TUTORIAL by Adam Coates *(Stanford University, USA)*

Deep Learning for Machine Vision

Abstract. Machine learning algorithms have freed practitioners from many error-prone, hand-engineered components for making decisions in common machine vision tasks such as object recognition. A major source of difficulty, however, is that such learning systems still rely on many hand-built components like sophisticated feature extractors that attempt to identify higher-level patterns in images that typical learning algorithms cannot discover on their own. "Deep learning" and "representation learning" algorithms aim to remove this hurdle by learning higher-level representations automatically from data and have led to recent successes in vision, speech, and language tasks. This tutorial will introduce the basic components of deep learning algorithms and practical techniques for debugging and applying these methods to machine vision problems. The first part of the tutorial will cover neural network models and basic training approaches including error back-propagation and numerical optimization methods, with image classification as a motivating application. The second part will cover additional (sometimes domain-specific) techniques to improve the performance of these algorithms and apply them to other vision tasks including detection and image segmentation. With these tools, audience members will understand how deep learning algorithms work and how they are used in practical applications with sufficient knowledge to complete a hands-on tutorial available on the web. We will conclude with a brief high-level overview of other important topics and results in deep learning research.

Biography. Dr Adam Coates received his PhD in Computer Science from Stanford University in 2012. He is currently a post-doctoral researcher at Stanford and a Visiting Scholar at Indiana University, Bloomington. His research focuses on scaling up machine learning and representation learning algorithms to enable machines to acquire knowledge from unsupervised experience. His interests cover related key topics in computer vision, reinforcement learning, and robotics. He has received Best Student Paper awards from ICML and ICDAR; but his favorite endorsement is still the Instrument Rating on his pilot's license.