Dynamic causal modelling for event-related analysis of EEG or MEG data :: Part I Theory

Quentin Huys

with Marta Garrido

Thanks to Stefan Kiebel

Methods for dummies, FIL, London

June 2006

Disclaimer

I'm not an expert – there are probably errors.

The slides are rather wordy, as I wanted them to be self-explanatory.

Please let me know of any errors: `qhuys ~at~ gatsby[.]ucl[.]ac[.]uk`

Outline

This talk (part I) will focus on the theory:

- EEG / EMG data – issues and aims
- Observation model: Linear superposition of potentials
- Dynamical system for EEG and MEG: neural mass models
- Inference in latent dynamical systems
- Discussion

Marta's talk (part II) will discuss an application to give some insight into the theory.

Data: what is measured

EEG: at each electrode, measure potential relative to a reference point
MEG: at each sensor, measure tangential parts of magnetic field (but not radial one).

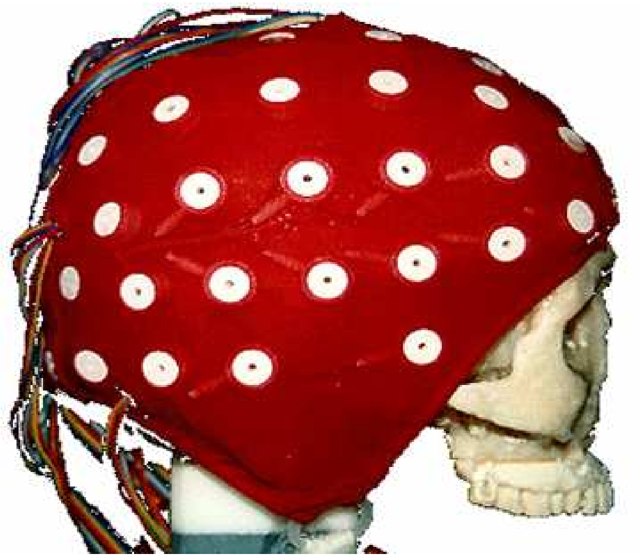

EEG electrodes

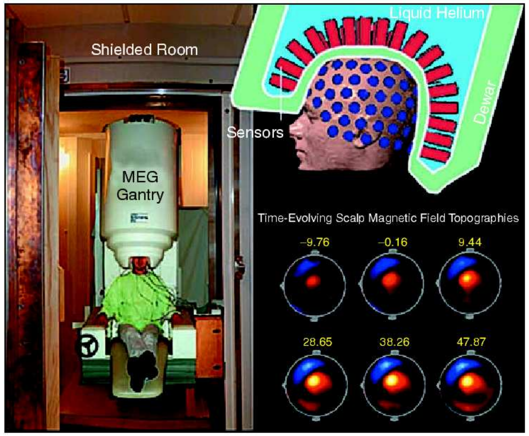

MEG machine

pictures from Stefan Kiebel and Baillet et al. (2001)

Data: what is measured, contd.

Both are due to cellular currents. Primary currents induce secondary currents, which arise due to current conservation law

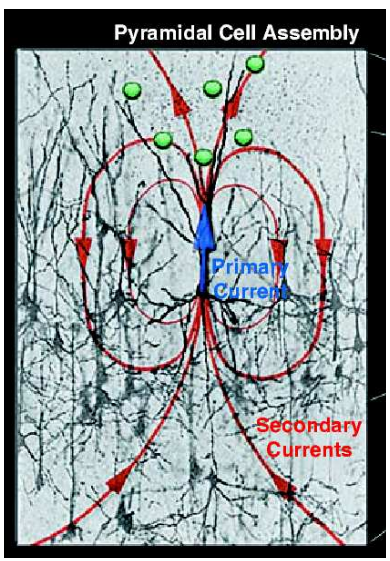

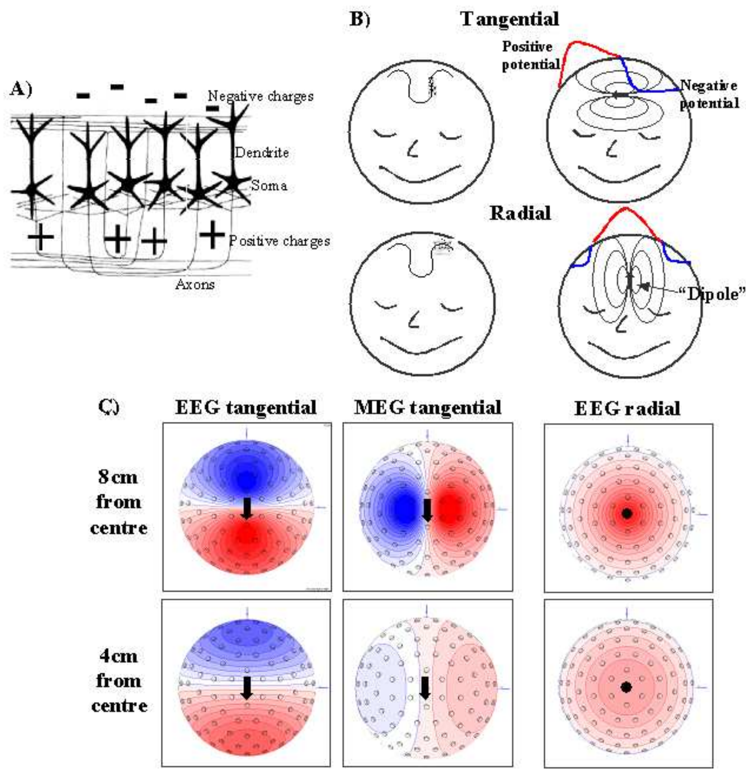

from Baillet et al. (2001) and MRC-CBU website.

EEG measures potentials due to secondary currents.
MEG measures magnetic fields due to primary currents.

Data: what is measured, contd.

Main current flow (axis along pyramidal dendrites) generates electric $\mathbf{E}$ and magnetic field $\mathbf{B}$. Given a current source $\mathbf{J}(\mathbf{r})$ at a point $\mathbf{r}$, $\mathbf{E}(\mathbf{r}')$ and potential $V(\mathbf{r}')$ at another point $\mathbf{r}'$ can be found by solving:

$$\nabla \cdot \mathbf{E} = \rho/\epsilon\epsilon_0 \qquad \nabla \cdot \mathbf{B} = 0$$
$$\nabla \times \mathbf{E} = 0 \qquad \nabla \times \mathbf{B} = \mu_0 \mathbf{J}(\mathbf{r})$$

which for a current dipole $\mathbf{q}$ (a vector in the direction of the current with its length the current strength) in a spherical head turns out to be very simple

$$\mathbf{B}_0(\mathbf{r}') \propto \mathbf{q} \cdot \frac{\mathbf{r}' - \mathbf{r}}{r||\mathbf{r} - \mathbf{r}'||^3}$$

which is linear in $\mathbf{q}$, but **very** nonlinear in the position of the dipole $\mathbf{r}$. Part of this is a result known as Biot-Savart law (Baillet et al., 2001).

Source localization

Very imporantly, contributions from several sources sum linearly. Thus we have that the data from all N sensors at time t $\mathbf{y}_t$ is a linear function of the activity of the M sources $\mathbf{s}_t$ at that time.

$$\mathbf{y}_t = \mathbf{L}\mathbf{K}\mathbf{s}_t = \mathbf{A}\mathbf{s}_t$$

where $\mathbf{L}$ is the lead field matrix. Each column of the lead field matrix incorporates the solution of the Maxwell equations for one source. $\mathbf{K}$ allows the source to change direction, and also to be hidden (if $\mathbf{K}_{ij} = 0$)

This is a nice regression which we can handle.

No, not that easy: we don't know where the sources are, so we don't know $\mathbf{A}$ *and* we don't know $\mathbf{s}$.

All right, then, let's just allow many positions (grid up the whole brain)

Source localization, contd.

No, this is not a good idea: There are **many** more sources than data points

$$M >> N$$

This means that many different source configurations that could give the same data. (Configurations now essentially are activations, as positions have been fixed)

Can also see this from Hessian $\mathbf{H}$ of problem which will be rank-deficient

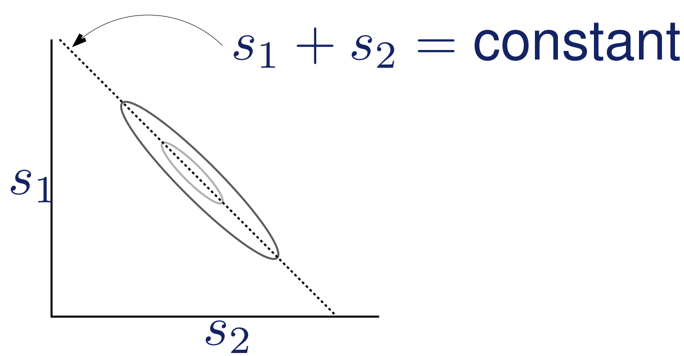

$$\begin{aligned}\hat{\mathbf{s}} &= \arg\min_{\mathbf{s}} ||\mathbf{y} - \mathbf{A}\mathbf{s}||^2 \\ &= \arg\min_{\mathbf{s}} (\mathbf{s}^T \mathbf{H} \mathbf{s} - 2\mathbf{s}^T \mathbf{A}^T \mathbf{y})\end{aligned}$$

Priors

We have to select one of these equivalent source configurations $\mathbf{s}$.

We want to choose these according to knowledge we may have from other experiments, other fields etc.

We encapsulate that knowledge in a prior.

The prior just makes sure we choose $\mathbf{s}$ that behave the way we want.

Dynamic causal modelling basically just imposes a sensible prior.

DCM = prior as a latent dynamical system

DCM assumes a set of source locations.

It specifies a *deterministic, dynamical* model for the evolution of the source activity in time

$$\frac{d\mathbf{s}}{dt} = F(\mathbf{s}_t, \mathbf{u}_t, \boldsymbol{\theta})$$

where $\mathbf{u}_t$ are inputs (eg +1 for stimulus on, 0 for stimulus off).

Given the $\mathbf{s}$, we know the sensor readings $\mathbf{y} = \mathbf{A}\mathbf{s}$.

However, $\mathbf{s}$ is not given, we have to infer it.

Overall, we have a latent dynamical system with parameters θ.

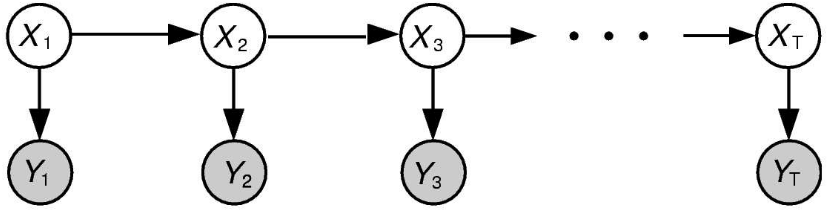

The DCM dynamical system

Assume there are a few ($\sim 5-10$) dipoles at spots of interest

For each of these, assume three neural populations: pyramidal cells, inhibitory interneurones and spiny stellate cells.

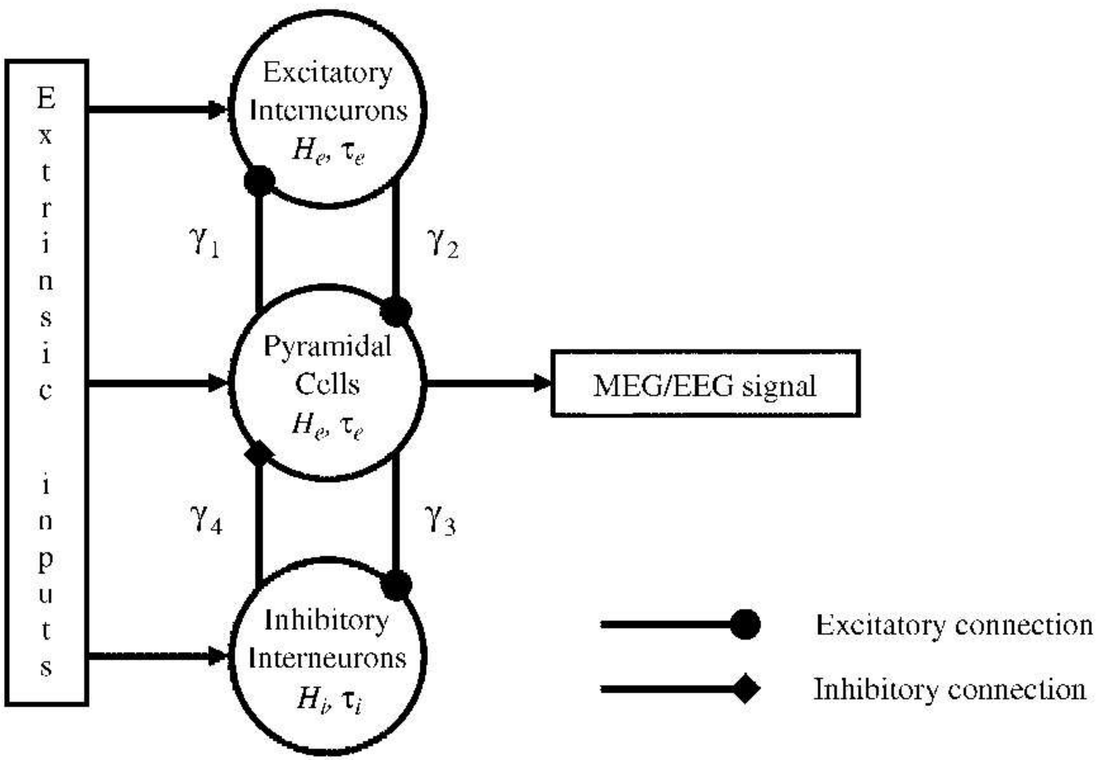

from David et al. (2005)

The DCM dynamical system

Why so many?

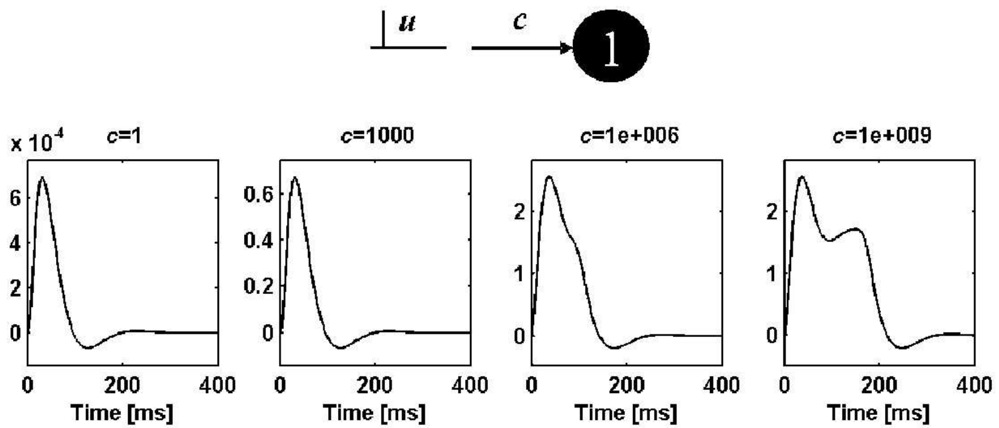

fromDavid et al. (2005)

The DCM dynamical system

Why so many?

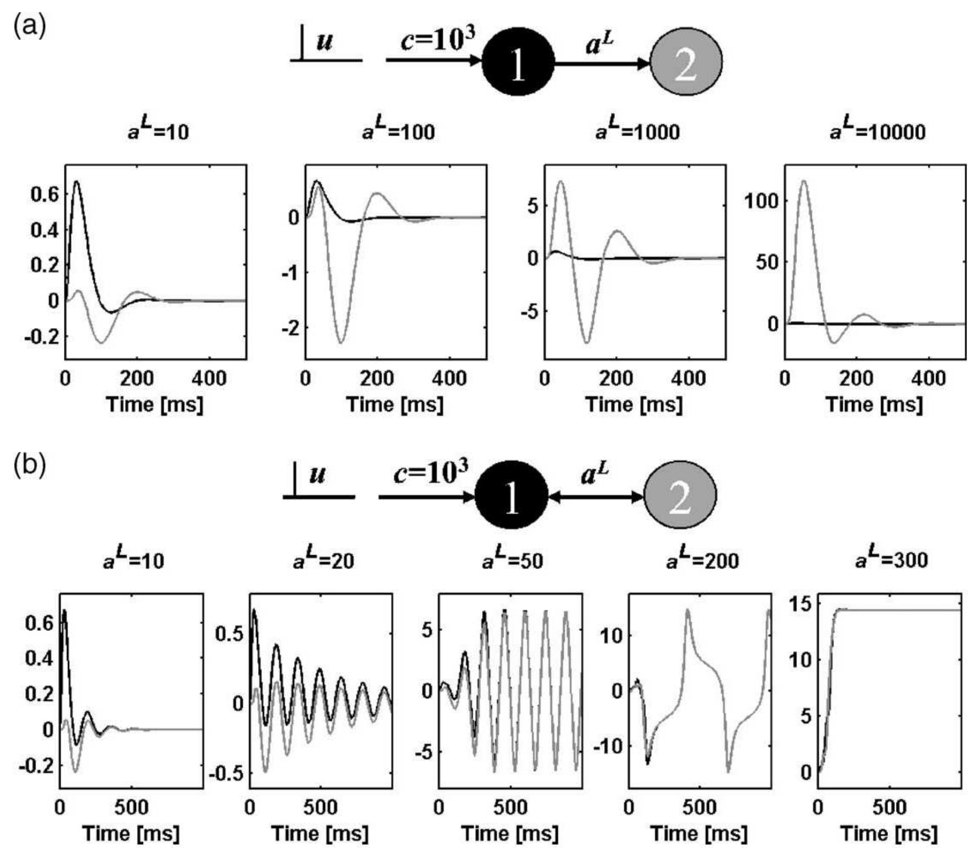

fromDavid et al. (2005)

DCM LDS

Ok, so let's run the hidden dynamical system and see what it'd generate

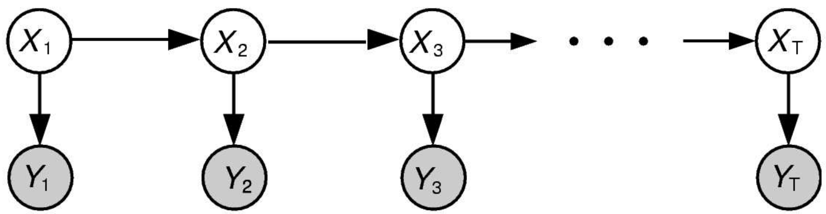

movie here

LDS Parametrisation

The parameters of interest will mainly be those related to effective connectivity:

Connection strengths between dipoles s_i and s_j in context k:

$$C_{ijk} = C_{ij} G_{ijk}.$$

The context k indicates for example whether the subject should be attending or not.

G_{ijk} is a gain matrix, which gets to modulate the interactions C_{ij} that are present in the baseline condition.

No other parameters are allowed to vary across contexts, so it's this modulation that has to explain any differences across conditions.

There are lots of other parameters, which are of less interest right now.

Inference in LDS

Inference in a LDS means we try to find the parameters that make the data most likely:

$$
\begin{aligned}
\hat{\theta} &= \arg\max_{\theta} p(\mathbf{y}|\theta) \\
&= \arg\max_{\theta} \int d\mathbf{s} p(\mathbf{y}, \mathbf{s}|\theta) \\
&= \arg\max_{\theta} \int d\mathbf{s} \underbrace{p(\mathbf{y}|\mathbf{s}, \theta)}_{\text{observation model}} \underbrace{p(\mathbf{s}|\theta)}_{\text{dynamical system}}
\end{aligned}
$$

Inference in LDS, contd.

$$\begin{aligned}
\hat{\theta} &= \arg\max_{\theta} p(\mathbf{y}|\theta) \\
&= \arg\max_{\theta} \log p(\mathbf{y}|\theta) \\
&= \arg\max_{\theta} \log \int d\mathbf{s} q(\mathbf{s}) \frac{p(\mathbf{y},\mathbf{s}|\theta)}{q(\mathbf{s})} \\
&\geq \arg\max_{\theta} \int d\mathbf{s} q(\mathbf{s}) \log \frac{p(\mathbf{y},\mathbf{s}|\theta)}{q(\mathbf{s})}
\end{aligned}$$

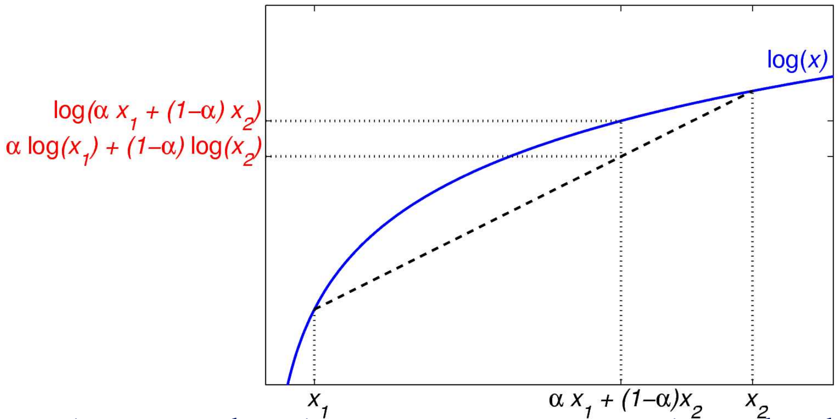

The last equation comes from Jensen's inequality.
$.5(\log(1) + \log(e)) = .5$
$< \log(.5(1+e)) = 1.85$

EM once again

Jensen's inequality provides us with a way of actually doing that maximisation. It proceeds in two steps.

In the M step, we assume we know the latents $\mathbf{s}$. Given those latents, it's now easy to infer the parameters. For example, we could infer $\mathbf{A}$ by setting $\mathbf{A} =< \mathbf{s}\mathbf{s}^T >^{-1}< \mathbf{s}\mathbf{y}^T >$ where the average is over time.

Note that we don't actually need to know $\mathbf{s}$. All we needed here was $< \mathbf{s}\mathbf{s}^T >$ and $< \mathbf{s}\mathbf{y}^T >$. This is generally true: We'll usually need some expected statistics (expectations of the latents) in the M step.

In the E step, we infer these expected statistics. In DCM, the latent dynamical system is deterministic, so given an initialization and the parameters, we just run it forward.

BUT: there are lots of local optima.

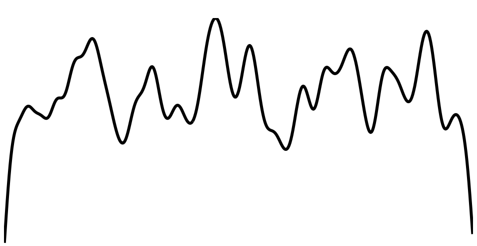

Discussion

DCM solves inversion by imposing various priors, in a way that allows us to access variables of interest.

It assumes only a few sources are present, and that their location is (at least approximately) known.

The main parameters of interest specify connection weights between these equivalent dipoles.

The observation model is linear.

There are many subtleties in the inference we've skipped, such as that the priors on the parameters have to be narrow, we need good initialisations.

References

Baillet, S., Mosher, J., and Leahy, R. (2001). Electromagnetic brain mapping. *Signal Processing Magazine, IEEE*, 18(6):14–30.

David, O., Harrison, L., and Friston, K. J. (2005). Modelling event-related responses in the brain. *Neuroimage*, 25(3):756–770.

David, O., Kiebel, S. J., Harrison, L. M., Mattout, J., Kilner, J. M., and Friston, K. J. (2006a). Dynamic causal modeling of evoked responses in eeg and meg. *Neuroimage*, 30(4):1255–1272.

David, O., Kilner, J. M., and Friston, K. J. (2006b). Mechanisms of evoked and induced responses in meg/eeg. *Neuroimage*.

Friston, K. J., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302.

Kiebel, S. J., David, O., and Friston, K. J. (2006). Dynamic causal modelling of evoked responses in eeg/meg with lead field parameterization. *Neuroimage*, 30(4):1273–1284.