34th Annual International Conference of the IEEE EMBS
San Diego, California USA, 28 August - 1 September, 2012

Probabilistic source separation for robust electrocardiography

R. Vullings

Abstract— **Blind source separation (BSS) techniques are widely used to extract signals of interest from a mixture with other signals. These methods, however, typically lack possibilities to incorporate any prior knowledge on the mixing of the source signals. Particularly for electrocardiographic signals, knowledge on the mixing is available based on the origin and propagation properties of these signals. In this paper, a novel source separation method is developed that combines the strengths and accuracy of BSS techniques with the robustness of an underlying physiological model of the electrocardiogram (ECG). The method is developed within a probabilistic framework and yields an iterative convergence of the separation matrix towards a maximum a posteriori estimation, where in each iteration the latest estimate of the separation matrix is corrected towards the physiological model. The method is evaluated by comparing its performance to that of FastICA on both simulated and real multi-channel ECG recordings, demonstrating that the developed method outperforms FastICA in terms of extracting the ECG source signals.**

I. Introduction

Blind source separation (BSS) techniques like independent component analysis (ICA) are widely used to extract signals of interest from a mixture of recorded signals [1], [2], [3]. The goal of BSS techniques is to 'unmix' the recorded signals as to obtain the original source signals and enable further processing or analysis of any particular source signal. BSS techniques come in a variety of implementations, each targeting specific source signal properties. Their common denominator is that the BSS techniques assume statistical independence of the source signals and assume virtually no specific knowledge on the mixing of the source signals: hence their description as being 'blind'.

In some problems, however, specific knowledge on the mixing of the source signals is available and not exploiting this knowledge impairs the source separation's accuracy and robustness [4]. In this paper, we develop a source separation technique aimed at extracting electrocardiographic signals from a mixture of electrophysiological signals and interferences. This technique is derived within a probabilistic framework and incorporates prior knowledge on the mixing matrix in terms of an electrophysiological model of the electrocardiogram (ECG).

Due to the prior knowledge, the developed source separation technique has the potential to outperform generic BSS techniques in terms of robustness during situations of low signal quality. In such situations, BSS techniques have been reported to provide independent source signals that correspond to any kind of interference, but not to an actual electrocardiographic signal [5]. With respect to merely unmixing recorded signals based on the ECG model, the developed source separation exploits the powerful mathematical aspects of BSS techniques and is hence expected to estimate the source signals more accurately.

The derivation and evaluation of the proposed source separation techniques is discussed in the following way. In Sec. II the ECG model is discussed and Sec. III focuses on the derivation of the source separation algorithm. In Sec. IV the source separation technique is evaluated and in Sec. V it is discussed.

II. ECG model

The contracting of the cardiac muscles during the beating of the heart is initiated by the propagation of action potentials across the myocardium. In a first-order approximation, the action potentials at any point in time can be described by a single dipole vector with stationary origin [6]. This vector varies in both amplitude and orientation during the beating of the heart.

The varying dipole causes circular currents to spread through the conductive tissues surrounding the heart, all the way to the cutaneous surface. Here, the skin impedance causes potential differences that fluctuate over time: the ECG. In this simplified model of the electrical activity of the heart, the potential difference between two separate locations on the skin can be regarded as the projection of the dipole vector onto the geometrical vector that describes the separate locations with respect to one another. This concept is illustrated in Fig. 1

When performing a bipolar, multi-lead ECG measurement with N recorded signals $\mathbf{V}(t)$ and relative electrode positions

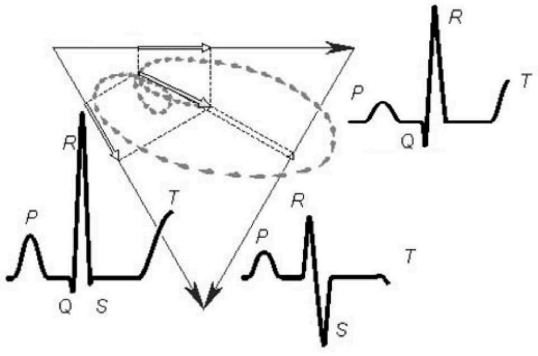

Fig. 1. The ECG is determined by the projection of the heart's dipole vector onto geometrical electrode vectors. The amplitude of the ECG is the length of the projected dipole vector.

R. Vullings is with the Faculty of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands r.vullings@tue.nl

D, the relation between the 3-dimensional dipole vector $\mathbf{S}(t)$ and $\mathbf{V}(t)$ can thus be described as

$$\mathbf{V}(t) = \mathbf{D}\mathbf{S}(t). \tag{1}$$

Here, $\mathbf{V}(t)$ is a $[N\times T]$ matrix, $\mathbf{D}$ is a $[N\times 3]$ matrix, and $\mathbf{S}(t)$ is a $[3\times T]$ matrix, with T the length of each signal.

III. Probabilistic source separation

The expression in Eq (1) resembles the problem of source separation. When we assume a recording of N signals, each comprising an unknown mixture of N source signals, the problem of source separation is to 'unmix' the recorded signals $\mathbf{x}(t)$ as to find the source signals $\mathbf{s}(t)$:

$$\mathbf{x}(t) = \mathbf{A}\mathbf{s}(t), \tag{2}$$

where $\mathbf{A}$ is the mixing matrix.

The probability that the source model of Eq. (2) is correct can be written in terms of the likelihood of the data and prior probabilities on the source signals and mixing matrix [4]

$$p(\mathbf{A},\mathbf{s}|\mathbf{x}) \propto p(\mathbf{x}|\mathbf{A},\mathbf{s})\, p(\mathbf{A})\, p(\mathbf{s}), \tag{3}$$

where $\mathbf{A}$ and $\mathbf{s}$ are assumed statistically independent, conform the assumption that the properties of signal propagation do not depend on the source signals and their magnitudes.

As mentioned previously, the goal of source separation is to determine the source signals $\mathbf{s}(t)$. Confer Eq. (2), inference of $\mathbf{A}$ also determines $\mathbf{s}(t)$. In addition, due to the typically smaller dimensions of $\mathbf{A}$ with respect to $\mathbf{s}(t)$, inference of $\mathbf{A}$ is computationally more efficient. Hence, the problem of estimating the source signals in source separation methods is typically translated to the problem of estimating the mixing matrix.

In the context of wanting to estimate $\mathbf{A}$, the sources signals $\mathbf{s}(t)$ can be treated a nuisance parameter. Marginalizing over $\mathbf{s}$ gives

$$p(\mathbf{A}|\mathbf{x}) \propto p(\mathbf{A}) \int p(\mathbf{x}|\mathbf{A},\mathbf{s})\, p(\mathbf{s})\, \mathrm{d}\mathbf{s} \tag{4}$$

A. Blind source separation

When we would assume no knowledge on the mixing matrix $\mathbf{A}$, its prior probability distribution $p(\mathbf{A})$ in Eq. (4) would be uniform. When we would furthermore assume that the mixing is noiseless, linear, and instantaneous [4], cf. Eq. (2) and that the source signals are statistically independent, the posterior probability distribution for $\mathbf{A}$ can be written as [2]

$$p(\mathbf{A}|\mathbf{x}) \propto \int \prod_i \delta\left(x_i - \sum_k A_{ik} s_k\right) \prod_j p(s_j)\, \mathrm{d}\mathbf{s}. \tag{5}$$

Taking the logarithm on either side and defining the separation matrix $\mathbf{W}$ as the inverse of the mixing matrix $\mathbf{A}$, Eq. (5) can be written as

$$\log p(\mathbf{A}|\mathbf{x}) = \log\det\mathbf{W} + \sum_j \log p\left(\sum_k W_{jk} x_k\right) + C, \tag{6}$$

where C is a constant.

B. Incorporating prior knowledge

In the separation of ECG signals, the mixing matrix $\mathbf{A}$ is not completely unknown. The model of Sec. II provides some information on how the mixing occurs. Lets redefine the matrix $\mathbf{D}$ of Eq. (2) into a $[N\times N]$ matrix for which the first 3 columns contain the geometrical positions of the electrodes in 3-D space. The other columns of $\mathbf{D}$ can be set to unity. Based on the model in Sec. II, the elements of mixing matrix $\mathbf{A}$ can be modeled as:

$$A_{ij} = \alpha_j D_{ij} + \eta_{ij}. \tag{7}$$

Here, α is a scaling that can vary between recorded signals and models differences in signal attenuation during propagation. Any imperfections in this model are captured by the (Gaussian) noise η_{ij} which is assumed to have zero-mean and variance σ_{ij}^2.

To keep the source separation analytically tractable, a few more assumptions are made: 1.) the elements of the mixing matrix are assumed independent: $p(\mathbf{A}) = \prod_{ij} p(A_{ij})$ [4]; 2.) the scaling α and geometrical matrix $\mathbf{D}$ are statistically independent; 3.) the noise variance σ^2 is user-defined, e.g. based on empirical results; 4.) our belief that $\mathbf{D}$ accurately describes the electrode positions is reflected by assigning a delta-function to $p(D_{ij})$. With these assumptions, the prior probability distribution $p(\mathbf{A})$ can be written as

$$\begin{aligned} p(\mathbf{A}) &\propto \prod_{ij} \iint p\left(A_{ij}\,|\,\alpha_j, D_{ij}\right) p(\alpha_j)\, p(D_{ij})\, \mathrm{d}\alpha_j \mathrm{d}D_{ij} \\ &\propto \prod_{ij} \iint \exp\left[-\frac{1}{2\sigma_{ij}^2}\left(A_{ij} - \alpha_j D_{ij}\right)^2\right] p(\alpha_j)\, \mathrm{d}\alpha_j. \end{aligned}$$

By expressing our ignorance on the scaling parameters α as a uniform distribution with minimum and maximum values a_1 and a_2, respectively [4]:

$$p(\alpha_j) = \begin{cases} (a_2 - a_1)^{-1} & a_1 \le \alpha_j \le a_2 \\ 0 & \text{otherwise} \end{cases}, \tag{8}$$

the prior probability becomes:

$$p(\mathbf{A}) = \prod_{ij} \frac{1}{2(a_2 - a_1) D_{ij}} \left\{\mathrm{erf}\left(f_{ij}(a_1)\right) - \mathrm{erf}\left(f_{ij}(a_2)\right)\right\}, \tag{9}$$

with $f_{ij}(a) = \frac{1}{\sqrt{2}\sigma_{ij}}(A_{ij} - aD_{ij})$ and $\mathrm{erf}(f)$ the error function: $\mathrm{erf}(f) = \frac{2}{\sqrt{\pi}} \int_0^f \exp[-t^2]\mathrm{d}t$.

C. Towards a solution

The goal of the proposed source separation is to maximize the posterior probability distribution in Eq. (6) with respect to the choice for $\mathbf{W}$. Including prior knowledge on the mixing matrix $\mathbf{A}$ [2] and taking the derivative with respect to W_{ij}, Eq. (6), the maximum-a-posteriori estimate for $\mathbf{W}$ follows from:

$$\frac{\partial}{\partial W_{ij}} \left\{\log\det W + \sum_l \log p\left(\sum_k W_{lk} x_k\right) + \sum_{kl} \log p\left(A_{kl}\right)\right\} = 0$$

Introducing $u_i = \sum_j W_{ij} x_j$, this expression can be further solved:

$$A_{ji} + \frac{\partial u_m}{\partial W_{ij}} \frac{\partial}{\partial u_m} \sum_l p(u_l) + \frac{\partial A_{mn}}{\partial W_{ij}} \frac{\partial}{\partial A_{mn}} \sum_{kl} \log p(A_{kl}) = 0$$

and further to:

$$\frac{\partial}{\partial \mathbf{W}} \log p(\mathbf{A}|\mathbf{x}) = \mathbf{A}^T + \left(\frac{p_i'(u_i)}{p_i(u_i)} \right) \mathbf{x}^T - \mathbf{A}^T \mathbf{M} \mathbf{A}^T \qquad (10)$$

with

$$\mathbf{M} = \frac{\partial}{\partial A_{mn}} \sum_{kl} \log p(A_{kl}) = \sqrt{\frac{2}{\pi \sigma_{mn}^2}} \frac{\exp\left[f^2(a_1)\right] - \exp\left[f^2(a_2)\right]}{\operatorname{erf}(f(a_1)) - \operatorname{erf}(f(a_2))}.$$

By post-multiplying by $\mathbf{W}^T\mathbf{W}$ this expression can be made covariant [7]. The optimum for $\mathbf{W}$ can subsequently be found using a stochastic gradient search with learning rate γ:

$$\mathbf{W}_{i+1} = \mathbf{W}_i + \gamma \left[\mathbf{W}_i + \left(\frac{p_i'(u_i)}{p_i(u_i)} \right) \mathbf{u}^T \mathbf{W}_i - \mathbf{A}_i^T \mathbf{M} \mathbf{W}_i \right]. \quad (11)$$

IV. Evaluation

To evaluate the method against an existing BSS technique, both the developed method and FastICA [3] are applied on the same sets of data. These data sets are generated in two ways: from simulations in which the source signals are known and from real recordings in which the signal to noise ratio of the ECG signals is low.

A. Simulated data

As starting point for the simulated data, three orthogonal leads of the ECG (known as the vectorcardiogram) are adopted from the MIT-BIH Physionet database [8]. A number of N supposed electrode positions for the simulated data are chosen by picking the positions from Gaussian distributions for which the means are randomly chosen. To mimic the uncertainty in electrode positions that is typically encountered in ECG recordings, the matrix $\mathbf{D}$ in Eq. (7) is filled with the means of the Gaussian distributions and not with the true (modeled) electrode positions (i.e. the positions that were picked from the Gaussian distribution). The set of recorded signals is generated by projecting the vectorcardiogram onto the vectors describing the electrode positions. In addition, each recorded signal is scaled by a randomly determined scalar and corrupted with additive noise and artificial powerline interference.

In Fig. 2, the three orthogonal ECG signals are shown. Also in figure, a simulated multi-channel ($N = 5$) ECG recording is shown. In Fig. 3 the results of the source separation are shown. In each diagram, the top graph shows the original input source signal. The middle graph shows the source estimated by the developed method. The bottom graph shows the source estimated by FastICA.

From Fig. 3 it can be seen that the overall estimation of all three sources by the developed method is more accurate than the estimation by FastICA. Quantitatively, the mean squared error between the original source signal and the estimated source signals is over 35% smaller for the developed method. Despite the ability of both methods to extract three ECG sources – although the source signal estimated by FastICA in Fig. 3(c) can be barely called an ECG signal – the separation of the T-wave in particular is mixed-up between the sources. Where in Figs. 3(a) and 3(b) the T-wave is overestimated, in Fig. 3(c) it is underestimated. An explanation for this mix-up can lie in the inaccuracy in the estimated electrode positions. With more accurately determined positions, the unmixing by the developed method would be more accurate. For the FastICA method, the total lack of knowledge on the source mixing yields an even more inaccurate unmixing of the source signals' T-waves.

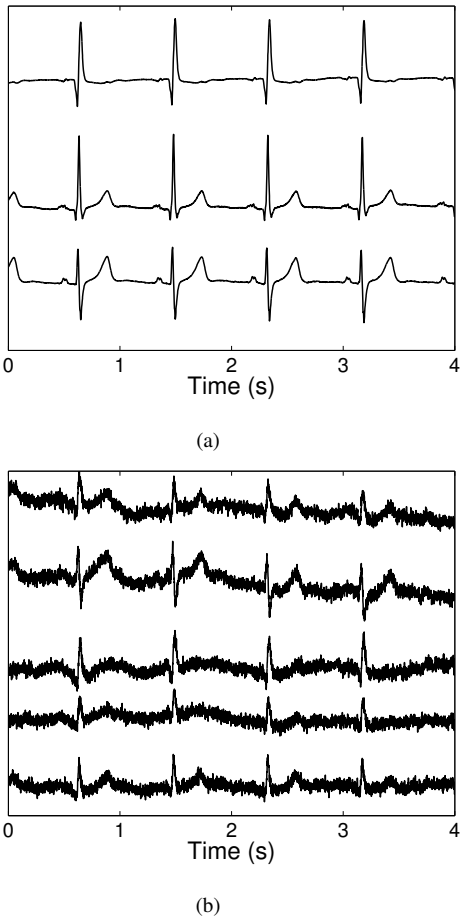

Fig. 2. Example of simulated recordings. In (a) the original sources used as input in the simulation are depicted. In (b) the generated ECG recordings are depicted for $N = 5$ ECG channels.

B. Real data

For the real data, a 45-minute, 8-channel recording performed on the abdominal surface of a pregnant woman is used. The gestational age of the foetus was 27+5 weeks and no complications were present. The mother had signed an informed consent. The ECG recordings were preprocessed by suppressing powerline interference using a fourth-order Butterworth notch filter centered around 50 Hz and suppressing the ECG of the mother using an adaptive template subtraction

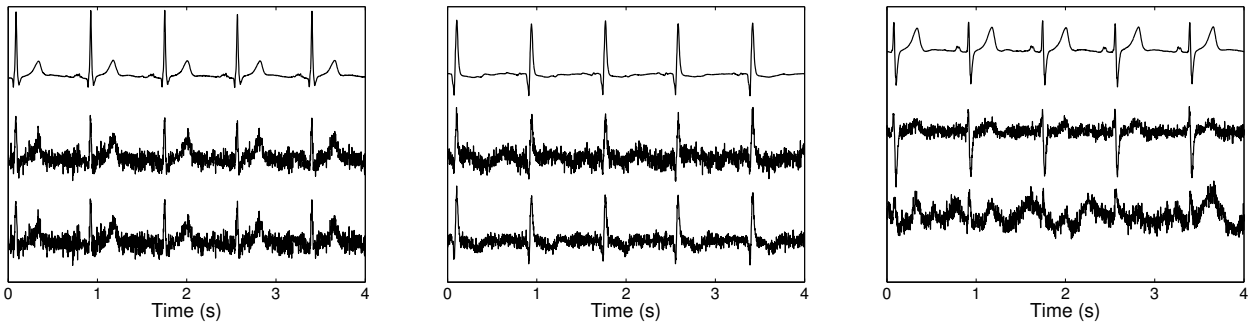

Fig. 3. The three ECG source signals (top graphs) depicted together with their estimates by the developed source separation (middle graph), and by FastICA (bottom graph).

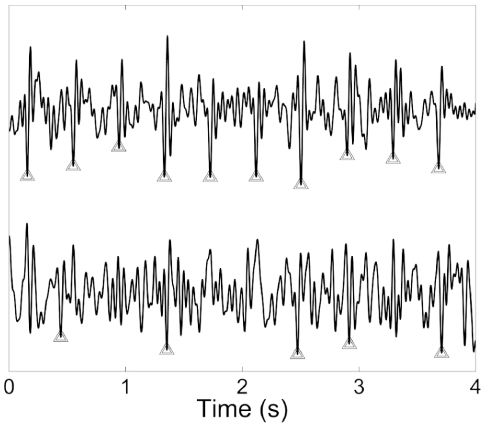

Fig. 4. The top graphs shows the fetal ECG source signals estimated with the developed source separation. The bottom graph shows the fetal ECG source signal estimated with FastICA. The triangles indicate detected fetal QRS complexes.

technique described in [9].

In Fig. 4 a fragment of 4 seconds long is shown for the fetal ECG source signals estimated by the developed method and by FastICA. Selection of the fetal ECG source signal was made based on visual inspection. It has to be noted here that in none of the source signals from FastICA a fetal ECG could be distinguished. To depict a source signal anyhow, the source signal was used that exhibited the largest correlation with the source signal estimated by the developed method.

It has to be noted here that, due to the presence of various other significant sources/interferences, only one fetal ECG source signal could be extracted from the abdominal ECG recording.

V. Discussion & Conclusions

In this paper a source separation technique for ECG signals was developed that exploits prior knowledge on the signal mixing. When critically examining the presented iterative solution to the source separation problem, it shows that the developed technique is similar to the Bell-Sejnowski [2] ICA approach, but with a correction in every iteration. This correction pushes the separation matrix $\mathbf{W}$ towards the physiological model. The confidence in this physiological model, quantified in the variance σ^2, is hence an important factor in the degree of correction. Little confidence, expressed by a large σ^2 causes the method to act as BSS technique with little to none pushing towards the physiological model. More confidence, on the other hand, leads to a technique that is more governed by the physiological model. The choice for σ^2 should hence be made based on the expectation that the electrode positions are determined correctly and the recorded ECG data conforms to the ECG model of Sec. II. In this paper, σ_{ij}^2 was set to 50% of the mean variance of the recorded signals (i.e. the mean of the variances of the individual recorded signals) for all i and j.

With respect to FastICA, the developed method performs better in retrieving the ECG sources in simulated data. In addition, based on visual analysis the developed method also performs better for real ECG data. More extensive evaluation of the developed method is however required to conclusively state about its performance in cases of poorly determined electrode positions, even lower signal quality, etc. This extensive evaluation will be subject of further studies, as is the choice for the variance σ^2.

References

[1] P. Comon, "Independent component analysis, a new concept?," Signal Processing, vol. 36, 1994, pp. 287–314.

[2] A.J. Bell, T.J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," Neural Comp., vol. 7, 1995, pp. 1129–59.

[3] A. Hyvärinen, "Fast and robust fixed point algorithms for independent component analysis," IEEE Trans. Neural Networks, vol. 10(3), 1999, pp. 626–34.

[4] K.H. Knuth, "Bayesian source separation and localization," In: A. Mohammed-Djafari (ed.), SPIE'98 Proceedings: Bayesian inference for inverse problems, SPIE vol. 3459, San Diego, 1998, pp. 147–58.

[5] S. Harmeling, F. Meinecke, K.-R. Müller, "Analysing ICA components by injecting noise," 4th Int. Symp. on Independent Component Analysis and Blind Signal Separation, Apr. 2003, pp. 149–54.

[6] E. Frank, "General theory on heart-vector projection," Circ. Res., vol. 2, 1954, pp. 258–70.

[7] S. Amari, "Natural gradient works efficiently in learning," Neural Comp., vol. 10, 1998, pp. 251–76.

[8] A.L. Goldberger, L.A.N. Amaral, L. Glass et al. "Physiobank, PhysioToolkit, and Physionet: components of a new research resource for complex physiologic signals," Circulation, vol. 101(23), 2000, pp. e215–20.

[9] R. Vullings, C.H.L. Peters, R.J. Sluijter et al. "Dynamic segmentation and linear prediction for maternal ECG removal in antenatal abdominal recordings," Physiol. Meas., vol. 30, 2009, pp. 291–307.