

Chapter 1

Introduction

PROBLEM OVERVIEW

Why is salt hard to model?

Oil producing regions like the Gulf of Mexico are known to have geologically complex salt body formations. These formations can act as seals, trapping hydrocarbons underneath, which makes them frequent targets of seismic imaging projects. However, salt bodies have two important properties that make them difficult to image. First, the P-wave velocity of salt is high, and contrasts sharply with that of surrounding sediment layers, making them very reflective to seismic energy. Second, salt bodies can adopt a variety of complex shapes and geometries. Because of these factors, the formations themselves act as lenses which focus or disperse seismic energy (Leveille et al., 2011; Etgen et al., 2016; Barnier and Biondi, 2015). Because salt bodies often have steep dipping boundaries, useful reflection energy may be directed along ray-paths that reach the surface far outside of the acquisition geometry extent (see Figure 1.1(b)). With less complex imaging subjects (Figure 1.1(a)) this is not so often the case. Lack of data capture can subsequently impact the imaging results by creating ‘blind’ spots near the base and flanks of the salt. Figure 1.2(a) demonstrates how flat portions of the salt boundary have good illumination, while the steeply dipping flank (Figure 1.2(b)) has poor illumination, making picking the salt boundary more difficult.

However, even when reflected energy is sufficiently recorded, strong reflectivity




Figure 1.1: a) Ray path from flat reflector. b) Ray path from salt reflector. Note that after passing through salt (black) the raypath in 1.1(b) reaches the surface outside of the acquisition zone. Receivers (green) and source (yellow) are shown on the surface.

[NR] `chapter1/. ray-path-no-salt,ray-path-with-salt`

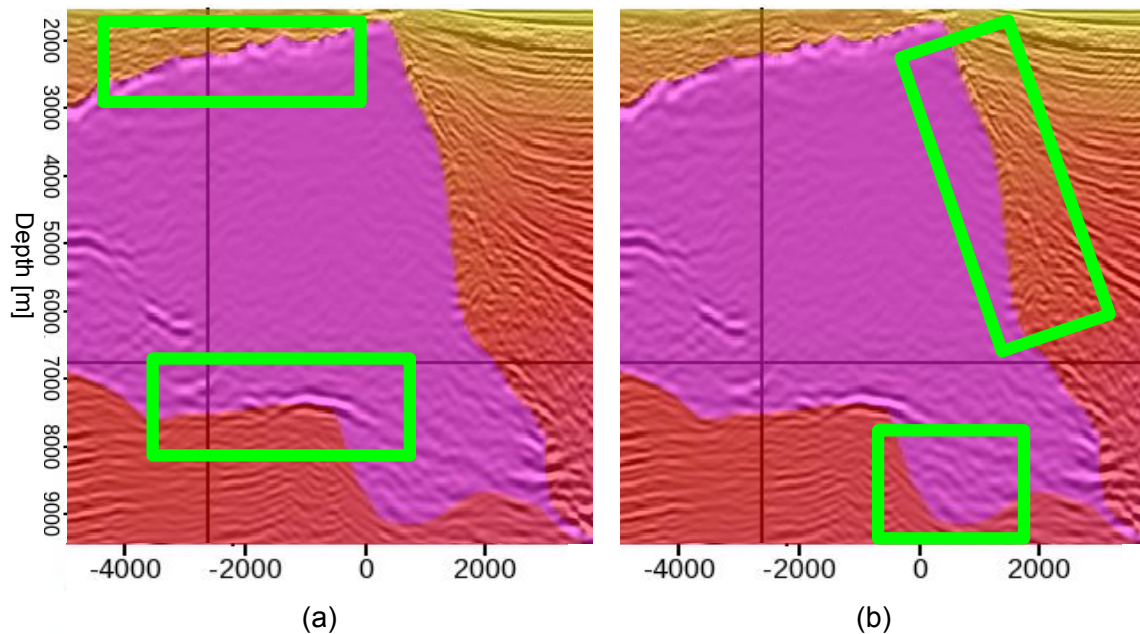


Figure 1.2: Salt model (pink) overlain on seismic image. Green boxes in a) highlight areas of good illumination, while boxes in b) highlight areas of poorer illumination.

[NR] `chapter1/. salt-picking-examp`

and lensing effects can amplify small errors in the salt model and negatively impact the resulting image. The boundaries of salt are typically improved in an iterative workflow (interpret salt from image, update salt model, then re-migrate image as in Figure 1.3). Subsequently, small initial errors may be compounded as velocity model development proceeds. Furthermore, these methods are time-consuming and tedious since expert input is necessary for either the actual picking, or the oversight and correction when picking is semi-automated. Each iteration of this approach can take days or weeks while coordinating entire teams of people. A robust method for further automating the salt model building procedure would greatly alleviate this bottleneck.

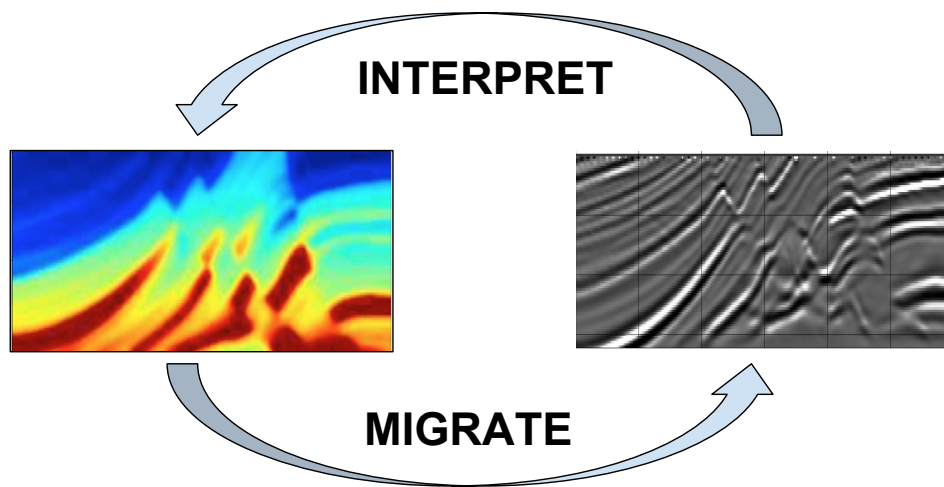


Figure 1.3: Basic velocity model building workflow concept. [NR]
[chapter1/. vel-building-workflow](#)

Approaches like full-waveform inversion (FWI) are intended to avoid the repeated manual interpretation necessary in the approach just described, but can be less than effective for a variety of reasons. Sometimes high frequencies in the data are of poor quality, causing difficulty reconstructing sharp velocity models. When high-frequencies are present with sufficient quality, multi-scale inversion workflows are typically used, which build a model starting from low and working up to high frequencies Bunks et al. (1995). However, to increase the sharpness of the model requires higher and higher frequencies, which become increasingly expensive to compute wave propagations for. To alleviate this expense, multi-scale inversion is sometimes done using wider frequency band increments, or even using a full band of frequencies from the beginning.

Unfortunately, using high frequencies too early in the inversion runs a higher

risk of converging to a local minima. This is because the FWI objective function is based on the data residual, and the size of the ‘well of convergence’ around the true model is based on the wavelengths in the data (see Figure 1.4(e), pink). The ‘well of convergence’ is the model-space region surrounding the true model where the objective function is convex. When a starting model is within this region, the objective function gradient will eventually direct convergence to the true model. When a shift in the modeled data is greater than a half wavelength from the true data, the objective function gradient will direct updating to a local minima, pushing the data to align with a wave cycle other than the one representing the true model (i.e. skipping a wave cycle). With low frequencies the wavelength is longer and the well of convergence is bigger (Figure 1.4(e)), and thus greater model error can be tolerated Virieux and Operto (2009). However, even using advanced FWI tools with starting models and frequencies that thoughtfully consider cycle-skipping, the final results typically lack sharp resolution around the salt edge (compare Figures 1.5(a), 1.5(b)), which can impact the clarity of seismic imaging near the salt.

One approach that is used to create sharp boundaries in velocity models is Total-Variation (TV) regularization, which makes use of the L1 norm to regularize the model parameters during the inversion (Maharramov and Levin (2015), Esser et al. (2018)). This creates a regularized version of the FWI objective function that includes an additional term, $\lambda|Am|_1$, where A is a Laplacian operator that takes the spatial derivative of the model space. There are several challenges with this approach. The first is that a parameter λ must be chosen which appropriately balances the regularization term with the data residual term. This is can be difficult to find, and may even need to change with iteration for the best inversion result. The second problem is that taking the derivative of this new term is difficult. This is because the L1 norm is not a differentiable function at zero. This means that an optimization method must be chosen for a non-smooth problem, such as the subgradient method (Maharramov, 2016), which can be much slower than Newton’s method on a smooth objective function.

Why are level sets the answer?

A key tool for addressing this problem is the concept of the level set. A level set is a contour of a higher dimensional surface used to track a discrete boundary. In the salt

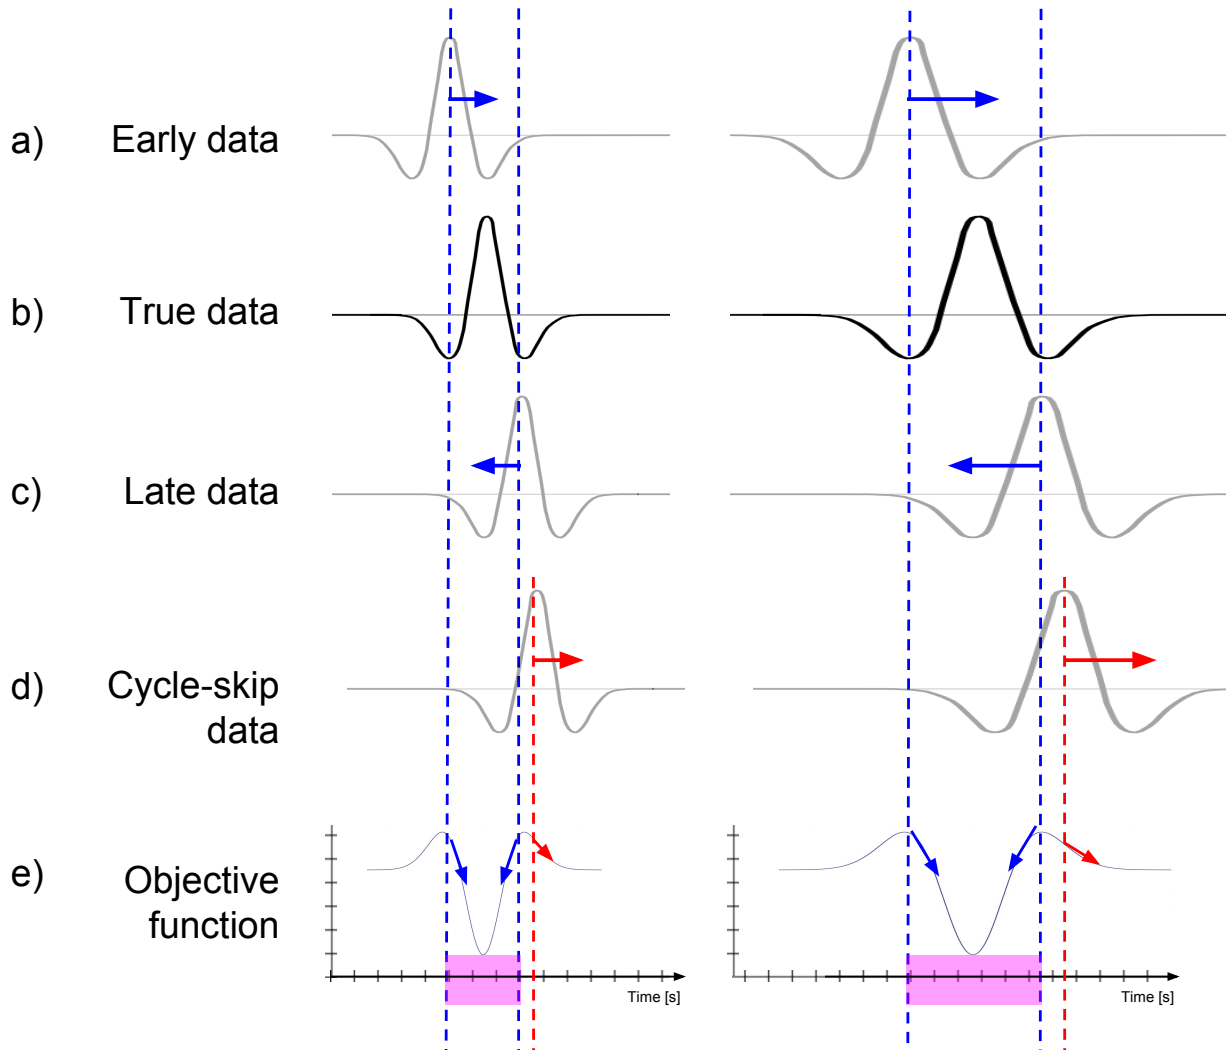


Figure 1.4: Left column shows higher frequency data while right column shows lower frequency data. True data (b) and shifted traces (a,c), with the objective function (e) showing the data residual norm measured at all shift positions. When a trace is shifted far enough from the true trace (d), the gradient of the objective function directs updating to a local minima, pushing the trace *further* from the true position (red arrows), and cycle skipping occurs. Blue arrows indicate convergence towards the true solution when the modeled trace shift (and thus model error) is within the well of convergence (pink). [NR] `chapter1/. cycle-skipping`





Figure 1.5: a) Legacy velocity model. b) Velocity model after FWI updating Shen et al. (2017). [NR] `chapter1/. xukai-true,xukai-fwi`


modeling problem, the salt boundary can be represented using a level set as in Figures 1.6(a) and 1.6(b). With this representation, one can easily modify the shape and topology of the level set *implicitly* by modifying its higher dimensional surface (for this reason, called an implicit surface). While explicit boundary tracking schemes break down for sharp edges or topology changes, level sets very elegantly manage these cases. For any problem where we wish to track the boundary of a homogeneous (or approximately homogeneous) body, level sets are an ideal tool. They have found use in fields like aerospace engineering for modeling fluid flow around airfoils (Xia et al., 2010), as well as image segmentation problems in medicine (Maciejewski et al. (2012), Tsai et al. (2003)). For the purposes of salt modeling, they give us the advantage of adjusting and updating a sharp boundary position (high spatial frequency content), even when the updates from our data contain only low frequency content (see Figure 1.7). Since cycle skipping is based on the spectra of the data and not the velocity model, we can gain the advantage of manipulating high spatial frequency boundaries while still mitigating the risk of cycle skipping by skewing the spectra of our data space to low frequencies. Furthermore, computing wave propagations for low frequencies is cheaper than propagating high frequencies.

What has already been done?

Level sets are an excellent tool for tracking boundaries, and their debut was ushered in by the work of Osher and Sethian (1988). They are especially useful for the class of problems called *shape optimization* that optimize an objective function based on shape properties. Santosa (1996) provides an early demonstration of solving these



(a)



(b)

Figure 1.6: Diagram showing a cross section (orange curve) through a 3D implicit surface that represents the salt boundary in the 2D model (bottom panel) for both a) circular salt and b) donut-hole models . [NR]
 chapter1/. levelset-cross-section,levelset-cross-section-donut

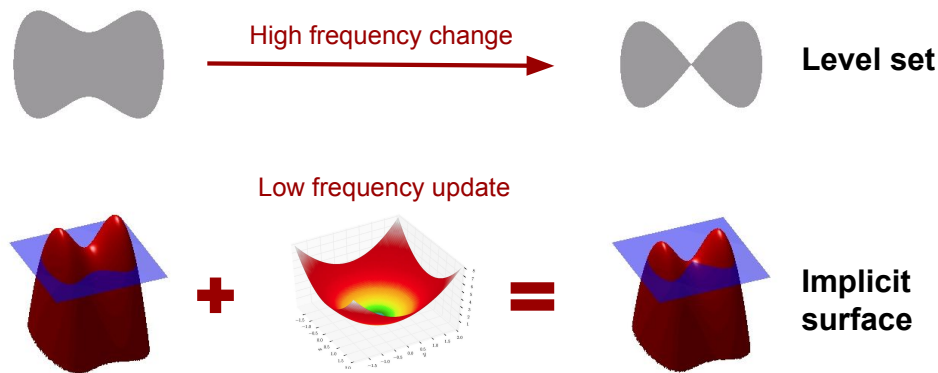


Figure 1.7: Diagram demonstrating how low spatial frequency updating leads to a change in a high spatial frequency boundary represented by a level set contour. [NR] `chapter1/. levelset-update`

types of problems with an L2 norm objective function, namely for deconvolution and diffraction-screen reconstruction problems. Burger (2003) continued to build a more unified theory of level set methods for inverse problems. Later work by Lewis et al. (2012) and Guo and de Hoop (2013) applied shape optimization to the FWI objective function for the purpose of segmenting salt bodies. Guo and de Hoop (2013) utilize this approach with frequency domain wave propagation to evolve a salt boundary and velocity model. However, these approaches use the full model domain to represent the level set, and do not explore the advantages of using second-order updating. More recently, Kadu et al. (2016) introduced the use of radial basis functions to sparsify the model space and speed computation for second order updating, but did not extend the work to 3D models or field data examples.

What does this thesis add?

In this work, I claim to demonstrate success using shape optimization with level sets to build 3D models of salt bodies at depth using field data. The differentiating aspects of this work include a number of advances. First, the use of radial basis functions (RBFs) to reduce the number of model parameters has been improved by distributing them with spatial variance according to a given likelihood of salt updating. I leverage the sparsified model that this representation yields to to speed our inversion of the Newton step that our Gauss-Newton Hessian approximates. I

further leverage interpreter guidance as input by molding the initial shape of our implicit surface according to the same spatial likelihood function that our RBF centers are based on. Last, I further extend the practice by applying this method to a 3D ocean bottom node (OBN) dataset.

