

Multi-voxel pattern analysis:

Decoding Mental States from fMRI Activity Patterns

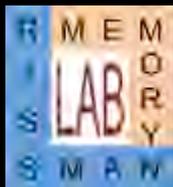


Artwork by Leon Zernitsky

Jesse Rissman, Ph.D.

Depts. of Psychology, Psychiatry and Biobehavioral Sciences,
Brain Research Institute, Integrative Center for Learning & Memory

UCLA



Core idea:

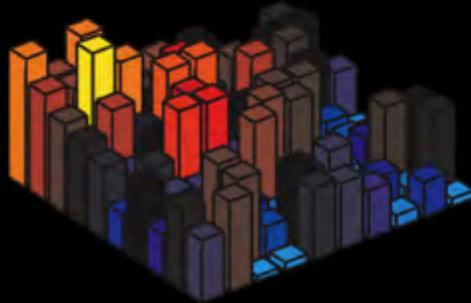
There is more informational content in BOLD activity patterns than typically detected with univariate analyses



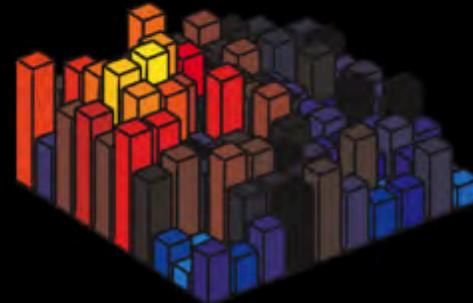
Core idea:

There is more informational content in BOLD activity patterns than typically detected with univariate analyses

Condition 1



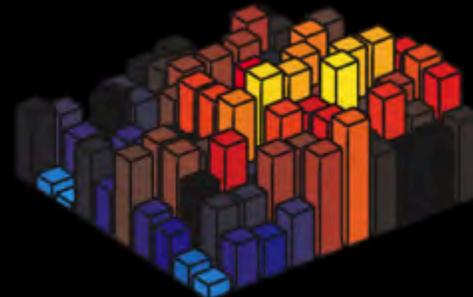
Condition 2



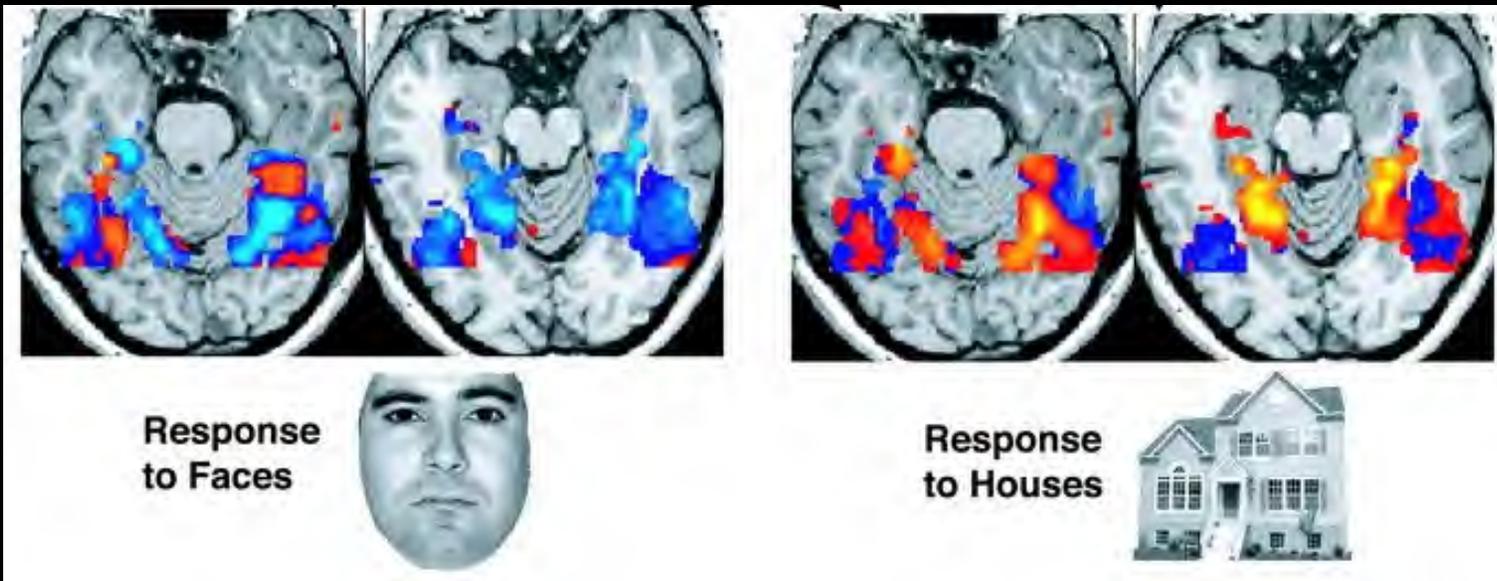
Condition 3



Condition n

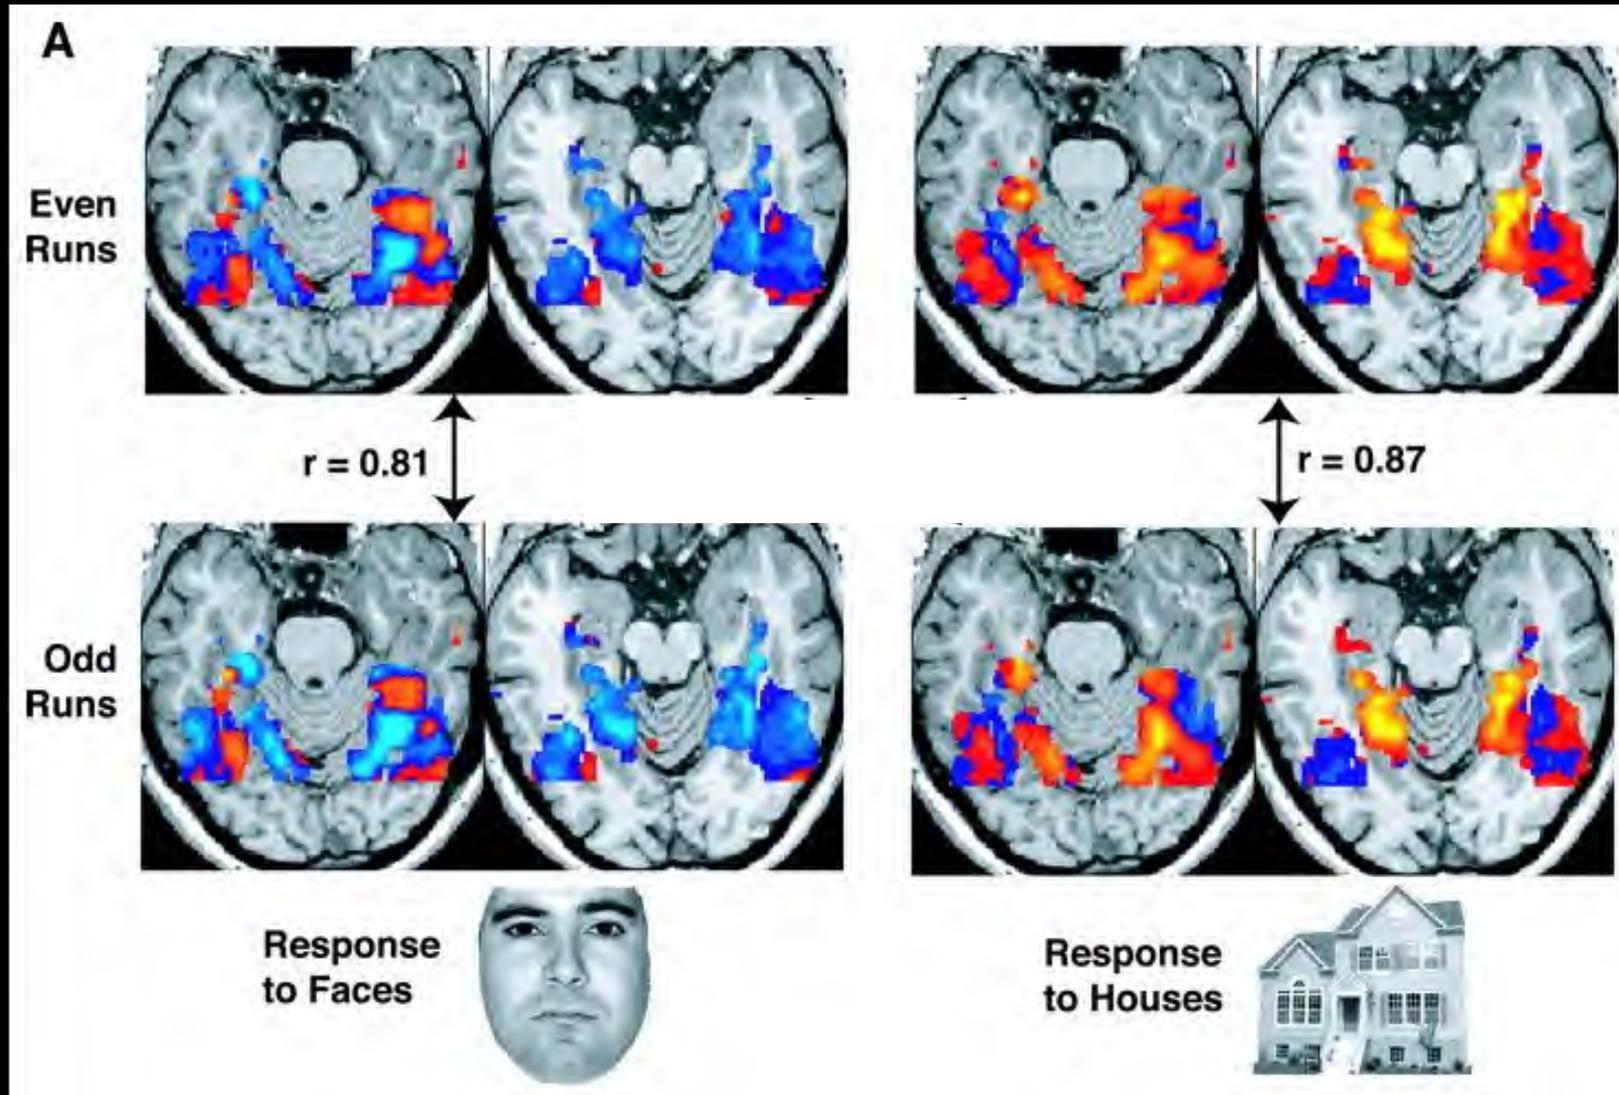


How it all began...



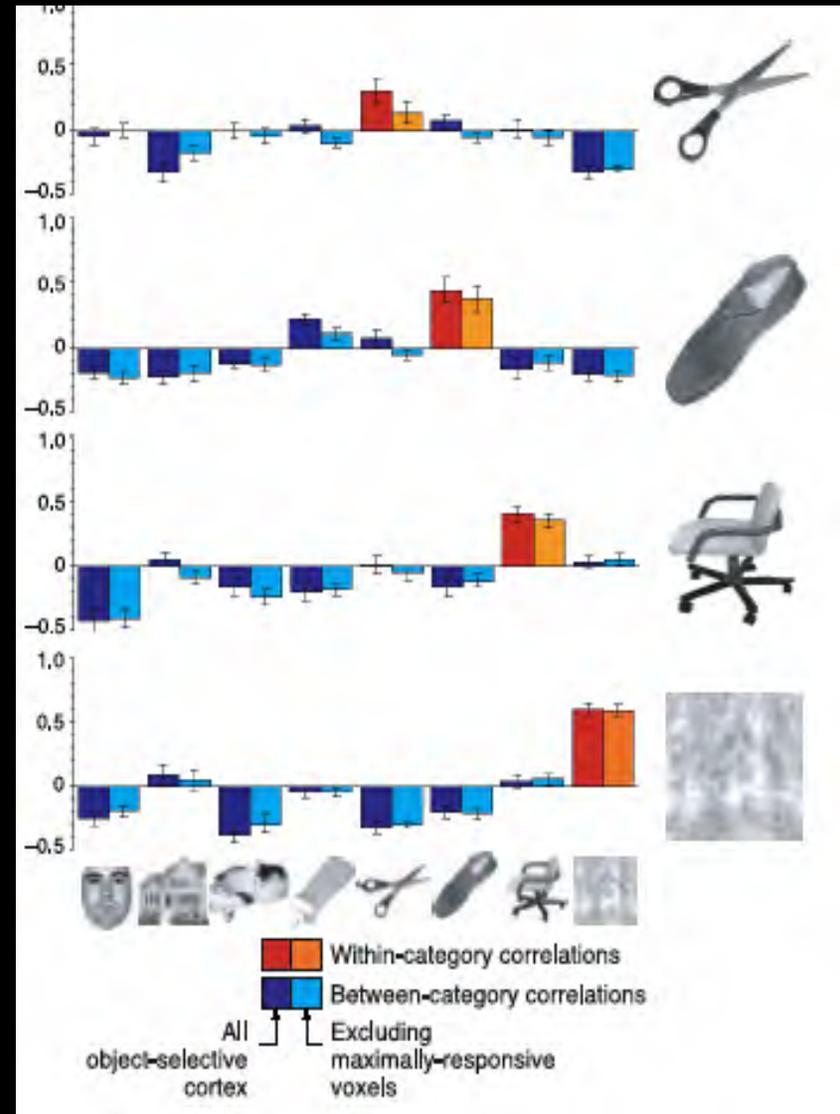
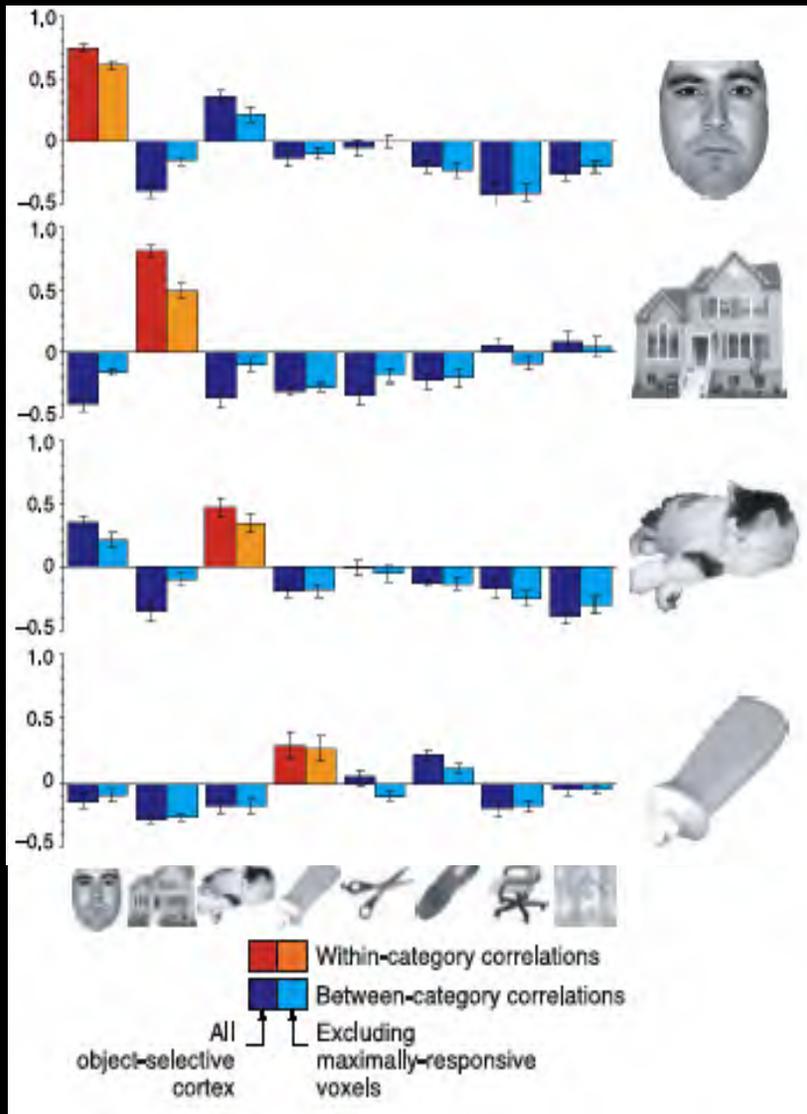
Haxby et al. (2001), *Science*

The Power of Patterns



Haxby et al. (2001), *Science*

The Power of Patterns



Haxby et al. (2001), *Science*

Adopting a Machine Learning Framework



ACADEMIC
PRESS

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

NeuroImage

NeuroImage 19 (2003) 261–270

www.elsevier.com/locate/ynimg

Functional magnetic resonance imaging (fMRI) “brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex

David D. Cox^{a,b,*} and Robert L. Savoy^{a,b,c}

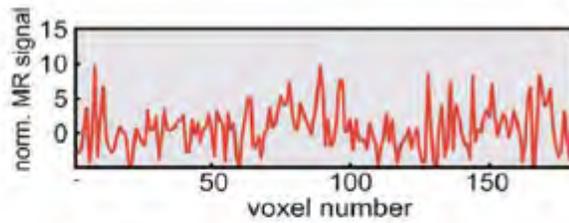
^a Rowland Institute for Science, Cambridge, MA 02142, USA

^b Athinoula A. Martinos Center for Structural and Functional Biomedical Imaging, Charlestown, MA 02129, USA

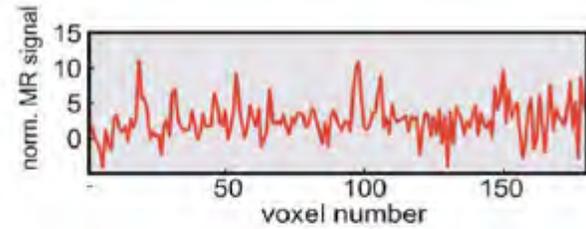
^c HyperVision, Inc., P.O. Box 158, Lexington, MA 02420, USA

Cox & Savoy (2003), *NeuroImage*

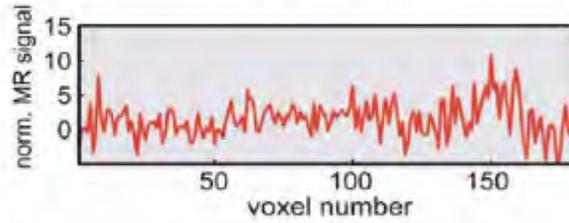
Baskets



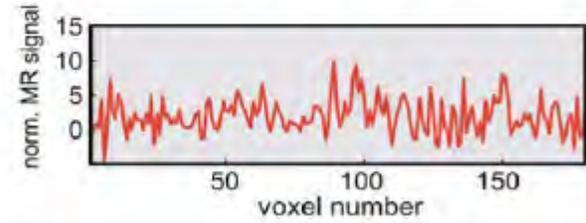
Birds



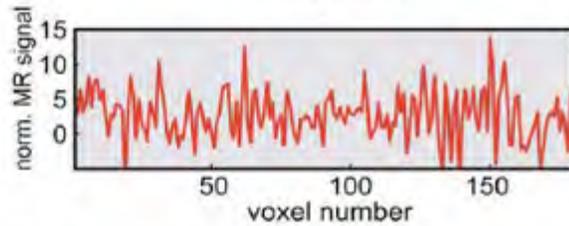
Butterflies



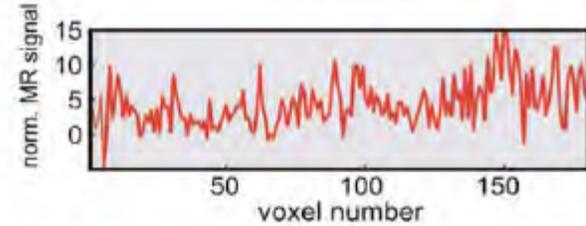
Chairs



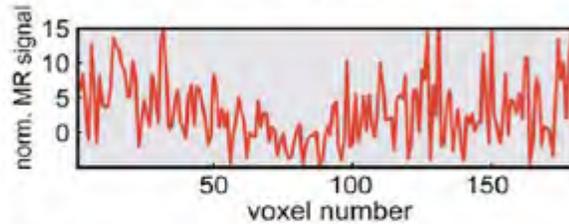
Cows



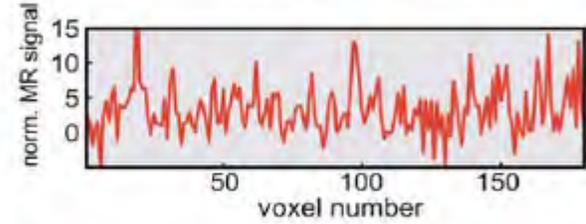
Fish



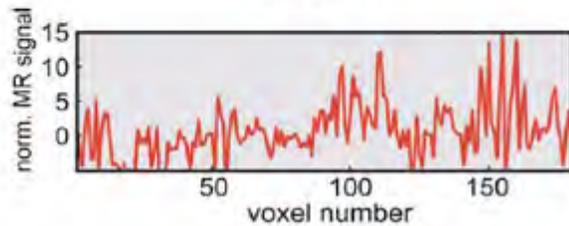
Gnomes



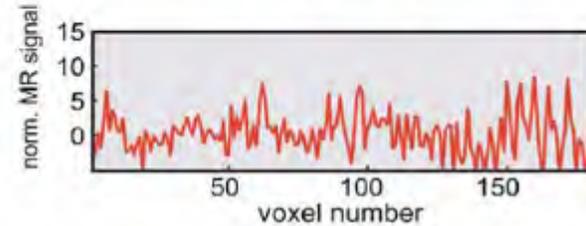
Horses



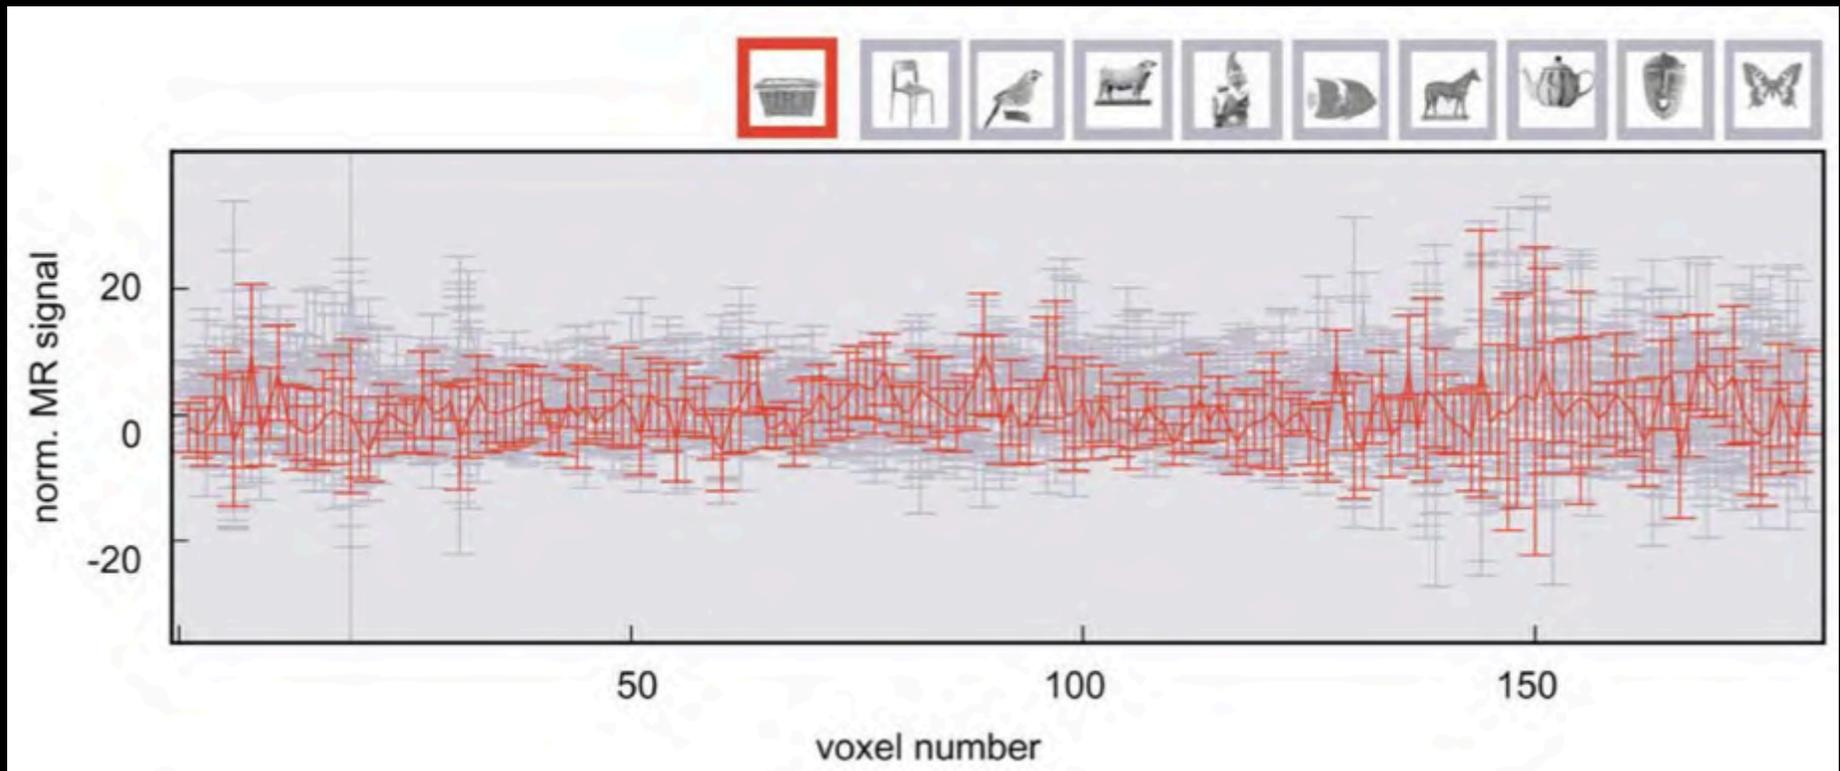
Masks



Teapots



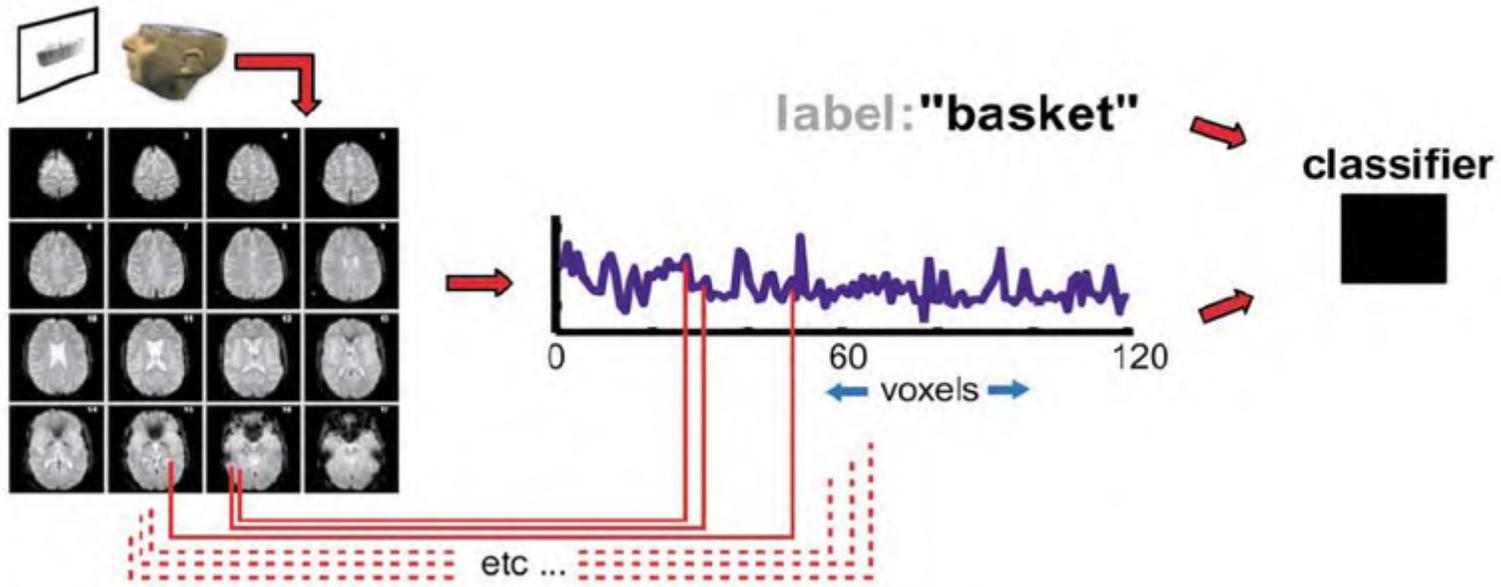
Cox & Savoy (2003)



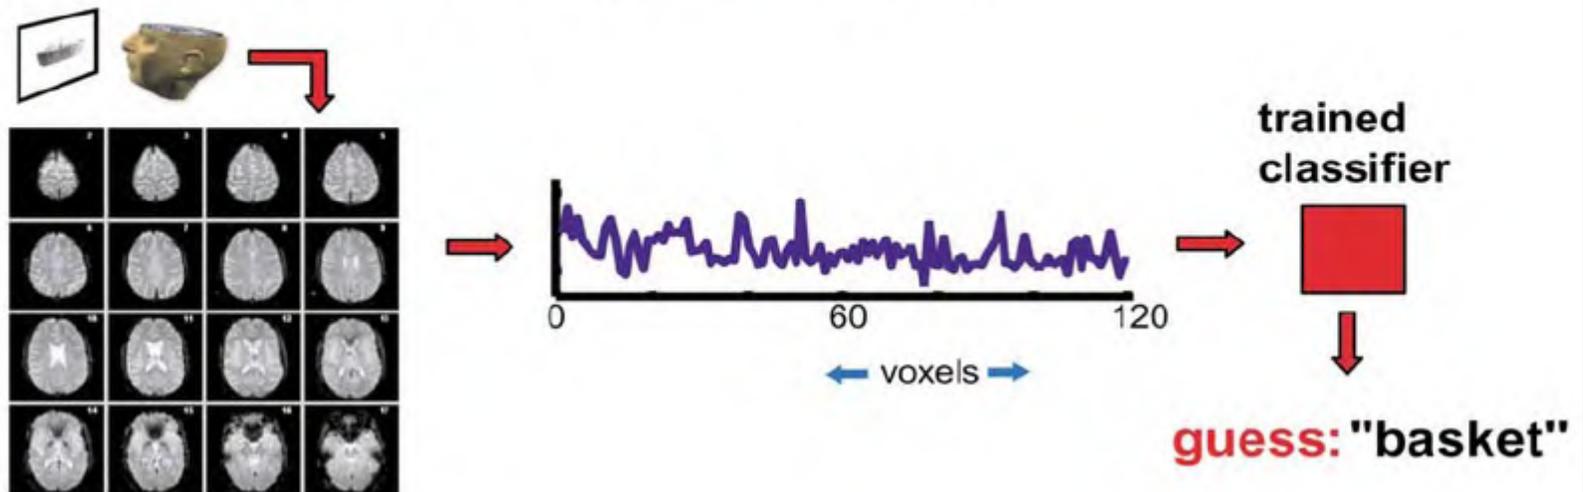
NOTE: no individual voxels show strong basket-specific response

Cox & Savoy (2003)

Training

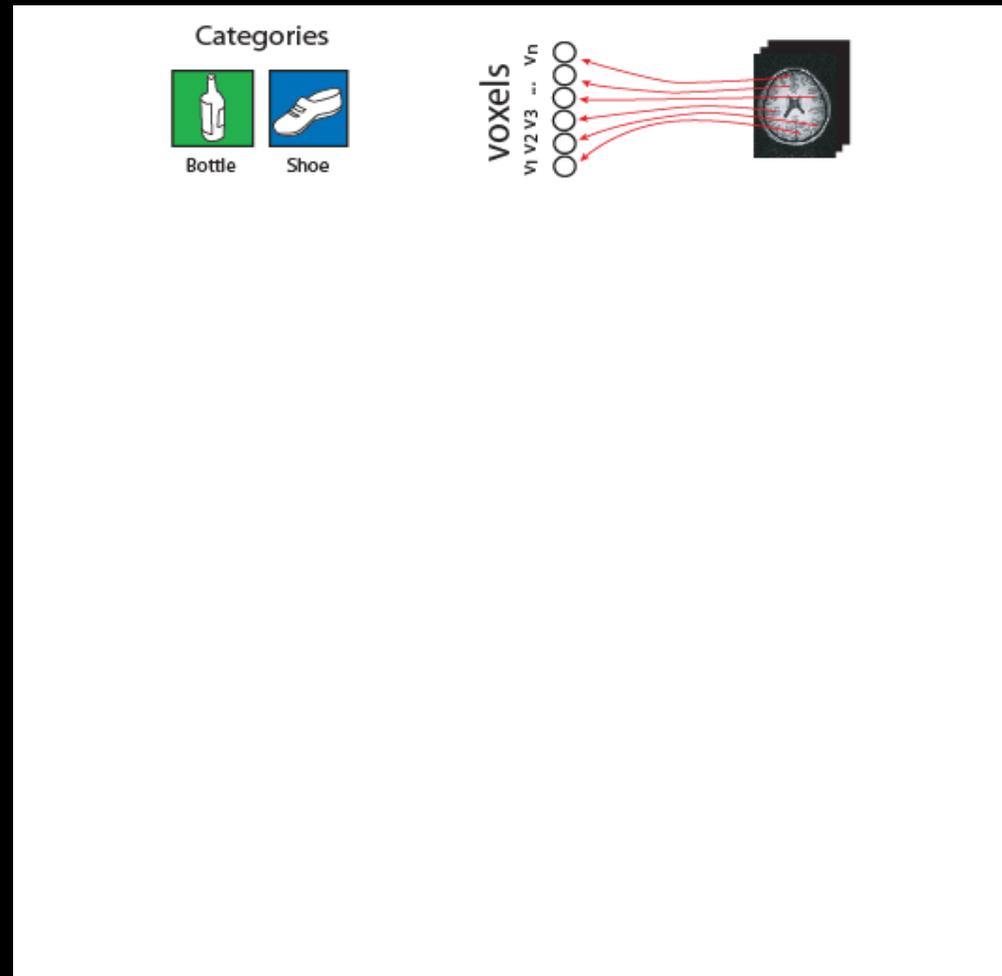


Classification (during a subsequent session)



The Multi-Voxel Pattern Analysis Approach

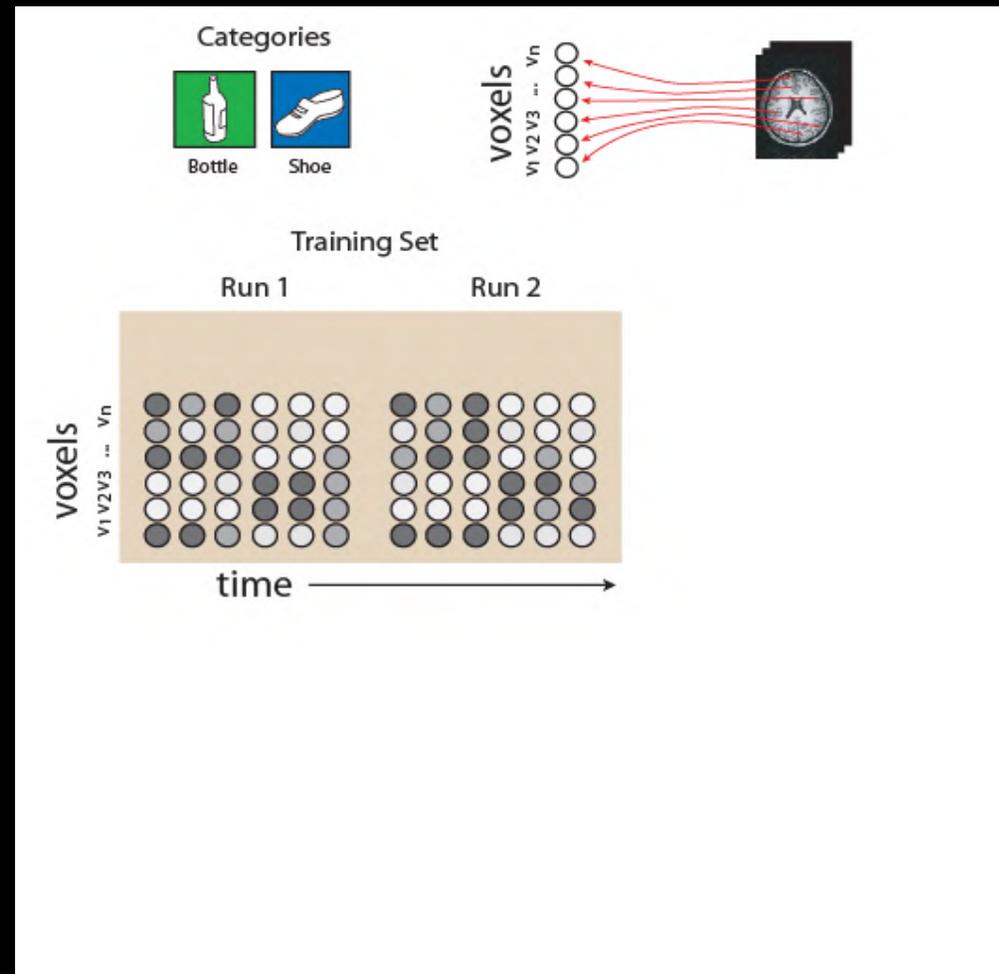
1. Acquire brain data while the subject is viewing shoes or bottles



The Multi-Voxel Pattern Analysis Approach

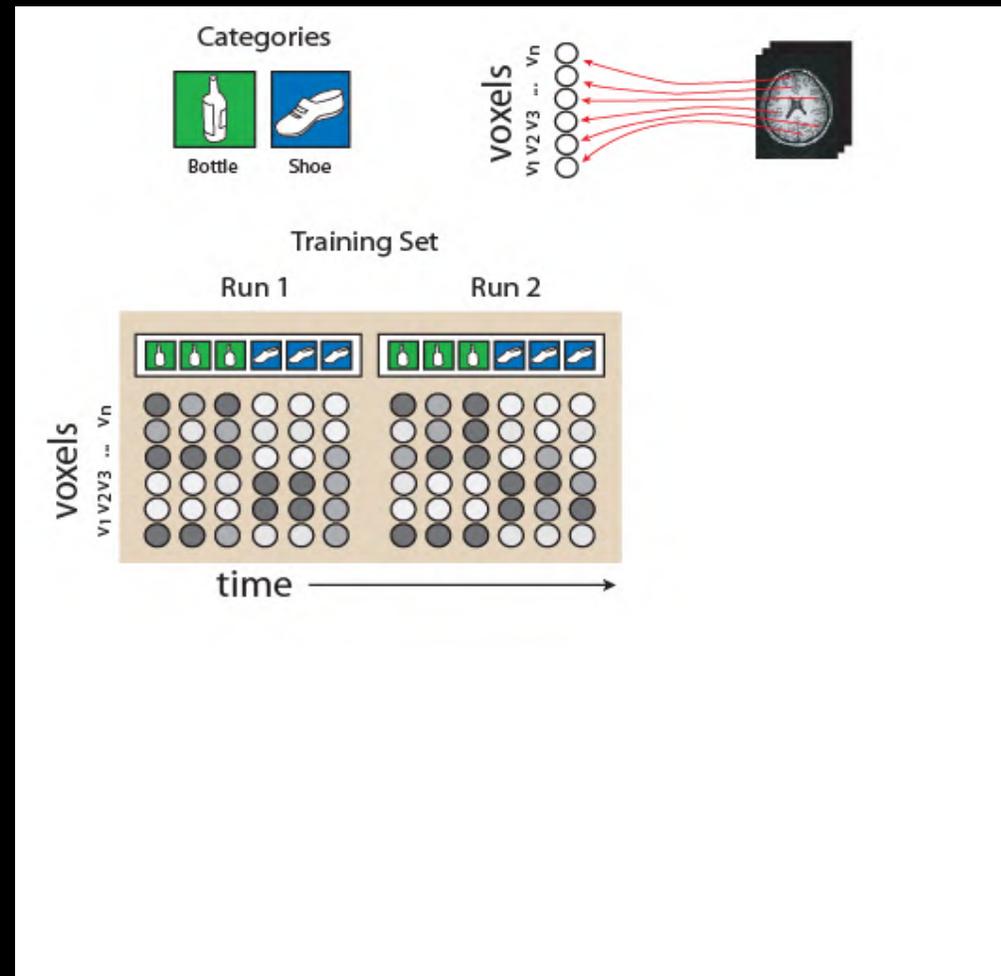
1. Acquire brain data
2. Convert each functional brain volume (or trial) into a vector that reflects the *pattern of activity across voxels* at that point in time

We typically do some kind of *feature selection* to cut down on the number of voxels



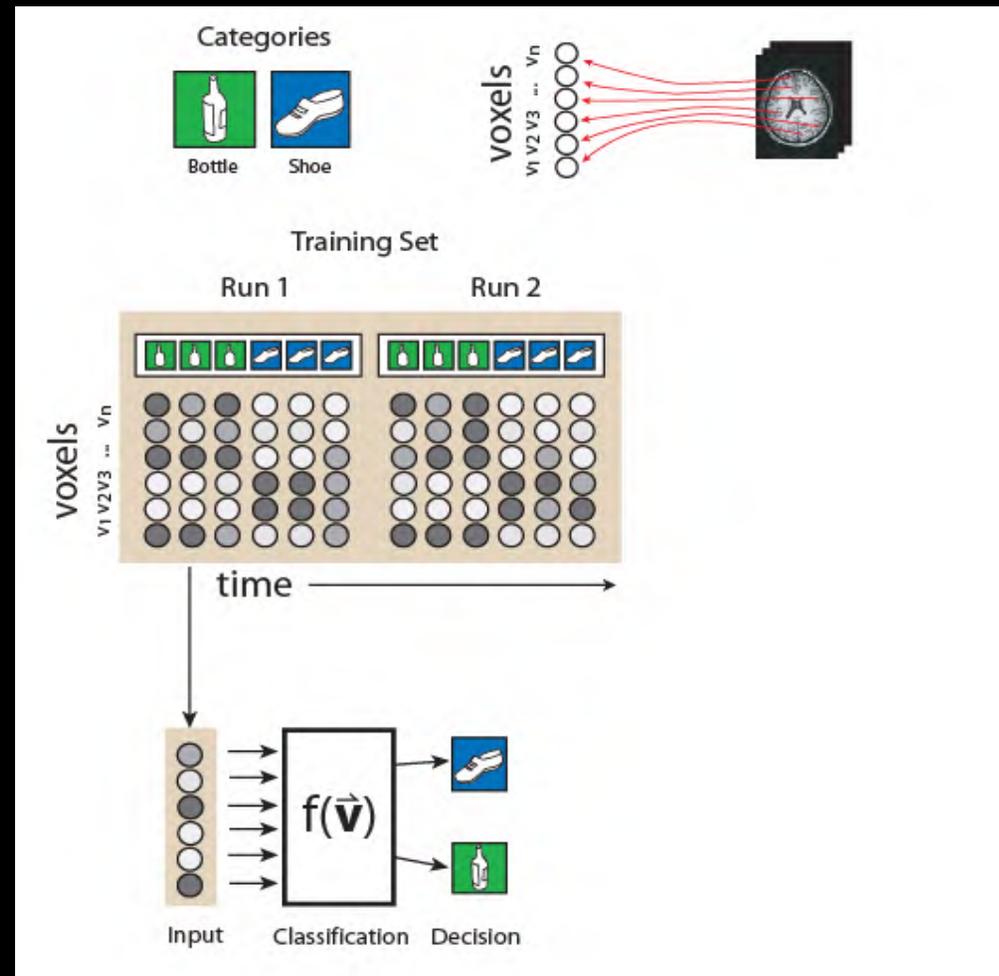
The Multi-Voxel Pattern Analysis Approach

1. Acquire brain data
2. Generate brain patterns
3. Label brain patterns according to whether the subject was viewing shoes vs. bottles (adjusting for lag in the blood flow response)



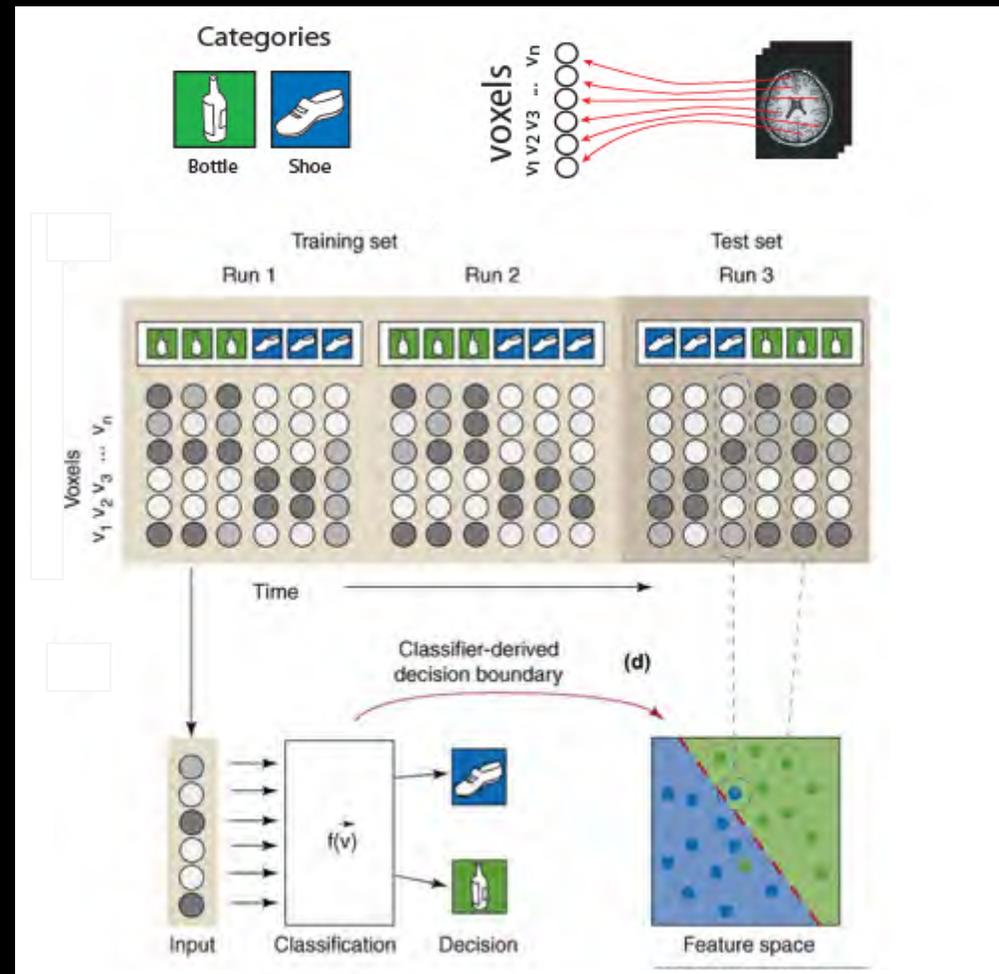
The Multi-Voxel Pattern Analysis Approach

1. Acquire brain data
2. Generate brain patterns
3. Label brain patterns
4. Train a classifier to discriminate between bottle patterns and shoe patterns



The Multi-Voxel Pattern Analysis Approach

1. Acquire brain data
2. Generate brain patterns
3. Label brain patterns
4. Train a classifier
5. Apply the trained classifier to new brain patterns (i.e., not included in training set)

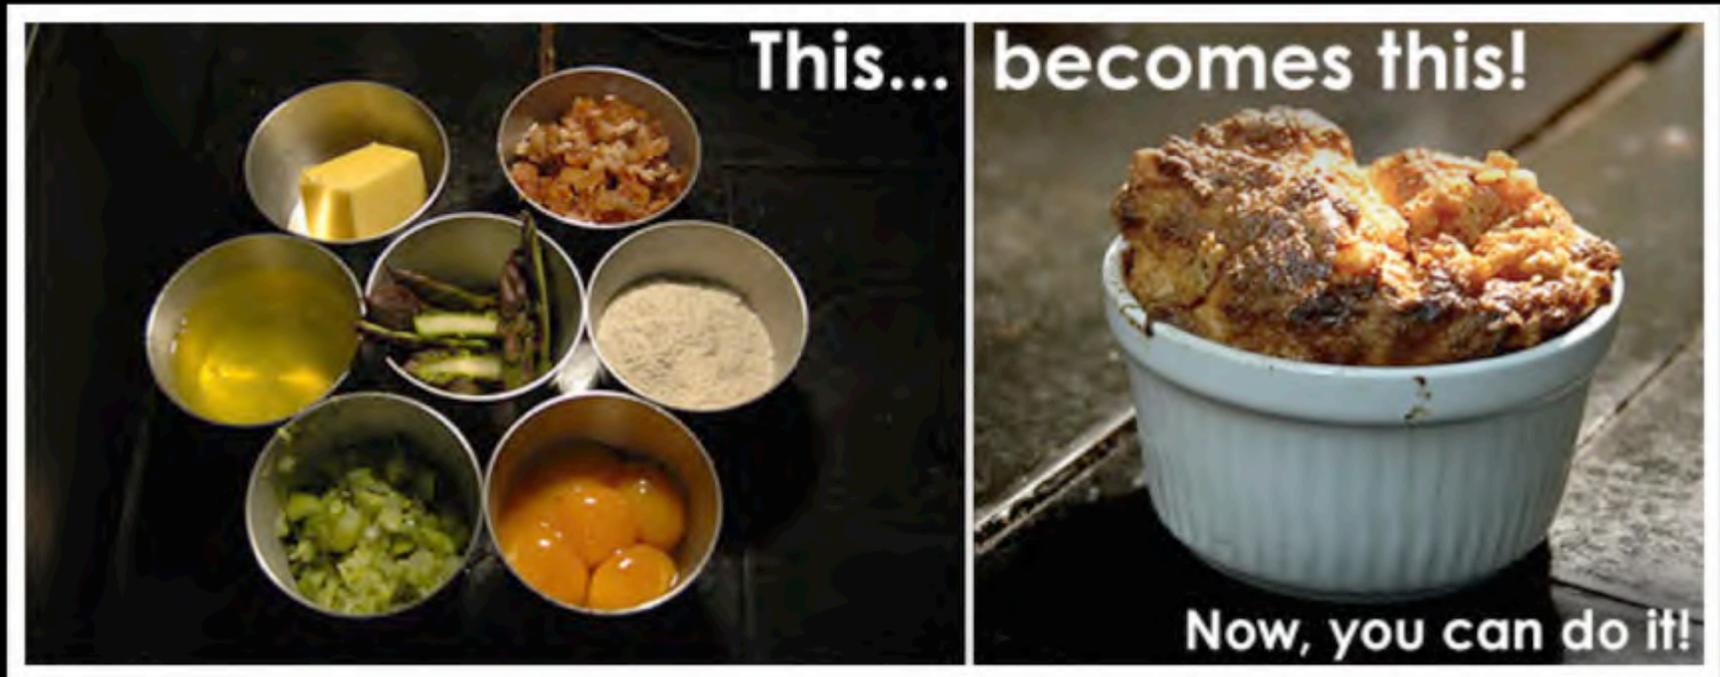


Goals of Multi-voxel Pattern Analysis

- Decoding percepts or thoughts (a.k.a. “mind-reading”)
 - What is a person perceiving, imagining, planning, or remembering?
- Decoding brain patterns (activity or connectivity) that distinguish individuals
 - Useful for diagnosis
- Characterizing the distributed cortical representations mediating specific cognitive processes
 - What features are extracted by which cortical areas or networks?
- Providing an index of the instantaneous activation level of specific mental representations
 - Can use this information to test psychological theories

MVPA Data Processing Protocol

- Cooking show-style demo





princeton-mvpa-toolbox

Multi Voxel Pattern Analysis Matlab Library by The Princeton Neuroscience Institute

[Project Home](#) [Downloads](#) [Wiki](#) [Issues](#) [Source](#)

[Summary](#) [People](#)

Project Information

 +1 Recommend this on Google

 Starred by 17 users

[Project feeds](#)

Code license

[GNU GPL v2](#)

Labels

[mvpa](#), [matlab](#), [princeton](#),
[fMRI](#), [voxel](#), [multivoxel](#),
[machinelearning](#), [statistics](#),
[science](#), [multivariate](#),
[neuroscience](#), [psychology](#)

Members

[bdsinger](#), [gdetre](#),
[gmcgrathredbluepill](#), [knorman](#)
[15 committers](#)

Your role

[Committer](#)

Featured

Wiki pages

[AssociatedPublications](#)
[Show all »](#)

Links

External links

[Get the Library](#)
[MVPA Manual](#)

Groups

[MVPA Discussion Group](#)

Princeton Multi-Voxel Pattern Analysis (MVPA) Toolbox

The MVPA Toolbox is a set of Matlab tools to facilitate multi-voxel pattern analysis of fMRI neuroimaging data.

The aim is to create a set of open source functions in a widely-used language to facilitate exploration of multi-voxel pattern analysis techniques and to reduce the 'startup costs' for knowledgeable users eager to apply pattern classification algorithms to their imaging data. By developing the toolbox in the Matlab environment, users are able to take advantage of the vast array of existing functions. The data structures used and generated by the toolbox are designed to facilitate exploration and further script development.

You can read the archive or join the [public discussion mailing list](#).

- [Download the Toolbox and Data-sets](#)
- [Getting started](#)
- [Wiki documentation](#)
- [Read/post to the Google Groups mailing list](#)
- [2006 Pittsburgh EBC Competition](#)

<https://code.google.com/p/princeton-mvpa-toolbox/>



CCN Center for Cognitive Neuroscience



Multivariate Pattern Analysis in Python

[PyMVPA Home](#) | [Sitemap](#) »

[next](#) | [modules](#) | [modules](#) | [index](#)

PyMVPA is a **Python** package intended to ease statistical learning analyses of large datasets. It offers an extensible framework with a high-level interface to a broad range of algorithms for classification, regression, feature selection, data import and export. It is designed to integrate well with related software packages, such as **scikit-learn**, and **MDP**. While it is not limited to the neuroimaging domain, it is eminently suited for such datasets. PyMVPA is free software and requires nothing but free-software to run.

PyMVPA stands for **MultiVariate Pattern Analysis (MVPA)** in **Python**.



[Installation](#)



[Tutorial](#)



[Documentation](#)



[Support](#)

News



studyforrest @studyforrest

Your chance to win 5000 Euro by just doing good science!

[studyforrest.org](#) #openscience

🔄 Retweeted by PyMVPA Team

Expand

27 May

Table Of Contents

- News
- Contributing
- License
- How to cite PyMVPA
 - Peer-reviewed publications
 - Posters
- Authors and Contributors
- Acknowledgements
 - Grant support
- Similar or Related Projects

Next topic

[PyMVPA User Manual](#)

Quick links

- [Source download](#)
- [Code repository](#)
- [Bug tracker](#)
- [Mailing list archive](#)

[Who is using PyMVPA?](#)
[Dataset Archive](#)

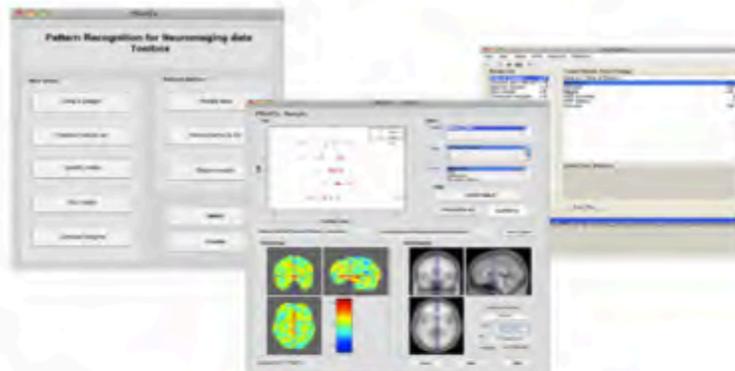
<http://www.pymvpa.org>

PRoNTTo Menu

- Introduction
- Software
- Documentation
- Courses
- Data sets
- Mailing list
- Credits



Pattern Recognition for Neuroimaging Toolbox (PRoNTTo)



PRoNTTo (Pattern Recognition for Neuroimaging Toolbox) is a software toolbox based on pattern recognition techniques for the analysis of neuroimaging data. Statistical pattern recognition is a field within the area of machine learning which is concerned with automatic discovery of regularities in data through the use of computer algorithms, and with the use of these regularities to take actions such as classifying the data into different categories. In PRoNTTo, brain scans are treated as spatial patterns and statistical learning models are used to identify statistical properties of the data that can be used to discriminate between experimental conditions or groups of subjects (classification models) or to predict a continuous measure (regression models).

<http://www.mlnl.cs.ucl.ac.uk/pronto/>

MVPA Data Processing Protocol

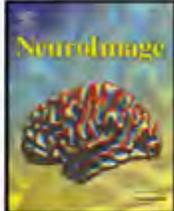
- A good starting reference:

NeuroImage 45 (2009) S199–S209

Contents lists available at [ScienceDirect](#)



NeuroImage



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

Machine learning classifiers and fMRI: A tutorial overview

Francisco Pereira ^{a,*}, Tom Mitchell ^b, Matthew Botvinick ^a

^a Princeton Neuroscience Institute/Psychology Department, Princeton University, Princeton, NJ 08540, USA
^b Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA

ARTICLE INFO	ABSTRACT
<p><i>Article history:</i> Received 15 September 2008 Accepted 18 November 2008 Available online 21 November 2008</p>	<p>Interpreting brain image experiments requires analysis of complex, multivariate data. In recent years, one analysis approach that has grown in popularity is the use of machine learning algorithms to train classifiers to decode stimuli, mental states, behaviours and other variables of interest from fMRI data and thereby show the data contain information about them. In this tutorial overview we review some of the key choices faced in using this approach as well as how to derive statistically significant results, illustrating each point from a case study. Furthermore, we show how, in addition to answering the question of 'is there information about a variable of interest' (pattern discrimination), classifiers can be used to tackle other classes of question, namely 'where is the information' (pattern localization) and 'how is that information encoded' (pattern characterization).</p>

© 2008 Elsevier Inc. All rights reserved.

MVPA Data Processing Protocol

- Preparing the data (preprocessing)
 - Remove signal artifacts
 - Detrend each run
 - High-pass filter each run
 - Z-score data from each run
- Parsing the data into “examples”
 - Block designs
 - Average timepoints from each block (↑ signal stability; ↓ examples)
 - Or treat each timepoint as an “independent” example
 - Event-related designs
 - Choose single post-stimulus timepoint from each trial or average several
 - Need long inter-trial intervals to prevent hemodynamic overlap

*** Always important to balance the number of examples from each condition ***

Pattern Analysis Approach: *Temporal Selection*

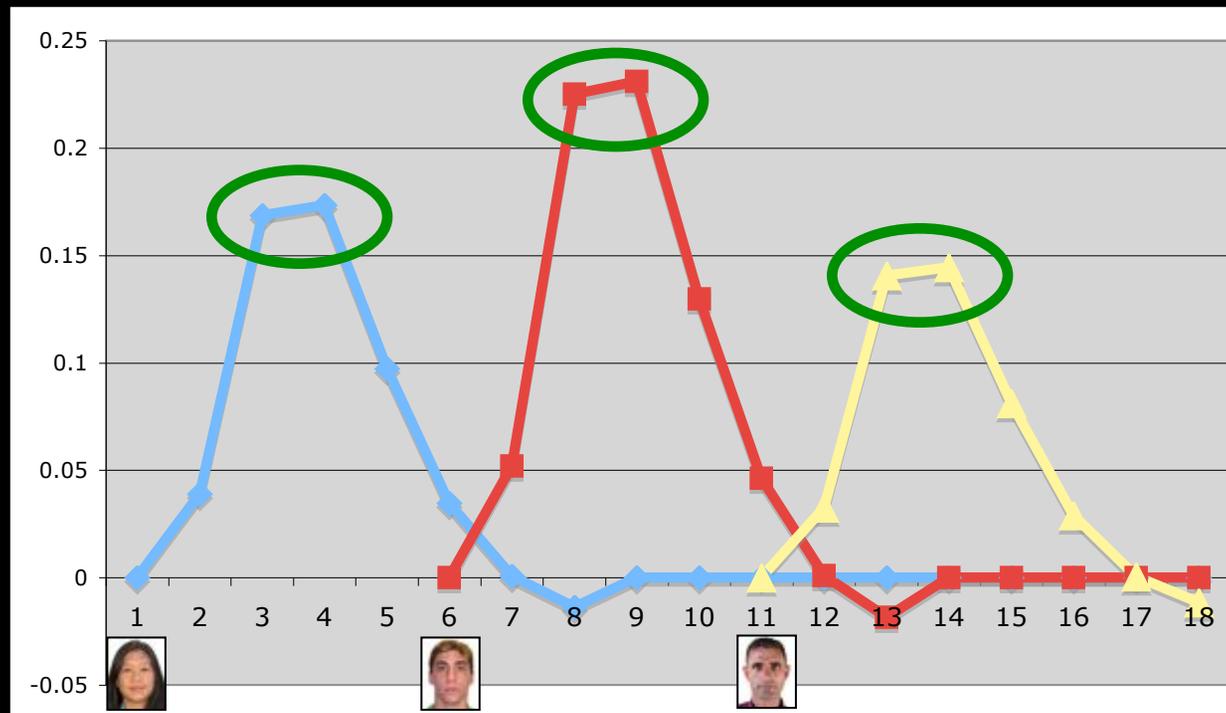
- Reduce full fMRI timeseries from having 5 values (TRs) per trial to having only 1 value per trial



- Average 3rd and 4th TR of each trial (e.g., 4-8 sec post-stimulus)

Pattern Analysis Approach: *Temporal Selection*

- Reduce full fMRI timeseries from having 5 values (TRs) per trial to having only 1 value per trial

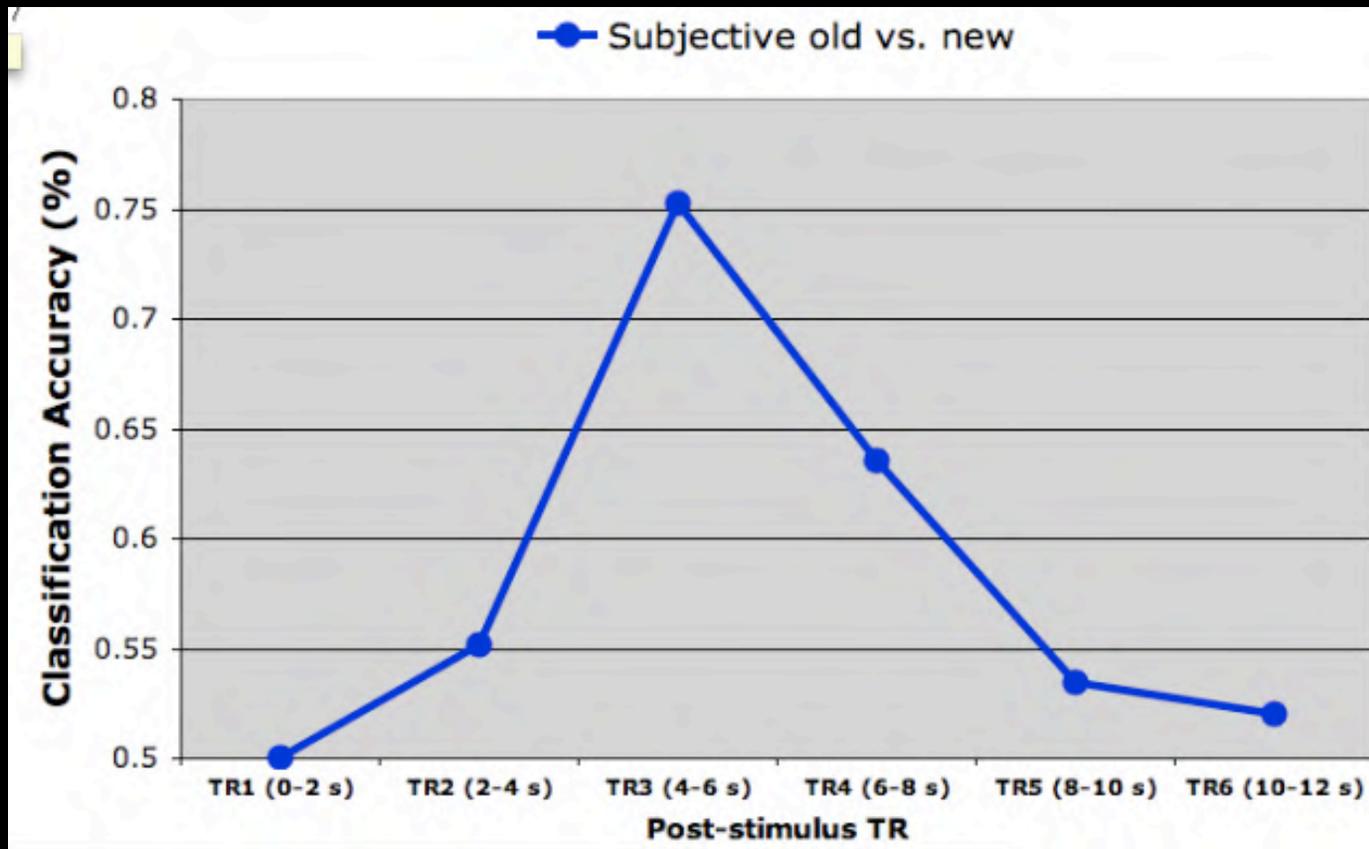


- Or run a new GLM that estimates a single parameter for each trial (i.e, beta-series approach; Rissman et al. 2004)

Pattern Analysis Approach: *Temporal Selection*

Alternative strategy

- Train and test separate classifiers using data from each post-stimulus TR

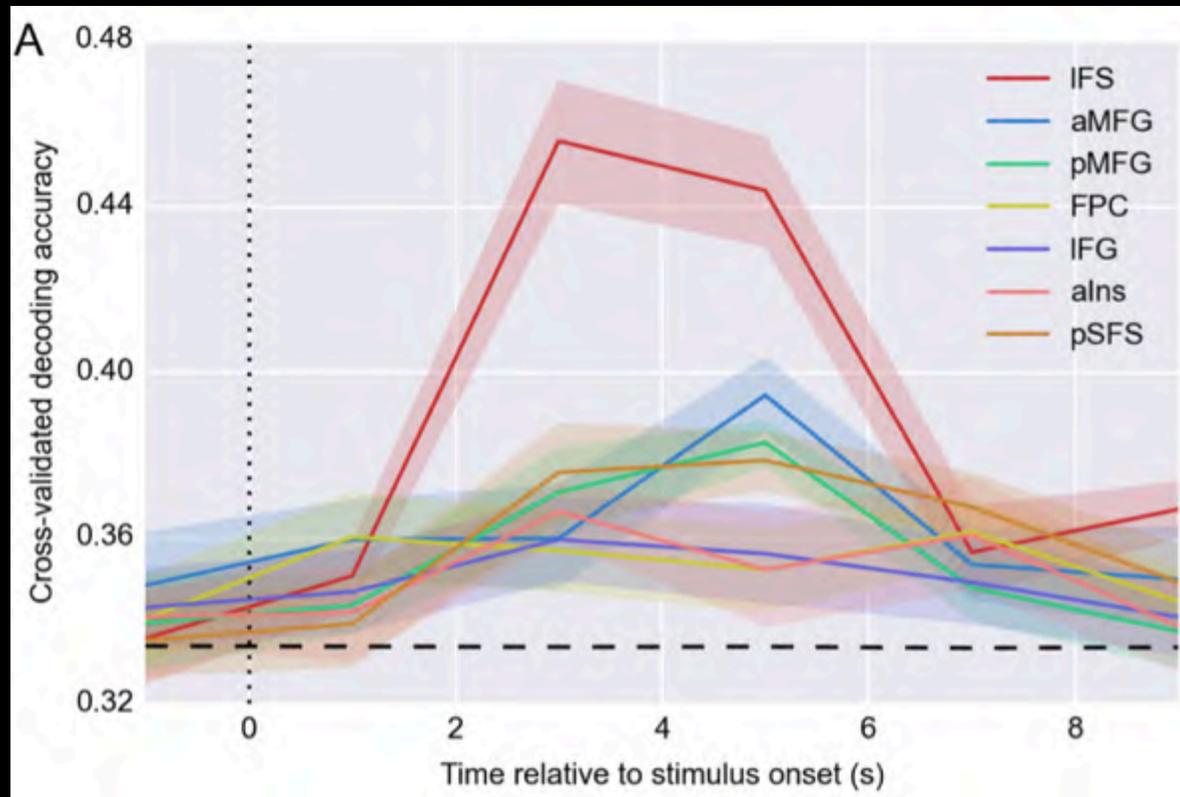


Classification timecourse for one subject

Pattern Analysis Approach: *Temporal Selection*

Alternative strategy

- Train and test separate classifiers using data from each post-stimulus TR

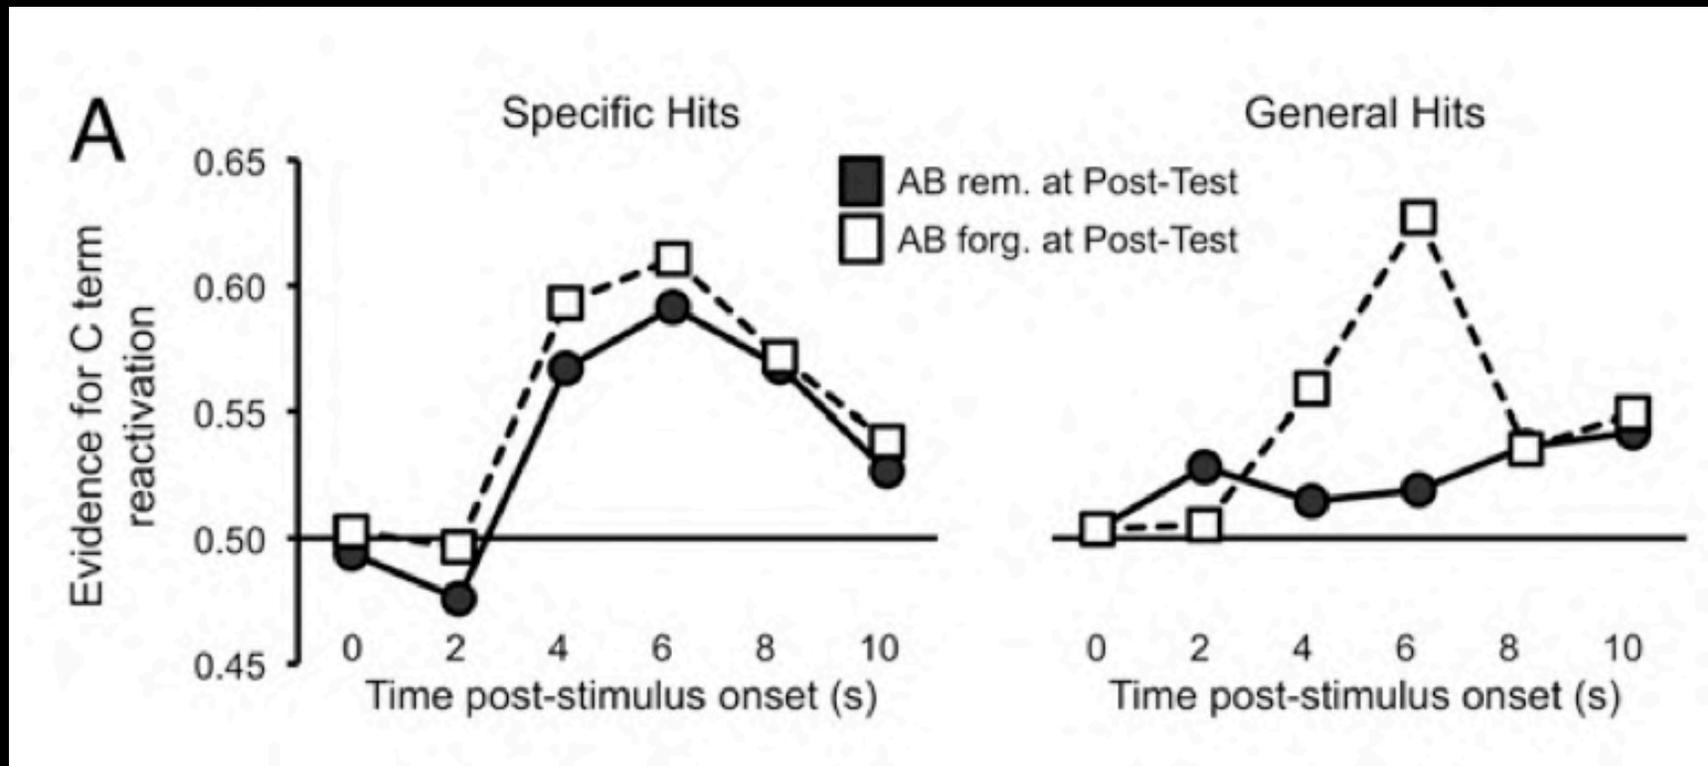


Waskom, Kumaran, Gordon, Rissman, & Wagner (in press) *J Neurosci*

Pattern Analysis Approach: *Temporal Selection*

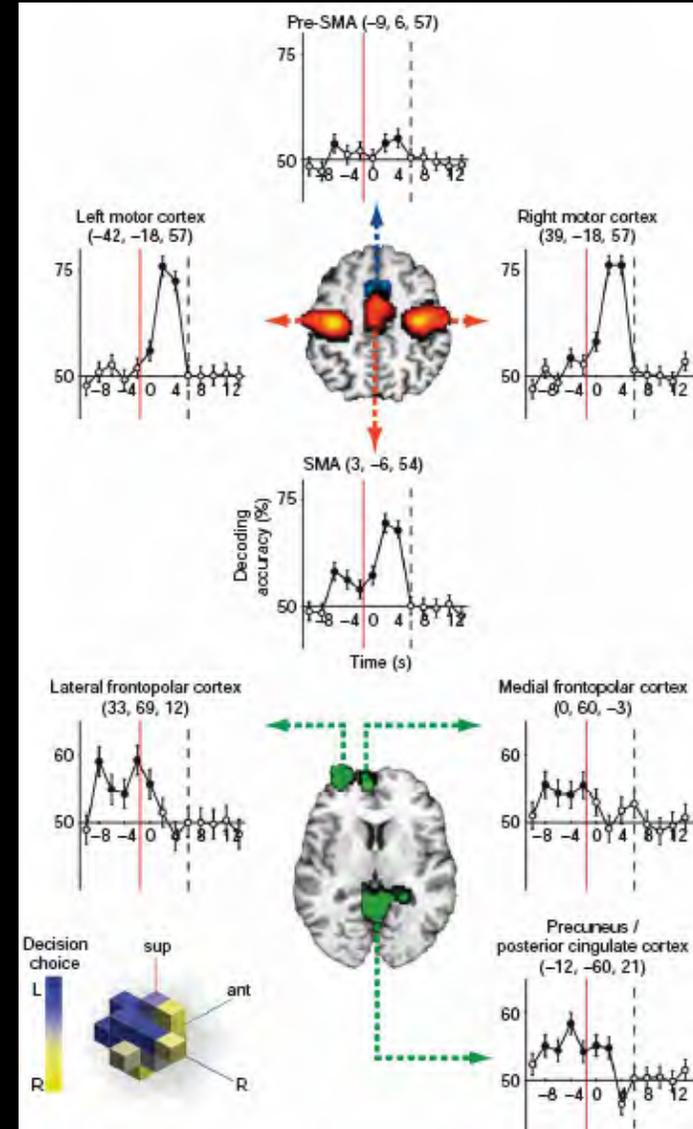
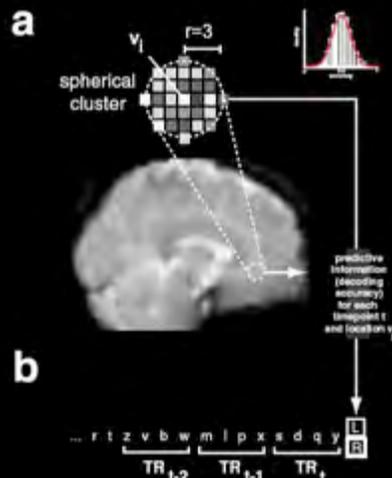
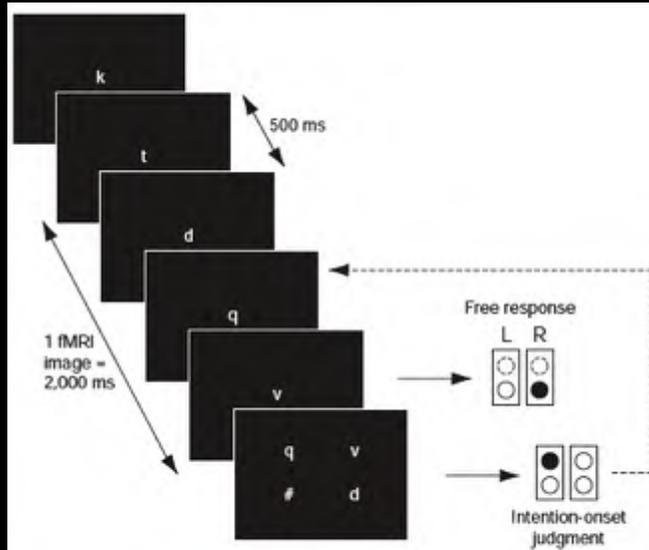
Alternative strategy

- Train and test separate classifiers using data from each post-stimulus TR



Kuhl, Rissman, Chun, & Wagner (2011) *PNAS*

Unraveling the Notion of Free Will?



Soon et al. (2008) *Nature Neuroscience*

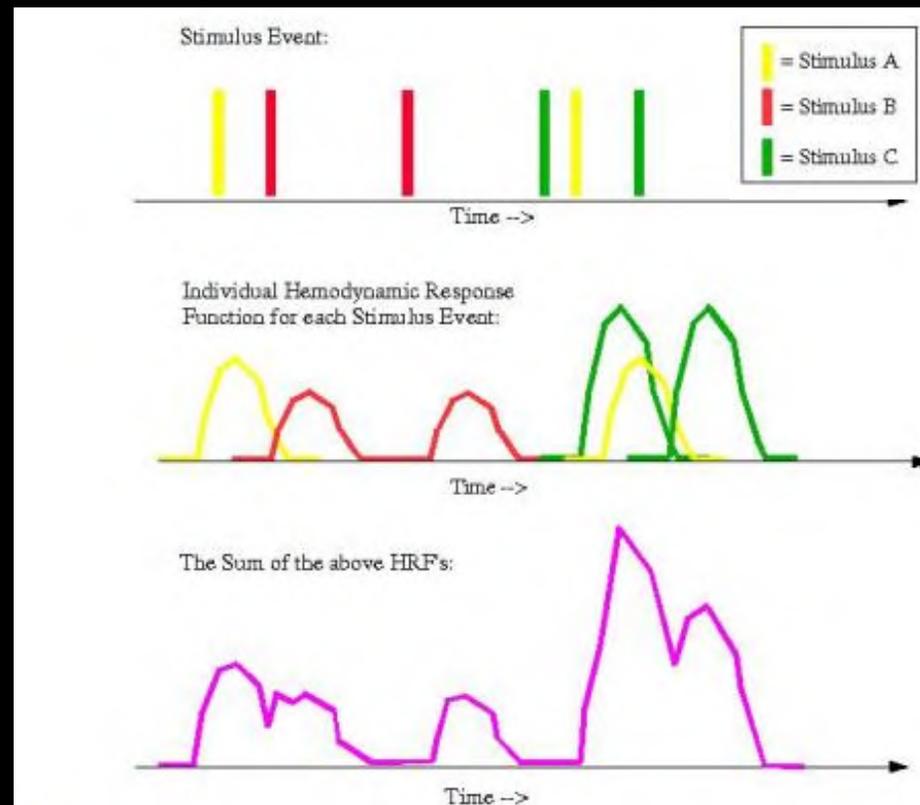
Pattern Analysis Approach: *Temporal Selection*

- NOTE: the previous examples assume that you are working with a slow event-related design (i.e., widely spaced trials with minimally-overlapping HRFs)
- What about rapid event-related designs?

Pattern Analysis Approach: *Temporal Selection*

One approach:

- For each TR, examine your design matrix and determine which condition has the maximal predicted activity
- Specify a threshold to exclude to ambiguous TRs



Pattern Analysis Approach:

Temporal Selection

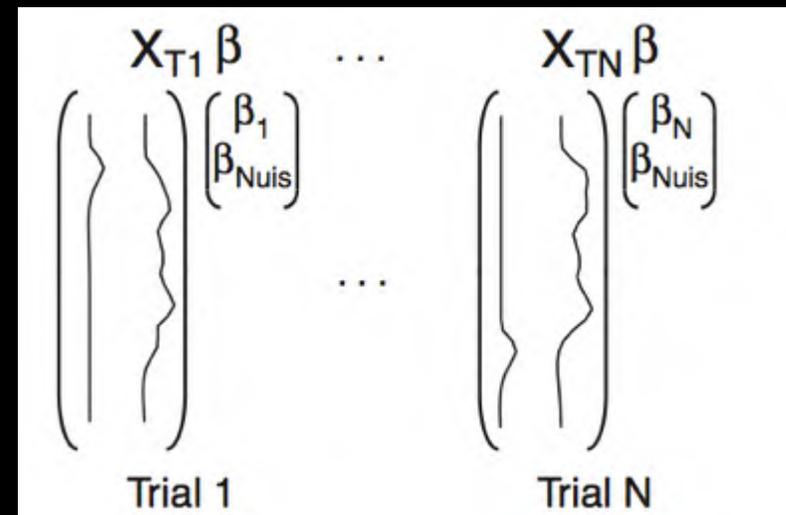
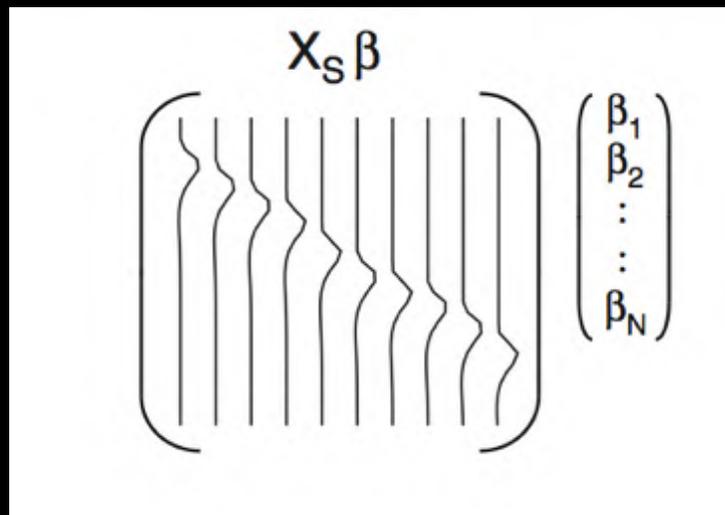
Another approach (Rissman et al., 2004):

- Estimate each trial's activity through a univariate GLM including one unique regressor for each event

Pattern Analysis Approach: *Temporal Selection*

A newer variant of this approach (Mumford et al., 2012):

- Estimate each trial's activity through a univariate GLM including one regressor for that trial as well as another regressor for all other trials.
 - Like beta series estimation approach, but involves running **many separate GLMs** (# of GLMs = # of trials)





Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses

Jeanette A. Mumford ^{a,*}, Benjamin O. Turner ^b, F. Gregory Ashby ^b, Russell A. Poldrack ^c

^a Department of Psychology, University of Texas, Austin, TX 78759, USA

^b Department of Psychological and Brain Sciences, University of California, Santa Barbara, CA 93016, USA

^c Departments of Psychology and Neurobiology and Imaging Research Center, University of Texas, Austin, TX 78759, USA

ARTICLE INFO

Article history:

Received 27 May 2011

Revised 6 August 2011

Accepted 23 August 2011

Available online 5 September 2011

Keywords:

Functional magnetic resonance imaging

Classification analysis

MVPA

Beta series estimation

Rapid event-related design

ABSTRACT

Use of multivoxel pattern analysis (MVPA) to predict the cognitive state of a subject during task performance has become a popular focus of fMRI studies. The input to these analyses consists of activation patterns corresponding to different tasks or stimulus types. These activation patterns are fairly straightforward to calculate for blocked trials or slow event-related designs, but for rapid event-related designs the evoked BOLD signal for adjacent trials will overlap in time, complicating the identification of signal unique to specific trials. Rapid event-related designs are often preferred because they allow for more stimuli to be presented and subjects tend to be more focused on the task, and thus it would be beneficial to be able to use these types of designs in MVPA analyses. The present work compares 8 different models for estimating trial-by-trial activation patterns for a range of rapid event-related designs varying by interstimulus interval and signal-to-noise ratio. The most effective approach obtains each trial's estimate through a general linear model including a regressor for that trial as well as another regressor for all other trials. Through the analysis of both simulated and real data we have found that this model shows some improvement over the standard approaches for obtaining activation patterns. The resulting trial-by-trial estimates are more representative of the true activation magnitudes, leading to a boost in classification accuracy in fast event-related designs with higher signal-to-noise. This provides the potential for fMRI studies that allow simultaneous optimization of both univariate and MVPA approaches.

© 2011 Elsevier Inc. All rights reserved.

Implementation of this approach is now a built-in feature in AFNI:
http://afni.nimh.nih.gov/pub/dist/doc/program_help/3dLSS.html

Pattern Analysis Approach:

Temporal Selection

One more commonly used approach:

- Run a standard univariate GLM analysis to derive condition-specific beta estimates for each scanning run
- Use these beta images as your patterns for classification
 - **Problem:**
 - If you only have 6 runs, then at best you'd only have 5 training examples for each condition
 - **Work-around:**
 - Subdivide each actual run into 2 or 3 mini-runs, and then run univariate GLM
 - More beta images to work with!



ELSEVIER

Contents lists available at SciVerse ScienceDirect

NeuroImage

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



The advantage of brief fMRI acquisition runs for multi-voxel pattern detection across runs

Marc N. Coutanche*, Sharon L. Thompson-Schill

Department of Psychology, University of Pennsylvania, 3720 Walnut Street, Philadelphia, PA 19104, USA

ARTICLE INFO

Article history:

Accepted 25 March 2012

Available online 3 April 2012

Keywords:

MVPA

fMRI

Multivariate

Runs

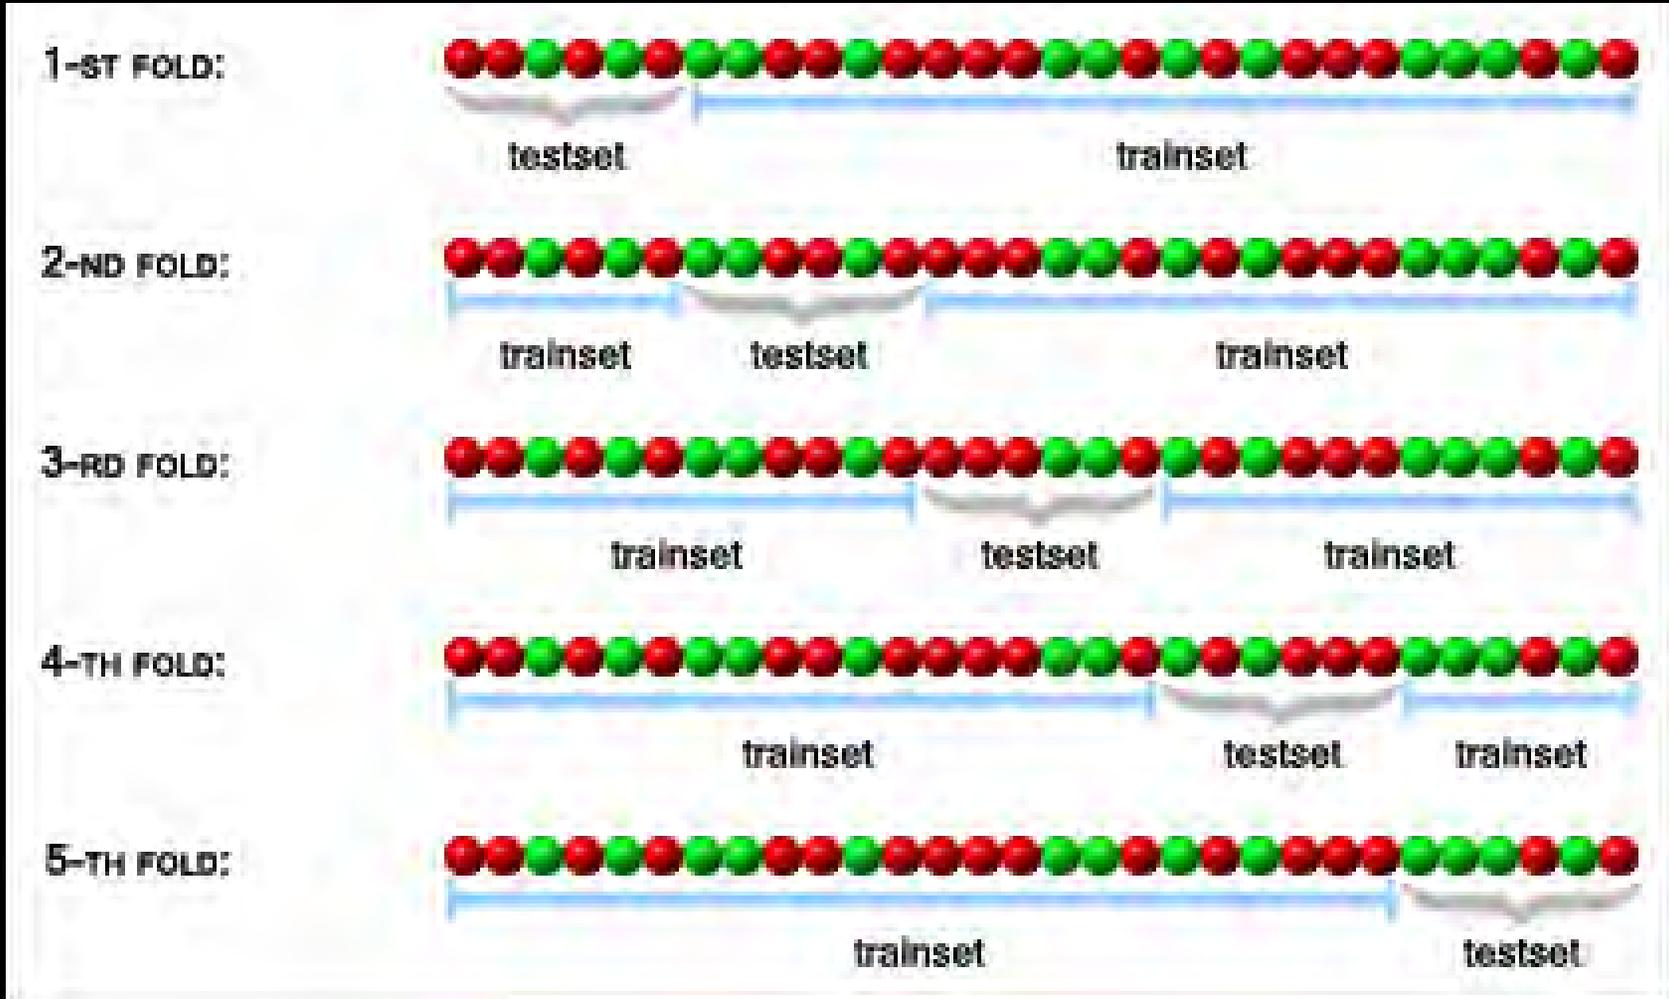
Design

Classification

ABSTRACT

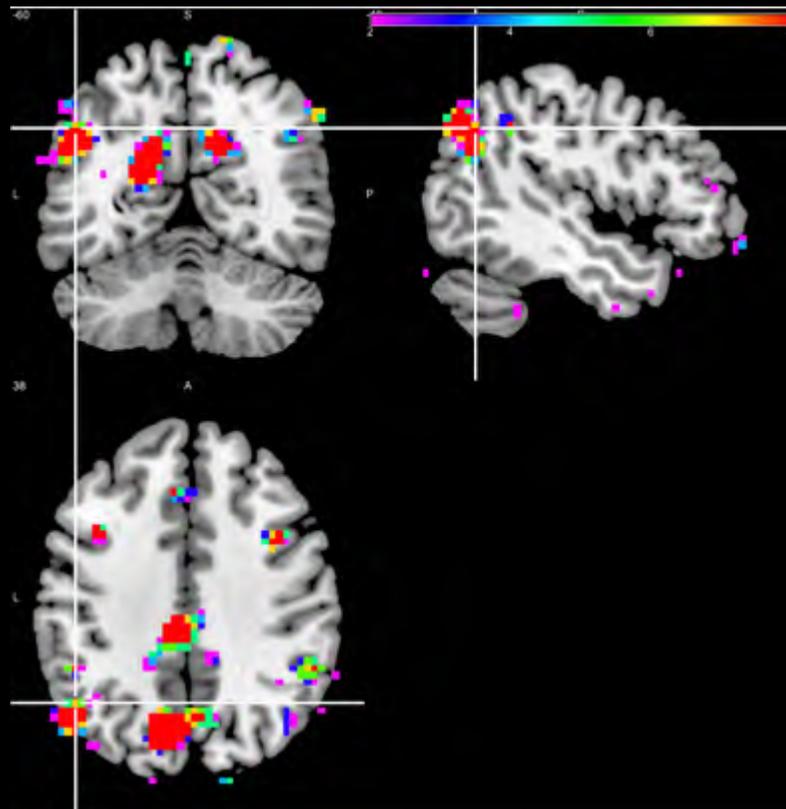
Functional magnetic resonance imaging (fMRI) studies are broken up into runs (or 'sessions'), frequently selected to be long to minimize across-run signal variations. For investigations that use multi-voxel pattern analysis (MVPA), however, employing many short runs might improve a classifier's ability to generalize across irrelevant pattern variations and detect condition-related activity patterns. We directly tested this hypothesis by scanning participants with both long and short runs and comparing MVPA performance using data from each set of runs. Every run included presentations of faces, places, man-made objects and fruit in a blocked 1-back design. MVPA performance significantly improved from using a large number of short runs, compared to several long runs, in across-run classifications with identical amounts of data. Superior classification was found across variations in the classifier employed, feature selection procedure and region of interest. Performance improvements also extended to an information brain mapping 'searchlight' procedure. These results suggest that investigators looking to maximize the detection of subtle multi-voxel patterns across runs might consider employing short fMRI runs.

Dividing up the data: *Cross-validation*



Feature Selection

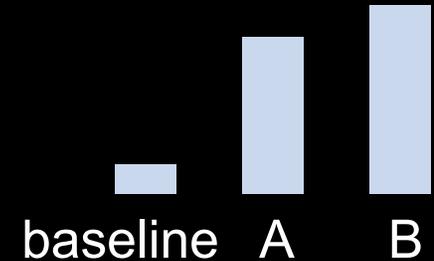
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing data



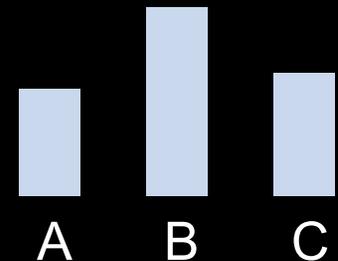
Feature Selection

- What criteria should define important voxels?

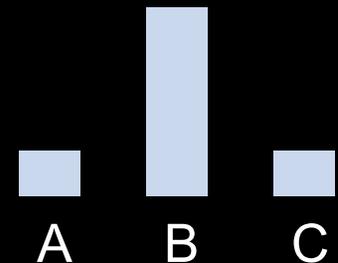
- difference from baseline



- difference between classes (e.g. ANOVA)

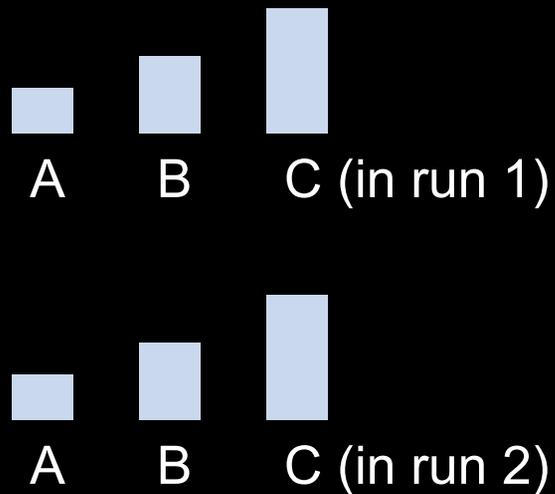


- preferential response to one class



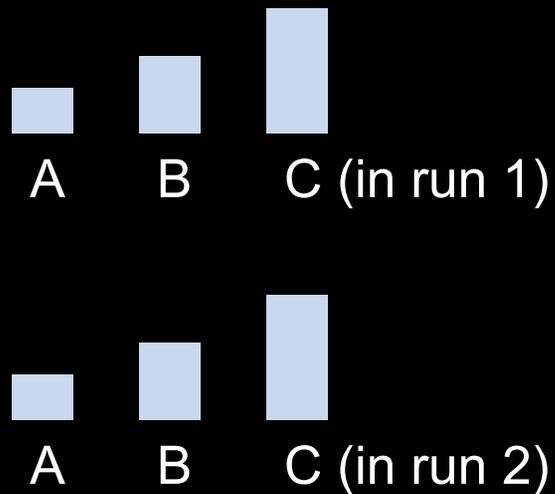
Feature Selection

- What criteria should define important voxels?
 - stability (i.e., across scanning runs)



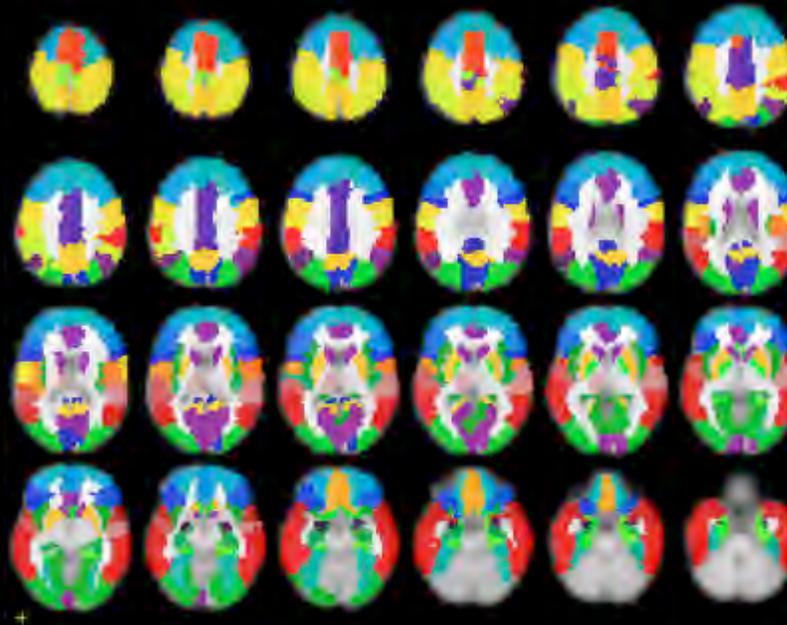
Feature Selection

- What criteria should define important voxels?
 - stability (i.e., across scanning runs)



Feature Selection

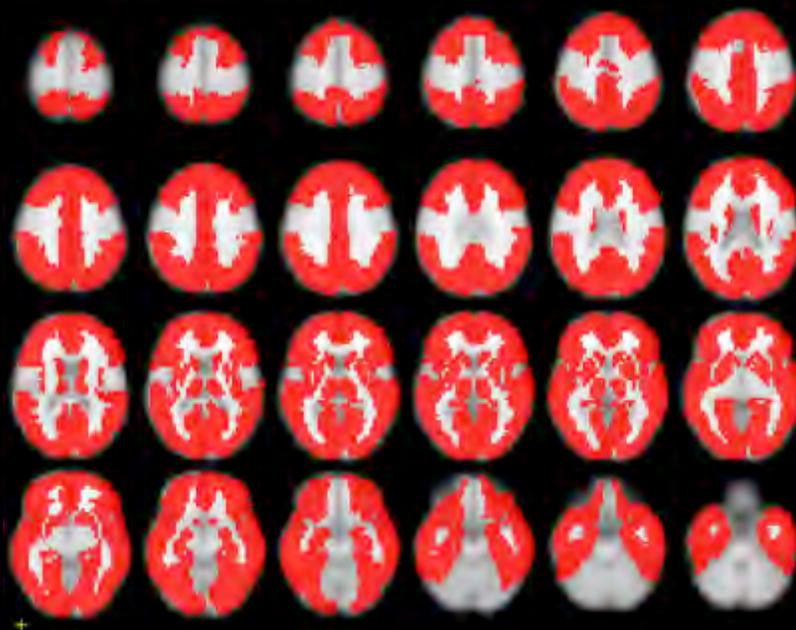
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing data
 - **Independently-defined ROIs**



ROIs from the
AAL atlas

Feature Selection

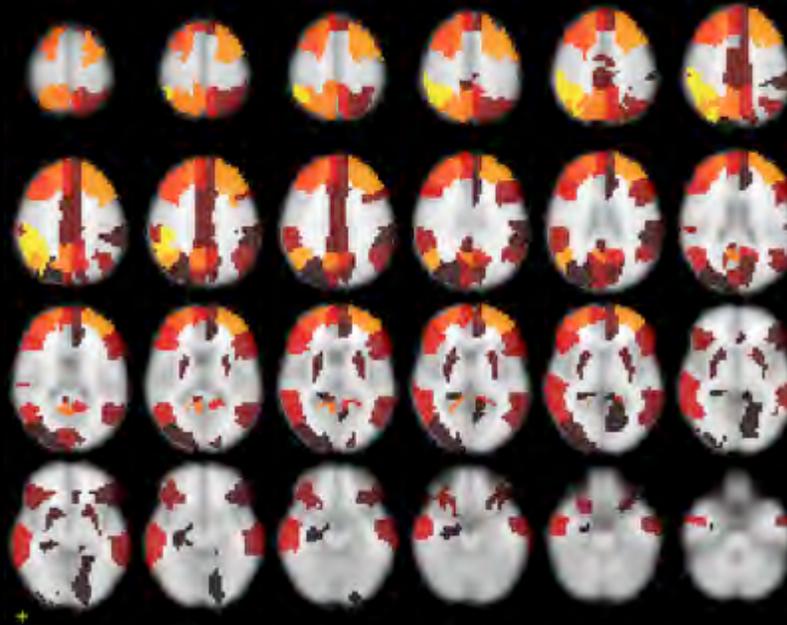
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing data
 - **Independently-defined ROIs**



You could combine all of these regions to make a large mask (in this case, excluding motor areas, white matter, and CSF)

Feature Selection

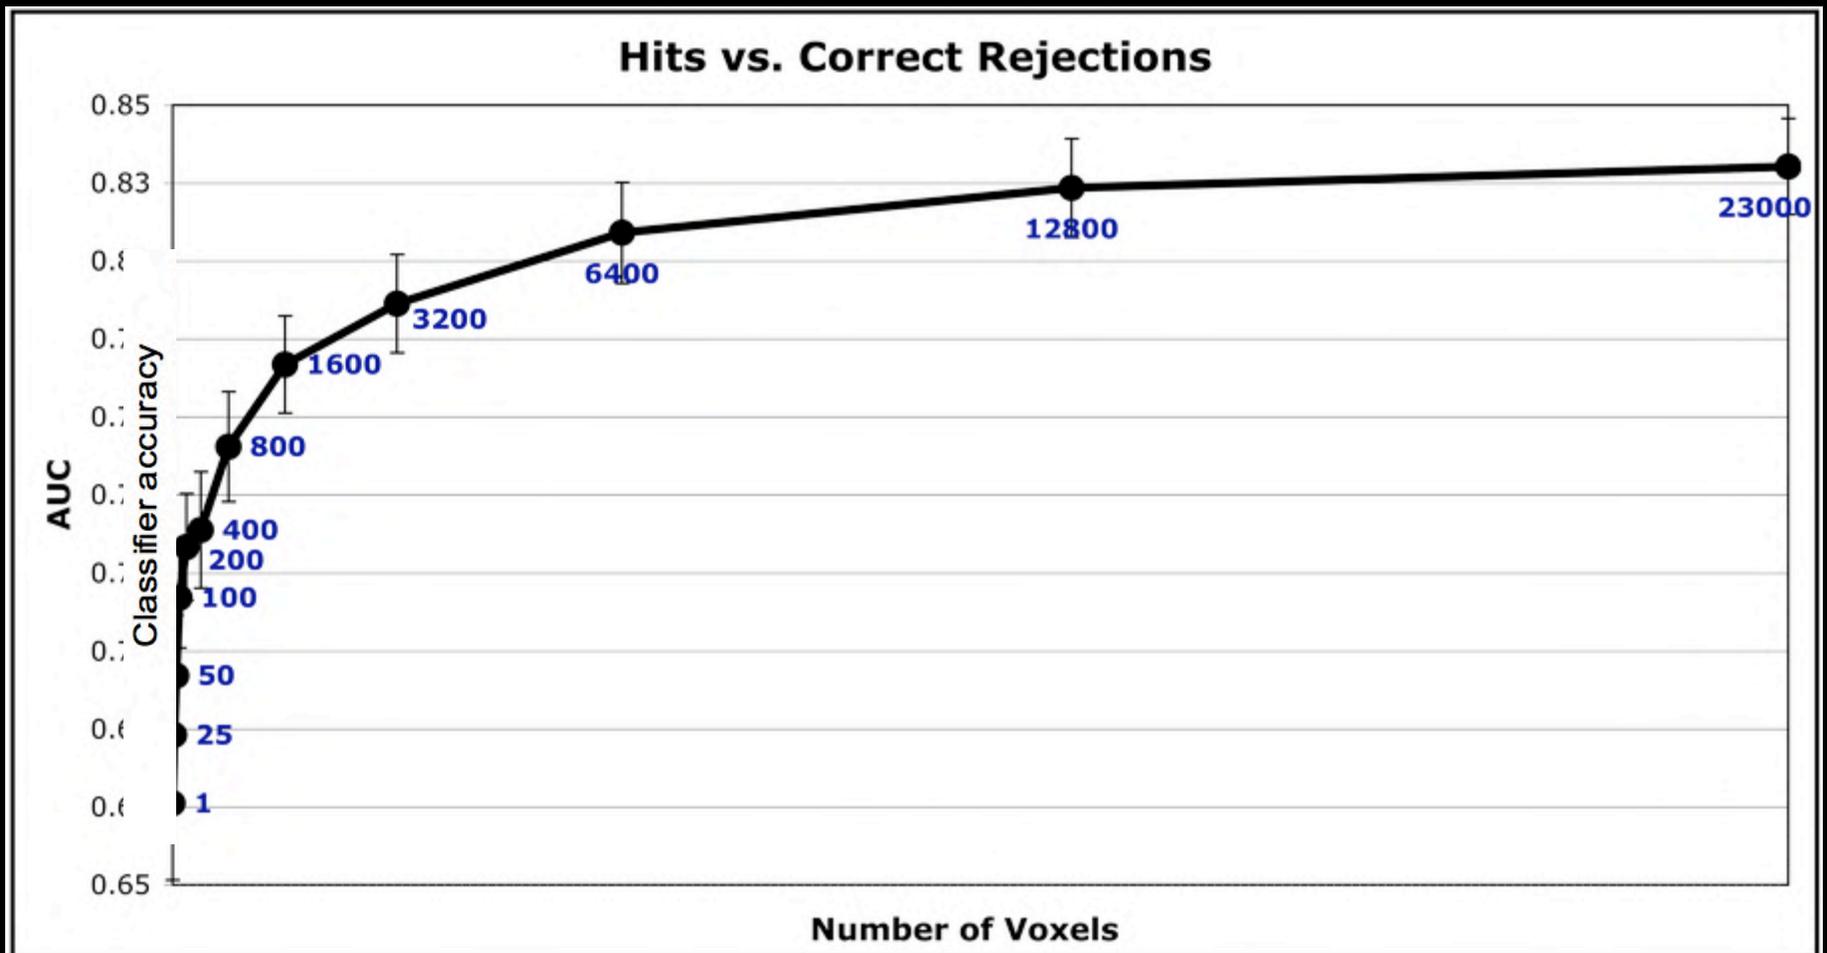
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing data
 - **Independently-defined ROIs**



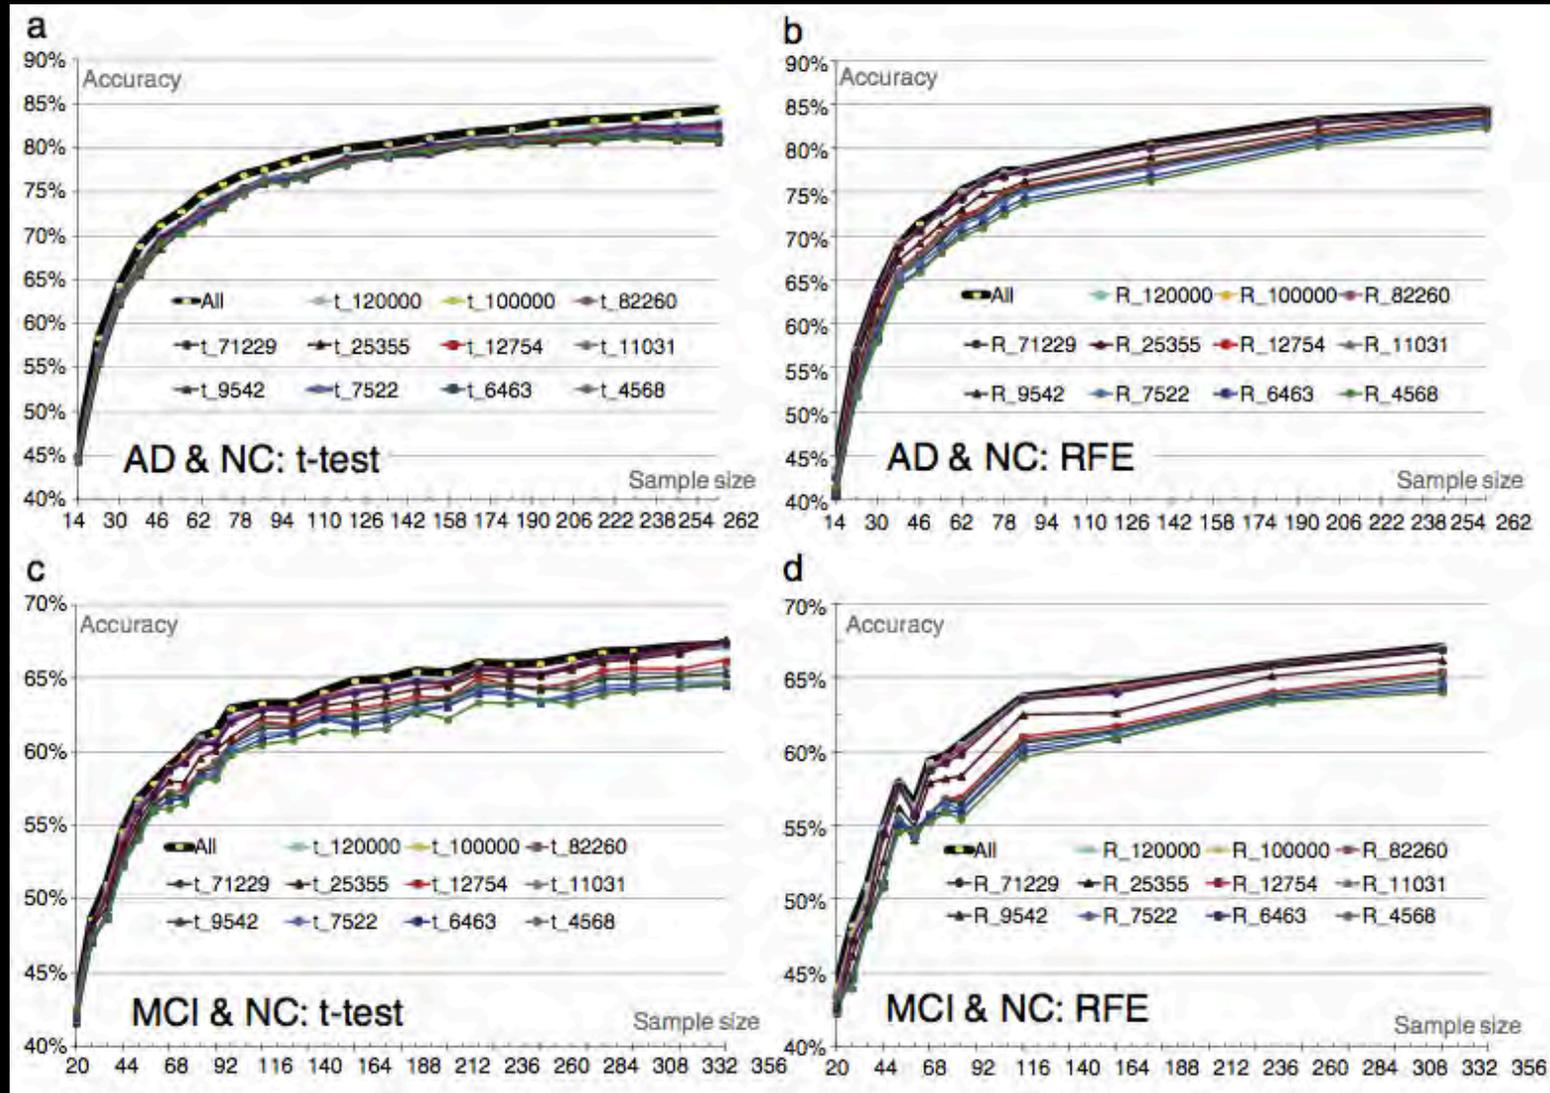
You also can
compute
classification
performance
within each ROI

How many features (voxels) to use?

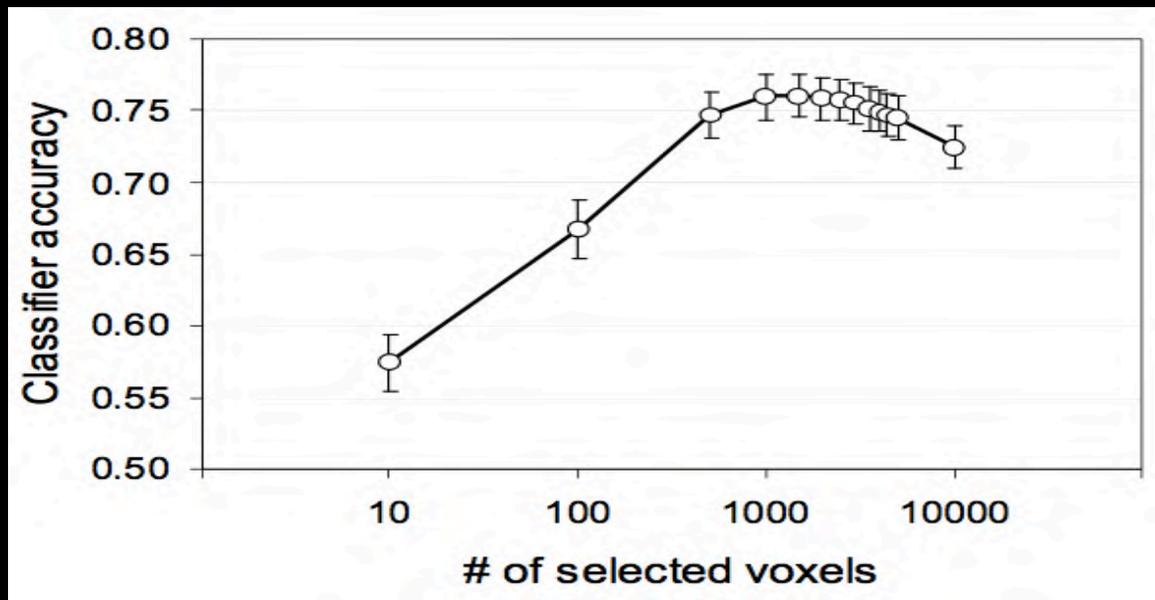
- Classification performance as a function of the number of voxels used by classifier (*ANOVA-based selection*)



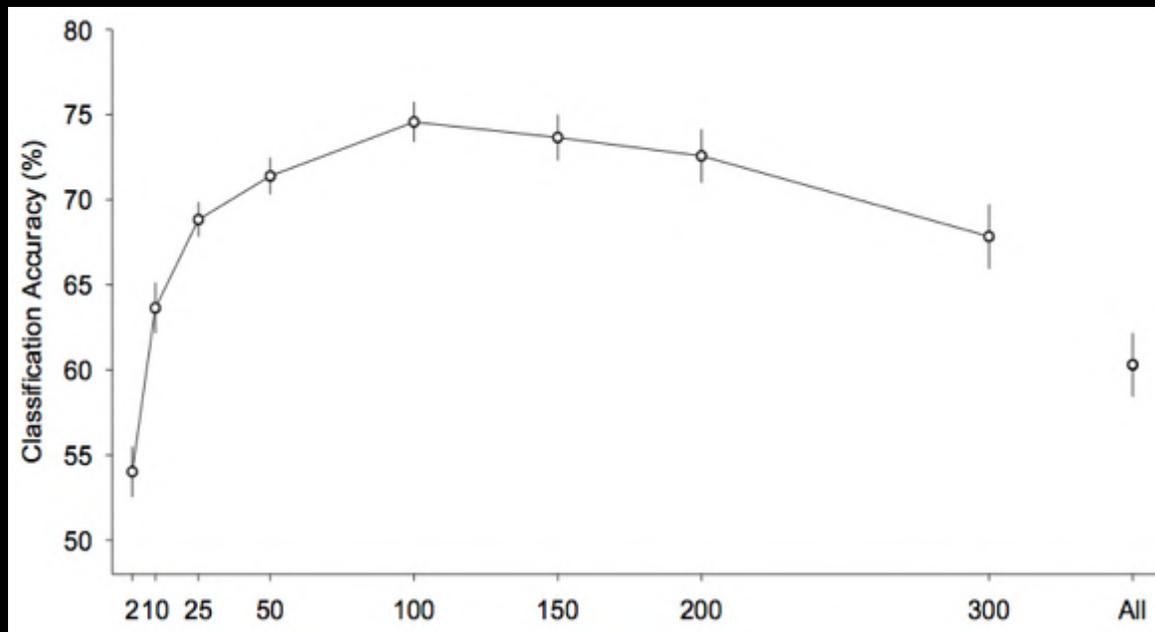
Feature selection doesn't always improve classification accuracy



Peak performance with ~1000 voxels
Johnson et al. (2009)

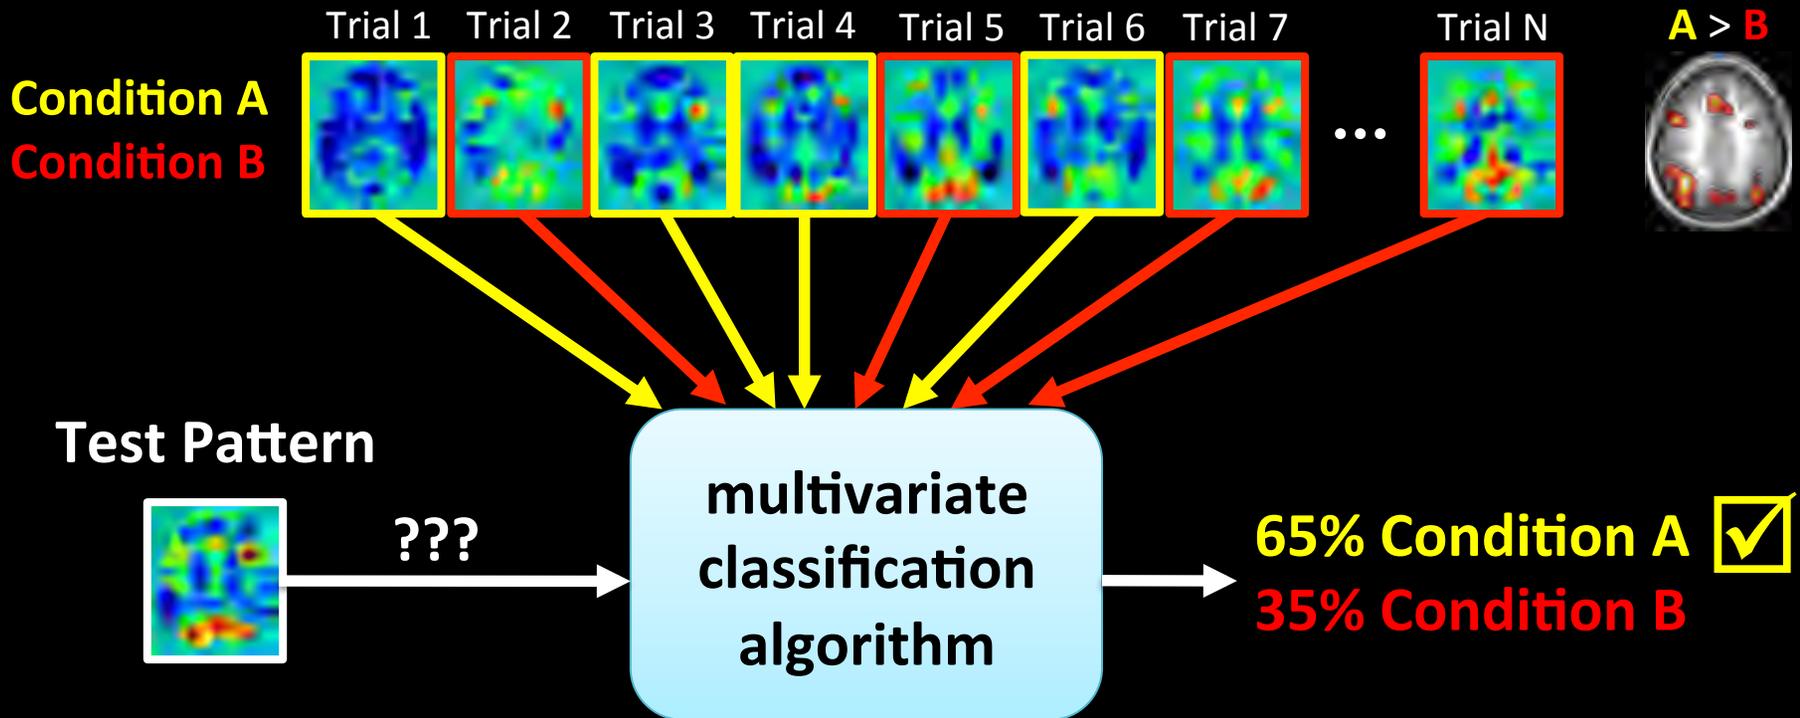


Peak performance with ~100 voxels
Esterman et al. (2009)



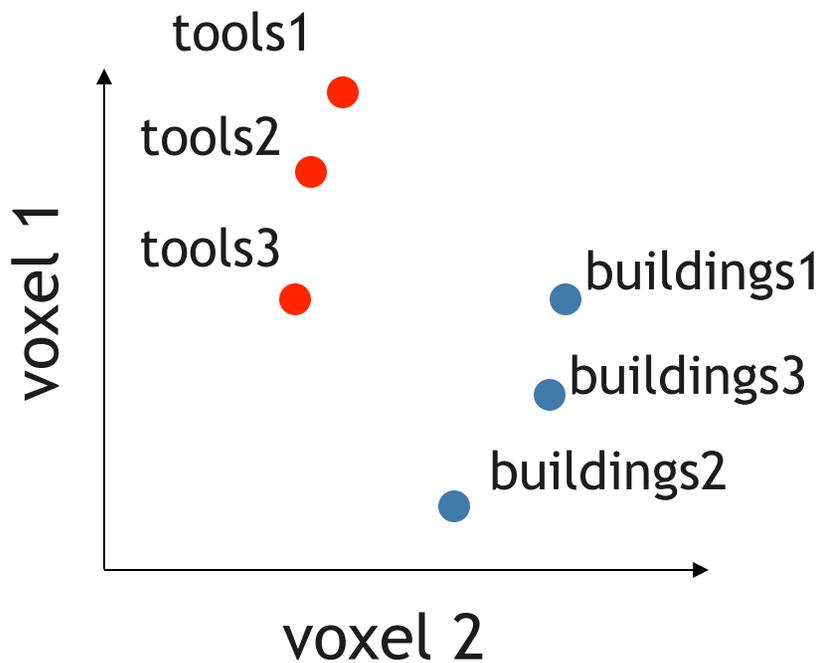
Training and testing the classifier

Training Patterns



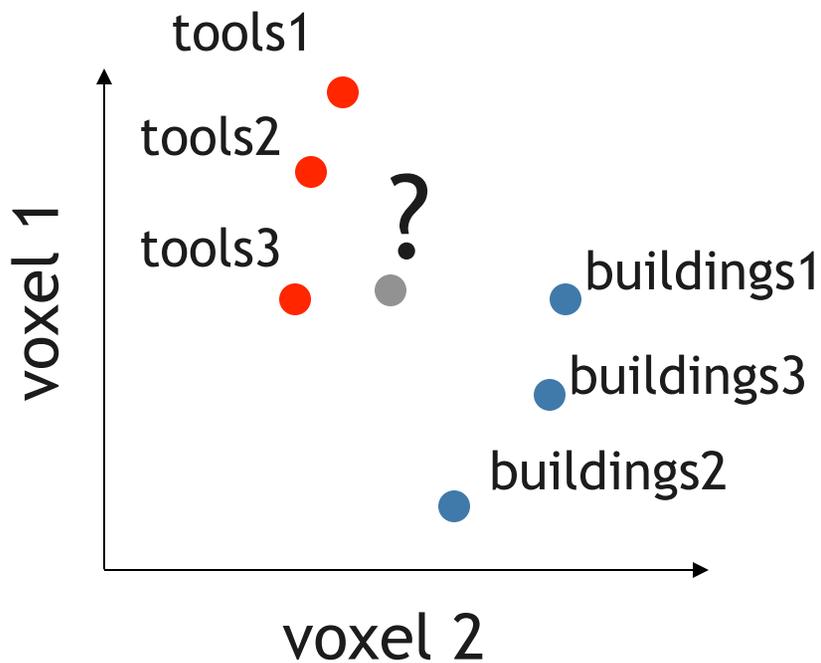
what is inside the box?

- simplest function is no function at all
- “nearest neighbour”



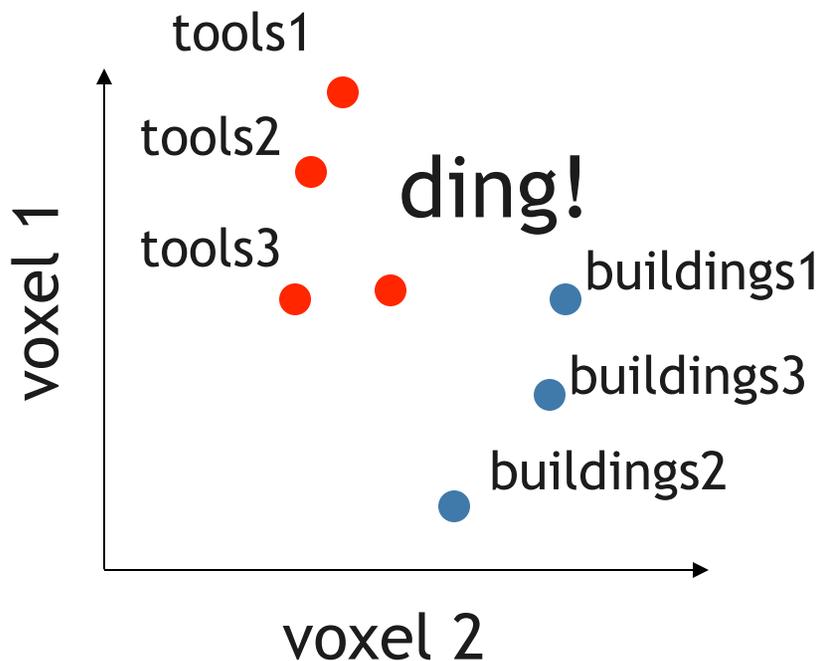
what is inside the box?

- simplest function is no function at all
- “nearest neighbour”



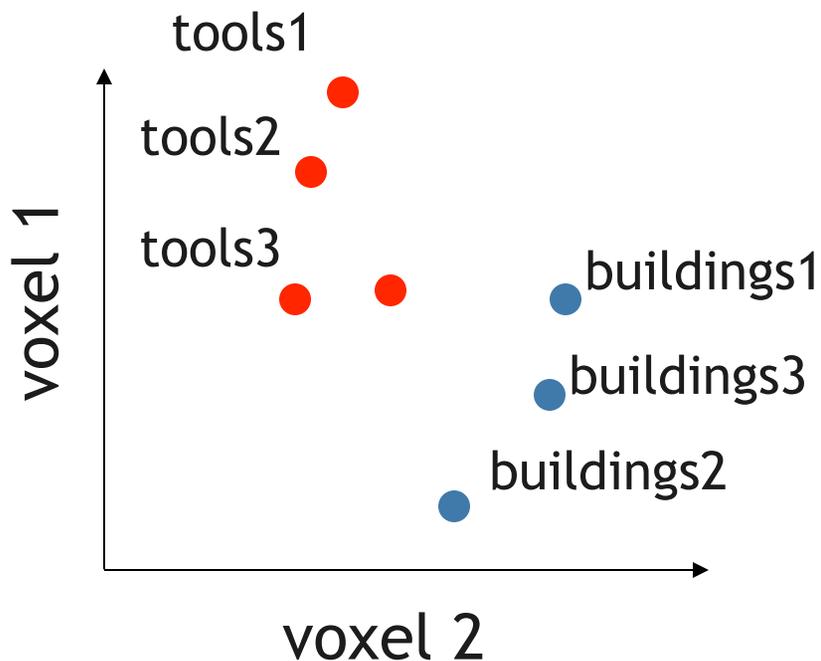
what is inside the box?

- simplest function is no function at all
- “nearest neighbour”

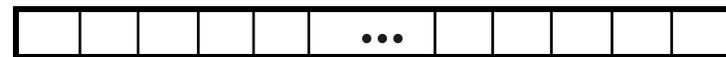


what is inside the box?

- simplest function is no function at all
- “nearest neighbour”



requires example similarity measure

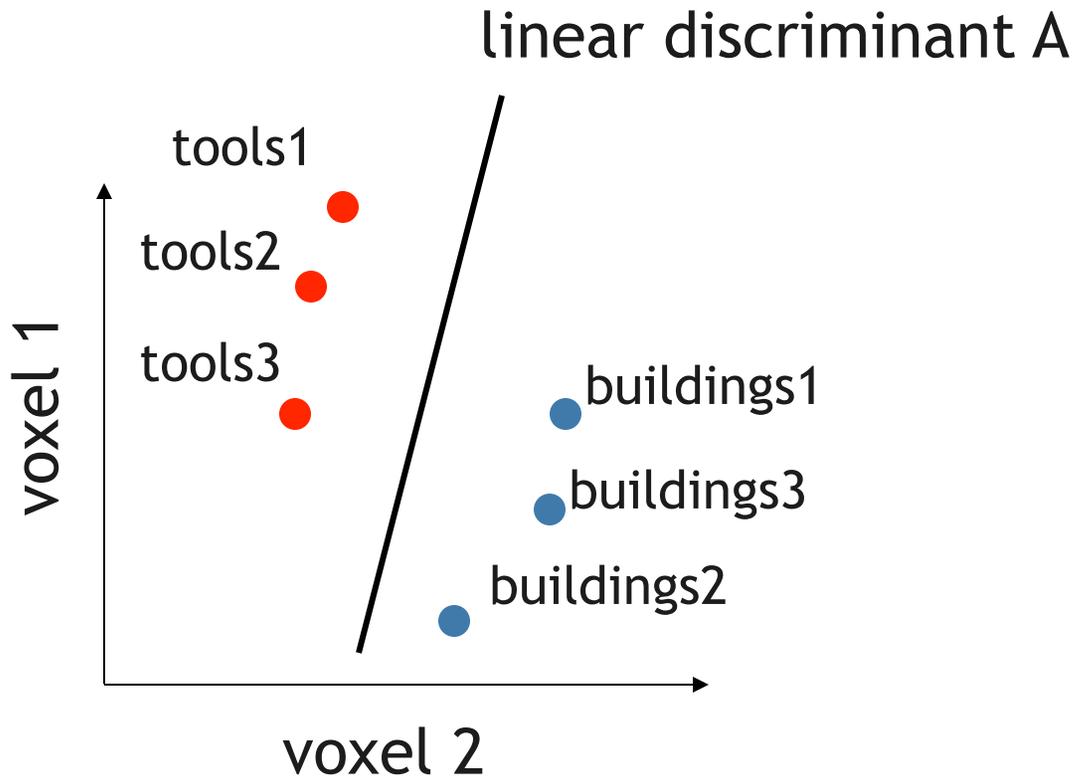


Euclidean dist., correlation, ...



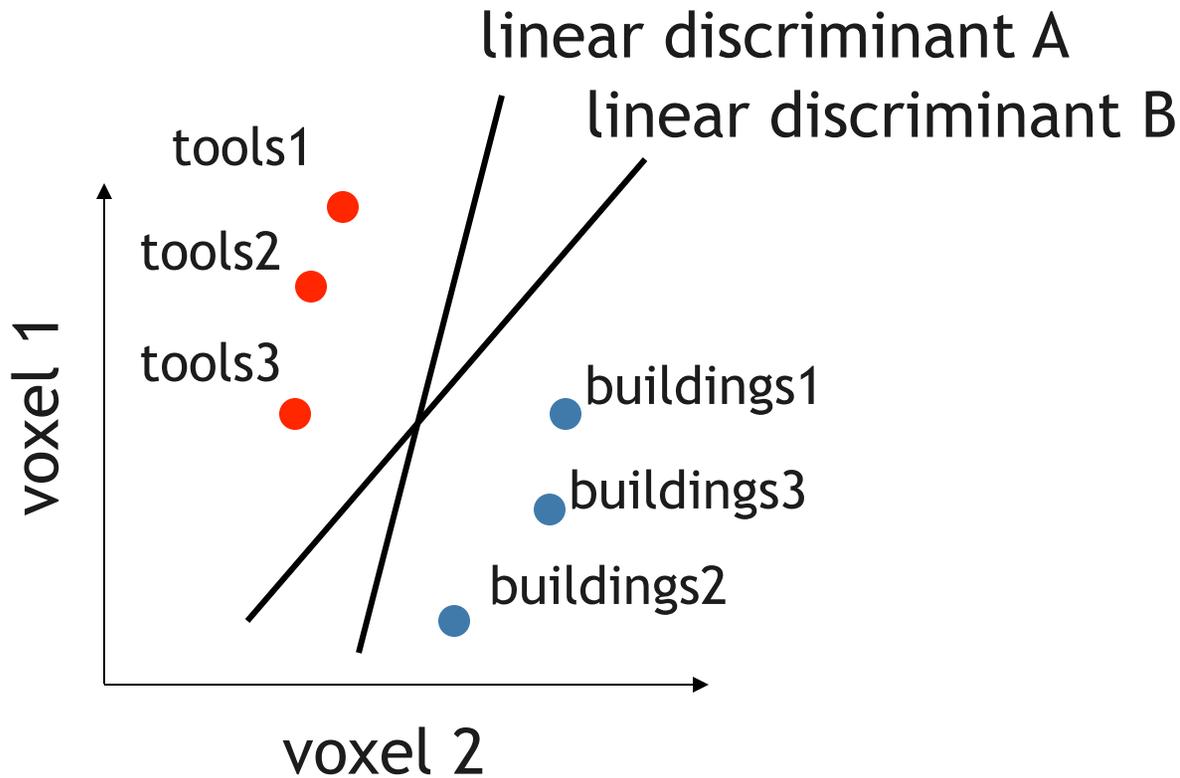
what is inside the box?

- next simplest: learn linear discriminant



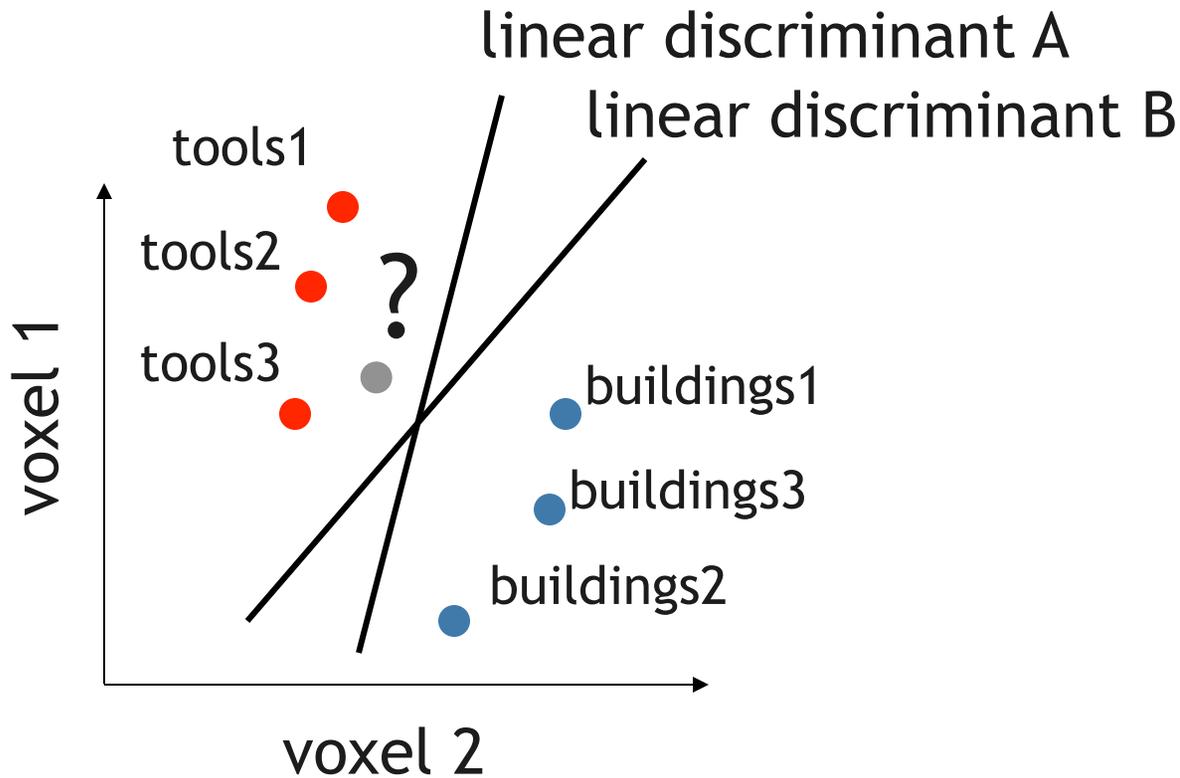
what is inside the box?

- next simplest: learn linear discriminant
- note that there are many solutions...



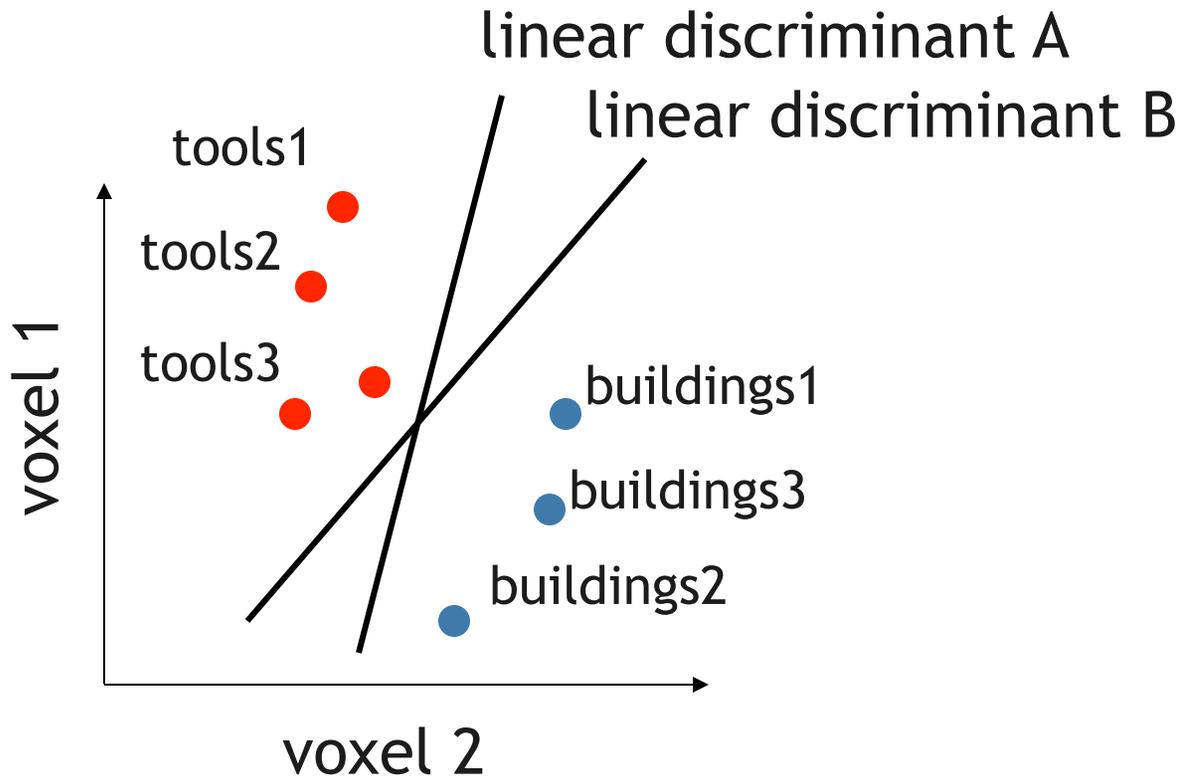
what is inside the box?

- next simplest: learn linear discriminant
- note that there are many solutions...

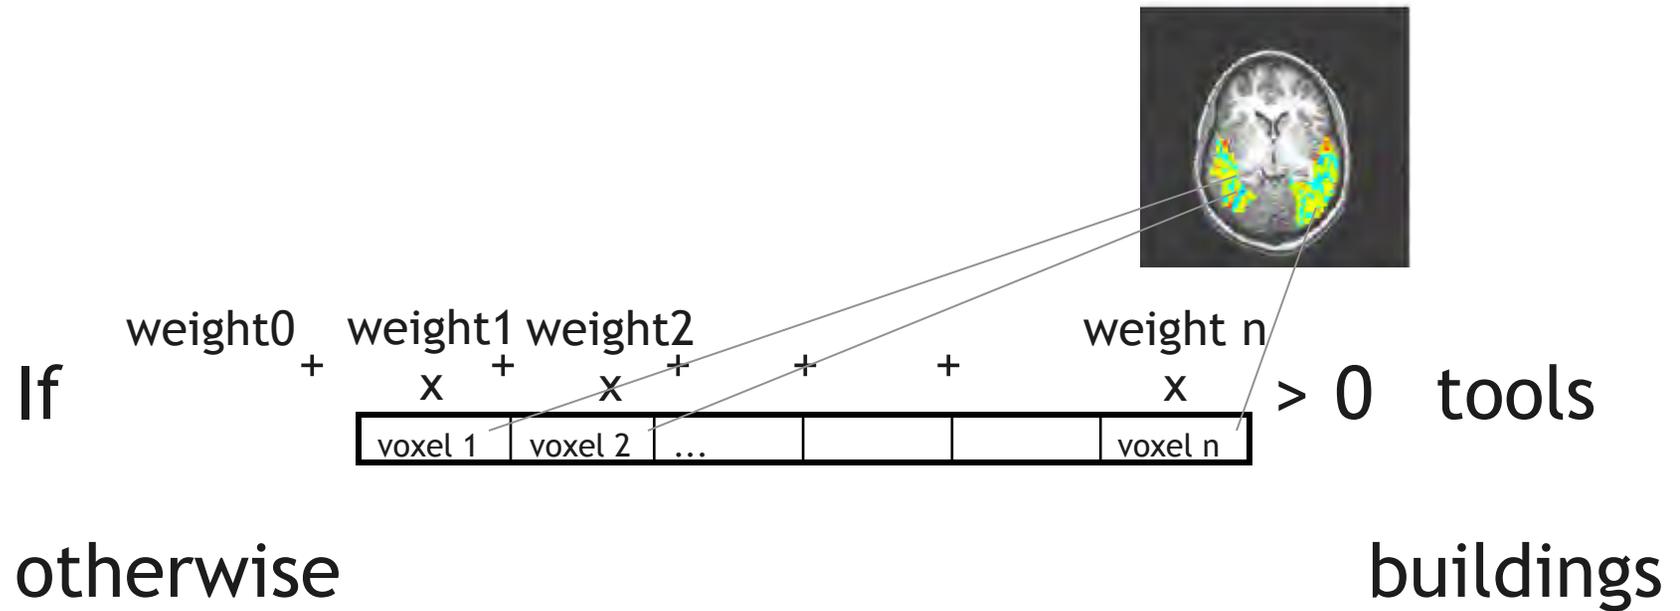


what is inside the box?

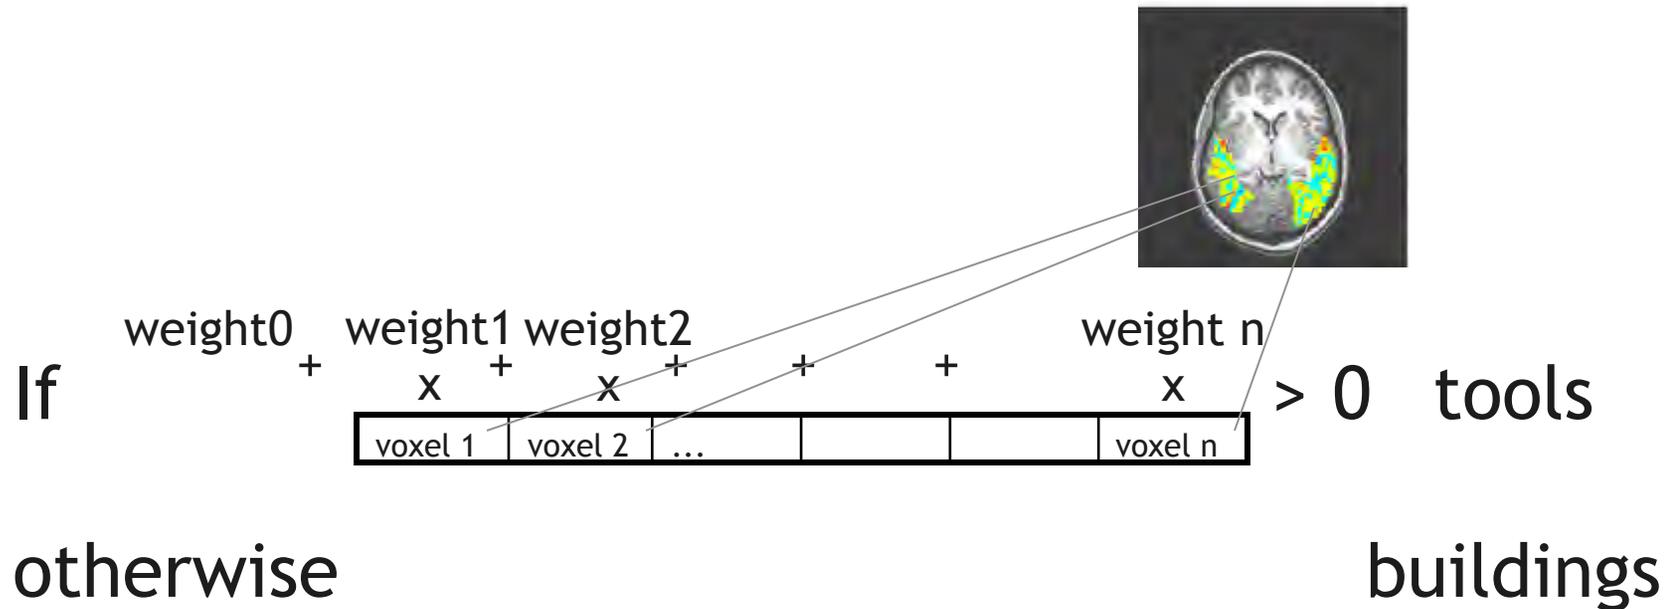
- next simplest: learn linear discriminant
- note that there are many solutions...



linear classifiers



linear classifiers



various kinds:

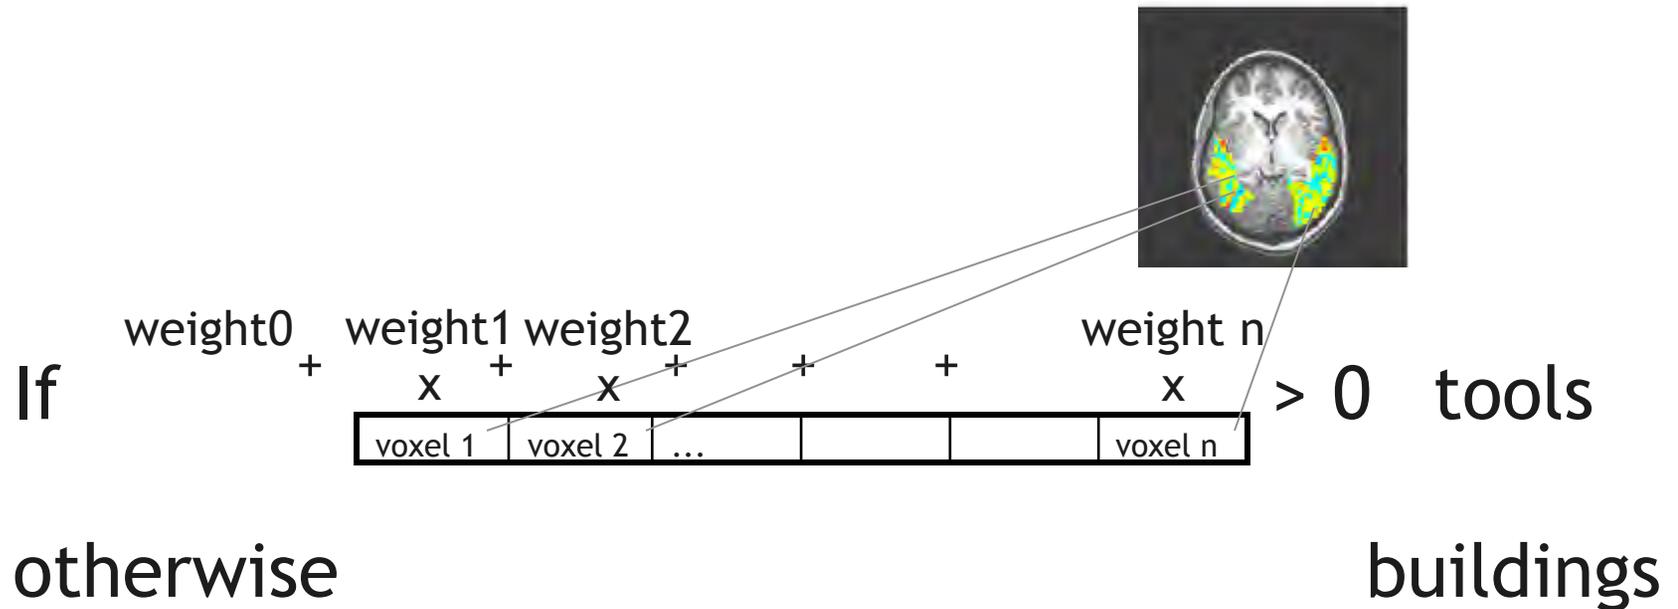
Gaussian Naive Bayes

Regularized Logistic Regression

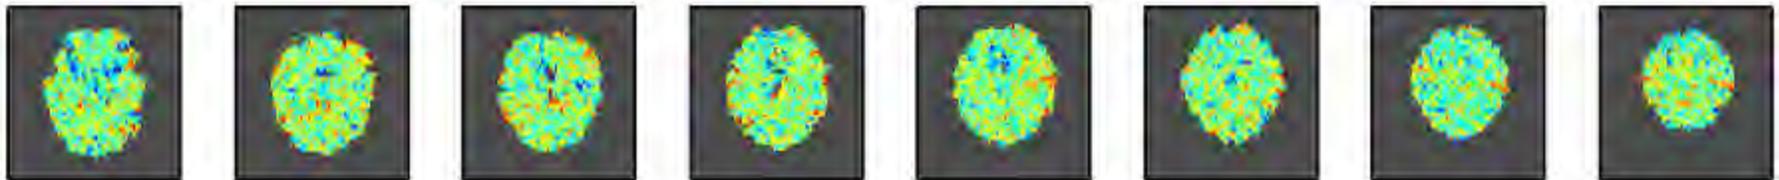
Linear Support Vector Machines (SVM)

differ on how weights are chosen

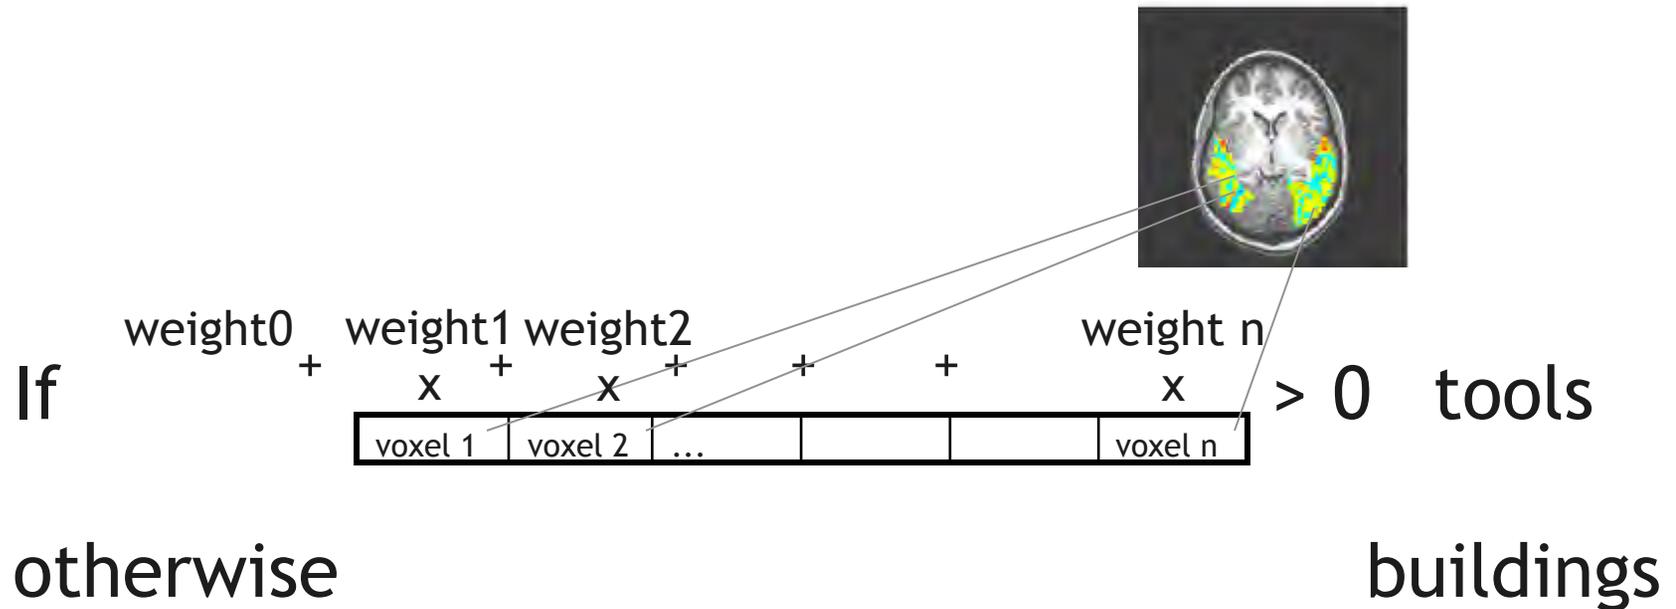
linear classifiers



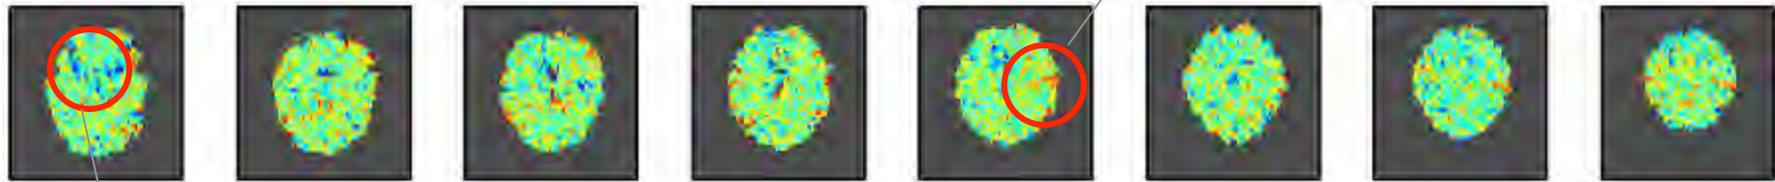
linear SVM weights:



linear classifiers



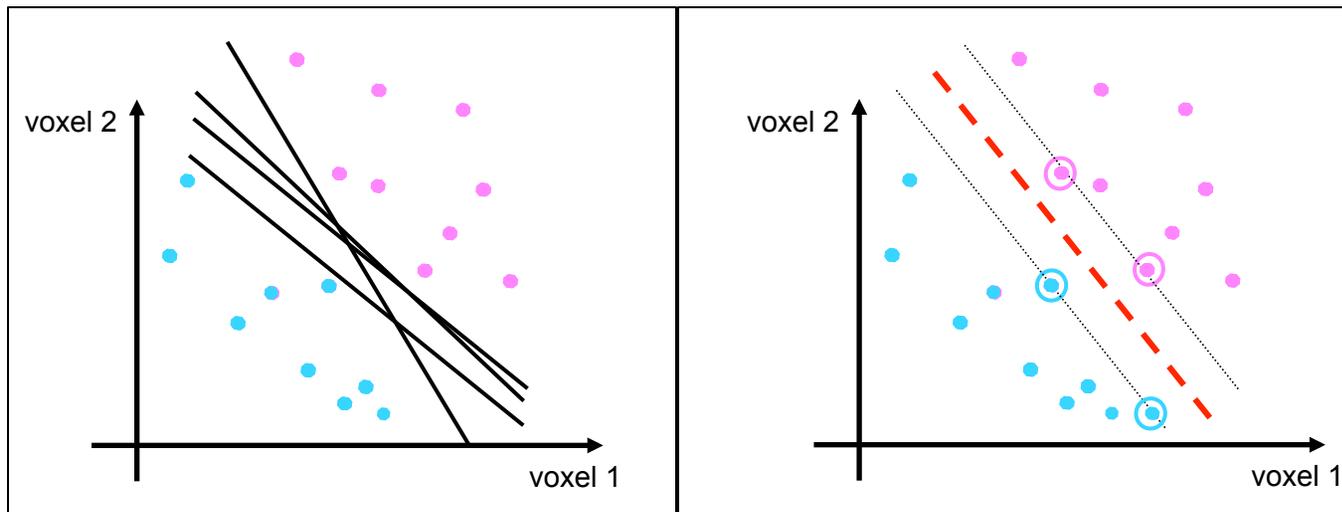
linear SVM weights:



weights pull towards buildings

weights pull towards tools

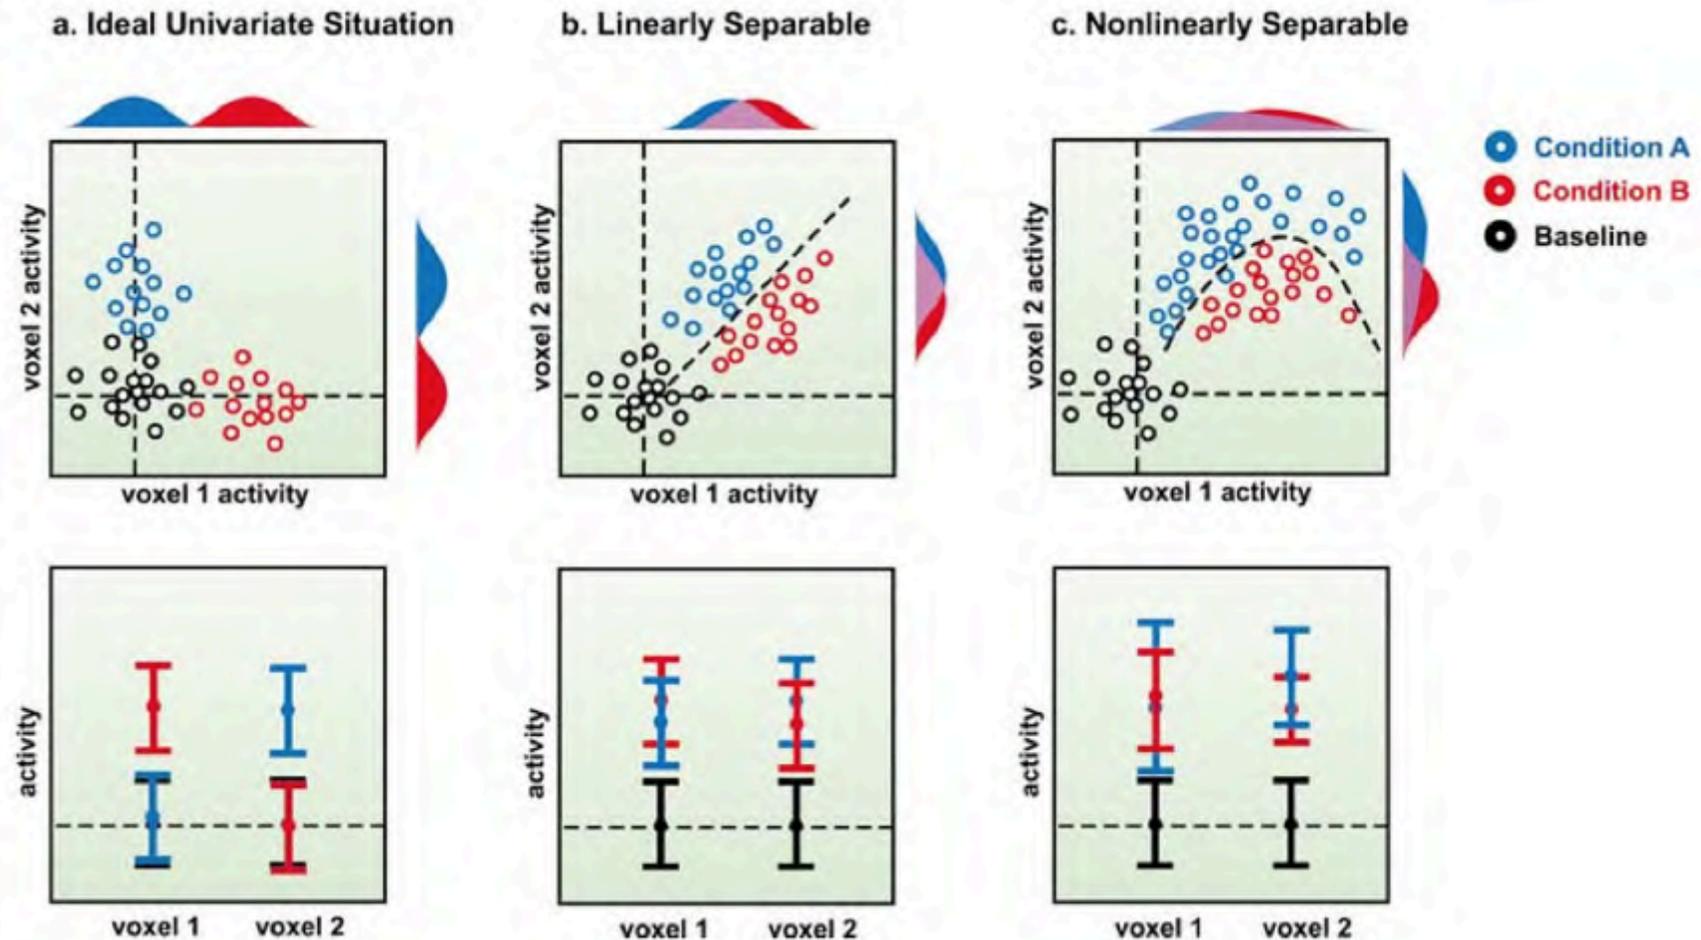
linear support vector machines



- Find linear decision boundary that maximizes the margin

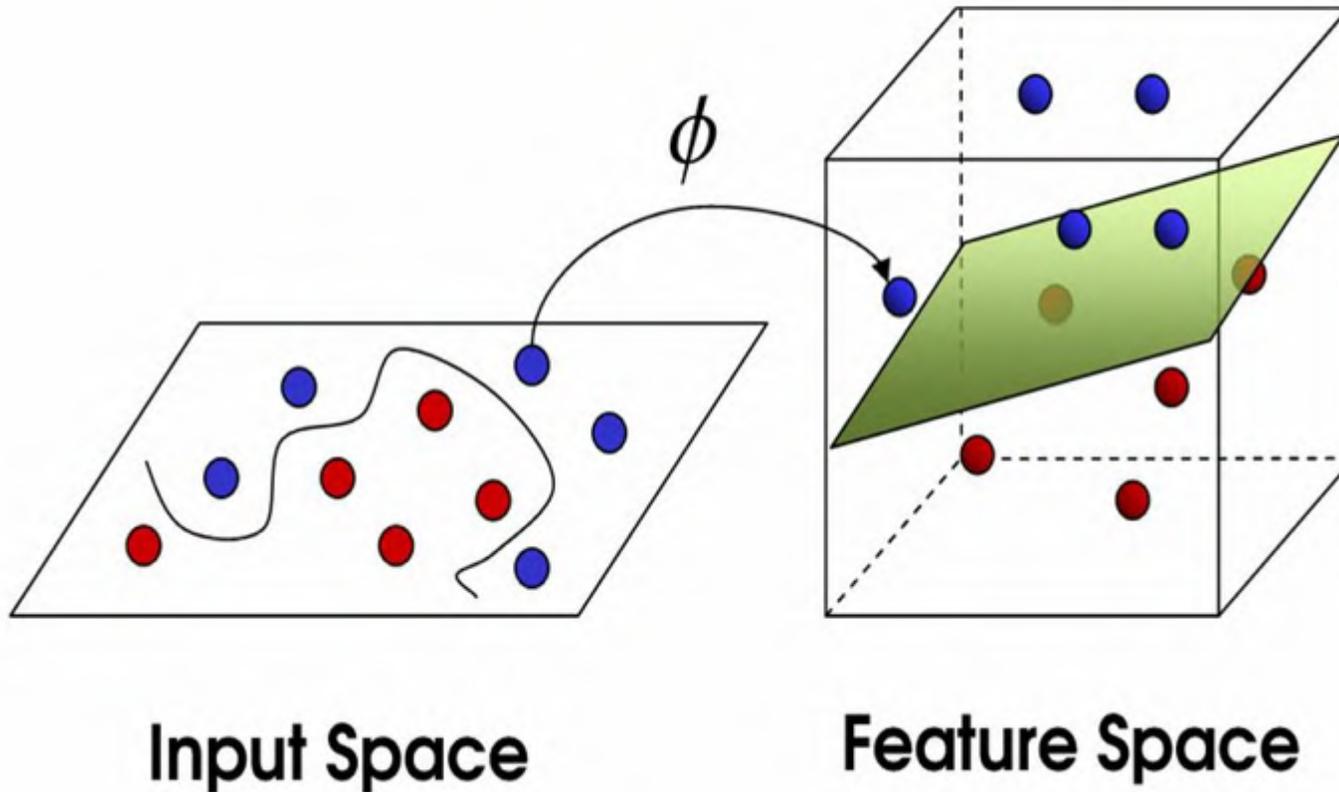
nonlinear classifiers

- sometimes, two classes will not be linearly separable



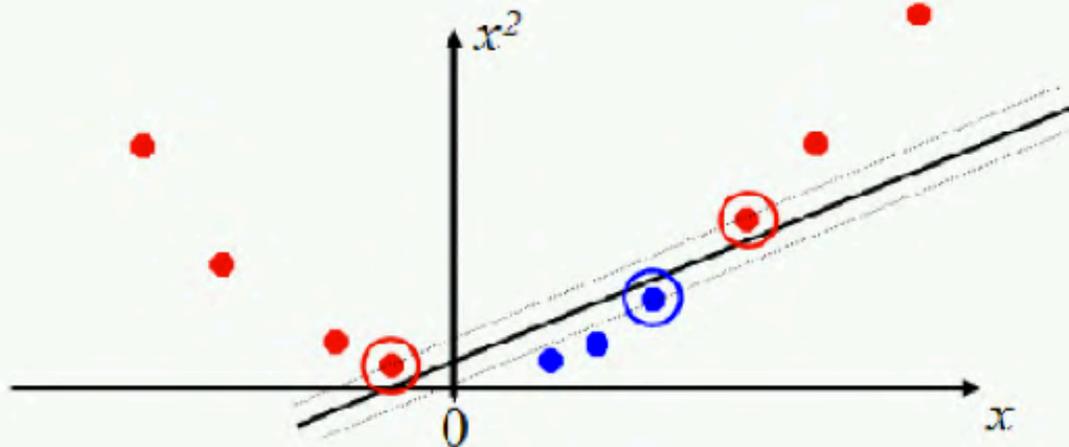
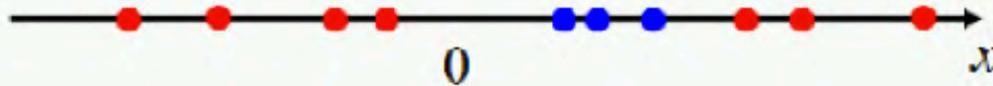
nonlinear classifiers

- Nonlinear decision boundaries can be represented as linear boundaries on a transformed feature space



nonlinear classifiers

- Nonlinear decision boundaries can be represented as linear boundaries on a transformed feature space



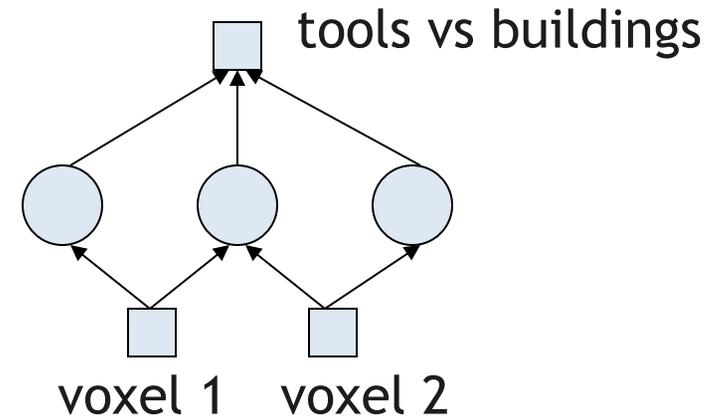
The “kernel trick”

(here a simple quadratic function creates the extra dimensionality)

Figure 15.6: Projecting data that is not linearly separable into a higher dimensional space can make it linearly separable.

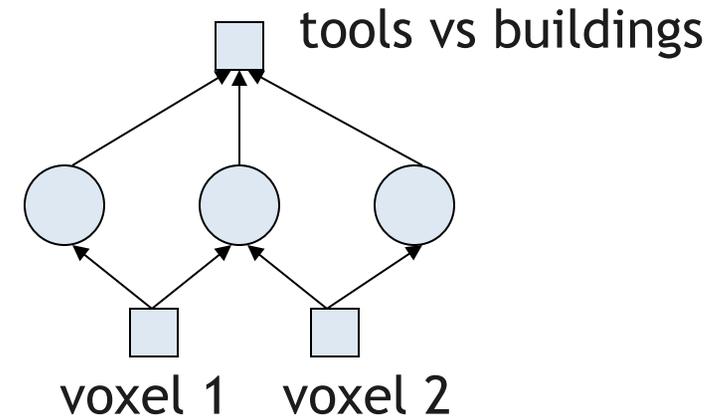
nonlinear classifiers

- **neural networks:**
new features are learnt
“hidden layer”

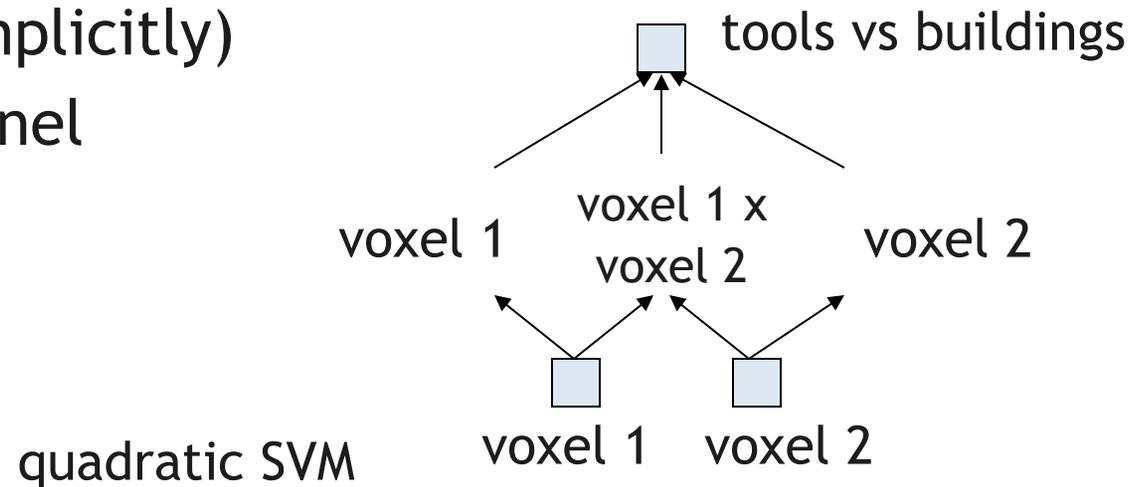


nonlinear classifiers

- **neural networks:**
new features are learnt
“hidden layer”



- **SVMs:**
new features are (implicitly)
determined by a kernel



nonlinear classifiers

reasons to be careful:

- too few examples,
too many features
- harder to interpret



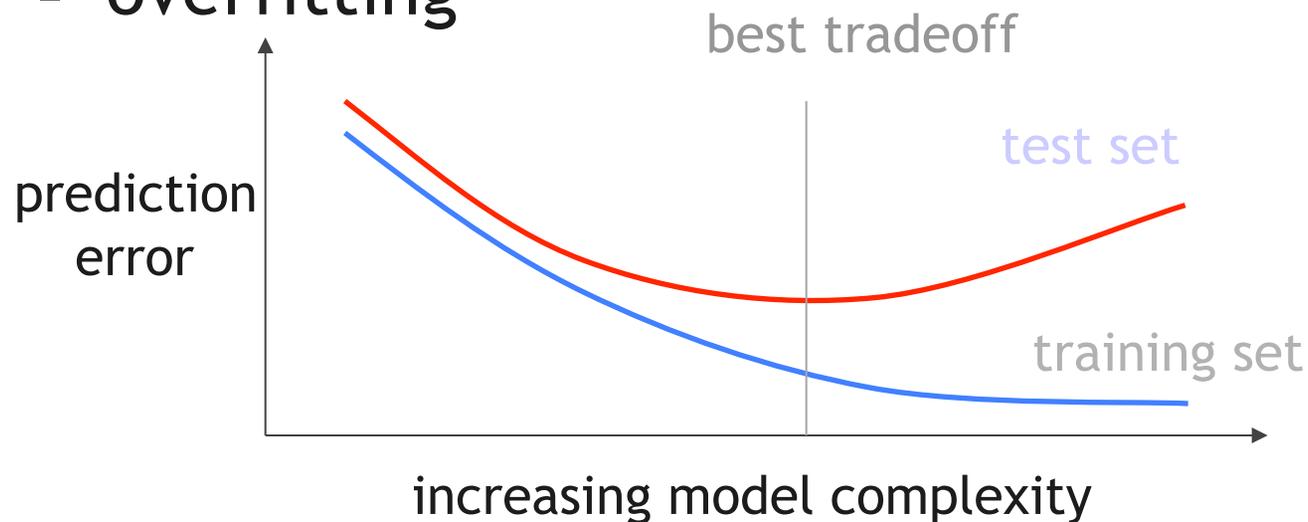
nonlinear classifiers

reasons to be careful:

- too few examples,
too many features
- harder to interpret



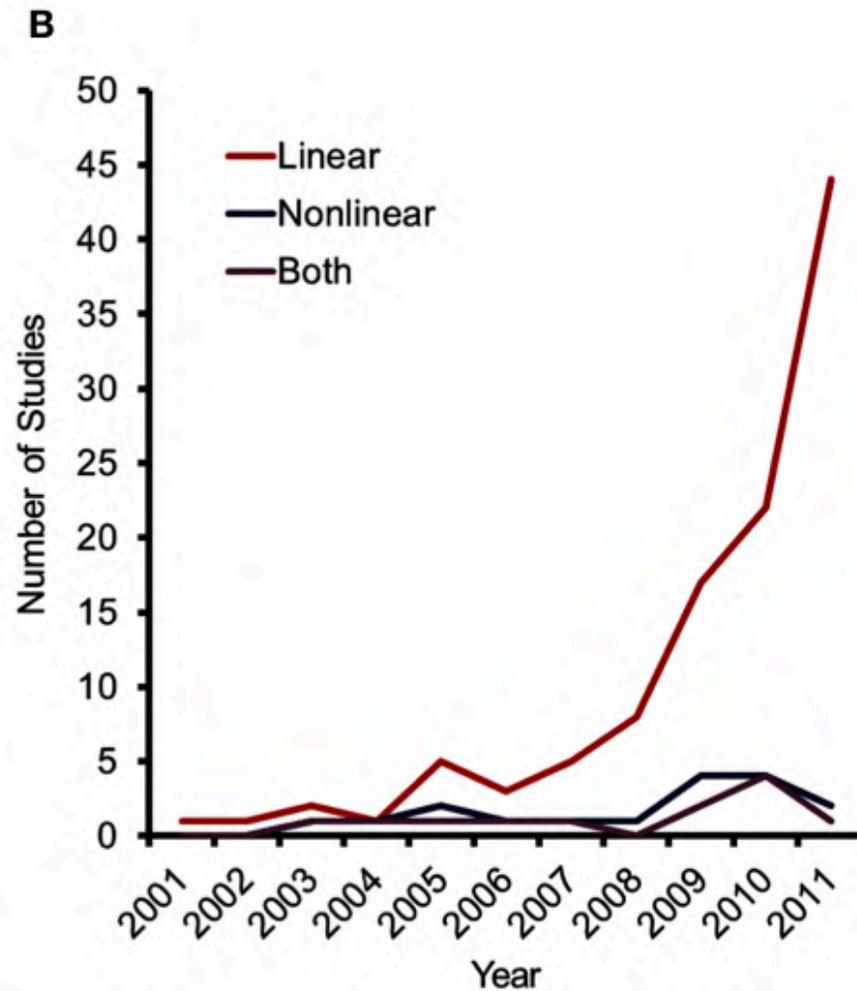
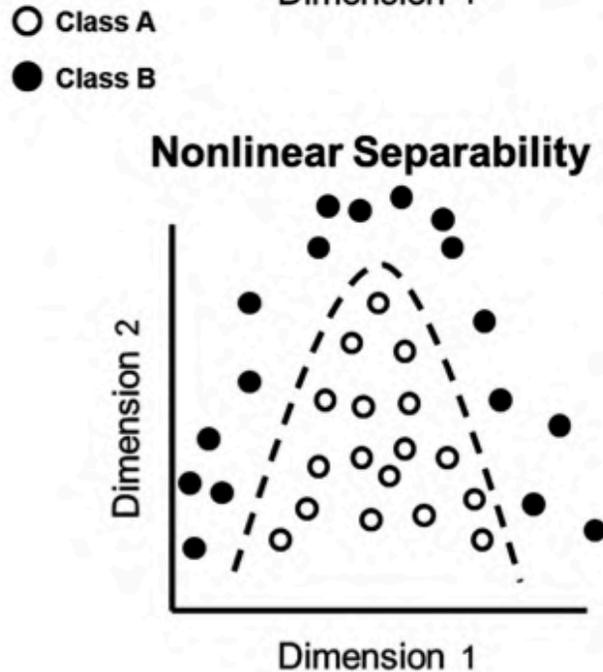
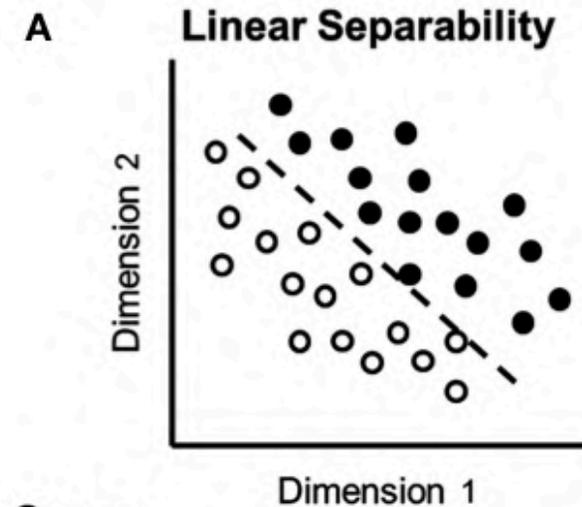
- overfitting



[from Hastie et al,2001]

Slide from Francisco Pereira

linear vs. nonlinear classifiers

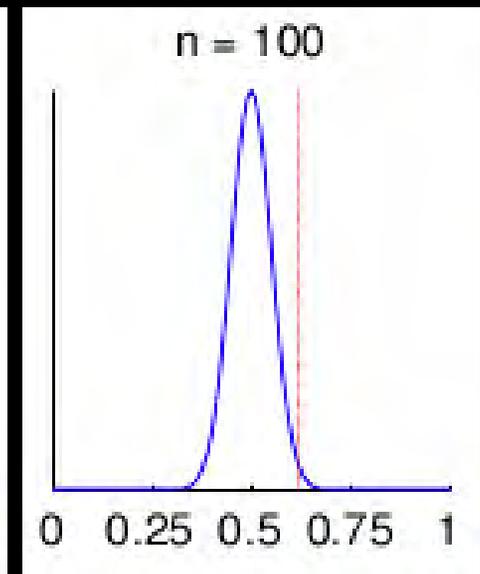
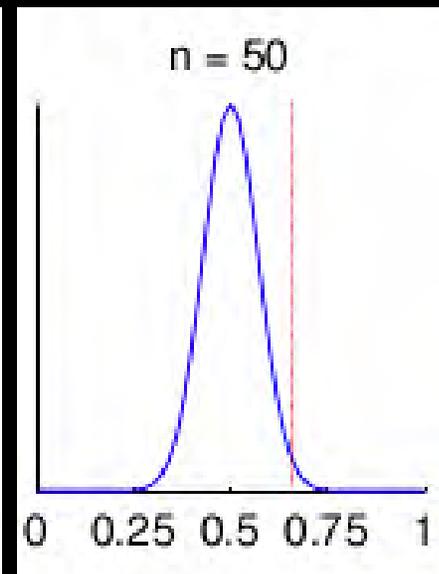
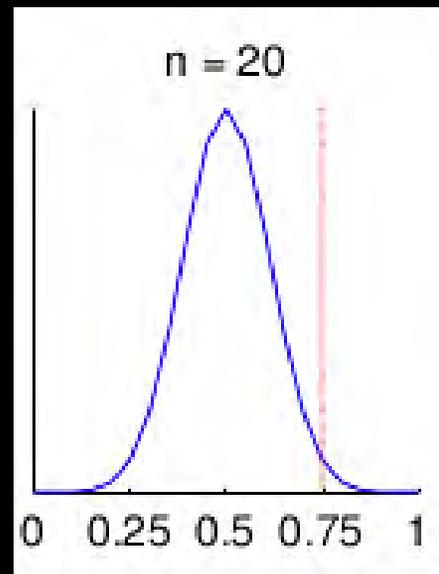
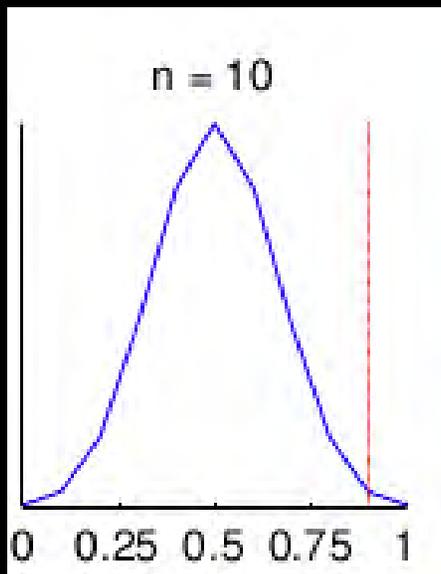


Determining whether classification accuracy is significantly better than chance

– Within-subjects

- Can compare to chance using binomial distribution

– e.g., $p_{chance} = .5$, trials = 120, successful = 75, $p_{observed} = .0039$



Determining whether classification accuracy is significantly better than chance

– Within-subjects

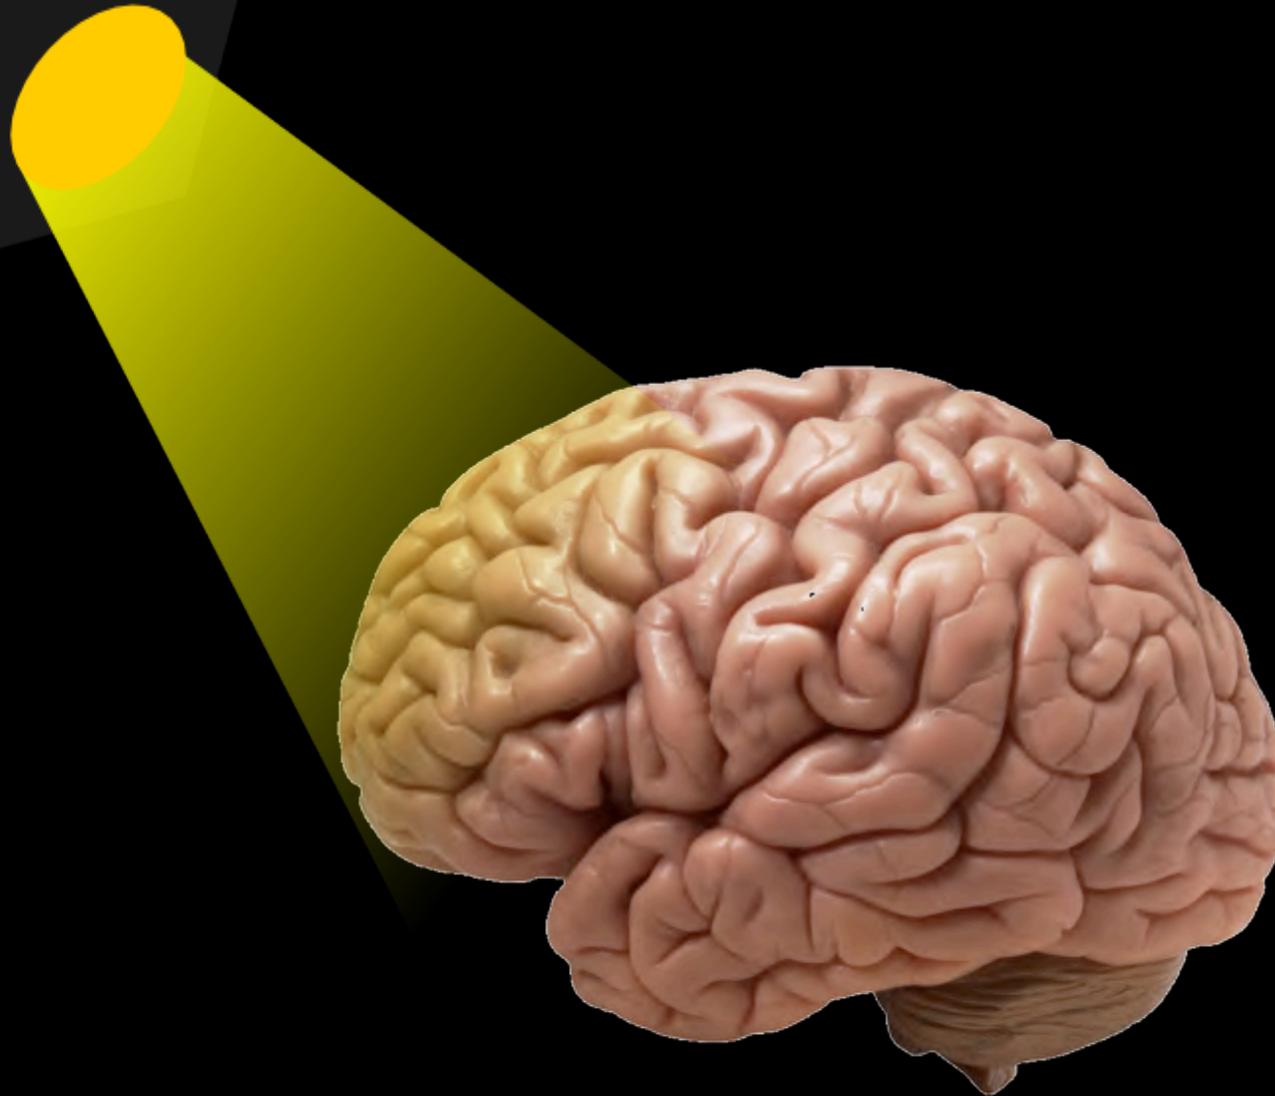
- Can compare to chance using binomial distribution
 - e.g., $p_{chance} = .5$, $trials = 120$, $successful = 75$, $p_{observed} = .0039$
- Or use permutation test to compare observed accuracy to distribution of performance generated by shuffling regressors many times (e.g., 1000 shuffles)

– Across-subjects

- Use one-sample t-test to compare subjects' mean classification accuracy to chance
- Commonly used, but places no requirement on mean accuracy (i.e., 53% correct could be highly significant)

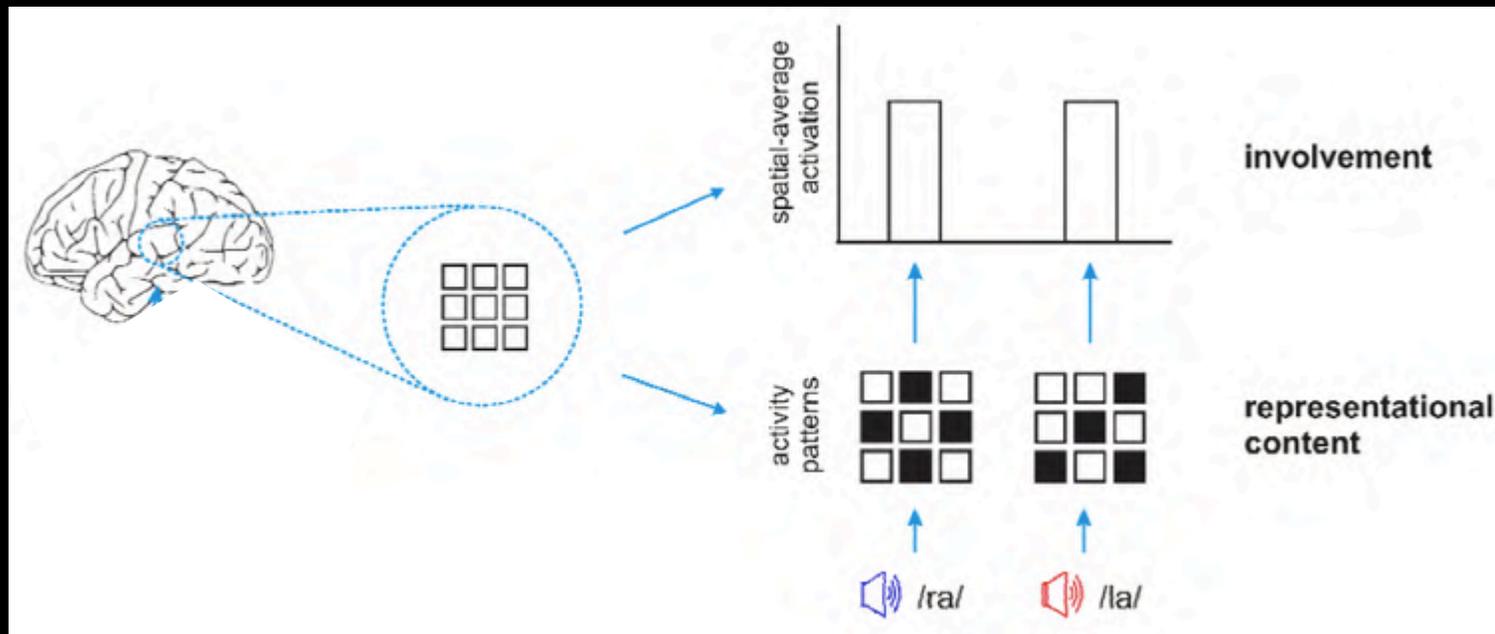
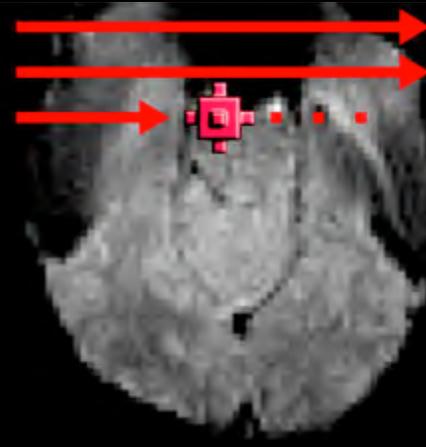
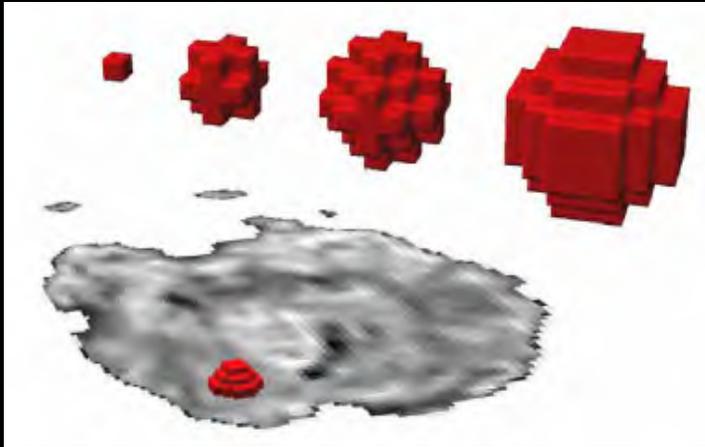
Information-based brain mapping

- Spherical searchlight mapping approach (Kriegeskorte et al. (2006), *PNAS*)



Information-based brain mapping

- Spherical searchlight mapping approach (Kriegeskorte et al. (2006), *PNAS*)



Information-based brain mapping

- Spherical searchlight mapping approach (Kriegeskorte et al. (2006), *PNAS*)

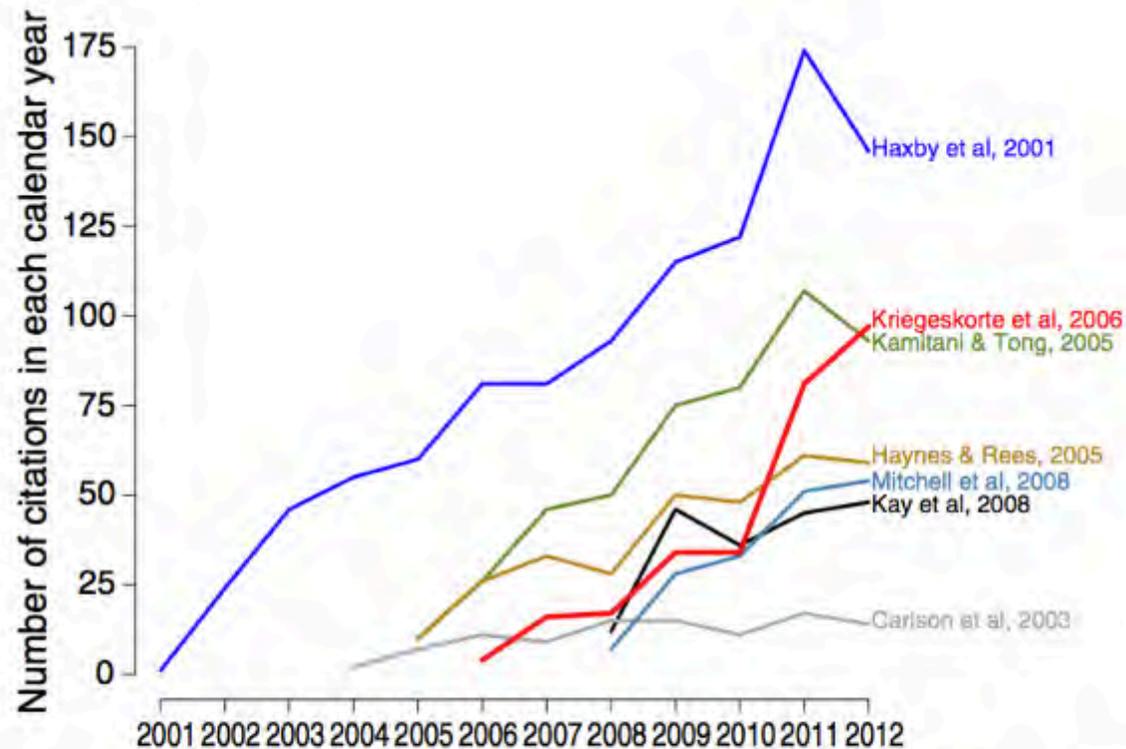
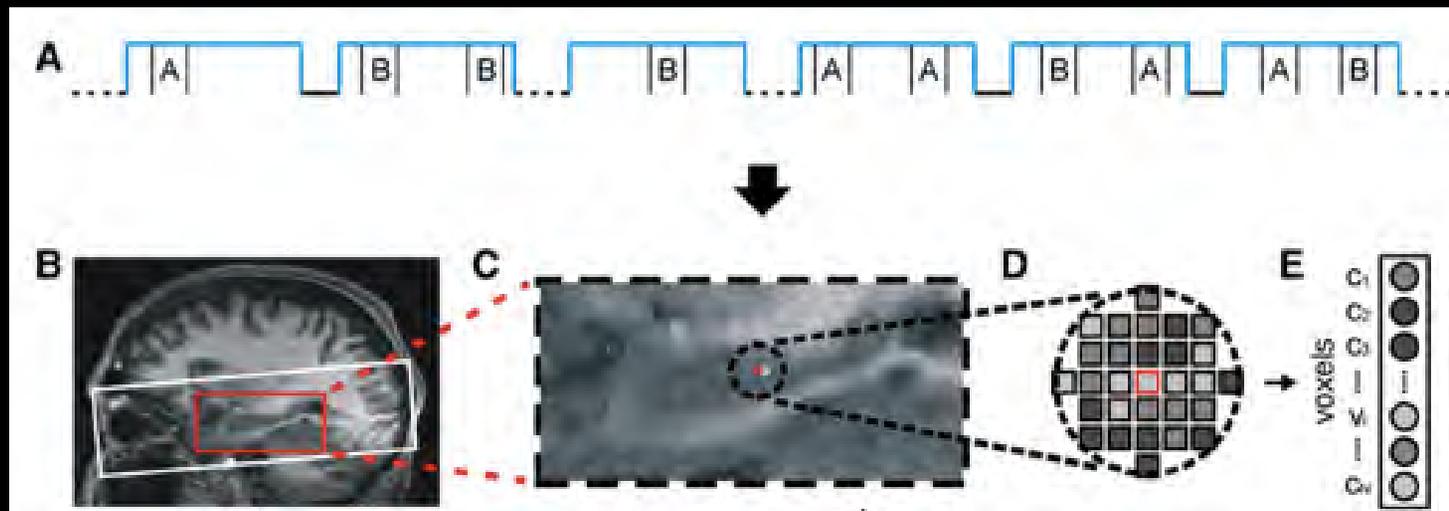
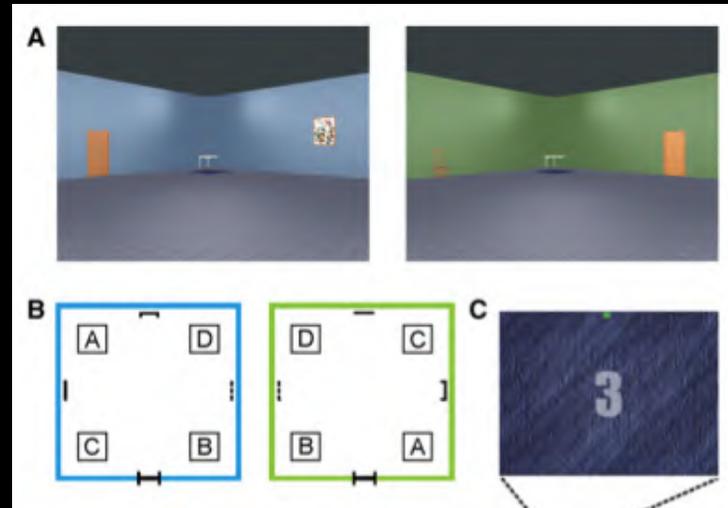


Fig. 1. Pattern-information fMRI is still a rapidly growing field, particularly searchlight analysis (note the rapid increase in papers citing Kriegeskorte et al., 2006). This figure follows Fig. 2 in Raizada and Kriegeskorte (2010), but uses the actual citation counts after 2008. The number of citations for each paper and year was obtained via Scopus (www.scopus.com) on 9 January 2013. (Carlson et al., 2003; Haxby et al., 2001; Haynes and Rees, 2005; Kamitani and Tong, 2006; Kay et al., 2008; Mitchell et al., 2008).

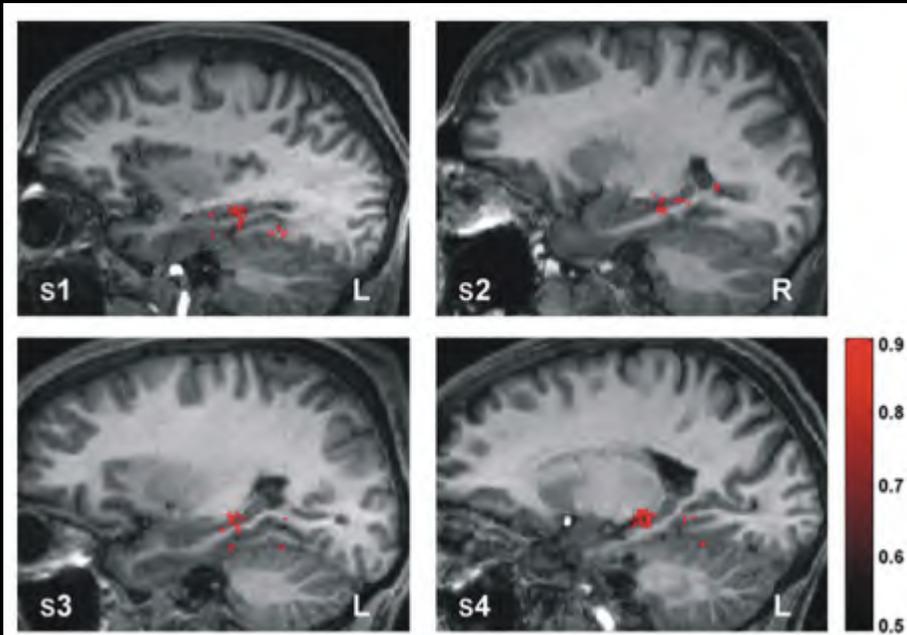
Decoding the representational content of medial temporal lobe activity patterns with high resolution fMRI data



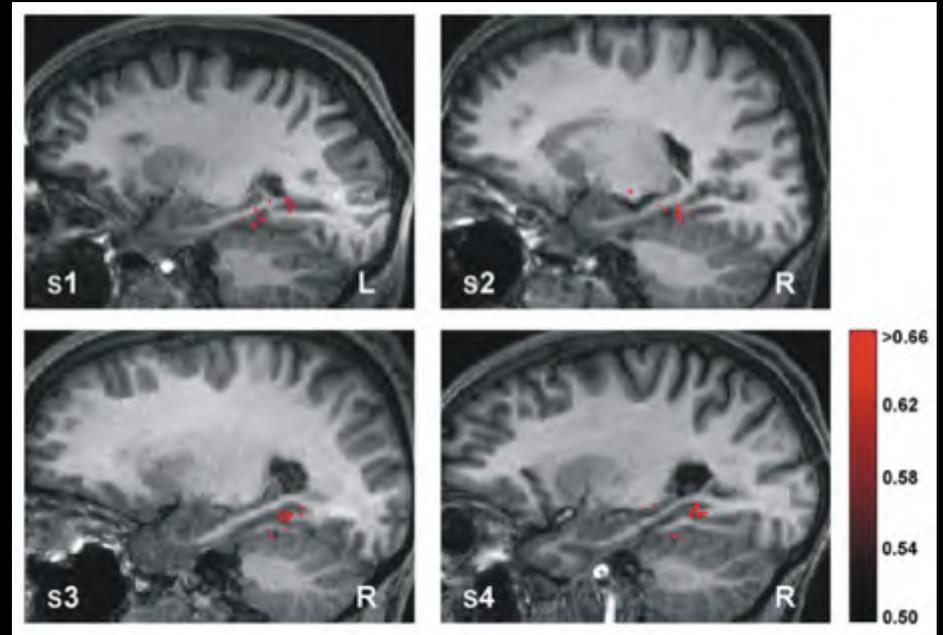
Hassabis et al. (2009) *Neuron*

Decoding the representational content of local activity patterns within the MTL

Which spatial quadrant?



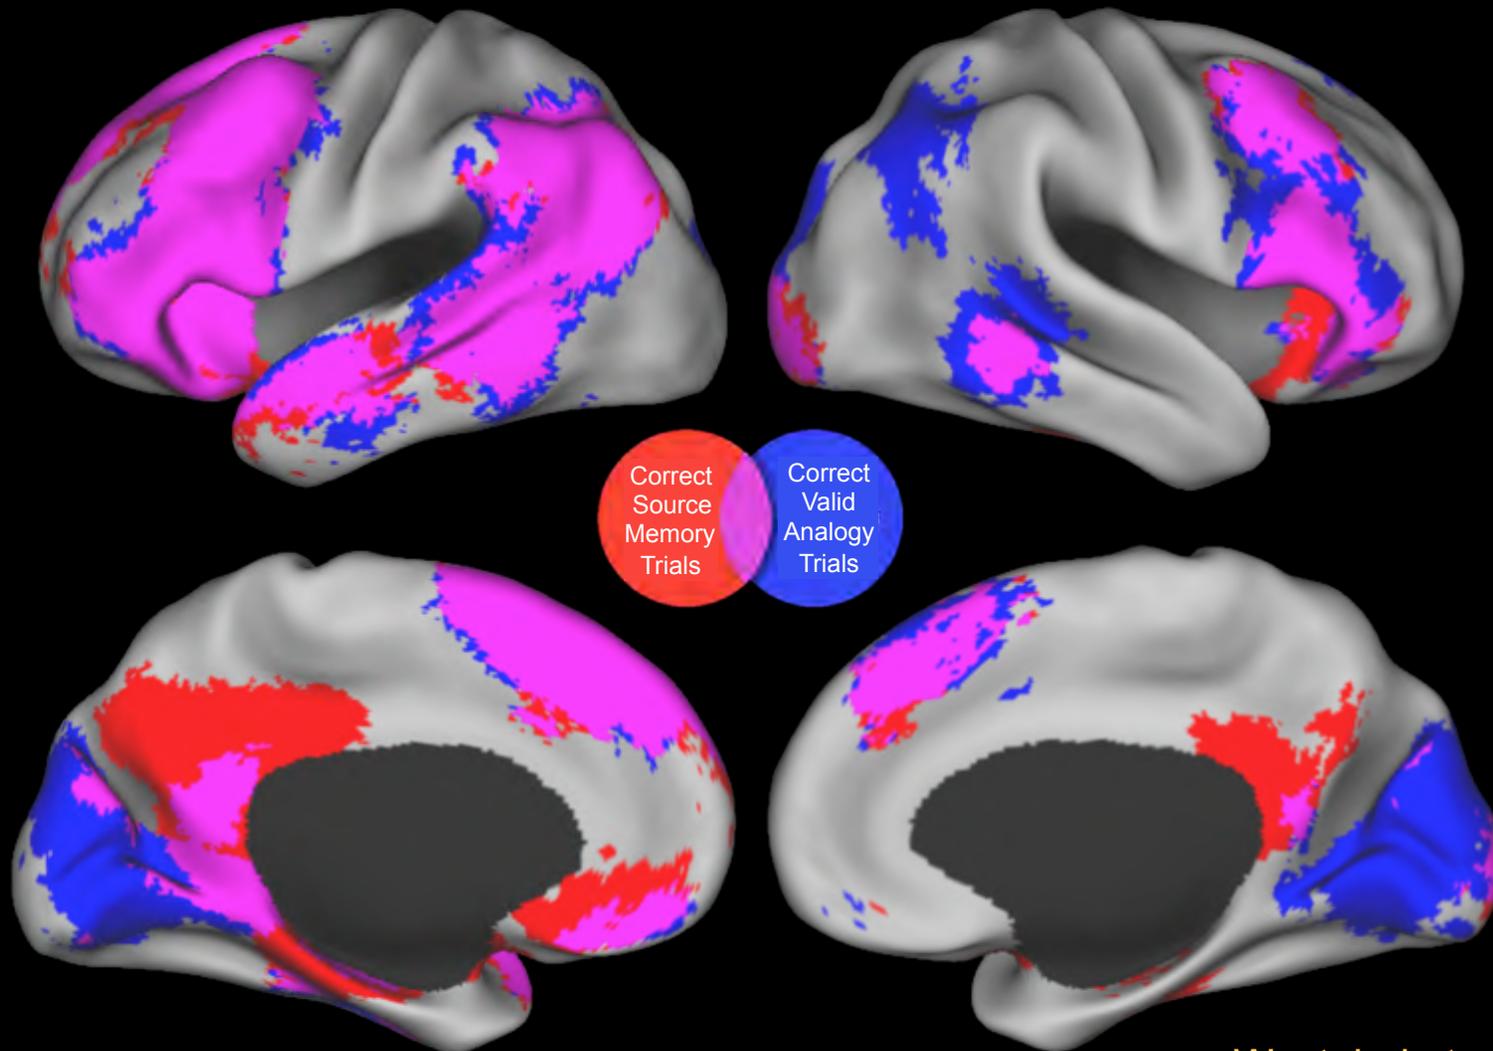
Which room?



Detecting subtle differences in
mental states despite common
univariate activity levels

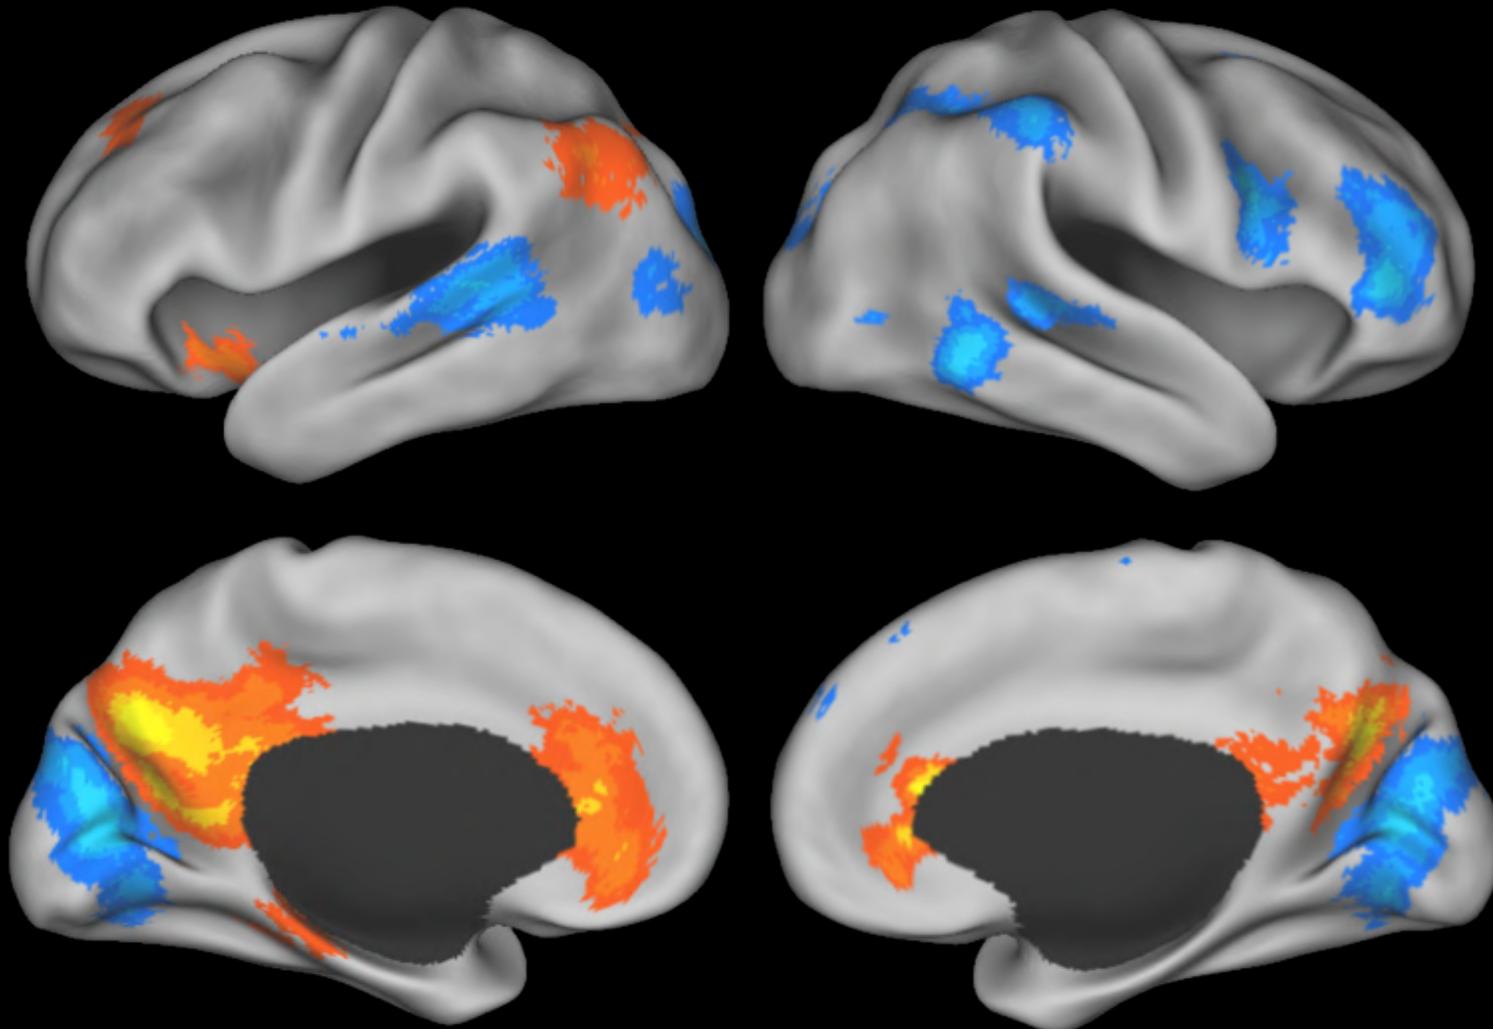
Retrieving Memories & Solving Analogies

Extensive overlap across lateral PFC



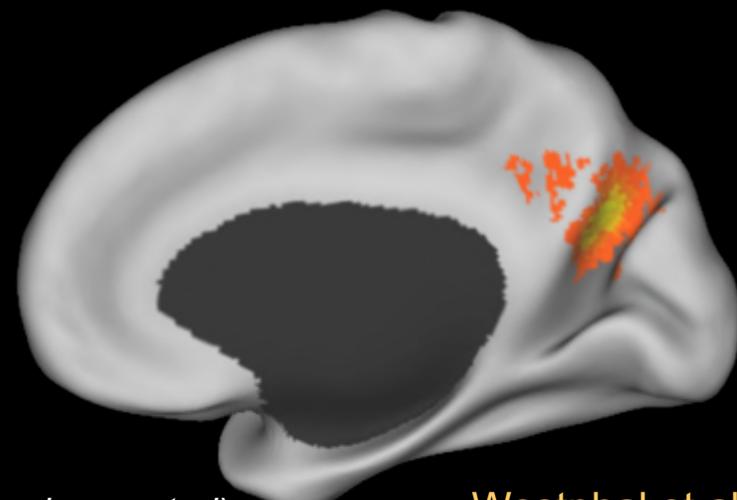
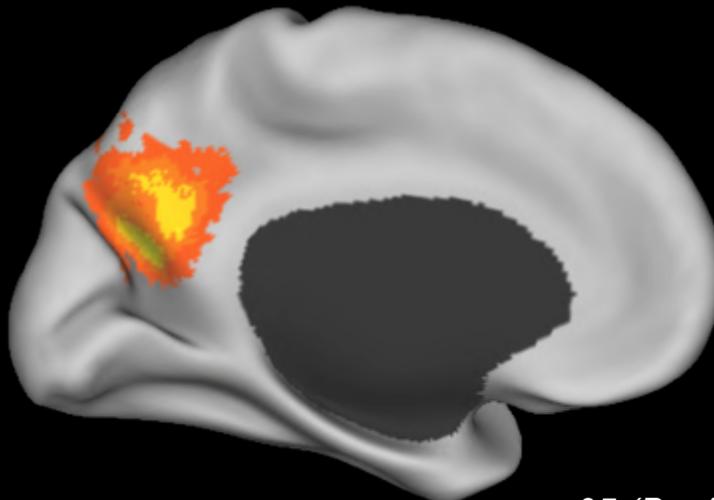
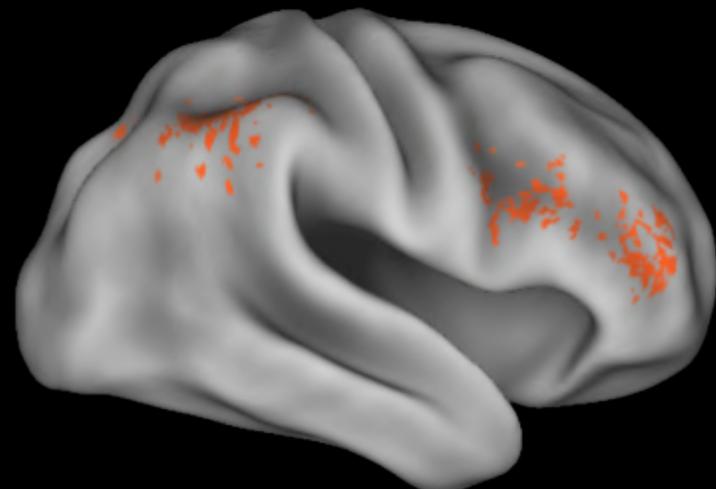
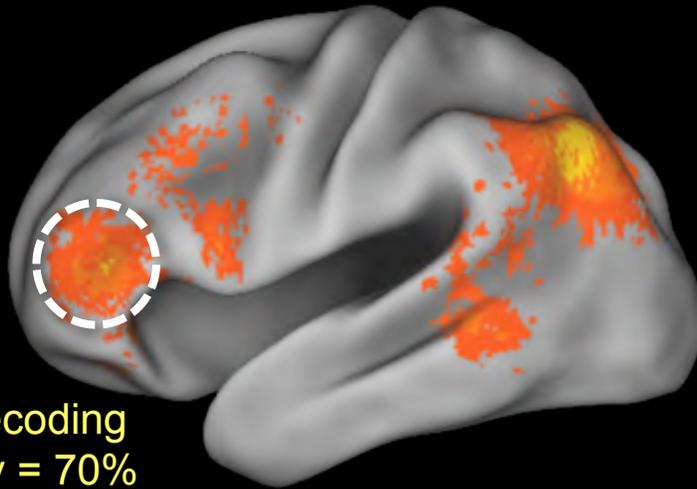
Source Memory vs. Analogy

No significant difference in Left Lateral PFC



Searchlight MVPA:

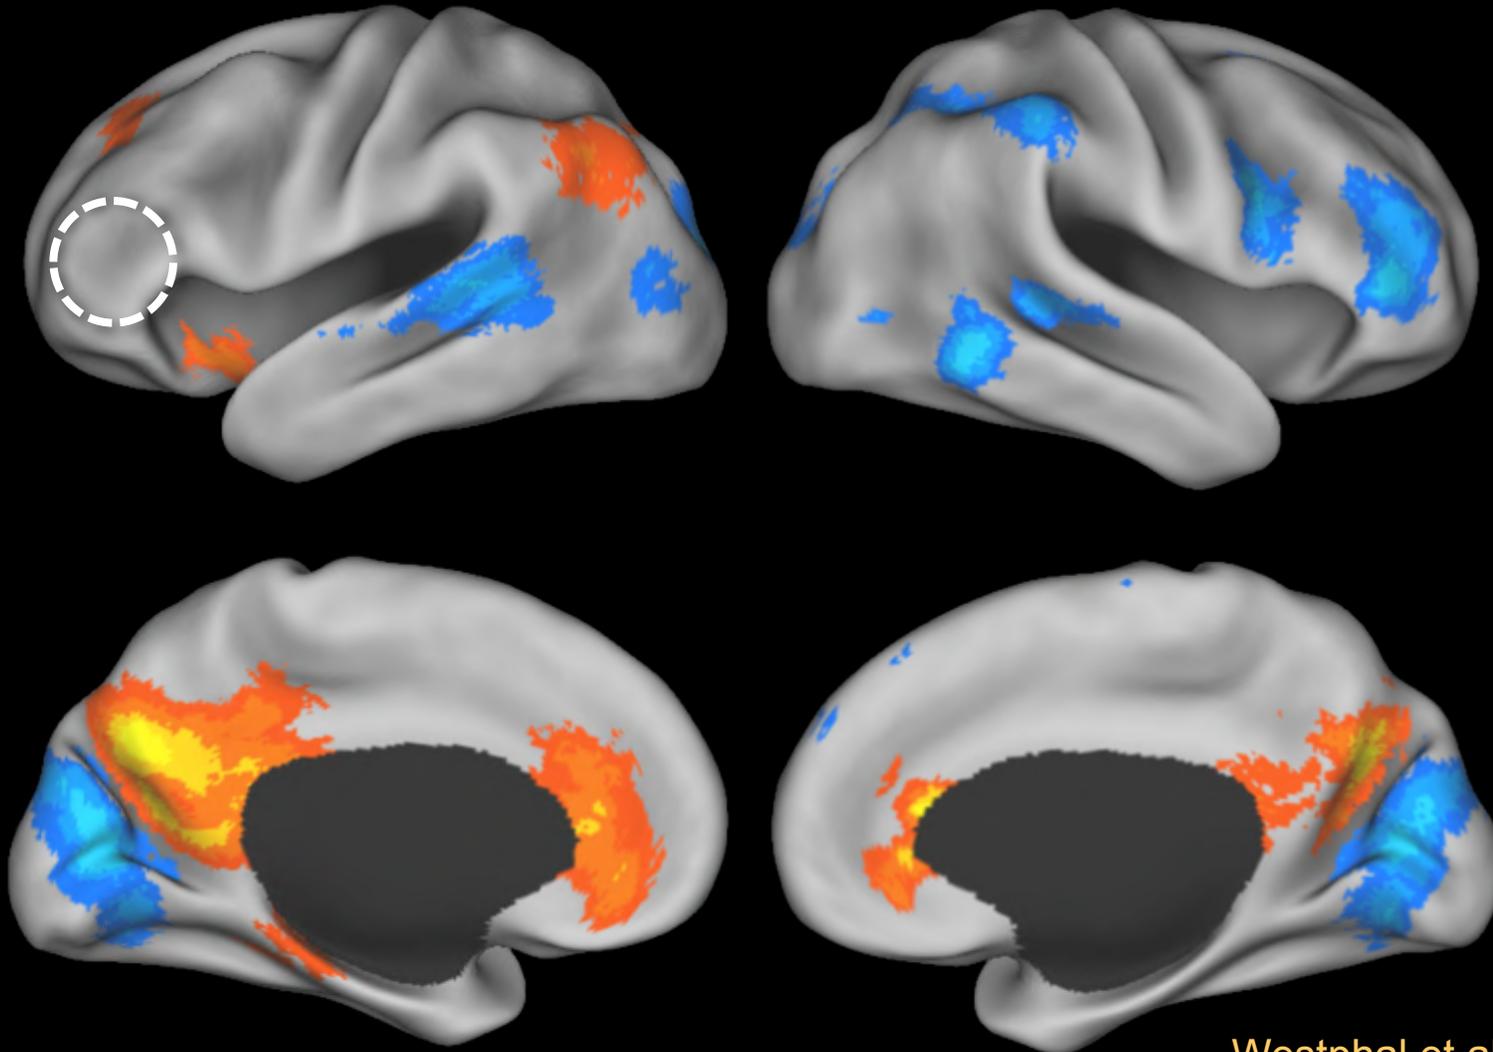
Decoding Source Memory vs. Analogy trials



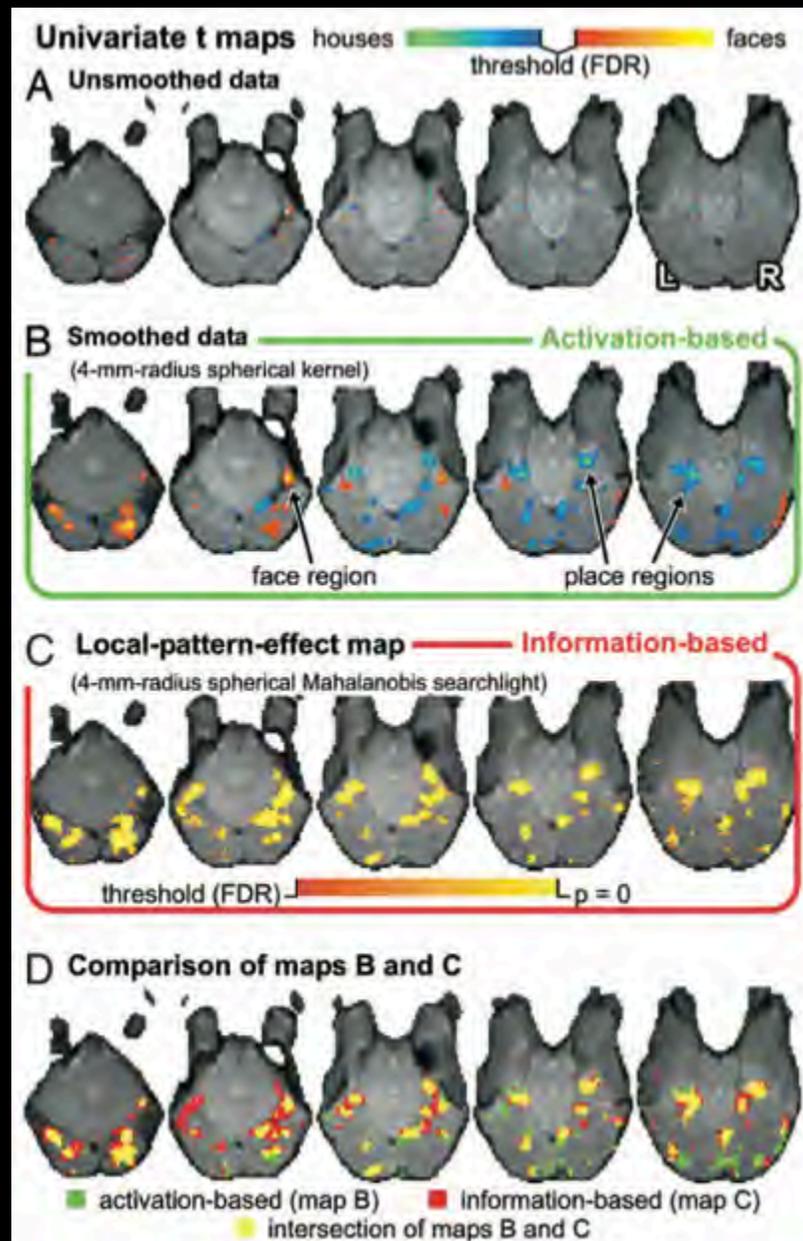
$p < .05$ (Bonferroni-corrected)

Westphal et al. (in prep)

Univariate Effects (revisited): *Source Memory* vs. *Analogy*

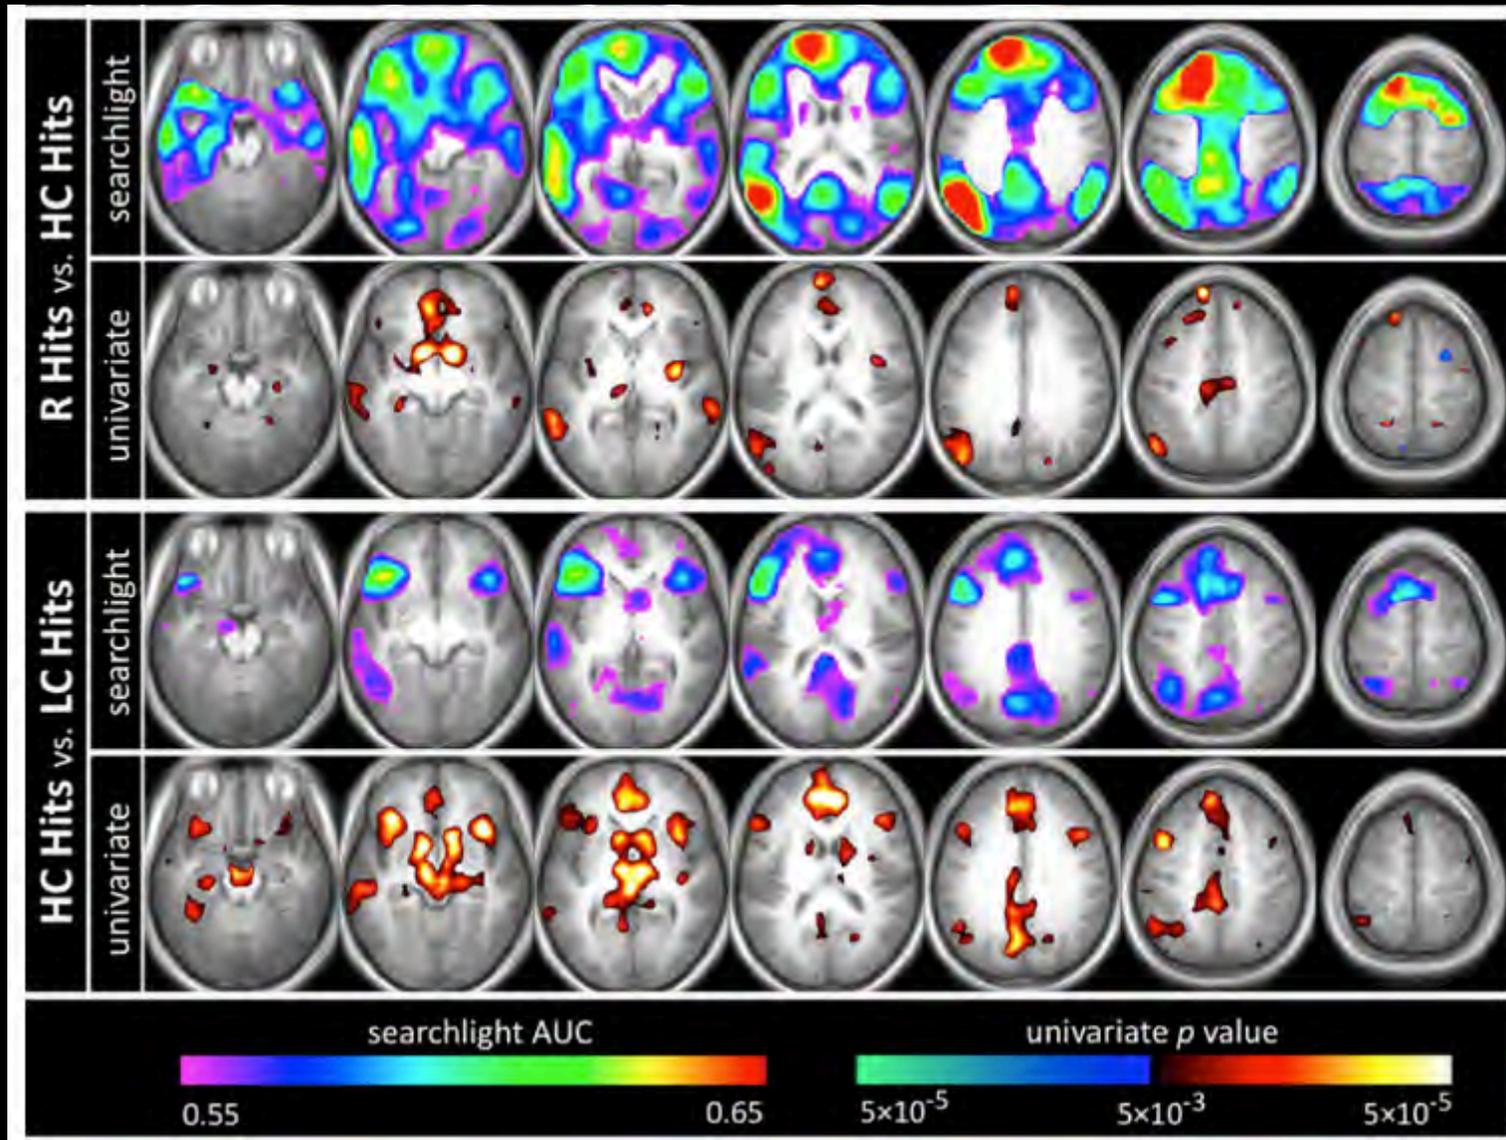


Univariate Mapping vs. MVPA Searchlight Mapping



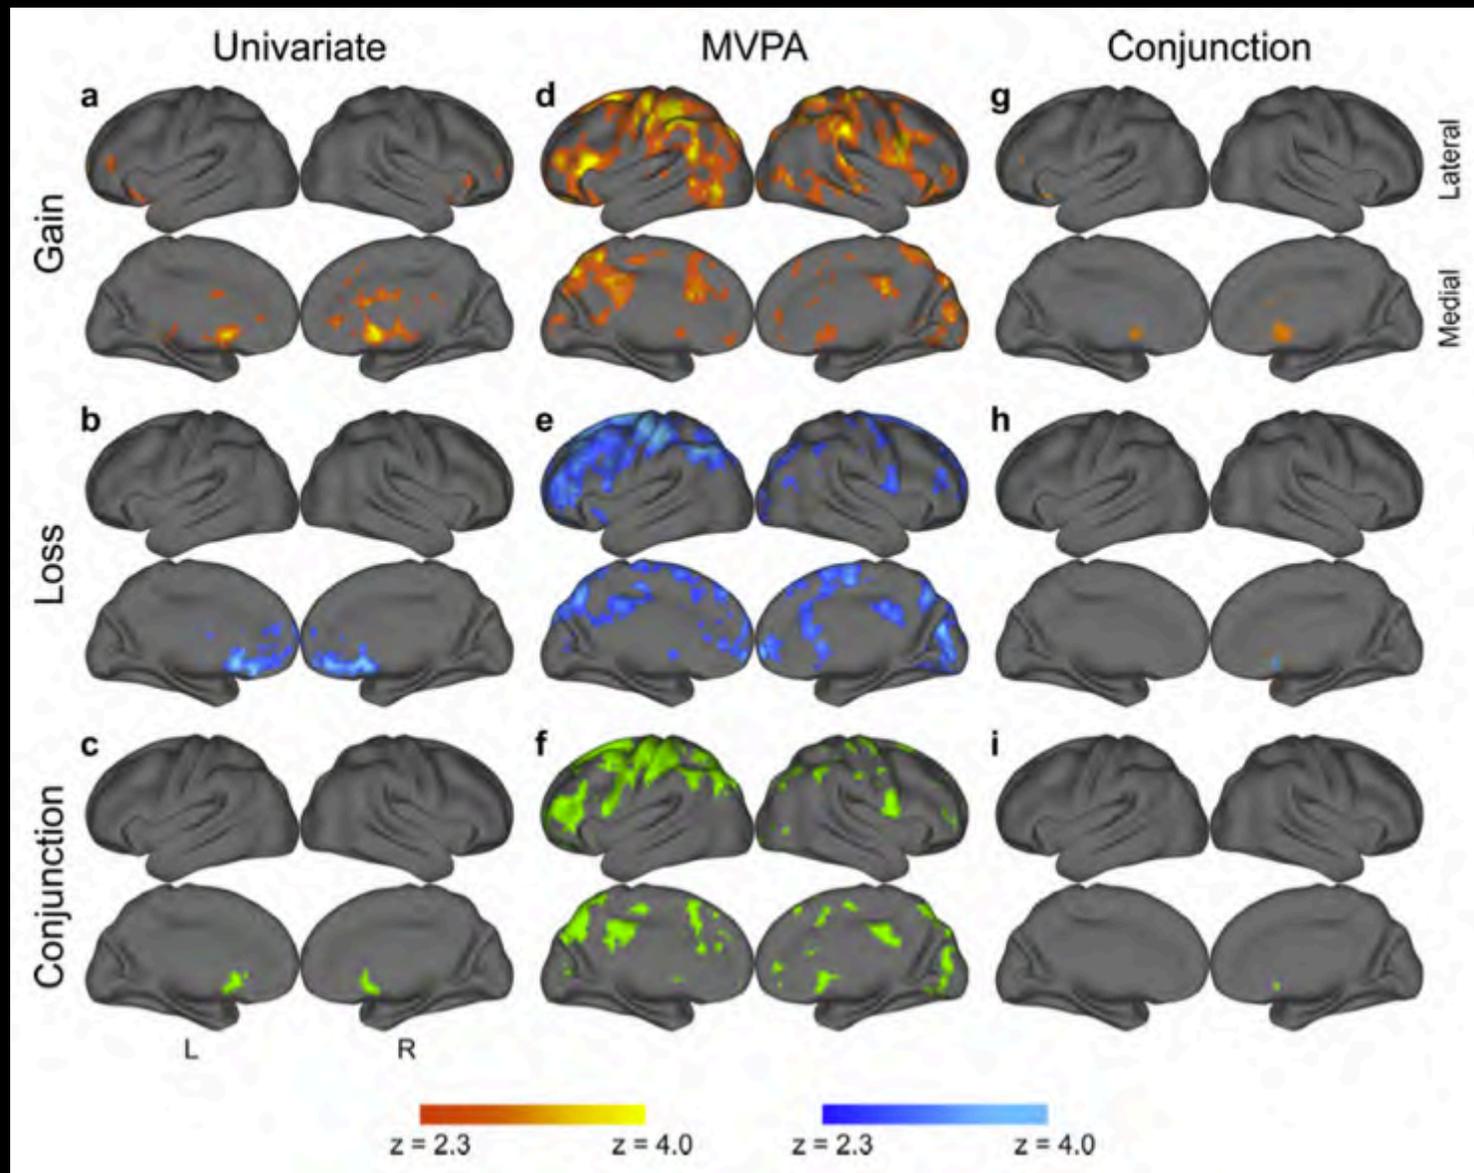
Kriegeskorte et al.
(2006), *PNAS*

Univariate Mapping vs. MVPA Searchlight Mapping



Rissman, Greely, & Wagner (2010) *PNAS*

Univariate Mapping vs. MVPA Searchlight Mapping



Jimura & Poldrack (2012) *Neuroimage*

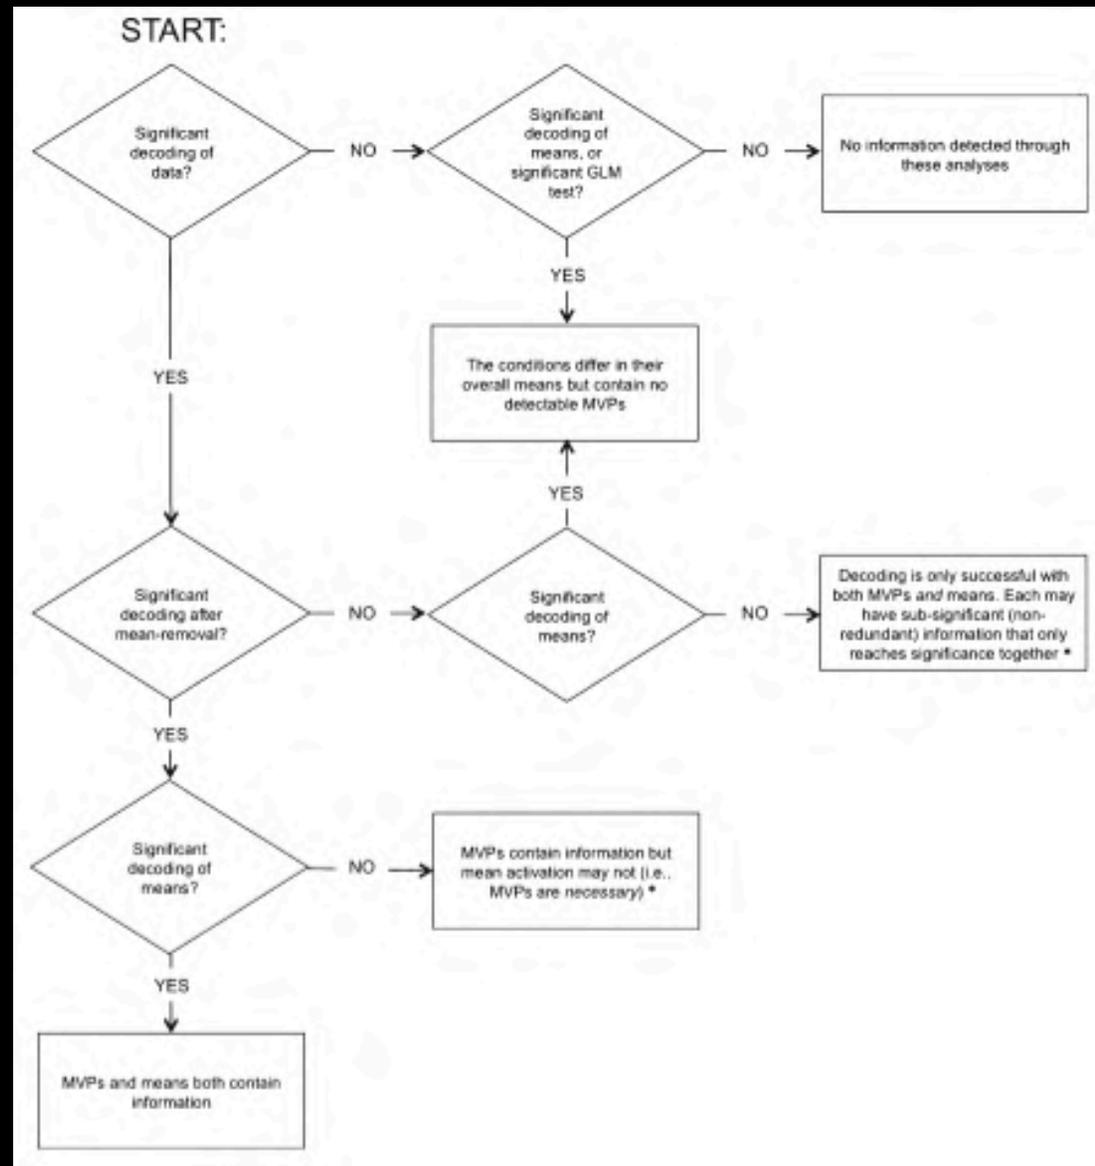
Univariate Mapping vs. MVPA Searchlight Mapping

Cogn Affect Behav Neurosci
DOI 10.3758/s13415-013-0186-2

Distinguishing multi-voxel patterns and mean activation: Why, how, and what does it tell us?

Marc N. Coutanche

Univariate Mapping vs. MVPA Searchlight Mapping



Univariate Mapping vs. MVPA Searchlight Mapping

NeuroImage 97 (2014) 271–283



Contents lists available at ScienceDirect

NeuroImage

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



What do differences between multi-voxel and univariate analysis mean? How subject-, voxel-, and trial-level variance impact fMRI analysis



Tyler Davis ^{a,*}, Karen F. LaRocque ^{b,**}, Jeanette A. Mumford ^{d,e,f}, Kenneth A. Norman ^{g,h},
Anthony D. Wagner ^{b,c}, Russell A. Poldrack ^{d,e,f}

^a Department of Psychology, Texas Tech University, USA

^b Department of Psychology, Stanford University, USA

^c Neurosciences Program, Stanford University, USA

^d Department of Psychology, University of Texas at Austin, USA

^e Department of Neuroscience, University of Texas at Austin, USA

^f Imaging Research Center, University of Texas at Austin, USA

^g Department of Psychology, Princeton University, USA

^h Princeton Neuroscience Institute, Princeton University, USA

Searchlight MVPA: *Too good to be true?*

NeuroImage 78 (2013) 261–269



ELSEVIER

Contents lists available at SciVerse ScienceDirect

NeuroImage

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



Comments and Controversies

Searchlight analysis: Promise, pitfalls, and potential

Joset A. Etzel^{a,*}, Jeffrey M. Zacks^b, Todd S. Braver^a

NeuroImage 77 (2013) 157–165



ELSEVIER

Contents lists available at SciVerse ScienceDirect

NeuroImage

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



Confounds in multivariate pattern analysis: Theory and rule representation case study

Michael T. Todd^{*}, Leigh E. Nystrom, Jonathan D. Cohen

DISCOVER

THE MAGAZINE OF SCIENCE, TECHNOLOGY, AND THE FUTURE

Search DIS

THE MAGAZINE | BLOGS | HEALTH & MEDICINE | MIND & BRAIN | TECHNOLOGY

BLOGS Visual Science | Gene Expression | The Crux | Collide-a-Scape | Out There
D-brief | Water Works | Field Notes | Body Horrors | Seriously, Science?

GET \$50 BACK*

Stay two consecutive weekends
back* when you pay with your



Neuroskeptic

» The Man With Uncrossed Eyes

For Preregistration in Fundamental Research »

fMRI: More Voxels, More Problems?

By Neuroskeptic | April 18, 2013 3:04 pm

A new paper could prompt a rethink of a technique that's become very hot in neuroscience lately: **Confounds in multivariate pattern analysis**

<http://blogs.discovermagazine.com/neuroskeptic/2013/04/18/fmri-more-voxels-more-problems>

The Neurocomplimenter

I'm starting a backlash against the fashionable anti-neuro backlash. After seven years of critical neuroblogging, it's time to highlight the positives.

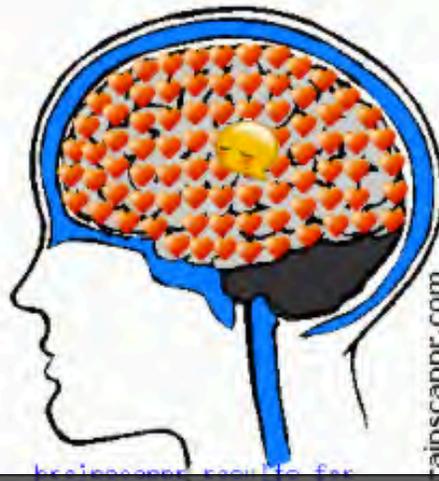
Thursday, July 4, 2013

What is this, anyway?

The [Neurocomplimenter](#) is a new project designed to counter gratuitous anti-neuroscience sentiment. It's part of my campaign to combat pop neurobashing profiteers.

The Neurocomplimenter

It's here!



<http://neurocomplimenter.blogspot.com>

Inaugural Posts

[Helicopter controlled by human EEG](#)
[The End of Language Embodiment](#)

Blog Archive

▼ 2013 (2)

▼ July (2)

[A New Slant on Frontal
Connectivity: the Frontal
A...](#)

[What is this, anyway?](#)

About Me



The Neurocritic

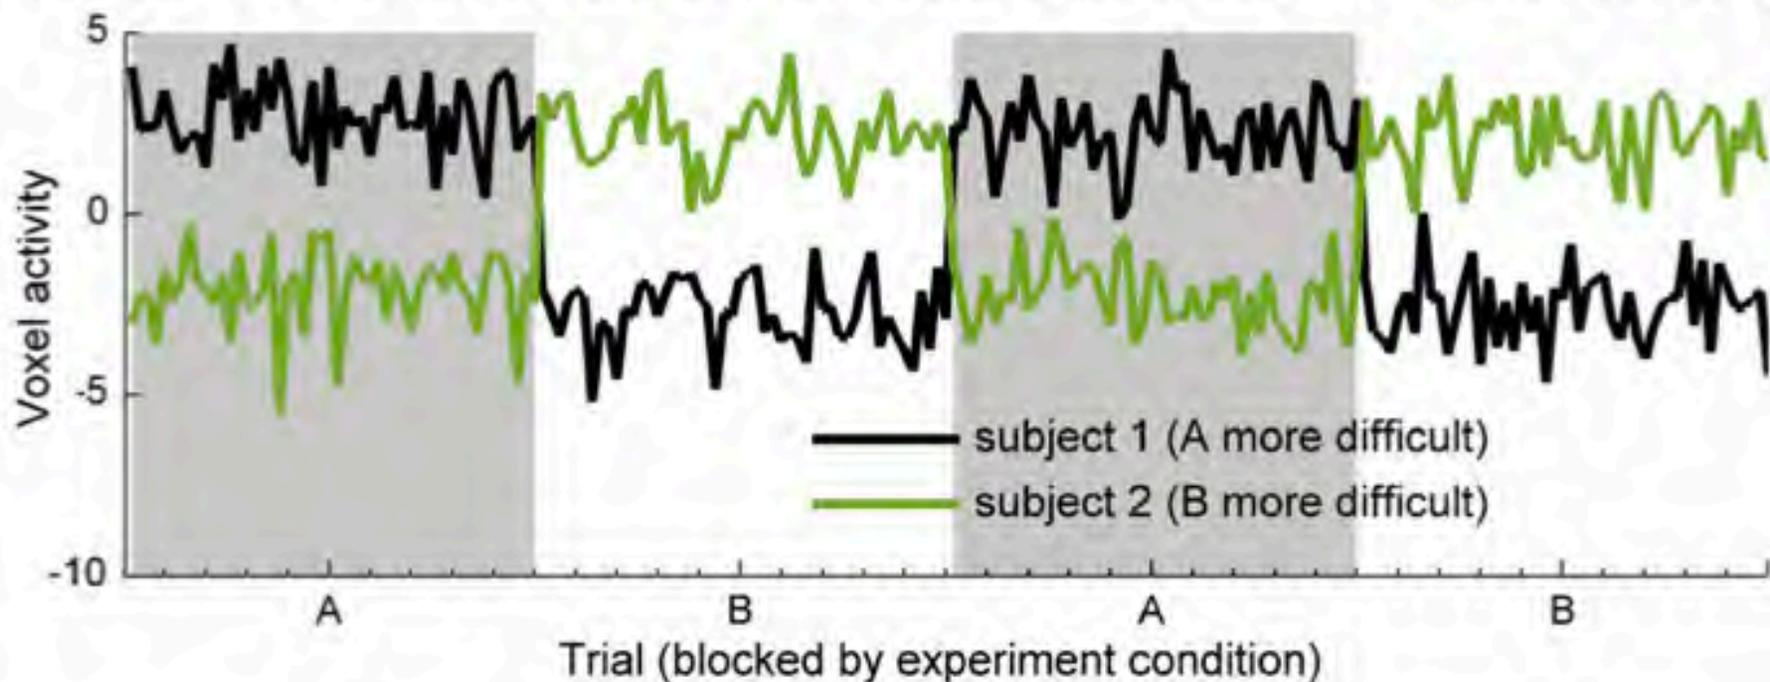
Born in West Virginia in 1980, The Neurocritic embarked upon a roadtrip across America at the age of thirteen with his mother. She abandoned him when they reached San Francisco and

when they reached San Francisco a
She abandoned him
across America at the age of thirteen

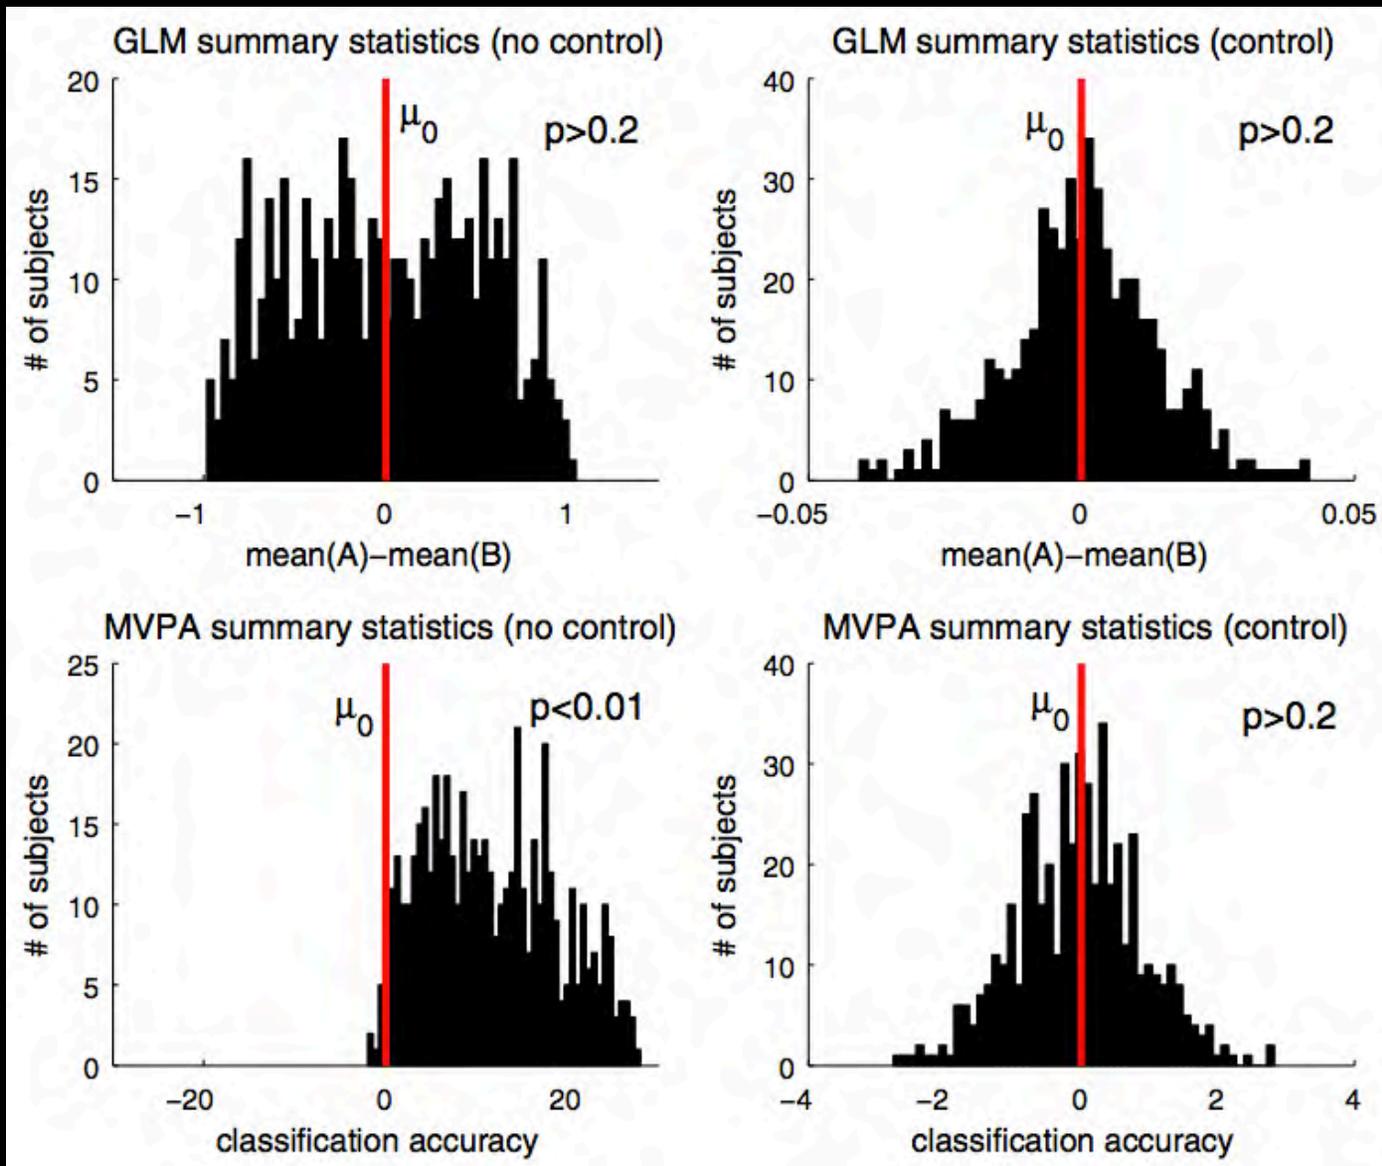
So what's the (potential) problem?

- An experimenter seeks to determine whether a neural signal differs across two experiment conditions (e.g., use of rule A vs. B to perform a task)
- Unknown to the experimenter, voxel activity is unresponsive to rule, but is responsive to difficulty.
- Furthermore, it happens incidentally that rule A is more difficult than rule B for some subjects, whereas the reverse is true for other subjects.
- Thus, at the individual-subject level, rule and difficulty are confounded.
- However, the task that is more difficult varies randomly across subjects, and therefore difficulty effects are approximately counterbalanced across subjects, and will cancel out when averaged in a group test.
- In contrast, MVPA will show robust discriminability at the individual subject level (due to brain activity effects related to task difficulty), so group-averaged maps will give the false impression of significant rule decoding.

Simulated example: experiment condition confounded with difficulty



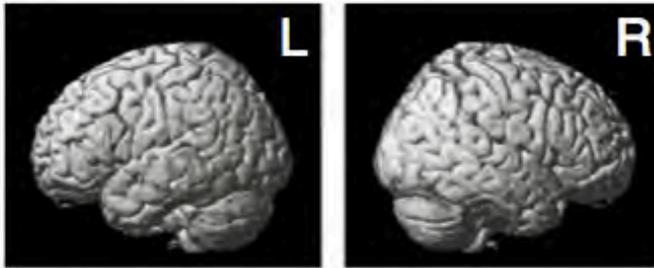
Controlling confounds with linear regression



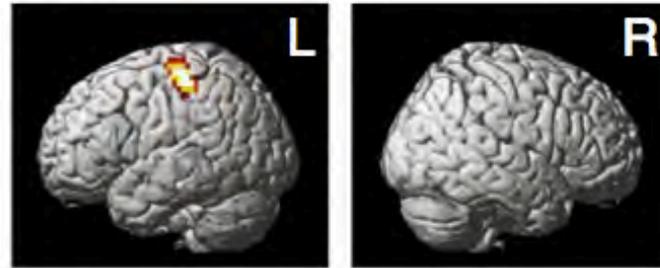
Task Rule (A vs. B)

Motor Response (Left vs. Right)

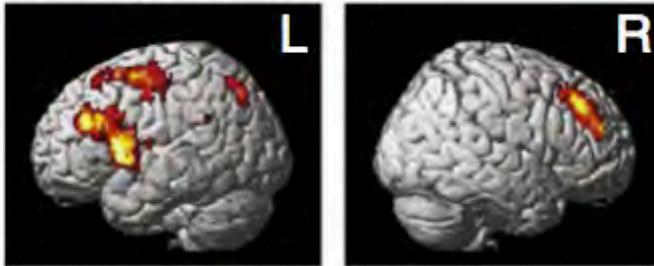
GLM



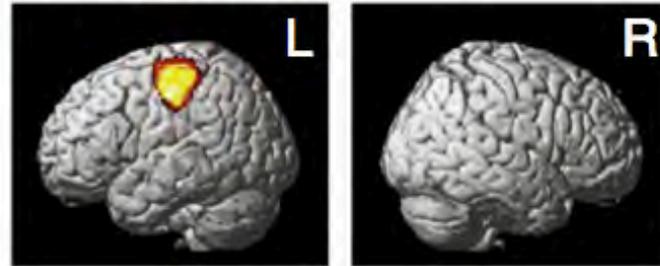
GLM



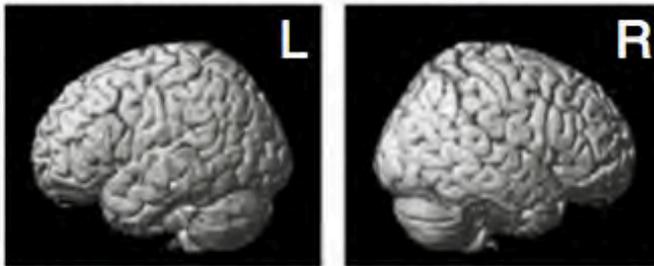
MVPA (no control)



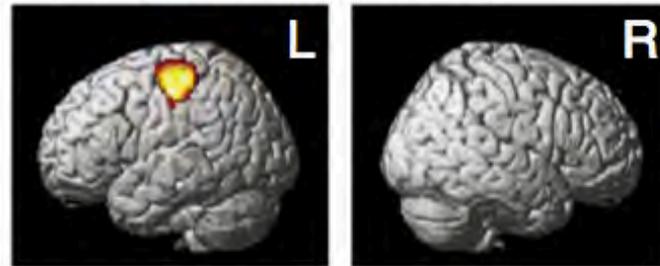
MVPA (no control)



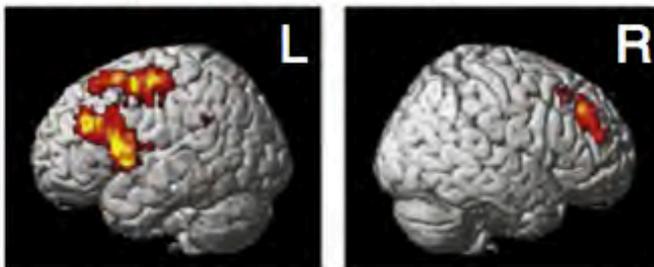
MVPA (after RT regression)



MVPA (after RT regression)



MVPA (after de-meaning)



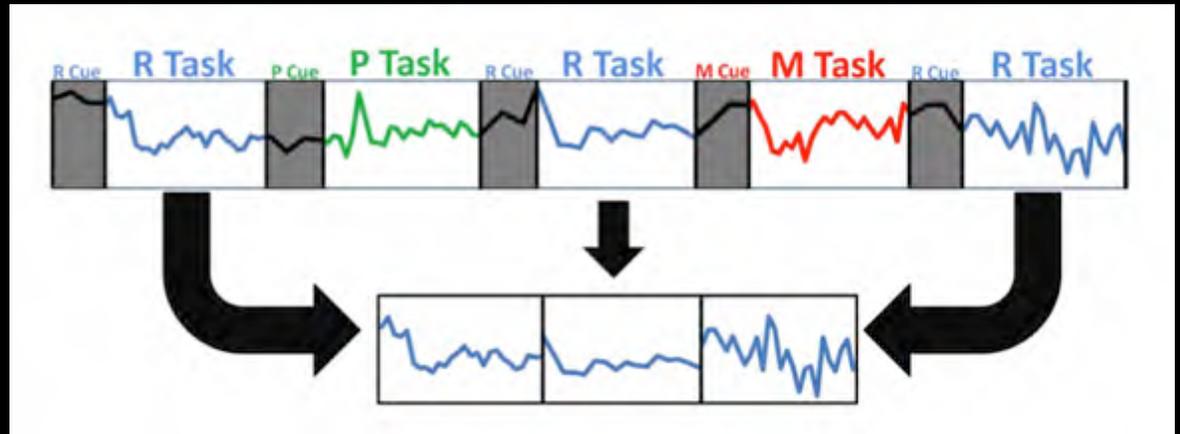
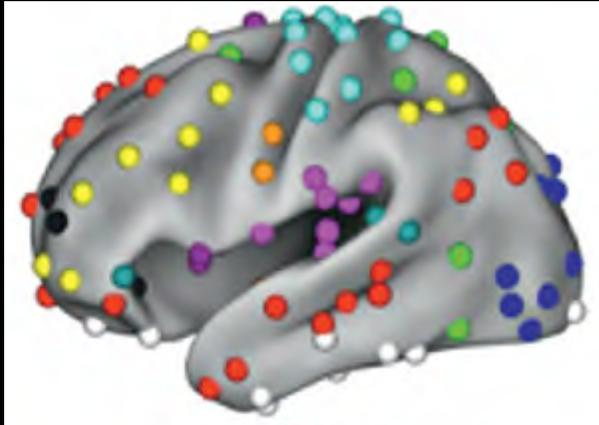
Todd et al.'s take-home message

- This concern is limited to cases in which a brain region revealed by MVPA is interpreted as representing a particular cognitive variable or type of information
 - “Brain region A represents information X”
- In contrast, claims of the following form are not problematic:
 - “Brain region A can predict behavior Y”
 - “The relationship between brain region's A and behavior Y follows model Z.”
- Recommended solution:
 - Item-by-item regression of RT (and/or other potential confounds) prior to MVPA

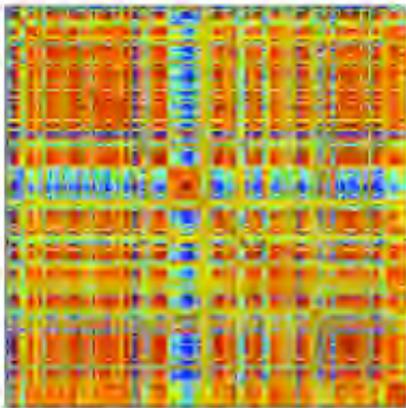
Taking the “V” out of MVPA

- Although the term MVPA implies that “Voxel” activity patterns must be the features used for decoding, the approach can also work well with **functional connectivity patterns**
- Cole, M. W., Reynolds, J. R., Power, J. D., Repovs, G., Anticevic, A., & Braver, T. S. (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature Neuroscience*, 16(9):1348-55.
- Dosenbach, N. U. F., Nardos, B., Cohen, A. L., Fair, D. A., Power, J. D., Church, J. A., et al. (2010). Prediction of individual brain maturity using fMRI. *Science*, 329(5997), 1358–1361.
- Ekman, M., Derrfuss, J., Tittgemeyer, M., & Fiebach, C. J. (2012). Predicting errors from reconfiguration patterns in human brain networks. *PNAS*, 109(41), 16714–16719.
- Fair, D. A., Nigg, J. T., Iyer, S., Bathula, D., Mills, K. L., Dosenbach, N. U. F., et al. (2012). Distinct neural signatures detected for ADHD subtypes after controlling for micro-movements in resting state functional connectivity MRI data. *Frontiers in Systems Neuroscience*, 6, 80.
- Heinze, J., Wenzel, M. A., & Haynes, J.-D. (2012). Visuomotor functional network topology predicts upcoming tasks. *The Journal of Neuroscience*, 32(29), 9960–9968.
- Pantazatos, S. P., Talati, A., Pavlidis, P., & Hirsch, J. (2012). Decoding unattended fearful faces with whole-brain correlations: an approach to identify condition-dependent large-scale functional connectivity. *PLoS Computational Biology*, 8(3), e1002441.
- Shirer, W. R., Ryali, S., Rykhlevskaia, E., Menon, V., & Greicius, M. D. (2012). Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex*, 22(1), 158–165.

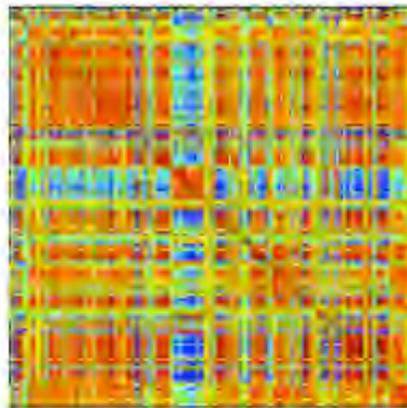
Connectivity-based decoding



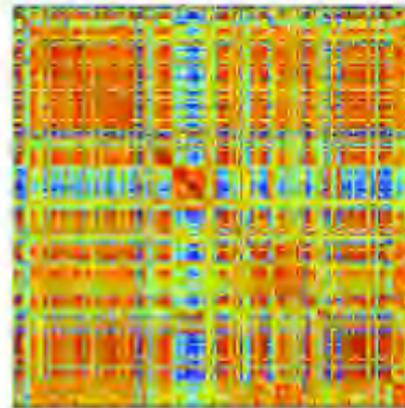
Memory



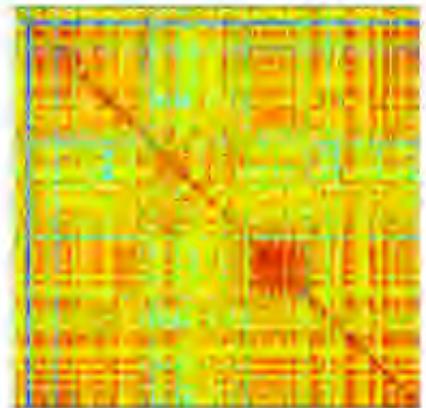
Reasoning



Perception

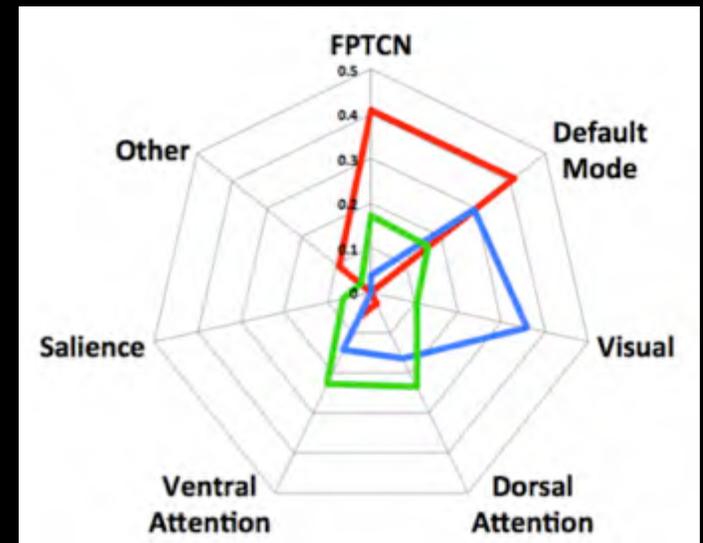
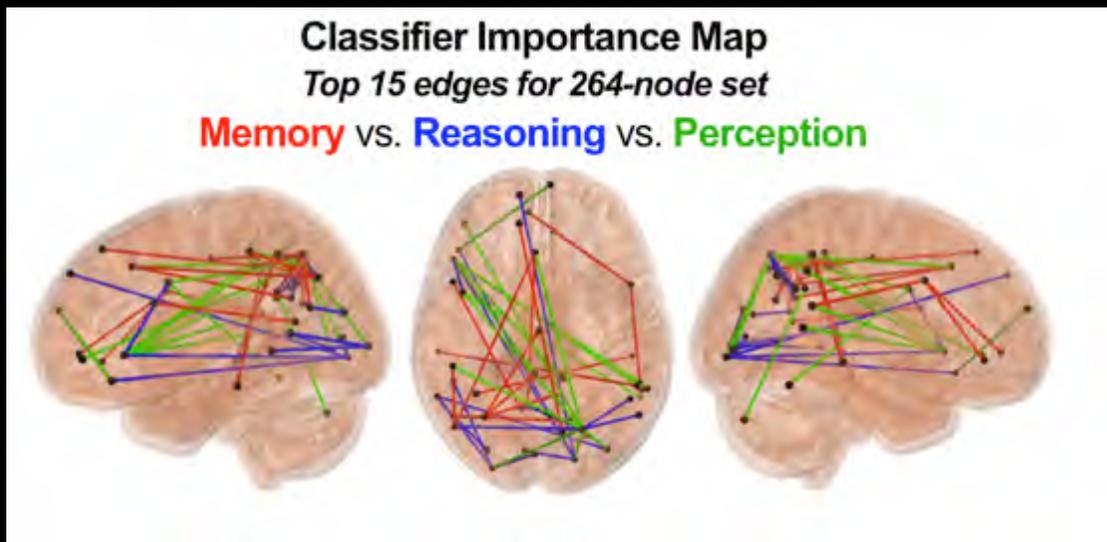


Resting



Westphal, Reggente, & Rissman (in prep)

Connectivity-based decoding



representational similarity analysis

goal:

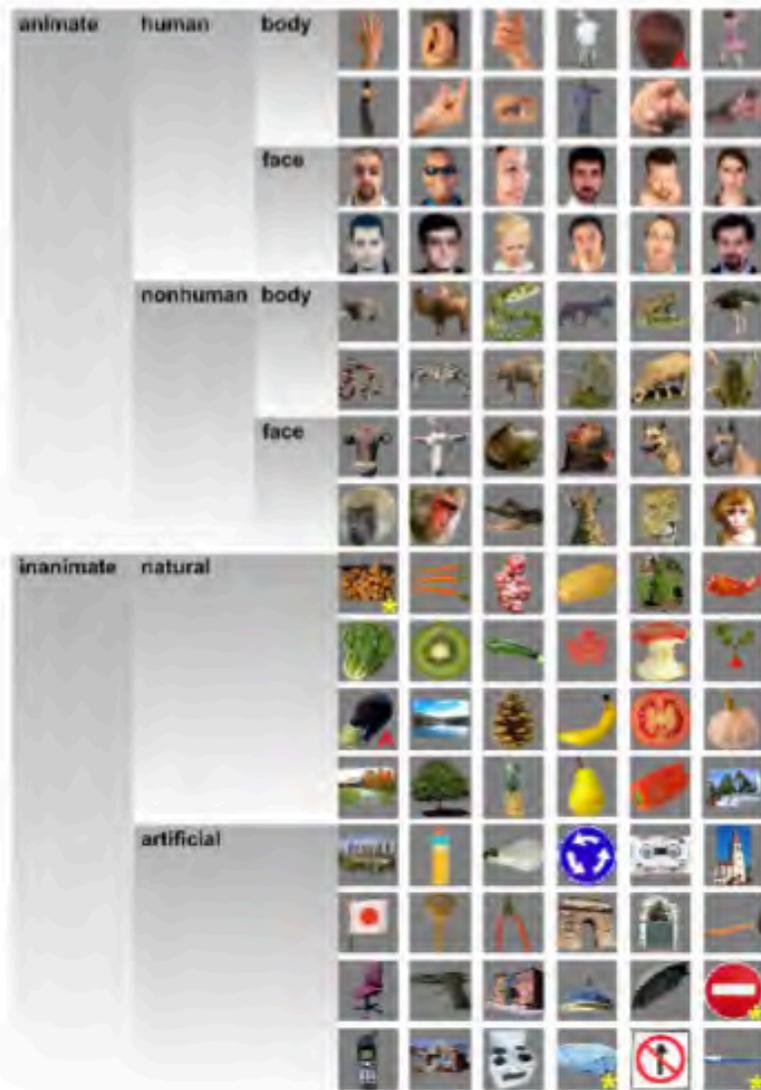
- compare fMRI patterns elicited by stimuli without assuming a priori structure
- relate similarity/distance in fMRI space to behavioural or hypothesized similarity/distance

- a range of ideas more than a technique

Kriegeskorte et al. (2008; 2009)

[slide from Francisco Pereira]

stimuli and similarity



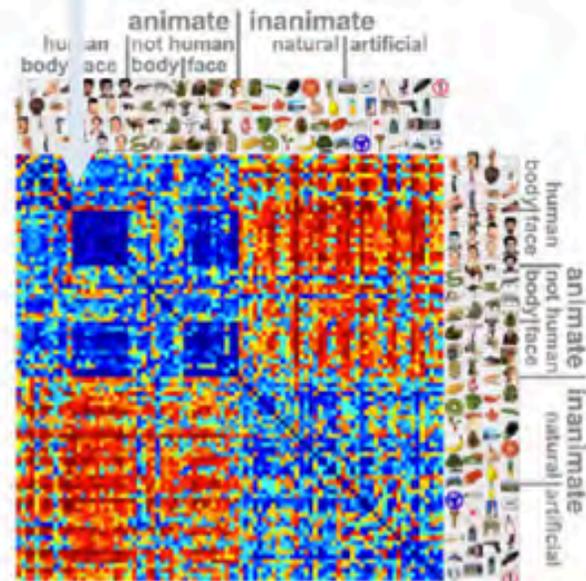
correlation



stimulus 1 pattern



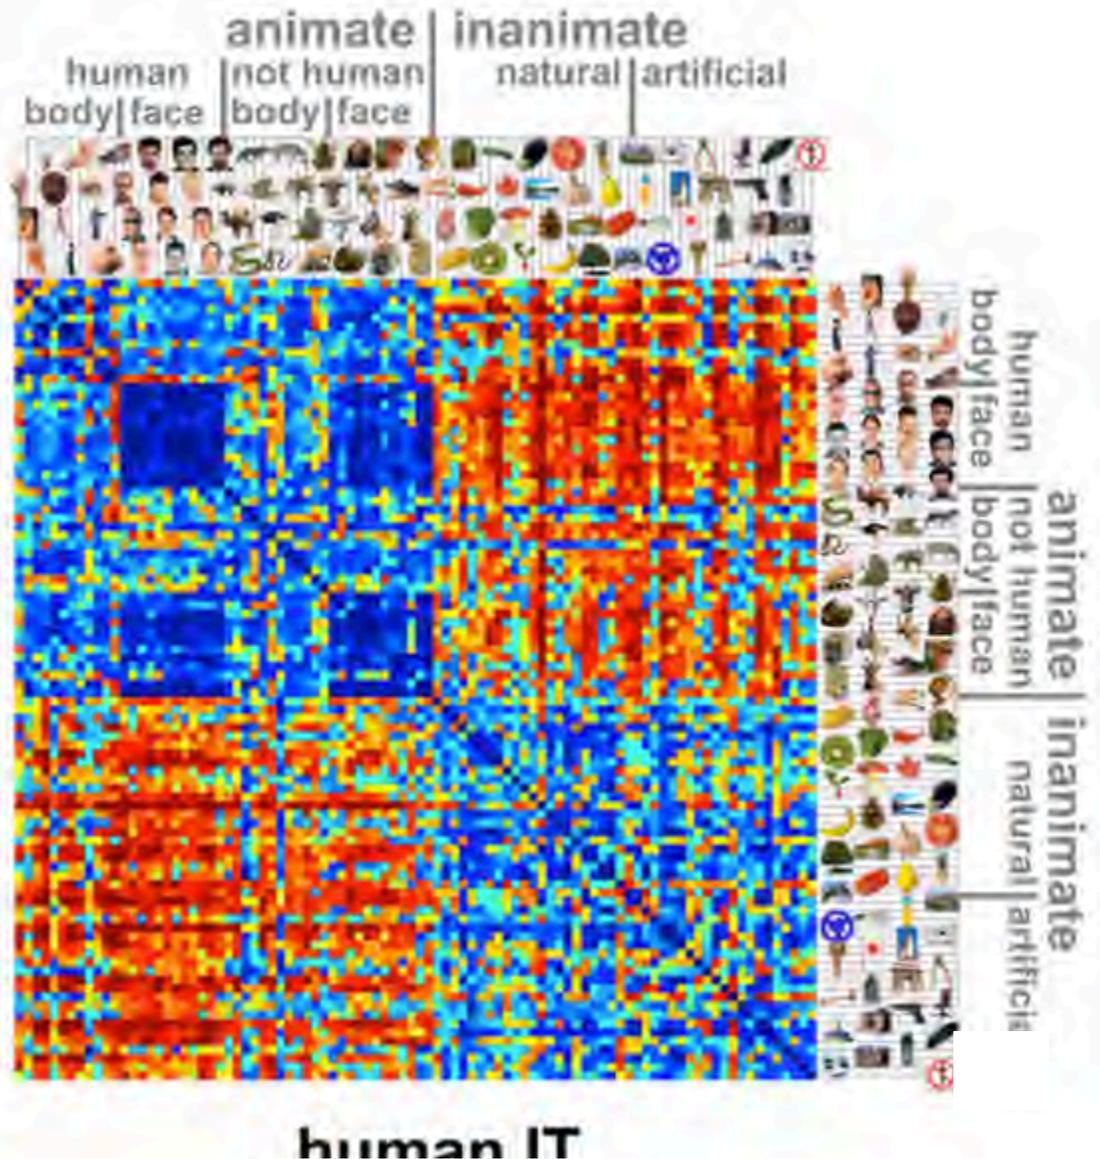
stimulus 2 pattern



human IT

stimuli and similarity

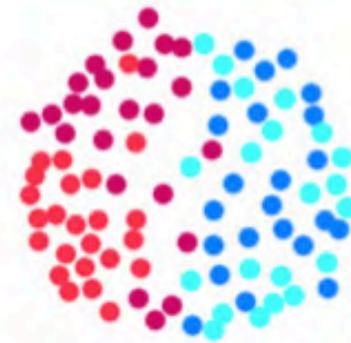
- summarizes a lot of information
- a bit hard to make sense of



Kriegeskorte et al. (2008; 2009)

[slide from Francisco Pereira]

multidimensional scaling



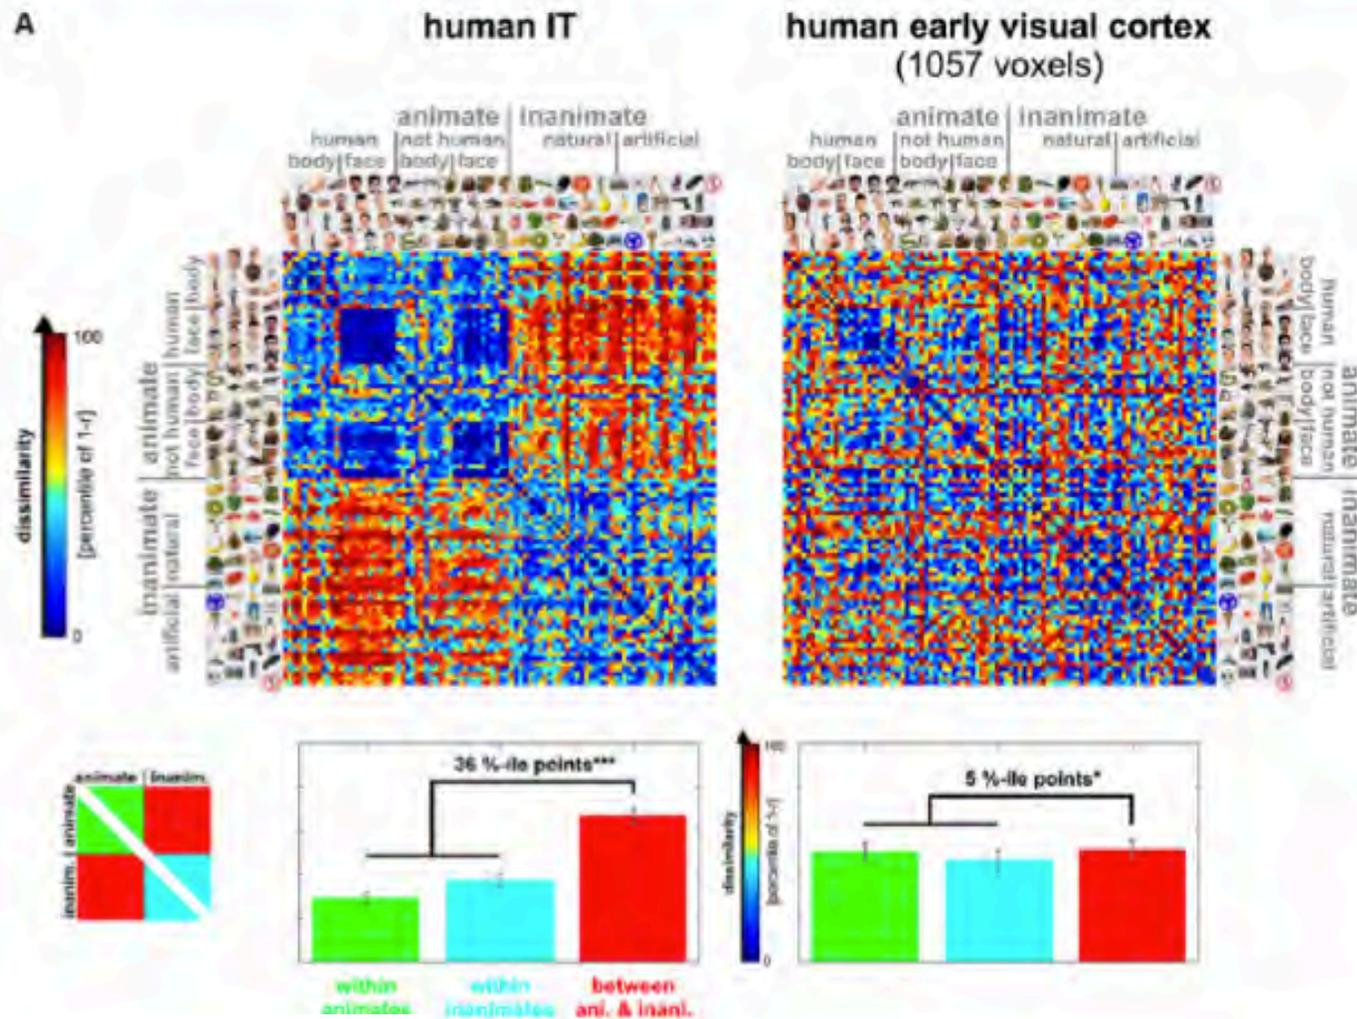
body
face
natural obj.
artificial obj.

human IT

Kriegeskorte et al. (2008; 2009)

slide from Francisco Pereira

compare multiple areas

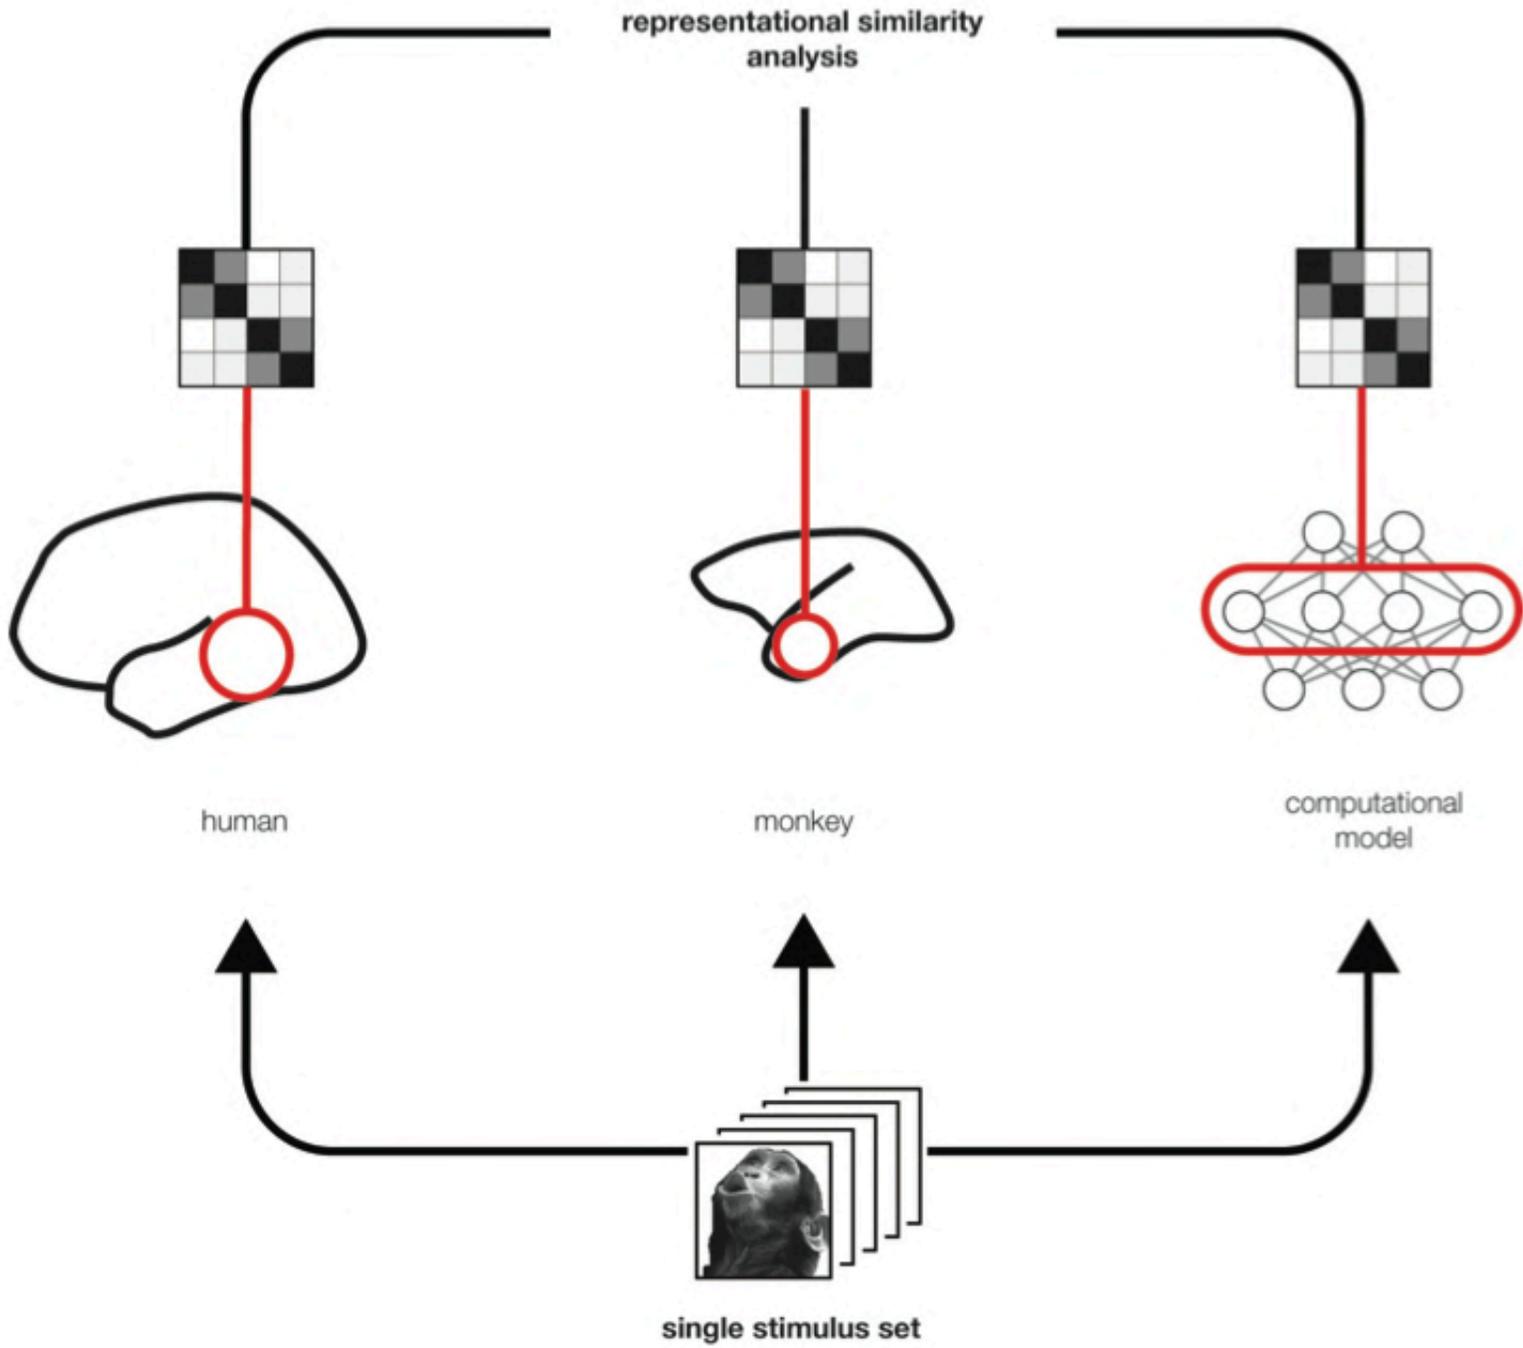


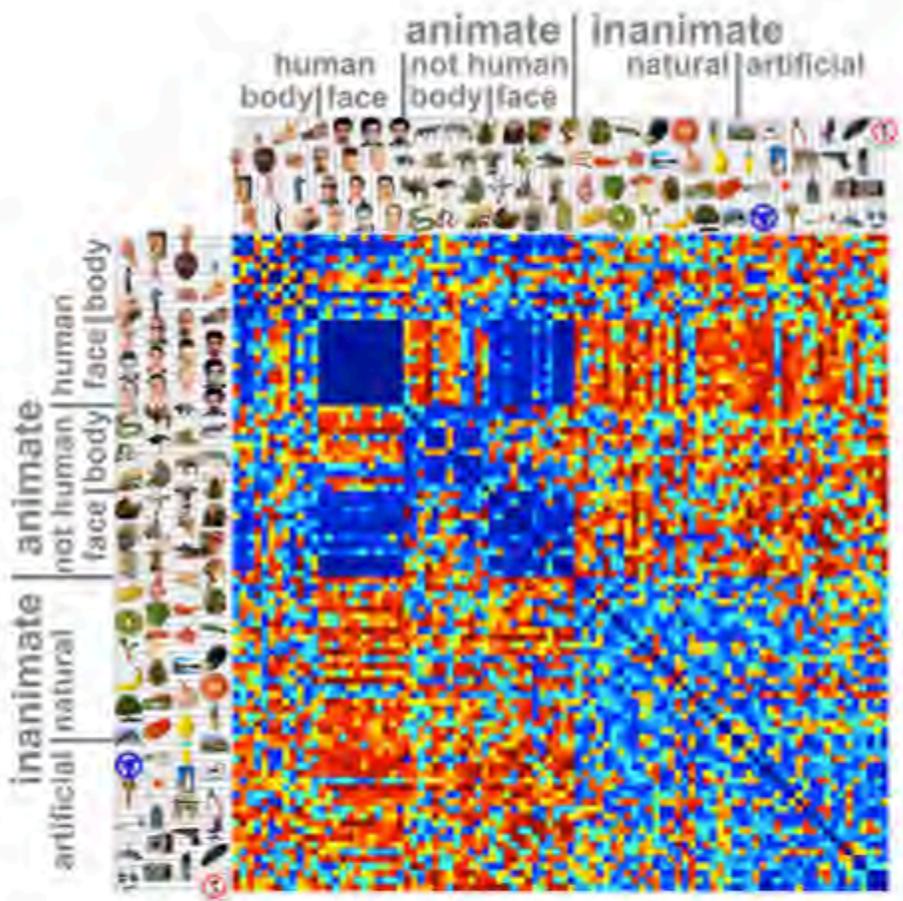
Kriegeskorte et al. (2008; 2009)

[slide from Francisco Pereira]

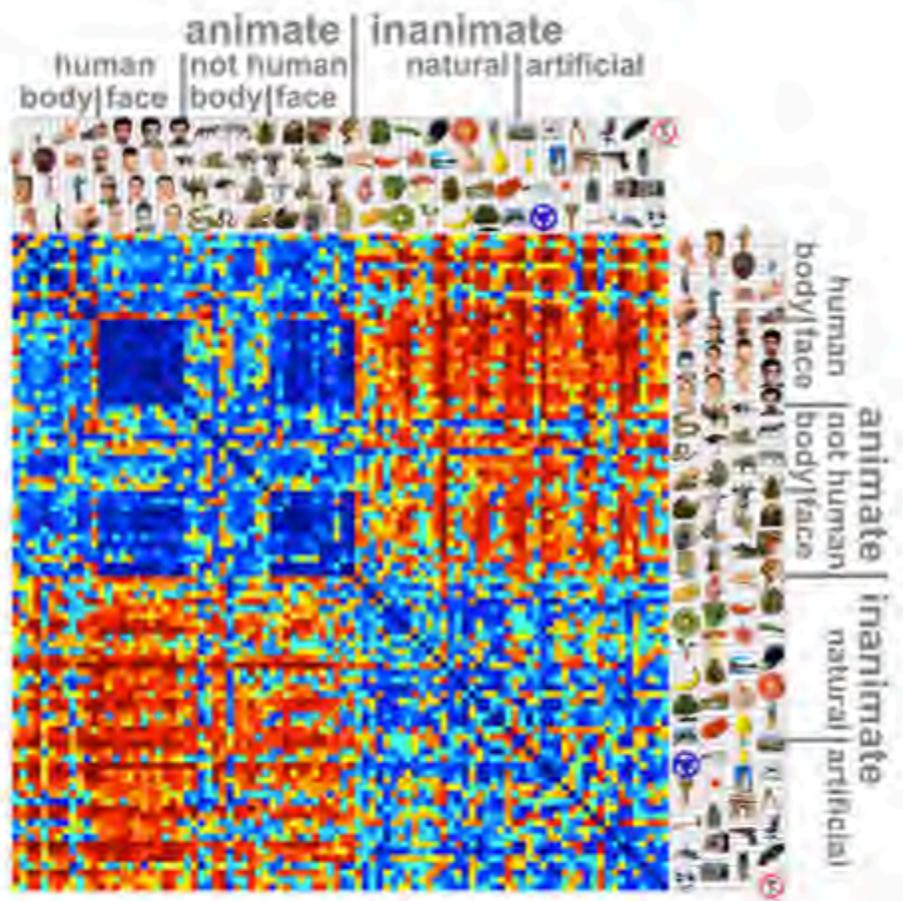
other aspects

- can compute matrices in different ROIs
- can look for matrices compatible with a certain hypothesis (or use it to compare competing ones)
- matrices matching behaviour, ratings, etc





monkey IT



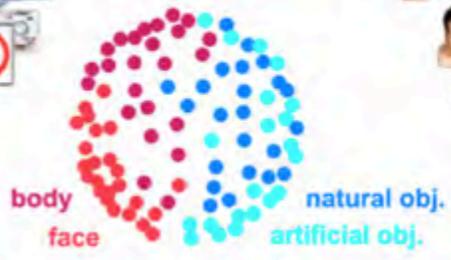
human IT



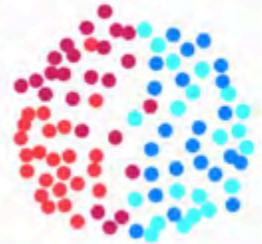
A

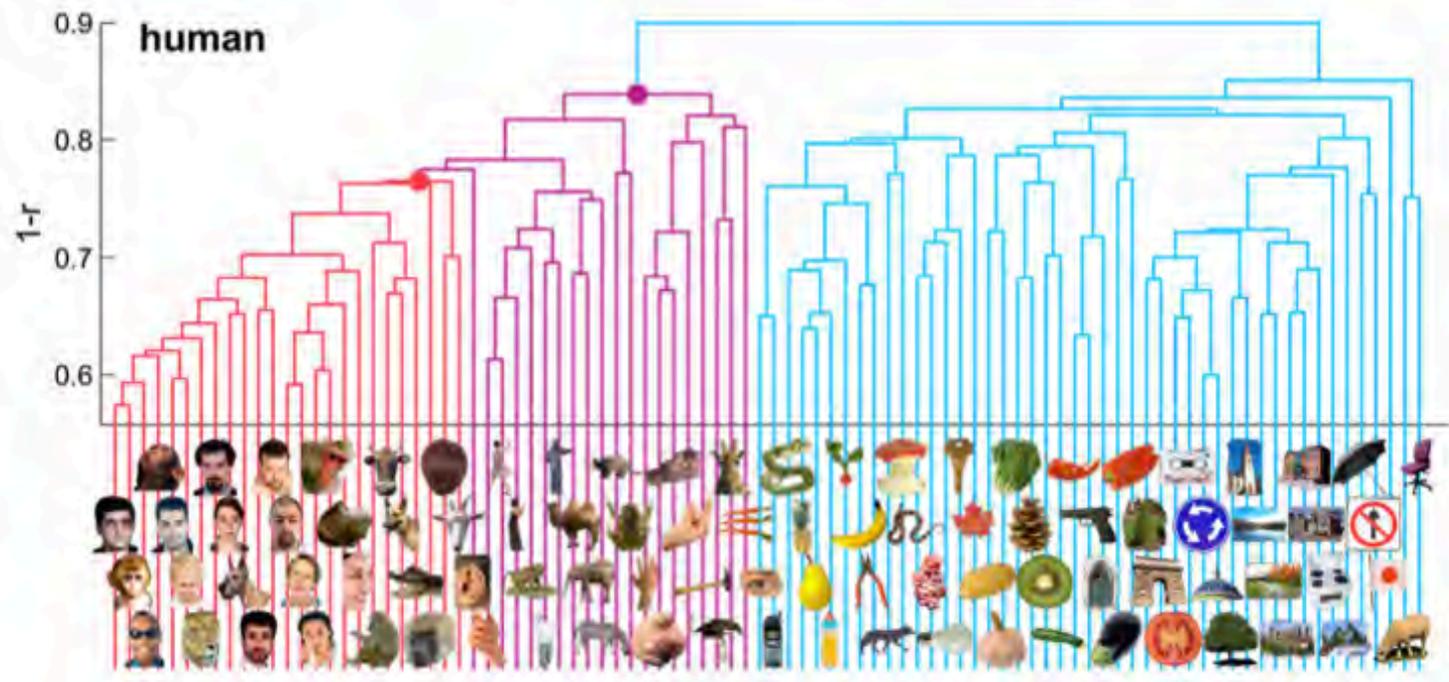
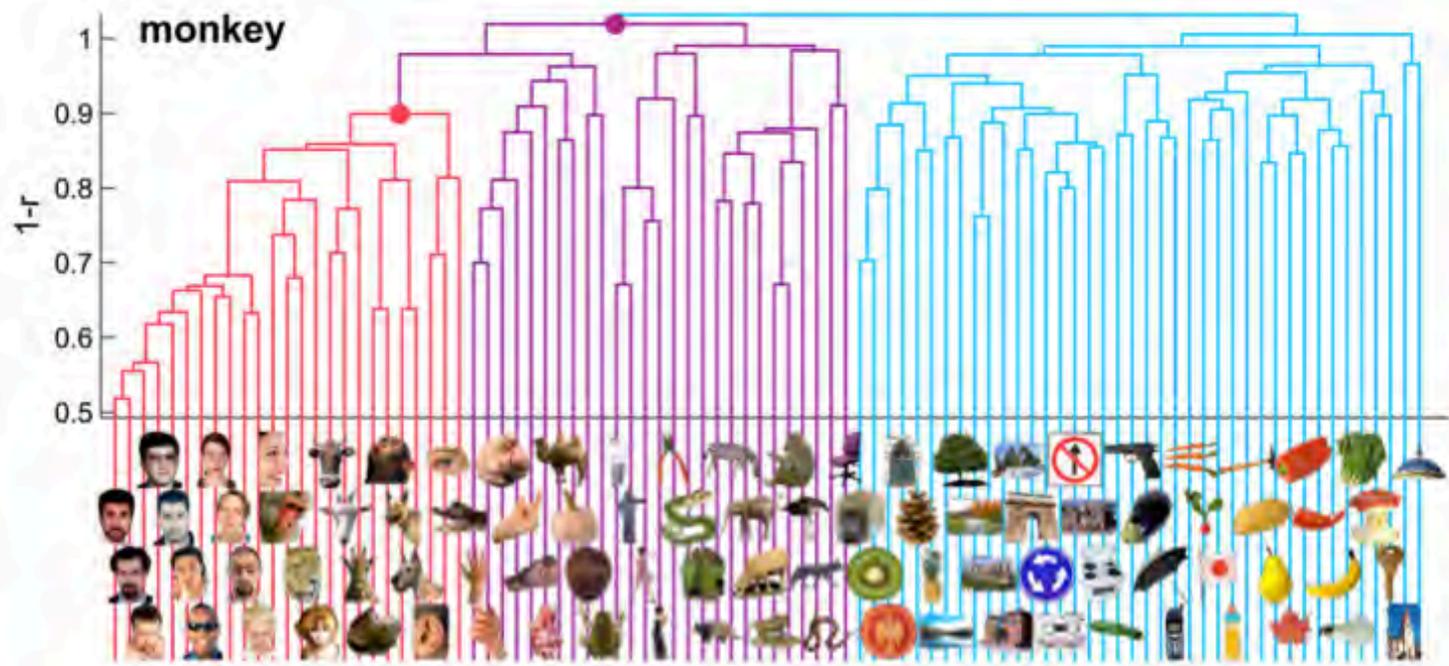


monkey IT



human IT





A Toolbox for Representational Similarity Analysis

**Hamed Nili^{1*}, Cai Wingfield², Alexander Walther¹, Li Su^{1,3}, William Marslen-Wilson³,
Nikolaus Kriegeskorte^{1*}**

¹ MRC Cognition and Brain Sciences Unit, Cambridge, United Kingdom, ² Department of Computer Science, University of Bath, Bath, United Kingdom, ³ Department of Experimental Psychology, University of Cambridge, Cambridge, United Kingdom

<http://www.mrc-cbu.cam.ac.uk/methods-and-resources/toolboxes/license/>