

Comparing the Performance of Solvers for a Bioelectric Field Problem

Marcus Mohr¹ and Bart Vanrumste^{2*}

¹ System Simulation Group of the Computer Science Department,
Friedrich-Alexander-University Erlangen-Nuremberg

Marcus.Mohr@cs.fau.de

² Electrical and Computer Engineering, University of Canterbury
b.vanrumste@elec.canterbury.ac.nz

Abstract. The model-based reconstruction of electrical brain activity from electroencephalographic measurements is of constantly growing importance in the fields of Neurology and Neurosurgery. Algorithms for this task involve the solution of a 3D Poisson problem on a complicated geometry and with non-continuous coefficients for a considerable number of different right hand sides. Thus efficient solvers for this subtask are required. We will report on our experiences with different iterative solvers, Successive Overrelaxation, (Preconditioned) Conjugate Gradients, and Algebraic Multigrid, for a discretisation based on cell-centred finite-differences.

1 Introduction

The electroencephalogram (EEG) is a major diagnostical tool to determine the state of the brain. In recent years the model-based analysis of such voltage measurements has substantially gained in importance. Given the measured data, the typical task is to reconstruct the sources inside the brain responsible for the potential field. The results are then used in the planning of brain surgery and even during the surgery itself. Figure 1 shows an application from epilepsy surgery [11]. Here an epileptic focus to be removed by surgery was reconstructed. Another application is to determine important brain areas, which must not be damaged during an operation, and to insert this information into the surgeon's microscope. The reconstruction is denoted the inverse EEG problem, because it is the inverse of the forward problem, where sources are known / given and the potential distribution inside the head and on its surface has to be computed. The forward problem constitutes a classical elliptic boundary value problem to be solved on a 3D representation of the patient's head. All solution approaches to the inverse EEG problem involve the solution of a considerable number of such forward problems.

* At the time of this research, the author worked for the Department of Neurology, University of Gent

Medical applications naturally require a high level of precision. It is becoming standard nowadays to base the numerical computations on the individual geometry of a patient's head, which can be obtained e.g. from magnetic resonance imaging (MRI). The electric conductivities of the different tissues are usually assumed to be constant and isotropic, but the determination of patient specific anisotropic values from diffusion weighted MRI is ongoing research, [10]. Thus there is growing interest in volume based discretisation techniques which are about to replace the traditional boundary element approach. This transition to volume discretised realistic head models is encumbered by the fact that there is still a lack of fast solvers for the forward problem, see e.g. [5].

In this paper we will report on results with different iterative solvers for a discretisation of the forward problem with a cell-centred finite difference scheme. The list of solvers includes Successive-Overrelaxation as an example of a stationary iteration method, Conjugate Gradients and Preconditioned Conjugate Gradients for Krylov subspace methods and a variant of Algebraic Multigrid. For a description of these methods see e.g. [6] and [7].

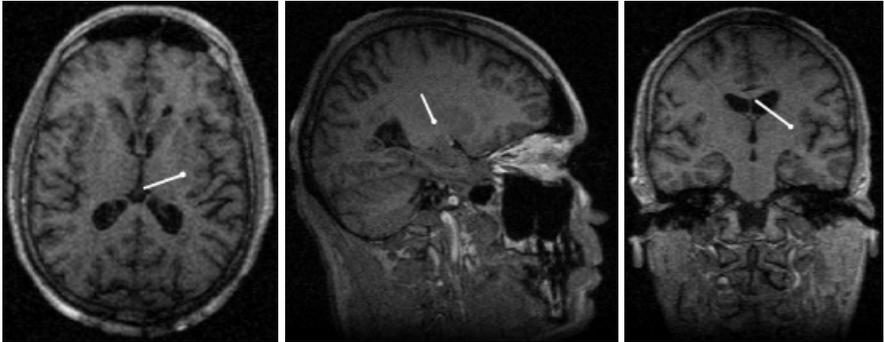


Fig. 1. Axial, sagittal, and cordial cuts through an MRI scan of a patient's head. Superimposed is an epileptic focus localised from EEG data.

2 Problem Formulation

In bioelectric field simulation the patient's head is modelled as a volume conductor, i. e. a contiguous, passively conducting medium. Inside the head we have neural current densities which constitute the cause for the electrical and potential field. Due to the temporal behaviour of the sources, which is typically < 1 kHz and the physiological conductivities (e.g. brain 0.2 S/m, skull 0.015 S/m [3]) one can assume a quasi-static behaviour. This leads to the central equation

$$\nabla \cdot (\sigma \nabla \Phi) = \nabla \cdot I_V \quad . \quad (1)$$

which relates the potential field Φ to the current densities I_V of the neural sources. The term σ denotes the conductivity tensor. Together with boundary conditions on the current flow through the scalp

$$\sigma \frac{\partial \Phi}{\partial \mathbf{n}} = g \quad (2)$$

this constitutes an elliptic boundary value problem. The classical task of numerical simulation would be to compute the potential field Φ for given source $\nabla \cdot I_V$ and boundary terms g . This is referred to as the *forward problem*. The related *inverse EEG problem* is to reconstruct from g and measurements of Φ at some electrodes either the sources or the potential at some internal interfaces, esp. on the cortex.

There are three standard approaches to the inverse EEG problem: *dipole localisation*, *current density analysis* and *deviation scans*, see e.g. [11, 2]. Important in our context is, that all of these involve the setup of a so called *lead field matrix*. This matrix relates the momentum and orientation of one or more dipoles at given locations to the potential values at measurement electrodes.

Assume that a dipole is given with position \mathbf{r} and orientation \mathbf{d} and let \mathcal{E} be a set of N electrodes. The vector $M(\mathbf{r}, \mathbf{d})$ of potentials generated by the dipole at the electrodes in \mathcal{E} is given by

$$M(\mathbf{r}, \mathbf{d}) = \hat{L}(\mathbf{r}) \cdot \mathbf{d} = R \cdot L(\mathbf{r}) \cdot \mathbf{d} \quad (3)$$

Here $\hat{L} \in \mathbb{R}^{N \times 3}$ is the lead field matrix. It can be split into two parts. A matrix $L \in \mathbb{R}^{(N-1) \times 3}$ that maps the dipole orientation onto potential differences between the electrodes and a referencing matrix $R \in \mathbb{R}^{N \times (N-1)}$ that turns these $(N-1)$ differences into N scalar potential values. Each row of L is of the following form

$$(\Phi_{AB}^x(\mathbf{r}), \Phi_{AB}^y(\mathbf{r}), \Phi_{AB}^z(\mathbf{r})) \quad (4)$$

where $\Phi_{AB}^k(\mathbf{r})$ denotes the potential difference between electrode A and B induced by a unit dipole at position \mathbf{r} which is oriented in k -direction. So in principle every entry of the lead field matrix requires the solution of a single forward problem. In all three approaches for the inverse EEG problem the lead field matrix has either to be assembled for a single dipole at a large number of varying positions or for a large number of different dipoles at fixed positions. This constitutes a major part of the computational work. Fortunately this amount of work can (in some cases) drastically be reduced with the help of the reciprocity theorem of Helmholtz, see e. g. [11, 12]. It allows the simple setup of the lead field matrix by interpolation from potential distributions computed in a preparatory step. Selecting $N - 1$ pairs of electrodes from \mathcal{E} this step consist of solving problem (1,2) with one electrode chosen as unit sink and the other one as unit source. This leave us with the solution of $N - 1$ Poisson problems with jumping coefficients on a complicated 3D domain.

3 Discrete Model

In order to solve the boundary value problem (1,2) in terms of the bioelectric potentials we need to pose the problem in a computationally tractable way. Modern medical imaging techniques make it feasible to employ the individual patient's anatomy in the creation of a volume conductor model of the head. Geometry information is typically derived from magnetic resonance imaging (MRI), while the conductivity is assumed to be isotropic and constant for each type of tissue and experimental values from literature are taken.

For our experiments in Sect. 5 we have used MR-images to form a 3D voxel model. In a segmentation process each voxel is assigned to one of the four compartment types scalp, skull, brain and air and constant isotropic conductivities are assumed in each of them. Note however, that the described approach is not limited to a certain number of compartments and in principle can also handle anisotropic conductivities. This option is important, since it is anticipated that in the near future diffusion weighted MRI will allow to approximate the locally varying and anisotropic conductivity tensors on a patient basis, see e.g. [10]. For details on the creation of the employed 3D voxel model see [11].

We discretise the boundary value problem (1,2) by means of Finite Differences. We use the box integration scheme and model the electrodes as point sources within the voxel that is nearest to the position of the centre of the respective electrode. This leads to homogenous Neumann boundary conditions. The resolution of our head models is fine enough that for the given spacing of the two electrodes in a pair will never lie within the same voxel. Let us denote by k_{source} and k_{sink} the index of the voxel containing the source and the sink. Then we get that the discrete approximation Φ^h of (1,2) has to satisfy

$$\left(\sum_{k \in \mathcal{N}_l} \gamma_k \right) \Phi_l^h - \sum_{k \in \mathcal{N}_l} \gamma_k \Phi_k^h = I(\delta_{k_{\text{source}},l} - \delta_{k_{\text{sink}},l}) \quad (5)$$

for every voxel V_l that belongs to the head. Here $\delta_{i,j}$ denotes the Kronecker symbol and \mathcal{N}_l the index set of the six voxels neighbouring V_l . The stencil coefficients γ_j are a mesh-size weighted harmonic average of the conductivities in the central voxel and the corresponding neighbour. For the eastern neighbour γ_j takes the form

$$\gamma_{\text{east}} = \frac{2h_y h_z}{h_x} \cdot \frac{\sigma_l \cdot \sigma_{\text{east}}}{\sigma_l + \sigma_{\text{east}}} \quad (6)$$

and analogously for the other five neighbours. Note that $\gamma_k \geq 0$ since σ is non-negative.

Let us denote by A the problem matrix of the linear system resulting from (5). Assuming that we have a contiguous model without isolated head cells it is easy to see from (6) and (5) that A will be a symmetric, positive semidefinite matrix. It is also a singular M-matrix in the classical notation of Berman & Plemmons. We know therefore that $\text{rank}(A) = n - 1$ and the kernel is of dimension one [1]. In fact, it is spanned by the vector $e = (1, \dots, 1)^T$. Let b be the right hand

side we get from (5), then $e^T b = 0$. Thus the problem $Ax = b$ is consistent and possesses infinitely many solutions differing only in an additive constant.

When trying to find the solution of the system, one can either work with the singular problem directly, or introduce an additional constraint, that fixes one element of the set of all solutions and solve the corresponding regular problem. The easiest approach is to fix the value of the potential Φ^h to 0 in one voxel. This leads to a problem with a regular M-matrix and its solution obviously solves the initial problem with $\Phi^h = 0$ in the respective voxel. We will compare these two possibilities.

4 Iterative Solvers

In the literature on bioelectric field problems typically Krylov subspace methods are mentioned for the solution of the related forward problems, see e.g. [2]. We consider representatives of this class and compare them to other methods:

- Successive Over-Relaxation (SOR)
- Conjugate Gradients (CG)
- Conjugate Gradients preconditioned by symmetric SOR (PCG)
- Algebraic Multigrid (AMG)

The SOR method is a representative of the classical stationary methods. It is known to be not the optimal choice as far as convergence is concerned, but it has a very simple structure. Thus it is a good candidate for an optimised implementation.

The CG method is the typical algorithm from the large class of Krylov subspace methods. The convergence of the CG method depends on the condition number, or more precisely on the spectrum, of the problem matrix. It is therefore seldom used without preconditioning. We have chosen a symmetric SOR preconditioner for this purpose.

The last contestant is an algebraic multigrid method. Multigrid methods in general are known to be very efficient solvers for elliptic boundary value problems. They employ a hierarchy of grid levels to treat individual problem components. Unfortunately finding the proper components, i.e. transfer operators, coarsening strategies, etc. can be quite tedious in the case of complex geometries and/or jumping coefficients. Therefore the idea of algebraic multigrid methods is again attracting increased attention. Here a “grid hierarchy” and inter-grid transfer functions are derived automatically from the problem matrix.

While these methods have been developed for regular linear systems, they can also be applied in our semi-definite case. In the case of a consistent right hand side convergence can be guaranteed for SOR and (P)CG, while for AMG theoretical results are more complicated, see [1, 9].

5 Numerical Experiments

For our numerical experiments we implemented all algorithms, except AMG, in a grid- and a matrix-based version. As AMG method we took *BoomerAMG*, see [7]. Implementation details can be found in [9].

We compare the performance with respect to two data sets, in the following denoted by data set *A* and *B*. They have been created from real patient MRI scans at the Epilepsy Monitoring Unit of the University of Gent, Belgium. The data sets have a different resolution and come from two different subjects. The edge length h of the voxels is the same in all three dimensions. For the conductivities we use the values given in [3]. We scale our linear system by $(\sigma_{\text{skull}}h)^{-1}$ and only use the relative values $(\sigma_{\text{scalp}} : \sigma_{\text{skull}} : \sigma_{\text{brain}} = 16 : 1 : 16)$. Note that the differences in the conductivities are only one order of magnitude. So we are faced with an interface problem with only moderate jumps. An overview of the properties of the two data sets and of the resulting linear system is given in Tab. 1.

For both patients an EEG was recorded using 27 electrodes, placed according to the international 10–20 standard [8], with three additional electrodes positioned on the temporal lobe areas on both sides. Thus the preparatory step, cf. Sect. 2, that provides the data needed for setup of the lead field matrix consists in the solution of 26 forward problems for as many electrode pairs.

From the four tested methods the CG algorithm is the only one that does not depend on the choice of a special parameter. In case of the SOR and PCG(SSOR) method we have to specify the over-relaxation parameter ω . For AMG a threshold parameter, which allows for the distinction of “weak” and “strong” inter-node connections, must be chosen. This distinction is critical for the construction of the grid hierarchy, see e.g. [7]. We have chosen two of the 26 electrode pairs to test the dependency of the algorithms on these parameters. In these and all following tests we accepted an approximate solution as soon as its residual measured in the Euclidean norm became smaller than 10^{-8} .

Figure 2 shows the dependency of the number of SOR iterations for the case of a singular matrix. We see that this number varies considerable even for the small interval $\omega \in [1.89, 1.99]$. The optimal values seem to be in the vicinity of 1.935 for data set A and 1.970 for data set B. We have employed these values for all subsequent tests. When we apply SOR to the regular matrix, we get basically

Table 1. Details of the two data sets and the corresponding singular linear systems

Data Set	A	B
number of voxels / cube dimensions	65^3	129^3
edge length of voxels	3.92mm	1.96mm
number of head voxels / unknowns	70,830	541,402
number of non-zero matrix entries	482,184	3,738,624
matrix sparsity	0.01%	0.001%

the same picture, with two differences. On the one hand the optimal ω is larger (about 1.9935 for dataset A and about 1.9977 for dataset B). On the other hand also the number of iterations needed to reach the desired accuracy is drastically larger. Even for the nearly optimal ω the numbers are on average 10 resp. 14 times larger, as can be seen in Tab. 2.

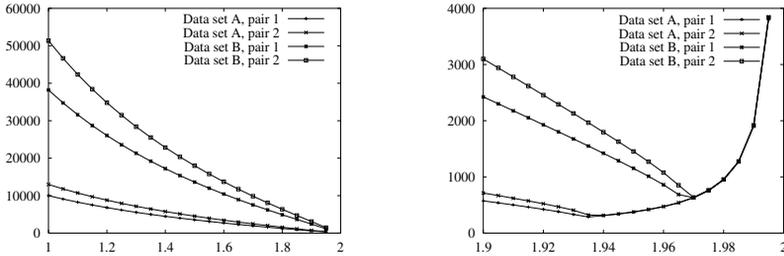


Fig. 2. Number of SOR iterations necessary to satisfy stopping criterion for the singular problem depending on choice of relaxation parameter ω .

The choice of ω appears to be not as important for the PCG method, as can be seen in Fig. 3. The number of iterations varies over smaller intervals and the valleys around the optimal value are more flat, thus choosing a reasonable ω is easier. The optimal values lie in the vicinity of 1.65 and 1.75 for the singular problem and 1.65 and 1.8 for the regular one. Note that again the number of iterations for the singular problem is smaller than for the regular case.

In the AMG algorithm there is quite a number of parameters and algorithmic variants, that have to be fixed, before the method can be applied. These include e. g. cycle type, number of smoothing steps, and so on. We have used a V-cycle with one pre- and one post-smoothing step. The relaxation scheme was a hybrid Gauß-Seidel / Jacobi method. Construction of the grid hierarchy was performed with Ruge-Stüben coarsening. Besides this, we left all other parameters at their default values and investigated the influence of the threshold value α , as shown in Fig. 4 for the singular case. The regular case is not shown, since the results are very similar.

We note two interesting facts. The first is, that the best convergence is achieved for comparatively small values of α . The second fact is that there is a sharp increase in the number of cycles in the interval $\alpha \in [0.50, 0.55]$. This is especially pronounced in the case of data set B. One can show [9] that for $\alpha > 0.55$ we get bad interpolation operators, since skull voxels on the fine grid are only interpolated from scalp / skull voxels on the next coarser grid. Our experiments also indicate that performance will decrease again, when α becomes too small. So, for all further experiments we have settled with $\alpha = 0.05$.

After having determined a set of reasonably optimal parameters for the different methods, we tested the number of iterations needed to satisfy the stopping criterion for all 26 electrode pairs. Table 2 summarises the mean values of all

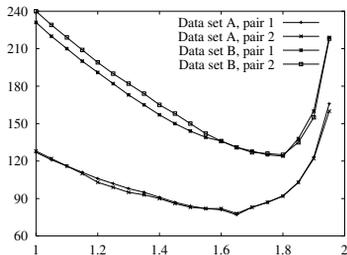


Fig. 3. Number of PCG iterations necessary to satisfy stopping criterion for the singular problem depending on choice of relaxation parameter ω .

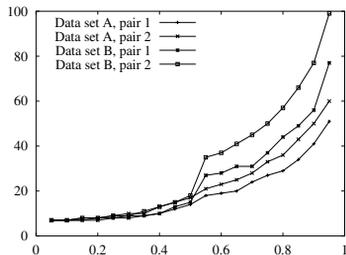


Fig. 4. Number of AMG cycles necessary to satisfy stopping criterion depending on choice of threshold parameter α .

pairs. To be able to better compare these values we have determined for each method an estimate of the total amount of arithmetic work involved. This estimate is given in Tab. 2 as a percentage of the amount of work for the SOR method in the semi-definite case.

We see that the CG method performs worst, due to a bad convergence rate. This is considerably improved by preconditioning. The number of iterations drops to roughly 18% of the unpreconditioned case. Due to the higher costs per iteration step this is still about two thirds of the reference case in the most favourable situation. The best performance is achieved by the AMG approach, which also shows the typical feature of a multigrid method, namely that the number of cycles remains constant, independent of the fineness of the discretisation.

Table 2. Iteration counts for different methods, matrix is singular (1) or regular (2)

Number of Iterations						
Method	Data set					
	A			B		
	Mean	% SOR	Std. Dev.	Mean	% SOR	Std. Dev.
SOR (1)	303.5	100	9.6	634.0	100	2.2
SOR (2)	2957.0	974	34.1	8839.0	1400	635.2
CG (1)	407.5	179	9.3	740.7	156	12.5
CG (2)	533.2	234	6.5	986.8	208	9.7
PCG(SSOR) (1)	80.3	88	1.3	126.7	67	1.8
PCG(SSOR) (2)	106.0	116	1.2	165.8	87	1.6
BoomerAMG (1)	7.0	16	0.0	7.0	7	0.0
BoomerAMG (2)	7.0	16	0.0	7.0	7	0.0

Besides convergence rates and amount of work, the property that is of primary interest to the user coming from the application side is run time. Factors determining the latter involve the numerical characteristics of an algorithm as well as its suitability for modern computer architectures and its implementation. We have tested the run times of our four algorithms on three different architectures, a 700MHz AMD Athlon, a 500MHz Alpha A21264 and a 1500MHz Pentium IV. More details of the machines can be found in [9]. In all cases, except for AMG, we compare the grid-based with the matrix-based implementation.

In Fig. 5 we present the user times for the complete problem. In each case we have measured times for 10 program runs and taken the mean value. Standard deviation was always less than 2%. The times consist of the time spent in the setup phase and for the solution of 26 forward problems. The setup times are always negligible except for the AMG case, where the construction of the grid hierarchy adds considerably to the total costs. This varies between 5 and 11% of the total time depending on architecture and problem size.

Concerning the measurements for BoomerAMG we should note that the hypr library, of which it is a part, was developed for solving large, sparse systems of linear equations on massively parallel computers. Thus it is not specially tuned for the sequential environment in which we used it. We also want to point out, that, although we have taken into account performance issues in the implementation of the algorithms, there are still numerous possibilities one could test for a further optimisation of the code. For some ideas on this, see e. g. [4].

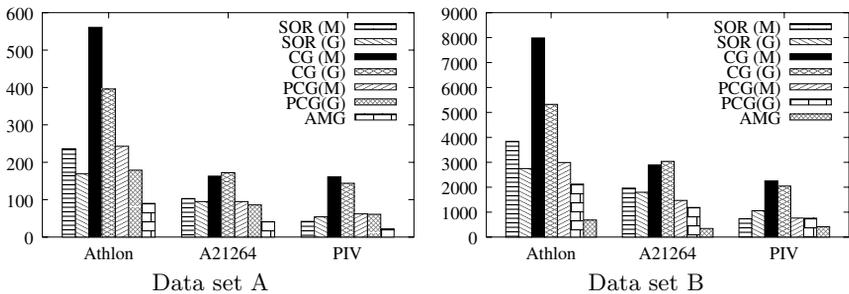


Fig. 5. Run times (user time in seconds) for the complete problem. This includes the setup phase and the solution of 26 singular forward problems.

6 Conclusions

Summing up the contents of this paper, we see three major points. The first one is, that the solution of inverse EEG problems can be sped up considerably by the application of multigrid methods for the solution of the forward problems involved. These appear to offer a much better performance than the Krylov subspace methods typically employed in this context. It remains to be seen

however, if this also holds for more sophisticated preconditioners like e.g. ILU, see [6].

The second aspect is that it seems pointless to transform the singular forward problem into an equivalent regular one utilising the approach described in Sect. 3. In our experiments the regularisation of the problem could not improve convergence of the tested methods and, in many cases, even led to a worse behaviour. However, this does not rule out the possibility that other regularisation approaches may lead to better performance.

The question, finally, whether a matrix- or a grid-based implementation of the methods is preferable in view of runtime behaviour, remains open. Our experiments in this context yielded results that vary from architecture to architecture and are ambiguous in themselves even on the same machine.

References

1. Berman, A., Plemmons, R.: *Nonnegative Matrices in the Mathematical Sciences*, SIAM, *Classics in Applied Mathematics* **9** (1994)
2. Buchner, H., Knoll, G., Fuchs, M., Rienäcker, A., Beckmann, R., Wagner, M., Silny, J., Pesch, J.: Inverse localization of electric dipole current sources in finite element models of the human head, *Electroencephalography and clinical Neurophysiology* **102** (1997) 267-278
3. Delbeke, J., Oostendorp, T.: The Conductivity of the Human Skull in vivo and in vitro, *Serving Humanity, Advancing Technology, Proceedings of The First Joint BMES/EMBS Conference* (1999) 456
4. Douglas, C.C., Hu, J., Kowarschik, M., Råde, U., Weiß, C.: Cache Optimization for Structured and Unstructured Grid Multigrid, *Electronic Transactions on Numerical Analysis* **10** (2000) 21-40
5. Ermer, J., Mosher, J., Baillet, S., Leahy R.: Rapidly re-computable EEG forward models for realistic head shapes, *Biomag 2000, Proc. of the 12th Internat. Conf. on Biomagnetism*, eds.: Nenonen, J., Ilmoniemi, R.J., Katila T. (2001)
6. Hackbusch, W.: *Iterative solution of large sparse systems of equations*, Springer (1994)
7. Henson, V.E., Meier Yang, U.: BoomerAMG: a Parallel Algebraic Multigrid Solver and Preconditioner, *Applied Numerical Mathematics* (in press), also available as Technical Report UCRL-JC-141495, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory (2000)
8. Jasper, H.: Report of committee on methods of clinical exam in EEG, *Electroencephalography and Clinical Neurophysiology* **10** (1958) 370-375
9. Mohr, M.: Comparison of Solvers for a Bioelectric Field Problem, Technical Report 01-2, Lehrstuhl für Informatik 10 (Systemsimulation), Friedrich-Alexander-Universität Erlangen-Nürnberg (2001)
10. Tuch, D., Wedeen, V.J., Dale, A., George, J., Belliveau, W.: Conductivity tensor mapping of the human brain using diffusion tensor MRI, *Proc. Natl. Acad. Sci. USA* **98(20)** (2001) 11697-11701
11. Vanrumste, B.: EEG dipole source analysis in a realistic head model, PhD thesis, Faculteit Toegepaste Wetenschappen, Universiteit Gent (2001)
12. Weinstein, D., Zhukov, L., Johnson, C.: Lead-field Bases for Electroencephalography Source Imaging, *Annals of Biomedical Engineering* **28(9)** (2000) 1059-1065