



City Research Online

City, University of London Institutional Repository

Citation: McKay, D., Buchanan, G. & Chang, S. H. (2015). Tyranny of Distance: Understanding Academic Library Browsing by Refining the Neighbour Effect. Paper presented at the 19th International Conference on Theory and Practice of Digital Libraries, 14-09-2015 - 18-09-2015, Poznan, Poland.

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <http://openaccess.city.ac.uk/12536/>

Link to published version:

Copyright and reuse: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Tyranny of Distance: Understanding Academic Library Browsing by Refining the Neighbour Effect

Dana McKay¹, George Buchanan², Shanton Chang¹

University of Melbourne, Parkville, VIC 3010, Australia
Centre for HCI Design, City University London, EC1V 0HB, U.K.
dmckay1@student.unimelb.edu.au; george.buchanan.1@city.ac.uk;
shanton.chang@unimelb.edu.au

Abstract Browsing is a part of book seeking that is important to readers, poorly understood, and ill supported in digital libraries. In earlier work, we attempted to understand the impact of browsing on book borrowing by examining whether books near other loaned books were more likely to be loaned themselves, a phenomenon we termed the neighbour effect. In this paper we further examine the neighbour effect, looking specifically at size, interaction with search and topic boundaries, increasing our understanding of browsing behaviour.

Keywords

Browsing, books, libraries, information seeking, classification systems, log analysis.

1 Introduction

Participants in a 2007 study comparing physical and digital libraries noted there is no digital analogue of library shelves for book seeking, particularly in terms of serendipitous discovery [1]; that same year Rowlands noted the dearth of literature on book selection [2]. Despite later research on book selection, reader behaviour at the library shelves remains largely mysterious.

Key models of information seeking include elements that focus on browsing and exploration [3, 4]. The literature on book seeking also shows that readers consider browsing an important part of book seeking [1, 5, 6]. Readers in both academic [5, 7] and public [8] libraries note lack of browsing support as a reason for avoiding the use of ebooks. Libraries increasingly offer ebooks alongside or in place of print, so avoiding them is likely to negatively affect readers.

Despite readers' insistence that browsing is important, until last year there was little literature on its impact on book use. Our prior work [9] leveraged two established digital library techniques—examining physical libraries [1, 6, 7] and transaction log analysis [10, 11]—to determine whether physical layout of library shelves affects book borrowing patterns. In that work we specifically examined whether proximity to a loaned book increases the chance a book will itself be loaned: we found a strongly significant increase in this likelihood that we termed the neighbour effect. This paper refines the neighbour effect by examining day-of-week variance, the distance between relatively nearby books borrowed on the same day (co-borrowed books), the impact of topic boundaries, and the distance between co-borrowed books in search results. Given the prevalence of complaints about digital library browsing systems [1, 6, 12], developing a better understanding of browsing is a vital step to improving users' experience of DLs. This paper thus aims to increase and improve our understanding of browsing. Section 2 presents the background literature, Section 3 our methodology. Section 4 presents the results, adfa, p. 1, 2011.

© Springer-Verlag Berlin Heidelberg 2011

which are compared to the literature in Section 5. We finally draw conclusions and suggest future work in Section 6.

2 Background literature

In this section we will first present the work on browsing in the context of information seeking behaviour, then cover browsing and browsing technologies. Next we will cover the literature on book selection; finally we summarise our previous work that this paper extends.

2.1 Browsing in Human Information Seeking Behaviour

Browsing is a central activity in many major models of information seeking (e.g. [3, 4]). The common models are primarily linear, with elements that address search and triage. Browsing in these models is closely interleaved with search and used for exploration and triage. Beyond these models, serendipity—the opportunity to discover information one otherwise would not—is a key attribute of browsing. Savvy information seekers leverage their surroundings to increase the likelihood of serendipity [13], a strategy also used at the library shelves [1, 6, 14].

2.2 Browsing and Browsing Technologies

Bates—who wrote the seminal work on browsing in electronic environments [15]—also provides a working definition of browsing grounded in a broad literature [16]. She described it as the process of glimpsing a ‘scene’ (a large collection of potentially interesting objects) sequentially examining objects of interest, and retaining or discarding these. Given this definition, library shelves are, as noted in earlier work on libraries [14], ideally suited to browsing.

More than 20 years since Bates’ seminal paper [15], online browsing systems remain under-researched. It may be argued that search results [17], and faceted search in particular [18], represent opportunities to browse online, however both the limited number of results presented, and the need to navigate away from the ‘scene’ (or search results) to examine items in depth contravene Bates’ definition of browsing [16]. These limitations also negatively affect the close interleaving of search and browse found in information seeking models [3, 4]. The need to prime search-based systems with query terms—a difficult task with imprecise information needs [19]—also limits their usefulness [20].

The literature on browsing-specific systems is limited. A 1995 paper presents an early shelf metaphor system intended for children [10], and a 2004 paper presents three browsing tools within the Greenstone DL system [21]. These tools are, however, only proofs-of-concept. Book browsing systems have been more common in recent research: e.g. exploiting a shelf metaphor [14]; using the non-bibliographic book features readers say they use in decision making [22], and Pearce’s iFish tool [23], which creates browsing sets based on user-specified preferences. Whichbook (<http://www.openingthebook.com/whichbook/>) is one of a number of increasingly common commercial systems. None of these tools however, is rooted in a detailed understanding of browsing behaviour, and none has yet been the subject of rigorous evaluation. The emergence of commercial systems both reinforces the need for research-based approaches, and renders the topic of browsing relevant and timely.

2.3 The Book Selection Process

Rowlands’ 2007 critique of the limited research on book selection [2] has been followed by a steady growth in that literature. The process of choosing a book in a library can be divided into

5 components: identifying a need for a book; searching (this step may be skipped); locating books of interest; choosing among them; and reading or otherwise using selected books. Catalogue book search, while it has major usability problems [11, 24], is well-studied elsewhere and not the focus of this work; we will not discuss it further. Similarly, the nature of reading—while an interesting open research question [25]—will not be further discussed here.

There is some work on how readers identify both fiction [8, 26] and non-fiction [27, 28] books of interest. Personal recommendations and shelf browsing are frequently reported, while search appears rare. There is a limited but growing literature on how readers identify and select books at the shelves. Research on children [29, 30] notes that they focus on eye-level shelves, and that shelf order affects the books they select. Adults are also affected by shelf height, though less than children are [7]. Like children [30], they struggle with some aspects of shelf layout in libraries [20], but adults also exploit it, using librarian created displays or recent returns [7, 8, 26] as information resources. Savvy library users value the shelves as a finding aid and relevance cue, and note there is no online equivalent [1]. Participants in numerous studies have said they value the opportunity to browse shelves, even giving it as a reason for avoiding the use of ebooks [1, 5, 7].

Decision-making at the shelves is progressive—looking, then looking more closely, then taking books off the shelf [7, 29]. Readers flip pages, and use index, blurb, and images to determine relevance [6, 7]. These cues are also used in ebooks, but other non-bibliographic cues that are not replicated online—such as dust, book size, and location—are also used [1, 6].

The process of examining shelves and choosing books mirrors Bates' definition of browsing [16]; indeed it seems likely that libraries, classification schemes and shelves have evolved to create a physical browsing engine [31]. The importance to users of browsing is supported by numerous small studies, including [1, 6, 7]. A larger study from 1993 demonstrates that browsing also affects the books readers select: over half of those who located one book identified by searching borrowed at least one further book [32]; a contemporaneous study noted that books near each other on the shelves were likely to be borrowed together [33]. Given the age of these studies and the ascendance of search in the intervening time, it seems reasonable to ask whether browsing still holds such sway: this paper investigates the impact of browsing on book loans.

2.4 Our Previous Work

The only recent study of the impact of browsing—or more specifically shelf location—on loans within academic libraries is our own [9]. We used a large publicly available circulation dataset from the OCLC [34], and selected six libraries based on a set of criteria to ensure their broad similarity. We created a shelf-sorted book set including circulation data for each library.

Circulation data only records the most recent loan for each book; given this limitation we used two tests to look for a neighbour effect. We compared the number of loans among the ten nearest neighbours (five either side) of loaned and unloaned books on the final date recorded in the data, and for randomly selected loans we compared the number of loans among the nearest neighbours on the loan date and on the day before. Both tests showed a strongly significant neighbour effect, supporting our hypothesis that browsing influences loan patterns. One significant limitation remains: without patron data, we cannot prove, for any co-borrowing, that it is the result of use by a single patron. There is a preponderance of evidence, though, in the form of observation [1, 6, 32], log analysis [33] and user self-reporting [5, 8, 26, 27] that makes browsing a logical explanation for a significant proportion of this activity.

The tests we used in this early work were a fairly blunt instrument: they examined a fixed number of books, and did not take search, day-of-week effects or topic boundaries into account. This work aims to refine our method, deepening our understanding of browsing along these axes.

3 Methodology

This paper extends the methodology of our previous work [9]. We will briefly review our dataset, then describe four new tests employed to examine the patterns of co-borrowing.

3.1 Dataset

We used the same six sample libraries as our previous work: Cedarville, Dennison, Case Western Reserve University (CWRU), Oberlin, Ohio Northern (ON) and Ohio State University (OSU). For this study, we omitted Oberlin's Dewey collection, though we retained Oberlin's LC collection: the former has a fragmented floor layout¹ limiting browsing opportunities. We mitigated against overwritten loans (circulation data stores only the most recent loan) by using data from the end of the dataset's collection period, in many cases the final week. The final week of data for OSU, however, had less than half the annual average number of loans per week (650, in comparison with 1592) so we used the penultimate week's data (1071 loans).

3.2 Tests

We conducted four separate tests on our datasets. Our tests primarily involve the ten closest books to an individual text: five to its left and five to the right (duplicates are discarded—they form <3% of any collection and usually <1%). We refer to this set, which we also used in our previous work, as N_{10} . The only study similar to ours used a width of two [33], or N_4 .

Shape of the Neighbour Effect: Books within the N_{10} set will usually be on the same shelf as a target book: 10 books is simply not enough to see the books above and below a borrowed book, nor is it wide enough to tell us when the neighbour effect disappears. Using the final week sets described above, we calculated for each borrowed book the distance to its nearest borrowed neighbour, and the number of books borrowed at each distance out to 150 books either side. We examined this data to see how rapidly the neighbour effect falls away, and whether there are secondary co-borrowing peaks that could account for books above and below a target book.

Day of the Week Effects: There is a known 'day of the week' effect on search behaviour in information retrieval [35, 36], representing changes in user behaviour based on context. The same effects in browsing data would reinforce our impression that the neighbour effect is a product of intentional human behaviour, and not random co-borrowing. We compared the strength of the neighbour effect at N_{10} on the busiest and least busy days of the week (based on the final year of loans data), and compared weekdays with weekends. To ensure that we had a representative time sample we ran this comparison over all complete weeks in 2008 (the final year of data collection—data ceases for all collections in late April or early May).

Comparing Search and Browsing: One explanation for co-borrowing is that users may identify books individually in search and select them independently of one another. While early studies [32, 33] do not support this, the increasing dominance of search [23, 28] suggests that search could be responsible for these results. We created a stop-worded log-rule search index [37] for the titles of the books (the only metadata we had, but data that is frequently used in book search [11, 24, 38]).

We then collated all the possible pairs of books borrowed on the same day in the final week. Those pairs within their N_{10} shelf-set were identified first to create the shelf co-borrowing set. We did two searches against the titles index: a title search, and a keyword search on their shared title words. Where one book and another loaned on the same day both appeared in the

¹ See <http://www.oberlin.edu/library/main/2.html>

top ten search results (the only results most users look at [37]), they were considered to be search co-borrowings. This created both a title search set and a shared title keyword set of co-borrowings. The difference in result ranking of two results was deemed to be their search distance. For any book pairs that had both shelf (within N_{10}) and search proximity, we determined which distance was shorter (and hence a better explanation for the co-borrowing).

Topic Boundaries: Earlier work suggests that shared topic is almost as important as shared shelf location, even where books are not co-located [33]. Topic and shelf location are intertwined in academic libraries [31]: to investigate this relationship we examined loans near topic boundaries. Topic boundaries were categorised as 3, 2, 1 or 0; rank 3 indicated a difference in the first digit of the call number, 2 the second, 1 the third and 0 none before this. We identified the borrowings near topic boundaries, then did a closer analysis where there was sufficient data.

4 Results

This section reports results for each of the four individual tests described in Section 3. We subsequently summarize our findings and point to limitations of our study.

4.1 Shape of the Neighbour Effect

To examine the width of the neighbour effect, we first calculated the distance to nearest co-borrowed book for all books loaned in the final week set. Not every visit to the library will result in co-borrowings, and borrowing rates per book are low (*annual* circulation ranged from 0.08 to 0.18 loans per book for all collections [9]). The data shows a log-log distribution often seen in information behaviour [37]. Nearest neighbour data is shown in Table 1, below:

Table 1. Shelf distance to nearest loan in number of books showing first three quartiles

	Median	Mean	Std.dev	Q1	Q2	Q3
Cedarville	571	1118	178	89	181	640
CWRU	2930	5437	1300	343	1297	3471
Dennison	1507	2177	898	189	919	2063
Oberlin	4917	9833	1910	189	1858	6198
ON	2351	3343	1056	170	992	2972
OSU	4281	8140	1795	471	1773	4585

The distribution of loans suggests a neighbour effect: many loans are close to their nearest borrowed neighbour. Browsing likely accounts for many of the close-by loans, but is unlikely to account for loans hundreds of books distant. There is one possible exception: examining nearest neighbour loans in groups of 50, four libraries (Dennison, Oberlin, ON and OSU) show unexpected peaks between 650 and 700 or 700 and 750, rather than the expected long flat tail. One possible explanation is borrowing from the ends of shelves (as seen in [7]), though this is speculative in the absence of detailed information about physical library layout.

After looking only at nearest neighbour information, we considered the data with respect to all co-borrowings at the micro level (within N_{20}) and the macro level (up to N_{300}).

Examining all neighbours at the micro-level reveals that the rate of co-borrowings falls rapidly with distance for four libraries, but Cedarville and OSU show flatter patterns (Fig 1). The macro-level view allows us to test for borrowing from above and below a target book. Of course the number of books per shelf will vary, even within libraries, due to variations in both shelf and book width. Averaging over a large enough sample, however, may reveal a secondary

peak representing the books above/below an index book. These books, while closely co-located, are likely to represent greater topic difference than the books beside an index book.

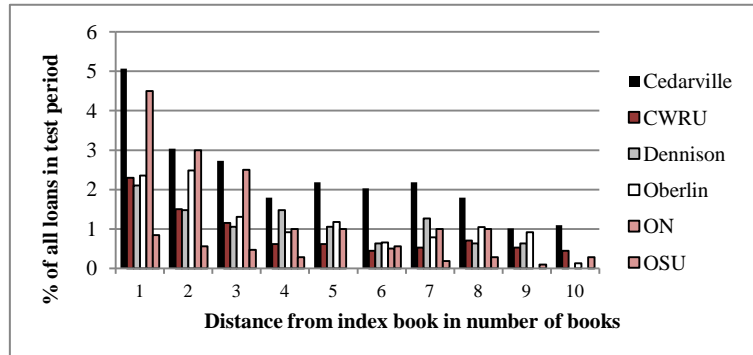


Fig. 1. All co borrowings within N_{20}

Studies of recommender systems show that users will choose items that fall a little outside what they expect [39]; and vertical neighbours on library shelves may have a similar effect. We examined all neighbouring loans (not just the nearest) as we didn't want to exclude patterns where users borrowed books both above or below an index book *and* on either side.

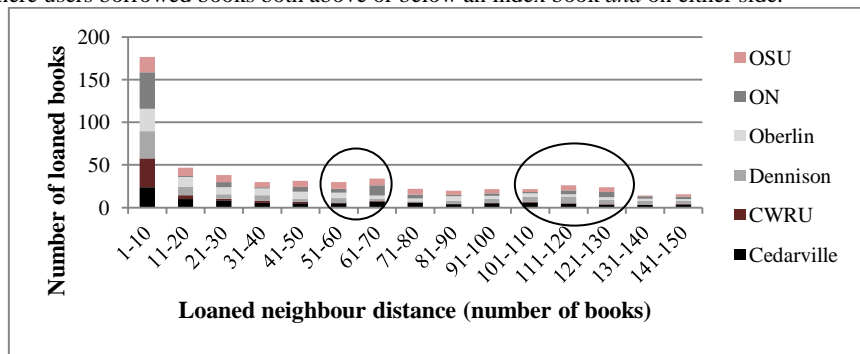


Fig. 2. Total borrowed books for all collections. Circles represent possible vertical browsing.

Figure 2 shows co-borrowing over 150 neighbouring books. Beyond 150, loans enter a long flat tail. As expected, there are secondary peaks (in all collections): most have two, one around the 50-70 mark, the next at 100-120. This suggests shelf wrap is c. 60 books, and supports previous evidence [5] that users browse 'three shelves' above and below a target book.

4.2 Day of the Week Effect

The literature on both search [35, 36] and academic library usage [40] evidences changes in usage between weekdays and weekends; we suspected that the same might hold true for browsing. In view of this we examined both the total number of loans on each day of the week for busy and quiet days (see Table 2), and compared weekdays to weekends (see Table 3). We also compared the prevalence of the neighbour effect in N_{10} against busy days and weekends.

In a χ^2 test of distribution there was clear non-random variance across days of the week at $p < 0.001$ for all libraries. For four libraries, the number of N_{10} loans was not consistent with the overall change in loan rate according to day (see Table 2). In each of those libraries neighbour

borrowing was more prevalent on quiet days than we would expect, showing a difference in borrowing behaviour: browsing is more likely on quiet days.

Table 2: Busy and quiet days, frequency of loans of the N_{10} neighbours of each borrowed book

	Busiest day	% of loans	Busy day N_{10} loans	Busy day loan totals	Quietest open day	% of loans	Quiet day N_{10} loans	Quiet day loans	P (B_{busy} vs Q_{quiet} N_{10})	χ^2 (df=1)
Cedarville	Mon	21.0	374	2466	Sat	6.0	128	834	<0.0001	80.49
CWRU	Wed	20.6	207	2846	Sat	4.8	93	668	<0.0001	26.85
Dennison	Mon	21.0	36	847	Sat	2.4	11	103	0.0110	6.46
Oberlin	Wed	17.8	121	14239	Sun	9.1	71	7329	0.4257	0.63
ON	Tue	21.5	43	4348	Sat	3.6	10	720	0.4424	0.59
OSU	Wed	19.3	317	47603	Sat	5.2	141	12719	<0.0001	25.06

A χ^2 test of distribution of loans between weekends and weekdays showed non-random variance of the overall loan rate at $p < 0.001$ in every library. Again, for four libraries the change in N_{10} borrowing was not consistent with the overall change in loan rate (see Table 3): N_{10} borrowing was more prevalent at the weekend, suggesting more browsing and less ‘grab and go’ borrowing happening then.

Table 3: Weekends vs weekdays for browsing: overall and neighbour loans

	W/day%	w/end N_{10} loans	w/end loans	w/day N_{10} loans	w/day loans	P (W/day vs w/end N_{10})	χ^2 (df = 1)
Cedarville	93.9	128	834	1484	10902	0.2063	1.60
CWRU	90.0	199	1364	1079	12329	<0.0001	44.63
Dennison	91.7	34	351	151	3871	<0.0001	23.39
Oberlin	75.3	194	1992	487	6073	<0.0001	816.77
ON	90.3	22	198	155	1846	0.2664	1.23
OSU	89.2	281	2641	1356	21887	<0.0001	880.79

Every library studied shows variance in N_{10} borrowing by day of the week that is inconsistent with overall change in loan rate. In each case this is an increase in N_{10} borrowing on a day that is otherwise ‘quiet’. This change suggests altered behaviour at the shelves—readers are more likely to browse during quiet times, perhaps because they have more time to do so.

4.3 The Impact of Search

One key argument against the neighbour effect is that search can account for much of it: users find items that are shelved close to each other during search, rather than from browsing at the shelf. As described in Section 3.2, we assess the overlap between the shelf and search co-

borrowing pairs (both title and shared keyword searches), and determined for each of these whether books were closer using search or browse (Table 4, left hand side).

While title and keyword search will not account for all possible search scenarios, they are the most common methods [24, 38] seen in libraries. In our data, search only accounts for a small proportion of shelf-based co-borrowings. The argument that search causes co-borrowings is therefore unsustainable. Shelf browsing remains the simplest explanation, accounting for more neighbouring and more total loans, both by rate and by closeness of the pairs.

Table 4: Shelf vs loan co-borrowing; search versus browse influence on co-borrowing

	Browse pairs	Title Search pairs	Keyword search pairs	Overlap	Shelf closer	Search closer	Ties	Neighbour loans explained by browse	Neighbour loans explained by search	Total loans	P (browse vs search)	χ^2 (df = 1)
Cedarville	374	212	97	95	54	24	17	333	189	1283	<0.0001	40.00
CWRU	140	71	54	36	28	7	1	132	71	1132	<0.0001	17.89
Dennison	68	32	40	25	12	5	8	51	22	475	0.0005	10.82
Oberlin	128	31	37	10	6	3	1	124	51	765	<0.0001	21.12
ON	46	24	15	11	8	1	2	43	14	200	<0.0001	14.00
OSU	46	14	6	1	1	0	0	46	13	1071	<0.0001	17.40

To further test whether search or shelf location was responsible for co-borrowing, we again took N_{10} sets and allocated any borrowed neighbours to search or browse, whichever formed the nearer pair. To maintain independent sets, ties were discarded. We then tested the search versus browse sets using a single-tailed χ^2 test (Table 4 right hand side). The data shows browsing having much more influence on co-borrowing than search. Given readers' penchant for convenience [5] this is unsurprising: books co-located on a shelf one is visiting require little extra effort, while going to another section of shelving takes additional time and effort [37].

4.4 Topic boundaries

We assessed the impact of topic boundaries on neighbouring loans to understand the interaction between topic and shelf layout. As described in Section 3, we defined three levels of topic boundary: from 3 (at the top level of classification), to 0. Our first discovery was that, notably, no book within 5 books of a top-level topic boundary had ever been loaned in any library. Loan numbers were also low when examining level 2 boundaries. We tested this distribution against the nominal likelihoods for each boundary, accounting for the number of topic boundaries and loan rates in each collection, using a two-factor chi-squared test: see the left of Table 5.

Topic boundaries appear to discourage loans: this may be related to how books are shelved close to boundaries: e.g. a level 3 boundary will often begin on a new shelf; given that users look less at top and bottom shelves [7, 29] it is not surprising that these are disadvantaged.

Even at the lowest topic boundaries, loans are limited. Dennison and ON have such small numbers of nearby loans—43 and 74 respectively—that we excluded them from further analysis. For the remaining four libraries we created samples of 100 loans for which N_{10} contained either a level-2 or level-1 boundary, then counted the number of co-borrowed books in each set (i.e. the total sample book count of each type was 1000). A two-tailed Fishers exact test was used to compare sets that included a level 1 boundary with sets with level 0 boundaries to see whether a topic boundary affects neighbouring loans (Table 5, right); results were significant.

Table 5: Number of loans occurring near classification boundaries at any time

	Loans near L2 boundary	Remaining loans	p	χ^2 (df = 2)	L1 Neighbour Loans	L0 Neighbour loans	p (L1 vs L0)
Cedarville	52	92092	0	19288.220	6	20	0.00089
CWRU	19	248644	0	6340.184	2	18	0.0004
Oberlin	28	236920	0	4532.100	1	9	0.0212
OSU	78	304510	0	11795.026	0	12	0.0005
Dennison	11	57830	0	1209.737	Discounted due to lack of data		
ON	71	36448	0	567.963	Discounted due to lack of data		

Clearly topic, as defined by shelf classification, is an important factor in co-borrowing: browsers are more likely to co-borrow within topic boundaries than across them, and are less likely to borrow near topic boundaries than in the rest of the collection.

4.5 Summary and limitations

Clearly this study shares some of the limitations of our earlier work: we cannot, from the data we have, prove that a single user is responsible for any co-borrowed pair. We are also not privy to users' motivations or behaviour at the shelf. Furthermore we do not have all the metadata required to fully address search; having only title data we cannot address author or other metadata. Nonetheless keyword searching and title searching are the dominant strategies in book search behaviour, and titles represent a linguistically user-friendly corpus [19].

We have examined different aspects of the neighbour effect to obtain new insights into browsing behaviour. Nearest neighbour loans have an overall power-law distribution. Secondary peaks in neighbouring loans were found at distances of 50-60 items. Search accounts for some co-borrowings on the shelf, but not most, and it plausibly generates its own co-borrowings. Across all our data, shelf-browsing accounts for more co-borrowings than title search. The amount of browsing varies by day of the week in all libraries. Finally, borrowing tends to occur within topic boundaries, suggesting that shelf layouts and topic divisions influence borrowing patterns. Each of these findings is novel, and represents a new understanding of human behaviour that could improve digital library design.

5 Discussion

Our findings can be divided into two major themes: the interaction of search, topic, and shelf, and the implications of reader context for browsing.

5.1 Shelf, topic and search

The dominant information discovery tool in every existing digital library system is search. Our data demonstrates, though, that search alone cannot fully meet the needs of information seekers. Search explains some co-borrowings in our data, so we would argue that it still plays a valuable role, but many more borrowings are poorly explained by search and well-explained by browsing. Search is extremely well covered by the DL literature (for a starting point see [37]); in contrast there is very little work on browsing. The existing work on browsing is very limited in terms of proposed solutions, and even where these have been proposed (for example [14, 20, 22, 23]) they have not been adequately assessed for their effectiveness. Search effectiveness is

assessed with standard metrics [37]; similar metrics do not yet exist for browsing. Our work demonstrates clearly that DLs need to support browsing, most currently do not.

Our data not only shows a need for browsing systems, it also suggests certain key characteristics of such systems. Earlier work suggests browsers check ‘three shelves before and after’ a target book—this is something in the range of 200-500 books. This is considerably more books than are typically shown in search results, and an order of magnitude more than is shown in some nascent browsing systems [23]. Research on recommender systems has shown that a few ‘unexpected’ recommendations improve user experience [39]. It is therefore likely that both volume and variety play a role in browsing.

Topic is another clearly important feature of books: our data shows that co-borrowing occurs primarily, though not exclusively, within topic boundaries. Conversely, both our data and the literature [7, 26] suggest that readers occasionally like to see ‘distant’ books. Physical shelves cannot be rearranged to meet individual user needs, yet we know rearranging shelves affects book selection [41]. Electronic shelves could—and should—leverage topic clustering that occurs across the boundaries of traditional classification schemes to ‘rearrange the shelves’ to aid discovery; the literature suggests this is likely to be useful [33]. The underlying data already exists within classification schemes [31], but no DL system has yet offered readers these options. Similarly DL systems could leverage topic classification schemes to offer readers a very few different-but-interesting books based on topic data: this approach would mimic the physical shelves, but, being data driven, would offer a higher chance of success.

5.2 Reader context

It is clear from our data that browsing, like search [35], is affected by day of the week. The way in which co-borrowing is affected by day of the week—increasing on quiet days, weekends, or both—suggests that reader context has a significant impact on browsing behaviour. This dovetails with earlier work: students report staying longer in the library when they do not have to rush to class [40], and academics report searching the catalogue inside the library when they are looking for inspiration, and outside the library when they are in a rush [6]. Similarly browsing the shelves [1, 6, 7] and serendipitous discovery [13] are activities that require time and attention, which, when a searcher is meeting an urgent information need, may not be available. This context-dependent approach to browsing has significant implications for DL design: users need to be able to ‘grab and go’ [7] when it suits them, but browsing facilities should be a visibly tempting way for users to spend any extra time they have.

6 Conclusions and Future Work

Not all who wander in libraries are lost: the library shelves afford browsing and serendipitous discovery in ways that simply do not exist in current DL systems. Our data demonstrates clearly that shelf arrangement and reader context—and thus browsing—have a clear impact on borrowing that cannot be explained by search. To be truly effective information resources, DL systems need to facilitate browsing; the literature shows that DLs are not meeting readers’ needs in this space. Our data further points to clear design implications: it is not enough to offer users a small number of books (as, for example in search results). The neighbour effect extends above and below a target book, as well as to the left and right—a range of over 200 books. Browsing systems also need to be optional: browsing happens under different circumstances to search, and users must be able to engage in the most appropriate information seeking strategy for their context. If they have time, they should be visibly tempted to linger and browse; if they do not, ‘grab and go’ should be an option.

Browsing is not just a necessity for DL systems, it is also an opportunity. DL systems have the potential to offer features that cannot exist in the library shelves: shelves can be rearranged

to reflect cross-classification topic clustering, for example. DL systems have the power to offer readers new and exciting paths to wander, and must leverage that if they are to provide adequate user experience.

Our study, of course, does not answer all (or even most) of the questions about browsing. To determine how much browsing occurs that does not result in loans, and what users motivations are requires a different kind of study; these questions remain future work.

Acknowledgements: This paper contains information from OhioLINK Circulation Data (<http://www.oclc.org/research/activities/ohiolink/circulation.htm>), made available by OCLC Online Computer Library Center, Inc. and OhioLINK under the ODC Attribution License (<http://www.oclc.org/research/activities/ohiolink/odcby.htm>). Travel to the conference was partially funded by a Google Travel Scholarship administered by the Department of Computing and Information Systems at the University of Melbourne.

7 References

1. Makri, S., Blandford, A., Gow, J., Rimmer, J., Warwick, C., Buchanan, G.: A library or just another informational resource? A case study of users' mental models of traditional and digital libraries. *JASIST* 58, 433-445 (2007)
2. Rowlands, I., Nicholas, D., Jamali, H.R., Huntington, P.: What do faculty and students really think about e-books? *Aslib Proceedings* 59, 489-511 (2007)
3. Kuhlthau, C.C.: *Inside the Search Process: Information Seeking from the User's Perspective*. *JASIST* 42, 361-371 (1999)
4. Marchionini, G.: *Information Seeking in Electronic Environments*. Cambridge University Press, Cambridge, UK (1995)
5. McKay, D.: Gotta keep 'em separated: Why the single search box may not be right for libraries. In: *CHINZ '11*, pp. 109-112. ACM, (2011)
6. Stelmaszewska, H., Blandford, A.: From physical to digital: a case study of computer scientists' behaviour in physical libraries. *IJDL* 4, 82-92 (2004)
7. Hinze, A., McKay, D., Vanderschantz, N., Timpany, C., Cunningham, S.J.: Book selection behavior in the physical library: implications for ebook collections. In: *JCDL '12*, pp. 305-314. ACM (2012)
8. Ooi, K.: *How Adult Fiction Readers Select Fiction Books in Public Libraries: A Study of Information Seeking in Context*. Master of Library and Information Studies. Victoria University of Wellington, Wellington, New Zealand (2008)
9. McKay, D., Smith, W., Chang, S.: Lend me some sugar: Borrowing rates of neighbouring books as evidence for browsing. In: *DL 2014*, pp. 145-154. IEEE (2014)
10. Borgman, C.L., Hirsh, S.G., Walter, V.A., Gallagher, A.L.: Children's searching behavior on browsing and keyword online catalogs: The Science Library Catalog project. *JASIS* 46, 663-684 (1995)
11. Lau, E.P., Goh, D.H.-L.: In search of query patterns: A case study of a university OPAC. *Inform Process Manag* 42, 1316-1329 (2006)
12. McKay, D., Hinze, A., Heese, R., Vanderschantz, N., Timpany, C., Cunningham, S.J.: An Exploration of ebook Selection Behavior in Academic Library Collections. In: *TPDL '12*, pp. 13-24. Springer, Heidelberg (2012)
13. Makri, S., Blandford, A., Woods, M., Sharples, S., Maxwell, D.: "Making my own luck": Serendipity strategies and how to support them in digital information environments. *JASIST* 2179-2194 (2014)
14. Kleiner, E., Rädle, R., Reiterer, H.: Blended shelf: reality-based presentation and exploration of library collections. In: *CHI 13*, pp. 577-582. ACM, (2013)
15. Bates, M.J.: The design of browsing and berrypicking techniques for the online search interface. *Online Inform Rev* 13, 407-424 (1993)

16. Bates, M.J.: What is browsing--really? A model drawing from behavioural science research. *Inform Res* 12, (2007)
17. Oksanen, S., Vakkari, P.: Emphasis on examining results in fiction searches contributes to finding good novels. In: *JCDL* 12, pp. 199-202. ACM, 2232855 (2012)
18. Kules, B., Capra, R., Banta, M., Sierra, T.: What do exploratory searchers look at in a faceted search interface? In: *JCDL* 09, pp. 313-322. ACM, (2009)
19. Borgman, C.L.: Why are online catalogs *still* hard to use? *JASIS* 47, 493-503 (1996)
20. McKay, D., Conyers, B.: Where the streets have no name: how library users get lost in the stacks. In: *CHINZ* 10, pp. 77-80. ACM, (2010)
21. McKay, D., Shukla, P., Hunt, R., Cunningham, S.J.: Enhanced browsing in digital libraries: three new approaches to browsing in Greenstone. *IJDL* 4, 283-297 (2004)
22. Thudt, A., Hinrichs, U., Carpendale, S.: The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In: *CHI* 12, pp. 1461-1470. ACM, (2012)
23. Pearce, J., Chang, S.: Exploration without Keywords: The Bookfish Case. In: *OzCHI* 2014, pp. 76-79. ACM, (2014)
24. McKay, D., Buchanan, G.: One of these things is not like the others: how users search different information resources. In: *TPDL* 11, pp. 260-271. Springer (2011)
25. Marshall, C.C.: *Reading and Writing the Electronic Book*. Morgan & Claypool, Chapel Hill, (2010)
26. Saarinen, K., Vakkari, P.: A sign of a good book: readers' methods of accessing fiction in the public library. *J Doc.* 69, 736-754 (2013)
27. Rowlands, I., Nicholas, D.: Understanding Information Behaviour: How Do Students and Faculty Find Books? *J Acad Libr* 34, 3-15 (2008)
28. Tenopir, C., King, D.W., Edwards, S., Wu, L.: Electronic journals and changes in scholarly article seeking and reading patterns. *Aslib Proceedings* 61, 5-32 (2009)
29. Reutzell, D.R., Gali, K.: The Art of Children's Book Selection: A Labyrinth Unexplored. *Reading Psychology* 19, 3-50 (1998)
30. Moore, P.: Information Problem Solving: A Wider View of Library Skills. *Contemporary Educational Psychology* 20, 1-31 (1995)
31. Svenonius, E.: *The Intellectual Foundation of Information Organization*. MIT Press, Boston, (2000)
32. Hancock-Beaulieu, M.: Evaluating the impact of an online library catalogue on subject searching at the catalogue and at the shelves. *J Doc* 46, 318-338 (1993)
33. Losee, R.M.: The relative shelf location of circulated books: A study of classification, users, and browsing. *Libr Resour Tech Serv* 37, 197-209 (1993)
34. O'Neill, E.T., Gammon, J.A.: Consortial Book Circulation Patterns: The OCLC-OhioLINK Study. *C&RL* 75, 791-807 (2014)
35. Sanderson, M., Dumais, S.: Examining Repetition in User Search Behavior. In: Amati, G., Carpineto, C., Romano, G. (eds.) *Advances in Information Retrieval*, vol. 4425, pp. 597-604. Springer (2007)
36. Jansen, B.J., Spink, A.: How are we searching the World Wide Web? A comparison of nine search engine transaction logs. *Inform Process Manag* 42, 248-263 (2006)
37. Baeza-Yates, R., Ribeiro-Neto, B.: *Modern information retrieval*. ACM Press NY (1999)
38. McKay, D., Buchanan, G.: Boxing clever: how searchers use and adapt to a one-box library search. In: *OZCHI* 13, pp. 497-506. ACM, 2541031 (2013)
39. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. *ToIS* 22, 5-53 (2004)
40. Fried Foster, N., Gibbons, S.: *Studying Students: The Undergraduate Research Project at the University of Rochester*. Association of College and Research Libraries, Rochester, NY (2007)
41. Saarti, J.: Feeding with the spoon, or the effects of shelf classification of fiction on the loaning of fiction. *Inform Serv Use* 17, 159 (1997)