Chapter 1

Introduction

Gaussian processes play an important role in a wide range of practical, large–scale statistical estimation problems. Moderately sized Gaussian inference problems are easily solved using standard linear algebraic techniques. However, for applications arising in such fields as image processing and oceanography, Gaussian priors are commonly used to model the statistical dependencies among tens or hundreds of thousands of random variables. In such cases, the storage and computational requirements of direct, brute–force modeling and inference techniques are intractably large. Instead, families of structured statistical models, and complementary classes of efficient estimation algorithms, must be developed.

In this thesis, we study Gaussian processes whose statistical properties are constrained by an associated graph. The nodes of the graph represent random variables, while edges specify their statistical interrelationships (see Figure 1-1). A wide range of stochastic processes can be compactly expressed by such graphical models [43, 48]. For example, linear state space models [44], Markov random fields [13], and multiscale autoregressive models [20, 30] are all defined using graphical constraints. When designing graphical models, graph structure naturally captures a fundamental tradeoff between the expressiveness and accuracy of the modeled distribution, and the complexity of statistical inference. At one extreme are tree–structured graphs: although they lead to highly efficient estimation algorithms, their modeling power is rather limited. The addition of edges to the graph tends to increase modeling power, but