ETI Annual Workshop, March 29 – 30, 2022

Self-Supervised Learning Captures Improved Spatial Variation from Satellite Imagery

David Carlson

Duke University

Presenter's Email Address: david.carlson@duke.edu

Abstract:

Convolutional Neural Networks (CNNs) are a promising technique to predict highly localized environmental measurements based on high-resolution satellite imagery. Unfortunately, CNNs typically require large amounts of supervised data to perform well, whereas many remote sensing applications have lots of unsupervised data (e.g., satellite imagery) and relatively sparse supervised data (e.g., measurements from ground sensors). A common approach to mitigate this challenge is to use transfer learning from another visual task to initialize the CNN weights; however, we hypothesize that standard transfer learning strategies would bias the CNN to focus on irrelevant details of the image for our applications. Instead, we develop a novel framework called Spatiotemporal Contrastive Learning (SCL) to pre-train the CNN. We test both regular contrastive learning and SCL on predicting environmental quantities from satellite images in two different cities and compare to CNNs with parameters initialized randomly and by transfer learning. Our results show that regular contrastive learning and our SCL frameworks both manage to better capture spatial variation compared to traditional initialization schemes, and that this performance gap increases as the number of ground sensors decreases, implying that the approach is more valuable in locations with fewer ground sensors. Our work demonstrates that contrastive learning is a powerful pre-training technique to build better spatial maps from remote sensing.