



Seeing what you want to see: Bayes and priors in vision problems

Andrew Fitzgibbon
Microsoft Research Ltd, Cambridge, UK

Abstract

Many vision problems are just too hard: the given data (a stereo pair of images, for example) allow many solutions that are all almost equally likely. Selecting from between the many equivalent explanations of the data is the task of Bayesian inference, and in particular, of the *prior* over what we expect to see. But isn't that just telling the system the answer before starting? How can we trust such an answer? Won't it, at the very least, be biased? In this lecture I will describe some of the main types of prior used in vision problems, and talk about techniques for estimation under the prior. Mostly I will deal with maximum a-posteriori (MAP) estimation, and will show how it is useful in problems of structure from motion, new-view synthesis, and stereo reconstruction.

Syllabus: Bayes, priors, MAP estimation, Markov random fields, image-based rendering, new-view synthesis, stereo, shape from silhouette.