

TOPSEG: A MULTI-SCALE TOPOLOGICAL FRAMEWORK FOR DATA-EFFICIENT HEART SOUND SEGMENTATION

Peihong Zhang, Zhixin Li, Yuxuan Liu, Rui Sang, Yiqiang Cai, Yizhou Tan, Shengchen Li

School of Advanced Technology, Xi'an Jiaotong-Liverpool University, Suzhou, China

ABSTRACT

Deep learning approaches for heart-sound (PCG) segmentation built on time–frequency features can be accurate but often rely on large expert-labeled datasets, limiting robustness and deployment. We present TopSeg, a topological representation-centric framework that encodes PCG dynamics with multi-scale topological features and decodes them using a lightweight temporal convolutional network (TCN) with an order- and duration-constrained inference step. To evaluate data efficiency and generalization, we train exclusively on PhysioNet 2016 dataset with subject-level subsampling and perform external validation on CirCor dataset. Under matched-capacity decoders, the topological features consistently outperform spectrogram and envelope inputs, with the largest margins at low data budgets; as a full system, TopSeg surpasses representative end-to-end baselines trained on their native inputs under the same budgets while remaining competitive at full data. Ablations at 10% training confirm that all scales contribute and that combining H_0 and H_1 yields more reliable S1/S2 localization and boundary stability. These results indicate that topology-aware representations provide a strong inductive bias for data-efficient, cross-dataset PCG segmentation, supporting practical use when labeled data are limited.

Index Terms— Heart Sound Segmentation, Topological Data Analysis, Phonocardiogram Signals

1. INTRODUCTION

Heart disease is the leading cause of mortality worldwide, responsible for approximately 17.9 million deaths annually [1], with auscultation serving as a primary, cost-effective diagnostic tool in clinical practice [2]. Accurate segmentation of phonocardiogram (PCG) signals into their primary components—first (S1) and second (S2) heart sounds, and the intervening systolic and diastolic intervals— aids physicians in cardiac diagnosis by helping identify timing abnormalities, detect murmurs, and assess valvular function [3]. Fig. 1 illustrates the segmentation task, where PCG signals are segmented into S1, S2, systolic, and diastolic phases.

Recent advances indicate that deep learning models achieve high performance in PCG segmentation by learning representations from spectro-temporal inputs such as spectrograms [4], mel-spectrograms [5], and scalograms [6], or directly from raw waveforms through end-to-end architectures [7]. However, when spectro-temporal features are used in deep neural networks, high performance typically relies on large-scale annotated datasets, and performance drops sharply in data-efficient scenarios [8, 9]. While handcrafted spectro-temporal features applied to machine learning models can reduce data requirements, they often fail to capture the complex temporal structures of PCGs, thereby limiting their overall effectiveness [7].

In addition to spectro-temporal features, some studies have explored envelope-based and statistical representations for PCG




Fig. 1: PCG signal segmented into S1, systole, S2, and diastole.

segmentation, which require fewer labeled data compared to spectro-temporal approaches [10]. A widely used approach is the logistic regression hidden semi-Markov model built upon Hilbert and homomorphic envelopes [11]. Variants have since incorporated zero-frequency filtering combined with Hilbert envelopes [12] to improve boundary detection. More recently, segmentation approaches have utilized clustering of envelope-derived features and wavelet envelopes [13] to adaptively locate S1 and S2 events. However, their effectiveness is constrained by inherent limitations: envelope and statistical features are highly sensitive to background noise, prone to misidentification in pathological or irregular heart sounds, and insufficient for capturing the complex temporal–structural patterns present in PCGs [14]. Consequently, while such features reduce data requirements, their lack of robustness and descriptive power restricts their utility for reliable segmentation in diverse clinical scenarios.

Despite these advances, PCG segmentation remains a challenging task. Annotated datasets are limited because labeling PCG signal requires the expertise of trained clinicians, which makes large-scale annotation difficult [5]. Moreover, the quality of recordings is often compromised by environmental noise and patient-specific factors, further complicating the process of reliable segmentation [7]. These challenges naturally raise the question: are there alternative feature representations that are robust and fully capture the structural characteristics of PCG signals under data-efficient conditions?

Topological data analysis (TDA) offers a promising direction to address these challenges by extracting features that capture the underlying structural organization of the data. Topological features are inherently robust to small perturbations and noise [15], while their topological structural priors reduce reliance on large annotated datasets [16], making them particularly suitable for data-efficient scenarios. This is especially relevant for PCG signals, which are often corrupted by respiratory noise and patient-specific artifacts and available only in limited annotated quantities. Although TDA has been explored in other biomedical domains [17], its application to PCG segmentation remains unexplored.

In this work, we propose TopSeg, a heart sound segmentation framework that leverages multi-scale topological representations to

capture the intrinsic structural patterns of PCGs. Our key insight is that cardiac cycles form persistent structural patterns across multiple physiological timescales, which remain stable under noise perturbations and can be reliably captured even with limited annotated data. To realize this idea, we extract multi-scale topological features at three complementary temporal resolutions: global rhythm patterns (2–8 s), individual cardiac cycles (≈ 500 ms), and fine-grained S1/S2 components (≈ 100 ms). These features are incorporated into multiple baseline segmentation models as complementary inputs, enriching their representation capacity. Extensive evaluations on several benchmark datasets demonstrate that integrating multi-scale topological features consistently enhances segmentation accuracy and improves improved generalization under data-efficient conditions.

In summary, our contributions are twofold:

- We propose TopSeg, a topological framework for PCG segmentation; to our knowledge, this is the first to operationalize multi-scale topological descriptors for this task. These descriptors can be used alone or seamlessly integrated into standard time–frequency baselines with minimal changes.
- The topological features extracted by TopSeg impart a geometry-aware inductive bias that promotes morphology and scale consistency, underpinning domain generalization and enabling robust cross-dataset transfer under limited data.

2. TOPOLOGICAL DATA ANALYSIS

TDA provides multi-scale topological invariants that are provably stable to small perturbations of the input [15, 18]. Its central construct, persistent homology (PH), assigns to a filtration $\{\mathcal{K}_\varepsilon\}_{\varepsilon \geq 0}$ of simplicial complexes the collection of birth–death pairs $D = \{(b_i, d_i)\}$, where b_i (resp. d_i) denotes the scale at which a homological feature appears (resp. disappears). Features with large persistence $d_i - b_i$ are typically interpreted as signal; short-lived features are attributed to noise. These stability and scale-selection properties make PH attractive for label-efficient signal analysis [16].

To obtain finite-dimensional, differentiable descriptors, we map a persistence diagram D to persistence landscapes [19]. Each pair $(b, d) \in D$ induces the tent function

$$f_{(b,d)}(\varepsilon) = \max(0, \min(\varepsilon - b, d - \varepsilon)),$$

and the k -th landscape is the pointwise k -maximum envelope

$$\lambda_k(\varepsilon) = \text{kmax} \{ f_{(b,d)}(\varepsilon) : (b, d) \in D \}.$$

Sampling $\{\lambda_k\}_{k=1}^K$ on a one-dimensional grid of size G yields a fixed-length representation that preserves PH stability in L^p norms and admits straightforward averaging, thereby facilitating integration with learning-based models [19].

3. PROPOSED TOPSEG FRAMEWORK

TopSeg has three stages (Fig. 3): (1) multi-scale topological encoder that extracts descriptors via time-delay embedding and persistent homology; (2) temporal decoder that maps the topological descriptors to framewise posteriors; and (3) inference-time convex refinement that enforces physiological consistency.

We operationalize TDA’s robustness by using multi-scale persistence descriptors as inputs and coupling them with a lightweight decoder plus a constraint-aware inference layer.




Fig. 2: TDA overview. (a) Time-delay embedding of a heart-sound segment; repeating cardiac cycles induce recurrent structures in the reconstructed phase space. (b) Filtration of simplicial complexes indexed by the scale parameter ε (illustrated with an α -complex), from which PH collects birth–death pairs for H_0/H_1 . (c) Persistence landscapes $\{\lambda_k\}$ obtained from the diagram provide stable, fixed-length descriptors amenable to learning.

Table 1: Multi-scale configuration. $W = (d - 1)\tau$. Global descriptors (2/4/8 s) are averaged into a single global stream.

Scale	τ (ms)	d	W (s)	Physiological target
Global-2s (60 Hz)	100	21	2.0	Multi-beat rhythm
Global-4s (60 Hz)	200	21	4.0	Long context
Global-8s (60 Hz)	200	41	8.0	Rhythm trend
Meso (600 Hz)	25	21	0.50	One cardiac cycle (60–120 bpm)
Fine (600 Hz)	10	11	0.10	S1/S2 morphology (80–150 ms)

3.1. Multi-scale Topological Feature Extraction

We extract topology-aware descriptors at three physiologically grounded scales: global (multi-beat rhythm), meso (single-cycle morphology), and fine (S1/S2 onsets).

For each scale $\ell \in \{\text{global, meso, fine}\}$, let $s_\ell(t)$ denote the prefiltered waveform. A delay-coordinate embedding with delay τ_ℓ and dimension d_ℓ is formed as

$$\Phi_\ell(t) = (s_\ell(t), s_\ell(t + \tau_\ell), \dots, s_\ell(t + (d_\ell - 1)\tau_\ell)),$$

covering an effective window $W_\ell = (d_\ell - 1)\tau_\ell$.

Scales and parameterization. W_ℓ is set by physiological targets. (τ_ℓ, d_ℓ) are initialized via average mutual information (AMI) and refined within a small neighborhood to match the intended scale. To balance context and cost, the global branch operates at 60 Hz, whereas meso/fine run at 600 Hz, shown in Table 1.

From embeddings to topological descriptors. Let $\mathcal{P}_\ell = \{\Phi_\ell(t)\}$ be the embedded trajectory. A window of length $L_\ell = mW_\ell$ slides with hop equal to one frame (global: 1/60 s; meso/fine: 1/600 s). For each window centered at time t , we construct a sparsified Vietoris–Rips filtration from a k -NN graph ($k = \lceil \sqrt{n} \rceil$, n points), clipping the filtration radius at the upper distance quantile $q \in [0.90, 0.99]$. Persistent homology is computed in H_0 and H_1 to obtain birth–death pairs $\mathcal{D}_\ell$, which are mapped to persistence landscapes with K layers sampled on a grid of G points, yielding a




Fig. 3: Illustration of the TopSeg framework.

fixed-length vector per window. For the global scale, the 2s/4s/8s landscapes are aggregated by element-wise mean into a single global descriptor. Window-centered descriptors are assigned to the center frame; linear interpolation across overlaps mitigates boundary drift and yields smoothly varying per-frame features. Concatenating the global, meso, and fine descriptors gives

$$\mathbf{f}_{\text{topo}}(t) \in \mathbb{R}^{D_{\text{topo}}}, \quad D_{\text{topo}} = 3 \times (2KG),$$

with each scale contributing $2KG$ dimensions (both H_0 and H_1).

3.2. Topology-Only Constraint-Aware Refinement

We refine the frame-wise class posteriors produced by the topological decoder via a small convex program that encodes physiological priors. Training optimizes the network on raw posteriors; refinement is applied only at inference.

Topology target. From fine-scale landscapes $\{\lambda_{\text{fine},h}^{(k)}(t)\}$ ($h \in \{0, 1\}$) we build

$$r_{\text{topo}}(t) = \text{Norm} \left(\sum_k \lambda_{\text{fine},0}^{(k)}(t) + \sum_k \lambda_{\text{fine},1}^{(k)}(t) \right) \in [0, 1],$$

where $\text{Norm}(\cdot)$ is percentile-based within a short sliding window. $r_{\text{topo}}(t)$ serves as a fixed, non-trainable guide.

Topology-derived reliability. To modulate the alignment strength using only topology, we define

$$\eta(t) = \sigma(\gamma[r_{\text{topo}}^{\text{EMA}}(t) - \tau]) \in [0, 1],$$

where $r_{\text{topo}}^{\text{EMA}}(t) = \rho r_{\text{topo}}^{\text{EMA}}(t-1) + (1-\rho)r_{\text{topo}}(t)$ is an EMA with decay ρ (initialized at $r_{\text{topo}}(0)$), τ is a mid-level threshold, and γ controls the sharpness.

Refinement objective (inference-only). Let $P_{\text{evt}}(t) = P_{S1}(t) + P_{S2}(t)$. For each recording we refine $\hat{P}$ by solving

$$\begin{aligned} \min_{\{P(t)\}} \quad & \sum_t \|P(t) - \hat{P}(t)\|_2^2 + \lambda_s \sum_t \|P(t) - P(t-1)\|_2^2 \\ & + \lambda_b \sum_t (P_{\text{evt}}(t) - \theta_{\max})_+^2 \\ & + \lambda \sum_t \eta(t) \|P_{\text{evt}}(t) - r_{\text{topo}}(t)\|_2^2, \end{aligned} \quad (1)$$

$$\text{s.t. } P(t) \in \Delta^C \ (C=4), \forall t,$$

Table 2: TopSeg hyperparameters (defaults and compact ranges).

Block	Symbol	Default	Notes / Range
Embedding (per scale)	W_ℓ L_ℓ	Tab. 1 mW_ℓ	physiology $m=2$ (1.5–2.5)
PH pipeline (per window)	k q (K, G)	$\lceil \sqrt{n} \rceil$ 0.95 (5, 128)	k-NN 0.90–0.99 landscape layers
Topology cues	$\eta(t)$	$\sigma(\gamma[r_{\text{topo}}^{\text{EMA}} - \tau])$	$\gamma=2.0, \tau=0.5, \rho=0.90$
Refine (1)	λ_s λ_b λ $\theta_{\max}$	1×10^{-2} 5×10^{-2} 5×10^{-2} 0.65	$[5 \times 10^{-3}, 2 \times 10^{-2}]$ $[2 \times 10^{-2}, 8 \times 10^{-2}]$ $[2 \times 10^{-2}, 8 \times 10^{-2}]$ event-mass cap
Solver	N_{iter}	8	monotone PGD

with $(\cdot)_+ = \max(0, \cdot)$. The terms respectively (i) keep P close to the network posterior, (ii) impose Tikhonov temporal smoothness, (iii) cap instantaneous event mass, and (iv) softly align P_{evt} with topology peaks under the topology-derived reliability $\eta(t)$. All penalties are convex in P ; with $\lambda_s > 0$ and simplex projection, the objective is strongly convex and admits a unique minimizer. We use projected proximal gradient; temporal coupling is tridiagonal, so each iteration costs $O(TC)$. A small fixed number of iterations is used per recording (Tab. 2).

4. EXPERIMENT AND RESULT

4.1. Dataset

We train on the PhysioNet/CinC 2016 database [20] and simulate data-efficient learning by randomly subsampling subjects at 5%, 10%, 25%, 50%, and 100% of the training pool; the corpus comprises 3,153 recordings from 764 subjects. For external validation, we use the CirCor DigiScope Dataset [21] with 5,272 recordings from 1,568 subjects. Since PhysioNet 2016 does not provide frame-wise segmentation for all files, we obtain S1/S2 labels by using the official logistic-regression HSMM segmenter released with the Challenge [11] and, where available, we preferentially adopt the hand-corrected training annotations.

4.2. Data Preprocessing

PCG recordings are first band-limited to 20–200 Hz using a zero-phase Butterworth band-pass filter [23], and then downsampled to 600 Hz with a polyphase anti-aliasing decimator to prevent aliasing. Signals are z-score normalized within each recording. For training we use 10 s segments: longer recordings are split into non-overlapping 10 s chunks, while shorter ones are looped to 10 s to preserve batch shape without changing local dynamics. The 20–200 Hz passband preserves the dominant energy of S1/S2 while suppressing low-frequency motion artifacts and high-frequency ambient noise.

4.3. Comparison Models

To demonstrate both the data-efficiency of our representation and the overall effectiveness of our pipeline, we compare along two complementary axes. First, in a feature-controlled setting, we train the same Temporal Convolutional Network (TCN) [24] decoder and a shallow MLP (1 hidden layer, 128 units) on three per-frame inputs—our topological landscapes, log-mel spectrograms (64 bins), and Hilbert

CirCor DigiScope external validation (macro-F1, %, 60 ms tolerance)					
Model	Train 5%	Train 10%	Train 25%	Train 50%	Train 100%
<i>Q1: Feature-controlled (same decoder; different inputs)</i>					
MLP (envelope)	43.0	48.1	54.3	58.2	62.5
MLP (log-mel)	48.4	53.2	60.5	64.7	67.9
MLP (topological, ours)	54.6	59.3	65.2	69.1	72.8
TCN (envelope)	55.2	61.3	67.7	71.6	75.4
TCN (log-mel)	58.1	64.2	71.3	76.4	80.2
TCN (topological, ours)	64.1	70.4	76.2	80.5	83.1
<i>Q2: End-to-end comparators (native inputs; same budgets)</i>					
LR-HSMM (envelope, native) [11]	54.3	57.2	60.3	63.2	66.1
CLSTM (raw audio, native) [7]	58.4	64.5	72.3	77.2	81.3
U-Net (time-frequency, native) [4]	59.2	65.3	73.1	78.4	82.5
FFT+CNN U-Net (time-frequency) [22]	60.5	66.3	74.2	79.4	83.2
Ours: Topological features + TCN (order/duration-constrained decoding)	66.7	71.9	78.4	82.1	85.3

Table 3: External validation on CirCor with models trained on PhysioNet under different data budgets.

envelope—thereby isolating the contribution of representations under identical model capacity. Second, at the architecture level, we include end-to-end models trained on their native inputs: the envelope-based LR-HSMM [11], a time-frequency CLSTM [7], a U-Net-style TF model [4], and an FFT+CNN U-Net [22]. All models use the same optimizer, subject-level splits, early stopping, and a 60 ms boundary tolerance; when indicated, we apply the same convex refinement at inference to enforce left-to-right ordering and minimal durations. This protocol tests (i) whether our representation yields superior accuracy in low-data regimes when decoded by matched-capacity models, and (ii) whether our full system outperforms representative architectures under identical data budgets.

4.4. Results under Data-Efficient Training

Table 3 summarizes external validation on CirCor (four-state macro-F1 with a 60 ms boundary tolerance), with all models trained on PhysioNet under subject-level subsampling and identical optimization.

Q1 (Representation). Under matched-capacity decoders (shallow MLP and the same TCN), our topological features dominate log-mel and envelope across all data budgets. The advantage is most pronounced in the low-data regime (5–25%) and remains positive at full data, indicating that the representation carries stronger inductive bias and learns reliably with limited supervision.

Q2 (Overall model). Our full pipeline—topological features decoded by a lightweight TCN with order/duration-constrained decoding—outperforms representative end-to-end baselines trained on their native inputs at every budget, with the largest margins when data are scarce and competitive performance at 100%. This pattern holds under cross-dataset evaluation (adult/mixed PhysioNet → pediatric, multi-location CirCor), supporting the claim that our approach is both data-efficient and robust to domain shift. As expected, absolute scores are below in-domain reports on CirCor, reflecting the stricter generalization setting rather than a weakness of the method.

4.5. Multi-Scale Ablation

We quantify the contribution of each temporal scale in our topological representation. Starting from the full three-branch design (global, meso, fine), we remove one branch at a time while keeping the rest of the pipeline unchanged. All ablations use the same decoder (the TCN in [24] with identical capacity and training schedule) and the same order/duration-constrained decoding at inference for fairness. Models are trained on 10% of PhysioNet/CinC 2016 (subject-level subsampling; three random draws) and evaluated on CirCor with four-state macro-F1 (%) under a 60 ms boundary tolerance.

Ablation at 10% training		macro-F1		
Variant	Overall	S1	S2	
Full: global + meso + fine	71.9	79.0	77.6	
w/o global	70.9	78.2	76.9	
w/o meso	70.5	77.1	75.9	
w/o fine	69.8	74.4	73.2	
H_0 -only (no H_1)	68.3	75.1	73.7	
H_1 -only (no H_0)	66.8	73.4	72.1	

Table 4: Multi-scale and homology ablations under the 10% budget.

Analysis. The ablation in Table 4 clarifies why our representation is data-efficient. (1) *Each scale is necessary*: removing any branch consistently degrades macro-F1, confirming that global (slow rhythm), meso (single-cycle structure), and fine (sharp transients) capture complementary cues. (2) *Fine → onsets*: dropping the fine branch produces the largest losses in S1/S2, indicating its role in precise valve-closure localization under noise. (3) *Meso → intra-cycle*: removing the meso branch mainly harms systole/diastole delineation, consistent with its coverage of one-few cardiac cycles. (4) *Global → stability*: the global branch improves long-range temporal consistency and reduces short-state oscillations; its removal yields smaller but systematic declines. (5) *H_0 and H_1 are complementary*: either homology alone underperforms the combined design, showing that both amplitude-driven morphology (H_0) and loop-level patterns (H_1) are informative.

Practicality. For deployment, topological features are computed once per recording and cached; test-time decoding uses the same lightweight TCN and the convex refinement from Sec. 3.2, adding negligible latency. Runtime scales near-linearly with sequence length and graph sparsity, and admits a smooth speed-accuracy trade-off by reducing landscape sampling density.

5. CONCLUSION

We introduced TopSeg, a data-efficient PCG segmentation framework that couples multi-scale topological descriptors with a lightweight TCN and a constraint-aware inference step. In cross-dataset evaluations with device and patient shifts, TopSeg consistently surpasses feature- and architecture-level baselines, especially with limited labels. Ablations trace these gains to complementary global/meso/fine scales and H_0/H_1 , which confer a geometry-aware inductive bias enabling domain generalization and robust transfer—making TopSeg practical and easy to integrate for real-world deployment under label scarcity and device variability.

6. REFERENCES

- [1] Theo Vos, Stephen S Lim, Cristiana Abbafati, Kaja M Abbas, Mohammad Abbasi, Mitra Abbasifard, Mohsen Abbasi-Kangevari, Hedayat Abbastabar, Foad Abd-Allah, Ahmed Abdelalim, et al., “Global burden of 369 diseases and injuries in 204 countries and territories, 1990–2019: a systematic analysis for the global burden of disease study 2019,” *The lancet*, vol. 396, no. 10258, pp. 1204–1222, 2020.
- [2] Maria Rosa Montinari and Sergio Minelli, “The first 200 years of cardiac auscultation and future perspectives,” *Journal of multidisciplinary healthcare*, pp. 183–189, 2019.
- [3] Rangraj M Rangayyan and Richard J Lehner, “Phonocardiogram signal analysis: a review,” *Critical reviews in biomedical engineering*, vol. 15, no. 3, pp. 211–236, 1987.
- [4] Yi He, Wuyou Li, Wangqi Zhang, Sheng Zhang, Xitian Pi, and Hongying Liu, “Research on segmentation and classification of heart sound signals based on deep learning,” *Applied Sciences*, vol. 11, no. 2, pp. 651, 2021.
- [5] Tharindu Fernando, Houman Ghaemmaghami, Simon Denman, Sridha Sridharan, Nayyar Hussain, and Clinton Fookes, “Heart sound segmentation using bidirectional lstms with attention,” *IEEE journal of biomedical and health informatics*, vol. 24, no. 6, pp. 1601–1609, 2019.
- [6] John Gelpud, Silvia Castillo, Mario Jojoa, Begonya Garcia-Zapirain, Wilson Achicanoy, and David Rodrigo, “Deep learning for heart sounds classification using scalograms and automatic segmentation of pcg signals,” in *International Work-Conference on Artificial Neural Networks*. Springer, 2021, pp. 583–596.
- [7] Yao Chen, Yanan Sun, Jiancheng Lv, Bijue Jia, and Xiaoming Huang, “End-to-end heart sound segmentation using deep convolutional recurrent network,” *Complex & Intelligent Systems*, vol. 7, pp. 2103–2117, 2021.
- [8] George Zhou, Yunchan Chen, and Candace Chien, “On the analysis of data augmentation methods for spectral imaged based heart sound classification using convolutional neural networks,” *BMC medical informatics and decision making*, vol. 22, no. 1, pp. 226, 2022.
- [9] Hiam Alquran, Yazan Al-Issa, Mohammed Alsalatie, and Shefa Tawalbeh, “Deep learning models for segmenting phonocardiogram signals: a comparative study,” *PloS one*, vol. 20, no. 4, pp. e0320297, 2025.
- [10] Amit Krishna Dwivedi, Syed Anas Imtiaz, and Esther Rodriguez-Villegas, “Algorithms for automatic analysis and classification of heart sounds—a systematic review,” *IEEE Access*, vol. 7, pp. 8316–8345, 2018.
- [11] David B Springer, Lionel Tarassenko, and Gari D Clifford, “Logistic regression-hsmm-based heart sound segmentation,” *IEEE transactions on biomedical engineering*, vol. 63, no. 4, pp. 822–832, 2015.
- [12] RaviShankar Prasad, Gurkan Yilmaz, Olivier Chetelat, and Mathew Magimai Doss, “Detection of s1 and s2 locations in phonocardiogram signals using zero frequency filter,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 1254–1258.
- [13] Xingchen Xu, Xingguang Geng, Zhixing Gao, Hao Yang, Zhiwei Dai, and Haiying Zhang, “Optimal heart sound segmentation algorithm based on k-mean clustering and wavelet transform,” *Applied Sciences*, vol. 13, no. 2, pp. 1170, 2023.
- [14] Youness Arjoune, Trong N Nguyen, Robin W Doroshov, and Raj Shekhar, “A noise-robust heart sound segmentation algorithm based on shannon energy,” *IEEE Access*, vol. 12, pp. 7747–7761, 2024.
- [15] Gunnar Carlsson, “Topology and data,” *Bulletin of the American Mathematical Society*, vol. 46, no. 2, pp. 255–308, 2009.
- [16] Samuel Leventhal, Attila Gyulassy, Mark Heimann, and Valerio Pascucci, “Exploring classification of topological priors with machine learning for feature extraction,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 7, pp. 3959–3972, 2023.
- [17] Yara Skaf and Reinhard Laubenbacher, “Topological data analysis in biomedicine: A review,” *Journal of Biomedical Informatics*, vol. 130, pp. 104082, 2022.
- [18] Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot, *The structure and stability of persistence modules*, vol. 10, Springer, 2016.
- [19] Peter Bubenik et al., “Statistical topological data analysis using persistence landscapes,” *J. Mach. Learn. Res.*, vol. 16, no. 1, pp. 77–102, 2015.
- [20] Chengyu Liu, David Springer, Qiao Li, Benjamin Moody, Ricardo Abad Juan, Francisco J Chorro, Francisco Castells, José Millet Roig, Ikaro Silva, Alistair EW Johnson, et al., “An open access database for the evaluation of heart sound algorithms,” *Physiological measurement*, vol. 37, no. 12, pp. 2181, 2016.
- [21] Jorge Oliveira, Francesco Renna, Paulo Costa, Marcelo Nogueira, Ana Cristina Oliveira, Andoni Elola, Carlos Ferreira, Alipio Jorge, Ali Bahrami Rad, Matthew Reyna, et al., “The circor digiscope phonocardiogram dataset,” in *IEEE Conference*, 2022.
- [22] Changhyun Park, Keewon Shin, Jinew Seo, Hyunseok Lim, Gyeong Hoon Kim, Woo-Young Seo, Sung-Hoon Kim, and Namkug Kim, “Enhancement of phonocardiogram segmentation using convolutional neural networks with fourier transform module,” *Biomedical Engineering Letters*, vol. 15, no. 2, pp. 401–413, 2025.
- [23] SM Debbal and Fethi Bereksi-Reguig, “Computerized heart sounds analysis,” *Computers in biology and medicine*, vol. 38, no. 2, pp. 263–280, 2008.
- [24] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager, “Temporal convolutional networks for action segmentation and detection,” in *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 156–165.