What makes a good multivariate model for fMRI-based decoding ?

Bertrand Thirion,
Inria, CEA, Université Paris-Saclay
http://parietal.saclay.inria.fr
bertrand.thirion@inria.fr

Brain activity *decoding*

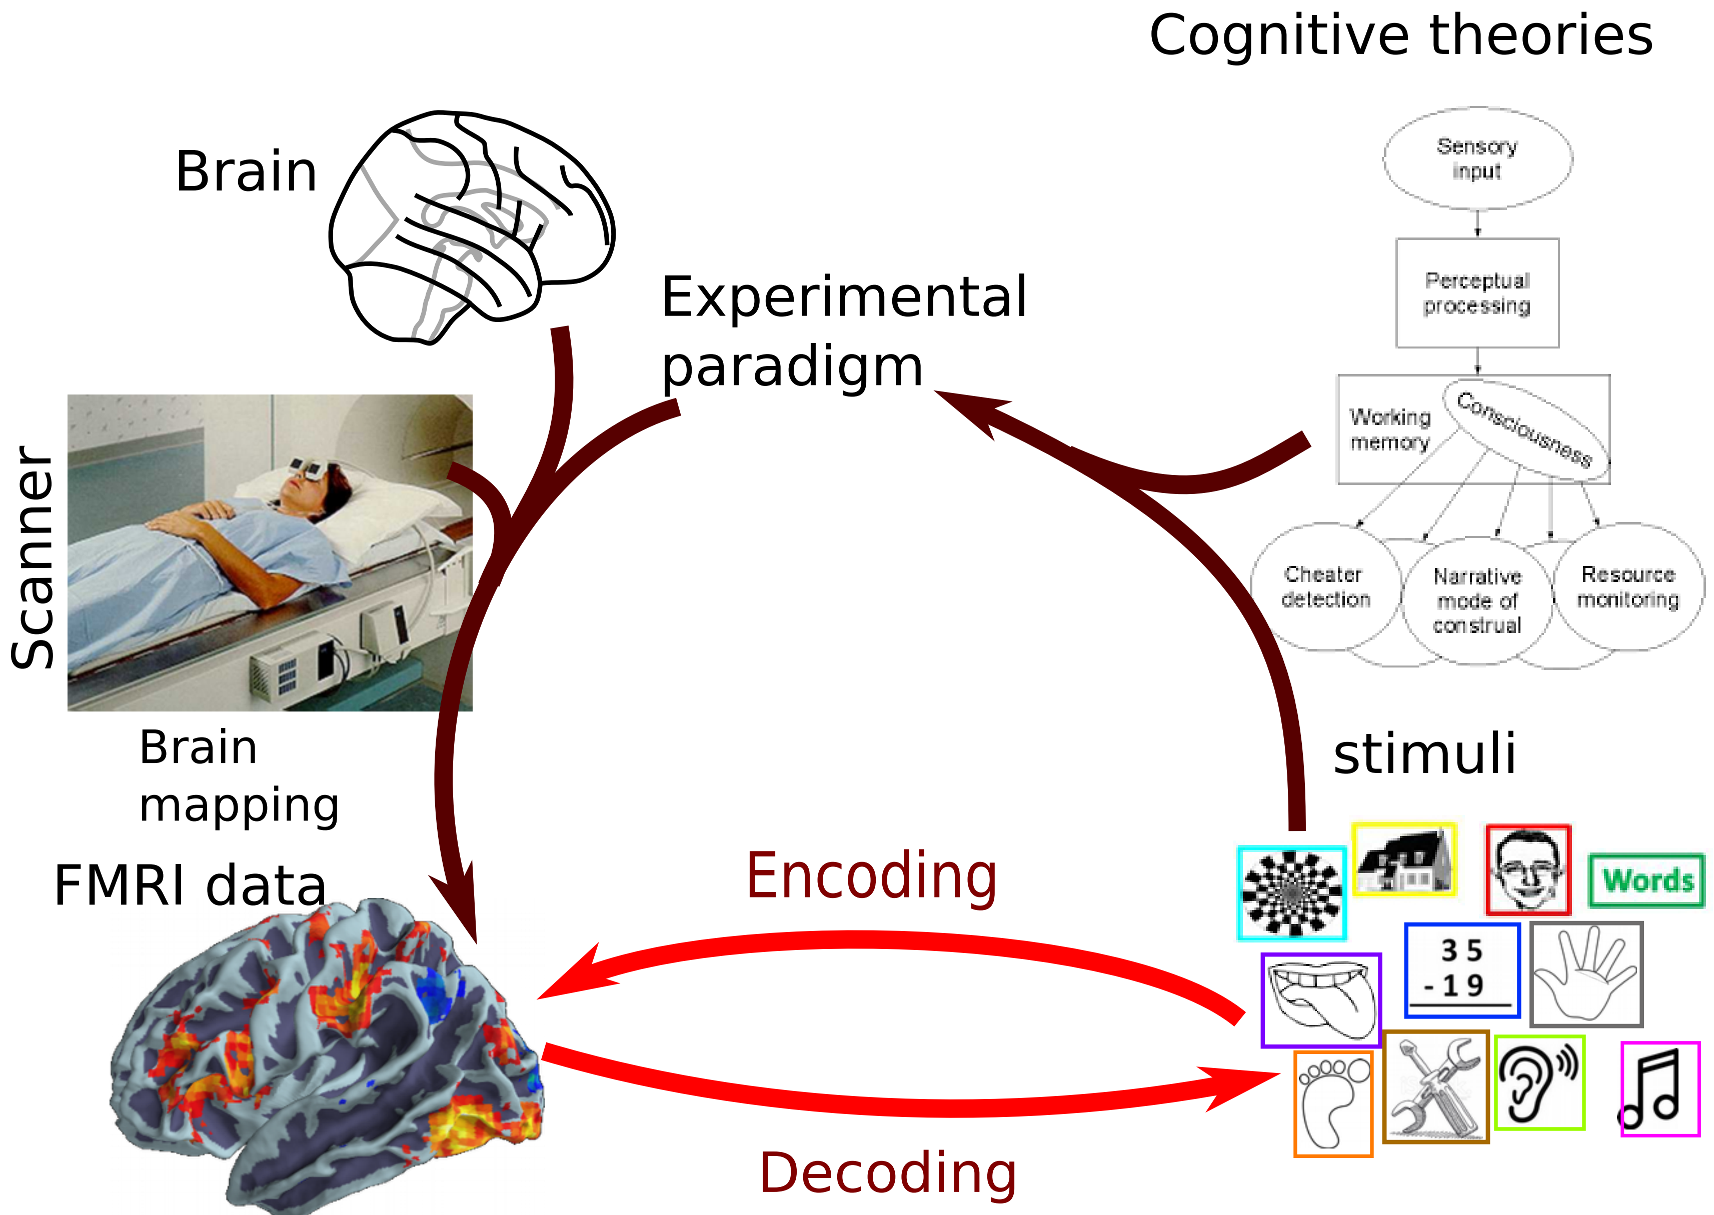

Outline

- Keep your decoding model simple
- Beware accuracy and significance
- Inference about decoding model parameters

Outline

- Keep your decoding model simple
- Beware accuracy and significance
- Inference about decoding model parameters

Keep your MVPA model simple

Can you explain what your model does to your grandmother ?

Can you explain her the validation metric you use ?

What is a simple model ?

- Linear models are simple
- Classification trees are somewhat simple
- Kernel methods not so simple
- Deep learning not simple at all

Complexity in the procedure:

- Feature selection
- Parameter tuning

Simplicity= few steps, few parameters to tune

Parameters tuning is hard

CV accuracy is overoptimistic

Unbiased accuracy By **nested cross-validation**

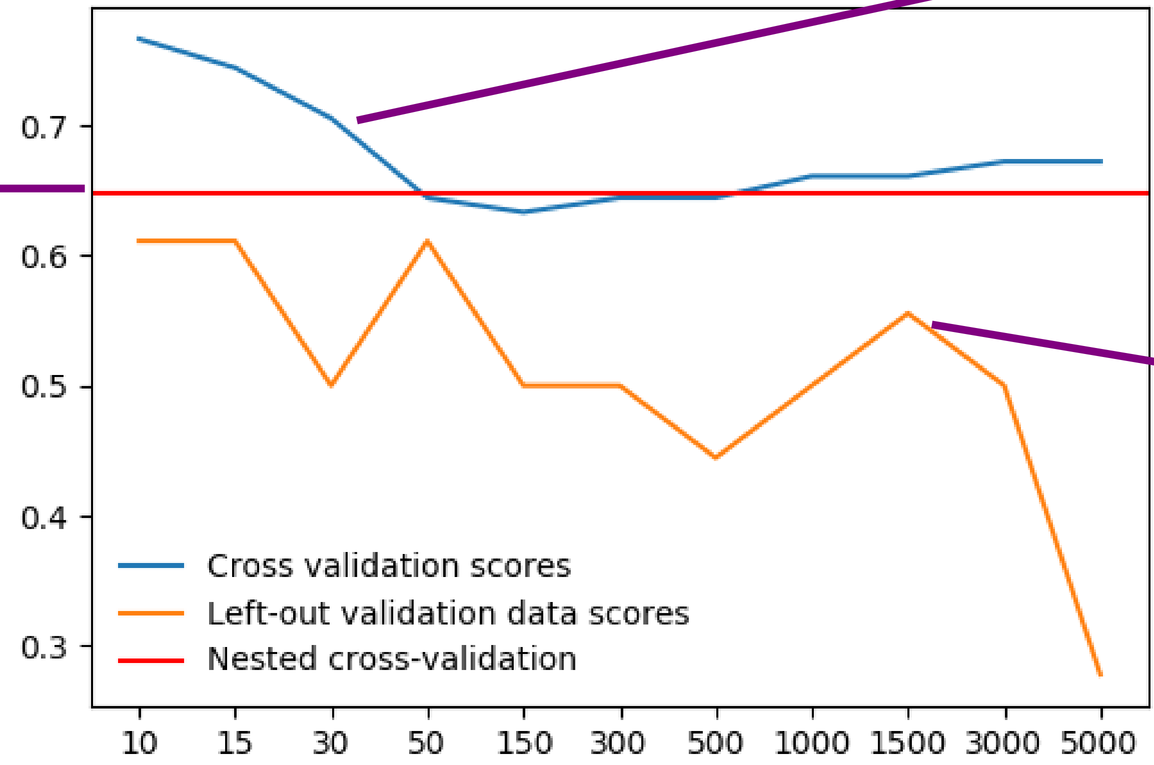

And misleading: generalizes poorly

Tuning a parameter: # selected voxels

http://nilearn.github.io/auto_examples/02_decoding/plot_haxby_grid_search.html

What is a simple model ?

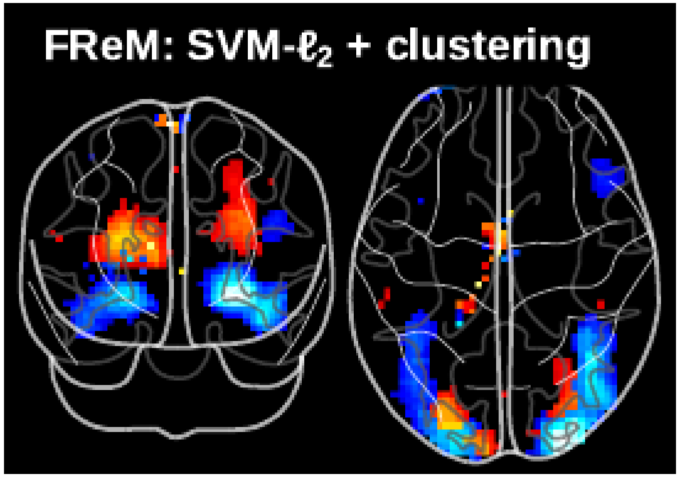

Rather than

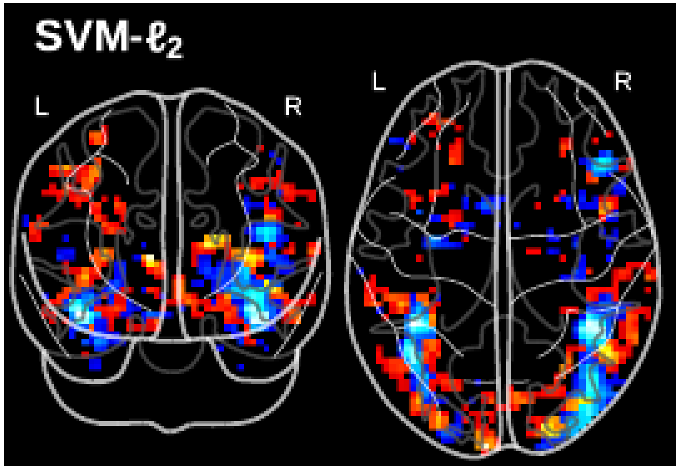

Both model equally good for discrimination,
but the first one is “simpler” and more stable
[Hoyos-Idrobo et al. NeuroImage 2017]

Lots of good news

- Linear models work well for neuroimaging data

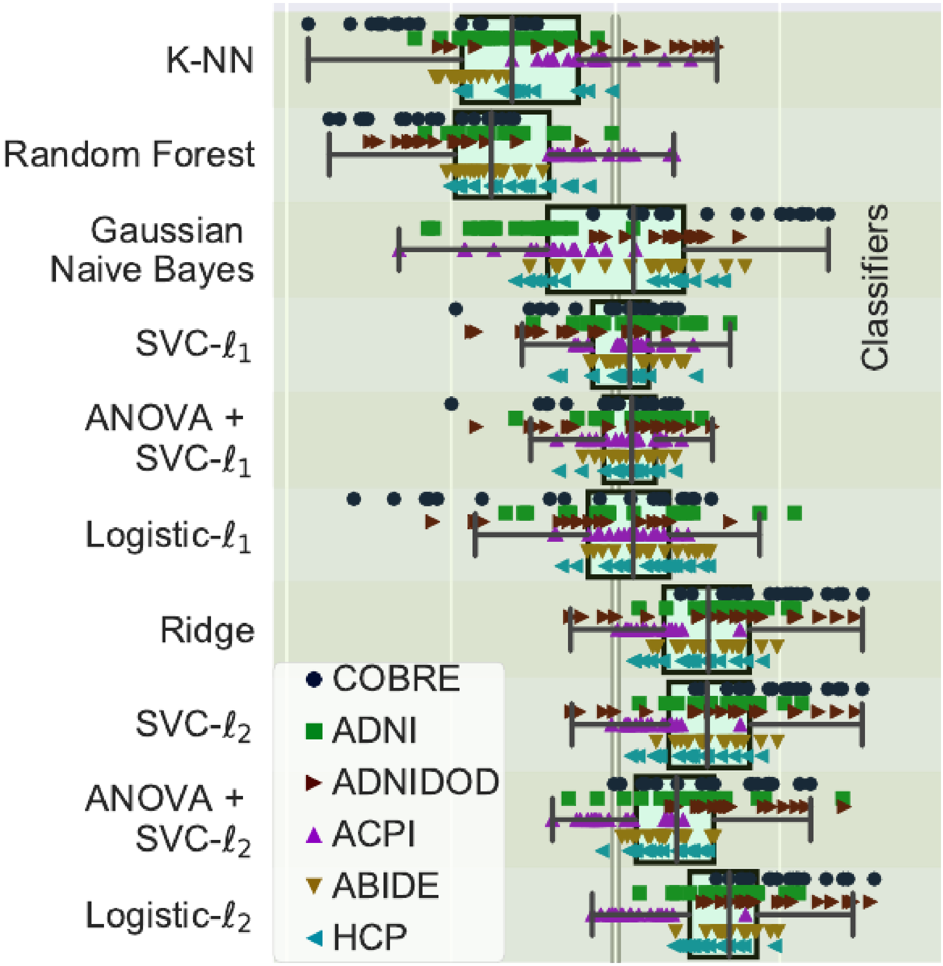

Example on connectome-based prediction [Dadi et al, NeuroImage 2019]: good old logistic regression outperforms alternatives

See also:
[He et al. NeuroImage 2019]

Lots of good news

- Linear models work well for neuroimaging data
- Standard libraries make machine learning plug-and-play

```
>>> from sklearn.datasets import load_diabetes
>>> from sklearn.linear_model import RidgeCV
>>> X, y = load_diabetes(return_X_y=True)
>>> clf = RidgeCV(alphas=[1e-3, 1e-2, 1e-1, 1]).fit(X, y)
>>> clf.score(X, y)
0.5166...
```

Scikit-learn example: parameter selection for Ridge regression

Lots of good news

- Linear models work well for neuroimaging data
- Standard libraries make machine learning plug-and-play
- High-level dedicated objects make it easier

Model specification

Provide data

```
decoder = Decoder(estimator='svc_l1', cv=cv,
                  mask=masker, scoring='roc_auc')
decoder.fit(task_data, classification_target, groups=session_labels)
mask_scores[mask_name][category] = decoder.cv_scores_[1]
```

obtain score

Now in nilearn.github.io

Caveat: linear decodability

- Does successful **decoding** w. *linear* classifier imply that brain activity **encodes** stimulus information *linearly* ?
- **No**
- Counter-example: position ((x, y) or (r, θ) coordinates) of an object in the visual field
  - not encoded linearly
  - can be decoded linearly

Caveat: linear decodability

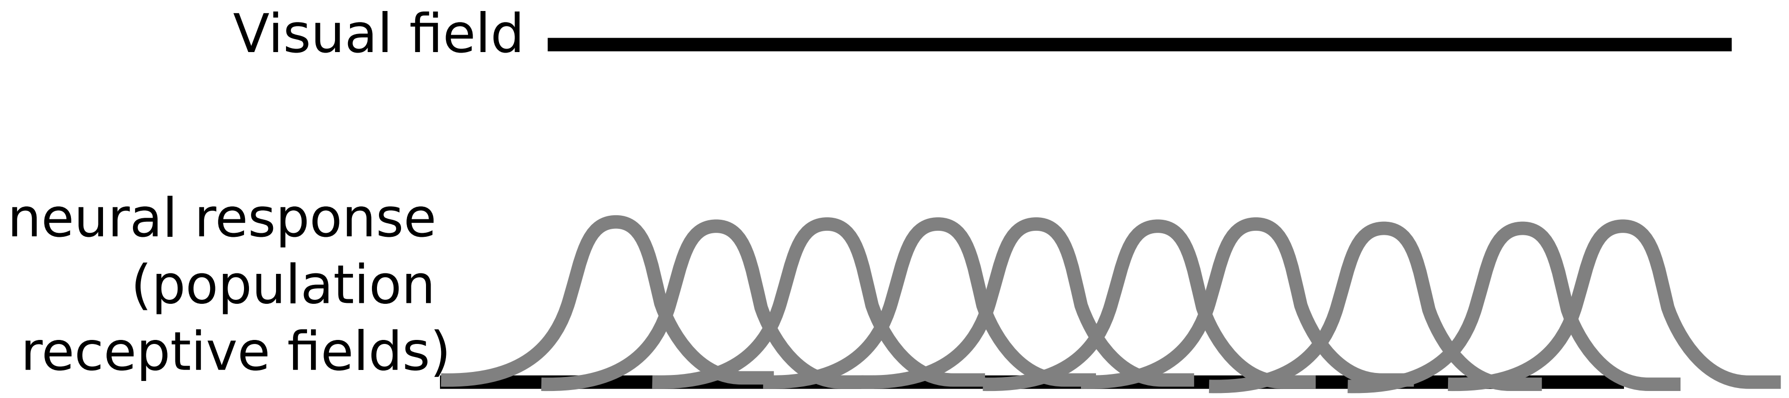

Caveat: linear decodability

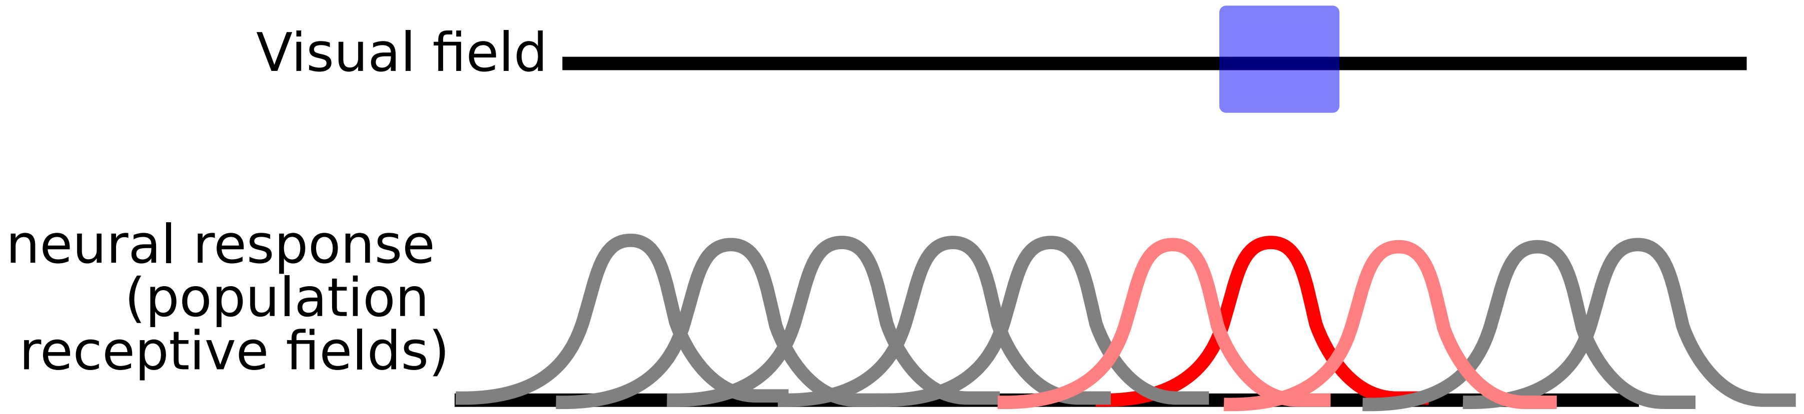

Caveat: linear decodability

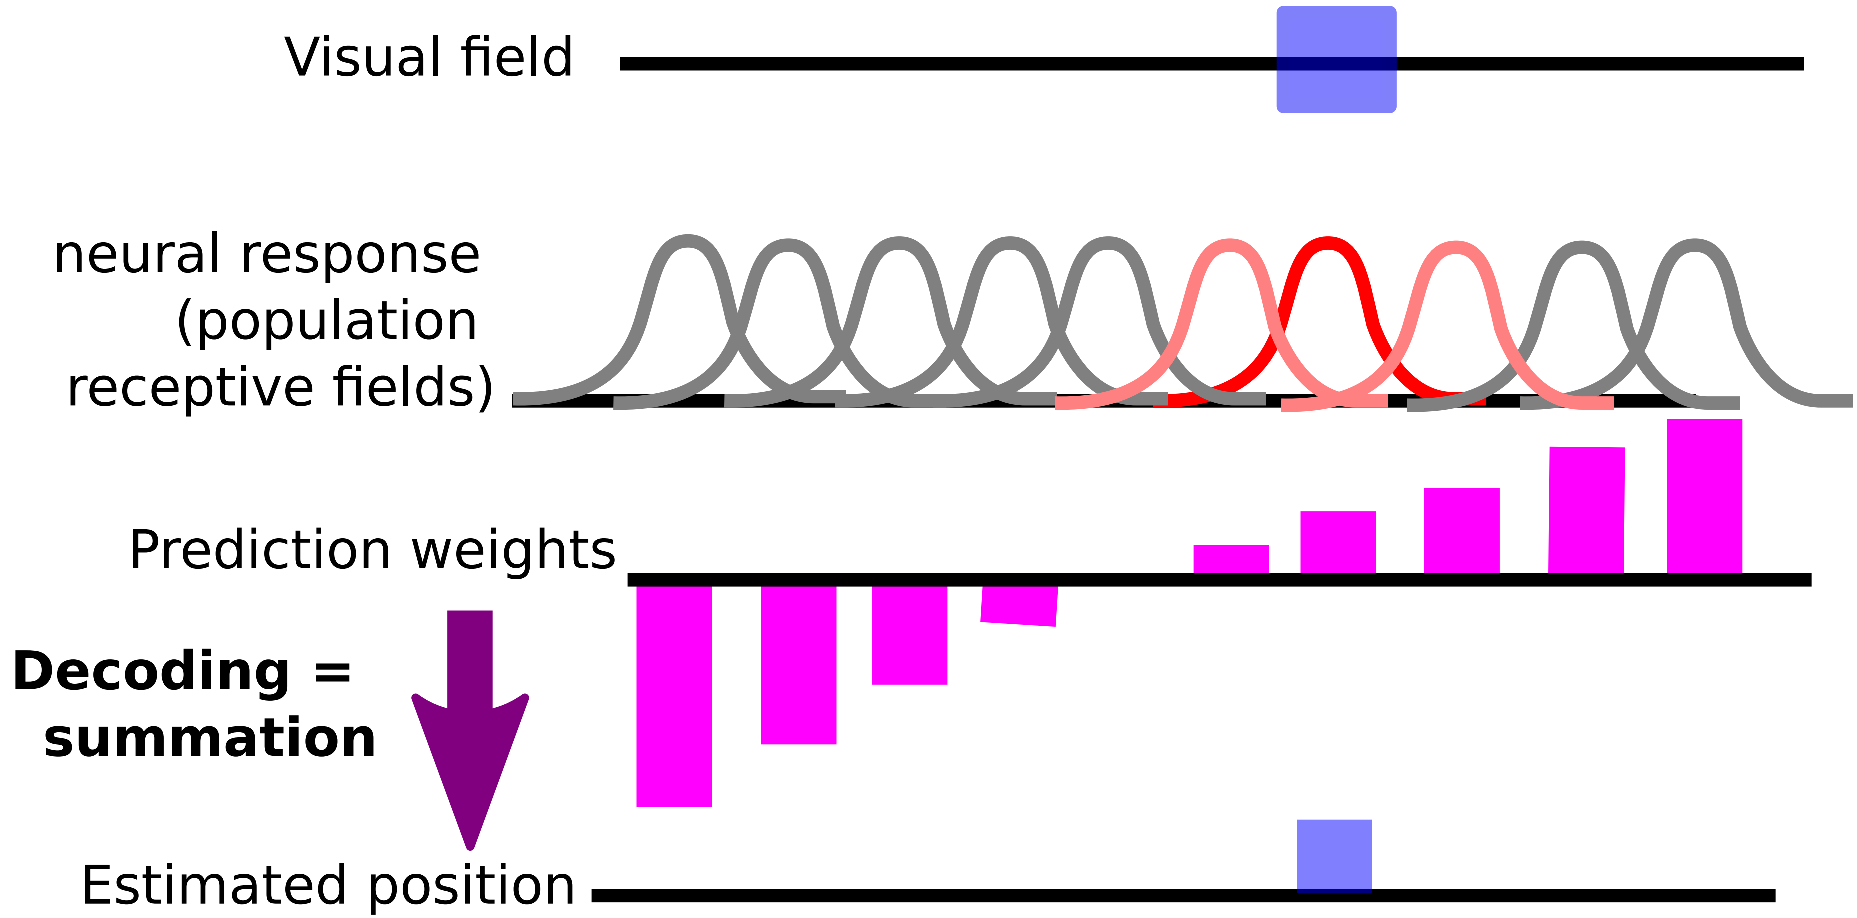

[Thirion et al. Neuroimage 2006]

A note on linear decodability

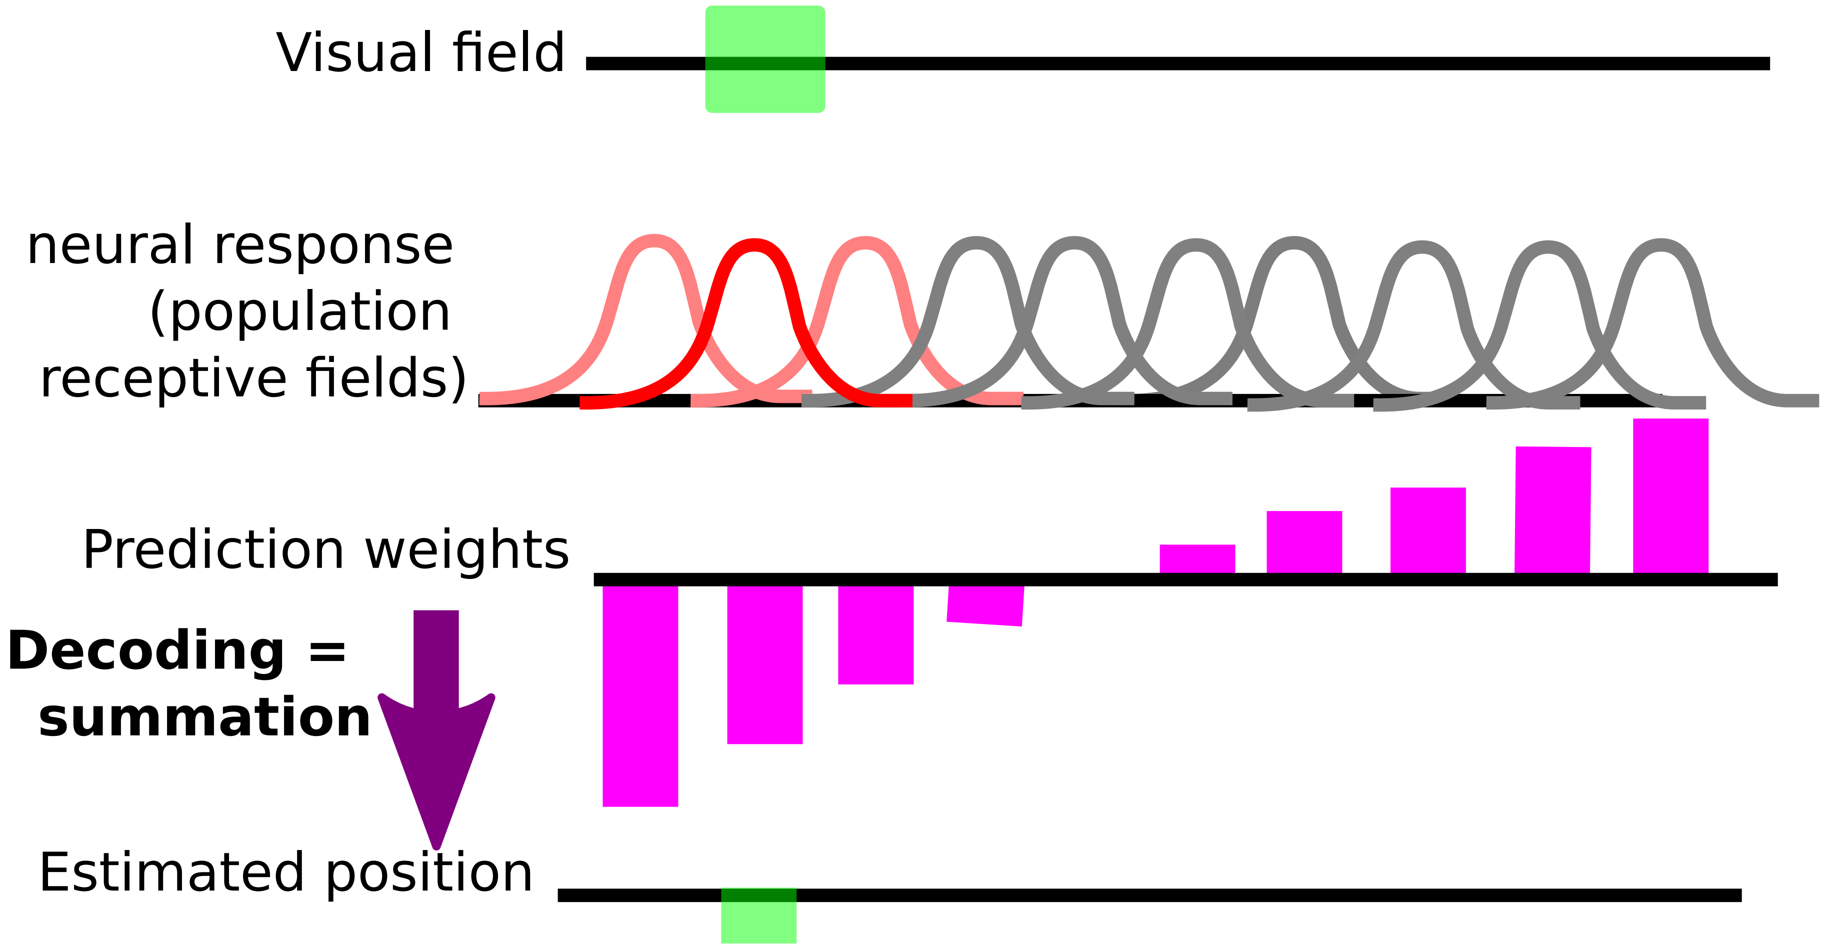

[Thirion et al. Neuroimage 2006]

See also:

https://nilearn.github.io/auto_examples/02_decoding/plot_miyawaki_reconstruction.html

Outline

- Keep your decoding model simple
- Beware accuracy and significance
- Inference about decoding model parameters

Need high accuracy

```
print("Classification accuracy: %.4f / Chance level: %f" %
      (classification_accuracy, 1. / len(conditions.unique())))
# Classification accuracy:  0.70370 / Chance level: 0.5000
```

Decoding doesn't work with 55% accuracy:

- Error bars are large (see next slide)
- Model may have picked any confound that helps predicting
- Statistical significance vs practical significance

Sample size and cross-validation

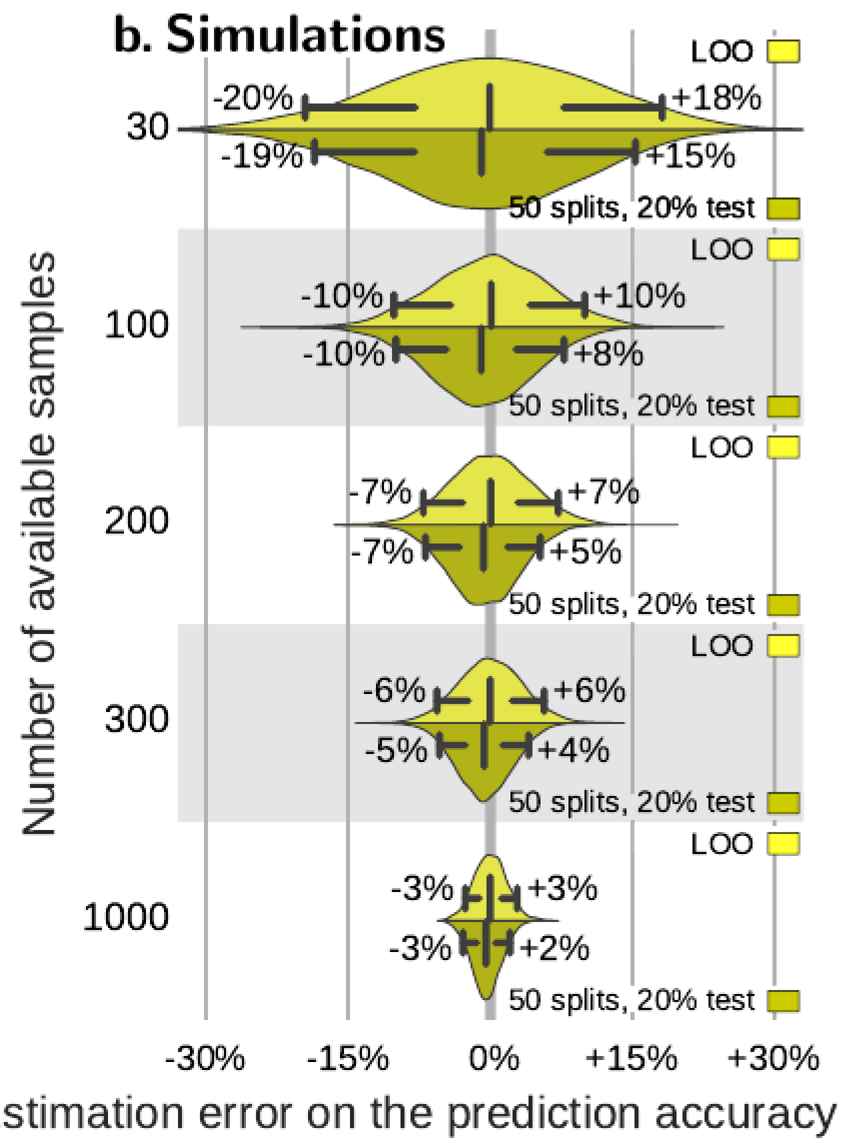

Rule of the thumb: uncertainty in prediction decreases with 1/ √n

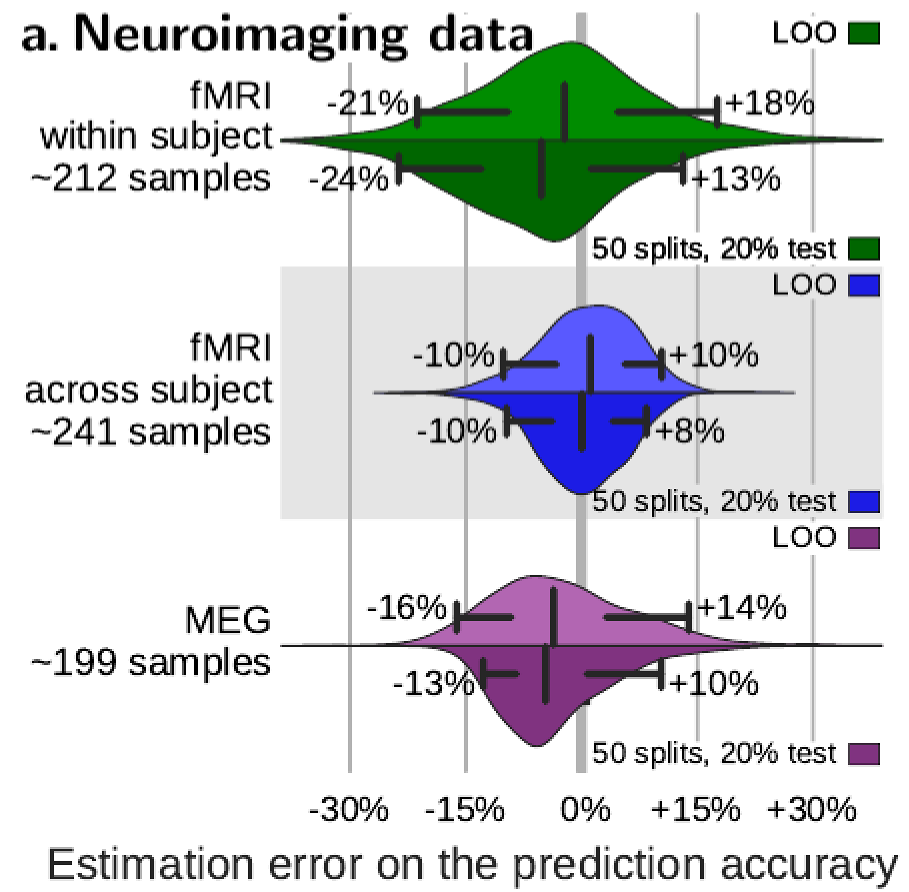

[Varoquaux et al NeuroImage 2017]

Learning curve: how prediction improves with n

- Predict the age of a subject given gray matter density maps (OASIS dataset, n=403)

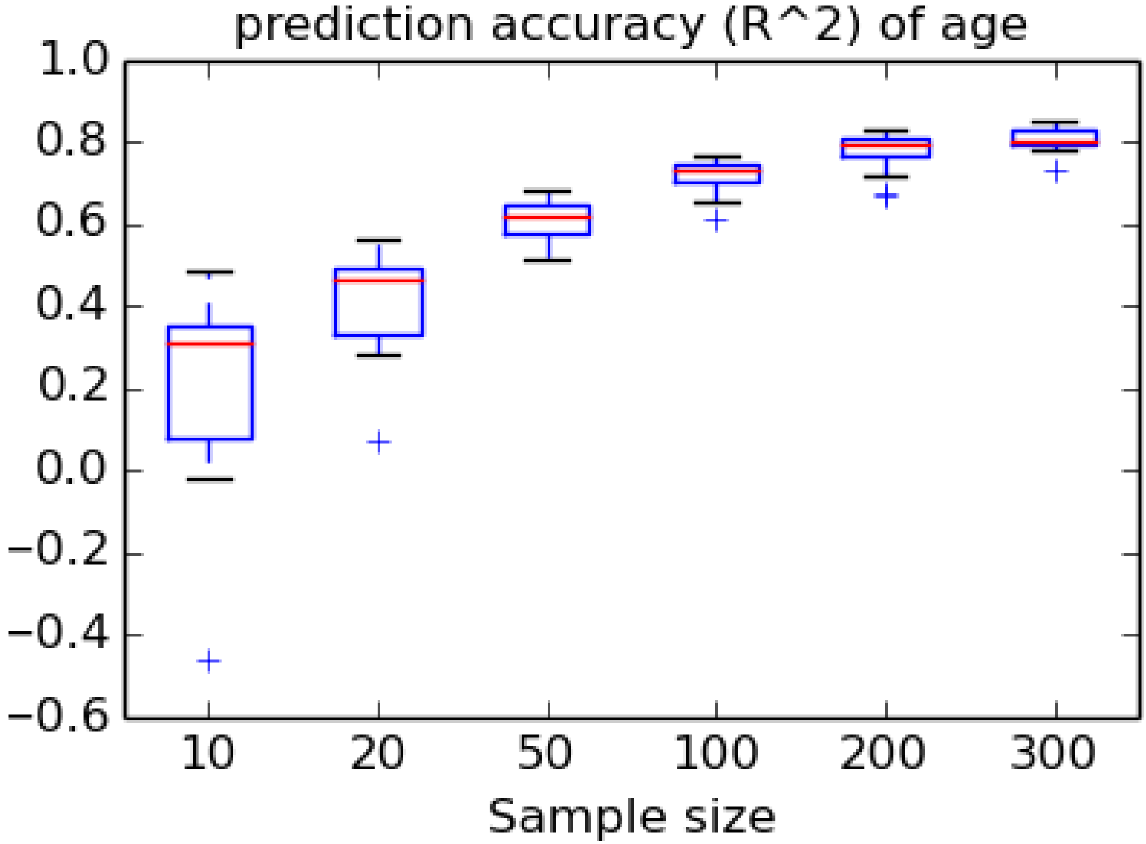

Weight maps for age prediction / OASIS

The weight map depends on the batch of subject considered (bootstrap):

One question, different datasets, different answers

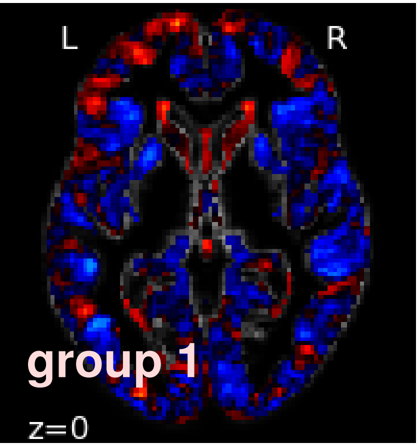

group 6

Variability actually worse than for univariate analysis !

Weight maps for age prediction / OASIS

The weight map depends on the batch of subject considered (bootstrap):

One question, different datasets, different answers

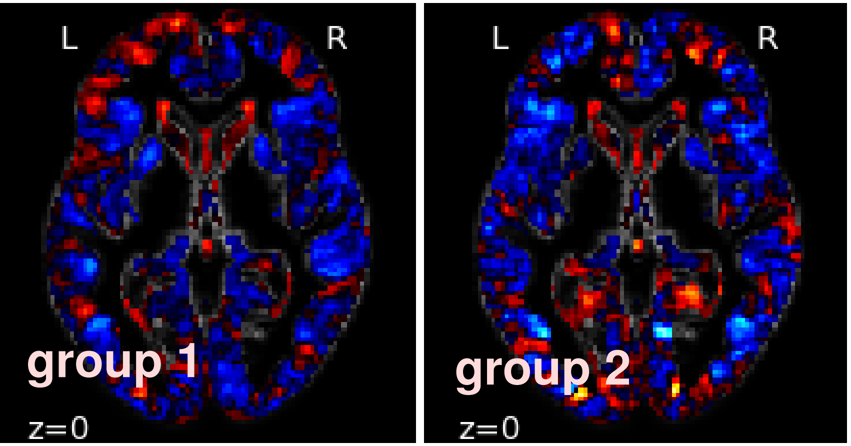

Weight maps for age prediction / OASIS

The weight map depends on the batch of subject considered (bootstrap):

One question, different datasets, different answers

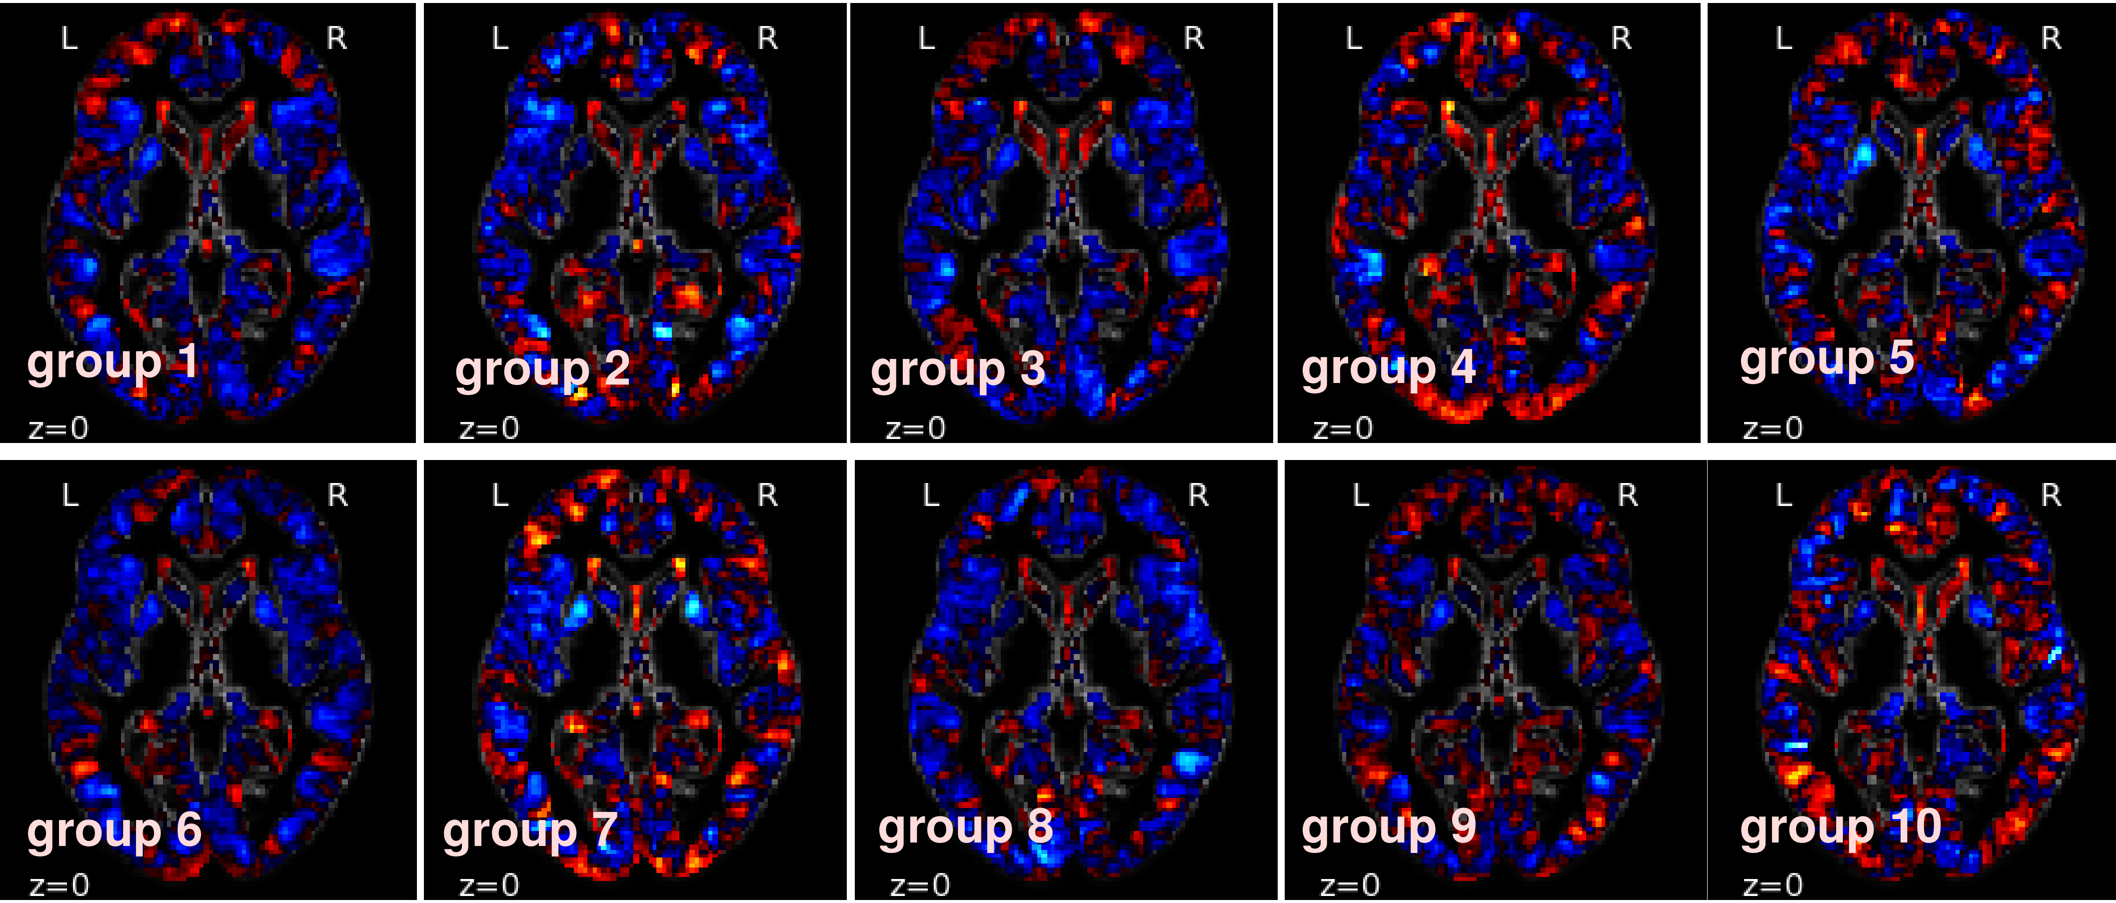

Variability actually than than univariate analysis

Weight maps for age prediction / OASIS

The weight map depends on the batch of subject considered (bootstrap):

One question, different dataset, different answers

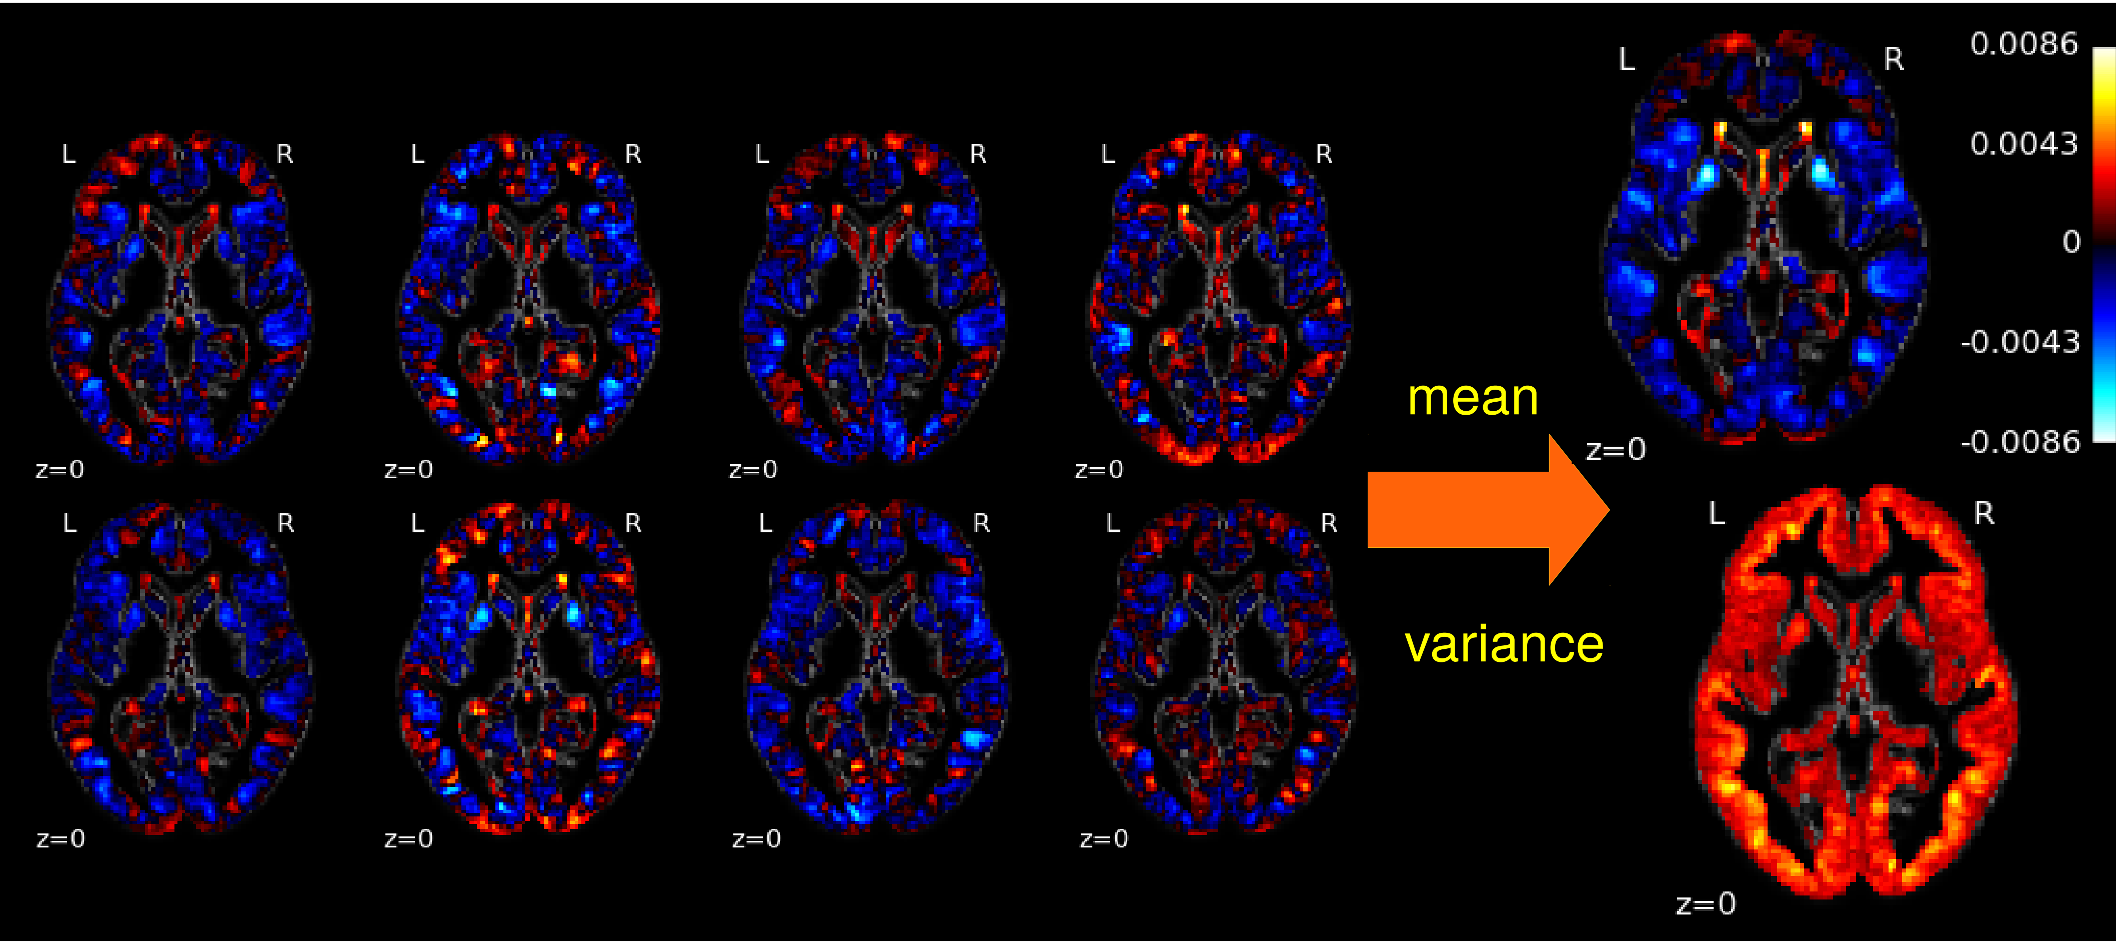

Weight maps for age prediction / OASIS

The weight map depends on the batch of subject considered (bootstrap):

One question, different dataset, different answers

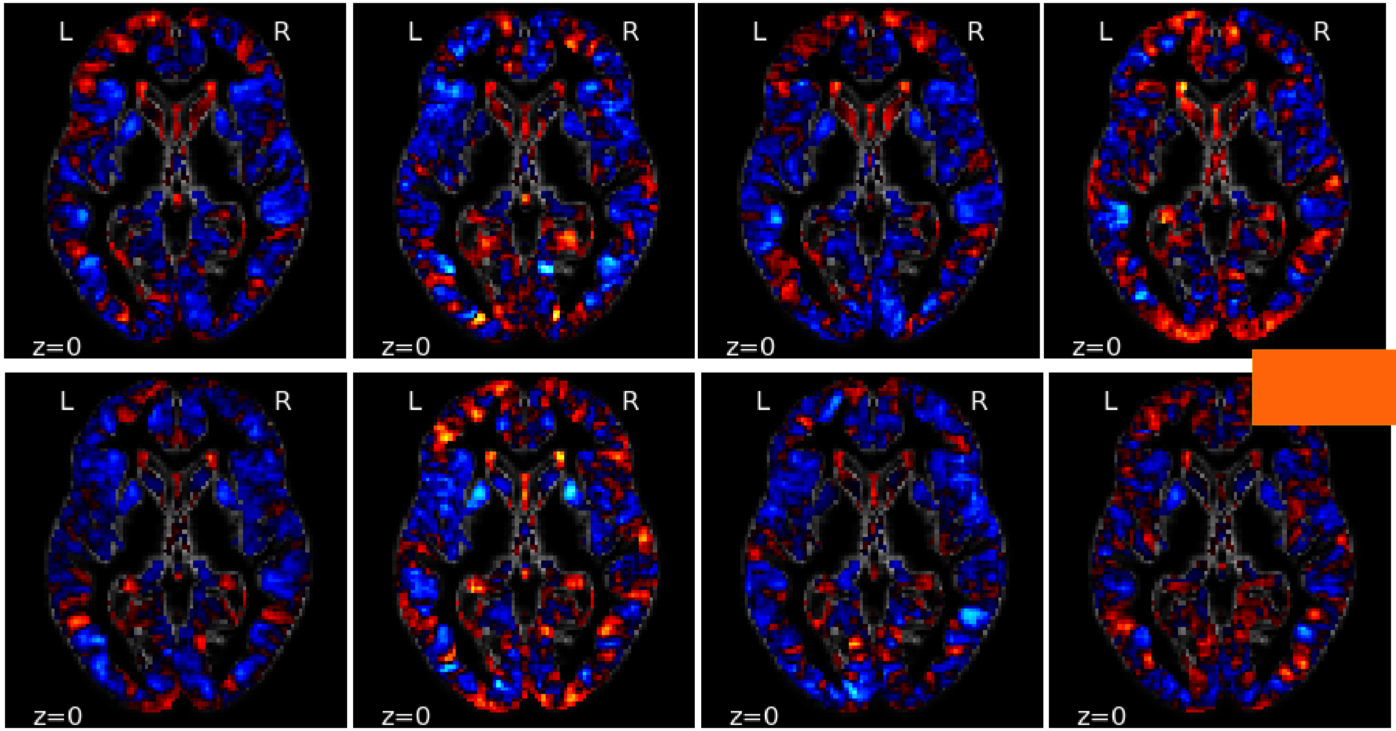

Summarized into a **d image**:

(effect size) / (effect std)

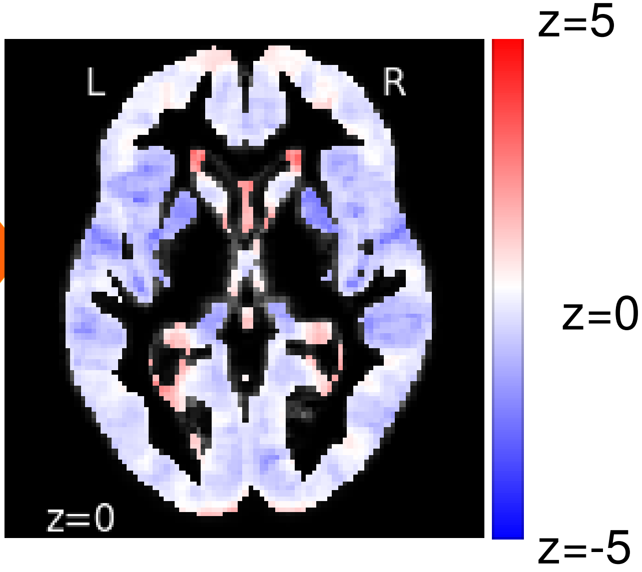

Weight maps for age prediction / OASIS

n=10

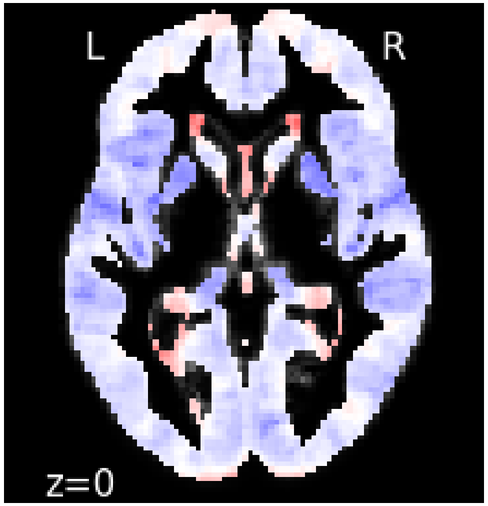

n=20

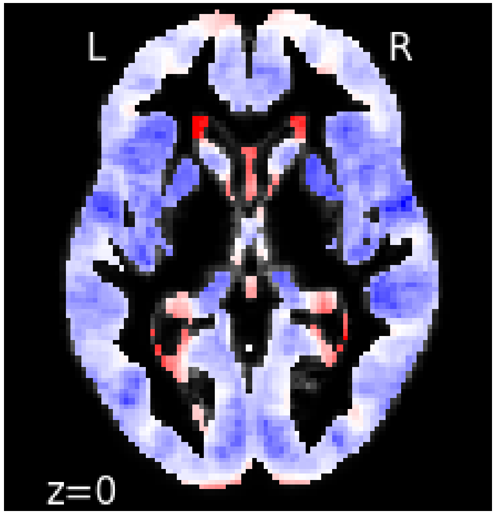

n=50

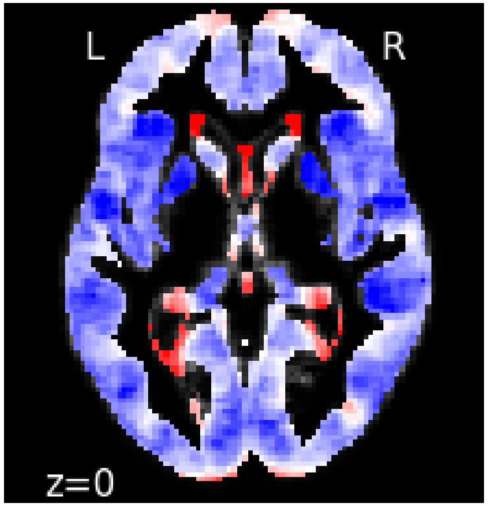

z=5

z=0

z=-5

(effect size estimated by bootstrap)

n=100

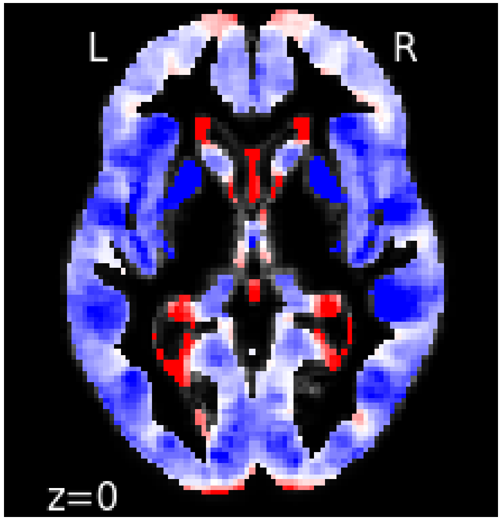

n=200

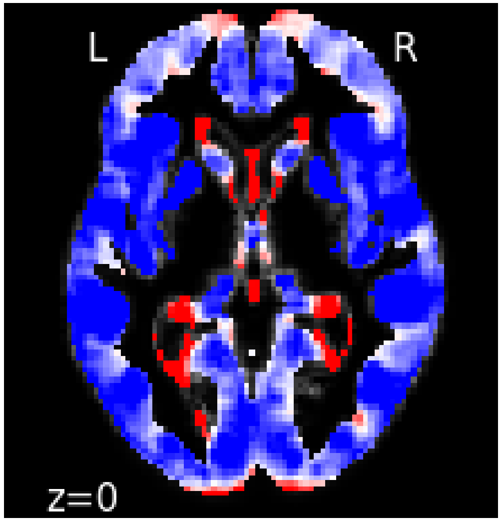

n=300

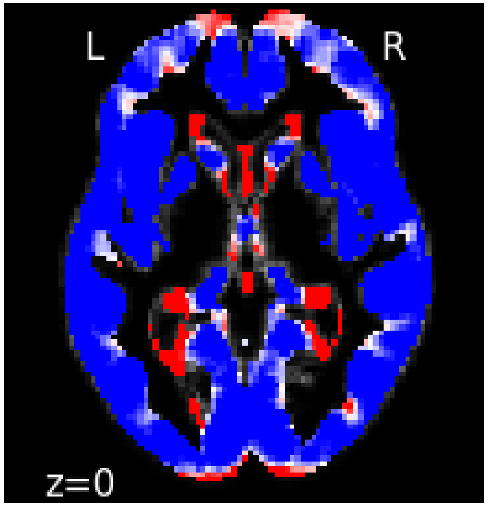

Better than chance ?

- Chance level sometimes unknown
  - e.g. imbalanced classes
- Use a dummy classifier to estimate it

```
>>> from sklearn.dummy import DummyClassifier
>>> null_cv_scores = cross_val_score(DummyClassifier(), fmri_masked, target, cv=cv)
```

```
>>> from sklearn.model_selection import permutation_test_score
>>> null_cv_scores = permutation_test_score(svc, fmri_masked, target, cv=cv)
```

- Spread of accuracy obtained across validation folds: accuracy > chance in 90% folds

Better than chance ?

- Chance level sometimes unknown
  - e.g. imbalanced classes
- Use a dummy classifier to estimate it

```
>>> from sklearn.dummy import DummyClassifier
>>> null_cv_scores = cross_val_score(DummyClassifier(), fmri_masked, target, cv=cv)
```

```
>>> from sklearn.model_selection import permutation_test_score
>>> null_cv_scores = permutation_test_score(svc, fmri_masked, target, cv=cv)
```

- Spread of accuracy obtained across validation folds: accuracy > chance in 90% folds

Scores across validation splits are not independent. Don’t use t-tests to assess significance of accuracy differences

Outline

- Keep your decoding model simple
- Beware accuracy and significance
- Inference about decoding model parameters

Interpreting a decoding model

Linear classifiers provide **weight maps**: Can we interpret them ?

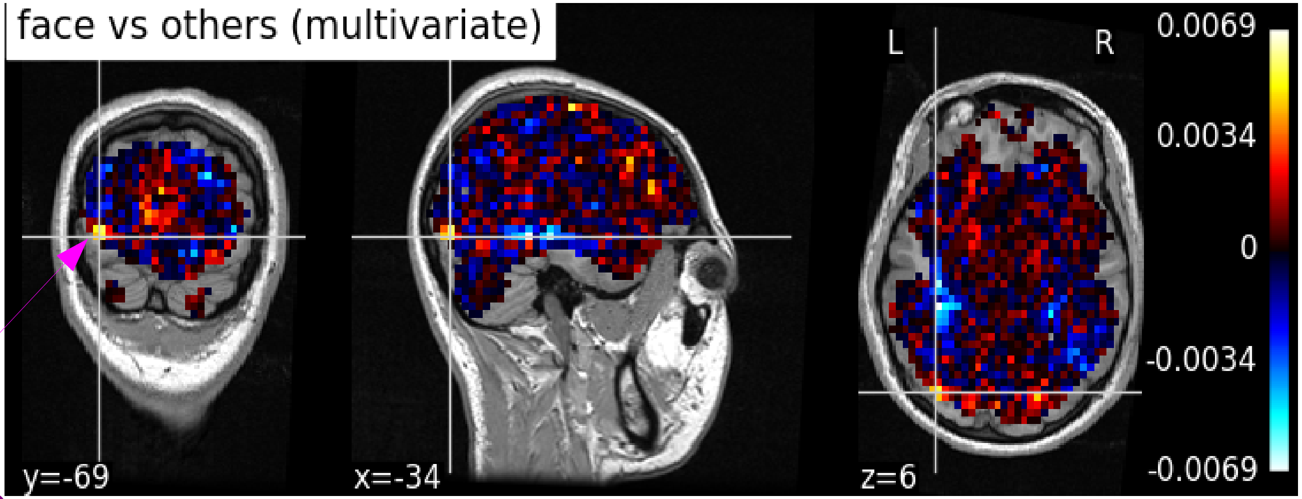

Is this spot encoding the "face" information ?

Problem: the **whole pattern** is predictive, hard to say anything about this particular region

Reverse inference

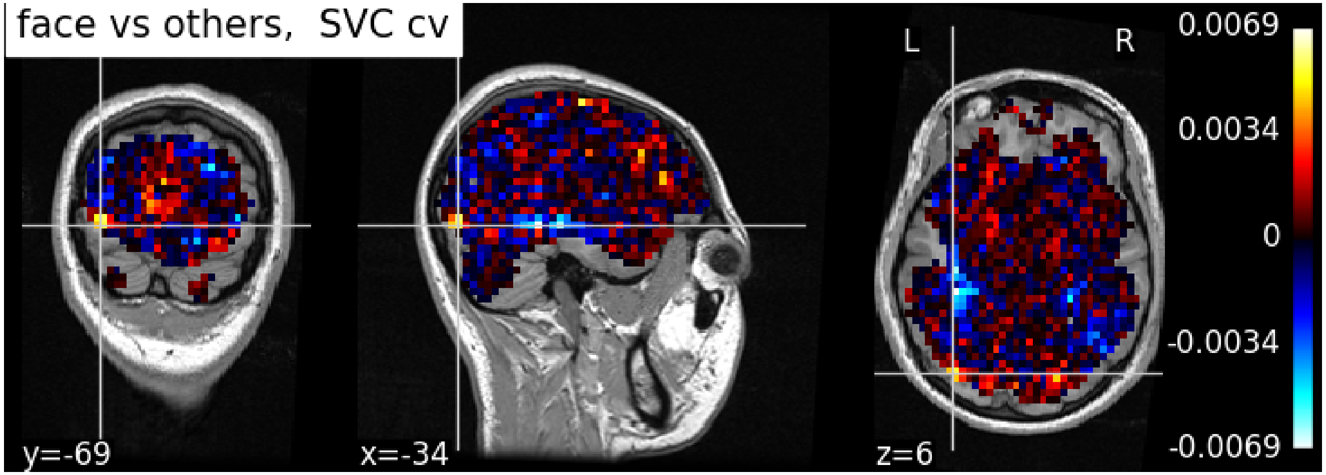

Conditional association test:

does $\mathbf{X}_i$ predict $\mathbf{y}$, **given** other signals $(\mathbf{X}_j)_{j \neq i}$

$\mathrm{Corr}\,(\mathbf{X}_i, \mathbf{y} | (\mathbf{X}_j)_{j \neq i}) \neq 0$

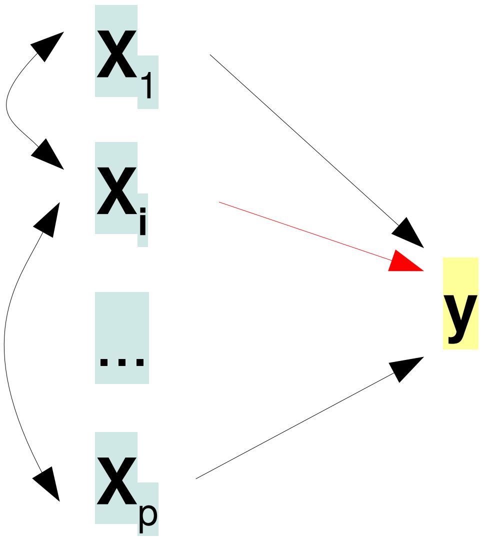

The importance of $\mathbf{X}_i$ depends on all other brain regions

Conditional and marginal inference

Multivariate linear models

$\text{Corr}\,(\mathbf{X}_\mathbf{i}, \mathbf{y}|(\mathbf{X}_\mathbf{j})_{j\neq i})$

- Hard to compute

  ill posed inverse problem
- Necessary to assert

implication of region

Univariate linear models

$\text{Corr}\,(\mathbf{X}_\mathbf{i}, \mathbf{y})$

- Easy to compute
- Unaware of other brain regions

Both are necessary [Weichwald et al. NeuroImage 2015]

Reverse inference is ill-posed

- Results depend on the prior used

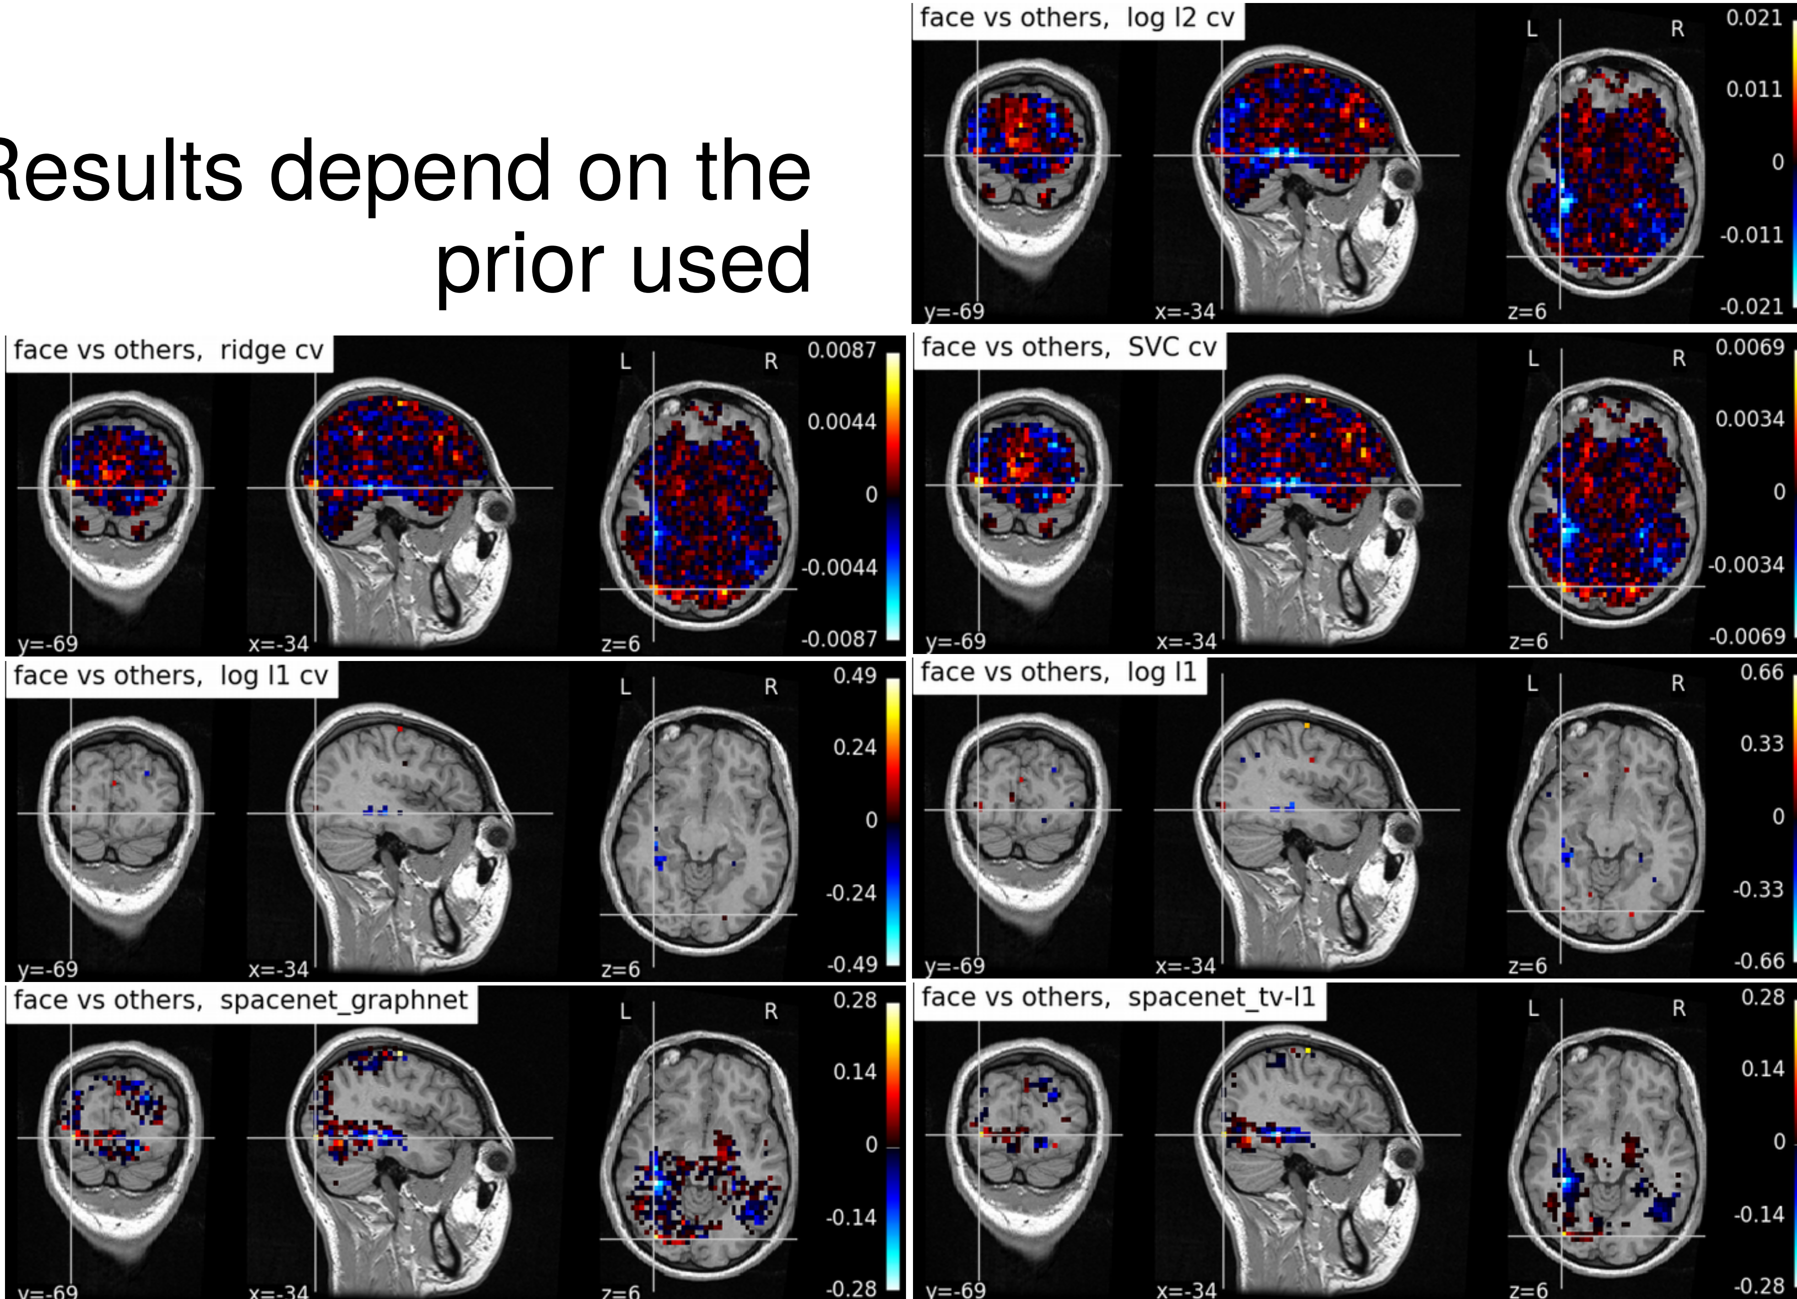

p-values on Predictive Patterns

[Chevalier et al. MICCAI 2018,
submitted to Neuroimage]

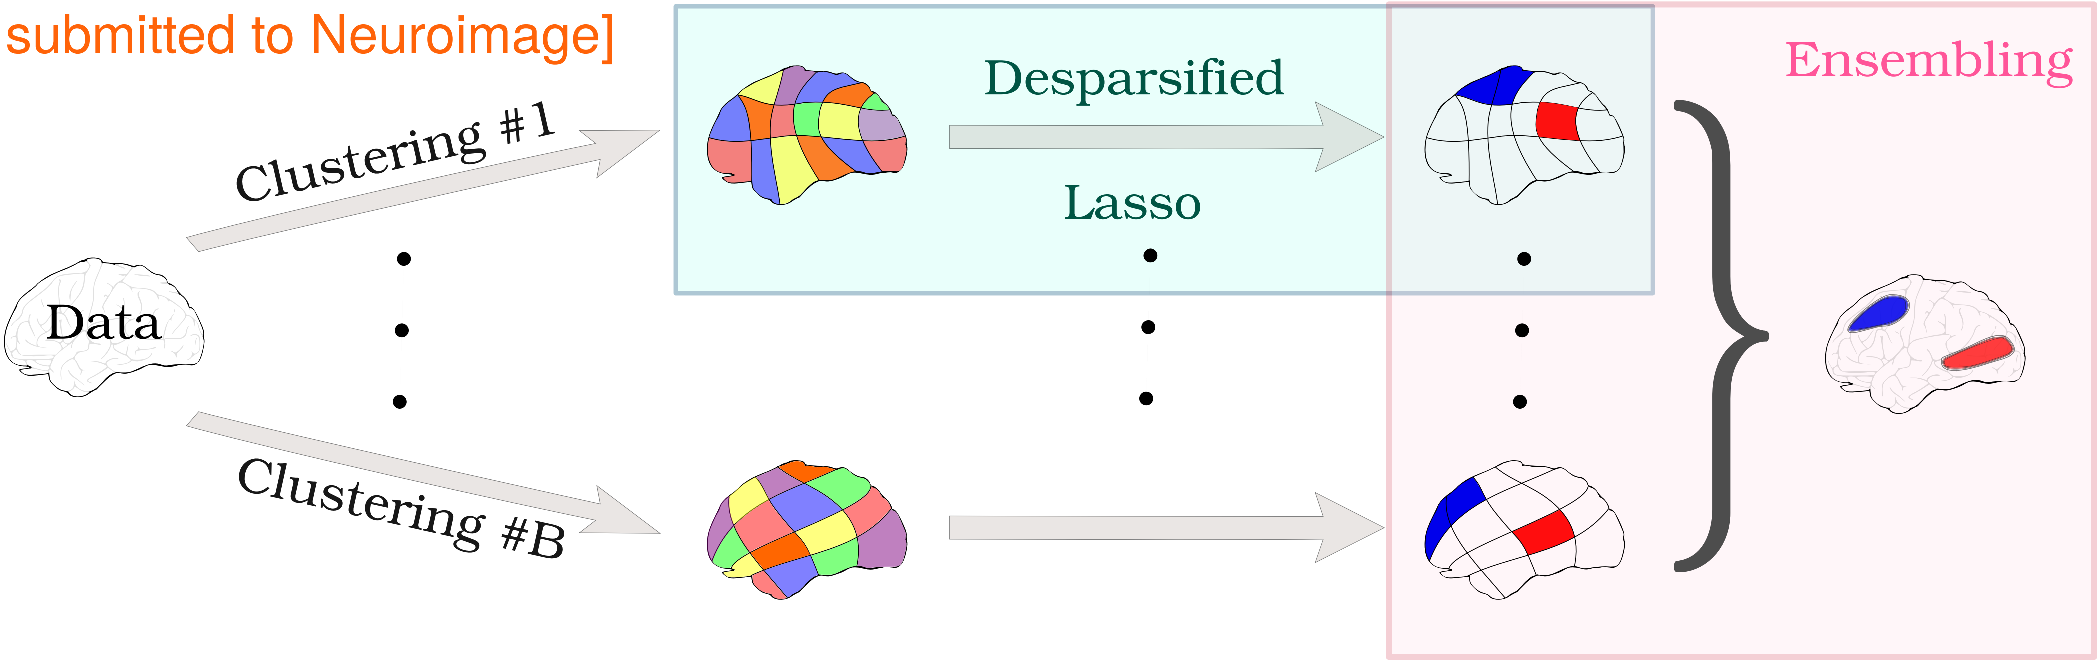

EnCluDL: estimate jointly the predictive pattern and its variance

→ p-values on pattern maps

→ type-1 error control

p-values on Predictive Patterns

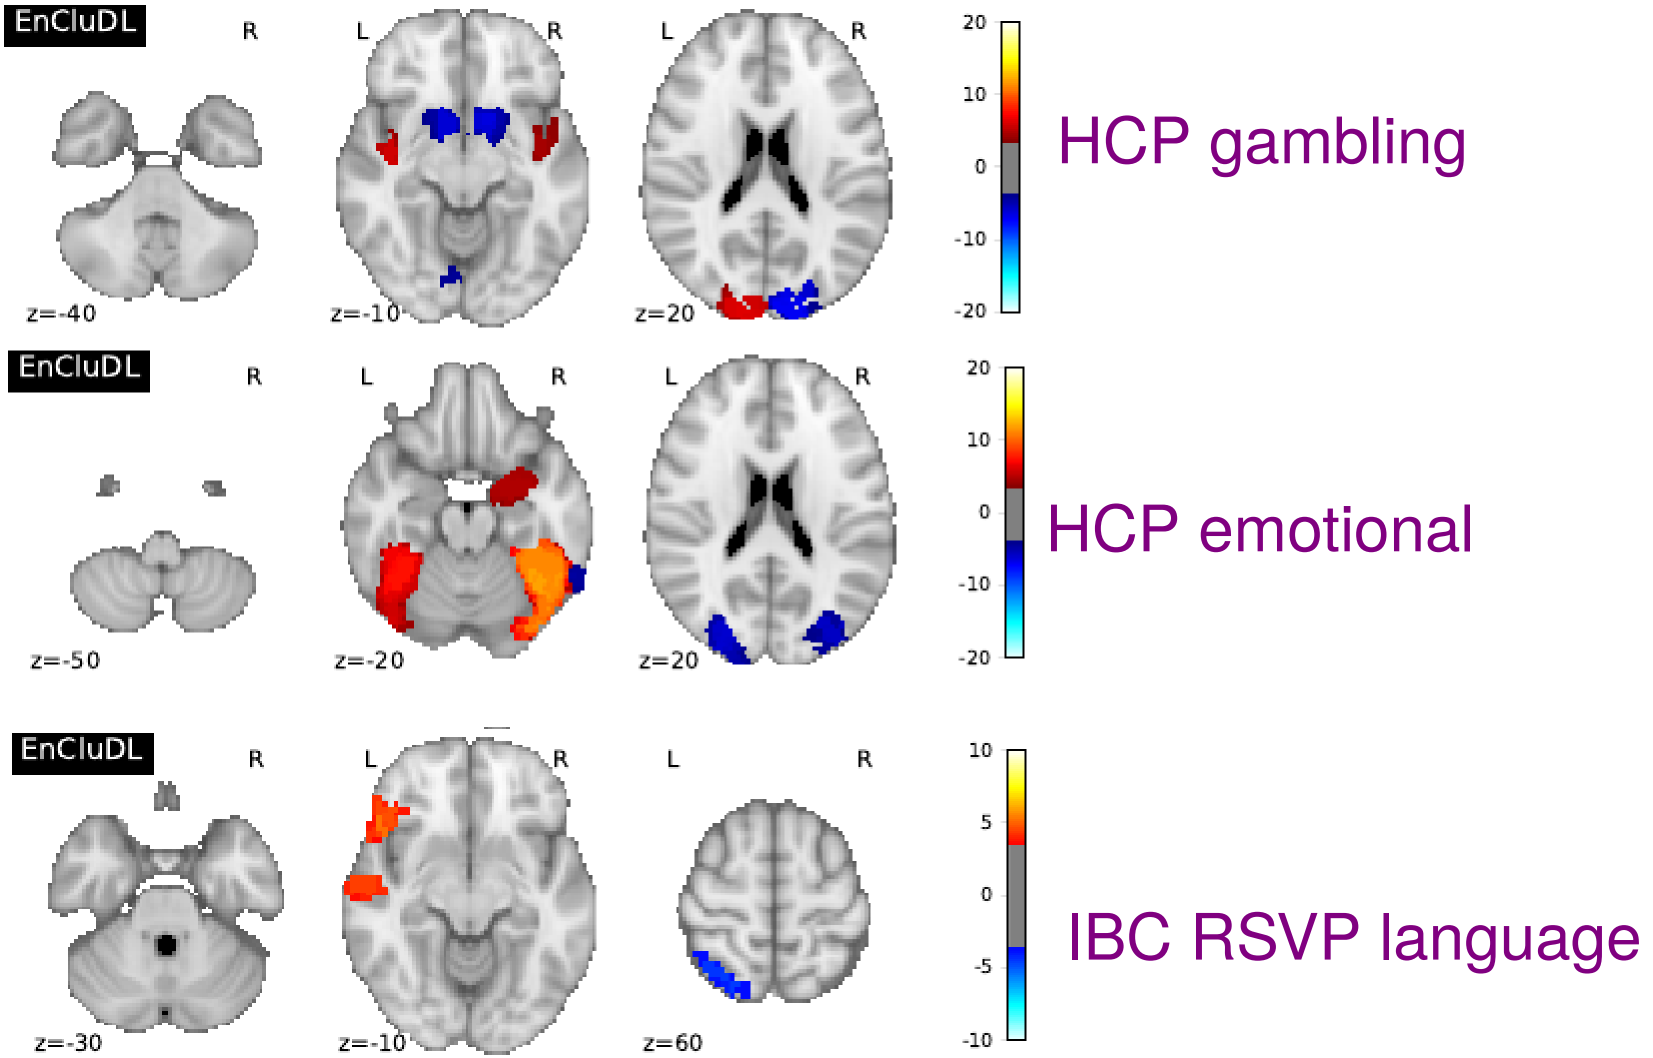

FWER ($p<0.05$) control on predictive patterns

Available here: https://github.com/ja-che/hidimstat

Functional specificity and MVPA

Discriminating patterns only reflect the categories present in the dataset

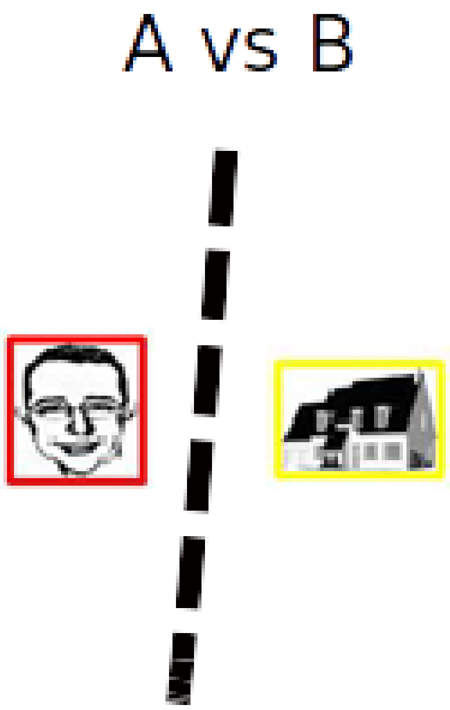

Functional specificity and MVPA

Discriminating patterns only reflect the categories present in the dataset

A vs B

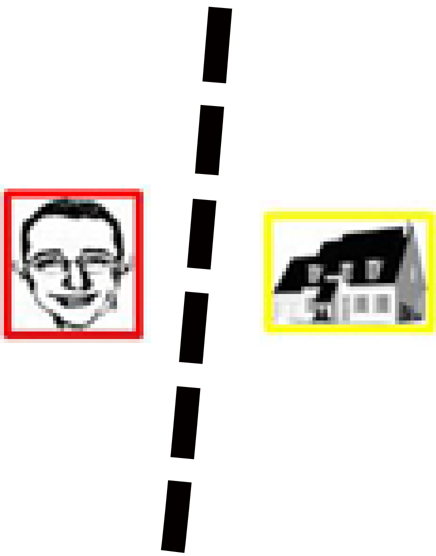

A vs B, C, D, E, F

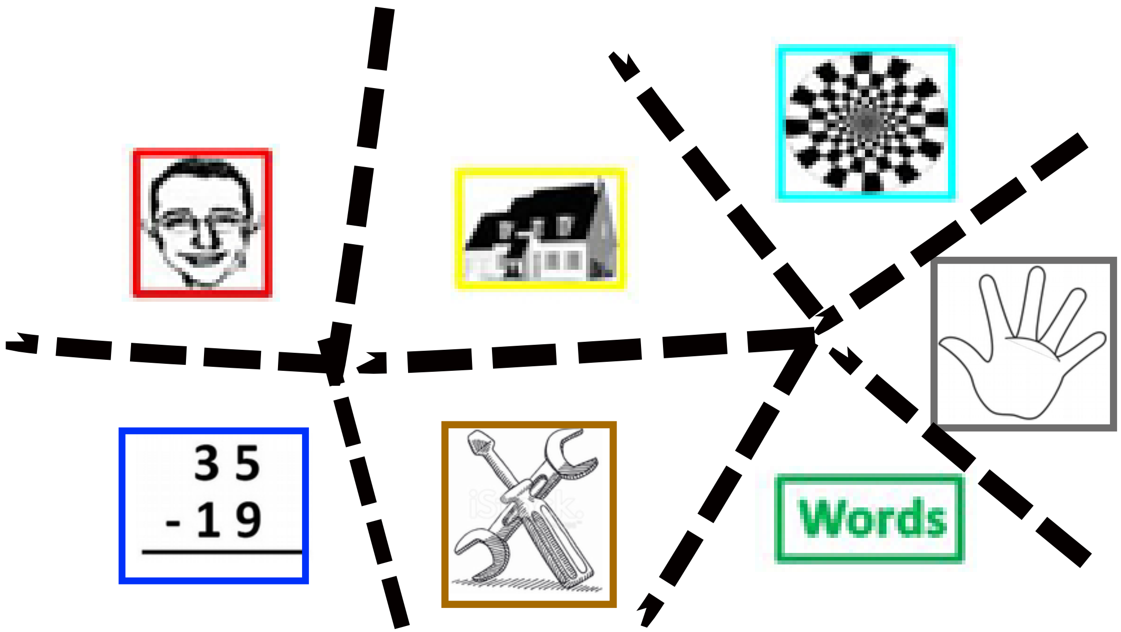

Functional specificity and MVPA

Discriminating patterns only reflect the categories present in the dataset

A vs B

A vs B, C, D, E, F

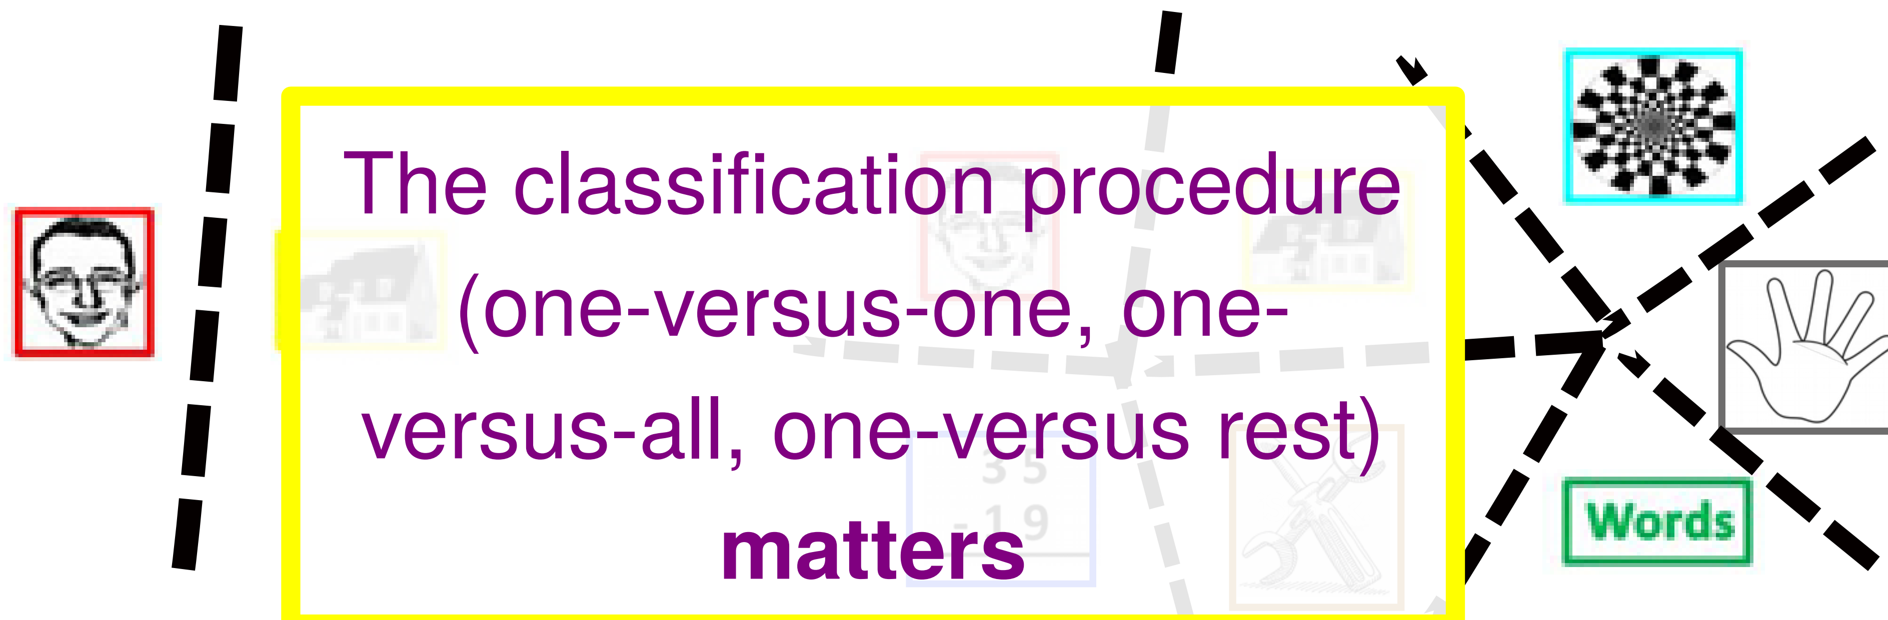

The classification procedure (one-versus-one, one-versus-all, one-versus rest) **matters**

Conclusion

- Limit complexity of Decoding procedures
  - Complexity come with major costs (time, statistical)
  - Linear models perform well
  - Use high-level objects
- Pattern maps are hard to interpret
  - Relative to many choices
  - Measure conditional associations
  - Novel estimators (model ensembles) with stat guarantees

The power of scikit learn for MVPA

- Machine learning for neuroimaging http://nilearn.github.io
- Scikit-learn-like API, BSD, Python, OSS
  - Classification of neuroimaging data (decoding)
  - Functional connectivity analysis
  - GLM analysis

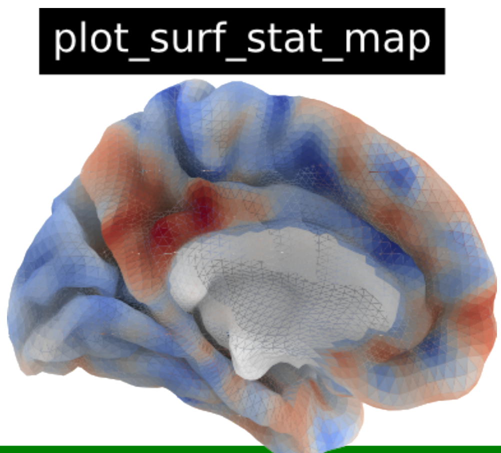

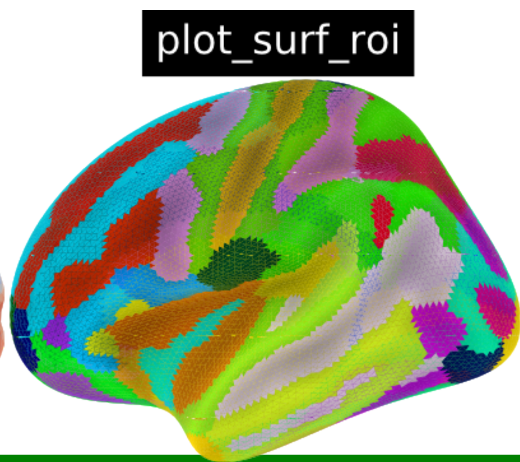

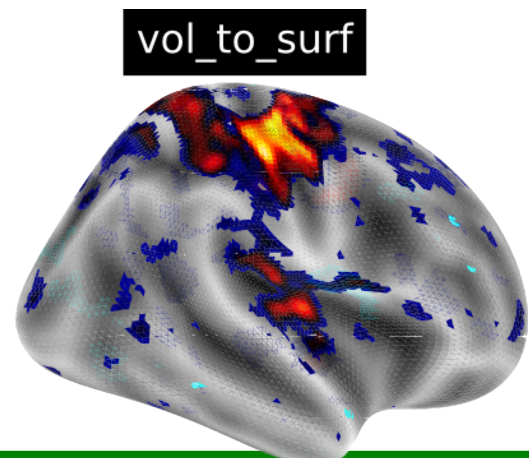

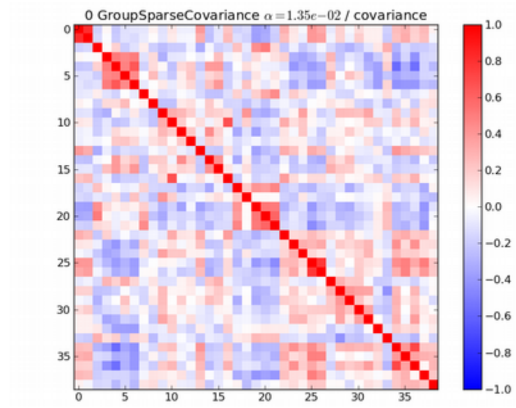

Parietal

G. Varoquaux,
A. Gramfort,
P. Ciuciu,
D. Wassermann,
D. Engemann,
B. Nguyen
A.L. Grilo Pinho,
E. Dohmatob,
A. Mensch,
J.A. Chevalier,
A. Hoyos idrobo,
D. Bzdok,
J. Dockès,
P. Cerda,
C. Lazarus
D. La Rocca
G. Lemaitre
L. El Gueddari
O. Grisel
M. Massias
P. Ablin
H. Janati
J. Massich
K. Dadi
H. Richard
C. Petitot

Acknowledgements

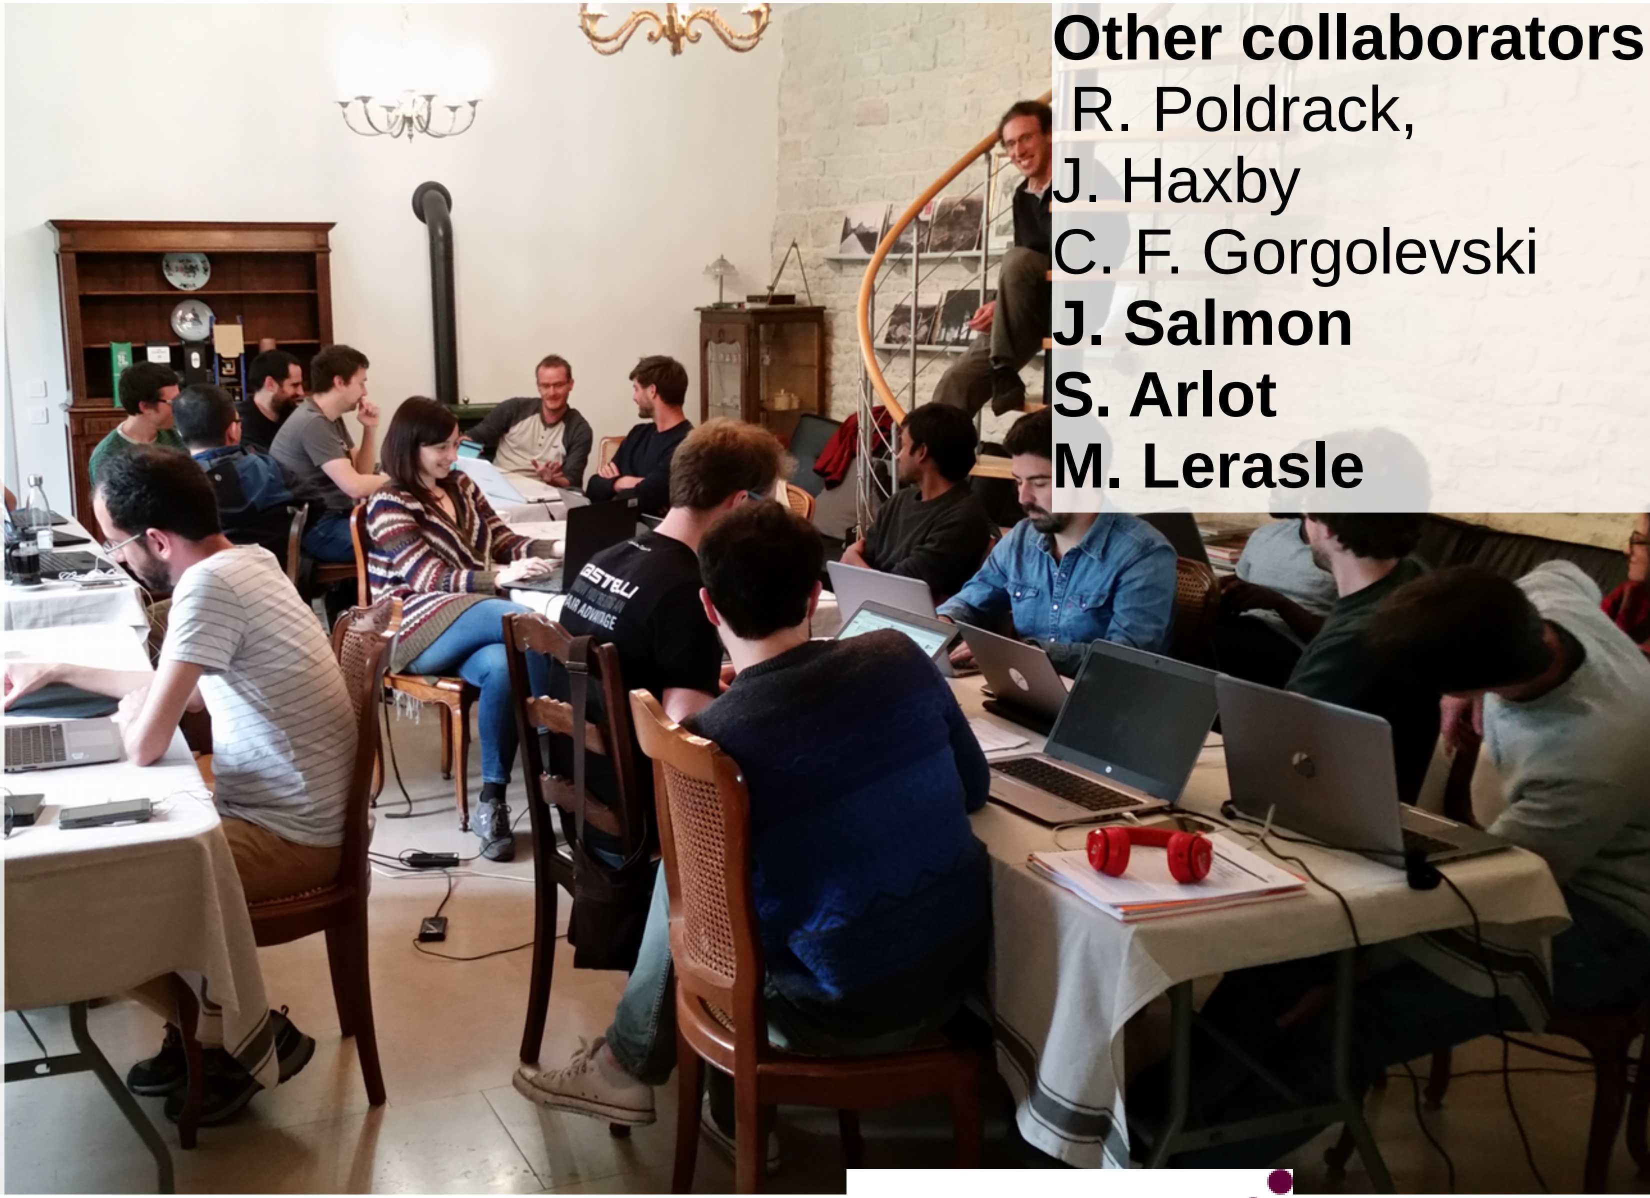

Other collaborators
R. Poldrack,
J. Haxby
C. F. Gorgolevski
J. Salmon
S. Arlot
M. Lerasle

Human Brain Project
université PARIS-SACLAY
AGENCE NATIONALE DE LA RECHERCHE ANR

To go further

- Toward a unified framework for interpreting machine-learning models in neuroimaging L Kohoutová, J Heo, S Cha, S Lee, T Moon, TD Wager, CW Woo Nature Protocols 15 (4), 1399-1435
- Encoding and decoding in fMRI. T Naselaris, KN Kay, S Nishimoto, JL Gallant. Neuroimage 56 (2), 400-410
- Predictive models avoid excessive reductionism in cognitive neuroimaging. Gaël Varoquaux, Russell Poldrack. Current Opinion in Neurobiology, Elsevier, In press, 55, 2018
- A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. JD Haynes. Neuron 87 (2), 257-270
- Atlases of cognition with large-scale human brain mapping. Y Schwartz, RA Poldrack, B Gauthier, D Bzdok, JB Poline, B Thirion. PLoS Computational Biology 14 (11), e1006565