A Re-Ranking Method Based on Irrelevant Documents in Ad-Hoc Retrieval

Rabeb Mbarek[1], Mohamed Tmar[2], Hawete Hattab[3], and Mohand Boughanem[4]

1 **Sfax University, Multimedia Information Systems and Advanced Computing Laboratory, Sfax, Tunisia**
rabeb.hattab@gmail.com

2 **Sfax University, Multimedia Information Systems and Advanced Computing Laboratory, Sfax, Tunisia**
mohamedtmar@yahoo.fr

3 **Umm Al-qura University, Department of Mathematics, Makkah, KSA**
hshattab@uqu.edu.sa

4 **University of Toulouse – IRIT lab France, Toulouse, France**
bougha@irit.fr

Abstract

In this paper, we propose a novel approach for document re-ranking, which relies on the concept of negative feedback represented by irrelevant documents. In a previous paper, a pseudo-relevance feedback method is introduced using an absorbing document $\widetilde{d}$ which best fits the user's need. The document $\widetilde{d}$ is orthogonal to the majority of irrelevant documents. In this paper, this document is used to re-rank the initial set of ranked documents in Ad-hoc retrieval. The evaluation carried out on a standard document collection shows the effectiveness of the proposed approach.

1998 ACM Subject Classification H.3.3 Information Search and Retrieval

Keywords and phrases Re-ranking, absorption of irrelevance, vector product

Digital Object Identifier 10.4230/OASIcs.SLATE.2016.2

1 Introduction

A commonly used strategy to improve search results is through feedback techniques, including relevance feedback [14, 15, 16], pseudo-relevance feedback [2, 5, 21] and implicit feedback [17]. A query is difficult if none of the top-ranked documents are relevant. In the case of difficult queries, if we can perform effective negative feedback when a user could not find any relevant document on the first page of the search results, we would be able to improve the ranking of the unseen results in the next few pages. It is clear that in this case of negative relevance feedback, we only have negative (i.e., irrelevant) documents. When a user is unable to reformulate an effective query (which happens often in informational queries due to insufficient knowledge about the relevant documents), negative feedback can be quite beneficial, and the benefit can be achieved without requiring extra effort from users (e.g., by assuming the skipped documents by a user to be irrelevant).

This work investigates the role of irrelevant documents in document re-ranking. In particular, our re-ranking strategy is based on a negative relevance feedback approach which takes into account irrelevant documents in the initial document ranking. The key idea behind our approach is to use the absorbing document [11], which fits the user's need and is orthogonal to the majority of irrelevant documents, to re-rank documents on the ground of their similarity with respect to the absorbing document. Generally, standard relevance feedback methods are able to handle negative feedback by subtracting information from the original query (for example the Rocchio's model [15]). The key issue of this approach is to quantify the side effect caused by information loss. To deal with this effect, we propose a negative feedback method based on absorbing document that is able to remove only the unwanted aspects pertaining to irrelevant documents. In our approach, documents are represented as vectors in a geometric space in which similar documents are represented close to each other. This space is the classical Vector Space Model (VSM).

We compare our strategy with other approaches. First, with the Baseline Model (the BM25 model [13]). Second, with the approach of Basile et al. [3].

How to identify irrelevant documents is an open question. We use two distinct approaches in our work proposed in [3]: the former exploits documents at the bottom of the rank, while the latter takes the irrelevant documents directly from relevance judgments. These approaches are thoroughly described in Section 3.

The paper is structured as follows. Related work are briefly analyzed in Section 2. Section 3 describes the two strategies used for re-ranking. Experiments performed for evaluating our approach are presented in Section 4. The last section concludes.

2 Related Work

There exist several groups of related work in the areas of document retrieval and re-ranking.

The first category performs re-ranking by using inter-document relationship [6, 7]. The idea is to build a document which represents the ideal response to the user's information need. In [6] documents in the result list are re-weighed according to a relevance function which reflects the distance between documents and the ideal document. Other researchers use inter-document similarities to combine several retrieved lists (see for example [7]). In this case, the idea of similarity is used to give support to documents with similar content highly ranked across multiple result lists.

A second category of work is related to recent advances in structural re-ranking paradigm over graphs. In the language modeling framework, the traditional cluster-based retrieval has been juxtaposed with document language model smoothing in which document representation incorporates cluster-related information [8, 9, 10].

An early attempt to model terms negation in pseudo-relevance feedback by quantum logic operators is due to Widdows [20]. In his work, Widdows has shown that negation in quantum logic is able to remove, from the result set, not only unwanted terms but also their related meaning. The concept of vectors orthogonality is exploited to express queries like *Retrieve documents that contain term A & NOT term B*. Widdows suggested that vectors which represent unrelated concepts should be orthogonal to each other. Indeed, orthogonality prevents vectors from sharing common features.

In [3], Basile et al. proposed a new re-ranking strategy based on a pseudo-relevance feedback approach which took into account both relevant and irrelevant documents in the initial document ranking. The key idea of this approach is to build an ideal document which fits the user's need, and then re-rank documents on the ground of their similarity with respect to the ideal document. The ideal document d^* is built using a geometrical space where d^* is computed as a vector close to relevant documents and unrelated to irrelevant ones. In this space the concept of relevance is expressed in terms of similarity, while the concept of irrelevance is defined by orthogonality (similarity equals to zero). Formally, Basile et al. [3]

computed the ideal document by the following logical operation:

$$d^* = d_1^+ \vee d_2^+ \vee ...d_n^+ \wedge NOT(d_1^- \vee d_2^- \vee ... \vee d_m^-) \quad (1)$$

where $D^+ = \{d_1^+, d_2^+, ..., d_n^+\}$ and $D^- = \{d_1^-, d_2^-, ..., d_m^-\}$ are the subsets of relevant and irrelevant documents respectively. Equation 1 consists in computing a vector which represents the disjunction of the documents in D^+, and then projecting this vector onto the orthogonal spaces generated by the documents in D^-. Disjunction and negation using quantum logic are thoroughly described in [20]. An overview of Quantum Mechanics for Information Retrieval can be found in [4]. The main problem of the approach of Basile et al. is the query drift problem related to the pseudo-relevance feedback approach. Query drift occurs when the documents used for relevance feedback contain few or no relevant documents.

In this paper the orthogonality is defined using the algebraic operator vector product[1]. Using this operator, we build an absorbing document which is orthogonal to the majority of irrelevant documents.

The idea to build a document which represents the response to the user's information need is not new. In [6] documents in the result list are re-weighed according to a relevance function which reflects the distance between documents and the "ideal document".

Whilst relevant documents have been successfully used in several approaches to improve Information Retrieval performance, irrelevant ones seem not to arouse researchers' interest. Singhal et al. [18] achieved an interesting result for the learning routing query problem: they showed that using irrelevant documents close to the query, in place of those in the whole collection, is more effective. Rocchio's original formulation explicitly includes a component of irrelevant documents [15]. In [12, 11], the authors showed that irrelevant documents can be used to extract better expansion terms from the top-ranking k documents. A successful use of irrelevant documents for negative pseudo-relevance feedback has been carried out in [19], where authors point out the effectiveness of their approach with poorly performing queries.

3 A Re-ranking Method Based on Irrelevant Documents

This section describes our re-ranking strategy based on irrelevant documents. The main idea is to build a document vector which attempts to model the absorbing document in response to a user query, and then exploit this vector to re-rank the initial set of ranked documents D_{init}. The absorbing document $\widetilde{d}$ should be orthogonal with each document in the set D^- of irrelevant ones. Identifying relevant documents is quite straightforward: we assume the top ranked documents in D_{init} as relevant, whereas identifying non-relevant documents is not trivial. To this purpose, we propose two strategies: the former relies on documents at the bottom of D_{init}, while the latter needs relevance judgments. The absorbing document vector $\widetilde{d}$ is exploited to re-rank documents in D_{init} on the ground of the similarity between $\widetilde{d}$ and each document in D_{init} in the Euclidean space (vector space equipped with an inner product).

3.1 Vector product

Let E be a vector space of dimension n and let u_1,..., u_{n-1} be $n-1$ vectors of E. For each vector x of E there exists a unique vector w such that:

$$det(u_1, ..., u_{n-1}, x) = w^T.x$$

[1] This operator, in a vector space, naturally models the orthogonality.

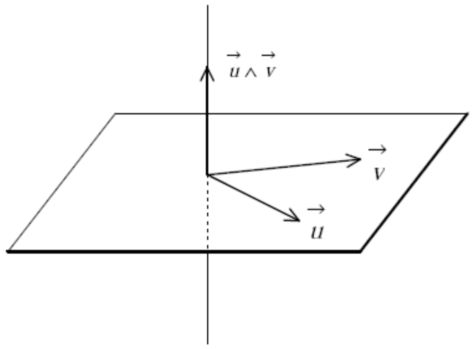

Figure 1 The cross product for $n = 3$.

where det is the determinant of n vectors, w^T is the transpose of w and $w^T.x$ is the classical inner product.

w is called *the vector product of* $u_1,..., u_{n-1}$ and is denoted by $u_1 \wedge ... \wedge u_{n-1}$ (for $n = 3$, see Figure 1). We have the following properties:

- the vector $u_1 \wedge ... \wedge u_{n-1}$ is orthogonal to each vector u_i.
- the vector $u_1 \wedge ... \wedge u_{n-1}$ is orthogonal to the subspace F of E generated by the family $(u_1,..., u_{n-1})$. Indeed, if u is a vector of F, there exists $n-1$ scalars $\alpha_1,..., \alpha_{n-1}$ such that $u = \alpha_1 u_1 + ... + \alpha_n u_{n-1}$.
- $u_1 \wedge ... \wedge u_{n-1} = \overrightarrow{0}$ if and only if $u_1,..., u_{n-1}$ are dependent.
- if $u_1,..., u_{n-1}$ are independent then $(u_1,..., u_{n-1}, u_1 \wedge ... \wedge u_{n-1})$ is a basis of E.

3.2 Scenario

Let n be the dimension of D_{init} as a vector space, n represents the number of indexing terms. Let $m < n$ be the number of linearly independent and representative documents of D^-, and let $u_1,...,u_m$ be these irrelevant ones. We eliminate $n - m - 1$ terms and so the dimension becomes $m + 1$. The *absorbing document* is:

$$\widetilde{d} = u_1 \wedge ... \wedge u_m \tag{2}$$

This document is orthogonal to the majority of irrelevant documents.

3.3 Compute of the absorbing document

To compute $\widetilde{d}$ it suffices to compute the vector product of $u_1,...,u_m$. Let $A = (u_1, ..., u_m)$ be the matrix of $m+1$ rows and m columns. Let A_i be the matrix obtained from the matrix A by deleting the ith row $(1 \leq i \leq m + 1)$. The vector product of $U_1,...., U_m$ is the vector:

$$u_1 \wedge ... \wedge u_m = \begin{pmatrix} detA_1 \\ -detA_2 \\ ... \\ ... \\ (-1)^m detA_{m+1} \end{pmatrix} \tag{3}$$

The Equation 3 generalizes the definition of vector product of two vectors in dimension 3. In the following, we give an example of vector product of three vectors in dimension 4:

if $u_1 = (1,0,1,-1)^T$, $u_2 = (0,2,1,1)^T$ and $u_3 = (1,3,1,0)^T$ are three vectors, then $u_1 \wedge u_2 \wedge u_3 = (4,-1,-1,3)^T$ and so $(4,-1,-1,3).(1,0,1,-1)^T = (4,-1,-1,3).(0,2,1,1)^T = (4,-1,-1,3).(1,3,1,0)^T = 0$.

3.4 An illustrative example

In this example we show how the absorbing document $\widetilde{d}$ help us to extract better expansion terms.

We consider four linearly independent irrelevant documents d_1, d_2, d_3, and d_4, selected from the bottom of the initial ranking of topic 351. These four irrelevant documents indexed by 5 expansion terms t_1, t_2, t_3, t_4 and t_5, selected from the 2-top relevant documents.

$$d_1 = (2,1,1,0,0)^T \quad d_2 = (1,0,2,0,0)^T \quad d_3 = (4,0,2,0,0)^T \quad d_4 = (0,1,0,2,1)^T \,.$$

The absorbing document $\widetilde{d}$ is the cross product of d_1, d_2, d_3, and d_4:

$$\widetilde{d} = (2,1,1,0,0)^T \wedge (1,0,2,0,0)^T \wedge (4,0,2,0,0)^T \wedge (0,1,0,2,1)^T = (0,0,0,-6,12)^T \,.$$

$\widetilde{d}$ is indexed by the terms t_4 and t_5. Note that d_4 is the only irrelevant document which is indexed by t_4 and t_5.

3.5 Strategies to select irrelevant documents

We use the two strategies proposed in [3] to select the set (D^-) of irrelevant documents:

- BOTTOM, which selects the irrelevant documents from the bottom of the rank; in other words we assume that the user selects the last m linearly independent irrelevant documents;
- RELJUD, which relies on relevance judgments provided by CLEF organizers. This technique selects the top m ranked documents which are irrelevant exploiting the relevance judgments. We use this strategy to simulate the user's explicit feedback; in other words we assume that the user selects the first m linearly independent irrelevant documents.

To select linearly independent irrelevant documents we use the Algorithm 1.

4 Experiments

In this section we give the different experiments and results obtained to evaluate our approach. The goal of the evaluation is to prove that our re-ranking strategy, which relies on the concept of negative feedback represented by irrelevant documents, improves retrieval performance and outperforms other methods. Moreover, we want to evaluate the performance of the BOTTOM strategy and RELJUD strategy.

4.1 Environnement

We set up a baseline system based on the BM25 multi-fields model [13]. The evaluation has been designed using the CLEF 2009 Ad-hoc WSD Robust Task collection [1]. The Robust task allows us to evaluate Information Retrieval System performance even when difficult queries are involved. The CLEF 2009 collection consists of 166,717 documents which have two fields: HEADLINE and TEXT. Table 1 shows the BM25 parameters, where b is a constant related to the field length, k_1 is a free parameter, and boost is the boosting factor applied to that field.

Listing 1 The set of linearly independent irrelevant documents.

```
Let n be the number of terms
Let A be the n × mm matrix of irrelevant documents
Let m be the rank of A
Let B be the n × m matrix of linearly independent irrelevant documents

for i = 1, ..., n
      b_{i,1} ← a_{i,1}
end for
k ← 1
for j = 2, ..., mm
      Let C be a vector
      for i = 1, ..., n
            c_i ← a_{i,j}
      end for
      for l = 1, ..., n
            b_{l,k} ← c_l
      end for

      Let p be the rank of B
      if p = (k + 1)
            k ← k + 1
      end if
      if k ← n
            break
      end if
end for

return B
```

Table 1 BM25 parameters used in the experiments.

Field	k_1	b	boost
HEADLINE	3.25	0.7	2
TEXT	3.25	0.7	1

In detail, the CLEF 2009 collection has 150 topics. Topics are structured in three fields: TITLE, DESCRIPTION and NARRATIVE. We used only TITLE and DESCRIPTION, because NARRATIVE field is the topic description used by assessors. Moreover, we used different boosting factors for each topic field (TITLE=4 and DESCRIPTION=1) to highlight terms in the TITLE.

For our approach, the experiments consist to re-rank documents (results of the baseline approach) on the ground of their similarity with respect to the absorbing document $\widetilde{d}$ (Equation 2). The retrieved documents are ranked by the inner product done by:

$$< \widetilde{d}, d >= \widetilde{d}^T.d \tag{4}$$

To evaluate the performance of our approach, we executed several runs using the topics provided by CLEF organizers. In particular, we took into account: m (the cardinality of D^-). We selected different ranges for parameter m: $[1, 5, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100]$.

For the approach of [3], the experiments consist to re-rank documents (results of the baseline approach) on the ground of their similarity with respect to the ideal document d^*

(Equation 1). The retrieved documents are ranked by the relevance score computed for each document d in D_{init} done by:

$$S(d) = \alpha * S_{D_{init}}(d) + (1 - \alpha).sim(d, d^*)$$

where $S_{D_{init}}(d)$ is the score of d in the initial rank D_{init}, while $sim(d, d^*)$ is the similarity degree between the document vector d and the ideal document vector d^* computed by cosine similarity.

To evaluate the performance of their approach, Basile et al. [3] executed several runs using the topics provided by CLEF organizers. In particular, they took into account: n (the cardinality of D^+), m (the cardinality of D^-) and the parameter α used for the linear combination of the scores. They selected different ranges for each parameter: n ranges in $[1, 5, 10, 20, 40]$, m ranges in $[0, 1, 5, 10, 20, 40]$, while α ranges in $[0.3, 0.4, 0.5, 0.6, 0.7]$. Table 2. shows the best five runs for BOTTOM and RELJUD strategies with respect to MAP and GMAP. For the both approaches, they set the cardinality of D_{init} to 1000. All the metrics have been computed on the first 1000 returned documents, as prescribed by the CLEF evaluation campaign.

4.2 Results

The experiments and the evaluations are as follow. Comparison between the Baseline Model (the BM25 multi-fields model [13]), the approach of Basile et al. [3], and our approach: re-ranking method using absorbing document (Equation 4), using Mean Average Precision (MAP) and Geometric Mean Average Precision (GMAP) over all the queries.

The results have been grouped by the number of irrelevant documents. Table 2 reports the results of the Baseline Model and the best performance obtained for the approach of Basile et al. [3] (the best five runs for BOTTOM and RELJUD strategies with respect to MAP values). Moreover, this table illustrates the best performance obtained for our approach (the best five runs for BOTTOM and RELJUD strategies where the number of irrelevant documents ranges in $[1, 5, 10, 20, 30, 40, 50, 60, 70]$). Improvements in percentage $\Delta\%$ with respect to the baseline are reported for MAP and GMAP values.

4.3 Analysis of results

Generally, BOTTOM strategy results are not significant improvements. This suggests that the BOTTOM strategy is not able to identify irrelevant documents. For this strategy, the highest MAP value for our approach is 0.476 (GMAP=0.234). Both values (MAP and GMAP) are obtained with 30 irrelevant documents. For the approach of Basile et al., The highest MAP value is 0.4384 (GMAP=0.1928). The MAP value is obtained with five irrelevant documents, while the GMAP is obtained with one irrelevant document.

The method RELJUD obtains very high results. For this strategy, The highest MAP value for our approach is 0.691 (GMAP=0.3328). Both values (MAP and GMAP) are obtained with 70 irrelevant documents. For the approach of Basile et al., The highest MAP value is 0.6649 (GMAP=0.3240). Both values (MAP and GMAP) are obtained with 40 irrelevant documents

For our approach, the performance of the two strategies (BOTTOM and RELJUD) increases if the number of irrelevant documents increases.

The experimental results are very encouraging. For our approach, both methods (BOTTOM and RELJUD) show improvements with respect to the baseline in all the approaches. The comparison between the results of our approach with the use of the two strategies

Table 2 Comparison between our approach, the baseline, and the approach of Basile et al.

Approach	Method	Run	n	m	α	MAP	Δ%	GMAP	Δ%
-	-	baseline	-	-	-	0.4139	-	0.1846	-
Basile et al.	BOTTOM	2.B1	1	5	0.6	0.4384	+5.92	0.1923	+4.17
		2.B2	1	10	0.6	0.4379	+5.80	0.1921	+4.06
		2.B3	1	1	0.5	0.4377	+5.75	0.1928	+4.44
		2.B4	1	5	0.5	0.4376	+5.73	0.1926	+4.33
		2.B5	1	20	0.6	0.4372	+5.73	0.1917	+3.85
	RELJUD	2.R1	40	40	0.7	0.6649	+60.64	0.3240	+75.51
		2.R2	40	40	0.6	0.6470	+56.32	0.3156	+70.96
		2.R3	40	40	0.5	0.6223	+50.35	0.3124	+69.23
		2.R4	20	40	0.7	0.6176	+49.21	0.2859	+54.88
		2.R5	20	20	0.7	0.6107	+47.55	0.2836	+53.63
Our approach	BOTTOM	B1	-	1	-	0.4	−3.36	0.17	−7.9
		B2	-	5	-	0.419	+1.23	0.185	+0.21
		B3	-	10	-	0.423	+2.2	0.191	+3.46
		B4	-	20	-	0.442	+6.78	0.212	+14.84
		B5	-	30	-	0.476	+15	0.234	+25.89
	RELJUD	R1	-	20	-	0.601	+45.2	0.272	+47.34
		R2	-	40	-	0.671	+62.11	0.331	+79.3
		R3	-	50	-	0.675	+63.08	0.3325	+80.11
		R4	-	60	-	0.687	+65.98	0.3327	+80.22
		R5	-	70	-	0.691	+66.94	0.3328	+80.28

(BOTTOM and RELJUD), the results of the classic BM25 model, and the results of Basile et al., shows that our approach improves the results of the two other approaches.

5 Conclusion and future work

This paper proposes a novel approach based on negative evidence for document re-ranking. The novelty lies on the use of the absorbing document to capture the negative aspects of irrelevant documents. This method has shown its effectiveness with respect to a baseline system based on BM25 and a re-ranking method based on the approach of Basile et al. [3]. Moreover, the evaluation has proved the robustness of the proposed strategy and its capability to absorb irrelevant documents. On the other hand our approach depends on a single parameter, while the other re-ranking approaches depend on many parameters. Moreover, the absorbing document is modelled by a vector product which is simply computed in a vector space model.

In a future work, we will apply this re-ranking approach with respect to a vector space basis which optimally separates relevant and irrelevant documents.

References

1 Eneko Agirre, Giorgio Maria Di Nunzio, Thomas Mandl, and Arantxa Otegi. CLEF 2009 ad hoc track overview: Robust-WSD task. In Carol Peters, Giorgio Maria Di Nunzio, Mikko Kurimo, Thomas Mandl, Djamel Mostefa, Anselmo Peñas, and Giovanna Roda, editors,

Multilingual Information Access Evaluation I, pages 36–49. Springer, Berlin, Heidelberg, 2010. doi:10.1007/978-3-642-15754-7_3.

2 Rony Attar, Fraenkel, and Aviezri Siegmund. Local feedback in full-text retrieval systems. *Journal of the ACM*, 24(3):397–417, July 1977. doi:10.1145/322017.322021.

3 Pierpaolo Basile, Annalina Caputo, and Giovanni Semeraro. Negation for document re-ranking in ad-hoc retrieval. In Giambattista Amati and Fabio Crestani, editors, *Advances in Information Retrieval Theory*, pages 285–296. Springer, Berlin, Heidelberg, 2011. doi:10.1007/978-3-642-23318-0_26.

4 Garrett Birkhoff and John Von Neumann. The logic of quantum mechanics. *Annals of Mathematics*, 37(4):823–843, 1936.

5 Chris Buckley, Gerard Salton, James Allan, and Amit Singhal. Automatic query expansion using SMART: TREC 3. In *In Proceedings of The third Text REtrieval Conference (TREC-3)*, pages 69–80, 1994.

6 Czesław Daniłowicz and Jarosław Baliński. Document ranking based upon Markov chains. *Information Processing and Management*, 37(4):623–637, July 2001. doi:10.1016/S0306-4573(00)00038-8.

7 Anna Khudyak Kozorovitzky and Oren Kurland. From "identical" to "similar": Fusing retrieved lists based on inter-document similarities. In Leif Azzopardi, Gabriella Kazai, Stephen Robertson, Stefan Rüger, Milad Shokouhi, Dawei Song, and Emine Yilmaz, editors, *Advances in Information Retrieval Theory*, pages 212–223. Springer, Berlin, Heidelberg, 2009. doi:10.1007/978-3-642-04417-5_19.

8 Oren Kurland. Re-ranking search results using language models of query-specific clusters. *Information Retrieval*, 12(4):437–460, 2009. doi:10.1007/s10791-008-9065-9.

9 Oren Kurland and Lillian Lee. Corpus structure, language models, and ad hoc information retrieval. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 194–201, New York, NY, USA, 2004. ACM. doi:10.1145/1008992.1009027.

10 Xiaoyong Liu and W. Bruce Croft. Cluster-based retrieval using language models. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 186–193, New York, NY, USA, 2004. ACM. doi:10.1145/1008992.1009026.

11 Rabeb Mbarek, Mohamed, Hawete Hattab, and Mohand Boughanem. Pseudo-relevance feedback method based on the cross-product of irrelevant documents. *International Journal Web Applications*, 8(1):8–16, March 2016.

12 Karthik Raman, Raghavendra Udupa, Pushpak Bhattacharya, and Abhijit Bhole. On improving pseudo-relevance feedback using pseudo-irrelevant documents. In *Proceedings of the 32nd European Conference on Advances in Information Retrieval*, pages 573–576, Berlin, Heidelberg, 2010. Springer-Verlag. doi:10.1007/978-3-642-12275-0_50.

13 Stephen Robertson, Hugo Zaragoza, and Michael Taylor. Simple BM25 extension to multiple weighted fields. In *Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management*, pages 42–49, New York, NY, USA, 2004. ACM. doi:10.1145/1031171.1031181.

14 Stephen E. Robertson and Karen Sparck Jones. Relevance weighting of search terms. *Journal of the American Society for Information Science*, 27(3):129–146, 1976. doi:10.1002/asi.4630270302.

15 Joseph J. Rocchio. Relevance feedback in information retrieval. In Gerard Salton, editor, *The SMART retrieval system - experiments in automatic document processing*, pages 313–323. Englewood Cliffs, NJ: Prentice-Hall, 1971.

16 Gerard Salton and Chris Buckley. Improving retrieval performance by relevance feedback. In Karen Sparck Jones and Peter Willett, editors, *Readings in Information Retrieval*, pages 355–364. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1997.

17 Xuehua Shen, Bin Tan, and ChengXiang Zhai. Context-sensitive information retrieval using implicit feedback. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 43–50, New York, NY, USA, 2005. ACM. doi:10.1145/1076034.1076045.

18 Amit Singhal, Mandar Mitra, and Chris Buckley. Learning routing queries in a query zone. In *Proceedings of the 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 25–32, New York, NY, USA, 1997. ACM. doi:10.1145/258525.258530.

19 Xuanhui Wang, Hui Fang, and ChengXiang Zhai. A study of methods for negative relevance feedback. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 219–226, New York, NY, USA, 2008. ACM. doi:10.1145/1390334.1390374.

20 Dominic Widdows. Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics*, volume 1, pages 136–143, Stroudsburg, PA, USA, 2003. Association for Computational Linguistics. doi:10.3115/1075096.1075114.

21 Jinxi Xu and W. Bruce Croft. Query expansion using local and global document analysis. In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 4–11, New York, NY, USA, 1996. ACM. doi:10.1145/243199.243202.