

# Segmentation-free Heart Pathology Detection Using Deep Learning

Erika Bondareva<sup>1</sup>, Jing Han<sup>1</sup>, William Bradlow<sup>2</sup>, Cecilia Mascolo<sup>1</sup>

**Abstract**—Cardiovascular (CV) diseases are the leading cause of death in the world, and auscultation is typically an essential part of a cardiovascular examination. The ability to diagnose a patient based on their heart sounds is a rather difficult skill to master. Thus, many approaches for automated heart auscultation have been explored. However, most of the previously proposed methods involve a segmentation step, the performance of which drops significantly for high pulse rates or noisy signals. In this work, we propose a novel segmentation-free heart sound classification method. Specifically, we apply discrete wavelet transform to denoise the signal, followed by feature extraction and feature reduction. Then, Support Vector Machines and Deep Neural Networks are utilised for classification. On the PASCAL heart sound dataset our approach showed superior performance compared to others, achieving 81% and 96% precision on normal and murmur classes, respectively. In addition, for the first time, the data were further explored under a user-independent setting, where the proposed method achieved 92% and 86% precision on normal and murmur, demonstrating the potential of enabling automatic murmur detection for practical use.

## I. INTRODUCTION

Cardiovascular (CV) diseases are the number one cause of mortality internationally, accounting for 17 million deaths per year according to the World Health Organisation [1]. Auscultation, the process of listening to the heart using a stethoscope, is a powerful screening tool due to its affordability, non-invasiveness, safety, and ease of administration; but it is a very difficult skill to master. Therefore, researchers are starting to explore the potential of digital technologies for both sound gathering and automatic data analysis.

To fully appreciate the capabilities of CV auscultation, it is important to understand the origin of the sounds. One full cardiac cycle consists of two sounds, commonly referred to as S1 and S2, which are caused by valve closure. Occasionally, there are additional sounds present, which are most often indicative of a pathology.

Conventional heart sound classification methods typically involve sound filtering, segmentation, feature extraction, and machine learning-based inference. However, the performance on a wider sample population is rarely a concern. It has been reported that many denoising approaches in the literature do not work on real-life data due to the algorithms developed

on artificially noisy data [2]. Additionally, there are questions and issues remaining which are discussed below and investigated particularly in this study.

*Limitations of segmentation:* One of the most common segmentation categories are envelope-based methods, including normalised Shannon average energy [3], [4], [5] and Hilbert transform and its variations [6], [7], [8]. These methods, however, are based on the premise that systole is shorter than diastole, which is only true for a narrow band of heart rates [9], and they become less effective for noisy signals [10], [11]. As a consequence, segmentation's reliance on the signal being clean within a narrow band of heart rates drastically limits its application, with the error introduced by incorrect segmentation propagating into heart sound pathology classification, and introducing considerable computational complexity into the algorithms [12]. To avoid such complexity and the underlying error propagation issue, in this work we propose a segmentation-free framework for heart sound classification. To the best of our knowledge, there are very few segmentation-free methods [13], and our method shows better performance.

*Lack of performance comparison:* To tackle the heart sound classification problem, a variety of models have been investigated. For instance, Support Vector Machines (SVMs) are arguably the most popular classifier owing to their mature theoretical foundation, with a plethora of research attempts to explore various feature inputs for heart sound classification. Some utilised SVMs with deep features as inputs, which yielded a sensitivity of 84.8% [14], and others applied sparse coding for feature extraction, achieving 86.5% accuracy [15], both on PhysioNet heart sound dataset. More recently, Deep Neural Networks (DNNs) have started to receive greater attention for this task [16], [17], [18]. While DNNs deal exceptionally well with multi-class problems, unlike SVMs, they require a large amount of training data. However, owing to the lack of fair performance comparison, no clear observations can be drawn as to the superiority of any of them. In our study, we compare the inferential capabilities of SVMs and DNNs on the same data set, and the results indicate that DNNs appear superior to SVMs for heart sound classification in most cases.

*Poor reproducibility issue:* Concerns about poor reproducibility of digital technologies have been raised in a variety of research areas, including biomedical engineering. In particular, the original PASCAL heart sound challenge [19] divided the data into train and test sets, but user-independent constraint was not maintained: two different sound samples from the train and test sets may have been collected from the same patient. In the present study, we address this

<sup>1</sup>The Department of Computer Science and Technology, University of Cambridge eb729@cam.ac.uk

<sup>2</sup>Department of Cardiology, University Hospitals Birmingham NHS Foundation Trust, 9 Birmingham and Institute of Cardiovascular Sciences, University of Birmingham.

\*This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/L015889/1 for the Centre for Doctoral Training in Sensor Technologies and Applications, and by ERC Project 833296 (EAR).

issue and re-evaluate the performance of our framework with user-independent constraints. This simulates the utmost authentic assessment in clinical practice and may increase the generalisability of our findings. Our results demonstrate that the proposed approach is able to detect heart murmurs from an unseen patient with 86% precision via a DNN model.

Motivated by previous encouraging studies and trying to tackle the aforementioned understudied issues for heart sound classification, in this work we propose a novel framework for heart sound classification. In particular, we focus on detecting heart murmurs among normal heart sounds and sounds with extrasystole. We develop a segmentation-free approach for the task at hand. Also, we carry out extensive experiments on the PASCAL heart sound dataset, aiming at fairly comparing the effectiveness and robustness of SVM and DNN models. More importantly, we evaluate the performance under a user-independent setting to verify our model’s generalisability to heart sound data from unseen users. Our results show that the proposed approach outperforms other state-of-the-art methods.

## II. METHODOLOGY

### A. Data and its limitations

The PASCAL Classifying Heart Sounds Challenge dataset (2011) [19] is the first large publicly available heart sound dataset. We used the PASCAL heart sound dataset B for this study, which was collected using a stethoscope. Dataset B contains three classes: normal, murmur, and extrasys (for extrasystole). Each class contains audio recordings of heart sounds, further referred to as samples. The labelled samples are split into two subgroups – noisy and “clean” (comparatively) samples. There are 149 noisy samples: 120 normal and 29 murmur, and a total of 461 clean samples: 200 normal samples, 66 with a murmur, and 46 with extrasys. Also, there are 195 unlabelled sound samples for testing. However, the user identification information can still be fetched from the unlabelled data, and thus can be re-labelled and exploited to enlarge the size of the data for training.

Worth mentioning that the patients whose samples are labelled as extrasys also have other samples in the dataset that do not have the extrasystole present and therefore are normal. Extrasys is essentially a normal heart sound, except at some point during the recording of any length there is an additional heart sound present. This knowledge was leveraged for data manipulation, described in section II-E.

### B. Data preprocessing

1) *Denoising*: Heart sound denoising was performed using a Finite Impulse Response (FIR) high-pass filter followed by a Discrete Wavelet Transform (DWT), utilising MATLAB’s Wavelet Toolbox. The high-pass filtering removed all frequencies below 60 Hz. Various parameters were tested for wavelet denoising, and the parameters that resulted in the removal of the most noise while keeping the rest of the signal intact were selected and applied to the whole dataset. Specifically, Daubechies 4 was selected as the mother wavelet, 6<sup>th</sup> decomposition level was used, heursure threshold selection

method, hard thresholding, and threshold rescaling using level-dependent estimation of level noise. Finally, before saving the denoised audio samples for feature extraction and classification, the signal was normalised in the range between -0.5 and 0.5, and centred so that the mean is equal to zero.

2) *Feature extraction*: The INTERSPEECH ComParE 2018 feature set (IS-18) [20] was extracted from the signal, resulting in a vector with 6373 features. This feature set has been applied successfully for many acoustic tasks. In order to enable extrasys detection, additional features were computed. These features were based on Shannon energy ( $E_{\text{Shannon}}$ ) calculated for the signal using the equation:

$$E_{\text{Shannon}} = -(x^2) * \log_{10} x^2, \quad (1)$$

where  $x$  denotes the signal. Then, a moving average was calculated with 100 samples window length. The resulting Shannon energy envelope was normalised, and peaks were extracted using a height threshold of 0.08 and distance of 400 samples, and the peak-to-peak distance was calculated. Finally, a number of features was extracted based on the peak-to-peak distance as well as peak timing: two largest and two smallest values, as well as a mean and a standard deviation. In addition, the total number of peaks was counted. The resulting 13 features were concatenated into a vector, and appended to the IS-18 feature vector.

While this approach still includes peak detection, it does not rely on the correct classification of S1 and S2 heart sounds, as the methods presented for extrasys detection tend to. Therefore, this approach is likely to be more robust to noise inherently captured with heart sounds.

3) *Feature reduction*: The resulting feature vector had 6386 components. The extracted features were standardised by removing the mean and scaling to unit variance, and the standardised features were then reduced by utilising principal component analysis (PCA). 460 components were retained, with variance ratio 99.99%.

### C. Classification

Python Tensorflow and Keras libraries were used for designing the classifiers and evaluating their performance. We decided to apply an SVM (with an RBF kernel) and DNN for classification, given the success reported in previous works. For the DNN, the design that yielded the best result was a 6-layer neural network, with 512 nodes in the first two layers, 256 in the middle two, and 128 nodes in the last two layers. We employed ReLU activation, the dropout 0.2 in the first layer and 0.5 on all the following layers, and Adam optimiser. The model was trained for a varying number of epochs, depending on the experiment, and the numbers are specified and explained in section II-F.

### D. Evaluation

The PASCAL challenge suggested a number of specific metrics for evaluation of the models:

- **Precision** per individual class — a measure of how many samples belonging to a specific class are labelled correctly, calculated as  $TP/(TP + FP)$ . When reporting

results, PN, PM, and PE stand for precision of normal, murmur, and extrasys respectively.

- **Sensitivity** (Sens = TP/(TP + FN)) and **specificity** (Spec = TN/(TN + FP)).
- **Youden’s index** ( $\gamma$ ) — a metric demonstrating the test’s ability to avoid failure, which can be calculated using  $\gamma = \text{Sens} + \text{Spec} - 1$ .
- **Discriminant power** (D) also summarises sensitivity and specificity:

$$D = \frac{\sqrt{3}}{\pi} \left( \log \left( \frac{\text{Sens}}{1 - \text{Sens}} \right) + \log \left( \frac{\text{Spec}}{1 - \text{Spec}} \right) \right)$$

- **Total precision** — the sum of precision values obtained per individual class. Maximum total precision for three-class problem is 3.0, and 2.0 for binary classification.

### E. Dataset manipulation

The dataset was manipulated for some of the experiments conducted as a part of this work, due to the concern that the split between train and test set suggested by the PASCAL challenge does not allow to perform cross-validation or user-independent testing. Keeping in mind the superior clinical importance of murmur detection, all extrasys files were relabelled into normal (or, rather, sounds with no murmur present), and the patient identification number was used to label all the unlabelled samples from the test set.

### F. Experiment design

Three experiments were designed for this study:

- **Exp 1 (Three-class Classification)** — trained and evaluated the model using the PASCAL dataset the way the challenge was intended: training on three classes (normal, murmur, and extrasys), and evaluating using the included locked spreadsheet that automatically calculates the suggested evaluation metrics. The deep learning model was trained for 500 epochs. Worth noting, that the train/test split is not user-independent, but the experiment was performed for method comparison with previous results reported on the PASCAL dataset.
- **Exp 2 (Binary Classification)** — relabelled all the extrasys samples to normal. As a result, instead of three-class classification, a *binary classification* (normal and murmur) was performed. Then, we also labelled all the unlabelled samples from the test set using the patient ID under the assumption that each patient has either normal or murmur heart sounds. Given that the classification problem became binary, only 50 epochs were required to train the DNN model. The main goal of this experiment was to observe the effect the increased amount of data has on the classifier performance.
- **Exp 3 (User-Independent)** — used the same data and labels as in Exp 2, but with *user-independent* train/test split. The DNN model benefited from a slight increase in the number of epochs to 100. This experiment aimed to investigate the classifier performance in a more realistic setting, on previously unseen users.

TABLE I  
RESULTS OF EXP 1.

	Previous works					Our method	
	[3]	[9]	[4]	[5]	[13]	SVM	DNN
PN	0.70	0.77	0.71	<b>0.82</b>	0.77	<b>0.82</b>	0.81
PM	0.30	0.37	0.33	0.59	0.76	0.70	<b>0.96</b>
PE	0.67	0.17	<b>1.00</b>	0.18	0.50	0.20	0.50
Sens	0.19	0.51	0.14	0.49	0.34	<b>0.54</b>	0.47
Spec	0.84	0.59	0.90	0.66	0.95	0.77	<b>0.99</b>
$\gamma$	0.02	0.01	0.04	0.15	0.29	0.31	<b>0.46</b>
D	0.04	0.09	0.09	0.15	<b>1.24</b>	0.33	0.98
TPr	1.67	1.31	2.04	1.58	2.03	1.72	<b>2.28</b>

For all of the experiments, both SVM and DNN classifiers were used for performance comparison.

Furthermore, for Exp1, we compared our approach with other state-of-the-art methods reported on the PASCAL dataset. In particular, the two winners of the challenge proposed employing segmentation-based features, with the first group utilising J48 and multi-layer perceptron (MLP) algorithms [3], while the second group proposed an algorithmic approach (no machine learning) [9]. In addition, discrete and continuous wavelet transform coefficients were explored as inputs to a random forest classifier in [4]. Segmentation-based features and a discriminant analysis classifier were used in [5]. Last, our approach was also compared with another segmentation-free method, which utilised autocorrelation features and diffusion maps [13].

## III. RESULTS AND DISCUSSION

The proposed methodology maintains high performance on the original challenge, while offering a significant improvement on the detection of murmur. It does not employ segmentation, which is highly desirable, considering segmentation’s susceptibility to noise and elevated heart rates. In addition, we demonstrate improved performance across all metrics by increasing the amount of data in the training set, managing to preserve the high performance for the user-independent split.

The results of Exp 1 and the comparison to other state-of-the-art results on the PASCAL dataset are presented in Table 1. DNN demonstrates better performance than the SVM, and the segmentation-free methodology presented in this paper yields superior results for the following metrics: the precision of murmur, the specificity of heart problem, Youden index of heart problem, and total precision. The metrics outperformed by other approaches are precision of normal, where our achieved precision is only marginally worse (1% reduction), the precision of extrasys, and discriminant power.

Exp 2 and 3 involved binary classification and increasing the amount of training data, which, as expected, resulted in increased performance (see Table 2). Worth noting that the results of the user-independent train/test split (Exp 3) yielded slightly poorer results, which is expected.

Comparing two classifiers, it could be argued that the deep learning model yields better performance in comparison to the SVM. While for the user-independent split SVM demon-

TABLE II

RESULTS OF EXP 2 AND 3. FOR EACH EVALUATION METRIC ITS MEAN AND ST DEV ARE REPORTED ACROSS FIVE-FOLD CROSS-VALIDATION.

	Exp 2		Exp 3	
	SVM	DNN	SVM	DNN
PN	0.88±0.01	0.92±0.03	0.88±0.01	0.92±0.02
PM	0.99±0.03	0.97±0.04	0.98±0.04	0.86±0.11
Sens	0.50±0.07	0.68±0.12	0.48±0.06	0.68±0.07
Spec	1.00±0.00	0.99±0.01	1.00±0.00	0.97±0.03
$\gamma$	0.50±0.07	0.67±0.12	0.48±0.07	0.65±0.09
D	N/A	N/A	N/A	N/A
TPr	1.87±0.04	1.89±0.04	1.86±0.05	1.78±0.13

It was impossible to calculate the discriminant power  $D$  across the five folds since in some folds specificity of murmur was 1, resulting in subsequent division by zero.

strates higher total precision and lower standard deviation across five folds, DNN yields significantly higher sensitivity of the heart problem.

The best performance in terms of the sensitivity of heart problem on user-independent training and testing set splits (Exp 3) was achieved by a DNN classifier, achieving 92% precision of normal and 86% precision of murmur, with 68% sensitivity of murmur.

Total precision slightly dropped for user-independent testing. This suggests that in previous works on the PASCAL dataset, which used the dataset with the data from the same patient appearing in both training and testing sets, according to the original design, inference models might have suffered from overfitting and might yield poorer performance on new samples from previously unseen patients, as a consequence of lacking generalisability and reproducibility.

While the proposed method achieves superior sensitivity of heart problem in comparison to all the approaches reported to date, it still could be deemed insufficient, given the importance of detection of heart murmurs. It is worth keeping in mind that in the PASCAL dataset every sample coming from the same patient with a murmur receives the same label, despite the murmur amplitude variation across auscultatory locations. This becomes especially problematic in patients with mild murmurs, where in distant locations murmur amplitude does not exceed the noise amplitude, masking the pathology, and thus needs further investigation.

#### IV. CONCLUSION

This paper presents a novel segmentation-free method for murmur and extrasystole detection in heart sounds. To the best of our knowledge, the proposed framework achieved the best performance on the PASCAL dataset. In addition, upon increasing the training set by manipulating the dataset and implementing user-independent testing to meet the practical needs, we reliably boost the performance by all the evaluation metrics.

Automated auscultation could become an exceptionally valuable tool for diagnostics and disease progression tracking, but the lack of publicly available granular heart sound datasets limits the research efforts in this field. The directions for future work would include expanding the models to more

detailed datasets, as well as considering all the data from a single patient combined to reduce the effect auscultatory location has on the model prediction.

#### REFERENCES

- [1] World Health Organisation, "Cardiovascular diseases (CVDs)," [www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-cvds](http://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-cvds).
- [2] D. Gradolewski and G. Redlarski, "Wavelet-based denoising method for real phonocardiography signal recorded by mobile devices in noisy environment," *Computers in Biology and Medicine*, vol. 52, pp. 119–129, 2014.
- [3] E. F. Gomes and E. Pereira, "Classifying heart sounds using peak location for segmentation and feature construction," in *AISTATS: Classifying Heart Sounds Workshop*, 2012.
- [4] C. C. Balili, M. C. C. Sobrepena, and P. C. Naval, "Classification of heart sounds using discrete and continuous wavelet transform and random forests," in *Proc. IAPR Asian Conference on Pattern Recognition*, 2015, pp. 655–659.
- [5] F. Chakir, A. Jilbab, C. Nacir, and A. Hammouch, "Phonocardiogram signals processing approach for PASCAL classifying heart sounds challenge," *Signal, Image and Video Processing*, vol. 12, no. 6, pp. 1149–1155, 2018.
- [6] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, and N. Yen, "The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis," in *Proceedings of the Royal Society*, 1998, vol. 454, pp. 903–995.
- [7] H. L. Cherif and S. M. Debbal, "Algorithm for detection of the internal components of the heart sounds and their split using a Hilbert transform," *Journal of Medical Engineering and Technology*, vol. 37, no. 3, pp. 220–230, 2013.
- [8] L. N. Sharma, "Multiscale analysis of heart sound for segmentation using multiscale Hilbert envelope," in *Proc. International Conference on ICT and Knowledge Engineering*, 2015, pp. 33–37.
- [9] Yiqi Deng, "A robust heart sound segmentation and classification algorithm using wavelet decomposition and spectrogram," in *Proc. AISTATS: Classifying Heart Sounds Workshop*, 2012.
- [10] A. K. Dwivedi, S. A. Imtiaz, and E. Rodriguez-Villegas, "Algorithms for automatic analysis and classification of heart sounds—a systematic review," *IEEE Access*, vol. 7, pp. 8316–8345, 2019.
- [11] M. H. Asmare, F. Woldehanna, L. Janssens, and B. Vanrumste, "Rheumatic heart disease detection using deep learning from spectro-temporal representation of un-segmented heart sounds," in *Proc. EMBC*, 2020, pp. 168–171.
- [12] P. Langley and A. Murray, "Heart sound classification from unsegmented phonocardiograms," *Physiological Measurement*, vol. 38, no. 8, pp. 1658–1670, 2017.
- [13] Shi-Wen Deng and Ji-Qing Han, "Towards heart sound classification without segmentation via autocorrelation feature and diffusion maps," *Future Generation Computer Systems*, vol. 60, pp. 13–21, 2016.
- [14] M. Tschannen, T. Kramer, G. Marti, M. Heinzmann, and T. Wiatowski, "Heart sound classification using deep structured features," in *Proc. Computing in Cardiology Conference*, 2016, pp. 565–568.
- [15] B. Whitaker and D. Anderson, "Heart sound classification via sparse coding," in *Proc. Computing in Cardiology Conference*, 2016, pp. 805–808.
- [16] T. Ölmez and Z. Dokur, "Classification of heart sounds using an artificial neural network," *Pattern Recognition Letters*, vol. 24, no. 1–3, pp. 617–629, 2003.
- [17] T. Nilanon, J. Yao, J. Hao, S. Purushotham, and Y. Liu, "Normal / abnormal heart sound recordings classification using convolutional neural network," in *Proc. Computing in Cardiology Conference*, 2016, pp. 585–588.
- [18] F. Noman, C. Ting, S. Salleh, and H. Ombao, "Short-segment heart sound classification using an ensemble of deep convolutional neural networks," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2019, pp. 1318–1322.
- [19] P. Bentley, G. Nordehn, M. Coimbra, and S. Mannor, "The PASCAL Classifying Heart Sounds Challenge 2011 (CHSC2011) Results," <http://www.peterjbentley.com/heartchallenge/index.html>.
- [20] B. Schuller, S. Steidl, A. Batliner, et al., "The INTERSPEECH 2018 computational paralinguistics challenge: Atypical and self-assessed affect, crying and heart beats," in *Proc. INTERSPEECH*, 2018, pp. 122–126.