Workshop Structure

We propose to largely copy the successful structure of the 1st Nobel-Turing Workshop in London 2020 which was based on discussions and breakout sessions:

October 3rd
p.m.
14.00

Introduction & Welcome (Ross)

Tour de Table Participants intros.

The Nobel Turing Challenge (Hiroaki)
Plenary discussion

15.00

Coffee Break

15.15

Foundation Models and their Implications for AI & Science
Plenary discussion
Breakout Sessions

Notes

- *Within the last year the world has been stunned by the success of Large Language Models (LLMs). They have achieved breakthrough performance on a wide range of tasks that require human intelligence.*
- *The success of LLMs is surprising, they have no explicit model of the world, no explicit internal symbols, or any grounding in the world.*
- *They have read the whole scientific literature, and can find previously unrecognised connections.*
- *"Hallucinations" are not necessarily a problem, as truth in science is tested by physical experiment.*
- *Integration with symbolic reasoning.*
- *Integration with experimental infrastructure – agency.*

Evening
17.00

Wine Tasting

19.00

Dinner

%%

October 4th
a.m.
9.00

Discovery Science History

Notes

- *Dendral/Meta-Dendral*
- *Bacon*
- *Eurisko*
- *Alphafold*

- *Nature magazine named ML as the 'breakthrough' of the 2010s decade.*
- *Science magazine named AlphaFold the 'breakthrough' of the year.*
- *The keyword 'machine learning' in PubMed now gives >128,000 refences, almost all of from the last decade.*
- *The use of AI, especially ML, is becoming characteristic of 21st century science.*

9.30

Progress in AI & Science since 2020

Plenary discussion

Notes

- *AI is making substantial impact on the research enterprise in disciplines spanning materials research, chemistry, climate science, biology, and cosmology.*
- *These advances have involved the development of AI-accelerated analysis of massive datasets, Laboratory automation to efficiently perform physical experiments, the fusion of many disparate data streams, integrated systems able to automate cycles of simple scientific research (AI scientists, Robot Scientists, Self-driving labs ,...) ...*
- *This discussion will highlight the recent achievements of AI in these domains to identify where AI is currently in the process of scientific discovery.*

10.30

Coffee Break

10.45

Research Goals for AI & Science

Plenary discussion

12.00

Breakout Sessions

Notes

- *High-Level Goals*

- *Demonstrating and formalizing expert human-level knowledge about specific areas of science.*
- *Robots and automation 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.*
- *Are integrated systems for long-term research efforts necessary to achieve scientific breakthroughs or whether focusing on individual components will suffice?*

- *For an AI system to have the knowledge, skills, and abilities to perform independent scientific discovery, it should be able to: 1) Determine research gaps from existing literature, 2) Form unique hypotheses (or intend to conduct hypothesis-generating research), 3) Develop and carry out an experimental plan, 4) Interpret the results, and 5) Draw conclusions both about the original research and how new research could follow on from these results.*
- *It is well-established that AI can be brittle and extrapolats performs poorly, and explainability and interpretability are other critical limitations.*
- *What are the limitations of AI in conducting independent scientific discovery and what could be done to mitigate these problems?*
- *Making the research process faster, vastly more experimentation.*
- *Addressing novel difficult scientific problems*
- *Increasing scale and throughput while decreasing cost and compute needs*

13.00

Lunch

p.m.

14.00

Research Goals for AI & Science

Plenary discussion

14.30

Breakout Sessions

Notes

- *What technologies does the community requires to advance the use of AI in research?*
- *LLMs*
- *Formal languages for scientific knowledge*
  - *Logic and ontologies*
  - *Probabilities*
  - *Mathematics*
  - *Turing machines*
- *Natural languages*
- *Hypothesis Formation and ML*
  - *Logical models*
  - *Deep Neural Networks*
- *Deduction and Model Simulation*
- *Designing Efficient Experiments*
- *Designing Novel Experiments*

15.00

Coffee Break

15.15

Research Goals for AI & Chemistry

Plenary discussion

15.45

Breakout Sessions

Notes

- *We are fortunate to have a number of world experts in Chemistry*
- *Can we identify any particular areas of Chemistry that are particularly suited for rapid progress through AI?*
- *Prediction/Reaction discovery?*
- *Synthesis planning?*

16.15

Nobel/Turing Prize Level Research

Plenary discussion

Breakout Sessions

Notes

- *Some Philosophers such as Kuhn believe that Nobel Prize quality science is qualitatively different from normal science. If this is correct*

- *Is the situation more like in chess?*
- *Do we need AGI?*

18.00 Live music

Evening

Dinner

%%%

October 5th

a.m.

9.00

Social Aspects of AI and Science

Plenary discussion

10.00

Breakout Sessions

Notes

- *The potential to automate scientific discovery opens up profound social and ethical questions.*
- *The increased use of AI in science and engineering creates new epistemic, methodological, l challenges for researchers.*
- *What are the legal/regulatory, social, and policy challenges associated with deploying artificial intelligence technologies to enable scientific/engineering research design and discovery?*
- *'The use of AI in the research process is blurring the boundaries among disciplines, rendering single-themed funding competitions, scientific reviews, and research programs less useful because domain expertise is unable to properly review the inter-disciplinary nature of AI-driven science.'*
- *In the scientific labour market some job displacement may be inevitable, and that the jobs of scientists and engineers may be transformed by the need to collaborate with AI systems.*
- *In the legal domain: 'several hurdles are emerging around the use of AI in science and engineering, most acutely in the domain of intellectual property (IP). Traditional IP frameworks are difficult to apply to machines capable of innovation or originality. Patents and copyrights have become key assets in an increasingly digital and intangible economy; however, these instruments were originally conceived as incentives with humans in mind'*

10.30

Coffee Break

10.45

Ethical Aspects of AI and Science

Plenary discussion

Breakout Sessions

Notes

- *The increased use of AI in science and engineering creates new ethical challenges for researchers*
- *AI capabilities are raising concerns about potential misuse, including the production of erroneous data or misleading results, the development of biological or chemical agents of concern, and the loss of human oversight.*
- *What could go wrong? Small and big*
- *Are autonomously interacting AI agents controllable?*
- *How can we direct these efforts in the direction of accelerating discovery in a responsible way? What governance principles could be developed?*
- *The increased deployment of AI systems for science and engineering risks perpetuating discrimination or biases both within the R&D environment and in broader society*

- *As AI evolves, it will interact with and influence our world in ways that are unexpected, some for the better and some for the worse. Over the past several years, it has become clear that machine learning methods have the potential to embed the biases that arise in society into their resulting models, such as the systemic discrimination and subjugation of certain populations.*

13.00

Lunch

p.m.
14.00

A Roadmap for AI and Science
Plenary discussion

14.30

Breakout Sessions

Notes

- *To succeed in the Nobel Turing Challenge do we need AGI?*
- *What are appropriate intermediate grand challenges?*
- *This would require novel AI research in areas such as harnessing sophisticated scientific knowledge, mastering advanced scientific processes and methods, and proficiency in scientific communication and collaboration.*
- *AI scientists are much more easily multiplied than human scientists. It is also potentially much easier to make them fully collaborate with each other. With AI scientists it is possible to completely capture and digitally curate all aspects of the scientific process: background theories, hypotheses, motivation for experiments, results, conclusions, etc. This formalisation would enable AI scientists physically on different continents to instantly communicate and collaborate, similar to the vision of self-driving cars.*

- *A Prize?*
  The most famous and prestigious prizes in science are the Nobel prizes. These are exclusively for human scientists: Alfred Nobel's will refers to 'persons' making discoveries and inventions.
  I suggest an annual scientific prize exclusively for AI systems: an annual prize for the AI system that made the most important discovery or invention in a field of science

  Raise the profile of the Challenge.
  It might be possible to get money for this from a large company or a foundation such as the Schmidt. Over time the quality of science may overtake that of the Nobel prize.
  The Nobel Turing Grand Challenge will be met when the AI prize is awarded to research that is recognised to be equal or better than that of the Nobel prize.
  It would be fitting if the AI system prize for science was based in Sweden.

15.00

Coffee Break

15.15

The Nobel-Turing Federation
Plenary discussion

Notes

- *The community is currently splintered.*

- *Hiroaki's 'Planetary Initiative'.*

The inspiration for this organization, and strong evidence for its future success, is the very successful RoboCup federation.
Catalyse a community.
Develop the organizational and governing structures of the initiative.
Promote infrastructure development and resource sharing.
Create opportunities for work across disciplinary boundaries.
Broker team formation.
Coordinate funding programs across continents. Workshops could also be organized outside of these regions when host organizations are identified.

Funding an International Community
Plenary discussion
Breakout Sessions - Swedish/EU funding, US funding, Japanese funding.

Notes

- *Existing US DOE initiatives?*
- *A possible US NSF initiative?*
- *Japanese initiative?*

p.m.
Evening

19.00
Dinner

%%

October 6th
a.m.
9.00
Laboratory Automation, Cloud Automation
Plenary discussion

10.00
Breakout Sessions

Notes

- *One of the core insights in developing the scientific method in the 17th century was the essential need to do physical experiments*
- *Many types of scientific work fundamentally involve physical experiments. In these fields, AI systems that are not able to directly run experiments may be quite limited.*
- *One of the powers of an AI system is its ability to simultaneously manage more ongoing experiments than a human can.*
- *Cloud automation.*
- *Automation of experimentation has therefore been a frequent topic over the years, most recently under the moniker of "self-driving labs".*
- *Most automated experimentation systems have significant limitations, can these be removed?*
- *Experimental Protocols and reproducibility*

10.30
Coffee Break

10.45

Human and AI Scientist Collaboration

Plenary discussion

12.00

Breakout Sessions

Notes

- *A key motivation for investing in AI for science is that AI systems 'think different'. Human scientists, at least all modern ones, are all educated and trained in much the same way. This probably imposes unrecognised cognitive biases is how they approach scientific problems. AI systems have very different strengths and weaknesses from human scientists.*
- *How should AI scientists should be designed to fit the science ecosystem and to respect the boundaries desirable in AI systems for science.*
- *AI scientists have complementary skills to human scientists: they can flawlessly remember vast numbers of facts. They can execute flawless logical reasoning, and near optimal probabilistic reasoning. They can learn from vast amounts of data. They are can read (with some understanding) every scientific paper, indeed everything ever written.*
- *Current AI systems lack deep models of the world, they don't know what they are or what they are doing. They lack of flexibility*
- *For the foreseeable future, human and AI scientists working together will probably be able to outperform either human scientists alone or AI scientists.*
- *How will the role of human vs machines change in the design, the execution and analysis of experiments?*
- *The analogy with chess is illuminating. In 1996 a computer (IBM's Big Blue supercomputer) for the first time beat the world's chess champion (Gary Kasparov) in a game under tournament conditions. Now, computer hardware has improved so much that your mobile phone can beat the world champion. However, humans can still complement chess computers. Advanced/centaur chess is a form of chess where humans and chess computers work together as an integrated team. The advantage goes to those centaur teams that best combine the strategic flexibility of the human player with analytical supremacy of the computer. Centaur chess are better than chess than computers are alone.*
- *It is often said, that the true test of understanding a concept is the ability to teach it. If an AI system could explain a novel scientific idea to a person then we would be happily credit the AI with the discovery.*
- *The AI pioneer Donald Michie distinguished in ML between non-symbolic and symbolic models, and proposed 'ultra-strong machine learning'; where a newly learnt hypothesis is expected to not only be accurate, but to also to demonstrably improve the performance of a human being provided with the learned hypothesis.*

13.00

Lunch