



Bayesian model selection for statistical analysis of neural data: Lessons from fMRI

Joram Soch, BCCN Berlin

Seminar: Current Topics in Computational Neuroscience

WS 2018/2019: Statistical Analysis of Neural Data

Wed, 13/02/2019, 10.15 a.m.

Outline

Part I: Theory

1. What is (Bayesian) model selection?
2. The (very few) formulas that we'll need
3. A (very brief) introduction to fMRI data analysis
4. Bayesian model selection for fMRI data analysis
5. Theory Q&A

Part II: Practice

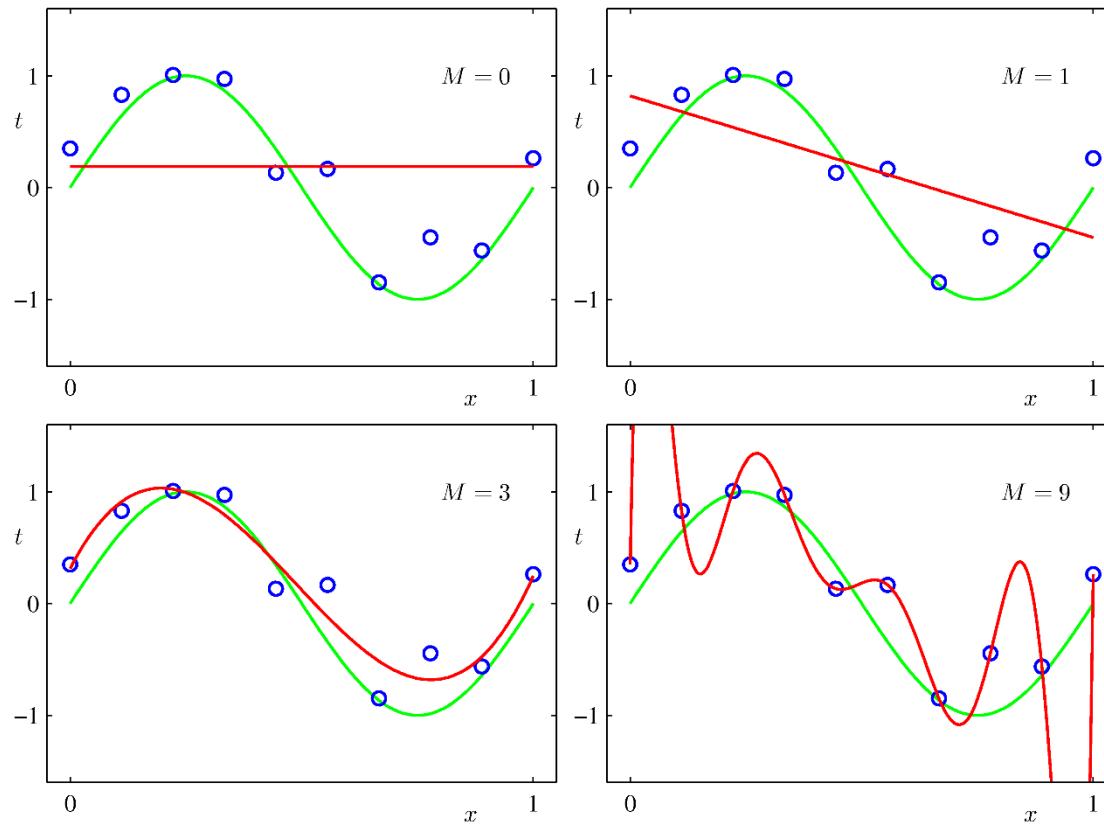
6. Code: PycvBMS – a Python module for model selection
7. Demo: Bayesian model selection for Neuropixel probes
8. What do I need to perform (Bayesian) model selection?
9. Practice Q&A

Part I: Theory

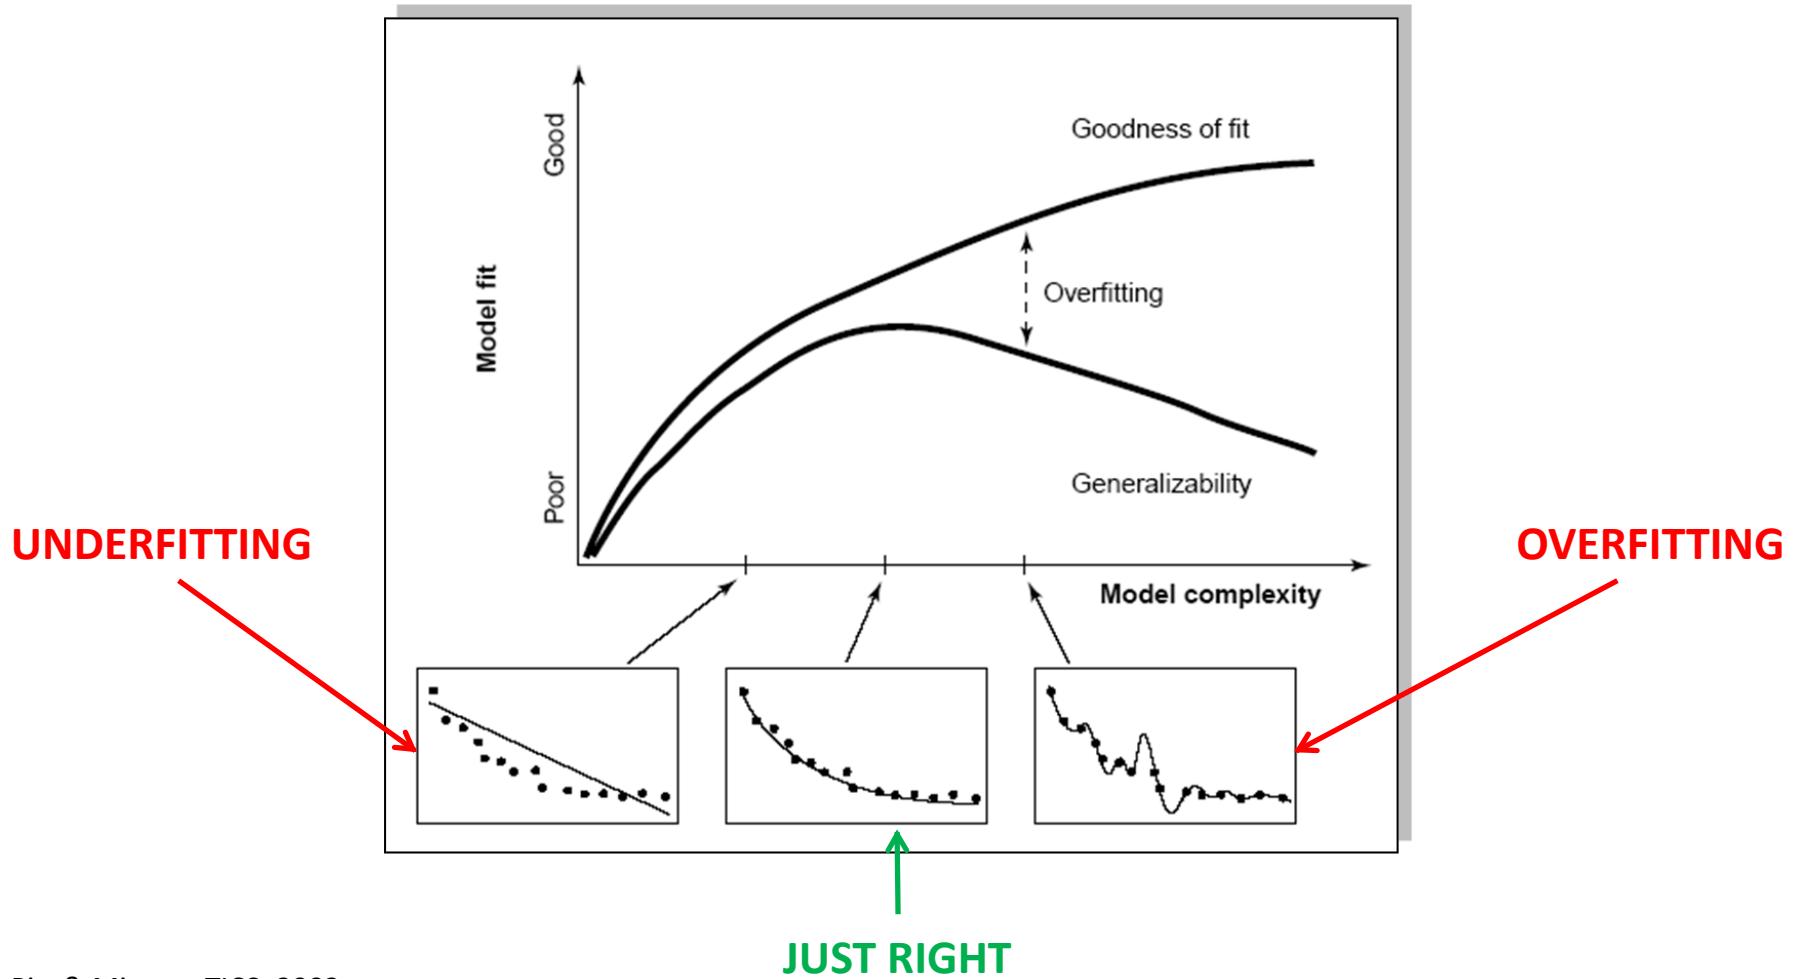
1.

What is (Bayesian) model selection?

Statistical model selection



Underfitting and Overfitting



General Principle

model quality = model accuracy – model complexity

2.

The (very few) formulas that we'll need

Classical information criteria

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

likelihood function

$$\hat{\theta} = \arg \max_{\theta} p(y|\theta, m)$$

maximum likelihood estimation

$$\text{MLL} = \log p(y|\hat{\theta}, m)$$

maximum log-likelihood

$$\text{IC} = -2 \text{MLL} + f_{\text{IC}}(n, k)$$

classical information criterion

$$f_{\text{AIC}} = 2k$$

$$f_{\text{BIC}} = k \log n$$

The Bayesian model evidence

posterior distribution	$p(\theta y, m) = \frac{p(y \theta, m) p(\theta m)}{p(y m)}$	likelihood function prior distribution (model) „evidence“
model evidence	$p(y m) = \int p(y \theta, m) p(\theta m) d\theta$	
log model evidence	$\text{LME}(m) = \log p(y m)$	

The cross-validated LME

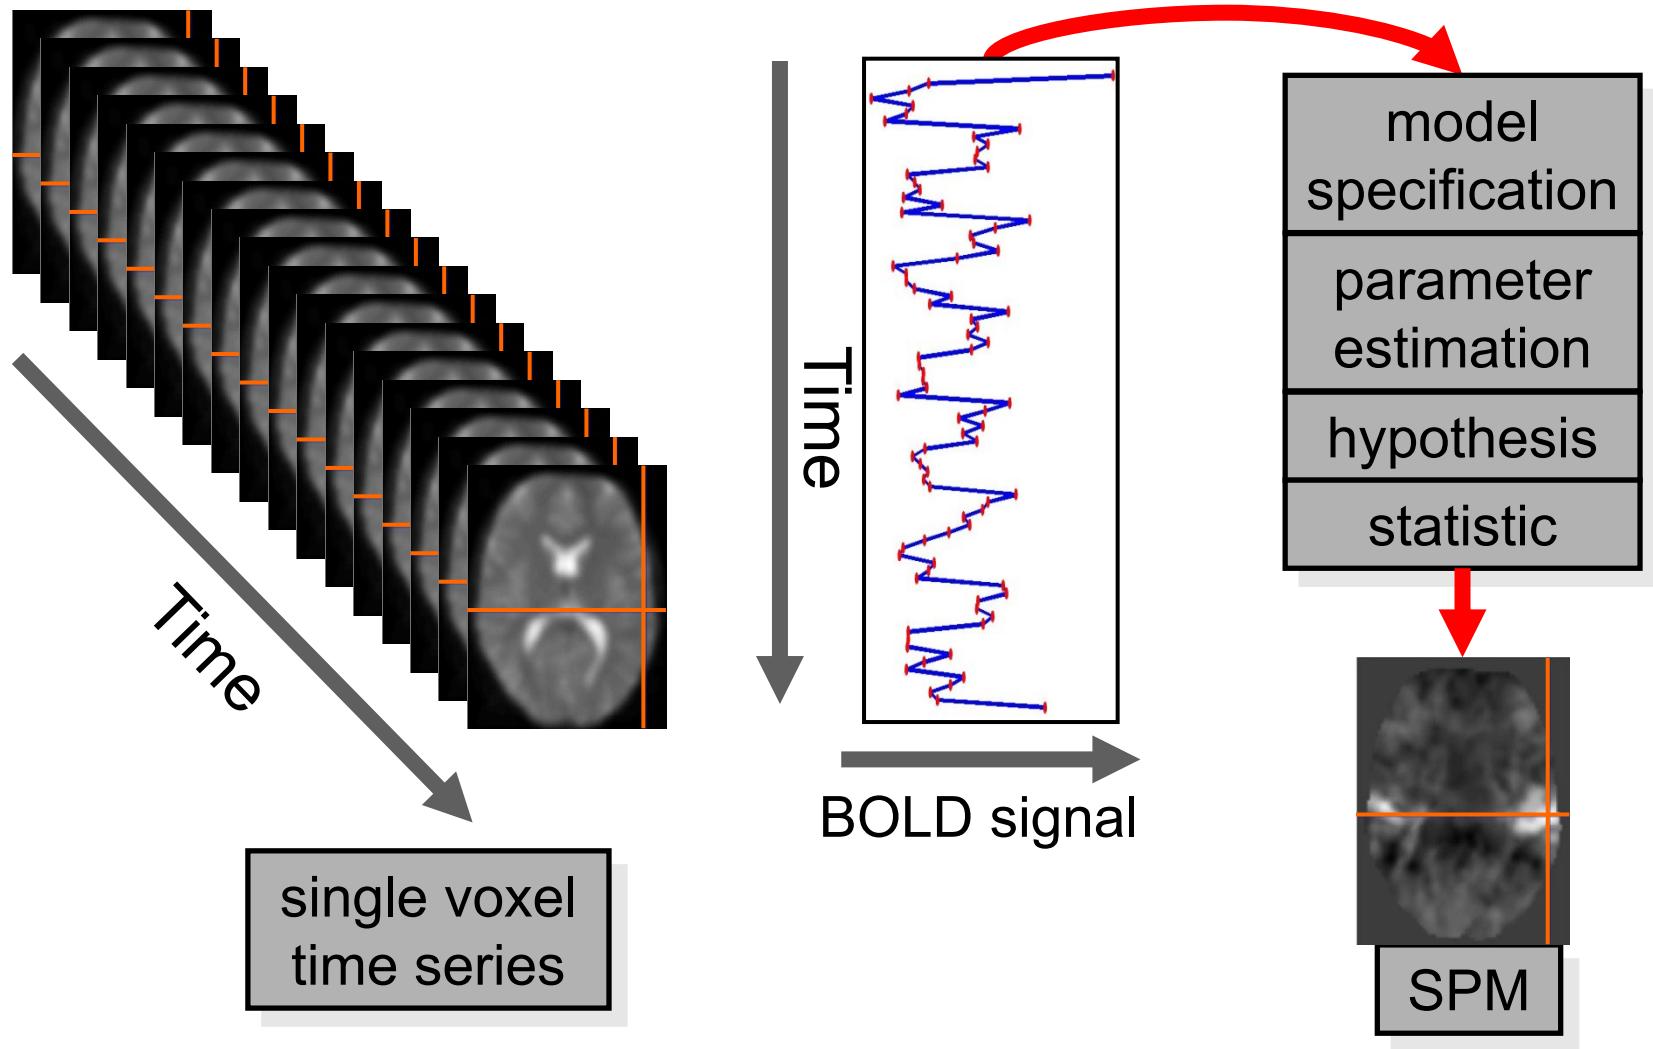
$$p(y|m) = \int p(y|\theta, m) p(\theta|m) d\theta$$

$$\text{cvLME}(m) = \sum_{i=1}^S \log \int p(y_i|\theta) p(\theta| \cup_{j \neq i} y_j) d\theta$$

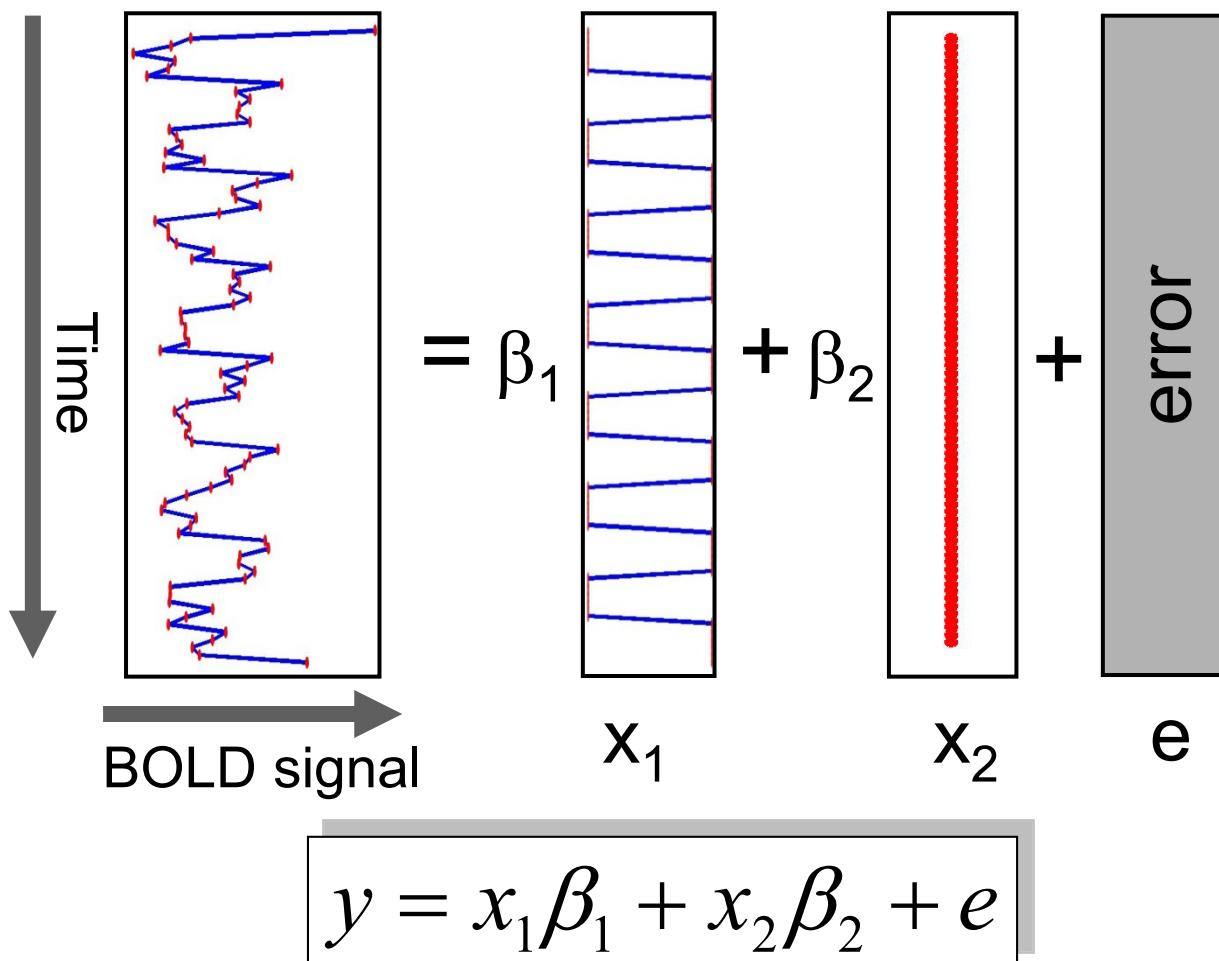
3.

A (very brief) introduction to fMRI data analysis

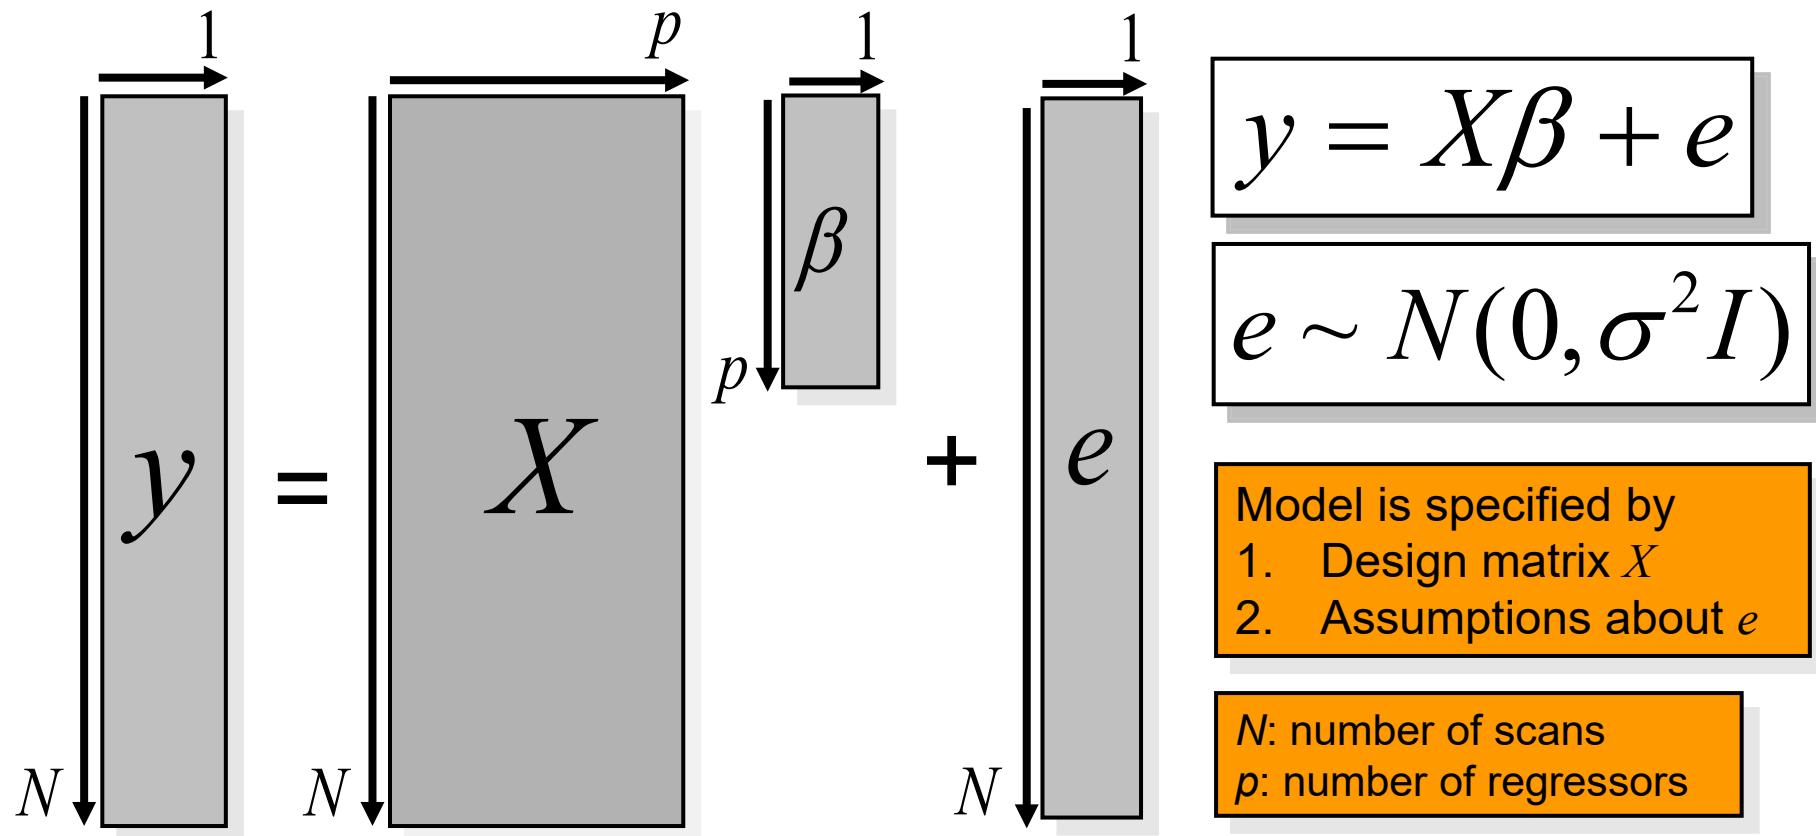
Voxel-wise time series analysis



Single-voxel regression model



Mass-univariate analysis: voxel-wise GLM



The design matrix embodies all available knowledge about experimentally controlled factors and potential confounds.

Ordinary least squares: beta estimates

$$X^T \hat{\varepsilon} = 0$$

$$X^T(y - X\hat{\beta}) = 0$$

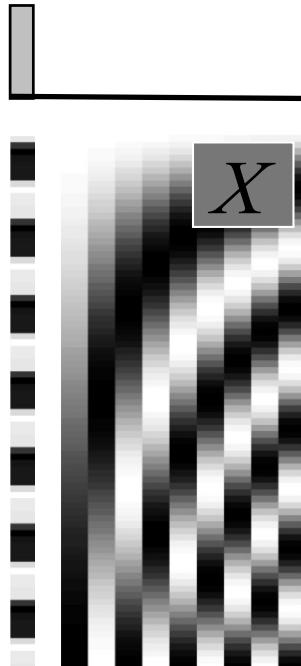
$$X^T y - X^T X \hat{\beta} = 0$$

$$X^T X \hat{\beta} = X^T y$$

$$\hat{\beta} = (X^T X)^{-1} X^T y$$

Contrasts & statistical parametric maps

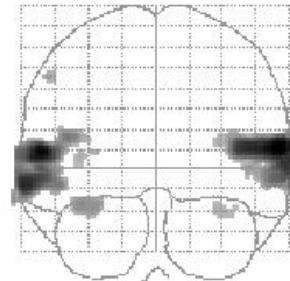
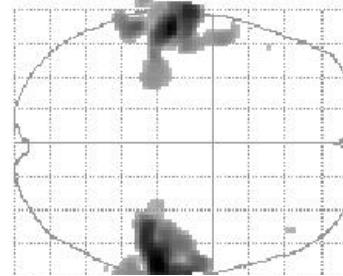
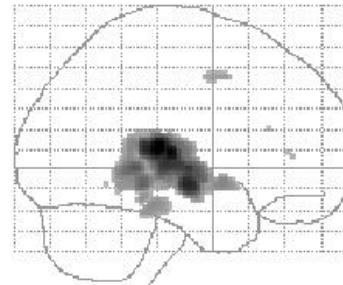
$c = [1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]$



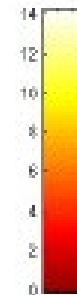
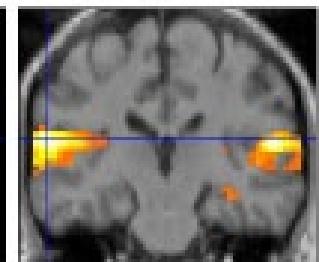
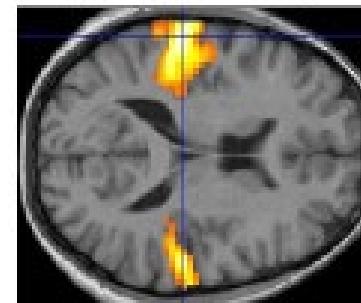
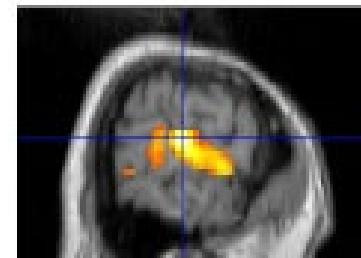
Q: activation during listening ?

Null hypothesis: $\beta_1 = 0$

$$t = \frac{c^T \hat{\beta}}{Std(c^T \hat{\beta})}$$



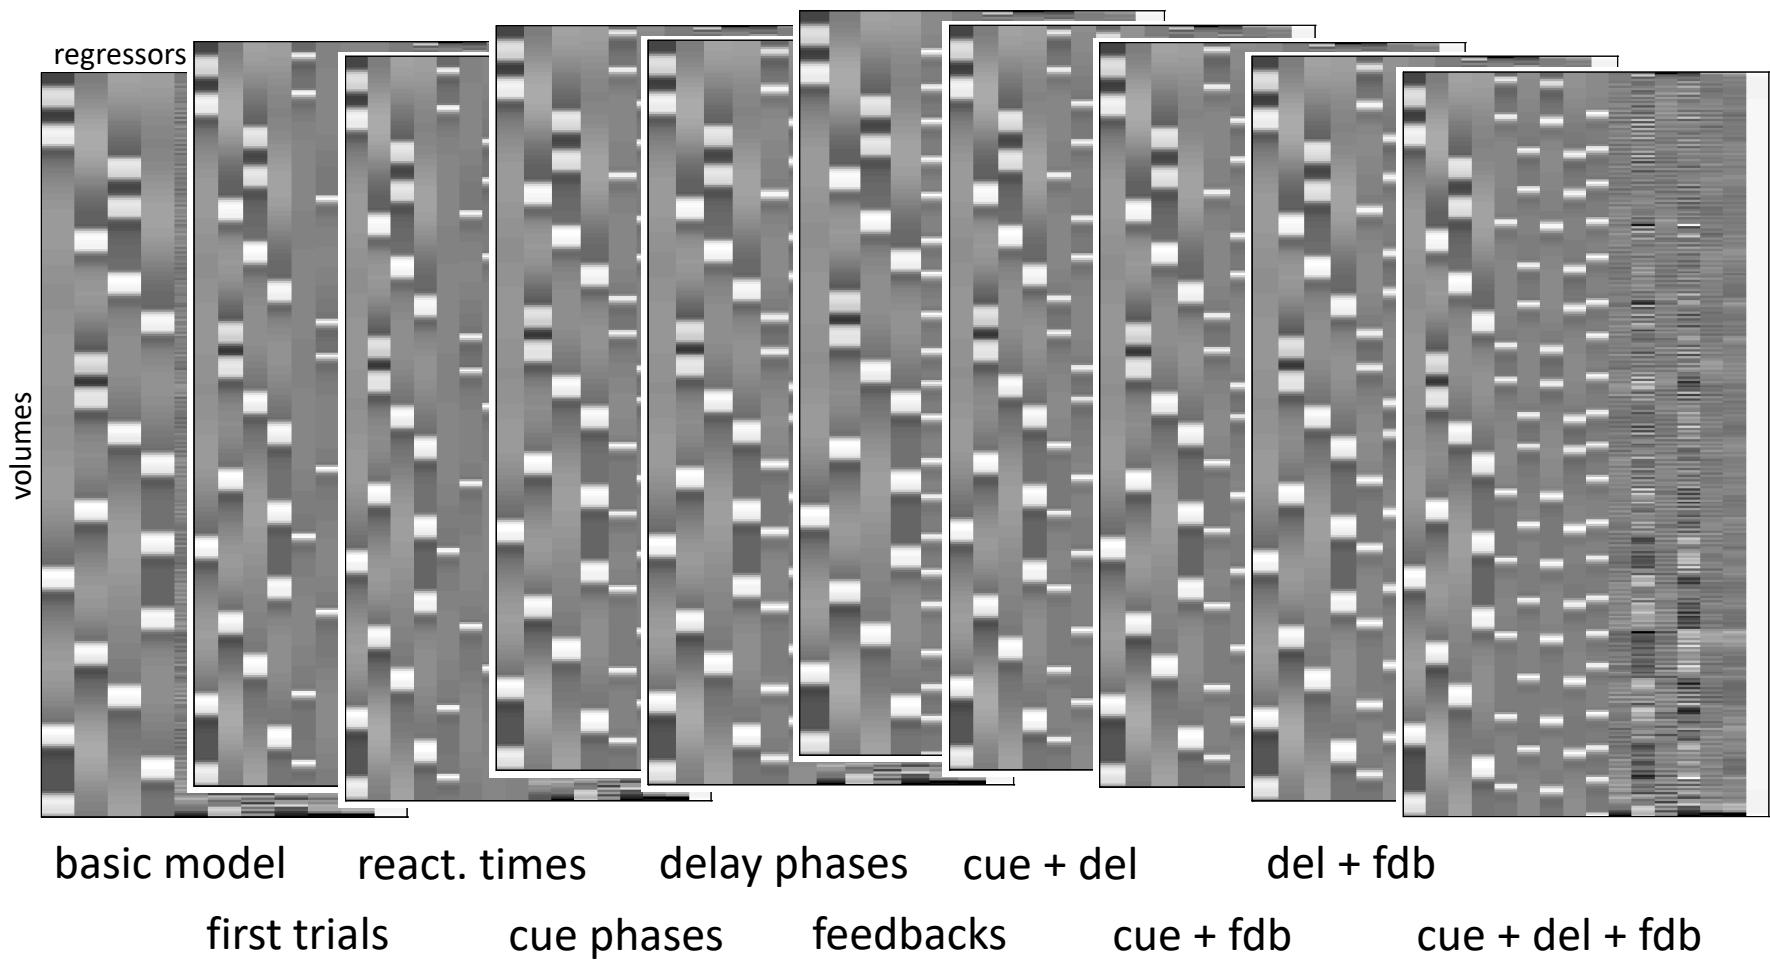
SPM{T₇₃}



4.

Bayesian model selection for fMRI data analysis

Many plausible GLMs for a given fMRI data set



cross-validated Bayesian model selection

First-level analysis:

$$y = X\beta + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2 V)$$

$$p(y|m) = \iint p(y|\beta, \sigma^2, m) p(\beta, \sigma^2|m) d\beta d\sigma^2$$

$$\text{cvLME}(m) = \sum_{i=1}^S \log \int p(y_i|\theta) p(\theta| \cup_{j \neq i} y_j) d\theta$$

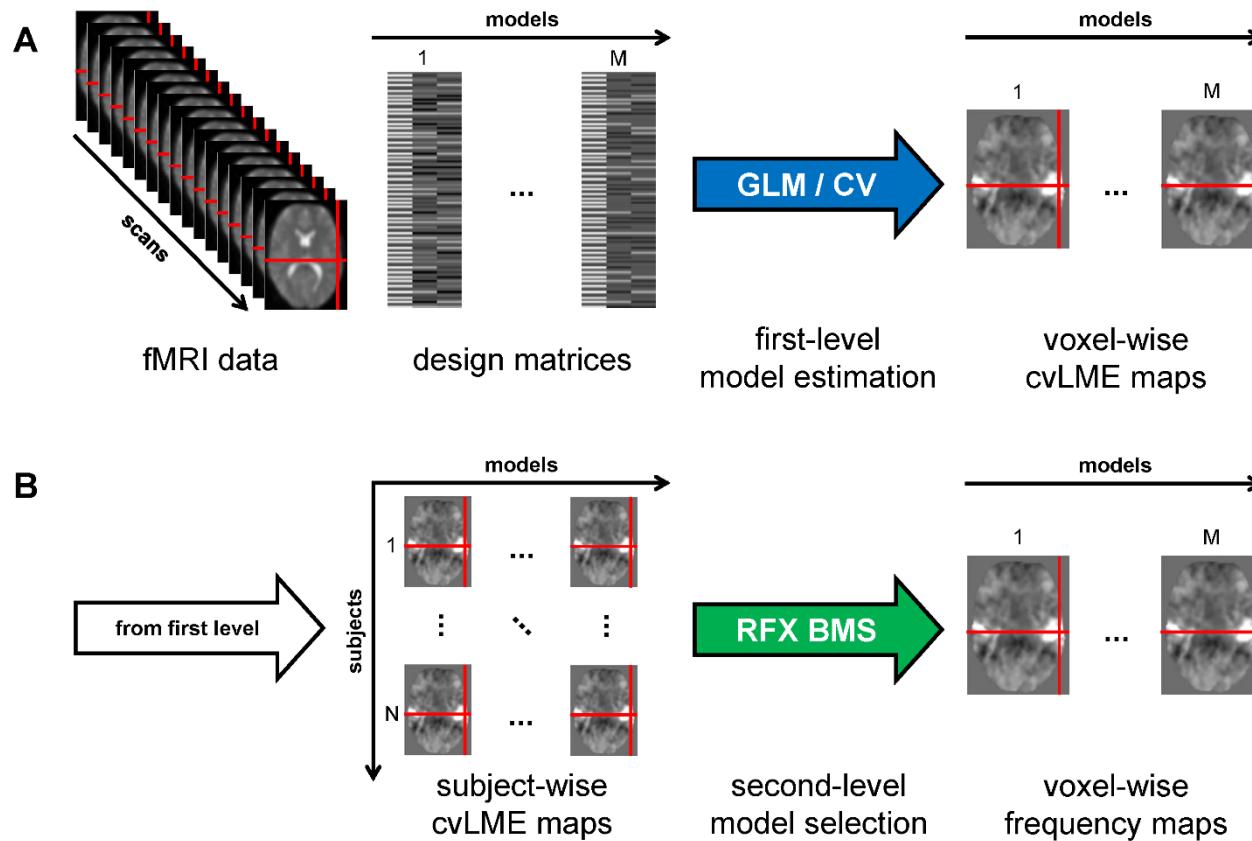
Second-level analysis:

$$p(y|m) = \prod_{i=1}^N p(y_i|m_i) = \prod_{i=1}^N \prod_{j=1}^M p(y_i|e_j)^{m_{ij}}$$

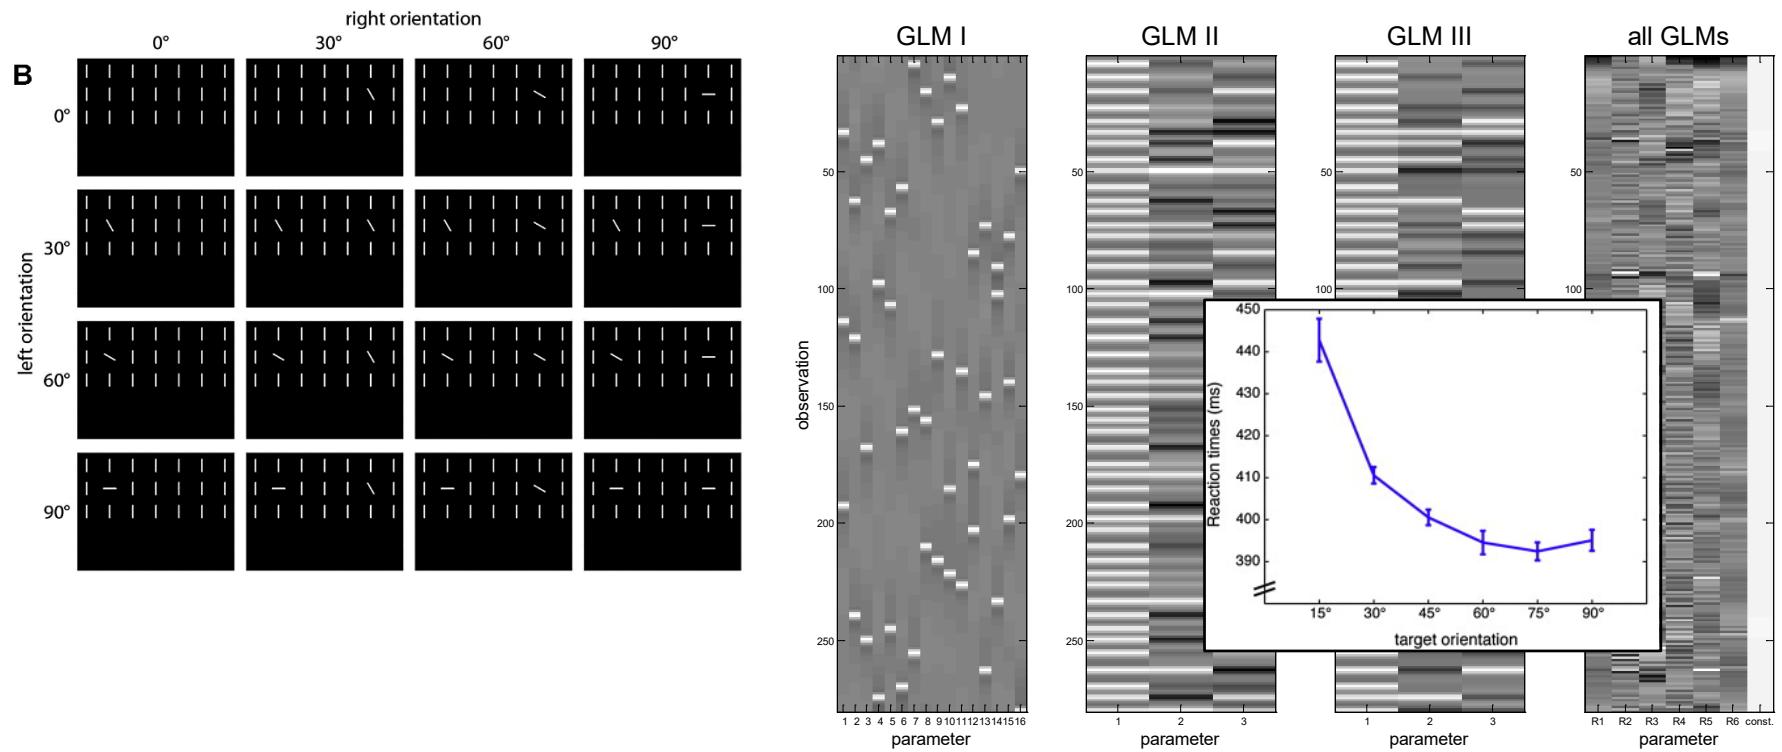
$$p(m|r) = \prod_{i=1}^N \text{Mult}(m_i; 1, r) = \prod_{i=1}^N \prod_{j=1}^M r_j^{m_{ij}}$$

$$p(r|\alpha) = \text{Dir}(r; \alpha) = \frac{\Gamma(\sum \alpha_j)}{\sum \Gamma(\alpha_j)} \prod_{j=1}^M r_j^{\alpha_j - 1}$$

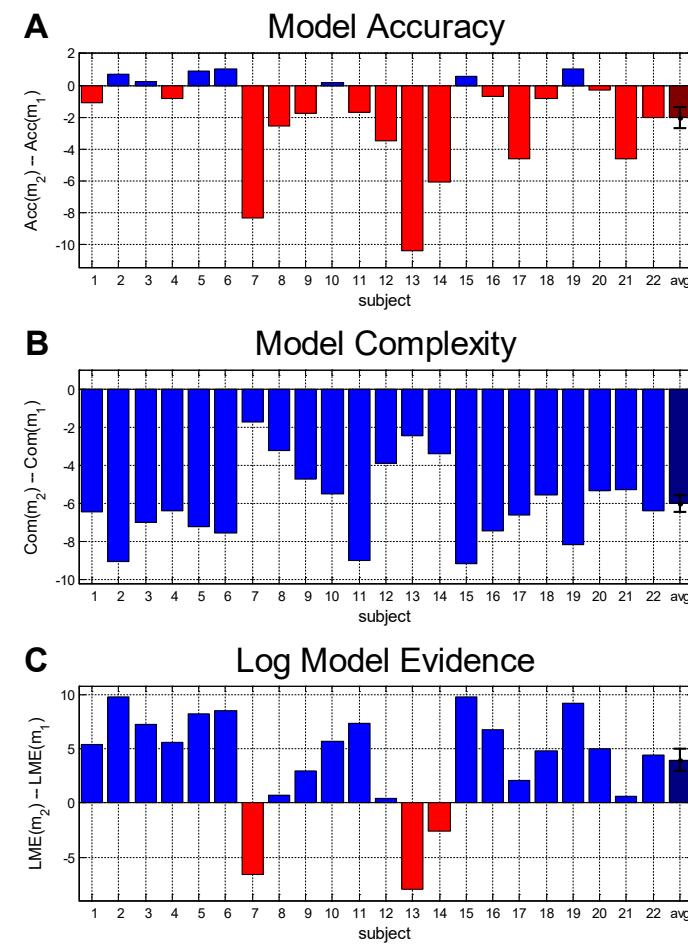
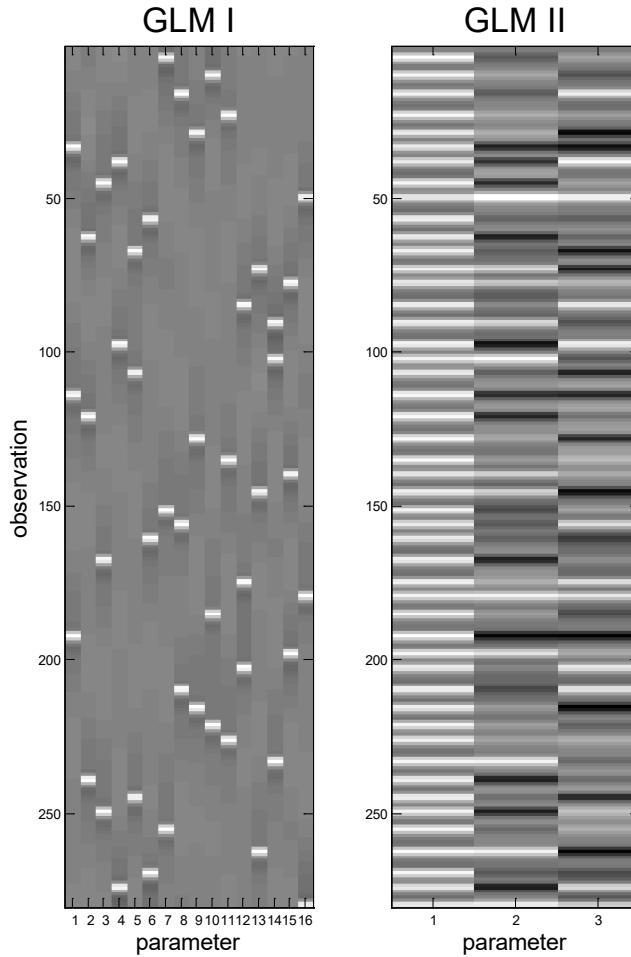
The cvBMS approach in practice



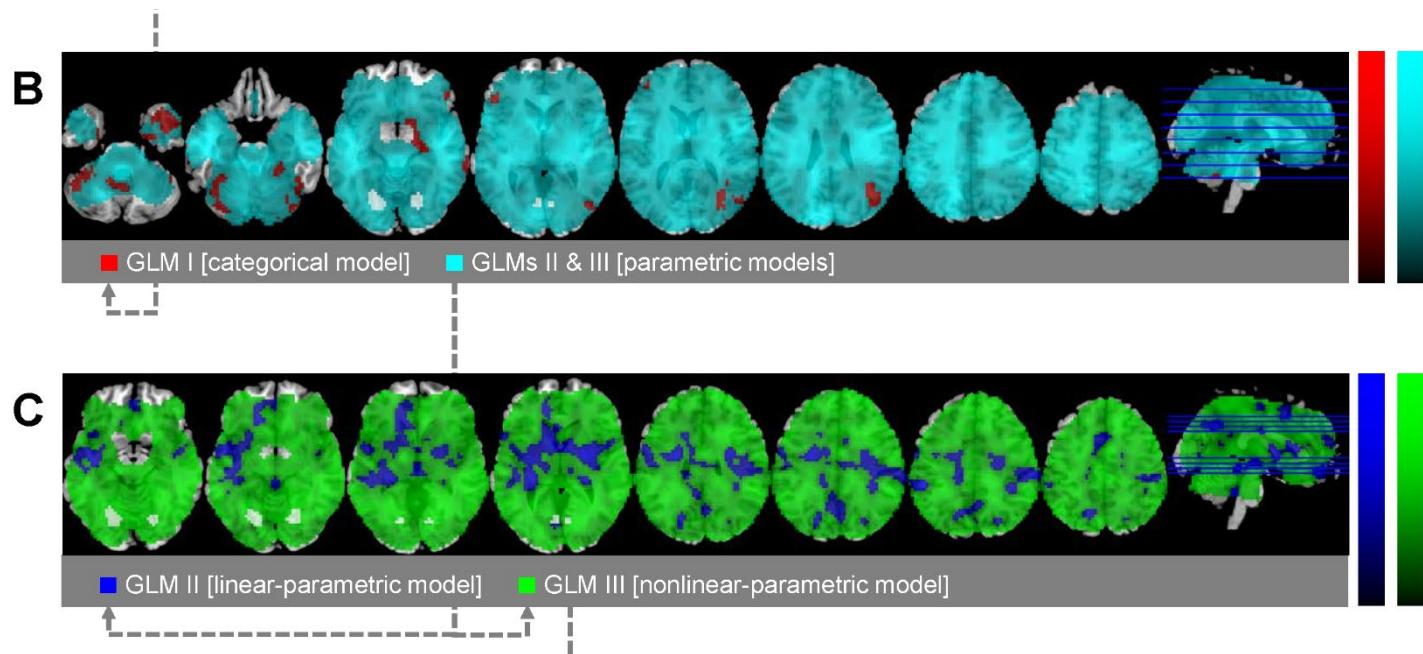
Empirical example: model space



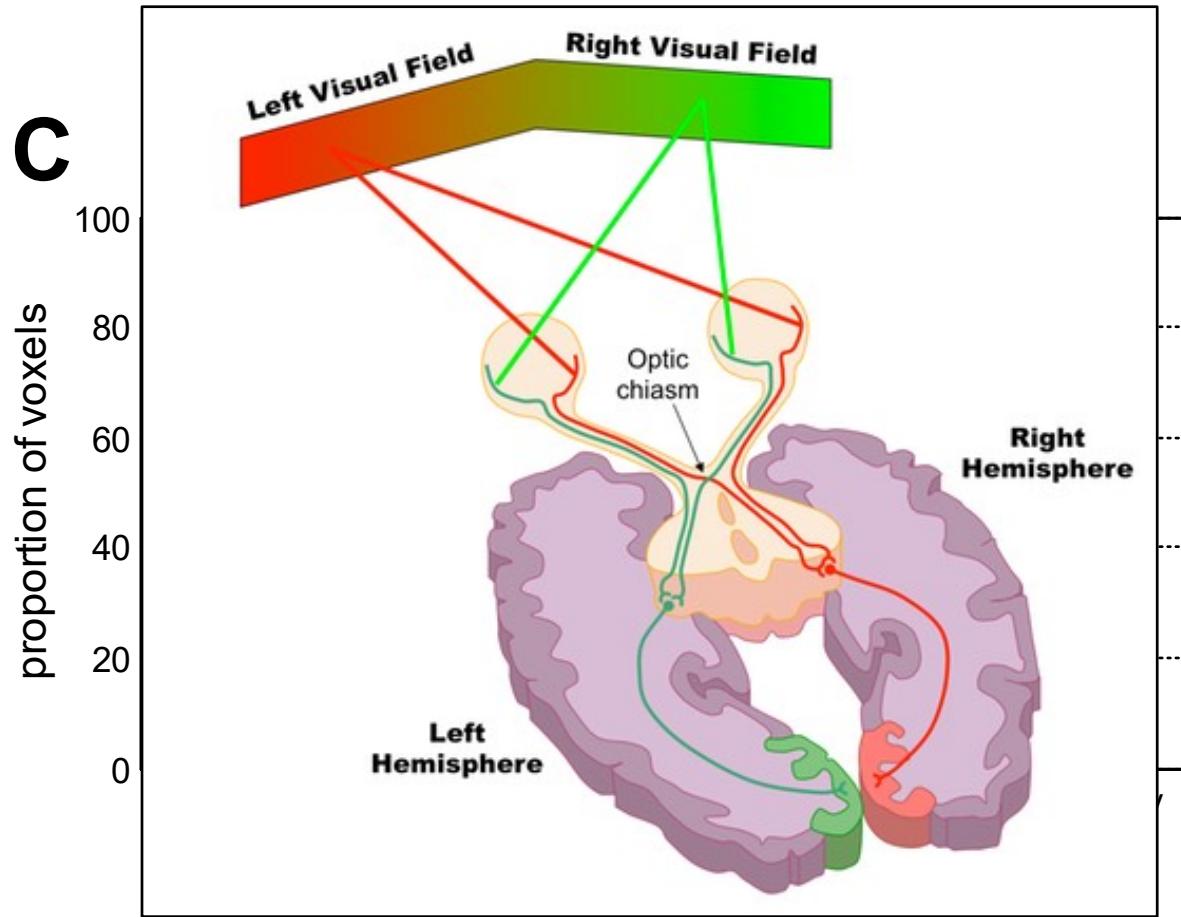
Model selection: GLMs I/II



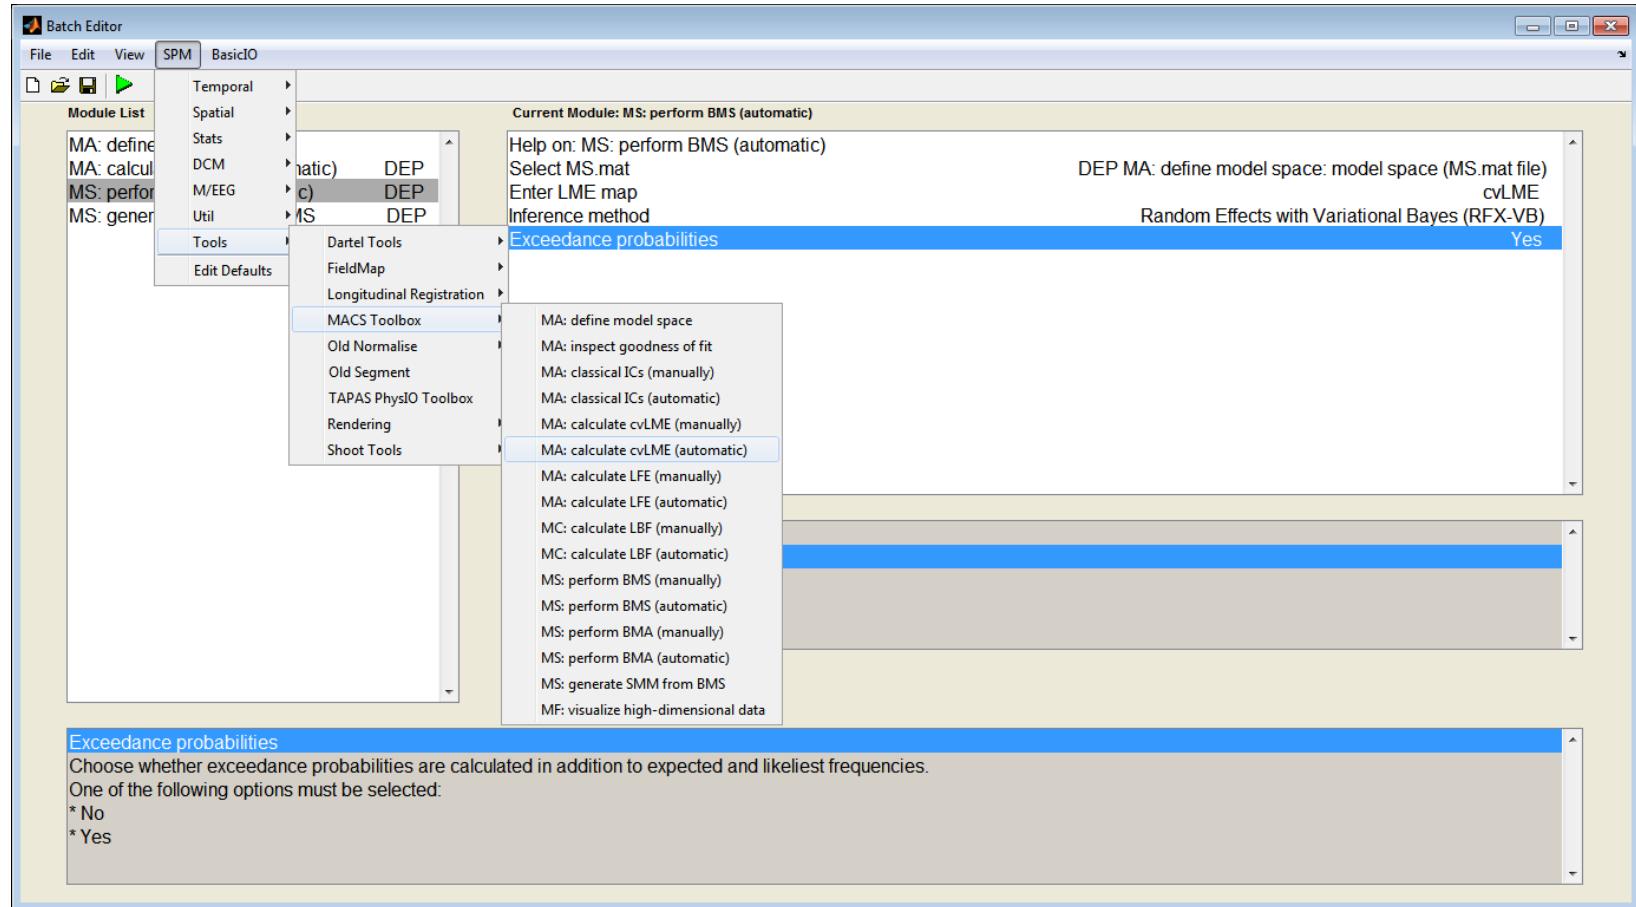
Model selection: GLMs I-III



Model selection: GLMs III-l/r



The MACS toolbox for SPM



5.

Theory Q&A



Part II: Practice

6.

Code: PycvBMS – a Python module for model selection

PycvBMS – available from GitHub

Python module for cross-validated Bayesian model selection

The screenshot shows the GitHub repository page for PycvBMS. At the top, it displays basic statistics: 3 commits, 1 branch, 0 releases, and 1 contributor. Below this, there's a dropdown for the branch (set to master), a 'New pull request' button, a 'Find file' button, and a prominent green 'Clone or download' button. The main content area shows a list of commits. The first commit is by JoramSoch, updating the README.md. The second commit is a project upload. The third commit is another update to the README.md. A large red note overlay covers the top half of the page, stating: 'Note: This repository is now obsolete. Please consider using the cvLME package!'. Below this note, the repository description reads: 'This module collects methods for calculating the cross-validated model selection (cvLME). Currently, it only features the general linear model (GLM); more functionality will come soon.' A 'Getting Started' section includes a code snippet:

```
import cvBMS
import numpy as np

# have your data ready
Y # an n x v data matrix
X1 # an n x p1 design matrix from one model
X2 # an n x p2 design matrix from another model
V # an n x n covariance matrix, probably V = np.eye(n)

# two models
m1 = cvBMS.GLM(Y, X1, V)
m2 = cvBMS.GLM(Y, X2, V)

# model space
ms = cvBMS.MS(np.r_[m1.cvLME(), m2.cvLME()])
PP = ms.PP() # posterior model probabilities
```

URL: <https://github.com/JoramSoch/PycvBMS>; <https://github.com/JoramSoch/cvLME>

PycvBMS – four easy steps

1. define GLM(s)
2. calculate cvLMEs
3. define model space
4. calculate PPs

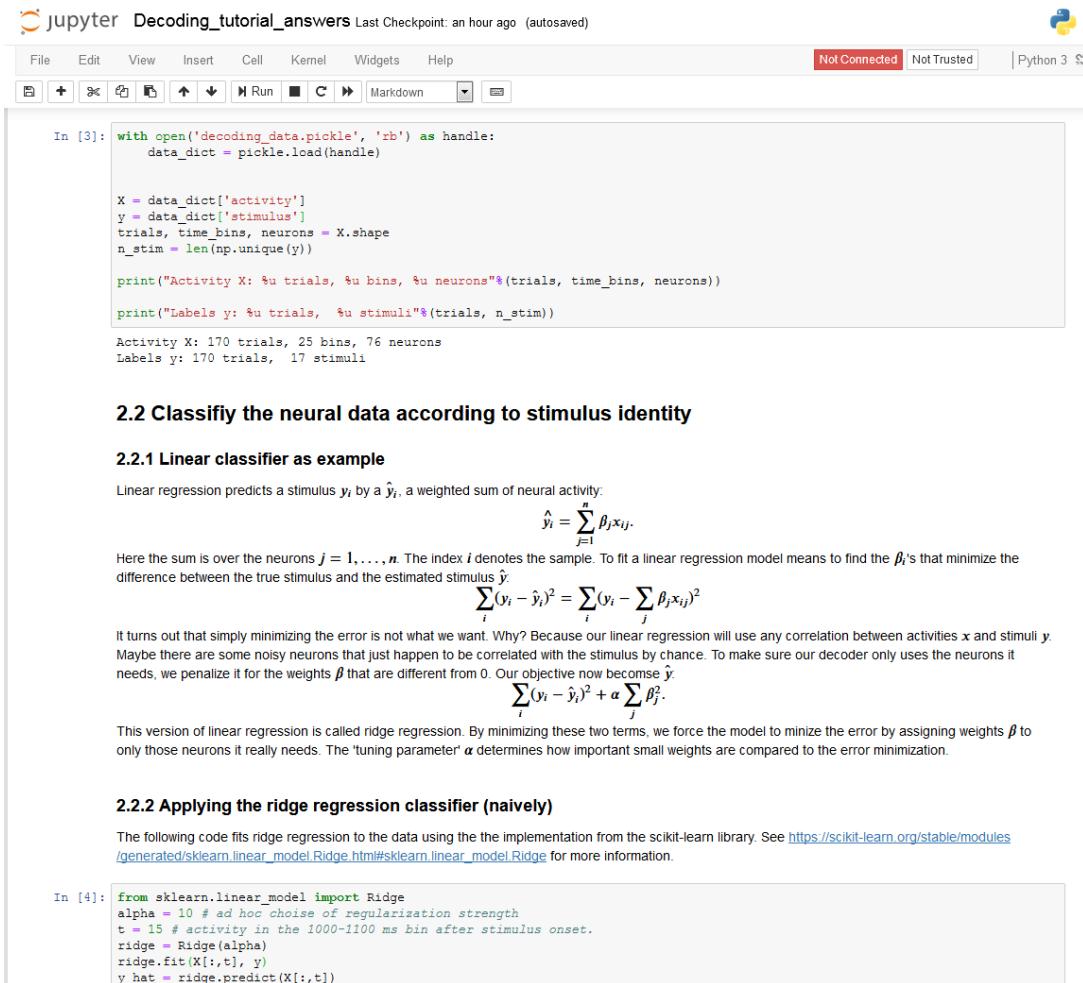
PycvBMS – a little demo

[switch to Anaconda/Spyder]

7.

Demo: Bayesian model selection for Neuropixel probes

What do we analyze?



The screenshot shows a Jupyter Notebook interface with the title "jupyter Decoding_tutorial_answers" and a status bar indicating "Last Checkpoint: an hour ago (autosaved)", "Not Connected", "Not Trusted", and "Python 3".

In [3]:

```
with open('decoding_data.pickle', 'rb') as handle:  
    data_dict = pickle.load(handle)  
  
X = data_dict['activity']  
y = data_dict['stimulus']  
trials, time_bins, neurons = X.shape  
n_stim = len(np.unique(y))  
  
print("Activity X: %u trials, %u bins, %u neurons" % (trials, time_bins, neurons))  
print("Labels y: %u trials, %u stimuli" % (trials, n_stim))  
  
Activity X: 170 trials, 25 bins, 76 neurons  
Labels y: 170 trials, 17 stimuli
```

2.2 Classify the neural data according to stimulus identity

2.2.1 Linear classifier as example

Linear regression predicts a stimulus y_i by a \hat{y}_i , a weighted sum of neural activity:

$$\hat{y}_i = \sum_{j=1}^n \beta_j x_{ij}.$$

Here the sum is over the neurons $j = 1, \dots, n$. The index i denotes the sample. To fit a linear regression model means to find the β_j 's that minimize the difference between the true stimulus and the estimated stimulus \hat{y} :

$$\sum_i (y_i - \hat{y}_i)^2 = \sum_i (y_i - \sum_j \beta_j x_{ij})^2$$

It turns out that simply minimizing the error is not what we want. Why? Because our linear regression will use any correlation between activities x and stimuli y . Maybe there are some noisy neurons that just happen to be correlated with the stimulus by chance. To make sure our decoder only uses the neurons it needs, we penalize it for the weights β that are different from 0. Our objective now become \hat{y} :

$$\sum_i (y_i - \hat{y}_i)^2 + \alpha \sum_j \beta_j^2.$$

This version of linear regression is called ridge regression. By minimizing these two terms, we force the model to minimize the error by assigning weights β to only those neurons it really needs. The 'tuning parameter' α determines how important small weights are compared to the error minimization.

2.2.2 Applying the ridge regression classifier (naively)

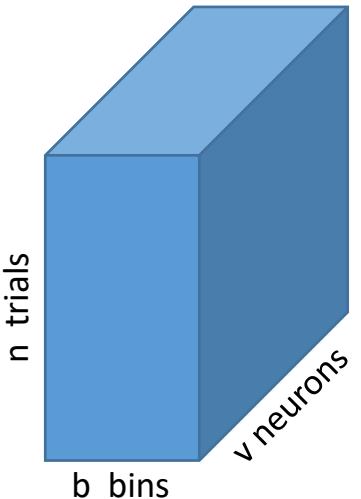
The following code fits ridge regression to the data using the implementation from the scikit-learn library. See https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.Ridge.html#sklearn.linear_model.Ridge for more information.

In [4]:

```
from sklearn.linear_model import Ridge  
alpha = 10 # ad hoc choice of regularization strength  
t = 15 # activity in the 1000-1100 ms bin after stimulus onset.  
ridge = Ridge(alpha)  
ridge.fit(X[:,t], y)  
y_hat = ridge.predict(X[:,t])
```

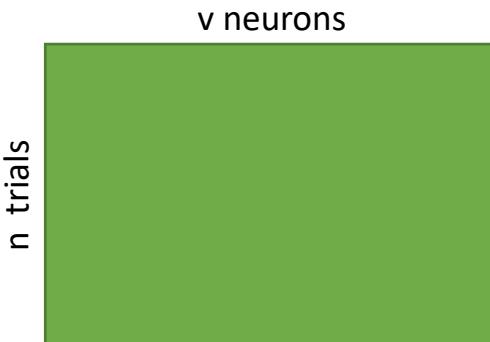
What data were acquired?

Data come as a



3D NumPy array of spike counts/100 ms

and we will analyze the



matrix belonging to the $k = 16$ -th time bin, i.e. 1000-1100 ms after stimulus onset (out of 1500 ms total stimulation time).

Which stimuli were presented?

Each trial is coded as ...

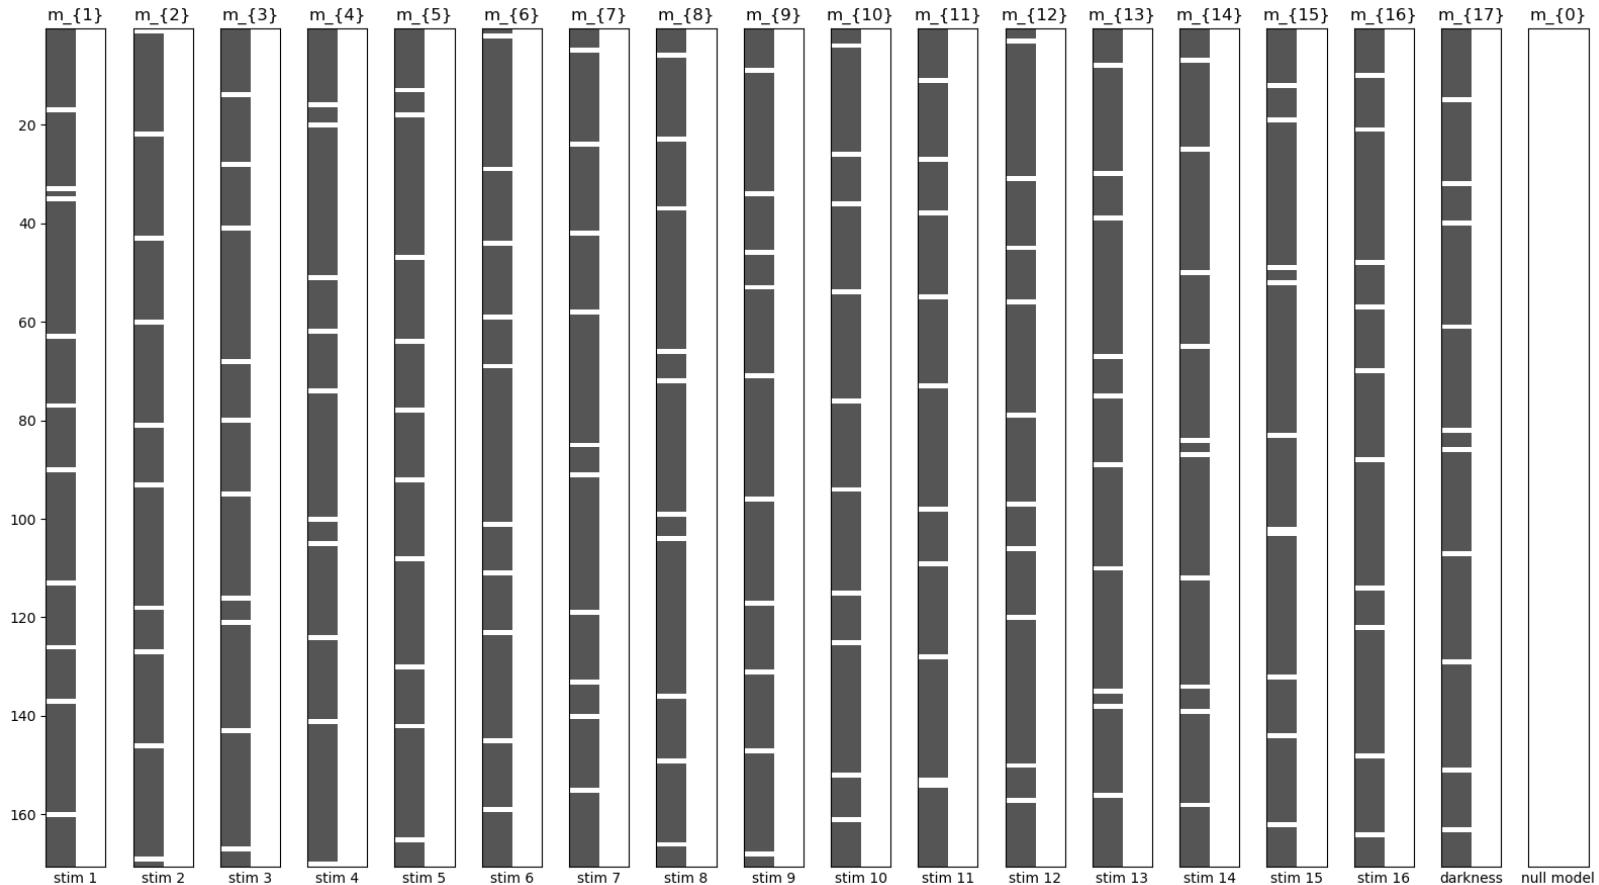
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
0°		45°		90°		135°		180°		225°		270°		315°		darkness

corresponding to ...

Research Question No. 1

„Which neuron reacts to which stimulus?“

Q1: model space

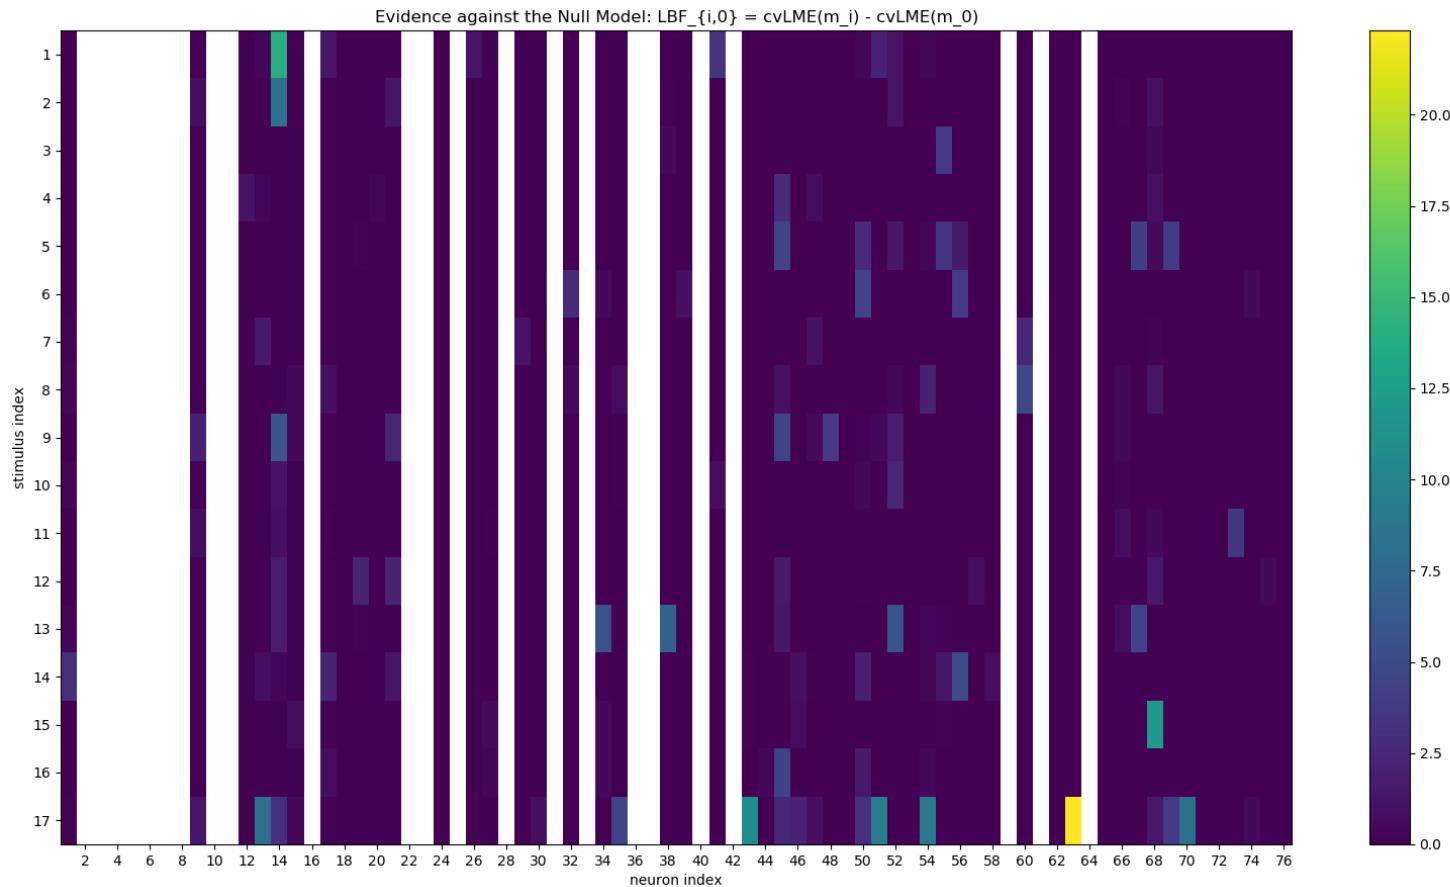


$$m_i : y_j = \beta_0 + \beta_1 s_{ij} + \varepsilon_{ij}$$

Q1: model selection tool

$$\text{LBF}_{12} = \log \text{BF}_{12} = \log \frac{p(y|m_1)}{p(y|m_2)} = \text{LME}(m_1) - \text{LME}(m_2)$$

Q1: model comparison



$$LBF_{i,0} = cvLME(m_i) - cvLME(m_0)$$

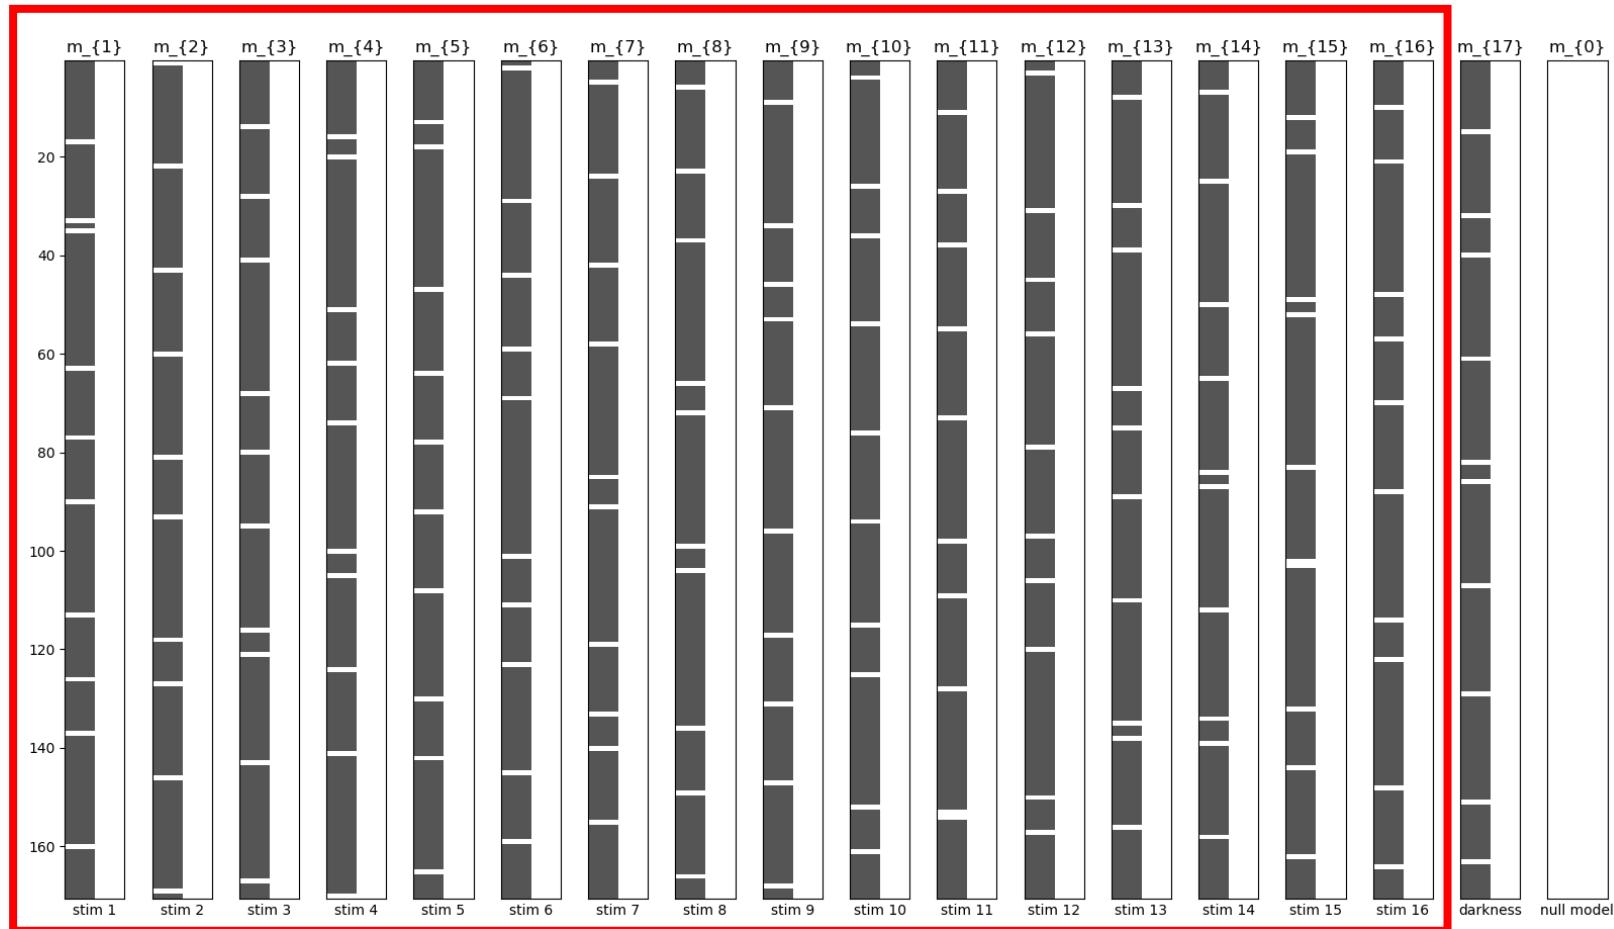
Research Question No. 2

„Does the neuron react to visual stimulation *per se*?“

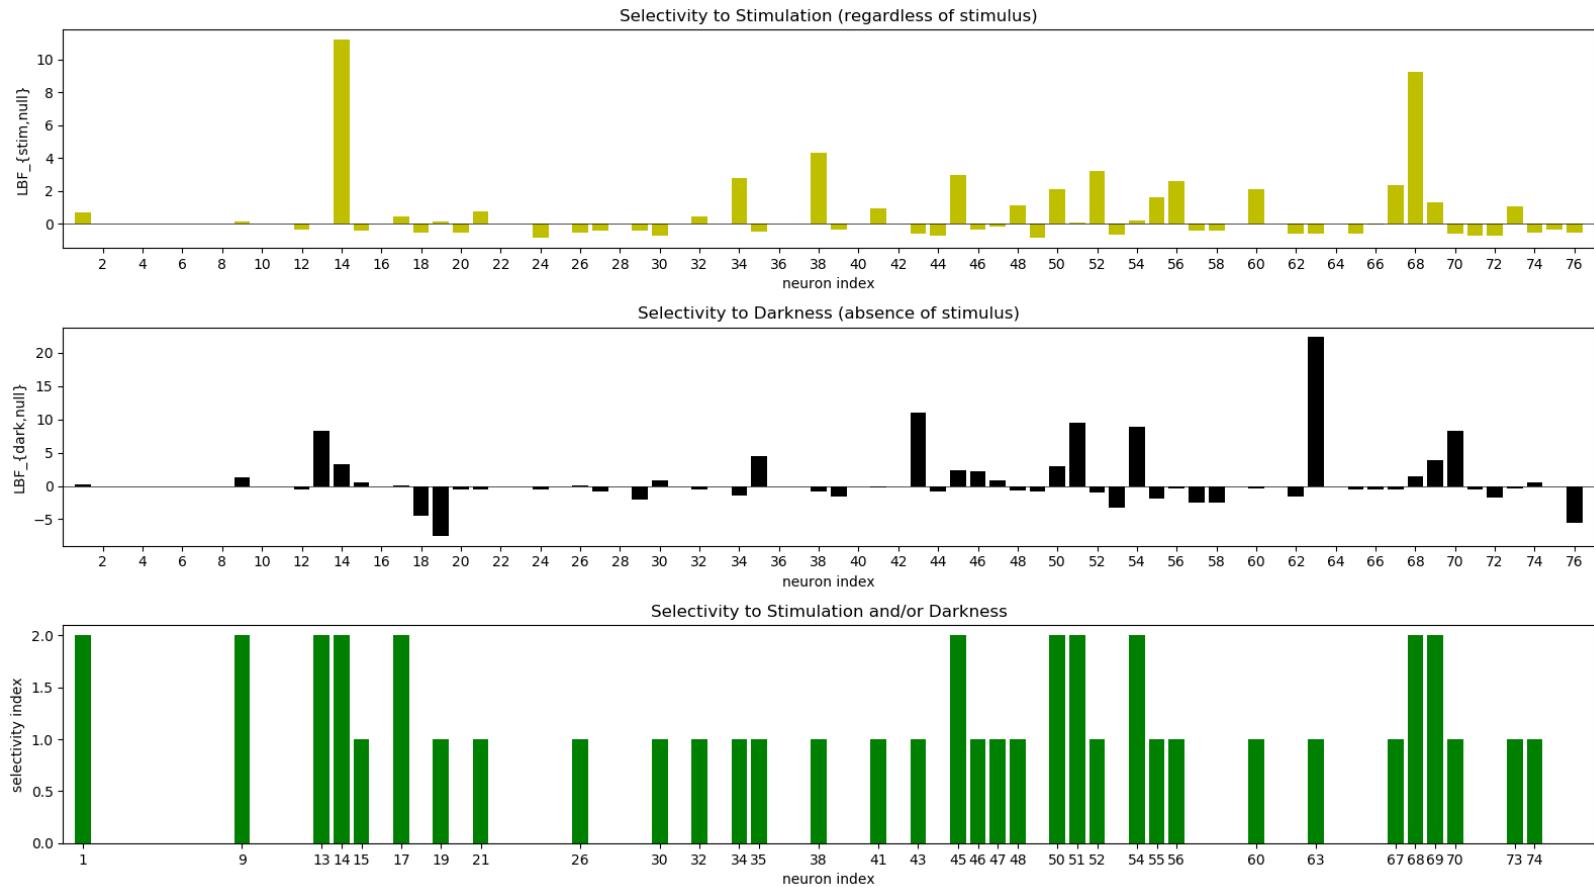
Q2: model selection tool

$$\text{LFE}(f) = \log p(y|f) = \log \sum_{m \in f} p(y|m) p(m|f)$$

Q2: model space



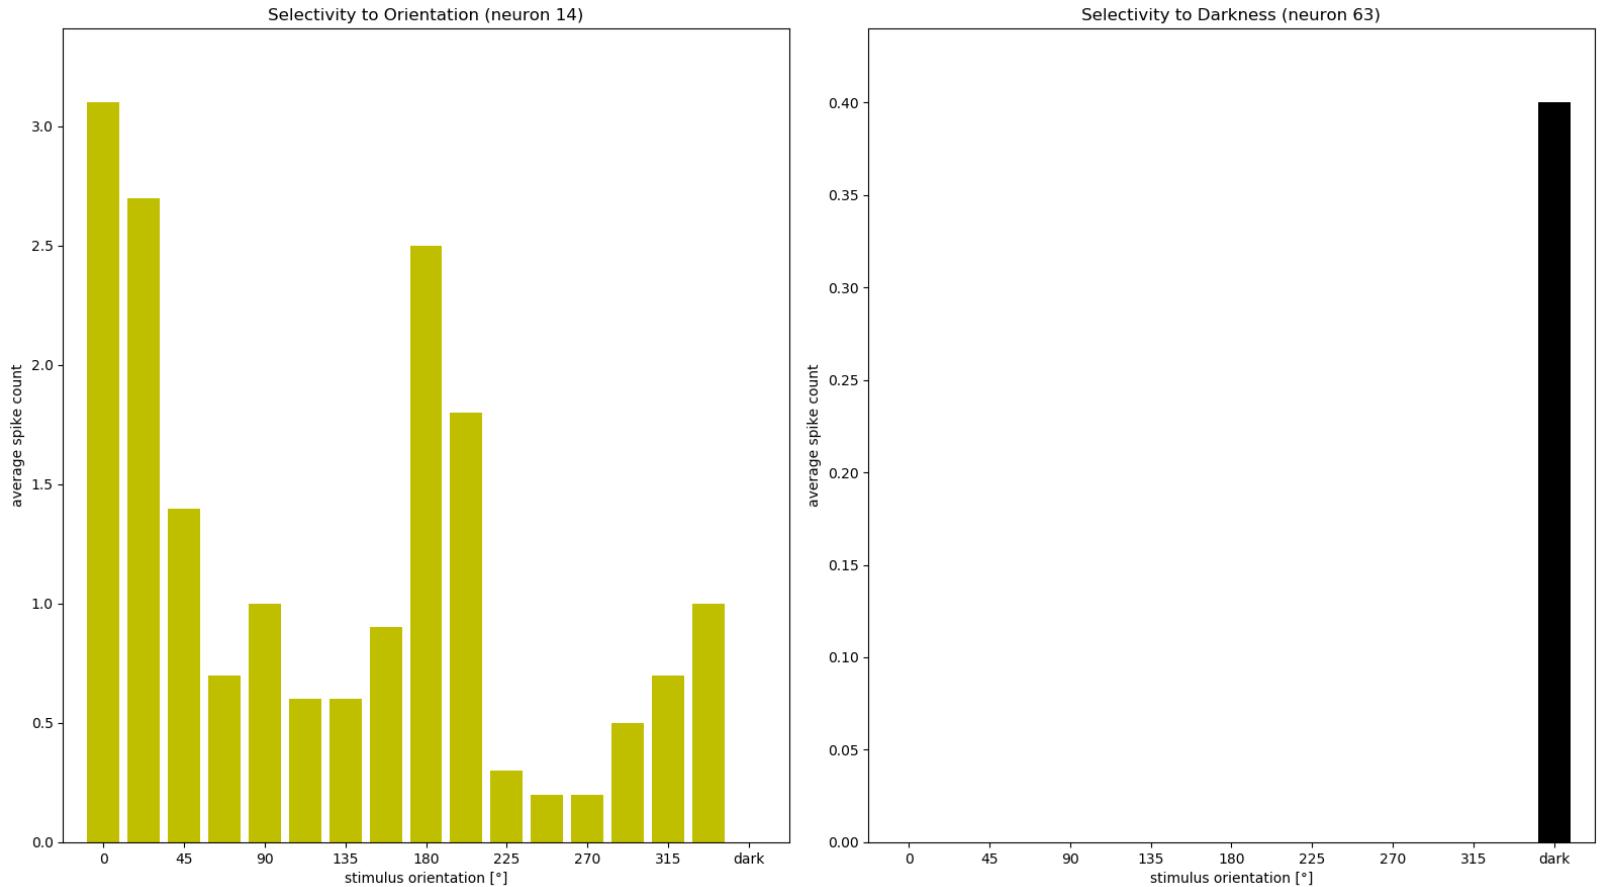
Q2: family comparison



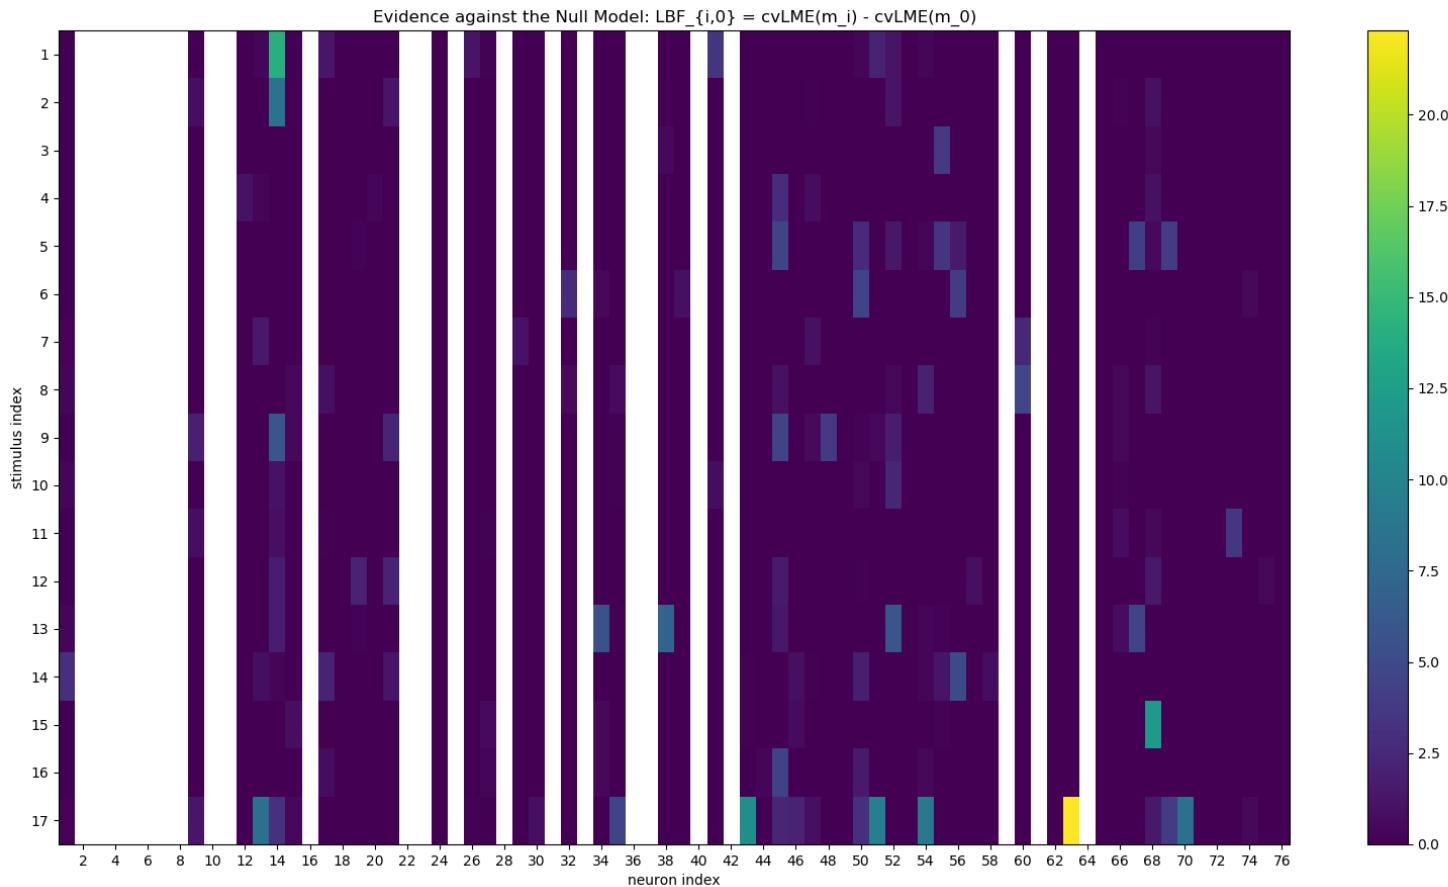
$$LBF_{\text{stim},\text{null}} = LFE(f_{1-16}) - \text{cvLME}(m_0)$$

$$LBF_{\text{dark},\text{null}} = \text{cvLME}(m_{17}) - \text{cvLME}(m_0)$$

Q2: data inspection



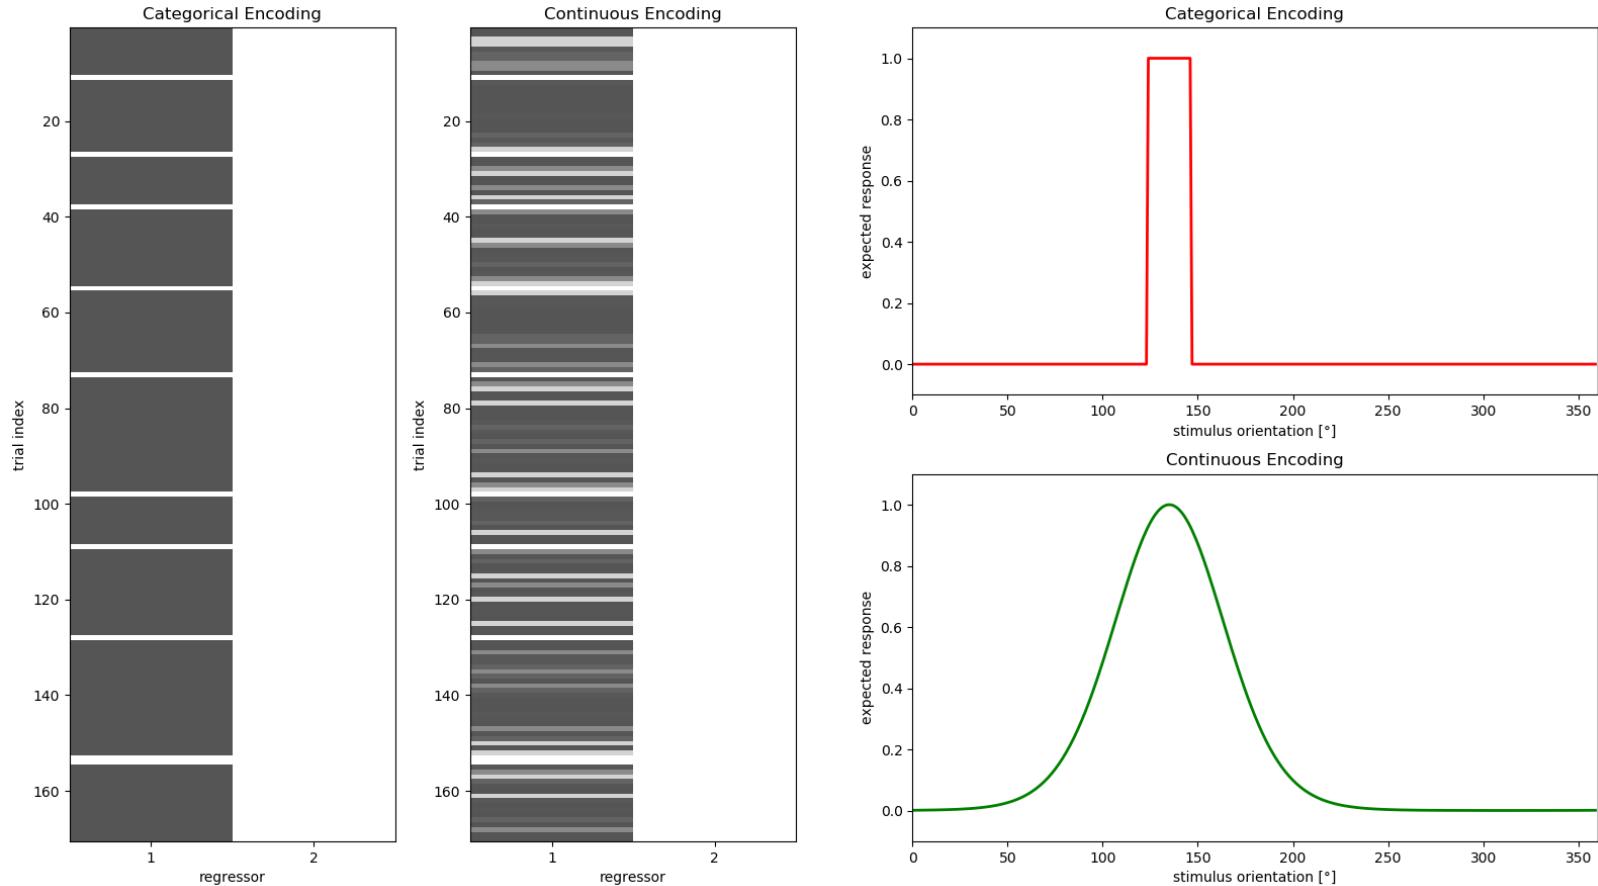
Q2: data inspection



Research Question No. 3

„Does the neuron respond categorically or does it exhibit some tuning function?“

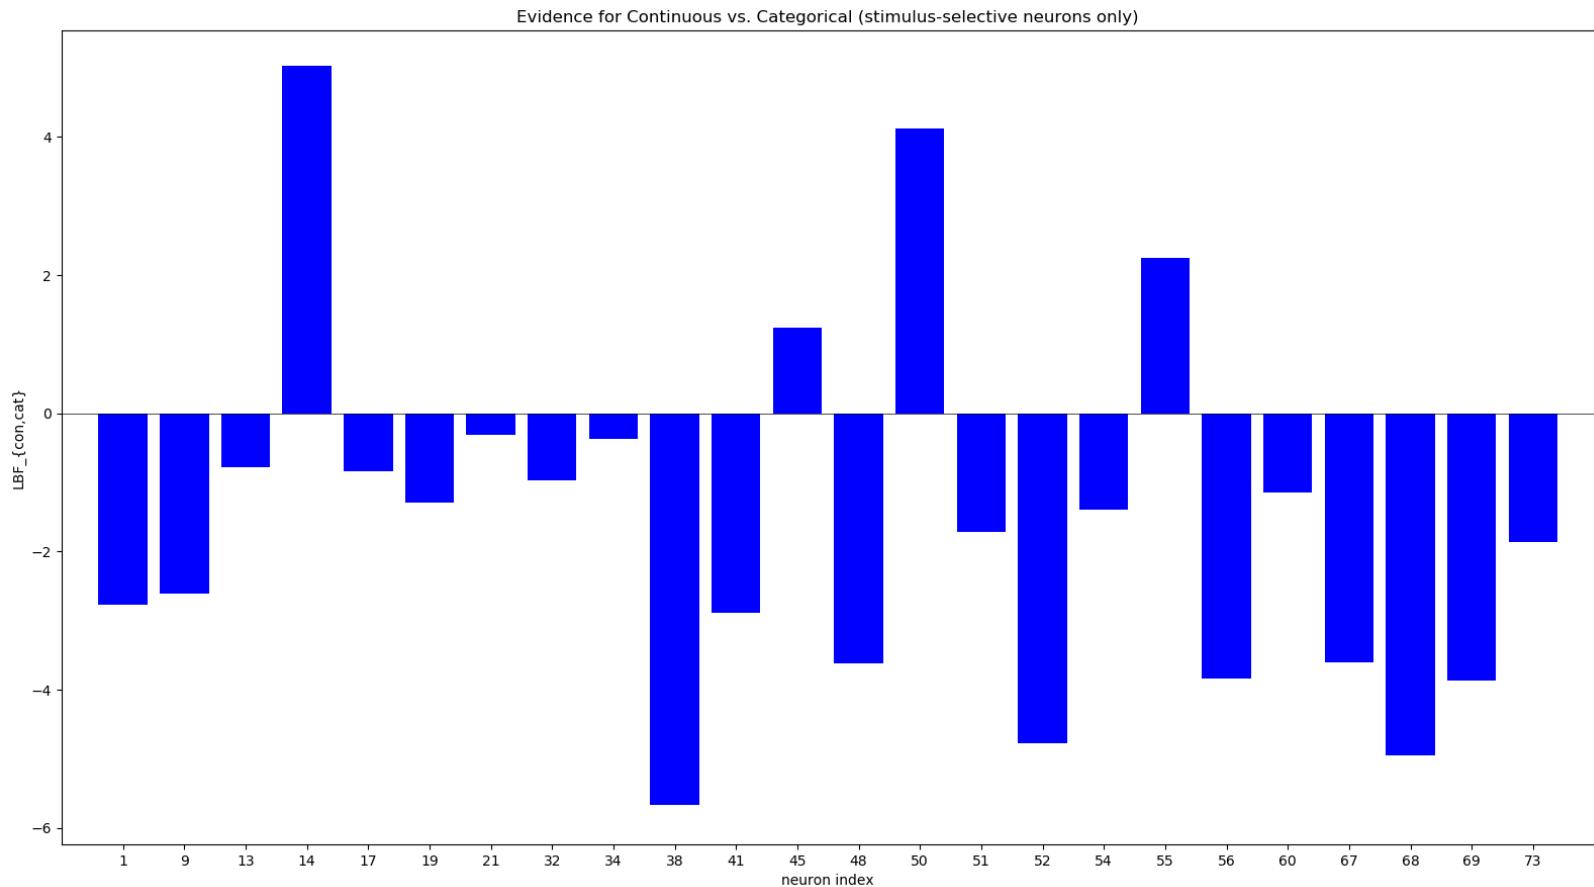
Q3: model space



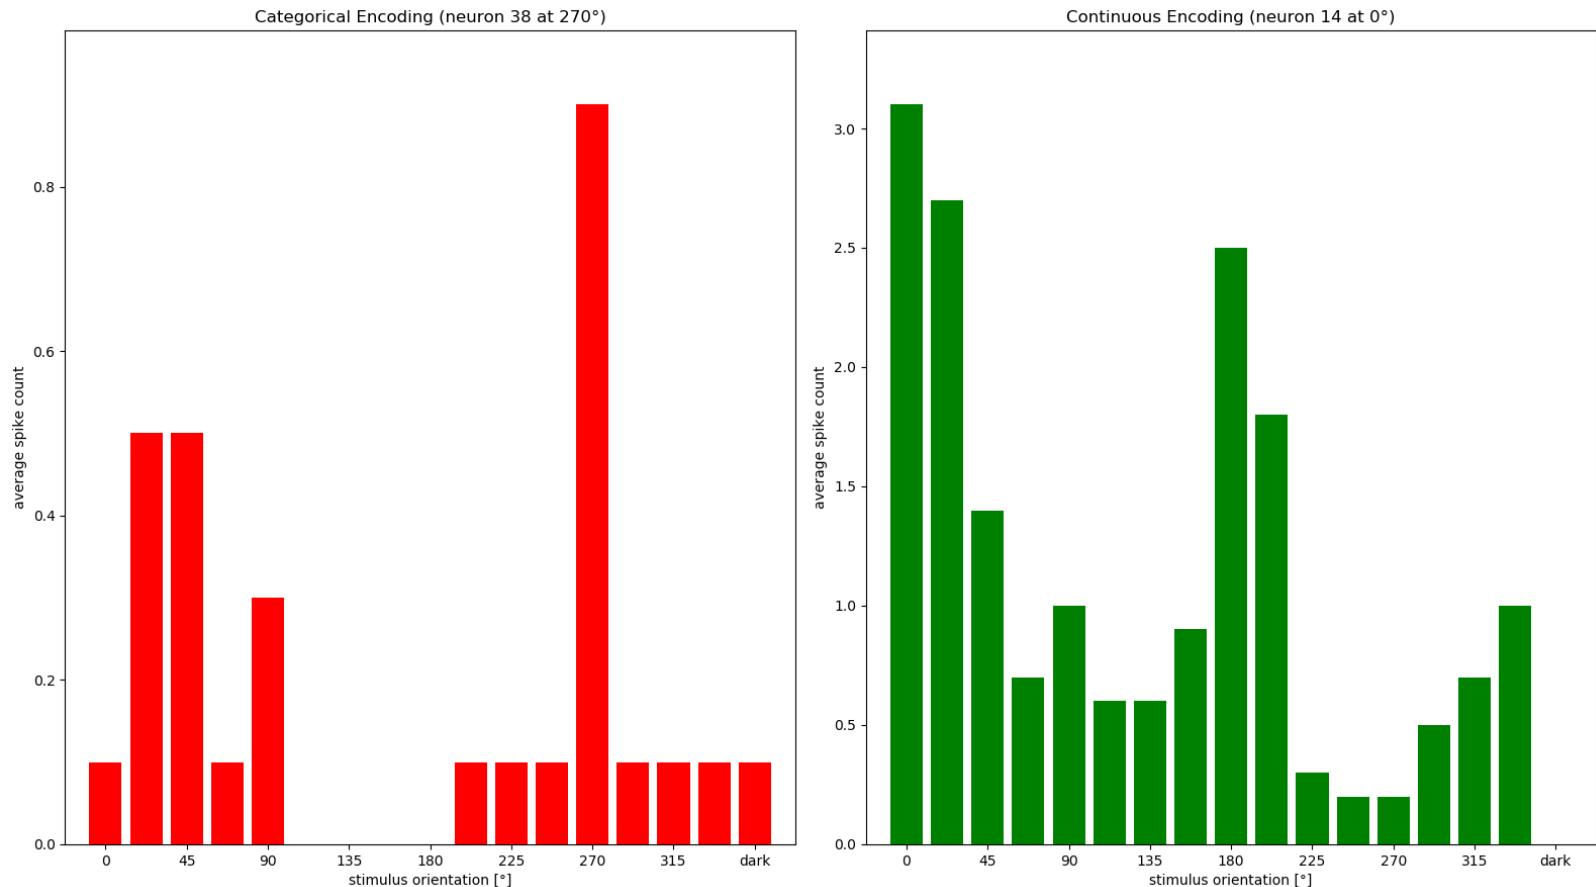
Q3: model selection tool

$$\text{LBF}_{12} = \log \text{BF}_{12} = \log \frac{p(y|m_1)}{p(y|m_2)} = \text{LME}(m_1) - \text{LME}(m_2)$$

Q3: model comparison



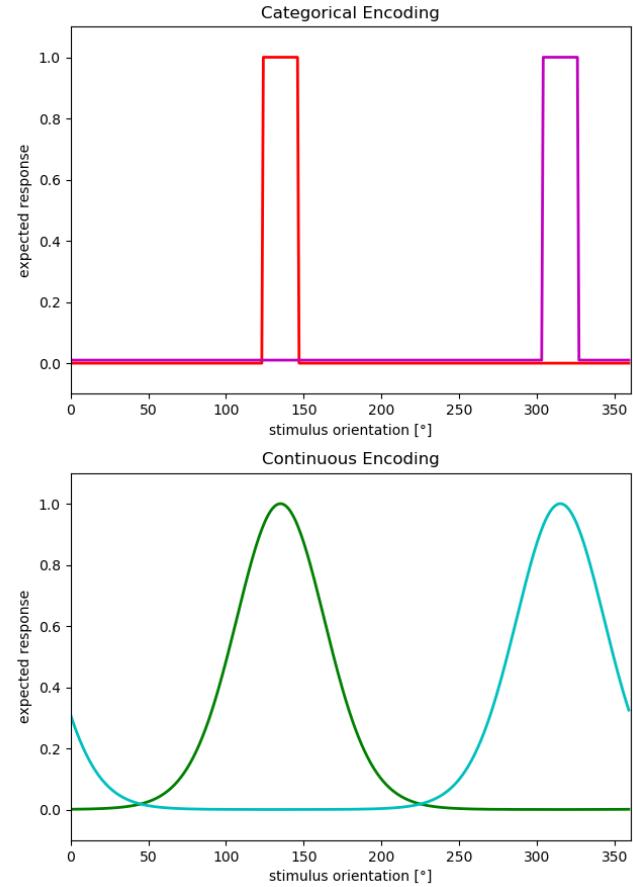
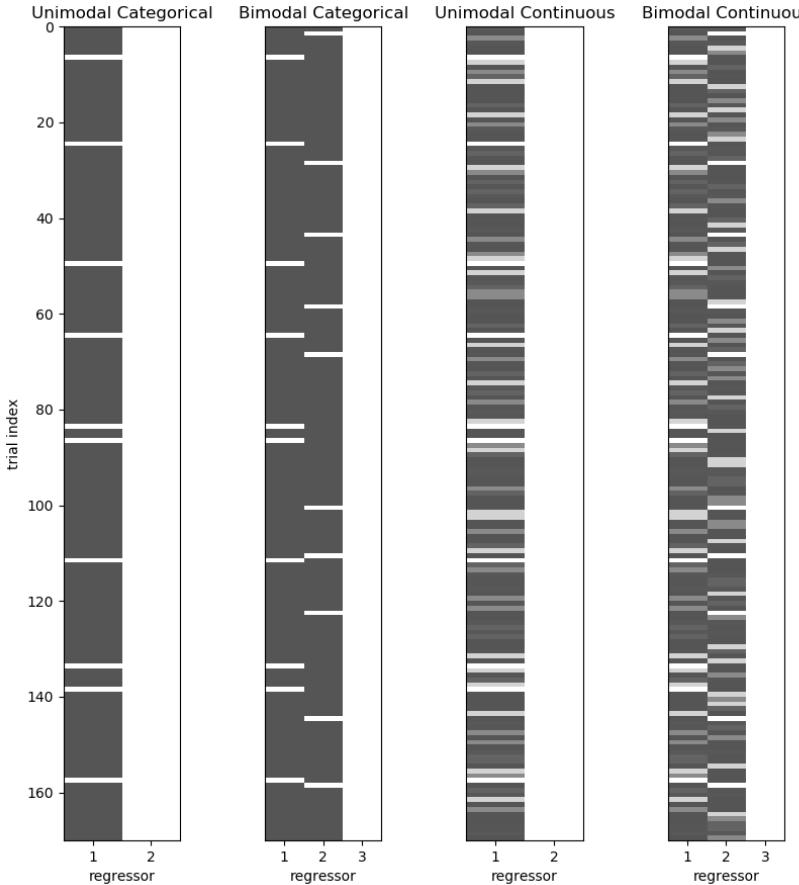
Q3: data inspection



Research Question No. 4

„Does the neuron also respond to the exact opposite of the preferred direction?“

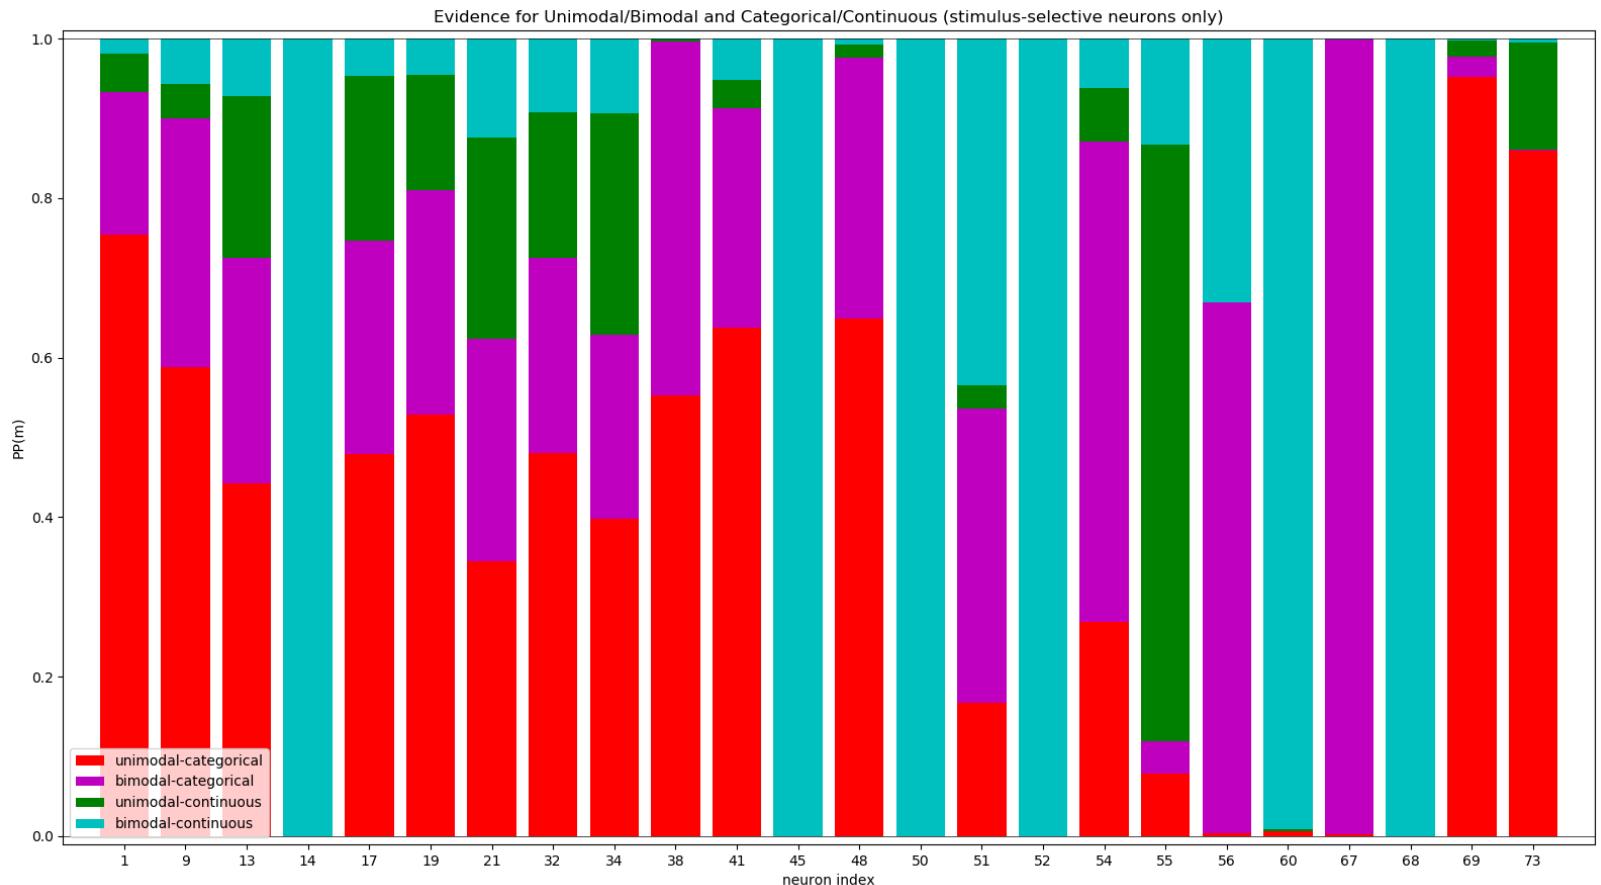
Q4: model space



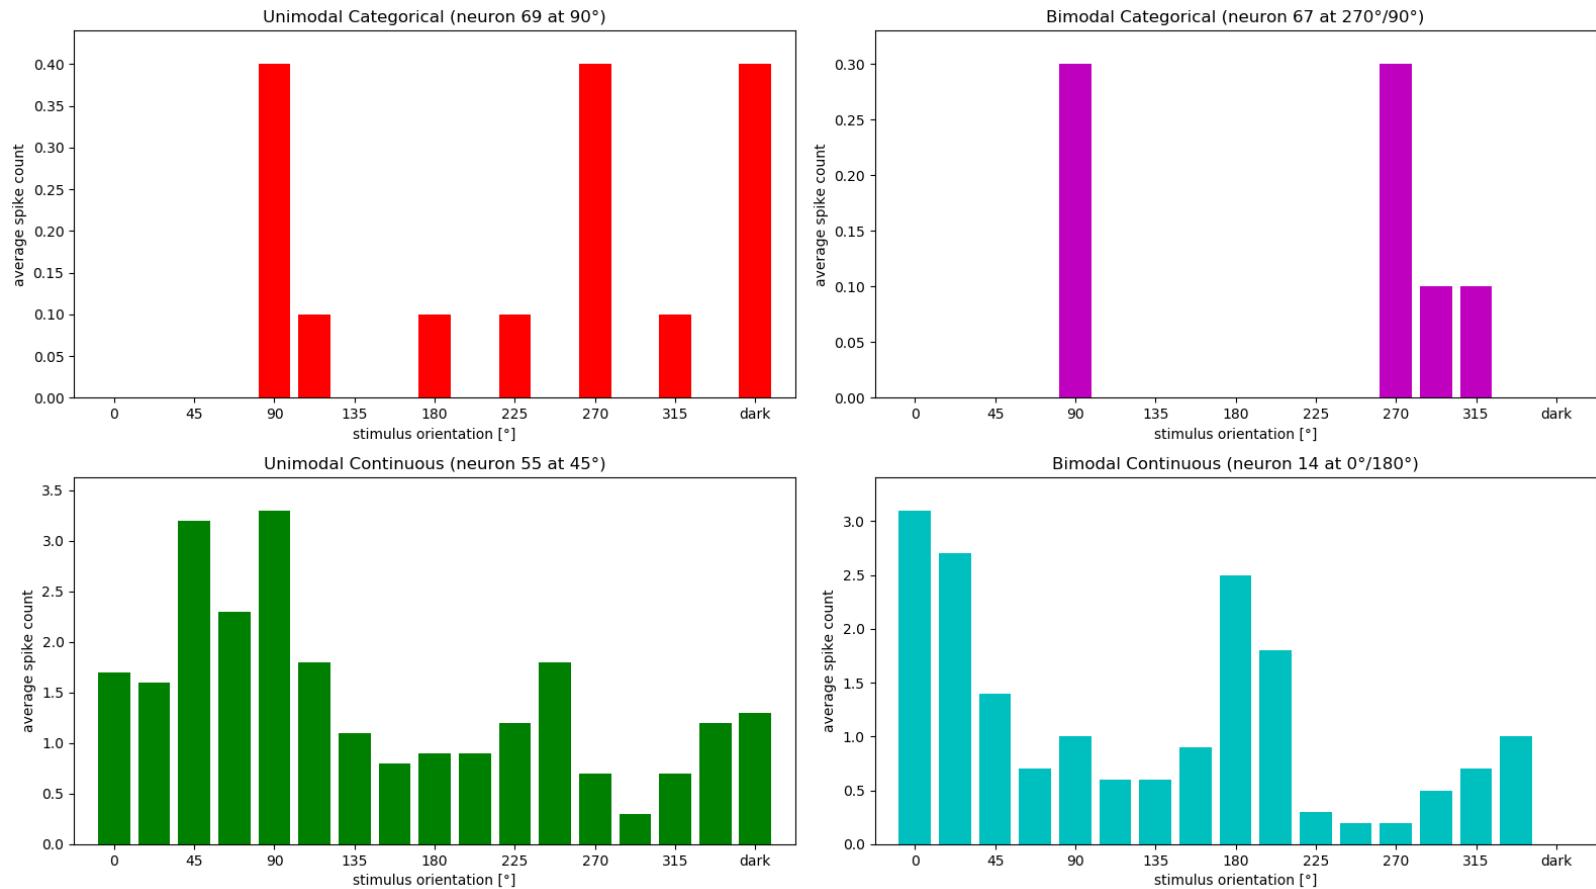
Q4: model selection tool

$$p(m_i|y) = \frac{p(y|m_i) p(m_i)}{\sum_{j=1}^M p(y|m_j) p(m_j)} = \frac{\exp[\text{cvLME}(m_i)] p(m_i)}{\sum_{j=1}^M \exp[\text{cvLME}(m_j)] p(m_j)}$$

Q4: model selection



Q4: data inspection



Possible Further Research Questions

„Does the neuron have a baseline response
or is it zero (i.e. no constant regressor)?“

„Are discrete spike counts normally distributed or
are they rather Poisson-distributed?“ (probably yes)

8.

What do I need to perform (Bayesian) model selection?

It depends! ;-)

If you have ...

... an objective function

$$\hat{\theta} = \arg \max f(y, \theta, m)$$

... a likelihood function

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

... plus a prior distribution

$$p(\theta|m)$$

... plus an analytic posterior

$$p(\theta|y, m) \propto p(y|\theta, m) p(\theta|m)$$

then you can use ...

... any measure of fit

$$\text{GoF} = g(y, \hat{y})$$

... classical information criteria

$$\text{IC} = -2 \text{MLL} + f_{\text{IC}}(n, k)$$

... the log model evidence

$$\text{LME}(m) = \log p(y|m)$$

... the cross-validated log model evidence

$$\text{cvLME}(m) = \sum_{i=1}^S \log p(y_i|y_{-i}, m)$$

9.

Practice Q&A





John-Dylan Haynes

Carsten Allefeld



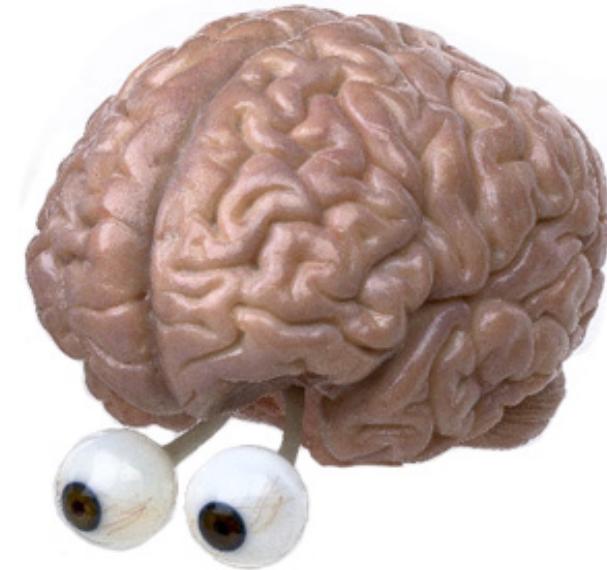
Carsten Bogler

Achim Meyer

THANKS TO MY COLLABORATORS!

**THANK YOU!
QUESTIONS?**

joram.soch@bccn-berlin.de



Appendix

Model quality control for fMRI

	cvBMS	cvBMA	MACS
<i>Poster</i>	<u>F1000 Research</u>	<u>OHBM 2017</u>	<u>OHBM 2018</u>
<i>Paper</i>	<u>NeuroImage</u>	<u>NeuroImage</u>	<u>JNeuMeth</u>
<i>Code</i>	obsolete	obsolete	<u>GitHub</u> <u>Demo</u>

Derivatives of the LME

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

$$p(m|y) = \frac{p(y|m) p(m)}{\sum_{j=1}^M p(y|m_j) p(m_j)}$$
LME(m) = $\log p(y|m)$

$$p(f|y) = \frac{p(y|f) p(f)}{\sum_{k=1}^F p(y|f_k) p(f_k)}$$

$$p(y|f) = \sum_{m \in f} p(y|m) p(m|f)$$
LFE(f) = $\log p(y|f)$

$$\text{BF}_{12} = \frac{p(y|m_1)}{p(y|m_2)}$$
LBF₁₂ = LME(m_1) - LME(m_2)

$$p(m|Y) = \frac{p(Y|m) p(m)}{\sum_{j=1}^M p(Y|m_j) p(m_j)}$$

$$p(Y|m) = \prod_{i=1}^N p(y_i|m)$$

$$\log p(Y|m) = \sum_{i=1}^N \log p(y_i|m)$$

$$\text{GBF}_{12} = \frac{p(Y|m_1)}{p(Y|m_2)}$$
LGBF₁₂ = $\log p(Y|m_1) - \log p(Y|m_2)$

Accuracy & Complexity in the LME

$$p(\theta|y, m) = \frac{p(y|\theta, m) p(\theta|m)}{p(y|m)}$$

$$p(y|m) = \frac{p(y|\theta, m) p(\theta|m)}{p(\theta|y, m)}$$

$$\log p(y|m) = \log p(y|\theta, m) - \log \frac{p(\theta|y, m)}{p(\theta|m)}$$

$$\log p(y|m) = \int p(\theta|y, m) \log p(y|\theta, m) d\theta - \int p(\theta|y, m) \log \frac{p(\theta|y, m)}{p(\theta|m)} d\theta$$

$$\text{LME} = \langle \log p(y|\theta, m) \rangle_{p(\theta|y, m)} - \text{KL} [p(\theta|y, m) || p(\theta|m)]$$

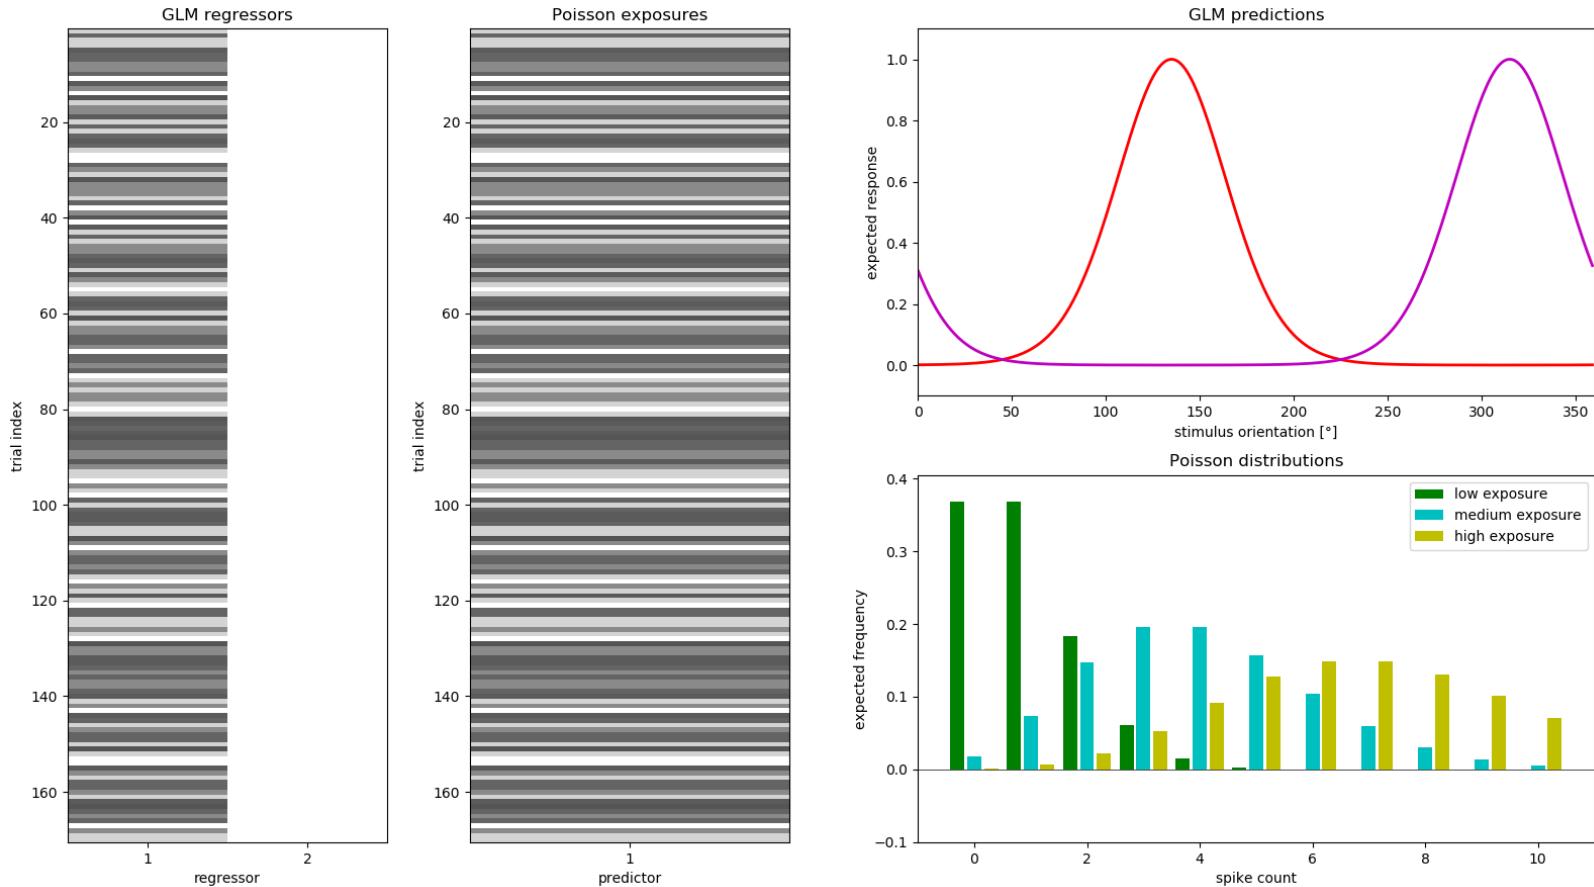
The pragmatics of model selection

- model dilution
 - model families
- sequential model selection
 - assume independence
- methodological control (Q3) vs.
scientific inference (Q1, Q2, Q4)

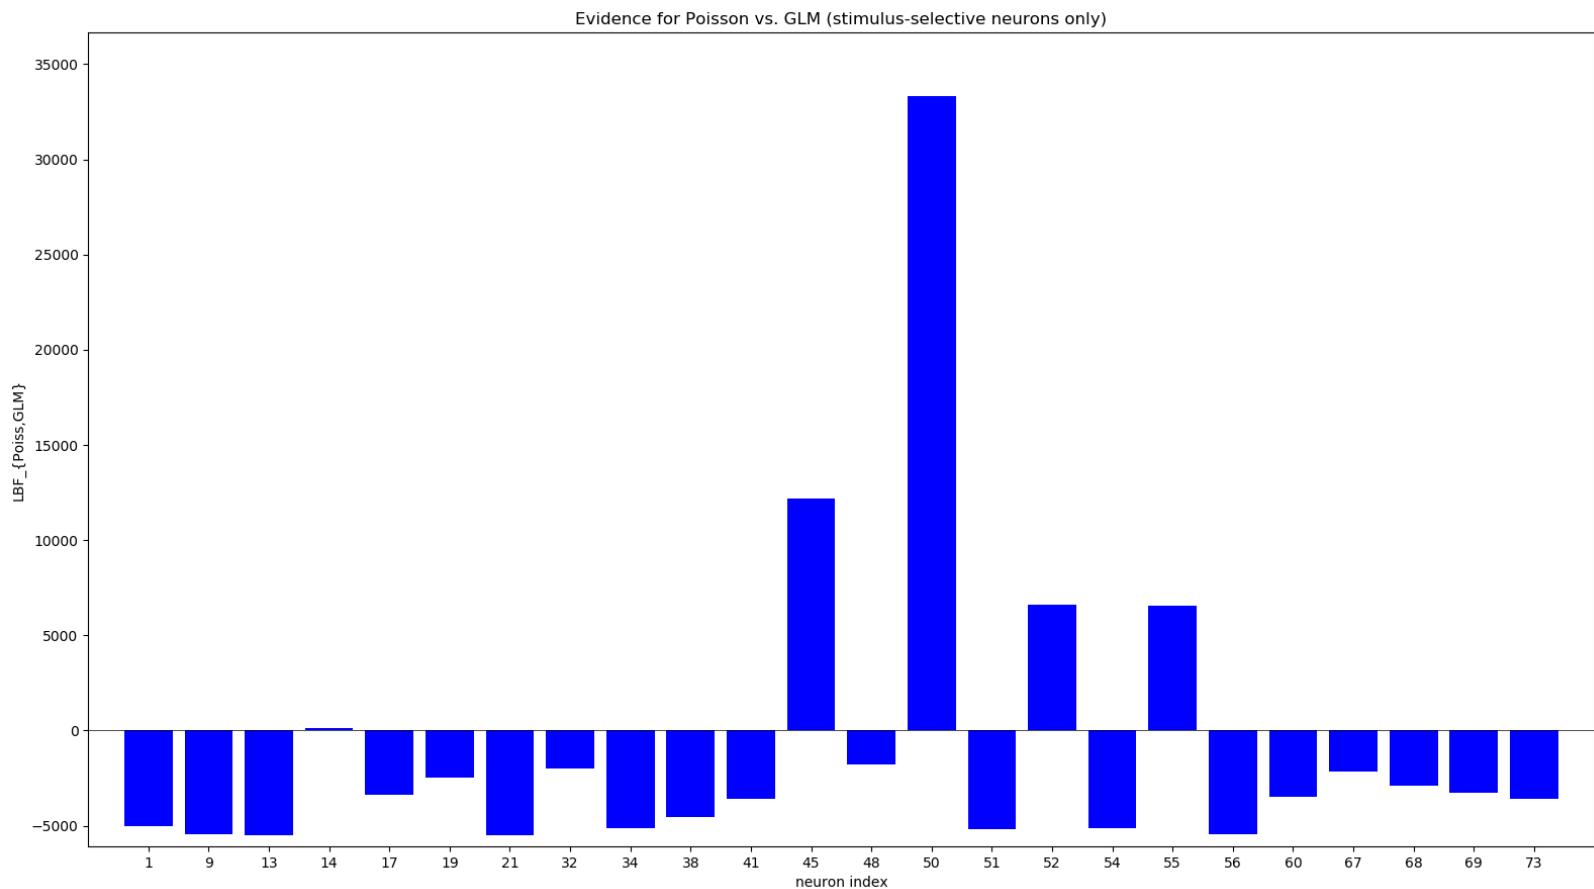
Research Question No. 5

„Are spike counts normally distributed
or are they rather Poisson-distributed?“

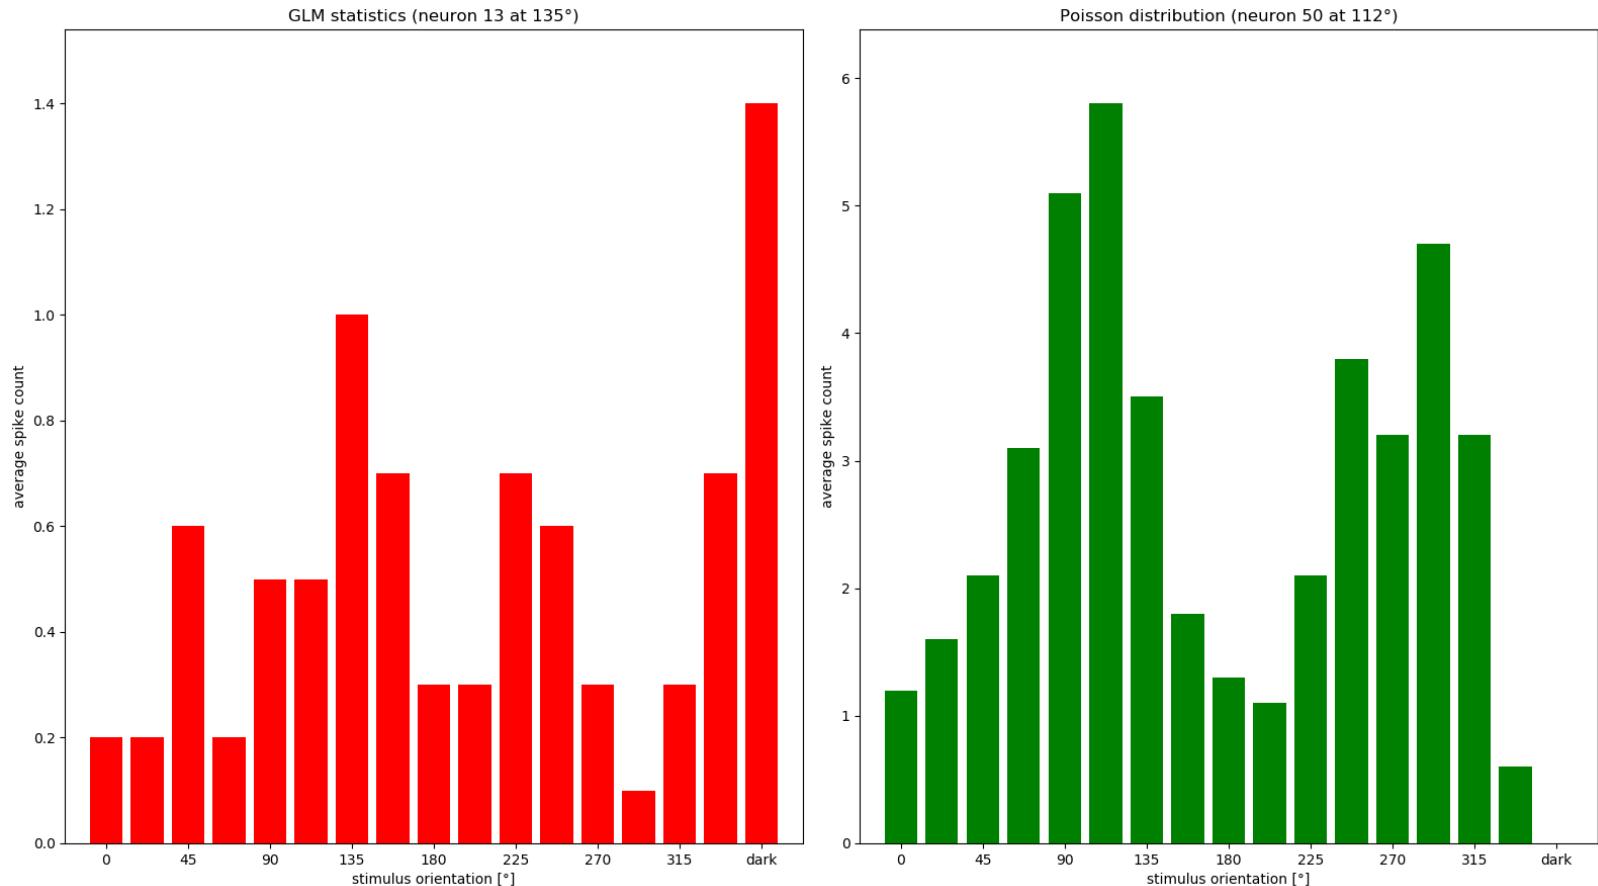
Q5: model space



Q5: model comparison



Q5: data inspection



Q5: data inspection

