Modeling User Feedback in Dynamic Search and Browsing

Jiyun Luo
Department of Computer Science
Georgetown University
Washington, DC, USA
jl1749@georgetown.edu

Categories and Subject Descriptors

•**Information systems** → **Retrieval models and ranking;** *Users and interactive retrieval;*

Keywords

Session Search; Dynamic Search; User Feedback Modeling

1. INTRODUCTION

Nowadays, searching for complicated information is a more and more common task. User's information needs are usually vague or consist of multiple sub-topics, and they must formulate many different queries to achieve their search's goal. A lot of work has been done to improve user satisfaction during the search process. Riccho etc. are developed to help find similar relevant documents. They help the user focus on and exploit the current search topic. xQuAD [5] etc. are efficient diversification algorithms that help the user explore multiple search topics. None of these approaches alone work well in session searches because none of them treat the search session as a whole. They can't answer the question of when to explore and when to exploit.

In our previous work [3], we argue that it is suitable to model session searches as a Dual-Agent Stochastic Game, which essentially is a Partially Observable Markov Decision Process (POMDP) [2] with two agents. Our model treats session search as a "trial-and-error" process and uses user feedback as learning signals to adjust its search strategies, such as exploration and exploitation.

The major feedback we considered in our previous works were query reformulations and user clickthrough datas. I first extend our work by introducing more user feedbacks into our framework. We implement a new search engine UI which allows users to explicitly mark out relevant passages and irrelevant documents. With these explicit feedbacks, I can improve the exploitation algorithm. I use the relevant text to reform a new query for retrieval. The irrelevant documents are then used to re-rank the retrieved documents. For exploration, my algorithm is inspired by xQuAD. The subtopics are found by the user during a search process. I run the algorithm n times in order to recommend n diversified documents. For each round, I sample one subtopic to recommend based on xQuAD.

The second research question I addressed is the decision of when to explore and when to exploit. In [6], the authors explore during the initial phase, and then choose one path to exploit based on user feedback. [4] applies exploration whenever it encounters a diversified query. I argue that their approaches only capture some specific user behavior models. [1] applies exploration as default behavior and interprets user click as a desire to exploit similar documents. However this assumption is not accurate. I use ϵ-greedy as a naive approach to balance Exploration and Exploitation. If the user has tried exploitation multiple times, then we should switch to exploration, and vice versa. Another possible solution could be to pick one search strategy at first, such as "exploitation", and if no positive feedback is received from the user, then we switch to the other search strategy.

The final challenge in this thesis is how to properly evaluate session search algorithms. We propose a new sophisticated evaluation metric, Cube Test. This new metric is able to emphasize subtopic coverage, novelty, and retrieval accuracy at the same time. It also emphasizes minimizing user effort by encouraging short sessions over long sessions. I plan to use this metric, MAP, nDCG, α-nDCG and nERR-IA to evaluate my algorithms as well as other well-known session based retrieval algorithms. I hope these metrics can reveal different aspects of retrieval algorithms and eventually help us to distinguish good session search algorithms from others.

SIGIR '16 July 17-21, 2016, Pisa, Italy

ACM ISBN 978-1-4503-4069-4/16/07.

DOI: http://dx.doi.org/10.1145/2911451.2911483

References

[1] C. Brandt, T. Joachims, Y. Yue, and J. Bank. Dynamic ranked retrieval. In *WSDM '11*.

[2] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.

[3] J. Luo, S. Zhang, and H. Yang. Win-win search: Dual-agent stochastic game in session search. In *SIGIR '14*.

[4] K. Raman, P. N. Bennett, and K. Collins-Thompson. Toward whole-session relevance: Exploring intrinsic diversity in web search. In *SIGIR '13*.

[5] R. L. Santos, C. Macdonald, and I. Ounis. Exploiting query reformulations for web search result diversification. In *WWW '10*.

[6] M. Sloan and J. Wang. Dynamic information retrieval: Theoretical framework and application. In *ICTIR '15*.