



Faculty Candidate Seminar

March 4th, 2026

10:00-11:15 am – Room: Conner Hall 202



Building AI that Reasons to Accelerate Real-world Discovery

Dr. Zhen Wang

Postdoctoral Fellow – University of California in San Diego

Abstract:

Recent advances in large language models and other foundation models have transformed what AI systems can do on benchmarks, especially in coding and math reasoning. Yet in the real world and scientific practice, the hardest problems are still about discovery, where we need AI systems that can systematically explore hypotheses and solution strategies, respect scientific structure, and work with humans inside high-stakes workflows.

In this talk, Zhen Wang will present a unified agenda built around three pillars for building AI systems that reason to accelerate real-world discovery. First, he will describe how to turn pretrained language models into a general reasoning engine for long-horizon exploration, using a Language-Agent-World (LAW) perspective that treats reasoning as controlled search over hypotheses, strategies, and plans retraining. This framework enables practical applications and procedural discoveries in natural language space, such as discovering effective prompts and solution strategies for new problems. Second, he will introduce how reasoning in the neuro-symbolic space can produce structure-level explanations that scientists can validate, illustrated by the first foundation model on cancer genomics mutations augmented by knowledge graphs. Third, he will show early AI co-scientist systems that embed these ideas into real analysis workflows in domains such as biology and biomedicine, where models collaborate with scientists to reason, run, and refine complex data analyses for discovery.

Together, these directions point toward a coherent path: augmenting foundation models with structured knowledge to build human-centered reasoning systems that are interpretable, exploratory, and genuinely useful for discovery in data-intensive domains.

Biography:

Zhen Wang is currently a Moore Foundation Postdoctoral Fellow, hosted at the Halıcıoğlu Data Science Institute in the University of California, San Diego, and affiliated with MBZUAI and CMU, working with Eric Xing and Zhiting Hu. He earned his PhD in Computer Science at The Ohio State University, working with Huan Sun. His research develops AI reasoning algorithms, agentic systems, and scientific foundation models that combine language models, agent models, world models, and structured knowledge to accelerate scientific discovery across biology, biomedicine, materials science, and social science. He led the first workshop on Language, Agent, and World models (LAW) for reasoning and planning at NeurIPS 2025. His research has been supported by the Gordon and Betty Moore Foundation Fellowship, an OpenAI Research Grant (sole-PI, one of 11 teams worldwide), and recognized with awards, such as Rising Star in Data Science (UChicago), NeurIPS Oral, Best Paper Award (SoCal NLP), Graduate Research Award (OSU), and Amazon Alexa Prize Winner, appearing in venues such as NeurIPS, ICLR, ACL, EMNLP, KDD, Bioinformatics, and Nature.