

Optimal design of simultaneous source encoding

Eldad Haber *

Kees van den Doel †

Lior Horesh ‡

This work was supported in part by IBM Research and MITACS

Open Collaborative Research project for Design in Inversion

<http://ocrdesign.wix.com/home>

July 22, 2013

Abstract

A broad range of parameter estimation problems involve the collection of an excessively large number of observations N . Typically each such observation involves excitation of the domain through injection of energy at some pre-defined sites and recording of the response of the domain at another set of locations. It has been observed that similar results can often be obtained by considering a far smaller number K of multiple linear superpositions of experiments with $K \ll N$. This allows the construction of the solution to the inverse problem in time $\mathcal{O}(K)$ instead of $\mathcal{O}(N)$. Given these considerations it should not be necessary to perform all the N experiments but only a much smaller number of K experiments with simultaneous sources in superpositions with certain weights. Devising such procedure would result in a drastic reduction in acquisition time.

The question we attempt to rigorously investigate in this work is: what are the optimal weights? We formulate the problem as an optimal experimental design problem and show that by leveraging techniques from this field an answer is readily available. Designing optimal experiments requires some statistical framework and therefore the statistical framework that one chooses to work with plays a major role in the selection of the weights.

1 Introduction

We consider discrete parameter estimation problems with data vectors d_j modeled as

$$d_j = A_j m + \epsilon_j \quad (j = 1, \dots, N), \quad (1)$$

where A_j is an $s \times n$ linear observation matrix that corresponds, for example, to the j^{th} configuration of an experiment, $m \in \mathbb{R}^n$ is the parameter vector to be recovered, and ϵ_j are zero-mean independent random vectors with covariance matrix $\sigma^2 I$. In addition, we assume that there is a large set of observation matrices A_j , which is the case of interest in many geophysical and medical applications such as seismic and electromagnetic (EM) exploration, diffraction and electrical impedance tomography [1, 4, 6].

*Department of Earth, Ocean and Atmospheric Sciences, University of British Columbia, Vancouver, Canada. haber@math.ubc.ca/kvdoel@cs.ubc.ca

†Department of Earth, Ocean and Atmospheric Sciences, University of British Columbia, Vancouver, Canada.

‡Department of Business Analytics and Mathematical Sciences, IBM TJ Watson Research Center, Yorktown Heights, NY, USA.

The linear systems in (1) can be aggregated into a single large system by defining

$$\hat{A} \equiv \begin{pmatrix} A_1 \\ \vdots \\ A_N \end{pmatrix} \quad \text{and} \quad \hat{d} \equiv \begin{pmatrix} d_1 \\ \vdots \\ d_N \end{pmatrix}. \quad (2)$$

In the well-posed case where $\hat{A}^\top \hat{A}$ has a stable inverse, the least-squares (maximum likelihood if ϵ_j are Gaussian) estimate \hat{m} of m is

$$\hat{m} = (\hat{A}^\top \hat{A})^{-1} \hat{A}^\top \hat{d}. \quad (3)$$

When $\hat{A}^\top \hat{A}$ does not have a stable inverse, a penalized least-squares estimate of the form

$$\hat{m} = (\hat{A}^\top \hat{A} + \lambda^2 M)^{-1} \hat{A}^\top \hat{d}, \quad (4)$$

is often used. Here M is usually a symmetric positive semi-definite matrix and λ is a positive regularization parameter. This choice will simplify analysis later on, although the proposed approach can be generalized for use of other regularizers of other functional forms.

We are particularly interested in large-scale problems where both the number of model parameters n and the number of observations N are large. For instance, in an electromagnetic or a seismic survey, the number of such observations, N , can be of the order of 10^4 or 10^5 . For such problems solution of the system (4) using direct methods is intractable computationally and instead iterative methods (typically, CGLS or LSQR [10]) are employed. The main computational bottleneck associated with such solvers is the computation of matrix vector products: When N is very large, the computation of \hat{A} times a vector can be prohibitively intensive by itself. We thus need to devise some data reduction methodology to make computations feasible.

For example, return to the well-posed case and assume the matrices A_j come from a parameter estimation problem where the observed phenomena can be described by a set of differential equations that are linear w.r.t. the source term. Such operators can be written as

$$A_j = SL^{-1}Q_j,$$

where S and Q_j are sparse selection matrices related to the sensors and the sources respectively, but L is a discretization of a partial differential equation on a mesh. Such problems arise in seismic and electromagnetic exploration as well as in impedance and diffraction tomography [1, 4, 6]. The least-squares estimate of the model m is

$$\hat{m} = \left(\sum_j Q_j^\top L^{-\top} S^\top S L^{-1} Q_j \right)^{-1} \sum_j Q_j^\top L^{-\top} S^\top d_j, \quad (5)$$

which requires multiple applications of L^{-1} making the procedure computationally expensive. Alternatively a far less expensive procedure relies upon a recently developed data reduction technique based on the idea of ‘simultaneous sources’ [2, 20, 11, 17, 19]. Here the estimate is based on a weighted combination of the form $d(w) = \sum w_j d_j$. The basic idea is to exploit the savings achieved by simultaneously collecting the data.

For example, if each experiment corresponds to a different source in seismic or electromagnetic exploration, it can take a rather long time to record many different experiments. If, on the other hand, it is possible to “shoot” simultaneously with all the sources, then the recording time can be significantly shortened as a collection of simultaneous sources enables direct measurement of linear combinations of the data. That is, for a fixed w one conducts an experiment that measures $d(w)$ directly. This acquisition strategy is of course, far more efficient than measuring each data vector by itself and then combining them with the weights w .

The linear system corresponding to $d(w)$ is

$$d(w) = A(w)m + \epsilon(w), \quad (6)$$

where $A(w) = \sum w_j A_j$ and $\epsilon(w) = \sum w_j \epsilon_j$. We assume the individual source noise levels ϵ_j are known. The estimates (3) and (4) corresponding to $d(w)$ are

$$\hat{m}(w) = (A(w)^\top A(w) / \|w\|^2)^{-1} A(w)^\top d(w) / \|w\|^2 \quad (7)$$

if $A(w)^\top A(w)$ has a stable inverse, and

$$\hat{m}(w) = \left(A(w)^\top A(w) / \|w\|^2 + \lambda^2 M \right)^{-1} A(w)^\top d(w) / \|w\|^2 \quad (8)$$

otherwise. Clearly, both solutions depend on w only through its unit-norm normalization $w/\|w\|$.

For example, let us return to the well-posed parameter estimation example with $A_j = SL^{-1}Q_j$. The penalized least-squares estimate (8) is

$$\hat{m} = \left[\left(\sum_j w_j Q_j^\top \right) L^{-\top} S^\top S L^{-\top} \left(\sum_k w_k Q_k \right) \right]^{-1} \left(\sum_i w_i Q_i^\top \right) L^{-\top} S^\top \left(\sum_m w_m d_m \right),$$

which requires much fewer applications of L^{-1} than (5). Of course, some information may be lost as (d_1, \dots, d_N) is not, in general, a sufficient statistic for m . We return to this point in Section 2.2 but it should be clear that some information might be sacrificed for the sake of computational or data collection savings.

Although the advantages of using simultaneous sources have been previously explored, the question of selecting appropriate weights w has not received sufficient attention. We are aware of two main conducts that have been utilized so far. In [7], the weights w_j are independent zero-mean random variables of variance one and are used to define the Hutchinson randomized trace estimator [13]. More precisely, let R be the residual matrix defined as

$$R = (A_1 m - d_1 \quad \cdots \quad A_N m - d_N).$$

Then,

$$\begin{aligned} \mathbb{E} \|A(w)m - d(w)\|^2 &= \mathbb{E} (w^\top R^\top R w) = \text{trace}(R^\top R) \\ &= \sum_j \|A_j m - d_j\|^2 = \|R\|_F^2, \end{aligned}$$

where $\|\cdot\|_F$ denotes the Frobenius norm. If $w^{(1)}, \dots, w^{(b)}$ are independent random weight vectors, then by the law of large numbers

$$\frac{1}{b} \sum_{k=1}^b \|A(w^{(k)})m - d(w^{(k)})\|^2 \rightarrow \|R\|_F^2,$$

as $b \rightarrow \infty$. Therefore the finite sum on the left hand-side can be used to approximate $\|R\|_F$. It is easy to see that the uniform distribution on $\{-1, 1\}$ minimizes the variance of the trace estimator (see for example, [22]). Methods to select appropriate values of b are discussed in [22].

In [21] the Frobenius norm of the residual matrix is replaced by the squared 2-norm and w is then chosen so that

$$\|A(w)m - d(w)\|^2 \approx \|R\|_2^2.$$

The paper is structured as follows. In Section 2 we derive the mathematical framework. In Section 3 we discuss efficient numerical techniques for the solution of the weight design problem. In Section 4 we demonstrate the effectiveness of our approach by performing numerical experiments and finally, in Section 5 we summarize the study.

2 Optimal selection of weights

In this section we consider statistical properties of the estimate \hat{m} to select optimal weights w . We also discuss methods to compensate for the information loss and to compare the loss achieved with different choices of weights w .

We begin our discussion with the well-posed problem and its solution (7). As can be observed, \hat{m} depends only on $w/\|w\|$; we may therefore assume that $\|w\| = 1$. Since \hat{m} is an unbiased estimator and $\text{Var}(\epsilon(w)) = \sigma^2 I$, it follows that its mean square error (MSE) is

$$\text{MSE}(\hat{m}) = \mathbb{E} \|\hat{m}(w) - m\|^2 = \sigma^2 \text{trace}((A(w)^\top A(w))^{-1}).$$

To find an optimal w , in the sense of minimizing the MSE, we need to solve the constrained optimization problem

$$\begin{aligned} \min_w \quad & \text{trace}((A(w)^\top A(w))^{-1}) \\ \text{s.t.} \quad & \|w\| = 1. \end{aligned} \quad (9)$$

Before discussing properties of the solutions, we consider the ill-posed problem.

2.1 The ill-posed case

Other than recent numerous exceptions (see [8, 5, 12, 15, 9, 14] and references therein), optimal experimental design of ill-posed problems has been somewhat an under-researched topic. In the case of ill-posed problems, the selection of optimal weights for (8) is more difficult because this estimate is biased and its bias depends on the unknown m : $\text{bias}(\hat{m}) = -\lambda^2 C(w)^{-1} M m$, where the inverse Fisher matrix is given by

$$C(w) = A(w)^\top A(w) + \lambda^2 M$$

(again, we may assume $\|w\| = 1$). The MSE of \hat{m} is then

$$\begin{aligned} \text{MSE}(\hat{m}) &= \mathbb{E} \|\hat{m} - m\|^2 = \|\text{bias}(\hat{m})\|^2 + \text{trace}(\text{Var}(\hat{m})) \\ &= \lambda^4 \|C(w)^{-1} M m\|^2 + \sigma^2 \text{trace}(A(w) C(w)^{-2} A(w)^\top). \end{aligned} \quad (10)$$

As in the well-posed case, the goal would be to find weights w that minimize this MSE subject to the constraint $\|w\| = 1$ but this is not possible because m is unknown. We need information on m to control the MSE.

If m is known to be in a convex set S , then we could consider a minmax approach where we minimize the worst $\text{MSE}(\hat{m})$ for $m \in S$. A less pessimistic approach is to minimize an average $\text{MSE}(\hat{m})$ over S . Of course, other choices are possible and may be useful as long their interpretation can be rationalized. In this note we minimize a weighted average of the MSE. We model m as random and since \hat{m} only depends linearly on m , it is enough to use prior second order moment conditions. For example, if $\mathbb{E}(m) = 0$ and $\text{Var}(m) = \Sigma_m$, then the Bayes risk $R_{\text{FB}}(w)$ under squared error loss of the frequentist estimate \hat{m} defined by the fixed w is

$$\begin{aligned} R_{\text{FB}}(w) &= \mathbb{E}(\text{MSE}(\hat{m})) \\ &= \lambda^4 \text{trace}(C(w)^{-1} M \Sigma_m M C(w)^{-1}) + \sigma^2 \text{trace}(A(w) C(w)^{-2} A(w)^\top), \end{aligned}$$

and we solve the optimization problem

$$\begin{aligned} \min_w \quad & R_{\text{FB}}(w) \\ \text{s.t.} \quad & \|w\| = 1. \end{aligned} \quad (11)$$

Finally, assume the distributions are Gaussian. That is, the conditional distribution of $d(w)$ given m is $\mathcal{N}(A(w)m, \sigma^2 I)$ and the prior distribution of m is $\mathcal{N}(0, \tau^2 M^{-1})$, with $\tau^2 = \sigma^2 / \lambda^2$. In this case, (8) is the posterior mean (and a maximum a-posteriori or MAP estimator) and the posterior covariance matrix is $\text{Var}(m | d(w)) = \sigma^2 C(w)^{-1}$. For each fixed w the Bayes risk of \hat{m} is therefore

$$R_{\text{B}}(w) = \sigma^2 \text{trace}(C(w)^{-1}).$$

The corresponding optimization problem is

$$\begin{aligned} \min_w \quad & R_{\text{B}}(w) \\ \text{s.t.} \quad & \|w\| = 1. \end{aligned} \quad (12)$$

Remark: So far we have only considered linear/affine reconstructions of the model, which of course do not include many recovering techniques based, for example, on total variation or sparsity control. However, if the reason for choosing one of these nonlinear procedures is the belief that its optimal Bayes risk is smaller than that of the linear reconstruction, then the optimal Bayes risk of the latter provides an upper bound for the optimal Bayes risk of the nonlinear procedure. Furthermore, if the upper bound is not too conservative then minimizing the Bayesian risk for an affine estimator can be a useful guide for other reconstruction techniques.

2.2 Compensating for the information loss

As we have explained, some information may be lost with the reduction to $d(w)$. If the loss is unacceptable, we may compensate by using more than a single reduction as follows. Let k be the number of independent data reductions and $w^{(\ell)} = (w_1^{(\ell)}, \dots, w_N^{(\ell)})^\top$. Define

$$d_j^{(\ell)} = A_j m + \epsilon_j^{(\ell)} \quad (j = 1, \dots, N, \ell = 1, \dots, k),$$

where the sequence $\{\epsilon_j^{(\ell)}\}_{j,\ell}$ is assumed to be i.i.d. $\mathcal{N}(0, \sigma^2)$. The ℓ^{th} data reduction is $d(w^{(\ell)}) = \sum_j w_j^{(\ell)} d_j^{(\ell)}$. We solve for m using the vector of k reductions $d(W_k)$ and matrix $A(W_k)$ defined as

$$d(W_k) = \begin{pmatrix} d(w^{(1)}) \\ \vdots \\ d(w^{(k)}) \end{pmatrix} \quad \text{and} \quad A(W_k) = \begin{pmatrix} A(w^{(1)}) \\ \vdots \\ A(w^{(k)}) \end{pmatrix}.$$

The associated Bayes risk is $R_B(W_k) = \sigma^2 \text{trace}(C(W_k)^{-1})$ with $C(W_k) = A(W_k)^\top A(W_k) + \lambda^2 M$. The optimization problem is

$$\begin{aligned} \min_{W_k} \quad & R_B(W_k) \\ \text{s.t.} \quad & \|w^{(1)}\| = 1, \dots, \|w^{(k)}\| = 1. \end{aligned} \tag{13}$$

There are some particular cases that may make this optimization easier and provide upper bounds for the optimal solution of (13). These bounds can be useful when we study the loss as a function of k . One particular case consists of using prior information about the experiments to classify the matrices A_j into k groups whose members are similar on a realistic set of models m . For example $d(w^{(1)})$ may include only A_1 and A_2 while $d(w^{(2)})$ uses only A_3, A_4 and A_5 , etc. A second simplification consists of assuming $w^{(1)} = \dots = w^{(k)}$ so that there is only one constraint but k independent experiments. The final and simplest approach is to find the optimal w^* in (12) and then repeat the experiment k times. This leads to the risk

$$R_B(w^*, k) = \sigma^2 \text{trace}[(k A(w^*)^\top A(w^*) + \lambda^2 M)^{-1}]. \tag{14}$$

2.3 Quantifying the information loss

Although in general the optimization problem does not have a unique minimum, a local minimum may still provide a sufficient reduction of the risk achieved with the reduced data $d(w)$. This does not mean, however, that the data reduction is useful for much information may be lost. It is therefore important to assess whether the computational savings obtained through the data reduction justifies the information loss.

To assess the information loss of the data reduction, we compare its associated risk to that of the complete data. The Bayes risk of the complete (i.e. exhaustive acquisition) problem is

$$R_0 = \text{trace}((\hat{A}^\top \hat{A} + \lambda^2 M)^{-1}),$$

while the risk associated with the reduced data $d(w)$ is

$$R_B(w) = \text{trace}((A(w)^\top A(w) + \lambda^2 M)^{-1}). \tag{15}$$

We quantify the information loss using the risk ratio

$$\text{Loss}(w) = \frac{R_B(w)}{R_0}. \tag{16}$$

This loss indicates the degradation of the average recovery when using the reduction $d(w)$ instead of the full data. If the loss is much greater than 1 then using a single set of weights may lead to significant deterioration in the recovery. In this case one may want to use a number of weights as explained in Section 2.2. The loss is then

$$\text{Loss}(W_k) = \frac{R_B(W_k)}{R_0}.$$




Figure 1: The loss as a function of the number of vectors

To determine an appropriate value of k , we may plot the loss as a function of k . The idea is to balance the trade-off between the cost of the experiment and the information loss. The following example illustrates this point.

Example 1 We generate $N = 50$ random matrices A_j of sizes 100×200 and use k random unit vectors (not optimal) $w^{(1)}, \dots, w^{(k)}$. Figure 1 shows a plot of $R_B(W_k)$ as a function of k . The curve has a classical L-shape and provides an upper bound for the optimal risks $R_B(W_k^*)$. The figure shows that using less than 10 experiments will result in significant loss of information but using more than 30 may be a waste of resources. Thus, the loss curve is important if we would like to design an effective and yet efficient experiment.

2.4 Numerical properties of the design function

We now study properties of the weights selected in the fully Bayesian framework; that is, solutions of (12). First, note that neither the constraint $\|w\| = 1$ nor the objective functions $R_B(w)$ and $R_B(W_k)$ are convex and consequently there is no guarantee of a unique solution. In fact, if w_* is a solution of (12), then so is $\tilde{w}_* = -w_*$. And any permutation of an optimal sequence $w_*^{(1)}, \dots, w_*^{(k)}$ is a solution of (13).

Example 2 Consider the following simple 2×2 well-posed case:

$$A_1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

It is easy to see that for $\|w\| = 1$,

$$\text{trace}((A(w)^\top A(w))^{-1}) = \frac{2}{1 - 4w_1^2 w_2^2} = \frac{1}{2(w_1^2 - 1/2)^2},$$

and the optimal solutions of (9) are $w^* = \pm(1, 0)$. That is, the optimal solution does not use the information provided by d_2 . To see how much information is lost, we compare the risk of the




Figure 2: The objective function as a function of ϕ for the weight $w = (\cos \phi \sin \phi)$ for two realizations of the matrices A_1 and A_2 from example 3. We also show the case $N = 3$.

full least-squares solution and the one based on $d(w^*)$. The two optimal weights lead to the same estimate $\hat{m}(w^*) = d_1$ with $\text{MSE}(w^*) = \sigma^2 \text{trace}(I) = 2\sigma^2$. For the full least-squares solution, we have $\hat{A}^\top \hat{A} = 2I$ and thus

$$\hat{m}_{\text{ls}} = \frac{1}{2} \begin{pmatrix} d_{1,1} + d_{2,1} \\ d_{1,2} - d_{2,2} \end{pmatrix},$$

with $\text{MSE}(w_{\text{ls}}^*) = \sigma^2$. Hence, in this case the information lost by the reduction to $d(w^*)$ leads to a MSE that is twice as large as that of the full least-squares estimate. To compensate, we may conduct do a single replication with the optimal w^* to match the MSE obtained without the data reduction.

Example 3 Next, we consider the problem (12) for the case $N = 2$ and $N = 3$. The matrices A_i are random 7×10 matrices with i.i.d. $N(0,1)$ entries, $M = I$ and $\sigma = 1 = \lambda$. We parameterize w in terms of the angle $0 \leq \phi \leq 2\pi$ by $w = (\cos \phi \sin \phi)$ for $N = 2$ and by spherical coordinates for $N = 3$. In Figure 2 the value of the objective function for two random realizations of the matrices A_1 and A_2 is plotted, and one realization for $N = 3$. The problem has multiple minima for the first instance, but a unique minimum for the second instance, as the two minima are equivalent. For the case $N = 3$ there are many more local minima at various places on the sphere. Local minima may not be too detrimental to our final goal. The value of MSE is not as important as long as we achieve a reduction of the risk. For the same reason, even if we do not have a global solution we may still find weights w that substantially decrease the risk one would have achieved by using other more naive choices.

3 Numerical solution

3.1 Direct approach

To solve the problem (12) we enforce the constraints by parameterizing w in terms of $N - 1$ spherical coordinate angles ϕ_i on the unit N -sphere:

$$\begin{aligned} w_1 &= \cos(\phi_1) \\ w_2 &= \sin(\phi_1) \cos(\phi_2) \\ w_3 &= \sin(\phi_1) \sin(\phi_2) \cos(\phi_3) \\ &\vdots \\ w_{N-1} &= \sin(\phi_1) \dots \sin(\phi_{N-2}) \cos(\phi_{N-1}) \\ w_N &= \sin(\phi_1) \dots \sin(\phi_{N-2}) \sin(\phi_{N-1}). \end{aligned} \tag{17}$$

We then use the BFGS method to solve the minimization problem iteratively. This requires the computation of the derivatives of the objective function. Rewriting the objective function as

$$\mathcal{J}(w) = \text{trace}((A(w)^\top A(w) + M)^{-1}) = \sum_j e_j^\top (A(w)^\top A(w) + M)^{-1} e_j$$

and setting

$$z_j = (A(w)^\top A(w) + M)^{-1} e_j \quad \leftrightarrow \quad (A(w)^\top A(w) + M) z_j = e_j$$

we have that

$$\nabla_w \mathcal{J}(w) = \left(\sum_j e_j^\top \nabla_w z_j(w) \right)^\top$$

To evaluate the derivative of z_j we use implicit differentiation

$$0 = \nabla_w ((A(w)^\top A(w) + M) z_j) = \nabla_w (A(w)^\top A(w) z_j^{\text{fix}}) + (A(w)^\top A(w) + M) \nabla_w z_j$$

We further define the matrices

$$\begin{aligned} G_F(v) &:= \nabla_w (A(w)v) = (A_1 v \quad \dots \quad A_N v) \\ G_T(u) &:= \nabla_w (A(w)^\top u) = (A_1^\top u \quad \dots \quad A_N^\top u) \end{aligned}$$

which leads to the calculation

$$G_j(w) := \nabla_w (A(w)^\top A(w) z_j) = A(w)^\top G_F(z_j) + G_T(A(w) z_j)$$

and finally

$$\nabla_w \mathcal{J} = - \sum_j G_j^\top(w) (A(w)^\top A(w) + M)^{-1} e_j. \quad (18)$$

The gradient w.r.t. the angles ϕ is then obtained by multiplication by the Jacobian of the transformation (17). We have used the BFGS [18] method to solve each optimization problem and we fix the termination tolerance of each problem, demanding that the norm of the gradient is reduced to 10^{-2} of its initial value. To address the problem of multiple minima we repeat the optimization procedure for a number of random initial guesses for w and select the solution with the smallest value of the objective function. The method to solve the problem (13) which involves multiple weights is completely analogous.

The techniques introduced in the previous sections work well for small-scale problems. However for large-scale problems, when A_j is large the computation of the objective function and its derivatives can be rather expensive. Here, we can use a stochastic trace estimation [13] and substitute the problem with

$$\begin{aligned} \min R_{FB}(w; v) &= \mathbf{E}_v [\sigma^2 v^\top (C(w)^{-1} M \Sigma_m M C(w)^{-1}) v \\ &\quad + v^\top (A(w)^\top C(w)^{-2} A(w)) v] \\ \text{s.t.} &\quad \|w\| \leq 1 \end{aligned} \quad (19a)$$

and

$$\begin{aligned} \min R_{FB}(w) &= \sigma^2 \mathbf{E}_v [v^\top (A(w)^\top A(w) + \sigma^2 M) v] \\ \text{s.t.} &\quad \|w\| \leq 1 \end{aligned} \quad (19b)$$

where the expectation value over random vectors v is approximated by sampling.

3.2 Heuristic weight selection

We can interpret minimization of (15) as maximizing the operator $A(w)^\top A(w)$. However in practice it only needs to be maximized over a set of “relevant” models m . Let us now make the bold step of maximizing over a single reference model m_0 , which can be for example a constant background model or a more sophisticated prior. (Note that we can not use the data to obtain m_0 as we consider experimental design here.) We now obtain

$$W = \arg \max_W m_0^\top A(W)^\top A(W) m_0.$$

Now $A(W)m_0$ is just the reduced data that would result from the model m_0 . Let us define and compute the m_0 data matrix \tilde{d} as in (2) but with d_i taken to be a row of s data (so \tilde{d} is $N \times s$) and denote by W the $k \times N$ matrix of weights (with unit norm rows). We get

$$W = \arg \max_W \text{trace } \tilde{d}^\top W^\top W \tilde{d}$$

Let us now perform a singular value decomposition (SVD) upon $\tilde{d} = U\Sigma S^\top$. We obtain

$$W = \arg \max_W \|WU\Sigma\|_2^2$$

and the solution for W (under appropriate constraints) is the sub matrix of U^\top corresponding to the truncated SVD, which is equivalent to obtaining the first k principal components from the data matrix \tilde{d} . For large scale problems, the first k left eigen vectors and the corresponding eigen values can be estimated effectively through Lanczos iteration. In the following sections we shall refer to this heuristic choice for W as the ‘‘SVD’’ choice. This heuristic approach works almost as well as the more optimal W in the numerical examples below while avoiding the intricate and expensive optimization process associated with the optimal weights problem. A similar approach using the SVD for faster data processing rather than experimental design was proposed in [23].

4 Numerical experiments

In this section we test the methods developed on two-dimensional seismic tomography problems. First we consider a simple toy model on a square domain to test and second we consider a more complicated reconstruction based on the Marmousi model [3].

We consider a rectangular domain Ω with the top boundary $\partial\Omega_1$ representing a ground-air interface while the remaining boundary $\partial\Omega_2$ represent an artificial boundary of the model.

The acoustic pressure field u is determined by the Helmholtz equation with appropriate boundary conditions:

$$(\nabla \cdot \nabla + \omega^2 m)u = q \text{ in } \Omega \quad (20a)$$

$$u = 0 \text{ on } \partial\Omega_1 \quad (20b)$$

$$(\partial_n + i\omega\sqrt{m})u = 0 \text{ on } \partial\Omega_2 \quad (20c)$$

where $m(x)$ is the slowness, i.e., $m = 1/c^2$ with c the propagation velocity and $q(x)$ represents sources of angular frequency ω . The Sommerfeld boundary condition (20c) prevents artificial reflections from the synthetic domain boundary $\partial\Omega_2$. We discretize (20) on a rectangular grid using a five point stencil for the Laplacian which is sufficient for numerical experiments and obtain

$$A^\omega \mathbf{u} = (-L^\omega(m) + \omega^2 \text{diag}(\mathbf{m}))\mathbf{u} = \mathbf{q},$$

where the Laplace operator L^ω depends upon m and ω through the Sommerfeld boundary condition. Note that \mathbf{u} is a complex valued vector.

The inverse problem is to determine \mathbf{m} from measurements of \mathbf{u} for several given sources \mathbf{q}_i , where $i = 1, \dots, N$ labels the experiments which can involve several frequencies. We assume further that all experiments measure \mathbf{u} at the same locations which are described by the linear sampling operator S , i.e., we measure $\mathbf{d}_i = S\mathbf{u}_i$. The predicted data for a given slowness model \mathbf{m} is $\mathbf{F}_i(\mathbf{m}) = SA^{-1}\mathbf{q}_i$ (suppressing implicit ω dependence) and the least squares error is

$$\phi_e(\mathbf{m}) = \frac{1}{2} \sum_{i=1}^N (\mathbf{F}_i - \mathbf{d}_i)^\dagger (\mathbf{F}_i - \mathbf{d}_i). \quad (21)$$

We add a Tikhonov regularization term to ϕ_e and solve the inverse problem as

$$\mathbf{m}^* = \arg \min_{\mathbf{m}} \phi_e + \frac{1}{2} \lambda^2 \|\mathbf{m}\|_\Delta^2 \quad (22)$$

where $\|\cdot\|_{\Delta}$ denotes the Laplacian weighted ℓ_2 norm.

We can apply the formalism developed above and consider instead of (21)

$$\hat{\phi}_e(\mathbf{m}) = \frac{1}{2} \sum_{k=1}^K \left(\sum_{i=1}^N w_i^{(k)} (\mathbf{F}_i - \mathbf{d}_i) \right)^\dagger \left(\sum_{j=1}^N w_j^{(k)} (\mathbf{F}_j - \mathbf{d}_j) \right) \quad (23)$$

for various choices of weights. The Bayes risk function (13) is defined by linearization around a reference model \mathbf{m}_0 . Because we are dealing with complex variables the weight vectors are now complex valued and satisfy $\mathbf{w}^\dagger \mathbf{w} = 1$ and there is an additional complex phase factor $e^{i\rho_k}$ for each component in (17).

4.1 A simple toy model

For the first problem we consider we take Ω to be the square domain $[-1, 1]^2$ which we discretize by a uniform 42^2 grid. Absorbing boundary conditions are assumed except at the top where we impose Dirichlet boundary conditions. Sources are placed on the left boundary and 20 uniformly spaced detectors are placed on the right. In Fig. 3(a) we depict the synthetic velocity model which we use to compute simulated data to which we add 2% Gaussian noise. Sources of $f = 0.5$ (i.e., one wavelength in domain) are placed on the left in two configurations. First we consider $N = 2$ experiments with a source at $1/3$ depth and at $2/3$ depth. Second one source is placed in the middle. The inversion is performed as described in Section 3 and the regularization parameter λ was chosen by the Morozov discrepancy principle [16]. We solved the optimization problem (22) using limited memory BFGS. We depict the resulting reconstructions in Fig. 3 (b-c) along with a relative error measure defined as follows:

$$\epsilon_r = \frac{\|m - m_{true}\|_2}{\|m_{true}\|_2} \quad (24)$$

It is clear that for $N = 2$ we capture the basic structure of the model, whereas $N = 1$ is insufficient.

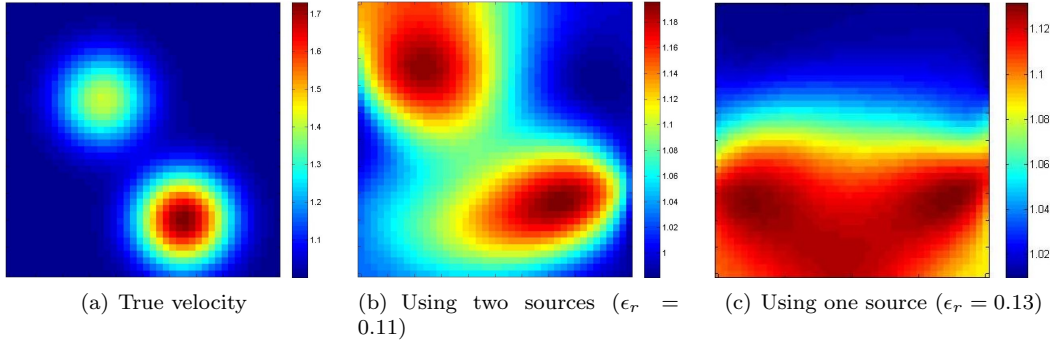


Figure 3: Reconstructions with their corresponding relative error $\epsilon_r = \frac{\|m - m_{true}\|_2}{\|m_{true}\|_2}$

Next we consider using a single simultaneous source by superposition of the two sources with a unit norm weight vector. In Fig. 4 we show reconstructions using various methods for selection of the weights (left column). The right column shows the absolute value of the pressure field u generated by the corresponding simultaneous source. We computed the optimal weight vector as described in section 3 using 32 different random initial guesses for w , and selecting the one with lowest risk. Second we used the method described in section 3.2 and finally we generated three random phase weight vectors. It is evident that the optimal weights yield the best reconstruction both visually and in terms of relative error as defined in 24. These optimal results were followed closely by those of the SVD method 3.2. The random weights offered a reasonable reconstruction in (g) but not in (e) and (i). By inspecting the fields in the rightmost

column is seems that the optimal choice has selected the source combination to provide the best possible domain coverage. The source of the SVD based approach did so to a lesser extent and the random choices essentially cover the domain randomly, as expected.

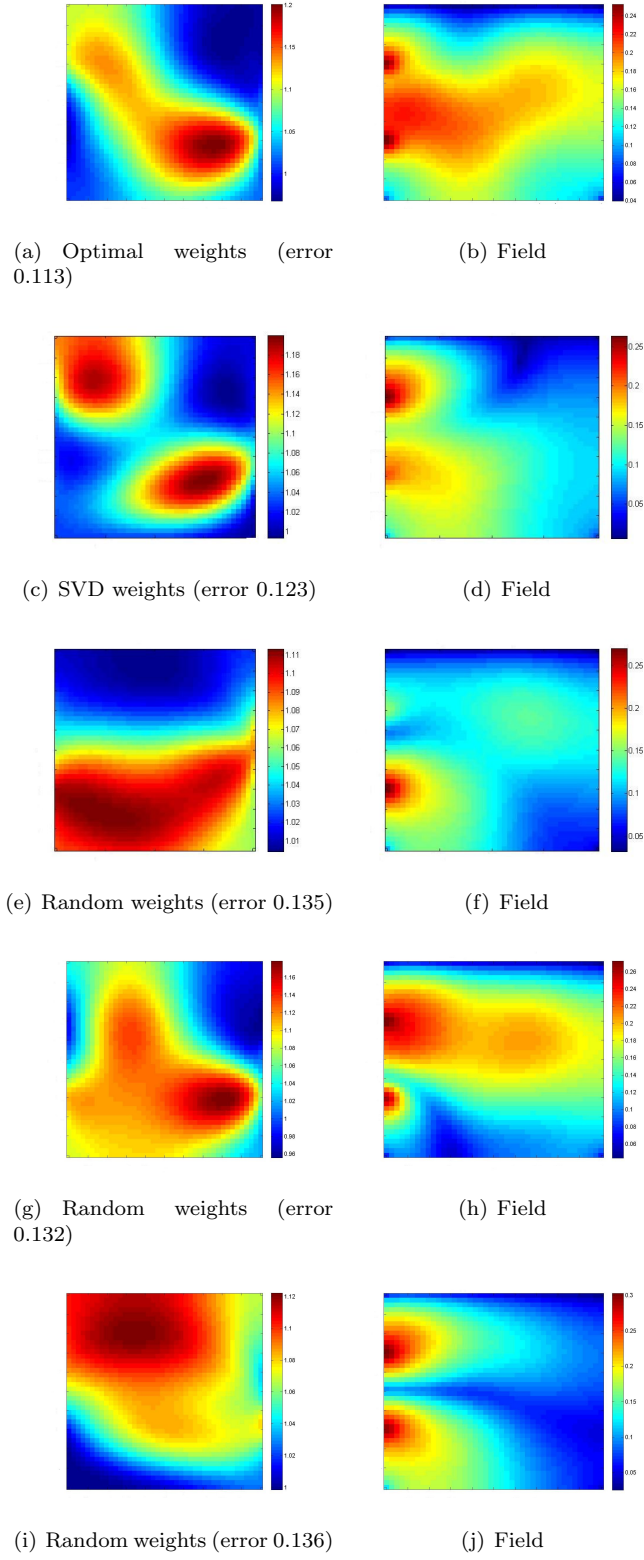


Figure 4: Reconstructions with a single combined source using various weight choices (left). On the right the fields generated by the sources are displayed.

4.2 Marmousi model

As a second numerical experiment we consider the true velocity model depicted in Fig. 5(f) of the Marmousi model. The domain is 3 km by 10 km and velocity is given in km/sec units. We use a 208×60 grid for the discretization. 40 emitters are placed at uniform distances on the surface and 39 receivers are placed in between the emitters. We generate data for 5 distinct frequencies $f = 0.2, 1.5, 3, 4.5, 6$ (in Hz), with a total of $5 \times 40 = 200$ experiments. We note that the low frequencies are needed to avoid local minima, but are not quite realistic. We added 3% Gaussian noise to the synthetic data. Optimal weights were precomputed before the inversion by linearization around a constant $m_0 = 2$ km/sec background model which was also used for the SVD method for selection weights.

First of all we present the reconstructions for a single frequency of 0.2 Hz using all 40 experiments and with $K = 2$ weight vectors, each a combination of the 40 sources. For the latter we show the results for the optimal weights, SVD based weight selection, and random phase weights in two realizations. It is clear from Fig. 5 that the random weight selection performs much worse than the other methods. In this case the SVD reconstruction is even better than using the optimal weights.

In Fig. 6 we depict the resulting pressure field from the two simultaneous sources. We observe again that the optimal and SVD weights result in a more or less uniform and balanced coverage of the domain whereas the random choices do not share this virtue.

To obtain more accurate reconstructions higher frequencies must be included. Because of the non-convexity of the problem we perform a frequency continuation by which we first solve (to full convergence) for the lowest frequency, then use the result as initial guess for the next stage where we add an additional frequency, etc. Because accuracy of the reconstruction increases we can expect to require more simultaneous sources as well as more and more frequencies are included. In this example we use $K = 2, 6, 12, 20, 30$, simultaneous sources corresponding to frequencies $f = 0.2, 1.5, 3, 4.5, 6$, again for the aforementioned 3 methods for weights selection. In Fig. 7 we depict the resulted reconstructions. Remarkably, the reconstructions with the SVD or optimal weights are no worse visually than the reconstruction using all the data, whereas the random weights do not offer comparable results. We note that the advantage of using simultaneous sources decreases with higher frequencies.

5 Summary and conclusions

We have introduced a novel methodology for optimal experimental design of acquisitions involving simultaneous sources. Based on an optimal Bayesian risk statistical framework and the assumption of a linear inverse problem we have formulated an optimality criterion for selection of the weights in which individual observations should be combined. We derived a heuristic guess for the optimal weights based upon the singular value decomposition and we performed numerical experiments to test the efficacy of the method. In all cases it was found that an informed choice of the weights generally outperforms a random weight choice. This is because our statistical framework incorporates information about the domain under consideration and the experimental setup.

Even though for most realistic applications the numerical models lead to nonlinear inverse problems, both weight selection based on a linearized version of the model was still found to perform quite well. Development of computationally viable optimal simultaneous source experimental design that accounts for the non-linear nature of the observation model is a subject for future work.

In this study, Gaussian assumptions were taken for the noise model and Tikhonov regularization was considered for the ill-posed case. These choices are by no mean exclusive, and in principle, derivation of optimal experimental design for simultaneous sources can be performed for other alternative noise models as well as for other forms of regularization.

So far incorporation of constraints was not considered (other than such that can be introduced in the form of exact penalty). Future research would be to generalize the proposed formulation to account for such.

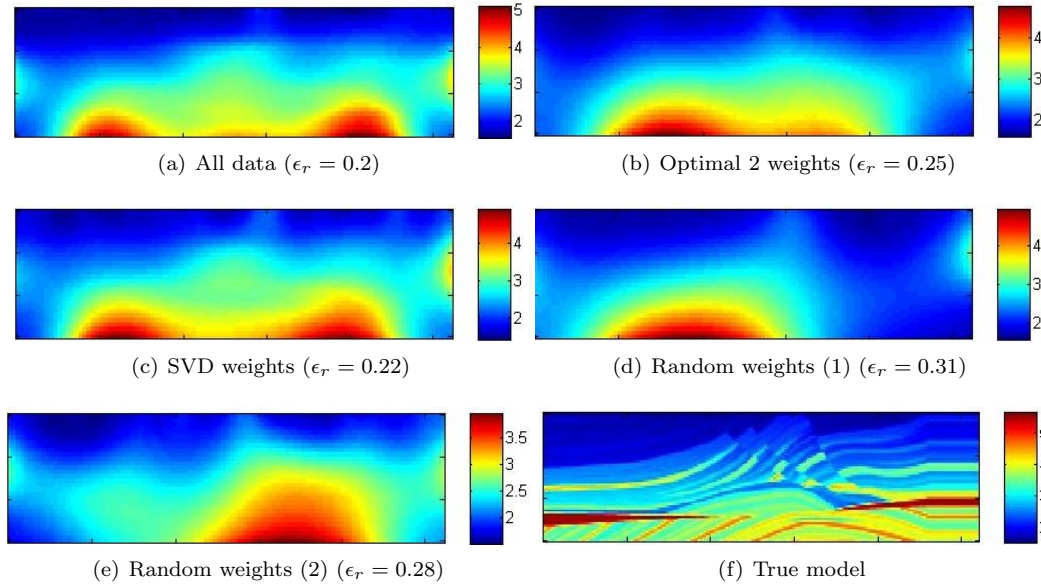


Figure 5: Reconstructions of the Marmousi velocity model using only the low frequency $f = 0.2$ data with all and with 2 simultaneous sources. We also indicate the relative reconstruction error $\epsilon_r = \frac{\|m_{rec} - m_{true}\|_2}{\|m_{true}\|_2}$.

References

- [1] S R Arridge. Topical review: Optical tomography in medical imaging. *Inverse Prob*, 15:41–93, 1999.
- [2] C.J. Beasley. A new look at marine simultaneous sources. *The Leading Edge*, 27(7):914–917, 2008.
- [3] A. Bourgeois, M. Bourget, P. Lailly, M. Poulet, P. Ricarte, and R. Versteeg. *The Marmousi experience*, volume 5-9, chapter Marmousi data and model. EAGE, 1991.
- [4] R. Casanova, A. Silva, and A. R. Borges. A quantitative algorithm for parameter estimation in magnetic induction tomography. *Meas. Sci. Technol.*, 15:1412–1419, 2004.
- [5] Darrell A. Coles and Frank Dale Morgan. A method of fast, sequential experimental design for linearized geophysical inverse problems. *Geophysical Journal International*, 178(1):145–158, 2009.
- [6] O. Dorn. A shape reconstruction method for electromagnetic tomography using adjoint fields and level sets,. *Inverse Problems*, 16:1119–1156, 2000.
- [7] E. Haber, M. Chung, and F. Herrmann. Solving pde constrained optimization with multiple right hand sides. *SIAM J. Optim.*, 22(3):739–757, 2012.
- [8] E Haber, L Horesh, and L Tenorio. Numerical methods for experimental design of large-scale linear ill-posed inverse problems. *Inverse Problems*, 24(5):055012, 2008.
- [9] E Haber, L Horesh, and L Tenorio. Numerical methods for the design of large-scale non-linear discrete ill-posed inverse problems. *Inverse Problems*, 26(2):025002, 2010.
- [10] P. C. Hansen. *Rank-Deficient and Discrete Ill-Posed Problems*. SIAM, Philadelphia, 1997.
- [11] F.J. Herrmann, Y.A. Erlangga, and T.T.Y. Lin. Compressive simultaneous full-waveform simulation. *Geophysics*, 74:A35, 2009.
- [12] L. Horesh, Haber E., and L. Tenorio. *Book Series on Computational Methods for Large-Scale Inverse Problems and Quantification of Uncertainty*, volume Optimal experimental




Figure 6: Fields for the 2 simultaneous sources for the case $f = 0.2$ using the various weight selection methods. Notice how the good weights cover the domain more or less uniformly.

design for the large-scale nonlinear ill-posed problem of impedance imaging. John Wiley & Sons, Ltd, 2009.

- [13] M.F. Hutchinson. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *J. Commun. Statist. Simul.*, 19, 1990.
- [14] Tom Lahmer. Optimal experimental design for nonlinear ill-posed problems applied to gravity dams. *Inverse Problems*, 27(12):125005, 2011.
- [15] Hansruedi Maurer, Andrew Curtis, and David E. Boerner. Recent advances in optimized geophysical survey design. *Geophysics*, 75(5):75A177–75A194, 2010.
- [16] V.A. Morozov. The discrepancy principle for solving operator equations by the regularization method. *Zh. Vychisl. Mat. Mat. Fiz*, 8:295–309, 1968.
- [17] S.A. Morton and C.C. Ober. Faster shot-record depth migrations using phase encoding. In *SEG Technical Program Expanded Abstracts*, volume 17, pages 1131–1134. SEG, 1998.
- [18] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, New York, 1999.
- [19] J. Romberg, R. Neelamani, C. Krohn, J. Krebs, M. Deffenbaugh, and J. Anderson. Efficient seismic forward modeling and acquisition using simultaneous random sources and sparsity. *Geophysics*, submitted, 2010.
- [20] L.A. Romero, D.C. Ghiglia, C.C. Ober, and S.A. Morton. Phase encoding of shot records in prestack migration. *Geophysics*, 65(2):426–436, 2000.
- [21] B. Symes. Source synthesis for waveform inversion. *Submitted*, 2010.
- [22] L. Tenorio, F. Andersson, M. de Hoop, and P. Ma. Data analysis tools for uncertainty quantification of inverse problems. *Inverse Problems*, 27:045001, 2011.

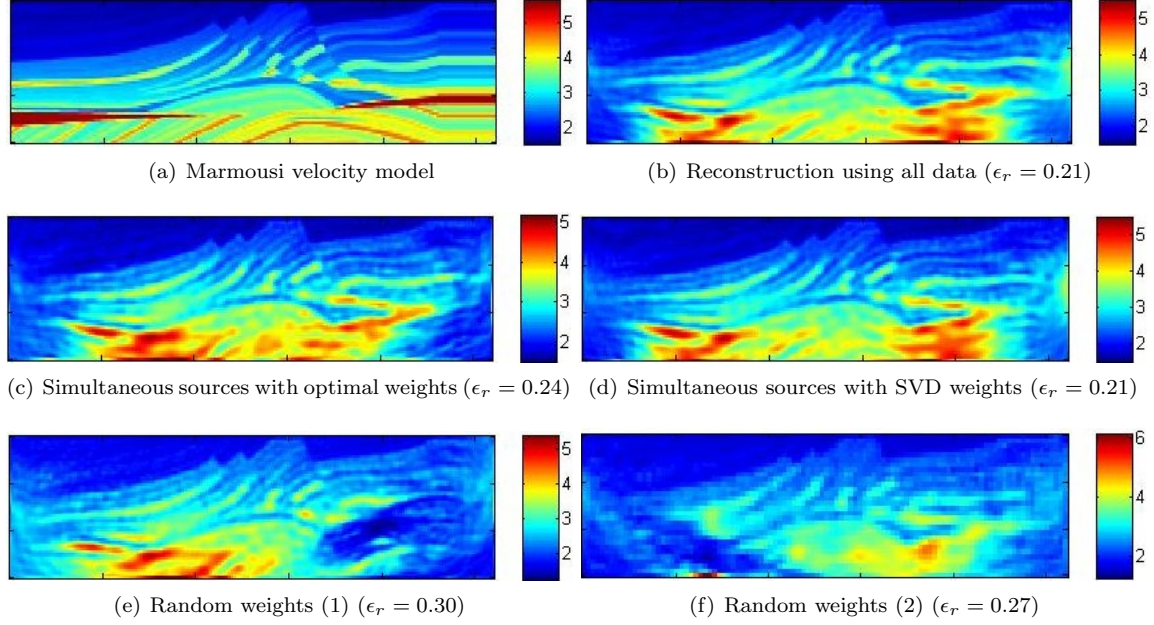


Figure 7: Reconstructions of the Marmousi velocity model using the protocol described. We compare the final result using all the data, and a simultaneous sources protocol with optimal, SVD, and random weights.

- [23] T.M.Habashy, A.Abubakar, A.Belani G.Pan, and Schlumberger. Full-waveform seismic inversion using the source-receiver compression approach. In *2010 SEG Annual Meeting, October 17 - 22, 2010*, volume 3353, Denver, Colorado, USA, 2010. Society of Exploration Geophysicists.