A FRAMEWORK TO DETECT AND CLASSIFY ACTIVITY TRANSITIONS IN LOW-POWER APPLICATIONS

Jeffrey Boyd

Arizona State University
Jeffrey.Boyd@asu.edu

Hari Sundaram

Arizona State University
Hari.Sundaram@asu.edu

ABSTRACT

Minimizing the number of computations a low-power device makes is important to achieve long battery life. In this paper we present a framework for a low-power device to minimize the number of calculations needed to detect and classify simple activities of daily living such as sitting, standing, walking, reaching, and eating. This technique uses wavelet analysis as part of the feature set extracted from accelerometer data. A log-likelihood ratio test and Hidden Markov Models (HMM) are used to detect transitions and classify different activities. A trade-off is made between power and accuracy.

Index Terms— Gesture Recognition, wavelet analysis, low power, inertial sensors, HMM

1. INTRODUCTION

This paper presents a framework for detecting and classifying activities in low-power devices used in the monitoring of stroke rehabilitation.

Survivors of stroke often have impaired mobility on one side of their body. They often develop ways to compensate for the affected limb by using their unaffected limb, shoulder, and torso. These compensatory habits only further weaken the affected limb because they simply don't use it. Usually rigorous physical therapy helps them relearn the use of their limb. An important measure of their functional recovery is the amount they use their impaired arm in activities of daily living. Monitoring such activities through a continuous observation either in person or via camera is expensive and raises serious privacy concerns. The issue can be addressed alternatively by a low-power monitoring device with inertial sensors such as accelerometers and gyroscopes. However, such a device needs to meet several important requirements: it needs to be unobtrusive and allows the patient to go about their daily activities unhindered, and it also needs to be small, lightweight and work continuously all day.

There is a considerable amount of literature devoted to tracking motion and detecting gestures, especially as it relates to rehabilitation. A good survey of this literature is [1]. The authors of [2] studied the effects of power vs. accuracy in a sound-based activity recognition system. They used a hand-worn microphone and accelerometers to recognize when a user uses specific kitchen appliances. Their work incorporates power-awareness in the design process from the beginning and specifically explored how they could maximize power while only slightly reducing accuracy by using data features that use less power.

Wavelet analysis has been used to analyze electroencephalogram (EEG) signals to detect seizures in [3]. A time-varying multivariate probability distribution function was used in [4] to detect seizures from EEG data as well. In this paper, we employ similar analysis, though applied to accelerometers and human activity data.

We propose a framework for power-efficient automatic segmentation and classification of simple human activities using inertial sensors. The idea behind our framework is that a person's activities can be segmented by first detecting the transitions between activities and then focusing most of the computations needed to classify the activities along these boundaries. The device would conserve energy by sampling at a high rate when a transition is detected then switch to a low sampling rate once a new activity has been classified.

In our framework, transitions are detected by decomposing accelerometer signals using Daubechies wavelets and measuring the likelihood that one block of the signal is different from an adjacent block. The coefficients of the wavelet decomposition are the inputs to this log-likelihood ratio test. These transitions can be detected at both high (100Hz) and low (10Hz) sampling frequencies. Once a transition is detected, a HMM or other classifier can be used to identify the new activity.

In our experiments we were able to detect transitions between simple activities such as sitting, standing, eating, reaching and walking at both 10Hz and 100Hz. A HMM classifier correctly identified 17 of the 18 test cases. We were also able to show that transitions could, on average, be detected within one second of the actual transition.

2. LOW-POWER FRAMEWORK

In order to achieve power-efficient activity segmentation and classification, our framework is modeled with two states: sleep and active, generally defined as low-power and high-power states, as depicted in Figure 1. This definition is intentionally general so implementations of the sleep and active states can device independent. For this paper, and in our experiments, we define our low-power state as a low sampling frequency and the active state as a high sampling frequency.

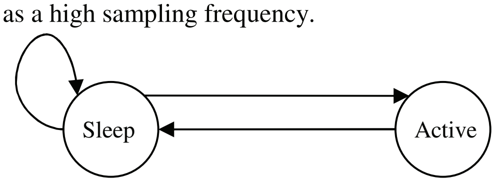

Figure 1. Sleep-Active Cycle

The concept still holds in more traditional implementations of sleep state in which a device may turn off components and enter a low-voltage or low-power state. As an example, consider the power consumption of the Intel Mobile Sensing Platform (MSP) [5]. According to [6], the device draws .390 mA in deep sleep mode and [7] claims as much as 137 mA in full active mode. If the device operates just 25% of the time in deep sleep mode rather than fully active, it will run approximately 33% longer. With a 1000 mAh battery, that translates to 7.3 hours vs. 9.7, approximately a 33% increase.

Our framework can be extended by defining multiple states in-between Sleep and Active where each state has a higher sampling frequency as it moves from Sleep to Active. Transitions into one of these intermediate states would be based on the log-likelihood ratio test. As the ratio increases the device enters a state with a higher sampling frequency.

3. TRANSITION DETECTION

This section discusses our transition detection method that enables the two-state activity monitoring process in our low-power framework. Transitions are the boundaries between activities where intense computations are needed to correctly identify the new activity and mark when a device can switch to a lower-power state. If we can correctly and efficiently detect transitions, then the device will spend most of it's time in a low-power state and thus conserve energy.

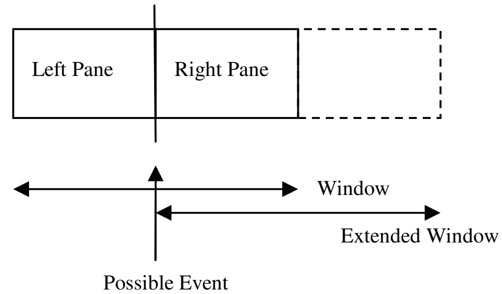

Figure 2. Window View

In this paper we explore the transitions between simple activities of daily living: sitting, standing, walking, reaching and eating. These activities were chosen because they are simple, repeatable, and their accelerometer signal signatures vary enough from each other that they can be easily distinguished.

One way to look at this transition detection scheme is to define time windows over parts of the signal as shown in Figure 2. We define several successive samples as a frame. A set of features is then calculated for each frame. A possible transition is detected by defining a window that stretches several frames before and after a sample. For each frame, a two-level, second-order Daubechies wavelet analysis was performed. Wavelet analysis was chosen for its ability to compactly represent both approximation and detail properties of a signal and because of their resemblance to the signal representation of the activity transitions of interest. The wavelet coefficients become the feature vector used in calculating the likelihood function for each window and for classification using HMMs.

Assuming our samples are independent, identically distributed Gaussian variables, a likelihood function can be defined for the left and right window "panes," and the whole window, as follows:

$$L(x_1, x_2, \ldots x_N) = \sum_{j=1}^{N} \ln\left(p(x_j \mid \mu, \Lambda)\right) \qquad \textbf{(1)}$$

where $x_1, \ldots, x_N$ denote the N observations from the left, right or the whole window, $\boldsymbol{\mu}$ and $\boldsymbol{\Lambda}$ are the mean feature vector and covariance matrix in the Gaussian model. The probability p is the multivariate Gaussian probability distribution function. For example, if we define a window as eight seconds and sample at 100 Hz, then each "pane" would be four seconds (or 40 frames) long and L would be calculated for the left and right panes and the entire window. Equation 2, a likelihood ratio test, can now be used to detect a transition.

$$\frac{L_{whole}}{L_{left} + L_{right}} \qquad \textbf{(2)}$$

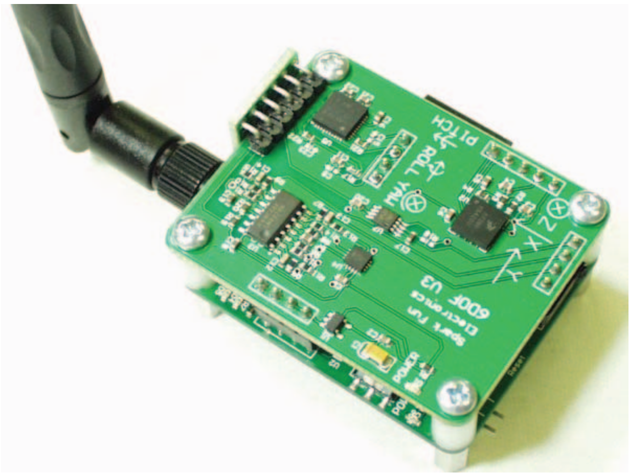

Figure 3. SparkFun Electronics IMU 6DOF v3. Image courtesy of SparkFun Electronics.

where L_{whole} , L_{left} and L_{right} use Equation 1 for the left and right window "panes," and the whole window, respectively. This ratio will be close to 1 for cases where no transition occurs and increase when there is a transition. Local maxima indicate the location of a transition. After the peak of a local maxima is detected, then an extended window is defined and a HMM is used to classify the new activity.

4. EXPERIMENTS

This section discusses the results of our experiments, which addressed three main questions: (a) Can our transition detection method properly segment the observed activities? (b) How well can we classify activities based on our transition detection method? (c) How do window sizes and sampling rates affect how close we can detect transitions?

To test the viability of the windowing method to detect activity transitions, we collected data using the SparkFun IMU v3 device, shown in Figure 3. This device features triple-axis gyroscopes and accelerometers, and a dual-axis magnetometer, though the magnetometers were not used in this experiment.

We collected data from simple activities: sitting, standing, walking, reaching for an object on a table, and eating, while the SparkFun IMU was attached to the author's right wrist. The sensors were sampled at 100 Hz. Four combinations of these activities were collected: sit-stand-walk-sit, sit-reach-walk, sit-eat-walk, and sit-walk-stand. The analysis in the next few paragraphs was all done using the 3-axis accelerometers exclusively. The magnetometer data was too noisy and including the gyroscope data resulted, in most cases, in a singular or near-singular covariance matrix. We were thus unable to compute the inverse of the covariance matrix needed in the multivariate Gaussian pdf. Also, only the approximation coefficients of the wavelet were used because using the detail coefficients would also result in a singular or near-singular covariance matrix.

Figure 4 shows representative sample of the raw values of the accelerometers and the results of using Equation 2 to detect activity transitions in the sit-reach-walk data. In the top figure one can see where the transitions take place in the raw data and where our framework detects them in the bottom figure. A windows size of eight seconds (four seconds before and after a possible transition) was used. The signal was tested for transitions every frame (10 samples). There are two transitions in this data set, indicated by the two peaks in the bottom graph.

Table 1. Classification Confusion Matrix

	sit	stand	reach	walk	eat
sit	5	0	0	0	0
stand	0	2	0	1	0
reach	0	0	2	0	0
walk	0	0	0	6	0
eat	0	0	0	0	2

Next we defined an extended window of 2 seconds after the transition. We used these two seconds as input to HMM models for each of the activities that had been trained using data from all the four transition scenarios previously described. These HMMs used five hidden states and ten output states. Table 1 shows the resulting confusion matrix.

The HMM correctly classified the activity in 17 of the 18 test cases. This speaks greatly to the usefulness of the window view technique to detect transitions and identify what activity a person has transitioned to.

Next we wanted to get a sense of how far off (in time) our transition detector was in detecting when the

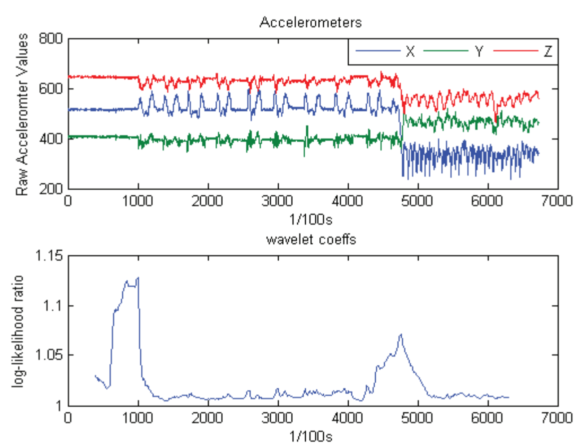

Figure 4. Original Signal & Wavelet Transition Detection. Window Size was 8 seconds.

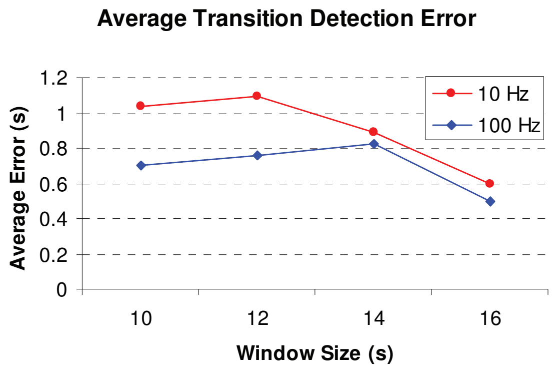

Figure 5. Average Detection Error

transition occurred and do this for both our low-frequency (10Hz) and high-frequency (100Hz) states using different sized windows. For these tests we again defined our frame to be 10 samples, this meant our transition detector should, in theory, be able to detect a transition with .1s when sampling at 100Hz and 1s at 10Hz. In other words, the system needs 10 samples to compute the wavelet coefficients and at 100Hz, this is every .1s, so the system checks for a transition every .1s. At 10Hz, it needs a whole second for 10 samples, so it tests for a transition once every second.

Detection of the transition always lags real-time because of the window used to compare the two "window panes." For example, if we define a window size of 8 seconds, the system would check if there was a transition 4 seconds ago, and only at .1s intervals (if sampling at 100Hz). Figure 5 summarizes the results for window sizes of 10, 12, 14, and 16 seconds. Window sizes of 8 or less seconds resulted in singular or near-singular covariance matrices in the 10Hz data. Error here is defined as the difference from the start of the transition to when our transition detector detected it. We used the peak value of the log-likelihood ratio test to define when the transition detector detected a transition.

The average improves with a larger window. This is most likely because the algorithm has more data points to average over a longer period. This figure also shows that sampling at 10Hz does almost as well as 100Hz. This means the system could detect transitions at 10Hz, and then switch to 100Hz to get enough samples to feed into a HMM. Another point of investigation is to use something other than just the peak value to define when the transition detector finds a transition. Some transitions had peak values very close to other values nearby.

5. CONCLUSION

This is work in progress toward power-efficient automatic segmentation and classification of simple human activities using inertial sensors. A practical outcome of this work is that devices designed with this framework can be smaller, lighter, and last longer than the current generation of technology.

In this paper we have demonstrated the viability of a software framework that can detect transitions in human activities and classify (when properly trained) what those activities are. We have also shown that a sample rate as low as 10Hz is sufficient to detect transitions within one second of when they occur. Future work involves creating a device or programming one to implement the transition model in Figure 1. More research is planned to find an appropriate set of features that can clearly detect activity transitions. More research is needed to find optimal window-lengths for different activities and classification algorithms that are computationally inexpensive. Different types of common activities need to be investigated, especially more complex activities that are made up of many smaller activities.

6. REFERENCES

[1] Huiyu Zhou, Huosheng Hu, "Human motion tracking for rehabilitation--A survey," *Biomedical Signal Processing and Control*, Volume 3, Issue 1, January 2008, Pages 1-18.

[2] Mathias Stager, Paul Lukowicz, Gerhard Troster, "Power and accuracy trade-offs in sound-based context recognition systems," *Pervasive and Mobile Computing*, Volume 3, Issue 3, June 2007, Pages 300-327.

[3] Zarjam, P.; Mesbah, M.; Boashash, B., "Detection of newborn EEG seizure using optimal features based on discrete wavelet transform," IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. Vol. 2, pp. 265-8. 6-10 April 2003.

[4] P.E. McSharry, T. He, L.A. Smith and L. Tarassenko, "Linear and non-linear methods for automatic seizure detection in scalp electro-encephalogram recordings," *Medical & Biological Engineering & Computing*, 40 (2002), pp. 447–461.

[5] Tanzeem Choudhury, Gaetano Borriello, Sunny Consolvo, Dirk Haehnel, Beverly Harrison, Bruce Hemingway, Jeffrey Hightower, Pedja Klasnja, Karl Koscher, Anthony LaMarca, James A. Landay, Louis LeGrand, Jonathan Lester, Ali Rahimi, Adam Rea, and Danny Wyatt. "The Mobile Sensing Platform: An Embedded System for Activity Recognition," *IEEE Pervasive Magazine - Special Issue on Activity-Based Computing*, vol. 7, no. 2, pp. 32-41, April-June 2008.

[6] iMote2 Data Sheet. Online. Available: http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/Imote2_Datasheet.pdf. Retrieved 18 December 2008.

[7] Louis LeGrand. Mobile Sensing Platform Email Archive. Online. Available: http://sourceforge.net/mailarchive/forum.php?thread_name=43563e40802260755p78ce4646qdb1d3bbb278260ea%40mail.gmail.com&forum_name=mob-sen-plat-users. Retrieved 18 December 2008.