Preface

Large, publicly available knowledge graphs (KGs) have given rise to many successful improvements on traditional information retrieval tasks such as ad hoc document retrieval. This workshop focuses on the end-to-end utilization of knowledge graphs and semantics in text retrieval and other IR-related applications. Its scope covers the *acquisition*, the *alignment*, and the *utilization* of knowledge graphs and semantic resources for the purpose of optimizing end-to-end performance of a system that responds to a user's information need. Examples of such technologies and applications include entity ranking, entity linking, entity-based retrieval models, entity recommendation, document filtering, knowledge graph population, and more.

Acquisition includes (but is not limited to) knowledge graph population and semantic resource construction with a special focus on enabling IR-related techniques and applications. Examples include domain/task-specific knowledge graph construction, knowledge representation, and query-time knowledge extraction.

Alignment includes (but is not limited to) the semantic annotation process such as entity linking of short keyword queries or relation extraction for satisfying information needs. It also includes information integration, relation extraction, ontology matching, entity search, and knowledge graph selection based on an information need.

Utilization includes (but is not limited to) using knowledge graphs and semantics in text-centric tasks. Examples are utilizing the knowledge graph to improve document retrieval, question answering, factoid search, dialogue systems, event tracking, and retrieval of complex answers.

This volume contains the papers presented at KG4IR: SIGIR 2017 Workshop on Knowledge Graphs and Semantics for Text Retrieval and Analysis held in conjunction with the ACM SIGIR conference on August 11, 2017 in Tokyo.

This workshop features discussions and presentations on innovative ideas for new methods, suggestions for shared tasks and benchmarks, position papers, as well as reports on practical experiences with KG technology both from academia and industry. We received a total of 11 submissions across two submission cycles. Each submission was reviewed by at least two program committee members and one organizer. The committee decided to accept 8 papers for presentation and poster. The program further includes 6 invited talks and a joint discussion panel with the OBKQA Workshop on the future on knowledge graphs for information retrieval applications.

This work was supported in part by National Science Foundation (NSF) grant IIS-1422676. Any opinions, findings, and conclusions in this paper are the authors' and do not necessarily reflect those of the sponsors. Collections of contributions was supported by EasyChair. Online proceeding service is provided by CEUR-WS.

July 31, 2017
Portsmouth, New Hampshire, USA

Laura Dietz
Chenyan Xiong
Edgar Meij