A comparison of greedy algorithms for 3D Radon transforms: matching pursuit vs orthogonal matching pursuit

Francesco Perrone, Warren S. Ross, Sergio Grion, Shearwater GeoServices*

Summary

Greedy algorithms based on matching pursuit have received a lot of attention in recent times because of their ability to produce accurate, sparse, high-resolution representation of signals even when the input data are aliased. However, given the heuristic nature of the algorithms of this family, similar performances in terms of accuracy and quality of the signal representation come with very different computational costs and runtimes. We compare two matching pursuit alternatives in the context of the computation of 3D linear Radon transform: the basic and the orthogonal matching pursuit algorithms. We compare the two approaches in terms of quality of the signal representation, sparsity in the transformed domain, and runtime. Despite the greater computational cost involved in solving a least-squares optimization for all coefficients in orthogonal matching pursuit, we observe faster convergence, which is due to a smaller number of matching pursuit iterations. The higher quality of the estimated coefficients leads to higher accuracy in signal representation, i.e. smaller final residual. Our results are consistent with previous comparisons of the two approaches already presented in the literature for the case of signal reconstruction from random measurements.

Introduction

Radon transform is a very important tool in seismic data processing. 3D input data are usually strongly aliased in the crossline direction, a condition that hinders the quality of the representation in the transformed domain. Greedy algorithms, such as matching pursuit, are designed to overcome the challenges faced by algorithms based on regular and periodic sampling of the input data. Successes include the computation of Fourier transforms for irregularly sampled data (Nguyen and Winnet, 2011), crossline interpolation (Hollander et al., 2012; Özbek et al. 2009; Vassallo et al., 2010), and the implementation of the Radon transform (Cao and Ross, 2017; Kamil and Özbek, 2017). The effectiveness of matching pursuit strategies is agreed upon, but a myriad of flavors and different variants of the basic matching pursuit idea exist: orthogonal matching pursuit (OMP) is a popular choice in the digital signal processing community because, by construction, it guarantees orthogonality of reconstructed signal and residual for each new model parameter added to the representation. Hollander et al. (2012) proposed OMP as a data interpolation strategy although, in their work, they suggested it should be 'relaxed' to decrease its computational cost. OMP is by construction more expensive than conventional MP since the size of the operator grows with the number of model parameters estimated and because a least-squares optimization problem involving all coefficients estimated up to the current iteration must be solved for each new parameter added to the representation. However, because of the imposed orthogonality between data and residual at each iteration, OMP requires fewer iterations to converge to the desired result. In this work, we compare two implementations of 3D linear Radon transform: one based on a conventional matching pursuit (Cao and Ross, 2017) and the other based on orthogonal matching pursuit (Hollander et al., 2012; Tropp and Gilbert 2007). We benchmark the accuracy and quality of the image representation and reconstruction for different frequency bandwidths and different percentages of the total number of model parameters to be estimated in the transformed domain.

The theory of matching pursuit

Matching pursuit denotes a family of greedy procedures to obtain a representation of signals as a linear combination of the elements of a predefined basis. The signal representation can be sparse by construction since only the most energetic components are matched, and the algorithm stops once the maximum number of desired basis vectors has been used or the error between the input and the approximation falls below a user-defined threshold. The matching pursuit strategy can be adopted for the computation of sparse 3D linear Radon transforms (Cao and Ross, 2017) using the basis function $g(p_x, p_y; \omega) = e^{-i\omega(p_x x + p_y y)}$. The algorithm iteratively picks the (p_x, p_y)-component that carries the most energy and thus approximates the input signal as

$$f(x, y; \omega) = \sum_j w_j e^{-i\omega(p_x^j x + p_y^j y)} \quad (1)$$

in which w_j represents the weight of the j-th (p_x, p_y)-component. The formulation in the frequency domain is standard for the computation of the linear Radon transform for efficiency reasons.

The main difference between matching pursuit and orthogonal matching pursuit flows are that the estimation of the weights of the identified components differ, otherwise they are very similar as follows:

- Initialize the residual with the input data $f(x, y; \omega)$
- Project the residual to the Radon (p_x, p_y) domain
- Pick the most energetic component and add it to the signal representation

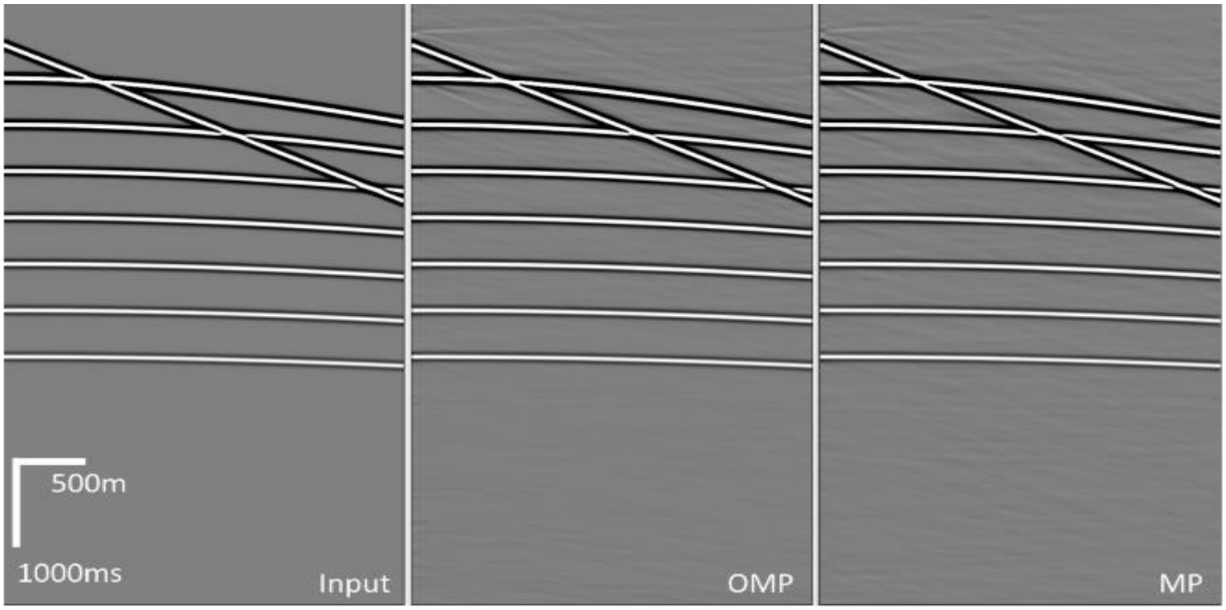

Figure 1: Input synthetic data (left), forward and inverse 3D Radon transform computed using orthogonal matching pursuit (center), and conventional matching pursuit (right). Notice the higher level of noise using the basic matching pursuit scheme. The display is clipped at 95% for each panel.

- If (**OMP**) solve a least-squares optimization problem for **ALL** the coefficients w_j of the components that have been picked so far;
- Compute the new residual by subtracting the current signal approximation from the input data
- Iterate from step 2. Unless the residual is small enough or the max number of iterations has been reached.

At each iteration, the orthogonal matching pursuit algorithm solves a least-squares inverse problem involving all the weights of all the components that have been identified up to that moment. The rationale of this approach is to minimize leakage during the greedy estimation of the components of the input signal $f(x, y; \omega)$, since the estimated residual is always orthogonal to the input signal because of the optimization step.

Examples

We use a synthetic example containing a limited number of events to assess the quality of the approximation obtained using MP and OMP. The optimization of the coefficients is done with a conjugate gradient (CG) scheme but because of the choice of the threshold, the algorithm is a *de facto* steepest descent. We tested reducing the threshold to run few iterations of CG but the results were marginally different and we decided to run all of the other comparisons with the least expensive setup. Figure 1 shows the comparison between the input data (left), the orthogonal matching pursuit (center) and the conventional matching pursuit round-trip (forward transform followed by the inverse transform) result. The parameterization is the same (30 Hz maximum frequency and 40% of the Radon coefficients to be used for the signal representation). The displays are clipped to 95% of the maximum amplitude in each panel. The orthogonal matching pursuit implementation produces a cleaner representation of the signal with fewer and weaker artefacts compared with the basic matching pursuit one. It is interesting to look at the signal in the Radon domain: matching pursuit allows us to automatically handle the sparseness of the signal, which is an attractive property, especially for 3D transforms of data that are not well sampled in at least one dimension. Figure 2 shows the Radon transforms of the input data (Figure 1, left) obtained with the orthogonal and basic Matching Pursuit implementation. The basic matching pursuit implementation (Figure 2, right) of the Radon transform appears “smoother” than the orthogonal matching pursuit one (Figure 2, left). However, the orthogonal implementation focuses better the linear (planar) event in the input gather and the hyperbolic events have more compact support, i.e. higher resolution.

In the second example, we use a common receiver gather extracted from the Volve dataset and shown in Figure 3. Volve is a field 200 km west of Stavanger in the Norwegian continental shelf, which was decommissioned in 2016 after 8.5 years in operation. The full dataset includes two ocean-bottom streamers acquisitions from 2002 and 2010. We use the PZ dataset from obtained from the 2002 survey. Figure 4 show with black squares the X and Y the positions of the

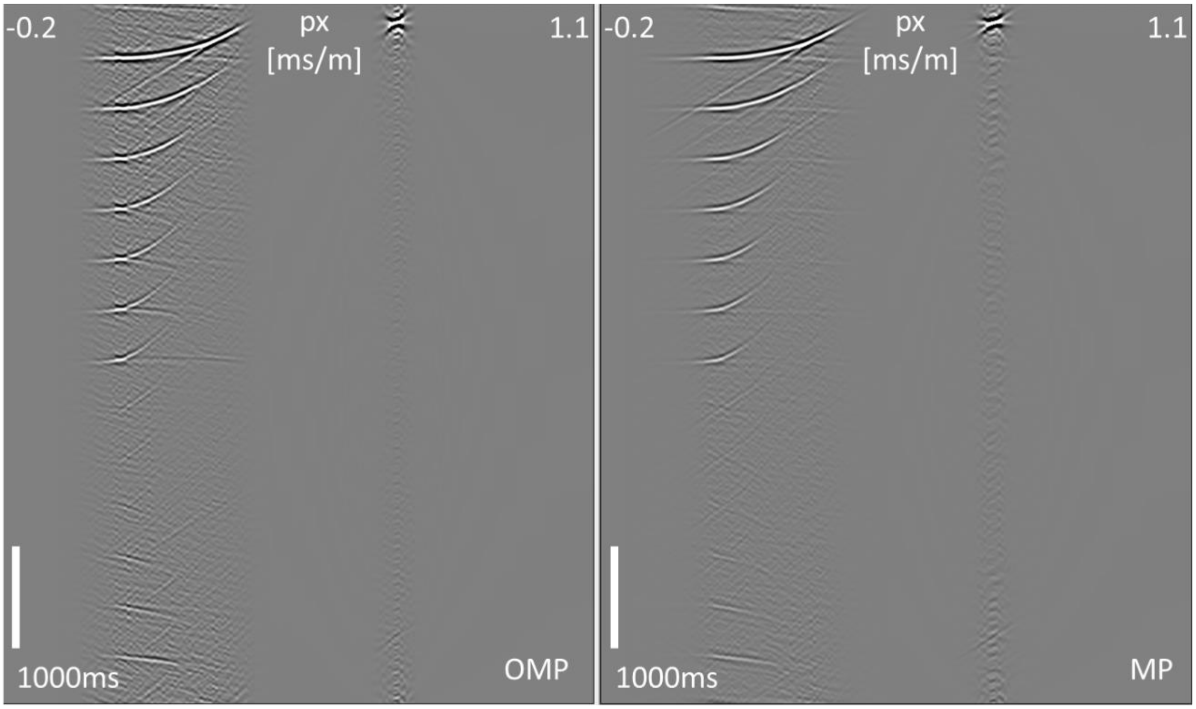

Figure 2: Forward Radon transform of the input synthetic gather: orthogonal matching pursuit (left), conventional matching pursuit (right)

sources and with a blue square the position of the receiver. Notice that the data have not been regularized. The gather contains 4579 traces and each trace is 7.5 s long and sampled at 4 ms. We compare the runtimes with different parameterizations with respect to a base case (20 Hz maximum frequency and 10% Radon coefficients) in Table 1. Both implementations are single-node and multithreaded. Table 1 reports the benchmark for the orthogonal and conventional matching pursuit implementations for maximum frequencies of 20 Hz and 40 Hz and different percentages of the total number of Radon coefficients used to represent the signal. The benchmarks show that, at fixed maximum frequency, for both conventional and orthogonal matching pursuit implementations, the runtime increases linearly as function of the percentage of coefficients. However, the orthogonal matching pursuit runtime increases slower. This is due to least-squares optimization of the coefficients, which guarantees the orthogonality between the input signal and the residual at the current step and thus allows the algorithm to reach the tolerance threshold faster than the basic matching pursuit version.

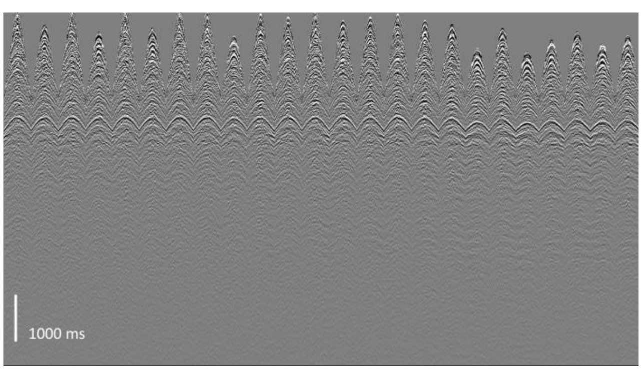

Figure 3: 'Flattened' 3D common receiver gather extracted from the Volve dataset. The horizontal axis is trace number.

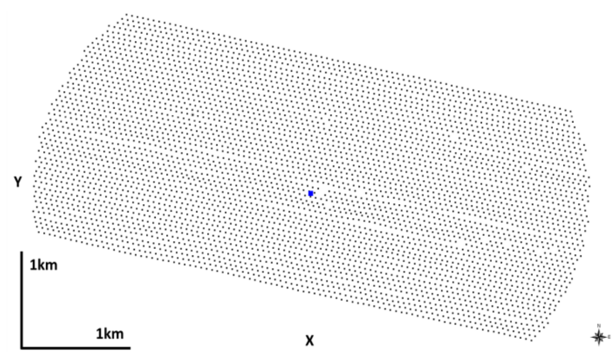

Figure 4: Acquisition geometry for the common receiver gather in Figure 3: the black squares represent the locations of the sources and the blue square the location of the receiver.

Max Frequency [Hz]	Coefficient [%]	OMP runtime	MP runtime
20	10	1.00	1.12
	20	1.57	2.23
	40	2.83	4.49
40	10	7.81	30.28
	20	9.99	49.00
	40	13.19	61.36

Table 1: Runtimes factors for 20 Hz and 40 Hz maximum frequency and different percentages of the number of Radon coefficients estimated for the Volve dataset using basic and orthogonal matching pursuit. The 20 Hz, 10% coefficient case represents the reference.

By doubling the maximum frequency, the OMP runtime increases by a factor slightly less than eight times: doubling the frequency range implies twice as many frequencies and it also decreases the sampling in the (p_x, p_y)-space, thus increasing by a factor 4 the number of coefficients in the Radon domain. However, the matching pursuit runtime is much higher because of the combined effect of the finer sampling of the (p_x, p_y)-space and the slow convergence of the algorithm that we already discussed for the comparison at equal maximum frequency. The slow convergence of the basic matching pursuit algorithm is highlighted again from a different perspective in Table 2, where we compare the residuals at equal maximum frequency for the two algorithms. The number in parentheses indicates the actual number of iterations required to reach the reported residual, which is expressed as a fraction of the initial value. The asterisk indicates that convergence (set at <0.1% of the initial residual energy) has been reached before the expected number of iterations (which depends on the percentage of coefficients and maximum frequency). It is interesting to observe that, for the case of Volve common receiver gather, orthogonal matching pursuit reaches convergence in essentially a fixed number of iterations. Even though each iteration is more computationally expensive than a matching pursuit iteration, there are significantly fewer iterations, and thus the total runtime accumulated over frequencies (see Table 1) reflects the faster convergence.

Max Frequency [Hz]	Coefficient [%]	OMP residual % (iterations)	MP residual % (iterations)
20	10	0.13 (1424)	15.5 (1424)
	20	0.099 (1440)*	3.88 (2849)
	40	0.099 (1440)*	0.24 (5698)
40	10	0.099 (1424)*	0.11 (5706)
	20	0.099 (1424)*	0.099 (5813)*
	40	0.099 (1424)*	0.099 (5813)*

Table 2: Convergence behavior of the OMP and MP algorithms for different maximum frequencies and percentages of the total number of coefficients to estimate. The * indicates that convergence was reached before the maximum number of iterations; the number in parentheses indicates the actual number of iterations performed at maximum frequency.

Conclusions

We compared the performances of two flavors of greedy algorithms applied to the computation of the 3D Radon transform: basic and orthogonal matching pursuit. Orthogonal matching pursuit performs better in terms of signal representation although the computational cost, i.e. the dimensionality of the operator grows with the number of iterations. Basic matching pursuit has a fixed computational cost per iteration, and every new coefficient to be estimated does not increase the computational complexity of the algorithm. However, the better reconstruction and data-fitting properties of orthogonal matching pursuit allow faster convergence in the case of complex real data, which leads to lower overall computational cost compared to the basic matching pursuit implementation. The better performance of orthogonal matching pursuit has been already documented in the literature for the case of signal recovery from random measurements (Tropp and Gilbert, 2007), this work confirms the superiority to the basic matching pursuit scheme for multidimensional Radon transform of seismic data.

Acknowledgments

The authors thank Shearwater GeoServices for permission to publish this work and Equinor and the former Volve Licence partners for making the Volve dataset publicly available.

References

Cao, W. and Ross, W. S., 2017, High resolution 3D Tau-p transform by matching pursuit: 87th Annual International Meeting, SEG, Expanded Abstracts, 4302-4306.

Hollander, Y., Kosloff, D., Koren, Z. and Bartana, A.,2012, Seismic Data Interpolation by Orthogonal Matching Pursuit: 71th Conference and Exhibition, EAGE, Extended Abstracts, B023.

Kamil, Y. I. and Özbek, A., 2017, IAA-MP: a novel greedy algorithm for severely underdetermined inverse problems: 87th Annual International Meeting, SEG, Expanded Abstracts, 4287-4291.

Nguyen, T. and Winnet, R., 2011, Seismic interpolation by optimally matched Fourier components: 81st Annual International Meeting, SEG, Expanded Abstracts, 3085-3089

Özbek, A., Özdemir, A. K. and Vassallo, M. , 2009, Interpolation by matching pursuit: 79th Annual International Meeting, SEG, Expanded Abstracts, 3254-3258.

Tropp, J. A. and Gilbert, A. C., 2007, Signal recovery from random measurements via Orthogonal Matching Pursuit: IEEE Transactions on Information Theory, 53, no.12, 4655-4666.

Vassallo, M., Özbek, A., Özdemir, K. and Eggenberger, K., 2010, Crossline wavefield reconstruction from multicomponent streamer data: Part 1 – Multichannel interpolation by matching pursuit (MIMAP) using pressure and its crossline gradient: Geophysics, 75, no. 6, WB53-WB67