EchoFlex: Hand Gesture Recognition using Ultrasound Imaging

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

Carol Phillips
Radiology Department
University Hospitals Bristol NHS Foundation
Trust, Bristol, UK
Carol.Phillips@UHBristol.nhs.uk

ABSTRACT
Recent improvements in ultrasound imaging enable new opportunities for hand pose detection using wearable devices. Ultrasound imaging has remained under-explored in the HCI community despite being non-invasive, harmless and capable of imaging internal body parts, with applications including smart-watch interaction, prosthesis control and instrument tuition. In this paper, we compare the performance of different forearm mounting positions for a wearable ultrasonographic device. Location plays a fundamental role in ergonomics and performance since the anatomical features differ among positions. We also investigate the performance decrease due to cross-session position shifts and develop a technique to compensate for this misalignment. Our gesture recognition algorithm combines image processing and neural networks to classify the flexion and extension of 10 discrete hand gestures with an accuracy above 98%. Furthermore, this approach can continuously track individual digit flexion with less than 5% NRMSE, and also differentiate between digit flexion at different joints.

Author Keywords
Gesture recognition; interactive ultrasound imaging; machine learning; computer vision

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

INTRODUCTION
Our hands are capable of skillful and intricate movements with multiple degrees of freedom, giving us the capacity to manipulate objects and enhance or replace verbal communication through the use of gestures. Gestures play a fundamental role in emerging Human-Computer Interaction (HCI) paradigms, and several wearables have been developed to capture gestures as a way of controlling devices naturally. However, many of these existing technologies and techniques to detect gestures have limitations which make them unsuitable for general applications.

These limitations include inaccurate or variable performance, ergonomic challenges of wearing the devices, and lack of robustness during regular use. Here, we explore the capabilities of ultrasound imaging as a gesture detection mechanism. Ultrasound imaging, or UltraSonography (US), has long been available as a medical technique for safe anatomical inspection for example to see developing fetuses or heart disease. While progress in the medical area has been significant, the use of ultrasound for interactive applications has not been explored within the HCI community.

Ultrasonography allows one to see the musculature within the body in real time. By observing the movement of the muscles, it is possible to infer hand movements. It is unlikely that this technique provides more accuracy than systems which directly measure the moving body parts, such as cameras or data gloves, but US has potential benefits over these existing wearable hand tracking techniques. Directly imaging the muscles is not subject to occlusion, unlike imaging the body parts externally which can be obscured by other body parts or be out of view depending on the mounting of the camera. In contrast to wearing instrumentation on the hand, imaging the muscles leaves the hands free to perform actions unencumbered by devices or sensors. This could be beneficial for sensitive or expert manipulation as in the case of fragile objects, surgical operations, or simply to keep the physical naturalness of a handshake.

Surface electromyography (EMG) captures the electric signals from the muscles with sensors on the skin. EMG is one of the main alternatives for non-invasive hand pose detection, but suffers from cross-talk. That is, it is difficult to differentiate between individual muscles particularly those which lie deeper in the forearm [35]. Techniques that measure wrist shape changes have been shown to work well for small numbers of wrist gestures but there is limited evidence that they are robust to a wide variety of hand movements.

Ultrasound imaging requires a small area of contact with the arm, whereas EMG electrodes are typically placed around the forearm at specific locations selected with a calibration procedure. With ultrasound, the probe that is used to obtain the images is considerably smaller in size, and in contrast with EMG, lends itself well to be integrated with smaller form factors. We discuss this and other wearable hand tracking tech-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI 2017, May 06-11, 2017, Denver, CO, USA.
© 2017 ACM. ISBN 978-1-4503-4655-9/17/05 ...$15.00.
DOI: http://dx.doi.org/10.1145/3025453.3025807

niques in related work below and suggest reasons why US is beneficial over existing methods.

In this paper, we analyse the performance of ultrasound to recognise discrete gestures, measure continuous angles of digit flexion and discriminate between different types of digit flexion. Recent work on US for gesture detection has shown very high accuracy. We explore issues of this technology in the context of HCI, and for the first time study the effects of probe location on accuracy, and find that ergonomic mounting positions are viable without significant variation in accuracy. We also provide a simple calibration mechanism to mitigate the effects of cross-session position shift, which is important for day-to-day use. We provide guidelines for the best locations depending on the gesture recognition application. The high recognition accuracy could make several applications possible, including robotic hand control, natural interactions and enhanced sports and musical instrument tuition. Given the trend in reduction of size and cost of ultrasonic imaging devices, we expect to see this approach applied for more interactive scenarios.

RELATED WORK

In this section we describe the different sensing technologies for hand gesture recognition, focusing on non-invasive wearable techniques, and review previous work on ultrasound to support interactive applications.

Optical

Optical methods are based on employing a camera to determine 2D or 3D hand position. The camera can be sensitive to the visible spectrum but alternatives exist in the infra-red or below. A review of these methods can be found in Rautaray et al. work [31]. An example is Digits [23] which uses an IR camera attached to the wrist to sense the distance to the fingers. The sensor is worn on the wrist and as with all optical approaches requires direct line of sight to the fingers, and suffers from occlusions of the fingers by the hand.

Wrist Deformation

When we flex our fingers and wrist the muscles inside our forearm change in size and shift around, causing a deformation of the wrist shape. Measuring this deformation is an indirect method to infer the hand gesture which caused this change in shape. One way of sensing the shape is with pressure sensors, WristFlex [10] uses 15 pressure sensors around the wrist and a Support Vector Machine (SVM) to classify five pinch gestures with an average accuracy of 80%.

Another method for measuring the distance between two parts of a wearable is to use capacitive sensing. The capacitance between two parallel plates is proportional to the distance and material between them. This technique was used by GestureWrist [32] to measure the deformation of the wrist and qualitatively detect some common gestures. Cheng et al. showed that this method can operate with accuracies of 77% for 36 gestures [9]. Electrical impedance tomography is a similar approach that has shown to be more effective in recent work by Zhang et al. [43], reaching accuracies above 94% for a variety of gestures. This technique is close to ultrasound imaging, the key difference is the images formed with IET lack detail of the muscles, meaning they cannot be tracked, but morphological changes of musculature is enough to classify discrete gestures with.

Another way of detecting the wrist shape is with an infra-red distance sensor [17]. Qualitative results suggest that it would be possible to differentiate between 10 gestures.

From these recent findings it is clear that improvements are needed for this to be a sufficiently accurate technique for many applications.

EMG: Electromyography

Electromyography consists of measuring the bioelectric signals produced by the muscles upon contraction. These signals can be measured with surface electrodes on the skin in a non-invasive way.

Tenore et al. [37] showed that with 32 electrodes around the forearm it is possible to detect individual flexion and extension movements of each fingers with an accuracy greater than 90% both in amputees and able-bodied subjects. An interesting and promising result is that detection accuracy between transradial amputees and able-bodied subjects showed no statistical difference. Castellini [7] showed that 10 electrodes is enough to differentiate 4 types of grasps as well as the force applied with an accuracy ranging from 75% to 90%. You et al. used 4 electrodes to detect flexion each of finger with 97.75% accuracy [42].

NASA developed a wearable band for astronaut suits made of 17 electrodes and accelerometers [39]. This wearable represents a way of enhancing communications with other astronauts or vehicles since the accuracy of gesture detection was greater than 96%.

Recent improvements came from using a dense array of 192 EMG electrodes to detect up to 27 gestures with 90% accuracy [2].

EMG tends to suffer from a trade-off between accuracy and the surface area used for measurements. Whilst EMG proves to be highly effective for a high number of electrodes, it performs poorly under placement and surface area restrictions, making it awkward to fit into existing wearable device form factors such as watch straps.

Smith at al. demonstrated that EMG can be used to detect continuous angle of the fingers [36] with an NRMSE of 11%. However, the nature of the biosignal creates particular difficulties regarding continuous angle detection, and so there have been very few attempts to challenge this task using EMG, and is therefore currently more suited to discrete gesture recognition.

Mechanomyography

Mechanomyography is a technique which records vibrations of the muscle fibres after contraction, which oscillate at their resonant frequencies and is characteristic of a specific muscle's activity [28]. Mechanomyography is a useful technique to support prosthesis and switch control since the placement of the sensors does not have to be precise and the change in the skin impedance due to sweating does not affect the performance [21]. Although this technique uses sound, it is not an imaging technique and has been under-explored within the HCI community (with the exception of [40]) despite its simplicity in hardware and robustness to sensor placement or skin condition.

Ultrasound

Sonomyography is the use of ultrasound imaging to create 1-dimensional scans (A-scans). An early study by Hodges et al. [19] proved that muscle contraction of the tibia, biceps and abdomen can be measured with an ultrasound sensor. Later, Zheng et al. [44] showed that it is also possible to infer the wrist angle with an mean error of 7.2%. They defined Sonomyography as measuring the dimensional change of muscles with ultrasound sensors.

Sonomyography has the advantage of using simple hardware because it only needs one transducer element per A-scan; current ultrasonic probes have more than 128 elements and are becoming the standard for ultrasound imaging. Sonomyography can also be used to continuously detect the opening and closing of the hand with an RMS of 12.8% [8] and recently to detect 6 common gestures with an accuracy of 74% [18]. The latter study used a wearable device and a bracelet with five transducers. Comparisons between EMG and sonomyography were favourable for the latter in detecting continuous wrist flexion [15] and hand grasping force [16].

Sonomyography is an appropriate method to detect a reduced set of gestures, but provides limited information and lacks the possibilities offered by modern phased arrays of ultrasound transducers for more detailed imaging.

Mujibiya et al. presented a device that is capable of detecting arm grasps and other on-body touch interactions, using a band of ultrasonic transducers around the forearm and fingers [27]. This device has the advantage of working at much lower frequencies and with fewer transducers than imaging devices, but the drawback is that it can only detect interactions when one hand is touching the arm of the other. Detecting hand movements involving only one hand is a requirement for certain situations [22].

Ultrasound Imaging

Ultrasound imaging has been used previously to infer the position of the foot by tracking the muscle displacement [26]. An image of the muscles and tendons was taken and the insertions of the tendons into the muscle were used as markers to track. A comparison between the estimated position and the ground truth (metallic plate inserted in the tendon) showed errors of less than 10 micrometers or the equivalent of 0.7 degrees on the ankle rotation. Markers were tracked using cross-correlation. Later, this technique was used to track the tendon in the wrist with an error of 80 micrometers [25].

Castellini et al. estimated continuously changing finger angles by analyzing ultrasonic images of the forearm [5] [6], obtaining an NRMSE of approximately 2%. This technique was also extended to determine the force that the fingers were exerting [14]. They captured ultrasound images at the wrist and divided them into regions. Then, the features from each region were used in a linear regressor to establish a relationship between the features and the angle or force of each finger. In more recent work, it was shown that it is possible to recognize 4 different grasps with 80% accuracy, and also 3 levels of strength with 60% accuracy [29].

Sikdar et al. focused their work on recognising discrete hand gestures by imaging the muscles mid-forearm. They divided the image and calculated the average brightness change per region to create different activity patterns for each gesture. A Nearest Neighbour search was able to classify input gestures correctly with 98% accuracy for the flexion of 4 fingers [35] and 91% for 15 gestures [1]. Another group of researchers used an optical flow algorithm to determine the movement of the extensor muscles [34], but they were only able to qualitatively detect different finger flexions.

The results from Castellini and Sikdar demonstrate the capabilities of modern ultrasound imaging in gesture detection, but did not explore probe location or apply recent machine learning techniques. Considering that the anatomy can vary a lot between different locations, there is likely to be some variation in performance. Cross-session accuracies have also not yet been studied, though it is vital for practical use. Given the importance of location to interactive wearables, further research is required to know whether probe location significantly effects the recognition accuracy. We shall also examine the effects of cross-session performance for each of these different locations.

ULTRASOUND IMAGING PRINCIPLES

Sound is a mechanical wave that travels by sequential compression and expansion (rarefaction) of the medium. The frequency describes the number of times the molecules expand and contract per second whereas the amplitude refers to the amount of compression.

Sound travels at a speed that depends on the temperature, pressure and the medium (e.g. 340m/s in air and 1500m/s in water under normal conditions). The product of the density of the medium and the speed of sound through it is called the acoustic impedance. When a sound wave passes from one medium to another with a different acoustic impedance, some of the energy is reflected at the boundary. The proportion that gets reflected is proportional to the mismatch in impedance between the two media. Our forearms are made up of many complex tissues with different acoustic impedances. The reflections from these tissues allow us to see the boundaries between them.

The most basic idea of ultrasonic imaging is to emit a short wave (pulse) and reading back the different reflections from the tissues. The pulse will get partially reflected at different depths of the body, the delay in arrival will be proportional to the distance whereas the intensity will indicate the type of tissue since different tissues have characteristic acoustic impedance.

An ultrasonic transducer or probe is a device which houses one or several piezoelectric elements. The elements can transform an electric pulse into a mechanical pulse to generate a wave. The elements can also transduce the mechanical energy from the reflected pulse into electric signals. These phenomena are referred to as the inverse and forward piezoelectric effect and support pulsing and reading with the same probe by quickly switching the electronic components. Modern ultrasound imaging devices use algorithms to form images, applying beamforming and harmonic imaging to increase resolution.

Ultrasonography typically uses sound waves with frequencies in the order of MHz. The higher the frequency, the better the resolution that can be attained, however the attenuation is higher and therefore imaging depth is reduced. In this paper we image musculoskeletal tissue using transducers with a central frequency ranging from 8 to 12MHz for good resolution and enough image depth.

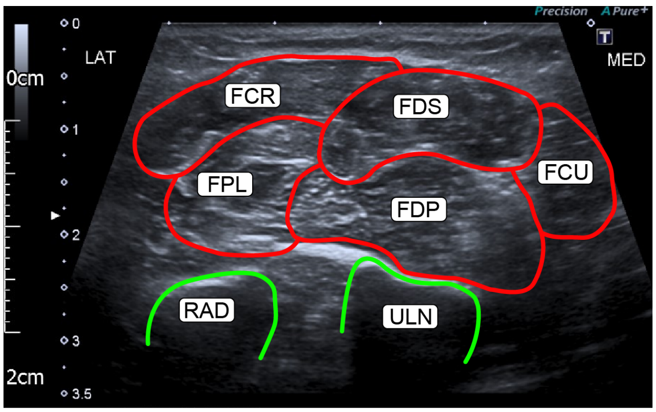

Figure 1: A transversal scan of the mid-forearm, anterior aspect. Labelled in red are the various muscles, and the bones in green.

HAND ANATOMY

The muscles which control the hand can be categorised into extrinsic muscles which originate in the upper/mid forearm, and intrinsic muscles which originate within the hand itself. The extrinsic muscles control wrist motions and some digit movements. The intrinsic muscles are used for finer motor control of the hand, for example pinch gestures and finger adduction/abduction. The extrinsic muscles can be further categorised into flexor and extensor muscles which bend or straighten the digits respectively and, in combination, move the wrist from side to side.

In this paper we exclusively image the extrinsic muscles at positions from the wrist to the forearm. These locations provide a position for the probe that will not interfere with hand movements. However, it will not be possible to capture hand movements that use intrinsic muscles, and this is a tracking limitation for all systems that measure muscle movement at the forearm level.

The muscles that are principally responsible for the flexion and extension of the fingers are the digitorum muscles (Fig. 1, FDS, FDP). These also assist with flexing and extending at the wrist, however the main muscles for the wrist are the carpi radialis/ulnaris (Fig. 1, FCR, FCU). In addition, the carpi radialis and ulnaris assist with adduction and abduction of the wrist. The thumb is controlled by the pollicis muscles (Fig. 1, FPL).

Hand movements change the musculature of the forearm due to contractions of the muscles. These contractions expand the size of the muscles and pull the tendons, and these changes are reflected in the US image. Furthermore, a particular hand movement changes the image at the specific areas where the involved tendons and muscles are located. Depending on the position of the probe (e.g. wrist or forearm) the observed muscles and tendons will be quite different.

STUDY DESIGN AND PILOT

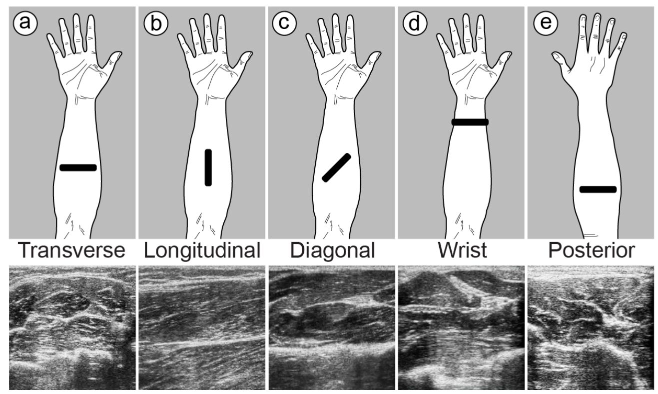

Figure 2: The mounting positions of the probe that were compared with their corresponding ultrasound image. a. Transverse, b. Longitudinal, c. Diagonal, d. Wrist, e. Posterior.

The objective of our study is to analyse the best mounting locations of a wearable device for discrete gesture recognition, cross-session accuracy and detection of continuous flexion angle. The wearable uses ultrasound imaging to track the muscles and tendons inside the forearm and with that infer hand pose. There were several variables to consider in the study; the main ones are locations and gesture sets. In order to make the study feasible we ran a pilot study to discard the options that were clearly inferior or not viable, which is presented after this section.

Locations

There are three parameters to consider when deciding the placement of the probe: orientations, proximity to the elbow, and anterior/posterior placement. The following explains each of these variables in more detail:

Orientation Clinical US scans are usually performed either **Transversally** (Fig. 2a) or **Longitudinally** (Fig. 2b) along the length of the forearm. In transverse mode, a larger range of muscles can be seen, however the muscle fibre displacement cannot be observed as well as in the longitudinal scan. It is medically uncommon to use the probe **Diagonally** (Fig. 2c) but we wanted to test if this position could image a large range of muscles and still observe the displacement of muscle fibres.

Proximal/Distal The muscles are more predominant at the **Proximal** mid-forearm (Fig. 2a), whereas the tendons are more visible at the **Distal** wrist location (Fig. 2d).

Anterior/Posterior The flexor muscles/tendons are located on the **Anterior** (inside) of the forearm and extensors on the **Posterior** (outside) of the forearm (Fig. 2e).

We chose 5 locations to initially test: Transverse, Longitudinal, Diagonal, Wrist and Posterior. These locations are illustrated in Fig. 2, along with their corresponding ultrasound images. The figure shows how different the US images look between positions. Although it cannot be seen from static images, the muscle fibre movement differs greatly with the orientation of the probe. Fig. 5a shows the fibre movement when the probe is oriented perpendicular to the flow of the muscle movement (Transverse), whereas Fig. 5b shows that the general movement is along the axis of the forearm in a Longitudinal scan.

Transverse mounting was used by Sikdar et al. [35] [1], Wrist was the location employed by Castellini et al. [6] [5] and a preliminary study tested Posterior [34]. Here, we compare those locations that were used in previous research and we introduce Longitudinal and Diagonal. Longitudinal is the only image that is completely parallel to the muscle and tendon movement. In a preliminary study for continuously detecting ankle angles it showed good accuracy [26] but we expect it to show poor performance for classifying a variety of gestures since there is less coverage of the muscles. We expected Diagonal to be a good compromise between continuous and discrete tracking.

For the purposes of dynamic assessment, clinical radiologists rarely look at images in the axial plane of the wrist, as it is difficult to identify individual tendons and the movement is not as easily appreciated as in the muscle. This was especially true of the the extensor tendons in the posterior compartment of the wrist, as the tendons here are too compact, and very little motion can be observed when compared to the anterior side. Since the wrist is normally considered to be the most ergonomic location to place a wearable device, we included the transverse distal anterior location to have at least one wrist location. If the tendons provide sufficient information for effective gesture recognition, this technique can be integrated with wrist-worn wearables.

In contrast to EMG, in US imaging we are also able to observe muscle relaxation. This means that it is not crucial to image both the anterior and posterior parts as it is possible to infer motion in both directions by inspecting either the flexor or extensor muscles. Consequently, we did not repeat posterior positions with anterior positions.

Gestures

In ultrasound images it is possible to observe the muscles and tendons which control the fingers, thumb and wrist at all locations, albeit more clearly in some locations than others. Consequently, for our gesture set we mixed a representative collection of single digit flexions, multi-finger flexion, wrist flexion and adduction.

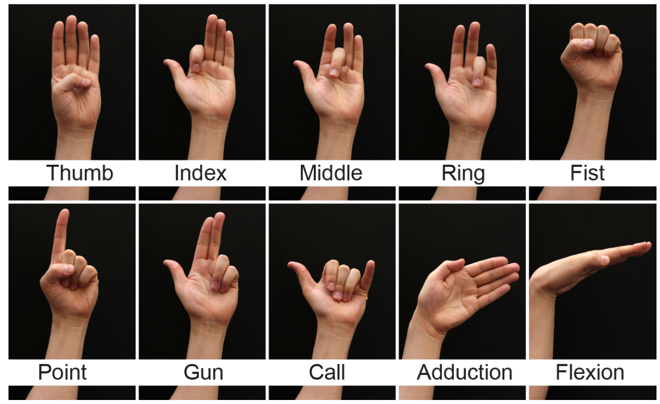

Figure 3: The set of gestures from top to bottom and left to right: thumb, index, middle, ring, fist, point, gun, call, wrist adduction and flexion.

We discarded movements that primarily involve intrinsic hand muscles such as pinch gestures and adduction of the fingers. Perhaps these gestures can be detected in the forearm because they are usually accompanied by other characteristic involuntary movements. However, we decided to remove those gestures from this study and focus on the gestures that involve a unique set of muscle movements that originate in the forearm. Fig. 3 shows the selected gesture set. We split each of these gestures into flexion and extension phases producing a total of 20 discrete movements.

There are numerous ways to flex a digit since they have at least two joints. We instructed the users to perform a natural flexion which usually involves all the joints but with more emphasis on the proximal interphalangeal joint. We thought that this would hinder the angle estimation since the finger is bent at several positions, different to previous studies in which the fingers were held with splints to force flexion at only one joint [6]. We wanted to avoid unnatural or uncomfortable gestures as they would not be used in a real situation. For the same reason, we also discarded individual pinky flexion.

With US imaging it is feasible to see the separation between the deep and the superficial muscles. For instance, the boundary between the flexor digitorum superficialis (FDS) and the flexor digitorum profundus (FDP) can be seen as a bright horizontal line in Fig. 2b . Consequently, it should be possible to differentiate between flexing of the digit at the metacarpophalangeal and the interphalangeal joints (ie different knuckles).

System

Hardware

We used a Toshiba Aplio 80 US imaging machine with a flat, linear probe with an 8MHz central frequency and 12MHz harmonic imaging mode. We used the default musculoskeletal settings with 1 focal point and 4cm of depth imaging for all the positions, except for the Wrist condition where we used 3cm. The machine was an ex-clinical machine but its deficiencies compared to the state of the art reflect the likely limitations of a mobile ultrasound device in resolution and image quality. In order to mount the probe onto the users forearm, we used a 3D printed structure to hold the transducer in place and to strap it around the user's forearm (Fig. 4a). We recorded the images from the US machine via a video capture card (Fig. 4d). Anagel US gel was applied for coupling between the probe and the user's skin.

In order to measure the ground truth of the hand pose, we created a sensor glove (Fig. 4b). We used 5 Spectra Symbol Flex Sensors [12] which were sewn onto each digit of the glove. We used an Arduino to read the resistivity values of each sensor and to transfer the data to the PC. This data was only used for the continuous angle tracking studies. We found that the sensor readings matched the amount of finger flexion linearly.

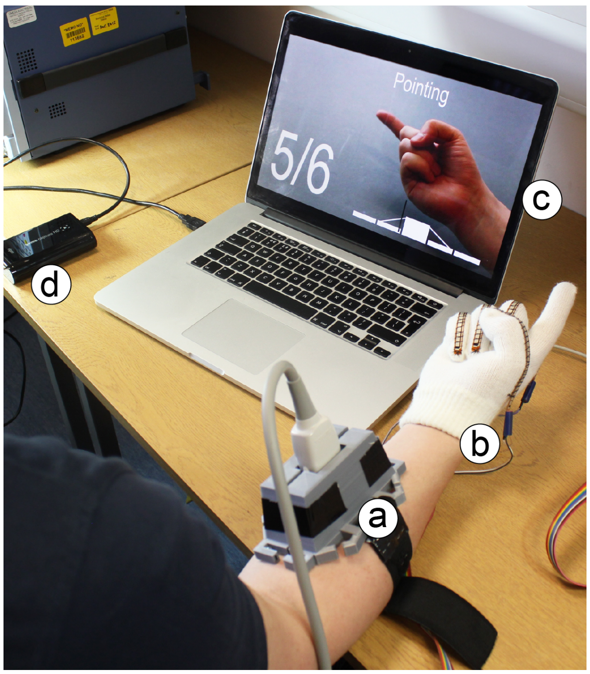

Figure 4: The setup of the experiment: a. Transducer, b. Sensor glove, c. Computer d. Video capture device.

Software

Since a gesture is actuated by a specific set of muscles, the gestures have unique patterns of activity within the US images. The type of motion depends on the location of the probe, and the observed motion is caused by reflections/scattering from the muscle fibres. In a longitudinal scan, the motion is primarily along the main axis of the probe, whereas transversal scans have a variety of different motions including rotations or elevations (Fig. 5). Nevertheless, there are patterns in the images wherever the probe is placed. The motion seemed therefore an appropriate feature to classify discrete gestures with.

Several operations must be performed with the ultrasound videos in order to classify discrete gestures or estimate finger angles. A reduced set of features are extracted from a collection of frames and then used to train a machine learning classifier. Later, this trained classifier determines the gesture or angles from a sample set of features derived from unseen images.

For the discrete classification, the first step is to segment the video stream into either flexions or extensions from the neutral position. Then, we extract features for each of the frames and these features are averaged per segment. In the continuous angle inference, the machine is a regressor that is trained with both the features at each frame and the angles from the data glove.

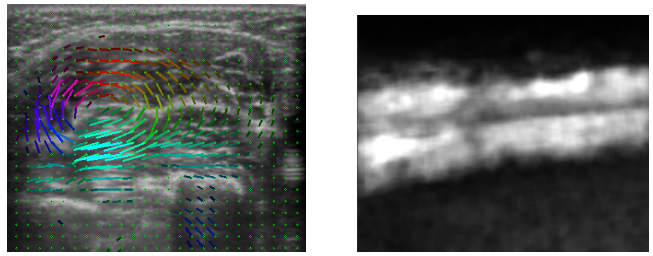

(a) Optical flow for the Transverse position during a thumb gesture.

(b) Summation of the magnitudes of the flow vectors representing the activity pattern of the index gesture.

Figure 5: Image processing algorithms.

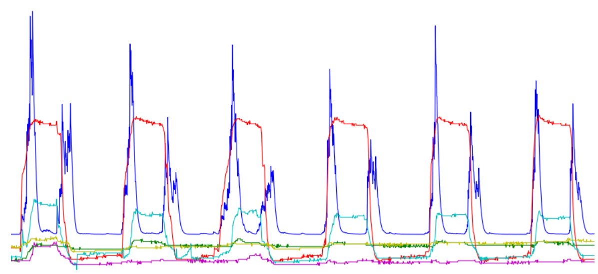

Figure 6: The blue line is a plot of the magnitudes of the flow vectors from an ultrasound video of the index finger repeatedly being flexed. The other lines represent the sensor data from the glove. Each channel is a 1D time signal.

Segmenting the videos for discrete classification was performed by calculating the sum of the magnitudes of the dense optical flow vectors for each pixel of every frame, using Farnebäck's algorithm [11]. A Gaussian blur with a kernel size of 15 pixels was applied beforehand. Fig. 5a shows the sampled points of the optical flow during the index flexing gesture and Fig. 6 plots the magnitude of the optical flow against the data glove values. The latter graph shows that the optical flow produces little or no activity during the period of time when the hand is not moving in the flexion hand pose, unlike EMG signals. The videos were split based on a plot such as this, where each gesture can easily be seen with two peaks at the beginning and end. Each gesture in the training set then has a video clip associated to show participants the gesture.

The machine learning algorithms cannot process all the information contained in an image and thus need a reduced feature set. For the discrete gesture classification, Sikdar et al. used the accumulated difference in brightness level [35] [1]. In our experiments we used the vertical and horizontal components of the optical flow. For each frame of a video clip, the dense optical flow was calculated, and then the horizontal and vertical optical flow values were averaged within each region based on a grid with 20 pixel spacing. Then the optical flow values were accumulated over every frame in the video clip, giving a final feature vector with a length in the order of hundreds of values. In continuous detection of finger angles, Castellini et al. used the first order surface as features [5], we calculated the moments for each frame (m00, m20 and m21-m12) which are equivalent but more standard in the computer vision community. The moments are computed for each frame, and the regressor trained on moments of each frame with the data glove data.

For machine learning in the discrete gesture recognition task, Sikdar et al. used Nearest Neighbour algorithms [35] [1]. We tested Support Vector Machine's (SVM) since it was shown that they gave better results for discrete gesture discrimination in another study [33]. We also tested a type of Neural Network called Multi-Layer Perceptrons (MLP) since they have often shown very good results and have yet to be applied in this technology. In the continuous angle detection, Castellini et al. used Linear Ridge Regression but suggested to use SVMs as they may yield better results [5]. In addition to testing SVMs for the regression task, we also test a multi-layer perceptron regressor. The accuracies of these algorithms are compared in the pilot study.

A shift correction was used to improve the accuracy for cross-session studies. These errors occur because the wearable can be placed slightly differently to the first time it was worn and thus the anatomical features are not in the same position. To correct this shift, we computed the optical flow between the first frame of a training set and the first frame on the current session. The flow was averaged to get a 2D translation and this transformation was applied to the current video to better align the training and sample features.

We used OpenCV [4] for processing the videos and Scikit [30] for the machine learning implementations.

Pilot Study

We eliminated some of the variables from our main experiment through a pilot study, including the least valuable locations, gestures and algorithms. For this pilot, 2 participants performed each of the 10 gestures (with both flexion and extension of each) 20 times at each of the 5 locations. We compared the accuracy of MLP and SVM, with different classifier parameters to find a suitably accurate configuration. We used 10-fold cross validation on the data set in order to have a simple and reliable measure of the accuracy for each case. Table 1 is a table of results for the average cross validation accuracy for each location, given a particular classifier, for the discrete gesture recognition case.

Classifier	Diag.	Long.	Wrist	Trans.	Post.
MLP	99.875	99.5	99.5	99.25	97.875
SVM	99.625	99.125	99.375	98.75	97.5

Table 1: Results for the average discrete gesture classification accuracies for each position, using different classification algorithms.

Classifier	Diag.	Long.	Wrist	Trans.
Participant 1	99.5	97.0	94.0	98.5
Participant 2	98.5	97.0	97.0	98.0
Average	99.0	97.0	95.5	98.25

Table 2: 10-fold cross validation classification accuracies for finger flexing at different joints.

Our initial results demonstrated that MLPs had a slight advantage over SVMs in every location, so our main study used MLPs. Our best MLP configuration had 15 neurons in the hidden layer, alpha=0.001 and BFGS for weight optimization. After initial exploration, it was very difficult to find a position for the posterior location that gave good coverage of all the required muscles. This is due to the arrangement of muscles and bones around the posterior side of the arm and the larger surface area, and is likely to be the reason for the Posterior location's weaker performance.

In our regression tests for this same data set, the same outcome was mirrored: Posterior location showed the worst performance, and the MLP neural network regressor surpassed the SVM. For this reason, we eliminated the Posterior position from our main study. The best MLP regressor configuration we used for this had 15 neurons in the hidden layer, alpha=0.0001 and Adam for weight optimization. The Posterior position is also the usual location for screens and interactive elements of wearables (e.g., watch face). Therefore our decision to ignore this location is appropriate for ergonomics as well as efficiency.

We also asked these participants to perform flexions for each of the 5 digits 10 times, flexing at the metacarpo-phalangeal joint and again at the interphalangeal joints. For each of these 10 video clips, they were again split into flexion and extension gestures. Analysis of this data indicated that it was in fact possible to differentiate between flexing at different joints, with an average accuracy of 97.4% across all positions. The different types of flexing exhibit unique areas of muscle activity across the superficial and deep muscles. This shows the potential for US to detect finer differences in hand pose, a feature that is difficult to achieve with EMG since it is difficult to differentiate between signals from the superficial and the deep flexors.

The ability to differentiate between different levels of pressure has been demonstrated in previous studies with ultrasound imaging [14] [29]. We also found discernible differences in the images and could qualitatively infer the amount of pressure with the changes in the ultrasound images.

USER STUDY

Procedure

12 participants aged between 20 and 50 years were recruited to take part in the user study; 8 were male and 4 female. The participants were asked to sit in front of a computer, wear the sensor glove and the probe mount on the right hand and then to follow the instructions shown on the monitor (Fig. 4c). The arm position was as depicted in (Fig. 4a). The video showed the user which gesture to perform and an indicator at the bottom gave visual cues for the timing of the gestures.

The conditions for the experiment were the locations of the probe: Diagonal, Longitudinal, Wrist and Transverse. Each of the 12 participants were assigned a pair of conditions (there are 12 ways of picking 2 out of 4 with order) and repeated the last condition. Therefore, each participant performed the gestures on three locations, the second and third locations were the same for testing cross-session performance. At each location, the probe was removed and placed again after a short rest, without any attempt to recalibrate using the US images. In this regard, the short time in between may not seem sufficient for a cross-session study. However, we presume the main issue for cross-session performance is sensor misalignment (as is the case with EMG [2]), as the significant changes in the US images are likely to cause classification problems. A short delay between sessions is enough to simulate the misalignment that would occur with extended delays, and this way we focus solely on the calibration mechanism without interferences from other variables.

At each location the participant had to perform a total of 10 gestures, 10 times each. Each 5-second gesture clip prompted the flexion and extension phases, which were then split into those two halves for a total of 10 flexion and extension gestures. Therefore the study involved: 12 participants X 3 locations X 10 gestures X 10 repetitions = 3600 gesture performances. The average study lasted 45 minutes including the time taken to explain the procedure to the participants and to equip the user with the sensors. The data collected from the participants were analysed offline. The data from the pilot study was not used in the main study.

Results

Different measures of accuracy were calculated for discrete recognition and continuous angle detection. All experiments are within-user.

Cross validation was performed on the data, using the MLP classifier described in the pilot study to produce classifications. We employed a 10-fold leave-one-out cross-validation strategy, with each fold containing one instance of every gesture, of which there are 20 of. Since we split gestures into flexion and extension, gestures instances of the same type were not temporally adjacent. The average classification percentages for each location are shown in Fig. 7. The confusion matrix for the Longitudinal location is shown in Fig. 9.

A one-way mixed Analysis of Variance (ANOVA) was conducted to compare the main effect of sensor location on 10-fold MLP classification performance in Diagonal, Longitudinal, Wrist and Traversal conditions. There was a significant effect of location, $F(3,6)=14.44$, $p<0.01$. Post hoc pairwise comparisons used t-tests with Bonferroni corrections to account for multiple comparisons. There was a significant difference between Diagonal (M=99.78, SD=0.09) and Longitudinal (M=97.94, SD=0.38), $t=5.67$, $p<0.01$; Diagonal and Transverse (M=98.94, SD=0.19), $t=4.08$, $p<0.05$; Longitudinal and Wrist (M=99.78, SD=0.12), $t=-4.30$, $p<0.01$; and Wrist and Transverse, $t=4.08$, $p<0.05$. There was no significant difference between Diagonal and Wrist, or Longitudinal and Transverse. From these results we can establish that for classification accuracy the Diagonal and Wrist conditions were best, followed by Transverse, with the Longitudinal condition last.

The cross-session accuracy is shown in Fig. 8 divided by location for both the raw videos and the shift-corrected videos. The cross-session accuracy is obtained by training the classifier with the data from the second session and then classifying the data from the third session which belonged to the same location.

For the continuous detection of finger angles, the main metric was the Normalised Root Mean Square Error (NRMSE) of the predicted angle compared to the real angle. The value was averaged over the five fingers since the errors were similar for all the digits. These values are shown for each location in Fig. 10.

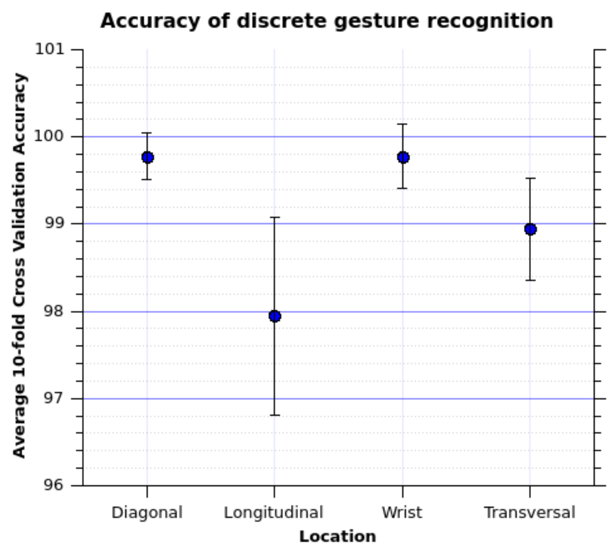

Figure 7: Accuracy of the discrete gesture classification. Average 10-fold cross-validation for each location across all participants.

DISCUSSION

The accuracy in discrete gesture recognition is very high, with the average of all locations being above 99%. The Longitudinal position has been statistically shown to perform the worst for this task. A more limited range of muscles are within the view of the probe at this position, supported by the fact that the gun and ring gestures were usually confused (for both flexion and extension - Fig. 9) since they involve a common

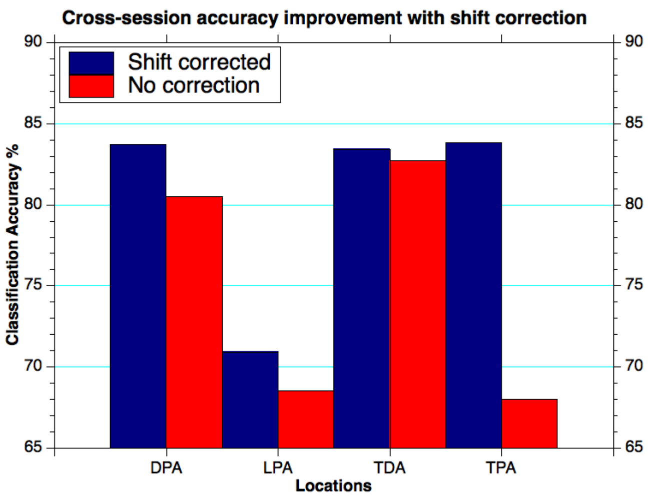

Figure 8: A graph plotting the average classification accuracies for the cross-session data at different locations.

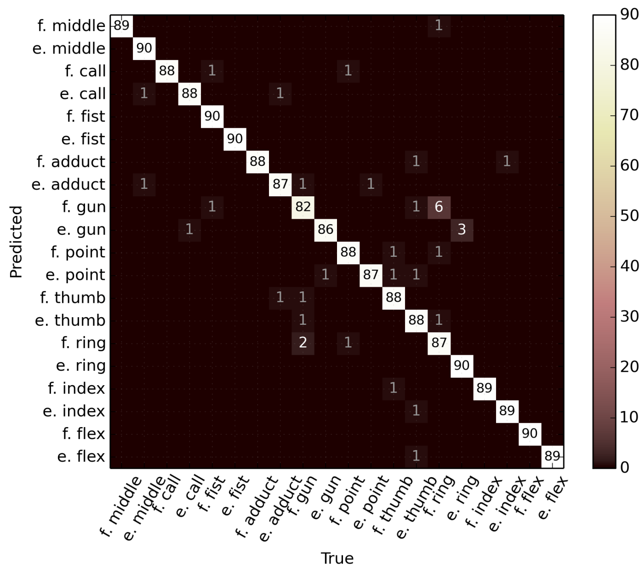

Figure 9: A confusion matrix for the Longitudinal position.

muscle. Other locations showed similar issues, however the issue was more marked in the Longitudinal images. This may also be indicative of inconsistent or insufficient information about the fifth digit muscle across all the locations. In our observations of the images, the thumb and pinky muscles were at the extremities of the images, and it is likely that the pinky could have been out of view. Unexpectedly, the Wrist location offered good results, even when the images seemed to show little movement in comparison to the proximal locations. One possible reason for this is the good coverage of the tendons by the probe, which can easily cover the central part of the wrist due to the smaller surface area. The tendons that move the fingers and wrist have to pass through the carpal tunnel in the wrist making it a concentrated area of information, in contrast

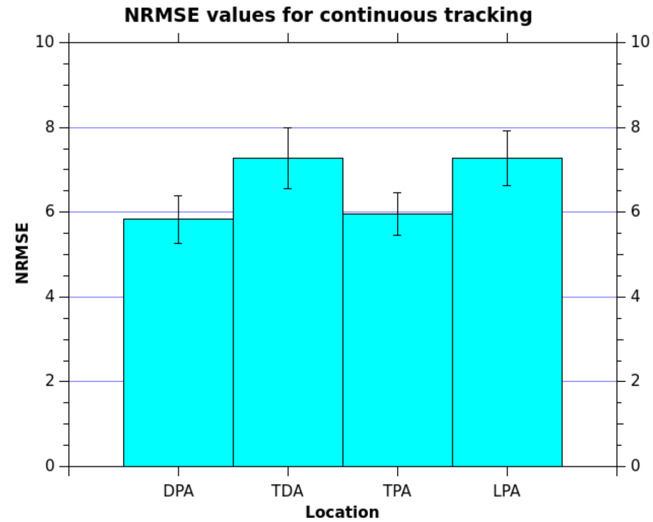

Figure 10: NRMSE values of the continuous digit flexion predictions averaged over all digits and for all participants.

to the proximal locations where the muscles and tendons are spread across a larger area.

The cross-session results show that the loss in accuracy due to probe displacement between sessions affects each location differently. Prior to shift compensation, Transverse and Longitudinal had worse cross-session accuracy when compared to the other locations. This suggests that large shifts occurred between sessions for these positions, and that it may be more difficult to place the probe in the same place between sessions. The shift compensation algorithm improves the cross session accuracy for every location, but this improvement is much larger for Transverse. In contrast, Wrist only improves slightly but that position already offered high cross-session accuracy. This is likely because the smaller size of the wrist allowed less variation in the placement of the probe. Although there were large cross-session errors for Transverse and Longitudinal, the shift can only be compensated well for Transverse. The image shift with Longitudinal cannot be recovered because the image changes differently when the probe shifts laterally, and a simple 2d translation is not enough to correct this change. The other locations are more robust to image changes as displacements perpendicular to the probe do not affect the images significantly.

In summary, the results from the study suggest that there is no significant difference between Diagonal and Wrist in discrete gesture classification, but these are superior to other locations. For ergonomic reasons, Wrist seems like the best location for discrete gestures. For continuous detection of digit angles both Diagonal and Transverse offer the best results. Therefore, the best position for both requirements is Diagonal but if ergonomics is an important factor the Wrist location could be selected for discrete gesture or Transverse for continuous angle detection.

LIMITATIONS AND FUTURE WORK

There are issues with practicality which are yet to be solved before this method is viable as a wearable gesture recognition device. In this section, we will address these limitations and provide information about the current state of research for each of these problems.

Gel is needed to couple the transducer to the skin and facilitate the transmission of ultrasound into the body. It may not be practical to use gel with an everyday wearable, but there have been recent developments which show that hydrogel pads can be used as an alternative coupling medium [20].

Participants were sedentary during our experiments. Different postures will change the muscular disposition inside the forearm and shift the features. There have been some preliminary studies that investigates classification rates for different arm positions, which show that they do not seem to significantly compromise reliability [1]. Further comprehensive research is required to analyse and mitigate the effects of arm movement during everyday activities. It may be possible to use a similar method to our cross-session calibration for this purpose. Also, with inferences of posture from other sensors (e.g. IMU), specific calibration schemes could be applied.

In our experiments, we used a decade old ex-clinical scanner, we are aware of the elevated price and bulky size of these machines. Recently, emergency point of care ultrasonography has created a need for cheaper and smaller pocket-sized scanners; in fact, several portable devices exist [38]. These handheld devices utilise the power of mobile phones to process the raw data. The image quality of these portable devices are comparable to the machine used in our studies. These devices are expensive (>$2k) but will decrease in price with acquisition boards integrated into a chip, and as research in transducer technology progresses [3]. As with most electronics, the cost of the raw materials of these devices is low.

Since the wrist requires a shorter imaging depth, it is possible to use a probe with a higher frequency. This higher frequency may yield better results with the increase in resolution. Higher frequency probes can be smaller in size, making it easier to integrate with wearables.

The probe that was used in this study was large and rigid, hindering its utilization in wearable scenarios. Transducer fabrication research has developed small and flexible thin-film probes, which should help reduce the size. Currently, it is possible to build a flexible probe with multiple elements that can be wrapped around the thumb [41]. This technology may also help to alleviate the aforementioned arm movement artifacts.

The US probe that we used is linear; that is, it has the elements arranged in one dimension. 2D probes that can image 3D volumes without moving them are becoming more available. These probes could improve robustness to shifts in the positioning of the probe, as well as providing more features to classify with.

Always-on gesture detection is desirable for real use. We do not have a dedicated study for proof of gesture spotting, but we found that high levels of activity can be measured by integrating the magnitudes of the optical flow vectors. While we only used this method to split our data, it may also be used as a variable for segmentation. However, this simple approach could be brittle and remains to be tested under realistic conditions, where motion artifacts may cause problems.

Power is a concern for wearable ultrasound imaging devices at the moment. The portable GE Vscan [13] lasts for $\tilde{1}$ hour of continuous use. Battery improvements are expected in the future, but a more interesting approach is to only activate full imaging during gestures, for instance pulsing ultrasound with fewer elements, effectively providing a low-res image to detect gesture onset. A more complete system could have other low-power sensors integrated, such as myoelectric sensors to initially spot the gesture using existing segmentation techniques.

All of the presented experiments are within-user. Much like electromyographic devices, anatomical differences between users creates an interesting challenge for cross-participant use of US wearables. Perhaps a more elaborate classifier and calibration scheme that takes into account the common anatomical features between users, could lead to an effective user-independent system.

Bio identification could be possible with ultrasound imaging. The veins and other anatomical features are detectable with high frequency probes; previously, it has been shown that the structure of the veins can be used as an unique characteristic for identification [24]. This technique could be used to identify the wearer and load their preferences, without the need for external validation.

CONCLUSION

We have presented a solution for detecting discrete gestures and tracking continuous angles of the fingers using ultrasound images captured from a probe mounted on the forearm. Our novel contributions include findings on the variation in performance between different mounting locations. In contrast to previous studies, we provide results for the cross-session accuracy decrease, and show that a simple calibration algorithm can improve the accuracy, and highlights the usefulness for future work. In conclusion, the performance variation across the tested locations vary somewhat insignificantly, meaning that an ergonomic location such as the wrist may be chosen as the desired location for a wearable US device. However, there are some differences in robustness and continuous angle recognition, and depending on the requirements other locations may be more desirable. Wearable applications using ultrasound imaging for gesture detection will become more feasible with the continued decrease in cost and size of US devices, and we hope that this work inspires research in this promising approach.

ACKNOWLEDGMENTS

This research was supported by EPSRC Doctoral Training funding through grant EP/M507994/1. We especially thank members of BIRCH within the UH Bristol NHS Foundation Trust. We also thank Nvidia for facilitating our research by providing hardware and Matthew Sutton for assistance with the video and pictures.

REFERENCES

1. Akhlaghi, N., Baker, C. A., Lahlou, M., Zafar, H., Murthy, K. G., Rangwala, H. S., Kosecka, J., Joiner, W. M., Pancrazio, J. J., and Sikdar, S. Real-time classification of hand motions using ultrasound imaging of forearm muscles. *IEEE Transactions on Biomedical Engineering 63*, 8 (Aug 2016), 1687–1698.
2. Amma, C., Krings, T., Böer, J., and Schultz, T. Advancing muscle-computer interfaces with high-density electromyography. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM (2015), 929–938.
3. Bhuyan, A., Choe, J. W., Lee, B. C., Wygant, I. O., Nikoozadeh, A., Oralkan, Ö., and Khuri-Yakub, B. T. Integrated circuits for volumetric ultrasound imaging with 2-d cmut arrays. *IEEE transactions on biomedical circuits and systems 7*, 6 (2013), 796–804.
4. Bradski, G., et al. The opencv library. *Doctor Dobbs Journal 25*, 11 (2000), 120–126.
5. Castellini, C., and Passig, G. Ultrasound image features of the wrist are linearly related to finger positions. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, IEEE (2011), 2108–2114.
6. Castellini, C., Passig, G., and Zarka, E. Using ultrasound images of the forearm to predict finger positions. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on 20*, 6 (2012), 788–797.
7. Castellini, C., and van der Smagt, P. Surface emg in advanced hand prosthetics. *Biological cybernetics 100*, 1 (2009), 35–47.
8. Chen, X., Zheng, Y.-P., Guo, J.-Y., and Shi, J. Sonomyography (smg) control for powered prosthetic hand: a study with normal subjects. *Ultrasound in medicine & biology 36*, 7 (2010), 1076–1088.
9. Cheng, J., Amft, O., Bahle, G., and Lukowicz, P. Designing sensitive wearable capacitive sensors for activity recognition. *Sensors Journal, IEEE 13*, 10 (2013), 3935–3947.
10. Dementyev, A., and Paradiso, J. A. Wristflex: Low-power gesture input with wrist-worn pressure sensors. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*, ACM (2014), 161–166.
11. Farnebäck, G. Two-frame motion estimation based on polynomial expansion. In *Image analysis*. Springer, 2003, 363–370.
12. Spectra Symbol Flex Sensor 4.5. `http://www.spectrasymbol.com/wp-content/themes/spectra/images/datasheets/FlexSensor.pdf`.
13. GE Vscan. `http://www3.gehealthcare.com/en/products/categories/ultrasound/vscan_portfolio/vscan`.
14. González, D. S., and Castellini, C. A realistic implementation of ultrasound imaging as a human-machine interface for upper-limb amputees. *Frontiers in neurorobotics 7*, 17 (2013).
15. Guo, J.-Y., Zheng, Y.-P., Huang, Q.-H., Chen, X., and He, J.-F. Comparison of sonomyography and electromyography of forearm muscles in the guided wrist extension. In *Medical Devices and Biosensors, 2008. ISSS-MDBS 2008. 5th International Summer School and Symposium on*, IEEE (2008), 235–238.
16. Guo, J.-Y., Zheng, Y.-P., Kenney, L. P., Bowen, A., Howard, D., and Canderle, J. J. A comparative evaluation of sonomyography, electromyography, force, and wrist angle in a discrete tracking task. *Ultrasound in medicine & biology 37*, 6 (2011), 884–891.
17. Hamid Muhammed, H., and Raghavendra, J. Optomyography (omg): A novel technique for the detection of muscle surface displacement using photoelectric sensors. In *10th International Conference on Bioelectromagnetism, ISBEM 2015*, vol. 10 (2015).
18. Hettiarachchi, N., Ju, Z., and Liu, H. A new wearable ultrasound muscle activity sensing system for dexterous prosthetic control. In *Systems, Man, and Cybernetics (SMC), 2015 IEEE International Conference on*, IEEE (2015), 1415–1420.
19. Hodges, P., Pengel, L., Herbert, R., and Gandevia, S. Measurement of muscle contraction with ultrasound imaging. *Muscle & nerve 27*, 6 (2003), 682–692.
20. HydroAid Ultrasound Hydrogel Pad. `http://www.civco.com/mmi/ultrasound/accessories/acoustic-standoffs/Disposable-AquaFlex-Acoustic-Standoff-Pads-610-1322.htm`.
21. Islam, M. A., Sundaraj, K., Ahmad, R. B., and Ahamed, N. U. Mechanomyogram for muscle function assessment: a review. *PloS one 8*, 3 (2013), e58902.
22. Kerber, F., Löchtefeld, M., Krüger, A., McIntosh, J., McNeill, C., and Fraser, M. Understanding same-side interactions with wrist-worn devices. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*, NordiCHI '16, ACM (New York, NY, USA, 2016), 28:1–28:10.
23. Kim, D., Hilliges, O., Izadi, S., Butler, A. D., Chen, J., Oikonomidis, I., and Olivier, P. 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*, ACM (2012), 167–176.
24. Kono, M., Ueki, H., and Umemura, S.-i. Near-infrared finger vein patterns for personal identification. *Applied Optics 41*, 35 (2002), 7429–7436.

25. Korstanje, J.-W. H., Selles, R. W., Stam, H. J., Hovius, S. E., and Bosch, J. G. Development and validation of ultrasound speckle tracking to quantify tendon displacement. *Journal of biomechanics 43*, 7 (2010), 1373–1379.
26. Loram, I. D., Maganaris, C. N., and Lakie, M. Use of ultrasound to make noninvasive in vivo measurement of continuous changes in human muscle contractile length. *Journal of applied physiology 100*, 4 (2006), 1311–1323.
27. Mujibiya, A., Cao, X., Tan, D. S., Morris, D., Patel, S. N., and Rekimoto, J. The sound of touch: On-body touch and gesture sensing based on transdermal ultrasound propagation. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces*, ITS '13, ACM (New York, NY, USA, 2013), 189–198.
28. Orizio, C., and Gobbo, M. Mechanomyography. *Wiley Encyclopedia of Biomedical Engineering* (2006).
29. Ortenzi, V., Tarantino, S., Castellini, C., and Cipriani, C. Ultrasound imaging for hand prosthesis control: a comparative study of features and classification methods. In *Rehabilitation Robotics (ICORR), 2015 IEEE International Conference on*, IEEE (2015), 1–6.
30. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. Scikit-learn: Machine learning in python. *The Journal of Machine Learning Research 12* (2011), 2825–2830.
31. Rautaray, S. S., and Agrawal, A. Vision based hand gesture recognition for human computer interaction: a survey. *Artificial Intelligence Review 43*, 1 (2015), 1–54.
32. Rekimoto, J. Gesturewrist and gesturepad: Unobtrusive wearable interaction devices. In *Wearable Computers, 2001. Proceedings. Fifth International Symposium on*, IEEE (2001), 21–27.
33. Rossi, M., Benatti, S., Farella, E., and Benini, L. Hybrid emg classifier based on hmm and svm for hand gesture recognition in prosthetics. In *Industrial Technology (ICIT), 2015 IEEE International Conference on*, IEEE (2015), 1700–1705.
34. Shi, J., Hu, S.-X., Liu, Z., Guo, J.-Y., Zhou, Y.-J., and Zheng, Y.-P. Recognition of finger flexion from ultrasound image with optical flow: A preliminary study. In *Biomedical Engineering and Computer Science (ICBECS), 2010 International Conference on*, IEEE (2010), 1–4.
35. Sikdar, S., Rangwala, H., Eastlake, E. B., Hunt, I. A., Nelson, A. J., Devanathan, J., Shin, A., and Pancrazio, J. J. Novel method for predicting dexterous individual finger movements by imaging muscle activity using a wearable ultrasonic system. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on 22*, 1 (2014), 69–76.
36. Smith, R. J., Tenore, F., Huberdeau, D., Cummings, R. E., and Thakor, N. V. Continuous decoding of finger position from surface emg signals for the control of powered prostheses. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, IEEE (2008), 197–200.
37. Tenore, F. V., Ramos, A., Fahmy, A., Acharya, S., Etienne-Cummings, R., and Thakor, N. V. Decoding of individuated finger movements using surface electromyography. *Biomedical Engineering, IEEE Transactions on 56*, 5 (2009), 1427–1434.
38. Wojtczak, J., and Bonadonna, P. Pocket mobile smartphone system for the point-of-care submandibular ultrasonography. *The American journal of emergency medicine 31*, 3 (2013), 573–577.
39. Wolf, M. T., Assad, C., Stoica, A., You, K., Jethani, H., Vernacchia, M. T., Fromm, J., and Iwashita, Y. Decoding static and dynamic arm and hand gestures from the jpl biosleeve. In *Aerospace Conference, 2013 IEEE*, IEEE (2013), 1–9.
40. Yamakawa, S., and Nojima, T. A proposal for a mmg-based hand gesture recognition method. In *Adjunct proceedings of the 25th annual ACM symposium on User interface software and technology*, ACM (2012), 89–90.
41. Yang, Y., Tian, H., Yan, B., Sun, H., Wu, C., Shu, Y., Wang, L.-G., and Ren, T.-L. A flexible piezoelectric micromachined ultrasound transducer. *RSC Advances 3*, 47 (2013), 24900–24905.
42. You, K.-J., Rhee, K.-W., and Shin, H.-C. Finger motion decoding using emg signals corresponding various arm postures. *Experimental neurobiology 19*, 1 (2010), 54–61.
43. Zhang, Y., Xiao, R., and Harrison, C. Advancing hand gesture recognition with high resolution electrical impedance tomography. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, UIST '16, ACM (New York, NY, USA, 2016), 843–850.
44. Zheng, Y.-P., Chan, M., Shi, J., Chen, X., and Huang, Q.-H. Sonomyography: Monitoring morphological changes of forearm muscles in actions with the feasibility for the control of powered prosthesis. *Medical engineering & physics 28*, 5 (2006), 405–415.