Joint APS/CNM Workshop 1: Leveraging AI and Large Language Models in Scientific User Facilities

Thursday, May 8, Morning

8:30 – 8:45 Workshop Organizers
Welcome and Opening Remarks

8:45 – 9:25 Luca Rebuffi (Advanced Photon Source, Argonne National Laboratory)
AI-driven Automatic Optimization of Nano-focused Beams and Wavefronts

9:25 – 10:05 Alexander Hexemer (Advanced Light Source, Lawrence Berkeley National Laboratory)
Emerging AI Tools and Workflows for Enhanced Scientific Discovery at User Facilities

10:05 – 10:15 Break

10:15 – 10:55 Vivek Thampy (Stanford Synchrotron Radiation Light Source)
AI-driven Discovery of High-performance Ferroelectric Materials for Energy-efficient Microelectronics

10:55 – 11:35 Kibaek Kim (Mathematics and Computer Science, Argonne National Laboratory)
Foundation Model for BCDI and Ptychographic Images

11:35 – 12:15 Apurva Mehta (SLAC National Accelerator Laboratory)
Towards Digital Twins for Risk-averse Control of Multielement Crystal Optics

12:15 – 1:30 Lunch Break

Thursday, May 8, Afternoon

1:30 – 2:10 Esther Tsai (Center for Functional Nanomaterials, Brookhaven National Laboratory)
Towards AI-embedded X-ray Scattering Experimentation

2:10 – 2:50 Chris Lu (OpenAI)
Towards Using AI for Fully Automated Open-ended Research

2:50 – 3:30 Daniil A. Boiko (Department of Chemical Engineering, Carnegie Mellon University)
LLM Agents in Chemical Sciences: Where Can We Get More Data?

3:30 – 3:40 Break

3:40 – 3:55 Aileen Luo (Advanced Photon Source, Argonne National Laboratory)
DONUT: Physics-aware Machine Learning for Real-time Nanodiffraction Microscopy Analysis

3:55 – 5:30 Xiangyu Yin (Argonne National Laboratory)
Empowering X-ray Science with Foundation Models and Agentic Workflow

Tutorial 1: LLM Assistants for Extracting and Organizing Scientific Data

1. *Overview of tools and techniques for leveraging LLMs to streamline data extraction and organization.*
2. *Hands-on exercises using real-world scientific datasets.*

Break

Tutorial 2: LLM Assistants for Experiment Workflow and Operation

1. *Demonstrating how LLMs can assist in planning, optimizing, and automating experimental workflows.*
2. *Interactive session with practical examples and simulations.*

5:30 Workshop Organizers
Closing Remarks and Networking Opportunity

5:35 Adjourn

AI-driven Automatic Optimization of Nano-focused Beams and Wavefronts

Luca Rebuffi[1]

[1]Advanced Photon Source, Argonne National Laboratory, Lemont, IL 60439

The Advanced Photon Source Upgrade (APS-U) represents a transformative leap in synchrotron radiation facilities, promising unprecedented x-ray brightness and coherence. This advancement opens new frontiers in material science, biology, and nanotechnology by enabling high-resolution imaging and real-time analysis of complex systems. However, to fully harness the potential of these cutting-edge capabilities, the precision and stability of the x-ray beam must be meticulously maintained. The Optics Group at the APS developed and deployed an AI-driven software suite, providing several methods to autonomously align and optimize x-ray nano-focusing optics with remarkable speed and accuracy, ensuring minimal wavefront distortion and maintaining the coherence of the beam. The software not only integrated any radiation detection and characterization technique available at the newly commissioned APS-U beamlines but is equipped to fully manage our real-time, reference-free wavefront sensing technique based on coded-mask, providing autonomous alignment of the optics based on the wavefront properties. By integrating machine learning algorithms and advanced optimization techniques, our system can adaptively respond to changes in beam properties and environmental conditions, providing real-time corrections and enhancing experimental efficiency. This capability not only conserves valuable experiment time but also significantly improves the reliability and reproducibility of scientific results. As the APS-U continues to push the boundaries of what is possible in synchrotron science, AI-driven systems will play a crucial role in realizing its full potential, paving the way for groundbreaking discoveries and innovations.

Emerging AI Tools and Workflows for Enhanced Scientific Discovery at User Facilities

Alexander Hexemer[1], Tanny Chavez[1], Wiebke Köpp[1], and Dylan McReynolds[1]

[1]Advanced Light Source, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Integrating artificial intelligence (AI) and machine learning (ML) is transforming experimental workflows at scientific user facilities. We present a suite of AI-driven tools addressing data analysis and scientific image processing challenges. Our work focuses on three main areas: AI-guided scientific image generation, automated data labeling pipelines, and web-based ML deployment tools. We have developed and fine-tuned AI models for generating domain-specific GISAXS and GIWAXS images, incorporating human-labeled (60%) and experimental (40%) data to ensure scientific accuracy. This approach is complemented by computer vision models that validate the realism of generated images.

Additionally, we have implemented a novel labeling pipeline with interconnected web-based interfaces (Data Clinic, MLCoach, and Label Maker) that enables efficient preparation of ML models for data reduction and classification. These integrated tools have succeeded in pattern recognition for large x-ray scattering datasets, remote RSoXS data analysis, and foundation model fine-tuning. To further enhance accessibility, we have developed a web-based segmentation application for tomography and SEM data that simplifies image segmenting and ML model deployment.

This work was performed and partially supported by the US Department of Energy (DOE), Office of Science, Office of Basic Energy Sciences Data, Artificial Intelligence and Machine Learning at the DOE Scientific User Facilities program under the MLExchange Project (award No. 107514). This research used resources of the Advanced Light Source, which is a DOE Office of Science User facility under contract No. DE-AC02-05CH11231.

AI-driven Discovery of High-performance Ferroelectric Materials for Energy-efficient Microelectronics

Vivek Thampy[1]

[1]Stanford Synchrotron Radiation Lightsource, SLAC National Accelerator Laboratory, Menlo Park, CA 94025

This talk highlights some recent advances in applying machine learning methods to enable closed-loop experimentation at SSRL with a particular focus on high-throughput exploration of functional materials. Synchrotron x-ray metrology provides crucial structural insights during materials synthesis and thermal processing that can be used to optimize the desired material properties. To accelerate the feedback loop, we are developing AI-guided autonomous systems that interpret x-ray scattering data in near real-time to dynamically guide experimental workflows. Specifically, we present an approach combining x-ray metrology with AI-driven closed-loop experimentation to investigate the flash lamp annealing (FLA) of ferroelectric HfO2-ZrO2 (HZO) thin films at SSRL BL17-2. By integrating high speed x-ray structural characterization with autonomous interpretation algorithms and Bayesian active learning, we aim to efficiently navigate the complex processing manifold and identify optimal compositions and annealing conditions that enhance the ferroelectric properties of HZO alloys.

Foundation Model for BCDI and Ptychographic Images

Kibaek Kim[1], Yijiang Li[1], and Mathew Cherukara[2]

[1]Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439
[2]X-ray Science Division, Argonne National Laboratory, Lemont, IL 60439

In this work, we present a machine learning approach that unifies Bragg Coherent Diffraction Imaging (BCDI) and ptychography under a single foundation model. Building on the success of foundation models in other vision domains, we adopt a physics-informed architecture capable of handling the multi-dimensional nature of BCDI and ptychographic data while supporting downstream tasks such as phase retrieval and reconstruction. We share preliminary results from both simulated and real datasets, comparing with existing technique-specific methods with respect to reconstruction accuracy. Moreover, we discuss our federated learning framework to leverage geographically distributed data, demonstrating stable and efficient training without direct data sharing.

This work was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357.

Towards Digital Twins for Risk-averse Control of Multielement Crystal Optics

Apurva Mehta[1], Matt Seaberg[1], and Aashwin Mishra[1]

[1]SLAC National Accelerator Laboratory, Menlo Park, CA 94025

To fully exploit the thousand-fold increase in spectral brightness of Linac Coherent Light Source (LCLS) from the LCLS-II-HE upgrade, significantly more complex monochromators and spectrometers are being proposed. These next-generation optical systems often contain multiple single-crystal optics spread over meters. Finding the optimal beam trajectory and maintaining it for several hours of an experiment requires rapidly finding the global optimum in a coupled multi-parameter manifold. Single crystal optics have a peculiar optimum in the scattering plane, characterized by a sharp rise followed by a gently sloping top. Finding and maintaining crystal-optics systems on top of these sharp, flat-topped optima requires submicron and nanoradian precision in multidimensional manifolds thousand times wider than the optimum in each dimension and presents a significant challenge not only for humans but also for traditional Bayesian optimization (BO) approaches.

In this presentation, we will illustrate these challenges and our modified BO approach for a split and delay (SnD) optical system composed of six crystals and twelve dimensions, which splits a pulsed beam, changes the temporal separation between the branches, and must bring the two beams together spatial with equal intensities. We will show preliminary simulation and experimental validation of rapid optimization and drift correction in eight coupled dimensions that surpasses human operators and traditional BO.

Trusting an AI agent to move two high-powered x-ray laser beams is highly risky for the machine and human researchers in the vicinity, which poses additional challenges. In the concluding section, we will outline our strategy of learning risk-averse recommendations from expert intuition and a longer-term vision of expanding recent successes into digital twins of complex optical systems for more precise, rapid, and trustworthy optimization and drift correction for high repetition rate, ultra-high brightness experiments at X-ray Free Electron Lasers and Diffraction Limited Storage Rings.

Towards AI-embedded X-ray Scattering Experimentation

Esther Tsai[1]

[1]Center for Functional Nanomaterials, Brookhaven National Laboratory, Upton, NY 11973

Synchrotron beamlines are often heavily oversubscribed, making efficient beamtime usage and sustainable operations crucial. X-ray scattering beamlines often utilize custom or *in-situ* setups and thus seamless integration and precise control are vital for efficient and productive beamtime. Recent advancements in artificial intelligence and machine learning, including natural language processing techniques, can enhance beamline operation efficiency to provide better user support as well as improve human-AI collaboration to foster accelerated scientific discovery.

Towards Using AI for Fully Automated Open-ended Research

Chris Lu[1]

[1]OpenAI, San Francisco, CA 94158

One of the grand challenges of artificial general intelligence is developing agents capable of conducting scientific research and discovering new knowledge. While frontier models have already been used as aides to human scientists, e.g. for brainstorming ideas, writing code, or prediction tasks, they still conduct only a small part of the scientific process [2]. I will present our work on the first comprehensive framework for fully automatic scientific discovery, enabling frontier large language models to perform research independently and communicate their findings. We call this system The AI Scientist [1], which generates novel research ideas, writes code, executes experiments, visualizes results, describes its findings by writing a full scientific paper, and then runs a simulated review process for evaluation. In principle, this process can be repeated to iteratively develop ideas in an open-ended fashion, acting like the human scientific community. We demonstrate its capabilities by applying it to machine learning research. Each idea is implemented and developed into a full paper at a cost of less than $15 per paper. To evaluate the generated papers, we design and validate an automated reviewer, which we show achieves near-human performance in evaluating paper scores. The AI Scientist can produce papers that exceed the acceptance threshold at a top machine learning conference as judged by our automated reviewer.

[1] Lu, Chris, et al. "The ai scientist: Towards fully automated open-ended scientific discovery." arXiv preprint arXiv:2408.06292 (2024).
[2] Lu, Chris, et al. "Discovering preference optimization algorithms with and for large language models." Advances in Neural Information Processing Systems 37 (2025): 86528-86573.

LLM Agents in Chemical Sciences: Where Can We Get More Data?

Daniil A. Boiko[1] and Gabe Gomes[1,2,3,4]

[1]Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213
[2]Department of Chemistry, Carnegie Mellon University, Pittsburgh, PA 15213
[3]Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213
[4]Wilton E. Scott Institute for Energy Innovation, Carnegie Mellon University, Pittsburgh, PA 15213

Large language models (LLMs) play an important role in scientific research by enabling AI-driven agents to design, plan, and execute complex experiments autonomously. In chemistry, their potential includes performing reaction optimization and implementing, conducting, and analyzing complex experimental protocols. However, advancements in the underlying foundational models are limited by the availability of high-quality datasets. This talk discusses strategies for expanding chemical reaction datasets and examines how high-throughput mass spectrometry can provide new opportunities for machine learning in synthesis and reaction modeling. By integrating AI-driven experimentation with scalable data generation, it is possible to advance autonomous chemical discovery.

DONUT: Physics-aware Machine Learning for Real-time Nanodiffraction Microscopy Analysis

Aileen Luo[1], Tao Zhou[2], Ming Du[1], Martin V. Holt[2], Andrej Singer[3], and Mathew J. Cherukara[1]

[1]Advanced Photon Source, Argonne National Laboratory, Lemont, IL 60439
[2]Center for Nanoscale Materials, Argonne National Laboratory, Lemont, IL 60439
[3]Department of Materials Science and Engineering, Cornell University, Ithaca, NY 14853

Coherent x-ray scattering techniques are critical for investigating the fundamental structural properties of materials at the nanoscale. Advancements in instrumentation and methodologies have made experiments more accessible, although real-time experimental feedback and analysis remain challenging and prone to artifacts. In particular, scanning x-ray nanodiffraction microscopy, used to spatially resolve structural heterogeneities in extended samples, data is convoluted by the angular range of the divergent diffracted beam. In this work, we propose DONUT (Diffraction with Optics for Nanobeam by Unsupervised Training), a physics-aware neural network for the efficient analysis of nanobeam diffraction data. Our approach incorporates a geometric diffraction model directly into the forward pass of an autoencoder, using automatic differentiation to train the neural network to predict crystal lattice strain and orientation. This method enables real-time analysis without relying on labeled datasets and pre-training. We demonstrate experimentally that DONUT accurately extracts all features within the data far more computationally efficiently than conventional data fitting methods.

Empowering X-ray Science with Foundation Models and Agentic Workflow

Xiangyu Yin[1]

[1]Argonne National Laboratory, Lemont, IL 60439

Foundation models (e.g., LLMs, VLMs) are revolutionizing x-ray science by enabling intuitive human–computer interfaces, fuzzy logic-based automation, and automated human-like insights. Deploying them effectively, however, requires frameworks that integrate advanced AI with existing tools as well as human guidance. Nodeology addresses this through a modular, graph-based architecture that integrates foundation models with established workflows, while maintaining crucial human oversight. Workflows can be shared, adapted, and versioned using Nodeology's template system, fostering collaboration and reproducibility. We have used Nodeology at the Advanced Photon Source to automate ptychography and x-ray fluorescence (XRF). In ptychography, AI-based workflows reduce trial-and-error effort by recommending reconstruction parameters, generating reconstruction code, and analyzing results iteratively. Similarly, XRF copilot program can control instrumentation for real-time optimization. Moreover, these AI-centric workflows can even facilitate autonomous research, from literature insights to exploration of novel algorithms.