Design considerations for fMRI studies using multivariate pattern analysis

Samuel A. Nastase

fMRI Brown Bag

April 25, 2014

DARTMOUTH

VOX CLAMANTIS IN DESERTO

1769

Outline

General considerations
- Types of questions
- Univariate vs. multivariate analyses

Experimental design (with examples)
- Participants
- Stimuli
- Block designs
- Slow event-related designs
- Rapid event-related designs
- Condition-rich designs for RSA
- Task

Particular issues
- Order effects
- Jittering

Constraints
- Input to multivariate analyses
- Cross-validation

References and resources

General considerations

Two fundamentally different types of *questions*

Univariate analysis:

- Measures *involvement* of a region in a certain task or process
- Concerned with *localization*—i.e., "*where*"
- To extent does activation predict behavior?

Multivariate analysis:

- Measures the *representational content* via pattern information
- Concerned with information representation—i.e., *"how"*
- Reflects distributed, probabilistic neuronal population codes

Conventional univariate analyses can be seen as a special case of the multivariate analytic framework

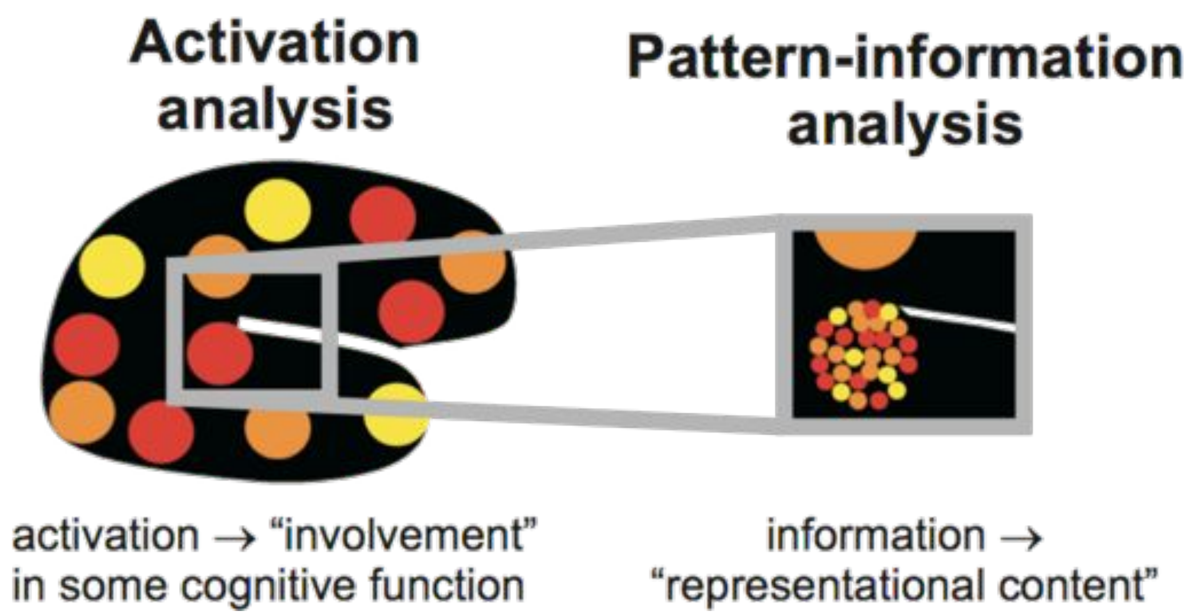

Raizada & Kriegeskorte, 2010

Why use multivariate analyses?

"Nothing is more multivariate than the brain" (Raizada & Kriegeskorte, 2010)
—i.e., *on principle*

Nothing in the brain operates independently or in isolation and the question of localization can only tell us so much

The multivariate framework is not new to psychological, cognitive, and neural science

- PCA, ICA, multidimensional scaling (e.g., Kruskal, 1964)
- Neural networks, connectionist models (e.g., Rumelhart & McClelland, 1986)
- Population coding in primate electrophysiology (e.g., Georgopoulos et al, 1986; Rolls & Tovee, 1995)
- Representational similarity structure (e.g., Edelman, 1998)
- Partial least squares in fMRI data analysis (e.g., McIntosh et al, 1996)

Comparing univariate and multivariate analyses

Univariate analyses:

- Detect only voxels with maximal responses to experimental manipulation
- Captures areas where all voxels show an effect in the same direction
- Typically involves spatial smoothing or averaging within an ROI
- Increased sensitivity to spatially-extended activation
- Interpretations sometimes resort to reverse inference; i.e., inferring mental states from localized brain activation (Poldrack, 2011)

Multivariate analyses:

- Exploit entire topography of responses, including submaximal responses
- Sensitive to fine-grained spatial pattern differences in absence of regional-average differences
- Pools across voxels and relaxes spatial contiguity constraints
- Benefits from high-resolution data
- More legitimately allows us to infer mental states from brain activity

Comparing univariate and multivariate analyses

Univariate analyses:
- Detect only voxels with maximal responses to experimental manipulation
- Captures areas where all voxels show an effect in the same direction

Multivariate analyses:
- Exploit entire topography of responses, including submaximal responses
- Sensitive to fine-grained spatial pattern differences in absence of regional-average differences

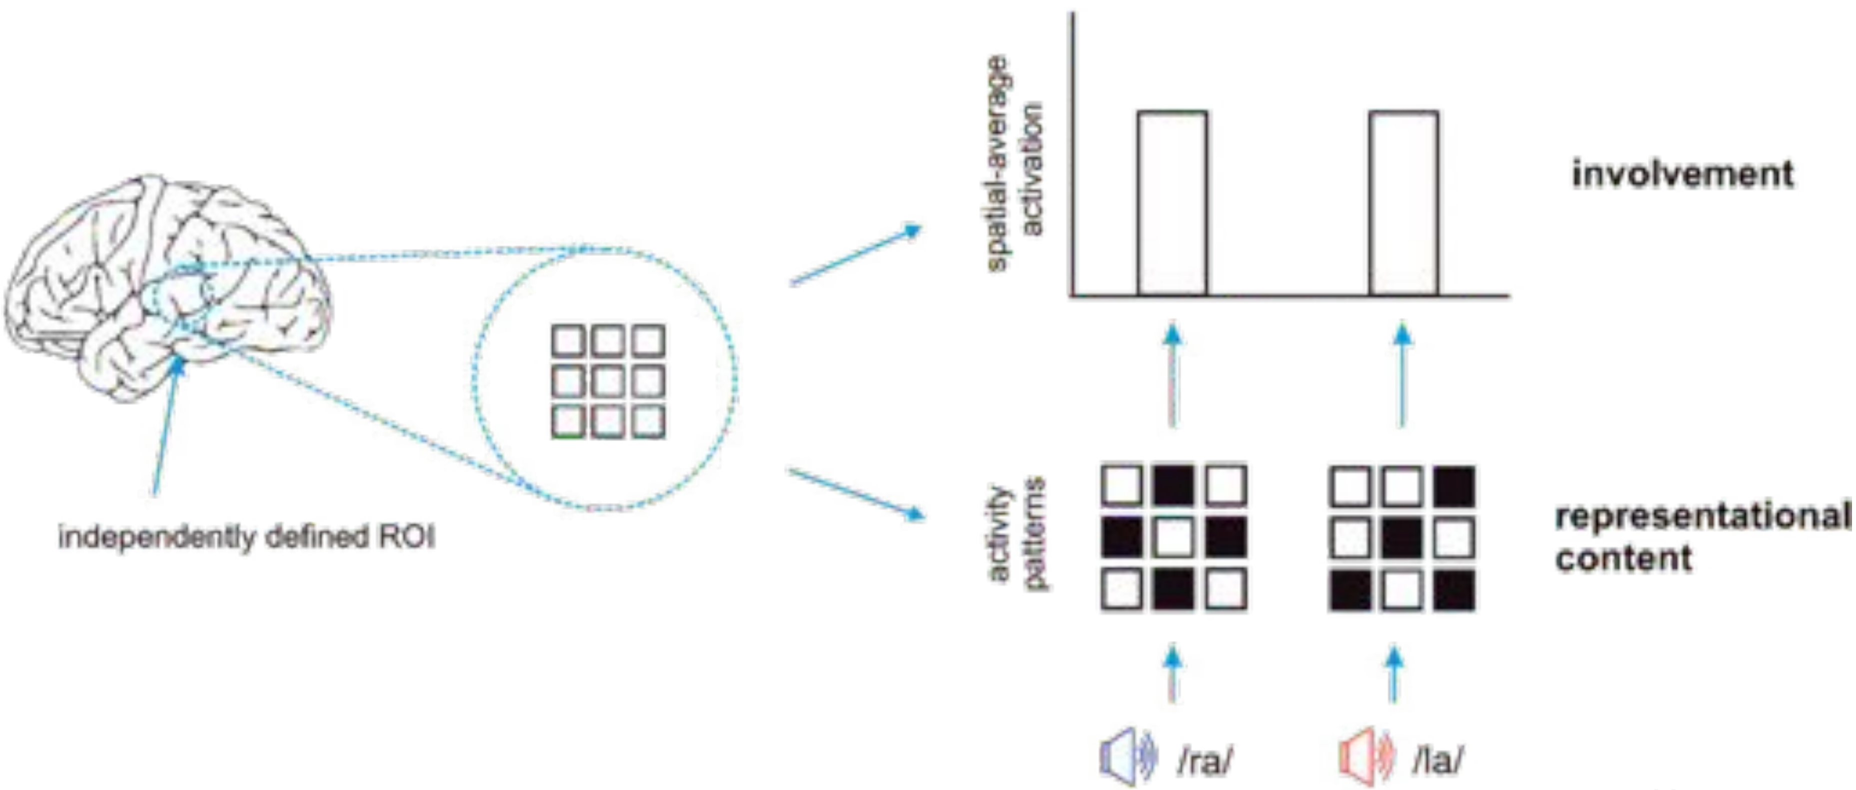

Mur et al, 2009

Multivariate approaches

The choice of multivariate method will heavily impact design considerations

Whole-brain information mapping vs. ROI analyses
- Searchlight methods (Kriegeskorte et al, 2006)
- Functional localizers (avoid double-dipping; Kriegeskorte et al, 2009)
- Anatomical ROIs

Classification/decoding
- Limited number of stimulus classes
- Highly sensitive (to both signal and noise), amenable to searchlights
- e.g., Haxby et al, 2001; Kamitani & Tong, 2005

Representational similarity analysis (RSA)
- Condition-rich designs
- Amenable to searchlights
- e.g., Kriegeskorte et al, 2008; Connolly et al, 2012

Stimulus encoding/decoding models
- Computational model of feature space, generalizes to novel stimuli
- e.g., Kay et al, 2008; Mitchell et al, 2008; Brouwer & Heeger, 2009; Huth et al, 2012

Multivariate approaches

The choice of multivariate method will heavily impact design considerations

Whole-brain information mapping vs. ROI analyses
- Searchlight methods (Kriegeskorte et al, 2006)
- Functional localizers (avoid double-dipping; Kriegeskorte et al, 2009)
- Anatomical ROIs

Classification/decoding
- Limited number of stimulus classes
- Highly sensitive (to both signal and noise), amenable to searchlights
- e.g., Haxby et al, 2001; Kamitani & Tong, 2005

Representational similarity analysis (RSA)
- Condition-rich designs
- Amenable to searchlights
- e.g., Kriegeskorte et al, 2008; Connolly et al, 2012

Stimulus encoding/decoding models
- Computational model of feature space, generalizes to novel stimuli
- e.g., Kay et al, 2008; Mitchell et al, 2008; Brouwer & Heeger, 2009; Huth et al, 2012

Multivariate approaches

The choice of multivariate method will heavily impact design considerations

What will be the input to my classifier?

What will I use for cross-validation?

…more on these later

Some vocabulary

Features: Variables representing the dimensions of your dataset; typically *voxels*, but sometimes surface nodes, time points, ROIs, etc

Samples: Observation or measurement of the values taken by these features; typically betas corresponding to certain condition or stimulus class; examples or exemplars

Targets: The class, category, experimental condition, or label with which any given sample is associated

Runs: Repeated, independent samples of the data; epochs, chunks; all runs typically include the same set of items

Classifier: A function that takes values of the *features* (independent variables) for a particular *sample* (a set of values taken by the IVs) and tries to predict the *target* (dependent variable) to which the sample corresponds.

Allows us to determine whether the features encode information about the class associated with a particular example.

The classifier is **trained** on a subset of the data (e.g., 9 out 10 total runs) and **tested** on the held-out subset of data (e.g., the 10th run).

Participants

Multivariate analysis capitalize on fine-grained spatial topographies

- Very susceptible to motion artifacts, noise
- Conventional classification does not generalize well across participants because fine-grained cortical anatomy is inhomogeneous (cf. hyperalignment; Haxby et al, 2011)
- Analyses are typically performed within-participants in native space, then pattern-information estimates/maps are aggregated at group level

Patient populations

- Classifier can (read: *will*) capitalize on *any* diagnostic differences
- Not necessarily experimentally interesting neural differences
- e.g., differences in head motion, physiology, etc

Stimuli

Controlled for low-level visual properties

Multiple exemplars per class/category for classification study

Numerous, non-repeated stimuli for RSA study (condition-rich design; Kriegeskorte et al, 2008)

Naturalistic stimuli (e.g., video clips)

Other ways to controlling for low-level visual properties

- Use multiple different exemplars per class and average responses
- Cross-validation across exemplars
- Use model of early visual processing in RSA (e.g., Connolly et al, 2012)

Block design

Characteristics of block design

- BOLD responses to different conditions do not overlap (separate by 8-12 s rest)
- Multiple exemplars of stimulus class presented within block
- Can use *normalized signal, mean signal*, betas (or t-values) from GLM (i.e., do not need to rely on GLM)
- Typically yields high functional contrast to noise ratio
- May yield better estimate of average response pattern

Haxby et al, 2001:

- 8 participants
- 8 stimulus classes
- Block design
- 1500 ms ISI
- 24 s trial blocks
- 12 s rest intervals
- 8 blocks per run
- 12 runs
- Betas input to classifier

Adapted from Haxby et al, 2001

Slow event-related design

Characteristics of slow event-related design:

- Single (or several) brief stimulus presentation(s) followed by 8-12 s rest
- Similarly to block design, intertrial interval avoids hemodynamic overlap/cross-trial contamination
- Can use normalized signal, mean signal—allows you to be “closer” to the raw data rather than relying on GLM
- e.g., use normalized signal lagged by 4-6 s to account for hemodynamic delay (e.g., Polyn et al, 2005), or average over several TRs (e.g., Pereira et al, 2009)

Rapid event-related design

Fast event-related design

- Requires GLM, betas or t-values as input
- Temporally overlapping hemodynamic responses
- Can afford greater number of conditions (Kriegeskorte et al, 2008)
- Yield more independent data points, which can improve classification performance, but noisier estimates of average response pattern for a condition

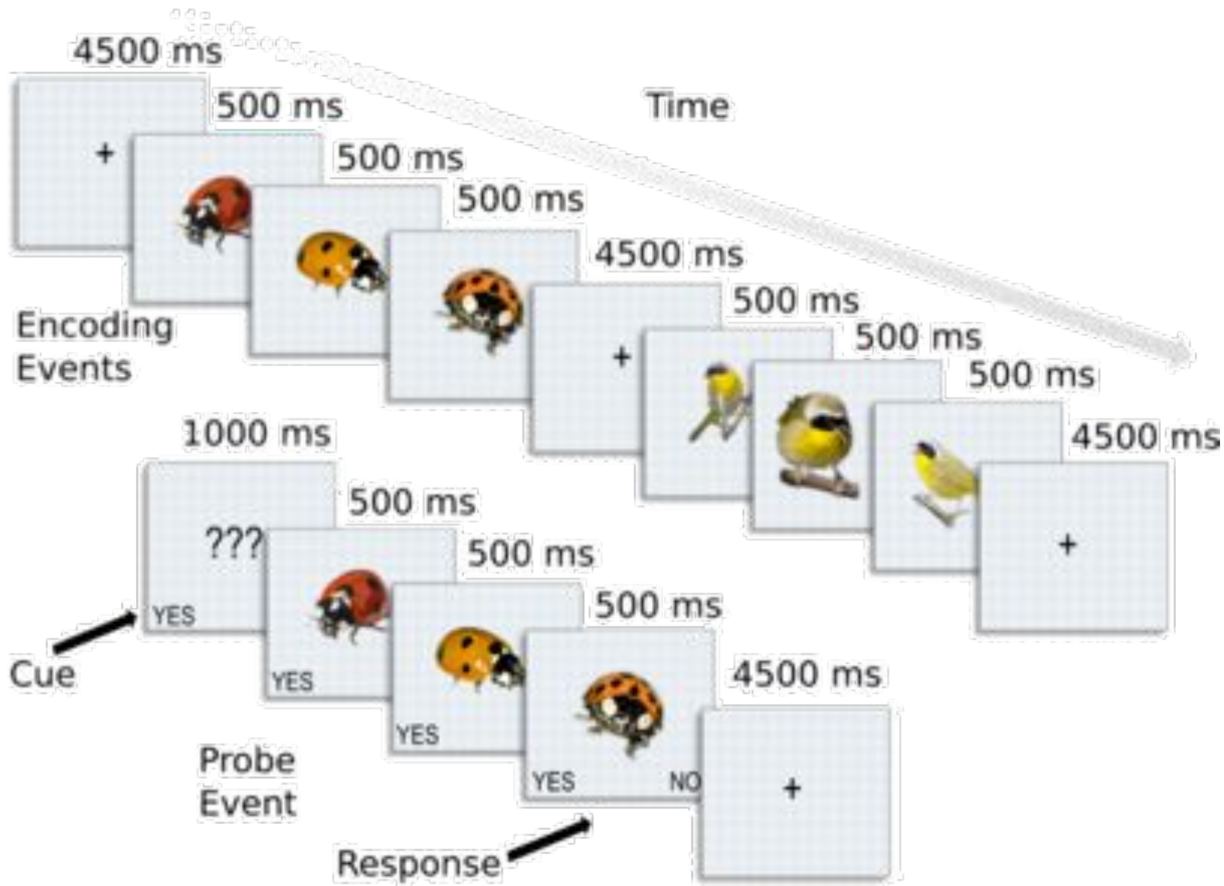

Connolly et al, 2012

Rapid event-related design

Fast event-related design

- Requires GLM, betas or t-values as input
- Temporally overlapping hemodynamic responses
- Can afford greater number of conditions (Kriegeskorte et al, 2008)

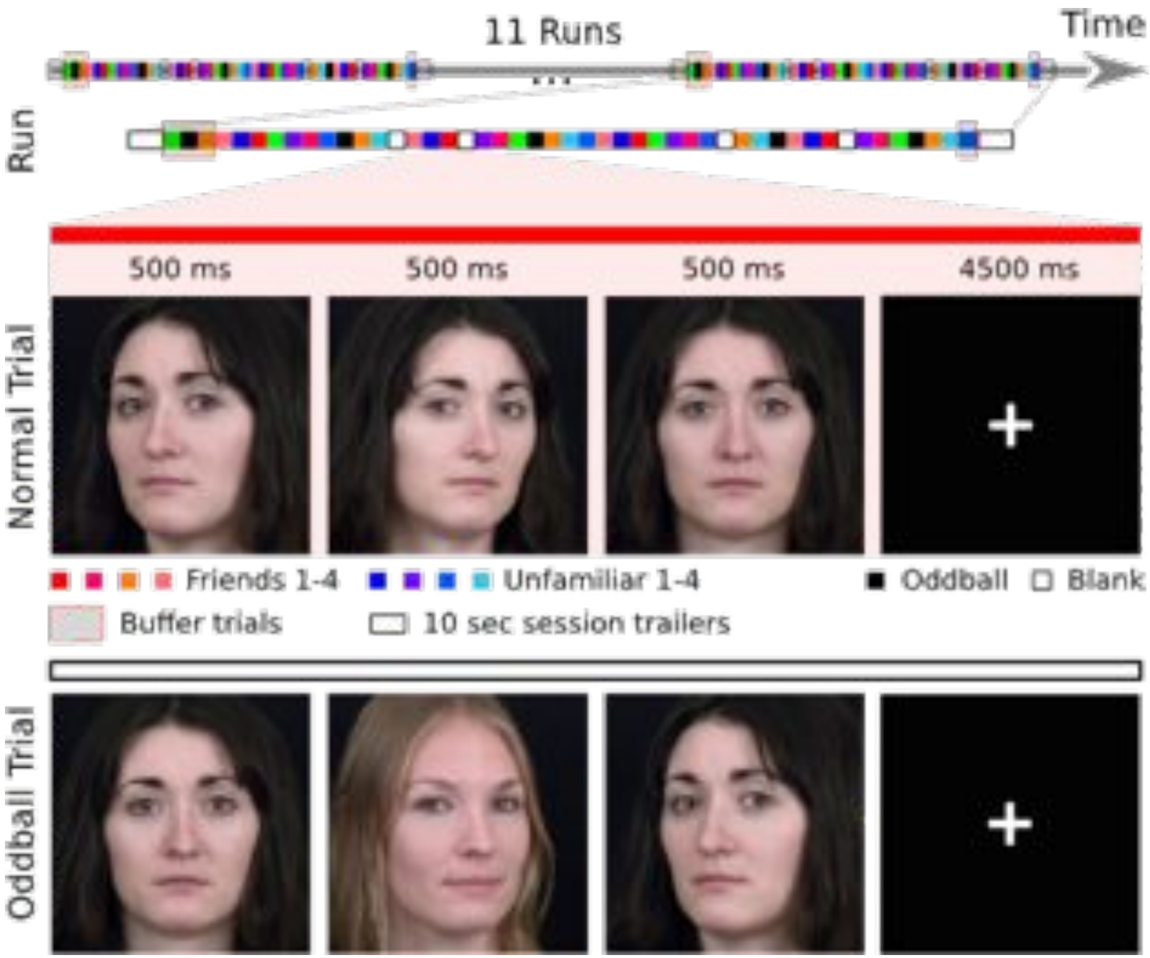

Halchenko et al, 2012

Rapid event-related design

Fast event-related design
- Requires GLM, betas or t-values as input
- Temporally overlapping hemodynamic responses
- Can afford greater number of conditions (Kriegeskorte et al, 2008)

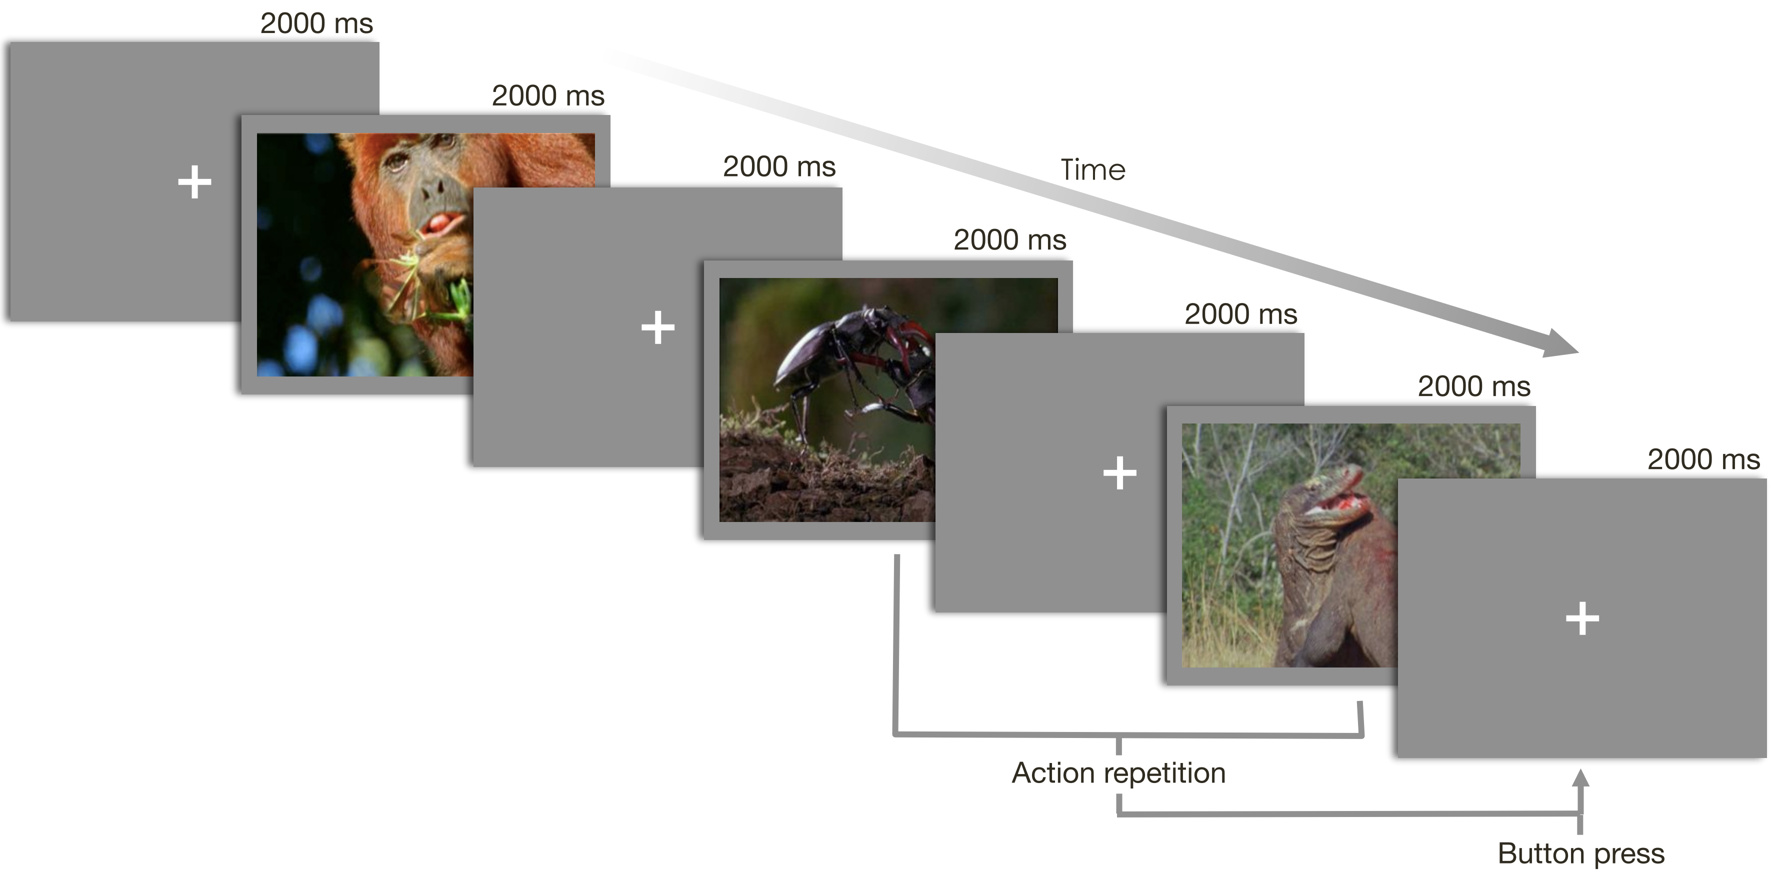

Nastase et al, in prep

Condition-rich design

Suited for representational similarity analysis (RSA; Kriegeskorte et al, 2008)

- Rapid events; 4 s trial onset asynchrony
- Overlapping but dissociable hemodynamic responses
- Requires the use of GLM

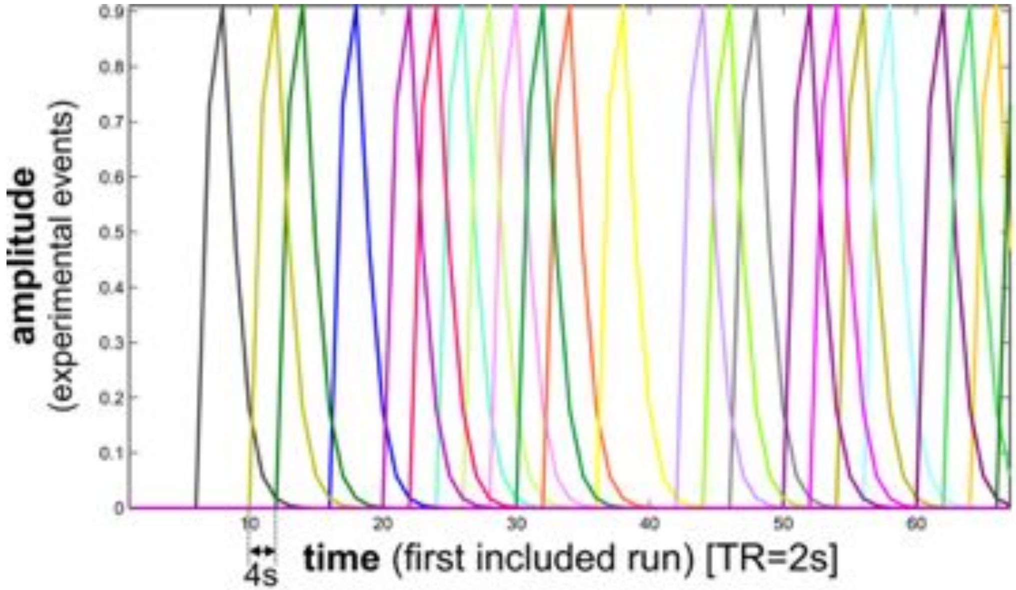

Design matrix for condition-rich ungrouped-events design

- 96 conditions
- Randomized trial order
- Each condition occurs once per run
- 4 runs

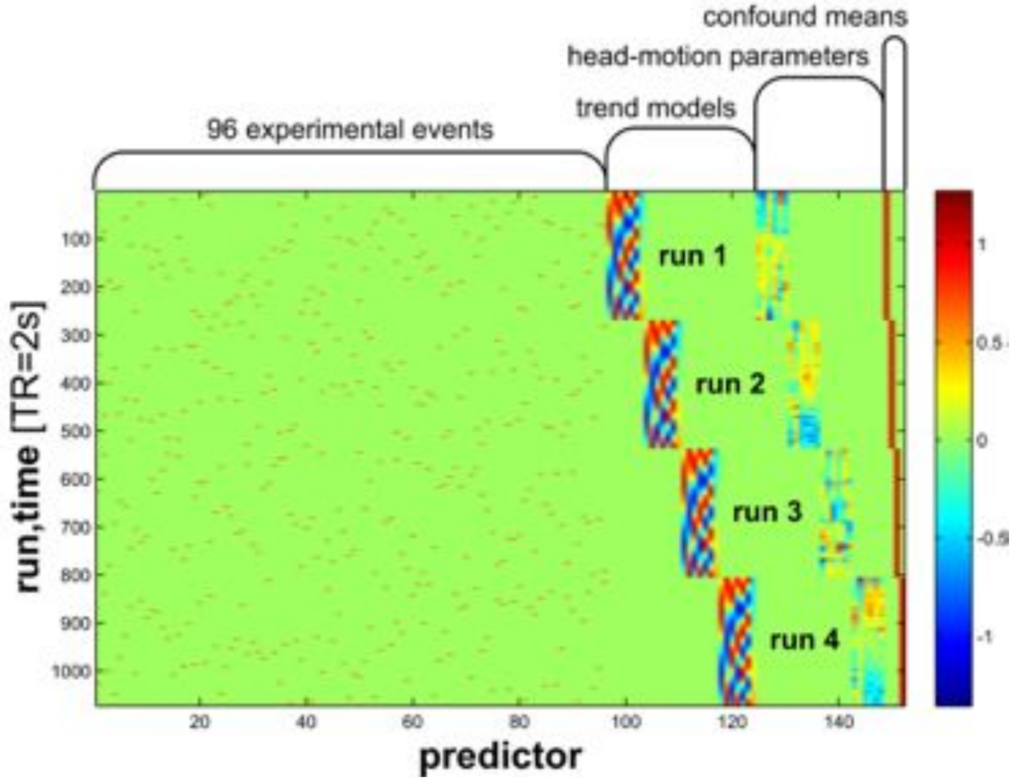

Kriegeskorte et al, 2008

Designs

Statistical efficiency as a result of trial onset asynchrony

- Statistical efficiency increases with more highly packed designs, despite greater hemodynamic overlap
- Relies on *linear* neuronal and hemodynamic response
- Less efficient for estimating a particular contrast
- Happy medium: 4 s TOA

Condition-rich design richly samples similarity space and avoids arbitrary category memberships

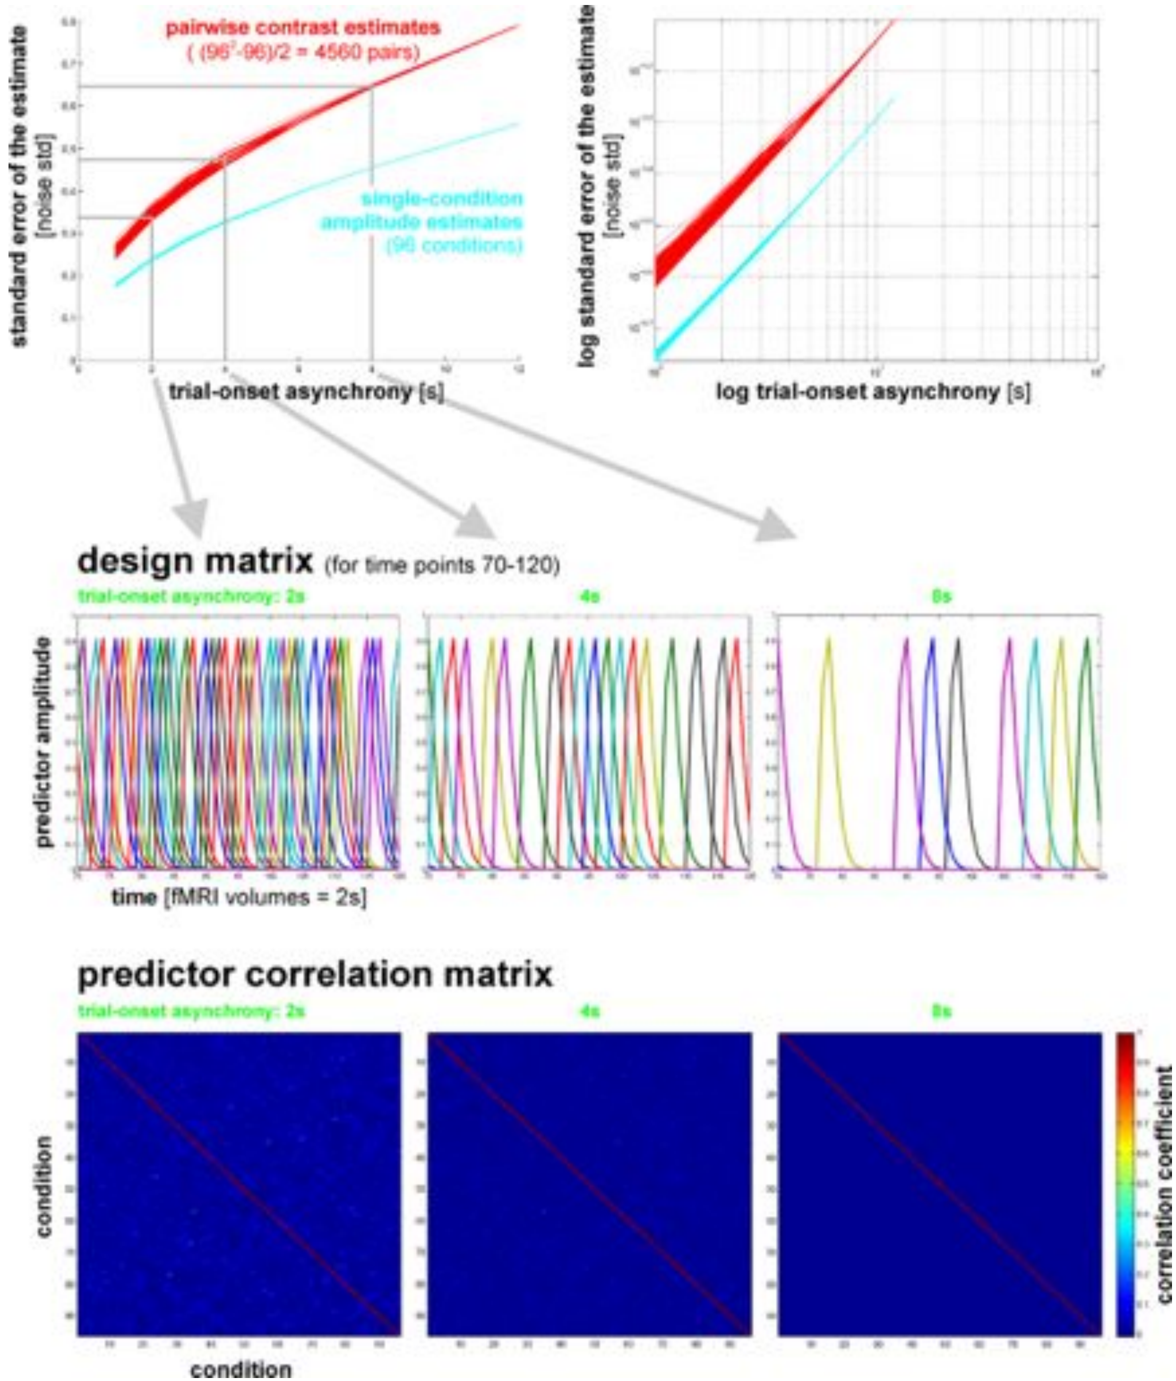

Kriegeskorte et al, 2008

Designs

Ungrouped-events design

- Randomized trial order rather than block or grouped event designs
- Continuous carry-over design (Aguirre, 2007)

Unique-events design

- Each condition presented only once

Time-continuous design

- Each volume is treated as a separate condition, resulting in time-point dissimilarity matrix (time points x time points matrix)
- May be useful for language, movies, video games, virtual reality
- Allows for greater ecological validity

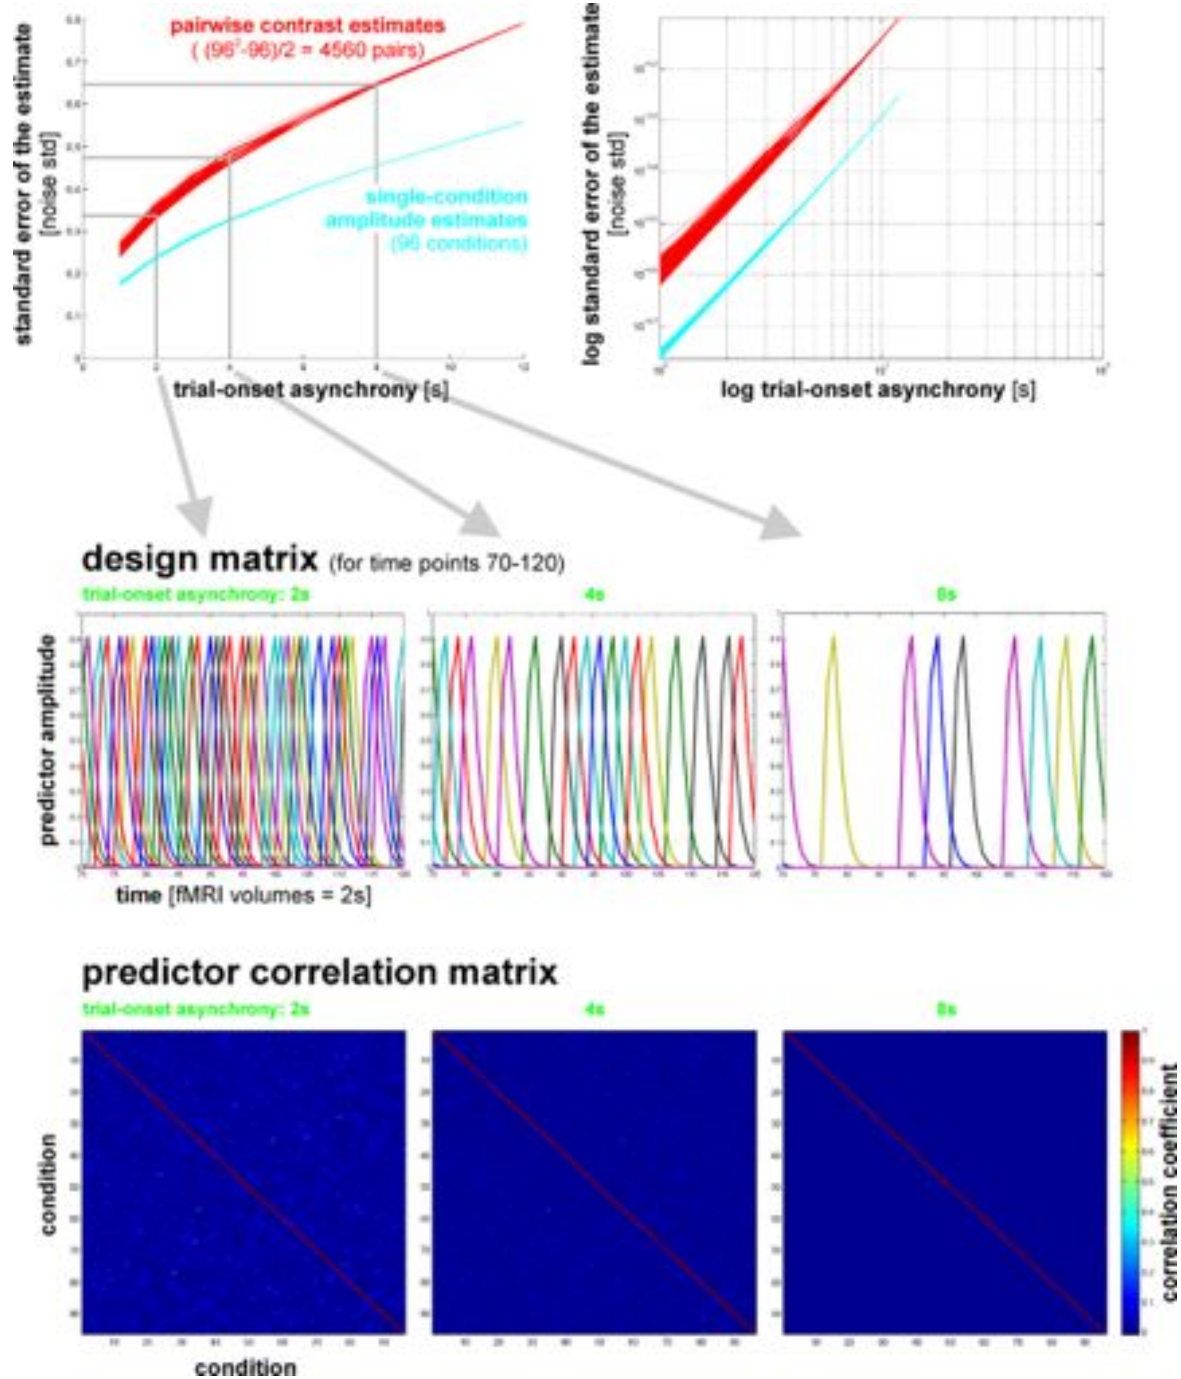

Kriegeskorte et al, 2008

Task

Tasks used to assure vigilance

- Catch trials, questions
- N-back repetition detection (usually 1-back)
- Button responses can be included as nuisance regressors
- Button responses must be balanced across classes and runs
- Must not systematically vary with conditions (e.g., always use same response hand)

More sophisticated task manipulations

- Harrison & Tong, 2009
- Haynes & Rees, 2005
- Polyn et al, 2005

Order effects

"With great power comes great responsibility" (Uncle Ben, 1962)

The greater sensitivity of multivariate analyses makes them more susceptible to order effects

- Random/pseudorandom and first-order counterbalancing may not be enough
- de Bruijn sequences

Run order

- Counterbalance across participants
- Use identical runs across participants for hyperalignment

Haxby Lab standard method:

- Generate 8-12 runs with internal trial order counterbalanced via de Bruijn sequence
- Counterbalance run orders across participants
- In analysis, rearrange runs to get same run order across participants—important for hyperalignment

Important:
Use the same number of samples/exemplars per class! Why?

- If uneven, classifier may be biased toward most frequent
- If averaging, more frequent classes will also be cleaner

Jittering

Should trials be temporally jittered?

- Jittering demands modeling responses with GLM, therefore input to classifier will be betas or *t*-values
- Numerous successful studies do not use jittering (e.g., Connolly et al, 2012; see Kriegeskorte et al, 2008, for discussion)

Jittering is particularly important when estimating the *shape* of the HRF, i.e., via finite-impulse-response model

- Important for spatiotemporal decoding (Turner et al, 2012; Mourão-Miranda et al, 2007)
- In the context of estimating response amplitudes using canonical HRF (single beta-of-interest per response), this is less important

Survey of study designs

Author	Year	Journal	N	Design	Trial	Runs	Conditions	Exemplars	Input	MVPA
Haxby et al	2001	Science	6	Block	36 s	12	8	12	Betas	kNN
Cox & Savoy	2003	NeuroImage	4	Block	20 s	8	10	12	Avg. signal	LDA, SVM
Polyn et al	2005	Science	9	Slow ER	11 s	4	3	10	Z-scored signal	NN
Kamitani & Tong	2005	Nature Neurosci	4	Block	16 s	20+	8	20+	Avg. signal	SVM
Haynes & Rees	2005	Nat Neurosci	4	Block	15 s	8	2	16	Signal	LDA
Haynes et al	2007	Curr Biol	8	Block	3-11 s	8	2	–	Betas	SVM
Kriegeskorte et al	2006	PNAS	11	Rapid ER	3 s	1	2	–	Betas	RSA
Rissman et al	2010	PNAS	16	Slow ER	8 s	10	10	40	Betas	LR
Carlin et al	2011	Curr Biol	18	Rapid ER	4 s	5	50	15	Betas	RSA
Nestor et al	2011	PNAS	8	Slow ER	10 s	17	32	17	Avg. signal	SVM
Connolly et al	2012	J Neurosci	12	Rapid ER	6 s	10	6	32 (64)	Betas (FIR, 6 s)	SVM, RSA

Classifier input

Normalized signal intensity values

Averaged interval of signal intensity values

Betas
- Canonical HRF
- Basis functions

T-values

Spatiotemporal

Cross-validation

Runwise cross-validation, better to have more numerous, short runs (Coutanche & Thompson-Schill, 2012)

- *k*-fold, leave-one-run-out cross-validation, *k* should equal 5-10
- Split half cross-validation
- Data from different runs are presumably independent

Samplewise cross-validation

- Better controls for low-level visual properties

Cross-validating across participants

- May require hyperalignment

Is my current (univariate) dataset amenable to MVPA?

Maybe!

- Are there well-defined classes/categories of stimuli for classification?
- Are the numbers of presentations for each class balanced?
- What will the inputs to my classifier be?
- Are there enough samples per class (>10)?
- Are there enough runs for runwise cross-validation?

Questions amenable to MVPA

- *Pattern detection:* Does the brain, or any region of the brain, contain information about the variable of interest? That is, do the response patterns discriminate stimulus classes/ experimental conditions

- *Pattern localization:* Where in the brain are distinctions among stimulus classes represented? Searchlight methods, feature selection and sensitivity measures

- *Pattern characterization:* How is information represented? Does the brain encode the stimulus classes according to particular relationships between the classes? RSA, confusion matrices

References and resources

Books:

Poldrack, R. A., Mumford, J. A., & Nichols, T. E. (2011). Multivoxel pattern analysis and machine learning. In *Handbook of Functional MRI Data Analysis* (160-172). Cambridge University Press.

Kriegeskorte, N., & Kreiman, G. (Eds.). (2012). *Visual population codes: Toward a common multivariate framework for cell recording and functional imaging*. MIT press.

Special issues:

Haynes, J. D. (2011). Multivariate decoding and brain reading [Special issue]. *NeuroImage 56*(2), 385-850.

Norman, K. A. (2012). Multivoxel pattern analysis and cognitive theories [Special issue]. *Neuropsychologia, 3*(4), 435-552.

Reviews:

Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: a tutorial overview. *NeuroImage*, *45*(1), S199-S209.

Mur, M., Bandettini, P. A., & Kriegeskorte, N. (2009). Revealing representational content with pattern-information fMRI—an introductory guide. *Social Cognitive and Affective Neuroscience*, *4*(1), 101-109.

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multivoxel pattern analysis of fMRI data. *Trends in Cognitive Sciences 10*(9), 424-430.

Haynes, J. D. & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience, 7*(7), 523-534.

O'Toole, A. J., Jiang, F., Abdi, H., Pénard, N., Dunlop, J. P., Parent, M. A. (2007). Theoretical, statistical, and practical perspectives on pattern-based classification approaches to the analysis of functional neuroimaging data. *Journal of Cognitive Neuroscience, 19*(11), 1735-1752.

Tong, F., & Pratte, M. S. (2012). Decoding patterns of human brain activity. *Annual Review of Psychology*, *63*, 483-509.

Articles:

Kriegeskorte, N., Goebel, R., Bandettini, P. A. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences, 104*(51), 20600-20605.

Misaki, M., Kim, Y., Bandettini, P. A., Kriegeskorte, N. (2010). Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *NeuroImage 53*(1), 103-118.

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*.

Connolly, A. C., Guntupalli, J. S., Gors, J., Hanke, M., Halchenko, Y. O., Wu, Y. C., Abdi, H., & Haxby, J. V. (2012). The representation of biological classes in the human brain. *The Journal of Neuroscience*, *32*(8), 2608-2618.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*(5539), 2425-2430.

References and resources

Articles, cont.:

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, *8*(5), 679-685.

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632-635.

Polyn, S. M., Natu, V. S., Cohen, J. D., & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science*, *310*(5756), 1963-1966.

Haynes, J. D., & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience*, *8*(5), 686-691.

Haynes, J. D., & Rees, G. (2005). Predicting the stream of consciousness from activity in human visual cortex. *Current Biology*, *15*(14), 1301-1307.

Mourao-Miranda, J., Friston, K. J., & Brammer, M. (2007). Dynamic discrimination analysis: a spatial–temporal SVM. *NeuroImage*, *36*(1), 88-99.

Cox, D. D., & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI)“brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage*, *19*(2), 261-270.

Haynes, J. D., Sakai, K., Rees, G., Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading hidden intentions in the human brain. *Current Biology*, *17*(4), 323-328.

Rissman, J., Greely, H. T., & Wagner, A. D. (2010). Detecting individual memories through the neural decoding of memory states and past experience. *Proceedings of the National Academy of Sciences*, *107*(21), 9849-9854.

Nestor, A., Plaut, D. C., & Behrmann, M. (2011). Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proceedings of the National Academy of Sciences*, *108*(24), 9998-10003.

Carlin, J. D., Calder, A. J., Kriegeskorte, N., Nili, H., & Rowe, J. B. (2011). A head view-invariant representation of gaze direction in anterior superior temporal sulcus. *Current Biology*, *21*(21), 1817-1821.

Coutanche, M. N., & Thompson-Schill, S. L. (2012). The advantage of brief fMRI acquisition runs for multi-voxel pattern detection across runs. *Neuroimage*, *61*(4), 1113-1119.

O'toole, A. J., Jiang, F., Abdi, H., & Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience*, *17*(4), 580-590

Aguirre, G. K., Mattar, M. G., & Magis-Weinberg, L. (2011). de Bruijn cycles for neural decoding. *NeuroImage*, *56*(3), 1293-1300.

Aguirre, G. K. (2007). Continuous carry-over designs for fMRI. *Neuroimage*,*35*(4), 1480-1494.

Poldrack, R. A. (2011). Inferring mental states from neuroimaging data: from reverse inference to large-scale decoding. *Neuron*, *72*(5), 692-697.

Turner, B. O., Mumford, J. A., Poldrack, R. A., & Ashby, F. G. (2012). Spatiotemporal activity estimation for multivoxel pattern analysis with rapid event-related designs. *NeuroImage*, *62*(3), 1429-1438.

Coutanche, M. N., & Thompson-Schill, S. L. (2012). The advantage of brief fMRI acquisition runs for multi-voxel pattern detection across runs. *NeuroImage*, *61*(4), 1113-1119.

References and resources

Articles, cont.:

Raizada, R. D., & Kriegeskorte, N. (2010). Pattern-information fMRI: New questions which it opens up and challenges which face it. *International Journal of Imaging Systems and Technology*, *20*(1), 31-41.

Hanson, S. J., Matsuka, T., & Haxby, J. V. (2004). Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a “face” area?. *NeuroImage*, *23*(1), 156-166.

O'toole, A. J., Jiang, F., Abdi, H., & Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience*, *17*(4), 580-590.

Said, C. P., Moore, C. D., Norman, K. A., Haxby, J. V., & Todorov, A. (2010). Graded representations of emotional expressions in the left superior temporal sulcus. *Frontiers in Systems Neuroscience*, *4*, 6.

Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *Journal of Neurophysiology*, *97*(6), 4296-4309.

Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience*, *12*(5), 535-540.

Kriegeskorte, N. (2011). Pattern-information analysis: from stimulus decoding to computational-model testing. *NeuroImage*, *56*(2), 411-421.

Halchenko, Y. O., Gors, J., Guntupalli, J. S., Haxby, J. V. & Gobbini, M. I. (2012). *Investigation of the familiar face recognition model.* Poster presented at the SfN 2012 Annual Meeting, New Orleans, Louisiana, USA.

McIntosh, A. R., Bookstein, F. L., Haxby, J. V., & Grady, C. L. (1996). Spatial pattern analysis of functional brain images using partial least squares. *NeuroImage*, *3*(3), 143-157.

Edelman, S. (1998). Representation is representation of similarities. *Behavioral and Brain Sciences*, *21*(04), 449-467.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*(4771), 1416-1419.

Rolls, E. T., & Tovee, M. J. (1995). Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *Journal of Neurophysiology*, *73*(2), 713-726.

Kruskal, J. B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, *29*(1), 1-27.

Rumelhart, D.E., J.L. McClelland and the PDP Research Group. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations*. MIT Press.