AI/ML Mini-Symposium
Room Ballroom BC - Session AIML-ThP
AI/ML for Scientific Discovery Poster Session

AIML-ThP-1 AI Agents for Semiconductor Processing: A New Benchmark for Autonomous Materials Synthesis, ***Angel Yanguas-Gil,*** Argonne National Laboratory

Over the past year there has been an increasing interest in leveraging foundation and large language models to design AI agents that can interact withexperiments to solve materials science and synthesis problems. One of the challenges of this approach is that testing the performance of these agents require access to automated labs. In contrast to benchmarks testing abilities such as knowledge, math skills, or reasoning, there is a lack of benchmarks that can help both design and evaluate agents without the access to dedicated experimental facilities.

In this work, we introduce Semibench, a benchmark to evaluate AI agents' ability to operate and solve synthesis challenges in the context of semiconductor processing. This benchmark introduces two core ideas: first, it introduces virtual tools that simulate the output of real life experiments. This allows us to test an agent's ability to solve a wide range of challenges involving different tool configurations, amount and nature of information that is accessible, and process complexity. Second, it focuses on the concept of microtasks, challenges designed to have a unique solution. This allows us to define quantitative performance metrics for the agent based on how far the proposed solution is to the ground truth. For Semibench, we have focused on three different techniques that are commonly used in the context of microelectronics: atomic layer deposition, sputtering, and reactive ion etching. For each challenge in the benchmark, agents are exposed to a collection of virtual tools and asked to solve specific questions by providing a sequence of synthesis steps. These steps involve selecting the right configurations for each of the tools, such as the precursor channels in the case of ALD, or the targets and power for sputtering, or the etching recipe for RIE.

We have applied this benchmark to agents based on leading large language models, such as OpenAI's o1 an o3 family of reasoning models. The results show that these agents can correctly identify the sequence of steps in a wide range of conditions. However, they struggle when they need to use quantitative data that is not provided explicitly to solve these challenges. These results provide useful information about how to design useful models and their limitations for thin film applications.

AIML-ThP-2 Domain Knowledge + AI for Chemically Accurate Potentials: Application to Diamond Surfaces, ***John Isaac Enriquez,*** Princeton University Plasma Physics Lab*; Yoshitada Morikawa,* Osaka University, Japan*; Igor Kaganovich,* Princeton University Plasma Physics Lab

Machine learning interatomic potentials (MLIPs) are powerful tools for accelerating atomistic simulations, but their reliability depends critically on training set construction. A common strategy is to build universal MLIPs from open-source databases, offering transferability but often sacrificing accuracy particularly in surfaces and interfaces with highly diverse chemical environments. These databases are dominated by equilibrium structures, leaving reaction pathways undersampled, forcing potentials to extrapolate in chemically critical regions—a limitation in catalysis, surface chemistry, and defect dynamics where reactive events dominate. Specialized MLIPs built via active learning can achieve higher accuracy but typically rely on molecular dynamics (MD) and committee models to sample configuration space. Because rare reactions are unlikely to appear within accessible timescales, discovery is left to chance, often requiring long simulations or many iterations. As a result, such MLIPs may fit training data well but fail to capture the most chemically relevant regions.

To address this limitation, we introduce DIAL (Domain-Informed Active Learning), a chemically targeted strategy that augments conventional active learning. Rather than relying solely on MD and uncertainty-driven sampling, DIAL incorporates both established reaction pathways and those identified via nudged elastic band (NEB) calculations. Training datasets are enriched with configurations along these pathways, particularly near transition states. By integrating data-driven active learning with domain knowledge of chemical processes, this approach ensures that the potential is trained on the chemically important regions of configuration space.

Using DIAL, we developed a specialized MLIP for diamond surfaces and interface reactions. The potential enabled large-scale molecular dynamics simulations that reproduced graphitized and oxidized surface morphologies and reaction products, while providing new insights relevant to diamond-based electronics and quantum technologies. In particular, the model captured thermal degradation mechanisms and suppression, facet-dependent oxidative etching, and suggested strategies for controlling surface termination to preserve quantum-relevant color centers. Although demonstrated on diamond, the DIAL framework is general and applicable to other reactive materials systems, including catalysis and battery interfaces.

These results demonstrate how DIAL bridges data-driven methods with domain expertise, highlighting the value of collaboration between materials scientists and AI specialists in advancing the next generation of materials discovery.

AIML-ThP-3 AI Operating System for Accelerating Semiconductor R&D Process Development, ***Suresh Ayyalsamy,*** *Manish Sharma,* Elucida Labs

Advanced plasma etch and deposition process development in semiconductor R&D requires the simultaneous optimization of dozens of interdependent parameters against stringent nanometer-scale metrics. Today, process engineers face a fragmented workflow characterized by siloed data, inefficient experimentation, manual analysis, and limited integration between process settings and physical outcomes. These bottlenecks slow discovery, drive up costs, and hinder knowledge transfer across teams.

We present Elucida Labs, an AI-native operating system designed to transform semiconductor R&D environments across both etch and deposition. Our system enables process teams to reduce experimental burden, shorten learning curves, and converge to target specifications faster. By embedding AI-driven intelligence directly into R&D workflows, Elucida Labs demonstrates how AI can amplify human expertise, accelerate innovation, and reshape the economics of semiconductor process development.

AIML-ThP-4 Physics-Informed Neural Networks for One-Dimensional Capacitively Coupled Plasma Physics Problems, ***Uvini Balasuriya Mudiyanselage,*** *Jesse Jing,* Arizona State University*; Abhishek Verma, Kallol Bera, Shahid Rauf,* Applied Materials Inc.*; Kookjin Lee,* Arizona State University

Physics-Informed Neural Networks (PINNs) offer a flexible framework for solving coupled partial differential equations by embedding physical laws directly into thetraining process. In this work, we develop and evaluate a PINN approach for modeling one-dimensional capacitively coupled plasma (CCP) discharges, governed by electron continuity equation under the drift-diffusion approximation and uniform ion density assumption, coupled with Poisson's equation for self-consistent electrostatic plasma description. The governing equations are non-dimensionalized to improve numerical stability and facilitate learning across disparate physical scales. The model consists of two separate fully connected networks—one for electron density and one for electric potential—augmented with Fourier Feature Mapping to capture multi-scale spatial variations and trained with exact Dirichlet boundary conditions enforced for both electron density and potential. Collocation points are sampled throughout the spatio-temporal domain to compute physics-based residuals directly. Our PINN approach successfully approximates the finite difference method (FDM) solution, achieving an average L^2 relative error of 3.55% for electron density and 3.89% for electric potential over spatio-temporal domain. To address training stiffness and gradient flow issues commonly observed in multi-equation PINNs, we are currently exploring adaptive loss balancing via gradient-based reweighting, as well as Neural Tangent Kernel (NTK)analysis. Preliminary results reveal a significant imbalance in the convergence rates of the two governing equations: the continuity equation loss decreases much faster than that of the Poisson's equation, necessitating disproportionately higher loss weights for the Poisson term to achieve balanced convergence. The model is currently being extended to include ion continuity and momentum conservation equations.

Author Index

Bold page numbers indicate presenter