Learning To Rank at Bloomberg

from Theory to Production

Diego Ceccarelli
Bloomberg LP
London
dceccarelli4@bloomberg.net

ABSTRACT

Learning to Rank exploits machine learning to improve the quality of search results in almost all commercial web search engines around the world. For the last three years, Bloomberg engineers have worked hard to add support for Learning to Rank in Apache Solr, and their contributions made Solr one the first open source engines to do it out of the box. But all that is for nought if you do not hunt down the necessary features, make it inter-operate with all the other functionality, and do this fast enough on a production system for such ranking to be feasible.

This talk is not just about how we build this functionality into Apache Solr: it is a war story of how the company's real-time, low-latency news search engine was tamed to learn how to rank. Join us on a journey that will teach you how to take your Learning to Rank system to clients, and more importantly, the many ways not to do it. There will be drama, excitement, and despair! Now grab that popcorn...

REFERENCES

[1] 2017. How Bloomberg Integrated Learning-to-Rank into Apache Solr. https://www.techatbloomberg.com/blog/bloomberg-integrated-learning-rank-apache-solr/.
[2] 2017. Learning To Rank, Lucene/Solr guide. https://lucene.apache.org/solr/guide/6_6/learning-to-rank.html.
[3] Diego Ceccarelli Michael Nilsson. 2015. Learning To Rank in Solr. Lucene/Solr Revolution - https://www.youtube.com/watch?v=M7BKwJoh96s.
[4] Diego Ceccarelli Michael Nilsson. 2017. Learning To Rank For The Win. Berlin Buzzwords - https://berlinbuzzwords.de/17/session/apache-solr-learning-rank-win.