Technical Report

TR-2008-016

Optimal Experimental Design for the Large-Scale Nonlinear Ill-posed Problem of Impedance Imaging

by

Lior Horesh, Eldad Haber, Luis Tenorio

Mathematics and Computer Science

EMORY UNIVERSITY

Optimal Experimental Design for the Large-Scale Nonlinear Ill-posed Problem of Impedance Imaging

Lior Horesh[1], Eldad Haber[1] & Luis Tenorio[2]

[1]Emory University
[2]Colorado School of Mines

0.1 Introduction

Optimal experimental design (OED) of well-posed inverse problems is a well established field (e.g., Pukelsheim (2006) and references therein) but, despite the practical necessity, experimental design of ill-posed inverse problems and in particular ill-posed nonlinear problems has remained largely unexplored. We discuss some of the intrinsic differences between well- and ill-posed experimental designs, and propose an efficient and generic framework for optimal experimental design of nonlinear ill-posed problems in the context of impedance tomography.

We use discrete data modeled as

$$d = F(m; y) + \varepsilon, \tag{1}$$

Computational Methods for Large-Scale Inverse Problems and Quantification of Uncertainty. Edited by People on Earth
© 2001 John Wiley & Sons, Ltd

© XXXX John Wiley & Sons, Ltd

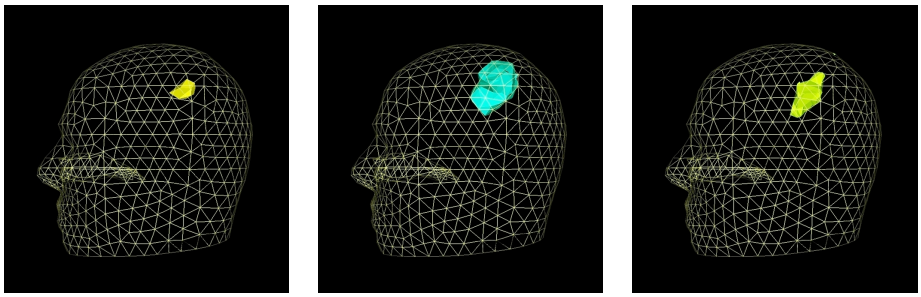

Figure 1 Impedance Tomography image recovery of the human head. The above models were recovered using three different data sets which corresponded to three different experimental designs. Despite the great difference in localization and intensity of the localized brain tumor, the data sets were in fact acquired from the same patient.

where d is a data vector, m is the model vector to be recovered, F is an observation model (i.e., forward operator), y is a vector (or matrix) of design parameters, and ε is a random noise vector.

We assume that even the noiseless problem is ill-posed in the usual sense (e.g., Hansen (1998)). In particular, there is no unique stable solution for the problem of finding m such that $F(m; y) = d$. Regularization techniques are used for finding stable and meaningful solution of such ill-posed problems. We consider a Tikhonov regularization approach where the solution $\widehat{m}$ is recovered by an optimization problem of the form

$$\widehat{m} = \underset{m}{\operatorname{argmin}} \left\{ \frac{1}{2} \left\| F(m; y) - d \right\|^2 + R(m) \right\}, \tag{2}$$

with R being a regularization functional that allows for the incorporation of a-priori information.

Most efforts in the field of inverse problems have been primarily devoted to the development of either numerical methods for the solution of the optimization problem (2), or methods for choosing a regularization functional $R(m)$. Little work has addressed the selection of the experimental settings y even though a proper design inherently offers better information extraction and bias control. Figure 1 demonstrates the importance of experimental design in impedance tomography. The images show three recovered models, each corresponds to a particular experimental design setting but conducted over the very same true model. If these models were acquired for presurgical assessment of brain tumor removal, for instance, the neurosurgeon would be faced the serious problem of deciding which region to operate.

While optimal experimental design for well-posed (over-determined) problems is well established (Pukelsheim 2006), the authors are aware of very few publications (e.g., Bardow (2008); Curtis (1999); Haber et al. (2008)) that address ill-posed (under-determined) problems and of even fewer that consider nonlinear problems. Yet, nowadays, many practical problems in biomedical imaging, geophysics, engineering as well as other physical sciences are intrinsically ill-posed and nonlinear.

In this paper we propose a novel formulation for optimal experimental design of ill-posed and nonlinear inverse problems and demonstrate its application to impedance tomography. The rest of the paper is organized as follows: In Section 0.2 we introduce the basic impedance tomography problem. In Section 0.3 we briefly review different optimal experimental design formulations for linear ill-posed problems and describe the shortcomings of these approaches in handling nonlinear ill-posed problems. In Section 0.4 we provide the rationale for our proposed formulation which is based on some basic statistical considerations. In Section 0.5 we describe the mathematical formulation of the suggested approach and explain the optimization framework developed for its solution. Section 0.6 provides numerical results for a realistic impedance tomography model problem. In Section 0.7 we discuss the numerical results as well as various technical matters related to the implementation of such formulation and close with some comments for future research.

0.2 Impedance Tomography

The inverse impedance tomography problem aims at the recovery of an impedance distribution in a domain given a set of boundary potential measurements for a known configuration of sources and receivers (Holder 2005). Acquisition is based on injection of low frequency currents from multiple sources while boundary potentials are being recorded by a set of receivers. Typically, the sources and receivers are both deployed over the domain boundary. In the context of the estimate defined by (2), the model m stands for the conductivity (or admittivity) distribution in the domain and the data vector d corresponds to the electric potentials measured on the boundary.

The observation (forward) model under consideration is derived from Maxwell's equations in the low frequency domain; they lead to the partial differential equation

$$\begin{array}{ll} \nabla \cdot (m \boldsymbol{\nabla} u) = 0 & u \in \Omega \\ m \boldsymbol{\nabla} u = Q & u \in \partial\Omega, \end{array} \tag{3}$$

where the model m is the conductivity, u the electric potential and Q are the sources.

Standard finite volume or finite element discretization (Ascher 2008; Vauhkonen et al. 1999; Vavasis 1996) leads to the construction of a discrete linear system. This system is solved for u given the conductivity model m:

$$A(m)u = D^\top M(m) D = P_Q Q, \tag{4}$$

where D is a sparse matrix holding the gradients of the basis functions (shape gradients) and $M(m)$ is a diagonal matrix that holds the conductivity values multiplied by the support (volumes) of the elements.

The forward model can be written as

$$F(m, V, Q) = V^\top P_V^\top A(m)^{-1} P_Q Q, \tag{5}$$

where the matrices P_Q and P_V map the source and measurement distributions onto the model grid.

In the context of the specific given forward model, another important physical restriction is current conservation. We assume that there is no dissipation of energy through the medium, and therefore the total amount of current driven in should be equal to the total amount of current drained out. This constraint implies that $e^\top Q = 0$, where e is a vector of ones.

0.3 Optimal Experimental Design - Background

We briefly review the basic framework of OED in the context of discrete linear well-posed inverse problems where $F(m; y) = K(y)m$.

The data are modeled as

$$d = K(y)\, m + \varepsilon, \tag{6}$$

where $K(y)$ is an $\ell \times k$ matrix representation of the forward operator that acts on the model m and depends on a vector of experimental parameters $y \in \mathcal{Y}$. The noise vector ε is assumed to be zero mean with iid entries of known variance σ^2. The experimental design question is the selection of y that leads to an optimal estimate of m.

In the well-posed case the matrix $K(y)^\top K(y)$ is nonsingular and for fixed y the least squares (LS) estimate $\widehat{m} = (K^\top K)^{-1} K^\top d$ is unbiased with covariance matrix $\sigma^2\, C(y)^{-1}$, where $C(y) = K(y)^\top K(y)$. One can then choose y so as to obtain a good LS estimate of m. Since $\widehat{m}$ is unbiased, it is common to assess its performance using different characteristics of its covariance matrix. For example, an A-optimal type experiment design prescribes a choice of y that minimizes the trace of $C(y)^{-1}$ subject to some reasonable constraints on y. If instead of the trace, the determinant or the ℓ_2-norm of $C(y)^{-1}$ is used, then the design is known as D- or E-optimal, respectively.

Designs that are solely based on the covariance matrix are unsuitable for ill-posed problems where estimators of m are most likely biased. In fact, the bias may be the dominant component of the error. Assume now that $K(y)^\top K(y)$ is singular or ill-conditioned. A regularized estimate of m can be obtained using penalized LS (Tikhonov regularization) with a smoothing penalty matrix L

$$\widehat{m} = \arg\min \frac{1}{2} \left(K(y)m - d \right)^\top \left(K(y)m - d \right) + \frac{\alpha}{2} \| Lm \|^2,$$

where $\alpha > 0$ is a fixed regularization parameter that controls the balance between the data misfit and the smoothness penalty. Assuming that $K(y)^\top K(y) + \alpha\, L^\top L$

is nonsingular, the estimator is $\widehat{m} = \left(K(y)^\top K(y) + \alpha\, L^\top L \right)^{-1} K(y)^\top d$, whose bias can be written as

$$\mathbf{Bias}(\,\widehat{m}\,) = \mathbf{E}\,\widehat{m} - m = -\alpha\, \left(K(y)^\top K(y) + \alpha\, L^\top L \right)^{-1} L^\top L\, m.$$

Since the bias is independent of the noise level, it cannot be reduced by averaging repeated observations. The noise level affects the variability of $\widehat{m}$ around its mean $\mathbf{E}\,\widehat{m}$. Thus, this variability and the bias ought to be taken into account when choosing an estimator of m.

Define $\mathcal{B}(y, m) := \| \mathbf{Bias}(\widehat{m}) \|^2$ and $\mathcal{V}(y) = \mathbf{E}\, \| \widehat{m} - \mathbf{E}\,\widehat{m} \|^2/\sigma^2$. The sum of these two error terms provides an overall measure of the expected performance of $\widehat{m}$. This is essentially the mean squared error (MSE), or risk, of $\widehat{m}$. More precisely, the MSE of $\widehat{m}$ is defined as $\mathbf{E}\, \| \widehat{m} - m \|^2$, which can also be written as

$$\begin{aligned}\mathcal{R}(\,\widehat{m}\,) &= \mathbf{E}\, \| \widehat{m} - \mathbf{E}\,\widehat{m} + \mathbf{E}\,\widehat{m} - m \|^2 = \| \mathbf{E}\,\widehat{m} - m \|^2 + \mathbf{E}\, \| \widehat{m} - \mathbf{E}\,\widehat{m} \|^2 \\ &= \| \mathbf{Bias}(\widehat{m}) \|^2 + \mathbf{E}\, \| \widehat{m} - \mathbf{E}\,\widehat{m} \|^2 = \alpha^2\, \mathcal{B}(y, m) + \sigma^2\, \mathcal{V}(y).\end{aligned}$$

The overall idea is then to define optimization problems so as to control a measure of the performance of $\widehat{m}$ that takes into account its bias and stochastic variability. A natural choice would be the MSE but this measure depends on the unknown m. However, in many practical applications it is possible to obtain examples of plausible models. For example, there are geostatistical methods to generate realizations of a given media from a single image. Also, in many applications, and in particular in medical imaging, it is possible to obtain a set of likely example models $\mathcal{M} = \{m_1, ..., m_s\}$. In medical imaging such models can be proposed based on prognosis studies. These studies provide useful statistics of common pathologies and their classes, and therefore can be used for the construction of models associated with likely pathologies. Basically, this set of example models enables one to implicitly define regions of interest and similarly regions of lack of interest; either of which provides valuable guidance for the experimental setup. For some applications, the designer may even have access to the distribution function of the desired model itself (or at least a trustworthy approximation of it).

Let $\mathcal{M} = \{m_1, ..., m_s\}$ be examples of plausible models which will be henceforth called training models. As with Bayesian optimal designs (Chaloner & Verdinelli 1995), it is assumed that these models are iid samples from an unknown multivariate distribution π, only that this time we use the sample average

$$\widehat{\mathbf{E}}_\pi\, \mathcal{B}(y, \mathcal{M}) = \frac{1}{s} \sum_{i=1}^{s} \mathcal{B}(y, m_i), \tag{7}$$

which is an unbiased estimator of $\mathbf{E}_\pi\, \mathcal{B}(y, m)$. We thus define the following approximation of the MSE

$$\widehat{\mathcal{R}}(\,y\,) = \alpha^2\, \widehat{\mathbf{E}}_\pi\, \mathcal{B}(y, m) + \sigma^2\, \mathcal{V}(y). \tag{8}$$

This type of empirical approach is commonly used in machine learning, where estimators are trained using iid samples.

0.4 OED for Nonlinear Ill-Posed Problems

Nonlinear problems are inherently more difficult. For example, there is no closed form solution of the Tikhonov estimate $\widehat{m}$ and its MSE cannot be easily separated into bias and variance components $\mathcal{B}$ and $\mathcal{V}$; linearity allowed us to find closed expressions for the expected value over the noise distribution. This time we estimate such expected value in two stages. First, we generate a noise sample ε_j that is added to $F(m_i; y)$; second, we solve the inverse problem and obtain an estimate $\widehat{m}_{i,j}$, which is then compared to m_i. We repeat this procedure for multiple noise realizations and for each training model. This sequence defines an empirical MSE estimate

$$\widehat{\mathcal{R}}(\, y\,; \mathcal{M}, \widehat{m}) = \frac{1}{2ns} \sum_{i=1}^{s} \sum_{j=1}^{n} \| \, \widehat{m}_{i,j}(y) - m_i \, \|^2. \tag{9}$$

Since the design problem may still be underdetermined, the design parameter y might oscillate wildly and may generate designs that overfit the training set. It is therefore important to provide additional constraints or preferences related to the design parameters.

There are several desirable properties one may consider for favoring one design over another. Depending on the application, some reasonable and popular preferences may be: to shorten acquisition time, to reduce the number of sources/receivers or to achieve higher resolution and spatial selectiveness (i.e., targeting energy towards regions of interest while avoiding application of radiation towards neighboring background tissues). All these preferences pursue , of course, the final objective of finding the best model estimate given the available resources.

For most geophysical problems as well as medical applications, control over the number of sources and/or receivers is desirable. Let Q and V be, respectively, the vectors of sources and receivers and set $y = \{Q, V\}$. A sparsity preference for the number of active sources and receivers can be formulated by the addition of an ℓ_1-norm penalty term to the empirical risk

$$\widehat{\mathcal{R}}_1(\, y; \mathcal{M}, \widehat{m}) = \frac{1}{2ns} \sum_{i=1}^{s} \sum_{j=1}^{n} \| \, \widehat{m}_{i,j}(y) - m_i \, \|^2 + \beta \, \|y\|_1 \, . \tag{10}$$

The design problem requires the minimization of this regularized empirical risk while maintaining the conventional inverse problem (2) and forward model feasible. This can be formulated as a bi-level optimization problem (Alexandrov 1994; Bard 2006; Colson et al. 2007; Dempe & Gadhi 2007) that reads as follows

$$\min_{y} \;\; \frac{1}{2ns} \sum_{i=1}^{s} \sum_{j=1}^{n} \| \, \widehat{m}_{i,j}(y) - m_i \, \|^2 + \beta \, \|y\|_1 \tag{11}$$

$$\text{s.t.} \;\; \widehat{m}_{i,j} = \operatorname*{argmin}_{m} \;\; \frac{1}{2} \, \|F(m; y) - d_{i,j}(m_i; y)\|^2 + R(m), \tag{12}$$

where $d_{i,j}(m_i; y) = F(m_i; y) + \varepsilon_j$. With a slight abuse of notation, we shall henceforth denote $\mathcal{R}_1$ by $\mathcal{R}$, and will write

$$\psi(m; m_i, y) = \frac{1}{2} \left\| F(m; y) - d_{i,j}(m_i; y) \right\|^2 + R(m).$$

Before we proceed with discussing an appropriate optimization framework for the solution of the above problem, we would like to make an important distinction. Beyond the fundamental difference between linear and nonlinear inversion that distinguishes this work from our previous work (Haber et al. 2008), another important difference is related to the design mechanism itself. In the previous work, we were trying to find the best subset out of a given set of excitation and measurement vectors, while here we predefine their number, but control their content by inducing sparsity. The current mechanism typically requires larger number of parameters, which, on the one hand offers greater solution flexibility and optimality but on the other hand, introduces higher computational demands.

In principle, it is also possible to employ Bayesian OED (Chaloner & Verdinelli 1995) for handling nonlinear OED problems of this sort. However, the strict necessity for complete knowledge of the pdf of the model is infeasible for most practical applications. Moreover, even when this knowledge is available, the great computational complexity of this approach makes it less appealing for our purposes.

0.5 Optimization Framework

0.5.1 General Scheme

General bi-level optimization problems are difficult because the inner optimization problem may have multiple solutions and/or may be non-convex. Here, we make the assumption that the inner level is convex (or at least convex over the set of y's of interest). This leads to simpler algorithms (Alexandrov 1994). Thus, we replace the inner optimization problem by the necessary conditions for a minimum solving

$$\min_{y, m=(m_{i,j})} \quad \mathcal{R}(y; \mathcal{M}) = \frac{1}{2ns} \sum_{i=1}^{s} \sum_{j=1}^{n} \left\| m_{i,j} - m_i \right\|^2 + \beta \left\| y \right\|_1 \tag{13}$$

$$\text{s.t.} \quad c_{i,j}(m_{i,j}, y) := \frac{\partial \psi_{i,j}(m; m_i, y)}{\partial m_{i,j}} = 0. \tag{14}$$

The necessary conditions for a minimum are

$$\mathcal{R}_{m_{i,j}} + (c_{i,j})^\top_{m_{i,j}} \lambda_{i,j} = 0 \tag{15a}$$

$$\mathcal{R}_y + \sum_{i,j}^{s,n} (c_{i,j})^\top_y \lambda_{i,j} = 0 \tag{15b}$$

$$c_{i,j}(m_{i,j}, y) = 0, \tag{15c}$$

where $\lambda_{i,j}$ are Lagrange multipliers. Although it is possible to develop an efficient method for the solution of the system for $m_{i,j}$ and y simultaneously, we have chosen an unconstrained optimization approach. The advantage of such approach is that storage of all $m_{i,j}$'s and $\lambda_{i,j}$'s is not required, and the complete optimization procedure can be executed by solving each inverse problem separately. More precisely, we first solve the $s \times n$ decoupled nonlinear systems (15c) for $m_{i,j}$ given y, then we solve the other $s \times n$ decoupled linear systems (15a) for $\lambda_{i,j}$. For each $m_{i,j}$ and $\lambda_{i,j}$ computed, the reduced gradient (15b) is updated. Thus, even when the number of training models and noise realizations is excessively large, such a design process can be conducted over modest computational architectures.

In order to avoid the non-differentiability property of the ℓ_1-norm at zero, we employ here the Iterated Reweighted Least Squares (IRLS) approximation (O'Leary 1990; Street et al. 1988). This approach has been successfully used for ℓ_1 inversion in many practical scenarios (Sacchi & Ulrych 1995; Vogel 2001; Whittall & Oldenburg 1992). In the IRLS approximation, the ℓ_1-norm is replaced by a smoothed version of the absolute value function $\|x\|_{1,\epsilon}$ where

$$|t|_\epsilon := \sqrt{t^2 + \epsilon} \quad \text{and} \quad \|x\|_{1,\epsilon} := \sum_i |x_i|_\epsilon.$$

As discussed in Haber et al. (2000), obtaining the sensitivities of $m_{i,j}$ with respect to y is straightforward; they can be written as

$$J = -\sum_{i,j}^{s,n} (c_{i,j})^{-1}_{m_{i,j}} (c_{i,j})_y.$$

The sensitivities are then used with IRLS (O'Leary 1990) to obtain an approximation of the Hessian

$$H = J^\top J + \beta \operatorname{diag}\left(\min(|y|^{-1}, \epsilon^{-1})\right),$$

which in turn is used to define the update

$$y \leftarrow y - \tau H^{-1} g, \tag{16}$$

where τ is a line search step size and $H^{-1}g$ is computed using the conjugate gradient method. Note that computation of the update step can be performed without the explicit construction of either the matrix $(c_{i,j})_{m_{i,j}}$ or its inverse. These matrices are only accessed implicitly via matrix vector products.

0.5.2 Application to Impedance Tomography

We now discuss a specific application of the approach presented above to impedance tomography. As previously discussed, we consider an experimental design problem of positioning and activating sources and receivers over a predefined region of permissible locations. Such region could be the entire boundary $\partial\Omega$ of a domain Ω as often occurs in medical imaging or, alternatively, any sub-domain $\Omega_r \subseteq \Omega$.

We consider the Tikhonov regularized estimate (2) with $F(m; y)$ equal to the forward operator (5) and $R(m) = \alpha \|Lm\|^2$, with L a discrete derivative operator that penalizes local impedance inhomogeneities.

At this point we comment on the choice of the regularization parameters α and β. The parameter α is chosen prior to conducting the inversion process; we choose it so as to balance the expected misfit with the norm of $R(m)$. For iid Gaussian noise the expected value of $\frac{1}{2}\|d - F(m; y)\|^2$ is $\ell\sigma^2$. Choosing $\alpha \approx \ell\sigma^2/R(m)$ yields reasonable model reconstructions. The choice of β is user-dependant. The designer may tune this parameter according to some sparsity preferences. The higher the sparsity of the active sources and receivers, the larger the value of β. Conversely better model recovery can be achieved with lower levels of β at the expense of lower sparsity. Further discussion regarding the choice of β for OED of linear problems can be found in Haber et al. (2008).

For our impedance tomography design problem, the optimality criterion prescribes the construction of an experimental design that minimizes the number of active sources and receivers, as well as minimizing the description error between the recovered and given training models. This configuration must comply with a feasible forward model and an inverse solution, and should be consistent with the acquisition noise level. Such design can be obtained with a sparsity requirement imposed over the sources and measurements vectors in the form of ℓ_1-norm penalty. Thus, we set $y := \{V, Q\}$ and accordingly denote the forward operator by $F(m; V, Q)$. We assume that a collection of feasible, representing models $\mathcal{M} = \{m_1, \ldots, m_s\}$ is at our disposal, for which we can compute a measurements set using the forward model $F(m; V, Q)$.

The design problem given in (11) - (12) can be formulated as follows

$$\min_{V,Q} \ \mathcal{R}(V, Q; \mathcal{M}, \widehat{m}) \tag{17a}$$

$$\text{s.t. } \widehat{m} = \operatorname*{argmin}_{m} \ \psi(m; V, Q), \tag{17b}$$

where the objective function (i.e., the regularized empirical risk $\mathcal{R}$) and the linearized constraints c, as in (13) and (14) respectively, are given by:

$$\mathcal{R}(V, Q; \mathcal{M}, m) = \frac{1}{2} \sum_{i,j}^{s,n} \|m_{i,j} - m_i\|^2 + \beta_1 \|V\|_1 + \beta_2 \|Q\|_1$$

$$c_{i,j}(m_{i,j}; V, Q) = S_{i,j}^\top V^\top P_V^\top \left((A(m_{i,j})^{-1} - A(m_i)^{-1}) P_Q Q - \epsilon_j \right) + \alpha L^\top L m_{i,j},$$

where $S_{i,j}$ is the sensitivity matrix, that is the Fréchet derivative (Jacobian) of the forward operator. This operator represents the sensitivity of the acquired data to small changes in the model $S_{i,j} := \partial d_{i,j}/\partial m$. For this formulation, $S_{i,j}$ can be derived by implicit differentiation using the relation (5) (Haber et al. 2000).

In order to solve (13) - (14), we shall now evaluate the remaining components of (15a) - (15b). Since the empirical risk is convex on $m_{i,j}$, we have $\mathcal{R}_{m_{i,j}} = m_{i,j} -$

m_i. The derivatives of the risk with respect to the design parameters P and Q are obtained using the IRLS relation given in (0.5.1):

$$\mathcal{R}_V = \beta_1 \operatorname{diag}\left(\frac{1}{|V|_{1,\epsilon}}\right) V, \quad \mathcal{R}_Q = \beta_2 \operatorname{diag}\left(\frac{1}{|Q|_{1,\epsilon}}\right) Q.$$

The derivative of the linearized constraint $c_{i,j}$ with respect to the model $m_{i,j}$ is in fact the conventional Hessian of the inverse problem (2)

$$(c_{i,j})_{m_{i,j}} = S_{i,j}^\top S_{i,j} + \alpha L^\top L + M_{i,j},$$

where $M_{i,j}$ stands for the second order derivatives that can be computed implicitly (Haber et al. 2000). The derivatives $(c_{i,j})_V$ and $(c_{i,j})_Q$ can be calculated from the structures of $S_{i,j}$ and $M_{i,j}$. The reduced space gradients $g(V)$ and $g(Q)$ are obtained using the relations

$$g(V) = \mathcal{R}_V + \sum_{i,j}^{s,n} (c_{i,j})_V^\top (c_{i,j})_{m_{i,j}}^{-\top} \mathcal{R}_{m_{i,j}}$$

$$g(Q) = \mathcal{R}_Q + \sum_{i,j}^{s,n} (c_{i,j})_Q^\top (c_{i,j})_{m_{i,j}}^{-\top} \mathcal{R}_{m_{i,j}}.$$

Next, the reduced Hessians H_V and H_Q can be derived as in (0.5.1) and, similarly, the updates for V and Q at each iteration of the reduced space IRLS as in (16).

Remark. We have used a conventional reparametrization of the model space in order to comply with the broad dynamic range requirement and therefore set $m := \log(\gamma)$, where γ is the conductivity.

0.6 Numerical Results

We now show results of applying the optimal experimental design framework described in Section 0.5.2 to a realistic design problem of positioning and activating sources and receivers for impedance imaging of a human head.

The setup involves two stages. In the first stage the source and receiver vectors are optimized to comply with the sparsity requirement using a set of given training models. In the second stage, several unseen (test) models are used to assess the performance of the obtained design.

For training purposes, three example models of impedance perturbations in the mid, left and right parietal brain lobe (marked in red on the head diagrams in Figure 4) are considered. The design problem consists of 16 receiver vectors V and 16 source vectors Q that were initialized at random.

Figure 2 shows the behavior of the empirical risk and MSE as a function of iteration number. The regularized empirical risk is reduced by an order of magnitude in the first six iterations, while the MSE itself was reduced by a factor of about two. On

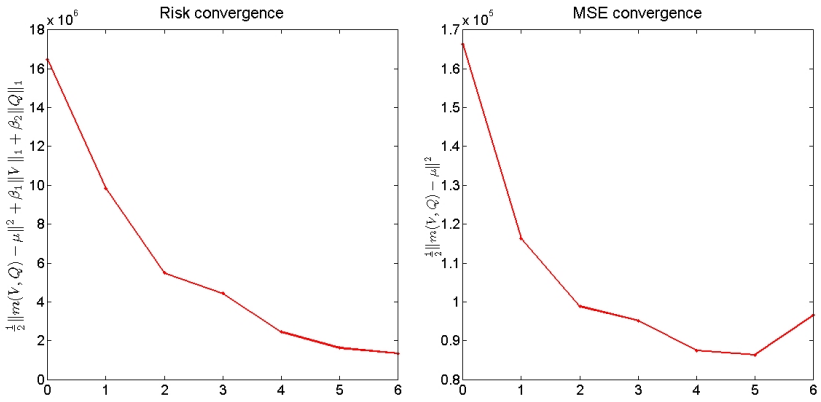

Figure 2 Experimental design of 16 receiver and 16 source vectors for three models. Left: Empirical risk convergence; Right: MSE convergence.

the 6^{th} iteration the MSE start to increase. At this point the tradeoff between sparsity and reconstruction accuracy begins to be evident. We stopped the optimization process when the relative change in the risk dropped below 1%.

Figure 3 shows the ensemble of source Q and measurement vectors V for the initial and optimized designs. The figure shows that the optimized vectors are clearly sparser than those obtained with the initial design. Another way to visualize the difference in the designs is shown in Figure 4, where 4 out of the 16 source vectors are plotted as colored circles at the active sources locations. Again, it can be clearly observed that the optimized design requires the deployment of a smaller number of active sources, which was the original objective.

While the first three figures illustrate the performance of the design procedure at the learning stage, the performance of the obtained design for an unseen (testing) model is shown in Figure 5. It presents a comparison between a 3D impedance tomography model that was recovered using the initial and optimal designs. The impedance intensities of the perturbations recovered with the optimized design are clearly better than those obtained with the initial design. The quantitative results for models obtained with the optimized design indicate that the recovered models were almost identical to the true model. Moreover, while models recovered using the initial design were cluttered with artifacts, models recovered with the optimized design show higher qualitative structural similarity with the true models.

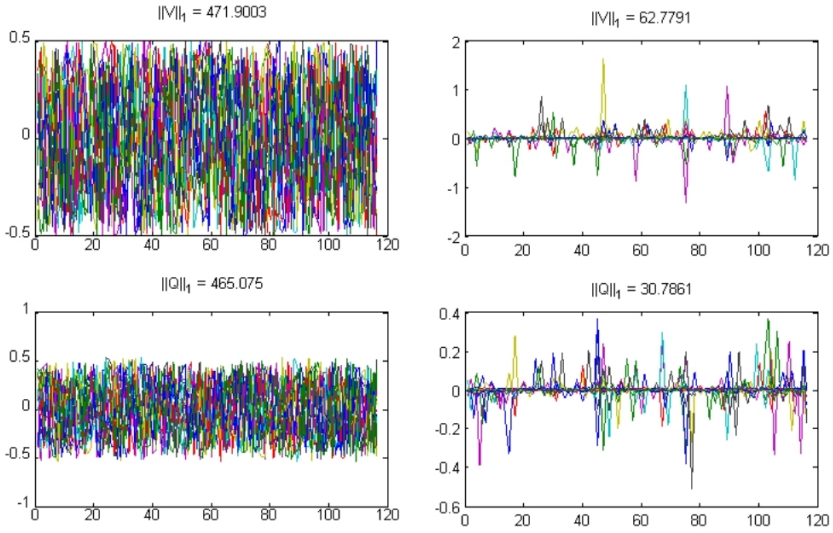

Figure 3 Sparsity pattern of 16 receiver vectors V (above) and 16 source vectors Q (below) for three training models. Left: Initial experimental design. Right: Optimized design after six iterations.

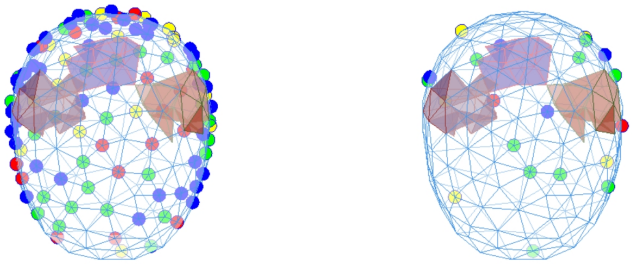

Figure 4 Human-head-shaped training models from a transversal projection. The red objects inside the head represent impedance perturbations as provided by the training example models. The colored dots located on vertices of the head mesh represent dominant locations for positioning sources (Only 4 out of the 16 vectors (colors) were presented here to avoid over-cluttering). Left: Initial design. Right: Optimized design.

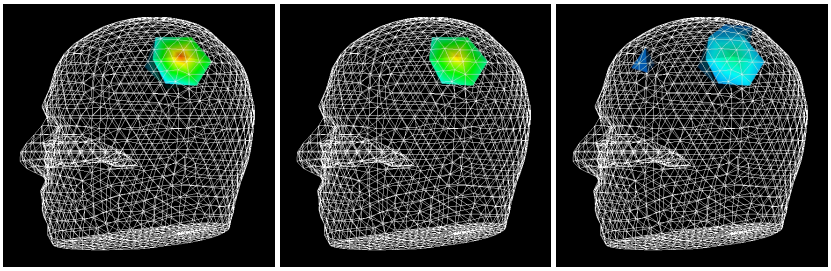

Figure 5 Testing set performance assessment for 16 receiver vectors V and 16 source vectors Q configuration. Left: True test model; Middle: Optimally designed recovered test model; Right: test model recovered using the initial design.

0.7 Discussion and Conclusions

We have presented a statistical formulation for optimal experimental design of nonlinear ill-posed problems. Such problems frequently arise in medical imaging and geophysics, but also in many other fields such as chemistry and aerospace engineering. For each of the different applications, the construction of an effective design procedure requires careful thought. Nevertheless, the rationale presented in this study is generic and therefore transferable to a broad range of important applications.

The numerical results have shown that the proposed optimal experimental design can provide substantial improvement in image quality, at both the quantitative and qualitative levels. But there are several outstanding issues that still need to be addressed. The first of which is the choice of regularized empirical risk measure. This choice is crucial for obtaining an effective design. We have used an ℓ_2-norm discrepancy between training and recovered models for the empirical risk measure, and an ℓ_1-norm penalty on y to promote sparsity. These choices may seem natural for the impedance design problem considered here but in general the choice will depend on the specific application.

Obviously, the choice of regularized estimate of the inverse problem is also important. We have chosen Tikhonov regularization with an ℓ_2 data misfit and smoothness penalty to promote smooth model estimates. Such smoothness may not be appropriate in some applications. In the numerical examples above, all model examples comprised sharp parameter changes, which may have impaired, to some extent, the overall performance of the design procedure.

One defining feature of the proposed formalism is that the choice of inversion method is embedded in the design problem. Thus, different inversion frameworks will inherently result in different designs. The fact that the design reflects the inversion settings may be regarded as a hindrance, as one may expect the design to be solely related to the physical properties of the problem. Yet, in reality the choice

of an inversion framework is mandatory. Moreover, since each inversion framework introduces distinctive a-priori information and individual degrees of belief one has attributed to the solution, the incorporation of such knowledge (if justified) into the design is essential.

Another fundamental issue is the source of training models. We have deliberately avoided a discussion for retrieving proper training examples as it is very much problem dependent. It was assumed that they were provided by a trusted source and as such were representative of the type of model to be recovered.

In addition to the source of models, their number is an important variable that determines how well the empirical risk approximates the true risk. A careful study of this question will require empirical process theory as in Vapnik (1995) and Rakhlin (2006).

Since several nonlinear inversion procedures are required at each design iteration, the design process is time consuming. Apart from the explicit construction of the forward model discretization matrix, the implementation of a bi-level optimization framework in this study relied on implicit calculation of matrices and parallel processing of the inversions of independent example models. This implementation enabled the processing of large-scale models despite memory limitations. Nevertheless, the overall design process is still computationally intensive and time consuming. For several applications, such as geophysical surveys, offline processing time of several hours or even a few days prior to a survey is certainly acceptable but, for some medical applications, development of faster design procedures may be essential.

Another issue that needs to be addressed is the sensitivity analysis. A better insight into the interplay between the design requirement and the sensitivity gain/loss can be obtained by observing changes in the sensitivity map that evolves throughout the design process. Such analysis can illustrate and explain how sensitivity and independence of measurable information originating in regions of interest (i.e., where perturbations are expected) form higher local sensitivity compared with irrelevant areas.

The practical aim of this work is mainly to help improve current designs and, therefore, whenever the designer is able to provide a good initial design, the formulation provided here will result in an improved experimental design. In many practical settings improvement of a design at hand may suffice and, therefore, an exhaustive solution of the design problem may not be needed.

Acknowledgements. The authors wish to thank Michele Benzi, Raya Horesh, James Nagy and Alessandro Veneziani for their valuable advice and criticism of this work. In addition we wish to express our gratitude to our funding bodies the DOE (DE-FG02-05ER25696-A002), NSF (DMS 0724759, 0724717 and 0724715, CCF 0427094 and 0728877) and NIH (HL085417-01A2).

Bibliography

Alexandrov N and Dennis JE 1994 Algorithms for bilevel optimization. In *AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, pp.810–816

Ascher UM 2008 *Numerical Methods for Evolutionary Differential Equations*. SIAM.

Bard JF 2006 *Practical Bilevel Optimization: Algorithms and Applications (Nonconvex Optimization and its Applications)*. Springer.

Bardow A 2008 Optimal experimental design for ill-posed problems, the meter approach. *Computers and Chemical Engineering* **32**, 115–124.

Boyd S and Vandenberghe L 2004 *Convex Optimization*. Cambridge University Press.

Chaloner K and Verdinelli I 1995 Bayesian experimental design: A review. *Statistical Science* **10**, 273–304.

Colson B, Marcotte P and Savard G 2007 An overview of bilevel optimization. *Annals of Operations Research* **153**, 235–256.

Curtis A 1999 Optimal experimental design: cross borehole tomographic example. *Geophysics J. Int.* **136**, 205–215.

Dempe S 2000 A bundle algorithm applied to bilevel programming problems with non-unique lower level solutions. *Comput. Optim. Appl.* **15**, 145–166.

Dempe S and Gadhi N 2007 Necessary optimality conditions for bilevel set optimization problems. *J. of Global Optimization* **39**, 529–542.

Haber E, Ascher UM and Oldenburg DW 2000 On optimization techniques for solving nonlinear inverse problems. *Inverse problems* **16**, 1263–1280.

Haber E, Horesh L and Tenorio L 2008 Numerical methods for experimental design of large-scale linear ill-posed inverse problems. *Inverse Problems*. In Press.

Hansen PC 1998 *Rank-Deficient and Discrete Ill-Posed Problems*. SIAM.

Holder DS 2005 *Electrical Impedance Tomography - Methods, History, and Applications*. Institute of Physics.

O'Leary DP 1990 Robust regression computation computation using iteratively reweighted least squares. *SIAM J. Matrix Anal. Appl.* **11**, 466–480.

Pukelsheim F 1993 *Optimal Design of Experiments*. SIAM.

Rakhlin A 2006 *Applications of Empirical Processes in Learning Theory: Algorithmic Stability and Generalization Bounds*. Phd Thesis, Dept. of Brain and Cognitive Sciences, Massachusetts Institute of Technology.

Sacchi MD and Ulrych TJ 1995 Improving resolution of Radon operators using a model reweighted least squares procedure. *Journal of Seismic Exploration* **4**, 315–328.

Street JO, Carroll RJ and Ruppert D 1988 A note on computing robust regression estimates via iteratively reweighted least squares. *The American Statistician* **42**, 152–154.

Vapnik V 1995 *The Nature of Statistical Learning Theory*. Springer.

Vauhkonen PJ, Vauhkonen M, Savolainen T and Kaipio JP 1999 Three-dimensional electrical impedance tomography based on the complete electrode model. *IEEE Trans. Biomedical Engineering* **46**, 1150–1160.

Vavasis SA 1996 Stable finite elements for problems with wild coefficients. *SIAM J. Numer. Anal.* **33**, 890–916.

Vogel C 2001 *Computational Methods for Inverse Problems*. SIAM.

Whittall KP and Oldenburg DW 1992 *Inversion of Magnetotelluric Data for a One Dimensional Conductivity*. SEG Monograph V.5.