Multileave Gradient Descent for Fast Online Learning to Rank

Anne Schuth[1,3] anne.schuth@uva.nl
Harrie Oosterhuis[1] harrie.oosterhuis@student.uva.nl
Shimon Whiteson[2] shimon.whiteson@cs.ox.ac.uk
Maarten de Rijke[1] derijke@uva.nl

[1] University of Amsterdam, The Netherlands
[2] University of Oxford, United Kingdom
[3] Blendle, The Netherlands

Dueling Bandit Gradient Descent (DBGD)

[1] T. Joachims. Optimizing search engines using clickthrough data. In KDD, 2002.
[2] Y. Yue and T. Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In ICML, 2009.

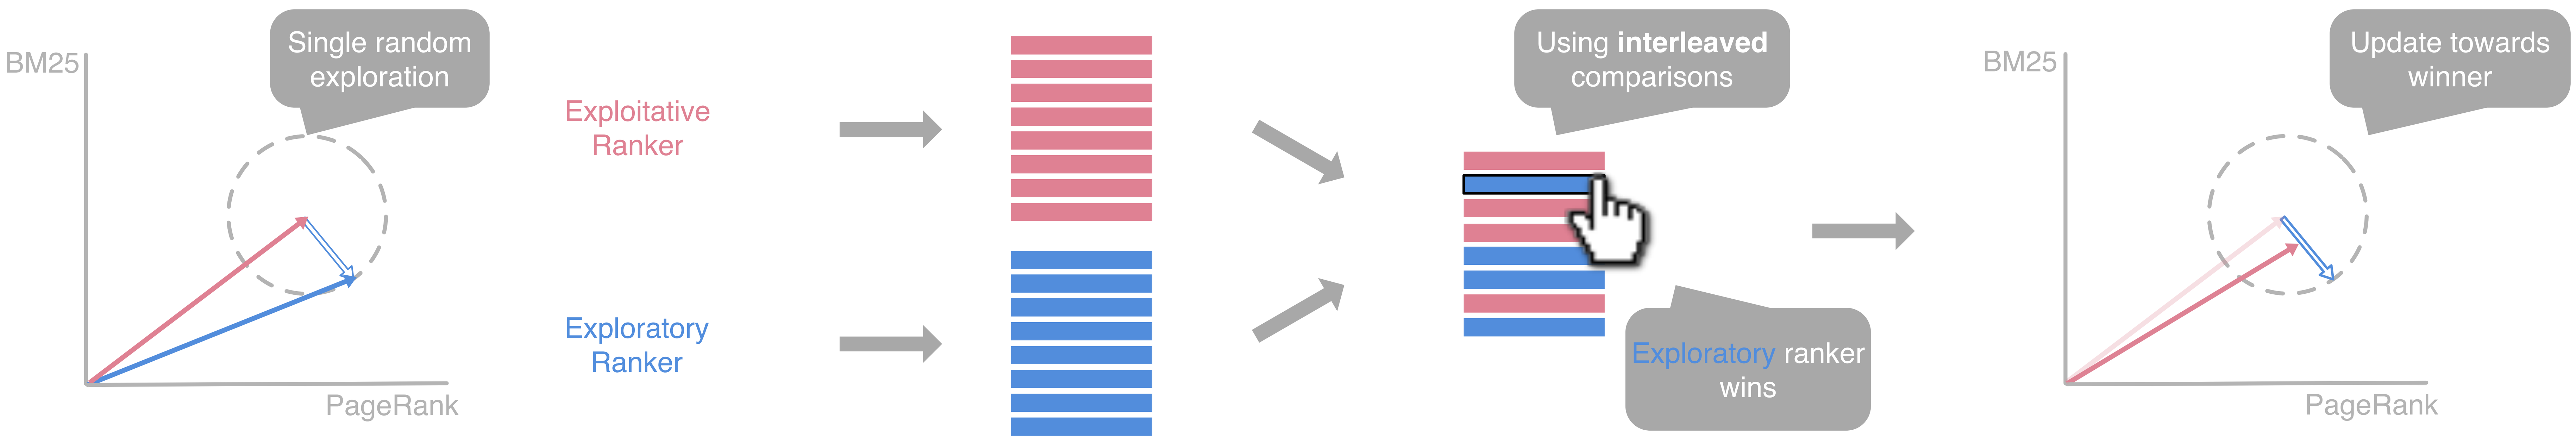

Multileave Gradient Descent (MGD)

[3] A. Schuth, F. Sietsma, S. Whiteson, D. Lefortier, and M. de Rijke. Multileaved comparisons for fast online evaluation. In CIKM, 2014.

Winner Takes All (MGD-W)

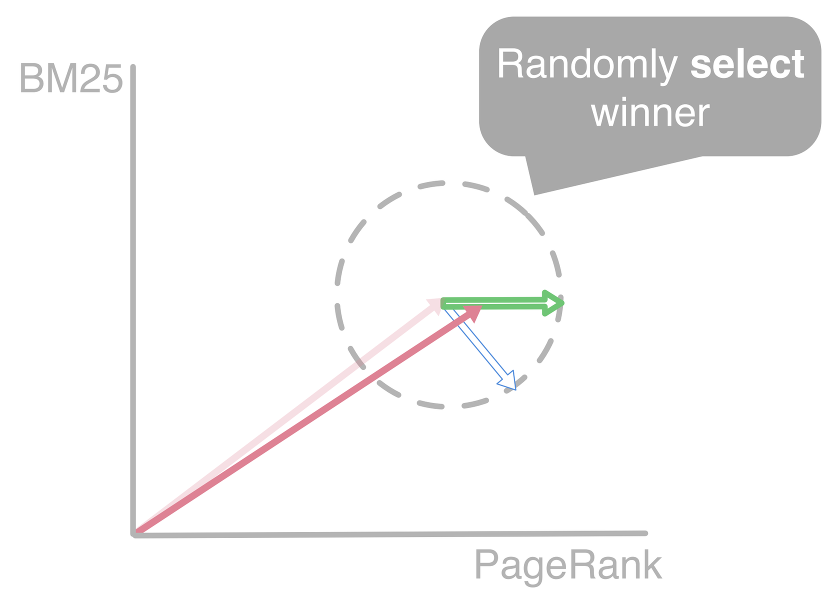

Mean Winner (MGD-M)

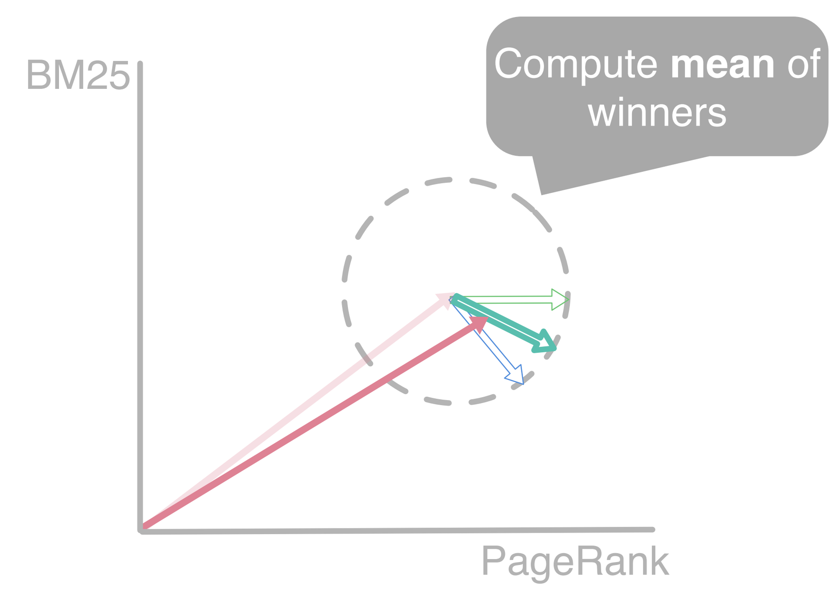

Experimental Results

[4] A. Schuth, K. Hofmann, S. Whiteson, and M. de Rijke. Lerot: An online learning to rank framework. In LivingLab Workshop at CIKM, 2013.

Noisy Feedback

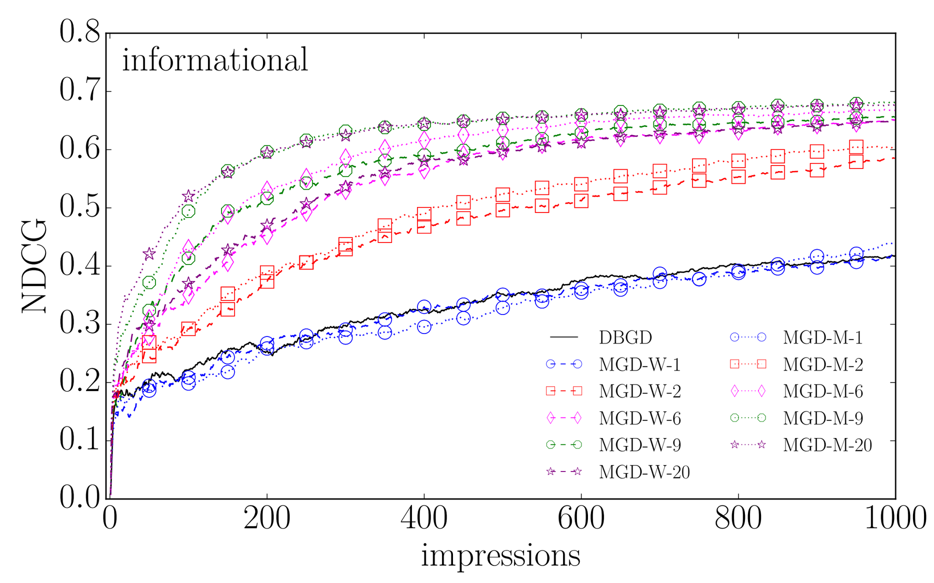

Long Run

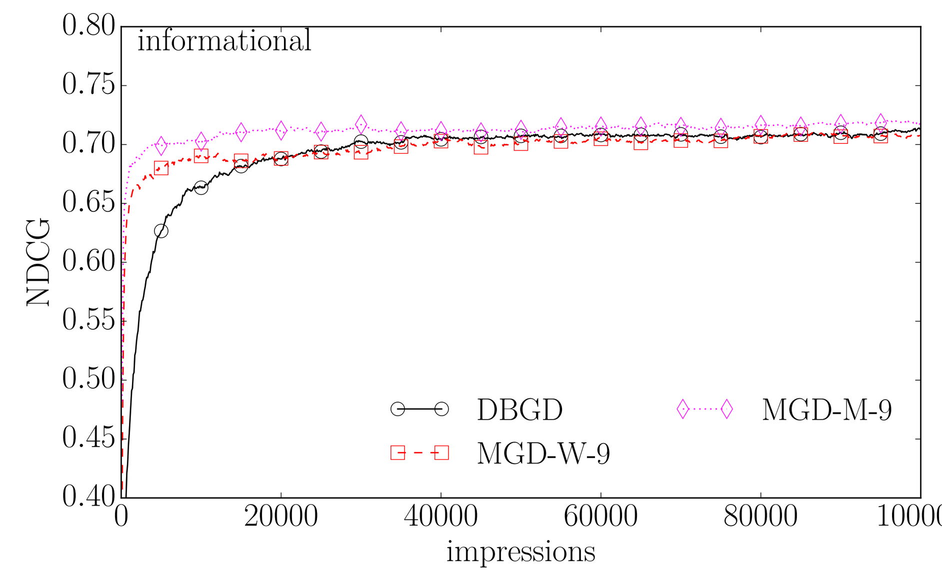

Learning Rate

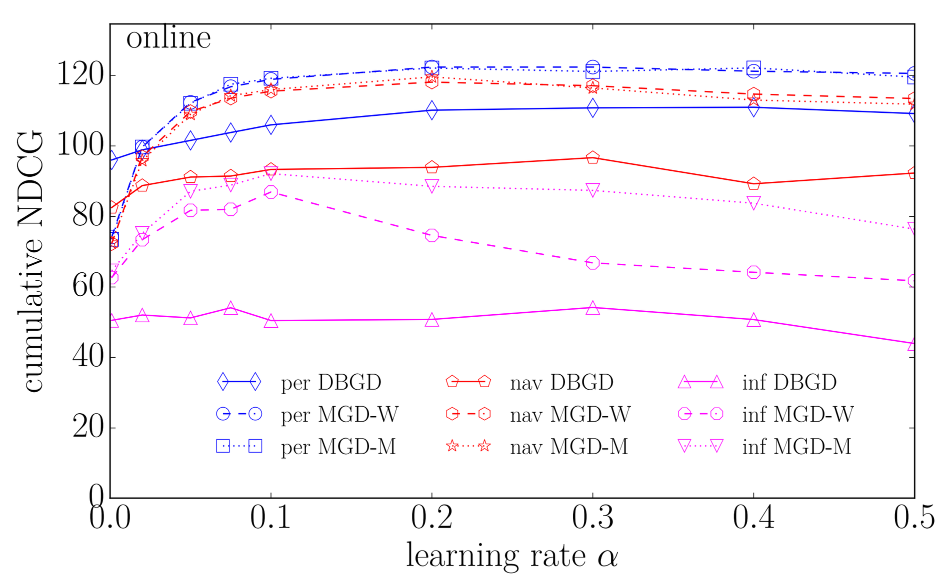

Conclusions

- Experiments show dramatic improvements over the baseline
- In particular with noisy feedback, MGD learns much faster than DBGD
- MGD-M performs equal or outperforms MGD-W
- Orders of magnitude less interaction data is required
- Far fewer user are exposed to inferior rankers

University of Amsterdam

WSDM 2016, San Francisco, US

This research was partially supported by Amsterdam Data Science, the Dutch national program COMMIT, Elsevier, the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement nr 312827 (VOX-Pol), the ESF Research Network Program ELIAS, the HPC Fund, the Royal Dutch Academy of Sciences (KNAW) under the Elite Network Shifts project, the Microsoft Research Ph.D. program, the Netherlands eScience Center under project number 027.012.105, the Netherlands Institute for Sound and Vision, the Netherlands Organisation for Scientific Research (NWO) under project nrs 727.011.005, 612.001.116, HOR-11-10, 640.006.013, 612.066.930, CI-14-25, SH-322-15, the Yahoo! Faculty Research and Engagement Program, and Yandex.