Local Propagation in Graphical Neural Networks

The indisputable success of deep learning mostly relies on vector-based representations of the inputs. Yet, many applications deal with non-Euclidean data that typically exhibit a graph structure. Examples come from very different domains, including social networks, molecular graphs in chemistry, and computer vision. In the last couple of years, the extension of neural computation to graphical domains, that was brought to the attention of the scientific community at the end of nineties, has come back to life thanks to a small community of scientists, who have significantly contributed to improve the algorithmic framework and, especially, to show remarkable experimental achievements in different application domains.

In this talk, we begin noticing that, apart from the above mentioned advances, the underlying idea behind the process of learning the weights is still based on an appropriate extension of Backpropagation to graphs. As such, we are still in front of a computational process that has been the source of many debates on its arguable biological plausibility. Then, we propose a novel reformulation of learning in graphical domains that is based on the description of the given graphs and of the neural network by a correspondent set of constraints that must be "parsimoniously satisfied." We propose a Lagrangian framework that gives rise to a biologically plausible local algorithm based on the search for saddle points in the *learning adjoint space* (LAS) composed of weights, neural outputs, and Lagrangian multipliers. This *Local Propagation* algorithm (LP) only involves local updates of the weights. Preliminary experiments are shown to illustrate the features and the performance of this novel local propagation learning algorithm. Interestingly, the learning of LP in the LAS also allow us to circumvent the classic problem of gradient vanishing in deep sequences.