

Multi-voxel pattern analysis:

Decoding Mental States from fMRI Activity Patterns



Artwork by Leon Zernitsky

Jesse Rissman

NITP Summer Program 2013

Mind Reading

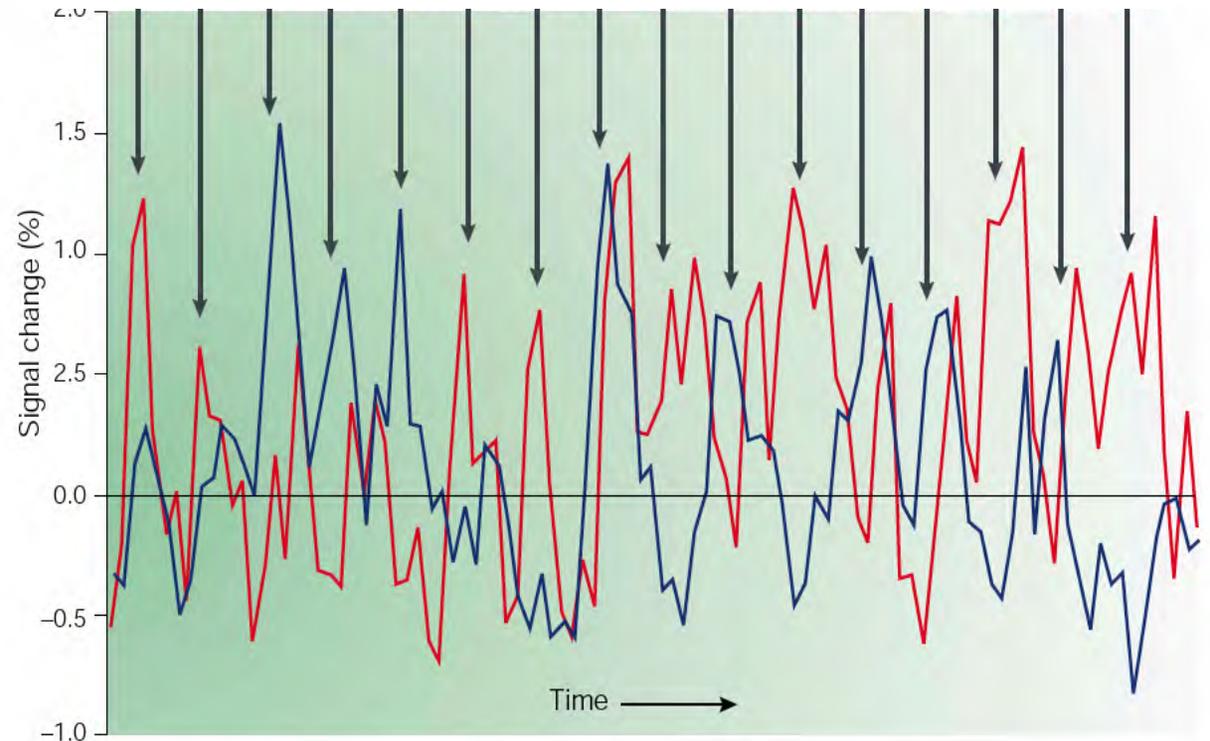
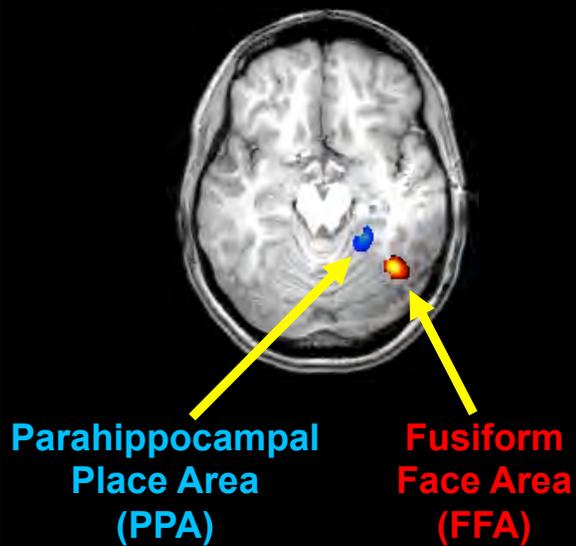
produced by
Shari Finkelstein



Some 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

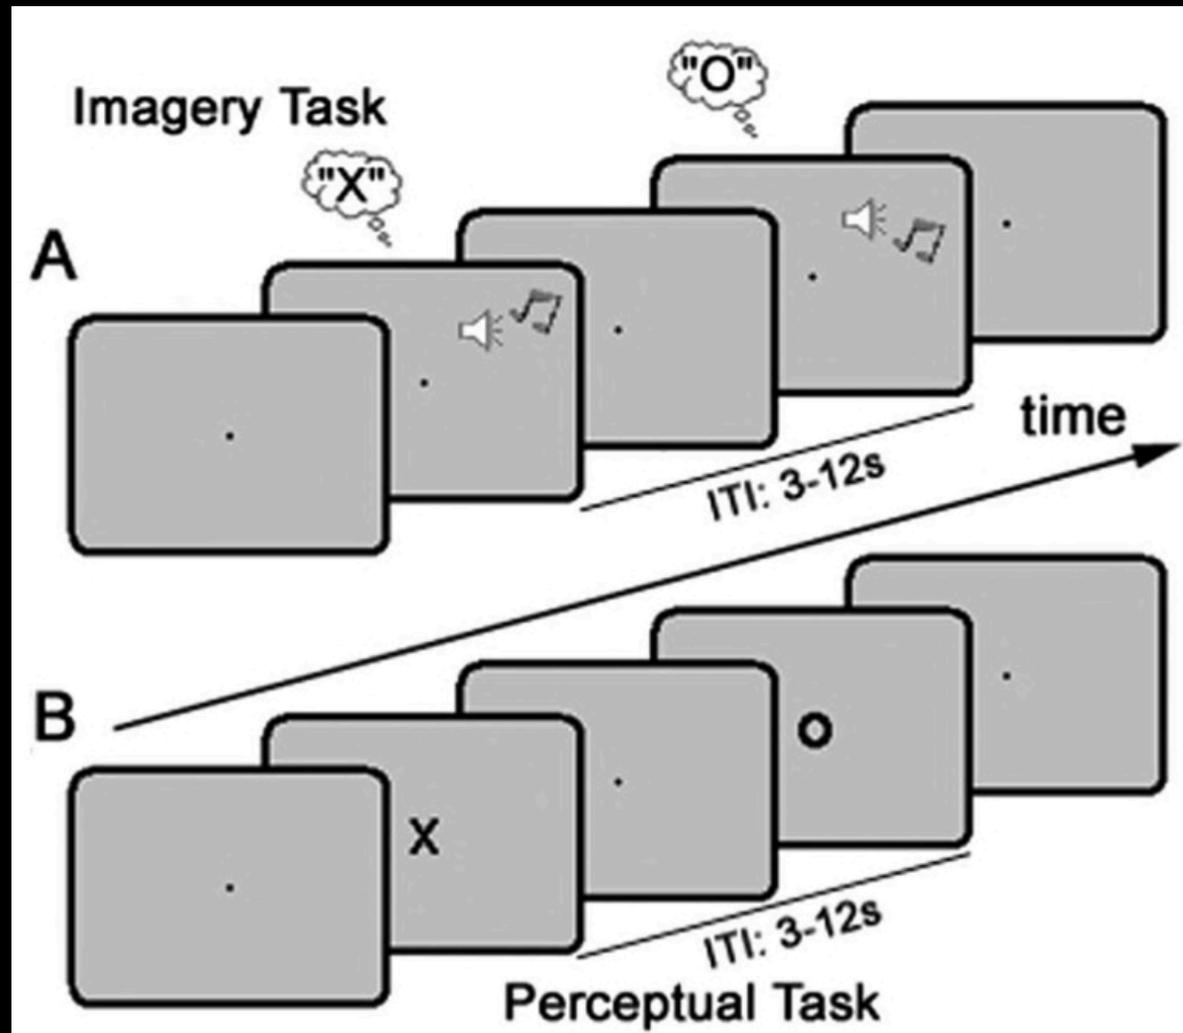
An Early Attempt to Decode Individual Thoughts



85% correct classification

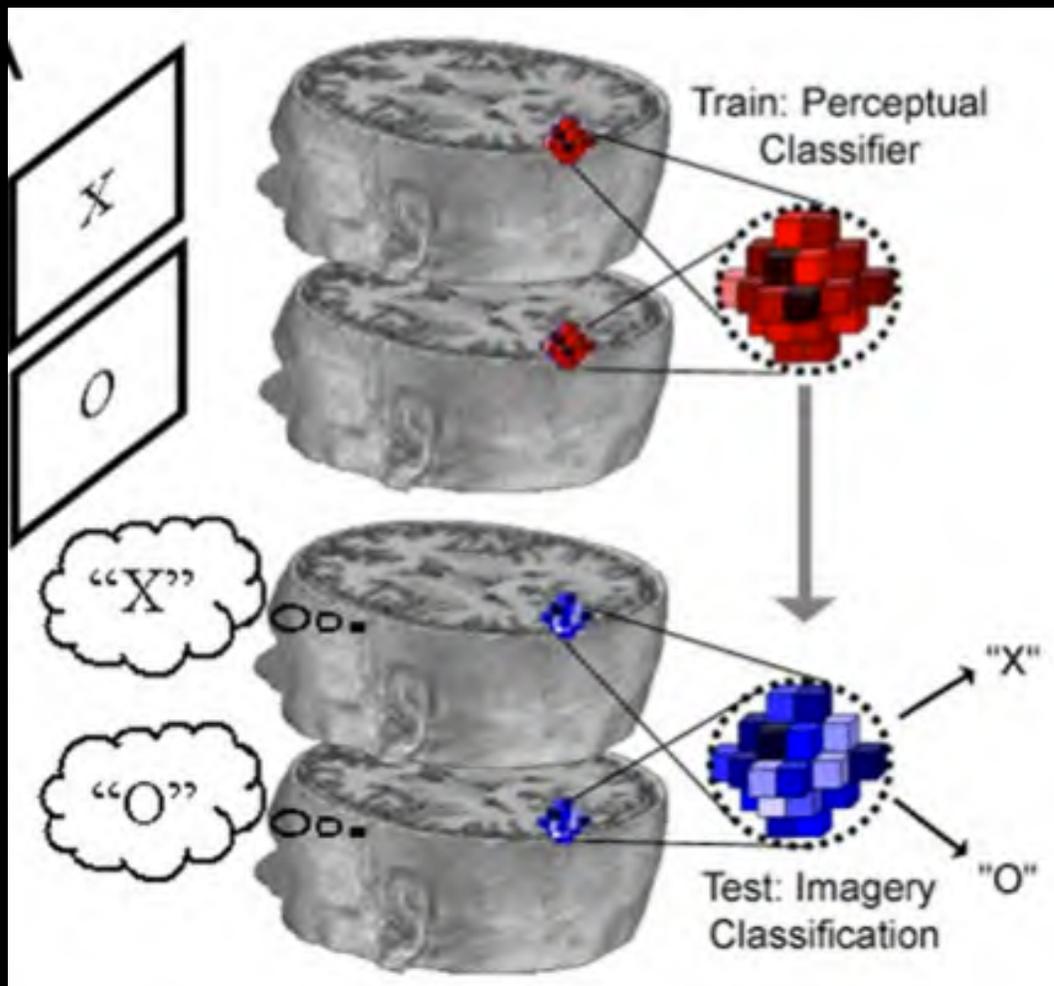
O' Craven & Kanwisher, 2000

Decoding Finer-grained Mental Images



The Power of Patterns:

Multi-voxel activity patterns differentiate mental imagery of X's vs. O's

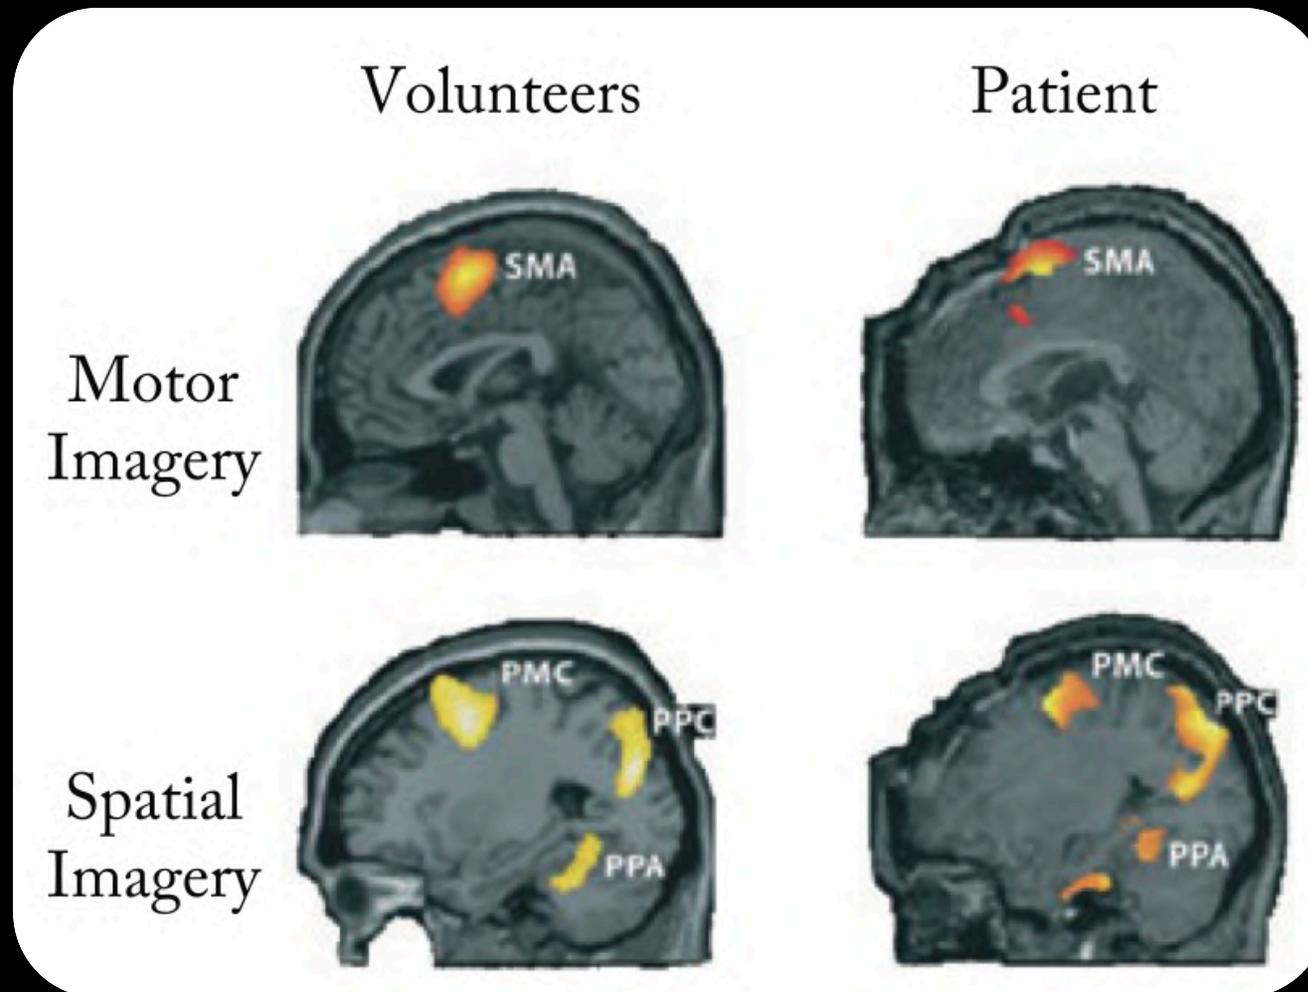


Activity patterns in left lateral occipital region could reliably predict what the subject was imagining

(well, at least 62% of the time)

Before diving into the details,
let me show you just two recent
applications of mental imagery
decoding...

Mental imagery-induced fMRI activity as a means of evaluating consciousness and facilitating communication in minimally conscious patients



Monti et al. (2010) *New England Journal of Medicine*

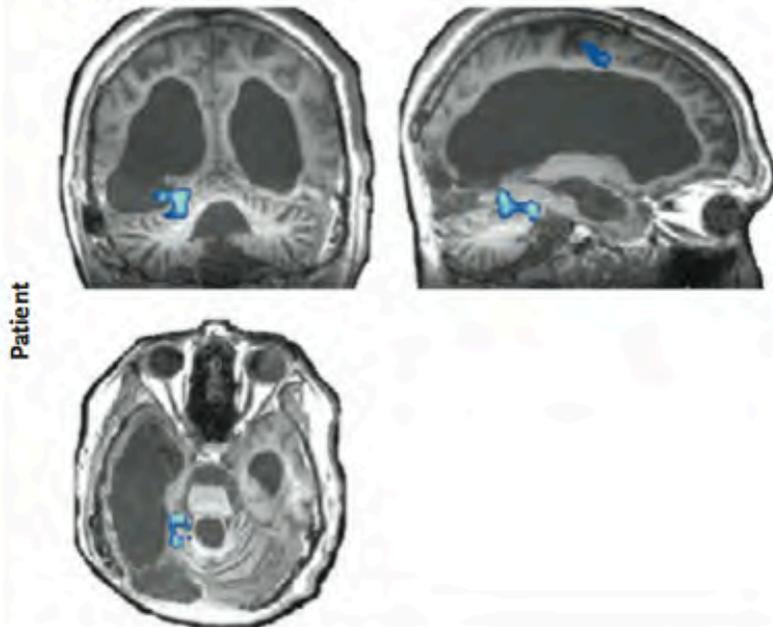
A "Is your father's name Alexander?" "Yes" response with the use of motor imagery



B "Do you have any brothers?" "Yes" response with the use of motor imagery



C "Is your father's name Thomas?" "No" response with the use of spatial imagery



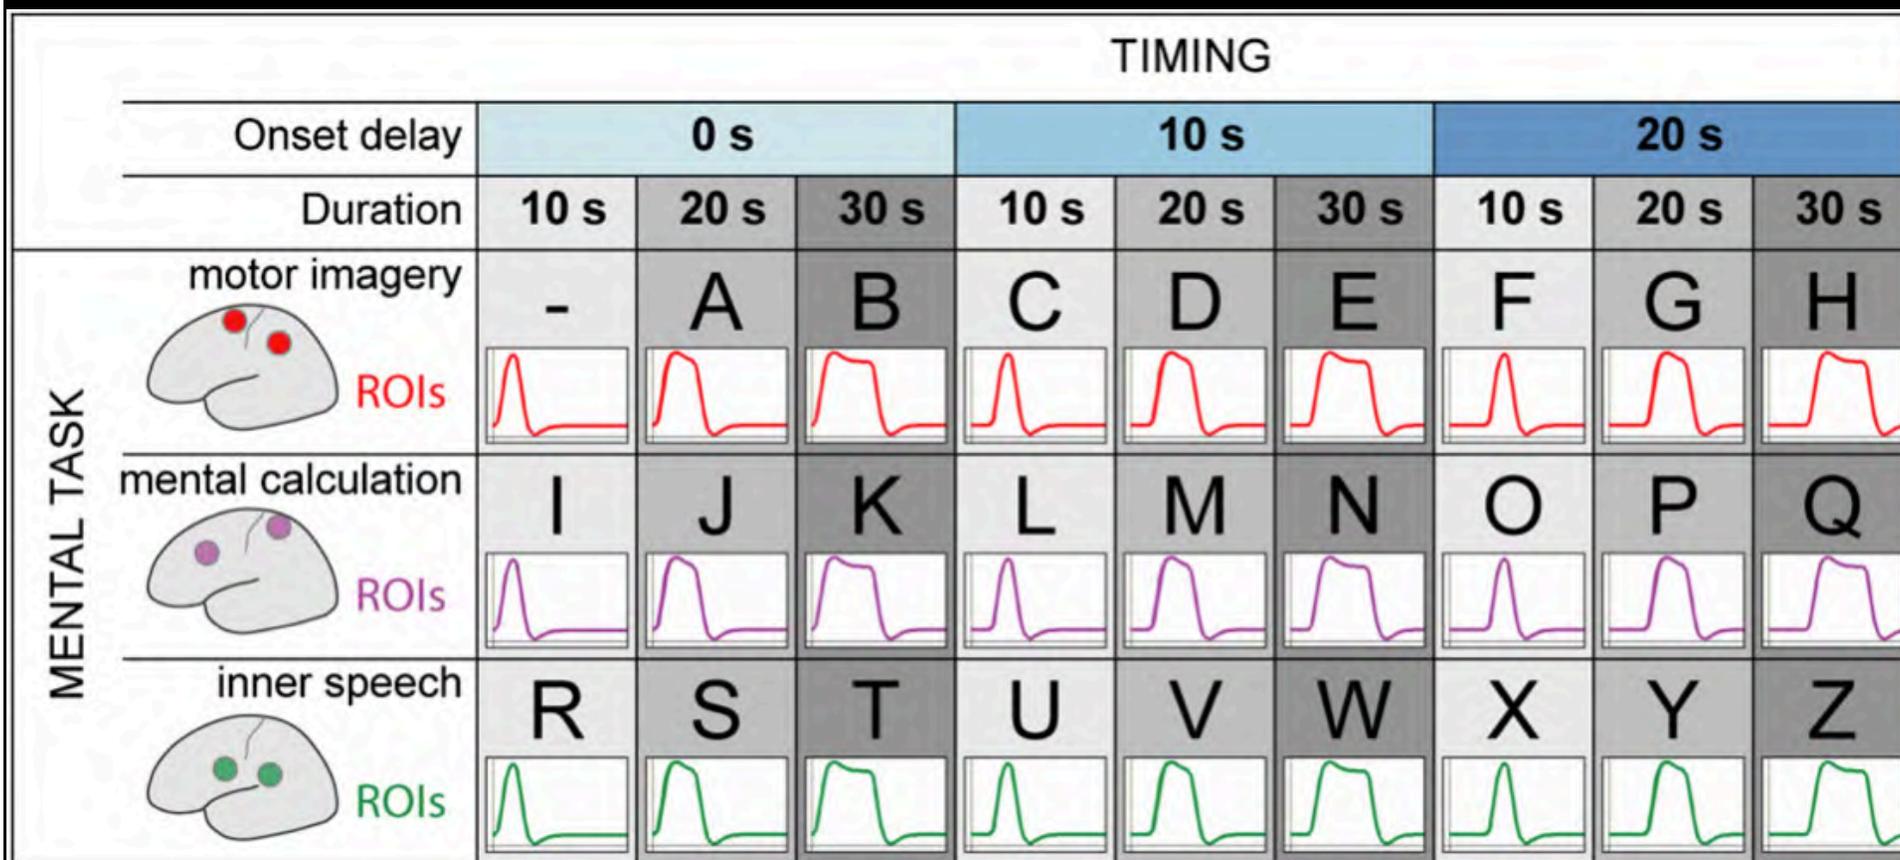
D "Do you have any sisters?" "No" response with the use of spatial imagery



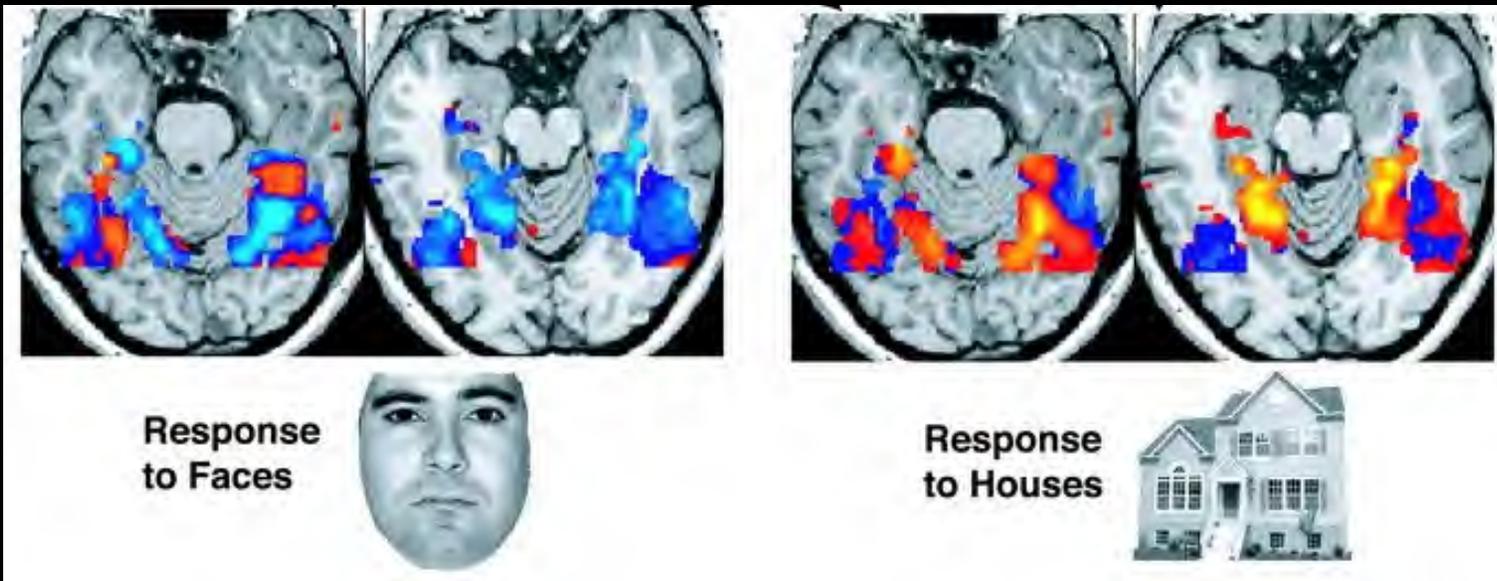
Related applications:
A real-time fMRI-based spelling device

Word encoding and automated letter decoding

New applications: A real-time fMRI based spelling device

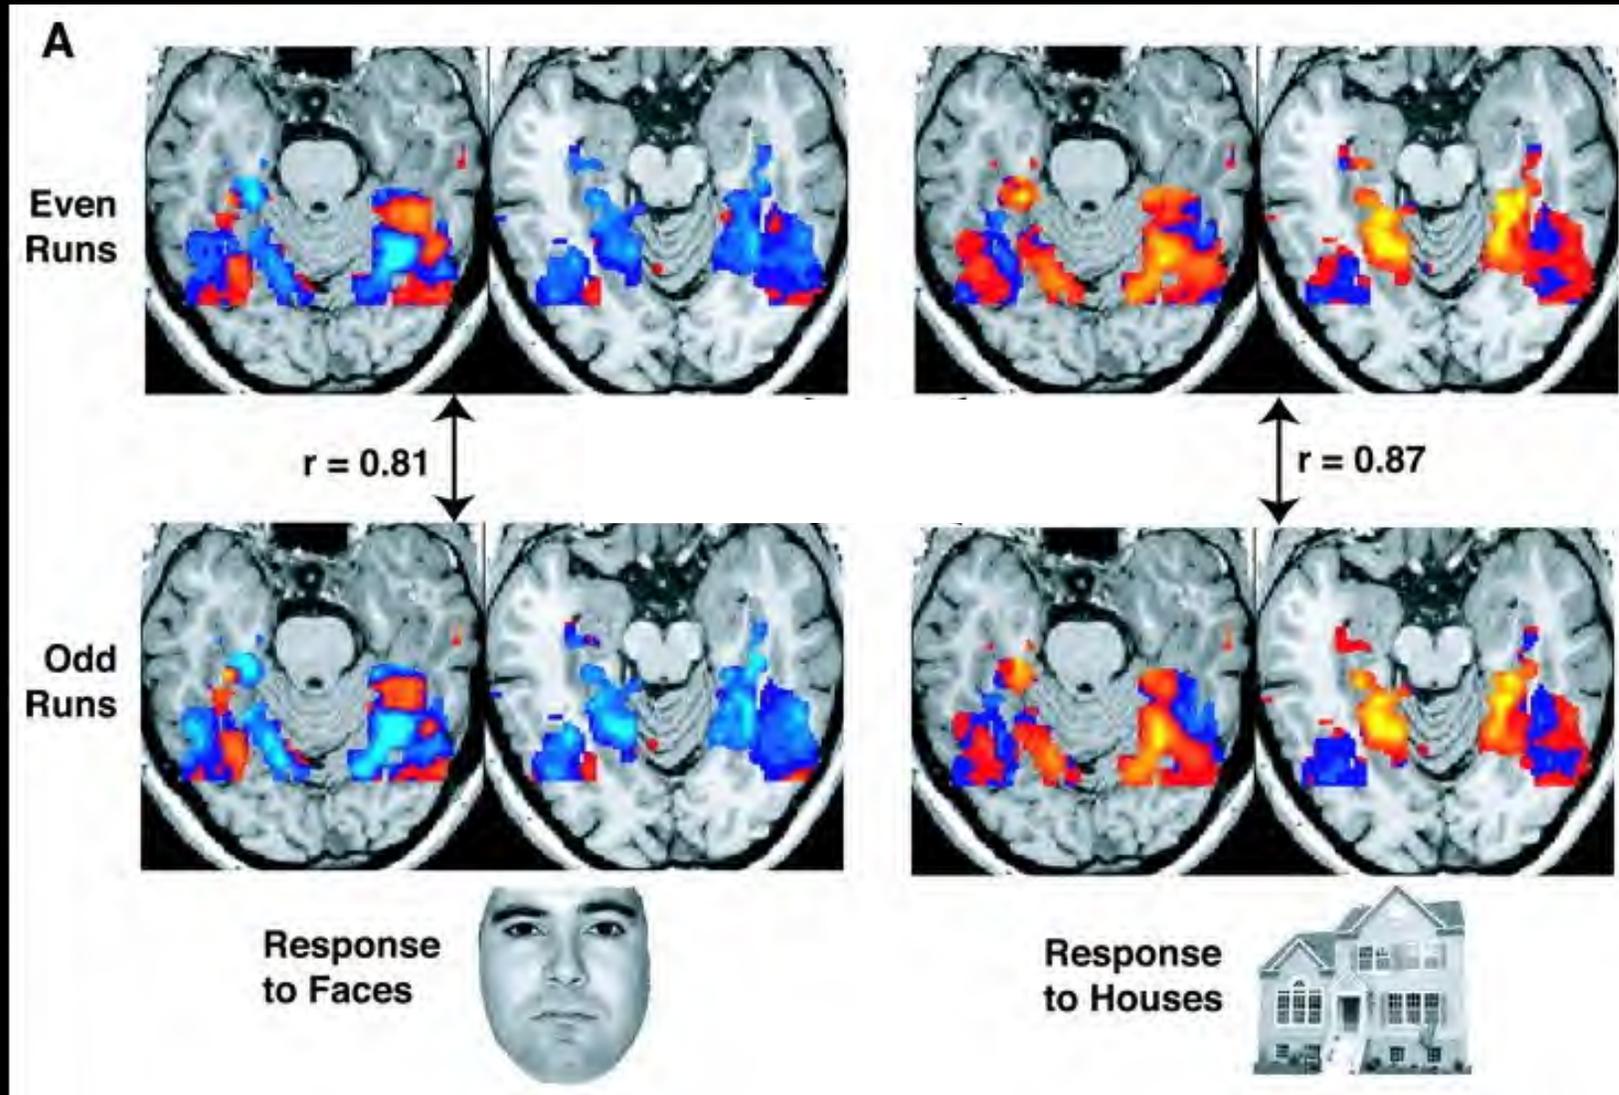


How it all began...



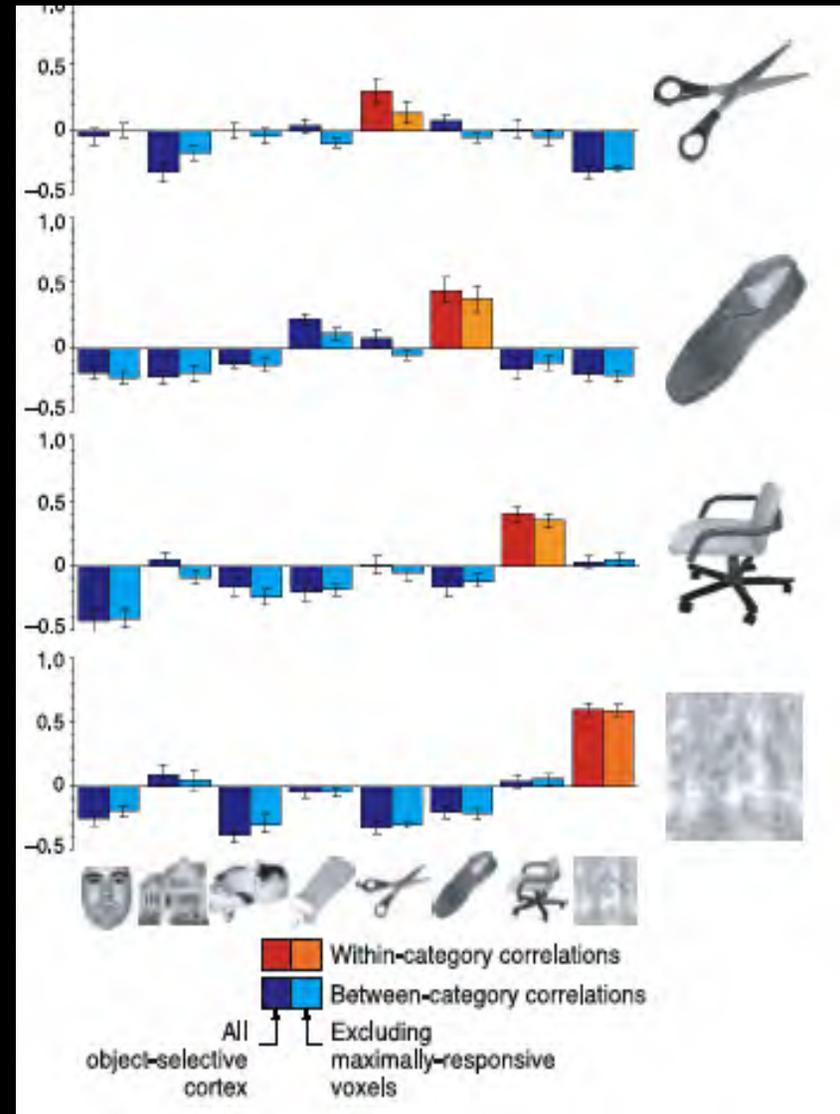
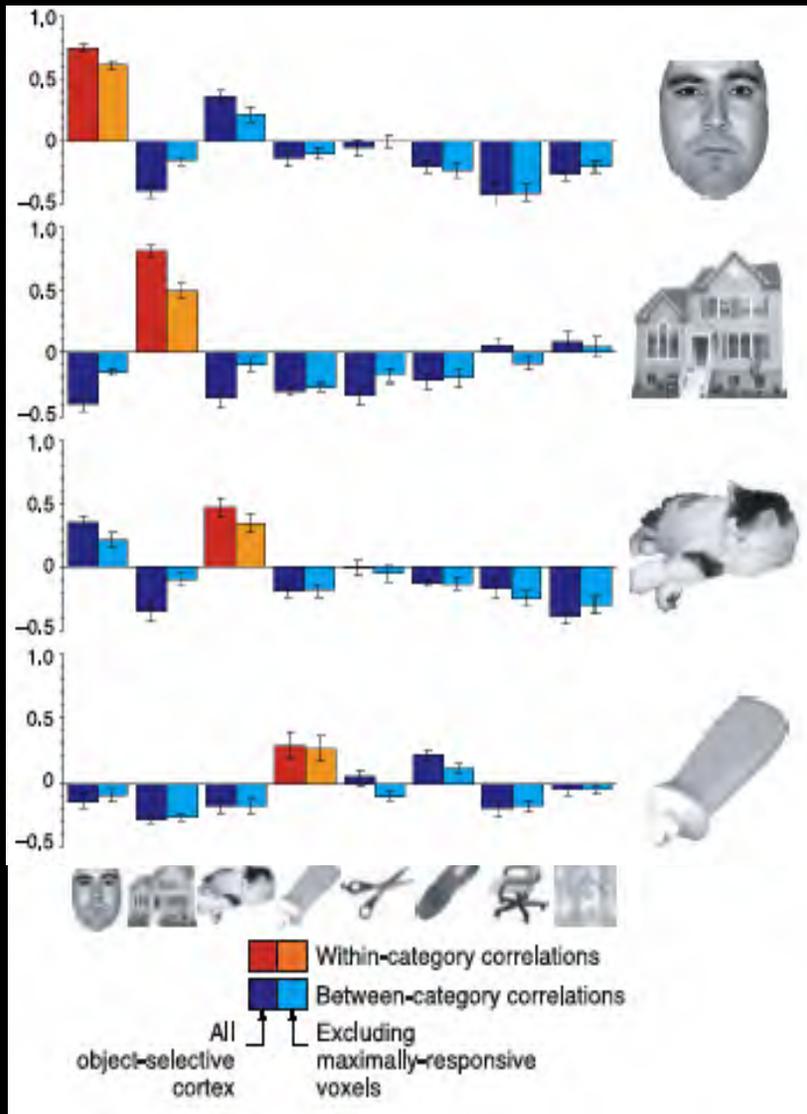
Haxby et al. (2001), *Science*

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}

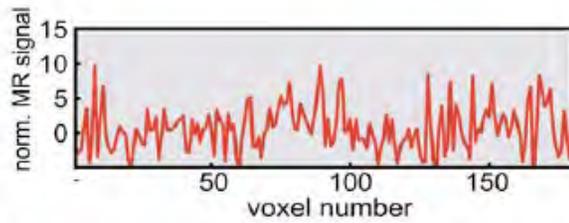
^a Rowland Institute for Science, Cambridge, MA 02142, USA

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

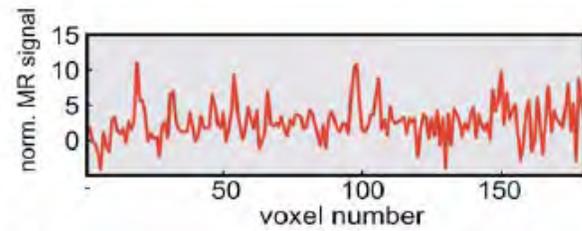
^c HyperVision, Inc., P.O. Box 158, Lexington, MA 02420, USA

Cox & Savoy (2003), *NeuroImage*

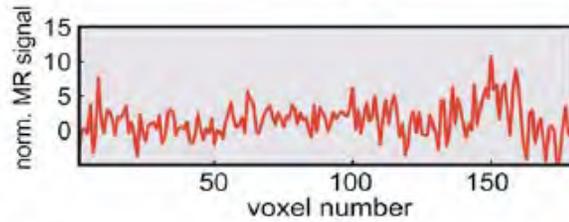
Baskets



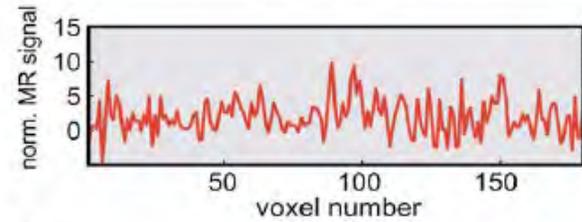
Birds



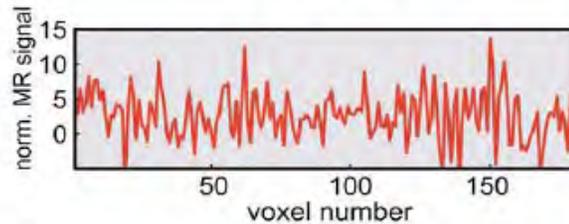
Butterflies



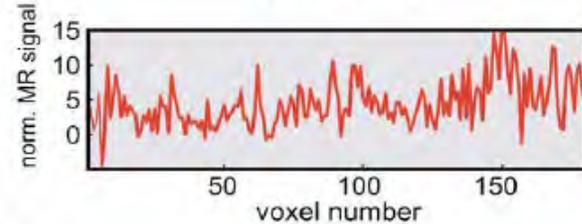
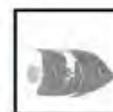
Chairs



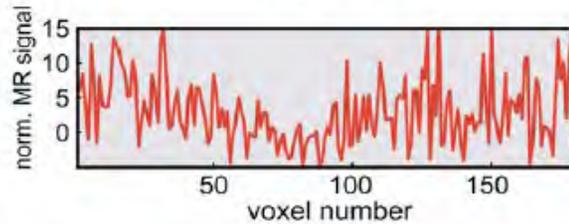
Cows



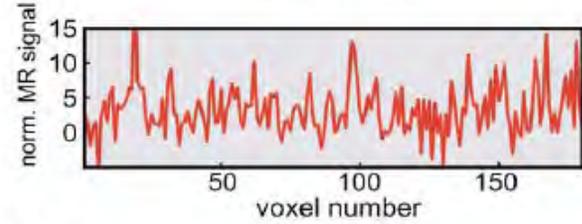
Fish



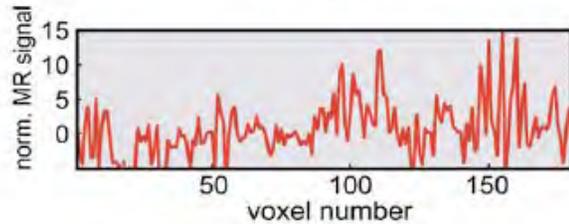
Gnomes



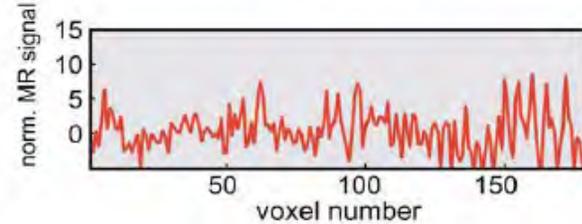
Horses

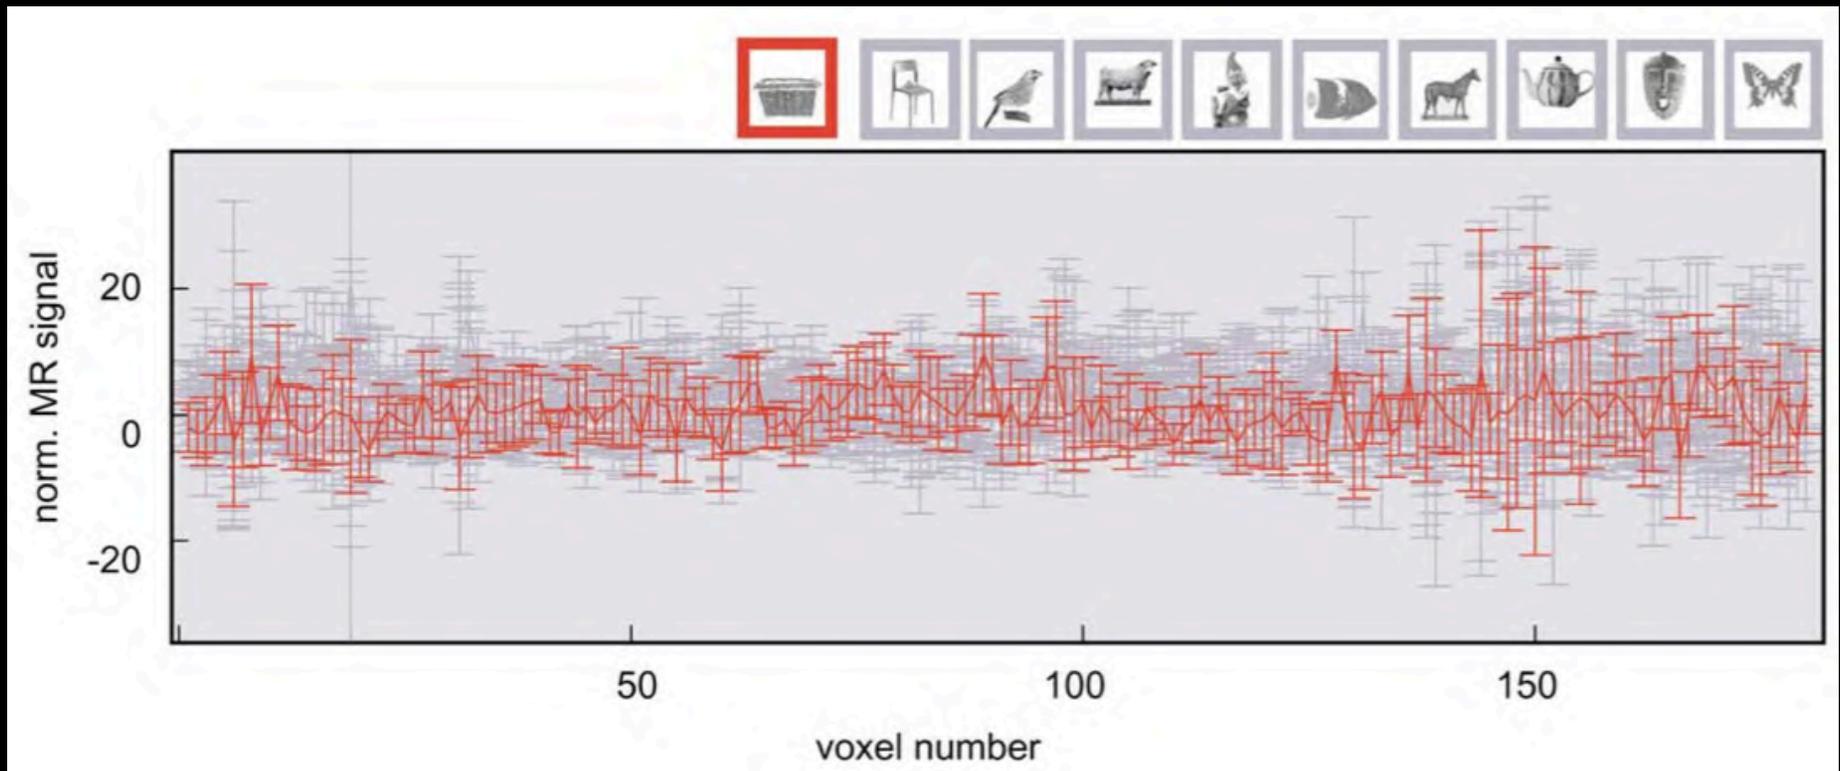


Masks



Teapots

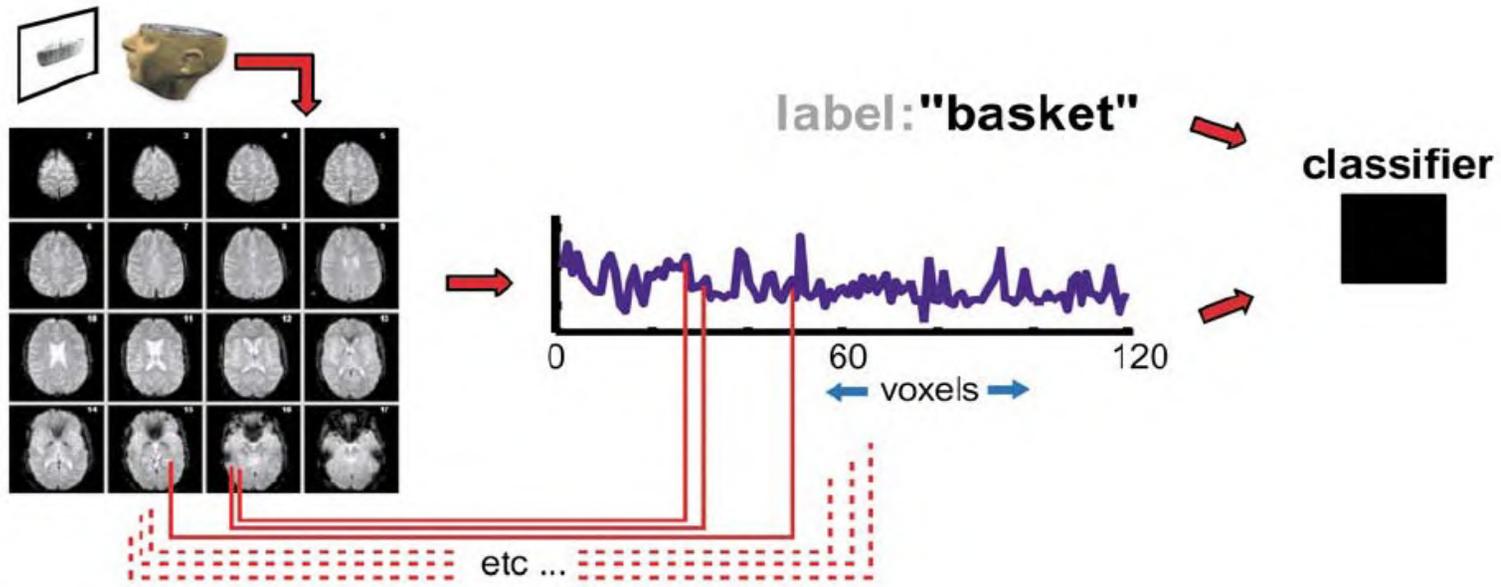




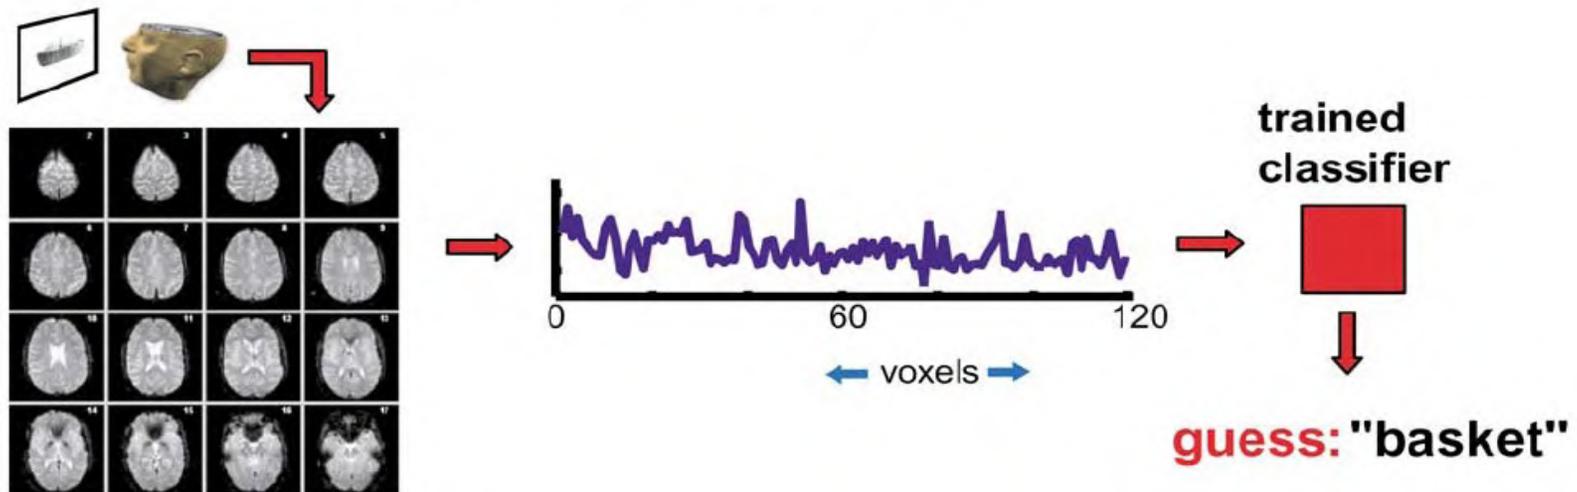
NOTE: no individual voxels show strong basket-specific response

Cox & Savoy (2003)

Training



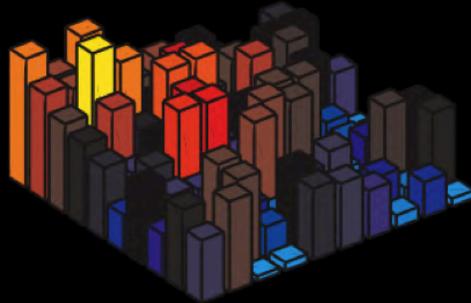
Classification (during a subsequent session)



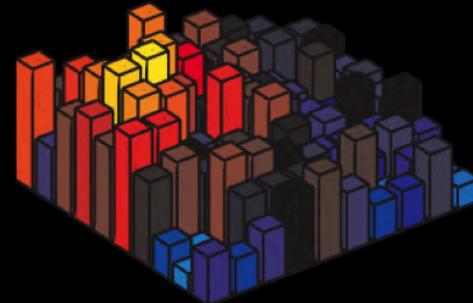
Core idea:

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

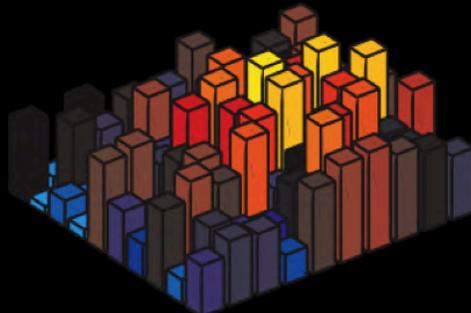
Condition 1



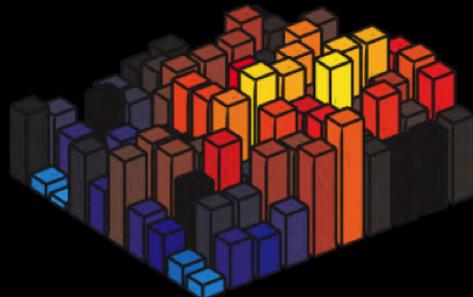
Condition 2



Condition 3



Condition n



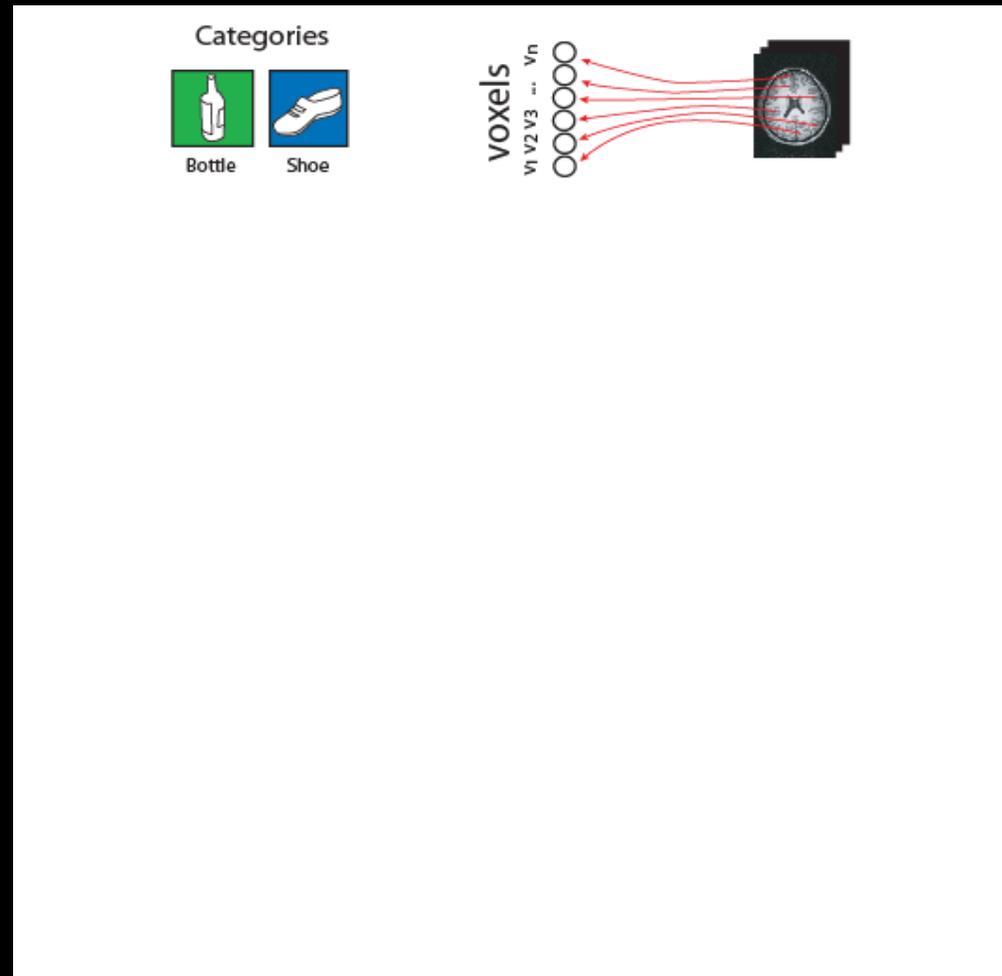
Core idea:

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



The Multi-Voxel Pattern Analysis Approach

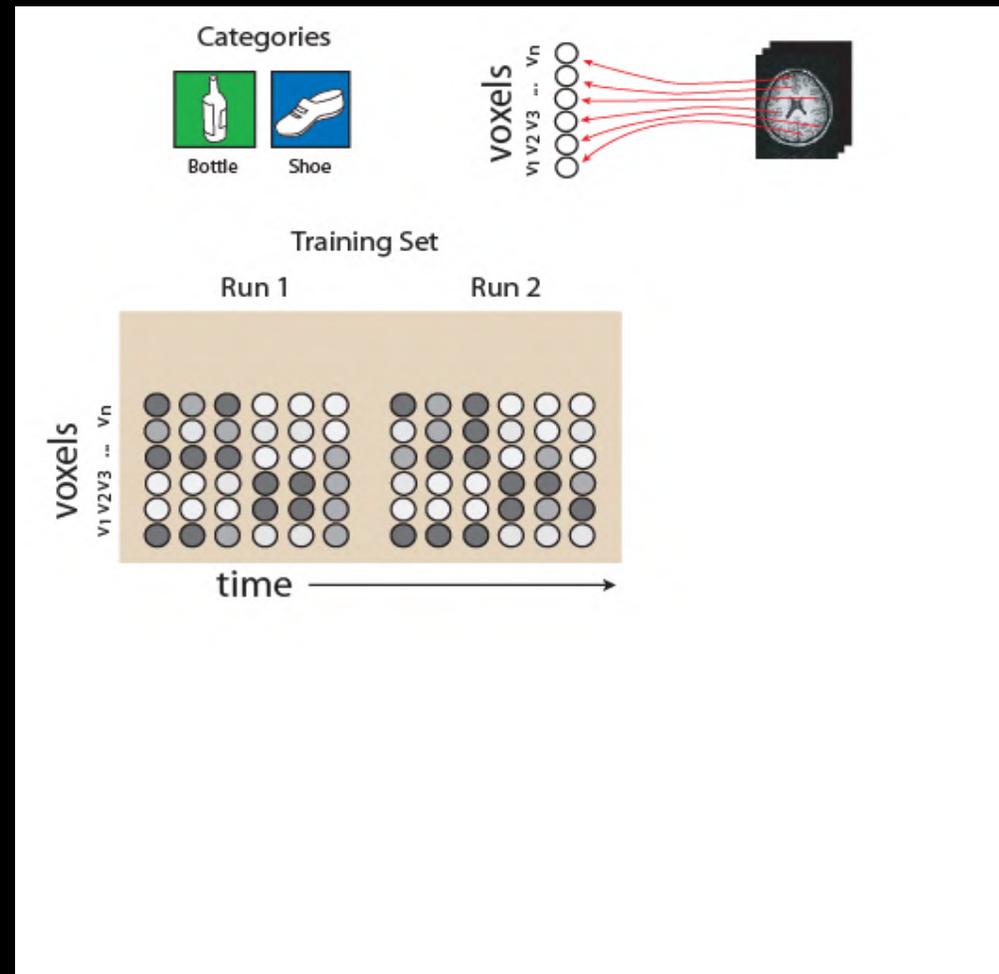
1. Acquire brain data while the subject is viewing shoes or bottles



The Multi-Voxel Pattern Analysis Approach

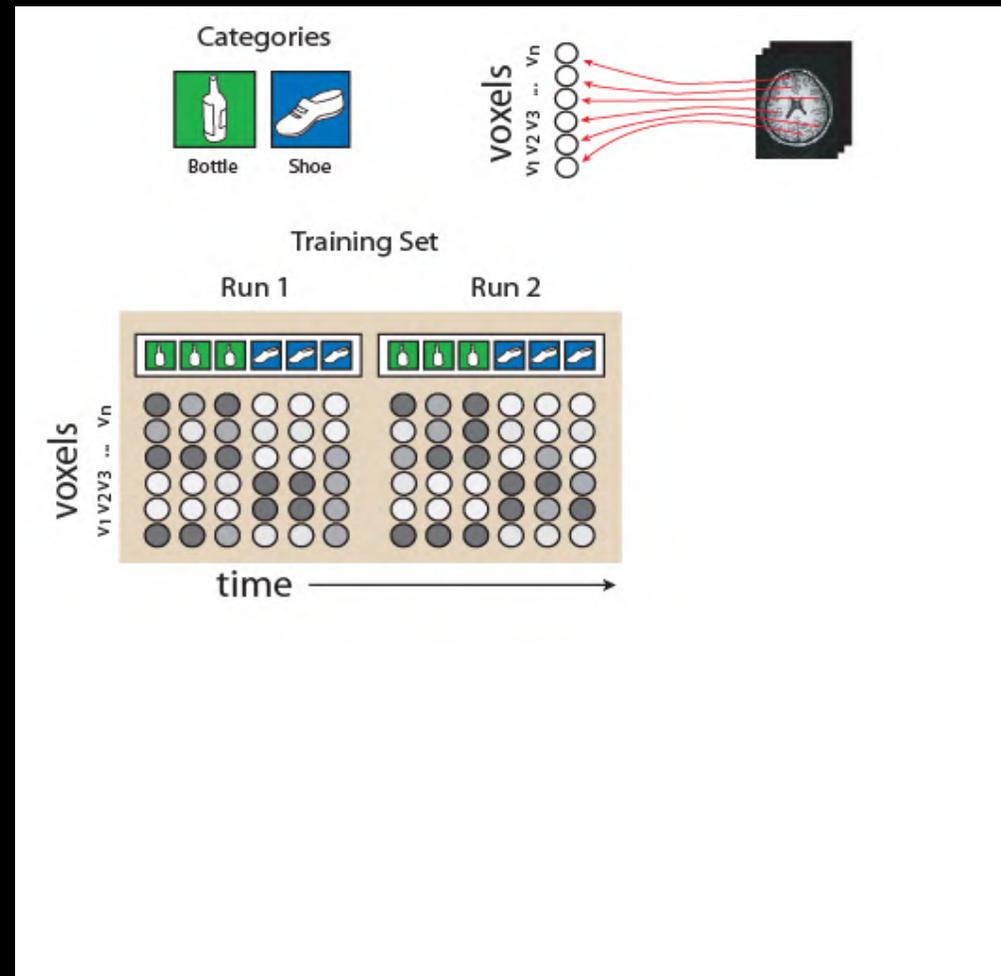
1. Acquire brain data
2. Convert each functional brain volume (or trial) into a vector that reflects the *pattern of activity across voxels* at that point in time

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



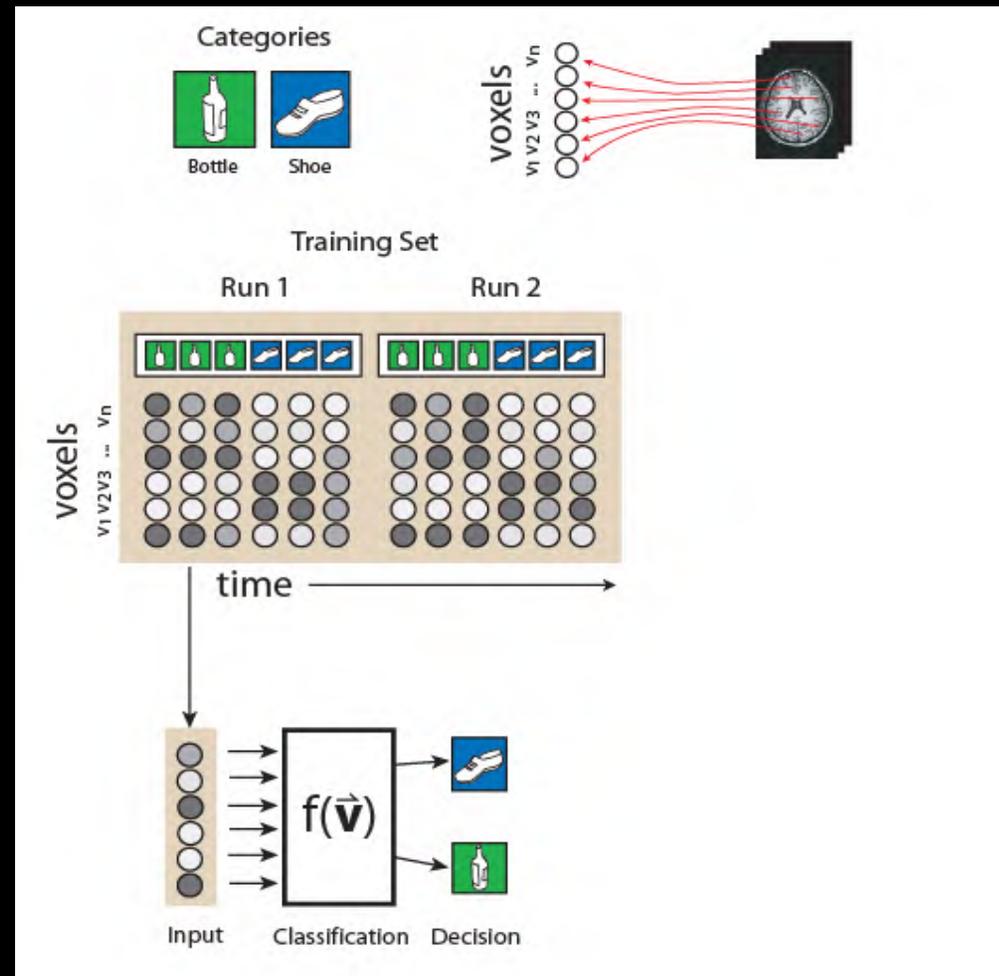
The Multi-Voxel Pattern Analysis Approach

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)



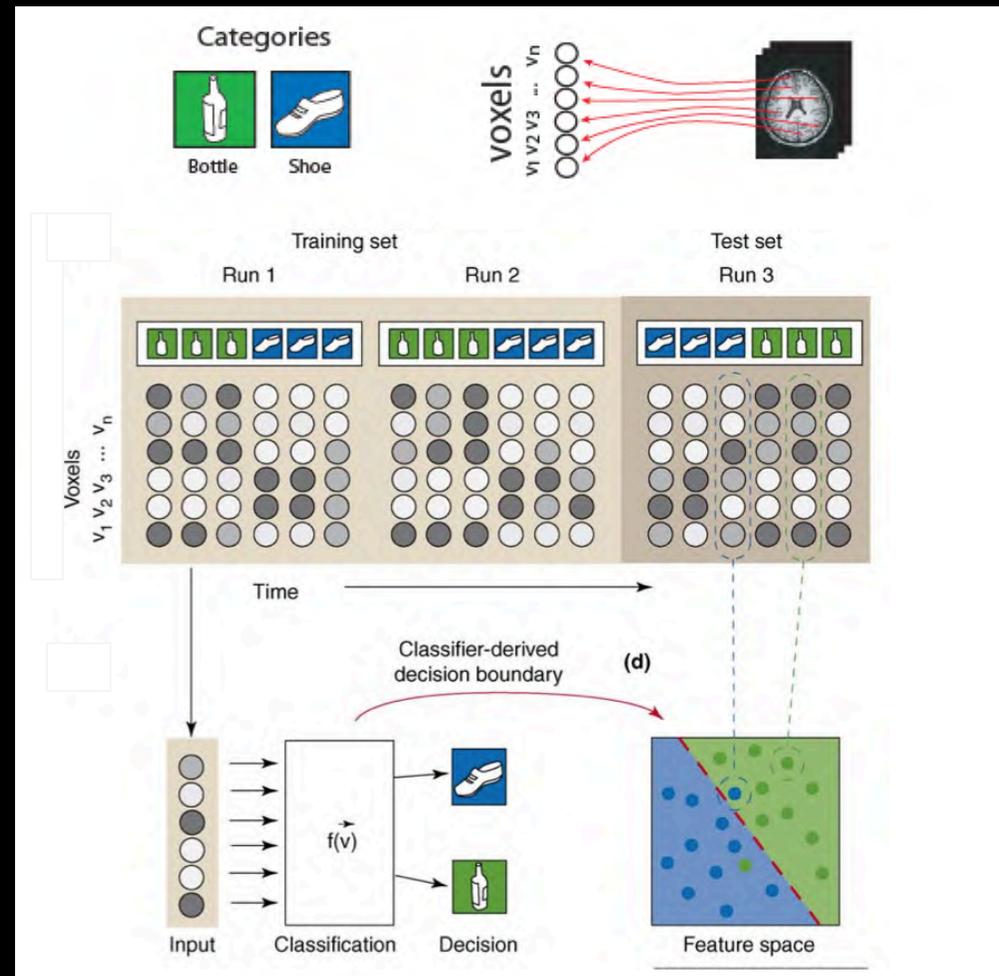
The Multi-Voxel Pattern Analysis Approach

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



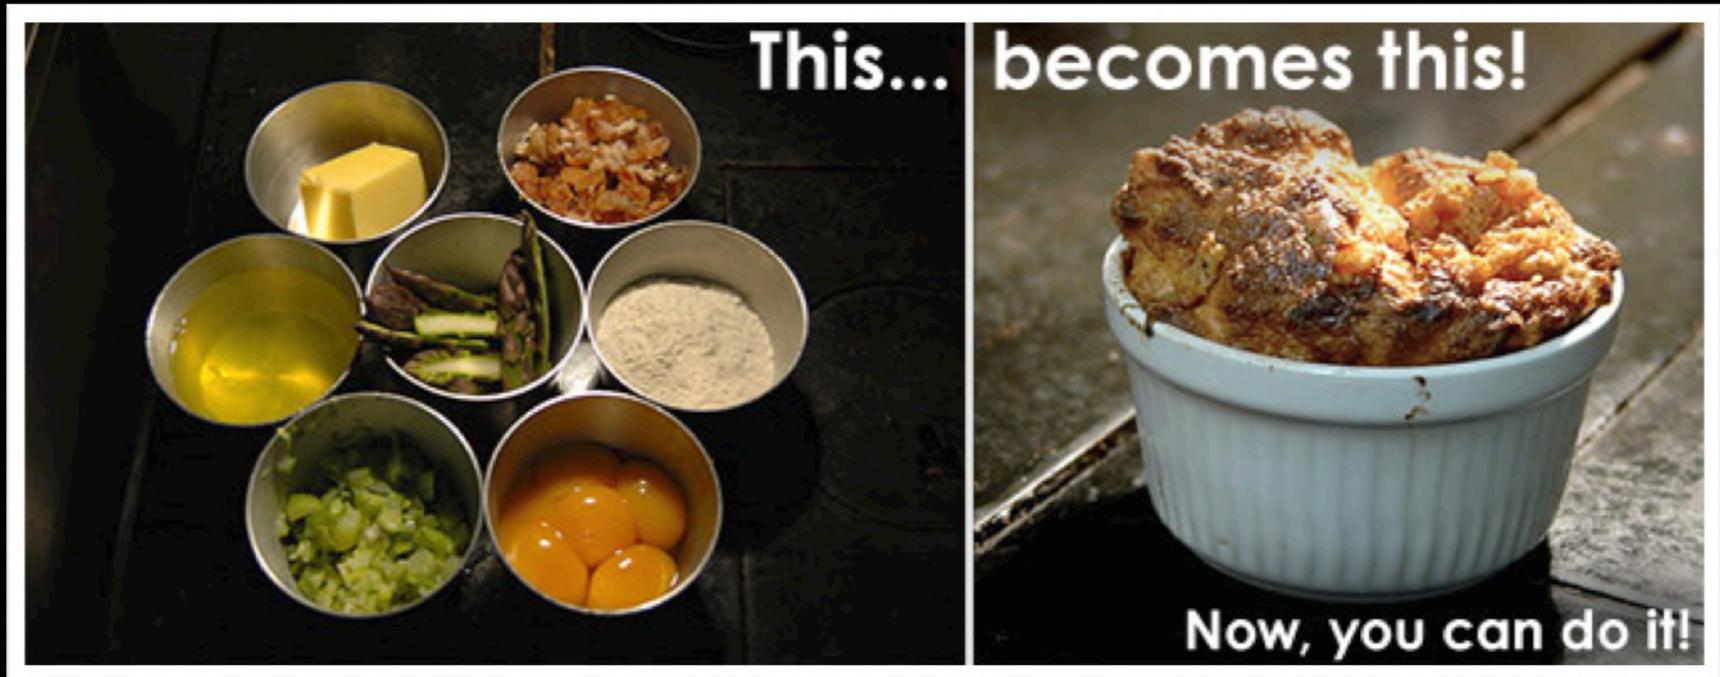
The Multi-Voxel Pattern Analysis Approach

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)



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](#)

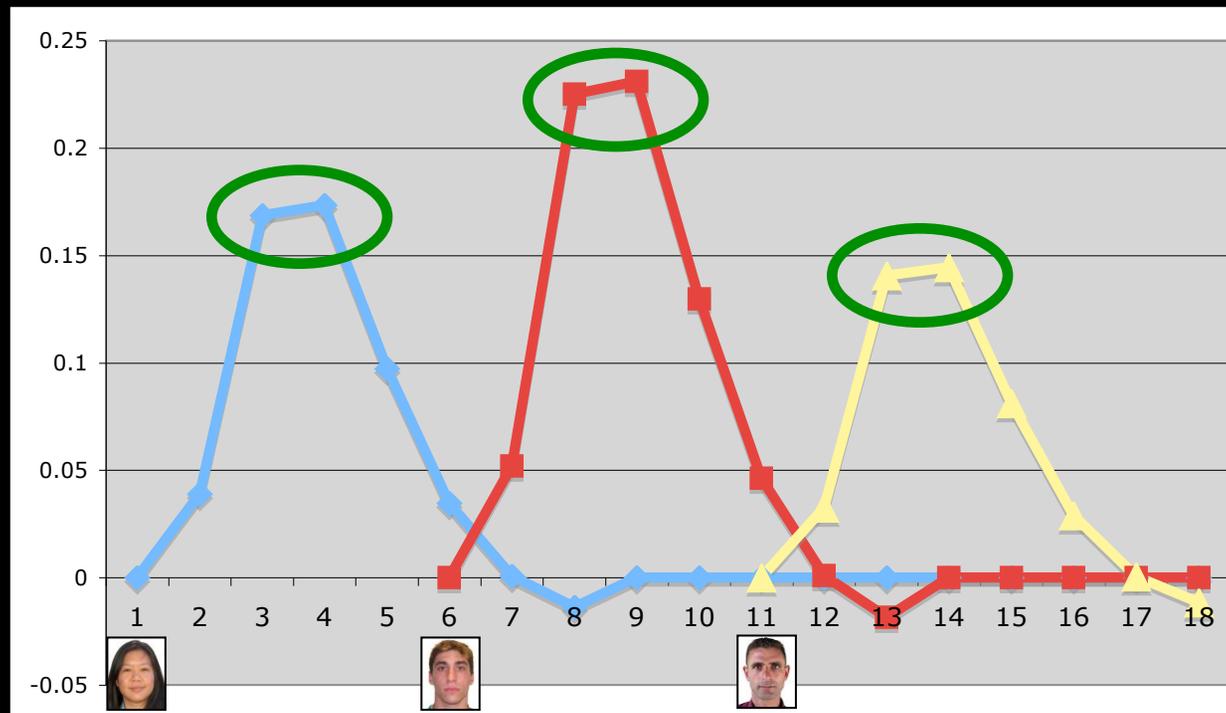
MVPA Data Processing Protocol

- **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*

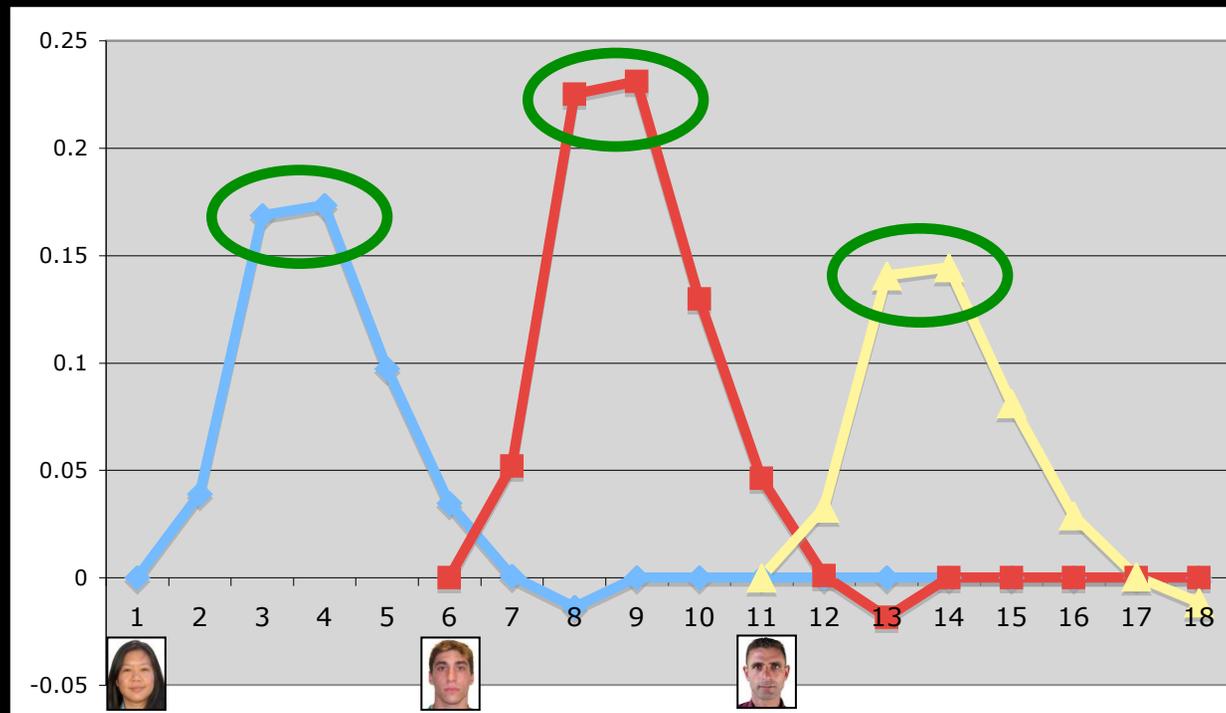
- Reduce full fMRI timeseries from having 5 values (TRs) per trial to having only 1 value per trial



- 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

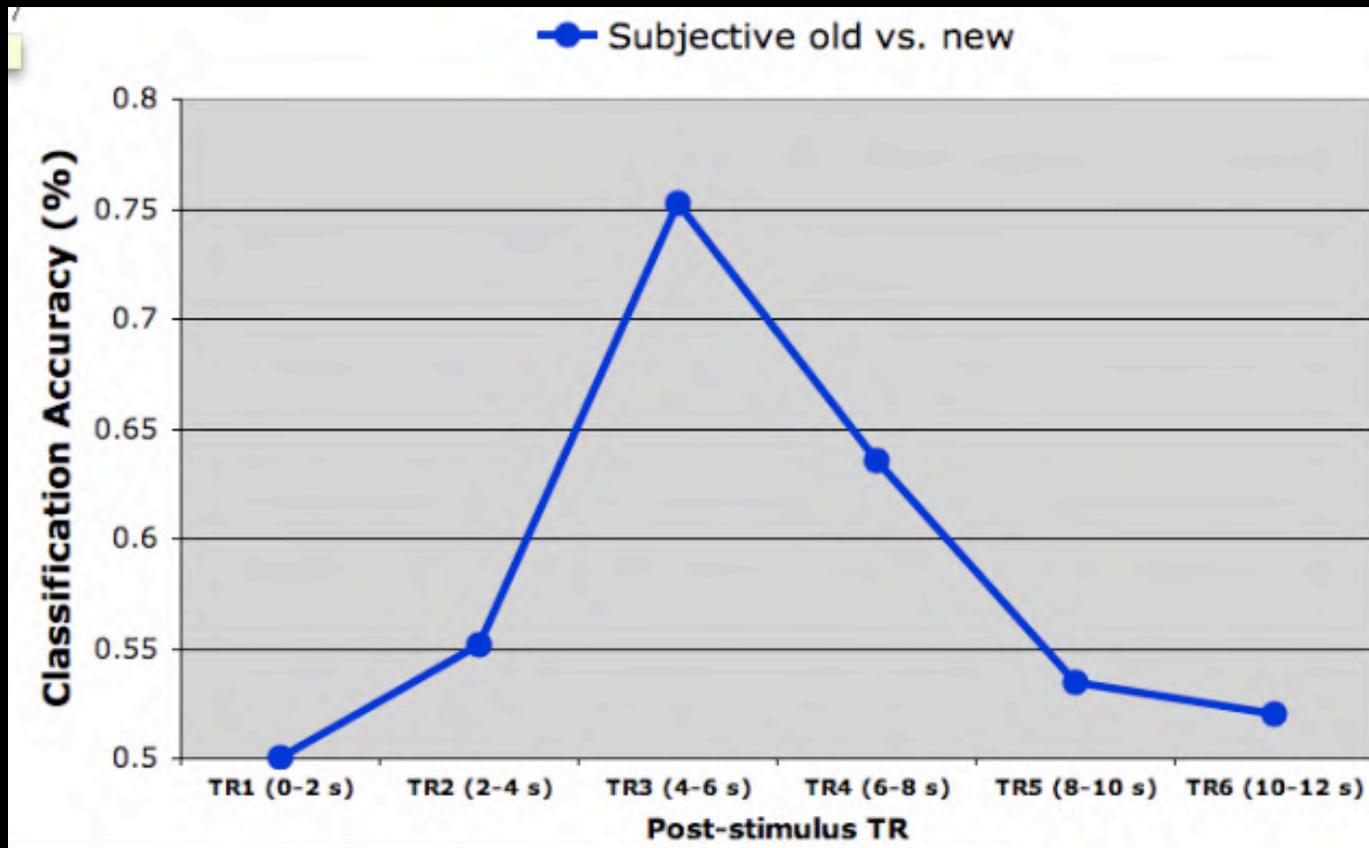


- Or run a new GLM that estimates a single parameter for each trial (i.e, beta-series approach; Rissman et al. 2004)

Pattern Analysis Approach: *Temporal Selection*

Alternative strategy

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

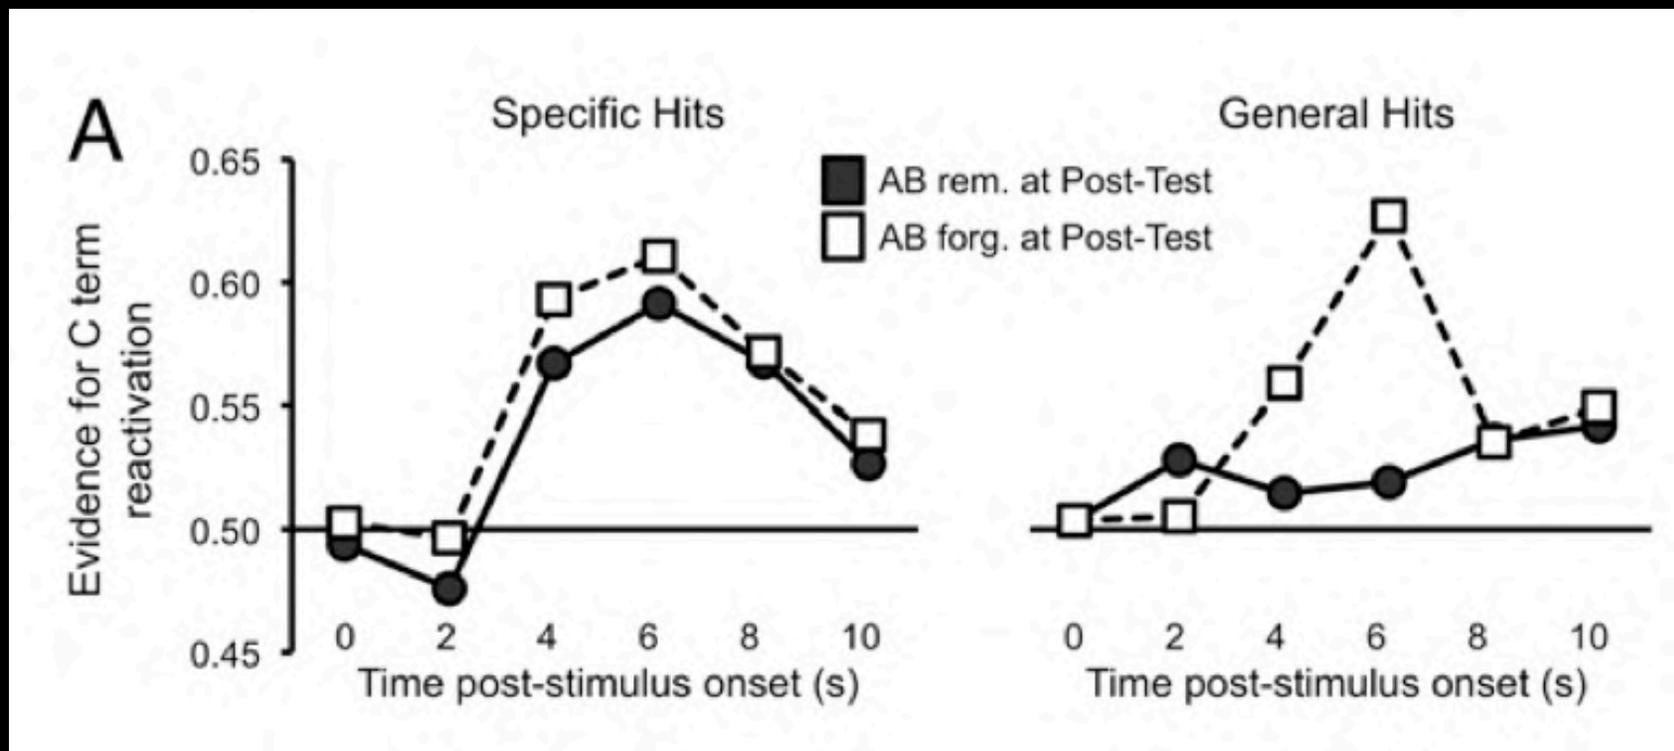


Classification timecourse for one subject

Pattern Analysis Approach: *Temporal Selection*

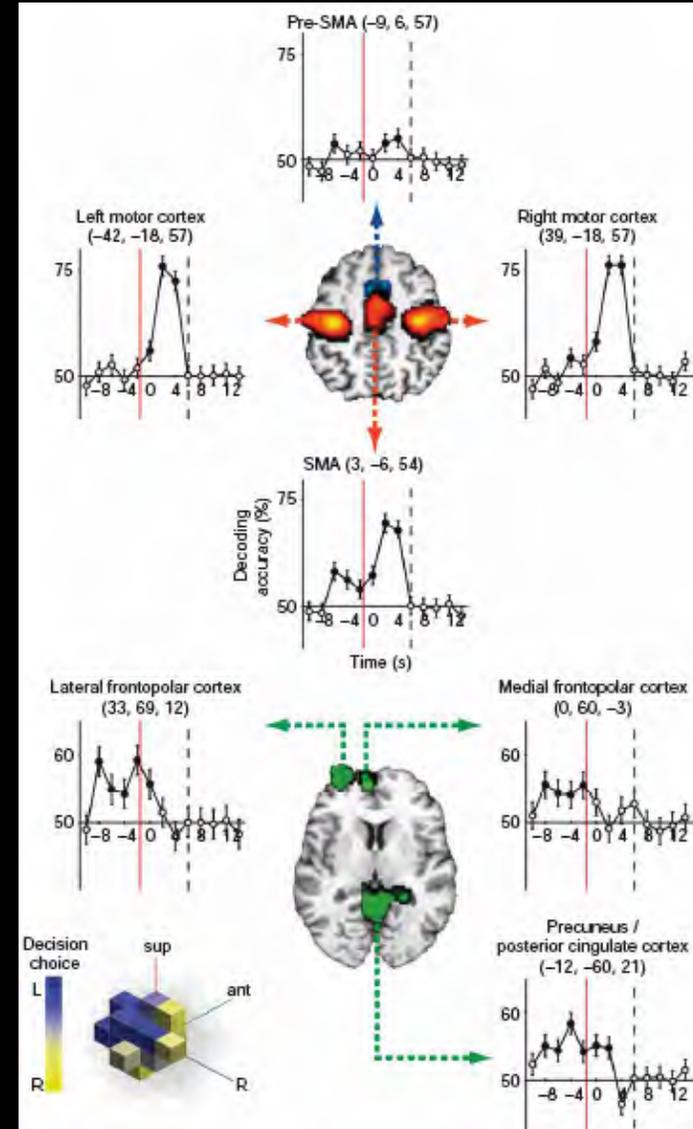
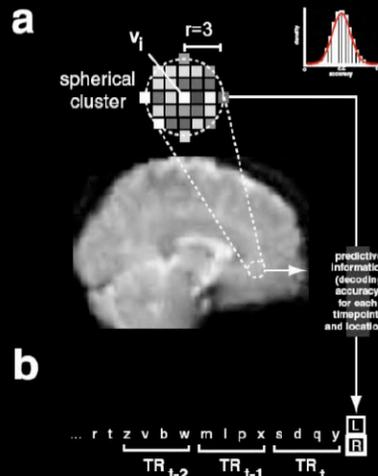
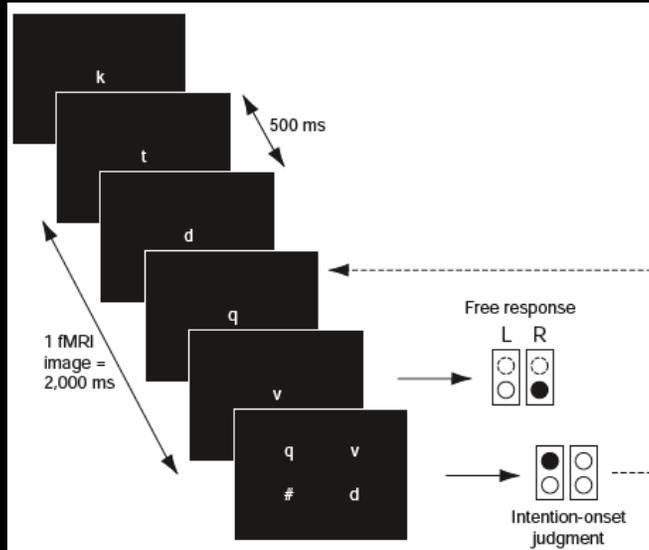
Alternative strategy

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



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

Unraveling the Notion of Free Will?



Soon et al. (2008) *Nature Neuroscience*

Pattern Analysis Approach:

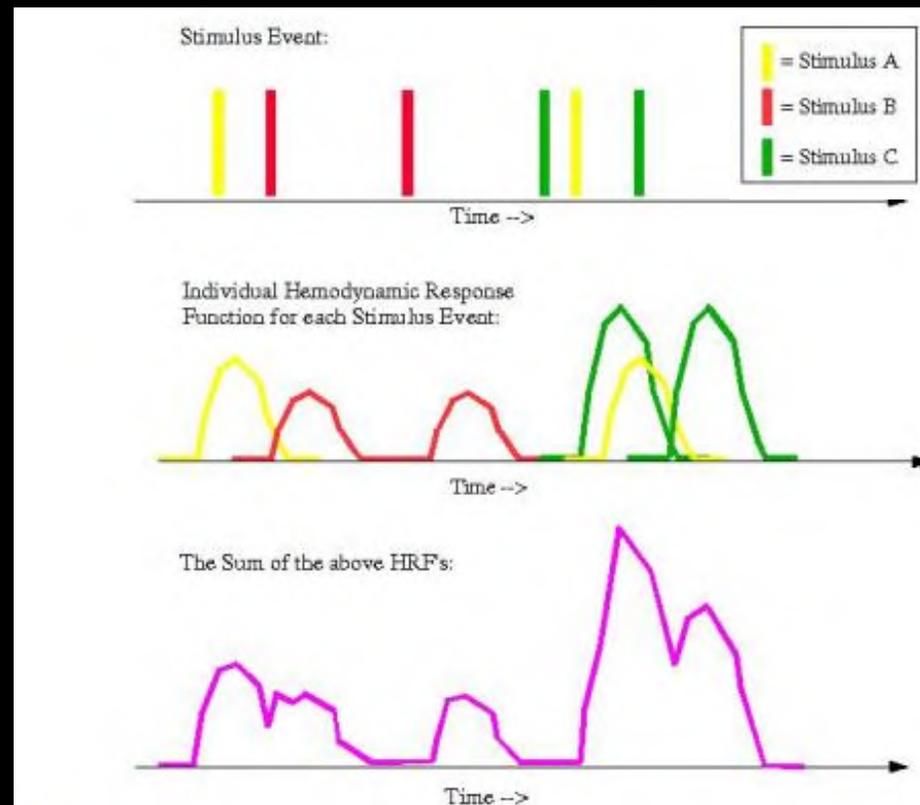
Temporal Selection

- 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:

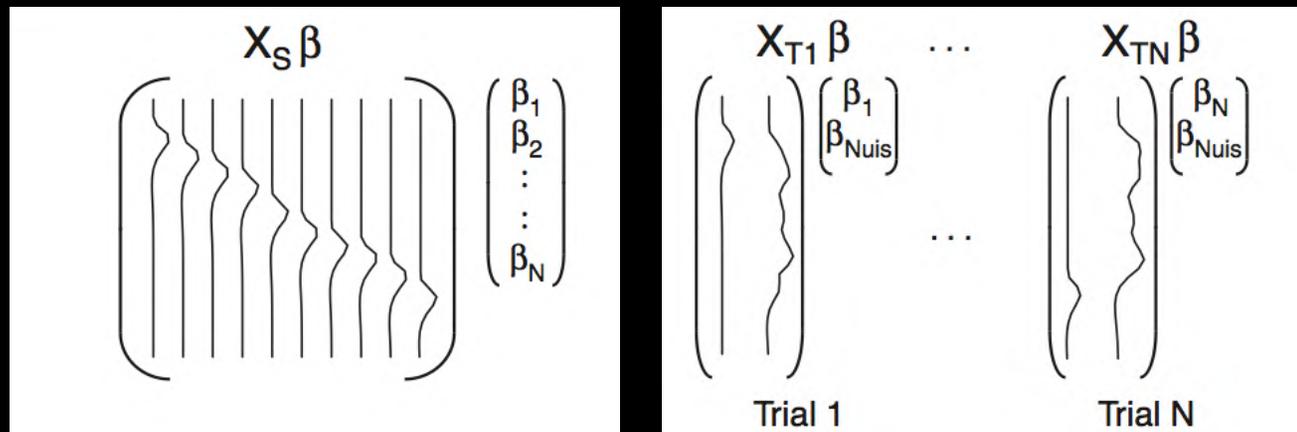
- For each TR, examine your design matrix and determine which condition has the maximal predicted activity
- Specify a threshold to exclude to ambiguous TRs



Pattern Analysis Approach: *Temporal Selection*

Another approach (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 Rissman et al. (2004) 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.

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!



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

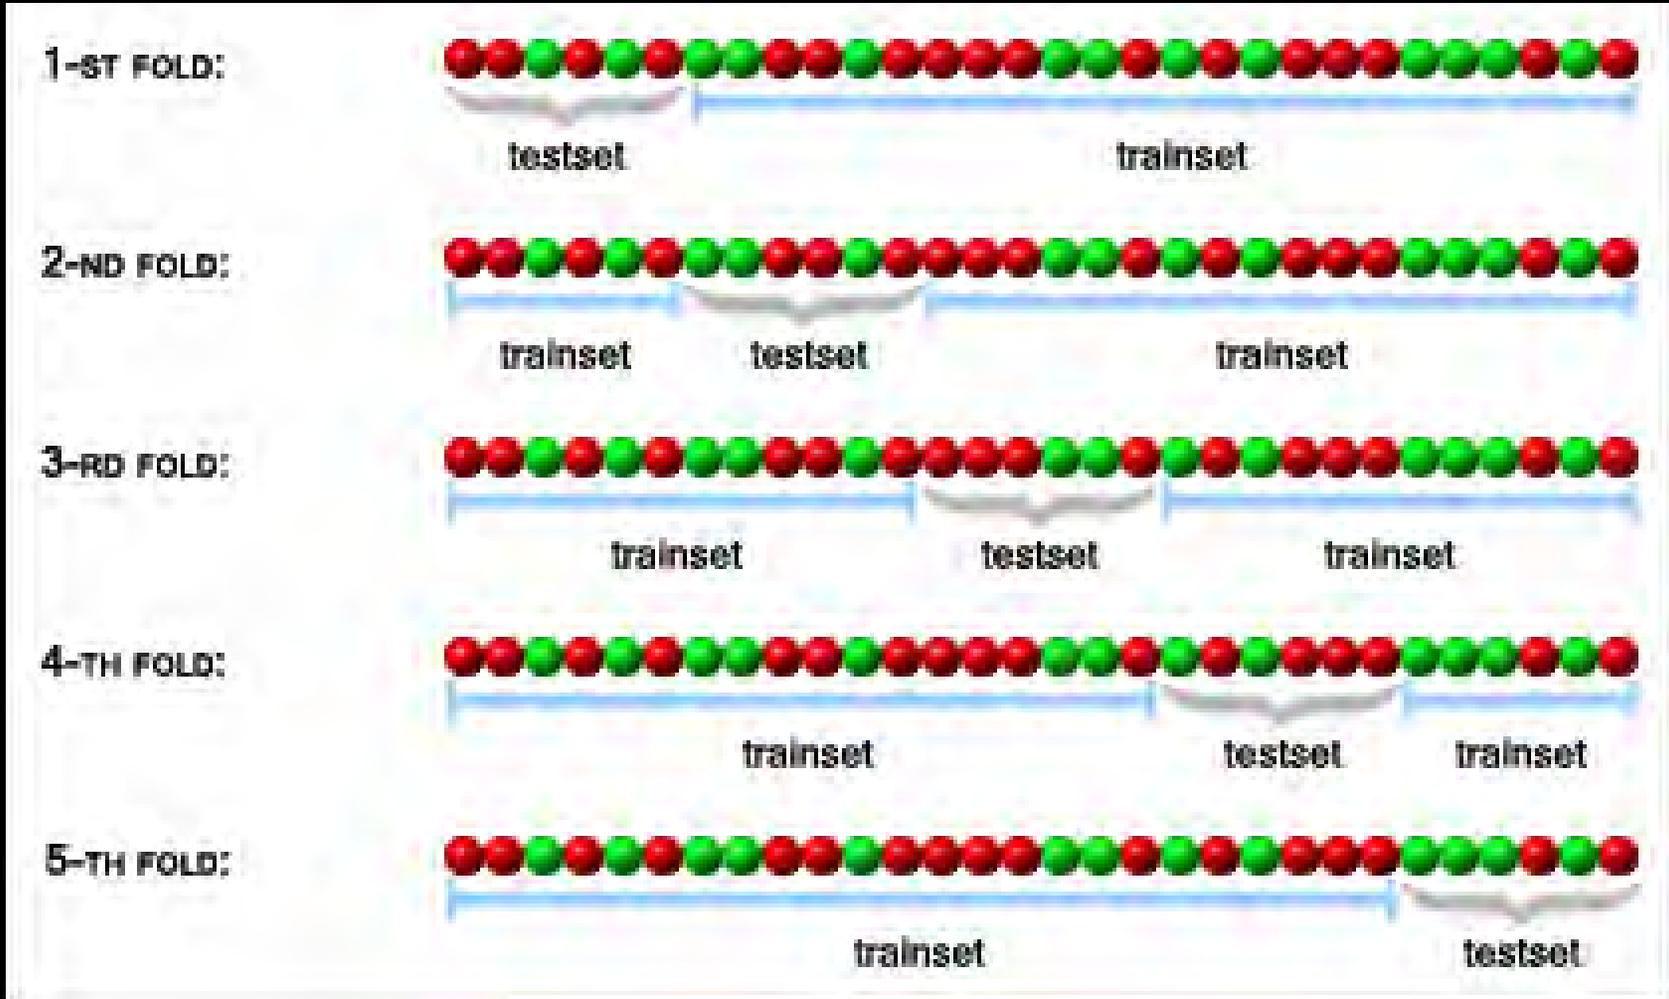
Design

Classification

ABSTRACT

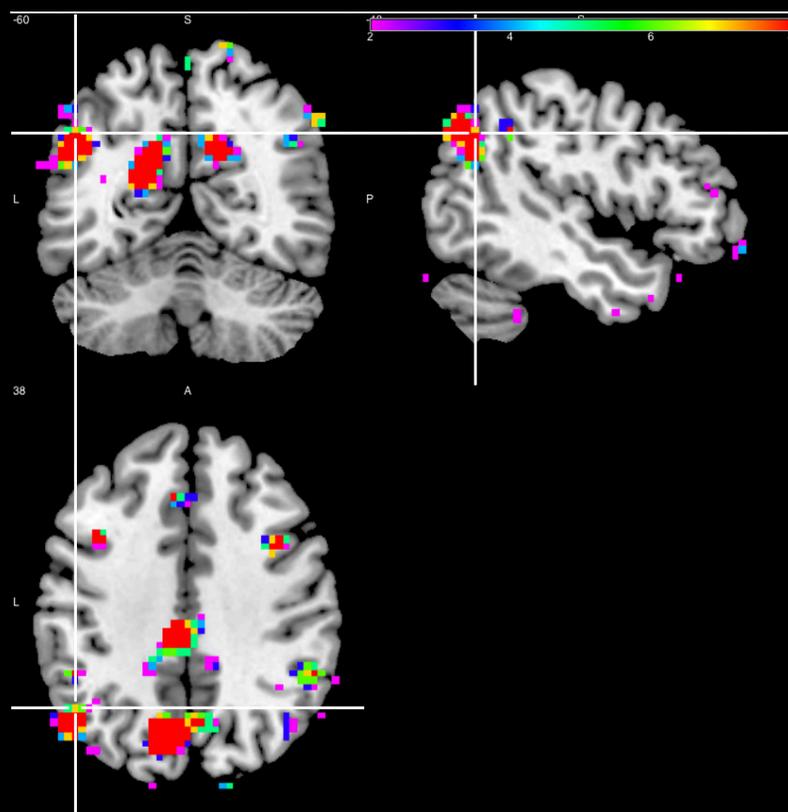
Functional magnetic resonance imaging (fMRI) studies are broken up into runs (or 'sessions'), frequently selected to be long to minimize across-run signal variations. For investigations that use multi-voxel pattern analysis (MVPA), however, employing many short runs might improve a classifier's ability 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

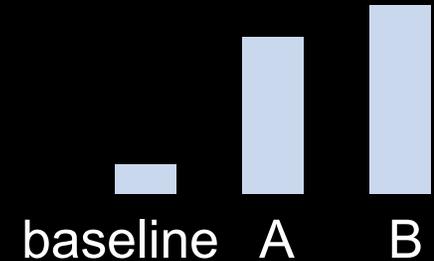
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing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



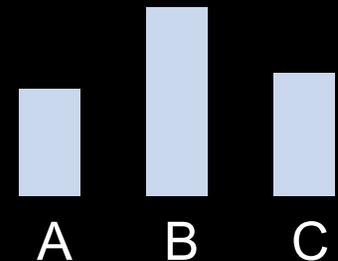
Feature Selection

- What criteria should define important voxels?

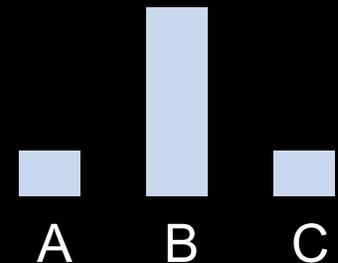
- difference from baseline



- difference between classes (e.g. ANOVA)

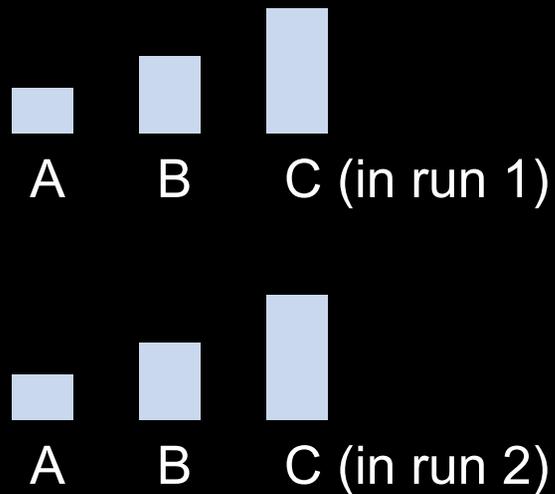


- preferential response to one class



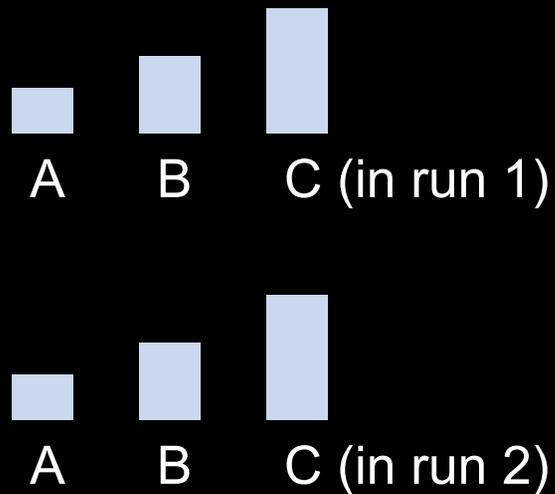
Feature Selection

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



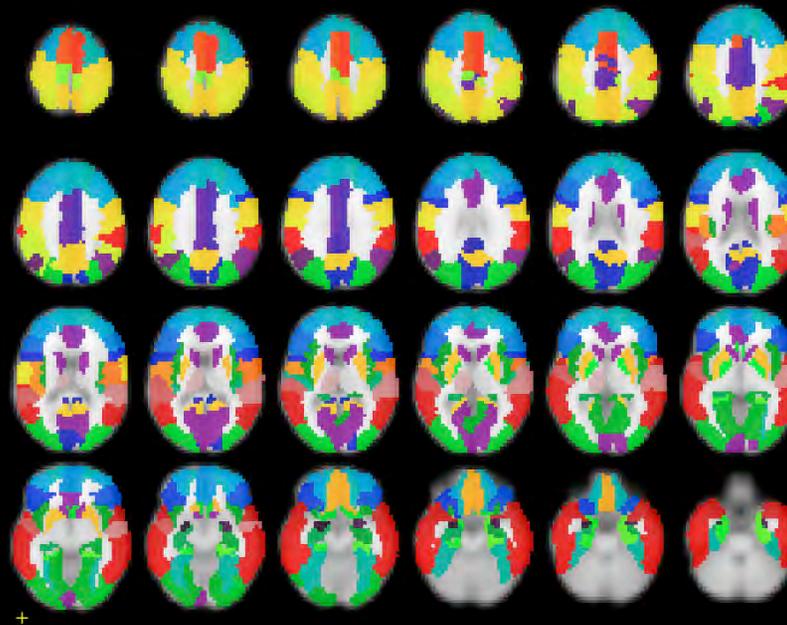
Feature Selection

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



Feature Selection

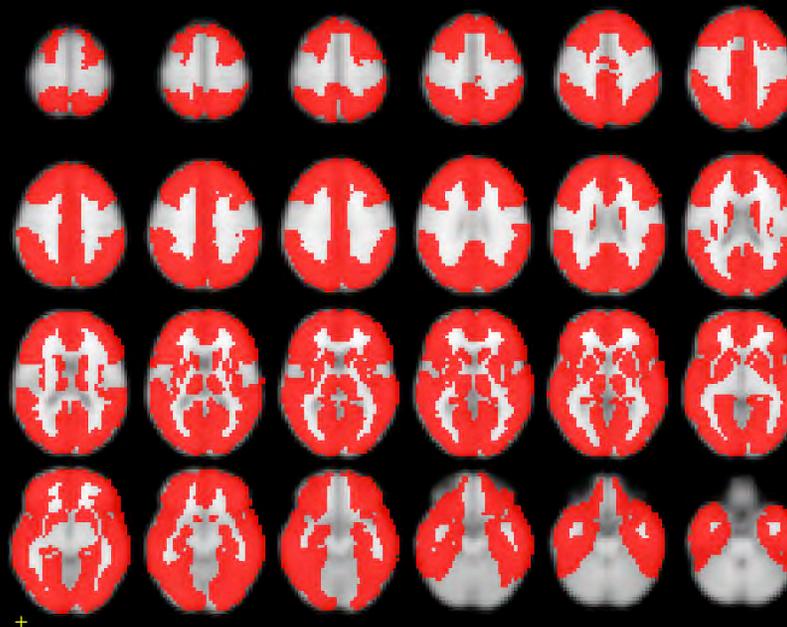
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing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



ROIs from the
AAL atlas

Feature Selection

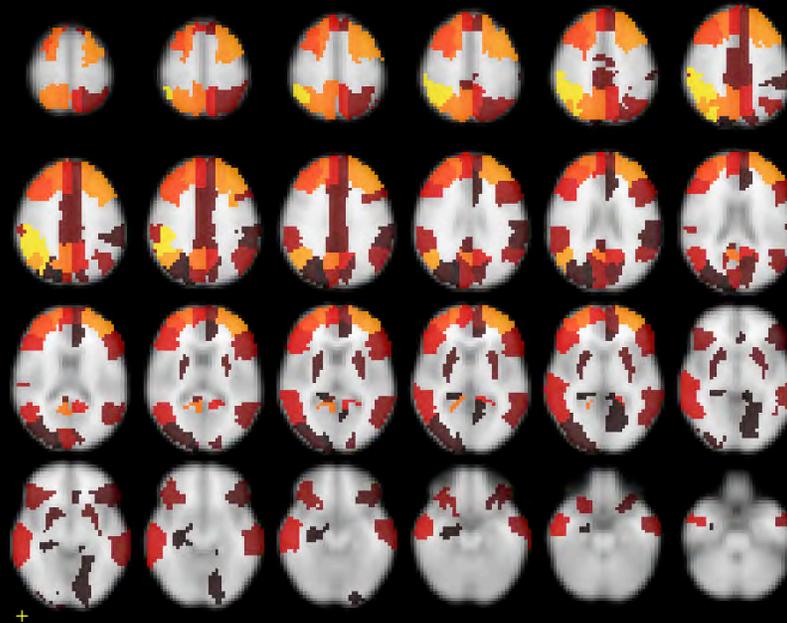
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing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

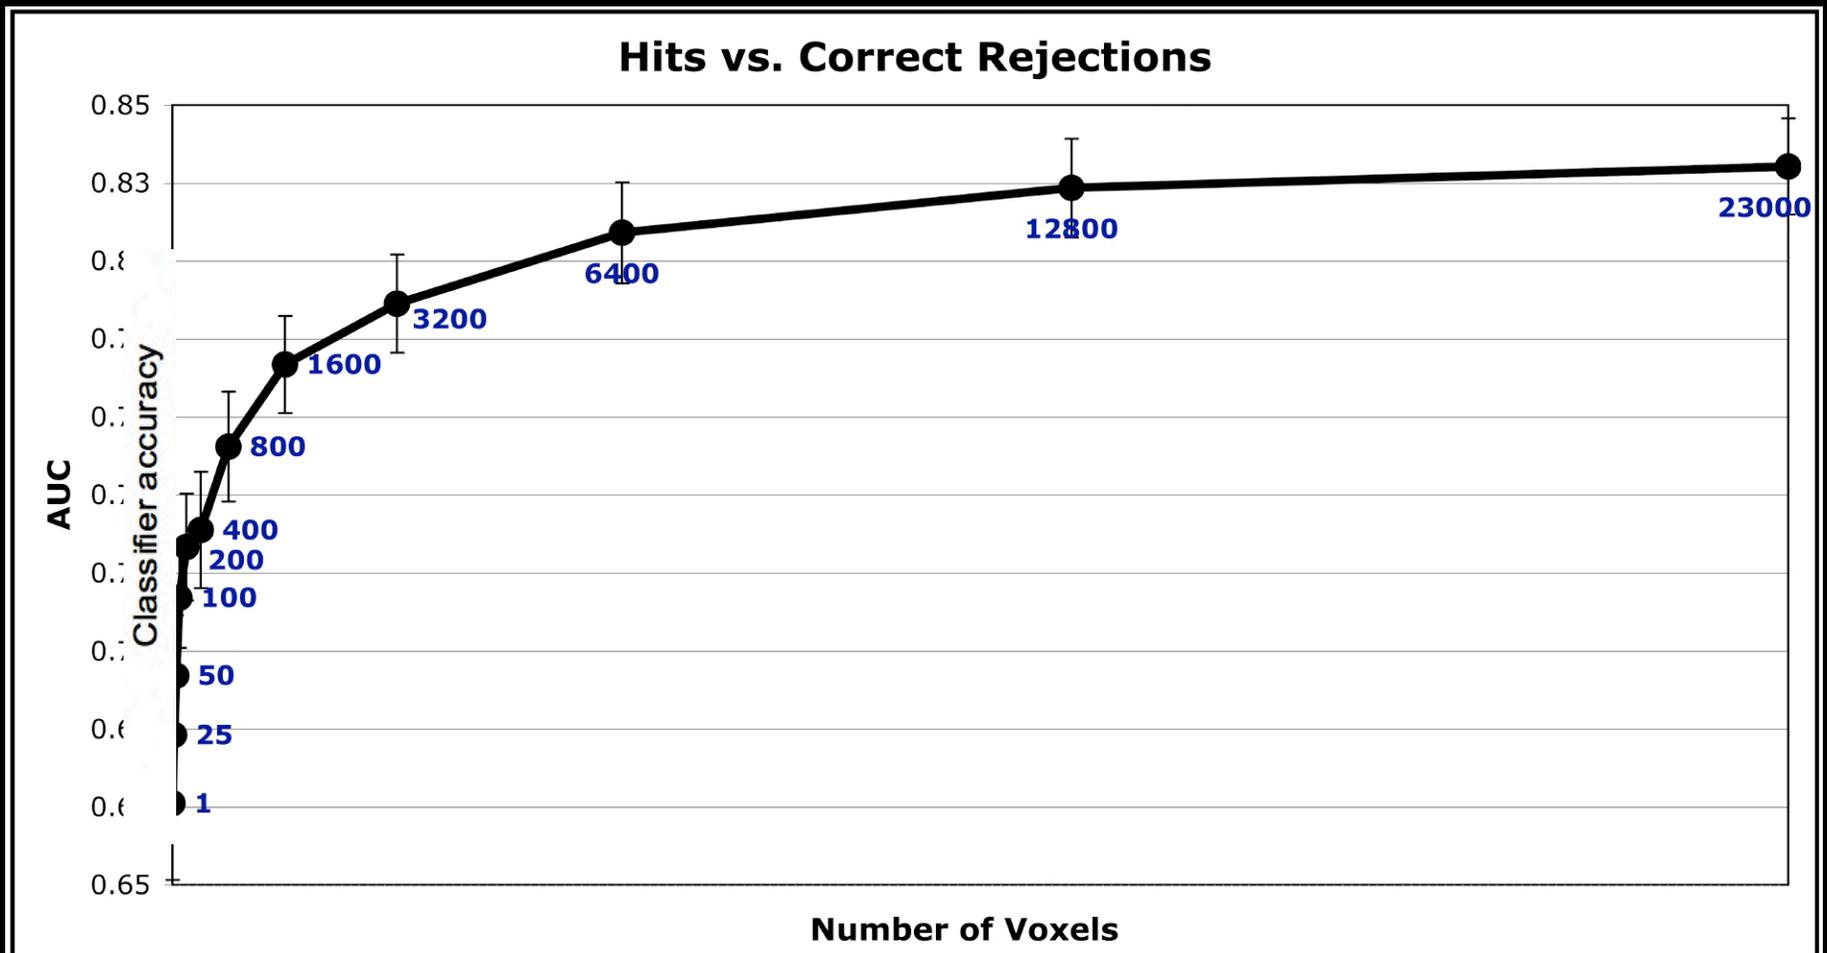
- Selecting which voxels to include in the analysis
 - **Univariate GLM (a.k.a. conventional brain mapping)**
 - Identify general task-responsive voxels (e.g., all conditions vs. baseline)
 - Identify task-selective voxels (e.g., Condition A vs. Condition B)
 - must be done without using “held-out” testing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



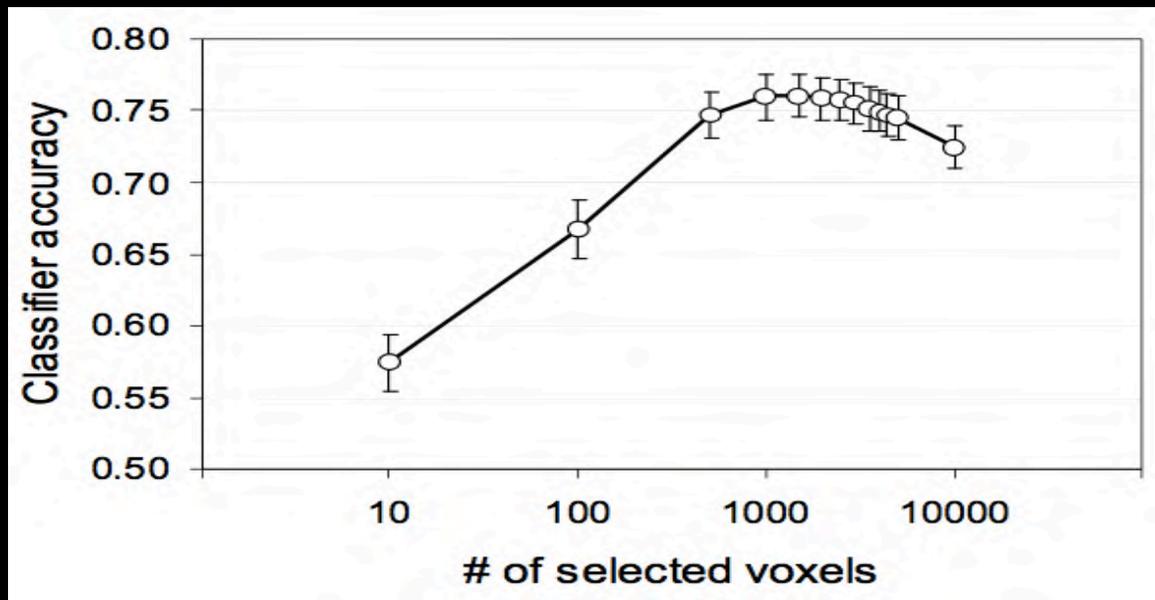
You also can
compute
classification
performance
within each ROI

How many features (voxels) to use?

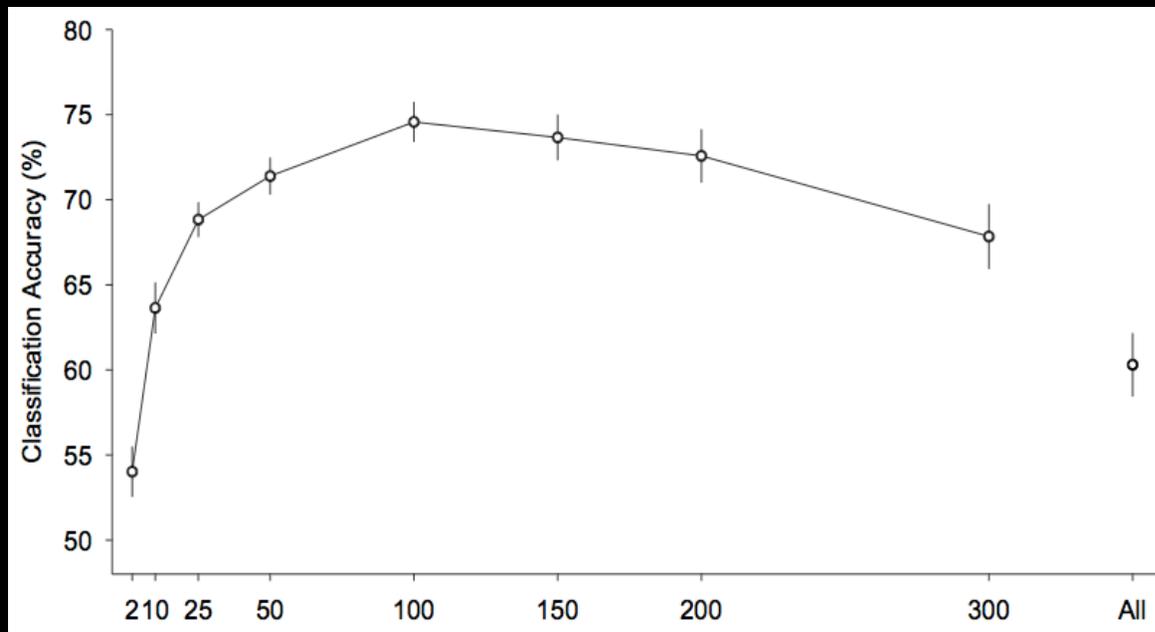
- Classification performance as a function of the number of voxels used by classifier (*ANOVA-based selection*)



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

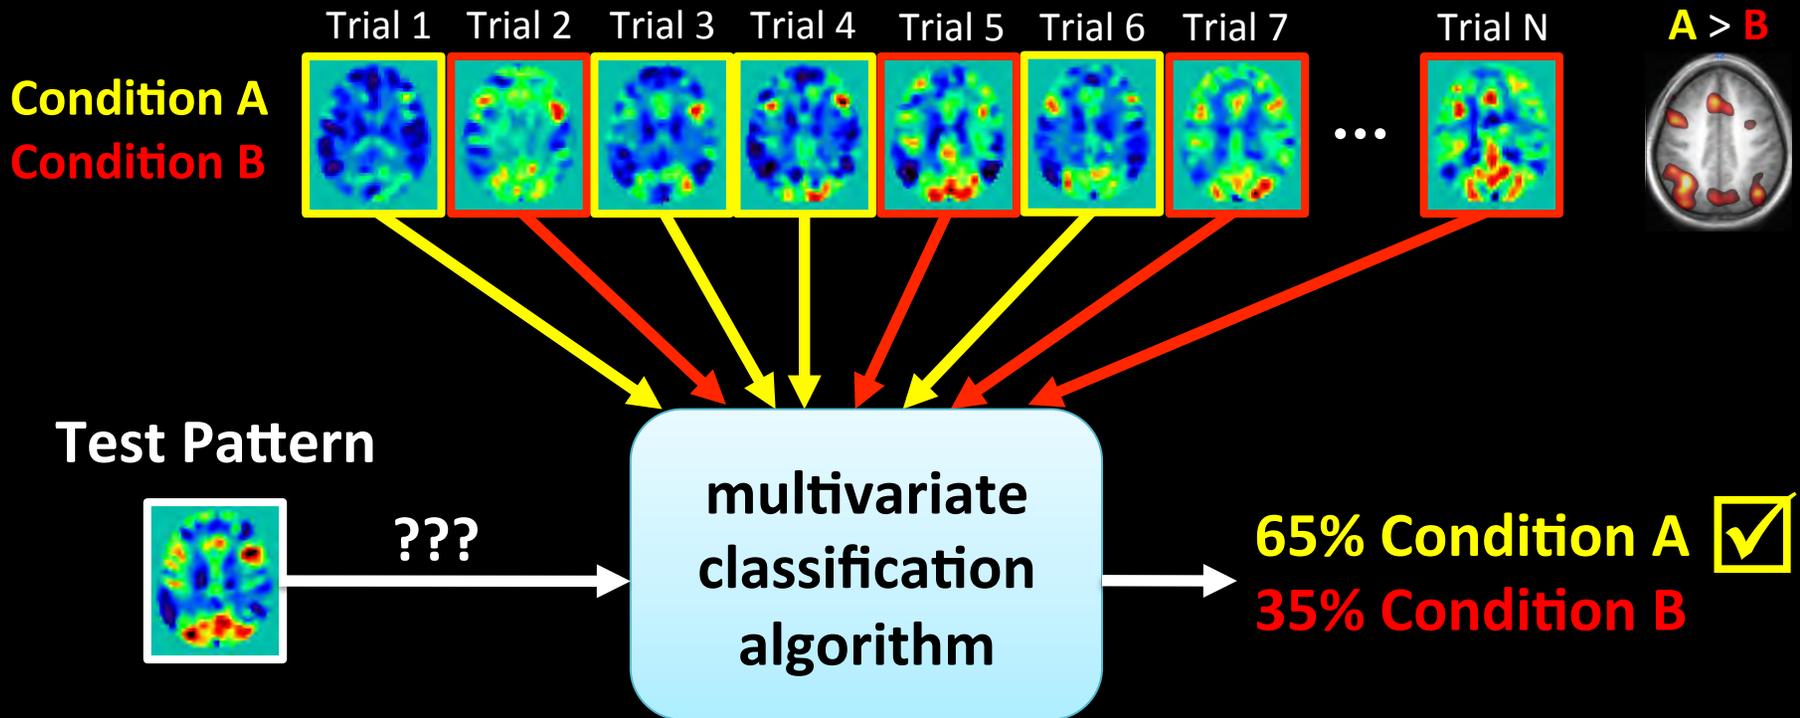


Peak performance with ~100 voxels
Esterman et al. (2009)



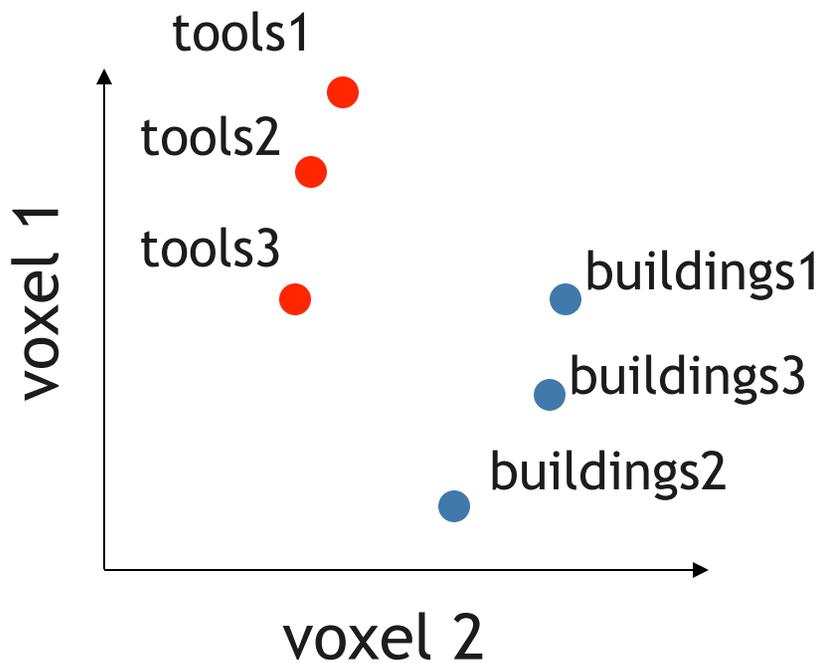
Training and testing the classifier

Training Patterns



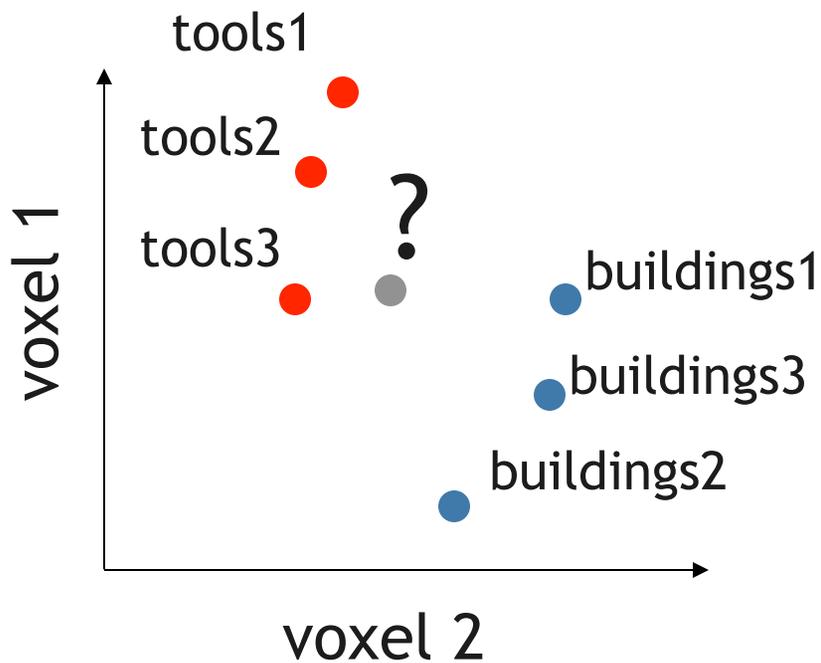
what is inside the box?

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



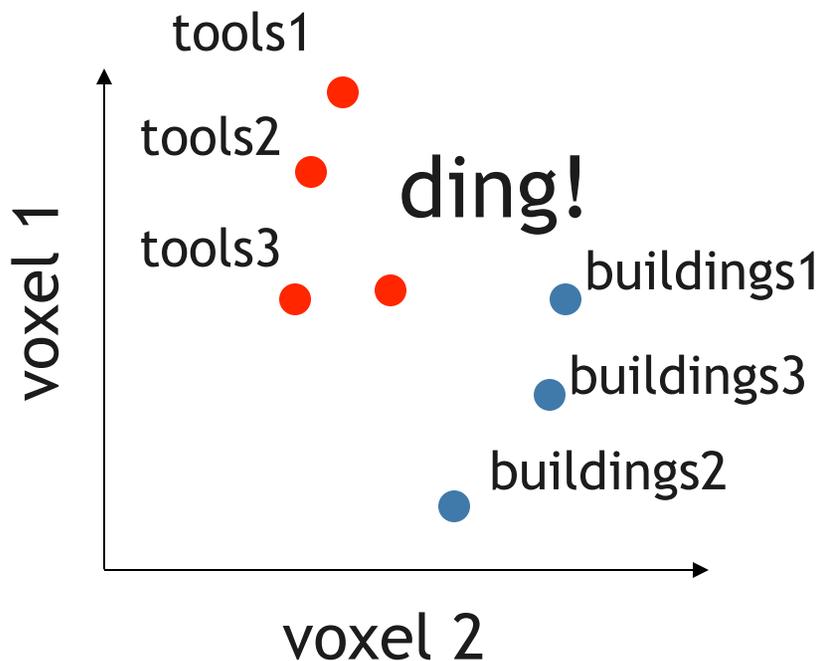
what is inside the box?

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



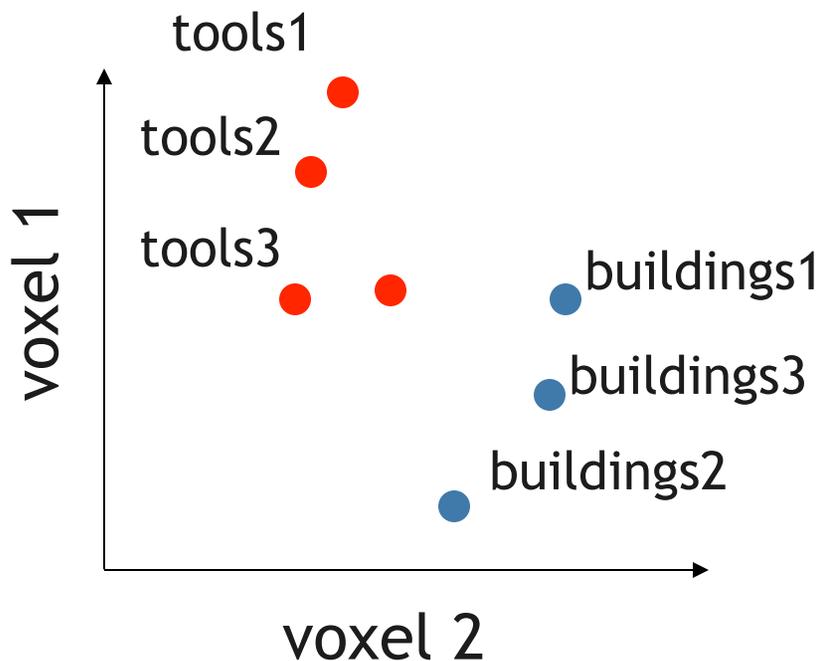
what is inside the box?

- 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

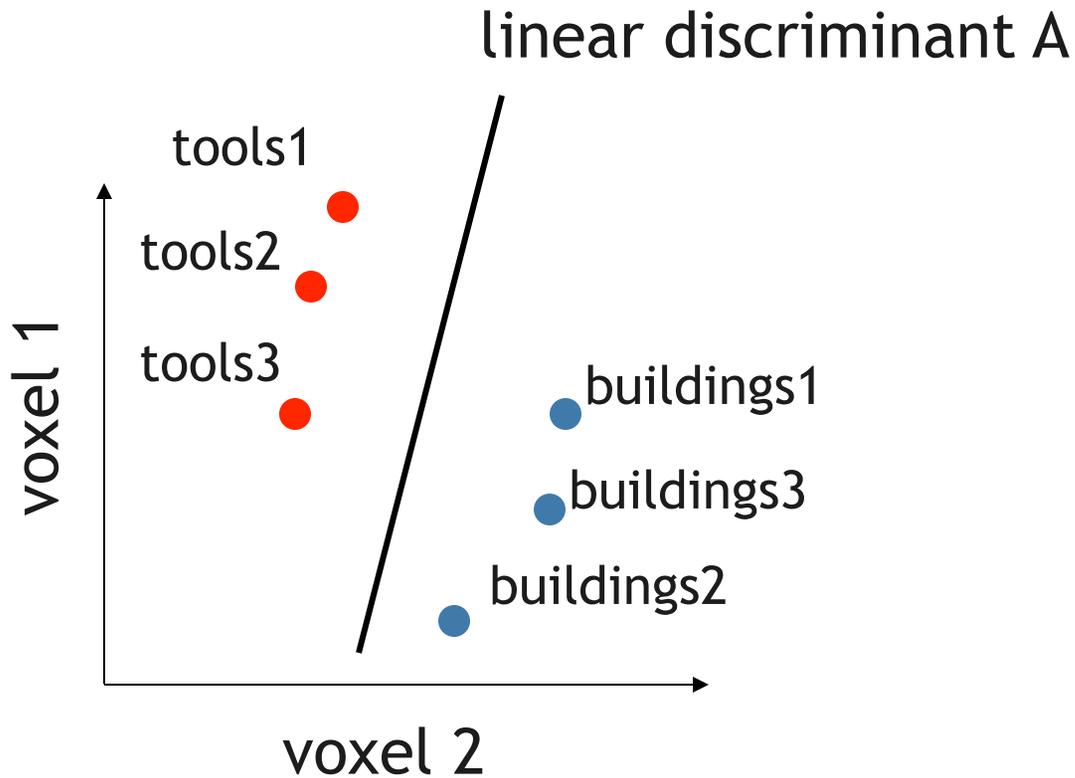


Euclidean dist., correlation, ...



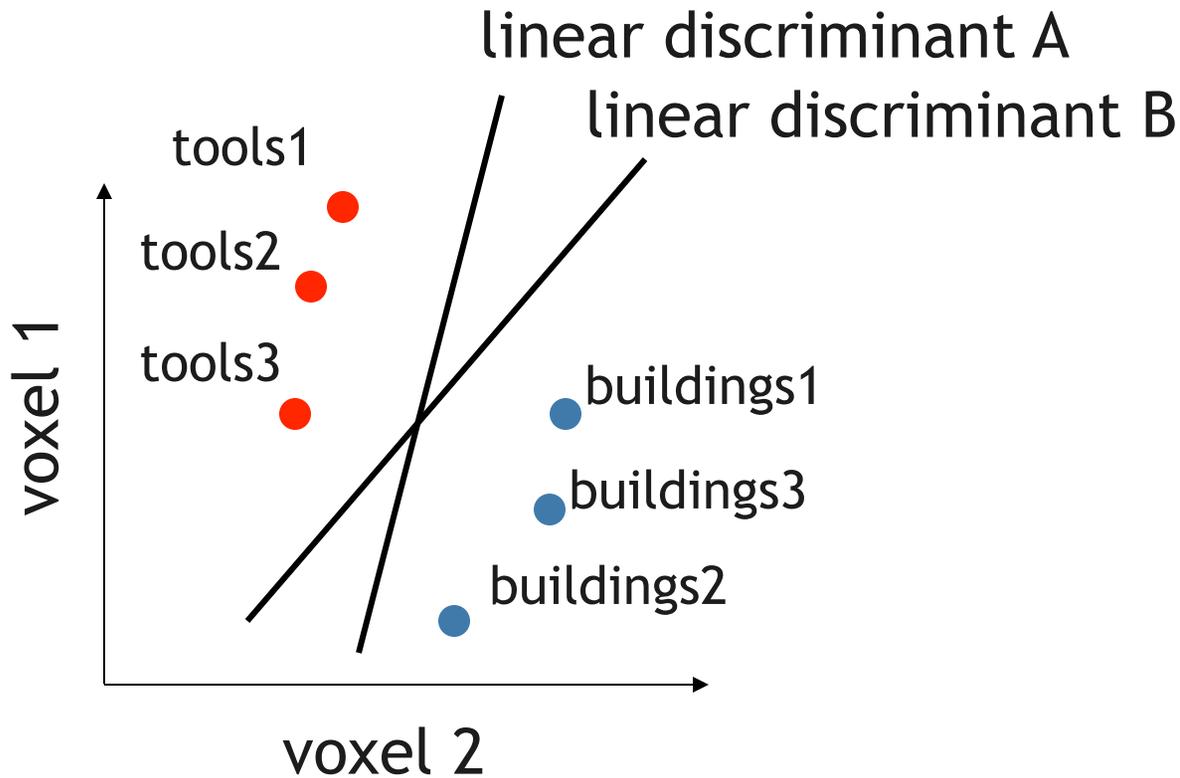
what is inside the box?

- next simplest: learn linear discriminant



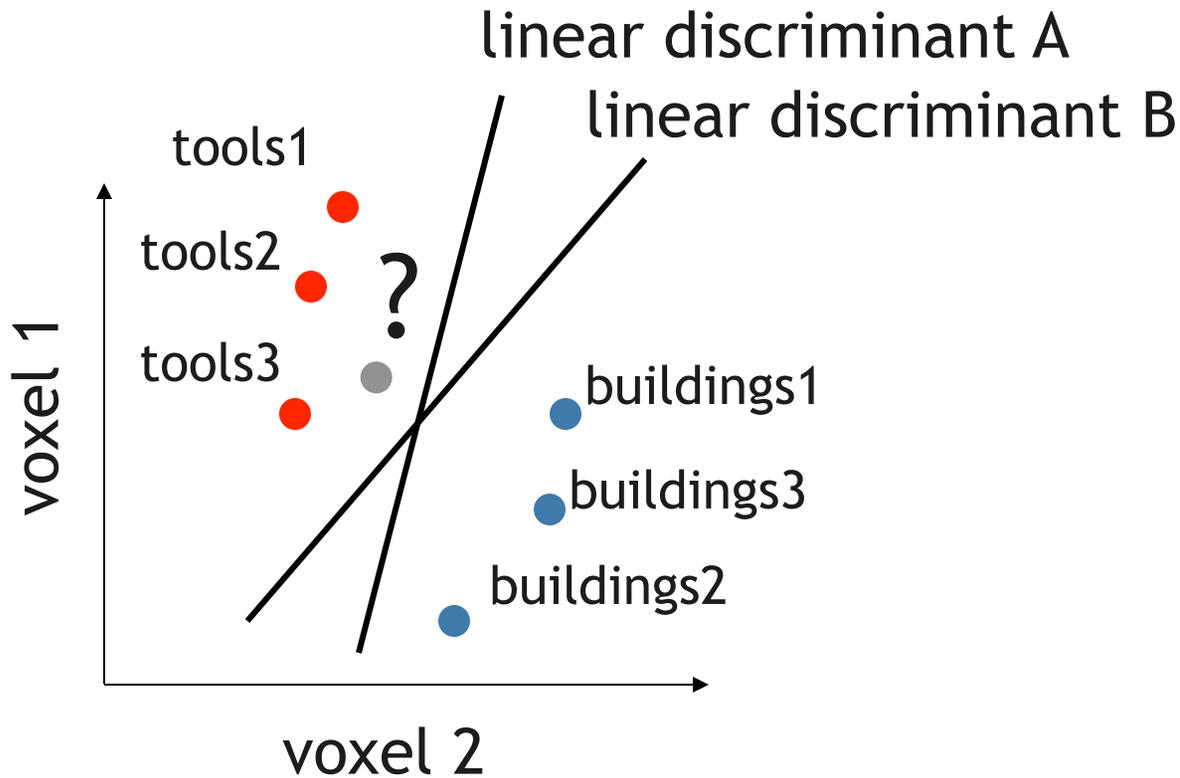
what is inside the box?

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



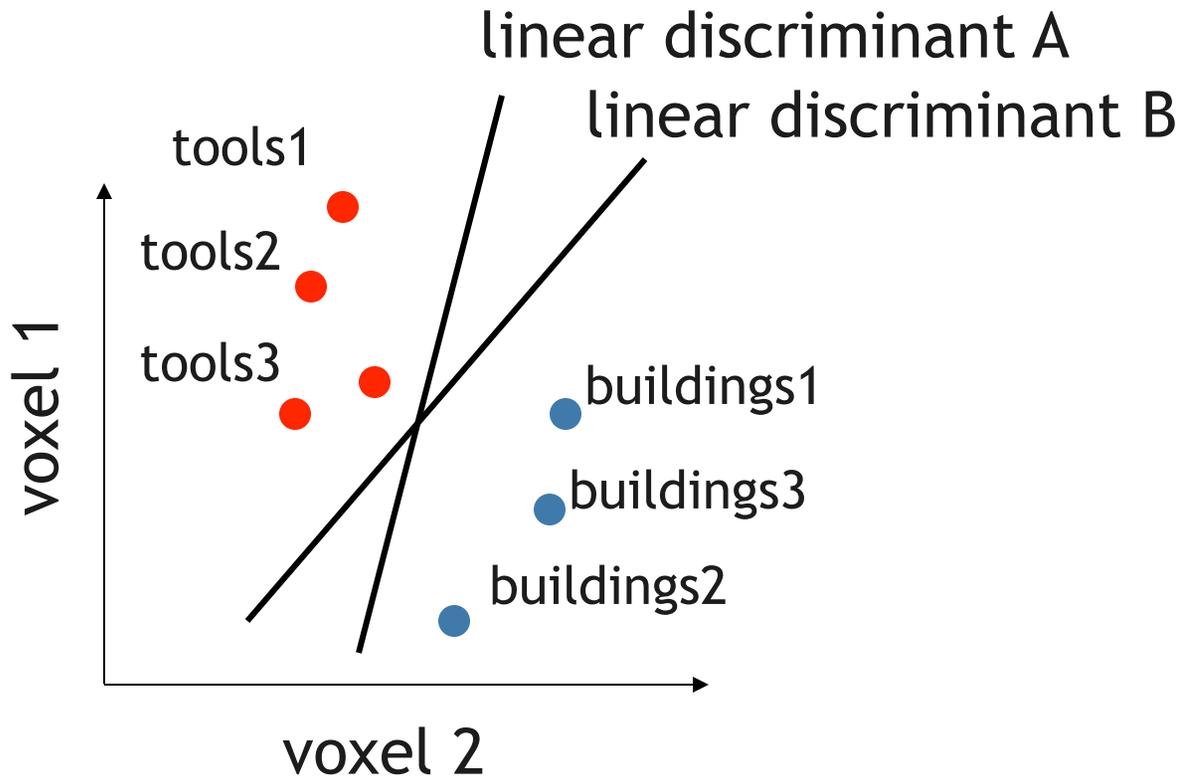
what is inside the box?

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

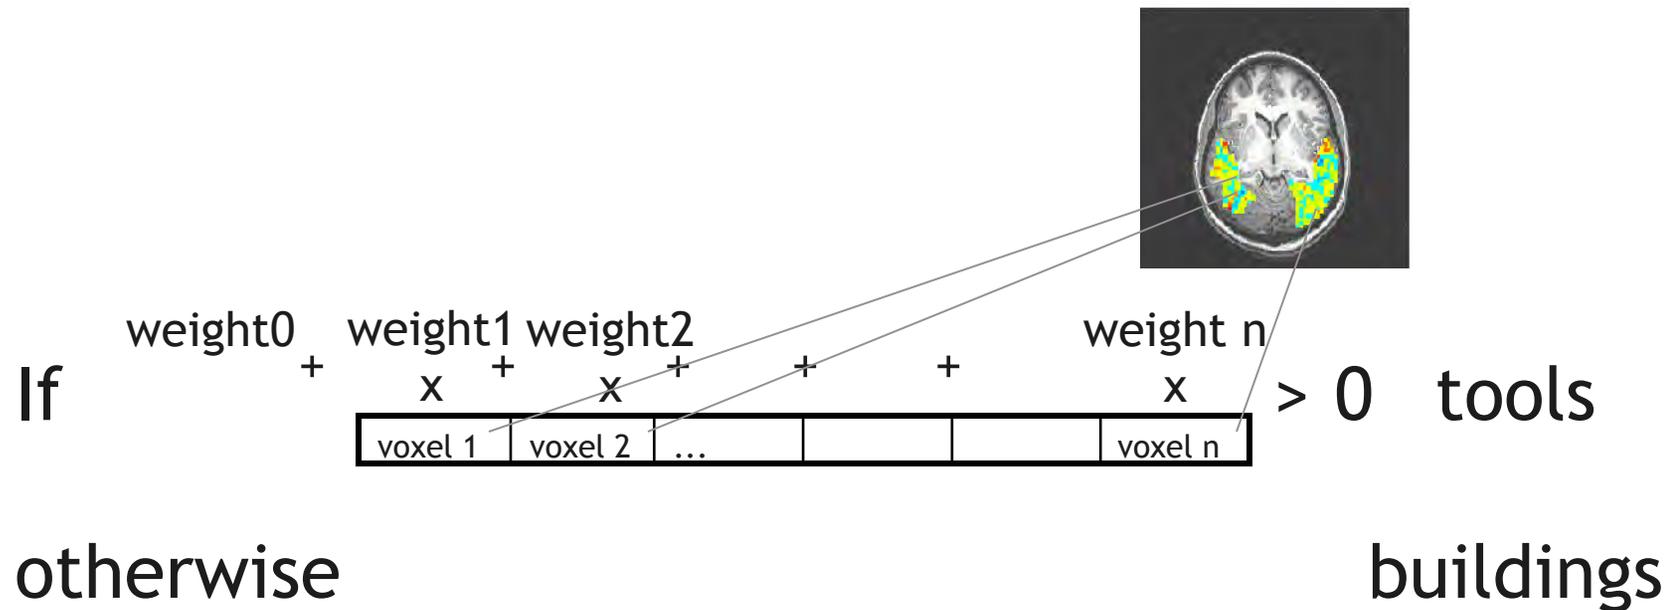


what is inside the box?

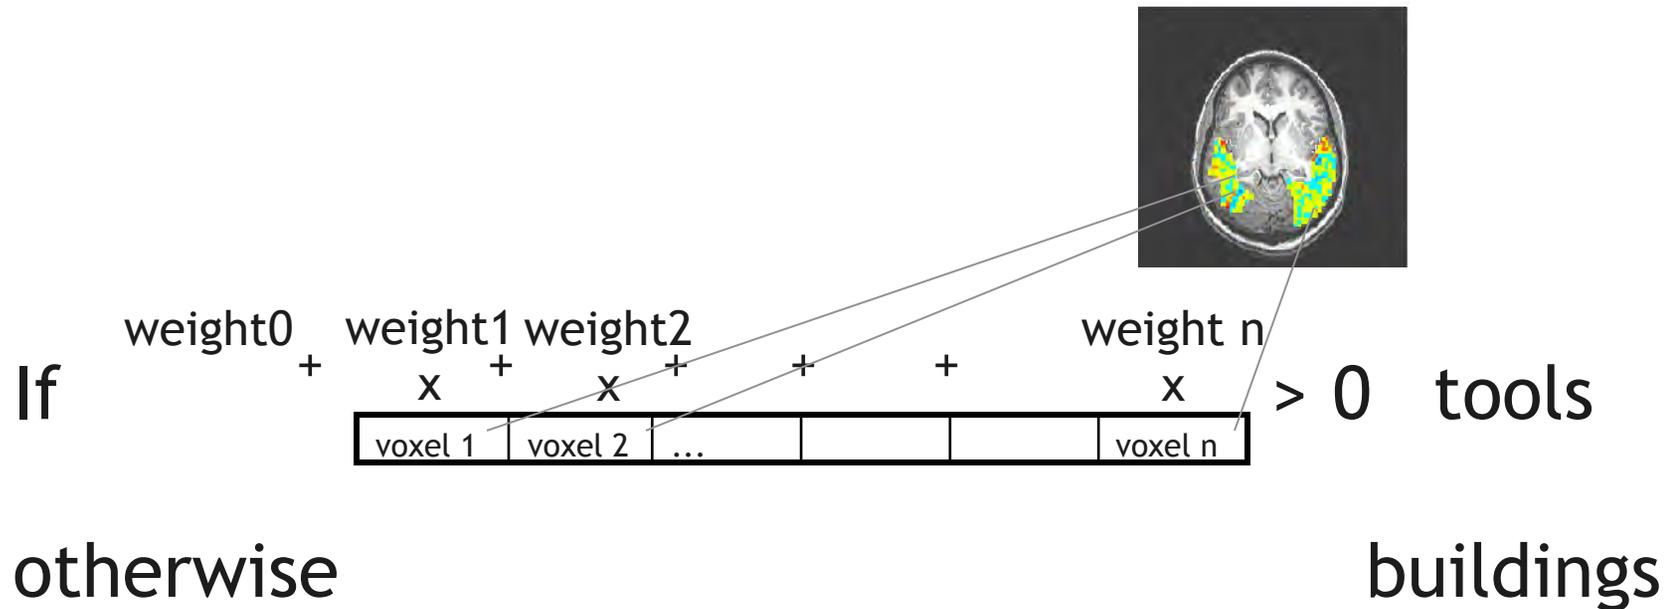
- next simplest: learn linear discriminant
- note that there are many solutions...



linear classifiers



linear classifiers



various kinds:

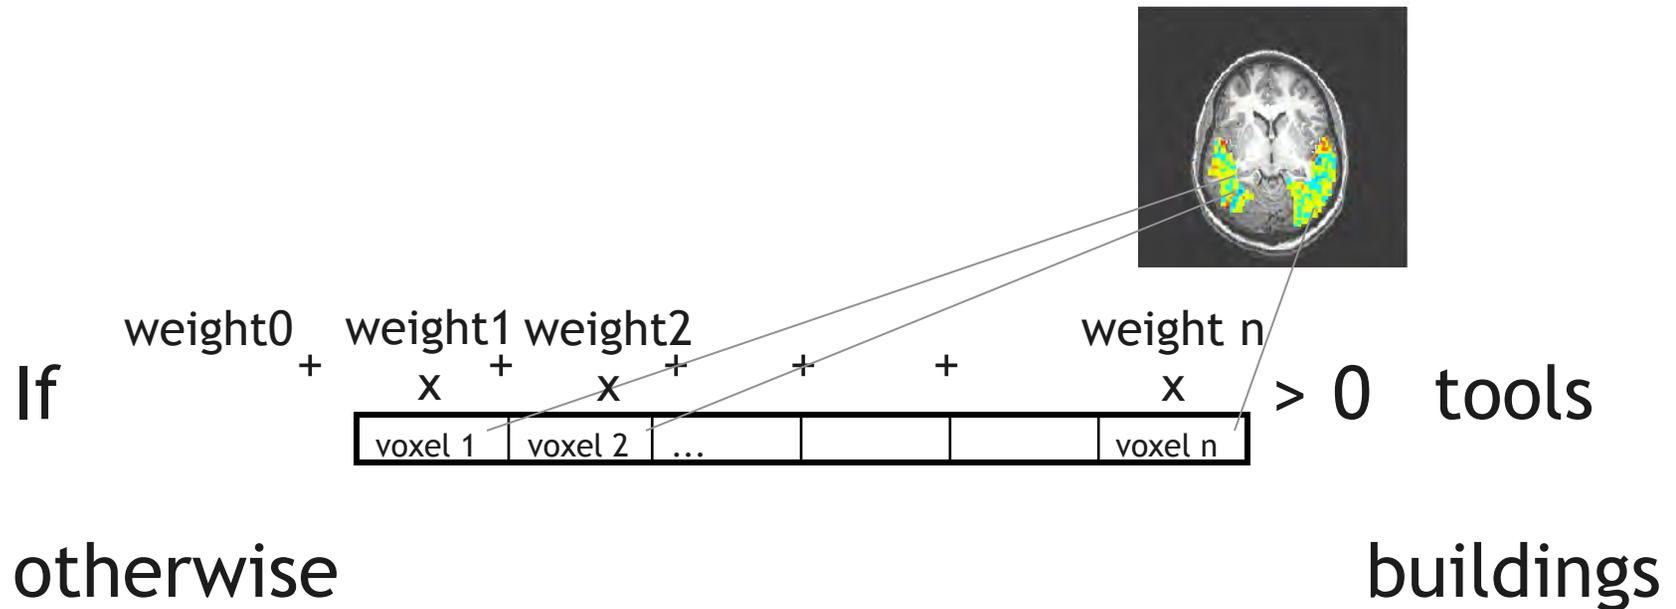
Gaussian Naive Bayes

Regularized Logistic Regression

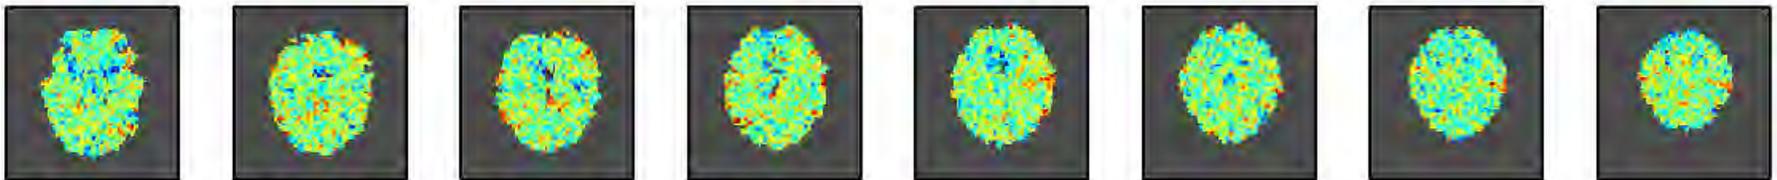
Linear Support Vector Machines (SVM)

differ on how weights are chosen

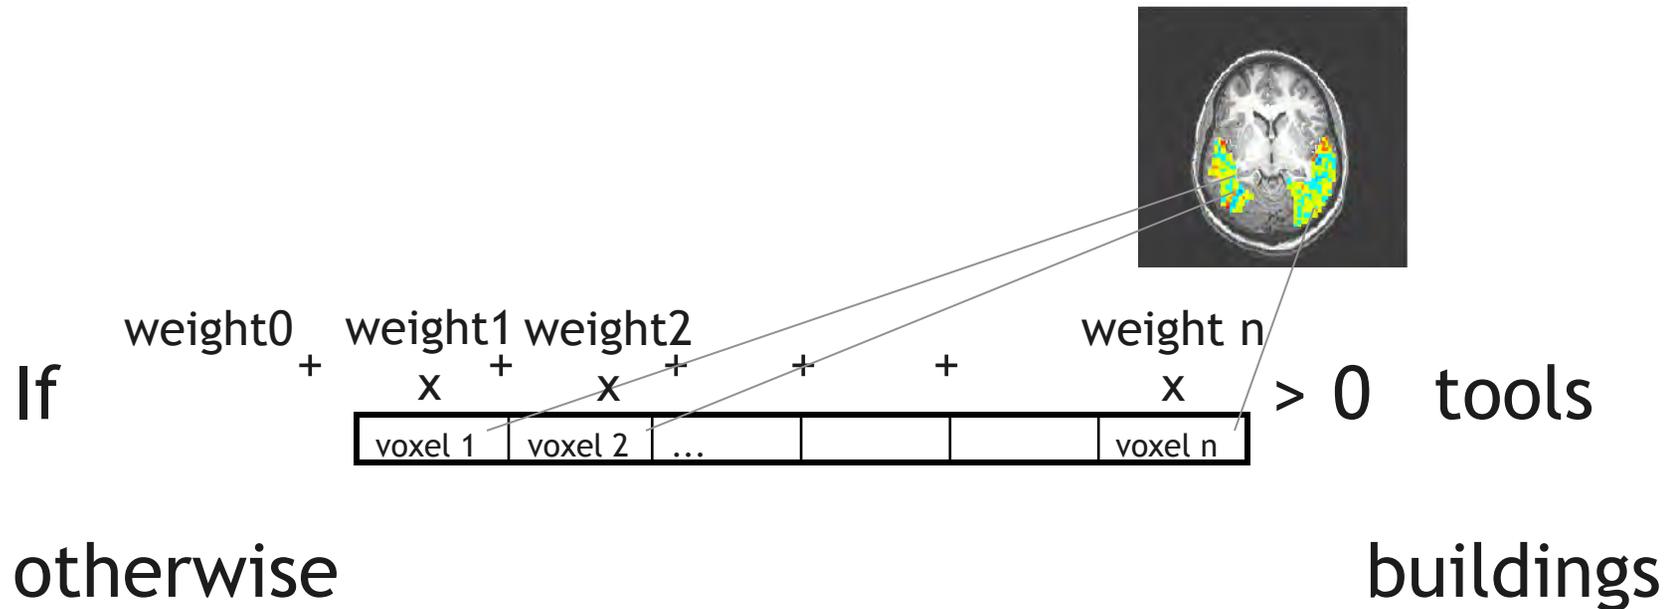
linear classifiers



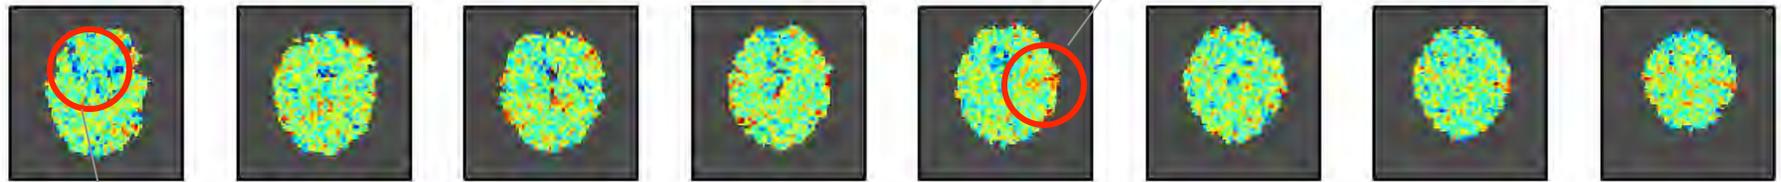
linear SVM weights:



linear classifiers



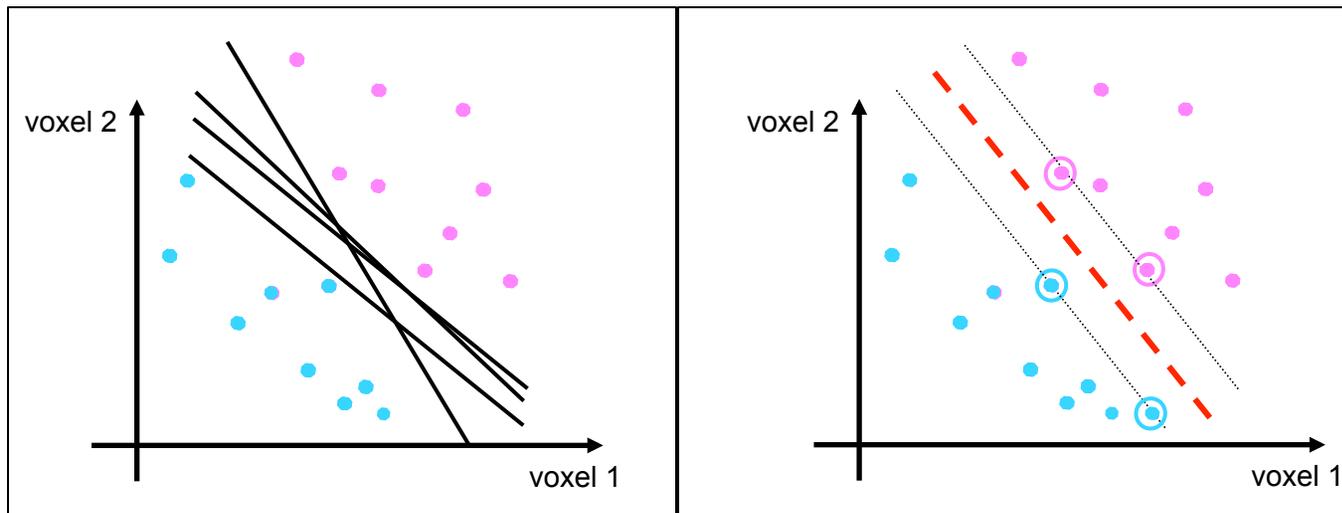
linear SVM weights:



weights pull towards tools

weights pull towards buildings

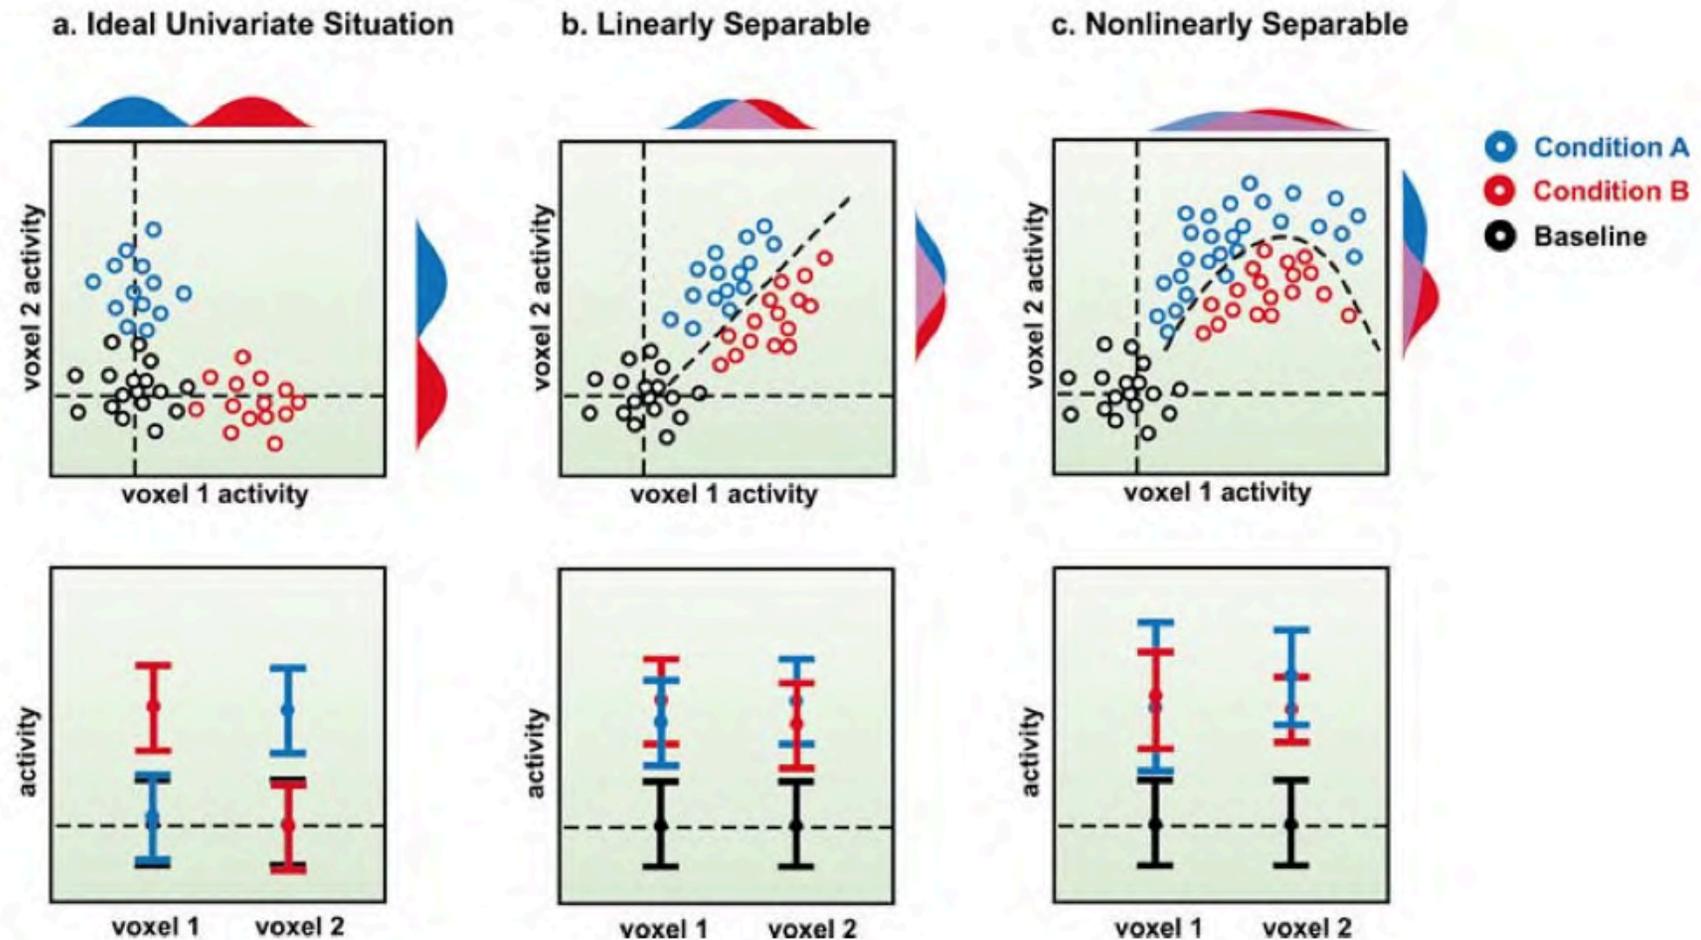
linear support vector machines



- Find linear decision boundary that maximizes the margin

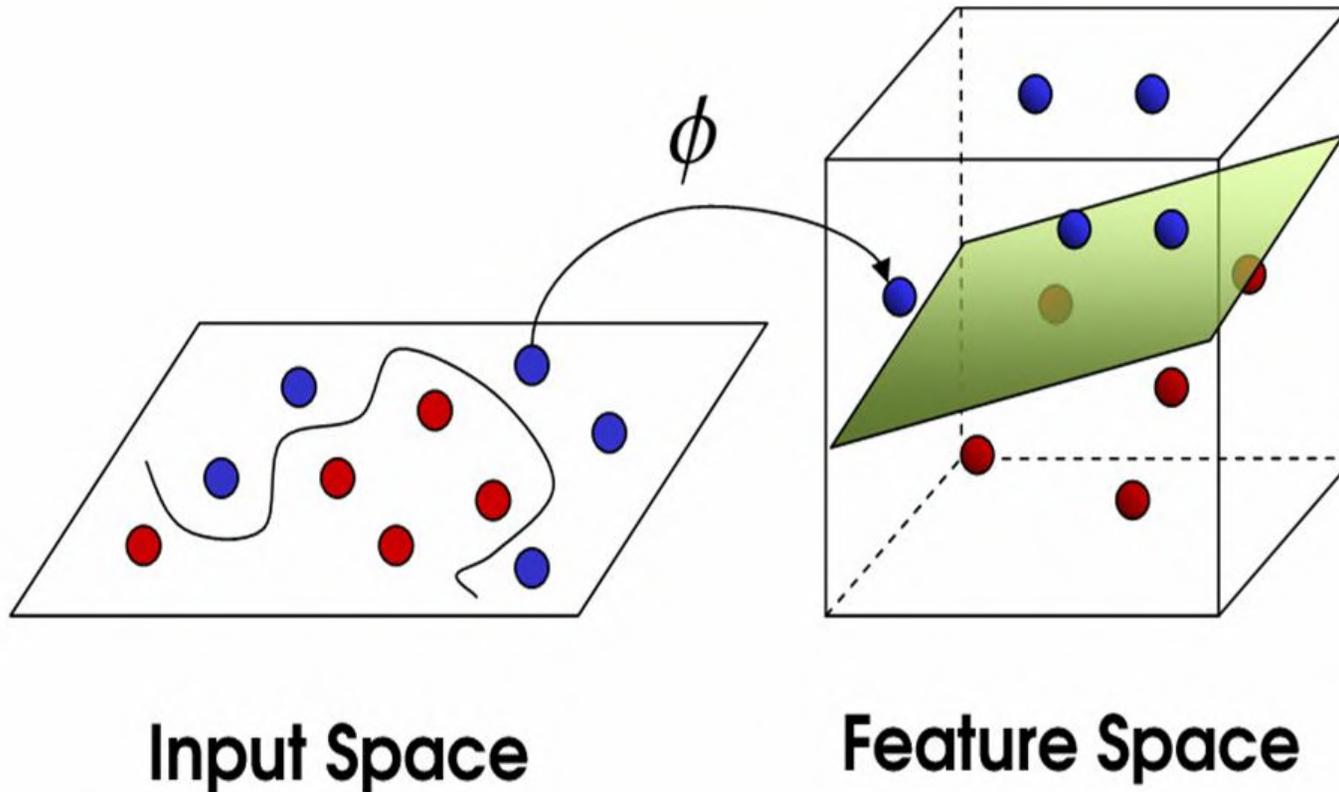
nonlinear classifiers

- sometimes, two classes will not be linearly separable



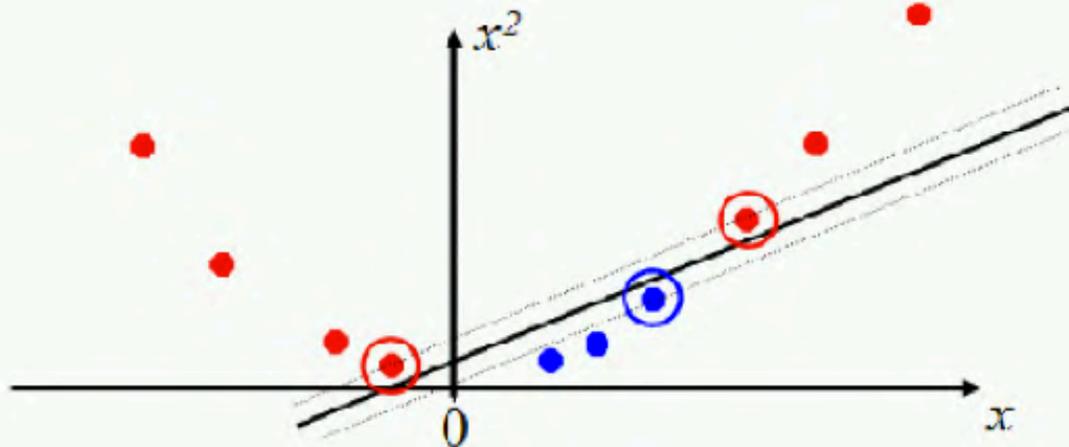
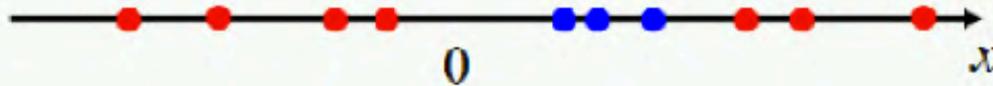
nonlinear classifiers

- 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



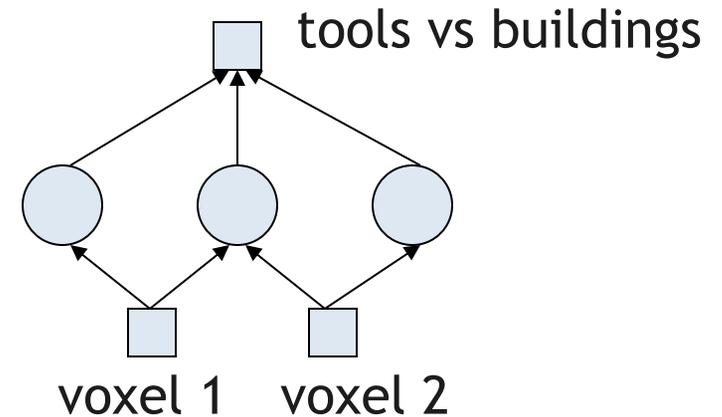
The “kernel trick”

(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.

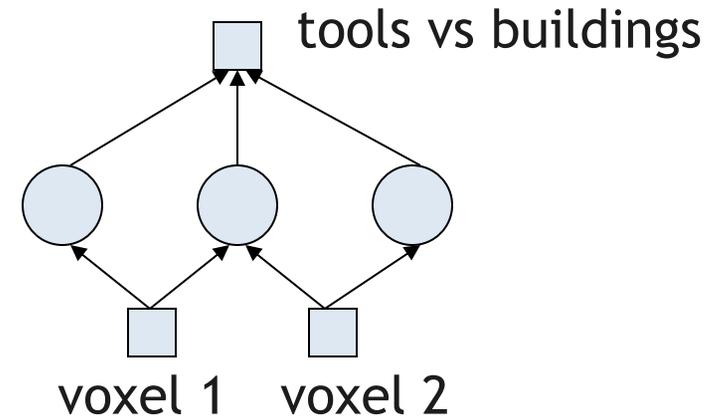
nonlinear classifiers

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

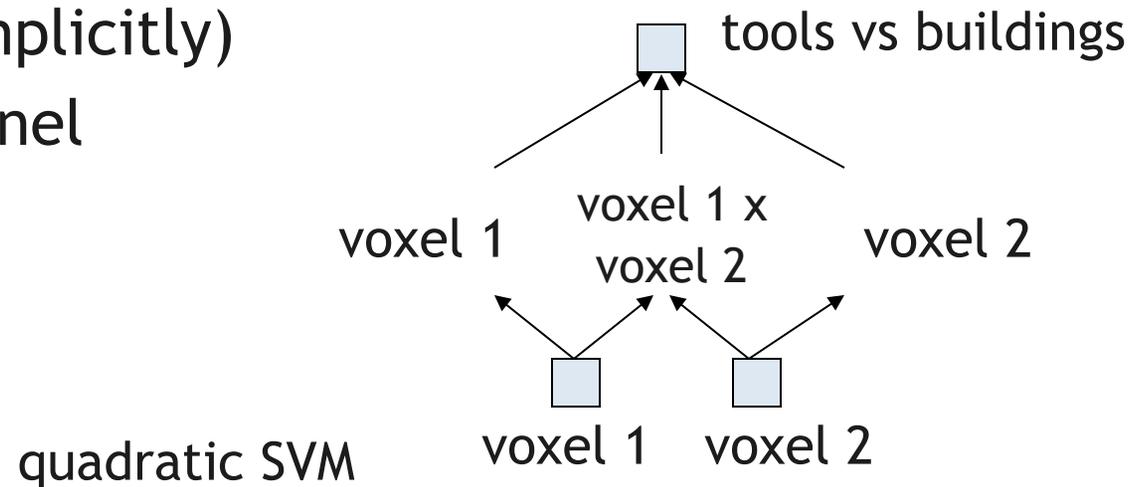


nonlinear classifiers

- **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



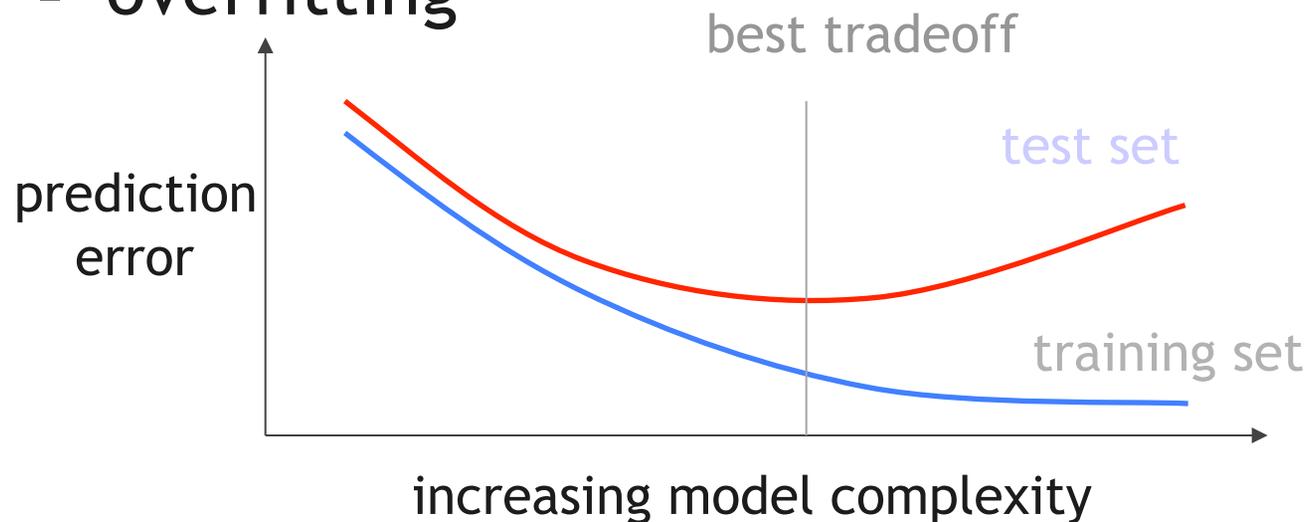
nonlinear classifiers

reasons to be careful:

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



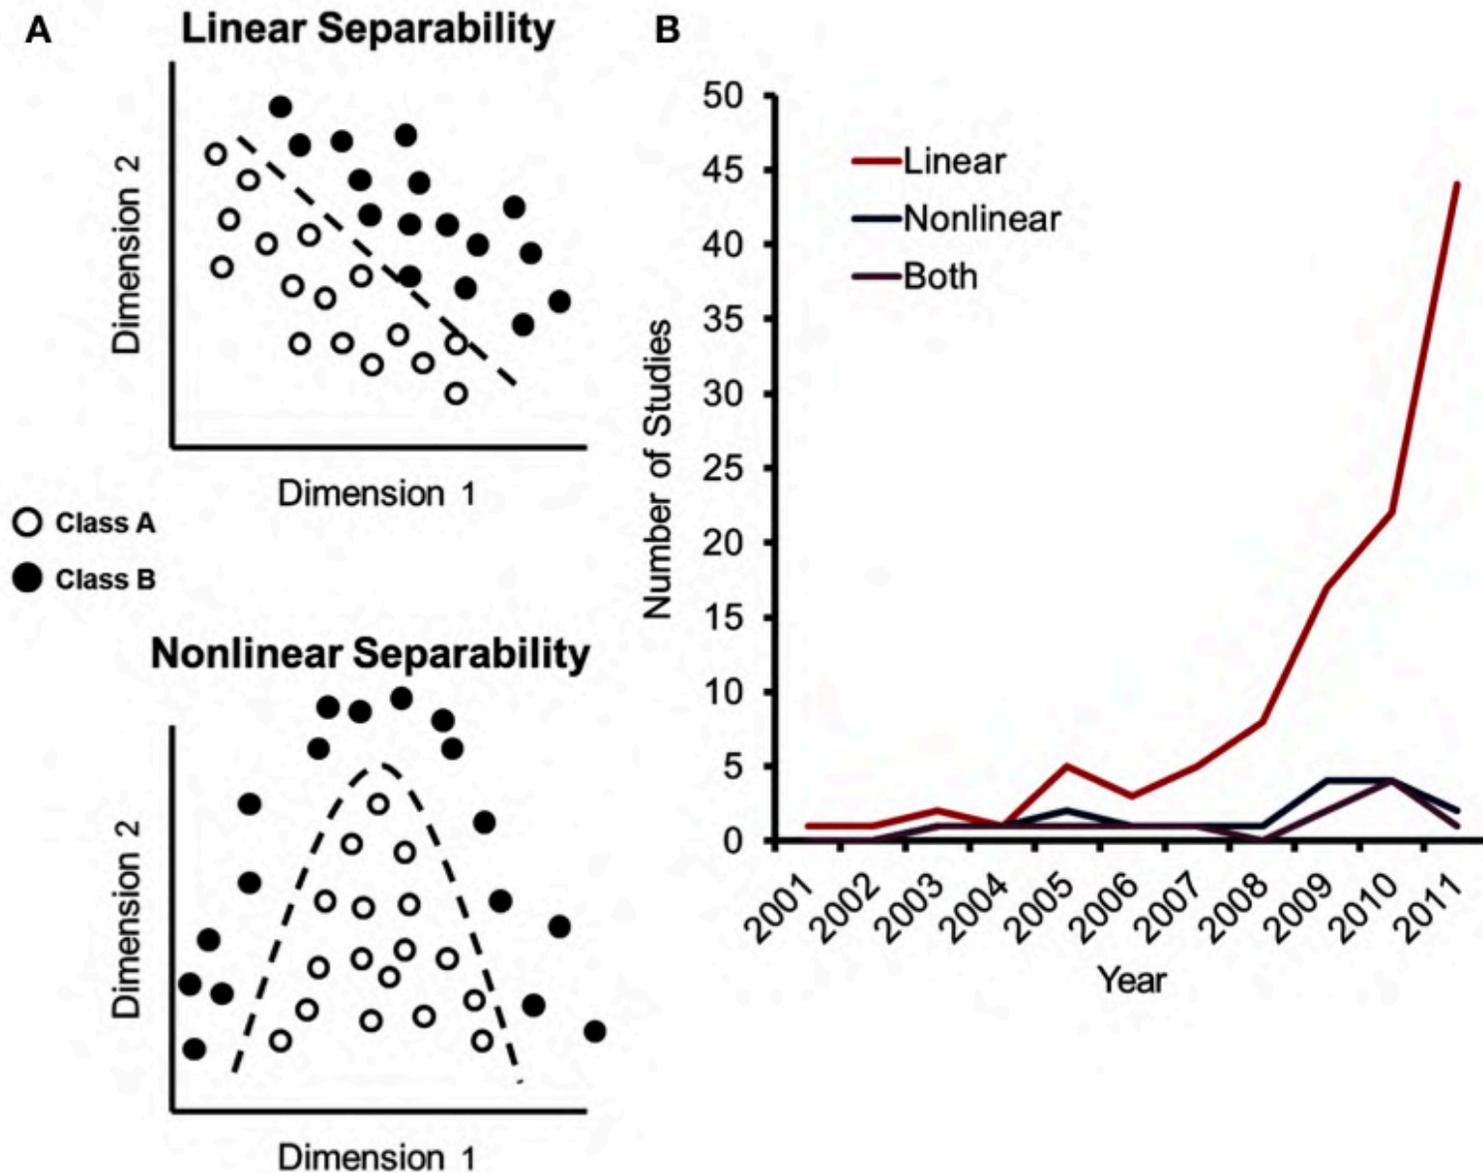
- overfitting



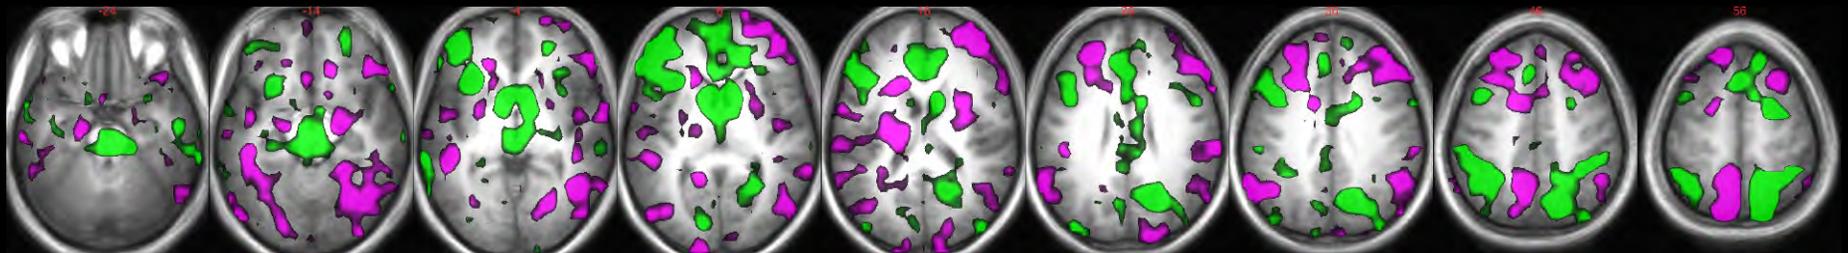
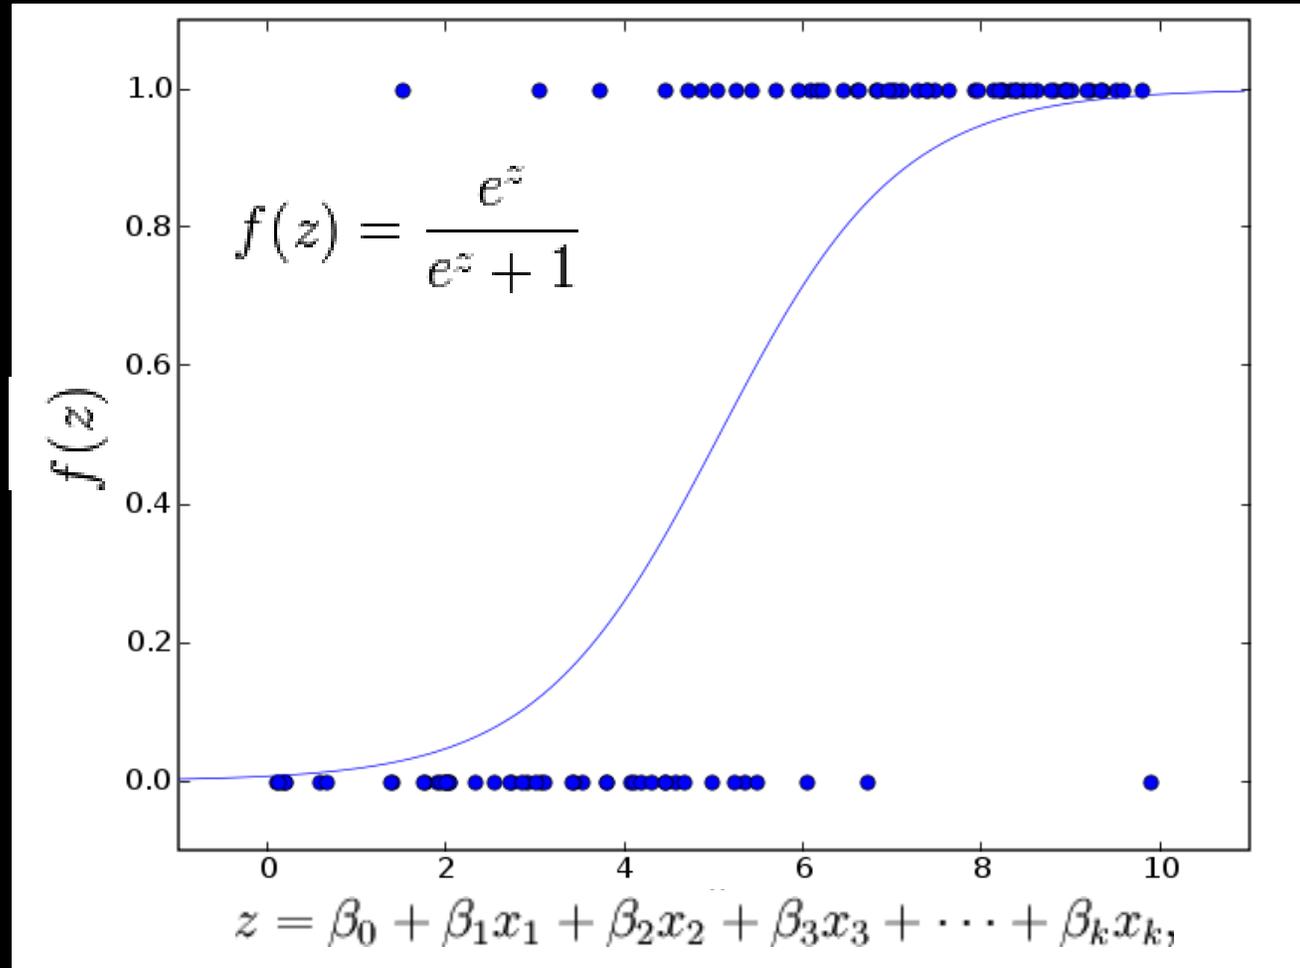
[from Hastie et al,2001]

Slide from Francisco Pereira

linear vs. nonlinear classifiers

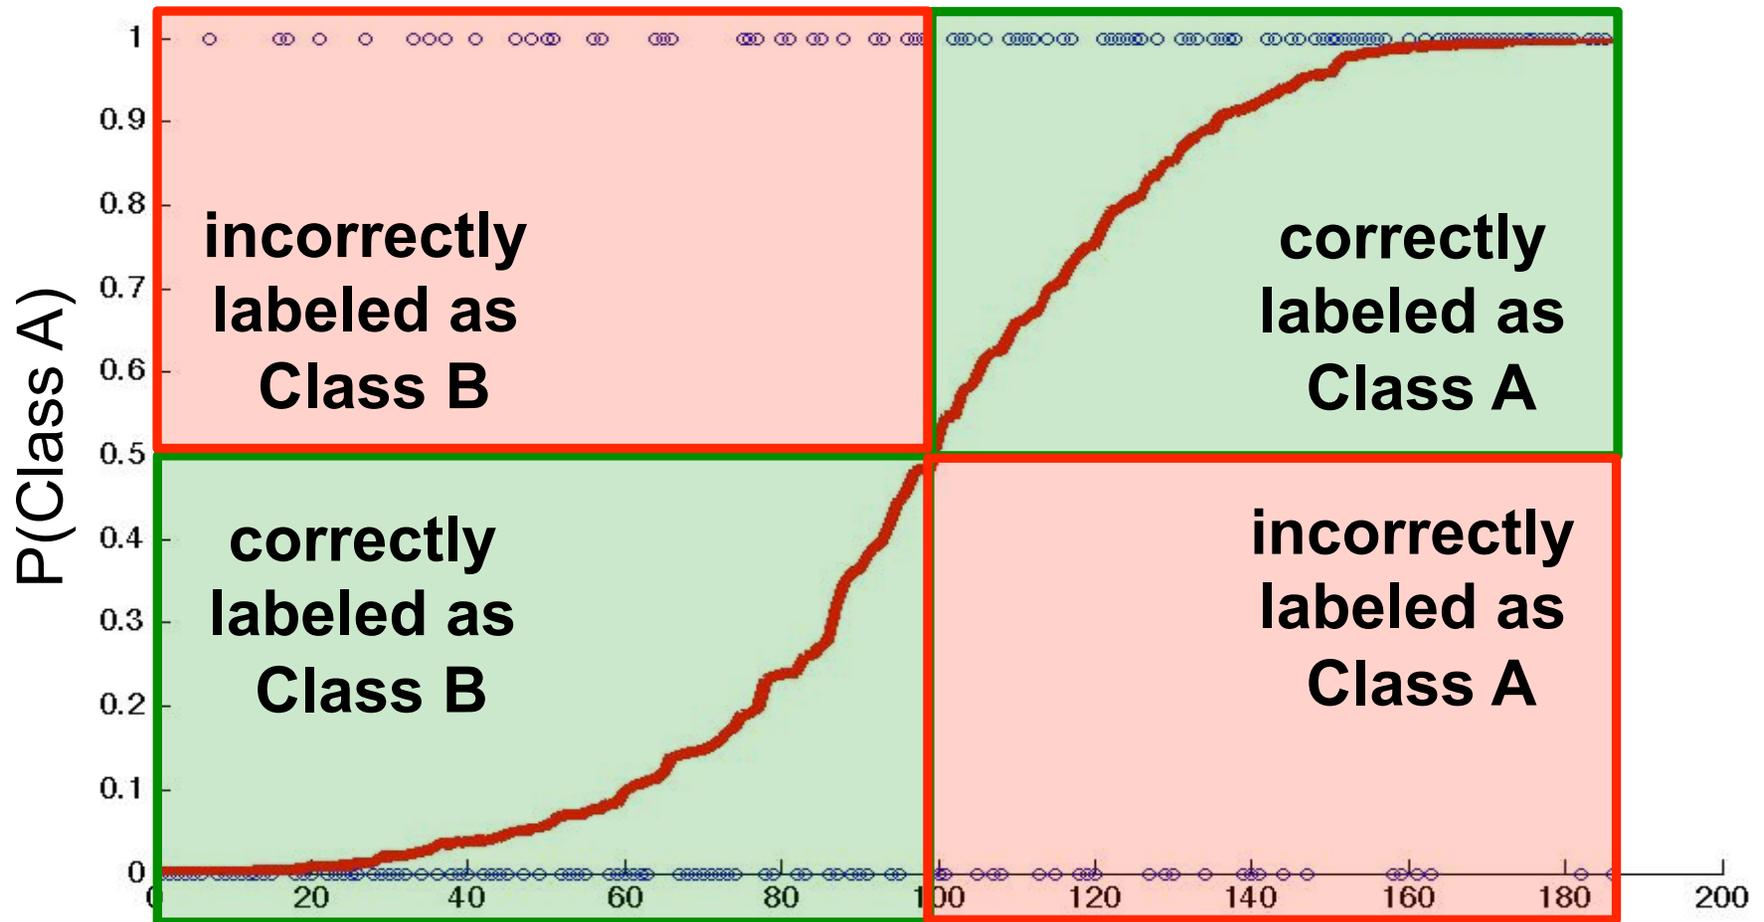


Logistic Regression Classifier



Testing the trained classifier:

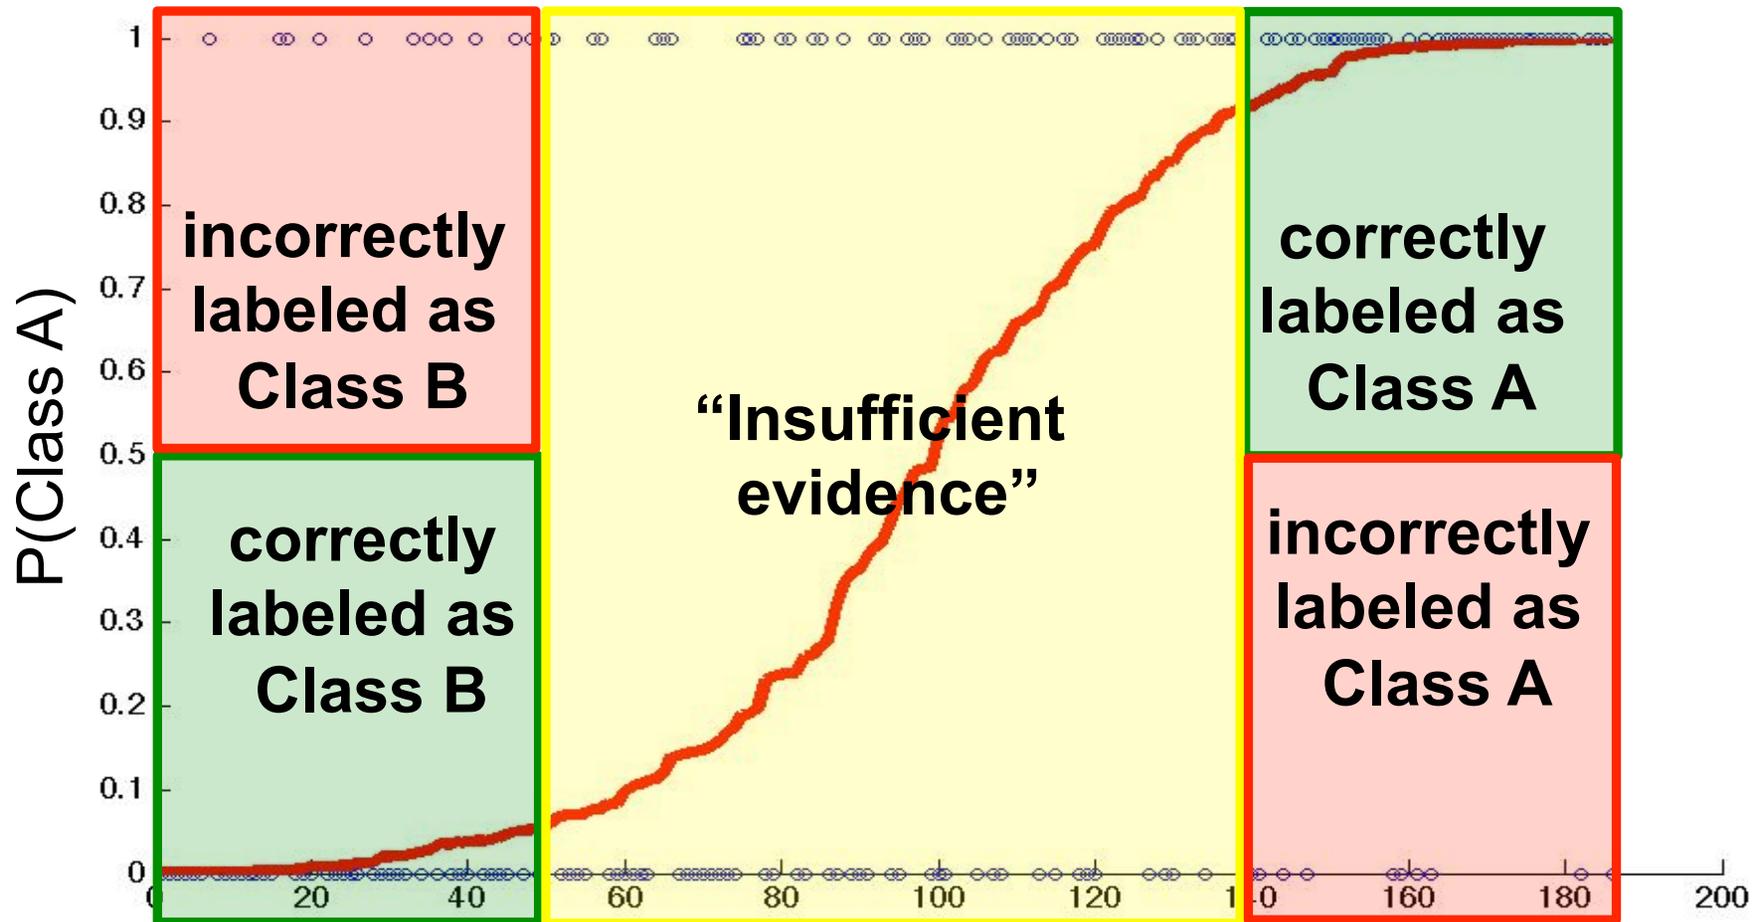
Ranked classifier predictions for individual test trials



70% accuracy

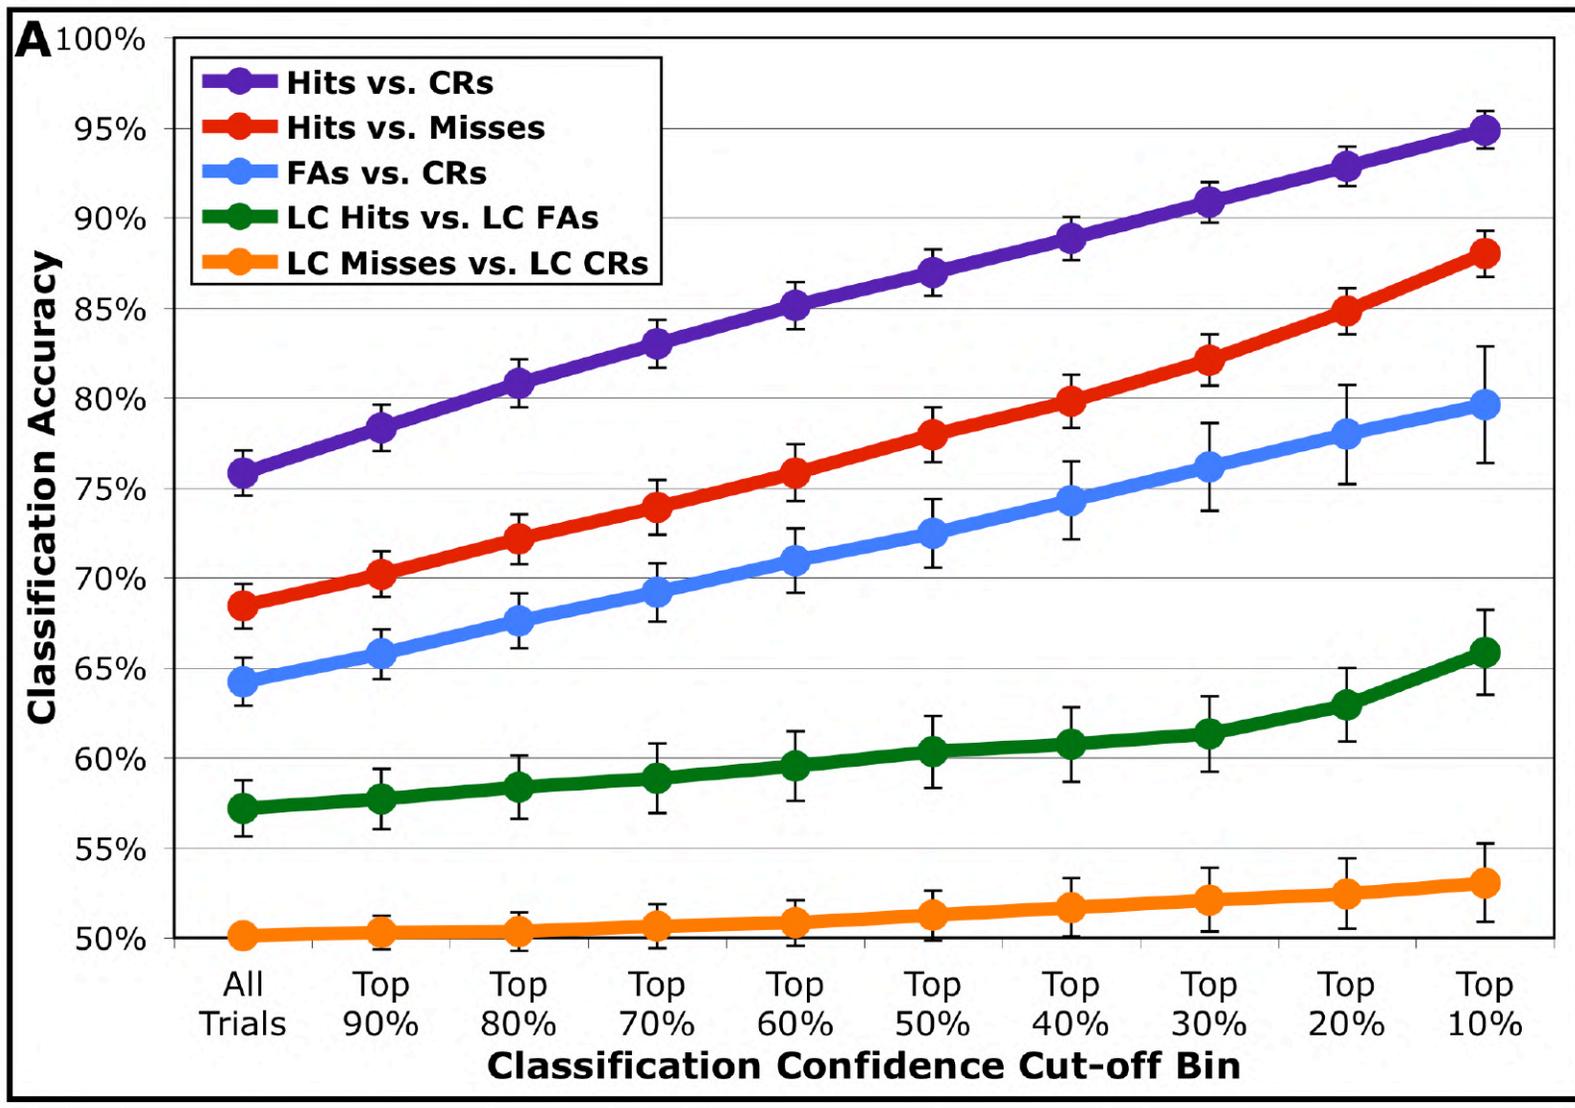
Testing the trained classifier:

Ranked classifier predictions for individual test trials

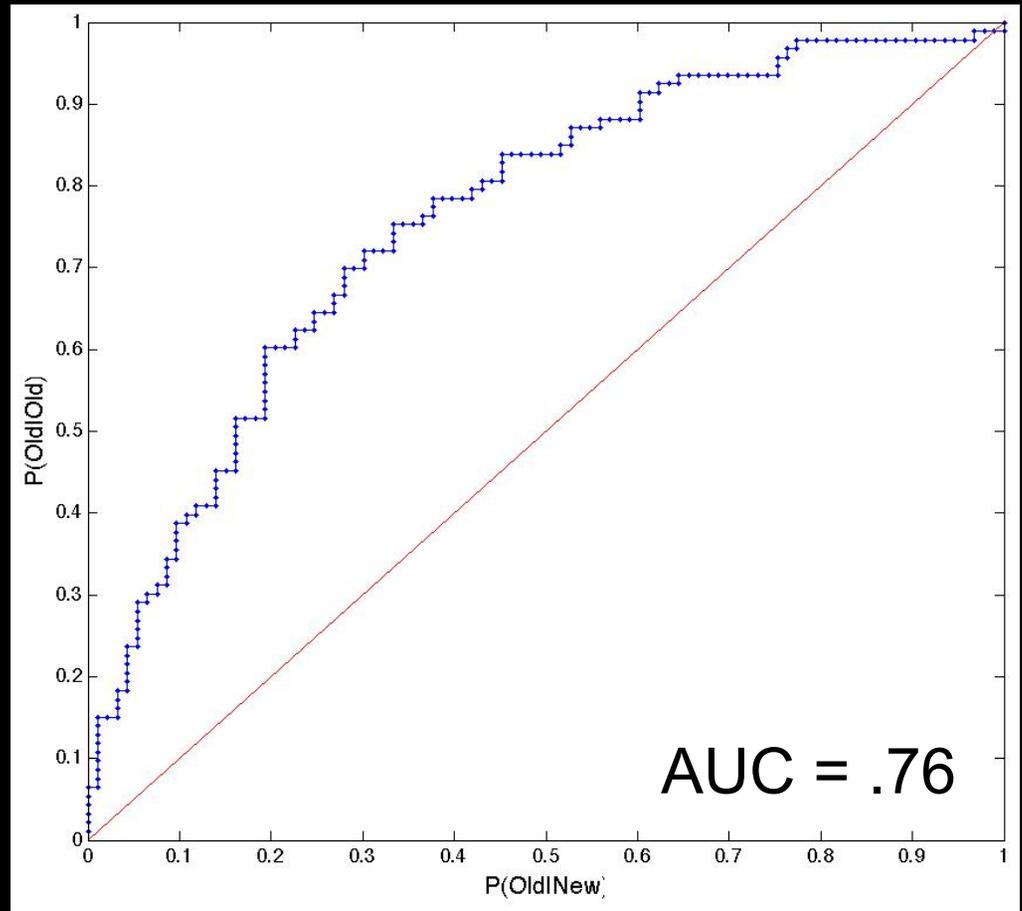
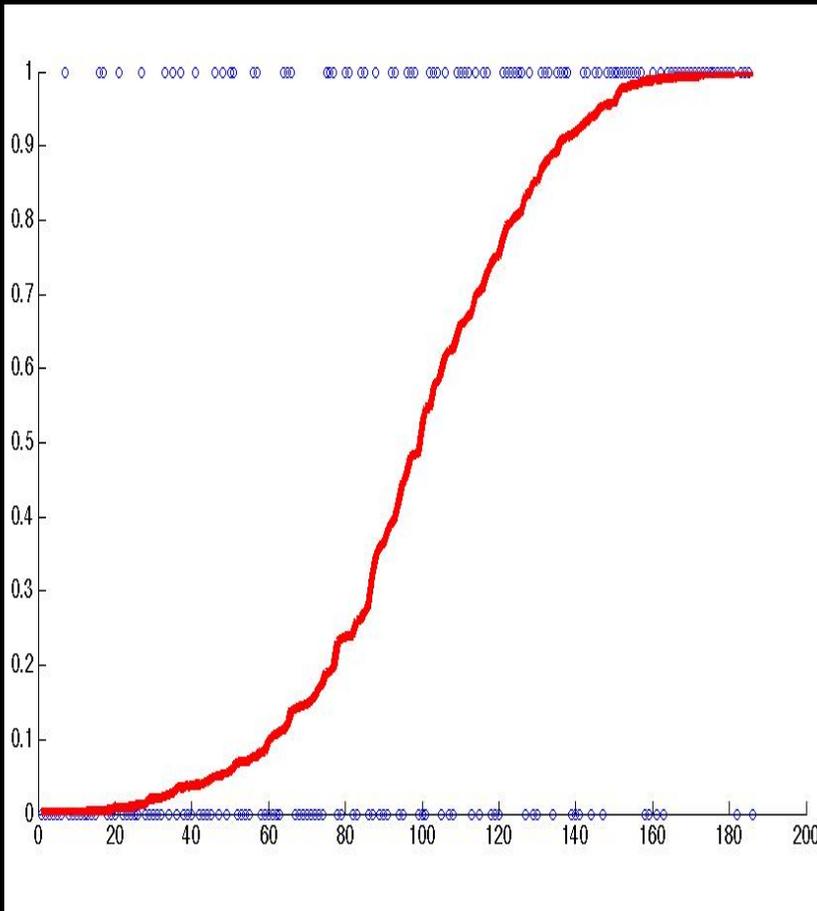


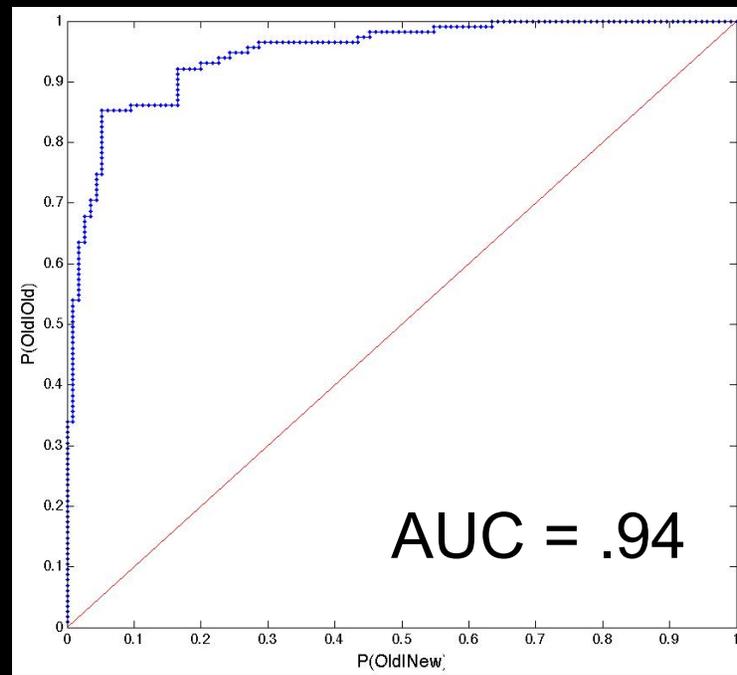
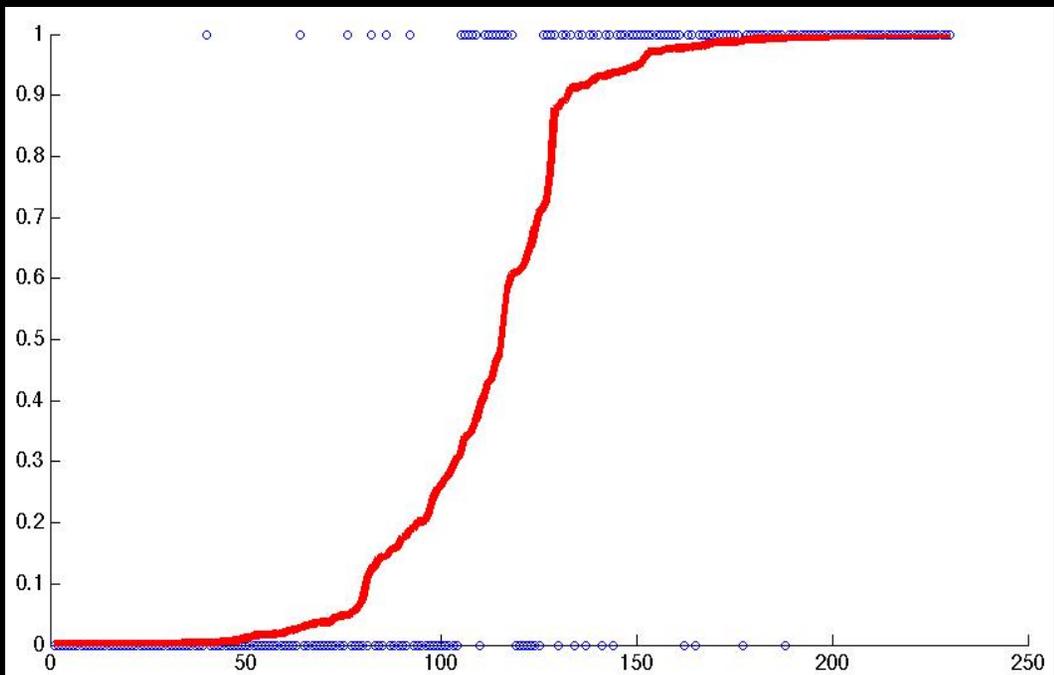
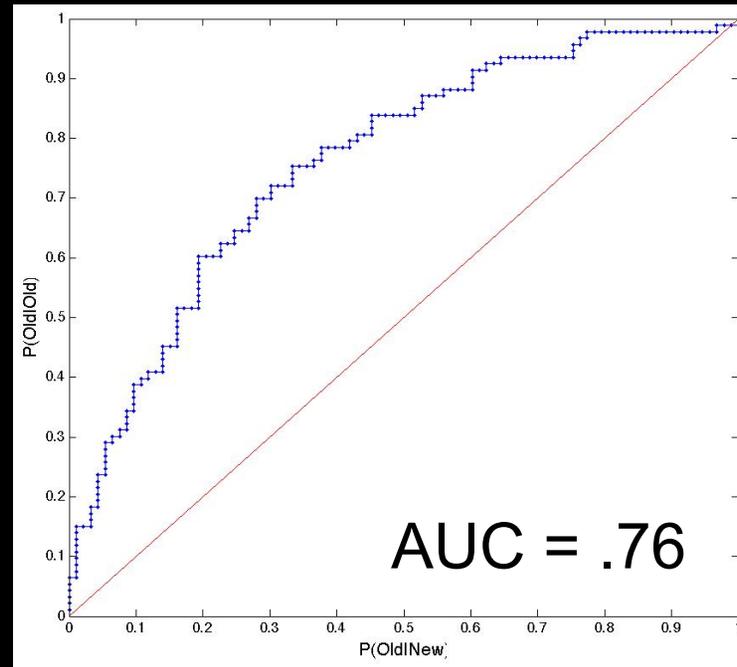
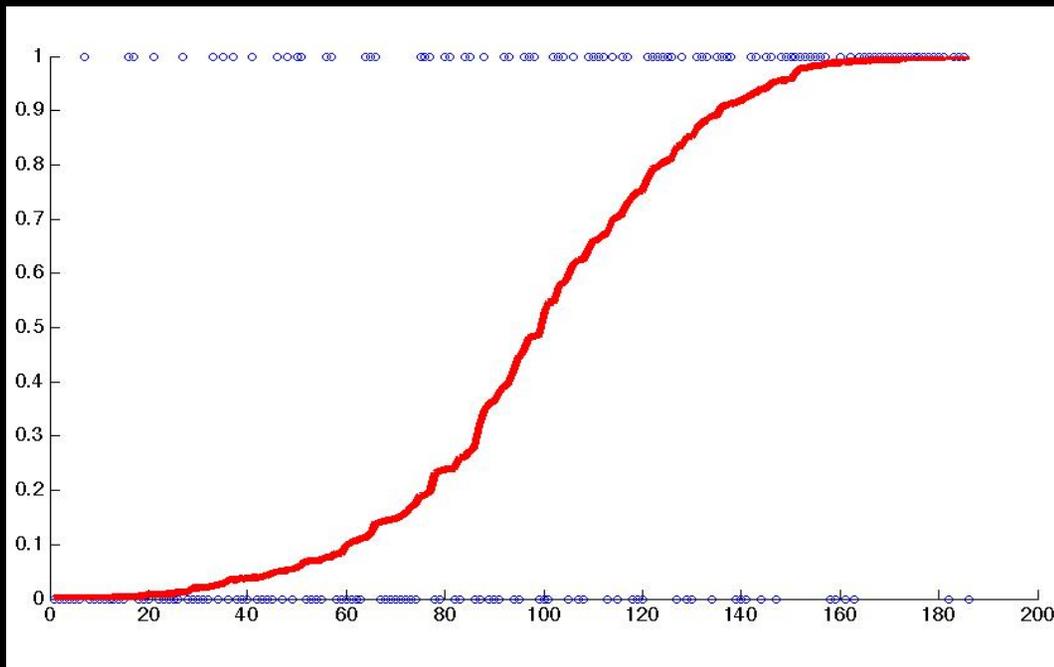
85% accuracy

Assessing a classifier's "confidence"

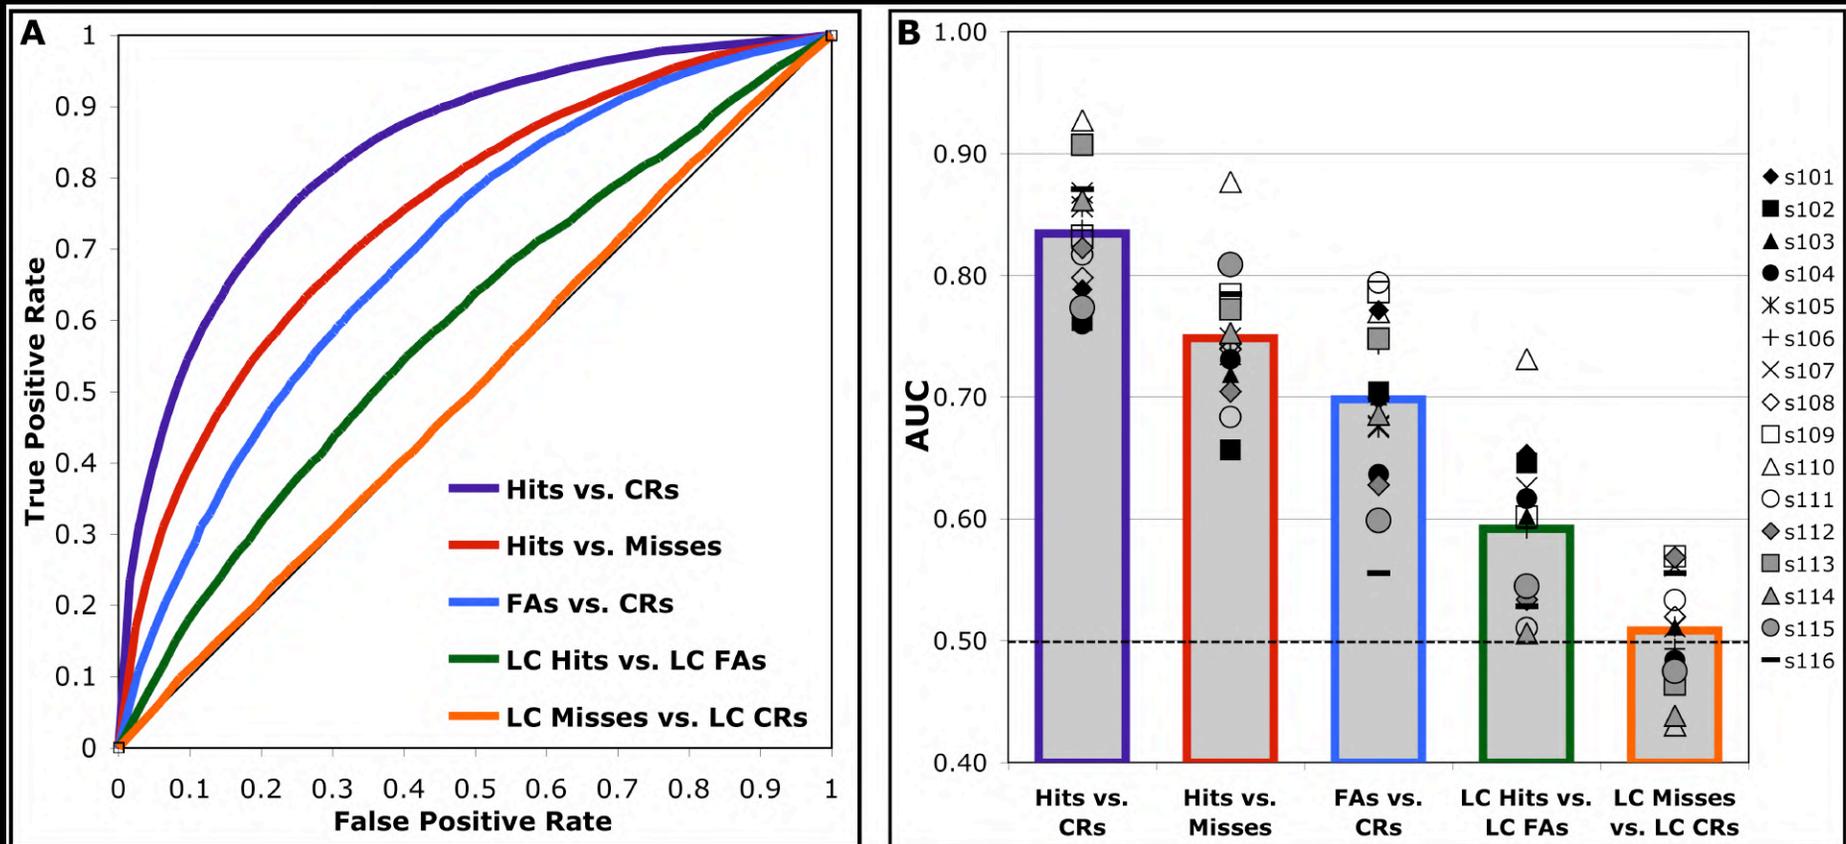


Representing classification performance with receiver operating characteristic (ROC) curves





Group summary of classification ROC data for a recognition memory experiment

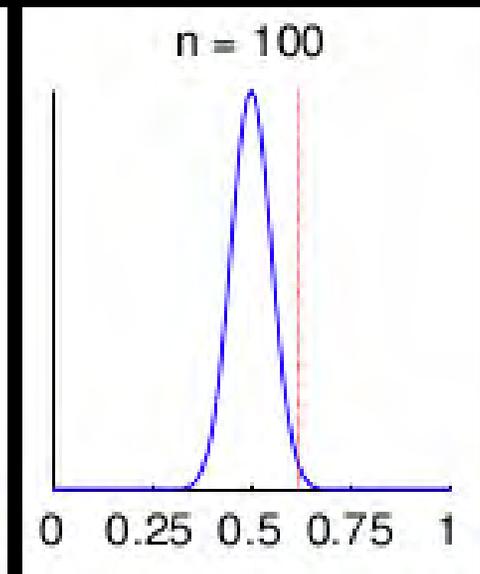
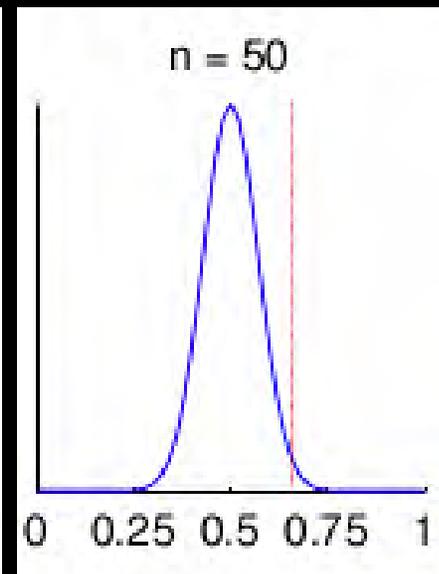
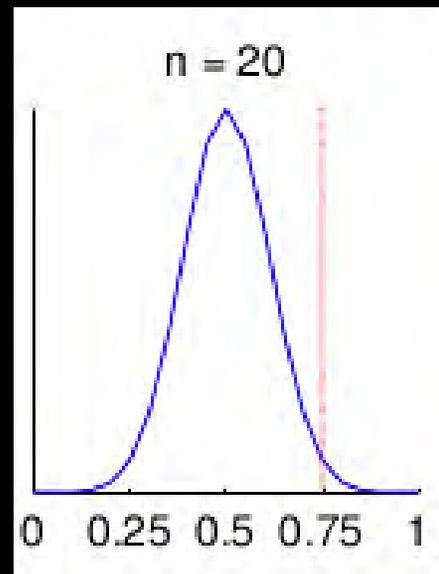
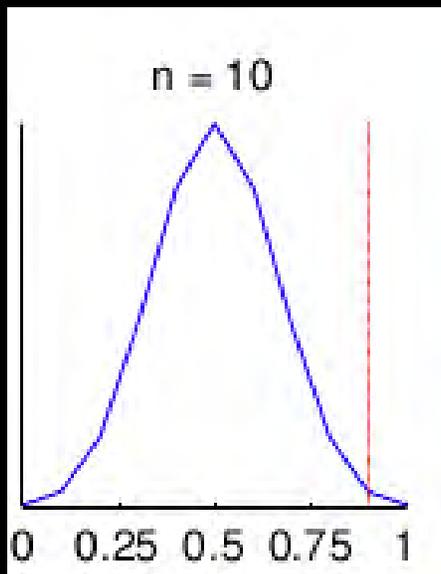


Determining whether classification accuracy is significantly better than chance

– 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

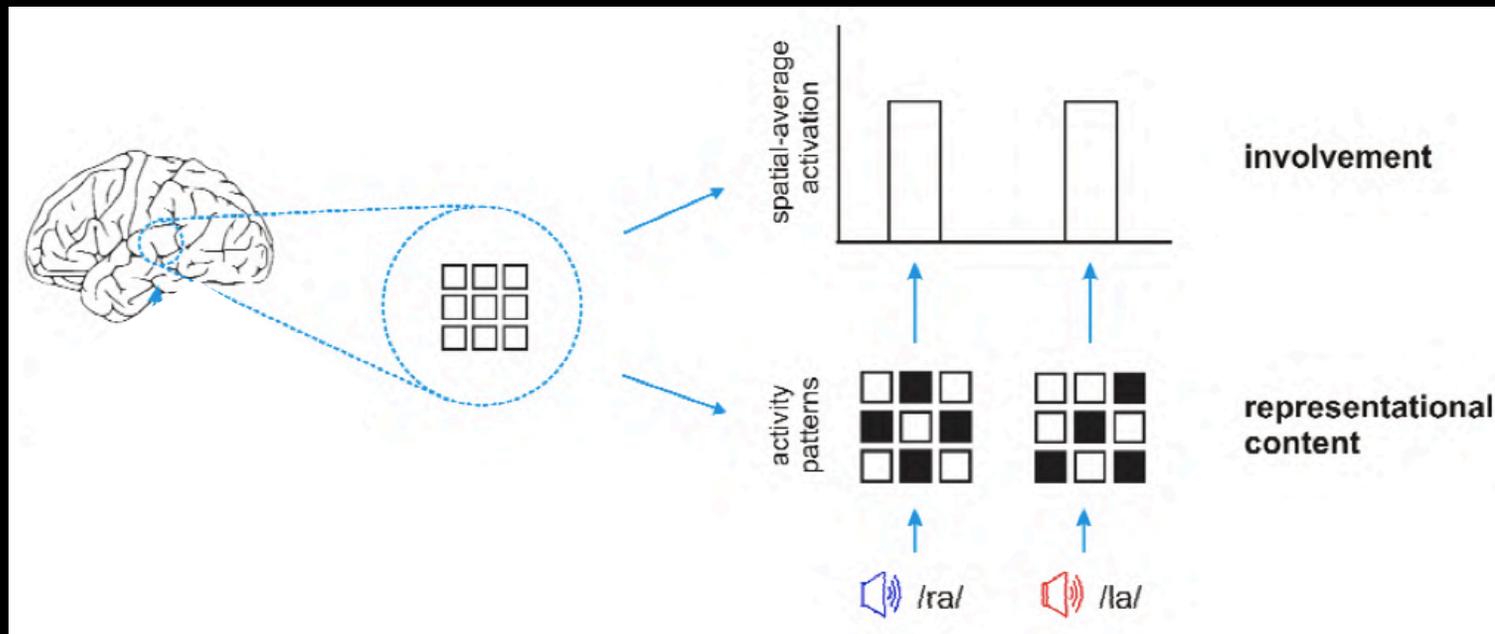
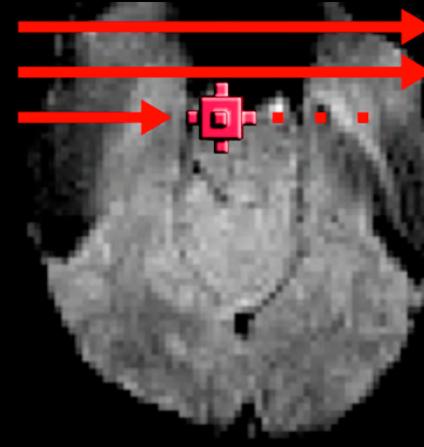
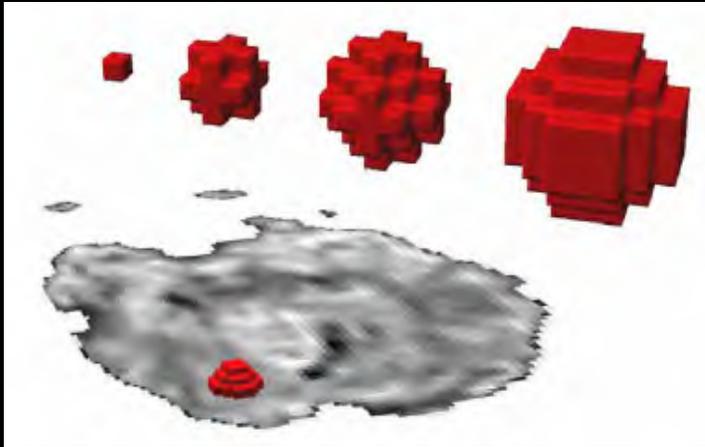
- Can compare to chance using binomial distribution
 - e.g., $p_{chance} = .5$, $trials = 120$, $successful = 75$, $p_{observed} = .0039$
- Or use permutation test to compare observed accuracy to distribution of performance generated by shuffling regressors many times (e.g., 1000 shuffles)

– 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*)

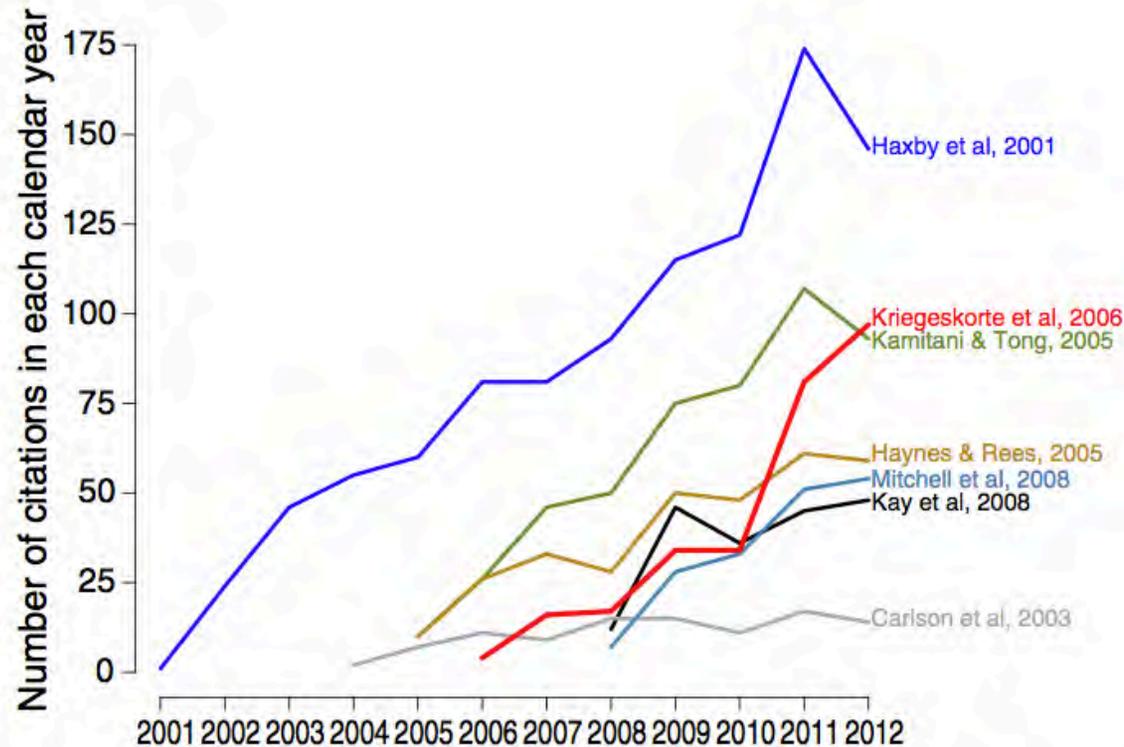
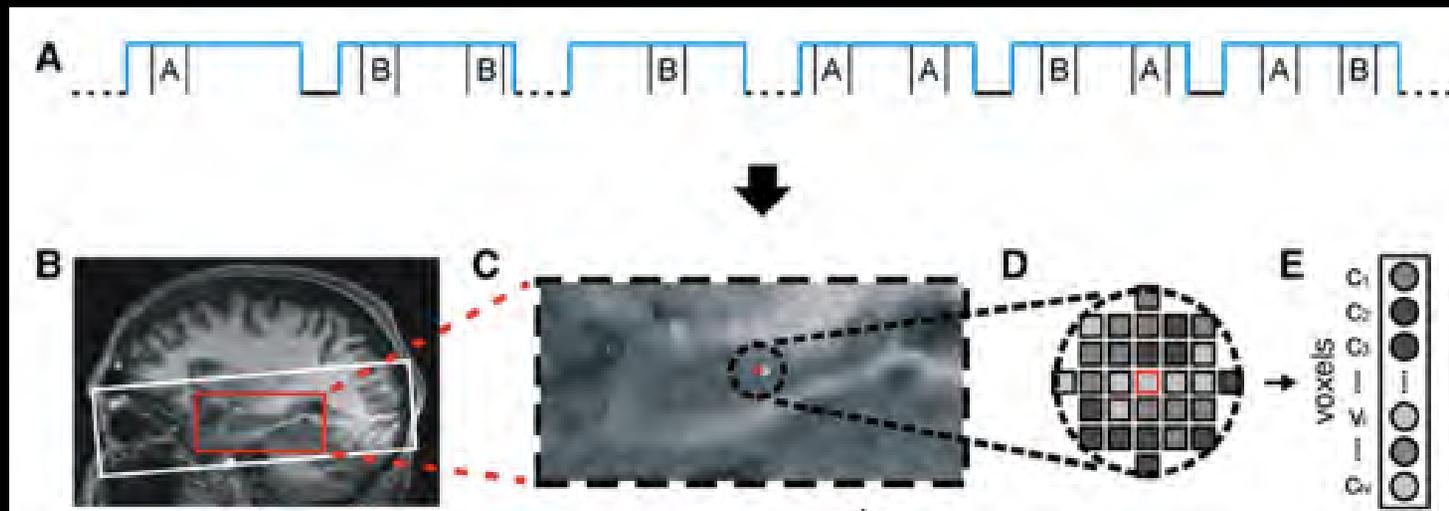
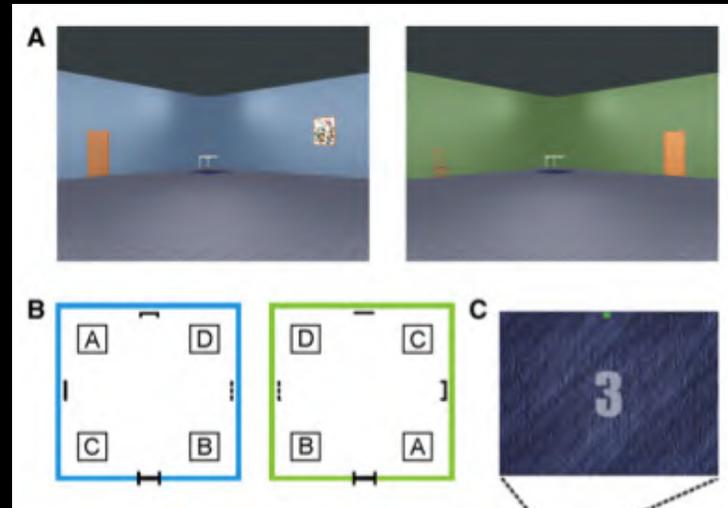


Fig. 1. Pattern-information fMRI is still a rapidly growing field, particularly searchlight analysis (note the rapid increase in papers citing Kriegeskorte et al., 2006). This figure follows Fig. 2 in Raizada and Kriegeskorte (2010), but uses the actual citation counts after 2008. The number of citations 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).

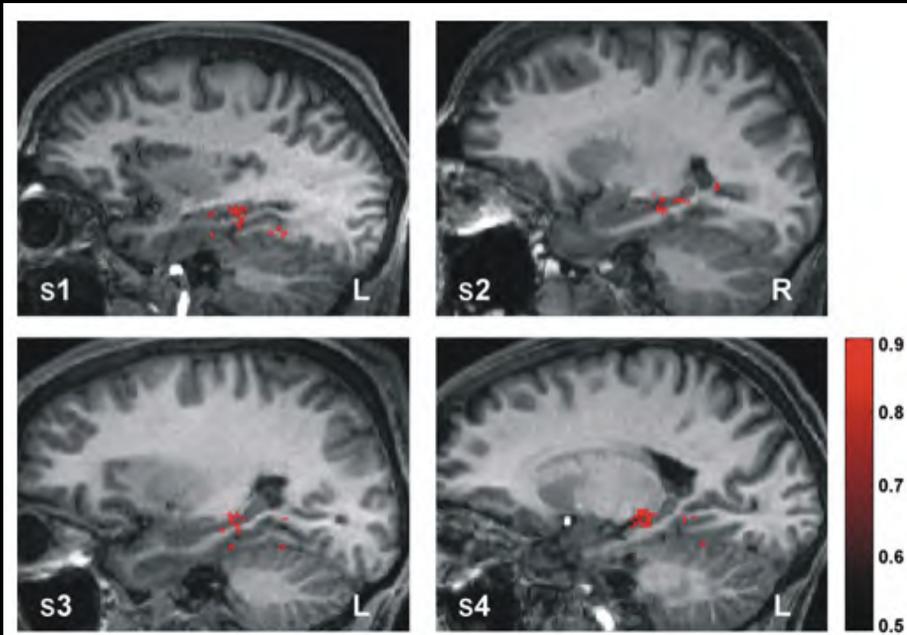
Decoding the representational content of medial temporal lobe activity patterns with high resolution fMRI data



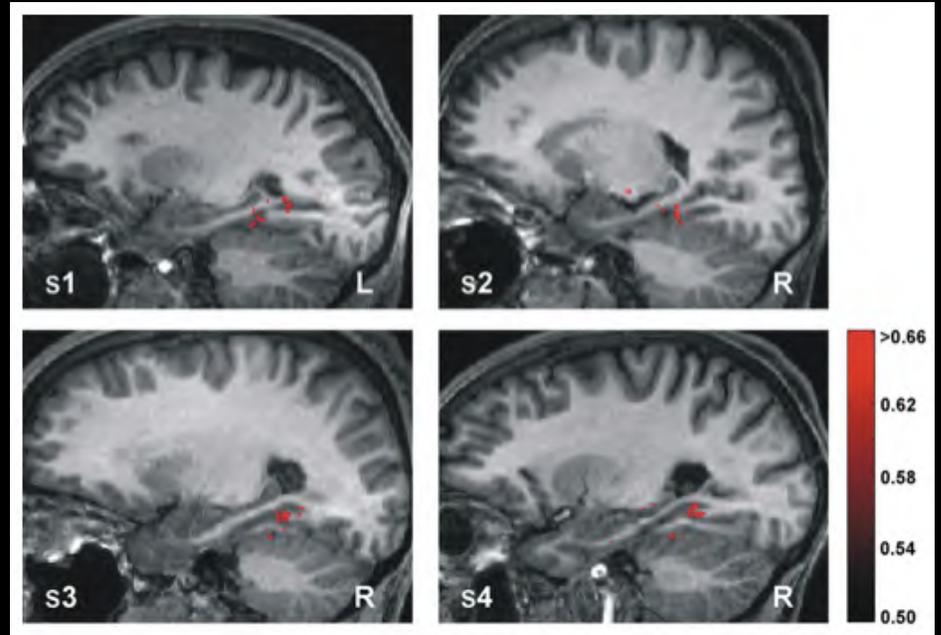
Hassabis et al. (2009) *Neuron*

Decoding the representational content of local activity patterns within the MTL

Which spatial quadrant?



Which room?





Tracking cognitive fluctuations with multivoxel pattern time course (MVPTC) analysis

Yu-Chin Chiu^{*,1}, Michael S. Esterman¹, Leon Gmeindl, Steven Yantis

Department of Psychological and Brain Sciences, Johns Hopkins University, United States

ARTICLE INFO

Article history:

Received 8 February 2011
Received in revised form 9 July 2011
Accepted 9 July 2011
Available online 23 July 2011

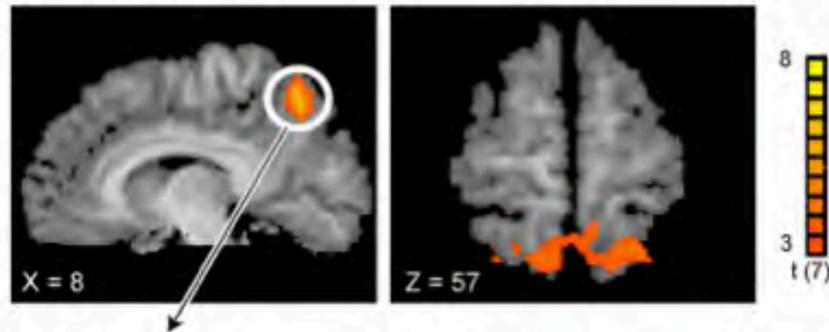
Keywords:

Attention
Cognitive control
Basal ganglia
Switching
MVPA
fMRI
SVM

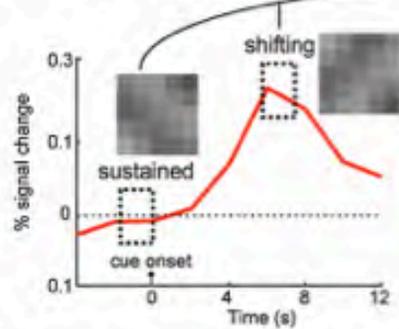
ABSTRACT

The posterior parietal cortex, including the medial superior parietal lobule (mSPL), becomes transiently more active during acts of cognitive control in a wide range of domains, including shifts of spatial and nonspatial visual attention, shifts between working memory representations, and shifts between categorization rules. Furthermore, spatial patterns of activity *within* mSPL, identified using multivoxel pattern analysis (MVPA), reliably distinguish between different acts of control. Here we describe a novel multivoxel pattern-based analysis that uses fluctuations in cognitive state over time to reveal inter-regional functional connectivity. First, we used MVPA to model patterns of activity in mSPL associated with shifting or maintaining spatial attention. We then computed a multivoxel pattern time course (MVPTC) that reflects, moment-by-moment, the degree to which the pattern of activity in mSPL more closely matches an attention-shift pattern or a sustained-attention pattern. We then entered the MVPTC as a regressor in a univariate (i.e., voxelwise) general linear model (GLM) to identify voxels whose BOLD activity covaried with the MVPTC. This analysis revealed several regions, including the striatum of the basal ganglia and bilateral middle frontal gyrus, whose activity was significantly correlated with the MVPTC in mSPL. For comparison, we also conducted a conventional functional connectivity analysis, entering the mean BOLD time course in mSPL as a regressor in a univariate GLM. The latter analysis revealed correlations in extensive regions of the frontal lobes but not in any subcortical area. The MVPTC analysis provides greater sensitivity (e.g., revealing the striatal-mSPL connectivity) and greater specificity (i.e., revealing more-focal clusters) than a conventional functional connectivity analysis. We discuss the broad applicability of MVPTC analysis to a variety of neuroimaging contexts.

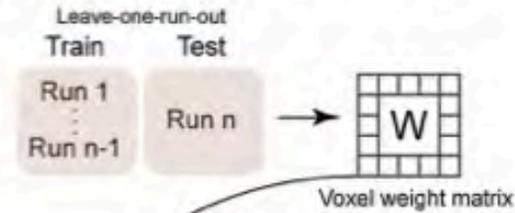
A. Define ROI: shift vs. hold contrast



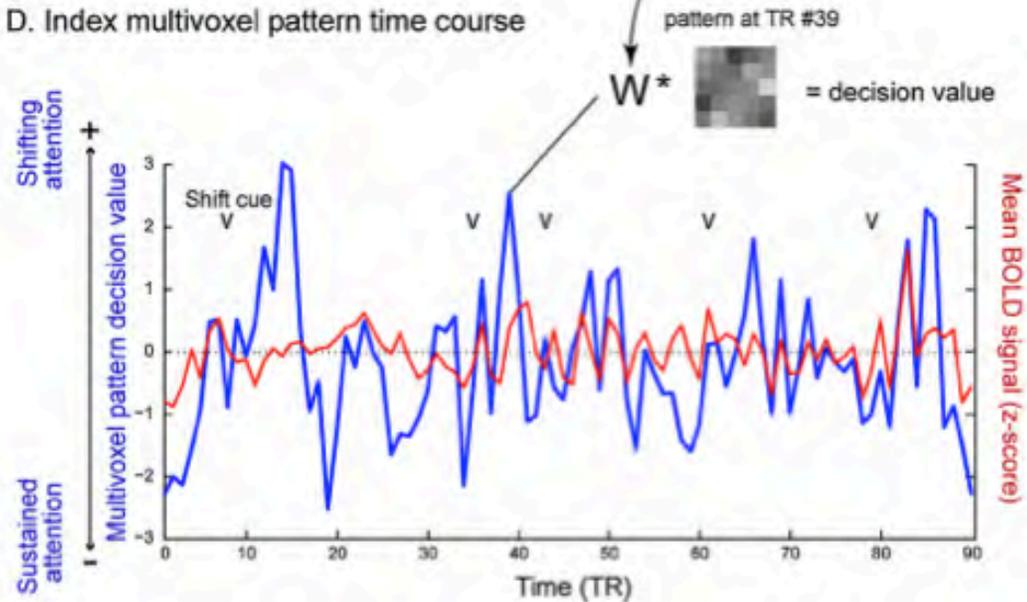
B. Extract multivoxel patterns from ROI



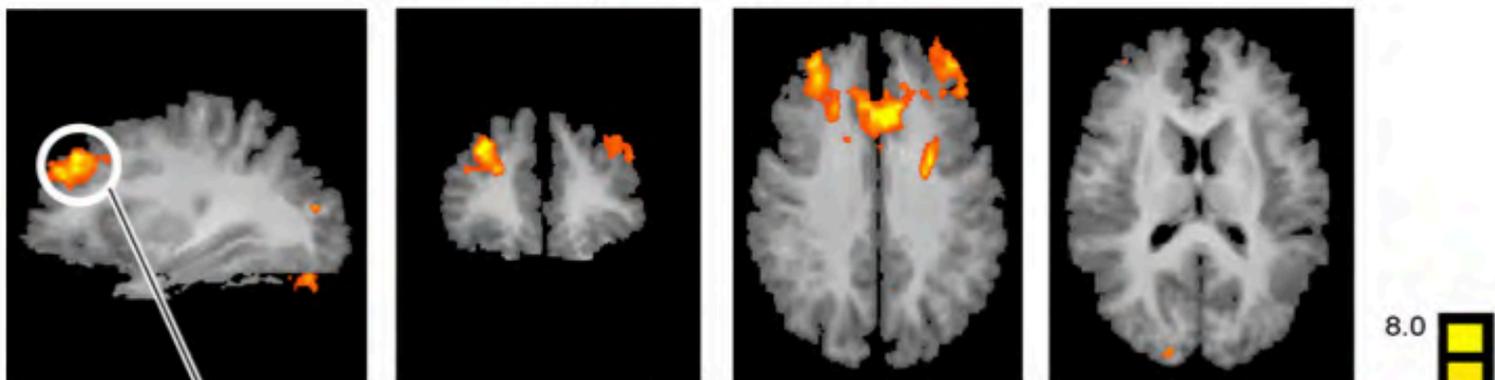
C. Classify multivoxel patterns using SVM



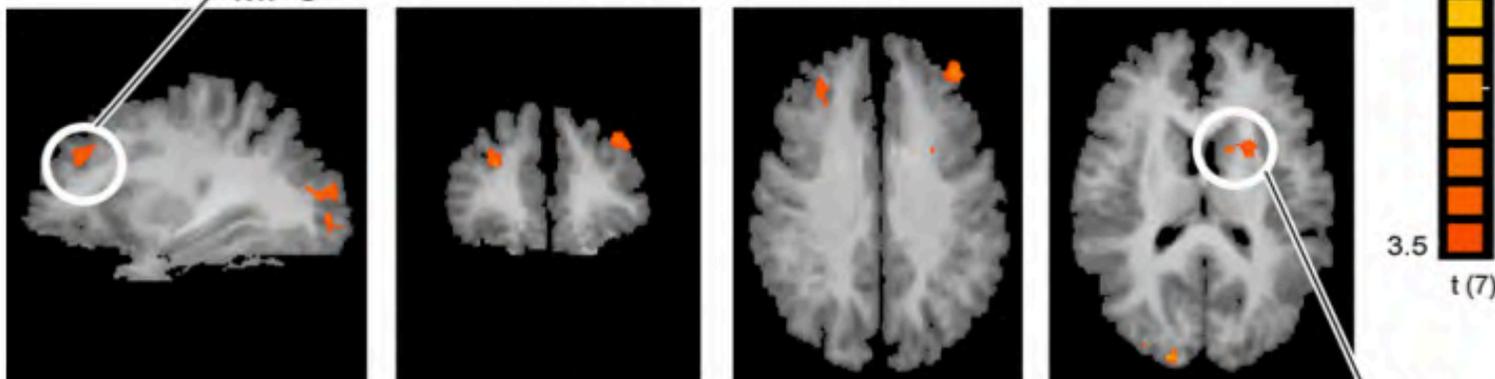
D. Index multivoxel pattern time course



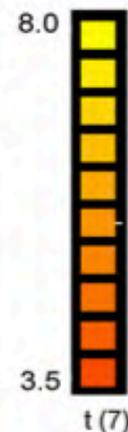
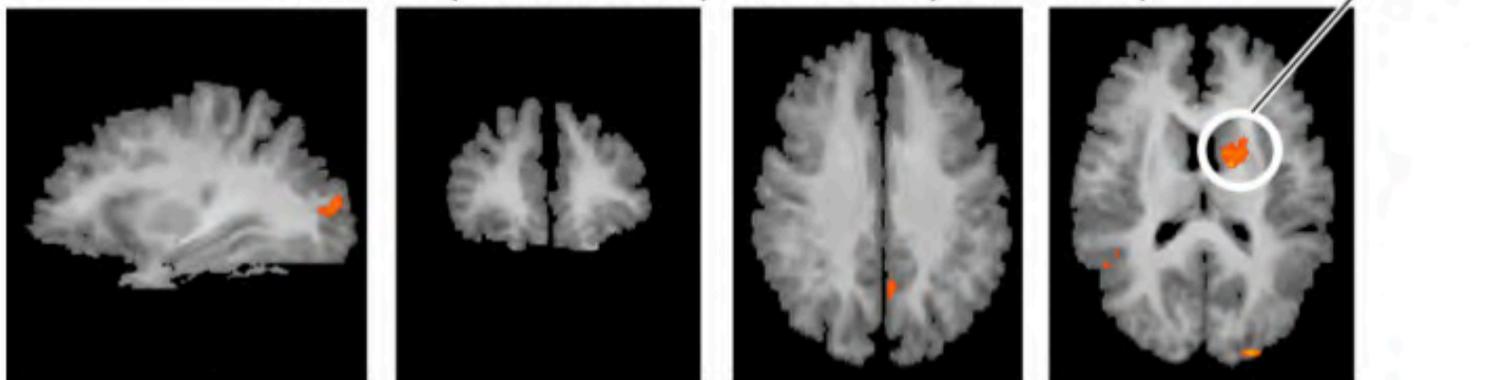
A. mSPL-mean model



B. mSPL-pattern model



C. mSPL-pattern model (mSPL-mean partialled out)



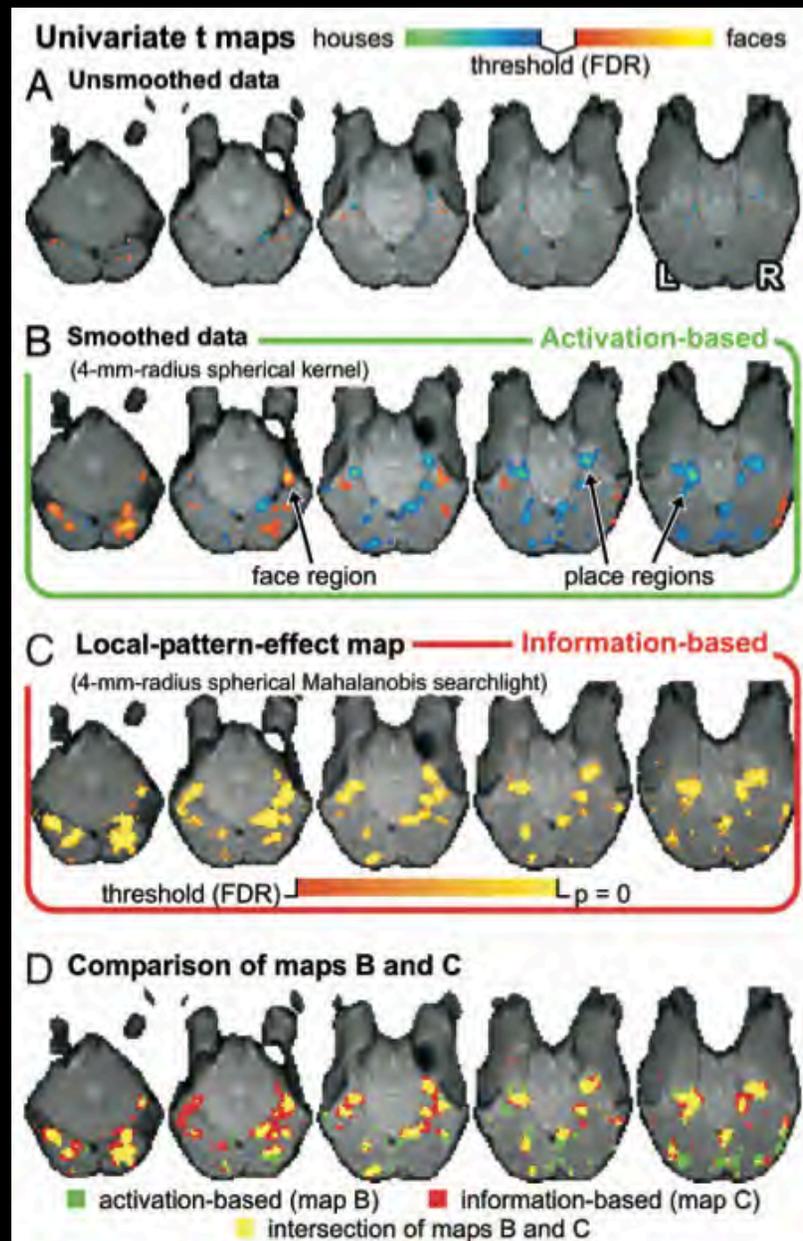
X = 30

Y = 36

Z = 30

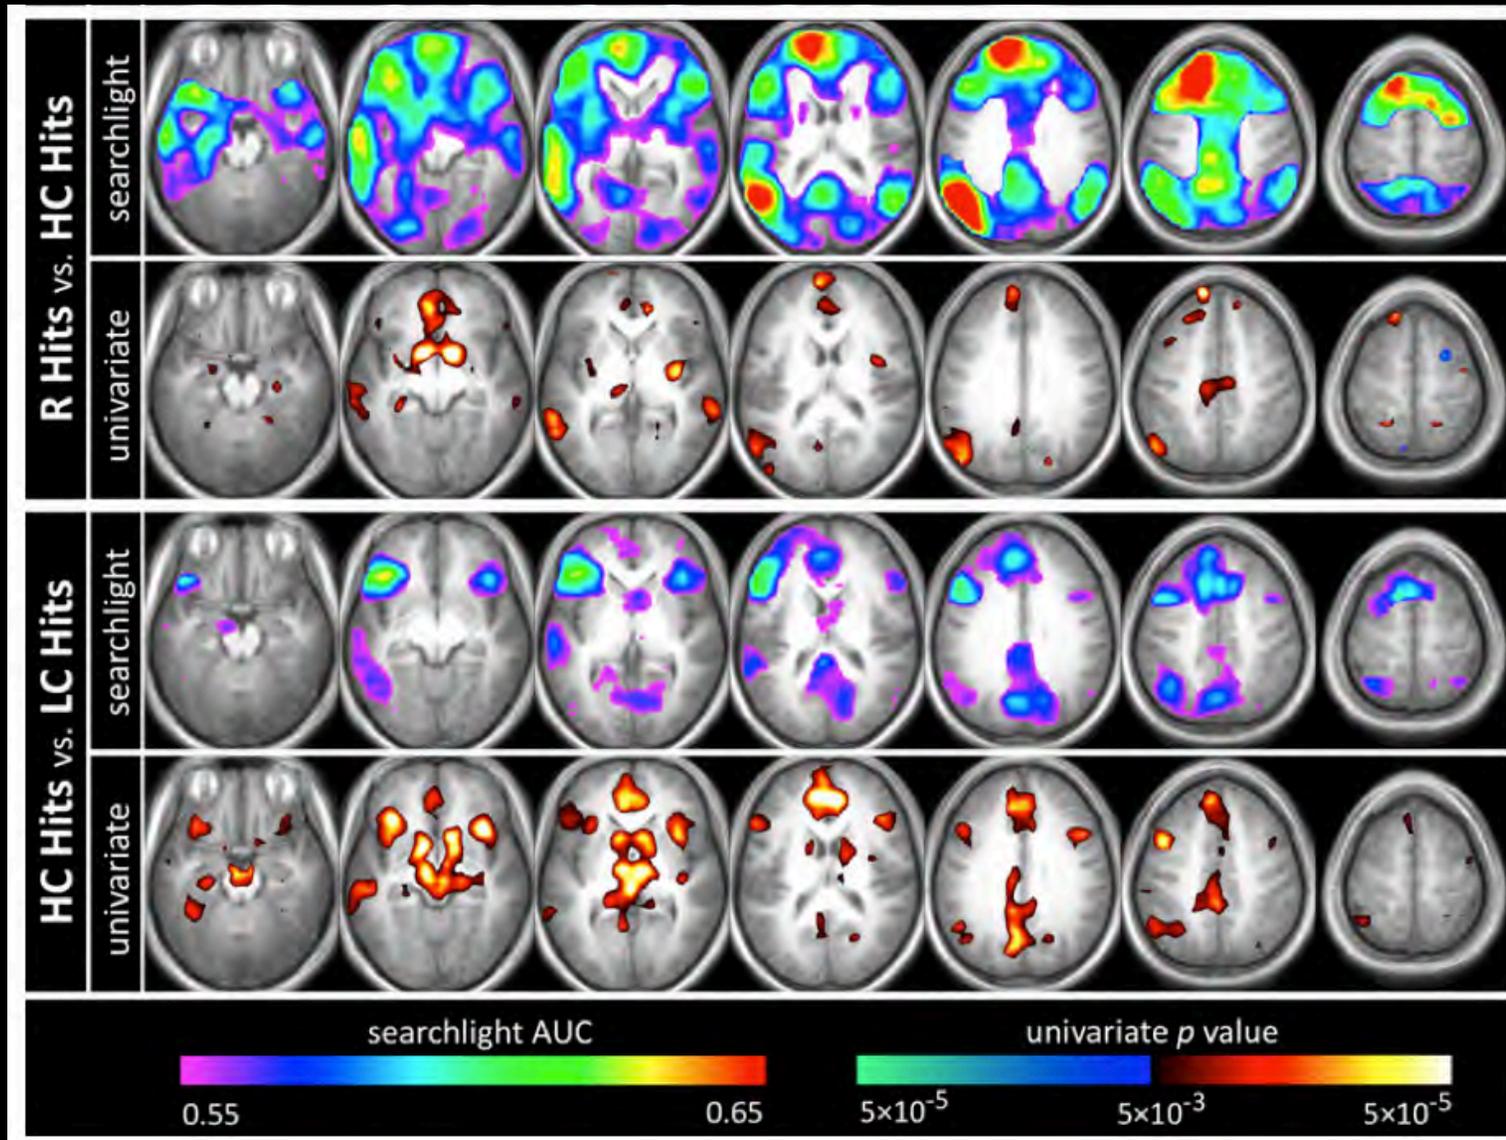
Z = 15

Univariate Mapping vs. MVPA Searchlight Mapping



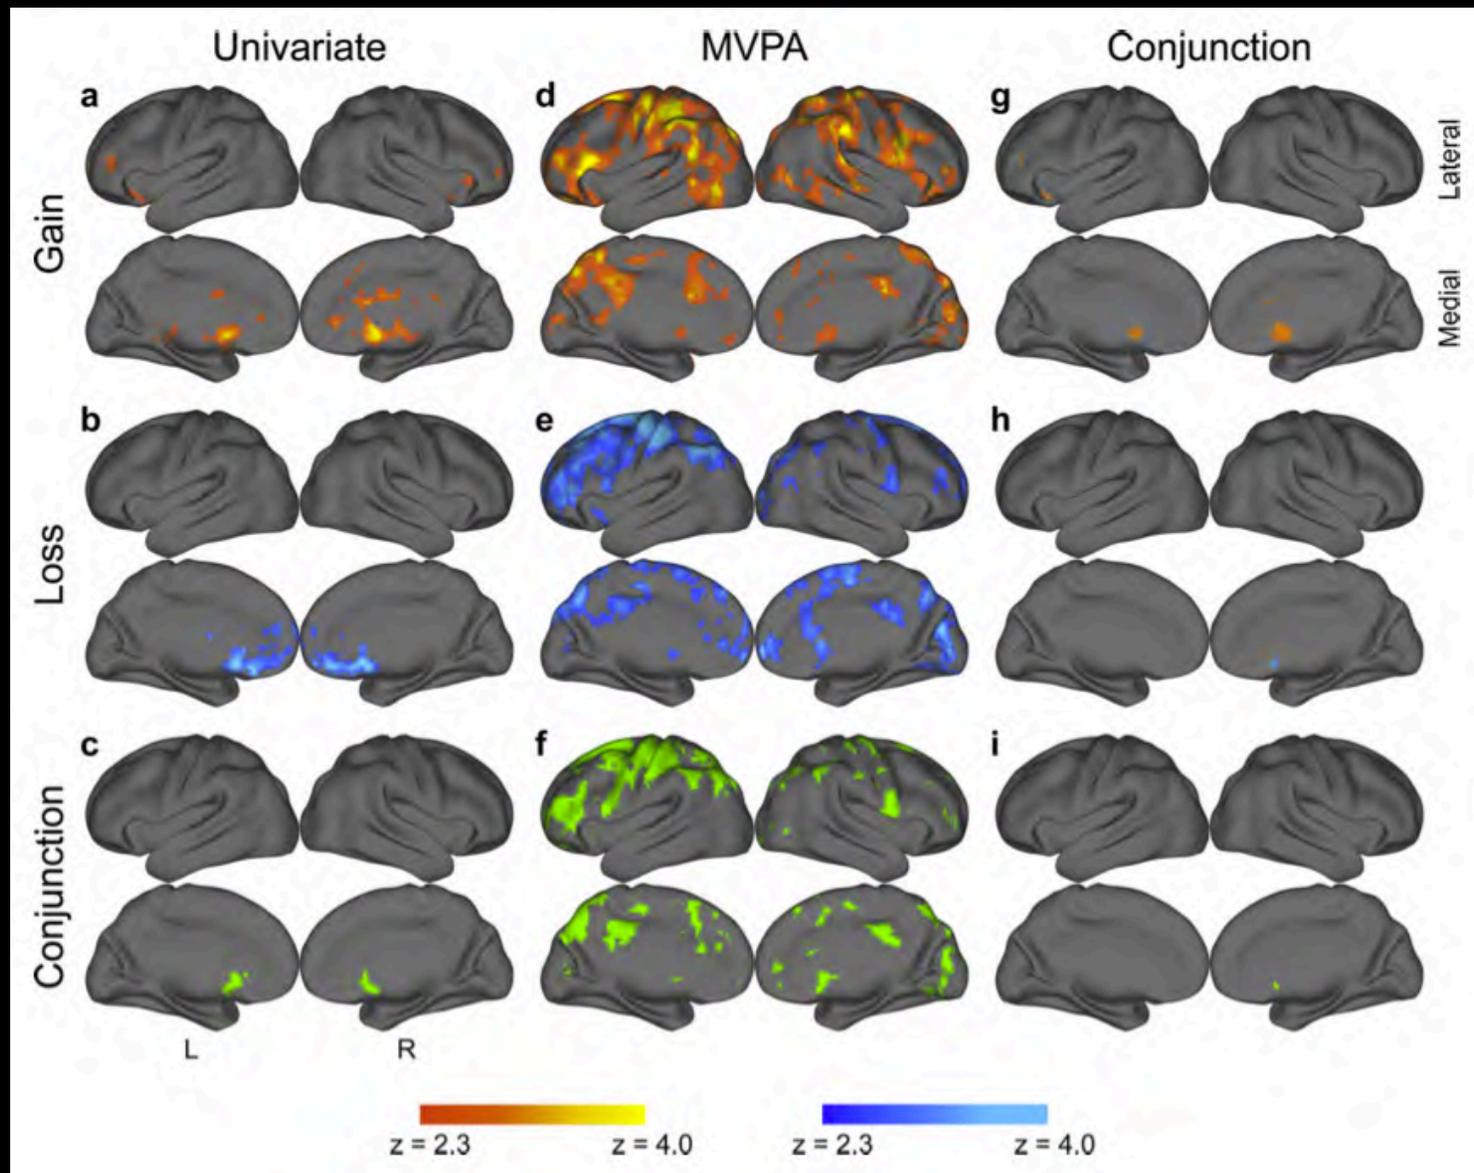
Kriegeskorte et al.
(2006), *PNAS*

Univariate Mapping vs. MVPA Searchlight Mapping



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

Univariate Mapping vs. MVPA Searchlight Mapping



Jimura & Poldrack (2012) *Neuroimage*

The agony and ecstasy of searchlight MVPA

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.

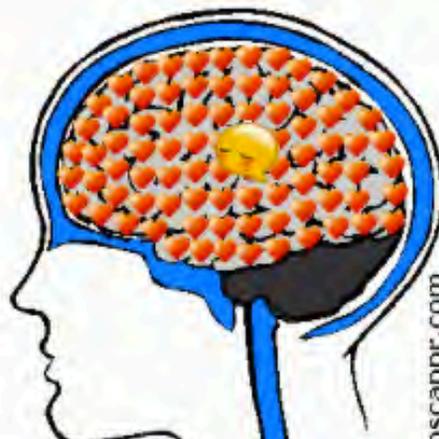
Thursday, July 4, 2013

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



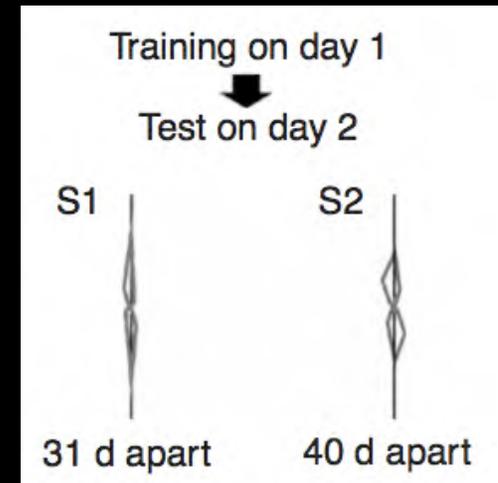
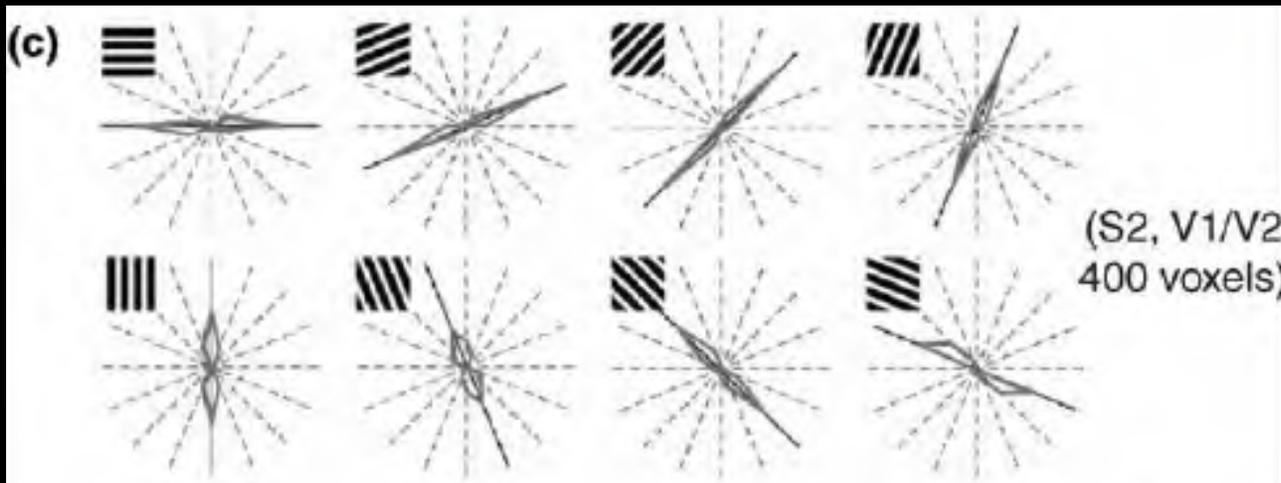
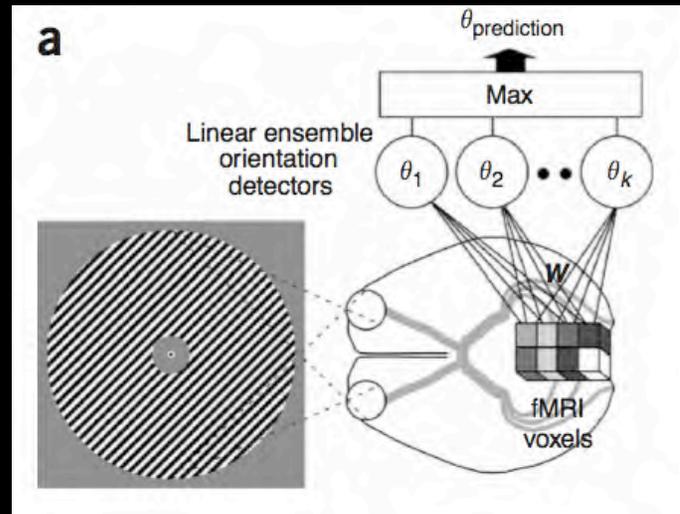
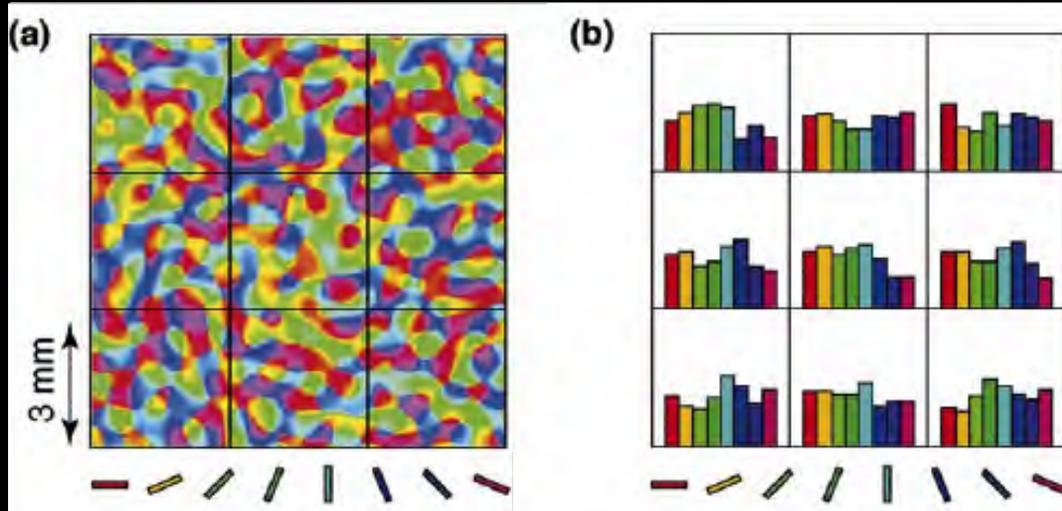
The Neurocritic

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 and she abandoned him across America at the age of thirteen

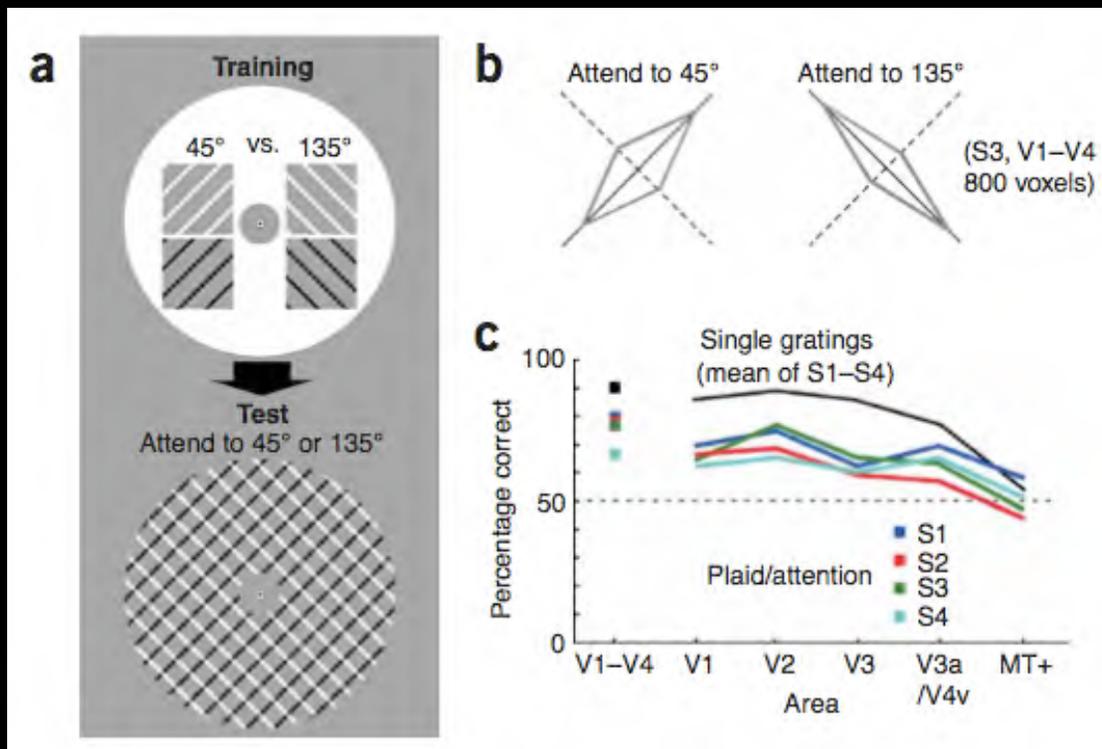
A few examples of studies where
classification analyses capture effects
that could not be observed via
univariate treatment of the data...

Decoding Visual Orientation



Kamitani & Tong (2005), *Nature Neuroscience*

Decoding Visual Attention



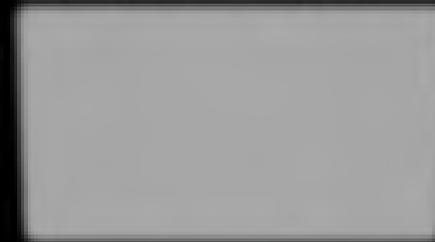
Kamitani & Tong (2005), *Nature Neuroscience*

Decoding the Contents of Working Memory

SEEING

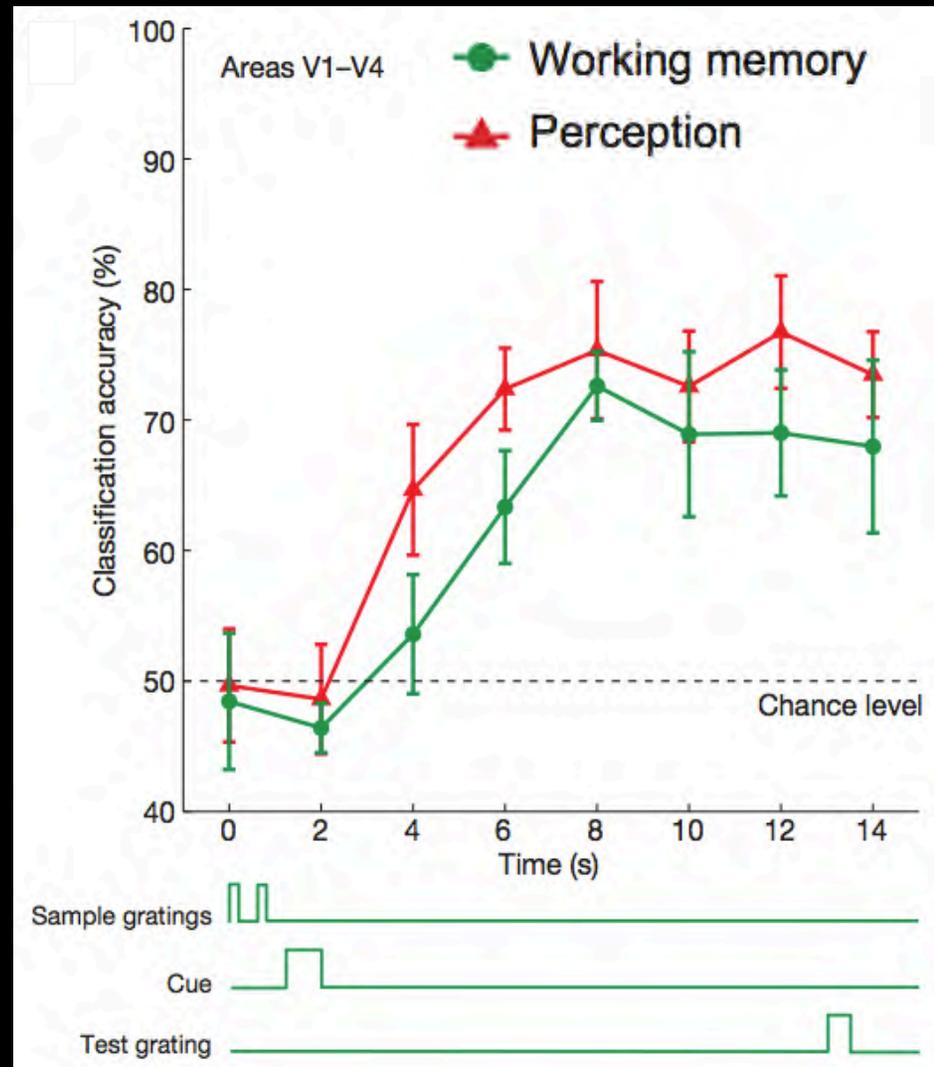
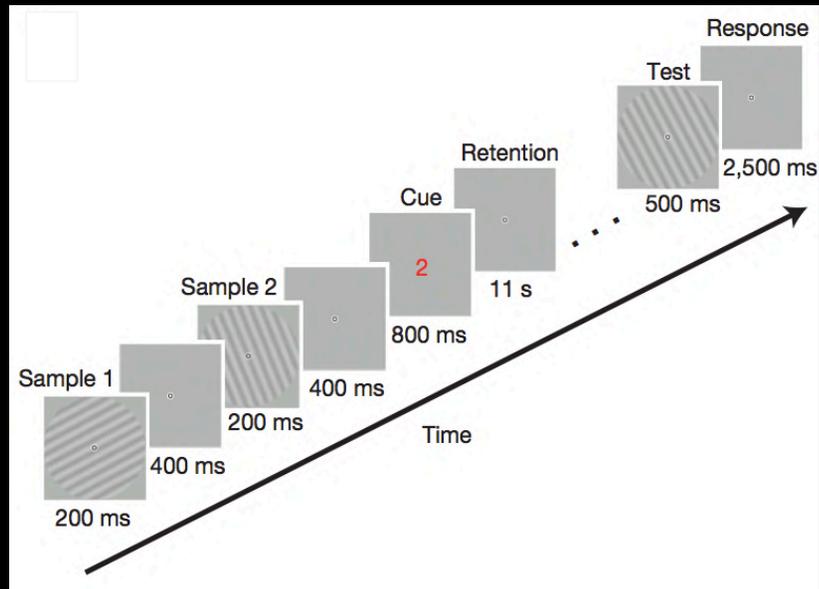


REMEMBERING



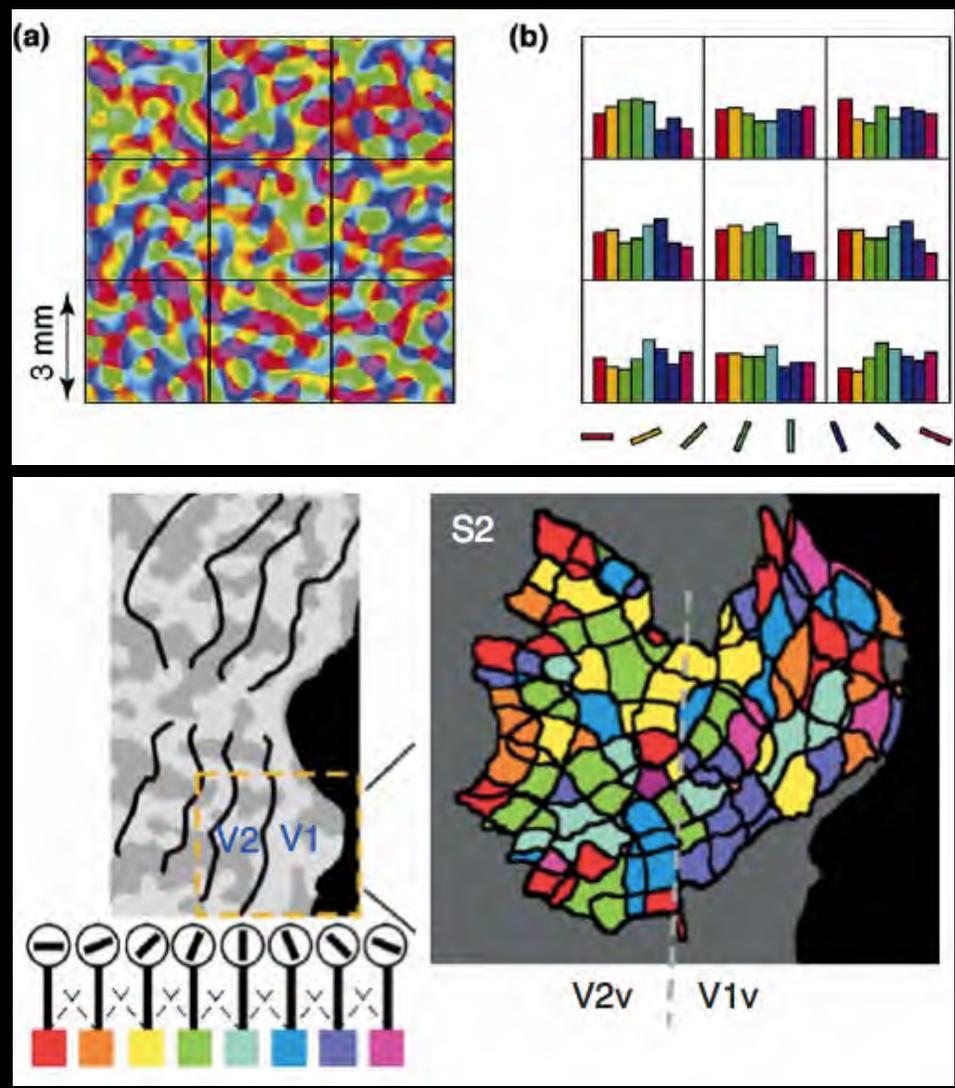
Harrison & Tong (2009) *Nature*

Decoding the Contents of Working Memory

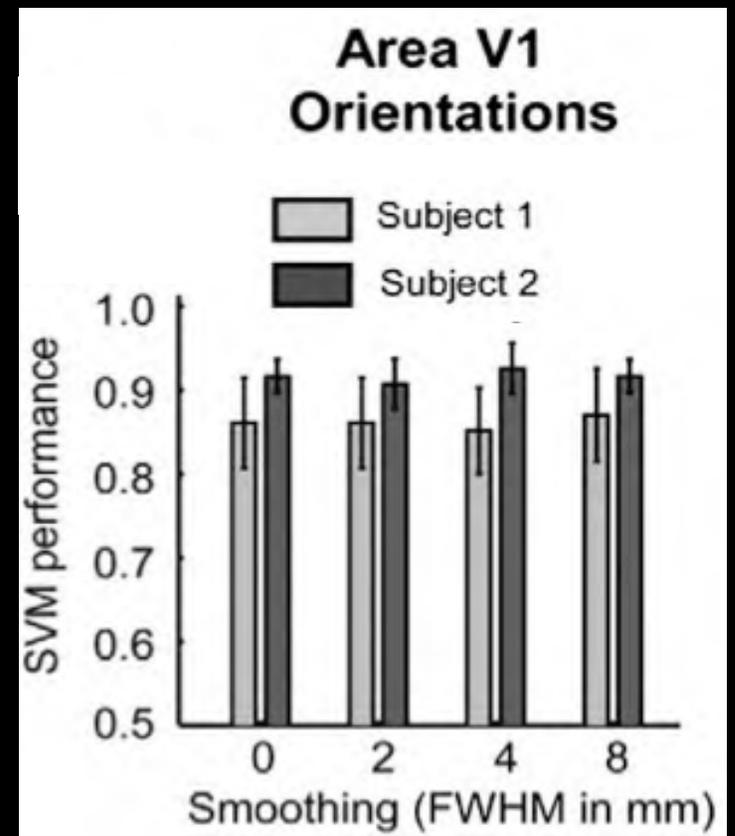


Harrison & Tong (2009) *Nature*

Do MVPA decoding techniques really capitalize on small-scale functional organization?



Kamitani & Tong (2005) *Nature Neuroscience*



Op de Beeck (2009) *NeuroImage*

Behavioral/Systems/Cognitive

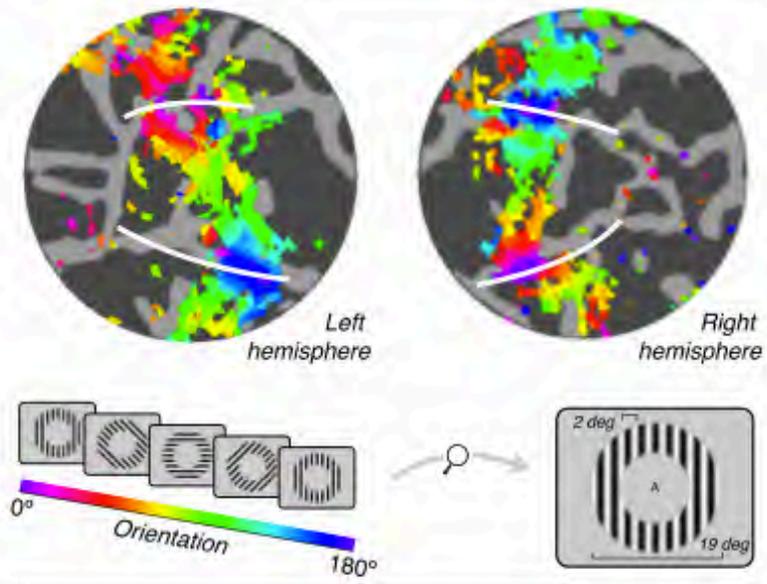
Orientation Decoding Depends on Maps, Not Columns

Jeremy Freeman,¹ Gijs Joost Brouwer,^{1,2} David J. Heeger,^{1,2} and Elisha P. Merriam^{1,2}

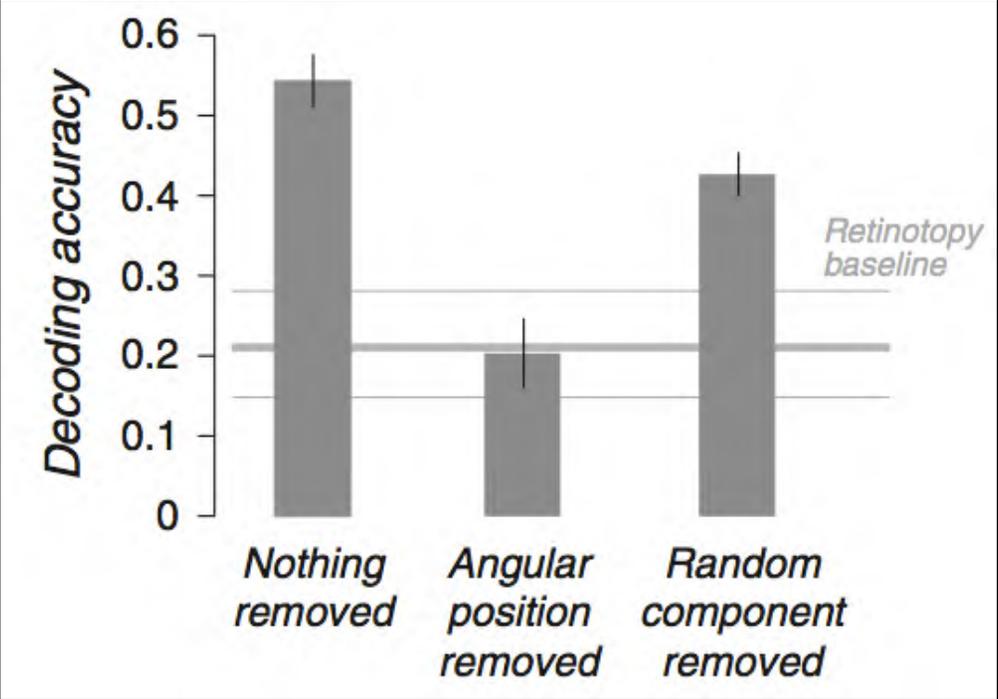
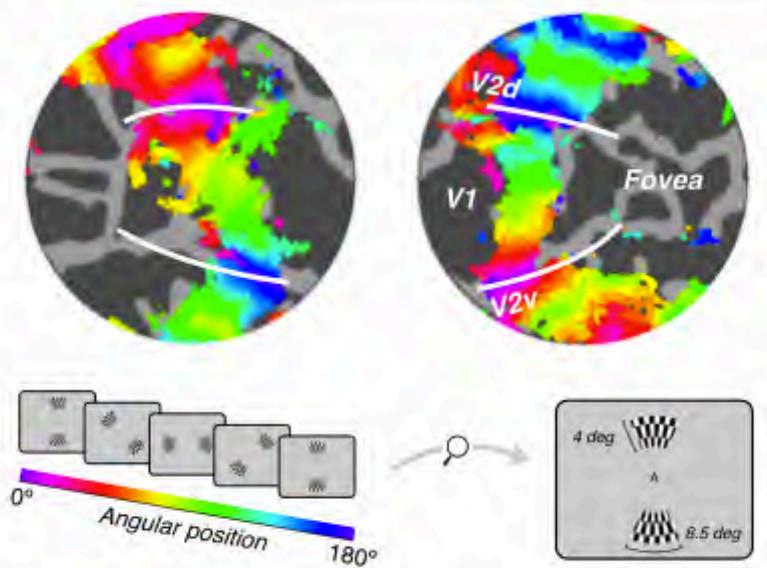
¹Center for Neural Science and ²Department of Psychology, New York University, New York, New York 10003

The representation of orientation in primary visual cortex (V1) has been examined at a fine spatial scale corresponding to the columnar architecture. We present functional magnetic resonance imaging (fMRI) measurements providing evidence for a topographic map of orientation preference in human V1 at a much coarser scale, in register with the angular-position component of the retinotopic map of V1. This coarse-scale orientation map provides a parsimonious explanation for why multivariate pattern analysis methods succeed in decoding stimulus orientation from fMRI measurements, challenging the widely held assumption that decoding results reflect sampling of spatial irregularities in the fine-scale columnar architecture. Decoding stimulus attributes and cognitive states from fMRI measurements has proven useful for a number of applications, but our results demonstrate that the interpretation cannot assume decoding reflects or exploits columnar organization.

A Orientation

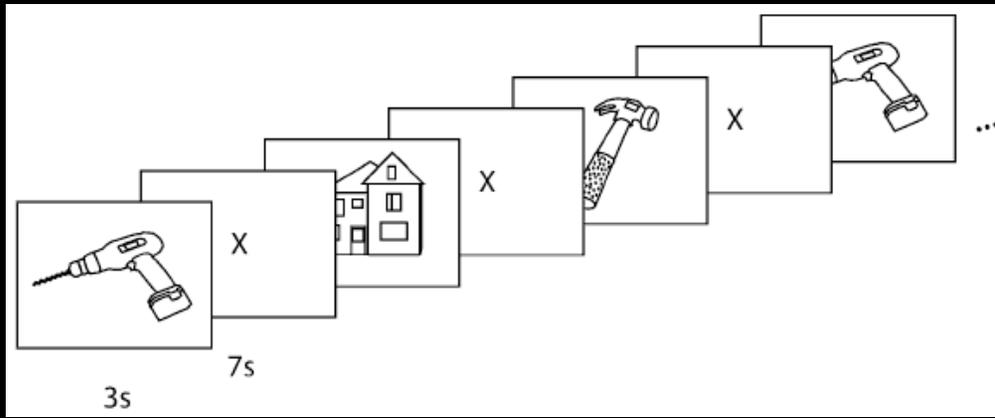


B Angular position



Freeman et al. (2011), *J Neuroscience*

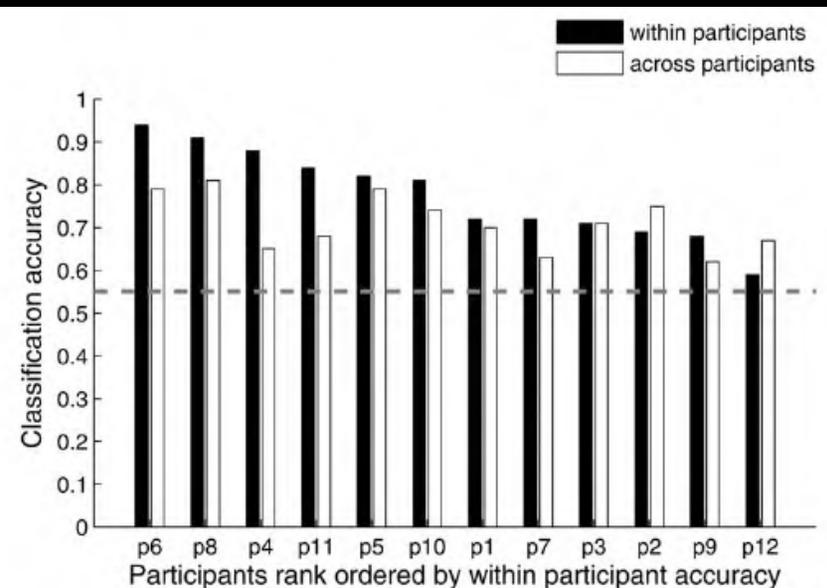
Classification Within and Across Subjects



- 10 objects total
- 5 tools (x6)
- 5 dwellings (x6)

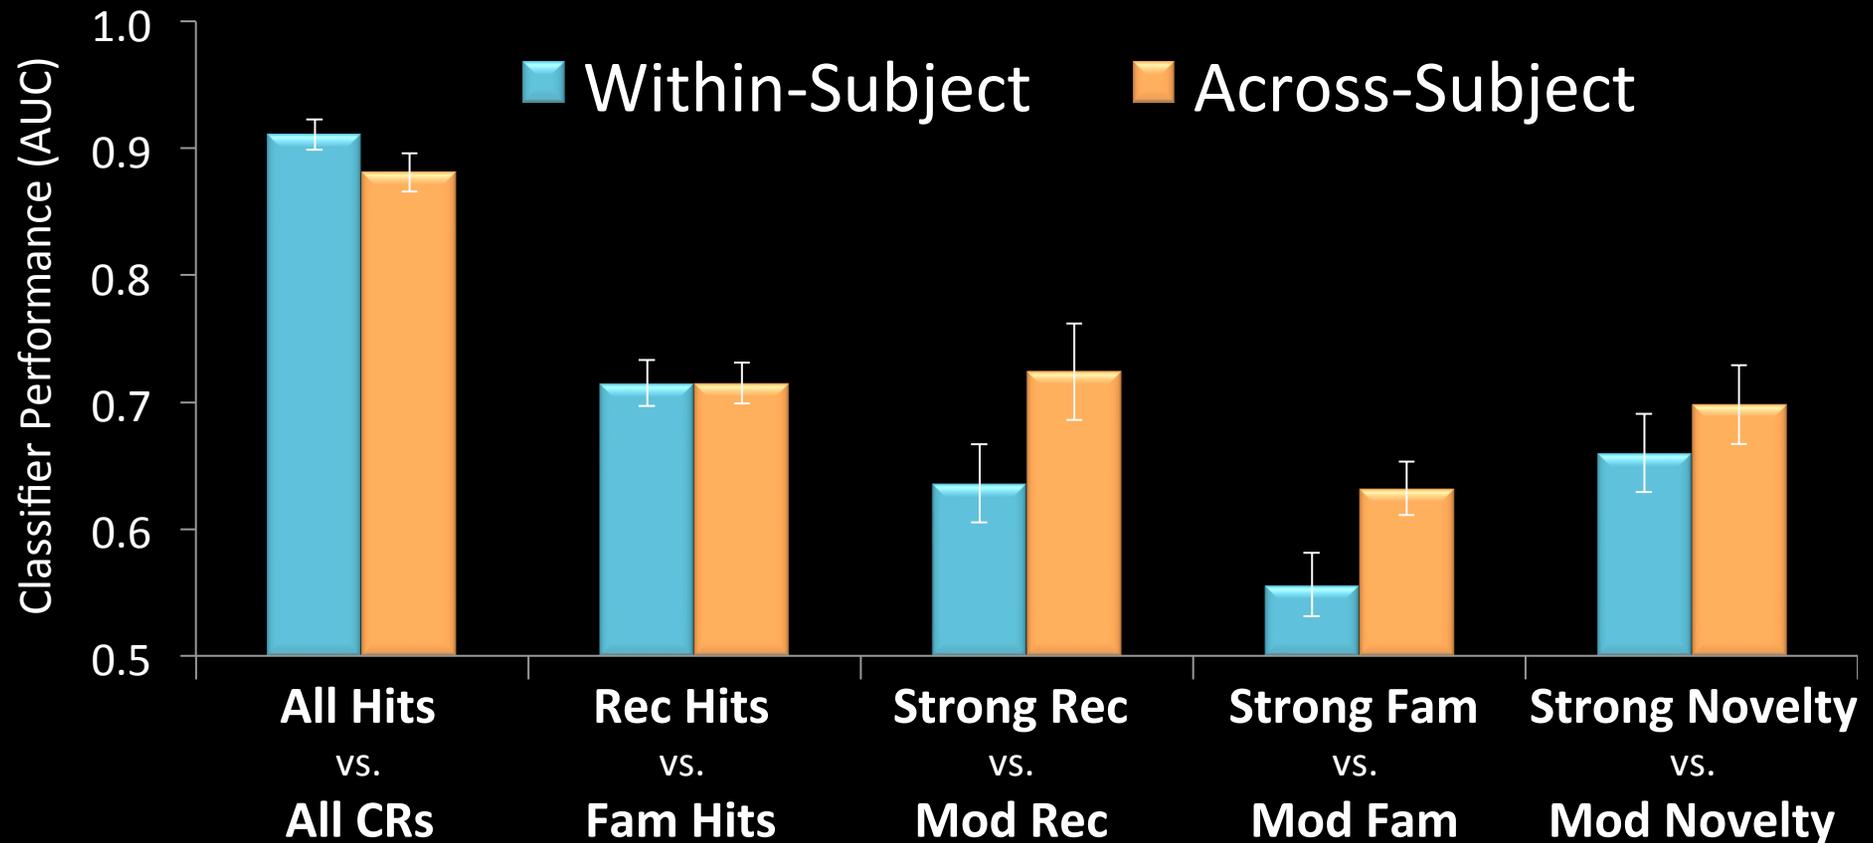
Classifier ranks likelihood that participant is viewing each of 10 objects

- visual regions EXCLUDED



Shinkareva et al. (2008), *PLoS*

Across-subject consistency of memory-related fMRI activity patterns



✓ Robust decoding performance even when classifier is trained and tested on different subjects

Across-experiment consistency of memory-related fMRI activity patterns

	TRAIN LAB-BASED ↓ TEST REAL-WORLD	TRAIN REAL-WORLD ↓ TEST LAB-BASED
All Hits vs. All CRs	67 %	76 %
Recollection vs. Familiarity	72 %	69 %

✓ Robust *across-experiment* classification performance

Generative Classifiers:

Predicting brain activity patterns for untrained exemplars

Predicting Human Brain Activity Associated with the Meanings of Nouns

Tom M. Mitchell,^{1*} Svetlana V. Shinkareva,² Andrew Carlson,¹ Kai-Min Chang,^{3,4}
Vicente L. Malave,⁵ Robert A. Mason,³ Marcel Adam Just³

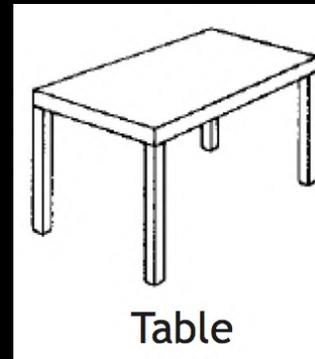
www.sciencemag.org SCIENCE VOL 320 30 MAY 2008

Generative Classifiers:

Predicting brain activity patterns for untrained exemplars

GOAL:

- understand how semantic information is represented in the pattern of activation

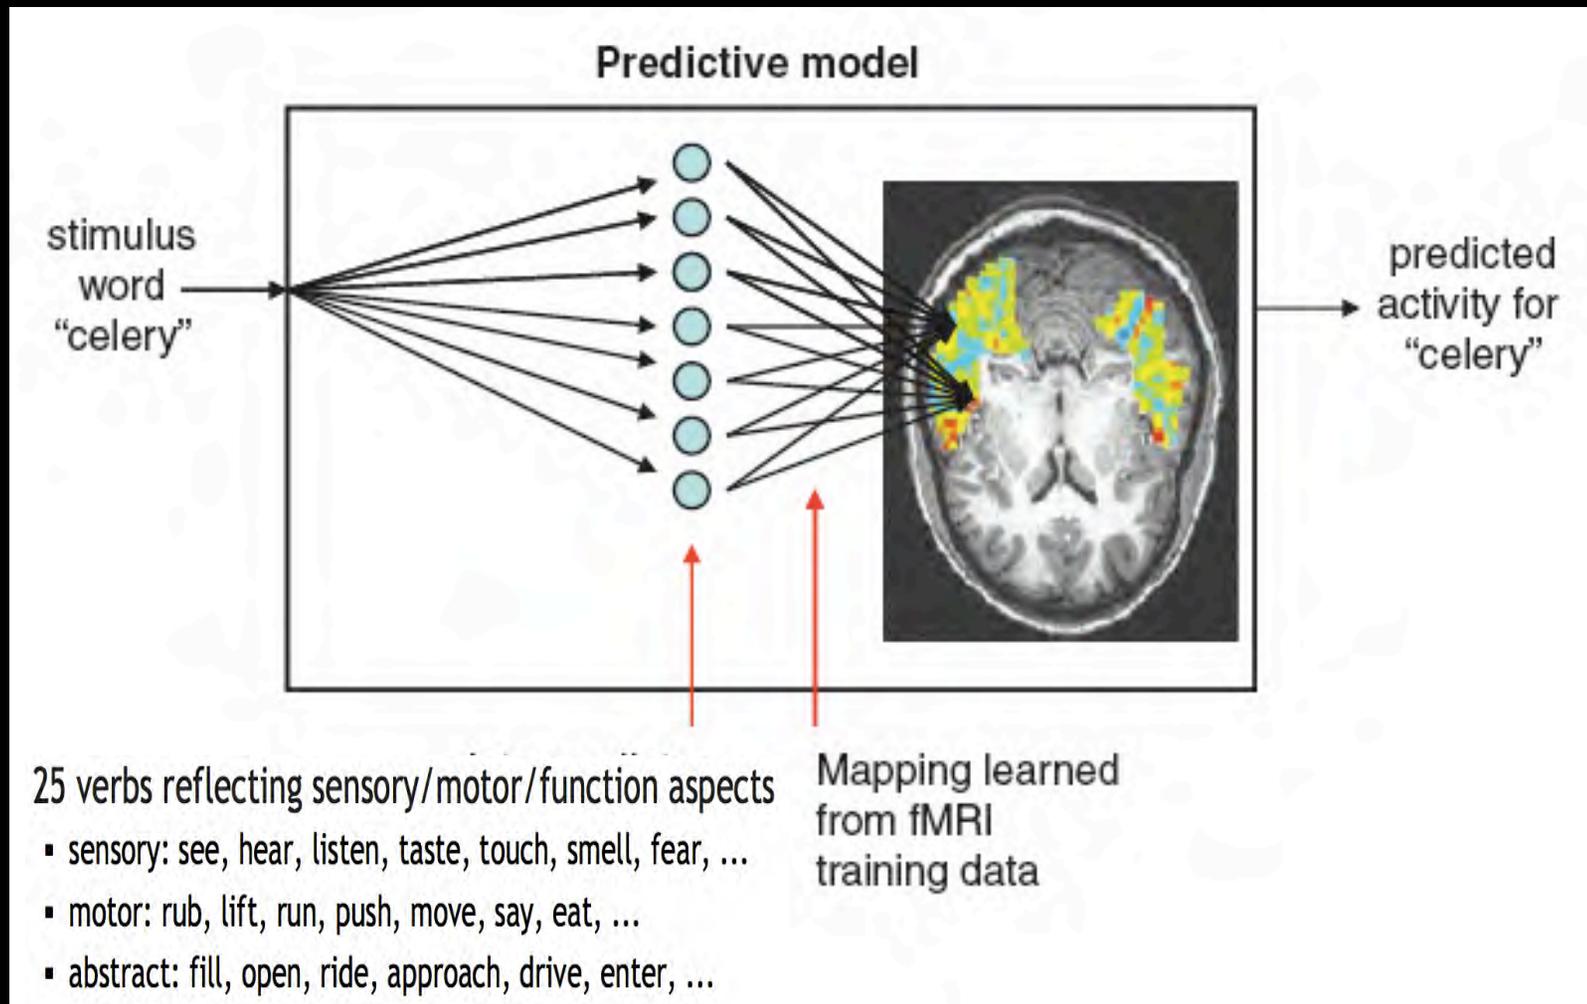


APPROACH:

- stimulus: word + drawing
- subject visualizes object, thinks of its properties, using it, etc.
- assume that each concept is represented in terms of features
- learn a mapping between each feature and the voxels it affects

Generative Classifiers:

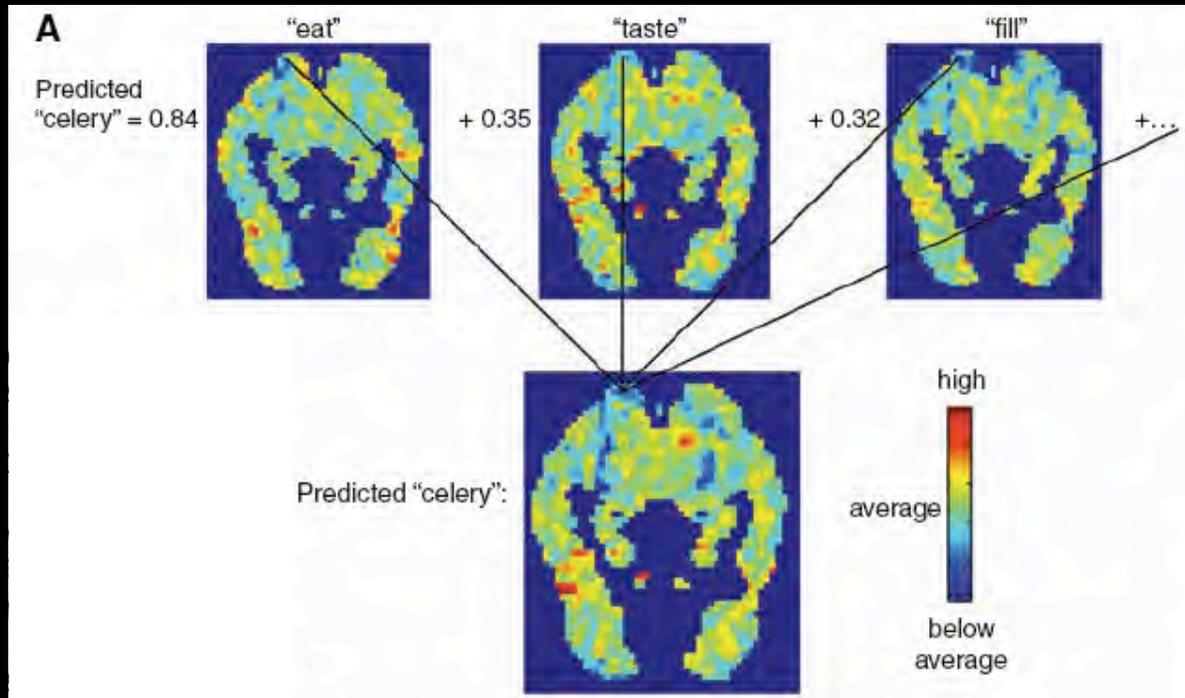
Predicting brain activity patterns for untrained exemplars



Mitchell et al. (2008), *Science*

Generative Classifiers:

Predicting brain activity patterns for untrained exemplars



60 nouns

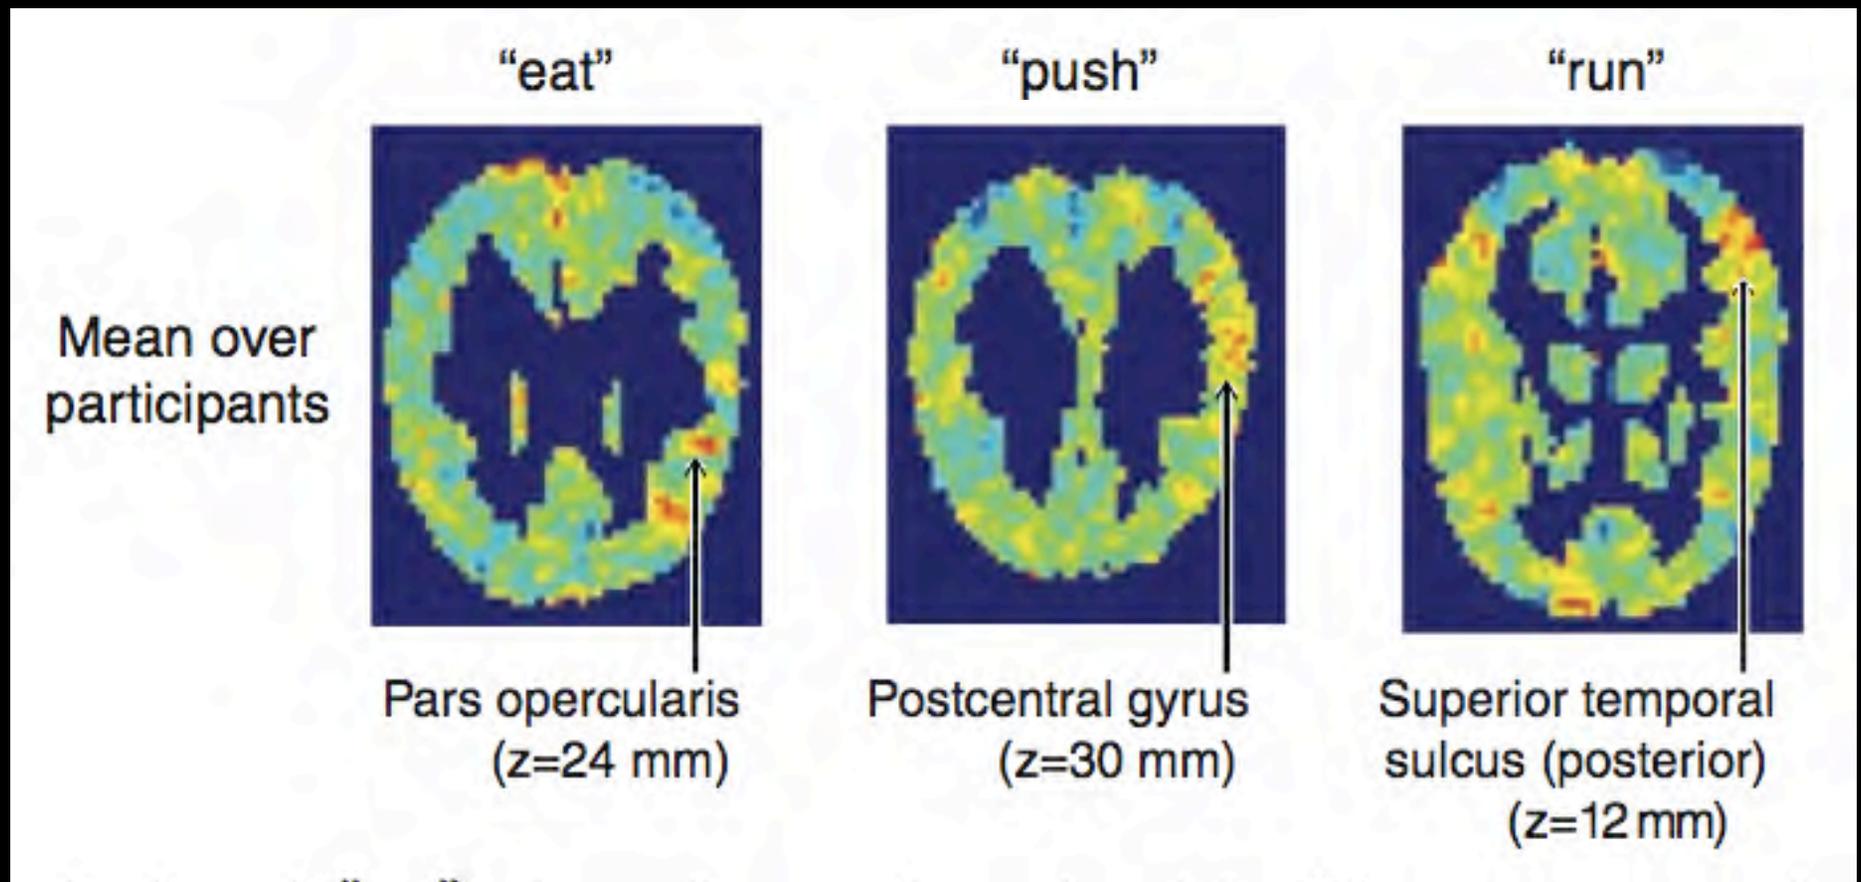
Mean accuracy = 0.77

Mitchell et al. (2008), *Science*

Generative Classifiers:

Predicting brain activity patterns for untrained exemplars

Interpreting the weight maps



Mitchell et al. (2008), *Science*

representational similarity analysis

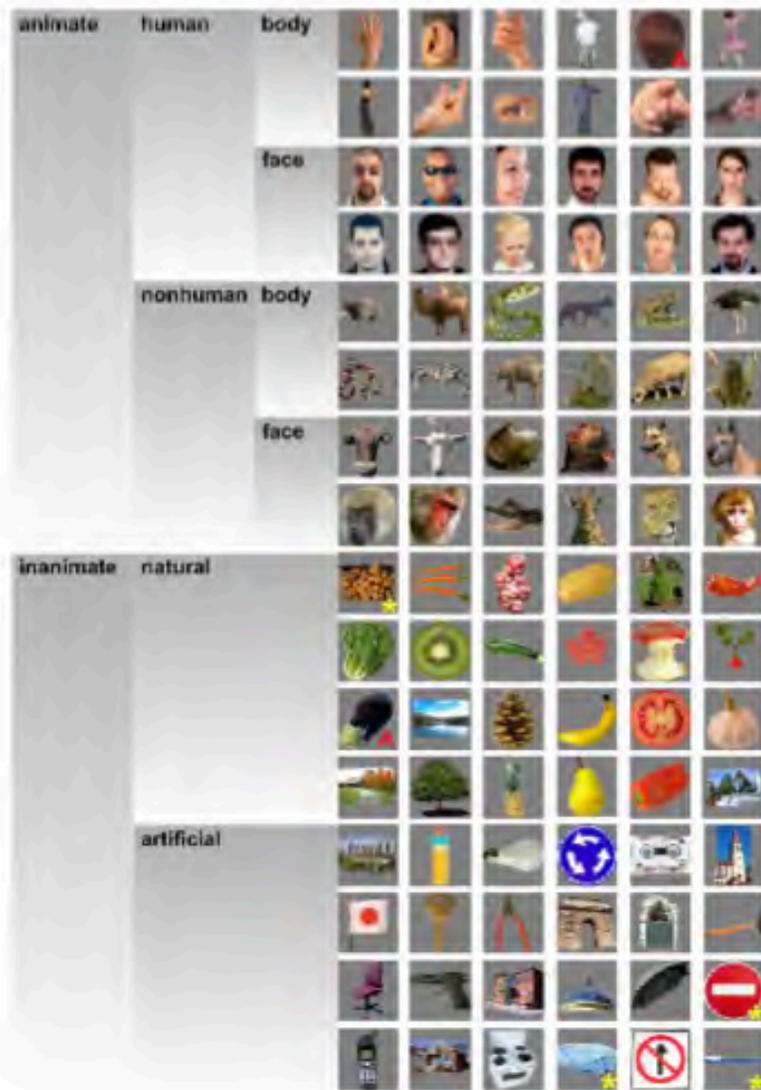
goal:

- 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



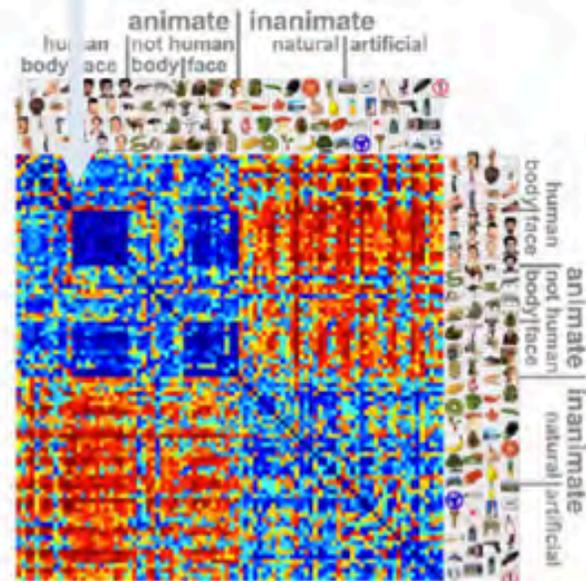
correlation



stimulus 1 pattern



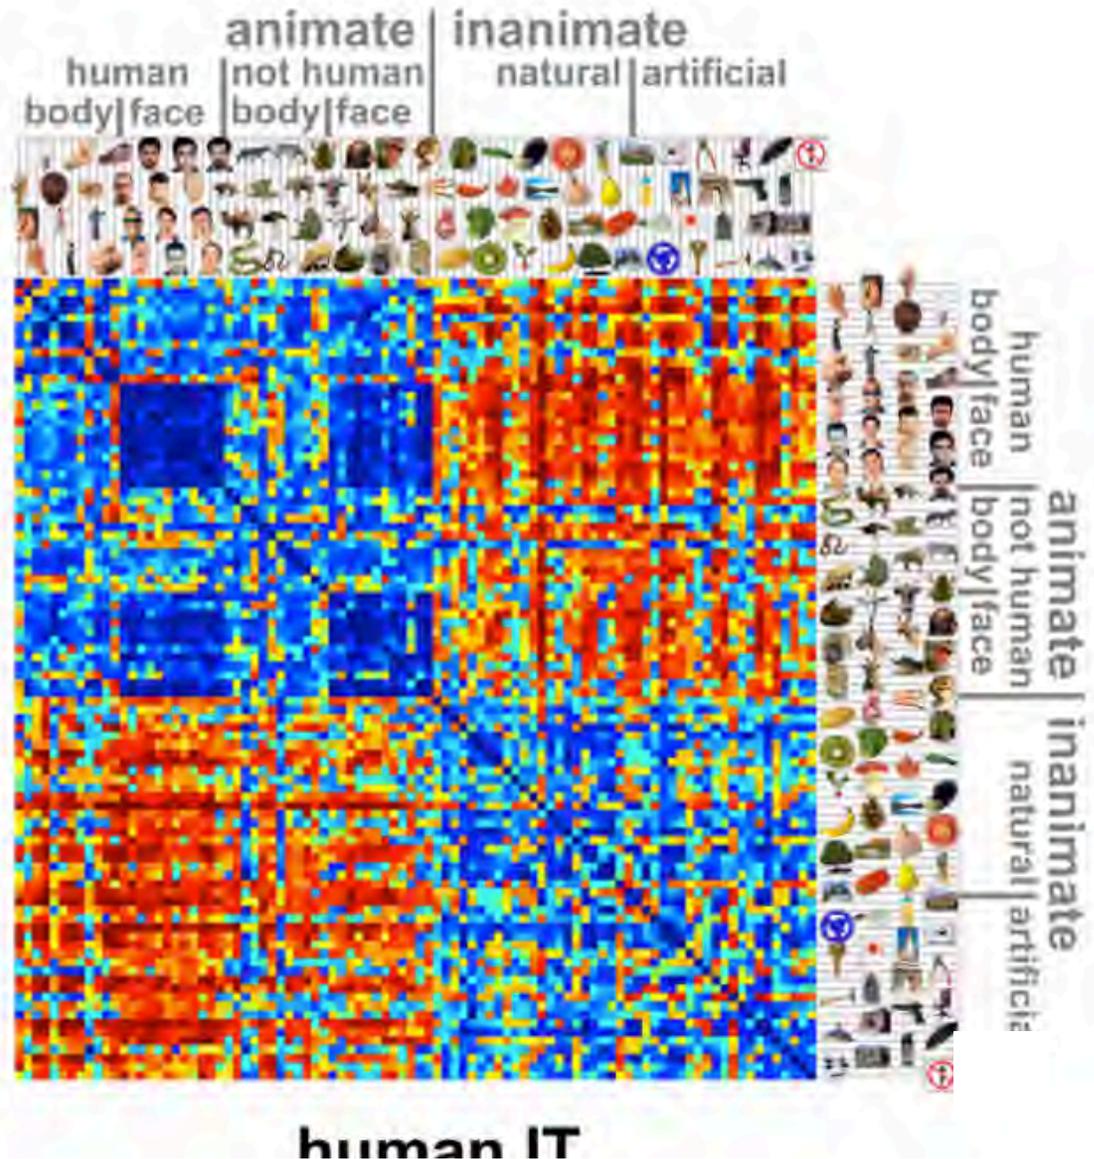
stimulus 2 pattern



human IT

stimuli and similarity

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



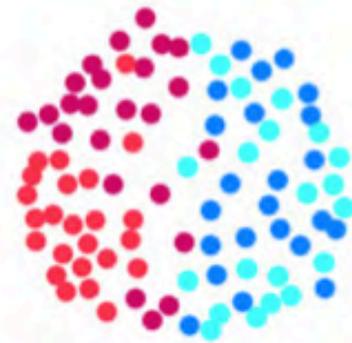
Kriegeskorte et al. (2008; 2009)

[slide from Francisco Pereira]

multidimensional scaling

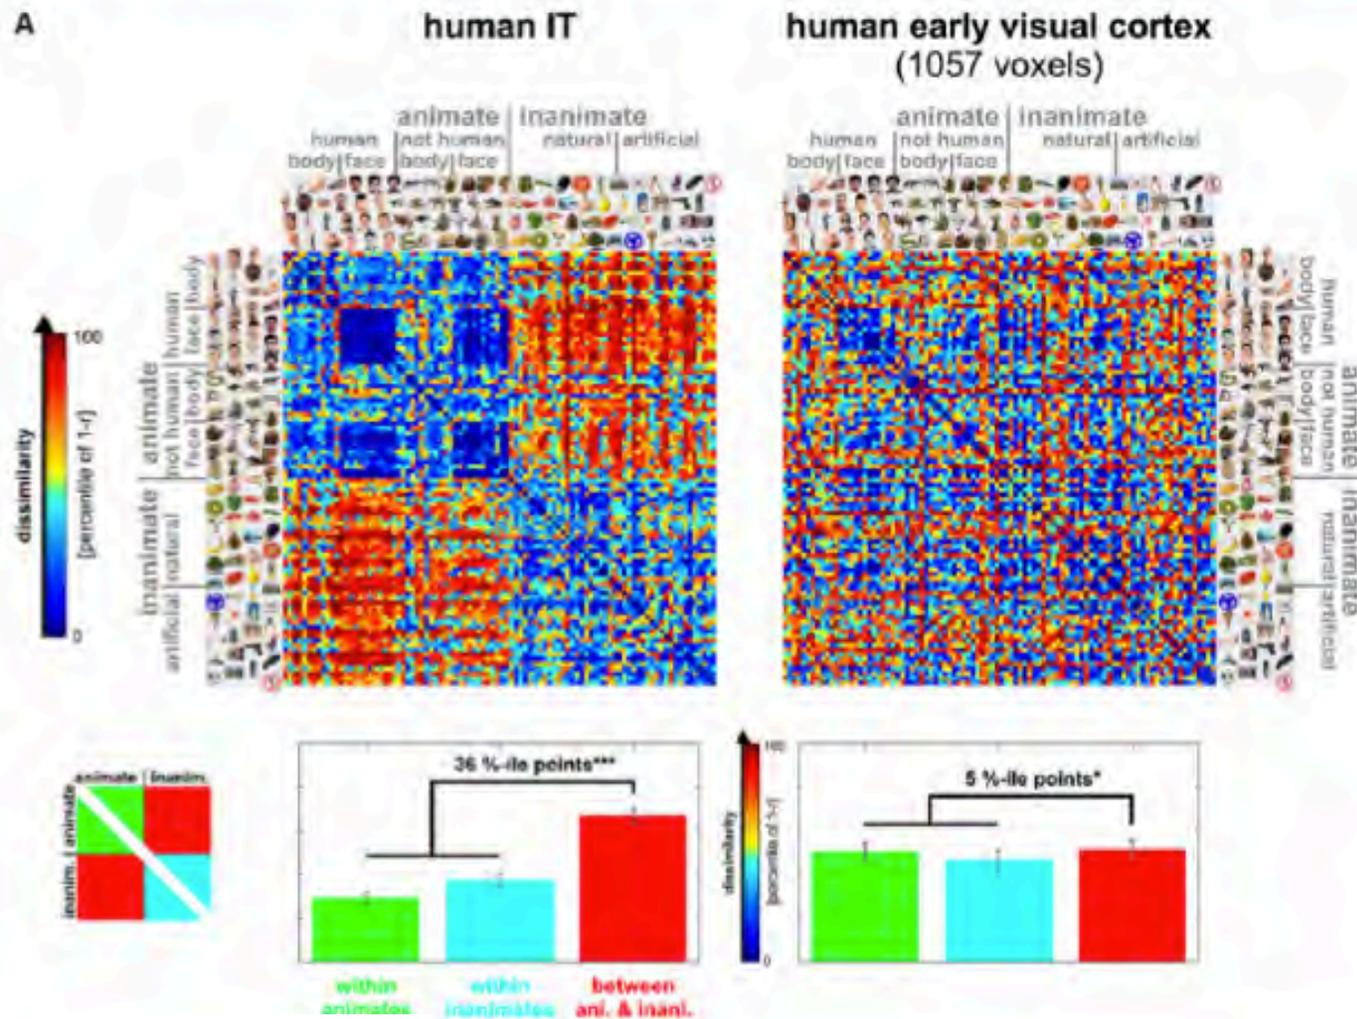


human IT



body
face
natural obj.
artificial obj.

compare multiple areas

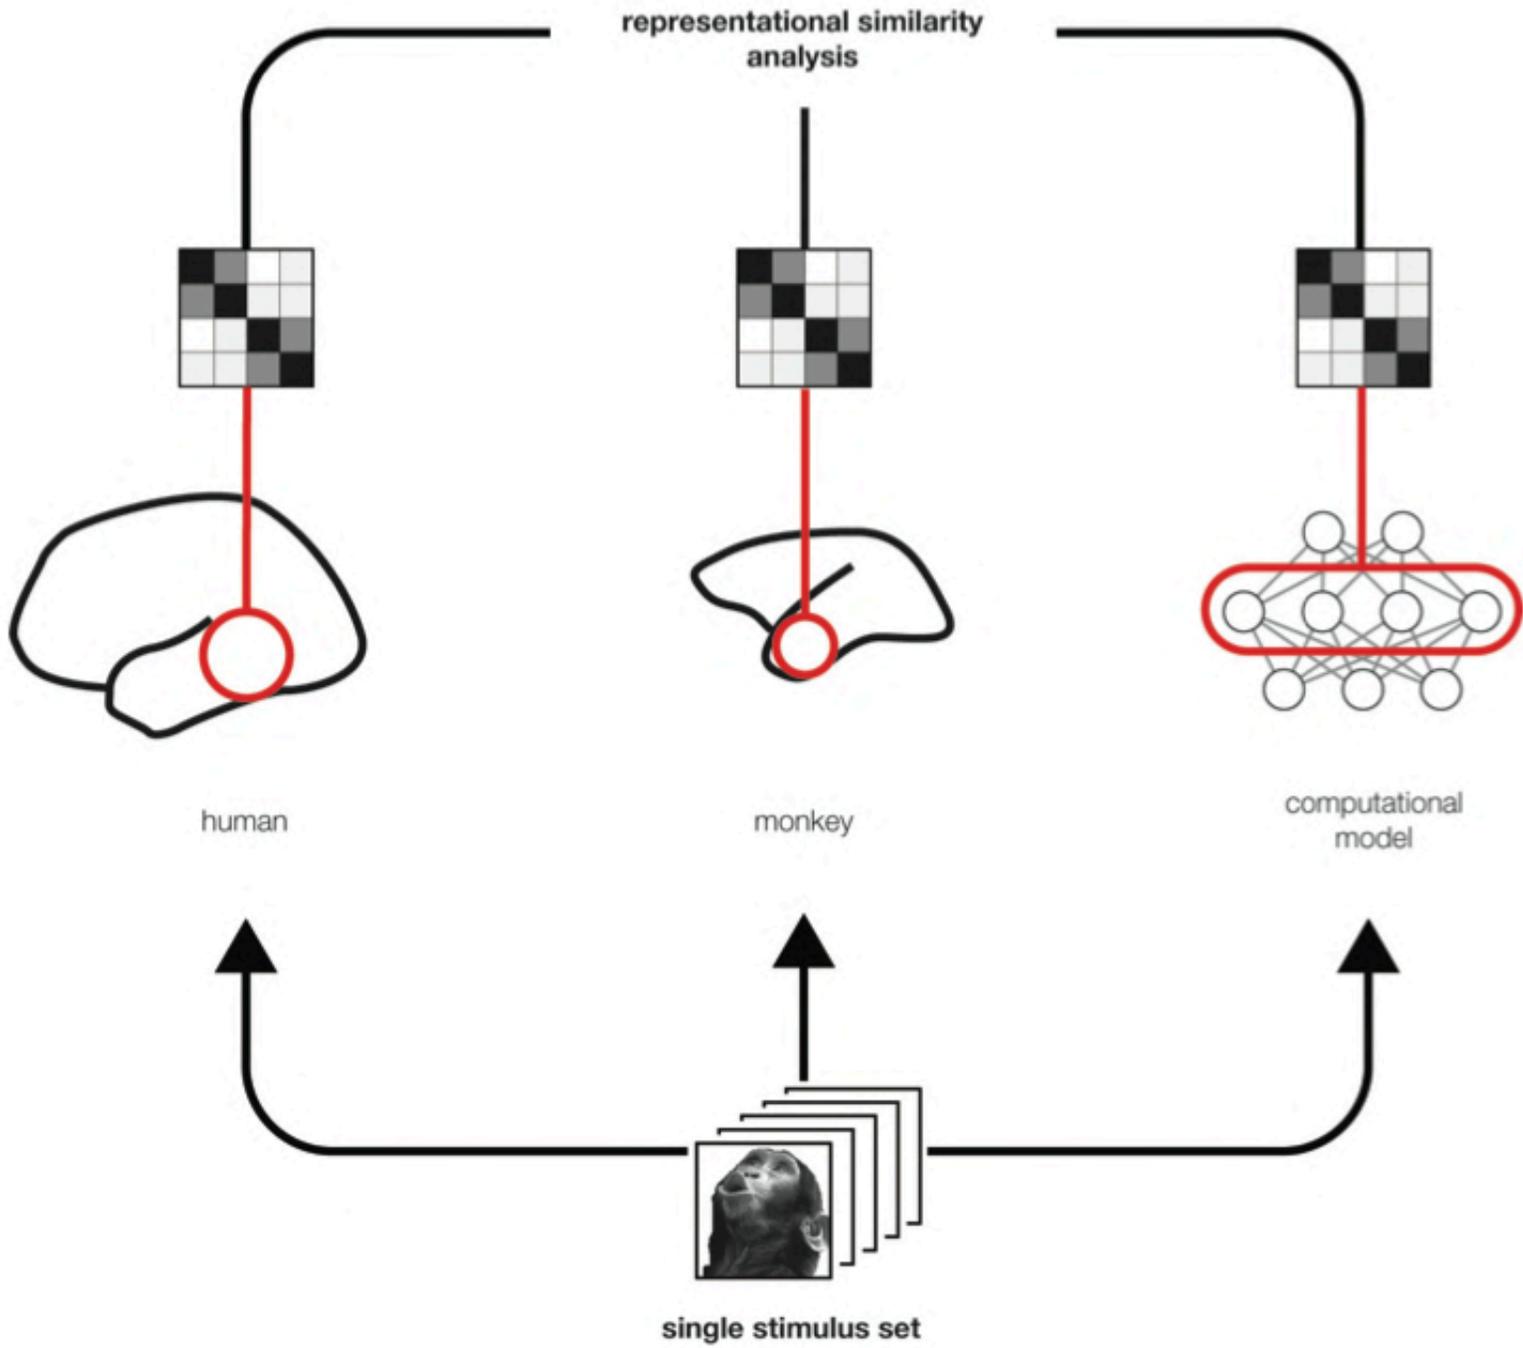


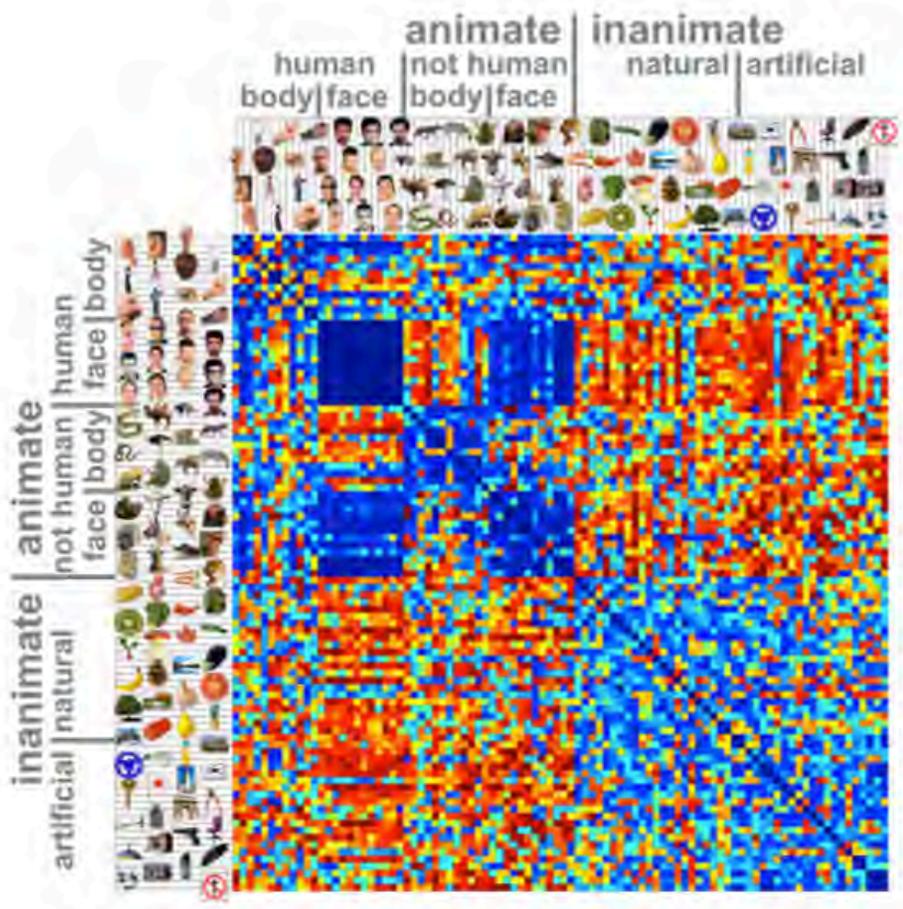
Kriegeskorte et al. (2008; 2009)

[slide from Francisco Pereira]

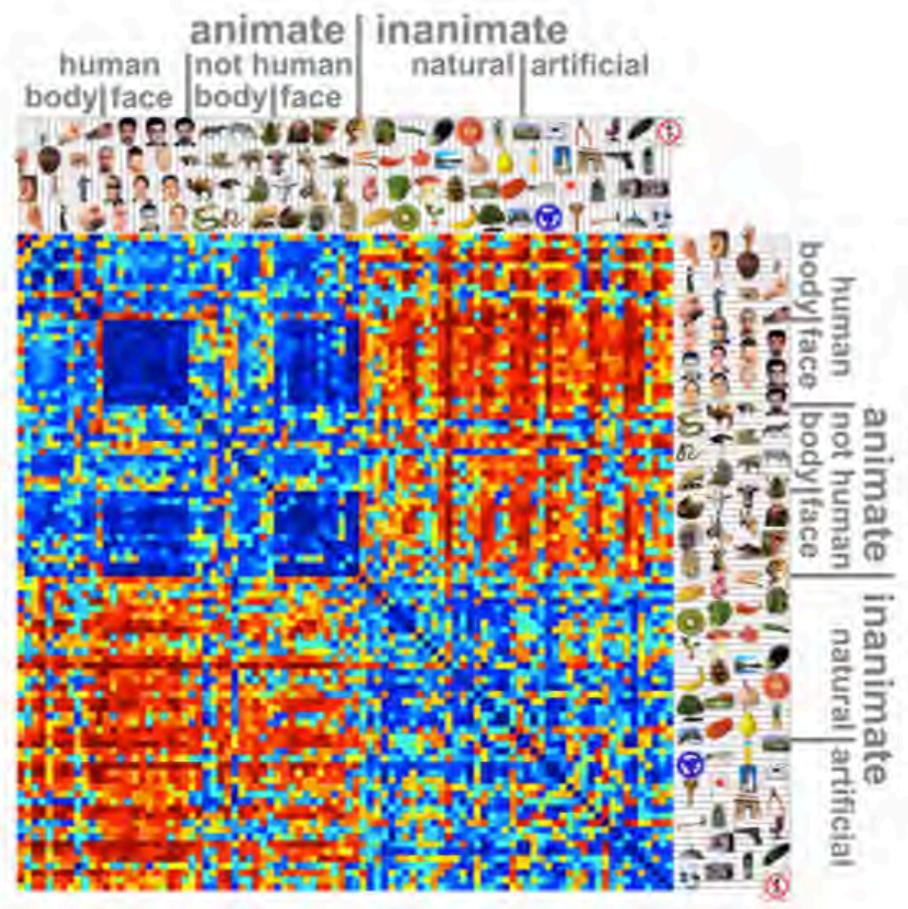
other aspects

- 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

