


Friday 2/7/2020, 2:15-3:15, P.O.S.T. 126

Naftali Tishby, Hebrew University of Jerusalem

The Information Bottleneck Theory of Deep Learning

In the past several years we have developed a comprehensive theory of large scale learning with Deep Neural Networks (DNN), when optimized with Stochastic Gradient Decent (SGD). The theory is built on three theoretical components: (1) Rethinking the standard (PAC like) distribution independent worst case generalization bounds - turning them to problem dependent typical (in the Information Theory sense) bounds that are independent of the model architecture. (2) The Information Plane theorem: for large scale typical learning the sample-complexity and accuracy tradeoff is characterized by only two numbers: the mutual information that the representation (a layer in the network) maintain on the input patterns, and the mutual information each layer has on the desired output label. The Information Theoretic optimal tradeoff between the encoder and decoder information values is given by the Information Bottleneck (IB) bound for the rule specific input-output distribution. (3) The layers of the DNN reach this optimal bound via standard SGD training, in high (input & layers) dimension.

In this talk, I will discuss two new surprising outcomes of this theory: (1) The computational benefit of the hidden layers, (2) the emerging understanding of the features encoded by each of the layers which follows from the convergence to the IB bound.



Prof. Dr. Naftali Tishby is a professor of Computer Science, and the incumbent of the Ruth and Stan Flinkman Chair for Brain Research at the Edmond and Lily Safra Center for Brain Science (ELSC) at the Hebrew University of Jerusalem. Tishby received his PhD in theoretical physics from the Hebrew University in 1985, was a research staff member at MIT and Bell Labs from 1985 to 1991 and more recently the head of Intel's Collaborative Research Institute for Computational Intelligence (ICRI-CI). His research is at the interface between computer science, statistical physics, and computational neuroscience. He pioneered various applications of statistical physics and information theory to computational learning theory and worked on the foundations of biological information processing and deep learning, and on the formal connections between dynamics and information. Prof. Tishby received several awards including the Landau Prize in Computer Science and the prestigious IBT award in Mathematical Neuroscience.