

UNIVERSITY OF CHICAGO
DEPARTMENT OF COMPUTER SCIENCE
PRESENTS:

Task-specific Recognition by Modeling Visual Tasks



Subhransu Maji

University of Massachusetts Amherst

Abstract:

The AI revolution powered in part by advances in deep learning has led to many successes in the last decade. I'll describe some of our work in this vein that has enabled us to infer detailed properties of objects in images such as their 3D structure or fine-grained category within a taxonomy, as well as study ecological phenomena at an unprecedented scale using data collected from RADAR networks.

Despite these successes the vast majority of important applications remain beyond the scope of current AI systems. One barrier is that existing algorithms lack the ability to learn from limited data. This is a fundamental challenge because real-world data is dynamic, heavy-tailed, where supervision can be hard to acquire. I argue that a principled framework for reasoning about AI problems can enable modular and data-efficient solutions. Towards this end I'll describe our framework for embedding computer vision tasks into a vector space that allows us to learn and reason about their properties. Our approach represents a task as the Fisher information of the parameters of a generic "probe" network. We show that the distance between these vectors correlates with natural metrics over tasks. It is also predictive of transfer, i.e., how much does training a deep network on one task benefit another, and can be used for model recommendation. On a portfolio of hundreds of vision tasks the recommended network using our approach outperforms the current gold standard of fine-tuning an ImageNet pre-trained network, especially when training data is limited. I'll conclude with some of the life-cycle challenges that we need to address to make AI systems widely applicable.

Bio

Subhransu Maji is an Assistant Professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst where he co-directs the Computer Vision Lab. He is also affiliated with the Center of Data Science and AWS AI. Prior to this he was a Research Assistant Professor at TTI Chicago, a philanthropically endowed academic institute in the University of Chicago campus. He obtained his Ph.D. in Computer Science from the University of California at Berkeley in 2011 and B.Tech. in Computer Science and Engineering from IIT Kanpur in 2006. For his work, he has received a Google graduate fellowship, NSF CAREER Award (2018), and a best paper honorable mention at CVPR 2018. He also serves on the editorial board of the International Journal of Computer Vision (IJCV).

Wednesday, March 11, 2020

10:30 am

Crerar 390

Host: Michael Maire