

Abstract

Deep neural network-based models have achieved unprecedented performance levels over many tasks in the traditional supervised setting and scale well with large quantities of data. On the other hand, improving performance in the low-data regime is understudied and the sheer number of parameters versus small amounts of data makes it a challenging task. The problem of learning from a few instances of data is known as few-shot learning. A related problem is that of metric learning over disjoint training and testing classes, where the task is to learn a function that maps a data point to a low-dimensional manifold such that similar points are closer in the low-dimensional space. In this thesis, we propose a couple of deep learning approaches for few-shot learning, and then extend them to the related problem of metric learning over disjoint training and testing classes.

We first argue that a more expressive pairwise similarity matching component is crucial for solving the few-shot learning problem. Towards that end, a network design with a learnable and more expressive similarity objective is proposed by extending the deep residual network. This residual pairwise network approximates a learned metric in the representation space and outperforms previous state-of-the-art on the challenging mini-ImageNet dataset for few-shot learning by getting over 54% accuracy for the 5-way classification task over unseen classes. We also evaluate the generalization behaviour of deep residual networks with a varying number of parameters over classes not observed during training.

Next, since regularization plays a key role in learning with small amounts of data, an additional generator network is proposed by extending the Generative Adversarial Network (GAN) framework for disjoint training and testing classes. This provides a strong regularization by leveraging the generated data samples. The proposed model can generate plausible variations of exemplars over unseen classes and outperforms strong discriminative baselines with L2 regularization for few-shot classification tasks. Finally, the idea of regularizing by adversarial generation is extended to the metric learning setting over disjoint training and testing classes. The proposed model is an efficient alternative to the hard-negative mining scheme necessary when training with triplets in the large margin nearest neighbours setting. The proposed model shows performance and efficiency improvement over models trained without negative mining with the triplet loss.