Proceedings of the 28th IEEE
EMBS Annual International Conference
New York City, USA, Aug 30-Sept 3, 2006

Sensors to Detect the Activities of Daily Living

Beth Logan and Jennifer Healey

Abstract— **We study the use of embedded and worn sensors to unobtrusively detect the activities of daily living (ADL). Our aim is to find the minimum set of sensors required to detect these basic tasks. In this exploratory work, we analyze the publicly available 'Intense Activity' dataset from the MIT PlaceLab project and study the classification of eating and meal preparation vs. other activities. We find that eating and meal preparation can be detected with an accuracy of 90% using less than 1/3 of the over 300 available sensors in the PlaceLab. If only 8 sensors are used, the accuracy is 82% which may be adequate for some applications.**

I. INTRODUCTION

Healthcare in western countries is currently in an unstable state. Rising costs and increasing numbers of elderly people will overburden existing systems before long. A proposed solution is to move health management away from high-cost, highly invasive acute care toward low cost, minimally intrusive preventative healthcare in the home.

Although some home healthcare systems will be tailored to specific chronic diseases, others will be more general 'ok-ness checking' systems. These systems unobtrusively monitor the daily activity patterns of people in their homes to assess the state of their health.

In this paper we explore how many and which sensors are required for simple ok-ness checking. Reducing the number of sensors has clear cost and maintenance advantages so it is interesting to examine the effects on accuracy. For some applications, reduced accuracy may be justified.

A. Prior Work

There are many existing commercial and research prototype systems for OK-ness checking, each with limitations in scope and usefulness. Some examples are given below.

1) Commercial Systems: A commercial example of a system that uses only one sensor is the Japanese 'i-pot' (information pot) which not only boils water for soup or tea but it also records the times this occurred [1]. This signal is sent to a server and caregivers can subscribe to a service to see recent records of i-pot usage. Although this system is non-invasive and easy to use, it provides only limited information and assumes that the user regularly uses the teapot.

A second type of commercial system provides ok-ness checking using small wireless motion sensors placed in key areas of the home [2]. Each sensor transmits information 24 hours a day, 7 days a week about the senior's activities to a base station which transmits over the telephone network to a centralized service. Changes in the senior's activities are then analyzed so caregivers can be alerted to problems. The information collected is quite extensive. A drawback however is that the motion sensors may be activated by pets.

Several companies market 'alert services' (e.g. [3]) These systems require the elder to push a button to receive assistance. Drawbacks include the fact that the elder must wear the device, be alert to push the button, and must want to push it.

2) Research Systems: Several research groups have developed systems to monitor activities in the home. The majority of these are based on worn or embedded sensors. In one experiment, Tapia et al. use around 80 wireless accelerometers installed on objects in the home to sense around 30 activities [4]. They report accuracies better than random although in many cases worse than 50% for data collected over a two week test period in two homes. A problem faced by these researchers was less than accurate labeling as the data was self-labeled by the study participants and no video or audio was recorded.

The smart home project at MARC has examined detecting activities using a combination of embedded and wireless sensors. For example Dalal et al. [5] applied automatic and hand-crafted rules to detect meal preparation from 23 embedded sensors including motion sensors, switches and temperature sensors. Results from a 37 day study with one participant show good performance with the hand-crafted rules out-performing the automatic technique. An analysis of the data indicated the best sensors for this task. However, again, the activities for this study were self-labeled by the study participant so the results may be inaccurate. Also, using hand-crafted rules may not scale to many users, homes and activities.

RFID sensors have proved useful for detecting a variety of fine-grained activities. Patterson et al. report accuracies of around 80% for recognition of 12 activities in a kitchen instrumented with 60 RFID tags [6]. A potential disadvantage of this system is that the user must wear a RFID reader, in this case a glove.

Typically video and audio are not used in in ok-ness checking systems since state-of-the-art technologies are currently inadequate for this task. Additionally most people find audio and video recording too invasive for home use.

B. Our Approach

We propose to analyze which sensors are best for unobtrusively monitoring the activities of daily living (ADLs). These are basic activities that healthy people undertake each day, such as eating, using the bathroom and the like. Our

Advanced Technology Group, Intel Digital Health, One Cambridge Center, Floor 11, Cambridge MA 02142 USA, beth.logan@intel.com

ultimate aim is to determine the best set of sensors according to various criteria such as cost, number of house occupants, presence of pets and the like.

For our experiment, we use the over 300 sensors installed in the MIT 'PlaceLab' [7]. This allows us to cover a large space of potential sensors. A key advantage of our study over previous work is that we have well annotated ground truth. Also, unlike the work of Dalal et al., all types of sensors are located in all rooms of the PlaceLab rather than just the kitchen[1]. We also use an order of magnitude more sensors than this earlier work.

II. Activities of Daily Living

'Activities of Daily Living' (ADLs) is a standard term for the basic everyday tasks of life such as eating, bathing, dressing and toileting [8]. Performance of these tasks is used as a measure for independent living; when people can no longer perform these tasks unassisted, help is required from caregivers. The tasks are as follows:

1) Bathing
2) Dressing
3) Toileting (uses toilet, brushes teeth, etc)
4) Transfer (moves in and out of bed and chairs)
5) Continence
6) Eating

III. MIT PlaceLab

The MIT PlaceLab is a residential condominium in Cambridge, MA which has been designed to be an observational facility for the scientific study of people in home environments [7]. The condominium is instrumented with over 300 sensors which are installed in nearly every part of the home. These sensors range from switch sensors on lights, cupboards, electrical appliances faucets and the like, to 'MITes' which are wireless sensors attached to many objects in the apartment. There are also environmental sensors such as temperature and pressure, worn MITes containing accelerometers and a worn heart rate monitor as well as cameras and microphones mounted throughout the home. The data from all of these sensors is relayed to a central processing and storage facility in the apartment.

IV. Intense Activity Dataset

The 'Intense Activity' Dataset (IAD) is a publicly available dataset collected in the PlaceLab [9]. The data was recorded over a 4 hour time period during which one person performed a number of typical household tasks such as cooking a meal, doing laundry, making phone calls, cleaning and the like. The data is extensively labeled with Body Posture, Location, Social Context and Activity information. The labels were determined by independent annotators who watched the multiple video streams collected. The activity information is quite fine-grained. For example, activities include 'Sweeping', 'Folding Laundry' and 'Brushing Teeth'. Full information, including an ontology of labels is given on the dataset website.

The IAD consists of data from over 300 sensors. Over 100 of these are associated with wired switches for which changes of state are recorded. The rest of the sensors record streaming audio, video, heartrate, acceleration, humidity, pressure, illumination, temperature, gas and electrical current data. This streaming data is sampled at rates varying from once per minute to 78kHz. All the data is relayed to a central server which records the time of each event or sample. For the video data, the video compression process adds some latency. Therefore, a synchronization file is provided.

For further information on the labeling and nature of the data, refer to the the dataset website.

V. Experiments

We conduct a series of experiments on the IAD. Our aim is to evaluate the difficulty of detecting the activities of daily living and to investigate the trade-offs when using various subsets of sensors. Of particular interest is whether a minimum set of sensors can be used.

In this preliminary study, we focus only on the 'eating' ADL. Specifically, we build classifiers to detect whether the PlaceLab occupant prepared a meal and ate it. We assume that we are monitoring a mentally alert person who does not forget to eat a meal that they spend time preparing. Our choice of only one ADL is due to the limited size and nature of the IAD.

We follow a very standard procedure for data analysis. We first convert the disparate data in the IAD into a series of feature vectors at a constant sampling rate. We then learn a variant of AdaBoost classifiers to separate eating and associated tasks from other activities. Our choice of classifier automatically chooses the best subset of sensors to include for a given model complexity. Thus we are able to examine trade-offs in performance. The following sections detail our experiments.

A. Data Preparation

In its raw form, the IAD data is not suitable for use with most data classification algorithms. We therefore converted it to a series of frames formed by concatenating all the data at 0.1s sampling rate. We chose this sampling rate as a good compromise between sampling too infrequently to observe switches changing state and sampling too frequently and generating too large a data set. We converted each type of data to components of this feature vector as follows.

For the event driven data and data at low sampling rates, we assigned one component of the feature vector to each sensor. We then sampled this data at the 0.1s sampling rate. We assumed that the data held the current state or value between samples.

The MITes accelerometer data is sampled at 200Hz thus averaging is required to convert it to the 0.1s sampling rate. In addition, prior work indicates that the time domain is not the best representation for this data. Following then the methodology described in [10] we converted each

[1] In [5] all the sensors were in the kitchen except for the 6 infrared sensors which were installed in all rooms of the house.

TABLE I
PLACELAB ACTIVITY LABELS FOR CLASS 0 (EATING, MEAL PREPARATION AND CLEANUP)

chopping/slicing	cleaning a surface
dishwashing miscellaneous	disposing garbage
drinking	drying dishes
eating meal/snack	grating
hand-washing dishes	loading/unloading dishwasher
meal preparation miscellaneous	measuring
mixing/stirring	preparing food background
preparing meal	preparing snack
putting away dishes	retrieving ingredients/cookware
washing ingredients	putting things away
combining/adding	

axis/channel of accelerometer data to three features: energy, entropy and correlation computed over 2.56s windows overlapped by 1.28ms. For the 3-axis accelerometers used, this resulted in 9 features per MITes sensor which were then concatenated with the rest of the data at the 0.1s sampling rate.

The audio and video data is captured from 9 color cameras, 9 infra-red cameras and 19 microphones in the PlaceLab. To save storage, the Place Lab's central processing unit uses standard image processing algorithms based on pixel differences between frames to select the 4 video streams and 1 audio stream that best capture an occupant's behavior as inferred from motion activity and the known camera layout. In addition to these raw video and audio streams, the dataset also includes several timestamped files identifying the chosen streams. For the video data, we are also provided with the number of pixels that changed for each camera since the last frame of data.

Since it was out intuition that simply knowing where the occupant was would provide useful and potentially sufficient information about the activity being performed, we chose to use information from these auxiliary files rather than processing the raw data. Specifically, for each camera we created a dimension in our feature vector and recorded the pixel difference data. For each microphone, we created a dimension and recorded '1' if it was the 'chosen' one, otherwise '0'. Thus our video data represents crude pixel changes and hence motion and our microphone data represents the likely position of the occupant.

We did not incorporate the heart rate data into this study.

B. Labeling Data

For this study, we divide the data set into two classes: 0) the occupant is preparing a meal, eating it or cleaning up, and 1) other activities. Since the data is labeled at a finer grain than required, we combined the labels for a number of activities to determine Class 0. Table I lists the labels that were combined for Class 0. All other labels were assumed to describe Class 1. This methodology resulted in 34% of the features being labeled as Class 0 (i.e. eating/meal preparation) for the 4 hour dataset.

TABLE II
ERROR RATES FOR 10 FOLD CROSS VALIDATION FOR VARYING NUMBERS OF WEAK LEARNERS.

Number of Weak Learners	Error Rate (%)
10	17.6
100	12.7
1000	9.8
10000	8.5

C. Data Analysis

We use a modified version of AdaBoost proposed by Tieu and Viola [11] to learn classifiers for eating and associated activities vs. other activities. A key advantage of this technique is that in addition to classifying the data with competitive performance, it also identifies which sensors are more informative for making the classification decision.

Briefly, AdaBoost [12] is an iterative algorithm which learns a series of 'weak' learners. Each weak learner focuses on data that the previous series of weak learners failed to classify correctly. This is achieved by weighting the training data according to how well it is classified. The final 'strong' classifier is a weighted sum of the weak learners.

In the variant of AdaBoost proposed by Tieu and Viola, the weak learners are simple linear classifiers based on only one dimension of the feature vector. On each iteration of AdaBoost, we learn the best such weak learner thus the best single feature for classifying the data. Since each iteration of AdaBoost focuses on data that is harder and harder to classify, the weak learners chosen earlier by AdaBoost are the most informative for classifying the data. In our case, each dimension of our feature vectors represents data from a single sensor (or a channel of a sensor in the case of the accelerometer features). Thus an examination of the weak learners learnt on our data will identify the best sensors for detecting eating and meal preparation.

VI. RESULTS

We used the implementation of AdaBoost in WEKA [13] to learn classifiers with varying numbers of weak learners for our data. For computational reasons we were not able to run AdaBoost on the entire data set. Instead, we sampled every 100th feature vector.

Table II shows the error rates for 10-fold cross validation. From this we see that there is a substantial reduction in error rate as we use more weak learners, with the best error rate being 8.5% for 10000 weak learners. Therefore, at first glance it appears that our hope of finding a small subset of required sensors may be in vain.

However, when we examined the sensors used by the weak learners, we noticed that often sensors were reused. That is, AdaBoost again chose these sensors as the best weak learners for the task on subsequent iterations of the algorithm. Table III shows the number of unique sensors used in each experiment. We see that for example in the experiment with 100 weak learners, only 25 unique sensors

TABLE III

NUMBER OF UNIQUE SENSORS IN EACH EXPERIMENT.

Number of Weak Learners	Number of Unique Learners (Sensors)
10	8
100	25
1000	68
10000	104

TABLE IV

SENSORS USED BY FIRST 10 WEAK LEARNERS AND THEIR WEIGHTS.

Sensor	Weight
1WireCurrent,Dishwasher	0.99
1WireHumidity,Master bath sink attic	0.93
VideoDifferencer,PlaceLab21 Video input (above sink)	0.78
MITesOnBodyPL17_12_5,living room	0.76
MITesOnBodyPL17_12_5,living room	0.65
1WireHumidity,Kitchen microwave cabinets attic	0.6
MITesOnBodyPL16_12_3,kitchen island	0.63
MITesStatic,bedroom - bed	0.53
1WireCurrent,Dishwasher	0.47
MITesStatic,kitchen - left of microwave - cabinet door	0.45

were required. Similarly, in the 1000 weak learner experiment, only 68 sensors were required. It seems than that there is a minimum set of sensors required to detect eating and associated tasks which is much less than the number of sensors in the PlaceLab.

We now examine which sensors were used by the first 10 weak learners. These are listed in Table IV. In this table, MITesOnBodyPLaa_bb_cc refer to worn accelerometer sensors. aa is the PlaceLab sensor identifier. bb is the channel with 12 referring to the sensor worn on the right wrist and 11 referring to the sensor worn on the left thigh. cc is the feature number for this sensor. This is a number from 0 to 8. 0, 1 and 2 refer to the energy, entropy and correlation for the x-axis accelerometer respectively. Similarly, 3,4,5 and 6,7,8 refer to features for the y and z axis accelerometers. The last part of the MITes identifier refers to which receiver recorded the data. For example, 'MITesOnBodyPL17_12_5,living room' refers to the correlation of the y-axis accelerometer in the MITesOnBody sensor 17 worn on the right wrist as received by the received in the living room.

From Table IV we see that a number of the sensors such as the current flowing to the dishwasher and the video activity above the sink reflect our intuition of good sensors to detect eating and associated activities. However, we also see that several sensors are not located anywhere near the kitchen. For example, the second sensor monitors activity in the master bathroom. Such sensors represent *negative* indicators of eating and associated activities. If the occupant is in the bathroom, they are unlikely to be preparing or eating food.

Another point to note is that several sensors are repeated in the table reflecting that they are chosen several times by the AdaBoost algorithm as the best weak learner on that iteration. Note also that the sensors chosen are of various types. Some are embedded and some are on-body. Finally note that it is unclear if the motion represented by the OnBody sensors is more important than the location of the receiver that detected them. It is unsurprising however that wrist accelerometer activity conveys more information that thigh activity for this task.

VII. CONCLUSIONS AND FUTURE WORK

We have analyzed the 'Intense Activity' dataset available from the MIT PlaceLab with a view to determining if eating and meal preparation can be detected and if so, whether there is a minimal set of sensors required. We found that eating and meal preparation can be detected with an accuracy of 90% using less than 1/3 of the over 300 available sensors in the PlaceLab. If only 8 sensors were used, the accuracy is 82%.

We examined the best sensors for this task and found that they were a mix of embedded and worn sensors. Also, while some were located in the kitchen, others were located in different rooms and were thus negative indicators of meal activity.

In the future, we plan to examine ADLs other than eating. We also hope to collect larger data sets augmented with additional sensors.

In the long term we hope to extend this work to *proactive systems* which not only monitor activity but seek to influence a person's behavior for the good of their health (e.g. see [14] for a review). These systems require very fine-grained activity detection to enable for example guiding a person through the steps required to make a cup of tea.

VIII. ACKNOWLEDGMENTS

The authors wish to thank Stephen Intille for many helpful discussions.

REFERENCES

[1] Zojirushi corporation. [Online]. Available: www.mimamori.net

[2] Quiet care systems. [Online]. Available: www.quietcaresystems.com

[3] Alert one. [Online]. Available: www.alert-1.com

[4] E. Tapia, S. Intille, and K. Larson, "Activity recognition in the home using simple and ubiquitous sensors," in *PERVASIVE 2004*, 2004.

[5] S. Dalal, M. Alwan, R. Seifrafi, S. Kell, and D. Brown, "A rule-based approach to the analysis of elders' activity data: detection of health and possible emergency conditions," 2005.

[6] D. Patterson, D. Fox, H. Kautz, and M. Philipose, "Fine-grained activity recognition by aggregating abstract object usage," in *ISWC'05*, 2005.

[7] S. S. Intille, K. Larson, J. S. Beaudin, J. Nawyn, E. M. Tapia, and P. Kaushik, "A living laboratory for the design and evaluation of ubiquitous computing technologies," in *2005 Conference on Human Factors in Computing Systems*, New York, NY, 2004.

[8] S. Katz, A. B. Ford, R. W. Moskowitz, B. A. Jackson, and M. W. Jaffe, "Studies of illness in the aged: The index of adl: A standardized measure of biological and psychosocial function," *Journal of the American Medical Association*, vol. 185, no. 12, pp. 914–919, 1963.

[9] Placelab "intense activity" test dataset. [Online]. Available: http://architecture.mit.edu/house_n/data/PlaceLab/PlaceLab.htm

[10] L. Bao and S. S. Intille, "Activity recognition from user-annotated acceleration data," in *PERVASIVE 2004*, 2004.

[11] K. Tieu and P. Viola, "Boosting image retrieval," in *International conference on computer vision*, 2000.

[12] Y. Freund and R. E. Shapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *Computational Learning Theory: Eurocolt '95*, 1995.

[13] I. H. Witten and E. Frank, *Data Mining: Practical machine learning tools and techniques*. San Francisco: Morgan Kaufmann, 2005.

[14] E. Dishman, "Inventing wellness systems for aging in place," *IEEE Computer*, 2004.