Geophysical Research Abstracts
Vol. 21, EGU2019-4410-1, 2019
EGU General Assembly 2019

Evaluating aleatoric and epistemic uncertainties of time series deep learning models for soil moisture predictions

Kuai Fang (1), Chaopeng Shen (1), and Daniel Kifer (2)

(1) Department of Civil and Environmental Engineering, Pennsylvania State University, (2) Department of Computer Science and Engineering, Pennsylvania State University

Recently, recurrent deep networks have been shown to be a powerful tool in harnessing newly available big data for accurate long-term soil moisture predictions. However, to be useful in forecasting applications, they need more than just accuracy - the deep networks must provide model uncertainty estimates as well. Uncertainty for statistical and machine learning models can be classified into two categories: aleatoric and epistemic. Aleatoric uncertainty is the inherent uncertainty due to measurement error and noise in the data - it can be thought of like the error obtainable by a correctly specified model with infinite training data. Epistemic uncertainty results from insufficient training data and/or model misspecification. In this study, we adapt and evaluate an efficient uncertainty estimation framework that was recently proposed in the deep learning community. This method simultaneously estimates both aleatoric and epistemic uncertainty. The aleatoric estimate is provided directly by the deep network by adding an additional output unit. The epistemic component is estimated by Monte Carlo Dropout (MCD), which has two interpretations: the variability observed in an ensemble of deep networks, and an uncertainty estimate derived from a Gaussian Process trained with variational inference. Neither interpretation is exact (both rely on approximations whose errors are not quantified); however, Monte Carlo dropout has shown promise in estimating uncertainty in computer vision tasks for non-recurrent deep networks.

Here, we carefully evaluate the quality of the aleatoric and epistemic uncertainty measures through a series of controlled experiments for soil moisture prediction. Using recurrent deep learning models trained on regional soil moisture time series, we show that the estimated uncertainties are indeed predictive of model errors. The MCD estimate of epistemic uncertainty clearly rises in response to cases that are different from the training data. Thus it can be used to detect when a model is being applied to a dataset for which the training data were not representative. Meanwhile, the aleatoric uncertainty estimate is accurate for representative training data, but, as expected, can be flawed if training data are not representative (a situation that can be detected Monte Carlo Dropout). However, we found that the two uncertainty estimates are not independent, suggesting that the network's aleatoric uncertainty estimate is also correlated with true epistemic uncertainty, e.g., the network learns not to trust its own predictions for cases that differ from the majority of the training data.