

Abstract

Daily rainfall is hard to forecast where most days are dry, a few days are very wet, and almost all of the rain arrives in one season. This dissertation studies day-ahead rainfall forecasting over West Bengal, India, using models that learn from both the temporal and the spatial layout of measuring stations. We first build up a forecasting model in stages, from simple models that look at one station’s past to a graph-based model that links nearby stations, and we identify the strongest deterministic forecaster among them. We then ask whether the choice of input data changes the story, by comparing rain-gauge observations against a reanalysis product on the same stations; the reanalysis turns out to be a smoothed stand-in that misses the sharpest days. Looking closely at the best model, we find it behaves like an estimator of the average rainfall for each day, which is why it cannot place the rare heavy events: almost all of the error comes from a small number of very wet days, and a squared-error objective drives the model to predict too little on exactly those days. We show that self-supervised pretraining and related representation tricks do not move this ceiling. Finally, we change what the model predicts: instead of a single number we predict a full range of possible rainfall values. This probabilistic model is well calibrated and produces useful, reliable warnings for heavy-rain thresholds that the single-number model could never flag, at the modest cost of a slightly worse typical-day error. The contribution is a clear, honest account of where standard spatio-temporal deep learning succeeds and fails for daily rainfall, and a simple change of formulation that recovers useful accuracy on the events that matter most.