

# Automatic Analysis of Lung Function Based on Smartphone Recordings

João F. Teixeira<sup>1</sup>✉, Luís F. Teixeira<sup>2</sup>, João Fonseca<sup>3</sup>, and Tiago Jacinto<sup>3</sup>

<sup>1</sup> Department of Electrical and Computer Engineering, University of Porto,  
Porto, Portugal

`jpfteixeira.eng@gmail.com`

<sup>2</sup> Department of Informatics Engineering, University of Porto, Porto, Portugal  
`luisft@fe.up.pt`

<sup>3</sup> Department of Health Information and Decision Sciences, University of Porto,  
Porto, Portugal

`jfonseca@med.up.pt, tajacinto@gmail.com`

**Abstract.** Over 250 million people, worldwide, are affected by chronic lung conditions such as Asthma and COPD. These can cause breathlessness, a harsh decrease in quality of life and, if left undetected or not properly managed, even death. In this paper, we approached part of the lines of development suggested upon earlier work. This concerned the development of a system design for a smartphone lung function classification app, which would only use recordings from the built-in microphone. A more systematic method to evaluate the relevant combinations of methods was devised and an additional set of 44 recordings was used for testing purposes. The previous 101 were kept for training the models. The results enabled to further reduce the signal processing pipeline leading to the use of 6 envelopes, per recording, half of the previous amount. An analysis of the classification performances is provided for both previous tasks: differentiation into Normal from Abnormal lung function, and between multiple lung function patterns. The results from this project encourage further development of the system.

**Keywords:** Asthma · Breath · COPD · Machine learning · Signal processing · Smartphone · Spirometry

## 1 Introduction

Chronic respiratory diseases such as Asthma and Chronic Obstructive Pulmonary Disease (COPD) are incurable, yet treatable and their early detection is crucial to provide a better quality of life. Major risk factors include air pollution, tobacco smoking and occupational environments containing dust and chemicals. The World Health Organization (WHO) estimates that over 250 million people suffer from asthma and COPD [15] and more than 3 million people died of COPD in 2005 [16].

Spirometry is the measurement of breath, i.e., is the most popular noninvasive set of timed tests that enables to measure the mechanical properties of the lungs,

also named pulmonary function [8]. The keystone test is the Forced Expiratory Maneuver (FEM) where the patient fully inspires and then forcefully exhales all the air available, as fast as possible.

The increasing use of smartphones has enabled the emergence of several health related systems. Their computational power is ever increasing and, equipped with multiple sensors, it is possible to develop disease prevention, diagnosis and monitoring applications.

This paper provides an extension of previous work [12], where the performance of several groups of methods, along with clinical parameters, were compared. The aim is to find the most relevant, most efficient and fastest combination to produce a smartphone app for measuring and classifying lung function. The system's input is restricted to the smartphone's built-in microphone, in order to avoid external components.

## 2 Background and Related Work

From the dawn of non-invasive lung function evaluation, which produced the standard spirometer, several improvements have been made throughout the years. The next logical step concerning portability and affordability involves lung function estimation with smartphones' microphones. Some studies have already been conducted in order to accurately measure the clinical parameters [5, 17], and also considering robustness to ambient noise [11].

In previous work [12], a lung function estimation system was proposed containing signal processing and machine learning algorithms. It was based on 101 recordings, collected from 61 patients performing the forced expiratory maneuver (FEM). Due to the reduced and unbalanced number of instances (label-wise), the evaluation of all models was done in a 5-fold cross-validation scheme.

## 3 Dataset

The dataset for this work comprises the previous 101 recordings and a new set of 44 recordings, collected in the same conditions of the previous ones. Some recordings were gathered on a controlled environment with low background noise, however more than 80 % of the recordings experienced background noise such as physicians giving verbal incentive, talking voices and small machine noises at a short distance. The recordings were made using a Samsung GT-I9000. The first 101 instances constituted the training set while the new set was used for testing.

Due to data collection constraints, the new recordings only included the 4 most popular parameters. However, this is not very problematic since, as previously shown, the remaining clinical indicators are very unreliable and extremely patient cooperation dependent. In fact, this enables the optimization process for the signal processing pipeline to focus on the most relevant values.

Each recording is accompanied by the patient's anthropometric parameters (age, height, weight and gender), clinical parameters, and classification of the patients lung function provided by the recording physician. The clinical parameters were obtained by performing the FEM to a spirometers available. The classification types are normal and abnormal (obstruction, restriction or mixed).

The included patients were part of the clinical study Control and Burden of Asthma and Rhinitis (ICAR), patients attending the allergology clinic from CUF Porto Institute (ICP) or from CUF Porto Hospital (HCP). Data collection occurred between April 3rd 2014 and February 1st 2015.

## 4 Algorithms and System Architecture

### 4.1 Signal Processing

The system's input consists of microphone recordings which are AC coupled, uncalibrated signals that represent air pressure. The signal processing pipeline can be divided in four portions: automatic signal segmentation, signal pre-processing, envelope generation and envelope processing. Figure 1 shows the initial architecture for the signal processing part without the signal segmentation.

**Automatic Signal Segmentation.** The audio input was initially segmented in order to remove non expiration sounds, such as the inspiration portion of the maneuver and ambient noise. The definition and cropping the beginning of the sound was accomplished using a modified version of the Back-Extrapolation algorithm [7].

First, an LPC envelope of the signal is obtained, as it will be described further on the paper. Then the minimum value between inspiration and expiration peaks is found and the initial part is removed. Afterwards, the zero-time back-extrapolation is performed by finding the instant corresponding to the envelope's peak (PEF time), calculating the Time-Volume curve, drawing the tangent at the PEF time and finding where the tangent crosses the abscissas, which is the initial instant.

The ending at noise level was detected using a sliding window algorithm (5 % of signal's length, 25 % overlap) based on the magnitude ratio threshold of the maximum value (2 %).

**Signal Pre-Processing.** The recordings are limited in excursion and patients need to perform the expiratory maneuver at an arms length to avoid microphone saturation. Therefore, it seems relevant to compensate the pressure lost between the lips ( $p_{lips}$ ) and smartphone ( $p$ ), using an Inverse Radiation Model. Furthermore, this model also atones the reverberation effect from sound reflections around a person's body. Afterwards,  $p_{lips}$  was converted to airflow at the lips ( $u_{lips}$ ), using a Pressure to Flow Conversion Model. Both models were developed in similar fashion to [5].

**Envelope Generation.** The third stage employed several methods to calculate the signal envelopes, approaching different sound characteristics to obtain a comprehensively robust feature extraction. The algorithms' input consisted of both the segmented audio and the two resulting signals from the pre-processing stage, as all of them can be considering roughly proportional to air flow.

*Generic Envelope Extraction.* To obtain an envelope based on a time domain approach two methods were used: the Hilbert Transform and Shannon curves [6]. The first approach consists of calculating the signal's harmonic conjugate, with the Hilbert Transform, and to add it back to the signal, resulting on an envelope. The second approach involves calculating the Shannon Entropy and Energy envelopes of the signal. They act as non-linear transformations focusing either on the higher (Energy) and lower (Entropy) intensities of the signal. Both approaches output highly noisy curves that need subsequent smoothing.

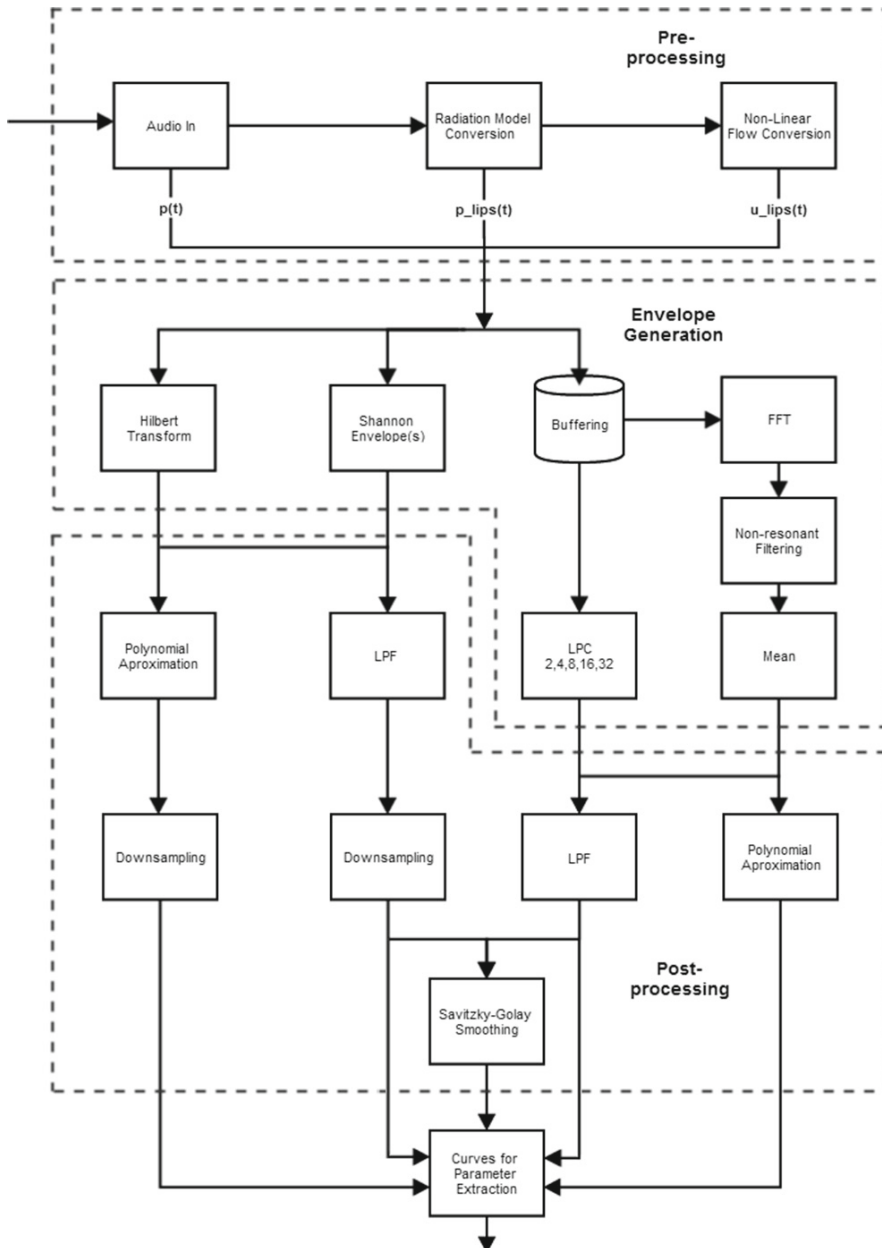
*Linear Predictive Coding.* The audio input is segmented in windows of 31.25ms, with 50 % overlap. The white noise variance, or power, is obtained from the LPC model outputs. While the LPC filters can approximate the vocal tract [14], the succession of power values should be proportional to the exhalation power at the respective time and constitute a sampled envelope of the signal. The implementation included models of degrees 2, 4, 8, 16 and 32, which represents increasing vocal complexity.

*Mean of Resonances.* Similarly to LPC, the signal was buffered into 31.25 ms frames, with 50 % overlap. Each frame underwent a 256-point FFT operation using a hamming window, producing a spectrogram. All spectrogram values lower than 20 % the respective frames' maximum were considered noise and were consequently discarded. Resonances over 250ms, within the respective frequencies' 2 bin neighborhood were kept, preserving only relatively large and long frequencies, and taking into account the natural occurring frequency shift. The envelope was obtained by averaging the frames' saved resonances.

**Envelope Post-processing.** The several envelopes obtained were processed using different settings in order to find the best combination for the application. The envelopes were smoothed by a regular low pass filter (LPF) and, in parallel, were also approximated by a 4th order polynomial. To obtain the same sampling rate as the buffered methods, the Hilbert Transform and Shannon envelopes' results were downsampled accordingly. The non-approximated envelopes were also further processed using a Savitzky-Golay filter (SG) with order 3 and size 11 [10], as depicted on Fig. 1.

## 4.2 Parameter Extraction

For each recording, the spirometry parameters were calculated from each of the final envelopes. The measurements extracted were PEF, FVC, FEV<sub>1</sub> and FEV<sub>1</sub>/FVC. The envelopes are viewed as Flow-Time curves, typical of spirometer reports.



**Fig. 1.** Initial signal processing pipeline (P0).

PEF is defined as the Peak Expiratory Flow or the global maximum of the audio envelope. By integrating the envelope with respect to time the Volume-Time curve can be obtained. FVC is defined as the total volume expired of a FEM. FEV<sub>1</sub> is the total volume expired during the first second.

### 4.3 Machine Learning

The system’s machine learning pipeline can be divided into two stages: the parameter regression and the classification. The first uses the parameters extracted from the curves to obtain an estimation of the respective clinical values, as given by spirometers. The second, devises models that can discern between the possible illness states, initially addressing the distinction of normal from abnormal lung function and then, normal from 2 types of pathologies.

**Regression Stage.** Every recording produces several envelopes and each one is used to extract clinical measurements. This information is used to produce a relatively robust estimation of the respective spirometer measurement. For instance, each set of PEF measurements computed from any recording is used as a batch input for the trained regression model to obtain an estimated PEF value. The process is then repeated for the other types of clinical measurements. The corresponding spirometer measurements acted as ground truth or regression targets. For this task, Regression Tree Bagging [1] and Random Forests (RF) [2] were used. The number of trees in the ensembles was reduced to 10 without significant loss in accuracy. Also, RF employed a selection size for the random feature subset of  $n/3$  out of the total  $n$  feature set.

**Classification Stage.** On this stage, the regressed parameters were the input of the learning models. Several different classification models were tested, namely: Decision Trees [3], either as one tree, Tree Bagging, Random Forest ( $\sqrt{n}$  subset) and AdaBoost [4], Support Vector Machines (SVM) [13] and Naïve Bayes [9]. Although the tree ensemble methods used 70 trees, only 10 trees were grown for AdaBoost to avoid overfitting.


## 5 Experimental Approach

### 5.1 Regression Experiments


Unlike in the previous work, we aimed to devise a more systematic approach concerning the removal of the signal processing blocks. Hence, the experiments are based on random forward selection and model improvement tracking. Basically, a random path (sequence of methods) of the system’s pipeline is used to predict the parameters’ values and, iteratively, more are added. After each prediction the mean error and standard deviation is saved as well as the improvement of the respective path’s addition (Fig. 2). The termination condition is adding 12 paths (maximum number of paths from previous work and to reduce cumulative path bias).

The  $Avg(Avg)$  and  $Avg(Std)$  refer to the averages of the mean value and standard deviation of the error, respectively. The *Score* (blue line) is calculated as an average of the parameter-wise weighted average between the mean ( $w = 3$ ) and standard deviation ( $w = 2$ ) of the error. The vertical lines indicate the minimum.


This process is repeated 100 times, ultimately providing a histogram of the relative influence of each path (Fig. 3). To evaluate which system blocks were more useful to include, the importance of a method was considered as the average



**Fig. 2.** Score progression with inclusion of system paths. Tree Bagging, initial dataset.



**Fig. 3.** Relative system path importance. Tree Bagging, initial dataset, 100 iterations. The 6 maximum values are represented as red dots (Color figure online).



**Fig. 4.** Relative system block importance. Tree Bagging, initial dataset, 100 iterations.

of the importance of each path that uses that method (Fig. 4). Each model generated used a 5-fold cross validation dataset and was later tested with the new dataset.

5.2 Classification Experiments

The classification process was first devised as a Normal against Abnormal classification problem, referred to as two label experiments (TLE). Then, multiple label experiments (MLE) were conducted, where the models tried to distinguish between Normal lung function and Obstruction and Mixed pathologies. For both problems, the experiments varied on the feature space used. The experiments used:

- 1. Set A<sup>1</sup>,
- 2. Set A and height,
- 3. Set A, height and age.

6 Results and Discussion

6.1 Regression Experiments

We proceeded to execute the analysis described in Sect. 5.1, several times. The ranking of relevance of the SP pipeline blocks, such as in Fig. 4, presented varying results for each 100 iterations run. The exception to these changes were the filtering options *LP-sGolay* and *Polynomial Fitting* that maintained a very high importance value, contrasting with using just the *LP* filter which tended

**Table 1.** Comparison of regression error average and standard deviation for the clinical parameters. P1 refers to results with the previous final SP pipeline and P2 to the new pipeline.

Task	SP	Model	Average Error (%)				Std. Dev. Error (%)			
			PEF	FVC	FEV <sub>1</sub>	Tiff	PEF	FVC	FEV <sub>1</sub>	Tiff
Train	P1	Bag	23.46	30.03	22.77	10.32	22.93	35.77	23.36	8.66
		RF	22.26	28.25	24.05	9.59	22.98	29.14	26.40	8.20
	P2	Bag	22.66	26.03	23.73	10.22	23.96	25.74	24.19	8.27
		RF	22.43	28.62	24.29	10.21	24.55	27.98	24.98	8.10
Test	P1	Bag	27.83	32.18	26.67	12.99	27.11	22.82	17.60	13.10
		RF	27.21	32.53	25.53	13.74	27.34	23.19	17.56	13.56
	P2	Bag	26.76	32.26	26.95	13.23	25.83	22.66	17.19	12.86
		RF	27.37	32.57	25.88	13.61	26.18	23.30	17.51	13.29

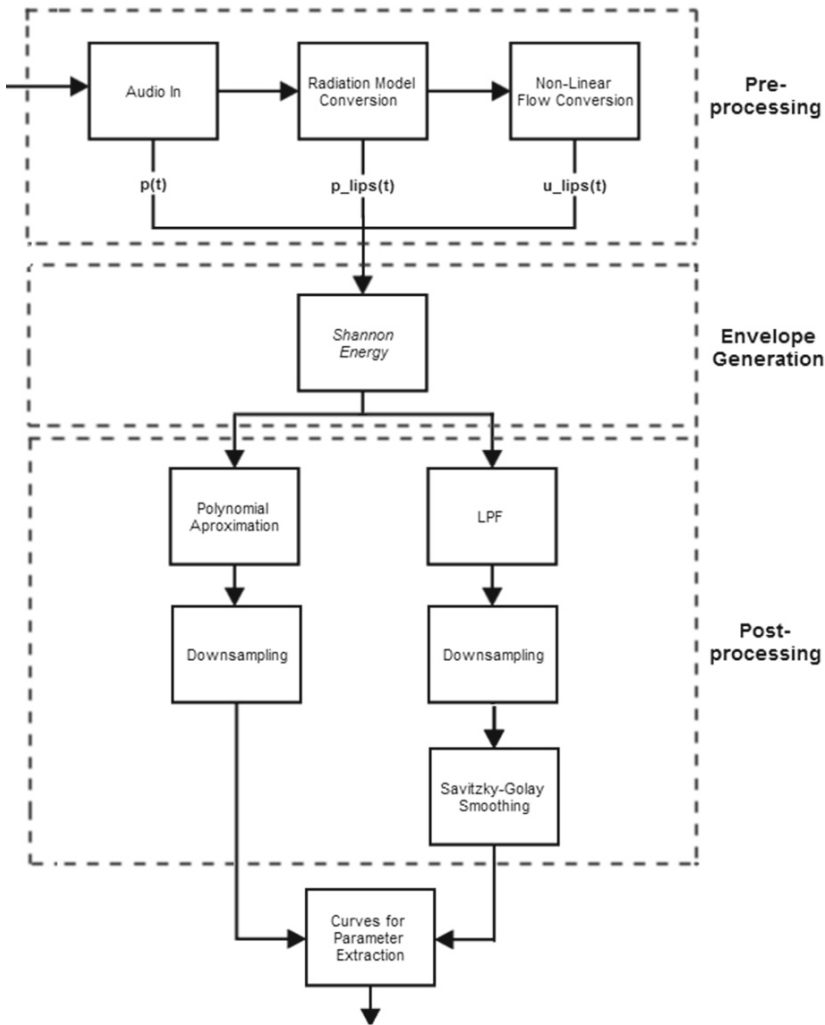
<sup>1</sup> Set A: PEF, FEV<sub>1</sub>, FVC, FEV<sub>1</sub>/FVC.



to be of very low importance. The remaining methods, pre-processing and envelope generating functions, changed often relative importance value and order, which suggests that neither is particularly relevant for the generation of good attributes/paths.

Based on these results, and in order to further reduce the complexity of the system to be implemented, we opted for the following system:

- Maintain all the filtering models (discarding the direct use of  $LP$  output)
- Maintain all the pre-processing outputs ( $p$ ,  $p_{lips}$ ,  $u_{lips}$ )
- Just use *Shannon Energy* (least complex envelope generator:  $x^2$ ,  $\sqrt{x}$  and  $\log_{10}(x)$ )



**Fig. 5.** P2 signal processing architecture.

This new pipeline (P2, Fig. 5), produces a total of 6 attributes per clinical parameter, half than the previous system (P1). With the pipeline selected, the models were retrained and tested. The results are shown in Table 1.

From these results a few remarks can be made. First, all values are very similar, either across models, signal processing pipeline and task. This is particularly interesting since, by halving the number of values to combine in regression, and consequently simplifying the pipeline, approximately the same results were obtained. Furthermore, the training version of Bagging with P2 pipeline for FVC managed to present significantly better results than the correspondent version for P1.

Another interesting fact is that the testing results quite resemble the training ones, with a maximum of 5 % increase in the average error rates and even significant improvement in standard deviation for FVC and FEV<sub>1</sub>.

Like in previous work, Bagging and Random Forest present very similar results and an option cannot be chosen for deployment with sufficient confidence. As before, the Random Forest was chosen to be used for the subsequent experiments, so to reduce overfitting and shorten regression time, along with the P2 pipeline.

## 6.2 Classification Experiments

The classification experiments using P2 presented interesting results, shown on Table 2. As before, through the inclusion of height and age attributes the models accuracies managed to improve. However, unlike in the experiments of previous work, including just height introduced significant bias towards the *Normal* class, despite improving the accuracy (Ac.). This apparent contradiction can be explained by the instance label unbalance. In fact, both testing results have shown a clear bias for Example 1 and 2, leaving the precision (Prec.) and recall (Rec.) for the *Abnormal* class mostly at zeros (not on table). Either way, similarly to previous work, using Set A in conjunction with height and age provides the best results and thus, mainly the Example 3 results are considered for model evaluation.

At this stage, the model that provides best results is Tree Bagging, which has the best accuracy (88.1 % / 65.9 %) and has one of the best precision-recall tradeoffs, in the sense that a higher precision is desirable, implying lower bias towards the *Normal* class (prec. 90.4 % / 78.2 %).

Overall, the training results resembled those from previous work, reaching a difference in error rate of 4 % for the best method, which is hardly significant for the dataset's size. Surprisingly, however, the models did not manage to produce testing results that matched training predictions. Models with P1 pipeline were also tested but did not show significant change (accuracy improvement below 2 %, not shown). This suggests that the testing set is somewhat different from the training set.

Table 3 presents the results of the multiple label classification problem, using P2. Like before, the training results seemed to improve with the succession of experiments reinforcing the importance of using the anthropometric values with

**Table 2.** Two label classification problem (TLP) results (%) for the best performing models and single tree classifier.

Method	Training						Testing					
	Exp. 1			Exp. 3			Exp. 1			Exp. 3		
	Ac.	Prec.	Rec.	Ac.	Prec.	Rec.	Ac.	Prec.	Rec.	Ac.	Prec.	Rec.
Single Tree	53.4	68.7	61.9	83.1	87.5	88.7	56.8	60.9	89.2	65.9	80.9	60.7
Tree Bagging	62.3	69.4	83.1	88.1	90.4	92.9	59.0	61.9	92.8	65.9	78.2	64.2
Random Forest	68.3	70.9	92.9	86.1	86.0	95.7	61.3	62.7	96.4	63.6	66.6	85.7
Adaboost	68.3	71.4	91.5	80.2	82.2	91.5	63.6	63.6	100	65.9	66.6	92.8
SVM	70.3	70.3	100	87.1	90.2	91.5	63.6	63.6	100	63.6	63.6	100

the clinical ones. On the other hand, the accuracy of the test results were overall slightly reduced. However this is caused by a reduction of bias towards the *Normal* class. In fact, on the Tree Bagging experiment evolution, the accuracy was maintained but the precision for the *Normal* class rose from 66.6 % to 84.6 % while precision and recall values for the other classes were changed from null values.

**Table 3.** Multiple label classification problem (MLP) results (%) for the best performing models and single tree classifier.

Method	Training		Testing	
	Exp. 1	Exp. 3	Exp. 1	Exp. 3
Single Tree	58.5	77.7	61.9	54.7
Tree Bagging	65.6	86.8	66.6	66.6
Random Forest	67.6	81.8	66.6	61.9
Adaboost	65.6	76.7	66.6	64.2
SVM	71.7	86.8	66.6	9.52

The testing results has also shown relatively low accuracy comparing to the training results, despite being more evident on Example 3 values.

Similarly to the TLP results, these also suggest that Tree Bagging is the best option among the tested. It provides the best accuracies in both training and testing and even slightly improves on the P1 results (accuracy 82.2 %) [12], even though not significantly.

## 7 Limitations and Future Work

### 7.1 Limitations

This project presented some issues concerning data collection that, once overcome, should enhance the learning models' performance and, consequently, the results.

A great portion of the recordings was gathered on a relatively fast paced clinical study where patients had to perform several respiration maneuvers before recording to the smartphone. This could have reduced the patient's cooperation level due to fatigue. On the other hand, the forced expiration maneuver itself is difficult to perform, specially when concerning these recordings where no mouth-piece was used. These factors also contributed to the reduced yield of properly executed recordings. The small dataset with little intra-patient samples is most likely the cause of the regression errors of over 20 %.

Additionally, since the spirometer and smartphone maneuvers were made separately there is no completely reliable ground truth.

## 7.2 Future Work

Further study of this technology is needed and some key features are proposed. In addition to collect further recordings, it is relevant to devise an algorithm to automatically detect poorly executed FEMs in order to immediately request a repetition during data collection. Also, an application based on the proposed architecture should be implemented.

## 8 Conclusion

In this paper, we approached part of the lines of development suggested upon earlier work. A more systematic method to evaluate the relevant combinations of methods was devised and an additional set of recordings was used for testing purposes.

The new evaluation method lead to a new pipeline (P2) which is simpler to implement and should be faster to execute, due to the decrease in envelope generating functions. This, in turn, further reduced the number of the envelopes processed to 6, half the amount of the previous system. The regression results from the P2 system are very similar to the previous ones (P1). Furthermore, models from either pipeline show test values with a maximum 5 % increase in error rate and even better standard deviation for some clinical parameters.

The classification experiments confirmed the usefulness of combining the clinical parameters with the patient's anthropometric data. The TLP and MLP results were comparable to previous work, even slightly improving on the MLP. The testing results were not expected to be so low and, considering the test results using the P1 pipeline, it suggests an issue with the testing data. Despite this, the remaining results encourage further development of the system.

**Acknowledgements.** This work was conducted with the support of the *Control and Burden of Asthma and Rhinitis* project (ICAR), with the grant PTDC/SAUSAP/119192/ 2010. The authors would like to thank Bernardo Pinho for the development of an enhanced recording app and Ivânia Gonçalves, Rita Silva and Daniela Santos for the added effort of recording the patients in addition to their tasks on the patient screenings.

## References

1. Breiman, L.: Bagging predictors. *Mach. Learn.* **24**(2), 123–140 (1996). <http://dx.doi.org/10.1007/BF00058655>
2. Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001). <http://link.springer.com/article/10.1023/A:1010933404324>
3. Buntine, W.: Learning classification trees. *Stat. Comput.* **2**(2), 63–73 (1992). <http://dx.doi.org/10.1007/BF01889584>
4. Freund, Y., Schapire, R.E.: Experiments with a new boosting algorithm. In: Saitta, L. (ed.) *ICML*, pp. 148–156. Morgan Kaufmann, Bari, Italy (1996). <http://web.eecs.utk.edu/~parker/Courses/CS425-528-fall12/Handouts/AdaBoost.M1.pdf>
5. Larson, E.C., Goel, M., Boriello, G., Heltshe, S., Rosenfeld, M., Patel, S.N.: SpiroSmart: using a microphone to measure lung function on a mobile phone. In: 14th ACM International Conference on Ubiquitous Computing, p. 10. Pittsburgh, Pennsylvania, USA (2012)
6. Liang, H., Lukkarinen, S., Hartimo, I.: Heart sound segmentation algorithm based on heart sound envelopogram. In: *Computers in Cardiology 1997*. vol. 24, pp. 105–108. IEEE (1997). <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=647841>
7. Miller, M.R., Hankinson, J., Brusasco, V., Burgos, F., Casaburi, R., Coates, A., Crapo, R., Enright, P., van der Grinten, C.P.M., Gustafsson, P., Jensen, R., Johnson, D.C., MacIntyre, N., McKay, R., Navajas, D., Pedersen, O.F., Pellegrino, R., Viegi, G., Wanger, J.: Standardisation of spirometry. *Eur. Respir. J.* **26**(2), 319–338 (2005). <http://www.ncbi.nlm.nih.gov/pubmed/16055882>
8. Pierce, R.: Spirometry: an essential clinical measurement. *Aust. Fam. Physician* **34**(7), 535–539 (2005). <http://www.ncbi.nlm.nih.gov/pubmed/15999163>
9. Russel, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*, 2nd edn. Prentice Hall, Englewood Cliffs (2002)
10. Savitzky, A., Golay, M.J.E.: Smoothing and differentiation of data by simplified least squares procedures. *Anal. Chem.* **36**(8), 1627–1639 (1964). <http://pubs.acs.org/doi/abs/10.1021/ac60214a047>
11. van Stein, B.: A Mobile Smart Care platform Home spirometry by using the smart-phone microphone. Master's thesis, Leiden University, Leiden, The Netherlands (2013)
12. Teixeira, J.F., Teixeira, L., Fonseca, J., Jacinto, T.: Lung function classification of smartphone recordings - comparison of signal processing and machine learning combination sets. In: *Proceedings of the International Conference on Health Informatics*, pp. 123–130. Lisbon, Portugal (2015). <http://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0005222001230130>
13. Vapnik, V.N.: An overview of statistical learning theory. *IEEE Trans. Neural Netw. Publ. IEEE Neural Netw. Council.* **10**(5), 988–999 (1999). <http://www.ncbi.nlm.nih.gov/pubmed/18252602>
14. Wakita, H.: Direct estimation of the vocal tract shape by inverse filtering of acoustic speech waveforms. *Audio Electroacoust. IEEE Trans.* **21**(5), 417–427 (1973)
15. Asthma: Fact sheet N307, November 2013. <http://www.who.int/mediacentre/factsheets/fs307/en/>. Accessed: 26–06–2014
16. Chronic obstructive pulmonary disease (COPD): Fact sheet N315 (October 2013). <http://www.who.int/mediacentre/factsheets/fs315/en/>. Accessed: 26–06–2014
17. Xu, W., Huang, M.C., Liu, J.J., Ren, F., Shen, X., Liu, X., Sarrafzadeh, M.: mCOPD. In: *Proceedings of the 6th International Conference on Pervasive Technologies Related to Assistive Environments - PETRA 2013*, pp. 1–8. PETRA 2013, ACM Press, New York, New York, USA (2013). <http://doi.acm.org/10.1145/2504335.2504383>, <http://dl.acm.org/citation.cfm?doid=2504335.2504383>