Interaction differences in web search and browse logs

Paul Thomas
CSIRO
Canberra, Australia
paul.thomas@csiro.au

Alex O'Neill
Australian National University
Canberra, Australia
alexoneill89@gmail.com

Cecile Paris
CSIRO
Sydney, Australia
cecile.paris@csiro.au

Abstract *We use logfiles from two web servers (public and internal), two corresponding search engines, and two user populations (public and staff) to examine differences in behaviour across users and sites.*

We observe similar overall characteristics to other browsing and searching logs, but differences in behaviour between staff and the public and between external and internal sites. Staff familiarity with organisational language and structure does not translate to more effective search or navigation, although staff do expend considerable effort looking for information and often look in the wrong place. This would not be apparent from logs covering only search or only browsing behaviour.

Keywords Log analysis; user behaviour; information retrieval

1 Introduction

Transaction logs, recording users' interactions with web servers or web search engines, provide a cheap and unobtrusive source of data on how people use the web. A good deal of analysis has been carried out on log data (see e.g. Silvestri [6] for a recent survey). However, although this data has been collected on a large scale—up to 15M queries and 7M sessions in the case of the MSN logs of 2006—it has generally been recorded from a single user population, using either a single website or a single search engine. Since logs have been collected at different times, with different users, and recording different information, direct comparisons are not generally possible.

Given log data from two distinct user populations, with the website and other factors held constant, it would be possible to compare the behaviours of these two populations; and, if there were salient differences, possibly to suggests ways a "one-size-fits-all" tool could be modified to best suit each population's needs. Similarly, if we had comparable log data from the same population interacting with two distinct data sources, it would be possible to compare users' behaviour across different websites or search engines.

Proceedings of the 15th Australasian Document Computing Symposium, Melbourne, Australia, 10 December 2010.

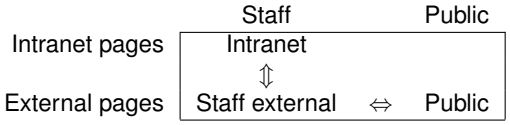

Figure 1: We have three classes of interaction data: from the intranet, from internal users of the external site ("staff external"), and from public users of the external site ("public"). Comparisons are possible across page sets (intranet vs staff external) and across user populations (staff external vs public).

We have explored these ideas with three data sets collected from the internal and external web servers and search engines of a large Australian government agency. Our three sets include two user populations (agency staff and the public) and two sites (one internal to the agency, one open to all). This collection allows two types of comparison (Figure 1):

- comparisons between behaviours of the same users, but on different sets of pages, by looking at staff use of the intranet and of the external webpages (we call this "intranet" and "staff external" data respectively);
- comparisons between behaviours of different users, on the same set of pages, by looking at staff and public uses of the public website (we call this "staff external" and "public" data).

Our experiments considered differences in both these dimensions (Sections 3–5); behaviours which predict when a user may switch from one source to another (Section 6); and morals for the design of web retrieval systems. We begin by describing the logfiles we used.

2 Data

Data was collected from two of the agency's web servers and an associated analytics service. One server was responsible for presenting intranet pages: these pages were only available inside the agency, or over a secure link, and we assume all users of this server were agency staff. The second server was responsible for the agency's public website, and presented pages for users both inside and outside the agency. The servers used different web publishing tools, but the same search engine.

The data covers a one month period (23 October to 22 November 2009), and were collated from search engine logfiles and metrics provided by the "Clicky" tracking service.[1] They represent 498,955 individual sessions and 1,595,180 page views, from 212,255 distinct IP addresses. This is approximately the same size as the logs from the Excite search engine [8] but under 3% the size of the more recent AOL and Microsoft logs.[2] For each session, on either server, we obtained the session start time and duration, the user's IP address (and in many cases their geographical region and network block), and platform details such as operating system and browser. Each session includes a number of "actions"—searches or page views—for which we have time, URL, and referrer data. For searches we also recorded query terms.

Users of the external pages were partitioned into "staff" and "public" on the basis of their IP address: those corresponding to agency domain names, or an agency network block, were assumed to be staff and the remainder were assumed to be the public. This produced 290,493 public sessions (58%); 26,591 staff external sessions (5%); and 181,871 intranet sessions (36%). The majority of staff sessions (87%) were on the intranet.

Note that the intranet search engine maintained logs which were not directly comparable with logs from other sources. We integrated this data by assigning searches to the closest intranet session, with matching user IP address, within 30 minutes of the search. (The 30 minute window is an arbitrary choice. Smaller values however presented similar results.)

Note also that not all of the agency's web presence, internal or external, is captured in our logs—in particular, we only have data from the main, agency-wide, servers. It is not clear how many sessions, on group intranets or groups' public web pages, are not accounted for here. We are confident, however, that we have a large fraction of interactions and those which represent a variety of behaviours and needs.

3 General characteristics

The broad patterns in our data are consistent with those seen in other log files and with commonsense expectations, with a moderate amount of searching and sessions of a few minutes' duration.[3] 61% of public sessions seem to be from Australia, according to hostname and/or geolocation information.

[1] `http://www.getclicky.com/`

[2] The "spring 2006 MSN search data asset" was not publicly released, but see e.g. `http://research.microsoft.com/en-us/events/searchsummit2007/`. Summary statistics are presented in e.g. Zhang and Moffat [9].

[3] We do not have data from any pages viewed prior to the agency's, and in particular we do not have complete data from external search engines which led people to the agency's pages. Search activity in the public and staff external sets will therefore be underestimated, but it is hard to know how much by.

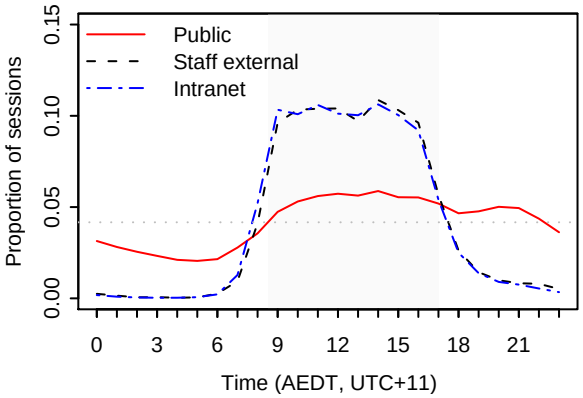

Figure 2: Sessions with each source, by start time.

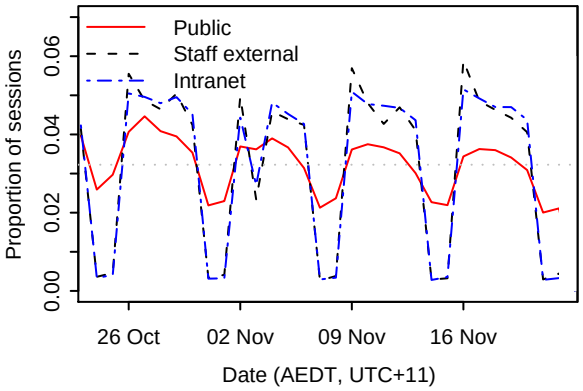

Figure 3: Sessions with each source, by day. Tuesday 3 November 2009 was a public holiday in the ACT and parts of Victoria.

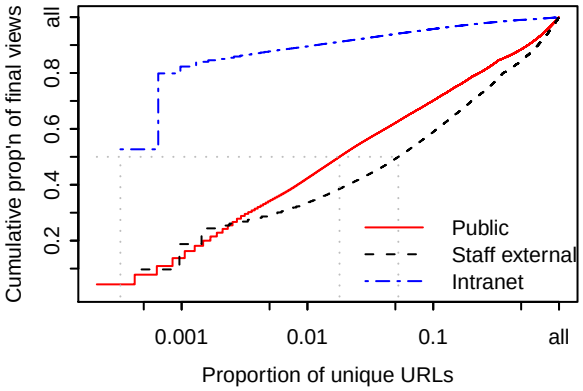

Figure 4: Distribution of final views. Note logarithmic scale on abscissa. From 0.03% to 5.3% of pages, depending on data set, account for 50% of final views.

Figures 2 and 3 summarise the volume of data in our three sets according to session start times. (All times are reported as if in south-east Australia.) There are clear trends: staff use both the internal and external web pages during work hours and on work days. Public usage is more uniform over time, which is consistent with more use from home and more use from outside Australia.

Session entry pages are largely uninteresting, since the distribution is dominated by the public and intranet home pages. Exit pages—the final URL viewed in a

session—are much more likely to have provided the information people require, and patterns here are more interesting.

We considered the distribition of final views across all sessions, except where those sessions were followed by a switch from the internal server to the external, or vice versa (see Section 6). Figure 4 summarises the distribution: reading left to right is an index of URLs, sorted by popularity so that the most popular ending URL is at the far left, and reading bottom to top is the cumulative proportion of final views.

Staff are clearly looking at a small number of intranet pages just before ending a session: the single most popular intranet page constitutes over 50% of final views, as do the two most popular pages from the staff external set. Public interests are more varied and it takes the 85 most popular pages from the public set to make 50% of final views, although this is still only 1.8% of all pages ever viewed. Examining the most popular final URLs shows public interest in the agency's projects and in job advertisements, while staff views are predominantly the web pages of organisational units.

Figure 4 confirms a common skewed pattern: there are some resources used extremely often, and a very long tail. The public data set, in particular, appears to follow a Zipfian distribution but (as is common) this is in fact a poor fit (fitted using the method of Clauset et al. [1], Komolgorov-Smirnov $D = 0.05$, $n = 4717$, $p > 0.05$).

Table 1 summarises a number of other statistics across the three data sets. Superscript [a] indicates significant differences between the public and staff external data (paired t test, $\alpha = 0.05$); that is, differences on the vertical axis of Figure 1 or between the first two columns in the table. Superscript [b] indicates differences between staff external and intranet data (horizontal axis of the figure, columns two and three of the table).

We also examined the proportion of queries (in a subsample) which used proper names; acronyms; and question words. There were only small numbers in each set, and no significant differences.

4 Navigation and effort

The combination of search engine and web server logs allows us to examine browsing and searching behaviour in tandem.

We expect to see some differences between the public and staff external sets, for several reasons: staff will be more familiar with the organisation and language of the agency; staff may be more familiar with the agency websites; and there are likely to be different sorts of tasks for each user group.

In particular, we expect staff to be more successful at navigation, which leads to our first hypothesis:

▷ NAVIGATION: Since staff have better knowledge of the agency's and the websites' structure, they will be more fluent at navigation. Staff sessions will be shorter than those of the public.

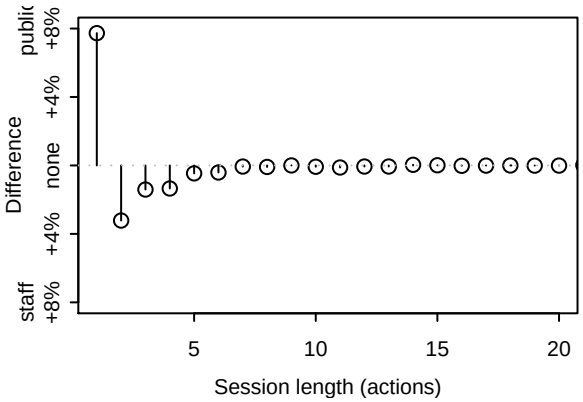

Figure 5: Differences in distributions of session length. Session length reads left to right; differences between public and staff external data sets read up (more common for public) or down (more common for staff).

We also expect that many public users may be browsing for fun, or to scratch an itch, whereas staff are more likely using the website for particular information relevant to their jobs. This leads to two further hypotheses, possibly contrary to the first:

▷ EFFORT #1: Staff are more likely than the public to be highly motivated—especially since they are mostly using the website during the working day, and may be looking for information essential to their work. We expect that staff will not give up if they can't find something immediately.

▷ EFFORT #2: We expect a similar effect as for effort #1, but more pronounced, between the staff external and intranet sets.

The session length statistics in Table 1 seem quite clearly to contradict our navigation hypothesis—rather than being shorter, sessions are instead longer for staff whether measured by actions (a 13% increase) or by time (an 87% increase). This could be explained by the navigation and effort #1 hypotheses together: a larger proportion of staff external sessions could be short, say length 1, while a longer tail (corresponding to extra effort) could drag out the mean.

Figure 5 tests this idea by illustrating the difference in distributions of session length. Session length (in actions) runs left to right; points above the axis are where public sessions of this length are more likely than staff external sessions of this length, and points below are the reverse.

The large spike above the axis for the shortest sessions, and the corresponding mass below the axis at sessions of moderate length, demonstrates that public users are more likely to have a very short session; and staff are more likely to have a longer one. (Less than 10% of sessions in either set are longer than six actions, so while there are similarities in the tail of each distribution these represent only a minority of cases.) Our

	Data set		
	Public	Staff external	Intranet
Scale			
Sessions in set	290,493	26,591	181,871
Actions in set	918,699	95,195	587,344
IP addresses in set	201,452	4,999	10,815
URLs in set	56,790	11,321	23,564
Sessions			
Session length (actions)[ab]	3.2 ± 0.01	3.6 ± 0.05	3.2 ± 0.01
Session length (min:sec)[ab]	$3{:}36 \pm 0{:}01$	$6{:}04 \pm 0{:}06$	$9{:}00 \pm 0{:}03$
Sessions with repeated page views[ab]	24.5%	33.5%	39.1%
Searches			
Sessions with searches[ab]	5.9%	7.9%	7.4%
Queries per session, where > 0[b]	2.5 ± 0.02	2.5 ± 0.08	1.7 ± 0.01
First query length (terms)[a]	2.2 ± 0.01	1.9 ± 0.02	1.9 ± 0.01
Overall query length (terms)[a]	2.3 ± 0.01	1.9 ± 0.01	1.9 ± 0.01
Query groups			
Number of queries per group[ab]	2.1 ± 0.02	1.8 ± 0.06	1.5 ± 0.01
Groups with only one query[ab]	56.9%	68.8%	76.2%
Reformulations			
Groups which added terms[a]	24.8%	14.4%	13.5%
Groups which removed terms[a]	20.1%	12.5%	11.8%
Groups with hypernyms	0.005%	0.002%	0.002%
Groups with hyponyms	0.004%	0.002%	0.002%
Groups with synonyms[a]	0.013%	0.005%	0.006%

Table 1: Summary statistics of the three data sets. [a] indicates significant differences between public and staff external; [b] between staff external and intranet. Shown is mean $\pm$ one standard error.

explanation is not supported; we conclude that in fact staff members' sessions are not made shorter by their better organisational knowledge.

The two effort hypotheses, however, are confirmed. Sessions are longer for the staff external set than the public set; they are longer again, in both actions and time, for the intranet set, although changes in layout and language may account for some of this difference.

There is also a significant increase, in both cases, in the number of sessions in which one or more URLs is visited more than once ("repeated page views" in Table 1). This "going in circles" suggests a serious attempt to find some information by navigation—staff are more likely to backtrack and try another path than are the public, and intranet users are more likely again.

We conclude that staff are not in fact more competent at navigation, on either the intranet or the public website, but do expend considerable effort trying to find information important to them. It is likely that the public site is poorly designed for staff use, with information scattered over many pages and forcing extra navigation for staff (but not the public); conventional site analysis or interviews with staff may help confirm or rebut this.

5 Queries, query groups, and reformulations

The combined logs also show differences in search engine use, and in the strategies different groups use to find information.

As before, it seems likely that staff will show some trace of better understanding the agency and the external website. This leads to two likely hypotheses:

▷ QUERIES PER SESSION #1: Staff will be more likely to use useful search terms; and staff will be more familiar with navigational aids and other landmarks in the websites. Therefore, when they do search they will issue fewer queries than the public do.

▷ QUERIES PER SESSION #2: Since the intranet is written for an internal audience, the effects of this familiarity with language and structure should be more pronounced for the intranet set.

The mean number of queries per session is the same for both user groups (Table 1), but the means hide a difference in distribution; staff are more likely to issue a single query in a session, and the public are more likely to issue two to five queries (Figure 6, read the same way as Figure 5). Staff also use slightly fewer terms per query. Overall, there is not strong evidence for the first hypothesis but it is at least plausible.

There is stronger evidence for the second hypothesis. Staff tend to issue fewer queries on the intranet than on the external site (31% fewer overall), and there is again a difference in distributions with intranet sessions around 12% more likely to include a single query and staff external sessions correspondingly more likely to include 2–8.

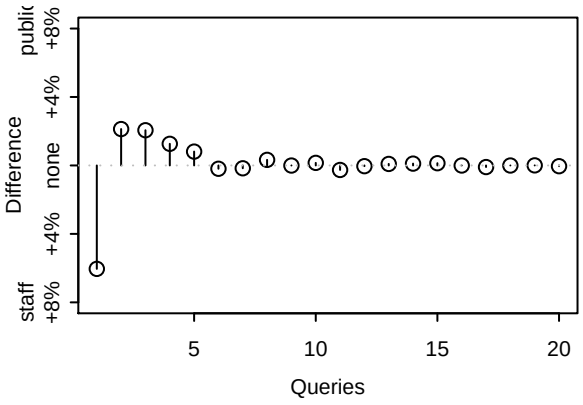

Figure 6: Differences in distributions of query counts. Query counts read left to right; differences between public and staff external data sets read up (more common for public) or down (more common for staff).

Since a session may include any number of searches, possibly for different topics or tasks, we also consider queries as part of "query groups". A "group" attempts to collect queries from the same session which correspond to a single topic. Given a sequence of queries, we created groups in three phases (Figure 7):

1. First, each query was considered in its own group.
2. Second, as in earlier work [9], term similarity was based on the Jaccard similarity between trigrams. Terms were considered "similar" where Jaccard similarity was greater than 0.25.

   Any queries which included similar terms were assumed to be on the same topic, and were grouped together.
3. Finally, any query that was bracketed by queries in some group was assigned to that group. This makes the simplifying assumption that users do not temporarily switch topics in the middle of a sequence.

If we take these query groups into account, there is a stronger effect for both public/staff external and staff external/intranet comparisons (see "query groups" in Table 1). Query groups tend to be shorter, in each case—there is a 15–20% decrease—and significantly more groups only include a single query.

Query reformulations, where a user has repeated a query with modifications, are likely to represent a failed search and are therefore of particular interest. We considered reformulations in each group by counting instances of five relationships between consecutive queries: adding terms; removing terms; including a term which is a hypernym of a term in an earlier query; including a term which is a hyponym of a term in an earlier query; and including a term which is a synonym of a term in an earlier query. Hyper-, hypo-, and synonyms were drawn from WordNet [5] via NLTK [4]. Note that more than one relationship might hold between any two queries: for example,

	Queries			
	butterflies	insects	butterfly flight	bricklaying
	↓			
1: Each query in its own group	A butterflies	B insects	C butterfly flight	D bricklaying
	↓			
2: "Butterflies" has 60% character trigram overlap with "butterfly"	A butterflies	B insects	A butterfly flight	D bricklaying
	↓			
3: "Insects" is between two queries in the same group	A butterflies	A insects	A butterfly flight	D bricklaying

Figure 7: Query grouping example. Four queries in a session are grouped to represent two information needs.

the sequence "vehicle allowance" and "vehicle rules" includes a deletion ("allowance") as well as an addition ("rules").

This data allowed us to test a further hypothesis:

▷ REFORMULATION: Staff will reformulate queries less often than the public do.

This assumes that staff will be more familiar with the agency's language, so will be more likely to choose useful terms early on. (The same effect would also be apparent if staff are so used to the agency's vocabulary that they find it hard to think of rephrasings.)

The reformulation hypotheses is borne out to some extent ("reformulations" in Table 1), and we see fewer reformulations in the staff external set than in the public set although the base rate for hyper-, hypo-, and synonym use is low. (Note that Wordnet's coverage of the terms used in queries was lower than expected, so the rate of hyper-, hypo-, and synonym use is probably higher than reported here.)

There is evidence then that staff are searching less than the public, on the external network, and are reformulating their queries less often. This would normally suggest that their extra knowledge of the agency's structure and language, and their experience with the website, is leading to greater search success. However, the two effort hypotheses were also borne out while the navigation hypotheses was not: and staff tend to have longer sessions, and fewer sessions of length one. Longer sessions with fewer searches do not in fact suggest search success; rather, they suggest that staff are using a different tactic. Recall too that staff were more likely than the public to go in circles, and view a page more than once—a likely sign of navigation, as well as search, failure.

On the basis of search engine log files alone, we may conclude that the smaller number of searches, and smaller number of repeated searches, mean that staff are generally successful; that the search engine is doing a good job. With the integrated logs, the opposite becomes clear: there are definite signs of increased effort (sessions are longer) and of staff abandoning the search engine in favour of navigation (fewer queries are issued). That navigation also seems unsuccessful a lot of the time (repeated views are high). Web server logs clearly add something important to the search engine logs and having only one or the other would tell only half the story.

6 Switching sources

Staff have access to both the internal and external web pages. It is possible therefore for staff to look for information from one source which is in fact available from the other: for example, it is possible to spend some time looking for street addresses in the intranet although they are generally published to the public web. Our last set of questions investigate occasions when staff switch sources—that is, when they stop using one data source in favour of another.

Session identifiers in the logs are not comparable across sources, so switches were determined by IP address: a switch was counted when the same IP address appears in both the staff external and intranet sets (in either order) with a gap of less than thirty minutes. Smaller gaps made minimal difference.

Table 2 summarises statistics for sessions which did switch, and sessions which did not switch, but could have (i.e. staff external or intranet sessions), in the same way as Table 1.

▷ LITTLE SWITCHING: As before, we expect that agency staff—who use the intranet and the public site on a daily basis—know where to look for different classes of information, at least at the site level. We expect to see little switching.

There were 10,117 instances of switching from 208,462 staff external or intranet sessions. This is a fairly small proportion (only 4.9%), but it does represent a significant effect. Sessions which end with a switch tend to involve more effort (mean 120% more actions) and be much longer (mean 112% more time). If the 17 minutes 31 seconds spent before switching sources to look somewhere else represents

	Sessions ending with	
	Switch	No switch
Scale		
Sessions in set	10,117	198,345
Actions in set	69,061	614,607
IP addresses in set	3,541	11,223
Sessions		
Session length (actions)[c]	6.8 ± 0.11	3.1 ± 0.01
Session length (min:sec)[c]	$17{:}31 \pm 0{:}17$	$8{:}15 \pm 0{:}02$
Sessions with repeated page views[c]	57.1%	39.8%
Searches		
Sessions with searches[c]	20.1%	6.8%
Queries per session, where > 0[c]	2.4 ± 0.06	1.8 ± 0.02
First query length (terms)[c]	2.0 ± 0.02	1.8 ± 0.01
Overall query length (terms)[c]	2.0 ± 0.02	1.9 ± 0.01
Query groups		
Number of queries per group[c]	1.8 ± 0.04	1.5 ± 0.01
Groups with only one query[c]	69.8%	76.0%
Reformulations		
Groups which added terms[c]	16.9%	13.1%
Groups which removed terms[c]	14.9%	11.4%
Groups with hypernyms	0.003%	0.001%
Groups with hyponyms	0.003%	0.001%
Groups with synonyms	0.006%	0.006%

Table 2: Summary statistics of sessions which do or do not end with a switch of sources. [c] indicates significant differences. Shown is mean $\pm$ one standard error.

wasted time, then this is over 140 hours per day of lost time—equivalent to eighteen full-time staff. This may be an underestimate since the logs used here only capture session switches, not sessions abandoned entirely in favour of external search engines or other methods entirely.

The sessions which end with a switch do show evidence of users struggling with both search and navigation. As well as being dramatically longer, switching sessions are 43% more likely to include repeated page views. They are three times more likely to involve querying, and when the search engines are used there are more queries, queries are slightly longer (although this is a small effect), and query groups are 26% more likely to include a second or subsequent query.

As well as trying harder, users who are about to switch sources show some signs of adopting different strategies. Query reformulation is also somewhat more common in switching sessions than otherwise, although still less common than in the public set.

If it were possible to predict which sessions will end in a switch, based on browsing and searching data like that in Table 2, it should be possible to offer extra help in these sessions. For example, if there is reason to believe a staff member is looking in the wrong place then we could expand the search engine's scope to include other sources. Depending on what sort of intervention is proposed, such a predictor will most likely need high precision but could sacrifice some recall.

Unfortunately whole-of-session features such as user IP address, session length, number of repeated page views, number of searches, and number of query groups are not promising in this regard. Neither J48 trees built with WEKA [2] nor SVMs built with SVMlight [3] have been able to predict a final switch with greater than 65% accuracy.

Rather than aggregate over the actions in a session, a more promising technique is to examine each separately and consider *sequences* of actions as cues that users might be struggling. As a first attempt, we sampled 400 sessions from the logs—200 which switched and 200 which were candidates, but did not—and coded each action as a search, a page view, or a repeated page view (i.e. a view of a page already seen in the same session). Each session was represented by all its subsequences of length two or more. For example, a session with the sequence search, view, view, re-view—"svvr"—would be represented by the subsequences { svvr, svv, vvr, sv, vv, vr }.

There were over 9000 subsequences in our sample which only occurred in switching sessions, and 243 which occurred in two or more switching sessions and no non-switching sessions. These included obvious subsequences such as "sssss"—five searches in a row, with no other page views—and long runs of repeated

views. 33 of our 400 sampled sessions included one or more of these 243 subsequences, so using their presence as a cue it would be possible to detect about 16% of switching sessions before the point of the switch, and with very high accuracy. By including other subsequences it would be possible to increase recall at the expense of precision.

These subsequences of actions could be used in combination with other techniques: for example, it may be possible to improve predictions by using the presence of certain subsequences as a feature in an SVM classifier. Analysis of maximal repeating patterns [7] may also offer insights, and we intend investigating these further.

7 Conclusions and further work

Combined logs from two servers, two search engines, and two user populations have allowed us to contrast users' behaviours across information sources and across populations. While the overall patterns are similar to other logs, there are differences in searching and browsing behaviour in both comparisons.

Although staff presumably are more familiar with the agency's organisation, language, and website they spend more effort finding information online. This may be because they are more motivated; but certainly their familiarity does not seem to translate to greater success at navigation or searching. This is likely due to poor site design, but further analysis or interviews are needed to better understand staff behaviour and to point to particular areas needing improvement.

With the combined logs we can also observe some sessions which end with a switch in source, from the intranet to the public site or vice versa, and these represent a considerable amount of staff time. Early indications are that we may be able to spot these sessions on the fly, with high precision, to offer extra help. Further work will pursue this idea.

Many of these observations would not be possible given only a search engine log, or only a web server log. The combination is useful and we hope to expand future logfile work to include both sources where possible.

Acknowledgements We are grateful for help and input from Tom Rowlands, and from the agency's information technology group. We also thank the anonymous reviewers for their useful comments.

References

[1] Aaron Clauset, Cosma R Shalizi and M E J Newman. Power-law distributions in empirical data. *SIAM Review*, Volume 51, Number 4, pages 661–703, 2009. arXiv:0706.1062v2.

[2] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann and Ian H. Witten. The WEKA data mining software: An update. *SIGKDD Explorations*, Volume 11, Number 1, pages 10–18, 2009.

[3] Thorsten Joachims. Making large-scale SVM learning practical. In Bernhard Schölkopf, Christopher J.C. Burges and Alexander J. Smola (editors), *Advances in kernel methods: Support vector learning*, pages 169–184. MIT Press, Cambridge, Massachusetts, 1999.

[4] Edward Loper and Steven Bird. NLTK: The natural language toolkit. In *Proc. ACL Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*, pages 63–70, Philadelphia, July 2002. Association for Computational Linguistics. arXiv:cs/0205028v1.

[5] George A Miller. Wordnet: a lexical database for English. *Comm. ACM*, Volume 38, Number 11, pages 39–41, 1995.

[6] Fabrizio Silverstri. *Mining query logs: Turning search usage data into knowledge*, Volume 4 of *Foundations and Trends in Information Retrieval*. Now Publishers, Delft, 2010.

[7] Antonio C Siochi and Roger W Ehrich. Computer analysis of user interfaces based on repetition in transcripts of user sessions. *Trans. Info. Systems*, Volume 9, Number 4, pages 309–335, 1991.

[8] Amanda Spink, Bernard J Jansen, Dietmar Wolfram and Tefco Saracevic. From e-sex to e-commerce: Web search changes. *Computer*, Volume 35, Number 3, pages 107–109, 2002.

[9] Yuye Zhang and Alistair Moffat. Some observations on user search behaviour. In *Proc. ADCS*, pages 1–8, 2006.