Report on the 9th ACM SIGIR / the 13th International Conference on the Theory of Information Retrieval (ICTIR 2023)

Masaharu Yoshioka
Hokkaido University
Japan
yoshioka@ist.hokudai.ac.jp

Mohammad Aliannejadi
University of Amsterdam
The Netherlands
m.aliannejadi@uva.nl

Julia Kiseleva
Microsoft
USA
julia.kiseleva@microsoft.com

Abstract

The 9th ACM SIGIR / The 13th International Conference on the Theory of Information Retrieval (ICTIR 2023) held in Taiwan co-located with SIGIR 2023. This brief report provides an overview of ICTIR 2023 and introduces the student reviewer program, which provides an opportunity for the students to understand and contribute to the conference review process.

Date: 23 July 2023.

Website: https://sigir.org/ictir2023/.

1 Introduction

The International Conference on the Theory of Information Retrieval (ICTIR) is the premier conference on theoretical information retrieval. ICTIR focuses primarily on the theoretical aspects of IR and interdisciplinary research that links IR to adjacent fields like artificial intelligence, cognitive science, digital humanities, machine learning, natural language processing, recommendation systems, and social computing. ICTIR also welcomes papers that aim to define novel tasks or apply fundamentally different ideas to information retrieval tasks.

The 9th ACM SIGIR / The 13th International Conference on the Theory of Information Retrieval (ICTIR 2023) held in Taiwan co-located with SIGIR 2023. ICTIR was originally held as an independent conference and changed its style as a co-located event of SIGIR since 2021. We would like to express our sincere gratitude to SIGIR 2023 local organizers.

ICTIR 2023 was a one-day conference held in a hybrid format. To avoid any scheduling conflicts with SIGIR's main conference sessions, we decided to use two parallel tracks this year. We encourage in-person participation but are accepting hybrid presentations due to visa issues.

2 Program

For the one-day program, we had a keynote speech entitled "Evaluating Parrots and Sociopathic Liars" [Sakai, 2023] by Professor Tetsuya Sakai and 30 oral paper presentations selected from
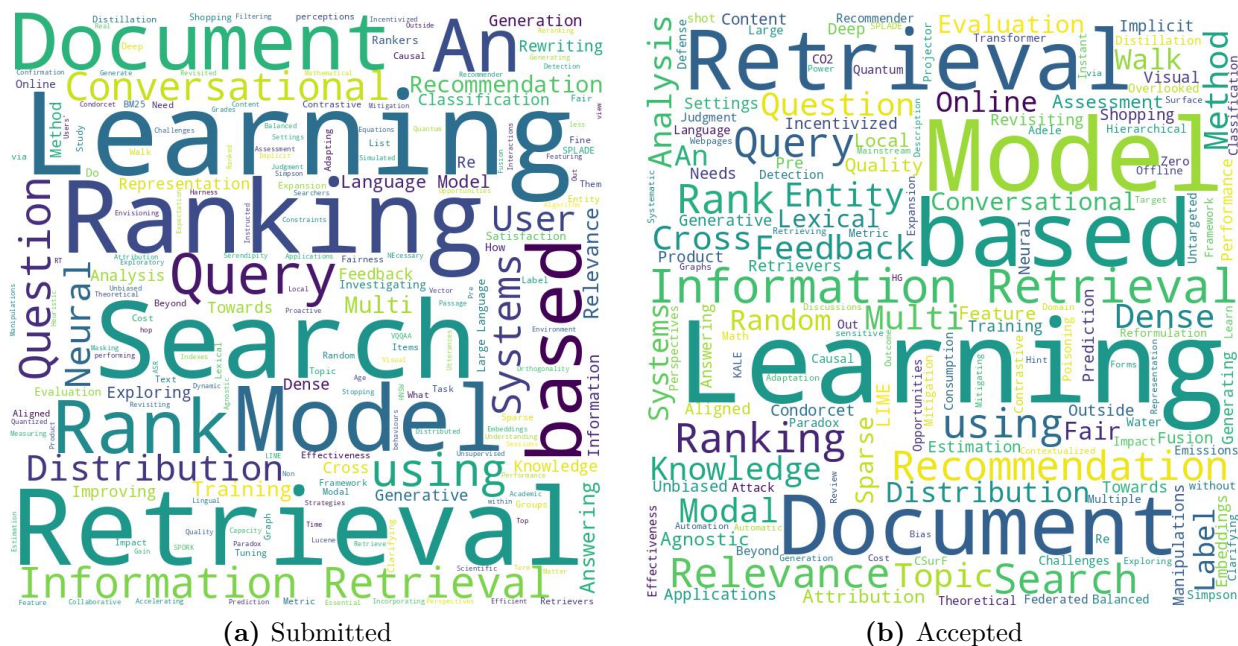
(a) Submitted

(b) Accepted

Figure 1. Word cloud of the titles of the (a) submitted and (b) accepted papers.

73 high-quality submissions (acceptance rate of 41%). Figure 1 shows word clouds of the titles of the submitted and accepted papers. We have chosen "Perspectives on Large Language Models for Relevance Judgment" [Faggioli et al., 2023] as the best paper and granted "Hierarchical Transformer-based Query by Multiple Documents" [Huang et al., 2023] an honorable mention.

3   Student Reviewer Program

This year, we conducted an experiment inviting students to volunteer and help in the review process. To do so, we invited the students through an open call from all over the world and asked them to volunteer to review papers. They had to fill out a form with some basic information about their experience in publishing papers in peer-reviewed conferences and/or journals, as well as their experience in reviewing papers. We received a variety of applications to join the PC, and adopted an inclusive approach, admitting almost all the applicants. We received 86 applications, from which we invited 80 students. This led to an addition of 68 student reviewers. Figure 2 shows the distribution of the applications in terms of their education, where we see that most of the invited volunteers are PhD students.

3.1   Paper Assignment

Our goal was to keep the review load of the students as low as possible (1–2 assigned papers per student). To avoid an unwanted negative impact on the review quality, we considered the student reviewers as additional reviewers for each paper. Therefore, we assigned three main reviewers to each paper with an addition of 1–2 student reviewers.

Table 1. Reviews and PC statistics.

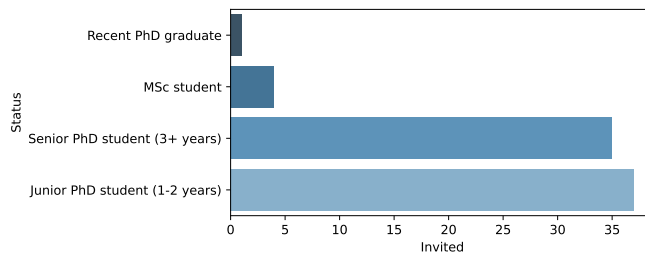|	Non-student	Student
# reviewers	48	68
# reviews	195	77
# comments	320	41
Avg. review length	425	533
Avg. review number	2.64	2.23
Avg. review score	0.07	0.31

Figure 2. Distribution of invited student reviewers' status.

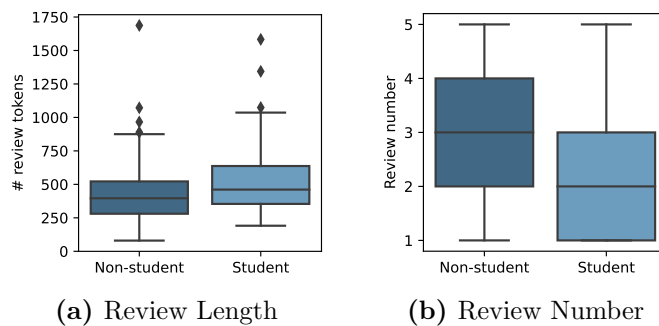**(a)** Review Length **(b)** Review Number

Figure 3. Comparison between non-student and student reviewers in terms of (a) review length and (b) review number (as indicated by EasyChair).
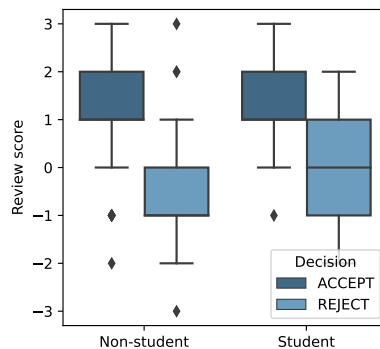

Figure 4. Reviewer score distribution for accepted and rejected papers, per student and non-student reviewers.

Table 2. Review score confusion matrix of non-student and student reviewers.

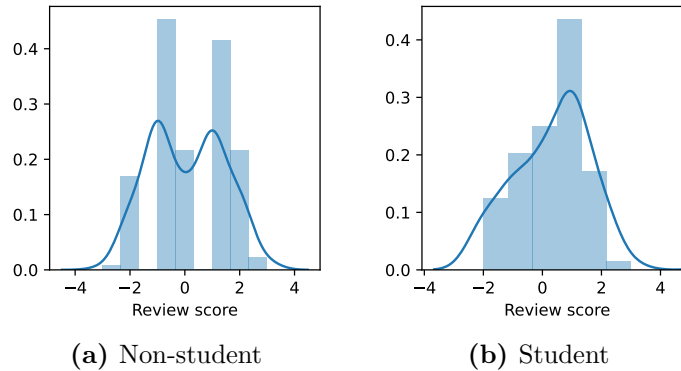|	Accept		Reject	
	Non-student	Student	Non-student	Student
Accept	65	25	17	7
Reject	21	15	95	31

(a) Non-student **(b)** Student

Figure 5. Review score distribution of (a) non-student and (b) student reviews.

3.2 Analysis

As part of this program, we aimed to understand if students exhibit a different behavior when reviewing papers, and whether a more inclusive approach of inviting junior students as PC can be beneficial. First, in Table 1 we report some general statistics of the reviewers and their reviews. As we see, the students wrote significantly fewer reviews in total, as we aimed to limit the student reviews to 1–2 reviews per paper. However, we see that they tended to write longer reviews. This is more evident in Figure 3a where we plot the distribution of review tokens for non-student and student reviewers, suggesting that the students can dive deeper into the technical details of the paper and bring new perspectives that more senior reviewers do not necessarily focus on. We also look at the average review number as an indicator of punctuality, where we see in Figure 3b that the student reviewers did indeed submit their reviews faster than non-student reviewers. In terms of discussion, however, we did not observe a great contribution. Next, we focus on the score distribution and their agreement with the final decisions. We see in Figure 4 for accepted papers, that the score distribution of non-student and student reviewers is very similar. For rejected papers, on the other hand, we observe that the students tend to give higher scores. Looking closer at the confusion matrix of the scores in Table 2 we see higher false negatives for student review scores, especially on the rejected papers. This is also corroborated in Figure 5 where we plot the score distributions and see that there is a clear positivity bias in the student scores, compared to the non-student reviewers. We think that a potential reason for this bias is the fact that the students were mostly assigned one paper to review, which leaves out the possibility of comparing the quality of different papers in the review pool.

4 Summary

We briefly introduce ICTIR 2023 and our new experimental trial of the student review program. We think this program can provide a good opportunity for the students who did not have any chance to understand the review process from the reviewer side and shed light on the potential influence they can have in the program of future conferences. We conclude that even though the student reviews can be less critical, they will provide an additional perspective to the reviews and can potentially spark interesting discussions about the papers. ICTIR 2024 will be held co-located with SIGIR 2024. We welcome submissions and participation in the conference.

Acknowledgement

References

Guglielmo Faggioli, Laura Dietz, Charles L. A. Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, and Henning Wachsmuth. Perspectives on large language models for relevance judgment. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, ICTIR '23, page 39–50, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400700736. doi: 10.1145/3578337.3605136. URL https://doi.org/10.1145/3578337.3605136.

Zhiqi Huang, Shahrzad Naseri, Hamed Bonab, Sheikh Muhammad Sarwar, and James Allan. Hierarchical transformer-based query by multiple documents. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, ICTIR '23, page 105–115, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400700736. doi: 10.1145/3578337.3605130. URL https://doi.org/10.1145/3578337.3605130.

Tetsuya Sakai. Evaluating parrots and sociopathic liars (keynote). In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, ICTIR '23, page 1, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400700736. doi: 10.1145/3578337.3605144. URL https://doi.org/10.1145/3578337.3605144.