

Need Accurate User Behaviour? Pay Attention to Groups!

Kasthuri Jayarajah, Youngki Lee, Archan Misra, and Rajesh Krishna Balan
Singapore Management University

kasthuri.2014@phdis.smu.edu.sg, {youngkilee, archanm, rajesh}@smu.edu.sg

ABSTRACT

In this paper, we show that characterizing user behaviour from location or smartphone usage traces, *without* accounting for the interaction of individuals in physical-world groups, can lead to erroneous results. We conducted one of the largest studies in the UbiComp domain thus far, involving indoor location traces of more than 6,000 users, collected over a 4-month period at our university campus, and further studied fine-grained App usage of a subset of 156 Android users. We apply a state-of-the-art group detection algorithm to annotate such location traces with group vs. individual context, and then show that individuals vs. groups exhibit significant differences along three behavioural traits: (1) the mobility pattern, (2) the responsiveness to calls / SMSs and (3) application usage. We show that these significant differences are robust to underlying errors in the group detection technique and that the use of such group context leads to behavioural results that differ from those reported in prior popular work.

Author Keywords

Groups; user behaviour; location; app usage; interruptibility

ACM Classification Keywords

J.4 Computer Applications: Social and Behavioural Sciences

INTRODUCTION

There is a rich history of UbiComp research that uses location, device usage or sensor traces to study and model a variety of behavioural attributes, such as a person's places of interest [21], her current activity [10], the likely dwell time at a place of interest [18] or the interruptibility level [22]. In recent years, the availability of a rich set of smartphone sensors has enabled such behaviour models to incorporate a wider set of contextual features, such as locomotive activity, ambient levels of noise and mode of transport. However, in almost all cases, the analysis treats each individual user's data as a singleton, ignoring the effect of physical world interactions with other users around them. We show, in this paper, that this distinction between group vs. individual context matters: *analyses that are oblivious to this dynamic context (whether the user is alone or in a group) lead to inference errors that can be significantly improved by making this distinction.*

Group interaction is generally recognised as a crucial context as people spend significant amounts of time in groups: a study on *American Time Use Survey* shows that people spend more than 8 hours per day with someone [16]. These groups can be of various sizes, ranging from 2-3 person groups (e.g., meeting friends at a cafe) to larger 10-15 person groups (e.g., tour groups). However, obtaining annotated high-fidelity, longitudinal traces is fairly challenging, as it requires simultaneous sensing of location traces from a large cohort of interacting people. (Not surprisingly, various recently popular data traces, such as the Nokia Mobile Data Challenge [15] or Device Analyzer [28], do not provide such observability.)

To study the importance of factoring in such group context into models of human behaviour, we analysed 4 months' worth of indoor location and mobile phone usage data (App usage, call and SMS logs) collected at our university as part of the LiveLabs large-scale mobile systems testbed [3]. We applied a state-of-the-art group detection system called *GruMon* [24] to partition the location traces into individuals and different sized groups. We then analysed (a) location (to study dwell time and next place transition behaviour), (b) SMS & call usage (to study interruptibility), and (c) App usage (to study patterns of content consumption) of these distinct partitions. Our primary goal is to identify and establish the statistical significance of those aspects of mobile user behaviour that appear to be affected by group context.

Our analyses helped us establish clearly that individuals behave *significantly differently*, across the above three behavioural properties, when alone vs. when in a group. Meaningful findings include: (1) Individuals' mobility patterns change starkly while in a group – they tend to spend significantly longer time at places and the larger the group becomes, the less likely they will move together to another place. (2) Individuals' propensity to respond to or initiate communication over the phone (e.g., calls and SMS) drops in the presence of friends / peers. (3) Individuals are likely to alter their app usage behaviour – for example, individuals in groups display *bursty* app usage, where they check their phones often, but restrict themselves to shorter usage durations.

Key contributions: Overall, our paper makes the following contributions:

Fine-grained, Longitudinal, and Large-Scale Group vs. Individuals Study: Prior work in capturing group vs. individual behaviour have mostly used coarse-grained cell tower traces, or observed the interaction among a relatively small set of users. Such studies suffer from sporadic and partial observability. In contrast, our observations span several thousand students on the university campus, whose locations were

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UbiComp '15, September 7–11, 2015, Osaka, Japan.
Copyright 2015 © ACM 978-1-4503-3574-4/15/09...\$15.00.
<http://dx.doi.org/10.1145/2750858.2804289>

tracked at $\pm 6 - 8$ meter granularity over an entire term (4 months) allowing us to establish meaningful group-related statistics.

Group vs. Individual: Key Differences: We employ statistical analysis to show that groups and individuals differ in (1) key mobility properties, including next-place transitions and residency times, (2) interruptibility propensity (measured in terms of length of SMS and phone call interactions, and responsiveness to notifications), and (3) App usage (measured in terms of session duration and frequency).

Demonstrating Statistical Difference from Past Results: We show that differentiating user behaviour using group vs. individual context results in behavioural patterns that are markedly different than those reported previously. If our entire corpus of user activity is viewed in a group-oblivious manner, our behavioural statistics (e.g., per-user call duration or App session duration) match previous results (such as De Melo et. al [8] and Falaki et. al [12]). However, if user behaviour is conditioned separately on group vs. individual context, the resulting distributions are significantly different, thus demonstrating the importance of *incorporating group context into models of user behaviour*.

Establishing Robustness of Results under Uncertainty in Group Context: Our underlying group partitioning is inherently noisy, given the medium-level accuracy of the underlying location system. To overcome the lack of ground truth (to establish that our claims of statistical validity are not affected by this underlying uncertainty in the inferred group context), we inject controlled amounts of random noise and show that our conclusions (both within the individual and group observations, and across the two sets) are robust to noise levels that are higher than the reported classification error rates of the group detection algorithm.

It is worth noting that the analyses presented in this paper focus only on aggregate behavioural characteristics (do individuals as a whole differ from the groups?), and not on establishing how individual-level attributes (e.g., demographics or social popularity) moderate such behavioural differences.

RELATED WORK

In recent years, substantial progress has been made on understanding various properties of online and physical behaviour of mobile users, such as call duration, app usage, mobility and conversational interactions [6, 7, 9, 11, 12, 20, 25, 28, 29, 30]. However, we have limited understanding of how users change their behaviours when they are part of groups. It is important to develop such knowledge, as group interactions constitute a significant portion of people's everyday lives, and such understanding can enable mobile services and interfaces to better adapt to the unique characteristics of group behaviours.

Social Sensing: Our work is closely related to a number of social behavioural studies based on smartphone data [11, 17, 19, 29]. The Reality Mining project [11] shed light on various social behaviours and dynamics of communities including the diffusion of opinions or evolution of social networks within a community. Recently, as part of the StudentLife project [29], the authors investigated if social features, i.e., frequency and

duration of conversations, can be used to estimate level of depression, stress, and mental wellbeing. While our work also focuses on studying the relationship between social interaction characteristics and user behaviours, we focus specifically on behavioural differences manifested as a function of underlying group interactions. We consider our work to have a wider coverage as we study the location traces of 6,000 users compared to 100 and 48 users in [11] and [29], respectively.

Mobility: Mobility patterns of human being using various location data has been well-studied [2, 4, 5, 23]. Our studies on mobility are distinct from such past work as: (1) we distinguish between group and individual mobility, (2) we focus on fine-grained movements (mostly indoors) within a small urban university campus compared to coarse-grained GPS or cell-tower data collected in large geographical areas, (3) we utilise a data-driven model to *empirically* understand how the movement of groups differs from that of individuals as opposed to basing movement models on assumptions as found in works related to ad-hoc networking protocols.

Works such as [7, 20] consider groups and social ties to better understand people's mobility. For instance, Brown et al. reported findings on group behaviour based on check-in data of location-based social networks [7]. Our work complements such insights to understand indoor movement patterns in a densely populated urban campus.

Phone usage. People's smartphone usage behaviour [6, 9, 12, 25, 28, 30] has been studied extensively. For example, based on a national-scale analysis of aggregate usage patterns, Xu et al. [30] showed how, when and where applications are used. Bohmer et al. [6] studied the spatiotemporal variation in the App usage behaviour of 4,100 Android users. In an effort similar to ours, Do et al. used the surrounding Bluetooth device density as a proxy for human density, and reported that mobile users use certain Apps in crowded environments [9]. Our work uses longer-term observations of location traces to explicitly identify groups, and thus avoids conflating crowdedness with group context (an important distinction in locations such as the food court, library and group study rooms).

Interruptibility. In the field of CHI, a long thread of research has attempted to understand and estimate human interruptibility, so as to derive better policies for notification management [13, 14, 22]. However, these studies do not account for the group context of the user, explicitly, although it is generally viewed as an interruption to receive notifications when people are physically together.

METHODOLOGY AND DATASET

Dataset

For this study, we used a dataset collected on the LiveLabs mobile lifestyle testbed [3] located at our university campus. Our university has about 7,000 undergraduate students and its campus comprises five 5-storey academic buildings. The dataset comprises indoor location data collected directly from the Wi-Fi controllers in the campus and phone usage data collected from a smaller number of undergrad student volunteers (over a period of 4 months, from Aug-Dec 2014). We describe each dataset in more detail below.

Label	Membership Size	Likely Interaction Type
Solo	1	By Themselves
Small	2–3	With Close Friends
Medium	3–7	In Project Groups
Large	>7	Class or CCA Activities

Table 1. Output partitions from *GruMon* used in the entire paper

Indoor location data: Our main dataset is indoor location data collected from over 6,000 devices connected to the campus Wi-Fi. The indoor locations are computed using a Wi-Fi fingerprinting approach and a RADAR-like [1] algorithm, based on AP-side measurements, to achieve room-level location accuracy (error within 4–8 meters) for the vast majority of the campus. Each device’s location is refreshed once every 2–3 minutes. All location traces were anonymised (using 1-way hashes of a device’s MAC address) and permission to use this anonymous data was provided by every campus user as part of their computer account signup process. **Note:** because we are using AP side measurements to compute locations, this approach captures all Wi-Fi connected client devices without installation of additional software and energy overhead, and thus eliminates resulting selection biases.

In this paper, we only consider mobile phone traces. To filter out non-phone traces from our location dataset, we use the following heuristic: we consider a location to be *transient*, if a device had spent less than 5 minutes at any location. A device is likely to be a laptop if it *teleports*, i.e., it moves between places but shows no intermediate *transient* locations. For each unique device in the location trace, we computed its *teleport ratio*, i.e., the the number of days the device was seen to ‘teleport’ divided by the total number of days it was observed on campus. Through empirical analysis, we found that, at a ratio of 0.3, the misclassification of known mobile devices (the 1,468 devices that that registered testbed participants had explicitly provided) was only 0.8%. Accordingly, for the rest of this paper, we considered only the subset of devices whose teleport ratio is lower than 0.3 and were seen on campus for at least two months out of the four-month observation period.

Phone usage data: We also collect detailed phone usage data from a smaller set of 156 Android smartphone users. These volunteer users, who provided informed IRB consent, installed our data collection App (which does not require rooting of the phone) on their primary smartphone. This App, runs in the background and logs, with accurate timestamps, the following information; (1) incoming and outgoing call and SMS logs, and (2) application start times and durations. These users were not told what analysis we would perform on the data (other than that they contributing data to a “long term data analysis project”). In addition, we collected their data over a long time period (the entire Fall 2014 semester), to minimise any bias caused by novelty effects.

This data is then regularly uploaded, when the phone is idling and connected to WiFi, to a central data collection server. We use the anonymised MAC addresses of these volunteer smartphone users to locate these individuals precisely in the location data described above. The demographics of the 156 users were Males – 61%, Female – 39%, 4th Year Student – 9%, 3rd Year – 15%, 2nd Year – 21%, and First Year – 55%.

Analysis Process

We used the three-step analysis to differentiate between individual and group behaviour: (1) application of a group labelling algorithm, (2) statistical hypothesis testing to identify differences, and (3) sensitivity analysis to determine the robustness of our results. We below describe each step in detail.

Step 1: Applying a Group Labelling Algorithm. We applied a group labelling algorithm on location traces for the entire Fall 2014 semester, to segregate the data into multiple mutually exclusive partitions – solo individuals and groups of various sizes (see below). This is not a straight-forward task as our dataset does not have any explicit group or individual labels added by either the data collection process or by the users. As stated earlier, we used the *GruMon* system [24] to do this partitioning.

GruMon extracts key features such as dwell time and place transitions from location streams, and declares a set of people as belonging to the same group if they have high feature correlations. It is pre-trained using a dataset of more than 250+ mobile users collected in other urban spaces, and is reported to have a >90% precision and >80% recall in detecting groups.

Using *GruMon*, we partitioned the dataset into 4 partitions; (1) *Solo*: a partition containing individuals moving by themselves, (2) *Small*: a partition containing groups with memberships of two to three people, (3) *Medium* a partition containing groups with memberships of four to seven people, and (4) *Large* a partition containing groups with eight or more people. Each partition contains the location traces for every person in that trace as long as the invariant for that partition holds. For example, *Solo* contains the trace for Person A as long as Person A is by themselves. The moment Person A becomes part of a larger group (say a group of size 2) and *GruMon* detects that, all of Person’s A’s traces thereafter (until the next change in membership occurs) will appear in the *Small* partition only. **Note:** The group sizes were chosen for the following reasons: *Small* was chosen to represent the common case of students hanging out with their close friends, *Medium* was chosen to represent the common project group sizes on our campus, while *Large* represent groups that formed as part of class or extra curricular (CCA) activities. Table 1 summarises the output of the group labelling process.

Step 2: Hypothesis Testing to Identify Differences. With the labelled outputs from *GruMon*, we now identify various hypotheses that are interesting to the UbiComp community. In particular, we explored the following three classes of hypotheses: (1) Do groups show significantly different mobility patterns from individuals? A positive result here could have deep implications for location-based systems that use mobility predictions as inputs to their operation. (2) Do groups have different interruptibility patterns from individuals? Again a positive result here could have implications for systems that determine the best time for content delivery. And finally, (3) Do groups use applications differently from individuals? A positive result here has strong implications on techniques that prefetch App-specific content so as to balance App responsiveness and energy overheads.

Note: when we says "Groups do X", we refer to the aggregate behaviour of the individuals who make up that group. In particular, we show that individuals who are by themselves behave differently from individuals who are part of a group – at an aggregate level. We do not perform any individual analysis comparing the behaviour of a specific individual when they are by themselves as compared to when they are in groups; this type of individual analysis is deferred to future work.

For the mobility hypotheses, we used the location trajectory data from over 6,000 users on our campus, while the interruptibility and App usage hypotheses used the partitions involving only the 156 participants who installed our data collection software (as these hypotheses required the corresponding collected data). For every result we obtained, we conducted Kolmogorov-Smirnov tests (KS test), with alpha set to 0.05, to determine if the differences were significant. We used the KS test as it does not assume any underlying distributions and is fairly robust to outliers. The results section in this paper report only the results that are significant.

Step 3: Determining the Robustness of the Results: An inherent problem with our approach is the errors caused by the labelling process using *GruMon*. The *GruMon* authors report a 91% precision and 82% recall. However, these values could result in 10% of the group data containing individuals and 20% of the solo data containing groups, thereby potentially invalidating the conclusions from our hypothesis testing.


To address this issue, we systematically subjected the *GruMon* output to various levels of random noise to determine robustness. Specifically, we flipped a random percentage, using values of 1%, 5%, 10%, 15%, and 20%, of the output of *GruMon*; i.e., if an element belonged to a group, it was now marked as Solo and vice versa.

We then tested these new "noisy" distributions, using the KS test, to determine if these "noisy" distributions were significantly different from the original data partition. We also compared the hypothesis analysis of "noisy" versions of the Solo and group data with the original Solo and group data to see if the results had changed significantly. In the rest of this paper, we only present results which were robust (i.e., there was no significant change observed) even under all five different noise levels. We present more detailed robustness check results at the end of the paper.

Limitations & Threats To Validity: There are two major threats to validity for this paper. We account and control for one but not the other. The first threat, that the errors inherent in our group detection algorithm could result in significant mis-classification errors, is handled by our random noise addition-based sensitivity test described above. The second threat to validity is that our dataset only captures the behaviour from an undergraduate population, observed only on the campus. Hence, how well the results generalise to other population demographics is currently unknown.

Time Spent in Groups

Before we present the results of our hypothesis, we first present an important initial results. Namely, we show that



Each boxplot shows the avg, 25 and 75 percentile, and stdev. values.

Figure 1. Time spent in various locations and in groups


users spend a significant proportion of their time in groups—and thus, reiterate the importance of studying behavioural changes in the group context. We observed the amount of time the users spent in various locations on campus, over the fall semester – differentiated by the group membership (Table 1).

Figure 1 shows the time spent by all the users in our location trace in four distinct campus locations – *GSRs* (Group Study Rooms that can accommodate up to 8 people), *Study Areas* (Public places on campus where student can congregate for project and study work), *Library* (study areas in the library), and *Food Court* (the main food court on campus). In addition, we show the percentage of time a user spent by themselves, and in *Small*, *Medium*, and *Large* groups respectively. We make the following observations: (1) students spend a large portion of their time on campus, 9.8 hours on average over the Fall term, and out of 9.8 hours, they spent 84% of time outside classes (45.29% in *Study Areas*, 14.50% in *Library*, 2.56% in *Food Court*, 3.48% in *GSRs*, and 18.17% in multiple other areas around the campus), (2) excluding class times, students spent 64.62% of their time in groups (in any size of group), indicating the importance of understanding group behaviour, and (3) students are engaged in various social activities with different groups sizes. Out of the total time spent in groups, they spent 23.43% of their time in *Small* groups, 16.31% in *Medium* groups and 24.88% in *Large* groups.

ANALYSIS OF PRIOR WORK

Before we present hypothesis-specific results, we first establish that, if group context was accounted for, results from prior work on mobile phone usage analysis would in fact differ significantly. To do this, we replicate the analysis of prior work on our complete dataset (both solo and group behaviour data together) and show that our results closely match prior work. We then show that the analysis results are quite different when we do account for solo versus group behavior. This confirms our assumption on the importance of accounting for group behaviour when analysing user behaviour.

We use two established well cited studies, De Melo [8], and Falaki et. al [12]. In the interests of space, we pick only one representative result from each paper to compare against. For this analysis, we used three datasets; *All* – Our entire dataset without any partitioning, *Solo* – a partition containing only



The figure show the odds ratio plot of call durations of one sample user. The three different lines represent the fitted regression lines for (1) *All* (adjusted R^2 0.9719), (2) *Solo* (adjusted R^2 0.9573) and (3) *Group* (adjusted R^2 0.9695).

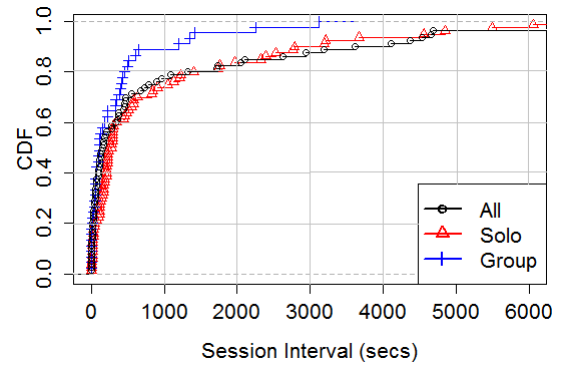
Figure 2. Extending the analysis of De Melo et. al

the location traces when users were by themselves, and *Group* – a partition containing only the location traces where users were in some group of any size. Note: *Solo* + *Group* = *All* with *Solo* and *Group* being mutually exclusive.

In De Melo et. al [8], the authors showed that the per-user call duration distribution follows a log-logistic distribution [26] (a continuous distribution that experiences an initial rate increase followed by a rate decrease). The distribution is truncated as the duration values can only take on positive values. Figure 2 shows the Odds Ratio plot for a sample user from the dataset. This is a graphical representation of the log-logistic distribution and each line represents the odds ratio (defined as $p/(1-p)$, where $p \in (0, 1)$) of the log-logistic distribution. A perfect fit would be a straight line (crossing at zero with a 45 degree gradient) in the log-log scale. The *All* dataset is the solid black line in the centre of the figure. We observe that this is a nice straight line (with an adjusted R^2 value of 0.9719) that closely matches the observations in De Melo et. al. However, we observe that the odds ratio plot changes significantly when we partition the data. In particular, the *Group* plot (the dotted red line) has a much steeper gradient than any other line while the *Solo* line (blue dotted line) has the shallowest gradient. This strongly suggests that *Solo* and *Group* behaviour are quite different from each other and then combining them leads to some sort of “averaged” behaviour.

In Falaki et. al [12], the authors show that the per-user app session intervals fits a Weibull distribution [27] (a flexible distribution that can effectively model many different data distributions). They also note that the *shape* parameters over all users to be less than one indicating that the longer the screen has been off, the less likely it is for it to be turned on by the user. We observe the same in our data but with strong variations when we account for *Solo* versus *Group* behaviour. In Figure 3, the CDF of a sample user’s observed session intervals are plotted against the fitted theoretical Weibull distribution. We observe that all three partitions have a good fit but they are significantly different. As with the previous result, this strongly suggests that *Solo* partition exhibit very different behaviour from the *Group* partition.

Although we limit ourselves to two results, we demonstrated two important points; (1) our overall dataset has very similar



The graph plots the CDF of session interval for all 3 partitions. The D statistic and p values for *All* are (0.0982, 0.231), *Solo* (0.1, 0.8186), and *Group* (0.1111, 0.944). The *shape* and *scale* parameters for the entire dataset are (0.513, 370.7754), *Solo* (0.5254, 435.72), and *Group* (0.5847, 279.92).

Figure 3. Extending the analysis of Falaki et. al

properties to the datasets used in previous studies, and (2) accounting for group interactions can result in very different results compared to treating the entire dataset as comprising solely individuals who do not interact with each other. We next delve into our results for our three classes of hypotheses.

MOBILITY

In this section, we present our observations on whether and how being in a group can affect the mobility pattern of individuals. In particular, we investigate two aspects of mobility; (1) dwell times and (2) semantic place transitions between the four distinct campus location earlier – *GSRs*, *Study Areas*, *Library*, and *Food Court*.

Dwell Times

We calculated the dwell times at all four places for every partition listed in Table 1. We first conjecture that the groups tend to stay for longer due to the increased interactions among group members. For example, in the case of *Food Court*, an individual’s objective is to grab a meal whereas when accompanied by a group of friends, conversations could result in prolonged dwell time. Hence, we hypothesize the following:

HYPOTHESIS 1A: *Groups Stay at Places Longer*

Overall, we found that individuals behaved different to groups, regardless of size, sizes tended to stay longer at all locations with the differences being statistically significant (p -value $< 2.2e-16$). In general, we observed that larger groups stayed for the longer durations. Figure 4 shows the dwell time CDF at a food court on campus for all four partitions. The difference is more drastic for *GSRs*; only 15% of *Large* groups spend less than 15 mins whilst almost 65% of *Solo* and *Small* spend less than 15 mins in *GSRs*. The *GSR* result also shows greater variability in the 15 to 120 min range after which all four configurations tend to merge exhibiting a long tail. In the case of *Study Areas*, 10% of *Solo* and 20% of *Large* stayed on for more than 2 hours, with the difference becoming weaker thereafter. Note that the dwell times are shown in 15 minute increments as we used 15 minute windows for detecting groups with *GruMon* [24].




Figure 4. Comparisons of dwell times across group contexts and locations (Hypothesis 1A)

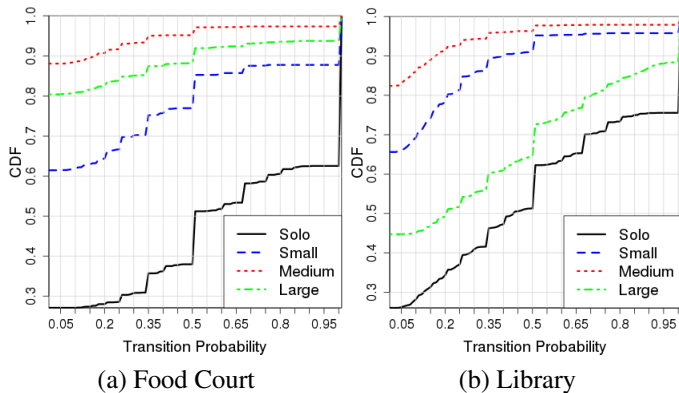


Figure 5. Comparisons of transition probabilities across group contexts from Study Area (Hypothesis 1B)

The *Library* showed the least difference between the partitions (lowest values of the D-statistic consistently). We believe that this could be due to the intrinsic nature of the library; although groups may visit the library together, the activity of studying or reading in the library, by itself, is not a group activity and hence limits the conversations or interactions amongst the group. This results in the dwell time distributions of individuals and groups at the library to be more similar than at other locations.

Semantic Place Transitions

Next, we consider the likelihood of groups versus individuals in making a transition from one place to another. For every pair of places ($p1$ and $p2$), we calculated probabilities that a user moves from $p1$ to $p2$ based on past history calculated separately for when that user was in a group and when they were alone. Thus, for each place pair, we have two sets of probability values, one for group and one for individual, for all users. We used these two probability value sets as inputs to a KS test which showed that our results was significant. Note that, for simplicity of explanation, we describe only two cases above (groups and individuals). In reality, we computed probabilities for all four partitions. With these four probability sets, we hypothesize the following:

Partition	Study Area	Library	GSR
<i>Solo</i>	0.0582 (0.140)	0.1234 (0.274)	0.0276 (0.101)
<i>Small</i>	0.034 (0.123)	0.1818 (0.281)	0.0108 (0.069)
<i>Medium</i>	0.0130 (0.086)	0.1228 (0.258)	0.0037 (0.048)
<i>Large</i>	0.030 (0.125)	0.1086 (0.233)	0.0101 (0.078)

Table 2. Prob. of next place transitions from Food Court (Hypo. 1C)

HYPOTHESIS 1B: Larger Groups Make Less Transitions

As hypothesized, we observed significant differences in the transition probabilities between different partitions and pairs of semantic places (p -values $< 2.2e-16$). Figure 5 shows the probability of transitions from *Study Area* to *Food Court*. We observe, for example, that at a transition probability of 65%, at least half of *Solo* had that probability of transitioning whereas only 14%, 3% and 8% of *Small*, *Medium*, and *Large* had a 65% probability of transitioning. Overall, we consistently observed that *Medium* and *Large* moved significantly less than *Solo* and *Small* (p -value of the KS-test between *Medium* and *Small* was $< 2.2e-16$). This means that although larger groups tend to spend more time together, they are less likely to move to different places together. An alternate explanation could be that the campus had fewer places that could accommodate larger groups (thus resulting in fewer locations to transition to) – but we ensured that all four places considered did in fact have the capacity to host at least *Medium* groups. Interestingly, we observed that transitions to *Library* were more popular by *Large* compared to *Small* and *Medium* (See Figure 5). However, even here, *Solo* was still the most likely to make this transition.

HYPOTHESIS 1C: Groups Have Different Place Transitions

Table 2 shows the next place transition probabilities (mean values with stdev in brackets) for all partitions from *Food Court*. We consistently observe that *Solo* has significantly different transition probabilities, compared to the three group partitions, to all three *next places* with p -values less than $2.2e-16$. However, we also observe similarities between larger groups – for instance, the transition probabilities between *Medium* and *Large* are not different statistically for transitions between *Food Court* to *GSR* (p -value 0.1433). We see similar “similarities” and “disagreements” between all the 4 partitions. Overall, this shows that being in a group (and the size of the group) does affect the probability of where you will go next.

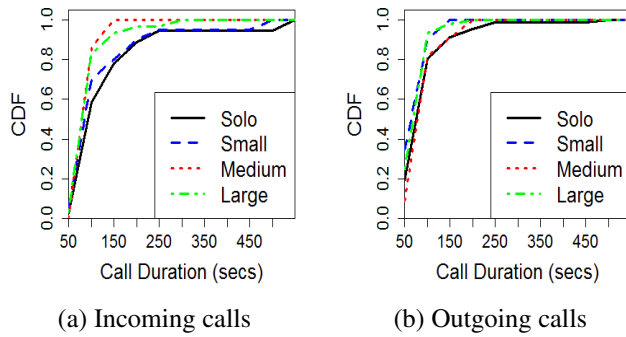


Figure 6. Call durations at *Food Court* (Hypothesis 2A)

Summary and Use Cases

Our observations lead to key lessons: (1) first, the amount of time people spend at places, where they go to next, and whether they will leave the place together have non-negligible dependency on the group context, and (2) second, such dependencies are pronounced depending on the semantics of the locations, and finally, (3) it not only matters whether the individual is in a group or not, but also, to a certain extent, the size of the group he/she is in. Below, we discuss example use cases where the additional consideration of the group context can lead to more accurate context-awareness.

Vacancy estimation: Indoor public areas (such as malls, airports, university campuses) consistently experience space crunches. In such crowded environments, a vacant-space prediction service can be quite useful. However, to predict accurately, it is important to model groups separately from individuals, along with the type of location and the size of the groups. For example, waiting times to be seated can be estimated accurately by computing the dwell times separately for the different sized groups currently at a restaurant operating at its full capacity. On the other hand, for places that do not exhibit strong group effects, such as the library, using a single dwell time distribution across all users may suffice.

Personalized advertising: Personalised coupons or recommendations triggered by various contexts (e.g., location) are popular. In particular, to accurately predict which coupon would be most valuable to a user, a good understanding of the next place(s) that user might visit is important. However, as our results show, groups transit far less than individuals. Accordingly, any group/bulk promotions or advertisements that do no account for this smaller probability may be treated as spam by members in a group. Also, as groups tend to stay much longer at a place before they move to another place, the right time to send promotions may be different for groups and individuals – this is especially true if an advertiser considers the time when people are about to move to a new place as the best time to send a promotion.

INTERRUPTIBILITY

Smartphone users are constantly subjected to information overload, as the applications on the device, all compete for the user’s attention. Many studies in the past have investigated delaying (or advancing) notification delivery [14, 13, 22] by understanding the intrinsic behaviour of users to explain the




Figure 7. SMS thread lengths at *Food Court* (Hypothesis 2B)

user’s willingness to engage with his/her smartphone. For example, depending on their context, users may feel *interrupted* when they receive a call or an incoming notification. This could be understood through observations such as whether the user engages in the call for shorter or longer durations. Instances where notifications arrive and the user does not respond may indicate the user’s unwillingness to respond.

To this end, we look at users’ call, SMS and app usage logs to understand the effect of group and location on interruptibility. Note that the results presented in this section pertain to data collected from 156 participants for whom fine-grained smartphone usage data is available in addition to their locations.

Call Logs

HYPOTHESIS 2A: Groups Have Shorter Phone Conversations.

We first investigate the above hypothesis on call durations. Figure 6 shows the CDF of incoming and outgoing call durations for all 4 partitions at a *Food Court*. It shows that when receiving incoming calls at a *Food Court*, 80% of group participants answered their phones for less than a minute (regardless of group sizes) whereas *Solo* spent significantly longer times on phone conversations; the average call duration (in minutes) for *Solo* was 72.73 seconds with a much longer tail while the average duration for *Large* was 38.03 seconds (with the p -value=0.0009 on the KS test). This indicates that physical groups influence how long users were willing to be interrupted to engage in remote conversations.

We observed similar patterns for outgoing calls. *Solo* was more willing to actively call other people than groups, indicating that groups are less willing to be interrupted.

SMS Logs

HYPOTHESIS 2B: Groups Have Shorter SMS Conversation Lengths

We next pose a hypothesis based on SMS thread lengths – we hypothesise that groups will engage in shorter SMS conversations. Figure 7 shows the CDF of SMS thread lengths (the number of messages exchanged with the same recipient) at the *Food Court*. From the graph, we observe clear differences between *Solo* and groups (of all sizes). More than 50% of *Solo* engaged in long conversations (of threads longer than




Figure 8. Response time to notifications (Hypothesis 2C)

40 messages) whereas no group engaged in such long conversations – in fact, we observe that individuals in the largest of groups tend to converse using messages the least.

Response Time to Notifications


HYPOTHESIS 2C: *Groups Respond Quicker to Smartphone Notifications*

Finally, we investigate response times to smartphone notifications as another proxy to the user’s interruptibility. We first observed how frequently users interacted with their phones as a response to incoming notifications. We observed, that 14.76% of the app sessions were *reactive* where an incoming notification or event caused the user to use the phone in response. Also, we noticed that of the 14.76% that responded to notification, we observed, across all four partitions, that users did not respond immediately, 50% of the time; to be specific 55.32%, 54.86%, 51.80% and 49.65% of *Solo*, *Small*, *Medium*, and *Large*, respectively. Here, we consider the response to be immediate if the user starts interacting with the phone before the phone goes back to sleep again (i.e., auto screen timeout).

Next, for sessions initiated by a notification, we observed the time it took for the user to respond, per partition. The response time is calculated as the difference in time between the time the notification was received and the time at which the user touches the phone. Figure 8 shows the CDF of the response times. Although we find the difference in response times between *Medium* or *Large* and *Solo* to be statistically significant (p-value of 0.001023 and 6.619e-06, respectively), the absolute differences were in the range of <2 seconds.

Summary and Use Cases

Our observations suggest interesting differences between solo vs groups in terms of interruptibility; we find that groups engage in shorter calls (both incoming and outgoing), and are less likely to initiate outgoing calls. Similarly, we find that groups tend to converse shorter through SMS with remote friends. This could be viewed as evidence of individuals conforming to social norms (where it is generally considered rude to be communicating on the phone while in the physical presence of others). Interestingly, the response time to notifications do not change much even when individuals are in groups



The figure is in log scale for visual clarity. The burstiness of *Medium* is shifted towards the right and is significantly different from *Solo* (p-value = 4.363e-14).

Figure 9. Burstiness factor (Hypothesis 3A)

– we suspect response time is decided based more on the content and urgency of the notifications regardless of presence of physical groups. We describe a relevant use case below.

Auto-configuration of Call and Notification Settings. Receiving a flood of notifications or untimely calls/SMSs can make mobile users very conscious about their call/notification settings. In particular, they would like to ensure that they are not interrupted by random messages at undesirable times (for example, when in the midst of a social group). There has been prior work that has proposed to automatically identify the right moments to deliver notifications (when users can be interrupted) [13, 22]. Complementary to such work, our findings suggest that understanding whether the user is by themselves or in a group can help make this interruptibility metric more robust and accurate.

APPLICATION USAGE

In this final data analysis section, we check if being in groups affects the types and durations of mobile applications used.

App Use Time and Frequency

We know that social norms can govern the behaviour of individuals in social groups. For example, in some cultures, when seated for a meal together, it is considered impolite to engage in other activities such as reading books, or listening to music, simultaneously. As a result, we hypothesise that in certain settings, individuals may change their smartphone app usage behaviour when in a group.

HYPOTHESIS 3A: *Groups Use Their Phones Less.*

We view app usage as two-fold: (1) the duration of app usage, and (2) the frequency of app usage. We use three metrics: (1) *session_duration*, (2) *session_intervals*, and (3) *burstiness factor* to understand app usage. An *app session* is a continuous segment of time for which a user interacts with his phone, and a *session_duration* is defined as the duration of each app session. *session_intervals* is defined as the period for which the phone is in sleep mode, or the time period between two consecutive app sessions. Finally, we define the **burstiness factor** as, $B_i = 1/(t_i * d_i)$ where t_i denotes the session duration for an app session i , and d_i is the previous session interval – i.e., the interval between the end of app session $i - 1$ and the start of the app session i . Checking the phone more frequently




Figure 10. Session interval vs. Session duration (Hypothesis 3A)

but for shorter durations suggests *bursty* usage; a burstiness factor close to 1 indicates *highly bursty* usage.

Figure 9 shows the PDFs of the burstiness factors for the 4 partitions. We see that *Solo*, *Small*, and *Large* have almost identical PDFs while in the case of *Medium*, users use apps more frequently, but for shorter durations. For the remaining three partitions, sessions may or may not be bursty. The PDF for *Medium* is shifted more to the right, i.e., closer to zero, which is equivalent to a burstiness factor of 1 in the normal scale. Additionally, Figure 10 shows the scatter plot of session durations (y-axis) versus session intervals (x-axis). We observe that the dispersion along either axis is much larger for *Solo* while the sessions are concentrated mostly towards the origin for *Medium*. The KS test proves that two distributions are statistically different with p-value 0.0005689. Hence, we reject the hypothesis for *Medium*.

SNS Usage

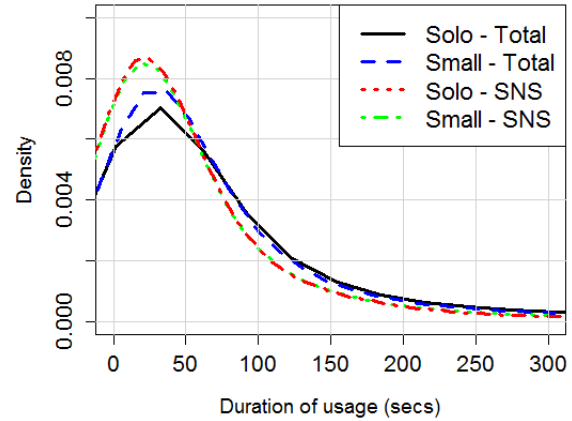
HYPOTHESIS 3B: *Groups Affect SNS Usage*

Social Network Services (SNS) apps are some of the most popular apps used today. As such, we decided to observe the duration of SNS app use per app session. We considered the fourteen most popular SNS apps amongst our users; the list included social networks such as Facebook and Twitter, OTT messaging services such as WhatsApp, and location based SNS such as Foursquare and Email. We did not consider calling and SMS here, as these were discussed separately.

Figure 11 shows both the SNS use durations and the total SNS app session durations. For the app usage duration, we observed a significant difference between *Solo* and all three group partitions; for example, p-values on the KS test between *Solo* and the *Small* and *Large* partitions was 0.006 and 0.034, respectively. On the other hand, we found that there was no significant difference in the use of SNS apps between *Solo* and *Small* (p value=0.1053). Overall, we find that only *Large* affects SNS app usage significantly.

Summary and Use Cases

Interestingly, we find that differences in app usage are limited to certain partitions. In particular, we believe that the type of group affects behaviour changes differently. For example, when individuals are in *Small* groups (mostly close friends), it seems their app usage behaviour including SNS app usage hardly change due to the level of comfort around each other. In the case of *Medium*, on the other hand, where



The SNS usage duration was not significantly different between the two cases (p value 0.1053)

Figure 11. Total and SNS usage duration per Session for *Solo* vs. *Small* (Hypothesis 3B)

the individuals are meeting for class projects and serious discussions, individuals do not spend too much time using their phones; instead, they tend to check their phones often, resulting in bursty usage. We below illustrate a use case based on this observation.

Adaptive energy management: More frequent app usage can cause higher energy drains as each time the device wakes up, radios and other chips need to be turned ON. In addition, if the device is woken up frequently with short user sessions, the ratio of utility to energy drain becomes worse for the user. Hence, it would be useful to model a user's usage patterns to synchronise wake up times with times the user would like to use an app. As we show, these wake up times can be different between *Solo* and *Medium* (for example) with *Medium* exhibited much more bursty app use behaviour. This type of app usage information can also be used to schedule good times to conduct periodic cloud synchronisation operations that can piggyback on top of other app usage.

DETERMINING THE ROBUSTNESS OF OUR RESULTS

Finally, we present precise statistics and one graph to show that our results are robust even with the underlying errors in our group detection algorithm.

What Effect Do Errors in Group Detection Labelling Have?

Because our group detection algorithm only has a precision of $\approx 90\%$ and a recall of $\approx 80\%$, there will be errors in the group detection labelling process. This means that a fraction of group data might actually contain solo data, and vice versa. To determine the robustness of our results even with these errors, we injected additional random errors into the output labels and tested if the KS test still outputs the same conclusions between the *noise-added group data* and the *noise-added solo data* distributions. More specifically, we introduced errors by flipping a designated percentage of labels (chosen randomly using 1%, 5%, 10%, 15%, and 20% data flip levels) from groups to solos or the other way. We chose this error range to match the various recall and precision levels of our group detector *GruMon*.

percentage of errors	D-statistic	p-value
Original	0.1295	<2.2e-16
1%	0.1268	<2.2e-16
5%	0.1186	<2.2e-16
10%	0.1076	<2.2e-16
15%	0.1034	<2.2e-16
20%	0.0963	<2.2e-16

Table 3. KS results bet. the noise-added group and solo data

percentage of errors	D-statistic	p-value
1%	2e-04	1
5%	1e-04	1
10%	0	1
15%	6e-04	1
20%	4e-04	1

Table 4. KS results bet. the original group and noise-added group data

Table 3 shows the p-values and D-statistic values from the KS test between the original group dataset and solo dataset, as well as the noise-added group datasets and noise-added solo datasets with different error percentages, for Hypothesis 1A. The table shows that the KS test still accepts the hypothesis even with 20% errors introduced; p-values for all the cases are much lower than 0.05, the well-known threshold of p-value for acceptance. Also, there is no meaningful difference across D-statistic values, which means the distances between the two distributions do not change noticeably with different levels of injected errors.

Does Noise Cause Deviations From the Original Dataset?

We further investigated if the noise-added group data or solo data kept their properties when compared to the original group data or original solo data, respectively. For this, we conducted the KS tests between the original group dataset and the noise-added group datasets (for all noise % error levels). Table 4 shows the KS-test results for stay time distributions at the *GSR* (Hypothesis 1A) and that the noise-added datasets do not show significant differences from the original dataset. Even for the comparison to the data with 20% errors, the p-value was 1 (in this case p-values > 0.05 indicate that the distributions are similar), which is a clear indicator that the two distributions are the same. Similarly, we compared the noise-added solo datasets with the original solo dataset, and found out that they are also not significantly different; all the p-values were above 0.05. Note that we have done these sensitivity analysis for all the other hypotheses, and presented only the results that are consistently accepted or rejected.

Probability Density Distributions of Noise-added Datasets

Finally, we also visually compared the differences among the original datasets and the noise-added datasets. Figure 12 plots the probability density distributions of the SMS thread lengths for the original solo data, original group data, and error-added solo datasets. The figure shows that the original solo data and all the noise-added solo datasets have similar distributions whereas the original group data shows noticeable difference. Note that all the noise-added group datasets (which we omitted from the Figure for clarity reasons) show similar distributions with the original group dataset. We also conducted KS tests between the different pairs of datasets and confirmed that our visual observation matches the results of the KS tests.




Figure 12. Probability density distribution of noise-added datasets

CONCLUSION AND FUTURE WORK

In this paper, we showed that individuals' behaviors can be *significantly different* when they are in *groups*, compared to when they are alone. This understanding is essential in characterising and predicting user behaviour since ignoring the interaction of individuals in groups can lead to erroneous results. We studied a unique large-scale, longitudinal dataset, collected over a period of 4 months from our university campus, containing both indoor location traces as well as fine-grained app usage data. We applied a state-of-the-art group detection algorithm to label the dataset with group vs. individual context, and then showed that individuals vs. groups exhibit significant differences across three behavioural properties: (1) dwell times and next place transitions (mobility patterns), (2) interruptibility and (3) application usage. In addition, we accounted for errors in the underlying group detection technique and established the robustness of our analyses.

The results presented in this paper are just the beginning toward accurate understanding of user behaviour by taking physical groups into account. In the future, we plan to extend our work in several interesting directions including; (1) segregate groups not just by size but also by social interaction strengths (using other sources of data such as social media, etc.) – for example, we can classify groups as composed of “near strangers” versus “frequent interacters”. (2) compare the group versus individual differences at an individual level (beyond an aggregate level as we did in this paper). (3) expand our study to other locations and user groups beyond a campus environment and student population. In particular, we plan to analyse other environments we have fine-grained location traces for such as a commercial airport, a resort island, and a convention centre.

ACKNOWLEDGMENTS

This material is based on research sponsored by the Singapore National Research Foundation under its International Research Centre @ Singapore Funding Initiative and administered by the IDM Programme Office, Media Development Authority (MDA). The authors gratefully acknowledge the support of the LiveLabs research centre for the indoor mobility and App usage datasets.

REFERENCES

1. Bahl, P., and Padmanabhan, V. RADAR: An in-building RF-based user location and tracking system. In *Proc. of INFOCOM* (Anchorage, Alaska, May 2000).
2. Bai, F., Sadagopan, N., and Helmy, A. Important: a framework to systematically analyze the impact of mobility on performance of routing protocols for adhoc networks. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, vol. 2 (March 2003), 825–835 vol.2.
3. Balan, R. K., Misra, A., and Lee, Y. LiveLabs: Building an in-situ real-time mobile experimentation testbed. In *Proc. of HotMobile* (Santa Barbara, CA, 2014).
4. Becker, R., Cáceres, R., Hanson, K., Isaacman, S., Loh, J. M., Martonosi, M., Rowland, J., Urbanek, S., Varshavsky, A., and Volinsky, C. Human mobility characterization from cellular network data. *Commun. ACM* 56, 1 (Jan. 2013), 74–82.
5. Bhattacharjee, D., Rao, A., C, S., Shah, M., and Helmy, A. Empirical modeling of campus-wide pedestrian mobility observations on the usc campus. In *Vehicular Technology Conference, 2004. VTC2004-Fall. 2004 IEEE 60th*, vol. 4 (Sept 2004), 2887–2891 Vol. 4.
6. Böhmer, M., Hecht, B., Schöning, J., Krüger, A., and Bauer, G. 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*, ACM (New York, NY, USA, 2011), 47–56.
7. Brown, C., Lathia, N., Mascolo, C., Noulas, A., and Blondel, V. Group colocation behavior in technological social networks. In *PLoS ONE* (2014).
8. De Melo, P. O. S. V., Akoglu, L., Faloutsos, C., and Loureiro, A. A. F. Surprising patterns for the call duration distribution of mobile phone users. In *Proceedings of the 2010 European Conference on Machine Learning and Knowledge Discovery in Databases: Part III, ECML PKDD'10* (2010).
9. Do, T. M. T., Blom, J., and Gatica-Perez, D. Smartphone usage in the wild: A large-scale analysis of applications and context. In *Proceedings of the 13th International Conference on Multimodal Interfaces, ICMI '11*, ACM (New York, NY, USA, 2011), 353–360.
10. Do, T. M. T., Kalimeri, K., Lepri, B., Pianesi, F., and Gatica-Perez, D. Inferring social activities with mobile sensor networks. In *Proceedings of the 15th ACM on International Conference on Multimodal Interaction, ICMI '13* (2013).
11. Eagle, N., and (Sandy) Pentland, A. Reality mining: Sensing complex social systems. *Personal Ubiquitous Comput.* 10, 4 (Mar. 2006), 255–268.
12. Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., and Estrin, D. Diversity in smartphone usage. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, MobiSys '10*, ACM (New York, NY, USA, 2010), 179–194.
13. Fisher, R., and Simmons, R. Smartphone interruptibility using density-weighted uncertainty sampling with reinforcement learning. In *Proceedings of the 2011 10th International Conference on Machine Learning and Applications and Workshops - Volume 01, ICMLA '11*, IEEE Computer Society (Washington, DC, USA, 2011), 436–441.
14. Hudson, S., Fogarty, J., Atkeson, C., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J., and Yang, J. Predicting human interruptibility with sensors: A wizard of oz feasibility study. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '03*, ACM (New York, NY, USA, 2003), 257–264.
15. Laurila, J. K., Gatica-Perez, D., Aad, I., J., B., Bornet, O., Do, T.-M.-T., Dousse, O., Eberle, J., and Miettinen, M. The Mobile Data Challenge: Big Data for Mobile Computing Research. In *Pervasive Computing* (2012).
16. Lee, Y., Ju, Y., Min, C., Kang, S., Hwang, I., and Song, J. Comon: Cooperative ambience monitoring platform with continuity and benefit awareness. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services, MobiSys '12*, ACM (New York, NY, USA, 2012), 43–56.
17. Madan, A., Cebrian, M., Moturu, S., Farrahi, K., and Pentland, A. Sensing the health state of a community. *Pervasive Computing, IEEE* 11, 4 (Oct 2012), 36–45.
18. Manweiler, J., Santhapuri, N., Choudhury, R. R., and Nelakuditi, S. Predicting length of stay at wifi hotspots. In *INFOCOM, 2013 Proceedings IEEE*, IEEE (2013).
19. Miluzzo, E., Lane, N. D., Fodor, K., Peterson, R., Lu, H., Musolesi, M., Eisenman, S. B., Zheng, X., and Campbell, A. T. Sensing meets mobile social networks: The design, implementation and evaluation of the cenceme application. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems, SenSys '08*, ACM (New York, NY, USA, 2008), 337–350.
20. Musolesi, M., and Mascolo, C. A community based mobility model for ad hoc network research. In *Proceedings of the 2Nd International Workshop on Multi-hop Ad Hoc Networks: From Theory to Reality, REALMAN '06*, ACM (New York, NY, USA, 2006), 31–38.
21. Nishida, K., Toda, H., Kurashima, T., and Suhara, Y. Probabilistic identification of visited point-of-interest for personalized automatic check-in. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (ACM Ubicomp'14)* (Seattle, WA, USA, September 2014).

22. Pejovic, V., and Musolesi, M. Interruptme: Designing intelligent prompting mechanisms for pervasive applications. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '14*, ACM (New York, NY, USA, 2014), 897–908.
23. Rhee, I., Shin, M., Hong, S., Lee, K., and Chong, S. On the levy-walk nature of human mobility. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE* (April 2008).
24. Sen, R., Lee, Y., Jayarajah, K., Misra, A., and Balan, R. K. Grumon: Fast and accurate group monitoring for heterogeneous urban spaces. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems, SenSys '14*, ACM (New York, NY, USA, 2014), 46–60.
25. Shin, C., Hong, J.-H., and Dey, A. K. Understanding and prediction of mobile application usage for smart phones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing, UbiComp '12*, ACM (New York, NY, USA, 2012), 173–182.
26. Virtual Laboratories in Probability and Statistics, University of Alabama in Huntsville. The log-logistic distribution.
<http://www.math.uah.edu/stat/index.html>.
27. Virtual Laboratories in Probability and Statistics, University of Alabama in Huntsville. The weibull distribution. <http://www.math.uah.edu/stat/special/Weibull.html>.
28. Wagner, D. T., Rice, A., and Beresford, A. R. Device analyzer: Large-scale mobile data collection. *SIGMETRICS Perform. Eval. Rev.* 41, 4 (Apr. 2014).
29. Wang, R., Chen, F., Chen, Z., Li, T., Harari, G., Tignor, S., Zhou, X., Ben-Zeev, D., and Campbell, A. T. Studentlife: Assessing mental health, academic performance and behavioral trends of college students using smartphones. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '14*, ACM (New York, NY, USA, 2014), 3–14.
30. Xu, Y., Lin, M., Lu, H., Cardone, G., Lane, N., Chen, Z., Campbell, A., and Choudhury, T. Preference, context and communities: A multi-faceted approach to predicting smartphone app usage patterns. In *Proceedings of the 2013 International Symposium on Wearable Computers, ISWC '13*, ACM (New York, NY, USA, 2013), 69–76.