Introduction

Reverse time migration (RTM) is a popular seismic imaging technique. Unfortunately, the computational cost of RTM scales *linearly* with the number of shots used in the survey. This poses a considerable challenge for modern large-scale surveys where the number of shots can run into hundreds of thousands. One approach to mitigate this challenge is to first compute a smaller number of linear combinations (or *encodings*) of shot records, and then to perform RTM on the blended shot records. There are numerous ways to choose the shot encoding operator; see for example Romero et al. (2000); Godwin and Sava (2011); Schmidt et al. (2013). Although the shot-encoding approach can reduce RTM's cost significantly, the final image inevitably contains some level of cross-talk that is introduced by the encoding process.

Alternatively, *decimation* of the shots avoids this cross-talk issue. Indeed, Godwin and Sava (2011) shows that simple decimation can be an effective method for reducing the number of shots while still obtaining an acceptable quality image. However, one could easily conceive of scenarios where uniform decimation is not the best strategy. For instance, consider a velocity model where most of the "interesting" subsurface features are localized in a small region in the domain of interest. In this scenario, only a small subset of shot location choices result in proper illumination of the interesting features, and a naïve decimation could potentially miss this set of important shot locations.

In this paper, we introduce two adaptive, *data-driven* strategies for shot selection in RTM. Informally, given a budget parameter B, our high-level goal is to deduce the optimal subset of at most B shot locations that will produce an acceptable image of a *given* (user-specified) region in the subsurface. This problem is combinatorial and computationally intractable, and hence we propose two tractable heuristics based on *greedy* optimization techniques. We report the results of several numerical simulations on the SEG-EAGE Salt Model and demonstrate improvement over conventional decimation in terms of image quality, while incurring only a small overhead in terms of computation. Therefore, our method can be beneficial for a host of shot-based subsurface imaging techniques, e.g. full waveform inversion.

Theory and Methods

At a high level, a single RTM experiment involves three steps: (i) simulate a forward wave propagation to obtain the *source* wavefield; (ii) use seismic data and simulate a backward wave propagation to obtain the *receiver* wavefield; and (iii) apply an imaging condition (e.g., by correlating source and receiver wavefields), and stack across time-steps to obtain an image. The final image is estimated by averaging across multiple RTM experiments corresponding to different shots. Formally, denote the i^{th} source and receiver wavefields as $\mathbf{W}_S(x,y,z,t,i)$ and $\mathbf{W}_R(x,y,z,t,i)$. Then, the standard RTM imaging condition (Kaelin and Guitton (2006)) is

$$A(x,y,z) = \sum_{i=1}^{S}\sum_{t=1}^{T} \mathbf{W}_S(x,y,z,t,i)\mathbf{W}_R(x,y,z,t,i)\ . \tag{1}$$

The above computational process can become expensive for large subsurface models with many shots, so we desire to obtain an image of comparable quality with fewer shots. Suppose that we are given a *budget* B of shot migrations that we are allowed to perform. As discussed above, budget-constrained imaging approaches include shot encoding and uniform/random decimation. However, these approaches are *oblivious* to the input velocity model. In contrast, we propose two new approaches for *data-adaptive* shot selection: (i) embedding with coverage maximization, and (ii) subsampling with sparse recovery.

Embedding approach

Our first approach is inspired by the *subsurface embedding* idea discussed in Bear et al. (2000); Liu et al. (2005). Overall, this approach can be described as a three-stage procedure.

1. *Embedded shot migration:* In the first stage, we identify an "interesting" region Ω in the velocity model, embed a small number of *virtual* shots, $j = 1,\dots,J$ (where J is a small fixed number) distributed within the region Ω, and run a forward modeling with these embedded shots to obtain

corresponding virtual wavefields $\widetilde{\mathbf{W}}_S(x,y,z,t,j)$. Using these virtual wavefields, we calculate the *surface energy profile* (similar to illumination maps) $f : X \to \mathbb{R}^+$:

$$f(x,y) = \sum_{j=1}^{J}\sum_{t=1}^{T} \widetilde{\mathbf{W}}(x,y,0,t,j)^2 .$$

2. *Shot selection via submodular optimization:* The surface energy profile is an indication of the distribution of energy observed at the surface that would emanate from the region of interest. Therefore, it makes sense to focus on the shots on the surface which correspond to the peaks of the surface energy profile. The naïve approach is to select the locations of the B largest coefficients in $f(x,y)$. However, this may not be the best choice; typically, for small values of B this provides a set of B shot locations that are spatially coherent. However, each shot possesses a certain range of "coverage" and there is not much benefit from selecting multiple shots with similar coverage. Instead, our goal is to impose some amount of diversity in our shot selection.
To achieve this goal, we model the selection problem as a *submodular coverage maximization* problem. Let w denote the coverage parameter (i.e., higher w denotes a larger area of coverage.) We define the coverage profile $g : X \to \mathbb{R}^+$ and its extension to sets of shot locations as follows:

$$g(x,y) = \sum_{\|(x',y')-(x,y)\|_2 \le w} f(x',y') . \qquad g(\mathscr{S}) = \sum_{\substack{\|(x',y')-(x,y)\|_2 \le w \\ \text{for some } (x,y)\in\mathscr{S}}} f(x',y') .$$

The goal is to maximize $g(\mathscr{S})$, subject to the budget constraint $|\mathscr{S}| \le B$. The problem is known to be NP-hard, i.e. theoretical intractable, but a greedy algorithm that iteratively picks the next-best shot achieves an (essentially tight) approximation ratio of $1-1/e$.

3. *Migration:* Once we have picked the set of B shot locations, we perform normal RTM using the set of chosen locations $\mathscr{S}$ and obtain the final stacked image using Equation 1.

Subsampling approach
Our second approach is based on the idea of *sparse recovery*, and also follows three stages.

1. *Low-resolution RTM*: We generate a low-resolution version of the velocity model by downsampling the spatial and temporal dimensions by a factor of α. Then, we run a low-resolution migration to obtain coarse versions of the source and receiver wavefields $\mathbf{W}_S^{\text{lo}}(x,y,z,t,j), \mathbf{W}_R^{\text{lo}}(x,y,z,t,j)$. We then correlate the wavefields to obtain a low-resolution image corresponding to each shot $j = 1,2,\ldots,S$, and stack these images to obtain a final low-resolution reflectivity estimate:

$$A_i^{\text{lo}}(x,y,z) = \sum_{t=1}^{T} \mathbf{W}_S^{\text{lo}}(x,y,z,t,j)\mathbf{W}_R^{\text{lo}}(x,y,z,t,j) , \qquad A^{\text{lo}}(x,y,z) = \sum_{i=1}^{S} A_i^{\text{lo}}(x,y,z) .$$

Computing these low-resolution estimates can be performed in a fraction α^3 (for 2D with equal subsampling factor per temporal and spatial dimension) of the time taken to perform normal RTM. Therefore, for small values of α (say, 0.25) this can be a significant speedup.

2. *Shot selection via sparse recovery*: For a given spatial region of interest Ω and a budget B, our aim is to identify the B shots that best approximate A_Ω, the final high-resolution image restricted to the region Ω. However, we do not have access to the final image, and therefore we use the *low-resolution* version as a proxy instead. Let vec denote the vectorization operator. Let $y := \text{vec}(A^{\text{lo}}(x,y,z))$ where $(x,y,z) \in \Omega$ and let $A = [a_1,\ldots,a_S]$ where $a_i = \text{vec}(A_i^{\text{lo}}(x,y,z))$ and $(x,y,z) \in \Omega$. Then our problem reduces to finding the best linear combination of at most B columns of A that is closest to y. This can be posed as a sparse recovery problem:

$$\min \|y - Ab\|_2 \quad \text{subject to } \|b\|_0 \le B . \tag{2}$$

This problem is again NP-hard in general. However, there are numerous heuristics to solve this problem very efficiently. One popular approach is to perform L1-relaxation, followed by a standard convex optimization method, e.g., the algorithm of van den Berg and Friedlander (2007).

Another approach is to use a greedy algorithm such as Orthogonal Matching Pursuit (OMP), discussed in Pati et al. (1993). We report results using both of these approaches. Finally, we obtain a sparse vector $\bar{b}$ whose B nonzeros correspond to the locations of the most informative shots.

3. *Migration:* Once we pick the set of B shot locations $l_1, l_2, \ldots, l_B$, we perform normal (full resolution) RTM using the set of chosen locations and obtain the final stacked image.

Examples

We perform several simulation experiments using the 2D SEG-EAGE salt model (of size 201×676), displayed in Figure 1(a). Our goal is to reliably image the rectangular region around the feature marked with $J = 6$ red dots (the dot locations can be interactively selected by the user and are used in the embedding approach only). Figure 1(b) shows the image formed by running RTM with a shot placed at every surface pixel location. This serves at the baseline image for numerical error calculation purposes. Figure 1(c) displays the image formed by running decimated RTM with $B = 12$ equispaced surface shots.

We now report simulation results using our proposed *embedding* approach. Figure 1(d) shows the resulting image for $B = 12$. The computational cost of this method includes J forward modelings and B migrations. We repeat the simulation experiment on the left sloping feature using the *subsampling* approach. Figure 1(e) shows the resulting image with $B = 12$ where the shot locations have been picked by solving Equation 2 using L1-minimization and picking the B largest coefficient locations. On the other hand, Figure 1(f) shows the resulting image for $B = 12$ where the shot locations have been selected using OMP. In all cases, our methods are able to reliably identify the features using only a small number of selected shots. Both the subsampled-L1 and subsampled-OMP costs are on the order of $B + \alpha^4 \cdot S$ migrations, where α is the subsampling factor and S is the total number of shots. In this case we have $\alpha = 1/4$ and $S = 676$, which gives a total cost of roughly 15 full-resolution migrations.

We provide quantitative comparisons of the different methods in Figure 2. Here, we plot the normalized mean square error (measured over the area of interest) with respect to the baseline RTM image using the different methods for a varying shot budget B. Our experiments suggests that data-driven migration can be an appealing alternative to (oblivious) decimation.

Conclusions

RTM, and by extension other shot-based techniques, is a very costly method that will become even more expensive with the introduction of complex physics, e.g. Orthorhombic media. Mitigation of the computational cost is a key endeavour for this community. In this work, we introduced three approaches that achieved such a goal without hampering the geophysical quality of the results. In fact, the results in the examples section show that our approach can recover the image with good accuracy while migrating only half (or even less) the number of shots needed by the decimation approach.

Acknowledgements

The authors thank Shell Intl. E&P Inc. for support and permission to publish this work.

References

Bear, G., Lu, C.P., Lu, R., Willen, D. and Watson, I. [2000] The construction of subsurface illumination and amplitude maps via ray tracing. *Leading Edge*, **19**(7), 726–728.

Godwin, J. and Sava, P. [2011] A comparison of shot-encoding schemes for wave-equation migration. *SEG Technical Program Expanded Abstracts*, 32–36.

Kaelin, B. and Guitton, A. [2006] Imaging condition for reverse time migration. *SEG Technical Program Expanded Abstracts*, 2594–2598.

Liu, S., Geng, J. and Feng, W. [2005] Controlled illumination and seismic acquisition geometry for target-oriented imaging. *Applied Geophysics*, **2**(4), 230–234.

Pati, Y., Rezaiifar, R. and Krishnaprasad, P. [1993] Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. *Asilomar Conference on Signals, Systems and Computers*, 40–44.

Romero, L., Ghiglia, D., Ober, C. and Morton, S. [2000] Phase encoding of shot records in prestack migration. *Geophysics*, **65**(2), 426–436.

Schmidt, L., Indyk, P., St-Cyr, A., Chen, C. and Hohl, D. [2013] Shot encoding with random projections. *SEG Technical Program Expanded Abstracts*.

van den Berg, E. and Friedlander, M.P. [2007] SPGL1: A solver for large-scale sparse reconstruction. `http://www.cs.ubc.ca/labs/scl/spgl1`.

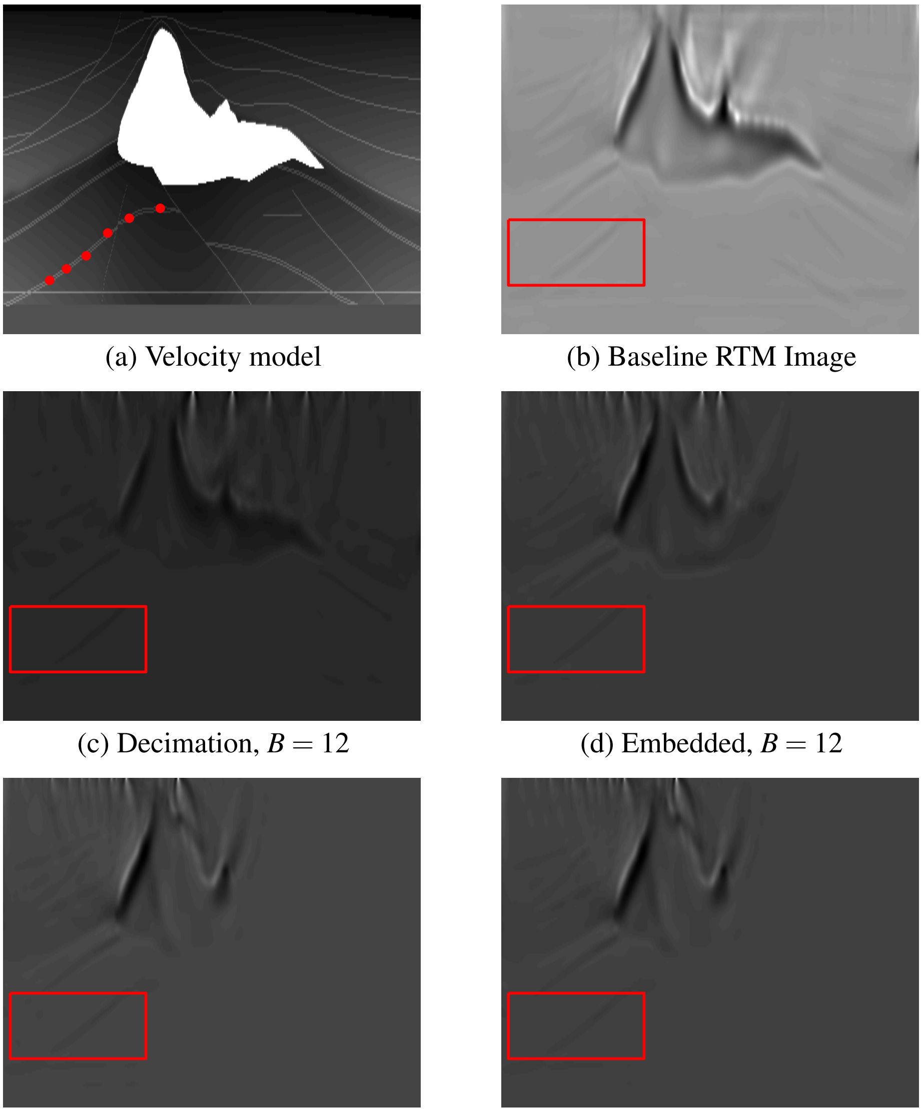

(a) Velocity model
(b) Baseline RTM Image
(c) Decimation, $B = 12$
(d) Embedded, $B = 12$
(e) Subsampled-L1, $B = 12$
(f) Subsampled-OMP, $B = 12$

Figure 1 *Example results. Subfigures (c) to (f) show the results of the four shot selection methods on the velocity model in subfigure (a). Subfigure (b) is the "ground truth" image obtained by running RTM with a large number of shots ($S = 676$). The red dots in subfigure (a) indicate the location of the subsurface shots used for the embedding approach. The red rectangles show the region of interest Ω which defines the objective function in the optimization procedures Subsampled-L1 and Subsampled-OMP.*

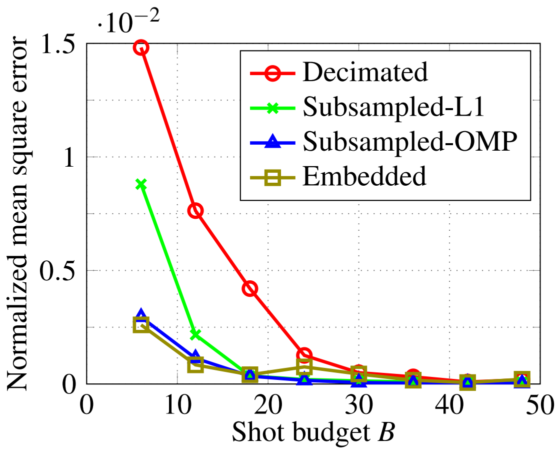

Figure 2 *Approximation errors for varying shot budgets. The approximation errors are computed on the region of interest (red rectangle in Figure 1) only. Our proposed methods (Subsampled-L1, Subsampled-OMP, and Embedded) give more accurate results for all values of the shot budget B. Subsampled-OMP and Embedded perform best, with Subsampled-OMP demonstrating somewhat more robust behavior. In order to match the performance of Subsampled-OMP, Decimated needs two to three times as many shots.*