Exploring iPhone Usage: The Influence of Socioeconomic Differences on Smartphone Adoption, Usage and Usability

Ahmad Rahmati¹, Chad Tossell², Clayton Shepard¹, Philip Kortum², and Lin Zhong¹

Dept. of Electrical and Computer Engineering, ² Dept. of Psychology, Rice University, Houston, TX {rahmati, chad.tossell, cws, pkortum, lzhong}@rice.edu

ABSTRACT

Previous studies have found that smartphone users differ by orders of magnitude. We explore this variability to understand how users install and use native applications in ecologically-valid environments. A quasi-experimental approach is applied to compare how users in different socioeconomic status (SES) groups adopt new smartphone technology along with how applications are installed and used. We present a longitudinal study of 34 iPhone 3GS users. 24 of these participants were chosen from two carefully selected SES groups who were otherwise similar and balanced. Usage data collected through an in-device programmable logger, as well as several structured interviews, identify similarities, differences, and trends, and highlight systematic differences in smartphone usage. A group of 10 lower SES participants were later recruited and confirm the influence of SES diversity on device usage. Among our findings are that a large number of applications were uninstalled, lower SES groups spent more money on applications and installed more applications overall, and the lowest SES group perceived the usability of their iPhones poorly in comparison to the other groups. We further discuss the primary reasons behind this low score, and suggest design implications to better support users across SES brackets.

Author Keywords

Smartphones; Socioeconomic Status; SES; Applications; Diversity; User Study; iPhone; Mobile.

ACM Classification Keywords

H.1.2 [Models and Principals]: User/Machine Systems

INTRODUCTION

Understanding user diversity is a central tenet of human-computer interaction (HCI) research [1]. With an understanding of how users vary, designers can better support a broad range of individuals with different backgrounds, capabilities, skills and interests. Smartphone users have been described as extremely diverse [2]. Yet, little research has moved towards understanding these differences in more precise ways. To this end, we contribute a naturalistic and longitudinal study of how different SES groups use their iPhones. The study leverages an in-device, programmable,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobileHCI'12, September 21–24, 2012, San Francisco, CA, USA. Copyright 2012 ACM 978-1-4503-1105-2/12/09...\$10.00.

continuously running logger that collects device usage, complemented by regular interviews with the participants.

Our study has two unique features. First, unlike prior work that has very limited information about the participants [3-5], it achieves more statistical control over potentially confounding variables. Our users are extremely similar in age, attend college at the same university, live in similar dorms, and have the exact same experience levels with their device. However, they differ in their SES backgrounds, which we show is important for explaining user variance.

Second, our study logs smartphone usage for a longer period of time. Data is collected from most of our iPhone users for twelve months. In contrast, prior work is based on studies lasting at most a few months on Android and Windows Mobile based smartphones. This yearlong study allows us to study the adoption and long-term evolution of user behavior, which has been previously impossible. We chose the iPhone as, at the time of the study, it represented the cutting edge of smartphone design for usability, accounting for over a third of the US mobile Internet traffic as of April 2010 [6]. Additionally, iPhone users have access to the largest number of third-party applications, with over 300,000 officially released apps as of October, 2010. This enables our study to paint a comprehensive picture of how iPhone users employ their devices in real environments and capture longitudinal trends that were previously missed.

Taken together, the two unique features highlighted above enable us to contribute unique findings in this paper instead of presenting mere usage statistics. After we describe the research that informs the current study in the Related Work section, we describe the logging methodology used to collect real-world usage data from 34 participants. Because our users are all new smartphone users, we present results longitudinally to show device adoption and trends. While smartphone users are known to be diverse [2], we look not only at what the users do with the iPhones, but the influence of SES on how users differ.

In the Diversity and Dynamics of Usage section, we describe aggregate usage patterns. This includes the first empirical look at how users install and uninstall applications from the Apple App Store. We demonstrate the importance of a try-before-you-buy App Store, and show that web based versions of applications often entice users to install the full application.

In the Effect of Socioeconomic Status section, we carefully examine the affects of SES by utilizing our two carefully selected participant subsets. They both attend the same small private college, have similar experience with their iPhones, are equal in gender distribution, and are similar in other regards as well; except they differ in SES. We find that SES has a significant impact on usage, suggesting different needs and preferences for these groups of users in their specific contexts. In particular, we find that the web usage of low SES users is more of an extension to their PCbased web access, and that users' disappointment with the browsing experience on smartphones decreases their usage. Our findings indicating that lower SES users spend more time on the phone and more money on applications, suggest that the positioning and marketing of manufacturers who position their devices in a low end (cheap) - high end (expensive) manner, where the lower end phones are unable to run the latest high-resource-usage applications and games, and/or have a lower quality camera / display, is unfruitful.

Since the above SES differences manifested across a number of interactive behaviors, in the Lower Income Smartphone Users section, we recruit ten more users from one of the most underprivileged regions in a major metropolitan city. Usage patterns logged from these individuals confirm the influence of SES is important to understand the variance in how smartphone technology is adopted and used in real environments. Furthermore, the lowest SES group perceived their iPhone usability as poor in comparison to the other groups. Limited battery lifetime was their primary reason behind this low score.

Our findings have strong implications not only for understanding smartphone users, but for device and application design, optimization, and evaluation. In the Implications section, we highlight the value of long-term user studies with carefully selected participants. While we show the feasibility and limitations of smartphones as a primary device for IT access, especially for cost effective IT access in underserved communities, our results strongly suggest smartphone users could benefit from a better web browsing experience. Last but not least, we assess the one-size-fits-all phone paradigm, and show that even among our limited set of participants, there are distinctively different usage patterns that would benefit from phones with different hardware and software configurations.

RELATED WORK

Although there has been little work understanding the influence of SES on device usage, human factors of mobile devices have been an active research area for more than a decade. Most HCI studies in mobile space employ either lab-based evaluation or a short period of field trials [7]. In the last few years, as smartphones began to be widely adopted, there have been several relatively long-term field studies that we build upon in the present report.

The MIT Reality Mining project [8] studied 100 users of Nokia Symbian 60 series phones for one year. In [9], we studied 12 high-school users of Windows Mobile smartphones for four month. Both of these studies used previous generation phones, and their usage does not generalize to current smartphones. More importantly, they were

severely limited in their data collection capabilities. The first recorded running application, currently associated cell tower, and visible Bluetooth devices. The second recorded the screen status and detailed network conditions.

In [2], the authors studied 33 users of Android smartphones for 7-21 weeks. The authors did not have access to the participants for interviews or have demographic information about them beyond several predetermined user types. The data was analyzed mostly for usage statistics in the form of distributions, and the authors concluded that smartphone users are very different, without providing insights into why. In contrast, our study employed participants selected in a controlled manner, and a much longer period of study. This enabled us to gain insight into the long-term evolution of smartphone usage, and illuminate the differences and similarities of smartphone users. Moreover, with a superset of usage data, we are able to also analyze many new aspects of smartphone usage, including App Store utilization, application usage, and web access.

Most recently, using the methodology described in a previous report [10], we characterized user differences in mobile space by examining how the Internet was accessed on smartphones over time [11]. We found systematic differences between users, showing some were more exploratory and others concentrated their usage on favorite resources. These findings motivated design suggestions to better support users across a behavioral spectrum. However, this previous study only examined native applications that accessed the Internet. Heavily-used applications such as text messaging and voice phone were not examined.

Previous studies have shown that SES differences are important to consider for the design of other technologies. In [12], Goel et al. revisited the digital divide and found, among other differences, that SES differences drove how frequently web pages were accessed. Individuals in lower SES brackets accessed the web more than their higher SES peers. Similarly, we logged usage from iPod Touch devices and found that lower SES users without access to other technologies used their iPod Touch substantially more for activities commonly used on PCs [13].

This study analyzes all native applications used by a set of new iPhone users over a period of one year. This data reveals important influences of SES on application installation and usage through a longitudinal lens. Additionally, user differences in the types of applications installed and used over time reveal important distinctions between individuals and suggest more tailored functionality and systems to support smartphone users.

FIELD STUDY AND DATA ANALYSIS

Field Study Participants

Our main 24 users were young college students with an average age of 19.7 years (SD = 1.1 years), and the study lasted from February 2010 to February 2011. These 24 balanced participants were recruited from two distinct SES groups from Rice University, Houston, TX. They all lived

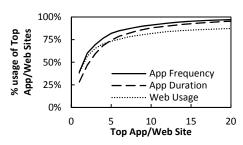


Figure 1: A small number of applications and websites dominate usage

on campus in dormitories. The low SES group was comprised of 13 students who received needs-based-scholarships and had a household income of under \$60,000. The 11 users in the high SES group did not receive scholarships and their household income was over \$80,000. Other factors, including their major, gender, race, PC access, and game console ownership, were balanced across groups. All had access to the university's computing labs, and a PC or laptop at their residence. All high, and 11 of the low SES participants had a personal laptop.

Every participant received a free iPhone for their participation, and was required to use the outfitted iPhone as their primary device for the entire year. Additionally, each participant received free service coverage, including 450 voice call minutes per month, unlimited data, and unlimited SMS, during the study. We helped all participants port their phone numbers to the iPhones.

Finally, approximately six months into the study, we added a third participant group of 10 students from a public community college in an underprivileged section of the same metropolitan area. We provided the same service plan and instrumented device to these participants, but for six months, from September 2010 to February 2011. The 10 users (M = 24.2 years, SD = 2.23 years) were in a lower SES bracket than the main 24 participants; they reported their annual income between \$35,000 and \$0. Yet, we note that there are other differences, beyond SES, between this group and the first two SES groups (e.g., occupation, age, and children). Therefore, we use this dataset carefully, and only in the Lower Income Smartphone Users section.

Data Collection

Logger Design and Implementation

The key component of the field study is an in-device, programmable logging software that collects almost all aspects of iPhone usage and context *in situ*. To run the logger in the background continuously, we had to jailbreak the iPhone 3GSs and exploit a setting provided by iOS 3.1.3, the latest at the beginning of our study. It starts the daemon process, as well as restarts it anytime it is killed. The main logger daemon is written as a shell script in bash and utilizes components written in various languages, including C, perl, awk, SQL, and objective C. Furthermore, the logger daemon is able to call built in functions, manage child processes, install and use programs from repositories, run custom programs, and add new features. We have implemented the

logger in a modular and robust fashion, thus a new iOS release may break individual components, but the main functionality is unaffected. In order to monitor and update the logger, it is programmed to report the logged data and, if necessary, update itself every day through an encrypted connection, via rsync [14], to a lab server. We employed several methods to limit energy consumption, and our measurements show that the logger consumed on average 5% of the battery per day.

While the logger recorded a plethora of context information, for this work we focus on logs regarding application installation, uninstallation, price, genre, and usage, as well as web usage. We note that given the nature of the study, there were short-term lapses in the log files of five users. These lasted from a few days up to over a month, and were caused by a number of reasons. These include bugs in our code, lost, stolen or damaged phones, travel, and phones that were accidentally erased by the users. We substitute data from missing days with the all time average of that user in order to maintain each user's uniqueness and to avoid magnifying the impact of short-term fluctuation in usage. We note that since there are only short periods of missing data on few users, we regenerate missing samples, and we analyze macro-dynamics (e.g., monthly usage) the overall effect of missing data is negligible.

Complementary Interviews

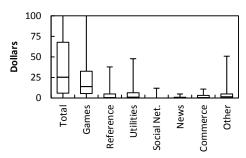
Since our study design allowed us to have access to the participants, we collected self-reported data alongside automated logging. The self reports were used to assess the participants' perceptions of their usage and their access to other IT resources, as well as help interpret the logged data.

Assuring Privacy

Collecting data from smartphones in the field naturally incurs privacy issues. We employ the following methods to protect privacy while retaining relevant information for research. First, we leverage one-way hashing to preserve the uniqueness of a data entry without revealing its content. For example, we hash the recorded phone numbers, names, and email addresses. With hashing, we can still construct call statistics without knowing actual phone numbers. Second, we perform information extraction in the device. For example, we extract emoticons from emails and text messages without collecting the raw content. Finally, we structure the research team so that the data analysis and logger development team do not directly interact with the participants, in order to avoid linking data to the actual users. A separate human factors team acts as the interface with our participants, but does not deal directly with the logger or access the raw data. This enables us to contact the participants in a privacy sensitive manner when necessary, for example to schedule impromptu interviews with users who exhibit a drastic change in behavior.

DIVERSITY AND DYNAMICS OF USAGE

In this section, we analyze the diversity and dynamics of application adoption and usage, for both built-in and App



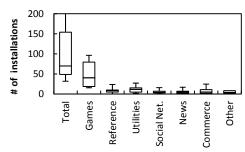


Figure 2: There is significant user diversity in paid application installations in terms of number (left) and cost (right). Broken down by category. Boxes: 2nd / 3rd quartiles. Whiskers: maximum / minimum. Horizontal lines: median

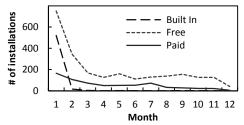


Figure 3: The ratio of paid to free application installations remained steady through the study, at $\sim 20\%$.

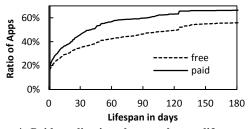


Figure 4: Paid applications have a shorter lifespan compared to free applications.

Table 1: We assigned categories to applications based on the genres reported by the App Store

Category	Genres	Notes
Games	Games, Entertainment, Media	Entertainment and media consumption
Utilities	Utilities and Productivity	Calculators, alarm clocks, todo lists
Reference	Books, Education, and Reference	Information resources
News	News, Sports, Travel, Weather	Contemporaneous information resources
Commerce	Business, Finance, Lifestyle (shopping)	Shopping or financial apps
Social Networking	Social Networking	Facebook, MySpace, Twitter
Other	Health, Navigation, Medical, Photography	Only a few (162) applications

Store applications. Our findings confirm the diversity of usage, yet, we will see in the later sections how we can extract order from this seemingly diverse usage.

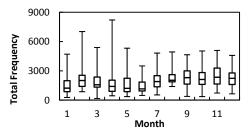
Application Usage

Our 24 participants installed over 3400 applications over the course of the study, of which over 2000 were unique. Our participants also purchased almost 750 applications, of which 500 were unique, from the Apple App Store, spending over \$1300. We were surprised to see 62% of the 3400 applications installed by our users were uninstalled during the study. We define the *lifespan* of an application as the time between its installation and its uninstallation. We notice that many applications have a short lifespan, (e.g., 20% were uninstalled within a day, and 31% within two weeks,) indicating that users tried but disliked them.

In order to analyze and more importantly present such a huge data set for behaviors and trends, we carefully assigned categories to applications, as shown in Table 1. The App Store already reports 20 genres for applications, but to the inconsistency of App Store genres and the fact that a certain application may be tagged by multiple genres, we had to carefully and manually categorize them. Further assisting us in analyzing the data set is the fact that each user's usage converges to a small set of applications. Figure 1 shows the median percentage of usage by each user's

monthly top applications. We can see that a small number of applications constitute a large share of our participant's usage in terms of frequency and duration. Approximately 40% of application usage is from the top application, and more than 90% is attributed to the top 10 applications.

The category with the most applications installed was games, accounting for over 50% of application installs, over 50% of money spent, and approximately 5% of application usage. In contrast, social networking applications, mostly being free, only accounted for less than 2% of money spent, but accounted for 8% of application usage. As expected, there was a wide variation between users in adopting paid and free applications. Our users spent a median of \$25 on 14 applications, as shown in Figure 2, and all but two users purchased at least one application. Figure 3 shows the total number of adopted applications during the study, broken in to built-in, free and paid applications. The ratio of paid to free applications stays relatively constant over time, at around 20%. Surprisingly, paid applications had a shorter lifespan overall, as shown in Figure 4. The large number of paid application with one day lifespan shows that users frequently purchase applications which they quickly determine they dislike, losing money in the process. The larger number of paid application uninstalls in the next months can be attributed to the large number of paid games.



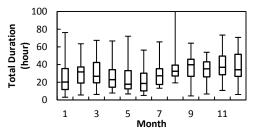


Figure 5: Application usage very diverse throughout the study, in terms of both frequency (left) and duration (right). Boxplots show $2^{nd}/3^{rd}$ quartiles. Whiskers: maximum / minimum. Horizontal lines: median

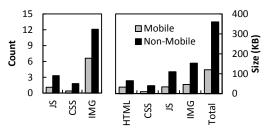


Figure 6: Mobile web pages are less content rich, in terms of the number of resources (right) and their sizes (right)

Even though categorizing applications allows us to analyze the application usage of our users, there is still a significant variation between application usage amount and frequency among different users. The differences, even among the second and third quartiles, highlight the fact that the average or median user alone is unable to serve as a benchmark for mobile usage. Instead, it is necessary to consider a wide variation of users and usage. Figure 5 shows a box plot of application usage by different users, for both frequency (left) and duration (right).

Web Usage

Similar to application usage, each participant's web usage converges to a small set of websites. As shown in Figure 1, the top website of a user accounts for 28% of web usage (median); and the user's top 10 websites accounts for 87% of usage. Compared to application usage, we found that users were more inclined to explore new web sites than applications, which is intuitive since visiting a new website requires much less commitment and time than installing an application. The key supporting evidence is the month-tomonth similarity of web usage, which is significantly lower than that of application usage, at (0.73-0.94), compared to (0.85-0.97) for application frequency.

While iPhone applications are developed for a smartphone environment, and are often tailored to the specific features of the smartphone platform, we expect web browsing to be an extension and supplement to users' regular browsing.

Our findings support this hypothesis, but strongly suggest users are disappointed with their web browsing experience. The users' responses to open-ended survey questions on their web browsing experience indicated that they generally were disappointed. Further, contrary to application usage, we observed a significant decrease in participants' web usage throughout the study. We hypothesize that the disappointment was due to poor web browsing usability caused

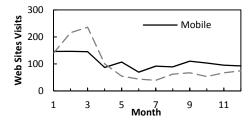


Figure 7: Median visits to mobile and non-mobile website per month

by factors such as connection latencies [15] and size of the device [11] according to the users' survey responses.

In order to further study the decrease in web usage, we analyze the web content browsed by our users. Our participants accessed both mobile and non-mobile websites. We classify web pages based on URL keyword matching; URLs that "m.", "mobile.", "iphone.", etc. are classified as mobile. Some popular websites, such as google.com, use the same URL for both mobile and non-mobile versions. In those cases, we assume the mobile version was used. Mobile web pages are less content rich than their non-mobile counterparts, in terms of styles (CSS), scripts (JS), multimedia content (IMG), and HTML size, as shown in Figure 6. Overall, the phone had to download 120KB for the typical mobile page and 3 times more, or 360KB, for the non-mobile page. Our results in Figure 7 indicate that the drop in web usage

through the study was due to a drop in non-mobile web usage, while mobile web usage, presumably better fit for mobile devices, remained relatively stable. The clear message for web designers is to develop a mobile version of their content. Indeed, this is even more important for users in lower SES brackets and new smartphone users as they transition and learn how to install native applications.

EFFECT OF SOCIOECONOMIC STATUS

Our second interest in this study is to assess the differences between SES groups in overall usage of their iPhones. We had expected differences to be minimal (e.g., how much they spent in App Store purchases) since both groups lived in the same dormitories on campus, and had no significant bias in their gender, major, PC access, or game console ownership. Surprisingly, our findings suggest stronger and broader differences in how they used their devices.

Application Usage

The low SES group consistently used more applications than their high SES counterparts, approximately 40% more,

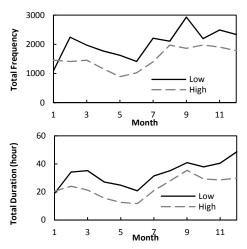


Figure 8: Median application usage was higher for low SES participants, in terms of both frequency (top) and duration (bottom)

both in terms of frequency and duration (Figure 8). For control, we assessed the differences between the low and high SES groups over the entire study period of one year. Visits to native applications and the web were combined for each user within each quarter. A 2 (SES: Low vs. High) x 4 (Time: quarters) analysis of variance (ANOVA) revealed that the low SES users more frequently launched applications compared to their higher SES peers (F(1, 22) = 9.73, p)= .01). A main effect for Time or interaction was not found. A similar main effect for SES was found for duration of usage (F(1,22) = 8.13, p = .02). The low SES users also consistently used a more diverse set of applications throughout the study, as shown in Figure 9 by the top 10 applications' smaller fraction of usage. The diversity is in part due to the low SES participants' higher variety of games used. Overall, the higher device usage and application variety in the low SES users suggests that the iPhones are used for both hedonic and utilitarian reasons (as defined in [15]) by the low SES group, but primarily for the latter for the high SES group. We hypothesize that this may be due to the low SES users having fewer or less interesting outside options, including access to entertainment devices.

Figure 10 is a radar chart showing application usage for each SES group, for the top 10 applications or application categories, normalized to the overall average usage of each application. Four of these applications or application categories revealed how SES groups differed: Facebook, phone, games, and utilities. Logistic regression confirms this statement; we compared the standardized logistic regression coefficients of each application or application category, as suggested in [16], to find the dominant predictors of SES. The results show that the top 3 dominant applications in frequency are utilities, games and phone; and top 3 in duration are Facebook, games and utilities, which comprise the exact same four applications.

Web Usage

By comparing the SES groups, we noticed that Web usage was initially much higher in the low SES group. However,

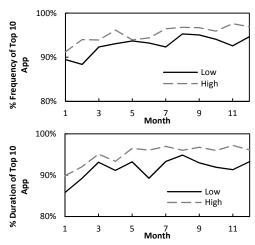


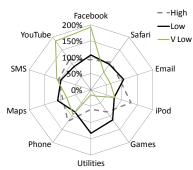
Figure 9: The top 10 applications contributed to a larger fraction of usage for high SES groups, in terms of both frequency (top) and duration (bottom)

the usage of both groups dropped, and their differences disappeared through the course of the study, as shown in Figure 11. Similar to the prior subsection, we used ANOVA for URL visits. Interestingly, it revealed that there were no main effects for SES or Time. However, a significant interaction showed that the lower SES group accessed the web more at the beginning of the study; over time, however, the differences in web use between SES groups attenuated, F(3, 66) = 4.60, p = .01. Correspondingly, duration of use followed similar patterns. The low SES users spent more time on the web early on; however, differences between the SES groups diminished as a function of experience.

While iPhone applications are developed for a smartphone environment, and are often tailored to the specific features of the smartphone platform, we can expect web browsing to be an extension and supplement to users' regular browsing. Higher initial usage for the low SES participants shows that smartphone web usage is more of an extension to their PC-based web usage. Indeed, interviews with both SES groups suggest that even though they both had access to personal and university PCs, the lower SES group owned older and lower-quality computers. In contrast with application usage, both SES groups had similar diversity in web usage, as shown in Figure 12. We attribute the similarity to the participants' previously established web browsing habits.

App Store Purchases

We had expected the low SES users to spend less on paid applications, but were surprised to find the opposite. They spent a median of \$31 on a median of 17 applications, compared to \$15 on 6 applications for their high SES peers (i.e., twice as much money on three times as many applications). Looking deeper into the data, we found that the low SES users were more money conscious and presumably more careful in their purchases compared to their high SES peers. This is shown by their significantly different prices paid per hour usage of paid applications. By dividing the total each user spent in the App Store by the total paid application usage duration, we calculate the cost per hour for paid ap-



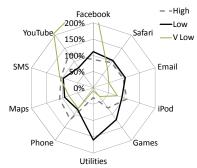


Figure 10: Application usage, relative to each application's average usage, for both SES groups in terms of frequency (left) and duration (right)

plications (price / duration). We found that the low SES users had significantly lower prices paid per hour (median: \$1.0 vs. \$2.6), which is substantial even considering the increased overall usage of the low SES users.

LOWER INCOME SMARTPHONE USERS

In the previous section, differences in SES influenced how users accessed their phones. The differences showed, among other things, the increased volume of usage for the lower SES group. In this section, we examine the use of iPhones in another group of community college students in one of the most underserved sections of a major metropolitan area. We limit the use of the data from this very low SES group, as there are other differences, beyond SES, between this group and the other two groups, such as occupation, age, and children.

Application Use

As expected, these ten users used their phone much more than the above users. However, the extent of the differences was surprising. Most comparisons between SES groups are presented as per use per day (PUPD) to account for the differences in study duration. Medians are used because of the large positive skew in the distributions. We found that the ten very low SES users from the community college accessed their devices 50% more frequently than our other users. As shown in Table 2, the very low SES group accessed applications at twice the rate of the high SES group. The differences in time consumed by applications were substantial as well. Interestingly, the differences between the frequency and duration of visits show that the nature of their interactions varied. The highest SES group spent less than a minute on average on each application. In contrast, though the lowest SES group yielded more app launches, the amount of time they spent within these launches was over one minute. This seemed to result from a wider diversity of resources used as shown by the amount of visits consumed by both the top 10% of applications and websites.

What applications drove these differences? The low SES students in the above section accessed the web, Facebook, iPod, and YouTube more than their high SES peers. This lower income group accessed Facebook and YouTube substantially more relative to their other applications (Figure 10). As reported more in the next section, this group did not access the web or utilities applications as much. In other words, for these users the device was much more for social network-

Table 2: Effect SES on application adoption and usage

SES:	High	Low	Very Low
App visits / day	57	74	96
Duration / day (min)	53	77	122
% Top 10 app visits	94	90	84
All app installs	0.30	0.37	0.64
Paid apps / month	0.5	1.4	2.3
\$ / month spent on apps	\$1.3	\$2.6	\$3.0
% Top 10 web visits	95	91	82

ing and passive entertainment. Interestingly, we again found that the SES level impacted the amount of apps installed and how much money was spent in the App Store. Overall, the very low users installed twice as many applications as the high SES group, which is significant even though most application installs were towards the beginning of the study. The former installed roughly two applications every three days. In contrast, both higher SES groups installed only one application during this same period. The lowest SES group spent more money on applications, contrary to our expectations. Table 2 shows the money spent by the very low SES users compared to the higher SES students. The dollar amounts are similar to the highest SES group, even though they used their phones for half the time.

Web Usage

We explored what users did on their web browsers to understand the above differences. Interestingly, even though the very low SES group showed very diverse usage, their adoption of the web was substantially slower than higher SES students. The very low SES users web usage generally increased during the six months, and their sixth month of usage significantly surpassed the main participants by 260%. Instead of applying a "PC mental model", the lowest SES users seemed to use their iPhones as mobile phones first. The fact that it was not until the second full month that the web was accessed regularly across users further attests to this finding. Once adopted, the web was accessed frequently. The very low SES group followed almost the reverse trajectory of the main participants, reported above.

We also notice a difference in the websites the very low SES group visited, compared to the 24 main participants. The URLs visited by each user were manually categorized into ten categories (Table 3). To increase reliability, we recruited three students to categorize the URL visits. The results yielded substantial agreement (*Kappa* = .83) [17]. The disagree-

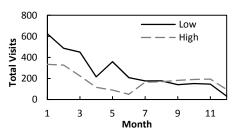


Figure 11: Median web usage was initially higher for low SES users, but became similar to high SES users

ments were reconciled by the authors. Most importantly, the main participants accessed their institution's website much more than the community college users, 24% vs. 2% of all web visits. Two other categories with large differences were religious and adult websites.

Usability

One aspect not usually captured in logged data is perceived usability of the device. We surveyed our users to assess the usability of the iPhone after the study was completed, using the system usability scale (SUS) [18]. The SUS is a ten question survey with scores ranging from 0 to 100, and has been validated extensively in a number of studies on a wide range of technologies [19].

The iPhone scored well overall across all users (M = 74.8, SD = 9.26). A one-way ANOVA comparing the three different groups in our study (high SES, low SES, and very low SES users)was significant showing that one of the group means was significantly different than the grand mean (F(2,31) = 8.51, p = .001). A Bonferroni post-hoc revealed that the low (M = 80.76, SD = 10.97) and high (M = 80.01, SD = 11.98) SES users did not reliably differ. However, the very low SES group (M = 61.4, SD = 14.04) was significantly lower than both other groups (p < .01).

What drove these differences? Open-ended questions revealed several factors. Most notable was that 9 of the 10 users mentioned their battery was either deficient or not functioning. Because these users were new smartphone owners, their perceptions of battery life did not match the actual battery life. 50% of the users complained to the researchers of bad batteries during the study period. The only other usability problems were related to page loading delays. This was mentioned by both groups; however, it was mentioned by 84% of the main participants (high and low SES), vs. 20% of the very low SES participants.

Energy Drain

These scores prompted us to assess the energy drain recorded in both groups. Recall that the lowest SES users not only used the phones more than others, but accessed YouTube and Facebook more frequently, relative to other applications. Such data and video intensive applications have increased power consumptions. This can be quantitatively seen in Figure 13. It shows the boxplot of each users average battery consumption per day, recharges (>1% charge) per day, and low battery warnings (at 20%) per day. Note that short duration syncing can count as a (partial) recharge, and that we normalize battery consumption with the energy capacity of a

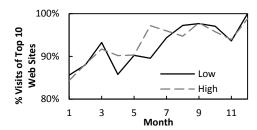


Figure 12: The top 10 websites contributed to a similar fraction of usage for both SES groups

Table 3: Per Use Per Day (PUPD) comparisons by SES

Category	High & Low SES (%)	Very Low SES (%)
Search	31	40
Institutional	24	2
Social / Blog	11	16
News / Sports	10	9
Commerce	10	8
Religion	4	9
Adult	4	10
Games/Movies	3	4
Health	2	1
Travel	1	1

full battery. We can see that the very low SES group consumes significantly more battery energy per day, and runs into more low battery situations. We note that over time, they appeared to learn the constraints of the battery and used their device more efficiently, reducing complaints.

IMPLICATIONS

We now elaborate on several implications of our finding on the design and evaluation of smartphones and smartphone applications. Clearly, SES differences influenced how iPhones were used. We controlled for experience with the device, type of device, temporal context, and other demographic factors (e.g., all students at the same university, age, gender, etc.) in the first study. Our findings suggest clear usage differences based on SES levels. In the second study, we found a very low SES group that accessed their device even more than the earlier groups and used it for more diverse functions. Of course, these users differ in other ways outside of SES (e.g., community college students, occupation, age, and children). We submit that SES, however, is at least a contributing factor in driving higher usage in the latter group and suggest several design implications to better support users at every SES level.

Application Development

First, our results provide insights into promoting third-party smartphone applications. Our findings regarding the application lifespan (Figure 4) show that users often try out applications for short periods, (e.g., a day). Unfortunately, neither the Apple App Store nor the Android Market offers trybefore-you-buy as a universal feature. Instead, users are typically expected to purchase applications based on reviews and word of mouth. However, our findings clearly indicate that users would benefit from a try-before-you-buy feature, such as the one introduced by the recent Windows Phone 7 plat-

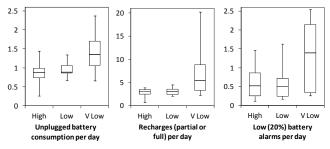


Figure 13: Very low SES users had significantly higher battery demands. Boxplot of users average battery consumption per day, recharges (>1% charge) per day, and low battery warnings (at 20%) per day.

form. This would enable them to waste less money, as well as potentially explore and purchase more applications. Additionally, real estate on iPhones is important and a try-beforeyou-buy store can facilitate users to quickly "clean house" if an application is not useful or engaging. Some operating systems, such as Windows Phone have already developed this feature. Our traces showing higher month-to-month diversity in web usage (Figure 12) highlight the fact that smartphone users are more comfortable exploring websites and web applications than downloading applications. Indeed, it is natural for users to be more adventurous in accessing different web sites than using applications; visiting a web site takes much less commitment than installing an application. This suggests that an application provider could reach a larger audience by providing a web service similar to its installation-based application when appropriate, so that first-time users can more easily assess them.

Designing Mobile Content

Second, many have envisioned feature-rich smartphones that provide cost-effective access to information technologies and entertainment, especially for users from underserved communities. This was one of the key motivations for our study to focus on SES. Our results support this vision: the low SES users tend to use smartphones more frequently and for more time than their high SES peers. Clearly, the web browser is more central to supporting low SES users adopt smartphone technology. When the low SES users first received their smartphones, they seem to use mental models developed through PC or laptop usage, manifesting as an increased reliance on the browser.

Over time, however, this reliance diminished in favor of native application use which was adopted earlier by their higher SES peers. In other words, low SES users, in addition to apps, require access to the mobile web to do things that could once only be done on PCs. Because many of these pages were not optimized for mobile use, it appeared they relied less on their browsers as a function of experience. Recall that the low SES users accessed more non-mobile sites which required more resources to load. The resulting page loading delays [20] have been noted as a primary cause of web usage declines on PCs [21]. Clearly, this is also a primary problem for the mobile web and this is especially problematic for low SES users. We note that since only a few top websites are

most commonly used, we suspect predictive capabilities can be leveraged to preload their most common resources and improve performance.

Smartphone Design

Third, based on the results of our SES comparisons, we identify several key groups of users that phones must cater to. We acknowledge that we observe these from a very narrow demography of smartphone users (college students), and that a broader user population likely has many more and different groups. Nonetheless, the significant differences in our narrow demographic strongly suggest that the one-size-fits-all paradigm fails to serve the best interest of users. Instead, multiple mobile platforms with appropriately selected features are more likely to compliment the needs of different user groups. While some features can be achieved through software and/or OS customizations, others require hardware changes (e.g., a hardware keyboard, game controller buttons, and small form factor).

Our very low SES users had much higher overall usage, placing greater requirements on the device's battery. Their web browsing was also shown to be more of an extension to their PC experience, increasing the value of larger screen sizes for them. Since both battery capacity and screen size come as a tradeoff to compactness, we hypothesize that, different users would significantly benefit from different choices in terms of these tradeoffs (e.g., higher capacity battery and larger displays for low SES users).

DISCUSSION

Field Evaluations

Our study provides important insights into how the field evaluation of smartphone and its service should be designed and carried out. First, our results demonstrate the importance of controlling for demographic factors to understand user differences. Prior work on smartphone usage was not particularly prudent in participant selection and, not surprisingly, failed to reveal any difference [2], or failed to provide conclusive evidence for speculated differences [9].

Second, our results demonstrated that extraordinary care must be taken in drawing conclusions from data collected by giving out devices and studying them in field for a short period of time (e.g., shorter than three months). Our results show that the first months see a significantly different degree of exploration and diversity in usage than in the remaining months (e.g., Figure 3, Figure 7, and Figure 9). Moreover, because usage continues to evolve even one year into the study, conclusions drawn using data collected from a short period of time should be generalized with care. Examples include the seasonal variation in usage, and applications losing appeal, as is often the case with games.

Third, our study demonstrated the value of following the same users for a long period of time. This is shown both by the significant usage changes in the later months of the study, and the SUS findings for the lower income users. However, this method is financially and administratively expensive, and therefore, can only be applied to relatively few partici-

pants. As a result, this method is complementary to those that gather data from a large number of users but only sparsely, such as [3].

User Perception vs. Actual Usage

Application usage patterns tell only part of the story. Our interviews provided complementary insights into what applications the users consider as the most important components of their iPhones, and the context in which applications are used. The interviews also gave us insight into why users utilize particular applications. The primary mismatch we found was what users perceived to be driving usability problems. Many of the lowest SES users reported their batteries were bad and requested new ones. This led to poor usability scores. Clearly, however, they used the device very frequently for a wide range of purposes. If we only gathered usage, we might suspect that their high volume reflects high perceived usability. On the contrary, the lowest SES users accessed their devices more than others at the same experience level. However, they also reported that the iPhone was not usable. The clear takeaway is that perceived usability reports are important to supplement logged data for a more holistic understanding.

CONCLUSION

We presented our findings from studying 34 iPhone 3GS users in the field. We highlighted the influence of SES on device usage, and revealed important differences between users that should impact how mobile content and technology is designed. Returning to our statement at the beginning of this report, one primary role of HCI is to understand user differences and better design technologies and content to support a wide range of users.

This study shows the influence of SES diversity in explaining how users access their devices. Indeed, users with the same device accessed it quite differently, though we contribute that they differ systematically. SES was one factor that influenced these differences. On one hand, the iPhone offered the lowest SES users access to technology for information and entertainment that was used very frequently, much more than others at higher SES levels. This suggests the device provided useful capabilities. On the other hand, the prevalent complaints about the battery life led to poor perceived usability. Thus, system designers should not only continue to work on energy efficiency, but in the mean time provide users with more options regarding the tradeoff between battery capacity and device bulk.

In particular, we found the device was poorly suited to serve the very low SES group. This suggests that different usage groups are best served by different phone designs. Our findings obviously don't cover all SES groups or smartphone users, and only focuses on a limited snapshot. Yet, our work raises questions and hypothesis beyond the scope of this paper. We hope that our findings will motivate researchers from multiple disciplines to work together toward answering them and, as a result, to offer even more insights into better mobile content and technologies.

ACKNOWLEDGEMENTS

This work was supported in part by the NSF grants NetSE-1012831, HCC-0803556, and CRI-0751173.

REFERENCES

- Shneiderman, B., Universal usability. Communications of the ACM, 2000.
- Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R. and Estrin, D., Diversity in Smartphone Usage. in *Proc. ACM MobiSys*, 2010.
- Huang, J., Xu, Q., Tiwana, B., Mao, Z., Zhang, M. and Bahl, P., Anatomizing Application Performance Differences on Smartphones. in *Proc. ACM MobiSys*, 2010.
- Falaki, H., Lymberopoulos, D., Mahajan, R., Kandula, S. and Estrin, D., A First Look at Traffic on Smartphones. in *Proc. Internet Measurement Conference (IMC)*, 2010.
- Bohmer, M., Hecht, B., Schoning, J., Kruger, A. and Bauer, G., Falling asleep with Angry Birds, Facebook and Kindle: a large scale study on mobile application usage. in *Proc. ACM MobileHCI*, 2011.
- AdMob. AdMob April 2010 Mobile Metrics Report http://metrics.admob.com/2010/05/april-2010-mobile-metrics-report/ April 2010.
- Hagen, P., Robertson, T., Kan, M. and Sadler, K., Emerging research methods for understanding mobile technology use. in *Proc.* ACM OZCHI, 2005.
- 8. Eagle, N. and Pentland, A., Reality mining: sensing complex social systems. *Personal Ubiquitous Computing*, 2006.
- Rahmati, A. and Zhong, L., Studying Smartphone Usage: Lessons from a Four-Month Field Study. *IEEE Transactions on Mobile* Computing, 2012.
- Shepard, C., Rahmati, A., Tossell, C., Zhong, L. and Kortum, P., LiveLab: measuring wireless networks and smartphone users in the field. SIGMETRICS Perform. Eval. Rev., 2010.
- Tossell, C., Kortum, P., Rahmati, A., Shepard, C. and Zhong, L., Characterizing web use on smartphones. in *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI)*, 2012.
- 12. Goel, S., Hofman, J.M. and Sirer, M.I., Who Does What on the Web: Studying Web Browsing Behavior at Scale. in *Proc. Int. Conf. on Weblogs and Social Media*, 2012.
- Tossell, C., Jardina, J.R., Kortum, P.T., Peres, S.C., Shepard, C.W., Rahmati, A. and Zhong, L., Effects of Socioeconomic diversity on iPod Touch Device Use in Real-World Environments. in *Proc. Human Factors and Ergonomics Society (HFES)*, 2011.
- 14. rsync. an open source utility that provides fast incremental file transfer http://samba.anu.edu.au/rsync/.
- Lee, I., Kim, J. and Kim, J., Use Contexts for the Mobile Internet: A Longitudinal Study Monitoring Actual Use of Mobile Internet Services. *Int. Journal of Human-Computer Interaction*, 2005.
- Menard, S., Six approaches to calculating standardized logistic regression coefficients. *The American Statistician*, 2004.
- Landis, J.R. and Koch, G.G., The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 1977.
- 18. Brooke, J., SUS-A quick and dirty usability scale. *Usability evaluation in industry*, 1996.
- Bangor, A., Kortum, P.T. and Miller, J.T., An Empirical Evaluation of the System Usability Scale. *International Journal of Human-Computer Interaction*, 2008.
- Wang, Z., Lin, X., Zhong, L. and Chishtie, M., Why are Web Browsers Slow on Smartphones? in *Proc. ACM HotMobile*, 2011.
- Sears, A. and Jacko, J.A. The human-computer interaction handbook: fundamentals, evolving technologies, and emerging applications. Lawrence Erlbaum, 2003.