Markov Random Fields
for Vision and Image Processing

EDITED BY ANDREW BLAKE, PUSHMEET KOHLI, AND CARSTEN ROTHER

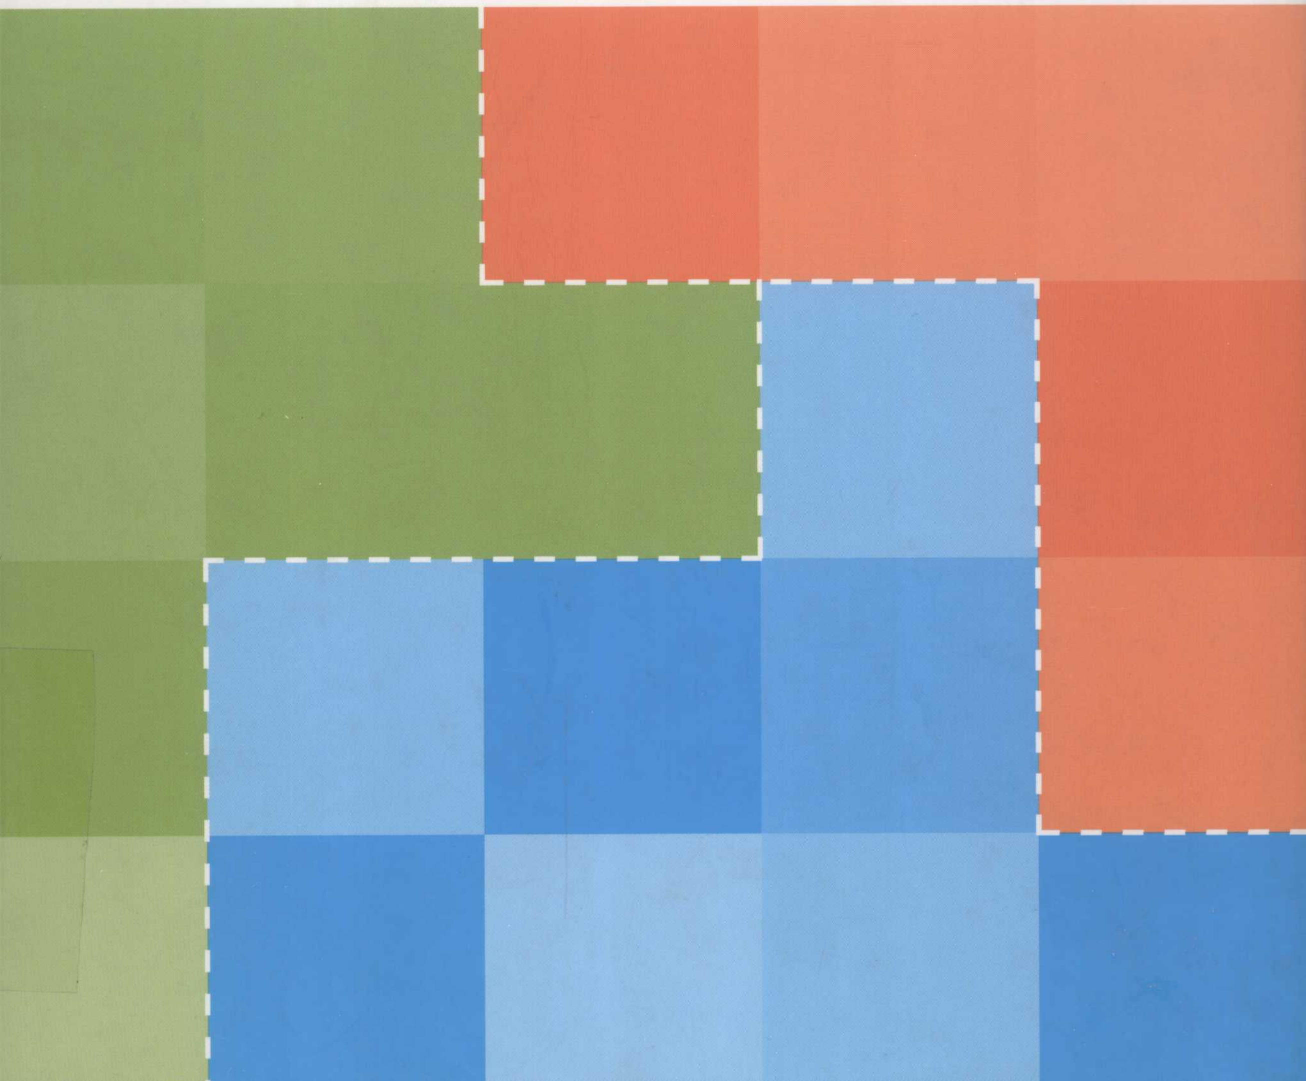

Markov Random Fields for Vision and Image Processing

edited by Andrew Blake, Pushmeet Kohli, and Carsten Rother

常州大学图书馆
藏书章

The MIT Press
Cambridge, Massachusetts
London, England

© 2011 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

For information about special quantity discounts, please email special_sales@mitpress.mit.edu

This book was set in Syntax and Times New Roman by Westchester Books Composition. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Markov random fields for vision and image processing / edited by Andrew Blake, Pushmeet Kohli, and Carsten Rother.
p. cm.
Includes bibliographical references and index.
ISBN 978-0-262-01577-6 (hardcover : alk. paper)
1. Image processing—Mathematics. 2. Computer graphics—Mathematics. 3. Computer vision—Mathematics. 4. Markov random fields. I. Blake, Andrew, 1956– II. Kohli, Pushmeet. III. Rother, Carsten.
TAI637.M337 2011
006.3′70151—dc22
2010046702

10 9 8 7 6 5 4 3 2 1

Markov Random Fields for Vision and Image Processing

Contents

1 Introduction to Markov Random Fields 1
Andrew Blake and Pushmeet Kohli

I Algorithms for Inference of MAP Estimates for MRFs 29

2 Basic Graph Cut Algorithms 31
Yuri Boykov and Vladimir Kolmogorov

3 Optimizing Multilabel MRFs Using Move-Making Algorithms 51
Yuri Boykov, Olga Veksler, and Ramin Zabih

4 Optimizing Multilabel MRFs with Convex and Truncated Convex Priors 65
Hiroshi Ishikawa and Olga Veksler

5 Loopy Belief Propagation, Mean Field Theory, and Bethe Approximations 77
Alan Yuille

6 Linear Programming and Variants of Belief Propagation 95
Yair Weiss, Chen Yanover, and Talya Meltzer

II Applications of MRFs, Including Segmentation 109

7 Interactive Foreground Extraction: Using Graph Cut 111
Carsten Rother, Vladimir Kolmogorov, Yuri Boykov, and Andrew Blake

8 Continuous-Valued MRF for Image Segmentation 127
Dheeraj Singaraju, Leo Grady, Ali Kemal Sinop, and René Vidal

9 Bilayer Segmentation of Video 143
Antonio Criminisi, Geoffrey Cross, Andrew Blake, and Vladimir Kolmogorov

10 MRFs for Superresolution and Texture Synthesis 155
William T. Freeman and Ce Liu

此为试读，需要完整PDF请访问： www.ertongbook.com

11 **A Comparative Study of Energy Minimization Methods for MRFs 167**
Richard Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall F. Tappen, and Carsten Rother

III **Further Topics: Inference, Parameter Learning, and Continuous Models 183**

12 **Convex Relaxation Techniques for Segmentation, Stereo, and Multiview Reconstruction 185**
Daniel Cremers, Thomas Pock, Kalin Kolev, and Antonin Chambolle

13 **Learning Parameters in Continuous-Valued Markov Random Fields 201**
Marshall F. Tappen

14 **Message Passing with Continuous Latent Variables 215**
Michael Isard

15 **Learning Large-Margin Random Fields Using Graph Cuts 233**
Martin Szummer, Pushmeet Kohli, and Derek Hoiem

16 **Analyzing Convex Relaxations for MAP Estimation 249**
M. Pawan Kumar, Vladimir Kolmogorov, and Philip H. S. Torr

17 **MAP Inference by Fast Primal-Dual Linear Programming 263**
Nikos Komodakis

18 **Fusion-Move Optimization for MRFs with an Extensive Label Space 281**
Victor Lempitsky, Carsten Rother, Stefan Roth, and Andrew Blake

IV **Higher-Order MRFs and Global Constraints 295**

19 **Field of Experts 297**
Stefan Roth and Michael J. Black

20 **Enforcing Label Consistency Using Higher-Order Potentials 311**
Pushmeet Kohli, Lubor Ladicky, and Philip H. S. Torr

21 **Exact Optimization for Markov Random Fields with Nonlocal Parameters 329**
Victor Lempitsky, Andrew Blake, and Carsten Rother

22 **Graph Cut-Based Image Segmentation with Connectivity Priors 347**
Sara Vicente, Vladimir Kolmogorov, and Carsten Rother

V **Advanced Applications of MRFs 363**

23 **Symmetric Stereo Matching for Occlusion Handling 365**
Jian Sun, Yin Li, Sing Bing Kang, and Heung-Yeung Shum

24 Steerable Random Fields for Image Restoration 377
Stefan Roth and Michael J. Black

25 Markov Random Fields for Object Detection 389
John Winn and Jamie Shotton

26 SIFT Flow: Dense Correspondence across Scenes and Its Applications 405
Ce Liu, Jenny Yuen, Antonio Torralba, and William T. Freeman

27 Unwrap Mosaics: A Model for Deformable Surfaces in Video 419
Alex Rav-Acha, Pushmeet Kohli, Carsten Rother, and Andrew Fitzgibbon

Bibliography 433
Contributors 457
Index 459

1

Introduction to Markov Random Fields

Andrew Blake and Pushmeet Kohli

This book sets out to demonstrate the power of the Markov random field (MRF) in vision. It treats the MRF both as a tool for modeling image data and, coupled with a set of recently developed algorithms, as a means of making inferences about images. The inferences concern underlying image and scene structure to solve problems such as image reconstruction, image segmentation, 3D vision, and object labeling. This chapter is designed to present some of the main concepts used in MRFs, both as a taster and as a gateway to the more detailed chapters that follow, as well as a stand-alone introduction to MRFs.

The unifying ideas in using MRFs for vision are the following:

- Images are dissected into an assembly of nodes that may correspond to pixels or agglomerations of pixels.
- Hidden variables associated with the nodes are introduced into a model designed to "explain" the values (colors) of all the pixels.
- A joint probabilistic model is built over the pixel values and the hidden variables.
- The direct statistical dependencies between hidden variables are expressed by explicitly grouping hidden variables; these groups are often pairs depicted as edges in a graph.

These properties of MRFs are illustrated in figure 1.1. The graphs corresponding to such MRF problems are predominantly gridlike, but may also be irregular, as in figure 1.1(c). Exactly how graph connectivity is interpreted in terms of probabilistic conditional dependency is discussed a little later.

The notation for image graphs is that the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consists of vertices $\mathcal{V} = (1, 2, \ldots, i, \ldots, N)$ corresponding, for example, to the pixels of the image, and a set of edges $\mathcal{E}$ where a typical edge is (i, j), $i, j \in \mathcal{V}$, and edges are considered to be undirected, so that (i, j) and (j, i) refer to the same edge. In the superpixel graph of figure 1.1), the nodes are superpixels, and a pair of superpixels forms an edge in $\mathcal{E}$ if the two superpixels share a common boundary.

The motivation for constructing such a graph is to connect the hidden variables associated with the nodes. For example, for the task of segmenting an image into foreground and background, each node i (pixel or superpixel) has an associated random variable X_i that

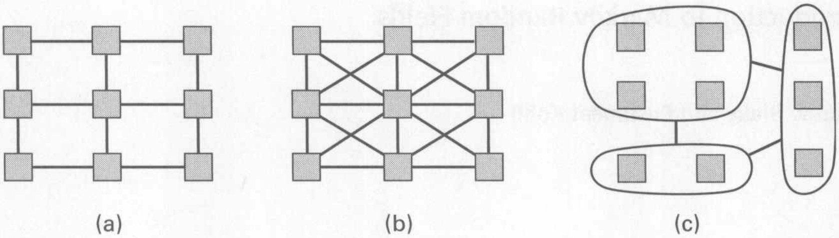

Figure 1.1
Graphs for Markov models in vision. (a) Simple 4-connected grid of image pixels. (b) Grids with greater connectivity can be useful—for example, to achieve better geometrical detail (see discussion later)—as here with the 8-connected pixel grid. (c) Irregular grids are also useful. Here a more compact graph is constructed in which the nodes are superpixels—clusters of adjacent pixels with similar colors.

may take the value 0 or 1, corresponding to foreground or background, respectively. In order to represent the tendency of matter to be coherent, neighboring sites are likely to have the same label. So where $(i, j) \in \mathcal{E}$, some kind of probabilistic bias needs to be associated with the edge (i, j) such that X_i and X_j tend to have the same label—both 0 or both 1. In fact, any pixels that are nearby, not merely adjacent, are likely to have the same label. On the other hand, explicitly linking all the pixels in a typical image, whose foreground/background labels have correlations, would lead to a densely connected graph. That in turn would result in computationally expensive algorithms. Markov models *explicitly* represent only the associations between relatively few pairs of pixels—those pixels that are defined as neighbors because of sharing an edge in $\mathcal{E}$. The great attraction of Markov Models is that they leverage a *knock-on effect*—that explicit short-range linkages give rise to implied long-range correlations. Thus correlations over long ranges, on the order of the diameters of typical objects, can be obtained without undue computational cost. The goal of this chapter is to investigate probabilistic models that exploit this powerful Markov property.

1.1 Markov Chains: The Simplest Markov Models

In a Markov chain a sequence of random variables $\mathbf{X} = (X_1, X_2, \ldots)$ has a joint distribution specified by the conditionals $P(X_i \mid X_{i-1}, X_{i-2}, \ldots, X_1)$. The classic tutorial example [381, sec. 6.2] is the weather, so that $X_i \in \mathcal{L} = \{sunny, rainy\}$. The weather on day i can be influenced by the weather many days previous, but in the simplest form of Markov chain, the dependence of today's weather is linked explicitly only to yesterday's weather. It is also linked *implicitly*, as a knock-on effect, to all previous days. This is a *first-order* Markov assumption, that

$$P(X_i \mid X_{i-1}, X_{i-2}, \ldots, X_1) = P(X_i \mid X_{i-1}). \tag{1.1}$$

This is illustrated in figure 1.2. The set of conditional probabilities $P(X_i \mid X_{i-1})$ is in fact a 2×2 matrix. For example:

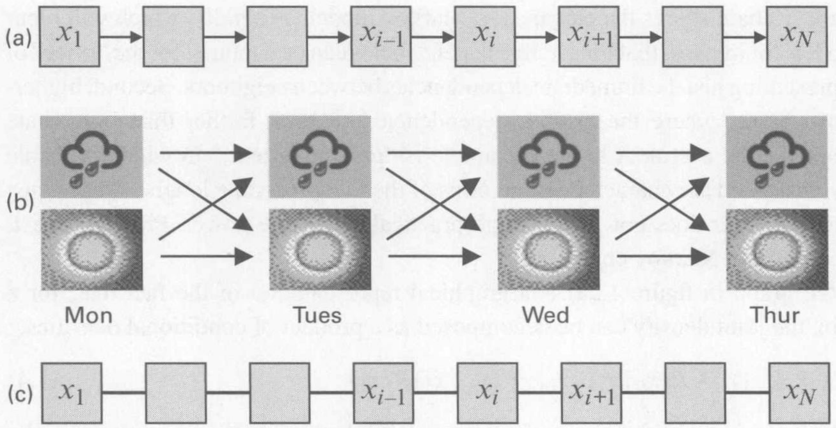

Figure 1.2
A simple first-order Markov chain for weather forecasting. (a) A directed graph is used to represent the conditional dependencies of a Markov chain. (b) In more detail, the state transition diagram completely specifies the probabilistic process of the evolving weather states. (c) A Markov chain can alternatively be expressed as an undirected graphical model; see text for details.

		Yesterday (X_{i-1})	
		Rain	Sun
Today (X_i)	Rain	0.4	0.8
	Sun	0.6	0.2

An interesting and commonly used special case is the *stationary* Markov chain, in which the matrix

$$M_i(x, x') = P(X_i = x \mid X_{i-1} = x') \tag{1.2}$$

is independent of time i, so that $M_i(.,.) = M_{i-1}(.,.)$. In the weather example this corresponds to the assumption that the statistical dependency of weather is a fixed relationship, the same on any day.

We will not dwell on the simple example of the Markov chain, but a few comments may be useful. First, the first-order *explicit* structure implicitly carries longer-range dependencies, too. For instance, the conditional dependency across three successive days is obtained by multiplying together the matrices for two successive pairs of days:

$$P(X_i = x \mid X_{i-2} = x'') = \sum_{x' \in \mathcal{L}} M_i(x, x') M_{i-1}(x', x''). \tag{1.3}$$

Thus the Markov chain shares the elegance of Markov models generally, which will recur later with models for images, that long-range dependencies can be captured for the "price" of explicitly representing just the immediate dependencies between neighbors. Second, higher-order Markov chains, where the *explicit* dependencies go back farther than immediate neighbors, can also be useful. A famous example is "predictive text," in which probable letters in a typical word are characterized in terms of the two preceding letters—taking just the one preceding letter does not give enough practical predictive power. Predictive text, then, is a second-order Markov chain.

The directed graph in figure 1.2a) is a graphical representation of the fact that, for a Markov chain, the joint density can be decomposed as a product of conditional densities:

$$P(\mathbf{x}) = P(x_N \mid x_{N-1}) \ldots P(x_i \mid x_{i-1}) \ldots P(x_2 \mid x_1) P(x_1), \tag{1.4}$$

where for simplicity, in a popular abuse of notation, $P(\mathbf{x})$ denotes $P(\mathbf{X} = \mathbf{x})$ and, similarly, $P(x_i \mid x_{i-1})$ denotes $P(X_i = x_i \mid X_{i-1} = x_{i-1})$. This convention is used frequently throughout the book. An alternative formalism that is commonly used is the *undirected* graphical model. Markov chains can also be represented in this way (figure 1.2c), corresponding to a factorized decomposition:

$$P(\mathbf{x}) = \Phi_{N,N-1}(x_N, x_{N-1}) \ldots \Phi_{i,i-1}(x_i, x_{i-1}) \ldots \Phi_{2,1}(x_2, x_1), \tag{1.5}$$

where $\Phi_{i,i-1}$ is a factor of the joint density. It is easy to see, in this simple case of the Markov chain, how the directed form (1.4) can be reexpressed in the undirected form (1.5). However, it is not the case in general, and in particular in 2D images, that models expressed in one form can easily be expressed in the other. Many of the probabilistic models used in computer vision are most naturally expressed using the undirected formalism, so it is the undirected graphical models that dominate in this book. For details on directed graphical models see [216, 46].

1.2 The Hidden Markov Model (HMM)

Markov models are particularly useful as prior models for state variables X_i that are to be inferred from a corresponding set of measurements or observations $z = (z_1, z_2, \ldots, z_i, \ldots, z_N)$. The observations z are themselves considered to be instantiations of a random variable Z representing the full space of observations that can arise. This is the classical situation in speech analysis [381, sec. 6.2], where z_i represents the spectral content of a fragment of an audio signal, and X_i represents a state in the time course of a particular word or phoneme. It leads naturally to an inference problem in which the *posterior* distribution for the possible states X, given the observations z, is computed via Bayes's formula as

$$P(\mathbf{X} = \mathbf{x} \mid \mathbf{Z} = \mathbf{z}) \propto P(\mathbf{Z} = \mathbf{z} \mid \mathbf{X} = \mathbf{x}) P(\mathbf{X} = \mathbf{x}). \tag{1.6}$$

Here $P(\mathbf{X} = \mathbf{x})$ is the prior distribution over states—that is, what is known about states $\mathbf{X}$ in the absence of any observations. As before, (1.6) is abbreviated, for convenience, to

此为试读，需要完整PDF请访问：www.ertongbook.com

$$P(\mathbf{x} \mid \mathbf{z}) \propto P(\mathbf{z} \mid \mathbf{x}) P(\mathbf{x}). \tag{1.7}$$

The omitted constant of proportionality would be fixed to ensure that $\sum_{\mathbf{x}} P(\mathbf{x} \mid \mathbf{z}) = 1$. Often multiple models are considered simultaneously, and in that case this is denoted

$$P(\mathbf{x} \mid \mathbf{z}, \omega) \propto P(\mathbf{z} \mid \mathbf{x}, \omega) P(\mathbf{x} \mid \omega), \tag{1.8}$$

where the model parameters $\omega \in \Omega$ may determine the prior model or the observation model or both. The constant of proportionality in this relation would of course depend on $\mathbf{z}$ and on ω.

The prior of an HMM is itself represented as a Markov chain, which in the first-order case was decomposed as a product of conditional distributions (1.4). The term $P(\mathbf{z} \mid \mathbf{x})$ is the *likelihood* of the observations, which is essentially a measure of the quality of the measurements. The more precise and unambiguous the measuring instrument, the more the likelihood will be compressed into a single, narrow peak. This captures the fact that a more precise instrument produces more consistent responses $\mathbf{z}$, under a given condition represented by the state $\mathbf{X} = \mathbf{x}$. It is often assumed—and this is true of the models used in many of the chapters of this book—that observations are independent across sites. The observation at site i depends only on the corresponding state. In other words:

$$P(\mathbf{z} \mid \mathbf{x}) = P(z_N \mid x_N) P(z_{N-1} \mid x_{N-1}) \ldots P(z_1 \mid x_1). \tag{1.9}$$

The directed graphical model for the conditional dependencies of such a first-order HMM is given in figure 1.3a). The figure captures the conditional dependencies both of the underlying

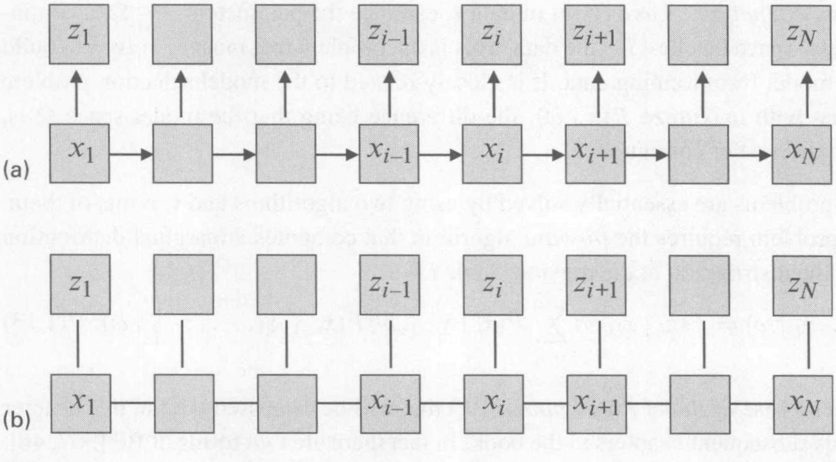

Figure 1.3
A first-order hidden Markov model (HMM). (a) A directed graph is used to represent the dependencies of a first-order HMM, with its Markov chain prior, and a set of independently uncertain observations. (b) Alternatively the HMM can be represented as an undirected graphical model (see text).

Markov chain and of the independence of the observations. Alternatively, an HMM can be expressed as an undirected graphical model, as depicted in figure 1.3(b), in which the prior is decomposed as in (1.5), and the likelihood is

$$P(\mathbf{z} \mid \mathbf{x}) = \Phi_N(x_N)\Phi_{N-1}(x_{N-1}) \dots \Phi_1(x_1), \tag{1.10}$$

where trivially $\Phi_i(x_i) = P(z_i \mid x_i)$.

Discrete HMMs, with a finite label set $\mathcal{L}$, are largely tractable. Rabiner and Juang [382] set out three canonical problems for HMMs, and algorithms to solve them. The problems are the following:

1. *Evaluating the observation probability* $P(\mathbf{z} \mid \omega)$ In this problem there is no explicit state dependence, because it has been "marginalized" out by summation over states:

$$P(\mathbf{z} \mid \omega) = \sum_{\mathbf{x} \in \mathcal{L}^N} P(\mathbf{z} \mid \mathbf{x}, \omega) P(\mathbf{x} \mid \omega). \tag{1.11}$$

The main application of this evaluation is to determine which of a set of known models fits the data best:

$$\max_{\omega \in \Omega} P(\mathbf{z} \mid \omega). \tag{1.12}$$

The quantity $P(\mathbf{z} \mid \omega)$ is also known as the *evidence* [328] for the model ω from the data $\mathbf{z}$.

2. *MAP estimation* Given a model ω and a set of data $\mathbf{z}$, estimate the most probable (maximum a posteriori) sequence $\mathbf{x}$ of states as the mode of the posterior distribution (1.8).

3. *Parameter estimation* Given a set of data $\mathbf{z}$, estimate the parameters $\omega \in \Omega$, a continuous parameter space that best fits the data. This is the problem that must be solved to build a particular model from training data. It is closely related to the model selection problem above, in that both maximize $P(\mathbf{z} \mid \omega)$, the difference being that the model space Ω is, respectively, discrete or continuous.

These three problems are essentially solved by using two algorithms and variants of them.

The first problem requires the *forward* algorithm that computes a marginal distribution node i from the distribution at the previous node $i - 1$:

$$P(x_i, z_1, \dots, z_i \mid \omega) = P(z_i \mid x_i, \omega) \sum_{x_{i-1}} P(x_i \mid x_{i-1}, \omega) P(x_{i-1}, z_1, \dots, z_{i-1} \mid \omega). \tag{1.13}$$

This is a special case of *Belief Propagation* (BP) that will be discussed later in this chapter and in various subsequent chapters in the book. In fact there are two forms of BP [367, 46], and this one is an example of *sum-product* BP. (The name derives from the summation and product steps in (1.13).) The other form is described shortly. In the case of the HMM, where the underlying prior model is simply a Markov chain, sum-product belief propagation is quite straightforward and is an exact algorithm for computing the marginal posteriors. After one complete forward pass, the final marginal distribution is $P(x_N, \mathbf{z} \mid \omega)$, and so finally

$$P(\mathbf{z} \mid \omega) = \sum_{x_N} P(x_N, \mathbf{z} \mid \omega) \tag{1.14}$$

can be computed as the evidence for a known model ω that solves problem 1 above. The forward pass (1.13) constitutes half of BP, the remaining part being a backward pass that recurs from node N back to node 1 (details omitted here, but see [382]). Using the forward and backward passes together, the full set of marginal posterior distributions

$$P(x_i \mid \mathbf{z}, \omega), \; i = 1, \ldots, N \tag{1.15}$$

can be computed. This is required for problem 3 above, in order to compute the expected values of the sufficient statistics that are needed to estimate the parameters ω by *expectation maximization* [121]—also known in the speech analysis literature as the Baum–Welch method [381].

The second algorithm is the Viterbi algorithm, a dynamic programming optimization algorithm applied to the state sequence $\mathbf{x}$. It is also equivalent to a special case of max-product belief propagation, which also is mentioned quite frequently in the book. The aim is to solve the second problem above, computing the MAP estimate of the state vector as

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} P(\mathbf{x} \mid \mathbf{z}, \omega) \tag{1.16}$$

via a forward recursion:

$$P_i(x_i) = P(z_i \mid x_i, \omega) \max_{x_{i-1}} P(x_i \mid x_{i-1}, \omega) P_{i-1}(x_{i-1}) \tag{1.17}$$

where P_i is defined by

$$P_i(x_i) \equiv \max_{x_1, \ldots, x_{i-1}} P(x_1, \ldots, x_i, z_1, \ldots, z_i \mid \omega). \tag{1.18}$$

Each forward step of the recursion can be viewed as a message-passing operation. In step $i-1$ of the computation, node $i-1$ sends the message $P_i(x_i)$ to node i.

After the forward steps are complete, the final component $\hat{x}_N$ of the MAP solution $\hat{x}$ can be computed as

$$\hat{x}_N = \arg\max_{x_N} P_n(x_N). \tag{1.19}$$

This is followed by a backward recursion

$$\hat{x}_{i-1} = \arg\max_{x_{i-1}} P(\hat{x}_i \mid x_{i-1}, \omega)\, P_{i-1}(x_{i-1}), \tag{1.20}$$

after which all components of $\hat{x}$ are determined.

The purpose of the discussion in this section has been largely to explain the nature of hidden variables in simple Markov models, as a precursor to later discussion of hidden variables in the more complex, two-dimensional kinds of models that are used in vision. However, even in vision the discrete HMM structure has some direct applications. It has proved useful for representing temporal problems that are somewhat analogous to speech analysis, but in which the audio input is replaced by a time sequence of visual features. Well-known examples include the recognition of American Sign Language [449] and the recognition of hand gestures for command and control [522]. This book deals mainly with discrete Markov models—that is, ones in which the states of each X_i belong to a finite set $\mathcal{L}$. However, in vision by far the greater application of timelike HMMs employs continuous state-space to represent position, attitude, and shape in visual tracking [333, 477, 50, 253, 209, 124]. In such continuous settings the HMM becomes a classical or nonclassical form of Kalman filter. Both exact solutions to the estimation problems that arise, and efficient approximate solutions, are much studied, but are outside the scope of this book.

1.3 Markov Models on Trees

In the following section 1.4, Markov Random Fields (MRFs) are defined as probabilistic models over undirected graphs. On the way there, we now consider undirected models on trees as intermediate in complexity between the linear graphs—chains and HMMs—of section 1.2, and graphs of unrestricted connectivity. Clearly the HMM graph (figure 1.3b) is a special case of an undirected model on a tree. Trees appear to be of intermediate complexity but, perhaps surprisingly, turn out to be closer to HMMs, in that inference can be performed exactly. The Viterbi and forward-backward algorithms for HMMs generalize to two different kinds of message passing on trees. However, once the nodes on two leaves of a tree are coalesced into a single leaf—for example, leaves **b** and **d** in figure 1.4—a circuit may be formed in the resulting graph, and message-passing algorithms are no longer an exact solution to the problems of inference.

As with Markov chains and HMMs, in undirected trees the topological structure conveys two aspects of the underlying probabilistic model. First are the conditional independence properties, that:

$$P(x_i \mid \{x_j,\ j \neq i\}) = P(x_i \mid \{x_j,\ (i,j) \in \mathcal{E}\}). \tag{1.21}$$

The set $B_i = \{j :\ (i,j) \in \mathcal{E}\}$ is known as the Markov blanket of i, its neighbors in the tree (or generally graph) $\mathcal{G}$. The second aspect is the decomposition of the joint distribution, the generalization of (1.5) and (1.10). How can a distribution with this independence property be constructed? The answer is, as a distribution that is factorized over the edges of the tree:

$$P(\mathbf{x}) = \prod_{(i,j)\in\mathcal{E}} F_{i,j}(x_i, x_j). \tag{1.22}$$

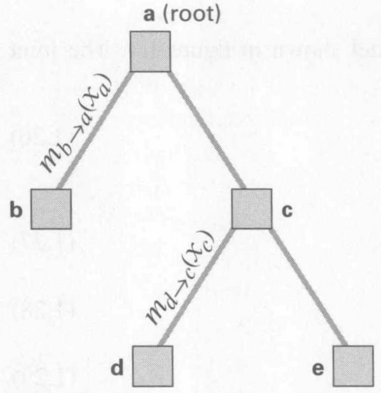

Figure 1.4
Message passing for MAP inference in tree structured graphical models. A graphical model containing five nodes (a, b, c, d, and e) connected in a tree structure.

1.3.1 Inference on Trees: Belief Propagation (BP)

The message-passing formulation of the Viterbi algorithm can be generalized to find marginal distributions over individual variables and the MAP estimate in a tree structured graphical model. The resulting algorithm is known as belief propagation [367] and has two variants: *max-product*, for computing the MAP solution, and sum-product (mentioned earlier), which allows computation of marginals of individual random variables.

Max-product message passing is similar in spirit to dynamic programming algorithms. Like the Viterbi algorithm, it works by passing messages between tree nodes in two stages. In the first stage, messages are passed from the leaf nodes to their parents, which in turn pass messages to their parents, and so on until the messages reach the root node. The message $m_{i\to j}$ from a node i to its parent j is computed as

$$m_{i\to j}(x_j) = \max_{x_i} P(x_j, x_i) \prod_{k\in\mathcal{N}_c(i)} m_{k\to i}(x_i) \tag{1.23}$$

where $\mathcal{N}_c(i)$ is the set of all children of node i. The MAP label of the variable at the root r of the tree can be computed as

$$\hat{x}_r = \arg\max_{x_r} \prod_{k\in\mathcal{N}_c(r)} m_{k\to r}(x_r). \tag{1.24}$$

Given the MAP label $\hat{x}_p$ of a variable X_p, the label of any of its children i can be found as

$$\hat{x}_i = \max_{x_i} P(\hat{x}_p, x_i) \prod_{k\in\mathcal{N}_c(i)} m_{k\to i}(x_i). \tag{1.25}$$

1.3.2 Example: Max-Product BP on a Five-Node Model

Consider the undirected, tree-structured graphical model shown in figure 1.4. The joint distribution factorizes as

$$P(\mathbf{x}) = P(x_a, x_b)P(x_a, x_c)P(x_c, x_e)P(x_c, x_d). \tag{1.26}$$

The messages computed by max-product are

$$m_{d\to c}(x_c) = \max_{x_d} P(x_c, x_d) \tag{1.27}$$

$$m_{e\to c}(x_c) = \max_{x_e} P(x_c, x_e) \tag{1.28}$$

$$m_{c\to a}(x_a) = \max_{x_c} P(x_a, x_c)m_{e\to c}(x_c)m_{d\to c}(x_c) \tag{1.29}$$

$$m_{b\to a}(x_a) = \max_{x_b} P(x_a, x_b). \tag{1.30}$$

The MAP labels can be found as

$$\hat{x}_a = \max_{x_a} m_{b\to a}(x_a)m_{c\to a}(x_a) \tag{1.31}$$

$$\hat{x}_b = \max_{x_b} P(\hat{x}_a, x_b) \tag{1.32}$$

$$\hat{x}_c = \max_{x_c} P(\hat{x}_a, x_c)m_{e\to c}(x_c)m_{d\to c}(x_c) \tag{1.33}$$

$$\hat{x}_d = \max_{x_d} P(\hat{x}_c, x_d) \tag{1.34}$$

$$\hat{x}_e = \max_{x_d} P(\hat{x}_c, x_e). \tag{1.35}$$

The sum product BP algorithm computes the marginal distributions $P(x_i)$ for all variables X_i. It essentially works in a way similar to max-product BP (1.23), except that rather than taking the max, a *sum* is performed over the different labels:

$$m_{i\to j}(x_j) = \sum_{x_i} P(x_j, x_i) \prod_{k\in\mathcal{N}_c(i)} m_{k\to i}(x_i) \tag{1.36}$$

where $\mathcal{N}_c(i)$ is the set of all children of node i. The marginal $P(x_i)$ can be computed by taking the product of the messages sent to the root node i:

$$P(x_i) = \prod_{k\in\mathcal{N}_c(i)} m_{k\to i}(x_i). \tag{1.37}$$

Now, by successively rearranging the tree so that each node i in turn becomes the root node, $P(x_i)$ can be computed for all nodes i.

此为试读，需要完整PDF请访问：www.ertongbook.com