Francesco Cagnetta: Learning hierarchical compositionality with deep convolutional networks: insights from a Random Hierarchy Model

While deep learning methods have achieved remarkable success in fields as diverse as language processing and protein structure prediction, understanding their inner workings remains a significant challenge. What are the properties of real data that make deep learning methods so successful and how do these methods learn them? Lacking an answer to these questions, we cannot even estimate the order of magnitude of the sample complexity, that is the number of training examples that a method requires to achieve good performance. In this talk, I will explore our approach to solving this critical challenge, which relies on studying a synthetic classification task---the Random Hierarchy Model---built to mimic the hierarchical and compositional structure of natural data. As the subject of an image (e.g. a dog) consists of features (head, body, limbs), themselves consisting of sub-features (eyes, nose and mouth for the head), each of the classes in our model is represented via a hierarchy of randomly-chosen composition rules, whereby high-level features are represented by a number of semantically equivalent strings of sub-features. Thanks to our specific model choice, we are able to predict the sample complexity in terms of the number of classes, the number of semantically equivalent sub-features and the number of composition rules. Intriguingly, our study reveals that the sample complexity is strictly related to the detectability of the correlations between low-level features of the data and their class. Furthermore, it corresponds to the number of data such that the internal representation of a trained deep network becomes invariant to the exchange of semantically equivalent sub-features. Both the existence of correlations between the low-level features of a datum and its label and the insensitivity of learned representations to aspects of the data irrelevant to the task hold well beyond the context of our synthetic classification task. Therefore, our analysis elucidates a general mechanism of deep learning methods.