The Knowledge Engineering Review, Vol. 26:2, 237–238.
doi:10.1017/S0269888910000275

Book Review

Probabilistic Graphical Models: Principles and Techniques by Daphne Koller and Nir Friedman, MIT Press, 1231 pp., $95.00, ISBN 0-262-01319-3

Over the 20 years or so since the publication of Judea Pearl's landmark book *Probabilistic Reasoning in Intelligent Systems*, there has been a steady stream of books on probabilistic graphical models, paralleling the steady rise in the acceptance of probabilistic techniques to the point where they are firmly in the mainstream of artificial intelligence. In my opinion, however, none of these books, excellent as many of them have been, can hold a candle to *Probabilistic graphical models*.

Partly, this is because the earlier books simply had less to say. When Pearl's book was published, the use of probabilistic graphical models was in its infancy, and his monograph—which summarized much of his writing on the subject at the time—contained the bulk of the existing theory. But the field was growing rapidly, and as a result later books both covered more ground and were quickly left behind by further developments. Given this progression, the latest book to come out always looks better than its predecessors since it is a more complete record of the subject.

I certainly think that *Probabilistic graphical models* is a more complete record of work in this area than any book that has come before. At a high level, it has a similar structure to several existing books—being split into four parts that cover representation in graphical models, inference, learning and decision-making—but each part is exhaustive in its account of these broad areas. For example, under 'inference' we not only find the usual algorithms for inference in networks with discrete valued variables but also algorithms for inference in networks with continuous and hybrid (mixed continuous and discrete) valued variables, and algorithms for inference in temporal networks. And we do not just have algorithms for exact inference; we have algorithms for approximate inference, and not just one approach to approximate inference, but a range of different approaches.

Completeness, however, is not the only thing that *Probabilistic graphical models* has going for it. More impressive, to my mind, is the job that the book does of tying this huge mass of material together so that it forms the closest thing to a coherent and absorbing narrative as one is likely to find in a technical volume. The reader is exposed to one view of a topic, led carefully through a treatment of the topic, and then introduced to a variant of the topic and its treatment in a way that feels very intuitive. For a while as I was reading the book, I wondered whether the reason I found the narrative so compelling was because as a researcher I grew up on the edge of this field, aware of many of the developments as they happened, and so had already fitted them into my mental model—so reading the book was just stepping through what I already knew. However, as I got further into the book, I realized that could not be the only reason—there is just too much material that I did not know before reading the book. Rather, it has to be that the authors have done an excellent job of arranging and explaining the material, and I am just in awe of the job that they have done.

Indeed, the only problem I can see with *Probabilistic graphical models* is its sheer size. When recommending it to researchers new to the area—as I have been doing even while still reading it—one has to bear in mind that a 1200-page book is going to be off-putting. Not everyone coming to the book is going to need to read it all, and the issue for such readers is going to be figuring out exactly which bits they need to read. The reader's guide in the introduction makes a good attempt to identify the various strands that run through the book, but I suspect that most readers will struggle to find what they need. However, the worst that is going to happen, if this is the case, is

that readers will end up reading more than they strictly need to. And when a book is this good, that is not much of a downside.

Reviewed by Simon Parsons
Brooklyn College, City University of New York, USA