

MAR16-2015-030028

Abstract for an Invited Paper
for the MAR16 Meeting of
the American Physical Society

Deep Learning, Group representations, and the Information-Bottleneck phase transitions.

NAFTALI TISHBY, Hebrew University

Deep Neural Networks (DNNs) are analyzed via the theoretical framework of the information bottleneck (IB). We first show that any DNN can be quantified by the mutual information between the layers and the input and output variables. Using this representation we can calculate the optimal information theoretic limits of the DNN and obtain finite sample generalization bounds. The advantage of getting closer to the theoretical limit is quantifiable both by the generalization bound and by the network's simplicity. We argue that both the optimal architecture, number of layers and features/connections at each layer, are related to critical points on the information bottleneck tradeoff line, namely, relevant compression of the input layer with respect to the output layer. The hierarchical representations at the layered network naturally correspond to the structural phase transitions along the information curve. An interesting class of solvable DNN's arise by applying this framework to the case of symmetries in the supervised learning task. The case of translation invariance leads to the familiar convolution neural networks. Other symmetry groups yield different types of bifurcation diagrams and network architectures, which correspond to information contained by irreducible representations of the group. These new insights also suggest new sample complexity bounds, architecture design principles (number and widths of layers), and eventually entirely different deep learning algorithms.

Based partly on works with Noga Zaslavsky and Ravid Ziv.