

I am a Smartphone and I Know My User is Driving

Hon Chu, Vijay Raman,
Jeffrey Shen,
Duke University

Aman Kansal, Victor Bahl
Microsoft Research

Romit Roy Choudhury
University of Illinois

ABSTRACT

We intend to develop a smartphone app that can tell whether its user is a driver or a passenger in an automobile. While the core problem can be solved relatively easily with special installations in new high-end vehicles (e.g., NFC), constraints of backward compatibility makes the problem far more challenging. We design a Driver Detection System (DDS) that relies entirely on smartphone sensors, and is thereby compatible with all automobiles. Our approach harnesses smartphone sensors to recognize micro-activities in humans, that in turn discriminate between the driver and the passenger. We demonstrate an early prototype of this system on Android NexusS and Apple iPhones. Reported results show greater than 85% accuracy across 6 users in 2 different cars.

1. INTRODUCTION

The synergy of sensing, computing, and communication on modern smartphones is enabling high resolution insights into human behavior. Recent research has attempted to leverage these insights for improved personal activity recognition, ranging from simple activities such as walking and running [3] to more sophisticated ones like, whether the user is laughing in a social gathering or is riding a car, bus, or a train [1, 7]. In this work, we intend to add *Is_Driver?* to this library of detectible activities. The goal is to discriminate whether a phone's user is the driver or the passenger in a car, thereby enabling a variety of vehicular applications on the smart phone. For instance, several car insurance companies are aiming to personalize the insurance rates paid by individuals. Their aim is to charge a premium in proportion to the number of hours he or she drives. Logging the hours of driving is an important step towards this direction. In a different application, a passenger in a car may be allowed to receive phone calls, location based notifications, or social network updates. However, such distractions may need to be suppressed for the driver to uphold driver safety [4, 6]. Our proposed *driver detection system* (DDS) would be an useful building block for these applications.

The core driver detection problem may be approached from multiple directions. Newer cars have a range of sensors and radios that could aid in detection. For instance, (1) *near field communication* (NFC) radios on the drivers door could be programmed to identify the mobile phone that is closest to it; the owner of this phone should be the driver. (2) Pressure

sensors on car seats could be used to measure the weight of the person on the driver seat; knowing the weights of people who are likely to drive that car could facilitate driver detection. (3) Given that most cars have audio speakers, a recent work from MobiCom 2011 [5] generated a sound signal from each of the four speakers in a programmed sequence – the phone recorded the sounds and triangulated its location to one of the four quadrants in the car. While all these approaches indeed identify the driver, they require some degree of modification to the car's inner workings. This paper targets the problem of driver detection with zero modifications to the car. Our goal is to leverage the rich suite of smartphone sensors, including the accelerometer, gyroscope, and microphone (in some infrequent cases). Users of our system can continue to use their existing cars, however old, making our system simple and backward compatible.

Our key idea behind the DDS system is simple. We hypothesize that the driver and the passenger perform non-identical micro-activities that can be captured and discriminated through multi-modal sensing. For instance, a driver inserts her right foot inside the car first, while the passenger does the opposite. Thus, if the inserting foot has a particular motion signature on the phone's accelerometer, then observing that signature reveals that the phone is either on the right pant pocket of the driver, or the left pant pocket of the passenger. Now, if its the former, then the driver should be pressing the gas pedals and breaks with her right foot – such motion signatures should be easily detectible. However, if such signatures are absent, it might be possible to conclude that the owner of the phone is a passenger.

Of course, the problem is actually more complicated because the user may carry the phone in her shirt pocket, or even in her purse. If the phone is in the shirt pocket, we observe that the driver is likely to turn slightly towards the left to pull the seat belt, while the passenger would again perform the opposite. Gyroscopes in modern smartphones are capable of measuring rotational along a vertical axis, called *yaw*. By observing the direction of the yaw, we should be able to perform the driver-passenger classification. Of course, the women's purse is harder to detect; nonetheless, we find that if the user throws her purse in with a reasonable motion, then the direction of the motion (left to right, or right to left) can be used for classification. Finally, if a car has some passengers in the rear seat, then its also necessary to

discriminate between the driver and the person sitting behind the driver. Since both their motion signatures could well be identical, we utilize the audio sound levels from the car stereo as a way to discriminate between them. DDS is an implemented system that aggregates a range of signatures and combines them to achieve consistent classification.

Of course, identifying signatures for each of the micro-activities is non-trivial – several discriminating patterns lay hidden in the raw sensor signals and needed to be extracted with precision. Moreover, the system needs to be developed with energy constraints in mind. Finally, the system needs to recognize when the confidence of classification is low, and offer this failure notification to the user. This may be important if the system is used for law-enforcement type applications. In such scenarios, it would be necessary to learn about the confidence of classification. DDS prototypes a functional system on the platform of iPhone OS 4.0 and Android OS 2.2, and tests the driver detection performance through 2 cars and multiple users. Performance results demonstrate a reasonably consistent accuracy of 85%, and failures occurring mostly when women carry the phone in their purses. While DDS is not ready for immediate deployment, we believe that the system invokes a sense of promise. Perhaps more importantly, the individual micro-signatures in DDS may themselves be useful building blocks to an emerging app store, dedicated to vehicles.

2. DRIVER DETECTION

Driver detection enables several useful applications. As mentioned in the previous section, one application is to control notifications based on user attention, since in-vehicle information delivery is highly sensitive to user state [6]. Driver detection may also feed into other systems such as personalized insurance, tracking carpooling, optimizing car-pool lanes and incentives. Combined with detecting transportation mode [10] it can be used to track users' commuting carbon footprints. For newly licensed teenage drivers, the detection helps track how many hours of actual driving have been performed, as distinct from total in-vehicle time spent that includes being driven around by parents. Such information may be used by parents to decide when the teenager may be provided their own car.

2.1 Problem Outline

Our goal, simply stated, is to enable the mobile phone to determine if its user, when in a car, is a passenger or a driver. Several alternatives are possible to realize this goal. The car itself has some sensors that could aid driver detection:

NFC: New phones have a near field communication (NFC) radio (e.g., for mobile payments). Some cars also have NFC radios for key-less entry. If the car's firmware can be modified to use the NFC channel to inform the phone which door was used, even if the user was not the one unlocking the car, the phone can determine if the user is a driver. However,

only a small fraction of existing and new cars have NFC and it is only installed on front doors.

Audio: Most cars have a speaker near each of the 4 corners of the car. Each time the car starts, it can generate a sound pattern that is emitted from the four speakers one after the other. The phone can use its microphone to triangulate its position with respect to the speakers. As with the previous approach this requires car firmware to be changed, though the hardware is more widespread than NFC.

However, given that phones are replaced every 2 or so years while the median age for cars is 9.2 years in the US, methods that require modifications to the car itself will be very slow to deploy. On the other hand, deploying additional software on smart phones is significantly easier. Thus, in DDS we use only the phone's built-in sensors. The key intuition behind our approach is that the phone sensors include an accelerometer that measures acceleration along three axes, and a gyroscope that measures rotation along three axis, denoted roll, pitch, and yaw. These sensors can be used to detect user movement patterns in an automobile to distinguish drivers from passengers.

2.2 Challenges

We show experimental data to help identify key challenges in detecting relevant user movements. One distinguishing micro-movement is that of the driver's right leg pressing the gas or brake pedals, which is absent in passengers. Sensor data for a short time duration around one such pedal press are shown in Figure 1. Several issues are apparent:

Multiple Movements: The data has significant noise. The motion of the vehicle causes the accelerometer and gyroscope outputs to change continuously. At a coarse granularity (Figure 1(a)), the signals are similar for the driver and passenger. The pedal related movement is not a simple sensor value but a specific pattern (Figures 1(b)) that varies somewhat with each pedal press and occurs interspersed with other patterns.

Phone Position and Orientation: The driver's leg movements can only be captured correctly when the phone is in the correct pocket, the one on the right leg. The phone does not capture these movements when placed in the left leg pocket (Figure 1(b)) since the left foot is typically not used for any pedals (except if the car is equipped with stick shift), or any other pocket. Even when the movement pattern is present, it can vary depending on how the phone is oriented within the pocket: Figure 1(c) shows that the pattern is more pronounced on the roll axis when the phone is carried in a horizontal orientation and on the pitch axis when the phone is vertical. Therefore, the solution must be flexible in terms of phone position and orientation. Clearly, multiple distinguishing signals may have to be employed.




Figure 1: Sensor data for driver and passenger movement patterns, for a brief time interval around a pedal press movement: (a) accelerometer (x-axis) and gyroscope (roll) about 60s surrounding the movement, (b) gyroscope roll zoomed to about 8s near a pedal press, and (c) gyroscope data with two phone orientations showing that the signature of interest can shift from one signal dimension to another.

Short-lived Signatures: The distinguishing movement occurs only intermittently and the sensor output is mostly similar for the driver and passengers due to the shared vehicular motion. Thus, driver detection cannot be performed on demand when a notification is to be provided. Instead, the movement patterns must be detected as they occur. Continuously sensing and processing the data to look for signals of interest would drain the battery at an unacceptable rate. Also, the output must be provided in soft-real time, rather than in an offline analysis mode, since the output is most relevant when the user is driving.

DDS systematically combines multiple sensor signal patterns to address the above issues.

2.3 Solution Design

The detection of multiple relevant micro-movements for various phone positions and orientations leads to multiple cases that DDS design should address. Fortunately, simplifications and commonalities across cases exist. Firstly, to overcome the variable phone orientation, we map all signals to a reference orientation. The phone compass and accelerometer data can be combined to get the earth’s gravity vector and magnetic field which suffice to generate this mapping. The Android OS and iOS provide APIs to support this calculation. Secondly, movement constraints help prune the search space of movement signatures in each case. For example, a driver enters the car using her right leg first while the front passenger enters using her left leg first. The pedal press should be searched only if the phone is in the leg pocket that enters first.

Figure 2 presents a systematic combination of multiple sensor signal signatures for various possible phone positions

(lower body pockets, upper body pockets, and handbag), to determine if the user is a passenger or a driver.




Figure 2: Decision flow: P = passenger, and D = driver, denote the outputs. The “RIGHT” output of Handbag L/R indicates passenger (omitted for clarity).

The sensor signatures used in the above diagram are:

Entry Swing: This block determines if the user is in the right or left side of the vehicle cabin, for the phone carried in a lower body pocket. The swing in the user’s lower body reveals the direction of entry (left vs. right). Furthermore, the movement can be separated into two significant parts to determine if this signature occurred first on the leg closest to the phone or the other one, revealing if the phone is in the pocket on the innermost leg (the one that entered the vehicle first) or not. Accelerometer, compass, and gyroscope data are used.

Seat-belt: This block detects right vs. left side for the phone carried in an upper body pocket. The users on the left will wear their seat belt by first turning left and then turning right as they pull the seat belt to the fastener. The motion is reversed for the other side. Accelerometer, compass, and

gyroscope data are used.

Pedal Press: This block determines if the user’s leg is being used for pressing the brake or gas pedals. The sensor data is matched with pedal press patterns using a clustering algorithm. This block is used only if the above block outputs that the user is in the left side of the cabin and the phone is in the user’s inner leg pocket. Accelerometer, compass, and gyroscope data are used.

Handbag L/R: This block detects left vs. right for the phone carried in a handbag, if the other two side detection blocks fail. The detection is based on nearness to the driver location and is similar to the front vs. rear detection described below.

Front vs. Rear: This block determines if the user is in the forward or rear portion of the vehicle cabin. Some of the branches from the above blocks directly reach the result state (P or D , denoting passenger and driver respectively) and this block is required only if the other branches were inconclusive, or if the user was found to be a driver (so that it can be discriminated against the rear passenger). This block uses communication with a back-end cloud service and could involve higher energy use. However, no direct communication is assumed among the phones and hence neither user is required to enter a password for Bluetooth pairing. Microphone data is used.

The entry swing and seat-belt blocks operate in parallel since the position of the phone on the user is unknown. If the entry swing signature detection fails to detect a signature, it is assumed that the phone is in the upper body area and hence, the seat belt signature may succeed, and vice versa. If neither succeeds, the handbag block is attempted. If that fails, but vehicle motion is detected, the result is inconclusive. The figure only shows the primary decision flow for clarity. Other redundant paths are possible and may help to reduce the error in overall operation. We assume a left hand drive vehicle but the flow is easy to change if the user is in a region where right hand drive vehicles are used, as may be determined from the location, time-zone, and language settings on the phone. The inference methods used for each block and the trigger events to activate the sensors in an energy aware manner are described next.

3. SENSING ALGORITHMS

The processing blocks shown in Figure 2 are described in detail below:

3.1 Building Blocks

Entry Swing: The entry swing detection is based on the intuition that for a phone in a lower body pocket, the sideways movement indicates the direction for the car entry. Figure 3(a) shows the gyroscope roll during a swing, for entry from left and right. The signal shape is similar for both right

and left entries, but the change in roll shows an opposite sign. The data also shows two peaks with a different absolute magnitude. These peaks in fact correspond to the two legs entering the car. The data shown here is for the phone placed in the inner leg pocket (the leg that enters the car first). For the phone in outer leg pocket (not plotted for brevity), the absolute magnitudes of the peaks are reversed. Effectively, this sensor signature allows determining both the *direction of entry* and if the phone is in the *outer or inner leg* pocket. Human movements are not reproduced exactly each time and vary across users and vehicles. To reliably identify the signature, we use a support vector machine (SVM) to classify input data as matching right-entry, left-entry, or nothing. The signal features used for training the SVM are: peak and trough counts of three axes gyroscope data, variance of gyroscope values, and their derivatives taken over short time windows. The peak and trough counts help distinguish the entry direction. The variance feature helps since variance is higher during the entry swing than at other times. The derivative is also helpful because the rate of change is high when the signature is present.

Seat-belt: The seat belt signature is based on the rotation of the user’s body for wearing a set belt. In the global reference frame, the rotation of the body maps to the gyroscope yaw. However, accelerometer and compass are also required to orient the data from an arbitrary phone orientation to the global reference frame. The rotation is only apparent for the phone carried in an upper body pocket such as the shirt pocket on the user’s chest; there is no significant rotation on the lower body. A sample data trace of gyroscope yaw for the time duration of a seat belt micro-movement is shown in Figure 3(b) for users in the left and right sides of the car.




Figure 3: Sensor data for driver and passenger movements: (a) entry swing for left and right entry (gyroscope roll), and (b) seat-belt for left and right vehicle sides (gyroscope yaw).

We use an SVM to classify input signals into left, right, and no-signature cases. The SVM is trained using the following features: first derivative of the yaw, variance of the yaw, and the variance of the three axes accelerometer data. Interestingly, even though the signal is most prominent along the gyroscope yaw, the accelerometer data still helps as it contains patterns related to getting and buckling the seat-belt which distinguish the signal from no-signature cases.

Pedal Press: This block detects the leg movement when pressing the brake or gas pedals. Sample data was shown in Figure 1(b). This micro-movement only occurs on the driver’s inner leg and is used only when the phone is suspected to be in this position. The complete signature consists of a sudden change in the pitch or roll when pressing the pedal, followed by a period of no activity for the period that the pedal is kept pressed, and then finally a change in the pitch or roll, as the pedal is released. The movement is more pronounced for the brake pedal than the gas pedal. The features used for the pedal SVM are: second derivative of the gyroscope pitch and roll and zero crossings for the second derivative. The second derivative was found to be most useful here because a pulse like pattern, corresponding to increase in pitch/roll, followed by no change in values and then a decrease of pitch/roll, is being detected.

Front vs. Rear: If the user is detected on the left, but the phone is not in the correct position to detect the pedal press, we need another way to determine if the user is the driver or the rear left passenger. Among multiple possibilities explored for this (some are summarized later under Discussions section), we found that comparing the amplitude of the turn signal sound worked most reliably. There are two key challenges here. Firstly, the absolute amplitude of the turn signal sound is not useful since it varies across cars and due to other attenuating factors. Rather, a comparison between the amplitude observed by the phones is needed. Since direct phone to phone communication requires user intervention (such as entering a Bluetooth PIN), we rely on a cloud server to perform the comparison. The back-end server must determine if two or more phones uploading data to it are from the same vehicle based on the sensor data itself. To this end, we include the phone location (coarse grained location if fine grained location is not available), timestamp, and accelerometer data, to aid in grouping phones from the same vehicle. Among the phones that start the trip in the same time window and region, accelerometer data helps identify if the overall movement pattern of the phones is the same, indicating the shared motion of the same car. Any of existing signal matching methods may be used to detect if the signals show the same pattern. In our implementation we use frequency domain correlation.

A second challenge is that the turn signal sound is mixed with other sounds such as road noise and music. These other sounds are often equally loud in the front and rear parts of vehicle cabin and do not have any distinguishing value. To minimize their effect, we use a band-pass filter to separate the turn signal sound. Of course, some frequency components of other audio such as music that lie within this band also pass through. Based on experimentally obtained recordings of the turn signals in multiple vehicles, the band pass filter designed is listed in Table 1.

Sampling	Stop band	Pass-band	Ripple
44.1 kHz	< 2.9kHz, > 3.1kHz	2.95-3.05kHz	0.057501127785

Table 1: Band-pass filter for turn signal.

Handbag L/R: If the phone is in a handbag, the side detection using entry swing and seat-belt is not applicable. If only one user in the car is carrying a handbag and all others have been identified to be passengers or driver, the phone in the handbag is easy to identify. If the handbag is placed on the rear seat or the trunk, DDS won’t work but in that case we assume that the phone is inaccessible to the user and therefore, driver detection is irrelevant. If the handbag is placed in the driver area (near the left foot or cup-holder area) and another user is also using a handbag, driver detection is important. To distinguish the phone in a handbag being closer to the driver area than the passenger area, we again use turn signal audio comparison similar to the above block.

Audio data is uploaded if the user state is not determined by the other blocks. Also, if the user is determined to be a driver, data is still uploaded to enable any other user in the car to correctly determine their state. If the back-end server receives only one upload from a vehicle, it is automatically assumed to be the driver.

3.2 Event Triggers

DDS also needs mechanisms to activate the sensing and inference at appropriate times since battery constraints prevent continuous sensing. We use a duty-cycled approach to detect when the first trigger event occurs and then activate sensors from that trigger event for a duration that ends based on another trigger event, indicating that sufficient data has been collected. The trigger events used are:

Walking and Pause: Detection of walking using accelerometer data has been used in several prior works and we use a simple amplitude based method to detect walking. The detection is duty cycled: the sensor is activated for 1s every 60s. More sophisticated methods may allow reducing the active time from 1s. Walking may not be detected if it lasts shorter than 60s, and this period may have to be reduced for users who walk less than 60s to reach their car. When walking is detected, the accelerometer is turned on continuously to detect when the user pauses (to open a car door if they were walking to a car).

Vehicular Motion: The start of vehicular motion is also easy to detect from accelerometer data and has been used in prior works [10].

These trigger events are used to start/stop sensor data collection and activation of detection blocks, as shown in Figure 4 in the form of a time line. The front vs rear block and hand-

bag L/R block are started after the time instant marked Start Audio.

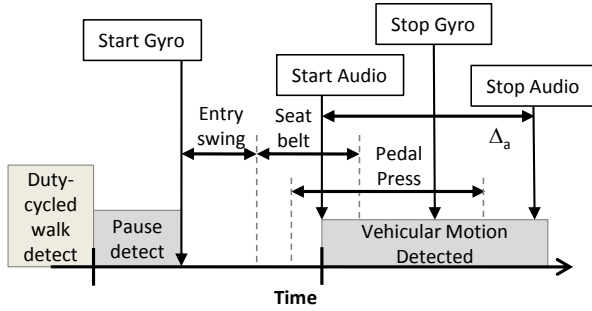


Figure 4: Time line for sensing based on trigger events.

Timeouts: In addition to the events, certain timeouts are employed to stop sensors and inference methods. The gyroscope data collection is stopped 15s after the vehicular motion is detected. The duration of 15s was chosen to ensure that the user wears their seatbelt, though in most cases, users wear the seat belt before the vehicular motion starts. Pedal press occurs within this time as well. Audio data collection is stopped after a duration Δ_a . The value of Δ_a is determined at run time based on the number of filtered audio samples above a threshold magnitude, to ensure a robust comparison. If the user was not walking to a car, no DDS signatures occur after the pause and the system shuts down after a timeout.

An example data trace from a phone carried in an upper body pocket is shown in Figure 5, illustrating the detection of walking and vehicular motion trigger events on accelerometer data and one of the sensor signals used for driver-detection, in this case the gyroscope yaw, indicating a seat-belt signature.




Figure 5: Trigger events and a sensor signature for phone carried in upper body pocket.

4. EVALUATION

Experimental Data Collection: The DDS system is evaluated with sensor data collected on Android OS 2.2 based Nexus Ss and iPhone OS 4.0 based iPods. The test data was collected in realistic settings from three distinct users with the phone carried in different positions on their body without controlling the phone orientation. The car was driven on routine routes that those users take. The in-vehicle music player was used. Multiple data traces were collected from each user, leading to 40 or more samples for each micro-movement.

4.1 Accuracy

One goal of the experiments is of course to quantify the error for each of the sensing blocks. Another is to verify if the inference methods trained on one user will allow the DDS system to work on another user, allowing it to be used without per user training. To validate this, all results presented are obtained through cross validation tests that exclude the test user’s data from training.

Entry Swing: Since this micro-movement is present in almost all traces collected, we have over 400 trials to test. As mentioned, training data from a given user was not included when testing test samples from that user. In each case 80% of the data is used for training and the remaining 20% for testing. To remove the effect of any bias in certain samples, 100,000 of the different possible splits among the samples are considered, each leading to one test run. The cumulative distribution function (CDF) of entry swing classification accuracy is shown in Figure 6 – average accuracy = 88.99%.




Figure 6: Entry-swing accuracy CDF.

Seat-Belt: The seat-belt signature is evaluated similarly. Again performing cross validation on all user samples collected, and measuring error over all runs of the data splits, we obtain the accuracy CDF shown in Figure 7. The accuracy is high across all runs and the mean accuracy is 91.08%, also marked in the figure.

Pedal Press: Following the same cross validation process, the CDF of correct detections is shown in Figure 8, and the detection has a mean accuracy of 89.78%.

Front vs. Rear: The first step for this block is detecting which uploads come from the same car. It is based on match-




Figure 7: CDF of correct seat-belt micro-movement detection, cross validated across all users.




Figure 8: CDF of correct pedal press micro-movement detection, cross validated across all users.

ing location area, time-window of the data timestamps, and signal shape of the accelerometer data. Matching time and location windows is trivial and hence we stressed the system by grouping only based on accelerometer. We take 100 accelerometer traces from different trips and treat each trace as an upload from a distinct phone. We divide these traces into groups of up to n simultaneous trips with 2 phones within the same vehicle for each trip. The matching algorithm attempts to identify the n correct pairs out of the $2n$ traces ($^{2n}C_2$ possible pairs). Measuring the error as the number of false positives and false negatives determined by the algorithm, Figure 9 plots the accuracy achieved with varying n . Two different lengths of the uploaded trace, 30s and 60s, are tested. With the accelerometer sampled at 50Hz, this amounts to 1500 samples and 3000 samples respectively.




Figure 9: Cloud server matching of uploaded data traces from multiple phones.

The second processing phase compares the filtered audio data amplitude between the driver and rear passenger phones. Incorrect results may occur due to several factors such as the driver’s phone being enclosed in a pocket made from thicker and more sound absorbent material than the passenger’s pocket. We collected test data with both the music on and not, as well as trips on freeways where road noise is higher. Denoting the cases with *Low* volume music and *Low* road noise as *LL*, *High* volume music and *Low* road noise as *HL*, and so on, we get four cases shown in Figure 10. In each case we consider the phones being inside pockets or outside of any pockets, as well as a worst case scenario for amplitude comparison: driver’s phone is inside the pocket but the passenger’s phone is outside (implying that the driver’s phone suffers extra attenuation). The ratio of the maximum absolute magnitudes measured by the front phone to that of the rear phone is shown in Figure 10. Correct detection results when this ratio is greater than one (horizontal line in the figure). The values shown for each case are averaged over multiple trials. The mean accuracy across all trials was 95.83%. In the *HH* scenario where the turn signal sound is almost hidden in the music and road noise, we did see a ratio lower than one in some of the trials, when the rear passenger’s phone is outside the pocket while the driver’s phone is inside.

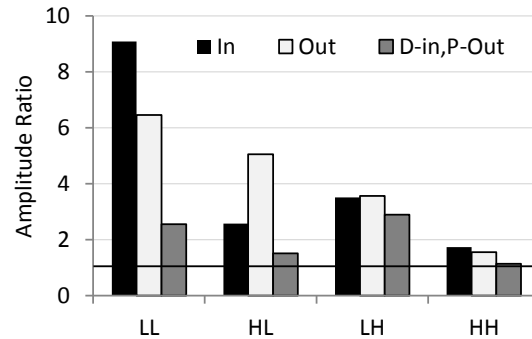


Figure 10: Front vs. Rear Detection: Ratio of audio amplitudes for driver and rear passenger in different scenarios. *L* indicates low and *H* indicates high; the first *L* or *H* refers to music and the second to road noise.

Handbag: As mentioned DDS ignores cases where the handbag is placed away from the driver (rear seat or trunk) as then the phone is inaccessible to the driver. Also, the handbag block is only used if the driver and at least one passenger carry their phone in a bag since otherwise, the correct detection on all other users will indicate the role for the phone in the handbag. In this scenario, the driver’s handbag can be placed either in the cup-holder area or on the floor near the driver feet. The phone may also be placed directly in the cup-holder without a bag.

Three cases arise based on the behavior of the passenger using a handbag for their phone: (i) the passenger handbag

is placed markedly farther from the driver console (turn signal sound source) compared to the driver’s handbag (e.g., driver’s phone in feet area, passenger phone on rear seat), (ii) both phones are a similar distance from the driver console though the passenger phone is farther (e.g., both driver and front passenger place their handbags near their left foot respectively), and (iii) the phones are placed together (e.g., both in the center cup-holder area).

Case (i) is expected to work accurately similar to the front vs rear detection block, as the audio intensities received at the two bags vary significantly. For evaluating case (ii), we once again collected the audio data, with music on and not, and with the car driven on roads with high and low road noise. Again using L and H respectively to indicate Low and High road noise or music volume, we tabulate the ratio of the audio intensity of the driver to that of the passenger in Table 2. We consider the cases when both the driver and passenger’s phones are inside a bag, both the phone’s are outside, and the worst case scenario: the driver’s phone is inside a bag and the passenger’s phone is outside. When the phones are inside a bag, they are placed near the left foot (for both driver and passenger), and when the phones are outside, the driver’s phone is placed in the cup-holder and the passenger’s phone is place on their lap. Ratios greater than one imply correct detection. We see that while most cases are addressed correctly, the detection fails in a couple of the worst case settings. Similar experiment for case (iii)

Music, Driving level	Amplitude ratio		
	In	Out	D-In,P-Out
LL	6.2305	1.3504	1.5535
LH	1.4006	2.2279	0.6498
HL	1.4477	1.0003	0.0532
HH	1.2301	8.2485	1.1887

Table 2: Audio intensities when phones far apart .

revealed that, as expected, this method cannot work reliably when both phones are equidistant from the driver console.

5. RELATED WORK

The widespread availability of sensors in mobile phones is enabling interesting sensing applications, including measuring traffic [11] and quantifying environmental impact [8]. Our work adds a new activity and application.

User activity sensing has been addressed in several works [1, 7]. We expand this repertoire to include the driving activity. Human movements have been sensed for multiple purposes in health-care [2]. However, their use driver detection is novel. Commercial products for disabling phone use in car are available [9] but they do not distinguish between drivers and passengers.

6. CONCLUSIONS

The design of DDS demonstrates how a small set of sensors, not mounted or worn in a special manner, and used in a noisy environment, can still be used to improve our information about an important user state. Experimental evaluations on real user traces showed that DDS achieves a practically useful accuracy in a majority of phone carrying positions, without requiring individual user training. The design also uses triggers to significantly reduce the energy overhead, illustrating a useful design principle for sensing applications that operate continuously.

7. REFERENCES

- [1] T. Choudhury, S. Consolvo, B. Harrison, J. Hightower, A. LaMarca, L. LeGrand, A. Rahimi, A. Rea, G. Borriello, B. Hemingway, P. Klasnja, K. Koscher, J. A. Landay, J. Lester, D. Wyatt, and D. Haehnel. The mobile sensing platform: An embedded system for capturing and recognizing human activities. *IEEE Pervasive Computing Magazine*, 2008.
- [2] B. Dobkin. Medical foundations and applications of human motion. In *Wireless Health*, October 2010.
- [3] M. Ermes, J. Pärkka, J. Mantyjarvi, and I. Korhonen. Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions. *IEEE transactions on information technology in biomedicine : a publication of the IEEE Engineering in Medicine and Biology Society*, 12(1):20–26, January 2008.
- [4] D. Hibberd, S. Jamson, and O. Carsten. Managing in-vehicle distractions - evidence from the psychological refractory period paradigm. In *2nd International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI)*, November 2010.
- [5] e. a. Jie Yang, Simon Sidhom. Detecting driver phone use leveraging car speakers. In *ACM MobiCom*, 2011.
- [6] I.-M. Jonsson and F. Chen. In-vehicle information system used in complex and low traffic situations: impact on driving performance and attitude. In *Proceedings of the 4th international conference on Universal access in human-computer interaction: ambient interaction*, UAHCI’07, pages 421–430, Berlin, Heidelberg, 2007. Springer-Verlag.
- [7] L. Liao, D. Fox, and H. Kautz. Learning and inferring transportation routines. In *Proceedings of the 19th national conference on Artificial intelligence*, AAAI’04, pages 348–353. AAAI Press, 2004.
- [8] M. Mun, S. Reddy, K. Shilton, N. Yau, J. Burke, D. Estrin, M. Hansen, E. Howard, R. West, and P. Boda. Peir, the personal environmental impact report, as a platform for participatory sensing systems research. In *Proceedings of the 7th international conference on Mobile systems, applications, and services*, MobiSys ’09, pages 55–68, New York, NY, USA, 2009. ACM.
- [9] Safe driving system 5080TGG. <http://www.skymall.com>.
- [10] A. Thiagarajan, J. Biagioni, T. Gerlich, and J. Eriksson. Cooperative transit tracking using smart-phones. In *ACM SenSys*, 2010.
- [11] A. Thiagarajan, L. Ravindranath, K. LaCurts, S. Madden, H. Balakrishnan, S. Toledo, and J. Eriksson. Vtrack: accurate, energy-aware road traffic delay estimation using mobile phones. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, SenSys ’09, pages 85–98, New York, NY, USA, 2009. ACM.