Abstract

Abstract:

Propelled by the success of deep learning, artificial intelligence (AI) is poised to transform science. Science strives to discover theories that predict the outcome of experiments and allow us to control the future. In harmony with these goals, deep learning is driving a revolution in how we build predictive and decision-making models, creating potential to alter and accelerate scientific discovery.

Historically, scientific theories have needed to be simple and compact enough to be written down and understood. This is changing as we explore the idea that scientific theories can be embodied in the weights of massive neural networks: our best theory for protein folding now lives in millions of parameters of a transformer network, largely inscrutable to humans. Hot on the heels of proteomics, deep learning is poised to transform many fields of science and mathematics.

The workshop will bring together leaders in deep learning with top scientists and mathematicians in their respective domains to understand advances in the physical sciences and mathematics driven by deep learning, such as protein folding, chemistry, quantum physics, geometry, and theorem proving. We will examine methodologies for embedding physical and mathematical intuitions into neural network architectures and algorithms: What are the best ways to build inductive biases into networks? Finally, we'll examine the potential of large language models to help us summarize, synthesize and analyze the scientific literature itself.

Schedule

Schedule:

Sun	9:30am−noon	First lecture	Neil Lawrence, Tim Lillicrap
	7:30pm−9pm	Neuroscience	Tim Lillicrap
Mon	9:30am−noon	Math and RL	He He, Siva Reddy, Sanjeev Arora
	7:30pm−9pm	AI Co-worker	Chris Pal, Laurent Charlin
Tues	9:30am−noon	AI Scientist	Roberta, Larry Zitnick
	7:30pm−9pm	Chemistry	Alán Aspuru-Guzik
Wed	9:30am−noon	Material Science	Santiago Miret, Alannah Hallas
	7:30pm−9pm	Protein Folding and Engineering	Yisong Yue
Thur	9:30am−noon	Discussion	Siva Reddy, Tim Lillicrap, Everyone
	7:30pm−9pm	Drug Discovery	Diego di Bernardo

Questions

1. **Where will scientific paradigms live in the age of AI?**
   As knowledge moves from textbooks and equations into simulations and large models, is this a genuine paradigm shift or just a new storage medium?
2. **What role should human understanding and agency play in AI-driven science?**
   If we can use models we don't fully understand, what kinds of understanding do we still demand from scientists and trainees?
3. **Omni-models or domain specialists?**
   Is the future "one model to rule them all" across domains, or a collection of highly specialized models, and what are the trade-offs?
4. **Is there an emerging playbook for AI-for-Science?**
   Are we converging on a standard recipe, accurate but slow simulators, amortized surrogates, differentiable pipelines, and how do we know when these surrogates truly generalize?
5. **Can AI grapple with open-ended discovery?**
   Beyond supervised prediction, to what extent can current or near-future systems generate meaningful hypotheses, concepts, and research directions?
6. **What tools do we have, and what tools are missing?**
   Which existing AI/ML capabilities are already reshaping scientific practice, and what critical tools (for uncertainty, causality, interpretability, or interfaces) are still absent?

Guidance for speakers

Guidance for speakers:

A typical talk could follow the format below.
Some talks won't easily follow this mould. That's not an issue.

1. Describe some key questions in the field of science that you're working in. In broad strokes, how could machine learning be applied to answer these questions?
    a. Assume that most of the audience will have the essential mathematics, but will not know your scientific discipline. Give them enough of an overview that they understand how machine learning is being used to try to solve problems in your field.
    b. Where possible, outline the key mathematical questions/issues in formal terms so that connections can be made at this level.
    c. Expect to spend ~20-30 minutes on #1a-c.
2. Describe in detail how you (or other labs) have used machine learning models in your scientific research.
    a. What "narrow" models have been created and deployed? Are they useful now? Do you expect them to be useful soon?
    b. Have you found "generalist" models (e.g. chatbots) useful for any stages of your research?
3. Are there any general lessons you've learned from applying (or trying to apply) ML models to your field? That is, are there ideas or patterns that you expect to generalize to other areas of science? For example, are there functional forms used for building models that might generalize to other areas?
4. Do you think that machine learning will fundamentally change your scientific discipline? If yes, how so? Over what time course? If no, then why not?

Info (dates etc)

Dates: March 13th-20th
Webpage: https://mcgill-nlp.github.io/barbados-workshop-2026/

Topics / People

Barbados Invites 2026

5x Long lectures
5x Medium lectures

Lecture ideas:

1. AI for Science Intro
    a. Common themes
    b. What could LLMs be useful for?
    c. Other kinds of specialized models
    d. Neil, Tim, Siva?
2. Quantum materials
    a. Alannah
3. Protein folding
    a. John Jumper (one last try) [INVITE]
4. Protein Engineering
    a. Yisong Yue (Caltech)
5. Pharma → point of view on what AI can/can't do
    a. Diego di Bernardo at Tigem
6. Chemistry/Material Science
    a. Santiago Miret [INVITE]
7. Foundation models in biology
    a. Jean-Philippe Vert
8. Math
    a. AlphaGeometry: He He
    b. IMO Gold: Could be Ed?
9. Physics
    a. Danilo Rendez? [INVITE]
10. New machine learning ideas for science
    a. C. Lawrence Zitnick?
    b. Roberta (DeepMind)
11. Framework for science with AI
    a. Neil Lawrence & Denis
12. Debate
    a. Neil Lawrence, Tim, Everyone
13. LLMs for Science
    a. Siva
    b. Chris Pal
14. LLMs working on math and reasoning {IMO / Math etc}
    a. Tim
15. Role of ML in neuroscience theories → debate about nature of science
    a. Tim

May 13 2025

Topics:

1. Materials science:
    a. Max Welling's startup;
    b. Alan Aspuru-guzik;
    c. Santiago Miret - Intel? (ask Siva?)
2. Protein folding:
    a. AlphaFold: John Jumper
    b. Small molecules: Josh Abramson (jabramson@google.com)
3. Chemistry
    a. Basic/Organic: ???
    b. Drug discovery:
        i. Kyunghyun Cho (NYU)
        ii. Jian Tang (ask Siva?)
        iii. Diego di Bernardo at Tigem? (ask Denis?)
4. Physics
    a. Quantum matter: Danilo Jimenez Rezende
    b. Quantum materials: Alannah Hallas
    c. Physics simulators?: (Tim to look for this; no rush)
5. Math
    a. Mathematics!: Terence Tao
    b. AlphaGeometry (Tim can find a name):
        i. Quoc Le
        ii. He He
    c. Theorem proving (formal in Lean) –
        i. Thomas Hubert
        ii. Sean Welleck
        iii. David Silver
    d. RL for Reasoning (informal / LMMs)
        i. Nathan Lambert
6. LLMs for Math and Reasoning:
    a. Doina Precup
    b. Joelle Pineau
    c. Literature synthesis:
        i. Chris Pal
        ii. Akari Asai (ask Siva)

Exclude:
Neuroscience, Health, Biology, Psychology, Medicine, softer sciences

Submission material

Submission Material

Please describe the academic focus of the workshop or conference, why the workshop or conference is timely, information about the anticipated attendees, and whether trainees will be included.

1. Academic Focus, Timeliness, and Attendees:
With the success of deep learning, artificial intelligence is poised to transform scientific discovery. This workshop will explore how advanced neural architectures can serve as frameworks for scientific theories across fields such as proteomics, quantum physics, and mathematics. By bringing together leaders in AI with scientists and mathematicians, we seek to accelerate breakthroughs in predictive modeling and decision-making, as recently demonstrated in protein folding (AlphaFold). The workshop is especially timely given the mounting potential of neural-based language models to help us reason, summarize, synthesize and analyze the scientific literature itself (e.g. R1 from Deepseek, LitLLM from Mila/ServiceNow and AI, Alpha co-scientist from Google).

Attendees will include top-level academics from institutions worldwide and world-caliber researchers from industry labs. They will discuss methods for building inductive biases and unifying scientific knowledge into novel neural network models. Graduate students from multiple universities will be integral participants, benefitting from direct mentorship and collaboration on leading-edge research. Our objective is to surpass our typical proportion of 20% of trainees. The workshop's immersive environment at Bellairs fosters intensive collaboration, open dialogue, and hands-on experimentation, providing junior researchers with an unparalleled opportunity to engage with global leaders and shape the next generation of AI-driven scientific innovation.

2. Additional Relevant Information:
Over its 35-plus-year history, our workshop at Bellairs has led to significant scientific contributions, including two books, fourteen sets of lecture notes used worldwide for graduate instruction, and dozens of high-impact papers, one recently published in *Nature Neuroscience*. Past attendees include a Fields medalist, Turing Award recipients, and winners of other prestigious honors such as the Abel and Gödel Prizes; one year ago John Jumper, who subsequently received the Nobel Prize in Chemistry, had accepted to be a speaker at our 2025 meeting, and we hope he will be able to come next year. Information about 33 of our past meetings can be found at https://www.cs.mcgill.ca/~denis/barbados.html. The current proposal is co-organized with Siva Reddy, early-career researcher at McGill School of Computer Science, and Timothy Lillicrap, Director of research at Google DeepMind. Several researchers who were contacted have expressed their high desire to join the meeting, such as Zemel (Columbia), Lawrence (Cambridge), Schölkopf (Max Planck), Arora (Princeton), O'Donnell (CMU), Ermon (Stanford).

Our unique approach fosters collaboration among established leaders and emerging scholars, with graduate students benefiting from close mentorship and networking. We also actively engage the local Barbados community and government; the Minister of Science and the Prime Minister's chief of staff have both visited recent workshops, and we are working with the University of the West Indies for next year's meeting, with Prof Simon Anderson, Director of the Chronic Disease Research Centre. Media outreach remains a priority: in 2024, a journalist from *La Presse* spent a week covering the workshop, resulting in five articles that highlighted Bellairs' distinctive atmosphere and the presentation of our scientific discussions.

<ABSTRACT>
AI for Accelerating Science

Propelled by the success of deep learning, artificial intelligence (AI) is poised to transform science. Science strives to discover theories that predict the outcome of experiments and allow us to control the future. In harmony with these goals, deep learning is driving a revolution in how we build predictive and decision-making models, generating obvious potential to alter and accelerate scientific discovery.

Historically, scientific theories have needed to be simple and compact enough to be written down and understood. This is changing as we explore the idea that scientific theories can be embodied in the weights of massive neural networks: our best theory for protein folding now lives in millions of parameters of a transformer network, largely inscrutable to humans. Hot on the heels of proteomics, deep learning is poised to transform many fields of science and mathematics.

The workshop will bring together leaders in deep learning with top scientists and mathematicians in their respective domains to understand advances in the physical sciences and mathematics driven by deep learning, such as protein folding, chemistry, quantum physics, geometry, and theorem proving. We will examine methodologies for embedding physical and mathematical intuitions into neural network architectures and algorithms: What are the best ways to build inductive biases into networks? Finally, we'll examine the potential of large language models to help us summarize, synthesize and analyze the scientific literature itself.

</ABSTRACT>

<INFO>
One difficulty is to present the case in full with only a minimum of arrogance.
For question 20, what about something along the following lines:

Over the 35+ years of the workshop, the scientific output has included the following:
- two books based on the meeting lectures (in 1993 and 2012),
- 14 sets of lecture notes that have been used in graduate courses in top-level universities around the world

- several dozens of scientific papers issued from collaborations that began at Bellairs, including a summary of the 2019 workshop published in Nature Neuroscience.

Attendees over the years have included top-level researchers from McGill, Canada and the rest of the world.
We have hosted one Fields medalist, several Turing Award winners, and awardees of other top prizes in mathematics and computer science, such as Abel Prize, Gödel prize, Nevanlinna Prize.
In May 2024, among the people who had accepted our invitation for the 2025 workshop were Avi Wigderson, who won the Turing Award last September, and John Jumper, who won the Nobel Prize in Chemistry last October. Unfortunately, the dates that were offered to us by Bellairs were conflicting with the main AI conference of the year, which meant we had to cancel our meeting. We fully expect that our 2026 workshop will be attended by people of the same quality as we typically have.
This of course represents very unique occasions for the many graduate students in attendance to interact with the very best researchers in the world in the very special context offered by Bellairs.

In recent years, an effort has been made to connect with the local community.
In 2023, the Minister of Science met at Bellairs with five workshop attendees, global leaders of the AI community,
and in 2024, the chief of staff of the Prime Minister participated in a workshop session.
We intend to repeat with members of the government and of the UWI community: in particular Professor Simon Anderson has already accepted our invitation.

Our 2024 workshop was attended by a journalist of La Presse who spend the week with us.
He then wrote a series of 5 articles in his newspaper presenting Bellairs as a very unique facility and describing in lay terms the scientific content of the workshop.
</INFO>

Brainstorming

Brainstorming

What does science look like now?

Is science about prediction?
Is science about control?
How is it different from engineering?
Experimental versus observational science.
Can we agree on this?

Review what is happening for AlphaFold
We have equations that we believe would/could predict the folds for proteins. If we could run them for high enough fidelity.

Quantum Field Theory (QFT): This is the gold standard for describing how fundamental particles (like electrons and quarks) interact with fields. It accounts for things like particle creation/annihilation and relativistic effects. The standard model is based on this.

The Schrödinger Equation: This is "standard" quantum mechanics. It describes how electrons move around nuclei without worrying about the high-energy "field" effects of QFT.

Molecular Dynamics (Classical): This treats atoms like balls on springs. This is how 99% of protein folding is simulated today.

The Computational Wall: Even with "much faster" computers, QFT is an exponential nightmare for protein-sized systems.

Consider a Small Protein: ~1,000 atoms.
To solve this with QFT, you aren't just tracking 1,000 points; you are tracking the values of multiple fields (electromagnetic, electron, etc.) at every single point in space occupied by the protein.

The Data: The number of degrees of freedom would be billions of times higher than standard simulations. **Even a computer a trillion times faster than today's supercomputers would struggle to simulate a single microsecond of folding using pure QFT.**

So what do we do instead? → AlphaFold

https://en.wikipedia.org/wiki/Standard_Model
What will science artifacts look like in the future? We're still indexed on papers and books.
We've barely started accepting Jupyter notebooks, videos, and blogs.
What about a model that you can speak to about a specific topic?

What will science look like in the future?
Simple/easy decisions are becoming possible.

Focusing the best humans on the key questions, distilling these into the best current model so that decisions are available in every lab.

https://www.worldgov.org/complexity.html

Notes for talk:

Notes

Mar 15, 2026

AI for Science:

Spectrum from formal → informal → social values

If we start to have models that we trust more than all the experts on the quantitative and the hard problems, this will bleed into the social/policy space.
How do we make sure to keep these separated out as much as possible?

2 eras:

1. Digestion of existing human knowledge.
2. Agents make up their own mind through experimentation.

Reform versus revolution
Things are working very well at the moment. Just difficult to see.

Notes for first lecture:

How long do we want it to be?: 1 hour?
Who wants to give it?:
What do we want to say? / Offer a framework?

Ideal output would be a deliverable:
E.g. Data Science & AI
=> Position paper contribution for special issue.
=> Could invite counter positions.

Intermediate goal:

- Interesting and compelling interactions
- Find a bridge between the areas of science linked by AI

Provocations / Questions:

1. Where will paradigms of science be stored? Books → Simulations/Computers → AI/LLMs
2. Is the introduction of AI a paradigm shift for science? Or, just business as usual?
3. What role does human understanding play in science?
    a. E.g. we don't understand how AlphaFold works.
    b. A biochemist may still understand how to use the tool.
    c. An ML person can understand how the tool is created.

4. To what extent is the future of science an “omni model” versus hyperspecific models for specific scientific disciplines?
5. What will happen to human agency in science? e.g. hypothesis creation and exploration
    a. Does it matter if humans aren’t involved in some aspects?
    b. If PhDs / trainees don’t get involved in the early bits, what happens in the long-run?
6. Is there a standard “recipe / playbook” for applying deep learning to fields of science?
    a. Build architectures inspired by local knowledge
    b. Learnability of architectures
    c. What is common across the fields?
    d. We have a accurate but slow simulator agreed upon for the field
        i. Use this to generate data
        ii. Build an ML model (amortized)
        iii. Key question:When are you sure of generalization? When do you need to backout to the original simulator?
        iv. Differentiating through ML simulators
7. ML could bring commonality => focus on prediction, inference, generalization across domains.
8. Is “one model to rule them all” even plausible?
    a. Proteins and Languages
9. Can AI grapple with open-endedness? Now? In future?
10. What are key AI/ML tools (that we have already)?
11. What are missing AI/ML tools that we need to make progress in science?

Talk structure:

1. What science has been / philosophy of science
2. What’s happening now?: Categorization of ways people are employing ML in science
3. Provocations / questions

Break out to groups to discuss a bit?