
ContextSense: Unobtrusive Discovery of Incremental Social Context using Dynamic Bluetooth Data

Zhenyu Chen^{1,4,5}, Yiqiang Chen¹, Lisha Hu^{1,5}, Shuangquan Wang¹, Xinlong Jiang^{1,5}, Xiaojuan Ma², Nicholas D. Lane³, Andrew T. Campbell⁴

¹Beijing Key Laboratory of Mobile Computing and Pervasive Device, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

²Noah's Ark Lab, Hong Kong, China

³Microsoft Research, Beijing, China

⁴Dartmouth College, Hanover, NH, USA

⁵University of Chinese Academy of Sciences, Beijing, China

{chenzhenyu, yqchen, hulisha,
wangshuangquan, jiangxinlong}@ict.ac.cn
xiaojuan.ma@huawei.com
niclane@microsoft.com
campbell@cs.dartmouth.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

UbiComp'14 Adjunct, September 13-17, 2014, Seattle, WA, USA

Copyright 2014 ACM 978-1-4503-3047-3/14/09...\$15.00.

<http://dx.doi.org/10.1145/2638728.2638801>

Abstract

User-centric ambient social contexts can be effectively captured by dynamic bluetooth data. However, conventional approaches for training classifiers struggle with social contexts that are incrementally constructed and continuously discovered in everyday environments. Incremental social contexts can confuse a classifier because it assumes that the number and composition of context classes is fixed throughout training and inference phases. To address this challenge we propose *ContextSense*, an ELM-based learning method for continuous and unobtrusive discovery of new social contexts incrementally from dynamic bluetooth data. Experimental results show that ContextSense can automatically cope with “incremental social context” classes that appear unpredictably in the real-world.

Author Keywords

Bluetooth; Incremental Social Context; Fuzzy Clustering; Class-Incremental Learning; Extreme Learning Machine

ACM Classification Keywords

H.5.m [Information Interfaces and Presentation]: Miscellaneous; I.5.2 [Pattern Recognition]: Design Methodology

Introduction

Context-aware computing has been drawing increasing interest from a wide variety of fields including location-based social networks and personalized recommendation because of the emerging importance of mobile applications and services like Google Now.

Existing research into context-awareness has extensively studied “place” recognition. For example, [5, 6] proposes several methods to identify semantically meaningful locations. Common approaches assume static radio signals, such as GPS co-ordinates and WiFi fingerprinting; but such techniques struggle when faced by signal instability. Alternative approaches have made use of light and sound fingerprints generated by cameras or microphones embedded in mobile phones [7, 1]. Such work has achieved some good results in clustering similar place categories or recognizing user meaningful events, but robustness is limited, as it is sensitive to the placement of the mobile device. Moreover, this type of data may fail to capture broader social context data including the dynamic group behavior of multiple people. In this paper, we explore the use of bluetooth data because it is readily available (from user’ smartphones) and can capture important social interactions as encoded by users being in close proximity to each other.

Importantly, in many key contexts (e.g., home or work) new “incremental social context” classes can appear unpredictably due to new behaviors and entities (objects and people). Social contexts might appear gradually with small differences to existing classes (i.e., context class drift). For example, through the introduction of a new device or person during an activity. Or alternatively, completely new social contexts might appear. For example, with a new action being performed or a new

social group forming in a particular location. Conventional classifiers for recognizing social context classes fail in such environments because they can not discover new social contexts by themselves; only social contexts present in training data can be detected.

In this paper, we focus on the challenge of coping with new classes of social context that unpredictably appear or change. We refer to these as “**incremental social context**”. To address this challenge, we propose a new approach called *ContextSense*. It aims to incrementally infer new classes of social context by only observing the patterns and changes of ambient bluetooth signals, and automatically update the classifier accordingly.

ContextSense Design

The aim of ContextSense is to recognize various kinds of ambient incremental social context by analyzing the dynamic bluetooth logs collected by smartphones using machine learning algorithms. Our **ContextSense** framework has four stages when analyzing bluetooth data. First, it extracts features from bluetooth data and trains a bootstrap extreme learning machine (ELM) classifier. Second, it uses a one-class classification algorithm to isolate and accumulate initially unknown but increasingly distinguishable social context data (Figure 1(a)). Third, it discovers new cluster candidates using a fuzzy clustering algorithm (Figure 1(b)). Fourth, it consolidates the newly discovered classes using an ELM based class-incremental approach (Figure 1(c)).

Feature extraction

We extract six features [2] from the bluetooth-related data collected by users’ smartphones [3] in each sampling period: {number of bluetooth devices detected in radio neighborhood, ratio of new and old devices within radio




Figure 1: ContextSense Data-flow.

neighborhood, the mean and standard deviation for the rate of change for new devices neighbors}.

One-Class Classification

NNDD (nearest neighbor data description) is a simple method in which a new sample is assigned to the class of its nearest neighbor. KNNDD is an extension to NNDD. This method is defined to compute a single distance value between a data point and its k -nearest neighbors: (1) the distance to the k -th nearest neighbor; (2) the distance to the average of the k -nn's; (3) average distance to the k -nn's. We use the first distance definition here. This one-class classification algorithm discriminates the samples of conventional social contexts by considering them as a single class, from the outliers which are potential incremental social context data points (Figure 1(a)).

Fuzzy ISODATA Clustering

F-ISODATA (Fuzzy Iterative Self-Organizing Data Analysis Techniques Algorithm) extends the K-means algorithm by adding merge and split operations. It also employs a mechanism for setting the control parameters for the clustering process. F-ISODATA can be seen as a transformation of ISODATA's hard partition into a soft partition. In other words, a probability is given to the cluster membership of each sample. We try to cluster the outliers using the fuzzy ISODATA algorithm, as shown in Figure 1(b). It forms several initial clusters via unsupervised clustering, and automatically learns the most suitable number of clusters by merging or splitting the initial groups. Finally, each sample falls under every one of the final clusters with a probability due to fuzzy clustering.

ELM-based Class-Incremental Classifier

Extreme Learning Machine (ELM) has been proposed for a Single Hidden Layer Feed Forward Network (SLFN) for both classification and regression problems [4]. Recently,

Class Incremental Extreme Learning Machine (CIELM) was proposed [8]. It has been shown to approximate the performance of batch learning ELM, and is employed to recognize both conventional social contexts and new incremental social contexts in this paper (shown in Figure 1(c)). The learning process of CIELM has three main steps: (1) collect and label a set of samples; (2) learn an initial ELM model on these samples; (3) update the initial model through the incremental learning algorithm with the incoming data labeled as known or new classes. After the adaptation phase, the model adds extra output neurons to the structure of ELM.

Experiment Results

Data Collection

We develop a program to collect bluetooth radio traces on Android smartphones. The sampling rate is set to 15 seconds. We recruit 13 participants to collect bluetooth radio traces in their everyday life for 3 to 4 weeks. We choose five representative social contexts for training: *meeting in the hall* (Label 1), *working indoors* (Label 2), *taking the subway* (Label 3), *shopping in the mall* (Label 4), *watching a movie in the cinema* (Label 5).

Parameter Settings

We consider a subset of social context classes as *targets* (Label 1, 2 and 3) and the remaining social contexts as *outliers* (Label 4 and 5). We treat *targets* as known classes and *outliers* as unknown classes. Figure 2 presents the ROC (receiver operating characteristic curve) when separating *targets* and *outliers* using the KNNDD algorithm. We select the threshold for the one-class classifier when the True Positive rate (TPr) is equal to 90%, because: (1) according to the *ContextSense* framework, *outliers* will be the input of F-ISODATA algorithm, and thus if TPr is low, some samples of the




Figure 2: ROC of one-class classification algorithm.




Figure 3: Threshold of fuzzy ISODATA clustering algorithm.




Figure 4: Performance comparison for batch-mode ELM and class-incremental ELM.

known classes will be incorrectly recognized, which might seriously impair the performance of the following steps; (2) if TPr is set too high, samples of the unknown classes (*outliers*) may be left out of the next step; (3) when TPr is 0.9, overall performance tends to be stable (Figure 2).

We use the fuzzy ISODATA (F-ISODATA) clustering algorithm to mine the number of clusters. According to F-ISODATA, when the fuzzy standard deviation of one cluster is larger than the separation threshold, this cluster will be divided into two clusters. In Figure 3, the separation threshold ranges from 0.001 to 0.8, and the number of clusters clearly decrease as the threshold increases. We find the five discovered clusters correspond to ground-truth when the threshold is 0.1.

Comparison

Figure 4 compares the testing accuracy of batch-mode ELM to CIELM. We again set Label 1, 2, and 3 as known classes and Label 4 and 5 as unknown classes. In the feature extraction phase of ContextSense, we use labeled samples of known classes to train an ELM model. This model is also treated as the bootstrap model for CIELM. Therefore ELM and CIELM models have initially the same structure and the inference accuracy (91.45%), assuming Label 4 and 5 are considered as the unknown classes by both methods. Under batch-mode ELM the testing accuracy declines greatly because the ELM model does not perform incremental learning and so can not recognize unknown classes. In comparison, our CIELM can almost maintain the classification performance even when recognizing one or two new classes. Specifically, the accuracy of batch-mode ELM is 61.90% and CIELM is 86.93% for one new class (Label 4). This accuracy falls to 53.90% (ELM) and 84.80% (CIELM) for two new classes (Label 4 and 5). In order for batch-mode ELM to cope

with unexpected classes it would need to be frequently re-trained; this requirement does not scale to real-world environments with high levels of class churn.

Conclusion

In this paper, we propose *ContextSense*, an ELM-based incremental learning method for continuous and unobtrusive discovery of new unpredicted social contexts based on dynamic bluetooth data. Our approach starts with a one-class classification algorithm to identify outliers of known classes of social context, before applying a fuzzy clustering algorithm to group outliers into new classes. Experimental results show that *ContextSense* can automatically and continuously recognize fresh new incremental social contexts from real-world environments.

Acknowledgements

This work is supported by Natural Science Foundation of China (61173066), Beijing Natural Science Foundation (4144085), Strategic Emerging Industry Development Special Funds of Guangdong Province (2011912030), National Science and Technology Major Project (2012ZX07205-005), open project of Beijing Key Lab of Mobile Computing and Pervasive Device.

References

- [1] Azizyan, M., et al. SurroundSense: Mobile Phone Localization via Ambience Fingerprinting. In *Proc. MobiCom'09*, ACM (2009), 261–272.
- [2] Chen, Y. Q., et al. Surrounding Context and Episode Awareness using Dynamic Bluetooth Data. In *Proc. UbiComp'12*, ACM (2012), 629–630.
- [3] Chen, Z. Y., et al. Inferring Social Contextual Behavior from Bluetooth Traces. In *Proc. UbiComp'13*, ACM (2013), 267–270.
- [4] Huang, G.-B., et al. Extreme Learning Machine: Theory and Applications. *Neurocomputing* 70 (2006), 489–501.
- [5] Kim, D. H., et al. Discovering Semantically Meaningful Places from Pervasive RF-Beacons. In *Proc. UbiComp'08*, ACM (2008), 21–30.
- [6] Kim, D. H., et al. SensLoc: Sensing Everyday Places and Paths using Less Energy. In *Proc. SenSys'10*, ACM (2010), 43–56.
- [7] Lu, H., et al. Soundsense: Scalable Sound Sensing for People-Centric Applications on Mobile Phones. In *Proc. MobiSys'09*, ACM (2009), 165–178.
- [8] Zhao, Z. T., et al. A Class Incremental Extreme Learning Machine for Activity Recognition. *Cognitive Computation* (2014), 1–9.