

Total-variation regularization strategies in full-waveform inversion

Ernie Esser^{†,1,2}, Lluís Guasch², Tristan van Leeuwen³, Aleksandr Y. Aravkin⁴,
and Felix J. Herrmann¹

¹The University of British Columbia, Department of Earth, Ocean and
Atmospheric Sciences;

²Sub Salt Solutions Limited;

³Mathematical Institute, Utrecht University;

⁴Applied mathematics, University of Washington;

Abstract

We propose an extended full-waveform inversion formulation that includes general convex constraints on the model. Though the full problem is highly nonconvex, the overarching optimization scheme arrives at geologically plausible results by solving a sequence of relaxed and warm-started constrained convex subproblems. The combination of box, total-variation, and successively relaxed asymmetric total-variation constraints allows us to steer free from parasitic local minima while keeping the estimated physical parameters laterally continuous and in a physically realistic range. For accurate starting models, numerical experiments carried out on the challenging 2004 BP velocity benchmark demonstrate that bound and total-variation constraints improve the inversion result significantly by removing inversion artifacts, related to source encoding, and by clearly improved delineation of top, bottom, and flanks of a high-velocity high-contrast salt inclusion. The experiments also show that for poor starting models these two constraints by themselves are insufficient to detect the bottom of high-velocity inclusions such as salt. Inclusion of the one-sided asymmetric total-variation constraint overcomes this issue by discouraging velocity lows to buildup during the early stages of the inversion. To the author's knowledge the presented algorithm is

the first to successfully remove the imprint of local minima caused by poor starting models and band-width limited finite aperture data.

[†]John "Ernie" Esser passed away on March 8, 2015 while preparing this manuscript. The original is posted here: <https://www.slim.eos.ubc.ca/content/total-variation-regularization-strategies-full-waveform-inversion-improving-robustness-noise>.

I. INTRODUCTION

Full-waveform inversion reconstructs high-resolution gridded models of subsurface medium parameters from seismic measurements by solving a large-scale inverse problem [1], [2]. The classic approach uses the following non-linear least-squares formulation:

$$\min_m \frac{1}{2} \sum_{j=1}^{N_s} \|F(m)q_j - d_j\|_2^2, \quad (1)$$

where the vector $m \in \mathbb{R}^M$ contains a discrete set of parameters describing the medium (e.g., spatially varying soundspeed), the vector q_j is the source term for the j^{th} experiment, $F(m)$ is the non-linear forward operator and $d_j \in \mathbb{C}^{N_r}$ are the corresponding measurements [1]. We denote the total number of measurements by $N = N_r \times N_s$. In the large-scale setting, problem size can vary from $M \sim 10^6$, $N \sim 10^4$ up to $M \sim 10^9$, $N \sim 10^6$.

Application of the forward operator typically involves solving a partial differential equation (PDE) with coefficients m and right-hand-sides q_j , $j = 1 \cdots N_s$. The operator $F(m)$ can be formally written $PA(m)^{-1}$, and in practice requires the solution of a PDE cast into a discretized linear system of equations $A(m)u = q$ (e.g. Helmholtz), followed by applying the measurement operator P to the solution $u \in \mathbb{R}^M$. Solving these PDEs is the main computational cost in evaluating the objective (1). Calculating the gradient of the objective can be done using the so-called *adjoint-state* approach and requires additional PDE solves for each experiment.

A well-known problem with such approaches is that the objective may have *parasitic* stationary points or local minima that are not informative about the true parameters. In seismic inversion in particular, the so-called loop- or cycle-skipping phenomenon is a common source of such local minima [3]. In practice this means that the inversion is very sensitive to the initial value of m . To mitigate this problem, many alternative formulations to the standard least-squares problem (1) have been proposed [4], [5], [6], extending the parameter space to avoid local minima. The general strategy is to design formulations that are useful in controlling the model space, moving away from these detrimental parasitic stationary points and toward models with realistic physical properties. The main challenge is to make such formulations *computationally efficient*, so that they can be used on large-scale problems.

Contributions. Our primary focus is to design robust formulations by controlling the model space. We propose a general framework for PDE-constrained optimization that allows multiple convex constraints to be applied. Though the overarching problem is non-convex, it is solved using a sequence of constrained convex subproblems. The generality of the scheme allows the simultaneous application of several constraints that turn out to be particularly well suited to seismic inversion: bound constraints on the model parameters (slowness squared), discontinuity preserving total-variation constraints, and a novel asymmetric variant developed specifically for seismic inverse problems. The best results obtained use all three constraints.

When given accurate starting models, the final scheme removes inversion artifacts and improves the delineation of high-velocity and high-contrast salt inclusions. More importantly, the asymmetric total-variation constraint prevents the buildup of detrimental artifacts related to parasitic local minima for poor starting models. To our knowledge, this scheme is a first instance of a *hands-free* inversion methodology that produces high-fidelity reproducible results where other wave-equation based inversions fail.

In the remainder of the introduction, we give a brief illustration of two key techniques that help control the model space: a penalized reformulation of (1), and the systematic incorporation of prior knowledge via multiple convex constraints.

Avoiding local minima. An important penalized reformulation of PDE-constrained optimization that helps avoid local minima was recently proposed by [7]:

$$\min_{m, \Delta q} \frac{1}{2} \sum_{j=1}^{N_s} \|F(m)(q_j + \Delta q_j) - d_j\|_2^2 + \lambda^2 \|\Delta q_j\|_2^2, \quad (2)$$

where $\Delta q = [\Delta q_1; \Delta q_2, \dots, \Delta q_{N_s}]$ can be thought of as slack variables that allow some freedom in fitting the data even for a wrong set of parameters. As $\lambda \uparrow \infty$, $\|\Delta q\|_2 \downarrow 0$, and (2) coincides with the original problem (1) [8]. The above problem can be solved by *projecting* out the slack variable [9], which can be done efficiently since Δq has a closed form solution for every fixed m . The result is a modified objective depending on m alone, and well-behaved as λ increases [10]. To illustrate the benefit of this extension, consider the following example.

We aim to retrieve the soundspeed of a medium from measurements of the response of a bandlimited source. The experimental setup is depicted in figure 1a; a single source (*) and three receivers (∇) are located as shown. The source emits a pulse that is subsequently

recorded, leading to three time series as depicted in figure 1b. The forward operator, acting on the temporal source signature $q(t)$, is defined as

$$F(c)q = \begin{pmatrix} \mathcal{F}^{-1}e^{i\omega\|x_1-x_s\|_2/c}\mathcal{F}q \\ \mathcal{F}^{-1}e^{i\omega\|x_2-x_s\|_2/c}\mathcal{F}q \\ \mathcal{F}^{-1}e^{i\omega\|x_3-x_s\|_2/c}\mathcal{F}q \end{pmatrix},$$

where \mathcal{F} denotes the temporal Fourier transform, ω the angular frequency, $i = \sqrt{-1}$, x_i is the receiver location, x_s is the source location, and c is the soundspeed of the medium.

Figure 1c shows the conventional ($\lambda \rightarrow \infty$) and extended ($\lambda = 0.1$) objectives as a function of the velocity c . The local minima are clearly visible in the conventional objective. By enlarging the search space, the extensions help to mitigate these local minima effectively.




Fig. 1: Example 1. (a) Snapshot of time-domain wavefield generated by a single source (*); (b) recorded time-domain data at three receiver locations (∇); (c) objective function for conventional ($\lambda \rightarrow \infty$) and extended ($\lambda = 0.1$) formulations.

In addition to exhibiting local minima, the above toy problem is severely ill-conditioned due to the band-limited nature of the data. Large- and small-scale variations of c are difficult to retrieve because low and high frequencies are missing. Full-waveform inversion is also hampered by missing spatial frequencies at both ends of the spectrum due to physical constraints on seismic surveys, which include restriction of sources and receivers to the

surface, bandwidth limitation of sources and receivers, and finite aperture (maximal distance between sources and receivers).

Since our main objective is to retrieve the global velocity structure in geologic areas with high-velocity and high-contrast (sharp) inclusions, we need to recover both the low frequencies (responsible for the kinematics and therefore placement of velocity perturbations), and the high frequencies to ensure accurate delineations of the high-velocity inclusions. A complicating factor is that all scales in c are intrinsically coupled through the wave equation—i.e., they appear as coefficients nonlinearly in the PDE, so failure to capture the large-scale variations in turn leads to a failure to retrieve useful information on the medium-scale features. As in many other inverse problems, invoking prior knowledge in the form of certain rudimentary constraints can have a significant impact on the quality of recovered models.

Incorporating prior knowledge using constraints. We aim to retrieve a one-dimensional function $c(z)$ from band-pass filtered measurements $d = F(c)q \equiv c * q$, where $*$ denotes convolution and q is the band-pass filter in the time domain. Figure 2a shows the ground-truth discrete velocities and corresponding data. A least-squares reconstruction is shown in 2b. This result clearly illustrates the failure to retrieve the large-scale trends in c . Since c represents a physical parameter, failure to reproduce these large-scale trends changes the physical interpretation (e.g. the two-way travel time) completely. We can incorporate prior information through constraints by imposing the lower bound $c_0 = 1500 \text{ ms}^{-1}$, and requiring that c increase monotonically. The resulting formulation is

$$\min_{c \geq c_0} \|F(c)q - d\|_2^2 \quad \text{subject to} \quad Dc \geq 0, \quad (3)$$

where D is the finite difference matrix. The results from this constrained formulation are shown in figure 2c. The additional constraints allow us to retrieve the ground-truth exactly. This example is simplistic, but effectively illustrates the important role constraints can play in restricting the feasible model space, with potential to recover broadband velocity profiles from bandwidth limited data.

Motivated by these ideas, we develop an inversion framework able to avoid local minima and mitigate ill-posedness by incorporating three kinds of constraints: (1) bounds, (2) total variation (TV), and (3) a novel one-sided TV extensions based on the hinge loss. Both the

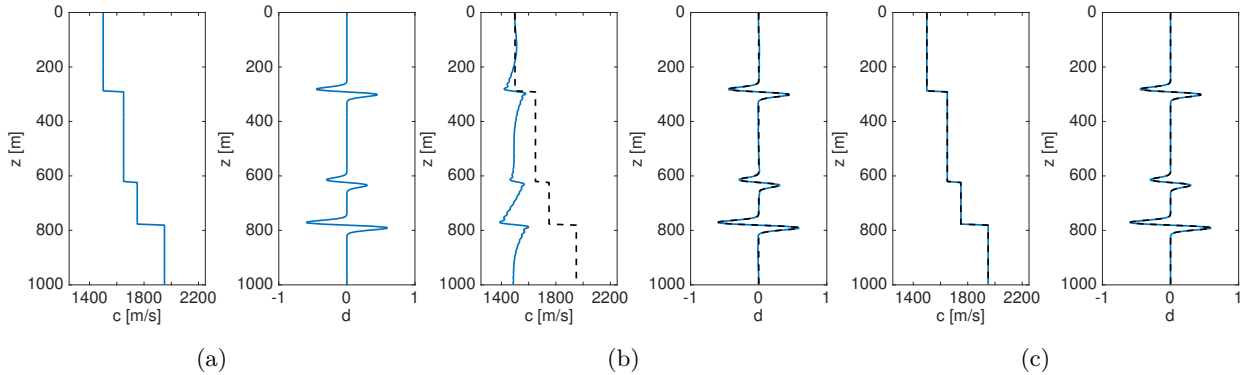


Fig. 2: Example 2. (a) Monotonously increasing velocity profile and observed data. (b) Unconstrained inversion result fits the observations but misses the global trend. (c) Constrained inversion results fits the data and recovers the velocity model accurately.

classic formulation (1) as well as the relaxed formulation (2) fit naturally into this framework. Formulation (2) has additional modeling advantages (see Fig. 1) as well as computational advantages highlighted in Section II.

The seismic setting comprises imaging sedimentary basins with the following features: (1) depth-increasing velocities, (2) high-velocity and high-contrast inclusions (i.e. salt bodies), and (3) potential over-pressured reservoirs that yield velocity lows. We show that the velocity structure over this complicated geology can be recovered using the proposed framework. The one-sided TV constraint is especially important, since it penalizes the negative vertical derivative more heavily than the positive vertical derivative making it less likely for the velocity to decrease with depth. The final workflow gradually relaxes this asymmetric constraint, recovering velocity structures with complex high-velocity and high-contrast inclusions from poor starting models.

Outline of the paper. The paper is organized as follows. In Section II, we specify a scaled gradient projection framework. In Section III, we show how to incorporate bound constraints and the TV constraint, and how to solve the resulting convex subproblems. Numerical experiments for bound and TV constraints are presented in Section IV. A new asymmetric TV variant, implementation details, and numerical results for joint convex constraints are

presented in Section V. Sections VI and VII present discussion and conclusions. All necessary details needed to reproduce the numerical examples are given in the appendix.

II. SCALED GRADIENT PROJECTION FOR LARGE-SCALE PROGRAMS

Considering formulations (1) and (2) with constraints such as (3), our goal is to develop a framework for large-scale optimization where we distinguish between relatively simple convex constraints and more computationally expensive PDE constraints, discussed later.

First consider problems of the following form:

$$\min_m f(m) \quad \text{subject to } m \in C, \quad (4)$$

where $f(m)$ is smooth, but may be nonconvex, and expensive to evaluate. The set C represents the convex constraint, designed to mitigate the ill-posedness of the inverse problem. Box constraints are common to seismic inversion, with $C = \{m : m_i \in [b_i, B_i]\}$ for $i = 1, \dots, M$, where the b_i 's are the lower and the B_i are the upper bounds for each entry in m . The set C may also be constructed as the intersection of several simpler sets. In the following sections,

The scaled projected gradient method for (4) is well-adapted for the PDE constrained setting, and can incorporate different constraint sets C (e.g. intersection of box and TV constraints). We first consider scaled gradient descent, and then describe scaled projected gradient.

Scaled gradient descent

Following [11], we start by considering the following unconstrained problem:

$$\min_m f(m),$$

and apply an iterative algorithm of the form

$$\begin{aligned} \Delta m &= \arg \min_{\Delta m} \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top H^n \Delta m \\ m^{n+1} &= m^n + \Delta m. \end{aligned} \quad (5)$$

In this quadratic minimization problem, $\nabla f(m^n)$ is the gradient of the objective with respect to the discretized model at the n^{th} iteration and H^n is a positive definite approximation to the Hessian of f . The Hessian approximation can range from a simple gradient descent with a step size of magnitude α if $H^n = \frac{1}{\alpha} I$, to more sophisticated forms such as $H^n = \nabla^2 f(m^n)$, which corresponds to Newton's method using the full Hessian (when it is positive definite).

Among positive definite Hessian approximations, the choice is typically shaped by a tradeoff between quality of approximation and efficiency of the scaled gradient iteration. Problem structure plays a key role here. For example, the Hessian for the conventional formulation (1) is a dense matrix whose evaluation requires PDE solves, while the extended formulation (2) produces accurate sparse Hessian approximations.

Scaled gradient projections

Adding constraints $m \in C$ to solve (4) requires a simple modification to the model updates Δm :

$$\begin{aligned} \Delta m &= \arg \min_{\Delta m} \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top H^n \Delta m \\ \text{s.t. } & m^n + \Delta m \in C. \end{aligned} \tag{6}$$

$$m^{n+1} = m^n + \Delta m.$$

This iteration is known as scaled gradient *projection* method [11], [12]. It ensures model iterates m^{n+1} are always in C , and admits provable convergence guarantees, discussed below. Note that solving for the unconstrained minimizer Δm in (5) and then projecting onto C

$$m^{n+1} = \Pi_C(m^n - (H^n)^{-1} \nabla(f(m^n))), \tag{7}$$

will not converge to a solution of (4) for general H^n [11].

We show that iterations (6) remain computationally tractable as long as C is easy to project onto, or is an intersection several simple simple convex constraints. In PDE constrained optimization, the main computational burden is obtaining the gradient and Hessian approximation, which remain fixed within (6).

The scaled gradient projection framework includes a variety of methods depending on the choice of H^n . For example, if $H^n = \frac{1}{\alpha} I$, we can solve (6) with (7) where the iterations correspond to projected gradients with stepsize of α . When the H^n are chosen to approximate the Hessian of f at m^n , we arrive at projected Newton-like methods that incorporate second-order information. A good summary of some of the possible choices can be found in [13], which uses a projected quasi-Newton method proposed based on a limited-memory Broyden-Fletcher-Goldfarb-Shannon (l-BFGS) approximation of the Hessian, and solves the convex subproblems

for each update with a spectral-projected gradient method. The extended formulation (2) fits naturally into this framework, and allows for accurate sparse Gauss-Newton Hessian approximations. We exploit this structure with the development of a dedicated solver.

Implicit trust-region approach

As we mentioned earlier, objective function evaluations in PDE-constrained optimization can be expensive. Therefore, we use an implicit trust region method, which avoids expensive linesearches. For this purpose, we replace H^n by $H^n + c_n I$ while adjusting the damping parameter c_n at each iteration adaptively and rejecting iterations that do not lead to a sufficient decrease in the objective.

When ∇f is Lipschitz continuous, i.e. for some constant K

$$\|\nabla f(x) - \nabla f(y)\| \leq K\|x - y\| \text{ for all } x, y \in C,$$

and H is a symmetric matrix, then c can be chosen large enough so that

$$f(m + \Delta m) - f(m) \leq \Delta m^\top \nabla f(m) + \frac{1}{2} \Delta m^\top (H + cI) \Delta m \quad (8)$$

for any $m \in C$ and Δm such that $m + \Delta m \in C$. When solving (6), $\Delta m = 0$ is feasible for the right-hand side, so if (8) is satisfied we have $f(m + \Delta m) - f(m) \leq 0$. If moreover the solution of (6) satisfies $\|\Delta m\| > 0$, we have

$$f(m + \Delta m) - f(m) < 0,$$

so strict descent (or stationarity) is guaranteed for method (6) under condition (8).

We can derive explicit conditions that ensure (8) is satisfied. Note that since

$$f(m + \Delta m) - f(m) \leq \Delta m^\top \nabla f(m) + \frac{K}{2} \|\Delta m\|^2,$$

it follows that

$$f(m + \Delta m) - f(m) \leq \frac{1}{2} (K - \lambda_H^{\min} - c) \|\Delta m\|^2 + \Delta m^\top \nabla f(m) + \frac{1}{2} \Delta m^\top (H + cI) \Delta m,$$

where λ_H^{\min} denotes the smallest eigenvalue of H , so choosing $c > K - \lambda_H^{\min}$ ensures (8). However, this extremely conservative choice of the damping parameter c may lead to a

slow rate of convergence, so instead one can choose c to be as small as possible while still decreasing the objective by a sufficient amount:

$$f(m + \Delta m) - f(m) \leq \sigma \left(\Delta m^\top \nabla f(m) + \frac{1}{2} \Delta m^\top (H + cI) \Delta m \right), \quad (9)$$

for some $\sigma \in (0, 1]$. Using the same framework as in [14], the resulting method is summarized in Algorithm 1.

Algorithm 1 A Scaled Gradient Projection Algorithm for (4)

$n = 0$; $m^0 \in C$; $\rho > 0$; $\epsilon > 0$; $\sigma \in (0, 1]$;

H symmetric with eigenvalues between λ_H^{\min} and λ_H^{\max} ;

$\xi_1 > 1$; $\xi_2 > 1$; $c_0 > \max(0, \rho - \lambda_H^{\min})$;

while $n = 0$ or $\frac{\|m^n - m^{n-1}\|}{\|m^n\|} > \epsilon$

$$\Delta m = \arg \min_{\Delta m + m^n \in C} \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m$$

$$\text{if } f(m^n + \Delta m) - f(m^n) > \sigma (\Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m)$$

$$c_n = \xi_2 c_n$$

else

$$m^{n+1} = m^n + \Delta m$$

$$c_{n+1} = \begin{cases} \frac{c_n}{\xi_1} & \text{if } \frac{c_n}{\xi_1} > \max(0, \rho - \lambda_H^{\min}) \\ c_n & \text{otherwise} \end{cases}$$

Define H^{n+1} to be symmetric Hessian approximation

with eigenvalues between λ_H^{\min} and λ_H^{\max}

$$n = n + 1$$

end if

end while

If ∇f is Lipschitz continuous and level sets of $f(m)$ intersected with C are bounded, then any limit point m^* of the sequence of iterates $\{m^n\}$ defined by Algorithm 1 is a stationary point of (4), i.e. $(m - m^*)^\top \nabla f(m^*) \geq 0$ for all $m \in C$.

To implement Algorithm 1, we need to specify both the Hessian approximation and a way to solve subproblem (6). In the remainder of the paper, we show how to incorporate (1) box, (2) TV and (3) one-sided TV (hinge-loss) constraints, and explain how to solve (6) in each

of the proposed formulations. We end this section with an important example where (6) is very simple to implement.

Example: diagonal Hessian approximation with box constraints

When H^n is diagonal and positive (e.g. approximation of the Gauss-Newton Hessian of (2)), Algorithm 1 simplifies. In this case, the subproblem

$$\begin{aligned} \Delta m &= \arg \min_{\Delta m} \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m \\ &\text{subject to } m_i^n + \Delta m_i \in [b_i, B_i] \end{aligned} \tag{10}$$

has the following closed form solution:

$$\Delta m_i = \max \left(b_i - m_i^n, \min \left(B_i - m_i^n, -\widetilde{\Delta} m_i \right) \right), \quad i = 1, \dots, M,$$

with $\widetilde{\Delta} m = (H^n + c_n I)^{-1} \nabla f(m^n)$. In this expression, b_i and B_i are the lower and upper bounds for m_i .

III. TOTAL-VARIATION CONSTRAINTS

In FWI, even with box constraints, the recovered model can still contain artifacts and spurious oscillations. As we will demonstrate below, inaccuracies in m can be reduced via convex constraints that bound the size of the total-variation norm ball to some positive value τ .

TV-norm regularization, via penalties, constraints, or objectives is widely used in image processing to remove noise while preserving discontinuities [15]. TV-norm minimization also plays an important role in a variety of other inverse problems, especially when the unknown model parameters can be represented by piecewise constant or piecewise smooth functions. Examples include electrical impedance tomography [16], inverse wave propagation [17], and more recently in FWI promoting blockiness [18], shape optimization [19], and time-lapse data [20]. These approaches use a TV-penalized formulation, while we introduce the regularization as a constraint. The constrained formulation keeps all iterates confined to a pre-defined convex set C , restricting the feasible model space. Penalized formulations do not offer this guarantee, which proves to be essential in FWI.

If we represent m as a $M_x \times M_y$ array, we can define the discrete TV-norm as

$$\begin{aligned} \|m\|_{TV} &= \frac{1}{h} \sum_{k,l} \sqrt{(m_{k+1,l} - m_{k,l})^2 + (m_{k,l+1} - m_{k,l})^2} \\ &= \sum_{k,l} \frac{1}{h} \left\| \begin{bmatrix} m_{k,l+1} - m_{k,l} \\ m_{k+1,l} - m_{k,l} \end{bmatrix} \right\|, \end{aligned} \quad (11)$$

which is the sum of the ℓ_2 norms of the discrete gradient vectors at each point in the discretized model. We assume Neumann boundary conditions so that these differences are zero at the boundary. We arrive at a more compact expression for $\|m\|_{TV}$ if we define the finite difference operator D such $(Dm)_i$ is the discrete gradient at location indexed by $i = 1, \dots, M$, where $M = M_x \times M_y$. We now define

$$\|m\|_{TV} = \|Dm\|_{1,2} := \sum_{i=1}^M \|(Dm)_i\|. \quad (12)$$

If we impose this TV-norm constraint in addition to the box constraints defined earlier, the set C is an intersection

$$C = [b_i, B_i] \cap \{\|m\|_{TV} \leq \tau\}.$$

The corresponding model updates (6) required by Algorithm 1 are given by

$$\begin{aligned} \Delta m &= \arg \min_{\Delta m} \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m \\ &\text{subject to } m_i^n + \Delta m_i \in [b_i, B_i] \text{ and } \|m^n + \Delta m\|_{TV} \leq \tau \\ m^{n+1} &= m^n + \Delta m . \end{aligned} \tag{13}$$

Example: Projecting the Marmousi model on the intersection of box and TV-norm constraints

Before we discuss how to minimize (4) in the seismic setting, we motivate the use of TV-norm constraints in seismic inverse problems. Consider projecting the Marmousi model [21], shown in Figure 3a, onto two sets with TV-norm constraints using decreasing values of τ . Let m_0 denote the original Marmousi model and let $\tau_0 = \|m_0\|_{TV}$. For the bound constraints on the slowness squared, set $B_i = 4.4444 \times 10^{-7} \text{ s}^2\text{m}^{-2}$ everywhere, which corresponds to a lower bound of 1500 ms^{-1} on the velocity. Taking advantage of the fact that these constraints can vary spatially, let $b_i = 4.4444 \times 10^{-7} \text{ s}^2\text{m}^{-2}$ in the water layer and $b_i = 3.3058 \times 10^{-8} \text{ s}^2\text{m}^{-2}$ everywhere else, which corresponds to an upper bound of 5500 ms^{-1} on the velocity. The orthogonal projection of m_0 onto the intersection of these box and TV-norm ball constraints is defined by

$$\begin{aligned} \Pi_C(m_0) &= \arg \min_m \frac{1}{2} \|m - m_0\|^2 \\ &\text{subject to } m_i \in [b_i, B_i] \text{ and } \|m\|_{TV} \leq \tau . \end{aligned} \tag{14}$$

Results with $\tau = .6\tau_0$ and $\tau = .3\tau_0$ are shown in Figure 3. The vertical lines at $x = 5000 \text{ m}$ indicate the location of the 1D vertical slices shown in Figure 4 for both slowness squared (a) and the velocity (b). As we decrease the size of the TV-norm ball, the feasible models that are close to the original model (m_0) in the 2-norm become more and more "cartoon like" with fewer and fewer unconformities. The size of the TV-norm ball (τ) controls the complexity of the model while still allowing for discontinuous unconformities. We exploit this property of the TV-norm extensively.

Solving the convex subproblems

The proposed approach requires solving the quadratic approximation in (6) with a given gradient and Hessian approximation. A computationally efficient approach for the convex




Fig. 3: Marmousi model (a) and projected Marmousi model for $\tau = .6\tau_0$ (b) and $\tau = .3\tau_0$ (c).

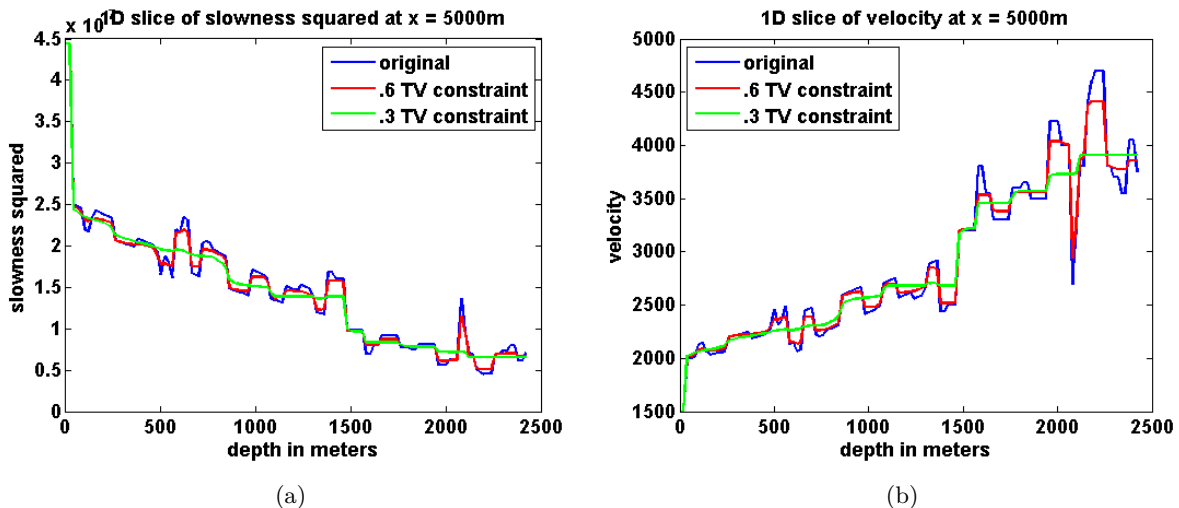


Fig. 4: Comparison of slices from the Marmousi model and its projections onto different TV constraints both in terms of slowness squared (a) and velocity (b).

subproblems of type (13) is the primal-dual hybrid gradient (PDHG) method [22] studied in [23], [24], [25], [26]. To develop the method, we first write down the Lagrangian for a dualization of (13):

$$\begin{aligned} \mathcal{L}(\Delta m, p) = & \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m + g_B(m^n + \Delta m) \\ & + p^\top D(m^n + \Delta m) - \tau \|p\|_{\infty, 2}. \end{aligned} \quad (15)$$

In (15), p is the vector of Lagrange multipliers, and g_B is an indicator function for the bound constraints

$$g_B(m) = \begin{cases} 0 & \text{if } m_i \in [b_i, B_i] \\ \infty & \text{otherwise.} \end{cases}$$

Here, $\|\cdot\|_{\infty,2}$ uses mixed norm notation to denote the dual norm of $\|\cdot\|_{1,2}$. It takes the max instead of the sum of the ℓ_2 norms so that $\|Dm\|_{\infty,2} = \max_i \|(Dm)_i\|$ in the notation of (12). This saddle point problem can be derived from the convex subproblem in (13) using the conjugate representation of the TV constraint:

$$\sup_p \left\{ p^\top D(m^n + \Delta m) - \tau \|p\|_{\infty,2} \right\}, \quad (16)$$

which equals the indicator function

$$\begin{cases} 0 & \text{if } \|D(m^n + \Delta m)\|_{1,2} \leq \tau \\ \infty & \text{otherwise.} \end{cases}$$

To find a saddle point of (15), the modified PDHG iterations are given by

$$\begin{aligned} p^{k+1} &= \arg \min_p \tau \|p\|_{\infty,2} - p^\top D(m^n + \Delta m^k) + \frac{1}{2\delta} \|p - p^k\|^2 \\ \Delta m^{k+1} &= \arg \min_{\Delta m} \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m \\ &\quad + \Delta m^\top D^\top (2p^{k+1} - p^k) + \frac{1}{2\alpha} \|\Delta m - \Delta m^k\|^2 \\ &\text{subject to } m_i^n + \Delta m_i \in [b_i, B_i]. \end{aligned} \quad (17)$$

These iterations can be written more explicitly as

$$\begin{aligned} p^{k+1} &= p^k + \delta D(m^n + \Delta m^k) - \Pi_{\|\cdot\|_{1,2} \leq \tau\delta} (p^k + \delta D(m^n + \Delta m^k)) \\ \Delta m_i^{k+1} &= \max \left(b_i - m_i^n, \min \left(B_i - m_i^n, \widetilde{\Delta m}_i \right) \right), \end{aligned} \quad (18)$$

where

$$\widetilde{\Delta m} = (H^n + (c_n + \frac{1}{\alpha})I)^{-1} (-\nabla f(m^n) + \frac{\Delta m^k}{\alpha} - D^\top (2p^{k+1} - p^k))$$

and $\Pi_{\|\cdot\|_{1,2} \leq \tau\delta}(z)$ denotes the orthogonal projection of z onto the ball of radius $\tau\delta$ in the $\|\cdot\|_{1,2}$ norm. This requires projecting the vector of ℓ_2 norms of the spatial gradient vectors onto a simplex. A simple approach to project z onto the unit simplex $\{x : x_i \geq 0, \sum_i x_i = 1\}$ is to use

bisection to find the threshold a such that $\sum_i \max(0, z_i - a) = 1$, in which case $\max(0, z_i - a)$ is the i^{th} component of the projection. An efficient linear time $O(n)$ implementation based on this idea is developed and analyzed in [27].

The step size restriction required for convergence is $\alpha\delta \leq \frac{1}{\|D^\top D\|}$ (see [23] for details). If h is the mesh width, then the eigenvalues of $D^\top D$ are between 0 and $\frac{8}{h^2}$ by the Gershgorin Circle Theorem, so it suffices to choose positive α and δ such that $\alpha\delta \leq \frac{h^2}{8}$.

The relative scaling of α and δ can have a large effect on the convergence rate of the method [23]. A reasonable choice for the fixed step size parameters is

$$\alpha = \frac{1}{\max(H^n + c_n I)} \text{ and } \delta = \frac{h^2 \max(H^n + c_n I)}{8} \leq \frac{\max(H^n + c_n I)}{\|D^\top D\|}.$$

However, this choice may be too conservative. The convergence rate of the method can be improved by using iteration-dependent step sizes as proposed by [24]. The adaptive backtracking strategy introduced by [28] can also be a practical way of choosing step size parameters that improve the convergence rate.

IV. NUMERICAL EXPERIMENTS

We consider four 2D numerical experiments based on synthetic data. The PDE in this case is a scalar Helmholtz equation

$$\left(\omega^2 m + \nabla^2\right) u = q \quad \text{in the interior of the domain}$$

with radiation boundary conditions

$$\left(\frac{\partial}{\partial n} - i\omega\right) u = 0,$$

where $\frac{\partial}{\partial n}$ denotes the derivative in the direction normal to the boundary. A standard finite difference discretization of the Helmholtz operator leads to a sparse banded matrix $A(m)$. The data are collected by an array of receivers for several point-sources and a range of angular frequencies ω . For details on discretization see the appendix.

In the experiments, we primarily use the extended formulation (2) and minimize the objective defined in (29). The corresponding gradient and Hessian expressions are given in (30) and (31). We refer to this approach as *Wavefield Reconstruction Inversion* (WRI). We also show the benefits of adding constraints to the conventional approach of minimizing the objective defined in (25), where we use the gradient and pseudo-Hessian are defined in (26) and (28). We refer to the latter approach as *Full-Waveform Inversion* (FWI). A standard approach in FWI/WRI is to start the inversion from low frequencies and gradually move to higher frequencies. This continuation strategy helps to avoid cycle-skipping related local minima [29].

To illustrate the performance of the proposed constrained formalism, we consider two examples derived from the 2004 BP velocity benchmark data set [30]. This data set was designed to evaluate the capabilities of velocity-analysis techniques for complex geologies that contain sedimentary basins, with increasing velocities, interspersed with high-contrast and high-velocity salt bodies.

The first experiment tries to recover the top left portion of this synthetic model from an accurate smooth starting model. Even in this situation, the resulting inverted model has noisy artifacts related to the applied source encoding, and the deeper part of the estimated model tends to have incorrect discontinuities. The TV constraint helps to remove some of

these artifacts while still recovering the most significant discontinuities. This example also demonstrates that relaxing the TV constraint over multiple passes through the frequency batches leads to improvements. A stronger TV constraint gives an oversmoothed model estimate, but this estimate then serves as a good initial model for future passes as the TV constraint is relaxed.

The second experiment tries to recover the top middle portion of the same BP velocity model from a poor starting model. This experiment illustrates that the combination of edge-preserving TV constraint and box-constraints are inadequate, and motivates an additional asymmetric TV-norm constraint, tailored to geologies with sedimentary basins with salt inclusions. Our approach here also uses a continuation strategy that at first strongly discourages downward jumps in the estimated velocity, and then gradually relaxing this constraint. This strategy helps prevent the method from getting stuck in bad local minima at early passes, eventually allowing downward jumps in velocity for a better data fit.

We first present details on frequency continuation and *simultaneous shot* strategies used in the experiments.

Frequency Continuation

We work with small batches of frequency data at a time, moving from low to high frequencies in overlapping batches of two. This frequency continuation strategy does not guarantee that we solve the overall problem after a single pass from low to high frequencies, but is far more computationally tractable than minimizing over all frequencies simultaneously. Moreover, the continuation strategy of moving from low to high frequencies helps prevent the iterates from tending towards bad local minima [29].

For example, if the data consists of frequencies starting at 3 Hz sampled at intervals at 1 Hz, then we would start with the 3 and 4 Hz data, use the computed m as an initial guess for inverting the 4 and 5 Hz data and so on. For each frequency batch, we will compute at most 25 outer iterations, each time solving the convex subproblem to convergence, stopping when

$$\max \left(\frac{\|p^{k+1} - p^k\|}{\|p^{k+1}\|}, \frac{\|\Delta m^{k+1} - \Delta m^k\|}{\|\Delta m^{k+1}\|} \right) \leq 1 \times 10^{-4}.$$

Since the magnitude of the data depends on the frequency, one can incorporate frequency dependent weights in the definition of the objective. However, as we only work with small frequency batches in practice, these weights do not have a significant effect.

In some of our experiments we perform multiple passes through the frequencies to further improve the results.

Simultaneous sources

The data are typically collected for many sources, resulting in a redundant data set. Since the number of PDE-solves required at each iteration is dictated by the number of sources, it makes sense to try to compress the data volume and thus reduce the computational cost. One way of doing this is by randomly projecting the data, also known as "source encoding", resulting in a smaller number of terms

$$\tilde{d}_i = \sum_{j=1}^{N_s} w_{ij} d_j, \quad \text{for } i = 1, 2, \dots, \tilde{N}_s, \quad \tilde{N}_s \ll N_s$$

where the w_{ij} 's are drawn from a standard normal distribution [31]. The source terms are similarly reduced (using the same weights)

$$\tilde{q}_i = \sum_{j=1}^{N_s} w_{ij} q_j, \quad \text{for } i = 1, 2, \dots, \tilde{N}_s.$$

Various alternative strategies have been proposed to choose the random weights and tailor the optimization to deal specifically with the resulting stochasticity in the problem [32], [33], [34].

Example – accurate starting model

To demonstrate the performance of the proposed formulation, we first consider a 3 by 12 km portion of the the 2004 BP velocity benchmark model with a gridspacing of 20 m, as shown in Figure 5a, for a relative accurate smooth starting model shown in 5b. We use a Ricker wavelet with a peak frequency at 15 Hz as a for the source signature. There are 126 sources every 80 m between 1 and 11 km at a depth of 40 m below the top of the model; 299 receivers are placed at 60 m below the surface and covering all the model in the horizontal dimension, in an equally spaced distribution (every 40 m). We employ frequency continuation by looping

through the frequencies ranging from 3 to 20 Hz from low to high in overlapping batches of two frequencies each. We define bound constraints on the slowness squared that correspond to minimum and maximum velocities of 1400 and 5000 ms^{-1} , respectively. The inversion result for the velocity after one pass through the frequencies is included in Figure 6a and shows that the velocity model is reasonably well recovered (cf. Figures 5a and 6a). This result is obtained using WRI with box constraints, working with only two sources ($\tilde{N}_s = 2 \ll N_s$). While encouraging, there are still visible noisy artifacts due to the simultaneous sources and from structural velocity lows in the salt. Without including additional constraints, the conventional adjoint-state approach to FWI is not able to produce a tangible result for this experiment.




Fig. 5: Top left portion of BP 2004 velocity model (a), and initial velocity (c).

Benefit of Multiple Passes

Despite its evident shortcomings, after one pass through the frequencies (Figure 6a) the WRI result can serve as input for a second pass. Using multiple multiscale passes is used frequently in scientific computing (e.g. V-cycles in multigrid), and [35] proposed this strategy for wave-equation based inversion. Comparing the inversion results after one and two passes shows significant improvements after the second pass (see Figure 6b) where the oscillation

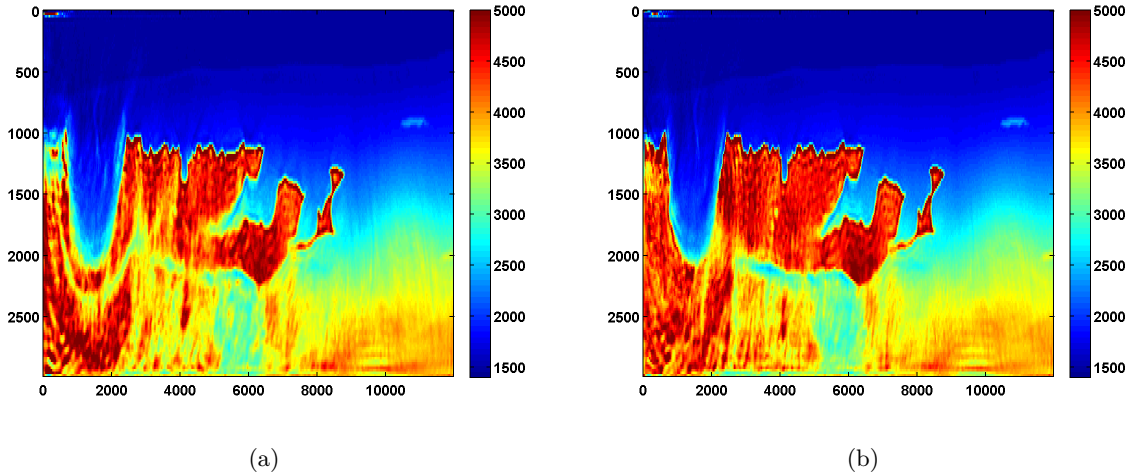


Fig. 6: Recovered velocity without TV constraint from a good smooth initial model after one pass (a) and after two passes (b) using small frequency batches from 3 to 20 Hz.

just below the top salt on the left side of the model mostly disappears after the second pass. The WRI approach with multiple passes allows useful results in situations where conventional FWI fails. However, the result remains noisy and suffers from artifacts near the boundary of the model due to poor illumination.

Including the TV constraint

Shortcomings of WRI with multiple passes include noisy artifacts lack of clear delineation at the top and bottom salt. As shown in Figure 7, the inversion results improve significantly when we use the TV-norm constraint $\|m\|_{TV} \leq \tau$, together with a relaxation strategy over multiple passes. For the first pass, we choose τ to be 0.9 times the TV-norm of the ground truth model. To reduce computation time, we again use two simultaneous shots, but now with Gaussian weights resampled every time the model is updated. The estimated model after 25 outer iterations per frequency batch is shown in Figure 7a. Using this result as a warm start for a second pass through the frequency batches yields the improved result in Figure 7b. The top salt is much better resolved compared to the examples without the TV-norm constraint. The second pass also leads to a significant improvement of the bottom

salt while preserving the velocity low below the salt related to an over-pressured reservoir. Because the model is off centre, issues with illumination remain but overall the interior of the salt body itself is well recovered and includes the small low-velocity inclusions that can be observed in the original model (Figure 5a).




Fig. 7: Recovered velocity with a TV constraint from a good smooth initial model after one pass with $\tau = .9\tau_{\text{true}}$ (a) and a second pass with $\tau = .99\tau_{\text{true}}$ (b) using small frequency batches from 3 to 20 Hz.

Example-poor starting model

The inversion results shown in Figure 7b relied heavily on having good initial models, and were obtained by smoothing the ground truth models. Consider the 3 by 12 km velocity model shown in Figure 8a, which is the top middle portion of the 2004 BP velocity benchmark data set, also sampled with a gridspacing of 20 m; our acquisition geometry as well as the source signature is the identical to the one defined in the previous example, but with better illumination. The approach we used to recover the top left portion of the model also works well here using a smoothed version of the true model as the starting point. However, starting with a poor initial model as in Figure 12a, the method obtains a poor inversion result,

possibly because of finding a parasitic stationary point. This happens despite the fact that the salt body is well illuminated.




Fig. 8: Top middle portion of BP 2004 velocity model (a), source and receiver locations (b) and initial velocity (c).

As shown in Figure 9, the WRI method with bound constraints alone yields a noisy result that fails to improve significantly even after multiple passes through the frequency batches. The effect of the initial updates, which decrease the velocity the after stepping into the salt, is persistent and can not be overcome. This behavior is typical for data that misses the low frequencies, as we also observed in Figure 2 of the introduction.

With a TV constraint added, the method still tends to get stuck at a poor solution, even with multiple passes and different choices of τ . Figure 10 shows the estimated velocity models after three passes, where increasing values of τ were used so that the TV constraint was weakened slightly after each pass. Inclusion of the TV constraint leads to accurate recovery of top salt, which is a significant improvement. However, the imprint of the velocity low is still too strong, and the results do not improve after three passes.

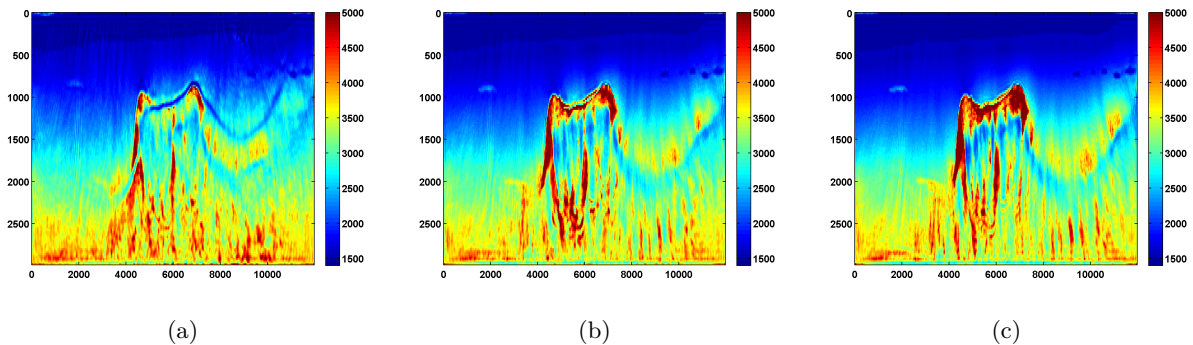


Fig. 9: Recovered velocity with a no TV constraint from a poor initial model after one pass (a) after a second pass (b) and after a third pass (c) through small frequency batches from 3 to 20 Hz.




Fig. 10: Recovered velocity without TV constraints from a poor initial model after one pass with $\tau = .75\tau_{\text{true}}$ (a), a second pass with $\tau = .825\tau_{\text{true}}$ (b) and a third pass with $\tau = .9\tau_{\text{true}}$.

V. ASYMMETRIC TV-NORM CONSTRAINTS

Since velocity generally increases with depth in sedimentary basins while remaining more or less constant in salt, it is natural to penalize downward jumps in velocity. This can be done with a one-sided, asymmetric total-variation constraint that penalizes increases in the slowness squared in the depth direction. Such a constraint naturally fits in the scaled projected gradient framework, and can be imposed along with TV-norm and box constraints.

Define a forward difference operator D_z that acts in the depth direction so that $D_z m$ is a vector of differences of the form $\frac{1}{h}(m_{k+1,l} - m_{k,l})$ for $k = 1, \dots, M_z$ and $l = 1, \dots, M_x$. To penalize the sum of the positive differences in m , we include the constraint

$$\|\max(0, D_z m)\|_1 \leq \xi, \quad (19)$$

where the max operation is understood in a componentwise sense so that

$$\|\max(0, D_z m)\|_1 = \sum_{k,l} \max(0, \frac{1}{h}(m_{k+1,l} - m_{k,l})).$$

This constraint, which limits the size of the asymmetric TV-norm ball to a value ξ , is the hinge loss penalty applied to $D_z m$. The hinge loss is commonly used in machine learning for support vector machines and support vector regression.

The constraint in (19) does not penalize model discontinuities in the horizontal direction, only in the depth direction. It is therefore likely to lead to vertical artifacts unless combined with additional constraints that penalize variations in the horizontal direction. We therefore combine this hinge-loss constraint with a TV constraint.

We can apply box constraints, TV-norm constraints, and one-sided TV constraints using the same scaled projected gradient framework. The set C in (4) is now the intersection of three constraints:

$$\min_m f(m) \quad \text{subject to} \quad m_i \in [b_i, B_i], \quad \|m\|_{TV} \leq \tau \quad \text{and} \quad \|\max(0, D_z m)\|_1 \leq \xi, \quad (20)$$

and the convex subproblem (6) now has an additional one-sided TV constraint:

$$\begin{aligned} \Delta m &= \arg \min_{\Delta m} \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m \\ &\text{subject to } m_i^n + \Delta m_i \in [b_i, B_i], \quad \|m^n + \Delta m\|_{TV} \leq \tau \\ &\text{and } \|\max(0, D_z(m^n + \Delta m))\|_1 \leq \xi. \end{aligned} \quad (21)$$

As before, we can use the PDHG algorithm. Analogous to (15), we want to find a saddle point of the Lagrangian

$$\begin{aligned} \mathcal{L}(\Delta m, p_1, p_2) &= \Delta m^\top \nabla f(m^n) + \frac{1}{2} \Delta m^\top (H^n + c_n I) \Delta m + g_B(m^n + \Delta m) \\ &\quad + p_1^\top D(m^n + \Delta m) - \tau \|p_1\|_{\infty, 2} \\ &\quad + p_2^\top D_z(m^n + \Delta m) - \xi \max(p_2) - g_{\geq 0}(p_2), \end{aligned} \quad (22)$$

where $g_{\geq 0}$ denotes an indicator function defined by

$$g_{\geq 0}(p_2) = \begin{cases} 0 & \text{if } p_2 \geq 0 \\ \infty & \text{otherwise} \end{cases}.$$

The additional terms in this Lagrangian come from the conjugate representation of the hinge loss constraint:

$$\sup_{p_2} p_2^\top D_z(m^n + \Delta m) - \xi \max(p_2) - g_{\geq 0}(p_2). \quad (23)$$

The reader can check that maximizing over p_2 we recover the indicator of the convex set defined by the 1-sided TV constraint:

$$\begin{cases} 0 & \text{if } \|\max(0, D_z(m^n + \Delta m))\|_1 \leq \xi \\ \infty & \text{otherwise.} \end{cases}$$

The modified PDHG iterations are similar to (18) and given below.

$$\begin{aligned} p_1^{k+1} &= p_1^k + \delta D(m^n + \Delta m^k) - \Pi_{\|\cdot\|_{1,2} \leq \tau \delta}(p_1^k + \delta D(m^n + \Delta m^k)) \\ p_2^{k+1} &= p_2^k + \delta D_z(m^n + \Delta m^k) - \Pi_{\|\max(0, \cdot)\|_1 \leq \xi \delta}(p_2^k + \delta D_z(m^n + \Delta m^k)) \\ \Delta m_i^{k+1} &= \max\left(b_i - m_i^n, \min\left(B_i - m_i^n, \widetilde{\Delta m}_i\right)\right) \end{aligned} \quad (24)$$

with

$$\widetilde{\Delta m} = \left(H^n + \left(c_n + \frac{1}{\alpha}\right) \mathbf{I}\right)^{-1} \left(-\nabla f(m^n) + \frac{\Delta m^k}{\alpha} - D^\top(2p_1^{k+1} - p_1^k) - D_z^\top(2p_2^{k+1} - p_2^k)\right).$$

The projection $\Pi_{\|\max(0, \cdot)\|_1 \leq \xi \delta}(z)$ is computed by projecting the positive part of z , $\max(0, z)$, onto the simplex defined by $\{z : z_k \geq 0, \sum_k z_k = \xi \delta\}$.

Including the asymmetric TV-norm constraint

Using WRI, bounds, and TV-norm constraints, our inversion results still suffered from missing low frequencies when starting from a poor initial model (see Figures 9 and 10). To discourage spurious downward jumps in velocity after entering the salt, we add the asymmetric TV constraint and use a continuation strategy in the ξ parameter. We start with a small value for the 1-sided TV-norm level set, and gradually increase it over each successive pass through the frequency batches. This continuation approach encourages the initial velocity

estimates to be nearly monotonically increasing in depth, a notion that corresponds to assuming we are in a sedimentary basin where we can step into the salt but not out of it. At later passes, the asymmetric TV constraint is relaxed to allow more downward jumps to fit the observed data. Starting with the poor initial model in Figure 12a, Figure 11 shows the progression of velocity estimates over eight passes. The sequence of ξ parameters as a fraction of $\|\max(0, D_z m_{\text{true}})\|_1$ is chosen to be $\{.01, .05, .10, .15, .20, .25, .40, .90\}$. We keep the τ parameter fixed at $.9\tau_{\text{true}}$ throughout. Although small values of ξ cause some vertical artifacts, the continuation strategy is surprisingly effective at preventing the method from getting stuck at a poor solution. As ξ increases, the bottom of the salt is recovered.

Both the model error and data misfit continue to decrease during each pass, as shown in Figure 13. The rugged shape of the objective function curve has two different scales that correspond to the relaxation passes and the increasing frequency content. The former causes the sudden drops in the functional every 9 iterations, which occurs because the data fit improves as we relax the TV constraints: due to the extra structure allowed in the model as we relax the constraints, the data-misfit reduces as more events can be generated and matched. The second, the objective function increase within each low to high pass through frequency batches is explained by the increase in frequency: there is more energy present at higher frequencies in the data and therefore the data-misfit increases accordingly.

Even with a poor initial model, the asymmetric TV constraint with continuation is able to recover the main features of the ground truth model. Poor recovery near the left and right boundaries is expected because the sources along the surface start about 1000 m away from these boundaries, which affects the illumination.

Conventional FWI approach

So far, all examples we discuss were based on WRI (29). Recall that unconstrained FWI was unable to produce meaningful results when given a poor starting model such as the one plotted in Figure 8a. However, if we impose the combination of the box, TV and asymmetric constraints in Algorithm 1, the conventional reduced adjoint-state method (cf. (25)) produces excellent results as long as we replace $H^n + c_n \mathbf{I}$ $c_n(H^n + \nu \mathbf{I})$ for some small positive ν and where H^n is defined by (28).




Fig. 11: Initial velocity (a) and recovered velocity with asymmetric TV continuation corresponding to $\frac{\xi}{\xi_{\text{true}}} = .01$ (a), .05 (b), .10 (c), .15 (d), .20 (e), .25 (f), .40 (g), .90 (h). Movies of the solutions with and without the asymmetric TV norm are available as ancillary material.

We use the same continuation strategy as was used to generate the results in Figure 11, and the results are nearly as good (see Figure 12). Compared to the WRI method, the

results are visually slightly worse near the top and sides of the model. Additionally, WRI finds a significantly better model error relative to ground truth (compare Figure 13b with Figure 14b), but it is encouraging to see that once again the error continues to decrease during each pass instead of stagnating at a poor solution.

In summary, the examples clearly demonstrate that continuation in the ξ parameter for the asymmetric TV constraint appears to be a promising strategy for preventing both the constrained WRI and adjoint-state FWI from stagnating in a poor solution when starting from a bad starting model.




Fig. 12: Initial velocity (a) and recovered velocity for the constrained adjoint state method with one-sided TV continuation corresponding to $\frac{\xi}{\xi_{\text{true}}} = .01$ (a), $.05$ (b), $.10$ (c), $.15$ (d), $.20$ (e), $.25$ (f), $.40$ (g), $.90$ (h).




Fig. 13: WRI objective function (data misfit) evolution (a). The model error —normalised RMS of the difference between the true and the current iterate— (b) increases initially, due to the effect the strict constraints coupled with the WRI updates have in the early iterations, but after the values are relaxed, we observe a descending trend that reduces the model misfit below 30%. The 72 iterations correspond to the cumulate effect of all the iterations for each frequency.




Fig. 14: FWI objective function (data misfit) evolution (a). In this case the model error —normalised RMS of the difference between the true and the current iterate— (b) behaves as expected initially but shows an increase after the first pass; due to the limitations of FWI (compared to WRI), the model error does not decrease as much as in the WRI experiment and stalls below 50%. Again, the 72 iterations correspond to the cumulate effect of all the iterations for each frequency.

VI. DISCUSSION

The work presented so far mainly concerned contributions from the first author Ernie Esser, who passed away preparing his work for publication. Our edits and additions were mostly stylistic and very much done in the spirit of Ernie’s original draft, which contained a section on "ongoing work". We took the liberty to rewrite and annotate this section to reflect the impact of Ernie’s contributions to seismic inversion for geologically challenging settings. Below, we give a summary of Ernie’s main points, along with brief descriptions of research progress Ernie’s work inspired.*

Ernie’s points

While the numerical experiments clearly demonstrate the potential benefits of including bound, TV- and asymmetric TV-norm constraints in wave-equation based inversions, the following issues remain:

- **Inversion crime.** The inversion results were all obtained from data generated with the same forward modelling kernel as the kernel used to invert the data. This practice, known as the *inversion crime*, is common when developing new methodologies. Ernie was aware of this and wrote “*The numerical examples should be recomputed without inversion crime.*” With the help of others, Ernie’s co-authors have redone the experiments and we are happy to reports that Ernie’s proposed method continues to perform well. These results have been presented at the 2016 EAGE meeting and will be included in a future publications.
- **Boundary conditions.** Ernie’s original work was in his words based on “*simplistic boundary conditions that may conspire to make the inversion problem easier.*” Ernie refers here to the fact that the inversions may inadvertently been helped by unnatural reflections emanating from the boundaries. These reflected waves could have illuminated bottom salt and the salt flanks. Again with assistance from others, we have been able to redo the experiments with a more sophisticated modeling including more accurate

*Available at <https://www.slim.eos.ubc.ca/content/total-variation-regularization-strategies-full-waveform-inversion-improving-robustness-noise>.

boundary conditions. The conclusion from these experiments is that simplistic boundary conditions had little to no effect on the final inversion results. Following Ernie’s suggestion to “*generate the data using more sophisticated modeling and boundary conditions while continuing to use the simple variants for inversion,*” we found that the improvements in the inversion results still stand.

- **Effect of randomness.** To accelerate the computations, Ernie used a source encoding technique where the required number of wave-equation solves is reduced via randomized projections. While this approach reduces the number of sources to only two, it creates noisy crosstalk. Quoting Ernie, “*To remove any effect of randomness, the examples should be recomputed using sequential shots instead of simultaneous shots with the random weights redrawn every model update. This will be more computationally expensive, but it’s doable and is not expected to significantly alter the results.*” We followed up on this suggestion and we found that working with all sources significantly improved the results (See [36]).
- **Continuation strategy.** Ernie wrote “*More practical methods for selecting the parameters and more principled ways of choosing continuation strategies are needed. Currently the τ and ξ parameters are chosen to be proportional to values corresponding to the true solution. Although the true solution clearly isn’t known in advance, it may still be reasonable to base the parameter choices on estimates of its total-variation (or one-sided TV). It would be even better to develop continuation strategies that don’t rely on any assumptions about the solution but that are still effective at regularizing early passes through the frequency batches to prevent the method from stagnating at poor solutions.*” This is still a topic of active research. However, our experience applying Ernie’s heuristic of relaxing the constraint in combination with warm restarts suggests to us that it is critical to relax the constraints slowly enough so that the algorithm steers free from the effects of parasitic local minima.
- **Convex subproblems.** With regard to imposing constraints on the model iterations Ernie pointed out “*There is a lot of room for improvement in the algorithm used to solve the convex subproblems. There are for example methods such as the one in [24] that have*

better theoretical rates of convergence and are straightforward to implement.”. We leave this suggestion, as well as extensions to 3D seismic models, for future work.

Geophysical impact

Inverting for velocity in sedimentary basins interspersed with high-velocity sharp-contrast salt bodies is perhaps one of the most challenging inversion problems in geophysics. Despite numerous efforts from academia as well as from industry, little progress has been made using fully automatic (i.e. void of extensive human interaction) methods to solve this inversion problem in a systematic and reproducible way. Instead the community relies on intricate workflows, which combine reflection tomography, manual *salt flooding* (where the salt is continued downwards into the basin), and full-waveform inversion. Quality of results depend rely on geophysical experience and manual intervention, and inversions are costly not reproducible.

To our knowledge, Ernie’s work is a first successful attempt to replace these labour intensive manual salt flooding workflows with an automatic heuristic: successive relaxations of the asymmetric TV-norm constraints with warm-started passes through frequency batches. In essence, Ernie encoded salt flooding [36] into a powerful automatic workflow, a breakthrough result.

Ernie’s continuation approach is able to steer free from parasitic stationary points. Numerical experience with the approach highlights the following important considerations:

- The lowest frequency should be sufficiently low to allow progress during each pass through the frequencies.
- The successively relaxed constraint sets should be large enough to prevent stalling (i.e. to allow model updates), but small enough to help avoid parasitic stationary points.

VII. CONCLUSIONS AND FUTURE WORK

We presented a scaled gradient projection algorithm for minimizing WRI and FWI formulations subject to additional convex constraints. We showed in particular how to solve the convex subproblems that arise when adding bound, total-variation (TV), and asymmetric TV constraints to the model space. The proposed framework is general, and the

convex constraints can be replaced or augmented by others, as long as the feasibility regions remain nonempty.

Synthetic experiments suggest that for sufficiently accurate starting models, TV constraints enhance the recovery by improving delineation of the salt and by eliminating spurious artifacts. The experiments also show that the asymmetric TV constraint, designed to encourage velocity to increase with depth, leads to major improvements in the recovery of salt structures from poor starting models. By pushing the bottom salt down into the basin in the beginning of the inversion procedure, the asymmetric constraint prevents the inversion from creating a velocity-low after the algorithm steps into the salt. In combination with a continuation strategy that gradually weakens the asymmetric TV constraint, we arrive at an approach that avoids getting stuck when starting with a poor initialization. Future work aims to study more realistic numerical experiments and investigate how to better take advantage of the proposed WRI framework.

VIII. ACKNOWLEDGEMENTS

Thanks to Bas Peters for insights about the formulation and implementation of the penalty method for full-waveform inversion, including the strategies for frequency continuation and computing multiple passes with warm starts. Thanks also to Polina Zheglova for helpful discussions about the boundary conditions.

IX. IN MEMORIAM: JOHN “ERNIE” ESSER[†] (MAY 19, 1980 – MARCH 8, 2015)

This paper is dedicated to its main author, Ernie Esser, who passed away in tragic circumstances during preparation of the manuscript. He was a very promising young scientist, positive, energetic, generous and talented, always a pleasure to work with. He is dearly missed by his family, friends and colleagues.



Fig. 15: Ernie Esser (May 19, 1980 – March 8, 2015). What Ernie (person on the left) loved most in life was hiking with friends, good espresso, and total-variation minimization. The above picture of Ernie with his long-time friend Justin Jacobs was taken on the Channel Islands in 2011. The total-variation constrained version of this picture was made by Ernie.

REFERENCES

- [1] A. Tarantola and A. Valette, “Generalized nonlinear inverse problems solved using the least squares criterion,” *Reviews of Geophysics and Space Physics*, vol. 20, no. 2, pp. 129–232, 1982.
- [2] A. Tarantola, “Inversion of seismic reflection data in the acoustic approximation,” *Geophysics*, vol. 49, no. 8, pp. 1259–1266, 1984. [Online]. Available: <http://link.springer.com/article/10.1007/s10107-012-0571-6><http://arxiv.org/abs/1110.0895><http://library.seg.org/doi/abs/10.1190/1.1441754>
- [3] W. W. Symes, “The seismic reflection inverse problem,” *Inverse Problems*, vol. 25, no. 12, p. 123008, dec 2009. [Online]. Available: <http://stacks.iop.org/0266-5611/25/i=12/a=123008?key=crossref.c0c22e66d22bb2cee818a94a2e839e74>
- [4] —, “Migration velocity analysis and waveform inversion,” *Geophysical Prospecting*, vol. 56, no. 6, pp. 765–790, nov 2008. [Online]. Available: <http://doi.wiley.com/10.1111/j.1365-2478.2008.00698.x>
- [5] B. Biondi and A. Almomin, “Tomographic full-waveform inversion (TFWI) by combining FWI and wave-equation migration velocity analysis,” *The Leading Edge*, vol. 32, no. 9, pp. 1074–1080, sep 2013. [Online]. Available: <http://library.seg.org/doi/abs/10.1190/tle32091074.1>
- [6] M. Warner and L. Guasch, “Adaptive waveform inversion-fwi without cycle skipping-theory,” in *76th EAGE Conference and Exhibition 2014*, 2014.
- [7] T. van Leeuwen and F. J. Herrmann, “Mitigating local minima in full-waveform inversion by expanding the search space,” *Geophysical Journal International*, vol. 195, no. 1, pp. 661–667, 2013. [Online]. Available: <http://gji.oxfordjournals.org/cgi/doi/10.1093/gji/ggt258>
- [8] —, “A penalty method for pde-constrained optimization in inverse problems,” *Inverse Problems*, vol. 32, no. 1, p. 015007, 2016. [Online]. Available: <http://stacks.iop.org/0266-5611/32/i=1/a=015007>
- [9] A. Y. Aravkin and T. van Leeuwen, “Estimating nuisance parameters in inverse problems,” *Inverse Problems*, vol. 28, no. 11, p. 115016, Jun. 2012. [Online]. Available: <http://arxiv.org/abs/1206.6532>
- [10] A. Y. Aravkin, D. Drusvyatskiy, and T. van Leeuwen, “Quadratic penalization through the variable projection technique,” *arXiv preprint arXiv:1606.02395*, 2016.
- [11] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific; 2nd edition, 1999. [Online]. Available: http://www.amazon.com/Nonlinear-Programming-Dimitri-P-Bertsekas/dp/1886529000/ref=sr_1_2?ie=UTF8&qid=1395180760&sr=8-2&keywords=nonlinear+programming
- [12] S. Bonettini, R. Zanella, and L. Zanni, “A scaled gradient projection method for constrained image deblurring,” *Inverse Problems*, vol. 25, no. 1, p. 015002, Jan. 2009. [Online]. Available: <http://stacks.iop.org/0266-5611/25/i=1/a=015002?key=crossref.31896d5825318968df4855dada3c1552>
- [13] M. Schmidt, D. Kim, and S. Sra, “11 Projected Newton-type Methods in Machine Learning,” *Optimization for Machine Learning*, no. 1, 2012. [Online]. Available: http://books.google.com/books?hl=en&lr=&id=JPQx7s2L1A8C&oi=fnd&pg=PA305&dq=Projected+Newton-type+Methods+in+Machine+Learning&ots=vccfyhm5Jb&sig=1JHG-pEdkJv_Yur0SVJQ8j_bw6A<http://books.google.com/books?hl=en&lr=&id=JPQx7s2L1A8C&oi=fnd&pg=PA305&dq=11+Projected+Newton-type+Methods+in+Machine+Learning&ots=vccfyhm8Bg&sig=zz6vgxjlcGhjMDFgaDnf774W0xU>

- [14] E. Esser, Y. Lou, and J. Xin, “A Method for Finding Structured Sparse Solutions to Nonnegative Least Squares Problems with Applications,” *SIAM Journal on Imaging Sciences*, vol. 6, no. 4, pp. 2010–2046, Jan. 2013. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/13090540X>
- [15] L. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: Nonlinear Phenomena*, vol. 60, pp. 259–268, 1992. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/016727899290242F>
- [16] E. T. Chung, T. F. Chan, and X.-C. Tai, “Electrical impedance tomography using level set representation and total variational regularization,” *Journal of Computational Physics*, vol. 205, no. 1, pp. 357–372, May 2005. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0021999104004711>
- [17] V. Akcelik, G. Biros, and O. Ghattas, “Parallel multiscale Gauss-Newton-Krylov methods for inverse wave propagation,” *Supercomputing, ACM/IEEE ...*, vol. 00, no. c, pp. 1–15, 2002. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1592877
- [18] A. Guitton, “Blocky regularization schemes for full-waveform inversion,” *Geophysical Prospecting*, vol. 60, no. 5, pp. 870–884, 2012. [Online]. Available: <http://dx.doi.org/10.1111/j.1365-2478.2012.01025.x>
- [19] Z. Guo and M. de Hoop, “SHAPE OPTIMIZATION IN FULL WAVEFORM INVERSION WITH SPARSE BLOCKY MODEL REPRESENTATIONS,” *gmig.math.purdue.edu*, vol. 1, pp. 189–208, 2012. [Online]. Available: <http://gmig.math.purdue.edu/pdfs/2012/12-12.pdf>
- [20] M. Maharramov and B. Biondi, “Robust joint full-waveform inversion of time-lapse seismic data sets with total-variation regularization,” *arXiv preprint arXiv:1408.0645*, 2014. [Online]. Available: <http://arxiv.org/abs/1408.0645>
- [21] A. Bourgeois, M. Bourget, P. Lailly, M. Poulet, P. Ricarte, and R. Versteeg, *The Marmousi experience*. EAGE, 1991, vol. 5-9, ch. Marmousi data and model.
- [22] M. Zhu and T. Chan, “An Efficient Primal-Dual Hybrid Gradient Algorithm For Total Variation Image Restoration,” *UCLA CAM Report [08-34]*, no. 1, pp. 1–29, 2008.
- [23] E. Esser, X. Zhang, and T. F. Chan, “A General Framework for a Class of First Order Primal-Dual Algorithms for Convex Optimization in Imaging Science,” *SIAM Journal on Imaging Sciences*, vol. 3, no. 4, pp. 1015–1046, Jan. 2010. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/09076934X>
- [24] A. Chambolle and T. Pock, “A first-order primal-dual algorithm for convex problems with applications to imaging,” *Journal of Mathematical Imaging and Vision*, 2011. [Online]. Available: <http://link.springer.com/article/10.1007/s10851-010-0251-1>
- [25] B. He and X. Yuan, “Convergence analysis of primal-dual algorithms for a saddle-point problem: from contraction perspective,” *SIAM Journal on Imaging Sciences*, pp. 1–35, 2012. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/100814494>
- [26] X. Zhang, M. Burger, and S. Osher, “A Unified Primal-Dual Algorithm Framework Based on Bregman Iteration,” *Journal of Scientific Computing*, vol. 46, no. 1, pp. 20–46, Aug. 2010. [Online]. Available: <http://link.springer.com/10.1007/s10915-010-9408-8>
- [27] P. Brucker, “An $O(n)$ algorithm for quadratic knapsack problems,” *Operations Research Letters*, vol. 3, no. 3, pp. 163–166, 1984. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0167637784900105>

- [28] T. Goldstein, E. Esser, and R. Baraniuk, “Adaptive primal-dual hybrid gradient methods for saddle-point problems,” *arXiv preprint arXiv:1305.0546*, 2013. [Online]. Available: <http://arxiv.org/abs/1305.0546>
- [29] C. Bunks, F. M. Saleck, S. Zaleski, and G. Chavent, “Multiscale seismic waveform inversion,” *GEOPHYSICS*, vol. 60, no. 5, pp. 1457–1473, sep 1995. [Online]. Available: <http://library.seg.org/doi/abs/10.1190/1.1443880>
- [30] F. Billette and S. Brandsberg-Dahl, “The 2004 BP velocity benchmark,” *67th EAGE Conference & Exhibition*, no. June, pp. 13–16, 2005. [Online]. Available: <http://www.earthdoc.org/publication/publicationdetails/?publication=1404>
- [31] J. R. Krebs, J. E. Anderson, D. Hinkley, A. Baumstein, S. Lee, R. Neelamani, and M.-D. Lacasse, “Fast full wave seismic inversion using source encoding,” *Geophysics*, vol. 28, no. 1, pp. 2273–2277, 2009. [Online]. Available: <http://link.aip.org/link/SEGEB/v28/i1/p2273/s1?Agg=doihttp://library.seg.org/vsearch/servlet/VerityServlet?KEY=SEGLIB&smode=stresults&sort=chron&maxdisp=25&threshold=0&pjournals=SAGEEP,JEEGXX&pjournals=JEEGXX&pjournals=SAGEEP&pjournals=GPYSA7,LEEDFF,SEGEB,SEGBKS&pjournals=GPYSA7&pjournals=LEEDFF&pjournals=SEGEB&pjournals=SEGBKS&possible1zone=article&possible4=krebs&possible4zone=author&bool4=and&OUTLOG=NO&viewabs=SEGEB&key=DISPLAY&docID=5&page=1&chapter=0>
- [32] E. Haber, M. Chung, and F. Herrmann, “An Effective Method for Parameter Estimation with PDE Constraints with Multiple Right-Hand Sides,” *SIAM Journal on Optimization*, vol. 22, no. 3, pp. 739–757, jul 2012. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/11081126X>
- [33] M. P. Friedlander and M. Schmidt, “Hybrid Deterministic-Stochastic Methods for Data Fitting,” *SIAM Journal on Scientific Computing*, vol. 34, no. 3, pp. A1380–A1405, may 2012. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/110830629>
- [34] T. van Leeuwen and F. J. Herrmann, “Fast waveform inversion without source-encoding,” *Geophysical Prospecting*, vol. 61, no. 2010, pp. 10–19, Jun. 2013. [Online]. Available: <http://doi.wiley.com/10.1111/j.1365-2478.2012.01096.x>
- [35] B. Peters, F. J. Herrmann, and T. van Leeuwen, “Wave-equation based inversion with the penalty method: adjoint-state versus wavefield-reconstruction inversion,” in *EAGE Annual Conference Proceedings*, 06 2014. [Online]. Available: <https://www.slim.eos.ubc.ca/Publications/Public/Conferences/EAGE/2014/peters2014EAGWeb/peters2014EAGWeb.pdf>
- [36] E. Esser, L. Guasch, F. J. Herrmann, and M. Warner, “Constrained waveform inversion for automatic salt flooding,” *The Leading Edge*, vol. 35, no. 3, pp. 235–239, 2016. [Online]. Available: <http://dx.doi.org/10.1190/tle35030235.1>
- [37] E. Haber, U. M. Ascher, and D. Oldenburg, “On optimization techniques for solving nonlinear inverse problems,” *Inverse Problems*, vol. 16, no. 5, pp. 1263–1280, oct 2000. [Online]. Available: <http://stacks.iop.org/0266-5611/16/i=5/a=309?key=crossref.98f435f9ee66231b63da02b10f82a60b>
- [38] J. Virieux and S. Operto, “An overview of full-waveform inversion in exploration geophysics,” *Geophysics*, vol. 74, no. 6, pp. WCC1–WCC26, Nov. 2009. [Online]. Available: <http://library.seg.org/doi/abs/10.1190/1.3238367>
- [39] M. Fisher, M. Leutbecher, and G. A. Kelly, “On the equivalence between kalman smoothing and weak-constraint four-dimensional variational data assimilation,” *Quarterly Journal of the Royal Meteorological Society*, vol. 131, no. 613, pp. 3235–3246, 2005. [Online]. Available: <http://dx.doi.org/10.1256/qj.04.142>

X. APPENDIX: PDE-CONSTRAINED OPTIMIZATION

We develop gradients and Hessian approximations for objectives defined in (1) and (2). These computations also appear in [37], [38].

Gradient and Hessian for the reduced adjoint-state method

The most commonly used formulation for FWI is obtained by eliminating the PDE-constraint, in which case the forward modelling operator can be explicitly written as

$$F(m)q \equiv Pu(m),$$

with P the sampling operator that models the measurement process. In this formulation, $u(m)$ is the discrete wavefield, computed by solving the discretized PDE

$$A(m)u = q.$$

After solving this system, the objective defined in (1) becomes

$$f(m) = \sum_{j=1}^{N_s} \frac{1}{2} \|PA(m)^{-1}q_j - d_j\|_2^2. \quad (25)$$

Derivatives with respect to the i^{th} entry of the model vector m are given by

$$\frac{\partial f}{\partial m_i} = \sum_{j=1}^{N_s} \left(\frac{\partial PA(m)^{-1}q_j}{\partial m_i} \right)^T (PA(m)^{-1}q_j - d_j) \quad (26)$$

$$= \sum_{j=1}^{N_s} q_j^T A(m)^{-T} \left(\frac{\partial A(m)}{\partial m_i} \right)^T A(m)^{-T} P^T (PA(m)^{-1}q_j - d_j) \quad (27)$$

with the symbol T denoting the adjoint. Introducing intermediate states

$$u_j = A(m)^{-1}q_j, \quad v_j = A(m)^{-T}P^T(Pu_j - d_j), \quad j = 1 \cdots N_s,$$

and the matrix

$$G_i(m) = \frac{\partial A(m)}{\partial m_i},$$

the derivative may be more succinctly written as

$$\frac{\partial f}{\partial m_i} = \sum_{j=1}^{N_s} u_j^T G_i(m)^T v_j.$$

The derivative formula requires solving both the forward PDE for each source term to obtain the state variables (wavefields) u_j , and the *adjoint* PDE to obtain adjoint state variables ("reverse-time" wavefields) v_j .

The full Hessian is dense and its elements are given by

$$\frac{\partial^2 f}{\partial m_k \partial m_l} = \sum_{j=1}^{N_s} u_j^T R_{kl}^T(m) v_j + u_j^T G_k(m)^T A^{-T} P^T P A^{-1} G_l(m) u_j,$$

with

$$R_{kl}(m) = \frac{\partial^2 A(m)}{\partial m_k \partial m_l} + 2G_k^T(m) A^{-1} G_l(m).$$

In practice, the individual elements of the Hessian are never computed for large-scale problems. Instead the action of the Hessian on a given vector can be evaluated at the cost of a few additional PDE-solves. While this approach avoids infeasible explicit storage of the Hessian, the additional PDE solves quickly become too expensive. To avoid these costs, the Hessian is often approximated by the *pseudo-Hessian* approximation, whose elements are given by

$$H_{kl} = \sum_{j=1}^{N_s} u_j^T G_k^T G_l u_j. \quad (28)$$

In our discretization, this approximation turns out to be diagonal, as G_k is a diagonal matrix with a single element at location (k, k) .

Gradient and Hessian for the extended formulation

While the adjoint-state method undergirds the majority of practical approaches to large-scale FWI, it requires relative accurate starting models to avoid getting stuck in local minima. Extended formulations, where the PDE-constraints are not eliminated but replaced by ℓ_2 -norm penalties, are less prone to these minima because they have more room to fit the data by optimizing over both the model m and state variables u_j , $j = 1 \cdots N_s$. In wave-equation based inversion, the extended formulation (2) has explicit form

$$f_\lambda(m) = \min_u \sum_{j=1}^{N_s} \frac{1}{2} \|P u_j - d_j\|_2^2 + \frac{\lambda^2}{2} \|A(m) u_j - q_j\|_2^2, \quad (29)$$

where we have made the variable substitution $\Delta q_j \rightarrow A(m) u_j - q_j$. Note the objective is written as an optimal value function of m . Evaluating f_λ for a given m requires solving an

optimization problem in u . However, since both terms are quadratic in the u_j 's, a closed-form solution is available:

$$u_{\lambda,j} = \left(A^T A + \lambda^{-2} P^T P \right)^{-1} \left(A^T q_j + \lambda^{-2} P^T d_j \right).$$

This expression can be seen as the equivalent of the PDE-solve needed to evaluate the regular objective but with the important distinction that solutions also aim to fit observed data. Approaches where PDE solves are combined with data-fit objectives are widely known in the field of data assimilation where PDEs appear as ℓ_2 -norm penalties, also known as "weak constraints" [39]. Indeed, as $\lambda \uparrow \infty$, we see that the expression reduces to $u_{\lambda,j} = A^{-1} q_j$. The derivatives of this objective are given by

$$\frac{\partial f_\lambda}{\partial m_k} = \sum_{j=1}^{N_s} u_{\lambda,j}^T G_k(m)^T v_{\lambda,j}, \quad (30)$$

and

$$\frac{\partial^2 f}{\partial m_k \partial m_l} = \lambda^2 \sum_{j=1}^{N_s} u_{\lambda,j}^T G_k^T G_l u_{\lambda,j} + u_{\lambda,j}^T R_{kl}^T v_{\lambda,j} - \left(A^T G_k u_{\lambda,j} + G_k^T v_{\lambda,j} \right)^T \left(\lambda^2 A^T A + P^T P \right)^{-1} \left(A^T G_l u_{\lambda,j} + G_l^T v_{\lambda,j} \right)$$

where

$$v_{\lambda,j} = \lambda^2 \left(A(m) u_{\lambda,j} - q_j \right).$$

Moreover, it can be shown that ∇f_λ is Lipschitz continuous, with the bound on the Lipschitz constant independent of λ [10], exactly as required by the implicit trust region framework. By ignoring higher order derivatives and the dependency of $u_{\lambda,i}$ on m , we obtain a positive definite approximation of this Hessian with elements

$$(H_\lambda)_{kl} = \lambda^2 \sum_{i=1}^{N_s} u_{\lambda,i}^T G_k^T G_l u_{\lambda,i}. \quad (31)$$

This approximation is typically sparse and does not involve additional PDE-solves.

Discretization

In the numerical experiments we use a finite-difference discretization of the Helmholtz operator with Robin boundary conditions, in which case A is block-diagonal matrix

$$A = \begin{pmatrix} A_1 & & & & \\ & A_2 & & & \\ & & \ddots & & \\ & & & & A_{N_f} \end{pmatrix}$$

with N_f blocks

$$A_i = \omega_i^2 \text{diag}(b) \text{diag}(m) - \omega_i \text{diag}(1 - b) \text{diag}(m^{1/2}) + L,$$

where ω is the angular frequency, b is a vector with $b_i = 1$ in the interior of the domain and $b_i = 0$ on the boundary and L is a 5-point discretization of the Laplace operator with Neumann boundary conditions. Both the wavefields u , v and the source vectors q are block-vectors with one block for each frequency as well. A single solve of the system $Au = q$ thus involves solving N_f systems of equations independently. The Jacobian matrix $G_k = \frac{\partial A}{\partial m_k}$ is a block-diagonal matrix with blocks

$$\frac{\partial A_i}{\partial m_k} = \omega_i^2 \text{diag}(b) \text{diag}(e_k) - \frac{1}{2} \omega_i \text{diag}(1 - b) \text{diag}(e_k) \text{diag}(m^{-1/2}),$$

where e_k is the k^{th} unit vector.