



Medical Imaging meets Machine Learning

Raquel Urtasun

University of Toronto, Canada

Lecture 1: Introduction to Convolutional Neural Networks

Deep learning algorithms attempt to model high-level abstractions of the data using architectures composed of multiple non-linear transformations. A multiplicity of variants have been proposed and shown to be extremely successful in a wide variety of applications including computer vision, speech recognition as well as natural language processing. In this talk I'll review the fundamentals of deep learning with a particular emphasis on convolutional neural networks (CNNs).

Lecture 2: Learning Deep Structured Models

Many problems in real-world applications involve predicting several random variables that are statistically related. Graphical models have been typically employed to represent and exploit the output dependencies. However, most current learning algorithms assume that the models are log linear in the parameters. In the first part of the talk I'll show a variety of algorithms that can learn arbitrary functions while exploiting the output dependencies, unifying deep learning and graphical models.

Supervised training of deep neural nets typically relies on minimizing cross-entropy. However, in many domains, we are interested in performing well on metrics specific to the application domain. In the second part of the talk I'll show a direct loss minimization approach to train deep neural networks, which provably minimizes the task loss. This is often non-trivial, since these loss functions are neither smooth nor decomposable and thus are not amenable to optimization with standard gradient-based methods. I'll demonstrate the applicability of this general framework in the context of maximizing average precision, a structured loss commonly used to evaluate ranking problems.

Deep learning has become a very popular approach to learn word, sentence and/or image embeddings. Neural embeddings have shown great performance in tasks such as image captioning, machine translation and paraphrasing. In the last part of my talk I'll show how to exploit the partial order structure of the visual semantic hierarchy over words, sentences and images to learn order embeddings. I'll demonstrate the utility of these new representations for hypernym prediction and image-caption retrieval.