

AUTO-ID LABS

Usage Analysis of a Mobile Bargain Finder Application

Stephan Karpischek, Darshan Santani, Florian Michahelles

Auto-ID Labs White Paper WP-BIZAPP-064

September 2012



Stephan Karpischek

Senior Researcher
University St.Gallen and ETH
Zürich



Darshan Santani

Student
ETH Zürich




Dr. Florian Michahelles

Associate Director of the Auto-ID
Labs Zürich/St. Gallen
Manager of the labs of Prof.
Fleisch at D-MTEC
ETH Zürich

Contact:

Dr. Florian Michahelles
ETH Zürich
Phone: +41 44 632 7477
Fax: +41 44 632 1740
Email: fmichahelles@ethz.ch
www.im.ethz.ch

This paper is already published in the 13th International Conference on Electronic Commerce and Web Technologies (EC-Web 2012), Vienna, Austria, September, 2012



Abstract


Mobile shopping applications for smartphones are popular among consumers. While mobile commerce research has focused on experimental prototypes and evaluation in small groups, only little is known about the real-world usage of these applications. Established tools and methods for analysis are missing. In this paper, we present the usage analysis of a mobile bargain finder application based on server logs of more than 98,000 users over a period of 6 months. We show that plots of the cumulative distribution function (CDF) are well suited to analyze the distribution of relevant parameters and present simple heuristics to identify active users. We can show that Pareto's law applies to the distribution of user requests. We also propose metrics to measure usage focus and find that active users tend to become more focused with increased usage. Finally, we combine the results from the log analysis with demographics from an online user survey.

Key words: Mobile, Shopping, Consumer, Application, Log, Usage, Analysis

1. Introduction

Mobile shopping assistants help consumers to make better buying decisions and have been a topic of interest in mobile commerce research for many years. Today shopping applications for mobile phones are used on a daily basis by millions of consumers worldwide. While research so far has mostly focused on experimental prototypes and evaluation in small scale user studies or laboratory environments, only little is known about the actual usage of these applications. Application developers and other stakeholders in the area of mobile commerce are interested in finding out how mobile applications are used in real-world scenarios. The goal of this paper is to better understand the real-world usage of a popular mobile commerce application for finding bargains in nearby supermarkets and to present methods and tools which are well suited for analyzing mobile commerce applications.

This paper follows a previous analysis of an earlier version of the mobile bargain finder application [8]. In this paper we analyze an updated version of the application and combine the log analysis with demographics from an online survey conducted among the application's users. We extend our analysis methodology and show that cumulative distribution function (CDF) plots are well suited to analyze the distribution of relevant parameters such as user requests or sessions. We also describe simple heuristics to identify active users. For the bargain finder application we can show that the distribution of user requests follows the Pareto principle. Then we propose metrics to measure a user's focus when using the application and show that the focus of active users increases with usage duration.



2. Related Work

Mobile applications which help consumers to make better buying decisions have been a topic of interest in research for many years: Early prototypes were customized hardware devices [1], later software prototypes were implemented on Personal Digital Assistants (PDAs) [5, 9, 13, 10], and beginning in 2003, software prototypes were implemented on mobile phones [15].

In 2009, Deng and Cox presented LiveCompare, a prototype using mobile camera phones for grocery bargain hunting through participatory sensing [4]. The system focused on crowdsourcing price information for grocery products and discussed the problem of data scarcity and data integrity. Obtaining data from retailers directly would help overcome these limitations. Our work studies the usage of a mobile application for finding bargains which gets its information on bargains directly from retailers.

Research on mobile shopping assistants so far has focused on prototypes which have not been widely deployed and evaluated on a large scale. While other mobile applications have been researched in the large [11, 12, 2], the evaluation in the reviewed work on mobile commerce applications used relatively small user groups and took place in controlled lab environments. Findings about the real-world usage of mobile applications by consumers are relevant for mobile commerce research and practitioners like retail companies and application developers but are missing so far.

Today the distribution channels for mobile applications to consumers on smartphones offer an interesting opportunity to deploy mobile shopping applications to large user groups and also to analyze real-world usage over a longer period of time. Our work focuses on this in-the-wild approach which has not yet been applied to mobile shopping applications for consumers. The contributions of this paper are a set of methods to analyze mobile commerce applications and findings about the real-world usage of a mobile bargain finder application.

3. Background Data

3.1. Mobile Application Overview

Comparis Shopper is a mobile commerce application for the iPhone which was first released in the iTunes App Store under the name Bargain Finder in 2009. Users can inform themselves about bargains and special offers in supermarkets nearby. They can choose to display bargains from specific product categories or retailers. Figure 1 shows some screenshots of the iPhone application.



Fig. 1: Comparis Shopper Screenshots (a) Top Bargains (b) Retailer (c) Product categories (d) Shop Finder

3.2. Log Data

In this section we give an overview of the anonymized query logs comprising of over 5 million requests which were collected from over 98,000 users over 191 days, from February 2011 until August 2011. As a user interacts with the application, hypertext transfer protocol (HTTP) requests are sent by the iPhone to the backend server application fetching relevant information for the user. Since there is no local caching on the iPhone, the server logs provide a detailed and precise representation of the spatiotemporal usage patterns of individual users.

Each request in the query logs contains a unique user identifier, originating IP address, the URI containing the application’s request for a specific function call, session identifier and the respective timestamp when the request was generated. With every fresh install of the iPhone application, the backend server generates a unique random number as user identifier which is stored locally on the device and appended every time the device sends a request. Thus every installed instance of the application can be identified in the query logs, and requests coming from the same instance and user can be grouped together. Similarly a session ID is generated to differentiate user sessions, where a session is defined as a series of requests from the same device within a certain period of time, commonly known as session delta [16]. If the device does not send a request within session delta of the last request, the session expires and next subsequent request is assigned a new session ID.

Request Types One of the interesting data embedded in the query logs is the URI string, which represents the mobile application’s request for a specific function of the server backend, and in most cases a user’s request for a specific information. In total, the query logs data contain 17 different request types, each representing a particular functionality. We can differentiate two basic groups of requests:

- **System Requests:** System requests are operational in nature and requested implicitly by the application in the background without any explicit action of the user, e.g., on startup the application checks if a newer version exists.
- **User Requests:** User requests are explicitly triggered by the user to satisfy a specific information need, e.g., the search for bargains from a specific retailer or for a specific product category.

Over 41% of the total requests are system requests, while the rest of 59% are user requests.

3.3. Survey Data

In addition to analyzing the query logs an online survey was conducted among the users of the application in July and August 2011. When starting the application during this time, users were presented with an invitation to fill out the survey online using the web browser of their iPhone. When accepting the invitation the iPhone's web browser launched the online survey with the unique user identifier from the application as parameter. Survey participants remained anonymous over the whole time. The only personal information collected was the email address which participants could enter after completing the survey if they chose to participate in a raffle. Email addresses were only used for the raffle and not in the analysis to preserve the participants privacy. In this paper we report results from combining the log analysis with demographics from 1,009 completed online surveys. More details about the survey and more results can be found in [6].

4. Analysis and Results

4.1. Preliminary

In this section we specify definitions related to users request types which will be used throughout the remainder of the paper. For a user i we denote the distribution of its requests across M request types ($q_1; q_2; \dots; q_M$) as a vector ($r_{i1}; r_{i2}; \dots; r_{iM}$). And we define the total queries requested by user i across all the request types during the entire period of analysis T as $R_i = \sum_{k=1}^M r_{ik}$.

With the given premise we now examine the fundamental properties and characteristics of usage behavior in the subsequent sections. The main goals of our analysis are (a) to understand how the application is used, (b) to explore whether application usage changes as users spend more time using the application, and (c) to develop a set of methods and tools which are well-suited to analyze mobile commerce applications.

4.2. Requests

First we analyze user activity by examining the total number of requests across all request types. Intuitively one would expect that different users have different access patterns – some are active users who frequently use the application on a regular basis, while others have sporadic application usage over their lifetime. The given dataset matches our intuition and we observe a similar behavior. Figure 2a shows a Complementary Cumulative Distribution Function (CCDF) of total requests R_i by users. The plot shows a heavy tail distribution where a handful of extremely active users send the most number of requests. In fact, 50% of all the users send less than 10 requests in total. Note that in the given figure (and subsequent analysis) we chose to use cumulative distribution function (CDF) or complementary cumulative distribution function (CCDF) as they show the relevant statistics and the underlying (probability) distribution more clearly and succinctly, in addition to providing a better visual aid to compare multiple distributions in the same plot.

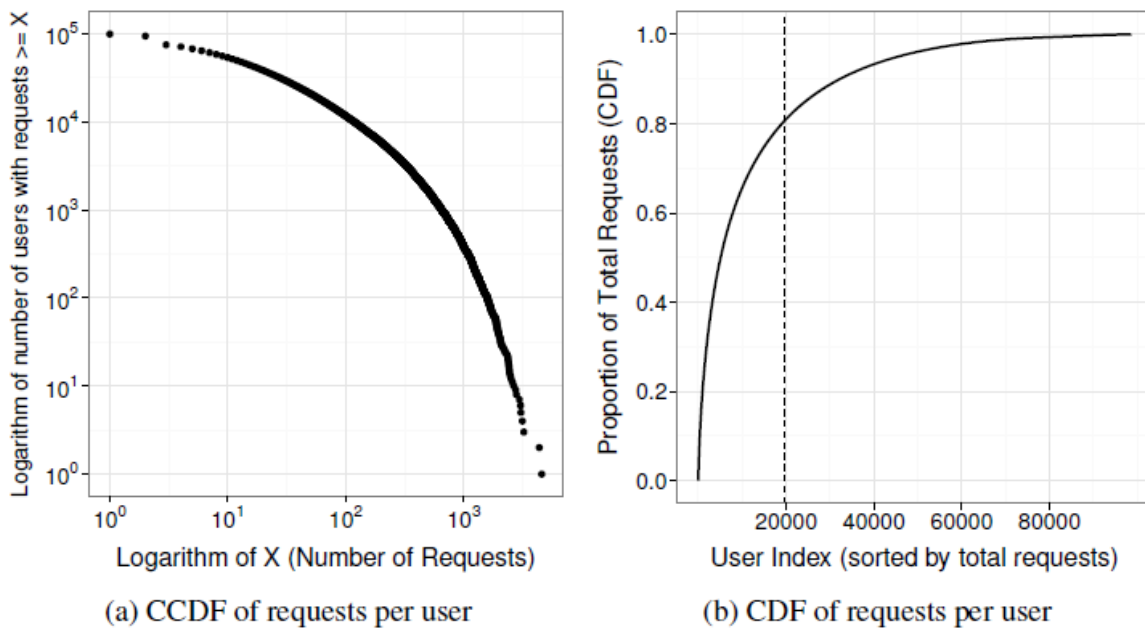


Fig. 2: Requests

To further investigate the skewness in the distribution of queries, we examine the applicability of the Pareto Principle, commonly known as the 80-20 rule. We analyzed the total queries for all the users, as shown in the cumulative distribution function (CDF) plot in Figure 2b, with the user rank on the horizontal axis and fraction of total requests from the corresponding set of users on the vertical. Users are ordered by the number of total requests. The distribution closely follows the Pareto Law, i.e., more than 75% of all requests are generated by top 20% of users.

4.3. Sessions

Now we turn our attention to investigating the user sessions, in particularly focusing on the number of sessions performed by each user and its respective session length. As explained in Section 3, a session is a period of constant activity where the user sends a series of queries within a certain time interval, usually termed as session delta [16]. If the user does not send a query within session delta of the last query, the current session expires and with the next subsequent request a new session is started. For our analysis, we have set the session delta to be 90 minutes.

In our dataset, we have observed a total of 561,707 sessions for all users. Similar to the heavy tailed distribution of request queries, total sessions per user also exhibit a heavy tail distribution, as shown in Figure 3a. Over 75% of users have less than 5 sessions in total, with the most active user having 359 sessions. It indicates that majority of the users are dormant users who just install the application for curiosity and exploration and then never use it again or even uninstall it. It is also interesting to note that in addition to total sessions, individual session lengths follow a similar distribution as shown in Figure 3b, where we have plotted individual session lengths on the horizontal axis and their corresponding CDF on the vertical axis. Over 75% of total sessions lasted for less than a minute, clearly suggesting low application usage and interactivity for majority of users.




Fig. 3: Sessions

4.4. User Classification

In addition to looking into individual session lengths, we also analyze the total session duration, i.e. the sum of the lengths of all sessions for one user. We observe that over 72% of users have used the application for less than 10 minutes in total. These cumulative patterns reveal significant variability in application usage and suggest a possible categorization of users as per their activity into “Inactive” and “Active” users. Inactive users are further classified as either “Dormant” or “One-Timers” users with heuristics shown in Table 1. Now we describe these user segments in more detail:

- **Dormant Users:** Dormant users typically install the application, but never interacts with it beyond opening the application’s landing page once in a while (which explains them having multiple sessions but with all sessions having length of 0)
- **One-Time Users:** One-time users are the ones who communicates with the application once in their lifetime but resulting in prolonged interactivity. In our dataset, one-timers constitute of over 31% of total users, who sends an average of 6 requests over their only session.
- **Active Users:** The rest of the user base i.e., which are not inactive, are classified as active users. Active users constitute of over half of the user population.

Table 2 shows basic statistics for the three user segments and in addition for the survey participants. Moreover, we perform a median split of the active user segment and list the relevant summary statistics for these two categories in Table 3.

User segment	Number of sessions	Usage duration
Dormant	≥ 1	0
One Timers	1	> 0
Active	> 1	> 0

Table 1: User Segments

User segment	% of Total users	% of Total requests	Average requests per user	Average sessions per user
Dormant	16.15%	0.59%	1.87	1.08
One-Timers	31.35%	3.93%	6.42	1.0
Active	52.50%	95.48%	93.14	9.94
Survey	1.02%	6.65%	332.34	31.75

Table 2: User Segments and Survey Participants Statistics

User segment	Requests per user			Sessions per user		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
Active Top 50%	164.2	86.0	228.91	16.47	9.00	22.34
Active Bottom 50%	22.13	18.00	16.97	3.40	3.00	2.35

Table 3: Active Users Median Split

4.5. Usage Metrics

Our goals are to better understand how the application is used, and if application usage changes with increasing time of usage. In this section we propose metrics to measure the focus of users and the variability of user requests. We begin by formally defining two key measures for each user i as: (a) Request Type Focus F_i (b) Request Type Entropy E_i as

$$F_i = \frac{1}{R_i} \max r_{ik} \quad (1)$$

$$E_i = - \sum_{k=1}^M \frac{r_{ik}}{R_i} \log \frac{r_{ik}}{R_i} \quad (2)$$

The idea of defining these metrics is mainly inspired by research work in the domain of information retrieval [7], large scale analysis of a popular online video service [3] and empirical analysis to understand content access behavior based on video-on-demand service [17]. It is important to note that while defining the focus and entropy measures for users, we take into account only the user specific requests and ignore the system requests. Now we describe these usage metrics in more detail.

Request Type Focus Given the distribution of a user's requests ($r_{i1}; r_{i2}; \dots; r_{iM}$) across all user specific request types, focus is the highest fraction of queries a user has sent for a single request showing a specific information need. In other words, focus metric measures how focused users are for their varied information needs, whether a user tends to access one request type category repeatedly than others or his access patterns are more spread across different request types.

Intuitively the focus value always lies between 0 and 1 as per definition, with higher values indicating low variability and more predictability towards user behavior, while low values describing the opposite. An F_i value of 1 shows that the user is interested in only one request type indicating a targeted and narrow application usage, while a value of 0 indicates that users requests are spread uniformly across all request types ($1/k; 1/k; \dots; 1/k$), hence highlighting high variability in application usage and providing little information about specific

information needs. Figure 4a shows the CDF of focus values with 50% of total users having focus values less than 0.5.

Request Type Entropy The idea of request type entropy is primarily inspired by Shannon’s work in information entropy [14]. Higher values of entropy indicate a usage pattern which is spread uniformly, while users with focused (and hence more predictable) usage are characterized by lower entropy values. Figure 4b shows the CDF of entropy values.

To give an example for how request type focus and request type entropy represent a user’s behavior consider two users: user A has 3139 total requests (with 1334 being user specific requests while the rest are system specific) in 267 sessions consisting of 4 different user request types: 1281 shopping list requests, 45 shopping list alert requests, 5 top bargains requests, and 3 retailer requests. This highly skewed distribution of request types results in a focus value of 0.96 and an entropy value of 0.27, representing a highly focused user. In contrast, user B has 3231 total requests (with 1725 being user specific ones) in 176 sessions. Her sessions consist of 8 different request types with a more even distribution (532 category bargains requests, 514 shopping list requests, 258 retailer requests, 223 shopping list alerts, 124 bargains given product identifier requests, 59 top bargains, 9 shop locator and 6 subcategory requests), resulting in a focus value of 0.39 and an entropy value of 2.16, representing a relatively less focused user.

As per definition focus and entropy values are inversely correlated. Higher values of focus relate to lower values of entropy and vice versa; the inverse correlation is clearly visible in Figure 4c, which has focus values on the horizontal axis and entropy on the vertical axis across all user segments. For the subsequent analysis and discussion, we choose to use request type focus as our usage metric because of the more intuitive range between 0 and 1.

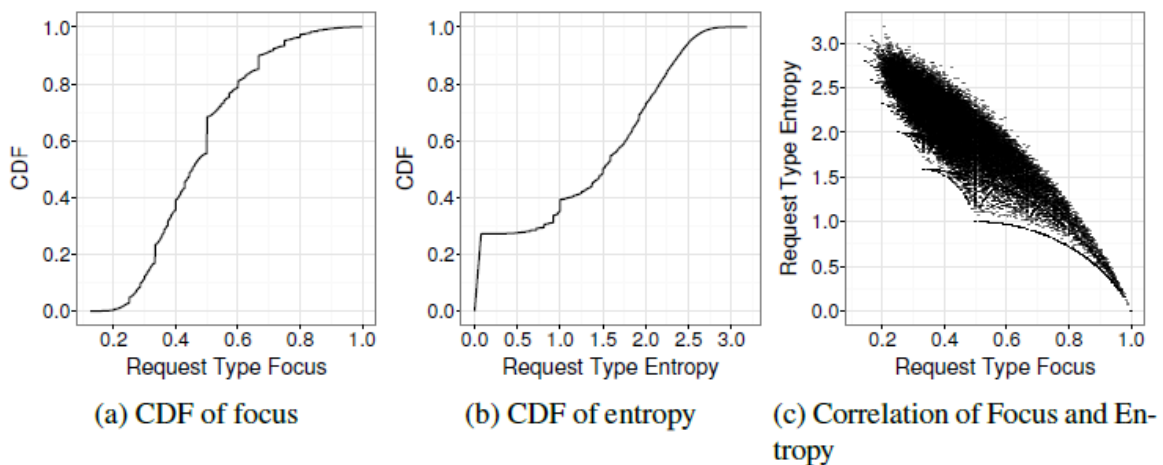


Fig. 4: Usage metrics

4.6. Request Type Focus and Session Duration

In order to understand whether the focus of users changes with increasing application usage, we examine the relationship between the request type focus values and total session duration for the active user segment. Figure 5a shows a scatter plot of total session duration, i.e. the total time of application usage in hours, for a given user on the horizontal axis and corresponding focus values on the vertical axis for the active top 50% users. Figure 5b shows the same plot with the usage time in minutes for the bottom 50% of active users.

4.7. User Demographics

From an online survey conducted amongst the users we collected demographics data including gender, age, education status and income levels. Since survey users are amongst the existing user base and the unique user identifier from the query logs is one parameter in the survey, we can link the survey with the log analysis in order to investigate the application usage of survey users. In this section, we analyze the request type focus across different demographic groups.

Of all 1,009 survey users, 253 of them are first time users, while the rest 756 users have used the application before. Even though we have a significant fraction of survey users (75%) who are familiar with the application, we don't know to which user segment the survey group belongs to. To investigate it, we compare the probability distribution of focus values for survey users with different user segment and as it turns out, survey users primarily consist of top 50% of active user base (plot omitted due to space constraints). Close to 90% of survey users are part of the active top 50% user group.

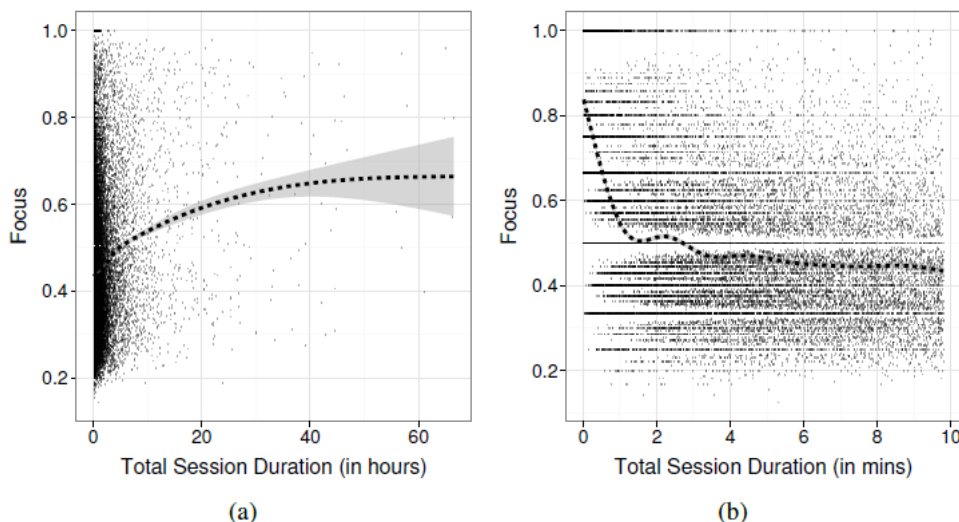


Fig. 5: Total session duration and corresponding focus values for active users (a) Top 50% (b) Bottom 50%. Note the different time units.

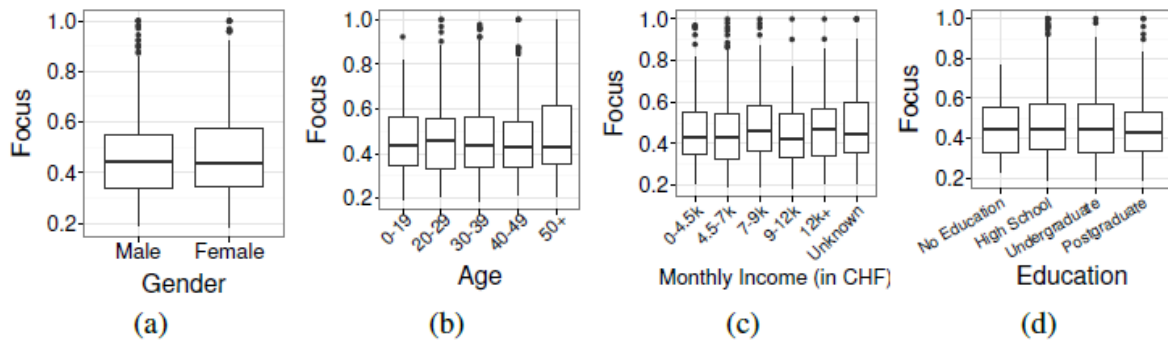


Fig. 6: Distribution of focus values for different user demographics (a) Gender (b) Age (c) Income (d) Education

Now to examine whether the focus values vary across different demographics, we analyze the distribution of focus values for different gender, age, income levels and education status. The box plots in Figure 6 show that the focus values do not vary significantly across these demographic groups.


5. Discussion

The total number of users is likely misleading when talking about real-world application usage as only a smaller fraction of users actually use the application. The CDF of user requests shows this very clearly. In our case we also find that the distribution of user requests follows the Pareto principle. It would be interesting to investigate if this is also the case with other mobile applications.

We consider the identification of active users an important part of usage analysis. We classified users into segments based on simple heuristics using the number of sessions and the total session duration. In the following we concentrated the analysis on the segment of active users, i.e. users with more than one session and a session duration greater than zero. The top 50% of active users is the user segment where most of the activity happens, and also the usage patterns in this group differ from the rest of users. As a result, we aim to understand their application usage in more detail.

We proposed two metrics to measure a user's focus when using the application. For us the request type focus seems to be more intuitive due to its range from 0 to 1. Figure 5a shows the trend that the top 50% of active users become more focused with increasing usage of the application. As active users interact more with the application, they seem to become more aware of which application functionalities best suit their respective needs and indulge in using only a few functions.

The same analysis for the bottom 50% of active users indicates a different trend – the focus values don't vary significantly with increased application usage. One of the reasons to



attribute this behavior to is that the bottom 50% active users spend a considerable less amount of time with the application compared to the active top 50% users: While the top 50% active users have an average of 92 minutes of total session duration, the bottom 50% users spend an average of 3.5 minutes with the application.

Combining quantitative log analysis with a qualitative user survey offers interesting opportunities for research, especially when survey results of individuals can be combined with corresponding usage logs and given that the privacy of users is preserved at all times. However, our analysis did not result in relevant findings as the user focus does not vary significantly over different user demographics.

6. Conclusions


In this paper we have analyzed the usage of a mobile bargain finder application for the iPhone using query logs from a period of 6 months. Using CDF plots we could show that the distribution of user requests follows the Pareto principle. We proposed simple heuristics to identify active users on which we concentrated our analysis. We also proposed metrics for measuring a user's focus when using the app and showed that active users tend to become more focused with increasing usage. We also combined the analysis of user focus with demographic data from an online survey and found no significant differences in user focus across demographic user segments. In future work we want to apply the same user segmentation and metrics to other mobile commerce applications and compare measurements and CDF plots for different applications. In the long term mobile commerce research could benefit from well established ways to measure active usage and user focus when analyzing mobile applications.

Acknowledgements We are grateful to Comparis for providing access to their dataset. We also want to thank Gilad Geron for conducting the online survey.

References

1. **Asthana, M. Crauatts, and P. Krzyzanowski.** An indoor wireless system for personalized shopping assistance. *Mobile Computing Systems and Applications*, IEEE Workshop on, 0:69–74, 1994.
2. **M. Böhmer, B. Hecht, J. Schöning, A. Krüger, and G. Bauer.** Falling asleep with angry birds, facebook and kindle: a large scale study on mobile application usage. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services, MobileHCI '11*, pages 47–56, New York, NY, USA, 2011. ACM.

3. **A. Brodersen, S. Scellato, and M. Wattenhofer.** Youtube around the world: geographic popularity of videos, 2012.
4. **L. Deng and L. P. Cox.** Livecompare: grocery bargain hunting through participatory sensing. In Proceedings of the 10th workshop on Mobile Computing Systems and Applications, HotMobile '09, pages 4:1–4:6, New York, NY, USA, 2009. ACM.
5. **A. E. Fano.** Shopper's eye: using location-based filtering for a shopping agent in the physical world. In Proceedings of the second international conference on Autonomous agents, AGENTS '98, pages 416–421, New York, NY, USA, 1998. ACM.
6. **G. Geron.** Mobile shopping assistant usage analysis - analysis of comparis smartshopper iPhone application. Master Thesis, ETH Zürich, 2011.
7. **M. Kamvar, M. Kellar, R. Patel, and Y. Xu.** Computers and iphones and mobile phones, oh my!: a logs-based comparison of search users on different devices. In Proceedings of the 18th international conference on World wide web, WWW '09, pages 801–810, New York, NY, USA, 2009. ACM.
8. **S. Karpischek, G. Geron, and F. Michahelles.** Towards a better understanding of mobile shopping assistants - a large scale usage analysis of a mobile bargain finder application. In Second Workshop on Research in the Large at UbiComp 2011, 18 September 2011, Beijing, China, 2011.
9. **V. Kotlyar, M. Viveros, S. Duri, R. Lawrence, and G. Almasi.** A case study in information delivery to mass retail markets. In T. Bench-Capon, G. Soda, and A. Tjoa, editors, Database and Expert Systems Applications, volume 1677 of Lecture Notes in Computer Science, pages 818–818. Springer Berlin / Heidelberg, 1999. 10.1007/3-540-48309-8 79.
10. **P. Kourouthanassis and G. Roussos.** Developing consumer-friendly pervasive retail systems. Pervasive Computing, IEEE, 2(2):32 – 39, april-june 2003.
11. **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.
12. **E. Miluzzo, N. D. Lane, H. Lu, and A. T. Campbell.** Research in the app store era: Experiences from the cenceme app deployment on the iphone. In First Workshop on Research in the Large at UbiComp 2010, September 26, 2010, Copenhagen, Denmark, 2010.
13. **G. Roussos, P. Kourouthanasis, and T. Moussouri.** Designing appliances for mobile commerce and retailtainment. Personal Ubiquitous Comput., 7:203–209, July 2003.
14. **C. E. Shannon.** A mathematical theory of communication. Bell Systems Technical Journal, 27:623–656, 1948.
15. **S. Shekar, P. Nair, and A. S. Helal.** iGrocer: a ubiquitous and pervasive smart grocery shopping system. In Proceedings of the 2003 ACM symposium on Applied computing, SAC '03, pages 645–652, New York, NY, USA, 2003. ACM.



16. **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.
17. **H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng.** Understanding user behavior in large-scale video-on-demand systems, 2006.