



BME I5100: Biomedical Signal Processing

Linear Mixtures and Independent Component Analysis



Lucas C. Parra
Biomedical Engineering Department
City College of New York





Schedule

Week 1: Introduction

Linear, stationary, normal - the stuff biology is **not** made of.

Week 1-4: Linear systems

Impulse response

Moving Average and Auto Regressive filters

Convolution

Discrete Fourier transform and z-transform

Sampling

Week 5-8: Random variables and stochastic processes

Random variables

Moments and Cumulants

Multivariate distributions

Stochastic processes

Week 9-14: Examples of biomedical signal processing

Probabilistic estimation

Harmonic analysis - **estimation** circadian rhythm and speech

Linear discrimination - **detection** of evoked responses in EEG/MEG

Independent components analysis - **analysis** of MEG signals

Auto-regressive model - **estimation** of the spectrum of 'thoughts' in EEG

Matched and Wiener filter - **filtering** in ultrasound

Linear Mixtures – Problem statement

$$X = A S$$

Q: Given X can you tell what A and S is?

A: Yes! Use **prior information** on A and S .

Linear Mixtures – Problem statement

Basic physics often leads to linear mixing where

- Rows in S are sources $s_i(t)$
- rows in X are sensors readings $x_j(t)$.
- rows in A are the amount the different sources i contribute to a sensor j due to a physical mixing process with coefficients a_{ji} .

Examples are:

X	=	A	*	S
Acoustic mic. array	=	room response	*	sound amplitude
Spectroscopy spectra	=	concentration	*	emission spectra
Hyperspectral image	=	abundance	*	reflection spectra
EEG electrical potential	=	elect. Potential	*	electrical coupling
MEG magnetic field	=	electrical current	*	magnetic coupling

Linear Mixtures – Prior information

Depending which prior information we get different results:

- Columns in A and rows in S orthogonal:

Principal Component Analysis (PCA)

- Rows in S statistically independent:

Independent Component Analysis (ICA)

- Rows in S orthogonal and white (or non-stationary):

Multiple Diagonalization

- A and S positive:

Non-negative Matrix Factorization (NMF)

Assignment 12: Read Lee & Seung Nature 1999, NIPS 1999

Linear Mixtures – Independent Components

Consider an invertible linear mixture

$$\mathbf{x}(t) = \mathbf{A} \mathbf{s}(t) \qquad \mathbf{s}(t) = \mathbf{W} \mathbf{x}(t)$$

where we denoted the inversion by $\mathbf{W} = \mathbf{A}^{-1}$.

In **Independent Component Analysis** (ICA) one assumes that the sources $\mathbf{s}(t)$ are **statistically independent**:

$$p(\mathbf{s}(t)) = p(s_1(t), s_2(t), \dots, s_d(t)) = \prod_{i=1}^d p(s_i(t))$$

To estimate the \mathbf{A} we use again Maximum Likelihood. The likelihood of i.i.d. observations $\mathbf{x}[n] = \mathbf{x}(t_n)$, $n=1, \dots, T$:

$$p(\mathbf{x}[1], \dots, \mathbf{x}[T] | \mathbf{A}) = \prod_{n=1}^T p(\mathbf{x}[n] | \mathbf{A})$$

Linear Mixtures – Independent Components

For any invertible transformation $\mathbf{s} = f(\mathbf{x})$,

$$p_{\mathbf{x}}(\mathbf{x}) = \left| \frac{d\mathbf{s}}{d\mathbf{x}} \right| p_{\mathbf{s}}(\mathbf{s})$$

In particular for $\mathbf{s} = \mathbf{A}^{-1} \mathbf{x} = \mathbf{W} \mathbf{x}$


The log-likelihood is then $p_{\mathbf{x}}(\mathbf{x}) = |\mathbf{W}| p_{\mathbf{s}}(\mathbf{s}) = |\mathbf{W}| p_{\mathbf{s}}(\mathbf{W} \mathbf{x})$

$$\begin{aligned} L(\mathbf{W}) &= \ln p_{\mathbf{x}}(\mathbf{x}[1], \dots, \mathbf{x}[T] | \mathbf{W}) = \sum_{n=1}^T \ln p_{\mathbf{x}}(\mathbf{x}[n] | \mathbf{W}) \\ &= T \ln |\mathbf{W}| + \sum_{n=1}^T \ln p_{\mathbf{s}}(\mathbf{W} \mathbf{x}[n]) \\ &= T \ln |\mathbf{W}| + \sum_{n=1}^T \sum_{i=1}^d \ln p_{\mathbf{s}}(\mathbf{w}_i^T \mathbf{x}[n]) \end{aligned}$$

Linear Mixtures – Gradient ascent

We can find the maximum of $L(W)$ with gradient ascent:

$$W_{t+1} = W_t + \mu \frac{\partial L(W)}{\partial W}$$



Linear Mixtures – Stochastic gradient ascent

We can find the maximum of $L(\mathbf{W})$ with gradient ascent:

$$\begin{aligned}\mathbf{W}_{t+1} &= \mathbf{W}_t + \mu \frac{\partial L(\mathbf{W})}{\partial \mathbf{W}} \\ &= \mathbf{W}_t + \mu \frac{\partial}{\partial \mathbf{W}} \sum_{n=1}^T L(\mathbf{x}[n]|\mathbf{W})\end{aligned}$$

Stochastic gradient ascent **updates for every sample n**

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mu \frac{\partial L(\mathbf{x}[n]|\mathbf{W})}{\partial \mathbf{W}}$$

making the assumption that that instantaneous gradient is a unbiased estimate of the full gradient.

Linear Mixtures – Independent Components

We can find the maximum of $L(\mathbf{W})$ with gradient ascent

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mu \left(\mathbf{W}^{-T} + \mathbf{u}[n] \mathbf{x}^T[n] \right)$$

where we have defined $\mathbf{u} = [u_1, \dots, u_d]^T$ with $u_i = \partial \ln p_s(s_i) / \partial s_i$


We can always multiply the gradient with a positive definite matrix, for instance $\mathbf{W}^T \mathbf{W}$

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mu \left(\mathbf{I} + \mathbf{u}[n] \mathbf{s}^T[n] \right) \mathbf{W}$$

If we assume high kurtosis signals (long tails) we can use Laplacian distribution

$$p_s(s) = \frac{\lambda}{2} \exp(-\lambda |s|) \quad u(s) = -\lambda \operatorname{sign}(s)$$


Linear Mixtures – ICA in MEG




$$X = A * S$$

X Magnetic fields measured in SQUID sensors
 A Magnetic coupling or attenuation due to geometry
 S Effective current flows in neuronal population

i^{th} column in A



stimulus locked $s_i(t)$



Linear Mixtures – Principal Components

Maximum Likelihood gives Principal Components if we assume:

- Sources are Gaussian, i.e. $p_s(s) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{s^2}{2\sigma^2}\right)$
- Mixing are rotations, i.e. $\mathbf{W}^{-1} = \mathbf{W}^T$

To see this set the gradient of the log-likelihood to zero:

$$0 = T \mathbf{W}^{-T} + \sum_{n=1}^T \mathbf{u}[n] \mathbf{x}^T[n]$$

For a Gaussian, $u = \partial \ln p(s) / \partial s = -s / \sigma^2$. With $\mathbf{\Lambda} = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$

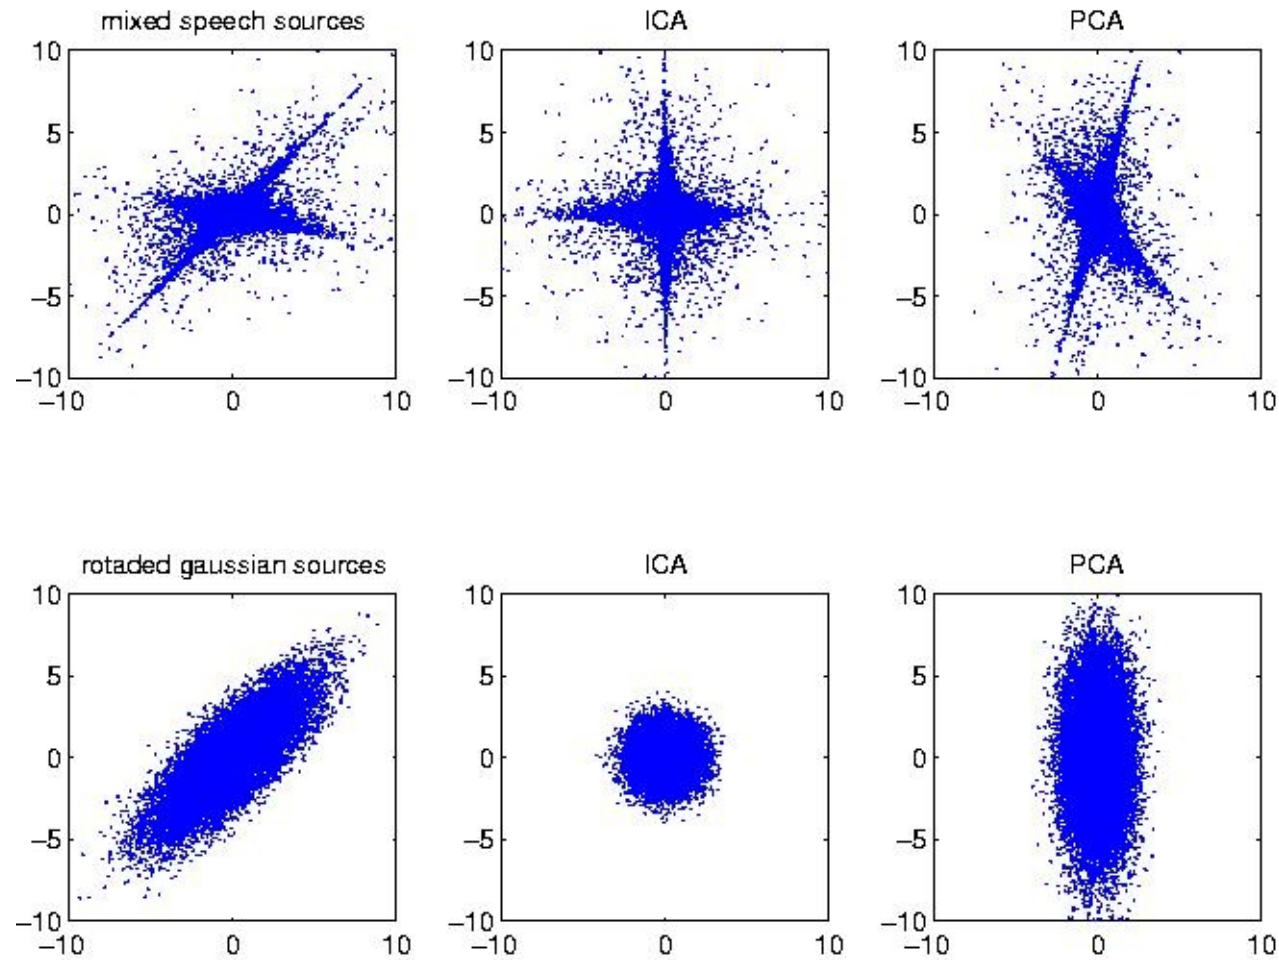
$$\mathbf{W}^{-T} = \frac{1}{T} \sum_{n=1}^T \mathbf{\Lambda}^{-1} \mathbf{s}[n] \mathbf{x}^T[n] = \frac{1}{T} \sum_{n=1}^T \mathbf{\Lambda}^{-1} \mathbf{W} \mathbf{x}[n] \mathbf{x}^T[n] = \mathbf{\Lambda}^{-1} \mathbf{W} \mathbf{R}_x$$

Using orthogonality we obtain PCA:

$$\mathbf{R}_x = \mathbf{W}^T \mathbf{\Lambda} \mathbf{W}$$

Linear Mixtures – ICA and PCA

Comparing results of ICA and PCA



Linear Mixtures – Multiple diagonalization

Statistical independence implies for all $i \neq j, t, l, n, m$:

$$E[s_i^n(t) s_j^m(t+l)] = E[s_i^n(t)] E[s_j^m(t+l)]$$

So far we have talked about same number of sensors than sources. In general for M sources and N sensors each t, l, n, m gives $M(M-1)/2$ conditions for NM unknowns in \mathbf{A} .

Sufficient conditions if we use multiple:

<u>use</u>	<u>sources assumed</u>	<u>resulting algorithm</u>
n, m	non-Gaussian	ICA
t	non-stationary	multiple decorrelation
l	non-white	multiple decorrelation

Linear Mixtures – Multiple diagonalization

For **stationary non-white sources** we have:

$$\mathbf{R}_s(l) = E[\mathbf{s}(t) \mathbf{s}^T(t+l)]$$

Which is diagonal assuming second order independent sources. The measured cross-correlation $\mathbf{R}_x(l) = E[\mathbf{x}(t) \mathbf{x}^T(t+l)]$ is then

$$\mathbf{R}_x(l) = \mathbf{A} \mathbf{R}_s(l) \mathbf{A}^T$$

Combing these equations for two time delays $l=l_1, l_2$ leads to a generalized eigenvalue problem for \mathbf{A} ,

$$\mathbf{R}_x(l_1) \mathbf{R}_x^{-1}(l_2) \mathbf{A} = \mathbf{A} \Lambda_s(l_1) \Lambda_s^{-1}(l_2)$$

Warning: Generalize Eigenvalue not robust to noise!

For increased stability diagonalize multiple delays l . This is known as the SOBI algorithms (second order blind identification)

Linear Mixtures – Multiple diagonalization

For **non-stationary white sources** we have:

$$\mathbf{R}_s(t) = E[\mathbf{s}(t) \mathbf{s}^T(t)]$$

Which is diagonal assuming second order independent sources.
The measured cross-correlation $\mathbf{R}_x(t) = E[\mathbf{x}(t) \mathbf{x}^T(t+l)]$ is then

$$\mathbf{R}_x(t) = \mathbf{A} \mathbf{R}_s(t) \mathbf{A}^T$$

Combing these equations for two time intervals $t=t_1, t_2$ leads to a generalized eigenvalue problem for \mathbf{A} ,

$$\mathbf{R}_x(t_1) \mathbf{R}_x^{-1}(t_2) \mathbf{A} = \mathbf{A} \Lambda_s(t_1) \Lambda_s^{-1}(t_2)$$

This is equivalent to “Common Spatial Pattern”, which diagonalizes two covariance matrices measured at different times)

Warning: Generalize Eigenvalue not robust to noise!

Linear Mixtures – Multiple diagonalization

For **non-Gaussian sources** we have:

$$E[s_i^u s_j^v] = E[s_i^u] E[s_j^v] \quad , \quad i \neq j$$

From this one can derive for a linear combination of 4-order tensors (cross-cumulants) with symmetric matrix \mathbf{M}

$$\mathbf{C}_s(\mathbf{M}) = E[\mathbf{s}^T \mathbf{M} \mathbf{s} \mathbf{s} \mathbf{s}^T] - \mathbf{R}_s \text{Trace}(\mathbf{M} \mathbf{R}_s) - 2 \mathbf{R}_s \mathbf{M} \mathbf{R}_s$$

The following diagonalization condition

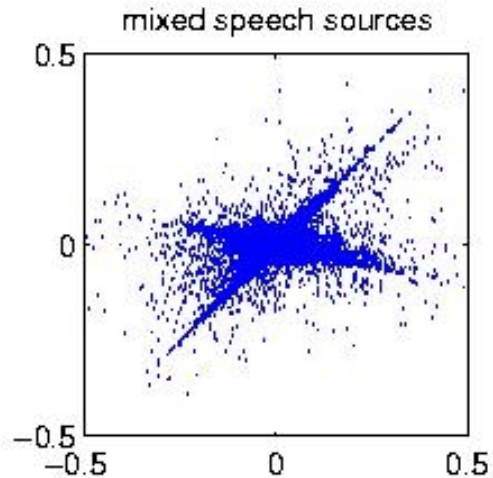
$$\mathbf{C}_x = \mathbf{A} \mathbf{C}_s (\mathbf{A}^T \mathbf{A}) \mathbf{A}^T$$

This can again - in combination with diagonal covariance - be combined to a generalized eigenvalue equation.

And again, for stability one should use use more than two cumulants, which leads to the “JADE” algorithms.

Linear Mixtures – Multiple diagonalization

Comparing different diagonalization criteria



% linear mix of sources S

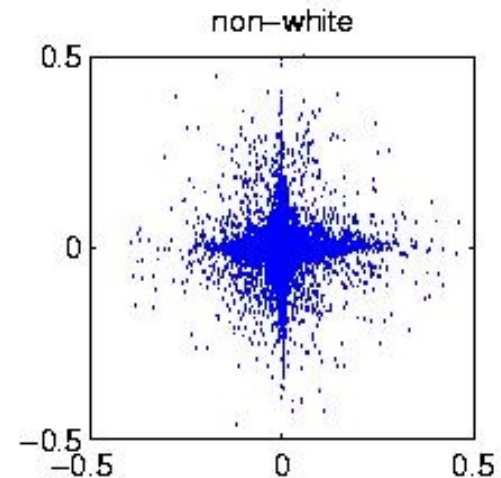
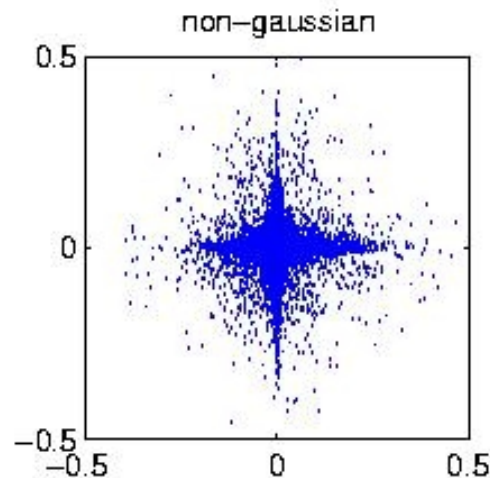

$$X=A*S;$$

% Separation based on Generalized Eigenvalues

$$[W,D]=\text{eig}(X*X',Q);$$

$$S=W' *X;$$

Results with Q assuming:




Linear Mixtures – Multiple diagonalization



$$X = A * S$$

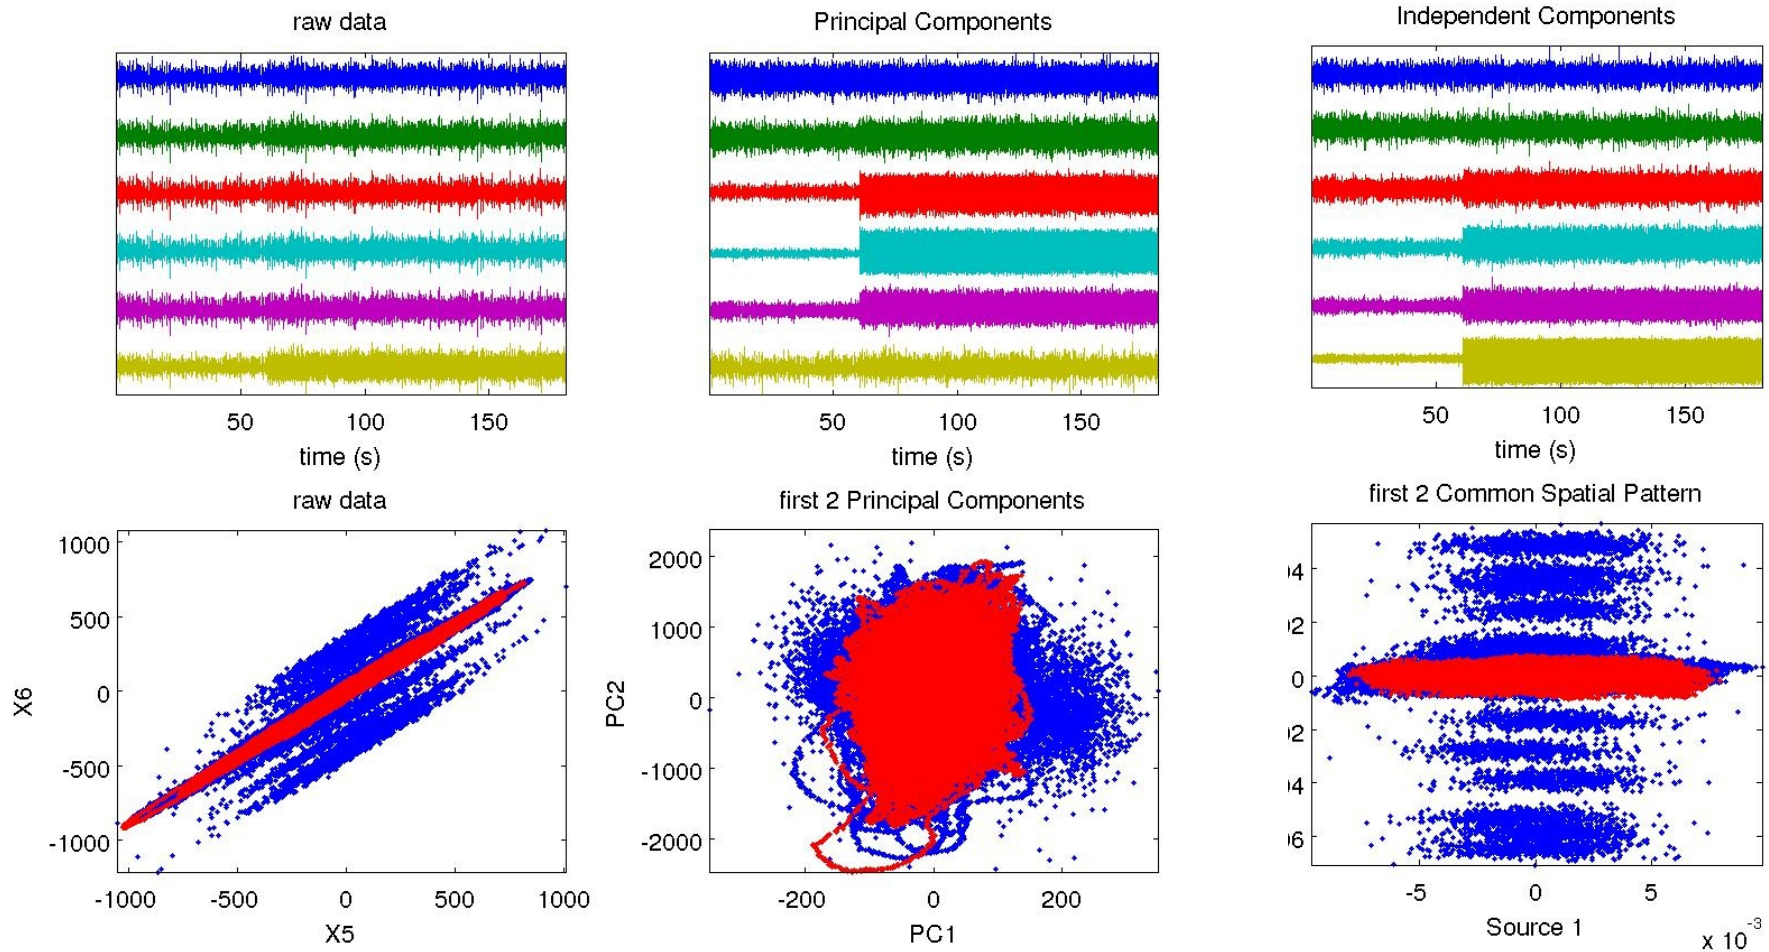
Electrical potentials on the skull surface Electric coupling or attenuation due to tissue resistance Large scale potential of neuronal population

Example of BSS on EEG using multiple diagonalization



Linear Mixtures – Multiple diagonalization

Example: 6D Local Field Potentials



$$R=X(t2,:) * X(t2,:)' ;$$

$$Q=X(t1,:) * X(t1,:)' ;$$

$$[W,D]=\text{eig}(R) ;$$

$$[W,D]=\text{eig}(R,Q) ;$$

$$S=W' * X ;$$

Linear Mixtures – Multiple diagonalization

Stimulus triggered evoked responses of these LFP (averaged over multiple repeats) for different spatial projections of the data:

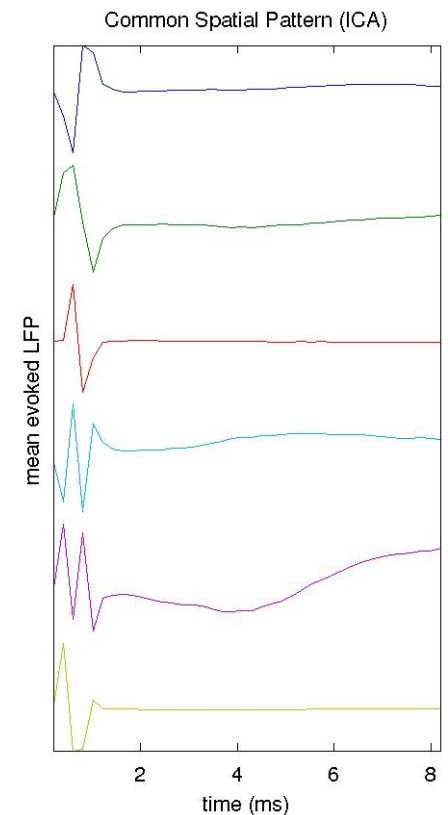



$Y = W^T X$, with baseline covariance Q and stimulus covariance R :

$W = \text{eye}(6)$

$[U, D] = \text{eig}(Q);$
 $W = (I - U(:, 6) * U(:, 6));$

$[U, D] = \text{eig}(R);$
 $W = U;$

$[U, D] = \text{eig}(R, Q);$
 $W = U;$



Linear Mixtures – Maximum SNR component

Assume we are looking for a linear response $r(t)$

$$r(t) = \mathbf{w}^T \mathbf{x}(t)$$


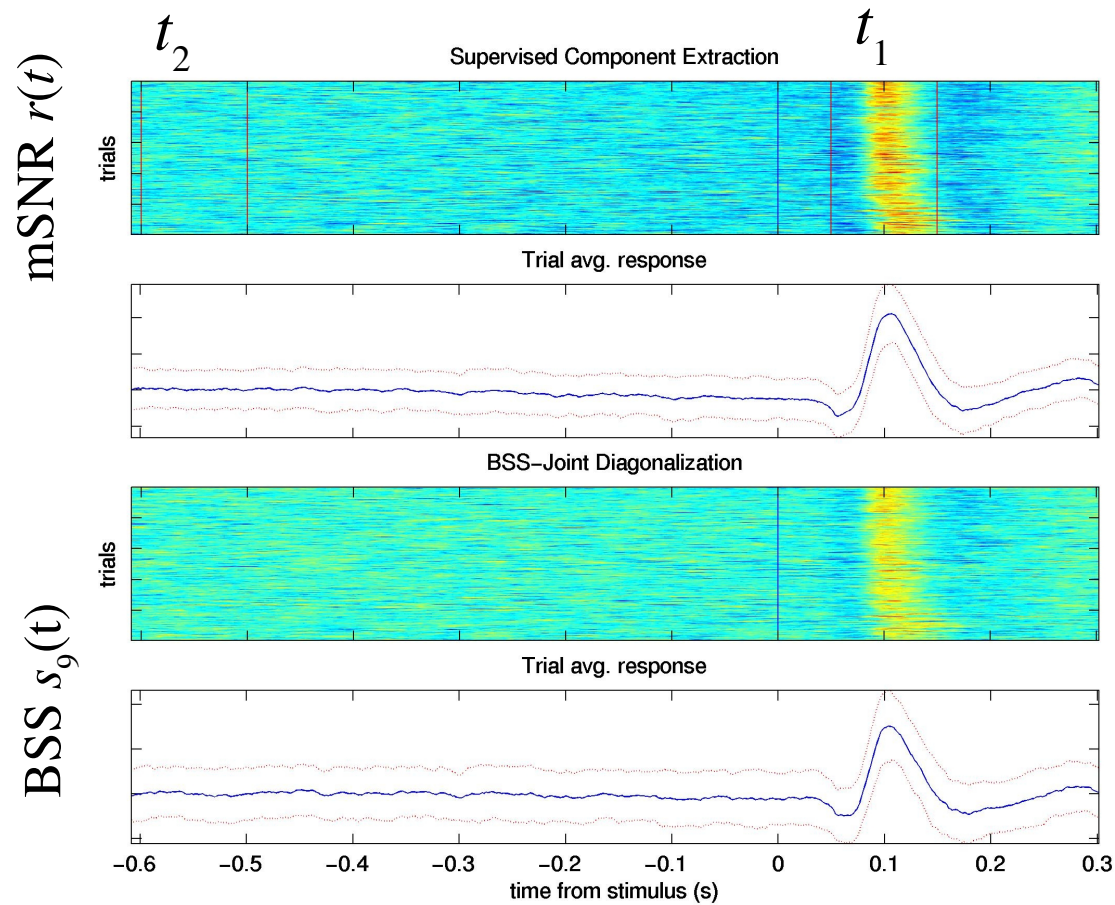
which has **maximal power** during specific times t_1 as compared to the **power** of baseline activity during times t_2 . That is, we are looking for a linear component \mathbf{w} that maximize the power ratio or signal to noise ratio (SNR)

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmax}} \frac{E[r^2(t_1)]}{E[r^2(t_2)]} = \underset{\mathbf{w}}{\operatorname{argmax}} \frac{\mathbf{w}^T \mathbf{R}_x(t_1) \mathbf{w}}{\mathbf{w}^T \mathbf{R}_x(t_2) \mathbf{w}}$$

With the usual definition, $\mathbf{R}_x(t) = E[\mathbf{x}(t) \mathbf{x}^T(t)]$. The solution is given by the maximum generalized eigenvector. Hence, the first component in the previous approach with two stationarity times has maximum SNR.

Linear Mixtures – Maximum SNR component

Comparing BSS based on non-stationarity and non-white in MEG



auditory locked

Assignment: source separation

Use PCA, and multiple diagonalization to generate various projections of EEG data.

1. Load the file `eeg-vep.mat` from the class webpage. It should have epoched data with 160 channels, 545 samples and 116 trials saved in variable `eeg`.
2. Stack up the trials so that you have a data matrix of channels by samples (160 x 1630).
3. Compute the covariance and from this principal components of this data. Show the eigenvalue spectrum (on a dB scale). Display first principal component (the eigenvector with the strongest eigenvalue) on the scalp using the `topoplot()` function using the corresponding location file provided on the website.
4. Project the data onto the first two principal component and display the result as an images of trials by samples (116x545 for each of the two components). Also show the average across trials for the two components (two curved using plot).
5. Repeat steps 2 through 4 using instead the first and second half of the samples (samples 1:272 and 273:545) to obtain two covariance matrices. Find the eigenvectors that diagonalize both these matrices (generalized eigenvectors).
6. Repeat step 5 but now using the covariance of all the data and as the second correlation matrix use the the cross-correlation of the data with a version of the data that is delayed by K samples (try different delays). Be sure to symmetrize this matrix with $R_{xy}=0.5*(R_{xy}+R_{xy}')$;

In total you should have 3 figures with results for 3 different techniques (steps 4, 5 and 6). For each method combine plots with subplot into single figure. Label all axis and use milliseconds on the time axis.