Connect the Dots by Understanding User Status and Transitions

Xuan Bao
Samsung Research America
75 West Plumeria
San Jose, CA, 95134

Yilin Shen
Samsung Research America
75 West Plumeria
San Jose, CA, 95134

Neil Zhenqiang Gong
Samsung Research America
75 West Plumeria
San Jose, CA, 95134

Hongxia Jin
Samsung Research America
75 West Plumeria
San Jose, CA, 95134

Bing Hu
Samsung Research America
75 West Plumeria
San Jose, CA, 95134

UbiComp'14, September 13–17, 2014, Zurich, Switzerland.

http://dx.doi.org/10.1145/2638728.2641704

Abstract
Human lives are composed by series of events and activities. Considerable research effort has been made to probe, sense, and understand them. In our research, we are interested in exploring the intrinsic string that connects all these events together, that is, user status and transitions. Such transitions can be reflected from multiple activity dimensions, ranging from our daily mobility trajectories, app usage sequences, to communication patterns and motion state switches. In this paper, we aim to identify whether a personalized model can be learned to capture various user states from different sensing dimensions and whether a unified view can be established to explain the state transitions that drive the changes in user context during day-to-day routines.

To this end, we have explored two types of traces – connected wifi sequences and cell location trajectories. We first model the states among these two individual dimensions. In the end, the identified states from both dimensions are linked together to recognize the spatial-temporal relationship between them. As we evaluate with the DeviceAnalyzer dataset, our method is able to recognize a range of states such as "at home", "working", "commute" and the trasitions between them, all in an unsupervised manner.

Introduction

Mobile devices have become an integrated and sometimes even the central part of our daily routines. People bring their devices to all types of locations and events - to workplaces, to classrooms, to business meetings, to social dinners, and to private parties. Aware of this opportunity, considerable research effort has been spent to probe, sense, and understand such events using observations obtained from mobile devices [3, 5]. In our research, we are interested in exploring the intrinsic string that connects all these events together, that is, human status and transitions. We believe that many of the activities people take can be explained using the current status of user as well as the transitions between the states.

Take daily routine activities as one example. Typical users maintain a balance between their work and private life. At workplaces, the main activities are centered around professional matters - executing office work, attending meetings, checking corporate emails, etc. All these activities are associated with the user's "working" state. After work, the user may go to other places to enjoy lunch, have a coffee break, or meet with friends – each time the individual events can be viewed as related to different leisure states. Then, the user may finally head home and enjoy various family activities – potentially relate to the "family life" state. Therefore, such "work", "family life", and the different "leisure" states are essentially different modes that a user exhibits and can be used to summarize and predict related activities. Moreover, the transitions between these states are also valuable. By analyzing a user's state transition pattern, a system may identify that whenever she is leaving work on Friday evening, she is more likely to go to a few favorite dinner places. In other words, transitions can quantify what the user tends to do next given the current state.

Besides daily routines, such states and transitions can be observed from many other activity dimensions, ranging from our mobility trajectories, app usage sequences, to communication patterns and motion state switches. In our work, we attempt to identify some of the most common user states from multiple information sources. Moreover, we aim to identify whether a personalized model can capture the states that drive user actions among all these different dimensions and more importantly whether a unified view can be established to explain all different dimensions together without having to rely on individual models.

Realizing this vision entails many challenges. First of all, the service needs to be light-weight to be running on the phone. This means using readily available data without additional retrieving cost and designing efficient algorithms with affordable execution time. With these considerations, we choose to use connected WiFi traces and cell location information as the start point. For any smartphone today, being a high-end flagship phone or a low-end one, the network functionalities must be built in for WiFi and cell tower connection. Therefore, these two data sources are already widely available with little overhead. Second, the service needs to identify semantic meanings without requiring human labels. In the real world, asking for manual labeling upfront demands much human effort and often annoys the user. Our design aims to provide most functionalities immediately and offer the opportunity to refine with human assistance later. Finally, due to privacy concerns, the service needs to function with anonymized data – for example, hashed cell IDs and WiFi names. Mobile users today are becoming more privacy-aware. Showing that the service respects their privacy in treating their data can help boost the acceptance of the service.

With these design requirements in mind, we present a series of methods to process and analyze connected WiFi traces and cell location sequences. In the end, we are able to identify different user states such as "at home", "working", "commute"and other user-specific significant states like "lunch" or "dining". A Hidden Markov Model is then used to model the transitions for each user and we show that various insights can be extracted.

Associate Semantics with Connected WiFi

The first step starts with analyzing connected WiFi traces. In the analyzed data set [1], users on average only connect with 10 WiFi access points (AP) during the entire study. This is because, though WiFi APs are available everywhere, most of them are password-protected from the public. Connecting to them requires the users' effort of inputing password and cautiously switching on/off WiFi to save energy. As a result, the APs that a user connects to usually represent meaningful places in her life - home, work, friends' place, favorite diners, etc. Our algorithm attempts to extract these information.

Using a similar approach as described in [1], we start by partitioning 24 hours of a day into eight uneven time windows as shown in Figure 1. These windows are designed to capture natural user behaviors such as sleeping, commuting, working, and dining. Later we show that a dynamic partition can be used to further optimize the performance.

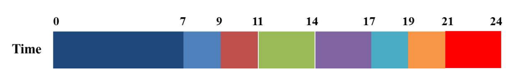

Figure 1: Initial time partition.

Then, our algorithm counts the number of appearances of each WiFi AP in each time window across the entire trace. The result is a $N \times 8$ matrix, where N is the number of APs. Then, a dual representation of the appearances is computed by normalizing each column (Time window (TW) view) and each row of the matrix (AP view). Figure 2 illustrates the results for a sample user's 8-month trace with rows denoting unique WiFi APs and columns showing time windows. The red color means probability close to 1 and the blue color corresponds to 0.

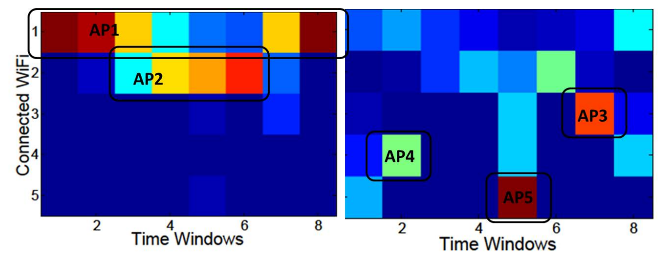

Figure 2: Use dual representations to identify (left, TW view) home, work and (right, AP view) other semantics. Blue to red denotes lower to higher probabilities.

The left TW view shows what the most frequent APs are during each time window while the right AP view shows which time window occurs the most for each AP. Interestingly, from the left figure, $AP1$ almost dominates time window 1, 2, 3, 7, and 8 - including morning and late night hours. These windows correlate with typical "home staying" behavior of most people. In contrast, $AP2$ dominates time window 4, 5, 6 and also shows up at 3 - typical working hours. Therefore, we make initial guesses here that $AP1$ and $AP2$ correspond to home and work AP respectively. For other APs, the right figure shows insights about what they might be. $AP3$, $AP4$ and $AP5$ are most correlated with window 7, 2, and 5 - common dining hours for dinner, breakfast, and lunch. Actually, $AP4$ appears across three time slots and could be a dining place that the user goes for different meals.

Notice that, user behaviors may differ – the initial time partition may not be applicable to everyone. The next

[1]Due to limited computation power, 300 users with 30MB to 300MB data are randomly selected for processing.

task is to dynamically change the time window boundaries to maximize the concentration of APs' appearances in each matrix in Figure 2 – e.g., make only $AP1$ appear during home hours. To measure such concentrations, we use a metric C similar to total variation – defined as

$$C(X) = \sum_i (x_{i+1} - x_i)^2 \qquad (1)$$

for each sorted probability mass function (PMF) for the columns in TW view and rows in AP view respectively. The metric measures whether one or a few elements in a PMF collectively contains most of the probability mass and equals 1 when one time window captures all the appearances. Using this metric, we can test different time partition schemes corresponding to shifting time window boudaries earlier and later. For the same sample user, Figure 3 shows C across 11 different settings. Setting 10 (late dinner) is a clear winner for $AP3$ and $AP5$ as most of their appearances are captured by one or few time windows.

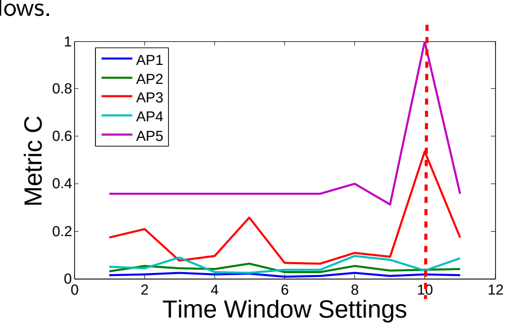

Figure 3: Concentration metric under different time settings.

After identifying optimal time partitions and associating semantic meanings with APs, we use a 3-state hidden markov model to represent the observed daily sequences. Unsurprisingly, Baum-Welch algorithm returns with $AP1$ and $AP2$ associating with 2 different states and other APs with the 3rd state. The transitions show that with $1/23$ chances the user goes somewhere else after work rather than going home directly and rarely returns to work after.

While connected WiFi is highly correlated with significant places, unfortunately, the user is connected only for a very limited period of time. On average, a user is connected to WiFi only 10.6% of the time. Therefore, relying on connected WiFi alone results in sparse and interruptive observations. To address this issue, we use cell sequences to fill in the gaps.

Fill in the Gap with Cell Locations

On contrary to connected WiFi, smartphones are almost always connected to cell towers for communication purposes. In analyzed traces, average users are connected to cell towers three times more than WiFi. We expect this is even an underestimate due to turning off phones and stopping DataAnalyzer. However, processing cell sequences is more challenging for several reasons. First of all, a cell sector can cover a large area so that a moved user may still stay within the same sector. Moreover, due to signal strength fluctuations, a user may observe cell tower handovers without physical movement. As a result, cell sector changes do not necessarily coincide with location changes. Compounding this problem, the number of unique cell IDs appearing in a trace (on avg. 1047) is significantly more than WiFi. Their correlation with time windows is also too weak for the most frequent cell sectors. Figure 4 illustrates this issue for the sample user. The top figure shows that more than two thousand cell IDs have appeared during the eight months and their appearances follow power-law distribution. However, the top cell sectors' appearances only weakly correlate with time windows as illustrated in the bottom figure. Most of the top 20 cells' apperance scatter across time windows, making the same treatment of WiFi sequence unfeasible.

Our solution is to propagate information learnt from WiFi to cell. Connecting WiFi and Cell traces together, our algorithm identifies the correlation between WiFi APs and cell IDs by analyzing their co-occurance pattern. Specifically, we are interested in whether a cell ID mainly appears with certain WiFi APs for which the semantic meaning is learnt previously. Interestingly, after normalization and processing, some of the most frequent cells show such strong correlation with WiFi as illustrated in Figure 5 for the sample user. Each column in the figure denotes the probability a cell appear at the same time with a WiFi AP among its appearance with all WiFi APs. 1 (red) means whenever the cell appears with an AP, the AP is the same one. Therefore, the high-value grids mean that some cells are found to be strongly related to certain APs so that we can associate the states (e.g., at home, working, etc.) found from APs to these cells.

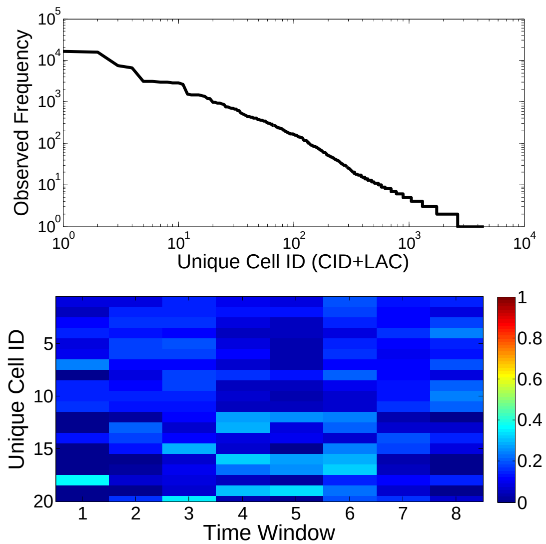

Figure 4: Cell appearances follow power-law distribution and only weakly correlate with time windows.

Another interesting property of cell sequences is its correlation with commute patterns [2]. Figure 6 shows the CDF of observed handovers during each half-hour in the sample user's entire trace. Intuitively, more observed handover may correspond to higher mobility.

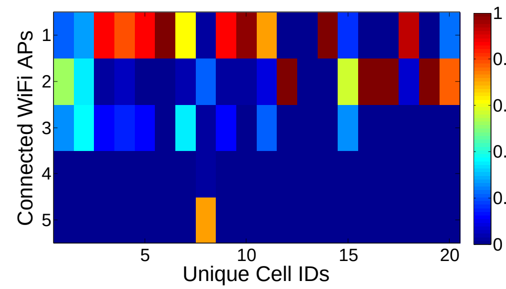

Figure 5: Co-occurance between WiFi and top cells.

Therefore, our algorithm selects a percentile as commute threshold using two criterions. First of all, it minimizes the detected commute during 0 to 5 AM since these are not typical commute hours. Second, it fits the data into a 2-state Gaussian Mixture Model, assuming typical users exhibit two home-to-work commutes per weekday, and attempts to identify a good fit. The identified number of commutes is shown in Figure 7. The false positives during 0 to 5 AM is suppressed and GMM identifies the two commute peaks (μ) at 10:20 AM and 7:40 PM respectively. In the future, our algorithm needs to accomodate other user types such as those living close to office or working from home.

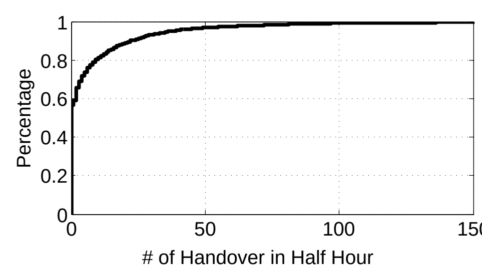

Figure 6: CDF of handovers in half an hour.

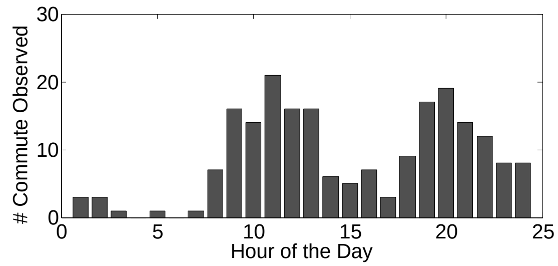

Figure 7: Commutes in the trace at different times.

Piecing all these clues together, Figure 8 shows the aggregated state information identified for the sample user. The x-axis shows the hour of the day and the y-axis counts how likely the user is observed at certain state. For the 265-day trace of this user, 530 observations are made for each hour (two per hour) of the day. The color coded states (from bottom up) corresponds to commute, no network record, home, work and others, respectively. Her working hour can be from 9 AM to mid-night and she is mostly at home during the evenings and in the morning.

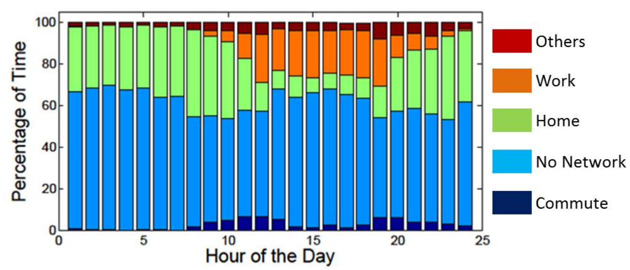

Figure 8: Identified states across hour of a day.

Using these identified states, we are able to compute the transition probability across all users in the 300-user dataset. On average, even on weekdays, users choose to go somewhere else after work with a probability of 56.8%, slightly more frequently than going home directly (43.2%). After spending time around, unsurprisingly, they choose to go home most of the time with a 92.7% probability but do continue to work from time to time (7.3%). Moreover, typical users commute twice a day around 8:24 AM and 6:10 PM (average) and a user's commute patterns may typically vary within a 1.5 to 2-hour range. This large variation in day-to-day behavior illustrates the advantage of sensing user states over using a fixed daily schedule all the time. Interestingly, the results are also consistent with U.S. census results [4], showing potentially common behavior patterns across large populations even within different regions.

Conclusion

Mobile phones are transforming to converged platforms to sense the human life. The observations are proven to be invaluable to various industries such as healthcare, city planning, customized search, targeted advertisement, and personal assistance. This paper explores identifying semantically meaningful states and transitions by combining readily available WiFi and cell tower traces. The algorithm mainly draws from common sense knowledge and relies on analyzing distributions in an unsupervised manner, thus does not require human labeling effort. In the future, we plan to expand the exploration into even more sensing dimensions, model behavior patterns as topics, and build the service into next generation of mobile platforms.

References

[1] Farrahi, K., and Gatica-Perez, D. What did you do today?: discovering daily routines from large-scale mobile data. In *ACM SIGMM* (2008).

[2] Isaacman, S., Becker, R., Caceres, R., Kobourov, S., Rowland, J., and Varshavsky, A. A tale of two cities. In *ACM HotMobile* (2010).

[3] Krause, A., Horvitz, E., Kansal, A., and Zhao, F. Toward community sensing. In *IEEE IPSN* (2008).

[4] Rapino, M., and Others. How can we best visualize worker movement throughout the day? In *U.S. Census* (2014).

[5] Wagner, D. T., Rice, A., and Beresford, A. R. Device analyzer: Understanding smartphone usage. In *ACM MobiQuitous* (2013).