Anatomy-Texture Aware Safe Gradient Guidance for Source-Free Domain Adaptive Echocardiography Video Segmentation

Jinrong Lv[1], Xun Gong[1] ⋆, Ning Cheng[1], Zhaohuan Li[2,3,4,5], and Weili Jiang[1]

[1] School of Computing and Artificial Intelligence, Southwest Jiaotong University, Chengdu, China
{lvjinrong, xgong, chenning,jiangweili}@swjtu.edu.cn
[2] TCM Hospital of Sichuan Province, Chengdu, China
[3] Sichuan Provincial People's Hospital, Chengdu, China
[4] Chengdu University of TCM, Chengdu, China
[5] University of Electronic Science and Technology, Chengdu, China
lzhhx2006@126.com

Abstract. Source-Free Domain Adaptation (SFDA) has emerged as a crucial technique for deploying medical image segmentation models across domains, particularly when source data is inaccessible due to privacy regulations. While standard SFDA methods based on entropy minimization self-supervised learning (EMSSL) have shown promise in natural images, they can be unreliable in echocardiography due to the modality's inherent speckle noise and artifacts. In this work, we observe that blindly minimizing entropy on noisy ultrasound data can induce optimization instability and increase the risk of catastrophic drift, in which the network becomes overconfident in artifacts rather than anatomical structures. To address this, we propose the Anatomy-Texture Aware Safe Gradient Guidance Source-Free Domain Adaptation (ATASGG-SFDA) framework tailored for robust echocardiography video segmentation. Our method incorporates two clinically grounded mechanisms: (1) an Anatomy-Texture Decomposition strategy that leverages the frozen source model as a coarse anatomical anchor while adapting to target-specific textures; and (2) a Safe Gradient Guidance mechanism that acts as a safety gate, dynamically suppressing EMSSL gradients when they conflict with anatomy-texture pseudo-label supervision learning (ATPSL). This encourages adaptation to be guided by structural fidelity rather than noise overfitting. Experiments on three multi-center echocardiography datasets with significant domain shifts (cross-device and cross-population) demonstrate that ATASGG-SFDA reduces this instability and improves over compared SFDA baselines. The code is available at `https://github.com/lvmarch/ATASGG-SFDA`.

Keywords: Source-Free Domain Adaptation · Echocardiography Video Segmentation · Gradient Alignment

⋆ Corresponding author.

1 Introduction

Echocardiography is a frequently used modality for cardiovascular disease diagnosis due to its real-time capability, portability, and lack of ionizing radiation [1,6,9,21,22]. Accurate segmentation of the Left Ventricle (LV) is fundamental for quantifying clinical indices such as Ejection Fraction (EF) [5,12,14,23]. However, the deployment of automated segmentation models in clinical practice is severely hindered by the domain shift phenomenon [18,26]. Echocardiography data is highly sensitive to variations in acquisition devices (scanner vendors), operator protocols, and patient demographics (e.g., pediatric vs. adult anatomies) [10,17,19]. A model trained on one hospital's data often degrades significantly when tested on another hospital's data. Furthermore, strict patient privacy regulations (e.g., GDPR, HIPAA) often prohibit the sharing of source domain data, rendering traditional Domain Adaptation (DA) methods infeasible. This creates an urgent need for Source-Free Domain Adaptation (SFDA) [3,11], where the model must adapt to the target domain using only pre-trained weights and unlabeled target data.

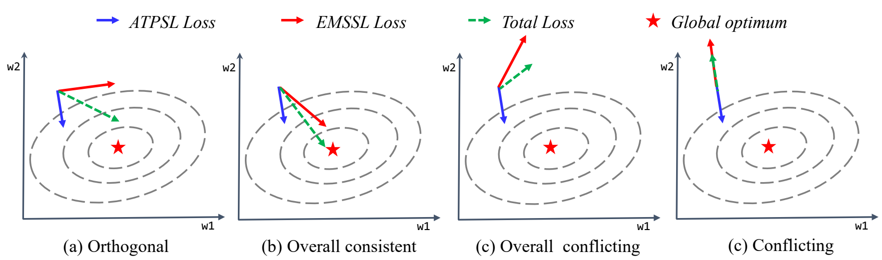

Fig. 1: Simplified gradient update direction schematic. When the gradient update directions of the ATPSL and EMSSL losses are orthogonal or overall consistent, they jointly guide the total update toward a desirable SFDA adaptation direction. Conversely, when these gradients conflict, the total loss gradient can deviate from that direction.

Most SFDA paradigms rely on self-supervised learning (SSL) [2,8,15,16,20,24,25], typically via prediction entropy minimization, to encourage high-confidence predictions [7,27]. While effective for clear natural images, this strategy can be risky for echocardiography. Ultrasound images are characterized by low signal-to-noise ratios, acoustic shadows, and speckle noise [4]. In this context, high-entropy regions often correspond to ambiguous boundaries or severe artifacts. Blindly minimizing entropy without anatomical constraints can force the model to make over-confident predictions on noise, increasing the risk of catastrophic drift or forgetting of the cardiac structure [27]. As illustrated in Fig. 1, this vulnerability arises when the gradient trajectories of the auxiliary entropy minimization self-supervised learning (EMSSL) objective conflict with those from the model's anatomy-texture pseudo-label supervision learning (ATPSL).

To safely mitigate domain shifts without succumbing to noise, it is crucial to distinguish between what to preserve and what to adapt. In echocardiography, the underlying anatomical topology (e.g., the shape of the left ventricle) remains relatively consistent across domains. In contrast, the image appearance (texture, speckle patterns) varies significantly [10, 17, 19]. A reliable SFDA framework should therefore be grounded in anatomical priors to maintain structural integrity while concurrently reconciling domain-specific textural discrepancies. Purely self-training on vanilla pseudo-labels often precipitates degenerative error accumulation due to the inherent noise in the target model's initial hypotheses. Consequently, architecting a robust pseudo-supervisory mechanism that explicitly decouples anatomy from texture constitutes a fundamental desideratum for reliable adaptation.

To tackle these challenges, we propose **Anatomy-Texture Aware Safe Gradient Guidance Source-Free Domain Adaptation (ATASGG-SFDA)** framework, designed specifically to ensure safety and robustness in ultrasound adaptation. First, to establish a reliable adaptation anchor amidst noise, we construct an anatomy-texture pseudo-label supervision learning (ATPSL) branch via a strategy that explicitly decouples robust structural priors from domain-specific texture variations. Second, leveraging this reliable anatomical foundation, our core insight is to treat the geometric relationship between gradients as a safety signal. We hypothesize that valid adaptation signals (from SSL) should align with, or at least not contradict, the fundamental anatomical features preserved by our decomposition strategy. Consequently, when the gradient from EMSSL conflicts with that from ATPSL, it indicates that EMSSL is likely overfitting to domain-specific artifacts, and our mechanism actively mitigates this risk.

Our contributions are summarized as follows:

- **Anatomy-Texture Decomposition Strategy**: A pseudo-labeling strategy that decouples coarse anatomical priors from domain-specific textures via a frozen source anchor and a momentum-driven adapter.
- **Safe Gradient Guidance Mechanism**: A novel mechanism that prevents model collapse by filtering EMSSL updates that conflict with reliable anatomical supervision, ensuring gradient-level optimization stability.
- **Strong SFDA Performance**: Comprehensive validation across cross-vendor and cross-population clinical shifts demonstrates that ATASGG-SFDA outperforms existing SFDA methods in boundary adherence and robustness.

2 Method

2.1 Problem Formulation and Clinical Motivation

In Source-Free Domain Adaptation (SFDA), we are given a segmentation model f_{θ_s} pre-trained on a source domain $\mathcal{D}_s$ (e.g., adult echocardiograms). The goal is to adapt this model to an unlabeled target domain $\mathcal{D}_t = \{x_t\}_{t=1}^{N}$ (e.g., pediatric or cross-vendor data) without accessing $\mathcal{D}_s$. The adapted model is denoted as f_{θ_t}, initialized with θ_s.

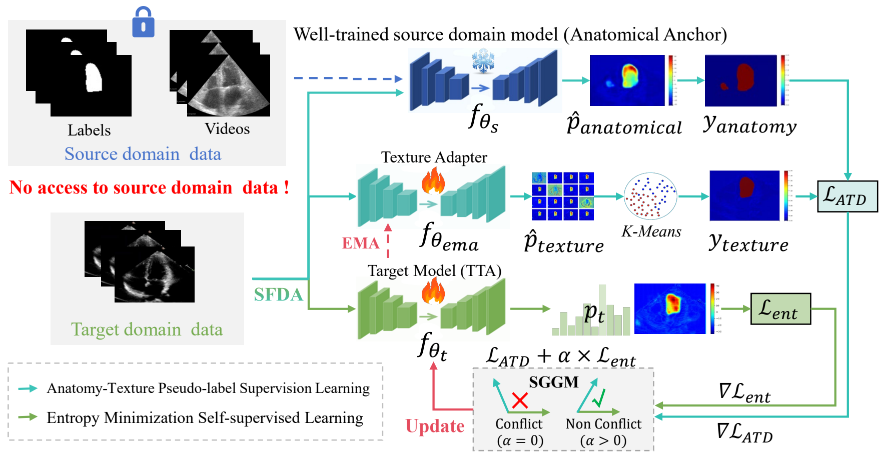

Fig. 2: **Overview of the proposed ATASGG-SFDA framework.** The framework encourages safe adaptation through two mechanisms: (1) Anatomy-Texture Decomposition Strategy (ATDS) uses a frozen Source Model as an Anatomical Anchor (f_{θ_s}) to preserve coarse structural priors ($y_{anatomy}$), while a Texture Adapter ($f_{\theta_{ema}}$) captures target-specific statistics ($y_{texture}$) via K-Means. (2) Safe Gradient Guidance Mechanism (SGGM) acts as a dynamic gate. It computes the geometric relationship between the entropy gradient ($\nabla\mathcal{L}_{ent}$) and the anatomical gradient ($\nabla\mathcal{L}_{ATD}$). If a conflict (obtuse angle) is detected, the safety gate (α) suppresses the entropy signal to reduce catastrophic drift risk.

The Echocardiography Dilemma. Unlike natural images, ultrasound video segmentation faces a unique challenge: the *anatomical structure* (e.g., Left Ventricle geometry) remains topologically consistent, while the *image appearance* (speckle noise pattern, contrast) varies drastically. Standard SFDA methods often rely on prediction entropy minimization as a self-supervised signal:

$$\mathcal{L}_{ent} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} p_t^{(i,c)} \log p_t^{(i,c)}, \tag{1}$$

where N is the number of pixels, C is the number of classes, and $p_t^{(i,c)}$ is the predicted probability derived from $\sigma(f_{\theta_t}(x_t))$. However, in echocardiography, high-entropy regions often correspond to acoustic shadows or artifacts rather than true boundaries. Unconstrained minimization of $\mathcal{L}_{ent}$ can encourage the model to overfit these artifacts, leading to unstable adaptation and possible catastrophic drift.

To address this, we propose **ATASGG-SFDA** (illustrated in Fig. 2), which enforces anatomical constraints through ATDS and ensures robust adaptation via an SGGM.

2.2 Anatomy-Texture Decomposition Strategy (ATDS)

To reduce catastrophic forgetting of cardiac structures, we decompose the pseudo-label supervision into two complementary components: an *Anatomical Anchor* (preserving coarse anatomy shape) and a *Texture Adapter* (adapting to target-domain appearance).

Anatomical Anchor (Domain-Invariant) The source model f_{θ_s} has learned a prior for cardiac shape. Although its pixel-wise accuracy may drop on the target domain, its predictions can still provide useful coarse LV structural cues. We freeze θ_s to generate conservative pseudo-labels $y_{anatomy} = \mathbb{I}(\sigma(f_{\theta_s}(x_t)) > 0.5)$. We employ the Dice loss to penalize structural deviations:

$$\mathcal{L}_{anatomy} = 1 - \frac{2 \sum p_t \cdot y_{anatomy}}{\sum p_t + \sum y_{anatomy}}. \tag{2}$$

Texture Adapter (Domain-Variant) To adapt to the target-specific intensity distribution, we utilize a Mean-Teacher approach. We maintain a momentum-updated model $f_{\theta_{ema}}$ with parameters θ_{ema}, updated via $\theta_{ema} \leftarrow \mu\theta_{ema} + (1-\mu)\theta_t$ at each step, where μ is a momentum coefficient. Instead of clustering one-dimensional probabilities, we apply K-Means clustering ($K = 2$) on the deep feature maps of the teacher model, denoted as $f^{feature}_{\theta_{ema}}(x_t)$, to dynamically separate the LV cavity from the background based on the target contrast. Here, $K = 2$ follows the binary LV/background setting rather than an arbitrary clustering choice:

$$y_{texture} = K - Means(f^{feature}_{\theta_{ema}}(x_t)). \tag{3}$$

Because $f_{\theta_{ema}}$ inherits its initialization from the source model and is updated by momentum, its features retain semantic cues during target-domain exploration, which helps mitigate semantic flipping. This generates a texture-aware pseudo-label, yielding the adaptation loss $\mathcal{L}_{texture}$ (computed similarly via Dice loss). The total supervised loss acts as our anatomical prior, balanced by a hyperparameter λ:

$$\mathcal{L}_{ATD} = \lambda\mathcal{L}_{anatomy} + (1-\lambda)\mathcal{L}_{texture}. \tag{4}$$

2.3 Safe Gradient Guidance Mechanism (SGGM)

While $\mathcal{L}_{ATD}$ provides stability, we leverage $\mathcal{L}_{ent}$ to enhance decision confidence. We hypothesize that a valid SSL gradient should lie within a Safe Subspace defined by the anatomical prior. Specifically, the gradient from entropy minimization ($g_{ent} = \nabla_{\theta_t}\mathcal{L}_{ent}$) should not contradict the update direction of the anatomical decomposition ($g_{ATD} = \nabla_{\theta_t}\mathcal{L}_{ATD}$).

Geometric Conflict Detection. We compute the cosine similarity between the gradients:

$$\cos(g_{ATD}, g_{ent}) = \frac{\langle g_{ATD}, g_{ent} \rangle}{\|g_{ATD}\|\|g_{ent}\| + \epsilon}, \tag{5}$$

where ϵ is a small constant for numerical stability. **Safety Violation ($\cos < 0$):** Gradients are obtuse. Minimizing entropy pushes the model *away* from the anatomical prior (e.g., overfitting to artifacts). **Safe Exploration ($\cos \geq 0$):** Gradients are consistent. Entropy minimization refines the boundary within anatomical constraints.

Dynamic Safety Gating. We introduce a dynamic gate α to filter conflicting components. The weight is formulated as:

$$\alpha = \min\left(1, ReLU\left(\frac{\langle g_{ATD}, g_{ent}\rangle}{\|g_{ent}\|^2 + \epsilon} \cdot \beta\right)\right). \tag{6}$$

Here, the term inside the ReLU represents the projection of anatomical gradients onto the entropy gradient direction, scaled by a factor β and clipped to $[0, 1]$ to control the magnitude of the SSL contribution. Crucially, if $\langle g_{ATD}, g_{ent}\rangle < 0$, α becomes 0, completely detaching the unsafe SSL branch. The final optimization objective is:

$$\min_{\theta_t} \left(\mathcal{L}_{ATD} + \alpha \cdot \mathcal{L}_{ent}\right). \tag{7}$$

Proposition 1 (Gradient Safety Condition). *Let $\mathcal{L}_{ATD}(\theta)$ be the anatomical loss function. An update step driven by the joint objective with entropy minimization is conservative (i.e., non-destructive) w.r.t $\mathcal{L}_{ATD}$ under the SGGM. The mechanism ensures that the SSL update locally strictly does not increase the anatomical loss.*

Justification. With the parameter update $\Delta\theta = -\eta(\nabla\mathcal{L}_{ATD} + \alpha\nabla\mathcal{L}_{ent})$ and learning rate $\eta \to 0$, the first-order Taylor approximation of the anatomical loss change is:

$$\Delta\mathcal{L}_{ATD} \approx \langle\nabla\mathcal{L}_{ATD}, \Delta\theta\rangle = \underbrace{-\eta\|\nabla\mathcal{L}_{ATD}\|^2}_{Standard\ Descent} \underbrace{-\eta\alpha\langle\nabla\mathcal{L}_{ATD}, \nabla\mathcal{L}_{ent}\rangle}_{SSL\ Interference}. \tag{8}$$

SGGM strictly restricts the interference term to be non-positive: if gradients conflict ($\langle\nabla\mathcal{L}_{ATD}, \nabla\mathcal{L}_{ent}\rangle < 0$), it assigns $\alpha = 0$; if synergistic, $\alpha \geq 0$. Consequently, $\Delta\mathcal{L}_{ATD} \leq -\eta\|\nabla\mathcal{L}_{ATD}\|^2 \leq 0$. The SSL signal either safely accelerates optimization or halts, firmly bounding the update within the anatomical Safe Half-space.

3 Experiments

3.1 Datasets and Evaluation Metrics

Datasets. We utilize three public echocardiography datasets: **CAMUS** (Source, 500 cases, Adult, GE Vivid E95 Device) [10], **EchoNet-Dynamic** (Target 1, 10030 cases, Adult, Multi-Vendor Device) [17], and **Echo-Pediatric** (Target 2, 3176 cases, Pediatric, Multi-Vendor Device) [19]. This setup allows us to assess robustness against two critical domain shifts: *equipment heterogeneity* (Source $\rightarrow$ Target 1) and *anatomical/population variation* (Source $\rightarrow$ Target 2).

Evaluation Metrics. We employ four standard metrics: Dice Similarity Coefficient (DSC), Intersection over Union (IoU) for region overlap; and Average Symmetric Surface Distance (ASSD), 95% Hausdorff Distance (HD95) for boundary precision. Given the clinical importance of accurate LV boundaries for functional quantification, we place particular emphasis on HD95 and ASSD. We report the mean and standard deviation across the test sets.

Table 1: Baseline performance of the pre-trained source domain model.

Model	Dataset	DSC (%)↑	IoU (%)↑	ASSD (px)↓	HD95 (px)↓
3D UNet [14]	CAMUS [10]	94.03 ±2.28	88.82 ±3.94	0.07 ±0.03	0.57 ±0.50

3.2 Implementation Details

All experiments were conducted on an NVIDIA RTX 4090 GPU using PyTorch. The 3D UNet source model [14] was pre-trained on CAMUS using standard supervised learning. As shown in Table 1, the model achieves high accuracy (DSC 94.03%) on the source domain, providing a reliable anatomical anchor for subsequent adaptation. For SFDA, we use the AdamW optimizer with a learning rate of 3e-4. The adaptation runs for one epoch over the target data. The key hyperparameters are empirically set to: $\lambda = 0.8$ (weight for anatomical anchor), $\beta = 1.5$ (scaling factor for SGGM). Data preprocessing and other implementation specifics adhering to the protocols established in [14].

Table 2: Quantitative comparison of state-of-the-art methods.

Method	CAMUS → EchoNet-Pediatric (Cross-Population)				CAMUS → EchoNet-Dynamic (Cross-Vendor)			
	DSC (%)↑	IoU (%)↑	ASSD (px)↓	HD95 (px)↓	DSC (%)↑	IoU (%)↑	ASSD (px)↓	HD95 (px)↓
Direct Test	73.77 ±17.44	61.00 ±18.86	2.24 ±3.30	10.63 ±11.37	75.63 ±12.64	62.32 ±15.03	1.57 ±2.09	8.79 ±10.20
UPL [13]	76.91 ±12.28	63.95 ±14.83	1.39 ±1.97	7.99 ±9.64	77.08 ±14.53	64.69 ±17.00	1.94 ±2.97	10.00 ±11.86
IAPC [3]	78.16 ±11.56	65.49 ±14.26	1.29 ±1.88	7.67 ±9.58	81.94 ±10.38	70.53 ±12.96	0.76 ±1.31	4.45 ±4.90
AIF [11]	80.66 ±11.16	68.88 ±13.95	0.85 ±1.47	4.82 ±5.06	81.24 ±11.45	69.79 ±14.40	0.88 ±1.51	4.95 ±5.45
POEM [2]	81.81 ±10.68	70.45 ±13.65	0.78 ±1.29	4.53 ±4.81	82.01 ±10.49	70.67 ±13.27	0.80 ±1.50	4.66 ±5.67
Ours	**82.74** ±10.42	**71.72** ±13.26	**0.75** ±1.38	**4.38** ±5.25	**84.63** ±8.85	**74.27** ±11.99	**0.46** ±0.84	**2.88** ±2.94

** Note: SD > Mean in ASSD/HD95 reflects right-skewed hard outliers, not negative distances.*

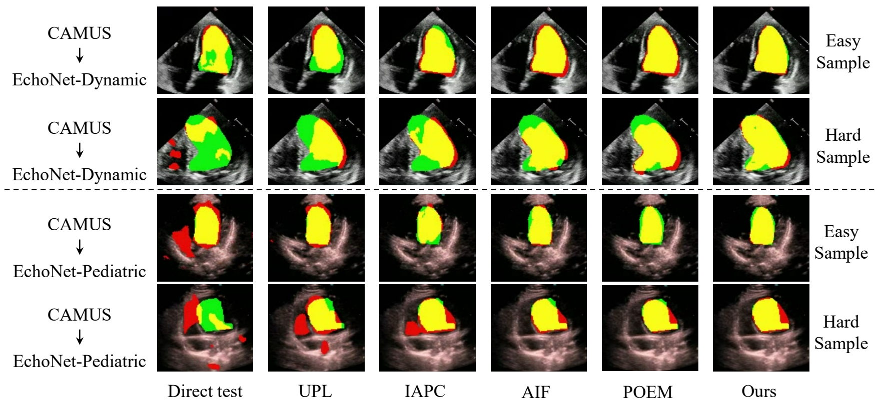

Fig. 3: Qualitative comparison. ATASGG-SFDA better adheres to boundaries.

Table 3: Ablation study on CAMUS → EchoNet-Pediatric.

Method	**DSC (%)↑**	**IoU (%)↑**	**ASSD (px)↓**	**HD95 (px)↓**
Direct Test	73.77 ±17.44	61.00 ±18.86	2.24 ±3.30	10.63 ±11.37
+ ATDS	79.86 ±12.62	68.06 ±15.20	1.12 ±1.94	6.02 ±7.29
w/o $\mathcal{L}_{anatomy}$	79.37 ±13.32	67.50 ±15.59	1.22 ±2.21	6.45 ±7.87
w/o $\mathcal{L}_{texture}$	76.02 ±16.17	63.62 ±17.96	2.10 ±3.38	10.10 ±11.71
ATDS + fixed α	79.87 ±13.22	68.20 ±15.71	1.21 ±2.28	6.30 ±8.18
ATDS + SGGM	**82.74** ±10.42	**71.72** ±13.26	**0.75** ±1.38	**4.38** ±5.25

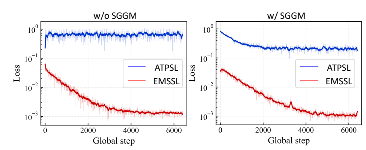

Fig. 4: Loss dynamics. **Left:** Unconstrained EMSSL. **Right:** safe gradient guidance.

3.3 Comparison with State-of-the-art Methods

We compared ATASGG-SFDA with state-of-the-art SFDA methods. The compared methods cover pseudo-labeling (UPL), prototype/calibration-based adaptation (IAPC), information-filtering adaptation (AIF), and protected entropy minimization (POEM). All baseline models were reproduced under identical experimental conditions using their officially released source code. Table 2 reports the results. ATASGG-SFDA improves over the compared methods, particularly in boundary metrics (HD95), supporting the utility of our ATDS and SGGM designs.

3.4 Ablation Study

Table 3 summarizes the ablation study on CAMUS → EchoNet-Pediatric. **Efficacy of ATDS:** Compared to the direct baseline, ATDS improves generalization. Decoupling the stable anatomical anchor ($\mathcal{L}_{anatomy}$) from the flexible texture adapter ($\mathcal{L}_{texture}$) mitigates catastrophic forgetting. **Impact of SGGM:** SGGM suppresses conflicting gradients when detected, reducing the risk of overfitting to noise artifacts and yielding the best overall performance. **Loss Dynamics:** As illustrated in Fig. 4, unconstrained EMSSL causes larger fluctuations in the anatomical ATPSL loss, whereas SGGM dynamically modulates EMSSL and stabilizes convergence.

4 Conclusion

We introduce ATASGG-SFDA, a safety-aware framework that mitigates catastrophic drift risk in echocardiography by decoupling anatomy from texture. Our approach enforces a geometric consistency constraint via SGGM, dynamically filtering noise-driven updates. Two cross-center SFDA evaluations demonstrate robust performance among the compared SFDA methods. Extending ATASGG-SFDA to other complex medical imaging modalities, such as CT and MRI, remains unvalidated and will require task-specific evaluation in future work.

Acknowledgments. This work was partially supported by the National Natural Science Foundation of China (62376231 and 62576296), the Natural Science Foundation of Sichuan Province (2026NSFSC0418), the Science and Technology Plan of Tibet Autonomous Region (XZ202601ZY0087), the Fundamental Research Funds for the Central Universities (2682026CX115), and the Doctoral Graduate Innovation Fund of Southwest Jiaotong University (CX2026YB31).

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

References

1. Anavekar, N.S., Oh, J.K.: Doppler echocardiography: a contemporary review. Journal of cardiology **54**(3), 347–358 (2009)
2. Bar, Y., Shaer, S., Romano, Y.: Protected test-time adaptation via online entropy matching: A betting approach. Advances in Neural Information Processing Systems (NeurIPS) **37**, 85467–85499 (2024)
3. Cao, Y., Zhang, H., Lu, X., Xiao, Z., Yang, K., Wang, Y.: Towards source-free domain adaptive semantic segmentation via importance-aware and prototype-contrast learning. IEEE Transactions on Intelligent Vehicles pp. 1–13 (2024)
4. Christensen, M., Vukadinovic, M., Yuan, N., Ouyang, D.: Vision–language foundation model for echocardiogram interpretation. Nature Medicine pp. 1–8 (2024)
5. Deng, X., Wu, H., Zeng, R., Qin, J.: Memsam: Taming segment anything model for echocardiography video segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9622–9631 (2024)
6. Feigenbaum, H.: Clinical applications of echocardiography. Progress in Cardiovascular Diseases **14**(6), 531–558 (1972)
7. Gong, T., Kim, Y., Lee, T., Chottananurak, S., Lee, S.J.: Sotta: Robust test-time adaptation on noisy data streams. Advances in Neural Information Processing Systems (NeurIPS) **36**, 14070–14093 (2023)
8. Haeusser, P., Frerix, T., Mordvintsev, A., Cremers, D.: Associative domain adaptation. In: Proceedings of the IEEE international conference on computer vision (ICCV). pp. 2765–2773 (2017)

9. Landgren, M., Overgaard, N.C., Heyden, A.: Segmentation of the left heart ventricle in ultrasound images using a region based snake. In: Medical Imaging 2013: Image Processing. vol. 8669, pp. 1132–1140. SPIE (2013)
10. Leclerc, S., Smistad, E., Pedrosa, J., Ostvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., Lartizien, C., Dhooge, J., Lovstakken, L., Bernard, O.: Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. IEEE Transactions on Medical Imaging p. 2198–2210 (Sep 2019)
11. Li, H., Li, H., Chen, J., Zhong, R., Niu, K., Fu, H., Liu, J.: Aif-sfda: Autonomous information filter driven source-free domain adaptation for medical image segmentation. Proceedings of the AAAI Conference on Artificial Intelligence (AAAI) **39**(5), 4716–4724 (Apr 2025)
12. Lin, J., Xie, W., Kang, L., Wu, H.: Dynamic-guided spatiotemporal attention for echocardiography video segmentation. IEEE Transactions on Medical Imaging **43**(11), 3843–3855 (2024)
13. Lu, Z., Li, D., Song, Y.Z., Xiang, T., Hospedales, T.M.: Uncertainty-aware source-free domain adaptive semantic segmentation. IEEE Transactions on Image Processing **32**, 4664–4676 (2023)
14. Maani, F., Ukaye, A., Saadi, N., Saeed, N., Yaqub, M.: Simlvseg: Simplifying left ventricular segmentation in 2-d+time echocardiograms with self- and weakly supervised learning. Ultrasound in Medicine Biology **50**(12), 1945–1954 (2024)
15. Niu, S., Wu, J., Zhang, Y., Chen, Y., Zheng, S., Zhao, P., Tan, M.: Efficient test-time model adaptation without forgetting. In: International conference on machine learning (ICLR). pp. 16888–16905. PMLR (2022)
16. Niu, S., Wu, J., Zhang, Y., Wen, Z., Chen, Y., Zhao, P., Tan, M.: Towards stable test-time adaptation in dynamic wild world. In: Internetional Conference on Learning Representations (ICLR) (2023)
17. Ouyang, D., He, B., Ghorbani, A., Yuan, N., Ebinger, J., Langlotz, C.P., Heidenreich, P.A., Harrington, R.A., Liang, D.H., Ashley, E.A., et al.: Video-based ai for beat-to-beat assessment of cardiac function. Nature **580**(7802), 252–256 (2020)
18. Reddy, C.D., Lopez, L., Ouyang, D., Zou, J.Y., He, B.: Video-based deep learning for automated assessment of left ventricular ejection fraction in pediatric patients. Journal of the American Society of Echocardiography **36**(5), 482–489 (2023)
19. Reddy, C.D., Lopez, L., Ouyang, D., Zou, J.Y., He, B.: Video-based deep learning for automated assessment of left ventricular ejection fraction in pediatric patients. Journal of the American Society of Echocardiography **36**(5), 482–489 (2023)
20. Saito, K., Ushiku, Y., Harada, T.: Asymmetric tri-training for unsupervised domain adaptation. In: International conference on machine learning (ICML). pp. 2988–2997. PMLR (2017)
21. Skalski, A., Turcza, P., Zieliński, T., Królczyk, J., Grodzicki, T.: Left ventricle usg image segmentation using active contour model. Procedia computer science **1**(1), 2723–2732 (2010)
22. Smistad, E., Ostvik, A., Haugen, B.O., Lovstakken, L.: 2d left ventricle segmentation using deep learning. In: 2017 IEEE International Ultrasonics Symposium (IUS) (Sep 2017)
23. Thomas, S., Gilbert, A., Ben-Yosef, G.: Light-weight spatio-temporal graphs for segmentation and ejection fraction prediction in cardiac ultrasound. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 380–390. Springer (2022)

24. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (December 2015)
25. Wang, D., Shelhamer, E., Liu, S., Olshausen, B.A., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization. In: International Conference on Learning Representations (ICLR) (2021)
26. Wu, H., Liu, J., Xiao, F., Wen, Z., Cheng, L., Qin, J.: Semi-supervised segmentation of echocardiography videos via noise-resilient spatiotemporal semantic calibration and fusion. Medical Image Analysis **78**, 102397 (2022)
27. Zhao, H., Liu, Y., Alahi, A., Lin, T.: On pitfalls of test-time adaptation. In: International Conference on Machine Learning (ICML) (2023)