Bayesian Inference in fMRI

Will Penny

Bayesian Approaches in Neuroscience

Karolinska Institutet, Stockholm

February 2016

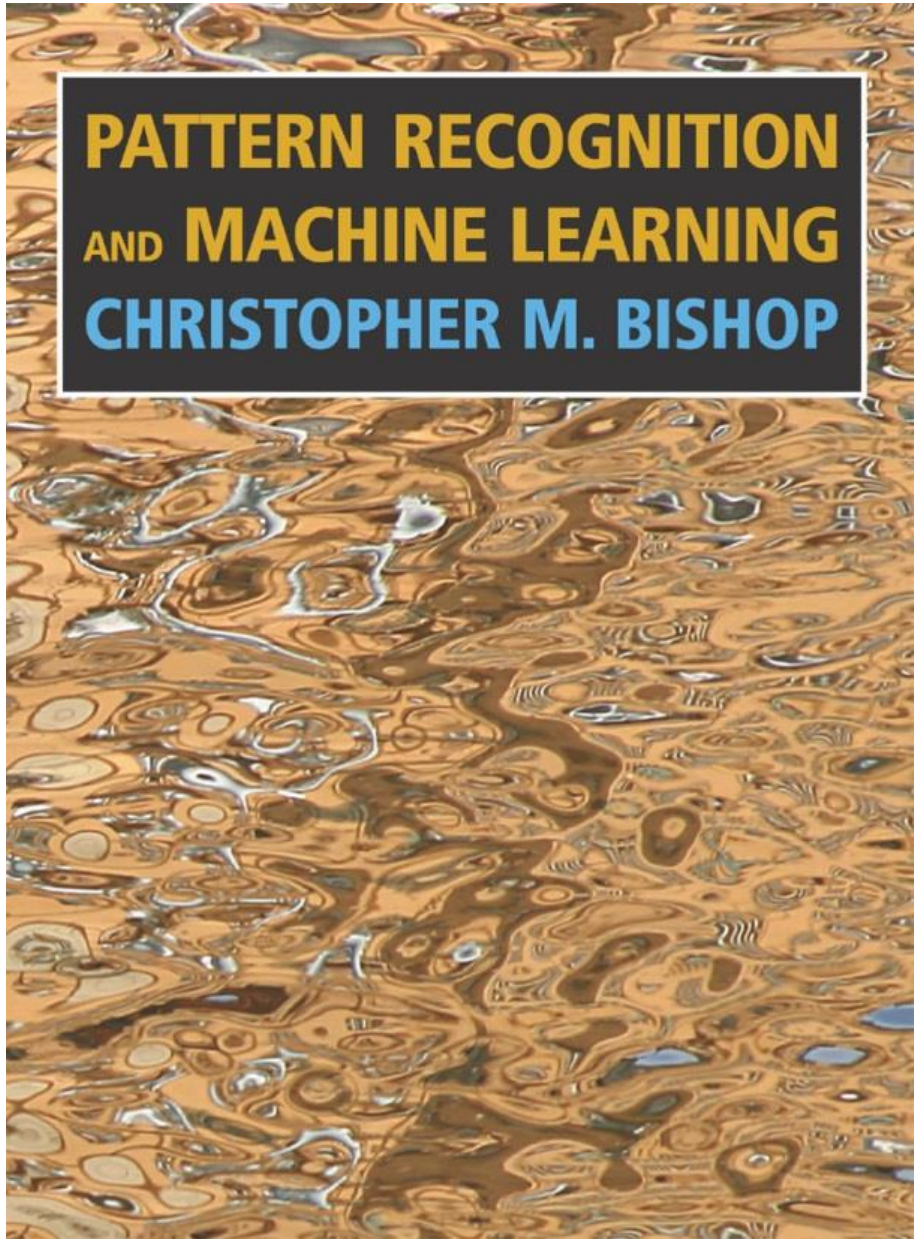
PATTERN RECOGNITION
AND MACHINE LEARNING
CHRISTOPHER M. BISHOP

NeuroImage xxx (2011) xxx–xxx

Contents lists available at SciVerse ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

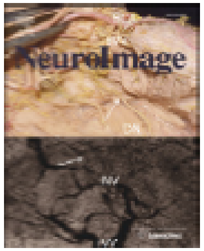

Review

Bayesian inference in FMRI

Mark W. Woolrich [a,b,*]

[a] *OHBA (Oxford Centre for Human Brain Activity), University of Oxford, UK*
[b] *FMRIB Centre, University of Oxford, UK*

ARTICLE INFO

Article history:
Received 4 August 2011
Revised 11 October 2011
Accepted 12 October 2011
Available online xxxx

Keywords:
fMRI
Bayes
Inference
Statistics
Probability
Priors

ABSTRACT

Bayesian inference has taken FMRI methods research into areas that frequentist statistics have struggled to reach. In this article we will consider some of the early forays into Bayes and what motivated its use. We shall see the impact that Bayes has had on haemodynamic modelling, spatial modelling, group analysis, model selection and brain connectivity analysis; and consider how these advancements have spun-off into related areas of neuroscience and some of the challenges that remain. Bayes has brought to the table inference flexibility, incorporation of prior information, adaptive regularisation and model selection. But perhaps more important than these things, is the ability of Bayes to empower the methods researcher with a mathematically principled framework for inferring on any model.

Overview

- Posterior Probability Maps
- Hemodynamic Response Functions
- Population Receptive Fields
- Computational fMRI
- Multivariate Decoding
- Dynamic Causal Modelling

Overview

- **Posterior Probability Maps**
- Hemodynamic Response Functions
- Population Receptive Fields
- Computational fMRI
- Multivariate Decoding
- Dynamic Causal Modelling

Bayes Rule for Gaussians

Likelihood and Prior

$$p(y \mid \theta^{(1)}) = N(\theta^{(1)}, \lambda^{(1)})$$

$$p(\theta^{(1)}) = N(\theta^{(2)}, \lambda^{(2)})$$

Posterior

$$p(\theta^{(1)} \mid y) = N(m, P)$$

$$P = \lambda^{(1)} + \lambda^{(2)}$$

$$m = \frac{\lambda^{(1)}}{P}\theta^{(1)} + \frac{\lambda^{(2)}}{P}\theta^{(2)}$$

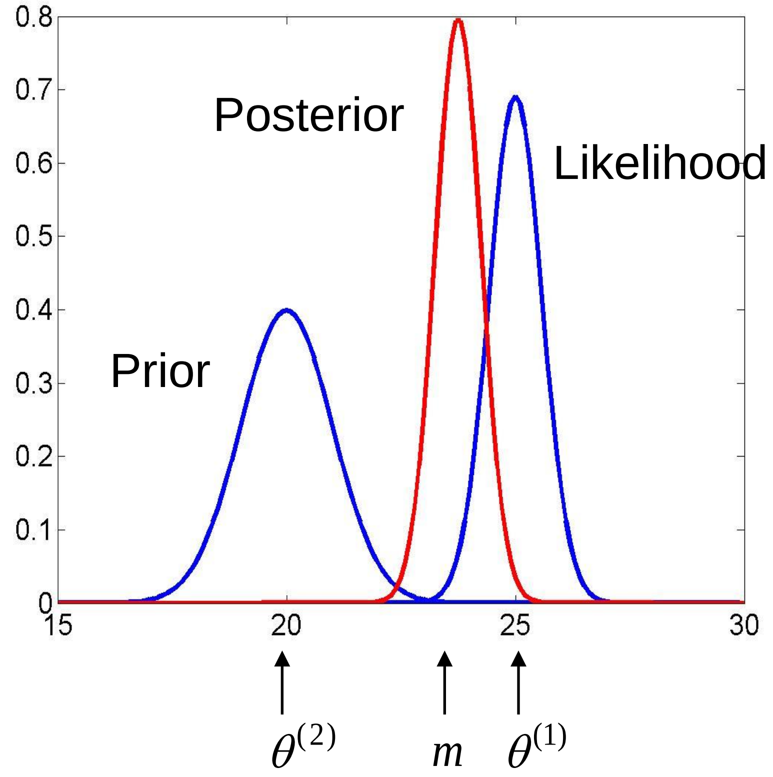

Relative Precision Weighting

Global Shrinkage Priors

$$p(\beta) = N(0, \alpha^{-1} I)$$

$$Y = X\beta + \varepsilon$$

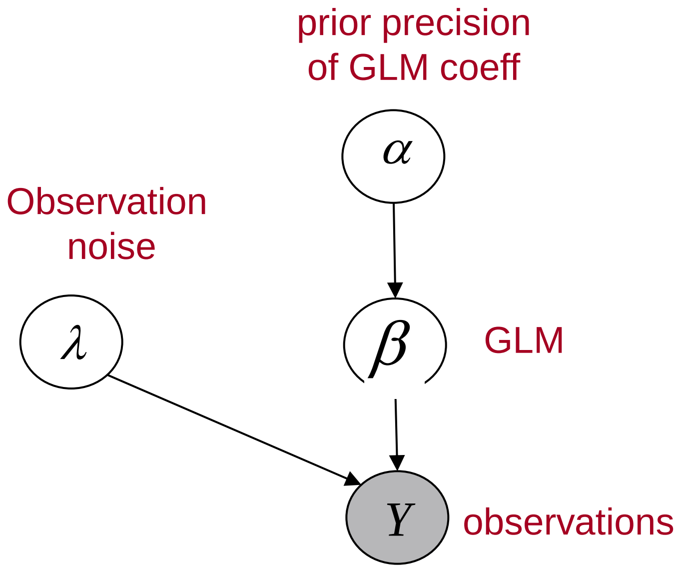

K.J. Friston and W.D. Penny. **Posterior probability maps and SPMs**. *NeuroImage*, 19(3):1240-1249, 2003.

Posterior

Posterior distribution: probability of the effect given the data

$$p(\beta \mid y)$$

Posterior Probability Map: images of the probability that an activation exceeds some specified threshold s_{th}, given the data y

$$p(\beta > s_{th} \mid y) > p_{th}$$

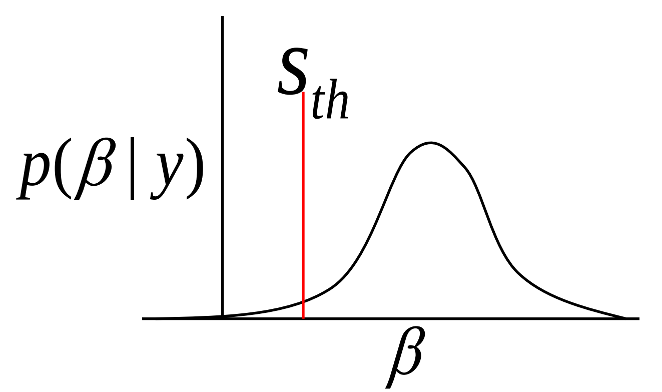

Two thresholds:

- activation threshold s_{th} : percentage of whole brain mean signal (physiologically relevant size of effect)
- probability p_{th} that voxels must exceed to be displayed (e.g. 95%)

PPM

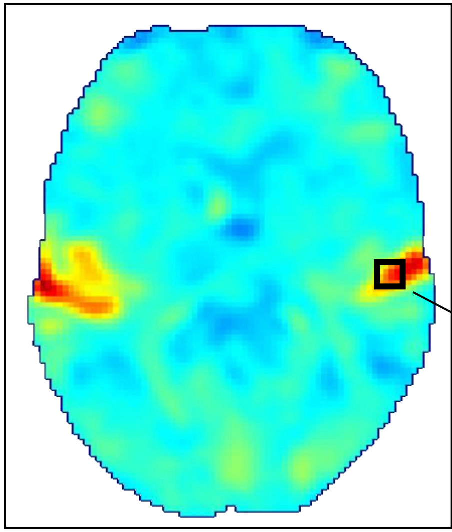

Mean *(Cbeta_*.img)*

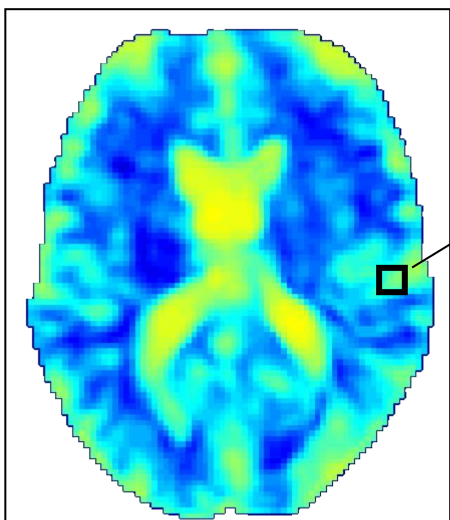

Std dev *(SDbeta_*.img)*

activation
threshold
s_{th}

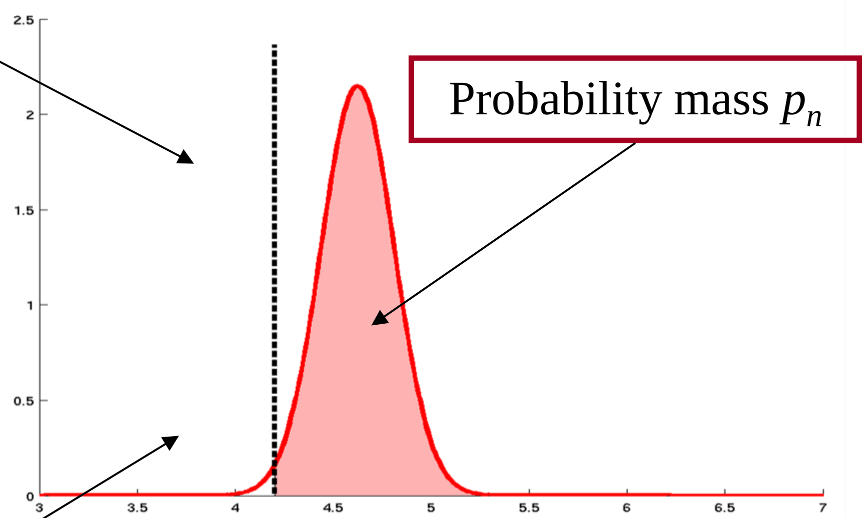

Posterior density $q(\beta_n)$

probability of getting an effect, given the data

$$q(\beta_n) = N(\mu_n, \Sigma_n)$$

mean: *size of effect*
covariance: *uncertainty*

Display only voxels
that exceed e.g. 95%

$$p > p_{th}$$

$$p = q(\beta > s_{th})$$

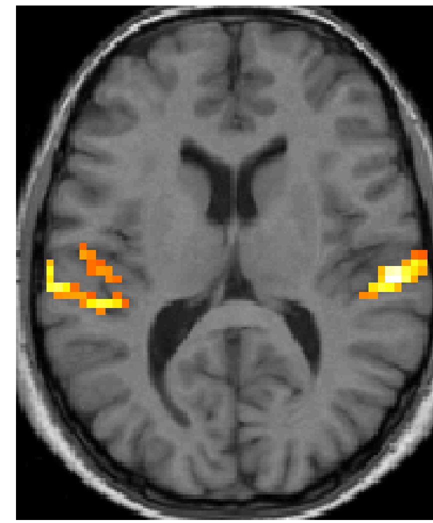

PPM *(spmP_*.img)*

Choice of Priors

Stationary smoothness:

W.D. Penny, N. Trujillo-Barreto, and K.J. Friston. **Bayesian fMRI time series analysis with spatial priors**. *NeuroImage*, 24(2):350-362, 2005.

Nonstationary smoothness:

L M Harrison, W Penny, J Daunizeau, and K J Friston.
Diffusion-based spatial priors for functional magnetic resonance images. *Neuroimage*, 41(2):408-23, 2008.

Global Shrinkage:

K.J. Friston and W.D. Penny. **Posterior probability maps and SPMs**. *NeuroImage*, 19(3):1240-1249, 2003.

Stationary Smoothness Priors

$$Y = X\beta + \varepsilon$$

$$p(\beta) = N\left(0, \alpha^{-1}L\right)$$

aMRI

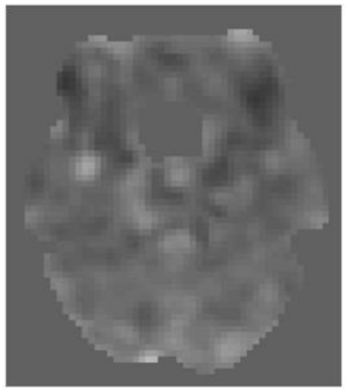

Smooth Y

ML

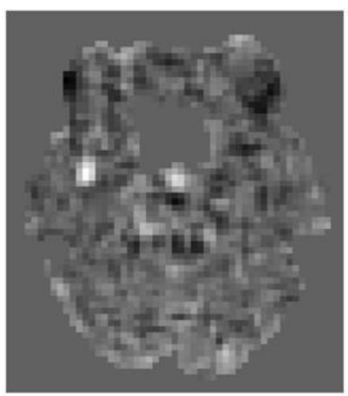

Posterior

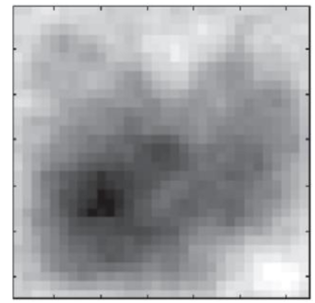

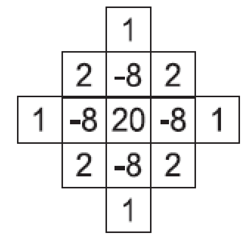

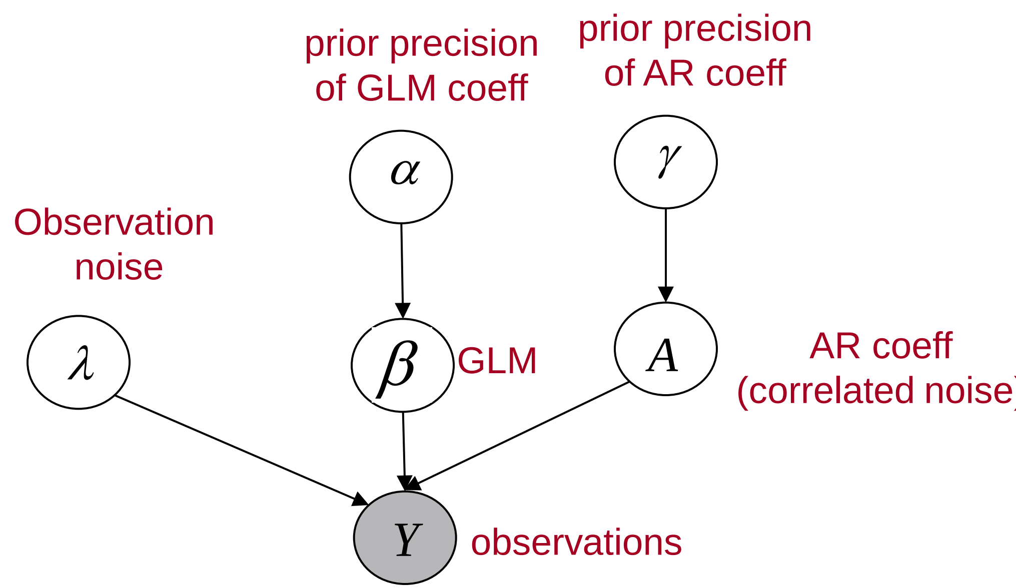

Overview

- Posterior Probability Maps
- **Hemodynamic Response Functions**
- Population Receptive Fields
- Computational fMRI
- Multivariate Decoding
- Dynamic Causal Modelling

K Friston et al. Event-Related fMRI: **Characterizing differential responses**, *Neuroimage* 7, 30-40, 1998

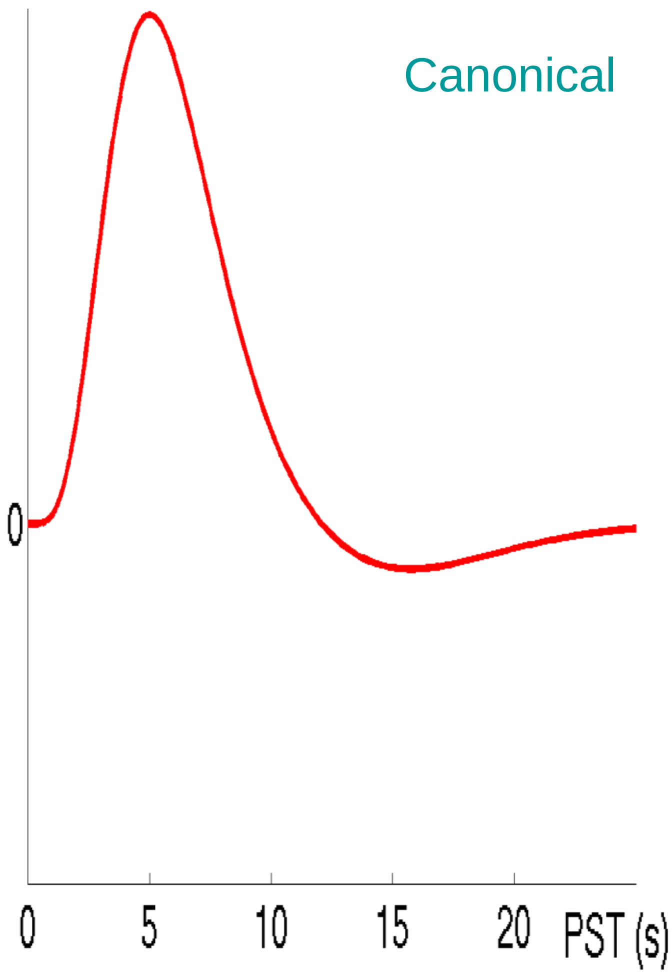

Two Gamma functions fitted to data from auditory cortex.

“Canonical” function f(w,t) with w width and t time.

Hemodynamic Response Functions

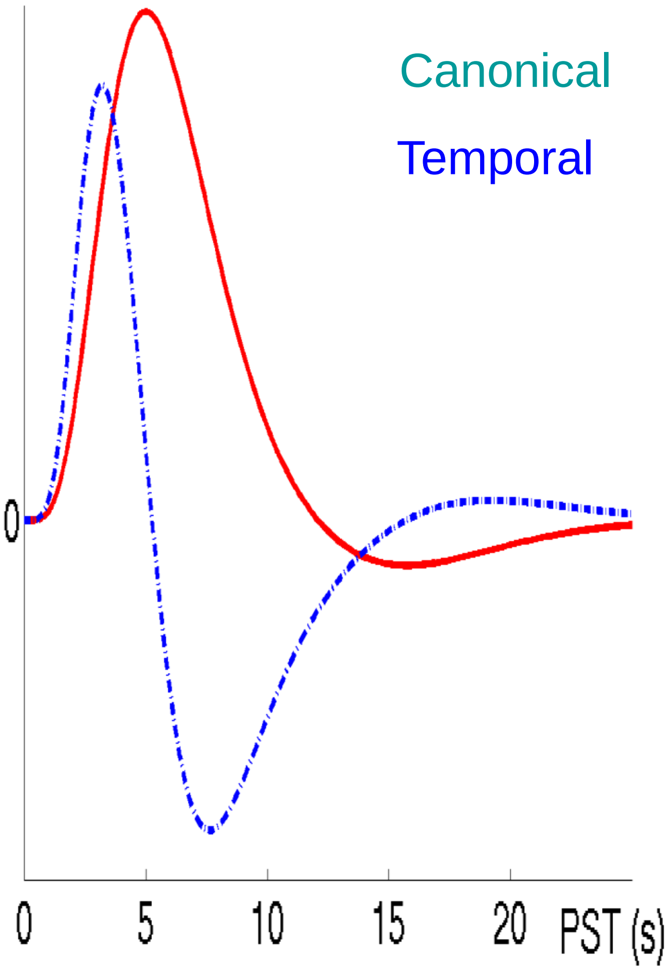

Temporal derivative, df/dt

Hemodynamic Response Functions

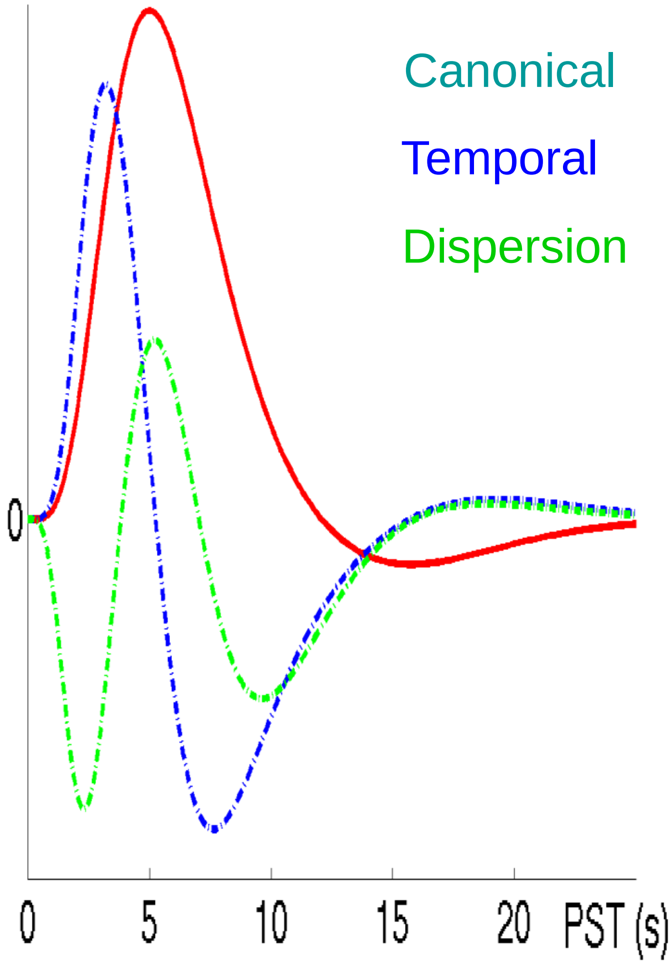

Dispersion derivative, df/dw

Hemodynamic Response Functions

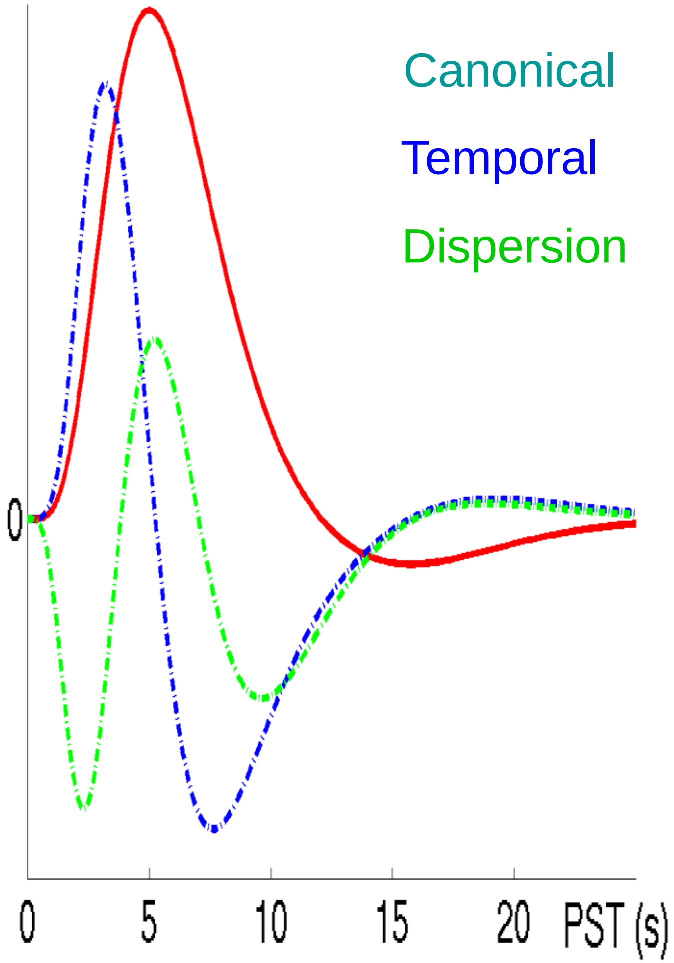

These three functions together comprise an "Informed Basis Set"

Finite Impulse Response (FIR)

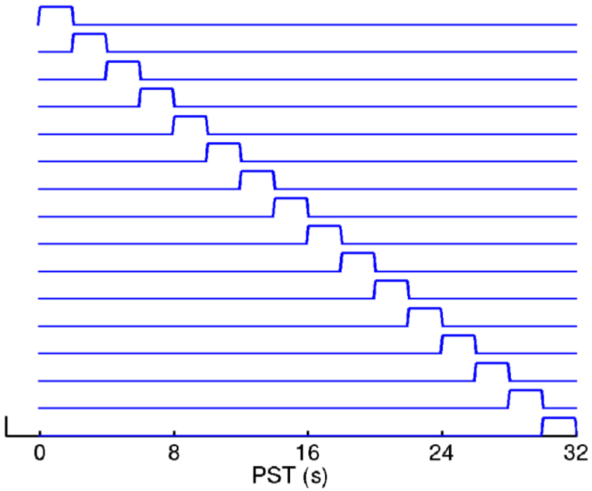

Fourier (F)

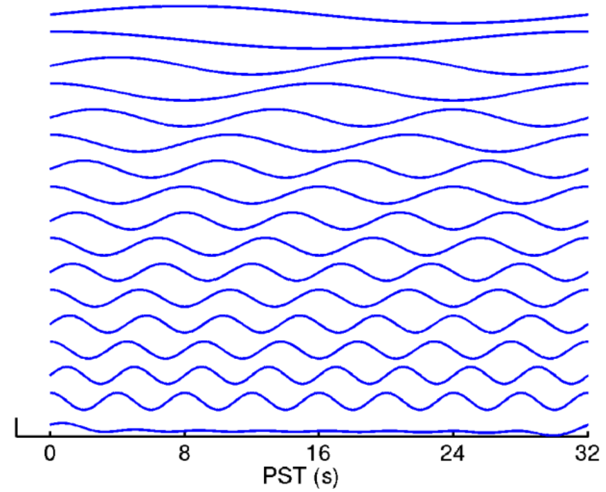

Gamma

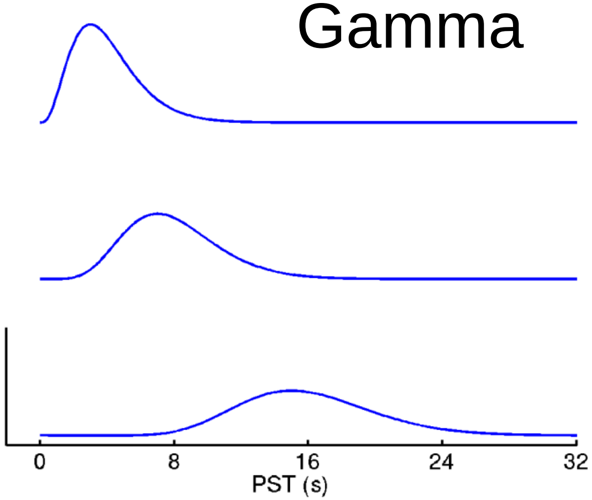

Informed (Inf)

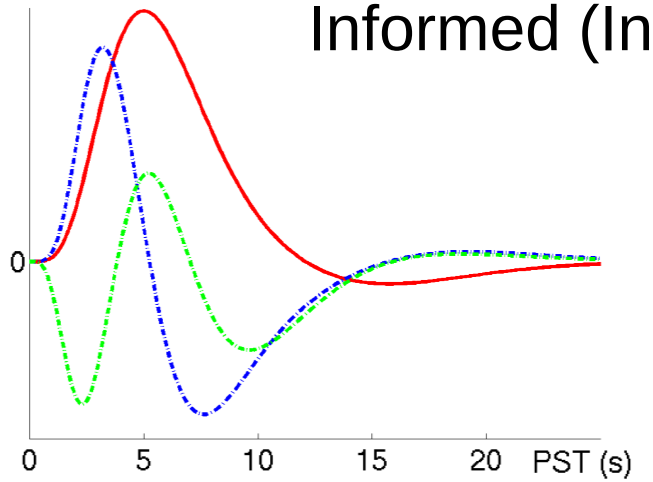

Hemodynamic Response Functions

W.D. Penny, G Flandin and N. Trujillo-Barreto. **Bayesian Comparison of Spatially Regularised General Linear Models** . *HBM*, 28:275-293, 2007.

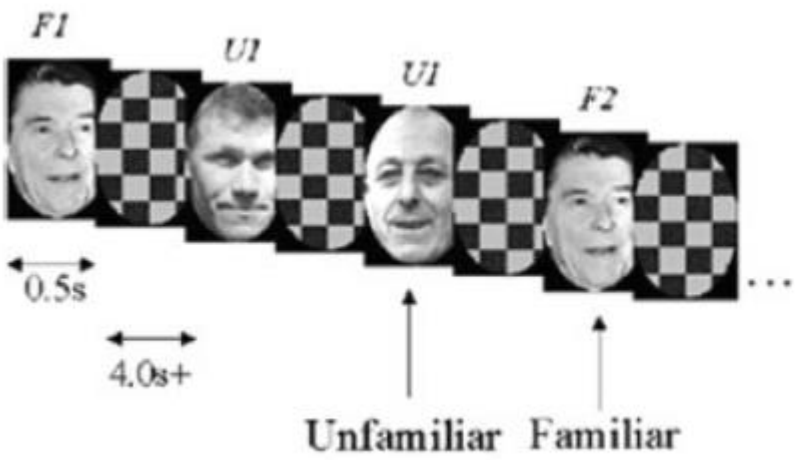

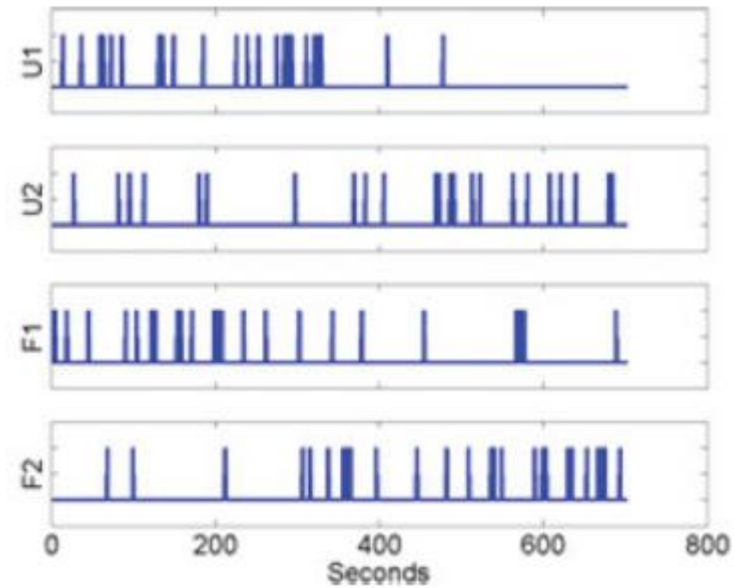

R Henson et al. **Face repetition effects in implicit and explicit memory tests as measured by fMRI**. *Cerebral Cortex*, 12:178-186.

Bayesian Model Comparison

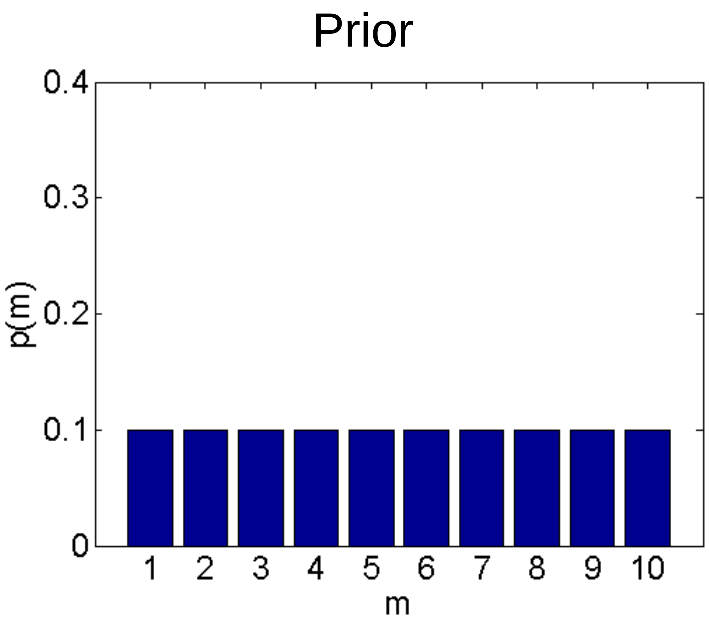

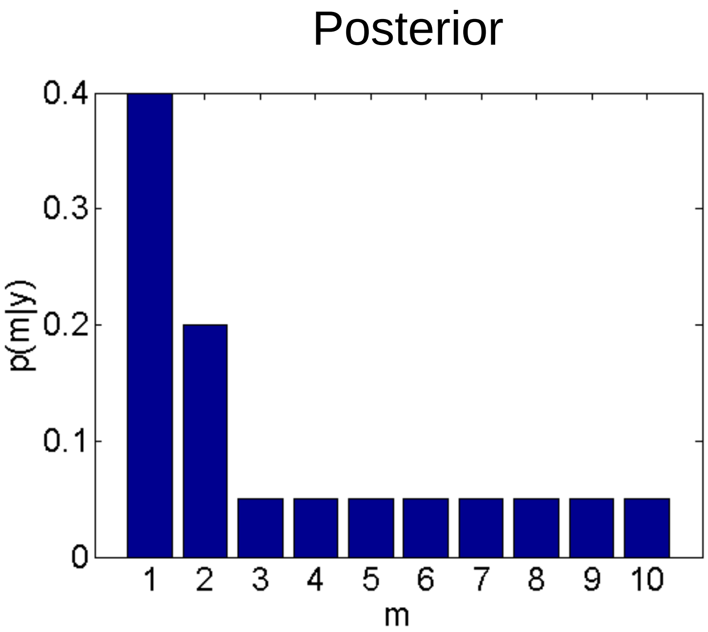

$$p(m|y) = \frac{p(y|m)p(m)}{\sum_{m'} p(y|m')p(m')}$$

Log Evidence = log p(y|m)

Hemodynamic Response Functions

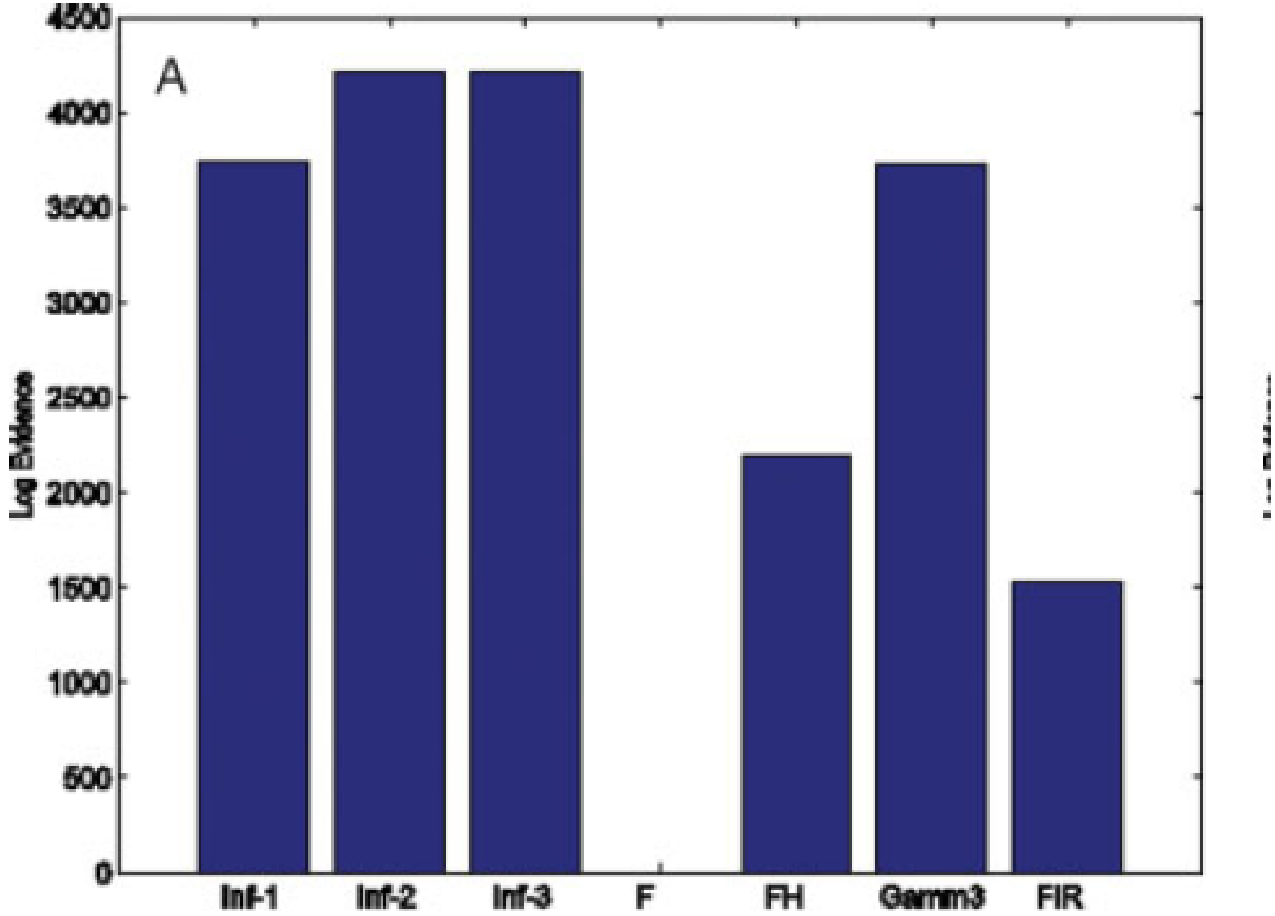

Left Occipital Cortex: Inf-2 is the preferred model

Hemodynamic Response Functions

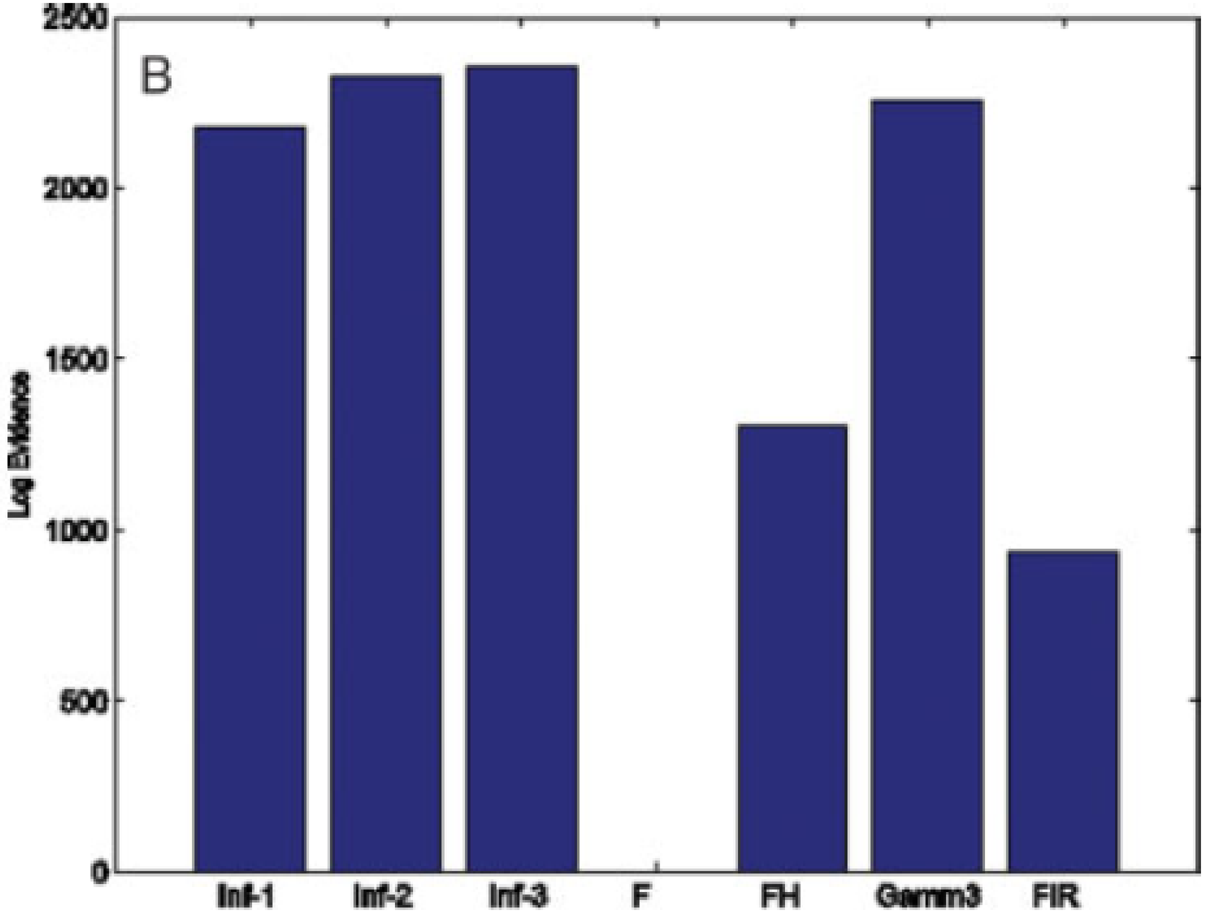

Right Occipital Cortex: Inf-3 is the preferred model

Hemodynamic Response Functions

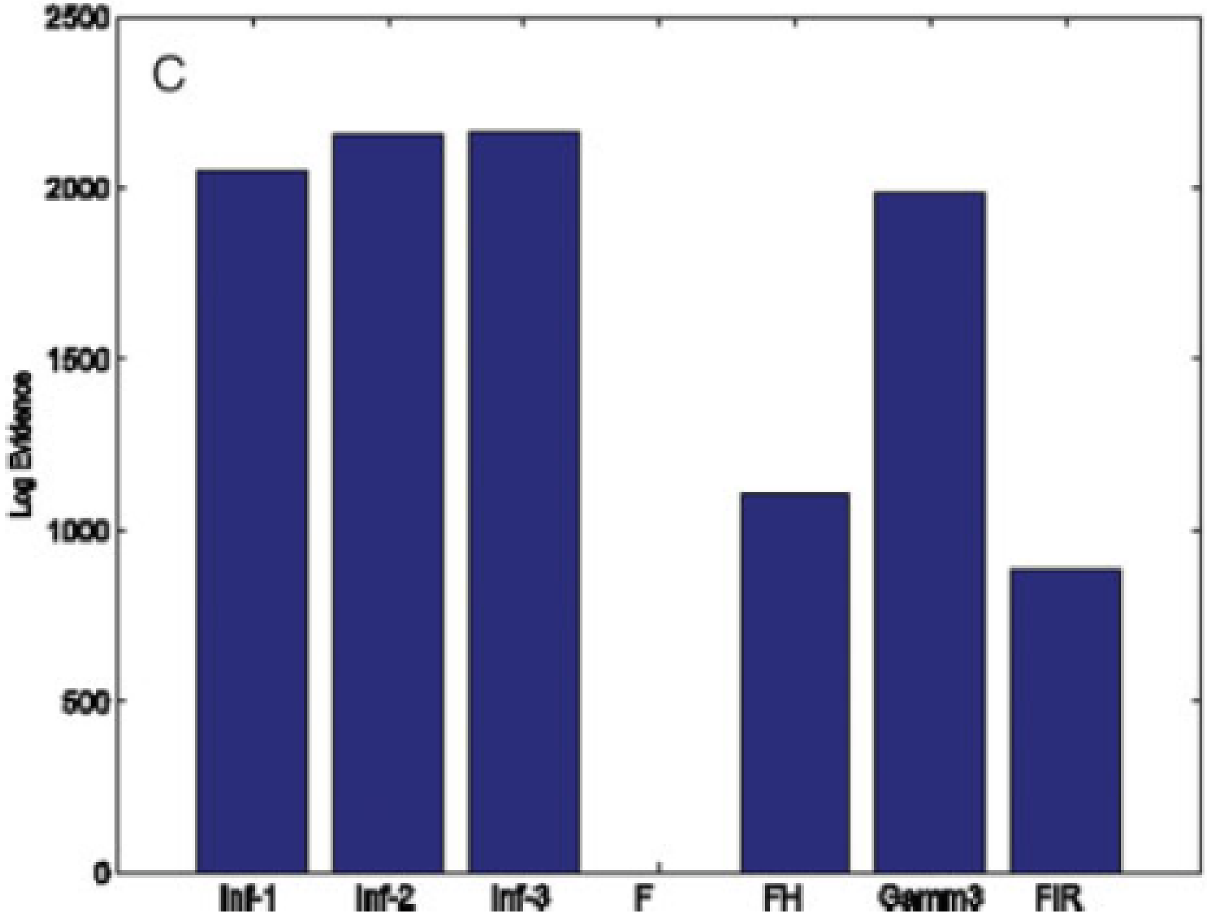

Sensorimotor Cortex: Inf-3 is the preferred model

K Friston. **Bayesian Estimation of Dynamical Systems: An application to fMRI**, *Neuroimage* 16, 513-530, 2002

Hemodynamic variables

$$\mathbf{x} = [s, f, v, q]$$

Dynamics

$$\dot{\mathbf{x}} = g(\mathbf{x}, z, \mathbf{h})$$

$$y = b(\mathbf{x})$$

Hemodynamic parameters

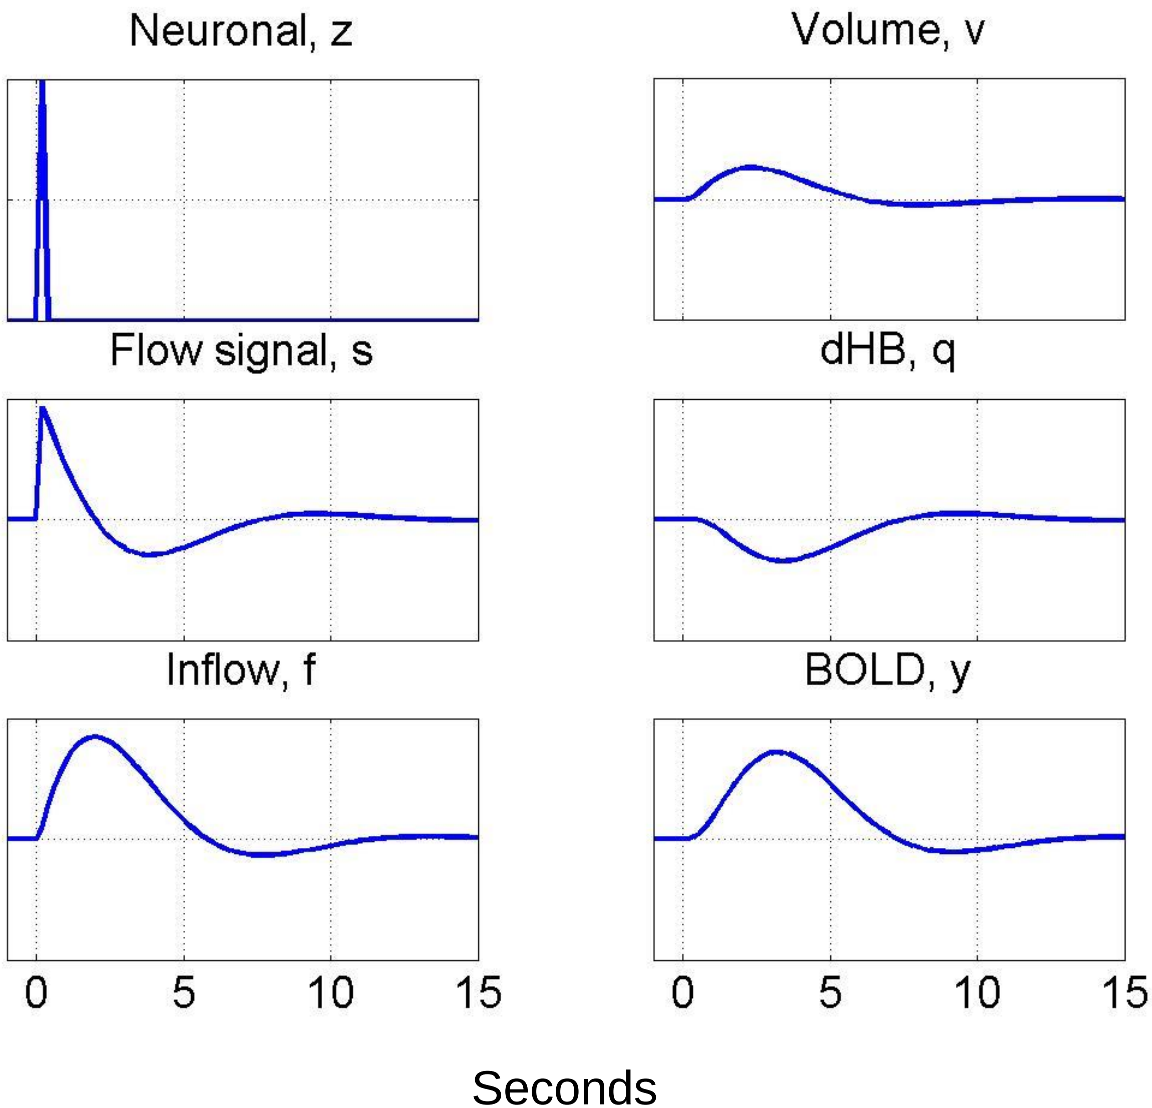

Hemodynamic Response Functions

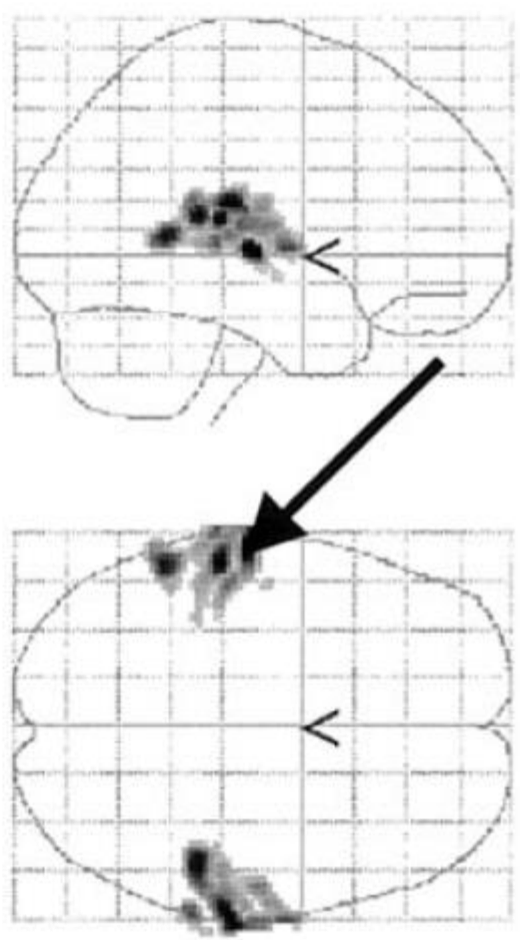

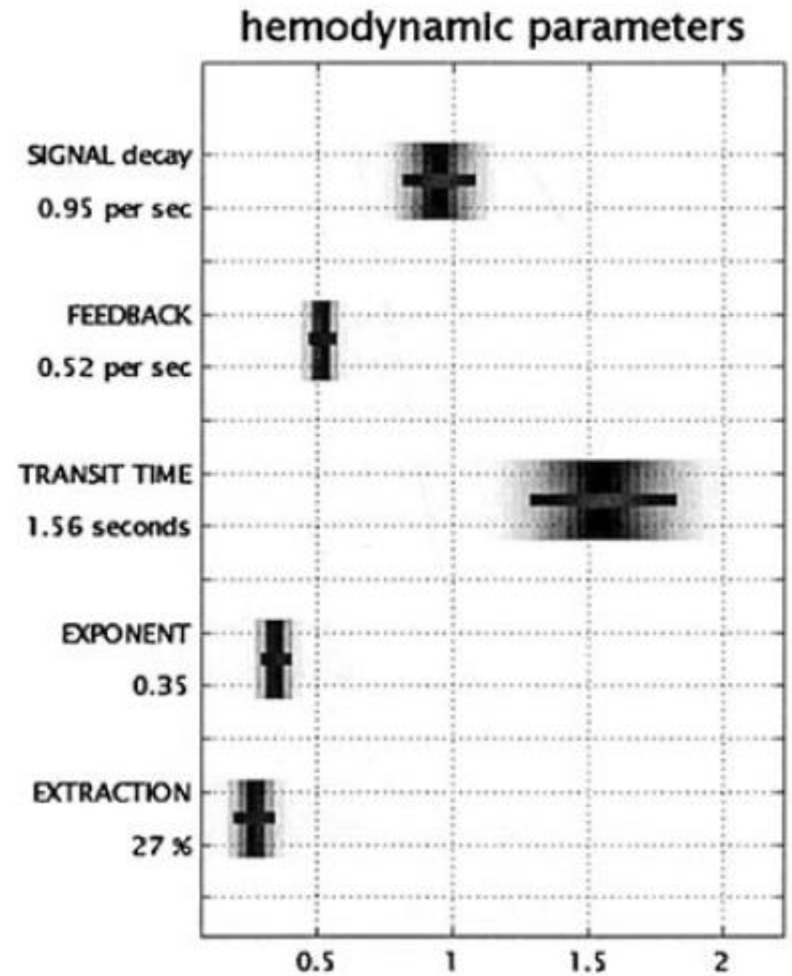

R Buxton et al. **Dynamics of Blood Flow and Oxygenation Changes During brain activation: The Balloon Model** , *Magnetic Resonance in Medicine, 39:855-864.*

Overview

- Posterior Probability Maps
- Hemodynamic Response Functions
- **Population Receptive Fields**
- Computational fMRI
- Multivariate Decoding
- Dynamic Causal Modelling

Population Receptive Fields

S. Kumar and W. Penny (2014). **Estimating Neural Response Functions from fMRI.** *Frontiers in Neuroinformatics, 8th May, doi: 10.3389/fninf.2014.00048.*

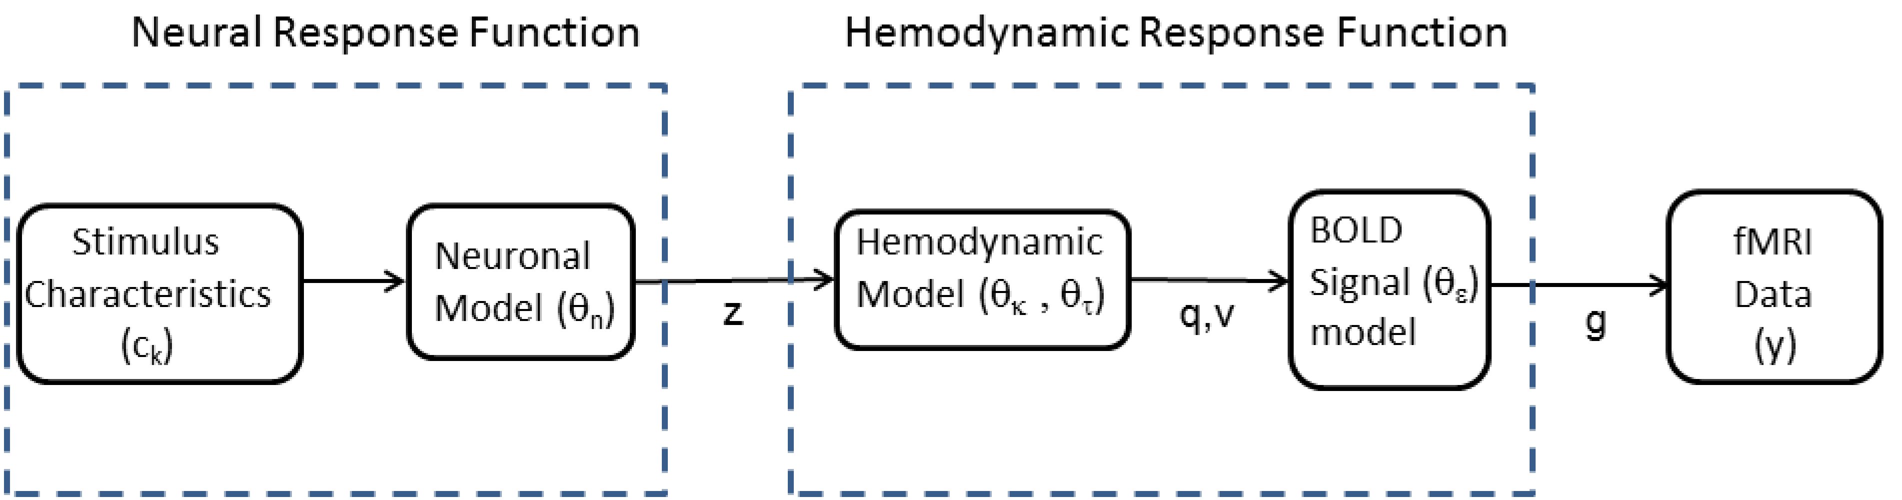

K Friston et al. (2007) **Variational free energy and the Laplace approximation**. *Neuroimage*, 34, 220–234.

Gaussian Population Receptive Fields

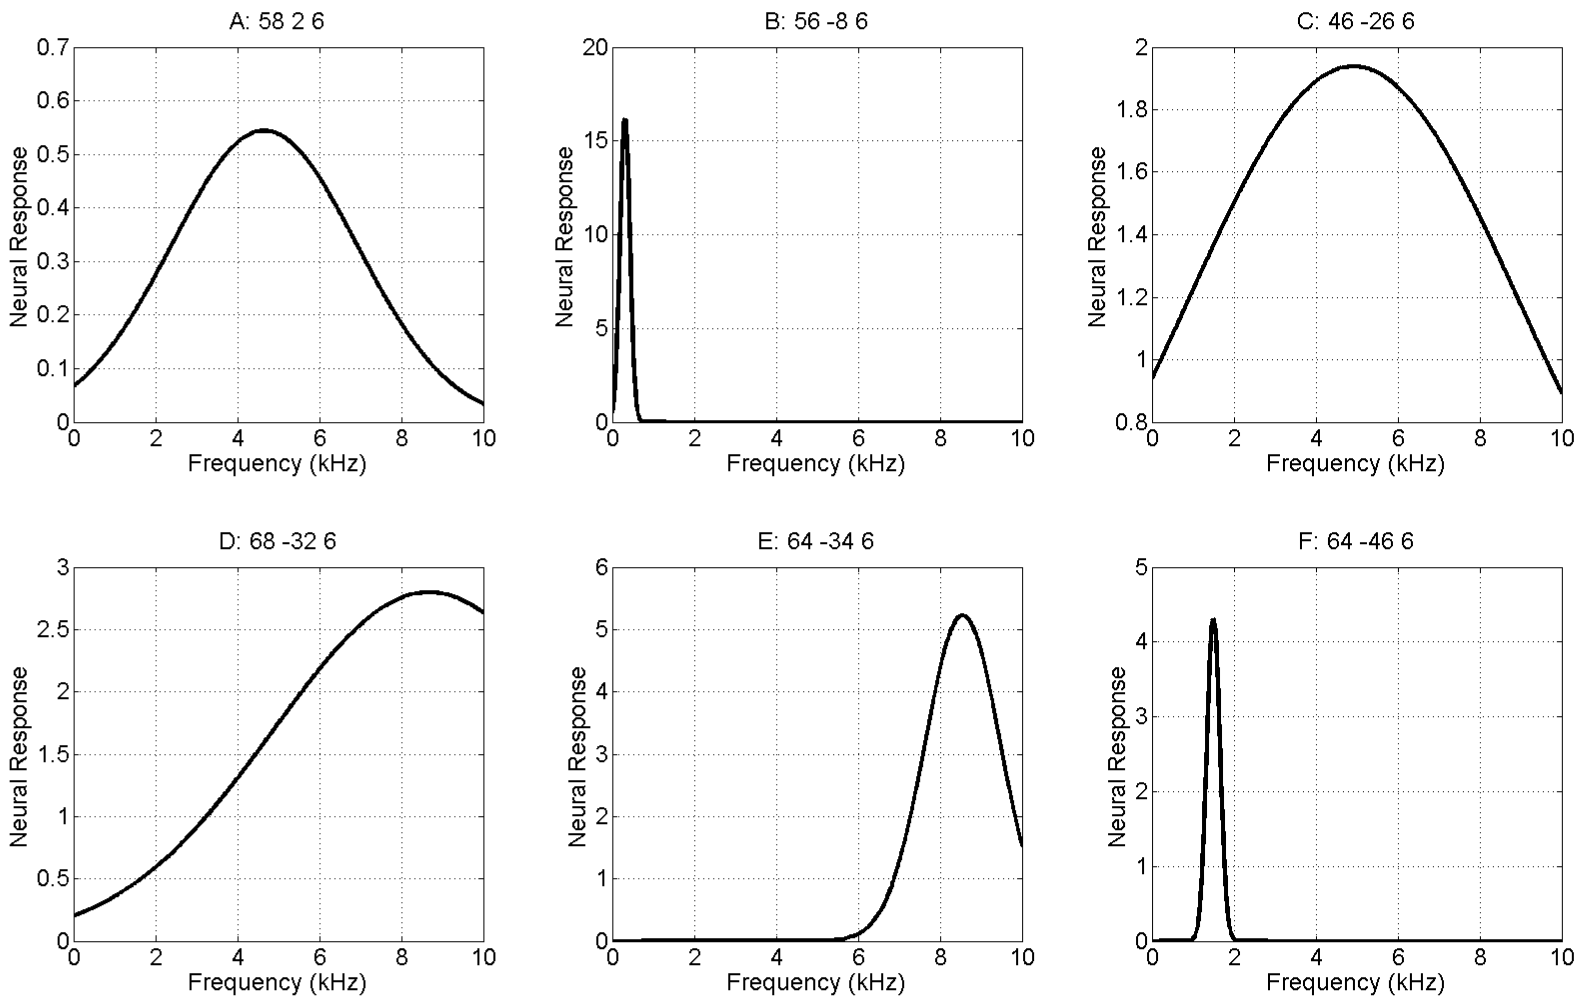

Gaussian Population Receptive Fields

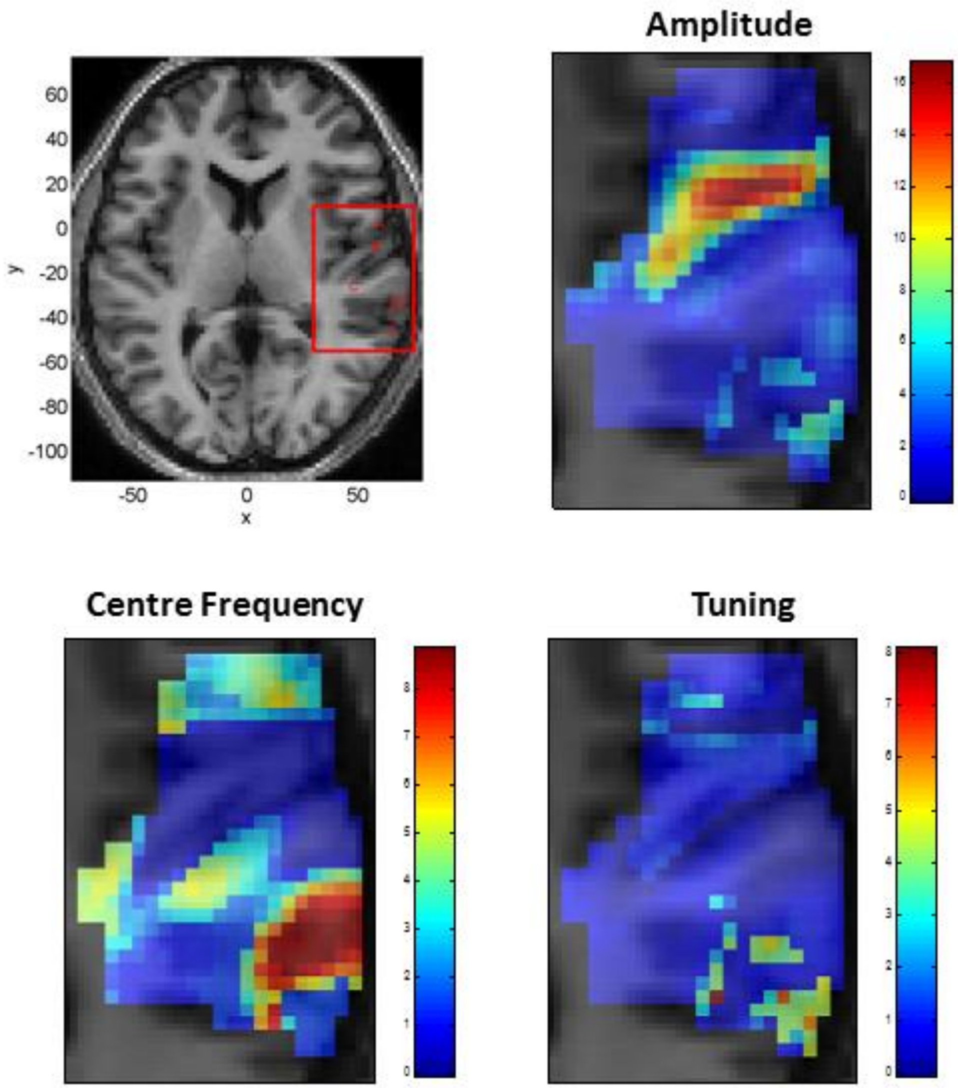

Mexican-Hat Population Receptive Fields

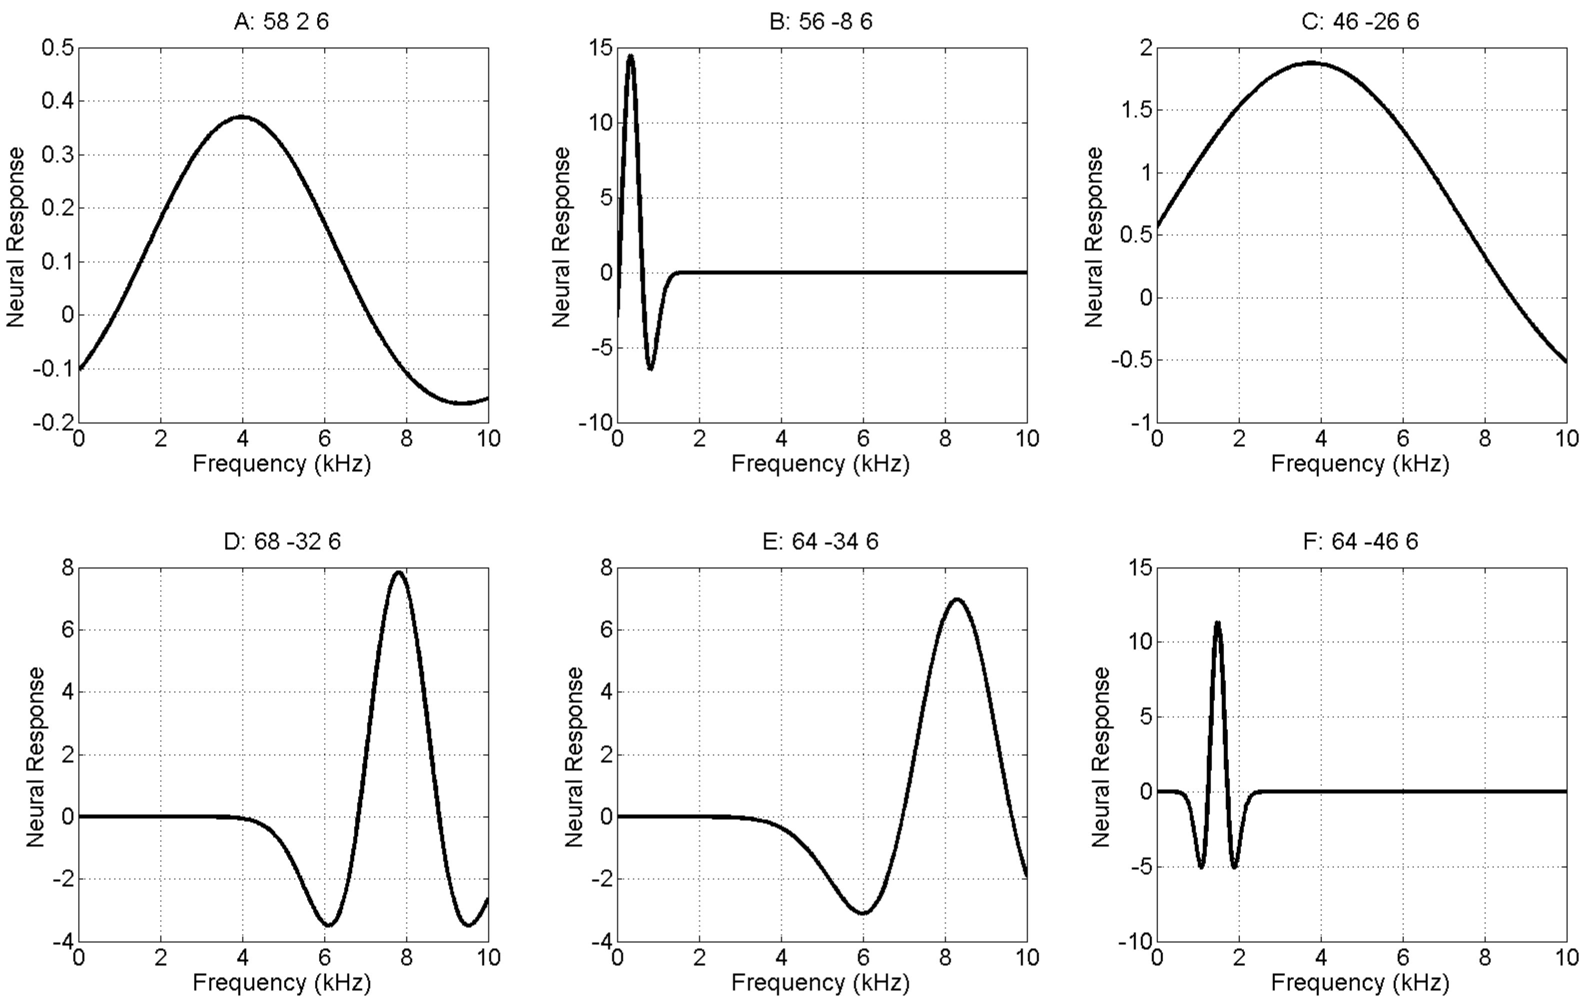

Mexican-Hat Population Receptive Fields

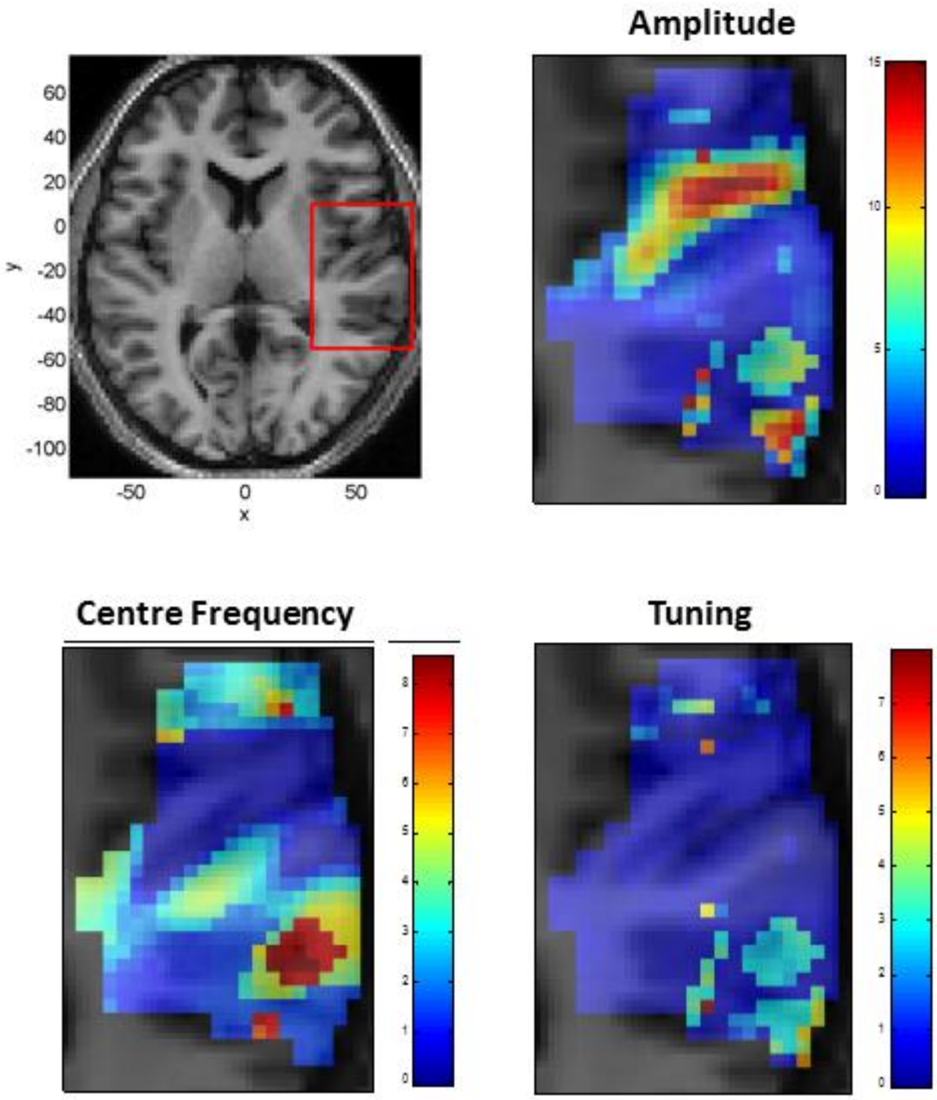

Which Parametric Function is a Better Descriptor ?

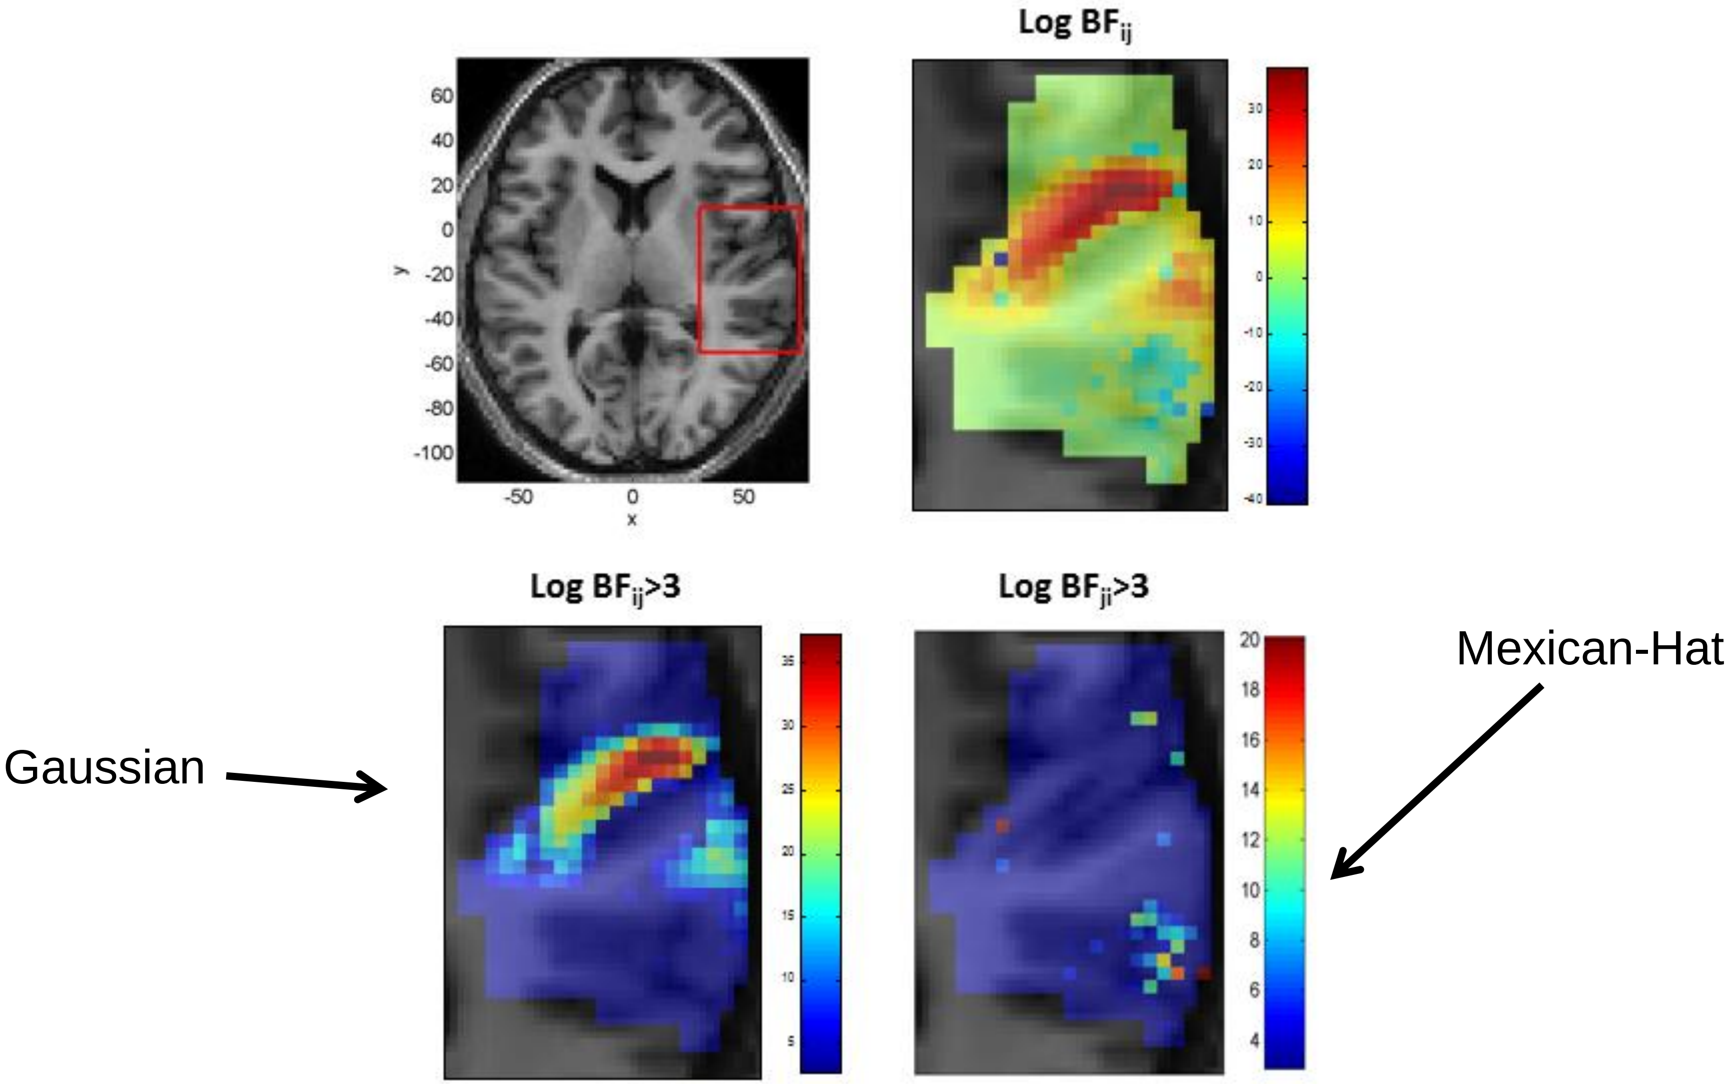

Overview

- Posterior Probability Maps
- Hemodynamic Response Functions
- Population Receptive Fields
- **Computational fMRI**
- Multivariate Decoding
- Dynamic Causal Modelling

Computational fMRI

Subjects pressed 1 of 4 buttons depending on the category of visual stimulus.

The 4 categories of stimuli occurred with different frequencies over a session.

Brain responses are then hypothesised to be proportional to the surprise, *S*, associated with each stimulus where *S=log(1/p)*.

But over what time scale is the probability *p* estimated ? And do different brain regions use different time scales ?

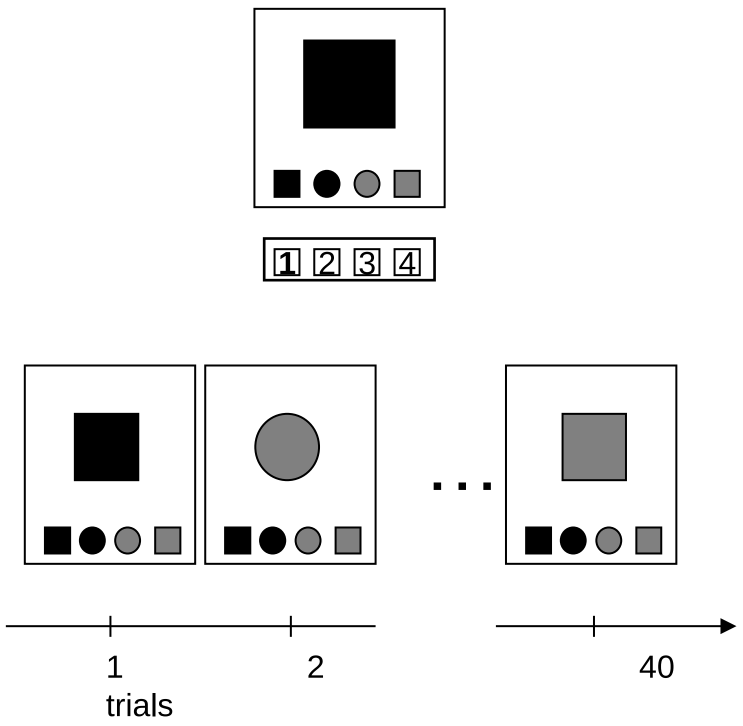

L. Harrison, S Bestmann, M. Rosa, W. Penny and G. Green (2011). **Time scales of representation in the human brain: weighing past information to predict future events.** *Frontiers in Human Neuroscience*, 5, 00037.

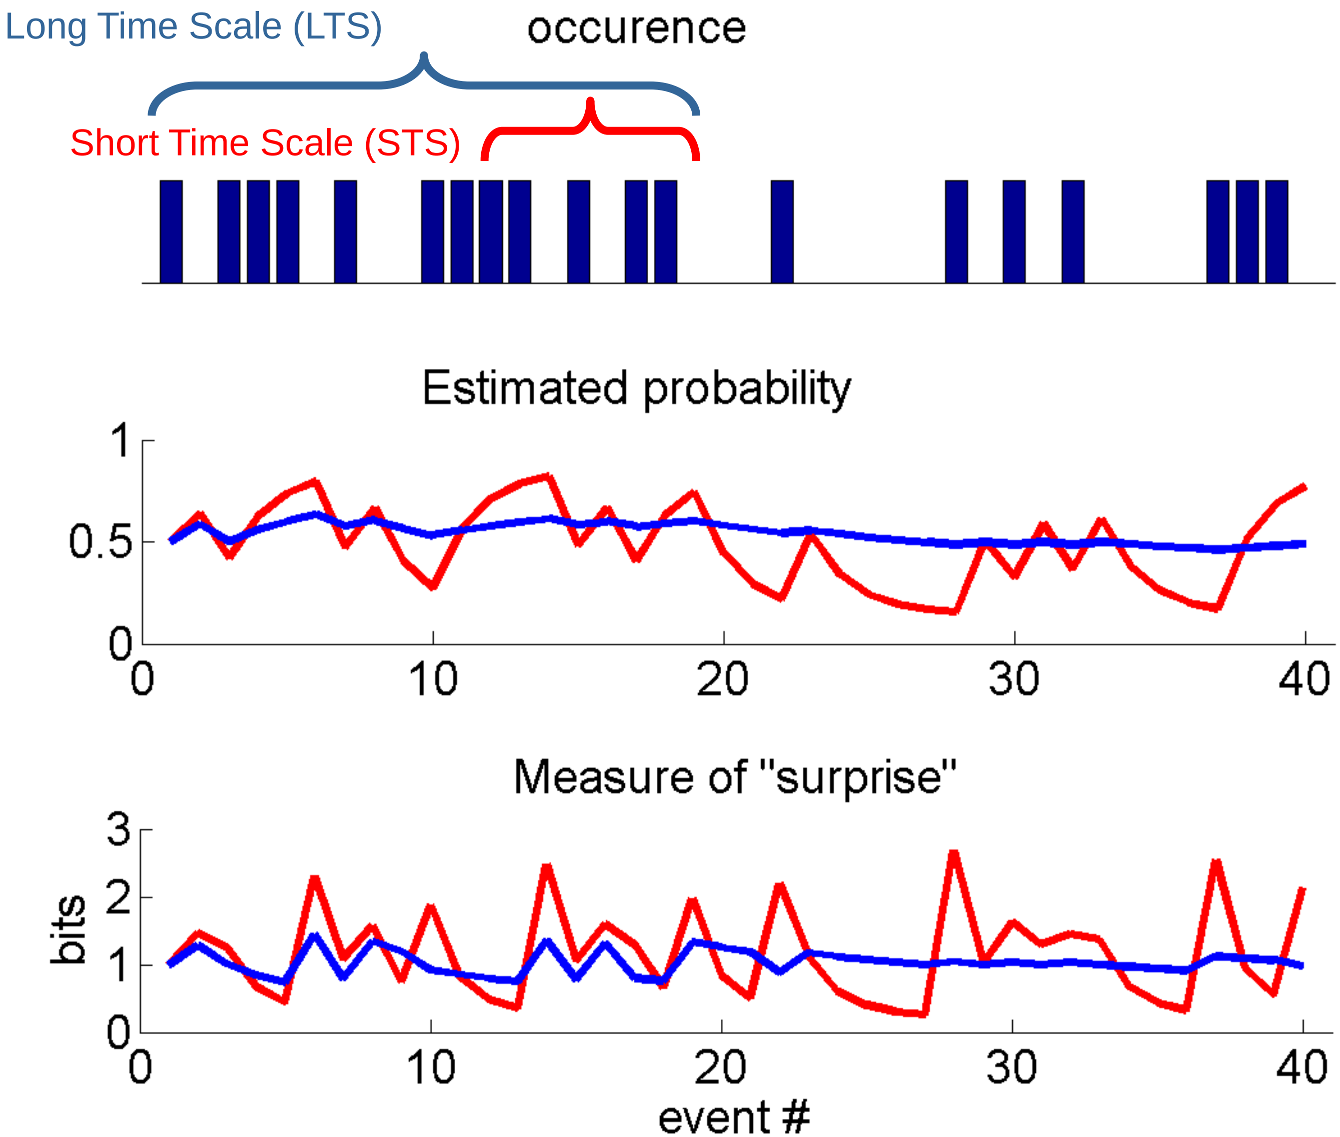

Enter surprise as a Parametric Modulator in first level GLM analysis.
Which surprise variable (STS or LTS) underlies the best model of fMRI responses?

M Rosa, S.Bestmann, L. Harrison, and W Penny. **Bayesian model selection maps for group studies**. *Neuroimage*, Jan 1 2010; 49(1):217-24.

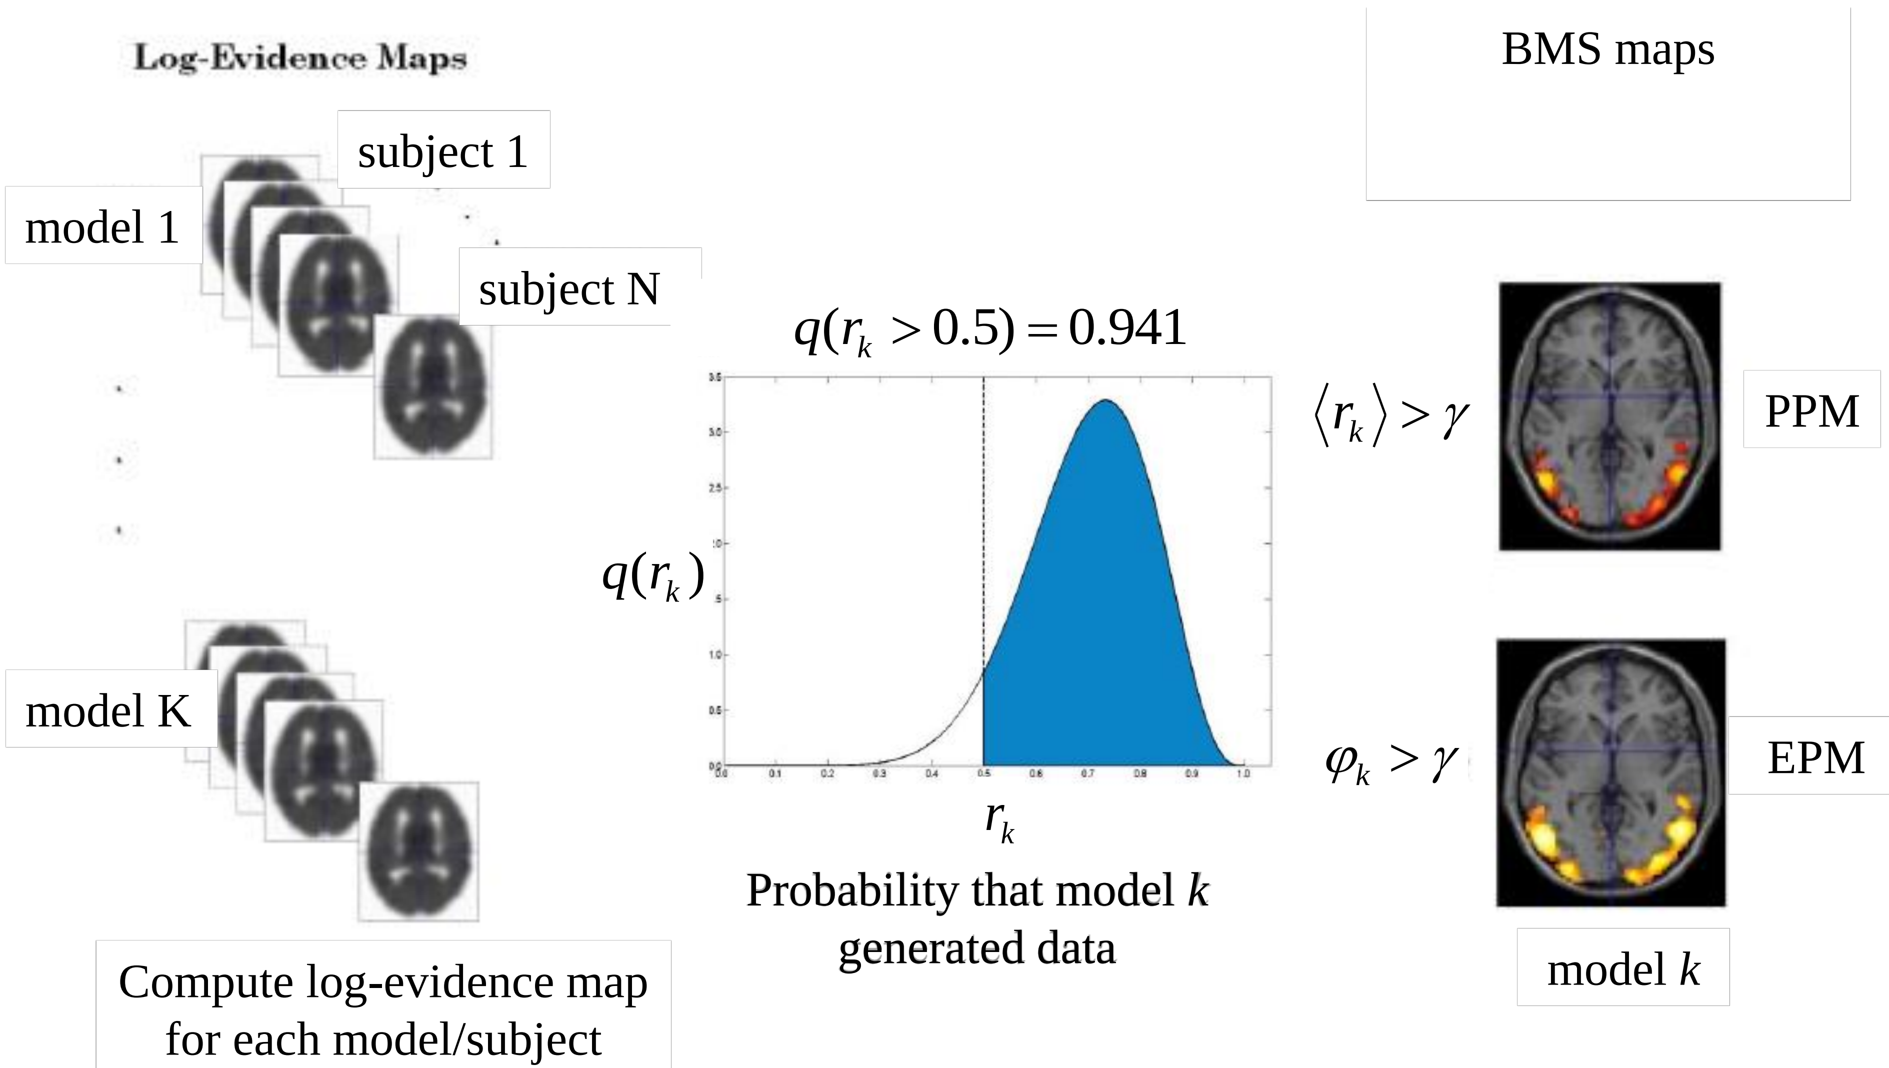

Exceedance Probability Maps

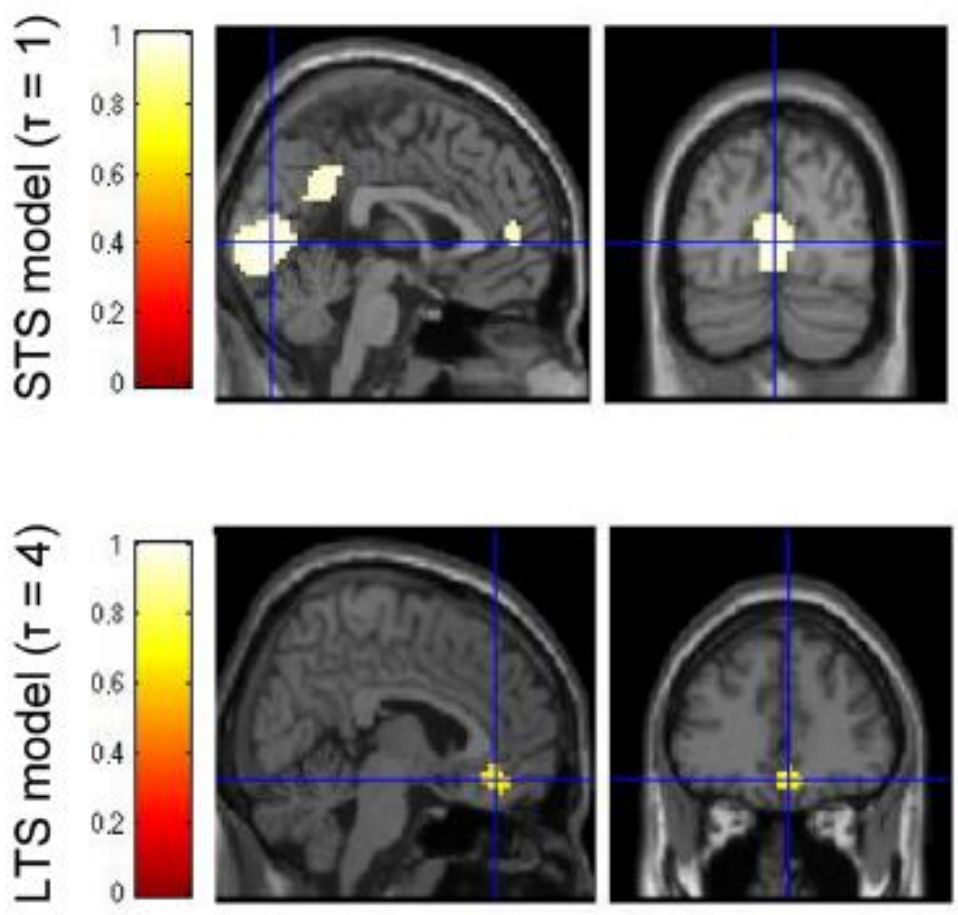

L. Harrison, S Bestmann, M. Rosa, W. Penny and G. Green (2011). **Time scales of representation in the human brain: weighing past information to predict future events.** *Frontiers in Human Neuroscience*, 5, 00037.

Overview

- Posterior Probability Maps
- Hemodynamic Response Functions
- Population Receptive Fields
- Computational fMRI
- **Multivariate Decoding**
- Dynamic Causal Modelling

K Friston et al. (2008) **Bayesian decoding of brain images**. *Neuroimage*, 39:181-205.

Relate Behavioural Descriptors, X, to fMRI data Y via voxel weights β

$$X = Y\beta$$

As the number of voxels in a region will likely exceed the number of time points in the fMRI time series, and only some combination of them will be useful for prediction we need to select 'features'

$$\beta = U\eta$$

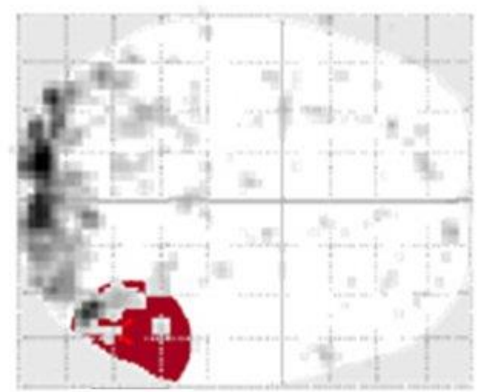

Which type of feature will be useful for decoding (1) voxels, (2) clusters, (3) singular vectors, (4) support vectors ?

Multivariate Decoding

Empirical (PEB) priors on voxel-weights

$$WX = RY\beta + \varsigma$$

$$\mathrm{cov}(\beta) = U\Sigma^{\eta}U^{T}$$

empirical priors

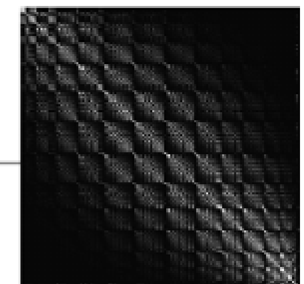

1.Voxels

2.Clusters

Null: $U = \varnothing$

Spatial vectors: $U = I$

Smooth vectors: $U(\vec{x}_i, \vec{x}_j) = \exp(-\tfrac{1}{2}(\vec{x}_i - \vec{x}_j)^2 \sigma^{-2})$

Singular vectors: $UDV^{T} = RY^{T}$

Support vectors: $U = RY^{T}$

Which type of feature will be useful for decoding (1) voxels, (2) clusters, (3) singular vectors, (4) support vectors ?

A Morcom and K Friston (2012) **Decoding episodic memory in ageing: A Bayesian Analysis of activity patterns predicting memory**. *Neuroimage* 59, 1772-1782.

Clustered Clusters

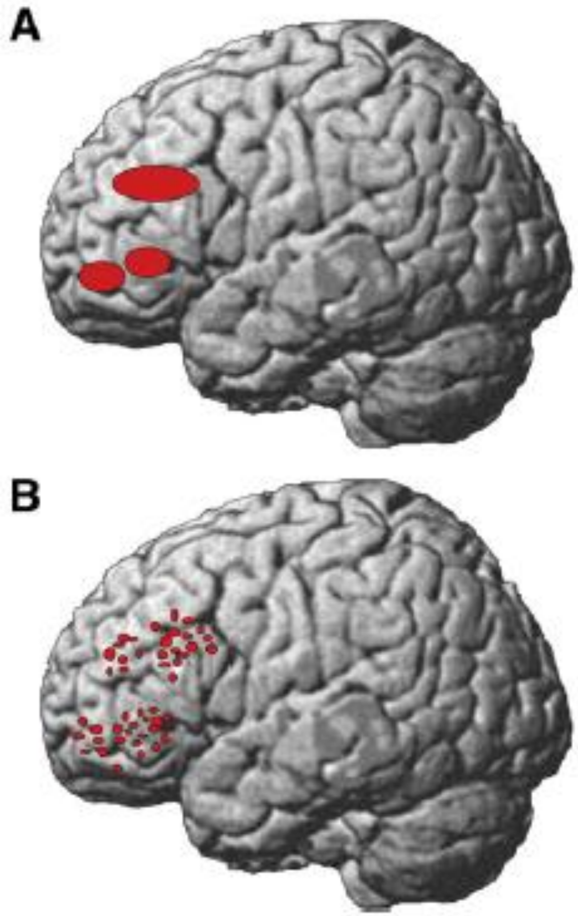

Distributed Voxels

A Morcom and K Friston (2012) **Decoding episodic memory in ageing: A Bayesian analysis of activity patterns predicting memory**. *Neuroimage* 59, 1772-1782.

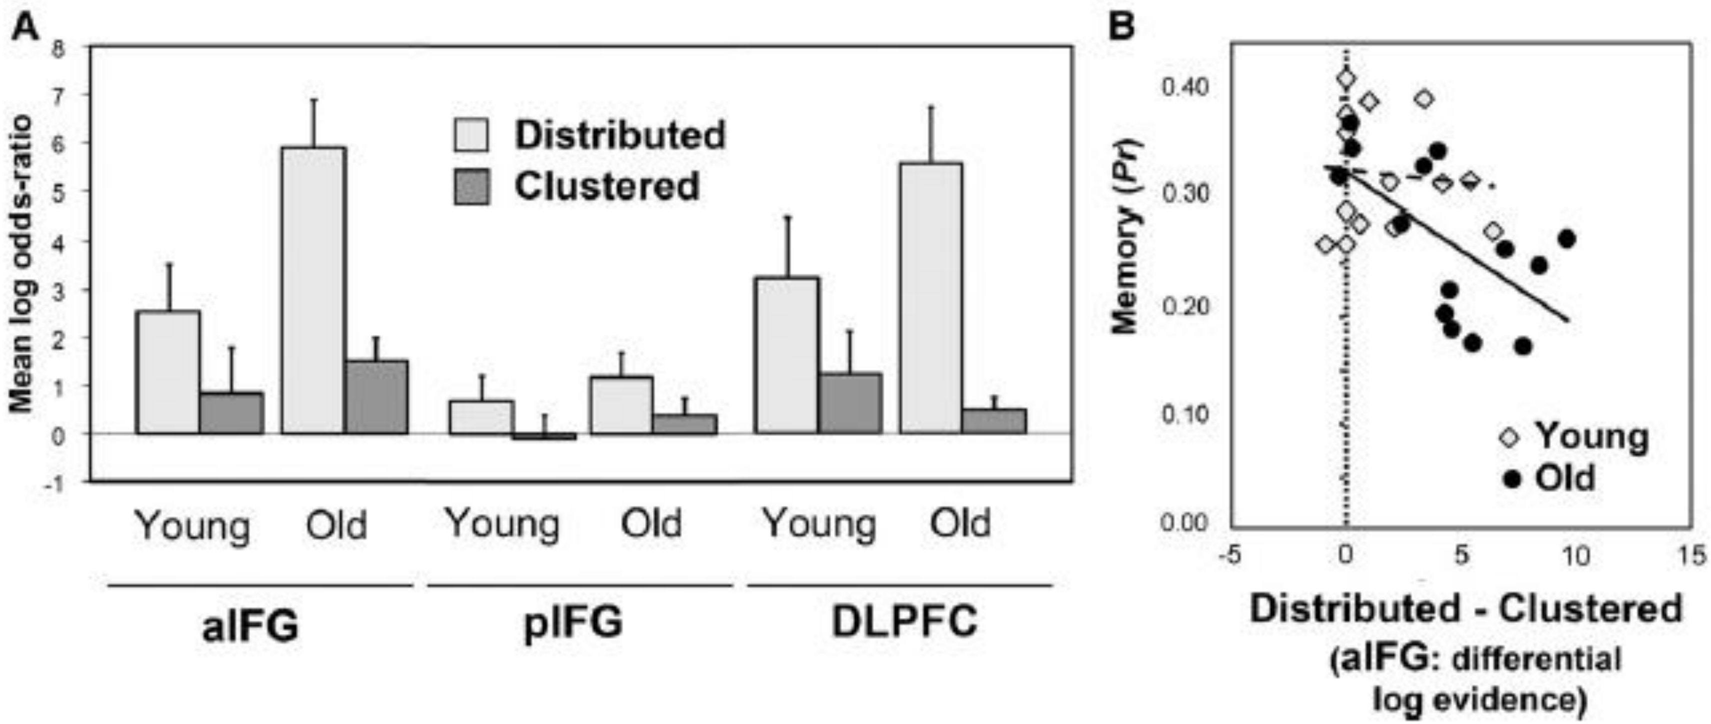

The more clustered the representation the better the memory

Multivariate Decoding

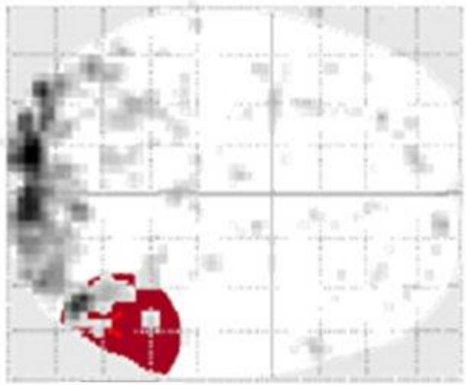

Q. With what sort of neural code is motion represented with in V5 ?

A. Voxels

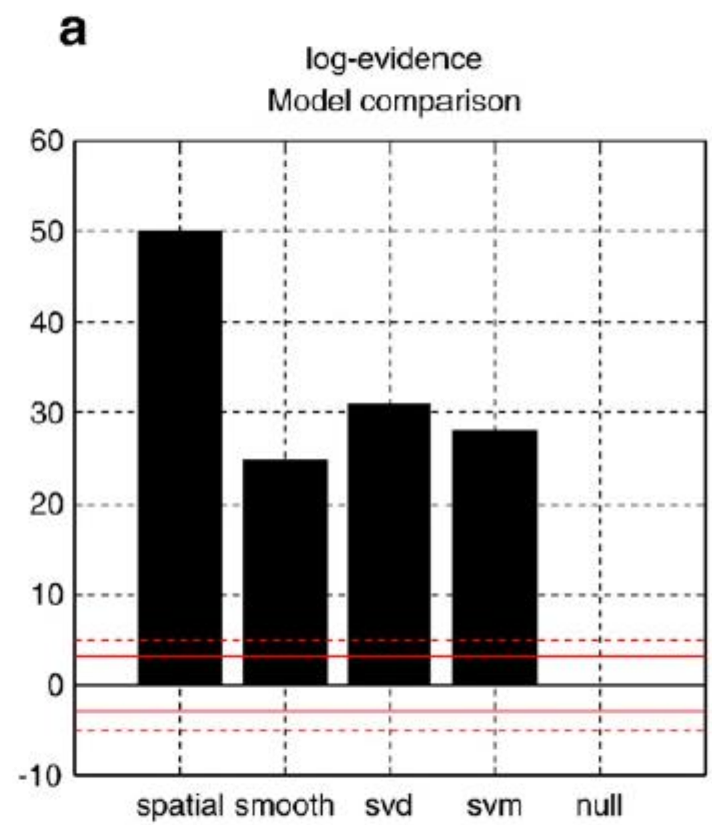

Voxels Clusters

Multivariate Decoding

Q. Which brain region can motion best be decoded from: V5 or PFC ?

A. V5.

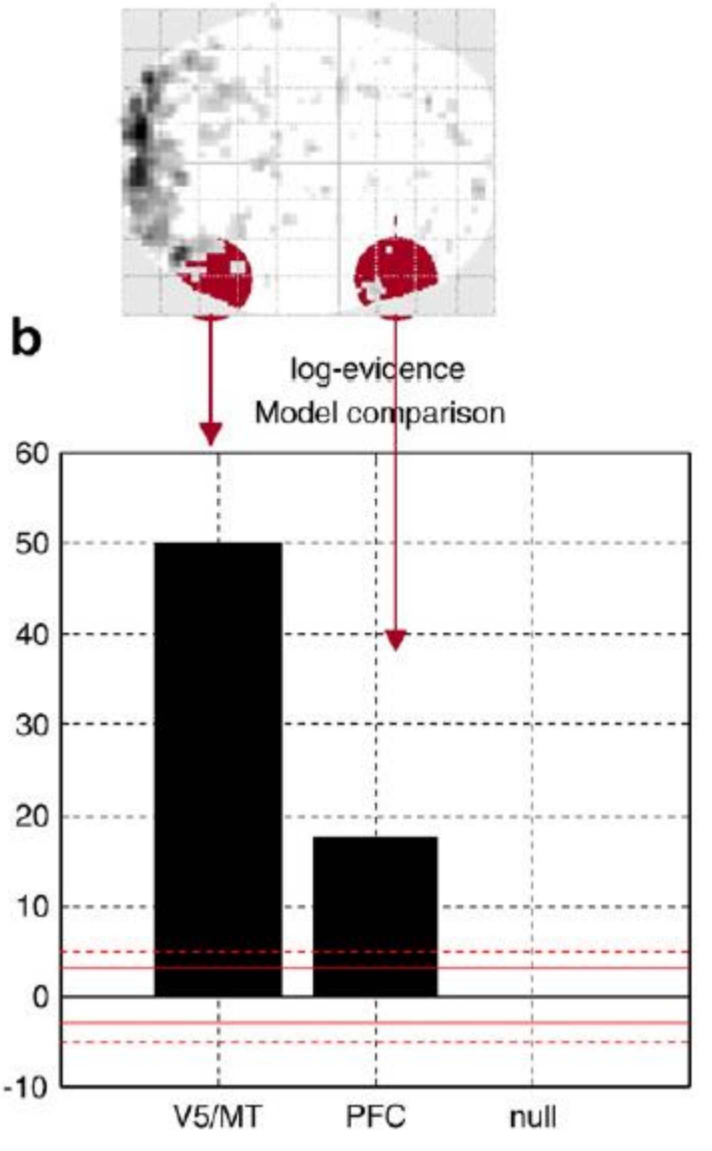

Multivariate Decoding

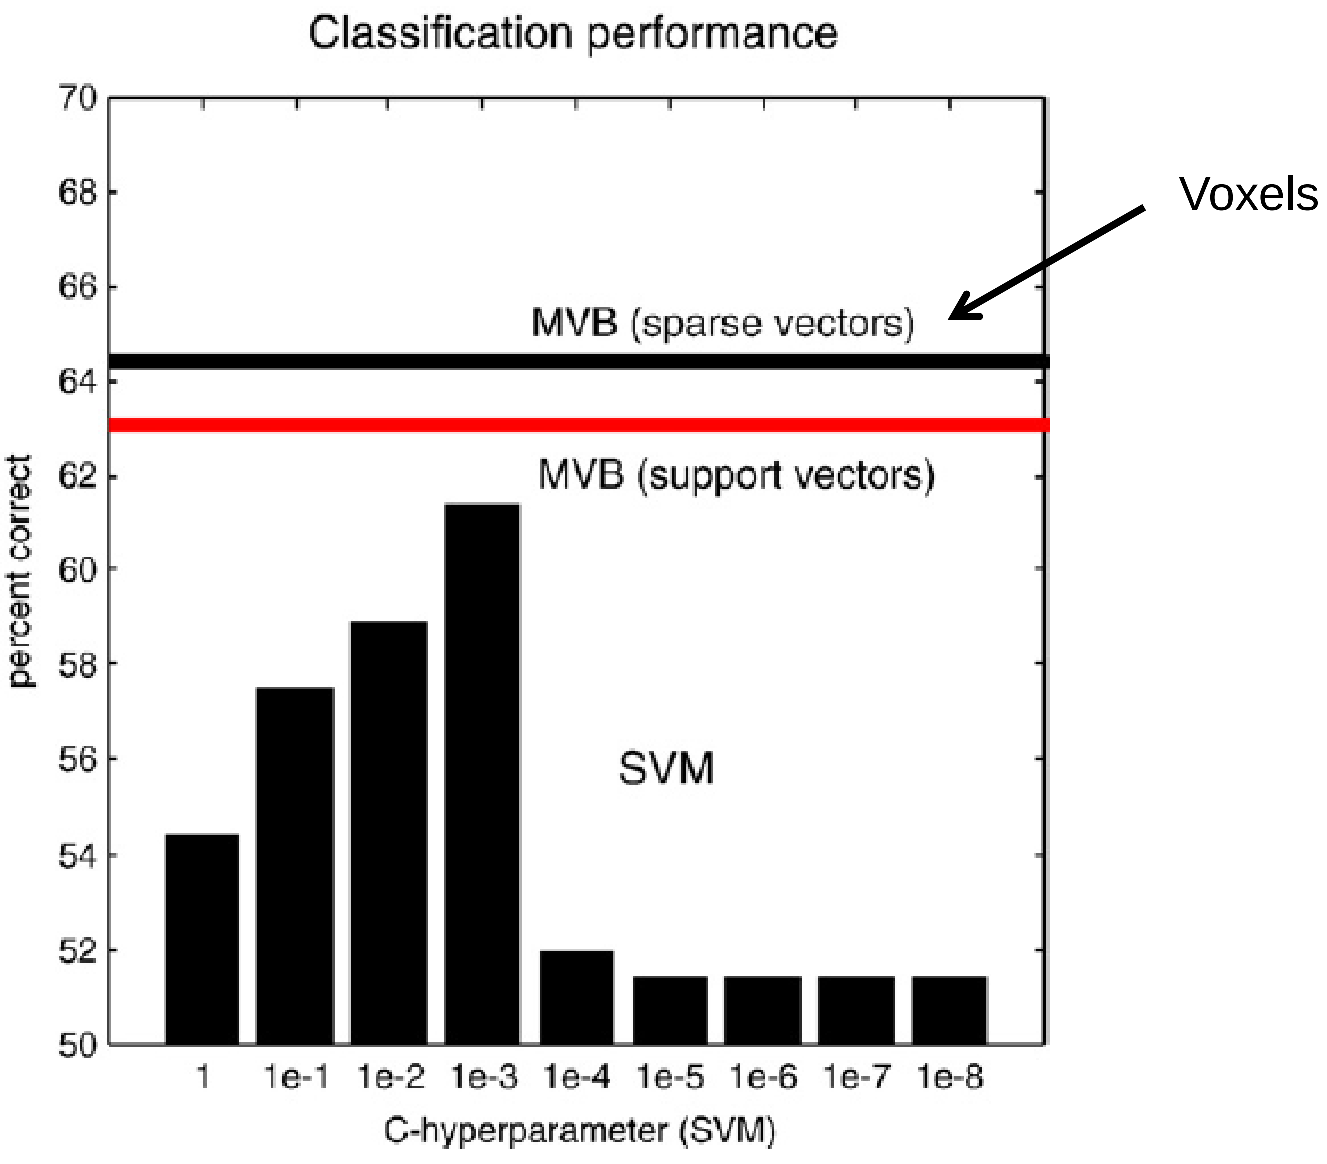

A Maas et al (2014) **Laminar activity in the hippocampus and entorhinal cortex related to novelty and episodic encoding** *Nature Communications,* 5:5547.

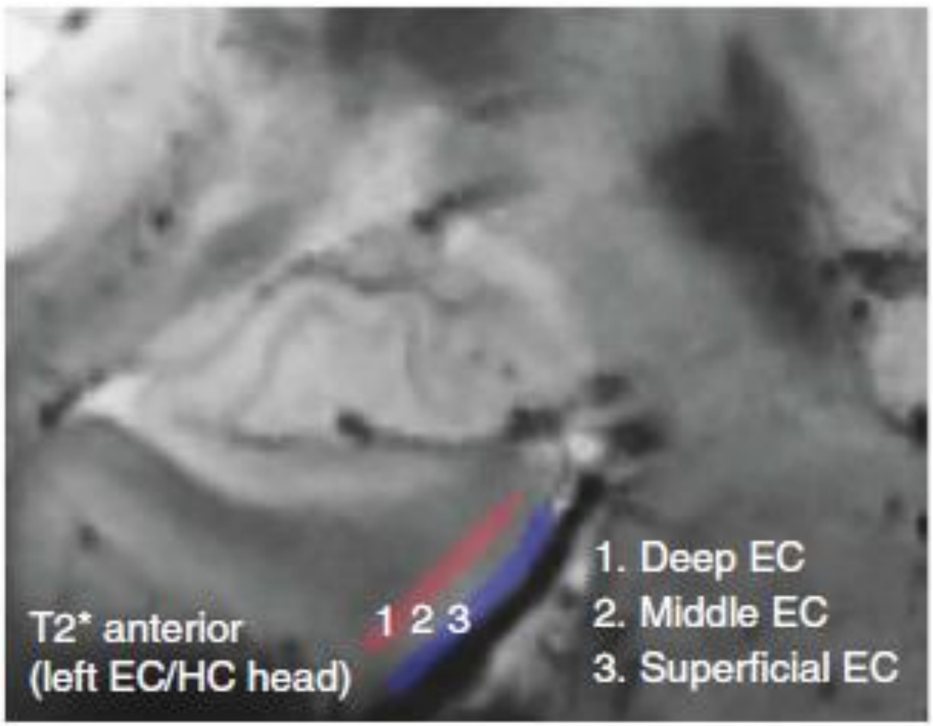

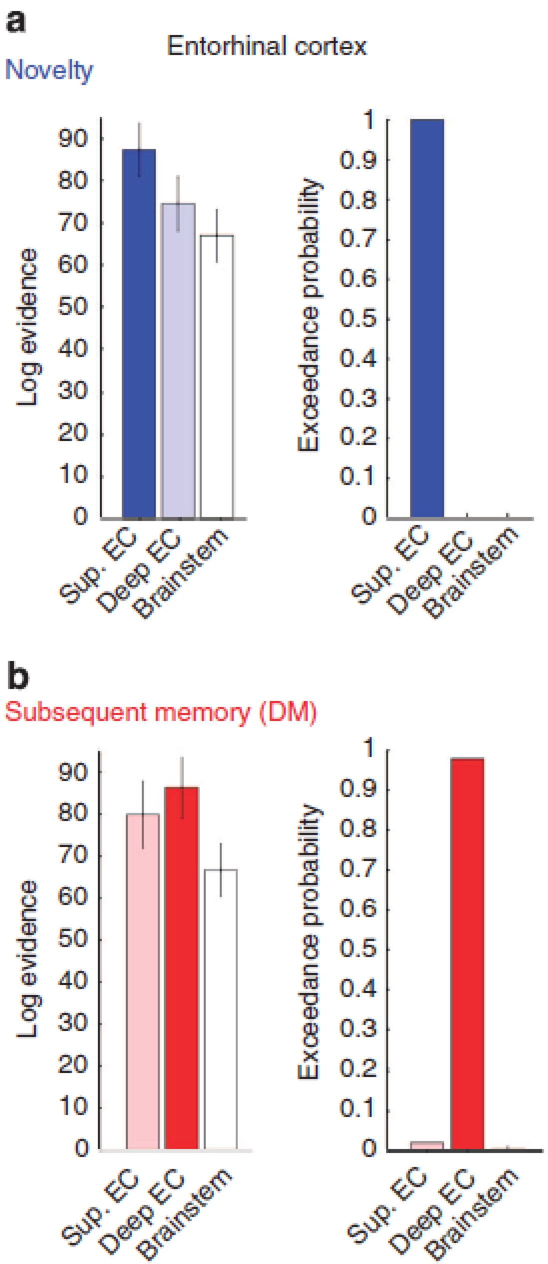

A Maas et al (2014) **Laminar activity in the hippocampus and entorhinal cortex related to novelty and episodic encoding** *Nature Communications,* 5:5547.

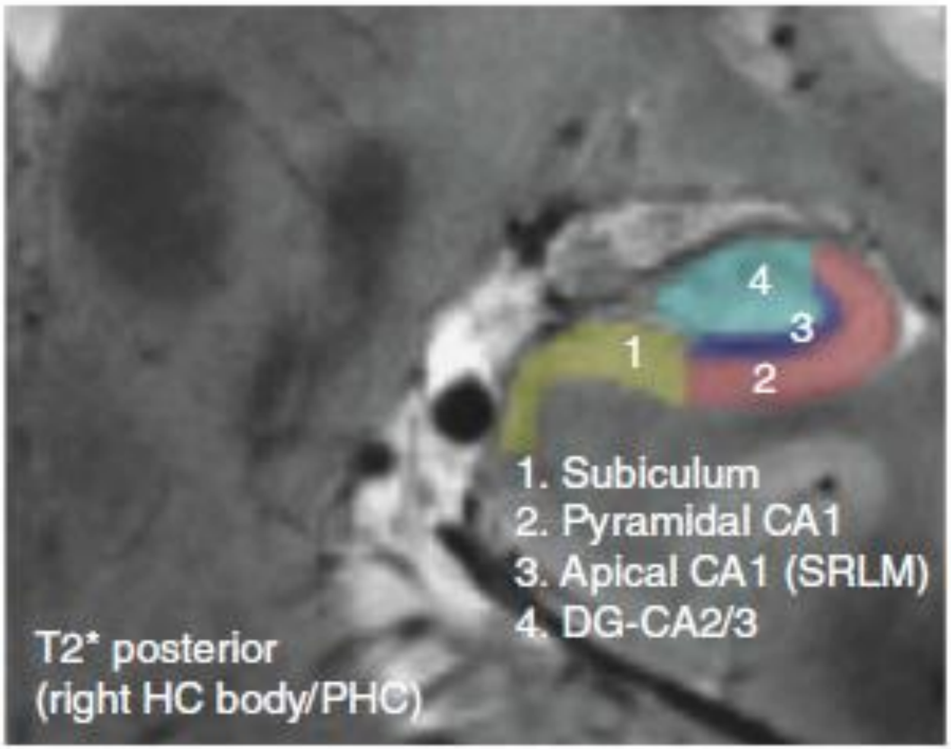

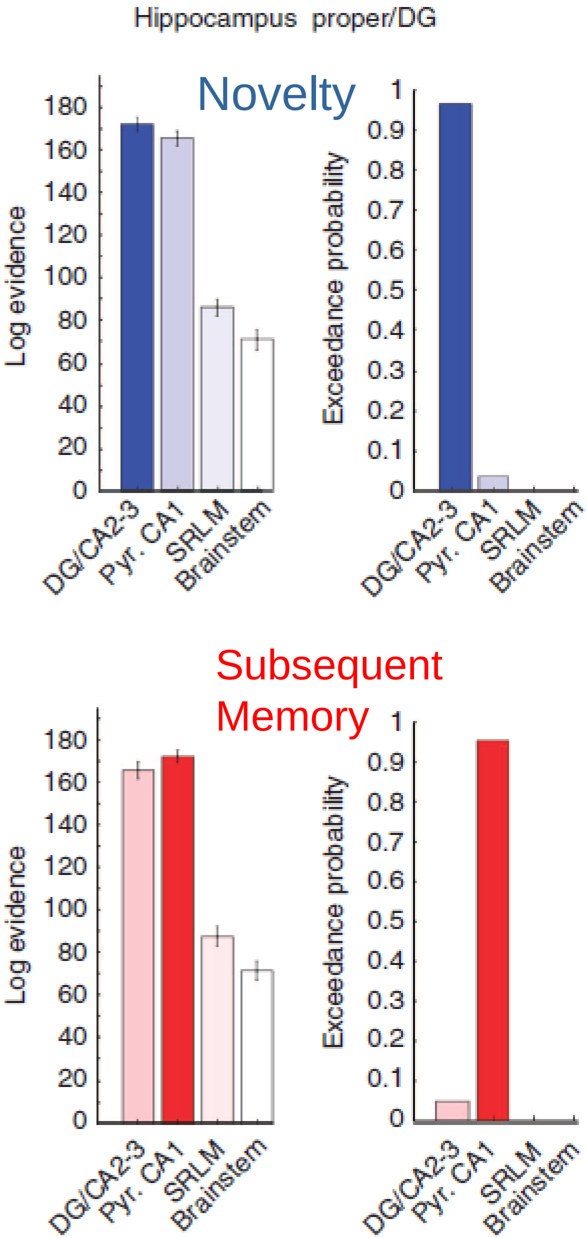

Overview

- Posterior Probability Maps
- Hemodynamic Response Functions
- Population Receptive Fields
- Computational fMRI
- Multivariate Decoding
- **Dynamic Causal Modelling**

Single region

$$\dot{z}_1 = a_{11} z_1 + c u_1$$

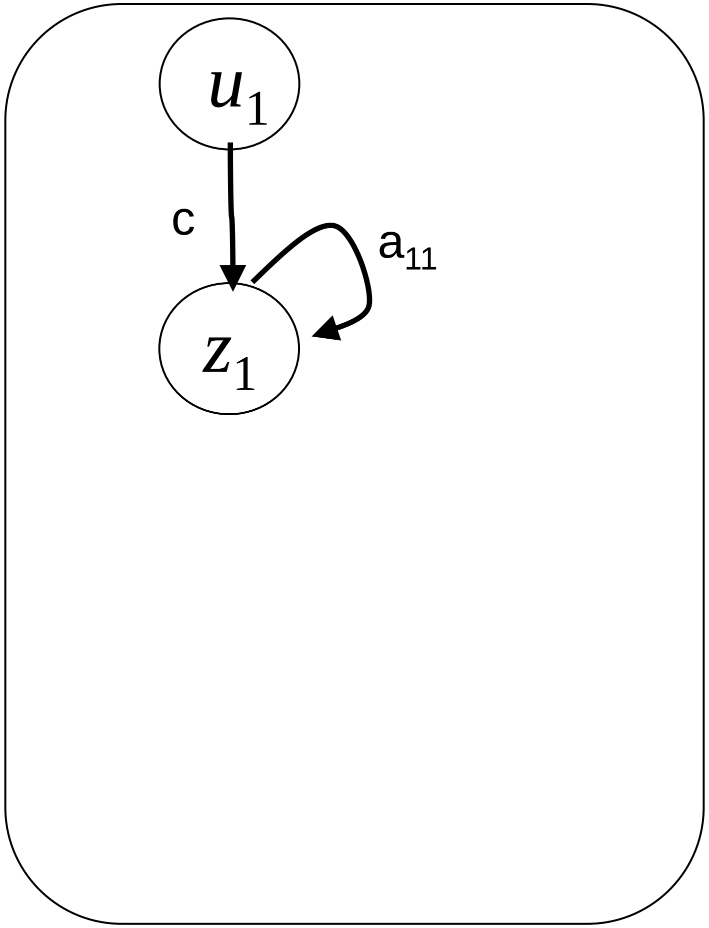

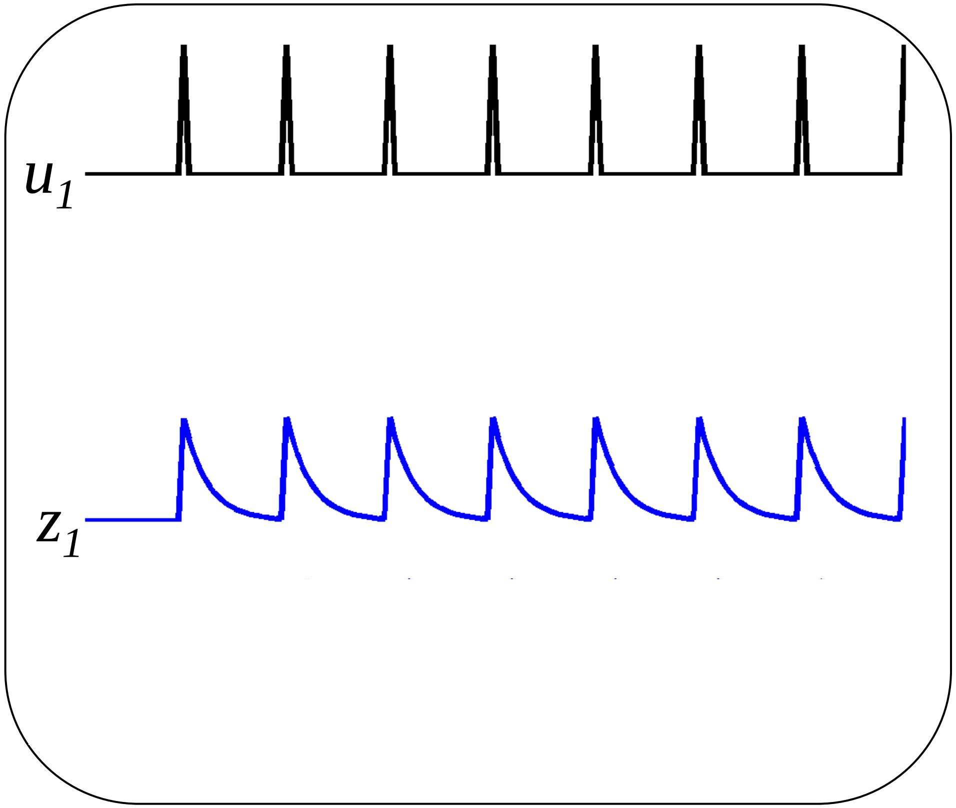

Multiple regions

$$\begin{bmatrix} \dot{z}_1 \\ \dot{z}_2 \end{bmatrix} = \begin{bmatrix} a_{11} & 0 \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + \begin{bmatrix} c \\ 0 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$$

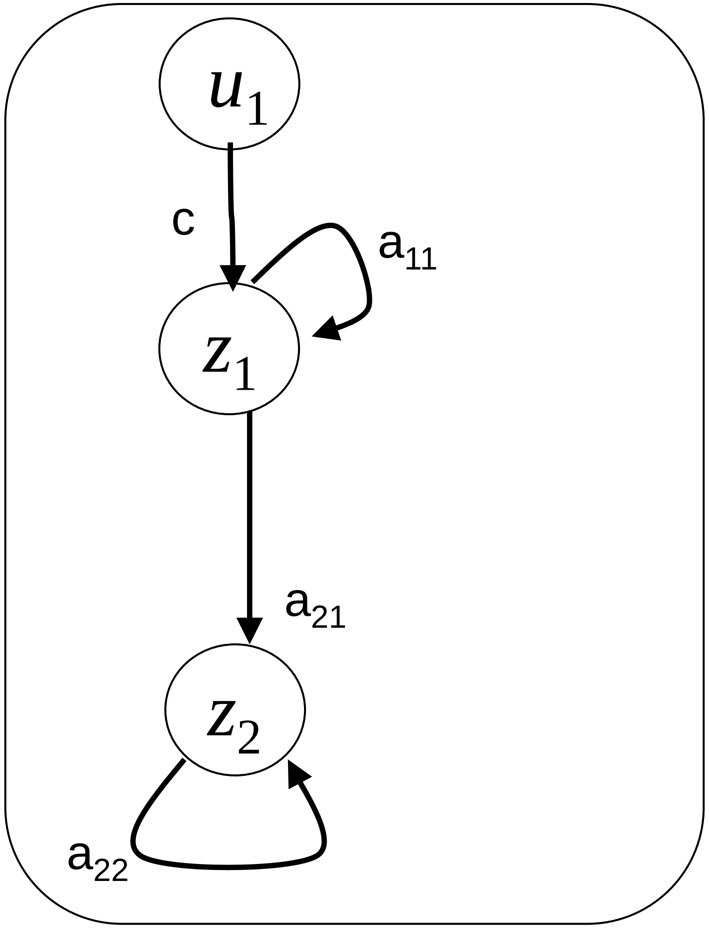

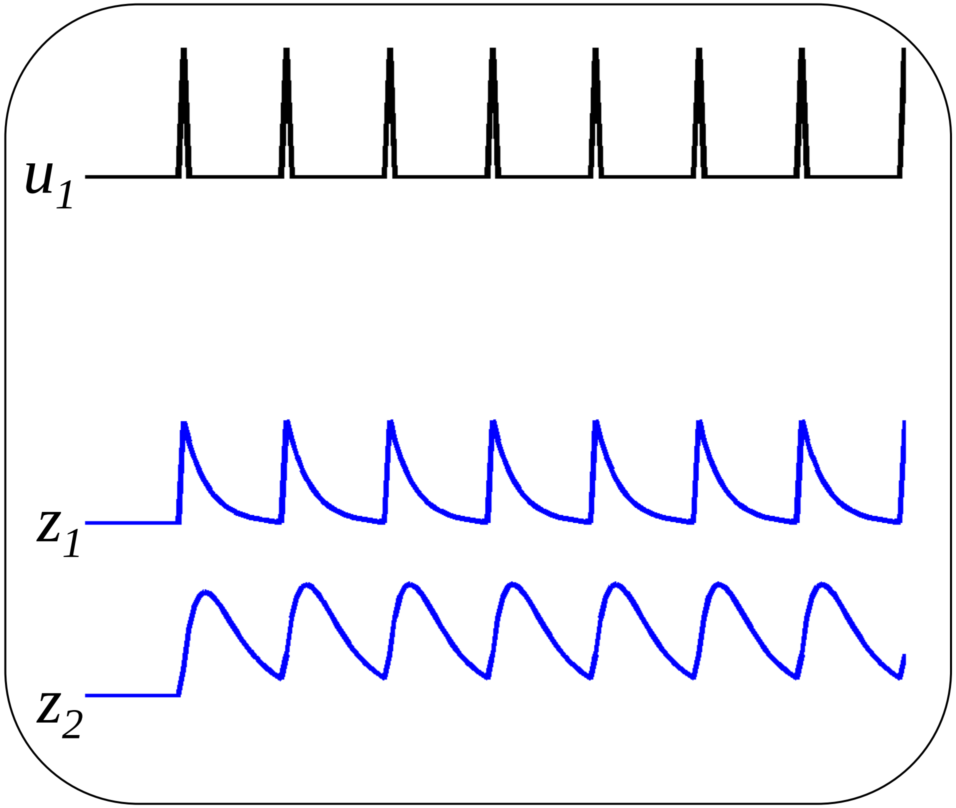

Modulatory inputs

$$\begin{bmatrix} \dot{z}_1 \\ \dot{z}_2 \end{bmatrix} = \begin{bmatrix} a_{11} & 0 \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + u_2 \begin{bmatrix} 0 & 0 \\ b_{21} & 0 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + \begin{bmatrix} c \\ 0 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$$

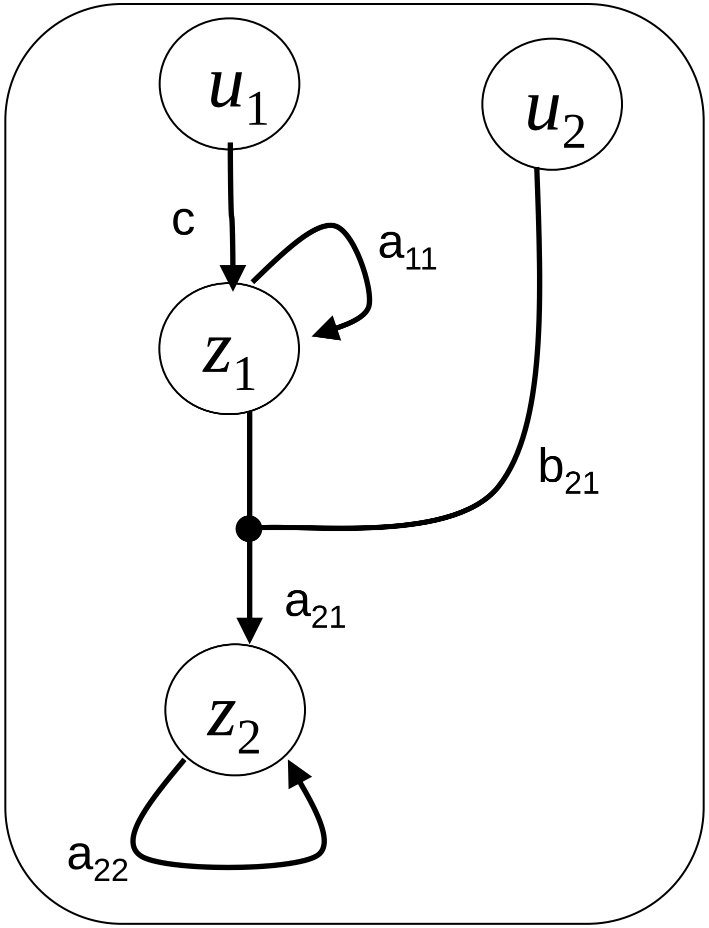

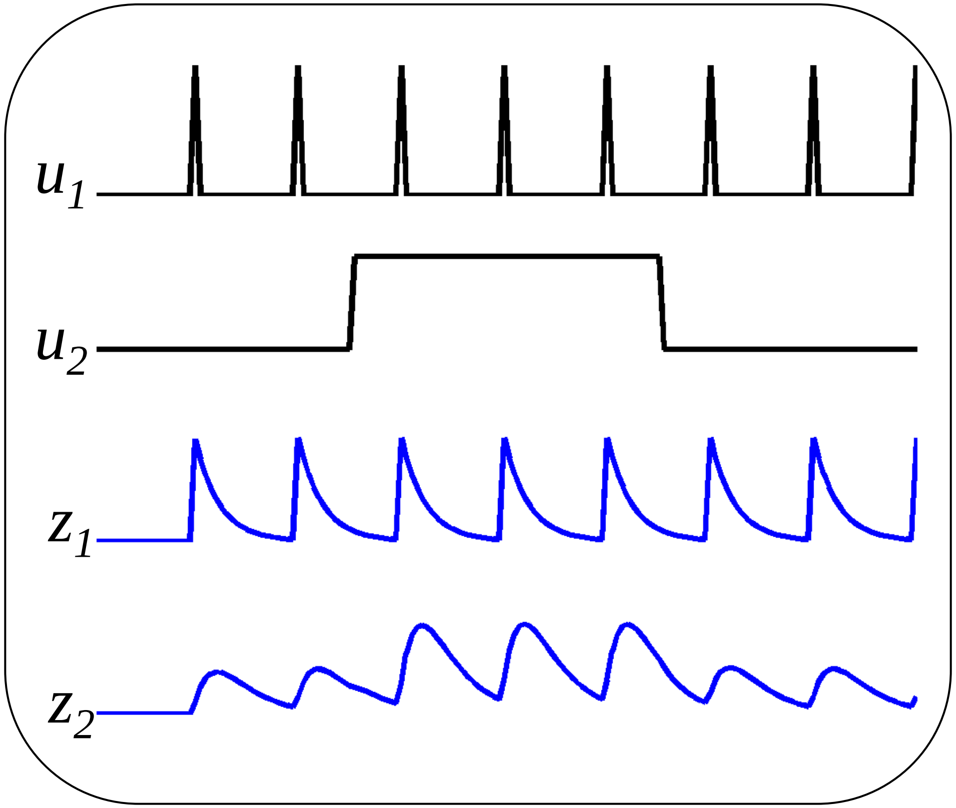

Reciprocal connections

$$\begin{bmatrix} \dot{z}_1 \\ \dot{z}_2 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + u_2 \begin{bmatrix} 0 & 0 \\ b_{21} & 0 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + \begin{bmatrix} c \\ 0 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$$

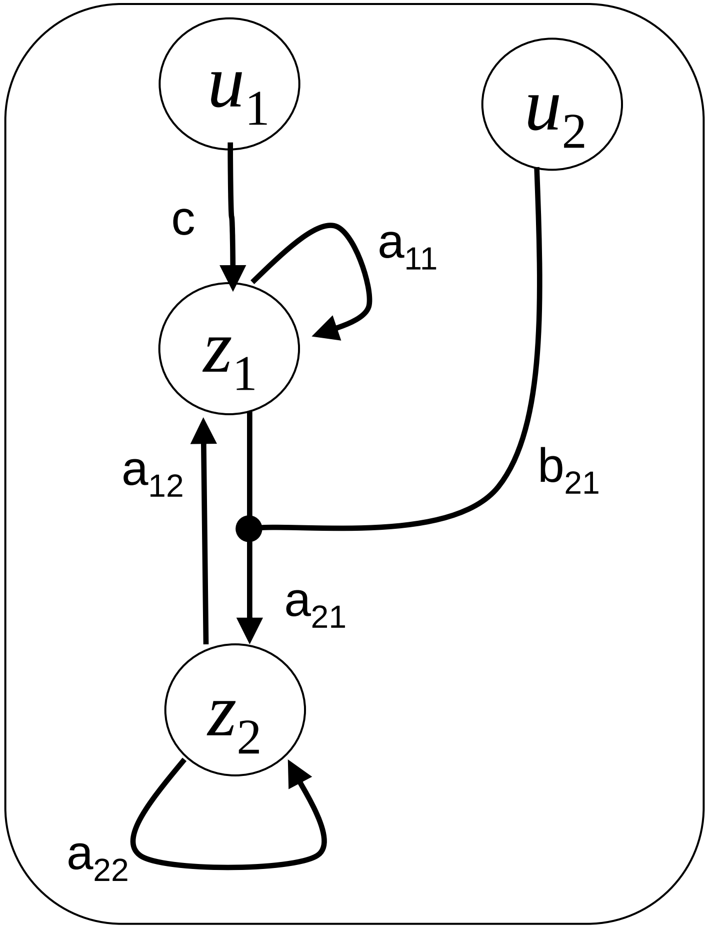

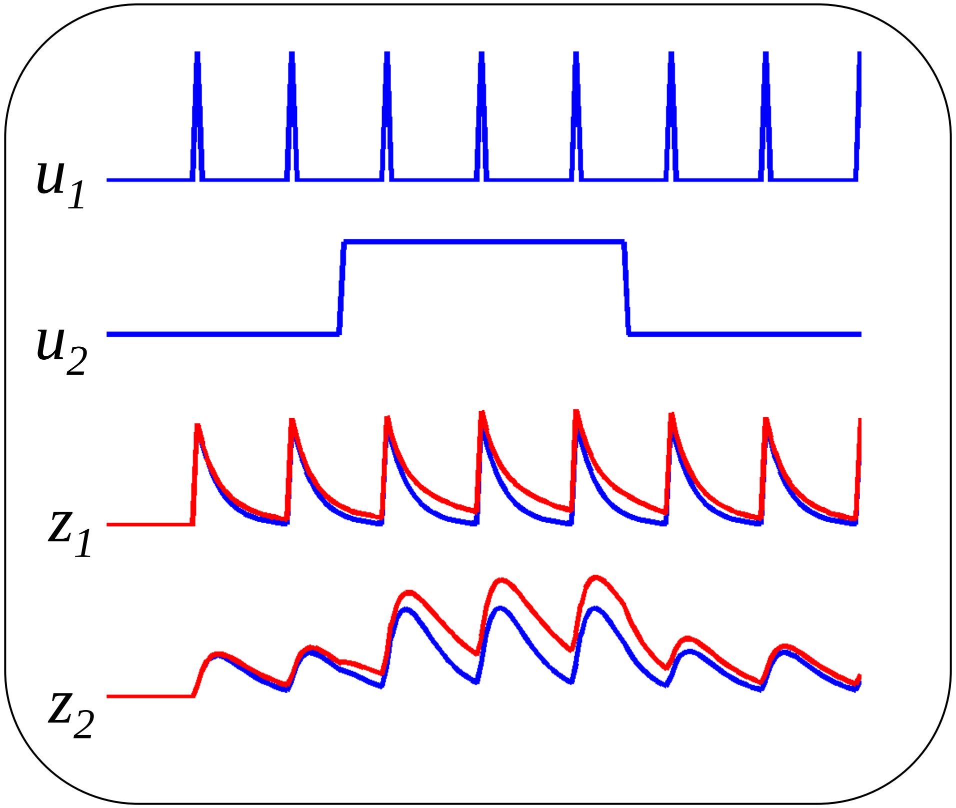

Neurodynamics

Neuronal Activity

Change in Neuronal Activity

Inputs

$$\dot{\mathbf{z}} = \mathbf{A}\mathbf{z} + \sum_i \mathbf{u}_i \mathbf{B}_i \mathbf{z} + \mathbf{C}\mathbf{u}$$

Input Connectivity Matrix

Intrinsic Connectivity Matrix

Modulatory Connectivity Matrices

Rowe et al. 2010, **Dynamic causal modelling of effective connectivity from fMRI: Are results reproducible and sensitive to Parkinson's disease and its treatment?** *NeuroImage, 52:1015-1026.*

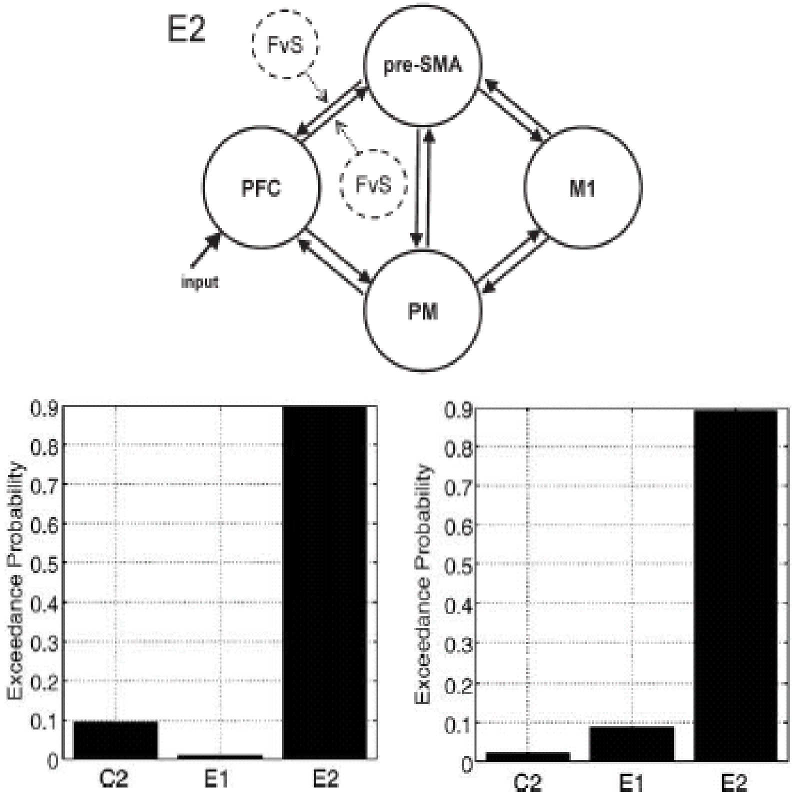

Age-matched controls

PD patients on medication

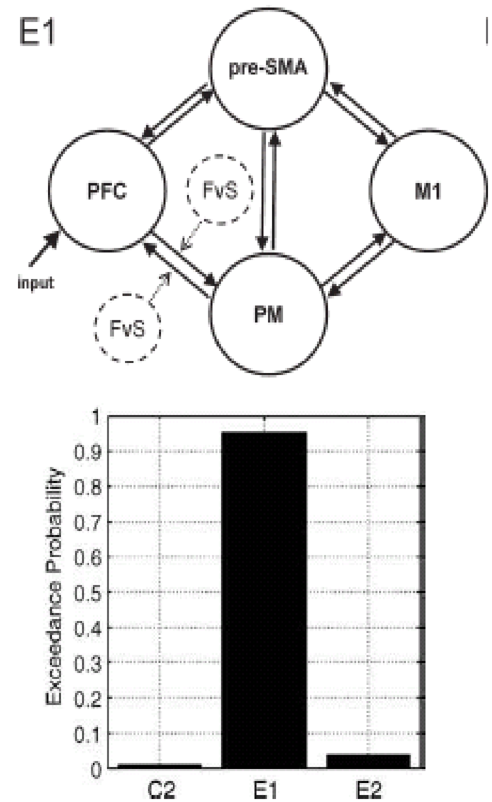

PD patients off medication

Selection of action modulates connections between PFC and SMA

DA-dependent functional disconnection of the SMA

Brodersen et al. 2011, **Generative Embedding for Model-Based Classification of fMRI data**. *PLoS Comput. Biol. 7(6):e1002079.*

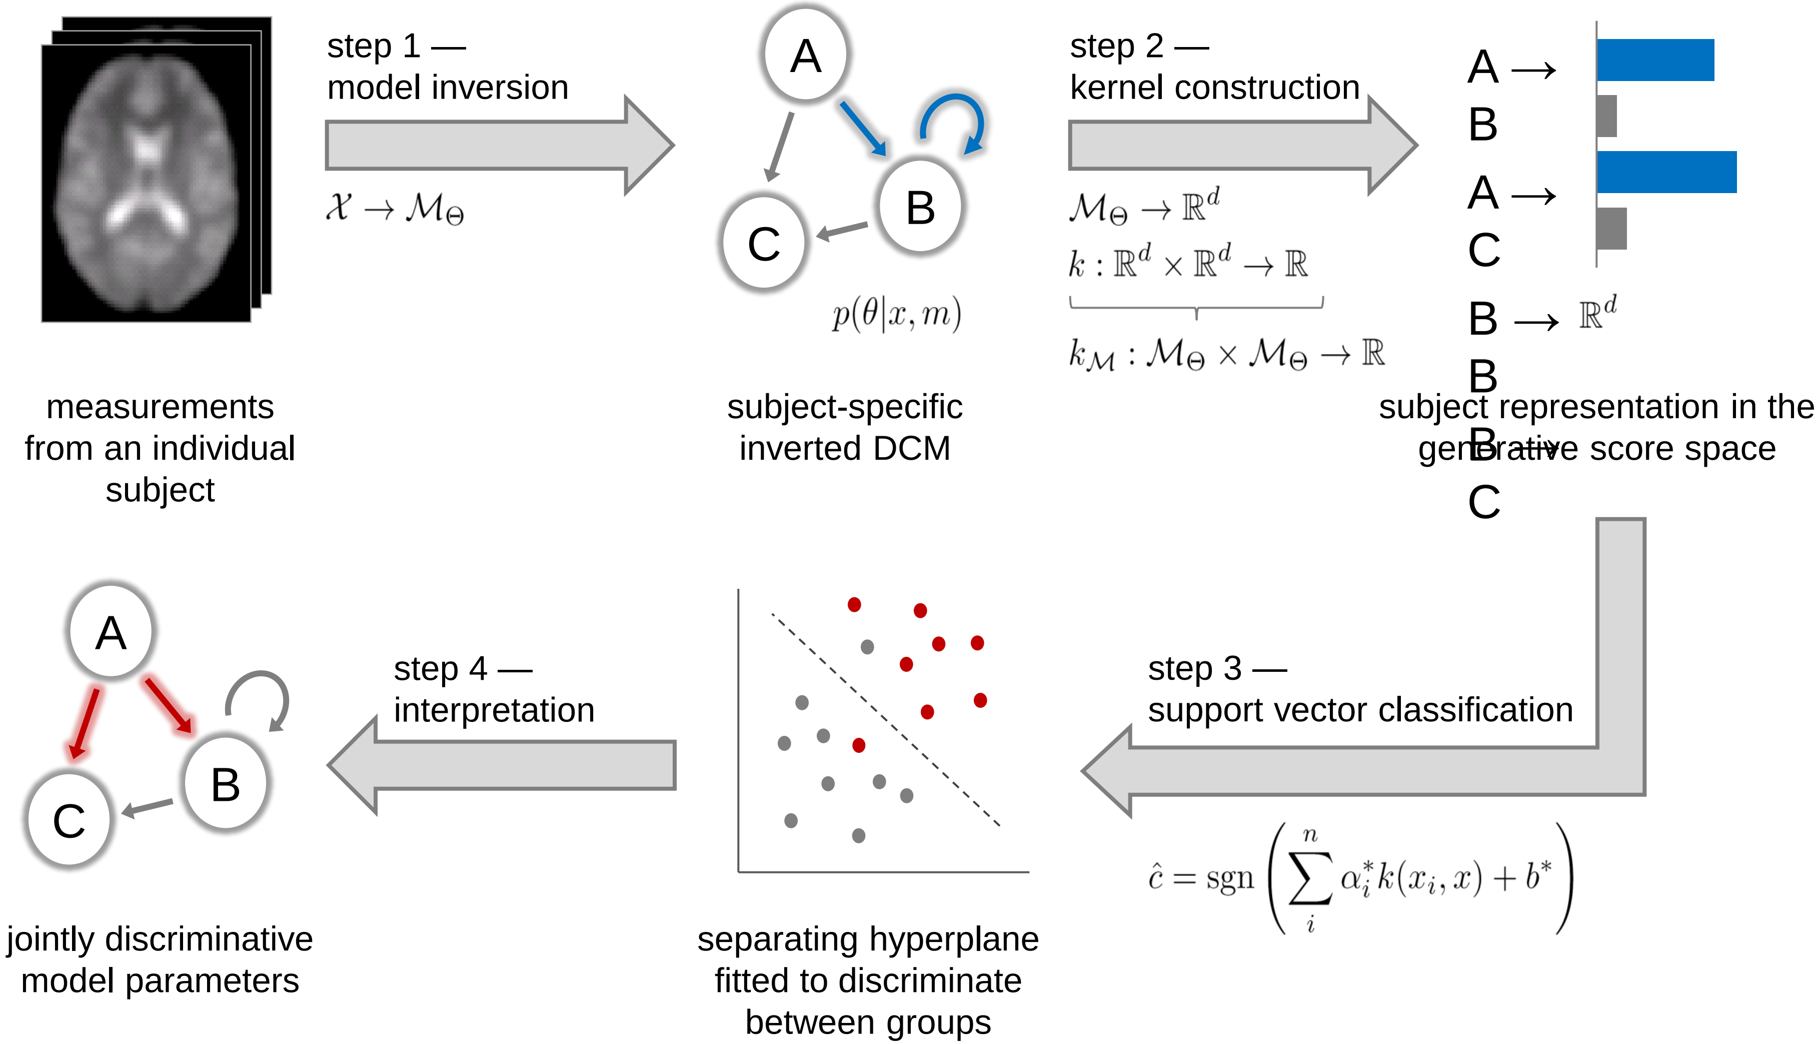

Model-based decoding of disease status: mildly aphasic patients (N=11) vs. controls (N=26)

Connectional fingerprints from a 6-region DCM of auditory areas during speech perception

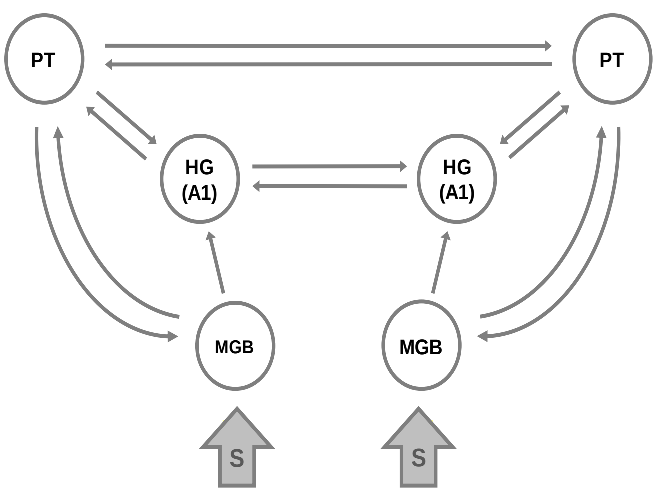

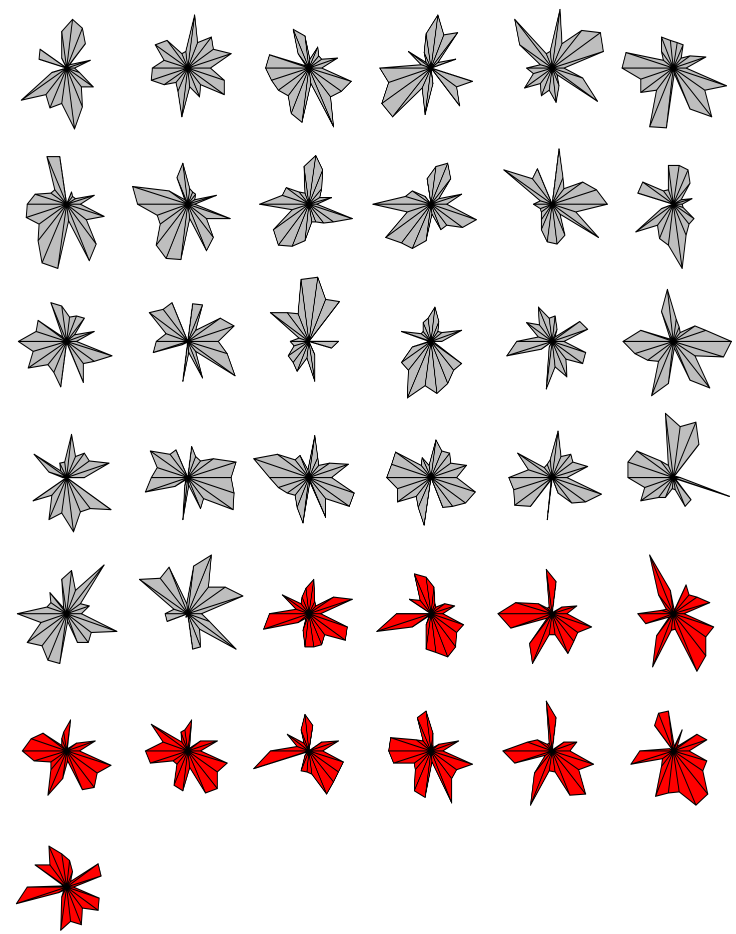

Model-based decoding of disease status: aphasic patients (N=11) vs. controls (N=26)

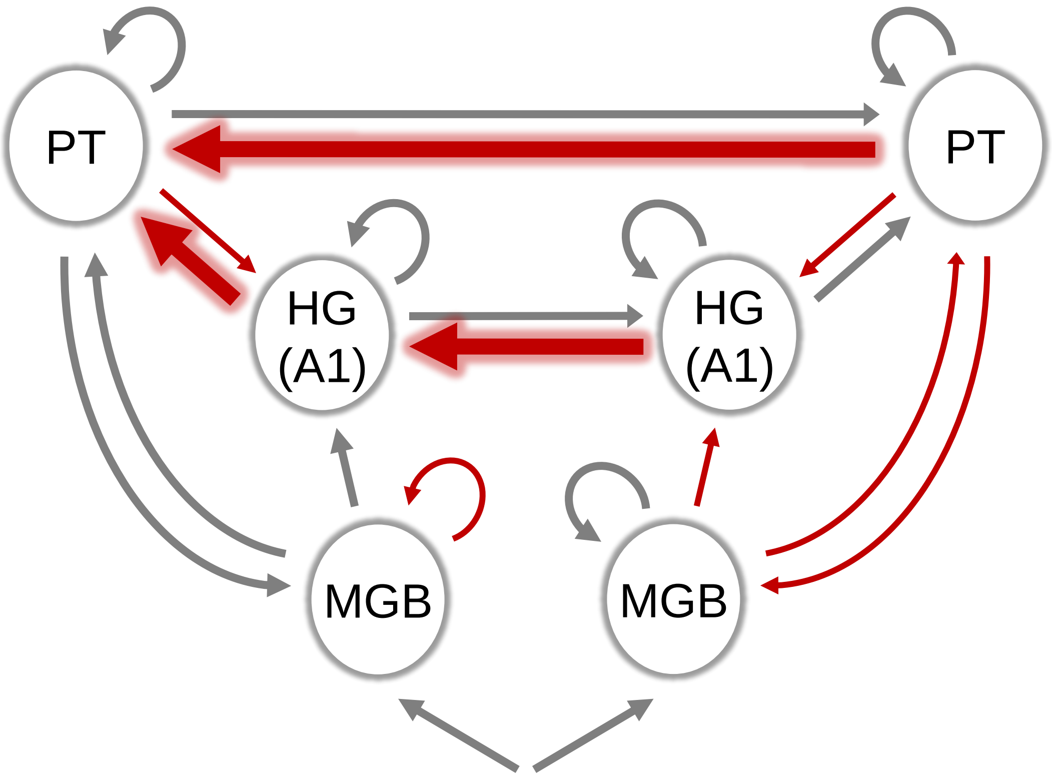

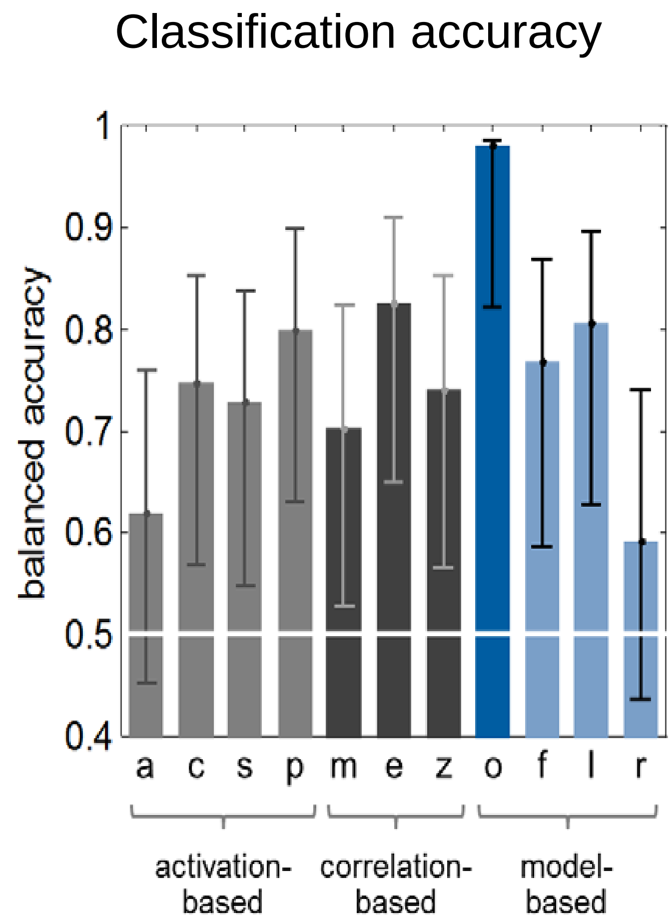

Multivariate searchlight classification analysis

Generative embedding using DCM

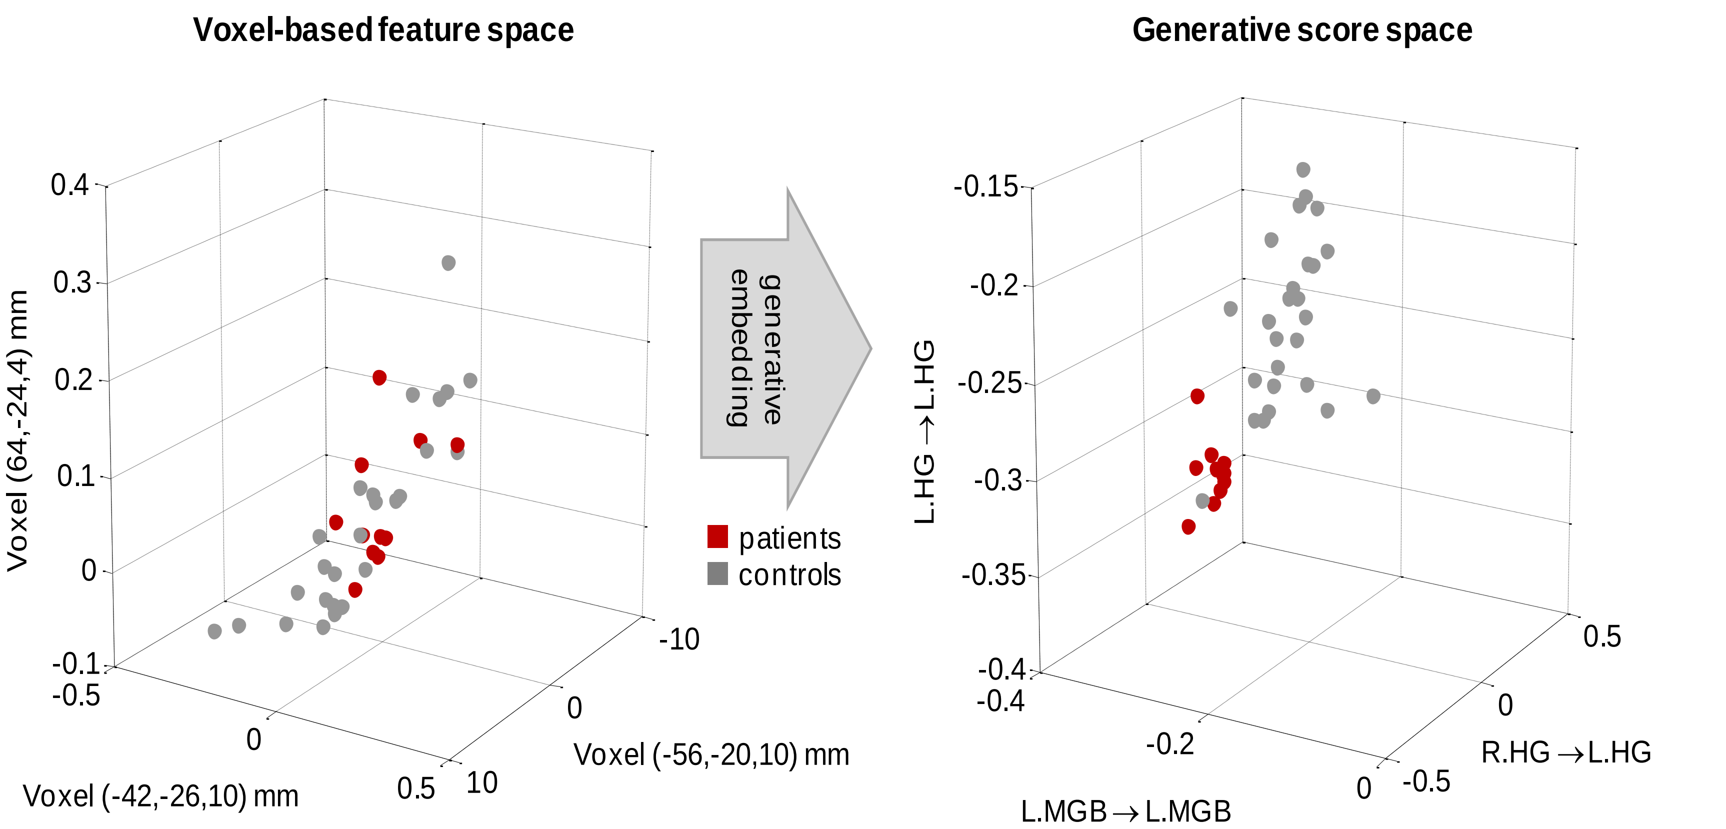

Summary

- Posterior Probability Maps
- Hemodynamic Response Functions
- Population Receptive Fields
- Computational fMRI
- Multivariate Bayes
- Dynamic Causal Modelling

Savage-Dickey Ratios

Bayesian equivalent of inference using F-tests implemented using Savage-Dickey approximations to the log Bayes Factor.

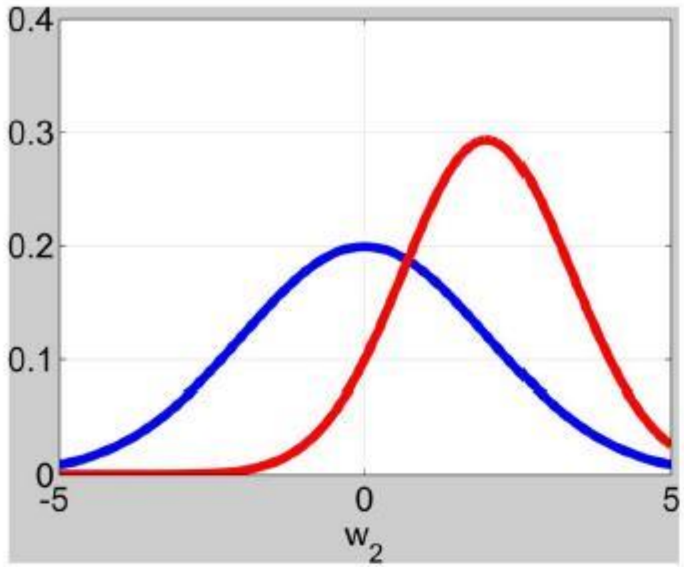

Figure 1. The figure shows the prior density $p(w_2|m_2)$ in blue and the posterior density $p(w_2|m_2,y)$ in red. Here $BF_{12}=0.5$, weakly favouring the more complex model m_2, since the parameter w_2 is half as likely to be zero after seeing the data than before.

W. Penny and G. Ridgway (2013). **Efficient Posterior Probability Mapping using Savage-Dickey Ratios.** *PLoS One* 8(3), e59655

Faces versus scrambled faces

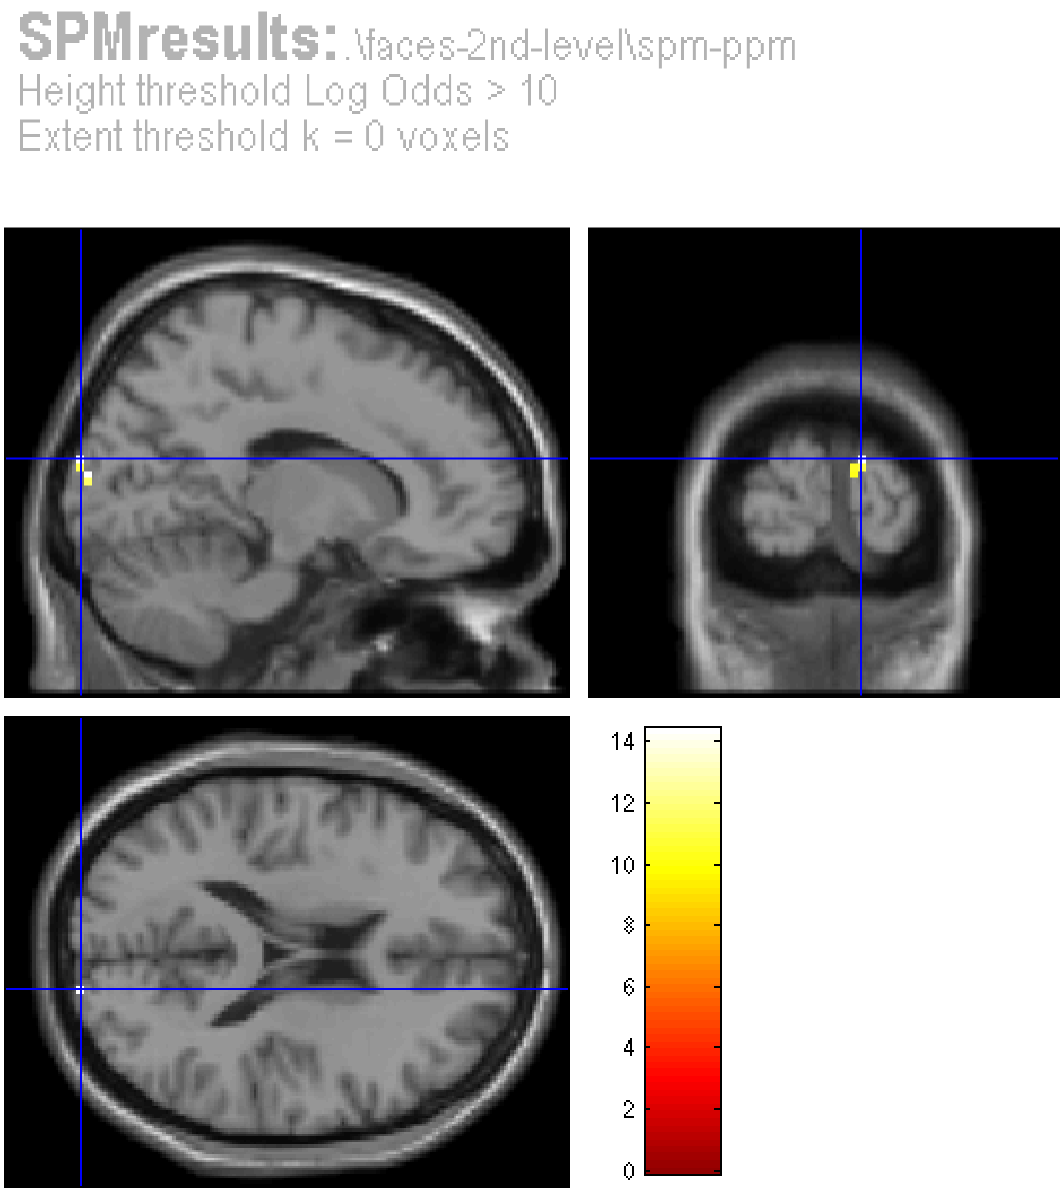

RFX analysis
on 18 subjects.

Data from
Rik Henson.

Faces versus scrambled faces: Evidence for Null

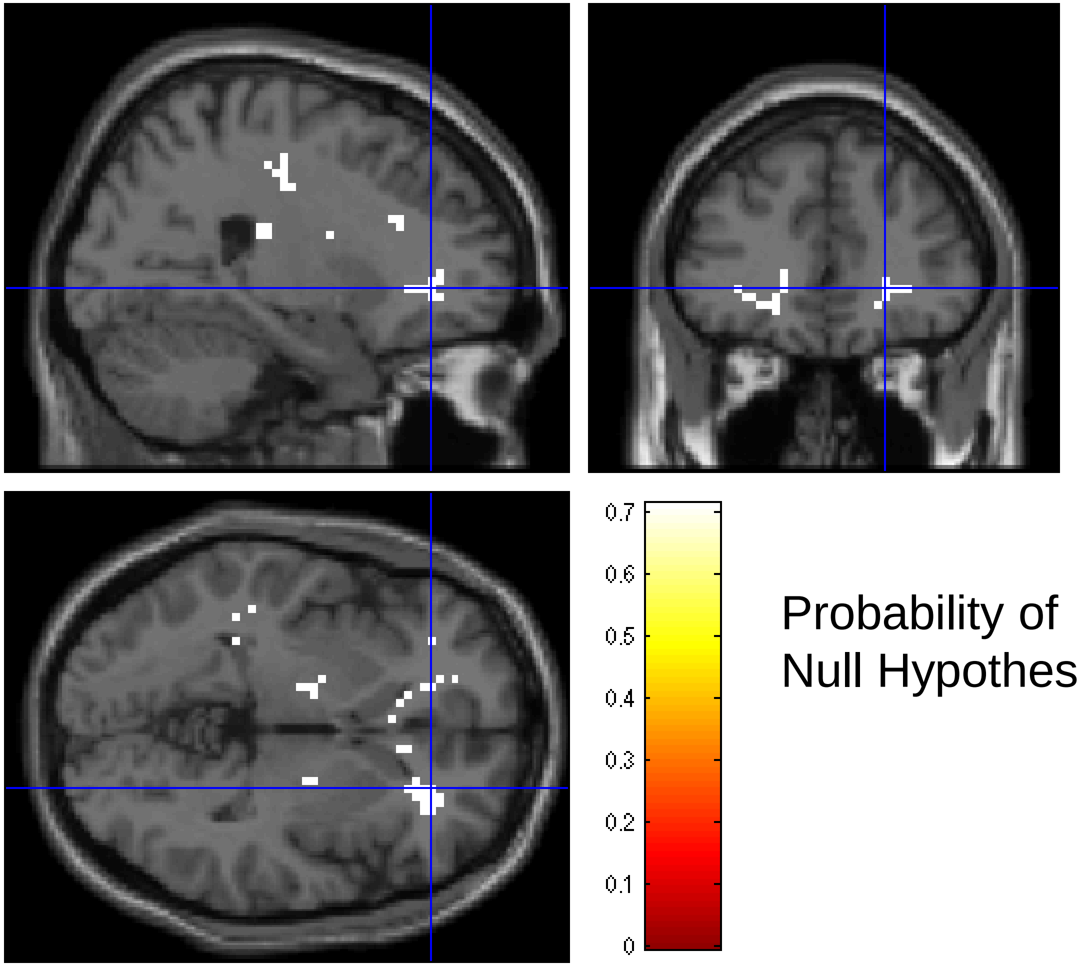

Probability of
Null Hypothesis

Using command line call to *spm_bms_test_null.m*

One parameter

Likelihood and Prior

$$p(y \mid \theta^{(1)}) = N(\theta^{(1)}, \lambda^{(1)})$$
$$p(\theta^{(1)}) = N(\theta^{(2)}, \lambda^{(2)})$$

Posterior

$$p(\theta^{(1)} \mid y) = N(m, P)$$
$$P = \lambda^{(1)} + \lambda^{(2)}$$
$$m = \frac{\lambda^{(1)}}{P}\theta^{(1)} + \frac{\lambda^{(2)}}{P}\theta^{(2)}$$

Relative Precision Weighting

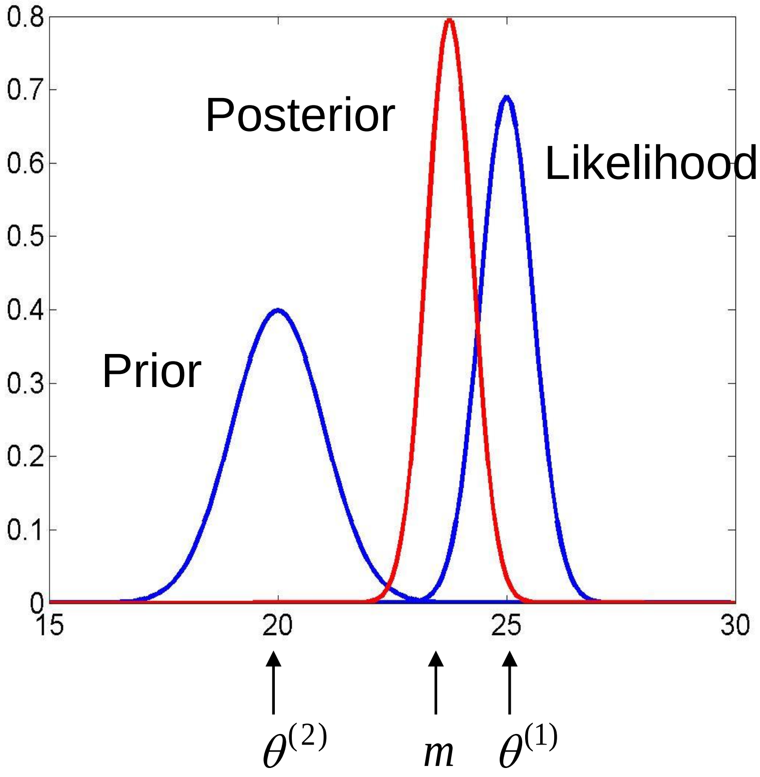

Two parameters

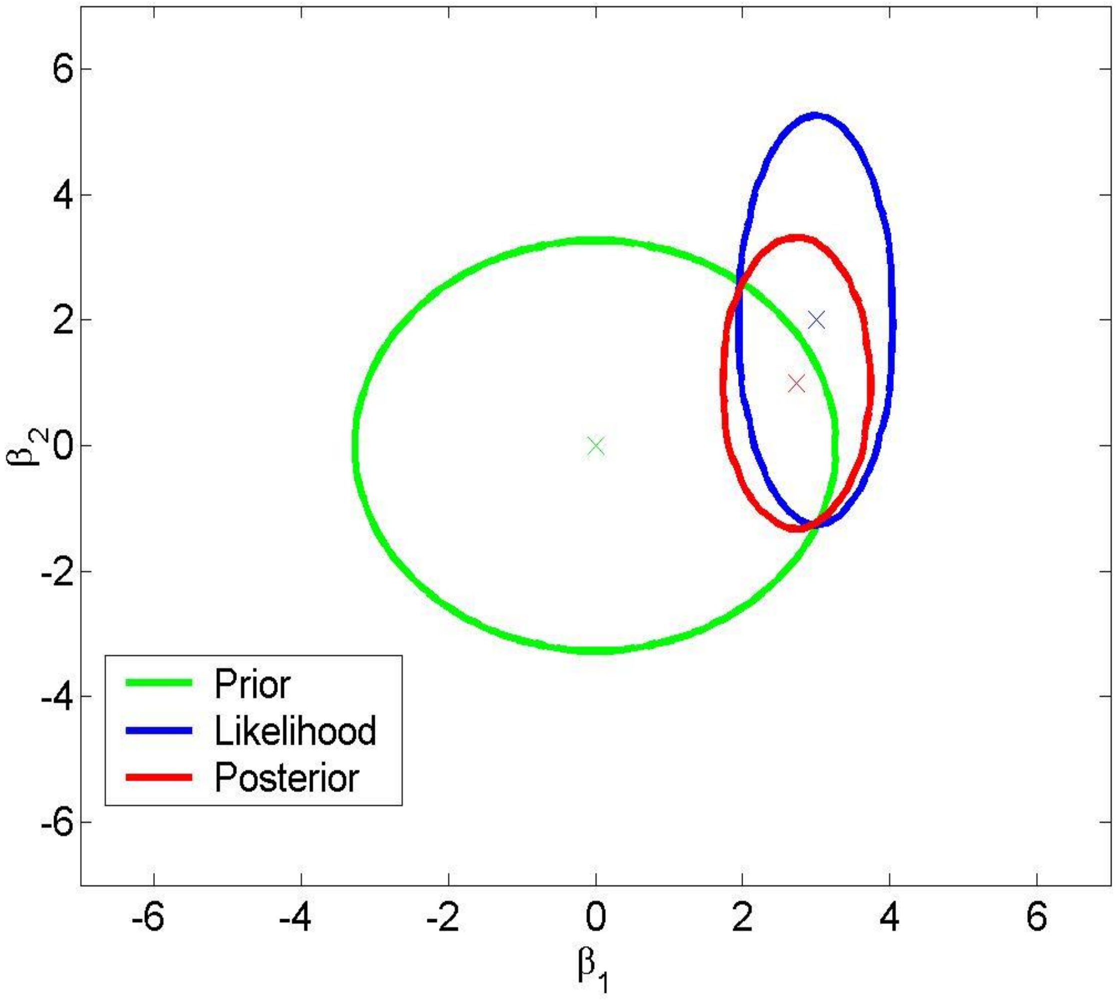

Bayes Rule for Gaussians

Likelihood: $p(y|\theta, m)$

Prior: $p(\theta|m)$

Bayes rule: $p(\theta|y, m) = \dfrac{p(y|\theta, m)\, p(\theta|m)}{p(y|m)}$

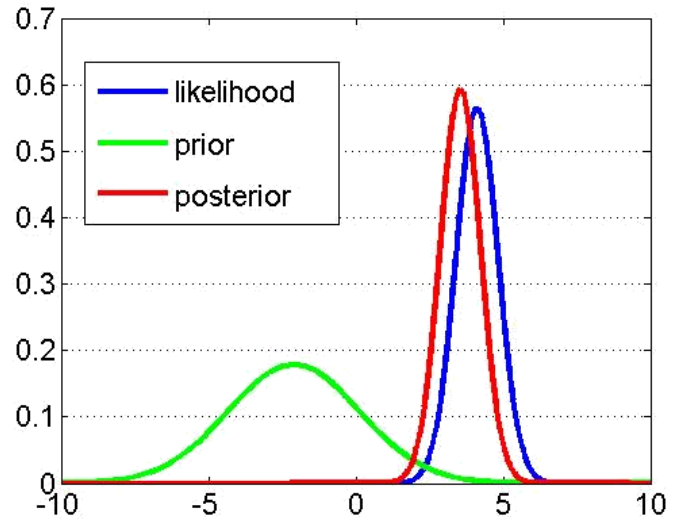

Model comparison

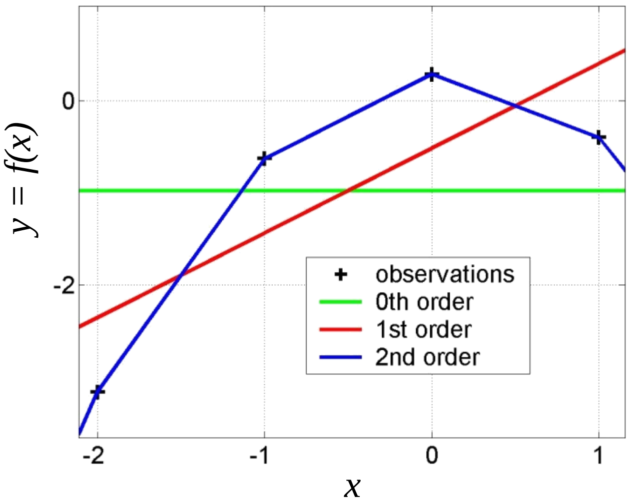

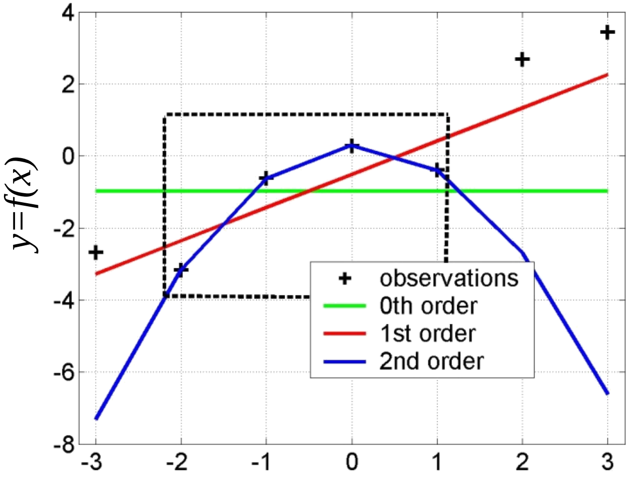

Model evidence:

$$p(y|m) = \int p(y|\theta, m)\, p(\theta|m)\, d\theta$$

"Occam's razor" :