Inferring transportation modes from smartphone sensors

Rob Bisseling (Utrecht University), Jason Frank (Utrecht University), Haris Gavranovic (University of Sarajevo), Jasper van Heugten (Radboud University Nijmegen), Anna Kruseman (Utrecht University), Daphne van Leeuwen (CWI), Christian Reinhardt (University of Leiden)

1 Introduction

Transportation in urban areas poses big challenges related to sustainability, safety and health of residents. A key step to improving policymaking in these respects is to collect and analyse data on how current resources are used, so that inefficiencies may be identified and addressed. The abundance of mobile devices makes it very attractive to harness the advanced data collection abilities of smartphones to tackle this question.

1.1 Mobidot problem statement

Mobidot b.v. develops software that can be implemented in smartphone applications to provide automated capturing and analysis of mobility traces of individuals via smartphone sensors. Their customers are businesses and government organisations interested in quantifying and improving the travel patterns of their employees or constituents, especially by providing those individuals with knowledge that motivates them to make safe, sustainable and healthy mobility choices.

The platform developed by Mobidot utilizes efficient real-time data acquisition using smartphone sensors, coupled with a central analysis server that cleans the raw data and compares it with available databases to infer travel trajectories and modes.

The problem formulated for the Study Group focuses on advancing the quality of data derivation. Mobidot infers the route, role, objective and mode of transportation from smartphone sensor data. Smartphones possess a variety of sensors, including GPS, mobile telephone (4G) and wi-fi signals, accelerometer and gyroscope sensors, etc. that could be used to determine the motion and position of the user, when coupled with geographic databases and known public transport tables.

The volume of monitoring and data recording must be weighed agains battery drain, and consequently a sensing strategy must be devised to optimise information gathering with minimal energy usage. The first objective of the problem posed to the Study Group was to optimise data measurement versus battery usage, by (1) devising and optimal scheduling plan for sensing, detecting mode changes, (2) developing a method to locally filter and compress relevant information at the mobile site, and (3) developing a method to infer motion from incomplete data. A second objective posed to the Study Group was to detect obvious errors in the trip analysis and minimise false inferences.

For the Study Group Mathematics with Industry, Mobidot provided relevant data to develop data deduction improvements and test approaches and methods. This included sample smartphone multi-sensor data, sensor energy usage stats and samples of resulting anonymised mobility profiles.

1.2 Work plan of the Study Group

The Study Group initially began work on several fronts and eventually converged on two most promising lines of investigation. The first of these was an improvement to transportation mode identification using high-resolution accelerometer measurements to try to identify transportation mode signatures from vibrational data (bicycling rhythms, motor frequencies, etc.). Related to this, the second investigation line was a sensing strategy that would accommodate taking such high-resolution accelerometer readings without excessively straining battery charge.

In this paper we propose approaches to dealing with the above challenges. In Section 2, we focus on devising a sensing strategy that is both energy efficient and provides enough data to sense multimodal traveling. This in particular necessitates the detection of changes in the transportation mode. To this end high resolution data is crucial to reliably deduce changes. We propose to supplement the currently used GPS/localization sensing with high-resolution accelerometer data in a three stage sampling procedure. This keeps both energy consumption low and prediction power high.

Once data is acquired it is necessary to infer motion signatures from it. This means that we need to process the time series data obtained from the measurements in a way that enables us to reliably distinguish between different modes of transportation. Current approaches use for example frequency analysis via Fourier transform methods. Here we investigate the use of Wavelet transform methods that provide local frequency information on time series signals. Using wavelet analysis, each time signal is efficiently converted into a distinct two dimensional signature, which in turn could be used to train a learning algorithm to distinguish different transport modalities. In Section 3 we discuss the use of Haar wavelets in some detail and stress how they are useful to detect characteristic changes in time series data. In Section 4 we show some sample applications of this method to accelerometer data taken from different modes of transportation. A powerful feature of the Wavelet approach is the great variety of available basis wavelets that enable one to look for changes with specific structures. We use as a second example the Mexican hat wavelet to analyse the same accelerometer data and discover characteristic structures for the different modes of transportation. The results suggest that further research might lead to powerful prediction tools via the 'right' choice of wavelet. Finally, in Section 5, we mention a method to determine the specific moment of modality change from accelerometer data stored on the phone.

2 Sampling algorithm

In subsequent sections we discuss ideas about improving modality sensing by incorporating high-frequency accelerometer data. Modality changes are rare, however, and it is unnecessary to continually probe and analyse accelerometer data throughout a commuter trajectory. At the same time, sampling should be done as sparsely as possible to minimize battery drain. In this section we discuss a possible scheme for acquiring high resolution accelerometer readings through irregular short bursts.

We propose 3 modes of data sampling:

- burst: high frequency sampling for a fixed short duration
- cruise: variable frequency sampling
- sleep: low frequency sampling, only when stationary

Switching between the modes is indicated in the scheme below:

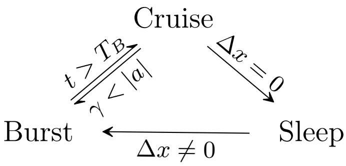

Burst mode is a new feature for Mobidot and its main purpose is to enable high frequency sampling for modality detection from accelerometer data. In this mode, both accelerometer data and GPS data are collected. The Burst mode lasts for a fixed amount of time T_B seconds and the time step between two consecutive measurements is short Δt_B. After Burst mode, the system always goes to Cruise mode.

Cruise mode is a medium frequency sampling mode to track the trajectory and to detect sudden changes, which call for a change back to Burst mode. The sampling frequency of cruise mode is adaptive, depending on the change in acceleration as derived from the GPS data. When Cruise mode is evoked after a Burst, the time step starts at Δt_{min}.

Each time step later, the stepping time is adapted: $\Delta t = \max(\Delta t_{min}, f(a))$. The function f is linear with acceleration a in m/s^2. So $f(a) := \Delta t_{max} - c * (a - a_{min})$, where $c = \frac{\Delta t_{max} - \Delta t_{min}}{a_{max} - a_{min}}$.
This is shown schematically in the figure below.

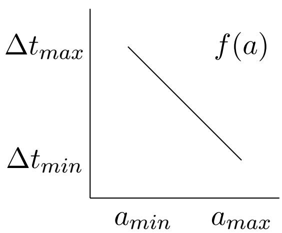

In Cruise mode only GPS data is collected. The system can leave Cruise mode when:

- location is stationary for a time span longer than T_S: go to Sleep mode
- there is a sudden change in acceleration $|a| > \gamma$: go to Burst mode

Sleep mode is only invoked when the device is stationary. This mode is characterised by low frequency sampling Δt_S. When significant location change is detected, the system goes immediately into burst mode, to detect the modality of the new trip.

The mode switching is demonstrated in the following algorithm:

if mode==Burst
    $\Delta t = \Delta t_B$
    if $T > T_B$
        mode =Cruise
    end
elseif mode ==Cruise
    $\Delta t = \max(\Delta t_{min}, f(a))$
    if $\Delta x == 0$ **and** $T_{stationary} > T_S$
        mode=Sleep
    elseif $|a| > \gamma$
        mode =Burst
    end
elseif mode==Sleep
    $\Delta t = \Delta t_S$
    if $\Delta x > 0$
        mode =Burst
    end
end

3 Discrete Wavelet Transform

In a practical application such as Mobidot's smartphone app, a continuous function f is often sampled at a finite set of discrete time points. Let us define the function values as $x_i = f(t_i)$ where the $t_i = i\Delta t$, $i = 0, \ldots, n-1$ are n equidistant sample times.

Time series analysis is a large field. The Study Group has investigated the use of wavelet analysis as a potentially efficient means of transforming discrete accelerometer data into a form suitable for training a learning algorithm to distinguish transport modality.

Analysis of a discrete time series requires a discrete version of the wavelet transform, and we will explain this for the simplest possible wavelet, the Haar wavelet. This wavelet is based on computing sums and differences of neighbouring function values, thereby representing both averages and differences. The *averages* give a smoothed version of the time series, with hopefully more reliable statistics, whereas the *differences* reveal whether a significant change has occurred over time.

One level of the wavelet transform is defined by

$$y_{2i} = x_{2i} + x_{2i+1}, \qquad y_{2i+1} = x_{2i} - x_{2i+1}, \qquad \text{for } i = 0, \ldots, n/2 - 1. \tag{1}$$

We denote this by $\mathrm{dwt1}(\mathbf{x}, n)$, where the output $\mathbf{y}$ overwrites the input $\mathbf{x}$,. We can also carry this out on the first k components of the vector $\mathbf{x}$ of length n, in which case we write $\mathrm{dwt1}(\mathbf{x}, k)$.

In the complete Discrete Wavelet transform (DWT), all differences are recorded and then they remain unchanged afterwards. For the sums, however, the procedure is repeated, but now with half the previous length. This is facilitated by first permuting the vector $\mathbf{y}$ by an even-odd sort, giving

$$z_i = y_{2i}, \qquad z_{i+n/2} = y_{2i+1}, \qquad \text{for } i = 0, \ldots, n/2 - 1. \tag{2}$$

We denote this by $\mathrm{sort}(\mathbf{y}, n)$, again assuming the output $\mathbf{z}$ overwrites the input $\mathbf{y}$. The complete DWT is given as Algorithm 1.

Algorithm 1 Discrete Haar wavelet transform

Require: $\mathbf{x}$: vector of length n, with $n = 2^m$.
Ensure: $\mathbf{y}$: vector of length n, $\mathbf{y} = \mathrm{DWT}(\mathbf{x}, n)$.

```
while n > 1 do
    dwt1(x, n) ;
    sort(x, n) ;
    n := n/2;
return x;
```

As a result of executing the DWT, we obtain an output vector $\mathbf{y}$ with

$$y_0 = x_0 + \ldots + x_{n-1}, \tag{3}$$

so that y_0/n equals the average of all the input values. The next value

$$y_1 = (x_0 + \ldots + x_{n/2-1}) - (x_{n/2} + \ldots + x_{n-1}) \tag{4}$$

indicates whether a significant change in mean value can be detected between the first half of the time series and the second half. The other values y_i give such change information at more detailed levels of accuracy.

The total number of additions and subtractions in the DWT algorithm for the Haar wavelet equals $n + n/2 + \cdots + 2 \approx 2n$, which is significantly smaller than the $5n \log_2 n$ floating-point operations (additions, subtractions, and multiplications) that would be needed for a standard radix-2 Fast Fourier Transform. For example, for $n = 1024$ the Haar wavelet transform is a factor 25 cheaper than the FFT. Note that the Haar wavelet does not need any multiplications (in contrast to the commonly used Daubechies wavelet), making it particularly cheap. Its use would lead to a much lower energy consumption in case the transform is computed on a smartphone.

4 Application of wavelet analysis

A Fourier transform of a signal in time only gives information on what frequencies are present in the total signal and thus the time-domain is lost. The wavelet transform provides a way to preserve the time-domain while also obtaining information about the frequency domain. The wavelet transform is most easily understood from the formula of the continuous wavelet transform

$$\Psi(\tau, s) = \frac{1}{\sqrt{s}} \int f(t) w\left(\frac{t-\tau}{s}\right) \mathrm{d}t,$$

where τ is the location of the window over which we integrate, s is the scale, and $w(t)$ is the wavelet, for example the Haar wavelet as shown in Figure 1.

In this section we apply the theory of the wavelet to the accelerometer data collected by Study Group. However, first we apply wavelet analysis to an example.

Consider the signal composed of two cosine functions

$$f(t) = 2\cos(20\pi t) + \cos\left(50\sqrt{2}\pi t\right). \tag{5}$$

This function is sampled on $t \in [0, 2]$ with steps of size 10^{-3}, leading to $n = 2000$ data points:

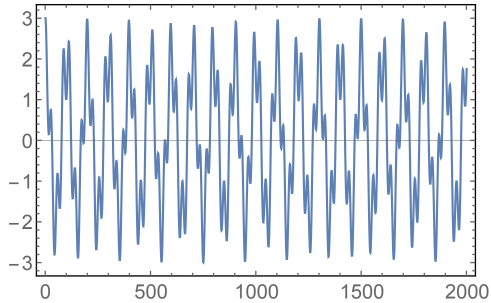

The wavelet transform is performed with the Haar wavelet using Mathematica and is shown in the lower left in Figure 1. Each box, numbered 1 to 7 in the plot, contains the same t-axis shown in the above image, namely $t \times 10^3 \in [0, 2000]$. As explained each level of refinement corresponds to convolution with the Haar wavelet at a different scale. The signal is thus compared to the Haar wavelet at each scale. In turn, each scale can be seen as a trade-off between frequency resolution and time resolution. In Figure 2 we compare the wavelet-transforms of each of the two cosines in the function $f(t)$ individually. Here we can see that at the sixth level of refinement there is a clear distinction between the two. The weight in the Haar transform of the complete function thus indicates that on this scale of the data the function contains a prominent slowly oscillating component. For each scale we can determine the energy fraction: $\{0.003, 0.012, 0.046, 0.138, 0.228, 0.385, 0.156, 0.0321\}$, where the energy of each scale is determined by the sum of squares of the values. The energy fraction can be used to rank the dominant contribution in the signal.

Although the Haar-wavelet transform is computationally very efficient, other wavelets can be tailored to find specific signals in data. For example, if we were to use the Meyer-wavelet the two oscillatory signals are clearly distinguishable as shown on the right-hand side of Figures 1 and 2. This suggests that we might use wavelets to pick out certain features in the Mobidot data specific to bicycles, trains or buses, by optimizing the wavelet basis for specific transport modes.

Next we apply the wavelet transform to the one-component of the 3-axis accelerometer data acquired by team member Jason Frank using a third party smartphone App. It is important to stress that a much more thorough data acquisition program is needed to characterise transport modes accurately. Here we provide only a random sample of time series data to indicate that differences can be discerned. Furthermore, our sampling rate was approximately 100–150 samples per second, which may be too low for detecting mechanical

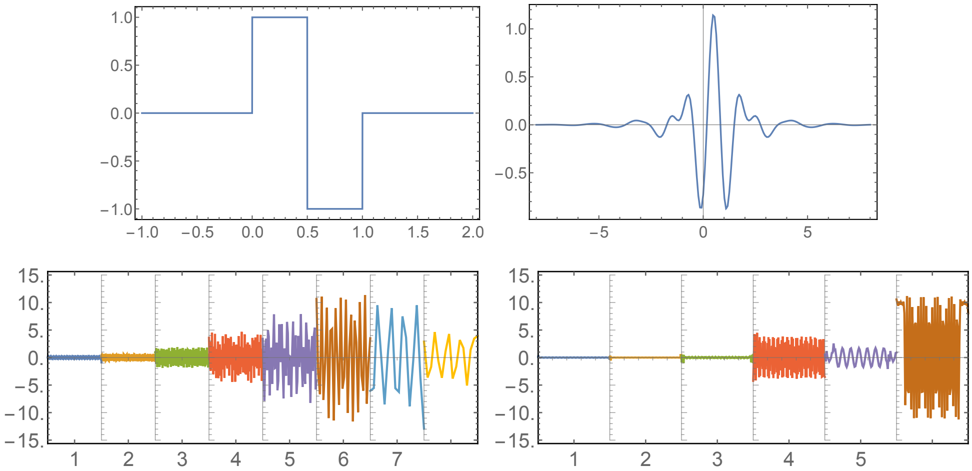

Figure 1: Left: The Haar wavelet (top) and the Haar transform of $f(t)$ (bottom). Right: The Meyer wavelet (top) and the Meyer transform of $f(t)$ (bottom).

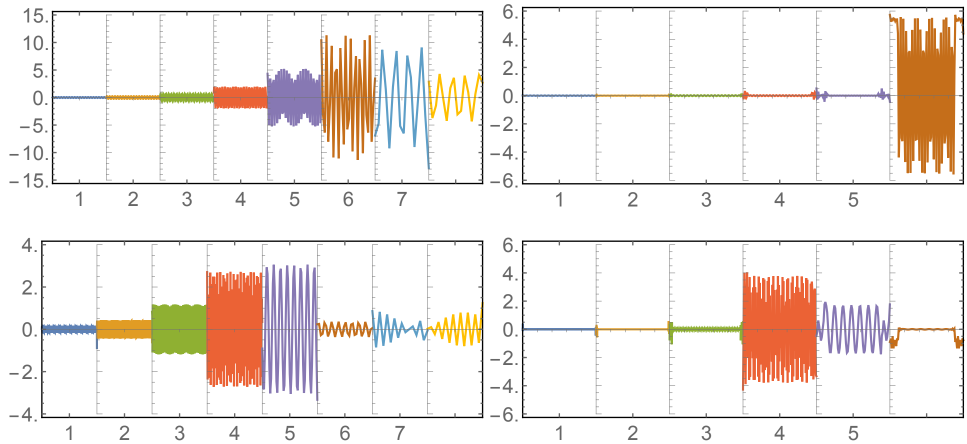

Figure 2: Left: The Haar-wavelet transform of the first (top) and second (bottom) term of $f(t)$. Right: Similarly, but for the Meyer wavelet.

vibrations in automobiles and buses. For visualisation purposes we will use the continuous wavelet transform with the Mexican-hat wavelet:

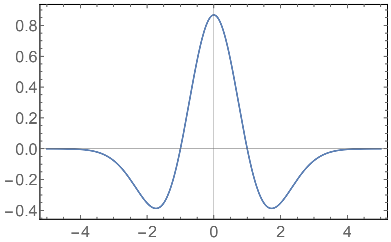

The time series data and its wavelet transform are shown in Figure 3. The results are shown with the scale (refinement) on the vertical axis and time on the horizontal axis. Each horizontal line in the image shows the absolute value of the wavelet transform with darker colours corresponding to larger magnitudes (cf. the amplitudes of the functions shown in Figure 2). Thus the image shows the dominant contribution to the wavelet transform at a particular scale.

It is clear that the bicycle shows distinctive features when analysed with this wavelet, which indicates that bicycles might be discernible. The car in general has high energies at low frequencies, while the train and bus are relatively quiet. A possible way to train mode-detection using wavelet analysis on the Mobidot data would be to record a sample of accelerometer data of fixed length of time, and subsequently apply the Haar transform to see if distinctions between different modes of transport can be identified based on which scales contain the largest energy fraction. If the Haar wavelet does not provide a clear distinction between modes of transport one could start to train on different wavelets. Using combinations of Haar wavelets one could potentially develop a training method which varies the wavelet form until an optimum wavelet is found. However, since Mobidot would like the wavelet transform with the optimum wavelet to be implemented on the mobile device, there will in general be a competition between the ease of detection and the computational cost.

5 Determine specific moment of transportation modality change

Continual recording of GPS location would prohibitively drain smartphone battery charge. Therefore Mobidot collects GPS data at regular intervals depending on detected movement. As a result it is hard to determine the exact starting time of a trip. If a change is detected one can only guess the starting

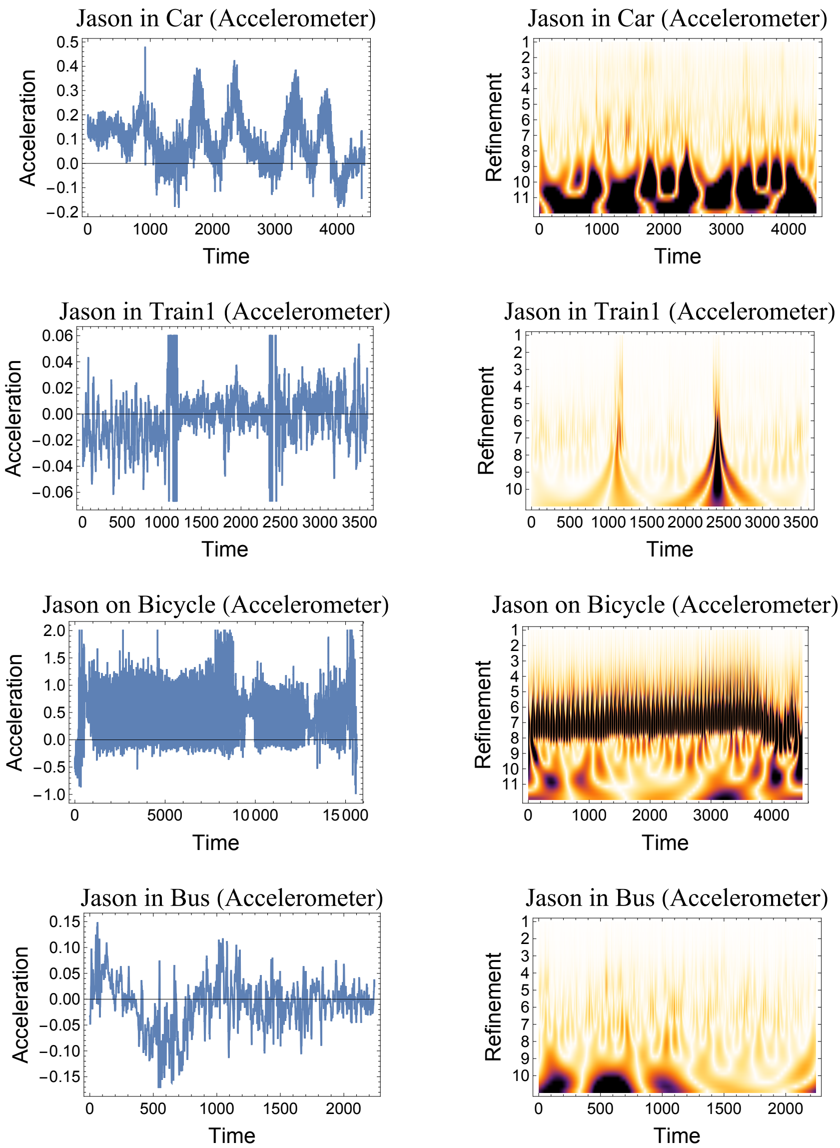

Figure 3: Accelerometer data and Mexican-hat wavelet transform for four different modes of transport.

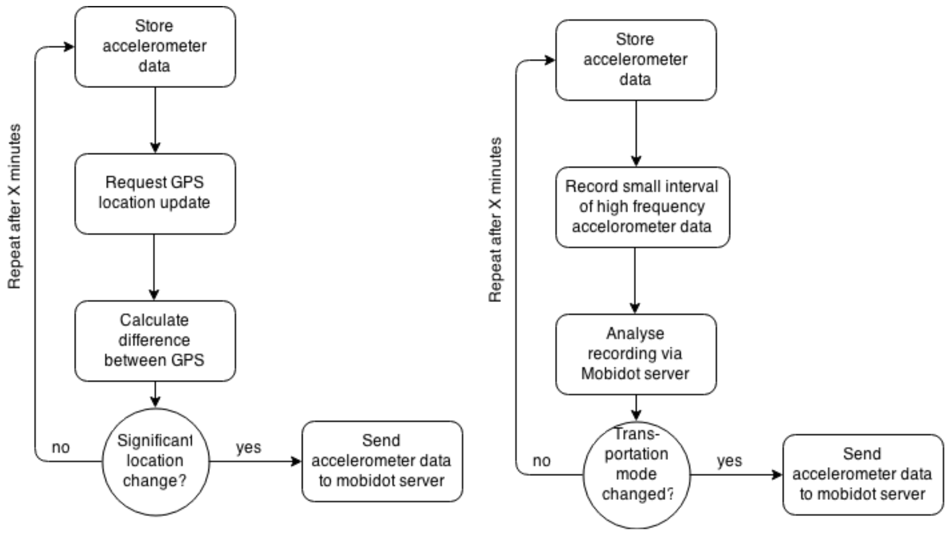

Figure 4: Visual representation of mode change detection via accelerometer sensor.

time via interpolation methods.

To overcome this problem the Mobidot application could augment GPS update requests with accelerometer data. The accelerometer in smartphones is a background application that is used for to detect gestures in the device. One can think of horizontal or vertical display switching or the natural 3D shadow effects for icons. Assuming the accelerometer is already running continuously, we can collect and store this data temporarily. After a certain time interval, the location detector can give an update about the current location. When no location change is detected, the accelerometer data can be deleted. However, if there is a significant change between the current and previous location we can analyse the accelerometer data of that time interval. This time interval can be analysed via the wavelet transform proposed in Section 4 to determine the moment of mode change.

This method can be used to detect the start of a trip, but it can also detect a change in transportation mode. To give a better overview of the method a graphical representation is shown in Figure 4. The left figure indicates the update method when the phone is in stationary mode. The right one represents a regular check to see changes in transportation type within a trip. Via small bursts of high frequency accelerometer data as described in Section 2 a mode change can be detected. If a change is detected we can use the temporarily stored accelerometer data from the phone and track at which moment the change took place.

6 Conclusions

We have reported on preliminary research to improve transport modality sensing using smartphone data acquisition. Our primary conclusions are:

- High resolution accelerometer data exhibits noticeable differences among different modalities such as bicycle, automobile, bus and train. Possibly this approach could be combined with currently used location service data to improve modality inference.
- Wavelet transforms offer an inexpensive and potentially powerful approach to obtain local frequency information on accelerometer signals. A more thorough multi-scale application of wavelets yields distinctive pictures of transport signals, that could be used to train learning algorithms.
- Accelerometer data may be effectively sampled in short, high-resolution bursts. These acquisitions can be incorporated in a multi-phase sensing strategy to preserve battery charge.
- The wavelet approach and multi-phase sensing strategy can be combined to improve the detection of mode changes during transit.