Unfavorable structural properties of the set of neural networks with fixed architecture

Mones Raslan

June 20, 2019

In this talk, we study the structural properties of the set of functions that can be implemented by neural networks with a fixed architecture. As it turns out, this set has many unfavorable properties: It is highly non-convex, except possibly for a few uncommon activation functions. Additionally, the set is not closed with respect to L^p-norms, $0 < p <\infty $, for all frequently used activation functions, and also not closed with respect to the L^∞-norm for all practically-used activation functions except for the (parametric) ReLU. Finally, the function that maps a family of parameters to the function computed by the associated network is not inverse stable for every practically used activation function. Overall, our findings identify potential causes for issues in the optimization of neural networks such as no guaranteed or very slow convergence and the explosion of parameters.