# Thinking Fast and Slow: An Approach to Energy-Efficient Human Activity Recognition on Mobile Devices

*Yifei Jiang, Du Li, Qin Lv*

■ *According to Daniel Kahneman, there are two systems that drive the human decision-making process: The intuitive system that performs the fast thinking, and the deliberative system that does more logical and slower thinking. Inspired by this model, we propose a framework for implementing human activity recognition on mobile devices. In this area, the mobile app is usually always on and the general challenge is how to balance accuracy and energy consumption. However, among existing approaches, those based on cellular IDs consume little power but are less accurate; those based on GPS/Wi-Fi sampling are accurate often at the costs of battery drainage; moreover, previous methods in general do not improve over time. To address these challenges, our framework consists of two modes: In the deliberation mode, the system learns cell ID patterns that are trained by existing GPS-/Wi-Fi-based methods; in the intuition mode, only the learned cell ID patterns are used for activity recognition, which is both accurate and energy efficient; system parameters are learned to control the transition from deliberation to intuition, when sufficient confidence is gained, and the transition from intuition to deliberation, when more training is needed. For the scope of this paper, we first elaborate our framework in a subproblem in activity recognition, trip detection, which recognizes significant places and trips between them. For evaluation, we collected real-life traces of six participants over five months. Our experiments demonstrated consistent results across different participants in terms of accuracy and energy efficiency and, more importantly, its fast improvement on energy efficiency over time due to regularities in human daily activities.*

In his Nobel Prize–winning work, Daniel Kahneman (2011) challenged the traditional rational model of human judgment and decision making by proposing a model of two systems: system 1 is fast, intuitive, and emotional, while system 2 is slower, more deliberative, and more logical. According to his theory, a human being sometimes makes decisions deliberately and sometimes relies on her intuition or gut feeling. On the one hand, intuition is directly proportional to the similarity of past experiences, relying on temporal and similarity relations to determine reasoning rather than an underlying mechanical structure. On the other hand, deliberation functions on logical structure and variables, basing upon rule systems to come to conclusions. In general, deliberation is slower and subject to conscious judgments, while intuition comes to mind quickly and effortlessly.

Kahneman's work has generated profound impacts on many fields, such as psychology, economics, medicine, management, and politics, that are related to how the two systems shape our judgments and decisions. In this article, we apply his model to a different type of decision making — how a mobile app can predict human activities based on sensory data collected on mobile devices. We propose to build mobile apps of this type in a framework that similarly consists of a deliberation mode and an intuition mode. As our results will reveal, the framework works well where human activities manifest regularities. We present the framework and evaluate its performance in this article.

  ISSN 0738-4602

## Human Activity Recognition

Human activity recognition has attracted considerable research attention in recent years, especially with the proliferation of sensor-rich smartphones. At a lower level, it aims to recognize the motion states of a user, for example, sitting, walking, running, biking, or driving. At a higher level, it aims to predict a user's locations and even plans, goals, and intents. A mobile app that performs human activity recognition usually runs in the background continuously to collect sensory data and make inferences, which often drains the battery if it is not designed in an energy-efficient manner.

For the scope of this article, we focus on a subproblem of activity recognition, trip detection, that aims to recognize significant places, where a person stays for long time durations or visits frequently, and the trips between those places. This problem is important in its own right as making trips is a common and regular activity in our daily life. According to Hu and Reuscher (2004), on average a person in the United States makes 4.09 trips every day. Interestingly, most trips are regular in that they are repeated between only a few highly frequented places, for example, the daily commute between home and work.

The proliferation of mobile devices such as cellphones makes it possible to understand these trips more easily and provide information services in context (Ashbrook and Starner 2003). Those services include single-user apps, for example, personalized ubiquitous advertising (Krumm 2011), itinerary recommendation and destination prediction (Yuan et al. 2010; Yoon et al. 2010; Krumm and Horvitz 2007), transportation mode (Patterson et al. 2003), user tracking for automatic travel diaries and emergency calls, personalized audiovisual narrations to museum visitors, contextualized reminders, and precaching of data. They also include collaborative apps, for example, road condition and traffic monitoring (Mohan, Padmanabhan, and Ramjee 2008), activity awareness (Bales, Li, and Griwsold 2011), and mining correlated behavior from multiple users' GPS data (Zheng et al. 2009; Zheng et al. 2008; Cao, Cong, and Jensen 2010).

Most of the above services require a fundamental function, trip detection, to recognize when the user is arriving at or departing a significant place. This information is useful for an application to determine when to start and stop its service. For example, the app can recommend an interesting product or coupon when the user is approaching a store and remove the recommendation or recommend something else when the user is leaving the store vicinity.

Trip detection is a challenging task for the following three reasons: First, trip detection must be automated. We cannot leave it to the user or require too much user input as it would be annoying and error prone. Second, it must be energy efficient. Since the user may start or end a trip anytime and anywhere, the function will be running backstage all the time. Third, the detection must be accurate and timely for the services to be useful to the user.

## Trip Detection Framework

In this article, we propose a framework for automatic, accurate, timely, and energy-efficient trip detection on mobile devices. We focus on the detection of when a trip starts and ends with regard to significant places such as home and work. We synergistically combine the merits of GPS-/Wi-Fi-based methods, which are more accurate yet energy consuming, and cell-ID-based methods, which are energy efficient yet less accurate. Compared with previous GPS-/Wi-Fi-based works, for example, Kang et al. (2005) and Kim et al. (2009, 2010), our approach exhibits a clear trend of decreasing energy consumption as a user visits the same places and repeats the same trips.

Our framework is inspired from Kahneman's theory on human decision making. More specifically, it consists of the following two distinct modes, which resemble in spirit the two systems in Kahneman's model, respectively.

In the deliberation mode, we learn cell ID patterns from trace data collected by GPS, Wi-Fi, and cell ID sensors and associate those patterns with places and trips learned using existing GPS-/Wi-Fi-based localization methods. The accuracy and energy expense will be at the same level as previous approaches that are based on GPS and Wi-Fi.

In the intuition mode, we use collected cell ID data and the learned cell ID patterns for determining whether a user is at a place, entering a place, or departing a place. When it is a repeated place or trip, the energy expense will be near zero because only cell ID data is collected. Meanwhile, because the cell ID patterns are trained in the deliberation mode, the accuracy will be comparable to previous GPS/Wi-Fi approaches on which the training is based.

The remainder of this article is organized as follows. In the next section, we survey related work. The Design Overview section motivates and overviews the framework. The Deliberation and Intuition section elaborates the deliberation and intuition modes and the transition between them. The Experiments and Performance Evaluation section evaluates the framework with real-life traces. We conclude the article with a summary of our contributions and future directions.

# Background and Related Work

Intelligence on mobile devices is manifested in multiple ways: First, mobile devices can perceive the user and the environment by collecting sensory data. They are equipped with an ever-increasing array of embedded sensors, including accelerometer, gyroscope, compass, microphone, camera, and location (GPS) as well as radios such as cellular, Wi-Fi, Bluetooth, and near field communication (NFC). Secondly, they are computationally powerful and can learn, reason, predict, and act on user activities, contexts, and trends, for example, as shown in Liao, Fox, and Kautz (2004) and Yin, Chai, and Yang (2004). Moreover, mobile devices are connected to other devices and the Internet, which greatly improves and multiplies the level of intelligence that they can achieve in isolation, for example, as shown in Zheng et al. (2010 and Peebles et al. (2010).

Modern AI puts more emphasis on statistic reasoning and learning. The works that are to be surveyed all corroborate this trend. Not to deviate from the scope of this article, however, we focus on works that are directly related to the use of GPS, Wi-Fi. and cell ID for energy-efficient trip detection. We acknowledge the existence of works that approach energy efficiency in mobile sensing, for example, by offloading data processing to the infrastructure (Cuervo et al. 2010) or low-power microcontrollers (Lin et al. 2012). They are orthogonal to our work and could be implemented together.

## Place Learning

There are generally two approaches to place learning: geometry and fingerprint. In geometry-based approaches, geocoordinates belonging to the same meaningful places are clustered. Ashbrook and Starner (2003) consider locations where the GPS signal is lost for some time as potential place candidates, exploiting the fact that GPS reception is poor indoors and around the so-called urban canyons. Kang et al. (2005) use time-based clustering and conceptually can work with any indoor or outdoor positioning technologies as long as they produce geocoordinates.

Among fingerprint-based approaches, Hightower et al. (2005); Laasonen, Raento, and Toivonen (2004); and Yang (2009) assume that radio fingerprints of places are stable and unique. By comparing similarity between two fingerprints, we can tell whether they are close or far apart. The BeaconPrint algorithm (Hightower et al. 2005) periodically scans Wi-Fi base stations and GSM cells to form fingerprints. According to Hightower et al. (2005), the visibility of a beacon is often a better metric than the received signal strength when constructing a fingerprint, an idea inherited in later works, for example, Kim et al. (2009; 2010). Works by Laasonen, Raento, and Toivonen (2004) and Yang (2009) address mass-market devices by making a harsher assumption that nothing but cellular radio is available. Both adopt simple algorithms to scan traces of time stamped cell IDs and cluster those seen as close by temporal correlations. Compared to other more advanced methods, they provide less accuracy but consume less energy and run on any cellular device.

## Place and Trip Recognition

Geometry-based algorithms recognize places by checking whether the device's current geolocation falls into the geometric shape of any place. Approaches in Laasonen, Raento, and Toivonen (2004) and Yang (2009) work by comparing the current GSM cell ID to each of the cell ID clusters. PlaceSense (Kim et al. 2009) proposes to recognize not only whether a user is at a place but also whether she is arriving at or leaving a place, that is, place entrance and departure. In addition to revising the original BeaconPrint algorithm in Hightower et al. (2005), PlaceSense reliably detects place departure by requiring all representative beacons associated with a place to have disappeared for some time and detects place entrance by buffering scanned fingerprints for parallel entrance and departure detection.

## Energy Efficiency for Location Sensing

Early work (Liao, Fox, and Kautz 2004; Yin, Chai, and Yang 2004) on GPS-based activity recognition targets a level higher than trip detection and does not specifically address energy efficiency. Location-sensing methods, relying on GPS or Wi-Fi, can provide accurate location information but are always power hungry. The following papers address energy efficiency using different approaches.

Chen et al.(2006) propose an algorithm for power-efficient location estimation by intelligently selecting a subset of available Wi-Fi access points. Paek, Kim, and Govindan (2010) proposed a rate-adaptive approach for GPS positioning, in which the GPS is adaptively turned on based on the estimated velocity of the user; further, their approach learns locations of GPS unavailability to avoid turning on GPS in those places.

Lin et al. (2010) proposed to select different location sensors to sense locations based on accuracy requirement, energy profile of sensors, and their accuracy at different locations. Kim et al. (2009, 2010) proposed an accelerometer-based approach, which uses acceleration readings to detect a user's movement and stops location sensing when the user is stationary. Zhuang, Kim, and Singh (2010) proposed an energy-efficient location-sensing framework that leverages other lower-power location sensors and uses an accelerometer to avoid unnecessary GPS sensing; it synchronizes the location sensing requests from multiple applications and tunes sensing parameters based on battery level.

To reduce the energy consumption in positioning and trajectory tracking on mobile devices, CAPS (Paek et al. 2011) and CTrack (Thiagarajan et al. 2011) also use cell ID traces, which are much more energy efficient to collect than GPS traces. CAPS estimates the current position based on the history of cell IDs and GPS position sequences that match the current cell ID sequence. CTrack tracks a user's trajectories using a two-pass hidden Markov model (HMM) that sequences cell ID fingerprints along with accelerometer and magnetic compass data. Neither work focuses on detection of starts and ends of trips, nor do they take a deliberation-intuition approach or the like to improve the energy efficiency over time, as does our work.

Most of the above approaches use accelerometers to reduce power-consuming GPS/Wi-Fi sampling. The flip side is twofold: (1) although accelerometers consume little energy compared to GPS or Wi-Fi, they still incur extra overhead when always on, and (2) movement detected by accelerometers does not necessarily indicate that the user is in place transition; for example, she may be walking around in her office.

### Comparisons

Our work aims to exploit cell IDs for trip detection and use GPS/Wi-Fi only for the learning (deliberation) phase. Its energy efficiency improves over time as places are revisited, a trend not demonstrated in previous works to our best knowledge. It resembles Kim et al. (2009) in addressing the similar problem of detecting the beginning and end of a trip; however, Kim et al. (2009) resort to periodic GPS/Wi-Fi scans and do not improve the energy efficiency over time. As in Kim et al. (2010) and Zhuang, Kim, and Singh (2010), we use accelerometers for energy-efficient data collection in the deliberation mode; however, we only use cell IDs in the intuition mode while at familiar places, thus avoiding the overheads and accuracy problems discussed previously. Techniques discussed in Lin et al. (2010), Paek, Kim, and Govindan (2010), and Zhuang, Kim, and Singh (2010) are complementary to ours, which could be leveraged in our future work for further improvements. Our approach differs from Laasonen, Raento, and Toivonen (2004) and Yang (2009) in that we use cell IDs for recognizing the beginning/end of a trip, by computing the time-variant cell ID flipping patterns (probabilities), instead of constructing time-invariant cliques as they do.

## Design Overview

Our work is motivated by the following four observations that are related to cellular devices. First, each device is assigned to one active cell tower at any time through which its basic telephone capability functions (or none in the absence of cellular signals). In a mobile phone application, reading the active cell ID is as cheap as reading a system variable. In contrast, it incurs extra costs to read other sensory data, such as GPS, Wi-Fi, and Bluetooth, that are often used for localization. If a similar level of accuracy can be achieved as that from using those alternative sensors, it would be more energy efficient to use cell ID for trip detection.

Second, for reasons such as load balancing and signal attenuation, cell towers usually have overlapping coverage. As a result, the active cell ID of a device often oscillates between several alternatives even when the device remains still at one place. However, the patterns typically differ when the device is moving within a place, for example, a building, and when it is moving between two places that are sufficiently far apart. That is, it is possible to use cell ID patterns for trip detection.

Third, it is fast to read the active cell ID information and detect its changes. Cell ID is almost always available, even where Wi-Fi and GPS information is unavailable or noisy. By comparison, Wi-Fi is not always available and takes time to scan; GPS is noisy in the urban canyons and often takes from tens of seconds to a few minutes to warm up, especially after a period of sleep or signal loss (Kang et al. 2005).

Finally, due to their coarse granularity, cell tower IDs alone cannot be used for accurate localization. We must use cell IDs in combination with more accurate measures, such as GPS and Wi-Fi beacons, or at least first use those measures to train the cell-ID-based mechanism so as to achieve the desired level of accuracy.

From these observations, it seems possible and advantageous to use cell ID patterns for detecting places and place boundaries (that is, the beginning and end of a trip). If this is feasible with high accuracy, then we can rely on cell IDs for energy-efficient trip detection. The questions are (1) how to obtain the cell ID patterns for places and their boundaries, (2) how to use these patterns for trip detection, and (3) under what conditions we can or cannot use these patterns.

### Approach and Scope

In the spirit of Kahneman's theory, we treat our mobile app as an intelligent agent, which distinguishes deliberation and intuition modes: In the deliberation mode, it uses more energy-expensive information (for example, Wi-Fi, GPS) with heavier computation to make decisions regarding trip detection. Meanwhile, the system will build "experiences," that is, cell ID oscillation patterns or probabilities. In the intuition mode, the system uses inexpensive information (the active cell ID) and "past experiences" to make decisions through lightweight probabilistic reasoning. Wherever past experiences are insufficient, the system falls back to the deliberation mode to accrue more experi-

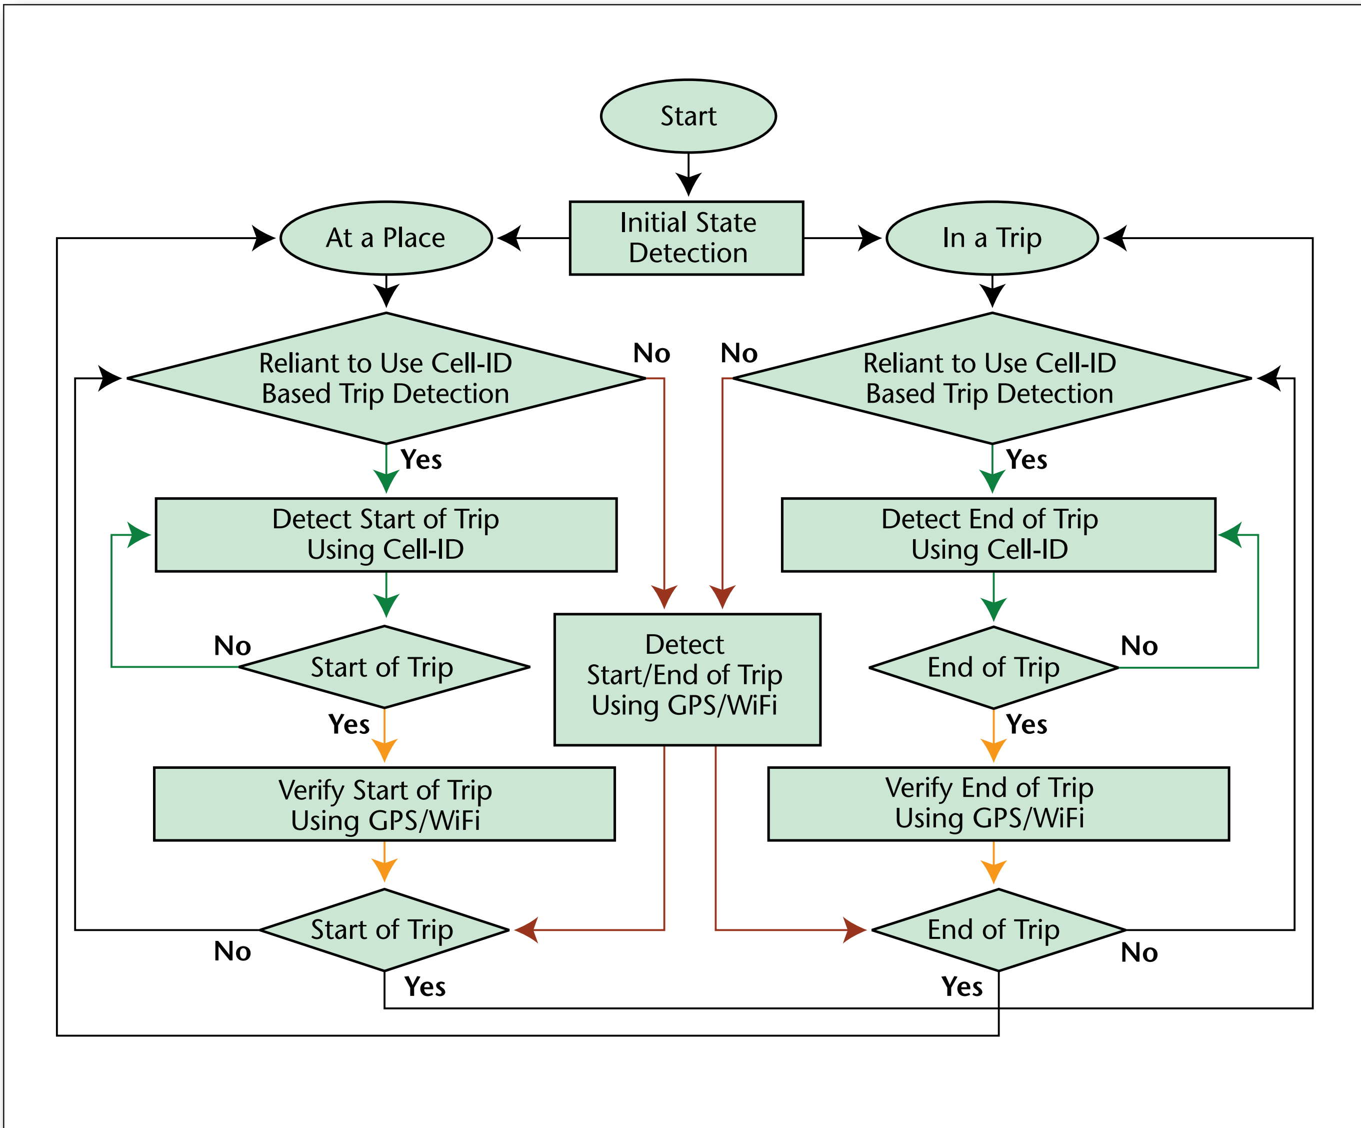


*Figure 1. Detecting the Beginning and End of a Trip Relative To a Significant Place.*

Use cell ID patterns for intuitive detection when possible, and use GPS/Wi-Fi for deliberate detection when necessary. Energy saving is achieved when a user visits the same places and repeats the same trips and, accordingly, the system works in the intuition mode.

ences by learning new patterns, reinforcing or correcting the learned patterns.

Figure 1 gives an overview of our trip detection approach as motivated above. Initially, our system determines the user's current state, whether she is at some significant place or on a trip. By definition, a significant place is where the user has stayed for more than some time (say 10 minutes) or where she frequently visits; a trip is the transition state between two places. In absence of knowledge such as an accurate schedule, when the user is at a place, the system has continuously to detect whether the user is leaving the place for a trip to another place, and, when the user is already on a trip, the system has continuously to determine whether she is arriving at some place. This is the so-called problem of trip detection, that is, to recognize the beginning and end of a trip (Kim et al. 2009).

In each of the states, at a place or on a trip, the system works in either a deliberation or an intuition mode. In the deliberation mode, sensory data such as GPS and Wi-Fi beacons are collected and used for trip detection and place recognition. Many existing methods could be leveraged here (Hightower et al. 2005; Kang et al. 2005; Kim et al. 2010). The method in our system is extended from Kang et al. (2005) and Kim et al. (2010), which use GPS/Wi-Fi to infer places/trips and accelerometer=based movement detection to reduce power consumption. After the significant places are recognized, we compute the patterns of cell IDs at those places.

In the intuition mode, we use cell IDs alone for trip detection (and place recognition). We keep a short history of recently observed active cell IDs and compute the probability of the history with regard to the patterns learned in the deliberation mode. If the user was last known to be at a place, she is likely departing the place when the probability is low relative to that place, or remains at that place when the probability remains above some threshold. Similarly, if the user was last known to be on a trip, she is likely arriving at a place when the probability is high relative to some known place, or she is still on the trip when no such place is found. It could be a new place when new cell ID observations are made; then the algorithm for place/trip learning will step in.

The two modes are not independent: The intuition mode makes trip detection decisions based on the cell ID patterns computed in the deliberation mode. The transition between modes is controlled by system parameters that gauge whether or not the patterns are sufficiently reliant. At a place or on a trip, when a new cell ID observation is made that has not been learned, the system falls back to the deliberation mode to collect more data and learn new patterns.

## Deliberation and Intuition

In this section, we describe the detailed design of the deliberation and intuition modes (phases). Specifically, we will discuss how to compute cell ID patterns in the deliberation phase, how to detect the beginning and end of a trip using cell IDs in the intuition phase, and when to switch from the deliberation phase to the intuition phase to harvest the energy benefits without sacrificing accuracy. We also highlight the key design parameters, whose values will be determined through experiments presented in the Experiments and Performance Evaluation section.

### The Deliberation Phase

This phase consists of three tasks: (1) learning significant places, (2) collecting cell ID traces associated with places, and (3) learning cell ID patterns. Leveraging the techniques presented in Kang et al. (2005) and Kim et al. (2010) for task (1), we assume that the interesting places and trips are already learned accurately. Hence we focus on tasks (2) and (3) in this work.

#### Collecting Cell ID Traces

In principle, cellular phones are able to see all available cells in the vicinity before selecting one to connect to, and some commercial phone models do provide such capability under the so-called field test mode. In reality, however, most phone operating systems only provide APIs to reveal the active cell to which the phone is currently connected. We assume that only the active cell information may be retrieved, to accommodate mass-market devices that lack a field test mode. The active cell information usually includes mobile country code, mobile network code, local area code, and cell ID. To protect user privacy, we hash each observed active cell tuple into an internal ID and refer to the hashed value as the "cell ID."

Each cell ID observation is associated with a time duration, which indicates how long that specific cell ID is continuously observed. For example, if connection to an observed cell ID *X* has lasted for a certain period, for example, 10 seconds, before the phone switches to another cell, we record this observation, the active cell ID associated with the active period. If that observation is immediately followed by another observation, the observations will appear in the trace consecutively and sequentially. Modern phone APIs allow an application to register a callback function that is called whenever the active cell changes. We can record the time difference between the moment the current active cell starts and the moment it ends (changes to a different cell). In absence of such APIs, we can read the active cell ID periodically, say every 2 seconds, and similarly record the duration of each observed cell ID.

#### Learning Cell ID Patterns

Given the cell ID traces collected at different places, task (3) aims to learn consistent patterns, which would allow us to determine (with high probability) individual places and trips between them. Our pattern learning algorithm learns the cell ID patterns from collected traces. In the input, we assume that each recognized place *Z* in the database is associated with a sequence of cell ID observations, that is, a list of (cell ID, time duration) tuples ordered through time. This provides the basis for supervised learning of cell ID patterns. We then calculate the following three sets of probabilities and store them in the database:

The prior probability of being at each place *Z*. It is computed as the ratio of total stay time at Z over the total stay time at all places.

The conditional probability of observing cell ID *X*, time duration of *X* at place *Z*, for each observation of cell ID *X*, time duration of *X* at each place *Z*. It is calculated by dividing the sum of durations of all cell ID, time duration instances by the total stay time at place *Z*.

The conditional probability of observing a new cell ID *Y* immediately following the observation of cell ID *X*, time duration of *X* at place *Z*, for each observation of cell ID *X*, time duration of *X* and each cell ID *Y* at each place *Z*, where $Y \neq X$. It is calculated by dividing the number of observation instances where *Y* appears immediately after cell

ID $X$, time duration of $X$, by the total number of instances observed at place $Z$.

Note that the time duration of $X$ in each observation is a real number and can be different every time cell $X$ is observed. To allow for meaningful and consistent probability calculation of cell ID patterns, the duration values are discretized into a small number of categories. The discretization scheme itself is a system parameter to be determined by experiments, which we discuss further in the next section

### The Intuition Phase

When operating in the intuition mode, the system continuously monitors cell ID observations in order to detect the start and end of a trip. The inputs of our trip detection process are the last known state $s$, either at-a-place or on-a-trip, and an up-to-date sequence $q$ of cell ID observations. The length of $q$ is fixed and we keep a sliding window of $|q|$ most recently observed cell IDs and their durations. Notation $q_i$ denotes the $i$th element in $q$, where $0 \leq i < |q|$.

When the last system state $s$ is at place $Z$, Our trip detection approach constantly monitors the active cell ID and calculates probability of $q$ belonging to place $Z$ every time a new cell ID observation is made and added to $q$. Note that the calculated probability should remain high if the user stays at $Z$, with a stable cell ID pattern. If the probability drops sufficiently low, say below a threshold $T_s(Z)$, the place has changed and a trip start is reported. As illustrated in figure 1, GPS/Wi-Fi scanning will be turned on to verify this detection. The system state $s$ is changed to on-a-trip if it is confirmed.

When the last state is on-a-trip, our trip detection approach continues to monitor the active cell ID, looking for potential arrival at another place $\hat{z}$, which has the highest probability of seeing cell ID pattern $q$ among all known places. If the probability of $q$ belonging to place $\hat{z}$ is above the threshold $T_e(\hat{z})$, a trip end is reported. Again GPS/Wi-Fi will be turned on for verification. If it is confirmed, state $s$ is changed to at place $\hat{z}$.

The calculation of the probability of $q$ at place $Z$ is based on Bayes's theorem. And in each step we assume a first-order Markov model, which is a widely made assumption. That is, each observation only depends on its immediate preceding observation. In the calculation, probability terms (1) the prior probability of being at each place $Z$, (2) the conditional probability of observing cell ID $X$, time duration of $X$ at place $Z$, and (3) the conditional probability of observing a new cell ID $Y$ immediately following the observation of (cell ID $X$, time duration of $X$) at place $Z$, can be directly retrieved from computation results in the cell ID learning phase at each place. We consider the sliding window size $|q|$ and, for each place $Z$, the two probability thresholds, $T_s(Z)$ (for detecting the start of a trip) and $T_e(Z)$ (for detecting the end of a trip), as parameters yet to be determined by experiments.

### Transition Between Two Phases

Obviously, we must deliberate the cell ID patterns before they could be used for intuitive trip detection. For the probabilities to be reliable, they must be computed from a sufficiently large set of data. Otherwise we must fall back to the deliberation mode for more data. The question is how much data would be sufficient and when to switch between modes.

Unfortunately, there is no obvious answer to this obvious question. At any place, for example, the set of cell IDs observed and their durations will be dependent on the local cell coverage as well as how active the user is. This has two practical implications: One is that the thresholds for the transition between deliberation and intuition are necessarily different from place to place. The other is that the value space is very large and it would be impractical to require the user to exhaust the value space for data collection.

To seek a balance between energy cost and accuracy, our system takes an incremental approach, which is illustrated in figure 2. Initially, the system works in a deliberation mode to aggressively collect data. After spending a few hours at one place, it computes the cell ID patterns for that place. Then, it starts trying to use the cell ID patterns for prediction, while staying in the deliberation mode. Hence there is a period (called dual mode) in which the two phases overlap. In this period, the results yielded from the intuitive predictor are compared to those from the deliberate predictor. When the accuracy of the intuitive predictor is sufficiently high (based on the results of the deliberate predictor), the deliberate predictor will phase out.

On one hand, there are heuristics by which we can estimate how long the user must stay before entering the dual mode. For example, it may deserve a trial if the user has accumulated two hours of data at a place. On the other hand, exceptions may still arise even after it has demonstrated sufficient accuracy at some point and is already in a pure intuition mode. This may happen, for example, when a new cell ID $X$ is observed at place $Z$, or the (discretized) duration $\delta X$ of some $X$ is not yet in the computed patterns. In such cases, our system will temporarily fall back to the dual mode.

## Experiments and Performance Evaluation

In this section, we evaluate the performance of the proposed framework for energy-efficient trip detection. Specifically, we will answer the following

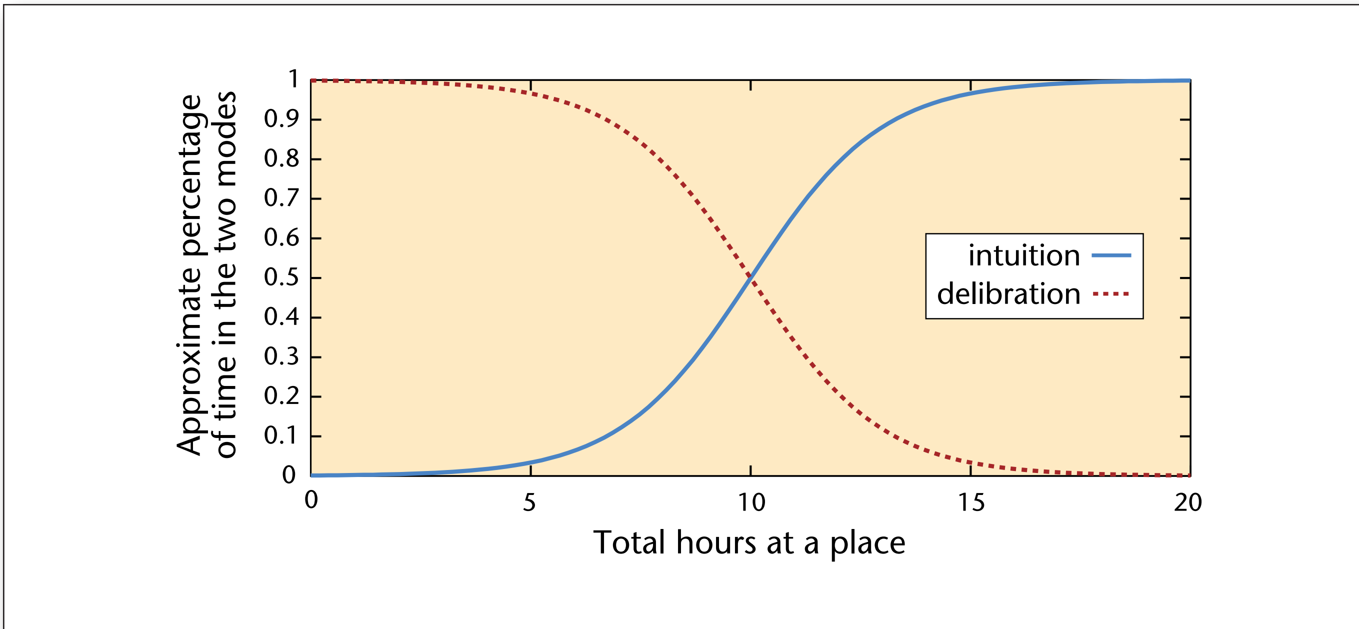


*Figure 2. The Energy-Consuming Deliberation Phases Out After Sufficient Experience Has Accrued for Intuitive Trip Detection.*

The two modes coexist for some period of time for desired accuracy.

three questions: How well does our approach perform trip detection in terms of accuracy, timeliness, and energy efficiency? How do the identified parameters affect system performance? How do we choose the "best" parameter values?

## Data Collection

We have developed a data collection program that collects GPS, Wi-Fi, accelerometer, and cell ID information. GPS and Wi-Fi data are collected every 2 seconds, while cell ID data are collected using the callback mode and the current cell ID is recorded every time it changes. Accelerometer data are sampled at 30 Hz. The data collection program has been deployed on six Nokia N900 devices, each carried by a user in our study. Four users ran it each for four months, and the other two users ran it each for one month. To obtain the ground truth, we asked the users to take notes on all the trips during the data collection period. All trips are also verified on maps using the Google Map API and we have manually marked the start and end of each trip in the collected data.

In total, we have collected 1755 trips, among which there are 139 unique trips and 83 unique places. To demonstrate the advantages of our approach, we ignore casual trips and focus only on regular trips, that is, trips that have been repeated at least once. This leaves us with 1701 trips, which contain 85 unique trips and 43 unique places. On average, a trip is repeated 4.17 times, ranging from 2 to 62 times. The trips range from 1 mile to 32 miles in distance (average 12 miles) and include driving, cycling, and walking trips. Trip duration varies from 5 minutes to 53 minutes, and the average is 24 minutes. The amount of time that our users stay at a place varies from 8 minutes to 16 hours and the average is 5 hours. The regularity in our data set is consistent with Hu and Reuscher (2004).

## Parameter Settings

There are three key parameters in our framework: the trip start and end thresholds, $T_s(Z)$ and $T_e(Z)$, for each significant place $Z$, the discretized duration length $\delta_X$ of each cell ID connection $X$, and length $|q|$ of observed cell ID sequence $q$ on which we compute probabilities.

### Length of Recent Cell ID Sequence

As described earlier, we compute the probability of a sequence $q$ of recent cell IDs. A larger $|q|$ value may yield more accurate probability estimation and better trip detection accuracy, but it may also reduce the timeliness of trip detection and incur higher computation overheads. As will be shown (figure 6), when $|q| = 3$, the algorithm demonstrates a desired performance trade-off, and there is no obvious improvement in accuracy when $|q| > 3$. Therefore, we set $|q|$ to 3 in the other experiments and our system.

### Discretization of Cell ID Connection Durations

Figure 3 shows the probability distribution of cell ID connection durations in users' places and trips. We see that the connection durations vary significantly from several seconds to a few hours. Equally treating all the possible duration values would result in many zero probabilities. Instead, we need

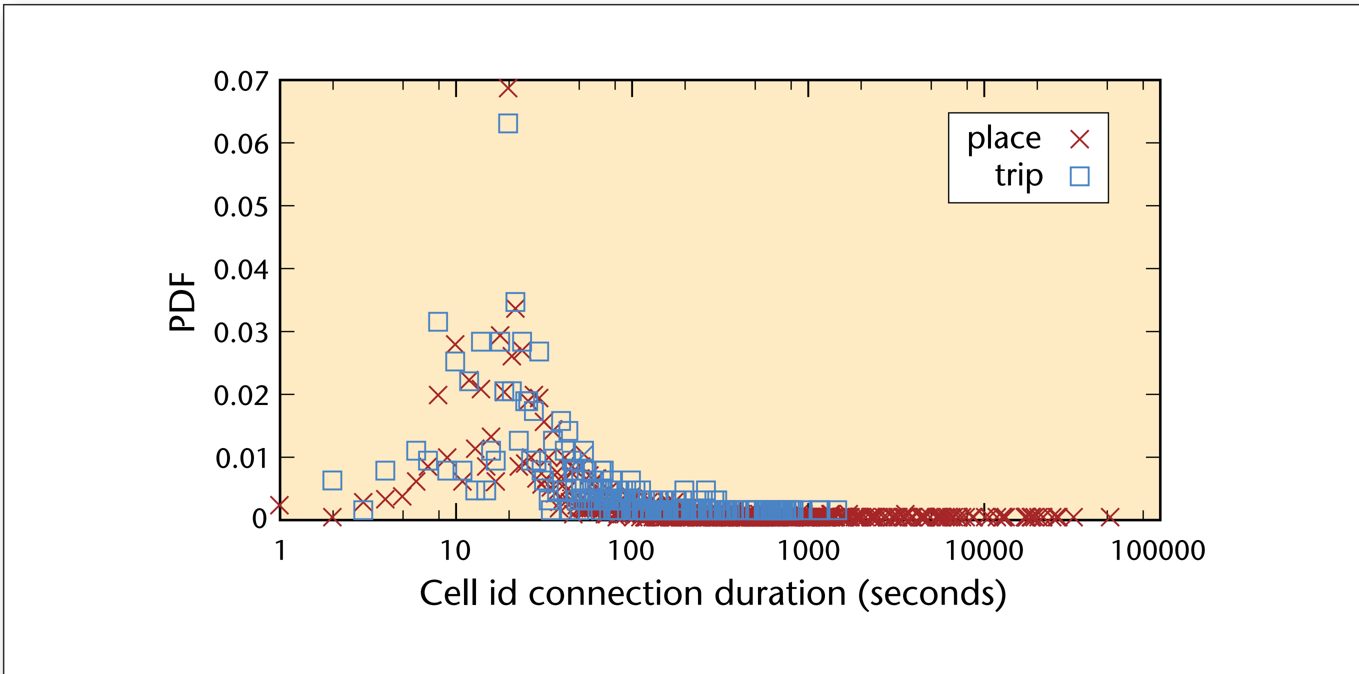


*Figure 3. Distribution of Cell ID Connection Duration (δ) When at Place and on a Trip.*

Accordingly, we define 6 discrete categories, that is, the class is 6 if δ > 243, and $[log_3\delta]$ if 0 < δ ≤ 243.

to discretize the duration values into multiple categories that best capture the ranges of connection durations that occur frequently. Our analysis reveals that (1) connection duration goes above 1000 seconds only when a user stays at a place (not during a trip), and (2) connection duration is most frequently around 20 seconds for both trips and places.

Let δ be the duration value of some cell ID connection. We define the discretization function as follows: If δ > 243, the class is 6; or otherwise if 0 < δ ≤ 243, the class is $[log_3\delta]$. Then we have finer classes for more frequent durations and coarser classes for less frequent connections. This helps to distinguish start/end of trips from places.

### Trip Start/End Thresholds

For each place $Z$, the thresholds $T_s(Z)$ and $T_e(Z)$) determine at what values of probability $Pr(q|Z)$ a user should be considered at a place or at the start/end of a trip. To decide their values, we studied the probability distributions of all users staying, departing, and arriving at their significant places and found that they demonstrated similar patterns.

As an example, figure 4 plots the probability distributions when one user stays at some place $Z$, leaves place $Z$ (from the user crossing the place boundary until $Pr(q|Z) = 0$), and arrives at place $Z$ (from $Pr(q|Z) > 0$ until the user crosses the boundary into $Z$). When the user stays at place $Z$, the curve shows higher probabilities for larger values of $Pr(q|Z)$; there is a long tail, however, when the user walks around the boundary of the place and/or some rare cell ID observations are made. When the user is leaving or approaching place $Z$, the probabilities of $Pr(q|Z)$ are relatively low as compared to those when staying at place $Z$. By the two curves, values of $-log(Pr(q|Z))$ are almost normally distributed around 4–10, where the user is nearby place $Z$. A few high probabilities happen when the user is at the place boundary and a few very low probabilities happen when some rare cell ID observations are made.

The goal is to find the optimal threshold $T^*_s(Z)$ or $T^*_e(Z)$ that separates the cases of staying at place $Z$ and leaving (or arriving at) place $Z$. In the spirit of the clustering threshold tuning approach proposed in Park et al. (2010), our optimal threshold $T^*_s(Z)$ or $T^*_e(Z)$ corresponds to the point at which the two corresponding probability distributions intersect. Our experimental results show that $T^*_s(Z)$ or $T^*_e(Z)$ may take different values for the same place $Z$. For each place $Z$, our system will recalculate the threshold periodically since both user behavior changes (for example, switching office room to another side of the building) and cell ID changes (for example, a new cell ID observation is made) may entail adjustment to the thresholds.

Monotonicity of $p(q|Z)$ when a user is leaving or approaching $Z$ is important because the verification would be triggered many times and thus not be energy efficient if the value of $p(q|Z)$ changed

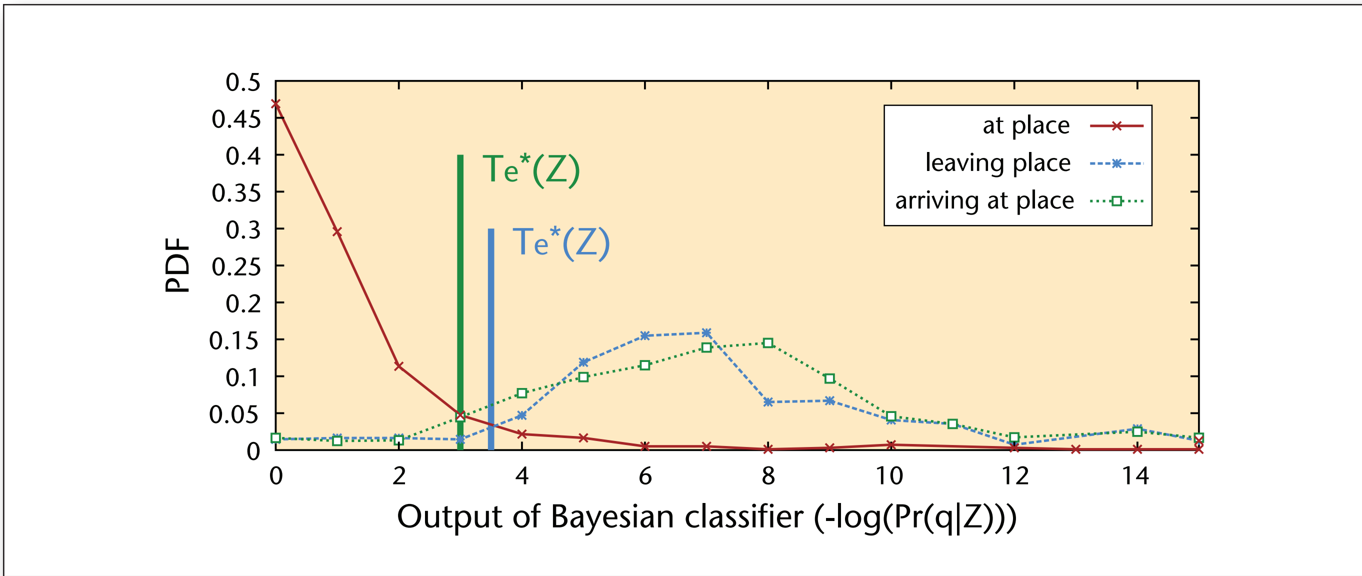


*Figure 4. Probability Distribution Plot.*

Determination of the probability thresholds $T^*_s(Z)$ and $T^*_e(Z)$ for place $Z$. Note that the probability $Pr(q|Z)$ decreases as $(-log(Pr(q|Z)))$ increases along the $x$-axis. To handle small probability values, we compute $-log(Pr(q|Z))$ instead of $Pr(q|Z)$.

drastically. However, we found that 85 percent of $p(q|Z)$ are monotonic when the user is leaving or approaching $Z$. We also found that 94 percent of the 15 percent nonmonotonic values happen below (when leaving $Z$) or above (when approaching $Z$) the threshold, which does not need verification by GPS/Wi-Fi.

## Performance Metrics

Figure 5 shows a typical trip in which a user travels from place $Z_1$ to place $Z_2$. Consider the following four phases:

Phase 1: Before a trip starts, the system should report that the user is at place $Z_1$. It is a false positive if the system reports a trip start ($S$). Although this type of error will be corrected by verification, a high false positive error rate will cause frequent verification and thus high energy cost.

Phase 2: A trip has started, which should be detected ($S'$); otherwise it is a false negative error if the system still reports at place $Z_1$. Ideally, the trip start should be detected before $Pr(q|Z_1)$ drops to zero so as to reduce the latency. This is where threshold $T_s(Z_1)$ comes into play.

Phase 3: During a trip, the system should report on a trip. Otherwise, if the system reports end of trip ($E$), it is a false positive error, which costs energy for verification. Practically, the report of $E$ is allowed to happen anywhere between when $S'$ and $E'$ are reported.

Phase 4: When arriving at place $Z_2$, the system should report end of trip ($E'$). Ideally, to have a timely report, $E'$ should be detected after $Pr(q|Z_2)$ turns greater than zero but before it reaches its maximum or equilibrium. This is where threshold $T_e(Z_2)$ comes into play. Reporting $E'$ either right before $d$ or right after $d$ is considered acceptable (accurate). However, if the system still reports on a trip after passing point $d$, it is a false negative error.

The performance of our trip detection solution can be measured using three key metrics: (1) accuracy measures — how often our system correctly detects a trip start or end; (2) timeliness measures — how quickly our system detects a trip start or end; and (3) energy efficiency measures — the energy overhead incurred by our trip detection system. Ideally, our system should be able to detect trip starts and ends with high accuracy, in a timely manner, and with high energy efficiency. However, these three metrics affect each other and it is our goal to seek an appropriate performance tradeoff. We will explore the three metrics in the following sections.

## Accuracy

Since we resort to existing localization methods in the deliberation mode, here we only evaluate the accuracy of our approach in the intuition mode. For each trip and place, we compare our decision with the ground truth in each of the four phases. Our system makes a trip detection decision every time the active cell ID changes. We define accuracy in any phase $y$ as the number of correct detections in phase $y$ divided by the total number of detections in phase $y$.

Figure 6 compares the trip detection accuracy when different cell ID sequence length values are used. We make the following observations:

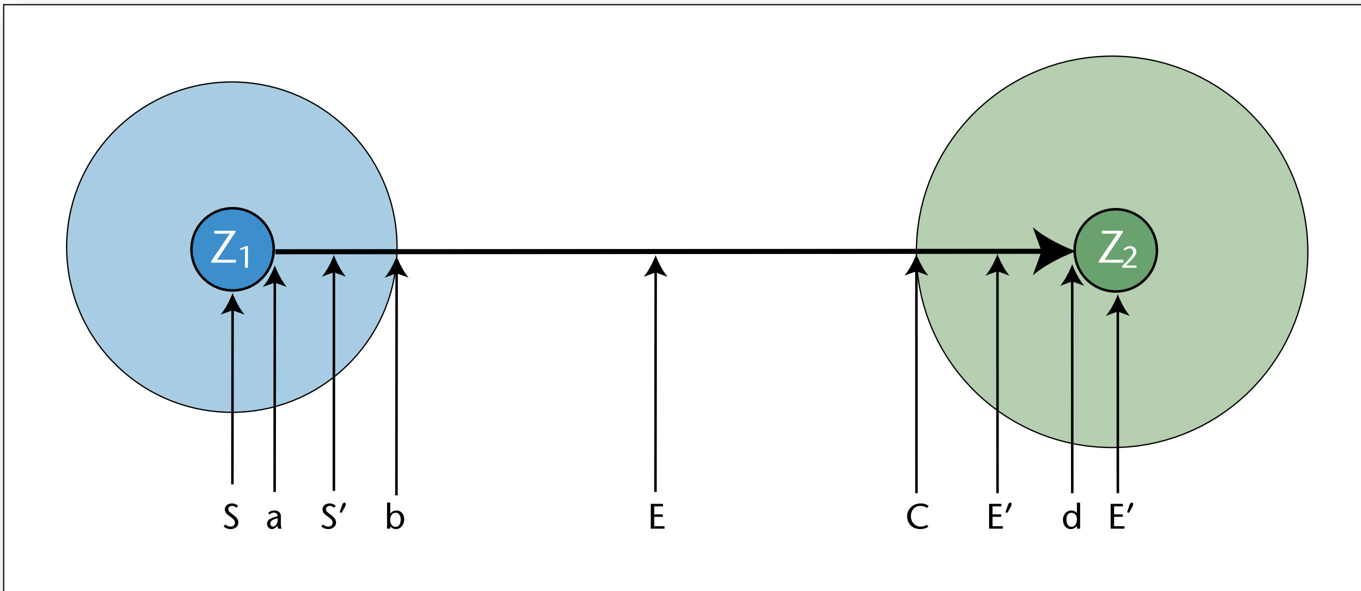


*Figure 5. A Trip from Place $Z_1$ to place $Z_2$.*

a/d is the start/end of the trip by ground truth; b/c is where the observation probability $Pr(q|Z_1)$ / $Pr(q|Z_2)$ turns zero when moving away from place $Z_1$ / $Z_2$.

Our chosen sequence length $|q| = 3$ achieves the best accuracy overall, and it represents a good balance among accuracy, energy efficiency, and timeliness.

Longer sequences help improve the accuracy of Phase 1, since people normally stay at a place longer and testing more cell IDs will help increase the confidence of decision.

Accuracy drops for Phase 2 and Phase 4 when $|q| > 3$, since longer sequences are more conservative and make it harder to detect trip start and end in a timely manner.

Longer sequences result in a slight decrease of accuracy in Phase 3, since people's trips tend to be shorter than when staying at a place, and increasing sequence length may introduce more noise and reduce detection accuracy.

Figure 7 shows the accuracy of the four phases using different classification thresholds. The $x$-axis is the percentage of shift from the two thresholds automatically calculated as explained before: zero means the original thresholds, 10 means that both thresholds are incremented by 10 percent, and so forth. We can make the following observations:

Our automatically determined thresholds balance the false positive and false negative for both start and end of trip detection. This is because the accuracy of Phases 1 and 4 increases while that of Phases 2 and 3 decreases when lowering the thresholds.

Under our automatically determined thresholds, higher accuracy is achieved for Phase 1 and Phase 2, that is, detecting a trip start is stabler than detecting a trip end. This is because the cell ID pattern of staying at a place is stabler, while trips may differ in the means of transportation, speed, traffic conditions, and routes.

For Phase 1 and Phase 2, lowering the thresholds leads to higher accuracy for Phase 2 (detecting trip start) but lower accuracy in Phase 1.

For Phase 3 and Phase 4, lowering the thresholds leads to lower accuracy for detecting end of trip in Phase 4, but higher accuracy in Phase 3.

In practice, we could adjust the thresholds to make trip detection more timely or more energy efficient.

Further, we also studied the overall accuracy of trip detection by using metrics of precision and recall. Recall is computed as the fraction of correctly detected trips among all trips that the users had, and precision is computed as the fraction of correctly detected trips among all trips that the algorithm believes to be correct. Based on our data set, we compared our method with four other methods: (1) GSM based, we use PlaceSense (Kim et al. 2009); (2) GPS based, we use Kang et al. (2005); (3) GPS + Accelerometer (or GPS+), we use SensLoc (Kim et al. 2010); (4) Wi-Fi + Accelerometer (or Wi-Fi+), we use SensLoc (Kim et al. 2010). The experimental results are plotted in figure 8.

As shown in figure 8, Wi-Fi- and GPS-based methods both have high precision and recall. For the Wi-Fi-based method, the errors in recall are because several buildings and open places in our data are not covered by Wi-Fi and hence trips from and to those places are not detected; the errors in precision are caused by noise of Wi-Fi signals; when the user walked within a building, some-

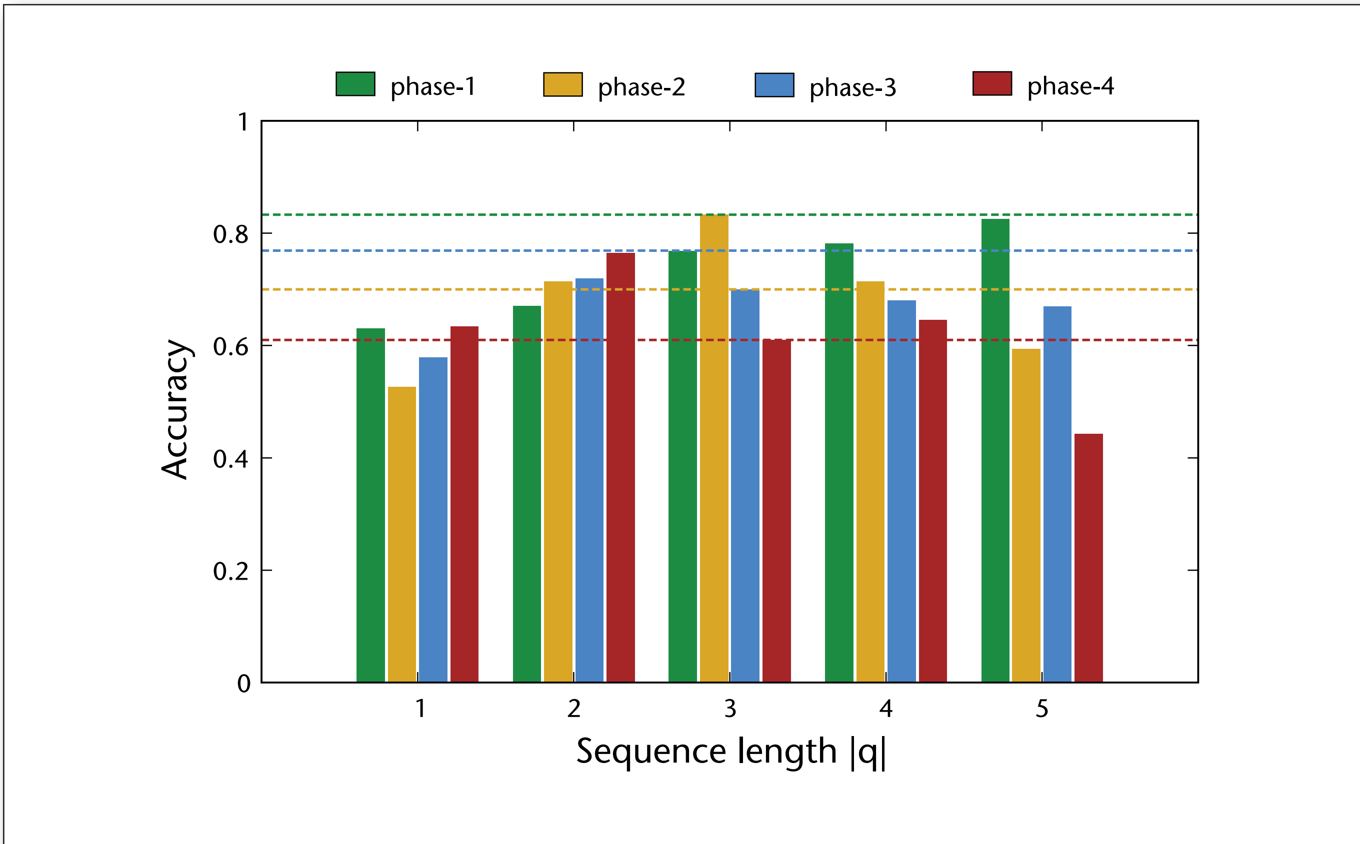


*Figure 6. Trip Detection Accuracy in Four Phases Using Different Cell ID Sequence Lengths.*

$|q| = 3$ represents a desired performance trade-off.

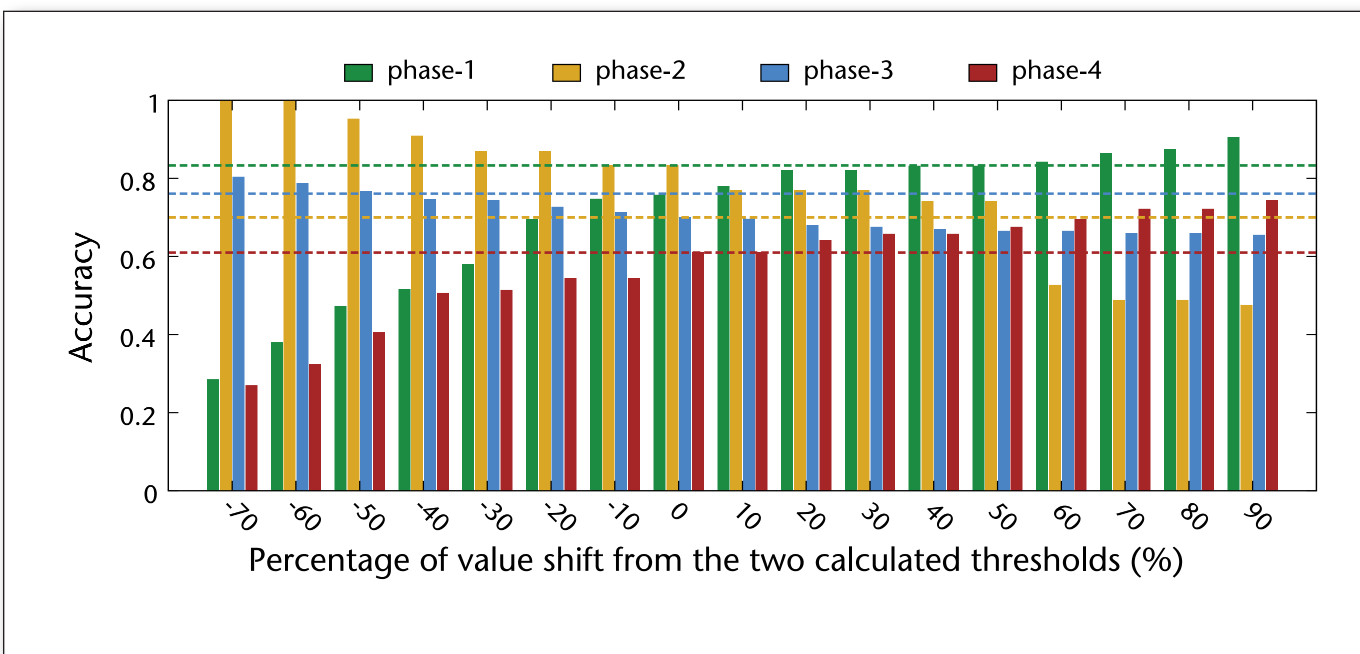


*Figure 7. Trip Detection Accuracy by Different Thresholds.*

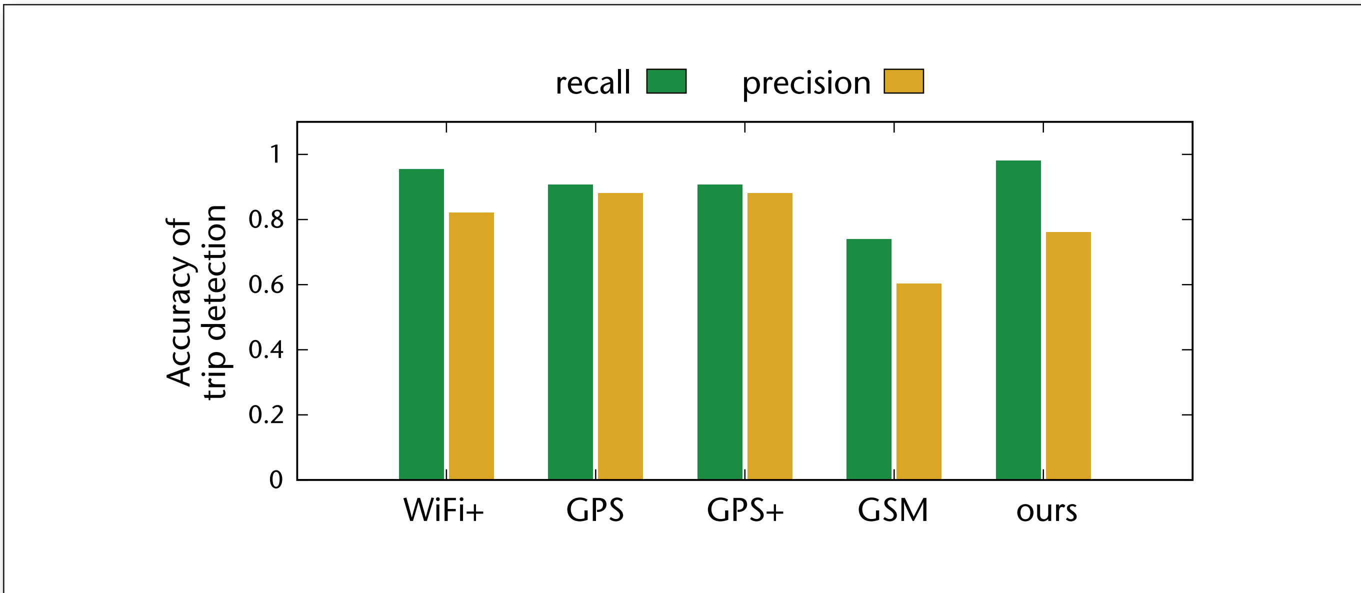


*Figure 8. Comparison of Overall Accuracy of Four Methods.*

times the algorithm treats it as a trip start. For GPS-based methods, the errors are due to the poor performance in urban canyon and indoor areas. GSM-based methods have poor accuracy because single-cell ID can only work at a very coarse level and it may not be able to detect short trips; oscillation of cell IDs at a place makes the precision low. Our method has better performance than Wi-Fi- and GPS-based methods in recall because the cell ID has better coverage than Wi-Fi and GPS. Although the precision of our method is slightly lower, our verification process using Wi-Fi or GPS compensates for the errors.

## Timeliness

Not only do we want to detect trips accurately, the detection must also be made in a timely manner. That is, when a trip starts or ends, our system needs to detect it quickly. The timeliness of trip detection depends on two factors: the first is how fast the cell ID changes when starting or ending a trip, and the second is how fast we make the decision. The ground truth of trip starts and ends is determined by the time when a user entered or left a building.

To evaluate how fast the active cell ID changes after a user starts or ends a trip, we examine all trips in the collected data. For each trip we extract time of trip start (or end) $t$ from the ground truth, and calculate the time of the first cell ID change $t_c$ after $t$. The difference between $t$ and $t_c$ is the inherent delay in our trip detection process. Figure 9a plots the cumulative distribution function (CDF) of all $t_c - t$ values calculated from our collected data for trip start and end. It shows that, for 80 percent (90 percent) of the time, the cell ID changes within 50 seconds after a user starts (ends) a trip.

We then evaluate the average delay in our intuitive trip detection. As shown in figure 9b, for 80 percent (trip starts) (90 percent [trip ends]) of the time, our system can detect the trip start (end) in less than 62 (68) seconds. It is clear that the actual delays, 12 or 18 seconds, are insignificant. They are caused by trading timeliness for accuracy in our work: often another cell ID change is awaited for $Pr(q|Z)$ to meet the desired threshold with regard to place $Z$. Our results are consistent with those in Kim et al. (2010) yet we use a more energy-efficient cell-ID-based method that is, however, insensitive to extra delays as possibly caused by the GPS/Wi-Fi scan period in their work.

We compared timeliness of different methods in Figure 10. As shown in figure 10, our method is comparable to the Wi-Fi-based method in that about 80 percent of trip starts can be detected within 60 seconds. The GPS- and GSM-based methods have longer delay on trip start detection because GPS takes time to find fixes when the user starts the trip and GSM has a coarse granularity in place recognition.

Figure 10 compares timeliness of trip end detection in different methods. Similarly to the Wi-Fi-based method, which can detect 80 percent of trip ends within 60 seconds, our method can detect 80 percent of trip ends within 40 seconds. The GPS and GSM methods can detect trip ends much earlier than our method — more than 90 percent of trip ends are detected 20–240 seconds ahead of time. However, they may lead to premature decisions when the user just passes by the place, which lowers the trip detection accuracy.

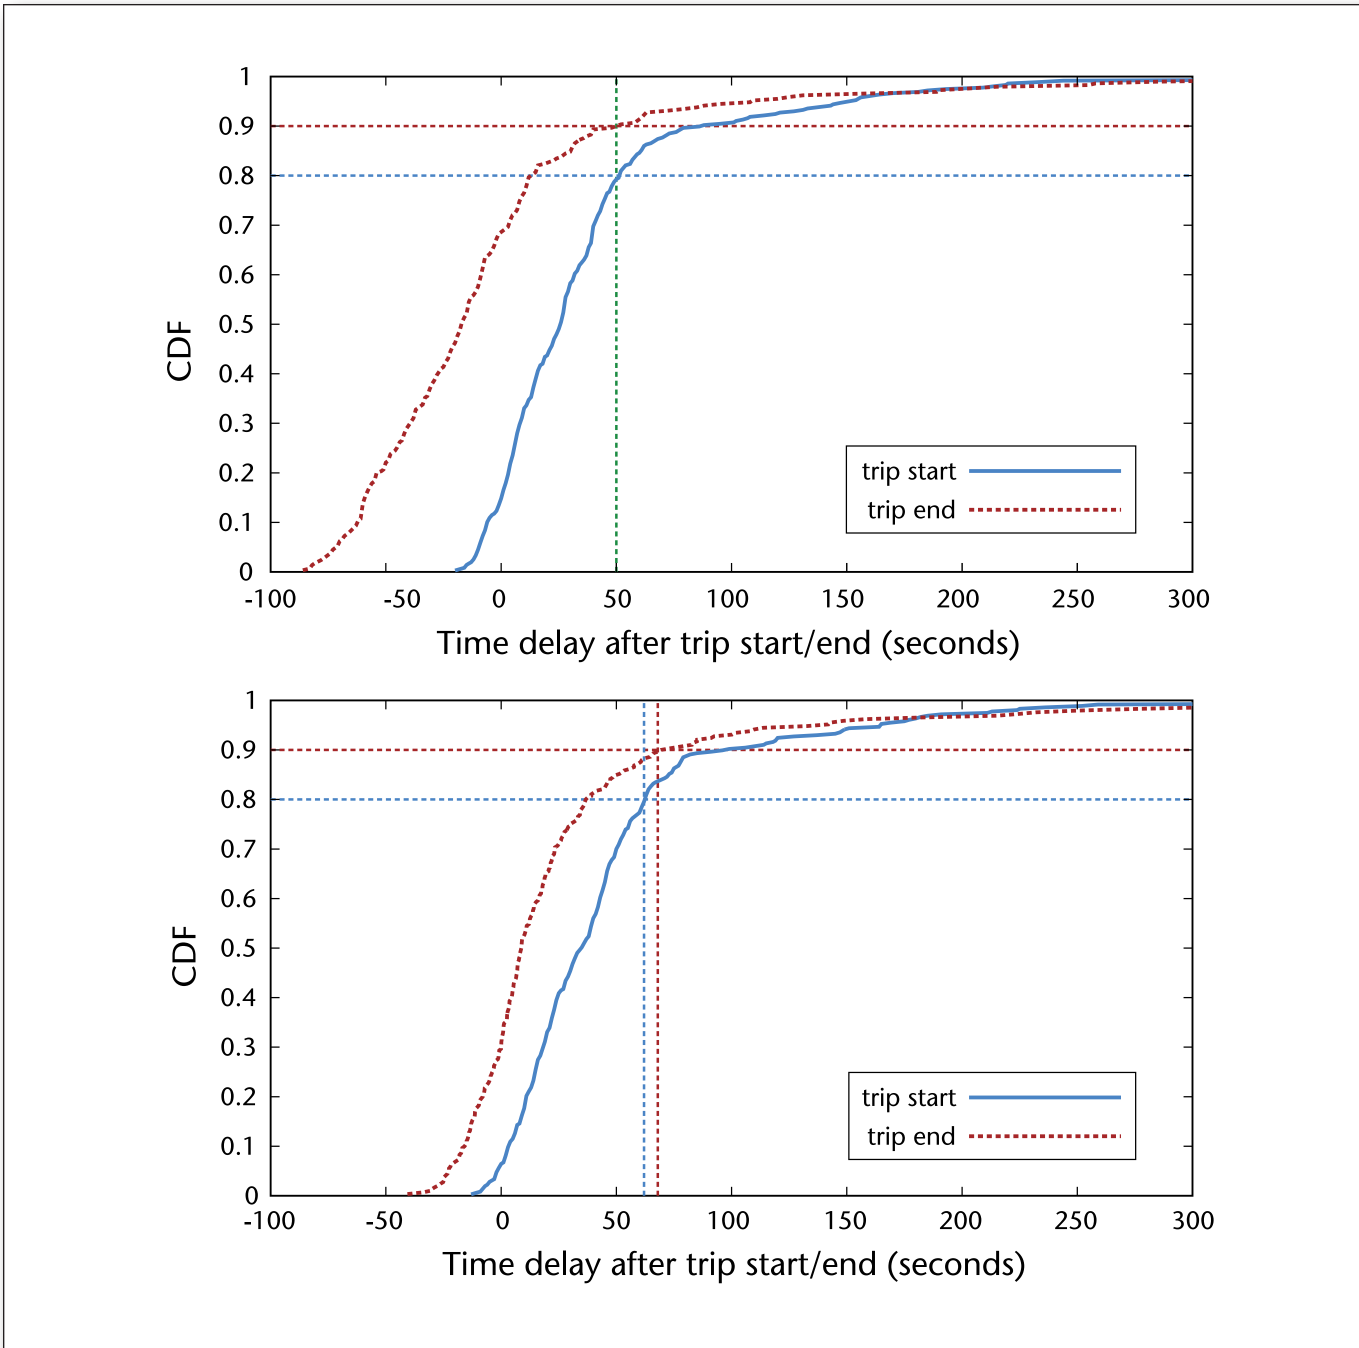


*Figure 9. Delays in Active Cell ID Changes and Intuitive Trip Detection.*

In figure 10, both CDF curves of our method have a long tail, which indicates that 7.5 percent (5 percent) of trip detections has a delay of more than 120 seconds. Those delays happened in remote areas due to sparse cellular coverage — the cell IDs did not change that frequently. To address the delays, currently we resort to traditional GPS-accelerometer-based methods (for example, Kim et al. [2010]) for energy saving at those places.

## Energy Efficiency

To estimate the energy use of our approach, we make the following assumptions as in Kim et al. (2010): In the deliberation mode, (a) the accelerometer is always on; and (b) GPS is used only when there is no Wi-Fi; (c) neither GPS nor Wi-Fi is on when the user is detected as stationary by accelerometer readings. In the intuition mode,

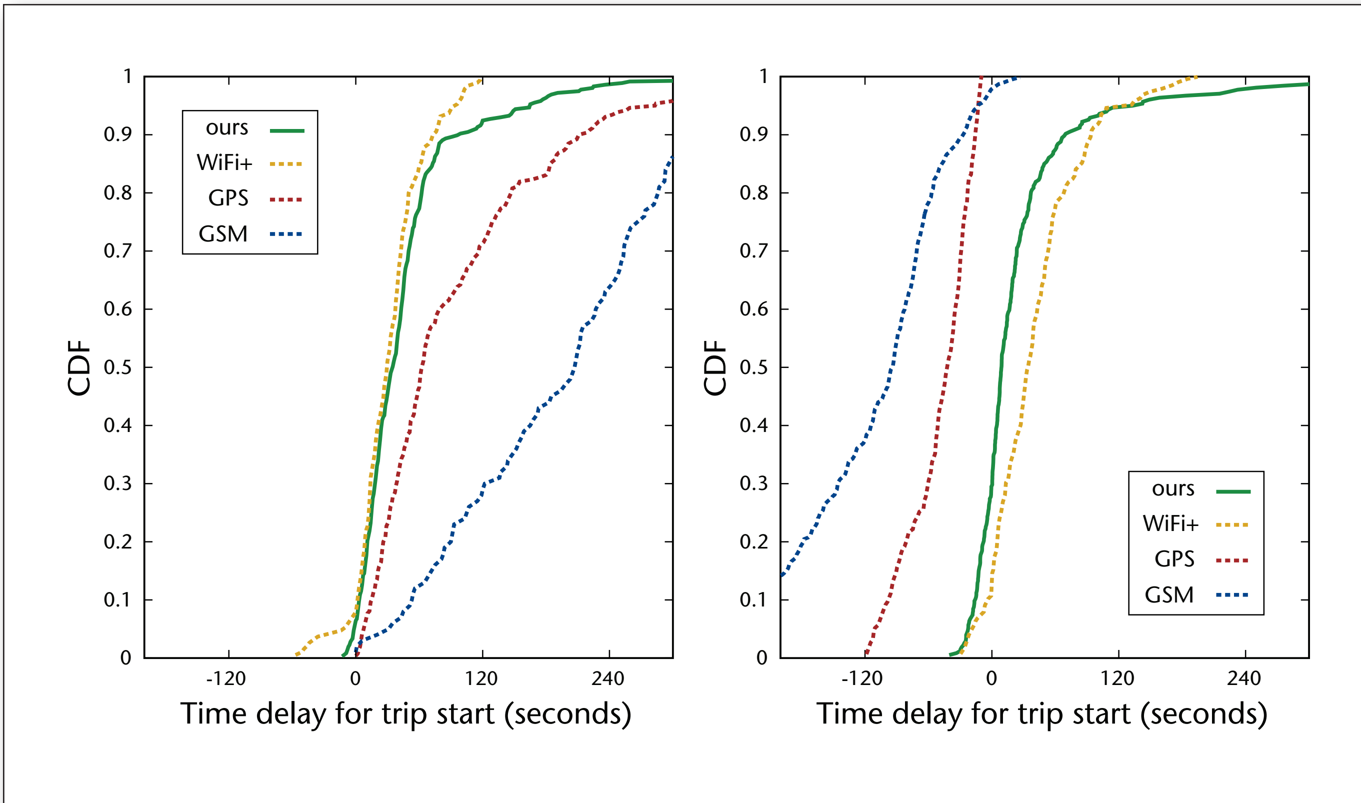


*Figure 10. Comparison of Trip Start/End Delays.*

(a) only cell ID information is collected; and (b) Wi-Fi and accelerometer are always off. We calculate average power consumption based on the N900 power consumption specification.[1] Specifically, reading cell ID costs 0.01 milliwatt each time, Wi-Fi scanning at 1/6 Hz costs 80 milliwatts, GPS reading at 1/2 Hz costs 325 milliwatts, and accelerometer reading at 30 Hz costs 25 milliwatts. In addition, we have conducted in-field power consumption measurements. These numbers are slightly higher than the spec but generally consistent, anf do not affect the overall energy-efficiency improvement we report here.

It is our observation from the data that a user's mobility heavily affects power consumption. For instance, if a user is stationary most of the time, the system consumes very little energy even when running in the deliberation mode. If a user stays at a place all day and patterns at that place are already learned, then our system will run in the intuition mode most of the time and the power consumption will also be very low. However, if a user visits many new places, then the cell phone will run in the deliberation mode most of the time and the power consumption will be very high.

Figure 11 shows the average percentage of sensing time based on the total daily system running time of six users. We consider the sensing time of cell ID, GPS, Wi-Fi, and accelerometer over the first 21 days. At first, the system runs mostly in the deliberation mode, and the accelerometer and Wi-Fi are turned on most of the time. After a few days, for some places, the system starts to work in the intuition mode; then the daily working time of Wi-Fi and accelerometer decreases generally, but still goes up when a user visits new places. The working time of GPS is dependent on how much time is spent in places without Wi-Fi coverage.

Figure 12 shows the average power consumption of six users over the first 21 days. It shows that, overall, the average power consumption decreases quickly over time in our method as the system learns about more places and spends more time in the intuition mode. On some days, the average power consumption increases, as a result of users making more trips or visiting new places. By comparison, other methods fail to demonstrate a similar trend of decreasing power consumption over time.

Over the first 21 days, the average power consumption is 49.1 milliwatts for the GPS-based method, 29.6 milliwatts for the Wi-Fi-based method, 1 milliwatt for the GSM-based method, and 14.5 milliwatts for ours. Hence, our approach saves 70 percent energy compared to the GPS-based method and 51 percent compared to the Wi-

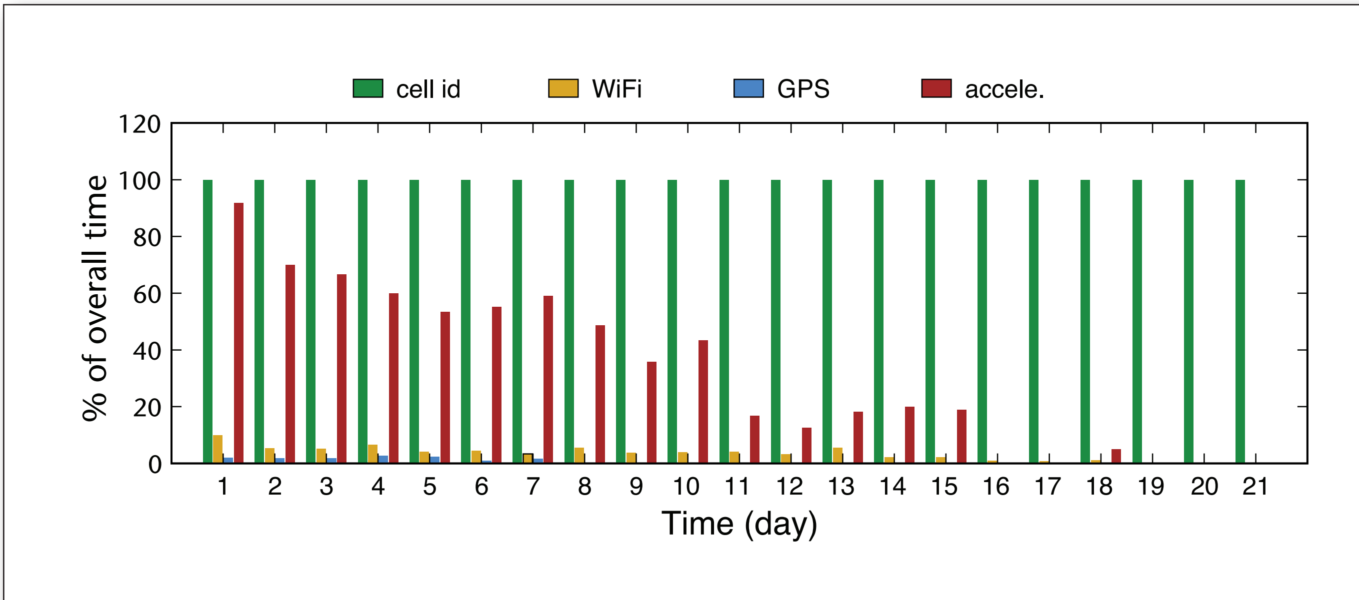


*Figure 11. Average Sensing Time of Cell ID, Wi-Fi, GPS, and Accelerometer by Six Users in the First 21 Days.*

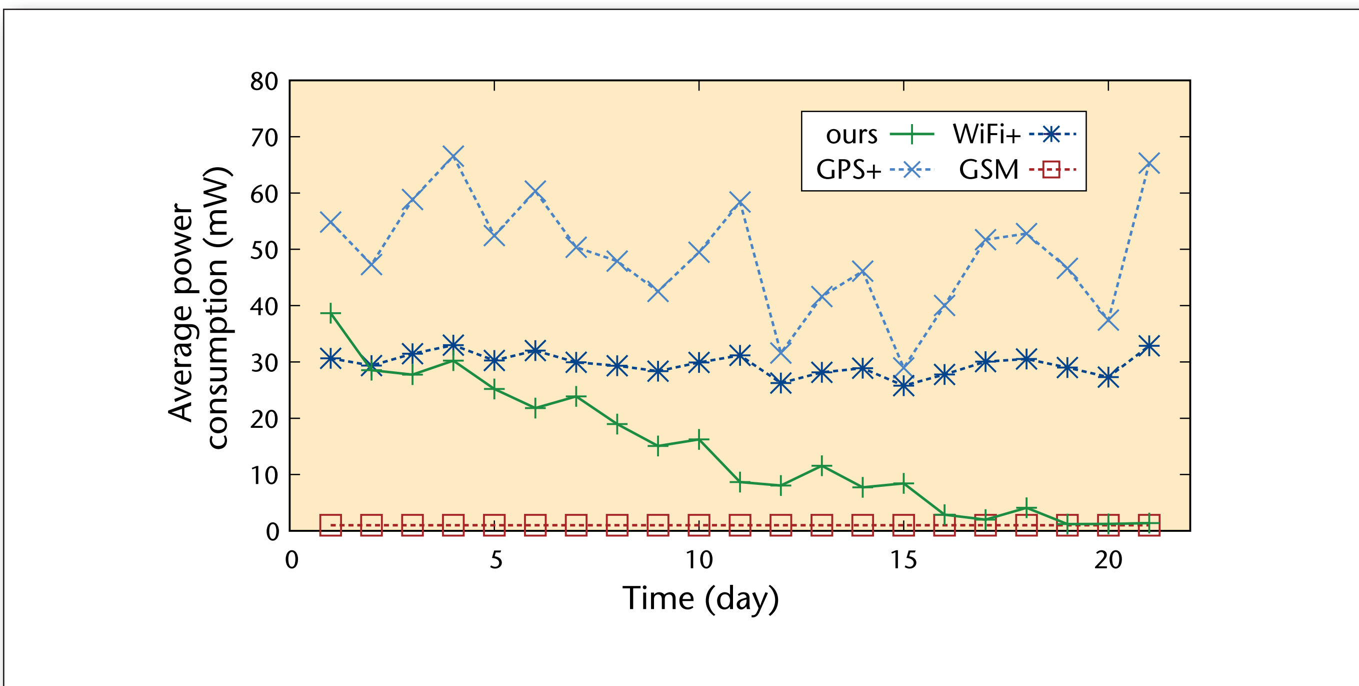


*Figure 12. Average Energy Consumption of Six Users in the First 21 Days, Which Shows Obvious Improvement over Time in Our Method d.*

Fi-based method. Furthermore, the longer a user runs our system, the more energy the user will save. In a long run, the power consumption of our system will be close to the GSM-based approach because only cell ID is used most of the time.

Figure 13 shows the average trip detection accuracy of six users in the first 21 days with our method. It shows that the precision and recall of trip detection remain high and stable over time although the system changed back and forth between deliberation and intuition modes. That is, there is no loss of accuracy despite a rapid decrease in energy consumption over time. The recall is around 95 percent because both cell ID (in intu-

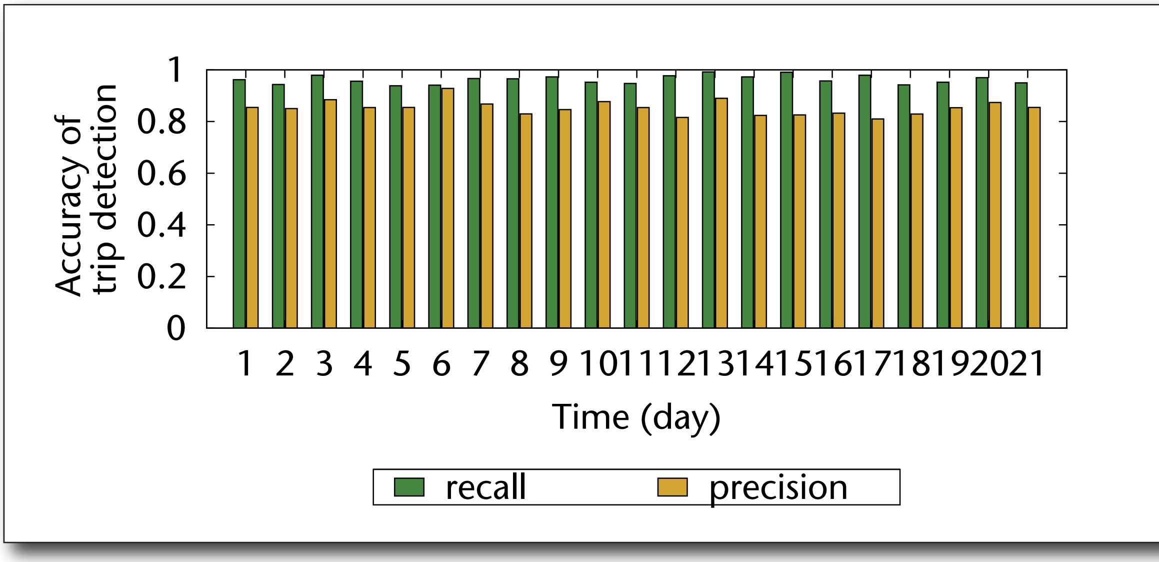


*Figure 13. Average Trip Detection Accuracy of Six Users in the First 21 Days, Which Reveals No Loss of Accuracy over Time.*

ition) and GPS + Wi-Fi (in deliberation) have good signal coverage for most of the places. The precision is around 85 percent because the system switches to intuition mode only when the accuracy is high enough.

## Conclusions and Future Directions

This article applies a well-known theory in psychology regarding human decision making (Kahneman 2011) to the new domain of human activity recognition on mobile devices. We propose a framework with a similar deliberation-intuition architecture to approach the chronic problem of balancing accuracy and energy efficiency from a new perspective. We first elaborate this framework in an important subproblem in activity recognition, that is, trip detection, that focuses on predicting the starts and ends of trips with regard to significant places. Specifically, this work makes the following two novel contributions.

We propose a two-phase framework for energy-efficient trip detection on mobile devices: the deliberation phase uses sensors such as GPS and Wi-Fi as in previous work for training a cell ID learner, and the intuition phase solely uses cell ID patterns for predicting start/end of a trip. The goal is to save energy by intuition when a user revisits familiar places or repeats familiar trips. We elaborated both the science and engineering aspects of this framework.

We have collected real-life traces from six users over five months and used the data to evaluate the framework in terms of accuracy, timeliness, and energy efficiency. Our experiments demonstrated a 51–70 percent energy savings over previous GPS-/Wi-Fi-based methods and, moreover, a clear trend in decreasing power consumption after a period of deliberation, with considerable accuracy and timeliness in intuitive trip detection. We also showed how the system parameters are tuned and how they affect the system performance trade-offs among the three metrics.

Our framework can be implemented either at the application level as a backstage service or at the operating system level as a system service. Most mobile platforms expose APIs for accessing cell IDs; the implementation is straightforward. In fact, we have prototyped our system on both Maemo and Android. The iOS platform, however, does not expose such APIs and hence the implementation is not convenient at the application level. Nevertheless, it can be done at the operating system level where the information is available.

We plan to look into the following directions in future research: First, it is possible for people who live or work nearby to share the learned cell ID patterns and collaboratively construct a community database of patterns. This would be useful, for example, for better coverage of data and faster bootstrapping of the system. It would be interesting to study the mechanisms for supporting this construction, reuse, and evolution process and related privacy issues.

Second, as is discussed in Kahneman (2011), the intuitive human decision process, although fast, is fundamentally biased and largely limited to the so-called availability heuristics. In the same light, a mobile app for human activity recognition like this and other works are also subject to the same limitation as imposed by the availability of data. Our approach performs well in urban areas where cellular coverage is typically dense and there are GPS and Wi-Fi signals to make use of. In remote areas such as parks and other recreational places, however, those signals may not even be available. Other sensors such as accelerometer, gyro, and compass could be exploited in those areas.

Third, in this article, we elaborate the framework for trip detection. The success seems attributable to the fact that different sensing methods (that is, cell ID versus GPS/Wi-Fi) have complementary strengths in energy efficiency and accuracy. It is interesting to explore further how this framework can generalize to other sets of sensing methods, different performance trade-offs, and application domains.

## Acknowledgements

We thank the editor and reviewers for useful feedback and suggestions. This work was supported in part by the National Science Foundation under Grant No. CNS-0910995.

## Note

1. For N900 hardware power consumption, see wiki.maemo.org/N900 Hardware Power Consumption.

## References

Ashbrook, D., and Starner, T. 2003. Using GPS to Learn Significant Locations and Predict Movement Across Multiple Users. *Personal Ubiquitous Computing* 7(5): 275–286.

Bales, E.; Li, K. A.; and Griwsold, W. 2011. Couplevibe: Mobile Implicit Communication to Improve Awareness for (Long-Distance) Couples. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing,* 65–74. New York: Association for Computing Machinery.

Cao, X.; Cong, G.; and Jensen, C. S. 2010. Mining Significant Semantic Locations from GPS Data. *Proceedings of the VLDB Endowment*. 3: 1009–1020. Provo, UT: Very Large Data Base Endowment, Inc.

Chen, Y.; Yang, Q.; Yin, J.; and Chai, X. 2006. Power-Efficient Access-Point Selection for Indoor Location Estimation. *IEEE Transactions on Knowledge and Data Engineering*. 18(7): 877–888.

Cuervo, E.; Balasubramanian, A.; Ki Cho, D.; Wolman, A.; Saroiu, S.; Chandra, R.; and Bahl, P. 2010. Maui: Making Smartphones Last Longer with Code Offload. In *Proceedings of the 8th Annual International Conference on Mobile Systems, Applications, and Services*, 49–62. New York: Association for Computing Machinery.

Hightower, J.; Consolvo, S.; Lamarca, A.; Smith, I. E.; and Hughes, J. 2005. Learning and Recognizing the Places We Go. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing,* 159–176. New York: Association for Computing Machinery.

Hu, P., and Reuscher, T. 2004. Summary of Travel Trends, 2001 National Household Travel Survey. Technical Report, U.S.F.H.A. U.S. Department of Transportation. Washington, DC: U.S. Department of Transportation.

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Kang, J. H.; Welbourne, W.; Stewart, B.; and Borriello, G. 2005. Extracting Places from Traces of Locations. *ACM Sigmobile Mobile Computing Communication Review* 9(3): 58–68.

Kim, D. H.; Hightower, J.; Govindan, R.; and Estrin, D. 2009. Discovering Semantically Meaningful Places from Pervasive RF-Beacons. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing,* 21–30. New York: Association for Computing Machinery.

Kim, D. H.; Kim, Y.; Estrin, D.; and Srivastava, M. B. 2010. Sensloc: Sensing Everyday Places and Paths Using Less Energy. In *Proceedings of the ACM Conference on Embedded Networked Sensor Systems,* 315–330. New York: Association for Computing Machinery.

Krumm, J., and Horvitz, E. 2007. Predestination: Where Do You Want To Go Today? *IEEE Computer* 40(4): 105–107.

Krumm, J. 2011. Ubiquitous Advertising: The Killer Application for the 21st Century. *IEEE Pervasive Computing* 10(1): 66–73.

Laasonen, K.; Raento, M.; and Toivonen, H. 2004. Adaptive On-Device Location Recognition. In *Pervasive,* ed A. Ferscha and F. Mattern, 287–304. Berlin: Springer.

Liao, L.; Fox, D.; and Kautz, H. 2004. Learning and Inferring Transportation Routines. In *Proceedings of the 19th National Conference on Artificial Intelligence,* 348–353. Menlo Park, CA: AAAI Press.

Lin, K.; Kansal, A.; Lymberopoulos, D.; and Zhao, F. 2010. Energy-Accuracy Trade-Off for Continuous Mobile Device Location. In *Proceedings of the 8th Annual International Conference on Mobile Systems, Applications, and Services*. New York: Association for Computing Machinery.

Lin, F. X.; Wang, Z.; Likamwa, R.; and Zhong, L. 2012. Reflex: Using Low-Power Processors In Smartphones Without Knowing Them. In *Proceedings of the 17th International Conference on Architectural Support for Programming Languages and Operating Systems,* 13–24. New York: Association for Computing Machinery.

Mohan, P.; Padmanabhan, V. N.; and Ramjee, R. 2008. Nericell: Rich Monitoring of Road and Traffic Conditions Using Mobile Smartphones. In *Proceedings of the ACM Conference on Embedded Networked Sensor Systems,* 323–336. New York: Association for Computing Machinery.

Paek, J.; Kim, J.; and Govindan, R. 2010. Energy-Efficient Rate-Adaptive GPS-Based Positioning for Smartphones. In *Proceedings of the 8th Annual International Conference on Mobile Systems, Applications and Services*. New York: Association for Computing Machinery.

Paek, J.; Kim, K.-H.; Singh, J. P.; and Govindan, R. 2011. Energy-Efficient Positioning for Smartphones Using Cell ID Sequence Matching. In *Proceedings of the 8th Annual International Conference on Mobile Systems, Applications and Services*, 293–306. New York: Association for Computing Machinery.

Park, J.-G.; Charrow, B.; Curtis, D.; Battat, J.; Minkov, E.; Hicks, J.; Teller, S.; and Ledlie, J. 2010. Growing an Organic Indoor Location System. In *Proceedings of the 8th Annual International Conference on Mobile Systems, Applications, and Services*. New York: Association for Computing Machinery.

Patterson, D. J.; Liao, L.; Fox, D.; and Kautz, H. 2003. Inferring High-Level Behavior from Low-Level Sensors. In

*Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing,* 73–89. New York: Association for Computing Machinery.

Peebles, D.; Lu, H.; Lane, N.; Choudhury, T.; and Campbell, A. 2010. Community-Guided Learning: Exploiting Mobile Sensor Users to Model Human Behavior. In *Proceedings of the 25th National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

Thiagarajan, A.; Ravindranath, L.; Balakrishnan, H.; Madden, S.; and Girod, L. 2011. Accurate, Low-Energy Trajectory Mapping for Mobile Devices. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation,* 20–20. Berkeley, CA: Usenix Association.

Yang, G. 2009. Discovering Significant Places from Mobile Phones: A Mass Market Solution. In *Proceedings of the Second International Conference on Mobile Entity Localization and Tracking in GPS-less Environments,* 34–49. Berlin: Springer-Verlag.

Yin, J.; Chai, X.; and Yang, Q. 2004. High-Level Goal Recognition in a Wireless Lan. In *Proceedings of the 19th National Conference on Artificial Intelligence,* 578–583. Menlo Park, CA: AAAI Press.

Yoon, H.; Zheng, Y.; Xie, X.; and Woo, W. 2010. Smart Itinerary Recommendation Based on User-Generated GPS Trajectories. In *Proceedings of the 8th Ubiquitous Intelligence and Computing International Conference,* 19–34. Berlin: Springer-Verlag.

Yuan, J.; Zheng, Y.; Zhang, C.; Xie, W.; Xie, X.; Sun, G.; and Huang, Y. 2010. T-Drive: Driving Directions Based on Taxi Trajectories. In *Proceedings of the 18th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems,* 99–108. New York: Association for Computing Machinery.

Zheng, Y.; Li, Q.; Chen, Y.; Xie, X.; and Ma, W.-Y. 2008. Understanding Mobility Based on GPS Data. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing,* 312–321. New York: Association for Computing Machinery.

Zheng, Y.; Zhang, L.; Xie, X.; and Ma, W.-Y. 2009. Mining Interesting Locations and Travel Sequences from GPS Trajectories. In *Proceedings of the 18th International Conference on World Wide Web,* 791–800. New York: Association for Computing Machinery.

Zheng, V. W.; Cao, B.; Zheng, Y.; Xie, X.; and Yang, Q. 2010. Collaborative Filtering Meets Mobile Recommendation: A User-Centered Approach. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence.* Menlo Park, CA: Press.

Zhuang, Z.; Kim, K.-H.; and Singh, J. P. 2010. Improving Energy Efficiency of Location Sensing on Smartphones. In *Proceedings of the 8th Annual International Conference on Mobile Systems, Applications and Services*. New York: Association for Computing Machinery.

**Yifei Jiang** is a Ph.D. candidate in the Department of Computer Science at the University of Colorado, Boulder. He received his M.S. degree from University of Colorado, Boulder in 2009 and his B.S. degree from Beijing University of Technology. His research interests span intelligent mobile systems, mobile context sensing, and social networks.

**Du Li** is a principal researcher at Ericsson Research, San Jose, California. He received his B.S. degree from Wuhan University in 1992, M.S. degree from Peking University in 1995, and Ph.D. degree from UCLA in 2000, all in computer science. He was on the faculty of Texas A&M University from 2000 to 2007 and won a CAREER award from the National Science Foundation in 2002. Before joining Ericsson, he worked for Nokia Research, Palo Alto from 2007 to 2011. His research experiences span collaborative systems, mobile and ubiquitous computing, internet of things, and human-computer interaction.

**Qin Lv** is an assistant professor of computer science at the University of Colorado, Boulder. She received her B.E. degree with honors from the Department of Computer Science and Technology, Tsinghua University in 2000, her M.A. and Ph.D. degrees from the Department of Computer Science, Princeton University in 2002 and 2006. Before joining the University of Colorado, she also spent one year at Princeton University as a postdoc, and one year in the Computer Science Department, Stony Brook University as an assistant professor. Her research integrates efficient system design with effective data analysis for the management and exploration of massive data. Her work has won the Computational Sustainability Award by the Computing Community Consortium (CCC) and Pervasive 2012, and the Best Paper Award nomination at ISLPED 2010. She has served on the technical program committees of SIGMETRICS, UbiComp, PerCom, and others. Her work has been cited more than 2000 times.