



ICVSS 2017

Sicily ~ 09-15 July

International Computer Vision Summer School

Adversarial Perceptual Representation Learning Across Diverse Modalities and Domains

Trevor Darrell

University of California, Berkeley, US

Abstract

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. New results in adversarial adaptive representation learning show how such methods can also excel when learning in sparse/weakly labeled settings across modalities and domains. I'll review state-of-the-art models for fully convolutional pixel-dense segmentation from weakly labeled input, and will discuss new methods for adapting models to new domains with few or no target labels for categories of interest. As time permits, I'll present recent long-term recurrent network models that learn cross-modal description and explanation, visuomotor robotic policies that adapt to new domains, and deep autonomous driving policies that can be learned from heterogeneous large-scale dashcam video datasets.