IEEE TRANSACTIONS ON
EMERGING TOPICS
IN COMPUTING

Received 2 December 2013; revised 25 March 2014; accepted 18 June 2014. Date of publication 8 July 2014;
date of current version 6 December 2017.

Digital Object Identifier 10.1109/TETC.2014.2335537

How Well Can a User's Location Privacy Preferences be Determined Without Using GPS Location Data?

TAKUYA MAEKAWA[1], NAOMI YAMASHITA[2], AND YASUSHI SAKURAI[3]

[1]Graduate School of Information Science and Technology, Osaka University, Osaka 565-0871, Japan
[2]NTT Communication Science Laboratories, Kyoto 619-0237, Japan
[3]Graduate School of Science and Technology, Kumamoto University, Kumamoto 860-0862, Japan
CORRESPONDING AUTHOR: T. MAEKAWA (maekawa@ist.osaka-u.ac.jp)

ABSTRACT The recent proliferation of GPS-enabled mobile phones has allowed people to share their current locations with others. Because disclosing one's location can be valuable but risky, many services and studies employ a user's GPS coordinates to determine automatically whether or not those coordinates can be disclosed by comparing them with handcrafted rules or privacy models trained using the user's actual preferences. However, these approaches that employ GPS coordinates constitute a drain on a phone's battery when the services assume continuous location sharing. In addition, recent positioning methods (assisted GPS and a WiFi-based positioning) rely on external location providers. That is, when a user's current location preference is determined using her coordinate point, her location information is disclosed to external providers even if this is not her wish. In this paper, we explore a way of learning a user's location privacy preference using sensors that are energy saving and that do not rely on external providers. This enables us to save energy and protect a user's privacy when she is unwilling to disclose her location. Note that the machine learning-based approach cannot deal well with a user's private situations that are not included in its training data. So, this paper proposes a new model that can determine a user's privacy preferences and handle such outlying situations.

INDEX TERMS J.9.a location-dependent and sensitive, J.9.d pervasive computing, K.4.1.f privacy.

I. INTRODUCTION

The recent proliferation of GPS-enabled mobile phones has enabled people to share their current locations with others. Various kinds of location sharing services including foursquare, Facebook, and Gowalla have taken advantage of such location information. However, concerns have been expressed about the privacy implications associated with location sharing [1], [2]. Therefore, a user should manually input her privacy preferences (determine whether or not her current location is disclosed) to protect her privacy information. However, because many of recent location sharing services and studies assume continuous location sharing, inputting her preferences manually places large burdens on her. (For example, the user should input her privacy preferences whenever she moves or at regular intervals.) Therefore, many studies attempt to automatically determine the user's privacy preferences to reduce the burdens. These studies employ the following two methods.

- *Static privacy settings*: Handcrafted rules are used to determine whether a user's location is disclosed or not. The user specifies rules that include location-based restrictions (e.g., my location is not disclosed when I am at home) in advance.
- *Machine learning*: Because people's privacy preferences related to context disclosure are diverse and complex, several studies have argued that privacy management should be a dynamic response to circumstances rather than automatic management based on static rules. Recently, several studies have tried to predict a user's privacy preferences by employing machine learning approaches [3], [4]. That is, these approaches obtain the user's teaching signals (whether she is willing to disclose her location or not) at several locations, and then construct her privacy preference model by using the training data. This paper also focuses on this approach.

The above two methods basically compare a user's current location (obtained with GPS) with specific locations,

e.g., home, included in handcrafted rules or training data. However, these approaches have the following problems.
1) There are many services and studies that assume continuous location sharing (e.g., Google Latitude and surveillance of remote family members). However, because a user's mobile phone should regularly measure the user's location with its GPS sensor to determine whether the location is disclosed or not, the phone's battery is rapidly exhausted.
2) Recent positioning methods such as A-GPS and WiFi-based localization require communication with external location providers. In such cases, even when a user does not want to disclose her location, her location(-related) information is disclosed to the external location providers if her location data are used to predict her location privacy preferences. The leaking of sensitive private information including location information by external service providers is frequently discussed in the pervasive computing community [5], [6]. When we determine whether or not a user's private information is disclosed, we should not rely on external providers.

The above facts motivated us to find a new way of predicting a user's privacy preferences. We assume that, when a person uses a location sharing system, the person is suffering from two types of privacy problems: (1) her location is disclosed to her acquaintance against her will and (2) her location is disclosed to a third party (location provider) when she wants to protect her location information. To solve these problems, we should correctly predict the person's privacy preferences without relying on the third party, i.e., without using GPS (or WiFi) coordinates. Our solution involves capturing features of her current location that may reflect her current privacy preferences (other than its coordinates) by using sensors on a mobile phone that are power-saving and that do not rely on external providers. For example, when people are in public places or attending social activities, they are reportedly willing to disclose their locations [7]. We consider that the microphone installed in the mobile phone may be used to capture features of private and public places. Also, people are reportedly willing to disclose their locations when they are exercising [7], but less willing when they are sleeping or resting. Such user activity information can be obtained from an accelerometer installed on the phone. In this paper, we focus particularly on a scanned WiFi signal (i.e., the unique MAC address of an access point and the signal strength from the AP) that may be useful for understanding the user's location privacy preference, and investigate how we infer the preference by analyzing WiFi signals. For example, we can simply model the user's preference by using MAC addresses of APs observed when the user did not disclose her location in the past. Also, we can easily know whether the user is at a private place (e.g., house) or not by using a MAC address of an AP at her house. In this paper, we investigate what kinds of location privacy-related features such energy-saving sensors can and cannot capture, and compare their performance with that of a GPS sensor. Our approach, which does not employ GPS coordinates, has the advantages described below.

- The participants in a survey undertaken by Khalil and Connelly [8] were unwilling to disclose their locations around 40%, 40%, 60%, and 80% of the time to their family members, friends, colleagues, and bosses, respectively. Therefore, we can greatly reduce battery consumption related to location sharing because we measure a user's location with a GPS sensor only when she is estimated to be willing to disclose her location. Fig. 1 shows phone battery drainage time series. (Loc: Measuring coordinates with GPS and WiFi. Acc: acceleration data sensing. MIC: sound sensing. WiFi: WiFi scanning. Acc+MIC+WiFi: acceleration sensing, sound sensing, and WiFi scanning.) When the locations of the phone were measured using GPS (Loc), the GPS data were obtained at about 0.01 Hz. Note that when the phone was indoors, its locations were measured using WiFi. The WiFi line in Fig. 1 shows a phone battery drainage time series when the WiFi data were obtained at about 0.01 Hz. That is, we assumed continuous location sharing and collected sensor data based on the application. Also, the acceleration data were obtained at about 16 Hz. Sound data were recorded by a microphone at 8 kHz. See section III for more detail about data sampling settings. As shown in Fig. 1, the power consumption of a GPS sensor is much higher than that of other sensors. The GPS power consumption was also higher than that of the WiFi, accelerometer, and microphone combined. In addition, Chon et al. [9] reported that the power consumption of GPS is about four times higher than that of WiFi when the sampling interval is one minute.

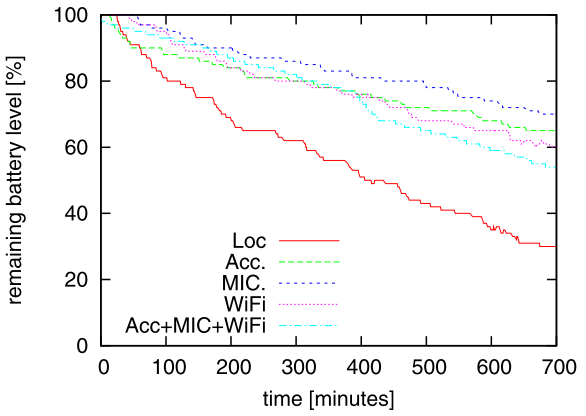

FIGURE 1. Mobile phone power consumption during continuous sensing.

- A user's location information is not disclosed to a third party against the user's will because our approach predicts her privacy preference by using only sensors that do not rely on external providers.

As described above, we attempt to construct a machine learning based model that determines a user's privacy preferences by using power-saving sensors. Note that the ML-based approach cannot deal well with situations that are not included in the training data. This is a fatal defect in

systems that deal with a user's private information. We explain this in Fig. 2, which shows a feature space including data points corresponding to the user's context (sensor data). The circles and squares respectively show data points corresponding to *disclose* and *not-disclose* classes included in training data. A discriminative classifier classifies a test data point into a probable class by using a hyperplane H computed from the training data points. The filled squares show test data points corresponding to *not-disclose*. The classifier can correctly classify point A that is close to the *not-disclose* training points. It can also correctly classify outlying point B because the point happens to be on the *not-disclose* side of the hyperplane even though it is far from the training points. However, the classifier cannot correctly classify outlying point C that is far from the training points and happens to be on the *disclose* side. This is because the ML-based classifier cannot compute a hyperplane by taking unseen data points (situations) into consideration. To cope with the problem, we propose a probability density based classifier that not only classifies test data points but also finds outlying data points in an unified framework.

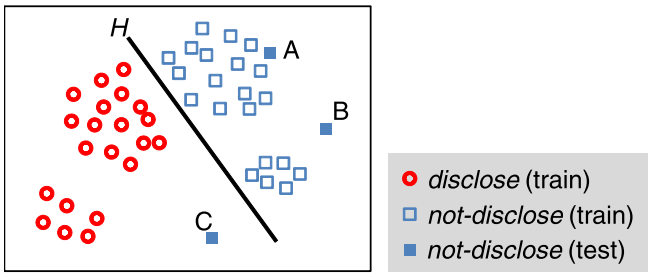

FIGURE 2. Handling feature vectors with discriminative classifier.

In the rest of this paper, we first introduce work related to privacy in context-aware systems. Next we describe our experimental study, in which we collected sensor data from the participants. We then explain our method for learning a user's privacy preferences, and employ 2753 hours of *in situ* sensor data obtained from the participants to devise an approach that allows mobile phone sensors to capture location privacy preferences. The technical contributions of our work are as follows: (1) We investigate how to extract sensor data features that are reflected by a user's privacy preferences with energy-saving sensors (esp. WiFi modules). (2) We propose a classifier that can find a user's outlying situations in addition to identifying a user's privacy preferences.

II. RELATED WORK

A. PRIVACY IN LOCATION-BASED SERVICES

Location sharing applications raise many privacy concerns for users. To cope with such concerns, many studies and services have employed static privacy management based on handcrafted rules [10], [11]. For example, Hull et al. realized privacy-conscious user data sharing by using handcrafted rules that define a user's privacy preferences based on sensor data obtained from her mobile device [10] (e.g., a user's data are not disclosed when she is at places within 3 kms of her home). While these services and studies are successful to some extent, several studies have argued that privacy management should be a dynamic response to circumstances rather than automatic management based on static rules because people's lives are not predictable [12], [13].

Several recent studies have investigated the possibility of realizing robust machine learning-based privacy management. For example, Sadeh et al. [3] implemented a web-based location sharing system that employs handcrafted rules to decide whether or not a user's current location is disclosed. One distinctive feature of the system is that it employs case-based reasoning and allows the user to correct rule-based decisions. Such input from the user could be utilized when the system makes a decision about a new location. When the user was usually willing to disclose her location at places close to the new location, the system decides to disclose the new location. Bigwood et al. [4] constructed a user's location privacy preference model by using several features extracted from GPS data, e.g., past privacy preferences obtained close to the current location by using the experience sampling method (ESM) [14] and type of current location (university, library, restaurant, etc.) inferred from GPS data and such location databases as Yell and Google Maps. (They also employ other features, e.g., those extracted from friend list data of social network services.)

In contrast, we try to learn a user's privacy preference model by using sensor data other than GPS location coordinates. Extensive research has been conducted with the aim of addressing people's privacy preferences, and several studies have investigated empirically the factors that affect people's willingness to disclose their locations [7], [8], [15]. For example, when people attend social activities, they are willing to disclose their contexts to friends. Also, people's location privacy preferences differ depending on whether they are at home or the workplace. We try to capture such features without using her current coordinates.

B. MOBILE PHONE SENSING TECHNOLOGIES

Recently mobile phones have been equipped with various kinds of sensors, and we can benefit from the context information that they provide. In particular, latitude and longitude information obtained from a GPS sensor is widely used in various applications. However, GPS-based positioning can not be deployed for indoor use because line-of-sight transmission between the GPS sensor and satellites is not possible in indoor environments. To determine positions indoors, WiFi modules incorporated in mobile phones are usually used. Many services and studies employ the fingerprint-based WiFi positioning method [16], [17]. The fingerprinting method refers to techniques that match the fingerprint of location-dependent characteristics, i.e., the received signal strength from a WiFi access point (AP). The fingerprints obtained at different locations are stored in a database associated with their coordinates, and compared with the current fingerprints to determine the user's current coordinates. That is, the

WiFi-based method discloses their exact coordinates to an external service provider. In our experiment (section III), our participants stayed outside of the range of GPS 88.5% of their daily lives. In such cases, people's coordinates are measured by using the WiFi-based fingerprinting method, and their exact coordinates were disclosed to the provider.

A scanned WiFi signal (i.e., the unique MAC address of an AP and the signal strength from the AP) is also useful for understanding the user's location context [18]. For example, it is easy to determine that the user is at home if the scanned MAC address matches that of an AP in her home. We employ the raw WiFi signals instead of GPS and WiFi coordinates. In addition, in many studies various kinds of context have been obtained from sensors mounted on mobile phones. Several studies have recognized a user's daily activities such as walking, sitting, and bicycling by employing an accelerometer on her mobile phone [19], [20]. Also, the system proposed in [21] recognizes sound events such as speech, music, and ambient sound from a microphone installed in a mobile phone. The system proposed in [22] attempts to identify the person a user is talking to by employing a mobile phone microphone and special hardware connected to it.

III. DATA COLLECTION

This section describes how we collect labeled sensor data from experimental participants to construct their privacy models. As in previous studies described in section II, we employed the ESM to obtain a user's actual privacy preferences. The ESM was originally used to intentionally interrupt an experimental participant in order to have her make notes about her current situation. In our approach, an application on the user's mobile phone regularly asks her whether she is currently willing or unwilling to disclose her current location to any given person (e.g., family and friends.). At the same time, the application collects sensor data from sensors on the phone. This approach enables us to obtain training data that consist of the user's preference (willingness to disclose her location) and its corresponding sensor data at a certain time. We then learn the relationship between the user's willingness to disclose her location and the sensor data features obtained at the same time.

A. ASSUMED ENVIRONMENT

A user (participant) carries a mobile phone that is equipped with sensors. We ask the participant to assume that her current location can be shared with other persons in her buddy list on a hypothetical application. The people in her buddy list can see her current location by using a web browser. In the application, an avatar corresponding to the participant is placed at her current coordinates on a Google Maps type interface. The people in her buddy list are categorized into several groups; significant other (SO), family members, friends, colleagues, and boss. We selected these categories based on settings used in previous studies on location disclosure [8]. We asked the participant to assume such an application.

B. DATA COLLECTION METHODOLOGY

Our implemented application, which runs on a mobile phone (Google Nexus One), collects and stores the following sensor data while the phone is on.

- Three-axis acceleration data obtained from an accelerometer at about 16 Hz.
- Sound recorded by a microphone at 8 kHz. It is expensive to continuously record sound at a high sampling rate, and so the application intermittently captures short periods of sound and then stores only features extracted from the sound. See the section IV-A for more detail.
- WiFi scan data (received signal strengths and MAC addresses of APs) obtained from a WiFi module at about 0.01 Hz.
- The user's current latitude and longitude obtained from GPS and WiFi using Android APIs at about 0.01 Hz. If her speed (obtained from a GPS sensor) exceeds 0, the sensor data are sampled at about 0.2 Hz.

The application also issues an ESM question about once a hour on average. The time intervals between the questions are randomly changed (50-70 minutes) for each question. When a question is issued, the mobile phone rings (or vibrates when it is in silent mode), and then the question is shown. The question message is very simple. ''Please select the categories of person to which you are unwilling to disclose your current location at the time the phone rang.'' A user selects any of five categories (SO, family members, friends, colleagues, and boss). (If the user does not have a significant other, the SO category is not shown.) The user can also ignore the question. After selecting categories, the user inputs the activity she was performing when the phone rang and why she was unwilling to disclose her location. This information is used to investigate the answers in detail. The ESM answer is saved in association with a timestamp that was obtained just before the phone rang.

C. PARTICIPANTS

We collected labeled sensor data from fifteen participants (eleven males and four females) consisting of four workers and two researchers in our laboratories, three of our family members, and six graduate students. We asked each participant to carry a mobile phone equipped with our application for about couple of weeks (15 to 25 days). We also asked the participants to turn the phone on when they woke up and turn it off when they went to bed. When answering the ESM questions, we asked each participant to assume actual individuals for each category of person. Because the participants knew that the location sharing application was hypothetical, they may have been less concerned about their privacy. However, employing the ESM to capture the participants' privacy preferences in real life and asking the participants to assume actual individuals are expected to minimize this bias. Before the data collection, each participant completed a demographic questionnaire including her home and workplace addresses, MAC addresses of APs at her home and workplace, and whether or not she had a spouse (or boyfriend/girlfriend).

D. DATA OVERVIEW

Here we introduce the collected ESM answers and sensor data. Note that the goal of this study is to learn people's privacy preferences. Previous surveillance studies have provided detailed investigations of people's privacy preferences [7], [8], [15]. We collected 2753 hours of sensor data and 2280 ESM answers from the participants. The participants answered about 78.8% of the questions that they received hourly. Reasons for failing to answer a question included showering, doing housework, and taking a nap. In 64.3% of the answers, the participants were willing to disclose their location to all categories of person.

The participants gave various reasons in their ESM answers for being unwilling to disclose their locations. Naturally, the participants were unwilling to disclose their locations when doing something that they wanted to keep secret. For example, a participant was unwilling to disclose his location to his wife when he went shopping without telling her. Also, when a participant walked to a station to save on bus fare, he was unwilling to disclose the location to his boss. Many of these cases were related to the participants' non-routine (outlying) activities. By contrast, many participants routinely protected their locations from certain people, for example their bosses outside business hours. However, even for such particular participants, their privacy preferences were not always static. For example, when a participant's work time was longer than usual due to a business trip, he disclosed this location to his colleagues and boss until late at night even though he was usually unwilling to disclose location to them after 7 p.m. We investigated all the ESM answers carefully and counted the number that were obtained in such outlying situations. Here, when a participant encountered situations different from the ordinary and changed his/her privacy preference, we define the obtained answer to be an outlier. We found that 5.8% of ESM answers were outliers.

Here we introduce the relationship between the ESM answers and the measured GPS (with WiFi) coordinates. People's privacy preferences depend greatly on whether they are at home or at their workplaces [7], [8]. Fig. 3 shows the relationship between the distance from the participants' home and the disclosure rate (the percentages of cases where the participants were willing to disclose their locations to all categories of person). A disclosure rate at x kms on the graph shows the disclosure rate of ESM answers given at places within x kms of the participant's home. As shown in the graph, the participants were more willing to disclose their locations the further away they were from their homes. Fig. 3 also shows the relationship between the distance from their workplaces and the disclosure rate. The participants were more willing to disclose their locations the closer they were to their workplaces.

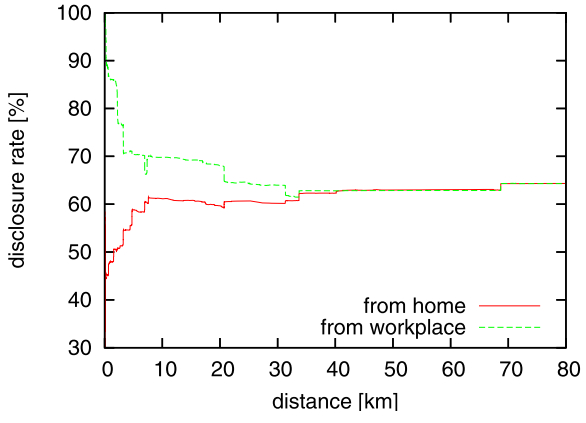

FIGURE 3. Relationship between disclosure rate and distance from home/workplace.

The user's privacy preference at her location of interest may strongly relate to her past privacy preferences at that location. Fig. 4 shows the relationship between the disclosure rate and distance between the user's location of interest and the closest *disclose* answer obtained in the past. A disclosure rate at x kms on the graph shows the disclosure rate of ESM answers whose past closest *disclose* answer is located within x kms of the ESM answer. We consider that, if the user disclosed her locations many times at places close to the location of interest, she may want to disclose her location of interest. As shown in the graph, a user's past privacy preferences at the same location are a strong clue as regards predicting the user's current privacy preferences. As above, the participants' location privacy strongly related to their location coordinates. In the next section, we clarify relationship between participants' location privacy and sensor data obtained from other sensors.

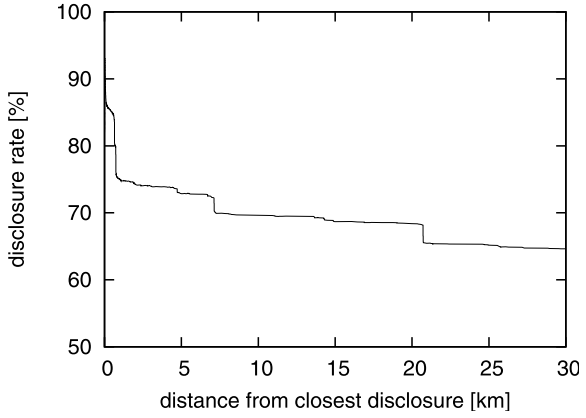

FIGURE 4. Relationship between disclosure rate and distance from the closest *disclose* answer.

IV. PREDICTING USER'S PRIVACY PREFERENCES

We explain how we construct a user's privacy preference model from sensor data and ESM answers obtained from that user. We first extract features from sensor data obtained just before the user answered each ESM question, and construct a feature vector that concatenates the extracted features for each answer. Then we learn the privacy preference model by employing pairs consisting of an extracted feature vector and its corresponding ESM answer. That is, we learn a model that classifies a feature vector in the *disclose* or *not-disclose* class.

Note that we construct the user's privacy preference model for each category of person by using answers related to the category because people's privacy preferences differ depending on to whom they disclose their location [7], [8]. In the following, we explain our approach in detail.

A. FEATURE EXTRACTION

We construct a feature vector for each ESM answer that concatenates features extracted from the following sensor data collected just before the ESM answer was obtained.

1) WiFi SCAN

As shown in Figs. 3 and 4, we can easily compute meaningful features from the distance between a current GPS coordinate point and a specific coordinate point (e.g., home or place where a user disclosed her coordinate point in the past). However, it is not easy to compute such features from a current WiFi scan and a specific scan (e.g, a scan that was obtained when a user disclosed her location) because WiFi scan data obtained at two different places are completely different even if the distance between the two places is only about 500 meters. Fig. 5 shows an example. Assume that a user disclosed her location when scan 1 was obtained. We may find a correlation and define the distance between scans 1 and 2 because both these scans include AP2. However, how do we find a correlation and compute the distance between scan 1 and scan 3 (or 4)? Also, the distance between scans 1 and 3 should be defined as being shorter than that between scans 1 and 4 because people's privacy preferences change according to the distance from a specific place as shown in Figs. 3 and 4. Here we assume that an end user does not have WiFi AP databases that list AP locations created by location providers (e.g., Skyhook and Google) because the databases are not public and very expensive. So, we solve the problem solely by using WiFi scans obtained from the user. Our solution for the problem is to construct an undirected graph (from the user's scans) whose vertices correspond to APs and the path length between vertices reflects the distance between the corresponding APs. The lower portion of Fig. 5 shows an example graph created from the four scans. In the graph, two vertices (APs) that are included in one scan are connected with an edge. We construct the graph from a user's WiFi scan histories by using Algorithm 1. The graph enables us to compute the distance between two APs by using Dijkstra's algorithm. For example, in Fig. 5, the distance between APs 1 and 3 becomes 2 and the distance between APs 1 and 5 becomes 3. (If we cannot find any path between two APs, we define its distance as d_{max}.) We can also compute the distance between two WiFi scans, e.g., scans 1 and 2. We simply define distance as the shortest distance between any pairs consisting of an AP in scan 1 and an AP in scan 2. The following features are computed from a user's scanned WiFi signal.

Algorithm 1 Graph Construction From WiFi Scans

Input: WiFi scans S
Output: $G(V, E)$, where each vertex $ap \in V$ corresponds to each AP

```
G = ∅
for each scan s ∈ S do
    for each AP ap in s do
        if ap ∉ V then
            Add vertex ap to G
        end if
    end for
    for each pair of APs (ap_i, ap_j) in s do
        if edge (ap_i, ap_j) ∉ E then
            Add edge (ap_i, ap_j) to G
        end if
    end for
end for
```

- Distances from home and workplace: We compute the distance between an AP at the user's home (workplace) and each AP in the WiFi scan of interest. The feature corresponds to the shortest distance. Fig. 6 shows the relationship between the distance from the participants' home (workplace) and the

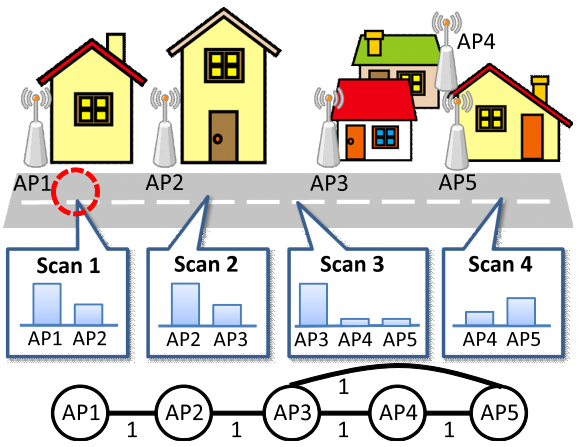

FIGURE 5. **Example of WiFi scans obtained at nearby places.**

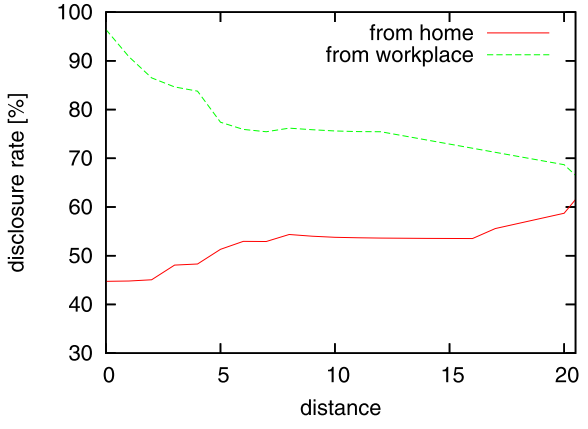

FIGURE 6. **Relationship between disclosure rate and distance from an AP at home/workplace.**

disclosure rate. We could extract features similar to those in Fig. 3 by using a graph made from participants' scans.

- Privacy preferences in the past: The user's privacy preference at her location of interest may strongly relate to her past privacy preferences at places close to that location. Before computing this feature value, we find the k-nearest ESM answers (WiFi scans) from the WiFi scan of interest. The feature value corresponds to the ratio of the number of *not-disclose* answers among the k answers ($k = 5, 10, 30$. We compute a feature value for each k.).

- Distance from closest *disclose* answer: We compute the distance between the WiFi scan of interest and that corresponding to each past *disclose* answer. The feature value corresponds to the shortest distance. This feature shows how close the WiFi scan (APs) is to that corresponding to a past *disclose* answer. Fig. 7 shows the relationship between the disclosure rate and the distance between the user's scan of interest and the closest *disclose* answer obtained in the past. We could also compute features relatively similar to those computed from GPS coordinates.

- Distance from closest *not-disclose* answer: We compute the distance between the WiFi scan of interest and that corresponding to each past *not-disclose* answer. The feature value corresponds to the shortest distance.

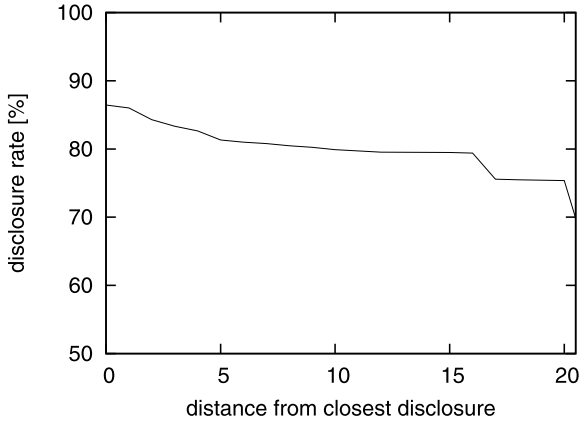

FIGURE 7. Relationship between disclosure rate and distance from the closest *disclose* scan.

- Difference between distances from closest *disclose* and *not-disclose* answers: The feature reveals whether the scanned WiFi signal is similar to that corresponding to a past *disclose* answer or *not-disclose* answer.

- Frequency of visiting AP: When an AP is detected, we can assume that the user visits the location at which the AP is installed. The frequency of a user's visit to a particular place may reflect her familiarity with the place. For example, the frequency related to her home may be high. On the other hand, the frequency related to a bar at which the user once stopped off after work will be low. For each AP included in the scanned signal, we compute the ratio of the number of past scans in which the AP was detected to the total number of scans. We employ the average ratio over the APs in the scan of interest as a feature. Fig. 8 shows the relationship between the frequency of visiting APs and the disclosure rate. A disclosure rate at x % frequency in the graph shows the disclosure rate of the ESM answers given at APs whose visit frequencies are lower than x %. As shown in the graph, when the frequency is small ($\leq$ 5%), the disclosure rate appears to be relatively low. Locations with a low frequency included restaurants and shops that the participants visited once. Locations with around 15% frequencies corresponded to their homes. Locations with a high frequency ($\geq$ 25%) corresponded to their workplaces.

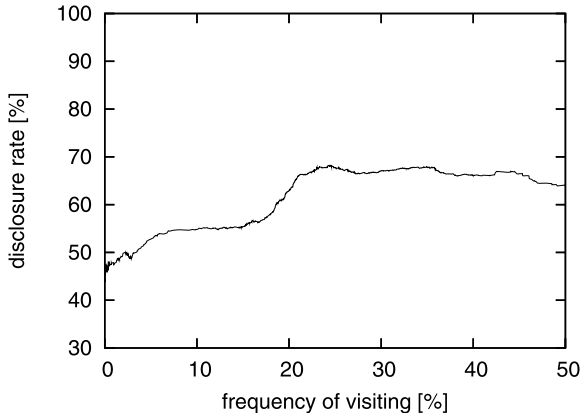

FIGURE 8. Relationship between disclosure rate and frequency of visiting APs.

- Entropy of AP (time of day): Visiting a location at a given time of day also reveals the characteristics of the location at which the AP is placed. For example, people spend most of each day at their homes and workplaces. However, people may visit their work canteens only at lunch time. We employ entropy as a feature for measuring the diversity of visits to AP ap by a user for every hour given by $H = -\sum_{t=0}^{23} p(ap, t) \log p(ap, t)$, where $p(ap, t)$ is the ratio of the number of visits to ap at t o'clock to the total number of visits to ap. Note that if multiple APs are included in the scan, we employ the average entropy over the APs as a feature. Fig. 9 shows the relationship between the entropy of an AP and the disclosure rate. The disclosure rates for APs with a small entropy were low. This indicates that the disclosure rates were low for locations where the participants visited only at a certain time period. These locations included shops and bars that the participants regularly visited after work and a hospital that one participant visited weekly at a scheduled time. On the other hand, the disclosure rates for APs with a large entropy (around 2.5) were very high. These APs correspond to those located at workplaces.

- Entropy of AP (day of the week): We also employ the diversity of visiting an AP on each day of the week. For example, a person may go shopping only on Saturday, and so the entropy of the shopping place becomes small.

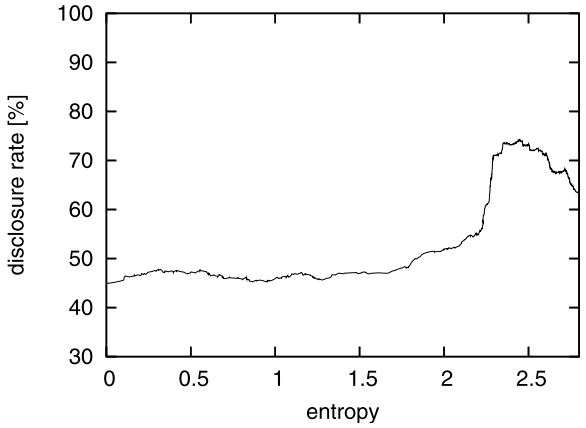

FIGURE 9. Relationship between disclosure rate and entropy of AP (time of day).

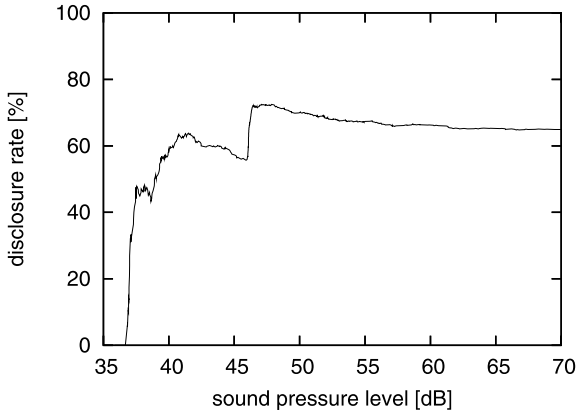

FIGURE 10. Relationship between disclosure rate and sound pressure level.

2) ACCELERATION DATA

A user's activity, e.g., walking, running, or resting, is reflected in her acceleration data [19]. Note that an acceleration signal in three orthogonal directions (X, Y, and Z) might be sensitive to the mobile phone placement, e.g., in the pants pocket or breast pocket. To compute features, we use the combined signal given by $R_i = \arcsin(\frac{Z_i}{\sqrt{X_i^2 + Y_i^2 + Z_i^2}})$ proposed in [23], where R_i is the ith combined signal. The following features are computed from the user's acceleration data collected just before the ESM answer was obtained.

- Variance: This feature value corresponds to the variance of the k-second combined acceleration data ($k = 5, 10, 30$). This feature shows the intensity of the mobile phone movement.

- Mean: This feature value corresponds to the mean of the k-second combined data ($k = 5, 10, 30$). This feature shows the posture of the mobile phone.

3) SOUND

Ambient sound tells us about the properties of places [20], [21]. When people are in public places, they are reportedly willing to disclose their locations [7]. Fig. 10 shows the relationship between the disclosure rate and the sound pressure level of ambient sound obtained from microphones. As shown in the graph, when people were in very silent places, they were unwilling to disclose their locations. These situations corresponded to sleeping and resting at their homes. As the volume increased, their disclosure rate also increased. The loud places included such public places as inside a train, a bar, and a shop (workplace of one participant), and the participants were willing to disclose their locations when they were at these places. In addition, we can obtain information about the user's activity from sound. For example, the sound of a vehicle, conversation, and running water may relate to a user's activities. In this study, we obtained approximately 50 milliseconds of sound recording per second. We computed the following features by employing k short sound recordings obtained just before we collected the ESM answer ($k = 20$).

- Sound pressure level: We compute the average and variance of the sound pressure level over k recordings, and use them as features.

- MFCC components: In [24], the Mel-Frequency Cepstral Coefficient (MFCC) is reported to be the best transformation scheme for environmental sound recognition. Also, Chen et al. achieved the highly accurate recognition of bathroom activities such as showering, flushing, and urination by using the MFCC [25]. We compute a 13-order MFCC of each short sound recording windowed by a Hamming window. We compute the average and variance values for each MFCC component over k recordings, and use them as features.

- Zero-crossings and auto correlation: Whether the user is alone or with someone is reported to affect her willingness to disclose her location [8], [15]. To determine when she is in conversation with someone, we employ the sound features that are usually used for speech detection [26].

Note that microphone data cannot provide us with an end user's precise privacy information. As mentioned above, people in public places are reportedly willing to disclose their locations. However, when a person visits a public place such as a hospital for health care related issues, for example, he/she may not want to disclose his/her location. Therefore, we consider that such data as microphone and acceleration data are supplementary information for the WiFi based privacy preference estimation.

4) WiFi INDOOR POSITIONING

Indoor positioning studies [27] attempt to estimate the precise indoor coordinates of a phone by using signal strength data based on fingerprinting techniques. Fingerprinting employs a training phase in which WiFi signals (i.e., the unique MAC addresses of APs and the received signal strengths from APs) are observed at known coordinates. A set of APs and their signal strengths become a fingerprint that is

unique to those coordinates. The fingerprints are stored in a fingerprint database on a server. In the positioning (test) phase, the observed WiFi signals at unknown coordinates (test points) are compared with the stored fingerprints in the fingerprint database on the server to determine the closest match. Many indoor positioning studies compute the distance between WiFi scans in order to find the closest match. We consider that we can employ the fingerprinting techniques for fine-grained privacy management. In this study, we compute the distance between scans S_1 and S_2 as follows.

$$dist(S_1, S_2) = \frac{1}{|A|} \sum_{i=1}^{|A|} |S_{1,i} - S_{2,i}|,$$

where A is a set of APs included in both S_1 and S_2 and $S_{1,i}$ is the received signal strength from the ith AP included in S_1. With the distance, the following features are computed.

- Distance from closest *disclose* answer: We compute the distance between the WiFi scan of interest and that corresponding to each past *disclose* answer. The feature value corresponds to the shortest distance. This feature shows how close the WiFi scan (APs) is to that corresponding to a past *disclose* answer.

- Distance from closest *not-disclose* answer: We compute the distance between the WiFi scan of interest and that corresponding to each past *not-disclose* answer. The feature value corresponds to the shortest distance.

- Difference between distances from closest *disclose* and *not-disclose* answers: The feature reveals whether the scanned WiFi signal is similar to that corresponding to a past *disclose* answer or *not-disclose* answer.

5) TIME

Each ESM answer includes a timestamp that shows when it was collected. The time strongly relates to people's privacy preferences. We compute the following features from the timestamp of the ESM answer.

- Hour of the day and days of the week: We employ the hour of the day and the days of the week of the timestamp as features.

- Past privacy preferences of the same time period: This feature is the ratio of *disclose* answers from previous days obtained at the same time period (within k minutes) as the timestamp of interest ($k = 30, 60$). Note that people's privacy preferences may differ on weekdays and weekends. If the timestamp of interest was obtained on weekends (weekdays), we compute the feature value using only answers from past weekends (weekdays).

6) LOCATION (OBTAINED FROM GPS AND WiFi)

To compare the performance of the above sensors with the performance of a GPS sensor in the evaluation section, we also extract features from GPS (with WiFi) coordinates and use them to predict a user's privacy preference. For comparison, we extract features from GPS data that are almost identical to those extracted from a WiFi signal (e.g., distances from home and workplace, frequency of visiting, and distance from past answers).

B. LEARNING PRIVACY PREFERENCE MODEL

With the above procedure, we obtain pairs consisting of a feature vector and an answer (*disclose* or *not-disclose*). From the pairs, we learn a user's personalized privacy preference model that classifies an extracted feature vector in an appropriate class (*disclose* or *not-disclose* class). We construct a model for each category of person. Note that, as mentioned in the introduction section, we design a model that can detect outlying feature vectors in addition to classifying those vectors into an appropriate class. To accomplish both the classification and outlier detection tasks, we learn the distributions of the *disclose* and *not-disclose* classes by using a statistical model. When we classify a feature vector, we compute the likelihood of each model (class) for the vector and can determine that the class with the highest likelihood is the classified class. On the other hand, when we determine whether or not a feature vector is an outlier, we can employ the estimated distances (likelihood values) between the vector and the distributions.

1) REQUIREMENTS

When we design a probabilistic density based model, we should take the following into consideration.

- When we model the distributions of the two classes, we should define the distances between feature vectors appropriately. The distances between vectors belonging to the same class should be small. In other words, the distances between vectors belonging to different classes should be large. By doing so, we can easily distinguish the *disclose* class from the *not-disclose* class. However, because we employ features extracted from many sensors, the scales of the values of the features are different. Assume that the scale of a feature closely related to a user's privacy preference is small, e.g., $[-1, 1]$, and the scale of another feature having little relation to the user's privacy preference is large, e.g., $[-100, 100]$. In this case, the value of the first feature is not reflected in the (Euclidean) distances between vectors. When we employ discriminative classifiers such as decision tree and SVM, we can disregard it because the classifier can find a class boundary or threshold independently of the scale. However, with the distribution based approach, we should define the distances between vectors carefully.
- As above, there are many vector dimensions. So, we suffer from the curse of dimensionality when computing the distances between the vectors. While it is necessary to reduce the dimensionality, we should take the above-mentioned distances between the vectors into consideration.
- In our preliminary investigation, we found that a user's privacy preference varies depending on various kinds of outlying situations. For example, the preference was affected by apparent outlying situations such as business trips. It may not be difficult to detect such situations because the corresponding vectors are very different from those of ordinary situations. On the other hand, we found many outlying situations that

were different from such apparent situations, e.g., a user is at home (or at her workplace) at a different time from usual and she holds a party at home (only the microphone and acceleration data are different from usual). In such cases, only the values of few particular features are different from those of 'not-outlying' feature vectors. That is, the distances between such vectors and 'not-outlying' vectors, which are computed from all the feature values, become small. We should cope with this problem.

2) CONSTRUCTING MODEL

We describe the procedure of our approach, which takes the above requirements into account.

[1. Computing importance of features]: First, we compute how much each feature contributes to the classification task by employing the concept of information gain. With this information, we define the distances between vectors below. The information gain is usually used to find distinguishable features of instances [28]. The better a feature classifies the instances, the greater the information gain of the feature becomes.

[2. Scaling]: By using computed information gain for each feature, we define the distances between the vectors. We first standardize the feature values of the vectors for each feature (dimension). Then, we multiply the feature values of each dimension by its corresponding information gain. By doing so, the distance between two vectors that have different values related to a distinguishable feature becomes large. That is, the distances between vectors belonging to different classes become large.

[3. Dimensionality reduction]: Although there are many dimensionality reduction techniques such as principal component analysis (PCA) and random projection, in these techniques a new dimension is represented as a combination of the original dimensions (e.g., by a weighted sum). However, because many outlying vectors have different values from not-outlying vectors in a few particular features (dimensions) as mentioned above, these techniques that employ combined dimensions make it impossible to distinguish such outlying vectors from the other vectors. In this work, we re-use the information gain computed above to reduce dimensionality. That is, we discard features (dimensions) whose information gain is small. In other words, we retain the top-k features. This simple approach can also approximately preserve the pairwise distances between feature vectors.

[4. Learning probabilistic density]: We employ a Gaussian mixture model (GMM) to model the feature vectors for each class. A GMM is a weighted sum of M Gaussian densities as given by $p(f|\lambda) = \sum_{i=0}^{M} \pi_i N(f|\mu_i, \Sigma_i)$, where f is a feature vector. π_i is the mixture weight of the ith Gaussian, and μ_i and Σ_i show the mean vector and covariance matrix of the Gaussian, respectively. These GMM parameters are collectively represented by λ. With the mixture of distributions, we can model user privacy policies that vary depending on the user's situations. We employ expectation maximization (EM) [29] to estimate the GMM parameters.

3) IDENTIFYING PRIVACY PREFERENCE

With the GMM prepared for each class, we classify a test feature vector into an appropriate class, and determine whether or not the vector is an outlier. How we deal with an outlying feature vector depends on the policy of the location sharing system. In this work, for safety we classify a vector into the *not-disclose* class when the vector is determined as an outlier.

[1. Classification]: We first compute the GMM likelihood for each class (i.e., $p(f|\lambda_{disclose})$ and $p(f|\lambda_{not\text{-}disclose})$) and then determine that the class with the highest likelihood is the classified class.

[2. Outlier detection]: As mentioned above, many outlying vectors have different values from not-outlying vectors in a few particular features (dimensions). To cope with this problem, we compute a vector's likelihood value for *each* dimension as follows and judge whether or not the vector is an outlier by using the minimum likelihood, i.e., $\min_n(p(f_n|\lambda))$.

$$p(f_n|\lambda) = \sum_{i=0}^{M} \pi_i \frac{1}{\sqrt{2\pi\sigma_{i,n}^2}} \exp\left(-\frac{(f_n - \mu_i^n)^2}{2\sigma_{i,n}^2}\right),$$

where f_n is the nth feature value of the vector, and μ_i^n and $\sigma_{i,n}^2$ are the mean and variance of the nth dimension related to the ith Gaussian, respectively. When the likelihood values of both the $\lambda_{disclose}$ and $\lambda_{not\text{-}disclose}$ models are smaller than a threshold, we determine the feature vector to be outlying. This approach can capture outlying vectors whose values are different from not-outlying vectors in a few particular features.

V. EVALUATION

A. EVALUATION METHODOLOGY

To investigate the performance of our approach, we compared the C4.5 decision tree (DT) with our method. We also investigated the performance of a method that employs simple handcrafted rules rather than extracted features. This method decides whether or not a participant's location is disclosed by using her handcrafted rules based on time and her location, e.g., ''my location is not disclosed to my boss when I am at places within 300 meters of my home'' and ''my location is not disclosed to my spouse between 7 p.m. and 9 p.m. on weekdays.'' This approach is similar to those used by existing location sharing services and studies. The rules used in our evaluation were made by the participants after the experiment. We evaluated the classification performance of the methods by using precision, recall, and F-measure ($\frac{2 \cdot \text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$). The ground truth for the evaluation was collected using the ESM survey mentioned in section III. Note that, if a participant responded with *disclose* (or *not-disclose*) for all the ESM questions related to a certain category of persons, i.e., she decided to disclose (or not to disclose) her location to the category every time, we did not evaluate her model for the category because it is very easy to predict her preferences.

We employed 'leave-one-day-out' cross validation. That is, we regarded one day's sensor data and ESM answers

as test data and the remaining days' sensor data and ESM answers as training data, and we computed the classification performance of the test data. We iterated the procedure so that each day's data were used as test data once. For example, assume that user A collects ten-day sensor data. We train a classifier by using nine days' data from user A and test the performance of the classifier by using the remaining one day's data from user A. That is, we use the nine days' data to compute features (e.g., constructing a WiFi graph and computing sound pressure levels) for training the classifier. We iterate the procedure ten times so that each day's data are used as test data once.

B. RESULTS

1) CONTRIBUTIONS OF EACH SENSOR

Fig. 11 shows the classification accuracy (overall F-measure) of DT when we computed the accuracy by employing only data from each sensor. The accuracies of Acc (accelerometer data) were much poorer than those of Loc (GPS and WiFi location). It was difficult to predict a user's privacy preferences using only an accelerometer. The accuracies of MIC (microphone data) were about 12% poorer than those of Loc. We could achieve relatively high accuracies with only microphone data. We investigated the prediction errors of Loc and MIC. The participants' privacy preference changed depending on whether they were at their homes or workplaces. We could not distinguish these places using MIC alone because their homes and workplaces were usually quiet. On the other hand, when we used GPS coordinates (and coordinate data for their homes and workplaces), we could distinguish these places. We also found that the participants' privacy preferences sometimes changed even when they were at the same places. For example, a participant basically disclosed her location when she was at her workplace. However, when she played sports after work at a sports facility in her workplace, she did not disclose her location to her boss.

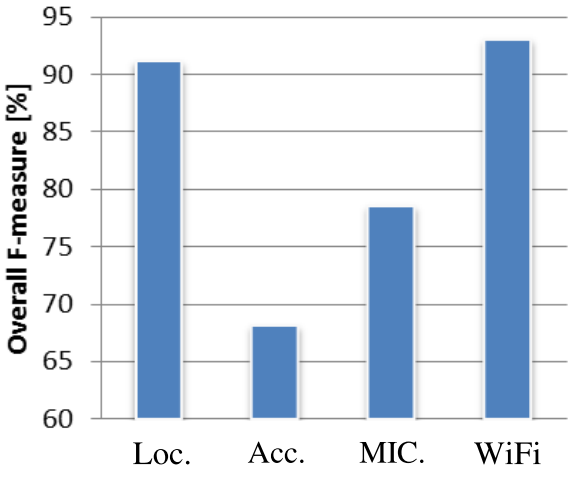

FIGURE 11. Classification accuracy when using only data from each sensor (DT: decision tree).

We could not model such preferences using only GPS coordinates. However, with MIC, we could successfully predict such preferences. This may be because we could capture her activities (working or playing sports) from sound data. Also, as mentioned in the section III, when a participant walked to a station to save on bus fare, he did not disclose his location. With MIC, we could successfully capture the participant's means of transportation (bus or on foot).

Surprisingly, the accuracies of WiFi (WiFi scan) were almost the same as those of Loc. This is because graphs constructed from WiFi scans enable us to compute features similar to those of Loc. We investigated the prediction errors of Loc and WiFi when we used DT. The WiFi method was usually unable to capture the participants' preferences when they were moving rapidly (esp. on trains). With GPS coordinates, we can easily compute the distance between a user's current location and her home/workplace even when she is moving rapidly. However, with the WiFi method, the geographical density of the WiFi scans becomes very low because the user moves rapidly, and we cannot find any path between the scans and an AP at a home/workplace. On the other hand, the WiFi method was able to capture the participants' room-level privacy preferences. As mentioned above, a participant did not disclose her location when she was playing sports after work at a sports facility in her workplace. With WiFi, we could capture different APs when she was working or playing sports. However, neither of these approaches could cope with situations where the participants' preferences changed even if the participants were in the same rooms. We show examples. One participant was unwilling to disclose her location to her friends *only* when she was at home during the daytime on weekdays. Also, one participant disclosed her location to her boss when she was engaged at her workplace (a shop). However, she did not disclose her location when she visited or passed by the shop outside of her working hours. The two methods could not deal well with such situations.

2) PERFORMANCE

We can achieve greater accuracies by combining sensor data obtained from several sensors and timestamps. Loc+T (DT) in Fig. 12 shows the accuracy (overall F-measure) when we used GPS (with WiFi) coordinates and timestamp data. MIC+WiFi+T (DT) shows the accuracy when we used microphone, WiFi scan, and timestamp data. Surprisingly, by combining MIC and WiFi data (MIC+WiFi+T DT), our method could outperform the GPS-based approach (1.0% higher). We could confirm a significant difference between their accuracy rates ($p < .01$) (In addition, as shown in Fig. 1, the MIC+WiFi approach has an advantage in terms of battery consumption over the GPS-based approach.) We found that the MIC data compensated for the weakness of the WiFi data. For example, we could successfully predict the participants' preferences when they were moving rapidly because a microphone could capture her means of transportation, e.g., train and car. On the other hand, the accuracy of Acc+MIC+WiFi+T (DT) was slightly poorer than that of

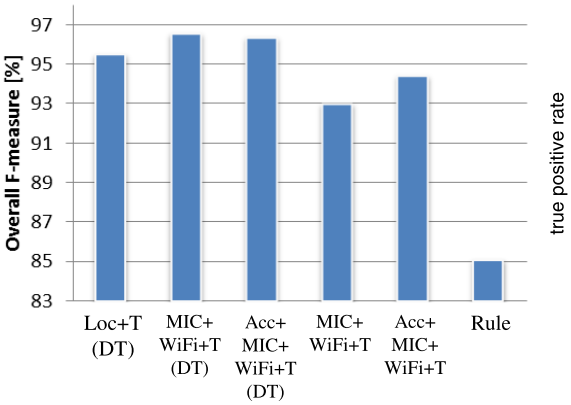

FIGURE 12. **Classification accuracy when using combination of sensors.**

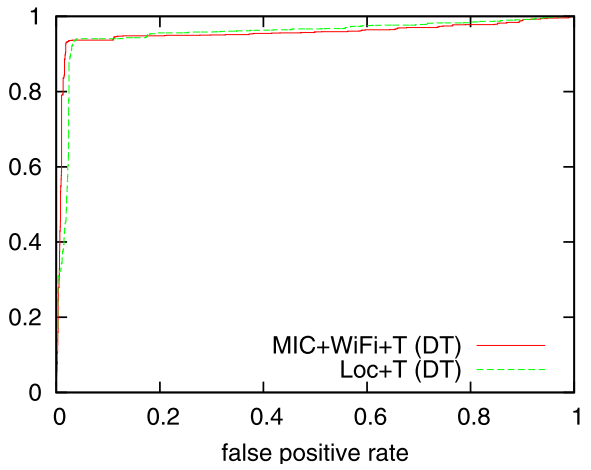

FIGURE 13. **ROC curves for MIC+WiFi+T (DT) and Loc+T (DT).**

MIC+WiFi+T (DT). This may be because the accuracy of MIC+WiFi+T (DT) was sufficiently high and acceleration data did not contributed.

By employing sound and WiFi data, we could capture the following privacy preferences that the GPS approach could not capture because a user's privacy preferences differ even if she is at the same coordinates. (1) A user's preferences sometimes change depending on the room. With WiFi scan data, we could capture her room-level preferences. (2) A user's preferences change according to her context, e.g., mode of travel and activities (working or playing sports). With sound, we were able to capture the user's context in detail.

Fig. 12 also includes the results of the rule-based method (Rule). The classification accuracy of Rule was only about 85% because it was difficult to deal well with unusual activity patterns, e.g., going on a business trip, participating in a party, a side trip, and dining out. It is difficult to make rules that specify such special events. About half of the errors were caused by this problem. In addition, we found that several results were misclassified due to GPS sensor data errors. When a user's GPS sensor acquires GPS signals from only small numbers of satellites, her position determined from the signals may not be accurate. Participants who were working by a window suffered from this problem. Rule is subject to errors because, for example, it determines whether or not a participant is at her workplace by using a static threshold.

3) IDENTIFYING *NOT-DISCLOSE*

We expect location sharing applications to exhibit high accuracy as regards *not-disclose* events in order to hide these events from other people. On the other hand, a binary classifier computes the score (probability) of an instance that exhibits the degree to which that instance is a member of a class. We can threshold this score to make a binary decision. By varying the threshold, we obtain an ROC curve for MIC+WiFi+T (DT) (and Loc+T) as shown in Fig. 13. Because we are now focusing on the *not-disclose* class, the true positive rate corresponds to the ratio of the number of *not-disclose* instances correctly classified in the total number of *not-disclose* instances. Also, the false positive rate is the ratio of the number of misclassified *disclose* instances in the total number of *disclose* instances. By changing the threshold, we can control the sensitivity of *not-disclose* event detection. As shown in the graph, when we want to find *not-disclose* events with a true positive rate of about 95%, the false positive rate becomes about 27%.

4) DETECTING OUTLYING SITUATIONS

Table 1 shows the results of our probabilistic density based classifiers (MIC+WiFi+T and Acc+MIC+WiFi+T). Our method (Acc+MIC+WiFi+T)) achieved almost the same performance as the Loc+T method. Our method also has performance comparable to the DT based approaches, even though for safety our method classifies all detected outlying ESM answers as *not-disclose*. The advantage of our method is that it prevents outlying *not-disclose* answers from being disclosed against a user's will. Table 1 shows the recalls of our method and the DT based approaches for outlying *not-disclose* answers in addition to the detailed classification results. A higher recall means that many outlying *not-disclose* answers are adequately protected. The recalls of our classifiers were much higher than those of the DT based approaches (about 10%). Note that when the log likelihood values of the models ($\lambda_{disclose}$ and $\lambda_{not\text{-}disclose}$) for a feature vector are smaller than -30, our method determines the vector to be outlying. By changing the threshold, we can control the sensitivity of the outlying event detection. Although we can detect many outlying *not-disclose* answers when we increase the threshold value, the overall classification performance decreases. Fig. 14 shows the relationship between the recall of outlying *not-disclose* and the average F-measure (of *disclose* and *not-disclose*) when we changed the threshold value. In this result, even when we want to find outlying *not-disclose* events with about 95% recall, the F-measure decreases only about 10%. Also, at that time, the recall related to *not-disclose* (true positive rate for both

TABLE 1. Classification accuracies of Loc+T, MIC+WiFi+T, and Acc+MIC+WiFi+T.

	Loc+T (DT)			MIC+WiFi+T (DT)			MIC+WiFi+T			Acc+MIC+WiFi+T		
	prec.	recall	F-measure	prec.	recall	F-measure	prec.	recall	F-measure	prec.	recall	F-measure
disclose	97.1	96.9	97.0	97.5	97.8	97.6	97.0	95.2	96.1	97.4	96.4	96.9
not-disclose	91.7	92.3	92.0	94.1	93.3	93.7	87.9	92.1	90.0	90.7	93.2	91.9
Average	94.4	94.6	94.5	95.8	95.5	95.7	92.5	93.7	93.0	94.1	94.8	94.4
not-disclose (outlier)	85.7	**80.5**	83.0	91.5	**76.4**	83.3	92.6	**88.8**	90.6	95.9	**87.3**	91.4

outlying and not-outlying) was 96.6%. Furthermore, the false positive rate (ratio of # of misclassified *disclose* instances in the total # of *disclose* instances) was only 15.0%. That is, our method could hide both outlying and not-outlying *not-disclose* events with the lower false positive rate compared with the DT based method (27%).

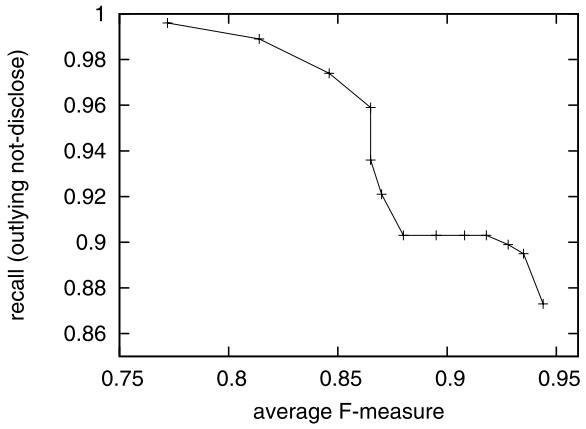

FIGURE 14. Relationship between recall of outlying *not-disclose* and average F-measure (Acc+MIC+WiFi+T).

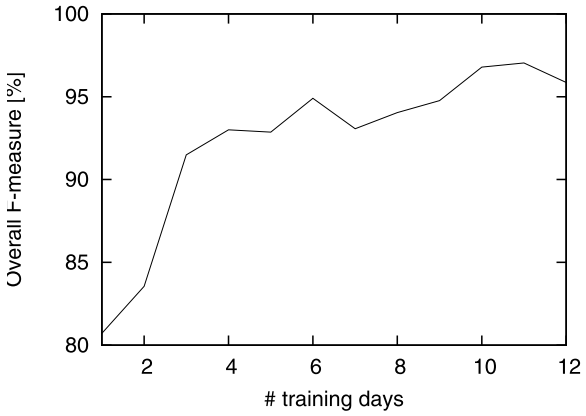

FIGURE 15. Transition of classification accuracy for Acc+MIC+WiFi+T when we increase # training days.

5) QUANTITY OF TRAINING DATA

Here we investigate the required quantity of training data. Fig. 15 shows the transition of the classification accuracy with Acc+MIC+WiFi+T when we change the number of training days. For example, when there are two training days, the method learns the preference model of a participant by using the first two of her sensor data and tests the model with the remaining sensor data. When there are seven training days, the accuracy reaches a high level (about 95%). This may be because the seven days data cover one week of data. Each participant responded to about 13 ESM questions a day. That is, seven days' data include about 90 ESM answers. Because we had to evaluate the classification performance, we collected ESM answers at a fixed interval (about 60 minutes) in this study. However, we should reduce the number of times ESM questions are presented when we implement an actual system. We can easily achieve this if we stop collecting ESM answers when the corresponding sensor data are similar to those of ESM answers collected previously. In our dataset, about half of the ESM answers were collected during "desk work" or "rest" activities. We may not require such a large quantity of ESM answers related to these activities. Therefore, we believe that we will be able to reduce the required ESM answers by about half if we do not collect so many ESM answers related to these activities.

6) EFFECTS OF WiFi SIGNAL STRENGTH

Here we investigate the effect of the features extracted by using the WiFi signal strength information. As shown in Fig. 16, by using WiFi signal strength information (SS), we could improve the classification performance. However, the improvement was small. This may be because the accuracies of Acc+MIC+WiFi+T and MIC+WiFi+T

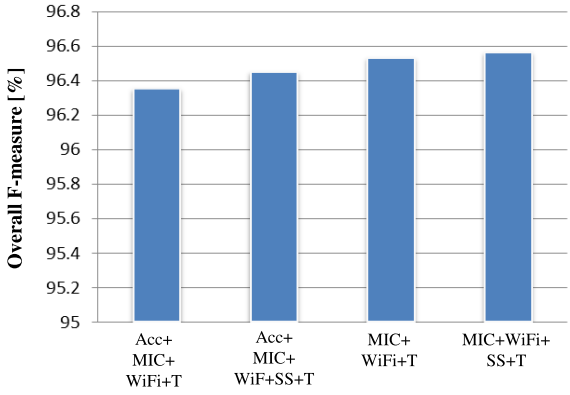

FIGURE 16. Effects of WiFi signal strength information (SS) when we used DT.

were sufficiently high. We confirmed that, by using the signal strength information, we could detect whether a participant was in a certain building or close to the building. It is difficult to detect such situation by using only a graph created from WiFi APs.

VI. CONCLUSION

In this work, we investigated how well can a user's privacy preferences be predicted by using sensors on a mobile phone such as a microphone and a WiFi module that are energy-saving and do not rely on external providers. By computing WiFi features based on a graph that represents the distance between WiFi APs, we could achieve a high level of performance comparable to that of the GPS-based approach. We also proposed a probabilistic density based classifier that identifies a user's privacy preferences in addition to detecting outlying situations that are not seen in training data.

REFERENCES

[1] V. Bellotti and A. Sellen, "Design for privacy in ubiquitous computing environments," in *Proc. 3rd Eur. Conf. Comput.-Supported Cooperat. Work (ECSCW)*, Sep. 1993, pp. 77–92.

[2] J. Krumm, "Inference attacks on location tracks," in *Proc. Pervas. Comput.*, May 2007, pp. 127–143.

[3] N. Sadeh *et al.*, "Understanding and capturing people's privacy policies in a mobile social networking application," *Pers. Ubiquitous Comput.*, vol. 13, no. 6, pp. 401–412, Aug. 2009.

[4] G. Bigwood, F. Abdesslem, and T. Henderson, "Predicting location-sharing privacy preferences in social network applications," in *Proc. 1st Workshop Recent Adv. Behavior Predict. Pro-Active Pervas. Comput.*, Jun. 2012.

[5] F. Durr, P. Skvortsov, and K. Rothermel, "Position sharing for location privacy in non-trusted systems," in *Proc. IEEE Int. Conf. Pervas. Comput. Commun. (PerCom)*, Mar. 2011, pp. 189–196.

[6] P. Skvortsov, F. Dürr, and K. Rothermel, "Map-aware position sharing for location privacy in non-trusted systems," *Pervas. Comput.*, Jun. 2012, pp. 388–405.

[7] S. Consolvo, I. E. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge, "Location disclosure to social relations: Why, when, & what people want to share," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2005, pp. 81–90.

[8] A. Khalil and K. Connelly, "Context-aware telephony: Privacy preferences and sharing patterns," in *Proc. 20th Anniversary Conf. Comput. Supported Cooperat. Work (CSCW)*, 2006, pp. 469–478.

[9] Y. Chon, E. Talipov, H. Shin, and H. Cha, "Mobility prediction-based smartphone energy optimization for everyday location monitoring," in *Proc. 9th ACM Conf. Embedded Netw. Sensor Syst.*, 2011, pp. 82–95.

[10] R. Hull *et al.*, "Enabling context-aware and privacy-conscious user data sharing," in *Proc. IEEE Int. Conf. Mobile Data Manag.*, Jan. 2004, pp. 187–198.

[11] L. Zavala, R. Dharurkar, P. Jagtap, T. Finin, and A. Joshi, "Mobile, collaborative, context-aware systems," in *Proc. AAAI Workshop Activity Context Represent., Techn. Lang.*, 2011.

[12] S. Lederer, J. I. Hong, A. K. Dey, and J. A. Landay, "Personal privacy through understanding and action: Five pitfalls for designers," *Pers. Ubiquitous Comput.*, vol. 8, no. 6, pp. 440–454, Nov. 2004.

[13] L. Palen and P. Dourish, "Unpacking privacy for a networked world," in *Proc. SIGCHI Conf. Human Factors Comput.*, 2003, pp. 129–136.

[14] R. Larson and M. Csikszentmihalyi, "The experience sampling method," *New Directions Methodol. Soc. & Behavioral Sci.*, vol. 15, pp. 41–56, 1983.

[15] D. Anthony, T. Henderson, and D. Kotz, "Privacy in location-aware computing environments," *IEEE Pervas. Comput.*, vol. 6, no. 4, pp. 64–72, 2007.

[16] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *Proc. IEEE 19th Annu. Joint Conf. IEEE Comput. Commun. Soc. (INFOCOM)*, vol. 2. 2000, pp. 775–784.

[17] P. Prasithsangaree, P. Krishnamurthy, and P. K. Chrysanthis, "On indoor position location with wireless LANs," in *Proc. 13th IEEE Int. Symp. Pers., Indoor Mobile Radio Commun.*, vol. 2. Sep. 2002, pp. 720–724.

[18] Y. Wang *et al.*, "A framework of energy efficient mobile sensing for automatic user state recognition," in *Proc. 7th Int. Conf. Mobile Syst., Appl., Services*, 2009, pp. 179–192.

[19] M. Berchtold, M. Budde, D. Gordon, H. Schmidtke, and M. Beigl, "ActiServ: Activity recognition service for mobile phones," in *Proc. Int. Symp. Wearable Comput.*, Oct. 2010, pp. 1–8.

[20] N. Lane *et al.*, "Enabling large-scale human activity inference on smartphones using community similarity networks (CSN)," in *Proc. 13th Int. Conf. Ubiquitous Comput.*, 2011, pp. 355–364.

[21] H. Lu, W. Pan, N. D. Lane, T. Choudhury, and A. T. Campbell, "SoundSense: Scalable sound sensing for people-centric applications on mobile phones," in *Proc. 7th Int. Conf. Mobile Syst., Appl., Services*, 2009, pp. 165–178.

[22] H. Lu, A. J. Brush, B. Priyantha, A. Karlson, and J. Liu, "SpeakerSense: Energy efficient unobtrusive speaker identification on mobile phones," in *Proc. 9th Int. Conf. Pervas. Comput.*, 2011, pp. 188–205.

[23] D. Gafurov, K. Helkala, and T. Sondrol, "Biometric gait authentication using accelerometer sensor," *J. Comput.*, vol. 1, no. 7, pp. 51–59, Nov. 2006.

[24] M. Cowling, "Non-speech environmental sound recognition system for autonomous surveillance," Ph.D. dissertation, Dept. Faculty Eng. & Inform. Technol., Griffith Univ., Nathan, Australia, 2004.

[25] J. Chen, A. Kam, J. Zhang, N. Liu, and L. Shue, "Bathroom activity monitoring based on sound," in *Proc. Pervas. Comput.*, 2005, pp. 47–61.

[26] S. Basu, "A linked-HMM model for robust voicing and speech detection," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Apr. 2003, pp. I-816–I-819.

[27] A. LaMarca *et al.*, "Place lab: Device positioning using radio beacons in the wild," in *Proc. Pervas. Comput.*, May 2005, pp. 116–133.

[28] I. Witten and E. Frank, *Data Mining: Practical Machine Learning Tools and Techniques*. San Mateo, CA, USA: Morgan Kaufmann, 2004.

[29] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc., Ser. B*, vol. 39, no. 1, pp. 1–38, 1977.

TAKUYA MAEKAWA is an Associate Professor with Osaka University, Suita, Japan. His research interests include ubiquitous and mobile sensing. He received the Ph.D. degree in information science and technology from Osaka University.

NAOMI YAMASHITA is a Senior Researcher with NTT Communication Science Laboratories, Kyoto, Japan. Her research interests include computer-supported cooperative work and human–computer interaction. She received the Ph.D. degree in informatics from Kyoto University, Kyoto.

YASUSHI SAKURAI is a Professor with Kumamoto University, Kumamoto, Japan. His research interests include indexing and search algorithms, data mining, and sensor data processing. He received the Ph.D. degree in engineering from the Nara Institute of Science and Technology, Ikoma, Japan.