Variational Inference for Uncertainty Quantification of 4D FWI Results

Fuchun Gao[1], Yiran Huang[2], Carlos Calderón-Macías[1] and Faqi Liu[1]
1: TGS; 2: Brown University

Summary

Four-dimensional full waveform inversion (4D FWI) has been increasingly applied in recent years in time lapse monitoring projects. Like many geophysical exploration techniques, 4D FWI involves inherent uncertainty because FWI is a non-linear ill-posed problem characterized by null space and non-uniqueness. Quantifying uncertainty of 4D FWI models is essential for informed drilling decisions, particularly in terms of cost efficiency and operational safety. Here we propose a method and workflow to quantify 4D FWI uncertainty using a Joint Stein Variational Gradient Descent (JSVGD) method, a tailored variational inference approach for 4D problems. Numerical tests demonstrate its computational affordability and practicality. For the demonstration of the proposed approach and its usefulness we utilize an example with non-repeatable acquisition geometries.

Introduction

Compared to conventional image-based 4D seismic processing, 4D changes in reservoir physical properties can be recovered more quantitatively and at higher resolution using 4D full waveform inversion (FWI). Numerous implementations of 4DFWI have been proposed to directly invert time-lapse model differences (e.g., Yang et al., 2013; Musa et al., 2015; Liu and Tsvankin, 2022; Dutta et al., 2023; Pintus et al., 2023; Gao et al., 2024), particularly in marine environments for reservoir monitoring. Because the resulting 4D model can strongly influence cost, operational safety, and reserve estimation accuracy, understanding and quantifying its uncertainty is essential.

Assessing uncertainty in a 4D model first requires identifying its potential sources. Uncertainty in 4DFWI arises from both the FWI algorithm and time-lapse–related factors. On the FWI side, uncertainties stem from characteristics and limitations of the acquisition survey geometry, field data noise, and the nonlinear nature of the inversion problem. Additional uncertainty arises from time-lapse effects (Zhou and Lumley, 2020), including discrepancies between baseline and monitor surveys that are not related to reservoir changes. These may involve differences in acquisition geometries, environmental conditions (e.g., tidal variations), receiver types and instrument responses, and background noise levels. Consequently, 4DFWI is subject to more sources of uncertainty than 3DFWI, making its quantification considerably more challenging.

Uncertainty in 3DFWI has been investigated using several approaches. It has been assessed in a limited manner by comparing the inverted model with well log information, or more qualitatively by visually examining flatness of common image gathers obtained through migration with the FWI model. More recently, 3DFWI has been formulated as a statistical inverse problem, with model uncertainty characterized by the posterior probability density function within a Bayesian framework. This is commonly achieved through variational inference (Zhang and Curtis, 2019; Lomas et al., 2023; Cen et al., 2024; Taufik and Alkhalifah, 2025) or through Markov Chain Monte Carlo (MCMC) sampling (Zhao and Sen, 2021). However, direct sampling of the posterior distribution for high-dimensional problems such as FWI is computationally prohibitive. As a more tractable alternative, Stein Variational Gradient Descent (SVGD), a variational inference method, has gained attention recently (Zhang and Curtis, 2019). Furthermore, because global posterior sampling remains computationally expensive, local uncertainty quantification methods such as the null-space shuttle (Keating and Innanen, 2021; Dai et al., 2024) have gained interest. These methods project carefully designed model perturbations into the null space of the forward modeling operator, assuming the FWI solution is sufficiently close to the global minimum. Given the significant challenges in quantifying uncertainty even for the case of 3DFWI, relatively little work has focused on estimating uncertainty in 4DFWI projects. In this study, we propose a new algorithm within the variational inference framework to quantify 4DFWI uncertainty. Specifically, we develop a 4D joint SVGD method and demonstrate that it provides a computationally affordable strategy for uncertainty quantification in 4DFWI.

We first review the theoretical background of 3D SVGD, followed by a detailed description of the proposed 4D joint SVGD algorithm. The method is then applied to synthetic 4D datasets based on a model representative of a typical marine geological setting characterized by thick salt bodies overlying a subsalt reservoir.

Method

In SVGD, the goal is to minimize the distance between two distributions: a distribution $q(m)$ represented by an artificially generated population of models, and the target posterior distribution $p(m \mid d_{\text{obs}})$of model parameters in 3D FWI. This distance is measured using the Kullback–Leibler (KL) divergence:

$$KL[q(m)||p(m|d_{obs})] = E[\log q(m)] - E[\log p(m|d_{obs})]. \quad (1)$$

Here, $p(m \mid d_{\text{obs}})$ is the posterior probability density of model m given the observed data d_{obs}, and $q(m)$ is a predefined distribution that is iteratively updated toward the target posterior by minimizing Eq. (1) using the SVGD algorithm (Liu and Wang, 2016). The distribution $q(m)$ is represented by a population of n models, referred to as particles.

The KL divergence is typically minimized iteratively using a local gradient method. The update direction at the i-th iteration is given by

$$\Phi_i(m) = \frac{1}{n}\sum_{j=1}^{n}\left[k(m_j^i, m)\nabla_{m_j^i} \log p(m_j^i|d_{obs}) + \nabla_{m_j^i} k(m_j^i, m)\right] \quad , \qquad (2)$$

where j is the particle index in the population of the n particles. The kernel function k is defined as

$$k(m_1, m_2) = e^{-\frac{\|m_1 - m_2\|^2}{h^2}} \qquad (3),$$

with scale factor h controlling the interaction strength between particles based on their separation.

Joint 4D SVGD

We propose a joint 4D SVGD algorithm designed to minimize the following objective function:

$$J(m_b, m_m) = KL_b + KL_m + \frac{1}{2}\alpha\|R(E(m_b) - E(m_m))\|^p \qquad (4),$$

where m_b and m_m are two populations of particles representing the baseline and monitor models, respectively. R is an optional 4D masking operator used to emphasize reservoir regions, α is a damping coefficient, and p is chosen as either 1 or 2. The first two terms correspond to the KL divergences for the baseline and monitor datasets (analogous to Eq. 1). The third term penalizes the difference between the expected models of $E(m_b)$ and $E(m_m)$, encouraging consistency between the two populations.

The two particle populations may be updated either simultaneously within each iteration or in an alternating fashion. Minimizing Eq. (4) yields two updated particle populations whose distributions approximate their respective posterior distributions, while also minimizing the difference between their expected models.

The 4D uncertainty is then computed as the standard deviation of the n^2 possible 4D model differences formed from the final particle populations of (m_b, m_m), assuming each population contains n particles (Figure 1).

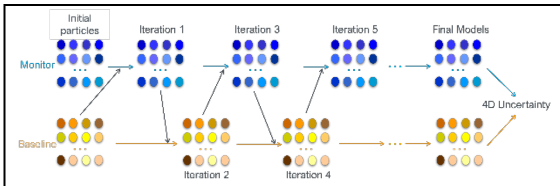

Figure 1: The flowchart in the 4D joint SVGD, showing how the two distinctive populations (blue and brown) of particles are updated in the 4D joint SVGD. The blue and brown arrows show the flow of monitor and baseline particles within each group, and the black ones show how one population interacts with the other.

From Eq. (2), it is evident that the 4D joint SVGD algorithm is at least twice as computationally demanding as standard 3D SVGD, and roughly $2n$ times more expensive than a conventional 3D FWI run. Strategies to mitigate this increased computational cost are available, and these will be discussed in the application section.

Numerical verification

4D uncertainty arises from many sources, but the dominant contributor is typically survey non-repeatability, particularly when the baseline and monitor surveys use different acquisition geometries, such as NAZ versus OBN. In such cases, the mismatch between survey types becomes the primary source of 4D uncertainty, whereas other non-repeatability factors (e.g., noise level differences) can generally be mitigated during preprocessing. It is therefore meaningful to first quantify the 4D uncertainty associated specifically with non-repeatable surveys. To do this, we apply the proposed 4D joint SVGD method to synthetic NAZ–OBN 4D datasets, which serve as a clean proxy for field data because they contain no noise and exclude other sources of non-repeatability.

Model and datasets

The baseline NAZ and monitor OBN datasets are simulated using true baseline and monitor models representative of a typical marine geological setting, featuring a large salt body beneath the seafloor and a subsalt reservoir (Figure 2, top). The water depth is approximately 300 m. In this test, we assume that the only 4D change occurs in V_p; no time-lapse variations are introduced in density or VTI anisotropy parameters. The bottom panels of Figure 2 show the true 4D model in inline and crossline directions, consisting of three anomalies with amplitudes within ±200 m/s.

The NAZ survey has a maximum inline offset of 8 km and a 700 m crossline cable span. The OBN survey also has a maximum offset of 8 km, with a nominal node spacing of 400 m. As noted earlier, 4D joint SVGD is computationally intensive. To reduce cost, we selected 80 OBN nodes located directly above the 4D anomalies. In contrast, the NAZ survey contains more than 50 times as many shots due to its much denser inline shot spacing (18.75 m inline with flip-flop shooting). This imbalance makes a direct application of 4D joint SVGD using the full NAZ dataset computationally infeasible and therefore requires additional strategies to reduce cost.

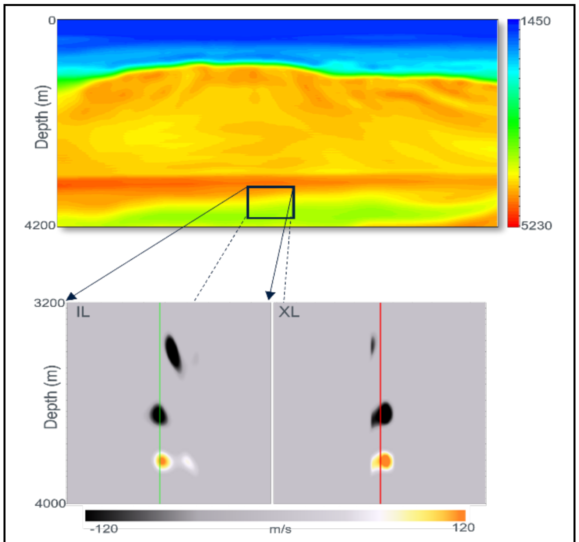

Figure 2: The inline vertical section of background Vp model (top) and the synthetic 4D model used in the 4D joint SVGD test, along inline (bottom left) and cross-line (bottom right) directions.

Two techniques—trace selection and gradient balance—have been developed for 4D FWI to approximately align the NAZ and OBN datasets in terms of their effective Hessians, with each being a function of respective source–receiver geometry, the background model, reflection opening angles, and illumination (Eq.(5)). Trace selection aligns the datasets in the pre-stack time domain by matching offset and azimuth ranges while accounting for elevation differences. Gradient balance further aligns the Hessians in the depth domain. Together, these techniques aim to satisfy the approximate relationship

$$f(x_s^{OBN}, x_r^{OBN}, v, \theta, illum_{OBN}) \approx f(x_s^{NAZ}, x_r^{NAZ}, v, \theta, illum_{NAZ}) \qquad (5)$$

Figure 3 shows that the 4D model from NAZ-OBN surveys with the two alignment techniques (Figure 3, middle) is fairly comparable to that from using OBN-OBN after traces are selected (Figure 3, bottom), in term of anomaly amplitude, position and shape in both inline and crossline directions, even though both are not as well reconstructed in term of amplitude, compared to that from using OBN-OBN datasets with no traces being discarded. That is reasonable since the OBN data without going through trace selection preserves full azimuth coverage, much better than that after trace selection to be aligned with NAZ data. Since the reconstructed 4D models between NAZ-OBN using the two aligning techniques and OBN-OBN after trace selection is fairly similar, we can now apply joint 4D SVGD using the selected OBN-OBN datasets to quantify the 4D uncertainty associated with the OBN-NAZ 4D surveys, which greatly reduces the computational intensity, making the application of 4D SVGD feasible.

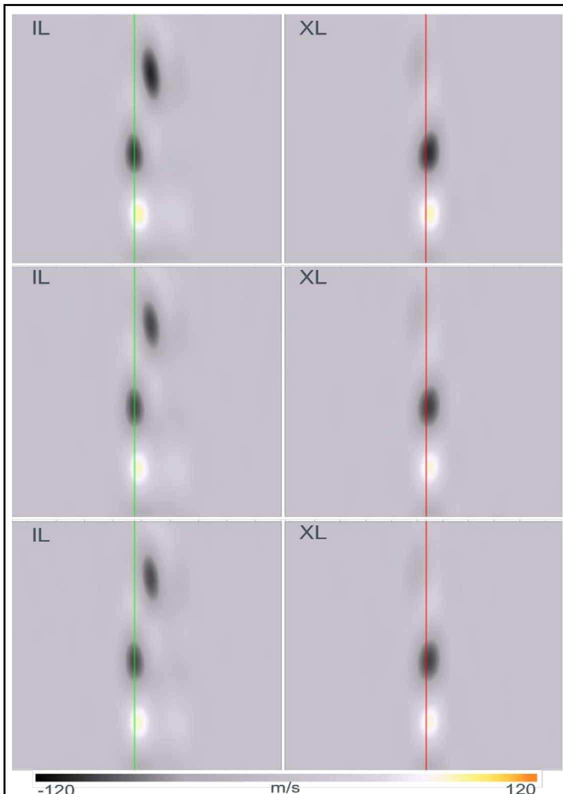

Fig 3: The reconstructed 4D model by 4DFWI: using OBN-OBN datasets as they are without trace selection(top), aligned NAZ-OBN datasets(middle) and OBN-OBN after alignment(bottom). The plots on left column show inline vertical section and on right the crossline vertical section.

Application of joint 4D SVGD

We apply the proposed 4D joint SVGD method to datasets in the low-frequency band (up to 6 Hz) to reduce computational cost. At these frequencies, the wavelengths are relatively large, so a population of twenty particles—with correspondingly large correlation lengths—is sufficient to capture the variability in the target subsalt reservoir zone below salt body. The particles for both the baseline and monitor models are generated independently using Gaussian random fields (GRFs) (Rasmussen and Williams, 2006), allowing us to incorporate prior covariance structure and geological dip information in a straightforward manner.

In this test, we use the dynamic matching objective function (Mao et al., 2020) to compute the FWI gradients required in Eq. (2). Figure 4 illustrates how data fitting evolves during the iterative particle updates for both the baseline and monitor populations, shown as histograms of DMFWI data matching. Initially, the matching values are widely scattered because the particles are randomly generated. As the iterations proceed, the data matching increases significantly and the distributions become increasingly concentrated around higher cross-correlation values. This behavior indicates that both particle populations progressively improve their data fit while still maintaining diversity—an essential condition for meaningful uncertainty quantification. This indicates that finally the two populations of particles explain both data well yet they are all different, therefore uncertainty can be derived.

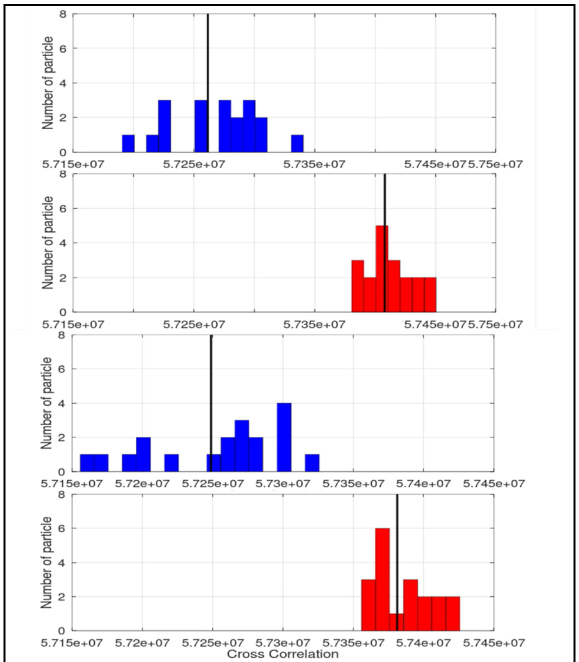

Figure 4: Evolution of data matching in 4D SVGD defined as summation of local cross correlation of synthetic and observed dataset, for baseline (top two plots) and monitor (bottom 2 plots). In each case, the blue histogram shows the distribution of data fitting using initial particles, and the red using final updated particles. In each plot, the black line is the mean data matching.

Figure 5 presents the resulting 4D uncertainty, computed as the standard deviation of the final particle populations. The left panel shows the uncertainty field within the reservoir, overlaid on an RTM image generated using the background model. Using this uncertainty estimate, we compute the confidence range along a vertical profile through the reservoir (Figure 5, right). In this example, the 4D uncertainty is less than 8.2 m/s—substantially smaller than the true 4D signal, which exceeds 100 m/s. This demonstrates that the proposed 4D joint SVGD method provides a reliable and quantitative measure of 4D uncertainty for this subsalt reservoir scenario.

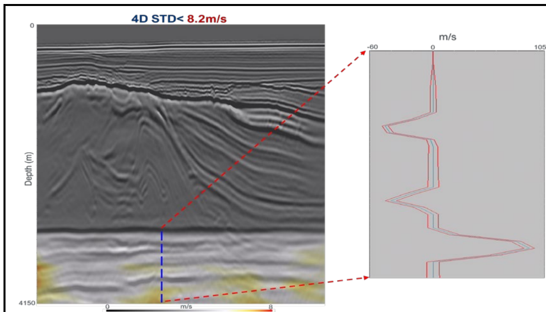

Figure 5: The final 4D standard deviation(STD) (left) and the confidence range centered at the mean 4D model(green curve, right), at a vertical profile (blue dashed line) using the STD.

Discussions and conclusions

We have introduced a joint 4D SVGD method for quantifying uncertainty in 4D full waveform inversion (4DFWI). The approach updates two populations of particles—one for the baseline model and one for the monitor model—so that each population progressively approximates its respective posterior distribution, while a coupling constraint ensures that the expected models remain consistent with one another as possible.

For 4DFWI involving hybrid NAZ–OBN acquisition, we proposed two techniques—trace selection and gradient balance—to align the effective Hessians of the two surveys. This alignment enables the use of joint 4D SVGD on a computationally manageable OBN–OBN equivalent dataset, while still capturing the uncertainty associated with the original NAZ–OBN configuration. The full workflow, including the alignment procedures and the joint 4D SVGD algorithm, has been validated through synthetic tests based on a realistic subsalt geological setting. These results demonstrate that the proposed method provides a practical and reliable strategy for quantifying 4D uncertainty in complex acquisition scenarios, including those involving non-repeatable survey geometries.

Acknowledgements

We thank Shuang Gao (from UT Austin) for her work on 3D SVGD at TGS, Matt Yates and his team for HPC support.

REFERENCES

Cen, H., K. Guo and D. Wang, 2024. FWI uncertainty analysis with Stein Variational gradient descent. 4th International Meeting for Applied Geoscience & Energy, 2024.

Dutta, G., E. Saragoussi, X. Cheng, Y. Zhai, C. Parekh and D. Vigh, 2023. EAFE Annual Meeting, 2023.

Gao, F., D. Cavalin, D. Davies, F. Liu, L. Saxton, C. Calderon, P. Tillotson and L. Wang, 2024. Time-lapse model by 4D dynamic matching joint full waveform inversion, SEG 4D Forum—Insights to Actions, Galveston, 2024.

Dai, W., Y. Yang, G. Kumar, S. Gerritsen, C. Perkins and V. Goh, 2024. Uncertainty quantification for visco-elastic multi-parameter full waveform inversion. 4th International Meeting for Applied Geoscience & Energy, 2024.

Keating, S. D. and K. A. Innanen, 2021. Null-space shuttles for targeted uncertainty analysis in full-waveform inversion. Geophysics, 86 (1), p. R63–R76.

Liu, Q. and D. Wang, 2016. Stein variational gradient descent: a general purpose Bayesian inference algorithm. Neural Information Processing Systems(NIPS), 29, 2016.

Liu, Y. and I. Tsvankin, 2022. Source-independent time-lapse full-waveform inversion for anisotropic media. Geophysics, 2022, Vo. 87, NO. 1, P. R111-122.

Lomas, A., S. Luo, M. Irakarama, R. Johnston, M. Vyas and X. Shen, 2023. 84th EAGE Annual Conference & Exhibition, 2023.

Mao, J., J. Sheng, Y. Huang, F. Hao and F. Liu, 2020, Multi-Channel dynamic matching full-waveform inversion, SEG Technical Program Expanded Abstracts.

Musa, M. B. Biondi and M. Meadows, 2015. Simultaneous TV-regularized time-lapse FWI with application to field data. SEG Annual Meeting, 2015.

Pintus, A., H. Ayadi and N. Salaun, 2023. Time-lapse FWI to improve understanding of superimposed reservoirs in deep offshore. EAGE Annual Meeting, 2023.

Rasmussen, C. E. and C. K. I. Williams, 2006. Gaussian Processes for Machine Learning. Appendix B.

Taufik, M. H. and T. Alkhalifah, 2025. Bayesian full waveform inversion using Stein variational gradient descent and manifold projections. Fifth International Meeting for Applied Geoscience & Energy, 2025.

Yang, D., M. Fehler, A. Malcolm, F. Liu and S. Morton, 2013. Double difference waveform inversion of 4D ocean bottom cable data: application to Valhall, North Sea, SEG Annual Meeting, 2013.

Zhang, X. and A. Curtis, 2019. Seismic tomography using variational inference methods. Journal of Geophysical Research: Solid Earth, 125.

Zhou, W. and D. Lumley, 2020. Non-repeatability on time-lapse 4D seismic full waveform inversion. SEG Annual Meeting, 2020.

Zhao, Z. and M. K. Sen, 2021. A gradient-based Markov chain Monte Carlo method for full-waveform inversion and uncertainty analysis. Geophysics, No. 86 and No. 1, p. R15-R30.