Research Statement:

As machine learning pipelines become widely adopted in safety-critical and high impact applications, strong reliability measures and uncertainty estimates become exceedingly important. For instance, uncertainty estimates are used actively for domains such as cosmological analysis that depend on Monte Carlo Markov chains. Moreover, they are also used as a way of assessing whether a neural network prediction can be trusted, saving scientists working in domains such as drug discovery from wasting time and money synthesizing molecules that don't have the desired properties that a neural network may have predicted. Beyond these examples, uncertainty quantification also shows a pronounced importance in the data-sparse regimes where equivariant neural networks—a class of neural networks that capture symmetry constraints such as rotations and permutations—tend to excel. While we have come to develop a very mature view of the tradeoffs between equivariant and non-equivariant networks, most of these tradeoffs are stated for traditional classification and regression tasks where the outputs are deterministic. My research has addressed the hitherto unstudied relationship between uncertainty and equivariance.

I have spent a significant portion of my undergraduate work dissecting this relationship with Professor Robin Walters. In particular, I've studied the relationship between equivariance and *calibration error,* which can be loosely defined as the discrepancy between a model's accuracy and purported confidence. My main contribution in this area was a theorem that placed lower and upper bounds on calibration error under various assumptions of symmetry-breaking between the model and the data. I also showed empirically how symmetry mismatch negatively impacts model calibration. **This led to a first author publication submitted to Transactions on Machine Learning Research (TMLR) and a paper at the NeurReps workshop at NeurIPS 2025.** My contributions to the project were especially significant considering that I proposed the main idea, what to prove and how, and how to substantiate the results experimentally.

Eager to demonstrate the applicability of my theory, I joined the Center for Astrophysics Harvard and Smithsonian as a visiting scientist. During this time, I showed that my previous work was applicable to the biomarkers retrieval problem. The biomarkers retrieval problem aims to identify signs of life on exoplanet atmospheres in the form of Carbon based macromolecules. Specifically, chemical properties and uncertainty information is used for Hamiltonian Monte Carlo (HMC) to estimate relative chemical abundances. I contributed to this problem by developing a rotation and translation invariant machine learning model that predicted these chemical properties *and* embedded uncertainty predictions into the prediction using a technique called evidential regression, allowing us to look for Carbon based macromolecules that may not be measured on Earth. I assessed the reliability of the uncertainties in terms of the calibration error bounds derived in my previous work. Once I deemed that the uncertainties were well calibrated, I used them to build the covariance information necessary to perform HMC. A paper is in preparation for ICML**.**

These projects have ignited my interest in uncertainty quantification and symmetry. In studying the relationship between the two, I have learned to balance theoretical investigations in generalization and approximation error of equivariant networks with applied case studies like the biomarkers retrieval problem. In addition to the biomarkers problem, I have also engaged in analogous case studies in cosmology with Professor Jacqueline McCleary. Initially motivated to develop a tool for characterizing the optical behavior of the James Webb Space Telescope, I became frustrated by data scarcity making traditional evaluation statistics inconsistent. Ultimately, I realized the data scarcity issues were exacerbated by clustering approximations used to make the calculation of the evaluation statistics computationally feasible. To quantify this, I proposed an algorithm that calculated the uncertainty due to clustering when computing 2-point and 3-point correlations. **This led to a first author publication in a top tier astrophysics journal, the Open Journal of Astrophysics.**

In summary, my research has established the first relationships between uncertainty and symmetry. Uncertainty quantification is fundamentally about trust that a system in deployment will warn users of potential failure. Similarly, symmetry is about trust that a system adheres to physics. Accordingly, my research vision is to understand the relationship between the two, and my long term goal is to become a professor at a R1 university. I have already tackled problems in this area as a researcher in cosmology, machine learning, and mathematics. I envision a future of machine learning where neural networks can warn us that they do not understand their physical constraints, and furthermore, a future where every correct prediction is accompanied by a certificate of confidence.