Does LambdaMART Do What You Expect?

Discussion Paper

Claudio Lucchese[1], Federico Marcuzzi[1,*] and Salvatore Orlando[1]

[1]Università Ca' Foscari Venezia, Via Torino, 155, 30170 Mestre, Venezia VE, Italy

Keywords
learning to rank, information retrieval, machine learning

Optimising information retrieval metrics is crucial in Learning-to-Rank tasks. One of the most effective approaches to do this is the well-known LambdaMART [1] algorithm. LambdaMART smoothly approximates for each document a gradient with respect to the IR metric to be optimised, the so-called lambdas. Intuitively, each lambda describes, with some degree of approximation, how much a document score should be pushed up or down to improve the ranking.

In this work, we show that lambdas can be incoherent with respect to document relevance: a document with high relevance can receive a downward push larger than a document with lower relevance. This behaviour goes far beyond the expected degree of approximation. In addition, despite the gain in training efficiency, directly optimising truncated evaluation metrics can exacerbate this discrepancy due to less number of lambdas contributions each document gradient receives. As a consequence each document gradient is only partially computed and this lead to worse model learning.

We analyse the idiosyncrasies of LambdaMART gradients and we introduce some strategies to remove or reduce gradient incoherencies. Specifically, we designed three selection strategies to compute the full gradient for only those documents that should be ranked in the top-k positions of the ranking. We empirically demonstrate on publicly available datasets that the proposed approach leads to models that can achieve statistically significant improvements in terms of NDCG [2] while maintaining the same training efficiently as optimising truncated metrics.

References

[1] C. J. Burges, From ranknet to lambdarank to lambdamart: An overview, Learning 11 (2010).

[2] K. Järvelin, J. Kekäläinen, IR evaluation methods for retrieving highly relevant documents, in: E. J. Yannakoudakis, N. J. Belkin, P. Ingwersen, M. Leong (Eds.), SIGIR 2000: Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, July 24-28, 2000, Athens, Greece, ACM, 2000, pp. 41–48.

IIR2023: 13th Italian Information Retrieval Workshop, 8th - 9th June 2023, Pisa, Italy

*Corresponding author.
claudio.lucchese@unive.it (C. Lucchese); federico.marcuzzi@unive.it (F. Marcuzzi); orlando@unive.it (S. Orlando)
0000-0002-2545-0425 (C. Lucchese); 0000-0002-8141-8294 (F. Marcuzzi); 0000-0002-4155-9797 (S. Orlando)