InSPIRE: Multiparameter Inversion for Ultrasound Computed Tomography via Sequential Position-based Implicit Representation

Xiaolu Zeng[1,2,⋆], Weicheng Yan[1,2,⋆], Zhaohui Liu[1,2], Yun Wu[1,2], Hongrui Tan[1,2], Wu Qiu[1,2], Ming Yuchi[1,2,✉]

[1] College of Life Science and Technology, Huazhong University of Science and Technology, Hubei, China
[2] Advanced Biomedical Imaging Facility, Hubei, China

Abstract. Simultaneous reconstruction of sound speed and acoustic impedance in Ultrasound Computed Tomography (USCT) is essential for quantitative tissue characterization but remains a severely ill-posed inverse problem. Conventional Full Waveform Inversion (FWI) methods often suffer from cycle-skipping and parameter crosstalk due to non-convex optimization landscapes and limited low-frequency data. We propose InSPIRE, a framework for robust, unsupervised multiparameter inversion via sequential position-based implicit representations. Unlike discrete pixel-based approaches, we parameterize physical properties as continuous spatial functions using implicit neural representations. We construct a geometry-aware positional encoding with radially modulated Fourier features, imposing spatially adaptive spectral bias to resolve high-frequency details in the central imaging region where ring-array illumination is restricted. To disentangle parameter crosstalk, we introduce a dual-branch architecture with frequency-adaptive coordinate attention that dynamically adjusts feature importance across frequency stages, enabling independent reconstruction of sound speed and acoustic impedance. We devise a frequency-progressive sequential optimization strategy stabilizing convergence by updating the reference state with implicit noise regularization during bandwidth transitions. Extensive validation Extensive validation demonstrates that InSPIRE outperforms conventional multi-scale FWI, yielding high-fidelity parametric maps with superior convergence and minimized crosstalk.

Keywords: Ultrasound imaging · Full Waveform Inversion · Multiparameter Reconstruction · Implicit Neural Representations.

1 Introduction

Ultrasound Computed Tomography (USCT) enables radiation-free breast imaging at sub-millimeter resolution. Unlike qualitative B-mode sonography, USCT

⋆ These authors contributed equally to this work.
✉ Corresponding author: m.yuchi@hust.edu.cn

reconstructs quantitative maps of sound speed and acoustic impedance [5,6]. Sound speed reflects tissue density and bulk modulus, while acoustic impedance reveals sharp boundaries essential for lesion diagnosis [7,25]. Reconstructing these parameters simultaneously is severely ill-posed. Full-Waveform Inversion (FWI) addresses this by iteratively minimizing the mismatch between simulated and measured wavefields [17,8], but faces two major challenges. Limited low-frequency data causes cycle-skipping, trapping optimization in local minima [4,2]. Additionally, velocity-density coupling induces parameter crosstalk, projecting impedance features as artifacts onto the velocity map [15].These challenges have limited multiparameter USCT reconstruction primarily to synthetic scenarios, with few methods demonstrating reliable performance on clinical data.

Deep learning provides effective regularization for such inverse problems [3,14]. While supervised methods perform well on synthetic data, they fail to generalize to in vivo USCT due to lack of ground truth. This has motivated unsupervised approaches using untrained networks. Recent work successfully adapted Deep Image Prior (DIP) to USCT sound speed inversion [21,26], showing that CNN structural bias can regularize reconstruction without training data. However, standard CNNs exhibit spectral bias toward low frequencies [18], oversmoothening details required for acoustic impedance. Implicit Neural Representations (INRs) [12,24,13], such as Positional-encoding Image Prior (PIP) [19], address this by mapping spatial coordinates to high-dimensional features. Polar coordinate-based INR has been applied to single-parameter sound speed reconstruction [23], yet multiparameter inversion remains challenging: these coordinate-based methods cannot simultaneously model smoothly varying sound speed and discontinuous impedance without crosstalk, and lack mechanisms to compensate for the spatially varying illumination inherent to ring-array geometry.

To address these limitations, we propose InSPIRE, a framework for unsupervised multiparameter INversion via Sequential Position-based Implicit REpresentations. InSPIRE combines implicit neural representations with U-Net backbone architecture, parameterizing physical properties as continuous spatial functions while retaining multi-scale contextual awareness for global wave propagation. Our contributions are threefold. First, we introduce geometry-aware positional encoding with radially modulated Fourier features to resolve high-frequency details in the central region where ring-array illumination is restricted. Second, we propose a dual-branch architecture with frequency-adaptive coordinate attention that dynamically disentangles velocity-density crosstalk. Third, we implement sequential frequency-progressive optimization that stabilizes convergence by updating the reference state with implicit noise regularization during bandwidth transitions.

2 Method

2.1 Preliminaries

The forward model describes acoustic wave propagation in a heterogeneous medium. We consider a bounded domain $\Omega \subset \mathbb{R}^2$, where the pressure field $p(\mathbf{x}, t)$

satisfies the second-order acoustic wave equation:

$$\frac{1}{\rho(\mathbf{x})c(\mathbf{x})^2}\frac{\partial^2 p}{\partial t^2} - \nabla \cdot \left(\frac{1}{\rho(\mathbf{x})}\nabla p\right) = s(\mathbf{x}, t), \quad \forall \mathbf{x} \in \Omega, t \in [0, T] \tag{1}$$

where $s(\mathbf{x}, t)$ is the source term. The medium is characterized by sound speed $c(\mathbf{x})$ and density $\rho(\mathbf{x})$. To recover both bulk tissue properties and structural interfaces, we define the target parameter space as $m = [c, z]^\top$, where $z(\mathbf{x}) = \rho(\mathbf{x})c(\mathbf{x})$ is the acoustic impedance. Alternatively, the wave equation can be expressed in terms of impedance and sound speed:

$$\frac{1}{z(\mathbf{x})c(\mathbf{x})}\frac{\partial^2 p}{\partial t^2} - \nabla \cdot \left(\frac{c(\mathbf{x})}{z(\mathbf{x})}\nabla p\right) = s(\mathbf{x}, t) \tag{2}$$

We estimate these parameters by minimizing the mismatch between simulated and observed waveforms:

$$m^* = \arg\min_m \mathcal{J}(m) := \frac{1}{2}\int_0^T \|\mathcal{F}(m) - d_{\text{obs}}\|_2^2 \, dt \tag{3}$$

where d_{obs} denotes the observed data from the USCT system. The forward operator $\mathcal{F}(m)$ maps the medium parameters m to simulated pressure data by solving Eq. (1) and sampling the field at receiver locations. Because $\mathcal{F}$ is nonlinear and the acquisition aperture is limited, this optimization is ill-posed and non-convex, making standard gradient-based methods prone to local minima.

2.2 InSPIRE Framework

Framework Overview. We present InSPIRE, an unsupervised framework for multiparameter acoustic inversion using position-based neural representations. As shown in Fig. 1, the framework addresses ring-array geometry constraints, parameter crosstalk, and convergence stability through geometry-aware position encoding with radially modulated Fourier features, a dual-branch architecture with frequency-adaptive coordinate attention, and sequential frequency-progressive optimization with noise regularization.

InSPIRE reformulates the inversion as a neural field optimization problem. Unlike conventional FWI methods discretizing the medium on fixed grids [22,17] or DIP approaches mapping latent noise to images [21], InSPIRE parameterizes physical properties as continuous functions of spatial coordinates. A coordinate-based network $\Phi_\theta : \Omega \to \mathbb{R}^2$ maps 2D coordinates $\mathbf{x} = (x, y)$ to acoustic parameters:

$$[c(\mathbf{x}), z(\mathbf{x})]^\top = \Phi_\theta(\gamma_{\text{geo}}(\mathbf{x})), \tag{4}$$

where $\gamma_{\text{geo}}(\cdot)$ is the geometry-aware positional encoding. At frequency stage k, we minimize:

$$\theta_k^* = \arg\min_\theta \mathcal{L}_k(\theta) = \mathcal{L}_{\text{data}}(\theta) + \lambda_{\text{TV}}\mathcal{L}_{\text{TV}}(\theta) + \lambda_{\text{ref}}\mathcal{L}_{\text{ref}}(\theta, \theta_{k-1}^*), \tag{5}$$

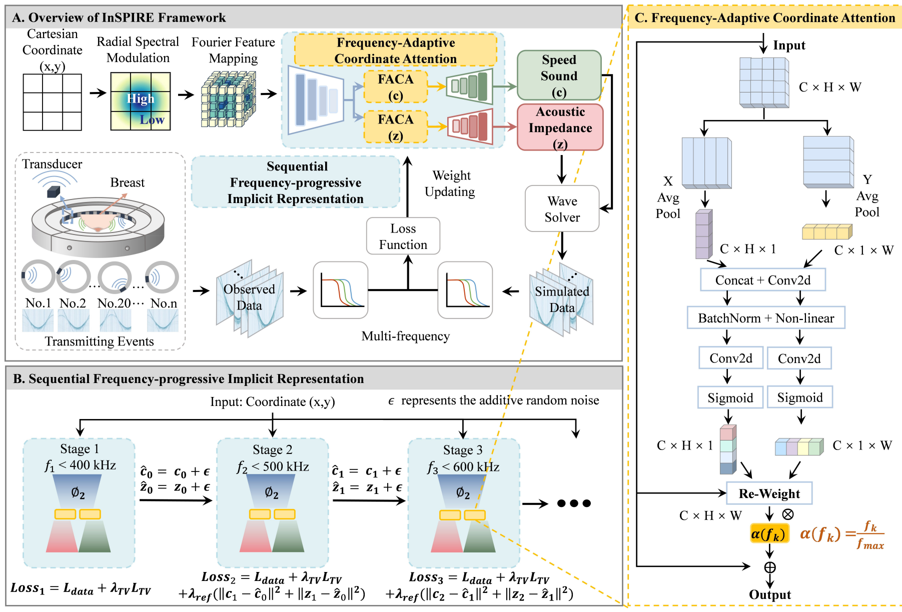

Fig. 1. Overview of our proposed InSPIRE framework

where $\mathcal{L}_{\text{data}}$ measures wavefield mismatch, $\mathcal{L}_{\text{TV}}$ enforces spatial smoothness, and $\mathcal{L}_{\text{ref}}$ maintains cross-stage consistency.

Geometry-Aware Positional Encoding. Standard Coordinate-Based Networks (CBNs) employ global Fourier feature mappings to mitigate spectral bias of MLPs [20]. A standard mapping $\gamma : \mathbb{R}^2 \to \mathbb{R}^{2L}$ projects coordinates $\mathbf{x}$ onto Fourier basis functions: $\gamma(\mathbf{x}) = [\cos(2\pi\mathbf{B}\mathbf{x}), \sin(2\pi\mathbf{B}\mathbf{x})]^\top$, where $\mathbf{B} \in \mathbb{R}^{L\times 2}$ is a Gaussian random matrix. While effective for natural images, this assumes spatially isotropic spectral distribution and ignores ring-array USCT geometry, where illumination and wavenumber coverage degrade centrally due to limited viewing angles.

To compensate for this, we introduce radially modulated Fourier encoding. Given normalized radial distance $r = \|\mathbf{x}\|_2$ from the array center, we define a spatial modulation function $\omega(r)$ to reweight spectral energy:

$$\gamma_{\text{geo}}(\mathbf{x}) = \omega(r) \odot \gamma(\mathbf{x}), \quad \text{with} \quad \omega(r) = 1 + \alpha \exp\left(-\frac{r^2}{2\sigma^2}\right), \tag{6}$$

where $\odot$ denotes element-wise multiplication, and α, σ control modulation intensity and bandwidth. This assigns higher weights to central features, imposing spatially adaptive bias to resolve high-frequency details. We set $\alpha = 2.0$ and $\sigma = 0.3$ based on the ring-array geometry.

Dual-Branch Architecture with FACA. Simultaneous reconstruction of sound speed and acoustic impedance suffers from parameter crosstalk due to velocity-density coupling in the wave equation. A shared representation often projects impedance boundaries as artifacts onto the sound speed map. To disentangle these modalities, we employ a dual-branch decoder with shared encoder $\mathcal{E}$ extracting geometric features from coordinate embeddings, followed by task-specific branches $\mathcal{D}_c$ and $\mathcal{D}_z$ generating independent feature maps:

$$\mathbf{f}_c = \mathcal{D}_c(\mathcal{E}(\gamma_{\text{geo}}(\mathbf{x}))), \quad \mathbf{f}_z = \mathcal{D}_z(\mathcal{E}(\gamma_{\text{geo}}(\mathbf{x}))). \tag{7}$$

We integrate a Frequency-Adaptive Coordinate Attention (FACA) mechanism inspired by [9]. FACA computes spatial and channel attention through separate pathways, then applies frequency-modulated recalibration:

$$\hat{\mathbf{f}} = \alpha(f_k) \cdot [\mathbf{f} \odot \sigma(\mathcal{A}_s(\mathbf{f})) \odot \sigma(\mathcal{A}_c(\mathbf{f}))] + \mathbf{f}, \tag{8}$$

where $\mathcal{A}_s$ and $\mathcal{A}_c$ denote spatial and channel attention modules, σ is the sigmoid function, and $\alpha(f_k) = f_k/f_{\max}$ is a normalized frequency-dependent scaling factor. As optimization sweeps from low to high frequencies, $\alpha(f_k)$ increases from 0 to 1, shifting network focus from global consistency to local boundaries.

Sequential Frequency-Progressive Optimization. Full-waveform inversion suffers from cycle-skipping when initialized with inaccurate velocity models, as high-frequency data can trap optimization in local minima. We adopt a sequential frequency-progressive strategy that gradually incorporates higher frequency components. Starting from low frequencies that provide long-wavelength updates, optimization progressively refines the model by incorporating higher frequencies to resolve finer details.

At each frequency stage k, we optimize network parameters using a bandwidth-limited data subset. The data fidelity term measures wavefield mismatch:

$$\mathcal{L}_{\text{data}}(\theta) = \frac{1}{2} \sum_{s=1}^{N_s} \int_0^T \|\mathcal{F}_s(\Phi_\theta(\cdot)) - d_{\text{obs}}^s\|_2^2 \, dt, \tag{9}$$

where N_s is the number of sources and $\mathcal{F}_s$ is the forward operator for source s. Total variation regularization enforces spatial smoothness while preserving boundaries:

$$\mathcal{L}_{\text{TV}}(\theta) = \int_{\Omega} (\|\nabla c(\mathbf{x})\|_2 + \|\nabla z(\mathbf{x})\|_2) \, d\mathbf{x}. \tag{10}$$

To stabilize transitions between frequency stages, we anchor the current solution to the previous stage with implicit noise regularization:

$$\mathcal{L}_{\text{ref}}(\theta, \theta_{k-1}^*) = \left\| \Phi_\theta(\cdot) - \left[\Phi_{\theta_{k-1}^*}(\cdot) + \epsilon \right] \right\|_2^2, \quad \epsilon \sim \mathcal{N}(0, \sigma^2 I), \tag{11}$$

where θ_{k-1}^* are the converged parameters from stage $k-1$. By perturbing the reference with Gaussian noise, this term prevents overfitting to the previous solution while retaining low-wavenumber information.

3 Experiments and Results

3.1 Experiment Settings

Datasets. For synthetic experiments, we construct numerical phantoms from MRI-derived breast tissue segmentations [11]. Ground truth sound speed and density maps are obtained from the segmentation masks, and acoustic impedance is computed as $z = \rho c$. Synthetic waveforms are generated using a ring-array configuration with 512 uniformly spaced transducers, of which 256 serve as active transmitters, excited by 500 kHz Ricker wavelets. To emulate realistic measurement conditions, we corrupt the simulated signals with additive Gaussian noise at a 40 dB signal-to-noise ratio. For in vivo validation, we use clinical USCT data acquired from three volunteers exhibiting varying compositions of glandular and adipose tissue.

Implementation Details. We validate the proposed InSPIRE framework on synthetic phantoms and in vivo breast data. All reconstructions are initialized from homogeneous background models for both sound speed and acoustic impedance. Training is performed using PyTorch [16] and Adam optimizer [10] with an initial learning rate of $1e^{-4}$. For synthetic data, the learning rate decays by a factor of 0.9 every 300 epochs over 5000 total epochs. For in vivo data, decay occurs every 200 epochs over 2000 total epochs. Regularization weights are set to $\lambda_{\mathrm{TV}} = 2e^{-6}$ across all experiments, with $\lambda_{\mathrm{ref}} = 1e^{-5}$ for synthetic data and $\lambda_{\mathrm{ref}} = 5e^{-5}$ for in vivo data, determined empirically. All experiments are conducted on an NVIDIA A100 GPU with 40GB memory.

Baselines and Evaluation Metrics. We compare InSPIRE against classical least-squares Full-Waveform Inversion (FWI) and its variants. For synthetic experiments, we evaluate four methods: FWI, multi-frequency FWI (mFWI), Deep Image Prior (DIP) [26], and the proposed InSPIRE. For in vivo experiments, we focus on mFWI, DIP, and InSPIRE due to their demonstrated robustness on clinical data. Reconstruction quality is assessed using Structural Similarity Index (SSIM), Peak Signal-to-Noise Ratio (PSNR), Learned Perceptual Image Patch Similarity (LPIPS), and Contrast-to-Noise Ratio (CNR).

3.2 Results

Synthetic Experiments. We conduct synthetic experiments using a ring-array configuration with radius 0.22 m. To satisfy the Courant-Friedrichs-Lewy stability condition [1], the computational domain is discretized into a 900×900 grid with 4000 time steps for finite difference forward modeling. We compare four approaches with FWI as baseline. Quantitative results are presented in Table 1, and reconstruction images with misfit maps are shown in Fig. 2. The multi-frequency FWI (mFWI) variant employs low-pass Butterworth filters with cutoff frequencies of 400, 500, 600 and 700 kHz to mitigate cycle-skipping.

In Vivo Experiments. For clinical validation, we use a 900×900 spatial grid with 4356 time steps. A solid gel coupling agent (SGCA) [27] serves as the acoustic coupling medium. All reconstructions are initialized from homogeneous

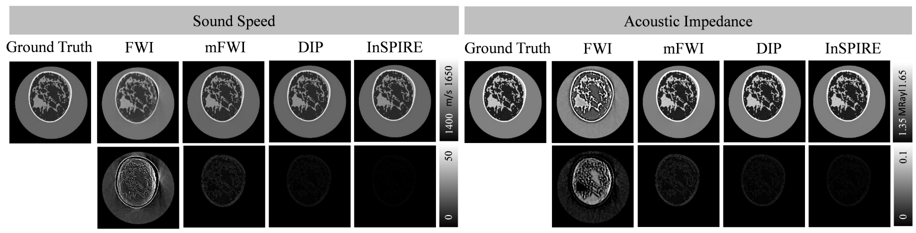

Fig. 2. Reconstruction results and misfit maps on synthetic data.

Table 1. Quantitative comparison on synthetic data.

Methods	Time	Sound Speed			Acoustic Impedance		
	(min)	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓
FWI	100	0.649	24.961	0.649	0.718	25.742	0.534
mFWI	85	0.916	32.982	0.338	0.922	33.601	0.325
DIP	170	0.920	33.128	0.319	0.929	33.874	0.317
InSPIRE	119	**0.952**	**34.255**	**0.189**	**0.964**	**34.792**	**0.162**

background models with a sound speed of 1500 m/s and an acoustic impedance of 1.5 MRayl, without additional prior information. Fig. 3 presents reconstruction results for three clinical cases across different methods. For Case 1 and Case 2, the region near the skin was chosen as the ROI and fat as the background. In Case 3, glands were selected as the ROI, while fat was designated as the background region. Table 2 presents the CNR results obtained from various methods.

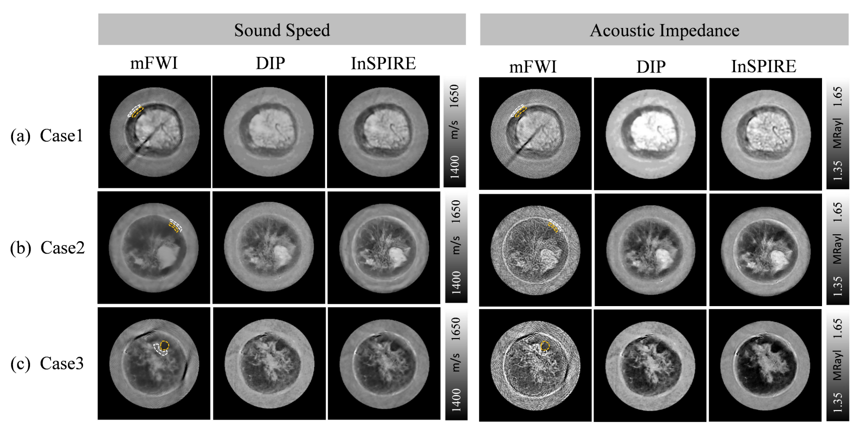

Fig. 3. Reconstruction results on in vivo breast data. ROI (white dashed lines) and background (yellow dashed lines) regions are outlined for CNR evaluation.

Table 2. CNR Results of the In Vivo Experiments.

Case	Sound Speed			Acoustic Impedance		
	mFWI	DIP	InSPIRE	mFWI	DIP	InSPIRE
Case 1	3.903	4.064	**5.459**	2.441	**5.730**	5.571
Case 2	3.153	2.915	**3.521**	1.898	4.052	**6.111**
Case 3	4.869	4.907	**5.018**	1.523	4.976	**7.123**

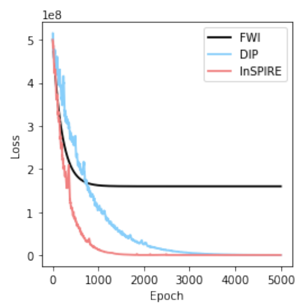

Fig. 4. Loss curve comparisons on synthetic data.

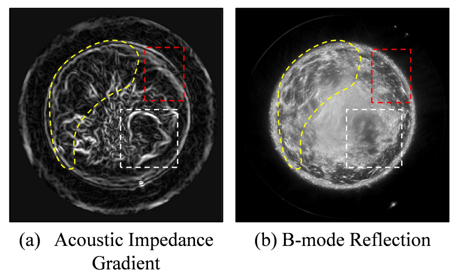

Fig. 5. Correspondence between acoustic impedance gradient and B-mode reflection imaging.

Table 3. Ablation study of the key components.

Experimental Setting				Sound Speed			Acoustic Impedance		
Base	GAPE	FACA	Freq-Progressive	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓
×	×	×	×	0.649	24.961	0.649	0.718	25.742	0.534
✓	×	×	×	0.920	33.128	0.319	0.929	33.874	0.317
✓	✓	×	×	0.922	33.285	0.298	0.932	33.912	0.295
✓	✓	✓	×	0.934	33.647	0.251	0.941	34.238	0.238
✓	✓	✓	✓	**0.952**	**34.255**	**0.189**	**0.964**	**34.792**	**0.162**

Ablation Studies. We evaluate the contribution of each component through incremental ablation experiments. The baseline uses a DIP network without Geometry-Aware Positional Encoding(GAPE) or attention modules(FACA). The transition from Row 2 (DIP baseline) to Row 3 includes shifting to coordinate-based input and adding GAPE. Table 3 shows that each module improves reconstruction performance. Fig. 4 demonstrates that InSPIRE method accelerates convergence. Although InSPIRE requires more iteration time than standard FWI, the substantial improvement in reconstruction quality justifies the increased computational cost.

4 Discussion and Conclusion

We introduce InSPIRE, an unsupervised coordinate-based framework for multiparameter acoustic inversion in USCT. By integrating geometry-aware encoding, frequency-adaptive coordinate attention, and frequency-progressive optimization, InSPIRE maintains physical consistency across scales, effectively resolving the velocity-density coupling and cycle-skipping inherent in traditional FWI. In the absence of ground truth for in vivo data, we validate structural consistency by comparing impedance gradients against USCT reflectivity maps (Fig. 5).

Unlike reflectivity maps that capture surface boundaries, InSPIRE reconstructs continuous material properties throughout tissue interior, providing complementary full-field information inaccessible to conventional imaging. The framework operates without paired training data or ground truth labels. Our results on healthy volunteers establish baseline references for early screening, laying the groundwork for future studies with patient cohorts and pathological correlation.

Acknowledgments. The authors are grateful to the High-Performance Computing platform of Huazhong University of Science and Technology and the Supercomputing Platform of Hubei Medical Devices Quality Supervision and Test Institute.

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

References

1. Abe, K., Higashimori, N., Kubo, M., Fujiwara, H., Iso, Y.: A remark on the courant-friedrichs-lewy condition in finite difference approach to pde's. Advances in Applied Mathematics and Mechanics **6**(5), 693–698 (2014)
2. Ali, R., Mitcham, T., Owolabi, I., McConnell, S., Duric, N.: Frequency-differencing strategy to kickstart full-waveform inversion without cycle skipping. JASA Express Letters **5**(1) (2025)
3. Arridge, S., Maass, P., Öktem, O., Schönlieb, C.B.: Solving inverse problems using data-driven models. Acta Numerica **28**, 1–174 (2019)
4. Bunks, C., Saleck, F.M., Zaleski, S., Chavent, G.: Multiscale seismic waveform inversion. Geophysics **60**(5), 1457–1473 (1995)
5. Duric, N., Littrup, P., Poulo, L., Babkin, A., Pevzner, R., Holsapple, E., Rama, O., Glide, C.: Detection of breast cancer with ultrasound tomography: First results with the computed ultrasound risk evaluation (cure) prototype. Medical physics **34**(2), 773–785 (2007)
6. Gemmeke, H., Ruiter, N.: 3d ultrasound computer tomography for medical imaging. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **580**(2), 1057–1065 (2007)
7. Greenleaf, J.F., Bahn, R.C.: Clinical imaging with transmissive ultrasonic computerized tomography. IEEE Transactions on Biomedical Engineering (2), 177–185 (2007)
8. Guasch, L., Calderón Agudo, O., Tang, M.X., Nachev, P., Warner, M.: Full-waveform inversion imaging of the human brain. NPJ digital medicine **3**(1), 28 (2020)

9. Hou, Q., Zhou, D., Feng, J.: Coordinate attention for efficient mobile network design. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13713–13722 (2021)
10. Kingma, D.P.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
11. Lou, Y., Zhou, W., Matthews, T.P., Appleton, C.M., Anastasio, M.A.: Generation of anatomically realistic numerical phantoms for photoacoustic and ultrasonic breast imaging. Journal of biomedical optics **22**(4), 041015–041015 (2017)
12. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021)
13. Molaei, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit neural representation in medical imaging: A comparative survey. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2381–2391 (2023)
14. Ongie, G., Jalal, A., Metzler, C.A., Baraniuk, R.G., Dimakis, A.G., Willett, R.: Deep learning techniques for inverse problems in imaging. IEEE Journal on Selected Areas in Information Theory **1**(1), 39–56 (2020)
15. Operto, S., Gholami, Y., Prieux, V., Ribodetti, A., Brossier, R., Métivier, L., Virieux, J.: A guided tour of multiparameter full-waveform inversion with multicomponent data: From theory to practice. The leading edge **32**(9), 1040–1054 (2013)
16. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems **32** (2019)
17. Pratt, R.G.: Seismic waveform inversion in the frequency domain, part 1: Theory and verification in a physical scale model. Geophysics **64**(3), 888–901 (1999)
18. Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the spectral bias of neural networks. In: International conference on machine learning. pp. 5301–5310. PMLR (2019)
19. Shabtay, N., Schwartz, E., Giryes, R.: Pip: Positional-encoding image prior. arXiv preprint arXiv:2211.14298 (2022)
20. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in neural information processing systems **33**, 7537–7547 (2020)
21. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9446–9454 (2018)
22. Virieux, J., Operto, S.: An overview of full-waveform inversion in exploration geophysics. Geophysics **74**(6), WCC1–WCC26 (2009)
23. Wang, Z., Yan, W., Liu, Z., Yuchi, M., Qiu, W.: P2INR-FWI: an Implicit Neural Representation Method for Speed of Sound Image Reconstruction in Ultrasound Computed Tomography . In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. vol. LNCS 15961, pp. 420 – 430. Springer Nature Switzerland (2025)
24. Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit neural representations for deformable image registration. In: International Conference on medical imaging with deep learning. pp. 1349–1359. PMLR (2022)

25. Wu, X., Li, Y., Su, C., Li, P., Lin, W.: Optimal transport assisted full waveform inversion for multiparameter imaging of soft tissues in ultrasound computed tomography. Ultrasonics **147**, 107505 (2025)
26. Yan, W., Zhang, Q., Wu, Y., Liu, Z., Zhang, H., Wang, Z., Yuan, J., Ding, M., Yuchi, M., Qiu, W.: Untrained neural network based full waveform inversion for breast sound speed imaging in ultrasound computed tomography. IEEE Transactions on Instrumentation and Measurement (2025)
27. Zhou, L., Zhang, Q., Wu, Y., Liu, Z., Wu, Y., Li, X., Qiu, W., Lou, C., Ding, M., Yuchi, M.: A coupling, stabilizing, and shaping strategy for breast ultrasound computed tomography (usct) with a ring array transducer. Ultrasonics **138**, 107212 (2024)