Heavy-Tailed Denoising Diffusion for High Frame Rate Ultrasound Coherent Plane-Wave Compounding

Sajjad Afrakhteh, Tristan S.W. Stevens, Ruud J.G. van Sloun
Eindhoven University of Technology (TU/e), the Netherlands

***Abstract*—Ultrasound coherent plane wave compounding (CPWC) enables high frame-rate imaging by coherently summing echoes acquired at multiple steering angles. While CPWC is fast, using only a small number of steering angles makes the images prone to strong speckle noise and angle-related artifacts, which can obscure important anatomical details. Traditional denoising methods often reduce noise at the cost of blurring fine structures. To address this limitation, previous work has been done based on denoising diffusion probabilistic models (DDPM), which rely on Gaussian assumptions. In this paper, we show that the noise, in the case of a small number of compounded angles, is not well described by the commonly assumed Gaussian model; instead, it has heavy-tailed, non-Gaussian characteristics that are better captured by a Laplacian distribution. Based on this observation, we incorporate this Laplacian noise model directly into the sampling process of a denoising diffusion implicit model (DDIM). We further propose a fast denoising method based on an adaptive diffusion time step estimation from noisy inputs. Applying the technique to both experimental and *in vivo* carotid datasets shows that the proposed method enhances contrast and cyst visibility, as measured by the generalized contrast-to-noise ratio (gCNR). In particular, the results indicate that the number of transmissions can be reduced by 90% while achieving a gCNR comparable to that obtained with the maximum number of plane waves (PWs).**

***Index Terms*—CPWC, diffusion models, Laplacian, ultrasound.**

I. Introduction

Plane wave imaging (PWI) has enabled ultrafast ultrasound acquisition by replacing focused transmissions with unfocused plane waves, followed by digital beamforming [1], [2]. Coherent plane-wave compounding (CPWC) improves image quality by summing echoes from multiple steering angles, enhancing contrast, resolution, and signal-to-noise ratio (SNR). However, the achievable image quality strongly depends on the number of compounded plane waves. In practical imaging scenarios, particularly in motion-sensitive applications requiring high frame rates and with limited data storage, the number of plane waves must be restricted, leading to CPWC images that are degraded by residual noise and artifacts [3].

This trade-off between frame rate and image quality has motivated extensive research into denoising and artifact suppression methods for CPWC imaging. To reduce the number of transmission angles, Austeng *et al.* [4] proposed the use of an adaptive minimum variance (MV) beamformer. Although this approach achieves a threefold increase in frame rate, this beamforming method significantly increases computational complexity [5]. Instead of directly applying adaptive beamforming, Hashemseresht *et al.* [6] introduced adaptive angular compounding applied to radio-frequency (RF) images obtained from PWI at multiple steering angles. In a related effort, Afrakhteh and Behnam [7] proposed the use of tensor completion (TC) to reconstruct missing RF data corresponding to untransmitted angles in CPWC. While this technique improves the frame rate, its performance strongly depends on the randomly selected transmission coding pattern.

Alternatively, coherence-based methods have been explored to enhance CPWC image quality [8]–[12]. These techniques primarily suppress noise and enhance contrast, but often at the expense of speckle preservation and reduced signal-to-noise ratio in deeper regions.

Advances in deep learning (DL) have enabled data-driven approaches that show strong potential for high-frame-rate CPWC imaging. [13]–[15]. Despite their promising performance, such methods typically require extensive training datasets, costly computational hardware, and large model architectures to achieve meaningful frame-rate improvements.

Diffusion models [16] have recently emerged as a powerful framework for image restoration and inverse problems, owing to their ability to learn rich data priors through iterative denoising. One approach is to employ diffusion models as priors within a Bayesian posterior sampling framework. This was, for example, effectively done for the problem of ultrasound dehazing [17], [18]. Asgariandehkordi *et al.* found that it is also possible to leverage the generative process directly and align the generative denoising process with the denoising of ultrasound data [19]. For example, in this study, they found that for plane wave imaging, the disparity between low- and high-compounded images approximately follows a Gaussian distribution. This conveniently corresponds to the corruption process of standard diffusion models. In contrast, we observe that the difference between a CPWC RF frame, beamformed with many plane waves, and one beamformed from only a few, follows a more heavy tailed distribution.

Based on this hypothesis, we posit that the dominant degradation observed in CPWC images acquired with a limited number of plane waves is not governed by Gaussian noise, but

This project has received funding from the European Union under the Horizon Europe Marie Skłodowska-Curie Action (GA 101204288).

instead arises from heavy-tailed reconstruction errors induced by angular undersampling and incoherent wave compounding. In the case of a smaller number of PWs, incomplete angular diversity, phase inconsistencies across plane waves, residual sidelobes, and speckle mismatch give rise to spatially localized and intermittently large deviations that are not adequately described by Gaussian models. Based on this hypothesis, we posit that reconstruction and denoising frameworks, particularly diffusion models, that incorporate Laplacian priors or ℓ_1-type noise assumptions, are better aligned with the underlying CPWC noise statistics and can therefore achieve more accurate and robust image recovery than conventional Gaussian-based approaches.

Therefore, we propose a diffusion-based denoising framework tailored to CPWC ultrasound imaging with a limited number of angles. By explicitly incorporating Laplacian noise modeling through ℓ_1-based objectives, the proposed method aligns the diffusion process with the statistical characteristics of CPWC reconstruction errors. Furthermore, by integrating the forward noise model directly into the generative diffusion process, we eliminate the need for explicit posterior sampling, which is generally intractable for diffusion models and typically requires approximate inference schemes. Experimental results demonstrate that the proposed method achieves superior noise suppression and edge preservation compared to Gaussian-based learning approaches, providing an effective solution for CPWC denoising.

II. Methods

A. Problem formulation

Let $\mathbf{x}_0 \in \mathbb{R}^{H\times W}$ denote clean CPWC ultrasound beamformed radio-frequency (RF) data reconstructed using a large number of plane-wave transmissions by a delay-and-sum (DAS) beamformer. Given the noisy CPWC beamformed RF data $\mathbf{y}$, reconstructed using a subset of the total angles used to reconstruct $\mathbf{x}_0$, the objective is to recover $\mathbf{x}_0$ by suppressing residual noise and interference artifacts while preserving anatomical structures and fine details.

Recent diffusion-based denoising approaches for ultrasound imaging have largely relied on denoising diffusion probabilistic models (DDPMs) and Gaussian corruption assumptions, and rely on a Markovian forward diffusion process whose increments are Gaussian [16]. However, as discussed in Section I, CPWC images formed from a small number of plane waves exhibit reconstruction errors with pronounced heavy-tailed behavior, which are not well captured by Gaussian noise models.

In this work, we introduce a Laplacian diffusion model for CPWC denoising that explicitly accounts for the heavy-tailed statistics induced by angular undersampling. Our method builds on the Denoising Diffusion Implicit Model (DDIM) framework [20], but replaces the Gaussian corruption process with a Laplacian one and adopts an ℓ_1 training objective.

Unlike the Gaussian case, the sum of independent Laplacian random variables does not follow a Laplacian distribution. Consequently, a DDPM-style formulation that incrementally adds small Laplacian perturbations at each time step does not yield a consistent Laplacian marginal distribution after multiple steps. As a result, it is not trivial to construct a reverse-time stochastic process for Laplacian noise, and the derivations used in DDPMs become inapplicable.

To address this limitation, we adopt an implicit diffusion formulation analogous to Denoising Diffusion Implicit Models (DDIMs) [20]. Rather than defining a Markovian forward process that accumulates noise over time, we specify the forward process only through its marginal distributions at each time step. Specifically, each noisy sample $\mathbf{x}_t$ is generated by directly perturbing the clean image $\mathbf{x}_0$ with appropriately scaled Laplacian noise.

B. Diffusion process with Laplacian noise

Following the DDIM formulation, we define a discrete-time, non-Markovian forward process through its marginal distributions. For each time step $t \in \{0, 1, \dots, T\}$ the diffused sample is generated as:

$$\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1-\alpha_t}\epsilon, \qquad \epsilon \sim \mathrm{Laplace}(0, b), \quad (1)$$

where $\alpha_t \in (0, 1]$ is a monotonically decreasing noise schedule. Contrary to regular denoising diffusion models, the noise variable ϵ is not Gaussian but Laplacian. Similarly, however, ϵ is i.i.d. over pixels, and the scale parameter is set to $b = \frac{1}{\sqrt{2}}$ so that the corruption has unit variance. This construction directly specifies the forward marginal $p(\mathbf{x}_t \mid \mathbf{x}_0)$ and allows the sampling of $\mathbf{x}_t$ at arbitrary time steps without requiring the intermediate latent variables.

We parameterize a time-conditioned noise prediction network $\epsilon_\theta(\mathbf{x}_t, t)$ with neural network parameters θ. The network is trained to estimate the Laplacian corruption added at each time step t using the following objective:

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x}_0,t,\epsilon_t}\big[\|\epsilon_\theta(\mathbf{x}_t, t) - \epsilon_t\|_1\big], \quad (2)$$

where $\epsilon_t \sim \mathrm{Laplace}(0, b)$. This ℓ_1 loss corresponds to maximum likelihood estimation under a Laplacian noise model and provides robustness to the heavy-tailed noise observed in CPWC imaging with few plane waves. The perturbed sample $\mathbf{x}_t$ is obtained using the forward model in (1).

At inference time, the predicted noise can be used to form an estimate of the clean image:

$$\hat{\mathbf{x}}_0 = \frac{\mathbf{x}_t - \sqrt{1-\alpha_t}\epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\alpha_t}}. \quad (3)$$

Reconstruction then proceeds via a deterministic implicit update rule:

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}}\hat{\mathbf{x}}_0 + \sqrt{1-\alpha_{t-1}}\epsilon_\theta(\mathbf{x}_t, t). \quad (4)$$

The question that remains now is where to initialize the diffusion trajectory, which we detail in Section II-C.

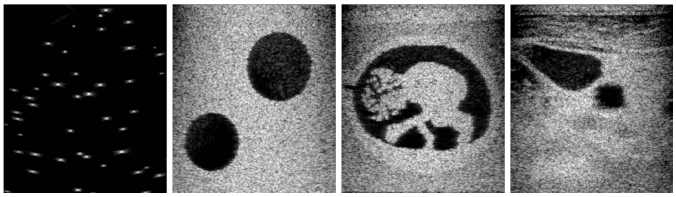

Fig. 1. Examples of beamformed images of simulated datasets using the zea simulator that are used for the training phase. The CPWC images were beamformed using 75 angles consistent with the PICMUS database which is used during inference.

C. Initial time step estimation

We expect the noisy observation $\mathbf{y}$ (measured CPWC RF data) to lie somewhere on the diffusion trajectory and thus correspond to a particular $\mathbf{x}_t$. To estimate at which diffusion time step t_τ we need to initialize, we compare $\mathbf{y}$ to a high quality reference image $\mathbf{x}_{\text{ref}}$, which corresponds to a CPWC image beamformed using all available angles (75 PWs). For each candidate time step t, an empirical noise residual is computed as $\hat{\mathbf{n}}_t = \mathbf{y} - \sqrt{\alpha}_t\,\mathbf{x}_{\text{ref}}$. Given that the noise component consists of unit variance Laplacian noise scaled by the diffusion schedule, we can equate the empirical noise variance to the theoretical variance $\mathrm{Var}(\hat{\mathbf{n}}_t) = (1-\alpha_t)$. To find the appropriate initial diffusion time step t_τ, we can find the time step that minimizes the difference:

$$t_\tau = \underset{t\in\{0,\dots,T\}}{\arg\min} \left\| \mathrm{Var}(\hat{\mathbf{n}}_t) - (1-\alpha_t) \right\|_1 . \tag{5}$$

Given the estimated time step t_τ, the reconstruction is initialized at x_{t_τ} and iteratively updated from $t_\tau \to 0$ using the implicit reverse sampling update rule in (4).

D. Datasets

1) Evaluation data: To evaluate the proposed technique, we use the standard *Plane Wave Imaging Challenge in Medical Ultrasound* (PICMUS) database [21]. Each dataset contains 75 steered plane waves covering the angle span from $-16°$ to $16°$. Each dataset is available in radio-frequency (RF) format. The PICMUS experimental and *in vivo* datasets are recorded using a 128-element linear array transducer. The height, width, and pitch of the transducer elements are $5\,\mathrm{mm}$, $0.27\,\mathrm{mm}$ and $0.3\,\mathrm{mm}$, respectively. Additionally, the aperture length is set to 38.4 mm. The center frequency and the sampling frequency are set to $5.2\,\mathrm{MHz}$ and $20.8\,\mathrm{MHz}$, respectively. All of the images are reconstructed using the zea toolbox [22].

2) Training data: To prevent any bias during training, the training data must be distinct from the evaluation data. For this purpose, we generate multiple CPWC images reconstructed at 75 angles ranging from $-16°$ to $16°$ using the zea simulator [22], which serve as clean training data. A diverse set of images is used as reflectivity maps in the simulations. Specifically, the grayscale images encode spatially varying scatterer reflectivity, which modulates the amplitudes of randomly distributed point scatterers within the imaging region. Initially, a large number of scatterers are randomly and uniformly distributed across the field of view. The corresponding image is then resized, mapped onto the same lateral–axial grid, and interpolated at each scatterer coordinate (x, z) to assign a local reflectivity value. An example of the resulting beamformed images generated with this approach is shown in Fig. 1.

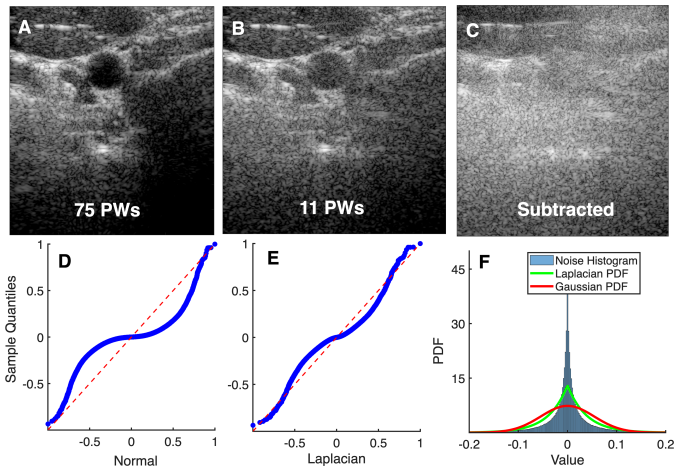

Fig. 2. CPWC noise distribution evaluation: **A)** Beamformed CPWC [75 PWs]. **B)** Beamformed CPWC [11 PWs]. **C)** Beamformed residual obtained by subtracting the beamformed RF data at 75 PWs from beamformed RF data at 11 PWs (The subtraction is performed in the RF domain, but figures are shown after log-compression and envelope-detection). **D)** Q-Q plot of the residuals against Gaussian distribution. **E)** Q-Q plot against Laplacian distribution. **F)** Log-likelihood test using Gaussian and Laplacian PDFs. The log-likelihood for Gaussian and Laplacian is obtained as 0.89 $\times 10^3$ and 1.1 $\times 10^3$, respectively.

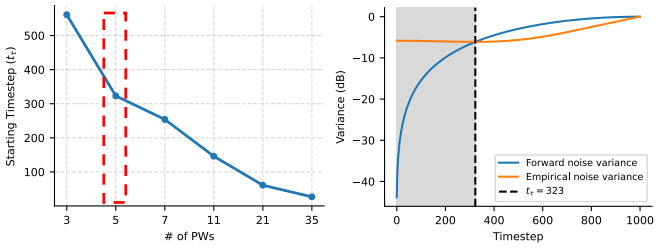

Fig. 3. Adaptive estimation of t_τ for *in vivo* carotid data. The left plot shows the starting time step fas a function of the number of compounded PWs. The right plot illustrates the 5 PW case as an example. The gray region shows the denoising interval used during the reverse diffusion process.

III. Results

To characterize the noise introduced by angular undersampling in CPWC imaging, we analyze the statistical distribution of reconstruction errors obtained with a reduced number of plane waves. First, a CPWC image beamformed using the maximum available number of angles (75 PW) is treated as a high-quality reference. For each lower PW setting, we subtract the corresponding beamformed RF data from this reference, such that the residual primarily captures the degradation caused by limited angular compounding.

Fig. 2 presents this analysis for *in vivo* carotid data from the PICMUS database, comparing 75 PW and 11 PW compounding. The resulting reconstruction error distribution is clearly

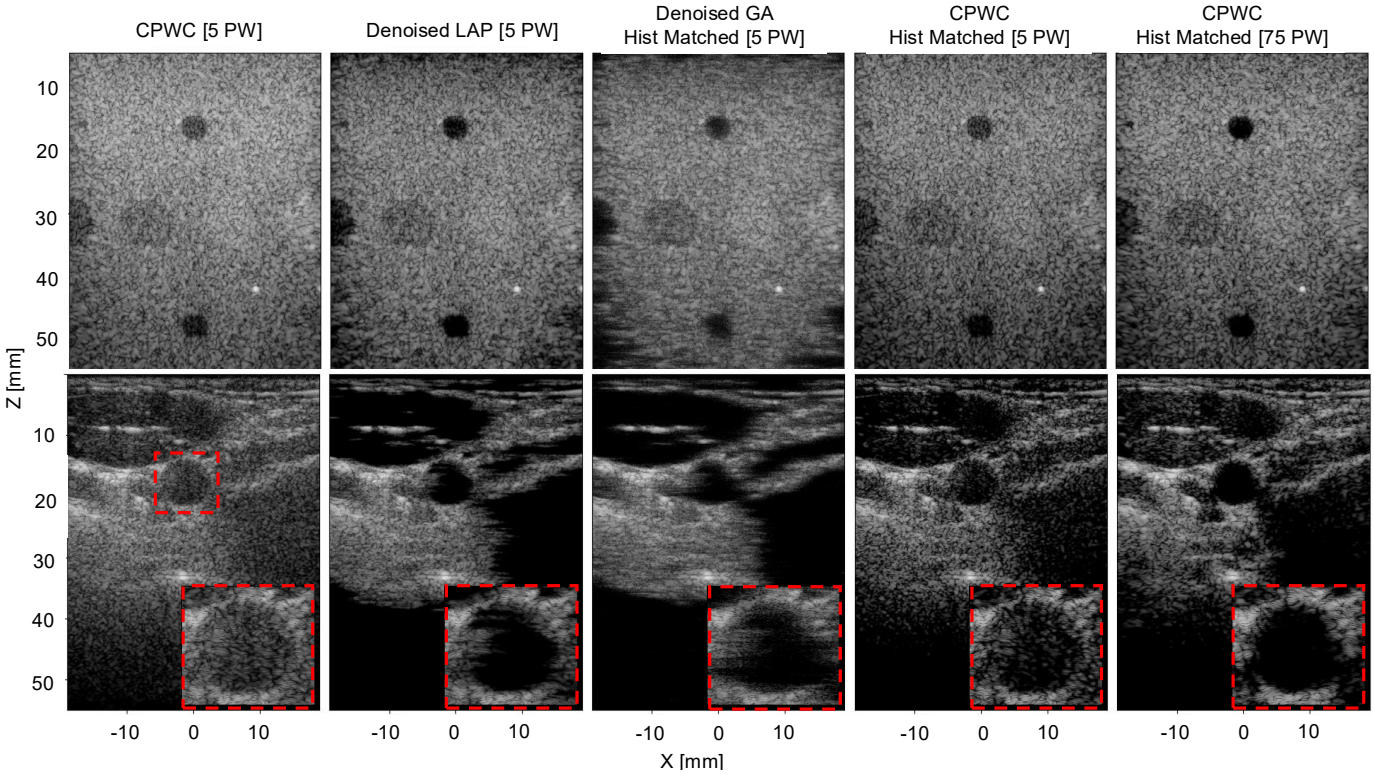

Fig. 4. Results of applying the proposed technique to the experimental cyst targets and *in vivo* carotid data in the PICMUS challenge data. All the beamformed images are shown in 60dB dynamic range. First column shows the noisy data while the second column shows the denoised when using the Laplacian priors. The third, fourth and fifth columns are shown histogram matched against the denoised Laplacian (second columns). The third clumn, shows the denoised image when using the Gaussian priors. The last two columns show the beamformed images for CPWC [5PW] and CPWC [75PWs] histogram matched with case denoised Laplacian (LAP: Laplacian, GA: Gaussian, PW: plane waves). A magnified view of the region indicated by the red box is presented for all beamformed images.

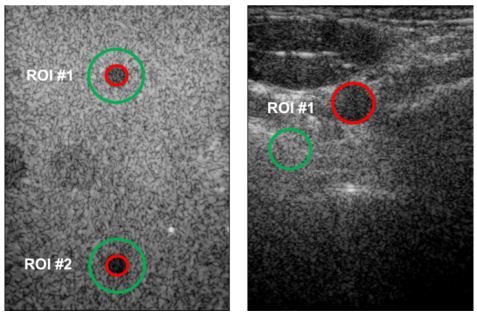

Fig. 5. ROIs for gCNR evaluation for experimental (first column) and *in vivo* carotid data (second column) in the PICMUS data. The the red and the green respectively indicate the foreground and background regions.

non-Gaussian and is more accurately described by a Laplacian distribution. This is evident from the strong deviations observed in the Gaussian Q–Q plot and the poor match between the empirical histogram and the Gaussian PDF. In contrast, the Laplacian Q–Q plot closely follows the theoretical quantiles, and the Laplacian density provides a much better fit to both the sharp peak around zero and the heavy tails of the noise distribution, as further supported by its higher log-likelihood. This observation directly motivates the use of diffusion-based denoising frameworks incorporating Laplacian priors or ℓ_1-type noise formulations for CPWC imaging.

We evaluated the proposed method on two examples within the PICMUS dataset. For efficient inference, we first estimated the optimal starting time step of the reverse diffusion process using (5). The estimated time step t_τ was computed for varying numbers of PWs. As shown in Fig. 3, for the *in vivo* carotid dataset, there is a clear monotonic relationship between the number of PWs and the estimated starting time step t_τ. This highlights the impact of the noise level on the proposed diffusion denoising process. Improved angular compounding reduces noise and enables the denoising process to begin earlier along the diffusion trajectory.

The denoised results and comparison with baselines are shown in Fig. 4. We compare CPWC images acquired with a low number of plane waves (5 PW) under different processing strategies while controlling for visual bias through histogram matching [23]. The proposed Laplacian denoised image is shown alongside the original 5 PW image, the Gaussian denoised result, and the fully compounded 75 PW reference. All images are histogram matched to the Laplacian denoised image. Under these conditions, the Laplacian-based approach produces visibly reduced background noise and sharper lesion boundaries. In contrast, Gaussian-based denoising yields smoother images but with diminished contrast and partial loss of fine structural details.

TABLE I
gCNR QUANTIFICATION FOR THE ROIS SHOWN IN FIG. 5.

	Experimental cyst targets		*In vivo* carotid
Method	ROI #1	ROI #2	ROI #1
CPWC [5PW]	0.59	0.58	0.74
Denoised LAP [5PW]	0.77	0.87	0.83
Denoised GA [5PW]	0.66	0.59	0.8
CPWC [75PW]	0.80	0.75	0.83

Quantitative evaluation was performed using the generalized contrast-to-noise ratio (gCNR) [24] over the regions of interest (ROIs) shown in Fig. 5. The results, summarized in Table I, confirm that denoising with a Laplacian noise model at 5 PW consistently leads to higher gCNR values for both experimental cyst phantoms and *in vivo* carotid data compared to the original CPWC [5 PW] and the Gaussian-denoised case, and in several instances approaches the performance of CPWC acquired with 75 PW. These results indicate that the acquisition frame rate can be increased by up to $15\times$ (from 75 PW to 5 PW) without compromising image quality in terms of gCNR.

IV. CONCLUSION

This work establishes that reconstruction errors in low-angle CPWC imaging exhibit pronounced heavy-tailed behavior that is inconsistent with conventional Gaussian noise models. By explicitly modeling this behavior using a Laplacian prior within a diffusion-based denoising framework, we align the generative process with the true statistics of CPWC degradation. Experimental results on *in vivo* data indicate improved noise suppression and enhanced lesion contrast, reflected by consistently higher gCNR values compared to Gaussian-based denoising. These findings suggest that heavy-tailed diffusion modeling is a promising and effective approach for improving image quality in low plane wave CPWC ultrasound imaging.

REFERENCES

[1] M. Tanter and M. Fink, "Ultrafast imaging in biomedical ultrasound," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 61, no. 1, pp. 102–119, 2014.

[2] J. Bercoff, G. Montaldo, T. Loupas, D. Savery, F. Mézière, M. Fink, and M. Tanter, "Ultrafast compound doppler imaging: Providing full blood flow characterization," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 58, no. 1, pp. 134–147, 2011.

[3] S. Afrakhteh, G. Iacca, and L. Demi, "A two-dimensional angular interpolation based on radial basis functions for high frame rate ultrafast imaging," *The Journal of the Acoustical Society of America*, vol. 154, no. 5, pp. 3454–3465, 2023.

[4] A. Austeng, C.-I. C. Nilsen, A. C. Jensen, S. P. Näsholm, and S. Holm, "Coherent plane-wave compounding and minimum variance beamforming," in *2011 IEEE International Ultrasonics Symposium*, pp. 2448–2451, IEEE, 2011.

[5] S. Afrakhteh and H. Behnam, "Low-complexity adaptive minimum variance ultrasound beam-former based on diagonalization," *Biomedical Signal Processing and Control*, vol. 62, p. 102110, 2020.

[6] M. Hashemseresht, S. Afrakhteh, and H. Behnam, "High-resolution and high-contrast ultrafast ultrasound imaging using coherent plane wave adaptive compounding," *Biomedical Signal Processing and Control*, vol. 73, p. 103446, 2022.

[7] S. Afrakhteh and H. Behnam, "Coherent plane wave compounding combined with tensor completion applied for ultrafast imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 68, no. 10, pp. 3094–3103, 2021.

[8] Y. Wang, C. Zheng, and H. Peng, "Dynamic coherence factor based on the standard deviation for coherent plane-wave compounding," *Computers in Biology and Medicine*, vol. 108, pp. 249–262, 2019.

[9] C. Zheng, H. Wang, X. Xu, H. Peng, and Q. Chen, "An adaptive imaging method for ultrasound coherent plane-wave compounding based on the subarray zero-cross factor," *Ultrasonics*, vol. 100, p. 105978, 2020.

[10] Y. Chen, Q. Kong, Z. Xiong, Q. Mao, M. Chen, and C. Lu, "Improved coherent plane-wave compounding using sign coherence factor weighting for frequency-domain beamforming," *Ultrasound in Medicine & Biology*, 2022.

[11] R. Mallart and M. Fink, "Adaptive focusing in scattering media through sound-speed inhomogeneities: The van cittert zernike approach and focusing criterion," *The Journal of the Acoustical Society of America*, vol. 96, no. 6, pp. 3721–3732, 1994.

[12] S. Afrakhteh and L. Demi, "Coherent plane wave compounding combined with euclidean distance transform for high frame rate and high contrast ultrafast imaging," *Ultrasonics*, p. 107759, 2025.

[13] J. Zhang, Q. He, Y. Xiao, H. Zheng, C. Wang, and J. Luo, "Ultrasound image reconstruction from plane wave radio-frequency data by self-supervised deep neural network," *Medical Image Analysis*, vol. 70, p. 102018, 2021.

[14] Y. Qi, Y. Guo, and Y. Wang, "Image quality enhancement using a deep neural network for plane wave medical ultrasound imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 68, no. 4, pp. 926–934, 2020.

[15] S. Goudarzi and H. Rivaz, "Deep reconstruction of high-quality ultrasound images from raw plane-wave data: A simulation and in vivo study," *Ultrasonics*, vol. 125, p. 106778, 2022.

[16] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.

[17] T. S. Stevens, F. C. Meral, J. Yu, I. Z. Apostolakis, J.-L. Robert, and R. J. van Sloun, "Dehazing ultrasound using diffusion models," *IEEE Transactions on Medical Imaging*, vol. 43, no. 10, pp. 3546–3558, 2024.

[18] T. S. W. Stevens, O. Nolan, J. Robert, and R. J. G. van Sloun, "Nuclear Diffusion Models for Low-Rank Background Suppression in Videos," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Barcelona, Spain*, 2026.

[19] H. Asgariandehkordi, S. Goudarzi, M. Sharifzadeh, A. Basarab, and H. Rivaz, "Denoising plane wave ultrasound images using diffusion probabilistic models," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2024.

[20] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," *arXiv preprint arXiv:2010.02502*, 2020.

[21] H. Liebgott, A. Rodriguez-Molares, F. Cervenansky, J. A. Jensen, and O. Bernard, "Plane-wave imaging challenge in medical ultrasound," in *2016 IEEE International ultrasonics symposium (IUS)*, pp. 1–4, IEEE, 2016.

[22] T. S. W. Stevens, W. L. van Nierop, B. Luijten, V. van de Schaft, O. Nolan, B. Federici, S. W. Penninga, N. I. Schueler, and R. J. G. van Sloun, "zea: A Toolbox for Cognitive Ultrasound Imaging," 2025.

[23] N. Bottenus, B. C. Byram, and D. Hyun, "Histogram matching for visual ultrasound image comparison," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 68, no. 5, pp. 1487–1495, 2020.

[24] A. Rodriguez-Molares, O. M. H. Rindal, J. D'hooge, S.-E. Måsøy, A. Austeng, M. A. L. Bell, and H. Torp, "The generalized contrast-to-noise ratio: A formal definition for lesion detectability," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 67, no. 4, pp. 745–759, 2019.