# Ultrafast Echocardiography Based on Dynamic Subspace-Guided and Temporal Adaptive Diffusion

Jingfeng Lu[1], Fabien Millioz[2], Jonathan Porée[3], Jean Provost[3], and Yi Zhang[1]*

[1] School of Cyber Science and Engineering, Sichuan University, Chengdu, China
[2] CREATIS, CNRS, Inserm, INSA-Lyon, University of Lyon 1, Villeurbanne, France
[3] Engineering Physics Department, Polytechnique Montréal, Montréal, Canada

**Abstract.** Ultrafast echocardiography based on diverging wave (DW) imaging enables extremely high frame rates, unlocking breakthrough insights into cardiac mechanics. However, high-quality DW cardiac imaging typically relies on motion compensation (MoCo) compounding from multiple transmissions, which leads to a fundamental trade-off between image quality and frame rate. In this study, we introduce a novel SpatioTemporal COnstrained Diffusion Echocardiography (ST-CODE) for high-quality cardiac imaging at high frame rate using only three-angle steered DW transmissions. ST-CODE formulates reconstruction as a dynamic subspace-guided and temporally adaptive diffusion process that embeds structural and temporal constraints into reverse denoising. A dynamic subspace guidance module is designed to capture sequence-level anatomical structure and enforce structural consistency across frames. In parallel, a temporal correlation sampling strategy is introduced to regulate the diffusion trajectory to propagate inter-frame dependency and suppress stochastic variations during inference. Extensive experiments conducted on *in vivo* cardiac data demonstrate that ST-CODE achieves image quality comparable to 32-angle MoCo-compounding reference and consistently outperforms existing reconstruction methods, leading to anatomically consistent and temporally coherent reconstructions without sacrificing ultrafast frame rates.

**Keywords:** Ultrafast echocardiography · Diverging wave · Diffusion reconstruction · Dynamic subspace · Temporal correlation.

## 1 Introduction

Echocardiography is a widely used non-invasive cardiac imaging modality due to its rapid applicability and low cost, and plays a key role in various clinical applications such as diagnosis, treatment planning, and patient monitoring [21]. Ultrafast echocardiography based on diverging wave (DW) imaging has gained considerable attention in recent years owing to its exceptionally high imaging

* Corresponding author: yzhang@scu.edu.cn

temporal resolution. Unlike conventional focused ultrasound, which scans the region of interest using successive narrow beams, ultrafast imaging insonifies a wide field of view with a single emission of DW. This approach enables frame rates up to several thousand frames per second, which unlocks the analysis of rapidly changing physical phenomena in the cardiovascular system such as myocardial strain [24] and natural mechanical waves [9].

Images obtained with ultrafast acquisition exhibit low image quality. A standard approach to improve the quality of ultrafast echocardiography is coherent diverging wave compounding (CDWC)[12], where consecutive steered transmit beams at different angles are transmitted and the backscattered echoes are coherently summed. In cardiac imaging, which involves fast-moving tissues, motion artifacts may ensue during a large number of transmit-receive events, which in turn degrades the quality of compounded images. Motion compensation (MoCo) methods have thus been introduced to provide high-quality echocardiography [15,10], while preserving the speckle patterns for carrying out efficient motion analysis [14]. However, compounding-based echocardiography relies on an increased number of ultrasound transmissions, which consequently reduces the achievable frame rate. This constraint leads to a fundamental trade-off between image quality and temporal resolution. Therefore, exploiting the full potential of ultrafast cardiac imaging in both frame rate and image quality remains an active area of research.

A wide variety of deep learning-based approaches have been proposed to improve the quality of ultrasound imaging in recent years. A major focus of these studies has been the reconstruction of high-quality images by learning a direct mapping from measurements to images [5,6,13,7,25]. However, such learning strategies rely on strong spatial correspondence assumptions, which are particularly challenging in cardiac imaging due to complex motion and acquisition variability [8]. Recently, deep generative models [19,22,4,23] have emerged as an alternative to direct reconstruction strategies by learning the underlying data distribution beyond pointwise measurement-to-image mapping. In particular, diffusion-based models [3,17] have attracted increasing attention due to their strong generative capability enabled by the well-conditioned decomposition of complex data distributions in a stepwise and traceable process. In the field of echocardiography, diffusion-based studies have focused on image synthesis tasks [18,20,16]. However, these methods primarily emphasize clinically realistic image generation with inherent stochasticity, while neglecting anatomical accuracy due to the lack of structural constraints. Moreover, in echocardiographic sequences, such limitations may impair inter-frame coherence, which is essential for reliable downstream analysis and diagnosis.

In this study, we consider both spatial accuracy and temporal coherence required for reliable cardiac imaging, and propose a SpatioTemporal COnstrained Diffusion Echocardiography (ST-CODE) for the high-quality reconstruction of ultrafast cardiac sequences. ST-CODE formulates reconstruction as a diffusion process that progressively recovers high-quality ultrasound frames through reverse denoising. To preserve anatomical consistency, we introduce a dynamic

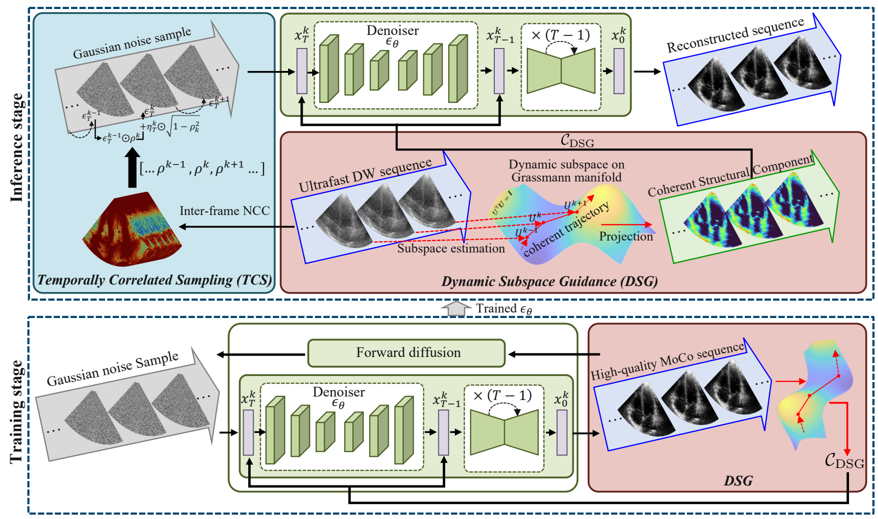


**Fig. 1.** Overview of our proposed ST-CODE framework.

subspace guidance (DSG) module that captures the sequence-level anatomy as a low-dimensional subspace on the Grassmann manifold, whose basis-invariant property provides a spatiotemporal coherent structural condition to guide the reverse diffusion process toward anatomically consistent reconstruction. Furthermore, we design a temporal correlation sampling (TCS) strategy to enhance temporal coherence in the reconstructed sequence. TCS introduces spatially adaptive inter-frame correlation at the initial noise sampling stage and enforces structured temporal dependency along the reverse denoising trajectory, thereby preserving consistent motion patterns against stochastic variations across frames. Extensive experiments conducted on *in vivo* cardiac data demonstrate the superiority of the proposed ST-CODE framework over existing reconstruction methods, validating its effectiveness in preserving anatomical consistency and enhancing temporal coherence in realistic cardiac imaging scenarios.

## 2 Method

### 2.1 Framework Overview

Figure 1 illustrates the proposed overall framework. The framework is formulated as a diffusion process, in which high-quality ultrasound frames are progressively reconstructed through a reverse denoising procedure. To ensure anatomical consistency in the reconstruction, the reverse denoising process is conditioned on a **dynamic subspace guidance (DSG)**, which encodes the dominant spatiotemporal anatomy at the sequence level. In addition, a **temporal correlation sampling (TCS)** strategy is introduced during inference to enforce inter-frame coherence across the reconstructed sequence, without altering the standard diffusion training procedure.

Specifically, given a frame $x$ from the high-quality target domain, the forward diffusion process gradually adds Gaussian noise according to a predefined variance schedule $\{\beta_t\}_{t=1}^{T}$. Denoting $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$, the noisy frame at timestep $t$ is

$$x_t = \sqrt{\bar{\alpha}_t}\, x_0 + \sqrt{1 - \bar{\alpha}_t}\, \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \tag{1}$$

where $x_0 = x$ denotes the original high-quality frame and $x_t$ represents its progressively noised version at timestep $t$. The reverse diffusion process is modeled as a conditional denoising Markov chain:

$$p_\theta(x_{t-1} \mid x_t, \mathcal{C}_{\mathrm{DSG}}) = \mathcal{N}\left(x_{t-1}; \mu_\theta(x_t, t, \mathcal{C}_{\mathrm{DSG}}), \mathbf{I}\right), \tag{2}$$

where $\mathcal{C}_{\mathrm{DSG}}$ denotes the external condition from the DSG module. The predicted mean $\mu_\theta$ is reparameterized using a noise prediction network $\epsilon_\theta(x_t, t, \mathcal{C}_{\mathrm{DSG}})$, and the model is trained by minimizing the conditional noise prediction objective:

$$\mathcal{L}_\epsilon = \mathbb{E}_{x_0, \epsilon, t}\left[\|\epsilon - \epsilon_\theta(x_t, t, \mathcal{C}_{\mathrm{DSG}})\|_2^2\right]. \tag{3}$$

While the learning process operates at the frame level, the framework leverages sequence-level structural guidance from DSG and temporal dependency introduced by the TCS strategy, which are detailed in the following sections.

### 2.2 Dynamic Subspace Guidance

Ultrafast echocardiography sequences are spatiotemporal data in which cardiac anatomy evolves continuously over time. Considering that the high-dimensional ultrafast sequence exhibits coherent structures across consecutive frames, the underlying anatomical structures can be captured by a subspace represented by multiple orthonormal bases, highlighting that the representation is invariant to the particular basis choice. In order to capture such invariance and ensure spatiotemporal coherence, we model the subspace on the Grassmann manifold, where each point corresponds to an equivalence class of orthonormal bases spanning the same subspace, inherently reflecting basis-invariance.

Let $\mathbf{X} = \{x^k\}_{k=1}^{K}$ denote the cardiac sequence, where $x^k \in \mathbb{R}^{H \times W}$ denotes the $k$-th frame, which is vectorized into a column vector in $\mathbb{R}^{HW}$, where $H$ and $W$ denote the image height and width, respectively. And let $\mathbf{U}^k \in \mathbb{R}^{(HW) \times m}$ be the subspace basis satisfying $(\mathbf{U}^k)^{\mathrm{T}} \mathbf{U}^k = I$, where $m$ is the subspace dimension. We then update the subspace along directions in the tangent space at $\mathbf{U}^k$, with the tangent direction $\boldsymbol{\Delta}^k$ constrained by $(\mathbf{U}^k)^{\mathrm{T}} \boldsymbol{\Delta}^k = 0$, which ensures that the updated subspace stays on the Grassmann manifold. The updated subspace is obtained via a retraction operation:

$$\mathbf{U}^{k+1} = \mathcal{P}_{\mathrm{Gr}}(\mathbf{U}^k + \lambda \boldsymbol{\Delta}^k), \tag{4}$$

where $\lambda$ denotes the step size and $\mathcal{P}_{\mathrm{Gr}}(\cdot)$ represents projection back onto the Grassmann manifold, implemented using a thin QR decomposition.

This tangent-space update gradually adjusts the subspace along the temporal dimension in a basis-invariant manner, which captures consecutive anatomical structures over time. In this context, the anatomical evolvement corresponds to the trajectories on the Grassmann manifold, whereas noise-induced perturbations tend to lie off the manifold. The resulting Grassmannian subspace representation encodes coherent spatiotemporal structure, serving as a robust structural prior for the reconstruction process.

The structural component is then obtained by projecting each frame $x^k$ onto its corresponding subspace spanned by $\mathbf{U}^k$, for $k = 1, \ldots, K$:

$$\mathcal{C}_{\mathrm{DSG}}^k = \mathbf{U}^k \big( (\mathbf{U}^k)^{\mathrm{T}} x^k \big), \tag{5}$$

The resulting $\{\mathcal{C}_{\mathrm{DSG}}^k\}_{k=1}^K$ captures the dominant anatomical structure, which then serves as the structural conditions in the reverse diffusion process to enforce spatiotemporal consistency.

### 2.3 Temporal Correlation Sampling

In standard diffusion-based inference, the initial noises of different frames are sampled independently. To enforce temporal coherence across reconstructed frames, we introduce correlation at the initial noise sampling step during inference. Specifically, given two consecutive frames $x^{k-1}$ and $x^k$, we compute a patch-wise correlation map $\hat{\rho}_k = \mathrm{NCC}(x^{k-1}, x^k)$ using a normalized cross-correlation (NCC) operator, and interpolate it to the image resolution via bilinear interpolation as $\rho_k \in [0,1]^{H \times W}$ to ensure the spatial alignment with the diffusion noise. The initial noise $\epsilon_T^k$ for the $k$-th frame is then sampled as

$$\epsilon_T^k = \rho_k \odot \epsilon_T^{k-1} + \sqrt{1 - \rho_k^2} \odot \eta_T^k, \quad \eta_T^k \sim \mathcal{N}(0, I), \tag{6}$$

where $\epsilon_T^{k-1}$ is the initial sampled noise for the $(k-1)$-th frame and $\odot$ denotes element-wise multiplication. This formulation preserves the marginal distribution of the noise at each spatial location while introducing cross-frame covariance at the noise level.

During the reverse diffusion process, the final reconstructed frame can be viewed as the mapping $x_0^k = f_\theta(\epsilon_T^k)$ from the initial noise, where $f_\theta(\cdot)$ denotes the reverse diffusion operator induced by the trained diffusion model. Under the first-order approximation that assumes the model mapping is locally linear with respect to the noise [2], the cross-covariance between consecutive output frames is propagated from the noise domain through the model mapping and satisfies

$$\mathrm{Cov}(x_0^k, x_0^{k-1}) = J_k \, \mathrm{Cov}(\epsilon_T^k, \epsilon_T^{k-1}) \, J_k^{\mathrm{T}}, \tag{7}$$

where $J_k = \frac{\partial f_\theta}{\partial \epsilon_T}\Big|_{\epsilon_T^k}$ denotes the Jacobian of the reverse diffusion mapping. Notably, this strategy is applied only during the inference stage and does not require additional processing in the training stage. The temporal correlation is captured at the initial noise sampling step and propagated through the reverse diffusion mapping, which is then adjusted according to the learned spatial structures to ensure both temporal coherence and spatial accuracy.

## 3 Experiments

### 3.1 Experimental Setup

**Data Acquisition.** The study was performed with the data acquired using a Verasonics scanner and a phased array probe (ATL P4-2; bandwidth, 2-4 MHz; center frequency, 3 MHz; pitch, 0.3 mm; kerf, 50 $\mu$m). Apical four-chamber views were examined from six informed and consenting normal subjects, with subject-independent data splits to ensure no overlap between training and testing. We performed two types of DW transmissions: i) 32-angle transmission to obtain references using coherent-weighted MoCo (CW-MoCo) [10], ii) 3-angle transmissions as the model input. In particular, for the test dataset, the two transmission types were interleaved to minimize the time interval between the corresponding reconstructions and make them comparable for quantitative evaluation.
**Learning Configuration.** Our model was implemented with PyTorch on a NVIDIA V100 GPU. The model was trained for 400 epochs using the Adam optimizer. The initial learning rate was set to $5 \times 10^{-5}$ and was fixed in the first 200 epochs, and linearly decayed to 0 in the final 200 epochs. In the training stage, similar to original DDPM [3], we set the number of diffusion steps to $T = 1000$. For the variance schedule in the forward process, we employed a linear schedule from $\beta_1 = 0.0001$ to $\beta_T = 0.02$. In the inference stage, we evenly sampled 100 steps from 1 to $T$ and then performed generation only on these 100 steps by employing the inference acceleration technique [11]. While such multistep diffusion inference introduced higher computational costs and runtime than conventional feed-forward models, it operated in the post-acquisition stage separate from ultrafast acquisition enabled by 3-angle DW transmissions.
**Model Evaluation.** We evaluated the spatial quality using peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), and normalized mutual information (NMI) using the MoCo reconstruction as the reference. Besides, we evaluated the temporal coherence by measuring endpoint error (EPE) and mean endpoint error (MEPE) on the inter-frame motion estimated with the reconstructed sequences. We compared our method with several deep learning-based ultrasound reconstruction models, including UNet [13], AUGAN [19], CCGR [25] based on a recurrent neural network (RNN) that takes multiple frames as input, and a DDPM model [1]. As CCGR required multi-frame input, we fed the 3-angle frames instead of their compoundings as the model input.

### 3.2 Experimental Results

**Image quality.** Fig. 2 presents a representative example of the reconstructed images obtained using different methods. Compared with the conventional CDWC

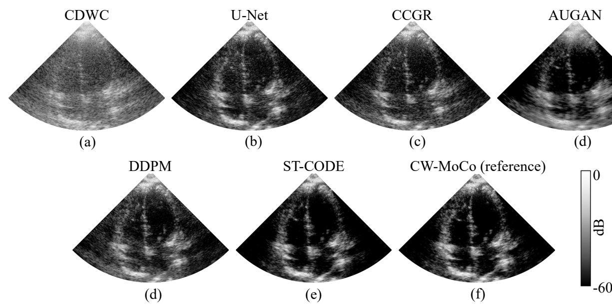


**Fig. 2.** Reconstructed images from a representative case. (a)-(e) Reconstructions obtained by different methods using 3 DWs. (f) CW-MoCo reconstruction using 32 DWs, serving as the reference.

**Table 1.** Evaluation metrics achieved by different approaches on the test set.

| Model | PSNR [dB](↑) | SSIM(↑) | NMI(↑) | MEPE [mm](↓) |
|---|---|---|---|---|
| CDWC | $25.97 \pm 0.39$ | $0.82 \pm 0.01$ | $0.31 \pm 0.01$ | $0.15 \pm 0.10$ |
| U-Net | $27.61 \pm 0.50$ | $0.85 \pm 0.02$ | $0.34 \pm 0.02$ | $0.17 \pm 0.14$ |
| AUGAN | $27.49 \pm 0.50$ | $0.87 \pm 0.01$ | $0.36 \pm 0.01$ | $0.14 \pm 0.06$ |
| CCGR | $28.28 \pm 0.55$ | $0.87 \pm 0.01$ | $0.37 \pm 0.01$ | $0.11 \pm 0.08$ |
| DDPM | $28.23 \pm 0.49$ | $0.86 \pm 0.01$ | $0.33 \pm 0.02$ | $0.16 \pm 0.13$ |
| ST-CODE | $\mathbf{30.44 \pm 0.47}$ | $\mathbf{0.93 \pm 0.01}$ | $\mathbf{0.44 \pm 0.02}$ | $\mathbf{0.09 \pm 0.06}$ |

reconstruction, all learning-based approaches yield noticeable improvements in overall image quality. However, it appears that U-Net and CCGR provide limited noise suppression, with the overall image quality remaining suboptimal. In contrast, the generative models, i.e., AUGAN and DDPM, provide improved denoising performance, yielding enhanced image contrast. Despite this improvement, one can observe impaired speckle patterns in the deep regions with AUGAN and hallucinated structures in the lateral regions with DDPM, which compromises anatomical structural consistency. The proposed method returns enhanced visual quality and maintains consistent anatomical structures, resulting in image quality very close to the reference. We report in the Table 1 the PSNR, SSIM, and NMI of each method using CW-MoCo as the high-quality reference. Consistent with the observation in Fig. 2, CDWC yielded the lowest scores in all metrics, and all learning-based models outperform the CDWC reconstruction, with ST-CODE leading the highest scores across all metrics.

**Temporal Coherence.** We computed tissue displacement and measured the EPE using the motion field extracted from the MoCo-based imaging as the reference. Fig. 3 illustrates the spatial distribution of EPE within the myocardium region during systole and diastole. UNet, AUGAN, and DDPM yielded relatively high EPE, while both CCGR and the proposed method reduce these errors by enforcing effective temporal constraints. Notably, the proposed method achieves

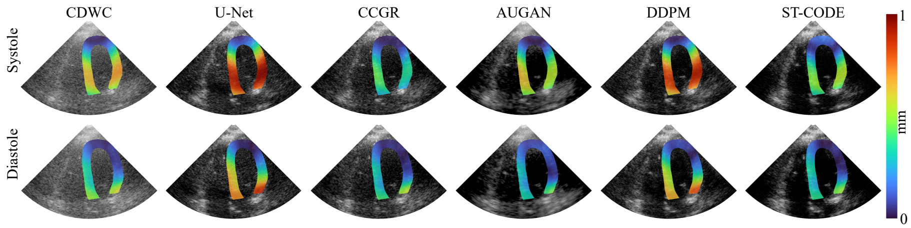


**Fig. 3.** Spatial distribution of EPE for motion estimated from the reconstructions of a representative case. From top to bottom: EPE distribution during systole and diastole

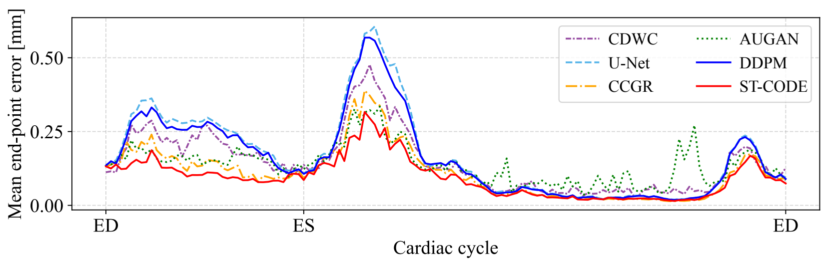


**Fig. 4.** MEPE obtained from reconstructions of different models over one cardiac cycle. Values were averaged across all cases in the test set at each time point.

**Table 2.** Quantitative results of the ablation study on the test set.

| Setting | DSG | TCS | PSNR [dB] ↑ | SSIM ↑ | NMI ↑ | MEPE [mm] ↓ |
|---|---|---|---|---|---|---|
| a | ✗ | ✗ | $28.23 \pm 0.49$ | $0.86 \pm 0.01$ | $0.33 \pm 0.02$ | $0.16 \pm 0.13$ |
| b | ✓ | ✗ | $30.08 \pm 0.56$ | $0.91 \pm 0.01$ | $0.40 \pm 0.02$ | $0.14 \pm 0.10$ |
| c | ✗ | ✓ | $29.13 \pm 0.70$ | $0.89 \pm 0.01$ | $0.39 \pm 0.01$ | $0.10 \pm 0.01$ |
| d | ✓ | ✓ | $\mathbf{30.44 \pm 0.47}$ | $\mathbf{0.93 \pm 0.01}$ | $\mathbf{0.44 \pm 0.02}$ | $\mathbf{0.09 \pm 0.06}$ |

the lowest EPE across the cardiac cycle, as consistently demonstrated with the MEPE curves in Fig. 4 and the quantitative values reported in Table 1.

**Ablation Study.** Table 2 summarizes the ablation results of the core components, i.e., DSG and TCS, in the proposed framework. The setting (a) without any of these modules demonstrates the worst performance in all metrics. With the addition of the DSG module in setting (b), clear improvements are observed in static image quality, which reflects its effectiveness in enhancing structural consistency. Moreover, TCS contributes to reductions in MEPE in setting (c), highlighting its role in improving temporal coherence. Finally, the complete model in setting (d) leads to the best performance in all metrics, which underscores their critical role in both spatial and temporal consistency.

## 4 Conclusions

In this work, we proposed a spatiotemporal constrained diffusion echocardiography framework, referred to as ST-CODE, for ultrafast cardiac imaging. By integrating dynamic subspace guidance and temporal correlation sampling in the diffusion framework, ST-CODE preserves spatial accuracy and enhances temporal coherence across frames. Extensive experiments on *in vivo* cardiac data demonstrate its superiority over existing reconstruction approaches. These results demonstrate that ST-CODE provides a reliable solution for anatomically consistent and temporally coherent reconstruction of cardiac imaging sequences under ultrafast DW acquisition conditions, and indicate the potential of diffusion-based reconstruction for high-quality ultrafast echocardiography.

**Acknowledgments.** This work was supported in part by the National Natural Science Foundation of China under Grant 62401381 and 62271335, in part by the Sichuan Science and Technology Program under Grant 2026NSFSC1404 and 2025ZNSFSC0470, and in part by the Sichuan Postdoctoral Research Program under Grant TB2024074.

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

## References

1. Asgariandehkordi, H., Goudarzi, S., Sharifzadeh, M., Basarab, A., Rivaz, H.: Denoising plane wave ultrasound images using diffusion probabilistic models. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **71**(11), 1526–1539 (2024)
2. Dockhorn, T., Vahdat, A., Kreis, K.: Genie: Higher-order denoising diffusion solvers. Advances in Neural Information Processing Systems **35**, 30150–30166 (2022)
3. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
4. Liang, W., Lu, J., Jing, J., Wang, H., Wang, A., Zhang, Y.: Unsupervised semantic constrained denoising framework for ultrasound plane-wave imaging. IEEE Transactions on Instrumentation and Measurement **74**, 1–13 (2025)
5. Lu, J., Liu, W.: Unsupervised super-resolution framework for medical ultrasound images using dilated convolutional neural networks. In: 2018 IEEE 3rd International Conference on Image, Vision and Computing (ICIVC). pp. 739–744. IEEE (2018)
6. Lu, J., Millioz, F., Garcia, D., Salles, S., Liu, W., Friboulet, D.: Reconstruction for diverging-wave imaging using deep convolutional neural networks. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **67**(12), 2481–2492 (2020)
7. Lu, J., Millioz, F., Garcia, D., Salles, S., Ye, D., Friboulet, D.: Complex convolutional neural networks for ultrafast ultrasound imaging reconstruction from in-phase/quadrature signal. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **69**(2), 592–603 (2022)

8. Lu, J., Millioz, F., Varray, F., Poree, J., Provost, J., Bernard, O., Garcia, D., Friboulet, D.: Ultrafast cardiac imaging using deep learning for speckle-tracking echocardiography. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **70**(12), 1761–1772 (2023)
9. Mohajery, M., Salles, S., Espeland, T., Wigen, M.S., Fadnes, S., Lovstakken, L.: The 3d estimation of mechanical wave velocities in the heart: methods and insights. IEEE Open Journal of Ultrasonics, Ferroelectrics, and Frequency Control **4**, 177–190 (2024)
10. Mougharbel, M., Porée, J., Lee, S.A., Xing, P., Wu, A., Tardif, J.C., Provost, J.: A unified framework combining doppler-based motion compensation, harmonic imaging, and angular coherence for high-quality b-mode in ultrafast echocardiography. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **72**(1), 141–152 (2024)
11. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International conference on machine learning. pp. 8162–8171 (2021)
12. Papadacci, C., Pernot, M., Couade, M., Fink, M., Tanter, M.: High-contrast ultrafast imaging of the heart. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **61**(2), 288–301 (2014)
13. Perdios, D., Vonlanthen, M., Martinez, F., Arditi, M., Thiran, J.P.: CNN-based ultrasound image reconstruction for ultrafast displacement tracking. IEEE Transactions on Medical Imaging **40**(3), 1078–1089 (2020)
14. Porée, J., Baudet, M., Tournoux, F., Cloutier, G., Garcia, D.: A dual tissue-doppler optical-flow method for speckle tracking echocardiography at high frame rate. IEEE Transactions on Medical Imaging **37**(9), 2022–2032 (2018)
15. Porée, J., Posada, D., Hodzic, A., Tournoux, F., Cloutier, G., Garcia, D.: High-frame-rate echocardiography using coherent compounding with doppler-based motion-compensation. IEEE Transactions on Medical Imaging **35**(7), 1647–1657 (2016)
16. Reynaud, H., Meng, Q., Dombrowski, M., Ghosh, A., Day, T., Gomez, A., Leeson, P., Kainz, B.: Echonet-synthetic: privacy-preserving video generation for safe medical data sharing. In: International conference on medical image computing and computer-assisted intervention. pp. 285–295 (2024)
17. Sivaanpu, A., Noga, M., Becher, H., Punithakumar, K., Le, L.H.: Denoising echocardiography with an improved diffusion model. In: 2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 1–4 (2024)
18. Stojanovski, D., Hermida, U., Lamata, P., Beqiri, A., Gomez, A.: Echo from noise: synthetic ultrasound image generation using diffusion models for real image segmentation. In: International Workshop on Advances in Simplifying Medical Ultrasound. pp. 34–43 (2023)
19. Tang, J., Zou, B., Li, C., Feng, S., Peng, H.: Plane-wave image reconstruction via generative adversarial network and attention mechanism. IEEE Transactions on Instrumentation and Measurement **70**, 1–15 (2021)
20. Van Phi, N., Duc, T.M., Hieu, P.H., Long, T.Q.: Echocardiography video synthesis from end diastolic semantic map via diffusion model. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 13461–13465 (2024)
21. Vray, D., Brusseau, E., Detti, V., Varray, F., Basarab, A., Beuf, O., Basset, O., Cachard, C., Liebgott, H., Delachartre, P.: Ultrasound Medical Imaging, p. 1–72 (2014)

22. Wang, A., Lu, J., Yang, Z., Liang, W., Zhuang, H., Zhang, Y.: Self-supervised content-consistent contrast enhancement framework for ultrafast plane-wave imaging. IEEE Transactions on Instrumentation and Measurement **73**, 1–12 (2024)
23. Wang, A., Lu, J., Zhang, Y.: CAW-CL: Cascaded adaptive weighted contrastive learning for unsupervised ultrasound plane-wave image reconstruction. In: ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1–5 (2025)
24. Wang, D., Chayer, B., Destrempes, F., Porée, J., Cardinal, M.H.R., Tournoux, F., Cloutier, G.: Ultrafast myocardial principal strain ultrasound elastography during stress tests: In vitro validation and in vivo feasibility. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **69**(12), 3284–3296 (2022)
25. Zhang, Z., Lei, Z., Zhou, M., Hasegawa, H., Gao, S.: Complex-valued convolutional gated recurrent neural network for ultrasound beamforming. IEEE Transactions on Neural Networks and Learning Systems **36**(3), 5668–5679 (2024)