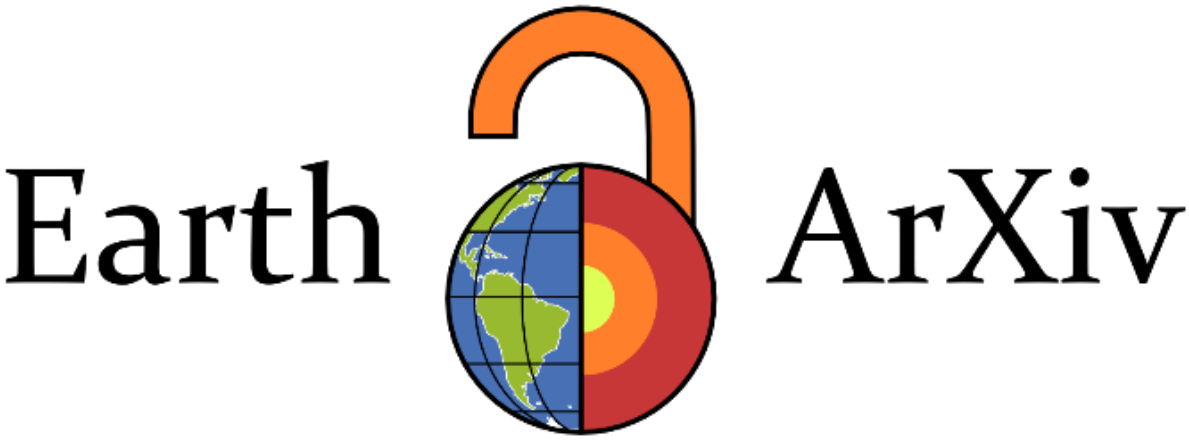

Comparison of probabilistic approaches to acoustic full-waveform inversion in compressed model and data spaces

Sean Berti[1], Fabio Macelloni[1], Mattia Aleardi[1], Eusebio Stucchi[1]

[1]Earth Sciences Department, University of Pisa, via S. Maria 53, 56126 Pisa, Italy

Abstract

Full-waveform inversion estimates subsurface properties by minimizing the misfit between observed and modelled data. However, conventional deterministic approaches are highly sensitive to noise, dependent on the starting model and prone to converging to local minima of the cost function. Bayesian approaches offer a viable alternative, enhancing solution space exploration and providing uncertainty quantification. This study compares three Bayesian inversion methods: Stochastic-Newton Markov Chain Monte Carlo, Ensemble Smoother with Multiple Data Assimilation and Annealed Stein Variational Gradient Descent. To mitigate ill-conditioning, reduce computational cost, and lower problem dimensionality, we incorporate Discrete Cosine Transform compression of model and data spaces. These methods are evaluated under three scenarios: optimal conditions as a baseline, a less informative prior to assess initialization sensitivity, and a more realistic setting with increased noise, missing traces and wrong assumptions about the source wavelet and noise level. These tests are conducted on a synthetic example based on a portion of the Marmousi model. Results indicate that the ensemble-based approach is the most computationally efficient, is robust to initialization, but strongly underestimates uncertainties. The Markov Chain method, while computationally demanding, is robust to initialization and effectively captures uncertainties. The Variational Inference approach, with intermediate computational cost, achieves good performance in complex scenarios, but it is more sensitive to initialization due to its deterministic nature. This study provides insights into the advantages and limitations of these three Bayesian approaches to acoustic full-waveform inversion and highlights the benefits of the Discrete Cosine Transform compression.

Keywords FWI · Bayesian approach · Discrete Cosine Transform · Acoustic

Article Highlights

- Three different Bayesian approaches to acoustic FWI are presented
- All methods are combined with a Discrete Cosine Transform compression of both model and data spaces
- Through a synthetic test, a comprehensive comparison between the performances of the three inversion approaches is performed

1. Introduction

Seismic full-waveform inversion (FWI) is a powerful geophysical technique used to estimate subsurface properties by minimizing the misfit between observed and synthetic seismic data generated through wave equation modelling (Lailly 1983; Tarantola 1984; Virieux and Operto 2009). Its ability to recover high-resolution models for subsurface compressional and shear wave velocities makes it an essential tool for seismic imaging. However, traditional deterministic approaches to FWI face significant challenges due to the inherent nonlinearity of the problem. These challenges include susceptibility to local minima, sensitivity to noise and a strong dependance on an accurate starting model. Moreover, these methods typically focus on finding a single best-fitting solution, often neglecting the quantification of uncertainties, which is critical for decision-making in exploration and monitoring applications.

To address these limitations, Bayesian FWI has emerged as a possible alternative, treating model parameters as random variables and inferring their posterior probability distributions (PPD) conditioned on the measured data. This framework enables a systematic exploration of the parameter space, providing comprehensive quantification of the solution uncertainties. However, Bayesian approaches introduce their own challenges, particularly the computational demands associated with exploring high-dimensional posterior distributions and the need for efficient sampling or approximation techniques to make the problem tractable.

For nonlinear inverse problems, a comprehensive characterization of the PPD typically requires sampling-based methods, such as Markov Chain Monte Carlo (MCMC). However, a significant challenge with MCMC methods is that convergence to the target posterior probability density heavily depends on the random perturbation applied to the current state of the chain. This perturbation, commonly referred to as the proposal distribution, must closely approximate the target posterior to ensure computational efficiency. Moreover, the efficiency of these methods significantly decreases in high-dimensional settings due to the so-called curse of dimensionality (Curtis and Lomax 2001). Traditional sampling strategies, such as the random walk Metropolis, may require billions of forward modelling evaluations to reach convergence, making them computationally prohibitive for problems involving expensive forward simulations, as in the case

of FWI (Aleardi and Ciabarri 2017). To address these limitations, various MCMC variants have been proposed to enhance sampling efficiency and mitigate the challenges associated with high-dimensional inverse problems. For example, the class of trans-dimensional MCMC methods, which enables exploration of model space by treating the number of parameters as an unknown itself, has garnered significant attention within the geophysical community. Among these, Reversible jump MCMC (RJ-MCMC), introduced by Green (1995), has been widely adopted in geophysics. Malinverno (2002) applied RJ-MCMC to electrical resistivity inversion, while Agostinetti and Malinverno (2010) extended its application to velocity model estimation using receiver-function data. Bodin and Sambridge (2009) further demonstrate the versatility of RJ-MCMC by employing it in seismic tomography using Voronoi cells. More recently, Ray et al. (2016) successfully applied RJ-MCMC to a frequency-domain 1D elastic FWI problem.

To enhance efficiency, gradient-based MCMC techniques have gained significant attention. A particularly promising variant is the Hamiltonian Monte Carlo (HMC; Neal 2011; Fichtner et al. 2019), which leverages Hamiltonian dynamics to incorporate gradient information, thereby accelerating convergence compared to gradient-free MCMC. This approach has been applied to various geophysical problems, including FWI (Sen and Biswas 2017; Fichtner and Simute 2018; Aleardi and Salusti 2020; Gebraad et al. 2020; Zunino et al. 2022).

Another advancement is represented by the Stochastic Newton MCMC, which constructs a proposal distribution based on both the local gradient and the Hessian of the negative log posterior. Originally introduced by Martin et al. (2012), this method has been later applied to acoustic FWI by Zhao and Sen (2021). Berti et al. (2023, 2024a, 2024b) further enhanced this approach by integrating it with a discrete cosine transform (DCT) reparameterization, reducing the dimensionality of both model and data spaces for acoustic and elastic FWI.

Although these methods provide accurate uncertainty quantification and despite these improvements to accelerate convergence and enhance efficiency, they remain computationally demanding, especially for large-dimensional problems involving costly forward modellings.

To alleviate this computational burden, an alternative Bayesian method, the ensemble-based (EB) inversion, can be applied. This class of methods provides a powerful framework for data assimilation (DA) by updating multiple system state realizations using the observed data (Carrassi et al. 2018; Evensen et al. 2022). The two main approaches within this framework are the Ensemble Kalman Filter (EnKF; Evensen 1994; Evensen 2009a) and the Ensemble Smoother (ES; van Leeuwen and Evensen 1996). As explained by van Leeuwen and Evensen (1996), a filter assimilates data only up to a given time, whereas a smoother incorporates all the available observations. An improvement of the ES consists of assimilating the available data multiple times, instead of performing a single global update. This results in a new algorithm called Ensemble Smoother with Multiple Data Assimilation (ES-MDA; Emerick and Reynolds 2013). Initially developed for oceanographic applications (Bertino et al. 2003), atmospheric modeling (Whitaker

et al. 2008) and weather prediction (Houtekamer and Mitchell 2005), these methods have gained attention also for geophysical applications. For example, EB methods have been employed in controlled-source electromagnetic (CSEM) inversion to identify large-scale structures (Tveit et al. 2015) and to monitor large-scale CO_2 injection by jointly inverting lower-resolution data (from CSEM or gravimetry) with higher-resolution amplitude-versus-offset (AVO) data (Tveit et al. 2020). Similarly, Aleardi et al. (2021) applied an EB approach to the inversion of electrical resistivity tomography (ERT) data, while Vinciguerra et al. (2023) extended this methodology for time-lapse ERT applications. In seismic inversion, Jin et al. (2008) and Gineste et al. (2019) used an EB approach to estimate elastic parameters in a 1D case. Additionally, several studies have explored EB techniques for history matching and reservoir characterization (Cominelli et al. 2009; Chen and Oliver 2013; Liu and Grana 2018a, 2018b, 2018c). More recently, Thurin et al. (2019) introduced an EB method for 2D visco-acoustic frequency-domain FWI, which was later extended to 3D time-domain FWI by Hoffmann et al. (2024).

While computationally more efficient than MCMC, EB methods exhibit limitations due to their reliance on an appropriate ensemble size and Gaussian assumptions on data and model parameters, which can introduce biases and hinder their ability to handle strong nonlinearities in the forward modelling operator and non-parametric prior assumptions.

An alternative class of approaches, variational inference (VI) methods (Blei et al. 2017), seeks to approximate the target posterior by minimizing the difference between a selected family of simple distributions (known as variational family) and the true posterior. Typically, this optimization process minimized the Kullback-Leibler (KL) divergence (Kullback and Leibler 1951). Advancements in computational power and deep learning frameworks such as PyTorch (Paszke et al. 2019) have facilitated the development and application of sophisticated variational algorithms. These methods either transform an initial set of samples to match the PPD deterministically (Liu and Wang 2016; Gallego and Insua 2020) or estimate the parameters of a chosen variational family that best approximates the posterior (Kingma et al. 2016; Kucukelbir et al. 2017). Mean field variational inference, for example, assumes complete independence among model parameters. Nawaz and Curtis (2018) applied this technique to infer subsurface geological facies and petrophysical properties from seismic data. However, its assumption of independent parameters limits its effectiveness for complex inverse problems. To overcome this limitation, Kucukelbir et al. (2017) proposed Automatic differential variational inference (ADVI), which employs a Gaussian variational family. ADVI has been successfully applied to seismic traveltime tomography (Zhang and Curtis 2020a) and earthquake slip inversion (Zhang and Chen 2022).

More recently, particle-based inference algorithms have gained attention for their ability to approximate complex posterior distributions without restrictive parameter assumptions. Among these, Stein Variational Gradient Descent (SVGD) has emerged as a promising alternative (Liu and Wang 2016). Unlike traditional MCMC methods that rely on randomness, SVGD iteratively

updates a set of particles—each representing a subsurface model—through deterministic transformations that minimize the KL divergence between the approximate and target distributions. This method has shown notable success in FWI applications, including transmission FWI (Zhang and Curtis 2020b), 2D FWI with realistic priors (Zhang and Curtis 2021) and 3D acoustic FWI using both synthetic (Zhang et al. 2023) and field data (Lomas et al. 2023). However, when applied to large-scale inverse problems and when using a relatively small number of particles, SVGD suffers of two main issues: the mode-collapse (Zhuo et al. 2018), with the particles collapsing into a small number of modes, and the variance collapse (Ba et al. 2022), with the estimated variance which can be significantly smaller than the variance of the target distribution. To tackle these issues, a variant of SVGD was introduced by D'Angelo and Fortuin (2021) under the name of Annealed SVGD and later applied to both acoustic and elastic FWI (Corrales et al. 2025, Berti et al. 2025).

This study compares three methods that belong to these distinct classes of Bayesian approaches applied to acoustic FWI: the Stochastic-Newton Markov Chain Monte Carlo (SN-MCMC), the Ensemble Smoother with Multiple Data Assimilation (ES-MDA) and the Annealed Stein Variational Gradient Descent (A-SVGD).

To mitigate the ill-conditioning and reduce the dimensionality of the FWI problem, all the considered methods incorporate a Discrete Cosine Transform (DCT) for model and data space compression. This reparameterization significantly reduces the number of unknowns while preserving essential model features, enhancing computational efficiency without compromising accuracy. For the SN-MCMC and ES-MDA inversions, the use of DCT is crucial to maintaining a feasible computational cost, as performing these algorithms in the full model space would be impractical. In the case of A-SVGD, the DCT allows us to satisfy the optimal condition for a SVGD method, where the number of particles should at least match the number of unknowns (Ba et al. 2022), all while ensuring computational efficiency.

In recent years, the use of DCT compression in geophysical applications has gained increasing attention. Lochbühler et al. (2014), Aleardi (2021) and Vinciguerra et al. (2023) applied it to ERT inversion, while Moghadas and Vrugt (2019) integrated Bayesian inversion with DCT-based image compression for subsurface electrical conductivity imaging. Zhou and Li (2013) used DCT for GPR compression and Aleardi (2020b) used it as a model compression technique in AVA inversion. Jafarpour and McLaughlin (2008) and Jung et al. (2017) employed the DCT in history matching applications. Jafarpour et al. (2009) proved the effectiveness of DCT compression in various problems, including facies reconstruction, crosswell traveltime tomography and porosity estimation. More recently, Kim et al. (2021) incorporated the DCT into acoustic FWI to compress the model space. Urozayev et al. (2021) integrated DCT in the forward modeling step within a variational Bayesian inversion framework to mitigate the challenges of high-dimensionality and ill-posedness. They applied it to two different problems: (i) reflectivity model reconstruction

using a linear forward model, as in Kirchhoff pre-stack depth migration, and (ii) 2D acoustic FWI, where DCT compression was employed in estimating the square slowness field. For additional insights into the use of DCT in the FWI problem we refer the reader to Berti et al. (2023, 2024b). Additionally, DCT plays a growing role in deep learning applications for geophysical problems (Aleardi 2020a; Rincón et al. 2025).

Several alternative reparameterization and compression approaches have also been explored in geophysics. Among them is the use of two grids in the inversion: a fine grid for the modeling and a coarse one for the model discretization (Mazzotti et al. 2016). Transdimensional approaches also offer a way to adapt the parameterization during inversion (Sambridge et al. 2013). Other frequently employed methods include Principal Component Analysis (PCA; Chen et al. 2014; Nuha et al. 2017; Corrales et al. 2024), Singular Value Decomposition (SVD; Liu and Grana 2018b), B-splines (Barnier et al. 2023) and wavelet transforms (Davis and Li 2011; Jafarpour 2010; Lin et al. 2012; Saad and Alkhalifah 2025). Furthermore, deep learning techniques, such as Convolutional Autoencoders, are becoming a valuable option to perform data and model compression in high dimensional inverse problems (Laloy et al. 2017; Liu and Grana 2018a; Canchumuni et al. 2019; Aleardi et al. 2021).

The performance of the three considered Bayesian methods is evaluated using a synthetic example based on the 2D Marmousi model, focusing on model prediction quality, data misfit, uncertainty quantification and computational workload. Three different tests have been performed: the first serves as a baseline (Test 1), the second to investigate the sensitivity of each approach to initialization (Test 2) and the third to assess their robustness when more realistic conditions are considered (Test 3). This work aims to provide insights into the performance of these Bayesian approaches to acoustic FWI and the benefits of DCT-based compression in addressing the challenges of high-dimensional inverse problems.

2. Methods

In this section, we first describe the DCT technique. Then, we present the three inversion methods employed in the study. Finally, we describe how the prior distribution is generated and used in the inversion framework.

2.1 Discrete Cosine Transform

The Discrete Cosine Transform (DCT) is a mathematical tool that represents a signal in terms of its energy distribution in the frequency domain. Like other Fourier-related transforms, DCT expresses a function as a sum of weighted cosinusoids of varying frequencies and amplitudes (Britanak et al. 2010). Among its different formulations, the most commonly used is the DCT-II

(Ahmed et al. 1974), which will be simply referred to as DCT in this work. The DCT is particularly advantageous in signal and image compression due to its ability to concentrate most of a signal energy into a small number of low-order coefficients (Jain 1989; Lochbühler et al. 2014). By setting high-order coefficients to zero, DCT enables efficient signal compression while preserving essential information.

DCT stands out as particularly well suited for our study due to several key properties. Unlike many alternative techniques, DCT is linear, computationally fast, easily extendible to two or three dimensions, and has a well-defined inverse transform. The implemented frameworks for SN-MCMC, ES-MDA and A-SVGD require several transformations from the original to the compressed domain, and vice versa. Indeed, all the inversions occur in the compressed space, but models must be projected back into the full domain for forward modelling calculations. Therefore, a reparameterization technique that ensures a stable, computationally efficient signal transformation, while providing strong compression capabilities, is crucial. The ability of DCT to significantly reduce the dimensionality of the problem, the overall computational workload and storage requirements makes it a valuable tool for high-dimensional inverse problems.

A key consideration in DCT-based compression is selecting the number of coefficients to retain, which balances the tradeoff between maintaining resolution and achieving high compression efficiency. To determine the optimal number of coefficients to compress both data and models, we adopt the approach proposed by Aleardi (2021), which evaluates how well a given DCT compression preserves the variability of the original signal. This variability is quantified as the ratio between the variances of the compressed and uncompressed signals, providing a measure of the retained information.

Let us consider a 2D velocity model $\mathbf{v}$, discretized on a grid with N_x nodes along the horizontal and N_y along the vertical direction, respectively. Its 2D DCT transformation can be expressed as:

$$\mathbf{V} = \mathbf{B}_y \mathbf{v} \mathbf{B}_x^{\mathrm{T}} , \tag{1}$$

where the matrices $\mathbf{B}_x$ and $\mathbf{B}_y$ contain the orthogonal DCT basis functions. Their dimensions are $N_x \times N_x$ and $N_y \times N_y$, respectively. The matrix $\mathbf{V}$ contains the DCT coefficients that represent the original model, with most of its variability captured by the low-order coefficients. To achieve compression, only the coefficients corresponding to the first p rows and the first q columns of $\mathbf{V}$ are retained: we refer to this selected portion of $\mathbf{V}$ as $\mathbf{V}_{pq}$. Therefore, the velocity model $\mathbf{v}$ can then be approximated as:

$$\mathbf{v}_{\boldsymbol{dct}} = (\mathbf{B}_y^p)^{\boldsymbol{T}} \mathbf{V}_{pq} \mathbf{B}_x^q , \tag{2}$$

where the model after compression $\mathbf{v}_{\boldsymbol{dct}}$ has the same dimensions as the original model ($N_y \times N_x$), but its actual parameterization is reduced to $p \times q$ unknowns. The matrices $\mathbf{B}_y^p$ and $\mathbf{B}_x^q$ contain the first p and q rows of $\mathbf{B}_y$ and $\mathbf{B}_x$, respectively.

In the framework of an inverse problem, this compression significantly reduces the dimensionality of the model space, as the velocity model is now represented in DCT domain by the numerical elements expressed by the matrix $\mathbf{V}_{pq}$, with these $p \times q$ parameters acting as unknowns. The velocity model in DCT domain can then be projected back to the full model space using the inverse transformation expressed in Equation 2. Figure 1 shows a schematic representation of the model compression. It highlights the concentration of the relevant information in low order DCT coefficients, i.e. in the top left corner of the matrix $\mathbf{V}$ represented in the center of the figure. On the right, we instead observe how the inverse transformation (IDCT) leads to a lower resolution version of the original model. This compression introduces an additional source of modelling error, as it limits the ability to reproduce high-wavenumber features, acting as a low-pass filter. This introduces a systematic component of the data misfit that we did not take into account in our inversion experiments. However, if needed, its effect can be considered and added to the data covariance matrix.

A similar approach is applied to the seismic data to efficiently reduce its dimensionality while preserving essential information. Given the large number of time samples and recording channels typically employed in large-scale seismic acquisitions, a higher number of DCT coefficients must be retained to prevent information loss and undesired numerical artifacts. Each shot gather is independently transformed into the DCT domain, compressed, and then reshaped into a single vector, ensuring an efficient yet accurate data representation. Figure 2 illustrates this procedure schematically.

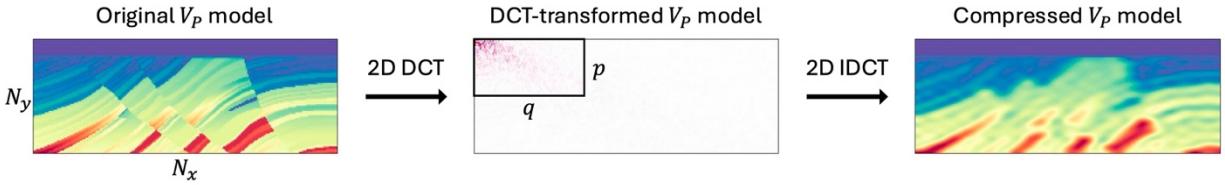

Fig. 1 *Example of application of the 2D DCT compression to a velocity model*

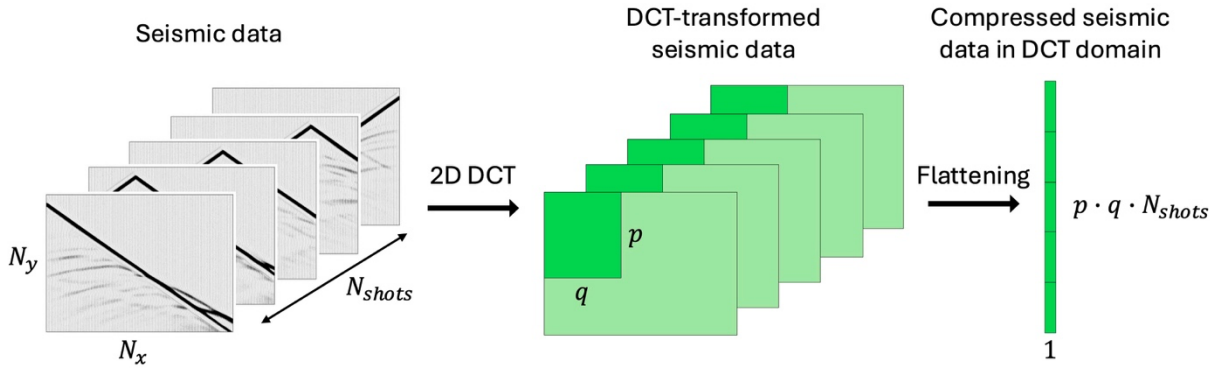

Fig. 2 *DCT compression of seismic data. The compression is applied to each shot gather, then the matrices containing the selected coefficients are flattened and stacked*

The primary advantage of this compression is the substantial improvement in computational efficiency during the inversion. By reducing the number of unknowns—representing the model in terms of DCT coefficients rather than physical parameters—the inversion operates in a lower-dimensional space, decreasing both computational time and memory requirements, and also mitigating the ill-posedness of the inverse problem. Furthermore, the reduction in data points simplifies the handling of key matrices involved in the inversion process. However, the choice of retained coefficients remains a tradeoff between data resolution, model resolution and computational cost, as a larger number of coefficients improves spatial resolution, data fitting, but increases computational complexity.

2.2 Stochastic-Newton Markov chain Monte Carlo

This version of MCMC sampling method builds upon deterministic gradient-based inversion techniques, which aim to minimize a misfit function typically formulated as a linear combination of data misfit and model regularization terms. Further details can be found in the works of Martin et al. (2012), who initially formulated this MCMC variant under the name of Stochastic Newton MCMC, and then successfully applied to acoustic FWI by Zhao and Sen (2021) and Berti et al. (2023).

In case of Gaussian assumptions for noise and model parameter distributions, the misfit function for a deterministic inversion can be formulated as (Tarantola 2005):

$$\boldsymbol{\Phi}(\mathbf{m}) = \frac{1}{2}\left\|\mathbf{C}_{\mathrm{d}}^{-\frac{1}{2}}\left(\mathbf{d} - G(\mathbf{m})\right)\right\|_2^2 + \frac{1}{2}\left\|\mathbf{C}_{\mathrm{m}}^{-\frac{1}{2}}\left(\mathbf{m} - \mathbf{m}_{\mathrm{prior}}\right)\right\|_2^2 \quad (3)$$

where $\mathbf{m}$ and $\mathbf{d}$ are, respectively, the model parameters and the observed data vectors; $\mathbf{C}_{\mathrm{d}}$ and $\mathbf{C}_{\mathrm{m}}$ are the data and prior covariance matrices, $\mathbf{m}_{\mathrm{prior}}$ represents the prior model and G is the forward modelling operator.

To minimize $\boldsymbol{\Phi}(\mathbf{m})$, a local quadratic approximation is made around the current model $\mathbf{m}_{\mathrm{k}}$:

$$\boldsymbol{\Phi}(\mathbf{m}) = \boldsymbol{\Phi}(\mathbf{m}_{\mathrm{k}} + \Delta\mathbf{m}) \approx \widetilde{\boldsymbol{\Phi}} = \boldsymbol{\Phi}(\mathbf{m}_{\mathrm{k}}) + \Delta\mathbf{m}^T\nabla_m\boldsymbol{\Phi}(\mathbf{m}_{\mathrm{k}}) + \frac{1}{2}\Delta\mathbf{m}^T\nabla_{\mathrm{m}}^2\boldsymbol{\Phi}(\mathbf{m}_{\mathrm{k}})\Delta\mathbf{m} + O(\|\Delta\mathbf{m}\|^3) \quad (4)$$

where $\Delta\mathbf{m} = \mathbf{m} - \mathbf{m}_{\mathrm{k}}$. The gradient $\mathbf{g}$ and Hessian $\mathbf{H}$ of $\boldsymbol{\Phi}(\mathbf{m})$ are given by:

$$\nabla_{\mathrm{m}}\boldsymbol{\Phi}(\mathbf{m}_{\mathrm{k}}) = \mathbf{g} = \mathbf{J}^T\mathbf{C}_{\mathrm{d}}^{-1}\Delta\mathbf{d}(\mathbf{m}_{\mathrm{k}}) + \mathbf{C}_{\mathrm{m}}^{-1}\left(\mathbf{m}_{\mathrm{k}} - \mathbf{m}_{\mathrm{prior}}\right) \quad (5)$$

$$\nabla_{\mathrm{m}}^2\,\boldsymbol{\Phi}(\mathbf{m}_{\mathrm{k}}) = \mathbf{H} \approx \mathbf{J}^{\mathrm{T}}\mathbf{C}_{\mathrm{d}}^{-1}\mathbf{J} + \mathbf{C}_{\mathrm{m}}^{-1} \quad (6)$$

where $\mathbf{J}$ is the Jacobian matrix containing the partial derivatives of data with respect to model parameters and $\Delta\mathbf{d}(\mathbf{m}_{\mathrm{k}}) = \mathbf{G}(\mathbf{m}_{\mathrm{k}}) - \mathbf{d}$ represents the data misfit.

In a Bayesian inversion, the final solution is the posterior probability density (PPD) function, which represents the probability of the model $\mathbf{m}$ given the observed data $\mathbf{d}$, by means of the Bayes' theorem:

$$p(\mathbf{m}|\mathbf{d}) = \frac{p(\mathbf{d}|\mathbf{m})p(\mathbf{m})}{p(\mathbf{d})} \tag{7}$$

In the previous equation, $p(\mathbf{m}|\mathbf{d})$ is the target PPD, $p(\mathbf{m})$ and $p(\mathbf{d})$ are the prior distributions for model parameters and data, respectively, and $p(\mathbf{d}|\mathbf{m})$ is the likelihood function.

The Bayesian inversion scheme can then be expressed in terms of $\mathbf{\Phi}(\mathbf{m})$, $\mathbf{H}$ and $\mathbf{g}$, assuming Gaussian distributions for data, noise and model parameters:

$$p(\mathbf{m}) \propto \exp\left(-\frac{1}{2}\left(\mathbf{m}-\mathbf{m}_{\text{prior}}\right)^T \mathbf{C}_{\text{m}}^{-1}\left(\mathbf{m}-\mathbf{m}_{\text{prior}}\right)\right) \tag{8}$$

$$p(\mathbf{d}|\mathbf{m}) \propto \exp\left(-\frac{1}{2}\left(\mathbf{d}-\mathbf{G}(\mathbf{m})\right)^T \mathbf{C}_d^{-1}\left(\mathbf{d}-\mathbf{G}(\mathbf{m})\right)\right) \tag{9}$$

Following the local approximation of the misfit function $\mathbf{\Phi}(\mathbf{m})$ described in Equation 4, we can retrieve a local Gaussian approximation of the PPD around the current model $\mathbf{m}_{\text{k}}$:

$$p(\mathbf{m}|\mathbf{d}) \propto \exp\left(-\tfrac{1}{2}\left(\mathbf{m}-(\mathbf{m}_{\text{k}}-\mathbf{H}^{-1}\mathbf{g})\right)^T \mathbf{H}\left(\mathbf{m}-(\mathbf{m}_{\text{k}}-\mathbf{H}^{-1}\mathbf{g})\right)\right) \tag{10}$$

The proposal distribution q used to generate a new MCMC sample is then defined as:

$$q(\mathbf{m}) \propto \exp\left(-\frac{1}{2}\left(\mathbf{m}-(\mathbf{m}_{\text{k}}-\alpha\mathbf{H}^{-1}\mathbf{g})\right)^T \frac{\mathbf{H}}{\beta^2}\left(\mathbf{m}-(\mathbf{m}_{\text{k}}-\alpha\mathbf{H}^{-1}\mathbf{g})\right)\right) \tag{11}$$

It is important to notice that we are approximating the target PPD as a Gaussian distribution, but the SN-MCMC method is theoretically able to sample any kind of posterior, considering that the PPD at the end of the sampling is independent from the proposal, that just influences the sampling efficiency.

In this formulation, α and β^2 are hyperparameters controlling the step size along the gradient direction and the magnitude of stochastic perturbations applied to the current model, respectively. The transition probability θ of moving from the current state of the Markov chain $\mathbf{m}_{\text{k}}$ to the next proposed state $\mathbf{m}_{\text{k+1}}$ is determined by the Metropolis-Hastings (M-H) acceptance rule:

$$\theta = p(\mathbf{m}_{\text{k+1}}|\,\mathbf{m}_{\text{k}}) = \min\left[1, \frac{p(\mathbf{m}_{\text{k+1}})}{p(\mathbf{m}_{\text{k}})} \times \frac{p(\mathbf{d}|\mathbf{m}_{\text{k+1}})}{p(\mathbf{d}|\mathbf{m}_{\text{k}})} \times \frac{q(\mathbf{m}_{\text{k}}|\mathbf{m}_{\text{k+1}})}{q(\mathbf{m}_{\text{k+1}}|\mathbf{m}_{\text{k}})}\right] \tag{12}$$

If the proposed sample $\mathbf{m}_{k+1}$ is accepted, the Markov chain updates to $\mathbf{m}_k = \mathbf{m}_{k+1}$; otherwise, $\mathbf{m}_k$ remains unchanged and a new proposal is generated in the next iteration.

To ensure an unbiased estimation of the posterior, the initial iterations of the Markov chain—known as the burn-in phase—are discarded. This allows the chain to reach the stationary state, mitigating the influence of the initial conditions. The retained samples are then used to compute key statistical properties, such as the mean and standard deviation of the PPD.

The Stochastic Newton MCMC method provides several advantages over other traditional MCMC approaches. It guides sampling towards high-probability regions of the parameter space, enhancing computational efficiency. The proposal distribution leverages local covariance structure of the target density, allowing for more informed model updates. The use of the approximated Hessian matrix captures parameter correlations, leading to improved convergence rates and more efficient exploration of the model space. This approach enhances sampling efficiency, accelerates convergence and ensures that the generated samples are both diverse and independent, all while maintaining high acceptance probabilities. The main issue related to this approach is associated with the computation of the Jacobian matrix $\mathbf{J}$, which is required for both the gradient vector $\mathbf{g}$ and the Hessian matrix $\mathbf{H}$ (Equations 5 and 6). A forward finite difference scheme can be employed to derive the Jacobian, which requires as many forward evaluations as there are model parameters, leading to extremely high computational costs. The DCT compression of the model and data spaces is here essential to ensure a more feasible inversion process, significantly reducing the number of forward computations needed for calculating the Jacobian and also decreasing the size—and consequently the memory storage requirements—of $\mathbf{J}$, $\mathbf{H}$ and $\mathbf{g}$.

2.3 Ensemble Smoother with Multiple Data Assimilation

In this study we adopt the Ensemble Smoother with Multiple Data Assimilation (ES-MDA) algorithm (Emerick and Reynolds 2013). The standard ES, originally introduced by van Leeuwen and Evensen (1996), performs a single global update by assimilating all available data at once. Reynolds et al. (2006) demonstrated that this process is mathematically equivalent to a Gauss-Newton iteration, where the sensitivity matrix is replaced by an ensemble-averaged approximation. To enhance convergence, particularly in highly nonlinear problems, Emerick and Reynolds (2012) proposed an extension in which the same data is assimilated multiple times. They showed that, under linear Gaussian assumptions, single and multiple data assimilation (MDA) yield equivalent results. The ES-MDA method thus applies small multiple corrections to each member of an ensemble of models, over multiple assimilation steps, improving stability and facilitating convergence. Each iteration effectively represents a Bayesian update of the current ensemble.

In ES-MDA, N model realizations are drawn from a Gaussian prior distribution and updated iteratively over a predefined number of iterations N_{it}. The final ensemble represents the posterior distribution, from which statistical properties (e.g. mean and standard deviation) can be derived. At each iteration l, the model update for the i-th ensemble member $\mathbf{m}_i$ (where $i = 1, \dots, N$) is given by:

$$\mathbf{m}_i^{l+1} = \mathbf{m}_i^l + \underbrace{\mathbf{C}_{\mathrm{md}}^l \left(\mathbf{C}_{\mathrm{dd}}^l + \alpha_l \mathbf{C}_{\mathrm{d}}\right)^{-1}}_{\mathbf{K}^l} \left(\tilde{\mathbf{d}}_i^l - \mathbf{d}_i^l\right) \tag{13}$$

where $\mathbf{d}_i^l$ is the synthetic data computed for the model $\mathbf{m}_i^l$, and $\tilde{\mathbf{d}}_i^l$ is a perturbed version of the observed data $\mathbf{d}_{\mathrm{obs}}$, defined as:

$$\tilde{\mathbf{d}}_i^l = \mathbf{d}_{\mathrm{obs}} + \sqrt{\alpha_l} \mathbf{C}_{\mathrm{d}}^{1/2} \cdot \boldsymbol{n} \tag{14}$$

Here, $\boldsymbol{n} = \mathcal{N}(0, \boldsymbol{I})$ represents Gaussian noise with zero mean and unit variance. In these equations, α_l is the inflation coefficient that controls the degree of data perturbation at each iteration, while $\mathbf{C}_{\mathrm{d}}$ is the data covariance matrix. The matrix $\mathbf{C}_{\mathrm{dd}}^l$ represents the covariance of the predicted data associated with the ensemble members at the previous iteration, while the cross-covariance $\mathbf{C}_{\mathrm{md}}^l$ describes the statistical relationship between model parameters and data. If N_m represents the number of model parameters (i.e., the unknown of the inversion) and N_d the number of data points, the dimensions of $\mathbf{C}_{\mathrm{dd}}^l$ are $N_d \times N_d$, whilst $\mathbf{C}_{\mathrm{md}}^l$ is a $N_m \times N_d$ matrix.
Given a model parameter j ($j = 1, \dots, N_m$) and a data point k ($k = 1, \dots, N_d$), the (j, k)-th entry of $\mathbf{C}_{\mathrm{md}}^l$ can be computed as follows:

$$\mathbf{C}_{\mathrm{md}}^l[j, k] = \mathrm{cov}\left(\mathbf{m}^{j,l-1}, \boldsymbol{d}^{k,l-1}\right) \tag{15}$$

Hence, each element of the cross-covariance matrix $\mathbf{C}_{\mathrm{md}}^l$ at iteration l, is the covariance between the j-th model parameter and the k-th data point. Instead, the entries of $\mathbf{C}_{\mathrm{dd}}^l$ are computed as the covariance between two different data points. The two matrices $\mathbf{C}_{\mathrm{md}}^l$ and $\mathbf{C}_{\mathrm{dd}}^l$ are computed from models, and corresponding data, of the ensemble at the previous iteration. Therefore, a proper number N of members in the ensemble is crucial to effectively estimate the covariances. In our framework, the dimensions of both covariance matrices $\mathbf{C}_{\mathrm{dd}}^l$ and $\mathbf{C}_{\mathrm{md}}^l$ are considerably reduced by DCT. Specifically, N_d corresponds to the number of retained DCT coefficients in the compression of a single shot gather multiplied by the number of sources, while N_m is the number of retained coefficients employed in the compression of the velocity model.
The matrix $\mathbf{K}^l$ in Equation 13, also known as the Kalman gain, determines the magnitude of the model update, balancing uncertainties between observations and predictions—assigning greater weight to the data when model uncertainty is high and vice versa.

Equation 13 represents the Bayesian step in this EB framework, under the aforementioned Gaussian assumptions, which restrict the applicability of the approach to weakly nonlinear problems. However, it is possible to apply this equation also in presence of nonlinearities, because the covariance matrices $\mathbf{C}_{\mathrm{md}}^{l}$ and $\mathbf{C}_{\mathrm{dd}}^{l}$ are inferred from the ensemble members and associated computed data at the previous iteration. This approach allows the covariance matrices to capture the nonlinear structure of the problem, with their update at each iteration enabling the method to adapt to the nonlinear relationship between model parameters and data.

Under linear Gaussian assumptions, to ensure equivalence between single and multiple data assimilation and to recover the correct posterior mean (Emerick and Reynolds 2013), the following condition must be satisfied:

$$\sum_{l=1}^{N_{it}} \frac{1}{\alpha_l} = 1 \tag{16}$$

This implies that the number of iterations must be predetermined. The simplest choice for the inflation coefficients, also adopted in this work, is to set $\alpha_l = N_{it}$ for $l = 1, \dots, N_{it}$, so that the previous condition is satisfied and α_l is kept fixed during the inversion.

A common issue affecting EB method is the excessive loss of variance within the ensemble, known as ensemble collapse: all the members tend to converge toward their mean, and consequently the uncertainties result extremely underestimated. This happens when the number of models is not sufficient to properly infer the relationships among different parameters and among model parameters and data points. The DCT compression, besides reducing the dimension of the Kalman gain and covariance matrices, helps mitigate variance collapse and decreases the ensemble size needed to ensure proper velocity model estimation.

2.4 Annealed Stein Variational Gradient Descent

Stein Variational Gradient Descent (SVGD) is a prominent algorithm within the class of Variational Inference (VI) methods, recognized for its deterministic approach to approximating complex posterior distributions. By iteratively minimizing the Kullback-Leibler (KL) divergence between an initial approximating distribution and the target posterior, SVGD ensures that the resulting set of particles aligns closely with the posterior probability distribution (PPD).

The algorithm begins with an initial set of N particles $\{\mathbf{m}_i\}_{i=1}^{N}$, which are updated through a smooth, invertible transform defined as $T(\mathbf{m}_i) = \mathbf{m}_i + \epsilon\phi(\mathbf{m}_i)$, where $\phi = [\phi_1, \dots, \phi_i]$ is a smooth vector function representing the perturbation direction, and ϵ is a step size controlling the perturbation magnitude. If T is invertible, the transformed distribution $q_T(\boldsymbol{m})$, corresponding to the initial distribution $q(\mathbf{m})$, evolves according to:

$$\nabla_\epsilon KL[q_T||p]|_{\epsilon=0} = -E_q\left[\mathrm{trace}\left(A_p\phi(\mathbf{m})\right)\right] \tag{17}$$

where p is the target distribution and A_p is the Stein operator, defined as $A_p\phi(\mathbf{m}) = \nabla_{\mathrm{m}} \log p(\mathbf{m})\, \phi(\mathbf{m})^T + \nabla_{\mathrm{m}}\phi(\mathbf{m})$. The optimal perturbation direction ϕ^* that minimizes the KL divergence is derived using kernel functions (Liu and Wang 2016).

In SVGD, the probability distribution q is represented by the set of N particles, allowing expectations to be approximated by sample means. The iterative update that minimizes the KL divergence is then given by:

$$\phi_l^*(\mathbf{m}) = \frac{1}{N}\sum_{i=1}^{N}\left[k\left(\mathbf{m}_i^l, \mathbf{m}\right)\nabla_{\mathbf{m}_i^l} \log p\left(\mathbf{m}_i^l \middle| \mathbf{d}_{\mathrm{obs}}\right) + \nabla_{\mathbf{m}_i^l} k\left(\mathbf{m}_i^l, \mathbf{m}\right)\right] \tag{18}$$

$$\mathbf{m}_i^{l+1} = \mathbf{m}_i^l + \epsilon^l \phi_l^*\left(\mathbf{m}_i^l\right) \tag{19}$$

where l is the iteration index and ϵ^l is the step size. For sufficiently small ϵ^l, the algorithm asymptotically converges to the target posterior as $N \to \infty$.

For the kernel function, SVGD typically employs the radial basis function (RBF), defined as:

$$k(\mathbf{m}', \mathbf{m}) = \exp\left[-\frac{\|\mathbf{m} - \mathbf{m}'\|^2}{2h^2}\right] \tag{20}$$

where h is the bandwidth parameter controlling the interaction strength between particles. In accordance with established practices (e.g., Liu and Wang 2016, Zhang and Curtis 2020a, Izzatullah et al. 2024), h is set to $\tilde{d}/\sqrt{2\log N}$, where $\tilde{d}$ is the median of pairwise distances between particles. The SVGD update consists of two key components:

- *Gradient-driven attraction*: the term involving $\nabla \log p\,(\mathbf{m})$ guides particles toward regions of high probability in the target density $p(\mathbf{m})$.
- *Repulsive force*: the term involving ∇k promotes particle dispersion, preventing mode collapse by ensuring that particles spread across multiple modes of the posterior distribution.

Despite its effectiveness, standard SVGD faces challenges in high-dimensional settings, where it can suffer from variance underestimation—a phenomenon referred to as variance collapse. Ba et al. (2022) attributed this issue to bias introduced by the deterministic update in the driving force term of SVGD. To address this limitation, they proposed a damped SVGD variant that modifies the driving force term to improve variance estimation.

A closely related approach, Annealed Stein Variational Gradient Descent (A-SVGD), was introduced by D'Angelo and Fortuin (2021). This method incorporates a heuristic temperature parameter $\tau(l) \in [0,1]$ to modulate the intensity of the driving force. Inspired by strategies used in MCMC methods to address mode-collapse (Neal 1996), this modification improves exploration while maintaining SVGD's deterministic framework (Berti et al. 2025). By integrating an annealing schedule, A-SVGD enhances the ability to explore the target distribution, mitigates initialization-

related challenges and ensures more robust performance when the number of particles is limited.

The update rule A-SVGD is expressed as:

$$\phi_l^*(\mathbf{m}) = \frac{1}{N}\sum_{i=1}^{N}\left[\tau(l) * k\left(\mathbf{m}_i^l, \mathbf{m}\right)\nabla_{\mathbf{m}_i^l}\log p\left(\mathbf{m}_i^l\middle|\mathbf{d}_{\text{obs}}\right) + \nabla_{\mathbf{m}_i^l}k\left(\mathbf{m}_i^l, \mathbf{m}\right)\right]. \qquad (21)$$

The parameter $\tau(l)$, ranging between 0 and 1, controls the balance between two distinct phases:

- *Exploration phase* $(\tau \approx 0)$: the repulsive force dominates, dispersing particles widely to ensure thorough exploration of the target distribution.
- *Exploitation phase* $(\tau \approx 1)$: the driving force focuses on concentrating particles around the modes of the posterior distribution.

Selecting an appropriate annealing schedule is crucial to maintaining SVGD's convergence properties. As the iterations progress, the method should smoothly transition from exploration to exploitation, ensuring effective adaptation to the target density.

Following D'Angelo and Fortuin (2021) and Corrales et al. (2024), this study employs an annealing schedule based on the hyperbolic tangent function, given by:

$$\tau(l) = \tanh\left[\left(\frac{1.3l}{T}\right)^r\right] \qquad (22)$$

where l is the current iteration index, T is the total number of iterations and r is a user-defined parameter that controls the transition rate between exploration and exploitation phases (typically between 1 and 5, in this work we choose 3 as in Corrales et al. 2024). As $l \to \infty$, the temperature parameter approaches unity, ensuring that the final particle distribution accurately represents the posterior density. In this context, DCT compression of the model space enables the use of a number of particles equal to the number of unknowns—matching the ideal condition for SVGD (Ba et al. 2022)—while keeping reasonable computational costs, which would otherwise be unfeasible in the full domain.

2.5 Model and data prior information

In the following inversion tests, we assume that the velocities follow a Gaussian prior distribution. Due to the linearity of the DCT transformation, both the prior mean vector and the prior covariance matrix can be analytically mapped onto the reduced DCT space (Aleardi 2020b).

Let $\mathbf{C}_{\text{m}}$ denote the prior model covariance in the original velocity domain (as in Equation 8). It incorporates a 2D stationary Gaussian variogram model, which defines the assumed spatial correlation of the velocity field. For example, the spatial correlation function for the velocity along the horizontal x-direction is expressed as:

$$\tau_x = \exp\left(-\frac{h_x^2}{a_x^2}\right), \tag{23}$$

where h_x represents the spatial lag of the autocorrelation function along the x-axis and a_x denotes the effective correlation range. The full covariance matrix $\mathbf{C}_\mathrm{m}$ in the original model space can be formulated as a double Kronecker product (denoted by the symbol ⊗) involving the prior variance and the two spatial correlation functions τ_x and τ_y:

$$\mathbf{C}_\mathrm{m} = \Big(\big(\mathrm{var}(\mathbf{v})\otimes\mathbf{T}(\tau_x)\big)\otimes\mathbf{T}\big(\tau_y\big)\Big), \tag{24}$$

where $\mathrm{var}(\mathbf{v})$ is the assumed prior variance of the velocity model, in our case the variance of the prior mean model $\mathbf{m}_\mathrm{prior}$, and $\mathbf{T}$ represents a Toeplitz matrix.

Using the notation introduced in Section 2.1, the corresponding covariance in the DCT transformed domain, denoted as $\mathbf{C}_\mathrm{r}$, is given by:

$$\mathbf{C}_\mathrm{r} = (\mathbf{B}_y^p)^T\mathbf{C}_\mathrm{m}\mathbf{B}_x^q, \tag{25}$$

where $\mathbf{B}$ is the transformation matrix that projects the model parameters from the full space ($N_y \times N_x$) onto the reduced DCT space ($p \times q$). The matrix $\mathbf{B}$ (see Equation 2) is constructed using the Kronecker product of two cosine basis functions matrices, each composed of the first p and q rows of the full DCT basis matrices in the vertical and horizontal directions, respectively.

In our inversion tests, the prior mean model $\mathbf{m}_\mathrm{prior}$ is defined as a smoothed version of the true model (Test 1 and Test 3) or as a homogeneous model (Test 2). This prior information is consistently applied across all three inversion methods, to ensure coherence and facilitate a fair comparison of their performances. Similarly, the data likelihood function is formulated consistently for all methods, based on the misfit between observed and predicted data, weighted by the inverse of the data covariance matrix (as expressed in Equation 9), thus assuming uncorrelated Gaussian noise affecting the data, with constant variance. Under these conditions, after DCT transformation, we still get Gaussian uncorrelated noise with constant variance in the compressed data space. However, given the high dimensionality of the full data space, directly computing the data covariance matrix in the reduced domain is computationally prohibitive. To address this challenge, we employ a Monte Carlo numerical simulation to project the data covariance matrix $\mathbf{C}_\mathrm{d}$ onto the compressed domain, following the approach of Aleardi et al. (2022). Assuming a diagonal covariance matrix in full domain, since the DCT rotates the noise without introducing off-diagonal correlations, this construction is mathematically consistent with the noise added to the data and the $\mathbf{C}_\mathrm{d}$ matrix remains diagonal also in the DCT space.

3. Results

To compare the performance of the three Bayesian FWI approaches, we employed a small portion of the 2D Marmousi model (Brougois et al. 1990), a standard benchmark for FWI applications (Figure 3). The synthetic seismic data were modelled using Deepwave (Richardson

2023) on a rectangular grid with 81 points in the vertical direction and 216 points along the horizontal direction, with a uniform grid spacing of 20 m. The acquisition setup consisted of five shots, evenly distributed along the free surface, recorded by 200 receivers placed at regular intervals. The seismic source signature for the observed data was modelled as a Ricker wavelet with a peak frequency of 5 Hz. The recording time was set to 3 s, with a sampling interval of 4 ms. In all inversion experiments presented in this study, the water layer was assumed to be perfectly known, with a constant depth of 260 m. As discussed in the previous section, all three FWI algorithms operate in the DCT compressed domain.

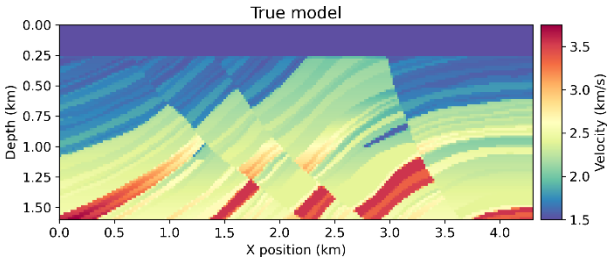

Fig. 3 *The true model employed in our inversion tests*

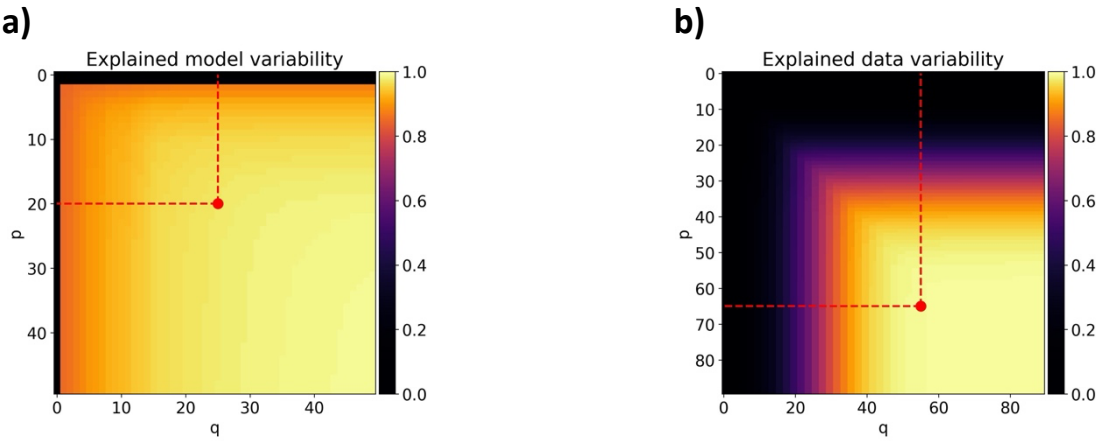

Fig. 4 *Explained variability of the true velocity model (a) and of the observed data (b). The two plots are close-ups of the corresponding matrices. The red dots show the chosen combinations of p and q for model and data compression*

Following the methodology of Aleardi et al. (2021) and Berti et al. (2023), we selected the combination of DCT coefficients based on the explained variability (as defined in Section 2.1) of the compressed signal (i.e., the velocity model or the seismic data) with respect to the original one in full domain. By computing the preserved variability for each possible combination of p and q, we generate a variability map, that helps determine the number of coefficients to retain in both horizontal and vertical directions of the signal while ensuring a given percentage of the

original variance is maintained. Figure 4 shows close-ups views of this explained variability plot for both the model and the data. The explained variability for the model is computed with respect to the true velocity field in our synthetic experiments, while for the data it corresponds to that of a single shot gather of the observed data. For the model, in case of an inversion test applied to a field dataset, the number of retrieved DCT coefficients can be estimated based on existing well log information or stack sections, for the vertical and lateral variability, respectively (Berti et al. 2024a).

As expected, compressing the data space requires a higher number of DCT coefficients compared to model space, because of the large number of time samples in the seismogram and the need to preserve the lateral continuity of events. The red dots in Figure 4 highlight the selected combinations of p and q for the number of retained basis functions along the vertical and horizontal directions, respectively. Regarding the model parameterization, we noticed that the combination $(p, q) = (20{,}25)$ allows to preserve about 97% of the variability in the original model. The latter, initially represented by

$$N_y \times N_x = 68 \times 216 = 14688$$

parameters, was reduced to

$$p \times q = 20 \times 25 = 500$$

DCT coefficients. The velocity model after the DCT compression is displayed in Figure 5 and this represents the target model for all the inversion tests.

Similarly, for the seismograms, with the choice $(p, q) = (65{,}55)$ about 99% of the variability is preserved after compression. Hence, the seismic data, originally containing

$$N_y \times N_x \times N_{shots} = 751 \times 200 \times 5 = 751000$$

samples, were compressed to

$$p \times q \times N_{shots} = 65 \times 55 \times 5 = 17875$$

DCT coefficients.

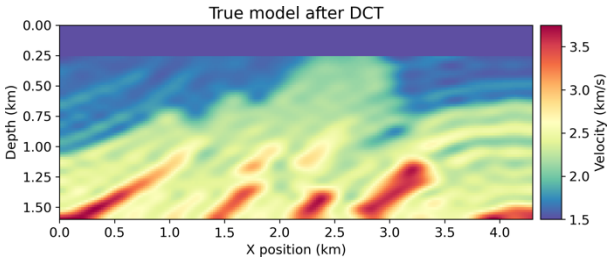

Fig. 5 *The true model after the DCT compression*

The three Bayesian FWI methods were configured as follows:

- *SN-MCMC*: we employed five independent chains, each running for 2500 iterations, with the first 1000 iterations designated as the burn-in period. This burn-in phase represents

the initial stage before convergence, and the corresponding samples were excluded from the statistical analysis of the posterior distribution.

- *ES-MDA*: the algorithm employed an ensemble of 10000 velocity models and 8 iterations, with the mean and standard deviation computed using the ensemble at the final iteration.
- *A-SVGD*: a set of 500 particles (chosen to be equal to the number of unknowns) was used, with the algorithm running for 350 iterations. At the final iteration, the mean and standard deviation were computed from the entire set of 500 particles.

The setting of some of these parameters comes from previous experiments and the need to ensure a fair comparison among the three methods. Therefore, the chosen settings were designed to facilitate a meaningful comparison of the results. For instance, the number of iterations in the SN-MCMC inversion was determined based on the length of the burn-in period and on the Potential Scale Reduction Factor value (PSRF, Gelman and Rubin 1992). The PSRF is a widely used statistic that assesses whether a Markov chain has converged to its equilibrium distribution by comparing its behaviour with other chains. This also justifies the use of multiple Markov chains, which enhances the exploration of the model space and helps preventing the sampling process from getting trapped in some local maxima of the PPD.

In the case of ES-MDA, we observed that a considerable number of ensemble members is required to properly estimate the velocity model for the considered problem. When using a smaller ensemble, artifacts tend to deteriorate the estimated posterior mean. The number of iterations was determined through trial and error: additional assimilation steps did not yield improvements significant enough to justify the resulting increase in computational cost. For A-SVGD, the number of particles was set equal to the number of unknowns (i.e., the number of retained DCT coefficients in model compression) to align with the optimal conditions of the algorithm. The number of iterations in this case was set based on the misfit evolution, stopping after reaching the convergence plateau where no significant changes were observed.

To evaluate and compare the three Bayesian FWI approaches, we designed three distinct tests:

- ***Test 1***, that serves as the baseline for the comparison of the three FWI approaches: we assume a perfectly known source wavelet, the right level of noise in the data and employ a smoothed version of the true model as the prior mean.
- ***Test 2***, to compare the effects of initialization of the inversions on the final solutions: we assume the same conditions used in the first test, but employing in this case a homogeneous velocity model as the prior mean.
- ***Test 3***, to test the robustness of each method simulating more realistic conditions: we employ the same prior mean used in the first test, an increased level of noise in the observed data, some missing traces and a phase-shift in the source wavelet when

simulating the predicted data. In addition, we made an erroneous assumption about the noise level contaminating the observation (see details in the following).

3.1 Test 1: baseline

This first test evaluates the performance of the three Bayesian FWI approaches under ideal conditions, where the source wavelet is perfectly known, and the observed data have a relatively high signal-to-noise ratio (SNR) of approximately 10 dB. The uncorrelated Gaussian random noise contaminating the data is perfectly modelled by the data covariance matrix.

For the ES-MDA and A-SVGD methods, the initial ensemble/particles were generated by drawing prior realizations from a Gaussian distribution, where a highly smoothed version of the true model served as the prior mean (Figure 6). In the SN-MCMC implementation, five randomly selected velocity models from the previously generated ensemble were used as the starting points for the five independent Markov chains.

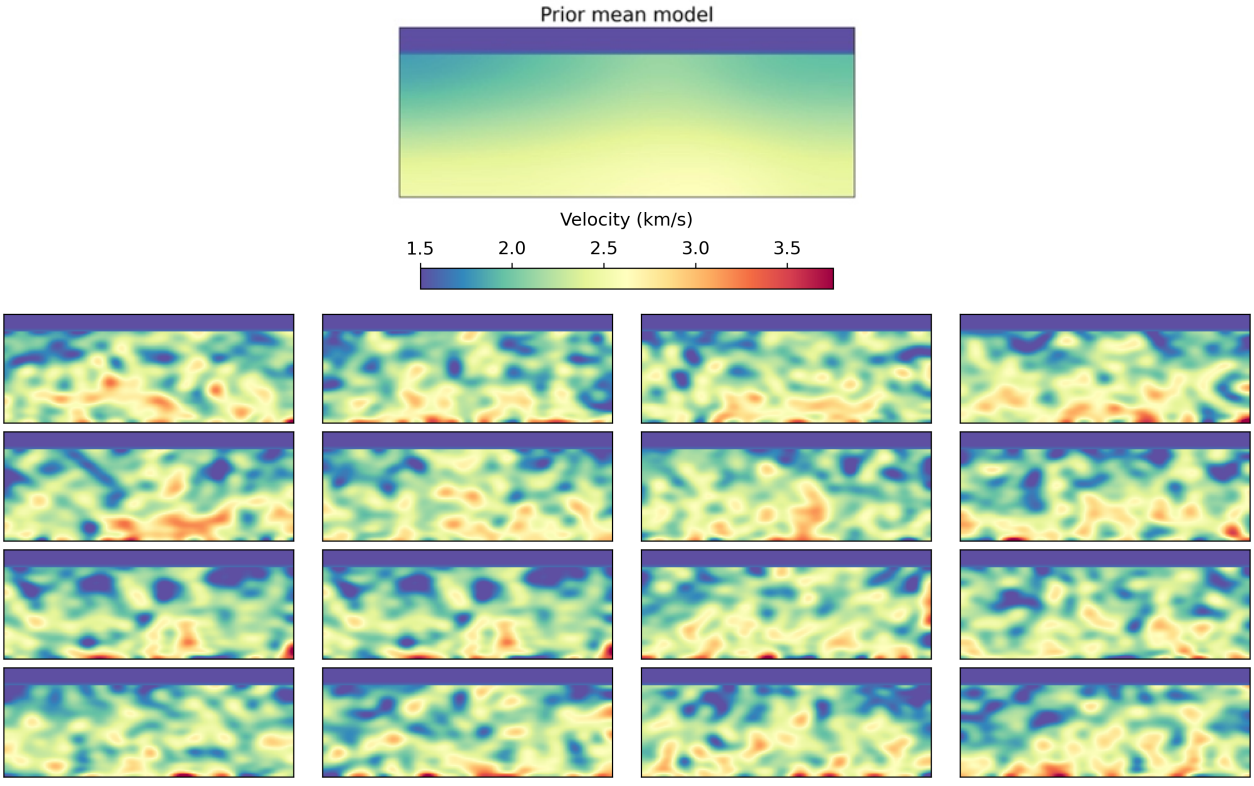

Fig. 6 *The smoothed version of the true model used as prior mean model in Test 1 and Test 3, and some realizations drawn from the prior distribution*

Figure 7 shows the evolution of the data misfit, normalized by the maximum initial value for all the three methods. The comparison highlights some differences among the considered approaches. We notice the rapid convergence of ES-MDA (Figure 7a), which requires only a few assimilation steps: this is due to the MDA scheme, which ensures better performance and data match with respect to computing a single global update, as in standard ES (Emerick and Reynolds 2013). For SN-MCMC (Figure 7b), the data misfit evolution is shown for all five chains, demonstrating that they all have converged to similar values after the burn-in period. The A-SVGD (Figure 7c) shows a period of stalled convergence between 50 and 100 iterations, compatible with the initial phase of the implemented annealed scheme, where the repulsive force is dominant (Equation 21).

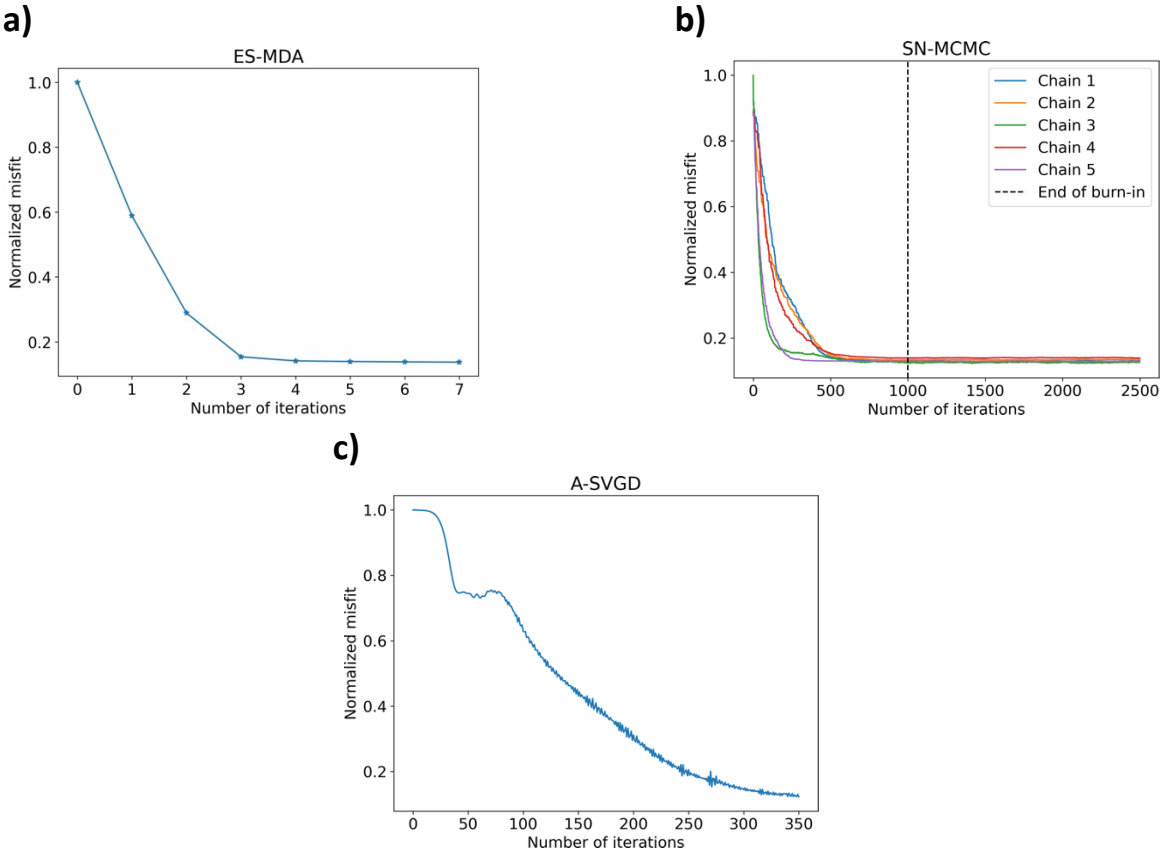

Fig. 7 *Evolution of the normalized data misfit for: (a) ES-MDA, (b) SN-MCMC and (c) A-SVGD.*

To assess the convergence and sampling quality of the Markov chains, we employed three standard MCMC diagnostics: the acceptance ratio, the autocorrelation length and the potential scale reduction factor (PSRF). Figure 8a shows the acceptance rates for each chain, computed as the ratio between accepted samples and number of iterations. All chains reach stable acceptance values of approximately 40% after about 1000 iterations, indicating a well-tuned proposal distribution and efficient exploration of the posterior space.

Figure 8b reports the mean autocorrelation lengths in the DCT parameter space for the five chains. The autocorrelation length measures the number of iterations required for sample to become effectively independent; shorter values indicate more efficient sampling and better mixing. After roughly 200 iterations, the autocorrelation lengths decrease to values close to zero, suggesting that the chains rapidly decorrelate and that successive samples provide large independent information about the posterior distribution.

Finally, Figure 8c presents the potential scale reduction factor (PSRF) values, which quantify the ratio between the averaged variance within a single chain and the variance across all chains. Each blue line corresponds to a model parameter in the DCT domain, while the red dashed line marks the commonly used threshold of 1.2, below which the sampling is considered to have reached a stable estimation of the PPD. As iterations progress, most PSRF values fall below this threshold, indicating that we have reached convergence for the corresponding model parameters, and the others are just above to the theoretical threshold.

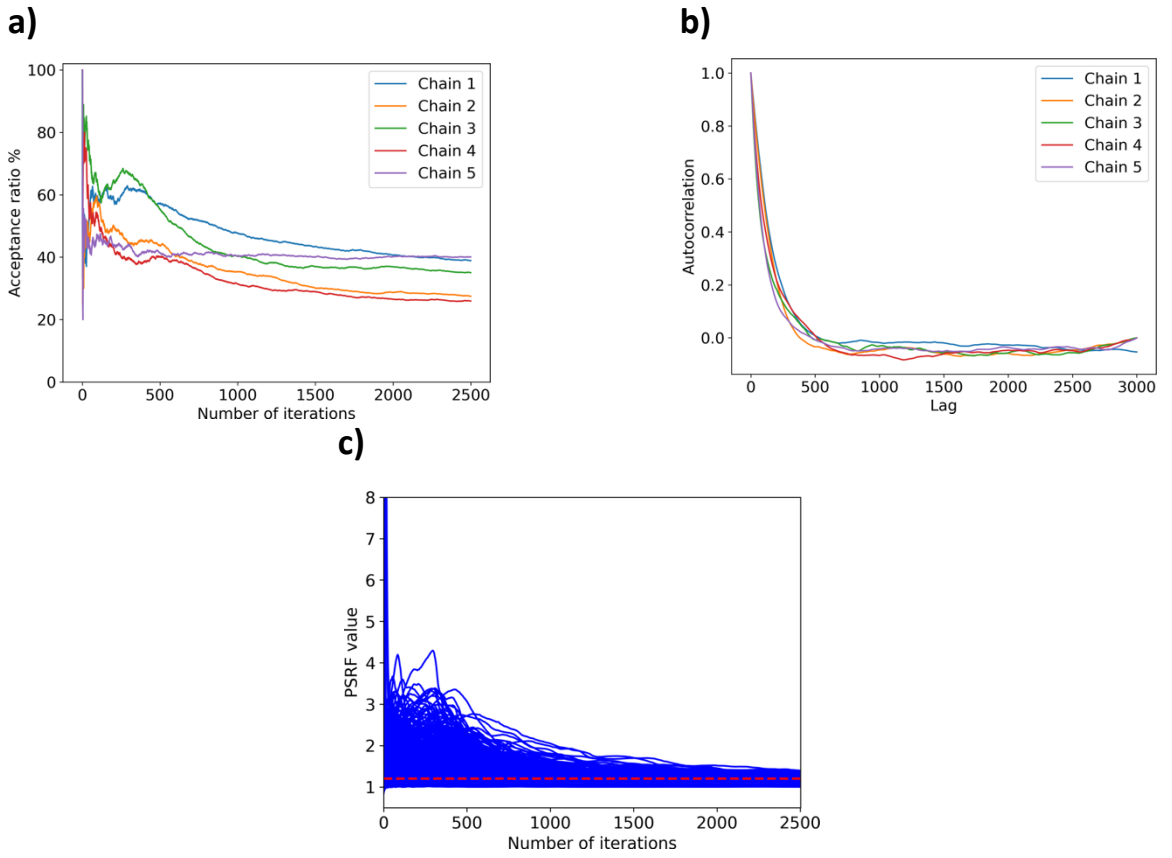

Fig. 8 *(a) Acceptance ratio (calculated as number of accepted models over number of iteration) for each chain employed for SN-MCMC inversion; (b) Mean of autocorrelation for all model parameters in DCT domain for all the five considered chains; (c) Evolution of the potential scale reduction factor (PSRF) values where each blue line refers to one of the retained DCT coefficients in model space, while the red dashed line shows the threshold for convergence (PSRF value of 1.2)*

Figure 9 presents the results of Test 1, displaying both the posterior mean velocity models (Figures 9a-c-e) and the standard deviation maps (Figures 9b-d-f). The mean models resulting from the SN-MCMC and ES-MDA methods are nearly identical, accurately capturing all the primary features of the true model. This observation is further supported by their high model SNR values, computed using the true model prior to DCT compression as the reference. In contrast, the A-SVGD solution exhibits a slight underestimation of the velocities in the deeper regions of the domain, leading to a lower SNR.

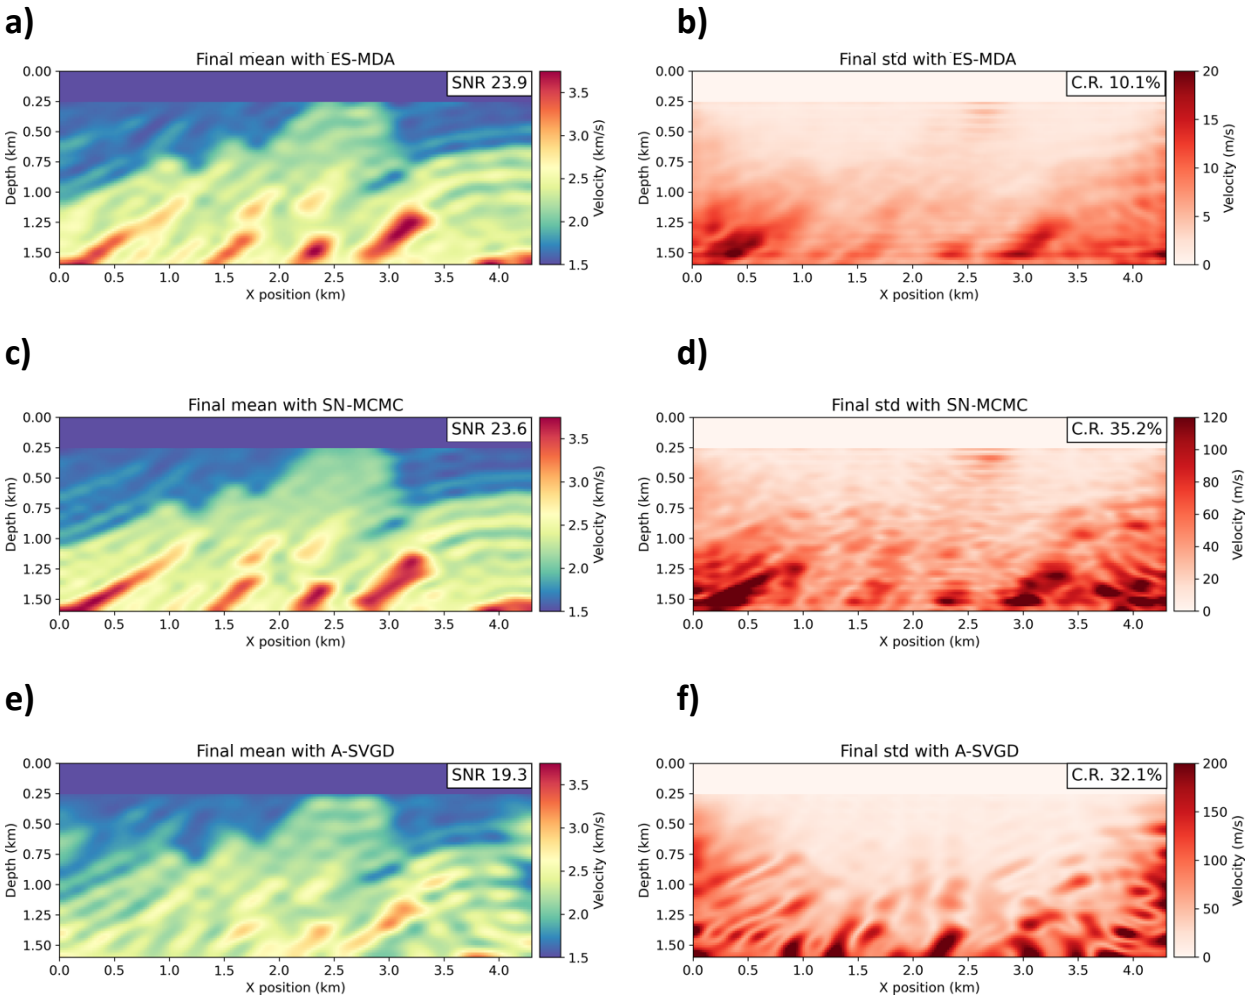

Fig. 9 *Results for Test 1: on the left column, the final mean models obtained with all three inversion approaches, with specified their SNR using the true model as reference; on the right column, the corresponding standard deviation maps with their 99 % coverage ratios (C.R.)*

The standard deviation maps from all three methods reveal similar patterns, showing increased uncertainty with depth and in correspondence of the highest velocity layers. However, the magnitude of the uncertainty estimates varies significantly across the different methods. In particular, the ES-MDA result exhibits a severe variance underestimation, an effect of the ensemble collapse issue affecting this approach. In contrast, while other particle-based methods like SVGD are also susceptible to variance collapse, this effect appears to be effectively mitigated when considering the A-SVGD result (Berti et al. 2025). This is confirmed by the standard deviation values, which closely align with those obtained from the SN-MCMC inversion, indicating a more reliable uncertainty quantification.

To better appreciate the differences between uncertainty estimations, we computed the 99% coverage ratios (C.R.), calculated as the percentage of parameters of the true velocity model (in the original domain) falling within the 99% confidence interval. This metric highlights the uncertainty underestimation in ES-MDA, emphasizing the impact of ensemble collapse on the final posterior distribution, while the coverage ratios for SN-MCMC and A-SVGD are comparable.

Figure 10 displays two 1D velocity profiles extracted from two distinct spatial locations (marked by the vertical magenta dashed lines in Figure 10a). These profiles compare the velocities derived from the true model before and after DCT compression with those obtained from the final mean models of the three inversion approaches. The shaded regions represent the 95% and 99% confidence intervals for each method, derived from their respective standard deviation values.

For the first profile (Figures 10b-d), corresponding to a horizontal position of 1.18km, all methods exhibit excellent agreement with the DCT-compressed version of the true model, which represents the optimal solution of the inversion problem. The uncertainty estimates follow a natural trend, increasing with depth due to reduced illumination, particularly for the SN-MCMC and A-SVGD methods. In contrast, the ES-MDA approach suffers from the ensemble collapse issue, resulting in overly narrow confidence intervals. In the second profile (horizontal position of 3.14km; Figures 10e-g), the velocity estimates from the ES-MDA and SN-MCMC inversions align almost perfectly with the true model after DCT compression, while A-SVGD shows some velocity underestimation at greater depths (as evidenced by Figure 9e). Once again, the confidence intervals associated to the SN-MCMC and A-SVGD inversions become much wider with depth. In the case of SN-MCMC, the true model generally falls within these uncertainty bounds, except at locations where abrupt velocity changes occur.

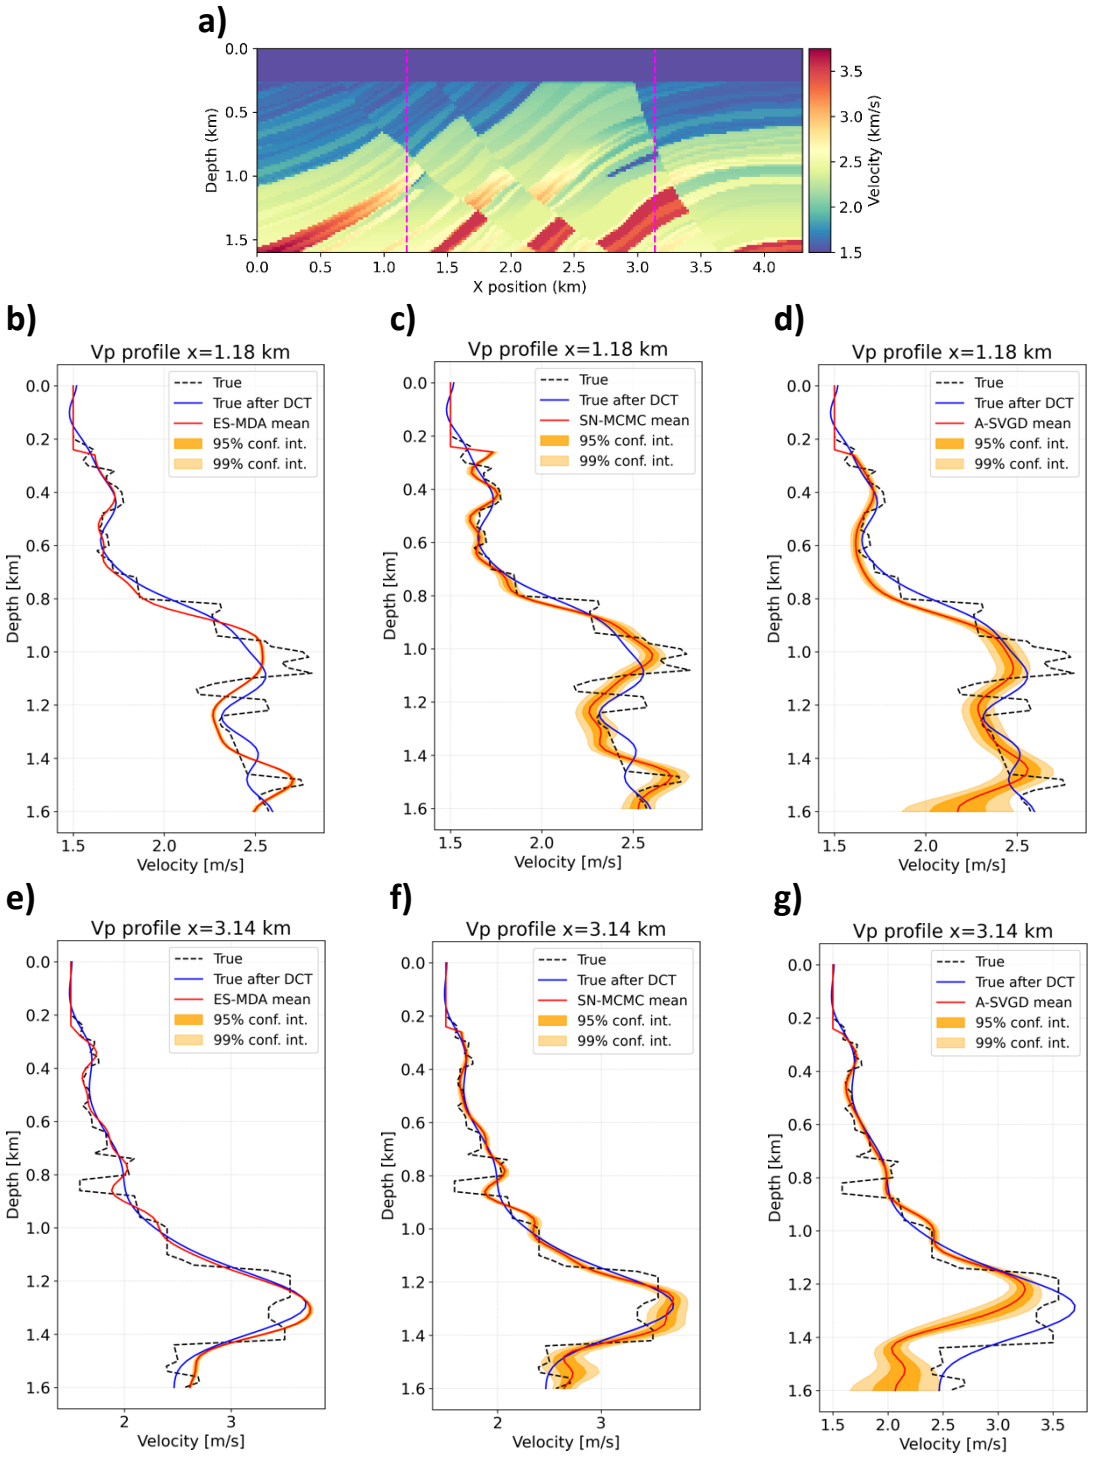

Fig. 10 *Pseudo-well velocity profiles associated to two different horizontal locations (dashed magenta lines in the true model), showing mean and two confidence intervals for all three inversion methods, together with the velocity values extracted from the true model before (dashed black) and after (blue) DCT compression*

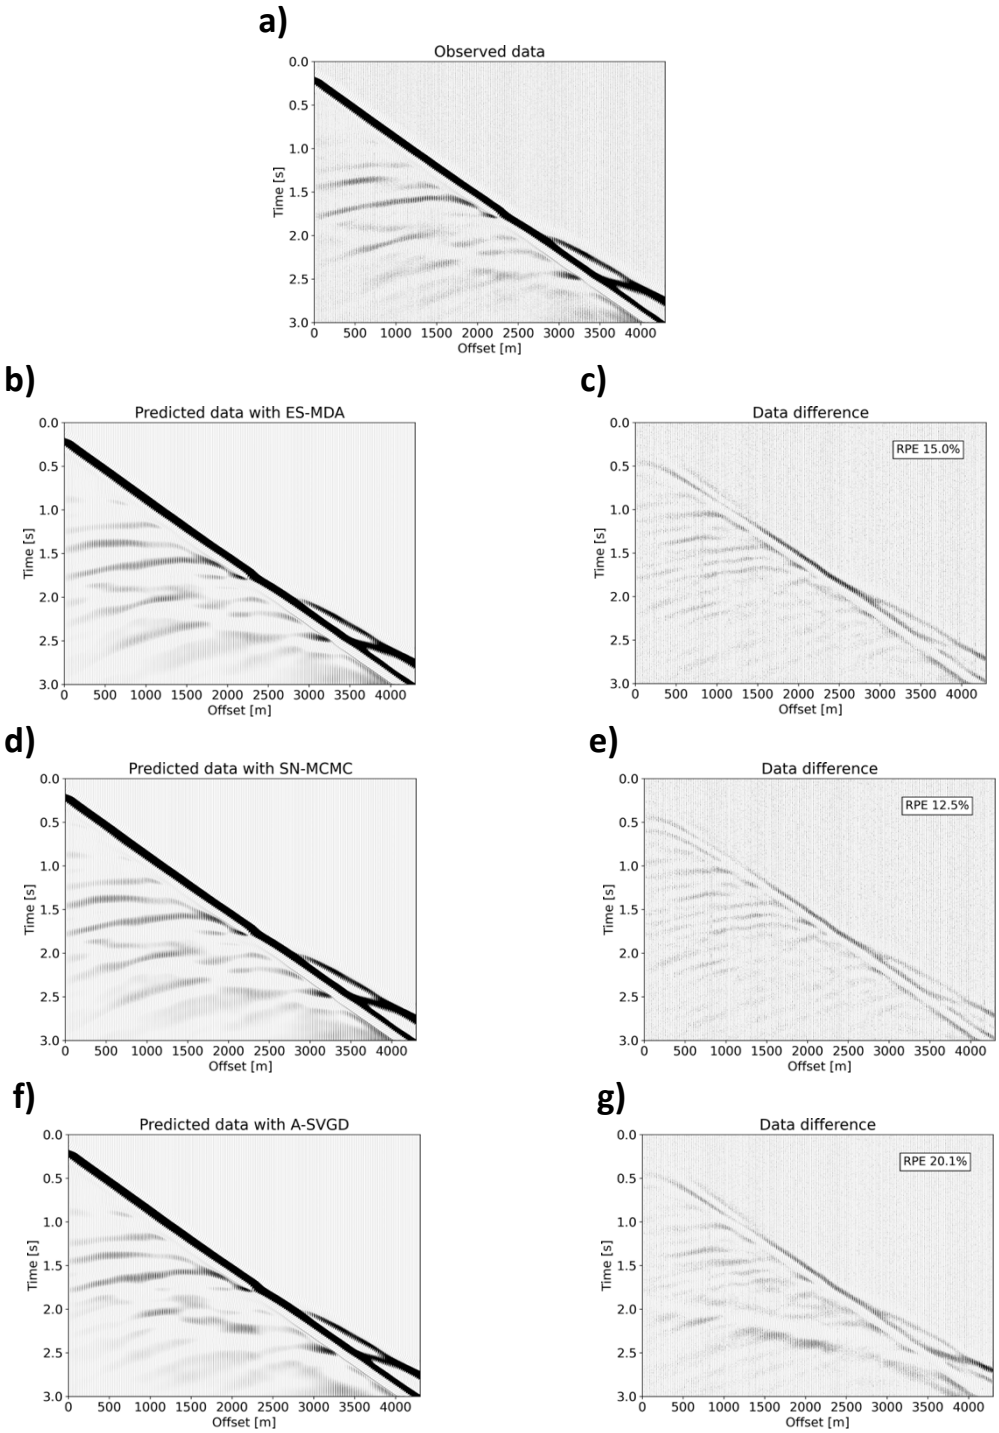

Fig. 11 *Comparison of the observed and predicted seismograms associated to the inversion results of Test 1, highlighting their relative percentage error with respect to the observed data*

Figure 11 compares the inversion results in terms of data prediction accuracy and relative data residual. Specifically, Figure 11a displays the first shot of the observed data, while Figures 11b-d-f show the same shot computed using the final mean models from Figure 9. Instead, Figures 11c-

e-g show the sample-by-sample differences between the predicted data and the observed one, for each inversion method. The ES-MDA and SN-MCMC solutions yield very similar relative percentage errors (RPE), confirming their comparable accuracy in reproducing the observed data. The residual differences visible in the seismograms can be partially attributed to the lower resolution of the model predictions with respect to the true model, a consequence of DCT compression. As expected, the A-SVGD result exhibits a slightly higher data misfit, with visible residual differences, particularly at larger offsets. These discrepancies are consistent with the velocity underestimation in the deeper regions, as shown in Figure 9e.

To further assess the performance of the three approaches, we analyzed local correlation maps corresponding to distinct subdomains of the velocity model. Specifically, we computed these maps for two arbitrary selected cells, marked by "+" symbols in Figure 12, each located at the center of a different subdomain (the colored rectangles in Figure 12a). These subdomains represent spatially distinct regions of the velocity model, allowing for a comparative evaluation of how each method captures correlations between a model parameter and the surrounding ones (Figures 12b-c).

The local correlation maps were generated based on the results of Test 1, considering the final 500 particles from the A-SVGD inversion, the samples drawn from each chain of the SN-MCMC inversion after the burn-in phase and all the models in the final ensemble of the ES-MDA inversion.

Each correlation coefficient, ranging between -1 and 1, provides insight into the relationships between model parameters. Values close to 1 indicate a strong positive correlation, suggesting that the parameters share similar physical properties. Conversely, values near -1 reflect a strong inverse relationship, implying that changes in one parameter correspond to opposite changes in the other. A coefficient close to 0 implies no significant correlation between the parameters.

Across the two subdomains, the correlation maps obtained from the three methods exhibit a high degree of consistency, reinforcing their ability to capture meaningful parameter relationships. In the first region (Figure 12b, corresponding to the magenta subdomain), the correlation maps clearly show negative correlations (in blue) aligned along the axis of a dipping high-velocity reflector. The boundaries of this feature are well-defined across all methods. In the second subdomain (black; Figure 12c), we observe predominantly positive correlation values across most of the map, except for the left side, where negative correlations emerge. This pattern closely follows the geometry of two dipping high-velocity layers visible in the true model. The reduced spatial resolution of the outcomes of this correlation analysis, when compared to the true model, is related to the effect of DCT compression.

The overall consistency observed in this analysis of the correlation maps generated by the three methods underscores their ability to capture complex relationships between model parameters, particularly those associated with geological features such as dipping layers and reflectors. This

is especially encouraging because it confirms that model parameters within the same interface or geological structure tend to exhibit similar correlation properties. Furthermore, the observed consistency enforces the reliability of the three inversion methods for uncertainty estimation and sensitivity analysis.

a)

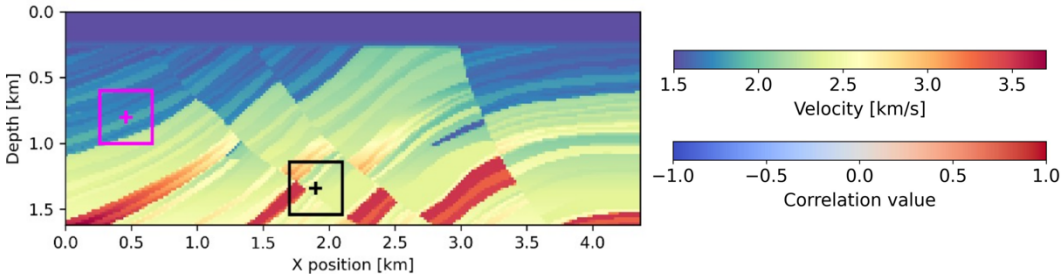

b)

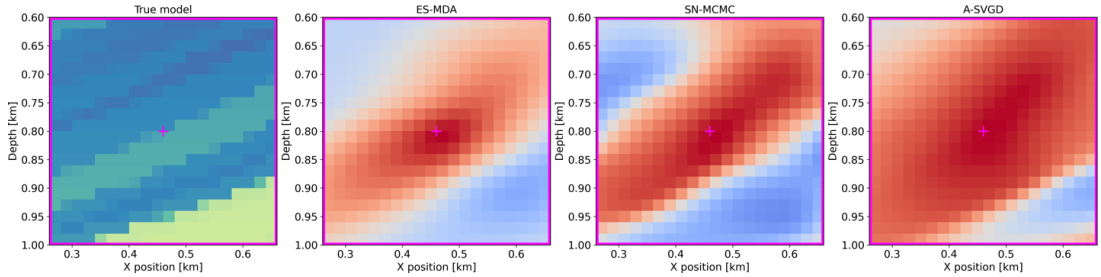

c)

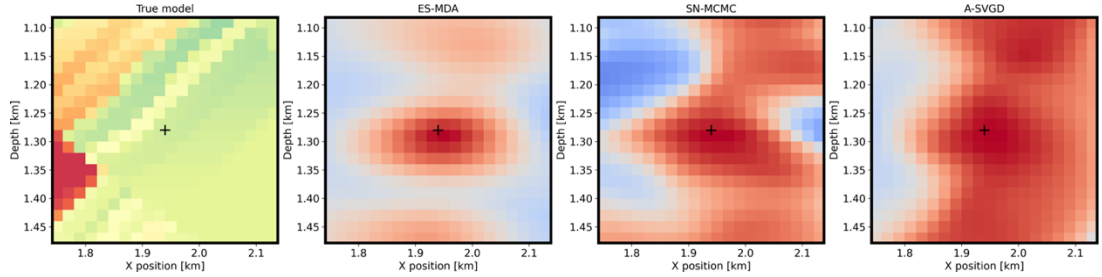

Fig. 12 *True model with highlighted two distinct subdomains (colored rectangles) and their relative local correlation maps for the model parameter highlighted by the "+" symbol, calculated using the three inversion approaches*

Table 1 provides a quantitative comparison of the results obtained with the three inversion approaches in this first test, using three different evaluations metrics: the SNR using as reference the true model shown in Figure 3, the RPE with respect to the observed data shown in Figure 11a and the 99% coverage ratios (C.R.). These results confirm the excellent performances of both SN-MCMC and ES-MDA in terms of prediction accuracy and data misfit, while the solution of the A-SVGD algorithm is affected by some velocity underestimations which reflect in a lower SNR and higher RPE. In terms of uncertainty estimation, the coverage ratio confirms the comparable

performances of the SN-MCMC and A-SVGD approaches, while the ES-MDA is affected by a severe underestimation, due to the previously mentioned ensemble-collapse issue.

To compare the computational workload of the three methods, Table 1 also reports the total number of forward model evaluations required to obtain the presented results, assuming five shot gathers per simulation. The large number of forward evaluations associated with SN-MCMC arises from the need to numerically compute the Jacobian matrix for each accepted model. Each Jacobian evaluation requires a number of forward simulations equal to the number of model parameters, although this cost is significantly reduced by the DCT compression. Moreover, a large number of iterations is required to ensure adequate sampling of the PPD. Using the acceptance rates shown in Figure 8a, we estimated the total number of forward evaluations for a single Markov chain, as reported in Table 1. For ES-MDA, the number of forward simulations is computed as the product of the ensemble size and the number of iterations. Similarly, for A-SVGD, the total number of forward evaluations is given by the number of particles multiplied by the number of iterations. In the subsequent tests, the number of forward evaluations for ES-MDA and A-SVGD remains the same as reported here, whereas for SN-MCMC it may vary slightly, since it depends on the number of accepted models.

Table 1 *Comparison of the results obtained with all methods for Test 1, using the three different metrics considered to quantify model prediction, data mismatch and uncertainty estimation. The number of forward evaluations for each method is also reported*

	SN-MCMC	**ES-MDA**	**A-SVGD**
SNR	23.6	23.9	18.3
RPE (%)	12.5	15.0	20.1
C.R. (%)	35.2	10.1	32.1
No. forward evaluations	$\sim 3.00 \cdot 10^5$ (for one chain)	$8.00 \cdot 10^4$	$1.75 \cdot 10^5$

3.2 Test 2: less informative prior information

In this second test, we analyse the impact of different initializations on the inversion results.

As in the first test, we assume a perfectly known source wavelet and introduce uncorrelated Gaussian noise to the observed data (SNR around 10 dB). The data covariance matrix is once again created assuming this same level of noise. The key difference in this test lies in the assumed prior mean model. Here, instead of using a smoothed version of the true model, we adopt a homogeneous velocity model with a constant velocity of 2 km/s as the prior mean for the

Gaussian distribution from which prior realizations are drawn. Figure 13 shows this prior mean and some models drawn from the corresponding prior distribution. Once again, for the SN-MCMC inversion, the five chains are initialized using five randomly selected velocity models from this ensemble.

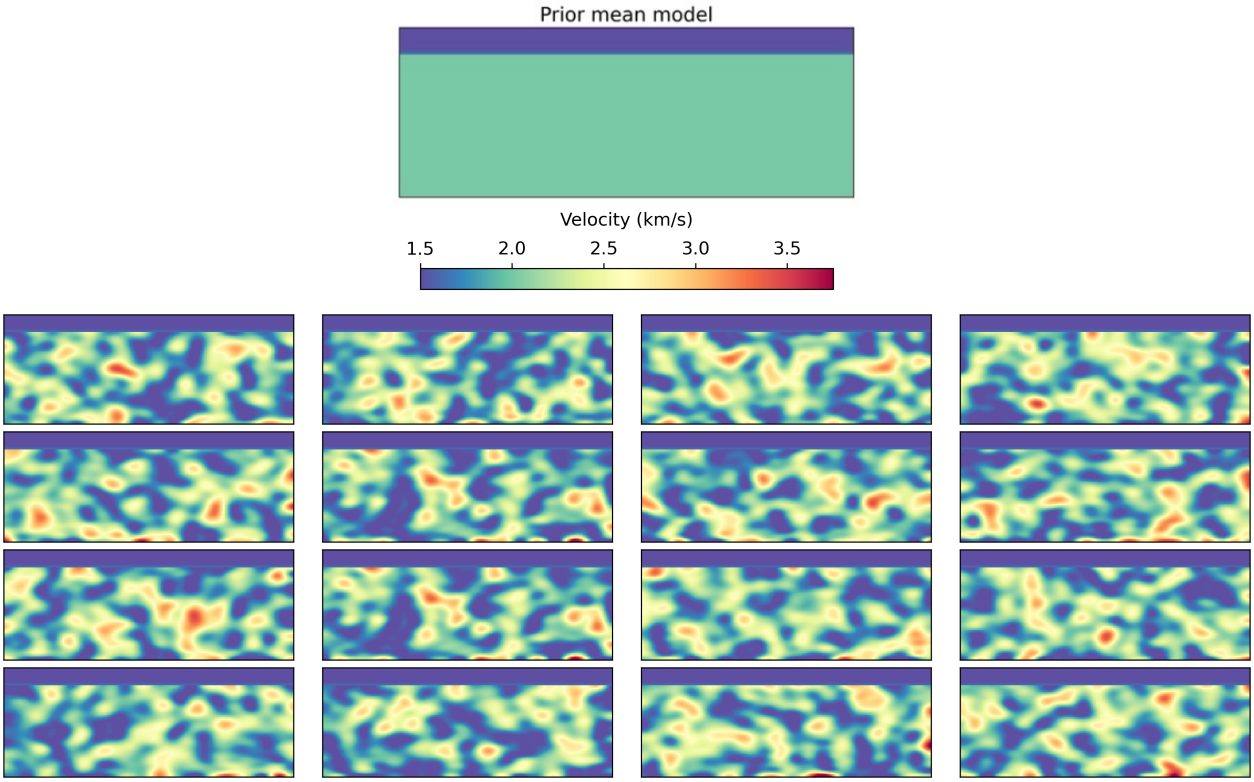

Fig. 13 *The homogeneous prior mean model in Test 2, and some realizations drawn from the prior distribution*

Figure 14 presents the results of this test for all three inversion approaches, illustrating the posterior mean models and corresponding standard deviation maps. The SN-MCMC mean model (Figure 14c) closely resembles the result obtained in Test 1 (Figure 9c), demonstrating the robustness of this approach in achieving consistent posterior estimates, regardless of the initialization. This behaviour is further validated by the corresponding standard deviation map (Figure 14d), which exhibits a spatial distribution and magnitude of uncertainties similar to those observed in the first test (Figure 9d)—characterized by increasing uncertainty with depth and higher standard deviations around high-velocity layers and the associated reflecting interfaces (with a 99% C.R. of approximately 50%).

For the ES-MDA inversion, the reconstructed velocity field (Figure 14a) successfully captures the primary features of the true model, with only minor artifacts appearing at greater depths. This results in a slightly reduced SNR compared to Test 1. The associated standard deviation map

(Figure 14b) again highlights a significant underestimation of uncertainties due to the ensemble collapse (which corresponds to a very low 99% C.R.), though the relative uncertainty trend remains similar to that of the SN-MCMC inversion (i.e. increasing uncertainty with depth and higher standard deviations around high-velocity layers and the associated reflecting interfaces). In contrast, the A-SVGD result (Figure 14e) shows a considerable deterioration in performance compared to Test 1 (Figure 9e). While the mean model accurately reconstructs the uppermost kilometer of the subsurface, deeper regions are dominated by low-velocity anomalies. Consequently, the SNR is substantially lower than that of the other two methods. Additionally, the standard deviation map does not exhibit the same behaviour seen in Figure 9f, with high uncertainties also arising in the shallower portions of the model. The 99% C.R. is considerably lower than the one shown in the previous test, due to the poor performance in terms of velocity predictions in the deeper regions of the model.

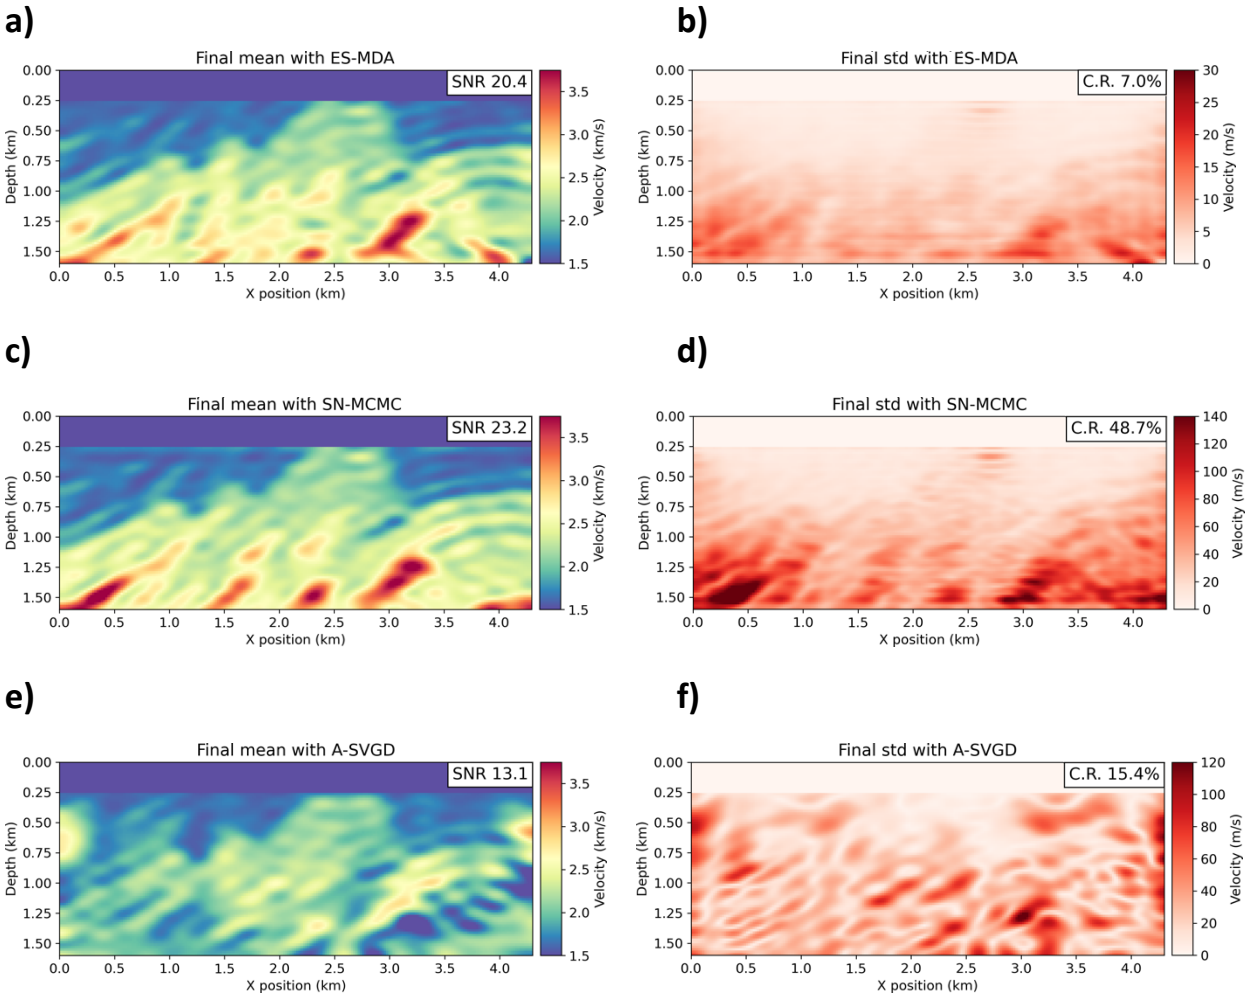

Fig. 14 *Results for Test 2: on the left column, the final mean models obtained with all three inversion approaches, with specified their SNR using the true model as reference; on the right column, the corresponding standard deviation maps with their 99 % coverage ratios (C.R.)*

Figures 15b-d presents the marginal posterior distributions obtained using the three inversion approaches for three selected cells, located at different depths and horizontal positions and highlighted by the magenta dots in Figure 15a, with the prior distribution centered around 2 km/s for all the cells. In the shallowest cell (Figure 15b), all three posterior marginal distributions are closely aligned with the true velocity value, indicated by the vertical black dashed line. This confirms that, in the shallower portions of the model, all three inversion methods show a good performance. For the deeper cell located at the center of the model (Figure 15c), the A-SVGD method underestimates the velocity value of approximately 400 m/s, the ES-MDA suffers from a slight overestimation of the actual velocity, while SN-MCMC provides the most accurate estimate, that is almost perfectly centered around the true value. For the cell located near the lateral edge of the model (Figure 15d), both ES-MDA and SN-MCMC yield distributions that are closely aligned with the true velocity, whereas A-SVGD again shows an underestimation of the actual velocity value.

As expected, the SN-MCMC marginal distribution exhibits the largest variance in all cases, while the ones associated to the ES-MDA inversion reflect the ensemble collapse effect with very low variances.

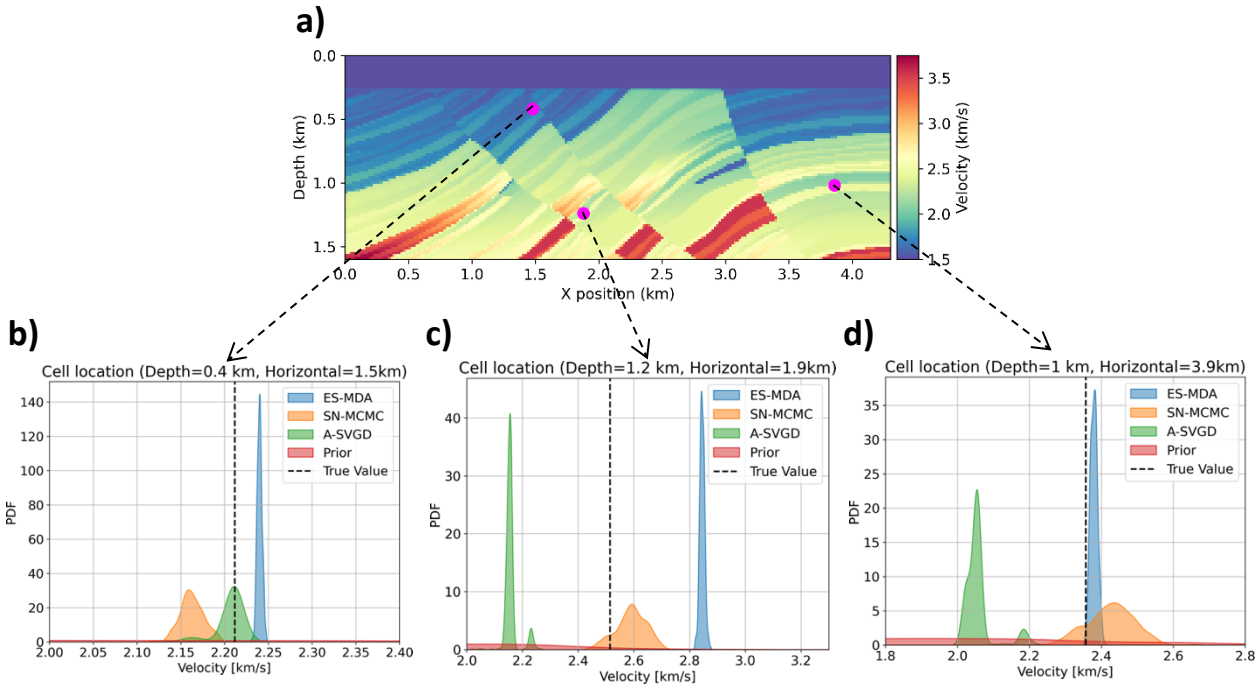

Fig. 15 *Marginal distributions associated with three different cells highlighted by magenta dots in the true model. Blue, orange and green represent the marginal posterior distributions for the ES-MDA, SN-MCMC and A-SVGD algorithms, respectively. Red indicates the prior distribution, and the black vertical dashed line represents the true velocity value*

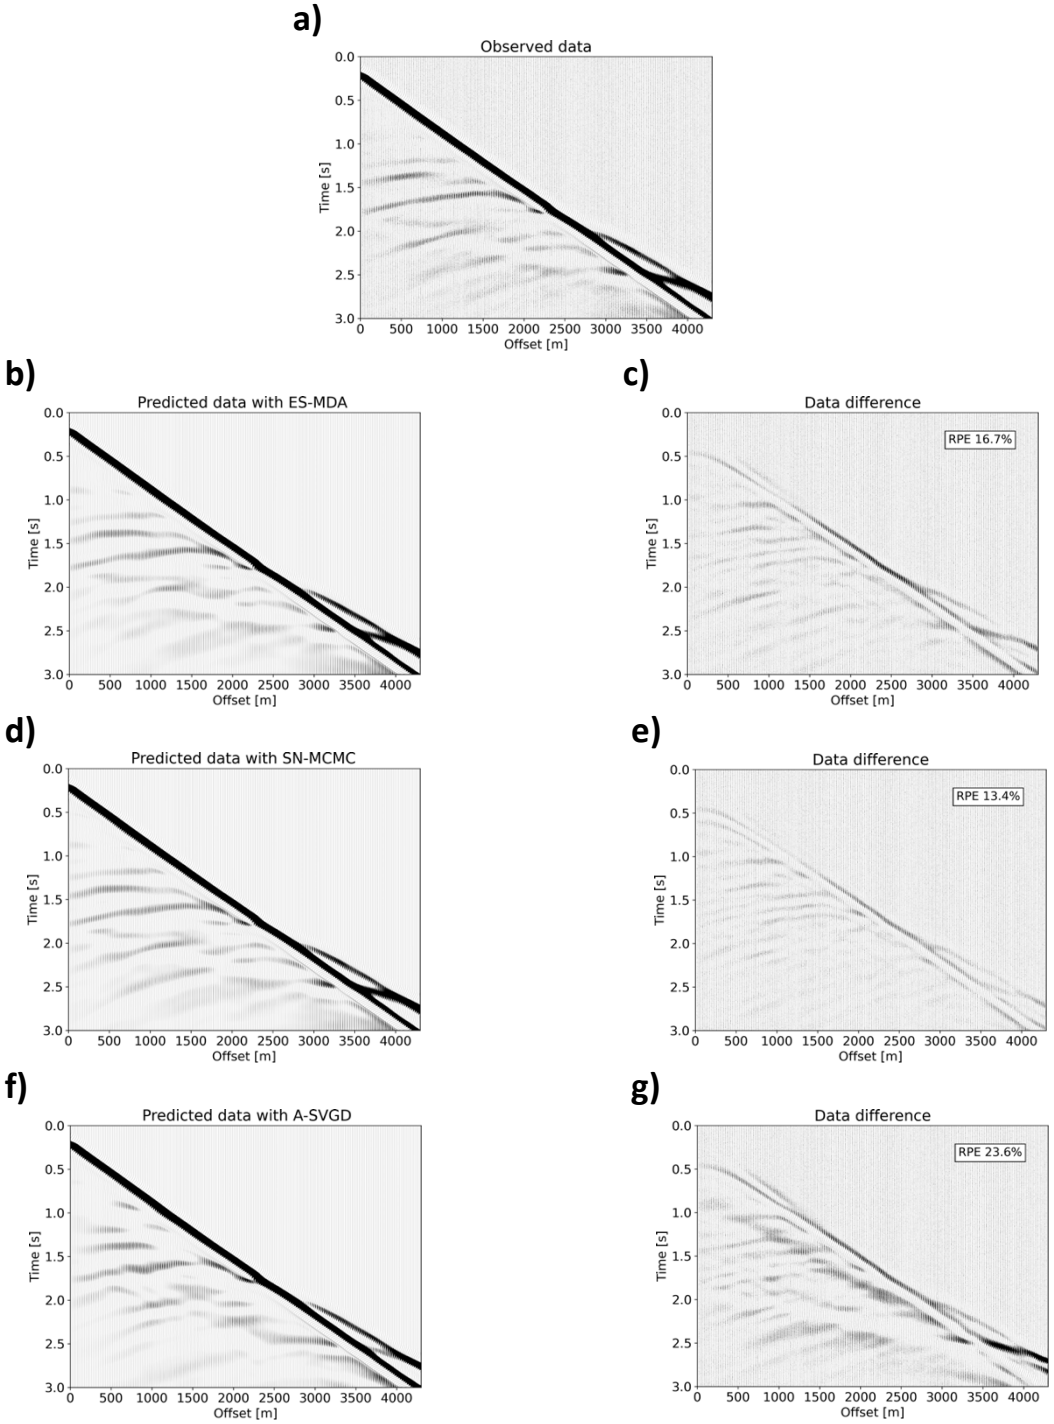

Fig. 16 *Comparison of the observed and predicted seismograms associated to the inversion results of Test 2, highlighting their relative percentage error with respect to the observed data*

Figure 16 presents the results in terms of data prediction, considering the first shot of the observed data (Figure 16a), the corresponding shot (Figures 16b-d-f) calculated on the posterior mean models of Figures 14a-c-d and their relative sample-by-sample differences (Figures 16c-e-g). The SN-MCMC and ES-MDA inversions yield excellent data predictions, comparable to those achieved in Test 1 (Figure 11), although the ES-MDA result exhibits slightly larger misfits at mid-to-far offsets. On the other hand, the A-SVGD inversion produces a significantly higher data misfit (with an RPE of 23.6%), with the largest discrepancies occurring at large offsets—consistent with the low-velocity anomalies observed in the deeper regions of the model.

Table 2 summarizes the performance of the three Bayesian algorithms evaluated using the same metrics employed for the first test. In this case, SN-MCMC achieves the highest SNR and the lowest RPE, almost identical to the ones shown for Test 1, proving again the robustness of this approach to the starting point of the inversion. The ES-MDA achieves accurate results, slightly worse than those obtained in the first test, but again significantly underestimated the model uncertainties. Finally, the A-SVGD shows the worst performance, proving its strong dependence on the initialization, likely due to its deterministic nature.

Table 2 *Comparison of the results obtained with all methods for Test 2, using the three different metrics considered to quantify model prediction, data mismatch and uncertainty estimation*

	SN-MCMC	**ES-MDA**	**A-SVGD**
SNR	23.2	20.4	13.1
RPE (%)	13.4	16.7	23.6
C.R. (%)	48.7	7.0	15.4

3.3 Test 3: realistic conditions

The third test evaluates the robustness of each inversion method under more realistic and challenging conditions. In this scenario, we introduced a higher level of Gaussian random noise, reducing the data SNR to 4 dB. The data covariance matrix was instead created assuming a slightly higher level of uncorrelated noise corresponding to a SNR of 3 dB. This practically means that we are overestimating the actual noise affecting the data. Additionally, to simulate missing data, 20% of the traces in each seismogram were randomly zeroed. To further increase the complexity, the simulated data were generated using a 20-degree phase-shifted version of the source wavelet, differing from the one used for the observed data.

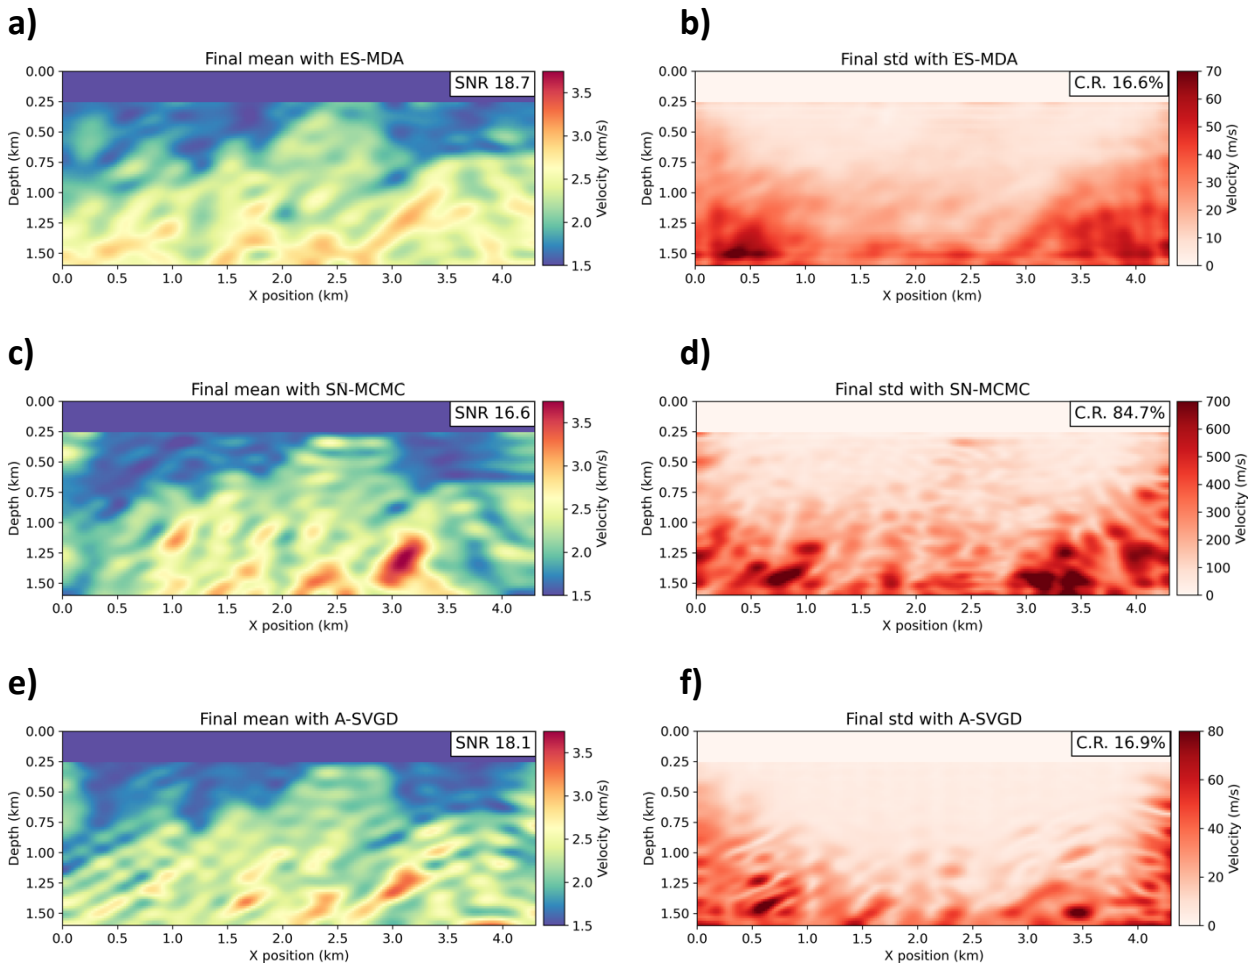

Fig. 17 *Results for Test 3: on the left column, the final mean models obtained with all three inversion approaches, with specified their SNR using the true model as reference; on the right column, the corresponding standard deviation maps with their 99 % coverage ratios (C.R.)*

Figure 17 presents the results of the three inversion algorithms in terms of posterior mean models and standard deviation maps. In this case, both SN-MCMC (Figure 17c) and ES-MDA (Figure 17a) inversions exhibit a general underestimation of velocity values across the entire model, with noticeable artifacts at the model's bottom and lateral edges. In contrast, the A-SVGD result (Figure 17e) remains almost identical to the one shown in Figure 9e, as confirmed by the similar SNR values. This indicates greater robustness of A-SVGD in handling degraded data conditions, in comparison to SN-MCMC and ES-MDA. This probably depends on the combination of the DCT compression, that mitigates the ill-posedness of the problem, and the fact that we are satisfying the ideal condition for SVGD (number of particles equal or greater than the number of unknowns).

The standard deviation maps for the ES-MDA (Figure 17b) and A-SVGD (Figure 17f) algorithms show similar patterns and magnitudes. Notably, the higher level of assumed noise in the data covariance matrix appears to have partially mitigated the ensemble collapse issue in ES-MDA, as evidenced by the comparable 99% C.R. values obtained with these two methods.

Since the inverse of the data covariance matrix is used in the model update, an increased noise level reduces the Kalman gain, leading to smaller updates to the velocity model. Indeed, the algorithm interprets the data as being less informative, hence the Kalman gain gives more weight to the prior producing small update in the parameters. This mechanism tends to preserve diversity among the ensemble members, resulting in higher estimated uncertainties at the end of the inversion.

Conversely, the SN-MCMC standard deviation map (Figure 17d) exhibits a similar spatial distribution of uncertainties but with significantly larger magnitudes (up to 700 m/s). This results in a substantial increase in the 99% C.R., exceeding 85%, indicating a higher level of uncertainty in the posterior distribution. Notably, this is associated with a posterior mean model that is less accurate than the SN-MCMC results in the other tests. Therefore, the strong increment in the 99% C.R. effectively supports that the model in Figure 17c is predicted with lower precision.

Figure 18 presents two 1D velocity profiles extracted from two different spatial locations (highlighted by the vertical magenta dashed lines in Figure 18a), comparing the velocities derived from the true model before and after DCT compression with those obtained from the final mean models of the three inversion approaches. The shaded regions indicate the 95% and 99% confidence intervals for each method, calculated using the corresponding standard deviation values.

As expected, the SN-MCMC method (Figures 18c-f) exhibits the widest confidence intervals, with both the true model and its DCT-compressed version lying almost entirely within these bounds. The uncertainty estimates also follow a natural trend, increasing with depth due to the reduced illumination. In contrast, the ES-MDA (Figures 18b-e) and A-SVGD (Figures 18d-g) approaches show significantly narrower confidence intervals, reflecting their lower estimated variance. Both methods also exhibit a slight underestimation of velocity, particularly at larger depths.

These findings align well with the previously discussed coverage ratios, where SN-MCMC achieves the highest value, while ES-MDA and A-SVGD yield significantly lower magnitudes.

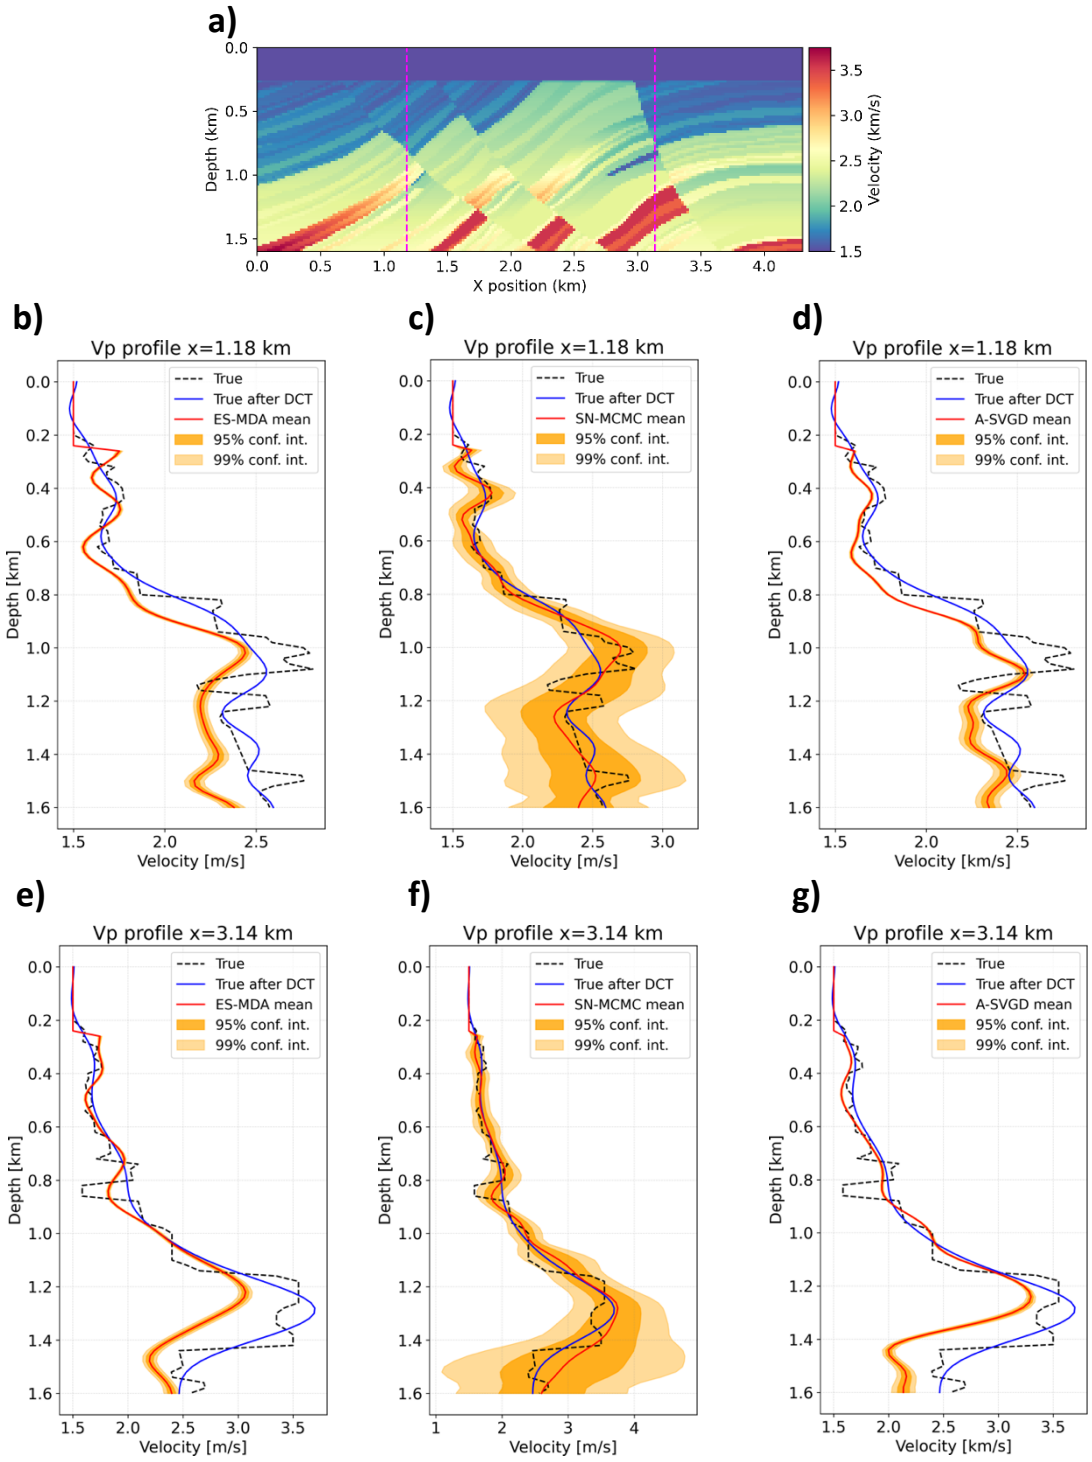

Fig. 18 *Pseudo-well velocity profiles associated to two different horizontal locations (dashed magenta lines in the true model), showing mean and two confidence intervals for all three inversion methods, together with the velocity values extracted from the true model before (dashed black) and after (blue) DCT compression*

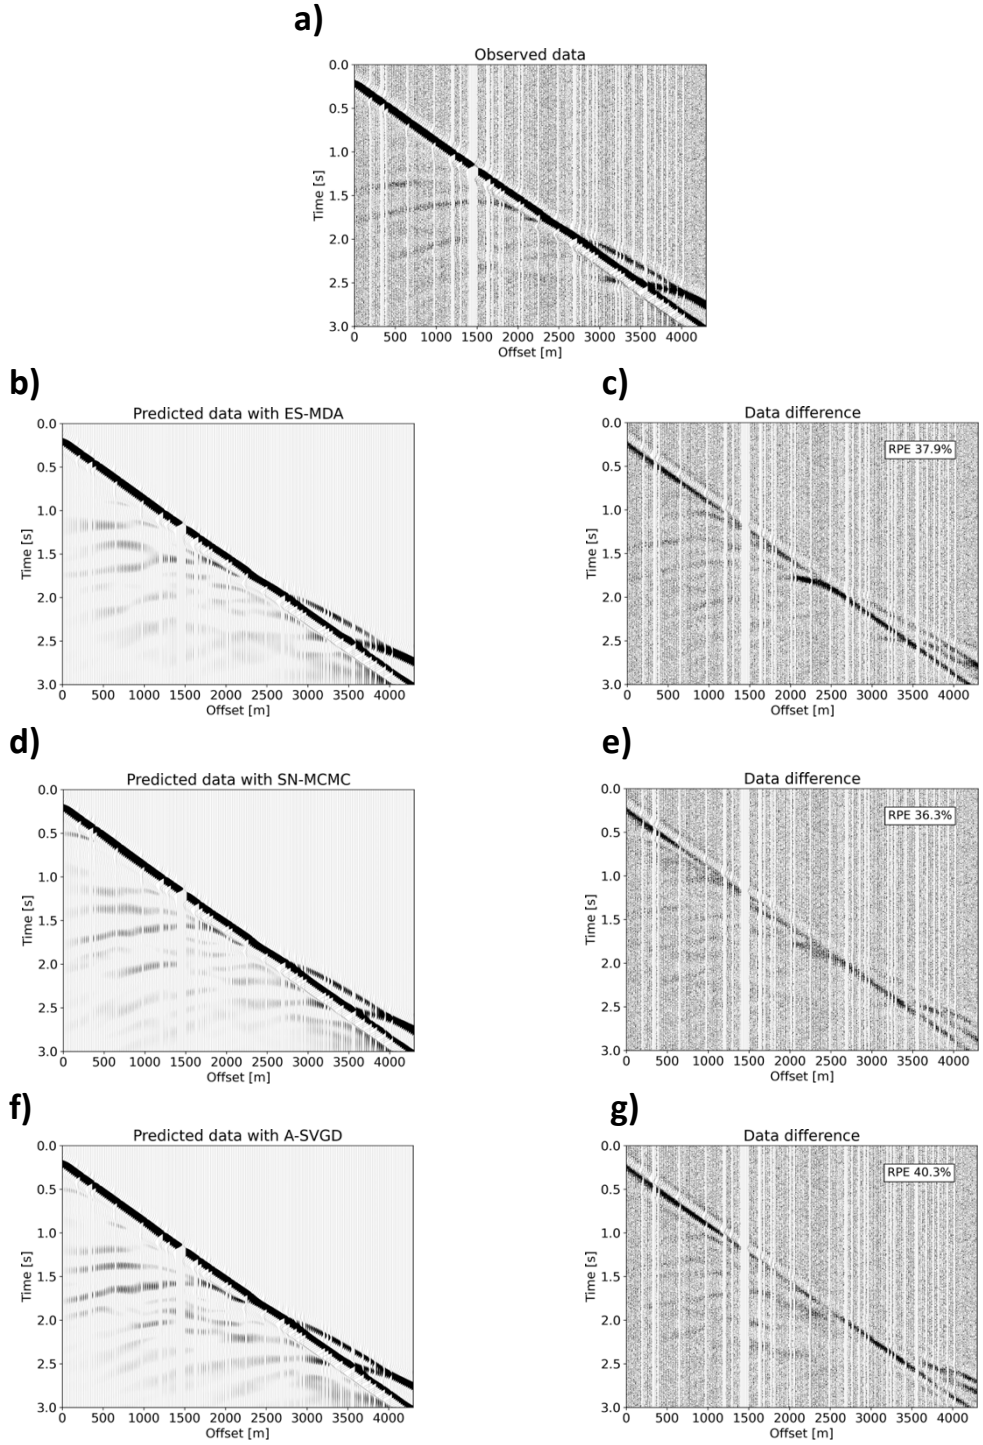

Fig. 19 *Comparison of the observed and predicted seismograms associated to the inversion results of Test 3, highlighting their relative percentage error with respect to the observed data*

In terms of data prediction, Figure 19 compares the first shot gather of the observed data (Figure 19a) and the corresponding shot gathers (Figures 19b-d-f) calculated using the final mean model of each inversion approach, shown in Figure 17a-c-e. Additionally, we display the sample-by-sample difference between the observed and predicted data (Figures 19c-e-g).

All three inversion methods achieve similar percentage errors (around 35-40%), which are significantly higher than those observed in the previous two tests, due to the higher level of noise and to the phase shift applied to the source wavelet employed for the simulated data. These discrepancies are particularly noticeable at mid-to-far offsets, but substantial mismatches consistently appear below the first arrival, primarily due to the wavelet phase shift.

As we did for Test 1, to further assess the performance of the three approaches, we analyzed local correlation maps corresponding to the same two subdomains of the velocity model used before. These maps were computed for the two cells marked by "+" symbols in Figure 20, each located at the center of a different subdomain, highlighted by the colored rectangles in Figure 20a. The local correlation maps were generated based on the results of Test 3, considering the final 500 particles from the A-SVGD inversion, the post-burn-in samples from each SN-MCMC chain and all the models from the final ES-MDA ensemble. These maps, displaying correlation coefficients ranging from -1 and 1, offer valuable insights into the relationships between different model parameters.

In the first subdomain (in magenta), located in a shallower region of the model, the correlation maps obtained from all three methods show a strong degree of consistency, even when compared with the same maps obtained for Test 1 (Figure 12b), showing negative correlation values (in blue) in correspondence of the dipping high-velocity reflector present in the true model. In the second subdomain (black), positive correlation values emerge near the analyzed cell in the center (as in Test 1, Figure 12c), while negative correlations appear toward the edges of the subdomain. This effect is particularly pronounced on the left side, which corresponds to the two dipping high-velocity layers visible in the true model.

The good consistency shown in this analysis of the correlation maps generated by the three methods, even when compared to those obtained in Test 1 under optimal conditions, underscores their capability to capture complex relationships between model parameters also in more challenging inversion scenarios.

Table 3 provides a comparison of the three Bayesian algorithms, in terms of the same metrics employed for the first two tests. We can observe the worse performance of SN-MCMC and ES-MDA compared to the previous tests, in terms of SNR and RPE, now comparable with those achieved using the A-SVGD algorithm, which yields a final solution similar to those shown in Tests 1 and 2. Notably, under these more challenging conditions, the SN-MCMC has reached a coverage ratio above 80%, due to the strong increase in estimated uncertainty, which reflects the lower precision associated with the inversion result, compared to the previous tests.

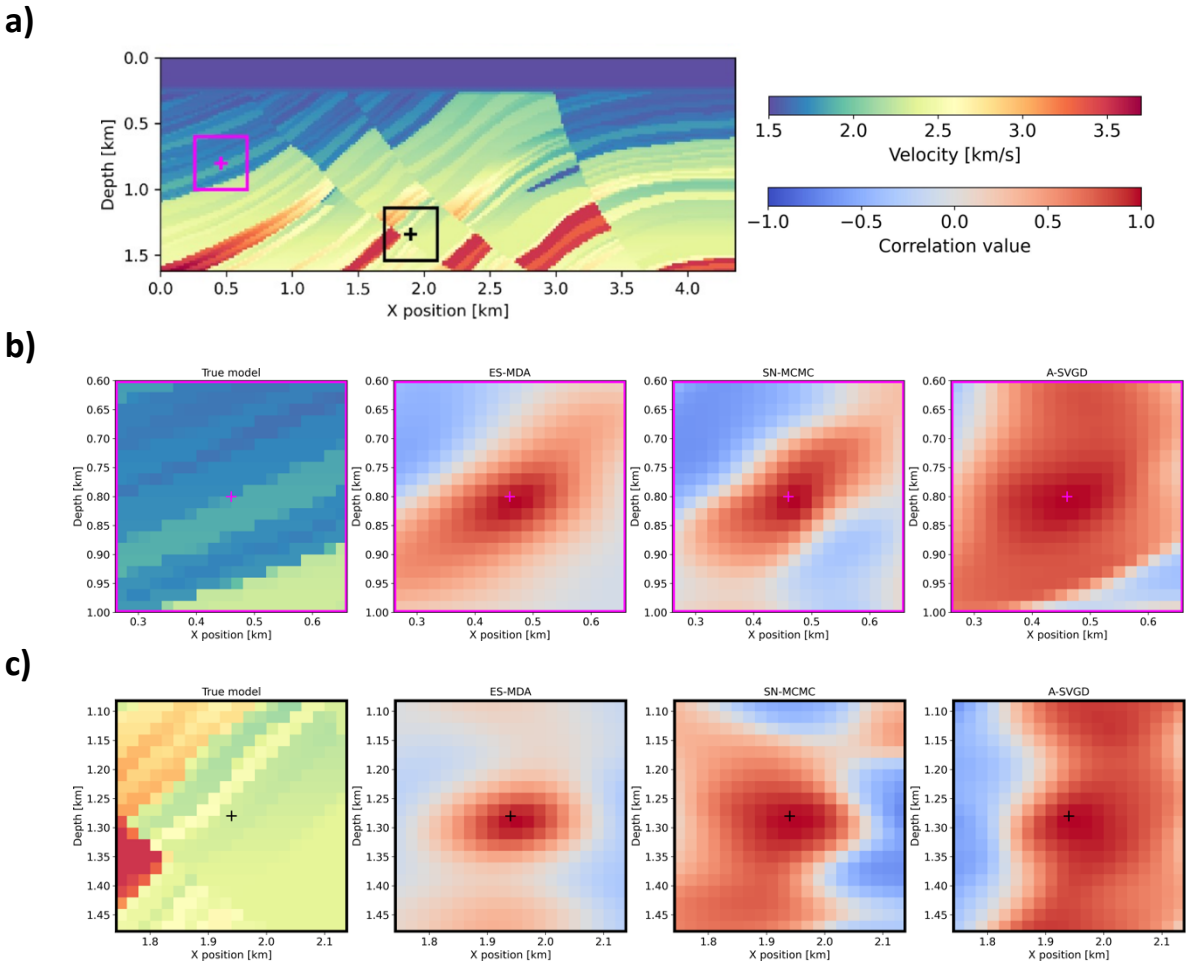

Fig. 20 *True model with highlighted two distinct subdomains (colored rectangles) and their relative local correlation maps for the model parameter highlighted by the "+" symbol, calculated using the three inversion approaches*

Table 3 *Comparison of the results obtained with all methods for Test 3, using the three different metrics considered to quantify model prediction, data mismatch and uncertainty estimation*

	SN-MCMC	**ES-MDA**	**A-SVGD**
SNR	16.6	18.7	18.1
RPE (%)	36.3	37.9	40.3
C.R. (%)	84.7	16.6	16.9

4. Discussion

This study evaluates and compares the performance of three Bayesian approaches, the SN-MCMC, ES-MDA and A-SVGD, for acoustic FWI under three different conditions. The first scenario assumes an optimal setting, where the source wavelet and noise level are perfectly known, and the starting points of the inversions are derived from a smoothed version of the true velocity model. In the second case, we assess the effect of a homogeneous prior mean on the inversion results, testing the sensitivity of each method to limited prior information. The third and most challenging scenario introduces more realistic conditions, including an increased level of noise, some missing traces in the data and incorrect assumptions for both wavelet and noise. The comparison focuses on the accuracy and precision of each method in terms of model prediction, data misfit and uncertainty quantification, sensitivity to the initialization and robustness across different inversion scenarios.

In all the considered tests, the assumed noise statistics are consistent with the noise generation procedure, and therefore the target posterior is unbiased within this controlled framework. However, in case of field applications, where data is dominated by temporally correlated noise and modelling errors, a diagonal covariance assumption would be inadequate.

Among the three methods, ES-MDA proves to be the most computationally efficient: A-SVGD requires approximately twice the number of forward modelling evaluations needed by ES-MDA, whereas SN-MCMC needs more than three times as many. ES-MDA does not need an explicit computation of Jacobians and Hessians (like SN-MCMC) or gradients (like A-SVGD), which are statistically estimated from the ensemble (van Leeuwen and Evensen 1996; Moyen and Gentilhomme 2021), leading to rapid convergence. Additionally, it demonstrates strong robustness against initialization, providing accurate results even when minimal prior information is available. However, the application of ES-MDA is constrained by its reliance on Gaussian priors and the assumption of relatively linear forward modelling operator. The primary limitation of this method is its tendency to significantly underestimate the uncertainty affecting the final solution (Chen and Oliver 2010, 2014), which is a critical drawback for a Bayesian approach to FWI.

To mitigate this issue, larger ensemble sizes can be used (Roe et al. 2016), which however would impact the main benefit of this approach, its computational efficiency. Some alternative strategies include covariance inflation (Anderson and Anderson 1999; Evensen 2009b), which artificially increases the ensemble spread to prevent underestimation of uncertainties, and iterative regularization (Kaltenbacher et al. 2008; Iglesias 2015; Chada et al. 2021) to stabilize the inverse problem. Another widely used is localization (Houtekamer and Mitchell 2001; Chen and Oliver 2010; Reynolds et al. 2014), which reduces spurious long-range correlations arising from limited ensemble sizes by limiting updates to local regions.

Localization is typically implemented through element-wise tapering of the covariance or Kalman gain using distance-based correlation functions such as Gaspari and Cohn (1999). However, in our case, spatial distance is not directly defined because both model and data are represented in the DCT domain. Alternative localization strategies that do not rely on physical distance have been proposed by Anderson (2007), and comparative assessment are provided by Lacerda et al. (2019).

The SN-MCMC method is the most computationally demanding approach because it requires repeated Jacobian evaluations, which are computed here through a finite difference approach. Its computational cost is strictly related to the acceptance ratio of the Markov chain, since a new Jacobian must be recomputed for each accepted model. To reduce this burden in cases of high acceptance rates (e.g., over 40% as in Test 1), we recompute the Jacobian only at fixed intervals after the burn-in phase (in our case, every 5 accepted samples), when all the sampled models are very similar between each other.

Despite its computational demands, SN-MCMC is the most robust method against initialization effects, consistently yielding optimal results in terms of both accuracy and precision, even with minimal prior information. Under more complex conditions (as in Test 3), the method produces less accurate models but compensates with significantly higher uncertainty estimates, leading to extremely high coverage ratio values. Furthermore, SN-MCMC is theoretically capable of handling any type of prior distribution, whether parametric or non-parametric.

A-SVGD, on the other hand, exhibits an intermediate computational cost compared to the other two methods. In case of complex scenarios (Test 3), it emerges as the most robust approach, yielding results almost identical to those obtained under optimal conditions, with comparable levels of accuracy and precision. However, A-SVGD is also the most sensitive to initialization, obtaining the poorest results when a less informative prior is used. This sensitivity is likely due to the deterministic nature of this approach, which behaves more like an optimization algorithm than a purely stochastic sampling method. The annealing schedule implemented in A-SVGD mitigates the mode collapse issue typically associated with standard SVGD methods, but it is not sufficient when the prior is overly restrictive (as observed in Test 2).

Regarding uncertainty estimation, A-SVGD provides more reasonable results compared to ES-MDA, achieving higher coverage ratio values. However, it still exhibits some underestimation of the variance, typical of SVGD methods. Further refinements in the annealing process or alternative repulsive interaction terms could potentially improve its ability to fully capture posterior uncertainty.

5. Conclusions

In this work, we presented a comparison of three Bayesian algorithms for acoustic FWI: Stochastic Newton Markov Chain Monte Carlo (SN-MCMC), Ensemble Smoother with Multiple Data Assimilation (ES-MDA) and Annealed Stein Variation Gradient Descent (A-SVGD).

A key aspect of our implementations is the use of Discrete Cosine Transform (DCT) to compress both model and data spaces, reducing computational demands and mitigating the ill-posedness of the considered inverse problem.

To assess each method's performance, we designed three distinct scenarios: an ideal baseline, a case with limited prior information and a more realistic setting, featuring a high level of noise, missing data and wrong assumptions on the source wavelet.

Among the tested methods, SN-MCMC proved to be the most reliable in capturing posterior uncertainties, though at a significantly higher computational cost. ES-MDA emerged as the most computationally efficient, but shows a strong uncertainty underestimation, due to the ensemble collapse issue, a well-known limitation of ensemble-based methods. A-SVGD offered a good-trade-off between robustness and computational effort across the different scenarios but showed sensitivity to the choice of the prior, due to its deterministic nature. These results serve as algorithmic benchmarks rather than definitive conclusions for field-scale FWI. Extending the comparison to large-scale applications and investigating acquisition-dependent uncertainty behavior remain important directions for future work.

Additional avenues for future research include exploring hybrid strategies that leverage the complementary strengths of these methods. Potential improvements may be obtained through the integration of ES-MDA with adaptive inflation strategies to mitigate variance collapse and relax the need to set the number of iterations before the inversion. Another possibility is to enhance A-SVGD with stochastic sampling techniques to improve posterior exploration while preserving its robustness.

Declarations

The authors have no relevant financial or non-financial interests to disclose.

References

Agostinetti NP, Malinverno A (2010) Receiver function inversion by trans-dimensional Monte Carlo sampling. Geophysical Journal International, 181(2), 858–872.

Ahmed N, Natarajan T, Rao K (1974) Discrete cosine transform. IEEE Transactions on Computers, C-23(1), 90-93.

Aleardi M, Ciabarri F (2017) Application of different classification methods for litho-fluid facies prediction: a case study from the offshore Nile Delta. Journal of Geophysics and Engineering, 14(5), 1087–1102.

Aleardi M (2020a) Combining discrete cosine transform and convolutional neural networks to speed up the Hamiltonian Monte Carlo inversion of pre-stack seismic data. Geophys. Prospect., 68, 2738–2761.

Aleardi M (2020b) Discrete cosine transform for parameter space reduction in linear and non-linear AVA inversions. Journal of Applied Geophysics, 179, 104106.

Aleardi M, Salusti A (2020) Hamiltonian Monte Carlo algorithms for target- and interval-oriented amplitude versus angle inversions. Geophysics, 85(3).

Aleardi 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.

Aleardi M, Vinciguerra A, Hojat A (2021) Ensemble-Based Electrical Resistivity Tomography with Data and Model Space Compression. Pure Appl. Geophys. 178, 1781–1803.

Aleardi M, Vinciguerra A, Stucchi E, Hojat A (2022) Stochastic electrical resistivity tomography with ensemble smoother and deep convolutional autoencoders. Near Surface Geophysics, 20(2), 160-177.

Anderson JL (2007). Exploring the need for localization in ensemble data assimilation using a hierarchical ensemble filter. Physica D: Nonlinear Phenomena, 230(1), 99–111.

Anderson JL, Anderson SL (1999) A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. Mon Weather Rev., 127(12), 2741-2758.

Ba J, Erdogdu MA, Ghassemi M, Sun S, Suzuki T, Wu D, Zhang T (2022) Understanding the Variance Collapse of SVGD in High Dimensions. International Conference on Learning Representations.

Barnier G, Biondi E, Clapp RG, Biondi B (2023) Full-waveform inversion by model extension: Theory, design, and optimization. Geophysics, 88(5), R579–R607.

Berti S, Aleardi M, Stucchi E (2023) A computationally efficient Bayesian approach to full waveform inversion. Geophysical Prospecting, 72(2), 580-603.

Berti 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, Vol. 65, no. 2, pp. 291-308.

Berti S, Aleardi M, Stucchi E (2024b) A probabilistic full waveform inversion of surface waves. Geophysical Prospecting, 72, 3448–3473.

Berti S, Ravasi M, Aleardi M, Stucchi E (2025) Bayesian Full Waveform Inversion of surface waves with Annealed Stein Variational Gradient Descent. Geophysical Journal International, 241(1), 641-657.

Bertino L, Evensen G, Wackernagel H (2003) Sequential data assimilation techniques in oceanography. Int. Stat. Rev., 71(2), 223-241.

Blei DM, Kucukelbir A, McAuliffe JD (2017) Variational Inference: A Review for Statisticians. Journal of the American Statistical Association, 112(518), 859-877.

Bodin T, Sambridge M (2009) Seismic tomography with the reversible jump algorithm. Geophysical Journal International, 178(3), 1411-1436.

Britanak V, Yip P, Rao KR (2010) Discrete cosine and sine transforms: general properties, fast algorithms and integer approximations. Amsterdam: Elsevier.

Brougois A, Bourget M, Lailly P, Poulet M, Ricarte P, Versteeg R (1990) Marmousi, model and data. European Association of Geoscientists and Engineers, ISBN 978-90-73781-01-6.

Canchumuni SWA, Emerick AA, Pacheco MAC (2019) Towards a robust parameterization for conditioning facies models using deep variational autoencoders and ensemble smoother. Computers & Geosciences, 128, 87-102.

Carrassi A, Bocquet M, Bertino L, Evensen G (2018) Data assimilation in the geosciences: An overview of methods, issues and perspectives. Wiley Interdisciplinary Reviews: Climate Change, 9(5), e535.

Chada NL, Chen Y, Sanz-Alonso D (2021) Iterative ensemble Kalman methods: A unified perspective with some new variants. Foundations of Data Science, 3(3): 331-369.

Chen C, Gao G, Honorio J, Gelderblom P, Jimenez E, Jaakkola T (2014) Integration of principal-component-analysis and streamline information for the history matching of channelized reservoirs. In: Proceedings of the SPE Annual Technical Conference and Exhibition, Amsterdam, The Netherlands, 27–29 October, Number SPE-170636-MS.

Chen Y, Oliver D (2010) Ensemble-based closed loop optimization applied to Brugge field. SPE Reserv. Eval. Eng., 13(1), 56-71.

Chen Y, Oliver D (2013) Levenberg-Marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. Computer Geoscience, 17, 689-703.

Chen Y, Oliver D (2014) History matching of the Norne full-field model with an iterative ensemble smoother. SPE Reserv. Eval. Eng., 17(2).

Cominelli A, Dovera L, Vimercati S, Naevdal G (2009) Benchmark study of ensemble Kalman filter methodology: history matching and uncertainty quantification for a deep-water oil reservoir. Proceedings of the International Petroleum Technology Conference, Doha, Qatar, 7-9 December, Number IPTC 13748.

Corrales M, Berti S, Denel B, Williamson P, Ravasi M, Aleardi M (2025) Annealed Stein Variational Gradient Descent for improved uncertainty estimation in full-waveform inversion. Geophysical Journal International, 241(2), 1088-1113.

Curtis A, Lomax A (2001) Prior information, sampling distributions and the curse of dimensionality. Geophysics, 66(2), 372-378.

D'Angelo F, Fortuin V (2021) Annealed Stein Variational Gradient Descent. arXiv, eprint:2101.09815.

Davis K, Li Y (2011) Fast solution of geophysical inversion using adaptive mesh, space-filling curves and wavelet compression. Geophysical Journal International, 185(1), 157-166.

Emerick AA, Reynolds AC (2012) History matching time-lapse seismic data using the ensemble Kalman filter with multiple data assimilations. Computational Geosciences, 16, 639-659.

Emerick AA, Reynolds AC (2013) Ensemble smoother with multiple data assimilation. Computers and Geosciences, 55, 3-15.

Evensen G (1994) Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. Journal of Geophysical Research, 99(C5), 10143-10162.

Evensen G (2009a) The ensemble Kalman filter. Springer.

Evensen G (2009b) The ensemble Kalman filter for combined state and parameter estimation. IEEE Contr. Syst. Mag., 83-104.

Evensen G, Vossepoel FC, van Leeuwen PJ (2022) Data assimilation fundamentals. Berlin: Springer.

Fichtner A, Simute S (2018) Hamiltonian Monte Carlo inversion of seismic sources in complex media. Journal of geophysical Research: Solid Earth, 123(4), 2984–2999.

Fichtner A, Zunino A, Gebraad L (2019) Hamiltonian Monte Carlo solution of tomographic inverse problems. Geophysical Journal International, 216(2), 1344–1363.

Gallego V, Insua DR (2020) Stochastic Gradient MCMC with Repulsive Forces. arXiv, doi: 10.48550/arXiv.1812.00071.

Gaspari G, Cohn SE (1999) Construction of correlation functions in two and three dimensions. Quarterly Journal of the Royal Meteorological Society, 125, 723–757.

Gebraad L, Boehm C, Fichtner A (2020) Bayesian elastic full waveform inversion using Hamiltonian Monte Carlo. Journal of Geophysical Research: Solid Earth, 125(3).

Gineste M, Eidsvik J, Zheng Y (2019) Ensemble-based seismic inversion for a stratified medium. Geophysics, 85 (1): R29–R39.

Green PJ (1995) Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination. Biometrika, Vol. 82, No. 4., p. 711-732.

Hoffmann A, Brossier R, Métivier L, Tarayoun A (2024) Local uncertainty quantification for 3-D time-domain full-waveform inversion with ensemble Kalman filters: application to a North Sea OBC data set. Geophysical Journal International, 237(3), 1353-1383.

Houtekamer PL, Mitchell HL (2001) A sequential ensemble Kalman filter for atmospheric data assimilation. Mon. Weather Rev., 129(1), 123-137.

Houtekamer PL, Mitchell HL (2005) Ensemble Kalman filtering. Q.J.R. Meteorol. Soc., 131, 3269-3289.

Iglesias M (2015) Iterative regularization for ensemble data assimilation in reservoir models. Computational Geosciences, 19(1), 177-212.

Izzatullah M, Alali A, Ravasi M, Alkhalifah T (2024) Physics-reliable frugal local uncertainty analysis for full waveform inversion. Geophysical Prospecting, 72(2), 2718-2738.

Jain AK (1989) Fundamentals of digital image processing. Englewood Cliffs: Prentice Hall.

Jafarpour B, McLaughlin DB (2008) History matching with an ensemble Kalman filter and discrete cosine parameterization. Comput. Geosci., 12(2), 227–244.

Jafarpour B, Goyal V, McLaughlin D, Freeman W (2009) Transform-domain sparsity regularization for inverse problems in geosciences. Geophysics, 74(5), R69-R83.

Jafarpour B (2010) Wavelet reconstruction of geologic facies from nonlinear dynamic flow measurements. IEEE Trans. Geosci. Remote Sens., 49 (5), 1520–1535.

Jin L, Sen MK, Stoffa PL (2008) One-dimensional pre-stack seismic waveform inversion Using Ensemble Kalman Filter. In SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, 1920-1924.

Jung H, Jo H, Kim S, Lee K, Choe J (2017) Recursive update of channel information for reliable history matching of channel reservoirs using EnKF with DCT. J. Pet. Sci. Eng., 154, 19–37.

Kaltenbacher B, Neubauer A, Scherzer O (2008) Iterative Regularization Methods for Nonlinear Ill-posed Problems, de Gruyter, Berlin.

Kim S, Chung WK, Lee JH (2021) Acoustic full waveform inversion using Discrete Cosine Transform. Journal of Seismic Exploration, 30: 365-380.

Kingma DP, Salimans T, Jozefowicz R, Chen X, Sutskever I, Welling M (2016) Improved Variational Inference with Inverse Autoregressive Flow. Advances in Neural Information Processing Systems, vol. 29, Curran Associates, Inc.

Kucukelbir A, Tran D, Ranganath R, Gelman A, Blei DM (2017) Automatic differentiation variational inference. Journal of Machine Learning Research, 18(1), 430-474.

Kullback S, Leibler RA (1951) On Information and Sufficiency. The Annals of Mathematical Statistics, 22(1), 79-86.

Lacerda JM, Emerick AA, Pires AP (2019) Methods to mitigate loss of variance due to sampling errors in ensemble data assimilation with non-local model parameters. Journal of Petroleum Science and Engineering, 172, 690–706.

Lailly P (1983) The seismic inverse problem as a sequence of before stack migrations. In: Conference on Inverse Scattering, Theory and application, Society for Industrial and Applied Mathematics, Philadelphia, edited by R. Bednar and Weglein, pp. 206–220.

Laloy E, Hérault R, Jacques D, Linde N (2017) Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network. Adv. Water Resour., 110, 387–405.

Lin Y, Abubakar A, Habashy TM (2012) Seismic full-waveform inversion using truncated wavelet representations. Expanded Abstracts, 82nd Ann. Internat. SEG Mtg., Las Vegas, 1-6.

Liu M, Grana D (2018a) Ensemble-based seismic history matching with data reparameterization using convolutional autoencoder. In SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, 3156-3160.

Liu M, Grana D (2018b) Stochastic nonlinear inversion of seismic data for the estimation of petroelastic properties using the ensemble smoother and data reparameterization. Geophysics, 83(3), M25-M39.

Liu M, Grana D (2018c) Ensemble-based joint inversion of PP and PS seismic data using full Zoeppritz equations. In SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, 511-515.

Liu Q, Wang D (2016) Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. Advances in Neural Information Processing Systems, vol. 29, Curran Associates, Inc.

Lochbühler T, Breen SJ, Detwiler RL, Vrugt JA, Linde N (2014) Probabilistic electrical resistivity tomography of a CO_2 sequestration analog. Journal of Applied Geophysics, 107, 80-92.

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 and Exhibition, volume 2023, pages 1–5.

Malinverno A (2002) Parsimonious Bayesian markov chain monte carlo inversion in a nonlinear geophysical problem. Geophysical Journal International, 675-688.

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, A1460- A1487.

Mazzotti A, Bienati N, Stucchi E, Tognarelli A, Aleardi M, Sajeva A (2016) Two-grid genetic algorithm full-waveform inversion. The Leading Edge, 35, 1068-1075.

Moghadas D, Vrugt JA (2019) The Influence of Geostatistical Prior Modeling on the Solution of DCT-Based Bayesian Inversion: A Case Study from Chicken Creek Catchment. Remote Sensing**,** 11(13), 1549.

Moyen R, Gentilhomme T (2021) Adaptive Ensemble-Based Optimisation for Petrophysical Inversion. Math Geosci **53**, 349–373.

Nawaz MA, Curtis A (2018) Variational Bayesian inversion (VBI) of quasi-localized seismic attributes for the spatial distribution of geological facies. Geophysical Journal International, Volume 214, Issue 2, Pages 845–875.

Neal R (1996) Sampling from multimodal distributions using tempered transitions. Statistics and computing, 6(4), 353-366.

Neal R (2011) MCMC using Hamiltonian Dynamics. Handbook of Markov Chain Monte Carlo, 11662. Chapman, Hall/CRC.

Nuha HH, Liu B, Mohandes M, Deriche M (2017) Seismic data compression using signal alignment and PCA. In: Proc. 9th IEEE-GCC Conf. Exhib. (GCCCE), May 2017, pp. 1–6.

Ray 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.

Reynolds AC, Zafari M, Li G (2006) Iterative forms of the ensemble Kalman filter. Moscow: European Association of Geoscientists and Engineers.

Reynolds AC, Le DH, Emerick AA (2014) Ensemble-Based Methods for Data Integration and Uncertainty Quantification. European Association of Geoscientists and Engineers.

Richardson A (2023) Deepwave, https://ausargeo.com/deepwave/.

Rincón F, Berti S, Aleardi M, Tognarelli A, Stucchi E (2025) Integrating deep learning and discrete cosine transform for surface waves full-waveform inversion. Geophysical Journal International, 240(1), 805–828.

Roe P, Almendral Vazquez A, Hanea R (2016) Distinguishing Signal from Noise in History Matching – Analysis of Ensemble Collapse on a Synthetic Data Set. 15th European Conference on the Mathematics of Oil Recovery, 29 August – 1 September, Amsterdam, Netherlands.

Saad O, Alkhalifah T (2025) Enhancing multiparameter elastic full-waveform inversion with a Siamese network. The Leading Edge, 44(5), 416a1-416a10.

Sambridge M, Bodin T, Gallacher K, Tkalčić H (2013) Transdimensional inference in the geosciences. Philos. Trans. R. Soc. A Math. Phys. Eng. Sci., 371.

Sen MK, Biswas R (2017) Transdimensional seismic inversion using the reversible jump Hamiltonian monte carlo algorithm. Geophysics, R119-R134.

Tarantola A (1984) Inversion of seismic reflection data in the acoustic approximation. Geophysics, 49, 1259–66.

Tarantola A (2005) Inverse problem theory and methods for model parameter estimation. Philadelphia, PA: SIAM: Society for Industrial and Applied Mathematics.

Thurin J, Brossier R, Métivier L (2019) Ensemble-based uncertainty estimation in full waveform inversion. Geophysical Journal International, 219(3), 1613-1635.

Tveit S, Bakr SA, Lien M, Mannseth T (2015) Ensemble-based Bayesian inversion of CSEM data for subsurface structure identification. Geophysical Journal International, 201(3), 1849-1867.

Tveit S, Mannseth T, Park J, Sauvin G, Agersborg R (2020) Combining CSEM or gravity inversion with seismic AVO inversion, with application to monitoring of large-scale CO_2 injection. Computational Geosciences, 24, 1-20.

Urozayev D, Ait-El-Fquih B, Hoteit I, Peter D (2021) A Reduced-order Variational Bayesian Approach for Efficient Subsurface Imaging. Geophysical Journal International. 229.

van Leeuwen PJ, Evensen G (1996) Data assimilations and inverse methods in terms of a probabilistic formulation. Monthly Weather Review, 124, 2898-2913.

Vinciguerra A, Aleardi M, Madsen LM, Bording TS, Christiansen AV, Stucchi E (2023) Stochastic inversion of time-lapse electrical resistivity tomography data by means of an adaptive ensemble-based approach. Geophysical Prospecting, 1-17.

Virieux J, Operto S (2009) An overview of full-waveform inversion in exploration geophysics. Geophysics, 74, no. 6, WCC1-WCC26.

Whitaker JS, Hamill TM, Wei X, Song Y, Toth Z (2008) Ensemble data assimilation with the NCEP global forecast system. Mon. Weather Rev., 136, 463-482.

Zhang X, Curtis A (2020a) Seismic Tomography Using Variational Inference Methods. Journal of Geophysical Research: Solid Earth, 125(4), e2019JB018589.

Zhang X, Curtis A (2020b) Variational full-waveform inversion. Geophysical Journal International, Vol. 222, Issue 1, 406-411.

Zhang X, Curtis A (2021) Bayesian full-waveform inversion with realistic priors. Geophysics, 86(5), A45–A49.

Zhang C, Chen T (2022) Bayesian slip inversion with automatic differentiation variational inference. Geophysical Journal International, Volume 229, Issue 1, Pages 546–565.

Zhang X, Lomas A, Zhou M, Zheng Y, Curtis A (2023) 3-D Bayesian variational full waveform inversion. Geophysical Journal International, Volume 234, Issue 1, Pages 546-561.

Zhao Z, Sen MK (2021) A gradient-based Markov Chain Monte Carlo method for full- waveform inversion and uncertainty analysis. Geophysics, 86, no.1, P. R15-R30.

Zhou H, Li Q (2013) Increasing waveform inversion efficiency of GPR data using compression during reconstruction. J. Appl. Geophys., 99, 109-113.

Zunino A, Ghirotto A, Armadillo E, Fichtner A (2022) Hamiltonian Monte Carlo probabilistic joint inversion of 2D (2.75D) gravity and magnetic data. Geophysical Research Letters, 49(20).