This April, I joined a call with rare disease organizations across the US to present TxAgent, an AI assistant for drug-related clinical research. I was moved by the people I met, who fought every day for patients dying from rare diseases. Modern drug development is too slow and expensive for rare disease settings. With 95% of rare diseases lacking treatments, we urgently need new methods to accelerate research across *all* diseases. This is why I want to develop general-purpose AI models that use tools and think like scientists, integrating chemistry, biology, and clinical data to generate promising hypotheses. I will graduate with a Statistics bachelor's and concurrent CS master's from Harvard, with significant research experience and coursework in structural biology and chemistry. This prepares me for my PhD, where I aim to build AI models that reason across scales, from molecular interactions to clinical outcomes. In this way, I hope to create AI-based therapeutic research assistants for all biomedical scientists, particularly those working with underserved patient populations.

Early Biomedical Research

Through my early research, I realized that computational methods can address global biomedical challenges. In my first year, I joined the Young-Pearse Lab as a Harvard Program for Research in Science and Engineering (PRISE) fellow to study molecular mechanisms of Alzheimer's. Using differential expression and functional enrichment analyses, I identified SMOC1 as a protein overexpressed in Alzheimer's brains and investigated its cellular pathways. This project showed me that computational tools like enrichment analysis allow me to study any biological system. I decided I want to *build*, and not just use, computational methods, so that I can apply them across different biological systems, diseases, and patient populations.

The AI Scientist: Agentic AI Approaches for Therapeutic Discovery

One of the first computational methods I worked on was an AI scientist, which is an AI agent system built from large language models (LLMs) that uses tools and multi-step reasoning to generate novel hypotheses. Unlike ordinary LLMs, AI scientists replicate human thinking and tool use, producing hypotheses that are interpretable and grounded in evidence. This helps us create more trustworthy AI research assistants. As a Harvard Summer Institute for Biomedical Informatics fellow in the Zitnik Lab, I co-developed an AI scientist, TxAgent, that uses 211 tools and multiple rounds of tool-calling and reasoning to answer biomedical queries.

I helped build, train, and evaluate TxAgent, contributing as second author on the project preprint [1]. Since its release in March, the preprint has been cited 21 times, underscoring its impact on AI scientist research. When developing TxAgent, I drew on my biomedical research experience to incorporate domain-specific tools like the Human Phenotype Ontology database into TxAgent's toolbox. I also created training data for a retrieval-augmented generation (RAG) model that helps TxAgent efficiently identify the right tool for a given query. We first evaluated TxAgent on questions about specific drugs. To test whether the model could reason more generally about therapeutics, I then created a benchmark of questions that provided only drug descriptions, with the names removed. On this more challenging task, TxAgent outperformed leading LLMs like GPT-4o. I also collaborated with 29 rare disease foundations to show that TxAgent can answer questions about a wide range of diseases with greater accuracy and better explanations than its base model. Through this project, I learned that teaching AI to think like humans, use tools, and integrate disciplines makes it useful across diverse diseases and drugs.

AI for Molecular Modeling

To allow TxAgent to think about drug molecules, I want it to not just use textual data, but also understand non-textual modalities like chemical structures. Thus, I am building AI tools for molecular modeling that let human and AI scientists manipulate biomolecules.

To better understand molecular modeling, I took a graduate MIT class where we used AI-based molecular modeling (diffusion models and graph neural networks) to design a novel drug for a rare muscular dystrophy in one semester. I was shocked by the speed of AI-based molecular modeling—these methods could shrink development timelines and costs for rare disease therapeutics. I dove deeper into AI for molecular modeling. I decided to tackle the challenge of designing proteins with a certain biological function because function is notoriously hard to engineer for. I co-developed MUTAGENIC, a method using vector representations of protein function from ESM3, a large language model specialized for protein design. Using ESM3's vector representations, MUTAGENIC mutates amino acids in a protein sequence so the final macromolecule performs a specific function. I published MUTAGENIC as a co-first-author at the 2025 Learning Meaningful Representations of Life workshop at ICLR [2].

Next, as an NSF Rosetta Commons REU fellow in the Gray Lab at Johns Hopkins, I worked on modeling protein-protein complexes. Modeling protein complexes is difficult because some proteins, such as antibodies, have highly flexible binding regions that are difficult for AI models to predict. The AI tool I developed outperformed a state-of-the-art method in identifying biophysically plausible protein-protein complexes, and I presented this first-author work at the 2025 Summer RosettaCon, 2025 Molecular Machine Learning conference, and 2025 Machine Learning for Structural Biology workshop [3]. I started the project with DFMDock, a diffusion model that proposes potential protein-protein complexes but isn't good at ranking these by biophysical plausibility. By fine-tuning DFMDock with the LambdaLoss, a ranking loss function from CS, I showed that DFMDock surpasses a state-of-the-art method on ranking protein-protein complexes. In molecular modeling, I saw how combining AI with structural biology creates new drug design methods. I wanted to teach AI scientists these methods so they could model therapeutic molecules.

Multi-Disciplinary AI Methods for Therapeutic Discovery

I incorporated molecular modeling tools into ToolUniverse, an open-source environment with 600+ tools that lets researchers build their own AI scientist. I then connected ToolUniverse with an LLM to build an AI scientist that can redesign a drug to reduce off-target effects. Thus, I successfully showed the potential of spaces like ToolUniverse that aggregate tools and expertise across disciplines, which allows AI scientists to tackle interdisciplinary challenges like drug discovery. Because of these contributions, I am a co-second author on the ToolUniverse preprint [4]. After the preprint was released last month, it created significant excitement in the field, with the GitHub receiving 700+ stars.

Ultimately, I aim to answer the question: How can we create useful AI research assistants that think like human scientists, integrating evidence across fields like biochemistry, structural biology, and medicine to generate promising hypotheses? Such methods can help identify better therapeutic candidates, reducing timelines and costs and making drug development accessible to underserved groups like rare disease communities. In the future, I aim to lead an interdisciplinary AI research lab that develops tools for therapeutic development that are useful across diseases and patient populations.

In HST, I can work with leaders in AI for multimodal biological modeling and drug discovery, including Dr. Marinka Zitnik for AI scientists and molecular modeling, Dr. Connor Coley for modeling drug molecules and reactions with AI, Dr. Sergey Ovchinnikov for protein design and evolutionary modeling, and Dr. Bonnie Berger for multimodal molecular design.

Moreover, through shadowing, interacting with veteran drug hunters, and working with rare disease collaborators, I have realized the clinical complexity of drug development. Current

AI methods often ignore this complexity, focusing on a single property like binding affinity. This neglects both human pathophysiology and the challenges of real-world patient care, like treatment adherence, adverse events, and healthcare access. How will a drug interact with the patient's comorbidities or other drugs? How will the delivery method—intravenous versus oral—affect treatment adherence? With TxAgent, I helped build an early clinical AI prototype that can answer questions about comorbidities and drug-drug interactions, but through HST, I hope to take the next step in developing clinically-informed AI. Through the Biomedical Sciences Core, I will study human pathophysiology alongside medical students, gaining a wider perspective on how therapeutic interventions affect every part of a patient's body. Through the Clinical Core (HST.201/202), I will care for patients and understand the messy reality of the clinic, which will help me design AI methods that propose hypotheses and drug candidates better aligned with patient needs. In my PhD, I want to build AI methods fully informed by the clinical needs and constraints that cause drugs to fail. With HST's coursework, clinical rotation, and community, I can create patient-driven AI that transforms how we do therapeutic research.

References

[1] Gao, S., **Zhu, R.**, Kong, Z., Noori, A., Su, X., Ginder, C., Tsiligkaridis, T., & Zitnik, M. (2025). TxAgent: An AI agent for therapeutic reasoning across a universe of tools. *arXiv preprint* arXiv:2503.10970.

[2] Pan, R.*, **Zhu, R.***, Lakshman, V., & Qu, F. (2025). Mutagenic: An embedding-based approach to protein masking for functional redesign. *Learning Meaningful Representations of Life Workshop at the International Conference on Learning Representations (ICLR). *denotes equal contribution*

[3] **Zhu, R.**, Xu, D., Chu, L., & Gray, J. (2025). Improving scoring functions for protein–protein docking with LambdaLoss. *Machine Learning in Structural Biology (MLSB) Workshop.*

- Also presented at *Molecular Machine Learning (MoML) Conference* and *Summer RosettaCon*.

[4] Gao, S., **Zhu, R.***, Sui, P.*, Kong, Z.*, Aldogram, S.*, Huang, Y., Noori, A., Shamji, R., Parvataneni, K., Tsiligkaridis, T., & Zitnik, M. (2025). Democratizing AI scientists using ToolUniverse. *arXiv preprint* arXiv:2509.23426. **denotes equal contribution*

- Also presented at *Harvard Dept. of Biomedical Informatics Science Day* **[2nd Place, Best Poster]**.