

Briefly Noted

Learning to Rank for Information Retrieval and Natural Language Processing

Hang Li

(Microsoft)

Morgan & Claypool (Synthesis Lectures on Human Language Technologies, edited by Graeme Hirst, volume 12), 2011, ix+101 pp; paperbound, ISBN 978-1-60845-707-6, \$40.00; ebook, ISBN 978-1-60845-708-3, \$30.00 or by subscription

This short volume gives an introduction and overview of current techniques for learning to rank objects. This is of great interest in information retrieval. Document retrieval systems, meta-search, and collaborative filtering all involve ranking in one way or another. Closer to the computational linguistics audience, the book points out that statistical machine translation typically relies on a re-ranking step to promote better sentence predictions. The author, Hang Li, is a well-respected researcher in information retrieval. He is one of the leading figures on the topic of learning to rank, and is in the core team that maintains the LETOR collection, an important benchmark in that field. Li is therefore in an ideal position to produce a short introduction to the topic.

The book itself is divided into seven very unbalanced chapters. The first two are an introduction to the field and a high-level overview of learning for ranking creation (as opposed to ranking aggregation, which has its own three-page chapter). The heart of this volume is Chapter 4, which lists and briefly describes a number of methods for learning ranking creation or aggregation. Out of 19 methods mentioned in Table 2.6, 12 are described further in that chapter. Additionally, the author covers ranking aggregation using Borda count, Markov chains, and Cranking. Although neither detailed nor exhaustive, this is certainly a fairly comprehensive treatment. A reader familiar with all these algorithms would certainly be well equipped to navigate the field. The book also provides many references in case the reader

would want to further her understanding of the various algorithms. The last three chapters provide a very concise coverage of applications, theory, and future work, respectively.

What is perhaps a bit disappointing with this volume is that in between and around the three main chapters (Chapters 1, 2, and 4), the rest of the book feels somewhat brief and superficial, almost like it was added to make a thorough review paper into a book. Another concern is that whereas the book lists a good number of methods and briefly described references, it is low on analysis and experimental results (apart from Section 2.3.4). A typical example is Chapter 4, which lists 15 methods but keeps their description to two to three high-level pages each. Also, and this is a rather surprising shortcoming of the publisher, the book suffers from an amount of typos and small mistakes that is unusual for this type of publication. For example, Table 2.2 introduces the LETOR benchmark as “LEOTR”!

Who is this book for? In the course of writing this review I have been struggling with this issue. For somebody with only a casual interest in the field, it may not be didactic enough. On the other hand, my impression is that the concise treatment of the various methods and topics covered is not detailed enough for the researcher aiming at implementing learning-to-rank methods and applying them to his needs, such as for example reranking target predictions in machine translation.

Computational linguists interested in the problem of learning to rank may indeed find in this volume a quick and fairly high-level description of a wide range of methods. If they are really interested in the topic, however, they may soon reach the limits of this book and have to turn to the actual papers describing the various methods in more detail in the relevant conferences and journals. The book’s bibliography could actually turn out to be a good starting point.—*Cyril Goutte, National Research Council Canada*

This book review was edited by Pierre Isabelle.