

Source localization in electromyography using the inverse potential problem

Kees van den Doel, Uri M. Ascher, and Dinesh K. Pai *

May 20, 2010

Abstract

We describe an efficient method for reconstructing the activity in human muscles from an array of voltage sensors on the skin surface. MRI is used to obtain morphometric data which is segmented into muscle tissue, fat, bone and skin, from which a finite element model for volume conduction is constructed. The inverse problem of finding the current sources in the muscles is solved using a careful regularization technique which adds a priori information, yielding physically reasonable solutions from among those that satisfy the basic potential problem. Several regularization functionals are considered and numerical experiments on a 2D test model are performed to determine which performs best. The resulting scheme leads to numerical difficulties when applied to large scale 3D problems. We clarify the nature of these difficulties and provide a method to overcome them, which is shown to perform well in the large scale problem setting.

1 Introduction

The aim of the methods considered here is the determination of the electrical activity of individual muscles in human limbs by using surface electromyography (sEMG) data from multiple recording sites. We call this computed myography (CMG) [7]. Potential applications include prosthetic control and muscle function assessment.

This paper has two major goals. The first is to determine an appropriate mathematical model to solve for the inverse problem, since it is well-known that the classical inverse potential problem does not provide a sufficient model without additional a priori information. Several approaches to adding such a priori information are possible, and one goal of the present paper is to determine which approach performs best. For this purpose we have tested various methods in the setting of a two dimensional (2D) model of a limb, leading to the conclusion that one particular modeling and regularization combination performs significantly better than any of the others.

The other goal of the present paper is to design an efficient computational method for solving the resulting discretized inverse problem. Since the practical problem of interest is

*Dept. of Computer Science, University of British Columbia, Canada. This work was supported in part by the Peter Wall Institute for Advanced Studies, Canada Research Chairs Program, NSERC, CFI, and BC KDF.

formulated in complicated three dimensional (3D) domains the resulting equations lead to large scale algebraic systems, and many powerful techniques applicable to small or medium size systems (such as the singular value decomposition, for example) cannot be used because of time and memory limitations.

Modern hardware and software enables accurate simulation of electrical signal propagation (the forward problem) in complex 3D geometries such as human limbs, see for example [23], but the inverse problem which is to determine the sources from surface measurements is much harder. Even when high density sEMG data is available [8], the inverse problem is known to have no unique solution without imposing further restrictions on the sources [17]. Thus, the task of reconstructing sources must be aided by additional information.

Discrete source models [26, 22, 21, 2] do not perform well, and since the number of sources in active muscles can be large a distributed underdetermined formulation of the inverse problem, also used in EEG [25], seems more promising.

In the CMG system proposed in [7] a finite element (FEM) discretization is constructed using MRI data and a distributed model of current activations in muscles is reconstructed using conductance based on the geometry, see Fig. 1.

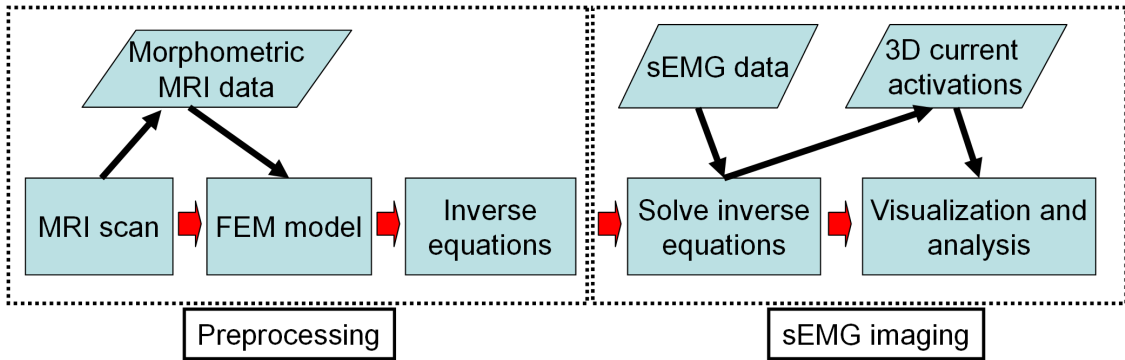


Figure 1: Diagram of the CMG system. The thick arrows indicate sequential processing steps, the black arrows indicate data produced or used, the rectangles represent processes, and the slanted rectangles represent data. The preprocessing is done once for a subject, and the resulting model can be used to monitor muscle function over a period of time as long as no significant change in morphology takes place.

The basic mathematical formulation of the CMG problem in [7] involves recovering the source term q of a generalized Poisson equation from boundary measurements of the electrical potential u . Thus, consider

$$-\nabla \cdot (\sigma \nabla) u = q, \quad \mathbf{x} \in \Omega, \quad (1a)$$

$$\frac{\partial u}{\partial \nu} \Big|_{\partial \Omega_1} = 0, \quad (1b)$$

$$\frac{\partial u}{\partial \nu} \Big|_{\partial \Omega_2} = -\Sigma u, \quad (1c)$$

where Ω is the domain representing the limb in 3D, $\frac{\partial}{\partial \nu}$ represents differentiation in the

outward normal direction, and the conductivity matrix $\sigma(\mathbf{x})$ is known (albeit only approximately). We seek to recover the source $q(\mathbf{x})$ from measurements of u on $\partial\Omega_1$. The rest of the boundary, $\partial\Omega_2$, represents the artificial boundary of our domain, e.g., the cut boundary of part of a limb. The Σ term is an effective resistance modeling the flow of currents in and out of our domain at its artificial boundaries.

However, it is well-known that this problem has no unique solution [17, 18, 20]. Thus, we need to regularize it in a way that introduces as much additional relevant information as possible, in the hope to determine a physically plausible solution in a stable fashion.

The sources q of the potential field u are polarization waves called intracellular action potentials (IAP) which travel outwards over the individual muscle fibers starting at the axon connection, roughly in the middle of the muscle fiber. Each motor unit (MU) is a collection of muscle fibers which is excited more or less synchronously. Typically there are 100-1000 MU's per muscle which may or may not be temporally dependent. The wave shape of the IAP is known, and is qualitatively a smeared tripole of size 2-40mm; standard models exist to describe it [24].

At any instant in time the tripoles are distributed more or less randomly over the active muscles, and the more active the muscle the denser they are. So at a small scale the source $q(\mathbf{x})$ consists of many localized tripole excitations, whereas at larger scales describing it by a spatially smooth density function becomes appropriate.

In order to formulate the inverse problem we have to decide on source representation and on a regularization functional. In Section 2 we develop and examine possible avenues of action for these two issues.

Our next task is to select which of the inversion models developed produces the best results for the problem at hand. This is examined in Section 3, using numerical experiments with synthetic data on a simplified geometry $\Omega \subset \mathbb{R}^2$ and utilizing relatively small meshes. The upshot is that the combination of (2b) with the regularization (3d) and $w \equiv 1$ (defined in Section 2 and referred to as L2D) yields the best reconstructions.

The resulting linear systems of equations are solved by an iterative method such as conjugate gradients (CG). This is essential for efficiency in the large scale case. Unfortunately, however, a straightforward implementation of a preconditioned CG (PCG) method runs into numerical difficulties, essentially because the method that produces better reconstructions involves a numerical differentiation matrix, which amplifies unstructured numerical errors. In Section 4 we demonstrate and explain these sources of difficulty, still using simple discretizations on simple domains in 2D, and show how they can be overcome. We end up recommending dynamic, or iterative regularization [16, 9, 5], even though a Tikhonov-type procedure is used in the construction of the mathematical inverse problem that gets approximately solved.

In Section 5 we then apply our method in a close to realistic experimental setting for the localization of muscles with a geometry obtained from segmenting MRI data from a human subject. Much better reconstructions than in [7] are obtained at a rather affordable cost. Section 6 concludes this article.

2 Model and regularization

Regarding the sought source representation, we could directly attempt to reconstruct q , or we could incorporate the information that the tripole is the basis excitation. In other words, we may seek a function $m(\mathbf{x})$ such that either

$$q \equiv m, \quad (2a)$$

or

$$q \equiv \mathcal{D}m, \quad (2b)$$

with \mathcal{D} standing for an operator converting to tripole basis functions. The latter yields an expression of the form $\mathcal{D} = \partial_{zz}^h$, where

$$\partial_{zz}^h f(x, y, z) = (f(x, y, z + h) + f(x, y, z - h) - 2f(x, y, z))/h^2 \quad (2c)$$

is a discrete second derivative operator in the z direction which we assume to be aligned along the direction of the muscle fibers. The length scale h is the expected size of the tripole, and any discretization of subsequent differential terms is mandated to be at least as fine as this. If f in (2c) is a delta function (monopole) then $\partial_{zz}^h f$ describes a tripole. The latter choice of variables introduces further a priori information that hopefully better resolves the nonuniqueness, i.e., it helps select a more physical solution. Note that the suitability of the choice (2c) does depend on which part of the human body the system gets applied to. In particular we assume that the sensors are not near the nerve endings or the tendons, as end effects (the generation and extinction of the action potentials) are not well described with this choice.

Further, we consider at this point introducing a priori information through a Tikhonov-type regularization, which penalizes some norm of m or ∇m . Utilizing $\|m\|_2$ is well-known to have a strong bias for sources near the sensors (as sources near the sensors can explain the data most “cheaply”) [15], which can be alleviated to some extent by a suitable weighting factor. Another candidate for a regularization functional is $\|m\|_1$, aiming for a sparse reconstruction. The justification for this is that the physical source is built out of a finite number of tripoles, and under certain conditions the 1-norm regularized solution corresponds to a sparse combination. This makes sense only for the choice (2b), not (2a). Note that this option leads to a nonlinear set of equations, so it should be considered only if the reconstruction results are significantly better. Finally, we may consider aiming for a smooth reconstruction by a penalty term proportional to the 2-norm of $|\nabla m|$, possibly weighted to compensate for bias against deep sources. The physiological basis for this well-known approach is that activity in nearby regions is correlated within muscles. Note that there is no reason to consider any variant of the popular total variation seminorm [1] here as we do not expect discontinuities in $m(\mathbf{x})$ inside a muscle. This regularization term is applied to the relevant regions of interest (ROI) which are the muscles. Outside the ROI (e.g., on regions with bone or fat) the requirement that $m(\mathbf{x})$ vanish is enforced by a large penalty on the size of m .

We write our penalty functional $R(m)$ as (an appropriate discretization of)

$$R(m) = R_{ROI}(m) + \frac{M}{2} \int_{\Omega} (1 - \kappa(\mathbf{x})) m(\mathbf{x})^2 d\mathbf{x}, \quad (3a)$$

where R_{ROI} is one of the following possibilities:

$$\frac{1}{2} \int_{\Omega} \kappa(\mathbf{x}) w(\mathbf{x}) m(\mathbf{x})^2 d\mathbf{x}, \quad (\text{N2}) \quad (3b)$$

$$\int_{\Omega} \kappa(\mathbf{x}) w(\mathbf{x}) |m(\mathbf{x})| d\mathbf{x}, \quad (\text{N1}) \quad (3c)$$

or

$$\frac{1}{2} \int_{\Omega} \kappa(\mathbf{x}) w(\mathbf{x}) |\nabla m(\mathbf{x})|^2 d\mathbf{x}. \quad (\text{L2}) \quad (3d)$$

Here M is a very large constant. The characteristic function κ is defined to satisfy $\kappa = 1$ on the ROI and $\kappa = 0$ outside. The (optional) weight function $w(\mathbf{x})$ is there to reduce bias. We shall refer to these regularizations as N2, N1 (for 2- and 1-norm), and L2 (for the Laplacian weighted 2-norm), to which we append D and/or W if nontrivial \mathcal{D} or w , respectively, are used.

The inverse problem can now be formulated as the optimization problem

$$\min_m \phi(m) = \phi_{LSQ} + \beta R, \quad (4a)$$

with the least squares error term

$$\phi_{LSQ} = \frac{1}{2} \|Qu(m) - b\|^2, \quad (4b)$$

where b is the measured data, u the potential generated by m , Q a projection operator which samples u at the electrode locations, and β the regularization parameter. After discretization, m , u , and q change from functions to n -vectors, where n is the (large) number of discretization variables, and linear operators become matrices. Because u depends linearly on m we can write

$$Qu(m) = Jm,$$

where the operator J becomes, upon discretization, a $d \times n$ matrix, with d the number of measurements.

We also entertain the following choice for the weight factor $w(\mathbf{x})$, advocated in [13, 10]. Note that the contribution of a single unit strength point source to ϕ_{LSQ} is proportional to $(J^T J)_{kk}$, where k is the discrete index corresponding to the location of the source. Since J is essentially the Green's function, this factor will be smaller for deep sources. Without weights w the penalty in the regularization term will be independent of depth, hence form a bias against deep sources. This can be compensated for by choosing the weights w as

$$w = \text{diag}(J^T J)$$

for N2 and L2, and

$$w = \sqrt{\text{diag} J^T J}$$

for N1. A statistical motivation for this choice of weight is given in [13, 10]. Note that this is a heuristic choice, based on single sources, and it is somewhat arbitrary, as there is obviously no perfect choice for w that will give an error free reconstruction (because the basic inverse potential problem is underdetermined).

3 Selecting the inversion method

In Section 2 we have defined ten inversion methods (three possibilities for R_{ROI} , each with or without a weight function and/or a nontrivial \mathcal{D} , except for the sparse method which requires \mathcal{D}). In order to determine the best one for our problem, we have constructed a 2D model resembling a cross section of an idealized upper arm, see Fig. 2. The numbered

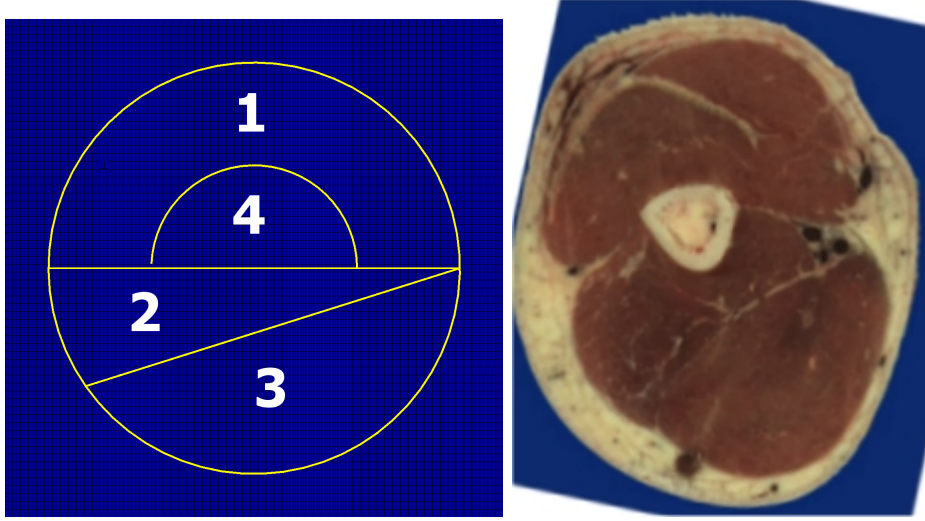


Figure 2: Our two dimensional upper arm test model is displayed on the left. The sensors are located on the boundary of the square. Artificial tripole sources oriented horizontally are placed in the muscle areas 1-4. Anatomically, area 2 corresponds to brachialis, 3 to biceps, and 1 and 4 to the triceps. We have divided the triceps into outside (1) and inside (4) regions, to test if activities in these regions can be distinguished. A cross section of a real upper arm is displayed on the right for comparison purposes.

regions represent four ROI, which we refer to as muscles, and the data is taken on the boundary of the square. The penalty term is applied outside the ROI, which includes the boundaries between muscles. The PDE (1a) is discretized in a standard way on an $N \times N$ square mesh (thus $n = N^2$). Synthetic data is generated by placing tripole sources randomly in the muscles. The spatial extent of the synthetic sources (the parameter h in (2c)) was varied between a single grid point $h = 1/64$ and 0.1, whereas the assumed tripole size h in the reconstruction was taken to be $h = 1/32$. No noticeable effect of the true source size was observed, presumably because the size of the tripole is not visible at the sensors due to the blurring effect of the volume conductor.

Data was computed by solving the forward problem with $N = 64$, to which we add 5% Gaussian noise. The inversion process is then performed using $N = 32$. This avoids “inverse crimes”. The regularization parameter β was chosen by the discrepancy principle [16] so as to make the misfit between predicted and measured data approximately 7.5%. The grid is small enough that direct methods can be used to solve arising linear algebraic systems and all matrices can be stored in memory.

Since the surface data, even under ideal circumstances, does not contain enough infor-




Figure 3: Reconstruction results for Experiment 1: sources only in Region 1. Fig. 3(c) is on a saturated scale in order to display the weaker sources. The source in the center region has a magnitude of 40.

mation to reconstruct the detailed source distributions, we cannot utilize an obvious error measure such as $\|m - m_{true}\|$ for the present comparative study, for m_{true} consists of discrete sources whereas m represents a “smeared” equivalent distribution.

Instead we consider the task of distinguishing overall activity in each of the four ROI, which we take to be the power density, i.e., the average of m^2 on a particular region. This would be useful information for example for prosthetic control or overall muscle strength assessment. Apart from these numerical measures we also display the resulting reconstructed model m , and use visual inspection to assess the relative quality of the various methods.

From the many computations performed we have selected to display here a set of three experiments designed to test if the interior and exterior of the “triceps”, i.e., Regions 1 and 4, can be distinguished. Artificial unit strength tripole sources were placed in Regions 1, 4, and in both. We refer to these as Experiments 1 to 3, respectively. These “true models” are depicted in Figs. 3–5 which also display reconstructions using ten regularization options.

In order to compare the activation reconstructions numerically the activations were normalized to sum to 1, as absolute values are not meaningful. The results were gathered into Table 1.

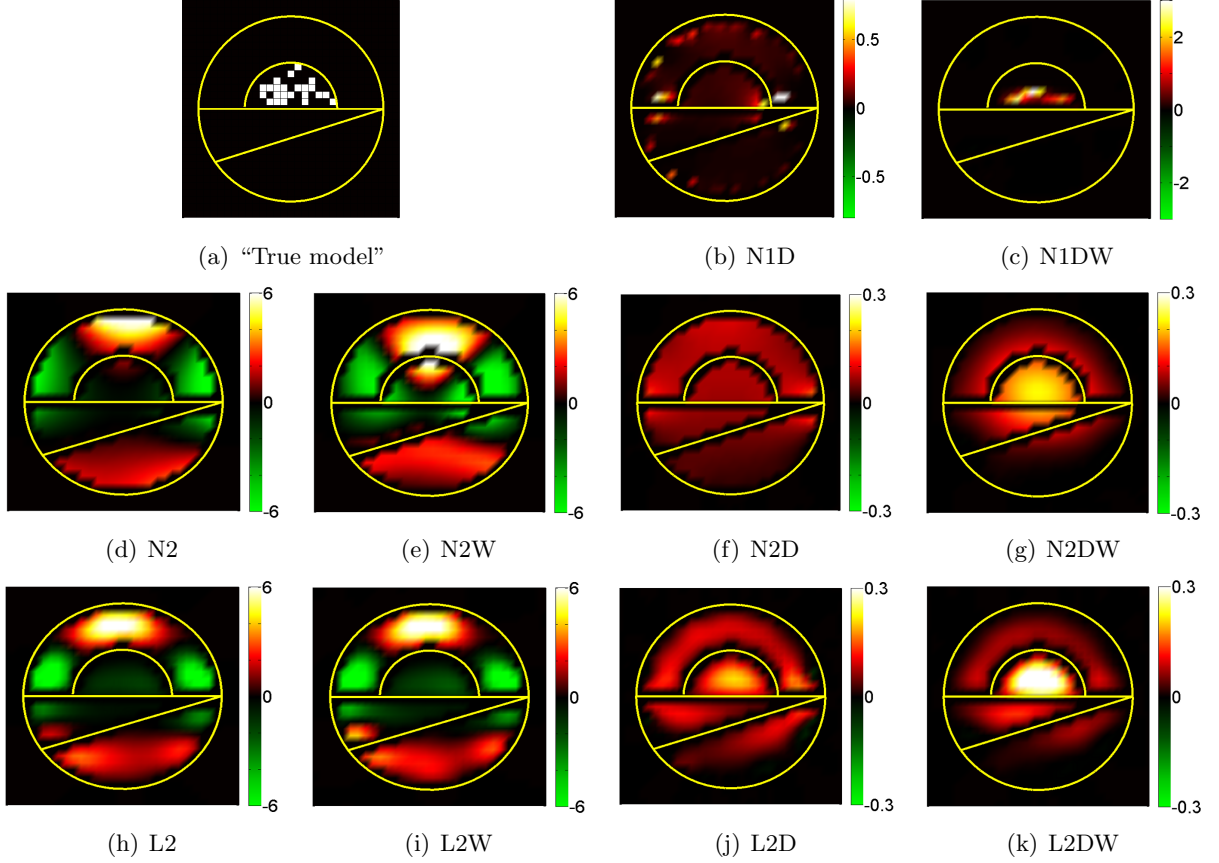


Figure 4: Reconstruction results for Experiment 2: sources only in Region 4. Note the different scale in (b) and (c).

Before discussing our findings, let us recall that all these reconstructions explain the data equally well: our task here is to decide which of the reconstructions are most useful for our purposes. Also, note that the inversions without D reconstruct the source q directly ($q = m$) whereas with D we reconstruct m , with $q = Dm$. Since these are different physical quantities the scales in the figures are rather different. Finally, note that in order to determine which method performs best, the results in Table 1 as well as Figs. 3–5 should be considered collectively rather than individually.

From Figs. 3–5 it is clear that the sparse reconstructions (N1) do not look physically reasonable. The sources for N1D are mostly arranged near the boundary, with a slight drift towards the center in Experiment 2. The sources reconstructed by N1DW are too few, but since they are sparse visual inspection may obscure things. Consulting the results in Table 1 it becomes clear that N1DW always puts the majority of sources in the center (Region 4). The numbers for N1D look somewhat better, but they are clearly insufficient to distinguish the activities. In short, a sparse prior is not adequate in the present context.

The methods without the D factor (N2, N2W, L2, L2W) produce reconstructions with positive peaks and negative valleys. This is because we are directly reconstructing the source q , and since it is made out of tripoles, they are reconstructed in the form of extended




Figure 5: Reconstruction results from Experiment 3: sources in Regions 1 and 4. Fig. 5(c) is on a saturated scale in order to display the weaker sources. The source in the center region has a magnitude of 40.

multipoles. The downside with these methods seems to be that significant activity “leaks” into inactive areas. From Table 1 we see that the N2(W) and L2(W) reconstructions are again unable to distinguish the two active regions. The results for L2 for detecting single region activation in Regions 1–3 (not shown here) are somewhat better, the main downside being the spurious activation of nearby regions. Finally, the results for L2D are clearly the best, both visually and numerically. See the highlighted entries in Table 1. The largest error is in Experiment 1, as 30% of the activity is erroneously attributed to the center Region 4. In Experiment 2 the converse happens, though the error is only 24%. However, since Region 4 is shielded by the other regions the ability to distinguish these conditions at all is limited, and this is probably the best one can do.

One could conceive of improving the reconstructions where D was not used by recovering a posteriori a corresponding tripole density s from q through solving $Ds = q$. However, our experiments indicate that this is not successful, resulting in even larger errors, since s is not properly regularized in this manner.

Numerous other source distributions were tested, resulting in the overall conclusion that the Laplacian weighted regularization term (3d), with no weighting (i.e., $w \equiv 1$), and

		Normalized activations										
	Muscle	True	N1D	N1DW	N2	N2W	N2D	N2DW	L2	L2W	L2D	L2DW
1	1	1.0	0.99	0.21	0.55	0.45	0.51	0.12	0.56	0.50	0.57	0.27
	2	0	0.004	0.11	0.28	0.28	0.12	0.12	0.27	0.3	0.07	0.1
	3	0	0.005	0.08	0.13	0.14	0.04	0.05	0.14	0.15	0.04	0.06
	4	0	0.005	0.98	0.04	0.14	0.32	0.38	0.028	0.06	0.31	0.58
2	1	0	0.33	0.001	0.54	0.39	0.33	0.15	0.51	0.44	0.24	0.12
	2	0	0.20	0.002	0.13	0.14	0.21	0.21	0.12	0.17	0.16	0.14
	3	0	0.13	0.001	0.24	0.19	0.16	0.06	0.16	0.19	0.14	0.06
	4	1.0	0.35	0.99	0.09	0.28	0.30	0.57	0.16	0.19	0.45	0.68
3	1	0.47	0.73	0.15	0.66	0.50	0.49	0.23	0.59	0.54	0.53	0.24
	2	0	0.06	0.06	0.13	0.16	0.13	0.13	0.25	0.28	0.08	0.09
	3	0	0.07	0.03	0.13	0.09	0.05	0.04	0.12	0.13	0.04	0.03
	4	0.53	0.14	0.88	0.07	0.25	0.33	0.60	0.03	0.04	0.35	0.64

Table 1: Normalized activations and their reconstructions in the three experiments of Section 3. Highlighted cells are discussed in the text.

including a nontrivial \mathcal{D} term as in (2c) (the combination referred to as L2D) performs better than any of the other methods. The 1-norm based regularization (3c) performs acceptably well only under unrealistic conditions of very small noise (less than 0.01%) and few sources (of the order of 10).

4 Numerical challenges in solving the L2D problem

In [7] most of the reconstructions were performed for the choice (2a), (3d), denoted L2 here, applying PCG iterations to the resulting large-scale linear system of equations. This led to acceptable runtimes.

In contrast, significant numerical difficulties were encountered when attempting to use L2D numerically in large scale problems such as those resulting from realistic modeling in three space dimensions. In this section we examine the source of these difficulties and show how they can be overcome.

The optimality conditions for (4a) after discretization read

$$(J^T J + \beta Z) m = J^T b, \quad (5)$$

where $Z = R''(m)$ is the Hessian of R , $J = QA^{-1}D$ with A the (symmetric positive definite) discretization matrix of the PDE operator in (1a)–(1c), D the (square symmetric) discretization matrix of the operator \mathcal{D} and Q the $d \times n$ projection matrix indicating where on the mesh the data is measured, and $b \in \mathbb{R}^d$ stands for the data. The regularization parameter β is chosen according to the discrepancy principle, such that the solution of (5) satisfies

$$\|Jm - b\| \equiv \|r\| \approx 1.5\epsilon\|b\|, \quad (6)$$

where ϵ is the noise level.

The differential operator that $A^{-1}D$ is a discretization of (with σ a frozen constant) has the symbol [27]

$$\frac{-\theta_3^2}{(-\sigma)(\theta_1^2 + \theta_2^2 + \theta_3^2)}, \quad (7)$$

for a Fourier mode $\exp\{i\mathbf{x} \cdot \boldsymbol{\theta}\}$ with frequency $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$. This symbol is bounded for any $\boldsymbol{\theta}$.¹

In principle, upon applying a CG-based iterative method for solving the problem (5), there is one a priori reason for expecting slower convergence with D compared to using (2a), and another a priori reason for expecting faster convergence. The reason for expecting slower convergence is that (7) is more complicated than when using $D = I$ (where the numerator is 1), leading to a more complicated spectrum, and since the convergence speed of CG is related to the number of eigenvalue clusters this would mean slower convergence. Also, without D the symbol tends to 0 in all three directions for high frequencies, whereas with the nontrivial D it does not. Since we are using Z as preconditioner, and the symbol of Z grows like $|\boldsymbol{\theta}|^2$, this means the preconditioner will be less effective for high frequencies with the nontrivial D . On the other hand, the nontrivial D is considered in the first place in order to enable a better regularization leading to a better solution, and this better solution (using the tripole basis) is smoother than the one typically obtained using the trivial $D = I$. Since CG iterations tend to gradually add more and more oscillatory singular vectors, obtaining a smoother solution that fits the data to a prescribed accuracy may require fewer iterations. In the sequel we see both of these conflicting a priori expectations reflected in the results.

The evaluation of the product Jv for a given vector v involves solving the linear system $Ay = Dv$ for y . (Thence, $Jv = Qy$.) This is done as in [5, 6] using PCG with incomplete Choleski (IC) decomposition as a preconditioner. In principle we would take the drop tolerance in the IC iteration as small as fast memory capabilities would allow, but we pretend that the discretization parameter n is so large that we may not use a direct method, hence take a nonzero drop tolerance. This gives rise to approximate evaluations of y and hence of Jv , and we may control the quality of such approximations in the usual way.

Now, in the CMG application A is necessarily a rough approximation due to the complexities of biological tissues, whose properties (such as the conductivity σ in (1a)) are often known inaccurately. Hence, for modeling purposes we do not require a relatively expensive accurate evaluation of Jv .

However, if we apply a PCG method to (5) where at each iteration matrix-vector products involving J and J^T are evaluated inaccurately, then $J^T J$ essentially looks different at each iteration! In this case the vector orthogonality that underlies the success of the CG method holds only roughly, and this is well-known to undermine the effectivity of the CG method [12, 14, 19]. Note that upon using inaccurate matrix-vector evaluations of Jv the variability in J is of a random sort that the presence of the differentiation matrix D arising from (2b)–(2c) amplifies, not of the structured sort that has a bounded effect as explained using (7). We may therefore expect numerical difficulties to arise when D is present when seeking inexpensive solutions for (5).

¹We thank Eldad Haber for pointing this out.

4.1 Using the Tikhonov method

Let m^β denote the solution of (5) for a given $\beta > 0$, with β_ϵ corresponding to equality in (6). The Tikhonov regularization has some desirable properties:

- As β is decreased, $\|Jm^\beta - b\|$ decreases monotonically.
- As β is decreased, $R(m^\beta)$ increases monotonically.
- In the absence of noise, and assuming that a consistent solution to $Jm = b$ exists, we obtain $m^\beta \rightarrow m^0$ such that $Jm^0 = b$.

However, for realistic noise levels (say 1% or 5%), β must be taken so small that the matrix in (5) is highly ill-conditioned. Fig. 6 shows disturbing results when D is present

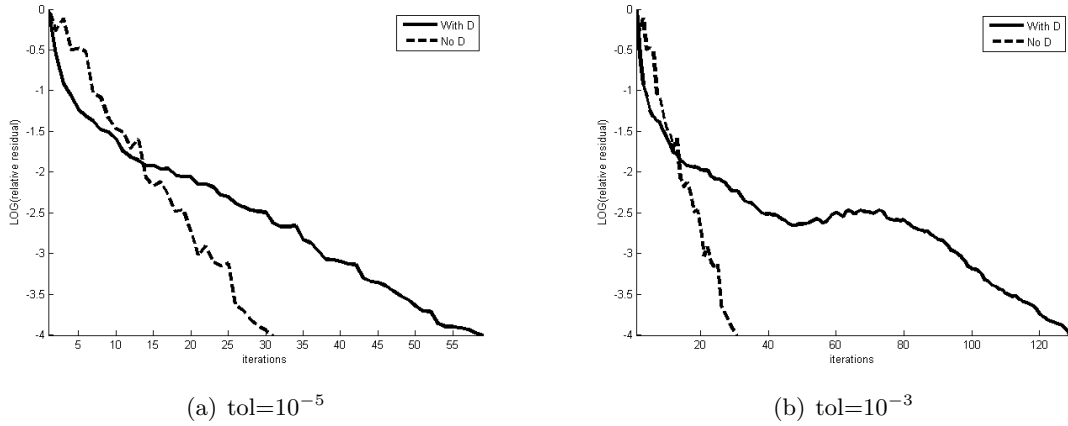


Figure 6: Iteration progress for the CMG problem using Tikhonov regularization on a 32^2 uniform grid with 5% noise in the data. Shown are relative residual plots with $D \sim \partial_{zz}^h$ and without D for $\beta = 10^{-3}$, evaluating J and preconditioner to tolerance level tol . The y -axis corresponds to the relative residual $\|(J^T J + \beta Z)m - J^T b\| / \|J^T b\|$.

and matrix-vector products involving J are not accurately evaluated. We apply the PCG method to (5) and solve to a relative tolerance of 10^{-4} for the outer iteration. For accurate forward solves in the inner CG loop involving A with tolerance 10^{-5} the equations with D converge half as fast as without, which is manageable. However, if the inner tolerance is increased to 10^{-3} the system with D is affected severely, whereas the effect without D is small.

The situation can be improved by explicitly evaluating and storing J^T , and subsequently using it for matrix-vector products. This evaluation (still involving an iterative procedure for A^{-1}) may not be very accurate, but it is the same matrix J that then gets used for different PCG iterations. Upon pre-constructing J^T and hence J , we are in effect applying a genuine CG method to a system of equations that has been perturbed because of the inaccuracy in constructing J . This option relies for its feasibility on the number of sensors in the experiment being small. Recall that J is $d \times n$, where for the CMG problem d is the number of sensors and n is the number of grid points. In particular, $d \ll n$. We

then construct J^T by solving d problems with A to form $A^{-1}Q^T$, and then forming $J^T = DA^{-1}Q^T$ upon multiplication by D . This approach is reminiscent of the adjoint method [3]. Note that it requires storage of d long columns. We still have to evaluate the preconditioner multiplications with an inner CG iteration.

The convergence behavior for the Tikhonov method is displayed in Fig. 7 for inner tolerances of 10^{-3} (same as in Fig. 6(b)) and 10^{-2} . The sensitivity in the iterative linear algebra method is now significantly reduced. Yet at a tolerance of 10^{-2} efficiency with D deteriorates again. In this case the cause is the inaccuracy in evaluating the preconditioner. The dotted curve in Fig. 7(b) depicts iteration progress when the preconditioner is evaluated more accurately, resulting in a well behaved iteration.

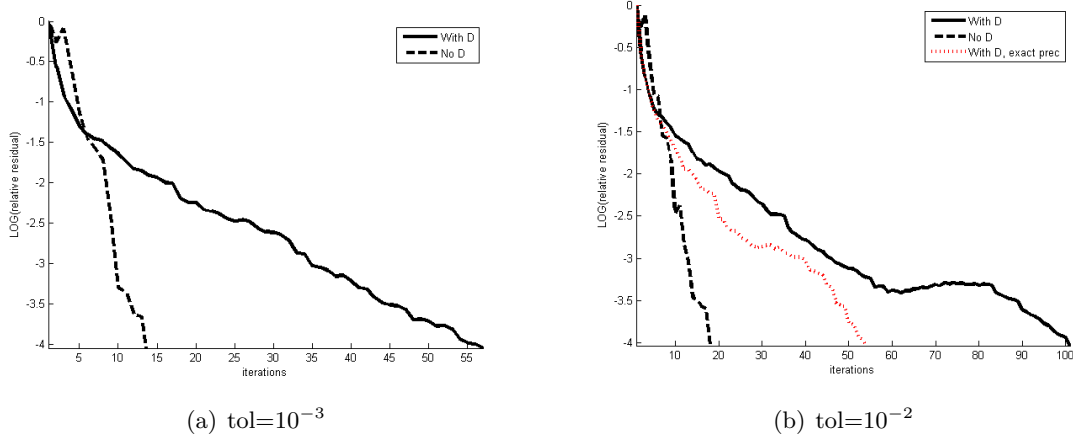


Figure 7: Iteration progress for the CMG problem using Tikhonov regularization on a 32^2 uniform grid with 5% noise in the data. Shown are residual plots as a function of iteration number with $D \sim \partial_{zz}^h$ and without D for $\beta = 10^{-3}$, pre-evaluating J to tolerance level tol , with the preconditioner computed fully iteratively to the same tolerance. The dotted line in (b) depicts a computation with a more accurate evaluation of the preconditioner to 10^{-4} . The y -axis corresponds to the relative residual $\|(J^T J + \beta Z)m - J^T b\| / \|J^T b\|$.

In terms of work units, where one work unit is the cost of solving one problem with A or Z , and it_{tikh} is the number of (outer) PCG iterations, the fully iterative method requires $3 \times it_{tikh}$ units, whereas the method that explicitly evaluates J^T requires roughly $d + it_{tikh}$ units. So, ignoring storage considerations, for a small number of sensors d the method of precomputation may be faster as well.

A remaining issue with the Tikhonov approach is finding β and satisfying a stopping criterion such as (6). Note that β_ϵ is unknown. If a given $\beta > \beta_\epsilon$ is too large then (6) may not be satisfiable at all. On the other hand, if β is significantly smaller than β_ϵ then the PCG iterations may yield a solution that satisfies (6) without properly solving (5).

The above difficulties can be circumvented using iterative regularization methods, which essentially achieve the same goals as a Tikhonov regularization. This entails setting $\beta = 0$ in (5), and instead using the regularizing effect of the PCG iterations directly. We next turn our attention to these methods.

4.2 Iterative (dynamic) regularization

Let us multiply (5) by Z^{-1} and set $\beta = 0$, obtaining

$$Z^{-1}J^T Jm = Z^{-1}J^T b. \quad (8)$$

We solve this by applying preconditioned CGLS to the underdetermined least squares problem

$$\min_m \|r\|, \quad r = b - Jm, \quad (9)$$

using Z as a preconditioner. The iteration is terminated when either $\|r\|$ stops decreasing or (6) holds, whichever occurs first. This is the method of *iterative regularization* [16]. It relates to the method of dynamic regularization of [4, 5, 6], but without an outer iteration and with a different stopping criterion.

Note that $Z = R''$ is a special preconditioner since it is related to the regularization operator in the formulation (5), (3a). Specifically, this preconditioner enforces a penalty on $R(m)$, similar to the penalty in the Tikhonov regularization.

Thus, whereas the Tikhonov formulation encounters serious trouble in selecting the regularization parameter, and also can lead to a condition (6) that is impossible to satisfy, the iterative regularization approach applied directly to the singular system elegantly avoids these thorny issues.

A few iterations of this PCG algorithm should suffice to obtain such a solution. Starting with the initial guess $m_0 \equiv 0$ and generating iterates m_k , $k = 1, 2, \dots$, the following can be shown [16]:

- The residual $\|Jm_k - b\|$ decreases monotonically in k , justifying our stopping criterion.
- The regularization penalty $R(m_k)$ increases monotonically in k .

Fig. 8 depicts the method in action and also shows the effect of inaccuracies in the inner CG process. In contrast to the Tikhonov plots in Figs. 6–7, the y -axis here and in Figs. 9 and 11 below depicts $\|Jm - b\|/\|b\|$, as this quantity can be controlled directly when using iterative regularization.

Note that, even without precomputing J , an inner tolerance up to 10^{-2} can be tolerated. Still, this method also breaks down once the forward solve is done to a relative tolerance of 1% or coarser, as is the inevitable fate of such an inexact CG method.

Compared to Tikhonov, the present method is faster by about an order of magnitude and somewhat more tolerant to roughly constructed matrix-vector products involving J and J^T . Also the effect of D is the reverse of its effect when utilizing Tikhonov, and convergence is faster with D than without. The reason for this seemingly paradoxical behavior is that the number of outer PCG iterations is not directly related to how large the penalty function $R(m)$ needs to be to satisfy the data., i.e., to the amount of regularization. It is easier to satisfy the data with \mathcal{D} because the tripole nature of the sources is already built in.

Upon pre-evaluating J^T as described earlier, rougher forward solves are tolerable, see Fig. 9. Still, overall pre-computing J is less crucial here than in the Tikhonov context.

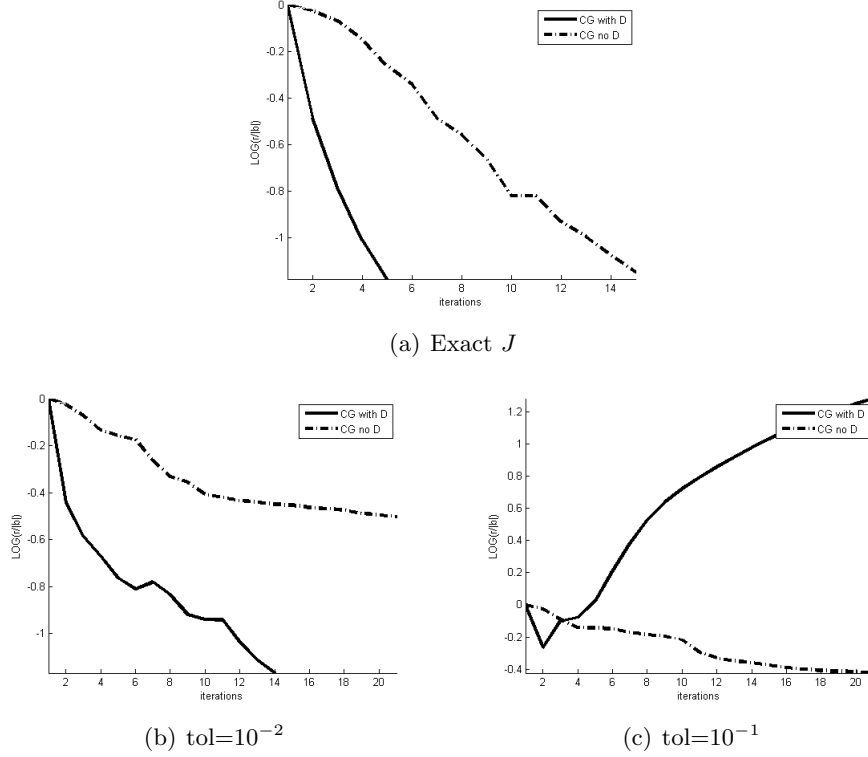


Figure 8: Iteration progress for the CMG problem using dynamic (or iterative) regularization on a 32^2 uniform grid with 5% noise in the data. Results with evaluation of J and preconditioner to different tolerances are displayed. The y -axis corresponds to the relative residual $\|Jm - b\|/\|b\|$.

5 Simulation results in three dimensions

A realistic 3D model of the upper arm was constructed from MRI data of a human subject and segmented into different anatomical regions: brachialis, biceps, triceps, fat and skin, and bone.

We applied a template registration technique to identify organ boundaries in MRI volumes. Our template consisted of discrete surfaces representing muscles (biceps brachii, triceps, brachialis), bones (humerus) and the skin, yielding ≈ 8000 vertices in total. To facilitate further meshing into finite elements, vertices were sampled on regularly spaced slices. The template was interactively registered to a reference MRI volume, to get a prior appearance model, and it was then automatically deformed to iteratively fit any patient-specific geometry. We applied the “shape matching” framework presented in [11].

The resulting segmented 3D geometry was then imported into the finite element software COMSOL. Conductivities were assigned to the different tissues according to the values reported in [23]. We assumed the muscle fiber directions to be aligned along the arm, and assigned the anisotropic conductivity tensor in the muscle tissue based on this. The 3D geometry was then meshed with tetrahedral elements, and the inverse model equations were

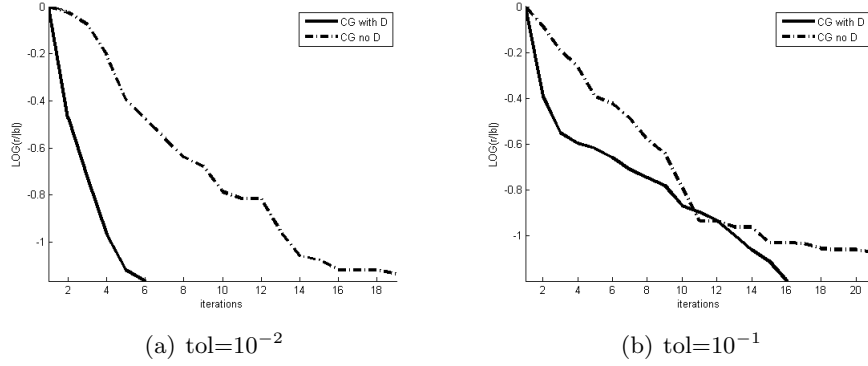


Figure 9: Iteration progress for the CMG problem using dynamic (or iterative) regularization on a 32^2 uniform grid with 5% noise in the data, and employing pre-evaluation of J to different tolerances. The y -axis corresponds to the relative residual $\|Jm - b\|/\|b\|$, and iteration is stopped once (6) is satisfied.

discretized in COMSOL and exported to MATLAB for solution. From this we can generate synthetic data, placing tripole current sources in the muscles, or we can use measured data. After obtaining the sEMG data we solved the equations within MATLAB and obtained reconstructed current sources in the muscles. We finally imported the result back into COMSOL for analysis and visualization. Fig. 10 shows a segmented slice of the MRI data used to create the upper arm model and the 3D geometry created after segmentation. The FEM grid used had about $n = 50,000$ degrees of freedom; for computation of the synthetic data the mesh was about twice as fine.

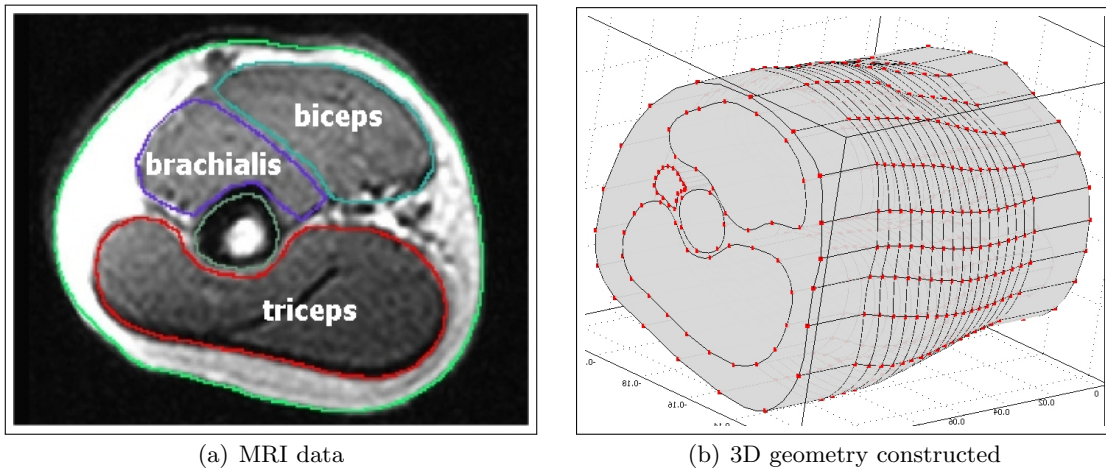


Figure 10: MRI data was segmented into a geometry model.

Surface data was taken using 21 simulated electrodes arranged in 3 concentric bands around the center of the arm segment. From each signal we subtracted the average signal of all sensors which thus functions as the ground level. This configuration is realizable by disposable XLTEK BlueSensor Snap Electrodes that we plan to use in the near future for

recording real data.

Various reconstructions were performed using our best method L2D from Section 3 and also, for comparison, using the L2 method. It was found that the numerical methods described in this article work very well in three space variables, and efficient reconstructions could be obtained also with the matrix D present. Fig. 11 displays the iteration progress using dynamic regularization.




Figure 11: Iteration progress for the realistic CMG problem using dynamic (or iterative) regularization with 5% noise in the data. Shown are efficiency plots with and without D , with J computed iteratively to tolerance 10^{-2} . The y -axis corresponds to the relative residual $\|Jm - b\|/\|b\|$.

		Normalized activations		
		True	L2	L2D
1	Biceps	1	0.73	0.91
	Brachialis	0	0.24	0.05
	Triceps	0	0.03	0.04
2	Biceps	0	0.16	0.02
	Brachialis	1	0.65	0.96
	Triceps	0	0.19	0.02

Table 2: Normalized activations and their reconstructions in the two examples of Section 5.

The quality of the reconstructions with \mathcal{D} was significantly better than without, mainly because less “spill-over” of activations into other muscle areas occurred.

We illustrate this with two examples. In both, 1000 artificial tripole sources with an extension of 1cm were used. They were placed at random locations in the biceps (Example 1) and in the brachialis (Example 2), respectively. The synthesized data was polluted by 5% Gaussian noise, and then reconstructions were performed. Matrix-vector products Jv were computed iteratively to a tolerance of 10^{-2} . Reconstructions using L2D converged in 7 iterations, whereas 14 iterations were required upon using L2. Thus, the “numerically troublesome” method is now the faster one, too.

In Table 2 we show the reconstructed normalized activities, demonstrating again the superiority of L2D over L2. Fig. 12 shows cross-sections through the central region of the power density of the reconstructed m . Visually the reconstructions with D are clearly superior. Further experiments with different synthetic sources support these conclusions.

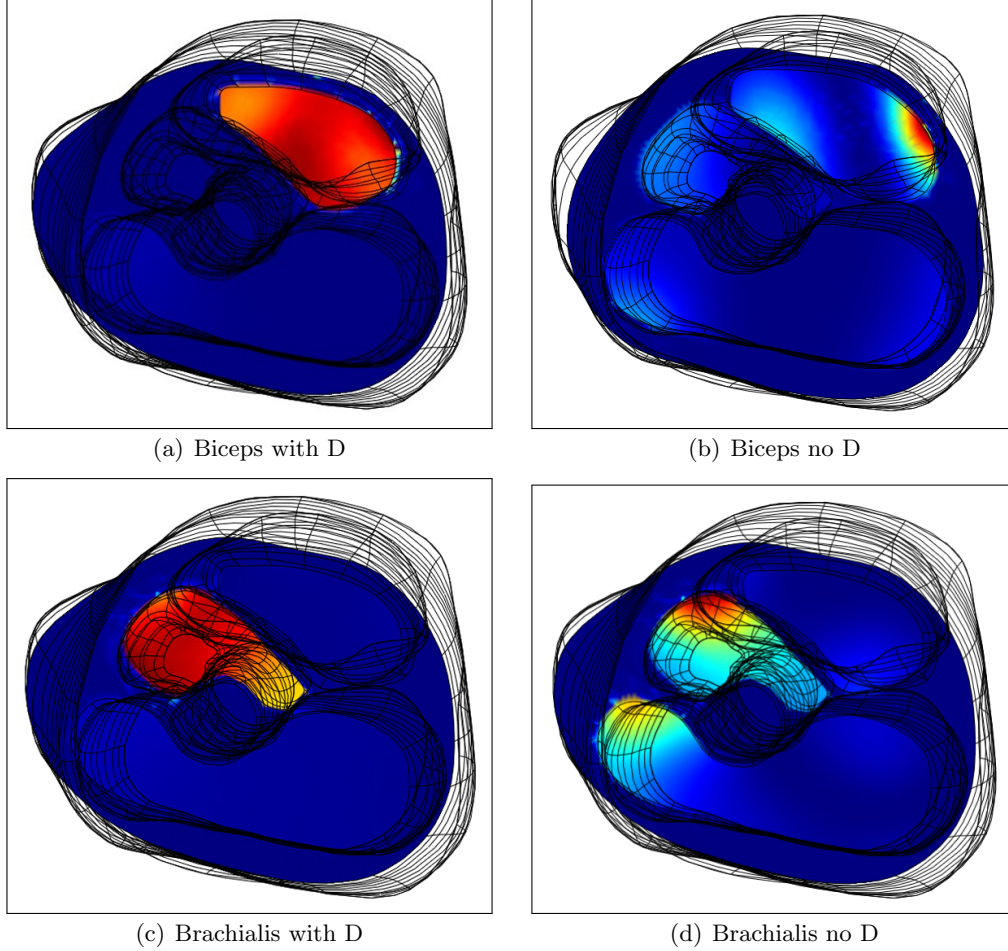


Figure 12: Reconstruction of activities with and without D with sources in biceps only and in brachialis only.

6 Conclusions

Several possible methods for inverting sEMG data were investigated. We considered reconstructing the current densities q directly, or instead through the tripole densities m such that $q = \mathcal{D}m$ where the operator \mathcal{D} behaves like a second derivative along the arm's length. Further, to add more a priori information and to regularize the problem several regularization candidates in a Tikhonov-type setting were investigated: the L_2 norm of m , the L_1 norm of m (to attempt a sparse reconstruction), and the L_2 norm of ∇m enforcing smoothness and yielding a Laplacian term in the corresponding Euler-Lagrange equations.

Since there is an inevitable bias against deep sources, we also tested a well-known weighting strategy to overcome such a bias.

Based on 2D simulated experiments it was found that one particular method performs better than the others. This method uses the tripole density function as the primary variable and regularizes the problem with a smoothness penalty, without depth bias correction (as this overcompensates). There is no advantage to sparse priors in our application.

The use of the tripole density function introduces an additional term in the large linear system to solve, and it was found that significant numerical difficulties arise if the matrix-vector products arising in iterative solution methods are computed without special care. The most efficient method turned out to be iterative (or dynamic) regularization with a special preconditioner and a careful stopping criterion.

Additional resilience to loose forward solves is obtained if this method is combined with a precomputed solution matrix for the adjoint problem. The latter is feasible or efficient only if the number of sensors d is small, since an $n \times d$ matrix needs to be precomputed and stored.

We then applied the method to an FEM forward model in three dimensions constructed from MRI data. The numerical behavior of the latter, larger system was found to be consistent with the two-dimensional results. In particular, the presence of the D term was found to significantly improve the quality of reconstructions in the simulated experiments, as compared to the results in [7], and these reconstructions were efficiently computed by our method.

Acknowledgement

We thank Benjamin Gilles for the MRI data and segmentation.

References

- [1] U. Ascher, E. Haber, and H. Huang. On effective methods for implicit piecewise smooth surface recovery. *SIAM J. Sci. Comput.*, 28:339–358, 2006.
- [2] E. Chauvet, O. Fokapu, and D. Gamet. Inverse problem in the surface EMG: A feasibility study. In *Proceedings of the 23rd Annual EMBS international Conference, Istanbul, Turkey*, pages 1048–1050, 2001.
- [3] G. Chavent and P. Leronnier. Identification de la non-linearité d’une equation parabolique quasi-lineaire. *Applied Math and Opt.*, 1:121–162, 1974.
- [4] K. van den Doel and U. Ascher. On level set regularization for highly ill-posed distributed parameter estimation problems. *J. Comp. Phys.*, 216:707–723, 2006.
- [5] K. van den Doel and U. Ascher. Dynamic level set regularization for large distributed parameter estimation problems. *Inverse Problems*, 23:1271–1288, 2007.
- [6] K. van den Doel, U. Ascher, and A. Leitao. Multiple level sets for piecewise constant surface reconstruction in highly ill-posed problems. *Journal of Scientific Computation*, 42, 2010.

- [7] K. van den Doel, U. Ascher, and D. Pai. Computed myography: three dimensional reconstruction of motor functions from surface EMG data. *Inverse Problems*, 24:1–17, 2009.
- [8] G. Drost, D. F. Stegeman, B. G. M. van Engelen, and M. J. Zwarts. Clinical applications of high-density surface EMG: A systematic review. *Journal of Electromyography and Kinesiology*, 16:586–602, 2006.
- [9] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Kluwer, 1996.
- [10] M. Fuchs, H. A. Wischmann, and M. Wagner. Generalized minimum norm least squares reconstruction algorithms. *ISBET newsletter*, 5:8–1, 1994.
- [11] B. Gilles and D.K. Pai. Fast musculoskeletal registration based on shape matching. *Proc. of MICCAI'08*, 2:822–829, 2008.
- [12] G. Golub and Q. Ye. Inexact preconditioned conjugate gradient method with inner-outer iteration. *SIAM J. Scient. Comp.*, 21:1305–1320, 2000.
- [13] R. E. Greenblatt. Probabilistic reconstruction of multiple sources in the neuroelectromagnetic inverse problem. *Inverse problems*, 9:271–284, 1993.
- [14] E. Haber and U. Ascher. Preconditioned all-at-one methods for large, sparse parameter estimation problems. *Inverse Problems*, 17:1847–1864, 2001.
- [15] M. S. Hämäläinen and R. J. Ilmoniemi. Interpreting magnetic fields of the brain—minimum norm estimates. *Med. Biol. Eng. Comput.*, 32:35–42, 1994.
- [16] P. C. Hansen. *Rank Deficient and Ill-Posed Problems*. SIAM, Philadelphia, 1998.
- [17] H. L. F. Helmholtz. Ueber einige Gesetze der Verheilung elektrischer Ströme in körperlicher Leitern mit Anwendung auf die thierisch-elektrischen Versuche. *Ann. Physik und Chemie*, 9:211–233, 1853.
- [18] F. Hettlich and W. Rundell. Iterative methods for the reconstruction of an inverse potential problem. *Inverse Problems*, 12:251–266, 1996.
- [19] H. Huang. *Efficient Reconstruction of 2D Images and 3D Surfaces*. PhD thesis, University of BC, Vancouver, 2008.
- [20] V. Isakov. *Inverse Problems for Partial Differential Equations*. Springer, 2006.
- [21] R.A. Jesinger and V.L. Stonick. Processing signals from surface electrode arrays for noninvasive 3D mapping of muscle activity. In *IEEE DSP Workshop Prodceedings*, pages 57–60, 1994.
- [22] R.A. LoPresti, E. F. and Jesinger and V.L. Stonick. Identifying significant frequencies in surface EMG signals for localization of muscle activity. In *IEEE EMBS Conference Prodceedings*, pages 967–968, 1995.

- [23] M. M. Lowery, N. S. Stoykov, A. Taflove, and T. A. Kuiken. A multiple-layer finite-element model of the surface EMG signal. *IEEE trans. on Biomedical Engineering*, 49(5):446–454, 2002.
- [24] R. Merletti and P. A. Parker. *Electromyography. Physiology, Engineering and Noninvasive Applications*. IEEE Press, John Wiley and Sons, 2004.
- [25] C. M. Michel, M. M. Murray, G. Lantz, S Gonzalez, L. Spinelli, and R. G. de Peralta. EEG source imaging. *Clinical neurophysiology*, 115:2195–2222, 2004.
- [26] J.T. Stonick, R.A. Jesinger, V.L. Stonick, and S.B. Baumann. Estimation and localization of multiple dipole sources for noninvasive mapping of muscle activity. *IEEE International Conference on IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP-96, Conference Proceedings*, 5:2912–2915, 1996.
- [27] U. Trottenberg, C. Oosterlee, and A. Schuller. *Multigrid*. Academic Press, 2001.