Pers Ubiquit Comput (2014) 18:355–368
DOI 10.1007/s00779-013-0655-1

ORIGINAL ARTICLE

Community Similarity Networks

Nicholas D. Lane · Ye Xu · Hong Lu · Shaohan Hu · Tanzeem Choudhury · Andrew T. Campbell · Feng Zhao

Received: 3 October 2012 / Accepted: 30 January 2013 / Published online: 22 May 2013

Abstract Sensor-enabled smartphones are opening a new frontier in the development of mobile sensing applications. The recognition of human activities and context from sensor data using classification models underpins these emerging applications. However, conventional approaches to training classifiers struggle to cope with the diverse user populations routinely found in large-scale popular mobile applications. Differences between users (e.g., age, sex, behavioral patterns, lifestyle) confuse classifiers, which assume everyone is the same. To address this, we propose *Community Similarity Networks* (*CSN*), which incorporates inter-person similarity measurements into the classifier training process. Under CSN, every user has a unique classifier that is tuned to their own characteristics. CSN exploits crowd-sourced sensor data to personalize classifiers with data contributed from other similar users. This process is guided by similarity networks that measure different dimensions of inter-person similarity. Our experiments show CSN outperforms existing approaches to classifier training under the presence of population diversity.

N. D. Lane (✉) · F. Zhao
Microsoft Research Asia, Beijing, China
e-mail: niclane@microsoft.com

Y. Xu · A. T. Campbell
Dartmouth College, Hanover, NH, USA

H. Lu
Intel Labs, Santa Clara, CA, USA

S. Hu
University of Illinois at Urbana-Champaign, Urbana, IL, USA

T. Choudhury
Cornell University, Ithaca, NY, USA

Keywords Smartphone sensing · Activity recognition · Community-guided learning

1 Introduction

The popularity of smartphones with embedded sensors is growing at rapid pace. At the same time, research in the area of mobile phone sensing is expanding the boundaries of mobile applications [14, 21, 25, 26, 31]. Advances in human centric sensing are being fueled by the combination of: (1) raw sensor data, which is now practical to sample from large user populations, and (2) classification models that extract human activities, events, and context from low-level sensor data.

As user populations of mobile sensing applications increase in size, the differences between people cause the accuracy of classification to degrade quickly—we call this the *population diversity problem*. In this article, we demonstrate that the population diversity problem exists and classification accuracy varies widely even as the user population is scaled to up as little as 50 people. To address this problem, we propose *Community Similarity Networks* (*CSN*). CSN is a classification system that can be incorporated into mobile sensing applications to address the challenge to robust classification caused by the population diversity problem. The conventional approach to classification in mobile sensing is to use the same classification model for all users. Using CSN, we construct and continuously revise a personalized classification model for each user over time. Typically, personalized models require all users to perform manual sensor data collection where users provide hand-annotated examples of them performing certain activities while their devices gather sensor data (i.e., labeling data). This is both burdensome to the user and wasteful as multiple users often collect nearly identical data but the training of each model occurs in isolation of each other. The key contribution of CSN is that it makes the personalization of classification models practical by significantly lowering the burden to the user through a combination of crowd-sourced data and leveraging networks that measure the similarity between users.

CSN exploits crowd-sourcing to acquire a steady stream of sensor data and user input, either corrections to classification errors or the more common hand-annotated examples of sensor data when performing an activity (i.e., labeling). Under CSN, training classifiers becomes a networked process where the effort of individual users benefits everyone. However, the use of crowd-sourced data must be done carefully. Crowd-sourced data must only selectively be used during training so the resulting model is optimized for the person using the model. CSN solves this problem by maintaining similarity networks that measure the similarity between people within the broader user population. We do this by proposing three different similarity metrics (i.e., physical, lifestyle/behavior, and purely sensor data driven) that measure different aspects of inter-person diversity which influence classifier performance. The CSN model training phase then utilizes forms of boosting and co-training to allow these different types of similarity to each contribute to improving the accuracy of the personalized classifier.

The contributions of this article are as follows:

- CSN is the first system to propose embedding inter-person similarity within the process of training activity classifiers. To the best of our knowledge, CSN represents the only activity recognition system designed specifically to cope with the population diversity problem, which would otherwise jeopardize large-scale deployments of mobile sensing systems.
- We propose similarity metrics and a classification training process that support: (1) the extraction of similarity networks from raw sensor data and additional end-user input and (2) a learning process that adapts generic classification models through careful exploitation of crowd-sourced data guided by similarity networks.
- We have evaluated our system using three large-scale mobile sensing datasets that range in size between 50 and 120 people. We measure the robustness of our classifiers and the ability of CSN to cope with population diversity.

This article is an extended version of a paper [18] presented at Ubicomp 2011. We begin in Sect. 2 by examining the challenges presented by the population diversity problem. Sect. 3 details how activity recognition is performed under the CSN approach. We present the evaluation of our system in Sect. 4 Finally, we discuss related work in Sect. 5 before presenting our conclusions in Sect. 6.

2 Community-scale classification

In this section, we discuss a key difficulty in realizing large-scale mobile sensing applications. Specifically, we examine how population diversity can cause classification to become unreliable and inaccurate.

2.1 One size does not fit all

As mobile sensing prototype systems are deployed to an increasing number of users, their diversity increases as well. These users differ from one another in a variety of ways, a concrete example being physical dissimilarities as measured by sex, weight, height, or the level of physical fitness. Beyond these visually obvious differences, there are differences based on lifestyle and background. People come from different ethnic and social-economic origins, live and work in different locations, and while they may perform the same core collections of activities (e.g., socializing, exercising, working), they may do these activities in significantly different ways.

Inter-personal differences can manifest as differences in the discriminative patterns contained in sensor data that are used to classify activities, events, and contexts. For example, the features from accelerometer data that allow classifiers to distinguish between the basic activities of walking and running can be completely different between a group of older adults (older than 65 years) and a group of people who are in their 20s and 30s. Figure 1 visualizes this difference when these two groups are walking. We plot the first two PCA components on each axis of the figure based on a range of already validated activity recognition accelerometer features [22]. The very clear distinction between sensor data sourced from these two groups is surprising, particularly given the homogeneity you would expect in a simple activity like walking.

To further quantify this problem, we build a LogitBoost classification model [10] and reuse the same previously validated activity recognition features. This model is trained using labeled data from the group of people in their 20s and 30s. Using this model classification accuracy while they walk and climbed stairs ranges between 80 and 90 % for each person. However, when this same classifier is used by the group of aged people, the average accuracy dropped to nearly 60 %. Clearly, a one size fits all approach to classification models will not scale to large user populations which will contain many such groups.

This effect is not only limited to strictly physical behavior (e.g., walking, running or climbing stairs) but extends to a broader range of behavioral inferences. We investigate the breadth of this problem by performing an experiment on two distinct mobile sensing datasets. The first dataset (obtained from the authors of [35, 36]) contains

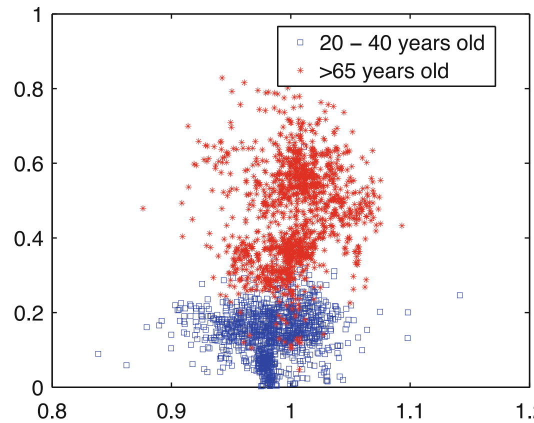

Fig. 1 We visualize the differences in features under an identical activity, walking, for two distinct community sub-groups. One group contains people over 65 years old, and the other group contains people in their 20s and 30s. Here we show just the first two components of the PCA of these features

GPS sensor data for 51 people performing 4 different transportation modes (e.g., driving a car or riding a bike). The second is comprised of multi-modal sensor data (e.g., microphone and accelerometer data) for 41 people performing a range of everyday activities (e.g., walking up stairs, exercising, brushing teeth). For both datasets, we train a single classifier and evaluate the accuracy for each user, by applying the classifier to test data sourced only from that user. Figure 2 is the Cumulative Distribution Function (CDF) of accuracy for these experiments and shows the spread of accuracy within the user population under both datasets. Accuracy levels for the transportation mode dataset are as low as approximately 40 % for the bottom performing 60 % of end-users and as high as 90 % for the top 7 %. Similarly, we find for the everyday activities dataset for 40 % of the users the accuracy is only 12 %, even for 80 % of users the accuracy raises only marginally to 55 %.

2.2 Limitations of current practice

The de-facto standard practice in incorporating classification into mobile sensing systems centers around a single unchanging classification model which is trained prior to deployment. Due to the reasons of population diversity, this model works for some people, but not others; the accuracy of the system remains difficult to predict and increasingly unreliable as the user population grows.

Ideally, the classifier would capture the distinctions between certain activities performed by different subgroups in the population as different activities entirely (e.g., walking when performed by two sub-groups could be two different classes), whenever these distinctions impact the classification process. However, this would significantly

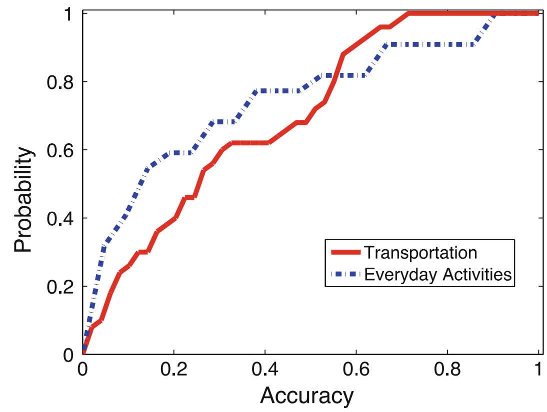

Fig. 2 Classification accuracy varies significantly within a large-scale user population for two datasets, one containing everyday activities and the other transportation modes

increase the amount of examples required for the same logical activity. Acquiring these examples is manually intensive (requiring careful labeling of data segments), making this approach impractical as it simply does not scale.

A promising direction being actively explored is the personalization of classification models to improve accuracy (e.g., [17, 23, 24, 32]). These models are tuned to sensor data generated or encountered by the individual. Typically, tuning occurs based on input from the user. For example, the user corrects classification errors or provides additional examples of activities by labeling sensor data with the ground-truth activity occurring during the sampling of the data. The classifier is then retrained using sensor data collected and labeled only by the user.

The limitation of such personalization of classification models is that accuracy *only* improves when and if people take the effort to manually provide additional sensor data examples. Independent of effort, it will also take time for people to encounter certain situations that are good discriminative examples to incorporate into the model. The key problem with this type of gradual improvement is that it leads to enormous amounts of redundant effort. Classification models are improved in *isolation* and each user potentially has to repeat steps that have already been done by other users to improve their own personal model.

3 Community Similarity Networks

In this section, we describe the system components and the core algorithms used at each stage of CSN. The CSN system is designed to construct and periodically update personalized classification models for each user. A key novelty of CSN is that it achieves personalization by using only a small amount of a specific user's training data and combining it with training data selectively recruited from others with whom the user shares similar traits.

3.1 Framework

Figure 3 illustrates stages of the CSN framework that produce personalized classification models for each user. Each of these stages occurs either in one of the two architectural components, the Mobile Phone Client software or the Mobile Cloud Infrastructure.

The Mobile Phone Client software samples sensor data to recognize human behavior and contexts by performing inference using classification models. While inference occurs locally, the models themselves are downloaded from the Mobile Cloud Infrastructure. The client software also collects training data comprised of both raw sensor data and data segments that have been labeled by CSN users with the ground-truth activity or context.

The Mobile Cloud Infrastructure is responsible for training classification models. Sensor data are used to construct similarity networks where network edges indicate the level of similarity between two users. CSN employs multiple dimensions of similarity (e.g., sensor data, physical, and lifestyle) to quantify the various ways users can differ. Several similarity networks are generated for each user, one for each of the similarity dimensions. Similarity-sensitive Boosting trains a classifier for each of the different similarity networks (that correspond to a different similarity dimension). Similarity Network Multi-training performs semi-supervised learning and improves every model by recruiting additional labels from the pool of unlabeled data. The final step of multi-training unifies each of the independent classifiers, trained by Similarity-sensitive Boosting, into a single ensemble classifier ready to be installed on the phone of the user.

3.2 Mobile Phone Client

In the following subsection, we describe the functions performed by the Mobile Phone Client, specifically we detail: (1) the classification pipeline, (2) implementation specifics, and (3) the collection of sensor data and ground-truth labels from users.

3.2.1 Classification pipeline

The classification pipeline includes sensor data sampling, feature extraction, and recognition of an activity, event or context.

We use the accelerometer, microphone, and GPS sensors to make a variety of proof-of-concept inferences. Our choice of features were based on observations made in prior work [24, 25, 35, 36]. For the accelerometer and microphone, we use the same feature set described in [25] which include a variety of time domain and frequency domain features effective for general activity recognition. For the GPS, we adopt features that were specifically designed in [35, 36] for transportation mode inference based on time-series GPS readings. Classification is done using a boosted ensemble [27] of naive Bayes classifiers [10]. Inference results are temporally smoothed using a simple Markov model. Although in our client the stages of the classification pipeline (i.e., features and the classification model) remain fixed, the parameters of the model are determined, and updated periodically, by the Mobile Cloud Infrastructure.

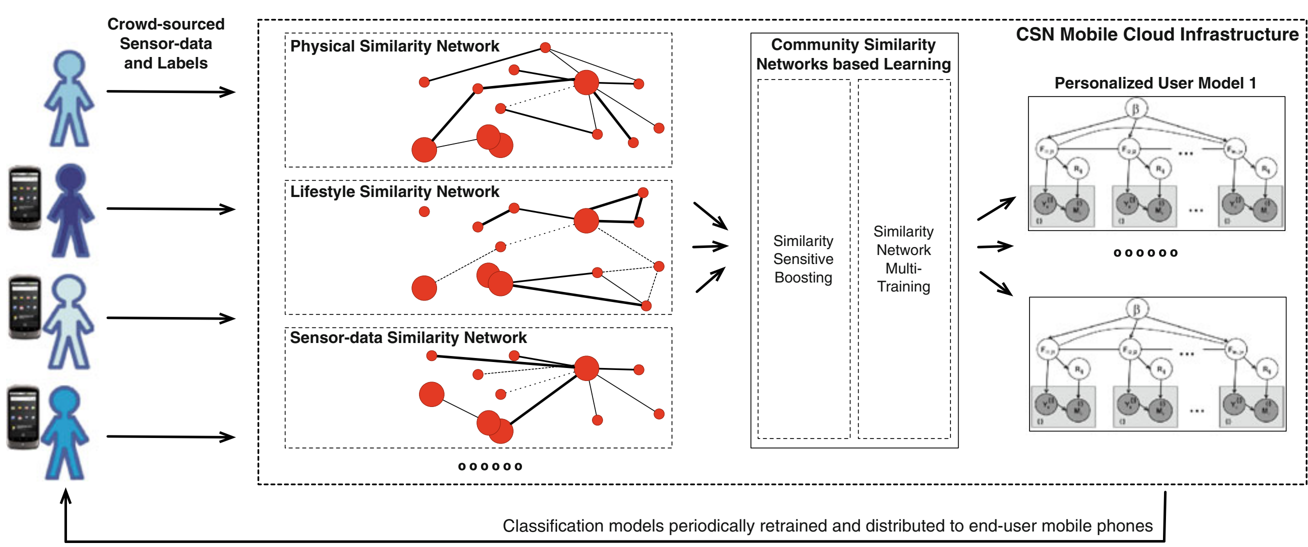

Fig. 3 The processing phases within Community Similarity Networks

3.2.2 Implementation

Our prototype client is implemented on the Google Android Nexus One [2], although other smartphone platforms also include the requisite capabilities needed by the Mobile Client [28]. The design of the phone client is split between a portable classification pipeline library written in C++ and set of device-specific supporting service (e.g., sensor sampling, end-user GUI). The library provides core classification pipeline components, including feature extraction and model inference. The device-specific components are written in Java and connected with the library via a JNI bridge.

3.2.3 CrowdSourcing

CSN exploits the crowd-sourcing of both sensor data and user input to improve classification models. User input provides the ground-truth activity to segments of sensor data. Two specific types of user input are supported. First, users can be asked to confirm or deny a class inference. For example, asking the user a question—*'Are you currently exercising?'*. Such responses are used later as positive or negative examples of certain activities. Second, users can explicitly label data as being an example of an activity or event. For example, users indicate the ground-truth activity when a segment of sensor data was sampled by selecting it from a list presented on the phone GUI. These types of interactions with users can be incorporated into applications in various ways. As an example, simple binary yes/no questions can be presented when the user unlocks their phone. Alternatively, more involved interaction, such as when users are selecting activities from a list, can be framed as software configuration or *calibration*. Similar forms of user interaction already occur in real products, for example, reading training sentences into speech recognition software or running for precisely one mile to calibrate a single activity recognition system like Nike+ [3].

3.3 Mobile Cloud Infrastructure

In this section, we describe how the Mobile Cloud Infrastructure: (1) computes similarity networks and (2) uses these networks to train personalized classification models that are distributed to all users.

3.3.1 Implementation

Our prototype implementation makes extensive use of Amazon Web Services [5] (AWS) which offer a number of generic components useful in building a distributed system. CSN utilizes the AWS message queues (SQS), binary storage (S3), and the simple queryable hash table service (SimpleDB). Each stage of the model training performed by the cloud is implemented either as python scripts or C++ modules depending on available library support given the required functionality. These stages run on a pool of linux machines as part of the Amazon Elastic Cloud product and interact with the individual AWS services as needed. Once a classification model is trained, it is serialized into a JSON-like format and written to the binary storage (S3), ready to be downloaded by the client.

3.4 Similarity networks

Each similarity network within CSN is constructed from the perspective of a single target CSN user. Nodes in the network represent other CSN users and edge-weights measure the degree to which the target user is similar to the other users. The CSN framework is designed to leverage multiple similarity measurements, which capture different dimensions of affinity between people. Depending on the activities or contexts that are to be recognized, different dimensions may be utilized. In this article, we propose the use of three dimensions of inter-person similarity: sensor data similarity, physical similarity, and lifestyle similarity. We demonstrate the effectiveness of each of these different dimensions in classifying different categories of activities and contexts in the Sect. 4. However, CSN is agnostic as to the exact similarity dimensions used.

We now describe the different similarity dimensions included within the CSN framework.

3.4.1 Physical similarity

Physical differences between people (e.g., weight, height, age, level of physical fitness, or well-being) will vary greatly from person to person within a large user population. Such differences can alter the way people move and perform certain physical activities. For example, as we detailed in the earlier section on community-scale classification, differences in age can effect the recognition of seemingly simple activities like walking upstairs or jogging.

To compute a single physical similarity value between a pair of users, CSN employs five types of physical information: age, height, weight, and the scores from two well-established physical well-being surveys (Yale Physical Activity Survey [15] and SF-36 physical activity score [4]). Each of these five values acts as an element in a vector that represents a single user. The physical similarity between users (i, j) is based on the mahalanobis distance between the two vectors for each person, as shown here,

$$\mathrm{sim}(i, j)^{\mathrm{phy}} = \exp(-\gamma(\boldsymbol{x}_i - \boldsymbol{x}_j)^{\top}\Sigma^{-1}(\boldsymbol{x}_i - \boldsymbol{x}_j)) \quad (1)$$

where $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ are the physical vectors for user i and user j, Σ is the covariance matrix and γ is an empirically determined scaling parameter.

3.4.2 *Lifestyle similarity*

The lifestyle similarity metric attempts to capture the diversity in how people lives their lives, examples of which include the following: occupation, diurnal patterns (e.g., are they an early morning person or active late at night), the distribution of activities performed, mobility patterns, and significant places [9] (e.g., where they work and live). Occupation and the location of work alter, for instance, the accelerometer and audio patterns occurring during social interactions (e.g., meetings and conversations). The time of day and significant places can effect the background context in which people perform activities. For example, late at night or early in the morning different locations will have different background activities that alter the sampled data (e.g., noise from people or cars). Collectively, these factors can change the distribution of features and shift discriminative boundaries for recognizing classes of activities.

We compute lifestyle similarity using three types of information: mobility and diurnal patterns in combination with the distribution of activities performed by users. Mobility patterns are based on GPS location estimates, which are tessellated into m distinct square tiles of equal size. Diurnal patterns are captured as a series of timestamps that are recorded whenever the user is inferred to be non-stationary by the classification pipeline. These timestamps are rounded and are represented as the particular hour in the week in which they occur (e.g., they range between hour 0 at the start of the week to hour 167 on the final hour of the final day). The distribution of activities is based on the duration users are inferred to be performing each activity classes (e.g, walking, socializing) detected by the classification pipeline. We construct three histograms for each of these types of lifestyle information for every user, normalizing the frequencies across all histograms. For each pair of users (i, j), we compute the lifestyle-based similarity by the following equation:

$$\mathrm{sim}(i,\ j)^{life} = \sum_{f \in \mathcal{F}} \mathbf{T}_f(i)^\top \mathbf{T}_f(j) \qquad (2)$$

where $\mathbf{T}_f(i)$ is a histogram vector for user i of type f and $\mathcal{F}$ contains each type of lifestyle histogram. Lifestyle similarity between two users is the sum of the inner product of the histograms for each type of lifestyle information used by CSN.

3.4.2.1 *Sensor data similarity*

Differences between users lifestyle, behavioral patterns, or location will likely manifest as differences in their sensor data. Sensor data similarity, unlike lifestyle similarity, does not require feature engineering to extract particular dimensions from the data. Similarly, unlike physical similarity, it does not require additional information to be provided by the user (e.g., age). Instead, measuring inter-person similarity based on sensor data is inherently purely a data-driven approach. Later, we report its effectiveness across a wide range of classification tasks. However, it requires much larger amounts of computation to determine similarity between users than computing lifestyle or physical similarity.

Computing similarity based on the raw sensor data will be effected by noise and capture too many insignificant variations in the data. Instead, we compute sensor data similarity between the features extracted from the raw data. For this purpose, CSN employs the same features used by the classification pipeline, described earlier in this section. Individual users will accumulate varying amounts of sensor data based on how frequently they use their device. Consequently, we compute “set” similarity whereby any duplicate feature vector for a user is ignored and only the unique vectors generated from the data of a person are used. For our similarity measurement, we adopt a commonly used formulation [34] where the similarity between two users is,

$$\mathrm{sim}(i,j)^{\mathrm{data}} = \frac{1}{N_i N_j} \sum_{l=1}^{N_i} \sum_{m=1}^{N_j} \mathrm{sim}(\boldsymbol{x}_{il}, \boldsymbol{x}_{jm}) \qquad (3)$$

where $\{\boldsymbol{x}_{il},\ l = 1 : N_i\}$ is the data of user i, and $\{\boldsymbol{x}_{jm}, m = 1 : N_j\}$ is the data of user j.

However, this pairwise computation quickly becomes impractical as the number of unique features per user increases. To cope with this problem, we adopt Locality Sensitive Hashing (LSH) [8] to construct a histogram to characterize the “set” of data from each user and then compute the similarity between a pair of users by employing this histogram representation. Our method obviates the need to compute the pairwise relations of data from two users as required by the traditional “set” similarity, which has a linear time complexity with the average data size of each user. The basic idea of the LSH method is that a hashing function family can capture the similarity between data. In other words, similar data have a high probability to share the same value after hash mapping.

$$\begin{aligned} Pr_{h \in \mathcal{H}}[h(\boldsymbol{x_1}) = h(\boldsymbol{x_2})] &= s_{\mathcal{H}}(\boldsymbol{x_1}, \boldsymbol{x_2}) \\ &= E_{h \in \mathcal{H}}[s_h(\boldsymbol{x_1}, \boldsymbol{x_2})] \end{aligned} \qquad (4)$$

Therein, $\boldsymbol{x_1}, \boldsymbol{x_2} \in \mathcal{X}$ are two data, $\mathcal{H}$ is a LSH family, h is the hash function sampled from $\mathcal{H}$, and $s_{\mathcal{H}}$ is a similarity measure of $\mathcal{X}$, which is induced by the LSH family $\mathcal{H}$ [8].

In CSN, we randomly choose B independent 0/1 valued hashing functions $\{\boldsymbol{h_i}\}$ from the *random projection for* $\mathcal{L}_2$

distance LSH family [8] and form a B-bit hash function $f = (\boldsymbol{h_1}, \boldsymbol{h_2}, \ldots, \boldsymbol{h_B})$. The number of functions B controls the trade-off between efficiency and accuracy [8].

We apply the B-bit hash function to build histograms for each user, whose size is 2^B. Now, we formalize how to construct a histogram from the features of the user. According to the description, let $\mathcal{X}$ be the data space, $\mathcal{F}$ be the B-bit hash functions family mapping from $\mathcal{X}$ to $\mathcal{D} = \{0, 1, \ldots, 2^B - 1\}$, and $\{\mathbf{e}[i] | i \in \mathcal{D}\}$ be the standard basis of the $|2^B|$-dimensional vector space. Hence, given $h \in \mathcal{H}$, the histogram T_f for any user i is defined as follows,

$$\mathbf{T}_f(i) = \sum_{\boldsymbol{x}_{il} \in i} \mathbf{e}[f(\boldsymbol{x}_{il})] \tag{5}$$

here, $\{\boldsymbol{x}_{il}, l = 1 : N_i\}$ is data of user i, and $\mathbf{T}_f(i)$ is determined by the hash function f sampled from $\mathcal{F}$.

Thus, each element of the histogram vector $\mathbf{T}_f(i)$ can be regarded as a bin to record the frequency at which data from user i is mapped into it. As the value of the hash function indicates the probability that two data share the same value after mapping, two users that have many "matched" values in the corresponding bins of histograms imply a high similarity between them. The inner product of the two histogram vectors is next applied to compute the similarity metric for the two users:

$$sim(i, j)^{data} = \mathbf{T}_f(i)^{\top} \mathbf{T}_f(j) \tag{6}$$

To estimate the expectation shown in Eq. 5, we construct several histograms $f \in \mathcal{F}$ for each user and compute an average value using Eq. 6.

The time complexity of computing the LSH-based similarity metric is linear with the average quantity of data for each user. Compared with the pairwise computing method shown in Eq. 3 which is quadratic, the LSH-based similarity metric is very efficient.

3.5 Community Similarity Networks-based Learning

Learning personalized classification models for each user occurs in two stages under CSN. First, Similarity-sensitive Boosting trains three separate classifiers, one for each type of similarity network that is maintained for every user. Each of the classifiers is personalized to the characteristics of the specific individual who will use them. Further, each has different strengths when recognizing specific categories of activity depending on the similarity network used (e.g., physical similarity performs well with physical activities like climbing stairs). Second, Similarity Network Multi-training occurs which: (1) uses a semi-supervised approach to recruit additional labels from the unlabeled pool of sensor data leveraging the different strengths of each separate classifier and (2) unifies the three classification models trained by Similarity-sensitive Boosting into a single ensemble classifier, ready to be installed on the phone of the user.

3.5.1 *Similarity-sensitive Boosting*

A personalized classification model emphasizes the particular characteristics found in a target user to increase accuracy. CSN accomplishes personalization using a modified online boosting algorithm [27]. Boosting is a common learning technique that builds a model that is a composite of several weak classifiers trained over multiple iterations. At each iteration, certain data segments are weighted higher than others. Under conventional boosting, these weights are only altered based on the classification performance of the weak learner trained during the previous iteration. Those data segments that were incorrect are weighted higher than others so the weak classifier produced in the next iteration will be better able to classify these previously incorrect segments. CSN modifies this process by imposing an additional term to the weight at the initial iteration. This weight is based on the similarity between the user i, whose personalized model is being trained, and the user which provides the data,

$$\mathrm{weight}^{(0)}(\boldsymbol{x}_k) = \mathrm{sim}(i, k) \tag{7}$$

where k indicates the user who produces the data $\boldsymbol{x}_k$. We define $\mathrm{sim}(i, k)$ as the edge weight between these two individuals within the similarity network being used during the boosting process. As a consequence, only data segments from user i or any users who are highly similar to user i will be weighted highly and so able to have strong influence over the learned classification boundaries. In subsequent iterations, the weighting of data segments is left to fluctuate based solely on classification performance. As boosting is an ensemble technique, the CSN framework remains flexible, the weak learner can be replaced with any alternative supervised classifier based on the requirements of intended classification task.

3.5.2 *Similarity Network Multi-training*

The three varieties of similarity networks currently used in CSN capture different dimensions of similarity between users. For example, some users being highly similar in terms of physical characteristics but polar opposites when it comes to lifestyle. Using Similarity-sensitive Boosting in conjunction with any of these different networks will result in different classification models. Each network will emphasize different partitions of the training data. This diversity is valuable as the different similarity networks produce models that are highly effective for some classes of activity but not others (see Sect. 4). A simple example of this being those activities that are closely connected to the physical characteristics of the person, e.g., running and exercising, benefits from a classification model trained using a physical similarity network.

CSN exploits the strengths of each similarity network by adopting the technique of multi-training (a variation of co-training proposed in [37]). Multi-training is a semi-supervised training algorithm designed to utilize multiple complementary views of the same labeled training data to generate additional labels, which are assigned to data segments within the pool of unlabeled data. This approach is appropriate for CSN given that crowd-sourcing generates large amounts of unlabeled data. People will only infrequently take the time to provide any manual user input; but since simply collecting data is transparent to the user, then large pools of unlabeled data quickly can accumulate. Employing a multi-training approach allows CSN to use the diversity provided by the different similarity networks to make use of a plentiful and otherwise wasted resource, unlabeled data.

The multi-training process begins by initially using the three classifiers trained by Similarity-sensitive Boosting. Each of these classification models maintains an independent logical copy of the labeled and unlabeled data. An iterative process is applied whereby the classification models are used to in turn to "label" unlabeled portions in the logical datasets maintained by each of the other models. At the end of each iteration, the classifiers are then retrained (using Similarity-sensitive Boosting) based on the combination of the labeled data from the previous iteration along with any new additional labels. Acquiring labels in this way can be an error-prone process, as a result labels are only accepted when there is agreement with more than half of the classification models. Judging the quality of a proposed new label, based on a majority decision, is only one of many ways that quality can be assessed. Multi-training continues to iterate for several rounds until a stopping condition is met. CSN uses currently a stopping condition based on how many labels are accepted at each iteration. If the number of recruited labels is too low for too many iterations, then multi-training stops.

4 Evaluation

In this section, we evaluate the effectiveness and design choices of CSN. Our experiments show that by incorporating similarity networks among users into the classification process, CSN is better equipped to cope with the population diversity problem, compared to existing techniques.

4.1 Methodology

To evaluate CSN, we use two large real-world datasets and three representative baselines.

4.1.1 Datasets

Our three datasets require a variety of the activity inferences frequently used in mobile sensing applications. The first dataset, `Everyday Activities`, contains a broad range of routine human activities that have been used to support application domains such as mobile health [14]. The remaining two datasets, `Transportation` and `Physical Activities`, are much more focused on single activity domains—transportation modes and motion-based user actions, respectively. These two categories are building blocks of various mobile applications, for example, applications that promote green transportation [16]. We collect the data for `Everyday Activities` as part of a series of internal experiments. The data comprise both simple activities: {*walk*, *run*, *stationary*} and high-level behaviors: {*meeting*, *studying*, *exercising*, *socializing*}. A total of 41 people contribute to this dataset using a Nexus One smartphone sampling sensor data from the accelerometer, microphone and GPS. People carry the device for variable lengths of time that range between one and three weeks. For `Transportation` and `Physical Activities` we use external sources. `Transportation` is collected by the authors of [35, 36], with the dataset containing only different transportation modes, specifically: {*bike*, *us*, *car*, *alk*}. This dataset comprises 51 people who carry for three months one of a variety of devices that are equipped with a GPS, including phones, PDAs, and personal navigation devices. `Physical Activities` are collected by the authors of [12] and contain activities closely associated with the physical actions of users, specifically: {*run*, *skip*, *stairs up*, *stairs down*, *stationary*, *walk*}. This dataset includes 120 people who contribute accelerometer data using a variety of iOS devices (i.e., different models of iPhones and iPods). People participate in this experiment for around 13 months.

4.1.2 Benchmarks

We compare the performance of CSN against three benchmarks, *single*, *isolated* and *naive-multi*. Our benchmarks use the same features and apply the same classification model as CSN but differ significantly in how they approach classifier training. The benchmarks of *single* and *isolated* correspond with the two types of common practice we detailed in the earlier section on community-scale classification. In *single*, the same generic model is provided to all users. Unlike CSN, after the release of the system, the model does not change and new training data is not collected. Under *isolated*, every user has their own model. Each user model is personalized by using training data sourced directly from the user. The weakness is that each classification model is considered in isolation of one another. No co-operation or sharing of training data occurs among users. Finally, *naive-multi* allows us to demonstrate the benefit of CSN solely attributable to the use of similarity networks. During training, *naive-multi* performs boosting and multi-training, the same techniques used in CSN. However, *naive-multi* uses the conventional versions of these learning techniques without the use of similarity networks. Specifically, the differences are as follows: (1) during boosting the weighting of training data at each iteration only changes based on classification performance instead of inter-personal similarity and (2) during multi-training the classifiers used are not based on dimensions of similarity but based on classifiers trained with equally sized random subgroups of the training data.

4.2 Robust classification with low user burden

Our first set of experiments finds, under both datasets, CSN provides more robust classification than any of the benchmarks. Not only is CSN able to achieve higher classification accuracy but we observe classification accuracy is also more evenly distributed throughout the user population. Under CSN, the burden to provide training data is lowered, and thus, accuracy numbers comparable to the benchmarks can be achieved with smaller quantities of data. In what follows, we report accuracy values resulting from fivefold cross-validation.

Figure 4 shows the results of experiments where we assume users contribute different amounts of labeled data. For each quantity of labeled data, we measure the average per person accuracy of classification for models trained under CSN and the three other benchmarks. Figure 4a uses `Everyday Activities`, and Fig. 4b, c repeats the experiment using `Physical Activities` and `Transportation`. In each figure, the accuracy of CSN outperforms all baselines for each quantity of training data tested. For example, Fig. 4a shows if 500 labeled data segments are used (approx. 15 min of training data per user) then CSN outperforms *naive-multi* and *isolated* by 22 %. Similarly, from Fig. 4b we discover if 2,000 labeled data segments are used (approx. 55 min of training data from each user) CSN exceeds the accuracy of the next best performing baselines, *single* by 20 %. Finally, in Fig. 4c, we see if 1.6×10^4 labeled data segments are used (approx. 137 min of training data from each user) the accuracy of CSN exceeds *naive-multi* by 32 % and *single* by 47 %.

From Fig. 4 we also learn that CSN is able to lower the user burden of contributing training data. As an example, Fig. 4a shows *isolated* requiring 36 min of training data from a user to achieve 74 % accuracy. CSN can provide approximately this same accuracy for only 15 min of training data, a data reduction of 58 %. Alternatively, if we consider Fig. 4c, *isolated* is able to perform with 77 % accuracy but requires 270 min of training data. Again, CSN can provide approximately this level of accuracy but with 49 % less data, only needing 137 min of data per user. Under CSN, users are better rewarded for contributing data due to it having a higher ratio of classification accuracy to crowd-sourced training data, than the other benchmarks.

Figure 5 presents CDFs of per-user accuracy. We illustrate the fraction of the user population who experience different classification accuracy under CSN and all benchmarks. Figure 5a uses `Everyday Activities` and assumes users each provide 15 min of training data. Figure 5b is based on `Physical Activities` under the expectation of 110 min of per-user training data. Finally, Fig. 5c assumes users provide 137 min of labeled data from `Transportation`. Ideally, all users should receive the same level of accuracy, otherwise classification performance will be unpredictable when deployed. Better performance is indicated in these figures by curves that are furthest to the right. We observe from each figure CSN has the most even distribution of accuracy compared to all benchmarks. For example, Fig. 5(a) shows for 75 % of users that CSN provides 82 % accuracy compared to just 65 % for *isolated*, 48 % for *single* and 52 % for *naive-multi*. Figure 5b, c reinforces this finding. Figure 5b indicates for again 75 % of users CSN provides 83 % accuracy instead of the 79, 68, and 43 % accuracy offered by *isolated*, *single*, and *naive-multi*, respectively. For the same fraction of users (75 %), Fig. 5c echoes this result and shows 77 % accuracy is attained by 75 % of CSN users relative to 68, 53, and 66 % accuracy reached by the three baselines.

4.3 Benefits of leveraging Similarity Networks

With the following experiments, we investigate the effectiveness of the similarity networks used by CSN.

To test whether the similarity networks used by CSN are capturing meaningful differences between people, we collected additional demographic information from 22 individuals who contributed to `Everyday Activities`. Figure 6a, b plots the result of applying multidimensional scaling (MDS) to two similarity matrices for these people using physical and lifestyle similarity. Distances between points in these figures are proportional to differences in similarity. Figure 6a shows two clear groupings. We find these groups correspond to people with similar physical characteristics. The people in the tight cluster near the left of the figure are all over 30 years old, all male, and have similar physical fitness levels. In contrast, the looser clump of people near the middle are in the same age range (22–26) but are all more diverse in

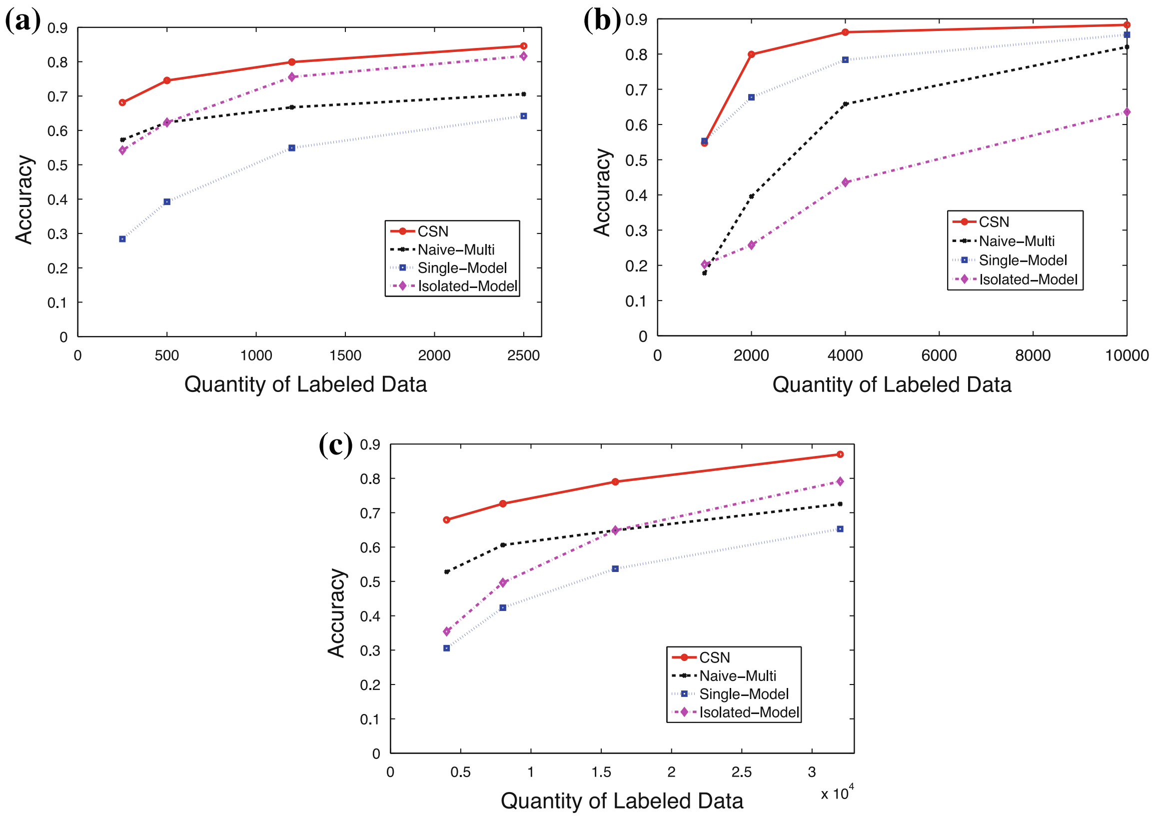

Fig. 4 Classification accuracy for CSN under different quantities of training data. **a** `Everyday Activities`, **b** `Physical Activities`, **c** `Transportation`

terms of sex and fitness. The outlier in Fig. 6a is a 50-year-old woman and is distinct due to her sex and exceptional fitness. The clusters in Fig. 6b also correspond to our interview ground truth. The tight cluster to the left is a small group of people who live off campus and maintain regular 9 a.m. to 5 p.m. working hours. They are in sharp contrast to the very loose cluster on the right of the figure. This cluster contains students who, although they live very close to each other, also have erratic sleeping and activity patterns which results in them being grouped but not as tightly as the nearby cluster.

In Fig. 7, we can see the value of using multiple similarity dimensions. The figure illustrates the different levels of classification accuracy achieved when using each of our three similarity dimensions to classify classes found in `Everyday Activities`. None of the three similarity metrics has the highest accuracy across all the activities. We find a similar pattern exists within `Transportation` and `Physical Activities`. By exploiting all of these forms of similarity, CSN is able to better handle a wide range of classification tasks. This result supports the design choice to use multiple dimensions of similarity and leverage them all when training classification models.

4.4 Cloud scalability with low phone overhead

Our remaining results report on the overhead to smartphones in adopting CSN, along with the ability for CSN to scale to large user populations.

We profile the computation and energy consumption of our CSN client on the Android Nexus One. We find resource consumption comparable to prior implementations of classification pipelines on phones (e.g., [25, 26]). As this overhead is not specific to CSN but found in any mobile sensing application, we do not report further details. Overhead specific to CSN includes the transmission of sensor data and the downloading of classifiers trained in the cloud. We find typical file sizes for our classification models are on the order of 1–2 KBs, which means the cost of downloading classification models is minor. However, a significant cost to the phone can accrue when uploading sensor data. To lower this cost, the default cloud interaction strategy of our client is to wait until the phone is recharging before uploading data, under this policy the battery burden due to uploading is eliminated.

Figure 8 justifies our choice of a conservative data upload and classifier re-training policy by examining potential benefits of a more aggressive uploading strategy.

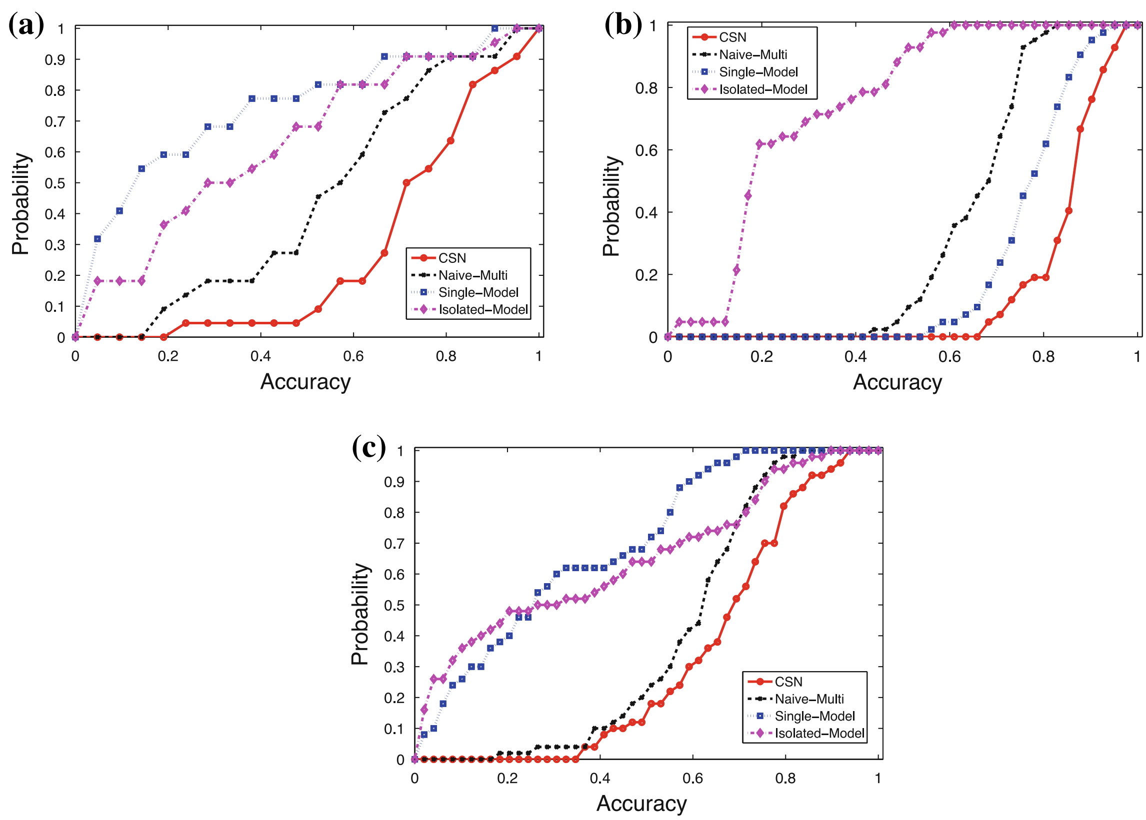

Fig. 5 DF of per-user classification accuracy for CSN. **a** `Everyday Activities`, **b** `Physical Activities`, **c** `Transportation`

In this experiment, we use only `Transportation` and assume a training set size of 1.6×10^4 labeled data segments. The dataset is replayed assuming that training test data enter the system based on the original experiment timestamps. Periodically, the classifiers for each user are re-trained based on the subset of the training data segments available. We repeat this experiment assuming different re-training frequencies, which in turn drive the rate at which data must be uploaded from clients. From this figure, we see that—for this particular scenario—classifier accuracy is not significantly impacted by the re-training frequency. Large amounts of client-side energy can be saved without lowering recognition accuracy. This figure also shows the energy consumed jumps significantly when the re-train frequency is equal or greater than 24 h. This is because at this level the client must upload collected data *intra-day* and so use the cellular radio; with less frequent re-training, the client is instead able to upload at night while recharging is occurring.

The computational demands of computing the three CSN similarity dimensions range from being light-weight to very demanding. We quantify this by profiling the computational overhead for computing similarity networks for all people within *Everyday Activities*. This raw dataset is more than 400 GB (mainly due to audio data). Using our CSN Mobile Cloud Infrastructure, configured with only one linux machine in the node pool, the computational time for each variety of similarity is, $\approx$200 min, $\approx$9 min, and $\approx$3 min, respectively, for sensor data, lifestyle, and physical similarity. The sensor data similarity is the most costly of these three as it requires pairwise calculations between users.

Personalized models are trained by CSN for each user; however, this can become a bottle-neck. The workload of the Mobile Cloud Infrastructure increases with population size due to: (1) the pairwise calculation of similarity between users and (2) each new user requires a new model to be trained. For this reason, we designed our Mobile Cloud Infrastructure to effectively leverage a variable sized pool of cloud nodes, so additional nodes could be added when required. Figure 9 illustrates the benefit of increasing the cloud node pool while computing the complete similarity network for `Everyday Activities` and `Transportation`. From this figure, we see the ratio between lower computation time and additional cloud nodes is similar for both datasets. Further, the largest gains occur when shifting from one machine to three.

We experiment with an alternative approach to addressing this same problem that requires a simple extension to CSN. Instead of training a model for each

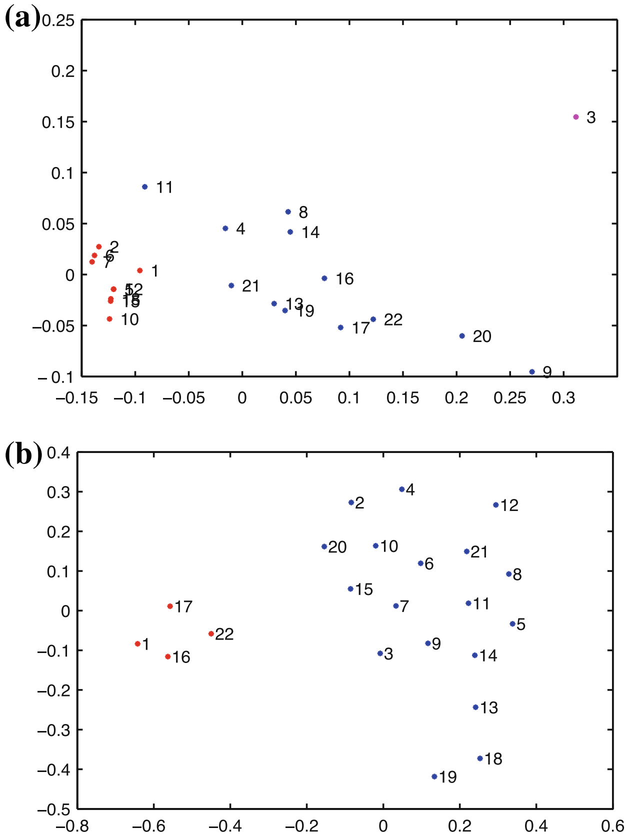

Fig. 6 MDS projection of physical and lifestyle similarity networks used by CSN. **a** Physical, **b** lifestyle

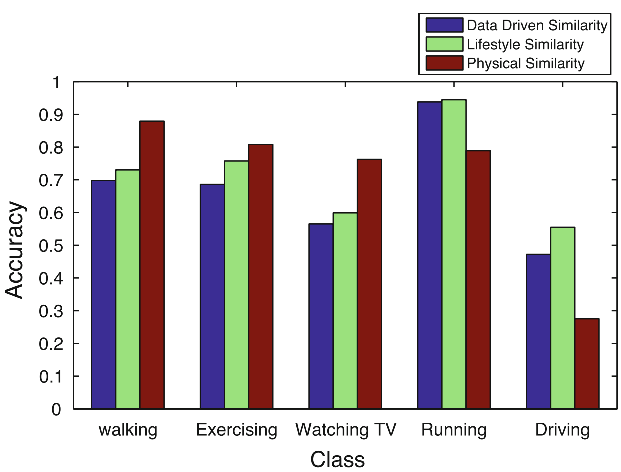

Fig. 7 The classification accuracy of each activity class under different dimensions of similarity using the Everyday Activities dataset. It shows different dimensions of similarity are effective for different activities

user, users are first grouped together by clustering. Similarity networks are then built not between people but between these groups, with a model trained for each group. We investigate this trade-off and cluster people with *k*-means using the least computationally costly similarity

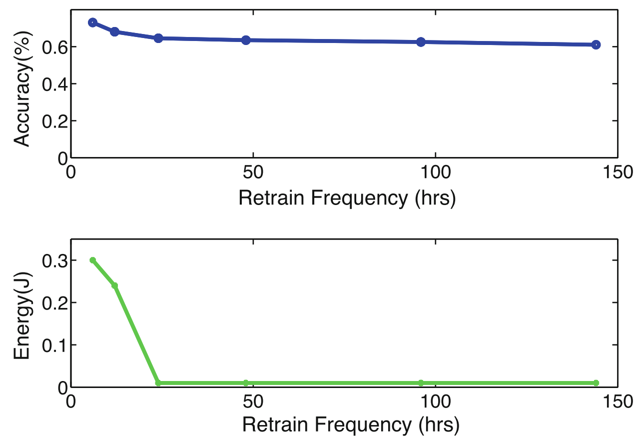

Fig. 8 Trade-off between client energy and classifier accuracy if classifiers are re-trained frequently

dimensions, lifestyle, and physical. By lowering the number of groups, we can reduce the Mobile Cloud Infrastructure workload. This trade-off is seen in Fig. 10. These figures illustrate how accuracy falls as the cluster size (the *k* in *k*-means clustering procedure) is reduced. Reducing the number of models dilutes the similarity between people in the cluster. Consequently, the model used by the entire group is less appropriate for everyone. Still, as the cluster number decreases the overhead to the mobile cloud is reduced, since fewer models need to be maintained. This approach allows us to regulate resource consumption by CSN irrespective of the size of the user population.

5 Related work

Applications that use mobile phone sensors have been steadily rising (e.g., [7, 11, 13, 14, 26]) and accurate classification of sensor data is becoming increasingly important.

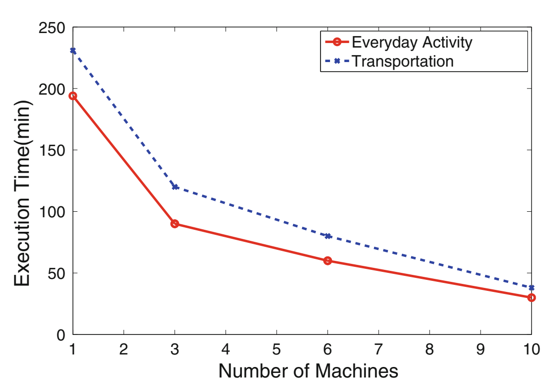

Fig. 9 Latency of processing training data decreases as the pool of cloud nodes is increased

Researchers investigating sensor-enabled mobile phone applications frequently encounter the limits of activity classification. It is becoming obvious that conventional approaches that rely on supervised learning and carefully controlled training experiments are not suitable. In recognition, researchers are considering alternatives. Current research directions point toward models that are adaptive and incorporate people in the process. Automatically broadening the classes recognized by a model is studied in [23] where active learning (where the learning algorithm selectively queries the user for labels) is investigated in the context of heath care. In SoundSense [24], a supervised classification process for a fixed category of sounds is augmented with a human-in-the-loop guided unsupervised process for learning novel sounds.

Research, such as [23, 24], focuses primarily on the individual to assist with classification. CSN leverages the user but also exploits communities of people (rather than just isolated individuals). How to precisely utilize communities is increasingly becoming an area of active research. Community-guided Learning (CGL) [29] models of human behavior are built with training data provided by non-expert mobile device users from the broader community. CGL overcomes the challenge presented by noisy labels being introduced to the training process by using data similarity in combination with the crowd-sourced labels. Lane et al. [19, 20] examines community-based techniques that leverage social ties that can coarsely proxy for the types of explicit similarity measurements used by CSN. Applications of community-awareness that can benefit more than just activity recognition are also beginning to emerge. For example, [6] investigates how mobile content dissemination can be improved by intelligently leveraging cross-community information present in multiple heterogeneous social communities.

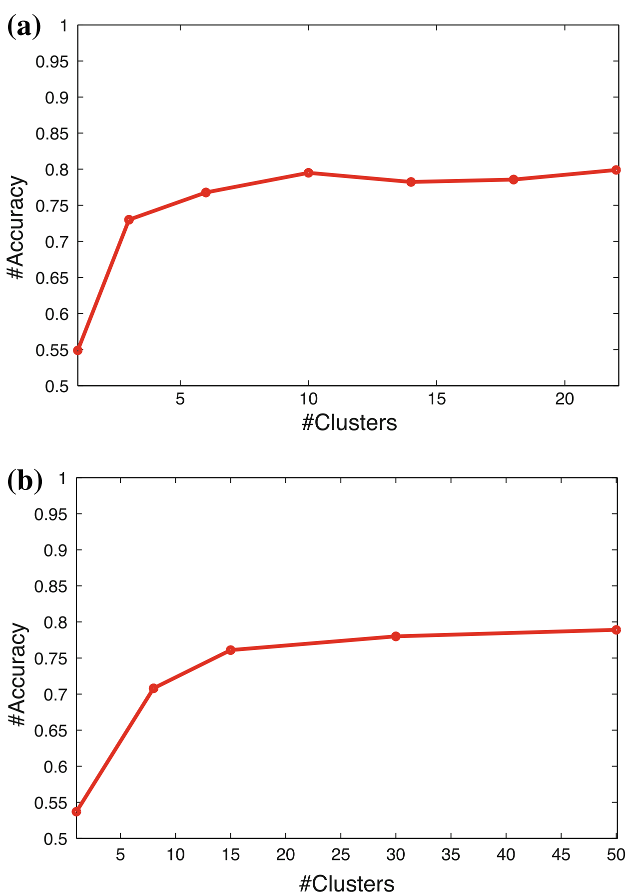

Fig. 10 The accuracy of CSN when we group the users into different number of clusters under both datasets. **a** `Everyday Activities`, **b** `Transportation`

The potential for crowd-sourcing has been long recognized with interest in the area being established by Luis Von Ahn [33]. Now, commercially available systems including Amazon's Mechanical Turk [1] have made it simple to exploit the power of using thousands of people. The use in CSN of crowd-sourcing builds directly on these existing directions. We see CSN as part of an exciting area of hybrid systems (e.g., CrowdSearch [30]) that intelligently combine the effort of the masses toward a task that neither computers nor humans can perform on their own.

6 Conclusion

In this article, we have proposed CSN, a classification system designed to address the population diversity problem. We demonstrated that the population diversity problem appears when using conventional techniques with as few as 50 users. CSN combines the crowd-sourcing of labels and sensor data with multiple similarity networks that capture user similarities across different dimensions. The similarity networks guide the process of selectively merging data from different individuals to produce personalized classifiers at a much lower per-user cost. Finally, the generality, flexibility, and effectiveness of CSN are demonstrated using three distinct mobile sensing datasets.

References

1. Amazon Mechanical Turk. http://www.mturk.com
2. Google Nexus One. http://www.google.com/phone/detail/nexus-one
3. Nike+. http://www.apple.com/ipod/nike/run.html
4. SF-36.org. A community for measuring health outcoming using SF tools. http://www.sf-36.org/tools/SF36.shtml
5. Amazon Elastic Cloud Computing. http://aws.amazon.com/ec2
6. Guo B, Zhang D, Yu Z, Zhou X (2012) Hybrid SN: interlinking opportunistic and online communities to augment information dissemination. In: Proceedings of the 9th IEEE international conference on ubiquitous intelligence and computing, UIC'12

7. Abdelzaher T, Anokwa Y, Boda P, Burke J, Estrin D, Guibas L, Kansal A, Madden S, Reich J (2007) Mobiscopes for human spaces. IEEE Pervasive Comput 6(2):20–29
8. Andoni A, Indyk P (2008) Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. Commun ACM 51(1):117–122
9. Ashbrook D, Starner T (2003) Using GPS to learn significant locations and predict movement across multiple users. Pers Ubiquitous Comput 7(5):275–286
10. Bishop CM (2006) Pattern recognition and machine learning (information science and statistics). Springer, New York
11. Burke J, Estrin D, Hansen M, Parker A, Ramanathan N, Reddy S, Srivastava MB (2006) Participatory sensing. In: Workshop on world-sensor-web: mobile device centric sensor networks and applications, WSW '06, pp 117–134
12. Hattori Y, Inoue S, Hirakawa G. (2011) A large scale gathering system for activity data with mobile sensors. In: Proceedings of the 15th IEEE international symposium on wearable computers, (ISWC 2011)
13. Campbell AT, Eisenman SB, Lane ND, Miluzzo E, Peterson RA (2006) People-centric urban sensing. In: Proceedings of the 2nd annual international wireless internet conference, WICON '06
14. Consolvo S, McDonald DW, Toscos T, Chen MY, Froehlich J, Harrison B, Klasnja P, LaMarca A, LeGrand L, Libby R, Smith I, Landay JA (2008) Activity sensing in the wild: a field trial of UbiFit Garden. In: Proceedings of the 26th annual SIGCHI conference on human factors in computing systems, CHI '08, pp 1797–1806
15. Dipietro L, Caspersen C, Ostfeld A, Nadel E (1993) A survey for assessing physical activity among older adults. Med Sci Sports Exerc 25:628–628
16. Froehlich J, Dillahunt T, Klasnja P, Mankoff J, Consolvo S, Harrison B, Landay JA (2009) UbiGreen: investigating a mobile tool for tracking and supporting green transportation habits. In: Proceedings of the 27th international conference on human factors in computing systems, CHI '09, pp 1043–1052
17. Kapoor A, Horvitz E (2008) Experience sampling for building predictive user models: a comparative study. In: Proceedings of the 26th annual SIGCHI conference on human factors in computing systems, CHI '08, pp 657–666
18. Lane ND, Xu Y, Lu H, Hu S, Choudhury T, Campbell AT, Zhao F (2011) Enabling large-scale human activity inference on smartphones using community similarity networks (CSN). In: Proceedings of the 13th international conference on ubiquitous computing, Ubicomp '11, pp 355–364
19. Lane ND, Lu H, Eisenman SB, Campbell AT (2008) Cooperative techniques supporting sensor-based people-centric inferencing. In: Proceedings of the 6th international conference on pervasive computing, Pervasive '08, pp 75–92
20. Lane ND, Xu Y, Lu H, Campbell AT, Choudhury T, Eisenman SB (2011) Exploiting social networks for large-scale human behavior modeling. Pervasive Comput 104:45–53
21. Lane ND, Mohammod M, Lin M, Yang X, Lu H, Ali S, Doryab A, Berke E, Choudhury T, Campbell AT (2011) BeWell: a smartphone application to monitor, model and promote wellbeing. In: Proceedings of the 5th international ICST conference on pervasive computing technologies for healthcare, Pervasive Health 11
22. Lester J, Choudhury T, Kern N, Borriello G, Hannaford B (2005) A hybrid discriminative/generative approach for modeling human activities. In: Proceedings of the international joint conference on artificial intelligence, IJCAI '05, pp 766–772
23. Longstaff B, Reddy S, Estrin D (2010) Improving activity classification for health applications on mobile devices using active and semi-supervised learning. In: Proceedings of ICST conference on pervasive computing technologies for healthcare, PervasiveHealth '10, pp 1–7
24. Lu H, Pan W, Lane ND, Choudhury T, Campbell AT (2009) Soundsense: scalable sound sensing for people-centric applications on mobile phones. In: Proceedings of the 7th international conference on mobile systems, applications, and services, Mobisys '09, pp 165–178
25. Lu H, Yang J, Liu Z, Lane ND, Choudhury T, Campbell AT (2010) The Jigsaw continuous sensing engine for mobile phone applications. In: Proceedings of the 8th international conference on embedded networked sensor systems, Sensys '10, pp 71–84
26. Mun M, Reddy S, Shilton K, Yau N, Burke J, Estrin D, Hansen M, Howard E, West R, Boda P (2009) PEIR, the personal environmental impact report, as a platform for participatory sensing systems research. In: Proceedings of the 7th international conference on mobile systems, applications, and services, MobiSys '09, pp 55–68
27. Oza NC, Russell S. (2001) Online bagging and boosting. In: Proceedings of the 8th international workshop on artificial intelligence and statistics, AISTAT '01, pp 105–112
28. Miluzzo E, Oakley JM, Lu H, Lane ND, Peterson R, Campbell A. (2008) Evaluating the iPhone as a mobile platform for people-centric sensing applications. In: Proceedings of international workshop on urban, community, and social applications of networked sensing systems, UrbanSense 08
29. Peebles D, Lu H, Lane N, Choudhury T, Campbell A (2010) Community-guided learning: exploiting mobile sensor users to model human behavior. In: Proceedings of the 24th national conference on artificial intelligence, AAAI '10, pp 1600–1606
30. Yan T, Kumar V, Ganesan D (2010) CrowdSearch: exploiting crowds for accurate real-time image search on mobile phones. In: Proceedings of the 8th international conference on mobile systems, applications, and services, MobiSys '10, pp 77–90
31. Rachuri KK, Musolesi M, Mascolo C, Rentfrow PJ, Longworth C, Aucinas A (2010) Emotionsense: a mobile phones based adaptive platform for experimental social psychology research. In: Proceedings of the 12th ACM international conference on ubiquitous computing, Ubicomp '10, pp 281–290
32. Stikic M, Van Laerhoven K, Schiele B (2008) Exploring semi-supervised and active learning for activity recognition. In: Proceedings of the 12th IEEE international symposium on wearable computers, ISWC '08, pp 81–88
33. von Ahn L, Maurer B, Mcmillen C, Abraham D, Blum M (2008) reCAPTCHA: human-based character recognition via web security measures. Science, pp 1465–1468
34. Rubner CTY, Guibas LJ (2000) The earth movers distance as a metric for image retrieval. IJCV 40(2):99–121
35. Zheng Y, Li Q, Chen Y, Xie X, Ma WY (2008) Understanding mobility based on GPS data. In: Proceedings of the 10th ACM international conference on ubiquitous computing, Ubicomp '08, pp 312–321
36. Zheng Y, Zhang L, Xie X, Ma WY (2009) Mining interesting locations and travel sequences from GPS trajectories. In: Proceedings of the 18th international conference on world wide web, WWW '09, pp 791–800
37. Zhu X, Goldberg AB (2009) Introduction to semi-supervised learning. Morgan and Claypool, San Rafael