Context Guided and Personalized Activity Classification System

James Y. Xu
Electrical Engineering Department
University of California, Los Angeles
56-125B Engineering IV Building
Los Angeles, CA 90095-1594
jyxu@ucla.edu

Yuwen Sun, Zhao Wang
Electrical Engineering Department
University of California, Los Angeles
56-125B Engineering IV Building
Los Angeles, CA 90095-1594
{sun8310111, wang2010}@ucla.edu

William J. Kaiser, Greg J. Pottie
Electrical Engineering Department
University of California, Los Angeles
56-125B Engineering IV Building
Los Angeles, CA 90095-1594
{kaiser, pottie}@ee.ucla.edu

ABSTRACT
Continued rapid progress in the development of embedded motion sensing enables wearable devices that provide fundamental advances in the capability to monitor and classify human motion, detect movement disorders, and estimate energy expenditure. With this progress, it is becoming possible to provide, for the first time, evaluation of outcomes of rehabilitation interventions and direct guidance for advancement of subject health, wellness, and safety. The progress in motion classification relies on both the performance of new sensor fusion methods that provide inference, and the energy efficiency of energy-constrained monitoring sensors. As will be described here, both of these objectives require advances in the capability of detecting and classifying the location and environmental context. Context directly enables both enhanced motion classification accuracy and speed through reduction in search space, and reduced energy demand through context-aware optimization of sensor sampling and operation schedules. There have been attempts to introduce context awareness into activity monitoring with limited success, due to the ambiguity in the definition of context, and the lack of a system architecture that enables the adaptation of signal processing and sensor fusion algorithms specific to the task of personalized activity monitoring. In this paper we present a novel end-to-end system that provides context guided personalized activity classification. With a refined concept of context, the system introduces interface models that feature a unique context classification committee, the concept of context specific activity classification, the ability to manage sensors given context, and the ability to operate in real time through web services. We also present an implementation that demonstrates accurate context classification, accurate activity classification using context specific models with improved accuracy and speed, and extended operating life through sensor energy management.

Keywords
context guided, activity classification, context specific classification, system architecture, real time classification

1. INTRODUCTION
The rapid advance in microelectronics has provided MEMS inertial sensors, low power processors, and low cost monitoring systems applicable to human motion classification. Many of the most urgent problems in health and wellness promotion, diagnostics and treatment of neurological condition and even athletic performance advancement are now possible. The wireless health community exploits this along with smartphone technology for an integration of monitoring and in field guidance for both advancing and evaluating treatment outcomes.

Recently developed solutions monitor a subject's physical activity, for example walking gait speed monitoring for recovering stroke patients in the field of wireless healthcare [1,2]. For many applications, there is also a need for personalized, targeted monitoring for specific activities, in specific environments. For example a stroke patient benefits from monitoring of gait speed while in the hospital and then at home to ensure that their mobility is sufficient to enable safe passage through urban areas. Also, these subjects benefit from monitoring and guidance for aerobic exercises while at home to maximize the effectiveness of recovery routines [3].

A large body of work has focused on the accurate detection of physical activities, using a diverse range of classification and feature extraction techniques [1,2,4,5,9,10-12,31]. These methods confront the challenge of classification of a specific, correct motion among many possible at any observation time. As the number of potential motion increases, classification reliability is degraded. In addition, lacking any other source of guidance for scheduling, these methods require continuous operation of all platform components. This latter requirement limits operating life of energy-constrained wireless platforms.

In fields including wireless sensor networking, pervasive computing, and others, the concept of context-awareness has been introduced with the objectives of improving human machine interaction, and enabling low energy operation while retaining system performance. Many architectures have been proposed to bring personalization and adaptation to a system [6], and recent attempts have been made to introduce context to activity classification [7,8]. These systems experienced limited success due to ambiguity in the definition of context, and a lack of an appropriate system architecture that is specific to the task of personalized activity monitoring.

In addressing the above-mentioned deficiencies, we propose a novel end-to-end system that provides context guided personalized activity classification. Our focus is on three major areas: 1) The ability to accurately detect context with multiple sensing modes; 2) The use of context to improve classification accuracy, speed, and energy usage; and 3) The ability to target specific physical activities of interest, given context. The development of this system also addresses the problems associated with the operation of experimental systems required for system training and validation. In support of the three areas of innovation above, this paper introduces some major contributions: 1) Context guided personalized activity classification; and 2) An architecture for real time end-to-end rapid development and operation.

2. RELATED WORK

Many investigations in medical science over the last decade have demonstrated the critical benefits of activity monitoring for applications ranging from health and wellness promotion to disease treatment, to performance advancement and injury risk reduction in athletics. One example is the use of motion and sound data sources in an application that provides telemonitoring for elderly individuals living independently [9]. Here, a method was developed that can detect when a user requires attention (as a result of a fall or long periods of inactivity). In another study, accelerometer sensor data sources and machine learning algorithms were applied for monitoring intervention effectiveness of acute stroke patients [1]. The technology provides physicians with the ability to directly measure a patient's activity level, even after discharge. This improves on the surrogate laboratory measurements, administered only in a clinical setting. An example of applications in athletics were presented in [10], where multiple accelerometers were used for ambulatory monitoring of elite athletes in both competitive and training environments. For swimmers, the characteristics of strokes can be captured and analyzed. For rowers, the addition of an impeller combined with accelerometer data was used to recover intra and inter stroke phases for performance analysis. This system was used by Australian Olympic athletes in training for competition in the 2004 Olympic Games.

Using sensors for activity monitoring has been studied extensively. In [11], a system using iPhone and Nike+iPod sport kit was proposed for classifying human activities. The activities considered include running, walking, bicycling, and sitting. In [4], a complex environment with many microphones, video sources and other sensors was designed. The study attempted to accurately track movements of arms and hands. Activities considered there are bathing, dressing, toileting, eating, and others. Results indicated that using one third of the 300 available sensors in the specially designed lab, tasks can be detected with an accuracy of 90%. A specially designed glove was introduced for activity classification [5]. The glove detects and records objects a user touches using an RFID reader. In this system, all the objects being monitored (such as utensils, toothbrushes, and appliances) need to have RFID tags instrumented.

Most of the studies above are restricted in the number of activities they can detect accurately. These systems are designed either for a specific set of activities that may not be easily modified, or have a high system installation cost with the requirement to modify environments and also monitor subjects only when they are present in these environments.

The recognition of user and environment context has been identified as a primary capability for advancing the performance and capability of human-computer interfaces in many fields [7]. Studies have emerged recently in wireless health that attempt to combine context and activity classification. In [12], a multi-sensor wearable system was proposed that enables a context that largely consists of physical activities. There, 30 sensors were embedded into a garment, with multiple processing nodes responsible for distributed processing of sensor data. This study treated physical activities as contexts, and focused on the sensor fusion development. A system for a context-aware mobile phone named Sensay was developed [8]. This included context defined as a set of user states (normal, idle, uninterruptable). By introducing light, motion and microphone sensors, Sensay is able to detect these contexts and manipulate ringer volume, vibration, and phone alerts.

It is important to note that the definition of context has varied between investigations. It is particularly important to define context with the requirement that context states do not themselves contain the very activities that are to be detected. Definition of context for new development reported here will be included below.

3. SYSTEM DESIGN

3.1 Context

When addressing context, many investigations use the important definition by Dey [13]. While powerful, the definition of a context that includes every characteristics of a given situation, in terms of both the environment and the user, is very broad. Useful for some applications, it is not suitable for leveraging context in monitoring physical activities, as in many cases a context contains physical activities that are underlying in the definition. There are a number of alternative definitions available in the field of pervasive computing, offering different selection of divisions, such as external and internal contexts [14,15]. These definitions are usually narrower, but still contain a mix of physical activities with other environmental attributes.

In this study, a context is defined thusly: "*a context is a subset of all attributes that characterizes an environment or situation, external to the user"*. This definition clearly distinguishes between the external environment, and the user's physical activities. By means of this context, we will be able to provide a guideline for deciding which attribute is associated with context and which is associated with physical activity. For example, a "meeting" environment is a context, and its characteristics may involve certain sound profiles and a set of possible locations. "Sitting in a meeting" in contrast is not a context, as it contains the user's physical activity of "sitting".

3.2 System Model

We have developed and report here a system that is able to provide context guided activity classification, with the capability for real time online operation. To enable a context guided system, we must provide ways of discovering a user's context and ways for this information to be consumed by activity classification systems. To enable guided training and online classification in real time, we must also provide a client for the end user and a corresponding web service to interface with the rest of the system. Figure 1 depicts the architecture of this new system.

This architecture provides context detection and activity classification, where the context information is utilized by an activity classification system, along with activity sensor data. The client application is used for collecting sensor data and labels from a user, and also for displaying results. A corresponding web service runs on the server, and acts as a gateway between the client and the core system. This provides a means for the client to transmit and access data through a network, in a structured manner.

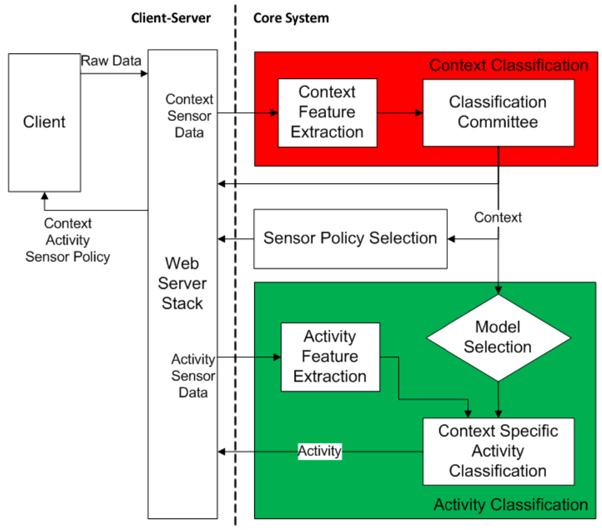

Figure 1 System model

3.3 Modular Object Oriented Design

It is important to note that at each step in design and implementation, individual subsystems are modeled as objects, and the entire system is defined by a set of interfaces and relationships. Each software interface is characterized by their public methods, defined by functionality, expected inputs and expected outputs. By implementing an interface, a class agrees to provide all the methods of that interface [16]. In this way, each subsystem is developed independently without the requirement to reveal its specific implementations, but only that it implements the required interface. This allows any part of our proposed system to be overridden by custom realizations, allowing for rapid prototyping and evaluation of various algorithms. Figure 2 demonstrates the interface model for our proposed core system.

As an example, consider the *IContextClassifierEx* interface, which covers the context classifier. Teams can develop classifiers independently and optimize them according to particular applications. As long as the classifiers provide the *getContext* method, they can be hot swapped into the system to adjust the classification system behavior without affecting the overall system.

3.4 Context Classification

Our definition of context can capture a large number of situations, so that users with different objectives can define their own sets of useful contexts. They can then identify required characteristics, and select necessary sensors. This generalization makes classification difficult, as we need to account for a diverse range of data sources such as GPS coordinates, wireless information, audio, and illumination level. For example, consider the following datasets (Table 1), demonstrating the result of an experimental system that combines both audio signal processing on sound detected in an environment, along with the Media Access Control (MAC) addresses associated with wireless access points detected in the same environment. Here, three locations were used, and MAC addresses and peak frequency of the audio power spectral density (PSD) were recorded.

Table 1 Example datasets

Label	Wireless MAC address	Audio PSD peak frequency
A	{30:46:9a:06:4d:e0, 00:0c:41:6e:1e:f6}	29.01KHz
B	{2e:25:b3:96:d5:f9, 00:b0:d0:86:bb:f7}	80.91KHz
C	{6e:51:f5:c1:11:00, 00:0c:f1:56:98:ad}	20.19KHz

To separate labels A, B and C in this example, a common method is to find thresholds that divide them, based on the MAC address and PSD peak frequencies. It is clear that the audio peak frequency data for environments can be assembled, and a distinct

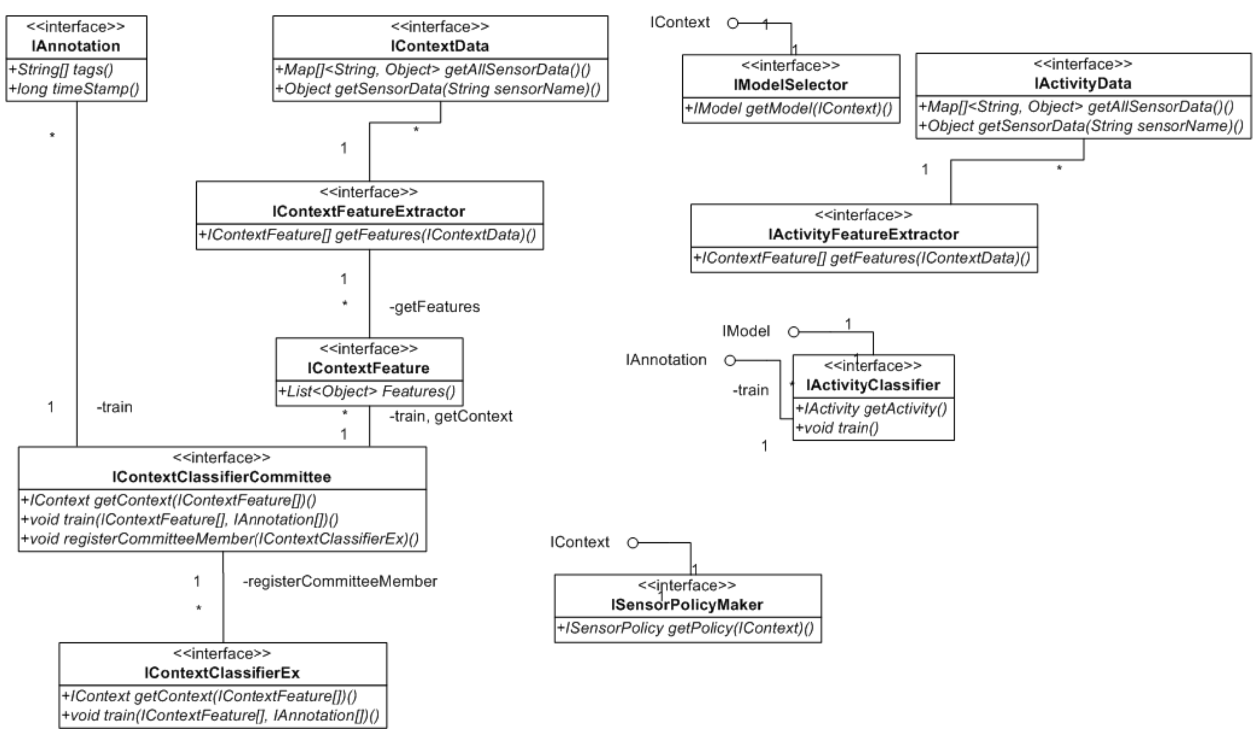

Figure 2 System interface model

separation between the labels can be found. The proper treatment of MAC addresses is less clear. It is challenging to represent these identifier values in a feature space, and to define a separation. This difference in data type determines the suitability of various classifiers. For example, classifiers based on SVM are not suitable for treatment of MAC addresses, whereas a method such as kNN has been used successfully [17].

In order to detect context based on a variety of data sources, there is a need to use multiple classifiers for different features. This paper describes the development of a classification committee consisting of n individual classifiers (Figure 3). The individual classifiers are trained separately, and after training they can be tested for individual accuracy. A voting weight (α) can be determined for each classifier, proportional to the perceived accuracy. When an unknown class is encountered, the committee performs a linear combination of the individual classifiers, and the context with the highest vote is the output.

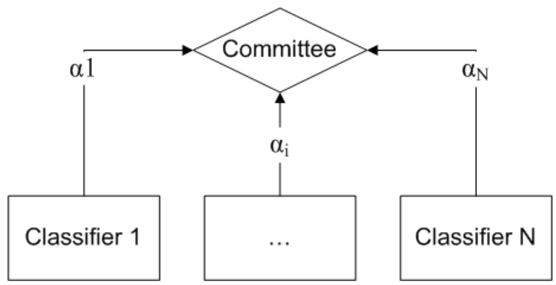

Figure 3 Classifier committee

This committee approach not only allows fusion of a number of sensors with various data types, but also allows for adaptation of context detection to individuals with varying habits. For the former, it is easy to see that different classifiers can be selected to compose the committee, depending on inputs. For the latter, suppose a habitual individual exhibits strong patterns in time of day relating to context. The weight of a classifier based on time of day will be adjusted during training so that the habitual subject would have a time classifier with higher vote weighting (compared to a subject that is irregular).

3.5 Context Guided Personalized Activity Classification

Once a context is obtained, the next step is to take advantage of this information in activity classification. We introduce the concept of a context guided classifier. These classifiers allow us to have specifically optimized models that each focus on the activities of interest, given context. Unlike conventional activity monitoring, there is no single list of comprehensive activities that needs to be built into a monolithic classifier. Instead, a basic set of activities common across all contexts can be chosen, and this set can then be extended or reduced should the need arise for a particular context. To illustrate this concept, consider some usage scenarios in Table 2.

Based on context information, the model selector would select an appropriate scenario (model), and the activity classifier can then make a classification based on the model. There are a number of benefits from using this system. First, we can improve classification accuracy and speed due to a simplification of feature space. Then, the system allows scenarios (models) to be determined by investigators, and a person's monitoring program can be modified. For example, physicians may wish to monitor a stroke inpatient's walking speed and also ensure they are intermittently sitting/standing to alleviate deterioration in exercise tolerance [3] (Scenario 3). Once discharged, physicians may then wish to monitor the patient to make sure that recommended daily activities are performed at home (Scenario 4). Another example is where personal trainers can prescribe personalized training plans, including activities and their duration and place (gym, home, office). Here the activity monitoring system can inform the user of his/her training progress. Finally, the system gives a user the ability to control his/her privacy. Unlike most other monitoring systems that are always on, a user can decide to only allow monitoring under specific contexts, for specific activities.

Table 2 Example scenarios

Scenario 1	Scenario 2
Basic Set • Walking • Sitting • Standing	Cafeteria • Walking • Sitting • Standing • Eating
Scenario 3	**Scenario 4**
Stroke patient (inpatient) • Standing • Sitting • Walking • Walking Slow • Walking Fast	Stroke patient (rehabilitation) • Aerobic activity • Walking • Walking Fast • Walking Slow

3.6 Sensor Control and System Training

By having context and context guided scenarios, we can also optimize sensor sampling rate and selectively enable or disable sensors to reduce energy demand. For example, there are no upper body motions from scenario 1 (Table 2), and the activities have a low rate of change. This means that upper body sensors could be disabled and sampling rate can be reduced on the lower body sensors at no loss in system performance. The benefits of this are an overall reduction of power, storage and communication usage.

Training starts with the user annotating two separate sets of data: current context with associated timestamps; and physical activity with associated timestamps. From here the training is split into two parts. The context classifier requires context sensor data and context label, while the activity classifier is trained using activity labels and activity sensor data.

3.7 Client and Server

Figure 2 describes the core context system. This includes description of the client and server architecture enabling real time operation.

As a part of the complete system, the client application guides a user in training mode, and then displays classification results in online mode. An ideal candidate here is a mobile application supported by a smartphone, as we will report in our experimental implementation. This is preferred for two primary reasons: First, mobile devices are pervasive, which makes the client accessible, and we can leverage services off existing network infrastructure that is available in the residential, workplace, and clinic environments, where the systems reported here are deployed; Second, mobile devices are high performance, so they are able to act not only as a user interface platform, but also as a wireless sensor hub that can log, process and store data from the wearable sensors.

Complementary to the client, a web service is provided in this system. The data payload expected is compact, as the complexity is in processing of data (feature extraction and classification). This also means that the web service interface is straightforward, as it is only required to send and receive data. These characteristics require only a straightforward web service architecture.

4. Experimental Implementation – Personalized activity monitoring

In this section, we present an initial implementation of the proposed system, and demonstrate many of the anticipated benefits obtained.

4.1 Data Acquisition and Processing

4.1.1 Data Acquisition

The data recorded in this implementation are: 6 triaxial accelerometer data sets along with activity labels (the sensors are located on both wrists and ankles, and both sides of the waist); wireless access point MAC addresses and signal strengths; audio data time series; and finally context labels.

Data acquisition not only involves collecting sensor data and corresponding annotations, but also includes post-processing analysis, where all annotations must be matched with corresponding data. Only then are they ready to be used by classifiers. While many studies in the area of activity classification provide detailed discussions on classifiers and features, they do not address the variety of issues related to data acquisition required for essential system training. Studies have shown a number of factors affecting data acquisition accuracy, ranging from end users being severely inconvenienced by the equipment they have to carry, to users not being able to record properly or meet the annotation demands using traditional pen and paper approaches [18-20]. Many solutions also do not scale well with increased subject numbers, and both lost and corrupted datasets are common due to human error in labeling or inaccurate recording. Another problem we observed in large measurement campaigns is that many time references for events are recorded based on different clock resources, depending on where the subject was at the time. These clock resources are mis-synchronized by more than several minutes with each other, and compared to the sensor system time-base. This phenomenon dramatically reduces the effectiveness of the labeling process.

4.1.2 Context and Event Data Acquisition (CEDA) System

We have designed and implemented a complete data acquisition and processing system that includes an Android based client and a centralized server-hosted labeling tool. Our system is developed against Android SDK version 2.2 and the target device may be any Android smartphone platform (we used an Archos 32 Internet Tablet, which has support for wireless and audio recording). The application follows a very simple flowchart with structured transitions. This means that all possible user inputs are predefined, thus lowering potential annotation errors.

The Context and Event Data Acquisition (CEDA) System is displayed in Figure 4. During training, a user can indicate the start of a context or activity by pressing "Start Context" or "Start State" respectively (Figure 4a). Once pressed, the user is then prompted by a selection list to choose a context or activity (Figure 4b and 4c).

4.1.3 CEDA Data Reader

To interpret the context information gathered from our Android application, the data reader need to parse the data into a format described by relevant interfaces, namely *IContextData* and *IAnnotation*. This module can be deployed on the same mobile device that runs the acquisition system for local processing, or it can be deployed across network. For example, in a server client scenario where the data logger simply transmits all data to a server for parsing, we would have the CEDA data reader running on a server, listening to traffic coming in from the network.

The *IContextData* type describes the context information collected, and consists of an array of timestamps, sound objects, and mapping of wireless MAC addresses and signal strengths. This interface is intended for use with a feature extractor implementing the *IContextFeatureExtractor* interface. The *IAnnotation* type simply describes the timestamp and label of recorded physical activities. This is used for establishing valid segments of accelerometer data to be used in the WHSFT module.

The CEDA data reader also serves as a filter to sanitize input. Some sensor data are corrupted due to memory write errors or low battery, and processing is done here to correct those (interpolation from previous values for our implementation).

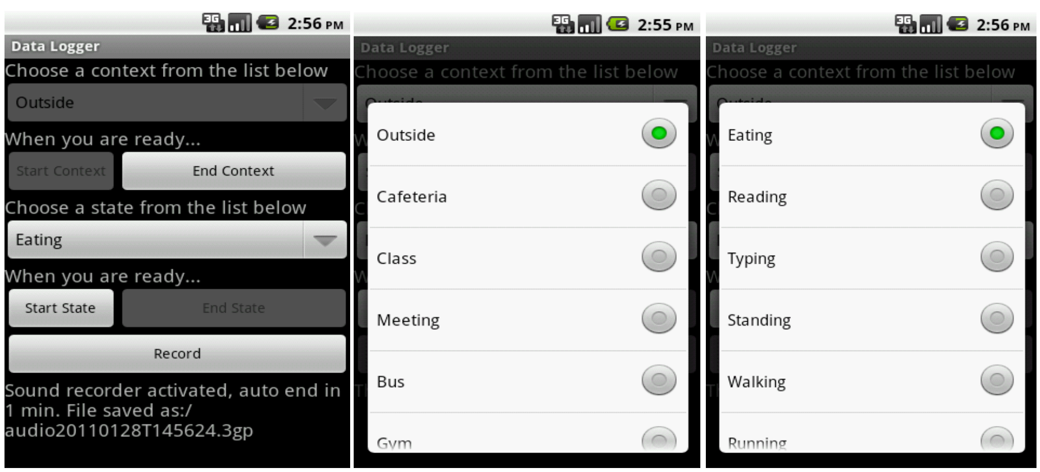

a. Main screen **b. Context selection** **c. Activity selection**

Figure 4 Data logger screenshots

4.1.4 Data Labeling

Once the data are processed, they must be collated and labeled according to user annotations. Our implementation includes a system that when used in conjunction with the data logger, allows us to move from a labeling process that is human intensive to one that only requires humans for verification. Because the Android application already records the start and end time of activities (annotations.txt), we can overlay timestamps on top of sensor data. Labeling from there is straightforward, and Figure 5 shows the results (activities are identified between black and red lines). Now human effort is only needed for quality checking. This drastically reduces the amount of time and effort required for organizing collected data.

Using this CEDA system, we have a robust means for supporting large campaigns, where users will be given a kit containing an Android (or other smartphone) application and sensors. Once a user receives the kit, the Android system can guide the user through training, and then provide individualized feedback. Notifications can also be provided automatically by the same system. This is a huge improvement over many current clinical trials, where notifications come in the form of phone calls or text messages [21]. For activities where the subjects cannot label data themselves (e.g. while in sports), a third party can use CEDA smartphone device to provide the labeling, and the same benefits over traditional approaches apply.

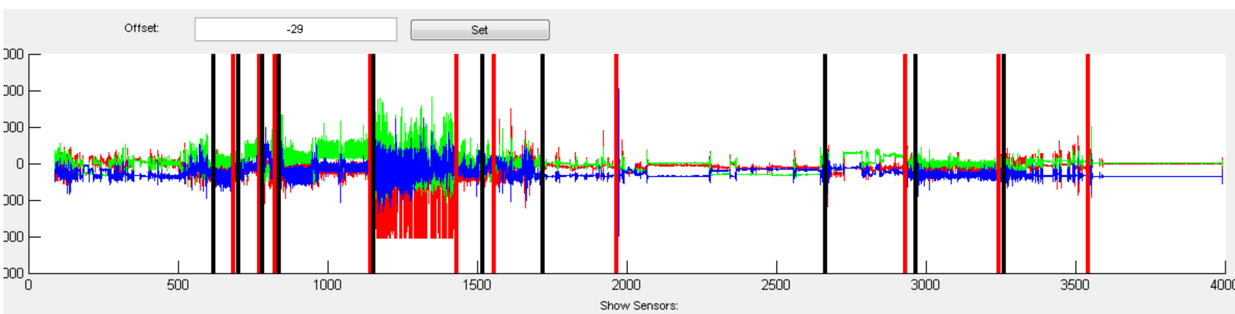

Figure 5 Regions of annotated activities

4.2 Context Detection

Based on the discussion of features and the choice of classifiers, there is a need to use separate classifiers to determine context based on different features. In this implementation, the committee is made up of 3 classifiers: kNN (k-nearest neighbors) with time as a feature; kNN with wireless MAC address and signal strength as features; and AdaBoost with audio peak frequency, peak energy, average power and total energy as features. These features are extracted from raw sensor data through a java program implementing the *IContextFeatureExtractor* interface.

The k-nearest-neighbors classifier is an instance based learner [22]. It is a lazy learner in that no real work is done when the training sequence is given during the training phase; they are simply stored by the classifier. When an unknown class is encountered, the classifier looks for the k nearest training samples to the unknown class, and a decision is made based on majority vote. Other than implementation simplicity, another major advantage of kNN is the ability to handle nominal data through custom designed distance functions. This is particularly important for data types including, for example wireless MAC address values. For our implementation, the kNN with time feature uses a simple absolute distance function, and the kNN with wireless features uses a custom distance function that looks for the closest k labels with overlapping MAC address sets, ranked by signal strength.

AdaBoost is a type of boosting method [23]. It is a meta learner, meaning that it is to be used in conjunction with a base learner. The base learner can be any classifier, and is usually straightforward to implement. They can also be very weak. In the binary case, a base learner needs only to outperform chance (50%). By forming multiple weak classifiers and weighting them on their accuracy, AdaBoost can combine the ensemble into a strong classifier. There are many papers describing the operation and algorithms of a number of AdaBoost variants [23,24]. For our implementation, we used the AdaBoost.MH algorithm with a decision stump base learner. AdaBoost.MH is an earlier variant, and one of the most popular. It is an extension of the earliest multiclass AdaBoost.M1 algorithm. A listing of the AdaBoost.MH algorithm can be found in [24].

4.3 User Activity Classification Using WHSFT

The Wireless Health Institute at UCLA has developed accurate classification methods for user activities under diverse situations and clinical settings [1,2,26]. In the process, we have developed a sensor fusion and classification toolkit: Wireless Health Institute Sensor Fusion Toolkit (WHSFT). It is a toolkit that provides a multimodal hierarchical classification system based on the Naive Bayes classifier.

Starting with raw data from multiple sensors, the WHSFT finds overlapping times for all the sensors, and combines streams of data into a single structure. Features such as short time energy, mean, and variance on three axes of the accelerometer are computed from the combined data structure. There are a vast number of diverse features providing freedom in selecting features that best suit each application.

From the features, we can then build a hierarchical structure that models the classification problem. This is a tree like structure that describes the activities we are trying to classify. At each level of the tree WHSFT uses a Naive Bayes classifier that separates unknown data into one of the branches. Once we reach a leaf node, a final classification is made. The personalization benefits of the proposed architecture are realized through these customizable hierarchical models. For example Figure 6 shows the model for "Cafeteria" context described in Table 2.

4.4 Web Services

The web service serves as a gateway between the client and server for real time classification. It also implements the whole automation process. While we only have the real time classification system implemented, the training system can easily be extended into the web service.

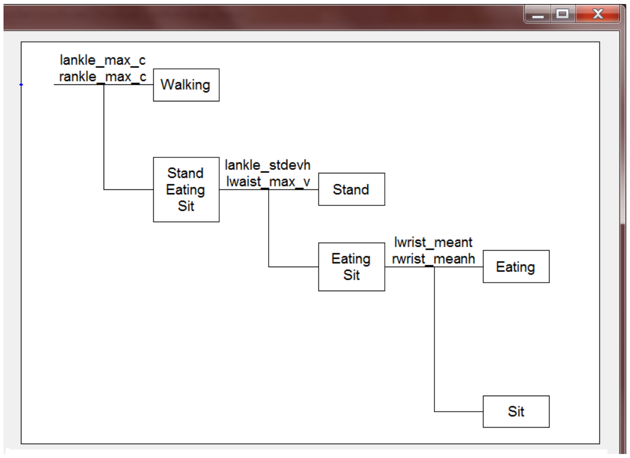

Figure 6 Context guided model for cafeteria

When the system is online, raw data is first collected using the CEDA data logger on the Android tablet. They are then passed to the server by invoking a web service. After receiving this data, the server will start the core classification system. Finally, both context and activity classification results are combined along with sensor policies, and returned to the Android tablet.

Following a RESTful web service architecture, we take advantage of existing HTTP infrastructure [25]. This approach requires a number of key components: a web server for providing the HTTP infrastructure; a platform for developing the actual web service; and a naming authority for redirecting requests to the web service. For the server, an Apache web server stack is deployed. This allowed us to develop the web services in PHP, and the naming authority in htaccess.

In this implementation, the system is able to respond in real time with both a context and an activity classification.

5. SYSTEM EVALUATION

5.1 Data Acquisition

Table 3 lists all models built for this experimental trial. Figure 6 and 7 shows two examples (Cafeteria and Bus) of the models running in our activity classification system (WHSFT). It is noted that due to our lack of wireless accelerometers, training and testing were conducted offline.

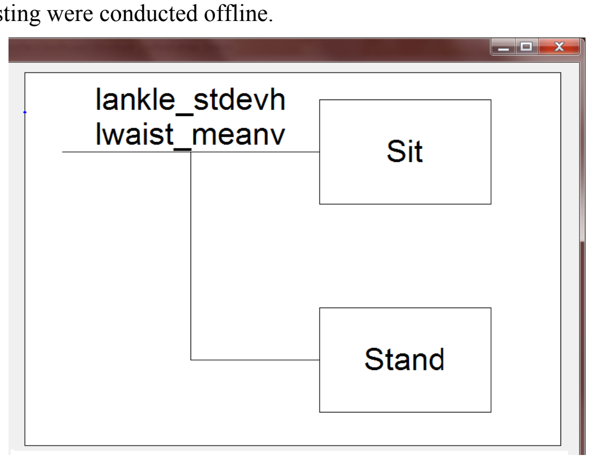

Figure 7 Context guided model for bus

Data acquisition was performed as follows: three subjects carried an Archos Internet Tablet and six Medical Daily Activity Wireless Network (MDAWN) devices. The tablet supports our Android client, and the MDAWN devices are placed on wrists, waist and ankles. The MDAWNs robustly provide trialaxial accelerometer data [26]. Each subject was asked to record two sets of data. The first set is for training the classifiers, where we performed each activity for at least 5 minutes. Each subject then spent 30 minutes in every context, and collected context data every 2 minutes. The second set is for training purposes, and each subject spent over eight hours across the contexts, collecting all data listed. Physical activity recordings are marked by signatures indicated by shaking of the right wrist every hour. This is used to confirm the functionality of the CEDA automatic labeling system.

All of the MDAWNs and Archos tablets are synchronized to the correct time by running a time synchronization tool for the MDAWNs, and setting the tablet's clock to match that of the PC that ran the synchronization tool.

5.2 Results

5.2.1 Context Classifiers

Our new context guided classification method includes the classifier committee system, and experimental results directly demonstrate its effectiveness, in the presence of complex classification challenges.

Table 4 summarizes the accuracies of the classifier committee and the individual classifiers in the committee.

We see that wireless kNN performs with insufficient accuracy for bus and outdoors. In the bus context case, the sensor system detects a large number of wireless access points that have not been incorporated into prior training due to the trajectory of the bus. In the outdoor context case, the system tends to detect access points that belong to one of the contexts at nearby indoor locations. For example, walking near a building causes the context to be classified as that of a context inside the building. Time KNN is also not sufficiently accurate for a number of contexts, and this is due to the varied nature when subjects visit these contexts. For example, subjects visited the cafe and outside at different times of the day. AdaBoost using sound features performed well for all contexts, but there are instances where a bus driving nearby causes a misclassification.

Table 4 Context classifiers accuracies

	AdaBoost	Time kNN	Wireless kNN	Committee
Bus	80	59	29	80
Cafeteria	100	35	80	90
Class	80	87	89	95
Meeting	100	73	100	100
Outdoors	85	33	25	85
Home	100	100	100	100

Clearly, this experimental evaluation provides a classification challenge for each individual classifier. However, as shown in Table 4, by combining the best of all classifiers, our committee is able to achieve high accuracy for all contexts.

5.2.2 Context guided Classification Accuracy

A critical benefit of context guided classification is a direct improvement in accuracy for each classifier. This is also demonstrated directly by experimental results here.

Table 3 Context guided models

	Walking	Running	Walking Upstairs	Walking Downstairs	Sitting	Standing	Writing	Eating
Outdoors	X	X	X	X				
Cafeteria	X				X	X		X
Home	X		X	X	X	X		
Class	X				X		X	
Meeting	X				X	X	X	
Bus					X	X		

Table 5 Context guided classification accuracy

Context	Generic	Specific	Improve	Context	Generic	Specific	Improve
Cafeteria				Meeting			
Standing	96.91	98.97	2	Sitting	91.67	100	9
Walking	84.81	100	17	Walking	97.83	100	2
Eating	1	96.77	95.77x	Writing	2.5	69.62	28.84x
Sitting	100	100	0	Standing	96.84	100	3
Outdoors				Home			
Walking	99.29	99.29	0	Sitting	100	100	0
Running	95.79	95.79	0	Standing	94.12	100	6
Upstairs	90.47	90.47	0	Walking	98.47	96.95	-1
Downstairs	97.30	97.30	0	Upstairs	100	100	0
				Downstairs	96.61	96.61	0
Bus				Class			
Sitting	95.94	100	4	Walking	98.56	100	1
Standing	81.29	86.33	6	Sitting	87.33	71.04	-20
				Writing	3.66	79.41	20.7x

The results are broken down by context, as shown in Table 5. The "Generic" column shows results from a standard classification tree using WHSFT, with all activities built in. The "Specific" column shows accuracy from context guided classifiers. All values are in percentage unless stated otherwise.

In nearly all examples, there is a substantial increase in classification accuracy resulting from the introduction of context guided classification, as a targeted model with fewer activities is presented (as opposed to the conventional approach where the classifier is presented with all possible activities for selection). In the case of writing and eating, a large increase in accuracy can be seen, from very limited accuracy for this complex upper body activity to acceptable accuracy. Under the class context, there is a decrease in accuracy for sitting, however, it is observed that classification of writing is available with almost 80% accuracy. The selection of new features and structures can further enhance accuracy through convenient and straightforward development.

5.2.3 Context guided Classification Speed

Context guided classification offers a direct advance in computational throughput that offers the possibility of real time classification.

Table 6 shows the computational speed advance that has been achieved. In all cases there is a significant increase in classification speed. This indicates that context guided classification can enable an online system capable of computing subject state with the accuracy reported here in real time. Again, "Generic" column shows the amount of time (in seconds) required to perform a classification with the full model, and "Specific" shows the amount of time (in seconds) required for the context guided system. "Improve" shows the improvement factor (in number of times).

Table 6 Speed increase of WHSFT-ca

	Generic	Specific	Improve
Bus	0.119	0.013	9.2x
Café	0.120	0.044	2.7x
Class	0.122	0.039	3.1x
Meeting	0.127	0.058	2.2x
Outside	0.128	0.033	3.9x
Home	0.119	0.050	2.4x

5.2.4 Context guided Classification Energy Usage

Context guided classifier now also offers the capability for selecting optimal sensors and schedules for energy and operating lifetime benefits. This also permits a minimum number of sensor systems to be selected (for user convenience) while maintaining classification accuracy.

Based on models constructed, we produced the sensor requirement chart in Table 7. Blank cells indicate that a sensor can be safely turned off without affecting the accuracy for a given context. For example, in the case of "Bus" (Figure 8), only the left waist sensor is required.

Table 7 Sensor requirement

	Left Ankle	Right Ankle	Left Waist	Right Waist	Left Wrist	Right Wrist
Bus			X			
Cafeteria	X	X	X		X	X
Class	X			X		X
Meeting		X	X			X
Outdoors	X	X		X	X	
Home	X	X	X			

Using this chart, our sensor policy selector (*ISensorPolicyMaker*) can determine which sensors can be shut down. To estimate the potential for energy reduction, our analyses are directed to determining the improvement in operation time by adopting sensor activation and sampling schedules, as determined by context. The analysis was performed based on recorded data, after the completion of data collection as our sensor systems do not allow real time scheduling. Results indicate the potential benefits of context guided sensor energy management and now may be applied to sensor systems that can be scheduled with real time coordination.

To indicate the operating time improvement over a range of subject behaviors, two cases were taken as examples, a graduate student subject and a subject remaining in a residential household. The typical profiles of their daily life are shown in Figure 8a, and the total operating time using continuous sensor system usage, in comparison to context guided sensor usage, is shown in Figure 8b.

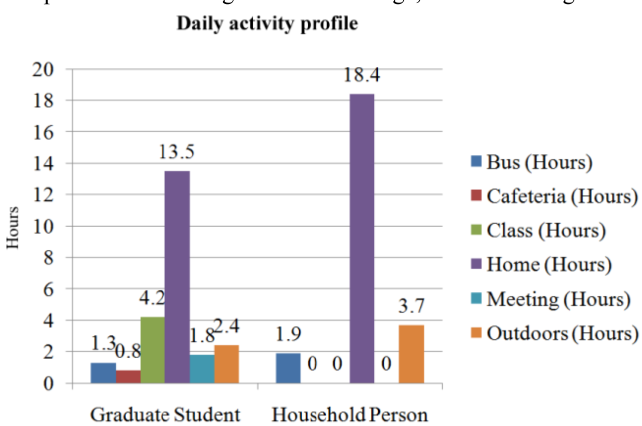

es

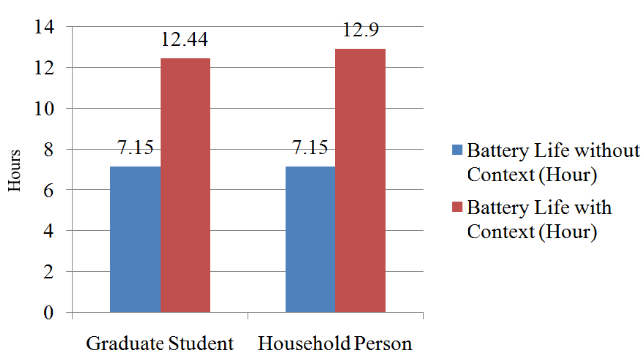

Figure 8 User profiles and their battery life comparison

6. CONCLUSION

Activity monitoring appears as a critical need and valuable source of disease intervention and guidance in healthcare, personal health and wellness promotion, workplace safety, and athletics. In this paper, we have described the design, implementation, and comprehensive evaluation of a novel end-to-end system that introduces context into activity classification.

On the architecture level, we first presented a refined definition of context, and a unique classification committee approach is described for detecting context of diverse forms. We then described how any current classification system can take advantage of the new context information through the concept of a context guided classifier. Finally, we described the potential for real time classification through the use of web services that exploit smartphone technology. The architecture also employs an interface model, providing great flexibility in the rapid implementation and integration of subsystems.

We also presented a realization of the above context guided classification system, where an Android client application addressed issues relating to robust data acquisition and large campaign support. For the core system, AdaBoost, kNN and hierarchical naive Bayes classifiers were all used for context detection and activity classification, demonstrating the inherent system flexibility. This has also demonstrated the important capability our system provides in enabling a matching of classifier systems to applications and the capability for the classification committee to properly combine these for optimization of classification accuracy.

Finally, through a series of experimental field evaluations sampling each of the diverse context examples and activities in multiple episodes by multiple subjects, the critical benefits of this system were demonstrated. First, it was demonstrated that context guided classification has enabled a substantial advance in classification accuracy for many activities including upper body motion. Second, it has been demonstrated that context guided classification offers a computational throughput advance that may be exploited for benefits including the support of real time, high accuracy classification. Finally, it was also demonstrated that the context classification capability can be applied to control the activation and selection of sensors. This benefit will be exploited in the immediate future to enable substantial operating lifetime extension for critical applications.

Future work in context guided classification will be directed to multiple areas. This includes the classification of complex motion for subjects afflicted with disease conditions causing motion disorders. This will then directly enable new rehabilitation methods that are now possible as a result of in field, accurate classification. The capabilities of the system will also enable a new series of investigations directed to the detailed characterization of classifier systems and their selection for specific applications. New motion sensor systems now under development will be included as well in the future. Finally, it is planned that context guided classification will be provided as a tool for the research community in collaborative development of wireless health applications.

7. ACKNOWLEDGEMENT

This material is based upon work supported by the National Science Foundation under Grant Nos. 0120778 (The Center for Embedded Networked Systems). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

The authors would like to thank Eric Wang, Xinda Lin, Lawrence Au, Celia Xu and Maxim Batalin, all from UCLA, for their discussions on the work presented in this paper.

8. REFERENCES

[1] B. H. Dobkin, X. Xu., M. Batalin, S. Thomas, and W. J. Kaiser, Accelerometry algorithms for activity recognition. *Stroke*, 2011 (in press).

[2] W.H. Wu, M.A. Batalin, L.K. Au, A.A.T. Bui, and W.J. Kaiser. Context-aware sensing of physiological signals. In

Annual International Conference of the IEEE Engineering in Medicine and Biology Society. pages 5271-5275, 2007.

[3] M. F. Gordon. Physical Activity and Exercise Recommendations for Stroke Survivors. *Stroke,* pages 1230-1240, 2004.

[4] B. Logan and J. Healey. Sensors to Detect the Activities of Daily Living. In *Engineering in Medicine and Biology,* pages 5362-5365, 2006.

[5] M. Philipose, K.P. Fishkin, M. Perkowitz, D.J. Patterson, D. Fox, H. Kautz, and D. Hahnel. Inferring activities from interactions with objects. *Pervasive Computing, IEEE,* vol. 3, pages 50-57, 2004.

[6] L. Han, S. Jyri, J. Ma, and K. Yu. Research on Context-Aware Mobile Computing. In *22nd International Conference on Advanced Information Networking and Applications Workshops 2008,* pages 24-30, 2008.

[7] A. Schmidt, M. Beigl, and H. Hans-W. There is more to context than location. *Computers and Graphics,* vol. 23, pages 893-901, 1999.

[8] D. Siewiorek, A. Smailagic, J. Furukawa, A. Krause, N. Moraveji, K. Reiger, and J. Shaffer. SenSay: a context-aware mobile phone, In *7th IEEE International Symposium on Wearable Computers, IEEE*, 2005.

[9] C. Doukas and I. Maglogiannis. Enabling human status awareness in assistive environments based on advanced sound and motion data classification. In *Proceedings of the 1st ACM international conference on Pervasive Technologies Related to Assistive Environments PETRA 08*, pages 1, 2008.

[10] D. A. James, N. Davey, and T. Rice. An accelerometer based sensor platform for insitu elite athlete performance analysis. In *Proceedings of IEEE Sensors 2004, IEEE,* vol. 3, pages 1373-1376, 2004.

[11] S. T. Scott, and L. Jonathan. iLearn on iPhone: Real-time human activity classification on commodity mobile phones. http://www.cs.washington.edu/homes/jlester/publications/UW-CSE-08-04-02.pdf, 2008.

[12] K. Van Laerhoven, A. Schmidt, and H. W. Gellersen. Multi-sensor context aware clothing. In *Proceedings Sixth International Symposium on Wearable Computers,* pages 49-56, 2002.

[13] A. K. Dey. Understanding and Using Context. *Personal and Ubiquitous Computing,* vol. 5, pages 4-7, 2001.

[14] T. Hofer, W. Schwinger, M. Pichler, G. Leonhartsberger and J. Altmann. Context-awareness on mobile devices: the hydrogen approach. In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences,* pages 292-302, 2002.

[15] P. Prekop, and M. Burnett. Activities, context and ubiquitous computing. *Special Issue on Ubiquitous Computing Computer Communications,* vol.26, pages 1168-1176, 2003.

[16] E. Gamma, R. Helm, R. Johnson, and J. M. Vlissides. Design Patterns: Elements of Reusable Object-Oriented Software, *Addison-Wesley Professional,* 1994.

[17] Y. C. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm. Accuracy characterization for metropolitan-scale Wi-Fi localization. *ACM,* pages 233-245, 2005.

[18] S. Pirttikangas, K. Fujinami, and S. Hosio. Experiences on Data Collection Tools for Wearable and Ubiquitous Computing, In *2008 International Symposium on Applications and the Internet, IEEE,* 2008.

[19] S. Patnaik, E. Brunskill, and W. Thies. Evaluating the accuracy of data collection on mobile phones: A study of forms, SMS, and voice. In *2009 International Conference on Information and Communication Technologies and Development ICTD,* pages 74-84, 2009.

[20] S. J. Lane, N. M. Heddle, E. Arnold, and I. Walker. A review of randomized controlled trials comparing the effectiveness of hand held computers with paper methods for data collection. *BMC medical informatics and decision making,* vol. 6, pages 23, 2006.

[21] C. V. Campen, H. Sixma, J. J. Kerssens, and L. Peters. Comparisons of the costs and quality of patient data collection by mail versus telephone versus in-person interviews. *European Journal of Public Health,* vol. 8, pages 66-70, 2005.

[22] I. H. Witten and E. Frank. Data Mining: Practical Machine Learning Tools and Techniques, Second Edition (Morgan Kaufmann Series in Data Management Systems), *Morgan Kaufmann,* 2005.

[23] R. E. Schapire, and P. Avenue. The Boosting Approach to Machine Learning An Overview. In *MSRI Workshop on Nonlinear Estimation and Classification,* pages 1-23, 2003.

[24] R. E. Schapire and Y. Singer. Improved Boosting Algorithms Using Confidence-rated Predictions. *Machine Learning,* vol. 37, pages 297-336, 1999.

[25] J. Y. Xu. Web-Based Billing System Exploits Mature and Emerging Technology. *IEEE IT Professional,* vol. 13, pages 49-55, 2011.

[26] X. Xu, M. A. Batalin, W.J. Kaiser, and B. Dobkin. Robust Hierarchical System for Classification of Complex Human Mobility Characteristics in the Presence of Neurological Disorders. In *Body Sensor Network,* 2011.