EMPress: Practical Hand Gesture Classification with Wrist-Mounted EMG and Pressure Sensing

Jess McIntosh, Charlie McNeill, Mike Fraser
Bristol Interaction Group,
University of Bristol
Bristol, UK
jm0152@bristol.ac.uk

**Frederic Kerber, Markus Löchtefeld,
Antonio Krüger**
German Research Center for Artificial
Intelligence (DFKI) and Saarland University
Saarbrücken, Germany
frederic.kerber@dfki.de

ABSTRACT

Practical wearable gesture tracking requires that sensors align with existing ergonomic device forms. We show that combining EMG and pressure data sensed only at the wrist can support accurate classification of hand gestures. A pilot study with unintended EMG electrode pressure variability led to exploration of the approach in greater depth. The EMPress technique senses both finger movements and rotations around the wrist and forearm, covering a wide range of gestures, with an overall 10-fold cross validation classification accuracy of 96%. We show that EMG is especially suited to sensing finger movements, that pressure is suited to sensing wrist and forearm rotations, and their combination is significantly more accurate for a range of gestures than either technique alone. The technique is well suited to existing wearable device forms such as smart watches that are already mounted on the wrist.

Author Keywords

Hand Gestures; Electromyography (EMG); Pressure; Force Sensitive Resistors; Practical Wearable Device Design.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

Increasing use of wearables has generated interest in novel gesture input techniques for mobile interactions. Of particular importance is integrating reliable and practical sensing into wearable devices that are ergonomically acceptable to designers and users. For smart watches and other wrist-mounted wearables, there are opportunities to detect interactions with the hand on the same side as the wearable [12]. However, current techniques can be impractical due to signal occlusion, requirement for additional sensors mounted independently of the wearable, and inability to sense different types of gestures.

We explore whether hand gestures can be accurately classified using sensors only on the wrist, the predominant location for existing wearable devices. Our gesture recognition approach is to explore the combination of Electromyography (EMG) and pressure, using Force-Sensitive Resistor (FSR) sensors mounted on the inside of a wrist strap. Our aim is to accurately classify a large range of hand gestures, detecting both finger and wrist movements, because many hand gestures include both these forms of movement. Throughout this paper we use the term 'gesture' to mean a discretely classified pose, rather than continuous measurement of hand movements.

Our pilot study found a high classification rate of finger gestures using wrist-mounted EMG sensors. We suspected that classification features derived from EMG signals were being modulated by variable pressure through the EMG sensors' contact with the wrist. This led us to build the EMPress system, which includes explicit pressure sensing in the gesture detection process, drawing on previous work on low-power gesture input with FSR sensors.

Our key contributions are:

- A novel design combining EMG and pressure data using machine learning to accurately detect and classify hand gestures.

- Two studies which identify and then quantify the performance of combining these sensors in the prototype EMPress system.

- Experimental evidence that these sensors are strongly complementary, emphasising EMG for detecting finger movements and FSR for detecting wrist movements.

Our first study investigates the effectiveness of wrist-mounted bipolar EMG electrodes for detecting single finger gestures. The results from this study exhibited surprisingly high EMG classification accuracy in the wrist condition. This high accuracy led us to suspect that variable wrist pressure on the EMG electrodes was modulating and enhancing our EMG finger gesture classification rate.

The pilot study therefore led us to an unexpected hypothesis that pressure data on the wrist strap could provide our classifier with features that enhance the EMG results. Our second study describes the EMPress technique which uses separate EMG and pressure sensors to isolate and quantify this effect.

Figure 1: Cross section of the forearm at the distal ends of the radius and ulna. Our final experimental prototype is seen around the wrist. The anterior of the arm is at the top. The muscles and tendons are coloured to represent groups with similar functionality.

The prototype we developed for our second study uses padded wet electrodes to moderate pressure effects on the EMG signal, and a cross-arm reference electrode to determine an upper bound of EMG performance. Our study demonstrates that pressure not only contributes to EMG classification of gestures under optimal conditions, but can actually parallel or surpass EMG performance across a range of gestures, and in complementary ways such that both techniques can be used together to further improve performance.

In the following sections we explain the relevant forearm anatomy which supports our wrist-based EMG approach. We then describe the different types of sensors which we use across the two studies in this paper.

BACKGROUND

Electromyography
Wearable devices have traditionally been worn on the wrist, normally considered the most ergonomic location for interactive wearables. When a muscle is contracted, an electrical potential difference is created by the electrically or neurologically activated muscle cells. Surface Electromyography (EMG) measures the difference with electrodes on the skin close to the muscles of interest, which can infer muscular activity. EMG data can be used to determine which muscles are active and even the amount of force they produce. Well-placed sensors are key to identifying patterns of EMG signals, which relate to specific movements of muscles. However, there are comparatively few muscle cells in the wrist compared to the proximal upper forearm (near the elbow), and the distal tendons (near the wrist) are more difficult to discriminate as they are more tightly packed. Thus there are increased challenges in performing EMG sensing of hand muscle movements by sensing at the wrist. As a result,

current off-the-shelf solutions such as the Myo armband typically capture EMG signals from muscles in the upper forearm [15].

Figure 1 shows a cross section of the wrist, the position for a typical wearable strap. There are three main groups of muscles which are responsible for the flexing and extending of the fingers. The flexor digitorum superficialis and flexor digitorum profundus (Fig 1, top) flex the fingers. The extensor digitorum communis, extensor indicis proprius, extensor medii proprius (Fig 1, bottom) extend the fingers, and aid a little to extend the hand. These muscles control all fingers except the thumb, the muscles responsible for this are known as the pollicis muscle group (Fig 1, left). As these muscles flex both the whole hand and the fingers it can be challenging to differentiate between gestures using EMG, for example flexion of the wrist and flexion of all fingers. The flexor carpi radialis, and the flexor carpi ulnaris (Fig 1, top left/right) control hand flexion at the wrist, and control hand abduction and adduction, respectively. Similarly, the extensor carpi ulnaris/radialis (Fig 1, bottom left/right) extend the hand at the wrist joint, and are also capable of abduction/adduction of the hand.

Commonly, bipolar electrodes are used to measure the electrical potential generated by the muscles in EMG systems. Normally, three electrodes are attached to the skin, two within close proximity of one another and another reference electrode to an area with less muscle activity. The signal can then be acquired by measuring the output of a differential amplifier, using the bipolar electrodes as input, removing any common noise that is measured by the reference electrode. Electrodes can be 'wet' (mounted onto the skin with adhesive and conductive gel) or 'dry' (without gel or adhesive). Recent designs for dry electrodes have comparable accuracy to wet electrodes, and are a more practical alternative [14].

Prototype EMG sensors can be designed in high density arrays worn on the forearm. These designs demonstrate excellent recognition rates for finger movements [3] and even for wrist movements [9][19]. However, these sensor arrays require a myriad of electrodes spread across a significant proportion of the arm's surface and may be more appropriate for integration into clothing than wearable devices. Finger muscle movement can also be captured using targeted wrist-mounted EMG sensors. For robustness these are commonly used in conjunction with additional EMG sensors mounted on the hand and/or distal forearm eg [6], making the overall configuration impractical for integration into a single wrist-mounted wearable. The use of EMG sensors in practical wearable scenarios typically also requires calibration in order to account for slipping watch straps and to align sensors with the anatomically optimal detection points. While we do not consider the issue of calibration in this paper, shift compensation algorithms [3] could further improve the results of our work.

Hand Gesture Sensing Techniques
Given the challenges of accurate EMG sensing, a number of other techniques have been applied to wearable gesture recognition. Inertial Measurement Units (IMUs) can provide

data to enhance gesture classification[19]. However, data collected from a wrist-mounted wearable can't be used to detect rotations about the wrist joint, because the hand can rotate separately to the forearm. Many gestures change meaning significantly based only on localised wrist rotation, for example the difference between pointing gestures in different directions, or the difference between a 'thumbs up' and a 'thumbs down'. Researchers have addressed this issue by placing accelerometers on the hand in rings or other hand-mounted wearables [10], but this requires additional hardware over a single wrist-worn wearable. Wrist-mounted IMUs can detect arm movement in a wider coordinate system, and can help to detect supination and pronation in the forearm but it is still difficult to differentiate between rotation of the whole arm and forearm, as the locally measured movements are the same.

GestureWrist [17] used a capacitance-sensing technique similar to localised EMG, to gather movement cues based on the shape of the wrist. The technique relies on providing a separate electrical source through the body which makes it difficult to combine with EMG sensing, as the local electrical source would drown out or interfere with the EMG signal.

Wrist-mounted computer vision approaches to gesture detection such as Digits [13] have recently become been integrated into wrist-mounted wearables and demonstrate excellent spatial accuracy. However, as with IMU approaches, the Digits system is currently restricted to the hand's coordinate system and so is unable to detect gestures which depend on wrist and forearm rotations. In general such computer vision techniques also suffer from intermittent occlusion which bounds the performance of the approach.

Finally, recent work has explored whether Force-Sensitive Resistors (FSRs) can provide a useful additional channel of data from a wrist-mounted wearable. FSRs have two copper traces that sandwich a special type of conductive polymer in between, which decreases in resistance as force is applied to it. By measuring the resistance of the resistor (e.g. by using a voltage divider), the amount of force applied can be inferred. Superficial tendons will move as hand gestures are performed, and because of their proximity to the surface of the skin, the movement can also recorded by FSRs to classify certain gestures. WristFlex [7] used an array of FSRs in a wearable wrist strap to detect finger pinch gestures based on the subtle tendon movements in the wrist. The authors found a high classification accuracy with low power consumption.

Gestures

The types of gestures that we wish to classify are shown in Figure 2. The coloured dots next to each gesture indicate the corresponding coloured muscles in Figure 1 that are predominantly responsible for producing the gesture, starting and finishing with the Palm hand posture. Most gestures use several different groups of muscles.

While there is no standardised hand gesture taxonomy, most gesture techniques described above evaluate against variations of general movements which draw on finger and wrist rotations. For our pilot study, we try to classify a set of ges-
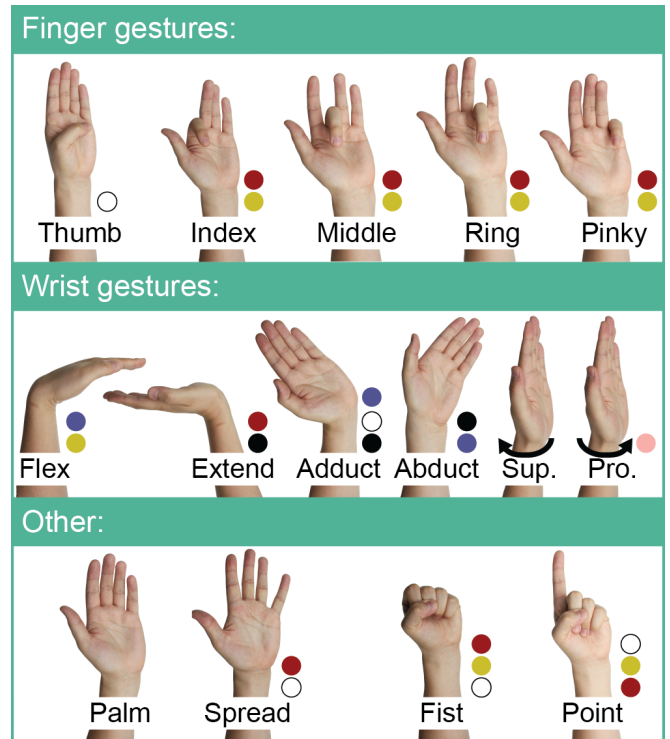

Figure 2: The set of gestures which we try to detect. The coloured dots represent the muscle groups shown in Figure 1 which are predominantly used during the gesture.

tures which contain several finger gestures, a fist and a palm gesture. Finger gestures and Wrist gestures are significantly different due to the different muscle groups which are required, illustrated by the colour variations between the gesture classes. For our follow-up study, we grouped and supplemented these gestures to explore a more challenging set of 15 gestures in three classes:

Finger Gestures These are gestures which only involve movements of the fingers. Anatomically, any movement which only involve rotations of the phalanges around their phalangeal joints will be classed as a finger gesture. In this class, we want to test whether individual fingers can be discriminated, and therefore include single finger flexions.

Wrist Gestures Hand movements which rotate the whole hand around the wrist joint are classified as wrist gestures. Although supination and pronation of the forearm occurs because of rotations at both the wrist and elbow joint, we include them as wrist gestures because they rotate the whole hand.

Other Gestures This set of gestures are not naturally focused only on single fingers or the wrist, consisting of whole-hand gestures that use multiple fingers simultaneously.

PILOT STUDY

Motivation
The aim of our pilot study was to compare the effectiveness of hand gesture recognition using EMG when using sensors located on the wrist in comparison to another device located on the proximal upper forearm. We were especially interested in how the relocation of sensors would affect the classification rates of finger gestures. We did not include wrist gestures in our pilot, as we expected to rely on existing studies that have measured wrist gesture performance with EMG [1]. The gesture set that we tested for included the following from Figure 2: Index, Middle, Ring, Fist, Palm. The palm gesture represents a relaxed state with no muscle activity.

Hardware
Hand-gesture recognition using electromyography is more challenging to implement around the wrist than if the sensors were located at the proximal end of the forearm, for the reasons explained previously, including smaller surface area, reduced muscle mass and the muscles being closer together.

Commonly in EMG, each muscle is individually measured. However, due to the constraints set by the surface area of the wrist and considering practically, we decided to use just two pairs of bipolar electrodes. The positioning of these electrodes is crucial, as they both need to have good coverage over different areas of the flexor digitorum muscles. The number of muscles in the forearm which we wish to detect are far greater than the number of data channels the device provides. The underlying principle for detecting which muscles are active in this situation is as follows.

Several muscles are measured by a single electrode. The sensitivity of an electrode to any particular muscle is proportional to it's distance away from the electrode's conductive centre. Since the values that are collected are roughly consistent each time a gesture is performed, a machine learning classifier can predict the gesture. The muscles of interest here are the flexor digitorum superficialis and profundus, and therefore we placed our 2 pairs of sensors orthogonally to these muscles, on the anterior of the forearm, so that each pair of sensors are more receptive to certain fingers. Figure 3 shows the difference in signals when flexing different fingers for two bipolar electrodes placed horizontally across the wrist, as per the placements in Figure 1. The graphs show that each sensor is indeed more receptive to a particular muscle, as confirmed by the differences in amplitude.

The reference electrode was placed to the side and on the posterior of the forearm, close to the ulna. Here there are fewer muscles and thus less muscle activity to be picked up by the electrode, making it the ideal for a wrist-mounted device to place a ground electrode. The electrodes are held in place with an elastic strap.

For this prototype, we created our own electrodes using conductive metal pins. It was necessary to apply gel to the electrodes for good electrical contact with the skin.

The electrodes are connected to circuit boards that apply a differential amplification, rectification and smoothing of the

(a) Index flexion EMG signals   (b) Ring flexion EMG signals

Figure 3: EMG signals from flexing a) the index finger and b) the ring finger

signals from the bipolar electrodes[18]. Each of the four boards are connected to an Arduino Uno. The Arduino is programmed to read analogue data from the sensors, which are then sent to the computer via serial communication for data collection and processing. The data is collected at approximately 60Hz.

Software
The data samples recorded for each channel from the device is of the form of a 1D time series. We extracted several time-domain features from each sensor in a given recorded sample. We used a support vector machine to classify the gesture data that we recorded. We chose to use an SVM over alternative classifiers because we knew our feature space would be small, and the kernel that the SVM uses will increase the dimensionality of the feature space. We also found that previous work on hand gesture recognition using EMG demonstrate that SVMs yield good performance [19][2]. We chose to use Libsvm, a simple yet efficient library for support vector machines[5]. After extracting the features, we normalised the values of each feature vector, since an SVM is not scale invariant. We used C-Support vector classification with a *radial basis function* kernel due to it's better classification performance over linear kernels[11], although a linear kernel would be faster if interactive feedback were required from a real-time wearable system with limited processing capabilities. The C-SVC classifier implemented by Libsvm uses a "one-versus-one" method for multi-class classification.

For the data collected with each device, we used the k-fold cross validation technique in order to find suitable parameters for the SVM classifier. Using these parameters to train the SVM on the training data, we could then use the classifier to predict each testing instance, and compare this to the training label to find out if the outcome was correct. The percentage accuracy of the classifier for each participant then equates to the number of correctly classified instances divided by the total number of instances.

Procedure
We used a within-subjects experimental design, with the independent variable being the placement of the electrodes, and dependent variable being the accuracy of the SVM classifier. We kept the feature set and the supervised learning implementation the same in each iteration of the study.

12 participants (9 male, 3 female) participated in the study, all participants were healthy with no known muscle impairment. The study consisted of a 30 minute session, during which the participant would perform gestures while wearing one device worn at the wrist, and then again with the second device on the upper forearm. During each of the 5 gestures, the device would record the data from each bipolar differential channel. The gesture would be repeated multiple times over the course of the session.

Results

A Shapiro-Wilk test showed that the classification sample was not normally distributed (W=0.731, p <0.05), so a non-parametric Wilcoxon Signed Ranks Test was applied to compare the devices' classification accuracy. This test indicated that the classifier was significantly more accurate for the data collected from the device worn on the wrist (mean 82% correctly classified) than for the device located on the upper forearm (mean 53%), Z = -2.359, p <0.01.

Discussion

Our pilot study demonstrated a significant increase in classification accuracy when the electrodes are worn on the wrist compared to the proximal forearm location. While we hoped that the new placement might prove comparable to the usual upper forearm placement for finger gesture detection, we were surprised that the difference was so significantly in favour of the wrist, which we believed should be more anatomically difficult to discriminate the EMG signals.

This result led us to explore other potential reasons for the increased performance of the classifier in the wrist condition. Looking at the differences in the system design, one key area was the home-made electrodes used under the wrist band. During some sessions, we found that the resting amplitude of the EMG sensor data varied within the same session. One reason for this could be a change in electrical contact with the skin. This change can be caused by a displacement of the electrodes, due to movements of the forearm. Upon further experimentation with the device, we came to believe that the increased performance could be attributed to changes in pressure on the EMG electrodes from the elastic wrist strap. The pressure was modulating the EMG signal in ways which we suspected may be providing additional gestural features to the classifier than pure EMG signals alone.

The hypothesis that wrist pressure was providing predictable features by modulating our EMG data was rather surprising. Existing studies in the literature directly using wrist pressure to detect hand gestures had selected specific movements to ensure classification accuracy, such as the use of a finger pinch in the WristFlex study [7]. Nonetheless, we hypothesised the wrist pressure was the most likely variable in the higher classification rate. In order to test this hypothesis, we conducted a second user study in which we isolated the pressure and EMG data collected only on the wrist, in order to quantify the effect of the pressure changes in the wrist on the gesture classification results.

Figure 4: An Image of the prototype worn around the wrist.

EXPERIMENT

This section describes the design process of the device for our main study. The following section describes the study and results.

Hardware

The changes in pressure applied to each electrode from the strap are present for a number of reasons. The key explanation is that muscles/tendons in the forearm become displaced upon contraction. Stretching of the skin can also affect the pressure between the sensor and strap. Factors of the strap design such as the elasticity, can significantly change how the movement of the skin affects the pressure, as we found out during trial and error of different strap types. Eventually we decided that a simple elastic band would suffice.

In order to measure the pressure between the strap and the wrist, we chose to use the same Force Sensitive Resistors that are used in the WristFlex prototype [7]. The FSR400 component [8] provides good performance in their study and our requirements are similar: inexpensive, small and highly sensitive.

In order to isolate the pressure data from the EMG readings, we chose to use a more robust EMG sensor (SKINTACT Ag/AgCl aqua-wet electrode, ref: FS-TF). 'Wet' ECG electrodes that are commonly used for medical purposes have a design which mitigates the effect of pressure, due to the use of highly porous foam beneath the sensor. The electrodes are connected to circuit boards via shielded cables with snap connectors. The circuit boards apply differential amplification, rectification and smoothing of the signals from the bipolar electrodes[18]. For each pair of bipolar electrodes this provides a single channel of EMG data. Each of the four signal processing boards are connected to an Arduino Uno.

The prototype for our second study extends the capabilities of the device used previously. We kept the same configuration of 2 bipolar EMG electrodes located on the anterior side of the forearm, because of the reasonable accuracy for single finger gesture recognition that our first study showed is attainable.

The increased gesture set which we aim to identify in this study includes gestures other than finger gestures. Some gestures require the use of the muscles located in the posterior compartment of the forearm (Figure 1, bottom), as indicated by the coloured dots in Figure 2. Chiefly, these are the Extend, Adduction, Abduction, Spread and Point gestures. Therefore to detect the activity of these muscles, we decided

Figure 5: Diagram of the hardware components used in the prototype.

to place two additional bipolar sensors close to the extensor muscles on the posterior of the forearm. This has the added effect of detecting muscle activity from the extensor digitorums, thus facilitating the classification of Finger gestures. The muscles that are active upon wrist Flex, Extend, Adduct and Abduct are the flexor/extensor **carpi** radialis/ulnaris.

The new EMG sensors we have chosen are unnecessarily large (excessive amounts of adhesive) and could be engineered at a fraction of their size with no loss of functionality. For our purposes, this limited the amount of space left to allocate to FSR sensors. We therefore chose to put the reference electrode on the other hand. In an ideal situation where the sensors are smaller, it should be possible to have the reference electrode placed close to the ulna as in our pilot study, without overlapping with any of the other sensors.

The force sensors we have used are much smaller than the EMG sensors. This allows us to easily place them on the sides, close to the tendons of the muscles which control thumb movement and hand abduction/adduction. We were also able to fit two more FSRs in between the two EMG sensors on the anterior and posterior of the wrist (Figure 1. These latter two are conveniently placed onto the tendons of interest for finger gestures: tendons of the flexor and extensor digitorum muscles. The sensors are spaced around an elastic band which is worn around the wrist. Our design supports adjustment of the sensors' placement on the band, to account for differently sized wrists. The elastic is required so that there is slight pressure exerted onto every force sensor, so that when the shape of the wrist changes, there is a change in pressure. In the absence of such an external force, the sensor would simply move with the surface of the skin, and there would be minimal change in signal. The force sensors are connected to a small circuit that consists of a digital multiplexer and a voltage divider which work in tandem. A resistor of $180k\Omega$ gave

us values within a suitable range. The output signal from the circuit is then connected to the Arduino's analogue pin.

The approximate positions of all the sensors are shown in Figure 1 and Figure 4. In total there are eight electrodes worn on the wrist, and there are also reference electrodes for each pair of bipolar electrodes attached to the other arm. Our aim is to be able to detect finger and wrist gestures with as few electrodes as possible while still maintaining a reasonable accuracy, as increasing the number of electrodes makes a device more impractical due to the size, power, computational load and cost.

The diagram in Figure 5 shows how each of the hardware components are connected. In the figure, only one of four sets of EMG components are shown.

Software

The data collection process that we use remains similar to that of the pilot study. Instead of receiving 2 data channels from the Arduino, we now have 4 FSR and 4 EMG channels. Features for EMG could theoretically be extracted from both the frequency and time domain. However, due to the limitations of our hardware, we could not collect sensor data at a high enough frequency for there to be any significant information in the frequency domain. Useful frequency information would have to sample data at a rate several orders of magnitude higher than our current rate of 60Hz [1][16]. We extracted the following features for each 1D time signal (2.4s):

- Root mean square (RMS) - This feature is correlated with the signal energy, and thus muscle activity in the case of EMG. This has been found to be a good feature for machine learning with EMG data[3]. Significant changes in signal energy also occur in the FSR data, as seen in the sample data shown in Figure 6. The muscles in the forearm have different sizes, and the larger they are, the higher the amplitude of the received EMG signals. Large muscles, such as the flexor carpi ulnaris, can influence several sensors due to their size. These properties facilitate the process of classification, as the muscles which could have produced the electric potential difference can be inferred.

- Standard deviation (SD) - Also correlated with the muscle activity, although this feature is invariant to amplitude offsets.

- Peak amplitude - Measures the maximum value of the sensor data. This is chosen to take into account the shape of the signal, as two signals with the same RMS could look entirely different.

We tested the accuracy of our classifier by using a leave-one-out 10-fold cross validation on our data sets. For each iteration, we use a stratified shuffle split on the training portion, and then used a grid search algorithm for selecting the hyper parameters, C and γ. Once suitable parameters were found, we then trained the SVM and tested the classification on the test fold. Using this method, we remove any bias from the parameter selection that would have otherwise occurred if the parameters were instead chosen using a grid search on the entire data set.

Figure 6: Graphs of electromyographic and pressure signals for select gestures. The horizontal axes show time, and the vertical axes show amplitude.

The time taken to read the data file, compute the features and classify the instance using the SVM, is less than 5ms on a desktop class Core i7 Intel processor , across all gestures and sensors. This is 3 times less than the sampling rate, which suggests the feasibility of real time classification using our system entirely possible, as further optimisations would reduce the time even more.

Our software suite also included an application that would show the participant a sequence of videos of someone performing certain gestures that they would then mimic. While the participant is performing the gesture, a second Python script would record 2.4 seconds of sensor data that was sent to the computer's serial port from the Arduino. It is then stored as comma separated value files for later analysis.

Participants
A total of 12 participants (different from the first study) took part in our experiment. There were 3 female, and 9 male participants. The average age across all participants is 33 years, with a standard deviation of 9.1 years. The circumferences of their wrists averaged 16cm, with a standard deviation of 1.3cm.

Procedure
An experiment began with the participant being asked to wear the device around their left wrist. We first adhered the electrodes for the electromyography onto the participants arm. The band with the force sensors was then placed around the wrist, making sure that each force sensor had good contact with the skin, with slight pressure on each of them to ensure the output signal was in a suitable range.

We then instructed them to mimic the gestures that were shown in a video clip on a screen in front of them. They were told to keep the timing of their hand movements in sync with the one in the video clip, as closely as possible. Each video clip is 4 seconds, and the gesture in every clip starts

and ends at the same point in time. This is to ensure that the gestures are performed consistently throughout. Each gesture video clip is shown 6 consecutive times, the first clip is to let the participant acknowledge that a new gesture has started, and so that they can practice it once. The data is not recorded during this period, only for the 5 subsequent gestures. The participants perform 15 different hand gestures in this manner, these gestures are all those shown in Figure 2. This process is repeated once more, to give a total of 10 data points per gesture. For each participant's data set, we tested 10-fold cross validation accuracies using a separate SVM to train each individual's gesture data.

In our study, we used three sensor conditions (EMG, FSR, Both) and three gesture sets (Fingers only, Wrist only, All), where the 'All' category combines the first two gesture classes with the 'Other' gestures shown in Figure 2. Since we vary both of these variables simultaneously, data is collected across a total of 9 different experimental conditions.

To re-iterate our hypothesis: we expected that incorporating pressure data increases the classification rate of gestures compared with using EMG alone on the wrist.

Results
Figure 7 shows the mean results for each experimental condition. For the full set of 15 gestures, average 10-fold cross validation classification rate for both EMG and FSR data across all participants is 95.8%. This gives us a classification rate for our overall device with respect to our gesture set. We also compared the Sensor and Gesture conditions to identify main effect and to identify any interaction between them.

A two-way repeated measures ANOVA indicated that the main effect of Gesture was significant, $F_{(2,10)} = 5.67$, $p < 0.001$. Post hoc analysis with Bonferroni correction accounting for multiple comparisons showed that Wrist Gestures were classified significantly better than All Gestures (p

Figure 7: Cross validation accuracies for each test condition, with standard error bars.

Figure 8: Confusion matrix across all participants, using both sensors and all gestures.

<0.05), but not significantly different to Finger Gestures (although approaching significance, p=0.06). The overall difference between Finger Gestures and All Gestures was also not significant.

The main effect of Sensor was also significant, $F_{(2,10)} = 17.70$, $p < 0.001$. Post hoc analysis, again with Bonferroni correction, indicated that Both sensors were significantly better than either EMG ($p < 0.001$) or FSR ($p < 0.001$), but that these did not significantly differ from one another.

The interaction effect between Gesture and Sensor was also significant $F_{(4,8)} = 11.69$, $p < 0.001$. The distributions of Gesture and Sensor classification rates suggest that this effect was caused by EMG data correctly classifying more Finger Gestures (90.5%) than Wrist Gestures (85.3%), while FSR data correctly classified more Wrist Gestures (96.1%) than Finger Gestures (86.5%). The EMG and FSR techniques in combination showed significant complementarity when classifying All Gestures, each technique alone providing fewer correct classifications (EMG 86.6%, FSR 89.4%) while together they classified All Gestures correctly to an accuracy of 95.8%.

The confusion matrix in Figure 8 shows the performance of the classifier for the case with Both sensor types and All gestures. There is no obvious misclassification for any particular gesture, although the main confusion appears in the finger gestures. The ring gesture has the highest number of false positives, wrongly classifying the middle, pinky, and palm gestures. Similarly, the SVM classified several index and pinky gestures as middle finger gestures, and fist as point gestures. This error could be attributed to a few possibilities.
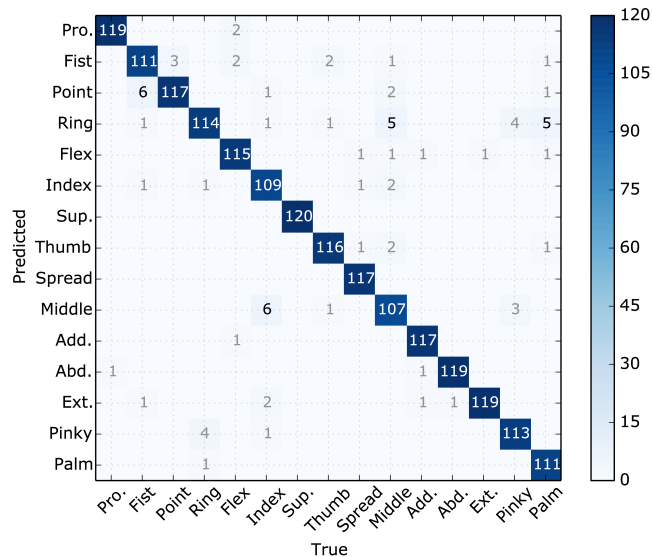
The most likely of them however, is that there is an insufficient number of sensors for the flexor/extensor muscles that control the fingers. The Finger gestures being the most incorrectly classified gestures can also be confirmed by the fact that the average cross validation result is worse than that of Wrist gestures, and also All gestures. The thumb gesture appears to have fewer false positives and false negatives, and this is likely due to do the fact that the thumb gesture uses muscles which are different from those the other finger gestures use.

We found the standard deviation to be much larger for Finger gestures (3.72%) than Wrist gestures (1.67%) and All gestures (1.98%), when all sensors are used. When the EMG sensors are only used, the deviation is 6.22% for the Finger gesture set. This is a much larger variance than Wrist gesture classification using only FSR (2.05%). One possible reason for this could be the variance in hand dexterity of participants, due to the somewhat difficult control of individual finger flexing; the flexor digitorum superficialis is principally responsible for this.

In theory, larger wrist sizes should increase the recognition rates of the classifier, since the muscles are more interspersed, and the larger muscle mass should increase the difference in electrical potential. Though our data does not suggest such a correlation between the circumference of the subjects wrists and the classification accuracies, our sample has too limited variance to determine this.

DISCUSSION

Our principal result is the accurate classification, just below 96%, demonstrated by the combined EMG and FSR data across our gesture set. This indicates that the EMPress technique is viable for hand gesture recognition, and that this approach is significantly better than either sensor type on its own.
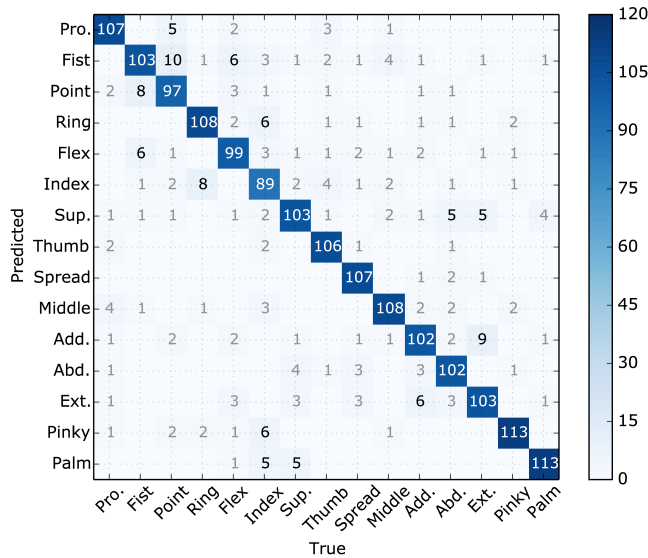
Figure 9: Confusion matrix across all participants, using only EMG sensors and all gestures.
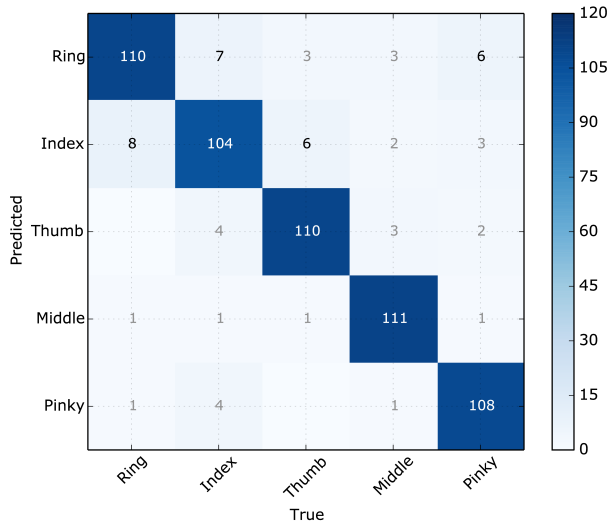

Figure 10: Confusion matrix across all participants, using purely EMG sensors for classifying finger gestures.


Figure 11: Confusion matrix across all participants, using purely FSR sensors for classifying finger gestures.

It is not possible to directly compare EMPress to other published gesture sensing methods, because gesture sets, number and type of sensors used, and anatomical position vary widely across the literature. Our overall classification rate for Finger gestures is 94.0%, while the rate for Wrist gestures is 97.5%. Confusion matrices show that the weakness in our system is generally due to the mis-classifications of adjacent fingers within the Finger gesture set. We suspect that this is partly due to the low electric potential generated by the few muscle cells that are present in the wrist. Sensing EMG spikes predictably at the wrist against low signal-noise ratios is the most challenging aspect of the EMPress technique. Nonetheless, our data shows that even with these challenges that EMG outperforms FSR when classifying Finger gestures, and its inclusion therefore remains an important component of our technique. Uniquely, we have also demonstrated that this high level of classification accuracy is possible without resorting to distributed arrays of EMG sensors across the forearm, instead localising the sensors to a wrist-mounted device. Following [3], even if constrained to spare real estate on the existing wrist strap, we expect that further EMG sensors should increase the classification rate.

While we expected pressure sensing to significantly supplement EMG gesture classification accuracy, our study demonstrates that pressure not only contributes to EMG classification of gestures, but can actually parallel or even beat EMG performance in complementary ways. The study data reveals that the overall FSR-only classification success rate is 89.4%, showing a strong predictability of wrist pressure over a wide range of gestures. While we expected pressure sensing to perform well under conditions of wrist movement (96.1%), perhaps most surprising is that pressure sensing alone was able to detect finger gestures to a classification accuracy of 85.3%. While WristFlex [7] already demonstrated a high predictability for a particular anatomically targeted pinch gesture using FSR sensing at the wrist, we believe the EMPress technique

is the first time such an approach has been shown to be potentially strong for a varied range of gestures without special anatomical targeting of the sensors, and particularly for combined wrist and finger movement classification.

Finally, our study shows that the EMG and FSR techniques are strongly complementary, with EMG significantly better for detecting finger gestures and FSR significantly better for Wrist gestures. We therefore propose that, for wrist-based wearables, the techniques are used in conjunction so that a high recognition rate can be achieved while locating sensors only at the wrist. Note that our device consists of wet electrodes with cross-arm references, this suggests that our findings apply under optimal EMG conditions. It is not necessarily the case that this complementarity will yield true for a system with dry electrodes with a same-arm reference electrode placed over the ulna, and further work will be required to verify this level of performance against dry and locally grounded EMG such as our pilot study used.

FUTURE WORK

We anticipate exciting future work in the design of the electrodes deployed in EMPress-enabled wearables.

Firstly, it may be possible to integrate the different sensing types. Having separated EMG and pressure sensing in order to separate their contribution to the overall effect, it is now possible to consider engineering a design which re-integrates EMG and pressure sensing. Designs would be possible which directly exploit the pressure-modulated EMG signal we observed, alongside a comprehensive exploration of optimal feature sets of the signal for classification. It would also be possible to retain the software design for our second study but integrate hardware within the EMG electrode itself, supporting pressure sensing while greatly increasing the space available on the wearable strap for a larger array of sensors around the wrist.

Moreover, with this approach the FSR component could then directly measure the pressure exerted onto the EMG sensor. In our final prototype, even with foam underneath the electrode, there is still a slight change in electrical contact as the foam is compressed. Dry electrode sensors such as those used in the Myo will also suffer from pressure changes due to the movement of the wrist strap. The use of conductive foam[4] may mitigate but not remove the effect of pressure. However, an integrated EMPress sensor would be able to measure the pressure exerted onto the EMG contact from the band. This would allow a predictive model of the modulating effect of pressure on the EMG signal to underpin an estimate of the 'true' EMG signal.

Our gesture set allows us to differentiate between Finger gestures and Wrist gestures. However, we have not tried to classify nuanced combinations of both Finger and Wrist gestures, nor have we attempted to combine multiple Finger gestures simultaneously. Detecting multiple Finger gestures with a reasonable accuracy will probably demand higher density of EMG sensors placed precisely above muscles of interest as described above. Since the sensors seem to be suited to recognising separate components of gestures, calibration al-

gorithms to map sensor-anatomy offsets [3] will be important. There is also the potential to biologically tailor the device so that additional pressure sensing can be targeted towards regions where there are limited surface EMG signals owing to wrist anatomy.

Finally, given existing wearables typically already include an IMU, our technique could potentially draw on acceleration data to detect many more gestures and/or with further increases in accuracy. In particular, gestures which require movements that stem from the elbow, arm, or whole body would be well-supported by such data, and we may observe some further improvement to Wrist gesture classification accuracy as well.

CONCLUSION

Our initial study planned to compare the effect of EMG sensor placement on the forearm. We found that placement of EMG around the wrist has comparable accuracy to when the sensors are located on the upper forearm. This unexpected outcome led us to believe that pressure exerted from the wristband which held the EMG sensors to the skin, modulated the signal in a semi-predictable manner.

Our main study confirmed our hypothesis that including the pressure around the wrist does indeed increase the classification rate for different kinds of gestures involving both finger and wrist movements. Furthermore, we found significant complementarity between the two types of sensors: the pressure sensing surpasses EMG for classifying Wrist gestures, and the reverse is true for Finger gestures. We believe this technique is sufficiently accurate and ergonomically practical to have significant potential for underpinning a new generation of wearable gesture sensing technology.

ACKNOWLEDGEMENTS

REFERENCES
1. S. A. Ahmad and P. H. Chappell. 2009. Surface EMG pattern analysis of the wrist muscles at different speeds of contraction. *Journal of Medical Engineering & Technology* 33, 5 (2009), 376–385. DOI: **http://dx.doi.org/10.1080/03091900802491246**

2. Ahmet Alkan and Mücahid Günay. 2012. Identification of EMG Signals Using Discriminant Analysis and SVM Classifier. *Expert Syst. Appl.* 39, 1 (Jan. 2012), 44–47. DOI: **http://dx.doi.org/10.1016/j.eswa.2011.06.043**

3. Christoph Amma, Thomas Krings, Jonas Böer, and Tanja Schultz. 2015. Advancing Muscle-Computer Interfaces with High-Density Electromyography. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 929–938. DOI: **http://dx.doi.org/10.1145/2702123.2702501**

4. Hyun Jae Baek, Hong Ji Lee, Yong Gyu Lim, and Kwang Suk Park. 2012. Conductive Polymer Foam

Surface Improves the Performance of a Capacitive EEG Electrode. *Biomedical Engineering, IEEE Transactions on* 59, 12 (Dec 2012), 3422–3431. DOI:
`http://dx.doi.org/10.1109/TBME.2012.2215032`

5. Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* 2 (2011), 27:1–27:27. Issue 3. Software available at `http://www.csie.ntu.edu.tw/~cjlin/libsvm`.

6. Xiang Chen, Xu Zhang, Zhang-Yan Zhao, Ji-Hai Yang, V. Lantz, and Kong-Qiao Wang. 2007. Hand Gesture Recognition Research Based on Surface EMG Sensors and 2D-accelerometers. In *Wearable Computers, 2007 11th IEEE International Symposium on*. 11–14. DOI:
`http://dx.doi.org/10.1109/ISWC.2007.4373769`

7. Artem Dementyev and Joseph A. Paradiso. 2014. WristFlex: Low-power Gesture Input with Wrist-worn Pressure Sensors. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*. ACM, New York, NY, USA, 161–166. DOI:
`http://dx.doi.org/10.1145/2642918.2647396`

8. Interlink Electronics. 2015. FSR400 short. (2015). `http://www.interlinkelectronics.com/FSR400short.php`.

9. Marco Gazzoni, Nicol Celadon, Davide Mastrapasqua, Marco Paleari, Valentina Margaria, and Paolo Ariano. 2014. Quantifying Forearm Muscle Activity during Wrist and Finger Movements by Means of Multi-Channel Electromyography. *PLoS ONE* 9, 10 (10 2014), e109943. DOI:
`http://dx.doi.org/10.1371/journal.pone.0109943`

10. Jeremy Gummeson, Bodhi Priyanth, and Jie Liu. 2014. An Energy Harvesting Wearable Ring Platform for Gesture Input on Surfaces. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '14)*. ACM, New York, NY, USA, 162–175. DOI:
`http://dx.doi.org/10.1145/2594368.2594389`

11. S. Sathiya Keerthi and Chih-Jen Lin. 2003. Asymptotic Behaviors of Support Vector Machines with Gaussian Kernel. *Neural Comput.* 15, 7 (July 2003), 1667–1689. DOI:
`http://dx.doi.org/10.1162/089976603321891855`

12. Frederic Kerber, Pascal Lessel, and Antonio Krüger. 2015. Same-side Hand Interactions with Arm-placed Devices Using EMG. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '15)*. ACM, New York, NY, USA, 1367–1372. DOI:
`http://dx.doi.org/10.1145/2702613.2732895`

13. David Kim, Otmar Hilliges, Shahram Izadi, Alex D. Butler, Jiawen Chen, Iason Oikonomidis, and Patrick Olivier. 2012. Digits: Freehand 3D Interactions Anywhere Using a Wrist-worn Gloveless Sensor. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*. ACM, New York, NY, USA, 167–176. DOI:
`http://dx.doi.org/10.1145/2380116.2380139`

14. Amanda C. Myers, He Huang, and Yong Zhu. 2015. Wearable silver nanowire dry electrodes for electrophysiological sensing. *RSC Adv.* 5 (2015), 11627–11632. Issue 15. DOI:
`http://dx.doi.org/10.1039/C4RA15101A`

15. Myo. 2015. How do you wear the Myo armband? (2015). `https://support.getmyo.com/hc/en-us/articles/201169525-How-do-you-wear-the-Myo-armband-`.

16. Angkoon Phinyomark, Pornchai Phukpattaranont, and Chusak Limsakul. 2012. Feature reduction and selection for {EMG} signal classification. *Expert Systems with Applications* 39, 8 (2012), 7420 – 7431. DOI:
`http://dx.doi.org/10.1016/j.eswa.2012.01.102`

17. Jun Rekimoto. 2001. GestureWrist and GesturePad: Unobtrusive Wearable Interaction Devices. In *Proceedings of the 5th IEEE International Symposium on Wearable Computers (ISWC '01)*. IEEE Computer Society, Washington, DC, USA, 21–. `http://dl.acm.org/citation.cfm?id=580581.856565`

18. Advancer Technologies. 2015. Muscle sensor v3. (2015). `http://www.advancertechnologies.com/p/muscle-sensor-v3.html`.

19. M.T. Wolf, C. Assad, A. Stoica, Kisung You, H. Jethani, M.T. Vernacchia, J. Fromm, and Y. Iwashita. 2013. Decoding static and dynamic arm and hand gestures from the JPL BioSleeve. In *Aerospace Conference, 2013 IEEE*. 1–9. DOI:
`http://dx.doi.org/10.1109/AERO.2013.6497171`