

# Evaluation of Relevance and Knowledge Augmentation in Discussion Search

Ingo Frommholz and Norbert Fuhr

University of Duisburg-Essen  
Duisburg, Germany  
[ingo.frommholz@uni-due.de](mailto:ingo.frommholz@uni-due.de)  
[fuhr@uni-duisburg.de](mailto:fuhr@uni-duisburg.de)

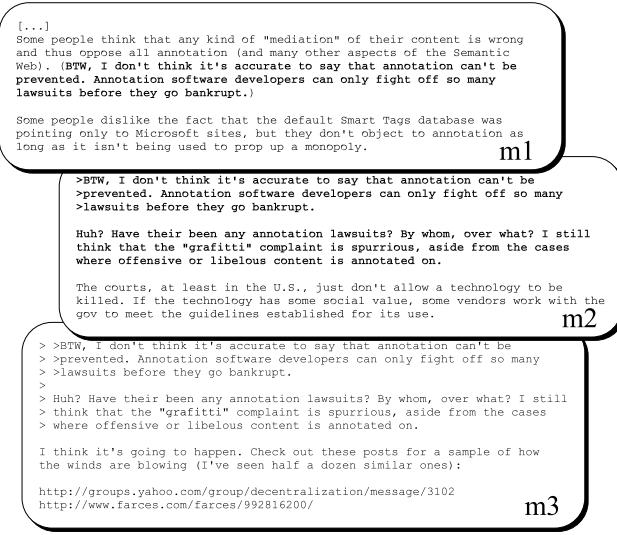
**Abstract.** Annotation-based discussions are an important concept for today's digital libraries and those of the future, containing additional information to and about the content managed in the digital library. To gain access to this valuable information, discussion search is concerned with retrieving relevant annotations and comments w.r.t. a given query, making it an important means to satisfy users' information needs. Discussion search methods can make use of a variety of context information given by the structure of discussion threads. In this paper, we present and evaluate discussion search approaches which exploit quotations in different roles as highlight and context quotations, applying two different strategies, knowledge and relevance augmentation. Evaluation shows the suitability of these augmentation strategies for the task at hand; especially knowledge augmentation using both highlight and context quotations boosts retrieval effectiveness w.r.t. the given baseline.

## 1 Introduction

Annotation-based discussions have been identified as an important concept for future digital libraries, supporting collaboration between users [3]. With annotations, a user can comment on the material at hand and others' annotations. As an example for an existing system, the COLLATE prototype uses nested public annotations as a building block for collaborative discussion in a community of scientists, with the purpose of interpreting the digital material at hand [13]. Other examples are web-based newswire systems like ZDNet News<sup>1</sup> which allow users to annotate published articles and other users' comments. In each of these systems, users can change their role from a passive reader to an active content provider. Stored discussion threads can be a helpful source for satisfying users' information needs: On the one hand, annotations can be exploited as auxiliary objects for *document search*, and on the other hand they are retrieval targets themselves in *discussion search*. It becomes clear that discussion search is an important means for uncovering valuable knowledge in information systems such as digital libraries.

---

<sup>1</sup> <http://news.zdnet.com/>



**Fig. 1.** A discussion thread

In this paper we present our discussion search approaches based on strategies called *knowledge* and *relevance augmentation*, respectively. The methods and results reported here continue the work and preliminary evaluation introduced in [4, 5]. In the next section we briefly present the test collection and our view on emails as annotations. We then introduce possible discussion search approaches in a probabilistic, logic-oriented framework. Subsequently we describe our experiments and discuss their results. We conclude after presenting some related work.

## 2 The Annotation View on Emails

In order to evaluate our discussion search approaches discussed below, we had to find a suitable test collection. Due to the lack of a “real” digital library testbed containing annotation threads, we participated in last year’s TREC Enterprise Track<sup>2</sup> in the discussion search task, where relevant emails had to be found [4]. The collection consists of 174,307 emails from several W3C discussion lists. Figure 1 shows an example excerpt of a discussion thread. Email replies usually consist of two different parts; the *quotations*, which are passages from the original text, and the *new part* containing the actual comments (annotations) of the email author. Quotations are usually prefixed by quotation characters like ‘>’; combinations of them determine the quotation depth. Quotations are thus the document fragments a comment belongs to. As an example, in m2 the comment

<sup>2</sup> <http://www.ins.cwi.nl/projects/trec-ent/>

“Huh?...established for its use” belongs to the fragment “BTW...go bankrupt” of m1. Applying the distinction between new parts and quotations as well as the thread structure extracted from email headers, we can transform email discussion threads into annotation threads with fragments (determined by quotations) as annotation targets. Due to the fact that whole new parts of emails were the primary target of the discussions search task, we applied one simplification. All quotations and all new parts of an email were merged, so that each email now consisted of one (merged) new part and at most one (merged) quotation part.

### 3 Discussion Search Approaches

We implement our retrieval functions in predicate logic, in particular probabilistic Datalog (pDatalog). We will briefly introduce pDatalog before discussing our retrieval approaches.

#### 3.1 Probabilistic Datalog

pDatalog [7] is a probabilistic variant of predicate logic. Similar to Prolog, its syntax consists of variables, constants, predicates and Horn clauses. Capital letters denote variables. Probabilities can be assigned to facts. Consider the following example program:

```
0.7 about(d1,"databases").  0.5 about(d1,"retrieval").
retrieve(D) :- about(D,"databases").
retrieve(D) :- about(D,"retrieval").
```

Probabilistic facts model extensional knowledge. The `about` predicate says that `d1` is about ‘databases’ with 0.7 probability and about ‘retrieval’ with 0.5 probability. Rules model intensional knowledge, from which new probabilistic facts are derived. The rule `retrieve` means that a document should be retrieved when it is about ‘databases’ or ‘retrieval’. With the given facts and rules, pDatalog would now calculate the retrieval status values of a document  $d$  w.r.t. the retrieval function `retrieve` as a combination of probabilistic evidence. In particular, if  $e_1, \dots, e_n$  are joint independent events, pDatalog computes

$$P(e_1 \wedge \dots \wedge e_n) = P(e_1) \cdot \dots \cdot P(e_n) \quad (1)$$

$$P(e_1 \vee \dots \vee e_n) = \bigoplus_{i=1}^n P(e_i) = \sum_{i=1}^n (-1)^{i-1} \left( \sum_{\substack{1 \leq j_1 < \dots < j_i \leq n}} P(e_{j_1} \wedge \dots \wedge e_{j_i}) \right) \quad (2)$$

For our example document `d1`, pDatalog would calculate

$$\begin{aligned} P(\text{retrieve}(d1)) &= P(\text{about}(d1, "databases") \vee \text{about}(d1, "retrieval")) \\ &= 0.7 + 0.5 - 0.7 \cdot 0.5 = 0.85 \end{aligned}$$

### 3.2 Simple Content-Based Approach

In this baseline approach, we do not apply any context at all for discussion search. Each document only contains new parts of an email, stripping all quotations. The approach can be expressed with the following datalog rules:

```
wqterm(T) :- qterm(T) & termspace(T).
about(T,D) :- term(T,D).
retrieve(D) :- wqterm(T) & about(T,D).
```

The `qterm` predicate contains the query terms (after stemming and stopword elimination). `termspace` contains the termspace, here regarding only the new part of an email as a document. `termspace` thus contains all terms appearing in new parts of emails. For each term  $t$  in the termspace, it is  $P(\text{termspace}(t)) = P(t)$  which is interpreted as an intuitive measure of the probability of  $t$  being informative.  $P(t)$  can be estimated based on the inverse document frequency of  $t$ ,  $\text{idf}(t) = -\log(df(t)/\text{numdoc})$ , with  $df(t)$  as the number of documents in which  $t$  appears and  $\text{numdoc}$  as the number of documents in the collection, and

$$P(t) \approx \frac{\text{idf}(t)}{\text{maxidf}} \quad (3)$$

with  $\text{maxidf}$  being the maximum inverse document frequency. The `wqterm` rule states that we weight a query term  $t$  according to  $P(t)$ . `term` relates terms to the documents they appear in. For each term  $t$  in document  $d$ ,  $P(\text{term}(t,d)) = P(t|d)$ , the probability that we observe term  $t$  given document  $d$ .  $P(t|d)$  is estimated as

$$P(t|d) \approx \frac{tf(t,d)}{\text{avgtf}(d) + tf(t,d)} \quad (4)$$

where  $tf(t,d)$  is the frequency of term  $t$  in document  $d$  and  $\text{avgtf}(d)$  is the average term frequency of  $d$ , calculated as  $\text{avgtf}(d) = \sum_{t \in d^T} tf(t,d)/|d^T|$  with  $d^T$  being the document representation of  $d$  (i.e. the bag of words of  $d$ ). We say that a document is about a term if the term appears in the document; this is modeled with the `about` rule. The `retrieve` rule is our actual retrieval function. A document should be retrieved if it contains at least one query term. The retrieval status value of  $d$  is determined by  $P(\text{retrieve}(d))$  which in turn depends on query and document-term weights, as described above. The result list presented to the user ranks documents according to descending retrieval status values.

### 3.3 Context Quotations

In the last subsection we were only considering new parts of email messages for retrieval. However, in an email discussion thread, we have the information about the targets that a comment addresses, given by the quotations. Quotations are an important source for determining what a new part is about, as can be seen

in message  $m_3$  in our example in Figure 1. If we only consider the new part of the message, as we do in the approach described above, the system could not infer that this part is actually about “annotation lawsuits”.  $m_3$  would not be retrieved for such a query, although it would be relevant. Quotations thus establish an important context for the new parts of messages; quotations are referred to as *context quotations* when regarding them as such a kind of context for new parts. We will now introduce our idea of exploiting context quotations for discussion search, beginning with the obvious choice, merging quotations and new parts by not distinguishing between them in email messages.

**Merging Quotations and New Parts** This simple approach sees whole emails as a document (instead of only new parts as in Subsection 3.2). Any further parsing of email messages to distinguish between quotations and new parts is not required here; all terms in a message are regarded as belonging to the corresponding document. In an annotation scenario, this is similar to the case where all annotation targets are merged with their respective annotation to form a new document. We apply exactly the same predicates and rules like those discussed in Subsection 3.2, except that the estimations of  $P(t)$  and  $P(t|d)$ , respectively, are now based on the view of a document being a full email message (resulting in different values for term and document frequencies).

**Knowledge Augmentation** While in the last approach context quotations were merged with new parts, the approaches discussed next regard context quotations as separate, virtual documents. Thus, from a message  $m$ , two new documents are created:  $d_m$ , containing the new part of  $m$ , and  $quot_m$ , containing  $m$ 's quotations.  $quot_m$  is regarded as “virtual” since it is not to be retrieved, but serves as an auxiliary document to determine the relevance of  $d_m$ . Furthermore, each virtual document does not contribute to the document frequency of a term. If a term  $t$  appears in both  $quot_m$  and  $d_m$ , only its appearance in  $d_m$  is counted and used in Equation 3.

We introduce a new predicate `quotedterm(t,d)` which says that the term  $t$  appears in the quotation  $quot$  belonging to document  $d$ . It is  $P(\text{quotedterm}(t,d)) = P(t|quot)$ , and the latter probability is estimated with Equation 4. We apply a *knowledge augmentation* approach by extending our `about` rule to

```
about(T,D) :- term(T,D).
about(T,D) :- acc("quotation") & quotedterm(T,D).
```

where `acc("quotation")` describes the event that a quotation is actually accessed when reading the unquoted part.  $P(\text{acc}("quotation"))$  is thus the probability that a quotation is considered. By extending the `about` rule like this, we augment our knowledge of what a new part is about with the knowledge of what the quotation is about. In this extended context, new terms are introduced which appear in quotations only, and the probability that a document is about a term

is raised according to Equations 1 and 2 if we also observe this term in the quotation. The analogy to the real world is that if a user reads the new part first and then the corresponding quotation, she augments her knowledge of what the new part is about. The `wqterm` and the `retrieve` rules are the same as before.

**Relevance Augmentation** Another augmentation strategy we are going to evaluate is *relevance augmentation*. Here, we augment the knowledge that a new part is relevant with the knowledge that its corresponding quotation part is relevant. The idea is that we infer to a certain degree the relevance of a new part with the relevance of its quotation part. This context-based relevance decision is performed by the system in two steps. First, the relevance of documents and context quotations w.r.t. the query is determined:

```
rel(D) :- wqterm(T) & about(T,D).
quot_rel(D) :- wqterm(T) & quotedterm(T,D).
```

In the second step, this knowledge is combined, taking into account the probability that we actually access the quotation:

```
retrieve(D) :- rel(D).
retrieve(D) :- acc("quotation") & quot_rel(D).
```

(`wqterm` and `about` are the same as in Section 3.2).

### 3.4 Highlight Quotations

When a user annotates a (part of a) document, it is assumed that she found it interesting enough to react to it. This means the annotation target is implicitly highlighted and considered important by the annotation author, reaching a kind of *n-way consensus* [9] of the significance of this part if *n* persons used it as annotation target. Examining the quotations and the quotation levels of emails, we can identify such highlighted parts of previous messages. A highlight quotation of a message *m* in another message *m'* is the part of *m* which is quoted by *m'*, where *m'* is a (direct or indirect) successor of *m* in the discussion thread. Consider the following simple example with 3 messages:

<pre>m1: line1.1       line1.2       line1.3</pre>	<pre>m2:  &gt; line1.2           &gt; line1.3           line2.1</pre>	<pre>m3: &gt;&gt; line1.3           &gt; line2.1           line3.1</pre>
--	---	--

*m1* consists of 3 lines (*line1.1 - line1.3*). *m2* quotes two of these lines, *line1.2* and *line1.3*. *m3* quotes a line from *m1* (*line1.3*) and from *m2* (*line2.1*). The quotation in *m2* containing *line1.2* and *line1.3* is a highlight quotation of *m1*. Our claim is that *line1.2* and *line1.3* are important due to the fact that they are quoted; *line1.3* seems to be even more important since it is quoted in *m3* as well. For an email message, we create a highlight quotation virtual document from each quotation containing a fragment of this email message. In our example we would create two highlight quotation virtual documents for *m1: high\_m1-m2*

consists of `line1.2` and `line1.3` (the part of `m1` quoted in `m2`), and `high_m1-m3` contains `line1.3` (the part of `m1` quoted in `m3`). For `m2`, one virtual document is created (`high_m2-m3` containing `line2.1`). We use highlight quotation virtual documents as a context for retrieval by performing knowledge and relevance augmentation again.

**Knowledge Augmentation** To add highlight quotations, we introduce a new predicate `highlightterm(t,d,high)` where `t` is a term, `d` a document and `high` the highlight quotation where `t` appears. It is  $P(\text{highlightterm}(t,d,\text{high})) = P(t|\text{high})$ , again estimated with Equation 4. *Knowledge augmentation* is applied by extending the `about` rule:

```
about(T,D) :- term(T,D).
about(T,D) :- acc("highlight") & highlightterm(T,D,H).
```

$P(\text{acc}(\text{"highlight"}))$  is the probability that we actually consider (access) highlight quotations. A short note on the evaluation of `highlightterm(T,D,H)` follows. In the second `about` rule, the variable `H` is free. For a possible valuation `D=d` and `T=t` to determine `about(t,d)`, pDatalog substitutes `highlightterm(t,d,H)` with a disjunction containing all possible values `H` can take. In our example above, let the term ‘`developers`’ appear in `line1.3`. Now, with `T="developers"` and `D=m1`, `highlightterm("developers",m1,H)` would resolve to `highlightterm("developers",m1,high_m1-m2) ∨ highlightterm("developers",m1,high_m1-m3)`. The probability of this disjunction is calculated and multiplied with  $P(\text{acc}(\text{"highlight"}))$  to gain a probability for the second `about` rule. `wqterm` and `retrieve` are the same as in Section 3.2 here.

**Relevance Augmentation** Relevance augmentation with highlight quotations is quite straightforward. Again, we need two steps;

```
rel(D) :- wqterm(T) & about(T,D).
high_rel(D,H) :- wqterm(T) & highlightterm(T,D,H).
```

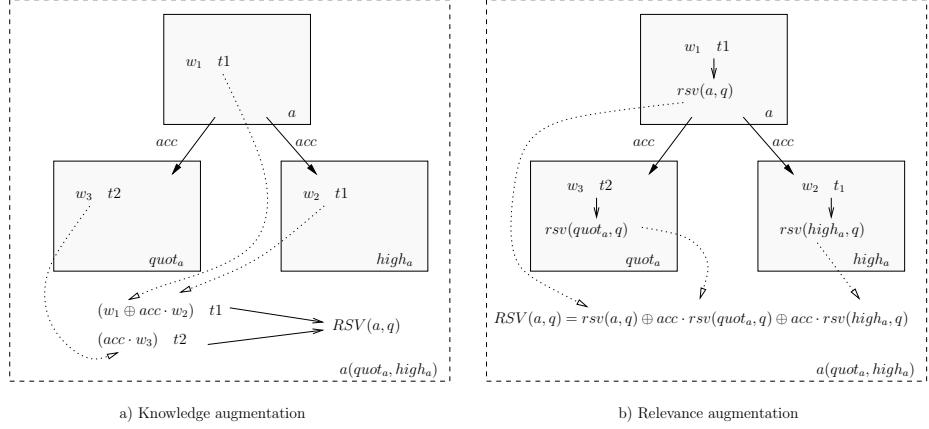
determines the relevance of documents and highlight quotations, and

```
retrieve(D) :- rel(D).
retrieve(D) :- acc("highlight") & high_rel(D,H).
```

combines this evidence in the actual retrieval rules. `wqterm` and `about` are the same as in Section 3.2.

### 3.5 Combination

We also conducted experiments where we combined the evidence gained from highlight and context quotations. For knowledge augmentation, we combined the corresponding `about` rules introduced in Sections 3.3 and 3.4 with `wqterm` and `retrieve` identical as in Section 3.2. For relevance augmentation, we combined the `rel`, `high_rel`, `quot_rel` and `retrieve` rules in Sections 3.3 and 3.4, respectively, with `wqterm` and `about` as before in Section 3.2.



**Fig. 2.** Knowledge and relevance augmentation

### 3.6 Non-probabilistic Formulation

The knowledge and relevance augmentation strategies are not bound to a probabilistic, logic-based formulation like the one we presented above with pDatalog. Consider the example in Fig. 2. Here we can see an example annotation  $a$  with a corresponding highlight quotation document  $high_a$  and a context quotation document  $quot_a$ .  $acc$  models the probability that  $high_a$  or  $quot_a$ , respectively, are accessed from  $a$ . With knowledge augmentation, the term weights ( $w_1$  and  $w_2$  for  $t_1$  and  $w_3$  for  $t_2$ ) are propagated to the supercontext  $a(quot_a, high_a)$  according to the access probability. The operator  $\oplus$  combines the weights from the subcontexts in the supercontext;  $\oplus$  can be a simple sum operator, or, as it is the case with pDatalog, formulated with the inclusion-exclusion formula in Equation 2. The calculated new term weights for  $t_1$  and  $t_2$  are then used to compute the final retrieval status value  $RSV(a, q)$  of  $a$  w.r.t. the query  $q$ . When applying relevance augmentation, we first calculate a local retrieval status value  $rsv(a, q)$ ,  $rsv(quot_a, q)$  and  $rsv(high_a, q)$ , respectively, for the subcontexts; these values are again combined in the supercontext  $a(quot_a, high_a)$  with the  $\oplus$  operator in order to compute the final retrieval status value  $RSV(a, q)$ .

## 4 Experiments and Results

The main goal of our experiments was to answer the question: can relevance or knowledge augmentation increase retrieval effectiveness, and which strategy should be preferred? Whereas for knowledge augmentation the first question has already been answered [4], we have as yet not conducted any experiments for relevance augmentation. We also provide the results of further runs for knowledge augmentation, applying different values for  $P(acc("quotation"))$

and  $P(\text{acc}("highlight"))$ , respectively. For both probabilities, we used global values ranging from 0.1 to 1.0, in steps of 0.1<sup>3</sup>. For our experiments, we used the W3C email lists described in Section 2 with 59 distinct queries. Topics and relevance judgements were given by the participants of the TREC 2005 Enterprise track. All runs were performed using HySpirit<sup>4</sup>, a pDatalog implementation. Table 1 briefly describes the experiments and their settings.

Experiment	Parameters	Description
baseline		The baseline, only new parts.
merged		Merged quotations and new parts
qknow- $x$	$P(\text{acc}("quotation")) = x$	Knowledge augmentation with context quotations
qrel- $x$	$P(\text{acc}("quotation")) = x$	Relevance augmentation with context quotations
hknow- $x$	$P(\text{acc}("highlight")) = x$	Knowledge augmentation with highlight quotations
hrel- $x$	$P(\text{acc}("highlight")) = x$	Relevance augmentation with highlight quotations
cknow- $x-y$	$P(\text{acc}("highlight")) = x$ $P(\text{acc}("quotation")) = y$	Knowledge augmentation with highlight and context quotations
crel- $x-y$	$P(\text{acc}("highlight")) = x$ $P(\text{acc}("quotation")) = y$	Relevance augmentation with highlight and context quotations

**Table 1.** Description of experiments

Some selected results of our experiments are presented in Table 2, where we show the mean average precision and the precision at selected numbers of documents retrieved. The latter values are important user-oriented ones: users tend to browse through the first 20 or even 30 top-ranked documents in a result list, but usually do not go deeper in the ranking. The other runs not presented here did not gain better results or considerably new insights. From the results we can see that both relevance and knowledge augmentation improve retrieval effectiveness: there are slight improvements with highlight quotations, and larger improvements with context quotations. To our surprise, the experiment with merged context quotations and new parts gains worse results than the baseline. The combination of highlight and context quotations further improves retrieval effectiveness. So we see that creating separate virtual documents from highlight and context quotations and linking them with a certain access probability to their corresponding document seems to be worth the effort. Regarding knowledge vs. relevance augmentation, the results clearly show that knowledge augmentation is

<sup>3</sup> We bear in mind that this is only a preliminary solution; more advanced ones might take evidence from the thread structure or given by users' preferences to estimate the access probability.

<sup>4</sup> <http://qmri.dcs.qmul.ac.uk/hyspirit.php>

to be preferred over relevance augmentation. In the case of context quotations, knowledge augmentation can possibly handle the vocabulary problem better (when query terms do not appear in the new part, but in the quotation), but the exact reasons are not yet clear and subject to further investigation. Figure 3 shows the interpolated recall-precision averages of selected runs.

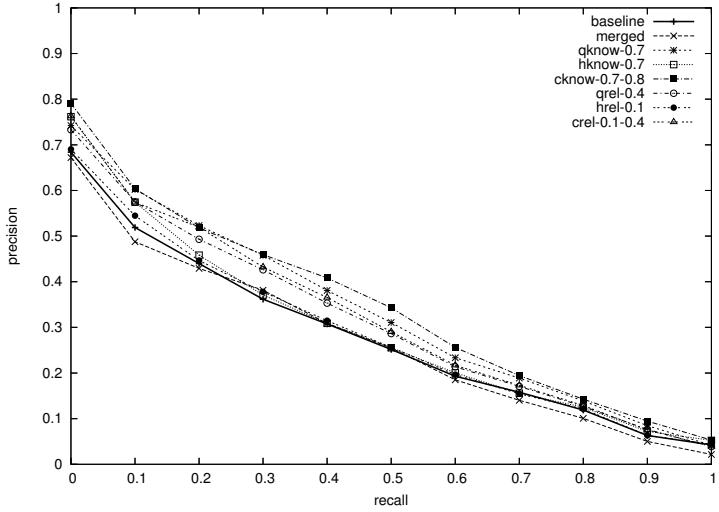
Run	MAP	P@5	P@10	P@20	P@30
baseline	0.2599	0.4441	0.4103	0.3695	0.3220
merged	0.2565	0.4678	0.3966	0.3458	0.3068
qknow-0.7	0.3162	0.5220	0.4678	<b>0.3983</b>	<b>0.3537</b>
qknow-0.8	0.3145	0.5186	0.4712	0.3915	0.3503
hknow-0.7	0.2784	0.5085	0.4390	0.3534	0.3124
hknow-0.4	0.2767	0.4542	0.4356	0.3737	0.3266
hknow-0.3	0.2726	0.4915	0.4458	0.3720	0.3260
qrel-0.4	0.2957	0.4746	0.4390	0.3847	0.3401
hrel-0.1	0.2669	0.4780	0.4288	0.3636	0.3192
cknow-0.7-0.8	<b>0.3298</b>	<b>0.5492</b>	0.4746	0.3890	0.3475
cknow-0.3-0.7	0.3252	0.5458	<b>0.4881</b>	0.3975	0.3520
crel-0.1-0.4	0.3024	0.4814	0.4424	0.3831	0.3367

**Table 2.** Mean average precision and precision at 5, 10, 20 and 30 documents retrieved for some selected runs. Best results are printed in bold.

## 5 Related Work

The studies performed by the Marshall group (see, e.g., [9, 10]) contain many results and conclusions relevant for designers of annotation systems, which have a strong impact on our work. The studies reported in Shipman *et al.* [12] focus on the identification of high-value annotations in order to find useful passages in a text. Agosti *et al.* examine annotations from a syntactic, semantic and pragmatic view [2].

There are several approaches for annotation-based information retrieval and discussion search. A relevance feedback approach where only highlighted terms instead of whole documents are considered is reported to be successful [8]. [1] reports on an approach where evidence coming from documents and the elements in the annotation hypertext is combined using a data fusioning approach. Xi *et al.* evaluate a feature-based approach for discussion search in [15]; their results show an increase in retrieval effectiveness when using the thread context. The proceedings of the TREC 2005 conference contain many other evaluations of discussion search approaches [14]. The idea of knowledge augmentation has its roots in structured document retrieval and is discussed thoroughly in [11].



**Fig. 3.** Interpolated recall-precision graph of selected runs

## 6 Conclusion

In this paper we presented some approaches for discussion search and their evaluation, using quotations in a special role as context and highlight quotations, respectively. Based on probabilistic datalog, we applied two strategies, knowledge and relevance augmentation. The results indicate that a knowledge augmentation strategy combining highlight and context quotations is preferable. Knowledge augmentation has another benefit: it is query independent to a certain degree, meaning that  $P(\text{about}(t, d))$  may be calculated offline as a post-indexing step, whereas the relevance augmentation strategy can only be applied during query processing. Based on the promising results gained so far, we proposed a probabilistic, object-oriented logical framework for annotation-based retrieval called POLAR in [5].

Future work will concentrate on further evaluation and discussion of our augmentation strategies using context and highlight quotations. As a third source of evidence, the content of annotations made to another annotation could also be used for augmentation, as discussed for relevance augmentation in [6].

## References

1. Maristella Agosti and Nicola Ferro. Annotations as context for searching documents. In Fabio Crestani and Ian Ruthven, editors, *Information Context: Nature, Impact, and Role: 5th International Conference on Conceptions of Library and Information Sciences, CoLIS 2005*, volume 3507 of *Lecture Notes in Computer Science*, pages 155–170, Heidelberg et al., June 2005. Springer.

2. Maristella Agosti, Nicola Ferro, Ingo Frommholz, and Ulrich Thiel. Annotations in digital libraries and collaboratories – facets, models and usage. In Rachel Heery and Liz Lyon, editors, *Research and Advanced Technology for Digital Libraries. Proc. European Conference on Digital Libraries (ECDL 2004)*, Lecture Notes in Computer Science, pages 244–255, Heidelberg et al., 2004. Springer.
3. Alberto Del Bimbo, Stefan Gradmann, and Yannis Ioannidis. Future research directions – 3rd DELOS brainstorming workshop report, DELOS Network of Excellence, July 2004.
4. Ingo Frommholz. Applying the annotation view on messages for discussion search. In Voorhees and Buckland [14].
5. Ingo Frommholz and Norbert Fuhr. Probabilistic, object-oriented logics for annotation-based retrieval in digital libraries. In *Proceedings of JCDL 2006*, 2006. In print.
6. Ingo Frommholz, Ulrich Thiel, and Thomas Kamps. Annotation-based document retrieval with four-valued probabilistic datalog. In Thomas Roelleke and Arjen P. de Vries, editors, *Proceedings of the first SIGIR Workshop on the Integration of Information Retrieval and Databases (WIRD'04)*, pages 31–38, Sheffield, UK, 2004.
7. Norbert Fuhr. Probabilistic Datalog: Implementing logical information retrieval for advanced applications. *Journal of the American Society for Information Science*, 51(2):95–110, 2000.
8. Gene Golovchinsky, Morgan N. Price, and Bill N. Schilit. From reading to retrieval: freeform ink annotations as queries. In *Proceedings of the 22nd International Conference on Research and Development in Information Retrieval*, pages 19–25, New York, 1999. ACM.
9. Catherine C. Marshall. Toward an ecology of hypertext annotation. In *Proceedings of the ninth ACM conference on Hypertext and hypermedia: links, objects, time and space – structure in hypermedia systems*, pages 40–49, 1998.
10. C.C. Marshall and A.J. Brush. Exploring the relationship between personal and public annotations. In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 349–357, New York, NY, USA, 2004. ACM Press.
11. Thomas Rölleke. *POOL: Probabilistic Object-Oriented Logical Representation and Retrieval of Complex Objects*. PhD thesis, University of Dortmund, Germany, 1998.
12. Frank Shipman, Morgan Price, Catherine C. Marshall, and Gene Golovchinsky. Identifying useful passages in documents based on annotation patterns. In Panos Constantopoulos and Ingeborg T. Sølvberg, editors, *Research and Advanced Technology for Digital Libraries. Proc. European Conference on Digital Libraries (ECDL 2003)*, Lecture Notes in Computer Science, pages 101–112, Heidelberg et al., 2003. Springer.
13. Ulrich Thiel, Holger Brocks, Ingo Frommholz, Andrea Dirsch-Weigand, Jürgen Keiper, Adelheit Stein, and Erich Neuhold. COLLATE - a collaboratory supporting research on historic european films. *International Journal on Digital Libraries (IJDLD)*, 4(1):8–12, 2004.
14. E. M. Voorhees and Lori P. Buckland, editors. *The Fourteenth Text REtrieval Conference (TREC 2005)*, Gaithersburg, MD, USA, 2005. NIST.
15. Wensi Xi, Jesper Lind, and Eric Brill. Learning effective ranking functions for newsgroup search. In Kalervo Järvelin, James Allen, Peter Bruza, and Mark Sanderson, editors, *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 394–401, New York, 2004. ACM.