Intrinsic versus extrinsic dimensionality of ground truths

Dennis Elbrächter
ETH Zurich
delbraechter@mins.ee.ethz.ch

Neural networks tend to be particularly successful in scenarios where the data points are high-dimensional. The learned neural network then, by it's nature, represents a function defined for any input in a high-dimensional space. Thus it is natural to also think of the ground truth as a function on a high-dimensional domain when considering it's approximation by neural networks. However, looking at, e.g., image data we can be quite sure that most potential inputs will never actually appear in any relevant task (e.g. in the case of images, only a tiny subset of all possible configurations of pixel values will produce a humanly meaningful picture).
In particular one might argue that interesting ground truths need to be of significantly lower complexity, i.e. intrinsic dimensionality, than the dimensionality of their input allows (even under other assumptions on their simplicity as, e.g., some kind of smoothness). One could now be so bold to go further and conjecture that, in fact, this combination of low intrinsic and high extrinsic dimensionality is a key prior which allows for successful learning by neural networks.
I will present my considerations on a formal notion of intrinsic dimensionality, which is designed to be particularly suitable for the study of approximation by neural networks, and moreover, has an empirical proxy which can be efficiently computed for finite sets of data points.

Joint work with: Helmut Bölcskei, Pavol Harar