A Study of Mobile Information Exploration with Multi-touch Interactions

Shuguang Han[1], I-Han Hsiao[2], and Denis Parra[3]

[1] School of Information Sciences, University of Pittsburgh, Pittsburgh, PA, United Sates
shh69@pitt.edu
2 The EdLab, Teachers College Columbia University, New York, NY, United Sates
ih2240@columbia.edu
[3] School of Computer Science, PUC Chile, Chile
dparra@ing.puc.cl

Abstract. Compared to desktop interfaces, touch-enabled mobile devices allow richer user interaction with actions such as drag, pinch-in, pinch-out, and swipe. While these actions have been already used to improve the ranking of search results or lists of recommendations, in this paper we focus on understanding how these actions are used in exploration tasks performed over lists of items not sorted by relevance, such as news or social media posts. We conducted a user study on an exploratory task of academic information, and through behavioral analysis we uncovered patterns of actions that reveal user intention to navigate new information, to relocate interesting items already explored, and to analyze details of specific items. With further analysis we found that dragging direction, speed and position all implied users' judgment on their interests and they offer important signals to eventually learn user preferences.

Keywords: Multi-touch interactions, implicit relevance feedback, mobile information seeking behaviors.

1 Introduction

The massive user adoption and the rapid improvements of mobile technologies have attracted researchers from academia and industry to investigate how people use mobile devices compared to the traditional desktop or laptop computers. For instance, several studies comparing user behavior during web search tasks [1,3,4] have found differences between desktop and mobile devices in terms of query length, user click patterns, and search time distribution. Other researchers have leveraged location information - which is more easily captured in mobile devices - to investigate how people search for near restaurants or tourism landscapes [5,6,7]. Touch-enabled devices provide additional features that can be used to augment the user search experience, such as drag, pinch-in (zoom in), pinch-out (zoom out), and swipe.

Though search is an important activity, there are other important activities such as information exploration. In search tasks, users usually have specific information needs and the search results are ranked by their relevance. However, in other user

W.G. Kennedy, N. Agarwal, and S.J. Yang (Eds.): SBP 2014, LNCS 8393, pp. 269–276, 2014.

activities such as reading of news or social media posts users mostly lack a clear goal and the items are commonly displayed in lists sorted by chronological order. Therefore, we argue that in this scenario the multi-touch interactions can be exploited to improve the user experience but they cannot be interpreted in the same way as in traditional search tasks, where the top results are assumed to be the most relevant ones. In this article, we investigate how users behave in multi-touch enabled mobile devices while exploring information, how differently they behave compared to search tasks, and more importantly, whether we can find a relation between multi-touch interactions and users' interests.

2 Related Works

To display the same amount of information as in desktop computers, users in mobile devices need to perform additional actions (e.g., page scrolls, zooms) because of the limited screen size. Implicit user feedbacks can be utilized to improve information filtering. Page dwell time [9] and mouse cursor movements [8, 10] have been already identified as important implicit feedback on desktop devices. Besides them, multi-touch feedback has been also found to be able to generate significant improvements [2] on ranking search results. For example, the zooming-in behavior may suggest that users are interested in the targeted content block, whereas the fast swiping behavior may indicate that they are only "skimming" the non-relevant content. Implicit feedback is also used to generate recommendations. For example, Hu et al. [11] introduced the concept of "confidence" in a matrix factorization model for weighting the implicit feedback, and Pan et al. [12] further considered weighting both positive feedback and negative feedback, treating it as One Class Collaborative Filtering problem.

To the best of our knowledge, implicit feedback is mostly used in information retrieval or recommendation tasks where the items are ranked by relevance to the users' queries or interests. Moreover, the traditional reading patterns over listed results learned from information retrieval systems, e.g., the F-Shape [14], are no longer useful when analyzing exploration tasks over items not sorted by relevance but rather chronologically such as a list of news or social media posts. We argue that finding the most relevant items from a list of search or recommendation results is significantly different from reading a list of items for exploration. Therefore, in this paper, we are interested in studying further implicit feedback in mobile devices: understand user behavior during information exploration, interaction patterns while performing this activity, and how those patterns can be used to infer users' interests.

3 Research Design

3.1 Task and Data Collection

The task was designed to resemble the reading of news or social media posts. However, we chose a task of academic content by displaying a list of scientific publications because: i) news and social media posts are usually augmented with images, videos or external links, which introduce variables that will influence user's preference on our analysis; and ii) we believe that a mobile application that supports exploration in the academic domain can generate future impact on scientific communication while it is still rare recent years. In this study, we focused on the computer science domain. Since the academic conferences have evolved to be the major channel for the dissemination of research achievements in this domain, the publications for this user study were chosen from three conferences in Information Retrieval (SIGIR), World Wide Web (WWW) and Human Computer Interaction (CHI). To minimize the bias that some users may have already read those publications, we considered the publications of those conferences during 2013, which are available online only less than three months before the time of our study. We have 739 publications (392 from CHI, 137 from SIGIR and 210 from WWW). The title, abstract, author and affiliation information was displayed in the user interface.

3.2 The Experimental System: ConfReader

We developed a mobile web application called ConfReader[1]. Fig. 1 presents a screenshot of the article list page, which provides a list of articles with title and briefed abstracts up to 180 characters. By quickly swiping left or right, users can bookmark an article, as shown in Fig. 3. By tapping an article, users will see the detailed information with full title, abstract, author and affiliation information, shown in Fig. 2. All the pages have a navigation bar fixed as a footer in order to facilitation navigation to the article list page or check the number of articles being bookmarked.

ConfReader is able to log the following multi-touch interactions on a mobile device: hold, tap, double-tap, drag up, drag down, drag left, drag right, pinch in, pinch out. Each interaction is defined as starting from touching the screen until the user releases the fingers. We implemented the functionalities based on an open source Javascript library for multi-touch gestures, the hammer.js[2].

Fig. 1. The article list page

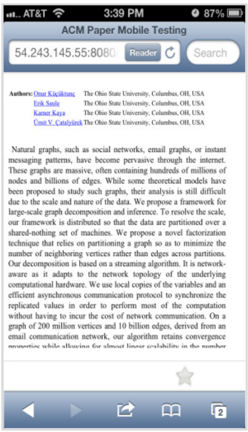

Fig. 2. The detail page

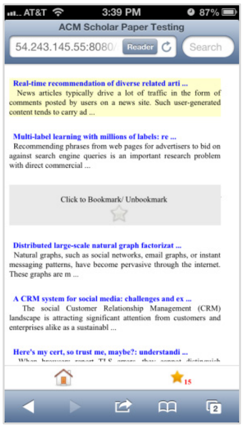

Fig. 3. Swipe to Bookmark

[1] http://54.243.145.55:8080/acmmw/index.jsp?uid=-1&section=1
[2] http://eightmedia.github.io/hammer.js

3.3 Participants and Procedure

Our target participants were PhD students majoring in computer or information science, considering that our data collection is from three conferences in computer science related domains. In total, 15 PhD students from 5 universities in China (2) and United States (3) were recruited. 9 are female and 6 are male.

The experiment began with an introduction to the study. Then, participants were given a 5-minute tutorial about the system, to continue with a training task around 10 minutes and to be familiar with the system on one of the 3 conferences. After the training, each participant was asked to finish three tasks in up to 10 minutes for each, each task corresponding to one conference. The order of those three tasks (conferences) was rotated based on the Graeco-Latin Square design. For each user task, we randomly assigned a list of 40 papers from one conference. They were asked to explore all those articles to interact with our experimental system, and to choose the top ten most relevant articles based on their personal interest. All of the four tasks were completed using our system with no interventions by the experimenters.

4 Result Analysis and Discussions

4.1 Descriptive Statistics on Multi-touch Interactions

In the study, we collected 3,519 multi-touch interactions (3,041 from article list pages and 478 from detailed pages) from 15 users. Within the given 30 minutes for exploration, the actual-time-on-task to interact with the article list pages spent 25.62±(18.83) minutes, which is double of the time spent on detailed pages, with only 13.44±(10.88) minutes. This is understandable due to the list pages had already provided a short snippet of the abstracts, users will tap into detailed pages only if they want to confirm. The action percentage of each page is presented in Table 1.

As for actions, users seldom used the pinch-in, pinch-out or double-tap. The result is consistent with the study of web search behaviors on mobile devices [2]. This is either due to the interface displayed proper font size and layouts resulted in smooth exploring process or because user has no preference on performing multi-touch behaviors. Since users are required to drag left/ right to bookmark articles in the article list pages, it was not surprising that dragging left/ right take high proportion in the article list page but not in the detailed pages. The actions are dominated by the tapping (39.12% in detail page) and dragging up/down (44.03% and 14.8% in the list page). In the detailed pages, users perform more tapping. In the detailed pages, users have more tapping. We think it may because most of the detailed pages are actually able to accommodate the full information in one page; however, users still conduct further interactions for reconfirming in case they missed some hidden content. Without scrollable content, the actions were counted as tapping. The dragging down behaviors in the article list page take around 15% of all actions which demonstrated a dynamic exploring process instead of linear browsing that goes from top to bottom. However, the deeper implications of dragging up and dragging down are still unclear, which are the focuses discussed in the following sections.

Table 1. Percentage of users' multi-touch interactions on different pages

	Article list page	Detail page
Drag up	44.03%	40.17%
Drag down	14.80%	6.28%
Drag left	6.61%	5.02%
Drag right	22.73%	7.95%
Tap	11.05%	39.12%
Pinch in/ Pinch out/ Hold/ Double-tap	0%	0 %

4.2 Dragging Down vs. Dragging Up

In our system, each logged interaction consists of a set of metadata, such as the distance (distance has moved on screen), centerX/Y (the X/Y axis positions when a touch begins), velocityX/Y (the touch speed on X/Y axis). To obtain a deeper understanding of the dragging up/ down actions on both the article list pages and the detail pages, we analyzed those metadata. Table 2 shows the statistics of the metadata information and the comparison between dragging up and down. We found that when dragging down on the article list page, users tend to put fingers at a higher position (in Y axis) comparing to the dragging up. It makes sense because a higher position allows more distances for dragging. Indeed, the dragging up in article list page has a longer average distance, although no significance is found. Since the article list pages contain more content than in detailed pages, the centerY is always bigger. A fair comparison may use the relative position in Y axis; however, our system cannot log such information because the limited API functionality from the open-source library. The reported mean of centerX is near the third quarter position of the screen.

In terms of velocity, the dragging down are significantly faster than dragging up, which suggests that users are more likely to exploring information when dragging up. While in dragging down, they are only skimming or relocating information to what they recalled from memory. We think that the consistent slow dragging up gesture may demonstrate users' attentive behavior in identifying articles for their interests, which we will test in the next section. We didn't compare between two pages because there are only few explorations on detailed pages. For several users, there are even no dragging up/ down in detailed pages, from which we cannot make fair comparisons.

Table 2. The comparsion of gesture metadata information (with standard deviation) for dragging up/ down on different pages. * means significance at 0.1 level, ** means 0.05 level and *** means 0.01 level. Significance tests are based on Generalized Linear Models (GLMs).

	Article list page		Detailed page	
	Drag down	Drag up	Drag down	Drag up
centerX	571 ± (240)	554 ± (266)	583 ± (227)	537 ± (284)
centerY	**4676** ± **(2861)***	4357 ± (2690)	917 ± (425)	982 ± (332)
velocityX	**0.382** ± **(1.721)****	0.168 ± (0.322)	**0.319** ± **(0.496)*****	0.170 ± (0.496)
velocityY	**1.294** ± **(4.660)****	0.760 ± (1.623)	**0.899** ± **(1.244)***	0.610 ± (1.354)
distance	65.78 ± (57.46)	60.51 ± (39.09)	51.49 ± (43.32)	53.70 ± (34.32)

4.3 Inferring Users' Interests from Dragging Up/Down

Dragging up/down is the dominating behavior in both article list and detail pages. In mobile devices, dragging up/down served the same functions as mouse scrolling up/down in desktop, except that users need to tap on a certain item to move up/down. When users drag on a specific item, we assume that there are N surrounding articles on users' reading zone. N is a small value considering the small screen size. As an initial step, we set N = 1 and assume that users will drag quickly if they are not interested in those articles; otherwise, they will slow down and read carefully. We would like to test whether there is a significant correlation between the dragging speed (on Y axis) and user interests.

Each user was required to bookmark articles on three different conferences, when studying users' interactions on one conference, the bookmarked articles on the other two are used to model users' real interest. User interest **I** is represented as a vector that aggregates bookmarked documents (**B**) using vector space model over all of the words in vocabulary, i.e. $\mathbf{I} = \langle \mathrm{w}_1, \mathrm{w}_2, \dots, \mathrm{w}_{|\mathrm{V}|} \rangle$. $|\mathrm{V}|$ is vocabulary size. Weight for each word in each document w_{ij} is represented as Formula (1), which denotes the TF-IDF for word w_{i} in d_{j}. To test the assumption, we calculate the cosine similarity between $\mathbf{D}_{\mathbf{r}}$ / $\mathbf{D}_{\mathbf{r+1}}$ / $\mathbf{D}_{\mathbf{r-1}}$ (i.e. the article users dragged on, the article above the dragged article and the article below) and user interest **I.** The similarity is used to measure the interestedness of the dragging area.

$$w_i = \sum_{d_j \in \mathbf{B}} \frac{w_{ij}}{|\mathbf{B}|} \tag{1}$$

Our system logged the dragging speed for each drag up/down action. Thus, we can compute the Pearson correlation between dragging speed (we use a log transformation because the speed distribution is highly skewed) and the interestedness of the dragged area. The correlation coefficients are -0.01 (no significance) for dragging down while **-0.1 (*p*<0.01**) for dragging up, which indeed provides evidence to our assumption that the dragging speed is reduced when users read content similar to their interest. In summary, this result can indicate that users mostly drag up to explore new items, whereas they drag down to relocate information relevant to their interest.

4.4 Predicting Users' Bookmarks Using Rich Interactions

With the goal of producing personalized recommendations, we aimed at evaluating whether identifying user interest based on touch interactions is comparable with identifying it based on bookmarks. In our study, each user needs to read a list of articles of 3 conferences and bookmark 10 for each. Using the logged actions at each conference, we performed an evaluation simulating an online recommender. To explain the protocol, let's assume a user **u1** is exploring WWW conference and at the moment of bookmarking each article, we produce recommendations based on the logged actions for CHI conference. The evaluation protocol is described as following:

- Given the user **u1** and the conference WWW, we follow the actions in sequential order and keep track of each article tapped in a vector **I_t**, and each article dragged up/down (and its immediate upper and lower articles in the list) in a vector **I_d**.

- When **u1** bookmarks a paper we add that paper to the vector **I_b**, and then we generate three sets of recommendations from the papers in the CHI conference based on their similarity to the user interests. The first recommended list **R_b** is generated using **I_b** (Pred_book), the second recommended list **R_t** using the tapped articles **I_t** (Pred_tap), and the third list **R_d** using the dragged articles **I_d** (Pred_drag). Since drag speed is negatively correlated with user interests, Pred_drag filters out those drags with high speed (speedY>2.0).
- The papers that the user **u1** bookmarked in CHI will be used as ground truth for evaluating the algorithms. We calculate the precision at 10 (P@10) and also Average Precision (AP) for each recommended list **R_b**, **R_t** and **R_d**.

We repeat the three previous steps for each article bookmarked at each conference, completing 10 evaluations, which are each of the ticks in the x-axis of Figure 4 and Figure 5. We also include a random prediction (Pred_Rand) by randomly guessing the top 10 recommendations. Since only 40 articles were displayed for each conference, the random baseline already achieved high performance. Comparing to the random baseline, the Pred_tap only improves the performance after the 3rd or 4th bookmarks, which is due to the data sparseness. Users may result in *judging-a-book-by-its-cover* even if they were interested in the article. Pred_drag has a smaller sparseness problem and works better than using the Pred_tap, which confirms the value of multi-touch interactions. The best performance is achieved by Pred_book, which considers users' explicit feedbacks. Though explicit feedback is indeed more valuable than implicit feedback, the former is usually scarcer in real systems. Given our results, we argue that multi-touch actions in mobile devices offer a promising alternative for non-invasive preference elicitation. We also observe that precision curves for both Pred_tap and Pred_drag have a convex shape, which indicates that implicit feedbacks might introduce noise when aggregated over time. Leveraging each type of implicit feedback and finding a better tradeoff among them are our next research focus.

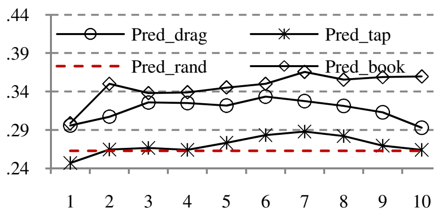

Fig. 4. The P@10 evaluation on the i-th (x-axis) bookmarked paper

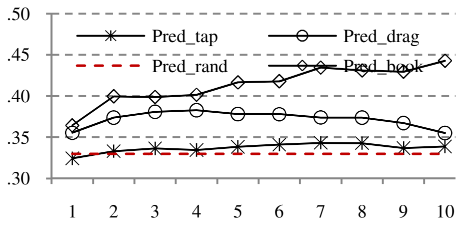

Fig. 5. The AP evaluation on the i-th (x-axis) bookmarked paper

5 Conclusion

In this paper, we studied the mobile information exploration behaviors through a controlled user study, in which users were asked to explore and bookmark a list of conference publications. Users tended to read and explore information when dragging up but relocating information when dragging down. We also found that dragging up/down actions' speed and position can be used to infer user interest over articles and also as additional feedback to predict users' future actions.

However, there are still several open issues to incorporate multi-touch interactions in predicting user interests. First of all, in this paper, we assumed that users were reading surrounding articles (the one being dragged, one upper and one lower). Follow up studies are needed for justifying this assumption, analyzing user's reading zones through eye-tracking. Secondly, we considered actions isolated. The drag down and tap actions could be combined with drag up actions for further improving the recommendation performance. For example, a slowly drag up with follow-up tapping on articles may further confirm user interest. Thirdly, with more exploration on the displayed articles, user's interests may change dynamically, particularly in the exploratory task where user lacks a clear goal. A possible solution is to consider weight decay on articles being explored based on temporal information, a recent explored article may receive more weight than the old one.

References

1. Song, Y., Ma, H., Wang, H., Wang, K.: Exploring and exploiting user search behavior on mobile and tablet devices to improve search relevance. In: WWW, pp. 1201–1212 (2013)
2. Guo, Q., Jin, H., Lagun, D., Yuan, S., Agichtein, E.: Mining touch interaction data on mobile devices to predict web search result relevance. In: SIGIR 2013, pp. 153–162 (2013)
3. Yi, J., Maghoul, F., Pedersen, J.: Deciphering mobile search patterns: a study of Yahoo! mobile search queries. In: WWW 2008, pp. 257–266. ACM (2008)
4. Kamvar, M., Baluja, S.: Deciphering trends in mobile search. Computer 40(8), 58–62 (2007)
5. Ricci, F.: Mobile recommender systems. Information Technology & Tourism 12(3), 205–231 (2010)
6. Brunato, M., Battiti, R.: Pilgrim: A location broker and mobility-aware recommendation system. In: PerCom 2011, pp. 5265–5272. IEEE Computer Society (2003)
7. Dean-Hall, A., Clarke, C., Kamps, J., Thomas, P., Voorhees, E.: Overview of the TREC 2012 Contextual Suggestion Track. In: Proceedings of the 21st NIST Text Retrieval Conference (2012)
8. Huang, J., White, R., Buscher, G., Wang, K.: Improving searcher models using mouse cursor activity. In: SIGIR 2012, pp. 195–204. ACM (2012)
9. Morita, M., Shinoda, Y.: Information filtering based on user behavior analysis and best match text retrieval. In: SIGIR 1994, pp. 272–281 (1994)
10. Kong, W., Aktolga, E., Allan, J.: Improving passage ranking with user behavior information. In: CIKM 2013, pp. 1999–2008. ACM, New York (2013)
11. Hu, Y., Koren, Y., Volinsky, C.: Collaborative filtering for implicit feedback datasets. In: Eighth IEEE International Conference on Data Mining, ICDM 2008. IEEE (2008)
12. Pan, R., Zhou, Y., Cao, B., Liu, N.N., Lukose, R., Scholz, M., Yang, Q.: One-class collaborative filtering. In: Data Mining, ICDM 2008, pp. 502–511 (2008)
13. Manning, C.D., Raghavan, P., Schütze, H.: An Introduction to Information Retrieval, p. 181. Cambridge University Press (2009)
14. Lorigo, L., Haridasan, M., Brynjarsdóttir, H., Xia, L., Joachims, T., Gay, G., Pan, B.: Eye tracking and online search: Lessons learned and challenges ahead. Journal of the American Society for Information Science and Technology 59(7), 1041–1052 (2008)