Towards a Better Understanding of Mobile Shopping Assistants – A Large Scale Usage Analysis of a Mobile Bargain Finder Application

Stephan Karpischek
Chair of Information Management, ETH Zürich
skarpischek@ethz.ch

Gilad Geron
ETH Zürich
ggeron@student.ethz.ch

Florian Michahelles
Chair of Information Management, ETH Zürich
fmichahelles@ethz.ch

ABSTRACT
Mobile shopping assistants have been subject to research in the field of ubiquitous and pervasive computing for many years. Now the wide adoption of mobile shopping applications for smartphones allows evaluation on a large scale. To study how consumers actually use these applications, we analyze server logs of a mobile bargain finder application for the iPhone used by 33,000 users over a period of six months. In this paper we discuss our approach, the methods we have used, and some challenges and limitations we have encountered. First results indicate that contrary to the focus of most research in the field the application is used rather from home than at the point of sale or on the go.

Author Keywords
Mobile, shopping, commerce, retail, bargain, consumer, app, analysis, usage mining

ACM Classification Keywords
J.7 Computers in other Systems: Consumer Products; H.5.m Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Documentation, Economics, Management.

INTRODUCTION
Mobile shopping assistants have been a relevant topic for pervasive and ubiquitous computing for many years. Brody and Gottsman presented the Pocket BargainFinder, a handheld device to scan barcodes of books and find the best available price online in 1999 [1]. Today, the wide distribution of mobile shopping applications on smartphones - e.g., ShopSavvy reports 16 million downloads of their price comparison app [9] - allows to study the usage of mobile applications on a larger scale and in their real contexts of use. This ”in the wild” approach has been proposed as a model for research in the field of ubiquitous and pervasive computing, e.g. by Korn [5], McMillan [7], and Michahelles [8], and promises insights into real-world usage of mobile applications compared to findings from research under laboratory conditions.

In order to increase the understanding of mobile shopping applications in the wild we have partnered with Comparis, a company providing an iPhone application for finding bargains in supermarkets. The application aggregates data about special offers and bargains from all local retailers and supermarkets on a daily basis and makes them available on the iPhone. Bargains are browsable by product category and retailer. In this paper we evaluate the application's usage by analyzing the log data of the server backend, studying service requests from more than 33,000 individual users over a period of more than six months. Our goals are to find out how the bargain finder application is used and what interesting research questions can be answered using the available data.

METHODOLOGY
Based on previous work on web usage analysis [10, 11] a usage mining process is applied for an explanatory analysis of the available log data to find out how the bargain finder application is used. The first step is to understand the application and the log data. Then the data are prepared and enriched for statistical analysis. The following subsections describe the application and its functionality, the available log data, and the process of data preparation and enrichment.

Description of the application
The iPhone application Bargain Finder was released on the iTunes App Store in March 2009 by Comparis, the national leading provider for price comparison of services and products in Switzerland. After starting the application the user sees a list of current top bargains for local retailers ordered by the percentage of potential saving. Users can select bargains by product categories and subcategories, e.g. sodas and lemonades in beverages. Users can also select bargains by retailer, combine retailer and product category selection, and locate nearby stores. A watchlist allows users to remember single bargains. Figure 1 shows screenshots. The application uses a webservice to request information about the current bargains from a server backend.

Large'11, September 18, 2011, Beijing, China.

Figure 1. Screenshots of the iPhone application with mapped request types (from left to right): top bargains, retailer selection, bargains for one retailer, category selection, subcategory selection in the category food.

Log Data

Every request from the iPhone application is logged on the server backend. The application does not cache any information about bargains but requests it from the server every time the user changes the selection of bargains. This means that the server logs show the actual usage of the application and interests of the users very precisely. Only adding bargains to the watchlist does not trigger a request to the server as the information is stored locally on the iPhone.

Data preparation and enrichment

The provider of the iPhone application made available more than one year of log data for analysis. An initial analysis shows that several weeks of log data are missing due to storage errors. For the further analysis the largest consistent part of the logs without gaps of missing data and representing full calendar weeks is chosen: from 22 November 2009 until 29 May 2010, in total more than six months of log data with over 1.2 million requests.

User and session identification

A persistent cookie ID in the logs allows tracking of individual instances over time. For the analysis it is assumed that one instance of the application is mainly used by the same user so this cookie ID can be considered to be a unique user identifier. The application does not collect any personal information from users so all analyzed data is anonymous.

Both the cookie ID and a session ID are generated on the server side upon the first request and then sent to the iPhone app. Thus, every first request of an instance and every first request of a session is missing the ID in the server logs. The timestamp and IP address of the requests are used to reconstruct this missing ID. Requests for which cookie or session ID can not be reconstructed are not included in the analysis.

Request types

Ten types of request originating from the iPhone application can be differentiated in the server logs and are shown in figure 1. To reduce the dimensions for analysis these types are grouped into four groups of high level request types: category specific requests (C), retailer specific requests (R), combined category and retailer requests (E), and generic requests (G) without a category or retailer specified.

Connection type

To get more meaningful location information originating IP addresses from the server logs are resolved using various data sources: The freely available GeoIP database is used to get a country of origin and the Internet service provider routing the IP address [3]. The Domain Name System (DNS) is used to do reverse-lookups in order to get hostnames. The WHOIS protocol is used to get more information about a range of IP addresses [2]. In combination this allows to determine whether a request origins from a WiFi Internet connection or a cellular network (3G or UMTS): Nearly all requests (98.5%) originate from the home country of the bargain finder service provider, Switzerland, where there are three mobile network operators (MNOs).

For two of the MNOs the range of IP addresses which are used for routing cellular Internet connections can be easily defined by looking at either the hostname, which is determined by a reverse DNS lookup, or by the output of a WHOIS query, which describes a whole class C network as being used for only mobile Internet connections. For the third MNO the range of cellular IP addresses could be identified using heuristics learned from the other two MNOs, i.e. much more different cookie IDs per time come from cellular IP addresses than from WiFi IP addresses.

Data analysis

The enriched data are analyzed using Knime, an open-source data analytics software [4]. Requests are grouped into sessions and users. A session is characterized by the type of Internet connection, day of the week, time of day, the number of requests, the total time from first request to last request, the ratios and sequences of request types, product categories and retailers. Individual user profiles are characterized by the total number of sessions and requests, the average number of requests per session, the average time of a session, the total time the application has been used, the average time spans between sessions, and the ratio of request types. The qualitative variance [6] of request types, retailers and product categories shows how focused a user's interest is.

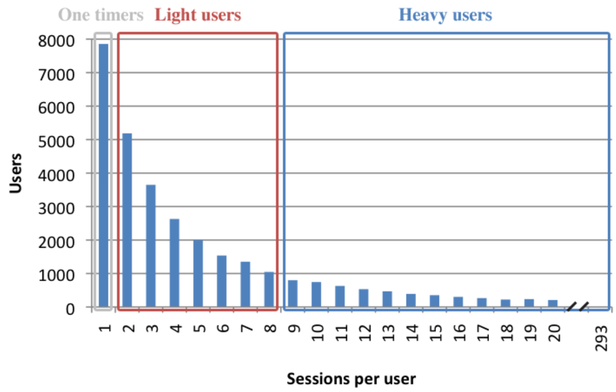

Figure 2. Distribution of sessions per user and user segmentation

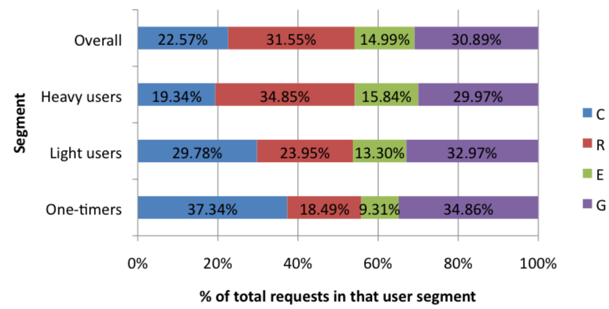

Figure 3. Distribution of high level request types

RESULTS

This section presents selected results of the statistical analysis of the log data.

How often is the application used?

1,218,775 requests can be mapped to a unique user and session ID, resulting in 258,702 unique sessions and 33,117 unique users over a time period of 27 weeks. On average, there were 6,508 requests (standard deviation 1,220) and 1,382 sessions (standard deviation 187) per day.

Figure 2 shows the distribution of sessions per user. Users have 7.82 sessions on average with a standard deviation of 13.78, the maximum number of sessions for a single user is 293, the median is 3. Using the number of sessions as a measure of usage intensity and to compare different user groups in the further analysis users are grouped into three segments: 23.73% are one-timers with only one session, 52.55% are light users with two to eight sessions, and 23.73% are heavy users with more than eight sessions. Heavy users account for 71.25% of all sessions.

When and from where is the application used?

Figure 4 shows the overall number of requests for the different connection types per hour of day. In total 61.42% of all requests come from WiFi Internet connections, while 38.58% come from cellular Internet connections. This ratio does not vary significantly over the user segments (Light users: 60.70%, Heavy users: 61.58% WiFi). In the evenings the ratio of WiFi connections increases.

Which functions are used most?

Figure 3 shows how the different high level request types are distributed for the user segments and overall. Heavy users show a higher percentage of retailer specific requests (R) while one-timers and light users have more category specific requests (C).

DISCUSSION

One of the main challenges we encountered was the balance between which questions are interesting for research and which questions can be answered from the available data. Another challenge when analyzing a single application is whether results can be generalized. In our case many results from the statistical analysis are meaningful for this particular application only, and thus are only interesting for the maker of the app and maybe local retailers, and probably not so much for the research community. Examples are the most popular product categories (wine, cosmetics, meat); popular days of the week (Tuesday); the development of usage numbers over time (stable); which features were rarely used (locate retail stores); and individual user profiles and usage patterns.

Some identified metrics are interesting to compare with other applications and industry reports. The percentage of one-time users can be compared to numbers reported from an U.S. app analytics company for more than thousand apps used in 2010. The bargain finder app with 23% one-timers performs slightly better than the reported average of 26% [12]. Other interesting metrics to compare could be the number of monthly active users, the average number of sessions per user, or the average time span between sessions.

One surprising and counterintuitive result is that most requests (61.42%) originate from a WiFi Internet connection. As retailers do not offer WiFi connections at the point of sale for customers, and in general WiFi connections are not open and publicly available in Switzerland, this also indicates that the application is not so much used while shopping at the point of sale and on the go, where mostly cellular connections could be used, but much more from a home or office environment. This is contrary to the focus of most related research in the field of ubiquitous and pervasive shopping where applications are mostly designed and evaluated for use at the point of sale or on the go. Our interpretation is that consumers use the iPhone bargain finder application mostly at home in order to plan shopping or to inform themselves about available bargains.

Grouping users by usage intensity shows that about 23% of the users are responsible for more than 70% of all requests and sessions. While one-timers and light users are more interested in specific product categories, heavy users are more interested in specific retailers. Figure 3 shows how the usage focus changes with increasing usage intensity. As the usage intensity increases, the ratio of retailer requests (R) and

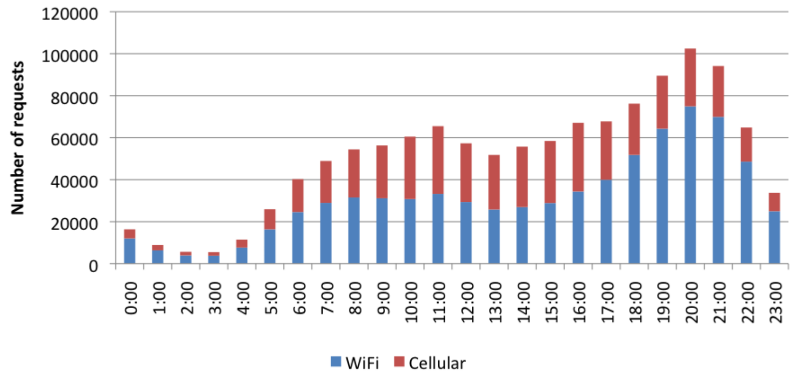

Figure 4. Number of requests split by connection type per hour of day

mixed requests (E) increases and the proportion of generic (G) and category (C) requests decreases. Our interpretation is that heavy users are more focused and make less generic requests.

CONCLUSIONS

We have presented our approach to analyzing a mobile bargain finder application. While the analysis can provide feedback for application developers, we found the results so far to be of rather limited value for research. A user survey is planned to overcome these limitations. The goal is to find out more about determinants for using the app and the effects on shopping behavior. Linking individual survey responses to the respective user profile from usage mining promises interesting results.

When designing an application it makes sense to consider evaluation options: The bargain finder application supports notifications which are pulled from the server and which can link to websites. This makes it easy to attract participants to an online survey. In order not to break the consistence of sessions in the server logs, the application could be improved by using separate requests to assign IDs to clients.

First results from this work in progress indicate that users tend to use the application from home rather than while shopping or on the go. This is contrary to related work where mobile shopping assistants are designed and evaluated for use at the point of sale or on the go. We also found that with increasing usage intensity users tend to use the application in a more focused way.

REFERENCES

1. A. B. Brody and E. J. Gottsman. Pocket bargain finder: A handheld device for augmented commerce. In *Proceedings of the 1st international symposium on Handheld and Ubiquitous Computing*, HUC '99, pages 44–51, London, UK, 1999. Springer-Verlag.

2. L. Daigle. WHOIS Protocol Specification. RFC 3912 (Draft Standard), September 2004.

3. MaxMind - GeoIP — IP Address Location Technology. `http://www.maxmind.com/app/ip-location`. Accessed 25 June 2011.

4. KNIME — Konstanz Information Miner. `http://www.knime.org/`. Accessed 25 June 2011.

5. M. Korn. Understanding use situated in real-world mobile contexts. In *Proceedings of the Workshop on Research in the large*, Ubicomp 2010, 2010.

6. J. Magidson. Qualitative variance, entropy, and correlation ratios for nominal dependent variables. *Social Science Research*, 10(2):177 – 194, 1981.

7. D. McMillan, A. Morrison, O. Brown, M. Hall, and M. Chalmers. Further into the wild: Running worldwide trials of mobile systems. In P. Floréen, A. Krüger, and M. Spasojevic, editors, *Pervasive Computing*, volume 6030 of *Lecture Notes in Computer Science*, pages 210–227. Springer Berlin / Heidelberg, 2010. 10.1007/978-3-642-12654-3_13.

8. F. Michahelles. Getting closer to reality by evaluating released apps? In *Proceedings of the Workshop on Research in the large*, Ubicomp 2010, 2010.

9. A. Muse. The shopsavvy numbers. `http://shopsavvy.mobi/2011/04/06/the-shopsavvy-numbers/`, April 2011. Accessed 25 June 2011.

10. C. Silverstein, H. Marais, M. Henzinger, and M. Moricz. Analysis of a very large web search engine query log. *SIGIR Forum*, 33:6–12, September 1999.

11. J. Srivastava, R. Cooley, M. Deshpande, and P.-N. Tan. Web usage mining: discovery and applications of usage patterns from web data. *SIGKDD Explor. Newsl.*, 1:12–23, January 2000.

12. B. Suthoff. First impressions matter! 26% of apps downloaded in 2010 were used just once. `http://www.localytics.com/blog/2011/first-impressions-matter-26-percent-of-apps-downloaded-used-just-once/`, January 2011. Accessed 25 June 2011.