

Mad Max: Affine Spline Insights into Deep Learning

Richard Baraniuk, Rice University

We build a rigorous bridge between deep networks (DNs) and approximation theory via spline functions and operators. Our key result is that a large class of DNs can be written as a composition of max-affine spline operators (MASOs), which provide a powerful portal through which to view and analyze their inner workings. For instance, conditioned on the input signal, the output of a MASO DN can be written as a simple affine transformation of the input. This implies that a DN constructs a set of signal-dependent, class-specific templates against which the signal is compared via a simple inner product; we explore the links to the classical theory of optimal classification via matched filters and the effects of data memorization. Going further, we propose a simple penalty term that can be added to the cost function of any DN learning algorithm to force the templates to be orthogonal with each other; this leads to significantly improved classification performance and reduced overfitting with no change to the DN architecture. The spline partition of the input signal space that is implicitly induced by a MASO directly links DNs to the theory of vector quantization (VQ) and K-means clustering, which opens up new geometric avenue to study how DNs organize signals in a hierarchical fashion. To validate the utility of the VQ interpretation, we develop and validate a new distance metric for signals and images that quantifies the difference between their VQ encodings.