Probabilistic models for understanding images

John Winn
Microsoft Research Ltd, Cambridge, UK

Abstract

Constructing a probabilistic model of an image allows us to use standard Bayesian inference methods to perform tasks like object recognition, segmentation, depth estimation, deblurring and so on. When designing the model, we must balance how accurately it represents the imaging process against how tractable the model is for inference. This tutorial will describe the compromises, pitfalls and challenges of image model design, as illustrated by example models of different aspects of the imaging process. In particular, we will compare approaches which model the image pixels (generative models) to those which model just the labels of interest (discriminative models), as well as looking at recent methods which try to achieve the best of both worlds.

Syllabus: Probabilistic models, graphical models, generative vs. discriminative models, Bayesian inference.