



HAL
open science

Ultrasound Imaging based on the Variance of a Diffusion Restoration Model

Yuxin Zhang, Clément Huneau, Jérôme Idier, Diana Mateus

► **To cite this version:**

Yuxin Zhang, Clément Huneau, Jérôme Idier, Diana Mateus. Ultrasound Imaging based on the Variance of a Diffusion Restoration Model. European Signal Processing Conference, Aug 2024, Lyon, France. hal-04622957

HAL Id: hal-04622957

<https://hal.science/hal-04622957v1>

Submitted on 24 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Ultrasound Imaging based on the Variance of a Diffusion Restoration Model

Yuxin Zhang*, Clément Huneau*, Jérôme Idier* and Diana Mateus*

*Nantes Université, École Centrale Nantes, LS2N, CNRS, UMR 6004, F-44000 Nantes, France

Email: yuxin.zhang@ls2n.fr

Abstract—Despite today’s prevalence of ultrasound imaging in medicine, ultrasound signal-to-noise ratio is still affected by several sources of noise and artefacts. Moreover, enhancing ultrasound image quality involves balancing concurrent factors like contrast, resolution, and speckle preservation. Recently, there has been progress in both model-based and learning-based approaches addressing the problem of ultrasound image reconstruction. Bringing the best from both worlds, we propose a hybrid reconstruction method combining an ultrasound linear direct model with a learning-based prior coming from a generative Denoising Diffusion model. More specifically, we rely on the unsupervised fine-tuning of a pre-trained Denoising Diffusion Restoration Model (DDRM). Given the nature of multiplicative noise inherent to ultrasound, this paper proposes an empirical model to characterize the stochasticity of diffusion reconstruction of ultrasound images, and shows the interest of its variance as an echogenicity map estimator. We conduct experiments on synthetic, in-vitro, and in-vivo data, demonstrating the efficacy of our variance imaging approach in achieving high-quality image reconstructions from single plane-wave acquisitions and in comparison to state-of-the-art methods. The code is available at: <https://github.com/Yuxin-Zhang/Jasmine/DRUSvar>.

Index Terms—Diffusion models, Inverse Problems, Ultrasound imaging

I. INTRODUCTION

Ultrasound (US) imaging finds extensive use in musculoskeletal, cardiac, obstetrical, and other medical diagnostic applications. In contrast to Magnetic Resonance or Computer Tomography, which are expensive or ionizing, ultrasound is real-time, affordable, portable, and minimally invasive. However, US imaging is affected by acoustic attenuation and artefacts (shadowing, reverberation, clutter), as well as electronic and speckle noise. The standard image reconstruction method, Delay and Sum (DAS) [1], converts time-domain signals into B-mode images. While this low-complexity approach enables fast reconstruction, it often results in poor image quality in terms of signal-to-noise ratio (SNR), contrast, and spatial resolution, especially with unfocused emissions.

In the last decade, model-based reconstruction methods [2], [3] have been proposed, providing physical plausibility but requiring high computational resources for convergence. Meanwhile, deep learning (DL) based techniques [4], [5] have excelled in image denoising and enhancement but face challenges with generalization and interpretability. To complement each

other, hybrid model-based DL approaches have gained interest. For instance, Chennakeshava *et al.* [6] solve a model-based plane-wave compounding problem by unfolding a classical optimization algorithm. Zhang *et al.* [7] propose a self-supervised beamforming approach enforcing explicit prior assumptions on the reconstruction through the loss function. Our work falls within this model-based DL family of approaches, but contrary to previous methods which are deterministic, we focus on stochastic methods based on the recent success of Diffusion Models [8]–[10], renowned for image synthesis.

Instead of conventionally training a “task-specific” neural network, diffusion-based inverse problem solvers use pre-trained neural networks to encode only the prior knowledge [11]–[16]. In a recent work [16], we adapted the Denoising Diffusion Restoration Models (DDRMs) framework [11] to incorporate the physics and constraints of US imaging through an approximate direct model, and with a self-supervised fine-tuning of the unconditional diffusion model. Also related, Asgariandehkordi *et al.* [17] apply diffusion denoising without considering a measurement model. Stevens *et al.* [18] focus on the specific task of dehazing cardiovascular US images.

The stochastic nature of a diffusion model leads to varying posterior samples under different noise initializations. Studies have demonstrated that the variance of these samples is non-uniform and tends to delineate edges or uncertain regions of the reconstructed objects [15], [18], [19]. Unlike such uncertainty interpretation, this paper proposes a novel perspective of the diffusion variance as an estimator for US images, following the nature of the multiplicative noise present in US reflectivity maps. The contributions of our work are:

- 1) Proposing an empirical model characterizing the stochasticity of diffusion-based US reconstructions.
- 2) Revealing that computing the variance of multiple diffusion reconstructed samples achieves higher SNR and contrast, and results in a despeckling without over-smoothing.
- 3) Confirming the superiority of the inverse-problem-informed diffusion reconstruction method over the denoising diffusion approach [17] through experiments on real data.

II. VARIANCE OF A DIFFUSION RESTORATION MODEL

A. Denoising Diffusion Restoration Models

Diffusion Denoising Probabilistic Models (DDPMs) are a class of parameterized Markov chains used to generate

This work has been supported by the European Regional Development Fund (FEDER), the Pays de la Loire Region on the Connect Talent scheme (MILCOM Project) and Nantes Métropole (Convention 2017-10470).

synthetic images from noise [8]–[10]. These models employ a fixed forward diffusion process and a learned backward generation process. During training, the forward process gradually adds Gaussian noise with variance σ_t^2 ($t = 1, \dots, T$) to a clean image \mathbf{x}_0 until it becomes random noise \mathbf{x}_T , while the backward process iteratively denoises the pure noise until the original clean image is reconstructed. Sampling involves simulating backward paths composed of multiple steps using either stochastic differential equations (SDEs) or ordinary differential equations (ODEs) to generate synthetic clean images.

Within the realm of model-based deep learning, a recent focus is made on leveraging learned prior knowledge from diffusion models to address inverse problems [11]–[16]. Unlike conventional task-specific learning approaches relying on paired datasets [4], [5], the above algorithms operate as Markov Chains conditioned on measurements through an *observation model*, allowing the utilization of pre-trained diffusion models across multiple tasks instead of task-specific training.

Assume an observation model $\mathbf{y}_d = \mathbf{H}_d \mathbf{x}_d + \mathbf{n}_d$, where \mathbf{y}_d denotes the measurements, \mathbf{H}_d is the degradation operator, \mathbf{x}_d is the sought quantity, and \mathbf{n}_d represents *i.i.d.* Gaussian additive noise with standard deviation σ_d . For such problems, Denoising Diffusion Restoration Models (DDRM) [11] serve as robust diffusion-based solvers operating in the spectral space of \mathbf{H}_d^1 . To this end, DDRM leverages the Singular Value Decomposition (SVD): $\mathbf{H}_d = \mathbf{U}_d \mathbf{S}_d \mathbf{V}_d^T$ with $\mathbf{S}_d = \text{Diag}(s_1, \dots, s_N)$, to decouple the dependencies between the measurements, and cast the original observation model as a denoising problem:

$$\bar{\mathbf{y}}_d = \bar{\mathbf{x}}_d + \bar{\mathbf{n}}_d$$

with $\bar{\mathbf{y}}_d = \mathbf{S}_d^\dagger \mathbf{U}_d^T \mathbf{y}_d$, $\bar{\mathbf{x}}_d = \mathbf{V}_d^T \mathbf{x}_d$, and $\bar{\mathbf{n}}_d = \mathbf{S}_d^\dagger \mathbf{U}_d^T \mathbf{n}_d$, where \mathbf{S}_d^\dagger is the generalized inverse of \mathbf{S}_d .

The ODE-based sampling in DDRM enables fast restoration by covering multiple diffusion steps at once, *i.e.*, directly transitioning from $\bar{\mathbf{x}}_t$ to $\bar{\mathbf{x}}_{t-k}$, where $k \geq 1$ denotes the number of skipped diffusion steps [20]. In practice, the value of $\bar{\mathbf{x}}_{t-k}$ is determined by linearly combining $\bar{\mathbf{x}}_t$, the transformed measurements $\bar{\mathbf{y}}_d$, the transformed current prediction $\bar{\mathbf{x}}_{\theta,t}$, and random noise using coefficients A , B , C , and D . These coefficients satisfy conditions related to noise and signal, specifically $(A\sigma_t)^2 + (B\sigma_d/s_i)^2 + D^2 = \sigma_{t-k}^2$ for the noise and $A + B + C = 1$ for the signal. The adjustment of these coefficients is governed by two hyperparameters, providing flexibility in the restoration process. After N_{it} denoising iterations, the final restored image \mathbf{x}_0 is obtained as $\mathbf{V}_d \bar{\mathbf{x}}_0$.

B. DRUS Variance Imaging

Having introduced the general principle of DDRM, this section introduces the adaptation of this framework to integrate a direct model describing the construction of an US image, incorporating both additive and multiplicative noise.

For a given imaged area, we denote $p(\mathbf{r})$ the echogenicity at position $\mathbf{r} \in \Omega$, which can be modeled as piecewise smooth. The tissue reflectivity can then be expressed as:

$$o(\mathbf{r}) = m(\mathbf{r})p(\mathbf{r}), \quad (1)$$

where m refers to the standard Gaussian multiplicative noise [4], [21]. To capture measurements from p , US imaging uses a transducer array with L elements, each emitting a pulse and receiving an echo signal. We consider the first-order Born approximation and no absorption within tissues to model the ultrasonic transmission-reception process as a linear time-invariant system.

The US wave emitted by the i th element passes through the object domain Ω and is received by the j th element. The received radio-frequency (RF) signal at time t is then:

$$y_{i,j}(t) = \int_{\mathbf{r} \in \Omega} h(t - \tau_{i,j}(\mathbf{r}))m(\mathbf{r})p(\mathbf{r})d\mathbf{r} + n_j(t), \quad (2)$$

where n_j represents the noise for the j th receiving element, h is a kernel resulting from the convolution of the emitted excitation pulse h_e and the transducer impulse response h_t .

For single plane-wave (PW) US imaging, the discretized model with N observation positions and K time samples for all L receiving elements can then be written as:

$$\mathbf{y} = \mathbf{H}(\mathbf{m} \odot \mathbf{p}) + \mathbf{n}, \quad (3)$$

where $\mathbf{y} = [\mathbf{y}_1^T, \dots, \mathbf{y}_L^T]^T \in \mathbb{R}^{KL \times 1}$, $\mathbf{H} \in \mathbb{R}^{KL \times N}$, $\mathbf{m} \odot \mathbf{p}$ is the componentwise product of \mathbf{m} by \mathbf{p} , both vectors being in $\mathbb{R}^{N \times 1}$, and $\mathbf{n} = [\mathbf{n}_1^T, \dots, \mathbf{n}_L^T]^T \in \mathbb{R}^{KL \times 1}$ includes electronic noise and model error. For simplicity, \mathbf{n} is assumed white Gaussian with standard deviation γ , which is reasonable for the plane wave transmission [22].

In the domain of medical US image reconstruction, some studies [7], [16] focus on restoring the reflectivity map, while others [23], [24] opt for reconstructing the echogenicity map. Our previous work, DRUS (Diffusion Reconstruction of US images) [16], falls into the former category. It projects the data in (3) using a beamforming matrix² $\mathbf{B} \in \mathbb{R}^{N \times KL}$:

$$\mathbf{B}\mathbf{y} = \mathbf{B}\mathbf{H}\mathbf{o} + \mathbf{B}\mathbf{n}, \quad (4)$$

and employs DDRM to estimate the reflectivity \mathbf{o} . The current paper explores a new application of DRUS, extending it to estimate US echogenicity \mathbf{p} maps.

Diffusion sampling in DRUS can yield varied $\hat{\mathbf{o}}$ due to stochasticity. Research has shown that the variance of diffusion posterior samples under additive noise tends to highlight edges of reconstructed objects [15], [18], [19]. However, due to the multiplicative noise inherent in US, the amplitudes of \mathbf{o} follow zero-mean Gaussian distributions with standard deviation depending on the intensity of \mathbf{p} , affecting neighboring pixel differences and thus diffusion variance. Therefore, the variance is intuitively linked to \mathbf{p} , and we introduce the following model to characterize this property:

$$\hat{\mathbf{o}}_c = \mathbf{m} \odot \mathbf{p} + \mathbf{p}^\beta \odot \mathbf{G}_c, \quad (5)$$

¹We use subscript d to refer to the original equations of DDRM.

²A weighted matched filter matrix of \mathbf{H} .

where \hat{o}_c represents the c th DRUS sample, \mathbf{G}_c follows a standard normal distribution to account for generative stochasticity, and β is an empirical parameter. It can be easily checked that $E[\hat{o}_c] = \mathbf{m} \odot \mathbf{p}$ and $\text{Var}[\hat{o}_c] = \mathbf{p}^{2\beta}$. Consequently, we define DRUSvar as an US echogenicity map estimator:

$$\hat{\mathbf{p}}_{\text{DRUSvar}} = \left(\frac{1}{C-1} \sum_{c=1}^C |\hat{o}_c - \hat{\mathbf{o}}_{\text{DRUSmean}}|^2 \right)^{\frac{1}{2\beta}}, \quad (6)$$

where $\hat{\mathbf{o}}_{\text{DRUSmean}} = \frac{1}{C} \sum_{c=1}^C \hat{o}_c$ is hereafter named DRUSmean. In our experiments, $\beta = 0.5$ (used in this paper) empirically yielded the most favorable results.

III. NUMERICAL ANALYSIS

We focus on reconstructing US images from a single plane wave (PW). This section demonstrates the feasibility of using DRUSvar as an US echogenicity map estimator under a controlled numerical analysis on two synthetic phantoms. As illustrated in Fig. 1, the occlusion phantom comprises 9 anechoic regions, while the second one contains 25 scatterers. For each phantom, we synthesize 9 reflectivity maps with various realizations of multiplicative noise, following (1).




Fig. 1. Echogenicity maps of the synthetic occlusion (left) and scatterer (right) phantoms. Metrics are calculated within the colored boundary regions.

To reproduce the behaviour of DRUSvar under a minimal observation model and emphasize on its pertinence in presence of multiplicative noise, we employ two 1-D convolution kernels instead of the full degradation matrix \mathbf{BH} in (4). Specifically, the lateral kernel is a Gaussian with a standard deviation of 0.17 mm, while the axial kernel is a cosine-modulated Gaussian with the same standard deviation to mimic a realistic US pulse-echo response (i.e., h in (2)).

The restoration process employs a US fine-tuned diffusion model at a resolution of 256×256 . The original open-source diffusion model (Downloaded from <https://github.com/openai/guided-diffusion>) was trained on ImageNet [25]. For fine-tuning, a dataset consisting of 2339 high-quality *in vitro* images was acquired from a TPAC Pioneer machine using the CIRS 040GSE phantom.

We conduct qualitative and quantitative evaluations on the above phantoms, considering various additive noise levels $std(\mathbf{n})$. All diffusion samples underwent $N_{it} = 50$ iteration steps and the DRUSmean/var images were constructed with $C = 10$ samples. The occlusion phantom is used to evaluate

the generalized Contrast to Noise Ratio (gCNR) [26] and the Signal to Noise Ratio (SNR), defined as:

$$g\text{CNR} = 1 - \int_{-\infty}^{\infty} \min \{g_{\text{in}}(v), g_{\text{out}}(v)\} dv,$$

$SNR = \mu_{\text{ROI}}/\sigma_{\text{ROI}}$, respectively, where the subscripts ‘in’ and ‘out’ indicate inside or outside the target regions, v denotes the pixel values, g refers to the histogram of pixels in each region, ‘ROI’ refers to the region of interest, and μ and σ denote the mean and the standard deviation respectively. The spatial resolution evaluated using -6 dB Full Width at Half Maximum (FWHM) is assessed on the scatterer phantom. A smaller FWHM value indicates a higher resolution.

The results in Figs. 2 (occlusion) and 3 (scatterer) comprise the quantitative scores and the qualitative images. The means and the standard deviations of the scores were calculated across all target regions from the 9 synthetic reflectivity maps. Fig. 2 demonstrates the significant superiority of DRUSvar over DRUSmean and By in terms of gCNR, SNR, and visualization. While DRUSvar exhibits slightly inferior performance compared to DRUSmean in terms of spatial resolution (Fig. 3), it remains stable even with large levels of additive noise.



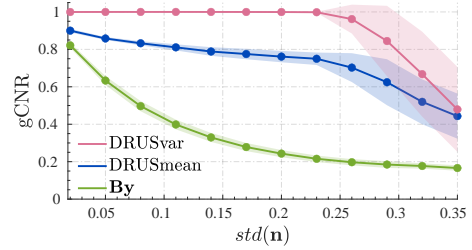


Fig. 2. Quantitative and qualitative comparison of the synthetic occlusion phantom-based images under varying levels of additive noise. Images are in decibels with a dynamic range [-60,0]. Left $std(\mathbf{n}) = 0.02$, right $std(\mathbf{n}) = 0.08$.



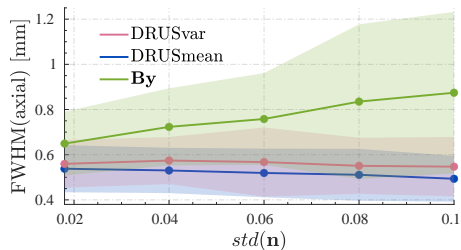


Fig. 3. Quantitative and qualitative comparison of the synthetic scatterer phantom-based images under varying levels of additive noise. Images are in decibels with a dynamic range $[-60,0]$. Left $std(n) = 0.018$, right $std(n) = 0.1$.

IV. REAL DATA RESULTS

We conducted quantitative and qualitative evaluations of our approach based on the publicly available experimental datasets from the Plane Wave Imaging Challenge in Medical UltraSound (PICMUS) [27]. The datasets were obtained with a 128-element L11-4v linear-array transducer. The transmit pulse has a central frequency of 5.208 MHz and a bandwidth ratio (BWR) of 67%. The sampling rate is 20.8 MHz. We name the 2 *in vitro* datasets *EC* (*Experimental Contrast*) and *ER* (*Experimental Resolution*), and the 2 *in vivo* datasets *CC* (*Carotid Cross-sectional*) and *CL* (*Carotid Longitudinal*).

The linear inverse problem model in (4) is leveraged in DRUSmean/var. The construction of the matrix \mathbf{H} necessitates the channel data acquisition parameters, the field of view, and the image resolution. The data acquisition parameters came from the PICMUS setup, the field of view spans from -18 mm to 18 mm in width and from 10 mm to 46 mm in depth, with the origin located at the transducer center. The image resolution was fixed at 256×256 . The construction of the beamforming matrix \mathbf{B} relies on the same parameters as \mathbf{H} , and the receive apodization weights defined by a window of Tukey0.25 and an *f-number* of 1.4.

The same diffusion model and evaluation metrics as in Sect. III are employed for *EC* and *ER*. *CC* and *CL* use a diffusion model fine-tuned on 1012 home-made *in vivo* images. $N_{it} = 50$ and $C = 10$. Quantitative results are presented in Table I. Fig. 4 and 5 depict the qualitative results. In addition to comparing DRUSvar with DAS (1PW

and 75 PWs) and DRUSmean, we also include a comparison against DENOmean with 10 DENO [17] samples, a state-of-the-art technique that denoises US images with diffusion models without solving an inverse problem. Furthermore, for despeckling comparison, DRUSvar is visually contrasted with DRUSmean+ADMSS [23]³, where ADMSS is an US despeckling method applied on beamformed images before log compression.

Table I shows that our approach, DRUSvar, significantly outperforms others in terms of contrast and SNR, while maintaining competitive spatial resolution. This is consistent with the numerical observations in Sect. III and the qualitative results in Fig. 4. Additionally, Fig. 5 demonstrates that DRUSvar mitigates the over-smoothing issue commonly seen in US despeckling methods like ADMSS [23]. However, for applications such as motion tracking, where speckle is useful, DRUSvar, being a despeckling method, is not suitable.

TABLE I
QUANTITATIVE COMPARISON ON THE PICMUS *in vitro* DATASETS. A AND L DENOTE AXIAL AND LATERAL DIRECTIONS RESPECTIVELY. BEST VALUES BOLDED, SECOND-BEST UNDERLINED.

		DAS	DENOmean	DRUSmean	DRUSvar
		75 PWs	1PW	1PW, [16]	1PW (proposed)
<i>EC</i>	gCNR \uparrow	0.95	0.87	0.95	0.97
	SNR \uparrow	1.92	1.97	1.93	1.87
<i>ER</i>	gCNR \uparrow	0.77	0.69	0.95	0.95
	FWHM A \downarrow	0.54	0.56	0.31	0.24
	[mm] L \downarrow	0.56	0.87	0.64	0.54
					0.32

V. CONCLUSION

This paper proposes an empirical model to characterize the stochasticity of diffusion reconstruction of US images, and introduces a novel US imaging approach, DRUSvar. The proposed approach leverages the empirical variance of multiple DRUS [16] samples to mitigate speckle noise, thereby enhancing contrast and SNR, without causing over-smoothing. The feasibility and competitiveness of DRUSvar are demonstrated in the task of single plane-wave US image restoration, using both synthetic and real-world datasets.

REFERENCES

- [1] V. Perrot, M. Polichetti, F. Varray, and D. Garcia, "So you think you can DAS?" *Ultrasonics*, vol. 111, p. 106309, 2021.
- [2] E. Ozkan, V. Vishnevsky, and O. Goksel, "Inverse problem of ultrasound beamforming with sparsity constraints and regularization." *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 65, pp. 356–365, 2018.
- [3] S. Goudarzi, A. Basarab, and H. Rivaz, "Inverse problem of ultrasound beamforming with denoising-based regularized solutions." *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 69, 2022.
- [4] D. Perdios, M. Vonlanthen, F. Martinez, M. Arditi, and J.-P. Thiran, "NN-based image reconstruction method for ultrafast ultrasound imaging." *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 69, pp. 1154–1168, 2022.
- [5] R. J. G. van Sloun, R. Cohen, and Y. C. Eldar, "Deep learning in ultrasound imaging." *Proc. IEEE*, vol. 108, pp. 11–29, 2020.
- [6] N. Chennakeshava, B. Luijten, O. Drori, M. Mischi, Y. C. Eldar, and R. J. G. van Sloun, "High resolution plane wave compounding through deep proximal learning," in *IEEE IUS*, 2020.

³Anisotropic Diffusion with Memory based on Speckle Statistics for Ultrasound Images. Code at MathWorks.

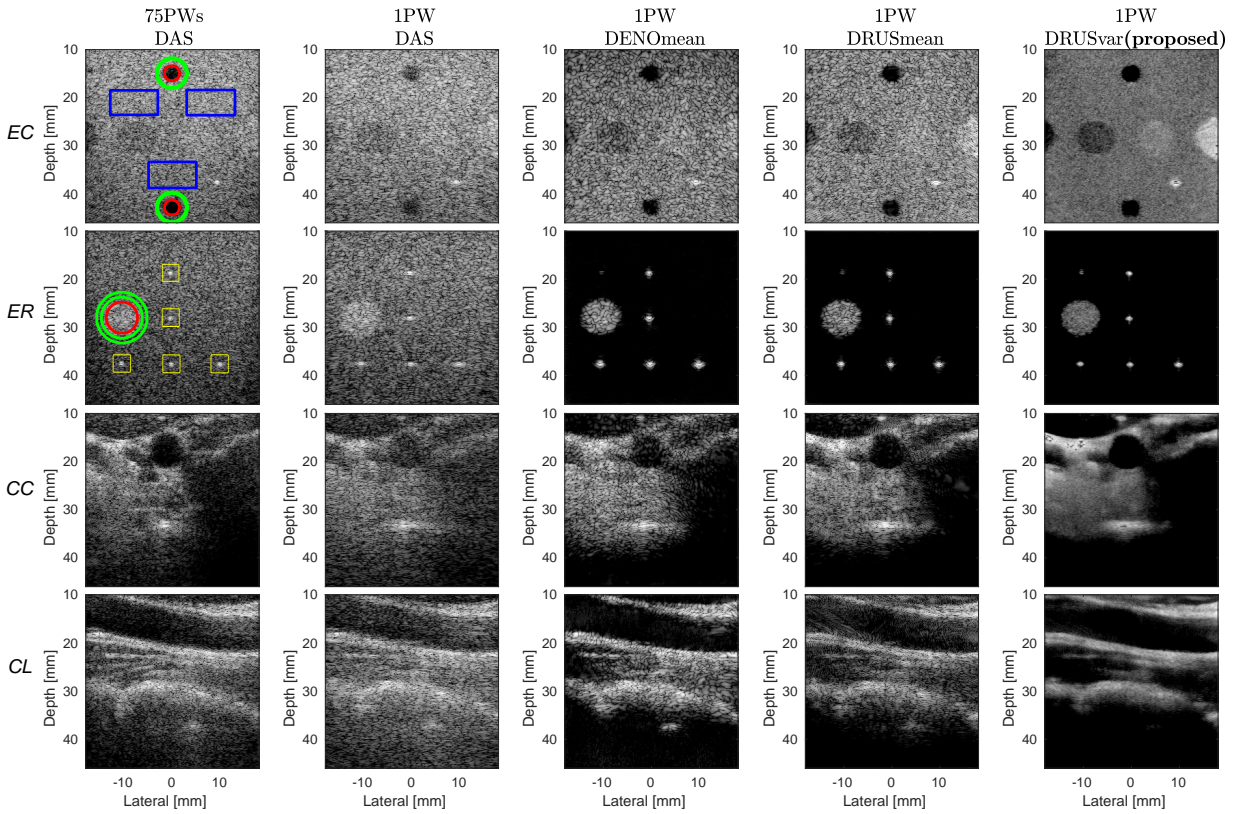


Fig. 4. Comparison of reconstructed images on the PICMUS *in vitro* (row [1-2]) and *in vivo* (row [3-4]) datasets. All images are in decibels with a dynamic range [-60,0]. The colored boundaries outline the regions where the *in vitro* evaluation metrics are calculated.




Fig. 5. Visual comparison of despeckled images on the *CC* dataset, in decibels [-60,0] dynamic range.

- [7] 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," *Med. Image Anal.*, vol. 70, 2021.
- [8] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *NeurIPS*, vol. 33, pp. 6840–6851, 2020.
- [9] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *ICML*, 2021, pp. 8162–8171.
- [10] P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," *NeurIPS*, vol. 34, pp. 8780–8794, 2021.
- [11] B. Kawar, M. Elad, S. Ermon, and J. Song, "Denoising diffusion restoration models," *NeurIPS*, vol. 35, pp. 23 593–23 606, 2022.
- [12] J. Song, A. Vahdat, M. Mardani, and J. Kautz, "Pseudoinverse-guided diffusion models for inverse problems," in *ICLR*, 2022.
- [13] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," in *ICLR*, 2022.
- [14] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," in *ICLR*, 2022.
- [15] H. Chung and J. C. Ye, "Score-based diffusion models for accelerated MRI," *Med. Image Anal.*, vol. 80, p. 102479, 2022.
- [16] Y. Zhang, C. Huneau, J. Idier, and D. Mateus, "Ultrasound image reconstruction with denoising diffusion restoration models," in *DGM4MICCAI*, 2023, pp. 193–203.
- [17] H. Asgariandehkordi, S. Goudarzi, A. Basarab, and H. Rivaz, "Deep ultrasound denoising using diffusion probabilistic models," in *IEEE IUS*, 2023.
- [18] 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 Trans. Med. Imaging*, 2024.
- [19] E. Horwitz and Y. Hoshen, "Confusion: Confidence intervals for diffusion models," *arXiv:2211.09795*, 2022.
- [20] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *ICLR*, 2021.
- [21] J. Ng, R. Prager, N. Kingsbury, G. Treece, and A. Gee, "Wavelet restoration of medical pulse-echo ultrasound images in an EM framework," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 54, 2007.
- [22] T. Chernyakova, D. Cohen, M. Shoham, and Y. C. Eldar, "iMAP beamforming for high-quality high frame rate imaging," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 66, 2019.
- [23] G. Ramos-Llordén, G. Vegas-Sánchez-Ferrero, M. Martín-Fernandez, C. Alberola-López, and S. Aja-Fernández, "Anisotropic diffusion filter with memory based on speckle statistics for ultrasound images," *IEEE Trans. Image Process.*, vol. 24, pp. 345–358, 2014.
- [24] H. Lee, M. H. Lee, S. Youn, K. Lee, H. M. Lew, and J. Y. Hwang, "Speckle reduction via deep content-aware image prior for precise breast tumor segmentation in an ultrasound image," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 69, pp. 2638–2650, 2022.
- [25] O. Russakovsky *et al.*, "Imagenet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, pp. 211–252, 2015.
- [26] A. Rodriguez-Molares *et al.*, "The generalized contrast-to-noise ratio: A formal definition for lesion detectability," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 67, pp. 745–759, 2020.
- [27] H. Liebgott, A. Rodriguez-Molares, F. Cervenansky, J. Jensen, and O. Bernard, "Plane-wave imaging challenge in medical ultrasound," in *IEEE IUS*, 2016.