

vertisers who bid for the exact user query can participate in its auctions. However, in real-life systems, a listing can participate in the “related” keywords’ auctions as well. For example, a listing that bids on “shoes” can participate in a user query “running shoes”. We can extend our method to such cases by forming a mapping between the bid keywords and user queries for an advertiser by using logs. In future, we plan to conduct rigorous experiments using our method’s extension for such listings. Finally, we used a first order DLM to capture trends and momentary peaks in query search volume. However, the DLM ignores information about query traffic features and only focuses on the total number of searches. We plan to address this limitation using dynamic Bayes net, that smoothly vary conditional probability distributions over each edge. Bayes net also allows for feature targeting: For example, if the advertiser has a certain geographical area, or a demographic group to which they would like to advertise, then by fixing the corresponding nodes in the Bayes net, we can generate the traffic corresponding to the targeting.

Acknowledgements

We would like to thank A. Kumaran, Krishna Leela Poola, B. Ashok, and Sayan Pathak for several discussions that helped in shaping core ideas of the paper and also for detailed comments on an initial draft of the paper.

7. REFERENCES

[1] S. Athey and D. Nekipelov. A structural model of sponsored search advertising auctions. In *Technical report, Microsoft Research*, May 2010.

[2] Y. Cui, R. Zhang, W. Li, and J. Mao. Bid landscape forecasting in online ad exchange marketplace. In *KDD*, pages 265–273, 2011.

[3] Q. Duaong and S. Lahaie. Discrete choice models of bidder behavior in sponsored search. In *WINE*, 2011.

[4] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1), March 2007.

[5] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *ICML*, pages 13–20, 2010.

[6] D. Koller and N. Friedman. *Probabilistic Graphical Models*. MIT Press, 2009.

[7] C. D. Manning, P. Raghavan, and H. Schtze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.

[8] F. Pin and P. Key. Stochastic variability in sponsored search auctions: observations and models. In *ACM Conference on Electronic Commerce*, pages 61–70, 2011.

[9] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25 (6):1163–1178, 2007.

[10] X. Wang, A. Z. Broder, M. Fontoura, and V. Josifovski. A search-based method for forecasting ad impression in contextual advertising. In *WWW*, pages 491–500, 2009.

[11] M. West and J. Harrison. *Bayesian Forecasting and Dynamic Models*. Springer Series in Statistics, 1997.