Topological Convolutional Neural Networks

Ephy R Love
elove4@vols.utk.edu
University of Tennessee
Bredesen Center

Benjamin Filippenko
benfilip@stanford.edu
Stanford University
Department of Mathematics

Vasileios Maroulas
vmaroula@utk.edu
University of Tennessee
Department of Mathematics

Gunnar Carlsson
carlsson@stanford.edu
Stanford University
Department of Mathematics

Abstract

This work introduces the Topological CNN (TCNN), which encompasses several topologically defined convolutional methods. Manifolds with important relationships to the natural image space are used to parameterize image filters which are used as convolutional weights in a TCNN. These manifolds also parameterize slices in layers of a TCNN across which the weights are localized. We show evidence that TCNNs learn faster, on less data, with fewer learned parameters, and with greater generalizability and interpretability than conventional CNNs.

Methods

The TCNN introduces two new types of convolutional layers. The first type, called the Circle One Layer (COL) or Klein One Layer (KOL), parameterizes the sets of input and output slices by a discretization of either the circle or Klein bottle and localizes weights with respect to a metric on the Klein bottle by fixing all weights between distant slices to zero. This significantly reduces the number of trained weights compared to a standard NOL layer. The second type of layer, called the Circle Filters Layer (CF) or Klein Filters Layer (KF), comes instantiated with fixed weights that do not change during training. These weights are given by the embeddings of the primary circle and Klein bottle into the space of image patches discovered in [1]. All of these layers (COL, KOL, CF, and KF) can be viewed as a form of regularization applied to a NOL.

Figure 1: Examples of TCNN Layers

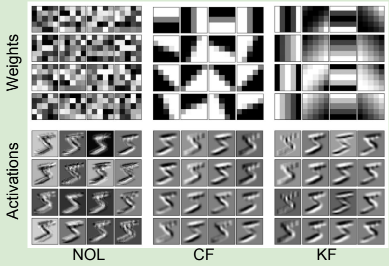

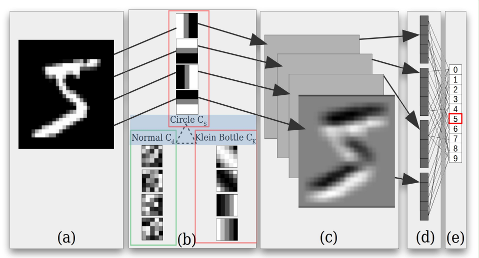

Figure 2: Visual guide to CNN (green rectangle) and TCNN (red rectangles) architectures. With (a) an input tensor, (b) spatially localized weights, (c) feature maps, (d) fully-connected layer(s) and output nodes (e). The TCNN modifies the typical CNN framework by specifying the weights in (b) to lie on a topological manifold such as the circle or Klein bottle.

Results

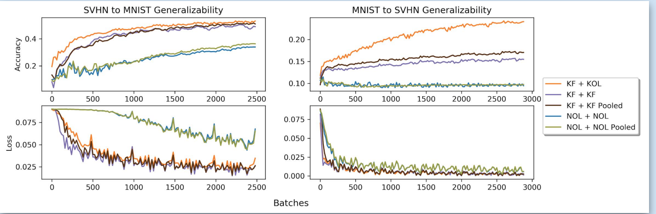

Figure 3: Example of TCNN's improved generalizability between relatively simple (MNIST) and more complex (SVHN) domains.

We compare several TCNN configurations with normal (NOL) CNNs. We present results in the domain of digit classification (Figure 3) and classifying dogs vs. cats. We demonstrate TCNNs' faster training rates (Figure 4). We also show the simplicity of interpreting the weights of a TCNN, as in Figure 1.

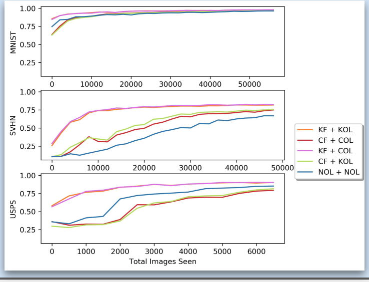

Figure 4: Example of TCNN's improved training speed in classification of 3 datasets of handwritten digits.

Conclusions

We find that TCNNs generalize better than CNNs in trivial domains such as classifying digits, but also more complex domains such as classifying cats and dogs. We find that TCNNs train faster and on less data than CNNs. We also find that topologically defined weights have an easier interpretation than stochastically arrived at weights. We are currently extending our experiments to new frameworks such as ResNets and new data domains such as video.

[1] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the Local Behavior of Spaces of Natural Images. International Journal of Computer Vision, 76(1):1–12, Jan. 2008.