The Geometry of Linear Convolutional Networks

Kathlén Kohn

Royal Institute of Technology, School of Engineering Sciences, Stockholm, Sweden

We study the function space of linear convolutional neural networks (LCNs), i.e., the family of functions represented by the network. The function space is a semi-algebraic subset of the set of linear maps from input space to output space. In contrast, the function space of a fully-connected linear network is an algebraic set. We observe that the functions represented by LCNs can be identified with polynomials that admit certain factorizations, and we use this perspective to describe the impact of the network's architecture on the geometry of the function space. For instance, for LCNs with one-dimensional convolutions having stride one and arbitrary filter sizes, we provide a full description of the boundary of the function space. We further study the optimization of an objective function over such LCNs: We characterize the relations between critical points in function space and in parameter space, and describe dynamical invariants for gradient descent. Overall, our theory predicts that the optimized parameters of an LCN, given generic training data, either yield the global minimum or filters that correspond to polynomials with repeated roots across layers.

This talk is based on joint work with Thomas Merkh, Guido Montúfar, and Matthew Trager.