

Preface

In recent years, with the fast growth of the World Wide Web and the difficulties in finding desired information, efficient and effective information retrieval systems have become more important than ever, and the search engine has become an essential tool for many people. The ranker, a central component in every search engine, is responsible for the matching between processed queries and indexed documents. Because of its central role, great attention has been paid to the research and development of ranking technologies. In addition, ranking is also pivotal for many other information retrieval applications, such as collaborative filtering, question answering, multimedia retrieval, text summarization, and online advertising. Leveraging machine learning technologies in the ranking process has led to innovative and more effective ranking models, and has also led to the emerging of a new research area named learning to rank.

This new book gives a comprehensive review of the major approaches to learning to rank, i.e., the pointwise, pairwise, and listwise approaches. For each approach, the basic framework, example algorithms, and their theoretical properties are discussed. Then some recent advances in learning to rank that are orthogonal to the three major approaches are introduced, including relational ranking, query-dependent ranking, semi-supervised ranking, and transfer ranking. Next, we introduce the benchmark datasets for the research on learning to rank and discuss some practical issues regarding the application of learning to rank, such as click-through log mining and training data selection/preprocessing. After that several examples that apply learning-to-rank technologies to solve real information retrieval problems are presented. The book is completed by theoretical discussions on guarantees for ranking performance, and the outlook of future research on learning to rank.

This book is written for researchers and graduate students in information retrieval and machine learning. Familiarity of machine learning, probability theory, linear algebra, and optimization would be helpful though not essential as the book includes a self-contained brief introduction to the related knowledge in Chaps. 21 and 22. Because learning to rank is still a fast growing research area, it is impossible to provide a complete list of references. Instead, the aim has been to give references that are representative and hopefully provide entry points into the short but rich literature of learning to rank. This book also provides several promising future research

directions on learning to rank, hoping that the readers can be inspired to work on these new topics and contribute to this emerging research area in person.

Beijing
People's Republic of China
February 14, 2011

Tie-Yan Liu