ME 507: Atomistic Modeling of Materials

Credits	3 Credits
Meeting time	T/Th 09:00–10:20am
Location	1018 Dow
Instructor	Thomas Swinburne, tomswinburne.github.io
Office hours	Th 11:00am–1:00pm, 3003A EECS

Course overview

Atomistic simulation aims to predict material properties from the ground up in two steps:

1. starting from quantum mechanics, build models for the energies of atomic interaction
2. starting from statistical mechanics, put $10-10^8$ atoms together, generate forces, and explore.

Such simulations can discover the atomic mechanisms responsible for many properties we measure in the lab; they are widely used across academia and industry both to understand existing materials and to design new ones. ME 507 introduces the theory and practice of atomistic simulation, with a deliberate focus on the scientific AI/ML tools that are now reshaping the field. At every point we ground theoretical concepts in practical application and give an overview of the research frontier.

You will acquire the theoretical foundations and practical fluency needed to build, run, assess, and extend modern atomistic simulations — from fine-tuning and deploying foundational machine learning models, predicting phase diagrams, investigating the structure and kinetics of crystal defects, and harnessing agentic systems.

Required background

We will review all core topics in class, but you should be comfortable reading scientific Python and willing to engage with math / phyiscs / chemistry topics as they arise. ME majors will have seen all applications in ME 382. Prior exposure to any of the following is helpful but not required:

- **Linear Algebra**: vectors, matrices, rotations, eigenvalues, regularization and iterative methods
- **Probability / Statistics**: random variables, expectation, uncertainty, outliers, Bayes' theorem
- **Data Science / ML**: (linear) regression, classification, neural networks, feature engineering.
- **Statistical Mechanics** partition function, classical/quantum harmonic oscillator, diffusion
- **Materials Science** Crystal structures, phase diagrams, material bonding, segregation
- **Mechanical Engineering** Elasticity, plasticity, dislocations, grain boundaries, fracture
- **Computation** Standard python stack: NumPy/Jupyter/matplotlib and the ASE package for simulation. Bonus points PyTorch, JAX, hardware acceleration, MPI, weak/strong scaling...

Delivery Style

- **In-person attendance is required**; accommodations will be made for valid absences.
- Lectures will be followed by take-home coding assignments which we will go through in class.
- We will host a few guest lectures from researchers expert in AI/ML-driven simulations

Grading

Component	Weight
Final presentation (a short oral presentation on any topic from the course)	10 %
Class attendance	30 %
Assignments	60 %

Honor Code

- AI or human collaboration on assignments is permitted but must be acknowledged.
- Whatever you produce, you are expected to be able to explain each concept every line of code.

Topics

Each topic could occupy the whole semester, so some strong selection is made. I will take a deep dive into machine learning potentials as the field is both mature and currently very active. All applications will have code, and assessment will primarily be coding assignments.

1. **Foundations of atomistic simulation**
   - History, scope, and goals; the ground-up paradigm
   - The fundamental limitations: accuracy, timescale, length scale
   - Structure of a molecular dynamics code; integrators and thermostats
2. **Mathematical and statistical preliminaries**
   - Linear algebra: vectors, tensors, eigendecompositions, SVD
   - Group transformations, representations, and invariance/equivariance
   - Regression, classification, and loss landscapes
   - Bayesian and Post-Bayesian uncertainty quantification
3. **From quantum mechanics to force fields**
   - Density Functional Theory and Jacob's Ladder
   - Computational effort and the scaling wall

4. **Machine learning interatomic potentials I: building the model**
    - Classical force fields and their failure modes
    - Why symmetry matters: invariant and equivariant models
    - Representing atomic environments; short-sightedness
    - Linear feature models and kernels (first-wave MLIPs)
    - Neural network, message-passing, and transformer-based models
5. **Machine learning interatomic potentials II: training the model**
    - Foundation models: architecture, training data, challenges
    - Distillation, fine-tuning, and low-rank adaptation
    - Active learning, extrapolation grade, and data selection
    - Charged systems and long-range electrostatics
    - Electrochemistry and grand-canonical descriptions
6. **Defects and mechanical behavior**
    - Point defects, diffusion, nucleation
    - Extended defects: dislocations, grain boundaries, plasticity
7. **Statistical mechanics for predictive modeling**
    - Ensembles, thermodynamic limits, MaxEnt
    - Free energy I: thermodynamic integration and related methods
    - Free energy II: stratified sampling and acceleration
    - Phase diagrams from simulation
    - Generative AI approaches for statistical mechanics
8. **The timescale problem**
    - Energy landscapes and the metastability problem
    - Transition state theory; weak vs. strong scaling
    - Classical accelerated methods: TAD, ParRep/ParSplice, hyperdynamics, metadynamics
    - ML-accelerated timescale methods
    - Generative AI approaches to overcome the timescale problem
9. **Outlook: autonomous and agentic simulation**
    - Agent-assisted high-throughput workflows