

FWI assisted by a U-net predicted reliability map for synthetic degraded data

Introduction

Full-waveform inversion (FWI) can recover subsurface velocity models by minimizing the mismatch between observed and simulated seismic waveforms. Its performance, however, depends strongly on data quality, frequency content, source and receiver geometry, and the accuracy of the starting model. These requirements become more restrictive when the target includes thin layers, strong velocity contrasts or localized structural discontinuities, because small or sharp features may be poorly recovered when the residual is dominated by corrupted or unreliable data samples.

Synthetic FWI experiments are often first tested using clean data and benchmark models. This is useful for verifying the numerical workflow, but it does not fully represent the data conditions expected in more realistic applications. Missing traces, amplitude dropouts, random noise, irregular acquisition effects and limited repeatability can all distort the waveform residual. In such cases, standard least-squares FWI may update the model using samples that should have contributed less to the objective function.

This work tests a U-Net-assisted acoustic FWI strategy in which the neural network estimates reliability weights for degraded shot gathers. These weights are then inserted into the waveform misfit. The method is evaluated in a controlled synthetic setting, with one possible long-term application being seismic monitoring for carbon capture and storage. The focus of the present study is methodological: determining whether learned reliability weighting can improve FWI results when the observed data are intentionally degraded.

Methodological Approach

The workflow uses two-dimensional acoustic time-domain FWI as a controlled first approximation. This formulation does not include elastic effects, attenuation, anisotropy, statics, source uncertainty or the full complexity of land seismic data. It is used here because it provides a transparent framework for testing how data degradation and data weighting affect inversion behaviour before introducing additional physical and acquisition-related uncertainty.

A personalized synthetic velocity model was constructed instead of using only a generic benchmark. The objective was to create a model with interpretable targets and known sources of inversion difficulty. The model includes a high-velocity basaltic cap, a thinner shale-carbonate interval, an intrusive sill and a localized discontinuity, it was based on geological challenges commonly encountered on seismic studies in the Paraná Basin. These features are represented only as acoustic P-wave velocity contrasts. The model is not intended to be a calibrated geological or petrophysical reconstruction, but a controlled numerical analogue for testing thin targets, strong contrasts and intrusive bodies within the same inversion experiment.

Synthetic shot gathers were generated through the true model and recorded at surface receivers. A smooth initial model was used as the starting model, retaining the broad velocity trend while excluding the sharp interfaces, sill and localized discontinuity. The FWI implementation uses Deepwave, a PyTorch-based wave-propagation package that allows acoustic modelling and automatic differentiation to be combined in the inversion workflow [Richardson, 2026]. A multiscale strategy was used, with lower frequencies inverted first before progressively introducing higher frequencies [Virieux and Operto, 2009].

Clean synthetic gathers were degraded using controlled random noise, missing receiver traces and low-amplitude trace dropouts. A two-output U-Net was trained to estimate both a reconstructed data

patch and a reliability map from degraded inputs [Ronneberger et al., 2015]. In the FWI experiment, the reconstructed output is not used as the observed data. The degraded gather remains the data to be matched, while the U-Net-predicted reliability map weights the residual. Therefore, the neural network does not replace the wave-equation inversion and does not directly infer velocity. Its role is limited to estimating which parts of the degraded data should contribute more or less strongly to the misfit.

Two degraded-data inversions are compared. In the first, standard acoustic FWI is applied directly to the degraded gathers using an unweighted waveform misfit. In the second, the same degraded gathers are inverted using a reliability-weighted misfit based on U-Net-predicted weights. The true model, initial model, acquisition geometry, source wavelet, frequency schedule and optimizer settings are kept fixed, so the comparison isolates the effect of the learned weighting.

Results and Discussion

The true model contains the high-velocity cap, the shale-carbonate interval, the intrusive sill and the localized discontinuity. The initial model contains only the smooth background velocity structure, so the sharper target features must be recovered from the waveform information (Figure 1). The degraded shot gather shows the imposed data-quality problem, while the target and predicted reliability maps show the weighting information used to condition the FWI objective (Figure 2).




Figure 1 Four panels, first row, from left to right: True acoustic velocity model with a high-velocity cap, an Irati-like interval and an intrusive sill, Smooth initial model. Second row from left to right: Standard FWI using degraded data and an unweighted waveform misfit, FWI using the same degraded data weighted by U-Net-predicted reliability.

The standard degraded-data FWI result recovers the broad structure of the velocity model, showing that the multiscale acoustic workflow remains sensitive to the main background organization even when the data are degraded. The result, however, also contains stronger artifacts and less focused recovery around the thinner high-contrast target interval. This behaviour is consistent with the scale of

the target: the Irati-like layers and sill are sharper than the background model and depend more strongly on bandwidth, illumination and reliable waveform amplitudes.




Figure 2 Three panels from left to right : Representative degraded shot gather, Synthetic target reliability map, U-Net-predicted reliability map.




Figure 3 Error-reduction map, where positive values indicate lower absolute velocity error for the U-Net-weighted inversion than for standard FWI.

The U-Net reliability-weighted inversion shows reduced error in parts of the target region, as indicated by the improvement map. This result should not be interpreted as the network learning the geological model. The comparison instead indicates that the predicted reliability map changed how the inversion used degraded samples. Corrupted traces and low-reliability regions contributed less to the objective function, allowing the wave-equation inversion to update the model from a more selective residual.

This strategy has two practical advantages in the present experiment. First, the degraded gathers remain the observed data, so the inversion is not forced to fit a fully network-generated seismic record. Second, the reliability map is an intermediate quantity that can be inspected and compared with a known synthetic target. This makes the U-Net contribution more transparent than a direct seismic-to-velocity prediction or an unconstrained denoising replacement.

The method also has limitations. The predicted weights are meaningful only to the extent that the training degradations represent the degradations used in the inversion. If the U-Net is trained on unrealistic noise or dropout patterns, the weighting may suppress useful arrivals or preserve misleading ones. The reliability map is also not a complete uncertainty estimate; it is a supervised synthetic weighting proxy. Finally, the present experiment is acoustic and synthetic, so it does not yet test elastic amplitudes, converted waves, attenuation, anisotropy, source uncertainty, statics or acquisition non-repeatability.

Conclusions and Implications

This work presents a U-Net reliability-weighted acoustic FWI workflow for degraded synthetic seismic data. The experiment tests whether learned data weighting can improve inversion behaviour when the observed gathers include controlled noise, missing traces and amplitude dropouts.

The results show that standard FWI can recover the broad velocity structure from degraded data, but thinner high-contrast targets remain more affected by artifacts and amplitude errors. When the same degraded data are inverted using U-Net-predicted reliability weights, parts of the target region show reduced model error. The main methodological result is that a convolutional network can be used as an auxiliary data-weighting module while keeping wave-equation inversion as the model-updating mechanism.

The workflow does not use the neural network to infer velocity directly and does not replace the degraded observations with denoised data. Instead, the network modifies the weighting of the waveform residual. This provides a conservative way to couple machine learning with FWI while preserving an interpretable connection between the observed data, the residual and the velocity update.

Future work will test the sensitivity of the method to different degradation recipes, reliability-map calibration, source uncertainty, elastic modelling and more realistic acquisition effects. Additional synthetic models will also be used to evaluate whether the reliability-weighted strategy remains stable across different geological structures, acquisition geometries and levels of target complexity.

Acknowledgements

This study was financed, in part, by the São Paulo Research Foundation (FAPESP), Brazil. Process Number #2025/21803-1. The work was developed as part of an undergraduate research/thesis project at the Institute of Astronomy, Geophysics and Atmospheric Sciences, University of São Paulo (IAG-USP).

References

- Kendall, A., Gal, Y. and Cipolla, R. [2018] Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7482–7491.
- Richardson, A. [2026] Deepwave. Zenodo. DOI: 10.5281/zenodo.3829886.
- Ronneberger, O., Fischer, P. and Brox, T. [2015] U-Net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention*, 234–241.
- Virieux, J. and Operto, S. [2009] An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6), WCC1–WCC26.