

Multi-parameter waveform inversion; exploiting the structure of penalty-methods

Bas Peters
June, 2014

Outline

- Challenging problem example (same as I presented on Tuesday) + extra details
- Waveform Reconstruction Inversion based multi-parameter algorithms

Wavefield Reconstruction Inversion [T. van Leeuwen & F.J. Herrmann, 2013]

Objective:

$$\bar{\phi}_\lambda(\mathbf{m}) = \frac{1}{2} \sum_{kl} \overset{\text{Data-misfit}}{\downarrow} \|P\bar{\mathbf{u}}_{kl} - \mathbf{d}_{kl}\|_2^2 + \frac{\lambda^2}{2} \overset{\text{PDE-misfit}}{\downarrow} \|H_k(\mathbf{m})\bar{\mathbf{u}}_{kl} - \mathbf{q}_{kl}\|_2^2$$

where $\bar{\mathbf{u}}_{kl} = \arg \min_{\mathbf{u}_{kl}} \left\| \begin{pmatrix} \lambda H_k(\mathbf{m}) \\ P \end{pmatrix} \mathbf{u}_{kl} - \begin{pmatrix} \lambda \mathbf{q}_{kl} \\ \mathbf{d}_{kl} \end{pmatrix} \right\|_2$

and λ is a tradeoff parameter between PDE-fit and data-fit

Wavefield Reconstruction Inversion [T. van Leeuwen & F.J. Herrmann, 2013]

Objective:

$$\bar{\phi}_\lambda(\mathbf{m}) = \frac{1}{2} \sum_{kl} \overset{\text{Data-misfit}}{\downarrow} \|P\bar{\mathbf{u}}_{kl} - \mathbf{d}_{kl}\|_2^2 + \frac{\lambda^2}{2} \overset{\text{PDE-misfit}}{\downarrow} \|H_k(\mathbf{m})\bar{\mathbf{u}}_{kl} - \mathbf{q}_{kl}\|_2^2$$

with gradient:

$$\nabla_{\mathbf{m}} \bar{\phi}_\lambda = \sum_{kl} \lambda^2 G_{kl}(\mathbf{m}, \bar{\mathbf{u}}_{kl})^* (H_k(\mathbf{m})\bar{\mathbf{u}}_{kl} - \mathbf{q}_{kl})$$

Non-linear waveform inversion

Example 1 (difficult):

- Lots of low frequencies missing, 24 frequency batches (15 iterations each) with intervals {5 6} ,{6 7},... ,{28 29} Hertz. Each interval contains 5 frequencies.
- We use 2 cycles through the batches: {5 6} ,{6 7},... ,{28 29}, {5 6} ,{6 7},... ,{28 29}
- Inaccurate initial model
- 103 sources and receivers near the surface, spread over the whole domain (6km).
- Source & receiver interval: 55m. Max. offset 6km.
- Shortest wavelength: 290m @ 5Hz. and 50m @ 29 Hz.
- Used Two-metric projection with L-BFGS Hessian for optimization with bound-constraints. [Bertsekas, 1982 ; Gafni & Bertsekas, 1982 ; Schmidt, Kim & Sra, 2009]

Why do we need more sophisticated algorithms?

- We need upper and lower bounds on the velocity for numerical reasons:
 - lower bound: sufficient number of grid points per wavelength
 - upper bound: very long wavelengths not absorbed by the PML
- Bounds can also incorporate a priori geological information
- Numerical experiments so far show that bound constraints are required for WRI.
- For pure gradient-descent, the model update can be projected onto the bounds (box).

Why do we need more sophisticated algorithms?


- Projecting l-BFGS/Newton updates onto the bounds is dangerous.
- The projected update may point in the opposite direction
- Projected-Newton methods exist, but can be computationally intensive.
- Multiple cheap solutions exist; the two-metric-projection (TMP) method is used here (with l-BFGS Hessian).
- TMP is a general concept, which can be used in conjunction with different Hessian approximations.

Two-metric-projection

- Project in a different norm than the one used for computing the update direction
- Split Hessian approximation in free and restricted variables


True, initial and final models

True model




after 1st frequency batch


Result FWI



Initial model




Result WRI, $\lambda = 1$



final result


True, initial and final models

True model




after 2nd frequency batch


Result FWI



Initial model




Result WRI, $\lambda = 1$



final result


True, initial and final models

True model




after last frequency batch


Result FWI



Initial model




Result WRI, $\lambda = 1$




final result

True, initial and final models

True model




Initial model




After 1st cycle

Result WRI, $\lambda = 1$




True, initial and final models

True model




Initial model




After 2nd cycle

Result WRI, $\lambda = 1$




True and final models

Result WRI, $\lambda = 1$




No noise

True model




True and final models

Result WRI with noise, $\lambda = 1$



$$\frac{\|\text{noise}\|}{\|\text{data}\|} = 0.04$$

True model




Initial phase-residuals

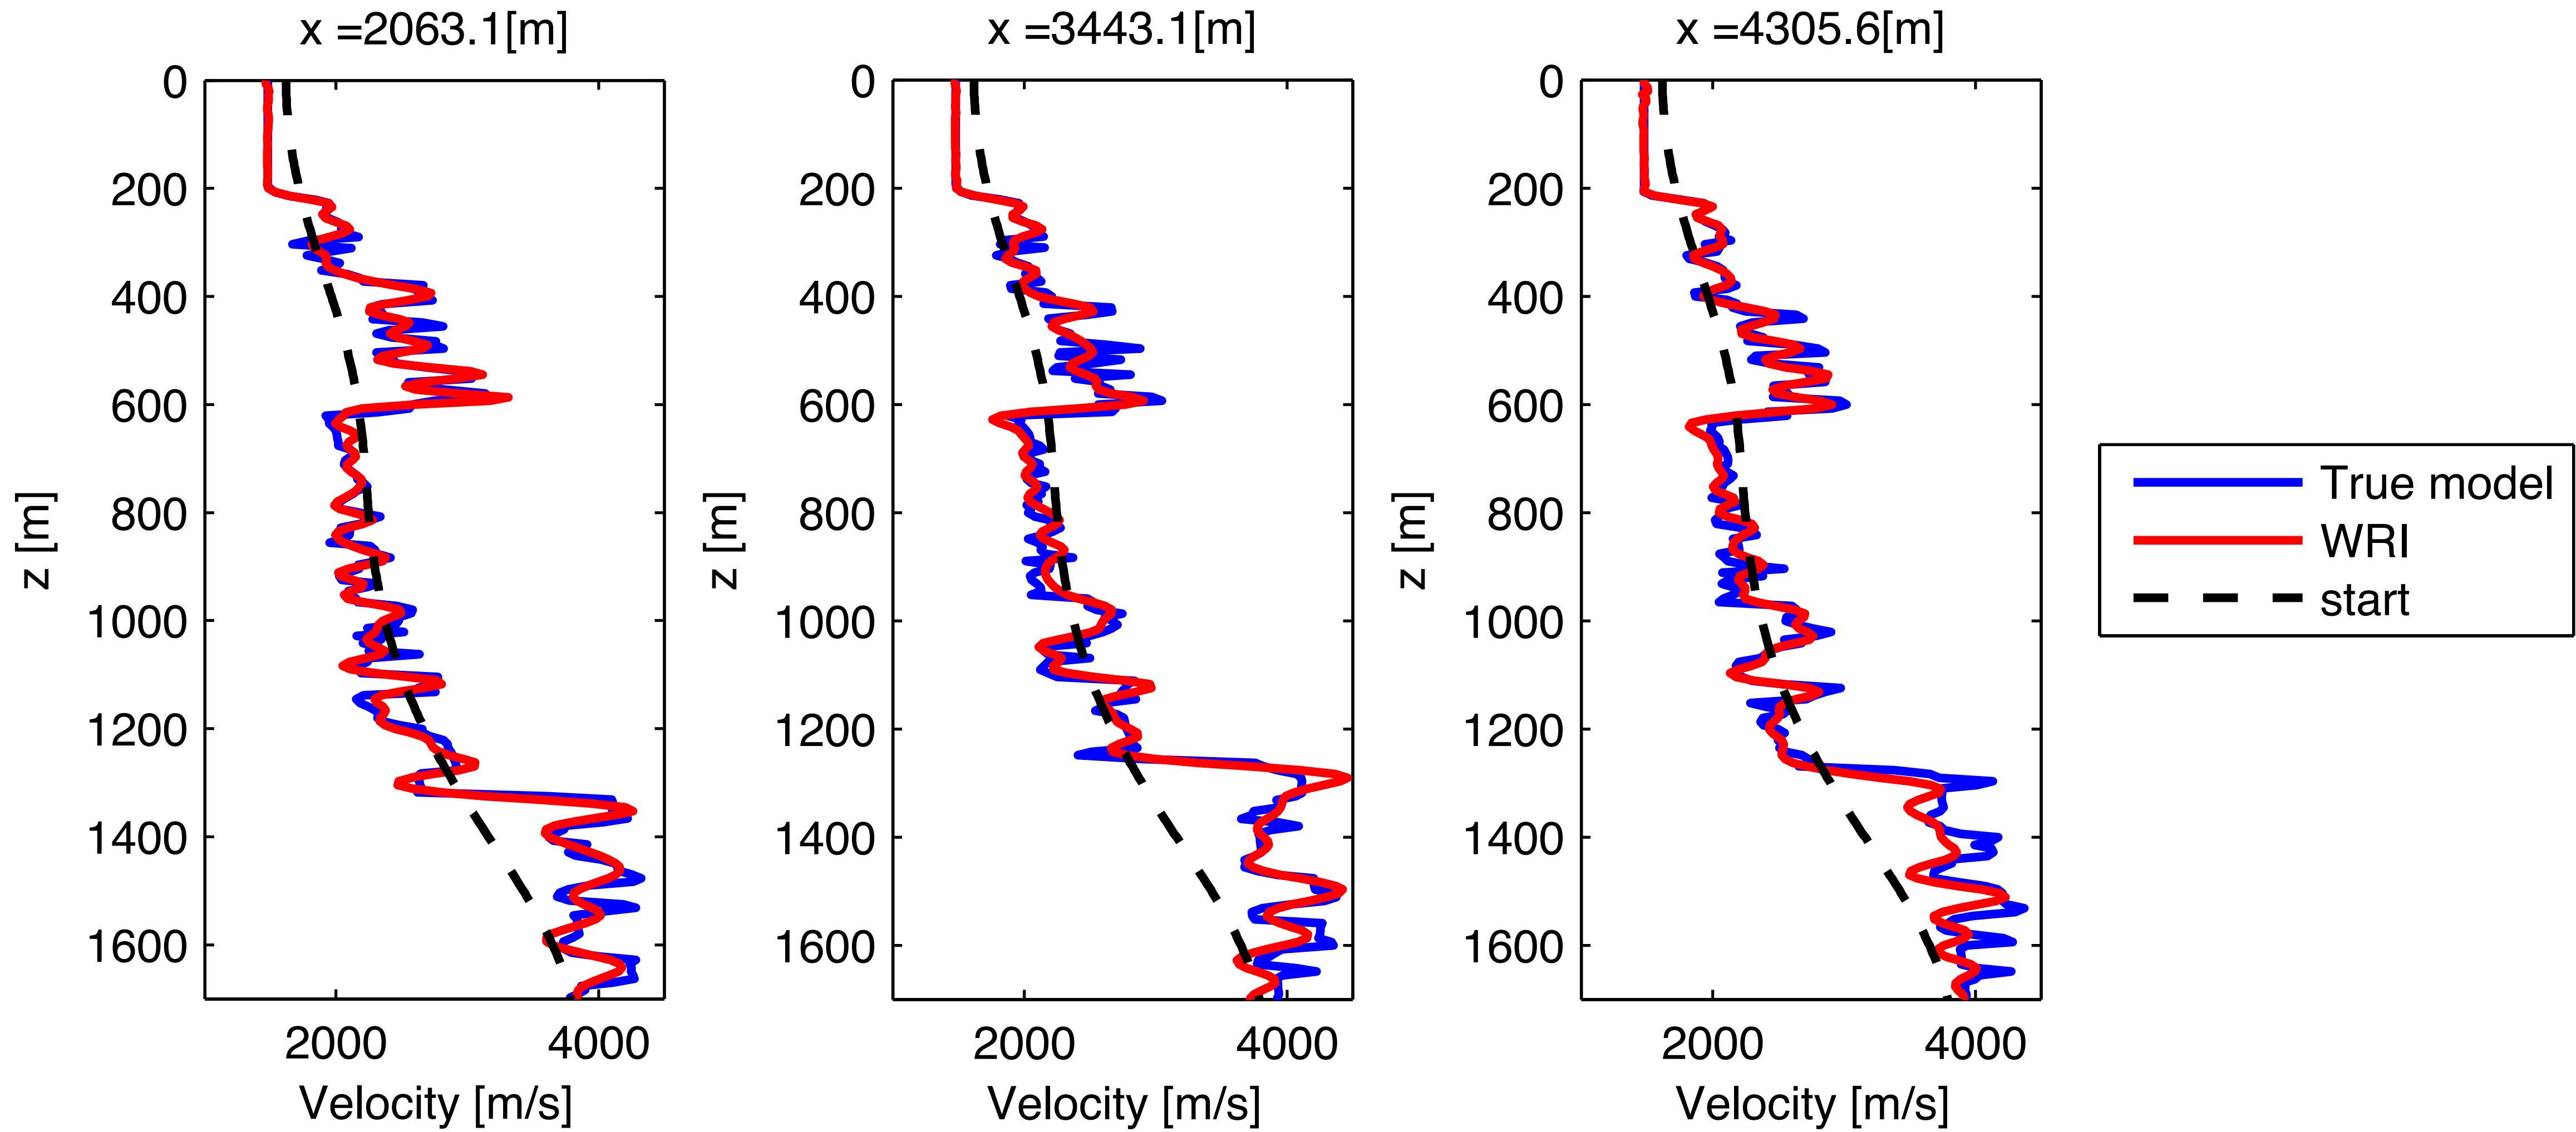
- Phase residuals computed using the Helmholtz equation in the start model
- WRI does not work with exact wavefields
- WRI uses the 'data-augmented' wavefield
- for λ very small, the phase residual will be 0.
- This will fit noise and cause ill-conditioning (not recommended!)
- Choose λ somewhat balanced (not difficult)

$$\bar{\mathbf{u}}_{kl} = \arg \min_{\mathbf{u}_{kl}} \left\| \begin{pmatrix} \lambda H_k(\mathbf{m}) \\ P \end{pmatrix} \mathbf{u}_{kl} - \begin{pmatrix} \lambda \mathbf{q}_{kl} \\ \mathbf{d}_{kl} \end{pmatrix} \right\|_2$$




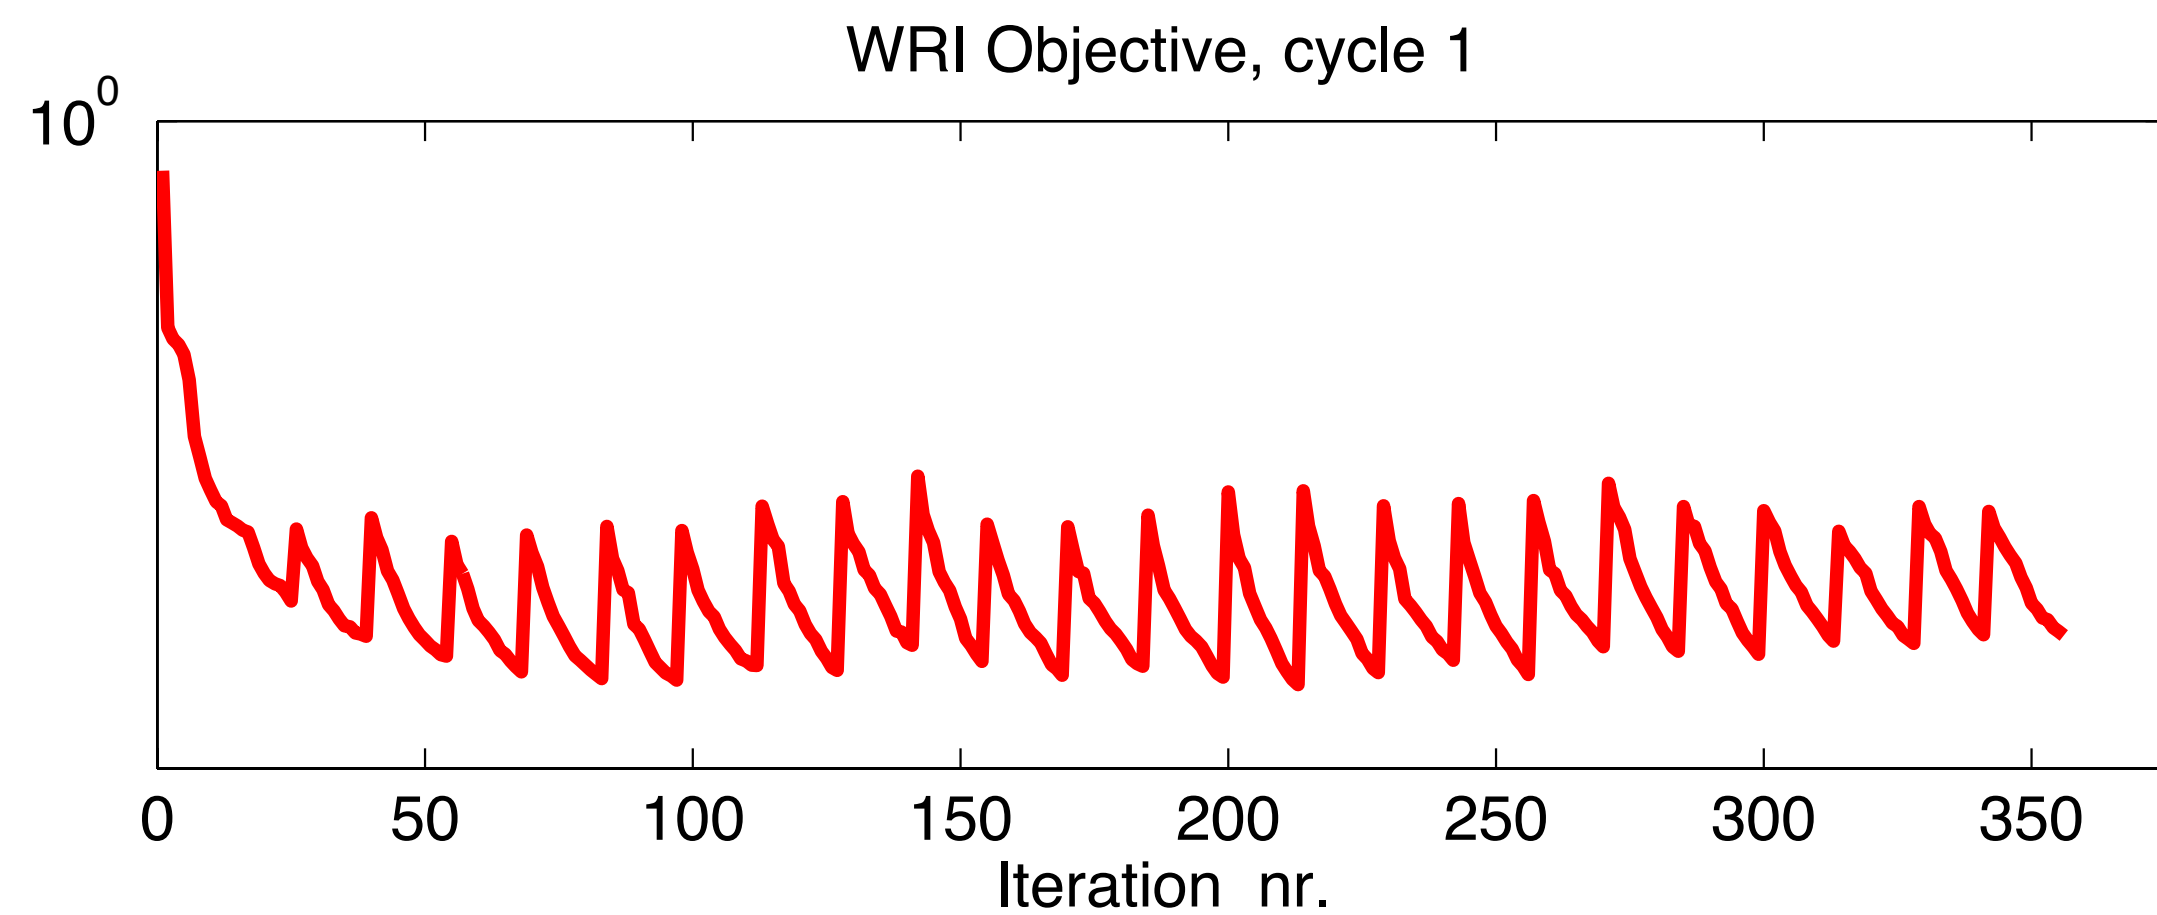
Phase residual in startmodel Ex2, 5 Hz.



Cross sections


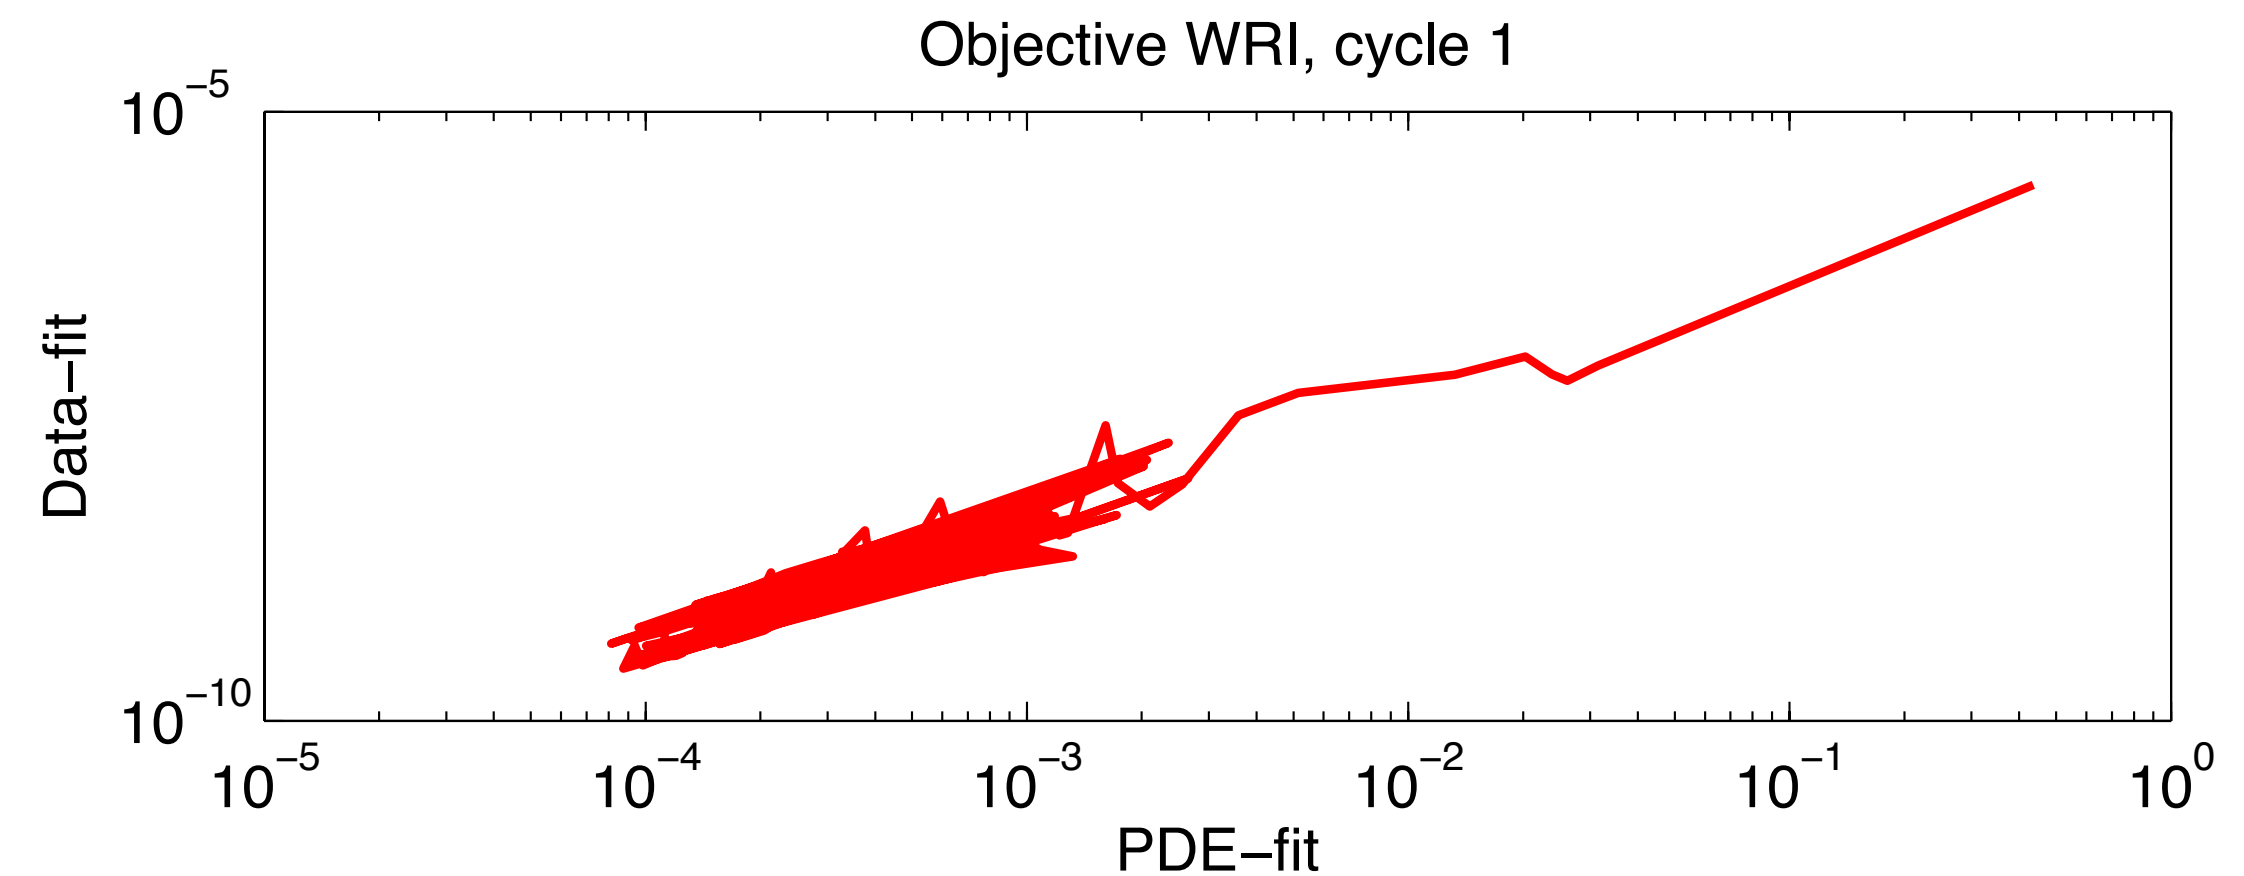


Objective and model error



Objective


$$\bar{\phi}_\lambda(\mathbf{m}) = \frac{1}{2} \sum_{kl} \overset{\text{Data-misfit}}{\downarrow} \|P\bar{\mathbf{u}}_{kl} - \mathbf{d}_{kl}\|_2^2 + \frac{\lambda^2}{2} \overset{\text{PDE-misfit}}{\downarrow} \|H_k(\mathbf{m})\bar{\mathbf{u}}_{kl} - \mathbf{q}_{kl}\|_2^2$$




- We can take a look at each part separately:
- Data-misfit can go up while iterating!

A look at the first update...


True model




First update FWI



Initial model




First update WRI, $\lambda = 1$




A look at the first update...


True model




First update FWI



Initial model




First update WRI, $\lambda = 1$



Different color scale


Model evolution

Cycle 1



Model evolution


Cycle 2




Effect of noise

Noise free

WRI Objective, cycle 1




WRI Objective, cycle 2

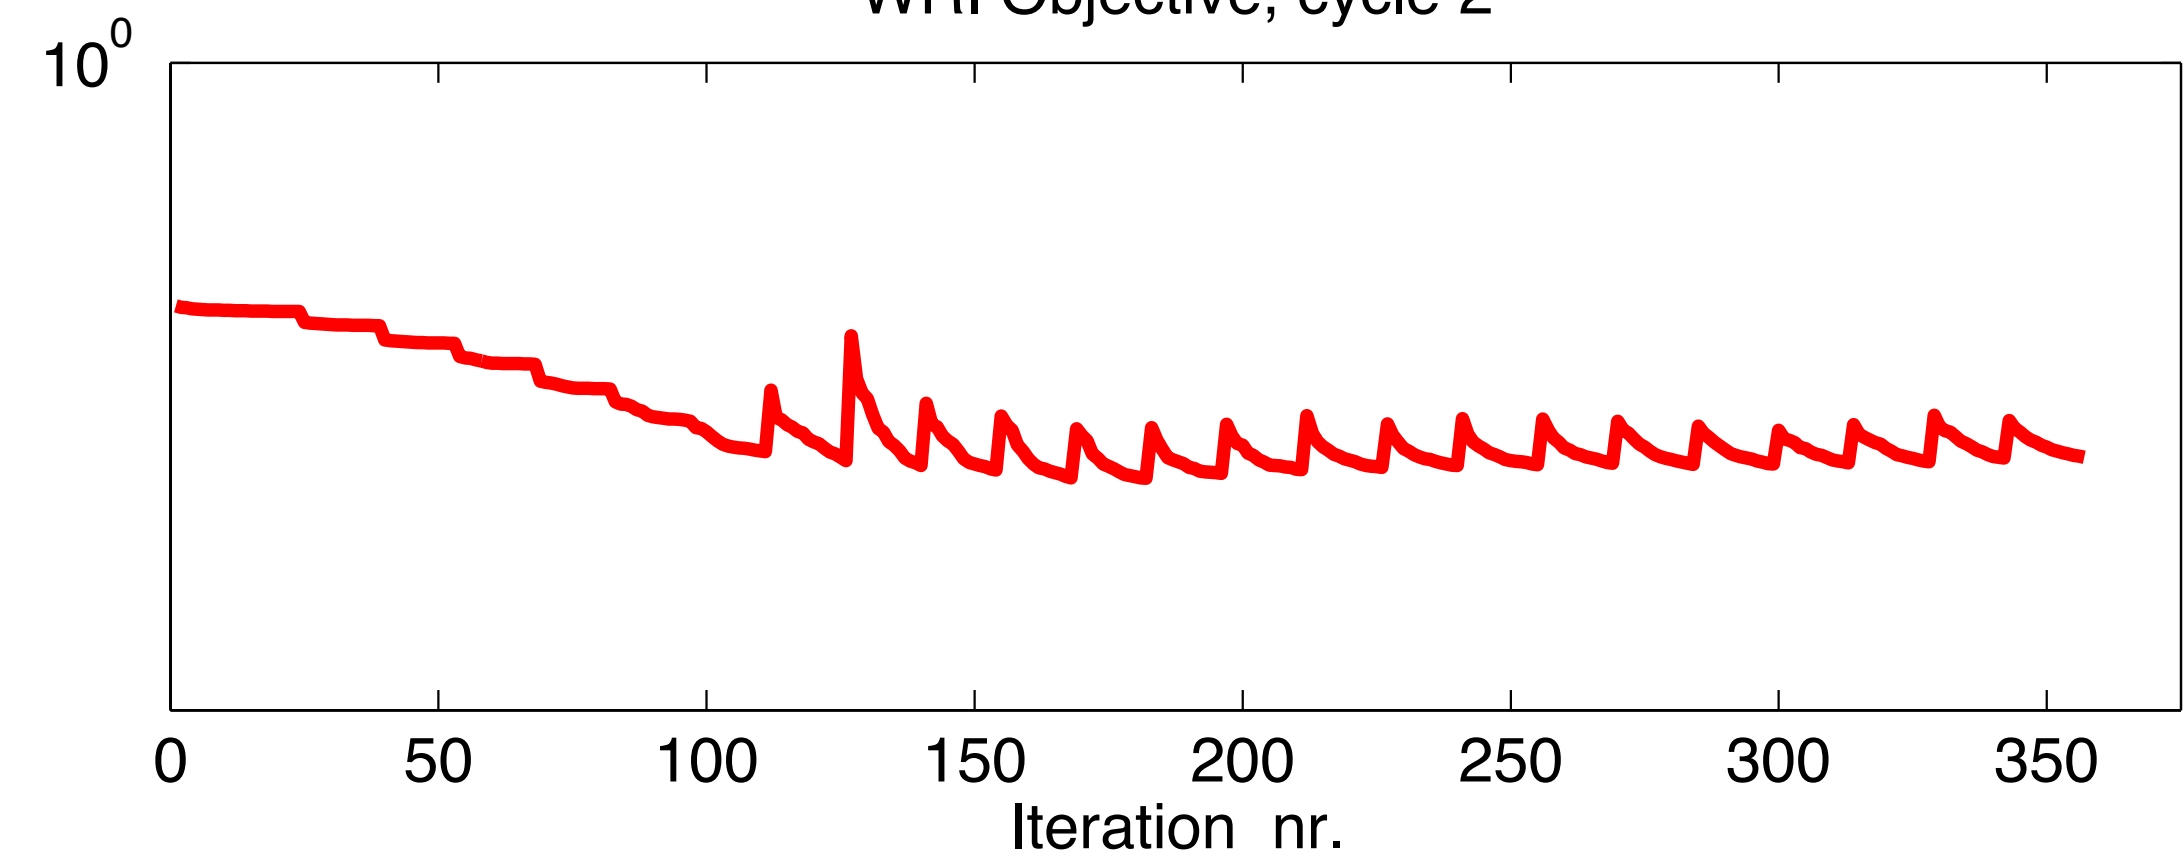


Noise

WRI Objective, cycle 1



WRI Objective, cycle 2



Observations about waveform inversion

- WRI performs much better for difficult problems
- WRI performs similar to FWI for not so difficult problems.
- Even for more difficult problems, only frequency continuation is required.
- No penalty parameter continuation was used, which can potentially increase quality and decrease the number of iterations.

Multi-parameter PDE-constrained optimization







$$\begin{aligned}\partial_k p + i\omega\rho v_k &= f_k \\ \partial_r v_r + i\omega\kappa p &= q\end{aligned}$$

- Difficulty 1: the compressibility and density occur in a coupled system of equations (or in 1 equation when the 2nd order form of the PDE is used)
- Difficulty 2: Different parameters can have very different magnitudes
- => Simple examples show >95.5% of the data can be fitted using only one of the two parameters.

Multi-parameter PDE-constrained optimization

Example: compute gradient of objective and naively plug into L-BFGS

Absolutely unsatisfactory result



Multi-parameter PDE-constrained optimization

- Multi-parameter problems of this sort have been solved with little generality
- Proposed solutions include:
 - sequential separate inversion
 - Structural similarity through regularization
 - Manual scaling of gradients corresponding to different parameters
 - Mitigate problems to limited extend by finding the ‘best’? parameterization
- All of these require fine-tuning and adaptation to different scenarios
- All of these were designed for gradient-based and quasi-Newton methods

Multi-parameter PDE-constrained optimization

- My approach: Use the actual Hessian
- The Hessian naturally provides information about ‘scaling’ and ‘coupling’.
- Multi-parameter problem:

$$\min_{\mathbf{b}, \boldsymbol{\kappa}, \mathbf{u}} \frac{1}{2} \|P\mathbf{u} - \mathbf{d}\|_2^2 \quad \text{s.t.} \quad H(\mathbf{b}, \boldsymbol{\kappa})\mathbf{u} = \mathbf{q}$$

- Penalty formulation:

$$\phi_\lambda(\mathbf{b}, \boldsymbol{\kappa}, \mathbf{u}) = \frac{1}{2} \|P\mathbf{u} - \mathbf{d}\|_2^2 + \frac{\lambda^2}{2} \|H(\mathbf{b}, \boldsymbol{\kappa})\mathbf{u} - \mathbf{q}\|_2^2 \quad [\text{T. van Leeuwen \& F.J. Herrmann, 2013}]$$

- Turns out to have additional benefits in this case

A reduced-space algorithm

$$\phi_\lambda(\mathbf{b}, \boldsymbol{\kappa}, \mathbf{u}) = \frac{1}{2} \|P\mathbf{u} - \mathbf{d}\|_2^2 + \frac{\lambda^2}{2} \|H(\mathbf{b}, \boldsymbol{\kappa})\mathbf{u} - \mathbf{q}\|_2^2$$

- Corresponding Newton-system (sparse)

$$\begin{pmatrix} \nabla_{\mathbf{u},\mathbf{u}}^2 \phi_\lambda & \nabla_{\mathbf{u},\boldsymbol{\kappa}}^2 \phi_\lambda & \nabla_{\mathbf{u},\mathbf{b}}^2 \phi_\lambda \\ \nabla_{\boldsymbol{\kappa},\mathbf{u}}^2 \phi_\lambda & \nabla_{\boldsymbol{\kappa},\boldsymbol{\kappa}}^2 \phi_\lambda & \nabla_{\boldsymbol{\kappa},\mathbf{b}}^2 \phi_\lambda \\ \nabla_{\mathbf{b},\mathbf{u}}^2 \phi_\lambda & \nabla_{\mathbf{b},\boldsymbol{\kappa}}^2 \phi_\lambda & \nabla_{\mathbf{b},\mathbf{b}}^2 \phi_\lambda \end{pmatrix} \begin{pmatrix} \delta_{\mathbf{u}} \\ \delta_{\boldsymbol{\kappa}} \\ \delta_{\mathbf{b}} \end{pmatrix} = \begin{pmatrix} \nabla_{\mathbf{u}} \phi_\lambda \\ \nabla_{\boldsymbol{\kappa}} \phi_\lambda \\ \nabla_{\mathbf{b}} \phi_\lambda \end{pmatrix}$$

write: $\begin{pmatrix} E & B \\ C & D \end{pmatrix}$

- When we solve for the fields as before, $(\nabla_{\mathbf{u}} \phi_\lambda(\bar{\mathbf{u}}, \boldsymbol{\kappa}, \mathbf{b}) = 0)$ we obtain

$$\begin{array}{ccc} \text{sparse} & \text{dense} & \\ \downarrow & \downarrow & \\ (D - CE^{-1}B) & \begin{pmatrix} \delta_{\boldsymbol{\kappa}} \\ \delta_{\mathbf{u}} \end{pmatrix} & = \begin{pmatrix} \nabla_{\boldsymbol{\kappa}} \phi_\lambda \\ \nabla_{\mathbf{b}} \phi_\lambda \end{pmatrix} - CE^{-1} \nabla_{\mathbf{u}} \phi_\lambda(\bar{\mathbf{u}}, \boldsymbol{\kappa}, \mathbf{b}) \end{array}$$

- Sparse-Dense splitting does not happen in the Lagrangian form.

A reduced-space algorithm

$$\begin{array}{ccc}
 \text{sparse} & \text{dense} & \\
 \downarrow & \downarrow & \\
 (D - CE^{-1}B) & \begin{pmatrix} \delta_{\kappa} \\ \delta_{\mathbf{u}} \end{pmatrix} = \begin{pmatrix} \nabla_{\kappa} \phi_{\lambda} \\ \nabla_{\mathbf{b}} \phi_{\lambda} \end{pmatrix} - CE^{-1} \nabla_{\mathbf{u}} \phi_{\lambda}(\bar{\mathbf{u}}, \kappa, \mathbf{b}) & = 0 \\
 \downarrow & & \\
 \begin{pmatrix} \nabla_{\kappa, \kappa}^2 \phi_{\lambda} & \nabla_{\kappa, \mathbf{b}}^2 \phi_{\lambda} \\ \nabla_{\mathbf{b}, \kappa}^2 \phi_{\lambda} & \nabla_{\mathbf{b}, \mathbf{b}}^2 \phi_{\lambda} \end{pmatrix} & \text{We can use this as an approximate Hessian} &
 \end{array}$$

So far I used only the diagonal blocks, yielding a sparse, SPD matrix

$$\tilde{H} = \begin{pmatrix} \nabla_{\kappa, \kappa}^2 \phi_{\lambda} & 0 \\ 0 & \nabla_{\mathbf{b}, \mathbf{b}}^2 \phi_{\lambda} \end{pmatrix} = \begin{pmatrix} G_{\kappa}^* G_{\kappa} & 0 \\ 0 & G_{\mathbf{b}}^* G_{\mathbf{b}} \end{pmatrix}$$

$$G_{\kappa} = \partial H(\mathbf{b}, \kappa) \bar{\mathbf{u}} / \partial \kappa \quad G_{\mathbf{b}} = \partial H(\mathbf{b}, \kappa) \bar{\mathbf{u}} / \partial \mathbf{b}$$

A reduced-space algorithm

Algorithm 1 Waveform inversion with a sparse Hessian approximation.

while Not converged **do**

$$1. \quad \bar{\mathbf{u}} = \arg \min_{\mathbf{u}} \left\| \begin{pmatrix} \lambda \mathbf{A}(\mathbf{b}, \kappa) \\ \mathbf{P} \end{pmatrix} \mathbf{u} - \begin{pmatrix} \lambda \mathbf{q} \\ \mathbf{d} \end{pmatrix} \right\|_2 \quad // \text{ Solve}$$

$$2. \quad \mathbf{G}_\kappa, \mathbf{G}_\mathbf{b}, \nabla_{\mathbf{b}} \bar{\phi}_\lambda, \nabla_\kappa \bar{\phi}_\lambda \quad // \text{ Form}$$

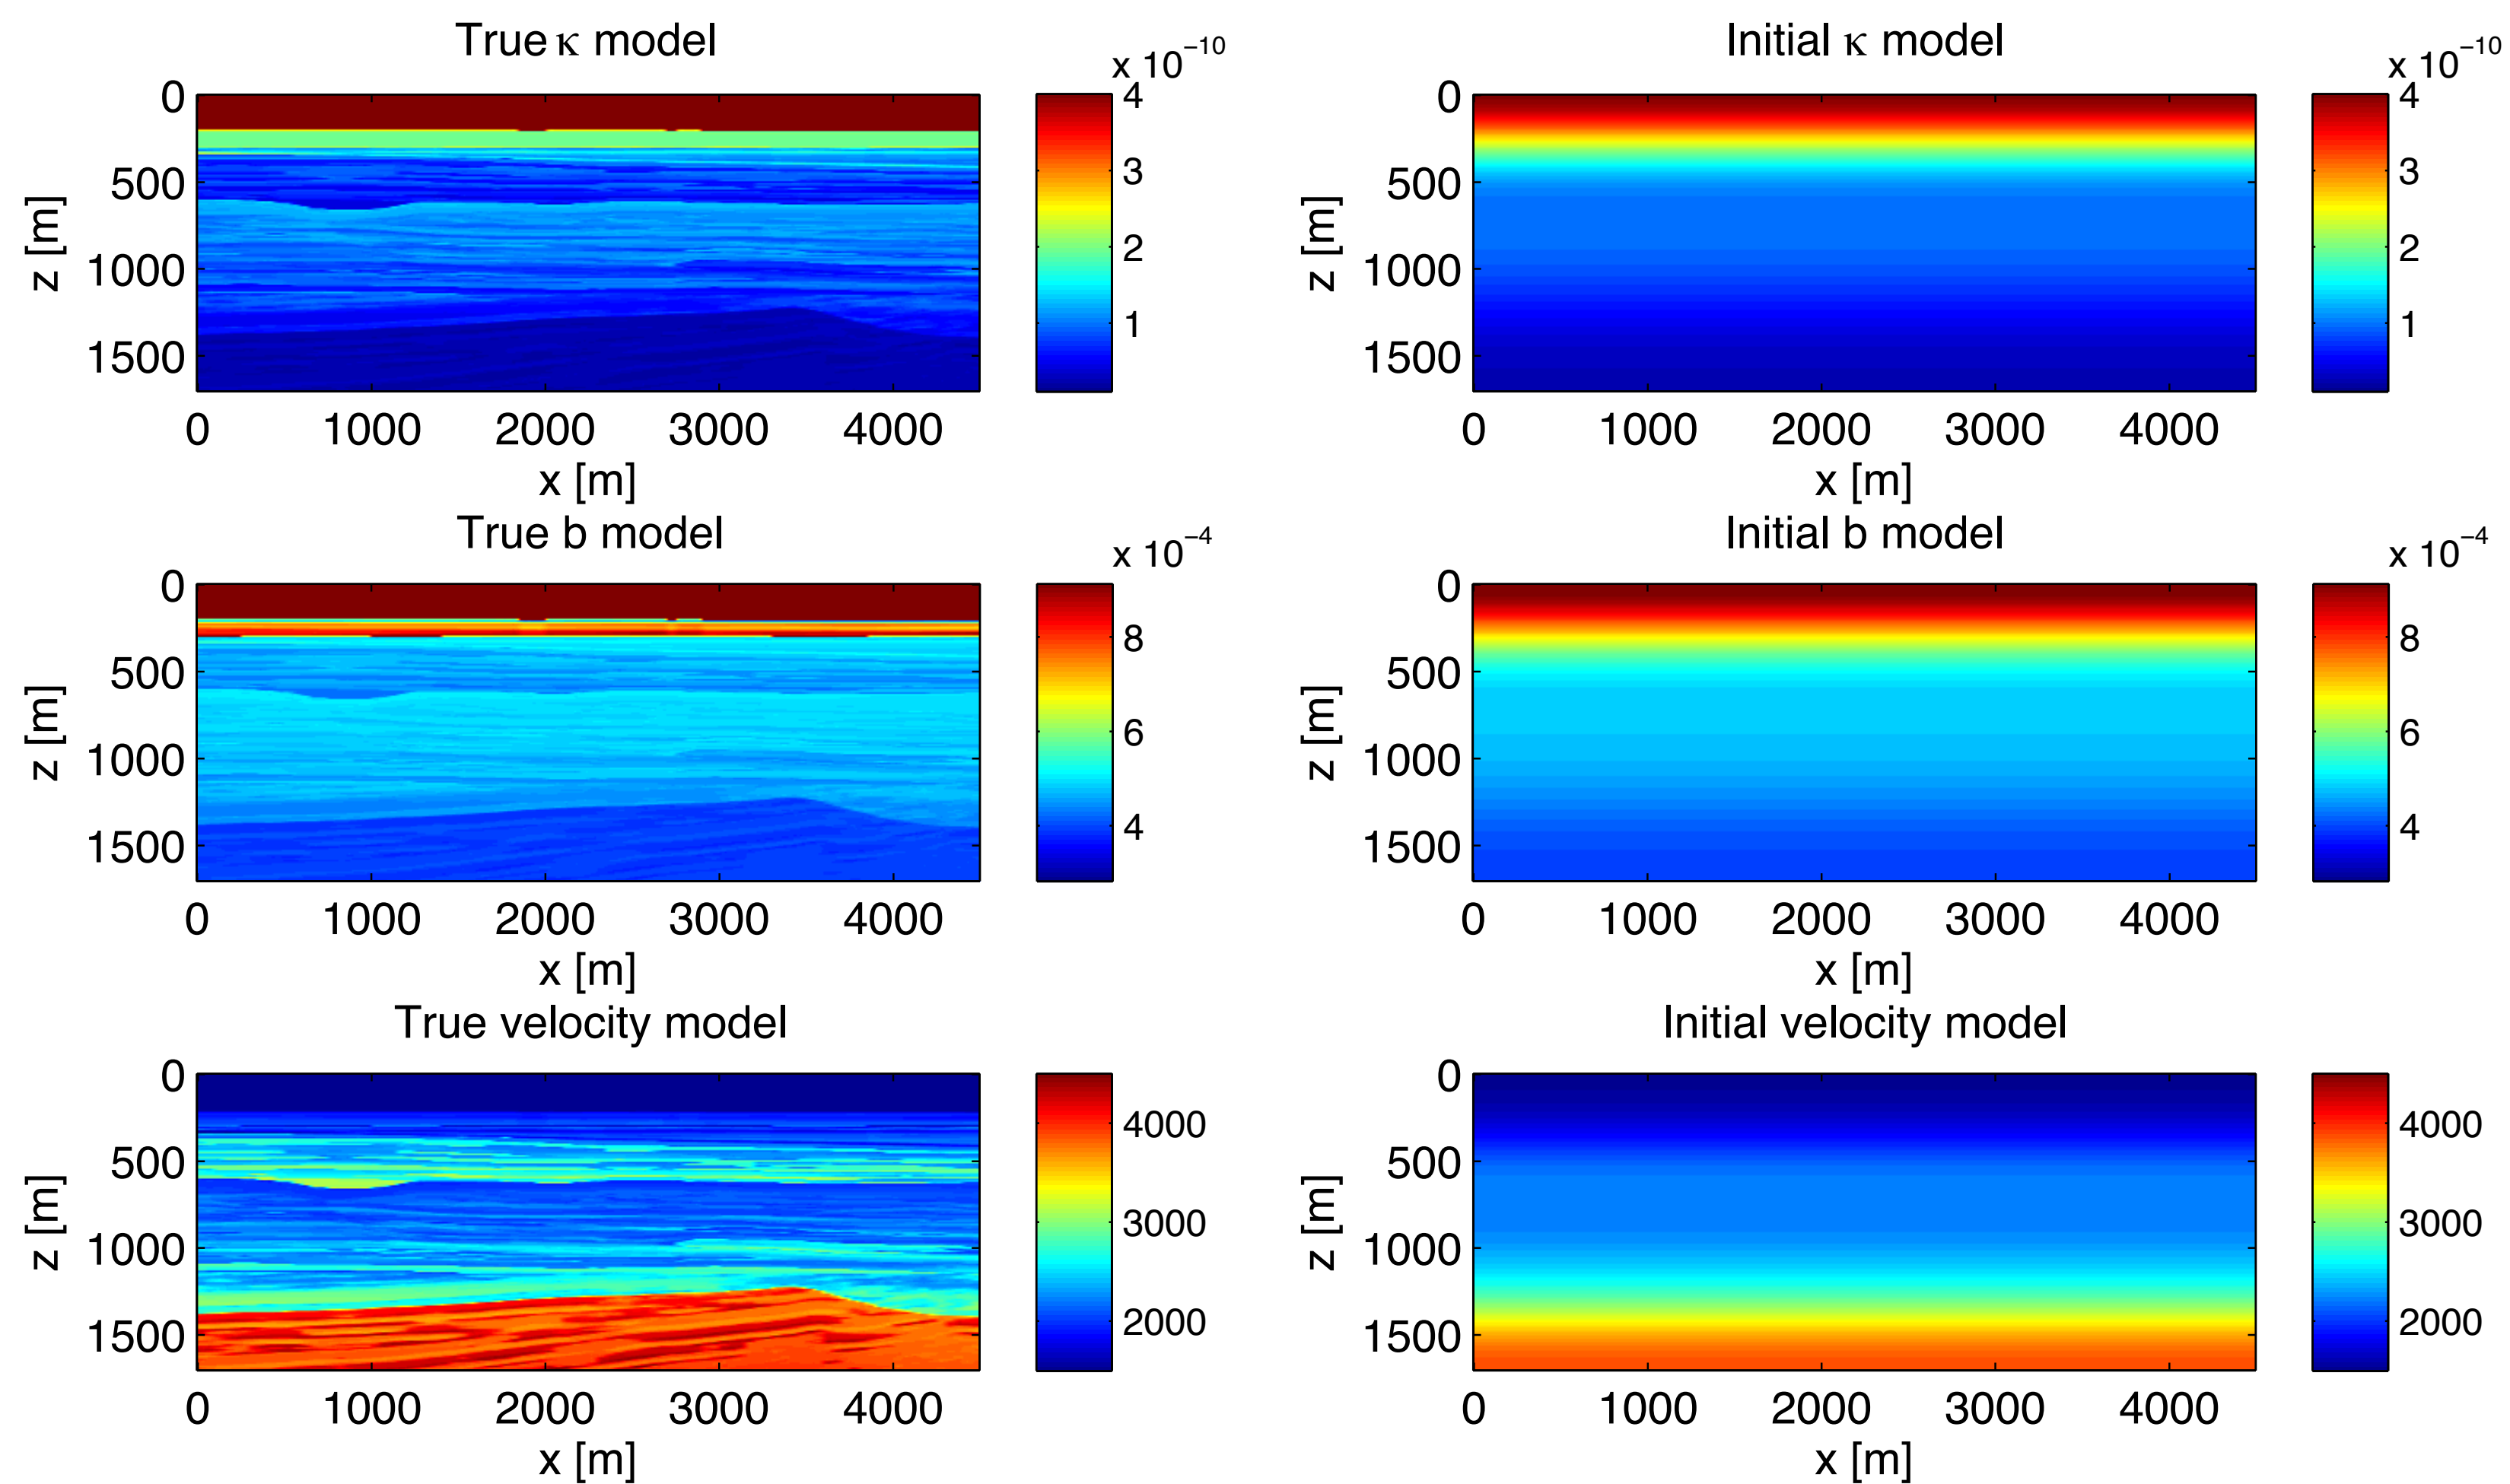
$$3. \quad \mathbf{p}_{gn} = \tilde{\mathbf{H}}^{-1} \mathbf{g} \quad // \text{ Solve}$$

$$4. \quad \text{find steplength } \alpha \quad // \text{ Linesearch}$$

$$5. \quad \mathbf{m} = \mathbf{m} + \alpha \mathbf{p}_{gn} \quad // \text{ update model}$$



end

A reduced-space algorithm; example







A reduced-space algorithm; example

New algorithm



basic L-BFGS



A reduced-space algorithm

- Sparse Hessian approximation seems to be quite effective
- Is it sufficient? Needs more experimenting...
- Alternatively, we can use a matrix-free Newton-CG algorithm including the dense part of the reduced Hessian.

$$\begin{array}{ccc}
 \text{sparse} & \text{dense} & \\
 \downarrow & \downarrow & \\
 (D - CE^{-1}B) & \begin{pmatrix} \delta_{\kappa} \\ \delta_{\mathbf{u}} \end{pmatrix} & = \begin{pmatrix} \nabla_{\kappa} \phi_{\lambda} \\ \nabla_{\mathbf{b}} \phi_{\lambda} \end{pmatrix} - CE^{-1} \nabla_{\mathbf{u}} \phi_{\lambda}(\bar{\mathbf{u}}, \kappa, \mathbf{b}) \\
 & & \begin{array}{c} =0 \\ \downarrow \end{array}
 \end{array}$$

Conclusions

- The use of the Hessian automatically scales the different parameter classes based on the PDE itself.
- Hessian information may turn multi-parameter inversion algorithms fully automatic.
- A quadratic-penalty formulation offers very cheap and effective Hessian approximations.
- Promising results using the multi-parameter WRI version.

Work in progress

- If you are willing to have more than 2 wavefields in memory, the door to many optimization algorithms is opened.
- For Lagrangian as well as penalty methods.
- Access to sparse Hessians, gradients without solving PDE's.
- Function evaluations do not require PDE solves either .
- Update the fields instead of solving.

$$\phi_\lambda(\mathbf{b}, \boldsymbol{\kappa}, \mathbf{u}) = \frac{1}{2} \|\mathbf{P}\mathbf{u} - \mathbf{d}\|_2^2 + \frac{\lambda^2}{2} \|H(\mathbf{b}, \boldsymbol{\kappa})\mathbf{u} - \mathbf{q}\|_2^2$$

Newton step \rightarrow

$$\begin{pmatrix} \nabla_{\mathbf{u}, \mathbf{u}}^2 \phi_\lambda & \nabla_{\mathbf{u}, \boldsymbol{\kappa}}^2 \phi_\lambda & \nabla_{\mathbf{u}, \mathbf{b}}^2 \phi_\lambda \\ \nabla_{\boldsymbol{\kappa}, \mathbf{u}}^2 \phi_\lambda & \nabla_{\boldsymbol{\kappa}, \boldsymbol{\kappa}}^2 \phi_\lambda & \nabla_{\boldsymbol{\kappa}, \mathbf{b}}^2 \phi_\lambda \\ \nabla_{\mathbf{b}, \mathbf{u}}^2 \phi_\lambda & \nabla_{\mathbf{b}, \boldsymbol{\kappa}}^2 \phi_\lambda & \nabla_{\mathbf{b}, \mathbf{b}}^2 \phi_\lambda \end{pmatrix} \begin{pmatrix} \delta_{\mathbf{u}} \\ \delta_{\boldsymbol{\kappa}} \\ \delta_{\mathbf{b}} \end{pmatrix} = \begin{pmatrix} \nabla_{\mathbf{u}} \phi_\lambda \\ \nabla_{\boldsymbol{\kappa}} \phi_\lambda \\ \nabla_{\mathbf{b}} \phi_\lambda \end{pmatrix}$$

Acknowledgements

The SLIM students & postdocs



This work was in part financially supported by the Natural Sciences and Engineering Research Council of Canada Discovery Grant (22R81254) and the Collaborative Research and Development Grant DNOISE II (375142-08). This research was carried out as part of the SINBAD II project with support from the following organizations: BG Group, BGP, BP, CGG, Chevron, ConocoPhillips, ION, Petrobras, PGS, Total SA, WesternGeco, and Woodside.