

Looking for Books in Social Media: An Analysis of Complex Search Requests

Marijn Koolen¹, Toine Bogers², Antal van den Bosch³, and Jaap Kamps¹

¹ University of Amsterdam, The Netherlands
{marijn.koolen,kamps}@uva.nl

² Aalborg University Copenhagen, Denmark
toine@hum.aau.dk

³ Radboud University, Nijmegen, The Netherlands
a.vandenbosch@let.ru.nl

Abstract. Real-world information needs are generally complex, yet almost all research focuses on either relatively simple search based on queries or recommendation based on profiles. It is difficult to gain insight into complex information needs from observational studies with existing systems; potentially complex needs are obscured by the systems' limitations. In this paper we study explicit information requests in social media, focusing on the rich area of social book search. We analyse a large set of annotated book requests from the LibraryThing discussion forums. We investigate 1) the comprehensiveness of book requests on the forums, 2) what relevance aspects are expressed in real-world book search requests, and 3) how different types of search topics are related to types of users, human recommendations, and results returned by retrieval and recommender systems. We find that book search requests combine search and recommendation aspects in intricate ways that require more than only traditional search or (hybrid) recommendation approaches.

Keywords: Book Search, Social Media, Evaluation, Recommendation.

1 Introduction

The rise of social media has had a major impact on how we search for and share information. For instance, it has radically changed the nature of book discovery, which has become easier than ever due to *social cataloging sites*, such as LibraryThing, GoodReads, Shelfari, BookLamp, Libib, and The Reading Room. We focus on LibraryThing¹ (LT), a popular social cataloguing site. The book collections shared on LT by its 1.8 million members cover over 8 million unique works in total. They describe not only the contents of those books, but also how the books engaged them, what their impact was, and how this related to other reading experiences. LT also offers a popular discussion forum (see Figure 1) for readers to discuss and review books, authors, and literature in general. A prominent use of the LT forum is book discovery: thousands of LT members use the forum to receive or provide recommendations for which books to read next. These book requests display a remarkable breadth, ranging from search-type requests for books on

¹ <http://librarything.com/>, last accessed January 11, 2015.

specific topics or for certain moods, to recommendation-type requests for books similar to what a member has already read.

The general aim of this paper is to investigate whether explicit information requests in such social media, in particular related to book search, can be used to gain insight in complex information needs, i.e., those that cannot be solved by a straightforward look-up search. We study this in the context of the INEX Social Book Search Track² [12, 14, 15]. In recent years, this track has focused on book requests posted on the LT discussion forums. This paper provides a more detailed investigation into the nature of such requests. In the forums anyone can ask for book recommendations for a specific topic and other members reply with book suggestions. These suggestions can be seen both as relevance judgments and recommendations. The search requests go beyond topical relevance [13] and include many subjective aspects such as quality, interestingness, engagement, and familiarity. Cosijn and Ingwersen [7] and Saracevic [20] are among many that argue for the existence of different types of relevance in addition to pure topical relevance, such as situational, motivational, and affective relevance. A comprehensive survey of different interpretations of relevance is given by Borlund [4]. In this paper, we explore the relevance aspects present in the book domain by annotating and analyzing a large set of book requests from the LT forums.

We aim to address the following research questions in this paper:

- RQ1.** How comprehensive are book requests on the LT forum in terms of explicit information on the information need, the context of use, and the context of the user?
- RQ2.** What topical and non-topical relevance aspects are present in book search requests on the LT forums?
- RQ3.** How do different types of topics relate to user characteristics, human recommendations, and retrieval and recommender system results?

The rest of this paper is organized as follows. Section 2 presents related work, followed by an overview of the rich contextual data about book requests we can extract from the LT discussion forums in Section 3. Section 4 analyses the book requests with respect to the topical and non-topical relevance aspects expressed in them. Section 5 explores how book requests relate to the context of the user, human book recommendations, and retrieval or recommender system results. Finally, in Section 6, we discuss our results and draw conclusions.

2 Related Work

The INEX Social Book Search Track [12, 14, 15] investigates book search in collections with both professional metadata and social media content. For evaluation they use book requests on the LT discussion forums as search topics and book suggestions by members as relevance judgments and recommendations. Koolen et al. [13] observed that these requests are complex and contain non-topical aspects, and found that the forum suggestions are different in nature than editorial relevance judgments with respect to system evaluation. In this paper we focus on the search requests themselves.

² All the data used in this paper are made available as part of the CLEF/INEX 2014 SBS Track.

The screenshot shows a forum post on the LibraryThing website. The page header includes the LibraryThing logo, navigation links (Home, Profile, Your books, Add books, Talk, Groups, Local, More, Zeitgeist), a search bar, and user information (MarinusFDT, Sign out, Help). The post is in the 'Political Philosophy' group, which has 212 members and 87 messages. The post title is 'Politics of Multiculturalism Recommendations?' and it has 11 messages. The first message is by user 'steve.clason' from Sep 26, 2010, at 11:32pm. The message content reads: 'I'm new, and would appreciate any recommended reading on the politics of multiculturalism. Parekh's Rethinking Multiculturalism: Cultural Diversity and Political Theory (which I just finished) in the end left me unconvinced, though I did find much of value I thought he depended way too much on being able to talk out the details later. It may be that I found his writing style really irritating so adopted a defiant skepticism, but still...'. Below this is a second message by user 'rsterling' from Sep 27, 2010, at 1:31am, which reads: 'Anyway, I've read Sen, Rawls, Habermas, and Nussbaum, still don't feel like I've wrapped my little brain around the issue very well and would appreciate any suggestions for further anyone might offer.' The post also includes a 'Reply | More' link and a 'Touchstones' section listing related works like 'Rethinking Multiculturalism: Cultural Diversity and Political Theory by Bhikhu Parekh' and 'Multicultural Citizenship by Will Kymlicka'.

Fig. 1. Book request on the LibraryThing forum

Ross [19] found that readers use a variety of clues to choose books. Reading a book is a substantial investment of time and energy, so readers look for recommendations from trusted sources for selection. Reuter [18] studied book selection by children and identify a list of 46 factors influencing their choices. Buchanan and McKay [5] investigated search activities of customers in bookshops. They find that enquiries often arise from cultural context—reading with others, references and reviews in media—and argue that customers’ mental models may deviate from the standard bibliographic metadata. Cunningham et al. [8] studied collaborative information behaviour in bookshops. They found that groups of customers use many different ways to share information about books, e.g., talking aloud, pointing, reading, and searching together, and that they use these interactions to achieve agreement on which books to select. The gap between their mental model and the access points for online book collections may be why users turn to the LT forum for requests.

A considerable amount of related work exists on forum search, where the focus is typically on retrieving results from the collection of threads in a single forum. Examples of such approaches include work by Elsas and Carbonell [10] and Bhatia and Mitra [3]. In contrast, we analyze the initial forum posts describing a user’s information need, in order to perform cross-collection search using these need descriptions. Our overall aim to use the forums to shed light on complex search requests, their context and relevance aspects, is related to a wealth of studies in information seeking. Some of the most comprehensive earlier studies predate the web and modern search systems (e.g., [21]). Our general approach is to tap into a new source of evidence for researching complex information seeking behavior.

3 Book Search Requests in LT Forums

In this section, we investigate RQ1: How comprehensive are book requests on the LT forum in terms of explicit information on the information need, the context of use, and the context of the user?

The LT discussion forums are used to discuss a broad range of topics, most of which are book-related. Many members turn to this forum asking for book suggestions and other members can reply and provide suggestions. In a random sample of 500 posts we found 67 (13.4%) containing an explicit book request. Given the massive scale of the forums with nearly 5 million messages and 3.5 million identified book mentions, this gives us access to a huge supply of real world complex search requests.³ For the more straightforward search tasks, LT users are likely to use book search engines available at e.g. LT, Amazon, or libraries. In contrast, the forum requests contain more complex search needs that LT members have, expressed in natural language.

For instance, the request in Figure 1 is highly complex, providing requirements about the content as well as examples of books and authors that the poster is already familiar with, and contextual cues on usage. The user name links to the profile of the user, which provides additional context such as their personal book catalogue. The example books mentioned introduce a form of query-by-example that could also be seen as a recommendation task. These forum threads provide us with an unobtrusive method of investigating realistic, complex search requests that go well beyond traditional query log analysis. Members are not limited by the functionalities of a search engine or recommender system when expressing their request, but only by the concreteness of their information need and their ability to express it in natural language. As a result, they typically leave rich descriptions of their information need as well as many contextual clues to ensure others can understand its complexity.

Moreover, the LT forums allow users to mark up the names of books and authors through a simple wiki-like syntax using so-called *touchstones*. The system then automatically identifies the correct book/author and links the marked-up text to the right LT entity. These suggestions are a form of human relevance judgements.

Summarizing, from the forums we can derive rich statements of requests, including explicit statements on the context of use and the context of the user, with example books and 'ground truth' human recommendations. We find that such forum data give a unique opportunity to study complex search requests, and that the requests exhibit an amazing variation in topical and non-topical aspects. This prompts us to investigate what relevance aspects are used in the next section.

4 Relevance in Forum Book Search

In this section we study RQ2: What topical and non-topical relevance aspects are present in book search requests on the LT forums?

4.1 Relevance Aspects

Our first step is to investigate the complexity of these book search requests and the kind of relevance aspects expressed in them. Reuter [18] collected data from a user study in a children's library and identified 46 aspects, grouped into seven broad categories. We use those categories as our guide for analyzing the relevance aspects of book search requests. Due to its prominence in the LT forums, we introduce *known-item* search as an additional aspect. This resulted in the following eight relevance aspects:

³ <https://www.librarything.com/zeitgeist>, last accessed on January 11, 2014.

Accessibility The language, length, or level of difficulty of a book.

Content Topic, plot, genre, style, or readability of a book.

Engagement Affective types of reading experiences evoked by books.

Familiarity Books similar to known books or related to a previous experience.

Known-item Descriptions of known books to identify the title and/or author.

Metadata Aspects like title, author, publication year and format.

Novelty Books that are unusual or quirky, or have novel content.

Socio-cultural Books related to the user's socio-cultural background or values, have (had) a particular cultural or social impact, or are popular or obscure.

4.2 Annotating Book Search Requests

To determine how prominent these different relevance aspects are on the LT forums, we annotated a sample of topic threads for relevance and other characteristics. We selected forum threads likely to contain requests for book recommendations using a simple regular-expression-based classifier, which filtered out all topics that did not contain one or more 'trigger' expressions, such as '*suggest*', '*looking for*' and '*which books*'. This resulted in a set of 9,403 topic threads containing touchstones. A random set of 2,646 of these topics were annotated by eight different Information Science students, three from the Royal School of Library and Information Science in Copenhagen, three from the Oslo and Akershus University of Applied Sciences, and two from Aalborg University Copenhagen. Each topic was annotated by a single annotator. We created a Web interface to help our annotators (1) identify topic threads as either *book requests* or *non-requests*; (2) annotate the requests by which relevance aspect(s) they express; and (3) annotate the suggestions provided by other LT members in the thread. This task included questions on whether the suggestion providers appeared to have read the suggested books and whether their recommendation was positive, negative, or neutral.

Of the 2,646 topics annotated by the students, 944 topics (36%) were identified as containing a book request (recall that 13.4% of a random sample contained book requests). For each identified book request, annotators could specify multiple relevance aspects. For example, for topic 99,309 on the "*politics of multiculturalism*" (partly shown in Figure 1), the topic starter asks for suggestions about a particular topic (*content* relevance), but also asks for books similar to what he has already read on the topic (*familiarity*), but written in a less annoying style (*engagement*).

4.3 Analysis

The distribution of relevance aspects in our annotated set of 944 book requests is shown in the left half of Table 1. The majority of book search information needs on the LT forums express *content* aspects (698 topics or 74%). *Familiarity* is the second most frequent aspect at 36%. These two aspects are often combined in a single book request: 267 topics (28%) express both aspects. An example of such a request is "*Can someone recommend a book that has all the joy, charm, numerous characters, pathos, adventure, love of language, etc. that the novel David Copperfield has?*" (topic 10392). The searcher wants recommendations based on the book *David Copperfield*, but also describes aspects of the book to base these recommendations on. This is querying by

Table 1. Aspect distribution and overlap in the 944 forum topics (left side) and the conditional probability $P(\text{column} | \text{row})$ (right side)

| | Aspect overlap | | | | | | | | Conditional probability | | | | | | | |
|----------------|----------------|-----|-----|-----|-----|-----|----|-----|-------------------------|------|------|------|------|------|------|------|
| | A | C | E | F | K | M | N | S | A | C | E | F | K | M | N | S |
| Accessibility | 152 | 109 | 44 | 50 | 15 | 39 | 8 | 27 | 1.00 | 0.72 | 0.29 | 0.33 | 0.10 | 0.26 | 0.05 | 0.18 |
| Content | | 698 | 172 | 267 | 100 | 176 | 26 | 99 | 0.16 | 1.00 | 0.25 | 0.38 | 0.14 | 0.25 | 0.04 | 0.14 |
| Engagement | | | 213 | 91 | 17 | 50 | 11 | 24 | 0.21 | 0.81 | 1.00 | 0.43 | 0.08 | 0.23 | 0.05 | 0.11 |
| Familiarity | | | | 338 | 12 | 83 | 17 | 45 | 0.15 | 0.79 | 0.27 | 1.00 | 0.04 | 0.25 | 0.05 | 0.13 |
| Known-item | | | | | 202 | 85 | 0 | 1 | 0.07 | 0.50 | 0.08 | 0.06 | 1.00 | 0.42 | 0.00 | 0.00 |
| Metadata | | | | | | 264 | 11 | 26 | 0.15 | 0.67 | 0.19 | 0.31 | 0.32 | 1.00 | 0.04 | 0.10 |
| Novelty | | | | | | | 34 | 10 | 0.24 | 0.76 | 0.32 | 0.50 | 0.00 | 0.32 | 1.00 | 0.29 |
| Socio-cultural | | | | | | | | 134 | 0.20 | 0.74 | 0.18 | 0.34 | 0.01 | 0.19 | 0.07 | 1.00 |

example as well as description, which is a form of querying that is not supported by any current systems.

Other frequently labeled aspects include *metadata* (28%), *engagement* (23%), and *known-item* (21%). On the LT forum, *metadata* is an interesting aspect. When searching a catalog, metadata is often used to find specific books or books by a certain author, but such straightforward lookup tasks are not typically posted on the forums. Of the 264 topics labelled with *metadata*, only 22 (8%) have no other relevance aspect. These topics typically ask for recommendations on which books to read from specific authors, publishers, or series, or for the proper sequence in which to read a set of books. In most cases, *metadata* is combined with other aspects, and is used to focus the suggestions. *Engagement* is something that is hard to express through a search engine query. For instance, how can a user search for text books that are ‘funny’ or for books that challenge the reader’s own views on a topic? Such complex relevance criteria may be a reason to ask for suggestions on the LT forum. The same holds for *known-item* topics where the user can only recall certain elements of the plot or attributes of certain characters. Most book search services are of limited use for such known-item topics, as they do not allow full-text search. Forum members, however, may be able to help out with such requests. *Accessibility*, *novelty* and *socio-cultural* aspects are less prominent in our sample.

The rest of Table 1 shows the distribution of the relevance aspects and their occurrences. We can see a pattern emerging of relevance aspects being combined with either *content*, *familiarity*, or both, forming groups of topics clustered around these two aspects. *Known-item* requests are an exception as they seem to be a separate group. *Content* requests tend to be more typical of search tasks, as they provide a specific description of the desired books. The *familiarity* aspect seems related to recommendation-oriented tasks. The other aspects are more contextual in nature: dealing with books for certain scenarios (e.g., waiting at an airport, selecting reading material for a book club), for certain age groups or personality traits (e.g., trying to get a spouse to pick up reading), or certain moods (e.g., books that are comforting or challenge ones views). Dealing with such contextual information is an active research topic for both search [9] and recommender systems [1].

Summarizing, in a large sample of book requests annotated by their relevance aspects we find that most requests combine multiple aspects. We observed the largest clusters

Table 2. Topic groups in terms of example books and requested prose genre

| Feature | KI | Cx | F | Co+F | Co | All |
|----------------------|------|------|------|------|------|------|
| Example books | 0.08 | 0.26 | 0.54 | 0.50 | 0.16 | 0.27 |
| Genre Fiction | 0.77 | 0.29 | 0.49 | 0.53 | 0.35 | 0.50 |
| Non-fiction | 0.06 | 0.06 | 0.10 | 0.15 | 0.26 | 0.16 |
| Mix | 0.03 | 0.21 | 0.10 | 0.13 | 0.15 | 0.12 |
| Uncertain | 0.13 | 0.44 | 0.30 | 0.19 | 0.24 | 0.23 |

around *content* and *familiarity* aspects, or both, and the *known-item* class. In the next section, we will divide the requests into different groups based on these relevance aspects and study them in more detail.

5 Impact of Content and/or Familiarity

In this section, we investigate RQ3: How do different types of topics relate to user characteristics, human recommendations, and retrieval and recommender system results?

5.1 Grouping Topics on Relevance Aspects

In the previous section we saw a prevalence of *content* and *familiarity* aspects, in isolation and in combination, suggesting a grouping of the requests based on these relevance aspects. With the *known-item* requests as a separate group and the four logical combinations of *content* and *familiarity* aspects, this results in the following five topic groups:

Known-item (KI) contains all 202 *known-item* topics. This is the most content-specific information need, but different from the rest in that the user wants a specific book.

Context (Cx) contains all 78 topics without *content*, *familiarity*, or *known-item*. There are no content-based aspects on which to base document similarity.

Familiarity (F) contains 66 topics with *familiarity*, but no *content*. Users search for books similar to a specific (set of) book(s) or genre(s). Document similarity is underspecified, i.e., the user gives no content aspects to base similarity on.

Content and Familiarity (Co+F) contains 260 topics with both *content* and *familiarity* aspects, articulating explicit and implicit topic aspects. The similarity of the desired books is expressed at the level of books as well as at the finer-grained level of specific textual aspects of the books.

Content (Co) contains 338 topics with *content*, but no *familiarity*. Users are searching for books matching specific content aspects. Here, document similarity is more explicit, corresponding to a more specific information need.

5.2 Analysis of Genre, Popularity, and Personal Catalogues

To understand how our groupings correspond to actual differences in the nature of topic groups, we compare them on characteristics of the request, the requester, and the suggested books: (1) the presence of example books in *touchstones*, (2) the genre of books

Table 3. Catalogue size in requester catalogue (median of each topic group)

| Feature | KI | Cx | F | Co+F | Co | All |
|------------|----|-----|-----|------|-----|-----|
| Pre-topic | 0 | 38 | 104 | 100 | 177 | 84 |
| Post-topic | 4 | 80 | 81 | 65 | 108 | 65 |
| Total | 16 | 155 | 195 | 201 | 415 | 197 |

they target, (3) the size of the requester's book catalogues, and (4) the popularity of books in the requester's catalogue or those suggested at the forum.

Providing Example Books. For some topics, requesters add example books to their initial post using *touchstones*. These examples can serve different purposes: (1) positive examples of what they want (more of); (2) negative examples that match some relevance aspect(s), but not all; or (3) examples of what they have already read. Out of the 944 topics in total, only 256 (27%) have example books in the initial request, as shown in Table 2. We expect that examples are common among **F** topics based on previous reading experiences, and rare among **KI** and **Co** topics. These expectations are supported by the relevance aspects: the majority of the **F** topics include examples (54%), whereas only 8% of the **KI** topics contain examples. This lends credence to our decision to split the topics into groups based on the *content* and *familiarity* aspects.

Genre. Our annotators indicated whether requests were for fiction, non-fiction, or both. Table 2 shows that, of the 944 topics in total, 469 (50%) asked for suggestions on fiction books, 150 (16%) on non-fiction, and 113 (12%) on both fiction and non-fiction. For 212 topics (22%) the annotator could not tell. Fiction was the most common prose genre for **KI** topics at 77%, whereas only 6% of the topics were non-fiction. Fiction was also common for the **F** group at 53%. **Cx** topics have no specific content aspects, so it makes sense that mixed-genre topics and ambiguous topics are more common. In contrast, the **Co** topics are focused on non-fiction books more frequently than the other topic groups. Intuitively, this makes sense, as the topical content is arguably the main reason for reading a non-fiction book. Requests for fiction books are more likely to refer to examples, because what one is looking for in fiction may be more difficult to express and less explicitly related to the topical content of the book. This provides further evidence that the criteria for the topic groups are meaningful for analysis.

Cataloguing Behavior. Next, we count how many books the topic creator catalogued before posting the request (pre-topic), after posting it (post-topic), and in total; results are listed in Table 3. **KI** topics are often posted by LT members who have no books in their catalogue. Private profiles are an unlikely explanation for this, as these are rare. It seems these LT members use the forums mainly as a search engine and discussion board instead of as a tool for managing their book collections. Requesters of **Cx** topics tend to have small pre-topic catalogues, but add more books afterwards. These may be relatively new users with limited reading experience and have difficulty describing in detail what books they are looking for. Instead, they describe the context in which they want to read books. **F** and **Co+F** topics tend to come from more active users who have

Table 4. Median book popularity in requester catalogue, forum suggestions, system results

| Feature | | KI | Cx | F | Co+F | Co | All |
|---------------------|----------------------|------|------|------|------|------|------|
| Requester catalogue | topic group median | 47 | 60 | 41 | 56 | 39 | 46 |
| | topic group mean | 91 | 98 | 76 | 92 | 72 | 86 |
| | topic group std.dev. | 101 | 97 | 78 | 90 | 79 | 89 |
| Forum suggestions | topic group median | 174 | 681 | 531 | 235 | 192 | 237 |
| Retrieval | Top 10 | 55 | 58 | 107 | 57 | 42 | 53 |
| | Top 1000 | 23 | 24 | 18 | 25 | 20 | 21 |
| Recommender | Top 10 | 5146 | 5685 | 6022 | 4028 | 5163 | 5997 |
| | Top 1000 | 1076 | 958 | 985 | 908 | 852 | 959 |

over 100 books pre-topic and remain active cataloguers post-topic. This suggests they know what they like and that their needs have become more specific, but are still broad enough that they only need to implicitly describe what they want by giving examples. We speculate that users with **Co** topics are typically heavy readers, who have large pre-topic catalogues and remain very active users. They can explicitly describe what they are looking for and may in fact leave out examples to avoid responders from picking up on the wrong similarity clues from those examples.

Book Popularity. Chandler [6] examined the different strategies of GoodReads users for discovering new books to read and how these relate to the popularity of discovered books. They found that the popularity distribution of books discovered through search has a long tail of less popular books, whereas for GoodReads recommendations the distribution is concentrated around the mid- to high-popularity books.

How are the five topic groups related to the popularity of books discussed on the LT forums? The popularity $Pop(d)$ of a book d is the number of users who have d in their catalogues in our profile crawl. The top half of Table 4 shows the median popularity of books in searchers' catalogues and the forum suggestions. The catalogues of requesters tend to have a mix of popular and obscure books—the topic group mean is higher than the median indicating the distribution is skewed with a minority of highly popular books. There is no big difference between the popularity distributions of requesters with **F**, **Co** and **Co+F** topics. For the forum suggestions we see larger relative differences between topic groups, however. Forum members suggest more popular books for **Cx** and **F** topics than for **Co** and **Co+F** topics. The popularity of suggested books diminishes as content-specificity increases. For **KI** topics—the group with arguably the highest content specificity—suggestions are even less popular. Relating this to the findings of Chandler [6], suggestions for **Co** and **Co+F** topics are closer to search-related discoveries and **F** and **Cx** closer to recommendation-related discoveries. In terms of suggestions, book search on the LT forums seems to have a mix of search and recommendation-oriented tasks.

5.3 Retrieval and Recommendation Results

We analyse the books returned by standard retrieval and recommender systems for the forum topics, and compare them to the actual suggestions given by LT forum members.

Table 5. Performance evaluation of retrieval, recommender and best case fusion results

| nDCG@10 | KI | λ | Cx | λ | F | λ | Co+F | λ | Co | λ | All | λ |
|----------------|-----------|-----------|-----------|-----------|----------|-----------|-------------|-----------|-----------|-----------|------------|-----------|
| Retrieval | 0.207 | | 0.086 | | 0.101 | | 0.050 | | 0.088 | | 0.095 | |
| Recommendation | 0.002 | | 0.007 | | 0.008 | | 0.009 | | 0.005 | | 0.006 | |
| Fusion | 0.215 | .85 | 0.101 | .70 | 0.106 | .85 | 0.056 | .70 | 0.090 | .80 | 0.098 | .75 |
| MRR | | | | | | | | | | | | |
| Retrieval | 0.249 | | 0.153 | | 0.188 | | 0.122 | | 0.161 | | 0.163 | |
| Recommendation | 0.003 | | 0.037 | | 0.033 | | 0.038 | | 0.018 | | 0.025 | |
| Fusion | 0.256 | .85 | 0.176 | .70 | 0.219 | .70 | 0.146 | .60 | 0.167 | .80 | 0.171 | .75 |

For the retrieval system, we use the Amazon/LibraryThing collection [2] that is also used in the INEX Social Book Search Track [15]. This collection contains book metadata for 2.8 million books, including formal metadata (title, author, publisher, publication date), professional subject metadata (subject headings, Dewey Decimal System codes) and user-generated content (Amazon user reviews, LT user tags). We use Indri [11] and index all content with Krovetz stemming and stopword removal using a list of 319 stopwords. Specifically, we use a standard Language Model run with Dirichlet smoothing ($\mu = 2500$), using a combination of the thread title, a query provided by the annotators and the name of the discussion group as a query. This combination gives the best performance with standard Language Model settings.

For the recommender system we use a set of 84,210 user profiles, with information on which books the user catalogued and when, to compute nearest neighbours. We represent each user by a vector of book IDs and compute tf-idf similarity using GenSim [17]. The recommendation score for a book is the sum of the similarities of the individual neighbours who catalogued that book. We use a standard k -NN model run with recommendations from the 100 nearest neighbours based on catalogue similarity.

The lower half of Table 4 shows the median book popularity of returned results of the two runs. The rankings show a strong popularity effect for the recommender system, with the retrieval systems picking up less popular books in general than the recommender system. The popularity effect of recommender systems is also known as the ‘‘Harry Potter’’ problem [16]. The top of the rankings show relative differences between content and familiarity topics, with especially recommender systems returning more popular books for **F** topics than for **Co+F** and **Co** topics. The query terms of **Co** topics target less popular books than query terms of **F** topics, and their users are similar to users with a smaller fraction of highly popular books in their catalogues. In terms of book popularity, forum suggestions are roughly as popular as retrieval results. This is in line with earlier results on recommendations from friends as found on GoodReads [6]. Even though forum members do not know the requester personally, the statement of request is comprehensive enough for them to target the right types of books.

Finally, we look at system performance of the retrieval and recommender systems on the forum requests and suggestions, based on the Qrels from the INEX 2014 SBS Track for evaluation. We focus on basic retrieval and recommender models to observe relative performance of the two approaches on the various request types under well-understood conditions. In future work we will explore different models and query and user models in more detail. In addition to the baseline, we assess the potential of hybrid systems

merging the results lists of both baselines using a weighted sum, $S_{fusion}(d, q) = \lambda \cdot S_{retrieval}(d, q) + (1 - \lambda) \cdot S_{recommendation}(d, q)$. The performance scores are shown in Table 5. As expected, the retrieval system outperforms the recommender system, but there are differences between the topic groups. The recommender system scores relatively well on the **Cx**, **F** and **Co+F** topics, while the retrieval system performs relatively better on **KI** and **Co** than on **Cx** and **Co+F**. More importantly, on all topic groups, fusion of the results lists leads to improvements. Topics with *familiarity* or non-content aspects show the largest relative improvement. Finally, as a combination of recommendation and retrieval aspects, the **Co+F** topic set shows through its λ value of 0.6 that a more balanced fusion produces the best results. This suggests that the type of request plays an important role in the design of book discovery systems.

Summarizing, we analyzed topic groups related to known-item search and the logical combinations of content and/or familiarity. We observed varying degrees of combinations of contextual search and recommendation aspects. In terms of cataloguing behaviour, the content-specificity of requests is related to the size of the requester's catalogue. In terms of popularity, forum suggestions for topics with *content* aspects are more similar to retrieval results, and those for topics with *familiarity* aspects more similar to recommendation results. We demonstrated there is room for improvement by combining retrieval and recommendation approaches.

6 Conclusions

The aim of this paper was to investigate complex search requests in social media, in particular focusing on book search as observed on the LibraryThing discussion forums.

First, we found that the LT forums provide an unobtrusive way to study realistic, complex book search requests, which show a broad variation in topical and contextual relevance aspects. Second, we annotated the relevance aspects expressed in book requests at the LT forums. We found that the two dominating aspects are the *content* of the book and looking for *familiar* reading experiences, while other aspects are more oriented toward the reading context. The combination of content, context, and examples in a search request is a form of querying that is not supported by any current systems. Third, we found that these topic groups based on content and familiarity aspects can be differentiated by whether the requesters provide example books, what genre they are looking for (fiction or non-fiction), their cataloguing activity, and the popularity of the suggested books. Retrieval systems can effectively use the content aspects of the search requests, and recommender systems can pick up signals in the requester's catalogue. We demonstrated the possibility for improvement when combining both approaches, in particular for topic groups where context and familiarity play a role. This suggests that the request type has an important role to play in the design of book discovery systems.

Our analysis was focused on the book search domain, yet similar rich profiles and contextual information is available in many modern search scenarios, in particular in mobile search and increasingly aggregated to mixed-device search scenarios. Research access to such mobile search logs and social media data is difficult due to privacy and commercial constraints, making the more constrained and less sensitive book search domain an attractive alternative to study many aspects of complex contextualized search.

We highlighted the diversity of complex search requests, and observed a mixture of content and context going beyond currently existing systems. This is an important first step toward the development of novel information access systems that blend traditional search and (hybrid) recommendation approaches into a coherent whole.

Acknowledgments. This research was supported by the Netherlands Organization for Scientific Research (README project, NWO VIDI # 639.072.601; ExPoSe project, NWO CI # 314.99.108).

References

- [1] Anand, S.S., Mobasher, B.: Contextual recommendation. In: Berendt, B., Hotho, A., Mladenic, D., Semeraro, G. (eds.) *WebMine 2007*. LNCS (LNAI), vol. 4737, pp. 142–160. Springer, Heidelberg (2007)
- [2] Beckers, T., Fuhr, N., Pharos, N., Nordlie, R., Fachry, K.N.: Overview and results of the INEX 2009 interactive track. In: Lalmas, M., Jose, J., Rauber, A., Sebastiani, F., Frommholz, I. (eds.) *ECDL 2010*. LNCS, vol. 6273, pp. 409–412. Springer, Heidelberg (2010)
- [3] Bhatia, S., Mitra, P.: Adopting Inference Networks for Online Thread Retrieval. In: *Proceedings of AAAI 2010*. AAAI Press (2010)
- [4] Borlund, P.: The Concept of Relevance in IR. *JASIST* 54(10), 913–925 (2003)
- [5] Buchanan, G., McKay, D.: The bookshop: examining popular search strategies. In: *Proceedings of JCDL 2011*, pp. 269–278. ACM (2011)
- [6] Chandler, O.: *How Consumers Discover Books Online*. Tools of Change for Publishing, O’Reilly (2012)
- [7] Cosijn, E., Ingwersen, P.: Dimensions of Relevance. *Information Processing & Management* 36, 533–550 (2000)
- [8] Cunningham, S.J., Vanderschantz, N., Timpany, C., Hinze, A., Buchanan, G.: Social information behaviour in bookshops: Implications for digital libraries. In: Aalberg, T., Papatheodorou, C., Dobрева, M., Tsakonas, G., Farrugia, C.J. (eds.) *TPDL 2013*. LNCS, vol. 8092, pp. 84–95. Springer, Heidelberg (2013)
- [9] Dean-Hall, A., Clarke, C.L.A., Kamps, J., Thomas, P., Simon, N., Voorhees, E.: Overview of the TREC 2013 contextual suggestion track. In: *Proceedings of TREC 2013*. NIST (2013)
- [10] Elsas, J.L., Carbonell, J.G.: It Pays to be Picky: An Evaluation of Thread Retrieval in Online Forums. In: *Proceedings of SIGIR 2009*, pp. 714–715. ACM (2009)
- [11] Indri. Language modeling meets inference networks (2014), <http://sourceforge.net/projects/lemur/>
- [12] Koolen, M., Kazai, G., Kamps, J., Preminger, M., Doucet, A., Landoni, M.: Overview of the INEX 2012 Social Book Search Track. In: *Proceedings of INEX 2012*. LNCS, Springer
- [13] Koolen, M., Kamps, J., Kazai, G.: Social Book Search: Comparing Topical Relevance Judgements and Book Suggestions for Evaluation. In: *Proceedings of CIKM 2012*, pp. 185–194. ACM (2012a)
- [14] Koolen, M., Kazai, G., Kamps, J., Doucet, A., Landoni, M.: Overview of the INEX 2011 books and social search track. In: Geva, S., Kamps, J., Schenkel, R. (eds.) *INEX 2011*. LNCS, vol. 7424, pp. 1–29. Springer, Heidelberg (2012)
- [15] Koolen, M., Kazai, G., Preminger, M., Doucet, A.: Overview of the INEX 2013 social book search track. In: *Proceedings of the CLEF 2013 Evaluation Labs and Workshop* (2013)

- [16] Linden, G.: Geeking with Greg: Early Amazon: Similarities (2006), <http://glinden.blogspot.nl/2006/03/early-amazon-similarities.html>
- [17] : Software Framework for Topic Modelling with Large Corpora. In: Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks, pp. 45–50. ELRA (2010)
- [18] Reuter, K.: Assessing aesthetic relevance: Children’s book selection in a digital library. *JASIST* 58(12), 1745–1763 (2007)
- [19] Ross, C.S.: Finding without seeking: the information encounter in the context of reading for pleasure. *Information Processing & Management* 35(6), 783–799 (1999)
- [20] Saracevic, T.: Relevance Reconsidered. In: Proceedings of COLIS 2, pp. 201–218 (1996)
- [21] Saracevic, T., Kantor, P.B.: A Study of Information Seeking and Retrieving. I, II, III. *JASIS* 39(3), 161–216 (1988)