MODEL-BASED NOISE REDUCTION FOR SINGLE TRIAL EVOKED POTENTIALS

Arthur Flexer[1], Herbert Bauer[2], Claus Lamm[2], Georg Dorffner[1]
[1] The Austrian Research Institute for Artificial Intelligence
Schottengasse 3, A-1010 Vienna, Austria
[2] Department of Psychology, University of Vienna
Liebiggasse 5, A-1010 Vienna, Austria

Abstract.

Two model-based techniques, Gaussian Mixture Models with integrated noise component and Principal Component Analysis, are applied to noise reduction for single trial evoked potentials which are buried in noise up to five times stronger than the signal. An empirical study using artificial data is presented and results are compared to the standard technique of averaging.

INTRODUCTION

An evoked potential (EP) is the electro cortical potential measurable in the human electro encephalogram (EEG) before, during and after sensoric, motoric or cognitive events. An EP is defined as the combination of the brain electric activity that occurs in association with the eliciting event and "noise", which is brain activity not related to the event together with inference from non-neural sources. For a set of N EPs, the assumed model for the ith EP is the following (see e.g. [5]):

$$\vec{z}_i(t) = \vec{x}(t) + \vec{y}_i(t); \quad i = 1, 2, \ldots N; \quad 0 \leq t < T \tag{1}$$

where $\vec{z}_i(t)$ is the ith recorded EP, $\vec{x}(t)$ the underlying signal, $\vec{y}_i(t)$ the noise associated with the ith EP, and T the duration over which each EP is recorded.

Since the noise contained in single trial EPs is up to five times stronger than the signal, it has to be reduced prior to any analysis of the signal. The common approach is to compute an average across several EPs recorded under equal conditions to improve the signal-to-noise ratio (SNR). However, there are two main drawbacks to averaging: (i) large numbers of single trials have to be recorded in order to estimate the signal; (ii) in the process of averaging all unique trial-to-trial information is lost; such varying latencies

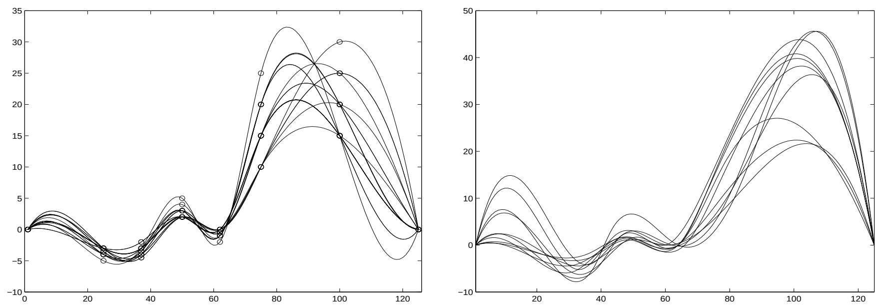

(a) (b)

Figure 1: (a) All $d = 22$ channels of an artificial EP plotted in one graph (x-axis time, y-axis amplitude). The underlying topographical patterns are shown as circles at $t = 1, 25, 37, 50, 62, 75, 100, 125$. Note the earlier onset of the $P300$ at more frontal electrodes relative to more posterior electrodes (i.e. amplitude peaks of $\sim 30\mu V$ at $t \approx 80$ and 100). (b) Ten artificial EP signals at electrode O2. Notice the between-trial variation in terms of latency and amplitude.

or amplitudes of components across trials also violate the basic assumptions behind averaging.

Most work so far has concentrated on various forms of FIR- and Wiener-Filtering for SNR enhancement (see e.g. [8] and [5]). Such filtering makes it possible to compute averages of high quality using less EP trials or even to estimate EP signals on a single trial basis. Our work is about new model-based approaches to noise reduction of single trial evoked potentials employing a Gaussian mixture model (GMM) with noise component and Principal Component Analysis (PCA). To our knowledge, model-based work on single trial EP analysis so far has concentrated on removing specific artifacts like eye movements or muscle tonus (see e.g. [4]) and not on reducing background noise as a whole. Our work is structured as follows: first we describe a testbed of artificial EPs which allows for systematic variation of noise characteristics; then we present all technical details concerning averaging, GMM with integrated noise component and PCA; in the results section we give quantitative measures for the achieved improvement in SNR; finally we discuss our approach by comparing it to other methods dealing with single trial estimation.

DATA

The goal of our work is to reduce the considerable amounts of noise contained in single trial EP signals. Since for real EP data the hidden signal is actually not known, we resort to simulated artificial EP data. Only the use of artificial data with known signal and noise proportions allows us to really quantify the gain in signal-to-noise ratio (SNR).

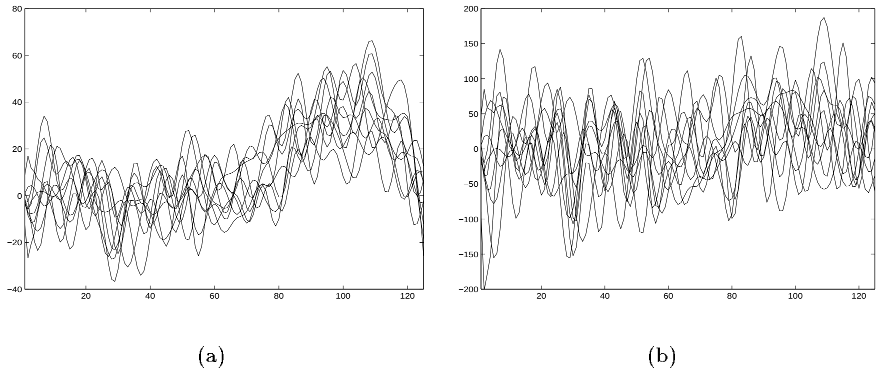

Figure 2: Ten artificial EPs at electrode O2 with noise added; (a) SNR was 1:1; (b) SNR was 1:5.

We simulated $500msec$ of a visually evoked EP recorded via $d = 22$ electrodes mounted according to the international 10-20 system. One artificial EP is of length $T = 125$ since we simulated a $250Hz$ recording. First we defined eight topographical patterns (each consisting of $d = 22$ values) corresponding to well known EP components like the $N100$ or $P300$. These topographical patterns have distinct amplitude values and temporal latencies (see Fig. 1 (a)). By computing a cubic spline interpolation for each signal dimension d we obtain values at all 125 sample points. To make our EP data more realistic we simulate between-trial variation by randomly shifting the topographical patterns both in time and amplitude prior to spline interpolation (change of temporal position: $t_{new} = t_{old} + 3 * \mathcal{N}(0,1)$, change of amplitude: $x_{new} = x_{old} + 0.33 * \mathcal{N}(0,1) * \mid x_{old} \mid$, $\mathcal{N}$ is the Gaussian distribution). The between-trial variation of ten EP signals is depicted in Fig. 1 (b).

We computed 5×100 of these artificial EPs with latency and amplitude variation. We added random Gaussian noise to each of these 5 EP sets at five different SNRs: 1:1, 1:2, 1:3, 1:4, 1:5. Gaussian noise was produced by limiting a $4000Hz$ Gaussian random signal to a bandwidth between 0 and $30Hz$ via FIR filtering and then down-sampling it to $250Hz$. A specific SNR was achieved by scaling the mean powers of signal and noise to yield the desired SNR. This yielded empirical mean *Input* SNR ratios of $3.34, -10.36, -18.45, -24.18$ and $-28.86dB$. Artificial EPs with noise added at ratios of 1:1 and 1:5 are depicted in Figs. 2 (a) and (b).

METHODS

Averaging

The common approach to the SNR is to compute an average across several EPs recorded under equal conditions. The average $\hat{\vec{x}}(t)$ over the sample of N EPs is used to estimate the underlying signal $\vec{x}(t)$:

$$\hat{\vec{x}}(t) = \frac{1}{N} \sum_{i=1}^{N} \vec{z}_i(t) = \vec{x}(t) + \frac{1}{N} \sum_{i=1}^{N} \vec{y}_i(t) \qquad (2)$$

Averaging will attenuate the noise $\vec{y}_i(t)$ and not the signal $\vec{x}(t)$ given that signal and noise linearly sum together to produce the recorded EPs $\vec{z}_i(t)$ and the evoked signal $\vec{x}(t)$ is the same for each recorded EP $\vec{z}_i(t)$ and the noise contributions $\vec{y}_i(t)$ can be considered to constitute statistically independent samples of a random process. If the above assumptions do not hold, the averaging will result in a biased, distorted estimate of the signal.

Gaussian Mixture Model with integrated noise component

The problem of noise reduction in single trial EPs is closely related to the estimation of speech model parameters in noisy acoustic environments. In [7] a general framework for estimation of integrated models of signal and background noise is presented. The authors distinguish between 3 processes: a process **Z** containing both the signal and the noise background, this is what usually can be directly observed (in our case this would be the single trial EP recordings); the signal **X** alone, in speech analysis this could be recordings with very low or zero background activity; the noise background **Y** alone, in speech analysis this could be segments of the recordings without any word utterances (at the beginning or in between).

The authors develop a general model around the assumption that all three processes can be modelled via mixtures of Gaussians. If **Z** can be observed it is necessary to either observe **X** or **Y** to be able to estimate the respective other "single" process. The authors present solutions for the case of additive Gaussian mixture models for **X** and **Y**. Both the signal $\vec{x}$ and the background noise $\vec{y}$ are modelled as mixtures of Gaussians, independent in time:

$$p(\vec{x}|\lambda_s) = \sum_{i=1}^{M} p_i b_i(\vec{x}) \qquad p(\vec{y}|\lambda_b) = \sum_{j=1}^{N} q_j a_j(\vec{y}) \qquad (3)$$

where the signal model $\lambda_s = \{p_i, \mu_i[d], \sigma_i[d]\}$ is a mixture of M Gaussian densities b_i and the background noise model $\lambda_b = \{q_j, \mu_j[d], \sigma_j[d]\}$ is a mixture of N Gaussian densities a_j, both with diagonal covariance matrices. The p_i and q_j are the mixing coefficients with $\sum_{i=1}^{M} p_i = 1$ and $\sum_{j=1}^{N} q_j = 1$.

If the observable noisy signal $\vec{z}$ is given via a general function of signal and background noise $\vec{z} = f(\vec{x}, \vec{y})$, than the likelihood of a noise corrupted observation $\vec{z}_t$ is given by:

$$p(\vec{z}_t|\lambda_s, \lambda_b) = \sum_{i=1}^{M} \sum_{j=1}^{N} p_i q_j p(\vec{z}_t|i, j, \lambda_s, \lambda_b) \qquad (4)$$

Our goal is now to obtain estimates of the underlying signal model by maximizing the likelihood of a set of noisy observations Z. Application of the Expectation-Maximization (EM) technique [2] yields the following general expressions for the estimation of the model parameters. Since all densities are assumed to have diagonal covariance matrices, each vector component can be estimated independently. The notational shorthands $\bar{p}_i$, $\bar{\mu}_i$ and $\bar{\sigma}_i$ therefore refer to arbitrary vector components.

$$\bar{p}_i = \frac{1}{T} \sum_{t=1}^{T} \sum_{j=1}^{N} p(i_t = i, j_t = j | z_t, \lambda) \tag{5}$$

$$\bar{\mu}_i = \frac{\sum_{t=1}^{T} \sum_{j=1}^{N} p(i_t = i, j_t = j | z_t, \lambda) E\{x_t | z_t, i_t = i, j_t = j, \lambda\}}{\sum_{t=1}^{T} \sum_{j=1}^{N} p(i_t = i, j_t = j | z_t, \lambda)} \tag{6}$$

$$\bar{\sigma}_i^2 = \frac{\sum_{t=1}^{T} \sum_{j=1}^{N} p(i_t = i, j_t = j | z_t, \lambda) E\{x_t^2 | z_t, i_t = i, j_t = j, \lambda\}}{\sum_{t=1}^{T} \sum_{j=1}^{N} p(i_t = i, j_t = j | z_t, \lambda)} - \bar{\mu}_i^2 \tag{7}$$

$$p(i_t = i, j_t = j | z_t, \lambda) = \frac{p(z_t | i_t = i, j_t = j, \lambda) p_i q_j}{\sum_{i=1}^{M} \sum_{j=1}^{N} p(z_t | i_t = i, j_t = j, \lambda) p_i q_j}. \tag{8}$$

If we assume an additive noise model with $f(x_t, y_t) = x_t + y_t$ and a background process which can be modelled by a single Gaussian (i.e. $N = 1$ in all above formulas), we get:

$$E\{x_t | z_t, i_t = i, \lambda\} = \frac{\sigma_i^2}{\sigma_i^2 + \sigma_b^2} \left[z_t + \left(\frac{\sigma_b^2}{\sigma_i^2} \mu_i - \mu_b \right) \right] \tag{9}$$

$$E\{x_t^2 | z_t, i_t = i, \lambda\} = \frac{\sigma_i^2 \sigma_b^2}{\sigma_i^2 + \sigma_b^2} + E\{x_t | z_t, i_t = i, \lambda\}^2 \tag{10}$$

$$p(z_t | i_t = i, \lambda) = \mathcal{N}[z_t, \mu_i + \mu_b, \sigma_i^2 + \sigma_b^2] \tag{11}$$

where $\mathcal{N}[z_t, \mu_i + \mu_b, \sigma_i^2 + \sigma_b^2]$ is the Gaussian distribution corresponding to the observed noisy signal. Since both the signal x_t and the background y_t are random variables with Gaussian distribution their sum is also a Gaussian distribution with the above parameters.

Once a GMM with integrated noise component is trained on a set of EPs we can use the component densities given in Equ. 11 to weigh the means of the Gaussians $\vec{\mu}_i$ at each time step t and obtain a denoised estimate $\hat{\vec{x}}_t'$ of the signal:

$$\hat{\vec{x}}_t' = \sum_{i=1}^{M} p(\vec{z}_t | i_t = i, \lambda) \vec{\mu}_i \tag{12}$$

Principal Component Analysis

Principal Component Analysis (PCA) (see e.g. [6]) finds a linear transformation of a matrix onto a set of orthonormal vectors accounting for the maximum possible fraction of the total sample variance. We use Singular Value Decomposition (SVD) to compute the prinicipal components of each single EP trial. A single EP trial can be seen as a t (time points)$\times d$ (channels) matrix X which can be decomposed into a product of three matrices:

$$X = USV^T. \tag{13}$$

S is a diagonal matrix of the same dimension as X with nonnegative diagonal elements in decreasing order which are the singular values. Both U and V are unitary matrices.

Let us assume matrix X describes a signal of intrinsic dimensionality $s \ll d$ which is contaminated with noise filling the rest of the $d-s$ dimensions in some uniform way. In this case we expect the first few singular values to be rather big and the rest of them to be all small and of almost equal size accounting for the "noise floor". If the differences between consecutive singular values are computed, the maximum difference indicates the drop from the big singular values to the "noise floor". In Fig. 3 histogram plots show at which difference these maxima occured. The majority of maximal differences can always be found between the first and second singular value, this result being clearer for data sets with higher SNR.

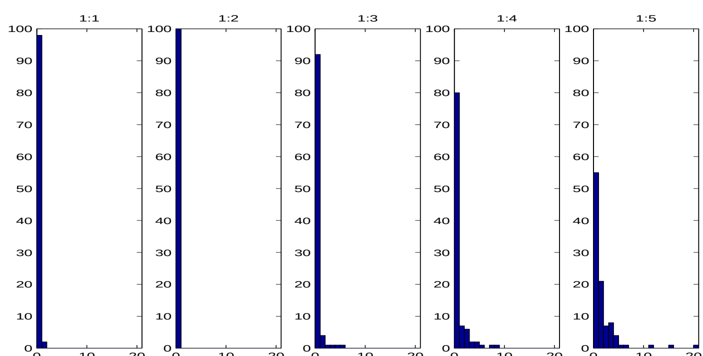

Figure 3: Histogram plots of indexes of maximum differences between singular values. The left-most bars in each plot indicate that the majority of maximum differences appeared between the first and second singular values. Plots are ordered for SNRs of 1:1 to 1:5 (left to right).

We therefore set all singular values in S except the first one to zero obtaining a new diagonal matrix $S_{denoise}$. This allows us to subtract the noise contributions and compute denoised estimates $\hat{X}$ for all single EP trials:

$$\hat{X} = US_{denoise}V^T. \tag{14}$$

Another technique based on a linear transformation that has found wide spread application to the analysis of EEG is Independent Component Analysis (ICA) (see [3] for a comprehensive introduction). Contrary to PCA, ICA components are not just uncorrelated but statistically independent and can therefore be non-orthogonal. ICA has been applied successfully to the removal of specific artifacts like eye movements or muscle tonus in single trial analysis [4]. Nevertheless we believe it is not applicable to the reduction of background noise as described above for the following reasons: the independent components must have non-gaussian distributions which is in clear contradiction to the model put forward in Equ. 1; there is no automatic way of deciding which of the obtained components is due to a signal and which to noise (viusal checking and selection of components is rather impractical for hundreds of EPs).

RESULTS

We applied averaging, Gaussian Mixture Models with integrated noise component (GMMnoise) and Principal Component Analysis (PCA) to all five artificial data sets. We used mixtures of $M = 10$ Gaussians for the signal $\vec{x}$ and a single Gaussian ($N = 1$) for the background noise $\vec{y}$. Using mixtures of $M = 20$ Gaussians yielded only minor improvements. As an estimate of the noise mean μ_b and variance σ_b^2 we used the respective parameters of the residual noise obtained via averaging [1]. For each set of 100 artificial EPs we computed an average according to Equ. 2 and subtracted it from all 100 artificial EPs. The remainder is the residual noise. We computed a PCA for each of the 5×100 single trials separately and denoised the signals by setting all but the first singular value to zero.

Since our data are artificial EPs we know exactly which signals $\vec{x}$ are hidden in the noisy signals $\vec{z}$. We can compare the denoised estimates obtained via GMMnoise and PCA as well as the common average given in Equ. 2 to these signals $\vec{x}$. We obtain noise residuals by subtracting the real signals $\vec{x}$ from the denoised signal estimates $\hat{\vec{x}}'_t$, $\hat{X}$ and $\hat{\vec{x}}$. The ratio in dB between the real signals $\vec{x}$ and these noise residuals are called *Output* SNRs and are given in Tab. 1 for GMMnoise, PCA and averaging. The same information is depicted in Fig. 4(a).

If the power of the underlying signal is approximately equal to the power of the background noise (data set “1:1”), GMMnoise and PCA are able to recover the information based on single EP trials as good as averaging which uses the full 100 trials of the data set. At lower SNRs, the performance of both GMMnoise and PCA decreases relative to averaging. But even for data set “1:5”, the improvement in SNR is around $25dB$ for GMMnoise and $20dB$ for PCA. Our results are comparable to those achieved with other more involved

[1] Since averaging does not fully remove the noise from the signal the variance of the noise residual is always underestimating the true noise variance. One could try to think about accounting for this to get more accurate estimates.

data set	1:1	1:2	1:3	1:4	1:5
Input SNR	3.34	−10.36	−18.45	−24.18	−28.86
Output SNR average	16.92	15.20	14.11	12.38	10.47
Output SNR GMMnoise	16.43	9.92	4.38	−0.02	−3.17
Output SNR PCA	17.63	10.08	2.15	−4.28	−9.94

Table 1: Results for averaging, GMMnoise and PCA applied to the five EP data sets given as SNR in dB.

approaches. In [1] an ARX (autoregressive with exogenous input) model for single trial estimation is employed using separate EEG and EOG (electro-oculogram) noise sources. The authors report an improvement in SNR of around $20dB$ for an Input SNR of $-30dB$ which is very similar to what GMMnoise and PCA achieve for data set "1:5". In [8] an a posteriori Wiener filter is employed which is time-varying and uses information from multiple electrodes to estimate single trial EPs. Contrary to both GMMnoise and PCA this is a non-stationary technique taking into account temporal variations of noise and signal statistics. In their study on simulated EP data, the best improvement in SNR achieved for their lowest Input SNR of $-9dB$ is around $12dB$. GMMnoise and PCA achieved improvements of around $20dB$ for an Input SNR of $-10.36dB$ (see Tab. 1, data set "1:2"). In Fig. 5 examples of the underlying EP signals plus their reconstructions via denoising are shown for all three methods at SNRs of 1:1 and 1:5.

Another application of GMMnoise and PCA is to use it as a preprocessing technique prior to averaging. The improvement in SNR might make it possible to compute averages of high quality using less EP trials. This would greatly benefit experimental work since recording of EPs is time and cost consuming.

We compared averaging, averaging after application of the GMMnoise

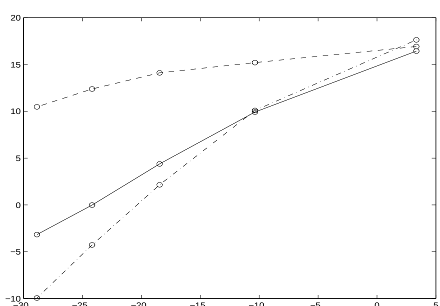

(a)

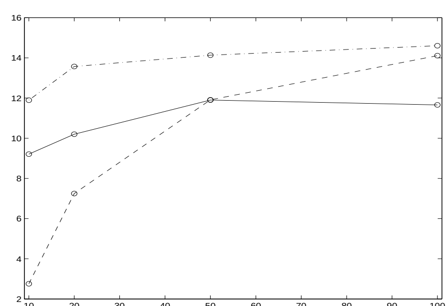

(b)

Figure 4: (a) Input SNR (x-axis) vs. Output SNR (y-axis) in dB for GMMnoise (solid), PCA (dash-dot) and averaging (dashed). (b) Number of EP trials (x-axis) vs. Output SNR (y-axis) for GMMnoise+ave (solid), PCA+ave (dash-dot) and averaging (dashed).

	average	GMMnoise+ave	PCA+ave
10 EP trials	2.76	9.21	11.90
20 EP trials	7.25	10.20	13.57
50 EP trials	11.91	11.90	14.13
100 EP trials	14.11	11.66	14.61

Table 2: Results in terms of Output SNR in dB for average, GMMnoise+ave, and PCA+ave applied to different numbers of EP trials (Input SNR was 1:3).

($M = 10, N = 1$ just as before) method (GMMnoise+ave) and averaging after application of PCA (PCA+ave) on subsets of the "1:3" data set containing 10, 20, 50 and 100 EP trials. Estimates of the noise means μ_b and variances σ_b^2 were based on the respective parameters of the residual noise obtained via averaging of 10, 20, 50 or 100 trials. Results are given in Tab. 2 and Fig. 4(b). Results achieved by averaging alone are worse for numbers of trials smaller than 50 when compared to results obtained by GMMnoise+ave and PCA+ave. Both GMMnoise+ave and PCA+ave show considerable improvements of SNR already using only 10 or 20 EP trials. PCA+ave is better than the other two methods regardless of how big the subset of the EP data set is.

CONCLUSION

We presented the application of two model-based techniques for the reduction of noise in single trial evoked potential (EP) analysis. Both Gaussian mixture models with integrated noise component (GMMnoise) and Principal Component Analysis (PCA) have been shown to be able to significantly improve the SNR in single trial EPs. Especially PCA seems to be suited as a preprocessing technique to be used prior to averaging enabling computation

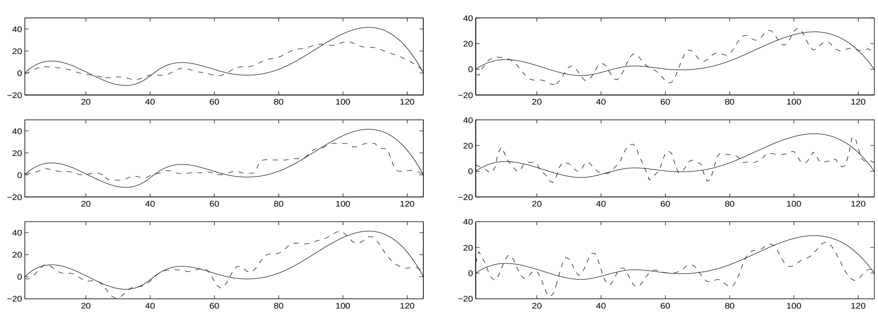

(a) (b)

Figure 5: Underlying EP signal (solid) at electrode O2 with denoised signal (dashed) obtained via averaging (top), GMMnoise (middle) and PCA (bottom); (a) SNR was 1:1; (b) SNR was 1:5.

of high quality averages based on less single trials.

GMMnoise is basically a mixture of Gaussian approach where one Gaussian kernel is reserved for noise modeling and its parameters are fixed prior to estimation of the whole model. The fact that the parameters of this "noise kernel" have to be estimated via common averaging as well as the rather heuristic choice of model order (number of kernels) are clear drawbacks of the method. PCA on the other hand is a self-contained method and the problem of model order (number of singular values set to zero for denoising) seems to be solved.

Although it was necessary to use artificial data to accurately measure the gain in SNR obtained via GMMnoise and PCA the ultimate test would of course be its application to real EP data.

Acknowledgements:

The Austrian Research Institute for Artificial Intelligence is supported by the Austrian Federal Ministry of Education, Science and Culture. This research is supported by the Austrian Scientific Research Fund (FWF), project P14100-MED.

REFERENCES

[1] Cerutti S., Chiarenza G., Liberati D., Mascellani P., Pavesi G.: A parametric method for identification of single-trial event related potentials in the brain, IEEE Trans. on Biomedical Engineering, vol. BME-35, pp. 701-711, 1988.

[2] Dempster A.P., Laird N.M., Rubin D.B.: Maximum likelihood from incomplete data via the EM algorithm, Journal of the Royal Statistical Society, B, 39:1-38, 1977.

[3] Hyvaerinen A.: Survey on Independent Component Analysis, Neural Computing Surveys, Vol. 2, 94-128, 1999.

[4] Jung T.-P., Makeig S., Humphries C., Lee T.-W., Mckeown M.J., Iragui V., Sejnowski T.J.: Removing electroencephalographic artifacts by blind source separation, Psychophysiology, 37, 163-178, 2000.

[5] Picton T.W.(ed.): Human Event Related Potentials, Elsevier Science Publishers, 1988.

[6] Ripley B.D.: Pattern Recognition and Neural Networks, Cambridge University Press, 1996.

[7] Rose R.C., Hofstetter E.M., Reynolds D.A.: Integrated Models of Signal and Background with Application to Speaker Identification in Noise, IEEE Transactions on Speech and Audio Processing, Vol. 2, No.2, pp.245-257, 1994.

[8] Westerkamp J., Aunon J.I.: Optimum multielectrode a posteriori estimates of single-response evoked potentials, IEEE Transactions on Biomedical Engineering, 34, 13-22, 1987.