Machine Learning (ML) and Artificial Intelligence (AI) are driving transformative changes across various domains, including networking. It is widely assumed that ML/AI-based solutions to complex security or performance-specific problems outperform traditional heuristics and statistical methods. However, this optimism raises a fundamental question: Can our current ML/AI-based solutions be used for high-stakes decision-making in production networks where errors can have serious consequences? Unfortunately, many of these solutions have struggled to fulfill their promises. The primary issues stem from the use of inadequate training data and an overemphasis on narrowly scoped performance metrics (e.g., F1 scores), neglecting other critical aspects (e.g., a model's vulnerability to underspecification issues, such as shortcut learning). The result has been a general reluctance among network operators to deploy ML/AI-based solutions in their networks.

In this talk, I will highlight our efforts to bridge this trust gap by arguing for and developing a novel closed-loop ML workflow that replaces the commonly used standard ML pipeline. Instead of focusing solely on the model's performance and requiring the selection of the "right" data upfront, our newly proposed ML pipeline emphasizes an iterative approach to collecting the "right" training data guided by an in-depth understanding and analysis of the model's decision-making and its (in)ability to generalize. In presenting the building blocks of our novel closed-loop ML pipeline for networking, I will discuss (1) Trustee: A global model explainability tool that helps identify underspecification issues in ML models; (2) netUnicorn: A data-collection platform that simplifies iteratively collecting the "right" data for any given learning problem from diverse network environments; and (3) PINOT: A suite of active and passive data-collection tools that facilitate transforming enterprise networks into scalable data-collection infrastructure. I will conclude the talk by discussing the roadmap of how we can develop a community-wide infrastructure to support this closed-loop ML pipeline for developing generalizable ML/AI models as key ingredients for the future creation of deployment-ready ML/AI artifacts for networking.