Self-correcting salt-model building from a highly inaccurate start model

Mike Warner, Christos Mavropoulos, Tenice Nangoo, Adrian Umpleby, Nikhil Shah: S-Cube, London
Dmitriy Tishechkin: Amazon Web Services

Summary

We demonstrate the feasibility of building an accurate velocity model, containing complex salt bodies, using only raw surface seismic data, advanced waveform inversion, and a start model that contains no salt and for which the assumed sedimentary velocities are systematically in error. We use ocean-bottom node data without exceptionally long offsets or unusually low frequencies, and we do not use reflection tomography, conventional migration, salt-boundary picking, explicit salt flooding or scenario testing. We are able to recover a model automatically that is sufficiently accurate for elastic conventional FWI. To achieve this, we use a self-correcting multi-stage inversion scheme, employing different implementations of adaptive waveform inversion as it proceeds. The earliest stages use an accelerated formulation, combined with simple constraints, to achieve the large velocity changes necessary using minimal compute. Use of Graviton 4 spot instances on the cloud further reduces cost. The final 10-Hz acoustic velocity model is both accurate and well resolved; it generates structurally simple FWI impedance reflectivity images beneath the salt and is ready for immediate multi-parameter elastic FWI run up to the highest frequencies in the dataset.

Introduction

Elastic, broadband, multi-parameter FWI has become an important tool for building high-resolution velocity and impedance models for salt-affected seismic datasets. However, the success of high-resolution elastic FWI is predicated on the accuracy of an initial velocity model which should contain both an accurate model of salt geometry, and an equivalently accurate model of acoustic velocity in the post and pre-salt sedimentary sequence, and within the boundaries of the salt itself. Building such a model using conventional workflows and least-squares FWI, that are sufficiently accurate to begin high-frequency elastic FWI, can be both expensive and error prone.

We have previously demonstrated (Warner et al., 2023), that we can build accurate salt models starting from a reasonably accurate one-dimensional sedimentary velocity trend. Here, we demonstrate that we can relax even that requirement, and can begin FWI successfully from a sedimentary section that is itself severely in error. To do this, we used a one-dimensional sedimentary compaction trend that was suspected to be systematically in error and that had likely local errors in sedimentary velocities in excess of 1500 m/s, increasing to 2500 m/s where there was salt. We also began from an isotropic velocity model to avoid introducing spurious anisotropy within the yet-to-be-determined salt bodies. For the lowest frequencies available in the field data, this starting model was cycle-skipped almost everywhere, often by several cycles.

We used unprocessed hydrophone-only data from a modern ocean-bottom node survey covering approximately 600 km², located in the Gulf of Mexico in a water depth of about 2 km. The node spacing was 400 m in the inline and crossline directions; the source interval was 50 m in both directions. Maximum offsets were about 30 km, but were only available for a few midpoints in the center of the survey area utilizing sources and receivers from opposite corners. For the great majority of midpoints and azimuths, maximum offsets were around or below 20 km. The source was a conventional airgun array.

Method

Figure 1 shows a schematic flow-diagram that outlines the process we used to convert a simple, highly inaccurate, 1D, isotropic starting model, through into a heterogeneous, accurate, 3D anisotropic model ready for subsequent high-resolution multi-parameter elastic inversion. The boxes shown with dashed lines are external inputs or outputs; the solid boxes are the principal stages through which the model passes, and the linking arrows are the FWI computational engines that turn one model into another. The workflow is automated and pre-defined; high-level QC-driven decisions are used to take the model around the central iterative loop that serves to correct the trend of the initially assumed sedimentary trend.

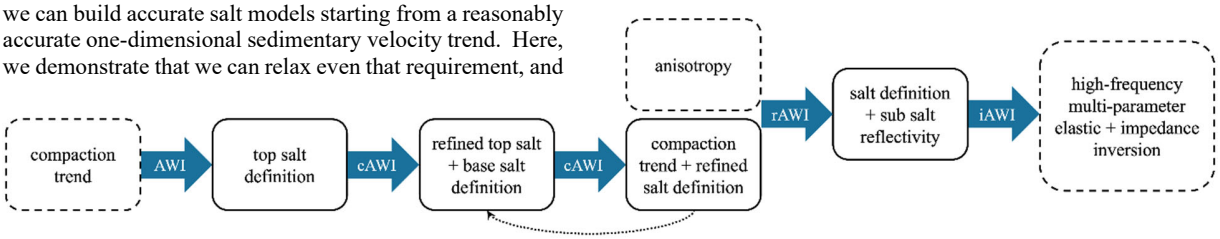

Figure 1: The inversion workflow: AWI (adaptive waveform inversion), cAWI (constrained), rAWI (reflection), and iAWI (impedance).

For waveform inversion to solve this problem efficiently and effectively, the workflow should meet two needs. It must: (1) make progress towards the true model despite the confounding effects of cycle skipping, and (2) take large steps when the model is far from the true answer. The former is required in order to proceed at all, and the latter is required in order to make that progress computationally affordable. The workflow is underpinned by adaptive waveform inversion (AWI), a family of FWI algorithms that do not have minima in their misfit functions associated with cycle skipping (Warner & Guasch, 2017; Guasch et al., 2020). Rather than using sample-by-sample differences to define the misfit between observed and model-predicted data, AWI derives matching filters that turn each trace in one dataset into its equivalent in the other. It then defines a measure of misfit that acts to drive those filters towards zero-lag band-limited delta functions.

There are many ways to define a misfit function for AWI. In this workflow, we used a more-aggressive accelerated misfit function in the early stages, relaxing this as inversion proceeded. In the early stages, AWI seeks to take large steps in broadly the right direction, but it is not guaranteed to get the finer details correct, and its corrections to the model may overshoot and are ultimately limited only by bounding constraints for the whole model. The principal role of these early AWI iterations is to push background velocities in the sedimentary section in the right direction, and to build a smooth salt model with the correct long-wavelength geometry and velocity. In the later stages, AWI employs a more-benign misfit function. This version is then unable to make such large changes to the macro model, but it does seek to get the finer details correct, and it acts to remove spurious local structure and to repair regions where the magnitude of earlier velocity corrections has been too large. To minimize computational cost, we run on ARM rather than X84 or GPU hardware, use the spot market on the cloud, and actively seek the lowest-cost instances across regions, using a formulation that is robust against preemptive loss of compute nodes.

In addition to this self-correcting evolution of model and AWI algorithm, the workflow has three other characteristics that promote rapid early progress. (1) We update the model consistently further than a conventional FWI step-length calculation would require. (2) We apply constraints and penalties to the evolving model that first encourage the rapid appearance and evolution of smooth salt bodies, followed by a restoration of the original sedimentary velocity trend below a smooth base of salt. (3) Part way through the workflow, we use the spatial variation of several independent measures of data misfit to indicate where in the model the original sedimentary velocity trend was in error. We use the phase misfit for turning rays and wide-angle reflections, to determine the smooth corrections that should be made to that trend; VTI anisotropy is also introduced at this stage.

Results

Figure 2 shows the evolution of the model. The model top is 1.5 km below sea level; its base is at 8.5 km depth. Crossline sections, on the left, extend for 24 km; in-line sections, on the right, extend for 20 km. The color scale clips at velocities of 1500 and 5000 m/s. Figure 2a shows the start model composed of an isotropic 1D sedimentary profile beneath an accurate seabed and water column. Figure 2b shows the results of early iterations of AWI without the application of constraints. This builds a thin high-velocity layer that represents the topmost portion of a salt body, and it begins to make changes in the shallow sedimentary section that are broadly correct, and deeper changes that are unlikely to be accurate at this early stage.

Figure 2c shows the model after the application of constrained AWI (cAWI) followed by the relaxation of those constraints. There are two constraints: (1) that the model should be smooth, and (2) that velocity should not decrease with increasing depth. This second constraint acts in a manner that is broadly similar to salt flooding, but it does not require the identification of the top-salt boundary, nor the specification of a salt-flood velocity. Rather, the data demands high velocities at the top of the salt body, and the constraints only allow those velocities to appear if they extend to the base of the model. As the constraints are relaxed, the model roughens, and velocities begin to reduce at around the position where base salt will subsequently appear. Note that the accelerated AWI misfit function used at this stage has overestimated the velocity within the salt.

Figure 2d shows the model after a second pass of constrained AWI without subsequent relaxation of those constraints. As before, the model is constrained to be smooth. However, the second constraint now requires that the model moves back towards the original compaction trend wherever that does not produce a significant misfit to the data. Algorithmically, this involves both a constraint on the model and a penalty on the data. The combined effect of constrained AWI without relaxation here is to produce a model composed of a smooth salt body superimposed upon a smooth sedimentary trend. The model in Figure 2d has also passed around the iterative loop in Figure 1. This has refined the original sedimentary trend, reducing the velocity in the shallowest sediments while significantly increasing sedimentary velocities deeper in the section. In keeping with our accelerated self-correcting strategy, this adjustment is large and approximate.

Figure 2e shows the model after the addition of anisotropy, and further passes of unconstrained AWI using both refracted and reflected energy (rAWI). The model shows the result immediately after the last iteration of the accelerated formulation of AWI. The macro model at this stage is accurate, but some of its intermediate-scale detail is not.

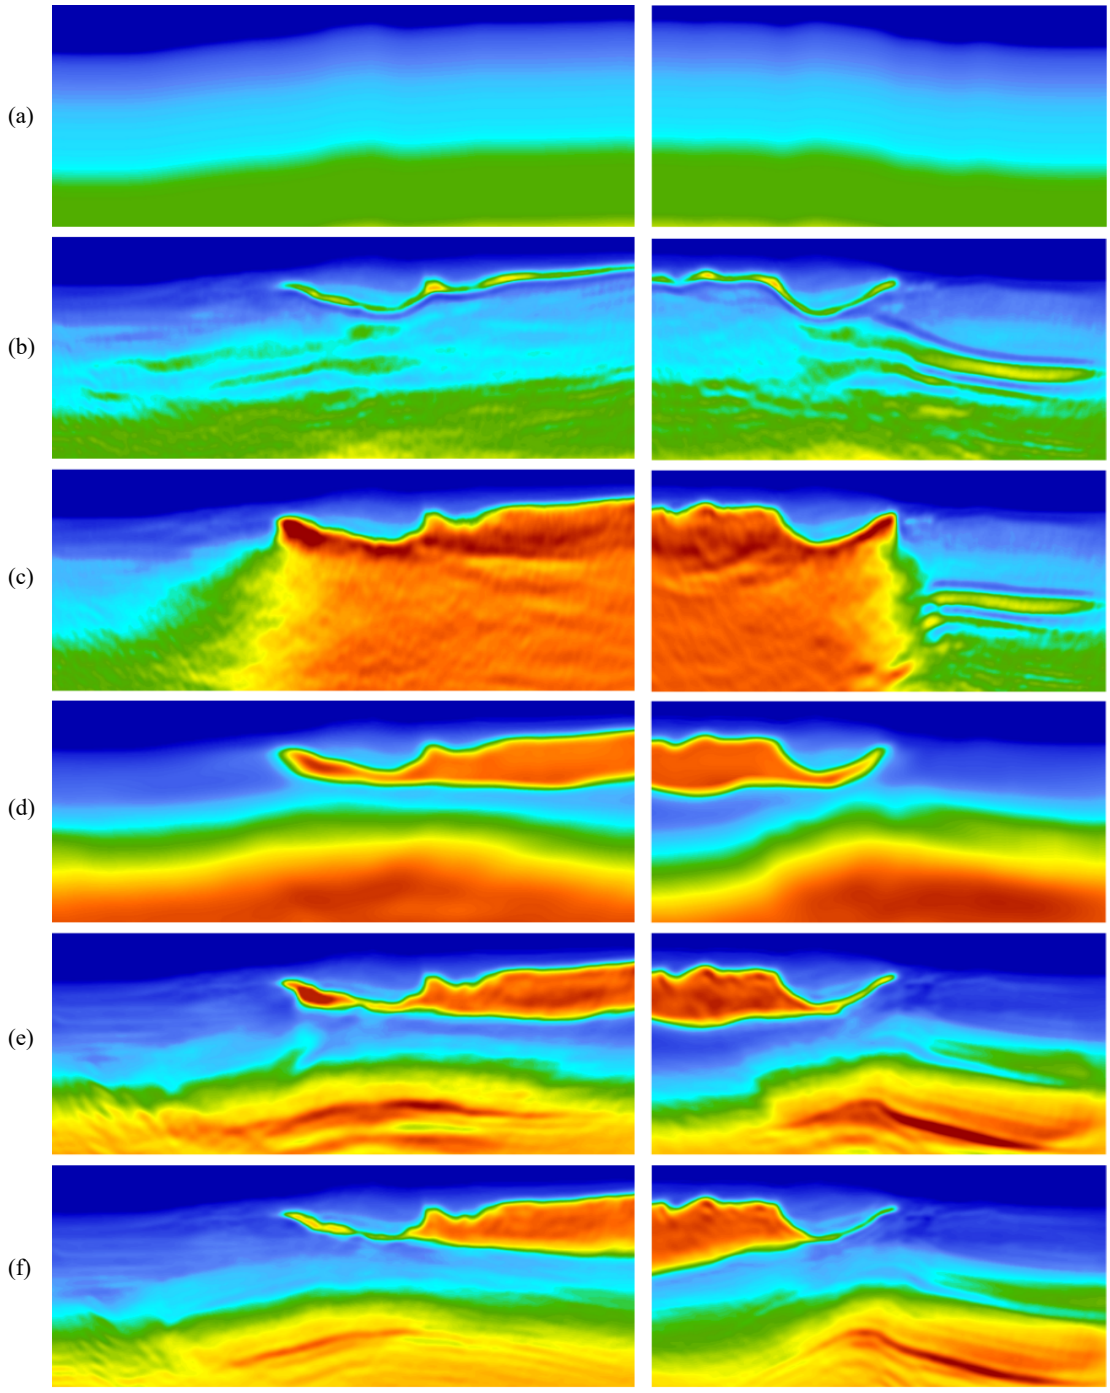

Figure 2: Model evolution. (a) Starting compaction model. (b) Unconstrained AWI builds top salt. (c) Constrained AWI first floods the model, then relaxes and begins to identify base salt. (d) Constrained AWI smooths the model and reimposes the compaction trend where the data allow. (e) Adding anisotropy and iteratively updating the compaction trend using an accelerated formulation of AWI builds sub-salt structure and sharpens the salt. (f) Finally, a more-benign formulation of AWI corrects artefacts and sharpens the model in preparation for high-frequency elastic FWI.

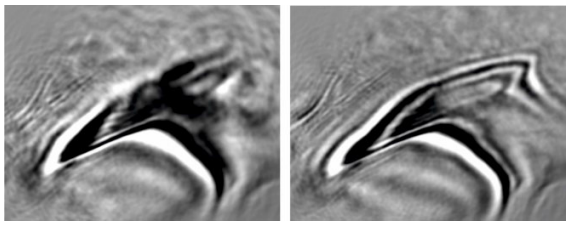

Figure 3: Horizontal reflectivity slices derived from models 2e (left) and 2f (right). Salt overlies the upper-right of these images.

Figure 2f shows the model after correction using the less-accelerated more-accurate AWI formulation. This acts to correct spurious structures that earlier, accelerated stages have introduced; it also increases the spatial resolution. Comparison of models 2e and 2f reveals many changes of detail. On the crossline in 2e, the extended nose of the salt is bulbous with an unusually high velocity; this is underlain by an upraised sedimentary feature. The inline for model 2e shows a step in the sedimentary section below the salt. These unusual features all disappear in model 2f. Figure 3 shows a horizontal slice through reflectivity volumes derived directly from models 2e and 2f. The most obvious differences coincide to the locations of the salt. This deep reflector is clearly better imaged by model 2f than 2e.

The upper panel in Figure 4 shows a 10-Hz impedance image for the final model 2f. This image was generated using a single iteration, inverting over a restricted offset range to mitigate elastic effects, and configured to update for acoustic impedance assuming no change in velocity (iAWI). Long wavelengths have been removed from the update using a Laplacian filter. This procedure generates a simple low-cost bandlimited RTM-like image directly from the unprocessed field data; it can serve as a rapid QC throughout the work-flow, and variations on that procedure can generate analogous FWI-derived CIGs. The lower panel in Figure 4 shows an image of acoustic reflectivity generated by differ-entiating the product of the velocity (and corresponding density) models from Figure 2f; the differentiation is performed perpendicular to the local velocity gradient.

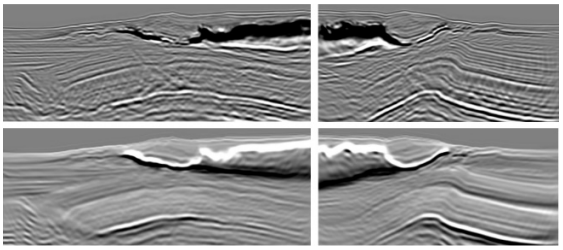

Figure 4: Upper panel: A 10-Hz, fixed-velocity, limited-offset, single-iteration, short-wavelength, FWI impedance image. Lower panel: Reflectivity model generated by differentiating Figure 2f.

The two panels show distinct physical properties; their derivation, purpose, effective wavelet, resolution, bandwidth and noise characteristics all differ. The acoustic velocity model in Figure 2f is now sufficiently accurate to be passed onward for higher-resolution refinement using conventional least-squared FWI employing any combination of higher-frequency, elastic, multi-parameter, or impedance inversion.

Conclusions

Elastic inversion is computationally expensive, especially at higher frequencies and for multi-parameter inversion. However, provided that elastic inversion can begin from a high-quality starting model, the total elapsed time and true total cost may not be especially large. In contrast, existing workflows for building velocity models, for salt-affected datasets, that are of sufficiently high quality for immediate high-frequency elastic inversion, can take significant elapsed time for data processing, tomography, horizon picking, scenario testing, and interpretation, and require repeated and costly manual intervention and expertise.

We have demonstrated that an automated, parsimonious, AWI-driven workflow can build a high-quality velocity model containing salt, for subsequent elastic inversion, beginning from a model in which the only real features are the seabed and water column. In its initial stages, this approach employs simple constraints on the model, together with an accelerated misfit function for AWI, in order to encourage the inversion to take large early steps. This configuration is highly cost effective, but is also liable to overshoot the true model leading to localised artefacts in the model. Subsequent iteration of AWI, employing a non-accelerated formulation, is however able to remove such features, leading to an accurate ,artefact-free, final model.

For many salt-affected datasets, a legacy velocity model, perhaps built at significant cost, already exists. Heavily smoothing such legacy models to remove all short wave-length features except the seabed, can be used to generate a model analogous to that shown in Figure 2d. The workflow described here can equally begin from such a starting point, circumventing the need for the earlier constrained iterations. While the aim of the workflow here is to build a model for subsequent elastic refinement at higher frequency, the acoustic models that it generates are often sufficient as a basis for accurate RTM imaging. Unlike pure FWI, AWI-derived acoustic models do not typically show a significant spurious halo around salt bodies; the sharpness of the salt model here is consistent with its derivation from 10-Hz data.

Acknowledgements

We thank Shell E&P for permission to publish, and thank Max Liu, Adriano Gomes and Yi Yang for their support.