

Saturday, March 23 Session 6: Deep Learning



Graham Taylor

University of Guelph and Vector Institute/
Google Brain Montreal

Title

Introduction to Feedforward Neural Networks, Computational Graphs, and Backpropagation

Abstract

Feedforward Neural Networks are the fundamental building block of deep learning and are at the heart of many sophisticated neural architectures. This tutorial will start with single artificial neurons, and discuss how to build networks with a single hidden layer and then "deep" networks with many hidden layers. It will introduce the concept of computational graphs, which underly the auto-differentiation capabilities of modern tools such as TensorFlow and PyTorch. It will then discuss how backpropagation is applied to computational graphs, and how parameter updates are performed with stochastic gradient descent. Time-permitting we will cover regularization and optimization strategies.

After working through this tutorial and the Jupyter Notebook exercises, students should be able to:

- Understand the background and basic elements of supervised learning with deep, feedforward networks.
- Apply gradient-based training using the backpropagation algorithm and stochastic gradient descent to feedforward architectures.
- Define a computational graph on paper and in code.
- Understand and apply the most common regularization and optimization

strategies.

Biography

Graham Taylor is a Canada Research Chair and Associate Professor of Engineering at the University of Guelph, and a member of the Vector Institute for Artificial Intelligence. His research aims to discover new algorithms and architectures for deep learning: the automatic construction of hierarchical algorithms from high-dimensional, unstructured data. He is especially interested in time series, having applied his work to better understand human and animal behaviour, environmental data (climate or agricultural), audio (music or speech) and financial time series. His work also intersects high performance computing, investigating better ways to leverage hardware accelerators to cope with the challenges of large-scale machine learning. He co-organizes the annual CIFAR Deep Learning Summer School, and has trained more than 50 students and research staff on AI-related projects. In 2016 he was named as one of 18 inaugural CIFAR Azrieli Global Scholars. In 2018 he was honoured as one of Canada's Top 40 under 40. He is currently a Visiting Faculty member at Google Brain, Montreal.

Graham co-founded Kindred, which was featured at number 29 on MIT Technology Review's 2017 list of smartest companies in the world and CB Insights AI 100 list, highlighting the most innovative artificial intelligence companies for 2018. He is the Academic Director of NextAI, a non-profit accelerator and founder development program for AI-focused entrepreneurs.