



UNIVERSITY OF CENTRAL FLORIDA
CENTER FOR RESEARCH IN COMPUTER VISION

Dr. Rob Fergus

New York University

“Deconvolutional Networks”

Wednesday, March 20, 2013 · HEC 101 · 2:00PM



ABSTRACT

We present a hierarchical model that learns image decompositions via alternating layers of convolutional sparse coding and a trainable form of pooling. When applied to natural images, the layers of our model capture image information in a variety of forms: low-level edges, mid-level edge junctions, high-level object parts and complete objects. To build our model we rely on a novel inference scheme that ensures each layer reconstructs the input, rather than just the output of the layer directly beneath, as is common with existing hierarchical approaches. This makes it possible to learn multiple layers of representation and we show models with 4 layers, trained on images from the Caltech-101 and 256 datasets. Features extracted from these models, in combination with a standard classifier, outperform SIFT and representations from other feature learning approaches.

Joint work with Matt Zeiler (NYU), David Eigen (NYU) and Graham Taylor (NYU, now at U.Guelph, Canada).

BIOGRAPHY

Rob Fergus is an Assistant Professor of Computer Science at the Courant Institute of Mathematical Sciences, New York University. He received a Masters in Electrical Engineering with Prof. Pietro Perona at Caltech, before completing a PhD with Prof. Andrew Zisserman at the University of Oxford in 2005. Before coming to NYU, he spent two years as a post-doc in the Computer Science and Artificial Intelligence Lab (CSAIL) at MIT, working with Prof. William Freeman. He has received several awards including a CVPR best paper prize (2003), a Sloan Fellowship (2011) and an NSF Career award (2012).