Sparse-spike deconvolution revisited: thin-layer solution via reinforcement learning

Saad Saleh and Sergey Fomel, The University of Texas at Austin*

ABSTRACT

We revisit the well-known yet unsolved sparse-spike deconvolution problem and offer a new promising approach to address challenges posed by thin layers. In its simplest noise-free formulation, the problem can be cast as an ill-posed matrix inversion problem, $s = Wr$, where we estimate the reflectivity vector r, *given* the trace vector s and the wavelet convolution singular matrix W, together with the prior knowledge that the number of non-zero elements in r is much smaller than the number of samples, n. A typical approach is to solve the constrained optimization problem

$$\min_{r \in \mathbb{R}^n} J(r) \doteq \|s - Wr\|_2 + \lambda g(r)$$

Where $g(r)$ is a regularization term, and $\lambda > 0$ is a control parameter. The direct approach is to take $g(r) = \|r\|_0$, which is the number of non-zero elements in r. A common approach, however, is to use $g(r) = \|r\|_1$ as a proxy because l_1 minimization leads to good sparse approximations while being more tractable numerically; e.g., see (Baraniuk, 2007). Recently, Torres and Sacchi (2023) addressed the limitations of the traditional approaches using deep learning.

In this paper we consider the direct approach based on l_0 minimization. The celebrated matching pursuit (MP) algorithm (Mallat and Zhang, 1993) offers a numerically tractable method based on sequential projections on 1-dimensional subspaces in an n-dimensional space. However, as Zhang and Castagna (2011) pointed out, MP tends to fail when dealing with thin layers.

This paper shows that MP can be readily extended to resolve the thin-layer problem. Specifically, we introduce a new algorithm, Variable Multidimensional Sequential Projections (VMSP), which produces excellent results for reflectivity series containing thin layers. The algorithm is similar to MP but differs in 2 crucial ways. First, it replaces projections on 1-dimensional subspaces in each iteration by projections on certain p-dimensional subspaces, where p is a small natural number, typically less than 10. This step extends the idea of a dictionary of basis vectors into a dictionary of "thin" $n \times p$ matrices. Secondly, VMSP allows for "variable" projections; i.e., the dimensionality of the optimal subspace computed in each iteration can vary from one iteration to the next. As such, the algorithm produces an optimal sequence of natural numbers $(p_1, \ldots, p_l)$ corresponding to the dimensions of various subspaces involved in the search for best estimate in each iteration.

The presentation will provide details of the VMSP algorithm. The framework can be cast as an optimal solution to a finite Markov Decision Process for computing the particular sequence of projections on a set of subspaces that result in the best approximation of the original sparse reflectivity series. As such, the problem can be formulated as a Reinforcement Learning problem (Szepesvári, 2022).

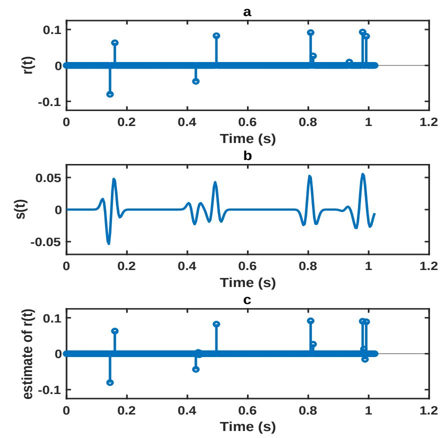

Figure 1: (a) reflectivity series, (b) trace, (c) deconvolved trace.

In Figure 1(a) we show a synthetic series of reflection coefficients consisting of 256 samples, 10 of which are non-zero. When convolved with a standard 20-Hz Ricker wavelet, the series yields the synthetic trace shown in Figure 1(b). The challenge is to use this trace as an input to an algorithm, together with the Ricker wavelet, to estimate the original sparse-spike reflectivity series. The user-specified parameters required for our algorithm are p, the highest dimension to use in the iterative projections, and l, the number of iterations to consider. For example, with $p = 5$ and $l = 5$, the VMSP algorithm found the optimal sequence (4,5,5,1,5). That is, the best estimate for r is found by an extended MP algorithm that searches for the optimal 4-dimensional projection in the first iteration, the optimal 5-dimensional projection in the second iteration, etc. The computed optimal estimate for r is shown in Figure 1(c). The algorithm has correctly computed all locations and approximate amplitudes of the 10 non-zero reflection coefficients in the original series. However, it produced a very small number of spurious coefficients that are nearly zero, which do not appear in the original series. The relative error in the final estimate is 0.1162. But the error can be driven essentially to zero if one applies a simple threshold-based denoising filter to the final estimate. As such, VMSP can provide surprisingly accurate sparse-spike deconvolution results.