Welcome to Day 2 of the AI+MPS Workshop!

March 24–26, 2025

Today’s Schedule (Tuesday)

9:00–10:30 am: Reports by Theme

- Interdisciplinary Research: Opportunities and Challenges
- Interdisciplinary Research: Resources Needed

10:30–11:00 am: Break

11:00 am–12:30 pm: Reports by Theme (cont’d)

- Education & Workforce Development
- Responsible AI

12:30–2:00 pm: Lunch

2:00–5:30 pm: Domain Breakouts

- Astronomical Sciences: Room 804
- Chemistry: Room 812
- Materials Science: Room 803
- Mathematical Sciences: Room 801 North (here)
- Physics: Room 801 South

Breakout Report: Interdisciplinary Research: Challenges and Opportunities

25 minutes presentation + 20 minutes Q&A

Lars Ruthotto, Sijia Dong, Chad Risko

Common Interdisciplinary Opportunities

- AI is a key for simulation-based inference and opens new avenues for **drawing inference of variables in models so complex that likelihoods are intractable or unknown**. This can bring new insights across disciplines (e.g., into galaxy formation, collider event modeling, …). AI's potential to model high-dimensional probability distributions can help scientists quantify and understand uncertainties of predictions.
- AI techniques such as reinforcement learning can help **scientists find needles in the haystack**. RL is a flexible framework to explore search spaces too large for humans. RL has potential to discover new drugs and materials and even construct counter examples for mathematical conjectures, for example. AI can also find patterns and lead to discovery in massive scientific datasets from space telescopes, LHCs and other facilities.
- **AI enables scientists to augment incomplete, simplified first-principle models with data to simulate**, understand, and control complex processes. AI can also be used to learn low-dimensional representations and build surrogate models for complex systems. Thereby, AI can help build bridges between the real and physical world and be incorporated in digital twins, self-driving labs, human machine interfaces.
- Large AI systems that are pre-trained on large datasets from observations and simulations could form **foundation models and enable breakthroughs in applications with limited physical models and small, noisy data sets** (e.g., study chemical reactions at equilibrium, construction defects in additive manufacturing).
- **AI can STEAR the scientific workflow** by STreamlining background research, Effectively collecting useful data, Automating design of experiments, Rapidly testing and validating hypotheses. AI can help translate research literature, reduce language barriers, and streamline education. **AI can fundamentally revolutionize how we do science.**

Common Interdisciplinary Challenges/Barriers - 1

- **Limited Availability and Flexibility of Funding Opportunities** in AI + MPS (e.g., single principal investigators, conference/workshops, seed funding, NIH style fellowships instead of projects, centers) Flexible review criteria are necessary, as benefits may not be equally distributed across all areas.
- **Lack of Qualified Reviewers and Communication Gaps**: A shortage of qualified reviewers for proposals and publications exists, coupled with insufficient communication between independent, domain-specific review panels evaluating identical proposals. Potential bias biases for or against AI integration.
- **Support for Interdisciplinary Faculty**: Institutions often lack adequate support for interdisciplinary faculty, particularly concerning tenure reviews, which can hinder the advancement of AI-integrated research.
- Challenges in **Interdisciplinary Collaboration**:
  - **Machine Learning (ML) + Science Domain**: E.g., attracting ML researchers and training domain scientists in ML methodologies, alongside the necessity for appropriately curated data.
  - **Science Domain 1 + Science Domain 2**: Difficulties arise from differing terminologies and communication styles between disciplines. ML and data science could potentially serve as a common language to bridge these gaps, facilitating cross-disciplinary collaboration.
- **Communication Across Domains**: Researchers may struggle to comprehend research questions, hypotheses, data, and quality of results from unfamiliar domains. For example, a computer scientist interpreting protein or single-cell data faces challenges due to the domain-specific nature of data curation. Overcoming personal and disciplinary biases is essential, akin to bridging the gap between LaTeX and Word users
- **Opaque Publication Landscape**: Inconsistent standards for model validation, an overwhelming volume of ML papers, difficulties in implementing published models, and the prevalence of closed-source models contribute to a lack of transparency in AI research publications.

Common Interdisciplinary Challenges/Barriers - 2

- **AI systems are often black box.** Need new insights from MPS disciplines to develop guidelines and theory that aid in designing and understanding of AI and its predictions.
- **AI models do not obey physical laws**, which is critical in some (but not all?) use-cases**.** Integrating fundamental physical principles (ODEs, PDEs, symmetries, invariances, equivariances, …) into AI models remains a significant challenge.
- **AI models are inefficient.** AI models often require substantial computational resources and large datasets for training. This high demand can lead to inefficiencies, especially when models are applied to complex scientific problems where data may be scarce or expensive to obtain.
- **AI models don’t know what they don’t know.** Reliable predictions in high-stakes applications require error estimates, out-of-distribution detection, …
- **RL models have trouble with extremely sparse rewards,** which is often the case in scientific applications.
- **AI requires hardware and support.** AI companies have teams of engineers to implement and experiment with models. Academic institutions lack this infrastructure.
- **AI has steep learning curve**
- **AI doesn’t always work as advertised.** Need computational and data resources for reproducibility

We need to develop the Science of AI to overcome current limitations of AI systems

Facilitating Interdisciplinary Collaboration

- Initiate Collaborations (funding opportunities targeting interdisciplinary groups, community-building conferences/workshops). Seed funding and cross-sectional meetings to define commonly beneficial projects and initiate interdisciplinary collaborations. Not all of them need cutting-edge AI: industry research has different objectives, so collab with industry research depends on finding similar objectives
- Maintain Collaborations - continue to host workshops, events, mixers, etc to foster a strong community. Promote flagship journals for ML + Science articles to foster community. Funding mechanisms could be 3 years + 2 years.
- Create online repositories to host short tutorials written by domain experts for AI researchers (e.g., about problem formulation and potential impact), and vice versa (e.g., explain how AI algorithm works using examples in a target domain). See, NAIRR
- Problem: AI & MPS researchers mutually benefiting from the research, but it may not be appropriate to require equal advancement in AI and MPS. Idea: Recognition of interdisciplinary research in academic reviews/hiring/promotion (Faculty positions specifically on AI+MPS, mainly evaluated by domain scientists). Establish new awards, prizes and other forms of recognition in AI + MPS
- Scientific reviewing process - database of reviewers. Acknowledge the role of trainees
- Look to success in Science + Tool Development (analytics, equipment, hardware, etc) interdisciplinary collaborations for examples of good communication across domains
- How to speak a common language across domains? Could we use the language of ML and data science to abstract a common lexicon away from specific domains? (eg domain1 has a novel solution to a classification problem and domain2 is looking for solutions to some other classification problem)
- Train domain scientists in AI. Have consistent definitions across AI.
- New conference series on AI for Science

Key Priorities/Recommendations

- **Create a program/call for AI for Science** that would specifically fund projects / scientists that pursue new science with AI and projects that develop the Science of AI, and recruit reviewers with sufficient knowledge of AI and the involved disciplines. When proposals are being reviewed by independent panels with different domain expertise, there should be a mechanism for interaction prior to decision.
- **Long-term funding opportunities for interdisciplinary projects** to allow for collaborations to build the capacity to bridge the language of the domain science(s) and AI expertise. These would also serve as training hubs for interdisciplinary research, and to educate institutions on AI.
- With domain-specific AI institutes, provide funding opportunities for external PIs (or smaller institutes) to possibly join / collaborate the larger institute.
- Educational programs (eg training programs/grants), events enabling collaboration,
- Inter-domain (physics, astronomy, chemistry, DMR, mathematics) + AI funding / projects. Also encourage intra-domain projects that are based on the use of AI.
- Enabling careers paths that cross disciplines. Create training opportunities, encourage curriculum development, offer trainee fellowships, incentivize joint hires, provide awards/recognition for ML+Science. Encourage, community appreciation for interdisciplinary research including promotion and tenure.

Common Interdisciplinary Challenges/Barriers

- Limited portfolio (single/few PI, centers, etc) of funding opportunities supporting AI+MPS research. Use flexible review criteria: benefits may not be equally distributed across both sides

- Lack of qualified reviewers for proposals / publications & insufficient communication between independent (domain specific) review panels that see the same proposals. Resistance from established communities, bias for/against AI.

- Support for interdisciplinary faculty at their institutions, tenure reviews

- Different kinds of interdisciplinary: ML + Science Domain or Science Domain 1 + Science Domain 2

    - ML + Science Domain: difficulties in attracting ML researchers and challenge to train domain scientists in ML, and appropriately curated data.

    - Domain1 + Domain2: different languages, efficient communication – could ML and data science itself be a common language? Methods can provide easy ways to cross disciplines!

    - Understanding data from a different domain. Ex: Comp Scientist trying to understand Protein/single cell data. Nature of data curation is/maybe very domain specific. LaTEX vs. Word - an analogy for needing to overcome personal / domain biases

- Opaque publication landscape: inconsistent standards for model validation, overwhelming quantity of ML papers, hard to implement published models, closed-source models.

- Communication: Domain problem articulation is poor, limiting interest from ML experts. Conversely, ML solution articulation is poor, limiting appreciation from domain experts.

Common Interdisciplinary Opportunities

Fundamentally revolutionize how we do science

- Augmenting first-principle models (PDEs, ODEs, symmetries, domain expertise,) with data. Imbue AI models with physical principles.
- Foundation models pre-trained on large datasets with downstream applications that potentially cross disciplines
- Create AI to handle small and/or noisy datasets (could be applications of previous two bullets)
- Accelerating experimental design, simulations, predictions, such as self-driving labs, using LLM to summarize & screen literature, data analysis (for high-throughput experiments); one way to achieve it is to use AI to assist coding
- AI for uncertainty quantification (simulation-based inference / prediction to bridge the gap between simulation and data); AI can model high-dimensional, complex probability distributions
- Reinforcement learning for discovery (drug discovery, materials, counterexamples in combinatorics, algorithmic discoveries), control; Leverage RL's potential to find 'needle in the haystack'
- Data mining, analyzing large-scale scientific data (experimental, simulations, …)
- Representation learning, surrogate modeling
- Synthetic experiments, digital twins (virtual/experimental). Collaborations with robotics engineers, etc., to develop human-machine interfaces.
- Real-time AI: integrates closed-loop systems with novel AI models, coprocessors and inference at scale to enhance data analysis, experimentation and dynamic physical system control to accelerate discovery.
- Use LLMs to translate research literature to reduce the "language" barrier between different fields.
- Education and workforce development - domain + data science training. Increasingly a demand in industry.
- Extracting new domain knowledge (physics, chemistry, etc.) from data. New representations for challenging hypotheses with human in loop.

ChatGPT generated summary

Here's a **summary of key takeaways** from the slides titled **"Common Interdisciplinary Challenges/Barriers, Opportunities, and Recommendations"** in the context of AI and physical sciences:

🔧 **Challenges & Barriers**

- **Funding & Review Process:**
- Limited funding opportunities for AI + physical sciences.
- Rigid review criteria; unequal benefits across disciplines.
- Lack of qualified interdisciplinary reviewers; siloed review panels.
- **Cultural & Institutional Resistance:**
- Bias for/against AI from traditional communities.
- Interdisciplinary faculty face difficulties in tenure and institutional support.
- **Communication & Integration Difficulties:**
- Cross-domain language barriers (e.g., science-to-science or ML-to-science).

Final Thoughts

-

Breakout Report: Interdisciplinary Research: Resources Needed

25 minutes presentation + 20 minutes Q&A

Andrew Ferguson, Pankaj Mehta,
Francisco Villaescusa Navarro

Funding

- AI for science directorate. Currently, not clear funding for interdisciplinary work
- Fellowships for AI+Science (e.g. Schmidt fellows)
- Funding to
  - Collect or generate data
  - Store data in a centralized place
  - Write and/or maintain code
  - Create benchmarks
- Train/teach faculty members
  - Simons Foundation Pivot program, data science - domain science
  - personnel “exchange programs”

Data and software

- "Good data" vs "more data". Is worth the cost to clean up the data?
- Make data publicly available (may require culture change; national labs vs universities
- Centralized data to store the data. Easy to access to it.
  - Huggingface for Science? Is HuggingFace good enough?
  - Invest in storage
- Create and maintain benchmarks
  - Results should be reproducible (industry standards may be different)
  - Math requires very rigourous benchmarks. In other fields hard to define.
- Investment in people who can
  - Write and maintain code (python2 vs python3)
  - Consolidate and distribute data
  - Avoid redundancy
  - Define a career (and competitive salary) for them. European model for engineers

Hardware

- Invest in common hardware (e.g. national facilities vs local clusters)
- Foundation models vs domain-specific models
  - Can we do this?
  - What is our role vs industry?
  - Carbon footprint
  - Embed our prior/physics knowledge (e.g. PINN)
- Diversify and invest in different types of hardware
- National support for LLM/foundation models
  - dedicated hardware for open models
  - subscription for closed models

ML and industry partnership

- Interdisciplinary projects are challenging
  - Difficult to capture interest from ML people
  - Clear funding paths
  - Interdisciplinary reviewers for grant committees

- Invest in partnership with industry
  - Promote workshops at top ML conference focused on science
  - Involve ML people in future editions of this workshop
  - Fund ML engineers to work on science programs
    - Encode program in the UK
  - Encourage ML people to participate in science hackathons

Education

- Bootcamps for teaching undergraduate how to code
- Building communities
  - Slack channels (e.g. training large-scale transformers)
- Publishing null results
- Outreach
  - Communicate our results
- Create spaces for cross-talks. Many similar problems, very different languages
- Using agents to (co-)teach students?
- Large vs small teams
  - Large teams can make very fast progress
  - We may not be used to this style in academia
- Encourage AI institutes to collabore broadly

Paradigm change?

- What do we expect from ML?
    - Better and faster science
    - Completely new way
- AI-agents
    - AI-scientist
    - Google co-scientist
    - Is academia the best place for this?
- Funding should be very different.
    - Do we even have the money?
    - Role of industry
    - Will we just generate data for industry?

Recommendations

1. NSF hosted infrastructure for data and models/code with professional support for longevity and usability

2. Upskilling fellowships / “exchange programs” for faculty to move between data science and domain science

3. benchmarks and scoreboards, where appropriate, but may not be a one-size fit

4.

Coming Up Next…

Today’s Schedule (Tuesday)

9:00–10:30 am: Reports by Theme

- Interdisciplinary Research: Opportunities and Challenges
- Interdisciplinary Research: Resources Needed

10:30–11:00 am: Break

11:00 am–12:30 pm: Reports by Theme (cont’d)

- Education & Workforce Development
- Responsible AI

12:30–2:00 pm: Lunch

2:00–5:30 pm: Domain Breakouts

- Astronomical Sciences: Room 804
- Chemistry: Room 812
- Materials Science: Room 803
- Mathematical Sciences: Room 801 North (here)
- Physics: Room 801 South

Breakout Report: Education & Workforce Development

25 minutes presentation + 20 minutes Q&A

Yuan-Sen Ting, Yaroslava Yingling, Shuwen Yue

Urgency

- **Emerging AI technologies are disrupting the educational process**: the time to act is now
- Graduates with hybrid skills in AI+STEM are already in high demand, reflected by a 28% increase in job postings explicitly requesting AI competencies within engineering fields[i]. By 2030, over 70% of global companies are expected to adopt AI-driven solutions, creating unprecedented demand for engineers skilled in applied AI[ii].
- **Need:** The creation of a well trained AI-ready workforce; to continue the leadership of US in AI+STEM

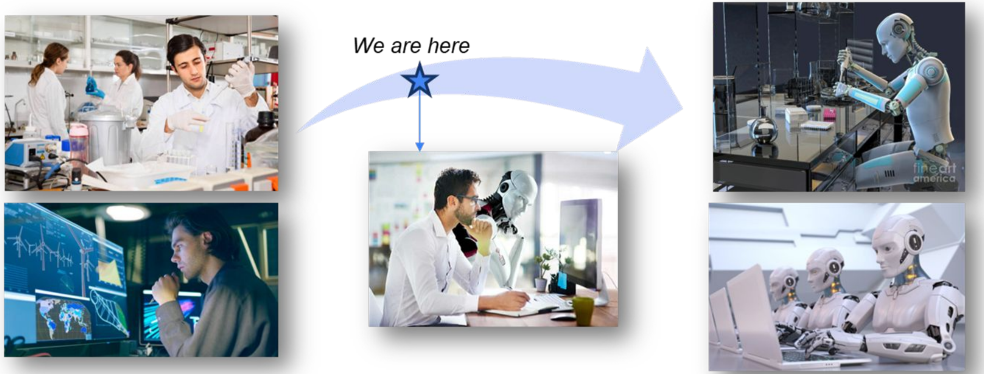

[i] Burning Glass Technologies & Business-Higher Education Forum. (2021). *AI Skills in Engineering Occupations.*
[ii] McKinsey Global Institute. (2022). *Artificial Intelligence Adoption Report.*

Key Priorities and Findings

- Integrate AI Across Curriculum: Incorporate AI into domain-specific courses and support flexible, modular learning pathways.
- Support and incentivize the development of innovative AI teaching tools and leverage AI tools to improve teaching quality and efficiency
- Strengthen student preparation in math and science to support meaningful AI engagement.
- Prepare an AI-Ready Workforce: providing a strong AI foundation that will have enduring effect and enhance public understanding of AI in MPS.
- Retrain Educators & Researchers: Provide AI and STEM upskilling opportunities for K–12 teachers and current faculty and researchers.
- Modernize Assessment & Credentials: Create new ways to assess AI-assisted learning and validate informal AI expertise.

Recommendation 1: Build Cross-Disciplinary AI Curriculum

- Support **dual-specialization programs** like AI+Physics or AI+Chemistry at undergrad, grad, and postdoc levels to develop deep, cross-cutting expertise.
  - Example: **co-teaching models** (e.g., CS faculty + domain faculty).
- **Incorporate AI into domain-specific courses** and support flexible, modular learning pathways.
  - Offer training and incentives to help faculty incorporate meaningful AI content into existing courses
  - Find a way to add sufficient amount of AI training into current domain curriculum:ensuring alignment with accreditation standards like ABET without overloading the curriculum
- Ensure non-CS students have **access to foundational AI tools**, concepts, and ethics
- Use AI tools like **LLM-agents to support learning**, not replace it
  - Design assessments that value understanding, not tool use, ensuring fairness for all students—AI users or not.
- Deprioritize publications as the main outcome for undergraduate research in AI; **emphasize skill-building and applied experience**.
- Offer **stackable credentials** and AI training modules for PhD students and postdoctoral researchers in academia and industry, that can be tailor to their research needs.
  - How to validate AI skills gained through hands-on research, not just coursework, in credentialing and job pipelines.

Recommendation 2: Build infrastructure for teaching materials/tools

- **Provide the curated set of evolving best practices** of how to incorporate AI in domain teaching
    - Create **AI teaching toolkits and on demand resources**
    - AI teaching infrastructure that reduces workload while enhancing engagement.
    - **LLM-based assistants** and instructional design support for faculty;
    - Support and **encourage collaborations** across various universities on best practices
    - Set-up a **national infrastructure for AI Education**, leveraging AI institutes
    - Support a **series of interrelated stackable summer schools** on AI and applications
- **Resources** for running the educational modulus: GPUs and servers
- Encourage a diverse range of hands-on AI+MPS training opportunities.
- Pursue multiple strategies to incorporate AI and related topics into MPS curriculum.

Recommendation 3: Continuing upskilling education in AI

- Offer flexible, **just-in-time training for faculty** and researchers to stay current with AI advances.
    - **Incentivise** the re-training: supplements to current NSF grants or NSF funded sabbaticals
    - encourage faculty to adapt to incorporate the latest developments in the courses and programs - this is particularly important in fast moving fields like AI.
    - Use AI tools to automate routine tasks and free time for higher-order thinking and creativity.
    - Encourage faculty provide feedback through documenting failures.
- Supporting programs that help faculty **AI-proof their courses.**
    - How do we balance this with modern pedagogical practices that favor authentic, untimed evaluations like projects
- Create more **AI+MPS postdoctoral fellowships**.
- **Reverse sabbatical:** Create mechanisms for **industry professionals** to teach short-term or part-time

Recommendation 4: Support the establishment of a common curriculum materials for AI Literacy

- AI can be a **great enabler** to emphasise the importance of STEM / meth education.
- Design flexible “**AI literacy pyramids**” that grow with student interest and depth.
    - Develop a **shared framework for AI literacy** across all stages of education and training.
    - Support the establishment of a **common curriculum materials**.
- Encourage **cross-departmental teaching** that blend computational & domain knowledge.
- Support Continuing Education: Offer **modular AI** upskilling opportunities.
- Students are entering college with less STEM background: run **STEM workshops** for K-12

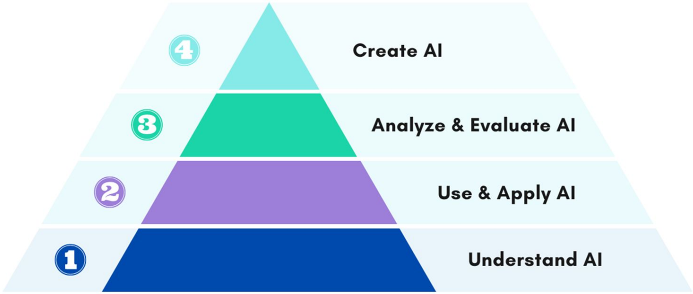

https://er.educause.edu/articles/2024/6/a-framework-for-ai-literacy

Recommendation 5: Assessment and evaluation of AI proficiency

- Assess outcomes that emphasize tangible skill development over traditional academic outputs such as papers. Prioritize collaboration, applied experience, and thoughtful application

- Formalize credentials for AI skills for AI+MPS researchers
  - For students who gain experience in research and coursework but not in AI/CS degree program
  - Deprioritize publications as the main outcome for undergraduate research in AI; **emphasize skill-building and applied experience**.
  - For postdocs who gain AI skills through applied and translational projects, not just formal coursework.
  - Faculty who upskill in AI will provide better education and training for their students
  - This is important for job pipelines, graduate admissions, tenure and promotions.
  - Consult with industry to understand desired credentials to hire AI+MPS researchers
- Shift focus to interdisciplinary problem-solving.
- Encourage the use of FAIR data sets/standards for AI+MPS education.

Key Priorities and Findings

- Integrate AI Across Curriculum: Incorporate AI into domain-specific courses and support flexible, modular learning pathways.
- Support and incentivize the development of innovative AI teaching tools and leverage AI tools to improve teaching quality and efficiency
- Strengthen student preparation in math and science to support meaningful AI engagement.
- Prepare an AI-Ready Workforce: providing a strong AI foundation that will have enduring effect and enhance public understanding of AI in MPS.
- Retrain Educators & Researchers: Provide AI and STEM upskilling opportunities for K–12 teachers and current faculty and researchers.
- Modernize Assessment & Credentials: Create new ways to assess AI-assisted learning and validate informal AI expertise.

Breakout Report: Responsible AI

25 minutes presentation + 20 minutes Q&A

Pratyush Tiwary, Becky Lindsey, Rene Vidal

Key Priorities and Recommendations

- **Responsible AI:** Broadly comprises openness, transparency, reliability, sustainability, & guardrails
  - Is very dependent on discipline, method, and capability, but recommendations can and still should be made for those specific contexts
- **Challenges** arise from both mistrust and blind faith; **can be addressed by:**
  - Funding research into generic design principles for responsible/ethical AI; better incentives, e.g. expanding “metrics of success” and funding priorities to include databases, benchmarks, software development/maintenance
    - Emphasizing research into interpretability, robustness, and UQ
  - Integrating ethics of AI course into curricula
  - Establishing programs to increase AI literacy in the community
  - Developing centaurs: training domain scientists on AI and AI students on domain science

Academic/Scientific Integrity

Traditional scientific integrity:

- Do not plagiarize; define methods used, cite sources, and provide interpretation of results.

Guiding question: Does AI require an updated definition of scientific integrity?

- What does plagiarism mean within the context of AI? e.g., are chat models just writing tools? What if these tools are used to facilitate ideation?
- Consider data lifecycle: How might it be used? Use documentation/FAIRing to fight against misuse, intentional or otherwise.
- Verify and validate models and predictions.
- Delineate what should non-expert users be on the lookout for; define limitations of the approach.

Interpretability, robustness, and UQ

General Observations:

- Interpretability, robustness and UQ are critical for the adoption of AI methods across many disciplines, especially in high-stakes domains such as medicine, thus it should be a priority.
- Interpretability, robustness and UQ take different forms depending on the application domain. Interdisciplinary work between AI and domain scientists is needed to address these topics.
- Recognize distinction between need for **both** new methodological developments **and** better integration of these approaches into applied AI research; collaboration could help
- When presented, new methods should include clear discussion of:
    - What is the domain of applicability and how do you measure it
    - Whether transparency/interpretability is relevant for the application and why/not?

Distrust of AI

Guiding Question: Is community distrust an issue? How do we address it?

- Scientists should approach any work with a healthy amount of skepticism.
- Increased emphasis on verification and validation is needed in certain subdomains (e.g., computational physics/chemistry/materials science), which begs: can succinct standards be set?
- Presentation of AI-derived results should be accompanied with clear discussion of validation or how results could be validated.
- Increased funding for establishing domain-specific benchmarks/benchmarking datasets.
- When developing or applying AI tools: Increased emphasis on collaborative efforts between AI- and domain science experts to establish constraints and/or develop *new* architectures based on physics equations; still work to be done here.

Ethical Considerations

General Observations:

- This is a very complex topic; ethical considerations vary widely by domain and type of AI.

Guiding question: what general notes can we share with the broader, community to facilitate ethical use of AI?

- Prepare “living” categorical statements (see NeurIPS ethics statement) that evolve with AI.
- Consider have/nots: Access to “pro” models; whose data is being learned on; role of publishers.
- Consider carbon footprint.
- Guidelines should be provided by sponsoring institutions on how AI tools should/n’t be used in publishing/proposal development/research activities.
- How does one properly credit for another’s work in the context of, e.g., LLM’s?

Final Thoughts

- Responsible AI is a broad umbrella - trust, reproducibility, do no harm, fair use and access to all
- We need to encourage thinking about AI interpretability, and if/when possible, implementing it. Decisions on this should be made at CHE/DMR/PHYS/AST/MATH level as well the program level
- Open datasets will help with reproducibility but easier to do so in domains without industry interest
- Need better incentive systems than "publish in highest journals" to move towards responsible AI including open datasets, well documented code, and careful benchmarking work when possible; NSF should encourage such incentives
- Ethics is complex, and we do not have a clear answer. Even something as simple as "do not plagiarize" is getting tested with use of chatbots and generative AI that have access to published/prior stuff on the internet

Interpretability, robustness, and UQ (EDITED)

General Observations:

- Very clearly a critical area for increased emphasis, but not all themes matter for every domain
- Existing UQ tools not always reliable/adapted to modern models/data; new methods needed
- When presented, new methods should include clear discussion of:
    - What is the domain of applicability and how do you measure it
    - Whether transparency/interpretability is relevant for the application and why/not?
- Recognize distinction between need for **both** new methodological developments **and** better integration of these approaches into applied AI research; collaboration could help
- Proposals should be required to discuss how these topics will be addressed in work; not necessarily needing *developments* in this space
- Funding opportunity language should directly address these points

Coming Up Next…

Today's Schedule (Tuesday)

9:00–10:30 am: Reports by Theme

- Interdisciplinary Research: Opportunities and Challenges
- Interdisciplinary Research: Resources Needed

10:30–11:00 am: Break

11:00 am–12:30 pm: Reports by Theme (cont'd)

- Education & Workforce Development
- Responsible AI

12:30–2:00 pm: Lunch

2:00–5:30 pm: Domain Breakouts

- Astronomical Sciences: Room 804
- Chemistry: Room 812
- Materials Science: Room 803
- Mathematical Sciences: Room 801 North (here)
- Physics: Room 801 South