BIOMETRIC IDENTIFI-CATION

WHITE PAPER ON METHOD COMPARISON

This document contains details about some biometric methods that are used for the purposes of identification verification based on lab experiments, including finger vein, bioimpedance, and electrocardiogram (ECG).

Taras Kurnyts'kyi & Vladyslav Tsybul'nyk

softserve

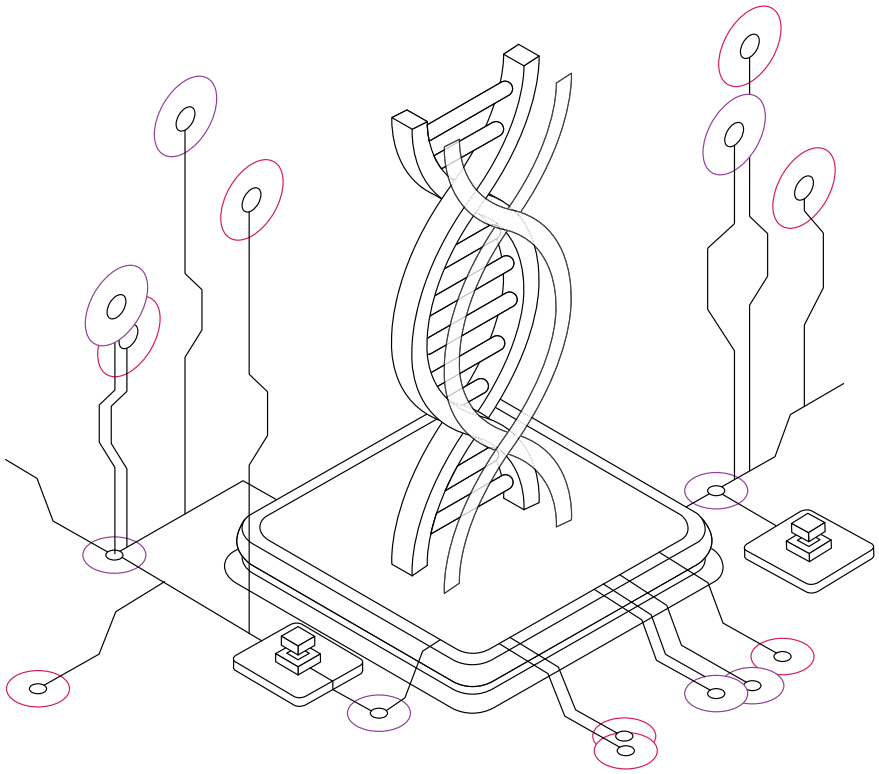

BIOMETRIC IDENTIFICATION. WHY NOW?

Biometric verification is any means by which a person can be exclusively identified by evaluating one or more distinguishing biological traits, such as their facial or body features, including finger veins, DNA, and signatures.

With the introduction of computerized databases and the digitization of analog data, biometric verification has become significantly more innovative, allowing for almost instantaneous identification. Therefore, the fundamentals of these biometric methods are hardly ever needed to be treated or classed as new or modern. For example, iris-pattern and retina-pattern authentication methods are already used in some bank ATM machines.

As devices have become more affordable and user-friendly, the use of biometric identification has become common, while at the same time software algorithms have become smarter and considerably more reliable too.

Times have certainly changed. A few decades ago, when referring to printed currency notes, people would have said “I have”, as in to say that they “have” the money with them. Nowadays, this can also be expanded to a “I have and I know” statement. What do we mean by this? Well, for example, let’s say we have a credit card and a password that accompanies it. Nowadays, that narrative of “I have” can be used to refer to the credit card, while the “I know” can be used to refer to the password that goes with the credit card (as in they “know” what the password is).

Possessing a credit card is definitely more convenient than carrying cash; however, there are still some disadvantages of carrying one, such as it is easy to lose one. But what if the "I know" part of the algorithm (i.e. a password in this case) is stolen? Then the problem becomes rather serious and complicated. This is where biometric identification can help because the information can store data and can help the user to get back any lost data.

Biometrics in personal identification has become quite attractive because it operates within the "I have" model. Effectively, biometric identification makes it easier for individuals to have their data carried along since it is naturally embedded and, thus, it is very difficult to lose. No matter what biometric methodology is used, the identification verification process remains the same wherever you go because a record of a person's unique characteristic is captured and kept in a database.

Whenever that stored biometric identification data is required (i.e. the "I have") , a new record is captured and compared with the previous record in the database. If, for whatever reason, the data in the new record matches that in the database record, the person's identity is confirmed.

Table 1 below shows a list of some of the most popular biometric identification features, and their types. This list is not comprehensive by all means.

FEATURE	TYPE	
	CHARACTERISTIC	BEHAVIORAL
Face	●	
Fingerprint	●	
Iris	●	
Retina	●	
DNA	●	
Finger vein / palm vein	●	
Voice	●	
Electrocardiogram	●	
Bioimpedance	●	
Gate		●
Keystroke		●
Signature		●

Table 1: List of the major biometric methods

Each biometric method has advantages and disadvantages, as well as different fields of application. For example, DNA profiling is undoubtedly the most accurate one, but it's not suitable for widespread applications because the identification process takes too long, and requires the usage of expensive equipment.

When we look at finger vein recognition, it's worth noting that finger veins are the tiny blood vessels inside your finger, and are laid out in a pattern which is unique to every individual. Finger vein based identity authentication systems might soon surge ahead based on the unique advantages they offer. Even though finger vein identification may be less accurate than DNA profiling, it still possesses other highly competitive features, such as a shorter identification time, and usage of user-friendly and inexpensive equipment.

Retina and iris scans are almost as accurate as DNA profiling but people are uncomfortable with this approach because they feel it's too intrusive. The point being that every biometric identification method comfortably finds its own application area.

Here is a list of key requirements that biometric methods must satisfy in order to be successful:

- **Uniqueness:** biometric feature must be unique for every person
- **Permanence:** should not change with time
- **Anti-forgery:** difficult to falsify
- **Enrollment time:** as minimal as possible
- **Identification time:** as minimal as possible
- **Accepted by a user:** should be user-friendly

Among these requirements, the first two, uniqueness and permanence, are probably the most important ones because they both must be intrinsically unique when it comes to biometric data, and the equipment must be able to permanently detect this uniqueness.

The same principle applies to the permanence feature. The only difference being that for this feature to be applied successfully, the measuring equipment must permanently produce a quality output as well, which is a serious requirement and isn't always met. For example, certain noises can influence ECG measurements, while different light conditions may provide a negative impact on the quality of the appearance of finger veins. These are minor but important issues which must be taken into account when developing biometric devices.

Preventing Forgery

For any biometric method to be successful, it must be difficult to forge. There has to be an ability to pick-up, detect, store, and reproduce any stolen biometric data so that it can be tracked down and prevented from happening again. While fingerprint images are easy to forge, finger vein recognition is challenging because it's not easy to pick up a finger vein image as it is easy to pick up a fingerprint. In any case, a compromise between key requirements of a biometric method determines its place, role, and success on the biometric methods field.

Enrollment and Identification Time

Enrollment time is the time needed to take and enter a person's data into a database, and the identification time is the time it takes to acquire and to compare an actual person's data with the one stored in a system's database.

Both enrollment and identification times differ for various biometrics—they are very short for fingerprint, finger vein, and face recognition, while longer for DNA profiling. Generally, the shorter this time span, the better it is.

Along with the enrollment and identification time, biometric method acceptance by a user is the other criterion determining a method's success. Most characteristic methods from the list provided in Table 1 have a rather high user acceptance rate, except for the iris and retina scans. Ultimately, the success of any specific biometric authentication system is determined by a goal to achieve and by a field of application as well.

FINGER VEIN RECOGNITION

Finger vein recognition is the latest and most reliable biometric method. This authentication system works by using the vein patterns in an individual's fingers to verify their identity. When an individual's finger is placed on the vein recognition device, the vein patterns are scanned and recorded onto, and then they are matched with the records in a database to confirm the identity of that individual.

As we have already seen, the performance of biometric methods is evaluated by a set of key factors. Here are some reasons why finger vein recognition is good for biometric identification:

- **Uniqueness:** research shows that finger vein patterns are unique for every person
- **Permanence:** research shows that the human finger vein pattern is very stable over the years. The only factor threatening this is if there is a finger injury. However, multiple fingers can be used for enrollment and identification.
- **Anti-forgery:** veins are hidden in the human body, so it makes it extremely difficult to forge a finger vein
- **Enrollment time:** around 30 seconds
- **Identification time:** only a couple of seconds

Unlike fingerprinting, where the biological information being scanned is on the outside of the body, finger vein recognition scans information on inside of the body and therefore makes forgery extremely difficult. Finger vein recognition thus serves as a highly secure

form of personal recognition.

This point alone makes it a credible point of why finger vein recognition is becoming popular as a reliable and a user-friendly biometric recognition method.

Approach

The main idea behind finger vein recognition technology lies in the fact that hemoglobin in the blood vessels intensively absorbs electromagnetic radiation in the range around 750-1000nm, known as the near-infrared radiation (NIR), while the finger tissue and bones do not.

A clear image of the veins can be seen by placing an infrared camera on one side of the finger, which is illuminated by the NIR source. When the image is acquired, it is then pre-processed by the correspondent algorithms to extract only the area of the finger (i.e. the so-called Region of interest, or else known as the ROI).

After that feature, extraction algorithms are applied to obtain feature values specific for a finger. Such sets of feature values for a person's fingers are stored in a database. Finally, to identify a person, the obtained finger vein image is matched with those stored in the database and eventually a final decision is made on the identity of the individual. The process for carrying out finger vein recognition is shown below in **Figure 1.**

Figure 1: A typical process for carrying out finger vein recognition

State-of-the-Art Technology

Imaging techniques

There are a couple of finger vein imaging techniques:

One of them uses different Near-Infrared Lights (i.e. LEDs) as a light source, and positions the infrared sensor camera (known as a CCD Camera) in various angles pointed towards the finger. The LED and the CCD Camera should be placed on both sides of the finger, and in this case the veins can detected by the transmitted light. Please see **Figure 2b.**

With the other technique, an image can be taken in a reflected light when the LED source, with the camera located on the same side with respect to a finger (please see **Figure 2a**). The latter approach makes installation more compact, though the quality of the image is normally lower compared to that taken in a transmitted light.

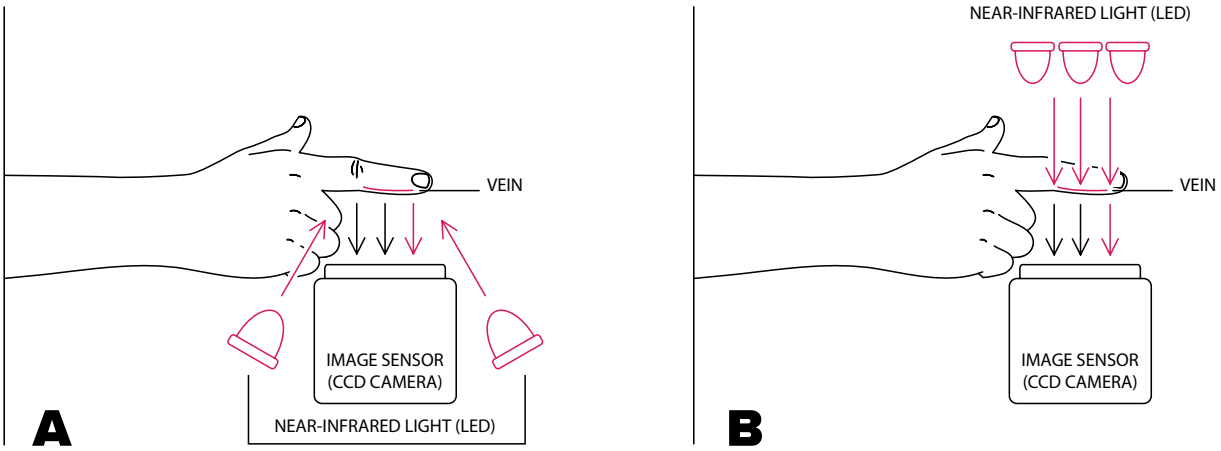

Figure 2: Finger vein image taken in a) reflected light, b) transmitted light

Identification algorithms

There are quite a number of advanced finger vein recognition algorithms. There are some great insights on the subject in the "A survey of finger vein recognition" report by the team at the School of Computer Science and Technology at Shandong University.**[2]**

Globally, these algorithms can be divided into three classes:

- Works with images taken from local data. This is the Local Binary Pattern (LBP) method and there are a number of its modifications that represent this class
- Works with an image as a whole and operates on the image's integral characteristics. Algorithms based on the Radon transform, wavelet transform, min/max curvature belong to this class
- Combines both the local and the integral approaches, and usually achieves better results

In recent years, more powerful machine learning algorithms, including neural networks, have also become key players in the field. Rapidly growing literature on the subject can be found in this paper published by the IEEE.**[1]**

Industrial solutions

There are a number of industrial systems on the market using finger vein recognition as an identification method. For example, industrial biometric security solutions developed by Hitachi and Fujitsu provide finger vein recognition technology that is used in the following applications:

- Logical access applications
- Applications for physical access through finger vein recognition systems
- Vein pattern applications for ATM & banking
- Embedded applications

Despite a number of ready industrial solutions on the market, our primary goal was to develop a simple and an inexpensive hardware solution that gives a true “feeling” of the subject.

Our methodology

Hardware setup

Figure 3 below shows a picture of a finger vein recognition setup. The bottom part of the device contains LED sources illuminating the finger from beneath. Infrared light from four 840 nm LEDs goes through the finger and is partially absorbed by finger veins caught by the infrared camera on the top. There is also a circuitry used in this device setup. The black box that contains the device is printed using our 3D printer.

Figure 3: Finger vein image taking setup

An inexpensive infrared CCD camera was used in our setup, along with an infrared filter and a “fish-eye” optic lens. The light was provided by 840nm LEDs. During the development of the device, the following important observations were made:

1. The intensity of the LEDs must be controlled so that it achieves the same brightness when we apply different parameters, such as if the thickness of the finger and its shape etc.

2. Better finger images are obtained if the LED lightens up the finger from above. In the device shown in Figure 3, the fingers are still lightened up from beneath.

In **Figure 4** below, raw images of two different fingers are shown. We can see that the quality of the images is good and quite sufficient for carrying out a successful finger vein recognition.

You can clearly see the network of veins. However, the main challenge is that our camera doesn't always generate high-quality images. Sometimes the images are either too bright or too dark, or sometimes they are just of an unacceptable quality. In any case, the camera cannot be controlled. This means that we should always use a camera of a high quality.

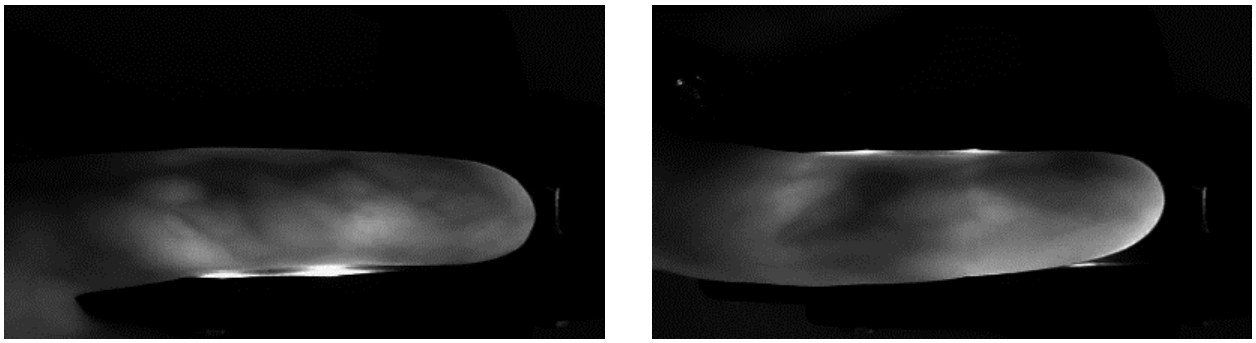

Figure 4: Raw images of different fingers

In **Figure 5** images of different fingers after pre-processing are shown. The pre-processing stage includes:

- Gray normalization (equalizing) of the image
- Extracting the region of interest (ROI)
- Image resizing to end up with a reasonable final image size (in our case it is 30x80 pixels)

It can be observed that pre-processed images are still of a high-quality because the veins are clearly visible.

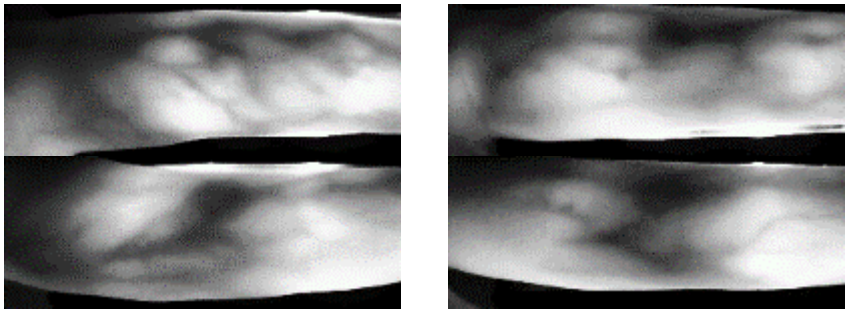

Figure 5: Pre-processed images of different fingers

PBBM and OpenCV C++

As briefly stated above, algorithms used in finger vein recognition are roughly classified as those that are analyzing local image features, integral features, or both. Local methods are good for initial analysis; however, technologically, high reliability can only be achieved with the methods that rely on local and integral approaches.

PBBM

After a number of experiments we made the decision to use a personalized best bit map (PBBM, see **[3]**) with some modifications.

The PBBM method partially addresses one of the key problems related to taking the correct photo of the finger, and how it differs considerably and can affect the end result. If the image of the finger is not exactly the same as in reality, then this eventually lowers the identification rates.

OpenCV C++

In time, we also developed a customized OpenCV C++ application implementing the full cycle of finger vein recognition. It starts by taking an image using a camera, followed by pre-processing, and eventually running through a trained model and making the final decision on the recognition of the finger vein. The application also allows a person's identification to be enrolled and stored.

The application was tested on a group of 12 people. Before the enrollment, the each person was asked to place their forefinger on the device. After an expected unsuccessful login, the individual was successfully enrolled. The enrollment included taking four images of the finger and lasted 20-30 seconds. After that, the person was asked to login into the system again.

It must be said that not many tests have been carried out so far.

However, the preliminary results are optimistic: there were a total of more than 200 attempts made by people to log in using their finger, by both enrolled and non-enrolled people.

In around 95% of the cases, the system reacted correctly. That is to say that access was granted for enrolled people and denied for those who were not enrolled.

We achieved a False Acceptance Rate (FAR) of 1% and a False Recognition Rate (FRR) of 5%. It should be noted that this was our first attempt, and because of this, there are areas in which we can improve. Namely in:

1. Hardware: consider having a high-quality camera that can take clear high-resolution images

2. Software: we can combine local and integral algorithms

Conclusion

Though a number of finger vein recognition devices are available on the market, in a rather short time we have developed a customized and inexpensive solution from scratch that produces good results. We also proposed some modifications to the existing algorithms that improved the identification rate. This proves that even with a limited time and resources we can develop a competitive customized solution.

BIOIMPEDANCE ANALYSIS

In general, the human body is able to conduct electric current and electromagnetic waves. An individual's reaction to this, that is to say a human body's resistance to electric current depends on individuals as everyone is different. If we are able to detect the difference in this resistance, we're potentially able to determine a person by their individual electric response.

This is where the Bioelectrical Impedance Analysis (BIA) technique, or sometimes known as the Bioimpedance Analysis, is used to estimate the body's composition. With BIA you can get a quick overview of the water and fat percentage in a body. Apart from being quick it is a safe technique and one that's become very popular to its ease of use.

The key reason why we paid attention to BIA is because it looked possible to implement it as an example of the so-called 'edge computing' when data was aquired.

Let's consider how BIA satisfies the key requirements of biometrics:

- **Uniqueness:** investigation shows that an individual's bioimpedance is not as unique as most other biometric characteristics

- **Permanence:** it's been research and proven that an individual's bioimpedance is not quite stable over time. It can change during the day and depends on what they eat and drink.

- **Anti-forgery:** it's quite difficult to forge any bioimpedance data since picking it up is quite an arduous task

- **Enrollment time:** 3-5 minutes
- **Identification time:** a couple of seconds
- **User acceptance:** high

Approach

Practical BIA consists of sending short electric pulses within a certain frequency range to a human body and measuring the response (i.e. the bioimpedance). The response is measured as a function of the specific tissue type (i.e. blood, muscle, bone, etc.) and its state, as well as anatomic configuration, which is determined by the orientation and quantity of the bone and tissues.

To cover the wider response area and thus increase the ability, the pulses are usually generated at a frequency range within 10-10000Hz. Then things such as the dissimilarities in bone structure, muscle density, fat content and the layout of blood vessels are expected to result in slight differences in the signals becoming weak at different frequencies This makes it possible to identify an individual.

Then the response at each frequency is subjected to the percentage of Fat-Free Mass (FFT), as well as active and reactive resistances (forming together a bioimpedance) are eventually calculated. If 50 different frequencies are chosen, then 100 samples (50 samples of active resistance and 50 samples of reactive resistance) make up a raw bioimpedance signal.

From the raw signal, a set of the selected features (usually up to ten) are extracted. Features are then inserted into a training model to determine the characteristic values for the certain individual. The decision to finally identify is made by matching the actual values during the identification session with those that are stored in the trained model.

State-of-the-Art Technology

BIA is mostly used in medicine for body composition measurements and healthcare assessment systems, and in particular for disease prognosis and the monitoring of vital status of a body.**[4]**

There are three types of BIA:

Single Frequency–pulse of a single selected frequency is used

Multiple frequency–pulses of a number of different frequencies are applied

Bioimpedance spectroscopy–deals with a broad range of frequencies

An individual's distinguishing ability is expected to be larger in the third case but the cost of the corresponding equipment is definitely higher. Choice of one of above BIA variations depends on the objectives of the specific case.

Currently a number of BIA devices are developed for medical purposes. They usually use many electrodes applied onto different parts of a body. Lately, body-area networks of wearable devices are becoming increasingly popular, especially in the healthcare, and entertainment industries.

These devices are able to discover each other, recognizing that they are on the same person, and are able to establish a communication channel. But the most important thing about BIA is that it must recognize a wearer.

There are a couple of solutions applying BIA in wearable devices (see sources **[6]** and **[7]**). One of them deals with a wrist device that has twelve electrodes embedded in it. The main part of the device is the impedance analyzer chip. The data collected from the electrodes is processed by the chip, and then "raw" bioimpedance samples are received as an output. A feature vector is extracted from the received raw data. Known classifiers, such as Naive Bayes, and the Support Vector Machine are eventually used for identification.

To sum up, it can be stated that using BIA in wearable or otherwise embedded devices is still in its infancy and the research carried out in this field is relevant.

Our research

We developed a hardware solution that consisted of having a STM32L476 ultra-low-power microcontroller (MCU) and an AD5933 impedance analyzer. Our goal was to have a small-sized hardware solution, preferably the size of a payment credit card. After a few iterations, we came up with an agreed design.

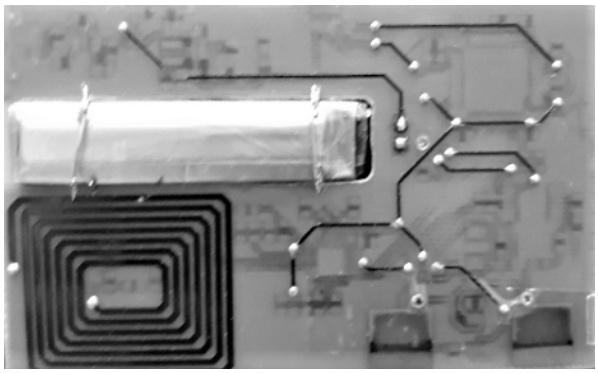

Figure 6: Credit card size device implementing full cycle of BIA

Due to the limited size of the card, it was fitted with only two electrodes. To start the bioimpedance data acquisition session, we attached one finger of each hand to the each of the electrodes on the card. The card detected the fingers and started to collect data. The impedance analyzer generated short pulses on 20 different frequencies in the range of 2-3 kHz and then analyzed the response. It applied a FFT analysis to respond to each frequency and eventually calculated the bioimpedance on the specific frequency.

The total response consisted of 20 active resistance samples and 20 reactive resistance samples. Hence, forming a raw biometric signal. Afterwards this data was sent to the MCU to be processed.

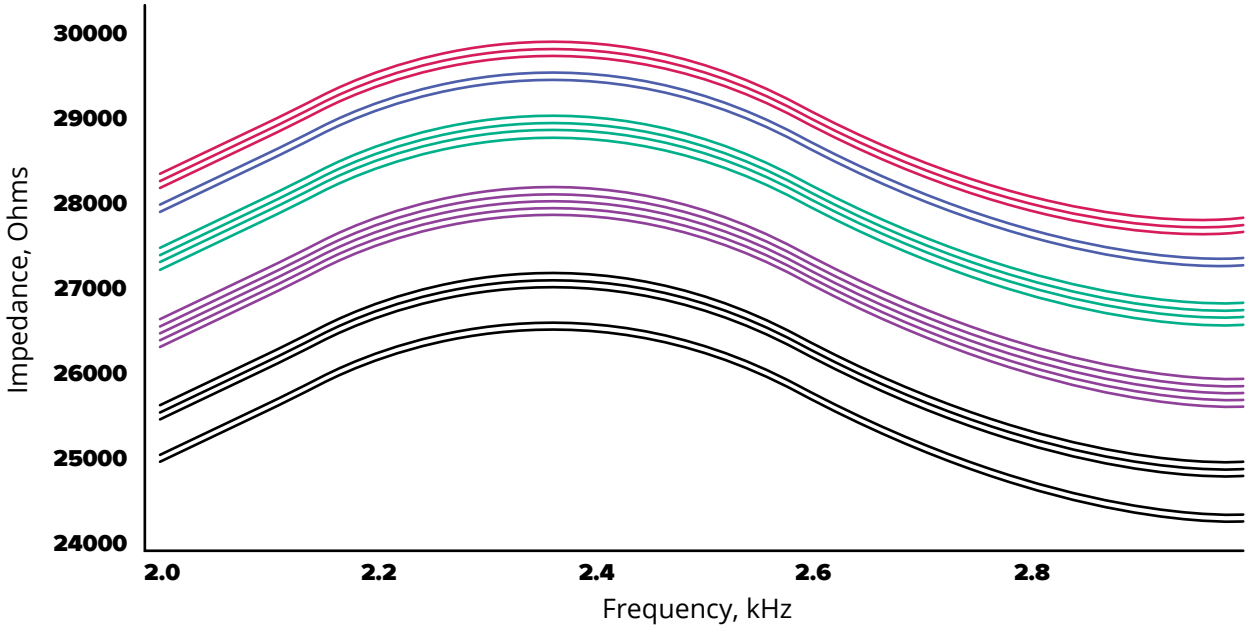

Figure 7: Impedance for different persons

Custom firmware running on a MCU was developed. It extracted the feature vector from the received raw bioimpedance signal, which were calculated for each frequency. In our case, the feature vector consisted of five components, namely minimum and maximum values of raw signal, signal power, and the kurtosis and the skewness.

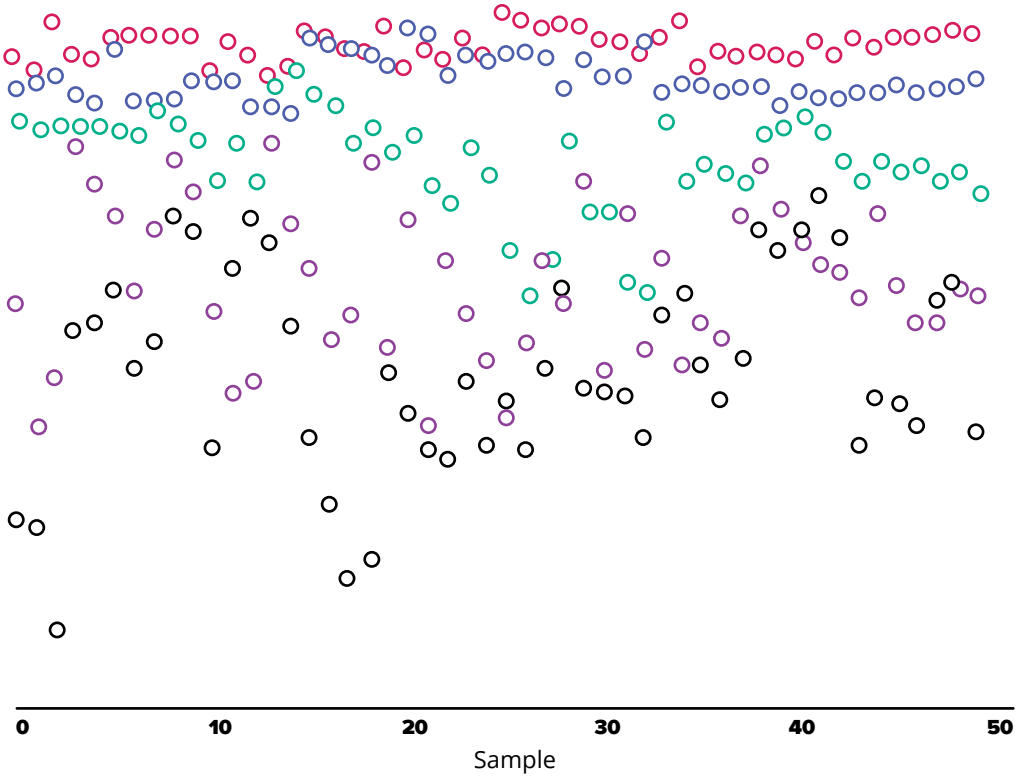

Figure 8: Features distribution

The training stage, which was carried out on PC to offload calculation from the card, included several sessions that collected sufficient raw bioimpedance data for each person. When the training was complete, the training model parameters were transferred to the card.

During the identification stage the feature vector extracted from actual bioimpedance signal runs through the classifier (the latter uses characteristics from the trained model) and eventually final decision is made. We tested several different classifiers, namely Support Vector Machine, Simple Perceptron and Naive Bayes. The last one proved to be the most effective.

Tests showed that the BIA is not as successful as finger vein recognition. The system reacted correctly in around 65% percent of the identification attempts. Effectively, access was granted for people who were enrolled (i.e. people whose data was in our database), and denied for those who were not enrolled. For finger vein recognition we had this rate leveling at 95%.

We found out that for a group of five individuals whose bioimpedance sufficiently differs, the BIA identification worked at an acceptable level. The larger the number of participants, the lower the rate of correct identification.

Conclusion

Two important observations regarding BIA were noticed. Either the bioimpedance data is not unique enough or the bioimpedance data is not stable enough.

Therefore, we can say that BIA by itself cannot reliably distinguish persons (see, for example **[5]**). It could be; however, used in a fusion of biometric solutions where BIA would play a secondary role in identification of an individual.

ELECTROCARDIOGRAM

The human heart is a complex muscular organ that uses electricity for activating different muscles. An electrocardiogram (ECG) is a recording of the electrical activity of the heart's rhythm. Sensors attached to the skin can be used to detect electrical signals produced by your heart each time it beats. The structure of the heart differs for each individual and the ability to distinguish ECG results for different people is limited by a lot of factors.

Let’s consider in more details how the ECG identification meets the principal requirements for biometric identification methods:

- **Uniqueness:** is not absolutely unique and has limited possibilities of precise measurement
- **Permanence:** is not stable over years, and could be significantly influenced by changes in lifestyle
- **Anti-forgery:** precise enough measurements could be done only with direct contact (or at least at a very short distance) during some period of time. Not easy to forge an ECG.
- **Enrollment time:** dozens of minutes
- **Identification time:** one heart beat (approx. 1 second), practically – dozens of seconds

ECG-based identification has the same advantages as vein recognition technology. Essentially, biometric data is not left imprinted on any device or everyday objects, and additionally it is completely passive as no additional electricity is applied to the human body.

Approach

The basis of ECG identification is a QRS complex–a name for the combination of three of the graphical deflections seen on a typical ECG heartbeat. It is usually the central and most visually obvious part of the chart. In other words, it's the main spike seen on an ECG line, as seen in **Figure 9.** It corresponds to the depolarization of the right and left ventricles of the human heart and contraction of the large ventricular muscles.

The peaking of the R is not fixed because it purely depends on the heart rate, and the waveform of the QRS complex itself is stable and could be used for identification.

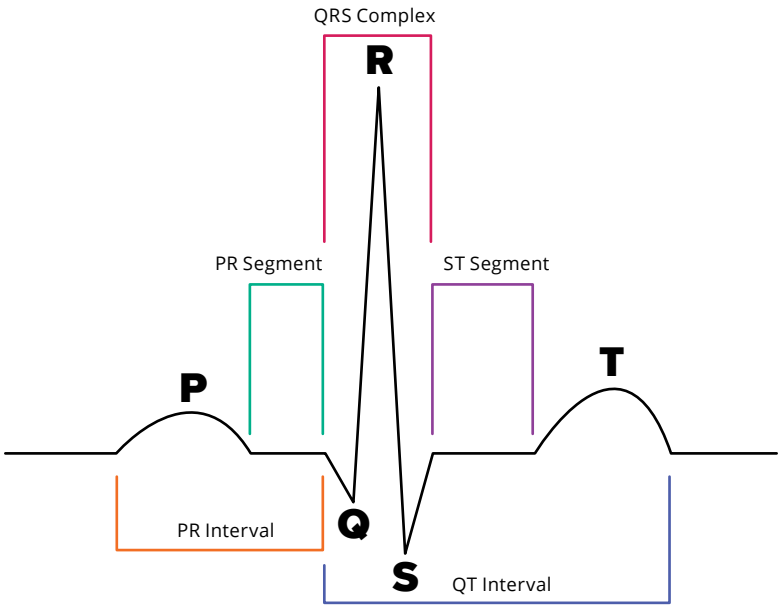

Figure 9. QRS complex

Theoretically, one 'beat' (i.e. a waveform) could be used for identification. Practically, electric noise is always present during the measurement stage (i.e. environmental noise, muscle noise, and equipment noise).

It should be noted that several beats are needed for the recording of a reliable waveform. Different methods for noise reduction should also be explored, and these can use machine learning **[8]** and other methods.

There are several methods that can be used for achieving a reliable and accurate identification, including traditional classifiers and deep neural networks.

Our research

We chose a case where it's possible to measure the body's ECG signal in an interesting way – by using a car driver's fingers, as shown in Photo 1 below. As a result, it was possible to identify a driver, which was the second factor together with having a key, and we continuously monitored the driver's presence to avoid a substitution.

We developed a prototype (Biolock) where the ECG was captured by electrodes embedded into the steering wheel. As an added advantage, no additional interaction with the user was required.

Photo 1: How we measured the ECG using a biolock in a steering wheel

The Biolock pipeline process is seen in **Figure 10** below:

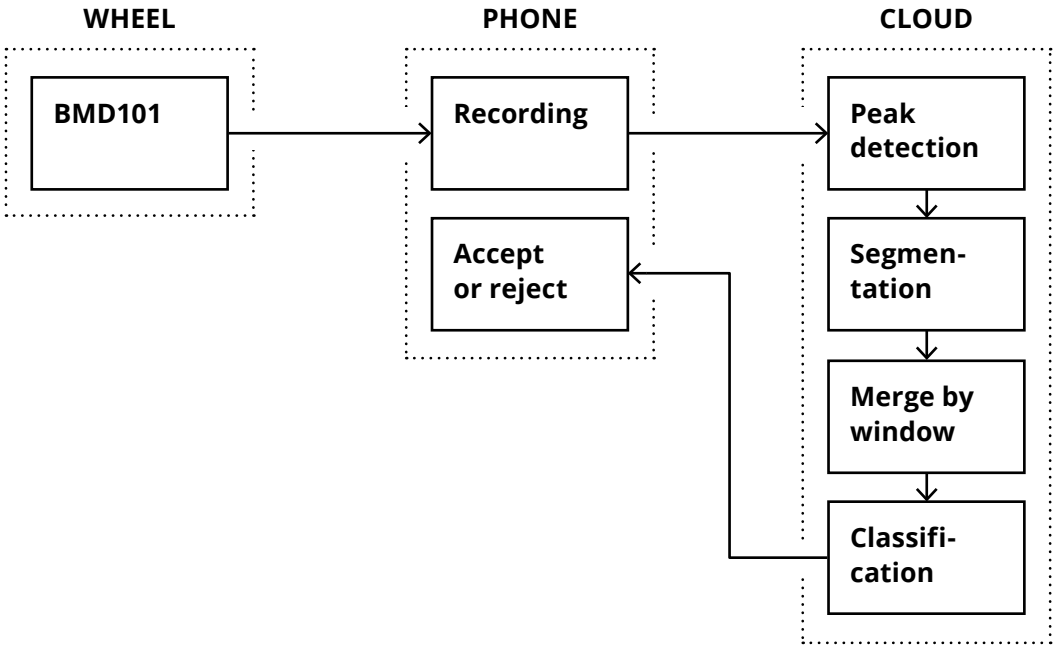

Figure 10: Biolock pipeline

A measured signal was streamed via Bluetooth to a mobile phone, where the application ran. The application collected the ECG in a 15-second buffer and sent data to the cloud for processing and recognition.

Firstly, R-peaks were detected using the Pan Tompkins algorithm. Then, the recorded signal was split on separate QRS-complexes (neighbor of R-peak, as shown on **Figure 7**). The extracted complexes were merged into one by averaging in order to reduce noise levels and artifacts.

After normalization was applied, we obtained QRS complex and passed to the neural network for training (enrollment stage) or for the prediction (verification stage).

As a result, we managed to satisfy two of our main criterias: the recognition time was less than 30 seconds and the accuracy was higher than 90% for a group of 20 people.

Our Biolock demo was one of the finalists at SXSW 2017 **[9]** in the “Smart City” category.

Conclusions

ECG identification has similar problems to Bioimpedance identification because:

- it is not unique in a certain way. Why? For example: if we have two people doing the same activity (e.g. running or hiking), for a long period of time (i.e. many years!), then their heart’s behaviors would become similar.

- it changes with time

Nevertheless, ECG identification works great in cases of continuous monitoring of an individual's presence.

Additionally, ECG could be measured without the need of sensors to be applied directly to the skin. So, the sensors can be placed on clothes or have them integrated into a car seat ("smart seating") etc.

Summary

It goes without saying that at the moment, there is not one single biometric identification method that allows the accurate identification of any individual. Each use case requires a thorough research on what method is the most appropriate in any specific conditions and limitations.

To improve the identification performance, so that we get accurate results, it requires a combination of neat blend of different methods (i.e. maybe two or three factor biometric identification methods) that can predict behaviors or possibilities to identify a user at any particular moment in time.

We believe that biometric identification will be a fundamental part of measuring all aspects of security in the future, with new and wider possibilities to recognize and feel the emotional state of individuals.

Hence, we also believe that the technologies will easily adapt to a specific individual's needs, even to the extent that they won't require any mediation. This all forms an important part of the "Ambient Intelligence" vision of the future.

References

1. Finger vein on IEEE
2. Survey on finger vein recognition
3. Personalized Best Bit Map (PBBM) method
4. Bioimpedance analysis in medicine
5. Brief notes on BIA
6. Generating secret keys from bioimpedance
7. A wearable system that knows who wears it
8. Artificial Intelligence for clean biosignals
9. 65 Finalists Announced for the 2017 Interactive Innovation Awards

ABOUT US

SoftServe is a digital authority that advises and provides at the cutting-edge of technology. We reveal, transform, accelerate, and optimize the way enterprises and software companies do business. With expertise across healthcare, retail, media, financial services, software, and more, we implement end-to-end solutions to deliver the innovation, quality, and speed that our clients' users expect.

SoftServe delivers open innovation—from generating compelling new ideas, to developing and implementing transformational products and services.

Our work and client experience are built on a foundation of empathetic, human-focused design that ensures continuity from concept to release.

We empower enterprises and software companies to (re)identify differentiation, accelerate solution development, and vigorously compete in today's digital economy—No matter where you are in your journey.

Visit our **website**, **blog**, **Facebook**, **Twitter**, and **LinkedIn** pages.

USA HQ
201 W 5th Street, Suite 1550
Austin, TX 75703
+1 866 687 3588

EUROPEAN HQ
One Canada Square
Canary Wharf
London E14 5AB
+44 (0) 800 302 9436

info@softserveinc.com
www.softserveinc.com

softserve