DOI 10.5162/Ultrasonic2025/P1.50

Enhancing Full Waveform Inversion for Ultrasound Imaging with Langevin Stein Variational Gradient Descent: Improved Accuracy and Uncertainty Quantification

Qiang Li[1], *Chengcheng Liu*[1], *and Dean Ta*[1]
[1] *College of Biomedical Engineering, Fudan University, Shanghai 200438, China*
corresponding.li_q@fudan.edu.cn; tda@fudan.edu.cn

Abstract: This study applies Langevin Stein Variational Gradient Descent (LSVGD) to Full Waveform Inversion (FWI) for medical ultrasound imaging. LSVGD enables joint reconstruction of speed of sound (SOS) and spatial uncertainty. Posterior accuracy is validated using a Gaussian mixture model. Applied to a 2D breast phantom, LSVGD achieves reconstruction quality on par with standard FWI, while highlighting uncertainty in heterogeneous regions. Results underscore LSVGD's utility for uncertainty-aware, high-resolution ultrasound imaging.

Keywords: Full waveform inversion, variational inference, uncertainty quantification, medical ultrasound imaging, Bayesian inference

Introduction

FWI is a high-resolution technique increasingly applied in medical ultrasound to recover acoustic tissue properties by minimizing the misfit between measured and simulated wavefields [1]. However, conventional FWI is deterministic and sensitive to the ill-posedness of inverse problems, yielding a single estimate without uncertainty quantification. The lack of uncertainty information limits clinical interpretability, particularly in heterogeneous biological tissues.

To address the limitations, Bayesian inversion reformulates FWI as a probabilistic problem, aiming to infer a posterior distribution over model parameters given the observed data [2]. Variational inference (VI) methods, particularly Stein Variational Gradient Descent (SVGD), approximate the posterior using particle-based updates without assuming a parametric form [3]. Yet, SVGD is prone to mode collapse, limiting its ability to capture multi-modal posteriors [4]. Langevin SVGD (LSVGD) enhances SVGD by introducing stochastic noise from Langevin dynamics, improving posterior exploration [4, 5].

This work introduces LSVGD to medical ultrasound FWI by deriving the log-posterior gradient. Posterior approximation accuracy is validated using a Gaussian Mixture Model. Application to a 2D breast phantom demonstrates accurate SOS reconstruction and pixel-wise uncertainty quantification. Results indicate LSVGD matches standard FWI in accuracy while enabling uncertainty-aware imaging.

Methods

FWI is increasingly being adapted for medical ultrasound imaging to reconstruct detailed maps of acoustic properties, such as SOS, density, attenuation, and others. FWI minimizes the difference between measured and simulated wavefields. The measured data consist of pressure fields from ultrasound transducers, while the simulated data are obtained by solving the acoustic wave equation forward in time. The acoustic wave equation in FWI is

$$m(\boldsymbol{x})\,\frac{\partial^2 p(\boldsymbol{x},t)}{\partial t^2} - \nabla^2 p(\boldsymbol{x},t) = f_s(\boldsymbol{x},t), \quad (1)$$

where $p(\boldsymbol{x},t)$ represents the pressure field at position $\boldsymbol{x}$ and time t, ∇^2 is the Laplacian operator, and $f_s(\boldsymbol{x},t)$ is the source term. The squared slowness $m(\boldsymbol{x})$ is related to the SOS, i.e., $m(\boldsymbol{x}) = 1/c^2(\boldsymbol{x})$. In 2D FWI imaging, the parameter vector $\boldsymbol{m}$ consists of $m(\boldsymbol{x})$ at each position $\boldsymbol{x}$. The vector $\boldsymbol{m}$ is iteratively adjusted to reduce the error between the measured data $\boldsymbol{d}_r^{\text{obs}}$ and the simulated data $\boldsymbol{d}_r^{\text{sim}}$ at each receiver r. The error quantified in FWI is the L_2 norm of the data misfit,

$$\Phi = \frac{1}{2}\sum_r \|\boldsymbol{d}_r^{\text{obs}} - \boldsymbol{d}_r^{\text{sim}}\|_2^2. \quad (2)$$

The gradient of the misfit is computed using adjoint-state methods, which allow for efficient updating of the model parameters.

Standard FWI is formulated as a deterministic optimization problem that aims to find a single best-fit

DOI 10.5162/Ultrasonic2025/P1.50

model by minimizing the misfit between $\boldsymbol{d}_r^{\text{obs}}$ and $\boldsymbol{d}_r^{\text{sim}}$. However, due to the highly nonlinear and ill-posed nature of the problem, conventional FWI yields a single estimate without quantifying uncertainty or confidence. To address this limitation, variational inference (VI)-based FWI reformulates the inversion as a probabilistic problem, aiming to infer a posterior distribution $p(\boldsymbol{m} \mid \boldsymbol{d}^{\text{obs}})$ over possible models given the observed data. VI enables a Bayesian framework,

$$p(\boldsymbol{m} \mid \boldsymbol{d}^{\text{obs}}) = \frac{p(\boldsymbol{d}^{\text{obs}} \mid \boldsymbol{m})\, p(\boldsymbol{m})}{p(\boldsymbol{d}^{\text{obs}})}, \tag{3}$$

where $p(\boldsymbol{m})$ denotes the prior distribution, $p(\boldsymbol{d}^{\text{obs}} \mid \boldsymbol{m})$ is the likelihood model, and $p(\boldsymbol{d}^{\text{obs}})$ is the marginal likelihood that normalizes the posterior. VI approximates the true posterior distribution $p(\boldsymbol{m} \mid \boldsymbol{d}^{\text{obs}})$ by finding a tractable surrogate distribution $q(\boldsymbol{m})$. To quantify the discrepancy between $q(\boldsymbol{m})$ and the true posterior, VI employs the Kullback–Leibler (KL) divergence and minimizes it to ensure the best approximation. The probabilistic formulation allows not only for reconstruction of the mean SOS map but also for pixel-wise uncertainty estimation via the standard deviation of the inferred posterior.

In VI-based FWI, the choice of the form of $q(\boldsymbol{m})$ directly influences the computational complexity and convergence rate. SVGD is a particle-based VI algorithm that approximates posterior distributions using a set of interacting particles $\boldsymbol{m}_i$. SVGD iteratively updates the particles to match the target posterior by minimizing the KL divergence through Stein's identity. SVGD combines the advantages of VI and deterministic optimization, leveraging insights from Stein's method, and does not assume a specific parametric form for $q(\boldsymbol{m})$. A common issue in SVGD is mode collapse, where particles cluster in a single mode, undermining the diversity of the approximation [4]. To address this, LSVGD enhances standard SVGD by introducing a stochastic noise term derived from Langevin dynamics, improving the exploration of the target posterior. LSVGD is a hybrid algorithm inspired by MCMC, combining the deterministic updates of SVGD with Langevin noise for better posterior approximation [5]. The update rule for LSVGD is

$$\begin{aligned}\phi(\boldsymbol{m}) =& \frac{1}{n}\sum_{j=1}^{n}\left[k\left(\boldsymbol{m}_j^t, \boldsymbol{m}\right)\nabla_{\boldsymbol{m}_j}\log p\left(\boldsymbol{m}_j^t \mid \boldsymbol{d}^{\text{obs}}\right) + \right.\\ & \left.\nabla_{\boldsymbol{m}_j} k\left(\boldsymbol{m}_j^t, \boldsymbol{m}\right)\right],\\ \boldsymbol{m}_i^{t+1} =& \boldsymbol{m}_i^t + \epsilon^t \phi(\boldsymbol{m}_i^t) + \sqrt{2\epsilon^t}\boldsymbol{\xi}_t,\end{aligned} \tag{4}$$

where $\boldsymbol{\xi}_t \sim \mathcal{N}(0, \mathbf{I})$ represents standard Gaussian noise and ϵ^t is the step size. The function $k(\boldsymbol{m}_j^t, \boldsymbol{m}) = \exp\left(-\|\boldsymbol{m}_j^t - \boldsymbol{m}\|^2/(2\sigma^2)\right)$ is the Radial Basis Function (RBF) kernel.

To incorporate the LSVGD update rule Eq. (4) into FWI, a critical step is the computation of the gradient of the log-posterior. According to Bayes' theorem Eq. (3), the gradient decomposes into the sum of the gradients of the log-likelihood and the log-prior. For the likelihood term, assuming a Gaussian likelihood yields

$$\log p(\boldsymbol{d}^{\text{obs}} \mid \boldsymbol{m}) \propto -\frac{1}{2}\sum_r \left\|\boldsymbol{d}_r^{\text{obs}} - \boldsymbol{d}_r^{\text{sim}}\right\|^2. \tag{5}$$

Eq. (5) indicates that the gradient of the log-likelihood can be efficiently computed using the adjoint-state method. In contrast, the gradient of the log-prior is analytically tractable when a specific prior distribution is assumed, such as a Gaussian or uniform prior.

LSVGD moves particles $\boldsymbol{m}_i$ to maximize the posterior distribution, which is equivalent to minimizing the negative log-posterior. Applying Bayes' theorem Eq. (3), the objective function is

$$\tilde{\Phi} = -\log p(\boldsymbol{d}^{\text{obs}} \mid \boldsymbol{m}) - \log p(\boldsymbol{m}). \tag{6}$$

For the Gaussian likelihood, Eq. (6) yields a data misfit term and a regularization term,

$$\tilde{\Phi} = \frac{1}{2}\sum_r \left\|\boldsymbol{d}_r^{\text{obs}} - \boldsymbol{d}_r^{\text{sim}}\right\|^2 - \log p(\boldsymbol{m}). \tag{7}$$

Results

To empirically validate the capability of LSVGD in approximating complex posterior distributions, a Gaussian Mixture Model (GMM) comprising three components is constructed as a toy example. The GMM, visualized as a gray curve in Fig. 1(a), is defined as

$$\begin{aligned}p(x) =& 0.2\mathcal{N}(x; -2, 0.5^2) + 0.6\mathcal{N}(x; 2, 0.7^2) + \\ & 0.2\mathcal{N}(x; 0, 1.5^2).\end{aligned}$$

The analytical mean and standard divation of the GMM are 0.8 and 1.831, respectively. LSVGD employs 100 particles initialized from a broad Gaussian distribution $\mathcal{N}(0, 5^2)$ and updates them over 500 iterations using a fixed step size of 5×10^{-2}. The updates utilize AdaGrad to adaptively adjust the learning rate. After reaching the maximum number of iterations, the particle distribution is displayed as green bins in Fig. 1(a). The approximated mean and standard deviation of the particles are 0.814 and 1.826, respectively. Fig. 1(b) illustrates the convergence behavior via the mean squared error (MSE) of both the estimated mean and standard deviation, with MSEs approaching zero. The experiment demonstrates LSVGD's ability to approximate multi-modal distributions, supporting its applicability to more complex inverse problems such as FWI in medical ultrasound imaging.

DOI 10.5162/Ultrasonic2025/P1.50

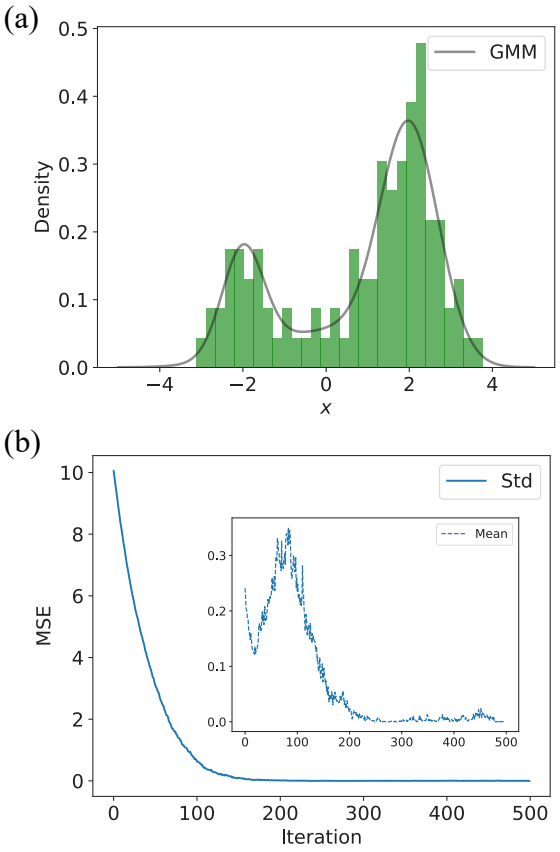

Fig. 1: LSVGD approximates the mean and standard deviation of a GMM. (a) Results of LSVGD with 100 particles. The gray curve represents the target GMM, and the green bins show the particle distribution after 500 iterations. (b) MSEs between the analytical and approximated mean and standard deviation over iterations.

To investigate the performance of LSVGD in FWI for medical ultrasound imaging, a two-dimensional breast phantom model [6] is employed. The background medium is water with a SOS of 1.5 km/s. The breast model includes five tissue types: fat, skin, glandular, ligament, and lesion, with SOS ranging from 1.411 km/s to 1.578 km/s. The simulation domain is discretized on a 351×351 grid with a spatial resolution of $\Delta x = \Delta y = 0.3$ mm. A Courant–Friedrichs–Lewy (CFL) number of 0.53 is adopted, resulting in a time step size of 1×10^{-4} ms. The imaging setup comprises 64 transmitters and 64 receivers, uniformly distributed along a circle of radius 0.51 cm centered within the model domain, as shown in Fig. 2. Each transmitter emits a Ricker wavelet with a center frequency of 0.5 MHz and duration of 0.1 ms.

For LSVGD-based FWI, the process is initialized from a pure water model with a constant SOS

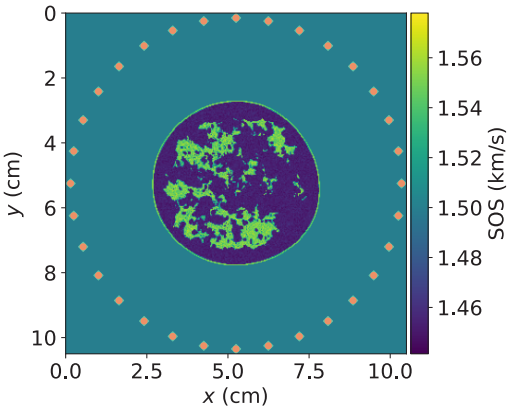

Fig. 2: A two-dimensional breast model. A total of 64 transmitters (orange dots) and 64 receivers (green diamonds) are uniformly distributed along a circle with a radius of 0.51 cm.

$\boldsymbol{c}_0 = 1.5$ km/s. Based on the squared slowness $\boldsymbol{m}_0 = 1/\boldsymbol{c}_0^2$, 20 initial particles are sampled from a Gaussian distribution with mean 0.444 s^2/km^2 and standard deviation 0.01. Given the SOS range of the breast phantom from $c_{\min} = 1.411$ km/s to $c_{\max} = 1.578$ km/s, the squared slowness is constrained within $m_{\min} = 1/c_{\max}^2$ and $m_{\max} = 1/c_{\min}^2$. The inversion employs a quasi-Newton optimization strategy using the L-BFGS-B algorithm with adaptive step size control. The maximum number of iterations is set to 50.

For the converged LSVGD particles, the final SOS distribution is computed for each particle as $\boldsymbol{c}_i = 1/\sqrt{\boldsymbol{m}_i}$. As shown in Fig. 3(a), the reconstructed mean SOS map resolves distinct tissue structures within the breast phantom. The relative error between the reconstructed and ground-truth SOS exhibits a maximum of 6.03% and an average of 0.27%. Fig. 3(b) presents the corresponding pixel-wise uncertainty, quantified by the standard deviation across particles. Notably, higher uncertainty values are observed within heterogeneous tissue regions exhibiting high SOS, indicating the algorithm's sensitivity to structural complexity. The maximum uncertainty reaches 0.056, and the average is 0.012. Areas with elevated uncertainty coincide with regions near the sources and receivers, where reconstruction errors are more pronounced in the mean SOS map.

To provide a baseline for comparison with LSVGD-based FWI, standard FWI is performed using the same experimental setup. The inversion starts from a homogeneous water model with a constant SOS of 1.5 km/s. The optimization is carried out using the L-BFGS-B algorithm, and the maximum number of iterations is set to 50. As shown in Fig. 4, the

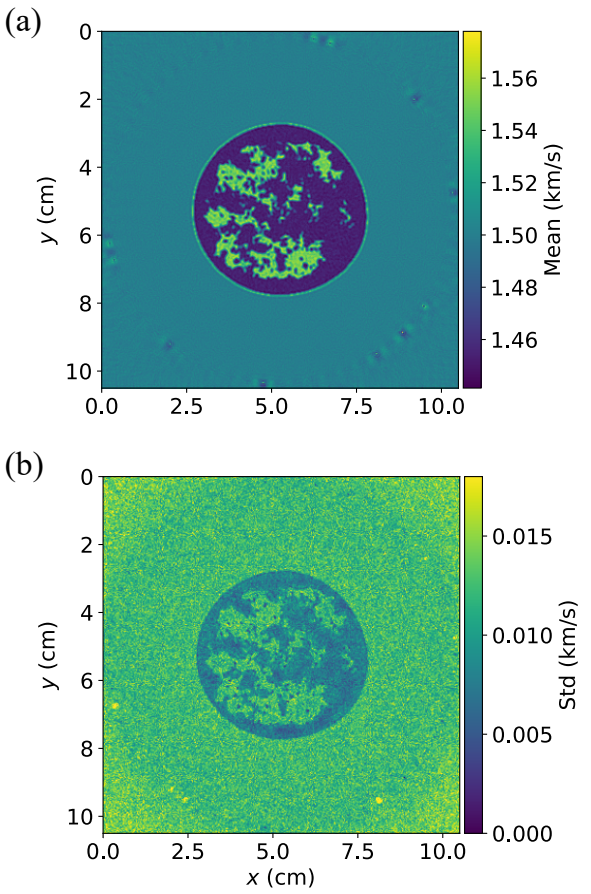

Fig. 3: SOS reconstruction using LSVGD with 20 particles. (a) Mean SOS map. (b) Uncertainty map quantified by standard deviation across particles.

reconstructed SOS map accurately resolves distinct tissue structures. The maximum and average relative errors are 5.06% and 0.21%, respectively. The LSVGD-based FWI attains reconstruction accuracy that is nearly indistinguishable from that of standard FWI, with the added benefit of pixel-wise uncertainty quantification.

Conclusion

This study derives the gradient of the posterior distribution with respect to model parameters and applies LSVGD to FWI in the context of medical ultrasound imaging. The effectiveness of LSVGD in approximating complex posterior distributions is first validated using a GMM. Subsequently, a two-dimensional breast phantom is used to reconstruct the spatial distribution of SOS. The reconstructed mean SOS obtained from LSVGD-based FWI achieves accuracy comparable to that of standard FWI. Moreover, LSVGD provides pixel-wise uncertainty quantification, with elevated uncertainty observed in heterogeneous tissue regions characterized by high SOS, highlighting the sensitivity

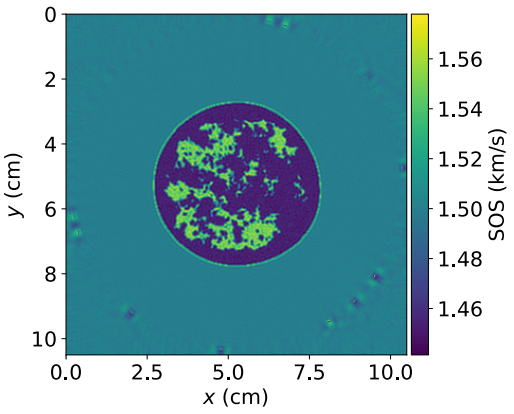

Fig. 4: Reconstructed SOS map using standard FWI.

to structural complexity. The results demonstrate the potential of LSVGD as a powerful tool for uncertainty-aware inversion in biomedical imaging.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant Nos. 12122403, 12034005, and 12327807) and the National Key R&D Program of China (Grant No. 2023YFC2410800). The authors also acknowledge the helpful suggestions provided by Dr. Xin Zhang, the first author of [5].

References

[1] L. Guasch et al. "Full-waveform inversion imaging of the human brain". In: *NPJ digital medicine* 3.1 (2020), p. 28.

[2] O. Bates et al. "A probabilistic approach to tomography and adjoint state methods, with an application to full waveform inversion in medical ultrasound". In: *Inverse Problems* 38.4 (2022), p. 045008.

[3] Q. Liu and D. Wang. "Stein variational gradient descent: A general purpose bayesian inference algorithm". In: *Advances in neural information processing systems* 29 (2016).

[4] D. Wang et al. "Stabilizing training of generative adversarial nets via langevin stein variational gradient descent". In: *IEEE Transactions on Neural Networks and Learning Systems* 33.7 (2020), pp. 2768–2780.

[5] X. Zhang and A. Curtis. "VIP - Variational Inversion Package with example implementations of Bayesian tomographic imaging". In: *Seismica* 3.1 (2024).

[6] F. Li et al. *2D Acoustic Numerical Breast Phantoms and USCT Measurement Data*. Version V1. 2021. DOI: 10.7910/DVN/CUFVKE. URL: https://doi.org/10.7910/DVN/CUFVKE.