Welcome to Day 3 of the AI+MPS Workshop!

March 24–26, 2025

Today’s Schedule (Wednesday)

9:00–10:15 am: Reports by Domain

- Astronomical Sciences
- Chemistry
- Materials Science
- Mathematical Sciences
- Physics

10:15–10:45 am: Break

10:45–11:45 am: Open Discussion

11:45 am–12:00 pm: Closeout

12:00–1:00 pm: *Lunch is served (optional)*

Breakout Report: Astronomical Sciences

10 minutes presentation + 5 minutes Q&A

Yuan-Sen Ting, Stella Offner, Uros Seljak


NETFLIX
FINAL
TRAILER

AI for Astronomy

- **Scientific Return from Surveys:** Efficiently realize the science returns of **multi-billion dollar $$$$ investments** in new experiments and surveys

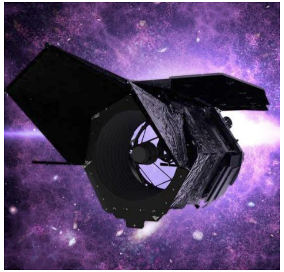

- AST is at **a moment of transition**, after a **community wide effort** to build interdisciplinary institutes - let's keep the momentum!

Successful Examples

Operation

- **Scalability of data processing:** working with large, complex data, e.g., cosmic ray rejection, efficient sub-sampling, amortized neural posteriors, **significant cost savings**
- **Astro-instrumentation**: **workflow optimization** - reinforcement learning (e.g., PFS follow-up)
- **Needle in the haystack:** Anomaly detection with generative models (e.g. gravitational lens)

Inference

- **Enabling parameter inference: Simulation Based Inference**, better cosmology constraint (SIMBIG 1Gpc^3 study, reduction of error bars, **exoplanet architectures**)
- **Emulators and surrogate models:** speed up simulations (N-body, cooling tables, image augmentations)
- Synthesis and analysis of complex signals: **new physical insights** (e.g. Microlensing)

Astronomy-Specific Priorities and Challenges

Priority : Reinforcing Our Successes

- **Inference with Imperfect Simulations:** Create reliable **uncertainty and bias quantification** and validation tests from AI models
- Build AST-specific architectures: Develop **AST specific AI models** (physics informed, **physics symmetries**, scalable to large data, **multi-scale information**)
- **Create data challenges** with simulations and/or real data: benchmarks both supervised (e.g. parameter uncertainty) and unsupervised (**unknown systematics**)

There might be more ~

- **Need clear transformative application** to show success, e.g., **true examples of discoveries** made with machine learning and AI
- Expedite AI-assisted discovery in and **analysis of the literature (STEAR)**

General Challenges

General challenges

- **Trust of AI** results within the community (e.g., review of proposals, hiring), very bimodal perception
- **Establish benchmarks — mitigate distrust** — especially on LLM reasoning ability.
- **Lagging education in AI literacy**, methods (see Education Theme)
- Lack of adoption of **responsible AI best-practices**, blind data analysis with AI

Tackling the general challenges

- AI expertise is dispersed: **Increase mobility** for students and postdoc to leverage on AI knowledge of these experts.
- Incentives for building **interdisciplinary collaborations between institutions**
- **Funding for internships**, exchanges for students, postdocs, and researchers
- Access to AI-resources (e.g., computing, data sharing, storage)

Astronomy for AI (Interdisciplinary Opportunities)

- Access to **open complex data sets** with a range of data volumes containing physics relations that have a **large dynamic range** and where there is **significant domain shifts** between the test and training data — **non-proprietary ‘safe’ data**.

- **Well curated literature with long time baseline** - with data and codes.

- **Understanding the Science of AI** (what information is encoded within a network and how we use that in getting new insights / building new architectures). Physicists are trained to **extract insights from complex phenomena** (e.g. neural scaling laws)

AI+AST Priorities (unedited)

- Efficiently realize the science returns of multi-billion dollar investments in new experiments and surveys
- Need to invest in the scientific inference on collected data
- Accelerate computationally intensive simulations and analyses
- Build AST-specific architectures: To extract maximal information we need development of AST specific AI models (physics symmetries, scalable to large 1d, 2d, and 3d data, multi-scale information)
- Create reliable uncertainty and bias quantification and validation tests from AI models
- Establish benchmarks — mitigate distrust — especially on LLM applications.
- Create data challenges with simulations and/or real data: compare against non AI baselines, benchmarks both supervised (e.g. parameter uncertainty) and unsupervised (search for unknown systematics)
- Establish funding mechanisms for well designed and open data challenges, validation tests
- Expedite AI-assisted discovery in and analysis of the literature (STEAR)
- Improvements are possible through AI that would lead to significant cost savings
- More funding for interdisciplinary research

Flagship / Successful Examples (unedited)

- Parameter inference: Simulation Based Inference on large scale structure (with validity of the simulations as key bottleneck) - higher order information - better cosmology constraint (SIMBIG 1Gpc^3 study, reduction of error bars, exoplanet architectures (predicting # planets /star –see Eric Ford))
- Scalability of data processing: working with large, complex data, e.g., cosmic ray rejection, efficient sub-sampling, amortized neural posteriors
- Needle in the haystack: Anomaly detection with generative models and unsupervised ML, gravitational lens detection
- Astro-instrumentation: workflow optimization - reinforcement learning success = cost reduction (e.g., PFS follow-up)
- Significant Speed: Speed up simulations (e.g, N-body, cooling tables, image simulations)
- Physical Inspired Models (Generative Models and beyond) – incentivise AI development along this area, learning physical SN parameters empirically
- Emulators and surrogate models
- Synthesis and analysis of complex signals: extract summary statistics, model exoplanet spectra, galaxy morphology, new physical insights from summarization of simulations in microlensing
- Multimodality in astronomy = opportunity in AI - e.g. spectra. Constrained complexity (simplicity of mapping between modalities) makes learning relations tractable
- Foundation models for image analysis (seems more like an opportunity)

Astronomy-Specific Challenges (unedited)

- Target areas in which challenges exist:
- Need to develop realistic simulations that include all of the real data complications (e.g. survey mask, point spread function, noise, astrophysical sources, including instrument) and have enough model flexibility to match data. Simulations are expensive and have a hierarchy of fidelities where the implications of adopted simplifications are often unknown.
- Generalizable models that extend to other areas (different subfields, different data), out of distribution data, hard to apply models trained on simulations to real data
- Need workflows and applications that work end-to-end (on large datasets): from input models to output predictions on real data; takes time and personnel.
- Astronomy is observational science (except lab astrophysics) of one complex experiment (the universe), we cannot do adaptive experiments, so conventional experimental design strategies do not apply, eg, unlike other MPS sciences, closed-loop robotic laboratory experiments are not generally possible
- General challenges
- Not enough mobility for students and postdoc to leverage on AI knowledge of some intsitutes. More exposure will accelerate the change of paradigm.
- Lack of funding for internships, exchanges for students, postdocs, and researchers to learn AI
- Trust of AI results within the community (e.g., review of proposals, hiring), very bimodal perception,
- Lagging education in AI literacy, methods (see Education Theme)
- Black-box methods, lack of adoption of responsible AI best-practices
- Blind data analysis with AI
- Efficient discovery of new methodologies and associated applications
- Inequities in access to AI-resources (e.g., computing, data sharing, storage)
- Uncertainty quantification, model specification
- Loss of talent to industry
- Need clear transformative application to show success, e.g., true examples of discoveries made with machine learning and AI

Astronomy Interdisciplinary Opportunities (unedited)

Access to open complex data sets with a range of data volumes containing physics relations that have a large dynamic range and where there is significant domain shifts between the test and training data, non-proprietary ‘safe’ data, collaborative

- Incorporating AI and domain knowledge (including uncertainty quantification, physics constraints, and laws) - e.g. building architectures based on physics relations, understanding the deviation from physics constraints when a given model “works”.
- Efficient use of data for discovery and inference. Improving forward models and simulations.
- Experimental design approaches for new experiments/instruments, target selection, and data discovery workflows utilizing existing infrastructures.
- Understanding the Science of AI (what is information is encoded within a network and how do we use that in getting new insights / building new architectures). Physicists are trained to extract insights from complex phenomenon (e.g. neural scaling laws)
- Developing multimodal analyses that incorporate literature, data, domain knowledge for hypothesis generation.
- AST is at a moment of transition, after a community wide effort to build interdisciplinary institutes - let’s keep the momentum!
- Well curated literature with long time baseline - with data and codes.
- Incentives for building interdisciplinary collaborations. Maybe more smaller institutes with a lot of focus on making links between big and small institutes and individual interdisciplinary researchers.
- How to relax generative models trained on simulations that we know are incomplete to cover a broader range of predictions that still satisfy all physics constraints. Models that are physical but also flexible enough to discover new physics (e. g. Automatic symmetry discovery)
- AI for instrumentation: RL, signal processing, experiment design
- Develop unsupervised anomaly detection techniques that are AST specific
- Unique opportunities for AI/ML researchers: large and complex open data, open models and software, physics-based models, broad public interest, well-defined literature, non-proprietary (‘safe’) data, where a low dimensional underlying physical model may be discoverable.

Breakout Report: Chemistry

10 minutes presentation + 5 minutes Q&A

Pratyush Tiwary, Brett Savoie, Shuwen Yue

AI+CHE Priorities

- **Data:** We need to find more ways to incentivize and professionalize data generation efforts. Funding mechanisms; industrial collaboration; broader recognitions for those that dedicate their groups to data generation.
- **Benchmarking:** Broad agreement that benchmarking is important and that we need more of it. Big competitions like CASP have been historically important. The downsides of gaming benchmarks need to be considered as well.
- **Research Culture:** AI is in the process of displacing interpretable methods with consequences that are still being revealed. Empirical methods development and physics-based development are often competing for the same funds and journal space. Can you can have discovery without interpretability?
- **Training/education:** More investments in training activities at all levels are a priority. In person workshops, training workshops for instructors at all levels, regional training specific to certain industries, and asynchronous training for the current workforce. The pace of change seems significant enough that there will be serious efforts to retrain/upskill the current workforce rather than just focus on the next.
- **The CHE/CS Interface:** Prior successful examples have mainly occurred through long-term center-scale awards. New mechanisms (hiring, training, funding structures) are required to better populate this interface.

Flagship Examples

- **AlphaFold**
  - The ChatGPT moment for a lot of biochemistry.
  - Accurate but non-interpretable methods are clearly useful as part of the knowledge generation.
  - Highlights the tension between data generators and model developers. No one from PDB won a Nobel.
  - Highlights the tension between data intensity for ML and data availability for most chemical applications.
- **Machine Learning Potentials:**
  - Orders of magnitude changes in accuracy, cost, and ease of application have changed the questions that can be asked.
  - MLPs also lay bare many of the tensions within the field: generality/narrow models; interpretability/predictability.
- **Generative models:**
  - Chemical structures and materials produced by generative models are increasingly being used as inputs to design loops.
  - Example of the centaur approach: inputs are AI generated, but the downstream validation process is often traditional.
- **CASP (Critical Assessment of protein Structure Prediction)**
  - A biennial structure prediction competition that was instrumental to the development of AlphaFold.
  - “Gold Standard” for a chemistry competition.
  - Benchmarks can be highly useful, but eventually they become saturated, so investments in new benchmarks are needed.
- **Shallow Methods**
  - Expert descriptors with relatively simple models have found many practical applications across chemistry.
  - Gaussian process regression as part of active learning optimization has been transformative in reaction condition optimization.
  - Getting shallow ML into the hands of practitioners in user friendly modes will be important.

Subdomain Considerations

- **Data**
  - Chemical data is highly contextual
  - Intrinsic scarcity for many properties
  - Uncertainty quantification is very important (experimental/simulation errors are expensive to assess).
  - Unlocking historical data seems like a solvable problem for chemistry with broad benefits.
- **Benchmarks**
  - Norms for model evaluation vary dramatically and are inconsistently applied.
  - How to improve efficiency and reproducibility of simulation algorithms?
  - Extrapolation vs interpolation
- **Culture**
  - How do we make sure tools are accessible and not being misused?
  - Do we need to teach more linear algebra or PDEs?
  - How do we incentivize data generation?
  - AI differs from simulations in the types of tasks it is being used for (e.g., idea generation and decision making versus data generation).
- **Education/Training**
  - How does AI differ from simulations.
  - Asynchronous training poses unique challenges for lab-based disciplines.
- **Gaps/Methods at the AI Interface**
  - How do we integrate physical constraints?
  - Should we integrate physical constraints?
  - How do we assess when a method is good enough?

Interdisciplinary Opportunities

- **Methods development**
  - CS is the home of a lot of methods/architecture development, but penetration to chemical problems requires a lot of work.
  - Foundation model concepts require cross-disciplinary integration.
- **Develop new benchmarks/joint-benchmarks**
  - We heard from CS colleagues that Chemistry doesn't care about their benchmarks. How do we cultivate shared interests and vocabulary at the AI/chemistry interface?
- **Training Opportunities**
  - Education/CS/Chem collaborations will be critical for effective instruction at the AI/chemistry interface.
  - Training programs that encourage jointly-advised students and cross-disciplinary training.
  - A lot more small regional conferences that target the AI/Chemistry interface and involve students, instructors, faculty, at all ranks and types of institutions.
- **Open-source all federally funded fundamental research content.**

Final Thoughts

- **Distinguishing signal from noise**
  - Enough breakthroughs have already taken place to convince our community of real opportunities.
  - Naive applications of AI are inevitable and contribute to cynicism.
  - Some things don't work until they suddenly do. Resource allocation will be challenging for those writing proposals at the AI/chemistry interface.
- **Automation**
  - Robotics (lots of opportunities, not directly MPS but impacts it)
  - Experimental design (used to be exclusively human, now human in the loop, next?)
  - Analysis (automating the interpretation of results, distillation to knowledge, hypothesis generation)

Breakout Report: Materials Research

10 minutes presentation + 5 minutes Q&A

Andrew Ferguson, Keith Brown, Chrisy Xiyu Du

Materials Research

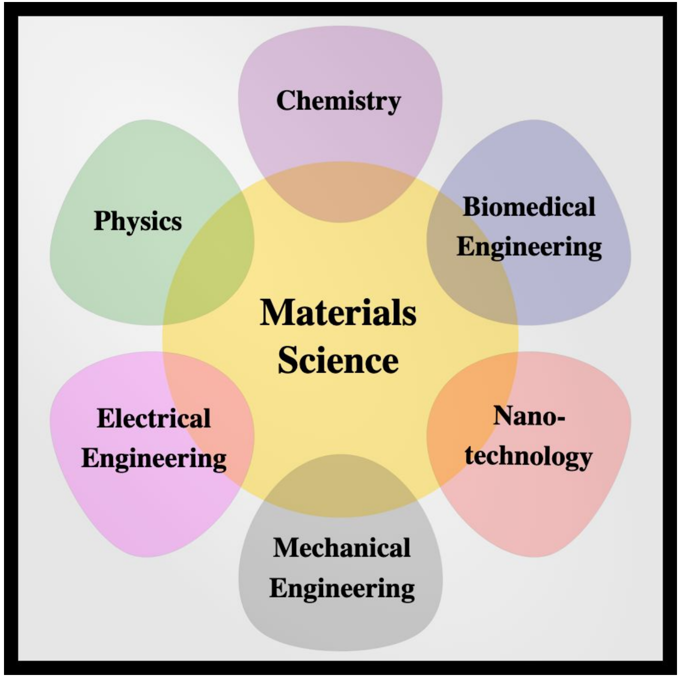

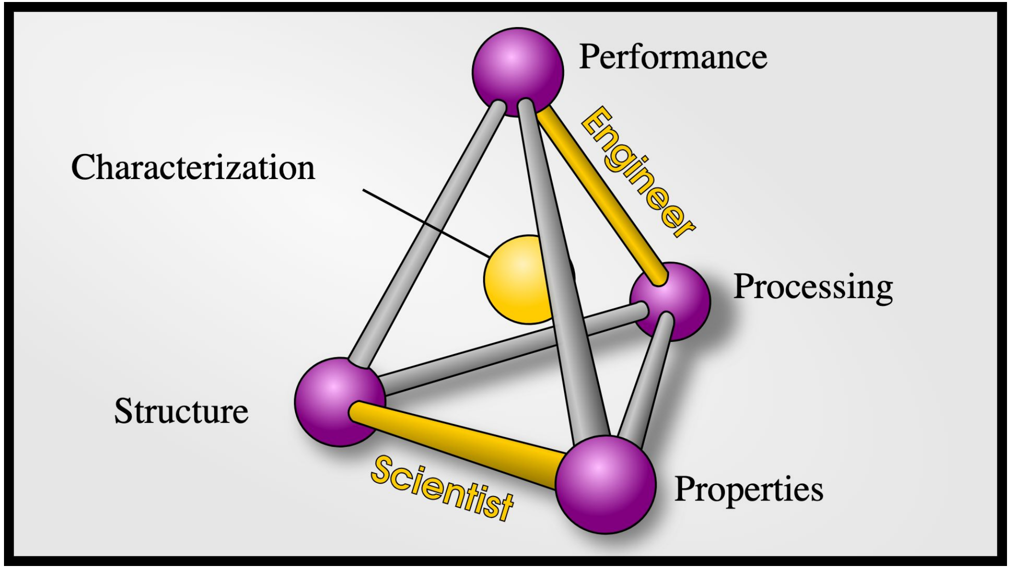

https://msestudent.com/what-is-materials-science-and-engineering-the-definitive-explanation/

AI+DMR Priorities

- Research that addresses small data
- Recognition of the heterogeneity in materials research
  - Hard materials, soft materials
  - Ordered, disordered, liquid systems
- Uncertainty quantification (especially systematic errors)
  - Same measurements from different equipments
  - Same simulation from different software packages
- Develop courses on domain + AI to increase students' interest and participation
- Identifying problems that are interesting to computer scientists
- Multiscale modeling

Flagship Examples

- Training ML potentials from DFT data
- Deep learning for photonic structures design
- Organic molecule design using ML
- Molecular discovery with genAI
- “chatGPT” for protein design with wet lab validations
- Utilizing data augmentation for small data materials discoveries
- Molecular dynamics with automatic differentiation
- Self-driving labs for materials mechanical properties

Subdomain Considerations

- Materials researchers come from all over the place (chemistry, chemical engineering, materials science, mechanical engineering…)
- We need guardrails (obeying physical or chemical laws) in our AI models
- Some big-data materials science works are not necessarily useful
- Data is very heterogeneous, cannot have a one-size-fits-all database
- Data is small, sparse and often wrong, this applies to metadata as well
- We need institutes like NSF to host datasets. Multiple valuable datasets recently became closed source
- Most interesting materials are messy, heterogeneous, disordered, and out-of-equilibrium

Interdisciplinary Opportunities

- Better data publishing practices
  - Metadata inclusion for both experiments and simulations
- Workshops to discuss failures and nuances of methods
- Middle ground between experimental and simulation lengthscales
- AI for materials and materials for AI
- Collaborate with applied math and statisticians to solve problems arising from small data
- Create defined projects that might be good for CS undergraduate capstone
- Materials science democratization
  - Build lower entry barrier tools for materials design
  - Open source software with good tutorials students can just download and use

Final Thoughts

- Some resources mentioned in our discussion
  - Chemical synthesis and materials discovery
  - Telluride Conferences
- Soft materials is hard, hard materials is also hard

Breakout Report: Mathematical Sciences

10 minutes presentation + 5 minutes Q&A

Soledad Villar, Melanie Weber, Zhuoran Yang

Overview

- AI is built on mathematics and statistics
  - fundamental research is key

- Mathematics and statistics will revolutionize AI by doing for AI
  - what geometry did for physics
  - what Maxwell's equations did for electronics / communication

- AI will revolutionize mathematics and statistics by doing for math and stats
  - what quantum mechanics did 100 years ago
  - what computers did 50 years ago

Moral: M + AI = O-mai

Flagship Examples: AI is built on math and stats

AI success relies on fundamental techniques and principles from mathematical sciences:

- Training: numerical linear algebra, 1st order optimization methods
- Generative AI: games, flows/optimal transport, and diffusions
- HP tuning/transfer: random matrix theory, interacting particles and mean field dynamics, PDEs, SDEs
- AI for Science: sampling, inference, hypothesis testing, estimation, UQ
- Architecture design: algebra, topology, geometry, physical priors

Moral: Trust the mathematical aesthetic

Subdomain Considerations: Math /Stat will revolutionize AI

- Model design – encoding physics/symmetries and ensuring trainability
- Math/stats principles more relevant with scale:
    - Model selection – What can we learn?
    - Model training – How to get the most out of your model?
    - Model deployment – How to quantize/prune?
    - Model assessment - How to ensure trust-worthy estimates
- Fundamental limits – guarantees, principled algorithms: privacy, robustness, UQ, and fairness
- Tools: random matrix theory, interacting particle systems, dynamical systems, statistical learning theory, probabilistic modeling, causal inference, conformal prediction, implicit likelihood inference (SBI)

Moral: Math/Stats : AI :: Transistor : Computer

Subdomain Considerations: AI will revolutionize Math/Stats

- Human / AI mathemacentaur (statistentaur?):
  - Democratizing research
  - Accelerating research
- AI will do for math what quantum mechanics did for math 100 years ago
- Non-convex optimization
- High-dimensional PDEs, SDEs, probabilities
- Data-driven modeling for scientific computing
- Optimizer and data-aware statistical learning
- Mathematical biology (protein structure prediction, systems biology modeling, synthetic biology)

MATH/STATS as the ultimate Turing Test?

Moral: AI will do for math what quantum mechanics did 100 years ago

AI+DMS Priorities

- Increasing funding at interface of Math / Stats / AI:
  - Person-centric over project-specific funding: a la Simons Investigators (e.g. NIH MIRA, NSFGRP, NSF Postdoc)
  - Increase funding for Math institutes:
    - invest in human and intellectual capital
    - provide space for interdisciplinary collaborations
  - Fund interdisciplinary students, postdocs, sabbaticals: consider more DMS+x cofunding
  - Summer schools, curricular materials for teaching mathematicians about AI
  - \epsilon compute = 1/\epsilon opportunity
  - Increase recognition at interface of Math / Stats / AI: new prizes/awards

- Enlarge DMS research priorities:
  - AI assisted theorem-proving: encourage formal verification, principled gen-AI/LLM use
  - help math / stats researchers coalesce around new fundamental problems in Math ∩ AI
  - Emphasize / accelerate integration with AI for Science: custom models and algorithms

Moral: the unreasonable effectiveness of math for AI is only more important as AI accelerates

Interdisciplinary Opportunities

- Emphasize / accelerate integration with AI for Science:
  - Custom models adapted to scientific domains
  - Custom algorithms for training these models
  - Joint funding opportunities
- Reward statistical work in AI for science and beyond
  - Data processing, reliable UQ with imperfect data/models, hypothesis testing, causal inference, anomaly detection
  - Statistical methods in AI for the social sciences
- Formal guarantees in AI models for the sciences

Moral: Again, the unreasonable effectiveness of math for the natural sciences still holds in the AI times

Final Thoughts

- AI is built on mathematics and statistics
  - fundamental research is key

- Mathematics and statistics will revolutionize AI by doing for AI
  - what geometry did for physics
  - what Maxwell's equations did for electronics / communication

- AI will revolutionize mathematics and statistics by doing for math and stats
  - what quantum mechanics did 100 years ago
  - what computers did 50 years ago

Moral: M + AI = O-mai

[EXTENDED] Flagship Examples: AI is built on math and stats

AI success relies on fundamental techniques and principles from mathematical sciences:

- Training: Mathematical optimization community long ago developed 1st order methods; adaptive optimization, momentum based methods. These are the workhorses of modern AI optimizers.
- Generative AI: GANs originally formulated as two player games, objective relied on Wasserstein distance; Diffusion models are literally normalizing flows and are build in the language of measure transport; Mean field interacting particle systems are THE language that has emerged for describing training dynamics of large neural networks. This has big applications to model scaling and hyperparameter tuning.
- Initialization: How to initialize and parameterize a model so that is it numerically stable enough to train? Formalizing these notions and getting answers that are now standard in all NN libraries used tools from random matrix theory (free probability and multiplicative ergodic theory), ideas in markov chain mixing times, insights from Gaussian processes, ODEs/SDEs as continuum limits of infinite deep MLPs and ResNets, gaussian processes as limits of neural networks at infinite width, kernel methods/RKHS techniques as a way to understand core phenomena discovered in neural networks e.g. benign overfitting, double descent. These phenomena continue to hold also for linear models.
- AI for Science: sampling, inference, hypothesis testing, estimation, UQ,
- Architecture design: spectral graph theory and graph algorithm to stability training of deep graph networks, imposing physical priors through architecture via invariant theory and other equivariant neural networks. Harmonic analysis as a tool to describe what a given architecture can learn, PDEs in spaces of probability measures as continuum limits of neural networks. Analyzing Bayesian inference with neural networks as a method to understand their “thermodynamic variables”

[EXTENDED] Subdomain Considerations: Math /Stat will revolutionize AI

- Model design – encoding physics/symmetries and ensuring trainability, neural architecture search (discovery of new architectures) requires
- Math/stats <u>more relevant with scale</u>:
  - Model selection – What can we learn? In what order does learning happen?
  - Model training – How to get the most out of your model? Predicting and optimizing scaling laws. Understanding effect of practical choices: learning rates, batch sizes, momentum coefficients, what is the implicit regularization of algorithms?
  - Model deployment – How to quantize/prune? How to make models smaller? What are the fundamental limits?
- Fundamental limits – guarantees, principled algorithms: privacy, robustness, UQ, and fairness, are adversarial examples unavoidable? Just how robust and in what sense can learning in high dimensions be? In what sense is learning universal?
- Tools: random matrix theory, interacting particle systems, dynamical systems, UQ, causal inference, spin glasses, approximate sampling, non-convex optimization,

Moral: DMS : AI :: Transistor : Computer

[EXTENDED] Subdomain Considerations: AI will revolutionize Math/Stats

- Human / AI mathemacentaur: democratizing research, accelerating research, proving the millennium problems?, running REUs on demand
- e.g. will drive research in: non-convex optimization, computationally/algorithmically constrained statistical learning, statistical to computational gaps, piecewise linear morse theory, non-linear functional analysis (spaces of Hilbert Spaces), manifold learning, interacting particle systems, random matrix theory, universality,
- Data-driven modeling for scientific computing
- Optimizer and data-aware statistical learning

MATH/STATS as the ultimate Turing Test?

Moral: AI will do for math what quantum mechanics did for math 100 years ago

[EXTENDED] Interdisciplinary Opportunities

- Emphasize / accelerate integration with AI for Science:
    - Custom models adapted to scientific domains: Physical priors, symmetries, conservation laws, incorporated into ML via representation theory, invariant theory, graph neural networks, modeling of multiscale physics. Define AI models for simulation-based inference with correct physical structure.
    - Custom algorithms for training these models: How do we use AI to produce simulations for simulation based-inference? How do we effectively train equivariant models?
    - Joint funding opportunities. Shared graduate students and postdocs. Encourage collaborations across disciplines through funding.
- Reward statistical work in AI for science and beyond
    - Data processing, UQ, hypothesis testing, causal inference, anomaly detection.
    - Statistical methods in AI for the social sciences: conformal prediction.
- Formal guarantees in AI models for the sciences. Stability and robustness guarantees, correctness guarantees. Post training data-driven guarantees.

Math and Statistics will revolutionize AI

- Model design – encoding physics/symmetries and ensuring trainability [NAS]
- Math/stats more relevant with scale:
    - Model selection – What can we learn?
    - Model training – How to get the most out of your model?
    - Model deployment – How to quantize/prune?
- Fundamental limits – guarantees, principled algorithms: privacy, robustness, UQ, and fairness
- Tools: random matrix theory, interacting particle systems, dynamical systems

Moral: DMS : AI :: Transistor : Computer

AI will revolutionize Math/Stats

- Human / AI mathemacentaur:
  - Democratizing research
  - Accelerating research
- AI will do for math what quantum mechanics did for math 100 years ago
- Non-convex optimization
- Data-driven modeling for scientific computing
- Optimizer and data-aware statistical learning

MATH/STATS as the ultimate Turing Test?

Moral: It’s a good day to revolutionize, Mr. Bond :-)

AI+DMS Priorities

- Increasing funding at interface of Math / Stats / AI:
  - Person-specific instead of project-specific funding: DMS version of Simons Investigators, funding for interdisciplinary students, postdocs, sabbaticals (e.g. NIH MIRA, NSFGRP, NSF Postdoc)
  - Increased support for graduate students and postdocs
  - Summer schools, curricular materials for teaching mathematicians about AI
  - \epsilon compute = 1/\epsilon opportunity
  - Increase recognition at interface of Math / Stats / AI: new prizes/awards
- Enlarge DMS research priorities:
  - AI assisted theorem-proving is extremely promising: encourage formal verification, principled gen-AI/LLM use
  - help math / stats researchers coalesce around new fundamental problems in Math ∩ AI
  - Emphasize / accelerate integration with AI for Science:
    - Custom models adapted to scientific domains
    - Custom algorithms for training these models

Moral: DMS is the home of timeless skills and training: only more important as AI accelerates

Math/Stats Problems from AI

New problems from AI: (analogy AI will do for math what quantum mechanics did 100 years ago). This will require work in all areas of math and statistics funded by DMS

1. Mathematical and statistical foundations for learning/search over nonlinear search spaces e.g.
   a. Explaining overparameterization, non-convex optimization, why do local optimization methods work so well? What do they learn? How to even talk about what they learn?
   b. Data and optimization-aware statistical learning theory: tight generalization guarantees?
   c. Emergent algebraic / topological structure on learned representations?
2. Which data distributions are learnable and by which architectures + algorithms? e.g.
   a. What exactly distinguishes between those distributions that we can learn by diffusion models and those that we probably cannot? Do we not violate no-go / free lunch theorems
3. Mathematical and statistical foundations of trustworthy AI (e.g. robustness, UQ, privacy, fairness)
   a. Privacy and fairness for LLMs
4. Optimal distributed algorithms at massive scale: statistical and computational optimality under communication (e.g. distributed optimization, siloed data, heterogeneous computational resources)

AI is changing the mathematical sciences research

The priorities in different fields changed based on techniques needed to improve and understand AI.

- How we think of optimization has significantly changed based on the current practice of training AI models.
- Data-driven modeling is moving from empirical equations to neural network architectures (e.g., Neural ODEs). This has changed how we do scientific computing.
- AI practice has expanded perspectives on classical statistical learning concepts (e.g., overparameterization).
- AI practice is expanding the applicability of classical statistical learning to large complex systems with heterogeneous data and low signal-to-noise settings.

Examples for “Hilbert List”

AI will revolutionize Mathematics and Statistics through new problems and accelerating of research.

AI for mathematical reasoning:

1. The automated grad student: take a random GSM textbook; correctly solve 99% of the exercises.
2. Human / AI centaur mathematician: (democratization of math)
    a. Formally verifiable proof of any millennium prize problem
    b. Conceptual proof of the classification of finite simple groups
3. Needle in the haystack: algorithms for finding “rare cases”, counter examples etc (Spurred by ease of data generation).
4. 10^6 Accelerating of formal verification: resolve AI integration bottleneck

Examples for “Hilbert List”

AI will revolutionize Mathematics and Statistics through new problems and accelerating of research.

AI for mathematical reasoning:

1. The automated grad student: take a random GSM textbook; correctly solve 99% of the exercises.
2. Human / AI centaur mathematician: (democratization of math)
   a. Formally verifiable proof of any millennium prize problem
   b. Conceptual proof of the classification of finite simple groups
3. Needle in the haystack: algorithms for finding “rare cases”, counter examples etc (Spurred by ease of data generation).
4. 10^6 Accelerating of formal verification: resolve AI integration bottleneclk

The Promise of Math for AI

AI will revolutionize Mathematics. Let's spur fundamental revolution in deep new mathematics:

1. AI is (yet another) example of how deep mathematics give the language for unexpectedly and profoundly useful applications
   a. Optimization – acceleration, momentum, scheduling, etc
   b. Diffusion models / flows – langevin
2. Mathematical methods for understanding robustness and
   a. Fundamental questions: adversarial example
   b. Privacy
   c. UQ in ML/ AI.
1. Leveraging mathematical tools to design more efficient models. Even small reduction in complexity could lead to very large savings:
   a. Fundamental principles for efficient search over spaces of architectures, optimizers, etc
   b. Hyperparameters transfer – scaling limits of neural networks (PDEs, random matrix theory, dynamical systems, mean field particle systems)
   c. Architecture design – equivariant architectures

Subdomain Considerations

- All subdomains can benefit from AI/ contribute to AI.
- Better than compartmentalization, combine ideas from different areas of Mathematics and Statistics to study Deep Learning. Successful research at intersection of Math and AI needs to cut across DMS subdisciplines?
- Foc on problems themselves rather than areas. This will better promote collaboration among areas. Examples, SAMSI, collaboration for postdocs coadvised by people with diverse backgrounds.

AI+DMS Priorities

- Create funding mechanisms for DMS PIs/postdoc/students to spend time in a lab from other MPS disciplines/CS to start interdisciplinary collaborations (mimicking NIH MIRA model).
- (Modest amount of compute resources largely sufficient for Math for AI/AI for Math, especially for prototyping the ideas…)
- DMS is the “glue” between fields: DMS could encourage AI-facilitated collaboration with the other sciences by developing cross-cutting collaboration programs using AI as “translators”
- Lens of statistical limits has been very fruitful in ML research. Lens of abstraction less utilized, but lots of potential: e.g., how to leverage compositionality in principled study of ML/AI.
- AI will influence every areas of math. DMS could initiate programs for using AI+PI collaboration to generate novel questions in Math ∩ AI.
- DMS could help build a community across mathematical disciplines to make a list of the most important problems in Math ∩ AI (cf. Hilbert)
- AI and theorem-proving is extremely promising: DMS can encourage development in everything from formal verification to principled gen-AI/LLM use.
- Develop mathematical tools to explain AI models and extract scientific explanations
- Augment AI with DMS to discover scientific models from data with explanation
- Summer schools for teaching mathematicians to use AI

Interdisciplinary Opportunities

- Various areas of DMS can work together and tackle some of the fundamental problems of AI understanding.
- Sustained collaborations between MPS domains. NSF AI Institutes with stronger focus on DMS disciplines.
- AI has rich structures that will lead to new mathematical understandings and tools. Math/stats can be used as a common language for various areas (e.g., physics and chemistry) to understand AI and to build trustworthy AI tools with theoretical guarantees.
- Math can come up with evaluation metrics for comparing AI objectives / methodologies. Math can provide tools for transparency/ trustworthiness.
- Various areas essentially care about similar problems but with different languages. Collaboration will be promising.

Interdisciplinary Opportunities

- Various areas within Mathematical Sciences (DMS) – including statistics, probability theory, and measure theory - can collaborate to tackle fundamental problems in AI understanding, interpretability, robustness, and verification.
- Mathematical concepts bridge terminological gaps between disciplines, transforming domain-specific AI problems in physics, chemistry, and other sciences into structured frameworks with common analytical approaches.
- Mathematical frameworks provide essential tools for AI evaluation, including rigorous uncertainty quantification, formal verification methods, and trustworthy statistical inference across the entire data science lifecycle.
- Sustained collaborations between Mathematical and Physical Sciences (MPS) domains through NSF AI Institutes with stronger DMS focus will create formal structures for interdisciplinary AI advancement.

Final Thoughts

- …

AI is built on mathematics and statistics

AI success relies on fundamental techniques from mathematical sciences.

- Optimization techniques are fundamental to train AI models, with new algorithms significantly reducing computational requirements, energy consumption, and environmental impact of large-scale models.
- Optimal transport, game theory, and stochastic differential equations are fundamental for generative models (e.g, matching flows, GANs, diffusion models).
- Probability techniques such as random matrix theory provide tools to model data and models.
- A big part of AI for scientific inference is based on sampling techniques and simulation-based inference, which relies on long standing statistical techniques for hypothesis training, parameter estimation, and uncertainty quantification.
- Geometric deep learning and equivariant machine learning enables implementation of AI models for science that incorporate the laws of physics into the design of the models (such as symmetries and conservation laws). These models employ techniques from representation theory, invariant theory and geometry to take into consideration data geometry and problem constraints.
- Learning Math is What can potentially make AI smart: Mathematical problems provide important benchmark for reasoning that has been driving AI model development

Opportunities: How mathematical sciences can help the development AI

AI will revolutionize Mathematics and statistics. Let's spur fundamental revolution in deep new mathematics:

- AI is (yet another) example of how deep mathematics give the language for unexpectedly and profoundly useful applications
    - Optimization. New optimization techniques can be developed to reduce computation, energy consumption.
    - Diffusion models can be improved to satisfy domain specific constraints and requirements. This will require different mathematical tools.
- Mathematical/statistical methods for guaranteeing robustness, reliability, privacy, safety, uncertainty quantification and fairness of AI.
- Leveraging mathematical/statistical tools to design more computational and data efficient models. Even small reduction in complexity could lead to very large savings:
    - Fundamental principles for efficient search over spaces of architectures, optimizers, etc
    - Hyperparameters transfer – scaling limits of neural networks (PDEs, random matrix theory, dynamical systems, mean field particle systems)
    - Architecture design – equivariant architectures that encode symmetries and physical law.

Breakout Report: Physics

10 minutes presentation + 5 minutes Q&A

Jesse Thaler, Michelle Kuchera, Gary Shiu

Physics (PHY)

- Physics spans a range of scales with different considerations for:
  - **Large-scale experiments**: Challenges to achieve end-to-end deployment
  - **Small-scale experiments**: Customized methods for bespoke devices (and tools to help)
  - **Theoretical investigations**: Turning statistical methods (e.g. AI) into math (e.g. symbolic regression)
- **Different PHY subfields are interfacing with AI in different ways**
  - Some fields have used machine learning for decades, while others are using off-the-shelf AI for the first time. Many are developing custom AI tools matched to the specific physics problem of interest.
  - AI is blurring the traditional boundary between “theory” and “experiment” (“data physicist”)
- Within the NSF, PHY research happens under the (somewhat arbitrary) headings of:
  - *Atomic, Molecular and Optical Physics*
  - *Elementary Particle Physics*
  - *Gravitational Physics*
  - *Nuclear Physics*
  - *Particle Astrophysics and Cosmology*
  - *Physics of Living Systems*
  - *Plasma Physics*
  - *Quantum Information Science*

Example of AI + AMO

Reinforcement learning implemented in experimental matter-wave interferometry with quantum degenerate atoms

C. LeDesma, K. Mehling, and M. Holland, Vector atom accelerometry in an optical lattice, arXiv:2407.04874.

- RL used to control quantum gases composed of Bose-Einstein condensed atoms stored in 3D optical lattices
- RL enabled atom optic elements that can be performed in the lab, beam splitter, mirrors, etc.
- Are combined to make atom interferometers for precision measurements of accelerometers, gyroscopes, gravimeters, gravity gradiometers

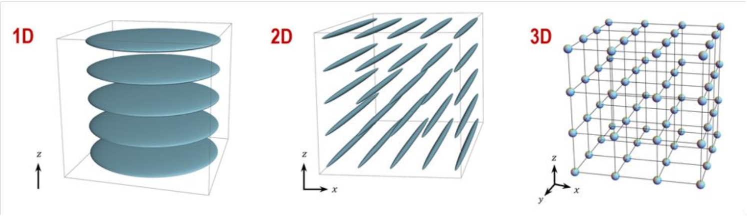

Optical Lattice:
Crystal made of light

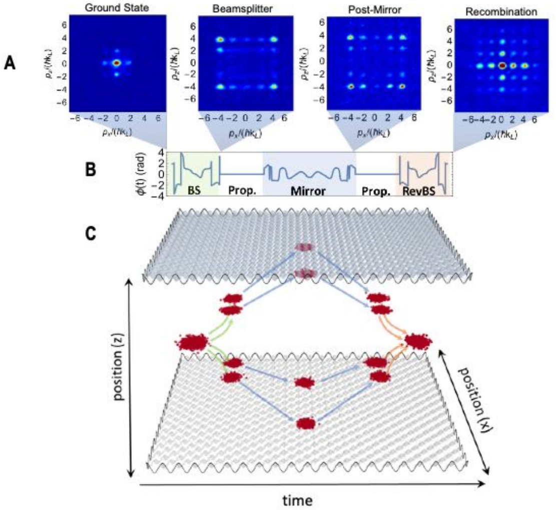

Vector
Accelerometer

Example of AI + Elementary Particle Physics

Anomaly detection for new physics discovery at LHC experiments:
from proof-of-concept to production

Offline analyses [PRL 125, 131801 (2020), CERN-EP-2024-291, PRL 132 (2024) 081801, PRD 108 (2023) 052009]

- Interdisciplinary collaboration: theorists and experimentalists

Real-time AI: anomaly detection trigger taking data in CMS experiment with 50 ns latency [CERN-CMS-DP-2023-079]

- Cross-disciplinary collaboration: scientists, CS, hardware experts, ML researchers
  → fast ML community

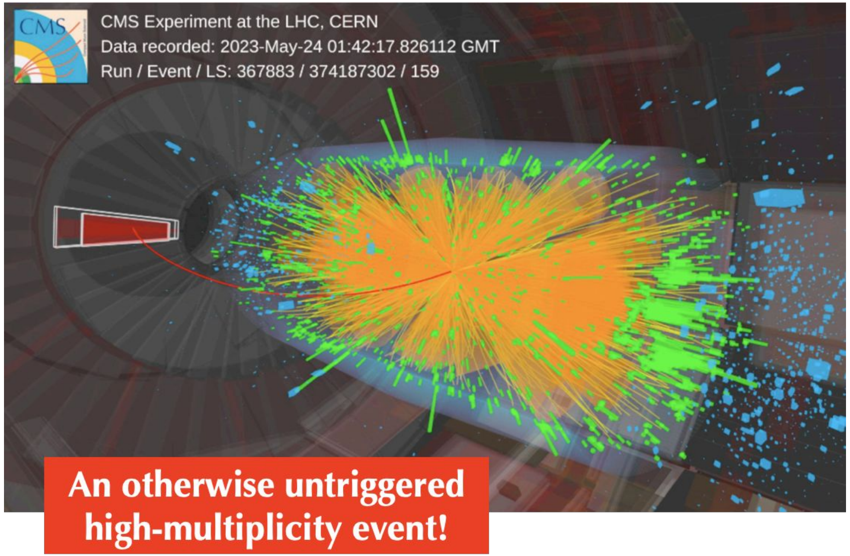

Example of AI + Gravitational Physics

- **Near real-time AI detection** of gravitational wave signals from compact binary mergers in LIGO data:
  - Enable alerts to the astronomical community for follow up.
- Enables multi-messenger astrophysics:
  - Estimate Hubble constant avoiding the cosmic ladder.
  - Probe neutron star equation of state.
  - Probe the r-process of formation of heaviest nuclei.

Simulation of a binary black hole merger.


Perspective | Published: 12 May 2022

Hardware-accelerated inference for real-time gravitational-wave astronomy

Alec Gunny, Dylan Rankin, Jeffrey Krupa, Muhammed Saleem, Tri Nguyen, Michael Coughlin, Philip Harris, Erik Katsavounidis, Steven Timm & Burt Holzman

Nature Astronomy **6**, 529–536 (2022) | Cite this article

Example of AI + Particle Astrophysics

IceCube Neutrino Observatory

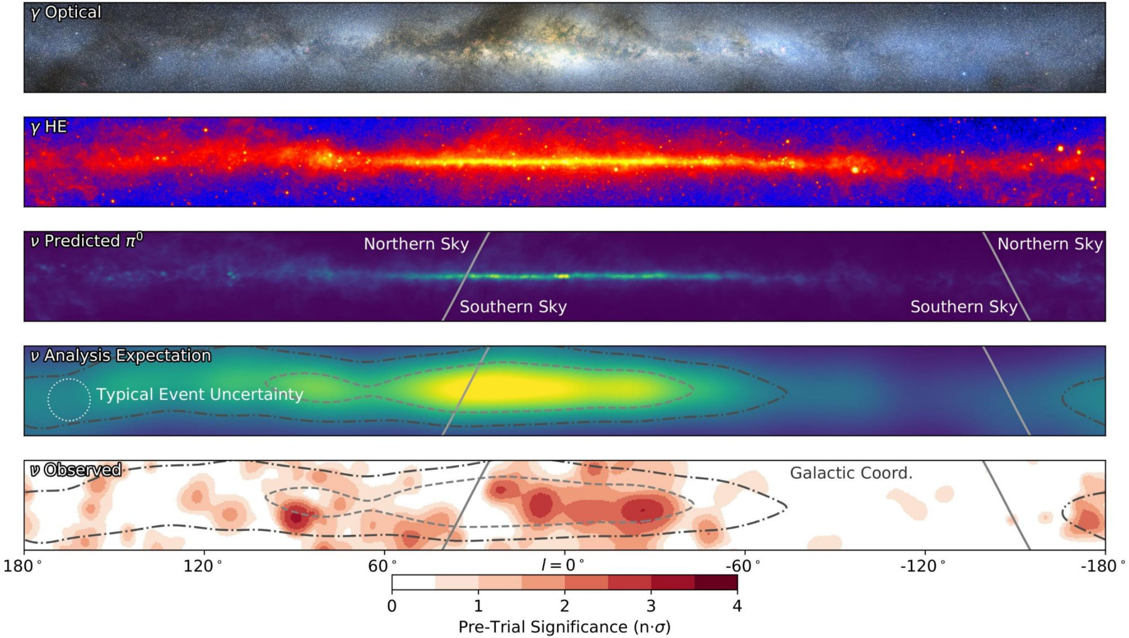

Imaging vs. Inference

"The observation of neutrinos from the Milky Way is a hallmark of the emerging critical value that machine learning provides in data analysis and event reconstruction in IceCube."

https://phys.org/news/2023-06-ghost-particle-image-milky-galaxy.html

Example of AI + Nuclear Physics

The Electron-Ion Collider is going to be the next flagship facility for Nuclear Physics, and will be the **first major NP experiment being designed in the AI era**

AI4EIC: Computing and Software for Big Science 8.1 (2024): 5.

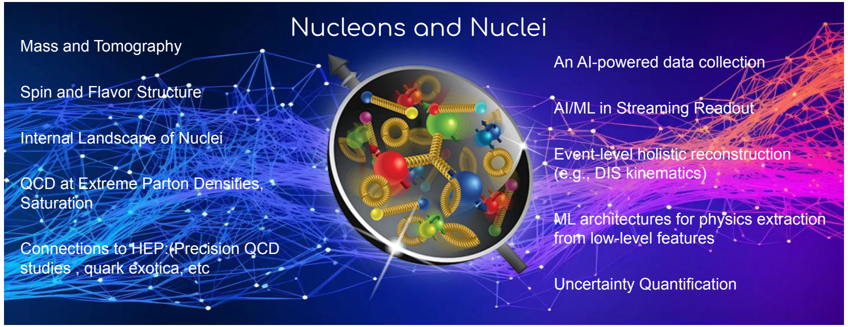

AI-assisted Detector Design for EIC (AID2E)

JINST15.05 (2020): P05009 (dRICH for EIC)
JINST 17.04 (2022): C04038.
NIM-A 1047 (2023): 167748.
AID2E Collaboration JINST 19.07 (2024): C07001

Scalable, distributed Multi-Objective Optimization

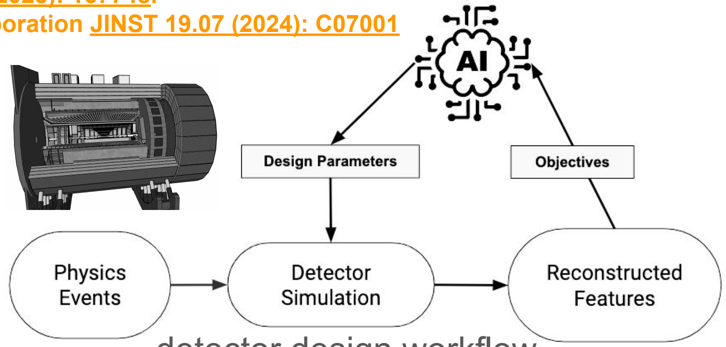

detector design workflow

Towards Agentic Workflows for EIC

JINST 19.07 (2024): C07006 (RAG-based summarization agent for EIC)

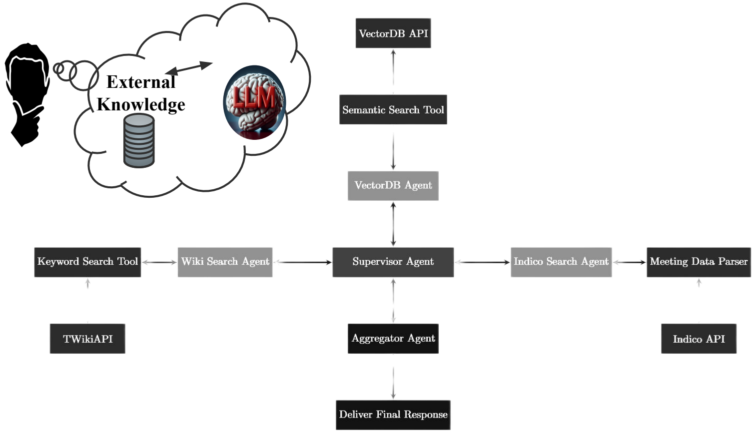

Examples of AI + Plasma Physics

Avoiding fusion plasma tearing instability with deep reinforcement learning
Jaemin Seo et al. Nature **626**, 746(2024)

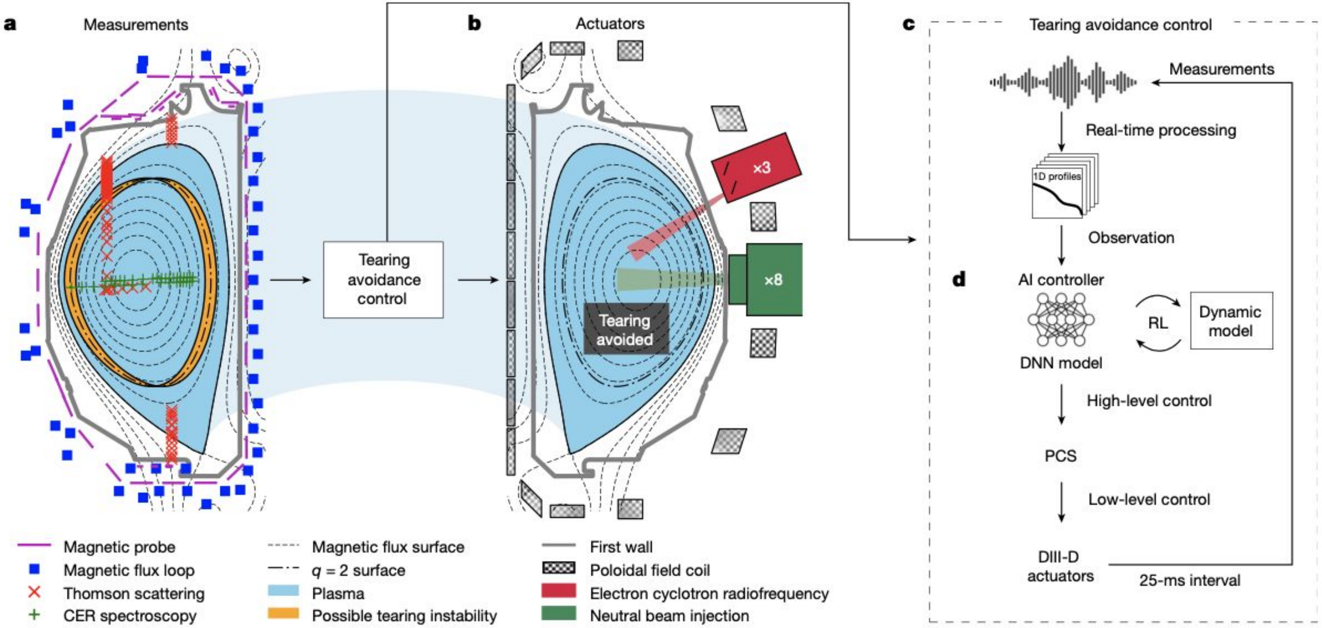

Example of AI "=" Physics of Living Systems

- Development of Deep Learning (Statistical physics + computational neuroscience)
  - Hopfield 2024 Nobel Prize
- Development of Diffusion Models motivated by statistical physics (Jarzynzki's inequality)

Deep Unsupervised Learning using Nonequilibrium Thermodynamics

Jascha Sohl-Dickstein JASCHA@STANFORD.EDU
Stanford University

Eric A. Weiss EWEISS@BERKELEY.EDU
University of California, Berkeley

Niru Maheswaranathan NIRUM@STANFORD.EDU
Stanford University

Surya Ganguli SGANGULI@STANFORD.EDU
Stanford University

Example of AI + Quantum Systems

The use of **neural networks as a variational ansatz** for many-body wavefunctions allowed us to learn new unexpected many-body physics (e.g. quantum tomography). There is now a large community working on **computationally modeling of emergent many-body physics**.

Letter | Published: 26 February 2018

Neural-network quantum state tomography

Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko & Giuseppe Carleo

Nature Physics **14**, 447–450 (2018) | Cite this article

35k Accesses | **600** Citations | **151** Altmetric | Metrics

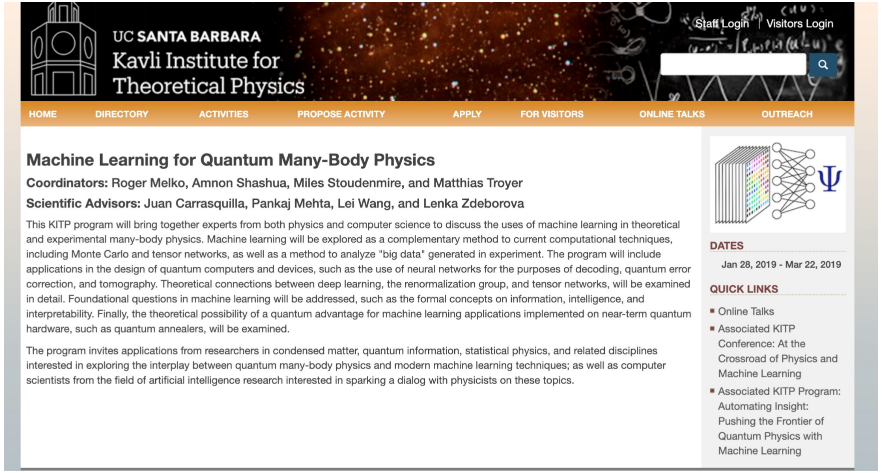

Machine Learning for Quantum Many-Body Physics

Coordinators: Roger Melko, Amnon Shashua, Miles Stoudenmire, and Matthias Troyer

Scientific Advisors: Juan Carrasquilla, Pankaj Mehta, Lei Wang, and Lenka Zdeborova

This KITP program will bring together experts from both physics and computer science to discuss the uses of machine learning in theoretical and experimental many-body physics. Machine learning will be explored as a complementary method to current computational techniques, including Monte Carlo and tensor networks, as well as a method to analyze "big data" generated in experiment. The program will include applications in the design of quantum computers and devices, such as the use of neural networks for the purposes of decoding, quantum error correction, and tomography. Theoretical connections between deep learning, the renormalization group, and tensor networks, will be examined in detail. Foundational questions in machine learning will be addressed, such as the formal concepts on information, intelligence, and interpretability. Finally, the theoretical possibility of a quantum advantage for machine learning applications implemented on near-term quantum hardware, such as quantum annealers, will be examined.

The program invites applications from researchers in condensed matter, quantum information, statistical physics, and related disciplines interested in exploring the interplay between quantum many-body physics and modern machine learning techniques; as well as computer scientists from the field of artificial intelligence research interested in sparking a dialog with physicists on these topics.

DATES

Jan 28, 2019 - Mar 22, 2019

QUICK LINKS

- Online Talks
- Associated KITP Conference: At the Crossroad of Physics and Machine Learning
- Associated KITP Program: Automating Insight: Pushing the Frontier of Quantum Physics with Machine Learning

Specific Considerations for PHY: Gaps

Lack of:

- **Robust “uncertainty estimation”**, which is crucial for reliable predictions in physics
  - A platform to streamline vetting of new AI methods in terms of reliability and reproducibility would make an impact on the field
- **Pipeline for dually trained scientists:**
  - AI+PHY requires double training in two highly complex fields
  - Domain-specific education is already full-time. Consider incorporating AI and computing throughout the major courses, or emphasizing computer programming at early stages.
  - Data science / AI / ML focused researchers may not fall into traditional categories (eg theory or experiment), a problem for hiring.
  - Need AI+PHY tenure-track faculty lines with support from both disciplines.
- **Common vernacular / foundation** when discussing across fields / subfields: e.g. UQ, interpretability
- **Valuing AI+PHY research output within traditional physics metrics**
- **Funding opportunities** to apply for AI+PHY funding (outside of AI Institutes)
- **Portable software stack** to go between heterogeneous AI computing resources

Specific Considerations for PHY: Challenges

- **Scaling Up from Proof-of-Concept to Deployment:** Can there be better mechanisms for theory-experiment communities to interface in large-scale experiments?
- **How to maintain research continuity/technical expertise?**
  - Especially challenging for small-scale experiments.
- **Need to understand the spectrum of "interpretability" requirements:** Physics thrives on understanding solutions, latent structures, coarse graining, effective descriptions
  - Can AI produce new latent representations that model the data better? How do we convert these representations into more familiar mathematical models/equations?
- **Some PHY fields need high-quality simulations**:
  - Need for high-quality simulations based on first principles and well-calibrated to match data
  - Do we need more support for authors of simulations?
- **Physics has unique data types**:
  - Low latency requirements, Multi-scale structures, Preciousness/Deluge of data, time series data, image and image-like data (eg in HEPX)
- **Quantum data/experiments are different from classical data/experiments**:
  - Destructive measurements, phase information, what features to train on (exponential growth)
- **The barrier to entry for much of ML is now much lower than the barrier for responsible use or innovation**

Specific Considerations for PHY: Challenges

- **Data:**
  - **Physics has unique data types**:
    - Low latency requirements, Multi-scale structures, Preciousness/Deluge of data, time series data, image and image-like data (eg in HEPX)
  - **Quantum data/experiments are different from classical data/experiments**:
    - Destructive measurements, phase information, what features to train on (exponential growth)
- **Diverse needs across PHY:**
  - **Need to understand the spectrum of "interpretability" requirements:** Physics thrives on understanding solutions, latent structures, coarse graining, effective descriptions
    - Can AI produce new latent representations that model the data better? How do we convert these representations into more familiar mathematical models/equations?
  - **Scaling Up from Proof-of-Concept to Deployment:** Can there be better mechanisms for theory-experiment communities to interface in large-scale experiments?
  - **Some PHY fields need high-quality simulations**:
    - Need for high-quality simulations based on first principles and well-calibrated to match data
    - Do we need more support for authors of simulations?
  - **How to maintain research continuity/technical expertise?**
    - Especially challenging for small-scale experiments.
- **The barrier to entry for much of ML is now much lower than the barrier for responsible use or innovation**

AI+PHY Priorities & Opportunities

- **Scale proofs-of-concept to production level**:
  - Streamlining (and improving the efficiency) of the process of scaling proof-of-concept to production should be a major priority. This can include better and more realistic datasets and benchmarks, better incentives for developing production level, more efficient processes within the large collaborations.
  - Integrating AI advances into near-term large-scale experiments (HL-LHC, Rubin, LIGO, DUNE, Ice-Cube, …): *e.g.* Real-time AI, high-dimensional deconvolution/simulation-based inference, anomaly detection
- **Developing reliable and reproducible uncertainty quantification methods is a priority.**
  - Opportunity for physics community to contribute to the methodology, given the emphasis on UQ in our field.
  - Include UQ in decision making
- **Leveraging Agentic AI for design/control and operation of experiments**
  - Experimental design of future experiments: EIC/ePIC, LISA, Higgs Factory
  - Automation and operation of experimental systems (HL-LHC, *etc.*)

AI+PHY Priorities & Opportunities

- **Developing Large Physics Models** (foundation models for all parts of physics, other domains) is a big opportunity for accelerating data analysis and scientific discovery at the LHC, LIGO, EIC, *etc.*
- **Integrating AI into Quantum 2.0 Revolution**: e.g. RL for quantum design/control, supervised learning for noisy quantum data, variational ansatz, designing quantum circuits for NISQ devices.
- **Learning Effective Physics Theories**: Use of AI to aid in the discovery of new physics models from data (experimental or simulation). Priorities are to support the development of physics-aware/physics-consistent AI. Use of AI to augment physics simulation with high fidelity, e.g.,multi-scale dynamics of complex physical systems or detector response.
- **Cultivate "Physics of AI"**: Build upon rich tradition of statistical physics/physics theory/experimental advancements for understanding and developing AI

Interdisciplinary Opportunities

- **Annual, intentionally designed conferences** across AI+MPS:
  - Cultivate common vernacular in presentation/goals
  - Create pathways of interdisciplinary connection
- **More institute-like structures** with multidisciplinary work surrounding AI (MPS+AI Institute)
- **AI collectives across scientific disciplines** (include individual group funding, larger, multi-PI collabs, etc):
  - **E.g. Real-time AI**: Developing tools and methods to embed robust AI on specific hardware. These challenges are interesting also outside the PHY domain and foster collaboration among domain scientists, ML researchers, engineers, CS and hardware experts.
  - **E.g. Curated datasets:** Creating benchmarks to communicate challenges to non-domain experts and involve them in innovation. Create repositories of open data for domain experts
  - **E.g. Foundation Models:** Community building foundation models within the context of their own domain, or bridging domains
  - **E.g. Simulation-based inference:** E.g. triggering, anomaly detection, surrogate modeling, These techniques are common across nuclear/particle physics (e.g. LHC) and cosmology (explosion of data in e.g. galaxy surveys), and computational biology.
- **Embedding domain knowledge into AI models:** E.g. symmetries, fundamental laws, etc. The development of general strategies to embed such constraints into AI models will be beneficial across MPS domains.

Interdisciplinary Connections

- **Multi-scale/multi-physics modeling**: AI could help address the grand computational challenge of modeling complex physical systems. E.g. recent success stories are AI-driven weather/climate modeling.
- **Collaboration with DMS+CISE on "Science of AI"**: statistical physics analysis of AI, generative models, reliable AI, new techniques for sampling and simulation
- **Collaboration between DMR, ENG, CISE, and AMO (PHY) for AI+Quantum 2.0:** quantum computing, quantum control, quantum simulation, ANISQ, quantum information. This gives new interesting problems for AI due to properties of quantum (i.e. quantum measurement);
- **Collaboration with CHE, stat mech (DMR), physics of living systems (PHY)**: all involve problems of inverse design, which can benefit from cross-fertilization of ML techniques and tools (RL, designing dynamical systems, algorithms for learning).
- **Collaboration with AST on multi-messenger astrophysics**.

Final Thoughts

Coming into this workshop, I expected many of the commonalities of our work and challenges surrounding AI to emerge, which they did.

However, I also realized that domain-specific challenges in AI+MPS provide unique and exciting opportunities for the advancement, development, and evaluation of AI.

Backup Slides

Flagship Examples in Gravitational Physics

- Use AI to detect gravitational wave (GW) signals from compact binary mergers near real-time in LIGO data and issue alerts to the astronomical community for follow up.
  - Enables multi-messenger astrophysics, where an event is observed in both GWs and light. Example: binary neutron star merger is also a short Gamma Ray Burst and a kilonova. Enables Hubble constant estimate avoiding the cosmic ladder; probes neutron star equation of state; probes the r-process of formation of heaviest nuclei.
  - Gunny et al, Nature Astronomy 6, 529–536 (2022)
  - M. Dax et al, Phys. Rev. Lett. 127, 241103 (2021)
- Use AI to remove contamination from GW data due to environmental and instrumental effects.
  - Effectively a complex filter, takes many witness channels (accelerometers, microphones…), and accounts for both linear and non-linear couplings
  - Ormiston et al, Phys. Rev. Research 2, 033066 (2020)
- Use AI to remove the compact binary signals from the data, since they form a foreground masking weaker signals e.g. from the early universe. This will leverage the above techniques, and also combine them in hierarchical Bayesian formalisms, such as the one discussed by Zhong et al, arXiv: 2406.10757
- Use AI to detect GW signals from unknown sources, unknown waveforms - anomaly detection in time series data
- Use AI to do inference in cases where signals have complex statistical distribution. Example is angular power spectrum of the stochastic GW background, which is complex, multivariate, chi-squared distributed, with non-trivial covariance. Approaches such as SBI may offer a way to extract information on cosmology or astrophysics that generates these signals.

Flagship Examples in Particle Physics

- What are **flagship examples** of AI+PHY that will help demonstrate where the field is now?
  - (What **key drivers** are advancing Physics? (i.e. Why was AI needed))

  - Real time AI:
    - Hls4ml: https://fastmachinelearning.org/hls4ml/
    - CMS anomaly detection trigger: https://cds.cern.ch/record/2876546?ln=en
    - Extreme edge for data compression/reduction: smart pixels (https://arxiv.org/abs/2406.14860) and high granularity calorimeter readout
  - Jet tagging: AI has accelerated jet tagging at LHC over the past ~7 years, the ROC curve for top tagging has improved by a large factor using SOTA deep learning like transformers and graph neural networks. One example of a major milestone: https://arxiv.org/pdf/2202.03772
  - Anomaly Detection
    - CMS Anomaly Detection Search https://arxiv.org/abs/2412.03747 (submitted to ROPP)
      - New model-agnostic search techniques
    - LHC Olympics 2020 anomaly detection data challenge (https://arxiv.org/abs/2101.08320)
    - Darkmachines anomaly score challenge (https://arxiv.org/abs/2105.14027)
    - Data quality monitoring: https://arxiv.org/abs/2311.04190
  - Simulation: CaloChallenge 2022: A Community Challenge for Fast Calorimeter Simulation [arxiv:2410.21611](https://arxiv.org/abs/2410.21611)
      - a comprehensive comparison of state-of-the-art generative models on calorimeter shower datasets, evaluating their quality, speed, and size.
  - Higgs boson Discovery in CMS (PU Jet Id/Photon Regression/Photon BDT/B-jet Regression/MVA Met) https://arxiv.org/abs/2212.05081
      - https://arxiv.org/abs/1207.7235
  - Single top Discovery

NSF AI+MPS Workshop: *March 24–26, 2025: MIT*

Flagship Examples in Nuclear Physics

Nuclear Physics:

- AI/ML is largely utilized in NP. Comprehensive review papers came out recently, e.g., Boehnlein, Amber, et al. "Colloquium: Machine learning in nuclear physics." Reviews of modern physics 94.3 (2022): 031003. Experiments at the next generation US-led nuclear facility, the Electron Ion Collider, are currently being designed leveraging AI and plan to become an "AI-powered data collection" (https://www.bnl.gov/eic/epic.php). A community paper has been published as Allaire, C., et al. "Artificial Intelligence for the Electron Ion Collider (AI4EIC)." Computing and Software for Big Science 8.1 (2024): 5. 2023 Long Range plan highlights AI/ML and computational resources as an important cross-cutting priority for the first time.

- High precision reconstruction of kinematics observables in DIS, SIDIS and other fundamental processes in ep-scattering experiments will be fundamental in achieving a more comprehensive nucleon tomography. DL-based methods could outperform traditional reconstruction as demonstrated by Arratia, Miguel, et al. NIM A: 1025 (2022): 166164 and Diefenthaler, M, et al. EPJ C 82.11 (2022): 1064. More recently, C. Fanelli, and J. Giroux. Machine Learning: Science and Technology 5.1 (2024): 015017 proposed a method to event-level uncertainty quantification (UQ) for DIS events utilizing Bayesian Neural Networks with Multiplicative Normalizing Flows. All these works utilized full simulations from HERA, and leverage full event reconstruction possible in ep-scattering experiments which typically have lower multiplicities, more "controlled" final states with "reduced" backgrounds compared to LHC (characterized by multi-partonic interactions, underlying event, QCD radiation, hence more complicated event-level reconstruction).

- Nuclear physics experiments (ongoing at JLab and future like at EIC) will leverage Cherenkov detectors for charged particle identification (for momentum range up to ~8 GeV/c); these methods within Deep(er)RICH (based on Swin Transformers and Normalizing Flows) outperform standard reconstruction methods C. Fanelli, Cristiano et al " Machine Learning: Science and Technology 6.1 (2025): 015028.

- A RAG-based summarization agent has been introduced in the EIC community Suresh, Karthik, et al. Journal of Instrumentation 19.07 (2024): C07006. – RAG4EIC project, towards agentic workflows

-

Flagship Examples in Plasma Physics

Plasma Physics:

- AI/ML to predict disruptions in Tokamak fusion devices (Kates-Harbeck 2019, Seo 2024)
- Development of surrogate models of plasma turbulent transport to predict macroscopic profiles of magnetically confined plasmas [Rodriguez–Fernandez 24]. Development of moment closures for fluid models of nonlinear kinetic plasma dynamics.

Coming Up Next…

Today's Schedule (Wednesday)

9:00–10:15 am: Reports by Domain

- Astronomical Sciences
- Chemistry
- Materials Science
- Mathematical Sciences
- Physics

10:15–10:45 am: Break

10:45–11:45 am: Open Discussion

11:45 am–12:00 pm: Closeout

12:00–1:00 pm: *Lunch is served (optional)*