Using Low Cost Embedded Systems for Respiratory Sounds Auscultation

João T. Fernandes[1], B. M. Rocha[1], R. P. Paiva[1] and Tiago J. Cruz[1]
[1]Center for Informatics and Systems, Department of Informatics Engineering
University of Coimbra
Coimbra, Portugal
Email: {joaonf, bmrocha, ruipedro, tjcruz}@dei.uc.pt

***Abstract*—With technology advances, trying to replace expensive devices with cheaper, but efficient, systems is a promising approach to pursuit in the future. The main objective of this paper is to compare the Littmann 3200 (state-of-the-art electronic stethoscope) with a prototype built for this purpose (with an electric microphone as the sensor, and an Arduino Mega as a controller), in terms of sound quality, cough detection, and costs. Regarding sound quality, the Littmann is better, especially because of all the technology behind the stethoscope. The costs of both devices are significantly different: the prototype costs 90% less than the Littmann. In terms of cough detection, the Littmann has a sensitivity of 75.4±32.0% and a specificity of 99.6±2.5%, and the prototype has a sensitivity of 71.7±32.1% and a specificity of 98.0±11.0%, but the differences are not significant.**

***Keywords*—Cough detection device, Respiratory sounds acquisition, Embedded Systems**

I. Introduction

Over time, the stethoscope has suffered several updates towards the improvement of sound's amplification, reflected with a higher percentage of correct diagnoses [1]. The stethoscope went through several improvements in the 20th century, making it easier to use without losing sound quality, for example. [2]. During that period, the first electronic stethoscopes that allowed to deepen the amplification of lung/heart sounds were launched. Since their existence, those devices are used to listen to lung and heart sounds. Presently, Littmann is one of the biggest companies that manufacture medical devices (stethoscopes included), and their Littmann 3200 electronic stethoscope is commonly used because of its multiple functionalities. However, the cost of the equipment is a major drawback.

With a new era of low cost and accessible Single-Board-Computers (SBC), these became suitable for embedded systems [3]. These cost-effective devices allow every object to connect to the Internet, the so-called Internet of Things (IoT). Some companies introduced new technology (i.e. new transfer protocols, like MQTT) and equipment (i.e. microcontrollers) that could improve life quality in many areas, including education, science, environment, and health. Now, Arduino LLC and Raspberry Pi Foundation are the main manufacturers of those devices and both have extra components (microphones, temperature sensors, humidity sensors, among others) to which we can connect and expand the system to obtain multi-sensor data.

The anatomy of the human body is fairly complex, containing several different components (organs, tissues, complex systems - respiratory, digestive, urinary, nervous, etc). Two of the main organs are the heart and the lungs. To auscultate the heart, there are sensors that can be used with IoT systems to analyze cardiac activity, e.g., SparkFun's Pulse Sensor, that can detect heart rate variability and the SparkFun's AD8232, a cost-effective board that can be used to monitor the electric activity of the heart. However, to the best of our knowledge there has not been any IoT system solely dedicated to the lungs, and, for that reason, It is crucial to have a system that can auscultate them, even if it is the only function that it does (it does not add extra complexity to the system and it is focused only on the lungs).

Given the cost associated with the acquisition of a stethoscope and knowing that embedded systems made possible to develop devices with some of the characteristics of a computer while simultaneously being more compact [3], it is pertinent to use embedded systems to create a cheaper and more efficient electronic stethoscope.

II. Related Work

Currently, there are several auscultation devices that can help collect and process data (sounds), but they have disadvantages. Some of them have hardware problems, others have incomplete documentation, and, consequently, it is not possible to understand their way of functioning.

With the objective of building a prototype to auscultate lung sounds, and get similar results to the state-of-the-art electronic equipment, both recent technology (embedded systems) and a traditional stethoscope were studied, according to specific parameters (described in each section).

A. Auscultation Systems

Littmann 3200 is one of the most established electronic stethoscopes. Two of its features are the possibility of transferring data, via Bluetooth, from the stethoscope to a computer, and being able to filter 85% of the background noise [4]. One

978-1-5386-5005-9/18/$31.00

of the major disadvantages is the price, which typically sets around 600$ (around 485€)[1]

On the other hand, there are wearable vests with integrated microphones, such as VitalJacket [5], NyxDevice Somnus Sleep Shirt [6] and Smartex Wearable Wellness System [7]. Those systems presented similar characteristics, including integration of data transmission methods (via Bluetooth or WiFi) and acquisition of heart and lungs sounds. Most recently, the Welcome Project aimed to build a vest that could integrate several sensors for monitoring of lung and heart sounds [8]. The idea was to implement an acquisition system able to collect physiological data that could provide information about chronic diseases, specifically, Chronic Obstructive Pulmonary Disease (COPD). Such system was able to monitor the estimated amount of oxygen in the blood (SpO2), electrical impedance tomography (EIT), adventitious respiratory sounds (crackles and wheezing) and high spatial resolution electrocardiograms (ECG). In these four cases, the major disadvantages are the price and their availability in the market. These projects have not had any updates for a long time, and it was not possible to find prices for those cases[2].

III. Materials and Methods

There are different methods used for lung auscultation that makes possible to record and process lung sounds and, in turn, can be used as decision support in diagnosis. This section includes a description of the acquisition protocol, the characterization of the prototype developed in this study, for lung sound auscultation, and the algorithm used for cough detection.

A. Acquisition Protocol

To achieve the best results, measurements needed to take place on the right side of the body, to minimize heart interference. By looking at Figure 1, locations 2, 4 and 6 are located near the heart and, for that reason, these places were discarded.

After testing the other three positions (1, 3 and 5), it was possible to conclude that position 1 presented higher amplitudes, regarding heartbeats. This is due to position 1 being near to position 2 and, consequently, near the heart (when compared to positions 3 and 5). Although positions 3 and 5 presented less heart noise than position 1, to make the sensor robust to different locations, positions 1 and 5 were the selected positions to perform the tests.

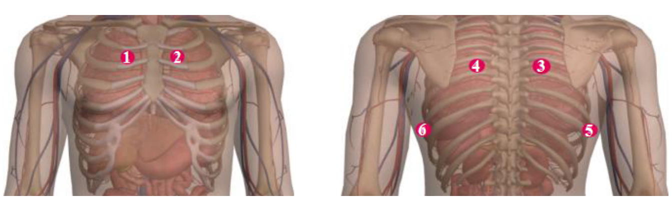

Fig. 1. Positions to acquire lung sounds (left image represents an anterior view and the right image represents a posterior view). The selected position are 1 and 5. Both images were taken from [9].

B. Prototype

To build the prototype, an Arduino Mega was used as a controller. To collect data, a modified SparkFun Electret Microphone Breakout was used. The microphone was unsoldered from the board and placed inside a stethoscope's diaphragm (the stethoscope used was a Logiko Echo DM130). The final montage is shown in Figure 2.

C. Algorithm

The algorithm for cough detection was previously developed by our research group. Rocha et al. presented an algorithm that takes into account spectral content descriptors and pitch-related features [9]. The authors preprocessed the data using a band-pass filter and removed near-silent segments. Then, the features were extracted (pitch and spectral) and classified (using the Logistic Regression algorithm). They achieved a sensitivity of 93.4% and a specificity of 83.4%. Given such performance, this algorithm was used to classify the new data, obtained with the previously described acquisition protocol, for both the prototype and the Littmann 3200.

IV. Experiments

Before performing the tests, each volunteer subject signed a consent form to comply with legal and formal requirements for data collection in humans. This section describes the test's conditions and the population. The systems used for the tests were the prototype, built for this purpose, and the Littmann 3200.

A. Subjects

For the tests, data was collected from 20 subjects, 14 male and 6 female, aging between 19 and 49 years old (the average age was 25 years old). From the 20 subjects, 19 do not smoke and never smoked and more than 50% practices exercise one or more days per week. The Body Mass Index (BMI) ranged from 19.1 to 30.5 (with an average of 23.63).

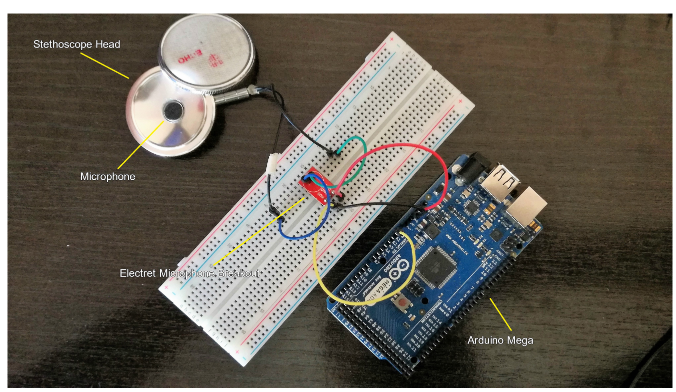

Fig. 2. The final version of the prototype using a microphone embedded in a stethoscope head, connected to a microphone breakout and to an Arduino Mega.

[1]The price was obtained from https://www.stethoscope.com/3m-littmann-model-3200-electronic-stethoscope-p/steth00510.htm - last consulted on June 19th, 2018.

[2]Until January 31st, 2018, the cost of the Smartex Wearable Wellness System was 398€, but the source is not available anymore.

B. Tests Conditions

For each subject, a total of 2 cough tests (for each system tested) were performed: one in the anterior chest and one in the posterior chest (a total of 40 records for each system). Every test lasted for 15s.

The subjects were in a sitting position and, depending on the sex and the system (Littmann or prototype), the tests performed on the anterior chest could be done with clothes on:

- **Male** - if the subject brought a shirt, the tests were made with the shirt opened (in the anterior part) and with the shirt over the device (when the tests were with the prototype) or with the shirt raised (with the Littmann; in the posterior part). If the subject brought anything that could not be opened, they took the piece of clothing off for the anterior part tests.
- **Female** - 5 of the 6 female subjects used a top, which allowed a better and easy way to collect the data; but with all the 6 subjects the data collection underwent in a similar way: the anterior tests were made normally, above the right breast, and the posterior tests were made with the shirt over the device (when the tests were with the prototype) or with the shirt slightly raised (with the Littmann).

On 4 of the 20 subjects, it was not possible to collect the sounds from the anterior part (with the prototype) due to the hairy chest. In total, 34 and 40 records were obtained for the prototype and for the Littmann 3200, respectively (before checking for sound quality).

V. Results

In this section, four components are going to be studied. In the first part, the sound quality will be analyzed according to 2 different parameters (silent and saturation), and also the different sampling rate from both systems. After that, the results regarding cough detection and costs will be presented and analyzed.

A. Sound Quality Assessment

Despite all the precautions with the electronic isolation, wire's length, and electronic signal losses, the sound quality is one of the main concerns regarding the validation of the prototype. For that reason, a simple algorithm was used to determine, based on some parameters (described below), if the recorded signal had acceptable quality for further analysis. To assess the quality of a given signal, first it was normalized between 1 and -1 (to unit variance) and a set of four parameters was evaluated:

- **thrSil** - Threshold below which no sound could be heard, i.e., silence (measured in terms of amplitude, in percentage). The value chosen was 1.5%, because higher values would discard most of the sounds (it works along with silLen).
- **thrSat** - Threshold about which the sound is considered saturated (measured in terms of amplitude, in percentage). For this parameter, a value of 50% was chosen, because the sounds do not have plenty of saturation (only during cough events and for a short amount of time). This value works along with satLen.
- **silLen** - Minimum duration (in seconds) of the silent segment below which that segment is discarded. The chosen value was 0.5s, because in most of the sound signals, when the thrSil is not overcome, it means that it is not possible to hear any sound through the entire signal.
- **satLen** - Minimum duration (in seconds) of the saturated segment above which that segment is discarded. The chosen value was 1.5s because cough events usually last less than this value.

Based on the evaluation of these parameters, a 90% threshold was defined to quantify the quality of a given sound signal (higher than 90% means record accepted).

As it is shown in Figure 3, the sound quality of 2 of the 34 records obtained with the prototype was below the threshold, resulting in a total of 32 records from prototype and 40 records from Littmann stethoscope. In order to compare the performance between the systems, only the records that were acquired in the same conditions were further analyzed (a total of 32 records for each system).

B. Sampling Rate

After performing the tests, it was possible to identify differences regarding the sampling rate. By looking at the Littmann, the tests have 60000 samples for every record, which indicates a steady frequency of 4000Hz and sufficient for cough detection, according to Nyquist theorem (cough frequency components lay between 500Hz and 1200Hz [10]). Regarding the prototype, the sampling rate varies between recordings. For instance, it is 3552Hz in one recording and 3723Hz in another. This happens because the Arduino works with variable bit rates, which means that storing a value of 0 (4 bits) can take 4 units of time less than storing a value of 16 (8 bits). With this operation mode, the prototype sampling rate depends, always, on the values collected, but the minimum sampling rate of the prototype (3552Hz) is still more than two

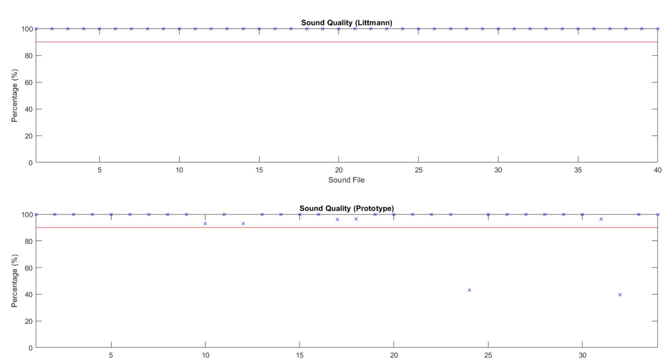

Fig. 3. Sound quality comparison between each signal, from the Littmann and the Prototype. The red line is the 90% threshold.

times higher than the maximum frequency (1200Hz). Although the prototype has a lower and variable sampling rate, when compared to the Littmann, it is still possible to use the sounds for cough detection because the Nyquist theorem is valid.

C. Cough Detection

Cough detection was performed using the algorithm in [9], and tested with the dataset collected, by dividing each record into 15 equal parts (i.e. windows of 1 second each). The objective was to discriminate cough events and non-cough events. An SVM model was trained with 28 features from 465 events, and it was validated using the Leave-One-Out (Patient) Cross-Validation (LOOCV), in [9]. Tables I and II show the results after running the algorithm on the data collected during the tests, where TP means True Positive (the algorithm correctly indicates a cough event), TN means True Negative (the algorithm correctly indicates a non-cough event), FP means False Positive (the algorithm wrongly indicates a cough event), FN means False Negative (the algorithm wrongly indicates a non-cough event), SS means Sensitivity (calculated with (1)) and SP means Specificity (calculated with (2)). For each record, a value for SS and SP was calculated. Subsequently, a non-parametric Mann-Whitney Rank Sum Test was applied, considering a significance level of 5% and the null hypothesis is: the data provided by both sensors are samples from continuous distributions with equal medians.

$$SS = \frac{TP}{TP + FN} \tag{1}$$

$$SP = \frac{TN}{TN + FP} \tag{2}$$

From a statistical point of view, the differences between the two systems were not significant (for SS, $p = 49.36\%$, and for SP, $p = 100\%$). By first looking at SP, the p-value is not surprising because both systems performed well regarding false positives (except in Littmann_20 and Prototype_11), with an average SP of 98.8% and a small standard deviation, in both sensors. Regarding SS, the results were mixed, with some records of the Littmann showing better performance than the prototype (such as Littmann_1 and Prototype_1), but also the opposite (such as Littmann_8 and Prototype_8). In the end, both systems performed similarly (the null hypothesis was not rejected).

Looking at both tables, there are two records that produced a difference in the total SP: Littmann_20 and Prototype_11. Starting with Prototype_11, this was the only record (in a total of 32) that contained 6 cough events. Cough events are, typically, preceded by inspiration. In this case, and after hearing the record, there are 6 cough events and 6 inspirations. These inspirations, unlike in other records, were longer and took almost a second to perform, meaning that the total cough event was longer, which might have led to a wrong evaluation of the algorithm. In other words, as the inspiration segments were more evident both visually and by hearing the records, they were misclassified and considered as cough events.

	TP	TN	FP	FN	SS	SP
Littmann_1	6	9	0	0	1.000	1.000
Littmann_2	4	11	0	0	1.000	1.000
Littmann_3	2	12	0	1	0.667	1.000
Littmann_4	3	12	0	0	1.000	1.000
Littmann_5	1	9	0	5	0.167	1.000
Littmann_6	3	12	0	0	1.000	1.000
Littmann_7	7	7	0	1	0.875	1.000
Littmann_8	4	7	0	4	0.500	1.000
Littmann_9	0	9	0	6	0.000	1.000
Littmann_10	3	10	0	2	0.600	1.000
Littmann_11	5	10	0	0	1.000	1.000
Littmann_12	3	7	0	5	0.375	1.000
Littmann_13	2	13	0	0	1.000	1.000
Littmann_14	0	13	0	2	0.000	1.000
Littmann_15	5	10	0	0	1.000	1.000
Littmann_16	6	8	0	1	0.857	1.000
Littmann_17	4	9	0	2	0.667	1.000
Littmann_18	4	10	0	1	0.800	1.000
Littmann_19	3	11	0	1	0.750	1.000
Littmann_20	5	6	1	3	0.625	0.857
Littmann_21	5	10	0	0	1.000	1.000
Littmann_22	6	9	0	0	1.000	1.000
Littmann_23	7	8	0	0	1.000	1.000
Littmann_24	8	7	0	0	1.000	1.000
Littmann_25	8	7	0	0	1.000	1.000
Littmann_26	6	8	0	1	0.857	1.000
Littmann_27	6	8	0	1	0.857	1.000
Littmann_28	7	8	0	0	1.000	1.000
Littmann_29	6	9	0	0	1.000	1.000
Littmann_30	1	11	0	3	0.250	1.000
Littmann_31	6	9	0	0	1.000	1.000
Littmann_32	2	8	0	5	0.286	1.000
				Average	0.754	0.996
				Standard Deviation	0.320	0.025

TABLE I
LITTMANN RESULTS FOR THE 32 RECORDS.

Regarding Littmann_20, the algorithm considered one window as cough, when that window actually corresponds to an inspiration segment. By hearing this record, the cough events can be distinguished, but the algorithm also classified the inspiration segment between cough as a cough event. Although a previous inspiration typically happens before the cough events, in this case, even though only the inspiration event was captured by the window, the algorithm classified the window as cough, as the inspiration segment was evident enough to be considered as cough. A few milliseconds later, the cough really happens, but not inside the window which the algorithm said. In this case, the difference is minimal and it missed only one time (regarding false positives) in all 32 records.

To summarize, the difference is not significant, but regarding SS, the results could be better. If the window size was smaller (for example, 250ms), the detail would be higher and the performance could increase. Also, with inconstant coughs (although in the protocol it was advised to keep the same number of coughs in every record), the Littmann had a total of 182 cough events (in 32 records; having correctly detected 138 of them, 75.82%), while the Prototype had 207 events (having correctly detected 147, 71.01%).

D. Costs

The total cost of the prototype (only taking into account the material's expenses) was 45€. The most expensive component

	TP	TN	FP	FN	SS	SP
Prototype_1	6	6	0	3	0.667	1.000
Prototype_2	3	10	0	2	0.600	1.000
Prototype_3	2	12	0	1	0.667	1.000
Prototype_4	3	12	0	0	1.000	1.000
Prototype_5	0	9	0	6	0.000	1.000
Prototype_6	0	11	0	4	0.000	1.000
Prototype_7	4	9	0	2	0.667	1.000
Prototype_8	8	7	0	0	1.000	1.000
Prototype_9	5	10	0	0	1.000	1.000
Prototype_10	0	9	0	6	0.000	1.000
Prototype_11	7	3	5	0	1.000	0.375
Prototype_12	2	9	0	4	0.333	1.000
Prototype_13	2	13	0	0	1.000	1.000
Prototype_14	3	12	0	0	1.000	1.000
Prototype_15	6	7	0	2	0.750	1.000
Prototype_16	4	6	0	5	0.444	1.000
Prototype_17	4	10	0	1	0.800	1.000
Prototype_18	8	6	0	1	0.889	1.000
Prototype_19	12	3	0	0	1.000	1.000
Prototype_20	4	7	0	4	0.500	1.000
Prototype_21	5	9	0	1	0.833	1.000
Prototype_22	2	6	0	7	0.222	1.000
Prototype_23	5	10	0	0	1.000	1.000
Prototype_24	8	7	0	0	1.000	1.000
Prototype_25	6	8	0	1	0.857	1.000
Prototype_26	9	6	0	0	1.000	1.000
Prototype_27	5	5	0	5	0.500	1.000
Prototype_28	7	6	0	2	0.778	1.000
Prototype_29	3	10	0	2	0.600	1.000
Prototype_30	3	12	0	0	1.000	1.000
Prototype_31	5	9	0	1	0.833	1.000
Prototype_32	6	9	0	0	1.000	1.000
				Average	0.717	0.980
				Standard Deviation	0.321	0.110

TABLE II
PROTOTYPE RESULTS FOR THE 32 RECORDS.

was the board (the Arduino Mega), costing around 35€. The difference between both sensor's cost is 90% (around 440€). It is a considerable difference, especially when observing that the performance of the sensors is not significantly different.

VI. CONCLUSION AND FUTURE WORK

By looking at the cough detection performance results obtained for both sensors (SS and SP), considering that no statistically significant differences were obtained and taking into account the costs associated with each item, it is possible to conclude that the developed prototype has potential to substitute a commercial stethoscope.

Despite the good results, some issues do not allow us to consider the prototype better than the Littmann 3200. For instance, the signal quality is worst and the sampling rate is unstable. The number of samples per second, in the Littmann 3200, is 4000Hz (60000 samples in 15 seconds), very stable and with no deviation in any recording. In the prototype, the samples varied between 3552Hz and 3723Hz for each second (average between 53280 and 55845 samples per record), not being uniformly sampled among all records.

The records obtained with the prototype showed considerable saturation every time that a cough occurred. Unlike the Littmann, the prototype does not contain any filter in the microphone and, for that reason, the records have a different morphology and more saturation. Regarding costs, the prototype cost was around 45€, but it was mostly soldered. During the tests, sometimes it failed to record and some tests needed to be repeated (because of random noise).

With satisfactory results, the next steps encompass improving the build by, for example, including filters to remove heart sounds interference. When considering the issues with sampling rate (leading to the collection of unevenly-sampled time-series and subsequently to records differing in the total number of samples) and the nonexistence of integrated filters in the prototype developed, it can be expected that analysis of data acquired with Littmann returns superior performance (it contains Bluetooth, computer application, and LCD screen).

This research was a preliminary work to understand if modern technology could replace older and most expensive one. Although it has promising results regarding sound analysis and cost, these devices are in an early stage of their development. Adding more features and work towards a new auscultation device for sound collecting (in this case, physiological sounds) are the next steps.

ACKNOWLEDGMENT

This work was carried out in the scope of the MobiWise project: From mobile sensing to mobility advising (P2020 SAICTPAC/0011/2015), cofinanced by COMPETE 2020, Portugal 2020 - POCI, European Union's ERDF, and the Portuguese Foundation for Science and Technology (FCT), and supported by the EU project WELCOME[3] (FP7-ICT-2013-10/611223) .

REFERENCES

[1] "American Diagnostic Corporation - Core Medical Device Manufacturer. Stethoscopes, Blood Pressure, Thermometry, and EENT," last visited on 2017-03-13. [Online]. Available: http://adctoday.com/learning-center/about-stethoscopes/history-stethoscope http://adctoday.com/learning-center/about-sphygmomanometers/history-sphygmomanometer

[2] "Stethoscope History," last visited on 2018-03-09. [Online]. Available: https://www.littmann.com/3M/en_US/littmann-stethoscopes/education-center/history/

[3] "What is Embedded System," last visited on 2017-03-13. [Online]. Available: http://internetofthingsagenda.techtarget.com/definition/embedded-system

[4] Littmann Stethoscopes, "3M Littmann Electronic Stethoscope Model 3200," 2009. [Online]. Available: http://www.littmann.com/3M/en_US/littmann-stethoscopes/products/ /3M-Littmann-Electronic-Stethoscope-Model-3200?N=5932256+8711017+3293188392&rt=rud

[5] VitalJacket, "VitalJacket Product," last visited on 2017-03-13. [Online]. Available: http://www.vitaljacket.com/?page_id=156

[6] "Nyx Devices," last visited on 2017-03-13. [Online]. Available: http://nyxdevices.com/product/

[7] Smartex, "Wearable Wellness System," 2012, last visited on 2017-03-13. [Online]. Available: http://www.smartex.it/en/our-products/232-wearable-wellness-system-wws

[8] "Welcome Project." [Online]. Available: http://www.welcome-project.eu/home.aspx

[9] B. M. Rocha, L. Mendes, J. Henriques, P. Carvalho, and R. P. Paiva, "Detection of Explosive Cough Events in Audio Recordings by Internal Sound Analysis," *Accepted to EMBS (to be presented)*, p. 4, 2017.

[10] J. Korpa and J. Sadlon, "Methods of Assessing Cough and Antitussives in Man Analysis of the Cough Sound : an Overview," pp. 261–268, 1996.

[3] http://www.welcome-project.eu/home.aspx