of the ACM, 62(1), 2019. http://doi.org/10.1145/3192335

[23] Willard, Jared, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. "Integrating Physics-Based Modeling with Machine Learning: A Survey." https://arxiv.org/abs/2003.04919