The field of atmospheric science has entered a period of rapid disruption. Data volumes from numerical simulations and observing platforms have grown far faster than the abilities of forecasters and researchers to analyze them. At the same time, an increasingly diverse group of stakeholders is demanding specialized weather and climate data at ever finer scales to aid in their decision making processes. Data science and machine learning (ML) techniques present a promising path forward for the atmospheric science community after enabling large tech companies like Google and Facebook to provide rapid data-driven services across the globe. In order to meet the demands of current and future stakeholders, the atmospheric science community needs to incorporate more data science and ML into their modeling, observation, and analysis tools. **Therefore, my research has focused on how to integrate ML throughout the atmospheric modeling and decision-making pipeline for improved predictions and increased understanding of physical processes.** My interdisciplinary experience in ML and meteorology has enabled me to develop ML prediction systems for different forms of high impact weather and evaluate them to understand the processes driving each one.

My dissertation research coupled ML and numerical weather prediction (NWP) model output for high impact weather forecasting. I developed a **storm-based day ahead hail size forecasting system** utilizing convection-allowing model output. The forecasting system identifies potential hailstorms within the NWP model and generates hail size predictions based on information about the storm environment. The ML hail forecasts improved on other hail forecasting systems for predicting both severe and significant hail. The ML hail model has produced forecasts in real-time during the 2015-2018 NOAA Hazardous Weather Testbed Experimental Forecast Program and as part of the NCAR MMM convection-allowing ensemble.The details of the ML hail model have been published in *Weather and Forecasting* and *Bulletin of the AMS*. I have participated in a NOAA Joint Technology Transfer Initiative grant, which has been recently renewed, to transition of the ML hail model to operational status at the Storm Prediction Center and investigate real-time deep learning hail forecasts.

I have also evaluated how well different **ML modeling frameworks produce renewable energy forecasts**. The amount of wind and solar power produced varies based on local weather conditions, so electric utilities require accurate forecasts to determine how much non-renewable power to keep in reserve. ML models can both correct biases in NWP model forecasts and convert relevant weather variables to power estimates. Through a collaboration with the NCAR Research Applications Lab, I developed and evaluated a gridded ML solar irradiance forecast system and found that gradient boosted regression consistently produced the best forecasts. The results were published in an article in *Solar Energy*. I have also consulted with Meso, Inc. to identify the best ML model configurations for their wind and solar energy forecasting systems.

During my postdoc, I investigated how **interpretable deep learning methods represent spatial atmospheric data** and whether deep learning provides more predictive information than traditional dimensionality reduction methods, such as principal component analysis (PCA). Deep learning is a rapidly growing area of ML focused on neural networks with many specialized layers that learn structural patterns from spatio-temporal data. Convolutional neural networks (CNNs) can learn optimal patterns from labeled spatial data, and generative adversarial networks (GANs) can generate realistic synthetic samples from low-dimensional representations. In a paper conditionally accepted to Monthly Weather Review, I compared CNNs with PCA and GANs based on their ability to extract spatial features related to hail prediction. Not only did the CNN perform better than the other models, but it was also able to encode different storm morphologies and environmental features, such as wind shear, instability, and confluence, that are associated with the growth of large hail. Deep learning interpretation techniques enabled me to identify neurons that encode certain storm types and then identify storms in my dataset that match those types. This kind of automated semantic analysis could be applied to any scale of feature identification problem in the atmospheric and oceanic sciences, enabling a wide range of possible collaborations.

As a Machine Learning Scientist as NCAR, my research has focused on **incorporating ML methods into sub-grid process parameterizations in earth system models**. ML can improve on existing parameterizations by either 1) emulating a computationally intensive parameterization scheme or 2) generating a new parameterization model from a large dataset of observations or direct simulations of sub-grid processes. For the first approach, I am developing a deep neural network parameterization of the microphysical process of converting cloud droplets to rain drops. The neural network is trained to emulate the output of a bin microphysics scheme but would be incorporated within a bulk scheme. The neural network has demonstrated that it can produce very similar tendencies to the bin scheme, while the original bulk scheme behaves much differently. For the second approach, I am developing machine learning models that can estimate surface layer momentum, sensible heat, and latent heat fluxes from long records of near-surface meteorological tower observations. This machine learning model would replace Monin-Obukhov similarity theory estimates. For both of these projects, I am utilizing model-agnostic interpretation techniques on both the ML models and the original parameterizations to evaluate how well the ML models are matching the behavior of the original parameterizations. As part of a third project, I have investigated how well a GAN can operate as a stochastic parameterization scheme within the Lorenz '96 framework at both weather and climate timescales. Performance depends heavily on the variance and correlation of the noise driving the stochasticity. Sustained success with these machine learning parameterizations could lead to an ongoing stream of projects targeting parameterization emulation of more complex processes in weather and climate models.

In my other current research collaborations, I am investigating ways to **apply deep learning to predicting and understanding a variety of hazardous weather phenomena**. On a project funded by the NOAA Hurricane Forecast Improvement Program, I am working with tropical cyclone experts to build CNNs to identify structures in the inner cores of tropical cyclones that could indicate higher chances of rapid intensification. I am building a lightning initiation and cessation deep learning prediction model based on GOES-16 satellite imagery and lightning data. As a member of a doctoral committee, I am mentoring a student on using deep learning to predict lake-effect snow and understand the synoptic and mesoscale processes driving different events. I am mentoring another student on how to use deep learning to produce downscaled predictions of temperature and precipitation in mountainous terrain.

Looking forward, I want to investigate how our modeling systems can better **synthesize information across scales**. Once information is synthesized, how should we present it to forecasters and stakeholders in order to maximize the capabilities of the human-machine team? How can we incorporate features from the synoptic, meso-, and micro-scales into a comprehensive machine learning prediction system? What long-range dependencies are important for predicting high impact weather ingredients days to weeks ahead? What kinds of interfaces and tools would be most useful for forecasters and decision-makers to interpret these models? How could the forecasters and decision-makers provide interactive feedback to produce more robust and relevant predictions? I will continue to develop interdisciplinary collaborations to answer these questions and train the next generation of atmospheric scientists to make the best use of these data science tools.