Preface

Overview

Probabilistic graphical models have become a powerful set of techniques used in several domains. This book provides a general introduction to probabilistic graphical models (PGMs) from an engineering perspective. It covers the fundamentals of the main classes of PGMs: Bayesian classifiers, hidden Markov models, Bayesian networks, dynamic and temporal Bayesian networks, Markov random fields, influence diagrams, and Markov decision processes; including representation, inference, and learning principles for all the techniques. Realistic applications for each type of model are also covered in the book.

Some key features are:

- The main classes of PGMs are presented in a single monograph under a unified framework.
- The book covers the fundamental aspects: representation, inference, and learning for all the techniques.
- It illustrates the application of the different techniques in practical problems, an important feature for students and practitioners.
- It includes some of the latest developments in the field, such as multidimensional Bayesian classifiers, relational graphical models, and causal models.
- Each chapter has a set of exercises, including suggestions for research and programming projects.

Motivating the application of probabilistic graphical models to real-world problems is one of the goals of this book. This requires not only knowledge of the different models and techniques, but also some practical experience and domain knowledge. To help the professionals in different fields gain some insight into the use of PGMs for solving practical problems, the book includes many examples of the application of the different types of models in a wide range of domains, including:

- Computer vision.
- Biomedical applications.

- Industrial applications.
- Information retrieval.
- Intelligent tutoring systems.
- Bioinformatics.
- Environmental applications.
- Robotics.
- Human–computer interaction.
- Information validation.
- Caregiving.

Audience

This book can be used as a text book for an advanced undergraduate or a graduate course in probabilistic graphical models for students of computer science, engineering, physics, etc. It could also serve as a reference book for professionals that want to apply probabilistic graphical models in different areas, or anyone who is interested in knowing the basis of these techniques.

It does not have specific prerequisites, although some background in probability and statistics is recommended. It is assumed that the reader has a basic knowledge of mathematics at the high school level, as well as a certain background in computing and programming. The programming exercises require some knowledge and experience with any programming language, such as C, C++, JAVA, Matlab, etc.

Exercises

Each chapter (except the introduction) includes a set of exercises. Some of these exercises are questions and problems designed to reinforce the understanding of the concepts and techniques presented in the chapter. There are also a few suggestions for research or programming projects (marked with "***") in each chapter, which could be used as projects for a course.

Organization

The book is divided into four parts. The first part provides a general introduction and motivation for PGMs, and reviews the required background in probability and graph theory. The second part describes the models which do not consider decisions or utilities: Bayesian classifiers, hidden Markov models, Markov random fields, Bayesian networks, and dynamic and temporal Bayesian networks. The third part

starts with a brief introduction to decision theory, and then describes the models which support decision making, including decision trees, influence diagrams, and Markov decision processes. Finally, the fourth part presents two extensions to the *standard* PGMs, one is relational probabilistic graphical models and the other causal models.

The *dependency relations* between the chapters are shown in Fig. 1. An arc from chapter X to chapter "Y", $X \rightarrow Y$, indicates that chapter X is required (or at least recommended) for understanding chapter Y. This graphical representation of the book gives a lot of information, in an analogous way to the graphical models that we will cover later.

From Fig. 1, we can deduce different ways of reading this book. First it is recommended that you read the introduction and the fundamental Chaps. 2 and 3. Then you can study relatively independently the different models in Part II: classification (Chap. 4), hidden Markov models (Chap. 5), Markov random fields (Chap. 6), and Bayesian networks (Chaps. 7–9). Before reading about learning Bayesian networks (Chap. 8), it is necessary to read Chap. 7—representation and inference; and both chapters are required before going into dynamic and temporal Bayesian networks.

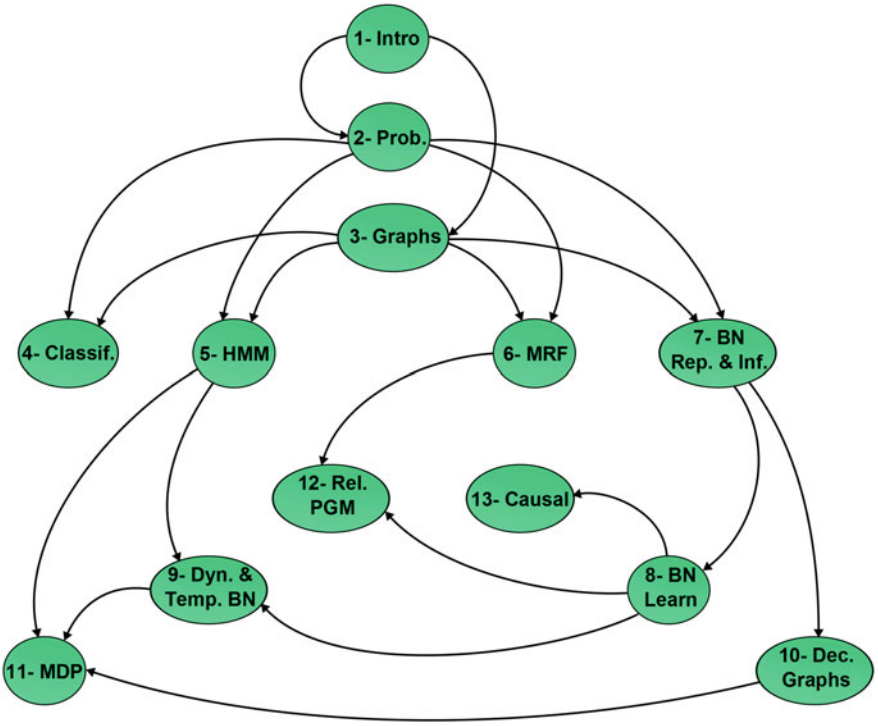

Fig. 1 This figure represents the structure of the book as a directed acyclic graph, showing which chapters are prerequisites for other chapters

The topics in Part III and IV require some of the chapters in Part II. For Chap. 10, which covers decision trees and influence diagrams, you should at least read the first chapter on Bayesian networks. For Chap. 11, which covers sequential decision making, it is recommended that you have covered hidden Markov models and dynamic and temporal Bayesian networks. Relational PGMs (Chap. 12) are based on Markov random fields and Bayesian networks; so Chaps. 6 and 8 are required. Finally, the causal models included in Chap. 13 are based on Bayesian networks including the learning techniques.

If there is not enough time in a course to cover all the book, there are several alternatives. One is to focus on probabilistic models without considering decisions or the more advanced extensions, covering Parts I and II. Another alternative is to focus on decision models, including Part I, the necessary prerequisites from Part II, and Part III. Or you can design your course a la carte, only respecting the dependencies in the graph. However, if you have the time and desire, I suggest you read all the book in order. Enjoy!

Puebla, Mexico
February 2015

Luis Enrique Sucar

http://www.springer.com/978-1-4471-6698-6

Probabilistic Graphical Models
Principles and Applications
Sucar, L.E.
2015, XXIV, 253 p. 117 illus., 4 illus. in color., Hardcover
ISBN: 978-1-4471-6698-6