The University of Amsterdam at the TREC 2007 Blog Track

Breyten Ernsting **Wouter Weerkamp** **Maarten de Rijke**

ISLA, University of Amsterdam
http://ilps.science.uva.nl/

Abstract: We describe our participation in the TREC 2007 Blog track. In the opinion task we looked at the differences in performance between Indri and our mixture model, the influence of external expansion and document priors to improve opinion finding; results show that an out-of-the-box Indri implementation outperforms our mixture model, and that external expansion on a news corpus is very benificial. Opinion finding can be improved using either lexicons or number of comments as document priors. In our approach to the feed distillation task we integrated time-based and frequency aspects into the retrieval model; we find that time-based retrieval improves results slightly, while frequency-based retrieval results in substantial improvements under the right circumstances.

1 Introduction

We describe our experiments for this year's edition of the Blog track. Our main aims were (1) to compare our inhouse mixture model against an Indri-based baseline for topical blog post retrieval, and (2) for the distillation task to examine the influence of time and frequency on retrieval effectiveness. In two largely independent sections we first discuss our work on the opinion finding task (Section 2) and then our work on the feed distillation task (Section 3). We conclude in Section 4.

2 Opinion finding

The opinion finding task aims at returning blog posts that contain an opinion regarding a certain topic [6]. The results of last year's opinion finding task indicate that a strong topical retrieval system is the single most important part of opinion finding. In Section 2.1 we present the models we use for topical retrieval and the usage of external expansion to improve topical blog post retrieval is discussed in Section 2.2. Section 2.3 shows the implementation of opinion finding indicators. Finally we present run details (Section 2.5) and results (Section 2.6).

2.1 Topic retrieval

Our baseline approach to the topical blog post retrieval task uses language models. The models we use in this track are similar to the models we use in the TREC Enterprise track and are described more fully in [1].

2.1.1 Indri

For comparison we use an out-of-the-box implementation of Indri 2.4[1] to process the 2007 topics. As preprocessing steps we strip all HTML tags and remove stopwords; we do not apply stemming.

2.2 External expansion

Our baseline query model $p(t|q)$ is estimated using Eq. 1:

$$p(t|q) = n(t,q) \cdot |q|^{-1}. \quad (1)$$

To improve topic retrieval performance we use relevance models [2]; the relevance models are constructed using feedback documents and return feedback terms with an associated weight. Instead of constructing the relevance models based on the top x documents of the blog collection we use insights from [5] stating that many queries issued on blog search engines are in fact news related. We use the (contemporaneous) AQUAINT-2 news corpus to construct relevance models; the top 10 terms are selected and their weights are normalized so that their sum equals 1. The weighted query is issued against the blog collection to retrieve the final set of blog posts.

2.3 Document priors

On top of our topical retrieval method, we implement opinion finding methods using query independent document priors ($p(d)$). We believe blog posts have a degree of opinionatedness regardless of the topic; besides, this approach does not have a negative impact on the speed of the retrieval system.

The main issue, then, is how to estimate the document priors. We compare two document priors for opinionatedness: (1) a lexical approach and (2) a comment-based approach.

[1] http://www.lemurproject.org/indri/

The lexical approach uses a list of opinionated words from the OpinionFinder system.[2] From this list we extract only the strong positive and strong negative words. The document prior is then estimated as given in Eq. 2:

$$p(d) = \sum_{i=1}^{w} c(w_i, d) \cdot |d|^{-1}, \quad (2)$$

where w is the list of opinionated words and $c(w_i, d)$ is the count of opinion word i in document d and $|d|$ is the document length in words.

Our second, comment-based approach is based on the intuition that opinionated blog posts are more likely to attract discussion. When a post contains an explicit point of view on a topic, readers are more likely to comment (either by agreeing or by expressing their own opinion on the topic). Using this idea, we estimate document priors as follows:

$$p(d) = \log(N_{comments,d}), \quad (3)$$

where $N_{comments,d}$ is the number of comments in document d. We use regular expressions like `\d+ (comment | response)` on the blog posts to identify the number of comments.

2.4 Polarity

New in this year's opinion task is the polarity subtask: given an opinionated post we need to identify whether it is either negative, positive or neutral. To address this task we experiment with two approaches. The first approach continues on the opinionated words list of the previous section. We use the two opinionated lists separately (positive and negative words) and use Eq. 4 to estimate the polarity of each post:

$$pol(d) = \begin{cases} positive & \text{if } r(d) > 0.01 \\ negative & \text{if } r(d) < -0.01 \\ neutral & \text{otherwise,} \end{cases} \quad (4)$$

where $r(d)$ is defined as

$$r(d) = (\sum_{i=1}^{n} c(n_i, d) - \sum_{i=1}^{p} c(p_i, d)) \cdot |d|^{-1}, \quad (5)$$

in which n is the list of negative words, p the list of positive words, $c(n_i, d)$ the number of times word n_i occurs in document d and $|d|$ the document length in words.

For our second approach we look at expressions of opinion other than words. The idea is that posts containing more expressive language tend to be more negative about the topic discussed in that post. To estimate this negativity, we use the following indicators: exclamation marks, question marks, ellipsis (...), and all caps strings of more than 3 characters. The ratio of these indicators is calculated for each blog post using Eq. 6, where $c(i, d)$ is the total number of occurrences of the above indicators in document d and $|d|$ is the document length in words. Polarity is estimated according Eq. 7:

$$r(d) = c(i, d) \cdot |d|^{-1} \quad (6)$$

$$pol(d) = \begin{cases} positive & \text{if } r(d) < 0.1 \\ negative & \text{if } r(d) \geq 0.1 \end{cases} \quad (7)$$

[2]URL: http://www.cs.pitt.edu/mpqa/

Table 1: Blog post retrieval results

	topic		opinion		polarity
run id	MAP	p@10	MAP	p@10	R-acc.
A	0.4342	0.6800	0.3281	0.4920	
C	0.3105	0.6060	0.2123	0.3840	0.1284
B	**0.4741**	**0.7620**	**0.3453**	**0.5620**	**0.1827**
D	0.1898	0.4500	0.1273	0.2520	0.0711
F	0.1865	0.4720	0.1459	0.3200	0.0840
E	0.1834	0.4620	0.1302	0.2900	0.0677

2.5 Runs

For the runs using the mixture model we use the 2006 topics and assessments to estimate the best mixture weights. Results show that only the title component has a positive influence on retrieval performance and the best mixture is estimated to be 0.15 title, 0.60 body text and 0.25 background.

(A) **`uams07indbl`** the baseline run uses an out-of-the-box Indri implementation.

(B) **`uams07topic`** the topic run also uses Indri out-of-the-box; queries are first rewritten to improve retrieval [4] and then expanded using the external corpus (as described in Section 2.2).

(C) **`uams07mmbl`** the baseline mixture model run: the best mixture as tested on the 2006 data without additional features.

(D) **`uams07mmq`** identical to the previous run; instead of the baseline query model relevance models are used based on feedback on a news corpus.

(E) **`uams07mmqcom`** identical to the previous run; to identify opinionated posts document priors based on the number of comments are included as discussed in Section 2.3

(F) **`uams07mmqop`** identical to run uams07mmq; document priors based on the ratio of opinionated words (Section 2.3) are used to estimate opinionatedness.

(G) **`polarity`** runs uams07topic, uams07mmbl and uams07mmqop use the opinionated words ratio as polarity indicator. Runs uams07mmq and uams07mmqcom use the punctuation-based polarity identifier (see Section 2.4).

2.6 Results

We look at the performance of our runs on topical retrieval, opinion retrieval and polarity identification. Table 1 shows the MAP and p@10 scores for all runs on the retrieval tasks and the R-accuracy on the polarity identification. The R-accuracy is the fraction of retrieved documents above rank R that are classified correctly, where R is the number of opinion-containing documents for that topic.

From Table 1 it is clear that run B (uams07topic, Indri with external expansion) performs best on all tasks.

Surprisingly, the mixture model runs with external expansion perform significantly worse than the baseline run, even though we see an improvement of external expansion on the Indri runs; it appears that the external expansion is not performed correctly, leading to expanded queries that miss original query terms. Topic 902 (*lactose gas*) provides an example: the expanded query contains the terms gas, contamination, water, underground, solution, tce, eaten, whey, study, atoms, but the original query term lactose is missing. The performance decreases dramatically from 0.4127 in the baseline run to 0.003 in the expanded runs. This error occurs in about half of the topics, making the final run scores (both on MAP and precision) much lower. An example of a topic that is expanded correctly, is topic 947 (*sasha cohen*). It achieves an AP score of 0.2897 in the baseline run, which increases to 0.6371 in the expanded runs. Similar effects are noticeable in other correctly expanded queries.

Both the lexicon and comment-based approach have a positive influence on opinion retrieval, with the lexicon approach outperforming the comment-based approach. Finally, polarity detection based on the difference between positive word ratio and negative word ratio performs only slightly better than the punctuation approach, even though the latter distinguishes only positive and negative posts and ignores neutral posts.

3 Feed distillation

The feed distillation task is a new task and it aims at finding feeds that are devoted to a given topic. The general idea is that a user can be presented with a suggested list of feeds that are worth reading, given the topic. Although the task is new, it resembles the Expert Search task in the Enterprise track, and the Topic Distillation task in the Web track.

Below, we present the models we used for topic distillation, and for incorporating aspects of time in the retrieval model, to improve the accuracy of the topic distillation.

3.1 Topic retrieval

The method we use to address topic retrieval for the feed distillation task is based on ranking individual posts contained in the feed; it is akin to the method in Section 2.1 (Eq. 1–5). However, $p(t|\theta_d)$ is estimated simply as follows:

$$p(t|\theta_d) = p(t|d), \tag{8}$$

where $p(t)$ is the maximum likelihood-estimate of the term t in the document collection.

3.2 Time-based retrieval

To improve the accuracy of our topical retrieval system, we incorporate query independent document priors which are based on the creation date of the documents. More recent documents are assumed to better reflect the current interests of a feed (blogger), and that these should therefore rank higher. We model this intuition using a time-based language model [3]:

$$p(d|q) \propto p(q|d)p(d|T_d), \tag{9}$$

where $p(d|T_d)$ is a time-based prior in the model. Since we're interest in *recency* and not a specific *event*, we use an exponential distribution to calculate the priors. This distribution is defined as follows:

$$p(d|T_d) = P(T_d) = \lambda e^{-\lambda(T_C - T_d)} \tag{10}$$

The optimal value of the parameter λ is determined in training experiments. T_C and T_d are measured in days; T_C signifies the most recent date of the documents in the collection, and T_d refers to the document being considered.

Using time-based document priors only works when using blog posts as the unit of retrieval. In order to derive a ranked list of feeds (as required) from a ranked list of blog posts, we take the score of the highest ranked blog post of a feed:

$$score(bl|q) = \max_{d \in bl} p(d|q), \tag{11}$$

where bl is a feed, and d ranges over blog posts in bl.

3.3 Frequency-based retrieval

Another way to improve the accuracy of our retrieval system is to consider the frequency of on-topic blog posts for feeds. We assume that the number of matching blog posts is proportional to the interest of a blogger in the topic. Consequently, we can derive a ranking of feeds according to the frequency of blog posts matching the topic in question. However, this ignores the fact that some feeds may contain a disproportionate number of blog posts, in comparison to other feeds. We therefore consider the number of on-topic blog posts in the feed versus the number of blog posts, either matching the topic or not, to score the feeds.

We incorprate this frequency-based score in our retrieval model using a linear combination of both the topic-based score and the frequency-based score. The topic-based score per blog post is calculated as detailed in Section 3.1 and the score per feed is calculated similar to Eq. 11 (i.e., we take the highest scoring blog post per feed). This leads to:

$$p(bl|q) = \lambda \cdot p_c(bl|q) + (1 - \lambda) \cdot p_f(bl|q), \tag{12}$$

where $p_c(bl|q)$ denotes the topic-based score for the feed given the topic, and $p_f(bl|q)$ denotes the frequency-based score for the feed given the topic; note that $p_c(bl|q)$ is the same as in Eq. 11, while $p_f(bl|q)$ is defined as follows:

$$p_f(bl|q) = \textstyle\sum_{d \in bl} R(d|q) \cdot |bl|^{-1} \tag{13}$$

Here, $R(d|q) = 1$ if $p(d|q) > 0$, and $R(d|q) = 0$ otherwise; $|bl|$ denotes the number of blog posts in the blog bl.

Table 2: Feed distillation results

run id	MAP	p@10	p@30
uams07bdtop	0.1589	**0.3111**	0.2141
uams07bdtblm	**0.1605**	0.3067	0.2156
uams07bdfreq	0.1248	0.2467	**0.2170**

3.4 Runs

The following runs were submitted:

`uams07bdtop` uses the topical blog post retrieval model as described in Section 3.1 and aggregates blog posts to a feed according to Eq. 11.

`uams07tblm` uses the time-based retrieval model as described in Section 3.2. Training experiments showed the $\lambda = 0.04$ was the optimal setting for the exponential function.

`uams07bdfreq` uses the frequency-based retrieval model as described in Section 3.3. $\lambda = 0.5$ was the setting used in this run.

3.5 Results

We now consider the performance of our runs for the feed distillation task. The results are displayed in Table 2, which shows, for each run, the MAP scores, as well as the p@10 and p@30 scores.

The results in Table 2 show that the run incorporating time-based priors is the best performing run, even if the results are only slightly better than the baseline topic only run. More interestingly, the frequency-based run performs worse than the baseline run. However, the p@30 score for this run is higher than the other scores for p@30, which leads us to believe that the parameter setting is sub-optimal, leading to the decrease in performance. We performed additional experiments after the runs were submitted to confirm this hypothesis. These experiments were done using leave one out validation and showed that the optimal setting was $\lambda = 0.85$, leading to a MAP score of 0.1850, outperforming the other two runs.

An additional observation of the results is that only the best scoring blog posts were considered; we believe that better scores can be realized by taking additional blog posts into account when calculating the score for topical relevance.

4 Conclusions

In this paper we described our participation in the TREC 2007 Blog track. Our aim for the opinion finding task was to experiment with Indri and a mixture model. Result show that Indri significantly outperforms the mixture model. External expansion using a news corpus leads to improvement over the Indri baseline run, although bugs in the implementation caused decreased performance in the mixture model. Opinion finding by means of document priors shows beneficial, especially in case of lexicons. Overall we can conclude that opinion finding is highly dependent on topical retrieval and that focus still should be on this aspect: opinion detection can be done using lexicons, but non-lexical features also show promising results.

As to the feed distillation task, results show that using time-based document priors improved slightly over the baseline run. Incorporating a frequency-based approach, results showed no improvement using a default parameter value. It did show that better results could be achieved by estimating a more appropriate parameter value. We also found that scoring for topical relevance on a single blog post was not optimal.

Acknowledgments

Maarten de Rijke was supported by NWO under project numbers 017.001.190, 220-80-001, 264-70-050, 354-20-005, 600.065.120, 612-13-001, 612.000.106, 612.066.302, 612.069.006, 640.001.501, 640.002.501, and by the E.U. IST programme of the 6th FP for RTD under project MultiMATCH contract IST-033104.

References

[1] K. Balog, K. Hofmann, W. Weerkamp, and M. de Rijke. The University of Amsterdam at the TREC 2007 Enterprise Track. In *This Volume*, 2007.

[2] V. Lavrenko and W. B. Croft. Relevance-based language models. In *Proceedings of the 24th Annual international ACM SIGIR Conference on Research and Development in information Retrieval (SIGIR '01)*, pages 120–127. ACM Press, 2001.

[3] X. Li and W. Croft. Time-based language models. *Proceedings of the 12th International Conference on Information and Knowledge Managment (CIKM)*, pages 469–475, 2003.

[4] D. Metzler and W. B. Croft. A markov random feld model for term dependencies. In *SIGIR '05: Proceedings of the 28th annual international ACM SIGIR conference on research and development in information retrieval*, pages 472–479, New York, NY, USA, 2005. ACM Press.

[5] G. Mishne. *Applied Text Analytics for Blogs*. PhD thesis, University of Amsterdam, 2007.

[6] I. Ounis, M. de Rijke, C. Macdonald, G. Mishne, and I. Soboroff. Overview of the TREC-2006 Blog Track. In *The Fifteenth Text REtrieval Conference (TREC 2006) Proceedings*. NIST, 2007.