Meta-learning for deep neural networks

Alfred Hero

University of Michigan

Deep neural networks (DNN) have attained unprecedented accuracy for signal processing applications including speech, text, and image processing. DNN's are continuously evolving for prediction tasks; they are learning to detect, classify, forecast, and represent ever more complex data. In this talk we address an important meta-learning problem: predict the best achievable classification accuracy from a given training set independently of any DNN model, architecture or training algorithm. Meta-learning deals with learning how difficult is it to learn, given a training dataset. We establish an optimization framework for this meta-learning problem, which we call benchmark learning. Benchmark learning leads to an accurate data-driven predictor of performance of a Bayes optimal classifier without having to construct the classifier and without assuming any parametric model for the data. The resultant predictor can be used to establish whether it is possible to improve classification performance of a specific classifier. It can also be used during DNN training to determine a stopping rule for iterating on minibatches, e.g., using SGD backpropagation training algorithms. The talk will cover relevant background, theory, algorithms, and applications of this meta-learning problem.