

Position Registration from Voltage Measurements

Fadil Santosa* and Carl Toews†

Abstract

We are interested in determining the position of an electrode within a bounded homogeneous region B . Across this region we apply three orthogonal voltage potentials, and for each potential collect data corresponding to voltages at the probe. Were the fields linear, the positions of the probe could be read directly from the voltage data. Unfortunately, the unknown conductivity of the medium surrounding B induces nonlinear fields, and the problem becomes an inverse problem for which the objective is to determine the probe positions and the fields.

To approach this task, we model each field as linear plus a perturbation term with a low order expansion in a known harmonic basis. We then propose and analyze an iterative algorithm to solve the problem in a least-squares sense, focusing on the behavior of the solution in the limit as the number of measurements becomes large. The method is assessed on simulated data.

1 Introduction

We are interested in using electric field measurements to register the position of an electrode within some compact region $B \in \mathbb{R}^3$. Figure 1 depicts one possible configuration for this problem: the region B is the cube $[-1, 1]^3$, the medium within B is homogeneous, and the surrounding medium is of variable and unknown conductivity. A voltage potential difference of 4 units is applied at $x = -2$ and $x = 2$, creating a nonlinear voltage potential within the cube. Suppose now we insert a two-sensor probe inside the cube. While we are able to access the cube with the probe, neither the position nor the orientation of this probe can be measured. What can be measured are the voltage potentials at the two electrodes. After recording these voltages, we next apply a voltage potential difference of 4 units at $y = -2$ and $y = 2$, record these voltages, and finally do the same thing with the potential difference applied at $z = -2$ and $z = 2$. Once the voltage measurements at the two electrodes are recorded for all three potential fields, the probe is moved, and the measurement repeated.

If the fields were linear, the position of the electrode and its orientation at each measurement location can be read directly from the voltages. When the fields are nonlinear and unknown, we cannot register the positions of the electrodes without solving an inverse problem. The inverse problem is to simultaneously determine the electrode positions and estimate the field nonlinearities.

The problem described here in some generality has a potential application in biomedical imaging, where one is interested in registering the position of an organ.

*School of Mathematics, University of Minnesota, Minneapolis, MN 55455, email: santosa@math.umn.edu

†Institute for Mathematics and its Applications (IMA), University of Minnesota, Minneapolis, MN 55455, email: toews@ima.umn.edu

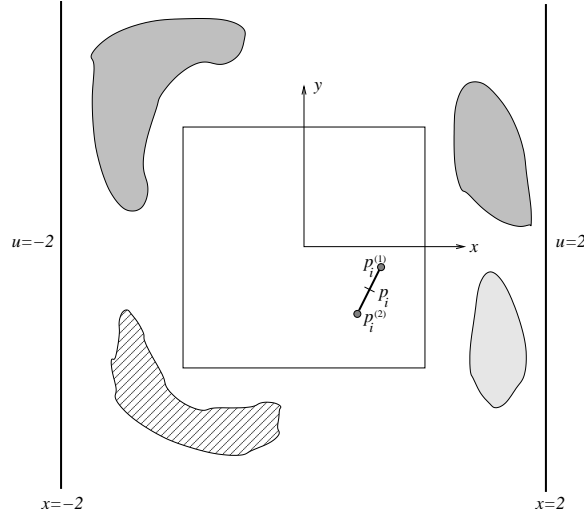


Figure 1: The measurement setup. The cube $B = [-1, 1]^3$ is shown as being surrounded by material with different conductivities. Voltage potential difference of 4 units is applied at $x = -2$ and $x = 2$. A two-sensor probe is inserted into the cube. We do not have access to its location and orientation within the cube, but we measure and record the voltages at the sensors. Next, the voltage potential difference is applied at $y = \pm 2$ and measurement is taken. Finally, the voltage potential difference is applied at $z = \pm 2$. Once the voltage measurements at the two electrodes have been recorded for the three potential field, the probe is moved, and the measurements are repeated.

To give a mathematical description of the measurement process, let $\mathbf{p}_i = (x_i, y_i, z_i)$ be the i th position of the centroid of the two-electrode probe. The electrodes are then located at

$$\mathbf{p}_i^j = \mathbf{p}_i \pm r(\cos \alpha_i \cos \beta_i, \cos \alpha_i \sin \beta_i, \sin \alpha_i) =: (x_i^{(j)}, y_i^{(j)}, z_i^{(j)}), \quad j = 1, 2. \quad (1)$$

for some orientation (α_i, β_i) . Here r is the half-distance between electrodes. The voltage field within the box for the three applied potentials are

$$\begin{aligned} u(x, y, z) &= x + \pi_x(x, y, z), \\ v(x, y, z) &= y + \pi_y(x, y, z), \\ w(x, y, z) &= z + \pi_z(x, y, z). \end{aligned} \quad (2)$$

The functions π_x , π_y and π_z are harmonic since the box contains no sources. At each probe location, we measure

$$(u_i^{(1)}, v_i^{(1)}, w_i^{(1)}, u_i^{(2)}, v_i^{(2)}, w_i^{(2)}).$$

The inverse problem is to use the equations

$$\begin{bmatrix} u_i^{(1)} \\ v_i^{(1)} \\ w_i^{(1)} \\ u_i^{(2)} \\ v_i^{(2)} \\ w_i^{(2)} \end{bmatrix} = \begin{bmatrix} x_i^{(1)} + \pi_x(x_i^{(1)}, y_i^{(1)}, z_i^{(1)}) \\ y_i^{(1)} + \pi_y(x_i^{(1)}, y_i^{(1)}, z_i^{(1)}) \\ z_i^{(1)} + \pi_z(x_i^{(1)}, y_i^{(1)}, z_i^{(1)}) \\ x_i^{(2)} + \pi_x(x_i^{(2)}, y_i^{(2)}, z_i^{(2)}) \\ y_i^{(2)} + \pi_y(x_i^{(2)}, y_i^{(2)}, z_i^{(2)}) \\ z_i^{(2)} + \pi_z(x_i^{(2)}, y_i^{(2)}, z_i^{(2)}) \end{bmatrix}, \quad i = 1, \dots, n, \quad (3)$$

to estimate the probe position (x_i, y_i, z_i) , probe orientation (α_i, β_i) , and voltage fields $\pi_x(x, y, z)$, $\pi_y(x, y, z)$, and $\pi_z(x, y, z)$.

If the fields in (2) were linear, that is $\pi_x = \pi_y = \pi_z = 0$, then the problem is entirely trivial. This is because the location of an electrode can be read immediately from the voltages measured by it. The complication comes from the fact that the linear fields are perturbed by unknown functions. The location of the electrodes cannot be determined until the perturbations are pinned down.

To cast this problem in the vocabulary of parameter estimation, we make the simplifying assumption that the perturbations π_x , π_y and π_z are well represented by a finite linear combination of known harmonic functions. That is, we assume

$$\begin{aligned} \pi_x(x, y, z) &= \sum_{l=1}^{m_u} a_l \phi_l(x, y, z), \\ \pi_y(x, y, z) &= \sum_{l=1}^{m_v} b_l \psi_l(x, y, z), \\ \pi_z(x, y, z) &= \sum_{l=1}^{m_w} c_l \chi_l(x, y, z), \end{aligned} \quad (4)$$

where ϕ_l , ψ_l , and χ_l are harmonic functions. The coefficients a_l , b_l , and c_l are unknown, and these, together with the probe positions and orientations at each measurements $(x_i, y_i, z_i, \alpha_i, \beta_i)$ give a total of $(m_u + m_v + m_w) + 5n$ unknown parameters. Each measurement at a probe location gives 6 equations. So to make the problem formally determined, we require

$$(m_u + m_v + m_w) + 5n \leq 6n.$$

This paper addresses the problem of solving (3) with model (4) in the least-squares sense. We begin in Section 2 by outlining a Gauss-Newton algorithm for solving the problem. The updating procedure in the iterative process is analyzed carefully in Section 3. In particular, we show that if we view the probe position as random, we are able to determine limiting behavior of the matrices involved in the update. The analysis in Section 4 reveals a special structure for the Gauss-Newton update subproblem, which allows us to solve it efficiently, and control the instability that may arise. Section 5 describes an algorithm for solving the Gauss-Newton update. In Section 6, we present numerical results for a simulation. A discussion section ends the paper.

2 The nonlinear least-squares

We propose a Gauss-Newton procedure [4] to solve the inverse problem (3)–(4). Let us set the notation first. We use \mathbf{S} to represent the unknowns, which are segmented into

$$\mathbf{s}^* := [a_1, \dots, a_{m_u}, b_1, \dots, b_{m_v}, c_1, \dots, c_{m_w}]^T, \quad (5)$$

representing the field perturbation coefficients, and

$$\mathbf{s}_i := [x_i, y_i, z_i, \alpha_i, \beta_i]^T, \quad (6)$$

representing the probe position and orientation at the i th measurement step. The full state vector \mathbf{S} is thus a column vector of length $m_u + m_v + m_w + 5n$ and can be written

$$\mathbf{S} = [\mathbf{s}^*; \mathbf{s}_1; \dots; \mathbf{s}_n].$$

The electrode positions at the i th measurement step can be found using (1),

$$\begin{aligned} (x_i^{(1)}, y_i^{(1)}, z_i^{(1)}) &= (x_i, y_i, z_i) + r(\cos \alpha_i \cos \beta_i, \cos \alpha_i \sin \beta_i, \sin \alpha_i) \\ (x_i^{(2)}, y_i^{(2)}, z_i^{(2)}) &= (x_i, y_i, z_i) - r(\cos \alpha_i \cos \beta_i, \cos \alpha_i \sin \beta_i, \sin \alpha_i). \end{aligned}$$

Note that knowing these six positions is equivalent to knowing the five dimensional auxiliary state \mathbf{s}_i . The forward map $f : \mathbb{R}^{m_u+m_v+m_w+5} \rightarrow \mathbb{R}^6$ which takes \mathbf{s}^* and \mathbf{s}_i to a single set of measurements can now be written in terms of the electrode positions:

$$f(\mathbf{s}^*, \mathbf{s}_i) = \begin{bmatrix} x_i^{(1)} + \pi_x(x_i^{(1)}, y_i^{(1)}, z_i^{(1)}) \\ y_i^{(1)} + \pi_y(x_i^{(1)}, y_i^{(1)}, z_i^{(1)}) \\ z_i^{(1)} + \pi_z(x_i^{(1)}, y_i^{(1)}, z_i^{(1)}) \\ x_i^{(2)} + \pi_x(x_i^{(2)}, y_i^{(2)}, z_i^{(2)}) \\ y_i^{(2)} + \pi_y(x_i^{(2)}, y_i^{(2)}, z_i^{(2)}) \\ z_i^{(2)} + \pi_z(x_i^{(2)}, y_i^{(2)}, z_i^{(2)}) \end{bmatrix}. \quad (7)$$

Denoting by \mathbf{o}_i the vector of all voltages observed at the i th probe placement,

$$\mathbf{o}_i = \begin{bmatrix} u_i^{(1)} \\ v_i^{(1)} \\ w_i^{(1)} \\ u_i^{(2)} \\ v_i^{(2)} \\ w_i^{(2)} \end{bmatrix},$$

the nonlinear least-squares problem becomes that of finding $\hat{\mathbf{S}}$, where

$$\hat{\mathbf{S}} := \arg \min_{\mathbf{S}} \sum_{i=1}^n \|f(\mathbf{s}^*, \mathbf{s}_i) - \mathbf{o}_i\|^2. \quad (8)$$

By concatenating the observations \mathbf{o}_i into a $6n$ dimensional observation vector

$$\mathbf{O} = \begin{bmatrix} \mathbf{o}_1 \\ \vdots \\ \mathbf{o}_n \end{bmatrix}$$

and writing

$$\mathbf{F}(\mathbf{S}) = \begin{bmatrix} f(\mathbf{s}^*, \mathbf{s}_1) \\ \vdots \\ f(\mathbf{s}^*, \mathbf{s}_n) \end{bmatrix}, \quad (9)$$

the nonlinear least-squares is given in short-hand as

$$\hat{\mathbf{S}} := \arg \min_{\mathbf{S}} \|\mathbf{F}(\mathbf{S}) - \mathbf{O}\|^2. \quad (10)$$

We will later specifically exploit the structure in (8) when we devise our algorithm. For now, we prefer the more general specification given in (10).

The Gauss-Newton iterations for (10) are as follows.

Start with an initial guess \mathbf{S}_0 ;

For $k = 0, 1, \dots$

- Calculate the Jacobian of \mathbf{F} at \mathbf{S}_k , $J(\mathbf{S}_k)$;
- Solve the linear least-squares

$$J(\mathbf{S}_k)^T J(\mathbf{S}_k) \delta \mathbf{S}_k = -J(\mathbf{S}_k)^T (\mathbf{F}(\mathbf{S}_k) - \mathbf{O}); \quad (11)$$

- Update $\mathbf{S}_{k+1} = \mathbf{S}_k + \delta \mathbf{S}_k$.

We implemented a version of this algorithm and will present a concrete numerical example in which we reconstruct probe information $(\mathbf{s}_1, \dots, \mathbf{s}_n)$, and field perturbation coefficients \mathbf{s}^* from synthetically generated data \mathbf{O} . First, however, we would like to understand how well we can solve (11). In the next section we exploit a block decomposition of J to recast (11) in terms of (small) submatrices, and the subsequent section uses these blocks to analyze the stability of the solution.

3 Gauss-Newton step

We will solve (11) for the update vector $\delta \mathbf{S}$ but we will exploit the structure of the nonlinear least-squares. Consider the forward map $\mathbf{F}(\mathbf{S})$ in (9). From (7), we see that for the i -th set of measurements, the function f depends on \mathbf{s}^* (an $m = m_u + m_v + m_w$ dimensional vector) and the probe information \mathbf{s}_i (a 5 dimensional vector). The Jacobian of $\mathbf{F}(\mathbf{S})$ will thus have a special block structure. To be more precise, write

$$\delta f(\mathbf{s}^*, \mathbf{s}_i) = U_i(\mathbf{s}^*, \mathbf{s}_i) \delta \mathbf{s}^* + V_i(\mathbf{s}^*, \mathbf{s}_i) \delta \mathbf{s}_i, \quad (12)$$

where U_i is a 6-by- m matrix and V_i is a 6-by-5 matrix. Putting these together in J , we have

$$J = \left[\begin{array}{c|ccccc} U_1 & V_1 & 0 & \cdots & \cdots & 0 \\ U_2 & 0 & V_2 & 0 & \cdots & 0 \\ \vdots & \vdots & & \ddots & & \vdots \\ \vdots & \vdots & & & \ddots & 0 \\ U_n & 0 & \cdots & \cdots & 0 & V_n \end{array} \right] = [J_1|J_2]. \quad (13)$$

Note that J has $6n$ rows and $m + 5n$ columns.

To further analyze the Gauss-Newton step, we will write

$$J^T J = \begin{bmatrix} A & B \\ B^T & C \end{bmatrix} \quad \text{and} \quad (J^T J)^{-1} = \begin{bmatrix} D & E \\ E^T & F \end{bmatrix},$$

where the dimensions of D , E , and F match those of A , B , and C , respectively, and the latter are given by

$$A = \sum_{i=1}^n U_i^T U_i \in \mathbb{R}^{m \times m}, \quad (14)$$

$$B = [U_1^T V_1 \quad U_2^T V_2 \quad \cdots \quad U_n^T V_n^T] \in \mathbb{R}^{m \times 5n}, \quad (15)$$

$$C = \begin{bmatrix} V_1^T V_1 & 0 & \cdots & \cdots & 0 \\ 0 & V_2^T V_2 & 0 & \cdots & \vdots \\ \vdots & & \ddots & & \vdots \\ \vdots & & & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & V_n^T V_n \end{bmatrix} \in \mathbb{R}^{6n \times 6n}. \quad (16)$$

The special structure in the Jacobian (13) allows us to calculate elements of $(J^T J)^{-1}$ in terms of U_i and V_i . The following sequence of lemmas shows that we can find the field update $\delta \mathbf{s}^*$ separately from the probe update $\delta \mathbf{s}_i$. This fact offers several advantages in implementation which we describe in the next section.

Lemma 1. *The matrix D can be computed by inverting*

$$D^{-1} = \sum_{i=1}^n U_i^T [1 - V_i(V_i^T V_i)^{-1} V_i^T] U_i. \quad (17)$$

Proof. Write D^{-1} as the Schur complement of C , i.e.

$$D^{-1} = A - BC^{-1}B^T.$$

Now substitute (14), (15), and (16) for the right hand side, and simplify. \square

Note that $V_i^T V_i$ is only a 5×5 matrix, so computing its inverse is relatively inexpensive. We next provide explicit formulas for the field perturbation update $\delta \mathbf{s}^*$.

Lemma 2. *The Gauss-Newton update for \mathbf{s}^* is given by*

$$\delta \mathbf{s}^* = D \sum_{i=1}^n U_i^T [1 - V_i(V_i^T V_i)^{-1} V_i^T] [\mathbf{o}_i - f(\mathbf{s}^*, \mathbf{s}_i)]. \quad (18)$$

Proof. Write $\delta \mathbf{s}^* = [D \ E] J^T (\mathbf{O} - \mathbf{F}(\mathbf{S}))$ and note that $E = -DBC^{-1}$. Now substitute (13), (15), and (16), and simplify. \square

Once the update for the field perturbation is known, it is a simple matter to calculate the update for the probe information:

Lemma 3. *The Gauss-Newton update for \mathbf{s}_i is given by*

$$\delta \mathbf{s}_i = (V_i^T V_i)^{-1} V_i^T [\mathbf{o}_i - f(\mathbf{s}^*, \mathbf{s}_i) - U_i \delta \mathbf{s}^*]. \quad (19)$$

Proof. Partition J as in (13) into $J = [J_1 | J_2]$, and partition the state vector as $\delta \mathbf{S} = [\delta \mathbf{s}^*; \delta \mathbf{s}']$, where $\mathbf{s}' = [\mathbf{s}_1; \dots; \mathbf{s}_n]$. Then the Gauss-Newton update is the minimizer of

$$\|J_1 \delta \mathbf{s}^* + J_2 \delta \mathbf{s}' - \mathbf{O} + \mathbf{F}(\mathbf{S})\|^2,$$

whereupon it follows then that

$$\delta \mathbf{s}' = (J_2^T J_2)^{-1} J_2^T (\mathbf{O} - \mathbf{F}(\mathbf{S}) - J_1 \delta \mathbf{s}^*),$$

the i -th 5-dimensional block of which is given by (19). \square

The results in the previous Lemmas, namely (17), (18), and (19), will form the basis of our algorithm. We next prove some properties of the Jacobian in the limit as the number of measurements becomes large. We put the theory in a probabilistic context.

4 Analysis of the Gauss-Newton update

Equation (12) shows that the submatrices U_i and V_i of the Jacobian depend on the probe locations \mathbf{s}_i . As the probe is moved randomly from location to location, these random position parameters induce a sampling of some distribution of the U_i and the V_i . It is assumed that the number of samples is large, so that we can say something probabilistic about the behavior of the Jacobian, and also about the Gauss-Newton update.

Let Ω denote a probability space equipped with probability measure γ , and suppose that

$$W : \Omega \rightarrow \mathbb{R}^{m \times m}, \quad \omega \rightarrow \begin{bmatrix} w_{11} & \cdots & w_{1m} \\ \vdots & & \vdots \\ w_{m1} & \cdots & w_{mm} \end{bmatrix},$$

is a measurable mapping, where the matrix W is formed by

$$W_\omega = U^T(\mathbf{s}^*, \mathbf{s}_\omega) \left[I - V(\mathbf{s}^*, \mathbf{s}_\omega) (V(\mathbf{s}^*, \mathbf{s}_\omega)^T V(\mathbf{s}^*, \mathbf{s}_\omega))^{-1} V(\mathbf{s}^*, \mathbf{s}_\omega)^T \right] U(\mathbf{s}^*, \mathbf{s}_\omega).$$

One can immediately see that D^{-1} in (17) is the sum of n realizations of the random matrix W_ω .

In order to analyze the random matrix D^{-1} , let us formally define the quantities

$$\mu_{ij} = \int_{\Omega} w_{ij} d\gamma(\omega), \quad i, j = 1, \dots, m$$

and

$$\sigma_{ij} = \int_{\Omega} (w_{ij} - \mu_{i,j})^2 d\gamma(\omega), \quad i, j = 1, \dots, m$$

as well as the associated matrices

$$M := \begin{pmatrix} \mu_{11} & \cdots & \mu_{1m} \\ \vdots & \vdots & \vdots \\ \mu_{m1} & \cdots & \mu_{mm} \end{pmatrix} \quad (20)$$

$$\Sigma := \begin{pmatrix} \sigma_{11} & \cdots & \sigma_{1m} \\ \vdots & \vdots & \vdots \\ \sigma_{m1} & \cdots & \sigma_{mm} \end{pmatrix}. \quad (21)$$

Letting \hat{d}_{ij} denote the (i, j) -th element of D^{-1} , we have

$$\frac{\hat{d}_{ij}}{n} = \frac{1}{n} \sum_{l=1}^n w_{ij}(\omega_l),$$

whereupon taking the expectation of both sides yields

$$\mathbb{E} \left(\frac{\hat{d}_{ij}}{n} \right) = \frac{1}{n} \sum_{l=1}^n \mathbb{E}(w_{ij}) = \mu_{ij}. \quad (22)$$

Similarly, the variance can be calculated as

$$\text{Var} \left(\frac{\hat{d}_{ij}}{n} \right) = \mathbb{E} \left(\left[\frac{\hat{d}_{ij}}{n} - \mu_{ij} \right]^2 \right) = \frac{1}{n} \sigma_{ij}. \quad (23)$$

Taken together, these relations are equivalent to the matrix equations

$$\begin{aligned} \mathbb{E}(D^{-1}/n) &= M \\ \text{Var}(D^{-1}/n) &= \Sigma/n, \end{aligned}$$

where the expectation and variance operators act element by element.

Lemma 4. For $\epsilon > 0$,

$$\Pr \left(\left\| \frac{D^{-1}}{n} - M \right\|_{\infty} \geq \epsilon \right) \leq \frac{\|\Sigma\|_F^2}{n\epsilon^2}.$$

Proof. The proof is an application of Chebyshev's Inequality [5]. Specifically, from (22), (23), and the one variable Chebyshev Inequality,

$$\Pr \left(\left| \frac{\hat{d}_{ij}}{n} - \mu_{ij} \right| \geq \epsilon \right) \leq \frac{\sigma_{ij}^2}{n\epsilon^2}.$$

The L_{∞} norm of a matrix is given by the largest absolute sum along the rows, and by the subadditivity of positive measures,

$$\Pr \left(\max_i \sum_{j=1}^n \left| \hat{d}_{ij}/n - \mu_{ij} \right| \geq \epsilon \right) \leq \sum_{i=1}^n \sum_{j=1}^n \Pr \left(\left| \hat{d}_{ij}/n - \mu_{ij} \right| \geq \epsilon \right),$$

Combining these two inequalities yields the desired result. \square

Theorem 1. *If M is invertible and Σ is finite, then given any $\epsilon, \delta > 0$, there exists an $N = N(\epsilon, \delta, M, \Sigma)$ such that $\forall n > N$,*

$$\|D - M^{-1}/n\| < \frac{\epsilon}{n} \quad (24)$$

with probability $P > 1 - \delta$.

Proof. Write

$$D^{-1} = n \left(\frac{D^{-1}}{n} \right) = n(M + \mathcal{E})$$

and take the formal inverse of both sides to yield

$$D = \frac{1}{n}(M + \mathcal{E})^{-1}. \quad (25)$$

By Lemma 4, \mathcal{E} can be made arbitrarily small with arbitrarily high probability by taking n large enough. Since M is invertible and matrix inversion is locally norm continuous [2], $(M + \mathcal{E})^{-1} \rightarrow M^{-1}$ as \mathcal{E} vanishes. Putting these fact together we see that there exists an N such that for $n > N$,

$$\|(M + \mathcal{E})^{-1} - M^{-1}\| < \epsilon, \quad (26)$$

holds with probability $P > 1 - \delta$. The estimate of the theorem now follows by substituting (25) into the left hand side of (24) and invoking (26). \square

Formally, the matrix D represents the covariance of the perturbation parameters. Theorem 1 may thus be interpreted as saying that in the limit as the number of measurements becomes large, estimation error induced by measurement noise goes to zero. Alternatively, since our solution hinges on inverting the matrix D^{-1} , the result may also be interpreted as meaning that if M is well-conditioned, the inverse problem for determining the perturbation coefficient is well-conditioned as well. Finally, it is interesting to note that although the perturbation covariance vanishes, the position covariance does not. This is the subject of the following Lemma, which we include for completeness:

Lemma 5. *If the covariance of $\delta\mathbf{s}^*$ vanishes, the covariance of $\delta\mathbf{s}_i$ converges to $\sigma^2(V_i^T V_i)^{-1}$ where σ^2 is the variance of the noise.*

Proof. The result follows from (19) and the fact that the second order statistics of $\mathbf{o}^i - f(\mathbf{s}^*, \mathbf{s}_i) - U_i \delta\mathbf{s}^*$ converge to those of $(\mathbf{o}^i - f(\mathbf{s}^*, \mathbf{s}_i))$ if the covariance of $\delta\mathbf{s}^*$ vanishes. \square

5 Computational Issues

The purpose of this section is to highlight several important algorithmic and computational aspects of this registration problem. We begin by developing explicit expressions for U_i and V_i .

Recall that the matrix U_i is the partial of the forward map f with respect to the field perturbation coefficients \mathbf{s}^* . Since the forward map (7) depends linearly on the coefficients

(5), the $\mathbb{R}^{6 \times m}$ matrix U_i is easy to calculate, and is given by

$$U_i = \begin{bmatrix} [\phi_1, \phi_2, \dots, \phi_{m_u}]_{\mathbf{p}_i^1} & 0 & 0 \\ 0 & [\psi_1, \psi_2, \dots, \psi_{m_v}]_{\mathbf{p}_i^1} & 0 \\ 0 & 0 & [\chi_1, \chi_2, \dots, \chi_{m_w}]_{\mathbf{p}_i^1} \\ [\phi_1, \phi_2, \dots, \phi_{m_u}]_{\mathbf{p}_i^2} & 0 & 0 \\ 0 & [\psi_1, \psi_2, \dots, \psi_{m_v}]_{\mathbf{p}_i^2} & 0 \\ 0 & 0 & [\chi_1, \chi_2, \dots, \chi_{m_w}]_{\mathbf{p}_i^2} \end{bmatrix}. \quad (27)$$

The subscript indicates that the vector is evaluated at the electrode location \mathbf{p}_i^j .

The matrix V_i is a little more cumbersome. Recall that this matrix is the partial of the forward map with respect to the probe position vector \mathbf{s}_i , which was defined in (6). Note that since both the linear and the perturbation parts of the forward map depend on \mathbf{s}_i , the matrix V_i will have two components, one from each. In order to write down an expression for this matrix, let

$$\nabla \phi_l = [\phi_{lx}(x, y, z), \phi_{ly}(x, y, z), \phi_{lz}(x, y, z)]^T$$

denote the gradient of the perturbation function ϕ_l . Then

$$V_i = \begin{bmatrix} 1 & 0 & 0 & -r \sin \alpha_i \cos \beta_i & -r \cos \alpha_i \sin \beta_i \\ 0 & 1 & 0 & -r \sin \alpha_i \sin \beta_i & r \cos \alpha_i \cos \beta_i \\ 0 & 0 & 1 & r \cos \alpha_i & 0 \\ 1 & 0 & 0 & r \sin \alpha_i \cos \beta_i & r \cos \alpha_i \sin \beta_i \\ 0 & 1 & 0 & r \sin \alpha_i \sin \beta_i & -r \cos \alpha_i \cos \beta_i \\ 0 & 0 & 1 & -r \cos \alpha_i & 0 \end{bmatrix} + \mathcal{Z}_i, \quad (28)$$

where the term on the left is the contribution from the linear part and the matrix \mathcal{Z}_i is the contribution from the non-linear part. To calculate \mathcal{Z}_i , we define

$$\mathbf{v}_1 = \frac{\partial}{\partial \alpha} [r \cos \alpha \cos \beta, r \cos \alpha \sin \beta, r \cos \alpha]^T, \quad \mathbf{v}_2 = \frac{\partial}{\partial \beta} [r \cos \alpha \cos \beta, r \cos \alpha \sin \beta, r \cos \alpha]^T.$$

Then \mathcal{Z}_i becomes

$$\mathcal{Z}_i = \begin{bmatrix} \sum_{l=1}^{m_u} a_l \cdot \left(\nabla \phi_l^T \mathbf{e}_1 \Big|_{\mathbf{p}_i^1} & \nabla \phi_l^T \mathbf{e}_2 \Big|_{\mathbf{p}_i^1} & \nabla \phi_l^T \mathbf{e}_3 \Big|_{\mathbf{p}_i^1} & \nabla \phi_l^T \mathbf{v}_1 \Big|_{\mathbf{p}_i^1} & \nabla \phi_l^T \mathbf{v}_2 \Big|_{\mathbf{p}_i^1} \right) \\ \sum_{l=1}^{m_v} b_l \cdot \left(\nabla \psi_l^T \mathbf{e}_1 \Big|_{\mathbf{p}_i^1} & \nabla \psi_l^T \mathbf{e}_2 \Big|_{\mathbf{p}_i^1} & \nabla \psi_l^T \mathbf{e}_3 \Big|_{\mathbf{p}_i^1} & \nabla \psi_l^T \mathbf{v}_1 \Big|_{\mathbf{p}_i^1} & \nabla \psi_l^T \mathbf{v}_2 \Big|_{\mathbf{p}_i^1} \right) \\ \sum_{l=1}^{m_w} c_l \cdot \left(\nabla \chi_l^T \mathbf{e}_1 \Big|_{\mathbf{p}_i^1} & \nabla \chi_l^T \mathbf{e}_2 \Big|_{\mathbf{p}_i^1} & \nabla \chi_l^T \mathbf{e}_3 \Big|_{\mathbf{p}_i^1} & \nabla \chi_l^T \mathbf{v}_1 \Big|_{\mathbf{p}_i^1} & \nabla \chi_l^T \mathbf{v}_2 \Big|_{\mathbf{p}_i^1} \right) \\ \sum_{l=1}^{m_u} a_l \cdot \left(\nabla \phi_l^T \mathbf{e}_1 \Big|_{\mathbf{p}_i^2} & \nabla \phi_l^T \mathbf{e}_2 \Big|_{\mathbf{p}_i^2} & \nabla \phi_l^T \mathbf{e}_3 \Big|_{\mathbf{p}_i^2} & \nabla \phi_l^T \mathbf{v}_1 \Big|_{\mathbf{p}_i^2} & \nabla \phi_l^T \mathbf{v}_2 \Big|_{\mathbf{p}_i^2} \right) \\ \sum_{l=1}^{m_v} b_l \cdot \left(\nabla \psi_l^T \mathbf{e}_1 \Big|_{\mathbf{p}_i^2} & \nabla \psi_l^T \mathbf{e}_2 \Big|_{\mathbf{p}_i^2} & \nabla \psi_l^T \mathbf{e}_3 \Big|_{\mathbf{p}_i^2} & \nabla \psi_l^T \mathbf{v}_1 \Big|_{\mathbf{p}_i^2} & \nabla \psi_l^T \mathbf{v}_2 \Big|_{\mathbf{p}_i^2} \right) \\ \sum_{l=1}^{m_w} c_l \cdot \left(\nabla \chi_l^T \mathbf{e}_1 \Big|_{\mathbf{p}_i^2} & \nabla \chi_l^T \mathbf{e}_2 \Big|_{\mathbf{p}_i^2} & \nabla \chi_l^T \mathbf{e}_3 \Big|_{\mathbf{p}_i^2} & \nabla \chi_l^T \mathbf{v}_1 \Big|_{\mathbf{p}_i^2} & \nabla \chi_l^T \mathbf{v}_2 \Big|_{\mathbf{p}_i^2} \right) \end{bmatrix},$$

where \mathbf{e}_i is the i th unit vector in \mathbb{R}^3 .

Although somewhat unsightly, these expressions for U_i and V_i are useful because they can be used to provide a closed form approximation for the matrix M of Theorem 1. To illustrate this idea in a simple setting, let us suppose that all three sets of basis functions are identical and consist of m function ϕ_i , $i = 1, \dots, m$. If we ignore \mathcal{Z}_i (whose elements are

small if the expansion coefficients are small), then a messy but straightforward calculation shows that

$$U_i^T (1 - V_i (V_i^T V_i)^{-1} V_i) U_i = \left(\begin{array}{c|c|c} \mathcal{M}_i^{11} & \mathcal{M}_i^{12} & \mathcal{M}_i^{13} \\ \hline \mathcal{M}_i^{12} & \mathcal{M}_i^{22} & \mathcal{M}_i^{23} \\ \hline \mathcal{M}_i^{13} & \mathcal{M}_i^{12} & \mathcal{M}_i^{33} \end{array} \right),$$

where $\mathcal{M}_i^{jk} \in \mathbb{R}^{m \times m}$ is the matrix whose (l, q) -th element is

$$\mathcal{M}_i^{jk}(l, q) = [\phi_l(p_i^1) - \phi_l(p_i^2)] \cdot [\phi_q(p_i^1) - \phi_q(p_i^2)] \cdot w_i^{jk}, \quad j, k := 1, 2, 3 \quad (29)$$

and where the w_i^{jk} are scalar multiples given by

$$\begin{aligned} w_i^{11} &= \cos^2 \alpha_i \cos^2 \beta_i & w_i^{12} &= \cos \beta_i \sin \beta_i \cos^2 \alpha_i \\ w_i^{22} &= \cos^2 \alpha_i \sin^2 \beta_i & w_i^{23} &= \sin \alpha_i \cos \alpha_i \sin \beta_i \\ w_i^{33} &= \sin^2 \alpha_i & w_i^{13} &= \sin \alpha_i \cos \alpha_i \cos \beta_i \end{aligned} .$$

Note that the right hand side in (29) can be viewed as a weighted product of finite differences of the expansion functions. Let $\Delta_\gamma \phi_l$ denote the directional derivative of ϕ_l in the direction γ , and define

$$\mathcal{M}^{jk}(l, q) := \int_{\Omega} \Delta_\gamma \phi_l(\mathbf{p}) \cdot \Delta_\gamma \phi_q(\mathbf{p}) w_{jk}(\gamma) d\omega, \quad j, k = 1, 2, 3, \quad l, q = 1, \dots, m, \quad (30)$$

where γ represents the unique direction defined by orientation angles α and β , and the integral is over parameter space (α, β, x, y, z) . Comparing (29) to (30), we see that if n is large and the electrode separation distance $2r$ is small, then

$$\frac{\sum_{i=1}^n \mathcal{M}_i^{jk}}{n} \approx \mathcal{M}^{jk},$$

where the approximation should be understood to include a constant of proportionality. The M of Theorem 1 can thus be approximated (again up to a constant) as

$$M \approx \left(\begin{array}{c|c|c} \mathcal{M}^{11} & \mathcal{M}^{12} & \mathcal{M}^{13} \\ \hline \mathcal{M}^{12} & \mathcal{M}^{22} & \mathcal{M}^{23} \\ \hline \mathcal{M}^{13} & \mathcal{M}^{23} & \mathcal{M}^{33} \end{array} \right), \quad (31)$$

which can be explicitly calculated once the basis and the sampling distribution are known. Any basis set that induces a singular or ill-conditioned M should be discarded out of hand.

We conclude this section with a summary of our algorithm. We start with the initial vector $\mathbf{s}^* = 0$, and use the voltage measurements at each probe location, \mathbf{o}_i , to obtain the initial electrode locations \mathbf{p}_i^j , $j = 1, 2$, by assuming linear fields x , y , and z . From the electrode location, we invert the relation (1) to obtain the initial vectors \mathbf{s}_i . The \mathbf{s}^* and the \mathbf{s}_i should all be stored, along with the observation data.

Each iteration now consists of the following:

1. For each $i = 1, \dots, n$, calculate Jacobian matrices U_i and V_i using (27) and (28).
2. Update D^{-1} using (17), and $\delta \mathbf{s}^*$ using (18).
3. For $i = 1, \dots, n$ obtain probe information update vector $\delta \mathbf{s}_i$ using (19).

4. Update $\mathbf{s}^* \leftarrow \mathbf{s}^* + \delta\mathbf{s}^*$; and $\mathbf{s}_i \leftarrow \mathbf{s}_i + \delta\mathbf{s}_i$, for $i = 1, \dots, n$.

Iteration is ceased when some stopping criterion has been reached.

Matrix inversions are involved in steps 2 and 3. Since the inverse problem is ill-conditioned, the instability is going to manifest itself through the conditioning of the matrices D^{-1} and $V_i^T V_i$. One can control the instability by using truncated SVD [1, 3]. We found that, in our numerical experimentation with the example problem described in the next section, $V_i^T V_i$ is generally well-conditioned. In fact, one can use (28) to show that $V_i^T V_i$ becomes ill-conditioned only when the elevation angle α_i is near 90° , in which case there is minimal observability in the azimuth direction and the matrix becomes unstable. The matrix D^{-1} , while having small singular values, is invertible. Therefore, for the computations discussed in the next section, the true inverses are used.

In principle, the SVD cutoff for forming the truncated SVD inverse of $V_i^T V_i$ can be chosen independently of that for D^{-1} . This is probably a good idea in general, and perhaps even necessary. Note that a non-zero cutoff in either case leads to an asymptotic matrix M that is slightly different from the one defined in (31).

6 Numerical example

In this section we describe numerical results obtained using the Gauss-Newton algorithm on a synthetically generated data set. We describe our implementation of the method and details of the numerical experiments.

Synthetic data

As in the opening paragraph, the region B is taken to be a square box of length 2 centered at the origin. The basis functions are taken to be polynomials of the form

$$x^i y^j z^k : i + j + k := d = 2,$$

and we assume that each of the three fields can be expressed in this basis. Thus the perturbation fields in (4) are

$$\begin{aligned} \pi_x(x, y, z) &= a_1 x^2 + a_2 xy + a_3 xz + a_4 y^2 + a_5 yz + a_6 z^2, \\ \pi_y(x, y, z) &= b_1 x^2 + b_2 xy + b_3 xz + b_4 y^2 + b_5 yz + b_6 z^2, \\ \pi_z(x, y, z) &= c_1 x^2 + c_2 xy + c_3 xz + c_4 y^2 + c_5 yz + c_6 z^2. \end{aligned} \tag{32}$$

Since the total number of 3-D polynomials of degree d is $(d+1)(d+2)/2$, we need to solve for a total of $6 \times 3 = 18$ coefficients. Note that our basis is not harmonic, and in consideration of the fact that conditioning of the problem is dependent upon the choice of basis, it is most likely to be also suboptimal. We have chosen to work with low order polynomials for ease of computation, but actual applications would doubtless entail a different choice of basis.

In our simulation, we fix the voltage differential as a linear sweep from -1 to 1, and merely adjust the perturbation coefficients. These coefficients a_l , b_l , and c_l in (32) are generated as zero-mean Gaussian random variables with a fixed variance. We consider 2 levels of nonlinearity in the field. These correspond to variances of 0.10 and 0.20. A typical realization with variance 0.20 produces a few coefficients as large as 0.5 in absolute value.




Figure 2: Slices of the field perturbations $u(x, y, z) - x$ at $z = -1, -0.5, 0, 0.5, 1$, over the region $|x| \leq 1$ and $|y| \leq 1$. Note that the perturbations are significant considering that the unperturbed field has maximum amplitude of 1.

The nonlinear field corresponding to a realization is shown in Figure 2, where we show slices of the field $[u(x, y, z) - x]$ for $z = -1, -0.5, 0, 0.5, 1$. Reading the values of the perturbation using color bar on the bottom of each figure gives an indication to the size of the nonlinearity in the field. It is clear that the deviation from the linear profile is quite significant. The size of the coefficient perturbations are meant only to give a qualitative measure of nonlinearity. The variance mentioned is related to the size of the second derivatives of the fields. A linear field would have zero second derivatives.

Once the fields $u(x, y, z)$, $v(x, y, z)$, and $w(x, y, z)$ are generated, we choose n , the number of measurements, and generate sensor locations \mathbf{p}_i , for $i = 1, \dots, n$. The locations are uniformly distributed in B . We assume that the two electrodes are separated by a distance of 0.1, and obtain their locations using (1) for each sensor position using random orientations (α_i, β_i) , also uniformly distributed. Once $(x_i^{(1)}, y_i^{(1)}, z_i^{(1)})$ and $(x_i^{(2)}, y_i^{(2)}, z_i^{(2)})$ are found, we sample the fields at these points to generate the data $\mathbf{o}_i = [u_i^{(1)}, v_i^{(1)}, w_i^{(1)}, u_i^{(2)}, v_i^{(2)}, w_i^{(2)}]^T$. Noise, uniformly distributed, is added to the data.

Solving the inverse problem

To solve the inverse problem, we need an initial guess. Assuming no a priori knowledge of the field perturbations, we set the initial vector of coefficients $\mathbf{s}_0^* = 0$. The initial position information for measurement i , the auxiliary state \mathbf{s}_0^i , is estimated from the data, with

$$\mathbf{p}_0^i = \left[\frac{u_i^{(1)} + u_i^{(2)}}{2}, \frac{v_i^{(1)} + v_i^{(2)}}{2}, \frac{w_i^{(1)} + w_i^{(2)}}{2} \right]^T.$$

and the angles

$$\alpha_0^i = \arctan \frac{w_i^{(1)} - w_i^{(2)}}{\sqrt{(u_i^{(1)} - u_i^{(2)})^2 + (v_i^{(1)} - v_i^{(2)})^2}}, \quad \beta_0^i = \arctan \frac{v_i^{(1)} - v_i^{(2)}}{u_i^{(1)} - u_i^{(2)}}.$$

With the initial guess in hand, we can proceed with the Gauss-Newton iterations as outlined at the end of Section 5.

In our implementation, matrix inversions will be done by calculating the inverse (not truncated SVD inverse). The matrices involved can be very ill-conditioned but in our computations, we never encountered non-invertibility. Since the solution method is iterative, we




Figure 3: The SVD of the full Jacobian at the initial guess for a given data with 120 measurements. The vertical scale is log base 10. Note the rapid fall of the singular values, indicating ill-conditioning.

need to set a stopping rule. In principle, we would stop the iterations when the norm of the gradient approaches zero. In practice, we proceed by calculating the relative residual (relative to the size of the data) at each iteration. When the difference between two successive residuals fall below a certain threshold, the algorithm is stopped. The threshold was set to 10^{-5} after many experimentations. (Note that although a small gradient guarantees small successive residuals, the converse is not true. We used the unprincipled successive residuals criterion merely for speed and ease of coding.) One final detail is that we implemented a step size $\tau \leq 1$ in Step 4 of the algorithm in Section 5. To be precise, we choose τ and update according to

$$\mathbf{s}^* \leftarrow \mathbf{s}^* + \tau \delta \mathbf{s}^* \quad \text{and} \quad \mathbf{s}_i \leftarrow \mathbf{s}_i + \tau \delta \mathbf{s}_i.$$

This is needed because the cost function can have narrow valleys. For the computations reported below we chose $\tau = 0.1$. A line search algorithm [4] is a more effective way to handle this difficulty.

Ill-conditioning

To get a sense of the ill-conditioning of the problem, we generated data and calculated the full Jacobian J (13) at the initial guess. This is for the case where the number of measurements is 120, so the numbers of unknowns and equations are 618 and 720, respectively. We next calculated the singular value decomposition of J and plot the base 10 logarithm of the singular values in Figure 3. We can see that the singular values drop very rapidly. The largest singular value is 10.95, whereas the smallest is 0.0029. The matrix is not singular but the presence of very small singular values indicate that there are directions in which the objective function changes very little if we move from the initial guess in one of these directions.

To give an interpretation to Theorem 1, we conducted a numerical experiment that serves to illustrate its implications. The theorem is a statement about the limiting behavior of the

matrix D in (24). It states that as the number of measurements become large, the variance in the reconstruction of $\delta\mathbf{s}^*$ goes to zero. To demonstrate this fact numerically, we start with field perturbation coefficients, generate the synthetic observation for the number of observations, add noise to the data, and run *one step* of the Gauss-Newton iteration. One step of the Gauss-Newton algorithm amounts to solving the linearized inverse problem, which the theorem addresses. The inversion gives an estimated $\delta\mathbf{s}^*$ from which we can calculate the error. Repeating this for different realizations of measurements and noise, we can collect the average variance of the error in the reconstruction. We choose small coefficients in the field perturbation so that the problem is nearly linear, which allows us to interpret the associated matrix D accordingly.

The experiment goes as follows. First we generate the field perturbation coefficients, which we fix for the rest of the computations. Next, for a fixed number of observations and a fixed level of noise, we choose 5 realizations of the electrode locations and noise. We invert the data, calculate the empirical variance of each element in $\delta\mathbf{s}^*$, and then collect the mean of these variances. Let us call this v_{calc} . During the run, we also calculate the mean of D for the 5 runs. Let us call this mean matrix \bar{D} . The ‘theoretical variance’ in the reconstruction error is

$$v_{\text{theo}} = \text{mean} [\text{diag} (D)] \cdot \text{noise variance}.$$

Theory predicts that v_{calc} should be close to v_{theo} and that as the number of measurements become large, both these numbers become small.




Figure 4: Comparison between theoretical and calculated variances. In this plot, we display v_{calc} (in squares) and v_{theo} (in dots). Theory predicts that they should be close to each other and they should approach zero as the number of measurements become large. The graphs, from bottom to top, are for data with Gaussian noise of 0.1%, 0.35% and 1.4%, respectively.

Our experiment is summarized in Figure 4 where we plotted v_{calc} and v_{theo} for different

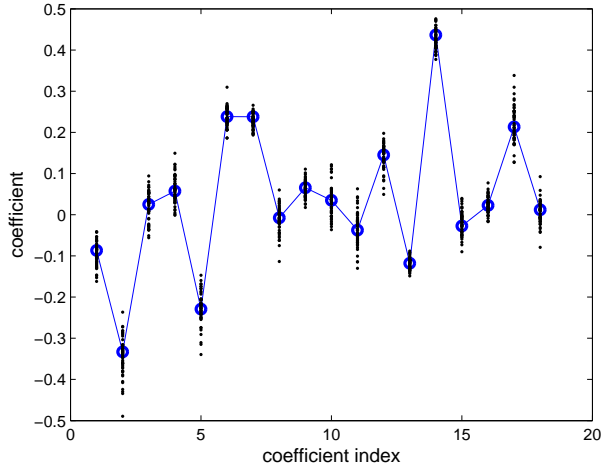


Figure 5: Scatter plot of the reconstructions of the coefficients for the 50 realizations of the noise. The true coefficients are denoted in circles while the reconstructions are shown in dots.

noise levels and different numbers of measurements. The noise levels are at 0.1%, 0.35%, and 1.4% of the signal, whereas the number of measurements range through 100, 200, 400, 800, 1600, 3200. Clearly the agreement between v_{calc} and v_{theo} is good, and we also see that both quantities approach zero for large number of measurements. If the problem were linear, this fact can be interpreted as stabilization of the estimation for large number of measurements. As the problem is nonlinear, its implication is not as clear cut. However, we observed stabilization in the nonlinear problem nonetheless. This is described below.

Noisy reconstructions

Coefficients for the polynomial perturbation field in (32) are generated. The $u(x, y, z)$ field corresponding to these coefficients, when the variance of the coefficients is 0.2, is displayed in Figure 2. We also generated another test coefficients with variance 0.1. Data for both cases are generated for 80, 120, and 160 measurements. Each of these measurement numbers leads to an overdetermined problem. Uniformly distributed noise is added to the data. The noise is measured in the ℓ_2 norm relative to the data size. The level of noise we chose is what one would expect in some medical applications [6].

To assess the accuracy of the reconstruction of the perturbation field, we use the relative ℓ_2 norm, that is if $\mathbf{s}^* = [a_1, \dots, a_6, b_1, \dots, b_6, c_1, \dots, c_6]^T$, the coefficient error is

$$\mathbf{s}_{\text{err}}^* = \frac{\|\mathbf{s}^* - \mathbf{s}_{\text{true}}^*\|}{\|\mathbf{s}_{\text{true}}^*\|}.$$

To measure the accuracy of the reconstruction of the voltage sensor locations, we calculate, for each placement of the probe, the Euclidean distance between the true sensor locations and the estimated ones

$$e_i^j = |\mathbf{P}_i^j - \mathbf{P}_{i,\text{true}}^j|, \quad j = 1, 2.$$

The sensor locations are calculated from the probe position and orientation through (1). Our measure of the accuracy of sensor locations is absolute and is the mean of the sum of




Figure 6: Histogram of the sensor position reconstruction error for the 50 realizations of the noise.

the difference of the Euclidean distances

$$\mathbf{p}_{\text{err}} = \frac{1}{n} \sum_i (e_i^1 + e_i^2),$$

where n is the number of probe positions. Below we report on the results of the reconstructions. For each choice of number of measurements, we run the inversion on 50 realizations of the noise. The average number of iterations needed to converge is just under 40.

Table 1 summarizes the results when the coefficients have variance of 0.1. The recovery of the coefficients is rather poor, however, note that the error in sensor location is quite small, considering that the box in which we probe is 2 units wide. It can be seen that increasing the number of measurements increases the accuracy of the reconstruction, and has a stabilizing effects. This latter is evidenced by the reduction in the variance of the error in both the coefficient and the position reconstructions.

num. meas.	$\mathbf{s}_{\text{err}}^*$		\mathbf{p}_{err}		$\ \mathbf{F}(\mathbf{S}) - \mathbf{O}\ /\ \mathbf{O}\ $	
	mean	var	mean	var	mean	var
80	0.5194	0.0139	0.1154	0.0006	0.0036	1×10^{-7}
120	0.3832	0.0081	0.0838	0.0004	0.0038	4×10^{-8}
160	0.3296	0.0087	0.0743	0.0004	0.0038	4×10^{-8}

Table 1: Reconstruction when the perturbation to the linear field is of moderate size. The polynomial coefficients in (32) have variance of 0.1. Note the improvements in the reconstructions when the number of measurements is increased.

The second table, Table 2, summarizes the results of the reconstruction for the case where the coefficients have variance 0.2. The reconstructions of the coefficients are consistently better for this case than when the coefficients are smaller. However, we notice very

little improvement in sensor location calculations. Other trends, such as better reconstruction and better stability when the number of measurements is increased can be seen for this experiment as well.

num. meas.	$\mathbf{s}_{\text{err}}^*$		\mathbf{p}_{err}		$\ \mathbf{F}(\mathbf{S}) - \mathbf{O}\ /\ \mathbf{O}\ $	
	mean	var	mean	var	mean	var
80	0.2504	0.0039	0.1188	0.0012	0.0036	9×10^{-8}
120	0.1911	0.0022	0.0922	0.0004	0.0038	6×10^{-8}
160	0.1596	0.0013	0.0767	0.0002	0.0038	5×10^{-8}

Table 2: Reconstruction when the perturbation to the linear field is of large size. The polynomial coefficients in (32) have variance of 0.2. In comparing with Table 1, note that the reconstructions of the coefficients are markedly better here. Again, we see improvements in the reconstructions when the number of measurements is increased.

To visualize the reconstructions, we display in Figure 5 a scatter plot of the reconstructed coefficients overlaid against the true coefficients in the 50 different inversions. The figure shows that the coefficients are reasonably well estimated. Figure 6 shows a histogram of the means of the sensor position errors \mathbf{p}_{err} for the 50 different inversions, and reveals that the majority of the inversions found the sensor locations to within 0.10 units in the 2 unit cube being probed.

In order to get a better sense of the reconstruction of the sensor locations, we plot the xy -coordinates of the sensors for one of the reconstructions. Both the left and the right plot in Figure 7 use solid circles to depict the true sensor locations and open circles to depict the estimates, with line between the circles used to indicate that the sensors are from the same probe position (midpoint between the circles). The estimates on the left are the initial guesses, those on the right the final reconstructions. We can see that there is a huge improvement over the initial guess, and that the reconstruction is quite accurate.

7 Discussion

We have described a problem of position registration from measurements of a non-linear voltage field. Our solution modeled the perturbations as low order expansions in a known basis, and we devised a simple Gauss-Newton method to simultaneously determine the expansion coefficients and the sensor locations. The Gauss-Newton updates are analyzed probabilistically, and we analyze the asymptotic behavior of these updates in the limit as the number of measurements becomes large.

The method is implemented and experimented on synthetically generated data. We found that the reconstruction of acceptable accuracy is possible in the presence of a small but realistic amount of noise. The stabilizing properties of a large number of measurements, predicted in theory, is observed in the numerical example.

Although the motive and thrust of this paper have been position registration, the principle results have been formulated for the wider class of problems whose state expands with the number of measurements. The primary theoretical result concerns the convergence of the formal covariance under the accretion of probabilistically parameterized measurements.




Figure 7: Reconstruction of the sensor locations. Solid circles depict the xy -coordinates of the true sensor locations and open circles depict the xy -coordinates of our estimates. The plot on the left shows our initial guesses, obtained by assuming that the field is linear, while the plot on the right shows the final reconstructions. A line is drawn between each pair of sensors to indicate that they correspond to a single probe location.

Practical contributions include a computationally tractable implementation scheme and a criterion for basis selection.

Although preliminary results suggest that this scheme can work, at least in the case of low order polynomial perturbations, much remains to be investigated. Three questions that stand out are:

- For a given application, is there an optimal choice of basis?
- Is there an analytic condition under which iterates can be guaranteed to converge?
- Is there an optimal regularization strategy?

The answers to these questions may well depend on the region and distortions of interest. In special cases, the distortions may be large but relatively well known, in which case one might start with either a special basis or a non-zero initial guess of the base state.

Acknowledgment

The research of FS is supported in part by the National Science Foundation under award DMS-0504185. We are grateful to Dr. Eric Voth of St. Jude who brought the problem studied in this work to the authors' attention, and with whom we had very helpful discussions.

References

- [1] Demmel, James, *Applied Numerical Linear Algebra*, SIAM, Philadelphia, 1997.

- [2] Golub, Gene and Van Loan, Charles, *Matrix Computations*, John Hopkins Press, Baltimore, 1983.
- [3] Kaipio, Jari and Somersalo, Erkki, *Statistical and Computational Inverse Problems*, Springer, New York, 2004.
- [4] Nocedal and Wright, *Numerical Optimization*, Springer, New York, 1999.
- [5] Papoulis, Athanasios, *Probability, Random Variables, and Stochastic Processes*, McGraw Hill, New York, 1984.
- [6] Eric Voth, personal communication.