Comparison of Temporal Dimensionality Reduction Methods for Constrained Inverse in Cardiac Electrical Imaging

⋆Jaume Coll-Font[1], ⋆Danila Potyagayo[2], Walther HW Schulze[2], Olaf Doessel[2], Dana H Brooks[1]

⋆ These authors contributed equally

[1] B-spiral group in the ECE dept. at Northeastern University, Boston, USA
[2] Karlsruhe Institute of Technology, Karlsruhe, Germany

Abstract

Cardiac electrical imaging, that is, reconstructing cardiac electrical activity from body surface measurements, is a technology with great potential. However, ill-posedness of this problem hinders its routine usage in clinical environment and continues to motivate the search for improvements on current methods. Messnarz et al. introduced an algorithm that constraints the reconstructed transmembrane potential (TMP) to be non-decreasing over time during QRS-complex. This physiologically meaningful constraint reduces the solution space of the problem and regularizes the solution. However, this approach is computationally extensive and can become prohibitive as spatial and temporal resolution of the problem increase. Here we compare three distinct options to reduce the computational load: downsampling the measurements in time, downsampling the measurements after filtering with an algorithm based on principal component analysis and non-linearly interpolating the potentials with a spline-based method. The data used were simulated TMPs that were forward propagated to the body surface in a densely sampled geometry. The resulting body surface potential simulations were corrupted with noise and the inverse computed using a much coarser mesh to take geometry errors into account. The results indicate that reducing the dimension of the signal in time does not reduce the quality of the solutions obtained, while the computational requirements decrease considerably, especially for the spline method.

1. Introduction

Electrocardiographic Imaging (ECGI) is a technology whose objective is to non-invasively image the electrical function of the heart. That is, it reconstructs the electrical potentials on the heart from the body surface potential (BSP) measurements and a mathematical model of the volume conductor that relates them. Its purpose is to detect abnormal behavior of the myocardial tissue, such as points of re-entry or ectopic beats, and has great potential as a tool for pre-interventional planning in ablation procedures. However, there are still some challenges to overcome. One of them is the ill-posedness of the inverse problem, which makes the solutions very sensitive to noise and requires the inclusion of prior knowledge to stabilize them. All ECGI methods introduce this prior in one way or another, so it is perhaps best to impose physiologically meaningful constraints. In this regard, Messnarz *et al.* introduced a method to solve for transmembrane potentials (TMP) that imposes a non-decreasing constraint on the solutions [1]. The estimated potentials obtained with this method are physiologically realistic and provide very good accuracy. However, the size of the problem grows quadratically with the number of unknowns and soon becomes computationally infeasible due to memory and time requirements. Therefore, there is a need to reduce the dimensionality of the problem for it to become practical. The dimension of the problem can be reduced in two ways: space and time. Reduction in the latter is possible due to the quasi static assumption, that states that the BSP at a time point are only affected by the instantaneous potentials on the heart. Thus, it is possible to reduce the number of temporal samples without affecting their inverse solutions.

In this work, we compare three dimensionality reduction strategies in time: straight downsampling, linear interpolation based on the SVD and a non-linear interpolation with splines. We use them in conjunction to the inverse method by Messnarz *et al.* and analyze their performance in terms of the solution quality and computational time for multiple sizes of the temporal space.

2. Methods

All the methods compared in this work can be decomposed in a three block pipeline: a dimensionality reduction step, an inverse solver and a temporal reconstruction of the inverse solutions.

The common block in all methods is the inverse method, which is the method that Messnarz *et al.* introduced in

[1]. This method solves the least squares problem in with a non-decreasing constraint in time on the TMP:

$$\min_{x(t)} \|y(t) - Ax(t)\|_2^2 + \lambda \|x(t)\|_2^2$$

$$\text{s.t.} \quad x(t+1) \geq x(t)$$
$$x(1) \geq -85 \text{ mV} \quad (1)$$
$$x(T) \leq 20 \text{ mV}.$$

Here $x(t)$ and $y(t)$ are the TMP and the ECG measurements at times $t \in [1, \ldots, T]$, and A is the forward matrix that models the volume conductor. This corresponds to a quadratic problem with a linear constraint which can be solved using CVX, a freely available software to solve convex optimization problems [2, 3].

The first and last blocks in the processing pipelines are determined by the temporal dimensionality method: temporal downsampling, principal component analysis (PCA) filtering with downsampling and spline interpolation.

Temporal downsampling:

This method consists of temporal decimation of the signal. We evaluated this simple approach to reduce the number of samples in time with downsampling rates (R) of 1, $1/2$, $1/5$ and $1/10$. This method yields inverse solutions whose size is $M * T * R$, for measurements with T time samples and spatial resolution on the heart of M nodes. After computing the inverse with the decimated samples, we reconstruct the complete temporal sequence with linear interpolation of the missing time instances.

PCA:

This approach is similar to the method from Huiskamp and Greensite to solve the inverse problem [4]. It consists in a linear deconvolution in time of the signal using its principal components – *i.e.* the eigenvectors of its covariance matrix. Ideally, we would have solved the inverse problem on a subspace defined by the most important principal components as in [4]. However, this approach did not lead to satisfactory results. We believe, that this was due to a basic assumption of the method which expects the signal to be separable in space and time, which does not hold for the TMP. Instead, we used the projection on the principal components as a filtering step similar to principal component analysis (PCA) filtering [5]. For all simulation cases, we set the truncation rank to be 40, which was determined to deliver the best results when evaluating with the considered setups. After filtering, we downsampled the measurements in time as previously described ($R = 1, 1/2, 1/5$ and $1/10$). As with the pure donwsampling approach, to reconstruct the complete temporal sequence we linearly interpolate the missing time instances. The objective of this pre-filtering step is twofold: one, improve the inverse solutions with the reduction of noise and, two, reduce the computational time of the inverse method with faster convergence to the solution.
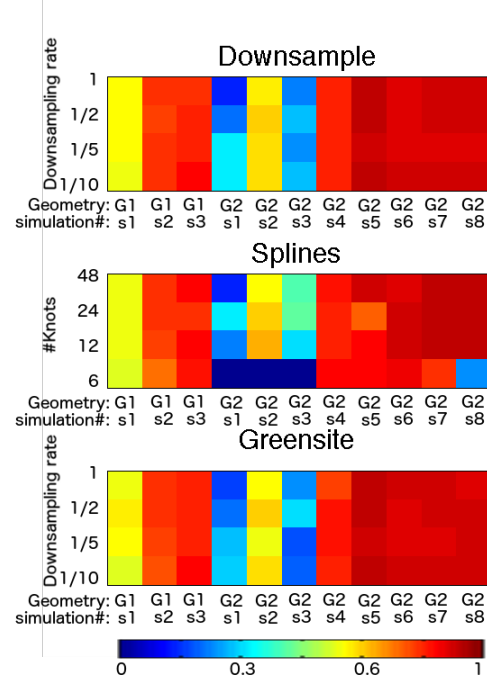

Figure 1: Correlation between true and estimated activation times. Each image contains the results of a dimensionality reduction method. Within each image columns show results of a single simulation (for both geometries) and rows show the different downsampling rates or number of knot points.

Spline interpolation:

This is a non-linear dimensionality reduction that uses spline interpolation to characterize the evolution of the signal in time [6]. This approach has the advantage that the parameters that are used in the interpolation (the knot points) are themselves potential distributions on the heart and thus they can be used as an input of an inverse method. Moreover, since the ordering of these knot points preserves the temporal ordering of the signal, the non-decreasing constraint remains a valid assumption in the inverse method. We evaluated this non-linear interpolation for 6, 12, 24 and 48 knot points. Finally, since the inverse solutions obtained are the equivalent knot points on the heart, we reconstruct the complete temporal sequence with the interpolation learned during the spline fitting on the measurements.

3. Experiments and Results

We tested each inverse pipeline described in section 2 on a dataset composed of two high resolution realistic torso geometries with multiple simulations of heart activations. Both geometries contained skeletal muscles, lungs, ventri-

cles with intracavitary blood, spleen, stomach, kidneys and liver as inhomogeneities.

For these geometries, we simulated ventricular ectopic activation sequences at different heart locations (3 and 8 locations per geometry respectively) with a rule-based cellular automaton. Then, we synthetically forward propagated the TMP of each simulation using a forward matrix obtained from the high resolution geometries in order to generate realistic BSP. The calculations of BSP were done with the finite-element method (FEM). For the inverse calculations we coarsened the meshes and added AWGN corresponding to SNR= 30dB to the BSP. For the first geometry we adopted the transfer matrix (which was also respectively coarsened). For the second geometry we calculated it with boundary element method (BEM) assuming homogeneous volume conductor.

For every simulation setup, the problem was solved using all three dimensionality reduction methods. From each solution we computed the activation times with a method that calculates the maximum dv/dt of the potentials weighted with the spatial gradient. Afterwards, we smoothed the resulting activation times with the method described in [6] to reduce errors created due to overly smooth TMP solutions. We determined the position of the first activation as the node with earliest estimated activation time.

The tests of each downsampling method were done in Matlab (12-Core Intel Xeon E5, 2.7 GHz processor and 64 GB memory).

Finally, we validated the inverse solutions in terms of the correlation between true and estimated activation times and the error in localization of the first activation on the heart. To evaluate the computational load, we stored the time it took to compute each inverse solution.

The results of the different methods are shown in the following figures. Figure 1 shows the correlation of the activation times and Figure 2 shows the localization error of the point of first activation on the heart. The computation time of the inverse methods is shown in Figure 3.

4. Discussion

Overall, the inverse algorithm by Messnarz *et al.* performs well in reconstructing the TMP on the heart. And, as the high correlation in activation times shows, the temporal pattern of most solutions captures the true activation on the heart. With the exception of three simulations – one in the first geometry and two in the second – where the correlation in activation times is much below 1. These correspond to simulations with an ectopic beat around the septal area, which has always been shown to be a challenge for the ECGI methods.

The error in localization of the first activation shows more varied results. In this case, there seem to be three
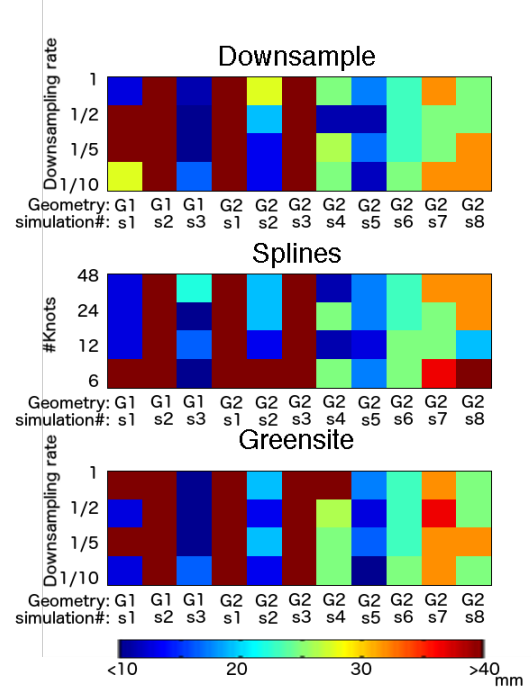

Figure 2: Localization error of first activation on the heart. Each image contains the results of a dimensionality reduction method. Within each image columns show results of a single simulation (for both geometries) and rows show the different downsampling rates or number of knot points.

type of results: simulations that do well (err $\leq$ 20mm), simulations that have somewhat increased error ($20 \leq$err$\leq$ 35 mm) and simulations that fail to detect the first activation (err$\geq$ 35 mm). The high variability of the results, is probably linked to the non-linearity introduced with the $\min$ operator in the search of the first activation, which is sensitive to small errors in the estimation of the activation times. These small post-processing errors are most likely to have a bigger impact for the second geometry, which is coarser than the first one. So, it explain the increased localization error in the group with $20 \leq$err$\leq$ 35 mm.

Comparing across temporal downsampling methods, it seems that the three approaches provide similar results. Correlation between true and estimated activation times is high for results of simulations starting on the free walls and worse for simulations starting around the septum. Although the localization error shows quite some variability, it does not appear to have any pattern of improvement for either method.

The factor that does make a difference across different downsampling methods is computational time, shown in Figure 3. Since this inverse method is a quadratic program with linear constraints, its computational burden increases exponentially with the number of unknowns and,
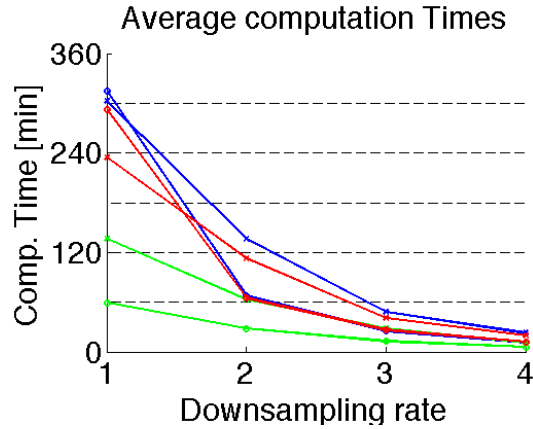
Figure 3: Average computational time for all downsampling methods and both geometries. Blue lines show results for regular downsampling, red lines for PCA filtering and green lines for the splines interpolation. Lines with 'x' show results for geometry 1 and 'o' for geometry 2.

with it, the computational time. Therefore, it is to be expected that the method with smallest temporal dimension, the spline interpolation, requires considerably less computation time than the other downsampled methods. That is especially the case when using 12 knot points, which takes less computation time than the other dimensionality reduction methods and slightly outperforms them. Another noticeable difference is between downsampling and PCA-filtering methods. Even though for those methods the number of unknowns is the same for every rate of reduction, PCA-filtering computes the solutions somewhat faster and attains a speedup of an hour for the most demanding inverses.

5.   Conclusions

The inverse method by Messnarz *et al.* provides good solutions for the TMP on the heart. However, it is computationally demanding and requires of dimensionality reduction techniques for it to be practical.

We have evaluated three temporal dimensionality reduction methods in terms of their capacity to reduce this computational burden while preserving the quality of the solutions. In this study we have not found any meaningful difference among their solutions. For most of the simulations the activation times have high correlation with the true solution and the localization of the first activation does not show particular improvement in any particular method. However, the computation times do indicate that the spline interpolation is much faster than the rest and that in case of using a decimation method, it is preferable to filter the signals to reduce computational load.

In this study we have also observed high variability in

localization of the first activation, which is sensitive to small errors in the estimation of activation times. It remains future work to study more robust methods to detect the location of the initial activation on the heart.

Acknowledgements

References

[1] Messnarz B, Tilg B, Modre R, Fischer G, Hanser F. A new spatiotemporal regularization approach for reconstruction of cardiac transmembrane potential patterns February 2004;51(2):273–281.

[2] Grant M, Boyd S. CVX: Matlab software for disciplined convex programming, version 2.1. `http://cvxr.com/cvx`, March 2014.

[3] Grant M, Boyd S. Graph implementations for nonsmooth convex programs. In Blondel V, Boyd S, Kimura H (eds.), Recent Advances in Learning and Control, Lecture Notes in Control and Information Sciences. Springer-Verlag Limited, 2008; 95–110. `http://stanford.edu/~boyd/graph_dcp.html`.

[4] Huiskamp GJ, Greensite F. A New Method for Myocardial Activation Imaging 1997;44:433–446.

[5] Bishop CM. Pattern Recognition and Machine Learning (Information Science and Statistics). Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006. ISBN 0387310738.

[6] Erem B, Coll-Font J, Orellana RM, St'Ovicek P, Brooks DH. Using transmural regularization and dynamic modeling for noninvasive cardiac potential imaging of endocardial pacing with imprecise thoracic geometry. IEEE Transactions on Medical Imaging 2014;33:726–738. ISSN 02780062.

Address for correspondence:

Jaume Coll-Font - jcollfont@ece.neu.edu,
B-spiral group in the ECE dept. at Northeastern University
360 Huntington Ave., Boston, Massachusetts 02115