Incorporating Prior Knowledge in Deep Learning Models

David Carlson

Assistant Professor

Duke University

david.carlson@duke.edu

Abstract:

Deep learning is a powerful tool in various applications, but it often requires a large amount of data for optimal performance. One way to overcome this limitation is by incorporating prior knowledge about the underlying patterns in the data. For example, convolutional neural networks are frequently used in remote sensing, which is subject to strong seasonal effects, and incorporating such prior knowledge would reduce necessary training data. Unfortunately, incorporating prior knowledge in typical neural networks is challenging because they are mainly designed to learn from data, rather than explicitly incorporating prior information. We propose a novel approach to incorporate prior knowledge in deep learning by blending the strengths of deep learning and Gaussian processes. Specifically, we design a composite kernel that combines a kernel implicitly defined by a neural network with a second kernel function that models known properties, such as seasonality. This composite kernel, called Implicit Composite Kernel (ICK), is approximated using a deep network and an efficient mapping based on the Nystrom approximation. We adopt a sample-then-optimize approach to approximate the posterior distribution. Our experiments show that ICK has superior performance and flexibility on both synthetic and real-world data sets. Overall, we believe that the ICK framework can be used to include prior information in neural networks for a variety of applications.