SIGIR WORKSHOP REPORT

Information Retrieval and Advertising

Mikhail Bilenko
Microsoft Research
mbilenko@microsoft.com

Evgeniy Gabrilovich
Yahoo! Research
gabr@yahoo-inc.com

Matthew Richardson
Microsoft Research
mattri@microsoft.com

Yi Zhang
University of California at Santa Cruz
yiz@soe.ucsc.edu

Abstract

The Information Retrieval and Advertising Workshop (IRA 2009) was held on July 23, 2009 in Boston, Massachusetts, in conjunction with the 32nd Annual ACM SIGIR Conference. The workshop covered theoretical and empirical issues in several research areas that span the intersection of computational advertising, information retrieval, and economics. The workshop consisted of 3 invited talks, 6 refereed paper presentations, 2 positions statement presentations and several discussion sessions.

1 Introduction

Online advertising has become the primary business model that supports a significant fraction of today's Web experience, including major Web search engines and numerous content-driven Web sites. Computational advertising systems employ many IR techniques alongside approaches developed in statistical modeling and machine learning, large-scale data processing, optimization, microeconomics, and human-computer interaction. Despite its commercial significance, computational advertising is a relatively young research discipline. The purpose of the SIGIR 2009 Workshop on Information Retrieval and Advertising was to provide a forum for discussing new issues, as well as to strengthen collaborations between industry and academia.

In the past several years, advertising-related workshops were organized at a number of conferences. Notably, in 2008, workshops focusing on various aspects of online advertising were held at the SIGIR, KDD, WWW, and EC conferences. However, depending on the host conference, these workshops were somewhat narrowly focused on small subsets of issues within the realm of online advertising, such as textual relevance, statistical modeling, or auctions and game theory. Given the recent advances in these individual directions as well as their convergence, we believed the time was ripe to address the interactions between these different aspects of online advertising, and to seek cross-pollination among the different fields. The broad reach of SIGIR made it an ideal venue to host such multi-disciplinary workshop. To achieve this goal, we also had a diverse representation of the relevant disciplines on the workshop program committee.

In order to bring together as many ideas as possible, we solicited submissions on a wide range of topics, from user behavior analysis, HCI, and social network advertising, to content and keyword analysis and ad matching. We sought a variety of submission types, including position papers, long papers, and short papers reporting on late-breaking results.

2 Technical Program

The workshop program included 6 papers (3 long and 3 short) and 2 position statements. We also invited three keynote speakers to talk about their respective area of expertise as related to computational advertising, representing state-of-the-art in academic and industrial research. Discussion time was provided during paper presentations and at the end of the workshop.

Invited Talks

- *Ads in Dynamic Query Suggestions*
  Yoelle Maarek
- *Online Advertising: Designing and Optimizing Marketplaces*
  Susan Athey
- *High Precision Text Mining for Product Search*
  Kamal Nigam

Position Statements

- *Better Query Modeling for Sponsored Search*
  Hema Raghavan
- *Flexible Advertising Is Gaining Momentum*
  Roumen Vragov

Papers

- *Get more Clicks!*
  Derek Hao Hu, Evan Wei Xiang and Qiang Yang
- *Sponsored Search for Political Campaigning during the 2008 US Elections*
  Eni Mustafaraj and Panagiotis Takis Metaxas
- *The Effect of Some Sponsored Search Auction Rules on Social Welfare: Preliminary Results from an Exploratory Study in the Laboratory*
  Roumen Vragov, David Porter and Vernon Smith
- *Exploring Collaborative Filtering for Sponsored Search*
  Sarah Tyler, Yi Zhang and Dou Shen
- *Towards Advertising in Social Networks*
  Maryam Karimzadehgan, ChengXiang Zhai and Manish Agrawal
- *A Preliminary Study on Dynamic Keyword Extraction for Contextual Advertising*
  Wen Ye, Wenjie Li, Furu Wei and Chunbao Li

2.1 Invited Talks

The invited talks were very well received by workshop participants: in the anonymous feedback forms gathered at the end of the workshop, four out of five positive comments were with respect to the invited talks.

The first invited talk was by Yoelle Maarek, who began by discussing the recent appearance of ads in Google's query suggestion box (which suggests queries as you type). This can be seen as an additional step in the trend toward richer interaction between a search engine and its users. Providing ads alongside query suggestions leads to several challenges, such as computational efficiency and effectiveness (relevance to the query being typed), new forms of user experience, and possibly new auction models. The problem is particularly difficult because of the small amount of context available (only a partial query), the limited amount of space for displaying ads, and users' high expectation on ads relevancy. On the positive side, query suggestion ads have exceptionally high clickthrough rates, so solving these problems can be beneficial for both search engines and advertisers.

The second invited talk was by Susan Athey, who extensively discussed auction systems and advertising market design. Tutorial-like in nature, Susan's presentation covered a wide range of economic concepts, including two-sided platform-based markets, the use of auctions for matching, and how such markets can tend to converge to monopolies (e.g., eBay). Susan also introduced the difference between market design and mechanism design, and explained why one should consider both aspects of the market, and how to design a market to encourage participation. Following the introduction of these concepts, she then applied them to online advertising markets, pointing out that due to the presence of multiple ad slots, the current search advertising auction model does not lead to truthful bidding. She proposed a model of how advertisers bid, which can be used to perform counterfactual modeling by simulating the changes in bids that would result from a change in the system. She concluded with some examples that demonstrated the importance of considering the long-term effects of market and mechanism changes. For example, using a model of bidding behavior allows the market maker to estimate the bidders' true value, and from that to estimate their reaction to market changes, such as increased competition. In some cases, a market change such as increased competition may initially lead to increased revenue, however, eventually it may lead to decreased revenue once the advertisers adjust their behavior accordingly. This talk was very beneficial to the audience as it offered the attendees a new perspective on the problem from the economics point of view, which is rarely represented in computer science venues.

The final invited talk was by Kamal Nigam, who discussed the techniques used in Google's product search. The data received by the product search team comes from both structured and unstructured sources (e.g., crawled from merchant or review websites). Processing the latter requires sophisticated text mining techniques. The data also requires significant processing to remove spam, duplicates, mistakes, etc. Kamal discussed their work on scaling up wrapper induction, which is a process to build a "wrapper" for extracting structured data from an unstructured Web site. Although several prior studies focused on this direction, the previous work does not scale to the extent necessary for extracting data for Google's product search. The key difference between previous work on wrapper induction and the work presented by Kamal is that the previous work is done per Web site, while the current work at Google does it per each vertical source. This allows the product search team to scale to many more Web sites, since the amount of work is proportional to the number of verticals (general shopping categories) covered, rather than the number of individual sites. Kamal also discussed other work related to product matching: de-duplicating product records, which are often noisy and contain missing data. Instead of traditional approaches based on the similarity between each pair of products, the proposed approach is based on IR indexing and a machine learning algorithm, which lead to high accuracy. By combining advanced wrapper-induction and text-mining techniques with advanced de-duplication methods, the system is able to perform at the high levels of precision and recall required for a commercial product search system.

2.2 Position Statements

In her position statement, "*Better Query Modeling for Sponsored Search*", Hema Raghavan discussed how some problems common in Web search are exacerbated in sponsored search advertising. Because of the relatively small corpus of content, advertising-based information retrieval performs even worse than Web-based information retrieval when dealing with long queries. For example, a query such as "Donna Karan New York in Boston" can cause difficulties with all three major search engines. Hema suggested that better entity detection and query "chunking" may help, along with better understanding of geographical intentions. Hema's position statement prompted an interesting discussion of differences between advertising and Web retrieval.

Roumen Vragov presented his position statement, "*Flexible Advertising is Gaining Momentum*" as a way to facilitate discussion about being more flexible in specifically tailoring advertising on a per-user basis, which may even mean giving the users the ability to opt-out of seeing advertisements at all. He introduced an experiment by Panasonic whereby users could push a button on their remote in order to skip a commercial for a small fee of $0.05 per commercial skipped. He proposed a similar model may work for sponsored search advertising, where users may be willing to pay in order to not see any ads, or alternatively, a search engine may give users the ability to temporarily opt-out of ads for no fee. This prompted a lively discussion about whether this would make users more likely to click on future ads, whether users feel they get value from ads already (e.g., sometimes the ads can be more relevant than the search results), and whether users would be willing to explicitly provide information about themselves to improve ad targeting.

2.3 Refereed Papers

The IRA 2009 Program Committee accepted six papers for presentation at the workshop.

Derek Hao Hu presented his work, titled "Get More Clicks". In this paper, Hu et al. evaluate the effectiveness of a number of commonly-held notions for improving advertisement clickthrough rates, such as having action terms like "get" or "shop", creating a sense of urgency, offering free products, etc. Using historical logs, the authors found that these features do indeed lead to higher CTRs, and incorporating these features into a ranking model allow it to better predict which ads will have high clickthrough rates.

Eni Mustafaraj presented her work, titled "*Sponsored Search for Political Campaigning during the 2008 US Elections"*. In this paper, Mustafaraj and Metaxas investigated what kinds of ads appeared for political candidate searches during the 2008 United States elections. They categorized the advertisers in four different categories: commercial, partisan, non-affiliated, and media. By analyzing the content of the collected ads, the authors discovered that the majority of them (63%) are commercial ads that have no political message, while the partisan group contributed only 14% of the ads. Furthermore, only 21 out of 124 monitored candidates were actively participating in sponsored search, by providing their own political message. One problem with online advertising for political campaigns is that FEC and other regulatory bodies may enforce laws requiring the ads to specify that they are political in nature and who is funding the advertisement. In a different line of research, Mustafaraj and Metaxas also investigated advertising behavior for the recent swine flu outbreak. They found that the initial ads were either highly relevant or accidentally mismatched ones. However, the spammers quickly arrived thereafter and filled the ad slots with irrelevant spam ads.

Roumen Vragov presented his work, titled "*The Effect of Some Sponsored Search Auction Rules on Social Welfare: Preliminary Results from an Exploratory Study in the Laboratory*". In this paper, Vragov et al. presented the findings from an empirical economic study that aimed to investigate the effect that variations on auction design may have on the overall revenue of the advertising platform. They sought to study the effect of ranking by CTR vs. by relevance, and also the effect of different payment methods (pay-per-click vs. pay-per-impression). They set up an artificial market with virtual goods, where advertisers are trying to sell the goods and buyers are clicking on the ads. In their study, the best results were obtained by using ranking according to historical clickthrough rates, and paying per click. However, they found that the choice of the payment method had a more pronounced effect than the choice of the ranking function.

Sarah Tyler presented her work, titled "*Exploring Collaborative Filtering for Sponsored Search*". In this paper, Tyler et al. investigated whether collaborative filtering can better predict which ads a user is likely to click on. They argued that advertisers do not necessarily bid on all terms that they should,

nor does an individual query capture all the context about a user. They suggested that by using a more advanced collaborative-filtering-based matching system one may be able to better match ads to users who are likely to click on the ads. The authors used the k-nearest neighbors algorithm and a latent factorization model, where the queries represent the user's information need. In re-ranking and filtering experiments, they found that collaborative filtering can help improve both CTR prediction and broad matching of advertisements.

Maryam Karimzadehgan presented her work, "*Towards Advertising in Social Networks*". In this paper, Karimzadehgan et al. gave a brief overview of the main issues facing advertising in social networks, including how to exploit social relationships, how to model users, and how to evaluate a novel system. Exploiting relations can be complex since they may change over time, and also vary in strength and influence. User modeling has been shown to help document ranking, so using the data in a social network, the evolution of user interests, and the interests of those connected to a user, may lead to better ad ranking. Evaluation can be difficult as there are multiple pertinent metrics (e.g., revenue, user experience), and creating a test set can be difficult due to the subjective measure of relevance. Finally, Maryam presented a news recommendation system implemented in Facebook, which was used to study social and community-based relevance.

Wen Ye presented his work, "*A Preliminary Study on Dynamic Keyword Extraction for Contextual Advertising*". In this work, Ye et al. studied the problem of extracting keywords from Web pages that are constantly changing, such as blogs or Web-based discussion forums. The authors considered two types of keywords: (a) a set of keywords characterizing the context of the entire Web page, and (b) keywords for each new posting in the discussion. The main challenges of this work are topic shift, observed when a discussion moves away from its original topic, and brevity, where some postings do not contain many relevant on-topic keywords. The proposed solution uses the linking relations that are built on top of replies and quotations. The authors proposed a dynamic extraction approach for both inter-post and whole page keyword extraction. Using evaluation that employed human judges, the authors showed that their approach results in more topically-relevant keywords.

3 Conclusions and Feedback

The SIGIR 2009 Workshop on Information Retrieval and Advertising was an overall success. We assembled an interdisciplinary technical program covering a wide range of topics. The program was quite tight, with 11 presentations throughout the day, yet it left sufficient time for discussion after each paper and at the end of the workshop. The interactions between researchers from different fields led to very interesting and informative discussions. The invited talks and position statements in particular inspired interesting debate and prompted a number of suggestions for possible future research directions.

According to the feedback we received from the IRA 2009 attendees, participants enjoyed the workshop and are willing to attend it again next year. Participants were excited about the invited talks, as well as the interesting discussions between academic and industrial researchers. Some participants also indicated that the amount of discussion could have been greater. For future workshops, we recommend to consider increasing the interaction between participants, such as an initial round of introductions, facilitating open discussion during lunch, and a panel.

4 Acknowledgments

We would like to thank the Program Committee members (see http://ira09.soe.ucsc.edu/#comittee) for their valuable reviews, the authors for their interesting papers, and the attendees for the inspirational discussions.