A Multimodal Late-Fusion of Weather Data and Satellite Imagery for Wildfire Prediction

Kim, Hajun (School: Stevenson School)

Early wildfire detection and prediction systems based on machine learning have typically relied on a single type of data, either meteorological data or satellite imagery, limiting their accuracy and applicability across diverse environments. Prior studies demonstrate that this unimodal approach can capture either environmental conditions or spatial patterns, but not both simultaneously. To address this limitation, this research develops a novel late-fusion multimodal framework that integrates the outputs of two independently trained models. A meteorological dataset containing 413 wildfire incidents was used to model environmental predictors. For visual classification, EfficientNetB0 was trained on approximately 6,300 satellite images representing wildfire-prone and non-wildfire regions. Feature embeddings were extracted from the pretrained EfficientNet backbone using Global Average Pooling. These high-level visual representations were then concatenated with scaled meteorological feature vectors and used as input to a fusion classifier. On a paired subset of aligned samples, the multimodal system achieved an AUC of 0.91, substantially outperforming the unimodal EfficientNetB0 model (AUC ˜ 0.60), in which multimodal model also reduced the number of missed fires. Although dataset alignment constraints and limited paired sample size remain challenges, the results indicate that late-fusion integration of heterogeneous data sources can enhance the reliability of early wildfire detection systems.

Awards Won: