

On the design of ultra-sparse seismic surveys for monitoring subsurface CO₂ storage sites using full waveform inversion

Xuebin Zhao^{*,} Andrew Curtis

School of Geosciences, University of Edinburgh, Edinburgh, United Kingdom

ARTICLE INFO

Keywords:

Subsurface CCS monitoring
Time-lapse imaging
Seismic survey design
Experimental design
Uncertainty estimation
Nonlinear information

ABSTRACT

Monitoring the dynamics of CO₂ in subsurface reservoirs allows the conformance of carbon capture and storage (CCS) projects to be assessed. Full waveform inversion (FWI) of data from dense, time-lapse seismic surveys can provide high resolution images of dynamic changes. However, FWI solutions remain highly non-unique, so uncertainties must be accounted for to ensure that conformance verification is robust. Time-lapse seismic FWI is therefore expensive because, first, dense surveys are costly to acquire, and second, quantifying realistic uncertainties requires extreme computational power and memory. We first introduce a significantly less costly method to quantify Bayesian uncertainties in the maximum a posteriori (MAP – most likely) solutions of time-lapse seismic velocity changes. The method embodies strong prior information from the baseline survey to inform inversions of monitoring surveys. In contrast to comparable methods, these uncertainty estimates are shown to be of a reasonable magnitude to inform subsequent decision-making. This method also allows the quality of *prospective* survey designs to be assessed in terms of expected confidence in time-lapse imaging results, at reasonable computational cost. We therefore perform a time-lapse seismic survey design study to assess the quality of more economically attractive surveys. We demonstrate for the first time that even if extremely sparse acquisition geometries are deployed, potentially even involving only a single seismic source and recordings on a single fibre-optic cable, reasonable images of subsurface time-lapse velocity changes are produced, and uncertainties remain sufficiently low to enable robust decision-making.

1. Introduction

Carbon capture and storage (CCS) within the Earth's subsurface has emerged as an essential strategy for mitigating atmospheric CO₂ buildup from fossil fuel combustion if global warming is to be limited to 1.5 or even 2 degrees above the pre-industrial average temperature (Lee et al., 2023). A critical aspect of any CCS project is to verify conformance of the reservoir performance with pre-injection forecasts. This requires that subsurface CO₂ migration is monitored, and any significant leakage is detected. Several approaches have been developed to track CO₂ migration, most involving numerical simulations of reservoir properties and dynamics (Shahkarami et al., 2014; Li et al., 2016), geochemical monitoring (Emberley et al., 2005; Harbert et al., 2020), and geophysical monitoring (Lumley, 2001; Fawad and Mondol, 2021). Among these, geophysical methods are considered to be the most effective for imaging the spatial heterogeneity of dynamic changes in subsurface properties.

While some geophysical methods exhibit a strong sensitivity to changes in the subsurface distribution of CO₂ (e.g., electrical and electromagnetic methods – Zhdanov et al., 2013; Gasperikova and Li,

2021), this does not translate into accurate images or maps of the changing distribution of CO₂ due to the well-known non-uniqueness in solutions found from such potential methods (Wei and Sun, 2021). Despite their relatively lower sensitivity to changes in CO₂ distribution, seismic monitoring methods produce high resolution, spatially resolved estimates of changes in properties of a reservoir and surrounding rock mass following CO₂ injection (Nakata et al., 2022, 2024), by solving inverse or inference problems based on surface or borehole seismic observations (Tarantola, 2005). The principal disadvantage of seismic methods is their cost: seismic data are usually expensive to acquire and process compared to some other geophysical methods, which may exceed budgets available for monitoring of CO₂ storage sites. This is particularly true when realistic estimates of uncertainties in the results are needed, in order to ensure that environmental and financial risks associated with operation or leakage of CO₂ from subsurface stores, can be mitigated. The computational cost of estimating realistic uncertainties is usually high, which may exceed budgets of available computational power (Gebraad et al., 2020; Levy et al., 2022; Siahkoobi et al., 2023; Zhang et al., 2023; Lomas et al., 2023; Zhao and

* Corresponding author.

E-mail address: xuebin.zhao@ed.ac.uk (X. Zhao).

Curtis, 2024b). In this paper we therefore introduce and test a method to estimate useful uncertainty estimates using resource-constrained seismic acquisition and processing methods, to be more consistent with budgets that may be available for future CO₂ storage monitoring.

Full waveform inversion (FWI) is a seismic imaging method that constructs high resolution seismic velocity models of the Earth's subsurface from seismic waveform data (Tarantola, 1984; Fichtner et al., 2009; Virieux and Operto, 2009). The method has been used for so-called time-lapse monitoring of subsurface CO₂ storage (Li et al., 2021; Sinha et al., 2024; Nakata et al., 2024), wherein seismic surveys are conducted periodically to collect data that reflect changes in the subsurface. Images of changes in seismic velocity are estimated using FWI of the data from multiple surveys.

A variety of time-lapse FWI methods have been proposed. The most straightforward approach is the parallel difference strategy, in which two (baseline and monitoring) datasets are inverted independently. Time-lapse changes are then determined by subtracting the inverted monitoring model from the inverted baseline model (Plessix et al., 2010). However, the convergence of FWI and final model for each dataset are affected by survey repeatability, data quality and computational parameters, which may vary between surveys. Consequently, this method often results in strong inversion artefacts (Asnaashari et al., 2015). The second strategy, known as the sequential difference method, uses the inverted baseline model both as the initial model and as additional prior information in the inversion of the monitoring data (Asnaashari et al., 2012). The assumption is that time-lapse changes in the subsurface are expected to be localised and to have small amplitude, making the baseline inversion solution an ideal starting point and prior constraint for the monitoring inversion. A third alternative is the double difference method, where only the differences between the two datasets are inverted to derive a differential subsurface model (Watanabe et al., 2004; Zheng et al., 2011; Zhang and Huang, 2013; Sinha et al., 2024). The main advantage of this strategy is that any data that are unexplained by the baseline inversion results may be disregarded in the monitoring inversion (see below), significantly reducing time-lapse imaging artefacts (Asnaashari et al., 2015; Yang et al., 2016). A fourth technique is the joint inversion of both baseline and monitoring datasets to obtain velocity models at two epochs simultaneously within a single inversion process. This approach helps to mitigate imaging artefacts caused by illumination differences and convergence issues, but can be extremely computationally expensive (Maharramov and Biondi, 2014; Maharramov et al., 2016).

Recently, Sinha et al. (2024) proposed to use a combination of time-lapse FWI and passive seismic imaging methods (e.g., dispersion inversion using seismic ambient noise data) to reduce monitoring cost in CCS projects. In this approach, passive methods are used to monitor subsurface changes routinely, serving as an early warning trigger system. If unexpected changes are detected, more detailed monitoring will be deployed using time-lapse FWI associated with active seismic surveys. However, it has been shown that ambient noise is likely to be insufficient to monitor some CO₂ storage reservoirs (Stork et al., 2018). In addition, two main issues still exist with the detailed monitoring methods of Sinha et al. (2024) as well as all of the aforementioned studies. First, the cost of conducting active seismic surveys are still high since dense acquisition geometries are normally used. Second, time-lapse inversion is executed within an optimisation framework in which a single optimal set of model parameters (subsurface properties) is estimated by minimising a misfit function defined between synthetic and observed waveform data. Given that seismic inverse problems are inherently under-determined, large uncertainties often exist in the solutions, implying that an infinite number of different models fit the observed data to within measurement uncertainties (Mosegaard and Sambridge, 2002; Tarantola, 2005; Gebraad et al., 2020; Zhang and Curtis, 2020; Khoshkholgh et al., 2022). The reliability and hence value of any single model estimate therefore always remains unclear. Our work is motivated by the need to address these two problems.

Specifically, in order to make fully risk-based decisions concerning containment of future CCS sites based on seismic data, it will be necessary to estimate uncertainty in the full, nonlinear FWI problem of imaging subsurface CO₂ migration, and to both acquire and process the seismic data at significantly reduced cost compared to existing methods.

Bayesian inference is commonly employed to account for uncertainties in an inverse problem, by estimating the so-called *posterior* probability density function (pdf), which describes all plausible model parameter values that are consistent with the observed data. Monte Carlo sampling methods have been used extensively in probabilistic geophysical inverse problems, in which an ensemble of samples (model realisations) is drawn from the posterior pdf to represent the solution's uncertainty (Mosegaard and Tarantola, 1995; Sambridge et al., 2006; Bodin and Sambridge, 2009; Sen and Stoffa, 2013; Fichtner and Simut , 2018; Khoshkholgh et al., 2021; Zunino et al., 2023). However, these sampling-based methods often suffer from high costs due to computational inefficiencies in high dimensional inverse problems, primarily due to the curse of dimensionality (Curtis and Lomax, 2001), and to difficulties in detecting statistical convergence of Monte Carlo algorithms (Gelman and Rubin, 1992).

Variational inference offers an alternative approach which uses optimisation methods to solve Bayesian inverse problems (Blei et al., 2017; Zhang et al., 2018). For certain geophysical inverse problems, variational methods have shown greater efficiency and scalability to higher dimensional problems (larger models), and they also allow convergence to be detected more easily than for Monte Carlo sampling methods (Nawaz and Curtis, 2018, 2019; Zhang and Curtis, 2020; Zhao et al., 2021; Siahkoobi et al., 2021, 2023; Levy et al., 2022; Zidan et al., 2022; Smith et al., 2022; Zhang et al., 2023; Sun et al., 2023; Lomas et al., 2023; Izzatullah et al., 2024; Sun et al., 2024).

Recent studies in time-lapse FWI have begun to emphasise uncertainty quantification. Kotsi et al. (2020a,b) conducted full waveform time-lapse imaging and uncertainty estimation using a local forward solver to reduce the computational cost associated with the forward simulations. Fu and Innanen (2022) used a Markov chain Monte Carlo (MCMC) method combined with a target-oriented inversion scheme, which reduced the dimensionality of the inverse problem by updating only a selected subset of model parameters during the inversion process. Zhang and Curtis (2024) implemented variational inference in Bayesian time-lapse FWI, comparing probabilistic versions of sequential difference and joint inversion strategies against a deterministic double difference method. Their results indicate that joint inversion provides the most precise estimates of velocity changes and their uncertainties. de Lima et al. (2024) investigated the sequential approach in a Bayesian framework using the Hamiltonian Monte Carlo (HMC) method, integrating probabilistic baseline inversion results as prior knowledge to enhance the monitoring estimates.

In this work, we develop and test a novel method for time-lapse FWI and uncertainty quantification aimed at monitoring subsurface CO₂ storage and detecting potential migration following injection, in a computationally tractable manner. We first deploy a fully nonlinear, probabilistic and computationally efficient variational Bayesian inversion method, specifically the physically structured variational inference method (PSVI – Zhao and Curtis, 2024c), to invert baseline waveform data for the subsurface seismic velocity structure. Strong prior information derived from those inversion results is then injected into the monitoring inversion as follows.

Given the localised nature of time-lapse changes within some vicinity of the injection point, the considerably smaller amplitude of these changes (typically one or two orders of magnitude lower than the magnitude of parameter values), and the fact that repeat surveys are often acquired such that they approximately span a subset of the acquisition geometry used in the baseline survey, it is expected that parameter values that provide good fits to the baseline survey should be close (in parameter space) to values that provide good fits to the monitor survey

data. This suggests that point-estimates of properties such as the pdf maxima for the monitor survey, should be different to those from the baseline but comparable. Therefore, to conduct the monitoring survey inversion, we generate a set of posterior samples from the baseline inversion results. Each sample is adjusted by the baseline data to obtain a local estimate of the *maximum a posteriori* (MAP) solution. The set of MAP solutions spans a manifold within parameter space, and the resulting set of models characterises uncertainty in the Bayesian MAP manifold for the baseline survey. This is illustrated schematically in Fig. 1a. We show that this uncertainty is significant, which implies that the standard method of using a single model as a reference for subsequent monitoring inversions is likely to lead to substantially biased results.

The set of baseline MAP samples is then used as strong prior information for the monitoring survey. This information is ‘strong’ in the sense that the MAP manifold is more focused on better-fitting baseline models than is the full baseline Bayesian posterior distribution (e.g., as used by de Lima et al. (2024)). Since velocity changes due to CO₂ injection are generally expected to be small (Kim et al., 2013), in the monitoring inversion we use each of the manifold samples as an initial model for a linearised (deterministic) optimisation which finds local models that fit data from the monitoring survey. Each of the models found in this way is a MAP sample of the Bayesian posterior distribution in the monitoring survey. Measures of uncertainty in the MAP time-lapse changes can therefore be estimated, by comparing each of the MAP baseline samples to their corresponding MAP monitoring samples. An illustrative example is provided in Fig. 1b.

Our method therefore estimates uncertainties in time-lapse changes, which are not equivalent to Bayesian uncertainties found in Zhang and Curtis (2024) and de Lima et al. (2024), since we estimate uncertainties in the time-lapse changes of the MAP manifold of models that provide a good fit to the data, rather than in the full Bayesian posterior pdf. We justify this approximation herein, first by showing that the computational cost reduction offered by this method is significant, and that the results are accurate and useful even when using sparse and cost-effective acquisition geometries. Second, introduction of the method is timely since the calculation of a full Bayesian solution of the baseline survey, and hence the ability to find representative samples of the full baseline MAP manifold (which differentiates this method from previous uses of MAP solutions), have only recently become computationally feasible, using methods discussed above. And finally, we show below that through careful reasoning, this method allows information about subsurface CO₂ distributions to be estimated even when monitoring using sparse data acquisition relative to the baseline survey.

Posterior uncertainties are implicitly conditioned on survey acquisition geometries since different geometries lead to observed datasets that contain different information about parameters of interest (Maurer et al., 2010; Strutz and Curtis, 2024). So after verifying that this method of uncertainty assessment is effective, we apply it to address the following question: can ultra-sparse acquisition geometries, and the efficient computational method above, be used to detect time-lapse changes sufficiently accurately to quantify practically useful uncertainties? This is of particular relevance to CCS monitoring, given that remuneration or financial reward expected for monitoring of CO₂ is expected to be limited. To answer this question, we apply the method to different acquisition geometries, varying the number of sources and receivers in the monitoring survey. We show that indeed, a very sparse acquisition geometry (even using only a single seismic source) together with the new processing method, may produce reliable, informative and useful results.

The rest of this paper is organised as follows. We first introduce the proposed time-lapse FWI methodology, followed by a numerical example demonstrating its effectiveness. We then apply the proposed method to time-lapse monitoring with sparse acquisition geometries. Finally, we present a brief discussion and draw conclusions.

2. Methodology

2.1. Variational Bayesian inversion of baseline data

The inversion of baseline data is performed under a probabilistic framework by calculating the *posterior* probability density function (pdf) using Bayes’ rule

$$p(\mathbf{m}_1|\mathbf{d}_1) = \frac{p(\mathbf{d}_1|\mathbf{m}_1)p(\mathbf{m}_1)}{p(\mathbf{d}_1)} \quad (1)$$

where $p(\cdot)$ stands for a probability distribution. Term $p(\mathbf{m}_1)$ represents the *prior* information about the unknown baseline model parameter vector \mathbf{m}_1 , and $p(\mathbf{d}_1|\mathbf{m}_1)$ is called the *likelihood*, defined to be the probability of observing the baseline data \mathbf{d}_1 given any value of \mathbf{m}_1 . Throughout this paper we use subscripts 1 and 2 to denote variables associated with the baseline and monitoring inversions, respectively. Term $p(\mathbf{d}_1)$ is a normalisation constant (it is independent of \mathbf{m}_1 and thus does not vary with the Earth model) called the *evidence*. Typically, a Gaussian distribution is used to represent data uncertainties in the likelihood function

$$p(\mathbf{d}_1|\mathbf{m}_1) \propto \exp \left[-\frac{1}{2}(\mathbf{d}_1 - \mathbf{d}(\mathbf{m}_1))^T \Sigma_{\mathbf{d}_1}^{-1}(\mathbf{d}_1 - \mathbf{d}(\mathbf{m}_1)) \right] \quad (2)$$

where $\Sigma_{\mathbf{d}_1}^{-1}$ is the covariance matrix of the data uncertainties, and $\mathbf{d}(\mathbf{m}_1)$ is the synthetic data predicted by a model vector \mathbf{m}_1 . In practice, $\Sigma_{\mathbf{d}_1}^{-1}$ is often assumed to be a diagonal matrix since there is currently no known way to determine reliable off-diagonal waveform covariances from standard seismic data sets. The resulting posterior distribution $p(\mathbf{m}_1|\mathbf{d}_1)$ describes the probability of any possible model \mathbf{m}_1 , calculated from its consistency with the combination of baseline data and prior information.

Estimating or characterising the probability distribution on the left of Eq. (1) is called Bayesian inference, or Bayesian inversion (since the conditional relationships in distributions on the left and right sides of the equation are inverted). In *variational* Bayesian inference, a family of probability distributions (known as the variational family) is defined, from which we select an optimal member that best approximates the unknown posterior pdf. This can be accomplished by minimising the Kullback–Leibler (KL) divergence (a measure of difference) between the variational and posterior distributions (Kullback and Leibler, 1951): this is an optimisation problem which provides a fully probabilistic result — an approximate posterior probability distribution.

We employ physically structured variational inference (PSVI – Zhao and Curtis, 2024c) for the baseline inversion. PSVI has previously been applied to Bayesian 2D FWI (Zhao and Curtis, 2024d) and 3D FWI (Zhao and Curtis, 2024b), demonstrating its efficiency and ability to deliver inversion results with reasonable uncertainty estimates. Our implementation uses a transformed Gaussian distribution with a specific covariance structure to approximate the true posterior pdf, characterised by a mean vector $\boldsymbol{\mu}$ and a covariance matrix $\boldsymbol{\Sigma}$. We emphasise that this is not the same as the Gaussian approximations that are often calculated using linearised inversion algorithms, since those produce only local approximations to the posterior density structure around a particular *maximum a posteriori* (MAP) solution. Instead, PSVI fits a transformed Gaussian to the full posterior probability density structure. The covariance matrix $\boldsymbol{\Sigma}$ obtained from PSVI therefore also provides different information from the inverse Hessian matrix estimated in linearised inversion. The former describes an approximation to the Bayesian posterior uncertainty of the solution, whereas the latter is simply a matrix of (second-order) derivatives evaluated at a single point in parameter space.

In our current implementation of PSVI, we include only specific (important) correlations in the posterior estimate of uncertainty in vector \mathbf{m}_1 . These consist of correlations between pairs of parameters that are spatially proximal, roughly within one dominant wavelength in FWI. This is to capture trade-offs in velocity estimates that occur locally, because far-field seismic wavefields only respond to wavelength-scale




Fig. 1. Schematic example of probabilities and samples of parameter values x and y , used to illustrate the proposed time-lapse FWI method. (a) Posterior pdf from the baseline survey data, which has three main modes. Darker colour represents higher probability; the baseline MAP manifold is shown in black. Orange and green dots denote a set of baseline posterior samples \mathbf{m}_1 and MAP samples \mathbf{m}_1^* , respectively. Pink arrows show the model updates found by solving Eq. (3), i.e., $\mathbf{m}_1^* - \mathbf{m}_1$. (b) Posterior pdf using data from an imagined monitoring survey performed using a fully repeated baseline survey design. The pdf is highly consistent with that in (a) because time-lapse changes are expected to be small relative to heterogeneity in the velocity structure itself. During the monitoring inversion, we would update baseline MAP samples \mathbf{m}_1^* (green dots from (a)) into the corresponding monitoring MAP samples \mathbf{m}_2^* (red dots) by minimising Eq. (4). (c) Posterior monitoring pdf using data collected from a sparse survey. Sparse monitoring data are relatively uninformative compared to the baseline data. The posterior pdf is thus relatively unconstrained — and the MAP manifold is broader than those in (a) and (b). Red stars denote the updated monitoring MAP samples of this pdf. Yellow arrows in (b) and (c) represent estimated time-lapse changes $\delta\mathbf{m} = \mathbf{m}_2^* - \mathbf{m}_1^*$.

averages of properties at sub-wavelength scales (Capdeville et al., 2003). All other inter-parameter correlations in \mathbf{m}_1 are ignored. This is accomplished by setting Σ to be sparse, with tunable parameters in specific off-diagonal elements, and fixing all other elements to be zero. The mean μ , and all non-zero parameters in Σ , are determined by optimisation: theoretical and implementational details of PSVI, together with results that verify its effectiveness, are provided in Zhao and Curtis (2024c).

In a Bayesian 3D FWI problem, Zhao and Curtis (2024b) estimated posterior distributions for a variety of different sets of prior information, using the method of prior replacement (Walker and Curtis, 2014; Zhao and Curtis, 2024d). This method is highly efficient but the results are approximate, so Zhao and Curtis (2024b) showed that they can be refined by invoking Bayes' rule again but with a limited additional computational budget. The updated results provide more accurate posterior statistics. We introduce a similar refinement step to the baseline inversion results in the current paper to improve estimates of time-lapse changes. We draw a set of samples from the baseline posterior pdf, and adjust each one locally, by maximising the logarithmic posterior probability value using a linearised inversion method with a low number of iterations:

$$\mathbf{m}_1^* = \underset{\mathbf{m}_1}{\operatorname{argmax}} \log p(\mathbf{m}_1 | \mathbf{d}_1) = \underset{\mathbf{m}_1}{\operatorname{argmax}} [\log p(\mathbf{d}_1 | \mathbf{m}_1) + \log p(\mathbf{m}_1)] \quad (3)$$

The updated sample \mathbf{m}_1^* is a local estimate of the *maximum a posteriori* (MAP) solution around the original sample \mathbf{m}_1 (we use superscript * to denote MAP estimates), and is therefore a sample of the manifold of baseline MAP solutions. We illustrate this idea with a 2-dimensional schematic example presented in Fig. 1a, in which model vector \mathbf{m} contains two unknown parameters x and y . The grey contour shows the baseline posterior pdf, with darker colour representing higher probability. Orange dots represent a set of posterior samples \mathbf{m}_1 , which are updated by solving Eq. (3), resulting in a set of baseline MAP samples \mathbf{m}_1^* , denoted by green dots. Pink arrows connecting each pair of orange and green dots quantify model updates in this step, i.e., $\mathbf{m}_1^* - \mathbf{m}_1$. As we will show below, this update step ensures that time-lapse estimates are calculated using two MAP model solutions obtained from baseline and monitoring inversions, respectively, which are directly comparable.

2.2. Deterministic update using monitoring data

In a subsurface CCS project, time-lapse variations are typically localised to some extent around and above the CO₂ injection point which is known *a priori*, and are often expected to be small (usually one or two orders of magnitude smaller than the actual model parameter values (Kim et al., 2013)). This implies that the true baseline model should closely resemble the monitoring model (Asnaashari et al., 2015). The baseline inversion results obtained in the previous section

should therefore provide both (1) reasonable initial or reference values of model parameters, and (2) additional prior information for the monitoring inversion.

Ideally, the inversion of the monitoring data would be performed in a fully probabilistic manner using Bayesian inversion methods, as detailed in Zhang and Curtis (2024) and de Lima et al. (2024). They used the sequential difference strategy in a probabilistic manner (i.e., the posterior pdf from baseline inversion is used as an initial value for monitoring inversion), and obtained posterior pdf's of both baseline and monitoring inversions. Posterior statistics of time-lapse changes, such as the mean and standard deviation models, are calculated from the two posterior pdfs. However, using a method that inverts for the posterior pdf of baseline and monitoring models in separate steps creates a problem: it is then not straightforward to draw samples of, nor to derive a parametric probability distribution to describe, velocity changes. These are crucial for interpretation and decision-making (Arnold and Curtis, 2018; Zhao et al., 2022; Zhang and Curtis, 2022), such as confirming or rejecting putative CO₂ migration or leakage. Joint Bayesian inversion (Zhang and Curtis, 2024) offers an alternative methodology to compute a joint probability distribution and samples of both the baseline and time-lapse models simultaneously. However, this approach potentially doubles the parameter and data space dimensionalities, and thus exacerbates the effects of the curse of dimensionality. In addition, considering the strong prior information derived both from the baseline inversion and from the small amplitude and localised nature of time-lapse changes in a CCS project, one might question whether a fully nonlinear Bayesian inversion of the monitoring data is necessary to quantify uncertainties in time-lapse variations.

We now introduce a novel method for conducting the monitoring inversion and evaluating uncertainties in the results. In the previous section, we obtained a set of MAP models \mathbf{m}_1^* from posterior samples of the baseline inversion results $p(\mathbf{m}_1 | \mathbf{d}_1)$ using Eq. (3). We assume that these samples are sufficient to characterise the most significant components of posterior uncertainty in MAP models of the baseline inversion. Considering that time-lapse changes expected during a CCS project are typically small, we update each baseline MAP model sample \mathbf{m}_1^* within the parameter space to become consistent with the monitoring data using linearised inversions. An issue occurs if a sparse monitoring survey is used, so the resulting monitoring data would be relatively uninformative compared to the baseline data. The posterior pdf from the posterior data alone will then be relatively unconstrained, as illustrated in Fig. 1c; clearly, estimating time-lapse changes by simply comparing the posterior distributions of baseline and monitoring surveys in panels (a) and (c) would result in extremely high uncertainties (Zhang and Curtis, 2024; de Lima et al., 2024). However, assuming that changes to parameter values describing the Earth model between the two epochs are small, the baseline MAP samples must be close to the MAP samples of the posterior pdf that *would* have been obtained in

the imaginary situation wherein the full baseline survey geometry had been repeated, represented in Fig. 1b. While the posterior pdf provided by data from a more economical, reduced-density survey design will be less constrained and hence may contain a much more extensive MAP manifold, it must therefore at least contain MAP values that are approximately consistent with the MAP values of that imagined survey, since both surveys interrogate the same true Earth structure. We use this reasoning to assume that if we update each baseline MAP sample (green dot in Fig. 1c) into the corresponding closest MAP sample of the monitoring survey (red dot), these will be approximately equal to the posterior MAP samples that we would have obtained given data from the imagined full survey (Fig. 1b). Thus, we use the baseline MAP samples, and the assumption that changes are small relative to background velocity heterogeneity, as strong prior information about the location of the MAP solutions of the imaginary survey's posterior distribution; this significantly reduces uncertainty in the time-lapse solution from that which would have been estimated from a sparse monitor survey by directly using the posterior distribution depicted schematically in Fig. 1c. Yellow arrows connecting those green and red dots illustrate model updates during this monitoring inversion, and are therefore interpreted as samples of possible time-lapse changes.

This update can be achieved using a deterministic (iterative linearised) inversion scheme. In this approach, each MAP sample \mathbf{m}_1^* is used as a starting point for the monitoring inversion. The goal is to find a locally optimal MAP model \mathbf{m}_2^* , which minimises the following regularised misfit function iteratively,

$$\chi(\mathbf{m}_2^*) = \frac{1}{2} \left[(\mathbf{d}_2 - \mathbf{d}(\mathbf{m}_2^*))^T \Sigma_{\mathbf{d}_2}^{-1} (\mathbf{d}_2 - \mathbf{d}(\mathbf{m}_2^*)) \right] + \frac{1}{2} \left[(\mathbf{m}_2^* - \mathbf{m}_1^*)^T \Sigma_{\mathbf{m}_2^*}^{-1} (\mathbf{m}_2^* - \mathbf{m}_1^*) \right] \quad (4)$$

where \mathbf{d}_2 and $\mathbf{d}(\mathbf{m}_2^*)$ are the observed monitoring data and the synthetic data predicted by model parameter \mathbf{m}_2^* , respectively. $\Sigma_{\mathbf{d}_2}^{-1}$ is the estimated covariance matrix of the monitoring data uncertainties (again usually assumed to be diagonal), and $\Sigma_{\mathbf{m}_2^*}^{-1}$ is a covariance matrix that encapsulates *a priori* uncertainty information about the difference between any pair of MAP model samples \mathbf{m}_2^* and \mathbf{m}_1^* . Assuming that $\Sigma_{\mathbf{m}_2^*}^{-1}$ is a diagonal matrix, each diagonal element stands for the prior variance of the corresponding model parameter difference between baseline and monitoring surveys. The first term on the right hand side of Eq. (4) calculates a data misfit value (weighted by $\Sigma_{\mathbf{d}_2}^{-1}$) which is equivalent to the logarithm of the likelihood function in Eq. (2), and is interpreted as a function of \mathbf{m}_2^* . The second term integrates prior information derived from the baseline MAP model sample \mathbf{m}_1^* , and can be interpreted as the logarithm of a Gaussian prior centred around \mathbf{m}_1^* with a covariance matrix $\Sigma_{\mathbf{m}_2^*}^{-1}$. Therefore, minimising Eq. (4) essentially updates the previously obtained MAP model \mathbf{m}_1^* which fits the baseline data, to create a new MAP model \mathbf{m}_2^* which fits the monitoring data (yellow arrows in Fig. 1b). This explains why we perform the refinement step for the baseline model sample \mathbf{m}_1^* in Eq. (3): by doing so we obtain paired MAP samples of baseline and monitoring inversions, such that the resulting time-lapse estimates are consistent and can be compared directly. Since we perform this inversion many times from different starting points (many different baseline samples \mathbf{m}_1^*), the spread of final models represents an estimate of uncertainty in the MAP model \mathbf{m}_2^* (multiple yellow arrows in Fig. 1b).

Given precise knowledge about the injection location and injection rate over time in a CCS project, and the assumption that time-lapse variations are localised around the injection point, we use a target-oriented strategy for the monitoring inversion (minimising Eq. (4)). In this approach we update only a predetermined target area around the injection location (Asnaashari et al., 2015; Fu and Innanen, 2022), the extent of which can potentially be chosen using fluid flow simulation (Kolditz et al., 2012; Guo et al., 2016) or by integrating external information from different imaging and inversion results (Raknes and

Arntsen, 2014). Conceptually, this target-oriented strategy is equivalent to setting the diagonal elements of $\Sigma_{\mathbf{m}_2^*}^{-1}$ that represent parameters outside of the target region to zero, effectively assuming that there is no prior uncertainty for those parameters. However, this is not actually implemented because then $\Sigma_{\mathbf{m}_2^*}^{-1}$ would no longer be positive semi-definite. Instead, we simply update parameters within the designated target area. Within this defined target area, we set the prior parameter variance to be relatively large around the injection location to allow for significant model updates, and the variance values decay for spatial locations further from the injection point to suppress potential imaging noise.

We use the double difference method for the monitoring inversion (Watanabe et al., 2004; Zheng et al., 2011; Zhang and Huang, 2013; Sinha et al., 2024). Instead of minimising the misfit between the observed and simulated data, this method minimises the misfit $\Delta \mathbf{d}$ of the difference between the two datasets \mathbf{d}_2 and \mathbf{d}_1 , yielding the following expression:

$$\Delta \mathbf{d} = (\mathbf{d}_2 - \mathbf{d}_1) - (\mathbf{d}(\mathbf{m}_2^*) - \mathbf{d}(\mathbf{m}_1^*)) = \mathbf{d}_{\text{composite}} - \mathbf{d}(\mathbf{m}_2^*) \quad (5)$$

where $\mathbf{d}_{\text{composite}} = \mathbf{d}_2 - \mathbf{d}_1 + \mathbf{d}(\mathbf{m}_1^*)$ is called the composite data, and can be viewed as a new dataset adjusted by the residual data of the baseline MAP model \mathbf{m}_1^* . This adjustment ensures that data which are unexplained in the baseline survey (those not fit by baseline model \mathbf{m}_1^* : $\mathbf{d}_1 - \mathbf{d}(\mathbf{m}_1^*)$) are disregarded in inversions of monitoring survey data. To implement the double difference method, we replace \mathbf{d}_2 in Eq. (4) by $\mathbf{d}_{\text{composite}}$, which allows the double difference residual in Eq. (5) to be minimised using the same regularised FWI scheme introduced above (Eq. (4)).

The proposed time-lapse FWI algorithm can therefore be implemented as follows:

1. Perform fully nonlinear Bayesian inversion of the baseline survey data and estimate the posterior pdf $p(\mathbf{m}_1|\mathbf{d}_1)$. In this work, we apply PSVI to accomplish this step.
2. Generate n random samples from $p(\mathbf{m}_1|\mathbf{d}_1)$, each of which is refined using Eq. (3) to obtain the corresponding MAP sample \mathbf{m}_1^* .
3. For each \mathbf{m}_1^* , perform regularised double difference FWI to obtain the updated MAP model \mathbf{m}_2^* . This is accomplished by minimising Eq. (4) and replacing \mathbf{d}_2 by $\mathbf{d}_{\text{composite}}$.
4. Retrieve samples of time-lapse changes by $\delta \mathbf{m} = \mathbf{m}_2^* - \mathbf{m}_1^*$ and estimate their uncertainties by considering $\delta \mathbf{m}$ obtained from all pairs of \mathbf{m}_1^* and its update \mathbf{m}_2^* .

Note first, that it is important that the samples from step 1 are distributed approximately according to the baseline posterior distribution so that the solutions in step 2 span the MAP manifold; and second, that the linearised inversions of each of the n model samples in steps (ii) and (iii) are independent, and can therefore be fully parallelised.

3. Examples

3.1. Problem setup

In this section, we present a synthetic experiment which we use to validate the proposed method for time-lapse monitoring and uncertainty assessment. As depicted in Fig. 2a, the target baseline velocity structure is a part of the Marmousi model (Martin et al., 2006), consisting of 110×250 grid cells in the vertical and horizontal directions, respectively, with a cell size of 20 m in each direction. A blue plus indicates the CO₂ injection location. For the baseline survey, we place 12 active sources (red stars in Fig. 2a) on the surface with a spacing of 400 m, and 250 receivers (along the white line in Fig. 2a) on the horizontal seabed (200 m depth) with a spacing of 20 m.




Fig. 2. (a) True baseline P-wave velocity model. Red stars and a white line denote source and receiver locations in the baseline survey, respectively. (b) P-wave velocity model at the monitoring stage. Values of velocity change in this panel are amplified by a factor of 20 to highlight the three areas with velocity variations. (c) True P-wave velocity changes due to the imagined injection and migration of CO₂. In each figure, a red box highlights the region in which the velocity model is updated in the monitoring inversion, and a blue plus denotes the imagined injection location of CO₂. Three numbers in (b) and (c) mark three areas of time-lapse changes considered in this example.

Fig. 2b illustrates the true velocity structure at the monitoring stage following imagined CO₂ injection and migration, and Fig. 2c displays the actual time-lapse changes resulting from the injection and potential migration of CO₂. These changes are identified in three localised areas marked with annotations 1, 2 and 3 on Figs. 2b and 2c. Area 1 corresponds to a region directly above the injection point, where the rising and laterally migrating CO₂ is imagined to have decreased the P-wave velocities. For clarity and to differentiate the time-lapse change from geological features presented in the baseline model and imaging artefacts that we show below, this area is deliberately shaped into a non-geological triangular form (this approach was shown to be effective in Zhang and Curtis (2024)). The other two areas (annotations 2 and 3) are imagined to have been caused by the migration of CO₂: in area 2 CO₂ is migrating along a fault, while area 3 shows CO₂ that migrated upwards into a low velocity zone representing a secondary reservoir. In this example, we assume no significant non-CO₂-related subsurface changes between the two surveys. A red box in each figure indicates a target region (the volume in which the velocity model is updated during the monitoring inversion) which encompasses all three areas. Note that the amplitude of the time-lapse changes is amplified by a factor of 20 in Fig. 2b to make them visible; their true magnitudes are approximately -80 m/s as shown in Fig. 2c. Areas 2 and 3 are above the depths where the phase transition to super-critical fluid is likely to occur for CO₂ (typically around 800 m depth), so the CO₂ in these areas would in reality be in gas phase producing larger velocity anomalies. We nevertheless assume the same order of magnitude of velocity anomalies in all three areas to represent a situation in which CO₂ is only partially and diffusely saturated in areas 2 and 3, to increase the difficulty of this test problem.

In this work, both observed and synthetic waveform data are generated by solving a 2D acoustic wave equation using a time-domain

pseudo-spectral method (Wang et al., 2019; Zhao et al., 2020) with a Ricker wavelet of 10 Hz dominant frequency. In addition, we add 1 percent uncorrelated Gaussian noise to the observed baseline and monitoring data. For both variational and least-squares inversions, the gradient of the data misfit with respect to the velocity model is calculated using the adjoint state method (Plessix, 2006; Fichtner et al., 2006).

3.2. Test of inversion method

We first test the effectiveness of the proposed time-lapse monitoring method in a relatively simple scenario in which the acquisition geometry used in the monitoring survey is exactly the same as that used in the baseline survey, as displayed in Fig. 2a. We first apply PSVI to invert the baseline survey data. Since it is often the case that we have limited prior knowledge about a target region before the baseline inversion, we adopt a non-informative uniform distribution over an interval of width 2500 m/s as the prior information for the baseline inversion (Fig. 3a). The likelihood function is characterised by a Gaussian with a diagonal covariance matrix (Eq. (2)), with a data noise level of 1%.

After inversion, we generate 500 samples from the posterior pdf. Fig. 4a shows statistics of these posterior samples. Each sample is then refined using Eq. (3) to find the *maximum a posteriori* (MAP) model solution. This adjustment is accomplished using the LBFGS optimisation method (Liu and Nocedal, 1989), with a maximum of 10 iterations. An early stop criterion, in which relative reduction of misfit function values between two adjacent iterations is smaller than 0.001, is applied in this step. The refinement of the 500 samples are parallelised using 1 Intel Xeon Platinum CPU with 48 cores. Differences between the two samples before and after the refinement step are also calculated. Figs. 4b and 4c display statistics of the 500 adjusted samples, and the corresponding samples of model differences, respectively. From top to bottom row, one random sample, the mean, and standard deviation maps of the three sets of samples are displayed. High consistency is observed between Figs. 4a and 4b, and model differences due to the refinement step in Fig. 4c are orders of magnitude smaller than the parameter values in Figs. 4a and 4b. This illustrates that the refinement step described in Eq. (3) does not change the mean or standard deviation statistics of the baseline inversion results significantly compared to the magnitude of the parameter values themselves, but that the change is of the same order of magnitude as the time-lapse velocity variations. The refined samples are used to perform the monitoring inversion.

For better comparison between the baseline and monitoring inversion results, we also display the baseline inversion results after refinement in Fig. 5a (i.e., Figs. 4b and 5a are the same). Since we use a uniform prior distribution (Fig. 3a) without injecting any additional correlation information into the inversion, all posterior samples are spatially rough with large velocity variations between neighbouring cells. The mean model offers a relatively accurate estimate of the true baseline model given that the observed data have a dominant frequency of only 10 Hz. The standard deviation map shows similar geometrical structures to the mean model. The uncertainties increase with depth since the sensitivity of seismic data decreases with depth. Similar patterns have also been observed in previous studies (Gebraad et al., 2020; Biswas and Sen, 2022).

We then perform the monitoring inversion using the method proposed above, in which each of the 500 baseline MAP model samples (after refinement) is updated deterministically to fit the monitoring data by minimising Eq. (4). The LBFGS method with a maximum of 30 iterations is again used to solve the optimisation problem (Liu and Nocedal, 1989). The same early stop criterion and computational resources are used as those used in the refinement step. During the monitoring inversion, we only update velocity values within the red box in Figs. 2 and 5 using the target-oriented scheme. We implement the double difference method and calculate the composite data set using Eq. (5). For the monitoring inversion, we define the likelihood using




Fig. 3. (a) Uniform prior distribution for the baseline inversion. Its upper and lower bounds vary with depth. (b) Prior standard deviation map of $\delta \mathbf{m}$ used to define $\Sigma_{\mathbf{m}_2}^{-1}$ in Eq. (4).




Fig. 4. Statistics of (a) posterior baseline inversion results, (b) baseline maximum a posteriori models after the refinement step in Eq. (3), and (c) pairwise differences between individual samples used to construct panels a and b. In each column, one random sample, the mean, and standard deviation maps are displayed from top to bottom. Note that panel b and Fig. 5a display the same results.

the same diagonal covariance matrix for data uncertainties as used in the baseline inversion. The prior model covariance matrix $\Sigma_{\mathbf{m}_2}^{-1}$ is also defined as a diagonal matrix. We assign the largest diagonal value ($100^2 \text{ m}^2/\text{s}^2$) to the model parameter at the CO_2 injection location (blue pluses in Fig. 2). This value decays as the spatial locations diverge from the injection point. We include the additional prior information that CO_2 , being less dense than brine, is likely to migrate upwards from the injection point over short time-scales (before significant dissolution occurs) after being injected into the subsurface (Ghosh and Ojha, 2020). This understanding suggests that time-lapse changes induced by the injection and migration of CO_2 are more probable above the injection area than below, so we reduce the corresponding diagonal elements in $\Sigma_{\mathbf{m}_2}^{-1}$ by a factor of 0.5 for model parameters located beneath the injection point. Fig. 3b displays the prior standard deviation map of velocity change used to define the diagonal elements of $\Sigma_{\mathbf{m}_2}^{-1}$ in Eq. (4).

Fig. 5b shows the results from the monitoring inversion. The first row displays the updated velocity model sample, starting from the baseline model sample displayed in the top row in Fig. 5a. The second and third rows represent the mean and standard deviation maps of the 500 updated MAP model samples obtained through deterministic (linearised) inversion. No significant differences are observed between Figs. 5a and 5b, since time-lapse changes are in this case 1 order of magnitude smaller than the actual velocity values of the target area.

We then calculate the MAP samples of time-lapse changes through $\delta \mathbf{m} = \mathbf{m}_2^* - \mathbf{m}_1^*$ for each pair of baseline and monitoring samples. Fig. 5c illustrates the corresponding results – the top row shows the difference between the two samples in the first row in Figs. 5a and 5b. The subsequent two rows show the mean and standard deviation maps of the MAP samples of time-lapse changes. Both the posterior sample and the mean map provide reasonably accurate estimates of the true velocity changes displayed in Fig. 2c. Some imaging artefacts which

reflect the geological features of the true velocity model are visible, but are much smaller than the actual time-lapse changes. Such artefacts are common in time-lapse FWI studies and have been noted in other research (Asnaashari et al., 2015; de Lima et al., 2024; Zhang and Curtis, 2024).

In the bottom row in Fig. 5c, the overall standard deviation values below the injection location are smaller than those above it (thus we observe obvious change in standard deviation values around the injection depth), due to the prior information injected during the monitoring inversion: that injected CO_2 is likely to migrate upwards. Higher uncertainties are present at locations where true velocity variations occur. In addition, since we impose strong prior information obtained from the baseline inversion and from the inherent characteristics of time-lapse problems (that time-lapse changes are small and localised), the standard deviation values in the bottom row in Fig. 5c are smaller than the values of the time-lapse changes themselves top two rows ensuring that this form of uncertainty analysis is useful for decision-making. This differs from the uncertainty results obtained from Bayesian time-lapse FWI presented in Zhang and Curtis (2024) and de Lima et al. (2024), in which significantly higher uncertainties are reported, because first they imposed weaker prior information than we impose here, and second they estimated Bayesian uncertainty on $\delta \mathbf{m}$ whereas we estimate nonlinear uncertainty on the MAP model of $\delta \mathbf{m}$.

To summarise, this example shows that the proposed method is able to detect time-lapse variations of the type and magnitude of those caused by CO_2 -related subsurface changes between baseline and monitoring surveys, and to quantify useful uncertainty information on the MAP given perfectly repeated acquisition geometries. If unexpected but significant subsurface changes occur, the estimated time-lapse changes might be biased.




Fig. 5. Statistics of (a) baseline MAP solutions, (b) monitoring MAP solutions, and (c) time-lapse changes. Note that panel a and Fig. 4b display the same results. In each column, one random sample, the mean, and standard deviation maps are displayed from top to bottom. Red box in each figure indicates the target region in which the model is updated using monitoring data. In this test, both baseline and monitor surveys have exactly the same acquisition geometry as that displayed in Fig. 2a.




Fig. 6. True velocity change models with different acquisition geometries in the monitoring survey. In each figure, numbers of sources and receivers are indicated in the title, and their locations are denoted by red stars and blue triangles, respectively. A red box illustrates the area in which velocity values are updated during the monitoring inversion.

Finally, we discuss the computational cost of the proposed method. For variational baseline inversion, we update the variational distribution for 5000 iterations, and 2 samples per iteration are used to estimate the gradients of the variational parameters, requiring 10,000 forward model evaluations. In the refinement and monitoring inversion steps, we update 500 model samples drawn from the baseline posterior pdf with 10 and 30 iterations, respectively. In addition, early stopping is applied during these two steps. In total, the number of FWI simulations should therefore be no more than 30,000. The reported computational cost could be further reduced by using fewer baseline posterior samples that are carefully designed so as to be representative of the key components of the probability distribution (Zhao and Curtis, 2024a). While these processing costs may seem high, below we show that they are compensated by the fact that ultra-sparse seismic acquisition geometries can be deployed, while still achieving accurate results with this processing method.

3.3. Time-lapse monitoring using sparse acquisition geometries

We now perform a set of tests designed to address critical questions regarding seismic survey configurations, to address the specific

questions: can sparse acquisition geometries detect time-lapse changes accurately? If so, what constitutes the simplest effective geometry, or in other words, the minimum numbers of sources and receivers that are required? Given the high costs associated with conducting seismic surveys, these questions are of practical importance: there is a clear need to minimise the numbers of sources and receivers in monitoring surveys while still achieving reasonably reliable time-lapse change estimates.

We consider 6 different source-receiver acquisition geometries. Starting from the densest geometry that replicates the configuration used in the previous test – i.e., 12 sources and 250 receivers (denoted S12-R250), we gradually reduce the numbers of sources and receivers used in the monitoring survey, respectively containing 6 sources and 50 receivers (S6-R50), 4 sources and 24 receivers (S4-R24), 3 sources and 12 receivers (S3-R12), 2 sources and 10 receivers (S2-R10), and 1 source and 6 receivers (S1-R6). The acquisition geometries are displayed in Fig. 6. This diminishes the amount of information provided by the monitoring data and thus the accuracy of time-lapse imaging, making it more difficult to discriminate true velocity variations from imaging artefacts and potentially leading to incorrect interpretation of




Fig. 7. Mean of MAP time-lapse changes obtained using different acquisition geometries in the monitoring survey. Only regions inside of the red boxes in Fig. 6 are displayed. In each figure, numbers of sources and receivers are indicated in the title. Dashed black lines in (a) display the locations of four vertical profiles used to compare the marginal pdf's of velocity changes in Fig. 10.




Fig. 8. Standard deviations of MAP time-lapse changes obtained using different acquisition geometries in the monitoring survey. In each figure, numbers of sources and receivers are indicated in the title.

CO₂ migration or leakage in the subsurface. Note that for simplicity, these acquisition geometries are defined heuristically without applying more sophisticated optimal experimental design methods (Maurer and Boerner, 1998; Curtis, 1999a,b, 2004a,b; Maurer et al., 2010; Bloem et al., 2020; Strutz and Curtis, 2024) since to-date these methods have only been developed and tested for FWI problems in which waveform physics is assumed to be linear (Maurer et al., 2017; Krampe et al., 2021; Mercier et al., 2025). We also hold the acquisition geometry of the baseline survey constant because a dense seismic survey is always likely to be performed initially, to identify the area targeted for CO₂ injection. The objective is to be able to monitor subsurface changes continuously over time at reduced monitoring survey cost by using fewer sources and receivers.

For each test involving different acquisition geometries, we update the same 500 posterior baseline samples with refinement (Eq. (3)) to perform the monitoring inversion. The same data covariance matrix and prior model used in the previous example are applied.

Fig. 7 illustrates the MAP mean of time-lapse changes obtained from different acquisition geometries in the monitoring survey. We only

display regions inside of the red boxes in Fig. 6 for clearer illustration. For each case, source and receiver numbers are indicated in the title of each figure, and their spatial locations are denoted by red stars and blue triangles in Fig. 6, respectively. Note that Fig. 7a (S12-R250) is exactly the same as that displayed in the second row in Fig. 5c. Although the shape of the true velocity changes is observed across all six mean maps, there is a noticeable decrease in the amplitude of the recovered time-lapse variations from Fig. 7a to 7f. Additional subsurface structures are observed as we simplify the acquisition geometry. These are imaging artefacts, which show similar amplitude to the true velocity changes but do not represent real changes. These phenomena become particularly obvious in Fig. 7f in which only one source is deployed. In this case it is difficult to distinguish between true velocity changes and imaging artefacts. As also shown in Zhang and Curtis (2024), these artefacts follow geological strata, so they can certainly bias interpretations of dynamics in time-lapse results. This suggests that employing a single source with sparse receivers may be insufficient for accurate and reliable monitoring of subsurface changes.

The corresponding standard deviation of the MAP models obtained using monitoring data from different acquisition geometries are shown




Fig. 9. Relative error maps of MAP time-lapse changes obtained using different acquisition geometries in the monitoring survey. The relative error is the absolute error between the mean and true velocity change models divided by the corresponding standard deviation at each point.




Fig. 10. Marginal pdfs of MAP time-lapse changes along four vertical profiles, obtained using different acquisition geometries in the monitoring survey. The locations of these profiles (top to bottom rows) are denoted by black dashed lines (left to right, respectively) in Fig. 7a. In each figure, black and red lines show the mean and true velocity changes, and two horizontal white lines show vertical locations of the inversion area in the monitoring inversion.

in Fig. 8 (again Fig. 8a and the last row in Fig. 5c display the same model). Overall in these maps, relatively higher uncertainties are present in areas with actual time-lapse variations, since these are the main sources of uncertainty in this time-lapse problem. Moreover, the magnitude of the standard deviations tends to decrease from Fig. 8a to 8f; this occurs because fewer data are provided by the monitoring

survey to update velocity models obtained from the baseline survey inversion, so the magnitude of estimated changes (and their variation between different samples) are smaller. Note that Fig. 8 displays the uncertainties in the velocity change $\delta \mathbf{m}$ rather than in the velocity model \mathbf{m}_2^* produced during the monitoring inversion. In an extreme situation where no monitoring data is available, the baseline velocities




Fig. 11. Acquisition geometries for the monitoring survey using one source and six receivers, but with different source locations. Key as in Fig. 6.




Fig. 12. Mean, standard deviation, and relative error maps of MAP time-lapse changes obtained using one source and six receivers, but with the different source locations displayed in Fig. 11.

would remain unchanged during the monitoring inversion. Then, both the mean and standard deviation of time-lapse changes calculated using the proposed method would be zero. This highlights the difference between this method and direct Bayesian inversion of velocity change $\delta \mathbf{m}$ in which posterior statistics of $\delta \mathbf{m}$ would be equal to the corresponding prior statistics if no monitoring data is provided. This arises because the second term in Eq. (4) regularises the inversion by encoding that \mathbf{m}_2^* should be close to \mathbf{m}_1^* unless the monitoring data dictates otherwise. In other words, the current method provides a conservative estimate of velocity change $\delta \mathbf{m}$: in effect, we are testing whether the monitoring data contains *evidence* of $|\delta \mathbf{m}| > 0$, given prior information about the range of $\delta \mathbf{m}$ that might be expected.

Fig. 9 shows the relative error maps, computed by dividing the absolute error between the true (Fig. 6) and mean (Fig. 7) time-lapse models by the standard deviation model (Fig. 8) at each point. Larger relative errors occur around the boundaries of true velocity variations due to the inability of low frequency waves to encode the location of an abrupt change in velocity, creating so-called uncertainty loops (Galetti et al., 2015). Such patterns prove that our method produces the expected characteristics of uncertainty results. Almost all relative errors are smaller than 3 which corroborates the accuracy of the results, except for those displayed in Fig. 9f.

To further analyse the results, in Fig. 10 we compare the marginal pdfs of MAP velocity changes from the six sets of results along four vertical profiles, with locations displayed by dashed black lines in Fig. 7a. From top to bottom, each row in Fig. 10 represents one profile from left to right in Fig. 7a. Red and black lines show the true and

mean of MAP velocity changes obtained using different acquisition geometries. The second profile is strategically placed to traverse two areas of velocity change which occur in close proximity, inherently challenging the ability of this method to distinguish between these nearby features. Similarly to those displayed in Fig. 7, 8 and 9, all velocity changes are observed clearly in these four profiles in Fig. 10. However, the results become less accurate with lower uncertainties (narrower marginal pdfs) as the numbers of sources and receivers decrease. In Fig. 10f, it becomes difficult to discriminate between the true velocity changes and imaging artefacts, especially in the second and third rows.

We perform two additional tests to explore the impact of using just one source in the monitoring survey, albeit with different source locations than those used above. All settings remain the same as those used to obtain the results displayed in Figs. 7f, 8f and 9f, other than the variation in source location displayed in Fig. 11. Fig. 12 displays three sets of results (Fig. 12b represents the results displayed above) obtained using different source locations in the monitoring survey, in which only the target region within the red boxes in Fig. 11 is displayed. We observe that none of these configurations successfully captures the true velocity variations, again proving that a single source (and a sparse array of receivers) is likely to be inadequate for accurate time-lapse monitoring. In addition, these tests underscore the importance of optimal experimental design (Maurer et al., 2010), particularly when employing sparse acquisition geometries in the monitoring survey: there is significant variation in results, which highlights the substantial




Fig. 13. Acquisition geometries for the monitoring survey using one source and 250 receivers, but with different source locations. Key as in Fig. 6: the 250 blue triangles are too close together to distinguish individually, mimicking the dense recordings on a fibre-optic cable.




Fig. 14. Mean, standard deviation, and relative error maps of MAP time-lapse changes obtained using one source and 250 receivers (to mimic fibre-based acquisition systems), but with the different source locations displayed in Fig. 13.

influence of source location on the accuracy of the subsurface velocity changes detected.

Below we additionally consider using 1 single source and 250 receivers along the seabed. This mimics a case in which we have a fibre cable that is fixed on the seabed, since fibre-based acquisition systems are becoming a standard and cheap way to create spatially dense sensors for geophysical monitoring. Figs. 13a and 13b show two different acquisition geometries considered, and the corresponding time-lapse results are displayed in Fig. 14. From top to bottom panel, each figure displays the mean velocity, standard deviation, and relative error maps of the MAP velocity changes. In Fig. 14a, the left time-lapse area is not recovered accurately due to the particular source location

used in this test. In Fig. 14b, the source location is the same as that considered in Figs. 7f, 8f and 9f. Compared to the results displayed above, the dense receiver layout provides significantly better time-lapse results with more accurate amplitudes. Some imaging artefacts are still observed in Fig. 14, which are slightly weaker than the changes that we wish to discriminate. This demonstrates that reasonably accurate monitoring with a single source would work only if a very dense array of receivers is used. However, this requirement is difficult to accomplish in realistic 3D monitoring.

These cumulative findings therefore prove that sparse acquisition geometries can indeed be deployed in the monitoring survey, and it is feasible to obtain accurate time-lapse estimates along with their

corresponding uncertainties. However, in order to use a single source it seems likely to be necessary to use a dense grid of receivers.

4. Discussion

In this work, we proposed a time-lapse FWI method that leverages strong prior information derived from baseline inversion results to enhance monitoring survey inversions, and a fully nonlinear baseline inversion in order to provide many samples of possible initial velocity structures and maximum a posteriori (MAP) solutions. In our method, the calculated time-lapse estimates represent model differences between each pair of MAP solutions from baseline and monitoring inversions. These MAP-to-MAP comparisons are therefore meaningful, and the results enable uncertainties in the MAP models of time-lapse changes to be estimated. Note that we do not apply much regularisation to the baseline and monitoring inversions individually. The posterior distributions of each of those surveys individually would therefore describe a strongly under-determined result, with many hyper (high-dimensional) ridges or plateaus of near-constant values in the posterior pdf. As a result, each adjusted sample is expected to converge to different points on the same ridge or plateau, on different ridges or plateaus, or more generally, on the posterior MAP manifold. That is to say, we adjust the baseline posterior samples to describe changes in the MAP manifold between baseline and monitoring surveys.

While weak prior information was used for the baseline inversion because limited prior knowledge about the target is usually available before conducting the baseline survey, after the baseline inversion is complete our understanding of the target region improves. We would know the exact location of injection points, and the amount of CO₂ injected into the subsurface, allowing more precise identification of potential locations of subsurface changes. Such strong prior information enables us to better constrain the monitoring inversion, so we obtain relatively small uncertainties in time-lapse changes, even though uncertainties associated with both the baseline and monitoring results remain significantly higher as illustrated in Fig. 5.

We noted above that in our approach, the second term in Eq. (4) regularises the time-lapse changes towards zero. Our method therefore tests whether the monitoring data contains sufficient evidence to estimate $|\delta \mathbf{m}| > 0$. An alternative approach would be to use each sample \mathbf{m}_1^* of the baseline MAP solution to generate a model \mathbf{m}_2' (or a distribution of such models) that encodes prior information about likely time-lapse velocity changes when starting from baseline model \mathbf{m}_1^* . In that case, the second term in Eq. (4) could be replaced by a term that regularises towards this new distribution of \mathbf{m}_2' ; in the limit of the monitoring data containing no relevant information, our estimate of $\delta \mathbf{m}$ would then reflect the prior information on $\delta \mathbf{m}$. However, the latter approach would not provide a conservative test of the information content provided by sparse acquisition geometries, which was one of our goals.

MAP models from the baseline inversion are updated by fine-tuning them to the monitoring data, using a deterministic (linearised) inversion, which is valid under the assumption that time-lapse changes are typically small, particularly in subsurface CCS projects (Egorov et al., 2017). Therefore, in our method we assume that uncertainties in the MAP monitoring models \mathbf{m}_2^* are similar to those in the MAP baseline models \mathbf{m}_1^* , and that a linearised inversion method (with a low number of iterations) for the monitoring inversion is sufficient to update our baseline MAP uncertainty estimates to monitoring MAP uncertainties. The updated samples are therefore treated as MAP samples of the monitoring inversion.

In contrast to our method, Fu and Innanen (2022) implemented a linearised deterministic FWI method for their baseline inversion, and that inversion result served as the starting point for multiple Markov chains that inverted the monitoring data in a probabilistic manner. Their choice of a deterministic baseline inversion was primarily to minimise computational demands of the MCMC algorithm. However,

relying on a single initial model can make it prohibitively expensive to thoroughly explore the parameter space (to quantify uncertainties) during the monitoring inversion using MCMC methods. Although efficiency is improved in the baseline inversion, the complexity and computational cost of the monitoring inversion are in principle increased. This is particularly true when the monitoring data are collected using sparse acquisition geometries, so larger hyper-volumes of parameter space remain consistent with the measured monitoring data.

We implement the double difference method in the monitoring inversion, which requires high repeatability between the baseline and monitoring surveys. In instances where this is not fully satisfied, for example if the baseline and monitoring surveys have slightly different locations, some pre-processing steps such as time-lapse binning, data interpolation and wavelet adjustment should be performed first (Fu et al., 2020). Even if obtaining both baseline and monitoring data from an identical acquisition geometry becomes unfeasible, Asnaashari et al. (2015) demonstrated that the sequential difference strategy with a target-oriented scheme can still provide similarly robust results to those obtained using the double difference strategy. Zhang and Curtis (2024) also showed that far greater robustness against larger changes in acquisition geometries is obtained by using nonlinear methods such as those used to process our baseline survey.

In the monitoring inversion, we use a target-oriented strategy to remove imaging artefacts outside of a predefined area. Nevertheless, forward and adjoint simulations are still performed in the same volume as in the baseline inversion. Several techniques can be implemented to improve this. For example, we can use some efficient yet accurate target-oriented wavefield solvers during the monitoring inversion (Robertsson and Chapman, 2000; van Manen et al., 2007; Vasmel et al., 2013), or we might focus on waveform data within a window specifically relevant to the expected time-lapse changes.

There are several differences between the method proposed herein and existing fully nonlinear Bayesian time-lapse FWI methods (e.g., Zhang and Curtis, 2024; de Lima et al., 2024). First, our method characterises the range of MAP solutions by minimising Eq. (4), rather than solving the monitoring inversion using Bayesian methods. This is done firstly because the limited data collected from sparse acquisition geometries are likely to be insufficient to constrain a fully nonlinear Bayesian inverse problem well in the monitoring inversion: the fewer data available for fully nonlinear inversions, the larger is the hyper-volume of parameter space that must be explored to characterise the solution. This is likely to make such inversions infeasible to solve numerically due to the curse of dimensionality (Curtis and Lomax, 2001). It nevertheless remains computationally feasible to use fully nonlinear methods for the (non-sparse) baseline survey, and to update each baseline MAP solution using (sparse) monitoring data with linearised methods. In addition, we inject strong prior information obtained from the baseline inversion and interpret differences between monitoring and baseline MAP models as models of time-lapse changes. Thus the obtained uncertainties are lower than those from Bayesian methods, making them easier to analyse in practise. This is what allows us to use sparse monitoring surveys to reduce cost. Secondly, our method provides samples of velocity changes directly. While a Bayesian framework provides full posterior pdfs for both the baseline and monitoring inversions, deriving posterior samples of time-lapse changes is not straightforward, and these are important for post-inversion decision-making processes (Arnold and Curtis, 2018). The joint Bayesian inversion approach (Zhang and Curtis, 2024) provides such samples simultaneously, yet the method combines the baseline and time-lapse model vectors, thereby doubling the dimensionality of the inverse problem, making it more difficult and expensive to solve. Thirdly, we need to note that results from our method do not represent Bayesian uncertainties, since monitoring MAP samples are obtained using deterministic (linearised) methods. Nevertheless, they represent a set of time-lapse models that fit the baseline and

monitoring data, and which are well distributed throughout the MAP model manifold.

In real applications, more complex but realistic factors might affect the performance of the proposed method. For example, it would be significantly more difficult to estimate accurate uncertainties in spatially 3D FWI problems (Zhang et al., 2023; Lomas et al., 2023; Zhao and Curtis, 2024b). The effects of shear waves and surface waves would influence the applicability of the method in real-world scenarios when elastic properties are considered (Gebraad et al., 2020; Berti et al., 2025).

In addition, accurate estimation of the target region is important since a large region would reduce the accuracy of the obtained time-lapse estimates, yet if a small target region is considered, a situation in which CO₂ migrates outside of the predefined target region might arise. We therefore risk ignoring this subjectively, leading to biased results. This will be more likely to happen when undetected faults exist within a survey region. In reality, this target region might be defined more accurately by running a fluid flow simulation model. For example for some CCS sites, the CO₂ storage reservoir is thin so that it extends far further in a lateral direction than our examples (Sinha et al., 2024). A reservoir simulation will still be able to estimate a reasonably accurate target region that captures possible leakage sites before the monitoring inversion or survey. The proposed method should therefore also work in that case.

In our examples, we did not employ a flow model explicitly. However, we did consider the injection explicitly in our additional prior information for the monitoring inversion: in a real CCS project we would know the exact injection location (e.g., blue pluses in Fig. 2) and the amount of injected CO₂, which will help us to better constrain subsurface changes that may be caused by CO₂ injection and migration. Emulating this process, we defined the prior uncertainties on velocity changes displayed in Fig. 3b.

Note that for some CCS projects, reservoirs and target regions would have simpler subsurface structures with less (lateral) heterogeneity than our test cases. Examples include the Sleipner site (Chadwick and Noy, 2015) and the Pelican site (Hoffman, 2018). We use the Marmousi model for our examples; this might be more geologically complex than some real cases, however, from the point of view of testing our inversion method, if the method works well for geologically complicated models, then it will certainly work with a much simpler model (and the results show that the method works well).

Our method produces samples of MAP models of the time-lapse changes. A next step could involve using the time-lapse samples to solve an interrogation problem, for example to answer questions such as whether CO₂ is leaking into particular subsurface volumes (Arnold and Curtis, 2018). Moreover, in our example we use 500 samples for the monitoring inversion, each updated through a linearised inversion. To reduce its computational cost, we could use a technique called boosting variational inference (Zhao and Curtis, 2024a), which selects tens of representative samples that effectively summarise the range of uncertainties required for the monitoring inversion.

We have shown that sparse acquisition geometries are feasible for monitoring subsurface changes in time-lapse FWI problems. However, numerical tests demonstrate that it appears to be insufficient to use one single source for such purposes, unless a very dense array of receivers is used such as in fibre-based acquisition systems. Nevertheless, our finding is particularly useful in situations in which a cost-effective solution for continuous monitoring is required. For example in a marine seismology setting, we could assume source-receiver reciprocity applies to the results shown herein, and obtain approximately similar survey designs by deploying one or two ocean bottom nodes (OBN) on the seabed and firing a number of marine sources, since the cost of deploying each OBN may be higher than that of deploying a source at one location. Optimal experimental design methods may also be important to improve the results presented previously, by identifying the minimum numbers of sources and receivers and their optimal locations that ensure a desired expected level of accuracy in results (Strutz and Curtis, 2024).

5. Conclusion

We introduce a method to perform time-lapse seismic full waveform inversion and quantify statistics of its uncertainties, with a particular focus on monitoring subsurface CO₂ migration using ultra-sparse acquisition geometries with low monitoring cost. We first use variational Bayesian inference to invert the baseline data. This provides strong prior information about the subsurface structure, that is subsequently injected into the monitoring inversion, executed using a deterministic inversion framework. To minimise imaging artefacts unrelated to actual subsurface changes, we employ a target-oriented scheme. In a 2D acoustic time-lapse FWI example, we simulate putative subsurface CO₂ storage scenarios in which velocity reductions are induced by CO₂ injection and migration into a fault and a secondary reservoir. The method provides accurate estimates of velocity changes and generates uncertainties in best-fitting models that are smaller in magnitude than the velocity changes themselves, as would be required to make these results useful for risk assessment applications. We also investigate the effectiveness of our methodology under various sparse acquisition geometries, involving different numbers of sources and receivers within the monitoring survey. The results demonstrate that sparse acquisition geometries can be used to reduce monitoring costs, without compromising the accuracy of the inversion, but if only a single source (receiver) is to be used then this must be compensated by a spatially dense array of receivers (sources). Overall, we conclude that our methodology is capable of detecting velocity changes and quantifying reasonable uncertainties in time-lapse imaging.

CRedit authorship contribution statement

Xuebin Zhao: Writing – original draft, Software, Data curation, Visualization, Methodology, Conceptualization. **Andrew Curtis:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Validation, Resources, Investigation, Formal analysis.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors thank Edinburgh Imaging Project (EIP - <https://blogs.ed.ac.uk/imaging/>) sponsors (BP and TotalEnergies) for supporting this research. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Data availability

Data will be made available on request.

References

- Arnold, R., Curtis, A., 2018. Interrogation theory. *Geophys. J. Int.* 214 (3), 1830–1846.
- Asnaashari, A., Brossier, R., Garambois, S., Audebert, F., Thore, P., Virieux, J., 2015. Time-lapse seismic imaging using regularized full-waveform inversion with a prior model: which strategy? *Geophys. Prospect.* 63 (1), 78–98.
- Asnaashari, A., Brossier, R., Garambois, S., Virieux, J., Audebert, F., Thore, P., 2012. Time-lapse imaging using regularized FWI: a robustness study. In: *SEG International Exposition and Annual Meeting*. SEG, pp. SEG–2012.
- Berti, S., Ravasi, M., Aleardi, M., Stucchi, E., 2025. Bayesian full waveform inversion of surface waves with annealed stein variational gradient descent. *Geophys. J. Int.* 241 (1), 641–657.
- Biswas, R., Sen, M.K., 2022. Transdimensional 2D full-waveform inversion and uncertainty estimation. *arXiv preprint arXiv:2201.09334*.

- Blei, D.M., Kucukelbir, A., McAuliffe, J.D., 2017. Variational inference: A review for statisticians. *J. Amer. Statist. Assoc.* 112 (518), 859–877.
- Bloem, H., Curtis, A., Maurer, H., 2020. Experimental design for fully nonlinear source location problems: which method should I choose? *Geophys. J. Int.* 223 (2), 944–958.
- Bodin, T., Sambridge, M., 2009. Seismic tomography with the reversible jump algorithm. *Geophys. J. Int.* 178 (3), 1411–1436.
- Capdeville, Y., Chaljub, E., Montagner, J.P., 2003. Coupling the spectral element method with a modal solution for elastic wave propagation in global Earth models. *Geophys. J. Int.* 152 (1), 34–67.
- Chadwick, R.A., Noy, D.J., 2015. Underground CO₂ storage: demonstrating regulatory conformance by convergence of history-matched modeled and observed CO₂ plume behavior using Sleipner time-lapse seismics. *Greenh. Gases: Sci. Technol.* 5 (3), 305–322.
- Curtis, A., 1999a. Optimal design of focused experiments and surveys. *Geophys. J. Int.* 139 (1), 205–215.
- Curtis, A., 1999b. Optimal experiment design: cross-borehole tomographic examples. *Geophys. J. Int.* 136 (3), 637–650.
- Curtis, A., 2004a. Theory of model-based geophysical survey and experimental design: Part 1—Linear problems. *Lead. Edge* 23 (10), 997–1004.
- Curtis, A., 2004b. Theory of model-based geophysical survey and experimental design: Part 2—Nonlinear problems. *Lead. Edge* 23 (11), 1112–1117.
- Curtis, A., Lomax, A., 2001. Prior information, sampling distributions, and the curse of dimensionality. *Geophysics* 66 (2), 372–378.
- de Lima, P.D.S., Ferreira, M.S., Corso, G., de Araújo, J.M., 2024. Bayesian time-lapse full waveform inversion using Hamiltonian Monte Carlo. *Geophys. Prospect.* 72 (9), 3381–3398. <http://dx.doi.org/10.1111/1365-2478.13604>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1365-2478.13604>.
- Egorov, A., Pevzner, R., Bóna, A., Glubokovskikh, S., Puzryev, V., Tertyshnikov, K., Gurevich, B., 2017. Time-lapse full waveform inversion of vertical seismic profile data: Workflow and application to the CO₂CRC otway project. *Geophys. Res. Lett.* 44 (14), 7211–7218.
- Emberley, S., Hutcheon, I., Shevalier, M., Durocher, K., Mayer, B., Gunter, W., Perkins, E., 2005. Monitoring of fluid–rock interaction and CO₂ storage through produced fluid sampling at the weyburn CO₂-injection enhanced oil recovery site, saskatchewan, Canada. *Appl. Geochem.* 20 (6), 1131–1157.
- Fawad, M., Mondol, N.H., 2021. Monitoring geological storage of CO₂: A new approach. *Sci. Rep.* 11 (1), 5942.
- Fichtner, A., Bunge, H.-P., Igel, H., 2006. The adjoint method in seismology: I. Theory. *Phys. Earth Planet. Inter.* 157 (1–2), 86–104.
- Fichtner, A., Kennett, B.L., Igel, H., Bunge, H.-P., 2009. Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods. *Geophys. J. Int.* 179 (3), 1703–1725.
- Fichtner, A., Simutė, S., 2018. Hamiltonian Monte Carlo inversion of seismic sources in complex media. *J. Geophys. Res.: Solid Earth* 123 (4), 2984–2999.
- Fu, X., Innanen, K.A., 2022. A time-domain multisource Bayesian Markov chain Monte Carlo formulation of time-lapse seismic waveform inversion. *Geophysics* 87 (4), R349–R361.
- Fu, X., Romahn, S., Innanen, K., 2020. Double-wavelet double-difference time-lapse waveform inversion. In: SEG International Exposition and Annual Meeting. SEG, D041S092R002.
- Galetti, E., Curtis, A., Meles, G.A., Baptie, B., 2015. Uncertainty loops in travel-time tomography from nonlinear wave physics. *Phys. Rev. Lett.* 114 (14), 148501.
- Gasperikova, E., Li, Y., 2021. Time-lapse electromagnetic and gravity methods in carbon storage monitoring. *Lead. Edge* 40 (6), 442–446.
- Gebraad, L., Boehm, C., Fichtner, A., 2020. Bayesian elastic full-waveform inversion using Hamiltonian Monte Carlo. *J. Geophys. Res.: Solid Earth* 125 (3), e2019JB018428.
- Gelman, A., Rubin, D.B., 1992. Inference from iterative simulation using multiple sequences. *Statist. Sci.* 7 (4), 457–472.
- Ghosh, R., Ojha, M., 2020. Prediction of elastic properties within CO₂ plume at sleipner field using AVS inversion modified for thin-layer reflections guided by uncertainty estimation. *J. Geophys. Res.: Solid Earth* 125 (11), e2020JB019782.
- Guo, B., Zheng, Z., Bandilla, K.W., Celia, M.A., Stone, H.A., 2016. Flow regime analysis for geologic CO₂ sequestration and other subsurface fluid injections. *Int. J. Greenh. Gas Control.* 53, 284–291.
- Harbert, W., Goodman, A., Spaulding, R., Haljasmaa, I., Crandall, D., Sanguinito, S., Kutcho, B., Tkach, M., Fuchs, S., Werth, C.J., et al., 2020. CO₂ induced changes in mount simon sandstone: Understanding links to post CO₂ injection monitoring, seismicity, and reservoir integrity. *Int. J. Greenh. Gas Control.* 100, 103109.
- Hoffman, N., 2018. The CarbonNet project's pelican storage site in the gippsland basin. In: 14th Greenhouse Gas Control Technologies Conference Melbourne. pp. 21–26.
- Izzatullah, M., Alali, A., Ravasi, M., Alkhalifah, T., 2024. Physics-reliable frugal local uncertainty analysis for full waveform inversion. *Geophys. Prospect.* 72 (7), 2718–2738. <http://dx.doi.org/10.1111/1365-2478.13528>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1365-2478.13528>.
- Khoshkholgh, S., Zunino, A., Mosegaard, K., 2021. Informed proposal Monte Carlo. *Geophys. J. Int.* 226 (2), 1239–1248.
- Khoshkholgh, S., Zunino, A., Mosegaard, K., 2022. Full-waveform inversion by informed-proposal Monte Carlo. *Geophys. J. Int.* 230 (3), 1824–1833.
- Kim, J., Nam, M.J., Matsuoka, T., 2013. Estimation of CO₂ saturation during both CO₂ drainage and imbibition processes based on both seismic velocity and electrical resistivity measurements. *Geophys. J. Int.* 195 (1), 292–300.
- Kolditz, O., Bauer, S., Böttcher, N., Elsworth, D., Görke, U.-J., McDermott, C.-I., Park, C.-H., Singh, A.K., Taron, J., Wang, W., 2012. Numerical simulation of two-phase flow in deformable porous media: Application to carbon dioxide storage in the subsurface. *Math. Comput. Simulation* 82 (10), 1919–1935.
- Kotsi, M., Malcolm, A., Ely, G., 2020a. Time-lapse full-waveform inversion using Hamiltonian Monte Carlo: a proof of concept. In: SEG Technical Program Expanded Abstracts 2020. Society of Exploration Geophysicists, pp. 845–849.
- Kotsi, M., Malcolm, A., Ely, G., 2020b. Uncertainty quantification in time-lapse seismic imaging: A full-waveform approach. *Geophys. J. Int.* 222 (2), 1245–1263.
- Krampe, V., Edme, P., Maurer, H., 2021. Optimized experimental design for seismic full waveform inversion: A computationally efficient method including a flexible implementation of acquisition costs. *Geophys. Prospect.* 69 (1), 152–166.
- Kullback, S., Leibler, R.A., 1951. On information and sufficiency. *Ann. Math. Stat.* 22 (1), 79–86.
- Lee, H., Calvin, K., Dasgupta, D., Krinner, G., Mukherji, A., Thorne, P., Trisos, C., Romero, J., Aldunce, P., Barret, K., et al., 2023. IPCC, 2023: Climate change 2023: Synthesis report, summary for policymakers. Contribution of working groups I, II and III to the sixth assessment report of the intergovernmental panel on climate change [core writing team, h. Lee and j. Romero (eds.)]. IPCC, Geneva, Switzerland..
- Levy, S., Laloy, E., Linde, N., 2022. Variational Bayesian inference with complex geostatistical priors using inverse autoregressive flows. *Comput. Geosci.* 105263.
- Li, D., Peng, S., Guo, Y., Lu, Y., Cui, X., 2021. CO₂ storage monitoring based on time-lapse seismic data via deep learning. *Int. J. Greenh. Gas Control.* 108, 103336.
- Li, C., Zhang, K., Wang, Y., Guo, C., Maggi, F., 2016. Experimental and numerical analysis of reservoir performance for geological CO₂ storage in the ordos Basin in China. *Int. J. Greenh. Gas Control.* 45, 216–232.
- Liu, D.C., Nocedal, J., 1989. On the limited memory BFGS method for large scale optimization. *Math. Program.* 45 (1), 503–528.
- 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. In: 84th EAGE Annual Conference & Exhibition, vol. 2023, (1), European Association of Geoscientists & Engineers, pp. 1–5.
- Lumley, D.E., 2001. Time-lapse seismic reservoir monitoring. *Geophysics* 66 (1), 50–53.
- Maharramov, M., Biondi, B., 2014. Joint full-waveform inversion of time-lapse seismic data sets. In: SEG Technical Program Expanded Abstracts 2014. Society of Exploration Geophysicists, pp. 954–959.
- Maharramov, M., Biondi, B.L., Meadows, M.A., 2016. Time-lapse inverse theory with applications. *Geophysics* 81 (6), R485–R501.
- van Manen, D.-J., Robertsson, J.O., Curtis, A., 2007. Exact wave field simulation for finite-volume scattering problems. *J. Acoust. Soc. Am.* 122 (4), EL115–EL121.
- Martin, G.S., Wiley, R., Marfurt, K.J., 2006. Marmousi2: An elastic upgrade for Marmousi. *Lead. Edge* 25 (2), 156–166.
- Maurer, H., Boerner, D.E., 1998. Optimized and robust experimental design: a non-linear application to EM sounding. *Geophys. J. Int.* 132 (2), 458–468.
- Maurer, H., Curtis, A., Boerner, D.E., 2010. Recent advances in optimized geophysical survey design. *Geophysics* 75 (5), 75A177–75A194.
- Maurer, H., Nuber, A., Martiartu, N.K., Reiser, F., Boehm, C., Manukan, E., Schmeltz, C., Fichtner, A., 2017. Optimized experimental design in the context of seismic full waveform inversion and seismic waveform imaging. In: *Advances in Geophysics*, vol. 58, Elsevier, pp. 1–45.
- Mercier, A., Boehm, C., Maurer, H., 2025. Designing full waveform inverse problems: a combined data and model approach. *Geophys. J. Int.* 241 (3), 1478–1493.
- Mosegaard, K., Sambridge, M., 2002. Monte Carlo analysis of inverse problems. *Inverse Problems* 18 (3), R29.
- Mosegaard, K., Tarantola, A., 1995. Monte Carlo sampling of solutions to inverse problems. *J. Geophys. Res.: Solid Earth* 100 (B7), 12431–12447.
- Nakata, R., Jang, U.-G., Lumley, D., Mouri, T., Nakatsukasa, M., Takanashi, M., Kato, A., 2022. Seismic time-lapse monitoring of near-surface microbubble water injection by full waveform inversion. *Geophys. Res. Lett.* 49 (24), e2022GL098734.
- Nakata, R., Nakata, N., Girard, A.J., Lumley, D., Ichikawa, M., Kato, A., Xue, Z., 2024. Time-lapse velocity change estimation by anisotropic elastic full waveform inversion for CO₂ sequestration at the Nagaoka CCS site. *Int. J. Greenh. Gas Control.* 136, 104176.
- Nawaz, A., Curtis, A., 2018. Variational Bayesian inversion (VBI) of quasi-localized seismic attributes for the spatial distribution of geological facies. *Geophys. J. Int.* 214 (2), 845–875.
- Nawaz, A., Curtis, A., 2019. Rapid discriminative variational Bayesian inversion of geophysical data for the spatial distribution of geological properties. *J. Geophys. Res.: Solid Earth* 124 (6), 5867–5887.
- Plessix, R.-E., 2006. A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophys. J. Int.* 167 (2), 495–503.
- Plessix, R.-E., Michelet, S., Rynja, H., Kuehl, H., Perkins, C., de Maag, J., Hatchell, P., 2010. Some 3D applications of full waveform inversion. In: 72nd EAGE Conference and Exhibition-Workshops and Fieldtrips. European Association of Geoscientists & Engineers, pp. cp-162.

- Raknes, E.B., Arntsen, B., 2014. Time-lapse full-waveform inversion of limited-offset seismic data using a local migration regularization. *Geophysics* 79 (3), WA117–WA128.
- Robertsson, J.O., Chapman, C.H., 2000. An efficient method for calculating finite-difference seismograms after model alterations. *Geophysics* 65 (3), 907–918.
- Sambridge, M., Gallagher, K., Jackson, A., Rickwood, P., 2006. Trans-dimensional inverse problems, model comparison and the evidence. *Geophys. J. Int.* 167 (2), 528–542.
- Sen, M.K., Stoffa, P.L., 2013. *Global optimization methods in geophysical inversion*. Cambridge University Press.
- Shahkarami, A., Mohaghegh, S., Gholami, V., Haghighat, A., Moreno, D., 2014. Modeling pressure and saturation distribution in a CO₂ storage project using a surrogate reservoir model (SRM). *Greenh. Gases: Sci. Technol.* 4 (3), 289–315.
- Siahkoobi, A., Rizzuti, G., Louboutin, M., Witte, P.A., Herrmann, F.J., 2021. Preconditioned training of normalizing flows for variational inference in inverse problems. *arXiv preprint arXiv:2101.03709*.
- Siahkoobi, A., Rizzuti, G., Orozco, R., Herrmann, F.J., 2023. Reliable amortized variational inference with physics-based latent distribution correction. *Geophysics* 88 (3), 1–137.
- Sinha, M., Saygin, E., Ross, A.S., Ricard, L., 2024. Seismic monitoring of CCS with active and passive data: A synthetic feasibility study based on pelican site, Australia. *Int. J. Greenh. Gas Control* 139, 104277.
- Smith, J.D., Ross, Z.E., Azizadneshehi, K., Muir, J.B., 2022. HypoSVI: Hypocentre inversion with stein variational inference and physics informed neural networks. *Geophys. J. Int.* 228 (1), 698–710.
- Stork, A.L., Allmark, C., Curtis, A., Kendall, J.-M., White, D.J., 2018. Assessing the potential to use repeated ambient noise seismic tomography to detect CO₂ leaks: Application to the Aquistore storage site. *Int. J. Greenh. Gas Control* 71, 20–35.
- Strutz, D., Curtis, A., 2024. Variational Bayesian experimental design for geophysical applications: seismic source location, amplitude versus offset inversion, and estimating CO₂ saturations in a subsurface reservoir. *Geophys. J. Int.* 236 (3), 1309–1331.
- Sun, C., Malcolm, A., Kumar, R., Mao, W., 2024. Enabling uncertainty quantification in a standard full-waveform inversion method using normalizing flows. *Geophysics* 89 (5), R493–R507.
- Sun, L., Wang, L., Xu, G., Wu, Q., 2023. A new method of variational Bayesian slip distribution inversion. *J. Geod.* 97 (1), 10.
- Tarantola, A., 1984. Inversion of seismic reflection data in the acoustic approximation. *Geophysics* 49 (8), 1259–1266.
- Tarantola, A., 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation*, vol. 89, SIAM.
- Vasmel, M., Robertsson, J.O., van Manen, D.-J., Curtis, A., 2013. Immersive experimentation in a wave propagation laboratory. *J. Acoust. Soc. Am.* 134 (6), EL492–EL498.
- Virieux, J., Operto, S., 2009. An overview of full-waveform inversion in exploration geophysics. *Geophysics* 74 (6), WCC1–WCC26.
- Walker, M., Curtis, A., 2014. Varying prior information in Bayesian inversion. *Inverse Problems* 30 (6), 065002.
- Wang, Y., Zhou, H., Zhao, X., Zhang, Q., Zhao, P., Yu, X., Chen, Y., 2019. Cu Q-RTM: A CUDA-based code package for stable and efficient Q-compensated reverse time migration. *Geophysics* 84 (1), F1–F15.
- Watanabe, T., Shimizu, S., Asakawa, E., Matsuoka, T., 2004. Differential waveform tomography for time-lapse crosswell seismic data with application to gas hydrate production monitoring. In: *SEG Technical Program Expanded Abstracts 2004*. Society of Exploration Geophysicists, pp. 2323–2326.
- Wei, X., Sun, J., 2021. Uncertainty analysis of 3D potential-field deterministic inversion using mixed Lp norms. *Geophysics* 86 (6), G133–G158.
- Yang, D., Liu, F., Morton, S., Malcolm, A., Fehler, M., 2016. Time-lapse full-waveform inversion with ocean-bottom-cable data: Application on Valhall field. *Geophysics* 81 (4), R225–R235.
- Zhang, C., Bütetage, J., Kjellström, H., Mandt, S., 2018. Advances in variational inference. *IEEE Trans. Pattern Anal. Mach. Intell.* 41 (8), 2008–2026.
- Zhang, X., Curtis, A., 2020. Seismic tomography using variational inference methods. *J. Geophys. Res.: Solid Earth* 125 (4), e2019JB018589.
- Zhang, X., Curtis, A., 2022. Interrogating probabilistic inversion results for subsurface structural information. *Geophys. J. Int.* 229 (2), 750–757.
- Zhang, X., Curtis, A., 2024. Bayesian variational time-lapse full waveform inversion. *Geophys. J. Int.* 237 (3), 1624–1638.
- Zhang, X., Lomas, A., Zhou, M., Zheng, Y., Curtis, A., 2023. 3D Bayesian variational full waveform inversion. *Geophys. J. Int.* 234 (1), 546–561.
- Zhang, Z., Huang, L., 2013. Double-difference elastic-waveform inversion with prior information for time-lapse monitoring. *Geophysics* 78 (6), R259–R273.
- Zhao, X., Curtis, A., 2024a. Bayesian inversion, uncertainty analysis and interrogation using boosting variational inference. *J. Geophys. Res.: Solid Earth* 129 (1), e2023JB027789.
- Zhao, X., Curtis, A., 2024b. Efficient Bayesian full waveform inversion and analysis of prior hypotheses in 3D. *arXiv preprint arXiv:2409.09746*.
- Zhao, X., Curtis, A., 2024c. Physically structured variational inference for Bayesian full waveform inversion. *J. Geophys. Res.: Solid Earth* 129 (11), e2024JB029557.
- Zhao, X., Curtis, A., 2024d. Variational prior replacement in Bayesian inference and inversion. *Geophys. J. Int.* 239 (2), 1236–1256.
- Zhao, X., Curtis, A., Zhang, X., 2021. Bayesian seismic tomography using normalizing flows. *Geophys. J. Int.* 228 (1), 213–239.
- Zhao, X., Curtis, A., Zhang, X., 2022. Interrogating subsurface structures using probabilistic tomography: an example assessing the volume of Irish sea basins. *J. Geophys. Res.: Solid Earth* 127 (4), e2022JB024098.
- Zhao, X., Zhou, H., Chen, H., Wang, Y., 2020. Domain decomposition for large-scale viscoacoustic wave simulation using localized pseudo-spectral method. *IEEE Trans. Geosci. Remote Sens.* 1–14.
- Zhdanov, M.S., Endo, M., Black, N., Spangler, L., Fairweather, S., Hibbs, A., Eiskamp, G., Will, R., 2013. Electromagnetic monitoring of CO₂ sequestration in deep reservoirs. *First Break* 31 (2).
- Zheng, Y., Barton, P., Singh, S., 2011. Strategies for elastic full waveform inversion of time-lapse ocean bottom cable (OBC) seismic data. In: *SEG Technical Program Expanded Abstracts 2011*. Society of Exploration Geophysicists, pp. 4195–4200.
- Zidan, A., Li, Y., Cheng, A., 2022. Regularized seismic amplitude inversion via variational inference. *Geophys. Prospect.* 70 (9), 1507–1527.
- Zunino, A., Gebraad, L., Ghirotto, A., Fichtner, A., 2023. HMCLab: a framework for solving diverse geophysical inverse problems using the Hamiltonian Monte Carlo method. *Geophys. J. Int.* 235 (3), 2979–2991.