I am applying to PhD program at UVA to explore applications of embedded and mobile systems and design secure mobile computing systems. Broadly speaking, I am also interested in designing secure CPS.

During my undergraduate study, **I worked with Dr. Gang Zhou to in designing and developing a power-auditing system against information leakage attack on IoT devices**. A novel covert channel attack on smart bulbs leaks user's private data by encoding them into shift in amplitude and emitting them through the bulb's infrared emission. Instead of monitoring the infrared channel, we proposed to detect the smart bulb's malicious data exfiltration behavior by monitoring its power consumption with a power-auditing sensor. One of my contributions is verifying and explaining that power consumption pattern reflects the behavior of the smart bulb. I used continuous wavelet transformation that transforms the raw time-series power consumption into a 2D image. I learned about deep learning models and designed, tested and tuned a two-dimensional convolutional neural network classifier which classifies the bulb's behavior with above 90% accuracy. The research resulted in a paper that was accepted to SenSys 2022. **This project confirms my interest in applying machine learning in sensor systems.**

For me, the LightAuditor is an effective application of machine learning in sensor system. **The idea inspired me to start my honor thesis research, which aims to use machine learning to infer IoT device activity based on the power consumption data from smart plugs**. The main research questions are 1) what information about smart home device activity can we infer from the data from smart plugs 2) how to design a real-time inference system. To answer 1), I formulated the problem of classification of IoT device type and device activity; I collected time-series power consumption data while performing tasks on smart home devices such as Echo Dot, and I labeled and preprocessed the data. To answer 2), I plan to first train a model offline and implement it on an Arduino board, which acts as an edge device. The research is still ongoing at the time this paragraph is written, and it will last throughout my senior year. I spent major efforts in formulating the problem as an ML classification and conduct experiments to test whether some features allow the model to make good inference. **The most important takeaway is learning how to formulate a problem, make hypothesis, and conduct experiments to test my hypothesis**.

Through these research projects, I noticed security & reliability issue related with sensors and IoT systems. For example, my LightAuditor relies on a deep learning model to make decision (telling whether the IoT device is leaking data). The deep learning model, though trained to be running at 90% accuracy on a test dataset, does not have a known criteria in deciding when to shut down the IoT device. The reliability of the system is questionable. As a student of mathematics, I tend to think rigorously about any concept I have known. **Dr. Feng's work in proving the safety of AI-enabled CPS with formal methods has interested me.** The *Predictive Monitoring with Logic-Calibrated Uncertainty for Cyber-Physical Systems* paper proposes Signal Temporal Logic with Uncertainty as a criteria to monitor sequential predictions. I appreciate how the work formally define a criteria for uncertainty estimation, and then selecting an uncertainty schema that cast deterministic model to a Bayesian RNN. Looking forward, I feel excited to rigorously define the safety criteria of a CPS and evaluate its decision making. I am also eager to learn about potential attacks on CPS, and thus becoming able to design more secure systems.