

Thumbs-Up

Wearable Sensing Device for Detecting Hand-to-Mouth Compulsive Habits

Asaph Azaria, Brian Mayton and Joseph Paradiso

Media Lab, Massachusetts Institute of Technology, 75 Amherst st, Cambridge, MA, U.S.A.

Keywords: Biomedical Computing, Bioimpedance, Sensor Applications, Artificial Intelligence, Wearable Computers.

Abstract: Thumbs-Up explores a novel sensing method for detection of hand-to-mouth compulsive habits. Using electrical bio-impedance spectroscopy and inertial measurement units, a prototype system was implemented. The system can easily be worn around the arm and may perhaps be integrated into future wearable devices. It recognises occurrences of the habits in real-time, allowing monitoring and immediate interventions. These have so far been very limited, impeding behavioural studies and the development of therapeutic treatments. Throughout this paper the method's feasibility is demonstrated and aspects of its performance are explored. We present an approach to process the bio-impedance signals and associate them with possible body postures. A positioning strategy optimises the device's sensitivity and increases its efficacy. Machine learning algorithms are leveraged to infer the hand-to-mouth detection. We achieve 92% detection accuracy for recurrent usage and 90% accuracy for users that have not been previously encountered.

1 INTRODUCTION

Hand-to-mouth compulsive habits, such as thumb sucking and nail biting, are surprisingly common. Studies found that 28-33% of children, 44% of adolescents, and 19-29% of young adults engage in nail biting alone (Tanaka et al., 2008). These compulsive behaviours expose those who exhibit them to multiple health risks. Threats range from expedited transmission of diseases, to dental malocclusion and even abnormal facial development (Baydaş et al., 2007).

Studies into the roots of these habits and their consequences have so far been limited to either subjective reports or lab settings. Treatment has been constrained to retrospective correction, relying on self awareness and manual monitoring (Azrin and Nunn, 1973). A wearable device, detecting hand-to-mouth habits in real-time, would mitigate these impediments, advancing therapeutic interventions and personal health monitoring. With aetiologies including anxiety, loneliness, frustration, and more, such a device could even serve as a diagnostic tool to reflect on quality of life impairments (Pacan et al., 2014).

This work presents a wearable hand-to-mouth classification system, consisting of an inertial measurement unit (IMU) and an electrical bio-impedance spectrometer. It can easily be worn around the arm. The device tracks the arm's orientation and electri-

cally excites it, discovering changes in its impedance properties. These are leveraged to infer hand-to-mouth behaviours in real-time.

To the best of our knowledge, this is the first system to demonstrate automated detection of hand-to-mouth habits. Furthermore, it is innovative in utilising bio-impedance spectroscopy for this goal. A related detection scheme is described in US Patent 6762687 B2. Rather than sensing physical contact, this scheme senses proximity between tagged body parts. Bio-impedance has been widely studied in the medical context for estimating biophysical quantities such as body composition, metabolic functioning, and cardiac activity. Recent research has demonstrated integrating bio-impedance measurements into wearable devices – monitoring biophysical markers (Seppä et al., 2007) and even estimating users' biometrics (Martinsen et al., 2007). None, however, leveraged bio-impedance technology for behaviour recognition, as our system uniquely does.

2 THEORY OF OPERATION

2.1 Electrical Properties of Biomaterials

Electrical bio-impedance (EBI) describes the electrical properties biological materials exhibit as current

flows through them. It is commonly measured by injecting a small sinusoidal alternating current (AC) into the tissue under study. The injected current induces an electrical field within the tissue and results in a measurable voltage drop across it. The AC version of Ohm's Law (Equation 1) relates the material's electrical bio-impedance Z to the ratio between the measured voltage V and the injected current I .

$$Z = \frac{V}{I} \quad (1)$$

It is a complex quantity, since biomaterials not only oppose current flow, but also store electrical charge, phase-shifting the voltage with respect to the current in the time-domain.

When current flows through a tissue, it passes through extracellular and intracellular fluids. These fluids are highly conductive as they contain salt ions that can easily be displaced by a potential difference. Conversely, the cells' lipid membranes are insulators. They act like capacitive plates, which prevent electrical charges from flowing through. Accordingly, the tissue's impedance reflects its chemical composition, membrane structures, and fluids distribution. For similar reasons, the specific cell types (blood, adipose, muscle, bone, etc.), the anatomic configuration (i.e., bone or muscle orientation and quantity), and the state of the cells (normal or osteoporotic bone, oedematous vs. normally hydrated tissue, etc.) greatly affect measured impedance quantities (Gabriel et al., 1996).

Most tissues display dispersive characteristics, i.e. their impedance varies with the frequency of the applied current. A typical dispersion curve is illustrated in Figure 1. It is displayed as a Cole-Cole plot which superimposes impedance measurements from a range of frequencies on the complex plane. At low frequencies, the cells' membranes block the current. Thus, the impedance corresponds only to the extracellular resistance. As frequency increases, more current passes through the intracellular capacitive path, and the phase angle accumulates. At high frequencies, the intracellular capacitance becomes negligible. The impedance is once again purely resistive, dominated by the intracellular and extracellular fluid resistances connected in parallel. The frequency at which the tissue's reactance reaches a peak is known as the centre frequency (f_c).

2.2 The Cole-Cole Equivalent Circuit Model

The simplest electrical circuit that can be used to model EBI response is presented in Figure 2. R_i and C_m represent the intracellular current branch and R_e

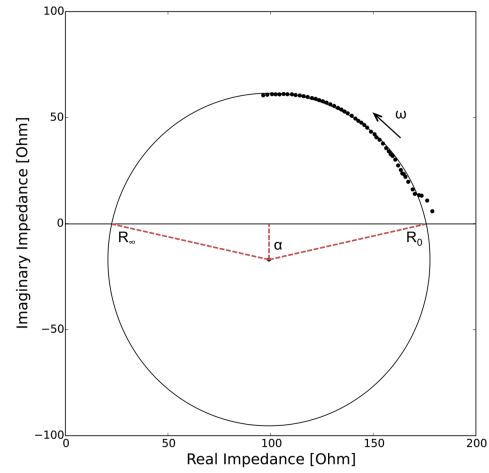


Figure 1: Typical Cole-Cole plot of biological tissue.

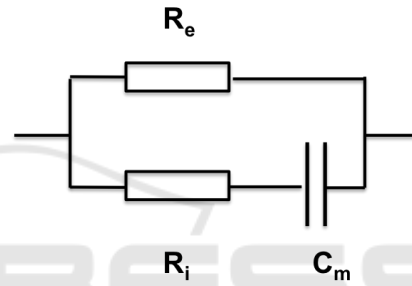


Figure 2: Simple circuit model for biological tissues.

represents the extra cellular one. This model results in a perfect semicircle in a Cole-Cole plot, with the centre of the circle on the resistance axis.

In real tissue, however, the cells' membrane is an imperfect capacitor. Moreover, the large variation in cell type, structure, and size causes a distribution of the cells' capacitive time constants (Cornish et al., 1993). Cole and Cole showed that when capacitive time constant distribution is added to the aforementioned circuit model, the impedance is related to the frequency by (Cole and Cole, 1941):

$$Z = R_\infty + \frac{R_0 - R_\infty}{1 + (j\omega\tau)^{1-\alpha}} \quad (2)$$

where,

$$R_0 = R_e \quad (3)$$

$$R_\infty = R_i || R_e \quad (4)$$

$$\tau = (R_e + R_i)C_m \quad (5)$$

This model preserves the circular shape in a Cole-Cole plot, but depresses the circle's centre below the resistance axis. α has a value between 0 and 1, and is proportional to the angle to the depressed centre.

2.3 Electrode Sensitivity Field and Configuration

Two types of electrode systems are commonly used to obtain EBI measurements. A 2-electrode system uses the same pair of electrodes to inject current (IC) and pick up (PU) the tissue's response. A 4-electrode system uses different pairs for excitement and pick up. The type of electrodes in use (needle or skin surface, gel or dry, etc.), and their configuration around the tissue affect the sensed impedance almost as much as the electrical properties of the tissue.

The electrode type determines how the electrical conductor in the measurement leads interfaces with the ionic conductor in the biological tissue. As current flows, substance concentration may change near the electrodes' interface, adding bias impedance called electrode polarisation. The skin-electrode contact introduces an additional resistive bias. The 4-electrode system is a robust setup that reduces the influence of these factors (Seoane et al., 2008). When voltage pick up is implemented with high-impedance differential amplifiers, such artefacts can be neglected.

The electrode configuration sets boundary values on the electrical fields that develop inside the tissue. Thus, it governs the fields' propagation and in effect, the relative contribution of internal tissue regions to the measured mutual impedance. Geselowitz formulated this idea for a 4-electrode system (Geselowitz, 1971). He established that the measured impedance Z resulting from the variable conductivity σ within a volume conductor can be evaluated by:

$$Z = \int_V \frac{1}{\sigma} S dv \quad (6)$$

Sensitivity S is a scalar field, determining the contribution of a local conductivity change ($\Delta\sigma$) to the overall potential. It can be calculated from the dot product of two current density fields:

$$S = \vec{J}_{IC} \cdot \vec{J}_{PU} \quad (7)$$

\vec{J}_{IC} represents the current density field generated by a unit current applied through the IC electrodes. \vec{J}_{PU} is the reciprocal current density field that would have been generated had the same current been injected through the PU electrodes.

Depending on the angle of the two fields, there can be regions where the sensitivity is zero, positive, or negative. Hence, the tissue regions, in which impedance changes are measured, can effectively be manipulated by the electrode configuration. Note that the measured mutual impedance is indifferent to interchanges between the IC and PU electrodes. In this context, this phenomenon is commonly referred to as the Reciprocity Theorem.

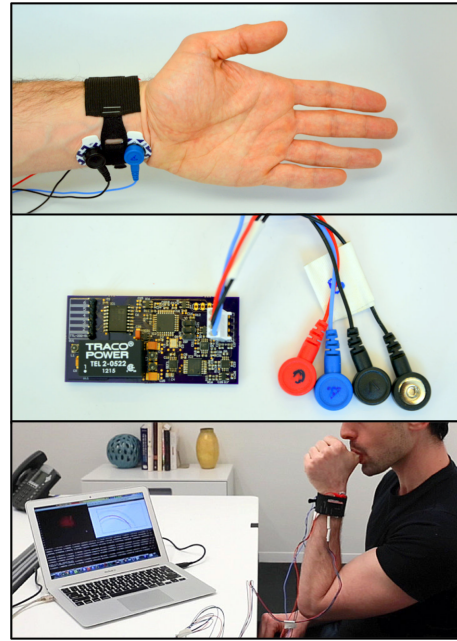


Figure 3: Images of the prototype system (From top: IMU wristband, EBI spectrometer, and PC Interface).

3 SYSTEM DESCRIPTION

Inspired by how pervasively electrical fields propagate inside biological tissues, we set out to examine if EBI spectroscopy can be applied to sense hand-to-mouth behaviours. We hypothesised that placing the hand inside the mouth will form different paths for the current to traverse inside the body, resulting in distinct impedance spectra. We presumed these can be analysed, detecting the placement of the hand inside the mouth. In light of the multiple factors that manifest in an EBI spectrum, we explored machine-learning techniques for making such distinctions. A similar technique, learning from capacitive signals, has previously been used to classify simple gestures for human-computer interfaces, suggesting promise to this approach (Sato et al., 2012).

To test this hypothesis and assess its validity, we built a prototype system. Images of the system are presented in Figure 3. We present the system's design in detail throughout the following section. When necessary, we discuss alternative implementation schemes and possible improvements.

3.1 Hardware

The hardware consists of an EBI spectrometer and an IMU, wired to a nearby personal computer via USB. The bio-impedance spectrometer was implemented on a printed circuit board (PCB) and was

based on Analog Devices' AD5933 on-chip network analyser. The analyser contains a frequency generator capable of outputting a sinusoidal excitation voltage. The response signal is picked up by the analyser and processed by an on-chip Discrete Fourier Transform. Real and imaginary 16-bit data words, proportional to the measured impedance, are returned for each output frequency. A frequency range between 1 to 100 kHz is supported, with a 0.1 Hz resolution.

We interfaced the AD5933 with the analogue front-end presented in (Seoane et al., 2008). This front-end, illustrated in Figure 4, converts the AD5933 from a 2-electrode to a 4-electrode measurement system. Hence, it cancels impedance contributions from electrode polarisation and skin-electrode coupling. An instrumentation amplifier, voltage follower, and a hardwired 20 kΩ resistor construct a voltage controlled current source (VCCS), which modifies the AD5933 excitation from being voltage to current driven. An auxiliary high-pass-filter (HPF) blocks DC voltage from flowing into the tissues. This way, current is controlled and regulated well below hazardous levels. Frequencies are limited above 3.5 kHz to fully comply with IEC-60601 safety guidelines. AC current of 174 μA (RMS) was confirmed by an independent measurement in our system.

An ATmega328P micro-controller controls the AD5933, via I²C, and obtains the impedance measurements. It is used to execute parts of our algorithms, assessing their performance on an embedded platform. We followed the layout by (Blomqvist et al., 2012) to power the board and communicate with the micro-controller over USB.

This layout does not restrict power consumption and requires 0.5W for its operation. It is directly wired to a PC, and so requires inefficient power and communication isolation, that will not be necessary in a battery powered wearable device. The basic components in use, however, are well suited for low-power applications. Typical power consumption of the AT-Mega328P and the AD5933 are 33mW and 20mW, respectively. Likewise, the analogue front-end can be adjusted to require minimal power. Future versions could easily revise power and communication to fit the requirements of a wearable device.

We used 1" adhesive gel electrodes (Covidien H124SG), commonly used in ECG and EMG, for our experiments. These surface electrodes allow non-invasive and precise fixture of the electrodes to the arm. The contact electrolyte gel is of less importance to our application. It is designed to mitigate the high impedance of the skin at lower frequencies, typically under 1 kHz. By removing the DC component from the excitation signal and alternating to a

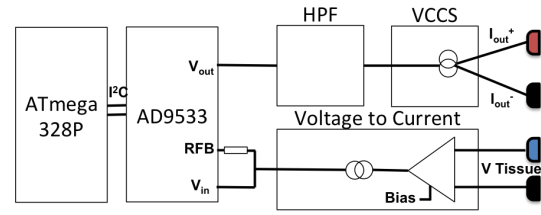


Figure 4: Block diagram of the analogue front-end described by Seoane et al.

4-electrode configuration, the polarisation characteristics of the particular electrodes in use can also be neglected (Seoane et al., 2008).

We therefore expect that a future change to dry electrodes – relying on sweat as the electrolytic solution – will not significantly affect our results. Dry electrodes will lower our system's cost and improve its unobtrusiveness and reusability. It will, however, have to assume the continuous natural perspiration of the skin. Otherwise, the sweat electrolytes, carrying our system's current, can be depleted over time.

The IMU is attached to a wristband-like strap and worn on the user's forearm. It streams measurements of the forearm's orientation and displacement. We used an MPU9150, which comprises a 3-axis accelerometer, a gyroscope, and a magnetometer in a single chip. The chip was placed on a breakout board and connected to an off-the-shelf Arduino with an identical ATmega328P micro-controller. In future versions, the IMU can easily be connected to the spectrometer's I²C bus, embedding all sensors on a single board.

3.2 Software

3.2.1 Calibration

The impedance measurements – real and imaginary 16-bit data words for each exciting frequency – are first scaled and calibrated. The calibration accounts for the complex impedance introduced by our system's various electronic components – between the AD5933 and the studied tissues.

A system **gain factor** scales the magnitude of each impedance measurement. Such scaling can be formulated as:

$$|Z_{tissue}(\omega)| = \frac{\text{Magnitude}}{\text{Gain Factor}} = \frac{\sqrt{\mathcal{R}(\omega)^2 + \mathcal{I}(\omega)^2}}{G(\omega)} \quad (8)$$

We used a 2R-1C circuit with known impedance, like in Figure 2, to derive the frequency-dependent gain factor of our system. Similarly, the **system phase** offset is derived from a circuit with pure resistance. It

factors in as a frequency dependent bias when calibrating the impedance measurements:

$$\angle Z_{tissue}(\omega) = \tan^{-1}\left(\frac{I(\omega)}{\mathcal{R}(\omega)}\right) - \angle_{system}(\omega) \quad (9)$$

The IMU allows programmable full-scale ranges to each of its components. Based on the arm's characteristic motion, we set these ranges to ± 250 dps, ± 2 g, and ± 1200 μ T for the gyroscopes, accelerometer, and magnetometer respectively. We measured offsets in the angular velocity and linear acceleration with a simple gimbal. Each of the 9 measurements were scaled and corrected accordingly.

3.2.2 Feature Extraction

Features for hand-to-mouth detection are extracted from calibrated EBI and IMU measurements. These features are listed in Table 1.

Table 1: List of extracted features.

EBI features	IMU features
Circle centre (X_c)	Quaternion orient. (q_0)
Circle centre (Y_c)	Quaternion orient. (q_1)
Circle radius (R)	Quaternion orient. (q_2)
Infinity impedance (R_∞)	Quaternion orient. (q_3)
Static impedance (R_0)	
Depression angle (α)	
Centre frequency (f_c)	
Fit accuracy ($\text{Var}\{R_n\}$)	

EBI Features

We chose the Cole-Cole model to compactly represent the EBI data, using its parameters as features. To calculate them, the calibrated impedance samples are mapped to the impedance plane in a Cole-Cole plot. Such mapping was exemplified in Figure 1. Each data point on the plane corresponds to the resistance and reactance measured at a single frequency. We seek to fit these points to the Cole equation – a perfect semi-circle with the centre depressed below the resistance axis.

We leverage the circular shape to estimate the parameters of the Cole-Cole model. Instead of solving directly for the model parameters (R_∞ , R_0 , τ , α and f_c), we estimate the circular curve that best represents our data. We do so by finding a point (X_c, Y_c) that minimises the variance in the distances R_n , measured from it to the N impedance data points.

$$\min \text{Var}\{R_n^2\} = \min \text{Var}\{|(X_c, Y_c) - Z_n|^2\} \quad (10)$$

This point will be the Cole-Cole model circle centre. A closed form solution for such a minimisation

is obtained by differentiating Equation 10. The Cole-Cole model parameters, R_∞ , R_0 and α , can then be solved geometrically from the circle centre. A detailed derivation is found in (Ayllón et al., 2009).

The centre frequency f_c is determined as presented in the work of (Cornish et al., 1993). Lengths of cords u_i and v_i , which respectively join each impedance data point i with R_∞ and R_0 , are calculated. The impedance points are then projected to a $\log(u_i/v_i)$ vs. $\log(\omega_i)$ plane, where they yield a line with slope $(1 - \alpha)$. The x-axis intercept of this line determines the centre frequency.

We calculate the variance in R_n^2 and affix it as an additional parameter to our EBI data representation. It serves as an indicator of how accurately our data fits to a circular curve. Later, it will be utilised when performing real-time classification, to identify movement artefacts and discard them.

IMU Features

We used an AHRS (Attitude and Heading Reference System) sensor-fusion algorithm to handle the IMU measurements. The algorithm produces a four-vector quaternion representation of the device's orientation in 3D. Rather than using computationally expensive Kalman-Filter implementations, we chose the sensor fusion algorithm developed by Madgwick (Madgwick et al., 2011). This algorithm derives simple gradient descent steps to conduct fusion and estimations iteratively, allowing it to operate on computationally constrained platforms compatible with wearable devices.

Execution

The extraction algorithms were designed to execute in real time. The orientation state is updated continuously, processing IMU measurements at a 400 Hz sampling rate. The EBI feature extraction algorithm processes a frequency sweep at a 0.64 Hz rate. We configured the frequency sweep to obtain 50 samples equally spaced between 4 kHz and 100 kHz. The frequency resolution can be modified to boost the sweeping rate at the expense of the fit accuracy. The sweeping rate determines the rate at which the feature vectors are generated. The EBI features are first extracted and then extended with the current four-vector quaternion orientation.

The feature extraction algorithms were implemented both for a personal computer and for the ATmega328P micro-controllers. We experimented with executing them on both. Running on the micro-controller reduces dramatically (factor 18) the data to be communicated to the next stage. It does so, however, at the cost of an increased computational load.

We experienced no performance issues executing the algorithms on the ATmega328P, when the sampling rates were set as previously mentioned. We leave the decision between the execution alternatives open for future applications.

3.2.3 Hand-to-Mouth Detection

The final stage of our software allows three modes of operation: **Collection mode** logs real-time collected data and extracted features into a local file for later analysis. A human observer labels the data points with a binary label. Positive and negative labels correspond to situations when the hand is inside and outside the mouth. **Enrolment mode** accepts previously collected labeled data either from a specific user or a group of users. It optimises a set of machine-learning algorithms and chooses the one which is most likely to perform best. **Prediction mode** makes real-time predictions on the hand's situation. It uses the optimal classifier from the previous mode to predict on each feature vector as it is produced. The predictions, as well as the arm's orientation, are visualised to the user as feedback.

Our system allows learning via any binary-classification model, as long as it implements a simple *fit*, *score* and *predict* interface. We embedded the *Scikit-learn* implementation of five classification models in our system: Random Forests (RF), Support Vector Machines (SVM), AdaBoost (AB), K-Nearest Neighbours (KNN) and Logistic Regression (LR) (Pedregosa et al., 2011). The models were chosen to intentionally differ in their underlying statistical assumptions and algorithmic implementations. Hence, they differ not only in the accuracy they may achieve, but also in the computational complexity their training or prediction involve.

We optimise the unique parameters of the classification models, by searching a grid of possible values. A classifier instance is trained for each parameter combination. How well it generalises is assessed using 3-iteration random subsampling cross-validation (Murphy, 2012). Once the optimal parameters per classification model are set, we choose among the models by their mean validation accuracy.

Prediction mode introduces two additional mechanisms to facilitate the real-time prediction and feedback. First, it discards data points whose Cole-Cole circular fit accuracy is below a threshold (i.e. high $\text{Var}\{R_n^2\}$). These were observed during natural movement, in singular cases – when by the time a frequency sweep is over, the user has already changed between postures with different EBI spectra. Clearly, fitting to the inconsistent measurements does not truthfully represent the electrical path taken by the

current. Discarding them, therefore, prevents unnecessary prediction errors. Second, a moving average window on the predicted labels was implemented, smoothing the visual feedback presented to the user.

4 EVALUATION

In this section we demonstrate the system's feasibility by evaluating various aspects of its performance. Note that the prototype system leaves many degrees of freedom in its application – from configuring the electrodes on the forearm, to when it is worn and by whom. We therefore tackle this challengingly broad space by breaking the evaluation into three stages. First, we present an approach to process the bio-impedance signals and verify they carry information about hand-to-mouth behaviours. We associate them with possible body postures, narrowing down the scope of movement for our experiments. Second, an electrode positioning strategy is proposed, optimising the system's sensitivity to conductivity changes around the fingers. Finally, the system's detection accuracy is assessed, analysing how it performs in two possible use cases with multiple subjects.

4.1 EBI of Hand-to-Mouth Behaviours

4.1.1 Method

The first stage explored how a wearer's hand-to-mouth behaviours manifest in the measured EBI spectra. To this end, we have conducted an experiment with a single subject in an indoor working environment. The subject wore the device for 5 consecutive hours, and data were recorded as he naturally moved in this everyday setting. Periodically, the subject was asked to place any of his 5 fingers in his mouth. An external observer annotated the recorded data with the subject's activity and body postures. Throughout this experiment, the electrodes were placed in a single configuration on the subject's forearm, marked (10,1) as will later be detailed in Section 4.2.

We utilised the recorded data to narrow down the scope of movement for the next stages of this work. This satisfies experimental control and reproducibility. The subject's imitation of natural movement captures most possible orientations of the forearm, as well as the variance of bioelectrical measurements. Special attention was paid to ensure that the dataset also included situations where different body parts come in contact with the subject's fingers. These are important as they create different paths through which

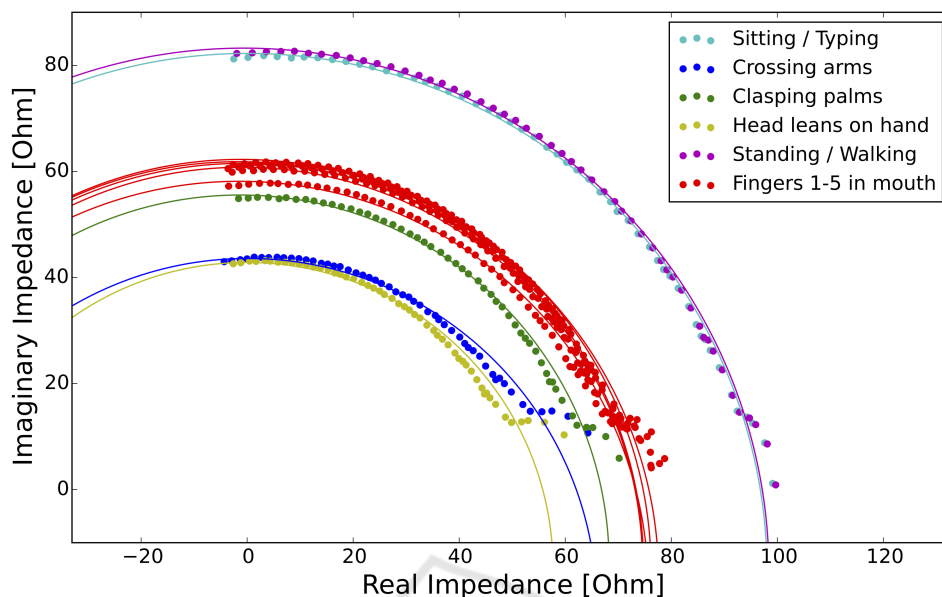


Figure 5: Sample EBI spectra generated by different body postures.

current can traverse, hence changing the characteristics of the EBI spectrum.

4.1.2 Results and Discussion

The EBI recordings supported our detection hypothesis. Noticeable differences were observed between situations where the subject’s hand was inside or outside the mouth. These manifest not only for specific exciting frequencies, but indeed throughout the entire spectrum. Fitting the spectra to a Cole-Cole model, we observe satisfactory differences in many of the model’s parameters. For example, the static impedance R_0 demonstrated a mean decrease of 32% when juxtaposing sitting with placing the thumb inside the mouth.

As mentioned, we were particularly interested in postures where different parts of the body come in contact with the fingers. We experimented with an assortment of these, in an attempt to identify ones which create unique EBI spectra. Our experimentation was limited to contact with body parts which were exposed wearing full clothing. We exemplify three such postures – crossing the arms, clasping the hands, and leaning the head on the hand – in Figure 5. EBI spectra from sitting, standing and placing any of the 5 fingers inside the mouth are overlaid for comparison.

Note that the EBI spectrum generated by clasping the hands resembles the ones generated by placing any finger inside the mouth. In some instances,

we noticed a circular spectrum almost identical to one from a hand-to-mouth posture. Our design trusts differences in the orientation components to settle these ambiguities.

Supporting this assumption, we analysed differences in the system’s detection accuracy, restricting it to rely only on a subset of components from the extracted features. Classification models were trained, as indicated in Section 3.2.3, for three cases: including only the orientation components, only the EBI components, or all available components. We then compared the models by their prediction accuracies, classifying an independent test set. This set was collected by repeating the aforementioned experiment with the same subject on a different day. Detection accuracies of 74.39%, 90.74% and 95.76% were obtained for each of the cases respectively. They imply the suitability of the different sensors for our detection task, and bolster our decision to combine them.

Finally, we studied the measured EBI spectra to select 10 representative body postures. The first five were tied with the placement of each of the wearer’s 5 fingers inside the mouth. The remaining postures were selected to generate the highest diversity in their corresponding EBI spectra. Crossing the arms, clasping the hands, and leaning the head on the hand were chosen by this criterion. Sitting still and standing were chosen due to their high rate of occurrence in the dataset. The last two represent EBI spectra similar to walking and typing which were the most frequently observed postures.

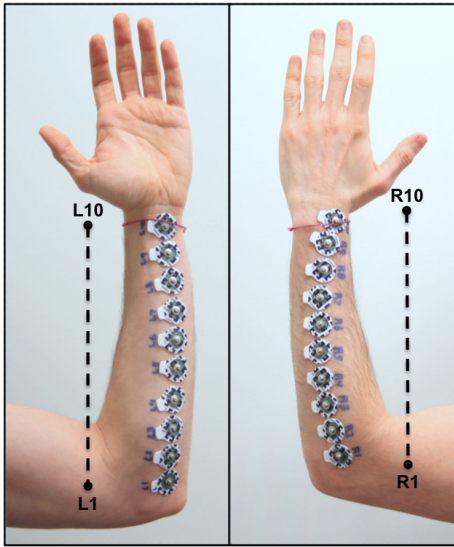


Figure 6: Electrode Instrumentation on a subject’s forearm used to evaluate an optimal positioning.

4.2 Positioning Optimisation

The second stage examined possible positioning of the electrodes on the arm. We limited the electrodes’ arrangements to the forearm, as only such arrangements would make practical sense for wearing or wiring in future applications.

4.2.1 Instrumentation

Figure 6 illustrates how a single subject was instrumented with 20 electrodes. In accord with the muscular anatomy of the forearm, the electrodes were lined in two bands. Electrodes following the posterior and anterior fascial compartments were marked with L and R respectively. Within each band, the electrodes were indexed from 1 to 10 starting from the elbow. The mean distance between the centre of consecutive electrodes was 2.66 cm.

Though each pair of electrodes is interchangeable up to a sign, and the PU and IC electrodes are interchangeable by the Reciprocity Theorem, it is still infeasible to exhaustively evaluate each of the 19,380 possible electrodes’ combinations. Therefore, we introduced a constraint that the PU and the IC electrodes are configured in a crossed alignment, in which the current density fields are intersecting with each other. Figure 7 illustrates such an arrangement. The IC (+) electrode and the PU (-) electrode were always assigned to the R band. Correspondingly, the IC (-) electrode and the PU (+) electrode were always assigned to the L band. Forming an intersection, the (+) pair and the (-) pair were restricted each to maintain the same index along the forearm. We use the elec-

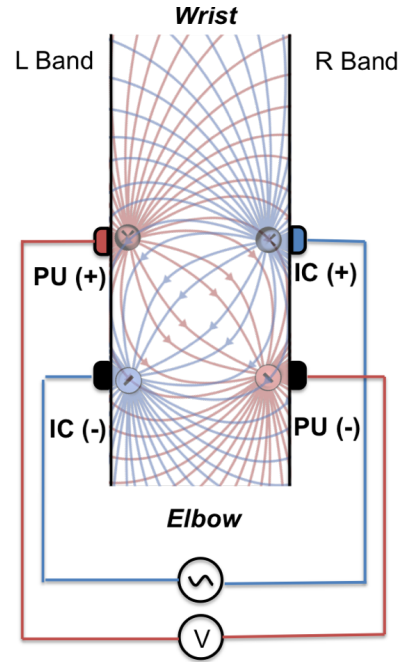


Figure 7: Schematic illustration of a crossed alignment of the electrodes. The propagation lines are a coarse approximation of the induced current density fields.

trodes’ indices to denote their configuration. Configuration (10,1), for example, marks an electrode configuration where that the (+) and (-) electrodes are in indices 10 and 1 respectively.

The assumption that the electrodes should be crossed is not verified within the scope of this work. However, it builds on our intuition of the current density fields (\vec{J}_{PU} and \vec{J}_{IC}) induced by the electrode system. A coarse approximation of these is visualised in Figure 7. Unlike most cases where conductivity changes are circumscribed by the electrodes, we would like to sense changes further away – in regions which are electrically coupled by the touch of the fingers. The crossed alignment, therefore, attempts to nullify the sensitivity field (dot product) inside the forearm by creating antiparallel components between the current density fields. On the other hand, it attempts to increase the sensitivity as we go further away – where the fields become more parallel.

4.2.2 Method

Our experiment varied the distance between the (+) and (-) electrodes, as well as their position along the forearm. Starting with configuration (10,1), we decreased the distance in three different ways, evaluating each configuration we encounter:

- Moving the (+) pair towards the elbow – from configuration (10,1) to (2,1).

- Moving the (-) pair towards the wrist – from configuration (10,1) to (10,9).
- Symmetrically moving both pairs towards each other – from configuration (10,1) to (6,5).

Each electrode configuration was evaluated by two separate metrics.

First, we compared the electrode configurations by their sensitivity to changes around the fingers. We calculated fluctuations in the measured impedance, as we alternated the body tissues that the fingers touch. By proxy, the more intense the fluctuations were, the more conductivity changes around the fingers contributed to the overall impedance. We used the 10 body postures to represent a fixed set of body tissues that come in contact with the fingers. For each electrode configuration, an identical number of samples were recorded in every posture. Fluctuations were then calculated by determining the coefficient of variation for the impedance magnitude in the recording:

$$\text{Coefficient of variation} = \frac{\text{std}\{|Z|\}}{\text{mean}\{|Z|\}} \quad (11)$$

Second, we assessed the electrode configurations by their predictive power. We held a second recording session with the same subject on a different day and re-instrumented his arm with the electrodes. The EBI spectra, generated by the various postures in this session, served as an independent test set. In this experiment, we only used the EBI features, discarding the orientation components. This way, our results only reflect the electrical dissimilarities between the electrode configurations. A classification model was trained and its optimal parameters were selected, using data only from the first recording session. The electrode configurations were compared with their prediction accuracy on the test set.

4.2.3 Results and Discussion

Figure 8 presents the coefficient of variation in the impedance magnitude obtained for each of the evaluated configurations. The presented results were derived from the impedance sampled with a single excitation frequency 52 kHz. Similar trends were observed measuring with other excitation frequencies within the system’s range.

It is evident that the coefficient of variation gradually increases as the electrodes in the pairs go further from each other. This finding is consistent with the theoretical sensitivity field model that was previously reviewed. The further the electrodes are, the vaster the current density fields spread inside the tissue, granting distant regions, such as around the fingers, higher relative contribution to the overall impedance. The

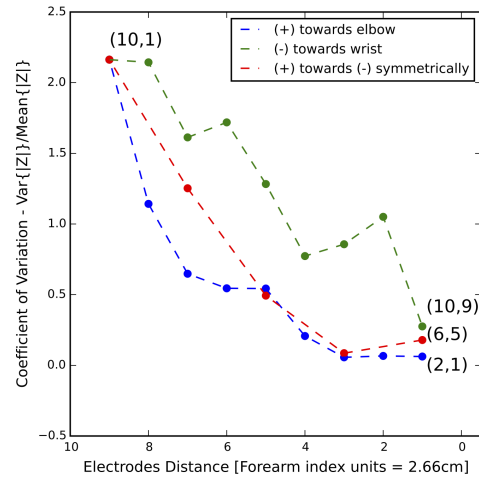


Figure 8: Coefficient of variation of the electrode configurations. The overlaid Indices pairs mark specific configurations of interest.

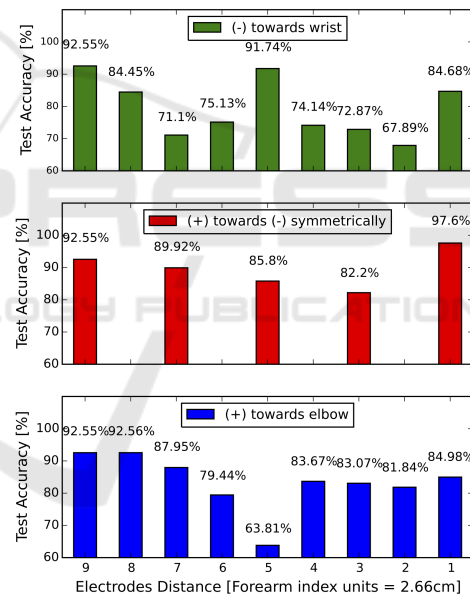


Figure 9: Predictive power of the electrode configurations.

region circumscribed by the electrodes also appears to have an effect on the coefficient of variation. This can be identified by studying the three graphs of Figure 8 against each other. Notice how the coefficient is higher when the circumscribed region is closer to the wrist rather than the elbow.

A second assessment was based on the configurations’ predictive power. Figure 9 plots the prediction accuracies achieved by the optimal classification algorithms after selecting their parameters. Configuration (10,1) yields powerful accuracy results as anticipated by its coefficient of variation. Nonetheless, the

best accuracy results (up to 97.6%) are encouragingly produced by configuration (6,5). As it is geometrically concentrated, such configuration can greatly facilitate a wearable design for practical applications.

We chose to proceed to the next stage of the evaluation with two electrode configurations: (10,1) due to its high coefficient of variation, and (6,5) due to its highest prediction accuracy. Detection accuracies from both are presented in the next section. We have also conducted experiments in which the IC and PU electrodes were not constrained to move symmetrically or together. These are not included here, since they revealed neither an observable trend nor higher sensitivity or accuracy.

4.3 Detection Accuracies

The final stage analysed aspects of our system's robustness, using data collected from multiple subjects on multiple days. Within the scope of this work, we explored two probable use cases – recurrent usage by the same user and encountering previously unseen users for the first time.

Recurrent Usage

First, we investigated situations when the same user wears the system on separate days. This captures plausible imprecisions from recurrent usage of the system, like inconsistent placement of the electrodes or misalignment of the IMU. Additionally, it accounts for transient changes in bioelectrical properties of the wearer's forearm. These may introduce perturbations of over 8% in an EBI spectrum (Gabriel et al., 1996). They are likely to stem from temporal variations in fluid volumes, blood pressure, body temperature, nervous activity, hydration, electro-dermal activity, etc. (Schwan, 1956)

Previously Unseen Users

Next, we examined how well the system generalises, performing on new subjects that have not been previously seen. This tests the system's tolerance to person-to-person differences, including differences in adipose, muscle and vasculature ratios, bone anatomy, sub-dermal water content and skin thickness. Such generalisation is of crucial importance for future design of practical applications. It involves implications on the system's capability to scale, deploying to a broad user base without a need to first collect their data and customise per individual.

4.3.1 Method

We recruited 15 subjects, 8 females and 7 males, to participate in a pilot user study. The study was approved by IRB protocol #1504007088. All subjects were between the ages 21 and 33 (mean=26.5; $\sigma = 3.6$). Their forearm length was measured and ranged between 22 cm and 27.5 cm (mean=24.8; $\sigma = 1.7$). They were asked for their weight, which varied between 47 kg and 77 kg (mean=58.3; $\sigma = 9.4$), and their height which varied between 155 cm and 185 cm (mean=169; $\sigma = 8.9$).

The subjects were invited twice to identical sessions on two separate days over the course of a single month. At the beginning of each session, they were instrumented with the electrode configurations selected in Section 4.2, on their left forearm. The subjects' forearm length was measured, and the placement of each electrode within the combination was scaled accordingly. Data were then recorded as the subjects were asked to pose in the postures from Section 4.1. Positive and negative labels were respectively assigned to postures when the subjects' hand was inside or outside their mouth.

An additional subtlety was introduced for postures in which the hand was placed inside the mouth. For each finger, measurements were taken both when the finger was fully sucked (up to the middle phalanx) and when its tip was only lightly touching the inner lip. This extension intends to simulate dissimilarities between thumb sucking and nail biting, and verifies the system's ability to detect both.

We framed the evaluation as a binary-classification problem, and report the system's detection performance in terms of classification accuracy. According to the evaluated use case, the dataset was partitioned for training, validation and testing purposes. Prior to any analysis, all data points were scaled using feature standardisation (Murphy, 2012). In total, we have collected 10,483 data points; Each of which is a Cole-Cole representation of an EBI spectrum, comprising 50 complex impedance measurements, and a corresponding orientation. 6,348 of the data points had a positive label and 4,134 had a negative one.

4.3.2 Results and Discussion

Recurrent Usage

First, we established how the system generalises when the same individual recurrently wears it. We did so by tailoring a classifier per user. The first recording session is randomly partitioned into training and validation datasets; Data recorded in the second session was held-out as a test set. We searched the param-

parameter space of each classification model, as indicated in Section 3.2.3. The combination of algorithm and parameters, which optimised the performance on the validation set, was selected as the subject’s personal classifier.

Table 2 reports the median, average and standard deviation of the test accuracy scores obtained by the personal classifiers. It determines that when recurrently worn by a user, our system is expected to accomplish a median accuracy of up to 92.95%, detecting hand-to-mouth behaviours. Encouragingly, electrode configuration (6,5) performed equally well as configuration (10,1). We regard the median statistic as the most relevant one for our evaluation, as it is the most robust to subjects which are potentially outliers in our user base.

We also examined a condition when the personal classifiers were constrained to train only with a specific classification model. In this case, only the model’s parameters were optimised per user. For configuration (10,1), similar median accuracies were obtained when using only KNN classifiers and when optimising with all possible models. The same emerges for configuration (6,5), constraining its training only to RF classifiers. This finding implies, that a single classification model may be used to generalise for recurrent usage. Relevant accuracy scores of these constrained cases are also presented in Table 2. We state, however, that a larger user base, which yields a smaller distance between the median and mean statistics, should be analysed to confirm such observations.

We highlight that these accuracies were obtained from training with a single usage. An alternative approach may suggest training with multiple usages to potentially increase detection accuracies. Despite its rigour, we chose the former, trusting that future applications would greatly benefit from requiring not more than a single usage for training. Otherwise, the users’ enrolment procedure can become cumbersome and impractical.

Previously Unseen Users

Next, we assessed the system’s performance when encountering subjects it has not previously seen. We partitioned our dataset, holding-out data points from 3 subjects (20%) for testing purposes. The remaining subjects were used for training and validation. Per subject, we included data points from the two recording sessions, diversifying the datasets with multiple usages. As in Section 3.2.3, the optimal combination of model and parameters was selected by its validation accuracy. Its test accuracy was derived from making predictions on the previously unseen test subjects.

Due to the relatively small number of subjects, the

Table 2: Detection accuracies for recurrent usage by the same user. Baseline accuracies were generated by a stratified dummy classifier.

Config.	Classifier	Test accuracies	
		Median	Mean (σ)
(10,1)	Optimal	92.95%	90.98% (6.3)
	KNN	92.15%	88.31% (6.5)
	Baseline	53.42%	53.4% (3.3)
(6,5)	Optimal	91.6%	90.96% (7.1)
	RF	91.6%	85.47% (9.2)
	Baseline	53.83%	52.9% (3.2)

particular assignment of test subjects is likely to affect the resulting accuracy. Therefore, instead of randomly selecting an assignment, we chose to exhaustively search through all possible ones. This approach gains a more credible insight into the system’s capability to generalise person-to-person differences. It protects from arbitrary bias that may mislead our evaluation. In Table 3 we report the median, average and standard deviation of the test accuracies obtained by the optimal classifiers as they predicted for their corresponding assignments.

Our results suggest that our system is likely to produce an accuracy of 87.5%, predicting hand-to-mouth behaviours for subjects it has never before seen. Regrettably, these high accuracy results were only produced for electrode configuration (10,1). Configuration (6,5), performed significantly worse, yielding median accuracy of only 79.69%. For two thirds of the subjects’ assignments, RF was selected as the optimal model regardless of the electrode configuration. This indicates that the RF model may be the fittest one for this kind of task.

We continued our analysis, exploring the possibility of adjusting the classifiers by the users’ physicality. We clustered our subjects into two groups, based on their gender, age, weight, height and forearm length. A 2-component Gaussian Mixture Model clustering technique was used, utilising Expectation Minimisation (EM) to fit the subjects’ physicality measurements. Except for a single subject, the resulting clusters overlapped with the subjects’ genders. We mark the clusters F and M respectively, according to the majority of females and males in their population.

We repeated the previous analysis for each of the clusters separately. To preserve the ratio from the former analysis, a single subject was held-out for testing from the M cluster, and a pair were held-out from cluster F. The median, mean and standard deviations of the test accuracies are also presented in Table 3 for comparison.

Successful improvements of up to 3% in the me-

Table 3: Detection accuracies predicting for users that have not been previously seen.

Config.	Dataset	Test accuracies	
		Median	Mean (σ)
(10,1)	All	87.5%	87.21% (5.1)
	Cluster F	90.52%	87.96% (6.3)
	Cluster M	90.72%	88.14% (6.1)
(6,5)	All	79.69%	79.81% (5.5)
	Cluster F	83.25%	84.7% (7.9)
	Cluster M	82.09%	82.82% (4.4)

dian accuracies were observed employing this approach. They suggest that person-to-person differences can be ameliorated by an a priori query of the user's physicality. It could be possible for future applications to optimise multiple classifiers based on physical typecasts. A mixture of those can then be employed for unseen users, consequently improving expected accuracy. The mixture should be weighted by the users' similarity to the physical typecasts produced while training. Future work should further explore this possibility with a larger user base.

5 CONCLUSION

This work has presented a wearable system detecting hand-to-mouth behaviours in real time. It demonstrated a novel sensing method, associating bioimpedance spectra with the wearer's behaviour. The relationship between the system's sensitivity and its electrode configuration has been investigated, optimising the latter to potentially increase detection efficacy. The system's performance has been evaluated, examining use cases where it is recurrently worn and where it encounters a user for the first time. It achieved a median detection accuracy of 92% and 90%, respectively. These, we aspire, may be sufficient to guide new directions in the treatment and monitoring of compulsive hand-to-mouth habits.

REFERENCES

- Ayllón, D. et al. (2009). Cole equation and parameter estimation from elect. bioimpedance spectroscopy measurements—a comparative study. In *Eng. in Medicine and Biology Soc., 2009. EMBC 2009. Ann. Int. Conf. of the IEEE*, pages 3779–3782. IEEE.
- Azrin, N. and Nunn, R. (1973). Habit-reversal: A method of eliminating nervous habits and tics. *Behaviour research and therapy*, 11(4):619–628.
- Baydaş, B. et al. (2007). Effect of a chronic nail-biting habit on the oral carriage of enterobacteriaceae. *Oral microbiology and immunology*, 22(1):1–4.
- Blomqvist, K. et al. (2012). An open-source hardware for elect. bioimpedance measurement. In *Electron. Conf. (BEC), 2012 13th Biennial Baltic*, pages 199–202. IEEE.
- Cole, K. S. and Cole, R. H. (1941). Dispersion and absorption in dielectrics i. alternating current characteristics. *The J. of Chemical Physics*, 9(4):341–351.
- Cornish, B. et al. (1993). Improved prediction of extracellular and total body water using impedance loci generated by multiple frequency bioelect. impedance analysis. *Physics in medicine and biology*, 38(3):337.
- Gabriel, S. et al. (1996). The dielectric properties of biological tissues: Ii. measurements in the frequency range 10 hz to 20 ghz. *Physics in medicine and biology*, 41(11):2251.
- Geselowitz, D. B. (1971). An application of electrocardiographic lead theory to impedance plethysmography. *Biomedical Eng., IEEE Trans. on*, (1):38–41.
- Madgwick, S. O. et al. (2011). Estimation of imu and marg orientation using a gradient descent algorithm. In *Rehabilitation Robotics (ICORR), 2011 IEEE Int. Conf. on*, pages 1–7. IEEE.
- Martinsen, O. G. et al. (2007). Utilizing characteristic elect. properties of the epidermal skin layers to detect fake fingers in biometric fingerprint systems - a pilot study. *Biomedical Eng., IEEE Trans. on*, 54(5):891–894.
- Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.
- Pacan, P. et al. (2014). Onychophagia is associated with impairment of quality of life. *Acta dermatovenereologica*, 94(6):703–706.
- Pedregosa, F. et al. (2011). Scikit-learn: Machine learning in Python. *J. of Machine Learning Research*, 12:2825–2830.
- Sato, M. et al. (2012). Touché: enhancing touch interaction on humans, screens, liquids, and everyday objects. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, pages 483–492. ACM.
- Schwan, H. P. (1956). Elect. properties of tissue and cell suspensions. *Advances in biological and medical physics*, 5:147–209.
- Seoane, F. et al. (2008). An analog front-end enables elect. impedance spectroscopy system on-chip for biomedical applicat. *Physiological measurement*, 29(6):S267.
- Seppä, V. et al. (2007). Measuring respirational parameters with a wearable bioimpedance device. In *13th Int. Conf. on Elect. Bioimpedance and the 8th Conf. on Elect. Impedance Tomography*, pages 663–666. Springer.
- Tanaka, O. M. et al. (2008). Nailbiting, or onychophagia: a special habit. *American J. of Orthodontics and Dentofacial Orthopedics*, 134(2):305–308.