

Reliable Multi-View Learning with Conformal Prediction for Aortic Stenosis Classification in Echocardiography

Ang Nan Gu¹, Michael Tsang², Hooman Vaseli¹, Teresa Tsang², and Purang Abolmaesumi¹

¹ Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, Canada
guangnan@ece.ubc.ca

² Division of Cardiology, Vancouver General Hospital, Vancouver, Canada **

Abstract. The fundamental problem with ultrasound-guided diagnosis is that the acquired images are often 2-D cross-sections of a 3-D anatomy, potentially missing important anatomical details. This limitation leads to challenges in ultrasound echocardiography, such as poor visualization of heart valves or foreshortening of ventricles. Clinicians must interpret these images with inherent uncertainty, a nuance absent in machine learning’s one-hot labels. We propose Re-Training for Uncertainty (RT4U), a data-centric method to introduce uncertainty to weakly informative inputs in the training set. This simple approach can be incorporated to existing state-of-the-art aortic stenosis classification methods to further improve their accuracy. When combined with conformal prediction techniques, RT4U can yield adaptively sized prediction sets which are guaranteed to contain the ground truth class to a high accuracy. We validate the effectiveness of RT4U on three diverse datasets: a public (TMED-2) and a private AS dataset, along with a CIFAR-10-derived toy dataset. Results show improvement on all the datasets. Our source code is publicly available at: <https://github.com/an-michaelg/RT4U>

Keywords: Echocardiography · Aortic Stenosis · Multi-View Learning · Conformal Prediction · Uncertainty Estimation

1 Introduction

Clinical motivation. Aortic Stenosis (AS) is a heart valve disease in which the restricted motion of the aortic valve leaflets (AVL) impairs blood circulation. AS is highly prevalent in older populations [4], posing a growing societal burden with global aging trends. Early detection, monitoring, and timely intervention are crucial due to the low survival rate among untreated AS patients [23]. AS diagnosis typically involves echocardiography (echo), with the decision on disease severity relying mostly on blood flow measurements from spectral Doppler [7].

** T. Tsang and P. Abolmaesumi are joint senior authors.

However, reliance on Doppler limits screening for AS to more specialized echo laboratories that have access to high-end cart-based echo machines. Acquiring spectral Doppler specifically requires very high level of expertise in echo imaging. Recent clinical [1,12,19] and machine learning [2,10,11,14,15,28] works show that AS classification with B-mode echo can achieve high-sensitivity and specificity, which may pave the way for developing screening tools that can be used by less experienced point-of-care practitioners. This work focuses on improving the automated screening of AS with B-mode echo input.

Challenges in interpreting echo exam data. Echocardiograms are 2-D (+time) visualizations of 3-D (+time) patient anatomy, captured by human sonographers. However, variations in patient size and operator skill can result in echo views that lack sufficient diagnostic information [8]. Directly training a model to predict AS severity from any cross-sectional view can be misleading due to limited AVL visibility. This limitation holds true even for the parasternal long-axis (PLAX) and parasternal short-axis (PSAX) views, for which acquiring the correct angle to image the AVL requires operator skill. Therefore, during inference, we would like the model to have robust uncertainty estimation (UE) to warn users and prevent failed predictions if the input is inadequate. To the best of our knowledge, simultaneously pursuing these two goals - learning with imperfect data and developing strong UE - is rarely done.

Conformal prediction. Neural networks predictions are often overconfident [13]. Temperature scaling can mitigate this effect, but does not have a rigorous connection between the validation and test set performance. In contrast, Conformal Prediction (CP) is a way of transforming class predictions $p(y|x) \in [0, 1]^K$ to prediction sets $C(x) \subset \{1, \dots, K\}$ (Fig. 1), where K is the number of classes. Crucially, the sets are guaranteed to “cover” the ground truth to some user-specified probability [6] (see section 2).

Contributions. Inspired by a technique for learning with noisy labels [9], we frame the problem as “learning with noisy inputs”. We consider the disparity between one-hot labels and inputs with some inherent ambiguity (e.g., from poor image quality or cross-sectioning) as noise, since these labels do not capture the ambiguity a clinician may encounter when analyzing the echo data. This type of noise, while uncommon in camera images, is ubiquitous in medical imaging due to the 2-D visualization of 3-D anatomy. We propose Re-Training for Uncertainty (RT4U), a training method that uses pseudo-labels to mitigate overfitting to noisy inputs. RT4U is a model-agnostic technique without hyperparameters which can be applied to any deep learning based classifier with minimal overhead. We integrate RT4U with existing state-of-the-art approaches for aortic stenosis classification, improving accuracy and reducing over-confidence.




Fig. 1. Conformal prediction returns a set $\{.\}$ of predictions. The set size correlates with the estimated uncertainty. Our **RT4U** method (green) enables better prediction and more adaptive set sizes compared to prior art (blue). Left: Inference on CIFAR-10 quadrants, orange box indicates the quadrant seen by the network. Right: Inference on echo cross-sections from the TMED-2 dataset. Ground truth is embedded in each echo image.

2 Background and Related Works

Automated detection of aortic stenosis. Deep learning approaches have been proposed for AS severity classification from 2-D echo. Huang et al. [16] used images from multiple standard-plane views, and predicted both the view and the severity. Dai et al. [10] and Holste et al. [15] separated severe AS cases from other cases with PLAX view videos as the input. Ginsberg et al. [11] proposed an ordinal regression based approach. Ahmadi et al. [2] used an attention-based architecture to identify key video frames for classification. Vaseli et al. [25] used an attention-based framework to identify important spatio-temporal components of the video. However, all the aforementioned works utilized one-hot labels. Secondly, none incorporated conformal prediction techniques to improve prediction reliability.

Conformal prediction. Conformal prediction originated from Vovk et al. [27]. Recent advancements such as LABEL [21], APS [20], and RAPS [5] have introduced CP algorithms with lower and more adaptive set sizes. In the medical domain, Lu et al. [18] validated CP on clinical data and extended RAPS to guarantee coverage for specific patient subgroups. Lu et al. [17] extended APS to ordinal regression problems. Stutz et al. [24] extended CP principles to data with multi-rater labels. CP principles have also been applied to skin disease [31] and histopathology images [29].

Mathematical results from [27] show sets produced by CP can be guaranteed to a user-defined coverage error rate α , e.g., 0.1. In essence, prediction sets satisfy the following equation, also known as the ‘‘Marginal Coverage Property’’:

$$1 - \alpha \leq \mathbb{P}(y_{test} \in C(x_{test})) \leq 1 - \alpha + \frac{1}{N + 1}, \quad (1)$$

where $\mathbb{P}(y_{test} \in C(x_{test}))$ is the coverage accuracy, and N is the number of datapoints used as calibration data. Higher coverage accuracy has a trade-off with higher prediction set size, denoted $|C(x)|$. Additionally, Eq. (1) is not always satisfied due to the finite nature of the calibration set; the realistic coverage is distributed around $1 - \alpha$ [26].

3 Methodology

3.1 Training dynamics with input-dependent noise

Deep learning models can overfit complex forms of noise [3]. This includes overfitting to non-informative ultrasound views, which can be considered as input-dependent noise. An observation by Chen et al. [9] shows that neural networks tend to fit simpler examples first. We confirm that this phenomenon also exists for ultrasound images (Fig. 2). Consequently, the evolution of model’s confidence during training can serve as a rough approximation of the ‘noise’ of a particular sample. We can take advantage of this to create new pseudo-labels using the history of model predictions during training (Eq. (2)), where T denotes the total number of epochs, σ denotes softmax, and $\hat{f}(\cdot)$ denotes the logits of the predictive model.

$$y'_i = \frac{1}{T} \sum_{t=1}^T \sigma(\hat{f}(x_i)_t) \quad (2)$$

3.2 RT4U algorithm

Our proposed RT4U algorithm can be applied to any iterative training approach. First, we train the model with $\mathcal{D}_{tr} = \{X_{tr}, Y_{tr}\}$ for T epochs, saving the prediction of training set $\hat{f}(x_i)_t \forall i \in X_{tr}, t \in T$. We form pseudo-labels $y'_i \forall x_i \in X_{tr}$ using Eq. (2). Subsequently, we train the model from scratch with (x_i, y'_i) pairs for T epochs. This re-initialization is necessary for eliminating the impact of overfitting caused by training in the initial round. Subsequently, CP algorithms can be applied to the trained model to generate prediction sets. We opt for the LABEL algorithm [21]. LABEL requires a calibration set $\mathcal{D}_{cal} = \{X_{cal}, Y_{cal}\}$ to find a conformal score quantile, $\hat{q}$. The $\hat{q}$ is then used as a threshold; classes with confidence exceeding it belongs to $C(x)$ for samples $x \in X_{test}$.




Fig. 2. Demonstration of the RT4U algorithm in three steps. 1: When fitting on one-hot labels, the performance over different epochs vary depending on the clarity of the image (blue line). 2: Using the average of the past predictions, we form a pseudo-label (orange dotted line). 3: Fitting on the new pseudo-label improves stability and the model’s uncertainty-awareness (green line).

4 Experiments

4.1 Simulated noise on CIFAR-10 dataset

We first demonstrate the efficacy of RT4U on a toy problem. We generate quadrant-label pairs from images in the CIFAR-10 dataset, creating what we term **CIFAR-Q**. For each image-label pair $(\mathcal{I}_i, y_i)$, we train the network on the four image quadrants $(x_{i1}, y_i), (x_{i2}, y_i), (x_{i3}, y_i), (x_{i4}, y_i)$ (see Fig. 1). CIFAR-Q simulates learning from potentially non-informative 2-D slices of 3-D objects, within a controlled environment constructed from a well-known dataset. We first split the CIFAR-10 training set randomly with a 95:5 ratio, then form X_{tr} and X_{val} from the quadrants of the splits.

We compare the following: Vanilla training with cross-entropy loss (CE), CE with temperature scaling on X_{val} (CE+Temp), CE with temperature scaling and RT4U training (CE+RT4U+Temp), Mean Absolute Error loss (MAE), and Active Negative Loss [30] (ANL) with hyperparameters $(\alpha, \beta, \delta) = (1, 1, 0)$.

4.2 Aortic stenosis classification using image and video data

We integrate RT4U into state-of-the-art methods for AS classification.

In the image domain, we compare with Huang et al. [16] on the public **TMED-2** dataset. This dataset contains 17270 images from 599 echo studies classified as no/early/significant AS. We use the training and validation splits from the provided DEV479 file, excluding studies lacking AS diagnosis labels.

In the video domain, we apply RT4U to enhance two existing approaches, R(2+1)D-18 [11] and ProtoASNet [25], on a private echo video dataset (**AS Private**), extracted from the echo database of a tertiary care hospital under ethics approval. Ultrasound videos were captured with Philips iE33, Vivid i, and Vivid E9 transducers. Each study was categorized as no/mild/moderate/severe AS based on Doppler echo guidelines [7], with moderate and severe grouped as significant. PLAX and PSAX videos were selected from each study by an experienced cardiologist. UI elements around the beam area were removed. The dataset contains 5055 PLAX and 4062 PSAX videos from 2572 studies. We create training, validation and test sets using an 80:10:10 split, ensuring mutually exclusive sets of patients. For R(2+1)D-18 [11], we use the cross-entropy variant. For ProtoASNet, we use hyperparameters from [25]. All experiments are conducted on an NVIDIA TITAN V GPU. Further detail on experimental settings is provided in the supplementary material.

4.3 Metrics

We evaluate approaches with the following metrics at both the instance-level and study-level. “Study” refers to an echo study consisting of multiple images/videos, or an entire image consisting of four quadrants for CIFAR-Q. $B\text{Acc}$ and $B\text{Cov}$ are class-balanced versions of the top-1 accuracy and coverage accuracy. For the metrics pertaining to CP ($B\text{Cov}$ and $C(x)$), we report the median from 100 random trials to confirm the marginal coverage property while accounting for randomness.

4.4 Decision aggregation

We use different approaches for aggregating multiple instance-level predictions into a study-level prediction. For CIFAR-Q experiments and R(2+1)D [11], we employ the method in Satopää et al. [22] to average the instance-level predictions in the logit space (Eq. (3)), where $x_j, j \in [1, M]$ denote the M instances within the same study.

$$\mathbb{P}(y|x_1, x_2 \dots x_M) = \sigma\left(\sum_{j=1}^M \hat{f}(x_j)\right) \quad (3)$$

For Huang et al. [16], we perform a weighted average based on the predicted view probability of PSAX or PLAX. We approximate the view probability with a ResNet-18 model. For ProtoASNet [25], we perform an average of the instance-level predictions $\in [0, 1]^K$ with consideration of the additional “uncertainty” class that is introduced by the ProtoASNet architecture. These aggregation approaches are consistent with their original works.

4.5 Results and discussion

Incorporating RT4U raises the $B\text{Acc}$ for each of the three AS severity classification methods (Tables 1 and 2). This is because RT4U reduces the loss induced

Table 1. Instance-level test set evaluation. We aim to maximize $B\text{Acc}$, minimize $|C(x)|$, and maintain $BCov$ above the user-defined coverage, $1 - \alpha$.

Dataset	Method	$B\text{Acc} \uparrow$	$BCov$	$ C(x) \downarrow$
CIFAR-Q $\alpha = 0.05$	ANL [30]	0.734	1.000	10.0
	MAE	0.735	0.950	2.90
	CE	0.743	0.950	2.38
	CE+Temp	0.743	0.950	2.37
	CE+RT4U	0.765	0.950	2.22
	CE+RT4U+Temp	0.765	0.950	2.22
AS Private $\alpha = 0.1$	R(2+1)D [11]	0.747	0.902	1.41
	R(2+1)D+RT4U	0.761	0.902	1.41
	ProtoASNet [25]	0.750	0.903	1.46
	ProtoASNet+RT4U	0.755	0.904	1.46
TMED-2 $\alpha = 0.1$	Huang et al. [16]	0.610	0.907	2.05
	Huang et al.+RT4U	0.661	0.905	1.89

Table 2. Study-level test set evaluation. We aim to maximize $B\text{Acc}$, minimize $|C(x)|$, and maintain $BCov$ above the user-defined coverage, $1 - \alpha$.

Dataset	Method	$B\text{Acc} \uparrow$	$BCov$	$ C(x) \downarrow$
CIFAR-Q $\alpha = 0.05$	ANL [30]	0.883	0.950	1.32
	MAE	0.887	0.949	1.28
	CE	0.901	0.950	1.19
	CE+Temp	0.901	0.950	1.18
	CE+RT4U	0.907	0.951	1.16
	CE+RT4U+Temp	0.907	0.951	1.14
AS Private $\alpha = 0.1$	R(2+1)D [11]	0.780	0.913	1.35
	R(2+1)D+RT4U	0.816	0.917	1.36
	ProtoASNet [25]	0.789	0.924	1.40
	ProtoASNet+RT4U	0.796	0.924	1.38
TMED-2 $\alpha = 0.1$	Huang et al. [16]	0.733	0.917	1.90
	Huang et al.+RT4U	0.742	0.913	1.99

by less informative inputs (for example, the PSAX view image with poor AVL visualization in Fig. 2). In contrast, training with one-hot labels leads the network to eventually overfit to samples which do not contain any class-revealing information. In both highly- and weakly- informative cases, RT4U sets the target to a more appropriate value.

Training with one-hot labels results in a high degree of over-confidence (Fig. 3). Conversely, training with RT4U makes the model slightly under-confident. This is because in RT4U the labels from highly-informative cases are also slightly smoothed to some value below 1.0. Additionally, we would like to emphasize that prior to temperature scaling, the model trained with RT4U has a lower expected calibration error (ECE). Secondly, both methods can use temperature scaling in post-processing to reduce ECE.




Fig. 3. Left: Calibration performance of R(2+1)D variants with RT4U and temperature scaling, evaluated at the instance-level on the AS Private dataset. Right: Histograms of network confidence for the predicted class. Compared to the baseline (blue), RT4U (green) has a lower expected calibration error, and the confidence distribution is closer to those from models with temperature scaling. Evaluations on CIFAR-Q and TMED2 in the supplementary material.

The noise-robust loss functions, MAE and ANL [30], underperform relative to the cross-entropy baseline. This suggests that these methods, which are designed to handle noise from “label flipping”, are susceptible to noise originating from the input. Furthermore, the model trained with ANL produces uncalibrated outputs, which compromises the construction of prediction sets.

Regarding the balanced coverage, our experiments confirm that the marginal coverage property is satisfied. The coverage is more conservative (i.e. a higher difference between $BCov$ and $1 - \alpha$) for AS Private and TMED-2, due to the smaller number of datapoints in their respective calibration sets. Furthermore, prediction sets for disease severity should aim to be “ordinal”, i.e. there are no discontinuities within the predicted range of severities. We do not enforce this constraint in any experiment. We observe that for all methods (both with and without RT4U), and across both AS tasks, that more than 98% of prediction sets are ordinal.

5 Conclusion and Future Work

We introduce RT4U, a novel training method aimed at tackling the issue of limited information within medical images. Firstly, the pseudo-labels from RT4U can mitigate overfitting to noisy inputs, and allow the network to associate these inputs to a high-uncertainty state. Secondly, in terms of conveying uncertainty, conformal prediction sets with a guaranteed accuracy offer a promising alternative to top-1 predictions and confidence values. Future improvements include introducing hyperparameters to refine the pseudo-label generation process, exploring the relationship between output logits and the size of prediction sets, and

designing models with the primary objective of generating informative prediction sets rather than solely focusing on maximizing top-1 accuracy.

6 Acknowledgements

This work was supported in part by the UBC Advanced Research Computing Center (ARC), the Canadian Institutes of Health Research (CIHR), and the Natural Sciences and Engineering Research Council of Canada (NSERC).

7 Disclosure of Interests

The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abe, Y.: Screening for aortic stenosis using physical examination and echocardiography. *Journal of Echocardiography* **19**(2), 80–85 (2021)
2. Ahmadi, N., Tsang, M., Gu, A., Tsang, T., Abolmaesumi, P.: Transformer-based spatio-temporal analysis for classification of aortic stenosis severity from echocardiography cine series. *IEEE Transactions on Medical Imaging* **43**(1), 366–376 (2024)
3. Algan, G., Ulusoy, I.: Image classification with deep learning in the presence of noisy labels: A survey. *Knowledge-Based Systems* **215**, 106771 (2021)
4. Ancona, R., Pinto, S.C.: Epidemiology of aortic valve stenosis (as) and of aortic valve incompetence (ai): Is the prevalence of as/ai similar in different parts of the world. *European Society of Cardiology* **18**(10) (2020)
5. Angelopoulos, A., Bates, S., Malik, J., Jordan, M.I.: Uncertainty sets for image classifiers using conformal prediction (2022)
6. Angelopoulos, A.N., Bates, S.: A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *CoRR* **abs/2107.07511** (2021)
7. Bonow, R.O., Carabello, B.A., Chatterjee, K., et al.: Acc/aha 2006 guidelines for the management of patients with valvular heart disease: a report of the american college of cardiology/american heart association task force on practice guidelines. *Journal of the American College of Cardiology* **48**(3), e1–e148 (2006)
8. Chamsi-Pasha, M.A., Sengupta, P.P., Zoghbi, W.A.: Handheld echocardiography: current state and future perspectives. *Circulation* **136**(22), 2178–2188 (2017)
9. Chen, P., Ye, J., Chen, G., Zhao, J., Heng, P.A.: Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 35, pp. 11442–11450 (2021)
10. Dai, W., Nazzari, H., Namasivayam, M., Hung, J., Stultz, C.M.: Identifying aortic stenosis with a single parasternal long-axis video using deep learning. *Journal of the American Society of Echocardiography* **36**(1), 116–118 (2023)
11. Ginsberg, T., Tal, R.e., Tsang, M., et al.: Deep video networks for automatic assessment of aortic stenosis in echocardiography. In: *Simplifying Medical Ultrasound: Second International Workshop, Strasbourg, France, September 27, 2021, Proceedings 2*. pp. 202–210. Springer (2021)

12. Gulič, T.G., Makuc, J., Prosen, G., Dinevski, D.: Pocket-size imaging device as a screening tool for aortic stenosis. *Wiener Klinische Wochenschrift* **128**, 348–353 (2016)
13. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *International conference on machine learning*. pp. 1321–1330. PMLR (2017)
14. Guo, X.: Predicting Aortic Stenosis Severity using Deep Learning. Ph.D. thesis, Massachusetts Institute of Technology (2021)
15. Holste, G., Oikonomou, E.K., Mortazavi, B.J., Coppi, A., et al.: Severe aortic stenosis detection by deep learning applied to echocardiography. *European Heart Journal* **44**(43), 4592–4604 (2023)
16. Huang, Z., Long, G., Wessler, B., Hughes, M.C.: A new semi-supervised learning benchmark for classifying view and diagnosing aortic stenosis from echocardiograms. In: *Machine Learning for Healthcare Conference*. pp. 614–647. PMLR (2021)
17. Lu, C., Angelopoulos, A.N., Pomerantz, S.: Improving trustworthiness of ai disease severity rating in medical imaging with ordinal conformal prediction sets. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 545–554. Springer (2022)
18. Lu, C., Lemay, A., Chang, K., Höbel, K., Kalpathy-Cramer, J.: Fair conformal predictors for applications in medical imaging. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 36, pp. 12008–12016 (2022)
19. Nemchyna, O., Soltani, S., Solowjowa, N., Schoenrath, F., et al.: Validity of visual assessment of aortic valve morphology in patients with aortic stenosis using two-dimensional echocardiography. *The International Journal of Cardiovascular Imaging* **37**, 813–823 (2021)
20. Romano, Y., Sesia, M., Candes, E.: Classification with valid and adaptive coverage. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 3581–3591. Curran Associates, Inc. (2020)
21. Sadinle, M., Lei, J., Wasserman, L.: Least ambiguous set-valued classifiers with bounded error levels. *Journal of the American Statistical Association* **114**(525), 223–234 (2019)
22. Satopää, V.A., Baron, J., Foster, D.P., Mellers, B.A., Tetlock, P.E., Ungar, L.H.: Combining multiple probability predictions using a simple logit model. *International Journal of Forecasting* **30**(2), 344–356 (2014)
23. Strange, G., Stewart, S., Celermajer, D., Prior, D., Scalia, G.M., Marwick, T., Ilton, M., Joseph, M., Codde, J., Playford, D., et al.: Poor long-term survival in patients with moderate aortic stenosis. *Journal of the American College of Cardiology* **74**(15), 1851–1863 (2019)
24. Stutz, D., Roy, A.G., Matejovicova, T., et al.: Conformal prediction under ambiguous ground truth. *arXiv preprint arXiv:2307.09302* (2023)
25. Vaseli, H., Gu, A.N., Ahmadi Amiri, S.N., et al.: Protoasnet: Dynamic prototypes for inherently interpretable and uncertainty-aware aortic stenosis classification in echocardiography. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 368–378. Springer (2023)
26. Vovk, V.: Conditional validity of inductive conformal predictors. In: *Asian conference on machine learning*. pp. 475–490. PMLR (2012)
27. Vovk, V., Gammerman, A., Shafer, G.: *Algorithmic learning in a random world*, vol. 29. Springer (2005)

28. Wessler, B.S., Huang, Z., Long Jr, G.M., Pacifici, S., et al.: Automated detection of aortic stenosis using machine learning. *Journal of the American Society of Echocardiography* (2023)
29. Wieslander, H., Harrison, P.J., Skogberg, G., Jackson, S., et al.: Deep learning with conformal prediction for hierarchical analysis of large-scale whole-slide tissue images. *IEEE journal of biomedical and health informatics* **25**(2), 371–380 (2020)
30. Ye, X., Li, X., Liu, T., Sun, Y., et al.: Active negative loss functions for learning with noisy labels. *Advances in Neural Information Processing Systems* **36** (2024)
31. Zhang, Y., Wang, S., Zhang, Y., Chen, D.Z.: Rr-cp: Reliable-region-based conformal prediction for trustworthy medical image classification. In: *International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging*. pp. 12–21. Springer (2023)