

Pulmonary Hypertension Detection from Heart Sound Analysis

Original

Pulmonary Hypertension Detection from Heart Sound Analysis / Gaudio, A., Giordano, N., Elhilali, M., Schmidt, S., Renna, F.. - In: IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING. - ISSN 0018-9294. - PP:(2025), pp. 1-13. [10.1109/tbme.2025.3555549]

Availability:

This version is available at: 11583/2999516 since: 2025-04-24T14:33:33Z

Publisher:

IEEE

Published

DOI:10.1109/tbme.2025.3555549

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Pulmonary Hypertension Detection from Heart Sound Analysis

Alex Gaudio, Noemi Giordano, Mounya Elhilali, Samuel Schmidt, and Francesco Renna

Abstract—The detection of Pulmonary Hypertension (PH) from the computer analysis of digitized heart sounds is a low-cost and non-invasive solution for early PH detection and screening. We present an extensive cross-domain evaluation methodology with varying animals (humans and porcine animals) and varying auscultation technologies (phonocardiography and seismocardiography) evaluated across four methods. We introduce PH-ELM, a resource-efficient PH detection model based on the extreme learning machine that is smaller ($300\times$ fewer parameters), energy efficient ($532\times$ fewer watts of power), faster ($36\times$ faster to train, $44\times$ faster at inference), and more accurate on out-of-distribution testing (improves median accuracy by 0.09 area under the ROC curve (auROC)) in comparison to a previously best performing deep network. We make four observations from our analysis: (a) digital auscultation is a promising technology for the detection of pulmonary hypertension; (b) seismocardiography (SCG) signals and phonocardiography (PCG) signals are interchangeable to train PH detectors; (c) porcine heart sounds in the training data can be used to evaluate PH from human heart sounds (the PH-ELM model preserves 88 to 95% of the best in-distribution baseline performance); (d) predictive performance of PH detection can be mostly preserved with as few as 10 heartbeats and capturing up to approximately 200 heartbeats per subject can improve performance.

Index Terms—Auscultation, Biomedical Signal Analysis, Pulmonary Hypertension, Machine Learning

I. INTRODUCTION

PULMONARY Hypertension (PH) is a hemodynamic condition characterized by an increased pulmonary artery pressure (PAP) and an increased afterload in the right ventricle [1]. PH affects both the cardiac and the pulmonary functionality, it can co-occur with several cardiovascular and respiratory diseases [2], including heart failure, and it is

We acknowledge support from the National Institutes of Health under Grant NIH 1R01HL163439. The work of Francesco Renna is co-financed by Component 5 - Capitalization and Business Innovation, integrated in the Resilience Dimension of the Recovery and Resilience Plan within the scope of the Recovery and Resilience Mechanism (MRR) of the European Union (EU), framed in the Next Generation EU, for the period 2021 - 2026, within project HfPT, with reference 41.

Alex Gaudio and Mounya Elhilali are with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218 USA (e-mail: agaudio2@jh.edu and mounya@jhu.edu).

Noemi Giordano is with the Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Torino, Italy (e-mail: noemi.giordano@polito.it).

Samuel Schmidt is with the Faculty of Medicine, University of Aalborg, Denmark (e-mail: sschmidt@hst.aau.dk).

Francesco Renna is with INESC TEC, Faculdade de Ciências da Universidade do Porto, Porto, Portugal (e-mail: francesco.renna@fc.up.pt).

associated with increased mortality [3]. Its prevalence was estimated to affect 1% of the global population [2], [4].

Two main challenges in PH screening include: (a) the need for reliable, low-cost, and non-invasive technology; and (b) the non-specific symptomatic presentation of the condition [3], [5]. While delayed PH diagnosis has been linked with decreased survival rate [6], [7], guidelines for PH diagnosis are designed to balance the benefits of early detection with the economic healthcare burden that early screening places on PH referral centers [5]. The diagnosis of PH, according to recent guidelines from the European Society of Cardiology (ESC) and European Respiratory Society (ERS) [2], is obtained via the right heart catheterization (RHC). The RHC is an invasive and expensive procedure to measure the systolic, diastolic, and mean pulmonary artery pressure (PAP), and it is therefore not suitable for screening.

The clinical need for reliable and non-invasive technologies to raise early suspicion of PH [2] is not well met by current technological capabilities. A reliable, low-cost, and non-invasive approach for raising suspicion of PH can reduce risks and costs, and improve awareness of PH in telemedicine, screening settings, developing countries, rural areas, and underprivileged settings. Existing technologies are used together to raise suspicion of PH and justify performing a final diagnosis by RHC [8]. These technologies include electrocardiography (EKG), blood tests, echocardiography (ECHO), chest radiology, magnetic resonance imaging (MRI), and pulmonary function tests such as spirometry and arterial blood gas [2], [5]. The 2022 ESC/ERS guidelines recommend ECHO as a main tool for the detection of PH, to be performed after a thorough physical exam including blood tests and resting EKG, and before final confirmation with RHC. The ECHO technology, however, has important drawbacks: (a) no single ECHO biomarker reliably informs about PH status [2] and it cannot estimate pulmonary pressure in 10-50% of patients [9]; (b) a normal ECHO does not exclude PH [5]; and (c) the ECHO procedure is not low resource cost because it requires a trained sonographer, an analysis by a cardiologist and typically also a referral to an imaging center or hospital. We propose an under-recognized but promising technology for raising suspicion of PH: digital auscultation and the computer analysis of heart sounds.

Contributions: This work contributes a PH-ELM predictive model for PH detection from heart sounds. The model has a very low resource footprint and higher prediction performance in out-of-distribution (OOD) tests when compared to competing heart sound analysis approaches. We also contribute

an extensive methodology and analysis for the in-distribution and out-of-distribution testing of three heart sound datasets and four predictive models. We study PH detection in different and changing data modalities, including heart sounds recorded with seismocardiography (SCG) and phonocardiography (PCG) technologies, and heart sounds from humans and pigs. For the science of heart sound auscultation in pulmonary hypertension, this work demonstrates that heart sounds auscultated via digital stethoscopes or via accelerometers can contain sufficient information for accurate PH detection. Moreover, we provide evidence that training datasets for PH detection can mix PCG and SCG signals or use either interchangeably, and that PH detectors can be trained on porcine heart sounds and evaluated on human heart sounds.

Section II reviews background material and surveys related works to contextualize and justify the proposed methodology. Section III describes the methods and datasets. Section IV shows results supporting our methodology and scientific claims. Section V offers insight into the scope of this work and the merits of the proposed modeling and evaluation methodologies.

II. BACKGROUND AND RELATED WORKS

Background: Within each heartbeat, the closing of the heart's four valves creates two main sounds colloquially referred to as the "lub" and the "dub" sounds. The "lub" sound, known as the first heart sound, or S1, arises from the closures of the two atrioventricular valves [17]. The second heart sound (S2) [18], or the "dub" sound, occurs as a result of the closure of the aortic valve (A2) and pulmonic valve (P2). The time delay and relative intensity of the P2 are known to be relevant for pulmonary hypertension detection [19], while the third and fourth heart sounds, known as S3 [20] and S4, are often not audible by a human ear but they can be useful for pulmonary hypertension detection [10], [12].

Related Works: Some of the existing PH detection methods create hand-crafted features that extract information from the S1, S2, S3, and S4 components of a heart sound recording, and then evaluate the features using correlation, significance testing, or linear regression. Yamakawa et al. [10] studied intensity, complexity, and strength features extracted from each of the four fundamental heart sounds to show that S2 complexity and S3 intensity are the most statistically significant features for PH classification. They also obtained auROC values by training and evaluating the predictive performance of single features for different kinds of PH. Huang et al. [11] used reference PAP estimates obtained from ECHO to partition human subjects into three levels of PH (mild, medium, severe) and a control group. They found that: S2 amplitude and frequency alone were predictive of PH; ratios of the S2/S1 energy, amplitude, and frequency were all higher in subjects with PH; and the amplitude, frequency, and energy of S1 sounds are not correlated with ECHO readings. In 2010, Dennis et al. [12] utilized a variety of features on all four fundamental heart sounds and the authors report good results while investing minimal effort on manual pre-processing. Kaddoura et al. [13] applied a Gaussian Mixture

Model to features extracted from the Mel-frequency Cepstral Coefficients (MFCC) of S2 segments. While their experiment was shown to outperform physician auscultation by a large margin, their relatively low overall performance (0.74 auROC) compared to the other works in Table I is possibly due to the choice of features. MFCC features are designed to approximate the human auditory system's response, but their result with human auscultators also suggests the human ear may not provide a good reference model for PH detection. While hand-crafted features can yield models with low computational resource requirements and useful interpretations, the predictive performance of models relying on these methods is often fundamentally limited by the availability of and ability to encode domain knowledge into the generated features. Our proposed fixed-weight approach bypasses limitations by randomly generating features and benefits from the associated computational improvements.

Data-driven machine learning techniques use a dataset to optimize model parameters, and they create or utilize features that are learned or automatically optimized to the data. Wang et al. [14] tested 10 different Deep Learning models on a set of heart sound features obtained from a magnitude spectrogram of the continuous wavelet transform and augmented with random noise. While the authors report 0.98 accuracy (averaged across multiple classifications including PH) and a similarly high F1 score for PH over a 10-fold cross-validation, they state that the experiment's dataset mixed the same subjects in the training and validation sets. When training data appears in the validation set, the empirical prediction performance is arguably over-optimistic due to overfitting. Ge et al. [16] extracted features from segmented S1 and S2 sounds by utilizing fixed-weight and data-driven methods, including hand-crafted features, time-frequency analysis with wavelets, and a convolutional neural network feature extractor. The study partitions 2415 recordings of 438 subjects into train, validation, and test sets. It was not specified if the subsets were stratified by subject. Our previous work [15], [19], [21] proposed a CNN-based PH detection method to evaluate an image representation of S2 segments. It incorporated a pre-processing method based on blind source separation to extract A2 and P2 waveforms from each S2 segment. The empirical analysis claimed that the availability of separated A2 and P2 components contributes significantly to prediction, but it was only cross-validated with one dataset, and 10-fold cross-validation performed only once.

III. MATERIALS AND METHODS

We propose a methodology for PH detection based on the extreme learning machine architecture that has the high prediction performance of a deep network with a small computational footprint. We present an evaluation approach that considers both in-distribution testing and out-of-distribution testing with two ranking metrics. Section III-A describes the proposed PH detector model. Section III-B describes the evaluation methodology, including an in-distribution analysis based on bootstrapped cross-validation and an out-of-distribution analysis across datasets. Section III-C describes the datasets

TABLE I: PH detection methods

Reference	Year	Data Analyzed	Ground Truth	Approach	Evaluation	Performance	Num Subjects
Yamakawa et al. [10]	2022	S1, S2, S3, S4	RHC	ANOVA	Not Specified	0.67 - 0.81 auROC	40
Huang et al. [11]	2023	S1, S2, S3, S4	ECHO	LR	Not Specified	0.78/ 0.83/ 0.88 auROC	209
Dennis et al. [12]	2010	S1, S2, S3, S4	ECHO	NB	Train:Test	0.78 auROC	20:31
Kaddoura et al. [13]	2016	S2	RHC	MM	5-fold CV	0.74 auROC	164
Wang et al. [14]	2022	S1, S2, S3, S4	Unspecified	DNN	10-fold CV	0.99 auROC	74
Gaudio et al. [15]	2022	S2	RHC	DNN	10-fold CV	0.95 auROC	42
Ge et al. [16]	2023	S1, S2	ECHO	XGBoost	Train:Val:Test	0.93 auROC	483

Approach: LR means Logistic Regression; NB is Naive Bayes; DNN is Deep Neural Network; MM is Mixture Model.

Evaluation: CV means cross-validation; Train:Test is a 1-fold CV with the validation set defined as the test set. No results in the table validated results with an out-of-distribution dataset for the detection of PH.

Performance: auROC means area under the receiver operating characteristic (ROC) curve.

195 used. Section III-D introduces the competing methods and
196 hyperparameters.

197 **A. Proposed Model based on Extreme Learning**
198 **Machine**

199 We restrict our focus to machine learning methods based on
200 deep learning and the extreme learning machine. The model
201 design, shown in Figure 1, can be described as a resource-
202 efficient extreme learning machine composed of three stages:
203 feature generation, compression, and classification.

204 **Input:** The input is a pre-processed image matrix containing
205 segmented S2 sounds, and it is visualized in Figure 1. The
206 image is obtained via the following analysis of a heart sound
207 audio recording: a 200 ms time window vector of the recorded
208 S2 heart sound is extracted by a dataset-specific segmentation,
209 and then the ordered set of time window vectors are stacked
210 as rows of a matrix. The output of the pre-processing step
211 for each sample is a matrix with one row for each heartbeat
212 and 200 columns that represent an aligned 200 ms window
213 of the S2 sound. We assume each image is a single-channel
214 grayscale image matrix, defined as $C = 1$ channel. Note that
215 our prior works [15], [19] utilized this image structure and
216 also introduced a $C = 3$ input image via a pre-processing
217 step. The shape of the input is (B, C, H, W) where $B = 1$
218 is the batch size¹, $C \geq 1$ channels, $W = 200$ audio samples,
219 and H heartbeats is variable.

220 **Feature Generation:** The feature generation step can be
221 interpreted as a kind of time-frequency transform with a fixed
222 set of random kernels and random fixed bias. A parallel
223 block convolution layer converts the shape (B, C, H, W) into
224 a shape $(B, O, 1, 1)$ using a series of convolution layers,
225 the ReLU nonlinear operator, and adaptive average pooling.
226 The convolution layer evaluates CO different 5×5 kernel
227 matrices with a stride of 1 and dilation of 2 to attain a shape
228 (B, O, H', W') , and it also uses a randomly initialized bias.
229 The bias and kernel parameters are randomly initialized from
230 uniform distributions with zero mean, and they are never mod-
231 ified after initialization. ReLU converts any negative values
232 to zero. Adaptive Average Pooling converts (B, O, H', W')

to shape $(B, O, 1, 1)$. Because adaptive pooling converts the
convolution feature matrix to a scalar number, it creates a
feature generator with fixed output that accepts variable size
input.

In the convolution layer, the kernel size and dilation can
be matched to the sample rate of the represented signal of
length W . At a sample rate of 1000 Hz, a kernel size of $k =$
 $(5, 5)$ and dilation $d = (2, 2)$ creates a filter with an effective
receptive field of $k + (k - 1)d = (13, 13)$ pixels, which defines
a 13 ms window over 13 consecutive heartbeats.

Adjustable RAM utilization: The parallel group convolution
also introduces a set of G groups, each of which computes
 $\frac{O}{G}$ outputs. Figure 1 shows a set of O white squares that are
divided into G groups. Each group of white squares is a single
convolution layer with $\frac{O}{G}$ outputs. In this work, we set $O =$
1000 and $G = 10$, therefore each group (of white squares) is
responsible for converting the input (B, C, H, W) into a subset
of the outputs $(B, O/G, H', W')$. Each group is computed in
series, and by tuning G , the total amount of utilized memory
is restricted; increasing G decreases memory utilization and
increases compute time.

Compression and Classification: The generated features
are compressed by a principal components analysis and clas-
sified via ridge regression. The training phase and inference
phase are subsequently described.

Training: The objective of the training phase is to solve
a ridge regression on a compressed feature representation. In
Equation (1), the feature representation of each input training
dataset image becomes a row of the matrix E_{CNN} , where
 $f(I_i)$ is the feature generation of the given input image. The
columns of E_{CNN} are reduced via SVD. We assume that
zero centering by ensuring each column has zero mean does
not affect downstream prediction performance because the
bias and kernel parameters have zero mean in expectation².
The first o columns of the projection matrix V are used for
compression and stored for inference on the validation set. The
value o is a hyperparameter that determines how many bases

¹Note that $B > 1$ implies the number of heartbeats is constant across images, or that all images are zero-padded or cropped to the same number of heartbeats. Zero-padding may introduce artifacts during feature generation due to the presence of a bias term in the feature generation convolution layers, and it is unknown how or if these artifacts impact performance.

²It as also found in [22] that zero centering paired with random CNN initialization did not meaningfully affect performance. However, our final implementation used zero centering.

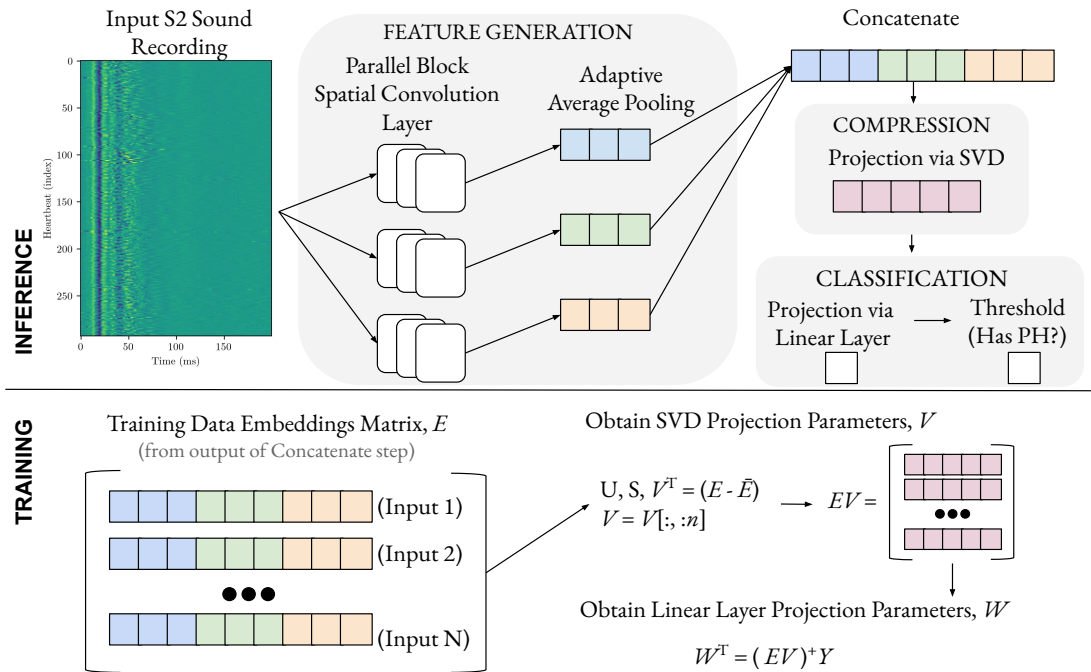


Fig. 1: PH-ELM, an extreme learning machine model architecture with small computational footprint designed with three parts: a fixed-weight, random and non-adaptive feature generation step, a compression step, and a classification utilizing ridge regression. Parallelized block convolutions, adaptive pooling, and singular value decomposition (SVD) convert each input heart sound recording into a compressed row vector. Training optimizes a SVD projection matrix and the weights for a final linear layer by direct analysis of a partially compressed representation of the training dataset.

270 to use for compression, and we use $o = 20$.

$$E_{\text{CNN}} = \begin{bmatrix} \mathbf{f}(I_1) \\ \vdots \\ \mathbf{f}(I_N) \end{bmatrix} \in \mathbb{R}^{N \times o} \implies \begin{aligned} U, S, V^T &= E_{\text{CNN}} - \bar{E}_{\text{CNN}} \\ V &= V[:, :o] \\ E &= E_{\text{CNN}} V \in \mathbb{R}^{N \times o} \end{aligned} \quad (1)$$

271 A ridge regression model generates a prediction of pulmonary
 272 hypertension with an objective defined in Equation (2). The
 273 objective, defined in Equation (2), does not include a bias term
 274 as a column of E because it is assumed that the bias is zero
 275 in expectation, as a result of the feature generation and SVD.
 276 The Lagrange multiplier $\lambda = 0.1$ is chosen to avoid degenerate
 277 cases that could occur if a value on the diagonal of the square
 278 matrix $E^T E$ equals zero, and $\mathbf{1}$ is an identity matrix.

$$\arg \min_{\mathbf{w}} \|\mathbf{y} - E\mathbf{w}\|_2^2 + \lambda \|\mathbf{w}\|_2^2 \quad (2)$$

$$\implies \mathbf{w} = (E^T E + \lambda \mathbf{1})^{-1} E^T \mathbf{y} \quad (3)$$

279 **Binarization:** In the ground truth, if $y \in [0, 1]$, then a
 280 threshold on the predicted value $\hat{y} > 0.5$ is a reasonable
 281 classification boundary. If $y \in [-1, 1]$, then $\hat{y} > 0.0$ is a
 282 reasonable classification boundary. The threshold could also be
 283 found by standard analysis in ROC or PR space of a validation
 284 set (not a test set).

285 Training therefore introduces hyperparameter o , the number
 286 of columns of the compressed output, and stores learned
 287 parameters, $\mathbf{w} \in \mathbb{R}^{(o,1)}$, $\bar{E} \in \mathbb{R}^{(1,o)}$ and $V \in \mathbb{R}^{(W,o)}$.

288 **Output Classification:** The scalar value $\hat{y} = (f(I) - \bar{E})V\mathbf{w}$
 289 determines an unnormalized score for PH detection that can
 290 be thresholded.

B. Evaluation

291 Figure 2 visually describes the evaluation methodology em-
 292 ployed. This section first introduces the two rank metrics uti-
 293 lized for evaluation, then presents the in-distribution analysis
 294 methodology, and last presents the out-of-distribution analysis
 295 methodology and shows how we combine in-distribution with
 296 out-of-distribution analysis.
 297

298 **Metrics for Empirical Evaluation:** We adopt the area
 299 under the ROC curve (auROC) and a variation of the area
 300 under the PR curve that we call AP*, described below and
 301 formally introduced in our concurrent work [23].

302 The auROC shows a change in score when the classifier
 303 is evaluated on datasets of varying imbalance ratios [24].
 304 Moreover, the auROC is insensitive to model prediction bal-
 305 ance, which means it penalizes all errors equally, regardless
 306 of whether the model outputs large prediction probabilities or
 307 small ones [25]. We adopt the standard implementation from
 308 the SciKit Learn library [26].

309 The AP* is designed to be sensitive to both class imbal-
 310 ance and model prediction imbalance because it penalizes
 311 prediction errors to the minority class more than errors to the
 312 majority. Given two models with identical auROC, the AP*
 313 gives a higher number to the model that is better at classifying
 314 the majority class, and it can therefore be useful in model
 315 selection. AP* is a variation of, and improvement upon, the
 316 average precision (AP). While the AP is non-symmetric in
 317 the presence of the class imbalance, the AP* and the auROC
 318 are symmetric. The AP* ranges from $[0,1]$, where the random
 319 classifier, all ones classifier, and all zeros classifier all have the

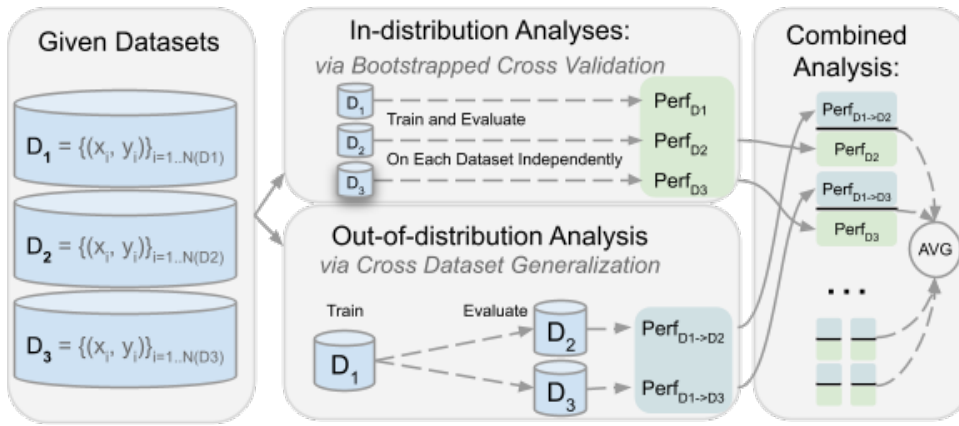


Fig. 2: Evaluation Methodology with in-distribution and out-of-distribution testing.

320 minimum value of $AP^* = \frac{\text{num samples in minority class}}{\text{dataset size}} \in [0, 0.5]$. It
 321 is defined as:

$$\mathbf{p}^+, \mathbf{r}^+, \mathbf{t} = \text{PR_curve}(\mathbf{y}, \hat{\mathbf{y}}) \quad (4)$$

$$\mathbf{p}^-, \mathbf{r}^-, \mathbf{t} = \text{PR_curve}(1 - \mathbf{y}, 1 - \hat{\mathbf{y}}) \quad (5)$$

$$\mathbf{p}^*, \mathbf{r}^* = \text{emin}(\mathbf{p}^+, \mathbf{p}^-), \text{emin}(\mathbf{r}^+, \mathbf{r}^-) \quad (6)$$

$$AP^*(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i \in \{1, \dots, |T|\}} (r_i^* - r_{i-1}^*) p_i^* \quad (7)$$

322 where $\mathbf{y} \in \{0, 1\}^n$ is a vector of the n ground truth PH
 323 annotation labels and $\hat{\mathbf{y}} \in \mathbb{R}^n$ is a corresponding vector of
 324 the model's predictions. The function $\text{emin}(\cdot, \cdot)$ computes the
 325 element-wise minimum between two vectors. The vectors \mathbf{p}^+ ,
 326 \mathbf{r}^+ , \mathbf{p}^- , \mathbf{r}^- contain the precisions and recalls computed at
 327 each of the thresholds (in vector \mathbf{t}) emitted by a precision-
 328 recall curve function on the defined input. If Equation (7)
 329 used \mathbf{p}^+ and \mathbf{r}^+ instead of \mathbf{p}^* and \mathbf{r}^* , it would be identical
 330 to the average precision score presented in the SciKit Learn
 331 library [26] and [27]. Whereas the average precision score
 332 assigns a larger penalty when the prediction probability is
 333 large, the AP^* assigns a larger penalty to the minority class.
 334 A detailed introduction to the AP^* is described in concurrent
 335 work [23].

In-distribution Analysis Methodology:

336 *Bootstrapped Cross-Validation* describes model perfor-
 337 mance using one dataset. On small datasets, such as all datasets
 338 used in this paper, cross-validation analysis (CV) alone may be
 339 insufficient because the way the folds are generated may have a
 340 significant impact on model performance. Bootstrapped cross-
 341 validation improves the reliability of the reported performance.
 342 By bootstrapping, we mean to compute K -fold CV multiple
 343 times independently, where each time, the folds are randomly
 344 generated using a different random seed. In the special cases
 345 of cross-validation with no variation in the folds, such as group
 346 cross-validation where each group is fixed, bootstrapping can
 347 capture possible random variations in model initialization or
 348 training. The following notation is used: in a given dataset D ,
 349 we define each k^{th} fold, $k \in \{1, \dots, K\}$ and each b^{th} bootstrap
 350 iteration, $b \in \{1, \dots, B\}$. For each fold and bootstrap
 351 iteration, we have two disjoint sets: the training set $D_{b,k,\text{train}} \subset D$
 352 and validation set $D_{b,k,\text{val}} \subset D$. After training the model

354 on $D_{b,k,\text{train}}$, the evaluation on $D_{b,k,\text{val}}$ yields a ground truth
 355 vector $\mathbf{y}_{k,b}$ and model predictions $\hat{\mathbf{y}}_{k,b}$. These vectors will be
 356 aggregated and evaluated with a scoring function, $s(\mathbf{y}, \hat{\mathbf{y}})$. The
 357 two scoring functions considered are auROC and AP^* .

358 The empirical in-distribution performance, $P_{\text{in-dist}, s(\cdot), f}(D)$,
 359 of the predictive model f on the given dataset D using scoring
 360 function $s(\cdot)$ is:

$$\mathbf{y}_{\text{micro-avg}, b} = \text{concat}(\mathbf{y}_{k,b} \forall k) \quad (8)$$

$$\hat{\mathbf{y}}_{\text{micro-avg}, b} = \text{concat}(\hat{\mathbf{y}}_{k,b} \forall k) \quad (9)$$

$$P_{\text{in-dist}, s(\cdot), f}(D) = \frac{1}{B} \sum_b s(\mathbf{y}_{\text{micro-avg}, b}, \hat{\mathbf{y}}_{\text{micro-avg}, b}). \quad (10)$$

Out-of-distribution Analysis Methodology:

361 To evaluate how well a model generalizes across datasets,
 362 we define a training dataset D_{train} and a test dataset D_{test} .
 363 Training and evaluation are repeated B times independently.
 364 Re-using notation from the previous section (and redefining a
 365 new number of bootstrap iterations B and $b \in \{1, \dots, B\}$),
 366 this bootstrapped out-of-distribution evaluation on the test set
 367 gives scalar numbers of the form $s(\mathbf{y}_{\text{test}, b}, \hat{\mathbf{y}}_{\text{test}, b})$ where $\mathbf{y}_{\text{test}, b}$
 368 contains the ground truth test dataset annotations, and $\hat{\mathbf{y}}_{\text{test}, b}$
 369 contains the trained model's test set predictions.
 370

371 We compare the model's performance to the performance of
 372 a separately trained in-distribution baseline model by comput-
 373 ing a ratio. This ratio is averaged over B bootstrap iterations:

$$P_{\text{out-of-dist}, s(\cdot), D_{\text{train}}, f}(D_{\text{test}}) = \frac{1}{B} \sum_b s(\mathbf{y}_{\text{test}, b}, \hat{\mathbf{y}}_{\text{test}, b}) \quad (11)$$

$$P_{\text{ood-ratio}, D_{\text{train}}, D_{\text{test}}, f} = \frac{\text{out-of-distribution score}}{\text{in-distribution score}} \quad (12)$$

$$= \frac{P_{\text{out-of-dist}, s(\cdot), D_{\text{train}}, f}(D_{\text{test}})}{P_{\text{in-dist}, s(\cdot), f}(D_{\text{test}})}, \quad (13)$$

374 where we have a given scoring function $s(\cdot)$ implementing the
 375 auROC or AP^* , a training dataset D_{train} , an evaluation dataset
 376 D_{test} , a predictive model f , and an in-distribution evaluation
 377 function $P_{\text{in-dist}, s(\cdot), f}$ from Equation (10).

378 Note that the denominator of Equation (11) is effectively
 379 a baseline against which the out-of-distribution score is com-
 380 pared. The baseline can be chosen arbitrarily, or in a way
 381 that aids interpretation. For instance, for model selection, the

TABLE II: Dataset Summary

Dataset	Domain	Mode	Subjects	Samples	PH Samples
H-PCG [15]	Human	PCG	42	42	29 (69%)
P-Both [28]	Porcine	PCG	10	32	16 (50%)
	Porcine	SCG	10	93	42 (45%)
H-SCG [29]	Human	SCG	73	82	63 (76%)

denominator can show how much performance is lost when compared to the best model. More specifically, given a set of models \mathcal{G} , where each model $g \in \mathcal{G}$ is trained on D_{train} and evaluated on D_{test} , we define the denominator as the maximum in-distribution score of the test dataset across all models. This variation of the score is used in our experiments, and is shown below:

$$P_{\text{ood-ratio}, D_{\text{train}}, f, G} = \frac{P_{\text{out-of-dist}, s(\cdot), D_{\text{train}}, f}(D_{\text{test}})}{\max_{g \in \mathcal{G}} P_{\text{in-dist}, s(\cdot), g}(D_{\text{test}})} \quad (14)$$

Given the predictive model f and scoring function, Equation (14) gives one measurement for each pair of datasets. To give a general sense of performance comparable across all datasets, we propose to compute the median over the pairs of datasets in order to have a single number for each model and scoring function. We define a set of pairs of train and test datasets as \mathcal{D} , and compute the median score, where $m(\cdot)$ is the median function:

$$P_{\text{ood-avg}, s(\cdot), f, G} = m(\{P_{\text{ood-ratio}, D_{\text{train}}, f, G} \mid (D_{\text{train}}, D_{\text{test}}) \in \mathcal{D}\}). \quad (15)$$

C. Datasets

We evaluate three different datasets. The datasets vary in recording modality, animal species and hospital location. The modalities considered include phonocardiography (PCG), and seismocardiography (SCG); PCG uses microphones while SCG uses accelerometers. Two datasets contain data from humans and one dataset is from pigs. All datasets are from different hospitals. All three datasets record subjects undergoing right heart catheterization. The datasets are described in the following paragraphs and summarized in Table II.

1) *Dataset H-PCG*: Human PCG data was acquired from 42 subjects at Centro Hospitalar Universitário do Porto, Portugal. 29 subjects have PH and 13 subjects do not have PH. The dataset was previously introduced with the name HSA [15] and it is not publicly available.

For each subject, ground truth pulmonary artery pressure was obtained from a right heart catheterization (Swan-Ganz catheter), and an accompanying five-minute PCG heart sound recording was collected. The recording was obtained in a quiet clinical setting with the patient supine and at rest. Auscultation was performed over the second left intercostal space using a custom cable stethoscope connected to a Rugloop Waves system. Heart sounds were recorded at a sample rate of 8 kHz and their amplitudes were quantized with 16-bit resolution. We sub-sample the signal to 1 kHz. The 200 ms time window vectors for each recording were obtained by the CNN-based method of Renna et al. [30]. The subjects in the dataset

are labeled as PH positive if have Mean Pulmonary Arterial Pressure (MPAP) above 25 mmHg, or Pulmonary Arterial Systolic Pressure (PASP) above 30 mmHg.

2) *Dataset P-Both*: Porcine PCG data (P-PCG) and SCG data (P-SCG) were acquired at Aalborg University Hospital, Aalborg, Denmark [28]. The pigs were sedated, ventilated, and subject to catheterism both in the aorta and the right ventricle using Swan-Ganz catheters. The SCG acquisition was carried out using an iWorx commercial system equipped with two triaxial accelerometers, located respectively over the fourth intercostal space next to the sternum and over the lower border of the sternum. The PCG acquisition was carried out using a multi-channel wearable system designed at Politecnico di Torino, which embeds 48 electret condenser microphones with a 12-mm spatial resolution [31]. A simultaneous ECG was collected by both instruments. The sampling frequency while recording was set to 5 kHz for SCG and to 1 kHz for PCG, and then we downsampled the SCG signal to 1 kHz using nearest-neighbor interpolation. A PH condition was triggered in the animal by subjecting it to either nitrogen asphyxiation (to cause hypoxemia) or carbon dioxide asphyxiation (to cause hypercapnia). The experiment was reversible, and therefore when the trigger was removed, the animal reached a baseline condition again and could be subjected to multiple experiments. A total of 59 experiments were carried out on ten pigs. The dataset for this study was created by considering, for each recording, two one-minute segments: one at the beginning of the recording, labeled as "no PH", and one at the peak of the effect of the trigger, labeled as "PH". Poor-quality recordings were manually discarded. There were 67 "no PH" and 58 "PH" samples. For the segmentation pre-processing steps, we applied a band-pass of 20 Hz to 200 Hz using a second-order IIR Butterworth filter, then the signals were segmented into heartbeats using the R-wave ECG, and then a 200-millisecond S2 segment was extracted from each heartbeat using time thresholding and peak detection. The full dataset P-Both is a union of the P-SCG and P-PCG datasets, where P-SCG and P-PCG share the same physiological basis but have different modalities. The P-Both dataset is analyzed only as a training dataset in the context of out-of-distribution experiments with human test datasets, and it enables the analysis of whether both modalities from the same set of subjects can improve generalization performance on humans.

The dataset is not publicly available.

3) *Dataset H-SCG*: The dataset was acquired at the catheterization laboratory at the University of California and is publicly available on PhysioNet [29] with the name SCG-RHC. The dataset includes 83 recordings of 72 patients referred for the hemodynamic assessment of their heart failure status. While the dataset includes data from a pharmacological experiment, we use only data from the end of the 10 minute patient resting period that occurred before any pharmacological intervention. ECG and triaxial SCG were simultaneously recorded using a wearable patch located over the sternum. The sampling frequency is 500 Hz. A subject was labeled as PH positive if the PASP was greater than 30 mmHg or if the MPAP was greater than 25 mmHg. The segmentation pre-

482 processing obtained 200 ms windows of each recorded S2 by
 483 the method used for the P-Both dataset, and the S2 segments
 484 were aligned in time by selecting a lag index that maximizes
 485 the cross-correlation between the homomorphic envelopgram of
 486 the segment and the homomorphic envelopgram of the first S2
 487 segment of the recording.

488 D. Hyperparameters and Competing Methods

489 A total of four methods are evaluated: PH-ELM,
 490 DenseNet121, Scattering-SVM, and STFT-SVM. The reason-
 491 ing for the choice of models and description of their construc-
 492 tion is presented below. The DenseNet121 deep network [32]
 493 was chosen for two reasons: a) our previous work found
 494 DenseNet121 to outperform the in-distribution performance of
 495 several other models on the H-PCG dataset [15]; and b) the
 496 PH-ELM model's fixed-weight convolutional feature generator
 497 is similar to a deep convolutional network. The Scattering-
 498 SVM consists of a wavelet scattering feature generator and
 499 support vector machine (SVM). The wavelet scattering com-
 500 ponent was chosen due to the similarity of scattering to the PH-
 501 ELM's feature generation, and the SVM was chosen because
 502 literature comparing SVM to ELM has found that SVM can
 503 outperform ELM models on small datasets [33]. The STFT-
 504 SVM model, previously introduced in [15], utilizes a short-
 505 time Fourier transform for feature generation followed by an
 506 SVM. It was chosen because it is a resource-efficient method
 507 that can perform well.

508 *Hyperparameters for DenseNet121:* The models utilizing
 509 the DenseNet121 architecture were randomly initialized (they
 510 were not pre-trained) and optimized with the AdamW opti-
 511 mizer [34] with learning rate of 1e-5 and weight decay of
 512 either 1e-5 (P-Both) or 1e-6 (H-PCG and H-SCG). Because
 513 the model cannot handle images with less than 61 rows or
 514 columns, the input to DenseNet121 was zero padded to a
 515 uniform size of 454 rows (H-PCG) or 157 rows (P-Both, P-
 516 SCG, and P-PCG). Padding on H-SCG required additional
 517 treatment, as some patients have too few rows, and others have
 518 too many rows. Each input image recording was zero padded
 519 to 400 rows, and recordings with too many rows were cropped
 520 to 400 contiguous rows. When cropping, the crop region was
 521 randomly chosen during training in order to use the entire
 522 dataset, and fixed to the last 400 rows of the recording (but
 523 before the H-SCG chemical experiment began) during testing.

524 *Hyperparameters for Scattering-SVM:* Since the SVM re-
 525 quires a fixed-size input that is the same in both training and
 526 testing, the input to Scattering-SVM was also zero padded. We
 527 adopt the padding approach and random cropping approach
 528 used for DenseNet121, but fix all datasets to 400 rows.
 529 Moreover, the columns of each padded (or cropped) image
 530 were normalized to variance 1 because we found empirically
 531 that this improved performance. The scattering layer used 2
 532 orders, wavelet scale $J = 2$, $L = 8$ angles, the Morlet wavelet
 533 family, and the standard implementation from the Kymatio
 534 library [35]. After scattering, each input was reshaped into a
 535 row vector. For training, all inputs are stacked into a matrix
 536 and each feature column was normalized to zero mean and
 537 unit variance. This matrix was passed to the support vector

538 regression from SciKit Learn [26] initialized with $C = 1.0$,
 539 RBF kernel, and $\gamma = \frac{1}{C(X)\sigma^2(X)}$ (also known as "scale"
 540 in SciKit Learn) where $C(X)$ is the number of training set
 541 columns ($C(X) = 200$ in all datasets) and σ^2 is the variance
 542 of the training dataset.

543 *Hyperparameters for STFT-SVM:* The padding methodology
 544 was identical to Scattering-SVM³. The STFT was performed
 545 by computing a spectrogram from the PyTorch Audio li-
 546 brary [36] for each row (each heartbeat), using FFT of size
 547 64 (giving 33 frequency bins), hop length 2, and power 1 for
 548 the magnitude spectrum. The spectra from all heartbeats were
 549 aggregated by computing the 98% and 100% quantiles across
 550 rows, reducing the size of each subject's recording from (H,
 551 200) to (33, 101), where rows are frequency bins, and columns
 552 are time. The columns of this matrix were normalized to unit
 553 variance and then flattened into a row vector. The SVM (from
 554 SciKit Learn) used $C = 1$, RBF kernel, and $\gamma = \frac{1}{C(X)}$ (also
 555 known as "auto" in SciKit learn). This model was previously
 556 introduced in [15].

557 *Hyperparameters for PH-ELM:* The PH-ELM's convolu-
 558 tional feature generator used $O = 1000$ output channels, kernel
 559 size 5×5 , dilation 2×2 , $G = 10$ groups, and 20 principal
 560 components. The linear regression used $\lambda = 0.1$. Padding was
 561 not necessary for this model because the feature generator
 562 standardizes all inputs to the same size.

563 IV. RESULTS

564 We evaluate the four PH detection models and three datasets
 565 described in Section III-C. Section IV-A and Section IV-B
 566 respectively present the observed in-distribution and out-of-
 567 distribution prediction performances, Section IV-C presents the
 568 computational performance, and Section IV-D examines the
 569 effect of varying heartbeats per subject.

570 A. In-distribution Prediction Performance

571 Table III is useful primarily to give a sense of the overall
 572 best prediction performance a model might expect to achieve
 573 on a given dataset. The calculation for auROC and AP* scores
 574 is defined by Equation (10). The DenseNet121, followed by
 575 PH-ELM, tends to have the highest prediction performances
 576 in both AP* and auROC across the three datasets. The PH-
 577 ELM model has the lowest standard deviation of all models
 578 on the H-PCG and H-SCG. Moreover, we can observe that
 579 the datasets can be ranked based on performance in this
 580 order: H-PCG (highest), P-Both, and H-SCG (lowest). Possible
 581 reasons for H-SCG's low performance might be related to
 582 the placement of the SCG sensor on the sternum, or to the
 583 dataset's 500 Hz sampling rate. The other datasets utilized the
 584 second and fifth intercostal spaces, respectively, and they were
 585 sampled at 1000 Hz. We do not make comparisons of PCG
 586 and SCG signals for PH detection based on this in-distribution
 587 analysis.

³The computational footprint tests in Table VI do not use padding, as padding is not necessary for this model and would increase the resource footprint.

B. Out-of-distribution Prediction Performance

Out-of-distribution testing shows the prediction performance when training a model on one dataset and evaluating it on a different dataset.

PH detection with SCG and PCG datasets: Table IV shows PH detection performance when training with PCG signals to evaluate SCG signals, and vice versa, in either humans (comparing H-SCG to H-PCG) or pigs (comparing P-SCG to P-PCG). While the human datasets are independent, the porcine datasets are generated from the same ten pigs. To ensure the train and test sets do not share the same physiological basis, we utilized a leave-one-pig-out cross-validation, with an additional constraint that the PCG and SCG data are disjoint across train and test sets (i.e., either $D_{\text{train}} \subset \text{P-PCG}$ and $D_{\text{test}} \subset \text{P-SCG}$, or $D_{\text{train}} \subset \text{P-SCG}$ and $D_{\text{test}} \subset \text{P-PCG}$), and the reported results are obtained from the *in-distribution* equation Equation (10). We identify the porcine experiments as out-of-distribution because we analyze generalization across modalities (PCG and SCG), but the non-independent nature of the datasets requires cross-validation and the in-distribution evaluation equation.

Regarding model selection, the PH-ELM model has the highest scores (in bold) for all pairs of datasets except when $(D_{\text{train}}, D_{\text{test}})$ is (H-SCG, H-PCG), in which case the model ranks second to Scattering-SVM (with auROC of 0.55 versus 0.83). DenseNet121 and STFT-SVM both become no better than random from H-SCG to H-PCG (auROC of $0.5 \pm .01$). Moreover, of the three instances in which a model gives *higher* out-of-distribution performance than in-distribution performance, two of these occur with the PH-ELM, and these are shown by an AP* Ratio or auROC Ratio greater than one. The results suggest that the PH-ELM generalizes well across PCG and SCG signals with little loss in prediction performance. In summary, the PH-ELM shows the best out-of-distribution generalization performance overall.

The table also offers some insight into dataset selection. Since the H-SCG dataset has the overall lowest in-distribution scores in Table III, it may not be a high-quality training dataset for PH detection, since three of the four models (DenseNet121, PH-ELM, and STFT-SVM) give performances that are effectively no better than random when training H-SCG to evaluate H-PCG. The fact that Scattering-SVM gave good performance on H-SCG, though, means the dataset does retain information about PH. On the other hand, P-SCG \rightarrow P-PCG shows comparably good generalization with three of the four models. Both results can be interpreted to suggest that SCG signals can provide a useful training dataset for PH detection in PCG signals. Finally, the evaluation PCG \rightarrow SCG, using either human or porcine datasets, gives useful models. In summary, the results show that both PCG and SCG data can be useful as training sets.

PH detection from pigs to humans: Table V shows that porcine heart sounds can provide useful training data to predict PH in human heart sounds. The performance from humans to pigs was excluded from the table because we assume that the generalization from humans to pigs is not relevant. The generalization to pigs has auROC values less than 0.6 in nearly

all cases, possibly due to the way the P-Both dataset defines PH as a part of its experiment design. From a model selection point of view, the PH-ELM is best for generalizing to H-PCG, and for generalization to H-SCG, it is unclear which model is preferable.

Regarding dataset selection, we interpret the results to suggest that more data can improve performance because using the full P-Both training dataset gives higher performance than either P-SCG or P-PCG alone. This result may suggest that combining modalities (both PCG and SCG) may improve performance.

The PH-ELM generalizes the best: Figure 3 summarizes the prediction performance of the evaluated models by showing box plots of the min, 25%, median, 75%, and maximum prediction performances over selected rows of Tables IV and V (the rows corresponding to the P-SCG and P-PCG datasets are excluded since they are covered by P-Both). The median is derived via Equation (15). The results show that the PH-ELM model has the highest median performance by +0.09 in both auROC Ratio and AP* Ratio over the closest model, DenseNet121. Indeed, the median auROC scores are 0.884 (PH-ELM), 0.815 (DenseNet121), 0.796 (Scattering-SVM), and 0.643 (STFT-SVM). In general, the PH-ELM model does the best job of preserving generalization performance that could be obtained by in-distribution analysis.

C. Computational Performance

Table VI shows the computational footprint of the four models. Two columns in the table count the total number of learned and fixed (not learned) parameters for each model. To count parameters in an SVM model, we count the number of coordinates across all support vectors of the SVM trained on the entire H-PCG dataset. Since the number of support vectors depends on dataset size, this measurement can vary, but there were 42 support vectors after training on H-PCG (one per subject in the dataset) for both STFT-SVM and Scattering-SVM. The values reported in columns Time to Train and Power to Train each report the average and standard deviation of three measurements. Each measurement was performed on the same dedicated Linux server, running no other jobs, with an NVIDIA RTX 2080 GPU and AMD Ryzen Threadripper 2920X 12-Core CPU. The time estimate for each measurement was obtained by training on H-PCG and evaluating on H-SCG (we assume the added overhead of time spent evaluating H-SCG is negligible). Independently of the timing estimates, each wattage estimate was obtained by sampling the Linux operating system (at location `/sys/class/powercap/*/energy_uj`) for CPU power usage and by sampling the NVIDIA graphics card (via the Linux tool `nvidia-smi`) for power usage. The watt numbers were recorded once per second during the period of time when the model was training (not evaluating) and then summed together. The Inference Time per Subject was computed by averaging the elapsed real time spent to generate a prediction of a single sample (i.e. batch size of one sample), using a notebook computer, Lenovo P1, with a 13th Gen Intel(R) Core(TM) i7-13700H CPU and NVIDIA RTX 4060 GPU. All variations in the table show one standard deviation.

The PH-ELM trains most quickly, requires the least amount of power to train, has the smallest number of learned parameters, the smallest number of parameters (0.047e6) overall, and is fast for inference. In contrast, the DenseNet121, which is the most competitive model with regards to classification performance, has over 300× more learned parameters, takes 36× longer to train on a decent quality GPU, and utilizes over 500× more watts of power to train. Note also that the PH-ELM is non-iterative. As the dataset size increases, the time and power required to train a deep network (or SVM model) will increase significantly due to the iterative nature of the algorithms, whereas the PH-ELM is bounded by the time needed to compute features from a single convolution layer. Finally, the STFT-SVM model is notable for extremely fast inference speed and nearly zero not learned parameters, but it ranks last by a large margin with regards to the median prediction performance shown in Figure 3.

D. Varying Number of Heartbeats

This section analyzes if there is a minimum number of heartbeats per subject necessary for PH detection. For instance, it would be beneficial to know if a clinician only needs a short heart sound recording for PH detection. Moreover, we can also ask whether a longer recording improves the quality of the training data.

Figure 4 shows the effect of reducing the number of heartbeats per subject in either the training set, validation set, or both. This test is performed on the H-PCG dataset via cross-validation with the PH-ELM model, and the figure reports the in-distribution auROC via Equation (10) and 12 bootstrap iterations. Each bar in the barplot reports the average and one standard deviation.

We observe that the PH-ELM model is somewhat insensitive to the varying number of heartbeats. While more heartbeats per subject may slightly increase the prediction performance, 40 heartbeats may offer an acceptable trade-off between recording time and performance. We interpret the figure to imply that the mean number of heartbeats in the training set is slightly more important than in the test set, since at 10 mean heartbeats per subject, the reported auROC drops more when reducing the training set (green bar) than the test set (orange bar).

V. DISCUSSION

The automated detection of Pulmonary Hypertension from heart sound analysis is a viable and robust technology that generalizes across domains and modalities. The results of our study show strong PH detection performance to human test data across all available combinations of three datasets. We applied a rigorous evaluation including two different modalities (PCG and SCG), two animals (humans and pigs), and in-distribution and out-of-distribution evaluation. The model with the best prediction performance, PH-ELM, is also computationally lightweight and well-suited to deployment on edge devices such as mobile phones or small portable PH detection machines. Our study offers evidence that a heart sound PH detector could be recognized in clinical guidelines as a useful and reliable technology for the detection of PH. A limitation

TABLE III: In-Distribution Prediction Performance – via Equation (10)

Dataset	Model	auROC	AP*	σ (auROC)	σ (AP*)
H-PCG	DenseNet121	0.932	0.728	0.029	0.038
H-PCG	PH-ELM	<u>0.922</u>	<u>0.726</u>	0.012	0.015
H-PCG	STFT-SVM	0.891	0.695	0.023	0.024
H-PCG	Scattering-SVM	0.827	0.573	0.019	0.047
H-SCG	DenseNet121	<u>0.547</u>	<u>0.274</u>	0.062	0.038
H-SCG	PH-ELM	0.650	0.335	0.024	0.021
H-SCG	STFT-SVM	0.541	0.271	0.032	0.012
H-SCG	Scattering-SVM	0.537	0.233	0.039	0.027
P-Both	DenseNet121	0.874	0.820	0.019	0.024
P-Both	PH-ELM	<u>0.816</u>	<u>0.765</u>	0.001	0.002
P-Both	STFT-SVM	<u>0.656</u>	<u>0.622</u>	0.000	0.000
P-Both	Scattering-SVM	0.694	0.574	0.000	0.000
P-PCG	DenseNet121	<u>0.924</u>	0.916	0.023	0.024
P-PCG	PH-ELM	0.894	0.872	0.007	0.019
P-PCG	STFT-SVM	0.720	0.741	0.000	0.000
P-PCG	Scattering-SVM	0.934	<u>0.903</u>	0.000	0.000
P-SCG	DenseNet121	0.860	0.784	0.024	0.045
P-SCG	PH-ELM	<u>0.784</u>	<u>0.732</u>	0.002	0.004
P-SCG	STFT-SVM	0.609	0.509	0.000	0.001
P-SCG	Scattering-SVM	0.697	0.546	0.000	0.000

Each row shows the average performance of 12 independent bootstrap iterations with cross-validation. P-Both had leave-one-pig-out CV (10 pigs), H-PCG had 10-fold CV, and H-SCG had 10-fold grouped CV to ensure that each subject's data was entirely in either the train or the validation set.

of the datasets considered is that all subjects underwent the right heart catheterization, which means the presentation of PH is sufficiently advanced to warrant an invasive and expensive procedure. The results of our work justify the need for studies of *early* PH detection, such as by following the long-term outcomes of presumably healthy individuals who are regularly subjected to PH screenings alongside existing technologies like ECHO, EKG, and blood tests.

A. In-distribution and Out-of-distribution Testing

Out-of-distribution testing offers benefits not available from in-distribution testing: The analysis of related work in Section II and Table I shows that none of the works surveyed have evaluated the empirical out-of-distribution performance with a separate dataset, possibly due to the challenge of obtaining PH datasets. Our present work therefore contributes improved evaluation techniques by utilizing both in-distribution and out-of-distribution analyses.

In-distribution performance reported on any single dataset is not representative of expected out-of-distribution results. While Table III reports an in-distribution standard deviation, the reported standard deviation can be made arbitrarily smaller by increasing the number of bootstrap iterations, a phenomenon explained by the Central Limit Theorem (CLT) [37]. This variability for a given dataset is not representative of the variability of the expected out-of-distribution performance when the dataset is used as a training or test set. Moreover, if the expected value of the in-distribution auROC was predictive of the out-of-distribution auROC, we would see an auROC ratio close to one, and we do not observe that in Tables IV and V; the same logic applies to AP*. In Table V, the expected

TABLE IV: SCG \leftrightarrow PCG, Out-of-Distribution Prediction Performance

D_{train}	D_{test}	Model	AP* Ratio	auROC Ratio	AP*	auROC	$\sigma(\text{AP}^*)$	$\sigma(\text{auROC})$
H-SCG	H-PCG	DenseNet121	0.401	0.506	0.292	0.471	0.042	0.068
H-SCG	H-PCG	PH-ELM	<u>0.452</u>	<u>0.553</u>	<u>0.329</u>	<u>0.515</u>	0.020	0.020
H-SCG	H-PCG	STFT-SVM	0.418	0.491	0.304	0.458	0.036	0.002
H-SCG	H-PCG	Scattering-SVM	0.746	0.828	0.543	0.771	0.068	0.094
H-PCG	H-SCG	DenseNet121	0.897	0.917	0.300	0.596	0.027	0.028
H-PCG	H-SCG	PH-ELM	1.082	0.978	0.362	0.636	0.004	0.005
H-PCG	H-SCG	STFT-SVM	0.819	0.912	0.274	0.593	0.000	0.000
H-PCG	H-SCG	Scattering-SVM	0.689	0.897	0.231	0.583	0.000	0.000
P-SCG	P-PCG	DenseNet121	0.817	0.944	0.748	0.872	0.017	0.007
P-SCG	P-PCG	PH-ELM	<u>0.650</u>	<u>0.821</u>	<u>0.595</u>	<u>0.759</u>	0.005	0.003
P-SCG	P-PCG	STFT-SVM	0.491	0.585	0.450	0.540	0.001	0.001
P-SCG	P-PCG	Scattering-SVM	0.545	0.636	0.499	0.588	0.009	0.015
P-PCG	P-SCG	DenseNet121	<u>1.136</u>	<u>1.057</u>	<u>0.891</u>	<u>0.909</u>	0.047	0.041
P-PCG	P-SCG	PH-ELM	1.213	1.125	0.951	0.967	0.030	0.021
P-PCG	P-SCG	STFT-SVM	0.509	0.530	0.399	0.456	0.000	0.000
P-PCG	P-SCG	Scattering-SVM	0.694	0.619	0.544	0.532	0.023	0.006

Each row shows the average performance of 12 independent bootstrap iterations for the given training dataset, test dataset, and model.

TABLE V: Pig \rightarrow Human, Out-of-Distribution Prediction Performance

D_{train}	D_{test}	Model	AP* Ratio	auROC Ratio	AP*	auROC	$\sigma(\text{AP}^*)$	$\sigma(\text{auROC})$
P-Both	H-PCG	DenseNet121	0.630	0.735	0.458	0.685	0.059	0.052
P-Both	H-PCG	PH-ELM	0.781	0.883	0.568	0.822	0.013	0.007
P-Both	H-PCG	STFT-SVM	0.576	0.642	0.419	0.598	0.001	0.001
P-Both	H-PCG	Scattering-SVM	0.468	0.555	0.341	0.517	0.000	0.000
P-PCG	H-PCG	DenseNet121	0.464	0.642	0.338	0.598	0.031	0.030
P-PCG	H-PCG	PH-ELM	0.544	0.682	0.396	0.635	0.006	0.006
P-PCG	H-PCG	STFT-SVM	0.430	0.473	0.313	0.440	0.000	0.000
P-PCG	H-PCG	Scattering-SVM	0.000	0.518	0.000	0.483	0.000	0.000
P-SCG	H-PCG	DenseNet121	0.519	0.649	0.378	0.605	0.069	0.080
P-SCG	H-PCG	PH-ELM	<u>0.517</u>	0.684	<u>0.376</u>	0.637	0.032	0.034
P-SCG	H-PCG	STFT-SVM	0.490	0.640	0.357	0.596	0.001	0.001
P-SCG	H-PCG	Scattering-SVM	0.468	0.555	0.341	0.517	0.000	0.000
P-Both	H-SCG	DenseNet121	0.897	0.991	0.300	0.645	0.030	0.026
P-Both	H-SCG	PH-ELM	0.928	0.954	0.311	0.620	0.020	0.004
P-Both	H-SCG	STFT-SVM	0.620	0.763	0.207	0.496	0.001	0.001
P-Both	H-SCG	Scattering-SVM	1.415	0.796	0.474	0.518	0.000	0.000
P-PCG	H-SCG	DenseNet121	0.821	0.896	<u>0.275</u>	<u>0.582</u>	0.019	0.018
P-PCG	H-SCG	PH-ELM	<u>0.822</u>	0.927	<u>0.275</u>	0.603	0.001	0.002
P-PCG	H-SCG	STFT-SVM	0.618	0.645	0.207	0.419	0.000	0.000
P-PCG	H-SCG	Scattering-SVM	1.415	0.796	0.474	0.518	0.000	0.000
P-SCG	H-SCG	DenseNet121	<u>0.815</u>	0.916	<u>0.273</u>	0.596	0.043	0.050
P-SCG	H-SCG	PH-ELM	0.813	0.885	0.272	0.575	0.013	0.014
P-SCG	H-SCG	STFT-SVM	0.792	0.849	0.265	0.552	0.001	0.001
P-SCG	H-SCG	Scattering-SVM	1.415	0.796	0.474	0.518	0.000	0.000

Each row shows the average performance of 12 independent bootstrap iterations for the given training dataset, test dataset, and model.

787 test set auROC of the PH-ELM model is 0.82, 0.64, 0.64,
 788 0.62, 0.60, and 0.58 (each also with small standard variation
 789 due to the CLT and twelve bootstrap iterations). Excluding the
 790 value 0.82, there is a relatively small variation in these results.
 791 The 0.82 auROC value appears when the dataset P-Both is
 792 the training set and H-PCG is the test set. The DenseNet121
 793 follows this trend, with its highest performance of 0.69 auROC
 794 when (P-Both, H-PCG) are the (training, test) sets. Since P-
 795 Both combines SCG and PCG data, this out-of-distribution
 796 result offers preliminary evidence that training with multiple
 797 modalities (PCG and SCG data) may improve performance
 798 on PCG test data. Comparing these reported results to the in-
 799 distribution results in Table III, we can observe that no single

in-distribution number reliably predicts the out-of-distribution
 performances.

In-distribution results can mislead model selection in a possible real-world setting. A realistic setting could arise where we have P-Both training data, and we wish to train and deploy a model to the hospital where H-PCG data is generated. Based on the in-distribution results with P-Both, we would choose DenseNet121. However, the corresponding out-of-distribution experiments (P-Both, H-PCG) in the first two rows of Table V show that the PH-ELM is preferable over DenseNet121 by 0.14 auROC. Model selection based on the reported in-distribution results can therefore be misleading.

The machine learning literature justifies the benefits of

TABLE VI: Computational Resource Footprint

Model	Learned Parameters	Not Learned Parameters	Time to Train ¹ (s)	Power to Train ¹ (W)	Inference Time per Subject ² (s)
DenseNet121 ³	6.95e6	0.084e6	318.8 ± 1.16	89,481 ± 92	0.062 ± 0.044
PH-ELM	0.02e6	0.03e6	8.80 ± 0.04	168 ± 17	0.0014 ± 0.0014
Scattering-SVM	17.8e6**	1.46e6	19.74 ± 0.14	2020 ± 36	0.0356 ± 0.0011
STFT-SVM	<u>0.14e6**</u>	64	<u>12.15</u> ± 0.05	<u>370</u> ± 8	0.00016 ± 0.00002

** The parameter count in SVM models depends on the number of support vectors, which depends on the training set and hyperparameter constraint C , and is therefore variable. The full H-PCG dataset was utilized to give parameter count estimates, and the SVM models learned 42 support vectors after training. The Scattering parts of Scattering-SVM contribute non-learned parameters.

1. Time to Train and Power to Train computed on a server with NVIDIA RTX 2080 GPU and AMD Ryzen Threadripper 2920X 12-Core Processor. Power is obtained by sampling the CPU and GPU watt usage every second, summing over the number of seconds of training time, and averaging across 3 runs.

2. Inference Time per Subject computed on a notebook with 13th Gen Intel(R) Core(TM) i7-13700H CPU and NVIDIA RTX 4060 GPU. SVM models are evaluated on the CPU, though the wavelet scattering for ScatteringSVM is computed on the GPU. Time may vary widely from system to system. Results average the performance across all samples in the H-PCG dataset.

3. DenseNet121 evaluated for 500 epochs. Training could be completed in as few as 150 epochs, at negligible performance loss, for most datasets.

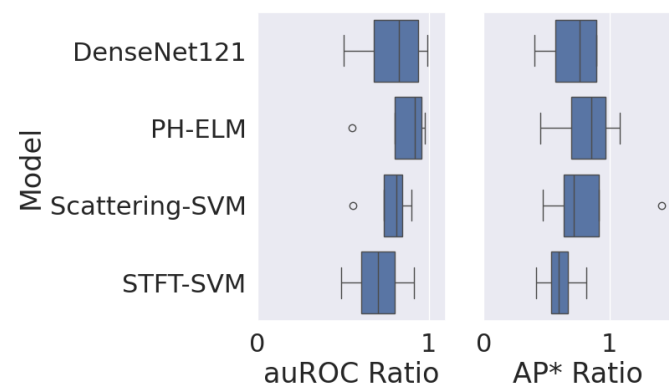


Fig. 3: Model Selection: All models have similar prediction performance, and the PH-ELM has the highest median out-of-distribution performance when evaluating H-PCG or H-SCG test sets with any of the training sets {P-Both, H-PCG or H-SCG}, with median auROC ratio and median AP* ratio of +0.09 (in both cases) over DenseNet121. Each boxplot shows 0%, 25%, 50%, 75%, and 100% percentiles of 12 bootstrap iterations. This figure summarizes rows from Tables IV and V.

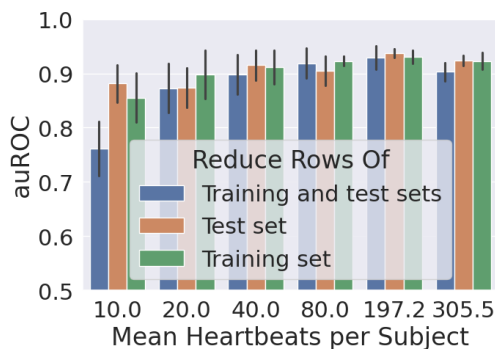


Fig. 4: Insensitive to a varying number of Heartbeats: In-distribution prediction performance with the H-PCG dataset and PH-ELM model is almost unchanged when either the test data or training data has as few as 10 heartbeats. Each bar shows the average and one standard deviation of 12 independently trained PH detectors.

OOD testing: A recent work claims that in-distribution and out-of-distribution performance can be inversely correlated, and therefore that “studies on OOD generalization that use ID performance for model selection (a common recommended practice) will necessarily miss the best-performing models, making these studies blind to a whole range of phenomena” [38], where ID means in-distribution. Similar works describe underspecification as a phenomenon of large variations in test time performance from a set of models that have equally good in-distribution performance [39]. Spurious correlation, as surveyed in [40], can also explain some kinds of test set variability. The literature therefore provides ample evidence that in-distribution and out-of-distribution performances can be completely different, and even inversely correlated. This literature shows that using cross-validation results, which is an in-distribution test, to perform model selection for application in an out-of-distribution setting, is not necessarily justified or reliable.

B. PH-ELM Model Design

The ELM architecture was originally designed for prediction tasks on tabular data [41], especially small datasets [42], using random and fixed random projection matrix followed by a nonlinear operator to generate features, and there are many variations of the basic architecture [43]. The application of ELM to images involves converting the 2-D or 3-D image data into a 1-D vector [44]. Previous works have trained a CNN using backpropagation and then utilized its outputs as a feature extractor for an ELM model, such as in the analysis of retinal fundus images for diabetic retinopathy [45], and to detect QRS complex in heart electrocardiograms [46]. In contrast, our approach does not train the CNN or use an iterative optimization algorithm. Similarly to our work, a fixed-weight convolutional layer followed by feature pooling was applied to images, where the convolutional kernel weights were initialized with Gaussian random parameters and then orthogonalized using SVD⁴ [47]. Some works have considered SVD compression and ELM; singular values were utilized as the input to an ELM classifier [48], and another work shows that using PCA to compress the data just before passing it to

⁴In [47], SVD was for parameter initialization, not for compression of the generated convolutional features.

an ELM classifier increases computational efficiency [49]. Our approach is distinct from these related works. We adopt a variation of the ELM architecture, where (a) our parallelized CNN implementation with an adaptive pooling function replaces the ELM's hidden layer, (b) we apply a PCA projection matrix to compress the CNN feature representation, and (c) we assemble the training data into a matrix and compute a pseudo-inverse like a standard ELM procure would.

Another machine learning pipeline that is similar to our ELM-based approach is the wavelet scattering network [50], which has been described as a convolutional network [51] with a fixed-weight architecture containing one to three convolutional structures, each followed by a non-linear operator. Wavelet scattering outputs a data structure that can be compressed, such as with PCA, and then used to train or evaluate a linear model [51]. Similar approaches based on wavelet packet compression [52] also use a single convolution layer and demonstrate no loss in prediction performance when parameters from the convolutional structure are removed or obscured. Convolution layers are linear functions by design, and they therefore fit into the ELM framework. These works therefore justify the PH-ELM's adoption of a convolutional feature-generating structure with fixed-weight parameters, and they also justify our comparative evaluations of the Scattering-SVM model.

Few works have previously applied the ELM architecture to heart sound analysis [53]–[56]. To the best of our knowledge, there are no existing published works that apply ELM to SCG data. The work of Liu et al. [53] utilized a standard ELM architecture on a vector of eleven features to predict Heart failure with preserved ejection fraction (HFpEF). For the detection of heart murmurs, an ELM was found to have performance similar to a support vector machine [54], and a Deep ELM network, or ELM with multiple hidden layers, was also successfully utilized for murmur detection [55]. Last, Ghosh et al. [56] propose a method that uses the ELM architecture, as an autoencoder, to reconstruct an image representation of heart sounds from a set of derived features, and the reconstructed image was then passed to a kernel sparse regression algorithm to classify any of five different diseases.

VI. CONCLUSIONS

For the automated detection of pulmonary hypertension from heart sounds, we proposed PH-ELM, a novel PH detection algorithm based on the extreme learning machine, that generalizes reliably across the analyzed datasets and is computationally efficient. We also developed a rigorous evaluation methodology based on both in-distribution and out-of-distribution evaluations. Our results on three datasets show that two of four ML algorithms generalize well across PCG and SCG modalities, as well as from pigs to humans. To the best of our knowledge, this work is novel for its use of multiple datasets in the evaluation of Pulmonary Hypertension detection from heart sounds.

REFERENCES

[1] F. Valentin, *Hurst's the Heart, (2 Volume Set)*. McGraw-Hill Medical, 2011.

[2] M. Humbert, G. Kovacs, M. M. Hoeper, R. Badagliacca, R. M. Berger, M. Brida, J. Carlsen, A. J. Coats, P. Escribano-Subias, P. Ferrari *et al.*, "2022 esc/ers guidelines for the diagnosis and treatment of pulmonary hypertension," *European Respiratory Journal*, vol. 61, p. 2200879, 2023.

[3] M. M. Hoeper, H.-A. Ghofrani, E. Grünig, H. Klose, H. Olschewski, and S. Rosenkranz, "Pulmonary hypertension," *Deutsches Ärzteblatt International*, vol. 114, no. 5, p. 73, 2017.

[4] M. M. Hoeper, M. Humbert, R. Souza, M. Idrees, S. M. Kawut, K. Sliwa-Hahnle, Z.-C. Jing, and J. S. R. Gibbs, "A global view of pulmonary hypertension," *The Lancet Respiratory Medicine*, vol. 4, no. 4, pp. 306–322, 2016.

[5] A. Frost, D. Badesch, J. S. R. Gibbs, D. Gopalan, D. Khanna, A. Manes, R. Oudiz, T. Satoh, F. Torres, and A. Torbicki, "Diagnosis of pulmonary hypertension," *European Respiratory Journal*, vol. 53, no. 1, 2019.

[6] M. Ginoux, S. Turquier, N. Chebib, J.-C. Glerant, J. Traclet, F. Philit, A. Sénéchal, J.-F. Mornex, and V. Cottin, "Impact of comorbidities and delay in diagnosis in elderly patients with pulmonary hypertension," *ERJ Open Research*, vol. 4, no. 4, 2018.

[7] E. M. Lau, M. Humbert, and D. S. Celermajer, "Early detection of pulmonary arterial hypertension," *Nature Reviews Cardiology*, vol. 12, no. 3, pp. 143–155, 2015.

[8] S. Janda, N. Shahidi, K. Gin, and J. Swiston, "Diagnostic accuracy of echocardiography for pulmonary hypertension: a systematic review and meta-analysis," *Heart*, 2011.

[9] J. Xu, L. Durand, and P. Pibarot, "A new, simple, and accurate method for non-invasive estimation of pulmonary arterial pressure," *Heart*, vol. 88, no. 1, p. 76, 2002.

[10] N. Yamakawa, N. Kotooka, T. Kato, T. Kuroda, and K. Node, "Cardiac acoustic biomarkers as surrogate markers to diagnose the phenotypes of pulmonary hypertension: an exploratory study," *Heart and Vessels*, pp. 1–8, 2022.

[11] J. Huang, W. Zhang, W. Fu, J. Le, Y. Qi, X. Hou, X. Pan, R. Li, and B. He, "Noninvasive evaluation of pulmonary hypertension using the second heart sound parameters collected by a mobile cardiac acoustic monitoring system," *Frontiers in Cardiovascular Medicine*, vol. 10, 2023.

[12] A. Dennis, A. D. Michaels, P. Arand, and D. Ventura, "Noninvasive diagnosis of pulmonary hypertension using heart sound analysis," *Computers in Biology and Medicine*, vol. 40, no. 9, pp. 758–764, 2010.

[13] T. Kaddoura, K. Vadlamudi, S. Kumar, P. Bobhate, L. Guo, S. Jain, M. Elgendi, J. Y. Coe, D. Kim, D. Taylor *et al.*, "Acoustic diagnosis of pulmonary hypertension: automated speech-recognition-inspired classification algorithm outperforms physicians," *Scientific Reports*, vol. 6, no. 1, p. 33182, 2016.

[14] M. Wang, B. Guo, Y. Hu, Z. Zhao, C. Liu, and H. Tang, "Transfer learning models for detecting six categories of phonocardiogram recordings," *Journal of Cardiovascular Development and Disease*, vol. 9, no. 3, p. 86, 2022.

[15] A. Gaudio, M. Coimbra, A. Campilho, A. Smailagic, S. E. Schmidt, and F. Renna, "Explainable deep learning for non-invasive detection of pulmonary artery hypertension from heart sounds," in *2022 Computing in Cardiology (CinC)*, vol. 498. IEEE, 2022, pp. 1–4.

[16] B. Ge, H. Yang, P. Ma, T. Guo, J. Pan, and W. Wang, "Detection of pulmonary arterial hypertension associated with congenital heart disease based on time–frequency domain and deep learning features," *Biomedical Signal Processing and Control*, vol. 81, p. 104451, 2023.

[17] A. N. Patnaik, "First heart sound," *Indian Journal of Cardiovascular Disease in Women-WINCARS*, vol. 4, no. 02, pp. 107–109, 2019.

[18] J. M. Felner, "Chapter 23 the second heart sound," *Clinical methods: the history, physical, and laboratory examinations*, p. 122, 1990.

[19] F. Renna, A. Gaudio, S. Mattos, M. D. Plumley, and M. T. Coimbra, "Separation of the aortic and pulmonary components of the second heart sound via alternating optimization," *IEEE Access*, vol. 12, pp. 34 632–34 643, 2024.

[20] N. J. Mehta and I. A. Khan, "Third heart sound: genesis and clinical importance," *International Journal of Cardiology*, vol. 97, no. 2, pp. 183–186, 2004.

[21] A. Gaudio, F. Renna, S. Schmidt, and M. T. Coimbra, "Explainable deep learning method for non-invasive detection of pulmonary hypertension from heart sounds," Worldwide Patent WO 2024 047 610A1, Mar 07, 2024.

[22] A. Gaudio, C. Faloutsos, A. Smailagic, P. Costa, and A. Campilho, "ExplainFix: Explainable spatially fixed deep networks," *WIREs Data Mining and Knowledge Discovery*, vol. 13, no. 2, p. e1483, 2023. [Online]. Available: <https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/widm.1483>

- [23] A. Gaudio and M. Elhilali, "AP*: An average precision score and modified precision recall space for class imbalanced datasets," *Authoria Preprints*, 2024.
- [24] L. A. Jeni, J. F. Cohn, and F. De La Torre, "Facing imbalanced data—recommendations for the use of performance metrics," in *2013 Humaine Association Conference on Affective Computing and Intelligent Interaction*, 2013, pp. 245–251.
- [25] M. McDermott, L. H. Hansen, H. Zhang, G. Angelotti, and J. Gallifant, "A closer look at auroc and auprc under class imbalance," *arXiv preprint arXiv:2401.06091*, 2024.
- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [27] J. Davis and M. Goadrich, "The relationship between precision-recall and roc curves," in *Proceedings of the 23rd International Conference on Machine Learning*, ser. ICML '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 233–240. [Online]. Available: <https://doi.org/10.1145/1143844.1143874>
- [28] A. Gaudio, N. Giordano, M. Coimbra, B. Kjaergaard, S. Schmidt, and F. Renna, "Cross-domain detection of pulmonary hypertension in human and porcine heart sounds," in *2023 Computing in Cardiology (CinC)*, vol. 50, 2023, pp. 1–4.
- [29] M. Chan, L. Klein, J. Fan, and O. Inan, "Scg-rhc: Wearable seismocardiogram signal and right heart catheter database," *PhysioNet*, 2023.
- [30] F. Renna, J. Oliveira, and M. T. Coimbra, "Deep convolutional neural networks for heart sound segmentation," *IEEE Journal of Biomedical and Health Informatics*, vol. 23, no. 6, pp. 2435–2445, 2019.
- [31] N. Giordano, S. Rosati, G. Balestra, and M. Knaflitz, "A wearable multi-sensor array enables the recording of heart sounds in homecare," *Sensors*, vol. 23, no. 13, p. 6241, 2023.
- [32] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4700–4708.
- [33] X. Liu, C. Gao, and P. Li, "A comparative analysis of support vector machines and extreme learning machines," *Neural Networks*, vol. 33, pp. 58–66, 2012.
- [34] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [35] M. Andreux, T. Angles, G. Exarchakis, R. Leonarduzzi, G. Rochette, L. Thiry, J. Zarka, S. Mallat, J. Andén, E. Belilovsky *et al.*, "Kymatio: Scattering transforms in python," *Journal of Machine Learning Research*, vol. 21, no. 60, pp. 1–6, 2020.
- [36] Y.-Y. Yang, M. Hira, Z. Ni, A. Astafurov, C. Chen, C. Pührsch, D. Pollack, D. Genzel, D. Greenberg, E. Z. Yang, J. Lian, J. Hwang, J. Chen, P. Goldsborough, S. Narenthiran, S. Watanabe, S. Hintala, and V. Quenneville-Bélaïr, "Torchaudio: Building blocks for audio and speech processing," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6982–6986.
- [37] K. J. H. Kwak Sang Gyu, "Central limit theorem: the cornerstone of modern statistics," *Korean Journal of Anesthesiology*, vol. 70, no. 2, pp. 144–156, 2017. [Online]. Available: <http://www.e-sciencecentral.org/articles/?scid=1156667>
- [38] D. Teney, Y. Lin, S. J. Oh, and E. Abbasnejad, "ID and OOD performance are sometimes inversely correlated on real-world datasets," in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 71 703–71 722. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2023/file/e304d374c85e385eb217ed4a025b6b63-Paper-Conference.pdf
- [39] A. D'Amour, K. Heller, D. Moldovan, B. Adlam, B. Alipanahi, A. Beutel, C. Chen, J. Deaton, J. Eisenstein, M. D. Hoffman, F. Hormozdiari, N. Hounsby, S. Hou, G. Jerfel, A. Karthikesalingam, M. Lucic, Y. Ma, C. McLean, D. Mincu, A. Mitani, A. Montanari, Z. Nado, V. Natarajan, C. Nielson, T. F. Osborne, R. Raman, K. Ramasamy, R. Sayres, J. Schrouff, M. Seneviratne, S. Sequeira, H. Suresh, V. Veitch, M. Vladymyrov, X. Wang, K. Webster, S. Yadlowsky, T. Yun, X. Zhai, and D. Sculley, "Underspecification presents challenges for credibility in modern machine learning," *Journal of Machine Learning Research*, vol. 23, no. 226, pp. 1–61, 2022. [Online]. Available: <http://jmlr.org/papers/v23/20-1335.html>
- [40] W. Ye, G. Zheng, X. Cao, Y. Ma, and A. Zhang, "Spurious correlations in machine learning: A survey," *arXiv preprint arXiv:2402.12715*, 2024.
- [41] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew, "Extreme learning machine: Theory and applications," *Neurocomputing*, vol. 70, no. 1, pp. 489–501, 2006, neural Networks. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S09252131206000385>
- [42] G. Huang, G.-B. Huang, S. Song, and K. You, "Trends in extreme learning machines: A review," *Neural Networks*, vol. 61, pp. 32–48, 2015.
- [43] J. Wang, S. Lu, S.-H. Wang, and Y.-D. Zhang, "A review on extreme learning machine," *Multimedia Tools and Applications*, vol. 81, no. 29, pp. 41 611–41 660, 2022.
- [44] Y. Huérfano-Maldonado, M. Mora, K. Vilches, R. Hernández-García, R. Gutiérrez, and M. Vera, "A comprehensive review of extreme learning machine on medical imaging," *Neurocomputing*, p. 126618, 2023.
- [45] M. Nahiduzzaman, M. R. Islam, M. O. F. Goni, M. S. Anower, M. Ahsan, J. Haider, and M. Kowalski, "Diabetic retinopathy identification using parallel convolutional neural network based feature extractor and ELM classifier," *Expert Systems with Applications*, vol. 217, p. 119557, 2023.
- [46] S. Zhou and B. Tan, "Electrocardiogram soft computing using hybrid deep learning CNN-ELM," *Applied Soft Computing*, vol. 86, p. 105778, 2020.
- [47] Z. Bai, L. L. C. Kasun, and G.-B. Huang, "Generic object recognition with local receptive fields based extreme learning machine," *Procedia Computer Science*, vol. 53, pp. 391–399, 2015.
- [48] J. Zhang, R. Jiang, B. Li, and N. Xu, "An automatic recognition method of microseismic signals based on EEMD-SVD and ELM," *Computers & Geosciences*, vol. 133, p. 104318, 2019.
- [49] A. Castaño, F. Fernández-Navarro, and C. Hervás-Martínez, "PCA-ELM: a robust and pruned extreme learning machine approach based on principal component analysis," *Neural processing letters*, vol. 37, pp. 377–392, 2013.
- [50] J. Bruna, "Scattering representations for recognition," Ph.D. dissertation, Ecole Polytechnique X, 2013.
- [51] J. Bruna and S. Mallat, "Invariant scattering convolution networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1872–1886, 2013.
- [52] A. Gaudio, A. Smailagic, C. Faloutsos, S. Mohan, E. Johnson, Y. Liu, P. Costa, and A. Campilho, "DeepFixCX: Explainable privacy-preserving image compression for medical image analysis," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 13, no. 4, p. e1495, 2023.
- [53] Y. Liu, X. Guo, and Y. Zheng, "An automatic approach using ELM classifier for HFpEF identification based on heart sound characteristics," *Journal of medical systems*, vol. 43, no. 9, p. 285, 2019.
- [54] X. Yang, F. Yang, L. Gobeawan, S. Y. Yeo, S. Leng, L. Zhong, and Y. Su, "A multi-modal classifier for heart sound recordings," in *2016 Computing in cardiology conference (CinC)*. IEEE, 2016, pp. 1165–1168.
- [55] P. R. Malaysia and B. Pahat, "An innovative machine learning framework for phonocardiography (PCG) using MFCC and deep extreme learning machine (DELM)," *Journal of Theoretical and Applied Information Technology*, vol. 102, no. 22, 2024.
- [56] S. K. Ghosh, R. Ponnalagu, R. Tripathy, and U. R. Acharya, "Deep layer kernel sparse representation network for the detection of heart valve ailments from the time-frequency representation of PCG recordings," *BioMed Research International*, vol. 2020, no. 1, p. 8843963, 2020.