

Research Article

Annealed Stein Variational Gradient Descent for Improved Uncertainty Estimation in Full-Waveform Inversion

Miguel Corrales¹, Sean Berti², Bertrand Denel³, Paul Williamson³, Mattia Aleardi², Matteo Ravasi¹

1. Physical Science and Engineering Division, King Abdullah University of Science and Technology, Saudi Arabia; 2. Department of Earth Sciences, University of Pisa, Italy; 3. Independent researcher

In recent years, Full-Waveform Inversion (FWI) has been extensively used to derive high-resolution subsurface velocity models from seismic data. However, due to the nonlinearity and ill-posed nature of the problem, FWI requires a good starting model to avoid producing non-physical solutions (i.e., being trapped in local minima). Moreover, conventional optimization methods fail to quantify the uncertainty associated with the recovered solution, which is critical for decision-making processes. Bayesian inference offers an alternative approach as it directly or indirectly evaluates the posterior probability density function using Bayes' theorem. For example, Markov Chain Monte Carlo (MCMC) methods generate multiple sample chains to characterize the solution's uncertainty. Despite their ability to theoretically handle any form of distribution, MCMC methods require many sampling steps; this limits their usage in high-dimensional problems with computationally intensive forward modeling, as is the FWI case. Variational Inference (VI), on the other hand, provides an approximate solution to the posterior distribution in the form of a parametric or non-parametric proposal distribution. Among the various algorithms used in VI, Stein Variational Gradient Descent (SVGD) is recognized for its ability to iteratively refine a set of samples (commonly defined as particles) to approximate the target distribution through an optimization process. However, mode and variance-collapse issues affect SVGD in high-dimensional inverse problems. This study aims to improve the performance of SVGD within the context of FWI by utilizing, for the first time, an annealed variant of the SVGD algorithm and combining it with a multi-scale strategy, a common practice in deterministic FWI settings. Additionally, we demonstrate that Principal Component Analysis (PCA) can be used to evaluate the performance of the optimization process and gain insights into the behavior of the produced particles and their overall distribution. Clustering techniques are also employed to provide

more rigorous and meaningful statistical analysis of the particles in the presence of multi-modal distributions (as is usually the case in FWI). Numerical tests, performed on a portion of the acoustic Marmousi model using both single and multi-scale frequency ranges, reveal the benefits of annealed SVGD compared to vanilla SVGD to enhance uncertainty estimation using a limited number of particles and thus address the challenges of dimensionality and computational constraints.

Miguel Corrales and Sean Berti equally contributed to this work.

Corresponding authors: Miguel Corrales, miguel.corrales@kaust.edu.sa; Sean Berti, sean.berti@unifi.it

1. Introduction

Full-Waveform Inversion (FWI) is a high-resolution imaging technique for estimating subsurface parameters from recorded seismic waveform data. Unlike methods that rely solely on the kinematic component of the recorded seismic waveforms (i.e., traveltimes), FWI exploits the entire wavefield information to invert for detailed subsurface models^[1]. However, the complex and nonlinear relationships between model parameters and seismic data — coupled with the oscillatory nature of the seismic waveforms, incomplete data coverage, and noise in the data — renders FWI an ill-posed inverse problem with a non-unique solution. In other words, many sets of model parameters can fit the data equally well within their inherent uncertainties; therefore, it is crucial to quantify the uncertainties to assess the confidence in the inverted models^[2].

FWI is typically addressed through optimization by minimizing a misfit function (e.g. L2 norm) between the observed and predicted seismograms^{[3][4]}. Due to the highly nonlinear nature of the problem and the multimodal landscape of the objective function, local optimization algorithms often get trapped in local minima. This challenge can be mitigated by enforcing specific requirements on the observed data, such as the presence of low frequencies or long offset, and/or a good starting model. Bozdağ et al.^[5] and Guo et al.^[6] have shown that a poor starting model can easily lead the inversion into a local minimum of the objective function, compromising the inversion outcome. Various misfit functions have been proposed to mitigate this dependency^{[7][8][9][10][11]}. Yet, when deterministic optimization algorithms are used alongside these misfit functions, no information about the uncertainty of the solution is provided.

Bayesian inference represents an alternative to deterministic optimization methods as it provides a probabilistic approach to quantify the uncertainty of the inverted models^{[12][11]}. Bayesian methods

embody Bayes' theorem to update our prior knowledge about the model parameters with new information obtained from the observed data. In this case, the posterior probability distribution (PPD), or the number of samples taken from such a distribution, represents the solution of the inversion process, thus Bayesian methods offer a more comprehensive solution describing all parameter values consistent with the observed data and quantifying their relative probabilities. Markov Chain Monte Carlo (MCMC) methods are commonly employed to characterize the PPD by constructing multiple chains of successive samples from the target posterior distribution through structured random walks in the parameter space. These samples form the basis for inferring valuable statistics of the PPD, and thereby enable the estimation of uncertainties that affect the recovered solution. One such method, the random walk Metropolis algorithm, has been applied across various geophysical problems, including electrical resistivity inversion^[13], travelttime tomography^[14], and gravity inversion^[12]. However, this algorithm faces significant computational challenges, as the curse of dimensionality^[15] restricts its applicability in high-dimensional problems and computationally expensive forward modeling operators, such as those encountered in FWI.

Over the past decade, by leveraging rapid advances in computing capabilities, researchers have revisited sampling-based methods to solve Bayesian FWI problems, developing sophisticated algorithms that aim to improve the efficiency for large-scale inversions. However, while more computationally feasible than standard MCMC approaches, these methods may also risk losing information due to implicit undersampling in high-dimensional spaces. These include Hamiltonian Monte Carlo (HMC)^{[16][17]}, stochastic Newton MCMC^[18], parallel tempering^[19], and gradient-based MCMC^{[20][21][22][23]}. Finally, trans-dimensional MCMC represent another class of MCMC techniques in which the number of model parameters is treated as an additional unknown^{[14][24][25][6]}. Despite their robustness and efficiency, when applied to problems such as FWI, MCMC methods typically require a large number of sampling steps ($> 10^6$) and a long burn-in period to achieve precise uncertainty estimations.

Variational Inference (VI) has emerged as an appealing alternative as it offers greater adaptability for approximating the posterior distributions with significantly lower computational overhead than MCMC^{[26][27][28]}. In VI, a set of simple probability distributions is defined (often called the variational family), and an optimal member of such a family is sought to approximate the true PPD. The Kullback-Leibler (KL) divergence measures the disparity between the two distributions, which enables potentially efficient and parallelizable optimization processes with well-understood convergence criteria. This can be achieved by either directly estimating the free parameters of the chosen distribution that best

approximate the true PPD^{[29][30]} or by deterministically modifying a set of samples from the proposal distribution to match the PPD^{[31][32]}. Variational approaches have been applied to various problems in geophysics, including traveltime tomography^{[33][34]}, seismic denoising^[35], seismic interpolation^[36], earthquake hypocenter inversion^[37], 2D FWI^{[38][39]}, and 3D FWI^{[40][41]}.

Recently, particle-based VI methods have emerged to bridge the gap between parametric VI and MCMC techniques. These methods utilize a specific number of samples, or particles, to represent the approximate distribution, akin to MCMC, whilst updating these particles through an optimization process similar to VI. This hybrid approach offers greater flexibility than parametric VI and is more particle-efficient than MCMC, as it fully leverages particle interactions. A notable example of this category is the Stein Variational Gradient Descent (SVGD) method^[32], which has already been applied to post-stack seismic inversion^[42], petrophysical inversion^[43], and FWI^{[41][44]}. This deterministic sampling algorithm iteratively minimizes the Kullback-Leibler (KL) divergence between the chosen approximate distribution and the target density to ensure that the final set of particles is distributed according to the desired posterior distribution. Despite the empirical successes of SVGD, its application to high-dimensional problems remains challenging, as it becomes computationally demanding to sample more particles than there are unknowns. SVGD can suffer from mode- and variance-collapse issues as the dimensionality of the problem increases. More specifically, variance collapse refers to the scenario in which the variance estimated by SVGD is significantly smaller than the true variance of the target distribution^[45]. This is undesirable because underestimation of the variance leads to a failure in explaining the uncertainty of the model predictions, which is a key benefit of Bayesian inference.

This study aims to enhance the performance of SVGD within the FWI framework by replacing the standard SVGD algorithm with an annealed variant. We conduct numerical experiments on a portion of the Marmousi model using both single- and multi-scale frequency approaches to evaluate the effectiveness of these methods to improve uncertainty estimation when working with a limited number of particles. Furthermore, we propose a number of additional strategies to augment our analysis of the particles after the SVGD process. First, Principal Component Analysis (PCA) is used to evaluate the performance of SVGD and gain deeper insights into the behavior and distribution of the particles. In addition to reducing the dimensionality of the solution space, PCA highlights the least confident model combinations, enabling more precise analysis of model uncertainties compared to simply examining the parameter variances (the diagonal of the covariance matrix). However, interpreting these results can be more challenging. Additionally, we employ clustering techniques to identify whether particles converge

to distinct modes, allowing for more rigorous and meaningful statistical insights by grouping particles into geological and non-geological clusters. Overall, this research addresses the challenges of high dimensionality and computational constraints by bridging the gap between theoretical potential and practical applications and, thereby, faster and more informed decision-making.

2. Theoretical Framework

FWI aims to estimate subsurface model parameters, such as P-wave velocity, represented by $\mathbf{m} \in \mathbb{R}^m$, from observed seismic data $\mathbf{d} \in \mathbb{R}^d$, where m and d denote the dimensions of model and data spaces, respectively. In order to capture the uncertainties inherent in this estimation process, any Bayesian estimation algorithms formulates this inverse problem using Bayes' rule:

$$p(\mathbf{m}|\mathbf{d}) = \frac{p(\mathbf{d}|\mathbf{m})p(\mathbf{m})}{p(\mathbf{d})} \quad (1)$$

where the probability density function (PDF) of the posterior, $p(\mathbf{m}|\mathbf{d})$, is determined by the likelihood $p(\mathbf{d}|\mathbf{m})$, describing the conditional probability of successfully modeling the seismic data given a seismic velocity model, and by our prior knowledge $p(\mathbf{m})$ of the model parameters, which reflects our initial confidence in the unknown model based on any available prior information. Lastly, a normalization constant, also known as the evidence, $p(\mathbf{d})$, ensures that the posterior distribution properly integrates to one over the entire parameter space.

2.1. Variational Inference

At the core of Variational Inference lies the idea of approximating this posterior distribution with a simpler, surrogate distribution, denoted as $q(\mathbf{m})$. This distribution should be selected from a family (called variational family) that is easy to sample and evaluate; a common choice is therefore the Gaussian distribution. The essence of this optimization process lies in minimizing the Kullback-Leibler (KL) divergence^{[46][27]}, which quantifies the discrepancy between the surrogate (approximate) distribution and the target (posterior) distribution. The KL divergence is expressed as:

$$KL(q(\mathbf{m})|p(\mathbf{m}|\mathbf{d})) = E_{\mathbf{m} \sim q}[-\log p(\mathbf{m}|\mathbf{d}) + \log q(\mathbf{m})] \quad (2)$$

By minimizing Equation 2, we are able to compute the expectation of the model parameters sampled from the surrogate distribution. Consequently, VI achieves an approximation of the posterior distribution through the minimization outlined below:

$$q^* = \underset{q}{\operatorname{argmin}} KL(q(\mathbf{m})|p(\mathbf{m}|\mathbf{d})) \quad (3)$$

2.2. Stein Variational Gradient Descent

Stein Variational Gradient Descent represents a distinctive approach in VI. Unlike traditional VI techniques, SVGD is a deterministic, particle-based inference algorithm that iteratively minimizes the KL divergence between the chosen approximate distribution and the target density. This innovative method leverages the concept of functional gradients to effectively transport a predefined set of particles toward the target distribution. The transport occurs within the Reproducing Kernel Hilbert Space (RKHS), guided by the gradient of the KL divergence (for a detailed derivation of the SVGD formulation, we refer the reader to Liu and Wang^[32]).

Given a collection of particles, the optimal update direction ϕ^* of Equation 3, for each particle, is given by:

$$\phi^*(\cdot) = E_{\mathbf{m} \sim q} [\underbrace{k(\mathbf{m}, \cdot) \nabla_{\mathbf{m}} \log p(\mathbf{m}|\mathbf{d})}_{\text{driving force}} + \underbrace{\nabla_{\mathbf{m}} k(\mathbf{m}, \cdot)}_{\text{repulsive force}}] \quad (4)$$

If we denote the particles that we are using to represent q as $\{\mathbf{m}_i\}_{i=1}^N$, the expectation in Equation (4) can be approximated using the sample mean over the particles. Thus, the KL divergence can be iteratively minimized as follows:

$$\begin{aligned} \phi_{q,p}^*(\mathbf{m}) &= \frac{1}{N} \sum_{j=1}^N [k(\mathbf{m}_j^l, \mathbf{m}) \nabla_{\mathbf{m}_j^l} \log p(\mathbf{m}_j^l|\mathbf{d}) + \nabla_{\mathbf{m}_j^l} k(\mathbf{m}_j^l, \mathbf{m})] \\ \mathbf{m}_i^{l+1} &= \mathbf{m}_i^l + \epsilon_l \phi_{q,p}^*(\mathbf{m}_i^l) \end{aligned} \quad (5)$$

where l denotes the current iteration, N is the number of particles, and ϵ_l is the step size. Assuming the step size to be sufficiently small, the process asymptotically converges to the target posterior as the number of particles tends to infinity.

Equation 4 comprises two distinct terms: the driving force and the repulsive force. The driving force aims to direct the particles toward higher probability regions. Conversely, the repulsive force has the crucial role of maintaining particle diversity, and actively prevents particle collapse by dispersing particles across the parameter space. This balance enables a comprehensive exploration and characterization of the target distribution.

Various types of kernel functions have been proposed over the last few years^{[32][47][38]}. In our study, we adopt the two most commonly employed kernels, namely the radial basis function (RBF) and the inverse multi-quadratic (IMQ). The RBF kernel is defined as:

$$k(\mathbf{m}, \mathbf{m}') = \exp\left(-\frac{\|\mathbf{m} - \mathbf{m}'\|^2}{2h^2}\right) \quad (6)$$

where h is the bandwidth, a scaling factor that controls the strength of the interaction between different particles based on their distances. As suggested by Liu and Wang^[32], we set $h = \tilde{d}^2 / \log N$, where \tilde{d} is the median of pairwise distances between all particles. It is worth noting that this parameter is recalculated at each iteration with limited heuristic justification, and there is some evidence that this non-linearity can generate instability. The IMQ kernel, instead, is defined as:

$$k(\mathbf{m}, \mathbf{m}') = \left(c^2 + \frac{\|\mathbf{m} - \mathbf{m}'\|^2}{2h^2}\right)^\beta \quad (7)$$

where c and β are two user-defined parameters. We set c and β to 1 and $-\frac{1}{2}$, respectively, as suggested by Gorham and Mackey^[47].

2.3. Annealed Stein Variational Gradient Descent

Despite the empirical success of SVGD, convergence guarantees are absent — except in the mean-field limit (where the number of particles $n \rightarrow \infty$, while the dimensionality d is kept fixed). Zhuo et al.^[45] showed that SVGD encounters degeneracy issues under finite particle conditions, which cause the particles to collapse into a small number of modes — the so-called mode collapse issue. On the other hand, as dimensionality increases (such that $d > n$), the variance estimated by SVGD may significantly underestimate the variance of the target distribution — a phenomenon known as variance collapse; this is an important concern in high-dimensional problems like FWI.

Ba et al.^[48] compared the SVGD update to the application of gradient descent to a Maximum Mean Discrepancy (MMD) objective function. In a high-dimensional example, they found that SVGD and MMD descent differ primarily in the driving force term that becomes increasingly problematic in higher dimensions. They have empirically demonstrated that removing the bias introduced by the deterministic update present in the driving force leads to more accurate estimation of the variance. Consequently, Ba et al.^[48] proposed modifying the driving force term of SVGD with a damped version, resulting in the damped SVGD approach. At a high level, this modification mirrors the approach taken by D'Angelo and Fortuin^[49] under the name of Annealed SVGD (A-SVGd), where a heuristic temperature parameter $\alpha(l) \in [0, 1]$ is introduced to adjust the intensity of the driving force. The updated rule is then simply expressed as follows:

$$\phi_{q,p}^*(\mathbf{m}) = \frac{1}{N} \sum_{j=1}^N [\alpha(l) k(\mathbf{m}_j^l, \mathbf{m}) \nabla_{\mathbf{m}_j^l} \log p(\mathbf{m}_j^l) + \nabla_{\mathbf{m}_j^l} k(\mathbf{m}_j^l, \mathbf{m})]. \quad (8)$$

Varying the α parameter within the $[0, 1]$ range induces two distinct phases. A first exploratory phase, dominated by a strong repulsive force (α close to 0) that disperses the particles from their initial positions, facilitating broad coverage of the target distribution. This is followed by a second exploitative phase, where the driving force dominates (α close to 1) and concentrates the particle distribution around different modes. The selection of the temperature parameter $\alpha(l)$ is crucial to maintain the convergence properties of SVGD, in order to ensure that the final iterations operate effectively on the target density, i.e. $\lim_{l \rightarrow \infty} \alpha(l) = 1$. In this work we will employ and compare the performances of two different annealing schedules proposed by D'Angelo and Fortuin^[49]. The first schedule is the hyperbolic tangent, defined as:

$$\alpha(l) = \tanh \left[\left(1.3 \frac{l}{N} \right)^p \right] \quad (9)$$

where l is the current iteration, N is the total number of iterations, and p is a user-defined parameter that controls the rate of transition between the two phases. The second is the cyclic schedule, which allows a sequence of exploratory and converging phases and can be defined as:

$$\alpha(l) = \left(\frac{\text{mod}(l, N/C)}{N/C} \right)^p \quad (10)$$

where C is the number of cycles, where the α value ranges from 0 to 1.

2.4. Principal Component Analysis

Principal Component Analysis is a dimensionality reduction technique, particularly effective when dealing with high-dimensional and highly correlated data. The main objective of PCA is to find a smaller set of features that can accurately represent the original data in a lower-dimensional space, while preserving as much information as possible^[50]. PCA operates under the assumptions of linearity (principal components are orthogonal to each other), and can be summarized as follows.

First, let us consider a set of velocity particles \mathbf{X} of size $n \times m$, where n is the number of particles (observations) and m is the number of dimensions (variables). PCA begins by computing the mean vector $\bar{\mathbf{X}} = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i$, which contains the mean of each column and allows us to standardize the data to have a zero mean:

$$\mathbf{Z} = \mathbf{X} - \mathbf{I}_{n \times 1} \bar{\mathbf{X}} \quad (11)$$

Next, we compute the covariance matrix \mathbf{C} of the mean-centered data \mathbf{Z} :

$$\mathbf{C} = \frac{1}{n-1} \mathbf{Z}^T \mathbf{Z} \quad (12)$$

We then perform eigenvalue decomposition on the covariance matrix \mathbf{C} :

$$\mathbf{C}\mathbf{V} = \mathbf{V}\mathbf{\Lambda} \quad (13)$$

where \mathbf{V} is the $m \times m$ matrix of eigenvectors (principal components) and $\mathbf{\Lambda}$ is the $m \times m$ diagonal matrix of eigenvalues. Finally, we sort the eigenvalues in descending order and reorder the eigenvectors accordingly to project the mean-centered data onto the new principal components:

$$\mathbf{Y} = \mathbf{Z}\mathbf{V} \quad (14)$$

where \mathbf{Y} is the $n \times m$ matrix of the transformed data.

In the present study, we apply PCA to the entire set of particles collected per iteration (after SGVD is performed) for a number of components equal to $n-1$, where n is the number of particles. We aim to obtain the explained variance per component, which quantifies the proportion of the total variability in the data that each principal component captures. Given a user-defined number of components, the explained variance for each component is determined by its respective eigenvalue:

$$\text{Explained Variance of the } i\text{-th component} = \frac{\lambda_i}{\sum_{j=1}^n \lambda_j} \quad (15)$$

where λ_i represents the eigenvalue of the i -th component, and $\sum_{j=1}^n \lambda_j$ denotes the sum of all eigenvalues, which accounts for the total variance. Analyzing the explained variance per component helps us understanding the behavior and distribution of the particles.

2.5. Clustering

Given the highly non-linear nature of FWI, the particles obtained at the end of the SVGD iterations may have converged to different modes (i.e., become trapped in various local minima). Whilst the particles that have reached the global minimum are likely to be geologically meaningful, this may not be the case for those that reached other basins of attraction. Consequently, it is important to identify clusters of particles within our high-dimensional model space and distinguish those that are geologically meaningful from the other ones.

In order to do so, we have employed the Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) to perform such a clustering operation^[51]. HDBSCAN is an advanced clustering

algorithm that extends Density-Based Spatial Clustering of Applications with Noise (DBSCAN)^[52]. In short, HDBSCAN uses the Mutual Reachability Distance, a metric that combines a density-based measure with pairwise distances to facilitate meaningful clustering. The algorithm creates a Minimum Spanning Tree (MST) from the mutual reachability distances, assembling the basis of the hierarchical cluster tree. Through Condensed Clustering, HDBSCAN extracts significant clusters by systematically removing edges in the MST. Finally, Stability-based Clustering selects the most stable clusters from the hierarchical tree. Unlike other clustering methods, HDBSCAN is well-suited for high-dimensional data and does not require the user to specify the number of clusters in advance. This flexibility makes HDBSCAN an ideal candidate for our analysis, as it allows for robust identification of clusters and noise.

3. Numerical examples

3.1. Synthetic data

In this section, we present a series of numerical experiments using a portion of the Marmousi model^[53] to evaluate the capability of the previously described SVGD approaches to estimate the uncertainty associated with FWI when working with a limited number of particles. Synthetic data are modeled using Deepwave^[54] in a rectangular uniform grid with dimensions of $n_z = 81$ and $n_x = 216$ and spacing of 20m in both directions, using five shots evenly distributed along the horizontal axis and recorded by 201 receivers placed at 1-meter intervals. The seismic source is modeled using a Ricker wavelet with a peak frequency of 7 Hz. The registration time is set to 3s and the sample interval used for the forward modelling is 1ms. To simulate more realistic conditions, colored noise matching the frequency spectrum of the signal is added to the synthetic data, resulting in noisy observed data with a Signal-to-Noise Ratio (SNR) of approximately 17 dB (see Figure 1).




Figure 1. Portion of the Marmousi model used in our numerical experiments. The left panel illustrates the velocity model in meters per second (m/s), with red dots indicating shot locations and white dots representing receivers. Vertical white lines represent pseudo-well logs, and black circles indicate pixel locations for marginal plots discussed in the results section and appendix. The right panel shows the shot gather for the source in the middle of the model.

3.2. Selection of hyperparameters

The primary objective of this work is to investigate the uncertainty associated with the modeling operator (the likelihood is assumed to be Gaussian) by deliberately excluding any influence of our prior knowledge, rather than imposing a uniform bound limits on velocities. Since the uncertainty associated with the FWI problem is expected to encapsulate both the scattering (high wavenumber) and transmission (low wavenumber) components of the model, the choice of perturbations of the initial particles is crucial. Drawing inspiration from Izzatullah et al.^[44], we aim to explore the uncertainties of these components by generating Gaussian Random Field (GRF) perturbations that introduce variability in both amplitude and scale.

In terms of the hyperparameters employed in this study, we opt for the Gaussian RBF and IMQ kernels. Bandwidth selection is performed using both the median trick and a fixed constant value. The constant value of the bandwidth is determined after a meticulous analysis of the bandwidth evolution in the scenario where the median trick is used. Finally, we use the Adam optimizer to update the particles at each iteration, given their gradient in Equation 5, with a constant learning rate of 100 and a fixed number of iterations equal to 600.

The experiments are conducted in two particle sets: one with a small number of particles (50) and the other with a larger number (200). To mitigate the effects of mode and variance collapse, we evaluate both the vanilla and annealed SVGD approaches. For the annealed SVGD, we consider two formulations for the

temperature parameter, namely hyperbolic and cyclic. For the hyperbolic formulation (Eq. 9), we set $p = 3$ and maintained $\alpha = 1$ for the last 20% of iterations. For the cyclic formulation (Eq. 10), we set $p = 2$ and $C = 8$, ensuring that the temperature remains at $\alpha = 1$ during the last two cycles.

The overall performance in terms of data misfit (L2 norm) and SNR for the various experiments with different hyperparameters is presented in Figures A1 and A2. The subsequent results and discussion focus on the vanilla and annealed SVGD (tanh) formulations, utilizing the RBF kernel with the median trick and a set of 200 particles, as these configurations showed the best performance.

3.3. Single-scale experiments

In the first set of experiments, we perform FWI using a single frequency band with a peak frequency of 7 Hz. Initial particles are generated by applying GRF perturbations with different variances to a highly smoothed version of the true Marmousi model. During the inversion process, we impose lower and upper-velocity bounds, allowing velocity values between 1500 m/s and 4370 m/s.

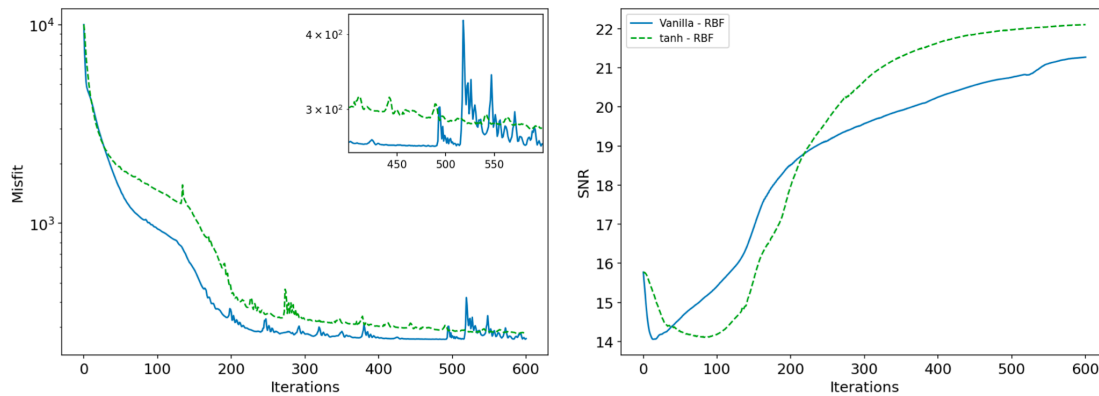


Figure 2. Data misfit (left) and SNR with respect to the true model (right) for vanilla and annealed (tanh) SVGD in the single-frequency scenario using 200 particles. The zoom window provides a clearer misfit comparison for the final 150 iterations.

Figure 2 illustrates the performances of the different methods in terms of data misfit and SNR with respect to the true model, computed from the mean over iterations for the various experiments. The vanilla and annealed SVGD with hyperbolic tangent demonstrate superior performance compared to the other methods (see Appendix A). Notably, the annealed SVGD using the cyclic formulation results in poorer data misfit and SNR; this may be due to its design, which aims to explore the parameter space

better and identify widely separated high-probability regions. As a result, the mean model may be perturbed, fitting the data less precisely compared to a more compact, but less comprehensive, posterior distribution.

A more detailed comparison of the mean and standard deviation for the 200-particle experiments is presented in Figures A3 and A4, respectively. Figure 3 shows the mean obtained after 600 iterations for the vanilla and annealed versions for the 200-particle experiment, along with the mean of the initial distribution. All predicted models are relatively similar in the shallower part (<1 km) and reproduce most of the main features of the true model. However, the model reconstruction is poorer in the deeper part, likely due to a lack of illumination (especially near the edges of the model). As stated before, the results obtained using the vanilla SVGD and the annealed SVGD with the hyperbolic tangent formulation (both with the RBF kernel) exhibit higher SNR and are less affected by artifacts. In contrast, the models predicted using other approaches show errors in terms of velocity magnitudes, particularly in the last 500 meters of depth.




Figure 3. Mean and standard deviation comparison of the experiments using 200 particles for vanilla SVGD and annealed SVGD using RBF kernel and median trick after 600 iterations in the single-scale scenario. The velocity values are expressed in m/s.

In addition, Figure 3 presents the standard deviation maps (expressed in m/s) obtained after 600 iterations using 200 particles. Both maps exhibit a similar expected pattern, with very low values in the shallower parts where illumination is greater and the values increasing towards the deeper portions of the model. These deeper regions, along with the lateral edges of the model, are expected to have higher uncertainties due to limitations in acquisition geometry and the physics of wave propagation within the

subsurface. Higher uncertainties are also observed in areas of high velocity and near the main velocity contrasts.

The standard deviation associated with the initial distribution is relatively low (< 150 m/s) across the entire model. We select a narrow proposal for the initial samples to prevent SVGD from rapidly repelling particles into undesirable modes at the early stages (particles converging to suboptimal local minima). By maintaining an initial low standard deviation, we aim to direct the convergence toward fewer and more geologically consistent local minima. In contrast, the values associated with the predicted models are substantially higher, reaching over 800 m/s. Notably, the standard deviation maps associated with the annealed version of SVGD show significantly higher values throughout the model compared to those associated with the vanilla SVGD. This indicates that the annealed approach allows for better exploration of the model space and reduces the variance collapse phenomenon affecting the vanilla SVGD.

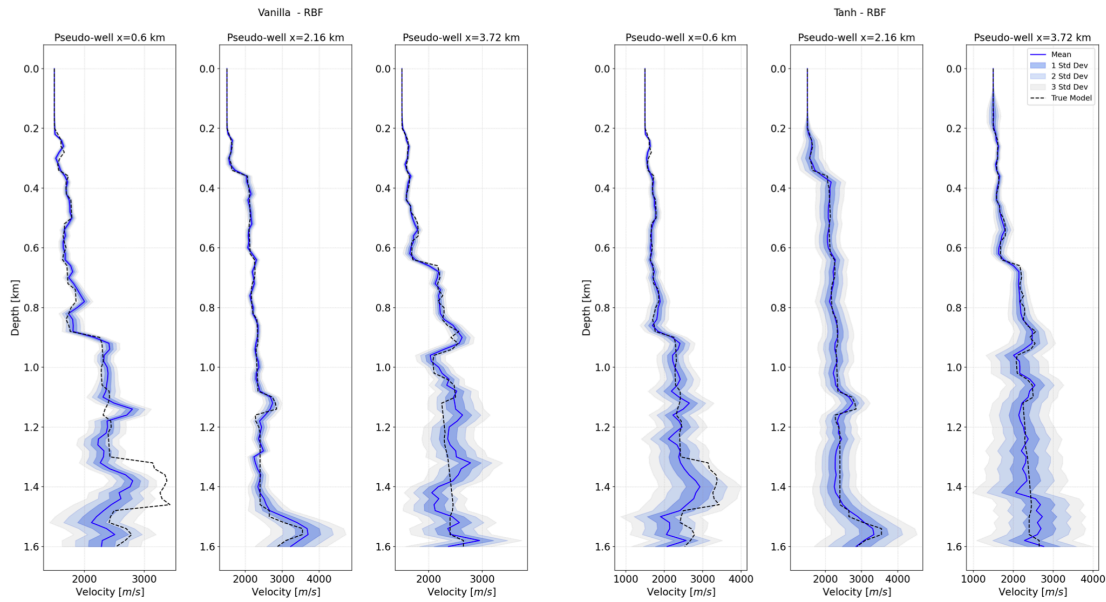


Figure 4. Pseudo-well marginal distributions showing mean and three confidence intervals for the experiment using vanilla SVGD (left) and annealed SVGD (right), using the RBF kernel with median trick and 200 particles in the single-frequency scenario.

Moreover, to better highlight the differences between the vanilla and annealed approaches, in Figures 4 (left) and 4 (right) we present three pseudo well logs corresponding to three distinct spatial locations, as indicated by the white vertical lines in Figure 1. For each position, we illustrate the true velocity varying

with depth, the velocities extracted from the mean model obtained using both approaches, and three confidence intervals (corresponding to one, two, and three standard deviations) based on the standard deviation maps shown in Figure A4. We observe that in both cases, the width of the confidence intervals increases with depth, indicating larger uncertainty at greater depths, as expected. Within the first 800 meters of depth, the mean model closely resembles the true model, with standard deviation values close to zero. The main differences between the two approaches become more evident at greater depths, where discrepancies between the predicted and true models are more significant. Specifically, the logs associated with the annealed SVGD show larger confidence intervals, with the values extracted from the true model falling almost entirely within these bounds. In contrast, the vanilla SVGD approach yields smaller confidence intervals, with the true values occasionally falling outside these bounds (e.g., the pseudo well at the spatial position of 0.6 km and around 1.4 km of depth). This comparison indicates that, while the annealed approach captures a broader range of uncertainty, the vanilla approach may sometimes underestimate the true variability at greater depths. Similar results are also obtained at three different pixel locations for vanilla (see figure 5) and annealed formulations (see figure 6).




Figure 5. Single-scale scenario; pixels marginals for vanilla SVGD with RBF kernel and 200 particles.




Figure 6. Single-scale scenario: pixels marginals for annealed SVGD (tanh) with RBF kernel, and 200 particles.

To further investigate the behavior and distribution of particles during the SVGD optimization process, we performed PCA. This analysis helps us understand the explained variance of the components, providing insights into how the particles evolve throughout the optimization. By examining the explained variance for $n - 1$ components (where n is the number of particles), it is possible to identify the dominant directions in which the particle positions vary the most. This allows us to understand the main variances captured by components and the structural dynamics of the particle distribution.

Figure 7 (left) illustrates the PCA variances for the vanilla SVGD case with 200 particles using the RBF kernel. We observe a relatively uniform spread of the components over the variance but erratic convergence, with a significant portion of the components converging towards mid- and low-variance values. In contrast, Figure 7 (right) shows the PCA for the annealed SVGD with the tanh formulation, which demonstrates a more stable convergence pattern and a more balanced distribution of variance between high and low variance directions. These PCA results highlight the effectiveness of the annealed SVGD to provide a more stable and informative representation of the particle dynamics compared to the vanilla SVGD approach.




Figure 7. Explained variance per component (using $n - 1$ components, where n is the particle number) for the experiments using (left) vanilla SVGD with the RBF kernel and (right) annealed SVGD with the tanh formulation and RBF kernel with 200 particles.

3.4. Multi-scale experiments

In this section, we present results for the multi-scale FWI experiments, focusing exclusively on the Vanilla and Annealed variants of SVGD (see Figure 8), as these have demonstrated superior performance in our single-frequency experiments. The multi-scale approach (introduced by Bunks et al.^[55]) has become a standard practice in FWI due to its ability to improve the inversion's convergence and accuracy, mitigating the cycle-skipping issue. It begins with the lowest frequencies available in the observed data and progressively incorporates higher frequencies, thereby mitigating the non-linearity of the inversion process by initially targeting large-scale features, which are more sensitive to low frequencies, and then refining the model with higher frequencies to capture higher-resolution details. Specifically, the inversion process is conducted in three stages: 200 iterations at a peak frequency of 4 Hz, 200 iterations at 7 Hz, and 200 iterations at 10 Hz, with learning rates of 100, 10, and 10, respectively. This step-wise frequency escalation ensures that each scale of the model is accurately resolved before moving on to the next, providing a robust framework for the inversion process.




Figure 8. Data misfit (left) and SNR (right) across different experiments for 200 particles in the multi-scale scenario. The close-up window provide a clearer comparison of the final 150 iterations.

As in the single-frequency experiment, the multi-scale approach using annealed SVGD demonstrates superior performance compared to the vanilla formulation in terms of SNR with respect to the true model (see Figure 8). It is essential to highlight that multi-scale SVGD not only achieves higher SNR values but also yields more meaningful and representative statistics compared to the single-frequency approach. This improvement is primarily due to the enhanced control over the particle refinement process, which prevents some particles from diverging and causing high standard deviation values. Consequently, the optimization process is more likely to converge toward geologically representative models, leading to more accurate and reliable results.

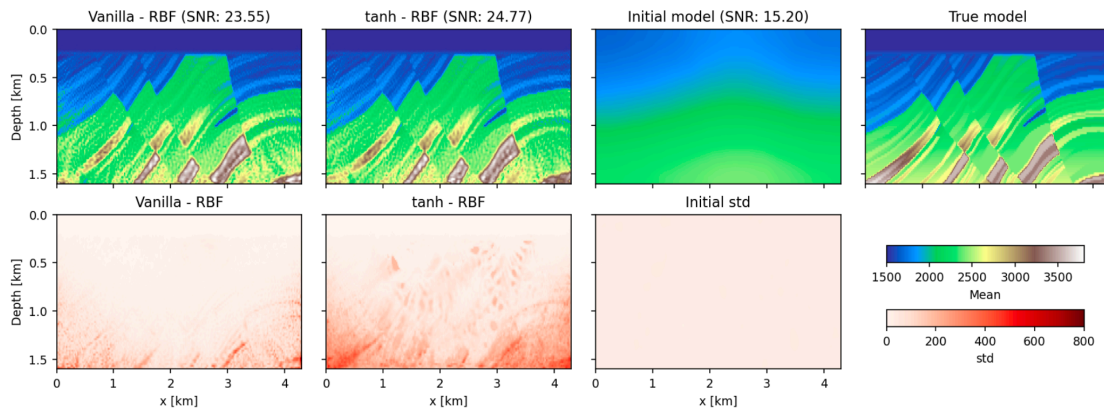


Figure 9. Mean and standard deviation comparison of the experiments using 200 particles for vanilla SVGD and annealed SVGD using RBF kernel and median trick after 600 iterations in the multi-scale scenario. The velocity values are expressed in m/s.

Appendices A and B provide supplementary results for both the single-scale and multi-scale approaches, respectively, which further strengthen the evidence supporting our findings.

3.5. Cluster Analysis and Statistical Evaluation

Our single-frequency experiments yield several key observations. A small number of principal components explain the majority of the variance (Figure 7, left). Consequently, many components capture small variances rather than meaningful patterns in the data. This is further supported by the presence of abnormal individual particles (see Appendix A) and artifacts in the shallower parts of the particles, which are evident in the respective mean and standard deviation maps (Figures A3 and A4). These observations suggest that the particles tend to converge to different modes, some of which may not be geologically meaningful. Although these particles fit the data term, they do not accurately represent the subsurface structure. Therefore, it is crucial to perform clustering to identify the presence of different modes and conduct statistical analysis within each cluster, rather than assuming all samples have converged to the same global minima.




Figure 10. HDBSCAN clustering for the experiment using annealed SVGD with tanh formulation and RBF kernel with 200 particles. The left panel displays clusters obtained in high-dimensional data, plotted after dimensionality reduction to two components. The right panel shows the distribution of each cluster after computing the SNR.

To illustrate the importance of identifying clusters in the final particles and use such information for any subsequent statistical analysis, we consider the experiment with annealed SVGD employing a tanh temperature parameter and RBF kernel for a set of 200 particles. After the optimization process (600 iterations), we apply HDBSCAN to the final particles. In this case, the clustering algorithm produces three distinct groups, labeled as -1, 0, and 1, of sizes 35, 163, and 2, respectively. For visualization purposes, we plot the particles in the 2D space defined by the first two components of the PCA and color-code the particles to indicate which cluster they belong to. (Figure 10 left). We then compute the SNR for each individual particle in the different clusters and display their distribution in Figure 10 (right). Overall, cluster 0 and cluster 1 contain particles with higher SNR values, although some of the particles in cluster 0 converge to similar SNR values. More importantly, cluster -1 is identified as the noisy cluster, representing particles not assigned to any other cluster, to be considered as outliers and geologically implausible. (Figure 11).

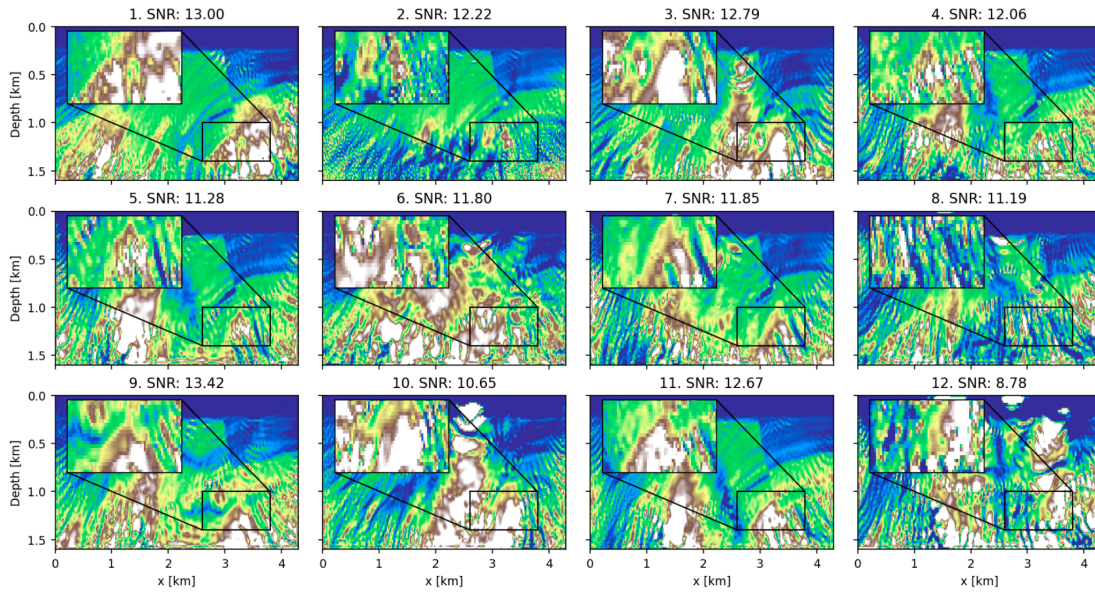


Figure 11. Noisy particles corresponding to cluster -1 for the experiment with annealed SVGD using the tanh formulation and RBF kernel with 200 particles.

Further analysis of the mean and standard deviation for each cluster (Figure 12) reveals that while the mean of cluster -1 appears reasonable, the individual particles are not representative, which leads to significant variations in the shallow parts of the model and, therefore, high uncertainty. This observation correlates with the artifacts detected in the shallow regions of the results in subsection 3.3. In contrast, clusters 0 and 1 display more consistent and plausible mean and standard deviation patterns, similar to those obtained with the multi-scale approach. This underscores the importance of performing clustering analysis of the particles produced by SVGD in order to discard non-representative particles and thereby obtain more accurate and geologically plausible outcomes.




Figure 12. Cluster statistics for the experiment with annealed SVGD using the tanh formulation and RBF kernel with 200 particles.

4. Discussion

In the context of FWI, addressing uncertainty quantification presents significant computational challenges. Given the high-dimensionality of our model space, we operate in a regime where the number of particles is much smaller than the number of unknown parameters, and therefore, it is only possible to provide a low-rank approximation of the posterior covariance, with the rank limited to at most n (particles) $- 1$. A key aspect of variance collapse occurs when the actual rank falls below this theoretical limit. This limitation prevents us from fully capturing uncertainty across all directions, thus our analyses produce only relative — but still meaningful — uncertainty estimates.

The SVGD algorithm stands out due to its flexibility, enabling optimization problems to be solved through standard gradient descent algorithms for a certain number of particles, whilst introducing inter-particle communication. However, it is crucial — yet laborious — to address the complexities of hyperparameter tuning. In our study, we experimented with two variants of SVGD (vanilla and annealed), Gaussian Random Field (GRF) perturbations to build initial particles, Gaussian kernels (RBF and IMQ), different bandwidth selection strategies (median trick and constant value), a constant learning rate, and varying particle counts. Moreover, our experimental evaluations were confined to single-scale and multi-scale scenarios and only assessed the uncertainties associated with the data misfit term (modeling operator). This methodology has the potential to produce nuanced standard deviation maps of velocities, though the integration of more informative prior information remains a subject for future exploration.

Our primary motivation was to apply annealed SVGD to mitigate mode- and variance-collapse issues that affect the vanilla SVGD approach. In the single-scale scenario, our findings reveal that annealed SVGD with the tanh formulation provide better model estimates (i.e., higher SNR, lower data misfit, and overall more meaningful uncertainty estimates) than the vanilla formulation. Higher standard deviation values are associated with high-velocity layers and areas of poor coverage due to the limited acquisition geometry. The vanilla approach yields smaller confidence intervals, sometimes following a different trend than the true model. Conversely, the annealed approach with the tanh formulation captures a broader range of uncertainty and increases the standard deviation values throughout the model. The single-scale outcomes present undesirable artifacts in the shallower parts, which should theoretically be well-illuminated areas. Under such conditions, there is no guarantee that all particles will converge to a unique local minimum. When projecting the particles onto a subspace spanned by $n - 1$ components, we observe that the annealed version recovers higher variances per component (reducing the variance collapse issue). Visualization of individual particles shows some particles fitting the data term but not representing the subsurface, which confirms the hypothesis that particles converge to different modes.

In the multi-scale scenario, both convergence and exploration are significantly enhanced. The mean and standard deviation estimates are improved and the shallower artifacts observed in the single-scale scenario are absent. This suggests that a sequential approach from low to high frequencies may mitigate, though not entirely eliminate, the mode-collapse issue. For what concerns the variance-collapse issue, the annealed version with the tanh formulation yields more reasonable standard deviation maps, as indicated by a more significant number of components explaining the majority of the data variance.

Given the non-linearity and high-dimensionality of the problem, it is challenging to ensure that all particles belong to a single mode. Therefore, it is essential to perform clustering analysis, regardless of the scenario. We opted for HDBSCAN due to its applicability in high dimensions. HDBSCAN enables the easy discovery of a noise cluster composed of non-geological particles. Independent statistical analysis per cluster in the single-scale scenario produced mean and standard deviation maps that are more realistic and do not suffer from shallower artifacts. The primary goal of using clustering analysis on the final set of particles as a post hoc technique is to quickly identify different modes and geologically meaningful particles. We prioritize this approach over incorporating a prior term, which can be mathematically challenging to formulate for filtering out geologically implausible features. A promising future direction could involve training a generative adversarial network (GAN) or variational autoencoder (VAE) to incorporate such a prior into the optimization process^[43].

While annealed SVGD does not entirely solve the problem of variance collapse, it seems a promising method for FWI, where gradient computations are usually computationally expensive. Furthermore, when annealed SVGD is combined with a multi-scale FWI scenario, reasonable estimates could be obtained with fewer gradient evaluations (iterations). This study highlights the potential of advanced SVGD methods to improve the reliability of FWI.

One promising direction for future work involves applying SVGD in a reduced or projected space^{[56][57]} to decrease the number of unknowns in our inverse problem and assess the impact on computational efficiency, convergence rates and quality of the inversion results. This approach would align with the ideal conditions for SVGD, i.e., when the number of particles equals or exceeds the number of unknowns. For example, a potential solution could be to run the SVGD FWI algorithms in a DCT compressed domain, following the works of Aleardi^[20], Berti et al.^[22] and Berti et al.^[58], who found that model compression through DCT effectively reduces the dimensionality of the problem. Also, we could use variational auto-encoders to compress both the model and data and apply SVGD in the compressed (latent) domain^[59]. This exploration holds the potential to significantly advance the application of SVGD in FWI, which could pave the way for more accurate and computationally feasible seismic imaging techniques.

5. Conclusions

This study demonstrates that annealed SVGD can significantly improve convergence and performance compared to vanilla SVGD in FWI applications, in scenarios where the number of particles is much smaller than the number of unknown parameters. Specifically, the annealed SVGD with the tanh formulation enhances the accuracy of mean estimates, leading to higher SNR, lower data misfit, and more reasonable standard deviation maps and thereby mitigating — but not eliminating — the variance-collapse issue. Additionally, applying multi-scale FWI with SVGD yields better mean and standard deviation estimates compared to the single-frequency scenario. Also, combining multi-scale FWI with annealed SVGD yields superior performance. The use of Principal Component Analysis to explain variance by component provides valuable insights into the behavior of the samples during the optimization process. This method provides additional understanding of how different components contribute to the overall variance to give a clearer picture of the sample distribution and convergence patterns. Finally, given the inherent complexity and high dimensionality of the problem, it is crucial to account for the possibility of multiple modes in the solution space. Consequently, statistical analysis should be conducted to identify clusters of particles; HDBSCAN is recommended for its ability to handle

high-dimensional data and identify noise clusters composed of non-geological particles. These strategies collectively offer a more robust and insightful approach to uncertainty analysis in FWI as they enhance the reliability of the results, provide a deeper understanding of the subsurface, and ultimately aid more informed decision-making in industrial applications.

Appendixes

Appendix A — Supplementary Results for Single-scale Experiments




Figure A1. Data misfit for 50 particles (left) and 200 particles (right) across different experiments in the single-frequency scenario. The zoom window provides a clearer misfit comparison of the final 150 iterations.




Figure A2. SNR for 50 particles (left) and 200 particles (right) across different experiments in the single-frequency scenario.

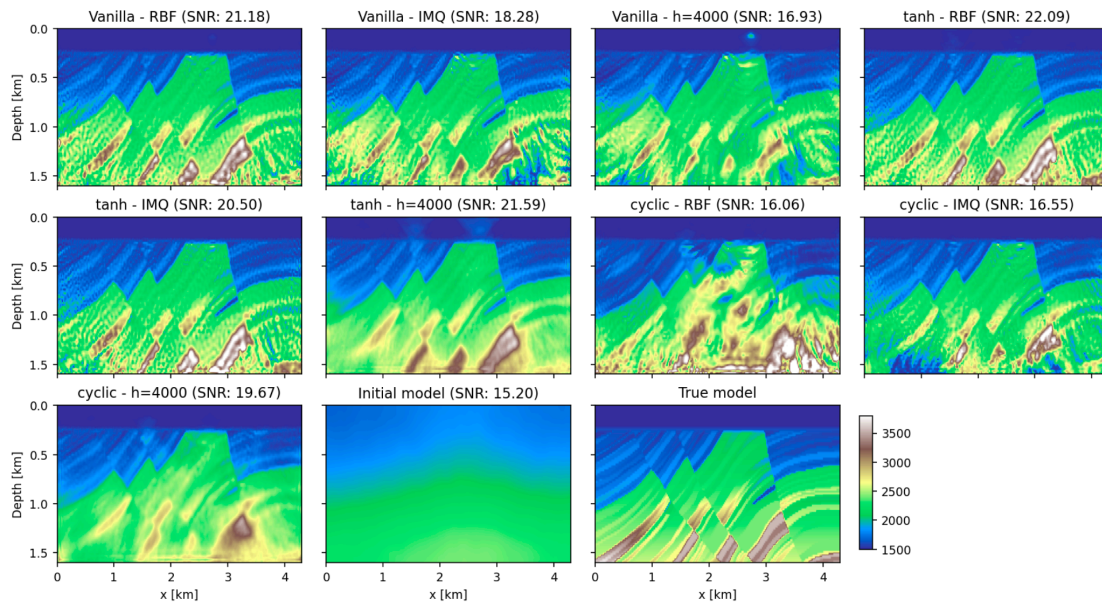


Figure A3. Comparison of the mean models from different experiments using 200 particles, highlighting various SVGD variants and hyperparameters after 600 iterations in the single-frequency scenario. The velocity values are expressed in m/s




Figure A4. Comparison of standard deviation from different experiments using 200 particles, highlighting various SVGD variants and hyperparameters after 600 iterations in the single-frequency scenario. The velocity values are expressed in m/s

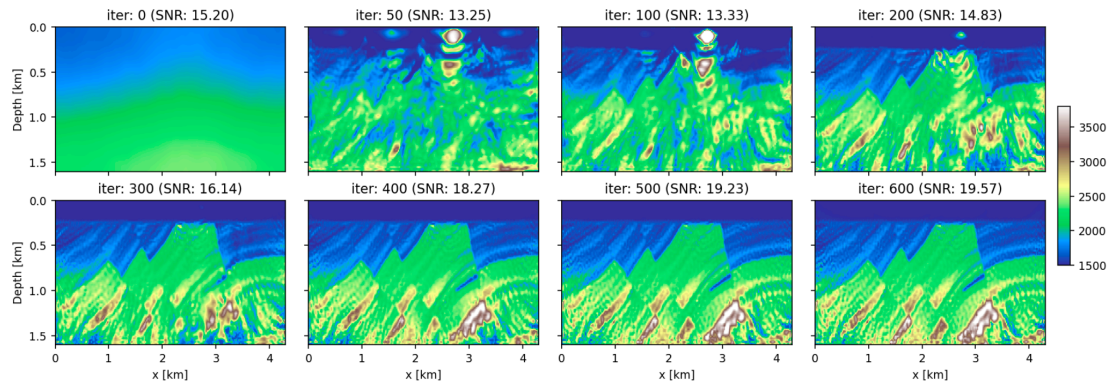


Figure A5. Evolution of the mean for the 200-particle experiment using Vanilla SVGD with RBF kernel.




Figure A6. Evolution of the standard deviation for the 200-particle experiment using Vanilla SVGD with RBF kernel.

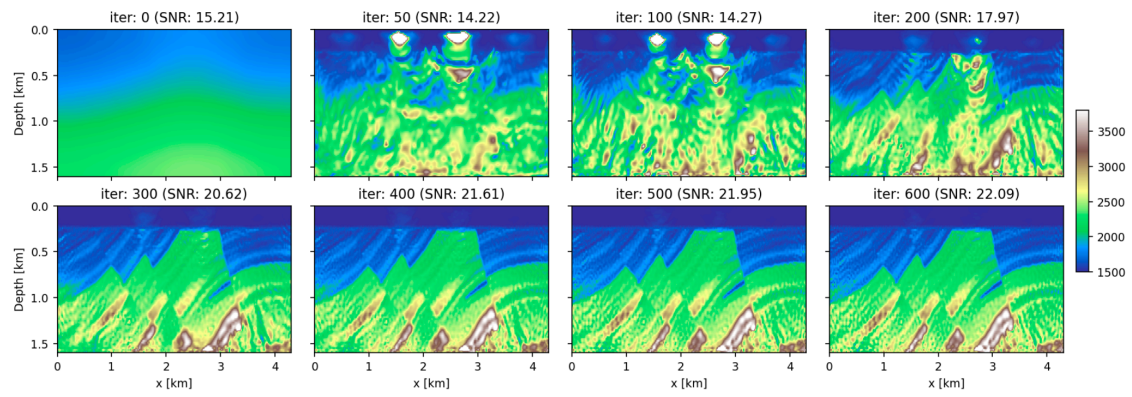


Figure A7. Evolution of the mean for the 200-particle experiment using Annealed SVGD (tanh) with RBF kernel.




Figure A8. Evolution of the standard deviation for the 200-particle experiment using Annealed SVGD (tanh) with RBF kernel.

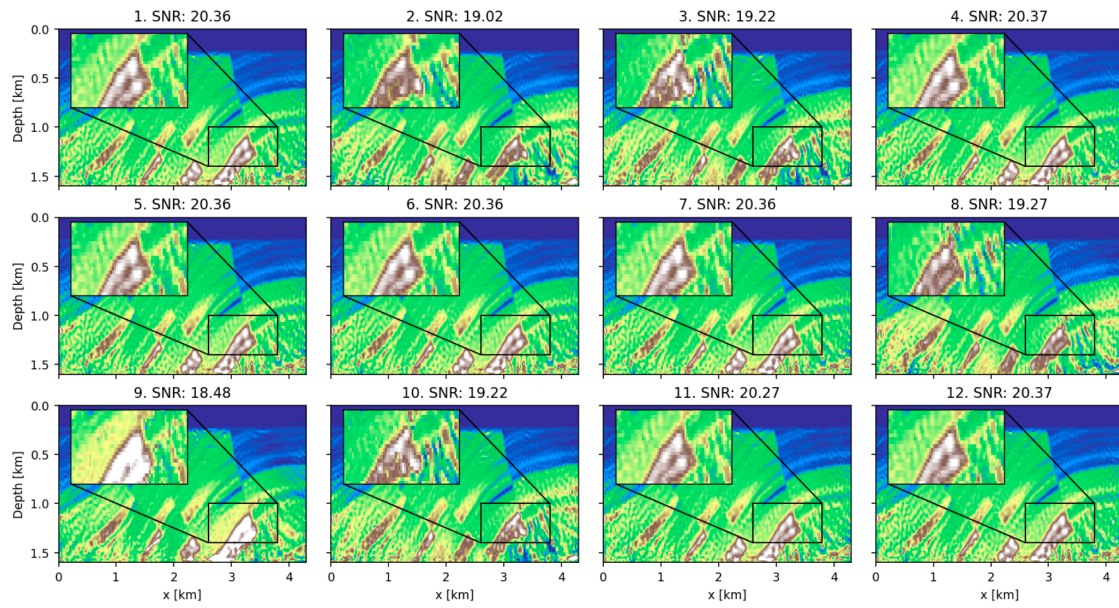


Figure A9. Single-frequency scenario: visualization of 12 particles from a 200-particle experiment using vanilla SVGD with RBF Kernel after 600 iterations.




Figure A10. Single-frequency scenario: visualization of 12 particles from a 200-particle experiment using annealed SVGD (tanh) with RBF Kernel after 600 iterations.

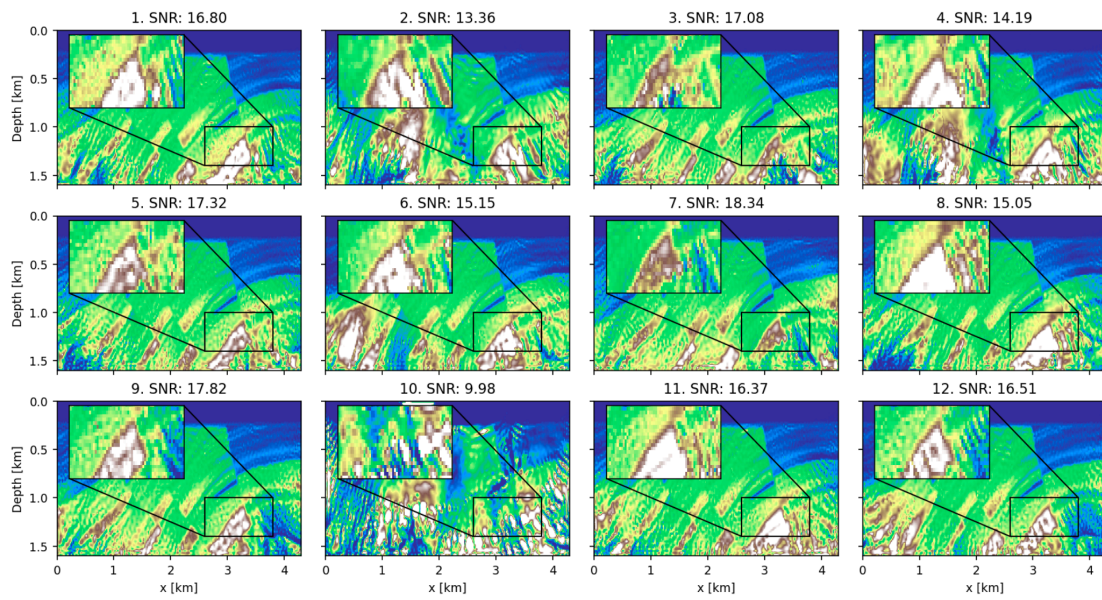


Figure A11. Single-frequency scenario: Visualization of 12 particles from a 200-particle experiment using Annealed SVGD, (tanh) with RBF kernel and constant bandwidth ($h = 4000$) after 600 iterations

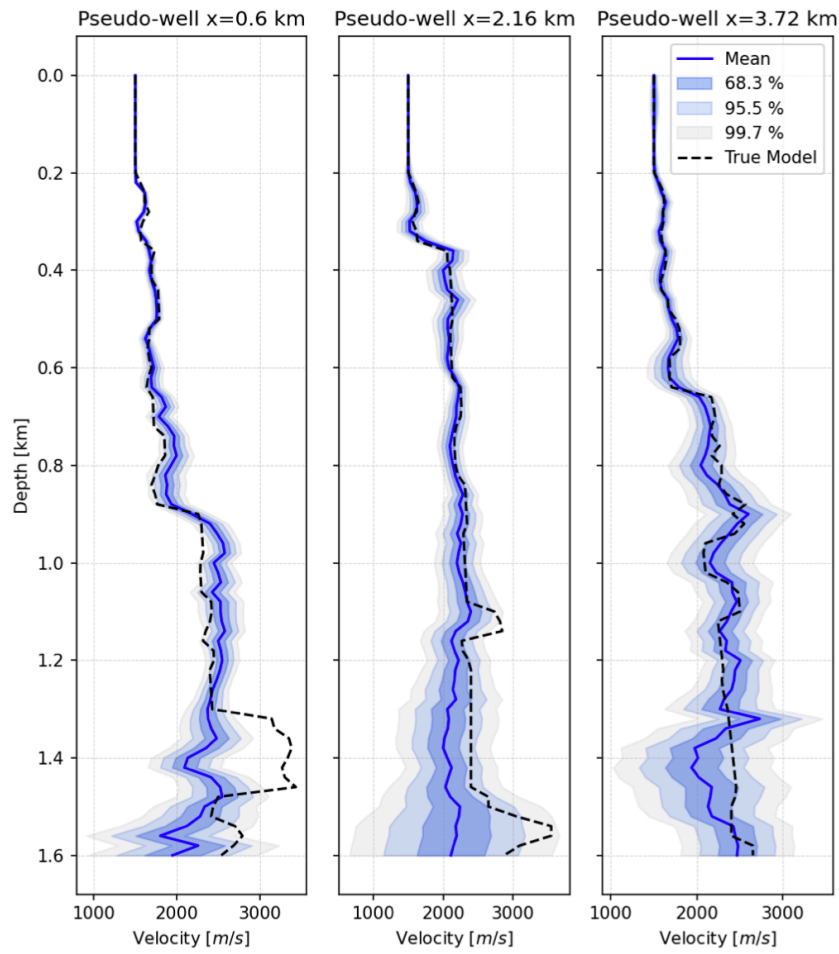


Figure A12. Single frequency scenario: pseudo-wells marginals from a 200-particle experiment using vanilla SVGD with RBF kernel and constant bandwidth ($h = 4000$) 200 particles.




Figure A13. Single frequency scenario: pseudo-wells marginals from a 200-particle experiment using annealed SVGD (tanh) with RBF kernel and constant bandwidth ($h = 4000$) 200 particles.




Figure A14. Single frequency scenario: pixels marginals for vanilla SVGD with RBF kernel and constant bandwidth ($h = 4000$), and 200 particles.




Figure A15. Single frequency scenario: pixels marginals for annealed SVGD (tanh) with RBF kernel and constant bandwidth ($h = 4000$), and 200 particles.

Appendix B — Supplementary Results for Multiscale Experiments




Figure B1. Multiscale scenario: Mean evolution for the 200-particle experiment using vanilla SVGD with RBF kernel.




Figure B2. Multiscale scenario: Standard deviation evolution for the 200-particle experiment using vanilla SVGD with RBF kernel.

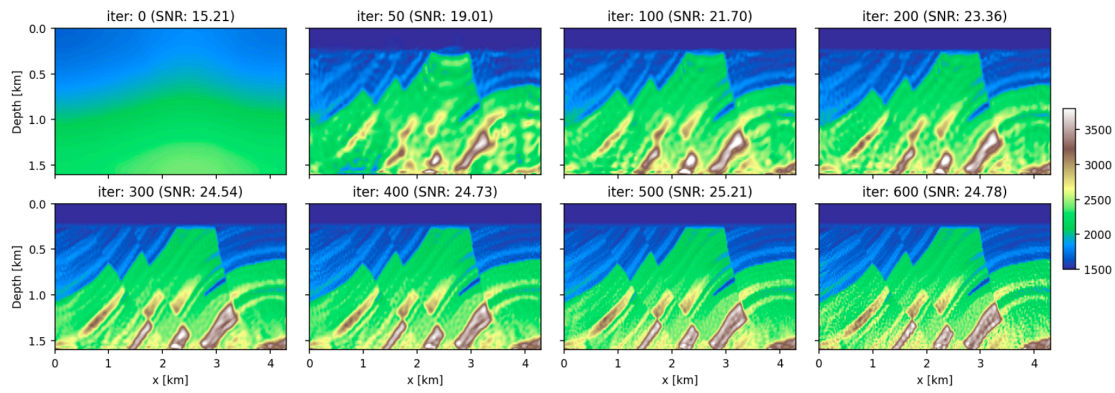


Figure B3. Multiscale scenario: Mean evolution for the 200-particle experiment using annealed SVGD (tanh) and RBF kernel.




Figure B4. Multiscale scenario: Standard deviation evolution for the 200-particle experiment using annealed SVGD (tanh) and RBF kernel.




Figure B5. Multiscale scenario: visualization of 12 particles from a 200-particle experiment using vanilla SVGD with RBF Kernel after 600 iterations.

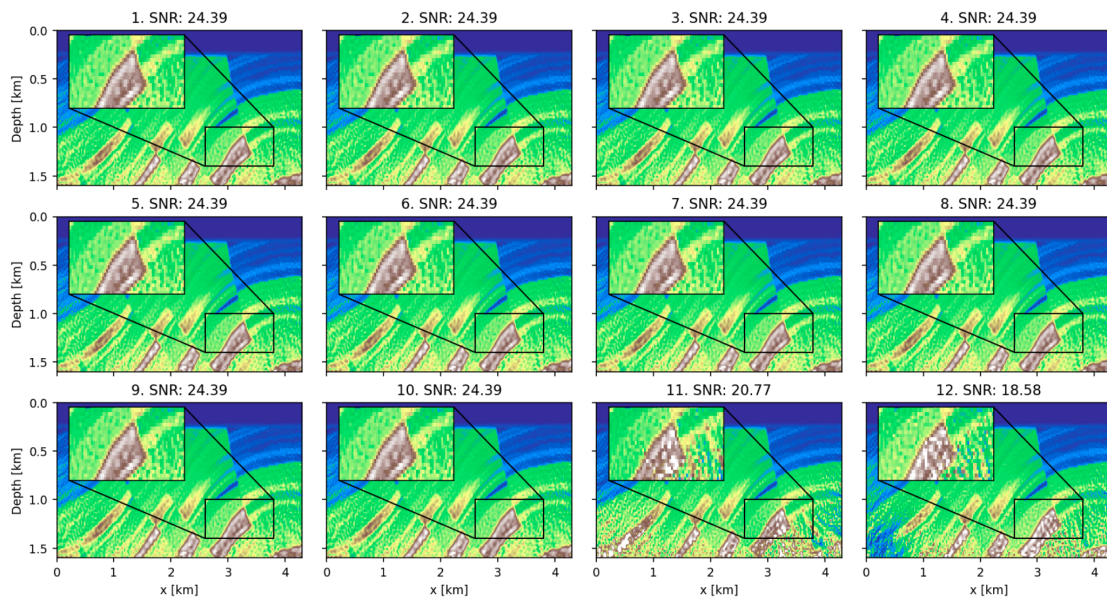


Figure B6. Multiscale scenario: visualization of 12 particles from a 200-particle experiment using annealed SVGD (tanh) with RBF Kernel after 600 iterations.

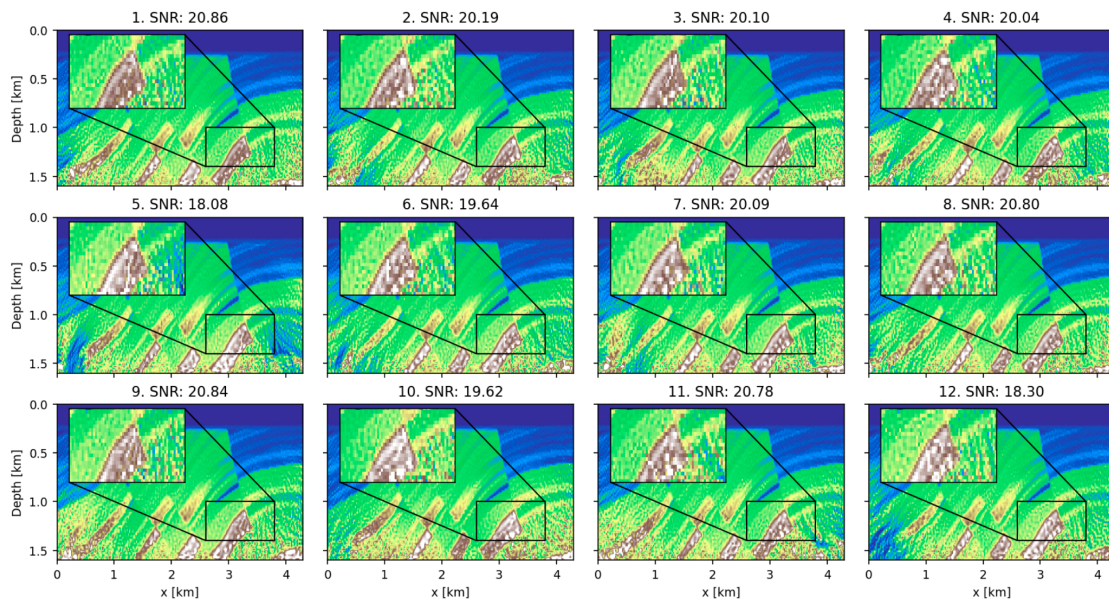


Figure B7. Multiscale scenario: visualization of 12 particles from a 200-particle experiment using annealed SVGD (tanh) with RBF kernel and constant bandwidth ($h = 2500$) after 600 iterations.

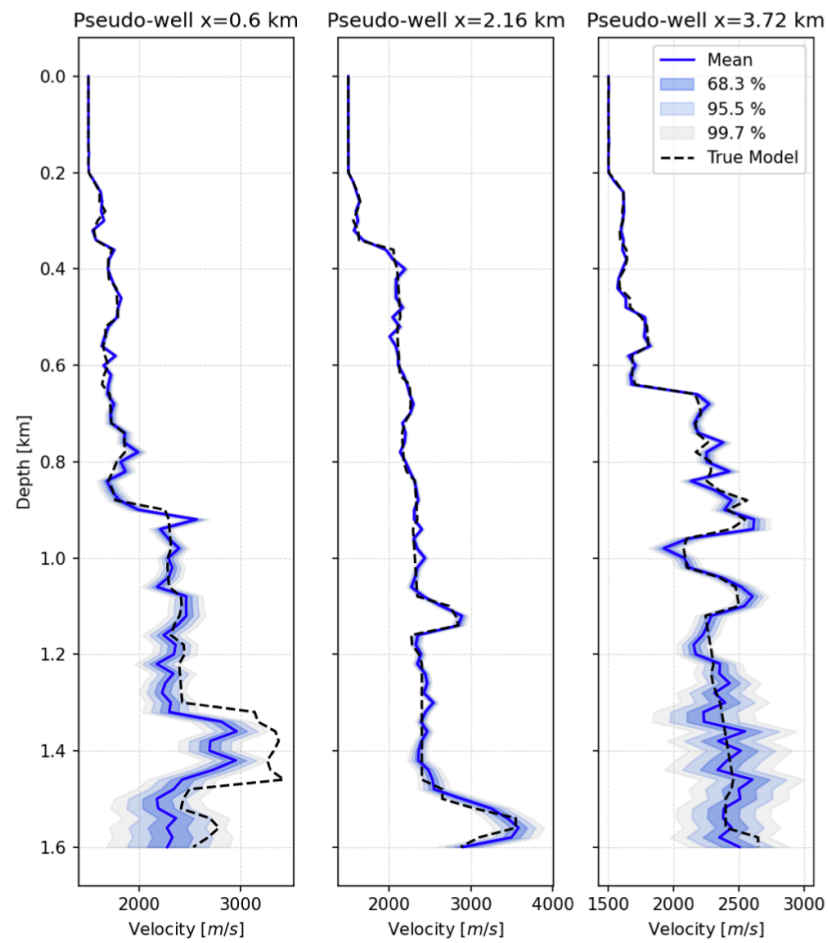


Figure B8. Multiscale scenario: pseudo-wells marginals from a 200-particle experiment using vanilla SVGD with RBF kernel.




Figure B9. Multiscale scenario: pseudo-wells marginals from a 200-particle experiment using vanilla SVGD with RBF kernel and constant bandwidth ($h = 2500$).




Figure B10. Multiscale scenario: pseudo-wells marginals from a 200-particle experiment using annealed SVGD (tanh) with RBF kernel.




Figure B11. Multiscale scenario: pseudo-wells marginals from a 200-particle experiment using annealed SVGD (tanh) with RBF kernel and constant bandwidth ($h = 2500$).




Figure B12. Multiscale scenario; pixels marginals for vanilla SVGD with RBF kernel, 200 particles.




Figure B13. Multiscale scenario: pixels marginals for vanilla SVGD with RBF kernel and constant bandwidth ($h = 2500$), 200 particles.




Figure B14. Multiscale scenario: pixels marginals for annealed SVGD (tanh) with RBF kernel, 200 particles.




Figure B15. Multiscale scenario: pixels marginals for annealed SVGD (tanh) with RBF kernel and constant bandwidth ($h = 2500$), 200 particles.

Statements and Declarations

Acknowledgments

This research was supported by King Abdullah University of Science and Technology (KAUST) and the DeepWave consortium. MC gratefully acknowledges TotalEnergies for the opportunity to conduct an internship in Pau, France, and SB extends thanks to KAUST for hosting an internship in Saudi Arabia. MC and SB equally contributed for this work.

Data Availability

The Marmousi dataset utilized in this study is publicly accessible through the SEG Wiki page, available at https://wiki.seg.org/wiki/Open_data.

References

1. [△]Virieux J, Operto S (2009). "An overview of full-waveform inversion in exploration geophysics". *GEOPHYSICAL*. 74 (6): WCC1–WCC26. doi:10.1190/1.3238367. Link.
2. [△]Tarantola A (2005). "Inverse Problem Theory and Methods for Model Parameter Estimation". Society for Industrial and Applied Mathematics. doi:10.1137/1.9780898717921. Link.
3. [△]Lailly P, Santosa F (1984). "Migration methods: partial but efficient solutions to the seismic inverse problem." In *Inverse problems of acoustic and elastic waves*, Vol. 51, pp. 1387–1403. Note: Publisher: SIAM Philadelphia Cited by: \$1.
4. [△]Tarantola A (1984). "Inversion of seismic reflection data in the acoustic approximation". *GEOPHYSICS*. 49 (8): 1259–1266. doi:10.1190/1.1441754. Link.
5. [△]Bozdağ E, Trampert J, Tromp J (2011). "Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements." *Geophysical Journal International*. 185 (2): 845–870. doi:10.1111/j.1365-246X.2011.04970.x.
6. [△][△]Guo P, Visser G, Saygin E (2020). "Bayesian trans-dimensional full waveform inversion: synthetic and field data application." *Geophysical Journal International*. 222 (1): 610–627. doi:10.1093/gji/ggaa201.
7. [△]Luo Y, Schuster GT (1991). "Wave-equation travelttime inversion". *GEOPHYSICS*. 56 (5): 645–653. doi:10.1190/1.1443081. Link.

8. [△]Brossier R, Operto S, Virieux J (2010). "Which data residual norm for robust elastic frequency-domain full waveform inversion?." *GEOPHYSICS*. 75 (3): R37–R46. doi:10.1190/1.3379323.
9. [△]Warner M, Guasch L (2014). "Adaptive Waveform Inversion – FWI Without Cycle Skipping – Theory". Vol. 2 014, pp. 1–5. doi:10.3997/2214-4609.20141092. Link.
10. [△]Métivier L, Brossier R, Méridot Q, Oudet E, Virieux J (2016). "Measuring the misfit between seismograms using an optimal transport distance: application to full waveform inversion". *Geophysical Journal International*. 205 (1): 345–377. doi:10.1093/gji/ggw014. Link.
11. [△][‡]Sambridge M, Mosegaard K (2002). "Monte Carlo Methods in Geophysical Inverse Problems". *Reviews of Geophysics*. 40 (3): 3–1–3–29. doi:10.1029/2000RG000089. Link.
12. [△][‡]Mosegaard K, Tarantola A (2002). "16 – Probabilistic Approach to Inverse Problems". In *International Geophysics*, W. H. K. Lee, H. Kanamori, P. C. Jennings, and C. Kisslinger (Eds.), *International Handbook of Earthquake and Engineering Seismology, Part A*, Vol. 81, pp. 237–265. Link.
13. [△]Malinverno A (2002). "Parsimonious Bayesian Markov chain Monte Carlo inversion in a nonlinear geophysical problem". *Geophysical Journal International*. 151 (3): 675–688. doi:10.1046/j.1365-246X.2002.01847x. Link.
14. [△][‡]Bodin T, Sambridge M (2009). "Seismic tomography with the reversible jump algorithm." *Geophysical Journal International*. 178 (3): 1411–1436. doi:10.1111/j.1365-246X.2009.04226.x.
15. [△]Curtis A, Lomax A (2001). "Prior information, sampling distributions, and the curse of dimensionality." *GEOPHYSICS*. 66 (2): 372–378. doi:10.1190/1.1444928.
16. [△]Fichtner A, Zunino A, Gebraad L (2019). "Hamiltonian Monte Carlo solution of tomographic inverse problems." *Geophysical Journal International*. 216 (2): 1344–1363. doi:10.1093/gji/ggy496.
17. [△]Gebraad L, Boehm C, Fichtner A (2020). "Bayesian Elastic Full-Waveform Inversion Using Hamiltonian Monte Carlo." *Journal of Geophysical Research: Solid Earth*. 125 (3): e2019JB018428. doi:10.1029/2019JB018428.
18. [△]Martin J, Wilcox LC, Burstedde C, Ghattas O (2012). "A Stochastic Newton MCMC Method for Large-Scale Statistical Inverse Problems with Application to Seismic Inversion". *SIAM Journal on Scientific Computing*. 34 (3): A1460–A1487. doi:10.1137/110845598. Link.
19. [△]Sambridge M (2014). "A Parallel Tempering algorithm for probabilistic sampling and multimodal optimization". *Geophysical Journal International*. 196 (1): 357–374. doi:10.1093/gji/ggt342. Link.
20. [△][‡]Alvardi M (2021). "A gradient-based Markov chain Monte Carlo algorithm for elastic pre-stack inversion with data and model space reduction." *Geophysical Prospecting*. 69 (5): 926–948. doi:10.1111/1365-2478.13081.

21. ^aZhao Z, Sen MK (2021). "A gradient-based Markov chain Monte Carlo method for full-waveform inversion and uncertainty analysis". *GEOPHYSICS*. 86 (1): R15–R30. doi:10.1190/geo2019-0585.1. [Link](#).
22. ^a^bBerti S, Aleardi M, Stucchi E (2024b). "A computationally efficient Bayesian approach to full-waveform inversion." *Geophysical Prospecting*. 72 (2): 580–603. doi:10.1111/1365-2478.13437.
23. ^aBerti S, Aleardi M, Stucchi E (2024a). "A bayesian approach to elastic full-waveform inversion: application to two synthetic near surface models." *Bulletin of Geophysics and Oceanography*. 65 (2): 291–308. doi:10.4430/bgo00442.
24. ^aRay A, Sekar A, Hoversten GM, Albertin U (2016). "Frequency domain full waveform elastic inversion of marine seismic data from the Alba field using a Bayesian trans-dimensional algorithm". *Geophysical Journal International*. 205 (2): 915–937. doi:10.1093/gji/ggw061. [Link](#).
25. ^aSen MK, Biswas R (2017). "Transdimensional seismic inversion using the reversible jump Hamiltonian Monte Carlo algorithm". *GEOPHYSICS*. 82 (3): R119–R134. doi:10.1190/geo2016-0010.1. [Link](#).
26. ^aJordan MI, Ghahramani Z, Jaakkola TS, Saul LK (1998). "An Introduction to Variational Methods for Graphical Models." In *Learning in Graphical Models*, M. I. Jordan (Ed.), pp. 105–161 (en). doi:10.1007/978-94-011-5014-9.5.
27. ^a^bBlei DM, Kucukelbir A, McAuliffe JD (2017). "Variational Inference: A Review for Statisticians." *Journal of the American Statistical Association*. 112 (518): 859–877. doi:10.1080/01621459.2017.1285773.
28. ^aZhang C, Bütetpage J, Kjellström H, Mandt S (2019). "Advances in Variational Inference". *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 41 (8): 2008–2026. doi:10.1109/TPAMI.2018.2889774. [Link](#).
29. ^aKucukelbir A, Tran D, Ranganath R, Gelman A, Blei DM (2017). "Automatic Differentiation Variational Inference." *Journal of Machine Learning Research*. 18 (14): 1–45. External Links: ISSN 1533-7928, [Link Cited by: S1](#).
30. ^aKingma DP, Salimans T, Jozefowicz R, Chen X, Sutskever I, Welling M (2016). "Improved Variational Inference with Inverse Autoregressive Flow." In *Advances in Neural Information Processing Systems*, Vol. 29. External Links: [Link Cited by: S1](#).
31. ^aGallego V, Insua DR (2020). "Stochastic Gradient MCMC with Repulsive Forces." *arXiv. Note: Comment: Extends the workshop version* External Links: [Link](#), [Document Cited by: S1](#).
32. ^a^b^c^d^eLiu Q, Wang D (2016). "Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm". *Advances in Neural Information Processing Systems*. 29. [Link](#).
33. ^aZhang X, Curtis A (2020a). "Seismic Tomography Using Variational Inference Methods". *Journal of Geophysical Research: Solid Earth*. 125 (4): e2019JB018589. doi:10.1029/2019JB018589. [Link](#).

34. ^a Zhao X, Curtis A, Zhang X (2022). "Bayesian seismic tomography using normalizing flows". *Geophysical Journal International*. 228 (1): 213–239. doi:10.1093/gji/ggab298. Link.
35. ^a Siahkoobi A, Rizzuti G, Louboutin M, Witte PA, Herrmann FJ (2021). "Preconditioned training of normalizing flows for variational inference in inverse problems". *arXiv*. doi:10.48550/arXiv.2101.03709. Link.
36. ^a Ravasi M (2023). "Multi-realization seismic data processing with deep variational preconditioners". doi:10.1190/image2022-3745255.1. Link.
37. ^a Smith JD, Ross ZE, Azizzadenesheli K, Muir JB (2022). "HypoSVI: Hypocentre inversion with Stein variational inference and physics informed neural networks". *Geophysical Journal International*. 228 (1): 698–710. doi:10.1093/gji/ggab309. Link.
38. ^{a, b} Zhang X, Curtis A (2020b). "Variational full-waveform inversion". *Geophysical Journal International*. 222 (1): 406–411. doi:10.1093/gji/ggaa170. Link.
39. ^a Urozayev D, Ait-El-Fquih B, Hoteit I, Peter D (2022). "A reduced-order variational Bayesian approach for efficient subsurface imaging". *Geophysical Journal International*. 229 (2): 838–852. doi:10.1093/gji/ggab507. Link.
40. ^a Lomas A, Luo S, Irakarama M, Johnston R, Vyas M, Shen X (2023). "3D Probabilistic Full Waveform Inversion: Application to Gulf of Mexico Field Data". 84th EAGE Annual Conference & Exhibition. pp. 1–5. doi:10.3997/2214-4609.202310720. Link.
41. ^{a, b} Zhang X, Lomas A, Zhou M, Zheng Y, Curtis A (2023). "3-D Bayesian variational full waveform inversion". *Geophysical Journal International*. 234 (1): 546–561. doi:10.1093/gji/ggad057. Link.
42. ^a Izzatullah M, Alkhalifah T, Romero J, Corrales M, Luiken N, Ravasi M (2024b). "Posterior sampling with convolutional neural network-based plug-and-play regularization with applications to poststack seismic inversion." *GEOPHYSICS*. 89 (2): R137–R153. doi:10.1190/geo2023-0035.1.
43. ^{a, b} Corrales M, Izzatullah M, Ravasi M, Hoteit H (2022). "Bayesian RockAVO: Direct petrophysical inversion with hierarchical conditional GANs." (en). doi:10.1190/image2022-3745255.1.
44. ^{a, b} Izzatullah M, Alali A, Ravasi M, Alkhalifah T (2024a). "Physics-reliable frugal local uncertainty analysis for full waveform inversion." *Geophysical Prospecting*. 72 (n/a): 2718–2738. doi:10.1111/1365-2478.13528.
45. ^{a, b} Zhuo J, Liu C, Shi J, Zhu J, Chen N, Zhang B (2018). "Message Passing Stein Variational Gradient Descent". In *Proceedings of the 35th International Conference on Machine Learning*, pp. 6018–6027. Link.
46. ^a Kullback S, Leibler RA (1951). "On Information and Sufficiency." *The Annals of Mathematical Statistics*. 22 (1): 79–86. External Links: ISSN 0003-4851, Link Cited by: §2.1.

47. ^a^bGorham J, Mackey L (2017). "Measuring Sample Quality with Kernels." In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1292–1301 (en). Note: ISSN: 2640–3498 External Links: [Link Cited by: §2.2, §2.2.](#)
48. ^a^bJ. Ba, M. A. Erdogdu, M. Ghassemi, S. Sun, T. Suzuki, D. Wu, and T. Zhang (2021) *Understanding the Variance Collapse of SVGD in High Dimensions.* (en). External Links: [Link Cited by: §2.3.](#)
49. ^a^bD'Angelo F, Fortuin V (2021). "Annealed Stein Variational Gradient Descent." *arXiv*. External Links: [Link, Document Cited by: §2.3, §2.3.](#)
50. ^ΔHotelling H (1933). "Analysis of a complex of statistical variables into principal components." *Journal of Educational Psychology*. 24 (6): 417–441. doi:10.1037/h0071325.
51. ^ΔCampello RJGB, Moulavi D, Sander J (2013). "Density-Based Clustering Based on Hierarchical Density Estimates." In *Advances in Knowledge Discovery and Data Mining*, J. Pei, V. S. Tseng, L. Cao, H. Motoda, and G. Xu (Eds.), Berlin, Heidelberg, pp. 160–172 (en). doi:10.1007/978-3-642-37456-2_14.
52. ^ΔEster M, Kriegel H, Sander J, Xu X, et al. (1996). "A density-based algorithm for discovering clusters in large spatial databases with noise." In *kdd*, Vol. 96, pp. 226–231. Cited by: §2.5.
53. ^ΔBrougois A, Bourget M, Lailly P, Poulet M, Ricarte P, Versteeg R (1990). "Marmousi, model and data." pp. cp (en). Note: ISSN: 2214-4609 External Links: [ISBN 978-90-73781-01-6, Link, Document Cited by: §3.1.](#)
54. ^ΔRichardson A (2023). "Deepwave". doi:10.5281/zenodo.8381177. [Link](#).
55. ^ΔBunks C, Saleck F, Zaleski S, Chavent G (1995). "Multiscale seismic waveform inversion." *Geophysics*. 60 (5): 1457–1473.
56. ^ΔChen P, Ghattas O (2020). "Projected Stein Variational Gradient Descent." In *Advances in Neural Information Processing Systems*, Vol. 33, pp. 1947–1958. External Links: [Link Cited by: §4.](#)
57. ^ΔLiu X, Zhu H, Ton J, Wynne G, Duncan A (2022). "Grassmann Stein Variational Gradient Descent". *arXiv*. doi:10.48550/arXiv.2202.03297. [Link](#).
58. ^ΔBerti S, Aleardi M, Stucchi E (2024c). "A probabilistic full waveform inversion of surface waves." *Geophysical Prospecting*. 72 (9): 3448–3473. doi:10.1111/1365-2478.13595.
59. ^ΔSun Y, Williamson P (2024). "Invertible neural networks for uncertainty quantification in refraction tomography". *The Leading Edge*. 43 (6): 358–366.

Declarations

Funding: This research was supported by King Abdullah University of Science and Technology (KAUST) and the DeepWave consortium.

Potential competing interests: No potential competing interests to declare.