M/EEG Source Analysis (in a Bayesian Flavour)

Ryan Timms
@blobsonthebrain

Overview

- Intro: What is the M/EEG inverse problem?
- Unifying all M/EEG inversion algorithms with prior assumptions
- Multiple Sparse Priors
- Validating source inversion attempts with model evidence

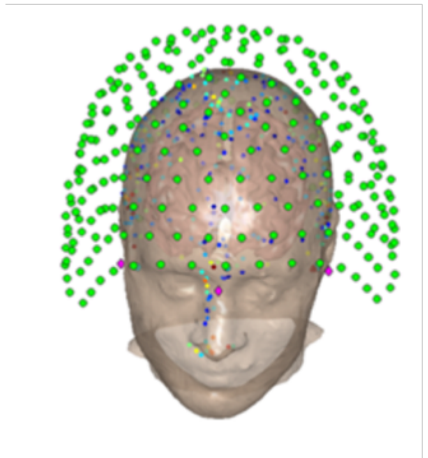

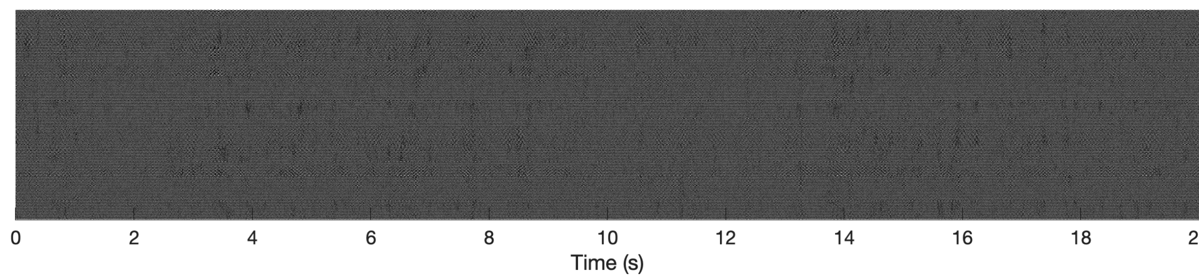

$N_{channels} \approx 300$

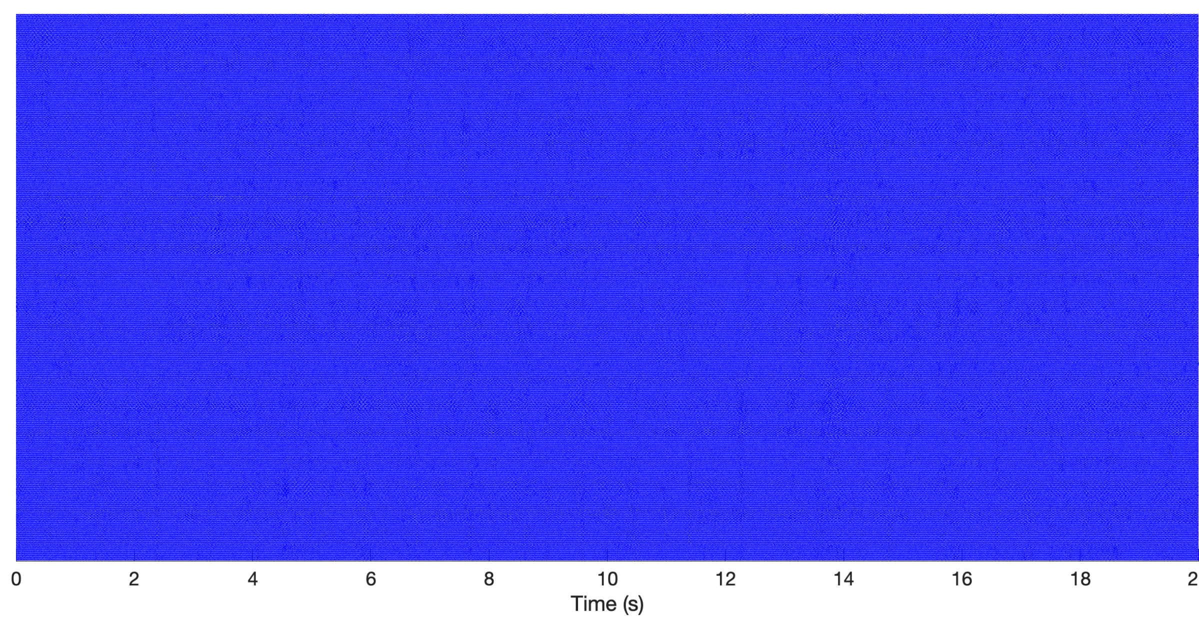

$N_{sources} \approx 10{,}000$

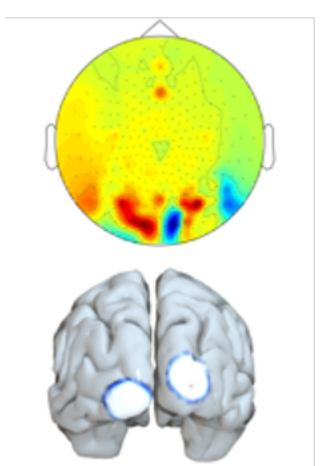

Interpretability ↑
SNR ↑
Spatial resolution ↑
Clinical viability ↑
Richer analyses facilitated

No unique solutions!
# of sources >> # of sensors
Forward model cannot (trivially) be inverted

Problem statement

Hard!

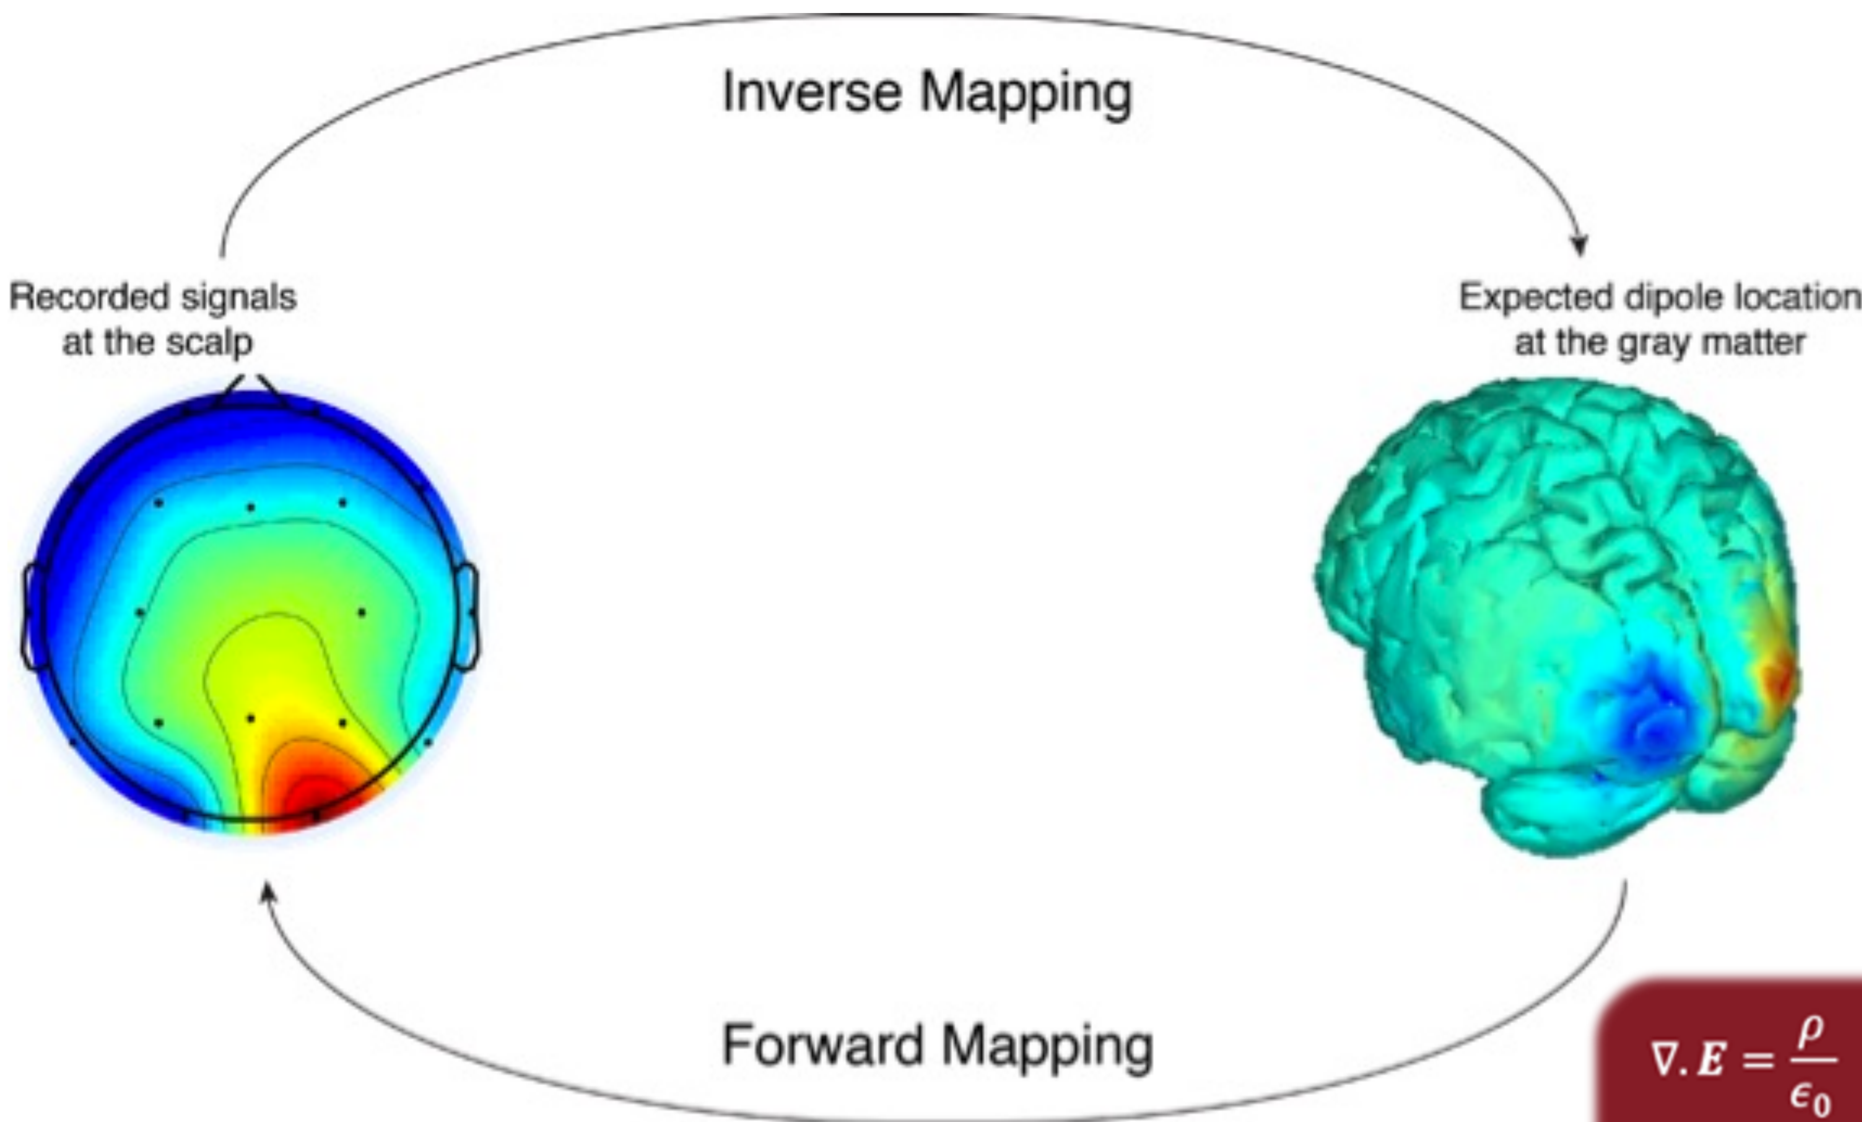

Easy!

$$\nabla . \boldsymbol{E} = \frac{\rho}{\epsilon_0}$$

$$\nabla . \boldsymbol{B} = 0$$

$$\nabla \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t}$$

$$\nabla \times \boldsymbol{B} = \mu_0 \boldsymbol{J} + \mu_0 \epsilon_0 \frac{\partial \boldsymbol{E}}{\partial t}$$

Figure credit: neuroelectrics.com

Formulation

$$\mathbf{Y} = \mathbf{HX} + \mathbf{E}$$

where:

$\mathbf{Y}$ - the sensor level data, sensors x time

$\mathbf{H}$ - the lead field matrix, sensors x sources

$\mathbf{X}$ - the neural signals, sources x time

$\mathbf{E}$ - non-brain signals, sensors x time

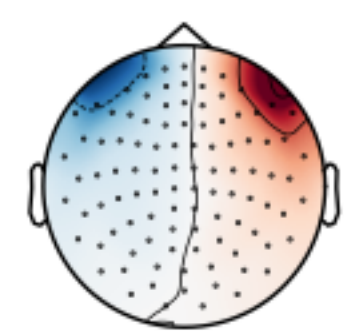

The lead field matrix

In MEG:

- Head model choice (single sphere, multiple spheres or single shell)

- Assumptions about the signal generators (layer V pyramidal neurons)

- Can be constrained to sources oriented perpendicular to cortex

- Physics - function of sensor/source displacement ***and*** orientation

In EEG:

- All of the above + choices about conduction models. Generally the same, but harder

Sensor/source position

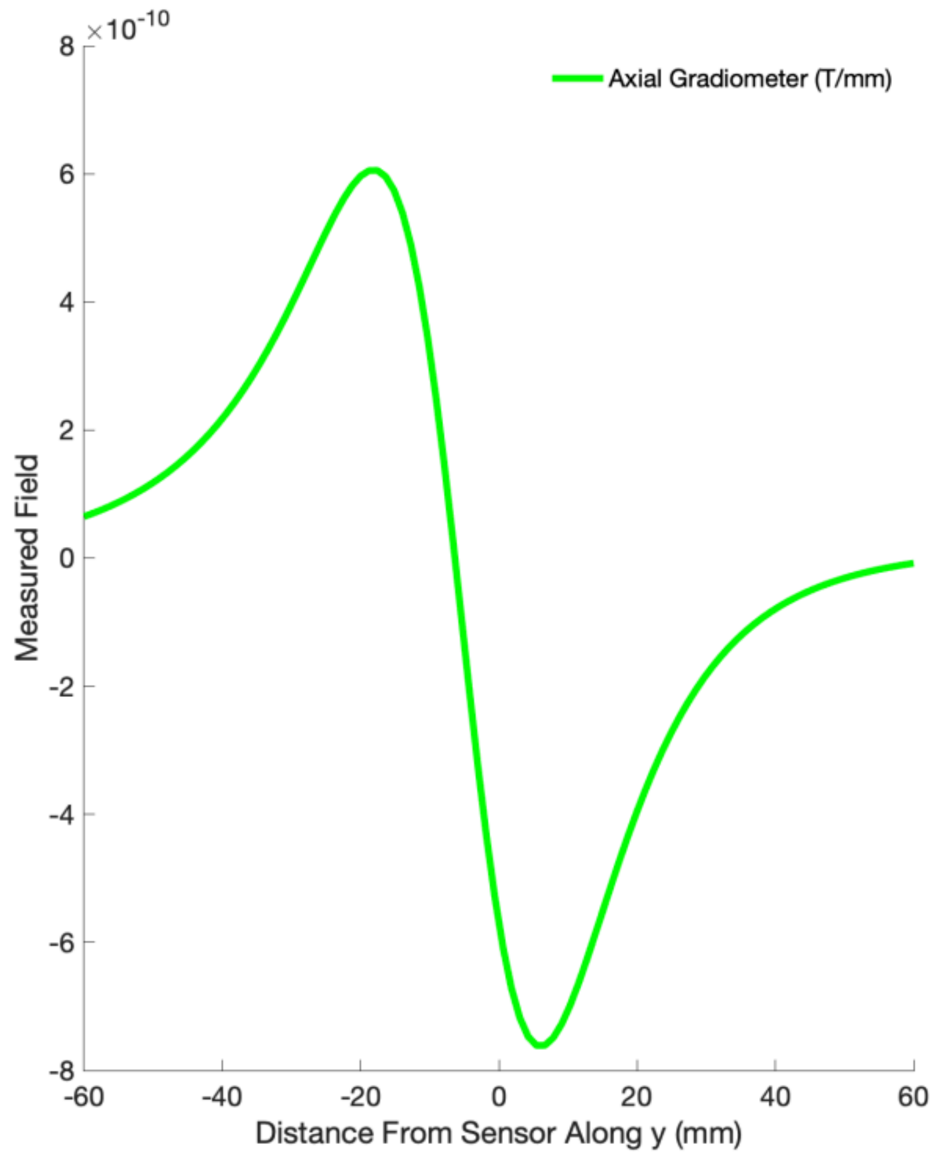

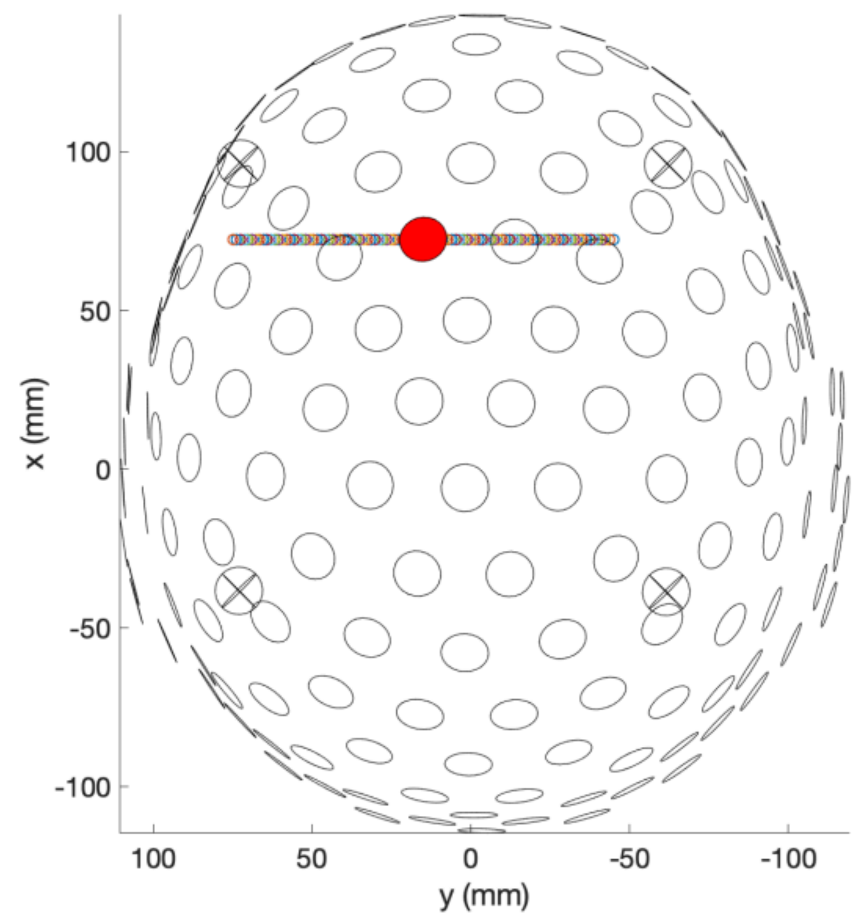

Sensor/source orientation

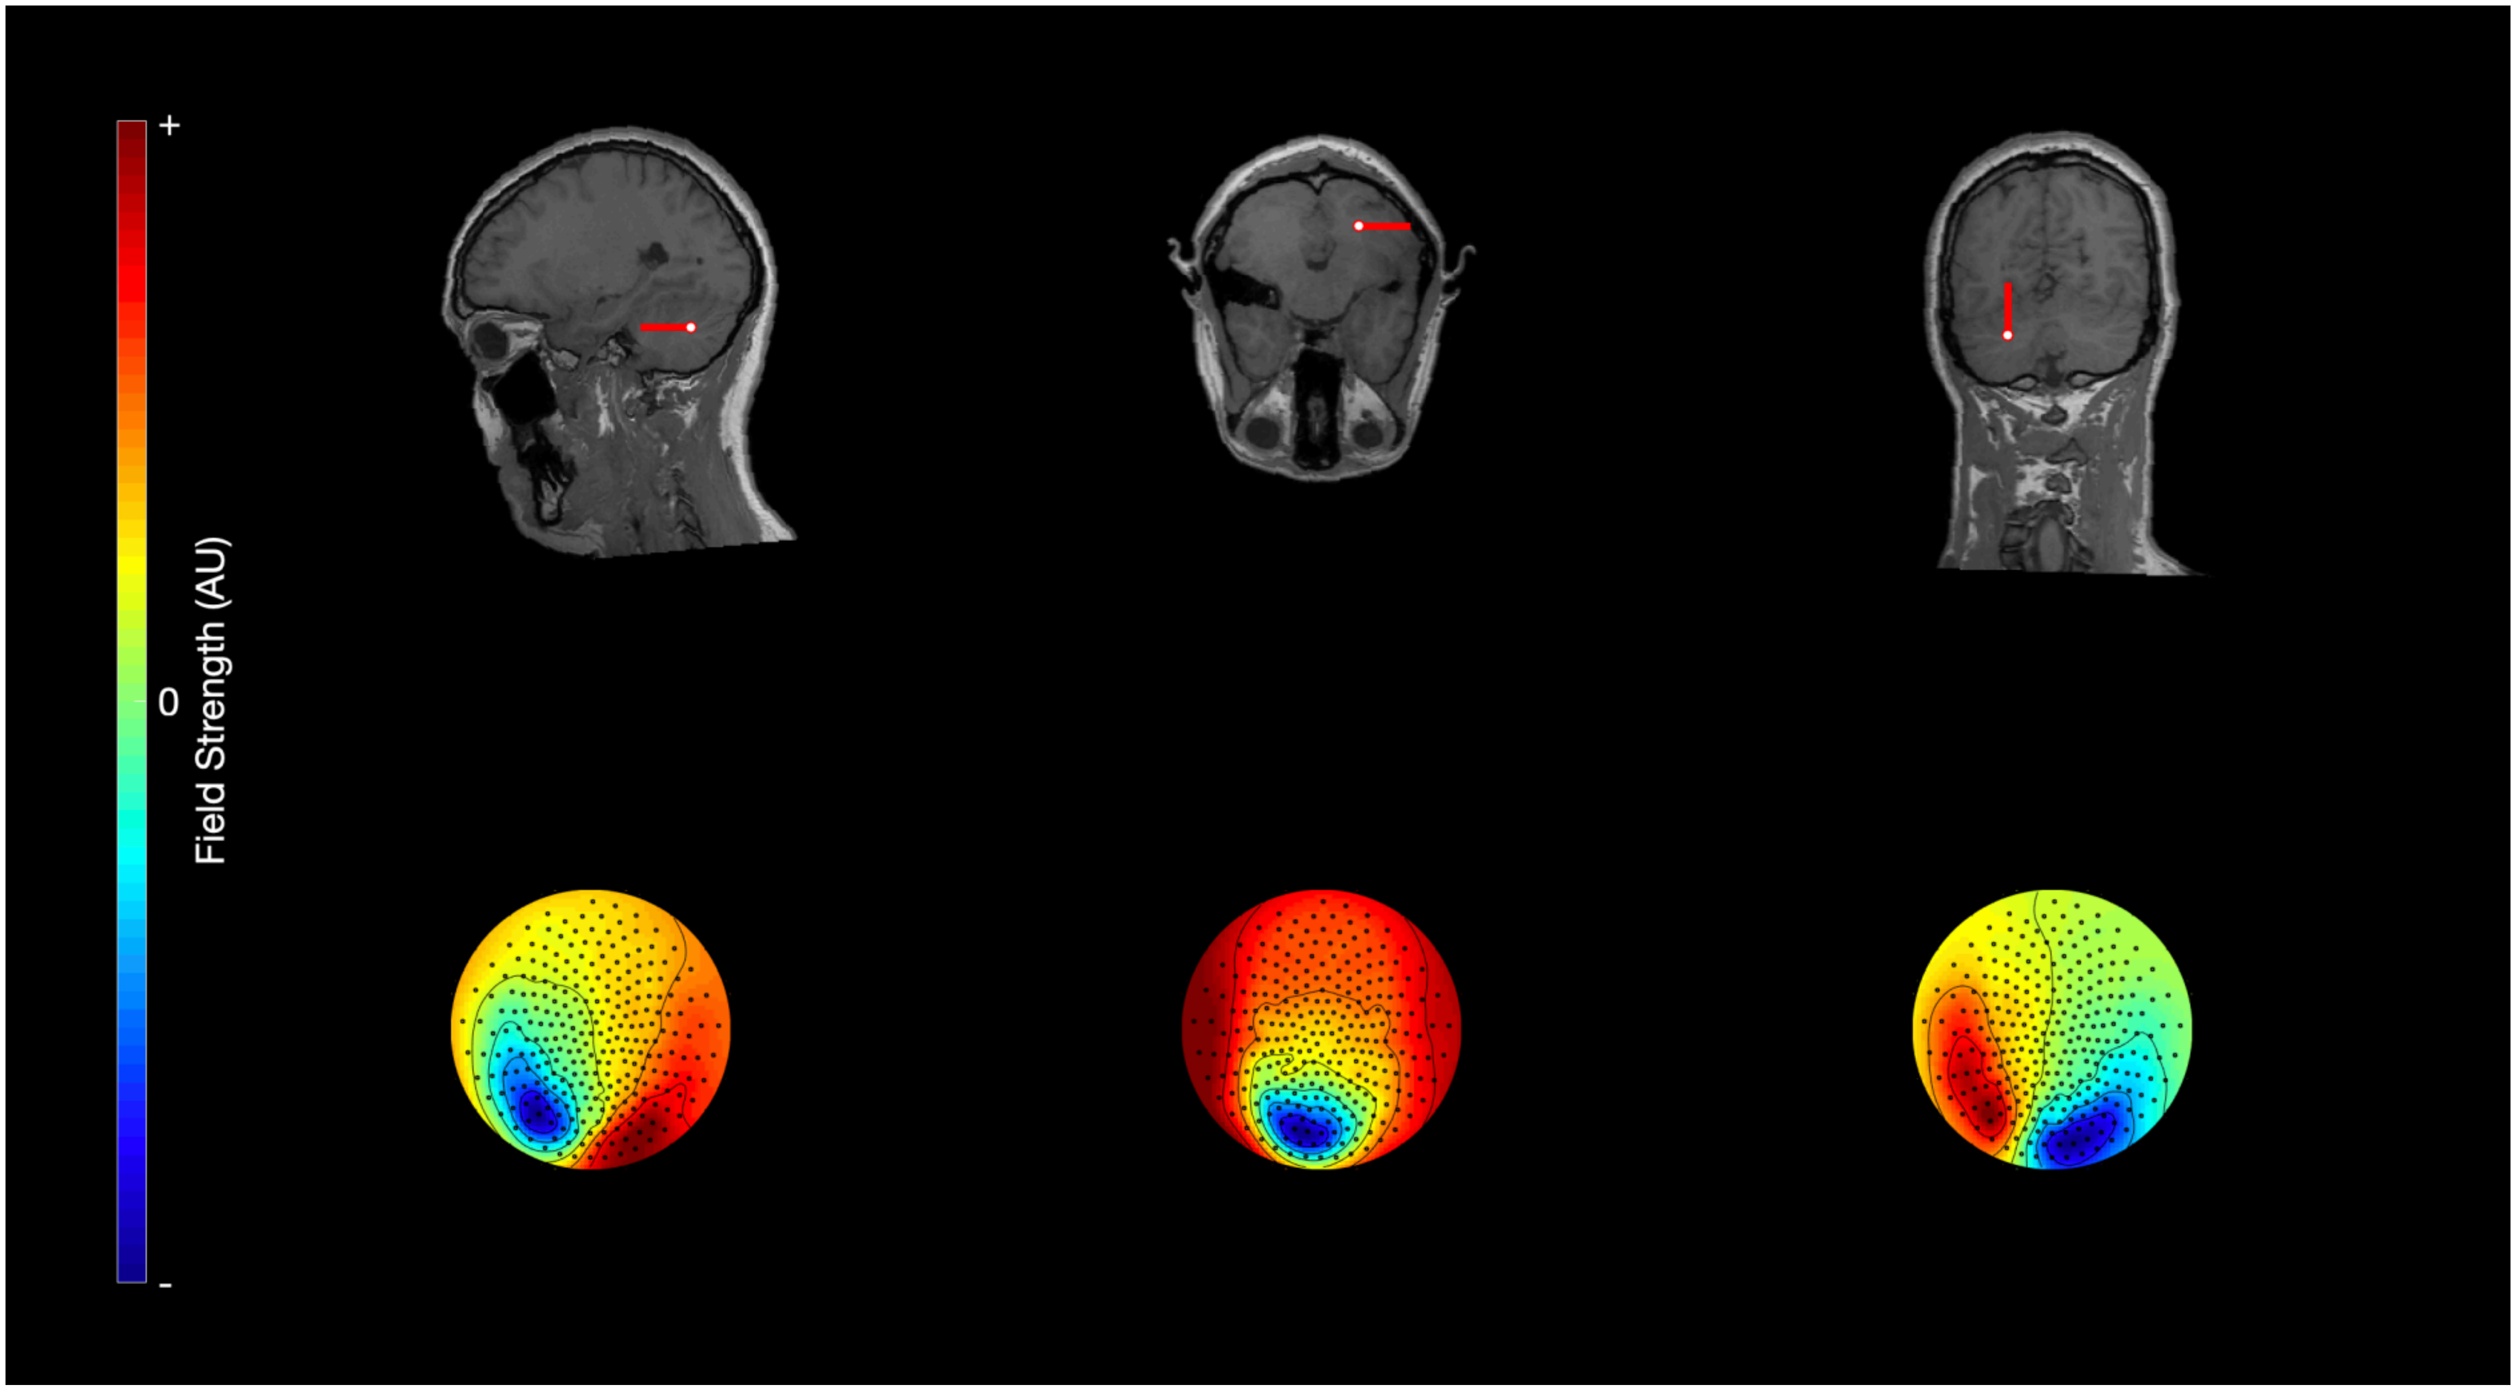

The lead field matrix

Once computed, the lead field is our projector from each source in the brain to each of the sensors.

In words, each column of the lead field matrix tells us what magnetic field we would expect to measure if a source was active at that part of the brain, given a fixed orientation.

$N_{sources} \approx 10{,}000$

$N_{channels} \approx 300$

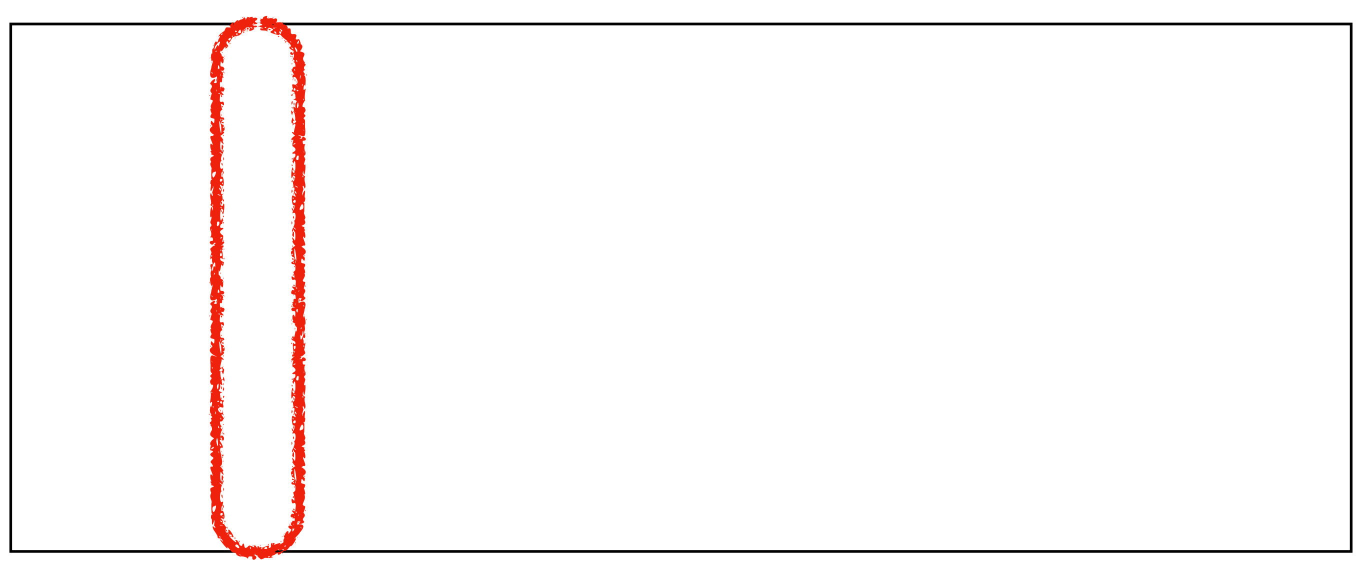

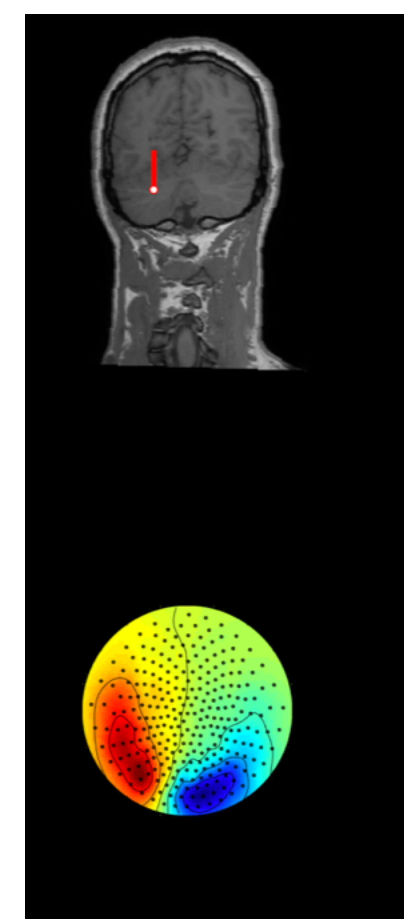

Source reconstruction

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathbf{E}$$

$$\hat{\mathbf{X}} = \mathbf{H}^{-1}\mathbf{Y}\ ?$$

No!

$N_{sources} >> N_{channels}$

$\mathbf{H}$ isn't square

$\mathbf{H}$ has a maximum rank of $N_{channels}$

Source reconstruction

Three philosophies to proceed

1) Dipole fits: I want to get around the ill-posed problem

2) Tomographic approaches: I want to reconstruct activity around the whole brain and explain variance

a.k.a Bayesian approaches

3) Spatial filters (beamformers): I want to reconstruct activity at a set of locations, based upon some other mathematical criterion

All of these approaches introduce some form of assumptions about the neural generators

Bayesian Formulation

$$p(X|Y) = \frac{p(Y|X)p(X)}{p(Y)}$$

Likelihood: $p(Y|X)$

Prior: $p(X)$

Evidence: $p(Y)$

Posterior: $p(X|Y)$

Bayesian Formulation

$$p(X|Y) = \frac{p(Y|X)p(X)}{p(Y)}$$

Likelihood: $p(Y|X)$

Prior: $p(X)$

Evidence: $p(Y)$

Posterior: $p(X|Y)$

For us:
Y are the recorded MEG data
X are the source currents

Bayesian Formulation

$$p(X \mid Y) \propto p(Y \mid X)p(X)$$

Likelihood: $p(Y \mid X)$

Prior: $p(X)$

~~Evidence: $p(Y)$~~

Posterior: $p(X \mid Y)$

Bayesian Formulation

$$p(X|Y) \propto p(Y|X)p(X)$$

Likelihood: $p(Y|X) = MVN(Y|HX, C_N)$

Prior: $p(X) = MVN(X|0, C_X)$

Posterior:

$$p(X|Y) \propto exp\left[-0.5(HX - Y)^T C_N^{-1}(HX - Y) - 0.5X^T C_X^{-1} X\right]$$

Bayesian Formulation

Posterior: $p(X \mid Y) \propto exp\left[-0.5(HX - Y)^T C_N^{-1}(HX - Y) - 0.5X^T C_X^{-1} X\right]$

Doing some maths (take the log of the above expression, and differentiating with respect to X), we find that the *maximum a posteriori* solution is given simply by

$$\hat{X} = C_X H^T \left[C_n + HC_X H^T\right]^{-1} Y$$

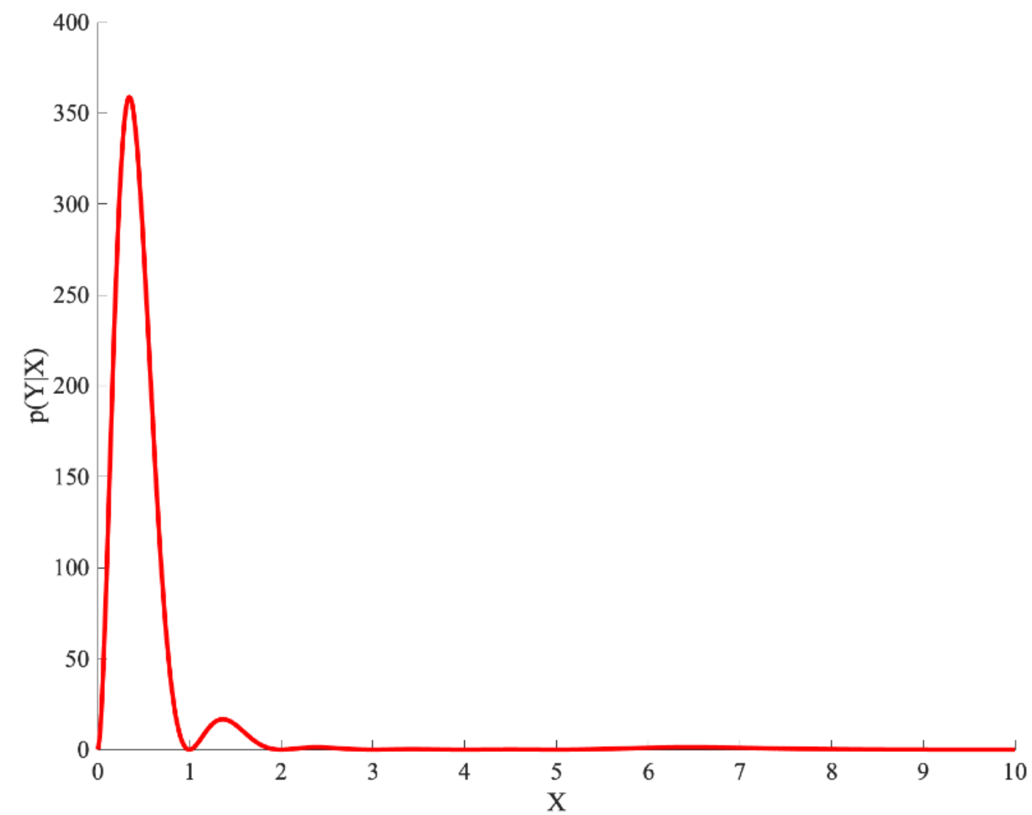

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

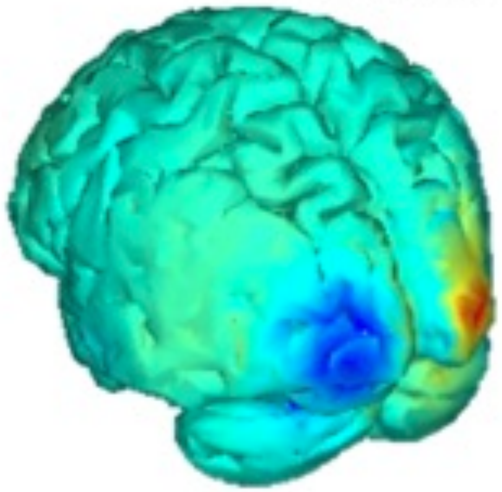

Current estimates = f (Data covariance, Forward model, Recorded data)

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T \right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

Prior noise covariance matrix

Empty room *Pre-stimulus baseline* *IID*

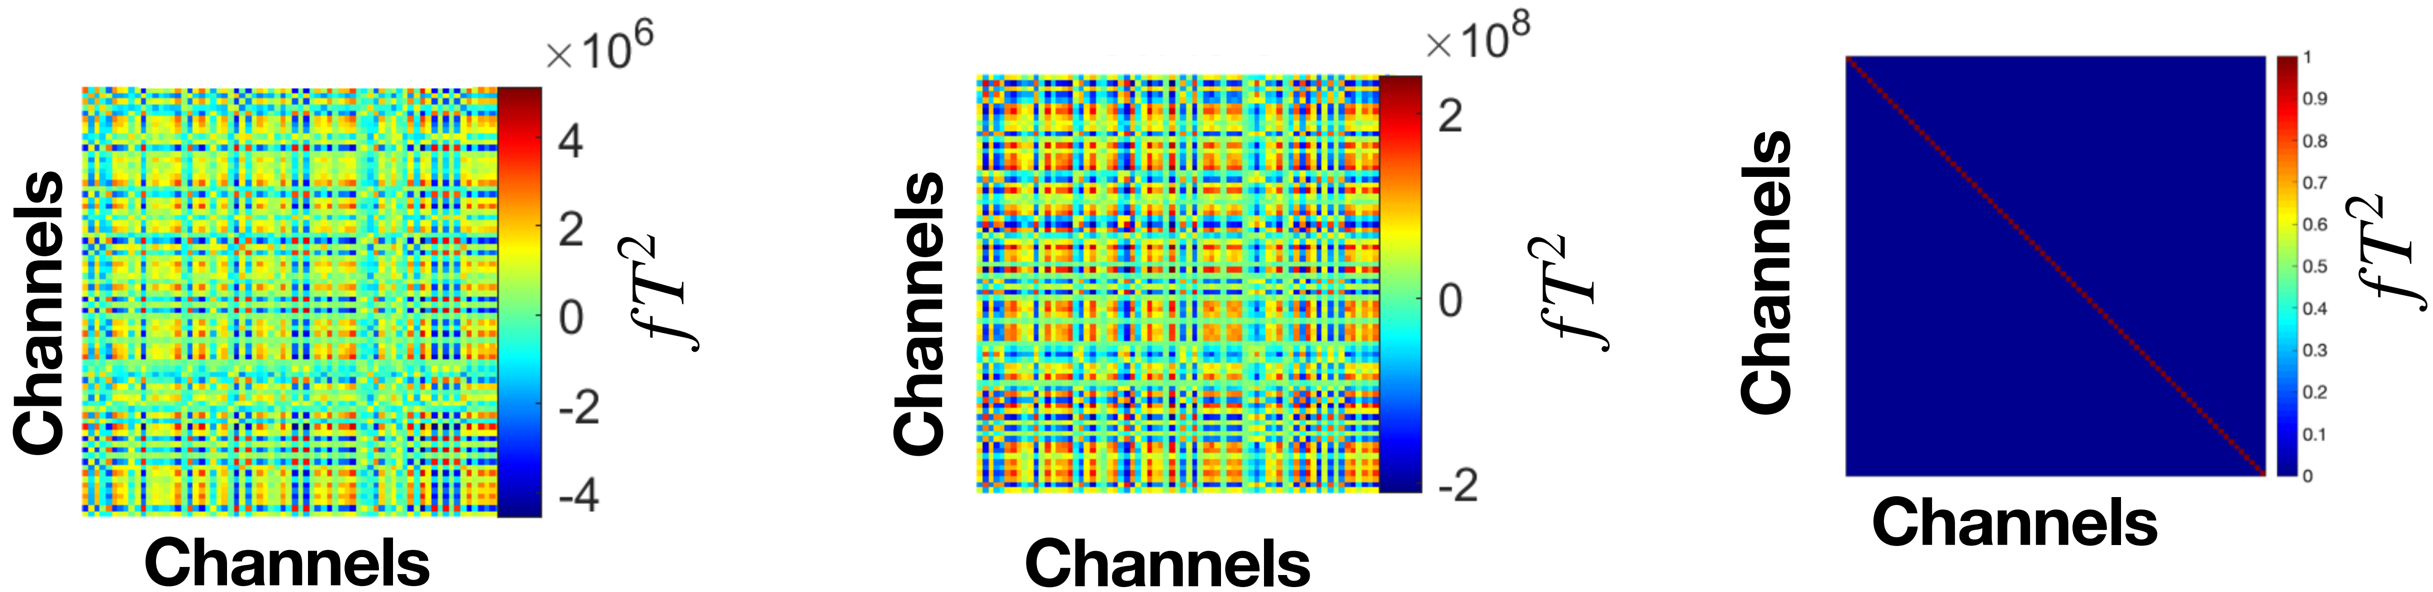

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T \right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

Source prior covariance matrix

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T \right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

Source prior covariance matrix

“What parts of the brain do I think are active during my recording?”

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

Source prior covariance matrix

$$C_X \in \mathscr{R}^{N_{sources} \times N_{sources}}$$

…big!

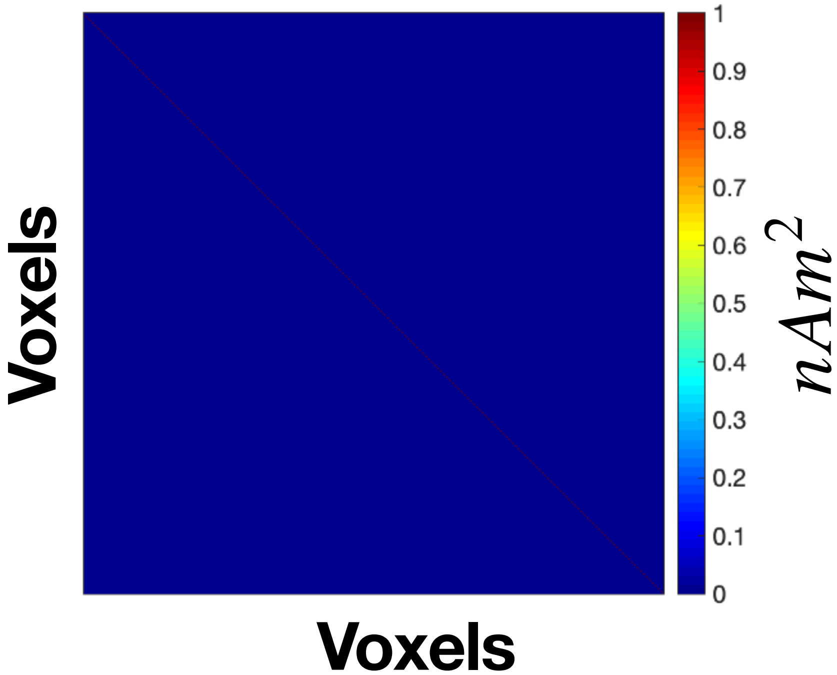

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

Source prior covariance matrix

IID, “minimum norm”

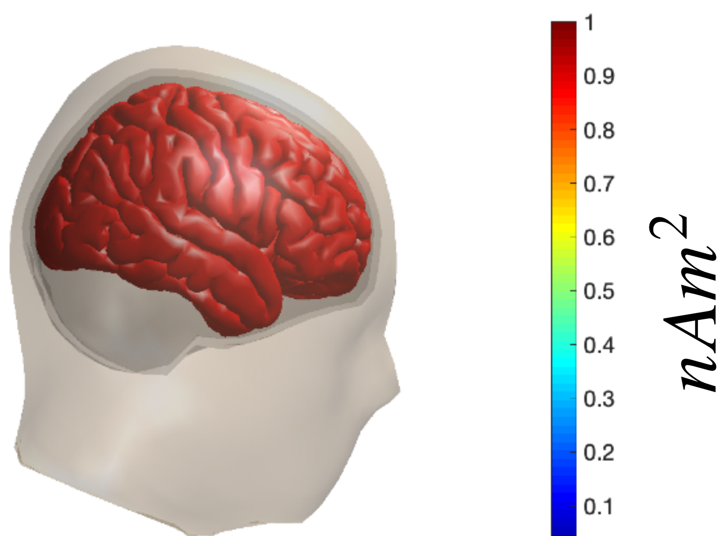

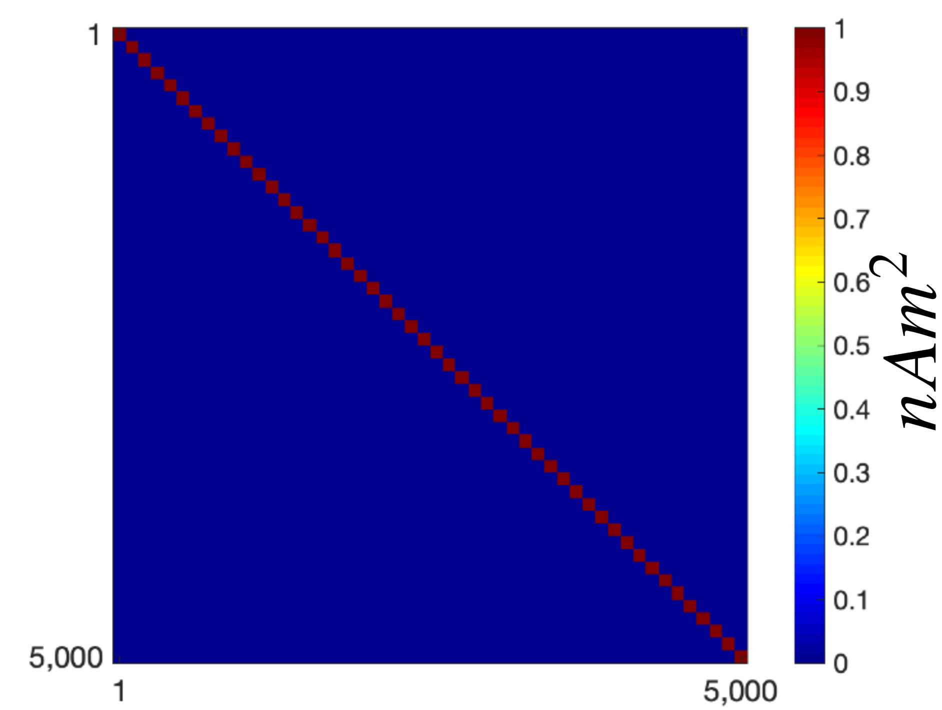

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

Source prior covariance matrix

Dipole fit

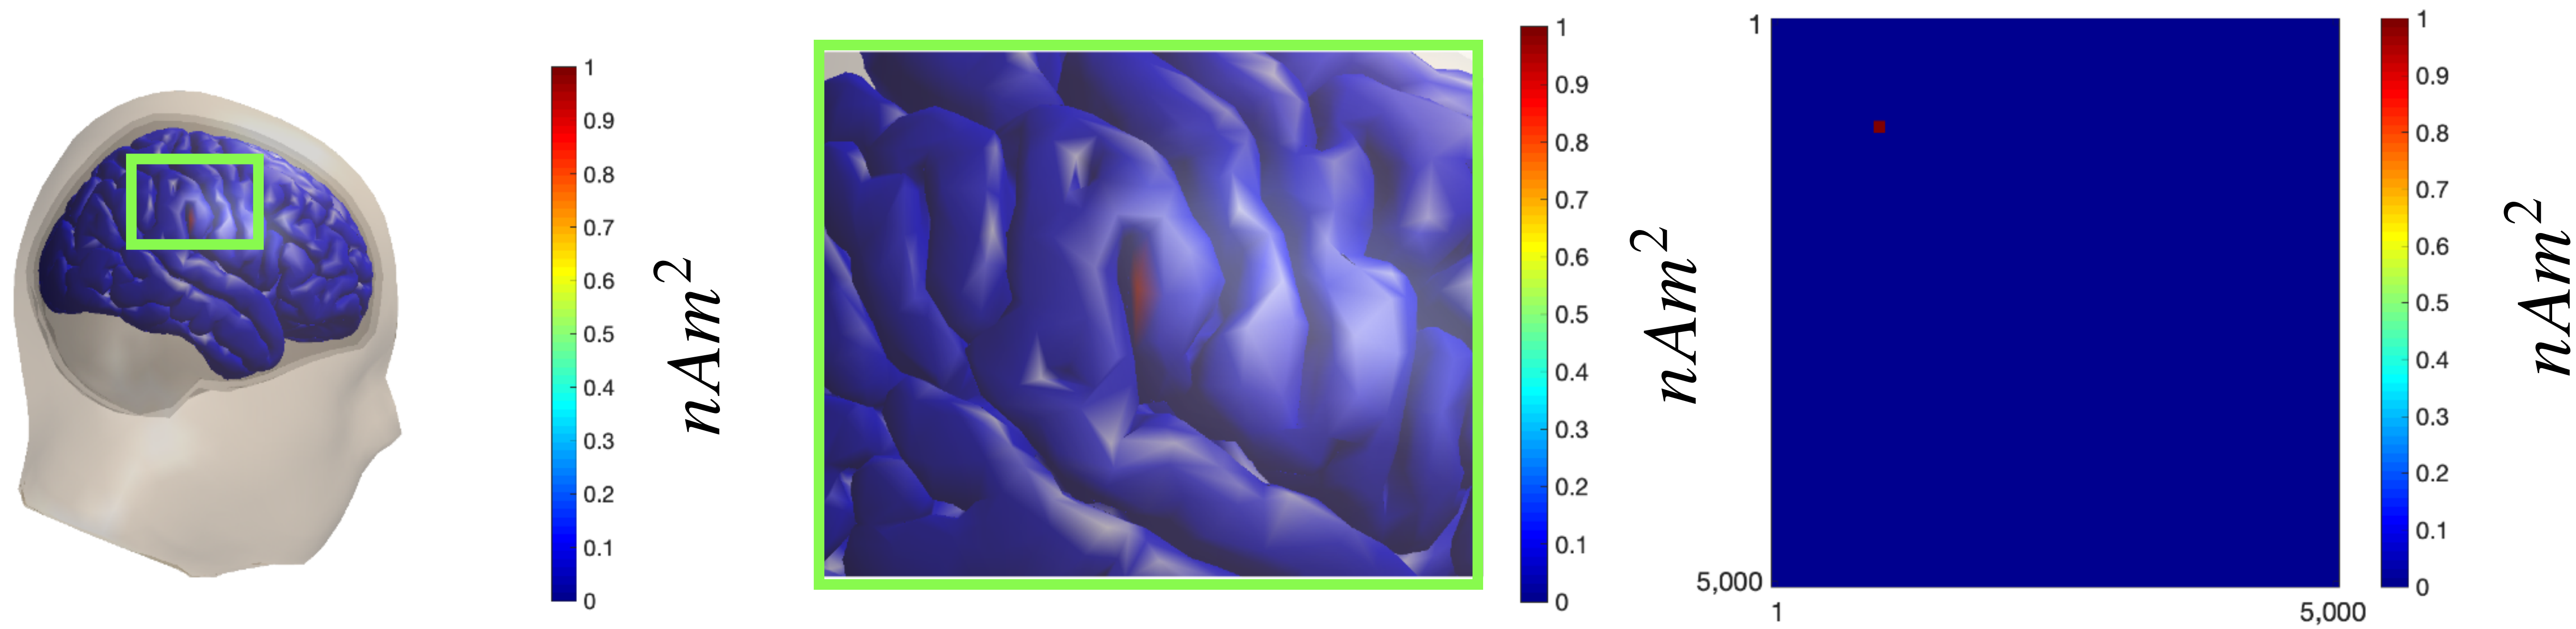

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

Source prior covariance matrix

"If my neighbour is active, I am also likely to be active"

eLORETA/sLORETA/local coherence

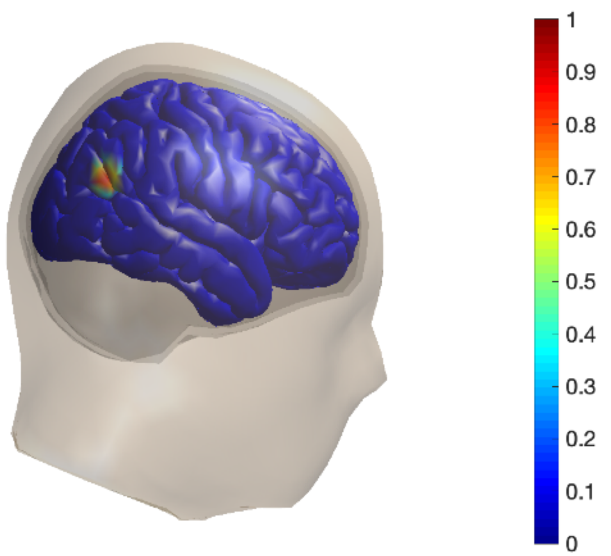

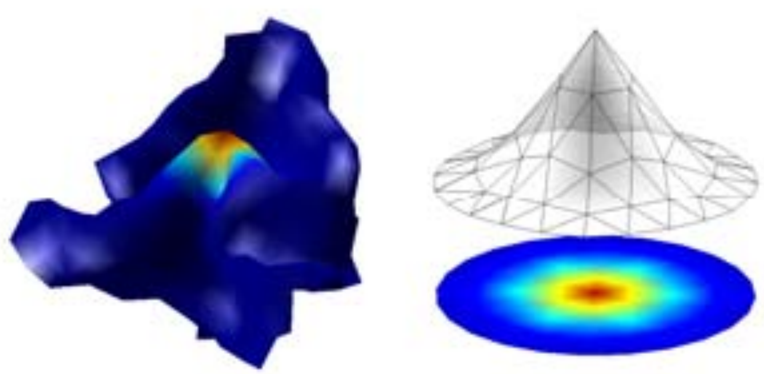

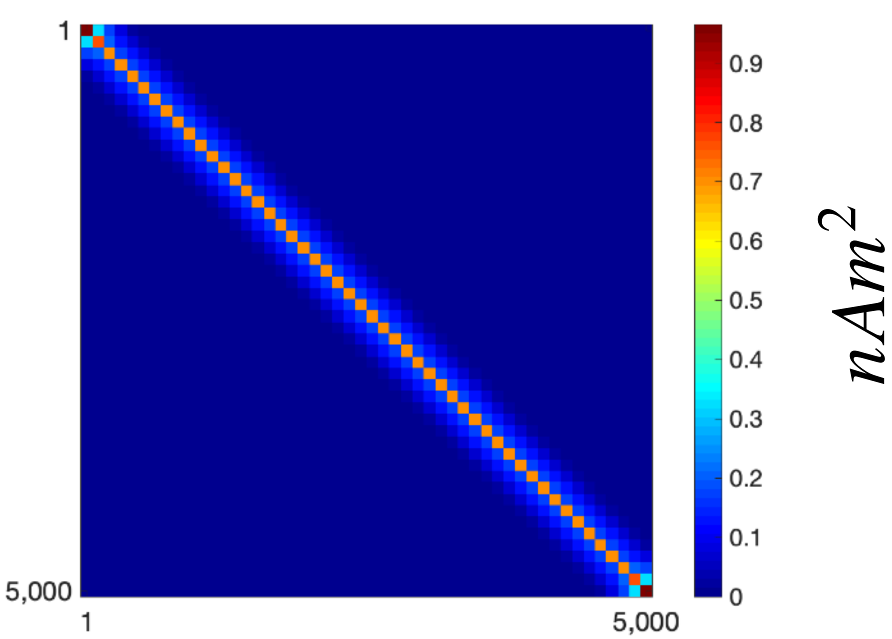

Dipole fit

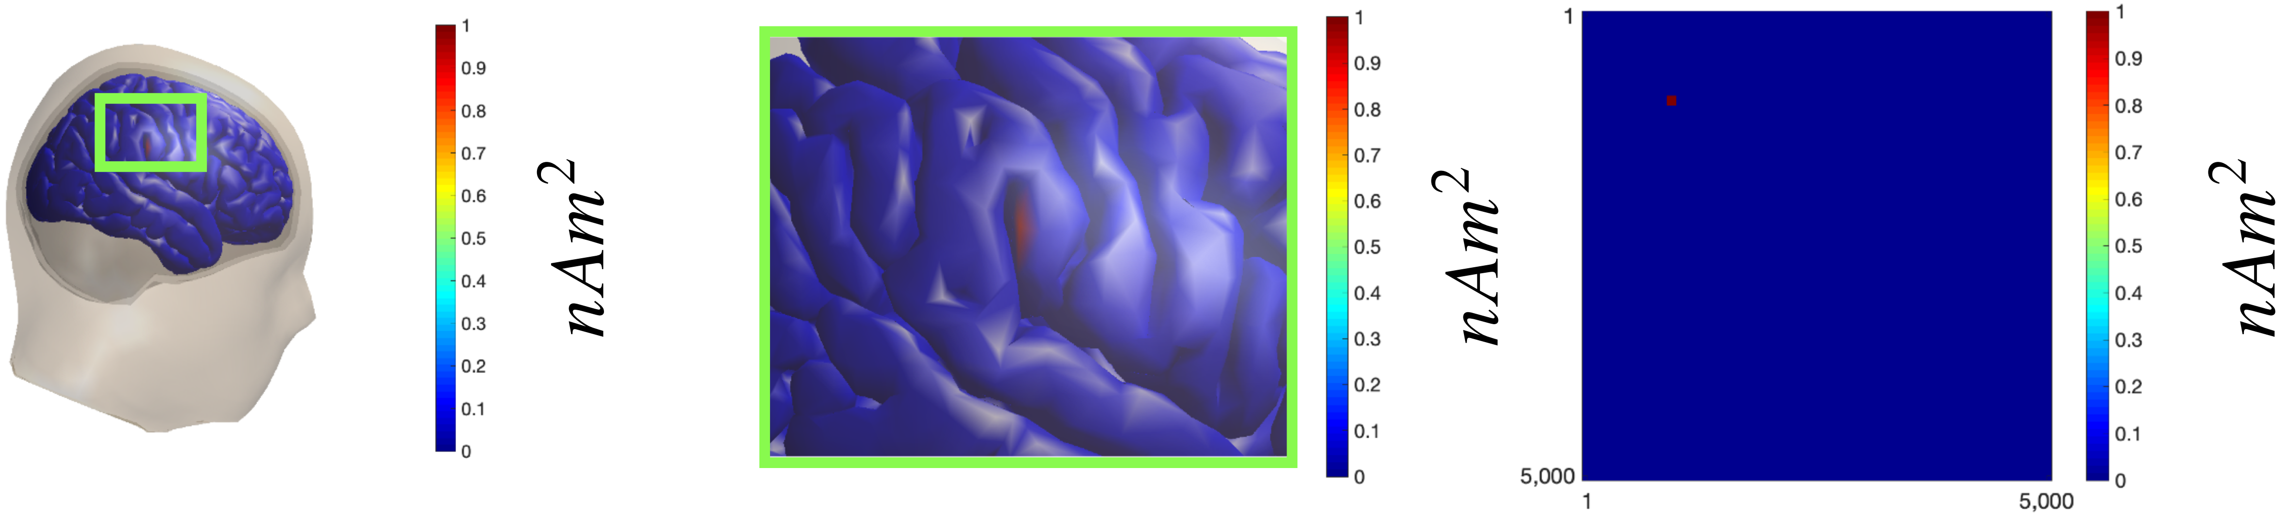

IID, "minimum norm"

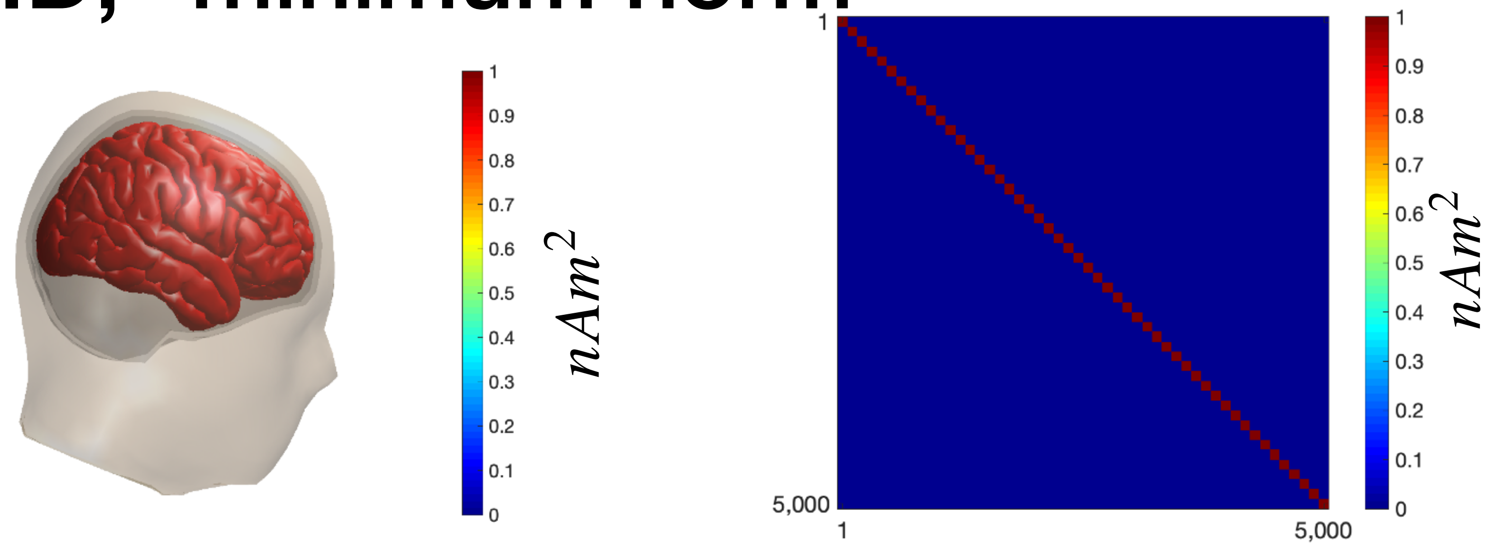

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T \right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

$$\mathbf{W} = f(\mathbf{H})$$

eLORETA/sLORETA/local coherence

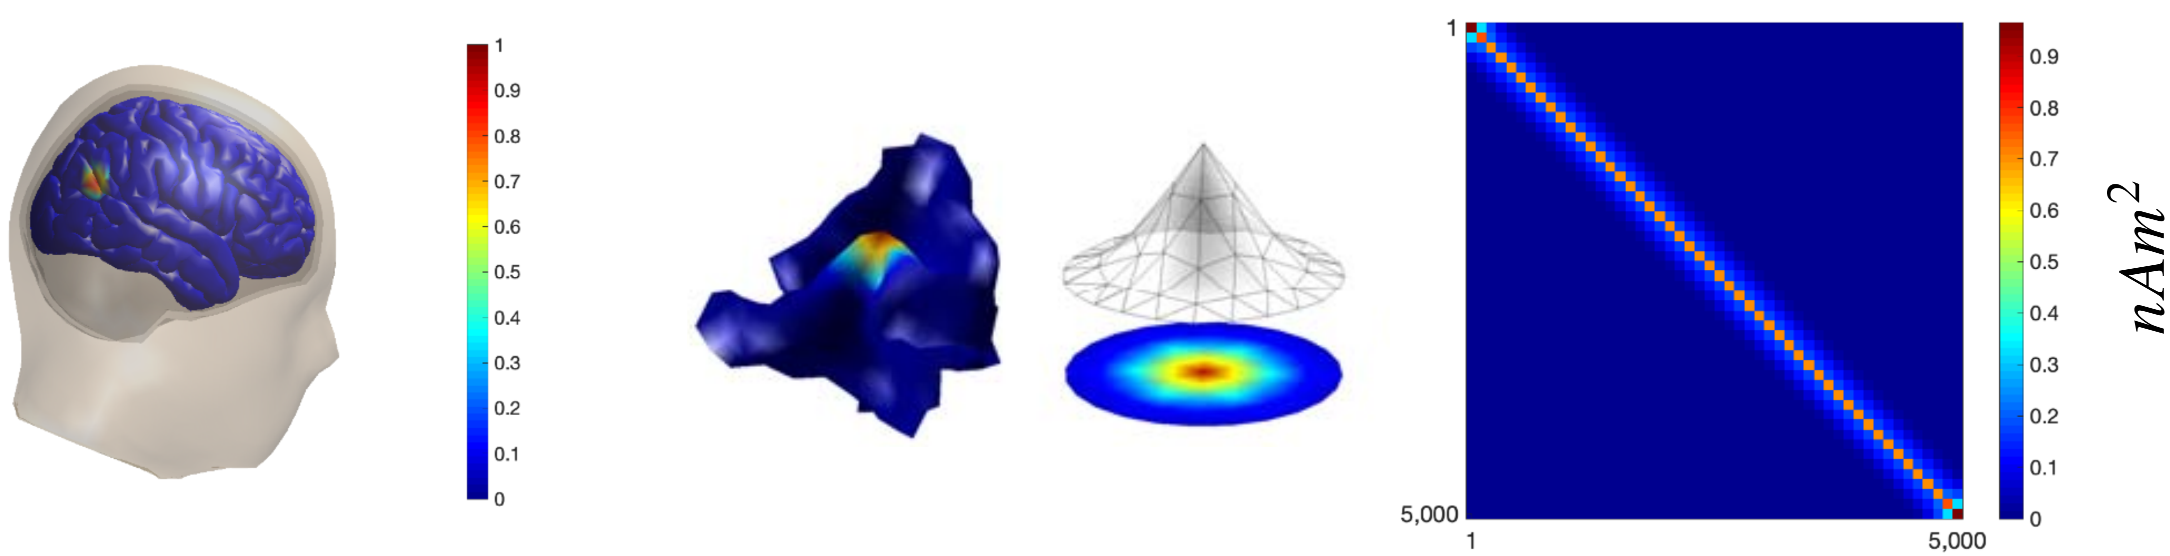

Multiple Sparse Priors

Dipole fit for dipole at vertex a

$$C_{X,ij} = \begin{cases} 1 & \text{if } i, j = a \\ 0 & otherwise \end{cases}$$

IID, "minimum norm"

$$C_X = I$$

eLORETA/sLORETA/local coherence

$$C_X = exp(\sigma G_L)$$

G_L is from the graph Laplacian of the cortical mesh, i.e. distances between vertices

σ controls the smoothness of the source space

Source prior covariance matrix

$$C_X = \sum_{i}^{K} \alpha_i \beta_i$$

Source prior covariance matrix

$$C_X = \sum_{i}^{K} \alpha_i \beta_i$$

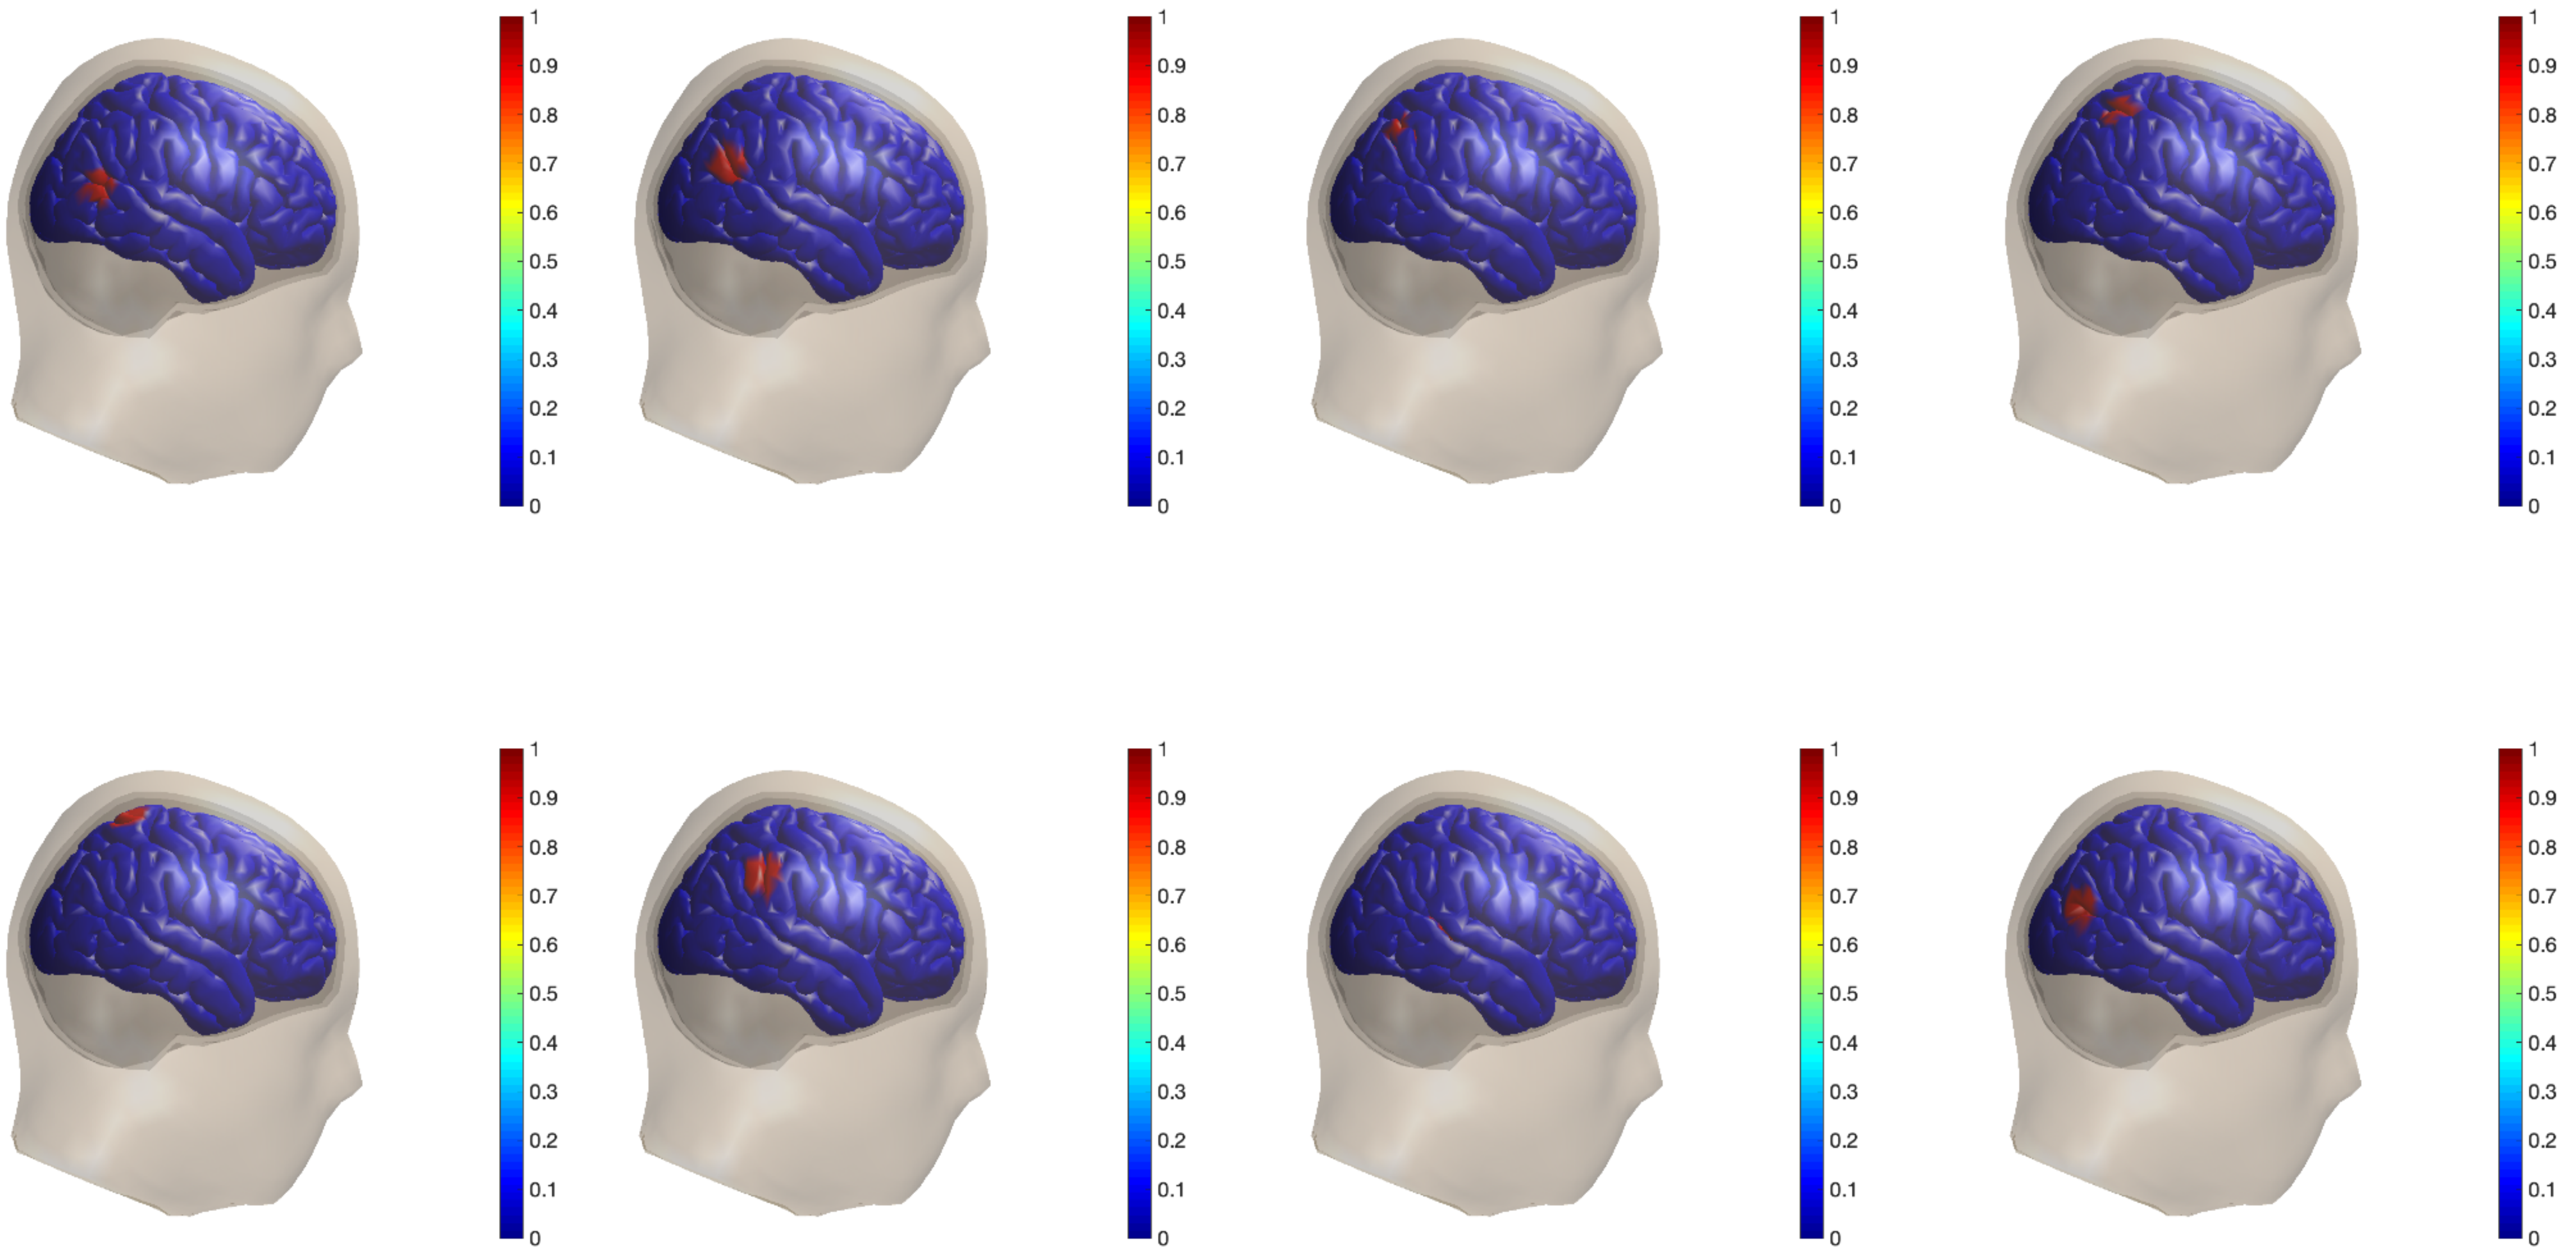

Source prior covariance matrix

$$C_X = \sum_{i}^{K} \alpha_i \beta_i$$

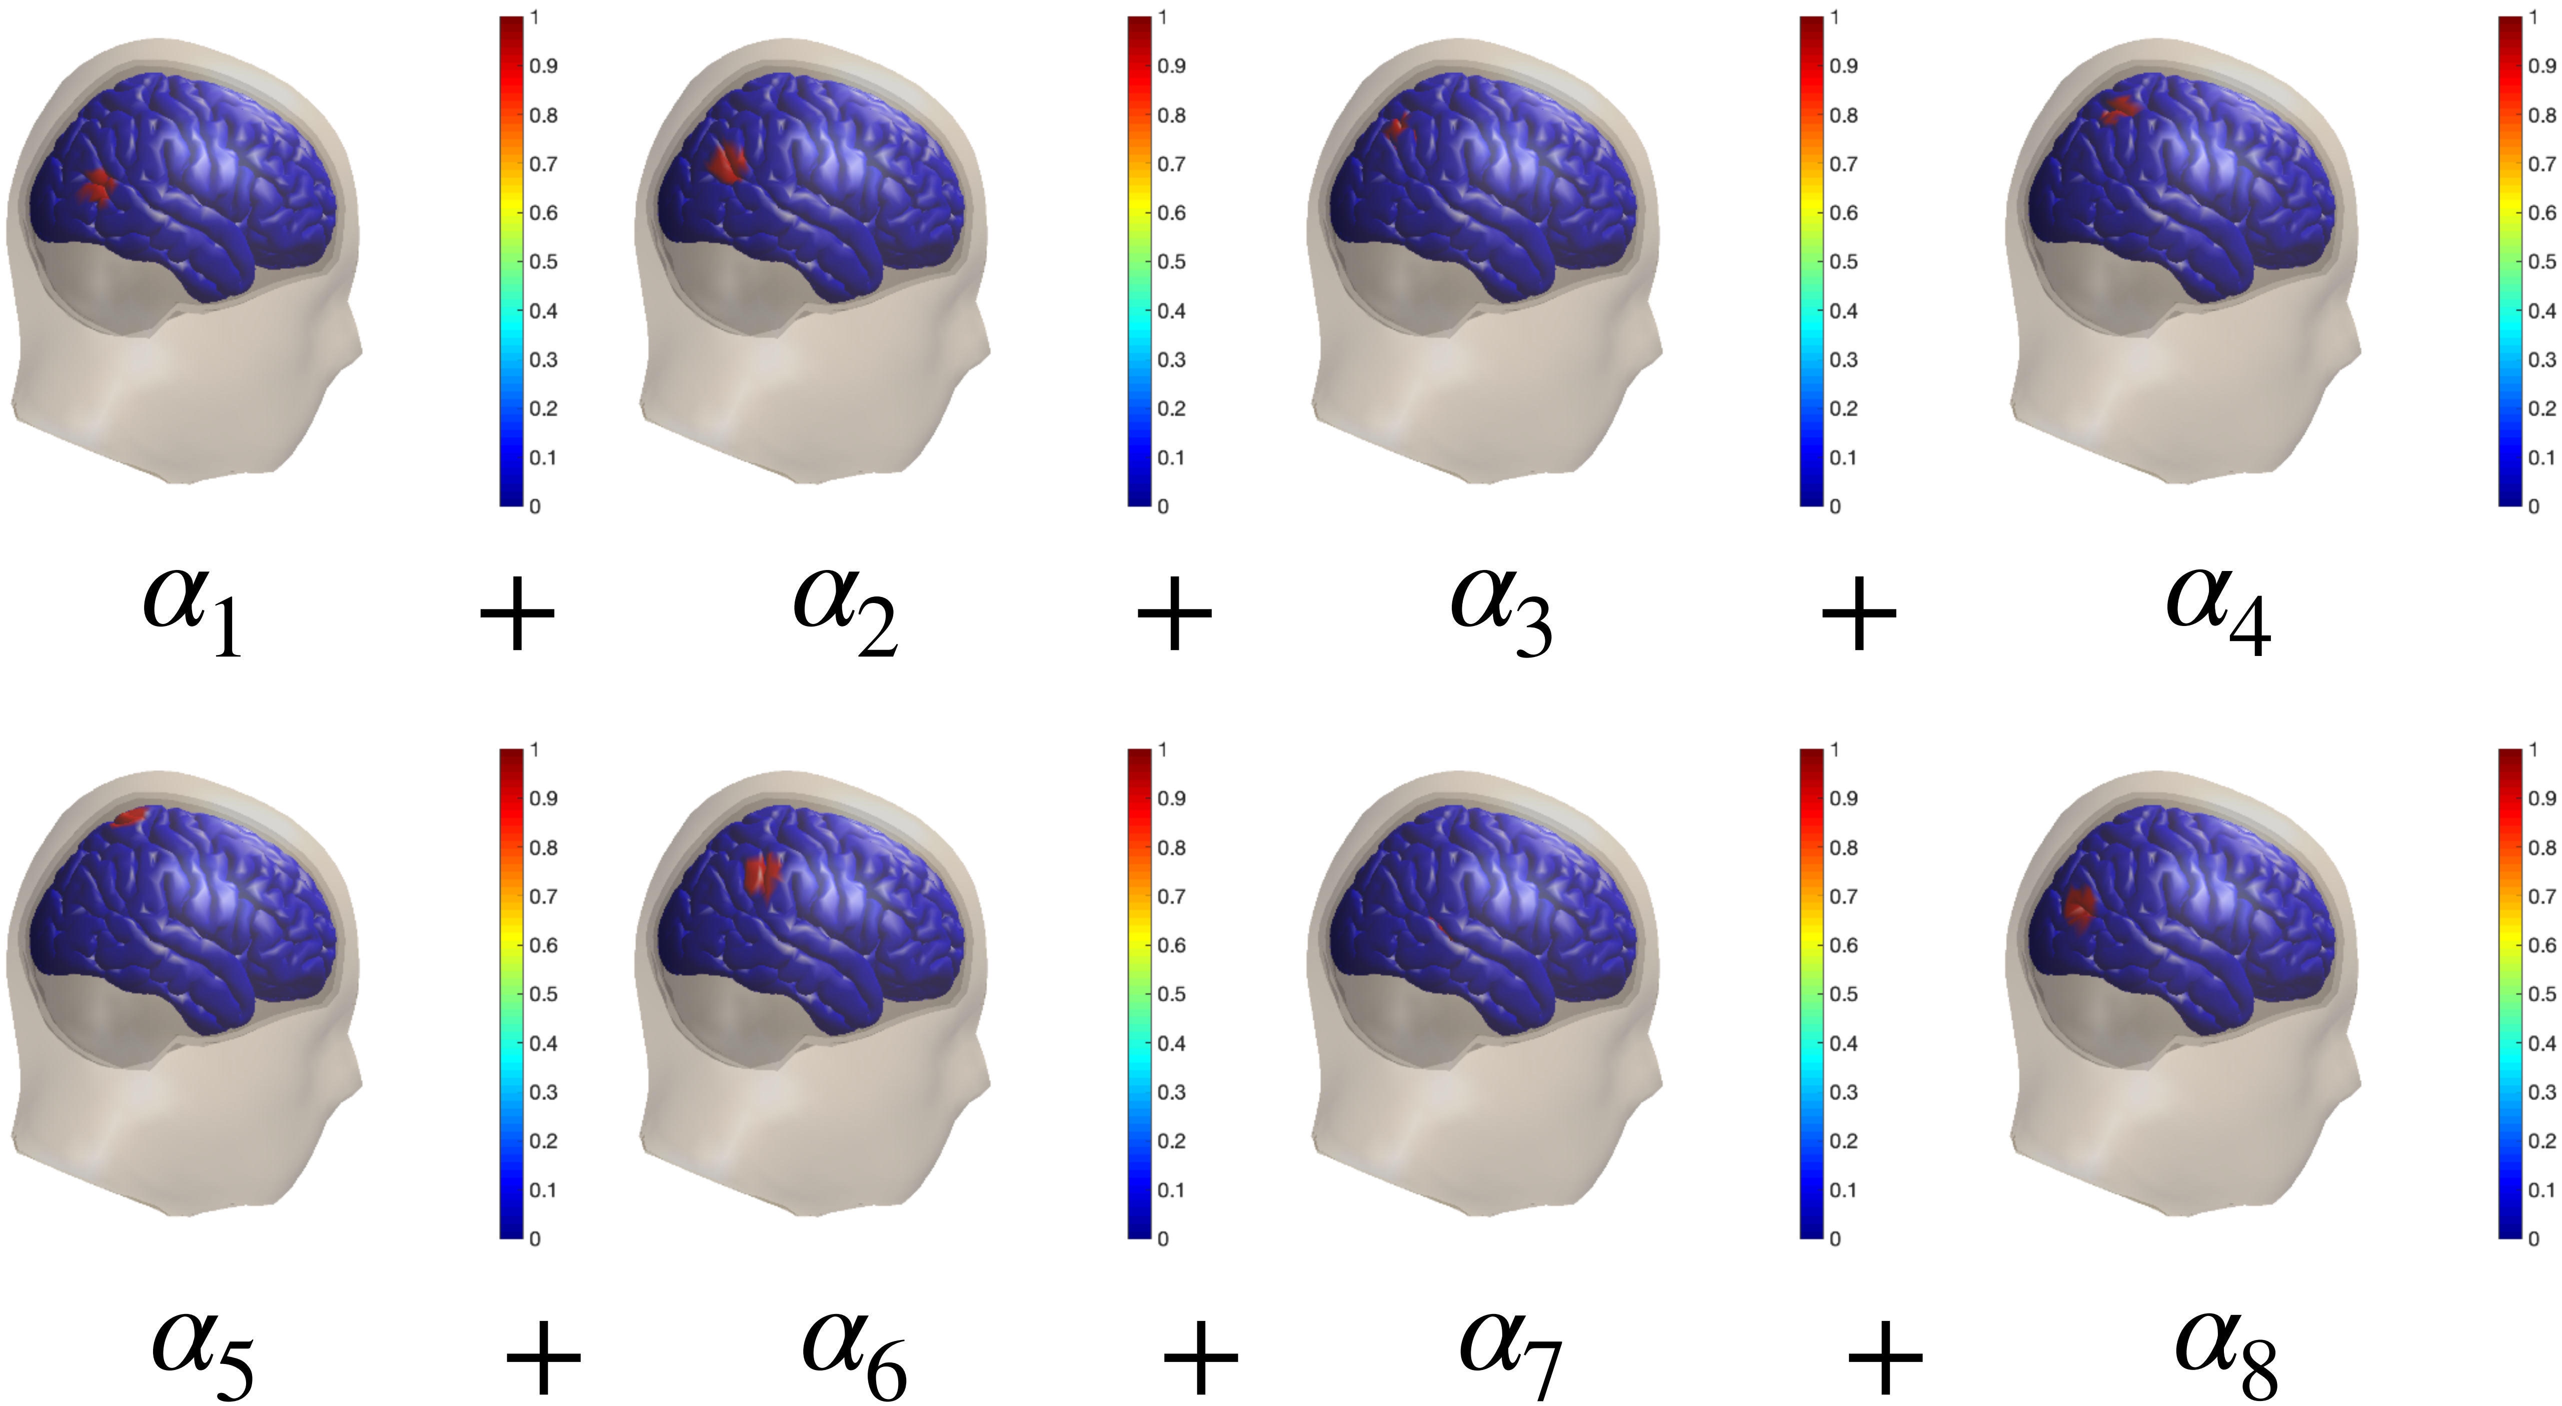

Source prior covariance matrix

$$C_X = \sum_{i}^{K} \alpha_i \beta_i$$

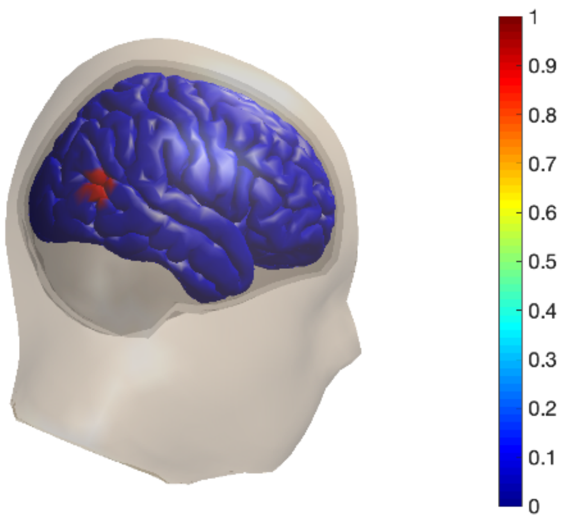

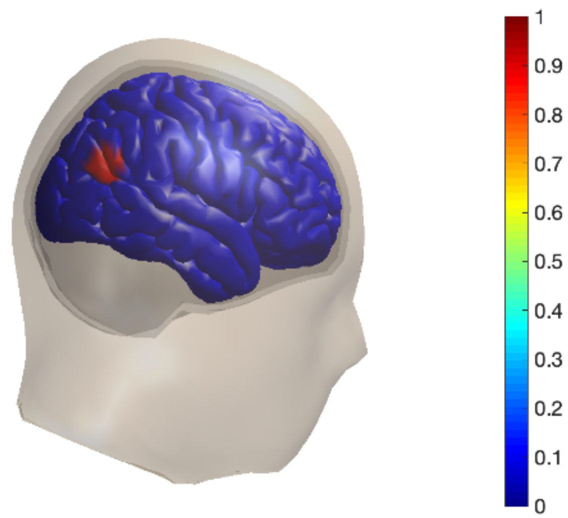

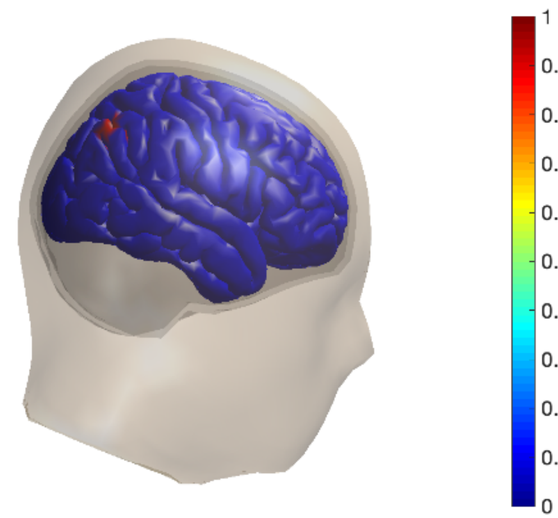

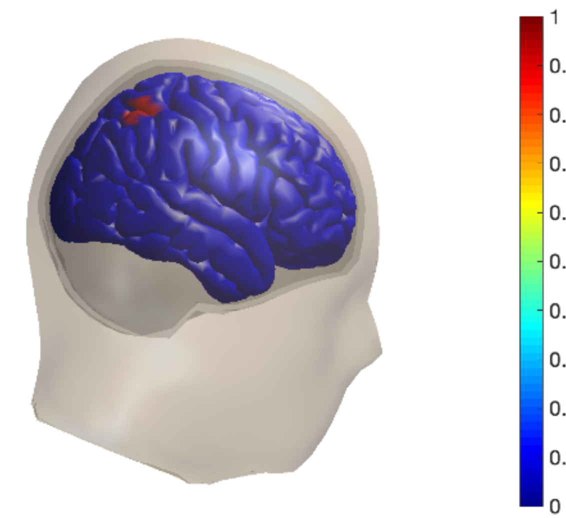

$$0.1 \quad + \quad \pi \quad + \quad \sqrt{2} \quad + \quad 0.665$$

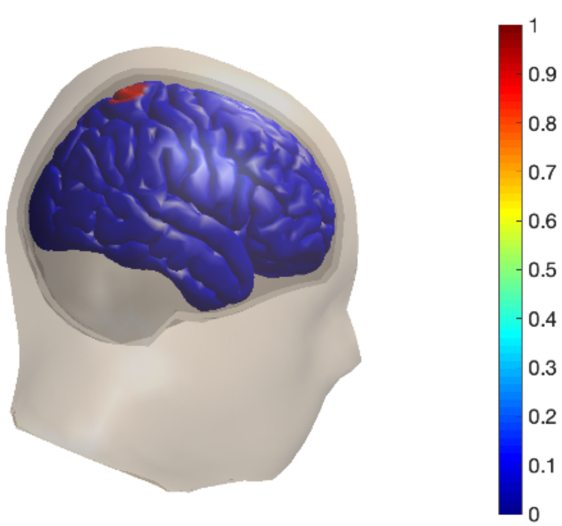

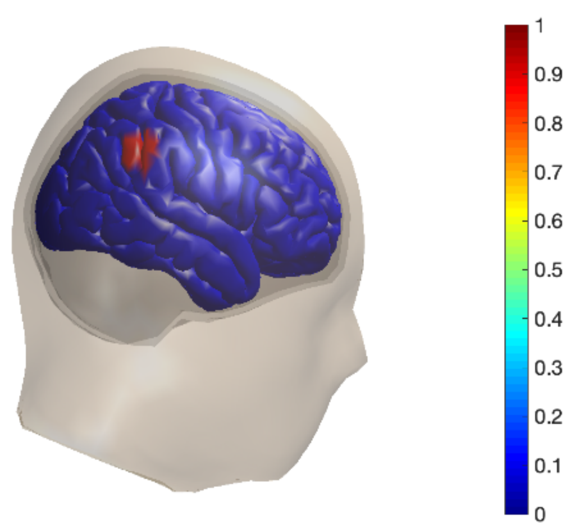

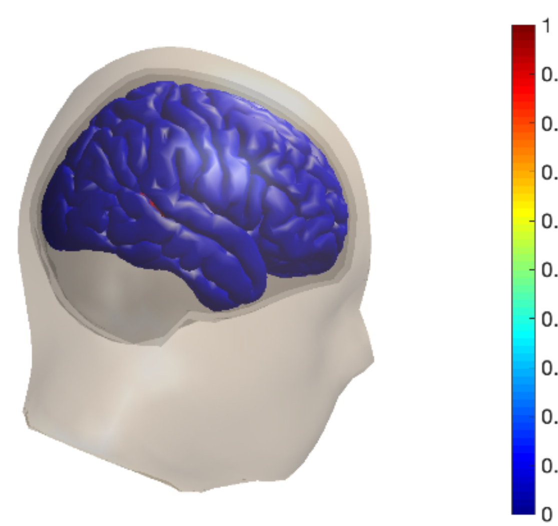

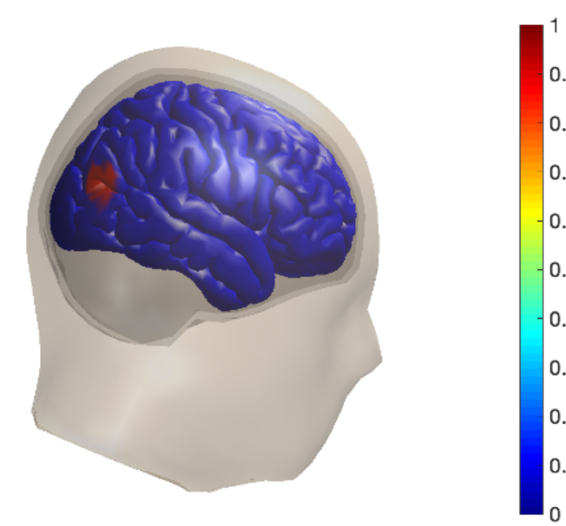

$$(e^{i\pi})^2 \quad + \quad 7.2 \quad + \quad 0.01 \quad + \quad 9$$

Source prior covariance matrix

$$C_X = \sum_{i}^{K} \alpha_i \beta_i$$

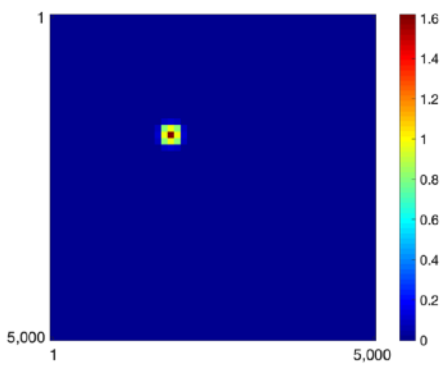

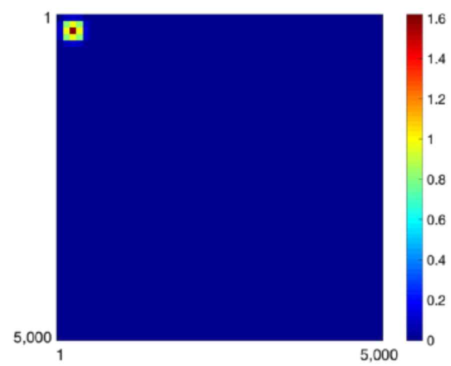

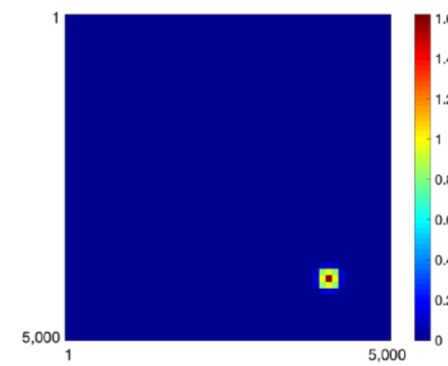

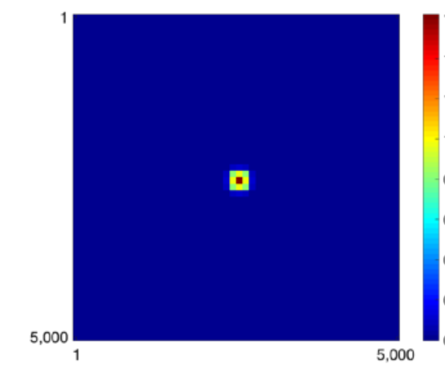

$0.1 \quad + \quad \pi \quad + \quad \sqrt{2} \quad + \quad 0.665$

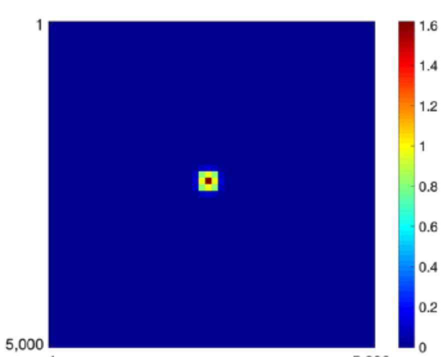

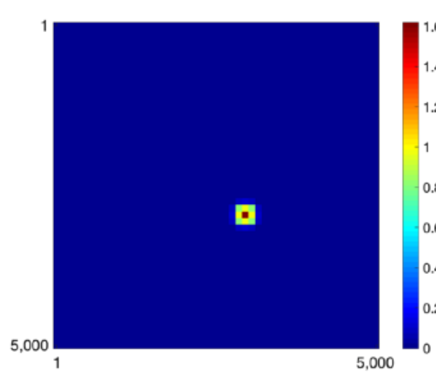

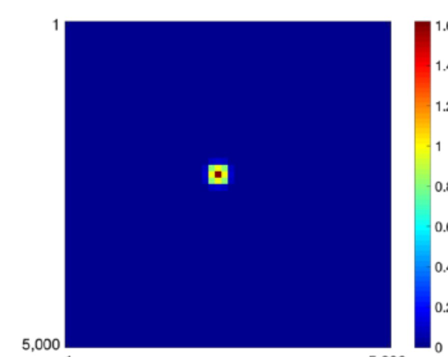

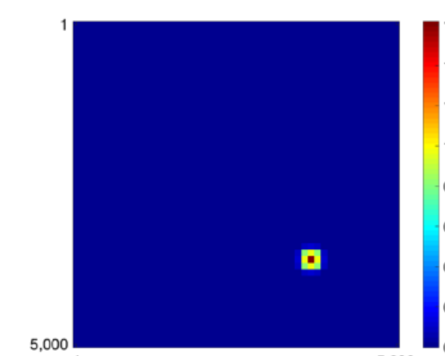

$(e^{i\pi})^2 \quad + \quad 7.2 \quad + \quad 0.01 \quad + \quad 9$

Source prior covariance matrix

$$C_X = \sum_{i}^{K} \alpha_i \beta_i$$

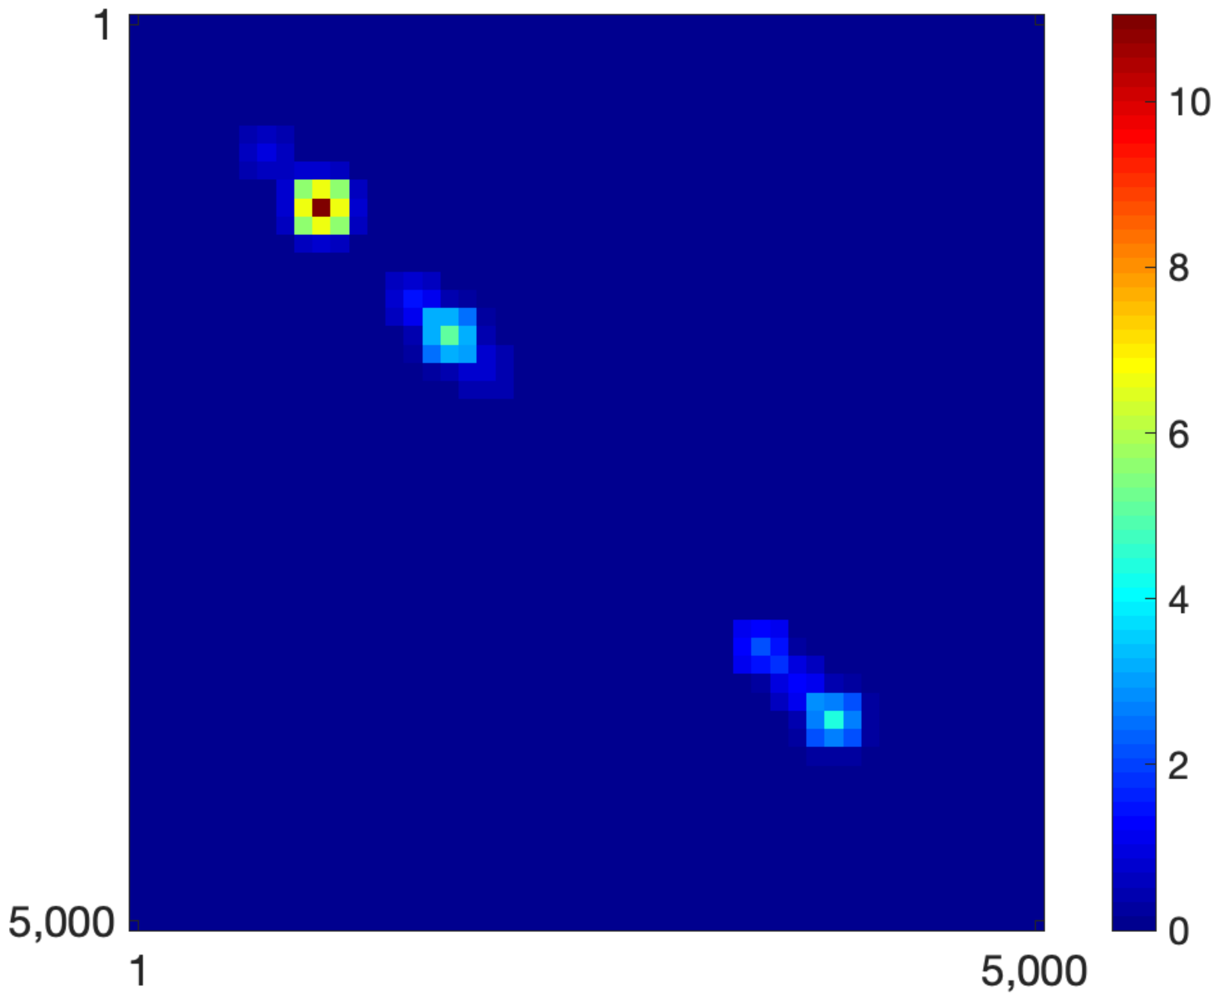

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T \right]^{-1} Y$$

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

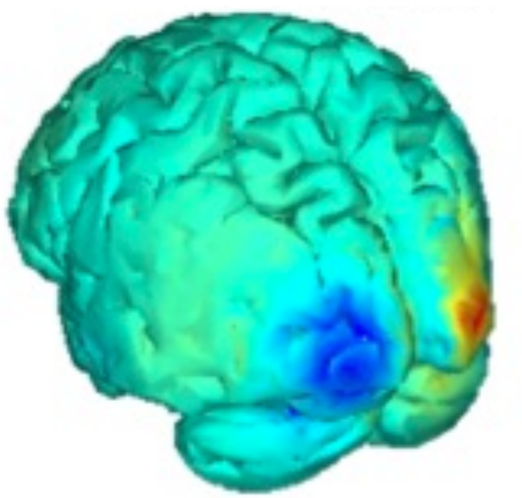

Current estimates = f (Data covariance, Forward model, Recorded data)

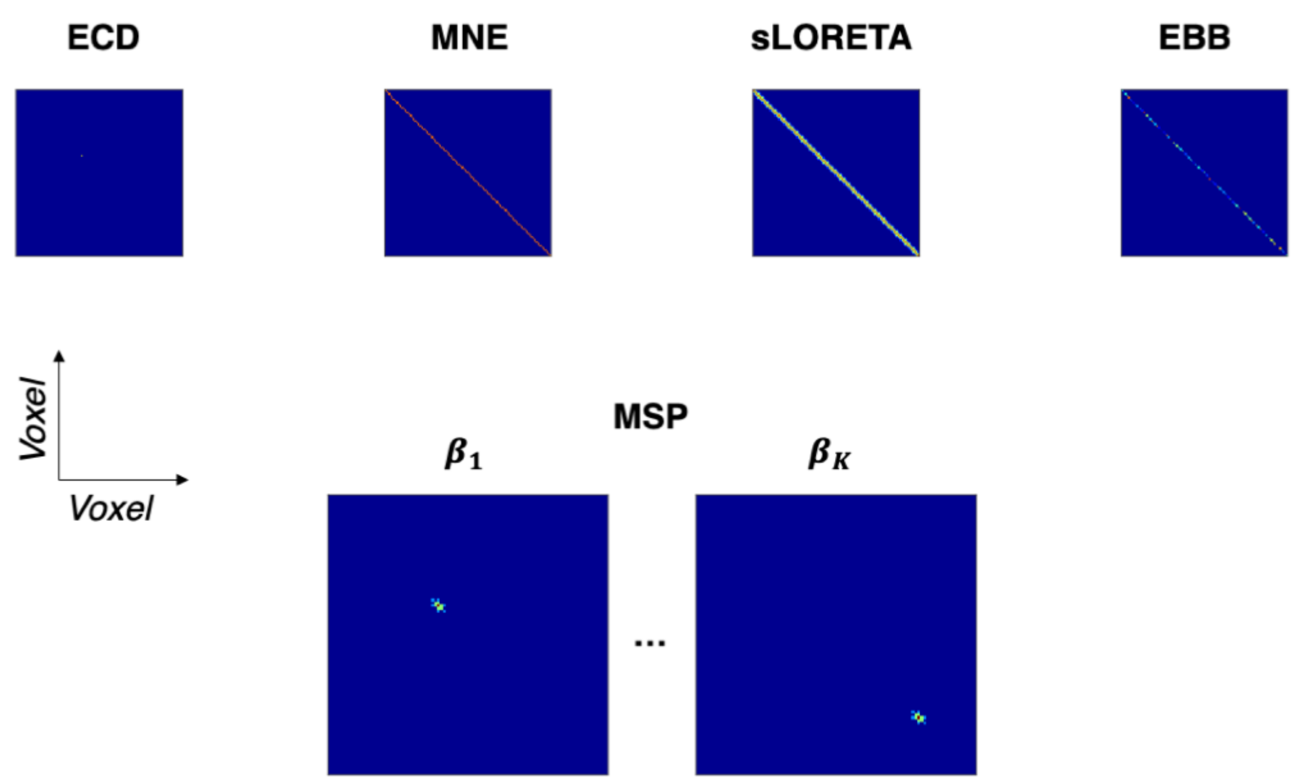

Source prior covariance matrix

$$C_X = \sum_{i}^{K} \alpha_i \beta_i$$

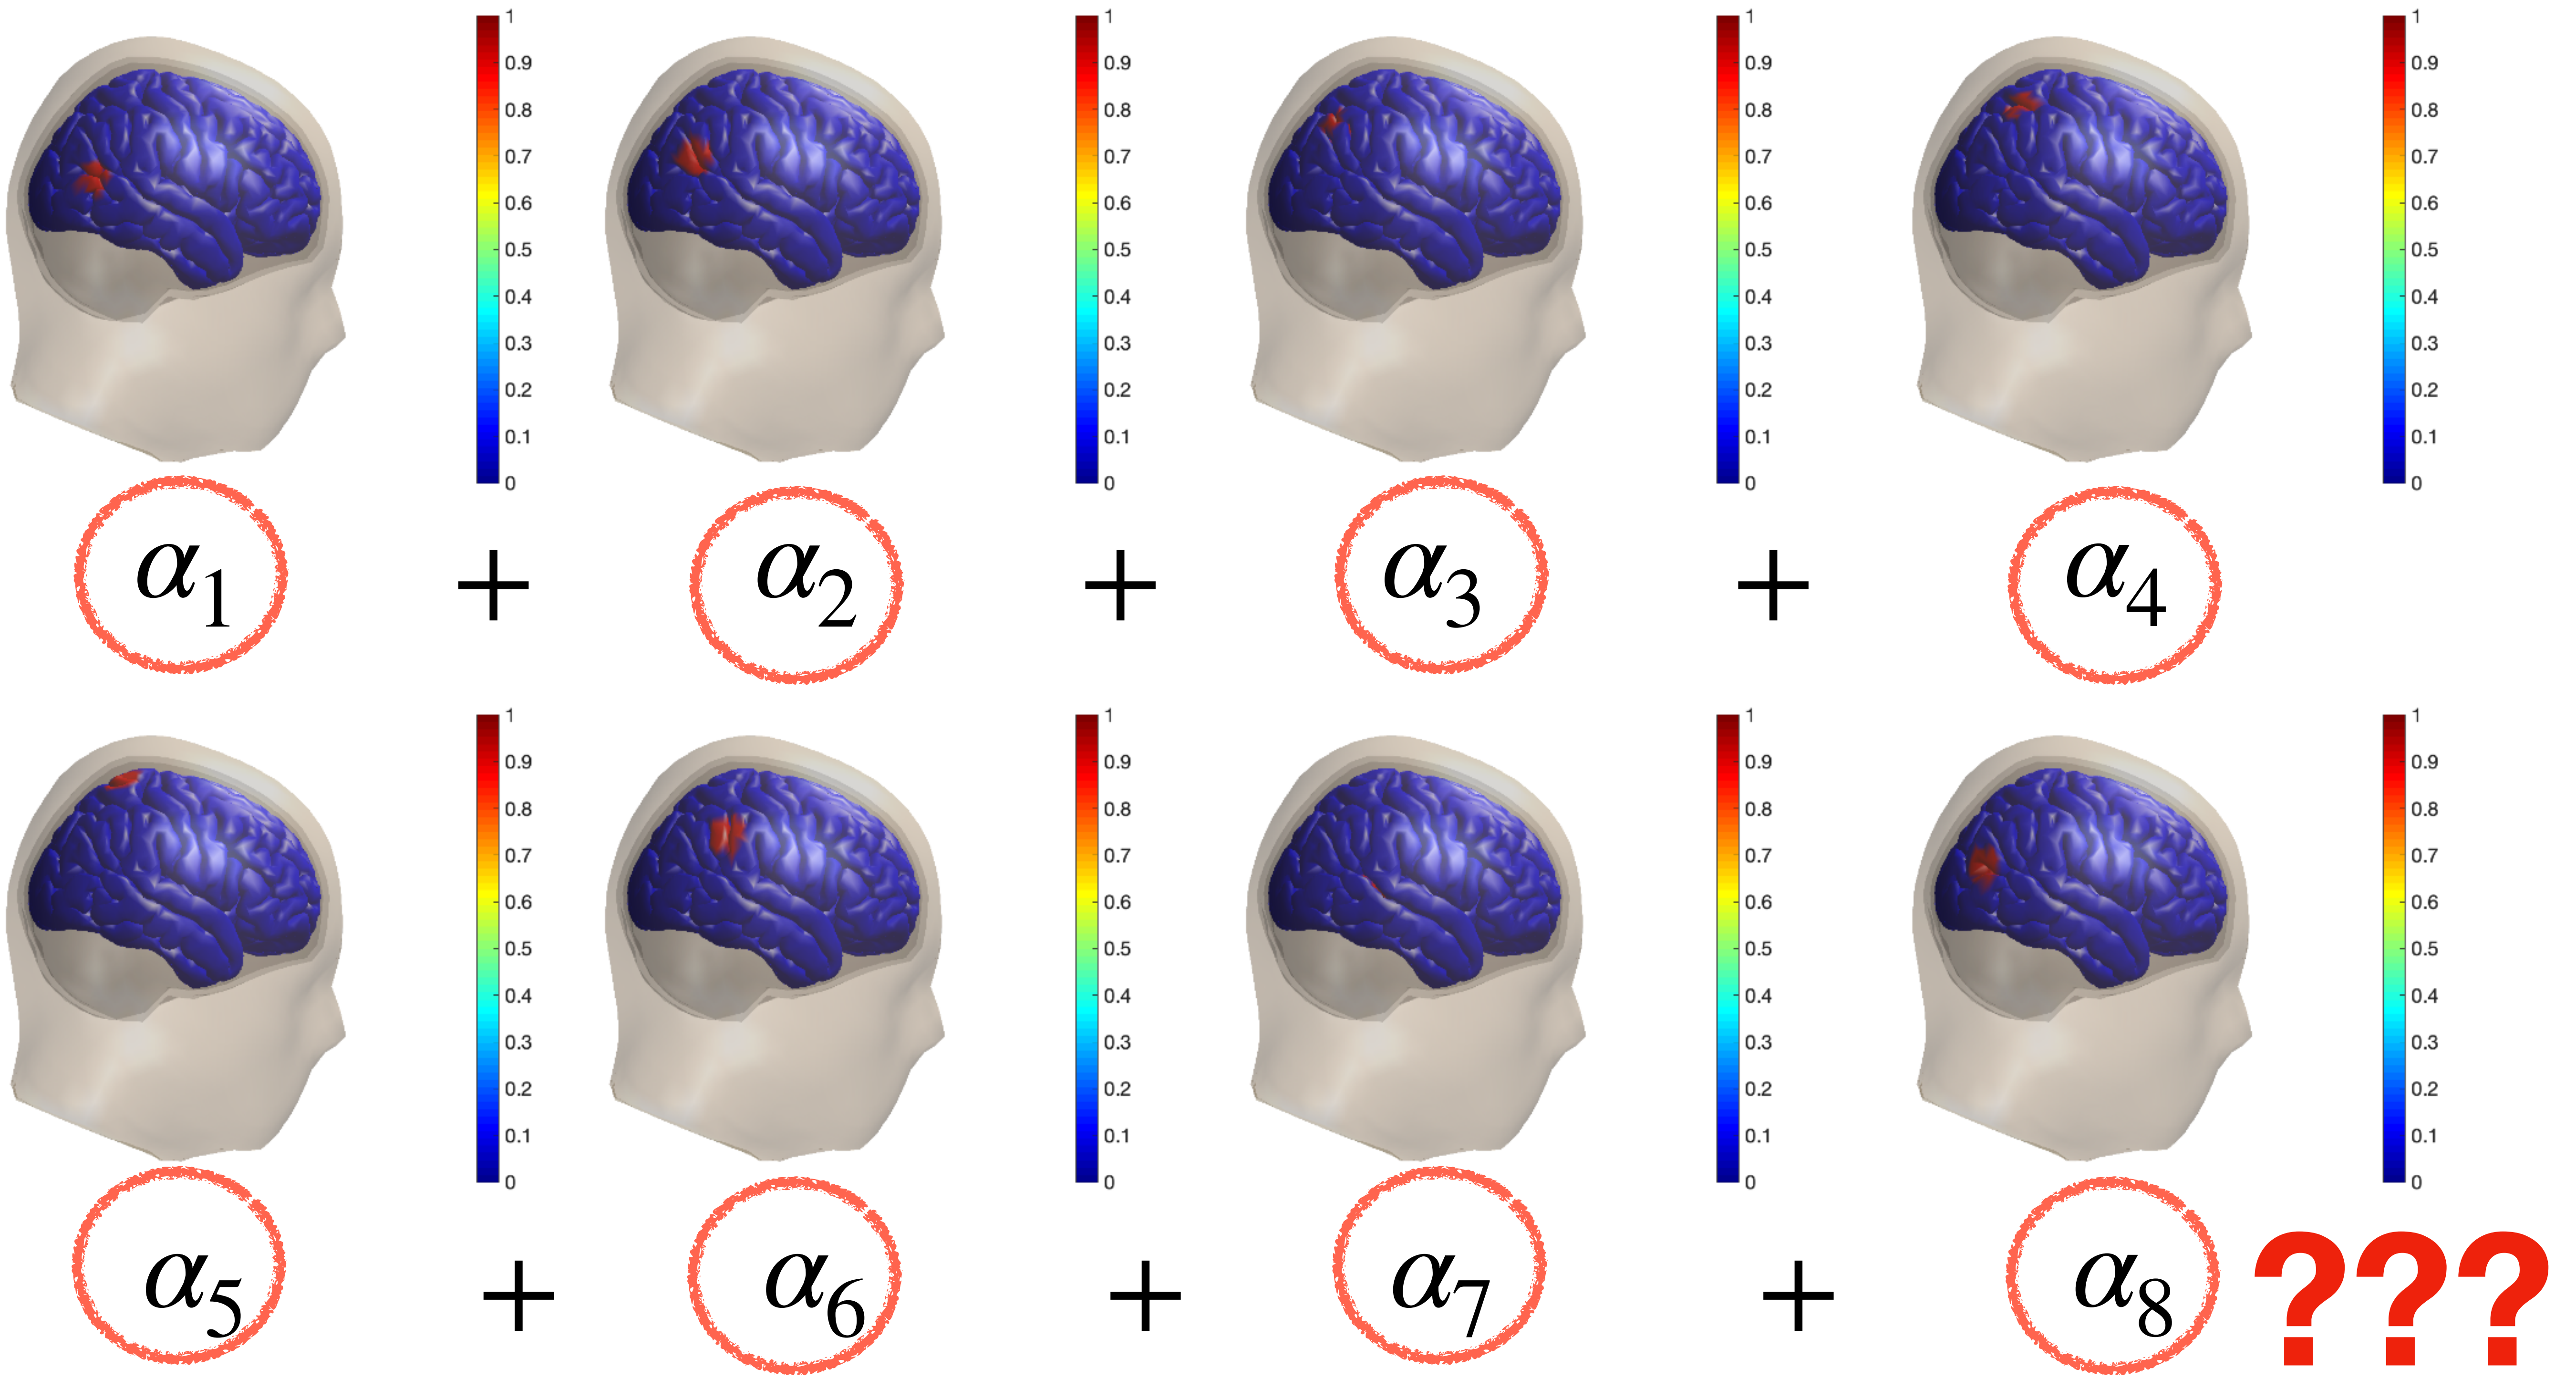

Each source prior has a representation at the sensor level

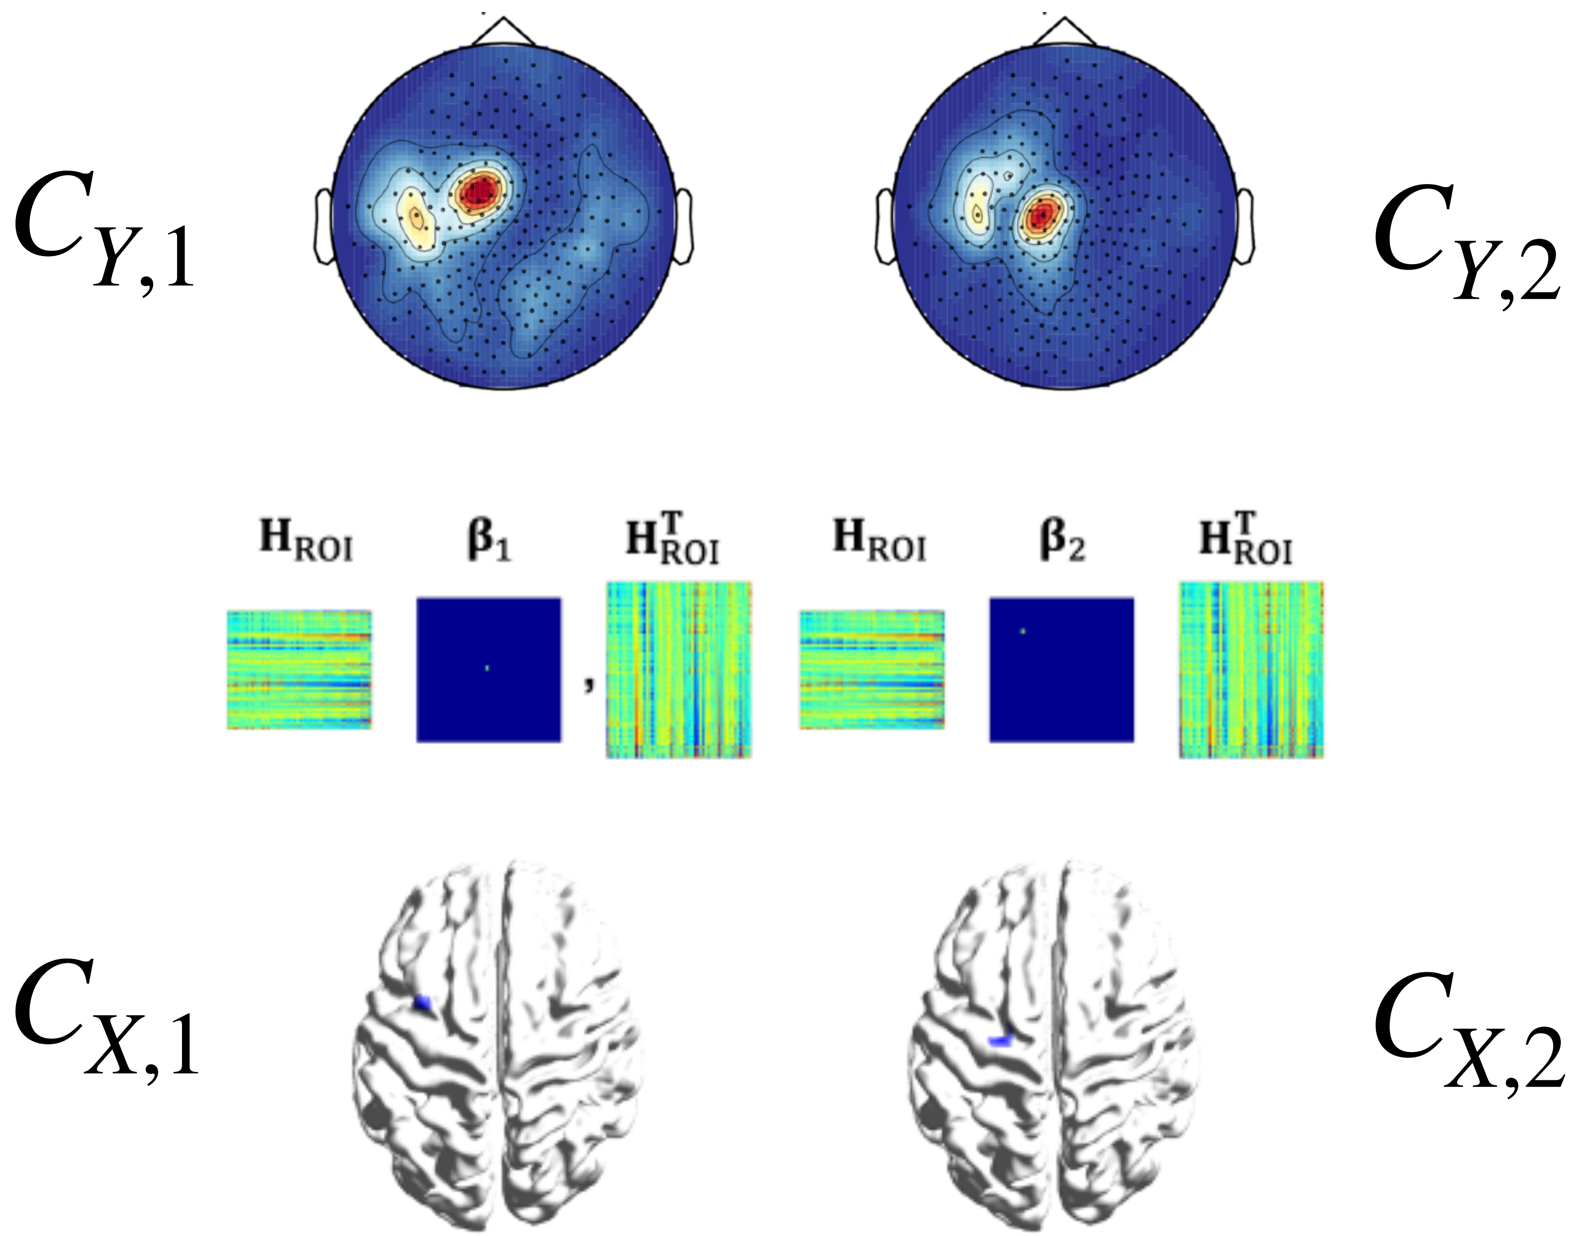

Problem: our posterior distribution (the thing that we want) is currently a function of the source space currents. This means our optimisation problem (i.e. learning the alphas) is in a very large space - the source space

$$p(X, \alpha \mid Y) \propto p(Y \mid X, \alpha) p(X \mid \alpha) p(\alpha)$$

Problem: our posterior distribution (the thing that we want) is currently a function of the source space currents. This means our optimisation problem (i.e. learning the alphas) is in a very large space - the source space

$$p(X, \alpha \,|\, Y) \propto p(Y \,|\, X, \alpha) p(X \,|\, \alpha) p(\alpha)$$

$$p(\alpha \,|\, Y) = \int p(X, \alpha \,|\, Y) dX$$

"Marginalisation"

Problem: our posterior distribution (the thing that we want) is currently a function of the source space currents. This means our optimisation problem (i.e. learning the alphas) is in a very large space - the source space

$$p(X, \alpha \mid Y) \propto p(Y \mid X, \alpha)p(X \mid \alpha)p(\alpha)$$

$$p(\alpha \mid Y) = \int p(X, \alpha \mid Y)dX$$

$$= \int p(Y \mid X, \alpha)p(X \mid \alpha)p(\alpha)dX$$

"Marginalisation"

$$= \int p(Y \mid X, \alpha)p(X \mid \alpha)dXp(\alpha)$$

...

Problem: our posterior distribution (the thing that we want) is currently a function of the source space currents. This means our optimisation problem (i.e. learning the alphas) is in a very large space - the source space

$$p(X, \alpha \,|\, Y) \propto p(Y \,|\, X, \alpha) p(X \,|\, \alpha) p(\alpha)$$

$$p(\alpha \,|\, Y) = \int p(X, \alpha \,|\, Y) dX$$

$$= \int p(Y \,|\, X, \alpha) p(X \,|\, \alpha) p(\alpha) dX$$

$$= \int p(Y \,|\, X, \alpha) p(X \,|\, \alpha) dX p(\alpha)$$

$$= \underline{p(Y \,|\, \alpha) p(\alpha)} \propto exp\left[-0.5tr\left(Y^T C_Y^{-1} Y\right)\right] p(\alpha)$$

$$C_Y = H C_X H^T$$

Before:

$$C_X = \sum_i^K \alpha_i \beta_i$$

Now:

Can just model the sensor level data covariance (i.e. the covariance of the data which we measure, $\mathbb{C}_Y$):

$$\mathbb{C}_Y \approx C_Y = HC_X H^T = H\left(\sum_i^K \alpha_i \beta_i\right) H^T = \sum_i^K \alpha_i \tilde{\beta}_i$$

These are the same alphas!

Each source prior has a representation at the sensor level

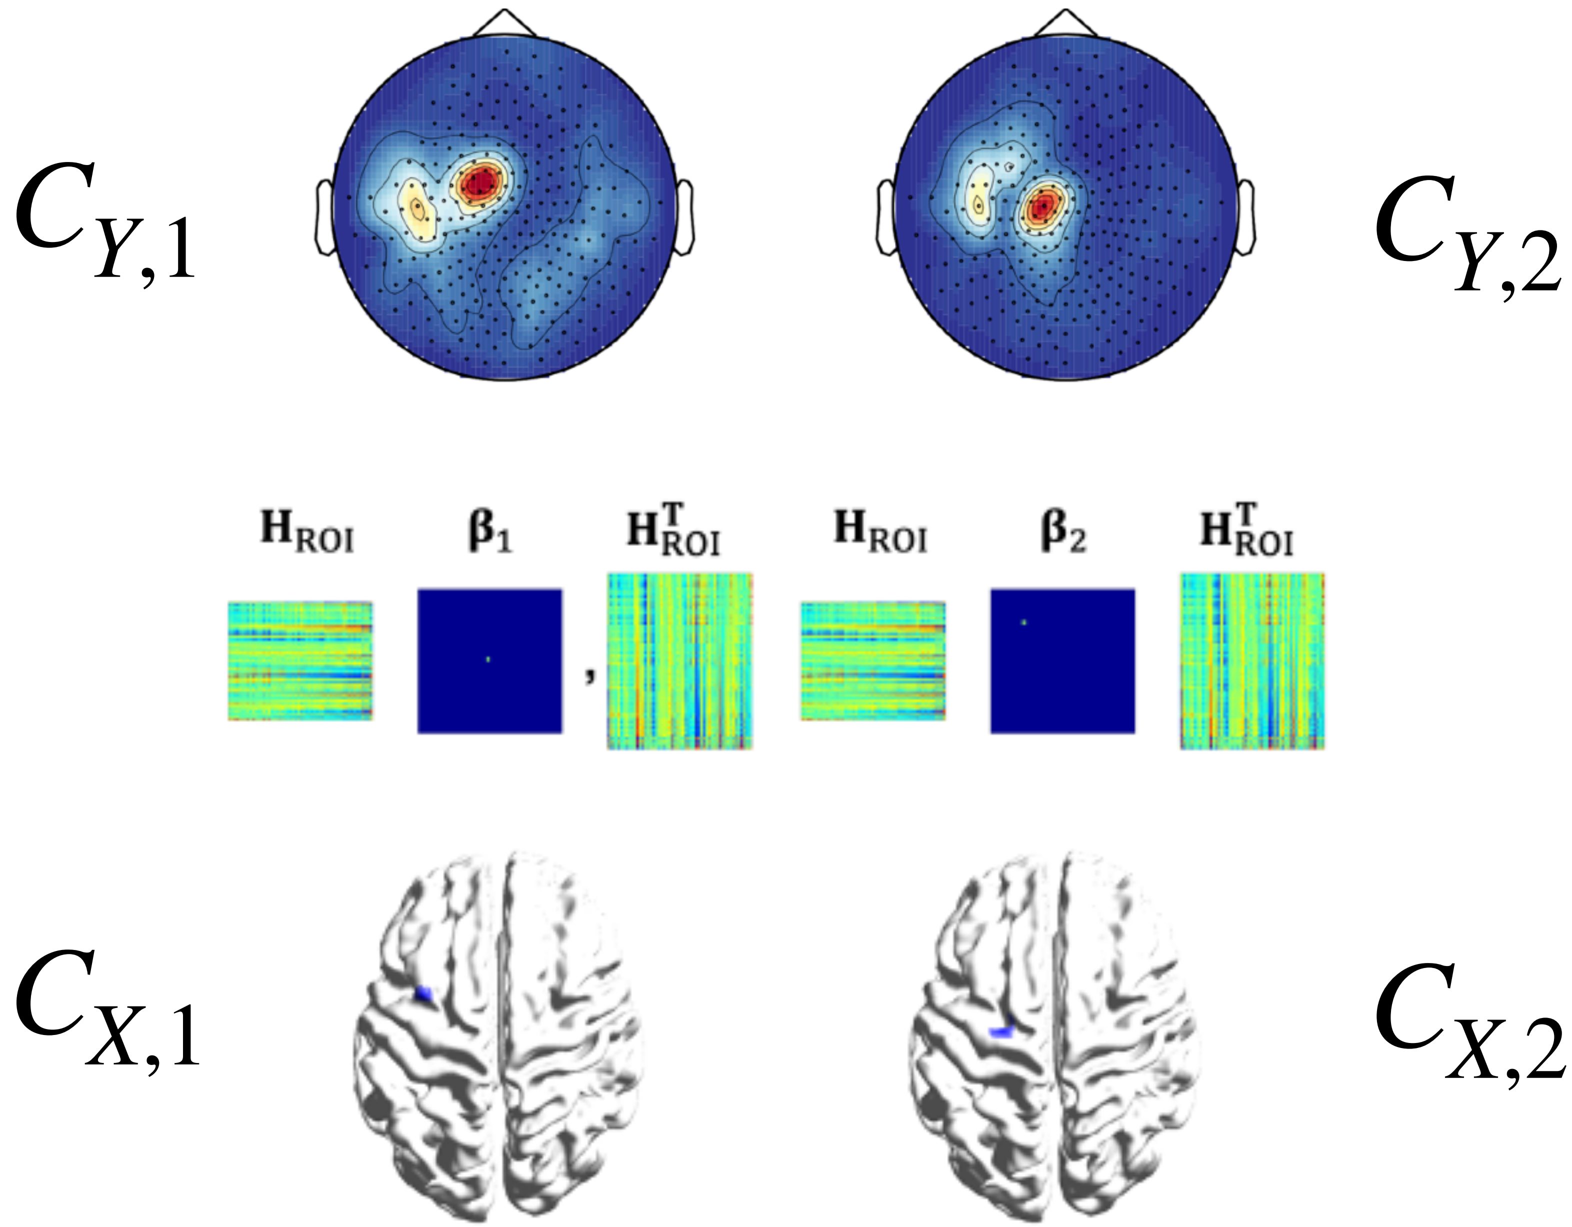

$$\mathbb{C}_Y \approx C_Y = HC_XH^T = H\left(\sum_i^K \alpha_i\beta_i\right)H^T = \sum_i^K \alpha_i\tilde{\beta}_i$$

Aside: an introduction to Variational Free Energy

Variational Free Energy

We would like to calculate the true (marginalised) posterior distribution, $p(\alpha \mid Y)$. This can be hard to calculate in practice.

Instead, we approximate the true posterior with some simpler parameterised distribution(s), $q(\alpha)$.

We then minimise the KL-divergence between the approximate posterior, $q(\alpha)$, and the true posterior distribution, $p(\alpha \mid Y)$.

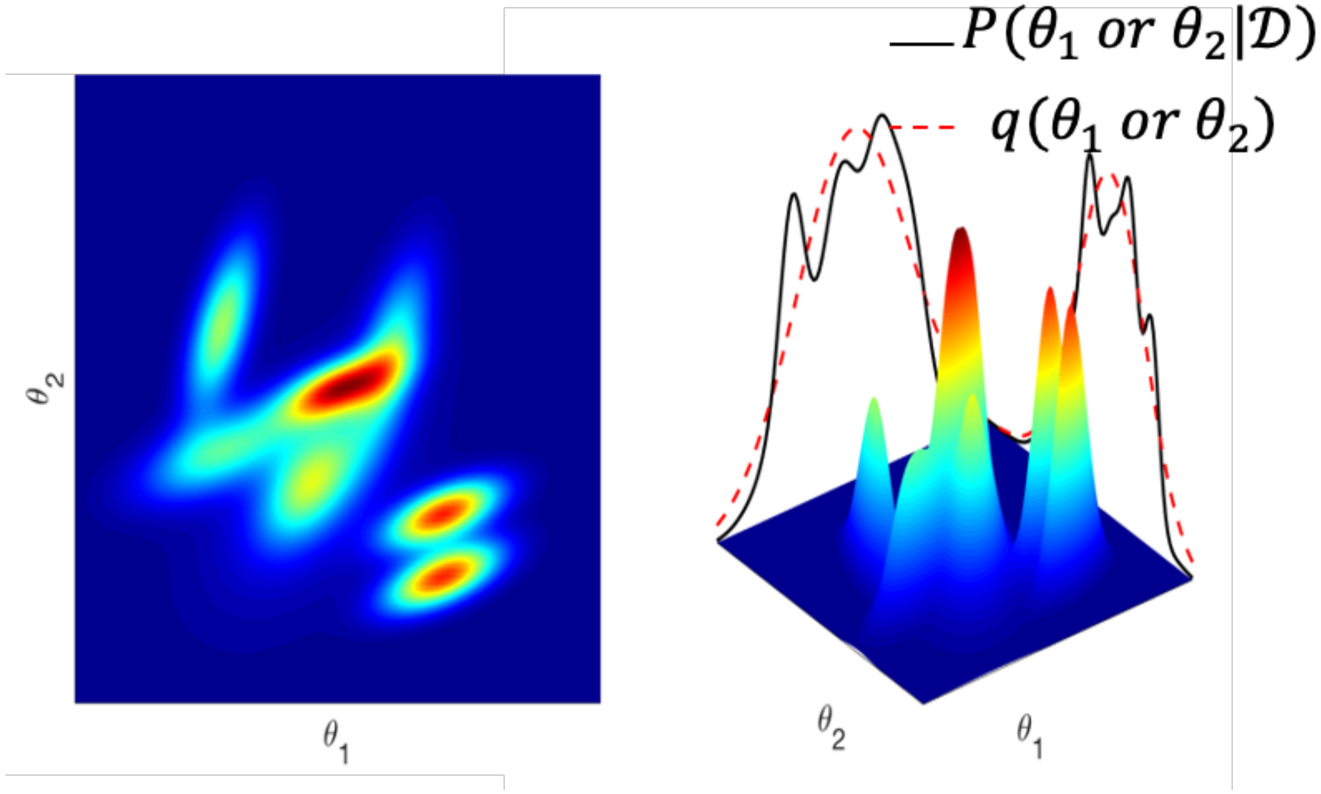

Variational Free Energy

We can show that the KL divergence between the true and approximate posterior can be written as

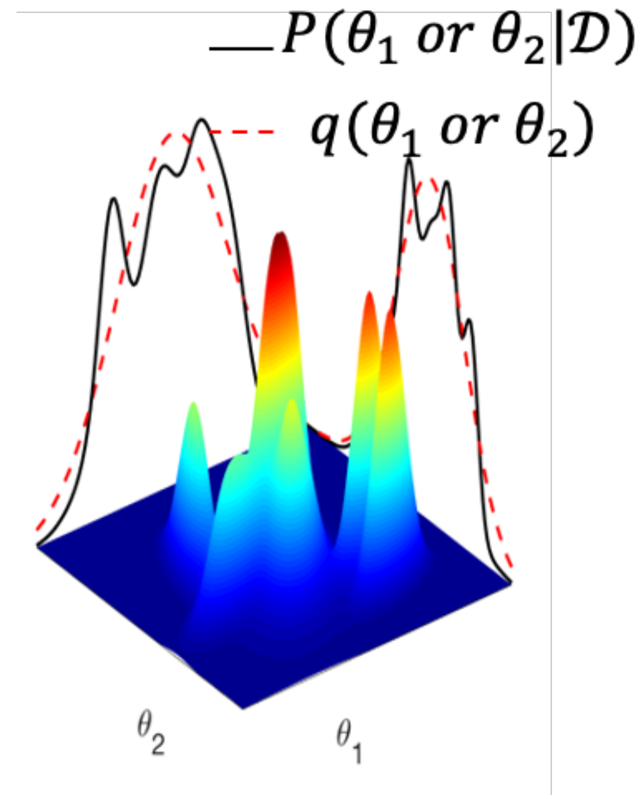

$$log \left[p(Y)\right] = F + KL \left[q(\alpha) \, || \, p(\alpha \, | \, Y)\right]$$

The KL divergence is strictly greater than or equal to zero. We would like to make this equal zero.

The log model evidence is a constant.

Hence maximising the variational Free Energy, F, is equivalent to minimising the distance between the true and approximate posterior distribution.

Variational Free Energy

$$F = log\left[p(Y)\right] - KL\left[q(\alpha)\,||\,p(\alpha\,|\,Y)\right]$$

Variational Free Energy

$$F = log\left[p(Y)\right] - KL\left[q(\alpha)\,||\,p(\alpha\,|\,Y)\right]$$

$$F = Log\ \ likelihood - KL\left[q(\alpha)\,||\,p(\alpha)\right]$$

$$F = Accuracy - Complexity$$

I want to explain my data well

But not at any expense

Choosing between solutions

- No way of knowing the ground truth on real data
- Can use variance explained as a means of choosing between source priors
  Dangerous!

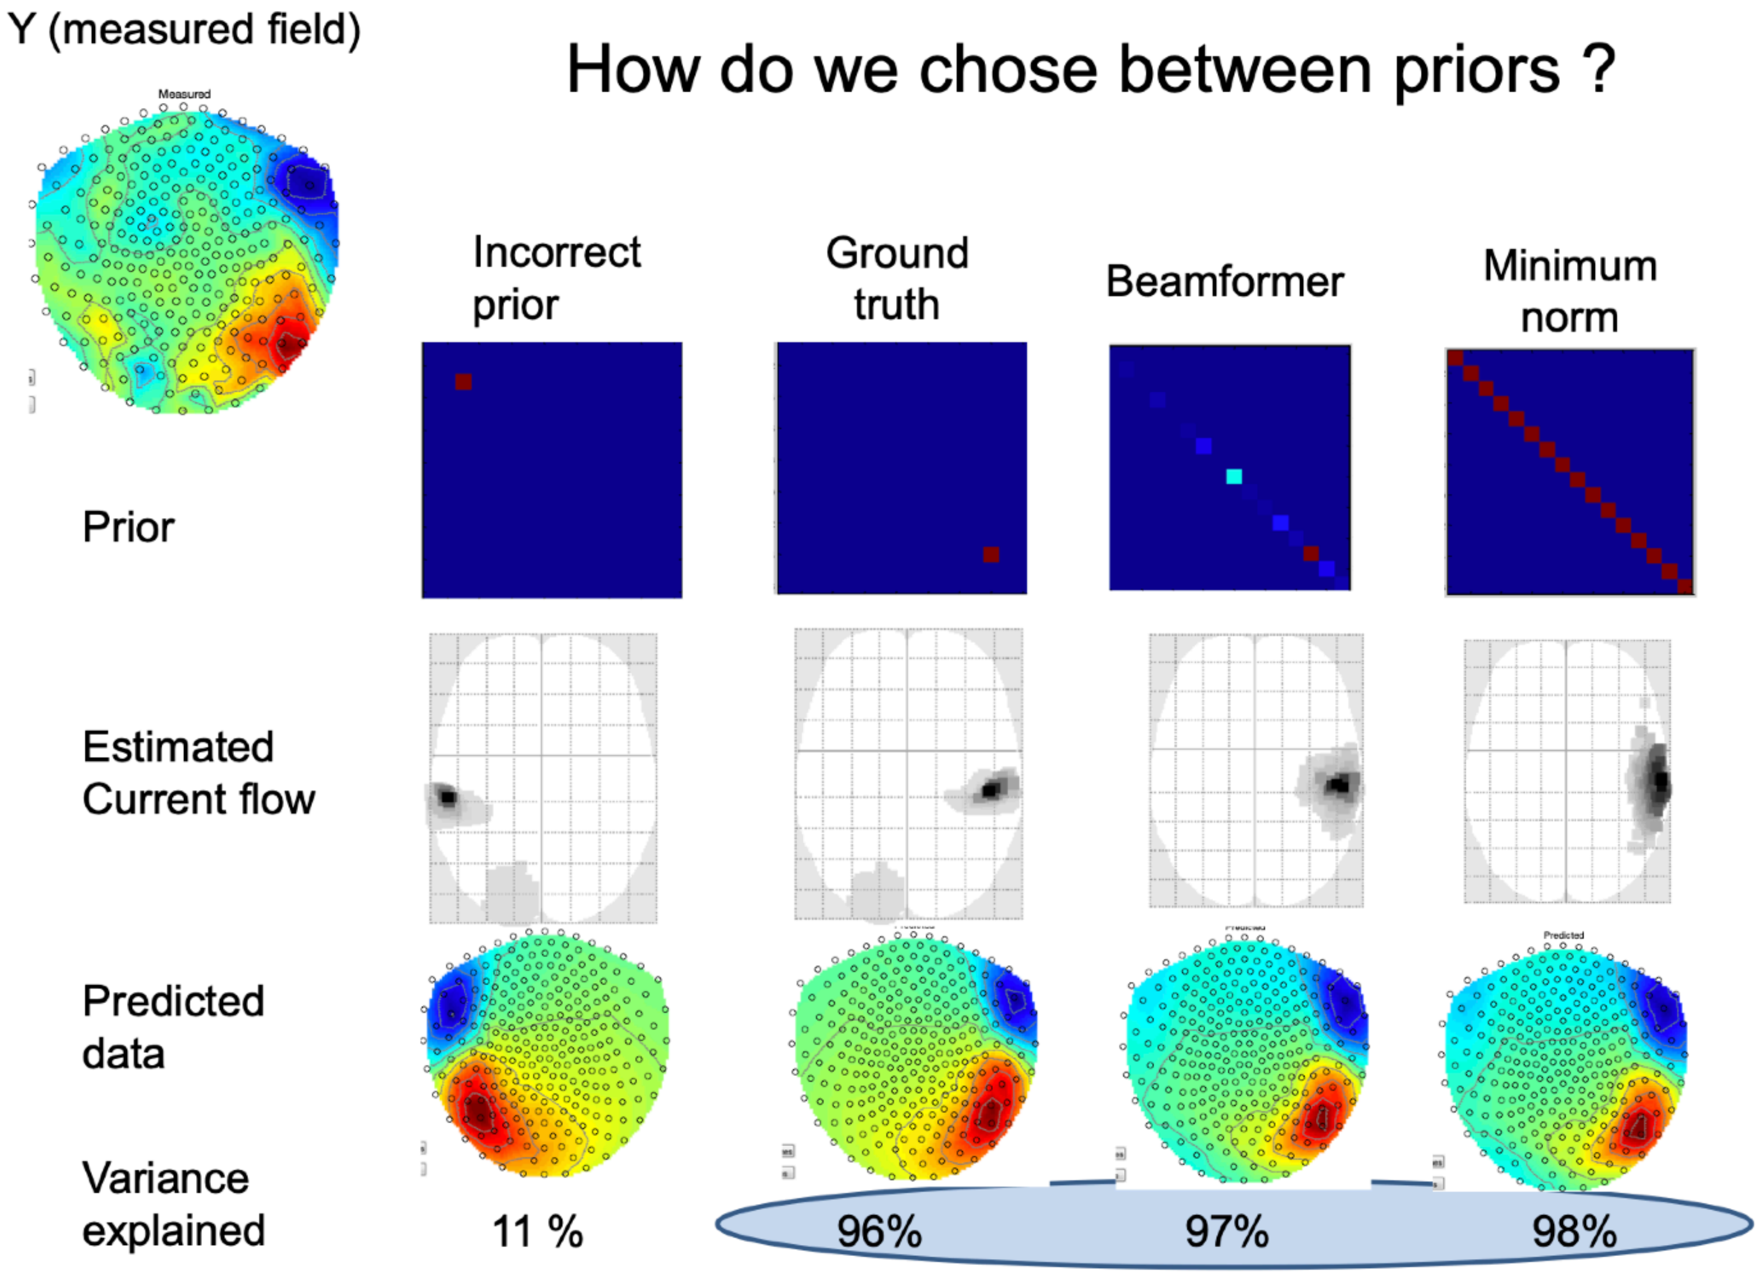

Figure credit: José David López

Choosing between solutions

- If I maximise the Free Energy, the KL divergence will have to decrease

$$F = log\left[p(Y)\right] - KL\left[q(\alpha)\,||\,p(\alpha\,|\,Y)\right]$$

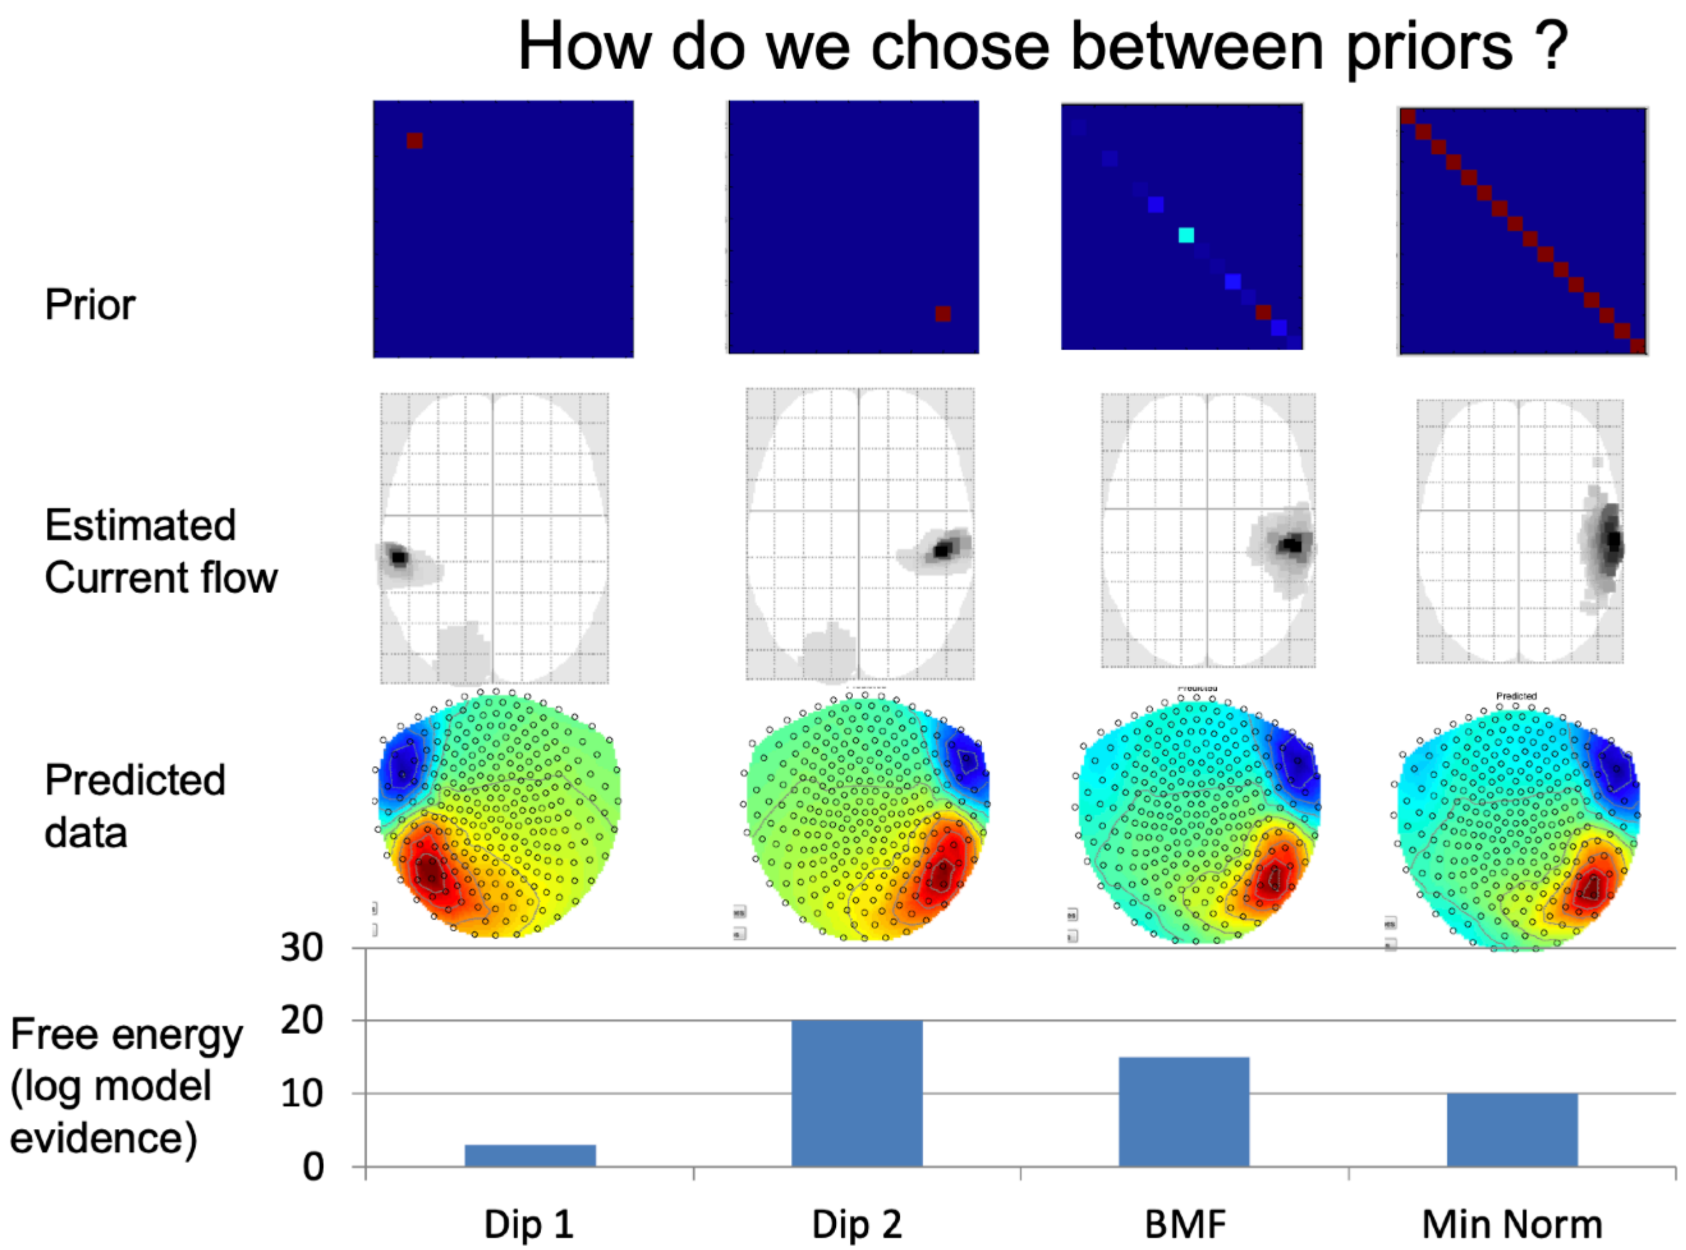

Figure credit: José David López

Multiple Sparse Priors

- Find the optimal set of α's which maximises the Free Energy

$$F = Accuracy - Complexity$$

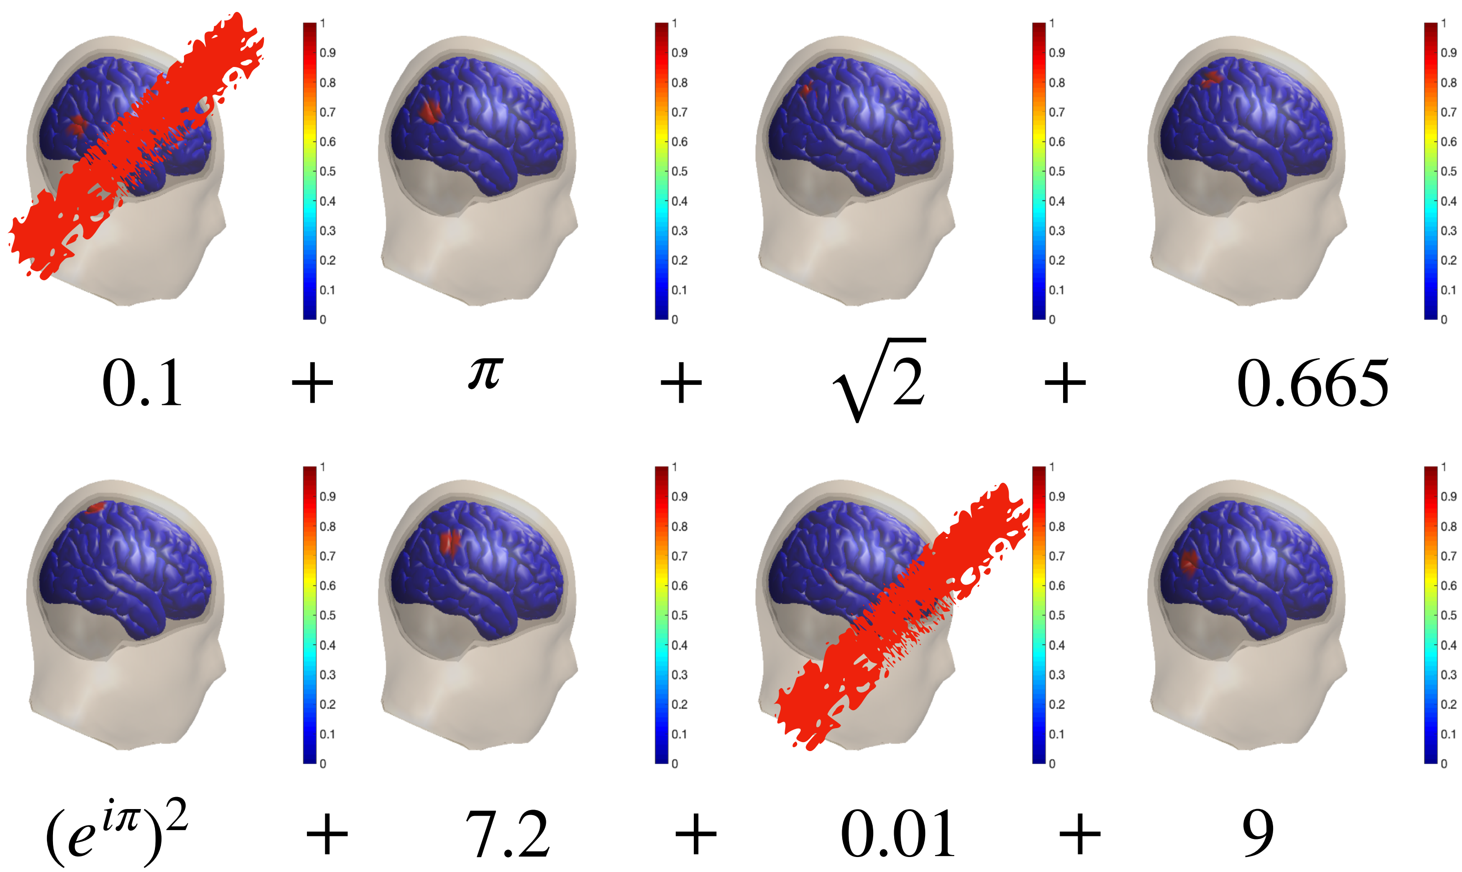

$$\mathbb{C}_Y \approx C_Y = HC_XH^T = H\left(\sum_i^K \alpha_i \beta_i\right)H^T = \sum_i^K \alpha_i \tilde{\beta}_i$$

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{X} = \alpha_1 C_X H^T \left[\alpha_2 C_n + \alpha_1 H C_X H^T\right]^{-1} Y$$

Point of note: we are always data driven in SPM, even when using an IID prior

i.e. we learn α_1 and α_2

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T\right]^{-1} Y$$

$$\hat{X} = \alpha_1 C_X H^T \left[\alpha_2 C_n + \alpha_1 H C_X H^T\right]^{-1} Y$$

Note!

These algorithms (in SPM) are designed to work on averaged data. Cannot apply to resting state etc.

Recap

Hard!

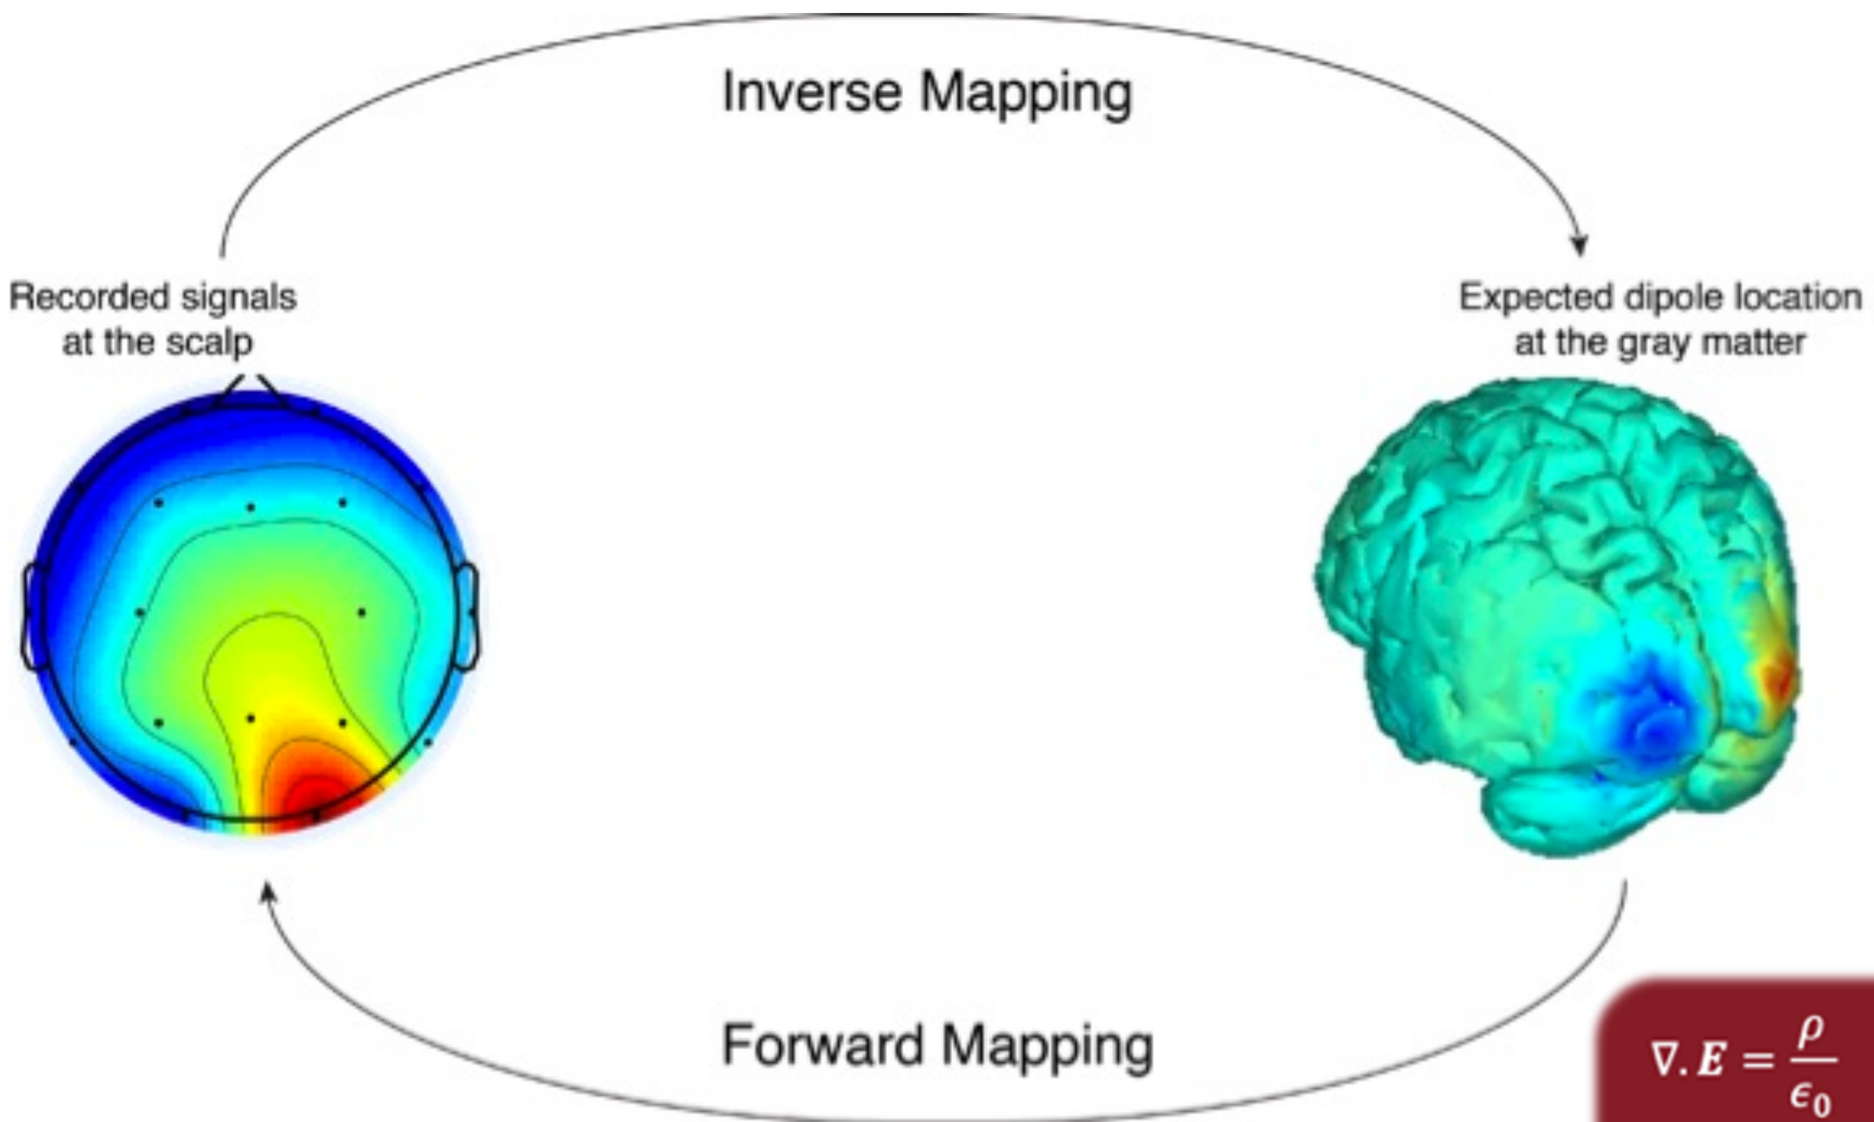

Easy!

$$\nabla . \boldsymbol{E} = \frac{\rho}{\epsilon_0}$$

$$\nabla . \boldsymbol{B} = 0$$

$$\nabla \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t}$$

$$\nabla \times \boldsymbol{B} = \mu_0 \boldsymbol{J} + \mu_0 \epsilon_0 \frac{\partial \boldsymbol{E}}{\partial t}$$

Figure credit: neuroelectrics.com

Recap

$$\hat{X} = C_X H^T \left[C_n + H C_X H^T \right]^{-1} Y$$

$$\hat{\mathbf{x}}(t) = \mathbf{W}\mathbf{y}(t)$$

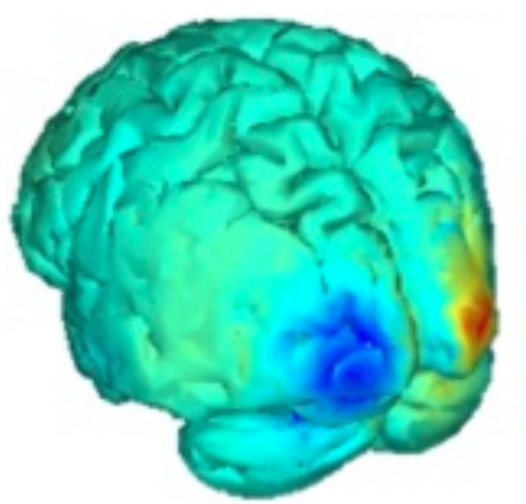

Current estimates = f (Data covariance, Forward model, Recorded data)

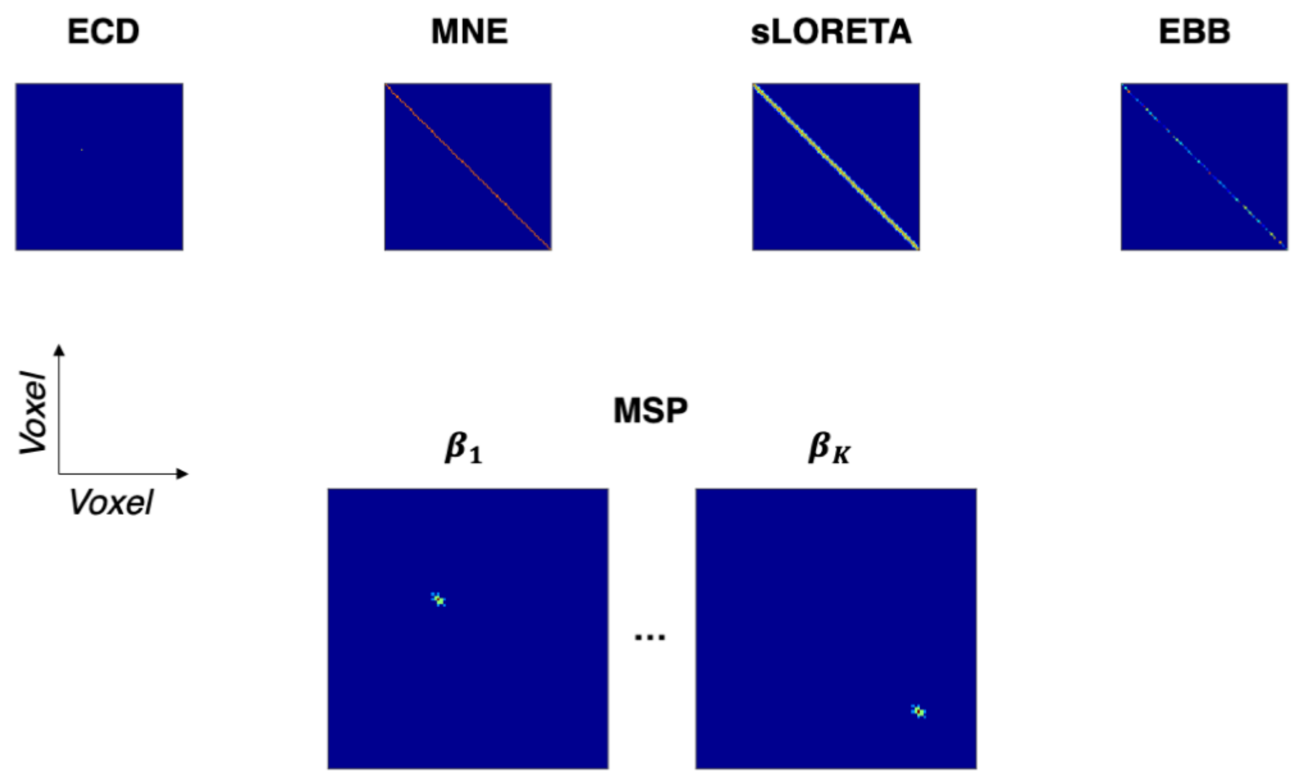

Recap

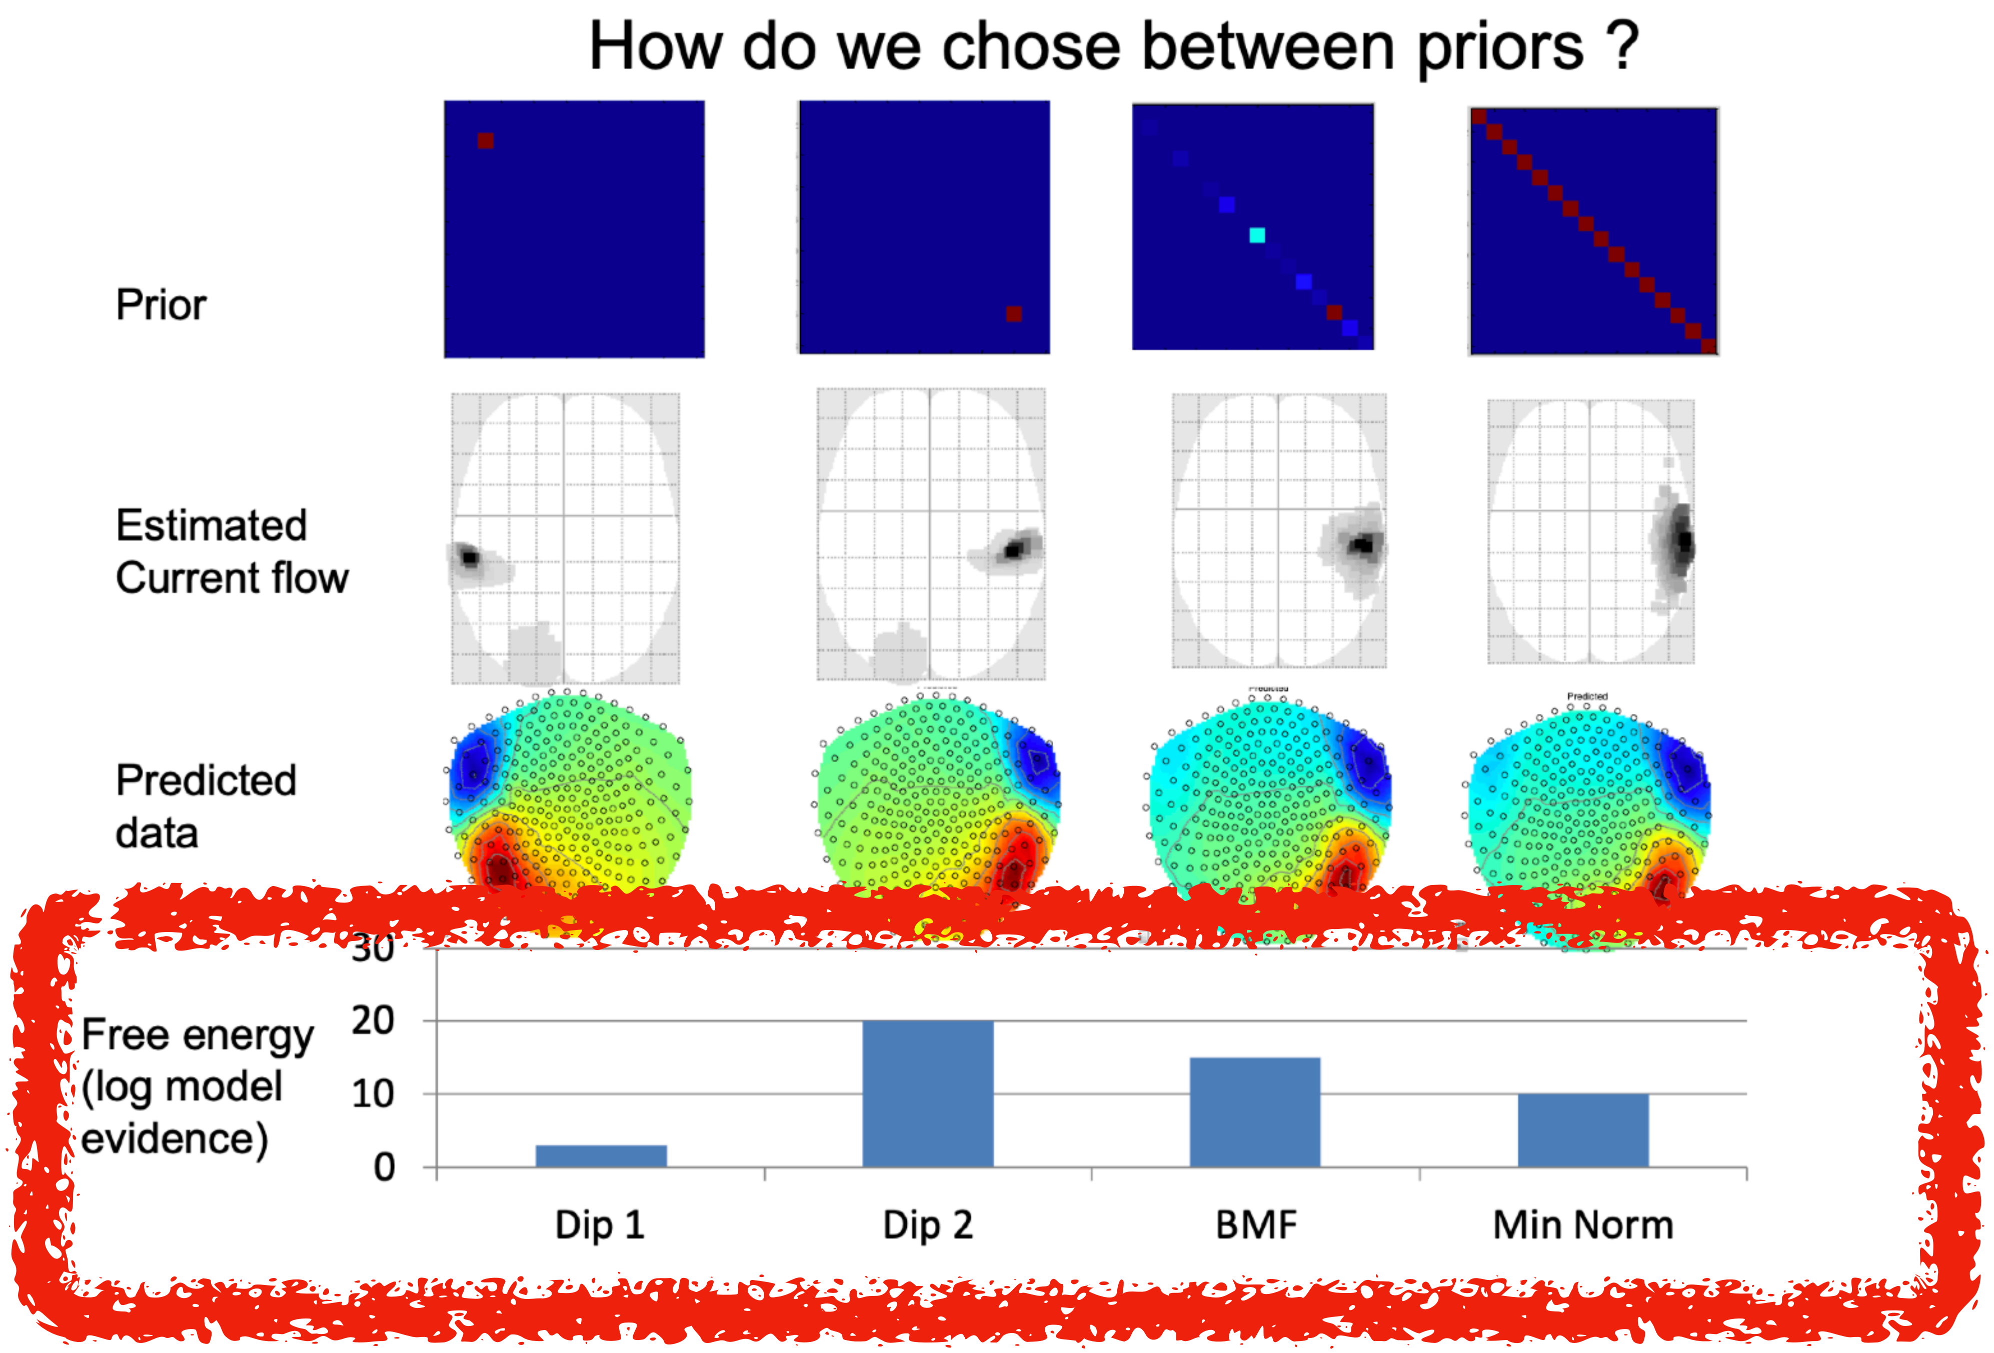
How do we chose between priors ?
Prior
Estimated
Current flow
Predicted
data
Free energy
(log model
evidence)
30
20
10
0
Dip 1
Dip 2
BMF
Min Norm

In practice

```
>> spm eeg
```

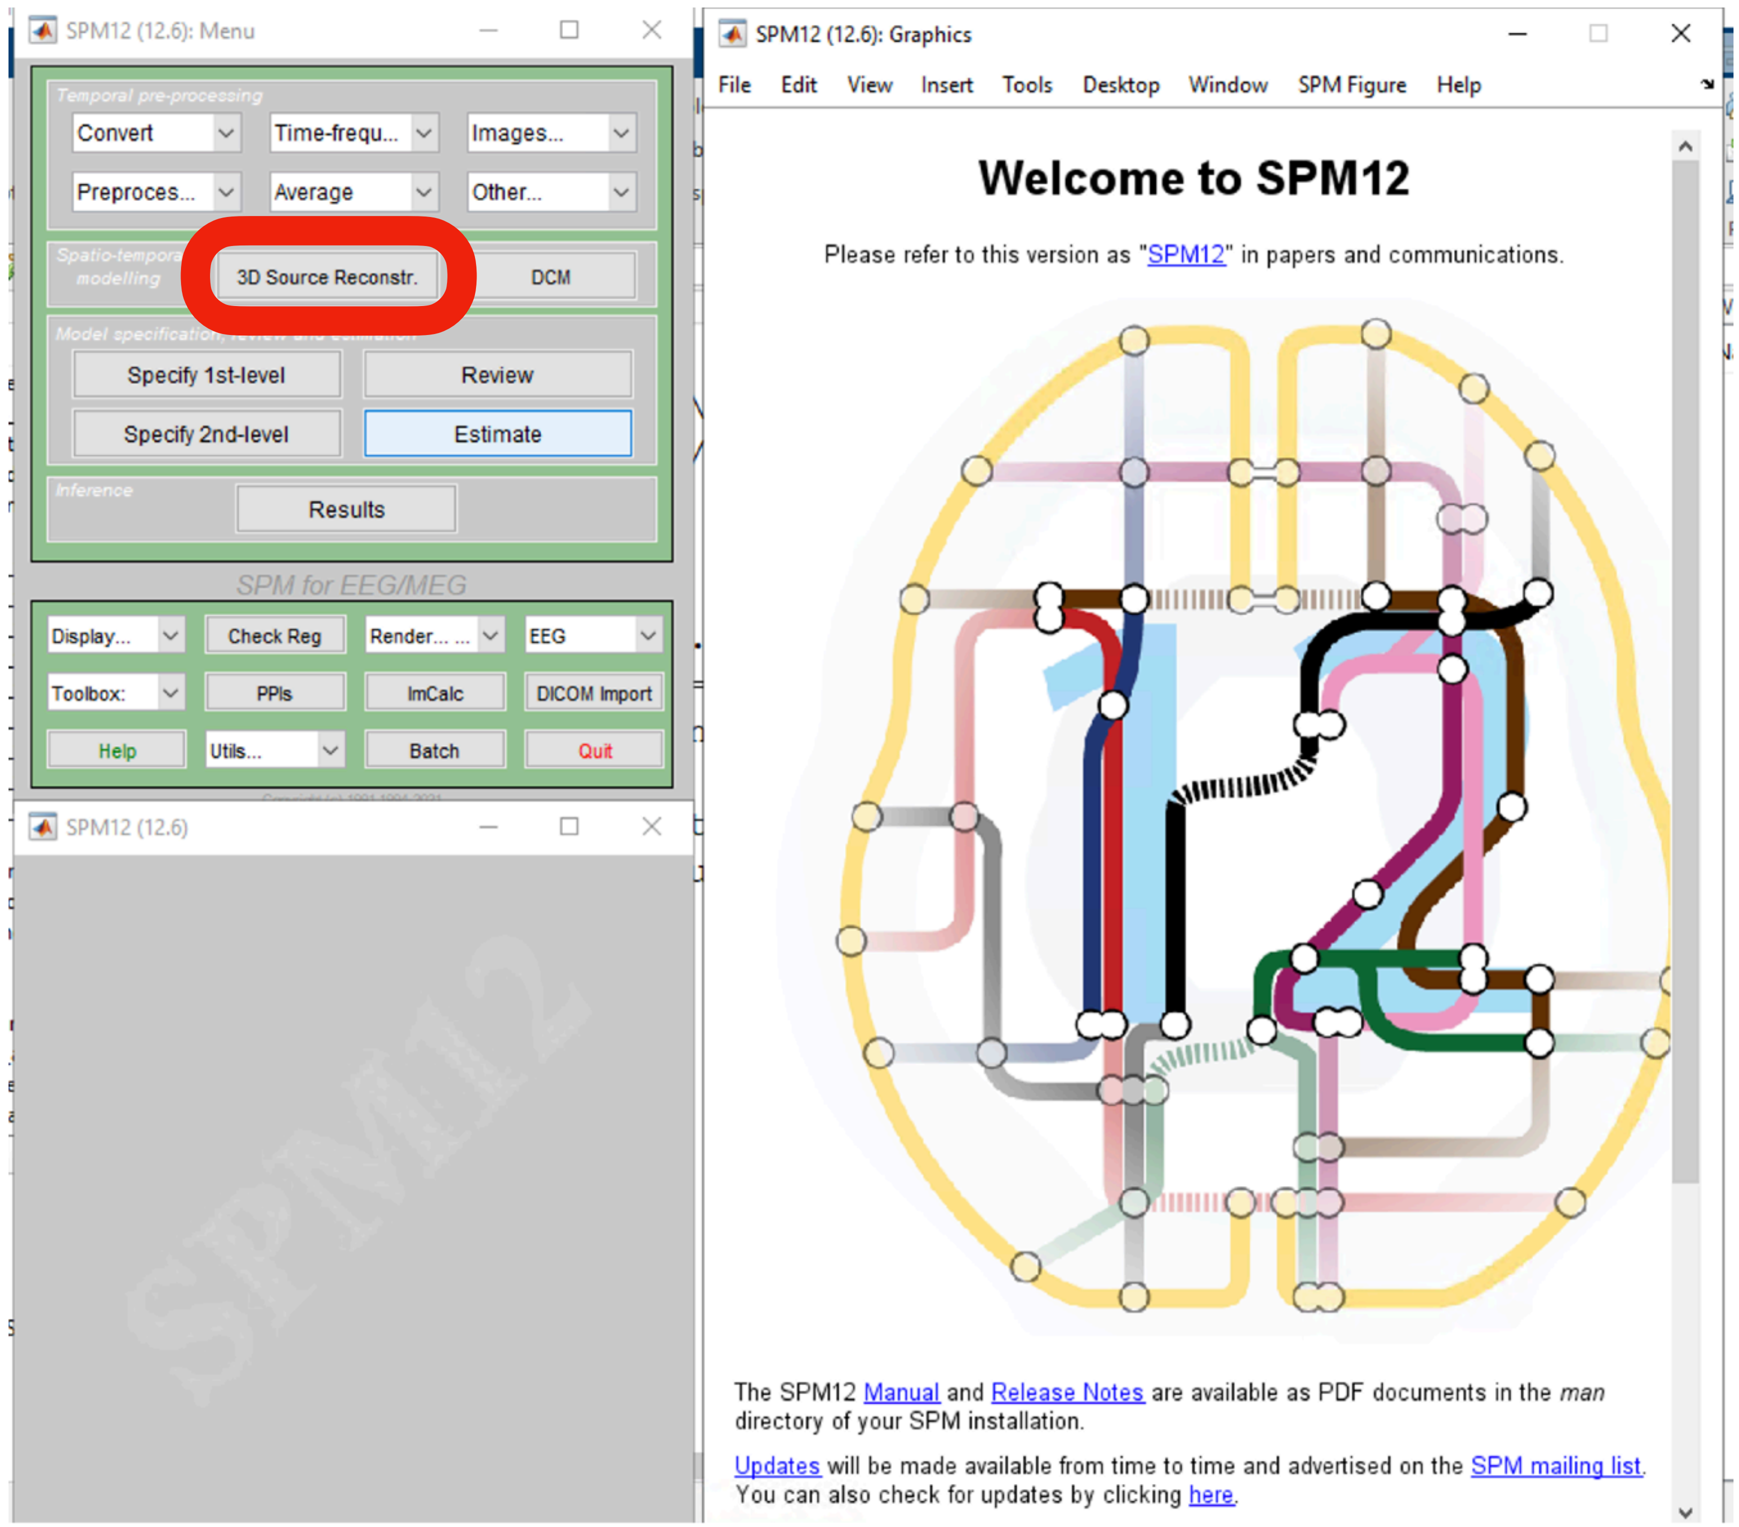

Not covered today

“Classic”/non-Bayesian source recon: beamformers, dipole fits, MUSIC etc

Group source reconstruction in SPM. See work by Wakeman and Henson.

Other software packages: MNE-Python, FieldTrip, EEGLab etc.

Other ways of quantifying performance. See work by Hauk et al., 2011

Practical pre-processing steps: coreg, forward model choices, exporting to NIFTI etc.

DAISS toolbox

References:

Bayesian
Spatial filters/other
Both

Brainstorm website: https://neuroimage.usc.edu/brainstorm/Introduction

YouTube tutorials (FieldTrip is good!)

López, J. D., Litvak, V., Espinosa, J. J., Friston, K., & Barnes, G. R. (2014). Algorithmic procedures for Bayesian MEG/EEG source reconstruction in SPM. *NeuroImage*, *84*, 476-487.

SPM12 Manual

Brookes, M. J., Vrba, J., Robinson, S. E., Stevenson, C. M., Peters, A. M., Barnes, G. R., ... & Morris, P. G. (2008). Optimising experimental design for MEG beamformer imaging. *Neuroimage*, *39*(4), 1788-1802.

Sekihara, K., & Nagarajan, S. S. (2015). Electromagnetic brain imaging: a bayesian perspective. Springer.

Sekihara, K., & Nagarajan, S. S. (2008). Adaptive spatial filters for electromagnetic brain imaging. Springer Science & Business Media.

Wipf, D., & Nagarajan, S. (2009). A unified Bayesian framework for MEG/EEG source imaging. *NeuroImage*, *44*(3), 947-966.

Friston, K., Harrison, L., Daunizeau, J., Kiebel, S., Phillips, C., Trujillo-Barreto, N., ... & Mattout, J. (2008). Multiple sparse priors for the M/EEG inverse problem. *NeuroImage*, *39*(3), 1104-1120.

Complexity

Thanks to:

Meaghan Spedden

Rob Seymour

Nic Alexander

Stephanie Mellor

Tim Tierney

George O’Neill

Gareth Barnes

Mansoureh Fahimi Hnazaee

Heather Payne

Vladimir Litvak

Karl Friston

The rest of the M/EEG SPM developers at the FIL

and you!

Questions: r.timms@ucl.ac.uk