

Robustness of Sequential Viscoacoustic Full Waveform Inversion to Data and Source Wavelet Perturbations

Introduction

In seismic imaging, conventional techniques such as travel-time tomography and migration-based methods exploit only a subset of the recorded data, discarding amplitude and waveform-shape information that constrains the subsurface. Full Waveform Inversion (FWI) overcomes this limitation by iteratively minimizing the discrepancy between observed and synthetic waveforms, achieving a spatial resolution approaching the seismic wavelength (Virieux and Operto, 2009).

However, FWI remains challenging: it is non-linear, ill-posed, and many implementations assume a perfectly elastic subsurface when in reality the Earth's subsurface is intrinsically anelastic. Mechanical energy is converted into heat through internal friction and scattering mechanisms, leading to amplitude decay and waveform dispersion quantified by the seismic quality factor Q_p (low Q_p meaning strong attenuation). Ignoring attenuation in the forward model introduces systematic errors in the predicted waveforms, which the inversion attempts to compensate for through spurious velocity updates. This effect is particularly pronounced in geologically complex settings such as gas clouds, where strong lateral contrasts in both the P-wave velocity (v_p) and the P-wave quality factor (Q_p) are present.

Viscoacoustic FWI jointly accounts for velocity and attenuation, offering a more physically consistent framework. However, the joint inversion of velocity and attenuation is inherently coupled and suffers from severe multi-parameter crosstalk, as their respective kinematic and dynamic effects on the wavefield are difficult to separate uniquely (Kamei and Pratt, 2013). Beyond these theoretical challenges, practical performance depends critically on data quality and the accuracy of the assumed source wavelet. Understanding how such errors differentially affect velocity and attenuation recovery is therefore essential before viscoacoustic FWI can be reliably applied to field data.

In this work, we implement a sequential viscoacoustic FWI workflow using the Standard Linear Solid (SLS) rheological model as the forward engine and evaluate the recovery of velocity and attenuation under data and source errors, with particular emphasis on the relative sensitivity of v_p and Q_p .

Method

Viscoacoustic wave propagation is modeled using a partial differential equation system based on the SLS mechanism. This formulation augments the acoustic pressure equation with a memory variable $r(\mathbf{x}, t)$ to account for the continuous dissipation of seismic energy, allowing for an effective constant- Q_p approximation over the relevant seismic bandwidth (Bai et al., 2014). The attenuation magnitude is controlled by the relaxation time $\tau \approx 2/Q_p$ (for $Q_p \gg 1$). The inversion minimizes the objective function $J(\mathbf{m})$, defined for a model vector $\mathbf{m}$ as the L_2 -norm of the data residuals between the observed pressure wavefields (d_{obs}) and the synthetic data (d_{syn}) generated by the forward operator, summed over N_s shots and integrated over the recording time T :

$$J(\mathbf{m}) = \frac{1}{2} \sum_{s=1}^{N_s} \int_0^T \|d_s^{syn}(\mathbf{m}) - d_s^{obs}\|_2^2 dt \quad (1)$$

To compute the gradients of the objective function with respect to the P-wave velocity (v_p) and the quality factor (Q_p), we employ the adjoint-state method (Plessix, 2006). The Q_p gradient follows by differentiating the Lagrangian under the approximation $\tau \approx 2/Q_p$:

$$\frac{\partial J}{\partial Q_p}(\mathbf{x}) \propto \frac{2}{Q_p^2(\mathbf{x})} \sum_{s=1}^{N_s} \int_0^T \bar{r}_s(\mathbf{x}, t) \frac{\partial p_s}{\partial t}(\mathbf{x}, t) dt \quad (2)$$

where $\bar{r}_s(\mathbf{x}, t)$ is the adjoint memory variable and $p_s(\mathbf{x}, t)$ is the forward pressure wavefield, both for shot s , and the $2/Q_p^2$ scaling reflects the diminishing data sensitivity as the medium becomes less dissipative. The forward and adjoint wavefields are propagated with Devito (Louboutin et al., 2019). To mitigate velocity-attenuation crosstalk, a sequential strategy is employed: Stage 1 recovers v_p with Q_p fixed to its initial value; Stage 2 subsequently updates Q_p with the recovered v_p held fixed. Stage 1

uses the Adam optimizer, whose adaptive step length accommodates the variable gradient magnitudes across perturbation experiments without per-experiment retuning. Stage 2 uses stochastic gradient descent (SGD) instead, since its gradient-proportional updates avoid overshooting the characteristically weak Q_p gradients. Gradients are smoothed with a Gaussian kernel prior to each update, and the water layer is masked during Stage 2, since Q_p has no physical meaning in non-attenuating water. Recovery quality is quantified by the normalized data misfit (J_{norm}), defined as the squared L_2 -norm of the data residuals normalized by that of the observed data, computed per shot and averaged over all shots (see Table 1). It is further assessed by the v_p and Q_p cloud errors (E_{v_p}, E_{Q_p}), computed as the relative L_2 model error between the recovered and true v_p and Q_p , respectively, within the gas cloud region.

Results

Experiments are conducted on a modified subsection of the BP 2004 benchmark model (Billette and Brandsberg-Dahl, 2004) spanning 5.6 km horizontally and 2.2 km in depth on a 450×350 grid and containing a highly attenuative gas cloud anomaly ($v_p = 1.5\text{--}2.0$ km/s, $Q_p = 19\text{--}22$) embedded in shallow sediments (Figure 1). A 2D surface acquisition geometry is used, comprising 10 equidistant sources and 201 receivers per shot deployed at 5 m depth along the surface, recording the pressure component, with a Ricker source wavelet at peak frequency $f_0 = 5$ Hz used throughout the inversion without frequency continuation. The data are recorded for 3 s at a sampling interval of 0.765 ms. The initial models are obtained by smoothing the true models. Spatio-temporally correlated noise is generated by combining two independent Gaussian random fields with exponentially decaying covariance — one across receivers, one across time samples — additively combined and scaled to a target signal-to-noise ratio (SNR) of 2 dB. For the source phase experiments, shifts of 10° and 45° with respect to the correct wavelet are applied to generate the predicted data.




Figure 1 a) True v_p model. b) True Q_p model. c) Initial v_p model. d) Initial Q_p model. e) Close-up on the true v_p gas cloud area, f) Close-up on the true Q_p gas cloud area.

Experiment 1: Baseline results

Under noise-free conditions with a perfectly assumed source wavelet, the baseline v_p inversion recovers the large-scale structure well with the gas cloud anomaly clearly resolved (Figure 2a, Figure 3a). This baseline serves as the reference case for the perturbation experiments below. The Q_p inversion captures the broad attenuation structure with the gas cloud correctly localized, though spatially smoother than the true model (Figure 2b, Figure 3b) and the contrast between the gas cloud and the surrounding sediments is slightly underestimated.




Figure 2 a) Recovered v_p model. b) Recovered Q_p model

The partial Q_p recovery reflects the fundamentally weaker sensitivity of the wavefield to attenuation – a frequency-dependent amplitude decay difficult to separate from geometrical spreading and reflection losses. The persistent low-attenuation structure at depth (Figure 2b) is a fundamental limitation of amplitude-based attenuation inversion in low-attenuation regions, where the gradient signal is too weak to drive any meaningful update away from the initial model.




Figure 3 a) Close-up on the recovered v_p in the gas cloud area. b) Close-up on the recovered Q_p in the gas cloud area.

Experiment 2: Sensitivity to noise

The v_p recovery remains structurally correct under correlated noise, though the gas cloud reconstruction is noticeably rougher than the baseline (Figure 4a). The spatial and temporal coherence of the noise introduces structured artifacts that are difficult to suppress during the inversion.




Figure 4 Close-ups of the recovered models in the gas cloud area under correlated noise. a) v_p model. b) Q_p model.

The Q_p recovery is significantly degraded at SNR = 2 dB (Figure 4b) with the gas cloud contrast reduced relative to the baseline. This degradation reflects the coherent projection of the noise onto the gradient: because this projection does not average out under summation over sources and time samples, it is partially mapped into spurious model updates, with the weaker Q_p gradient affected far more than the velocity gradient.

Experiment 3: Sensitivity to source phase errors

The v_p recovery in the gas cloud remains relatively robust under source phase errors in both cases. At 10° , the recovered gas cloud closely matches the baseline (Figure 5a). At 45° , the gas cloud boundary is slightly blurred, and artifacts appear within the close-up region (Figure 5b), though the gas cloud structure is reasonably well retrieved.




Figure 5 Close-ups of the recovered v_p models (a, b) and recovered Q_p models (c, d) in the gas cloud area under source wavelet phase shifts of 10° and 45° , respectively, within each pair.

The Q_p recovery shows a clear sensitivity to the magnitude of the phase error. At 10° , a faint low- Q_p anomaly is discernible within the gas cloud region (Figure 5c). Both Stage 1 and Stage 2 J_{norm} confirm that small phase shifts have a limited impact on the inversion, while larger shifts (45°) are substantially more destructive (see Table 1), leaving the gas cloud anomaly essentially unrecovered (Figure 5d). In both cases, velocity errors introduced in Stage 1 propagate into the Q_p gradient through the inherited velocity model, but at 45° this degrades the attenuation recovery far more severely.

Table 1. v_p and Q_p gas cloud errors (E_{v_p} , E_{Q_p}) and J_{norm} for all experiments.

Experiment	E_{v_p} (%)	E_{Q_p} (%)	Stage 1 J_{norm} (%)	Stage 2 J_{norm} (%)
Baseline	1.71	8.30	1.67	0.03
Correlated noise	3.37	17.19	39.58	38.73
Phase 10°	3.28	16.36	2.90	1.24
Phase 45°	4.66	28.12	34.77	27.60

Conclusions

This work presents a sequential viscoacoustic FWI workflow and evaluates its robustness to correlated noise and source phase errors on the BP gas cloud benchmark. The sequential strategy recovers both v_p and Q_p under ideal conditions but Q_p recovery is largely limited to the strong, shallow gas cloud anomaly; the weaker low-attenuation structure at depth is left unrecovered — an intrinsic limitation of amplitude-based attenuation inversion, independent of data quality. Across all experiments, Q_p recovery is consistently more sensitive than v_p , owing to the weak data sensitivity to attenuation. Reliable Q_p recovery is therefore contingent on an accurate Stage 1 velocity model, since v_p errors propagate directly into the attenuation update. Correlated noise degrades the attenuation recovery while largely sparing the velocity structure, and small source phase errors are absorbed by the inversion, whereas the large ones degrade the Q_p recovery. Incorporating source wavelet estimation into the inversion would mitigate this sensitivity directly, while joint rather than sequential inversion of v_p and Q_p could remove the inter-stage error propagation.

References

- Bai, Jie, Yingst, David, Bloor, Robert, and Leveille, Jacques. Viscoacoustic waveform inversion of velocity structures in the time domain. *Geophysics*, 79(3): R103–R116, 2014. Doi: 10.1190/geo2013-0120.1.
- Billette, Frédéric J. and Brandsberg-Dahl, Sverre. The 2004 BP velocity benchmark. 67th EAGE Conference & Exhibition, Extended Abstracts, 2004.
- Kamei, Rie and Pratt, R. Gerhard. Inversion strategies for visco-acoustic waveform inversion. *Geophysical Journal International*, 194(2):859–884, 2013. Doi: 10.1093/gji/ggt109.
- Louboutin, M., Lange, M., Luporini, F., Kukreja, N., Witte, P. A., Herrmann, F. J., Velesko, P., and Gorman, G. J. Devito: an embedded domain-specific language for finite differences and geophysical exploration. *Computers & Geosciences*, 127:79–91, 2019. Doi: 10.1016/j.cageo.2019.02.008.
- Plessix, R.-E. A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophysical Journal International*, 167(2): 495–503, 2006. Doi: 10.1111/j.1365-246X.2006.02978.x.
- Virieux, Jean and Operto, Stéphane. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6): WCC1–WCC26, 2009. Doi: 10.1190/1.3238367.