

Beat-wise Uncertainty Learning for Murmur Detection in Heart Sounds

Xingyao Wang¹, Foli Fan¹, Hongxiang Gao¹, Shuo Zhang¹, Chenxi Yang¹, Jianqing Li¹, Chengyu Liu¹

¹ State Key Laboratory of Bioelectronics, School of Instrument Science and Engineering, Southeast University, Nanjing, China

Abstract

This paper introduces a murmur detection solution (Team SeaCrying) to the PhysioNet Challenge 2022. The method is based on beat-wise uncertainty learning for heart sounds. The target task is to distinguish the present and absent state for murmur, with an outlier situation indicated as unknown in the challenge. Two uncertainties induced by outlier noise and fuzzy sounds are addressed while beat segmentation and murmur discrimination, respectively. In beat segmentation stage, we employ a confidence branch trained by a frame-level noise contrastive framework to quantify the uncertainty for out-of-distribution episodes. Then we transmit the groups of five effective heart beats to the murmur discriminator and each beat is concatenated by a systole (containing S1 and S2) and a diastole. To alleviate the issue of disability for the model learning unknown sounds, we adopt an uncertainty estimation module on the basis of binary classification for murmur detection. The unknown samples will lead to a highly uncertainty score. As well a cross-beat decision strategy is designed for the same phonocardiogram recording in the final stage. Our contribution submitted for evaluation achieved a weighted accuracy of 0.601 on the test set in murmur detection and 0.837 for the outcome classification task. The cost metric score the test set is 12126 for murmur detection and 15083 for outcome classification.

1. Introduction

Cardiac auscultation, for identifying heart sounds, is commonly the first step and the most cost-effective measure for screening the various heart dysfunction, even though the final diagnosis is based on the combined analysis from a series of electrophysiologic study or ultrasound recordings. Heart sounds can reflect the hemodynamic processes of the heart and identify some representative symptoms of different diseases, including arrhythmia, valve disease, pulmonary hypertension, heart failure, among other issues [1]. However, only about 20% of medical interns can effectively detect heart conditions using

auscultation[2], and extensive training is necessary for human expert evaluation. Automatic and accurate analysis of the recording of heart sounds (the phonocardiogram, or PCG) can be useful for auxiliary diagnosis in clinical applications, and it can potentially assist interns with less developed skills.

The aim of the Challenge [3] is to identify the presence, absence, or unclear cases of murmurs and the normal vs. abnormal clinical outcomes from heart sound recordings collected from multiple auscultation locations on the body using a digital stethoscope. Through the observation of the recordings annotated with unknown murmur, we have found out that these recordings can be regarded as ambiguous data for murmur detection. Previously, many researchers were dedicated themselves to extract features from both the time and frequency domain [4, 5] as well the morphology characteristics for anomaly detection in heart sounds. Nowadays, deep learning (DL) have also been applied in murmur detection [6–8], yet the DL-based methods cannot quantify the confidence for predicted results of the ambiguous heart sounds.

Different from the PhysioNet/CinC Challenge 2016, this year's challenge introduces unknown annotations provided by experts as the third class besides presence and absence for murmur. On account of this change, we introduce uncertainty measurement in our method. There are two types of uncertainty - epistemic uncertainty (EU) and aleatoric uncertainty (AU). By definition, epistemic uncertainty is inherent to the model, caused by a lack of training data and aleatoric uncertainty originates from the inherent noise in data and hence irreducible with more data [9]. Therefore, the EU can be utilized to screen the outlier noise without heart sound information while distinguishing the ambiguous recordings with the AU. In this paper, we conduct a trial to decompose the challenge into two stages, beat segmentation and murmur detection. While beat segmentation, we screen out the noisy episodes referring to the estimated EU since the criterion to decide whether a physical signal belongs to heart sounds or not is to find successive heartbeats therein. The AU is adopted in murmur detection to recognize unknown recordings.

2. Method

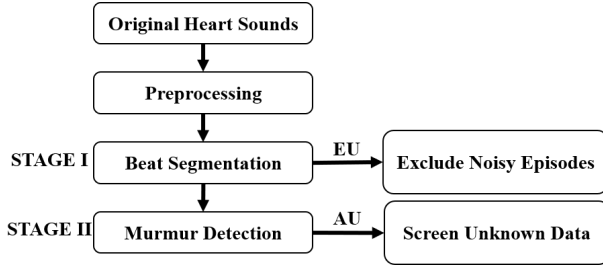


Figure 1. Flowchart diagram of the proposed method for the detection of murmurs.

To address the unknown data, we introduce uncertainty measurement in our murmur detection framework, which is separated as beat-segmentation and murmur-detection stages. In effect, there are two factors inducing the unclear cases of murmurs, one is the noisy episodes and the other is ambiguous sounds caused by non-standard auscultation or mixed breath. Here invalid heart sounds can be deemed as the episodes without identifiable beat information, which contains the characteristics of rhythm and morphology. Thus we measure the chaos of the beat-segmentation results via epistemic uncertainty to exclude the noisy episodes. Then we will select five beats with the highest confidence among the valid episodes provided by the beat-segmentation stage for murmur detection. Similar uncertainty estimation is deployed in murmur-detection stage to screen the beats with high aleatoric uncertainty. A strong assumption is set in this stage that the unknown heart sounds represent the suspected murmur that is ambiguous in the experts’ cognition. Figure 1 outlines the architecture of our proposed algorithm.

2.1. Preprocessing




Figure 2. Illustration for beat annotations in heart sounds.

In this challenge, 3163 heart sound recordings from 942 patients are utilized for training and each recording has an uncertain length [10]. The public training set contains heart sound recordings, routine demographic information, murmur-related labels (presence, absence, or unknown), outcome-related labels (normal or abnormal), an-

notations of the murmur characteristics (location, timing, shape, pitch, quality, and grade), and heart sound segmentations.

All the recordings are downsampled at a sampling rate of 1000 Hz. Then we implement a bandpass filter with 15-300 Hz pass-band on the input heart sounds. An adaptive Wiener filter is adopted in the next step, which has been proved owning capacity to suppress the in-band noise varied in distribution induced by individual difference in our previous study [11]. As it is difficult for training the model with non-identical length of heart sound recordings, in the preprocessing phase, each heart sound record was divided into segments with a length of 5 seconds for beat segmentation. As above depicted, 5 beats with the lowest uncertainty are chosen for the next stage, murmur detection. We unify the beat length as 1000 points and each data point for murmur-detection model is a $1000 * 5$ matrix.

For training the beat segmentation model, Training-A in the 2016 PhysioNet/CinC Challenge [12] is utilized since it is the only database which contains simultaneously recorded PCGs and ECGs. Except for Training-A, we also integrate the recordings with segmentation annotation in this year’s challenge. Eventually, 15366 5-second episodes with state annotation are prepared and 10% are divided as evaluation set. Note that each beat contains the three states (S1, systole, S2) according to the general annotation for heart sound segmentation, which is a complete cardiac cycle. Therefore, we distribute a binary label for each frame, among 1 represents S1, systole and S2 and 0 represents diastole. Figure 2 gives a specific describe for the beat annotation. For murmur detection, 90% of the individuals with the murmur-related labels of presence, absence and unknown are divided for training and the remaining 10% for validation.

2.2. Network structure

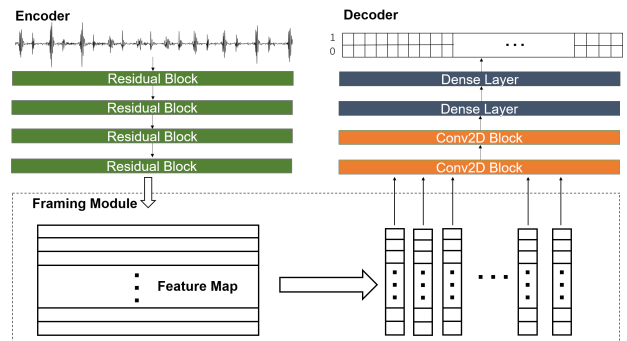


Figure 3. Encoder-Decoder architecture for beat segmentation in heart sounds.

As shown in Figure 3, we simplify the temporal-framing adaptive network (TFAN) in [11] to acclimatize the beat

segmentation task with binary annotations. For murmur detection, we design a light-weight Squeeze-and-excitation network (SENet) [13] as the Encoder and implement two fully connected layers as Decoder. For facilitating the description of uncertainty estimation in the following section, we define the Encoder as f and the Decoder as g .

2.3. Uncertainty Estimation

In order to disentangle the epistemic uncertainty (EU) and aleatoric uncertainty (AU), we refer to the formalization defined by Malinin and Gales [14]:

$$\underbrace{IG[y; \theta | \mathbb{D}]}_{\text{EU}} = \underbrace{H[E_{p(\theta|\mathbb{D})}[P(y | \theta)]]}_{\text{Total Uncertainty}} - \underbrace{E_{P(\theta|\mathbb{D})}[H[p(y | \theta)]]}_{\text{Expected AU}}, \quad (1)$$

where IG indicates the information gain (IG) brought by the θ , the model parameters, which always need complex Bayesian neural networks or ensemble models to approximate the ground-truth distribution. Considering the methods in the challenge should not consuming redundant calculation, we implement uncertainty estimation in a deterministic network through rethinking IG.

Recall that the definition of IG between the latent representation z and the label y is

$$\begin{aligned} IG[z; y | \theta] &= IG[z; y] \\ &= KL\left(\sum_{\phi} p(z | \phi, y)p(\phi | z, y) \| P(z)\right), \end{aligned} \quad (2)$$

where ϕ is the assumed class-biased transformation conditioned on the label y . In this paper, we decompose the three-class detection for murmur into two binary classification tasks, beat segmentation and presence/absence detection for murmur. Therefore, only two transformations are needed for uncertainty estimation and we named the two transformations as ϕ^+ and ϕ^- , respectively. Thus, EU can be approximated by

$$\begin{aligned} EU &= D_{KL}(P(\phi(z)) \| P(z)) \\ &\simeq \sum_{\phi} g(\phi(z)) \log \frac{g(\phi(z))}{g(z)}. \end{aligned} \quad (3)$$

And for estimating AU, according to Equation 1, the straightforward approach is to calculate the expected entropy, which is

$$\begin{aligned} AU &= -(g(\phi^+(z)) \log g(\phi^+(z)) \\ &\quad + (1 - g(\phi^-(z))) \log(1 - g(\phi^-(z)))). \end{aligned} \quad (4)$$

According to the above definition of the class-biased transformations, ϕ should be able to maintain the feature

invariance of the corresponding class. Aiming at this purpose, we design the two loss functions for ϕ^+ and ϕ^- as follows:

$$\mathcal{L}_{\phi^+} = \mathcal{D}(\phi^+(\ddagger), \nabla^+ | \ddagger = \infty) + \mathcal{D}(\phi^+(\ddagger), \ddagger | \ddagger = r), \quad (5)$$

$$\mathcal{L}_{\phi^-} = \mathcal{D}(\phi^-(\ddagger), \nabla^- | \ddagger = r) + \mathcal{D}(\phi^-(\ddagger), \ddagger | \ddagger = \infty), \quad (6)$$

where r^+ and r^- are the approximate marginal distribution for the latent representations belong to $y = 1$ and $y = 0$. Here we adopt the Monte-Carlo sampling method to estimate the marginal distribution among a random-sampled batch in each training epoch. Cosine distance is chosen in this work to measure the distance, D .

3. Results

Table 1. The overall murmur scores of the proposed method on the independent hidden test set. W-Acc here indicates the weighted accuracy defined by the challenge.

Dataset	AUROC	AUPRC	F_1	W-Acc
Training	0.913	0.746	0.562	0.736
Validation	0.783	0.626	0.478	0.609
Test	0.793	0.612	0.461	0.601

Table 2. The murmur scores for each class of the proposed method on the independent evaluation set divided from the training set.

Classes	Present	Unknown	Absent
AUROC	0.987	0.825	0.934
AUPRC	0.948	0.331	0.973
F_1	0.879	0.257	0.672
Acc	0.950	0.838	0.509

The proposed method is tested on the training, validation and test sets respectively. The results show that the proposed method owns high sensitivity on murmur, yet relative lower specificity for the absent and unknown states. The weighted accuracy of our method is 0.609 on the validation set and 0.601 on the test set, decreased by 0.127 and 0.135 comparing to the overall performance on the training set.

4. Discussion and Conclusion

In this work, we introduce uncertainty measurement to screen out the unknown heart sounds for murmur detection while identifying absence and presence of murmur for the first time. Different from the general supervised learning, the proposed method will not learn the direct mapping from input heart sounds to the unknown class. Instead, the ambiguous recordings are recognized based on

AU. Here the ambiguous recordings are assumed as the suspected murmur that is difficult for doctors to confirm. The results show that our method has a certain accuracy in detecting murmur, yet tends to produce more false positives facing recordings with and without murmur. Due to the time limit of the challenge, we have not conducted a detailed study of the data and methods. In the future, we will further study how to model experts' perception on unknown murmur based on uncertainty estimation.

Acknowledgment

This work is supported by the National Natural Science Foundation of China (62001111, 62171123, 62071241 and 81871444), the National Key Research and Development Program of China (2019YFE0113800) and the Natural Science Foundation of Jiangsu Province (BK20200364, BK20190014 and BK20192004).

References

- [1] Tavel ME. Cardiac auscultation: a glorious past—but does it have a future? *Circulation* 1996;93(6):1250–1253.
- [2] Mangione S, Nieman LZ, Gracely E, Kaye D. The teaching and practice of cardiac auscultation during internal medicine and cardiology training: a nationwide survey. *Annals of Internal Medicine* 1993;119(1):47–54.
- [3] Goldberger AL, Amaral LA, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *Circulation* 2000;101(23):e215–e220.
- [4] Ahmad MS, Mir J, Ullah MO, Shahid MLUR, Syed MA. An efficient heart murmur recognition and cardiovascular disorders classification system. *Australasian Physical Engineering Sciences in Medicine* 2019;42(3):733–743.
- [5] Levin AD, Ragazzi A, Szot SL, Ning T. Extraction and assessment of diagnosis-relevant features for heart murmur classification. *Methods* 2022;202:110–116.
- [6] Dominguez-Morales JP, Jimenez-Fernandez AF, Dominguez-Morales MJ, Jimenez-Moreno G. Deep neural networks for the recognition and classification of heart murmurs using neuromorphic auditory sensors. *IEEE Transactions on Biomedical Circuits and Systems* 2017; 12(1):24–34.
- [7] DeGroff CG, Bhatikar S, Hertzberg J, Shandas R, Valdes-Cruz L, Mahajan RL. Artificial neural network-based method of screening heart murmurs in children. *Circulation* 2001;103(22):2711–2716.
- [8] Vakamullu V, Mishra M, Mukherjee A. Real-time heart murmur classification using attention based deep learning approach. In *2022 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*. IEEE, 2022; 1–6.
- [9] Kendall A, Gal Y. What uncertainties do we need in bayesian deep learning for computer vision. In *Neural Information Processing Systems*. 2017; .
- [10] Oliveira J, Renna F, Costa PD, Nogueira M, Oliveira C, Ferreira C, Jorge A, Mattos S, Hatem T, Tavares T, et al. The circor digiscope dataset: from murmur detection to murmur classification. *IEEE Journal of Biomedical and Health Informatics* 2021;26(6):2524–2535.
- [11] Wang X, Liu C, Li Y, Cheng X, Li J, Clifford GD. Temporal-framing adaptive network for heart sound segmentation without prior knowledge of state duration. *IEEE Transactions on Biomedical Engineering* 2020;68(2):650–663.
- [12] Clifford GD, Liu CY, Moody B, Springer D, Silva I, Li Q, Mark RG. Classification of normal/abnormal heart sound recordings: The Physionet/Computing in Cardiology Challenge 2016. In *2016 Computing in Cardiology Conference (CinC)*. 2016; 609–612.
- [13] Hu J, Shen L, Sun G. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018; 7132–7141.
- [14] Malinin A, Gales M. Predictive uncertainty estimation via prior networks. *Advances in Neural Information Processing Systems* 2018;31.

Address for correspondence:

Chengyu Liu
 The School of Instrument Science and Engineering
 Southeast University, 2nd Sipailou
 Nanjing, Jiangsu, China
 E-mail: chengyu@seu.edu.cn