

Physics-Informed Neural Networks for Ultrasound-Based Varicose Vein Screening

Fan Yang¹ and Qian Mao²

¹Department of Applied Physics, The Hong Kong Polytechnic University, Hong Kong, 999077

²School of Design, The Hong Kong Polytechnic University, Hong Kong, 999077

ABSTRACT

Early screening of varicose veins using non-invasive and low-cost sensing techniques remains a practical challenge in community and home-based healthcare settings. Conventional ultrasound imaging systems are accurate but require bulky hardware, skilled operators, and clinical environments, which limit accessibility for large-scale screening. This study proposes a physics-informed deep learning framework for one-dimensional ultrasound A-scan analysis to enable lightweight and interpretable varicose vein screening. A convolutional neural network with an encoder–decoder reconstruction branch and a classification head was developed. The reconstruction branch incorporates physics-motivated constraints, including second-order temporal smoothness and signal energy consistency, to regularize the learned waveform representation. These constraints aim to preserve physically meaningful propagation characteristics while suppressing noise and overfitting. A stratified train–test split with a guaranteed test size of 50 samples was employed to improve statistical reliability under small-sample conditions. Experimental results on 75 labeled A-scan segments demonstrated stable convergence and high screening performance. The final model achieved approximately 98% test accuracy with zero false negatives, indicating strong sensitivity for detecting abnormal vascular conditions. The physics-informed reconstruction produced smoother yet structurally consistent waveforms compared with raw signals, suggesting improved interpretability of learned representations. The findings indicate that physics-guided learning can enhance robustness and clinical relevance in small-sample ultrasound screening tasks. While further validation with participant-level separation and larger cohorts is required, the proposed framework provides a feasible direction for portable, AI-assisted vascular screening in non-clinical environments.

Keywords: Physics-informed neural networks, Ultrasound screening, Varicose veins, Health informatics, Biomedical signal analysis

INTRODUCTION

Ultrasound (US) imaging is among the most widely used diagnostic imaging modalities in routine clinical practice and continues to expand due to its portability and real-time capability (Liu et al., 2019). Varicose veins are a prevalent vascular disorder associated with venous insufficiency and may lead to discomfort and long-term complications if not managed early. Current clinical pathways for confirming superficial venous incompetence

rely heavily on duplex ultrasound (DUS), which is supported as the gold-standard diagnostic test in evidence syntheses and guideline-oriented reviews (Farah et al., 2022). However, DUS examinations are typically performed in specialized clinical settings and require trained operators and dedicated equipment, which limits accessibility for large-scale community screening and home-based monitoring, particularly for ageing populations.

Recent advances in artificial intelligence (AI) and low-cost medical sensing have enabled lightweight diagnostic approaches based on simplified signal acquisition. In pulse-echo ultrasound systems, a one-dimensional amplitude line (A-mode/A-scan) represents echo amplitude as a function of depth/time and constitutes the fundamental building block from which B-mode images can be formed (Powers & Kremkau, 2011). Compared with full imaging modalities, A-scan sensing offers advantages in portability, cost, and data efficiency, making it attractive for point-of-care and at-home screening scenarios. Nevertheless, extracting clinically meaningful information from raw A-scan signals remains challenging because of speckle/noise, inter-subject variability, and limited labeled datasets—issues that are widely recognized in medical ultrasound analysis (Farah et al., 2022). Deep learning methods have shown strong potential for ultrasound analysis tasks such as classification, detection, and segmentation (Farah et al., 2022). However, purely data-driven models may overfit and exhibit limited interpretability in small-sample biomedical settings. Incorporating physical knowledge into learning has therefore emerged as a promising strategy to improve robustness and generalization. Physics-informed neural networks (PINNs) and related physics-guided regularization frameworks introduce constraints derived from governing laws or physically plausible priors, thereby guiding models toward physically consistent representations and improving data efficiency (Raissi et al., 2019). Motivated by these developments, this study investigates a physics-informed deep learning approach for ultrasound A-scan-based varicose vein screening, aiming to develop a lightweight and more interpretable screening framework suitable for portable sensing environments. The objective is to develop a lightweight, interpretable, and robust screening framework suitable for portable sensing environments.

Methodology

The success of deep learning can be attributed to its ability to achieve excellent learning performance by leveraging a large number of labeled training samples. The effectiveness of deep learning models is largely dependent on the availability of sufficiently large and well-annotated datasets. In medical ultrasound imaging, however, the acquisition of labeled data is constrained by several practical factors, including the requirement for expert annotation, inter-observer variability, and the relatively low prevalence of certain pathological conditions. These limitations make the direct training of deep neural networks prone to overfitting and unstable generalization.

Consequently, improving model robustness under limited training samples has become a central issue in medical ultrasound image analysis. Addressing this challenge requires methodological strategies that go beyond conventional

supervised learning, including model regularization, knowledge transfer, and the incorporation of domain-specific priors.

Model optimization in convolutional neural networks aims to enhance predictive performance and generalization capability through architectural design, parameter regularization, and training strategy refinement. Regularization techniques such as dropout, weight decay, and early stopping are commonly employed to mitigate overfitting by constraining the effective capacity of the network and preventing excessive adaptation to training data.

In the present study, several stabilization mechanisms were integrated into the CNN framework:

- Weight decay was applied through AdamW optimization to limit parameter magnitude and improve generalization.
- Dropout layers were introduced in the classification head to reduce co-adaptation of latent features.
- Gradient clipping was adopted to maintain numerical stability during optimization.
- Stratified data splitting with a fixed minimum test size was implemented to improve statistical reliability in performance evaluation under small-sample conditions.

These strategies collectively contribute to a more stable learning process and reduce the risk of optimistic bias caused by insufficient evaluation data.

Transfer learning provides an alternative pathway for addressing data scarcity by leveraging knowledge learned from large-scale datasets. Pre-trained models can capture generic low-level and mid-level feature representations that remain useful across different imaging domains. Fine-tuning such models on small medical datasets can significantly improve convergence speed and predictive stability.

However, in one-dimensional ultrasound A-scan analysis, the domain discrepancy between natural image datasets and temporal acoustic signals limits the direct applicability of conventional transfer learning. Therefore, instead of relying on external pre-training, the present work focuses on physics-guided inductive bias as a complementary mechanism for improving generalization in small-sample scenarios.

To enhance robustness and interpretability, a physics-informed deep learning framework was introduced. Rather than treating ultrasound signals as purely abstract data, the model incorporates signal-level priors derived from physical propagation characteristics.

Given an input A-scan signal $x(t)$, the network predicts both a class label \hat{y} and a reconstructed waveform $\hat{u}(t)$:

$$(\hat{y}, \hat{u}(t)) = f_{\theta}(x(t)). \quad (1)$$

Two physics-motivated constraints are imposed.

Ultrasound wave propagation in biological tissue exhibits continuity in time.

This property is enforced using a discrete second-order derivative penalty:

$$\mathcal{L}_{smooth} = \frac{1}{T-2} \sum_{t=2}^{T-1} (u_{t+1} - 2u_t + u_{t-1})^2. \quad (2)$$

This constraint suppresses high-frequency noise while preserving structural echo patterns.

Physical signal propagation approximately conserves energy within short temporal windows. Therefore, an energy matching loss is defined as:

$$\mathcal{L}_{energy} = (\|u\|_2^2 - \|x\|_2^2)^2. \quad (3)$$

This prevents trivial attenuation of reconstructed signals and maintains amplitude realism.

The final optimization objective integrates classification, reconstruction, and physics constraints:

$$\mathcal{L} = \mathcal{L}_{cls} + \lambda_{rec} \mathcal{L}_{rec} + \lambda_{phys} \mathcal{L}_{phys}, \quad (4)$$

where

$$\mathcal{L}_{phys} = \mathcal{L}_{smooth} + 0.5 \mathcal{L}_{energy}. \quad (5)$$

This formulation enables the network to learn physically meaningful latent representations while maintaining discriminative capability.

Beyond signal-level analysis, clinical interpretation of medical images inherently depends on patient-specific contextual information, including demographic characteristics, clinical history, and diagnostic findings. Conventional CNN-based classification pipelines typically rely solely on image labels, which may limit diagnostic reliability and personalization.

In the proposed framework, patient informatics is incorporated during the labeling and interpretation stage of ultrasound A-scan data. By associating each signal segment with clinically derived labels informed by patient context, the training data more accurately reflects real diagnostic conditions. This integration aims to:

- Improve label reliability and clinical relevance
- Provide implicit contextual priors for model learning
- Facilitate future extension toward multimodal and three-dimensional ultrasound analysis

The inclusion of patient informatics represents a step toward context-aware medical AI, enabling more precise and personalized diagnostic modeling.

Further research and exploration are warranted to fully exploit the benefits of integrating patient informatics into CNN-based medical image analysis frameworks.

Data and Medical Imaging Label

The dataset comprises 75 one-dimensional ultrasound A-scan segments, each represented by 2046 temporal samples acquired at a sampling frequency of 264 MHz, as shown in Figure 1.

3D Waterfall A-scan Comparison (5 MHz probe, fs = 264 MHz)




Figure 1: Representative A-scan waveform comparisons for normal and varicose conditions. Physics-informed reconstruction generates smoother signals while preserving key echo structures, supporting the interpretability of learned representations and the stability of classification performance.

Prior to model training, each A-scan segment underwent per-sample normalization to reduce inter-sample amplitude variability and stabilize optimization. Specifically, the signal mean was removed to suppress DC bias, and the waveform was scaled by its root-mean-square (RMS) magnitude to obtain a consistent amplitude range across samples. This normalization was performed independently for each segment to avoid leakage of global statistics into the evaluation set.

Unlike generic image classification tasks, ultrasound-based vascular screening is inherently linked to clinical interpretation. In the present setting, labels were treated as medically grounded supervision signals reflecting vascular status, rather than purely visual categories. This design aligns model training with the practical screening objective: identifying abnormal vascular patterns that may warrant further assessment.

It is noted that the dataset exhibits class imbalance (59 varicose vs. 16 normal). Therefore, evaluation metrics beyond accuracy (e.g., sensitivity and specificity) are reported to better reflect screening utility.

To reduce optimistic bias under small-sample conditions, a stratified train–test strategy was adopted with an explicit emphasis on evaluation stability. The split follows two principles:

1. Stratification to preserve the label distribution across subsets.
2. A minimum test set size of 50 to obtain a more statistically meaningful estimate of generalization performance compared with very small test sets.

The resulting split was: Training set: 25 segments; Test set: 50 segments. The results are shown as Figure 2.

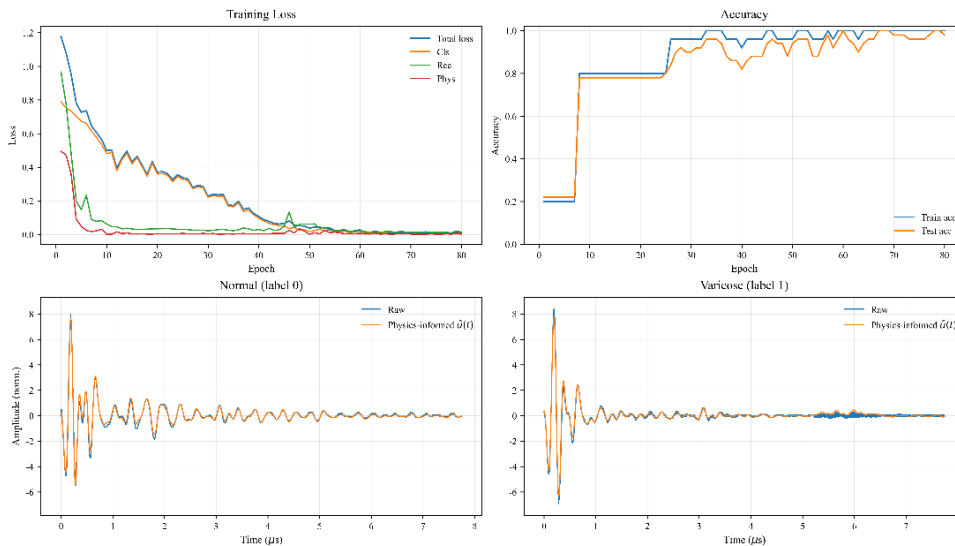


Figure 2: Training loss and accuracy curves of the physics-informed network. The model exhibits stable convergence of the composite loss (classification, reconstruction, and physics-informed terms) and achieves consistently high test accuracy under a stratified split with a fixed test size of 50 samples.

This configuration intentionally prioritizes test reliability over training set size. Given the high variance typically associated with small medical datasets, using a sufficiently large held-out set helps ensure that reported performance is less sensitive to sampling noise.

The physics-informed network exhibited stable convergence during training. The optimization objective consisted of (i) cross-entropy classification loss, (ii) reconstruction loss, and (iii) physics-informed regularization loss. Across epochs, the combined loss decreased smoothly without numerical instability, indicating consistent gradient behavior under the adopted regularization and clipping settings.

The accuracy curves showed the following pattern:

- Training accuracy approached high values early and gradually converged toward saturation.
- Test accuracy increased steadily and remained stable at a high level, without late-stage degradation.

This training behavior is consistent with a model that learns discriminative features while maintaining stable generalization under limited training samples.

Performance was evaluated on the test set of 50 samples using the confusion matrix:

- TP = 39
- TN = 10
- FP = 1
- FN = 0

From these values, the screening metrics were derived:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (6)$$

$$\text{Sensitivity} = \frac{TP}{TP + FN} \quad (7)$$

$$\text{Specificity} = \frac{TN}{TN + FP} \quad (8)$$

Substituting the observed values yields:

- Accuracy ≈ 0.98
- Sensitivity = 1.00
- Specificity ≈ 0.91

From a screening perspective, the absence of false negatives (FN = 0) is clinically important because it indicates that the model did not miss abnormal cases within the evaluated set. The presence of a small number of false positives (FP = 1) is generally more tolerable in screening contexts, as such cases can be referred for confirmatory assessment.

In addition to classification, the framework produces a reconstructed waveform $\hat{u}(t)$ constrained by physics-informed priors. Qualitative inspection of reconstructed A-scan signals indicates the following characteristics:

- Reduced high-frequency fluctuations, consistent with the temporal smoothness constraint.
- Preserved echo structure and dominant reflections, suggesting that reconstruction does not collapse discriminative morphology.
- Stable temporal decay behavior, supporting the role of energy consistency in preventing trivial amplitude attenuation.

These observations imply that the model does not solely rely on arbitrary latent features for classification. Instead, the reconstruction branch encourages the internal representation to remain consistent with physically

plausible signal behavior, which improves interpretability relative to purely data-driven CNN classifiers.

CONCLUSION

This study presents a physics-informed deep learning framework for lightweight ultrasound A-scan-based varicose vein screening. By integrating temporal smoothness and energy consistency constraints into a convolutional encoder–decoder architecture, the method improves robustness and interpretability under small-sample conditions. Experimental evaluation demonstrates high screening accuracy and zero false negatives on a stratified test set of 50 samples. The reconstructed waveforms maintain physical plausibility, supporting the clinical relevance of the learned representations.

Although the dataset size remains limited and participant-level separation is required for further validation, the results indicate the feasibility of portable AI-assisted vascular screening using one-dimensional ultrasound signals. Future work will focus on large-scale clinical validation, subject-independent evaluation, and integration with wearable or handheld ultrasound devices.

ACKNOWLEDGMENT

This work was supported by the Hong Kong Polytechnic University, Angel Fund (Grant number: TSSSUO-26-033).

REFERENCES

- Farah, M. H., Nayfeh, T., Urtecho, M., Hasan, B., Amin, M., Sen, I., Wang, Z., Prokop, L. J., Lawrence, P. F., & Gloviczki, P. (2022). A systematic review supporting the Society for Vascular Surgery, the American Venous Forum, and the American Vein and Lymphatic Society guidelines on the management of varicose veins. *Journal of vascular surgery: venous and lymphatic disorders*, 10(5), 1155–1171.
- Liu, S., Wang, Y., Yang, X., Lei, B., Liu, L., Li, S. X., Ni, D., & Wang, T. (2019). Deep learning in medical ultrasound analysis: A review. *Engineering*, 5(2), 261–275.
- Powers, J., & Kremkau, F. (2011). Medical ultrasound systems. *Interface focus*, 1(4), 477–489.
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378, 686–707.