



Medical Imaging meets Deep Learning

Trevor Darrell

University of California at Berkeley, USA

Adaptive Deep Learning for Perception and Explanation

Learning of layered or "deep" representations has provided significant advances in computer vision in recent years, but has traditionally been limited to fully supervised settings with very large amounts of training data, where the model lacked interpretability. New results in adversarial adaptive representation learning show how such methods can also excel when learning across modalities and domains, and further can be trained or constrained to provide natural language explanations or multimodal visualizations to their users. I'll present recent long-term recurrent network models that learn cross-modal description and explanation, using implicit and explicit approaches, which can be applied to domains including fine-grained recognition including part-level structures.