

ROBUST INVERSION VIA SEMISTOCHASTIC DIMENSIONALITY REDUCTION

Aleksandr Y. Aravkin^{*†} Michael P. Friedlander[†] Tristan van Leeuwen^{*}

^{*} Department of Earth and Ocean Sciences, University of British Columbia

[†] Department of Computer Science, University of British Columbia

ABSTRACT

We consider a class of inverse problems where it is possible to aggregate the results of multiple experiments. This class includes problems where the forward model is the solution operator to linear ODEs or PDEs. The tremendous size of such problems motivates the use of dimensionality reduction (DR) techniques based on randomly mixing experiments. These techniques break down, however, when robust data-fitting formulations are used, which are essential in cases of missing data, unusually large errors, and systematic features in the data unexplained by the forward model. We survey robust methods within a statistical framework, and propose a sampling optimization approach that allows DR. The efficacy of the methods are demonstrated for a large-scale seismic inverse problem using the robust Student's *t*-distribution, where a useful synthetic velocity model is recovered in the extreme scenario of 60% corrupted data. The sampling approach achieves this recovery using 20% of the effort required by a direct robust approach.

Index Terms— inverse problems, seismic inversion, stochastic optimization, robust estimation

1. INTRODUCTION

Consider a generalized data fitting scheme where we conduct m experiments and record the corresponding measurements. We encode the parameters for each of the m experiments in the matrix $Q = [q_1, q_2, \dots, q_m]$, and encode the measurements corresponding to each experiment in $D = [d_1, d_2, \dots, d_m]$. We then fit model parameters x via the model

$$D = F(x; Q) + \mathcal{E}, \quad (1)$$

where \mathcal{E} captures the discrepancy between observations and predictions. We assume that F is linear in Q , which is true for all models where F is the solution operator to a linear differential equation with boundary conditions, and the columns of Q comprise the right-hand sides of these equations.

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, BPG, BP, Chevron, Conoco Phillips, Petrobras, PGS, Total SA, and WesternGeco. The authors thank British Gas for providing the synthetic velocity model used in the experiments. January 17, 2012

The maximum-likelihood approach to estimation can be generically expressed by the data-fitting problem

$$\underset{x}{\text{minimize}} \quad \psi(R(x)), \quad (2)$$

where

$$R(x) = D - F(x; Q) := [r_1(x), \dots, r_m(x)]$$

is the residual, and ψ is the negative log-likelihood of the underlying distribution for \mathcal{E} . One interpretation of this optimization problem is to choose parameters x that maximize the likelihood that $R(x)$ are samples of the random variable \mathcal{E} . Regularization terms may be needed, depending on the application.

We are motivated by the full-waveform inversion (FWI) problem in geophysics, where F is the solution operator for the Helmholtz equation, the vectors q_i encode information for the i th source experiment, and the vectors d_i contain the corresponding measurements. A typical survey in exploration seismology contains thousands of large-scale experiments. The aim is to recover the velocity field on a grid over a 2- or 3-dimensional volume. This leads to an overwhelming amount of data that consumes months of CPU time on industrial-sized clusters.

Dimensionality reduction (DR) is a technique where entire groups of experiments are fused into “super” experiments (e.g., “super shots” or “random source encoding”, in the seismic context) with the overall effect of reducing the problem size [1, 2, 3]. A least-squares (LS) fit for (1) is a popular choice, largely because of the implicit assumption that the error \mathcal{E} is normally distributed. A notable algorithmic benefit of an LS fit is that the DR technique can be naturally interpreted as a form of stochastic optimization. However, the LS approach is unsuitable for noisy or missing data, as often encountered in practice. Similarly, the LS formulation breaks down in the face of systematic features in the data unexplained by the model F . As a result, tremendous effort goes into data cleaning and increasing the complexity (and cost) of the model.

Our approach is based on a robust fit of (1), where \mathcal{E} is modeled using a distribution with heavy tails, e.g., a Student's *t*-distribution, first used by [4] for data fitting. This approach is based on robust statistics, where the aim is to inoculate the fitting process against outliers [5, 6, 7]. However, the

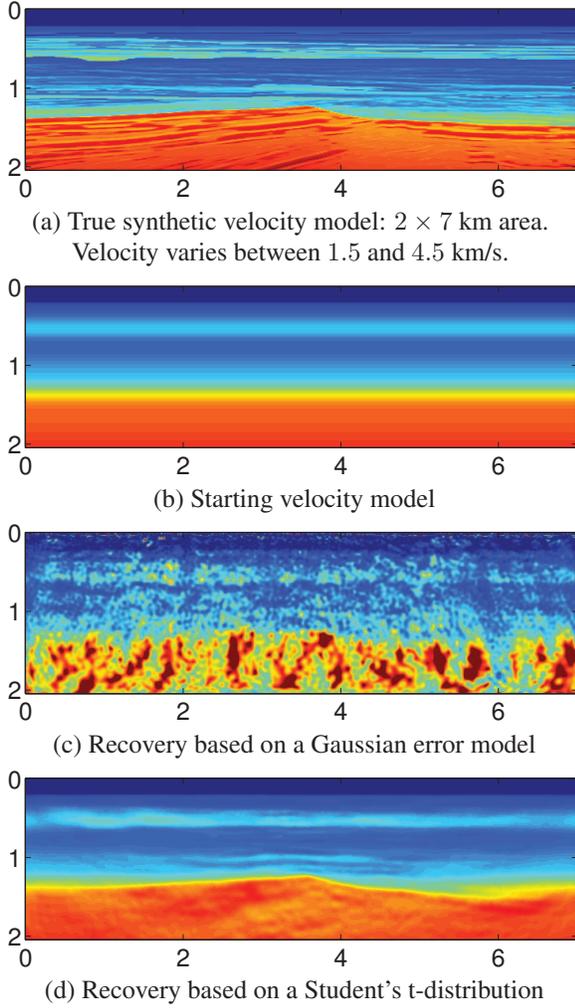


Fig. 1. Seismic full waveform inversion with 60% corrupted data.

special form of least-squares, which allowed us to use DR as an algorithmic framework, is not available in the robust context. Instead, we use a semistochastic method based on randomizing batches of experiments.

2. SEISMIC INVERSION

The FWI problem is classically formulated as

$$\underset{x}{\text{minimize}} \quad \psi(R(x)) := \|R(x)\|_F^2, \quad (3)$$

which is a nonlinear LS problem. In this application,

$$F(x; Q) = MU \quad \text{and} \quad H(x)U = Q, \quad (4)$$

where $H(x)$ is the Helmholtz operator, U is the wavefield induced by the source experiments encoded in Q , and M restricts U to the receiver positions on the surface [8].

The terrific computational cost in solving this problem can be illustrated by describing the work involved in evaluating a single gradient. The adjoint-state method requires the solution of the second equation in (4), and a corresponding adjoint equation involving H^* . This translates to the solution of $2m$ PDE solves per frequency. The simple test problem, based on a synthetic velocity model shown in Fig. 1, has $m = 141$ and 17 frequencies, for a total of 4,794 PDE solves for a single gradient evaluation. Industrial problems easily have thousands of experiments over tens of frequencies.

Fig. 1 illustrates the main weakness of LS inversion. In this example, we use (3) to recover the velocity model shown in Fig. 1(a) when 60% of the data is corrupted. This might correspond to the situation in practice where many receivers malfunction. The velocity model shown in Fig. 1(c), recovered using a quasi-Newton method, which represents the state-of-the-art in this application [8], is clearly not a good representation of the underlying model.

3. ROBUST ESTIMATION

The problems encountered in the LS inversion of the last section can be interpreted in a statistical framework by considering the “tail heaviness” of the underlying distribution for the errors \mathcal{E} . The Gaussian density, implied by the choice of the LS objective function, decays as $\exp(-\tau^2)$, where τ is the distance from the mean. As a result, outliers are extremely rare in the Gaussian model, e.g., the probability of seeing an event more than 8 standard deviations from the mean is of the order 10^{-15} . The inversion process, therefore, unnecessarily weighs outliers at the expense of good data.

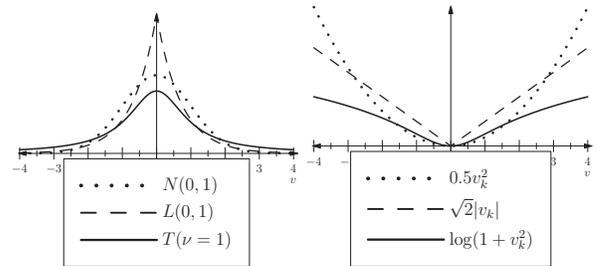


Fig. 2. Densities and penalties.

A robust approach to inverse problems is derived by assuming a distribution that has heavier tails. Examples include the ℓ_1 -Laplacian and Student's t- distributions, among others; see Fig. 2. Among the distributions that are supported on the entire space, the heaviest tails possible are proportional to $1/\tau^{d+2}$, where d is the problem dimension; the Student's t is one such example.

The parametric form of the chosen distribution, through the maximum-likelihood approach, determines the particular form of the inverse problem analogous to (3): Gaussian corresponds to the 2-norm, ℓ_1 -Laplacian corresponds to the 1-norm,

and Student's t corresponds to the objective $\log(1 + \tau^2)$. The first two functions are convex; the last is nonconvex. In many applications it is crucial to preserve the convexity of the inverse problem. But when F is nonlinear, as for example FWI in §2, convexity of the inverse problem is highly unlikely. In such cases, we can be free to ignore issues related to nonconvexity, and instead focus on the additional robustness afforded by the tail heaviness of the Student's t-distribution. See [9] for a comparative study with other heavy-tailed distributions. Figs. 1(c) and (d) show the benefits of the robust approach based on the Student's t against one based on a Gaussian model.

4. STOCHASTIC OPTIMIZATION

Stochastic optimization is a technique especially suited for large scale problems like FWI, where gradient computations with full data are prohibitively expensive. Recently developed methods exploit the special structure of (3) in order to develop a sampling scheme that randomly aggregates the experiments $\{q_i\}$ and corresponding observations $\{d_i\}$ into a set of super experiments [3]. As a result, the reduced problem involves only a fraction of the computational effort for each gradient evaluation.

The linearity of $F(x; Q)$ in Q allows us to exploit this approach because $R(x)W = DW - F(x; QW)$. This corresponds to aggregating observations and experiments into matrices DW and QW with k columns each, where $k \ll m$. This aggregation sampling scheme meshes well with the Frobenius norm, since it is induced by an inner product. Define the sampled function by

$$\psi_w(R(x)) := \psi(R(x)W) = \|R(x)W\|_F^2, \quad (5)$$

where W is a matrix of k mixing vectors, where k can be as small as 1. If $\mathbb{E}[WW^T] = I$, then

$$\begin{aligned} \mathbb{E}_w[\psi_w(R(x))] &= \mathbb{E} \langle R(x)W, R(x)W \rangle \\ &= \mathbb{E} \langle R(x)^T R(x), WW^T \rangle \\ &= \langle R(x)^T R(x), \mathbb{E}[WW^T] \rangle \\ &= \|R(x)\|_F^2 = \psi(R(x)). \end{aligned} \quad (6)$$

Also, by the linearity of the gradient,

$$\mathbb{E}_w[\nabla \psi_w(R(x))] = \nabla \mathbb{E}_w[\psi_w(R(x))] = \nabla \psi(R(x)). \quad (7)$$

In words, the expectation of the approximate objective and gradient are equal to the true objective and gradient. As a result of this property, this sampling scheme can be used within the framework of a stochastic optimization method.

In general, stochastic gradient methods for minimizing differentiable functions f can be generically expressed by the iteration

$$x_{k+1} = x_k - \alpha_k (s_k + e_k),$$

where s_k is any sufficient descent direction for f , i.e., $s_k^T \nabla f(x_k) \leq -\mu \|\nabla f(x_k)\|^2$ and μ is a positive constant,

and e_k is a random error. If the function has a Lipschitz gradient, $\mathbb{E}[e_k] = 0$, $\mathbb{E}[\|e_k\|^2]$ is bounded, and $\sum_k \alpha_k = \infty$ and $\sum_k \alpha_k^2 < \infty$, then $\nabla f(x_k) \rightarrow 0$; see [10, Prop. 3].

We now consider search directions based on the gradient of the sampled function (5), which in the stochastic optimization framework corresponds to $s_k + e_k = \nabla \psi_w(R(x_k))$. By (7), we have $\mathbb{E}_w[\nabla \psi_w(R(x))] = \nabla \psi(R(x))$, which implies $s_k = \nabla \psi(R(x_k))$ —which is clearly a direction of sufficient descent—and $\mathbb{E}[e_k] = 0$, as required.

This motivation breaks down if we consider more general distributions for \mathcal{E} , which necessarily changes the definition of ψ in (3). However, as long as the errors made across experiments are independent and identically distributed, then ψ is separable across the columns of $R(x)$, i.e.,

$$\psi(R(x)) = \sum_{i=1}^m \rho(r_i(x)) \quad (8)$$

for some function ρ . For example, a Gaussian distribution for \mathcal{E} gives $\rho(r_i(x)) = \|r_i(x)\|_2^2$, as in (3), Student's t gives $\rho(r_i(x)) = \sum_j \log(k + r_{ij}(x)^2)$, and ℓ_1 -Laplacian gives $\rho(r_i(x)) = \|r_i(x)\|_1$.

To design a DR technique that works for this general case, we again consider the approximation (5). For general ψ , however, the assumption $\mathbb{E}_w[WW^T] = I$ no longer suffices to prove (6) and (7). Fortunately, we can still recover these results by choosing a particular type of distribution for the mixing-matrix W , as shown in the next theorem.

Theorem 1. *Suppose that the mixing-matrix W (with $k \leq m$) is chosen to be a uniform random selection of columns of the scaled $m \times m$ identity matrix $\frac{m}{k}I$. Then for any objective ψ given by (8),*

$$\mathbb{E}_w[\psi_w(R(x))] = \psi(R(x)) \quad (9a)$$

and

$$\mathbb{E}_w[\nabla \psi_w(R(x))] = \nabla \psi(R(x)). \quad (9b)$$

Proof. Note that

$$\begin{aligned} \mathbb{E}_w[\psi_w(R(x))] &= \mathbb{E}_w[\psi(R(x)W)] \\ &= \mathbb{E}_w \left[\sum_{j=1}^k \rho((R(x)W)_{j \cdot}) \right] \\ &= \sum_{j=1}^k \mathbb{E}_w[\rho((R(x)W)_{j \cdot})] \\ &= \frac{m}{k} \sum_{j=1}^k \mathbb{E}_j[\rho(r_{i(j)}(x))], \end{aligned}$$

where $i(j)$ is randomly selected without replacement. Since $i(j)$ is equally likely to be any particular column index, we necessarily have

$$\begin{aligned} \frac{m}{k} \mathbb{E}_j[\rho(r_{i(j)}(x))] &= \frac{m}{k} \frac{1}{m} \sum_{i=1}^m \rho(r_i(x)) \\ &= \frac{1}{k} \psi(R(x)). \end{aligned}$$

Because there are k terms in the sum, and each summand is equal to $\frac{1}{k} \psi(R(x))$, (9a) holds. The relationship in (9b) follows immediately from the linearity of the gradient. \square

Note that the particular choice of W in Theorem 1 satisfies $\mathbb{E}_W[WW^T] = I$, which was the property we used in deriving (6) and (7).

The subsampling strategy suggested by Theorem 1 gives us a range of choices: at one extreme, we sample only a single column of $R(x)$ at each iteration, which corresponds to $k = 1$; a less extreme choice is to choose a k that is small relative to m . Both of these cases can be considered under the umbrella of incremental gradient methods [10, 11] for minimizing general functions f . In that framework, even if we make the favorable assumption that f is strongly convex, it is only possible to guarantee that this class of methods generates iterates x_k that converge sublinearly to a solution x^* , i.e.,

$$\|x_k - x^*\|^2 \leq \mathcal{O}(1/k).$$

An alternative approach that gradually increases the sample size as the iterations progress is proposed in [12]. The result is a sampling method that interpolates between the one-at-a-time incremental gradient method at one extreme, and using the full gradient at the other. By choosing k to grow at a certain rate, it is possible to guarantee that the iterates x_k converge linearly, i.e.,

$$\|x_k - x^*\|^2 \leq \mathcal{O}(\gamma^k) \quad \text{for some } \gamma < 1.$$

In practice, it is virtually impossible to guarantee convexity of $\psi(R(x))$ unless ψ is convex and $R(x)$ is linear; this is not the case for many inverse problems. Moreover, the loss functions corresponding to heavy-tailed distributions, such as Student's t , are nonconvex [9, Theorem 2.1]. Nonetheless, the theory for the strongly convex case is also supported by empirical evidence, where sampling strategies tend to outperform basic incremental gradient methods [12].

Fig. 3 compares the sampling and direct approaches by showing how the relative model error $\psi(R(x^k))/\psi(R(x^*))$ decreases as a function of effective passes through the data. The result in Fig. 1(c), obtained using the sampling approach, consumed 1/5th of the computational effort of a direct approach.

5. REFERENCES

- [1] J. R. Krebs, J. E. Anderson, D. Hinkley, R. Neelamani, S. Lee, A. Baumstein, and M.-D. Lacasse, "Fast full-wavefield seismic inversion using encoded sources," *Geophysics*, vol. 74, no. 6, pp. WCC177–WCC188, 2009.
- [2] F. Herrmann, M. P. Friedlander, and O. Yilmaz, "Fighting the curse of dimensionality: compressive sensing in exploration seismology," Tech. rep., University of British Columbia, 2011.
- [3] E. Haber, M. Chung, and F. J. Herrmann, "An effective method for parameter estimation with pde constraints with multiple right hand sides," Tech. Rep. TR-2010-4, UBC-Earth and Ocean Sciences Department, 2010.
- [4] K.L. Lange, R.J.A. Little, and J.M.G. Taylor, "Robust statistical modeling using the t distribution," *Journal of the ASA*, vol. 84, no. 408, pp. 881–896, 1989.
- [5] P. J. Huber, *Robust Statistics*, John Wiley & Sons, Inc., New York, 1981.
- [6] R. A. Maronna, D. R. Martin, and V. J. Yohai, *Robust Statistics: Theory and Methods*, Wiley, 2006.
- [7] A. Tarantola, *Inverse Problem Theory and Methods for Parameter Estimation*, SIAM, 2002.
- [8] J. Virieux and S. Operto, "An overview of full-waveform inversion in exploration geophysics," *Geophysics*, vol. 74, pp. 127–152, 2009.
- [9] A.Y. Aravkin, M.P. Friedlander, F.J. Herrmann, and T. van Leeuwen, "Robust inversion, dimensionality reduction, randomized sampling," *Submitted to Math. Prog.*, November 2011.
- [10] D. P. Bertsekas and J. N. Tsitsiklis, "Gradient convergence in gradient methods with errors," *SIAM J. Optim.*, vol. 10, no. 3, pp. 627–642, 2000.
- [11] A. Nedic and D. Bertsekas, "Convergence rate of incremental subgradient algorithms," *Stochastic Optimization: Algorithms and Applications*, pp. 263–304, 2000.
- [12] M. P. Friedlander and M. Schmidt, "Hybrid deterministic-stochastic methods for data fitting," Tech. Rep., Univ. of British Columbia, April 2011, revised September 2011.

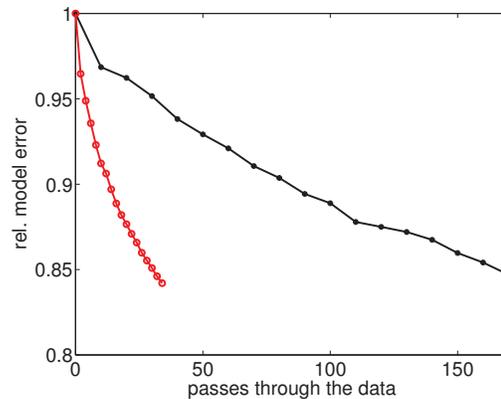


Fig. 3. Performance of the sampling method (bottom curve) compared to a full gradient method (top curve).