# Multi-Objective Optimization for Information Retrieval[1]


Joost van Doorn[a] Daan Odijk[ac] Diederik M. Roijers[ab] Maarten de Rijke[a]

[a] *University of Amsterdam, Informatics Institute*
[b] *Vrije Universiteit Brussel, Artificial Intelligence Laboratory*
[c] *Now at Blendle*


The primary goal of information retrieval (IR) systems is to satisfy the information need of a user. To achieve this most search engines are optimized to rank documents based on their topical relevance to the query. However other relevance criteria, such as credibility, reputability and readability can play a role depending on the information need of the user [7]. The importance of these criteria may differ, and it might not be known a priori how these criteria need to be balanced. E.g. a child might need more easily readable documents in comparison to other users.

We propose to mitigate this by viewing multiple relevance criteria as objectives and learning a set of rankers that provide different tradeoffs w.r.t. these objectives. We combine the multi-objective technique *Optimistic Linear Support (OLS)* [5] with multiple utility-based metrics in a learning-to-rank setting.

**Scalarization function** If we can measure a value for every relevance criterion, then $V_i$ is the value for relevance criteria $i$. Following [4] we use a *scalarization function* $f$, that collapses the value vector to a scalar utility: $f(\mathbf{V}, \mathbf{w})$, where $\mathbf{w}$ is a vector that parameterizes $f$. As we cannot observe the set of all possible rankers directly, as such we aim to find a *coverage set* [4] of rankers, that contains an optimal trade-off for each possible preference (i.e., $f$ and $\mathbf{w}$) that a user might have, see Figure 1.

We assume that $f$ is linear: $f(\mathbf{V}, \mathbf{w}) = \mathbf{w}^T\mathbf{V}$, i.e., the utility for the user is a *weighted* sum ($\sum_i^C w_i = 1$) over relevance criteria.

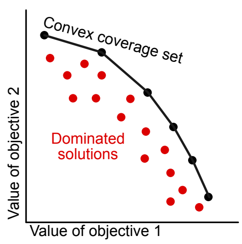


Figure 1: The points on the line represent solutions in the coverage set, the others are dominated.

**Metrics as objectives** To formulate our own scalarization function we can follow the standard notion of gain based utility functions [1], and adapt it to a scalarized value function:

$$V_i = \frac{1}{N_i} \sum_{k=1}^{K} gain_i(doc_k) \times discount(k) \quad (1)$$

Where $k$ is the rank, $discount(k)$ lowers the weight of lower ranked documents, and each of the $C$ relevance criteria has a separate gain function, which are based on expert annotations.

**Convex coverage set** Because each criterion contributes positively to the scalarization function, and we are interested in the *relative* importance of each criterion, we can assume that $\mathbf{w}$ is a positive vector that sums to 1 in order to determine a coverage set. A coverage set that covers all possible linear scalarizations is called a convex coverage set (CCS) [4]. To compute the CCS, we use OLS (see [5]); OLS computes a CCS by solving a multi-objective problem as a series of single-objective problems, i.e., problems that are scalarized using different $\mathbf{w}$. At each iteration, OLS tries to find a new ranker, thereby incrementally building a CCS. OLS finishes when no new rankers are found. Figure 2 illustrates

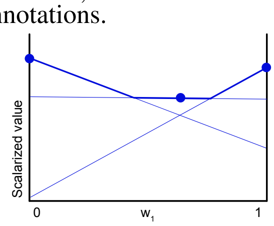


Figure 2: CCS with three rankers (dots), and their respective scalarized value functions (lines).

[1]Compressed contribution of the paper "Balancing Relevance Criteria through Multi-Objective Optimization" in *SIGIR '16*, pages 769-772. 2016. [8]

the CCS with three rankers. We can use existing single-objective optimization techniques to find rankers for a given $\mathbf{w}$. We use Dueling Bandit Gradient Descent (DBGD) [9]. As each iteration of OLS requires re-running the learning to rank algorithm, an easy way to improve convergence speed is by reusing previous solutions [6]. For each new corner weight, we multi-start DBGD, initializing the parameters based on the rankers that were found at the three closest corner weights so far. Additionally we employ Iterated Local Search [2] to avoid getting stuck in local minima, as OLS requires the single-objective optimization method to be approximately optimal [6].

In practice the ranker presented to the user is selected based on: $\max_{\mathbf{V} \in S} f(\mathbf{V}, \mathbf{w})$, where $S$ are the values of the rankers, and $\mathbf{w}$ is determined based on the user profile, or set in the user interface.

To demonstrate how multi-objective optimization for balancing multiple relevance criteria works in practice, we perform experiments on two datasets: (i) balancing readability and topical relevance in a health setting (CLEF eHealth 2015 task 2 [3]), and (ii) balancing diversity and topical relevance in a web search dataset annotated for sub-topic relevance (TREC 2012 Web Track diversity task).

On CLEF eHealth 2015 task 2, both topicality, and readability is evaluated using nDCG. While on TREC 2012 diversity task, we simultaneously optimize nDCG and $\alpha$-nDCG. The CCSs, constructed using OLS, are shown in Figure 3 and Figure 4. Fewer solutions were found for the CCS compared to the readability task, suggesting a large correlation between the metrics nDCG and $\alpha$-nDCG, from which we conclude that this setting is less suitable for our method. Whereas for readability and relevance more different alternative rankings can be presented to the user.

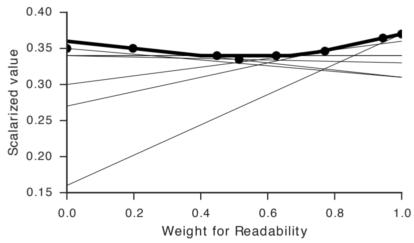


Figure 3: The CCS found on CLEF eHealth 2015.

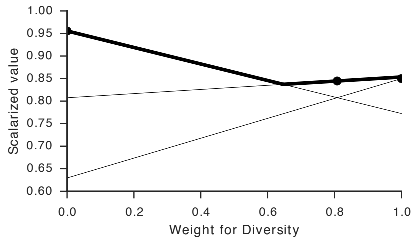


Figure 4: The CCS found on the TREC 2010 and 2011 datasets.

We demonstrated how to optimize rankings for multiple objectives. Using this approach, we have found multiple optimal rankers on the CLEF eHealth 2015 task 2 and on the TREC diversity task that offer different trade-offs w.r.t. different relevance criteria. These multiple optimal rankers are more flexible than a one-size-fits-all ranker produced by a standard learning to rank approach, and our work therefore forms an important step for flexibly optimizing search when multiple criteria are in play.

# References


[1] Ben Carterette. System effectiveness, user models, and user utility: a conceptual framework for investigation. In *SIGIR'11*, pages 903–912. ACM, 2011.

[2] Helena R. Lourenço, Olivier C. Martin, and Thomas Stützle. *Iterated local search*. Springer, 2003.

[3] Joao Palotti, Guido Zuccon, Lorraine Goeuriot, Liadh Kelly, Allan Hanbury, Gareth JF Jones, Mihai Lupu, and Pavel Pecina. CLEF eHealth evaluation lab 2015, task 2: Retrieving information about medical symptoms. In *CLEF '15*. Springer, 2015.

[4] Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.

[5] Diederik M. Roijers, Shimon Whiteson, and Frans A. Oliehoek. Computing convex coverage sets for faster multi-objective coordination. *Journal of Artificial Intelligence Research*, 52:399–443, 2015.

[6] Diederik M Roijers, Shimon Whiteson, and Frans A Oliehoek. Point-based planning for multi-objective POMDPs. In *IJCAI'15*, 2015.

[7] Linda Schamber and Judy Bateman. User criteria in relevance evaluation: Toward development of a measurement scale. In *ASIS'96*, volume 33, pages 218–25. ERIC, 1996.

[8] Joost van Doorn, Daan Odijk, Diederik M. Roijers, and Maarten de Rijke. Balancing relevance criteria through multi-objective optimization. SIGIR '16, pages 769–772. ACM, 2016.

[9] Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *ICML'09*, pages 1201–1208. ACM, 2009.