ONLINE SUBJECT EVALUATIONS ARE NOW OPEN

http://web.mit.edu/subjectevaluation

- You have until Monday, Dec. 16 at 9 AM
- Please evaluate all subjects in your list
- Don’t forget your TAs
- Write comments

Your feedback is read and valued!

Also, class projects due December 11th
email them to: shimon.ullman@gmail.com

Using machine learning to understand biological vision and learning

Ethan Meyers

Vision and learning: computers and brains

	Computers	Brains
Vision	Mike Jones Yann LeCun	Winrich Freiwald Elias Issa Charles Cadieu
Learning	Andrew Barto 9.520	Yael Niv Larry Abbot Haim Sompolinsky

Tomaso Poggio

Shimon Ullman

Shihab Shamma

Vision and learning: computers and brains

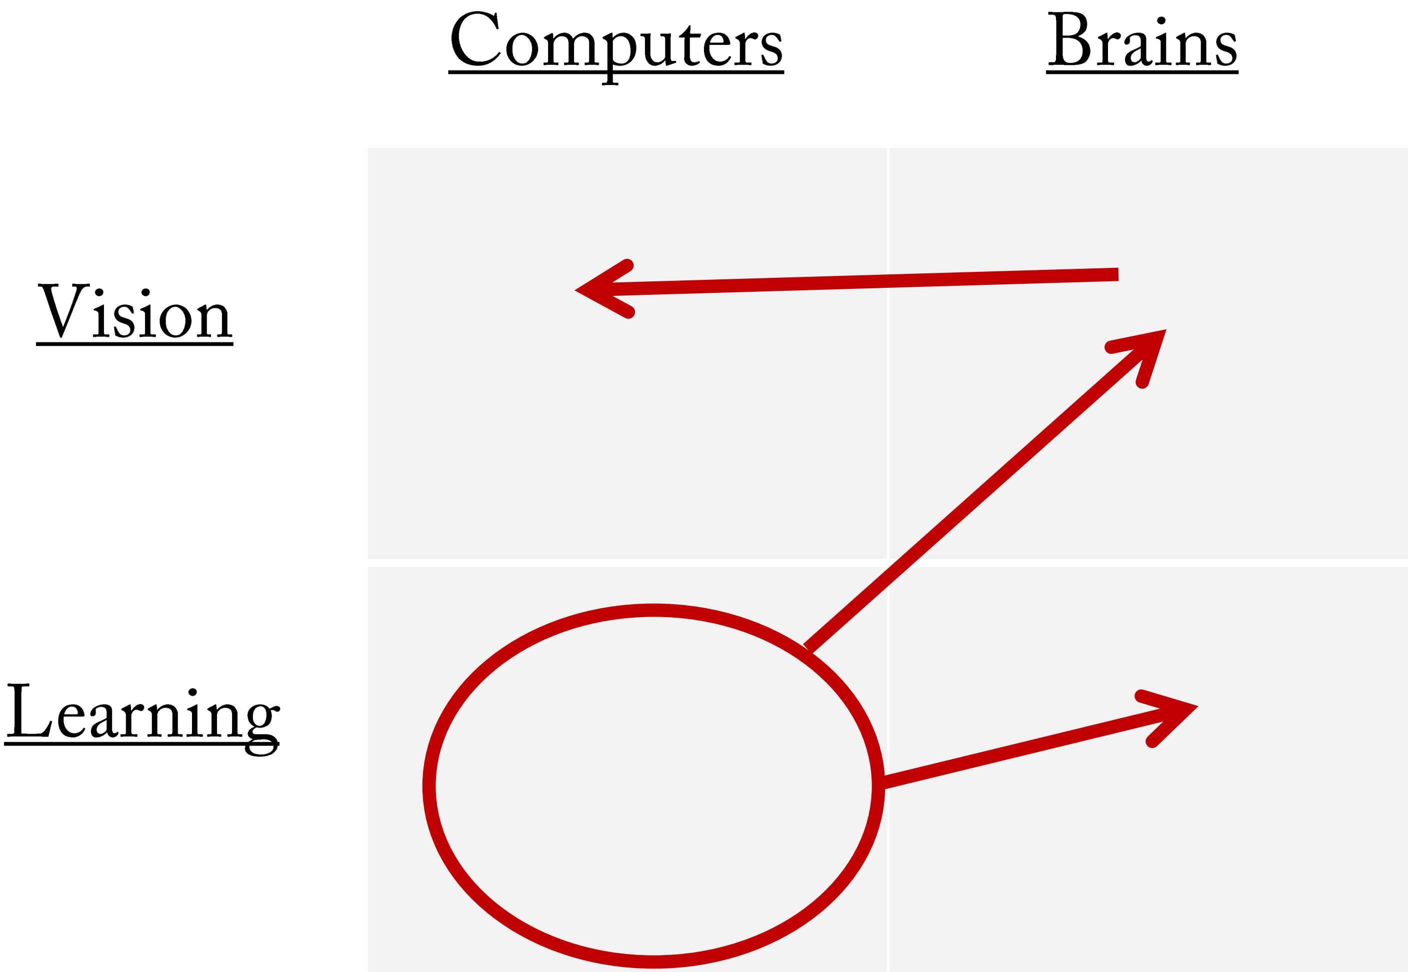

Machine learning applied to neural data

Decoding (readout/MVPA):

stimulus = f(neuronal response); $P(S|R)$

Encoding:

neural response = g(stimulus); $P(R|S)$

Decoding was used to analyze the data (training the classifier)

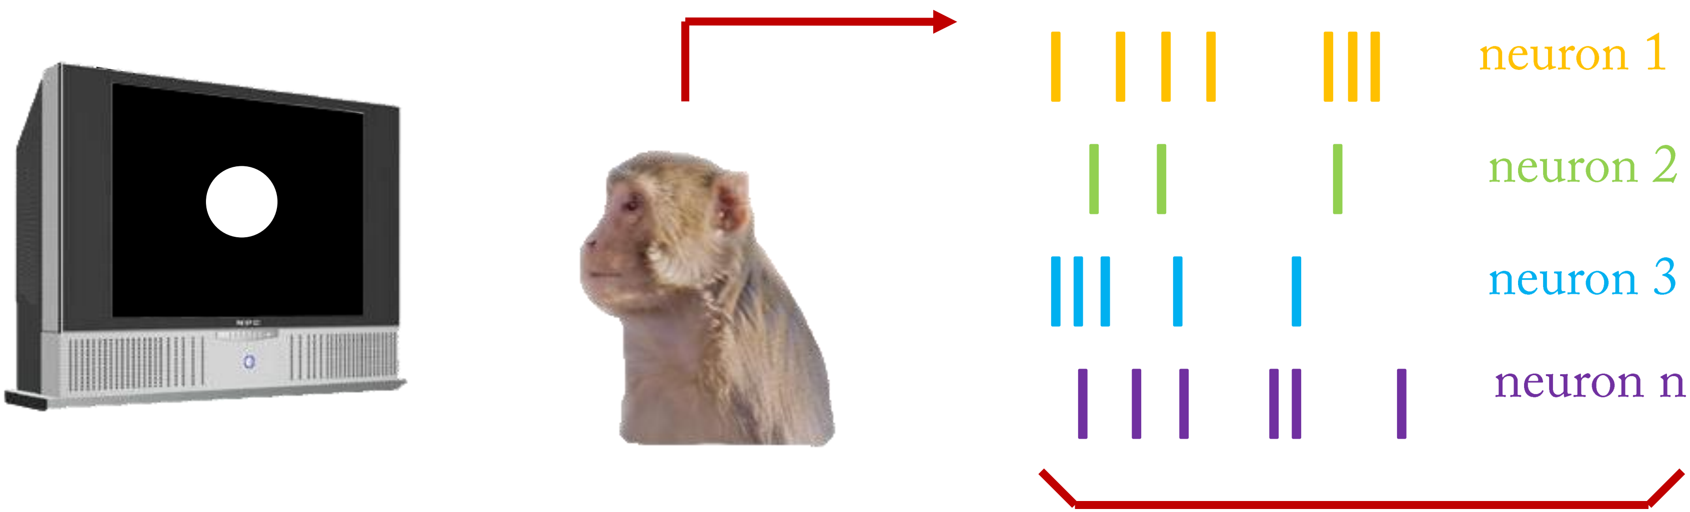

Learning association between neural activity an image

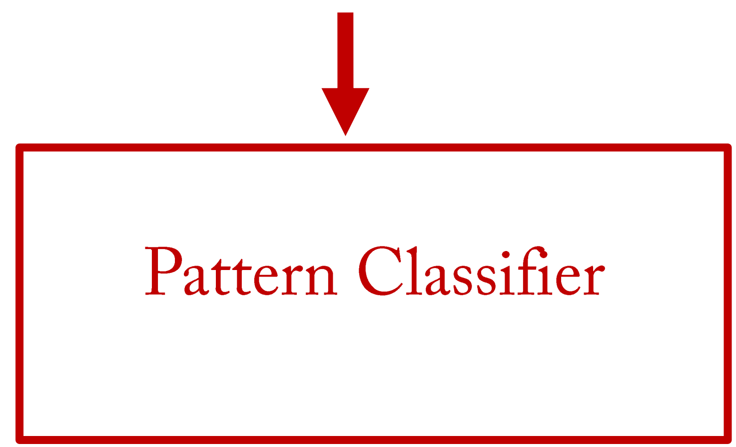

Decoding was used to analyze the data (training the classifier)

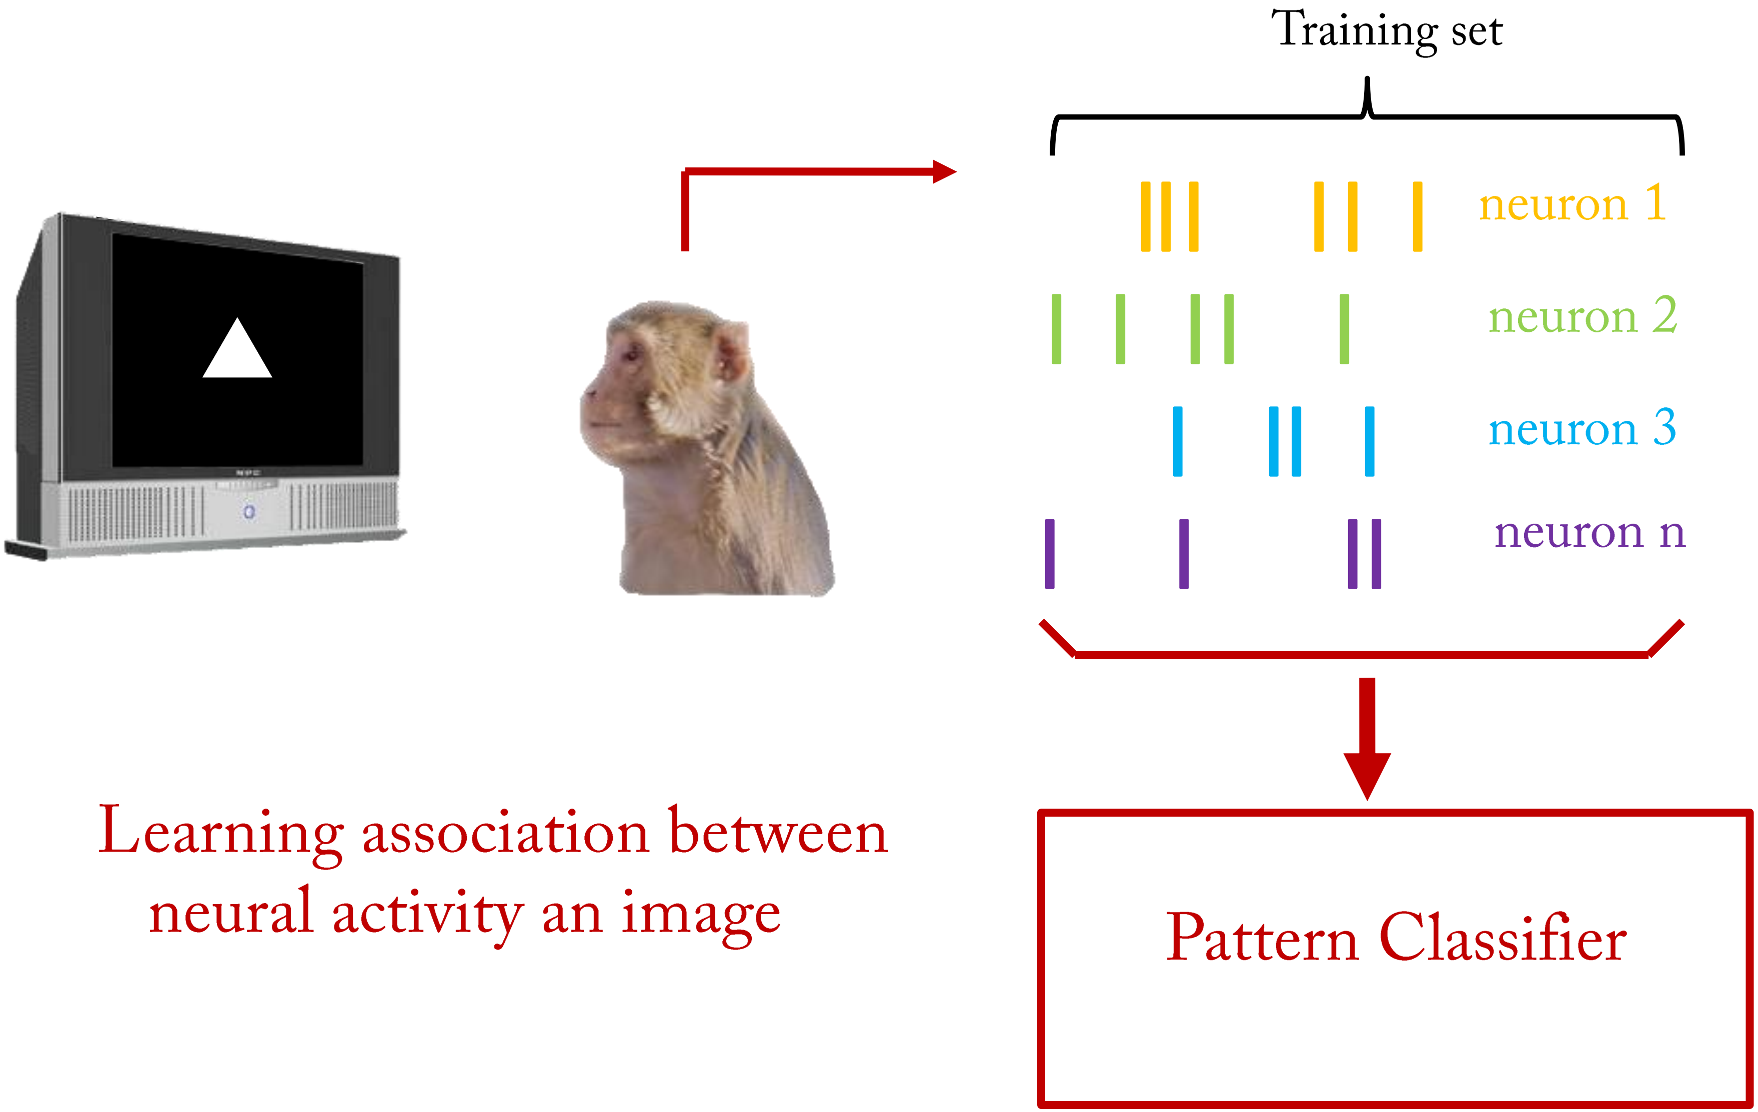

Learning association between neural activity an image

Decoding was used to analyze the data (testing the classifier)

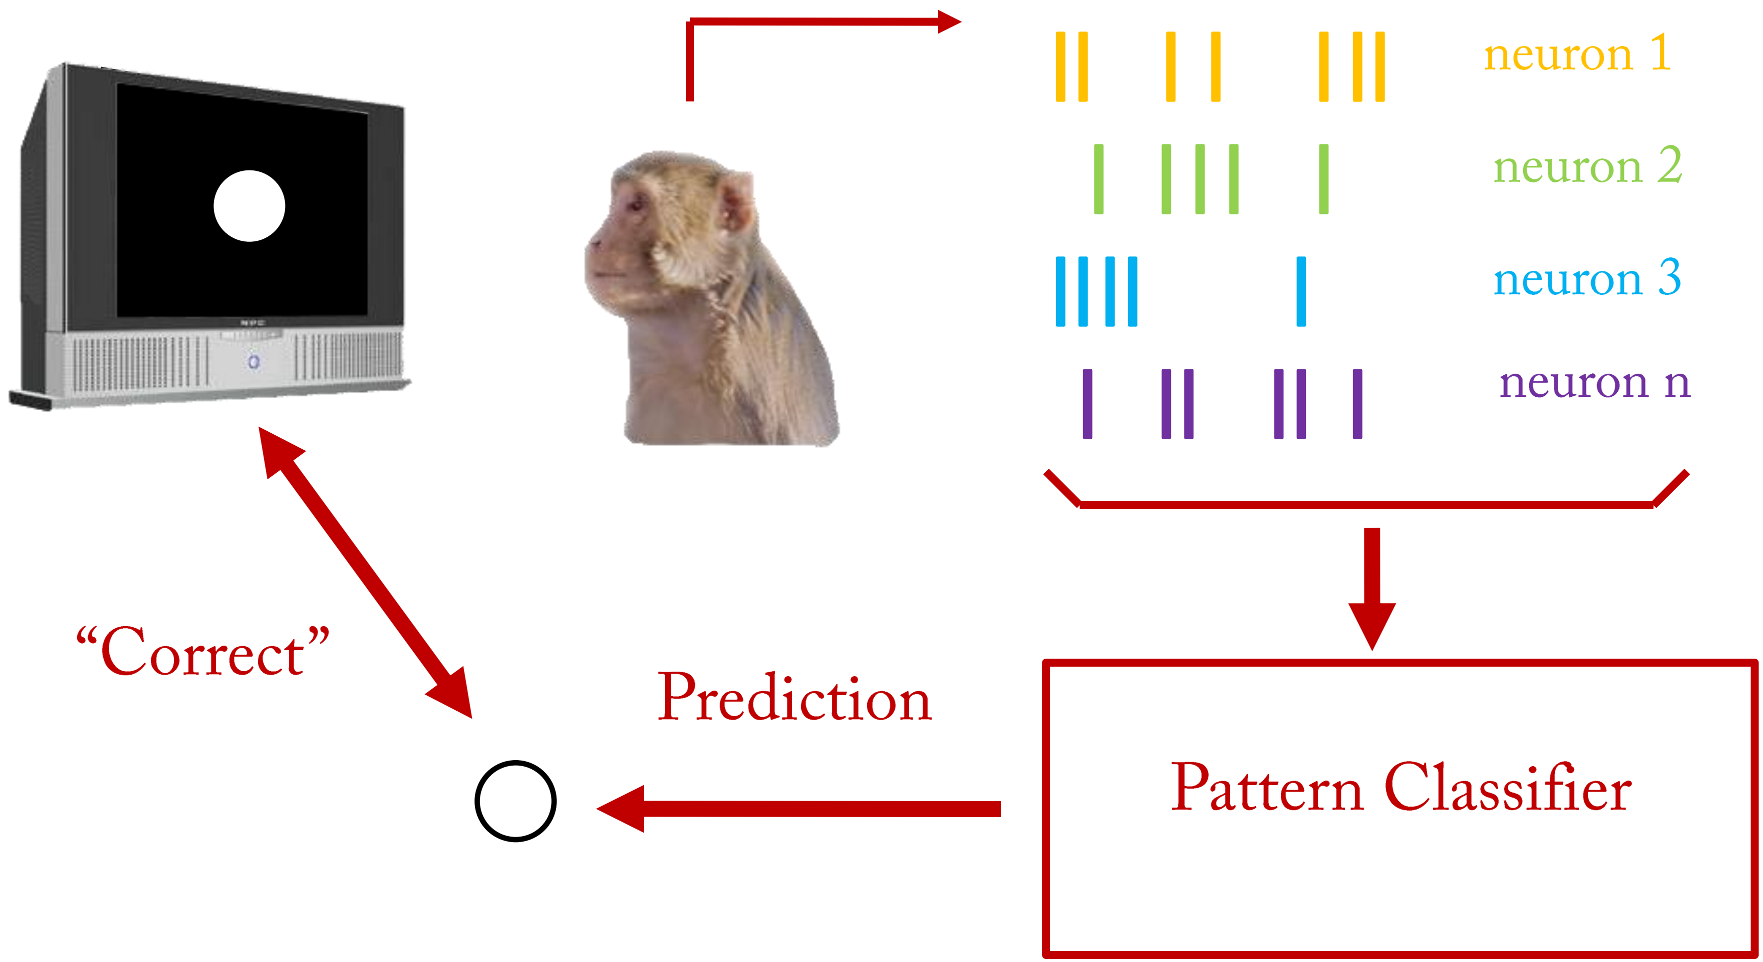

Decoding was used to analyze the data (testing the classifier)

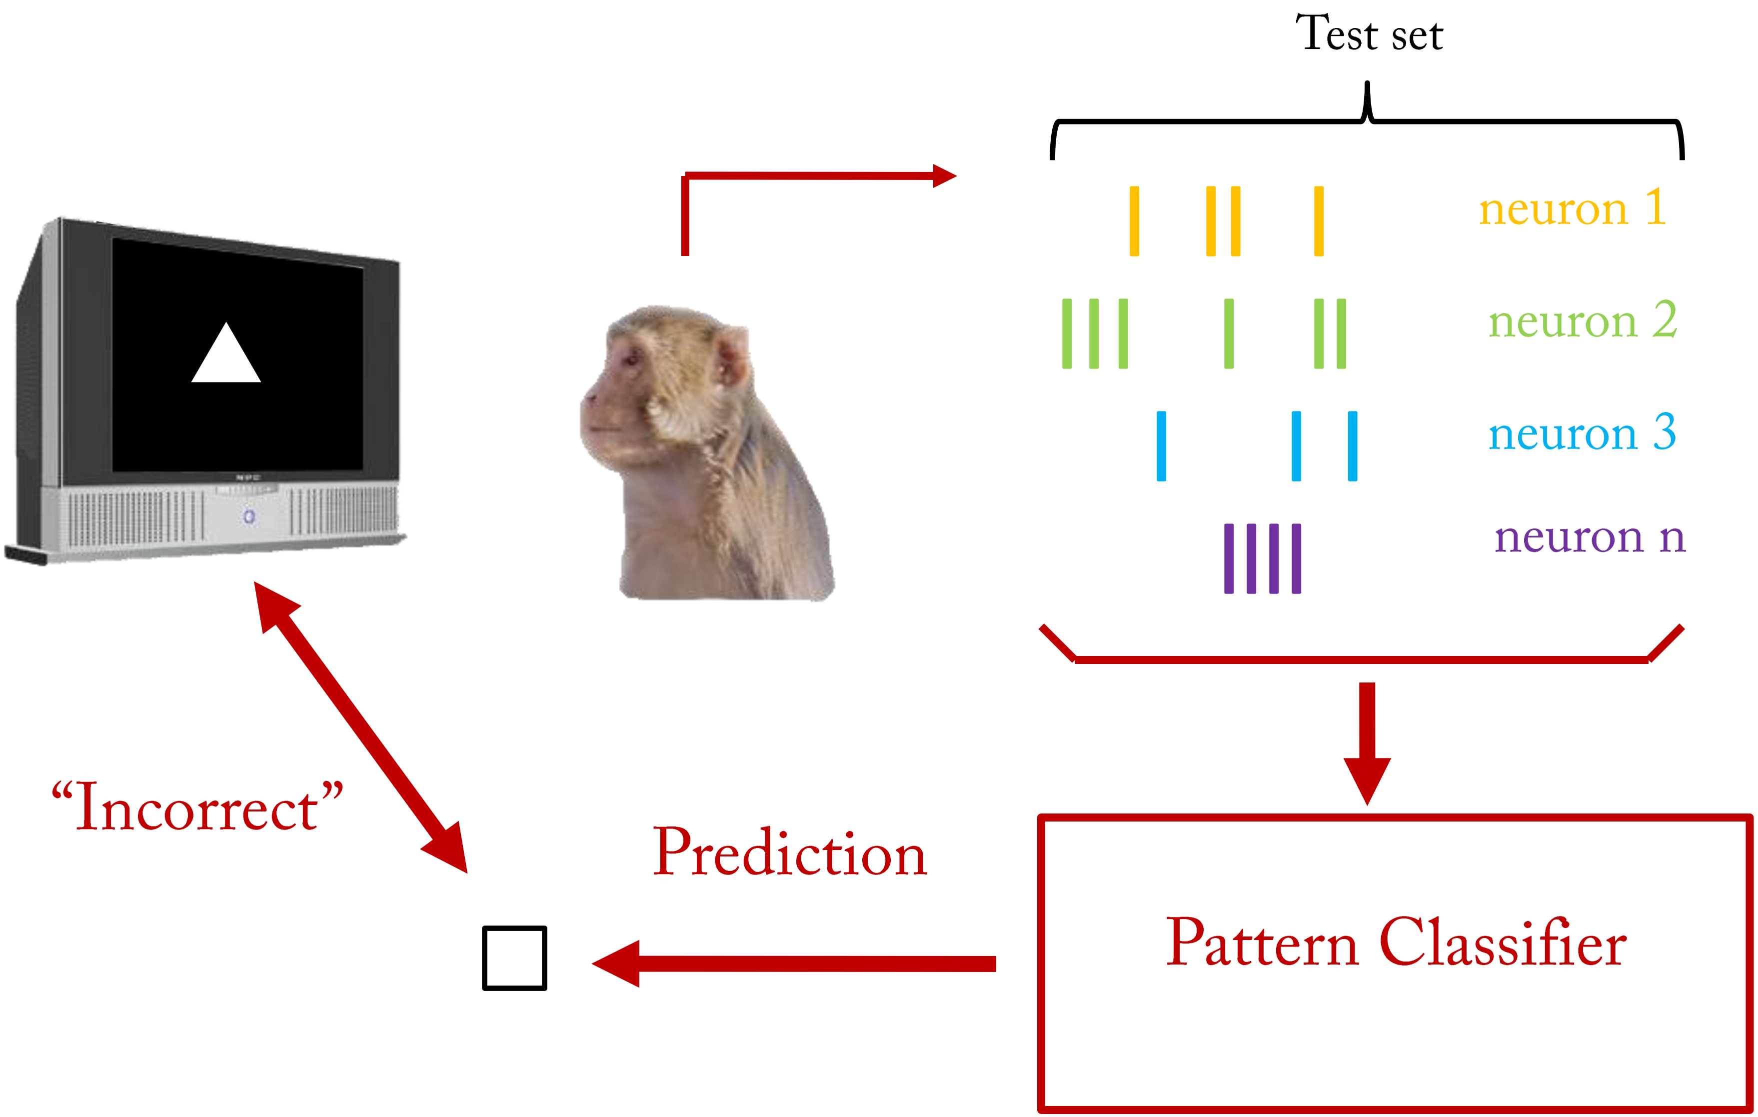

Outline

- BCI to study learning in the motor cortex
- MVPA to study vision in human fMRI data
- Population decoding to study high level learning and vision in macaque monkeys

Outline

- BCI to study learning in the motor cortex
- MVPA to study vision in human fMRI data
- Population decoding to study high level learning and vision in macaque monkeys

The motor cortex

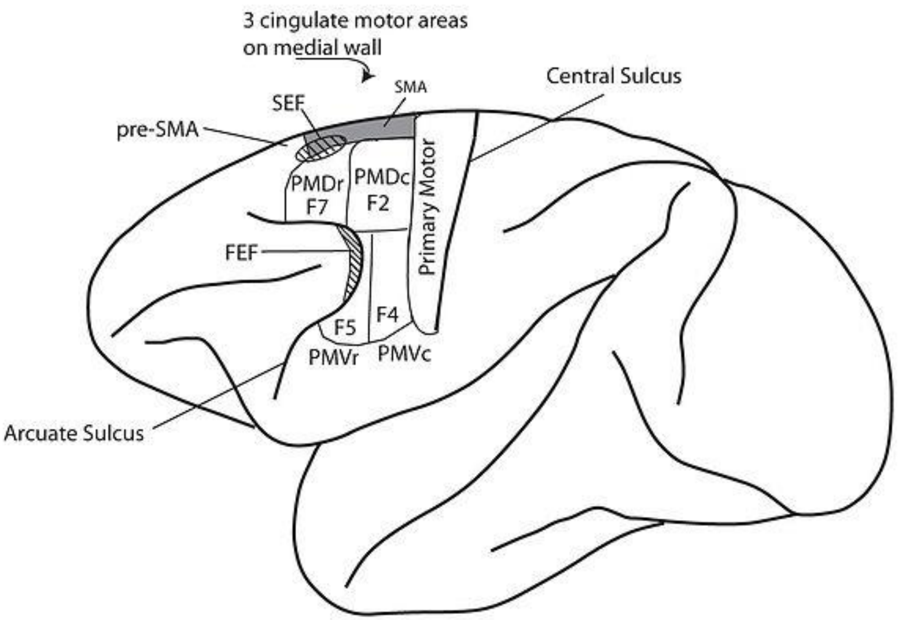

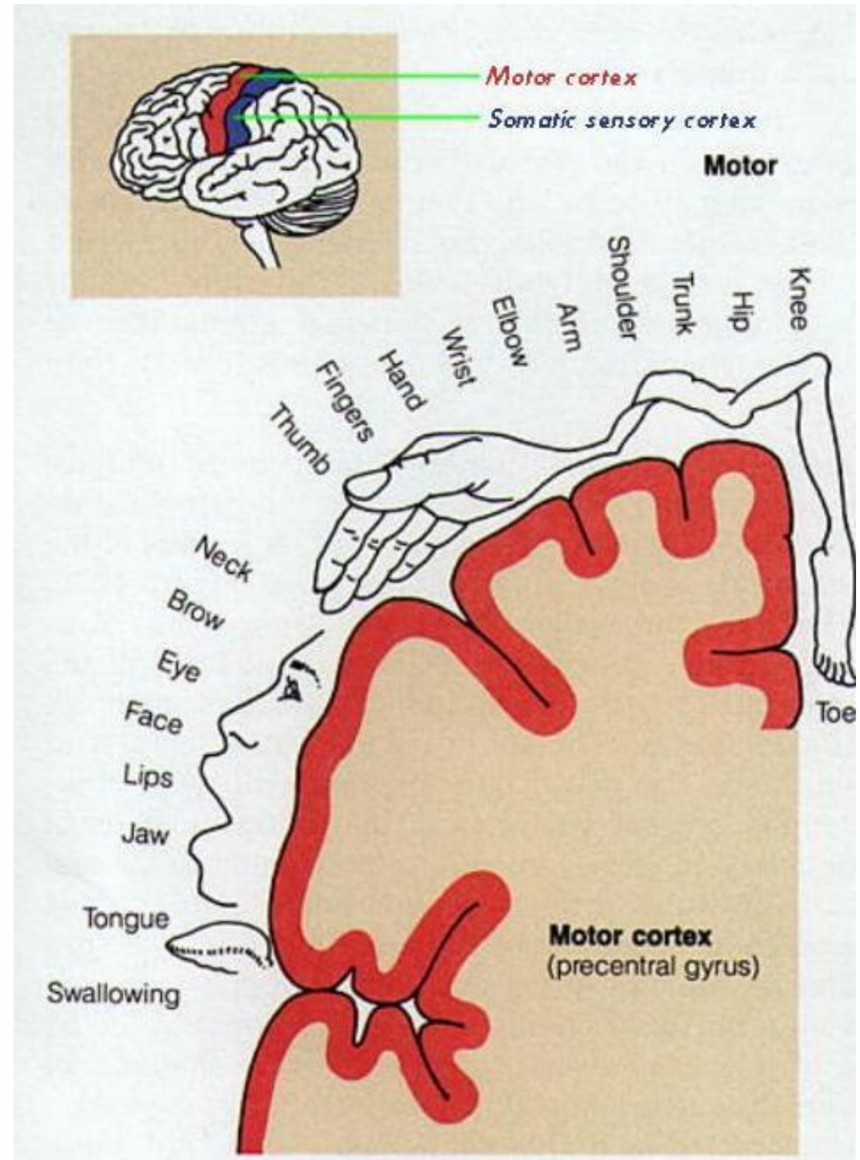

Ferrier (1874),
Penfield (1937)

What is coded by the motor cortex?

Muscle/joint activation

- Evarts (1968)
- Scott and Kalaska (1995)

Direction of movement

- Georgopoulos et al (1982)
- Moran and Schwartz (1999)

Complex motor sequences

- Graziano et al (2002)

Direction tuning in cortex

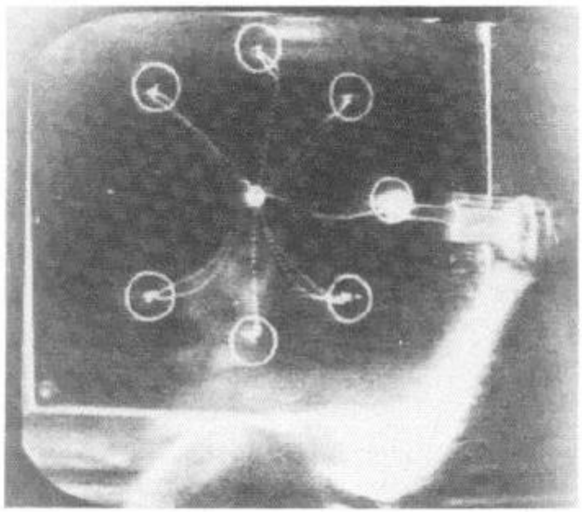

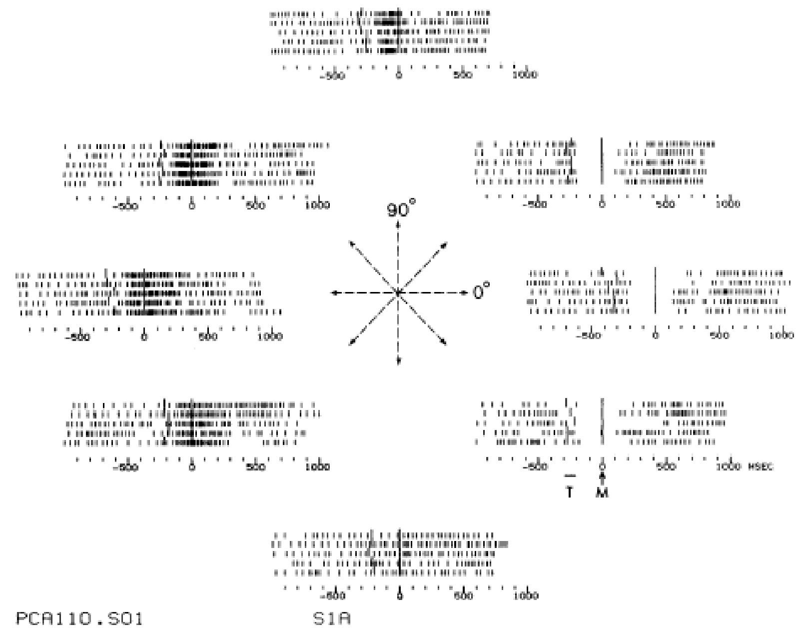

$$f(\boldsymbol{v}) = b_0 + b_x v_x + b_y v_y$$

$$= b_0 + \boldsymbol{b}^T \boldsymbol{v}$$

$$= b_0 + \|\boldsymbol{b}\| \|\boldsymbol{v}\| \cos \theta_{vb}$$

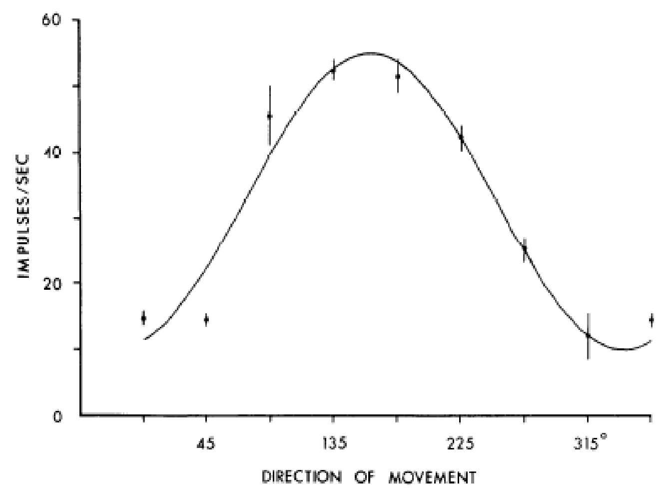

Georgopoulos et al. (1982)

Population vector - offline decoding

$$f_i(\boldsymbol{v}) = b_{i0} + b_{ix}v_x + b_{iy}v_y$$

$$= b_{i0} + \boldsymbol{bi}^T\boldsymbol{v}$$

Call $\boldsymbol{b_i}$ the '*preferred direction*' of neuron i

Decoded movement direction at time t is:

$$\boldsymbol{v}(t) = \sum_i^n (f_{\boldsymbol{i}}(t) - b_{i0})\, \boldsymbol{b_i}$$

Georgopoulos et al. (1986)

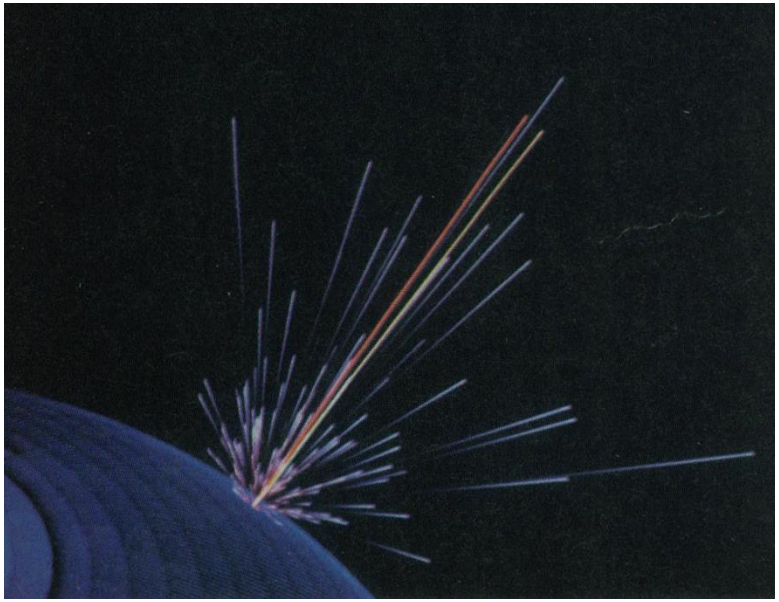

Real-time decoding *Wessberg, et al, Nature 2000*

MITnews

search

go

engineering science management architecture + planning humanities, arts, and social sciences campus press video connect

Monkey controls robotic arm using brain signals sent over Internet

Elizabeth A. Thomson, News Office

today's news

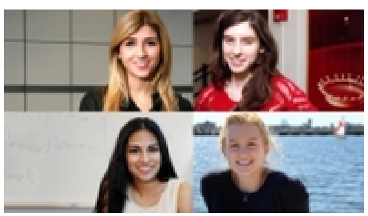

Four MIT seniors win Marshall Scholarships

Kate Koch, Colleen Loynachan, Kirin Sinha, and Grace Young will study for two

December 6, 2000

Share

Monkeys in North Carolina have remotely operated a robotic arm 600 miles away in MIT's Touch Lab -- using their brain signals.

The feat is based on a neural-recording system reported in the November 16 issue of Nature. In that system, tiny electrodes implanted in the animals' brains detected their brain signals as they controlled a robot arm to reach for a piece of food.

According to the scientists from Duke University Medical Center, MIT and the State University of New York (SUNY) Health Science Center, the new system could form the basis for a brain-machine interface that would allow paralyzed patients to control the movement of prosthetic limbs.

multimedia

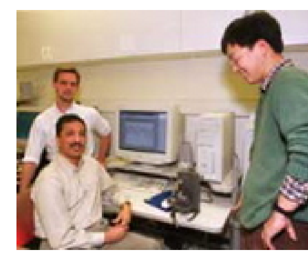

James Biggs, postdoctoral associate in the Research Lab of Electronics (left), Professor Mandayam Srinivasan, director of MIT's

PLAY TRAILER

Surrogates (2009)

TOMATOMETER | All Critics | Top Critics

39%

Average Rating: 5.4/10
Reviews Counted: 112
Fresh: 44 | Rotten: 68

Though it sports a slick look and feel, Surrogates fails to capitalize on a promising premise, relying instead on mindless action and a poor script.

AUDIENCE

38%

liked it

Average Rating: 3.1/5
User Ratings: 293,681

MY RATING

WANT TO SEE IT | NOT INTERESTED

Add a Review (Optional)

POST

http://www.youtube.com/watch?v=Zl_h9RaL0es

Closed-loop decoding

Differences between hand control and brain control

Taylor et al, Science 2002

Decoding in humans

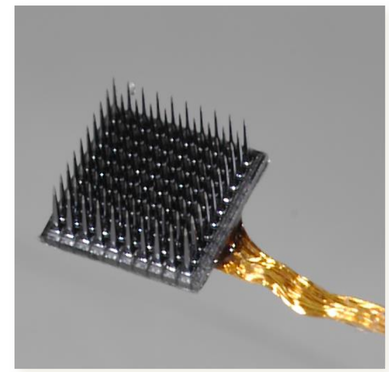

Blackrock array

Video1 News 6

Hochberg et al. (2006)

Using BCIs to studying learning

Optimal decoder (PDs) is different for hand and brain control

Is it possible to change the coding properties of neurons?

Operant conditioning to control neurons

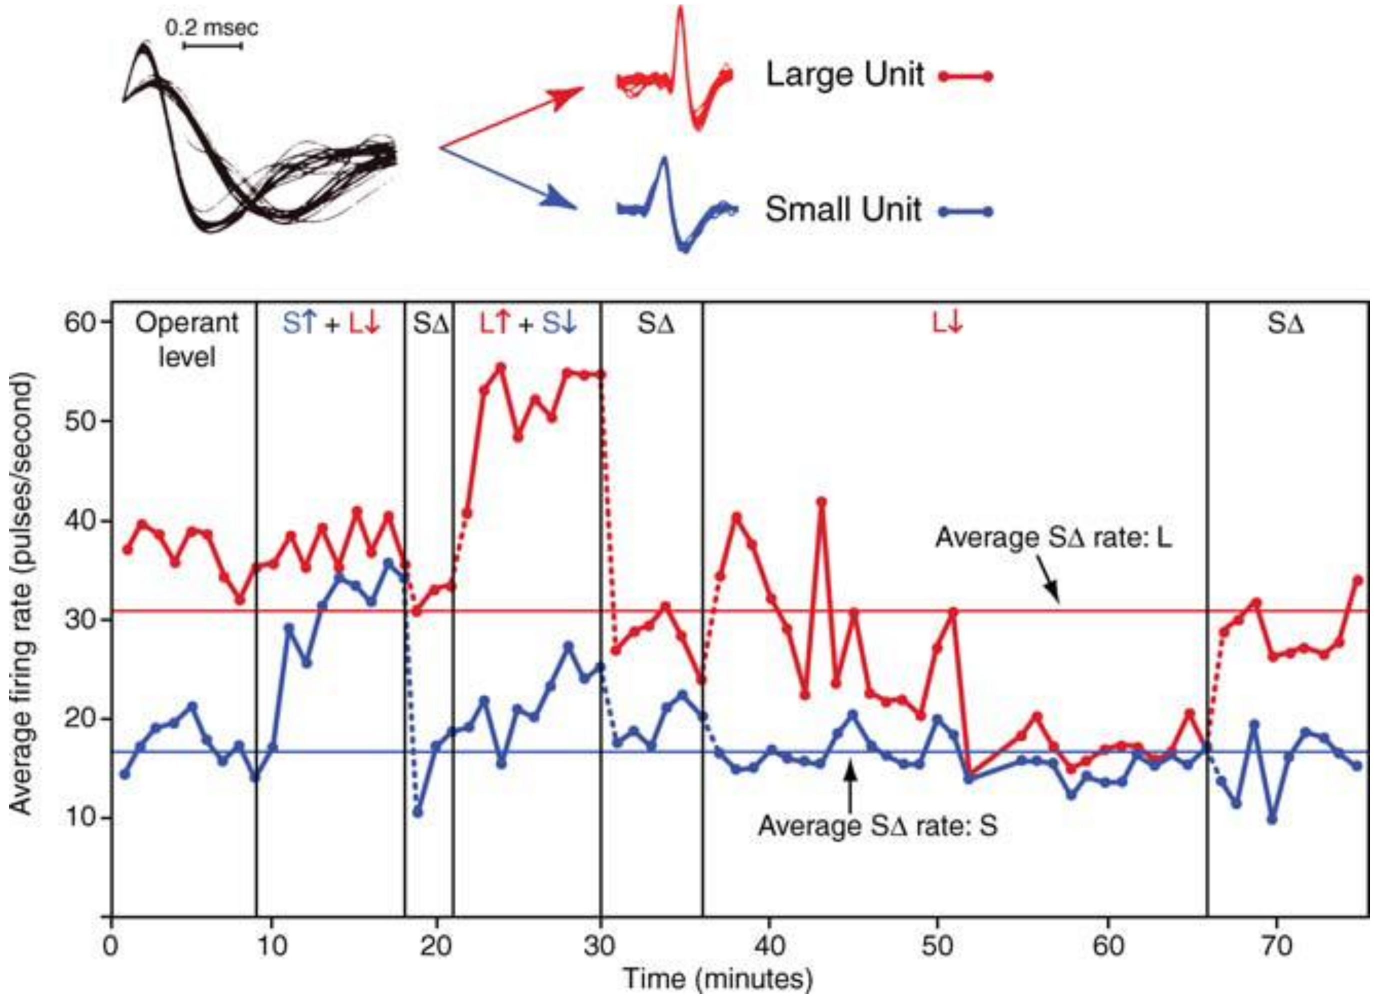

Fetz and Baker (1973)

Using BCIs to study learning

Jarosiewicz, Chase, Fraser, Villiste, Kass and Schwartz, PNAS, 2008

Ganguly and Carmena, PLoS Biology, 2009

Plasticity in BCIs: Jarosiewicz et al, 2008

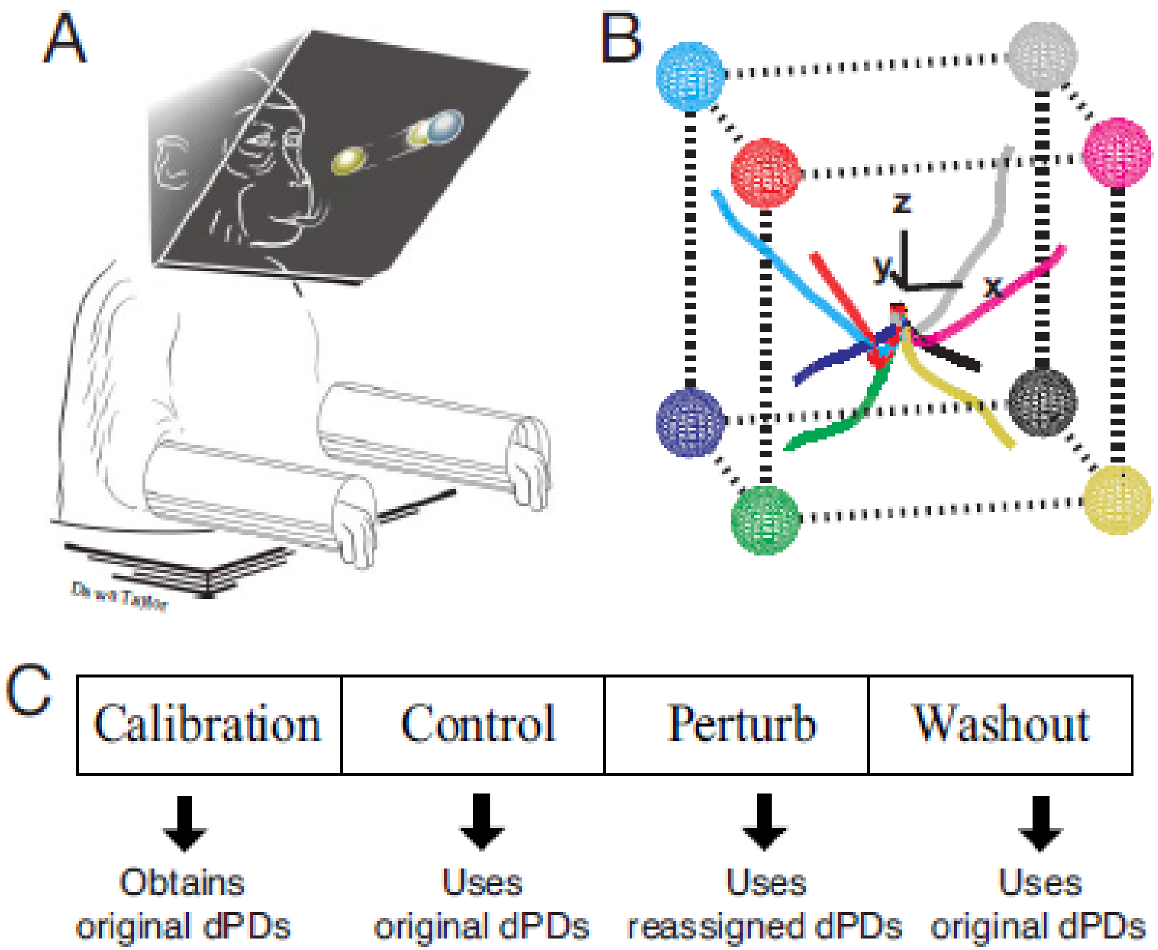

Expected effect of perturbation on cursor movement

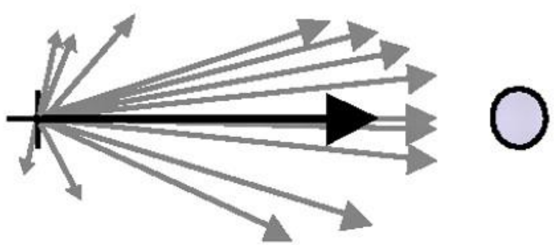

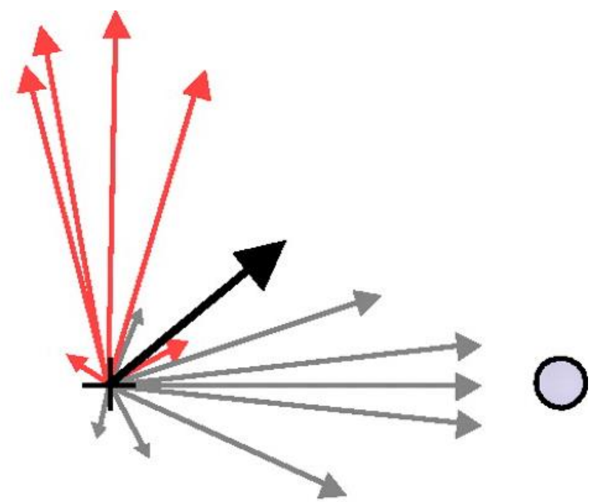

Effect of perturbation on cursor movement

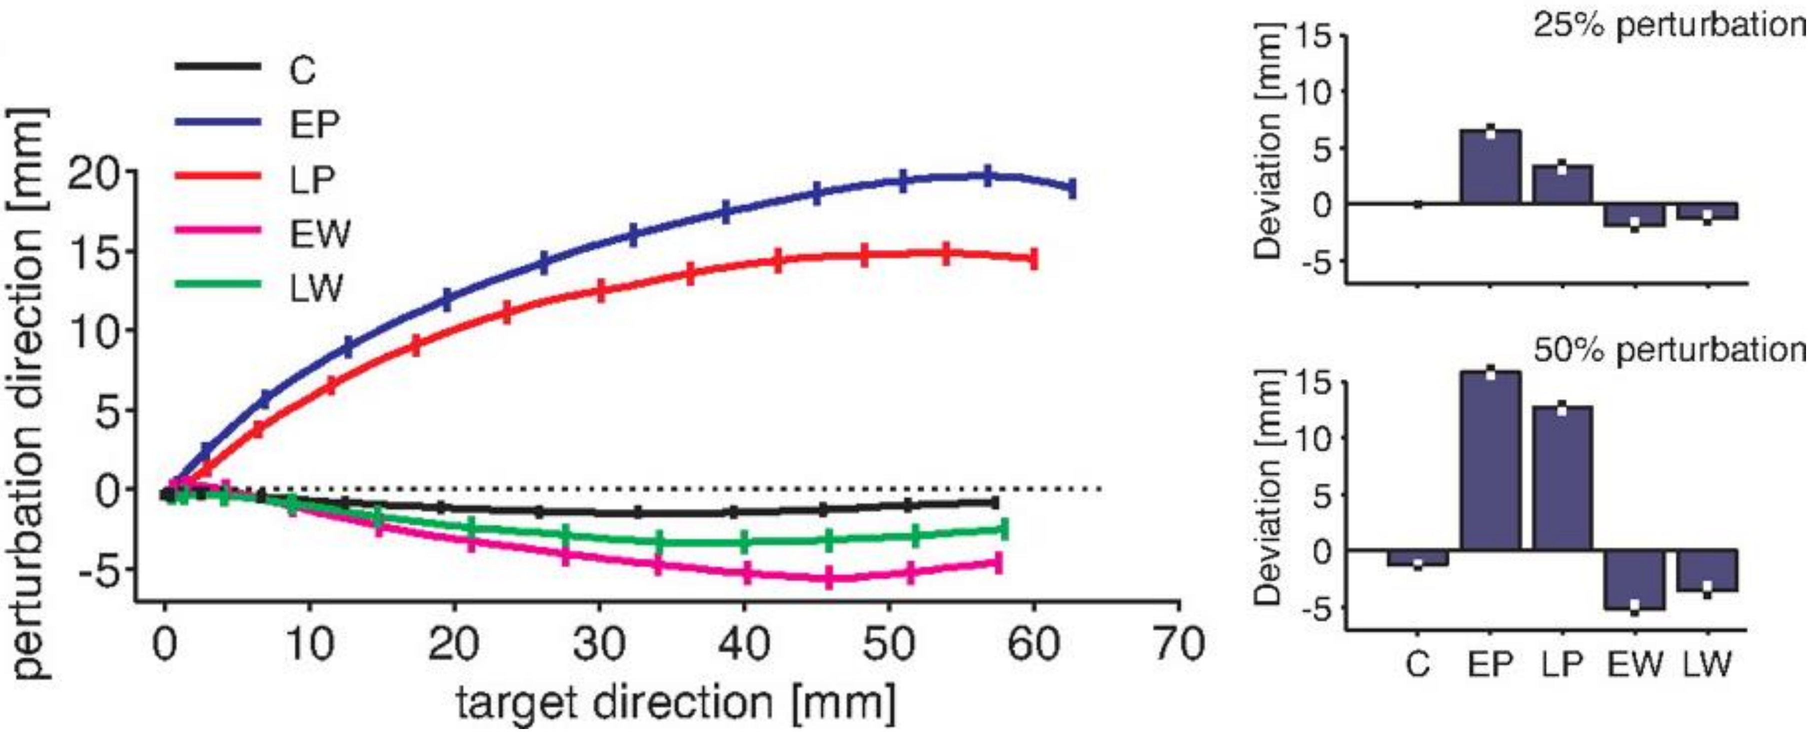

Possible neural compensation mechanisms

1) **<u>Re-aiming:</u>** the monkey could have aimed the cursor to offset the perturbation caused by the reassignment, disregarding the relative contributions of the rotated vs. non-rotated units to the error.

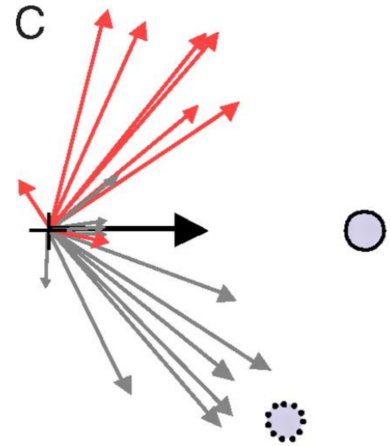

Possible neural compensation mechanisms

1) **Re-aiming:** the monkey could have aimed the cursor to offset the perturbation caused by the reassignment, disregarding the relative contributions of the rotated vs. non-rotated units to the error.

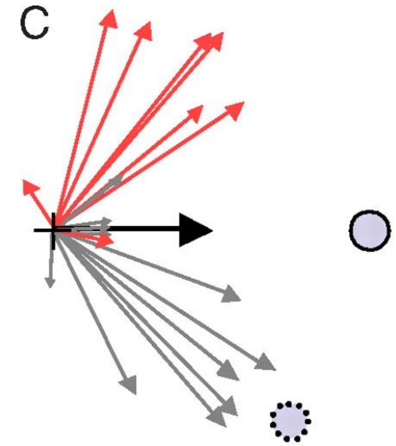

2) **Re-weighting:** the rotated units could have suppressed their contribution to the population vector by firing at baseline rate everywhere; i.e. by **decreasing their modulation depths**

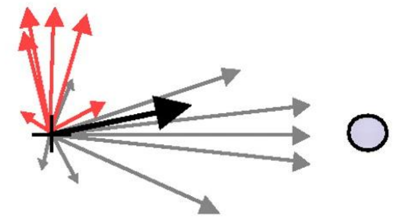

Possible neural compensation mechanisms

1) **Re-aiming:** the monkey could have aimed the cursor to offset the perturbation caused by the reassignment, disregarding the relative contributions of the rotated vs. non-rotated units to the error.

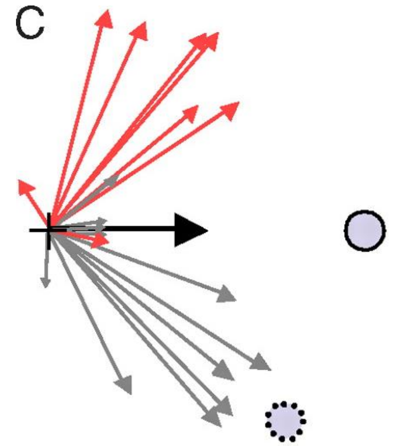

2) **Re-weighting:** the rotated units could have suppressed their contribution to the population vector by firing at baseline rate everywhere; i.e. by **decreasing their modulation depths**

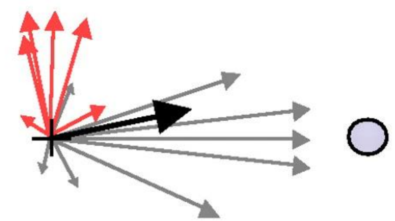

3) **Re-mapping:** the rotated units could have **shifted their actual PDs** (activation functions) **toward their reassigned dPDs** (labels).

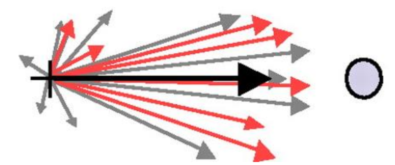

Testing for these possibilities

$$f = b_0 + b_x v_x + b_y v_y + b_z v_z + \varepsilon$$

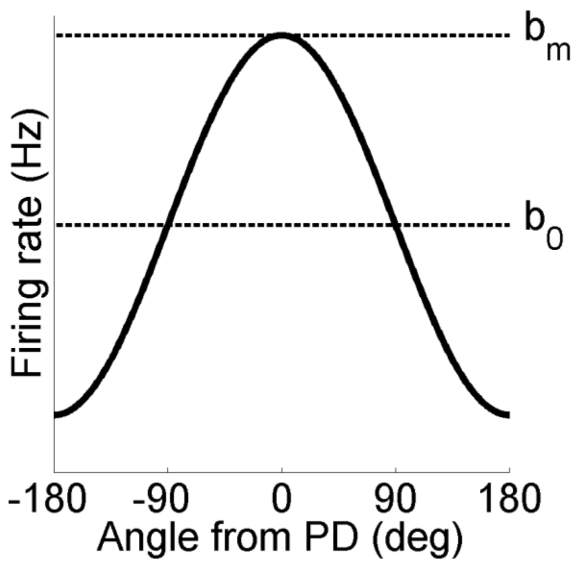

Recalculate (offline) neuron's PD during the perturbation session

Evidence for re-weighting

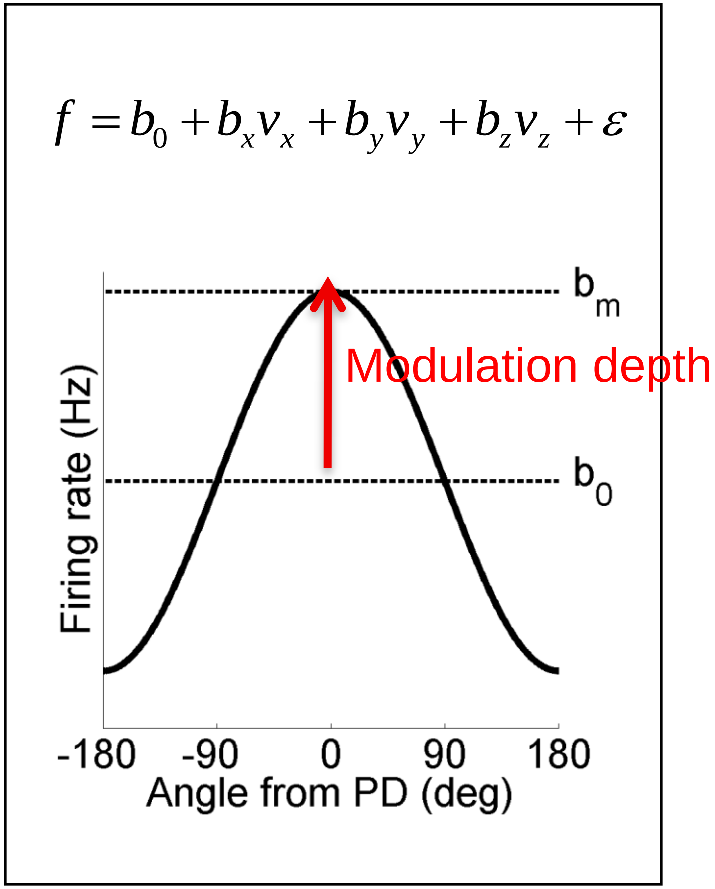

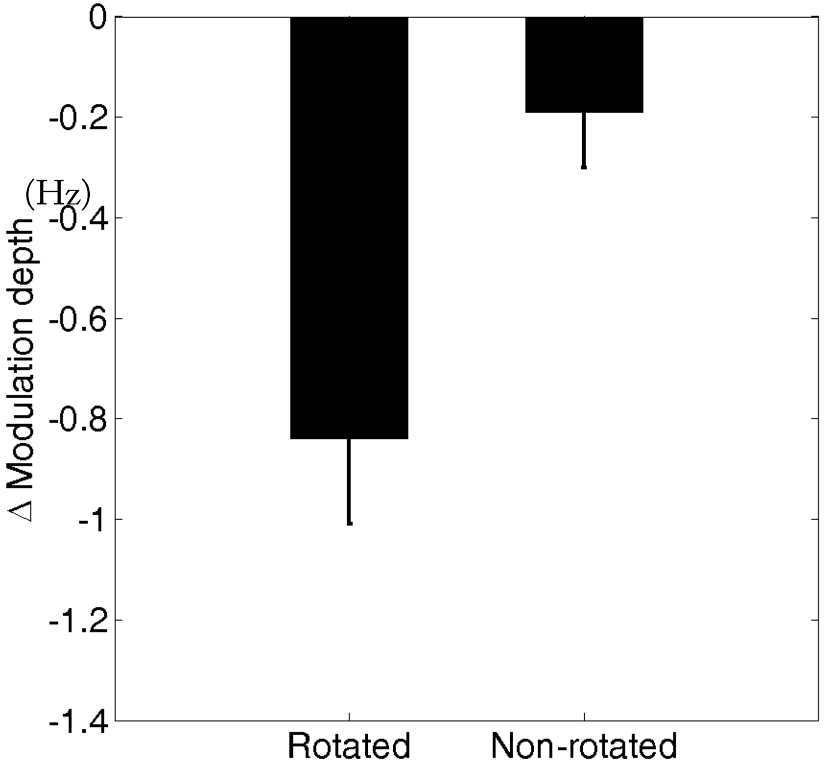

Evidence for re-aiming and re-mapping

Re-aiming appears as a shift in **all** PDs in the direction of the perturbation:

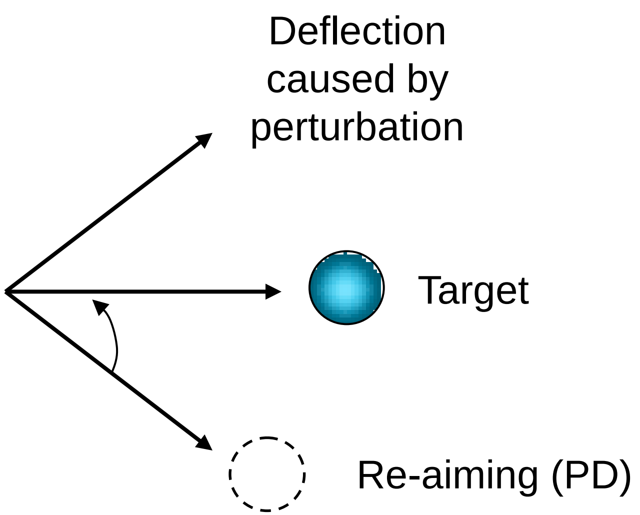

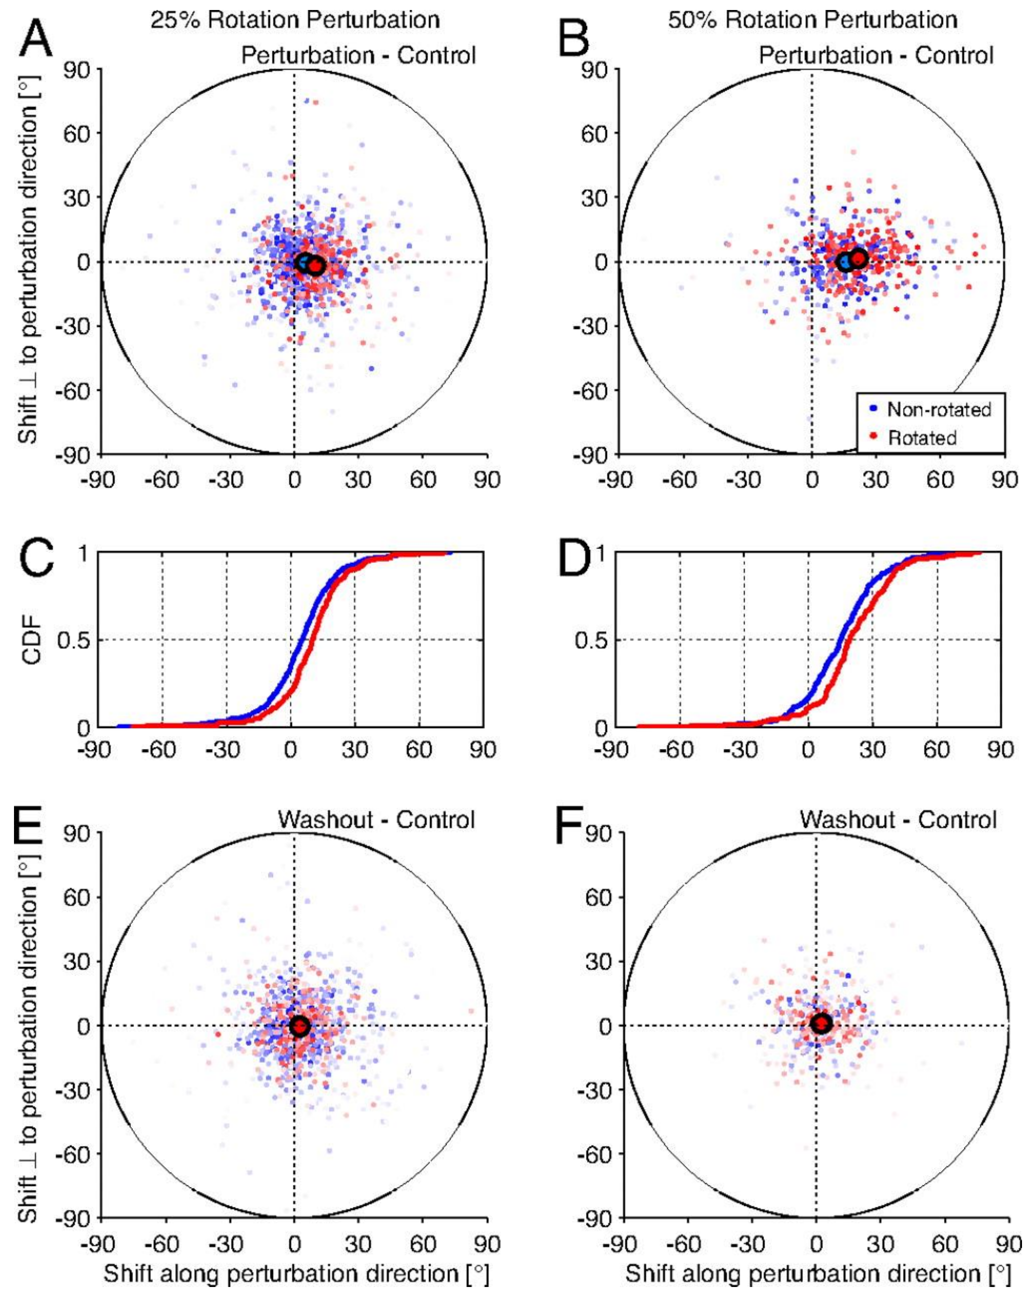

Summary: Jarosiewizc et al

Improved brain control of the cursor after the perturbation was due to: re-weighting, re-aiming, and re-mapping

Ganguly and Carmena 2009

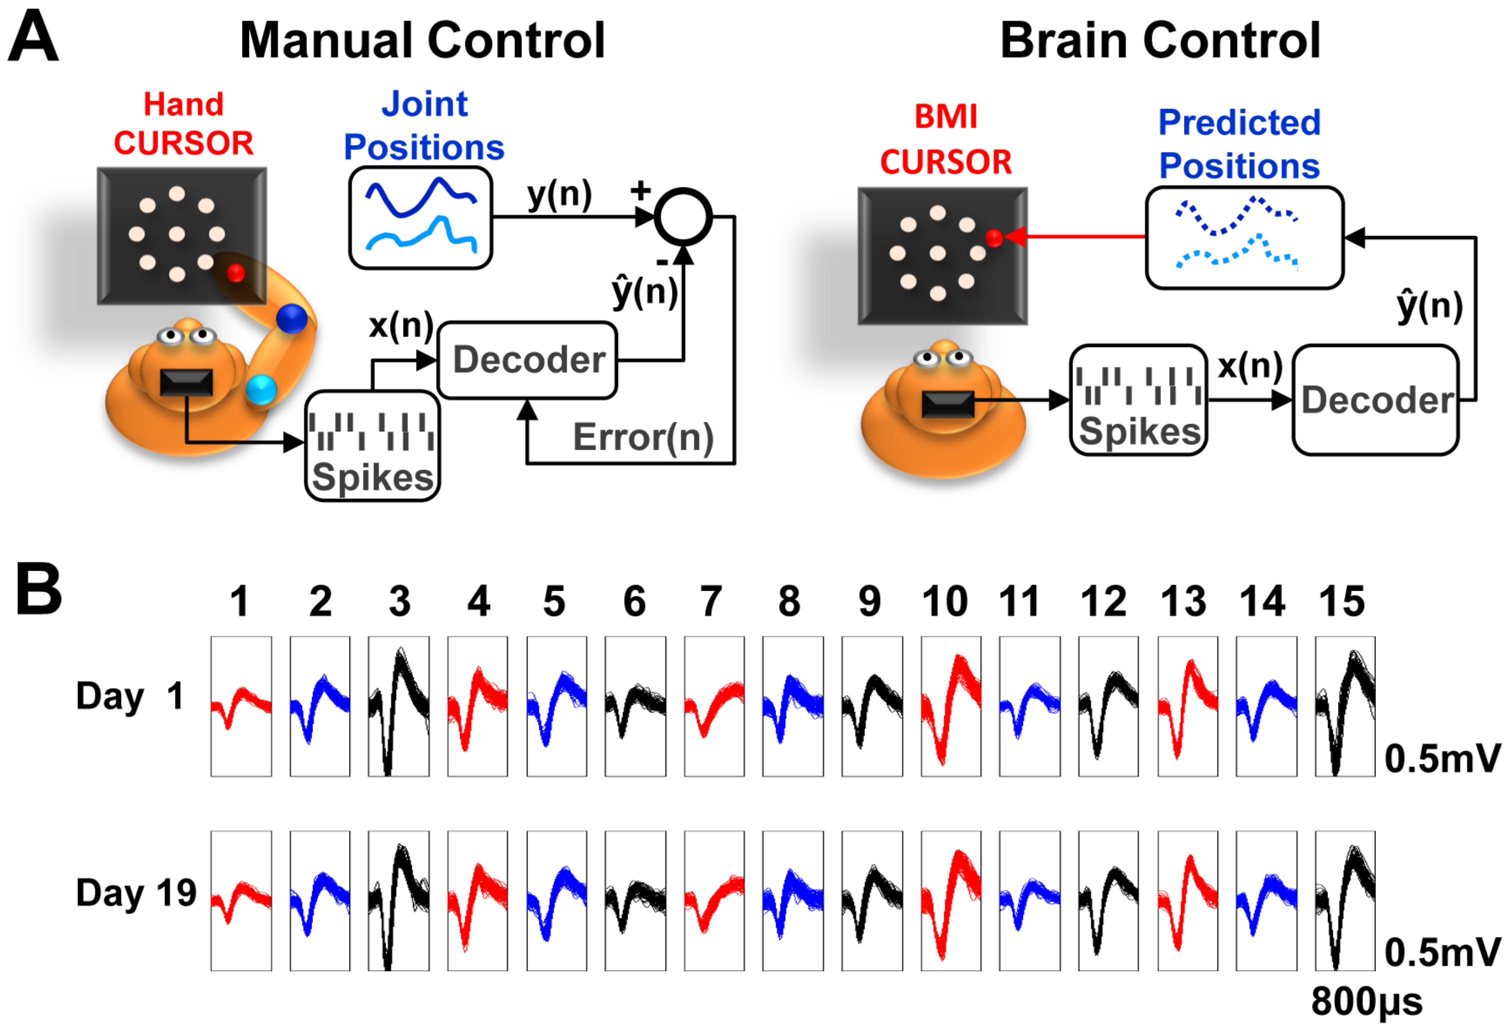

Performance increased over days

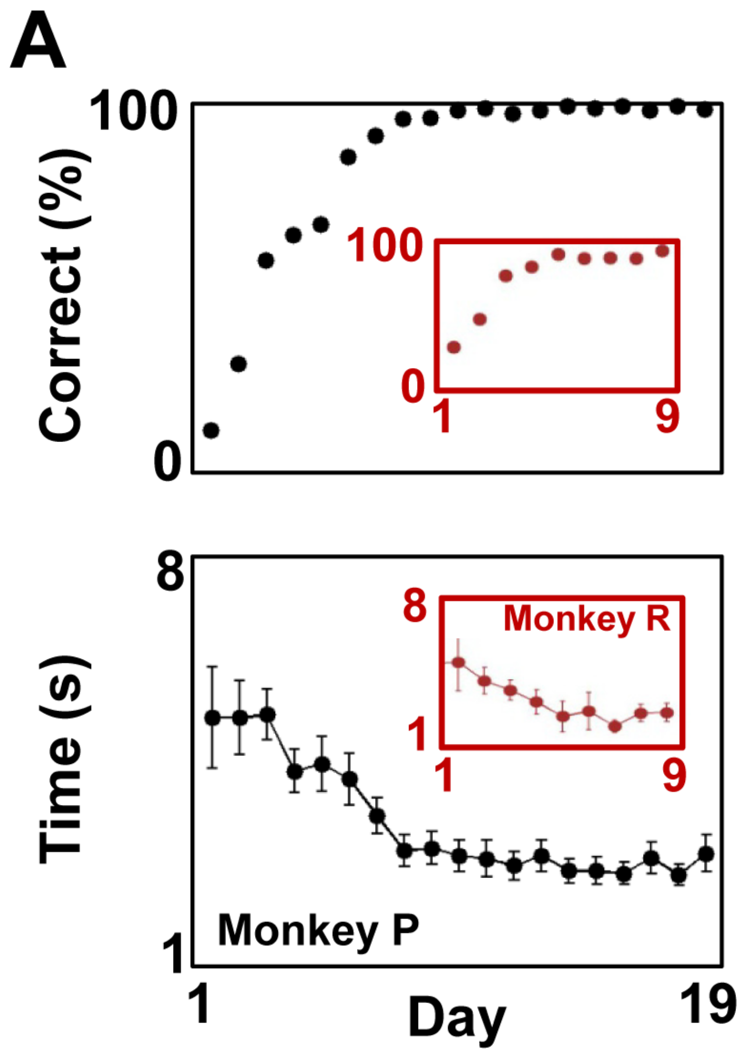

Accuracy

Speed

Neural activity formed a stable map

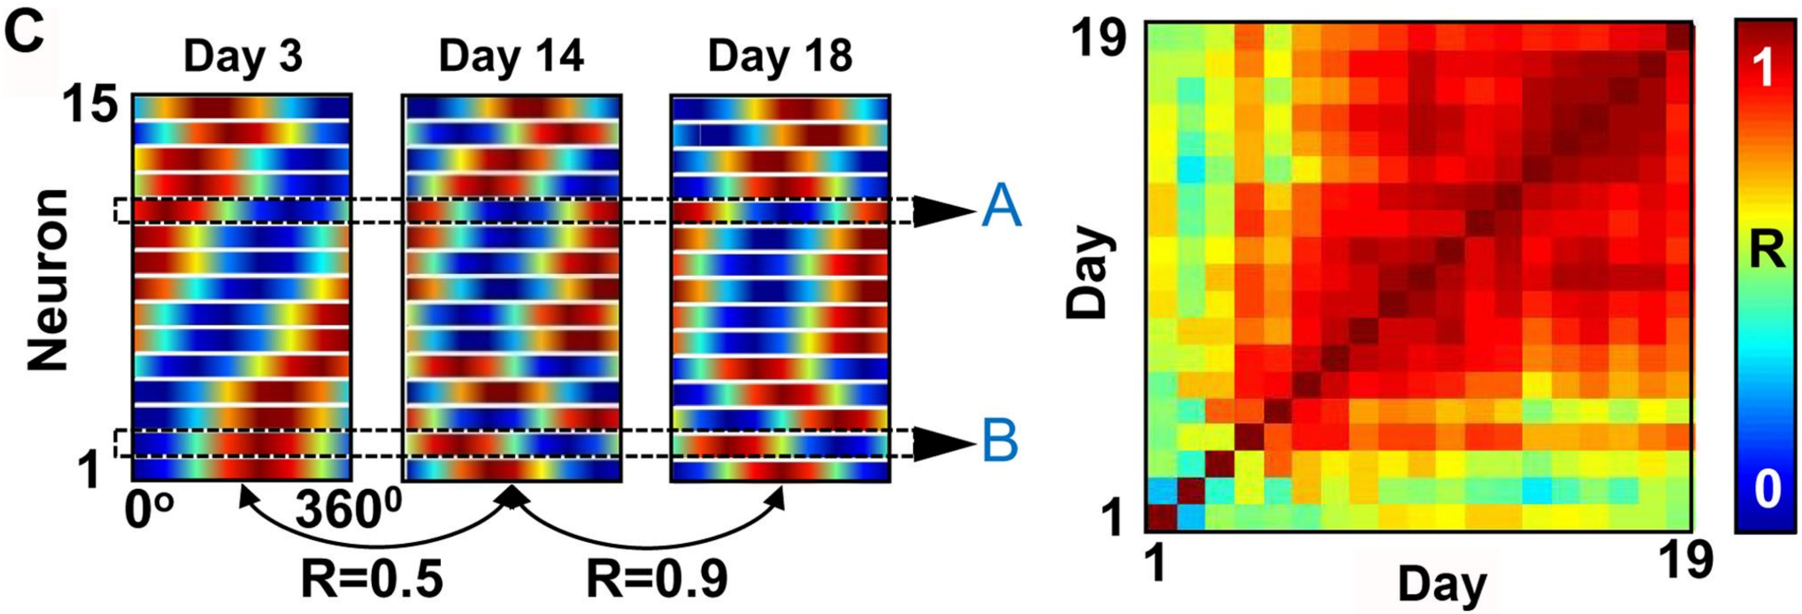

There was a high correlation between the increase in decoding performance and the neural activity stabilization

Using a random decoding

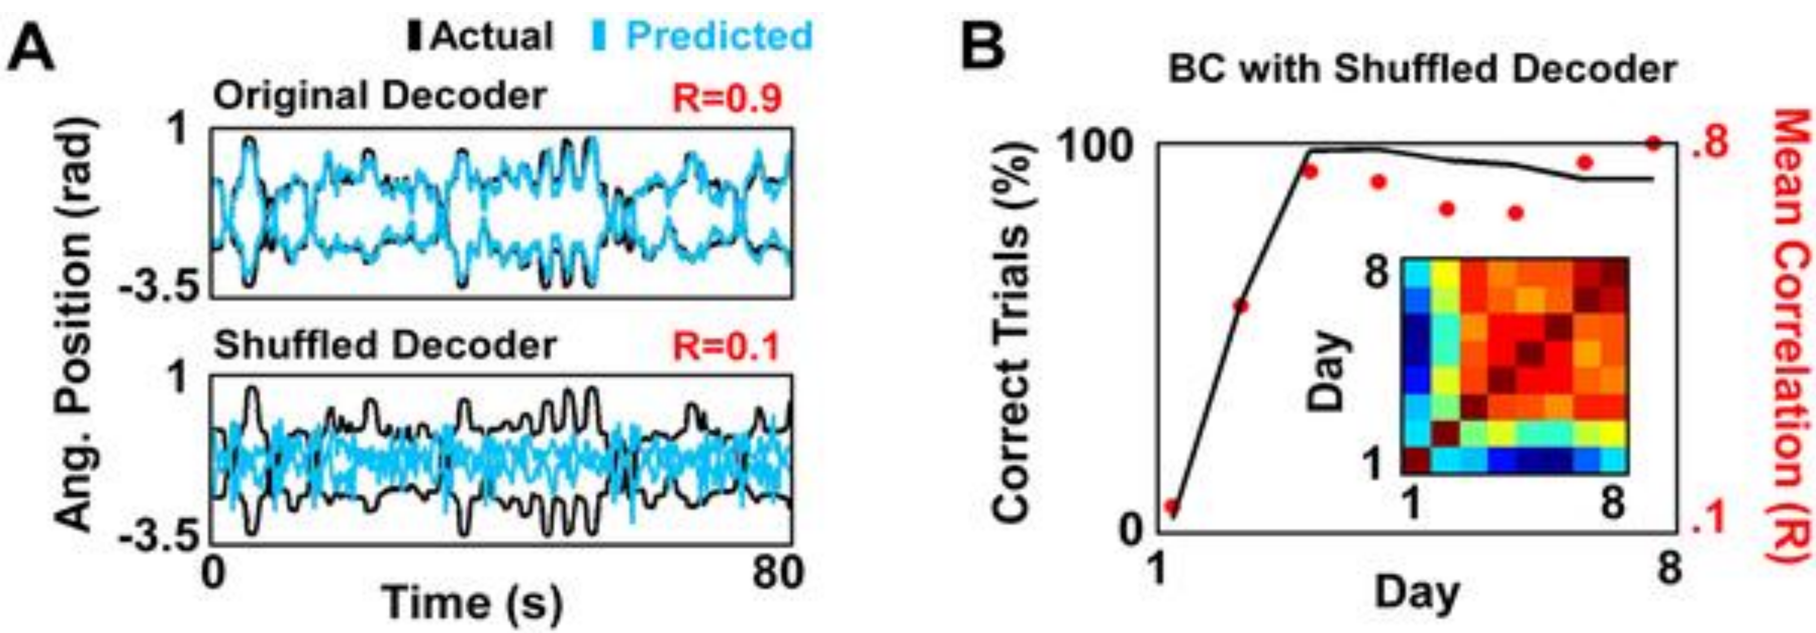

A Created a random decoder.

- Performance on motor control was poor on day 1

B Performance on the random decoding improved over time

Switching between two different decoders

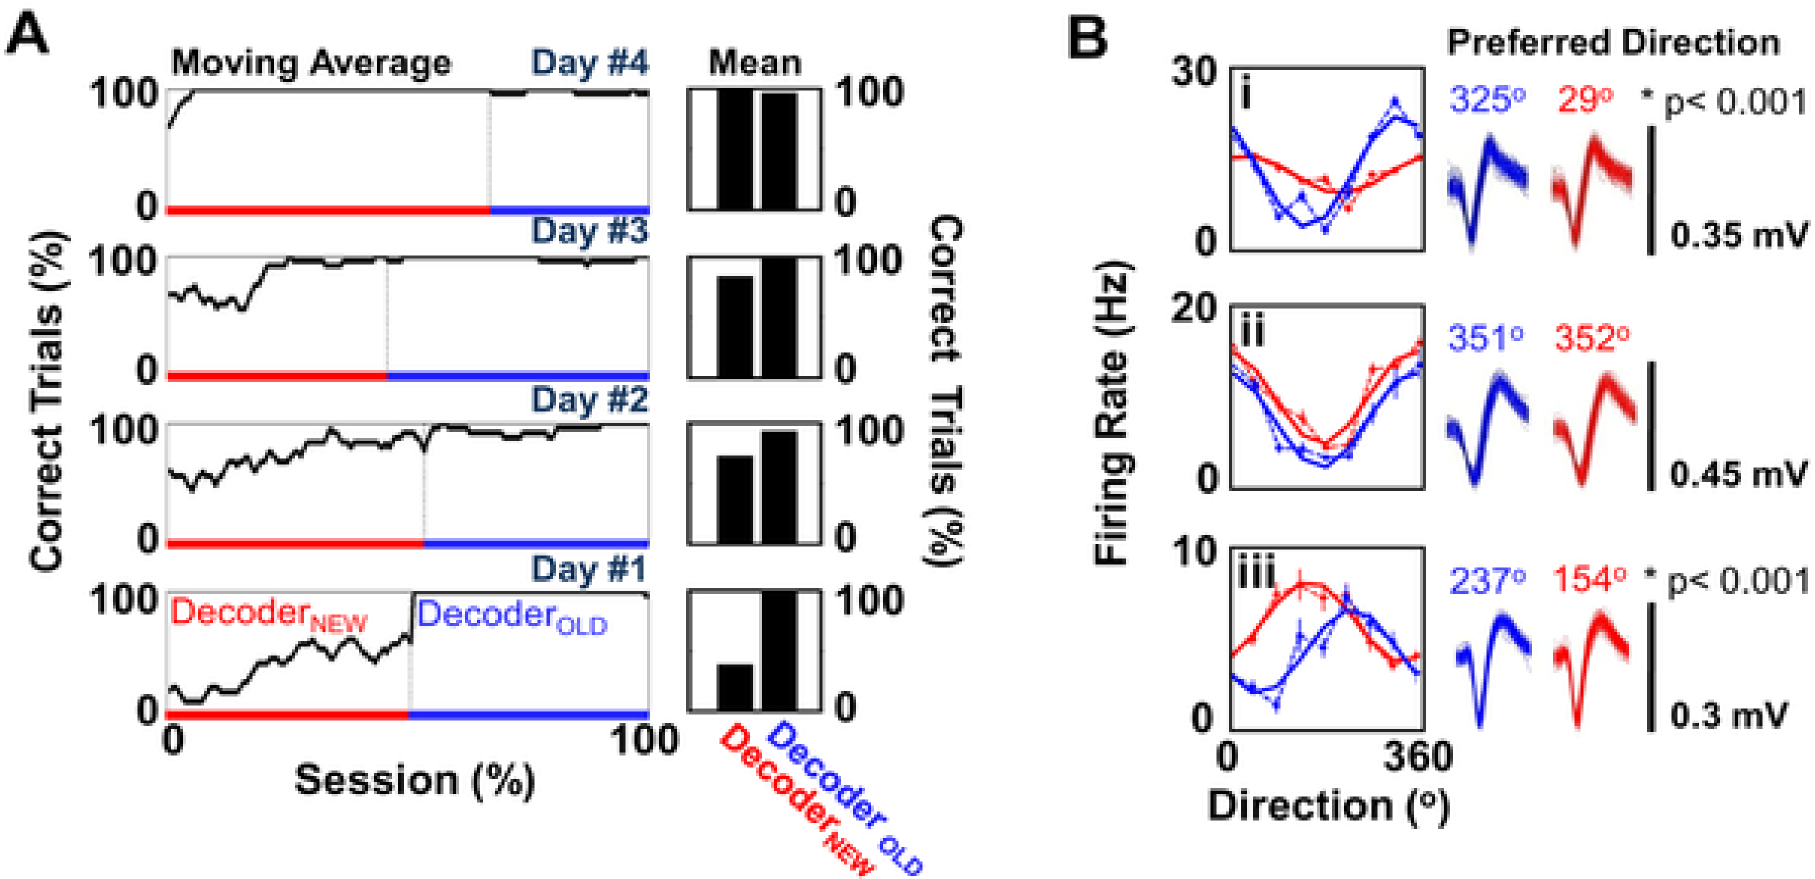

Summary: BCI to study plasticity

- It is possible to decode neural activity to control external devices
  - i.e., closed loop brain computer interfaces work

- Neurons can change their tuning properties to improve their performance on BCI tasks

- Perhaps reinforcement learning mechanisms are involved (see Yael Niv's talk)
  - Also see:
    Koralek, et al, Nature, 2012

Outline

- BCI to study learning in the motor cortex
- MVPA to study vision in human fMRI data
- Population decoding to study high level learning and vision in macaque monkeys

How MVPA works (Haxby et al, Science 2001)

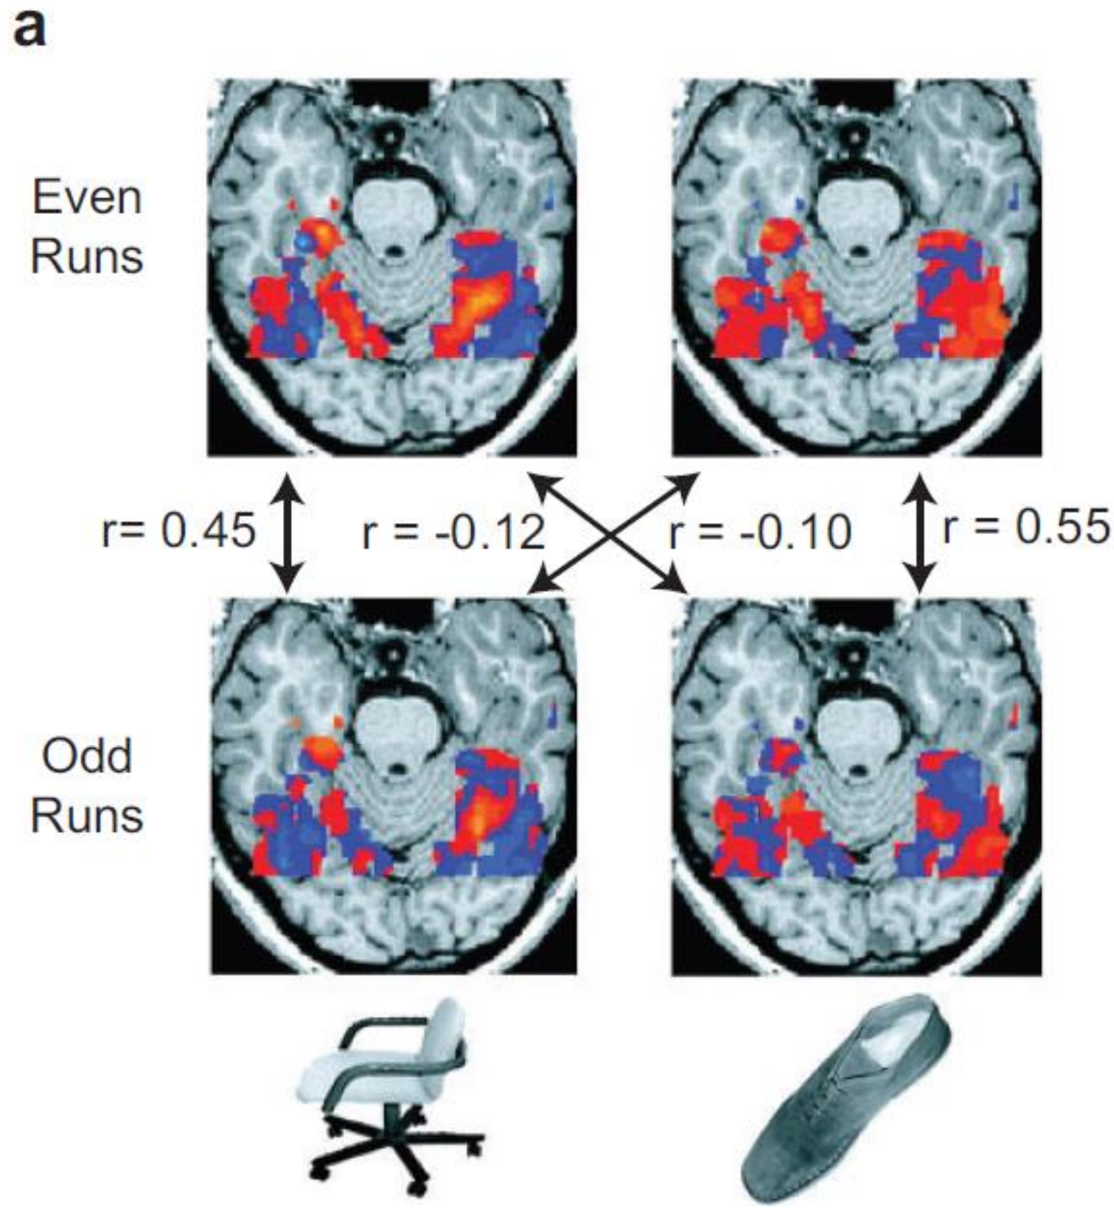

Detailed visual information can be extracted

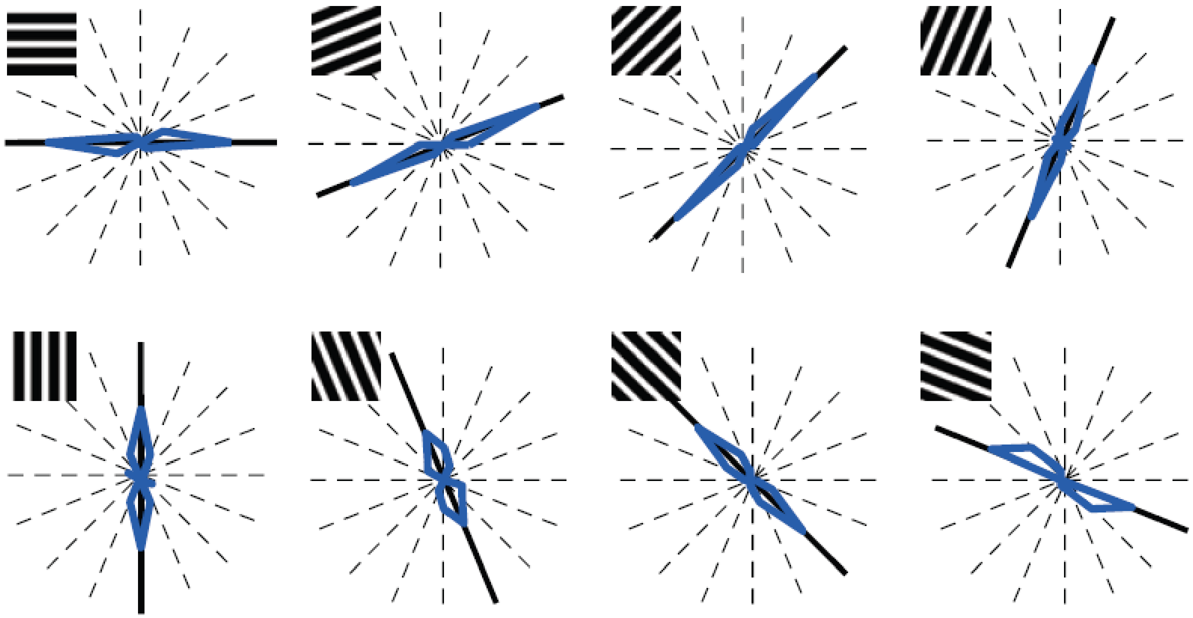

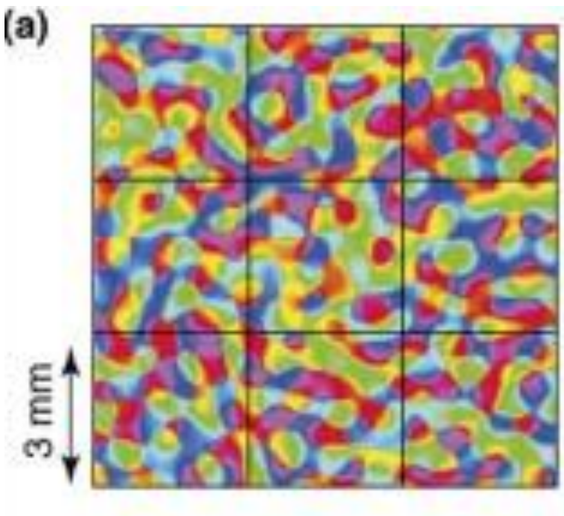

Kamitani and Tong (2005)

Using encoding models of visual information

- Encoding models can estimate voxel parameters (e.g., voxel's retinotopic location, spatial frequency and orientation)
  - Kay et al, 2008, was able to decode 820 of 1000 novel images correctly
  - Miyawaki et al, 2008 could reconstruct images

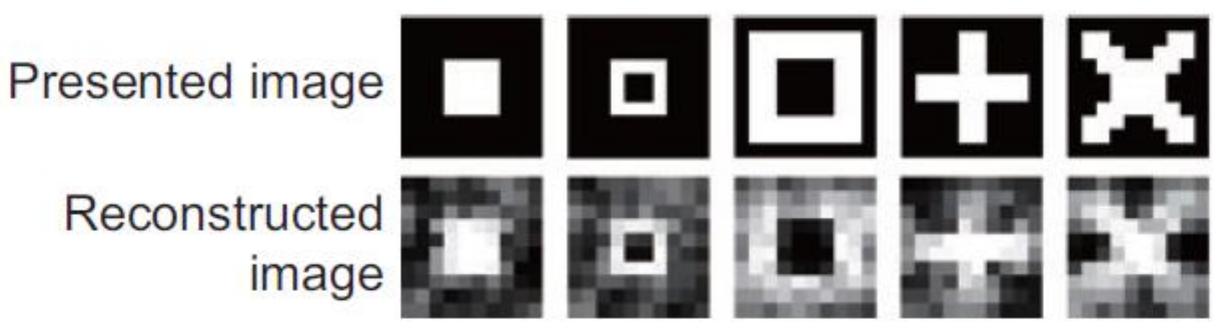

Nishimoto et al, Current Biology, 2011

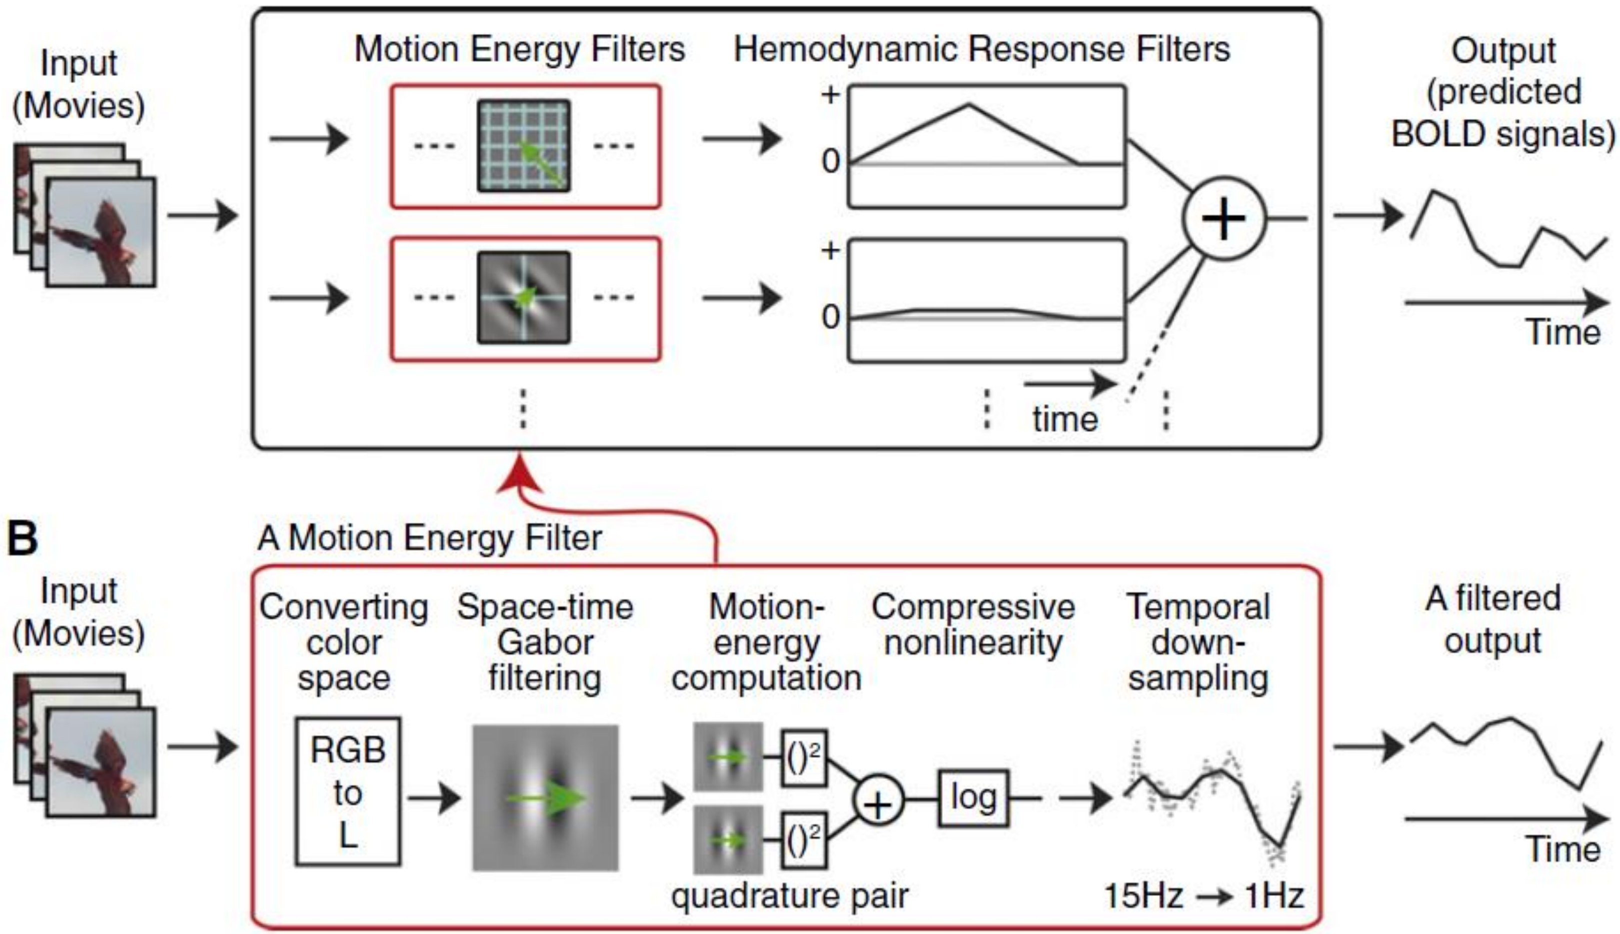

http://www.youtube.com/watch?v=HVL8GrUs_-E

Beyond detailed visual information

- Decoding what people are:
  - Attending to
  - Imagining
  - Recalling from memories
  - Assessing semantic content of nouns

Decoding what subjects are attending to

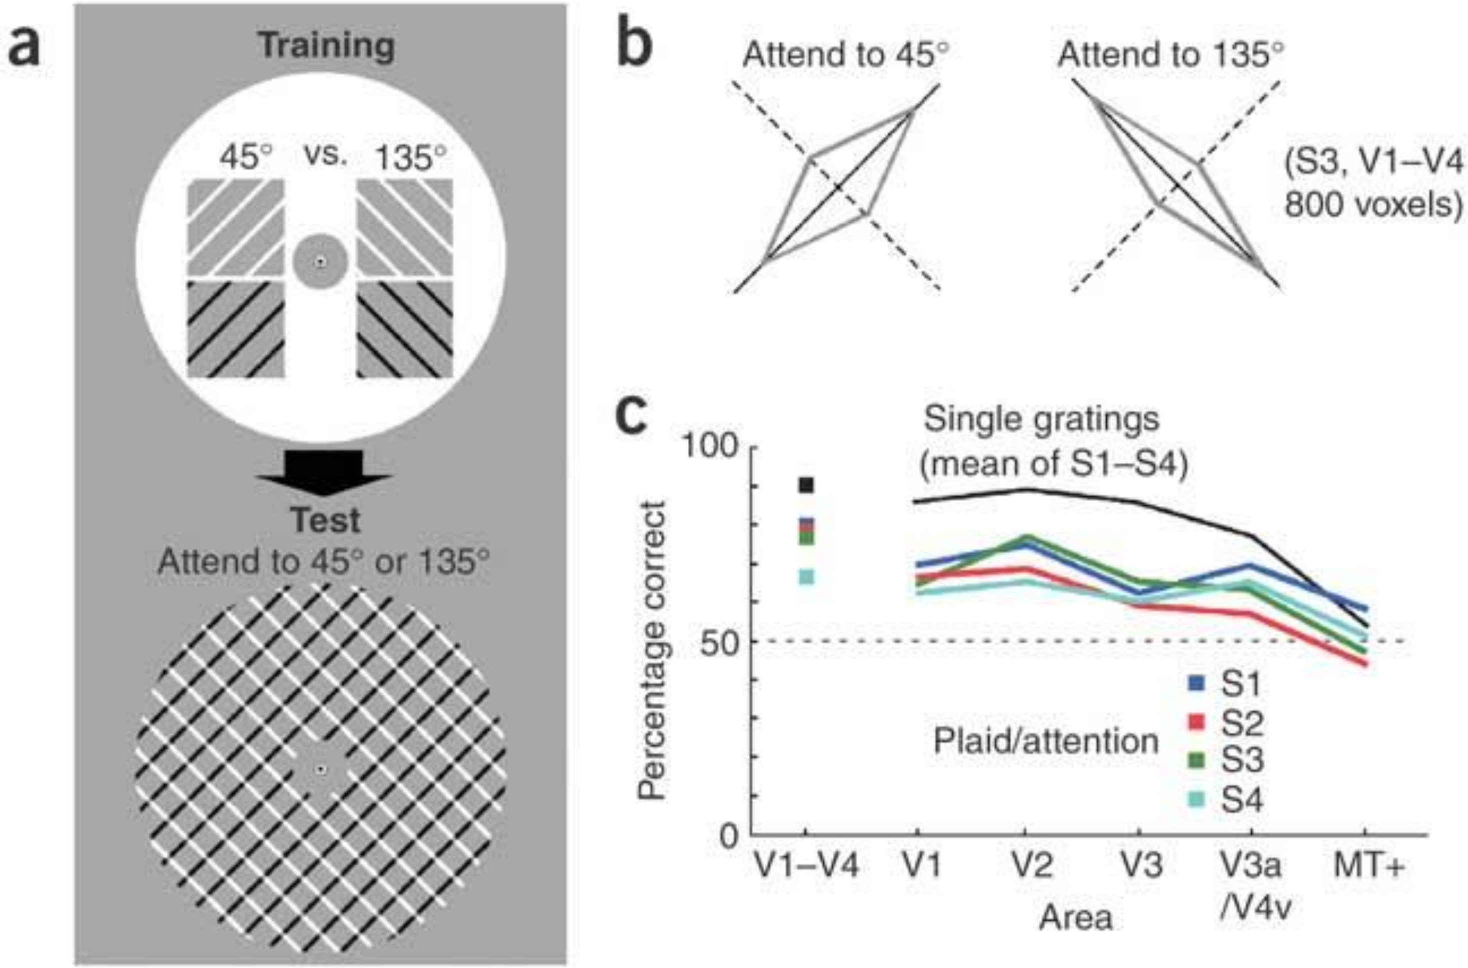

Kamitani and Tong (2005)

Decoding imagined categories

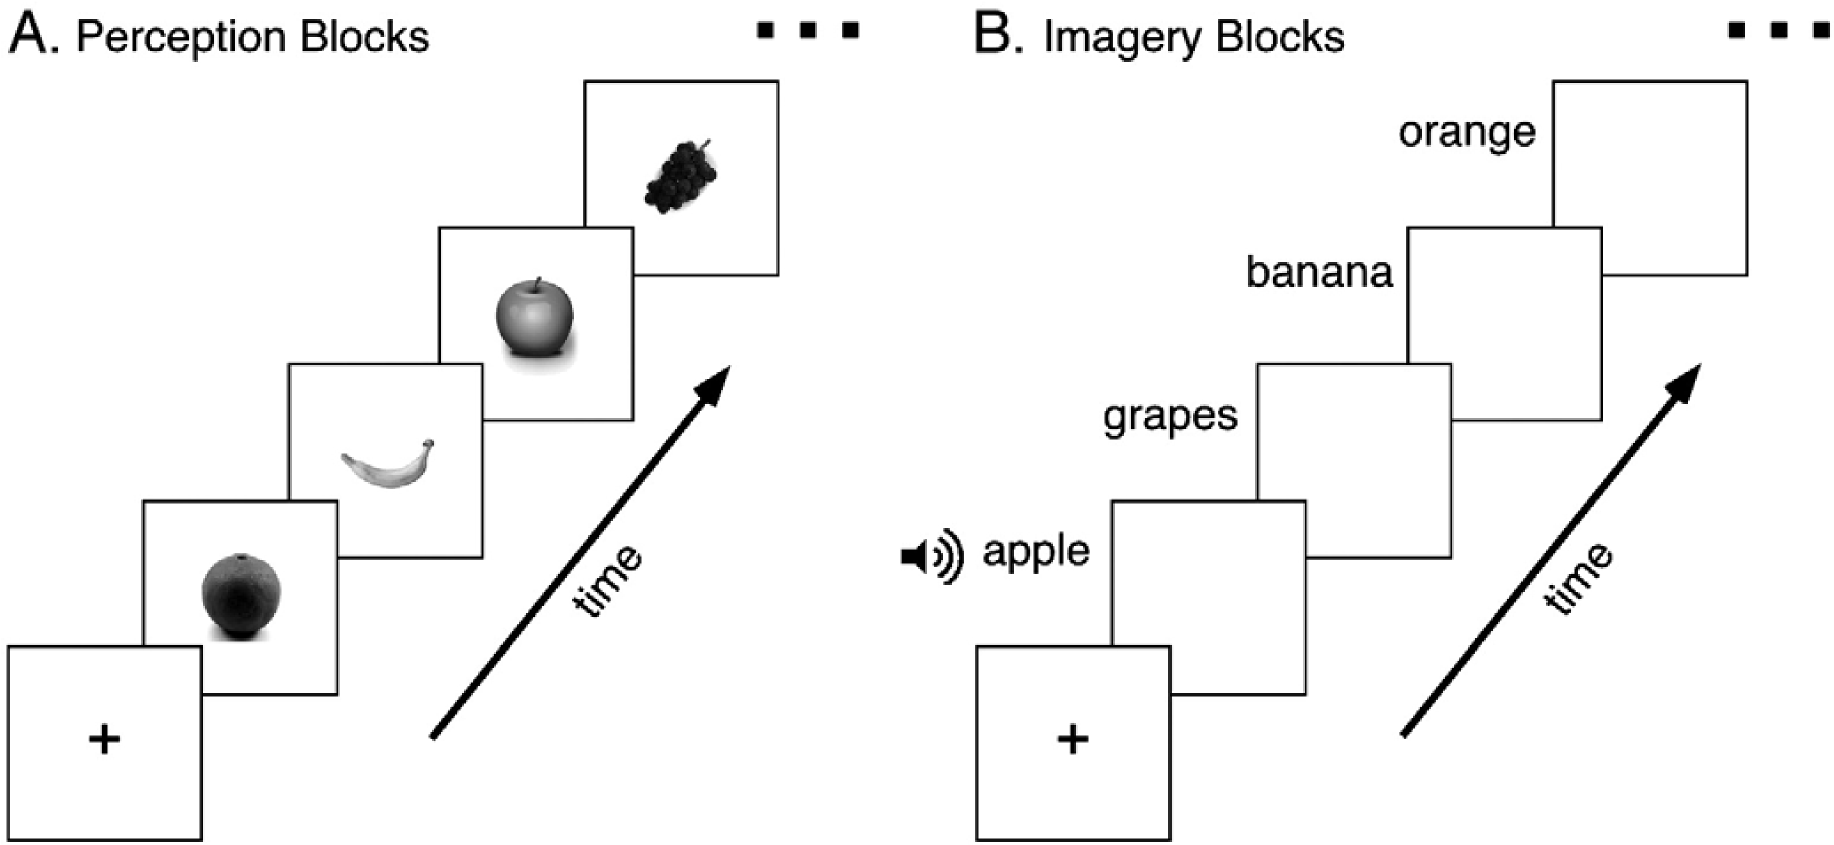

Reddy et al (2010)

Decoding imagined categories

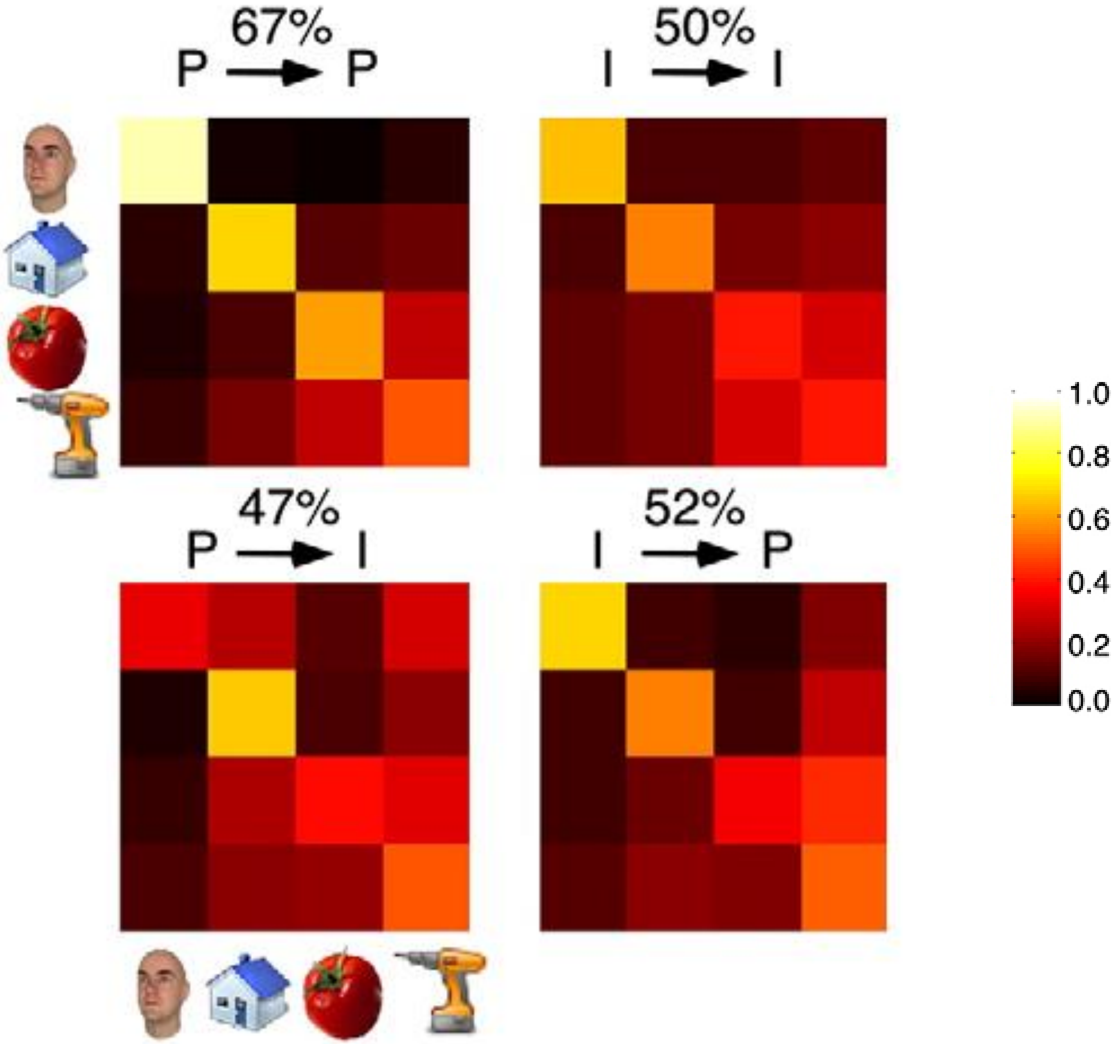

Reddy et al (2010)

Recalled episodic memory

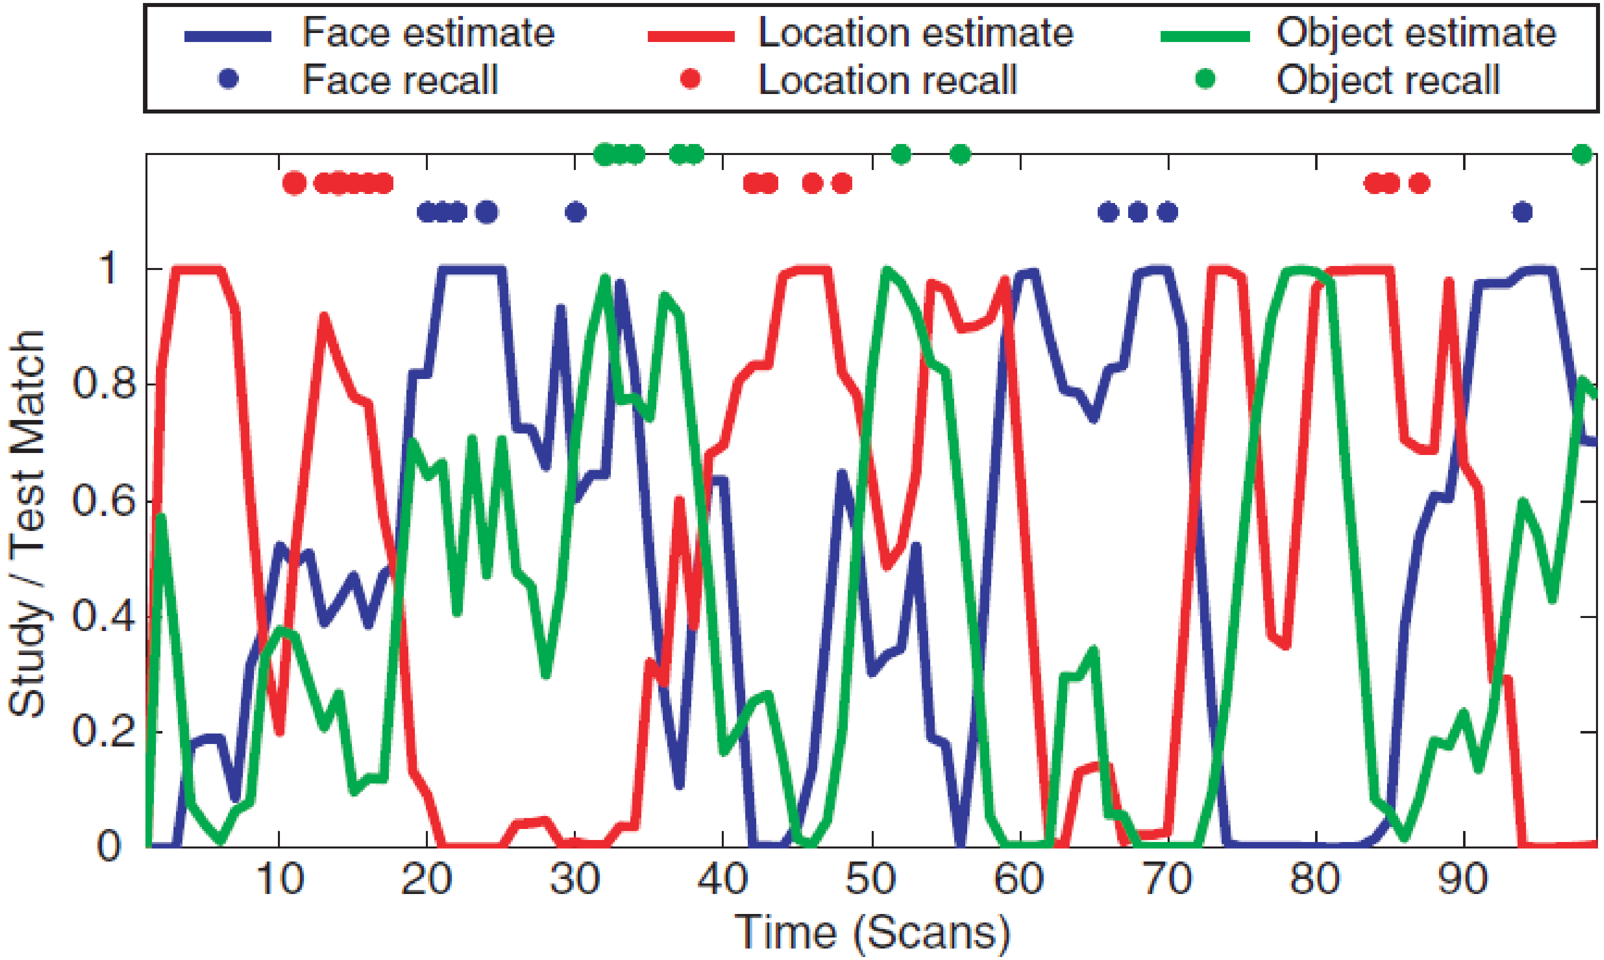

Polyn et al (2005)

Predicting novel nouns

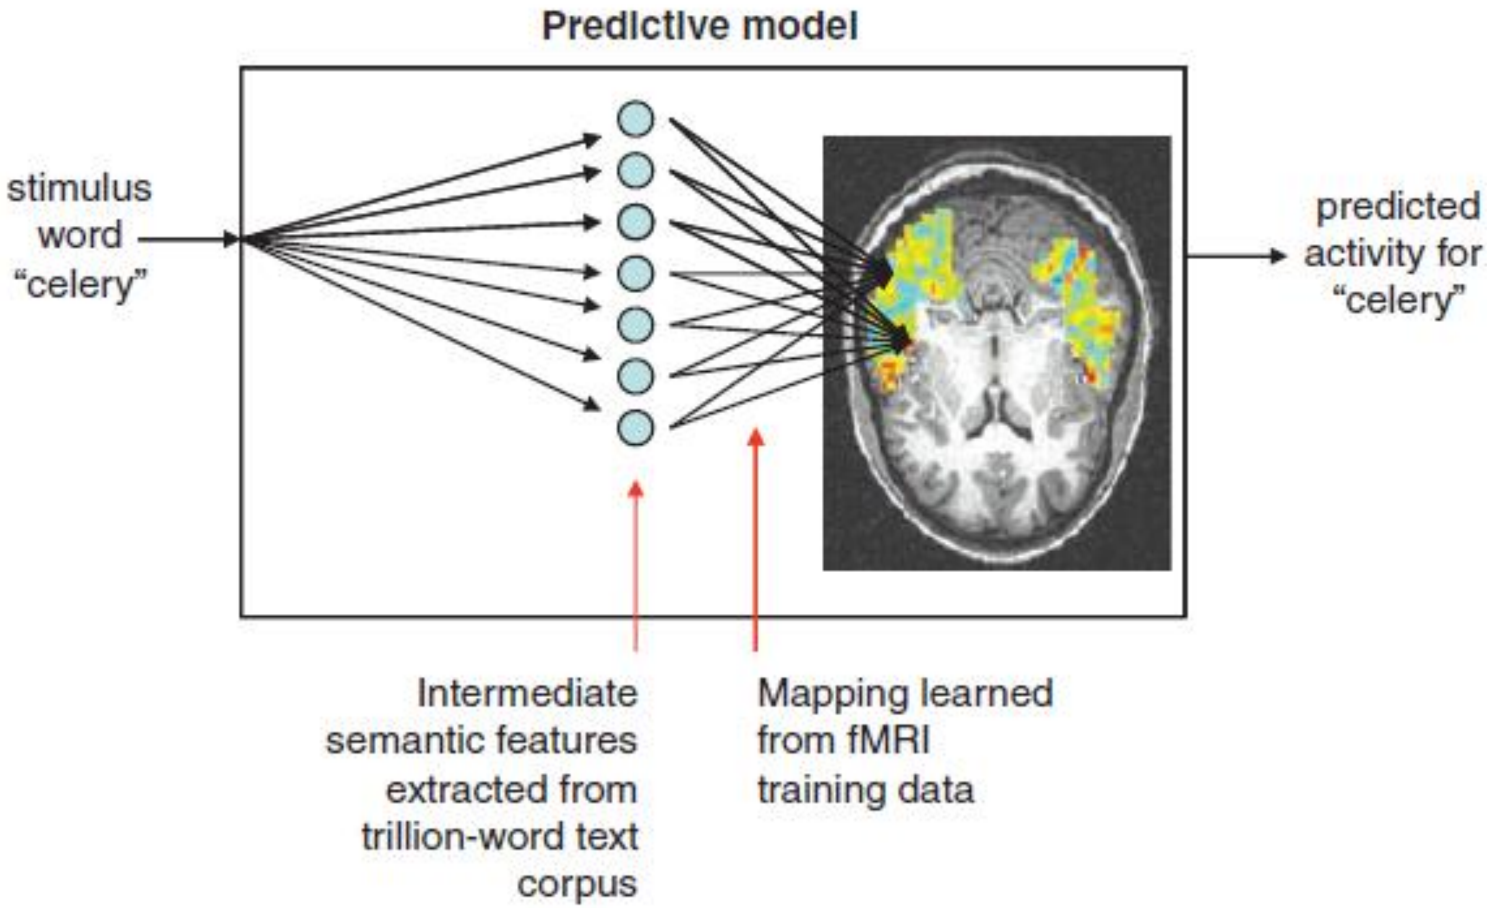

$$y_v(w) = \sum_{i=1}^{25} c_{iv} f_i(w)$$

Mitchell et al (2008)

Predicting novel nouns

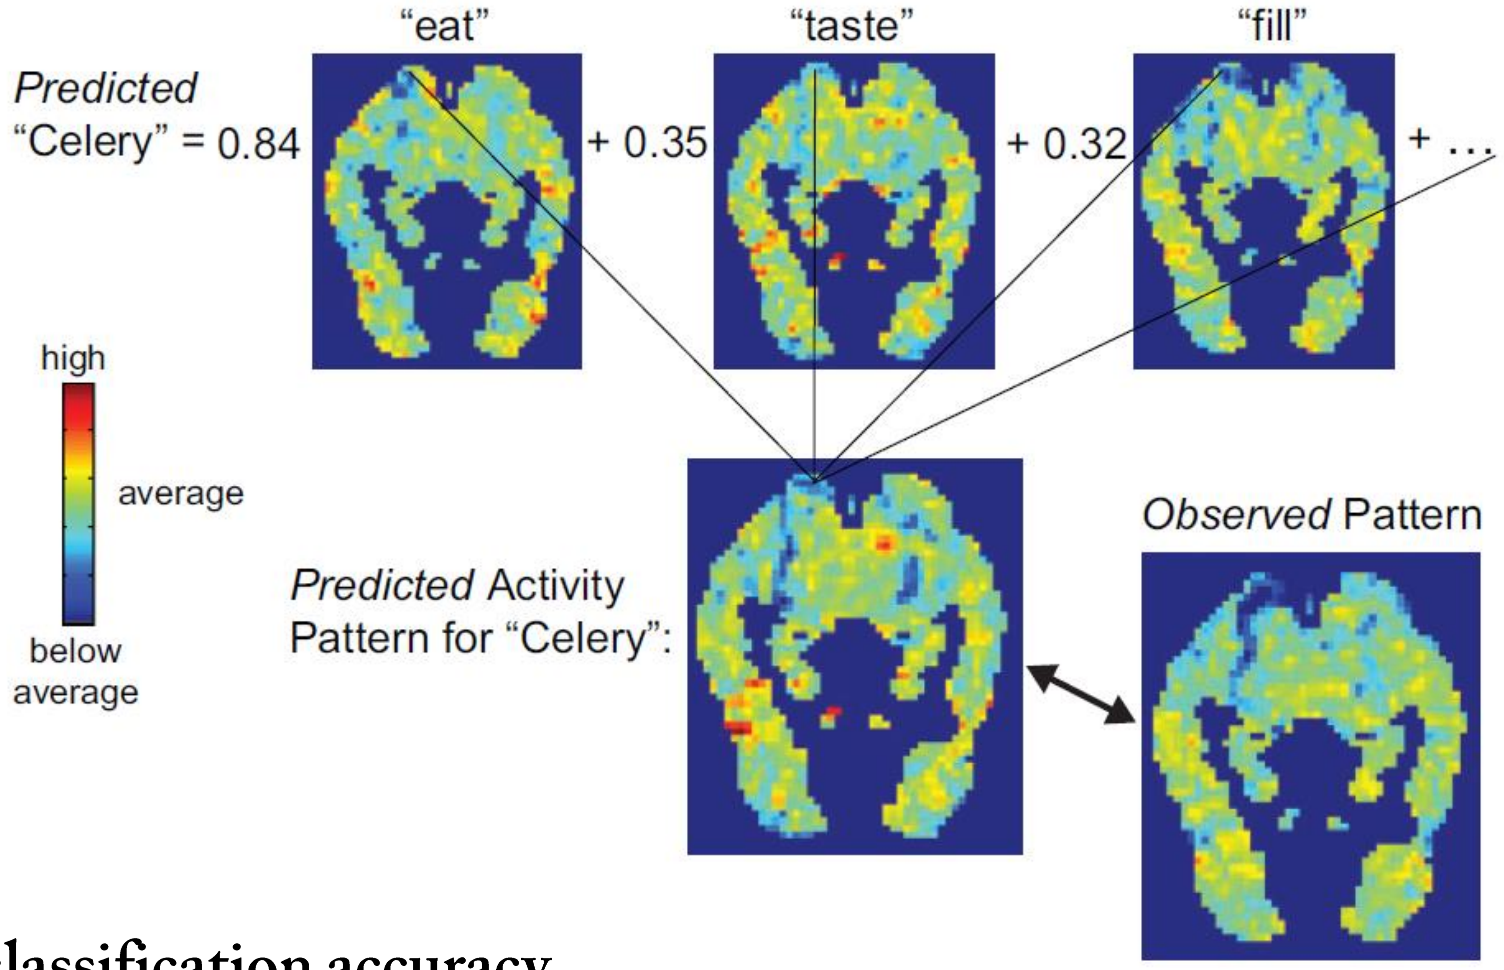

77% classification accuracy
(chance 50%)

Mitchell et al (2008)

Summary

- MVPA can be used to:
  - Extract detailed visual information
  - Decode:
    - Attended objects
    - Freely Recalled memory
    - Novel nouns
- Reliable method for gaining insight into a range of neuroscience questions

Outline

- BCI to study learning in the motor cortex
- MVPA to study vision in human fMRI data
- Population decoding to study high level learning and vision in macaque monkeys

Decoding visual information from IT

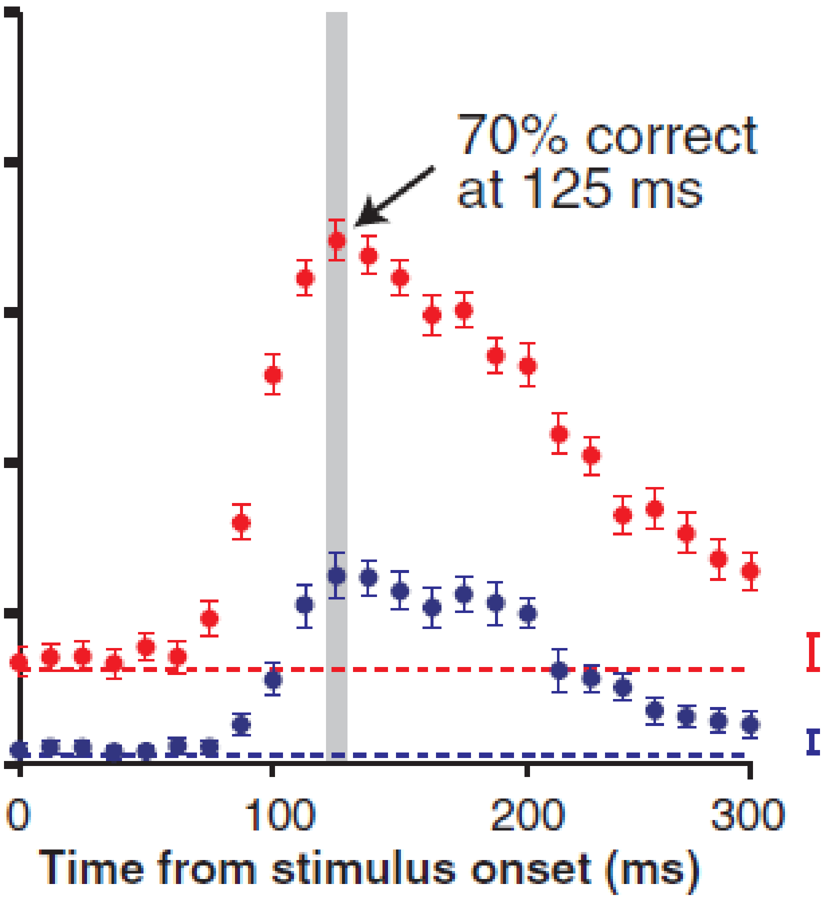

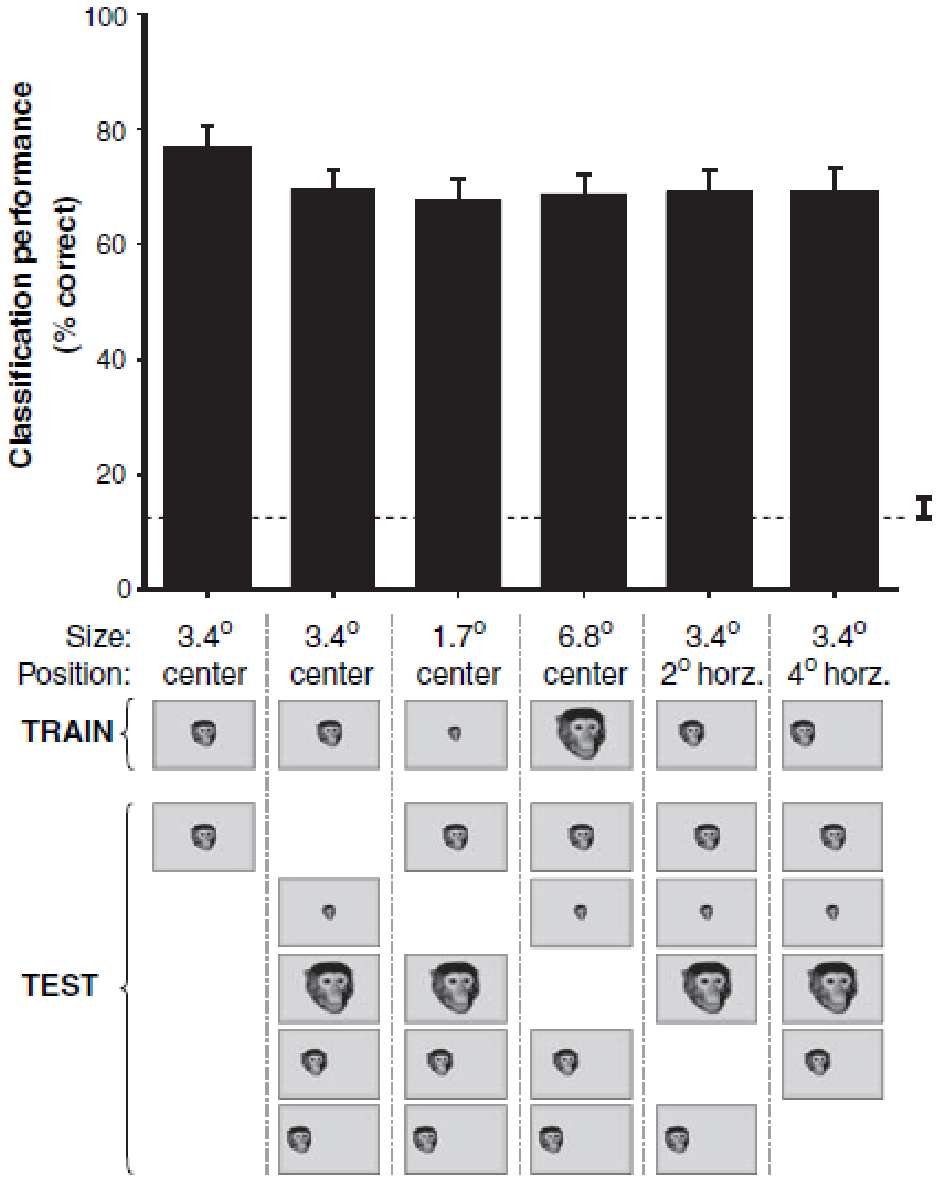

Hung et al (2005)

Work with Freiwald & Tsao

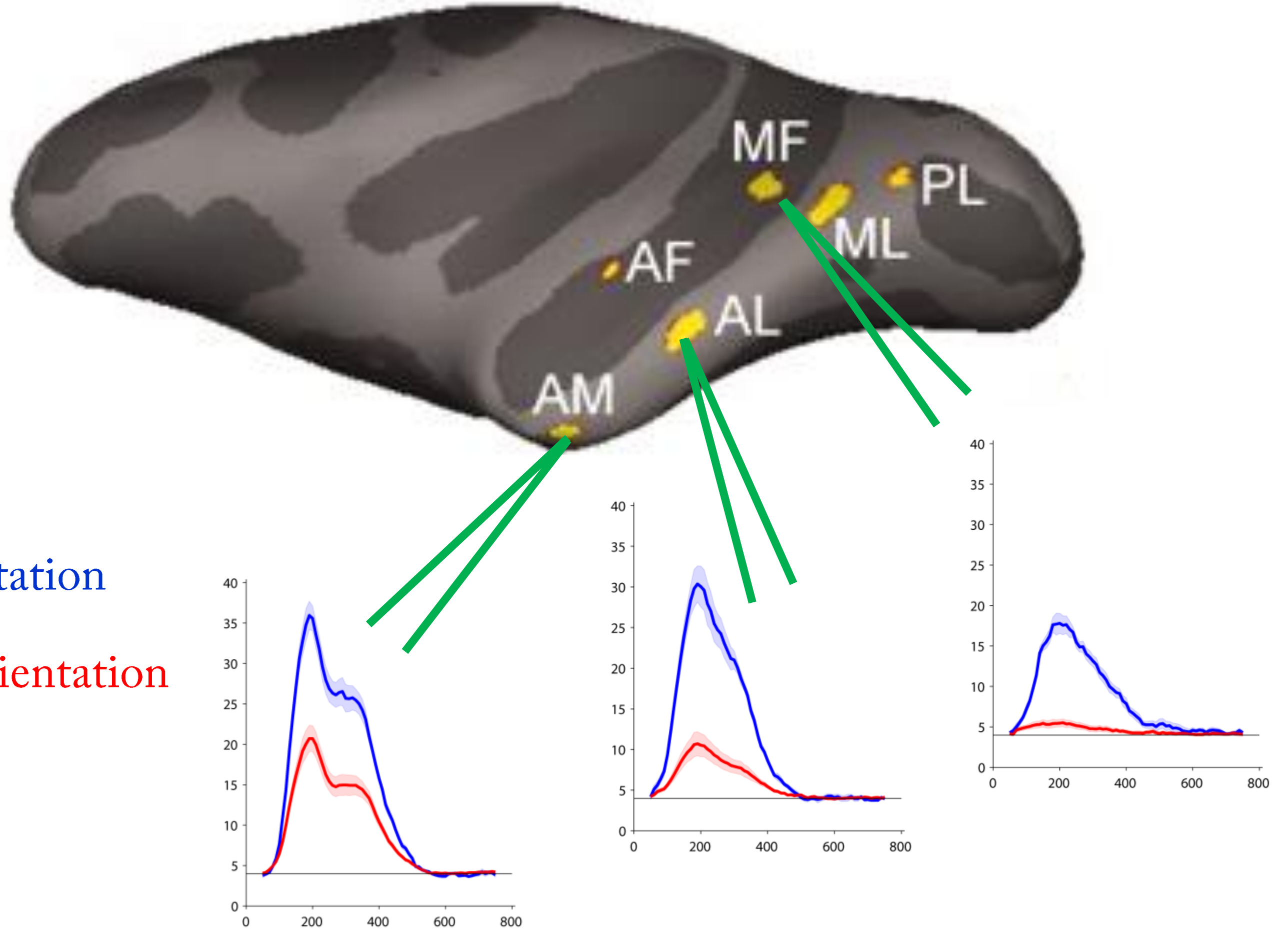

Same orientation

Different orientation

Incorporation of new information into prefrontal cortical activity after learning working memory tasks

Ethan M. Meyers[a,1], Xue-Lian Qi[b], and Christos Constantinidis[b]

The prefrontal cortex (PFC) is involved in working memory, task learning, and executive function

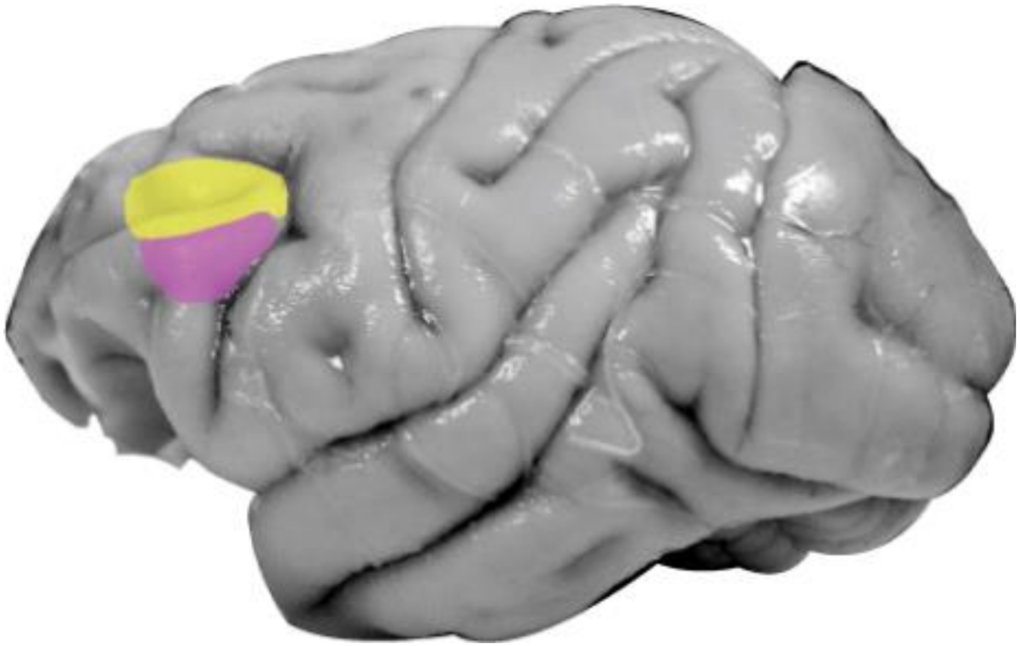

Miller 2000; Miller and Cohen 2001; Duncan 2001

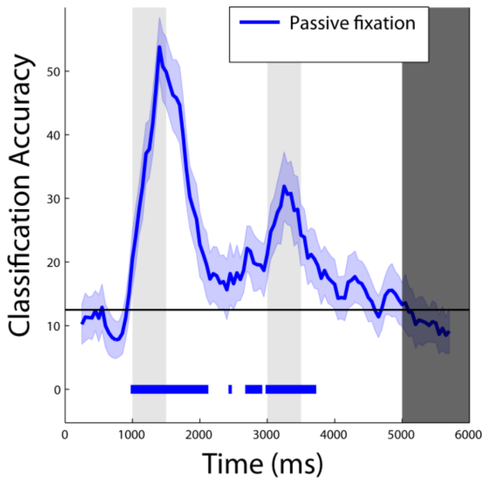

1. How does the information content change after learning a new task?

2. How is information coded in neural activity?

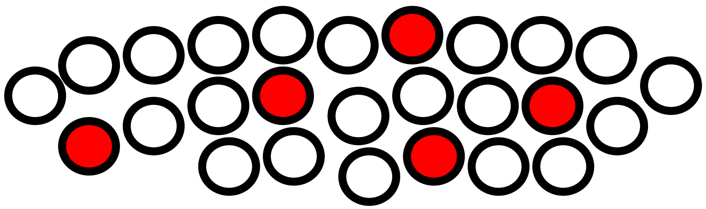

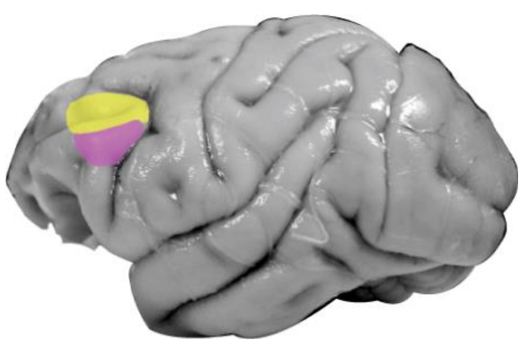

3. Are there regions differences within dorsolateral PFC?

Monkeys were first trained to passively fixate

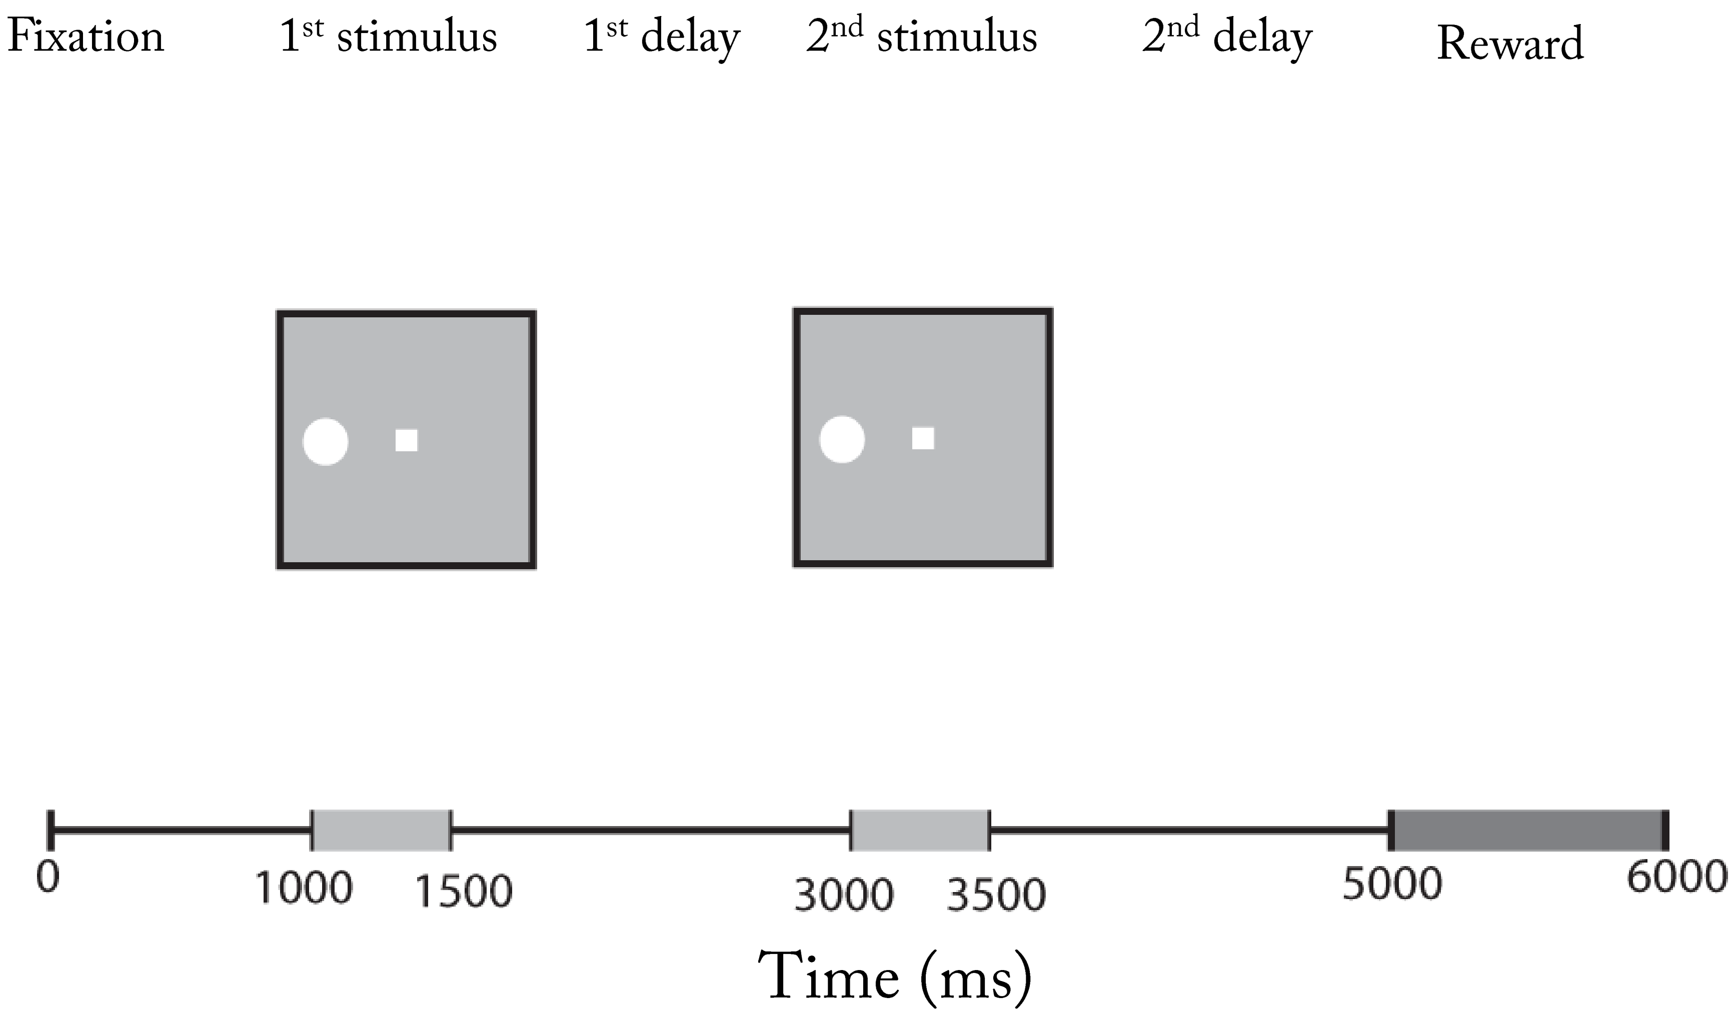

Monkeys were first trained to passively fixate

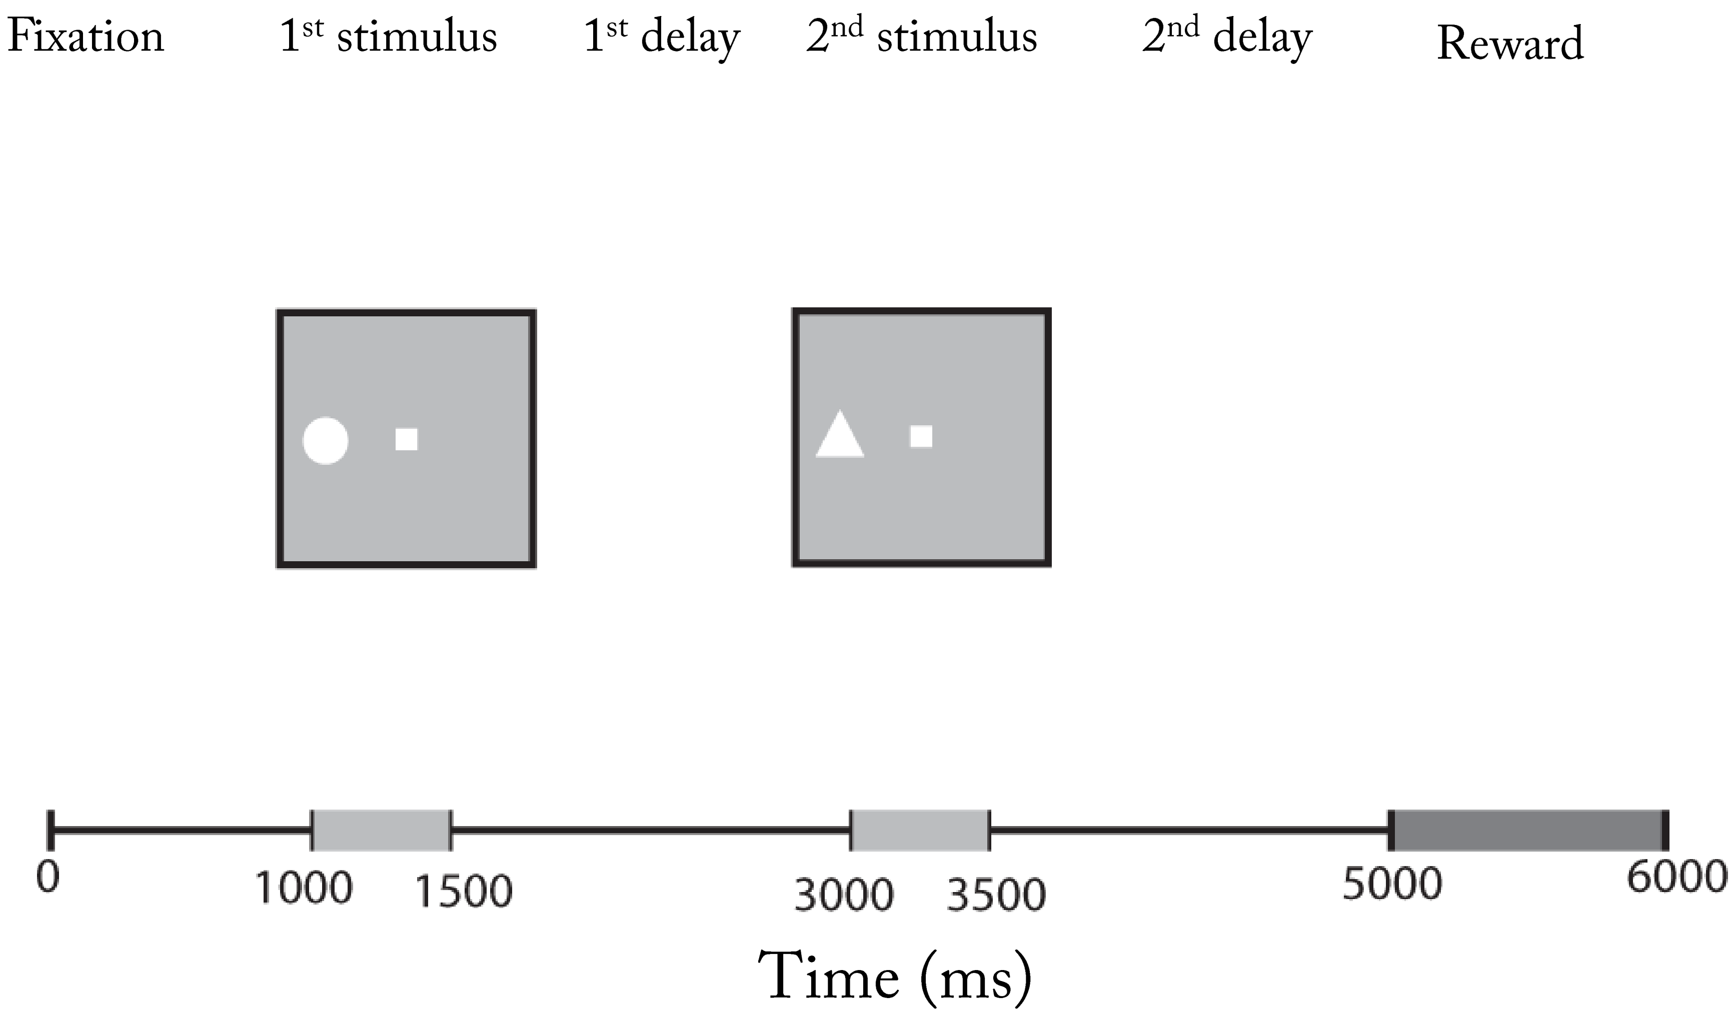

Monkeys then engaged in a delayed-match-to-sample task (DMS task)

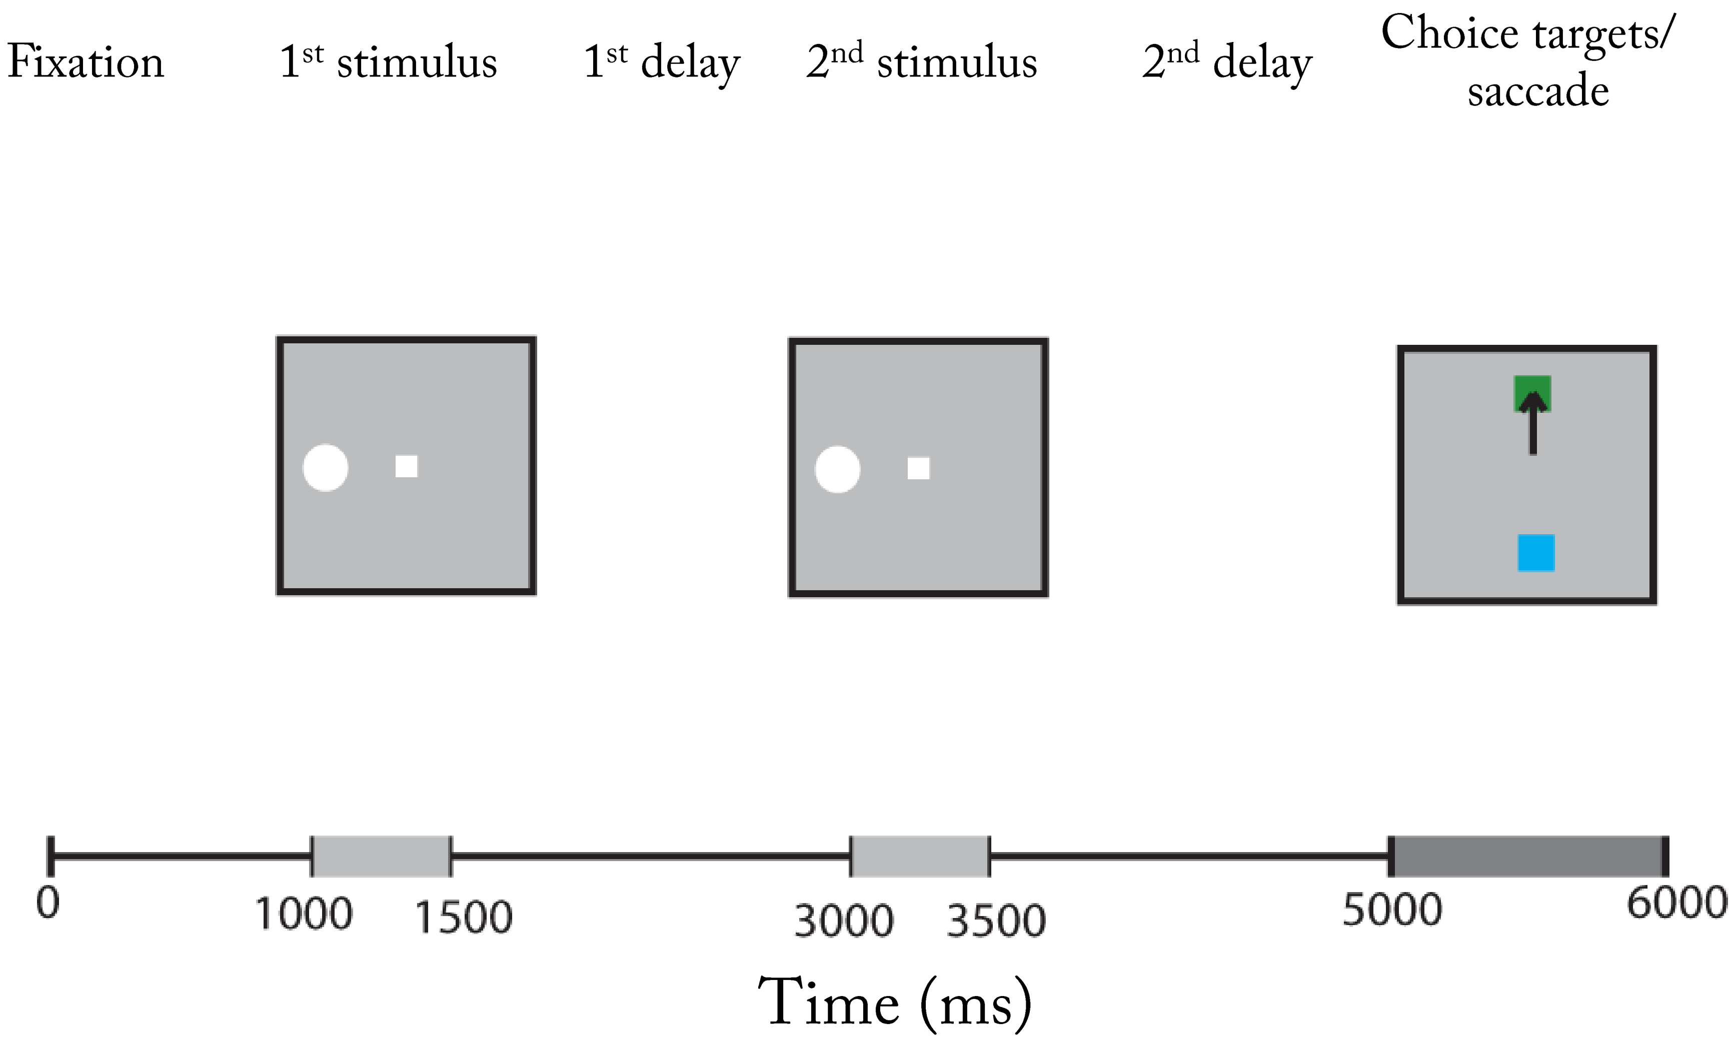

Monkeys then engaged in a delayed-match-to-sample task (DMS task)

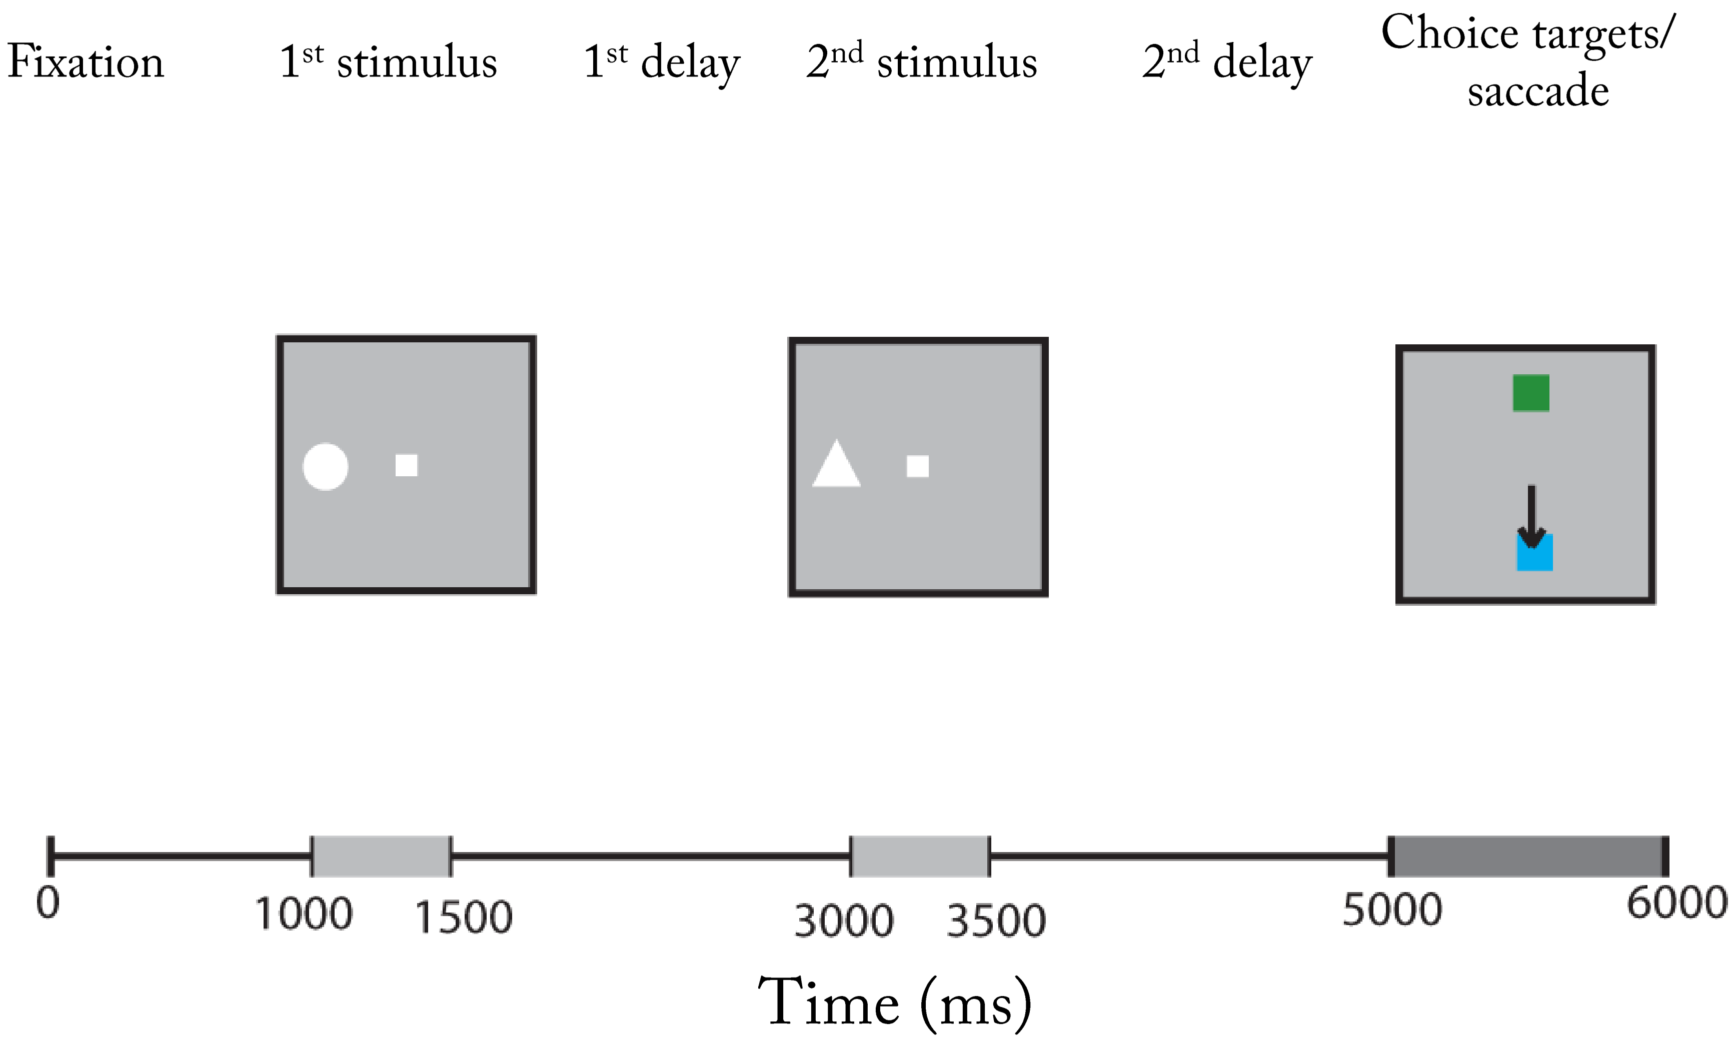

Spatial task: DMS task
Spatial task: Passive fixation

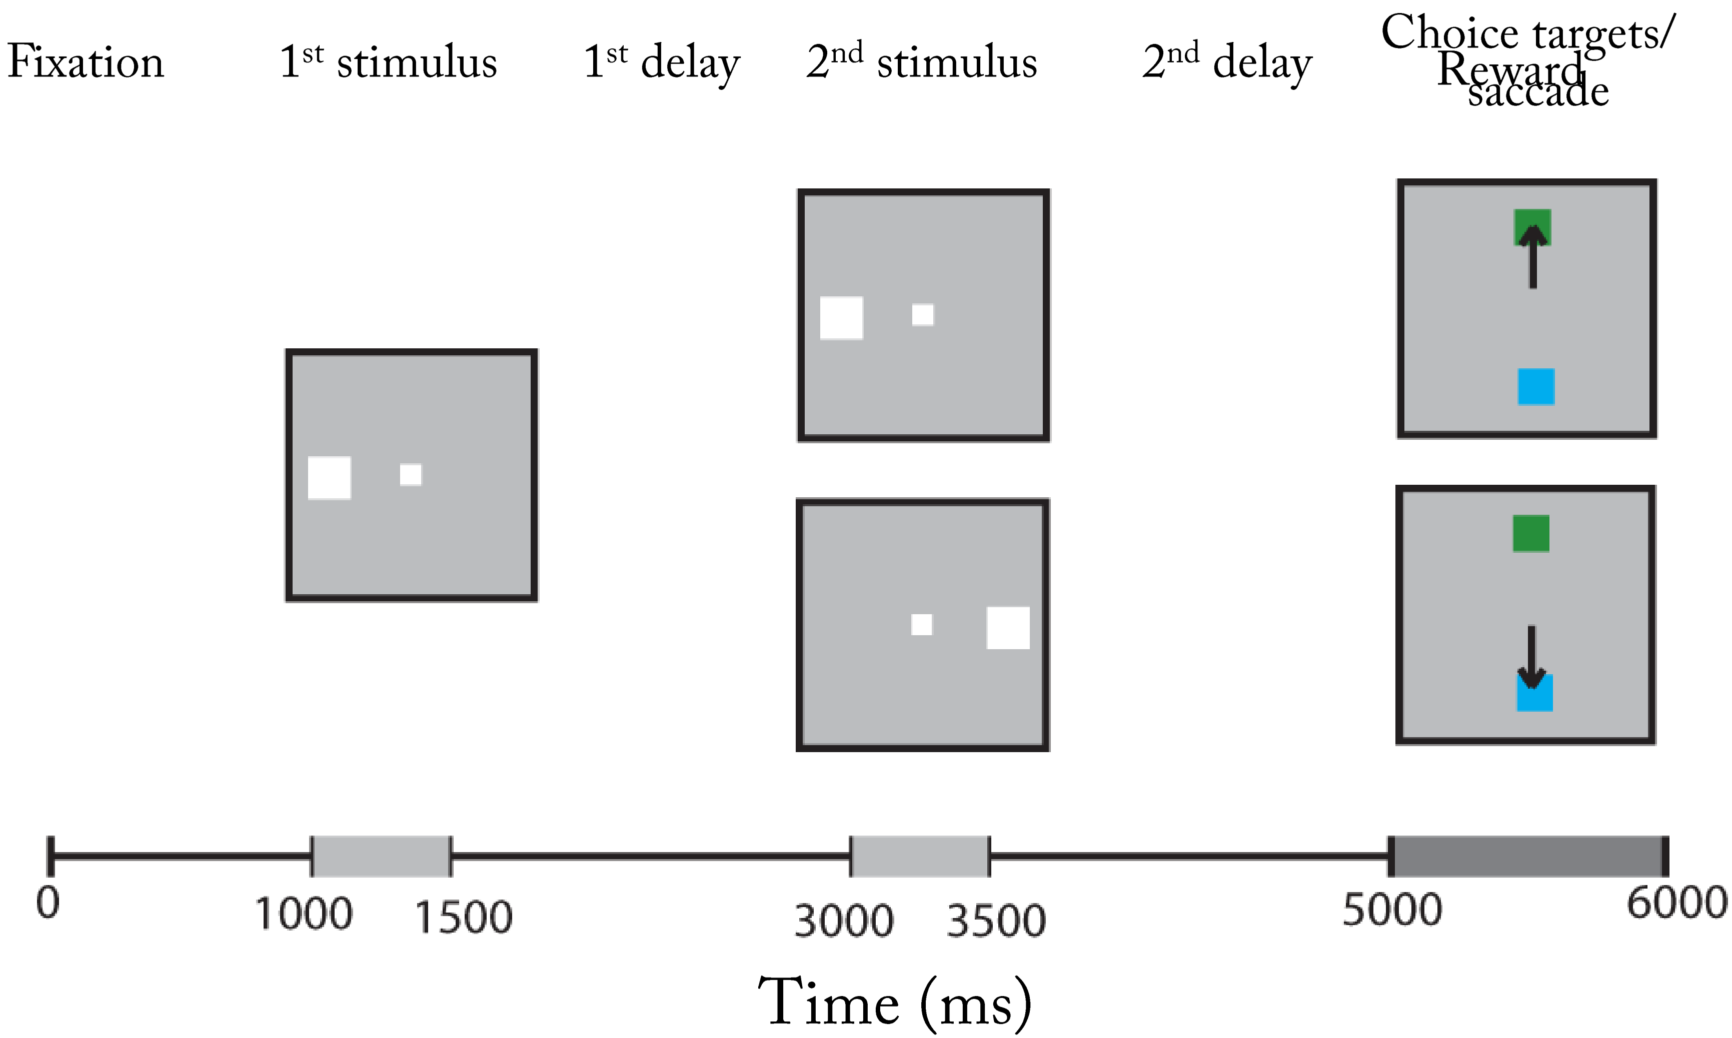

Population Decoding

Decoding applied

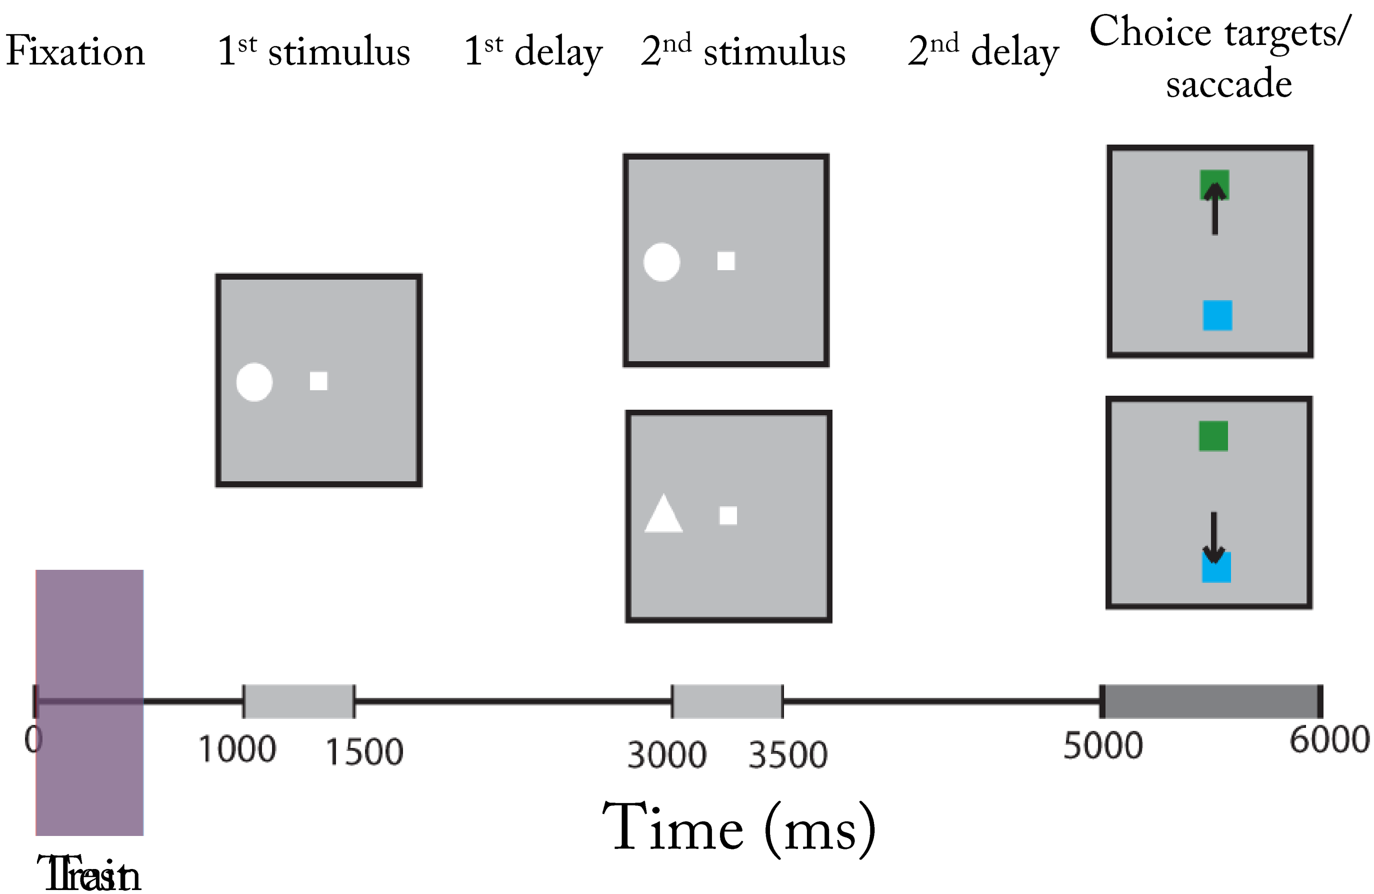

Decoding applied

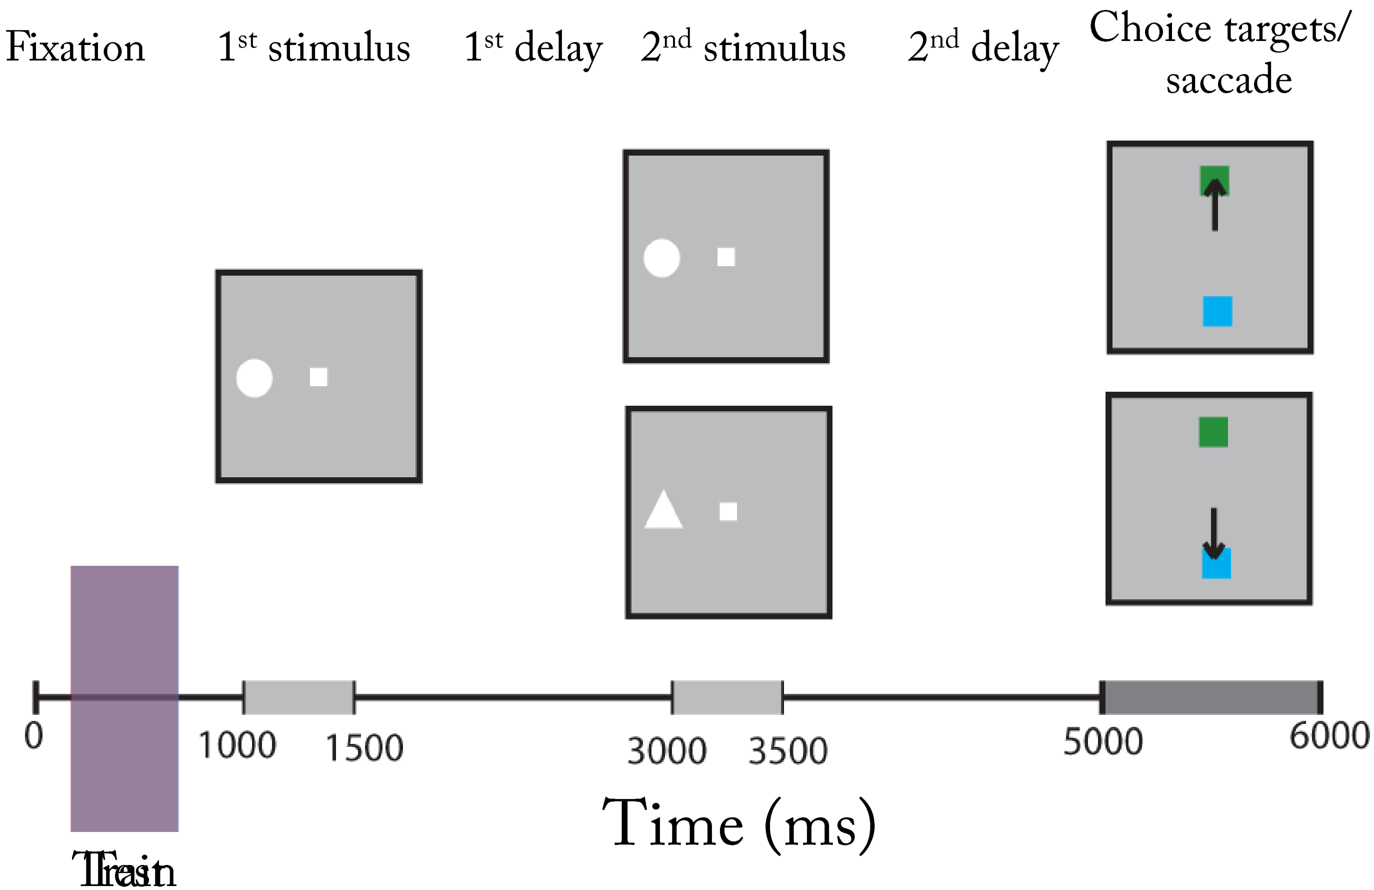

500 ms bins, sample every 50 ms

What types of information are in PFC?

Visual information

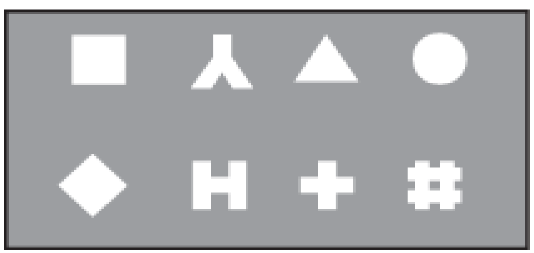

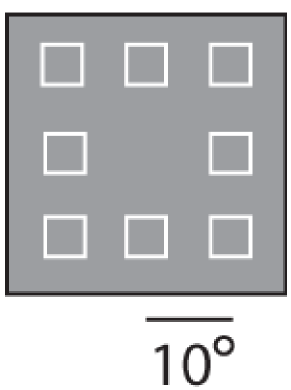

Match/nonmatch information

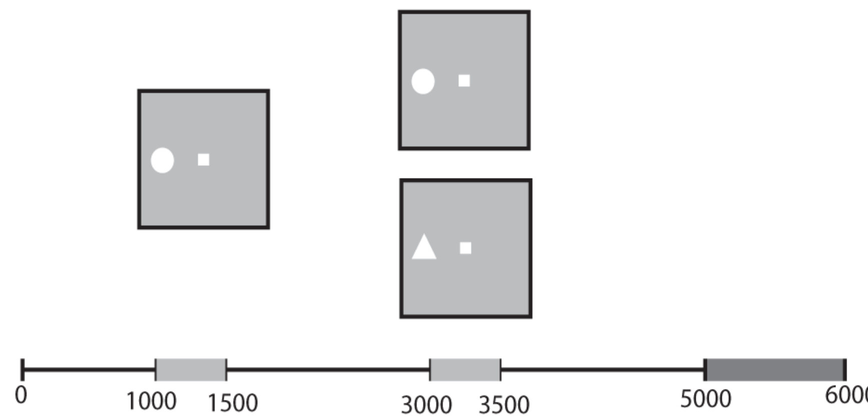

Analyzed pseudo-population of 750/600 neurons

Decoding visual information (feature task)

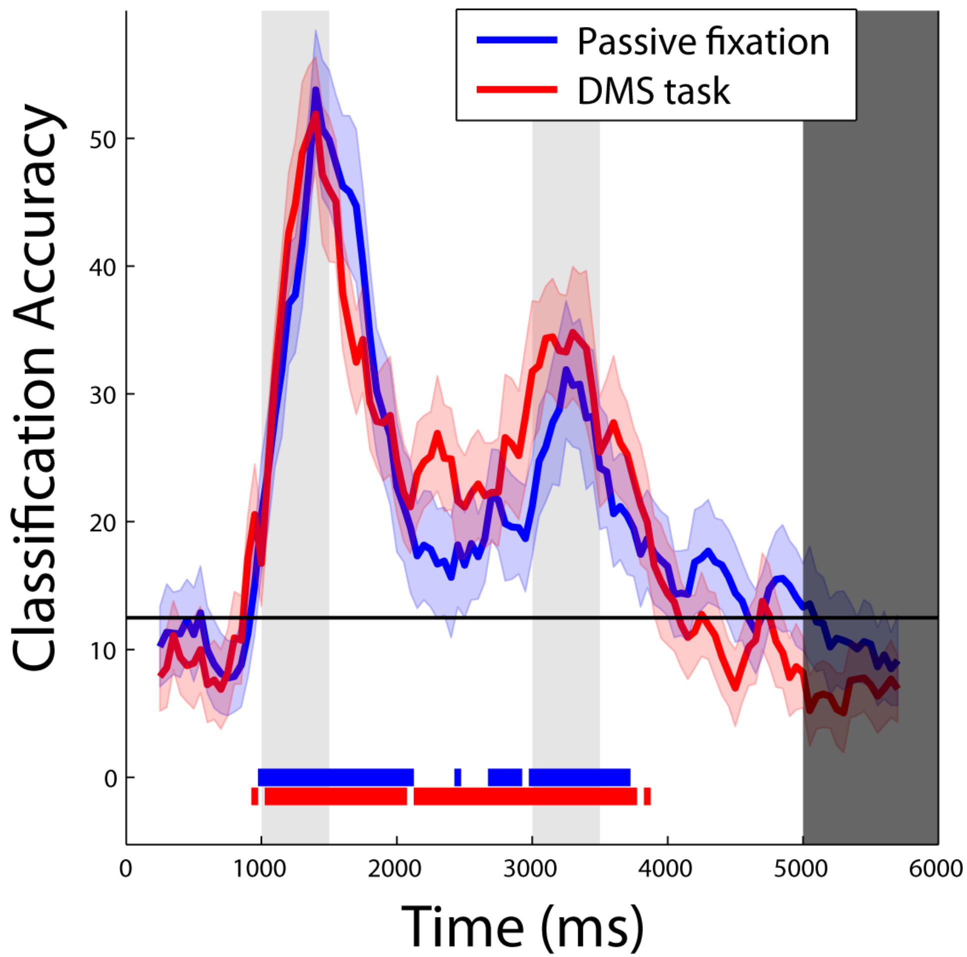

Decoding visual information (spatial task)

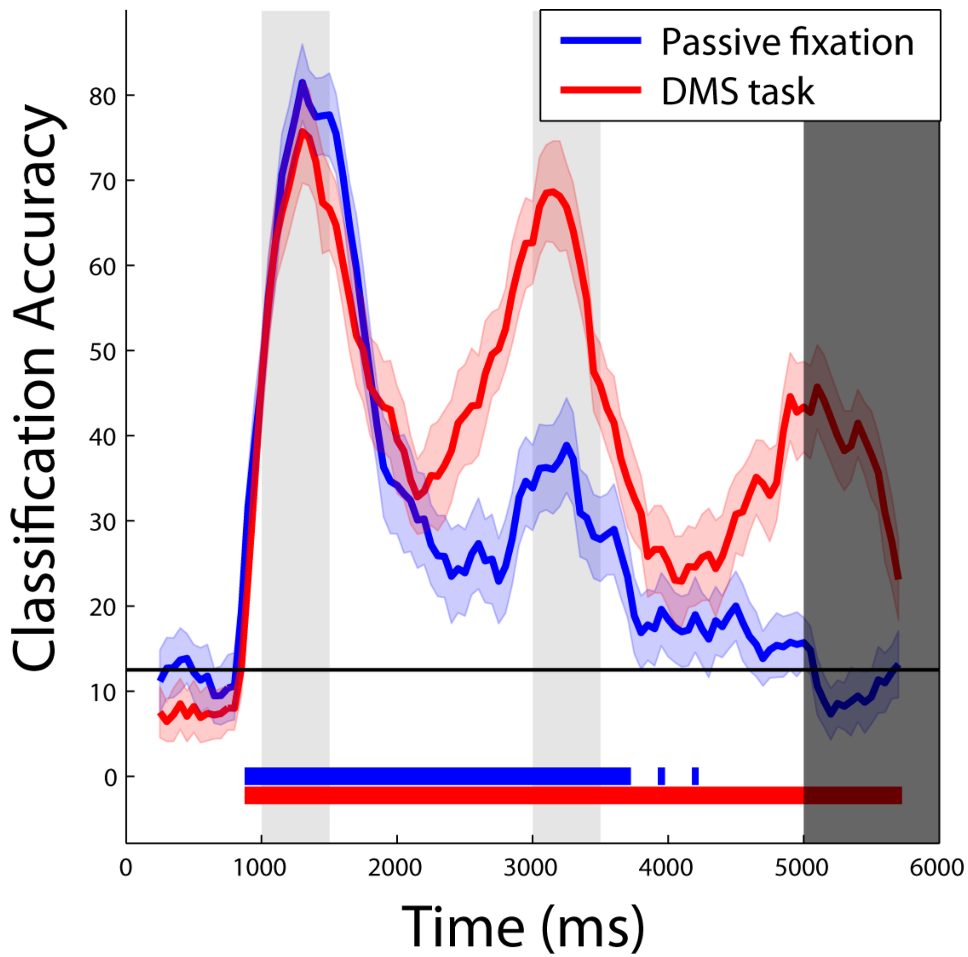

Decoding match/nonmatch information

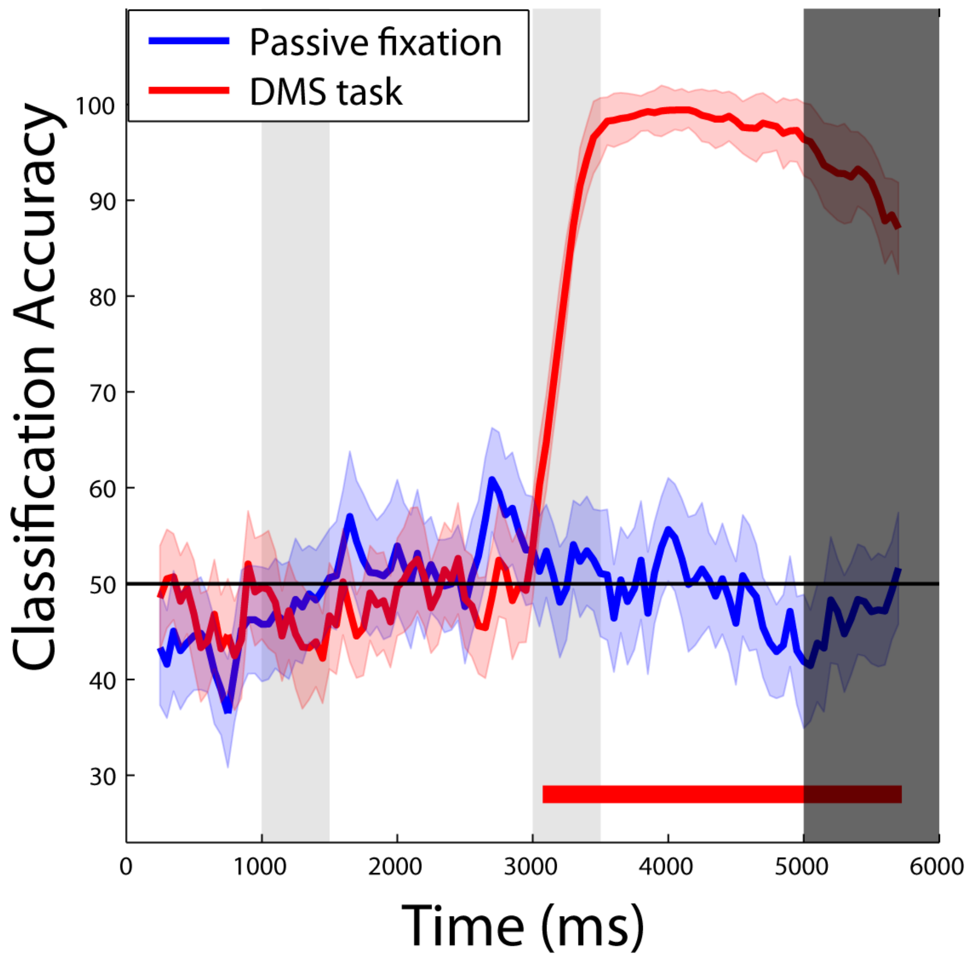

Decoding match/nonmatch information

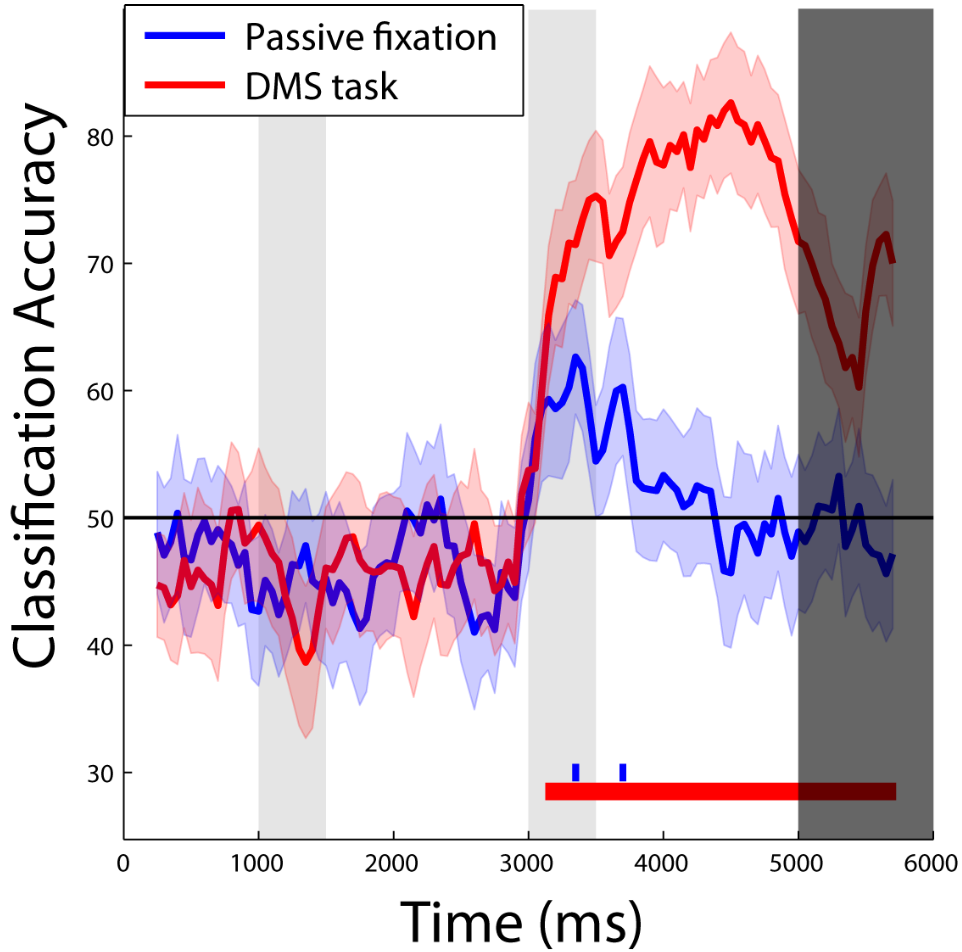

1. How does the information content change after learning a new task?

<u>Answer:</u>

Visual information remains largely unchanged

There is a large increase in task-relevant information

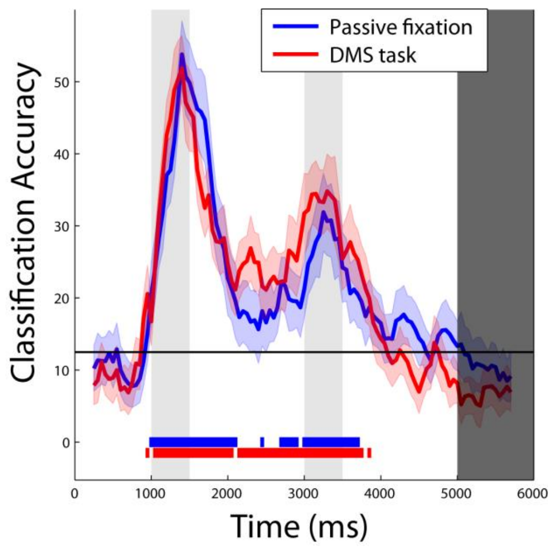

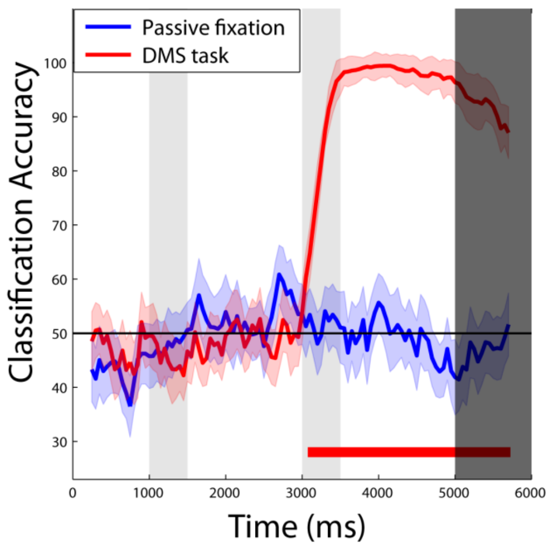

How is information coded in neural activity?

a) Sparse vs. distributed information

b) Dynamic vs. static population coding

c) Dedicated neurons vs. multiplexing

Is the new information widely distributed?

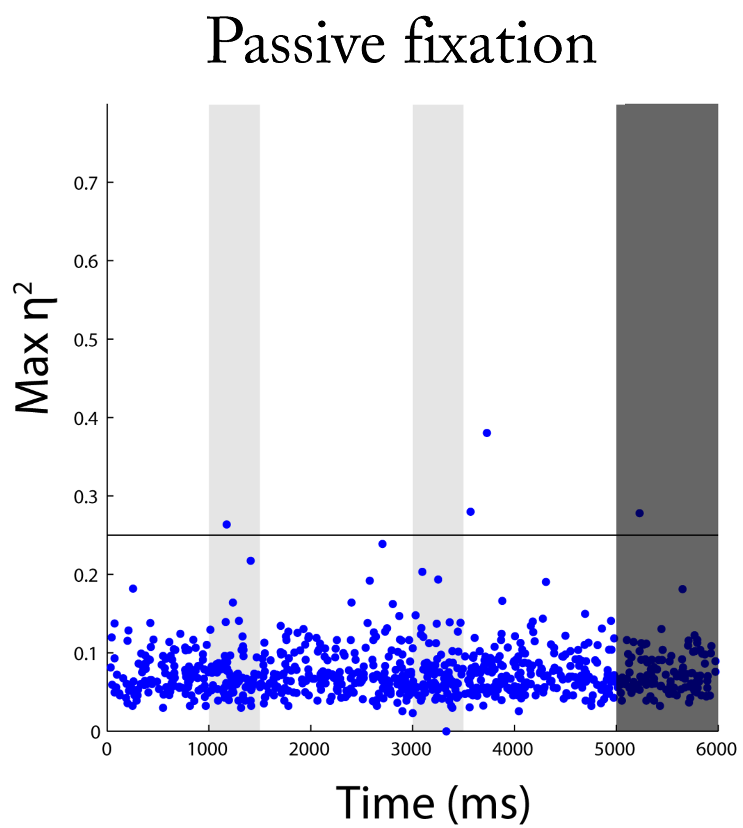

Feature task

Is the new information widely distributed?

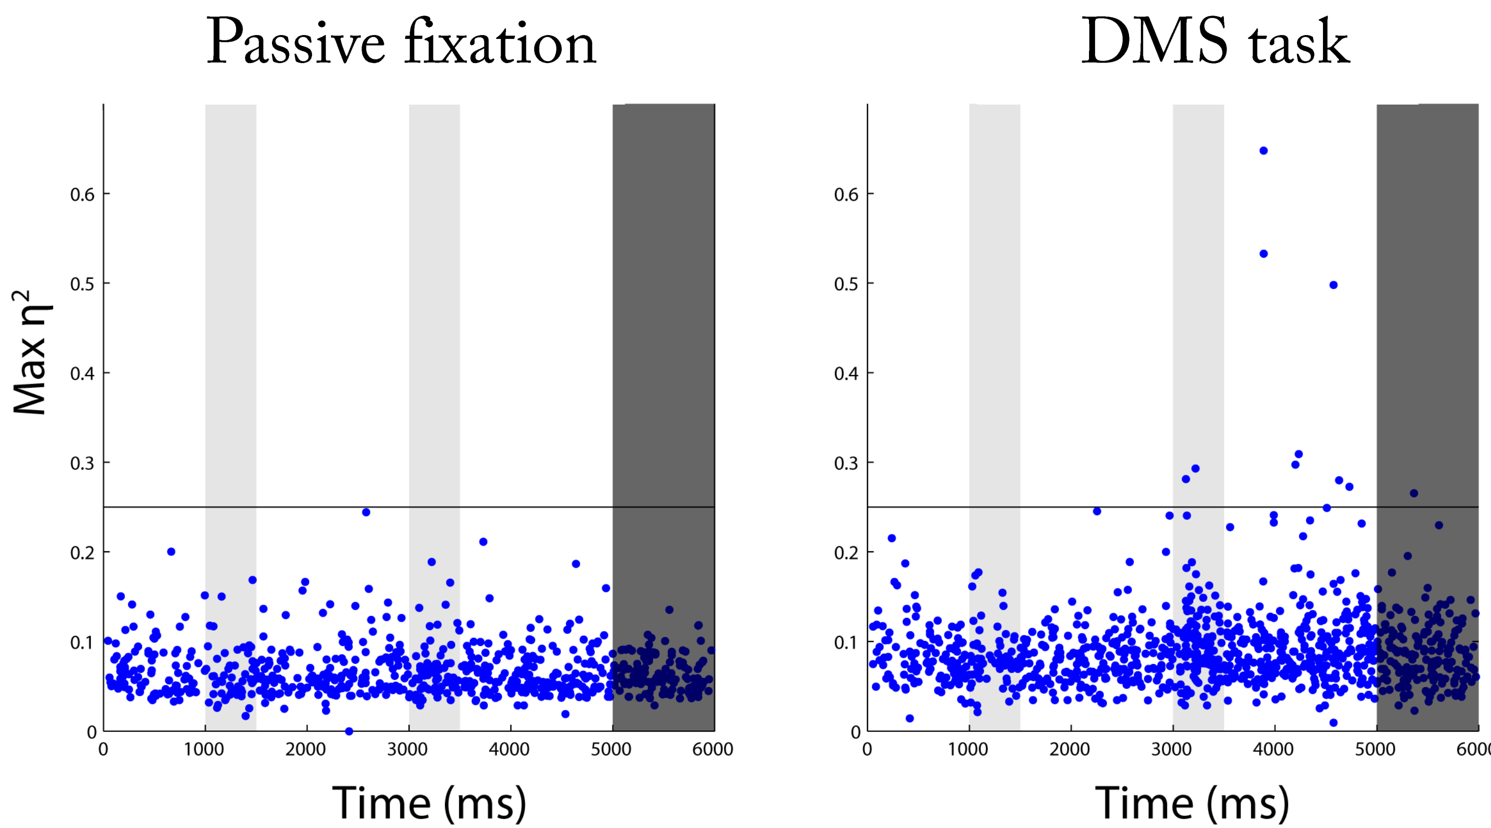

Spatial task

Is the new information widely distributed?

Using only the 8 most selective neurons

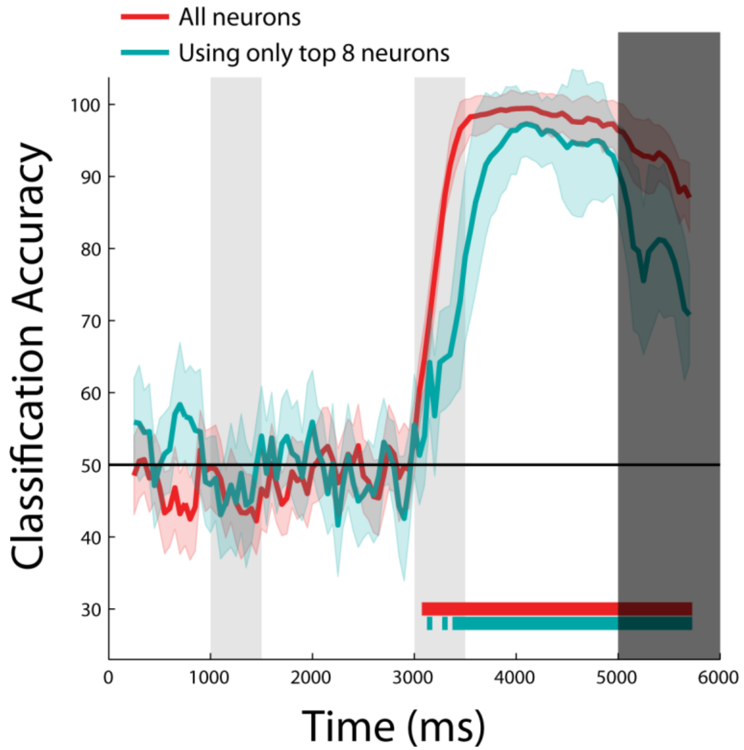

Feature task

Is the new information widely distributed?

Using only the 8 most selective neurons | Excluding the 64 most selective neurons

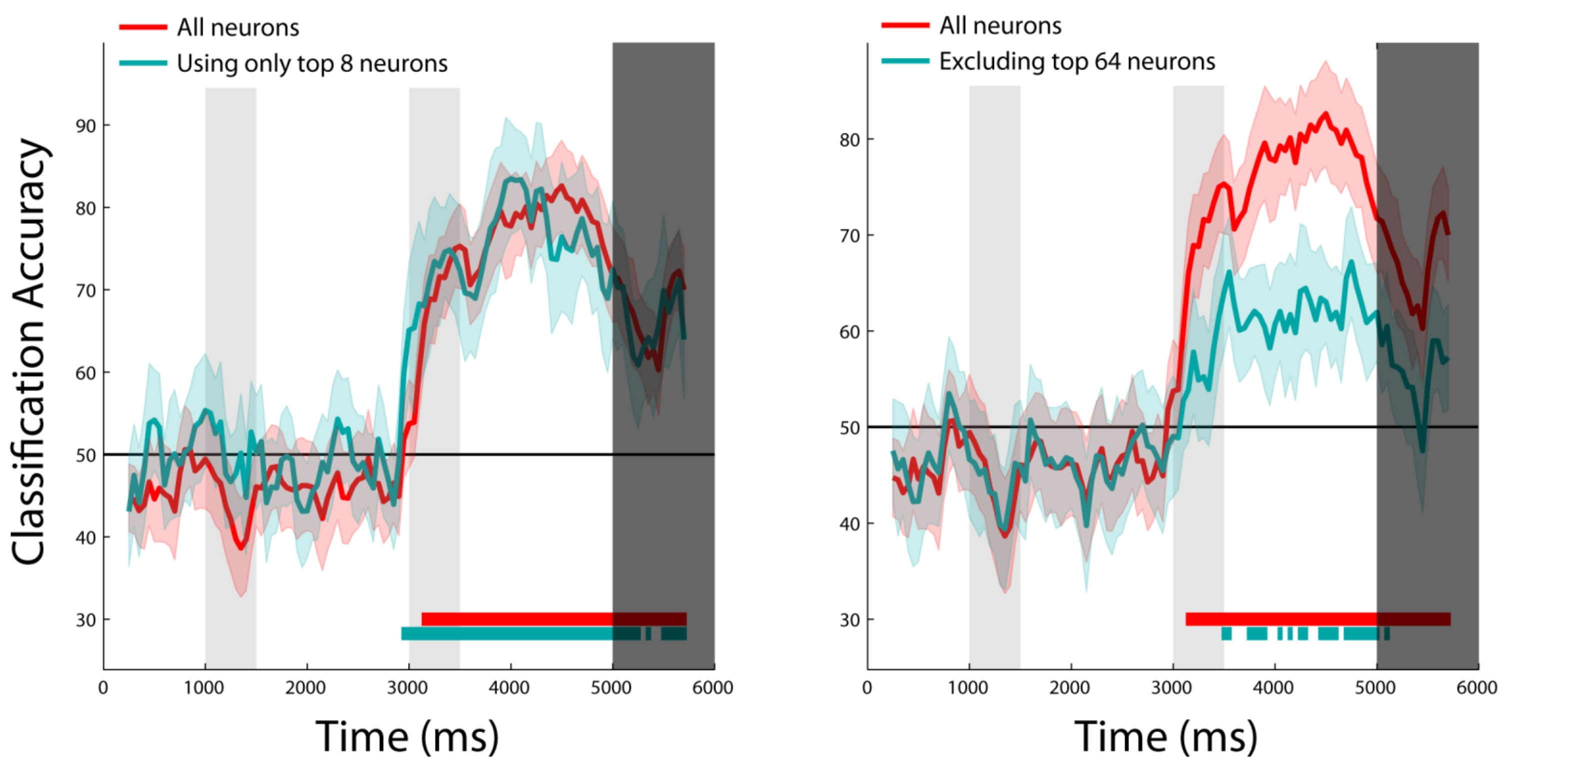

Spatial task

How is information coded in neural activity? (dynamics)

DMS Task

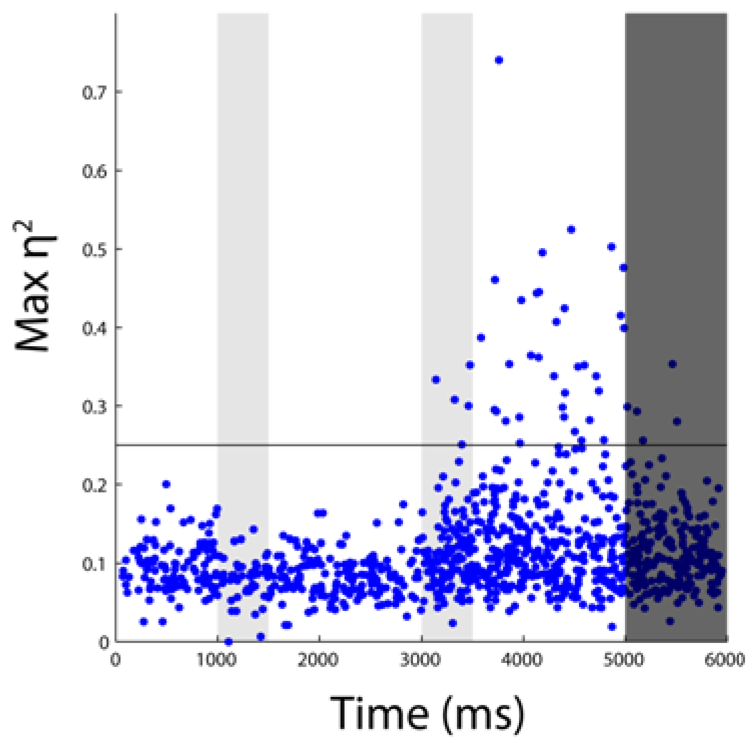

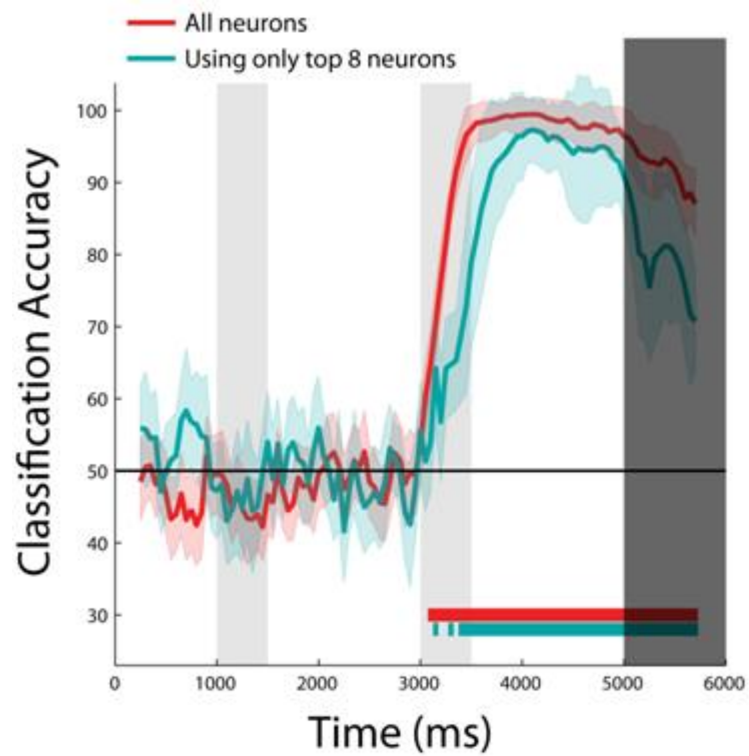

Meyers, et al., 2008 Add more refs

Is information contained in a dynamic population code?

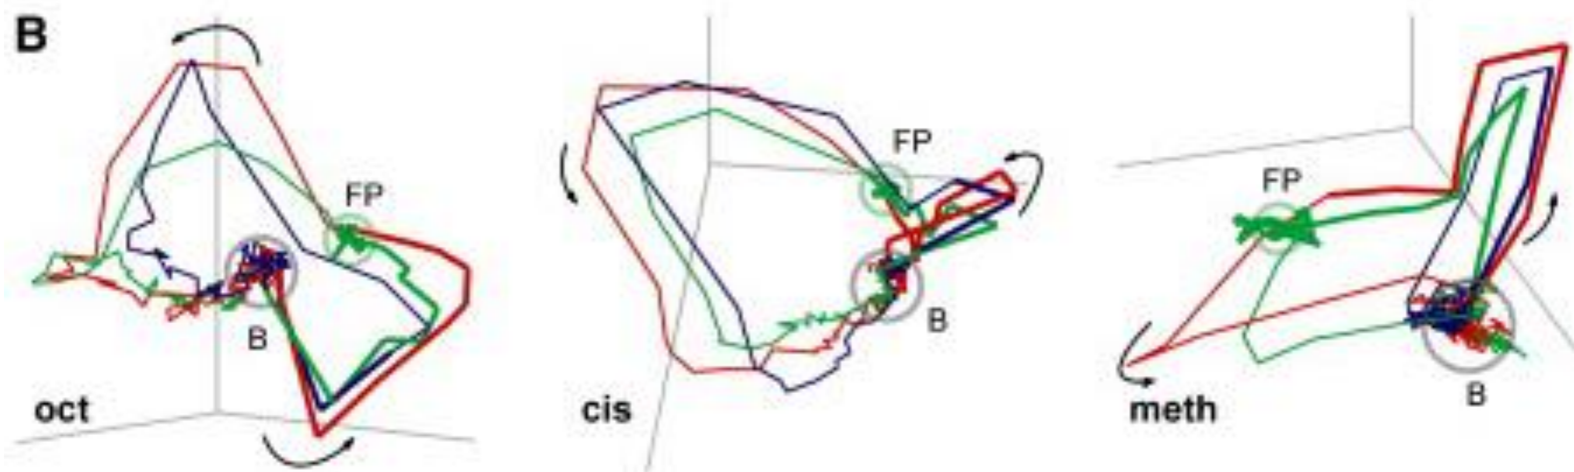

Mazor and Laurent 2005, Meyers et al 2008, Harvey et al 2012

Decoding applied

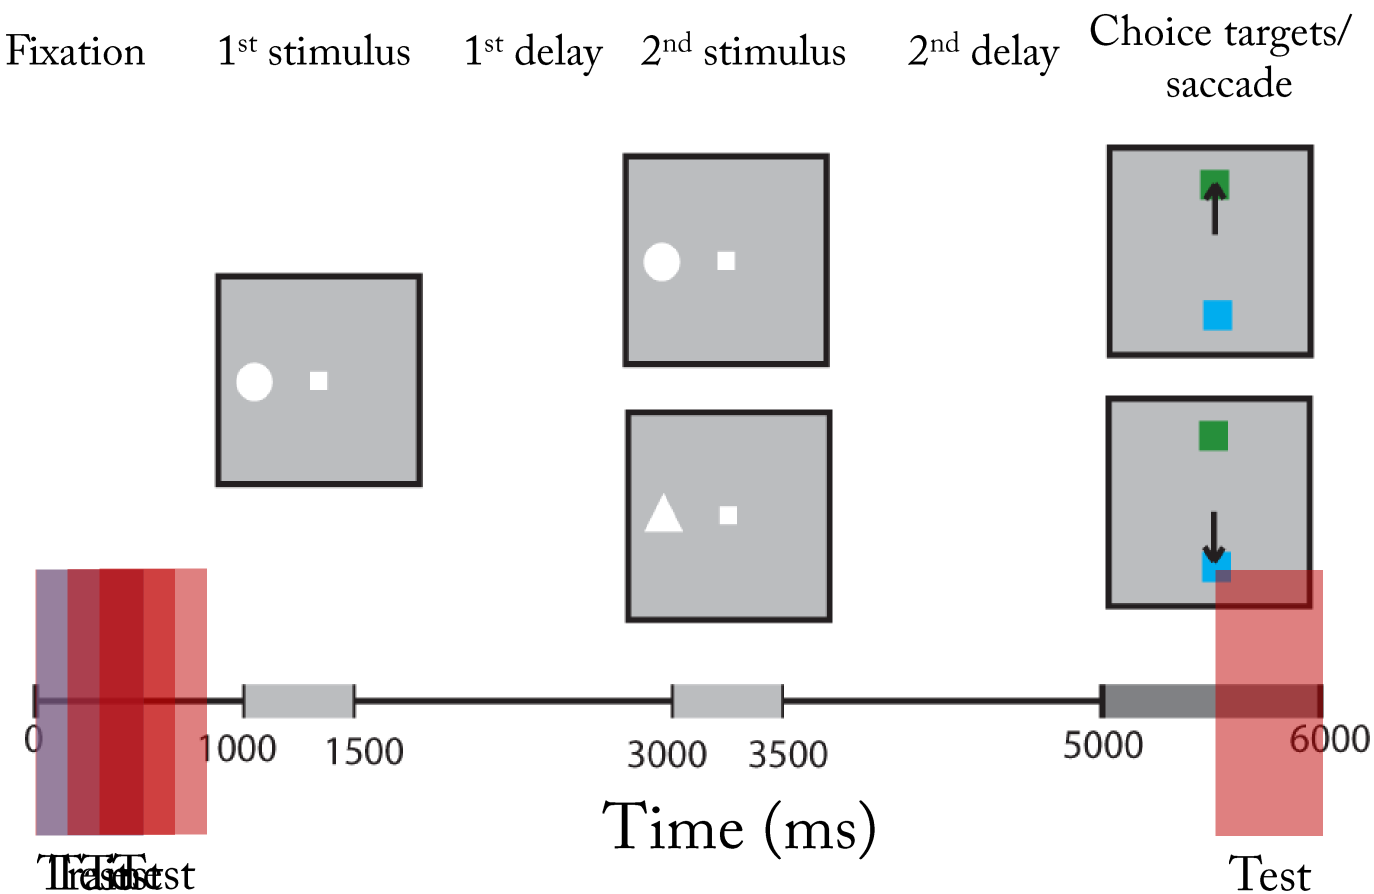

Decoding applied

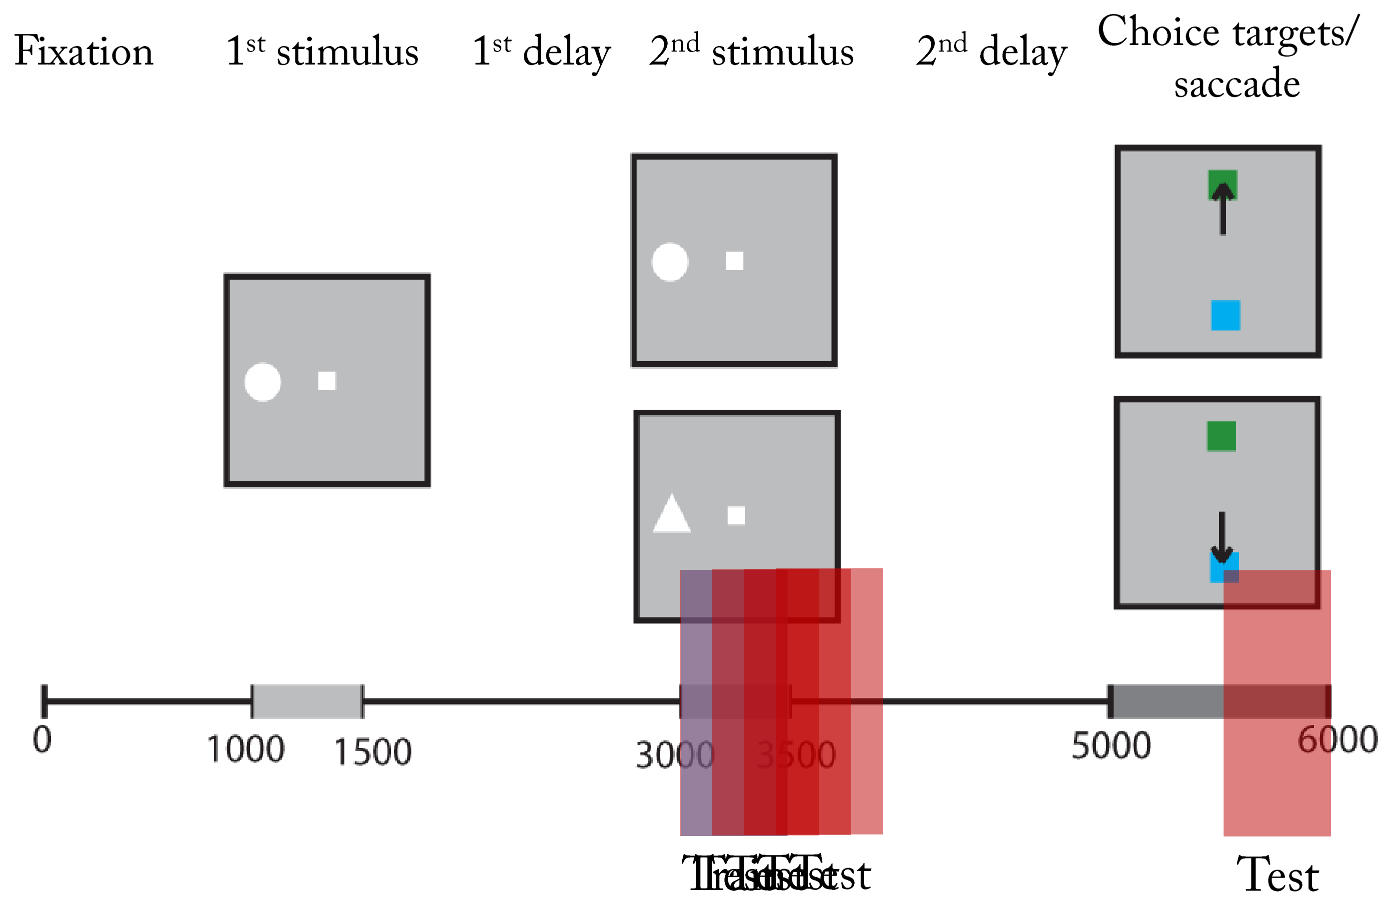

Dynamic population coding

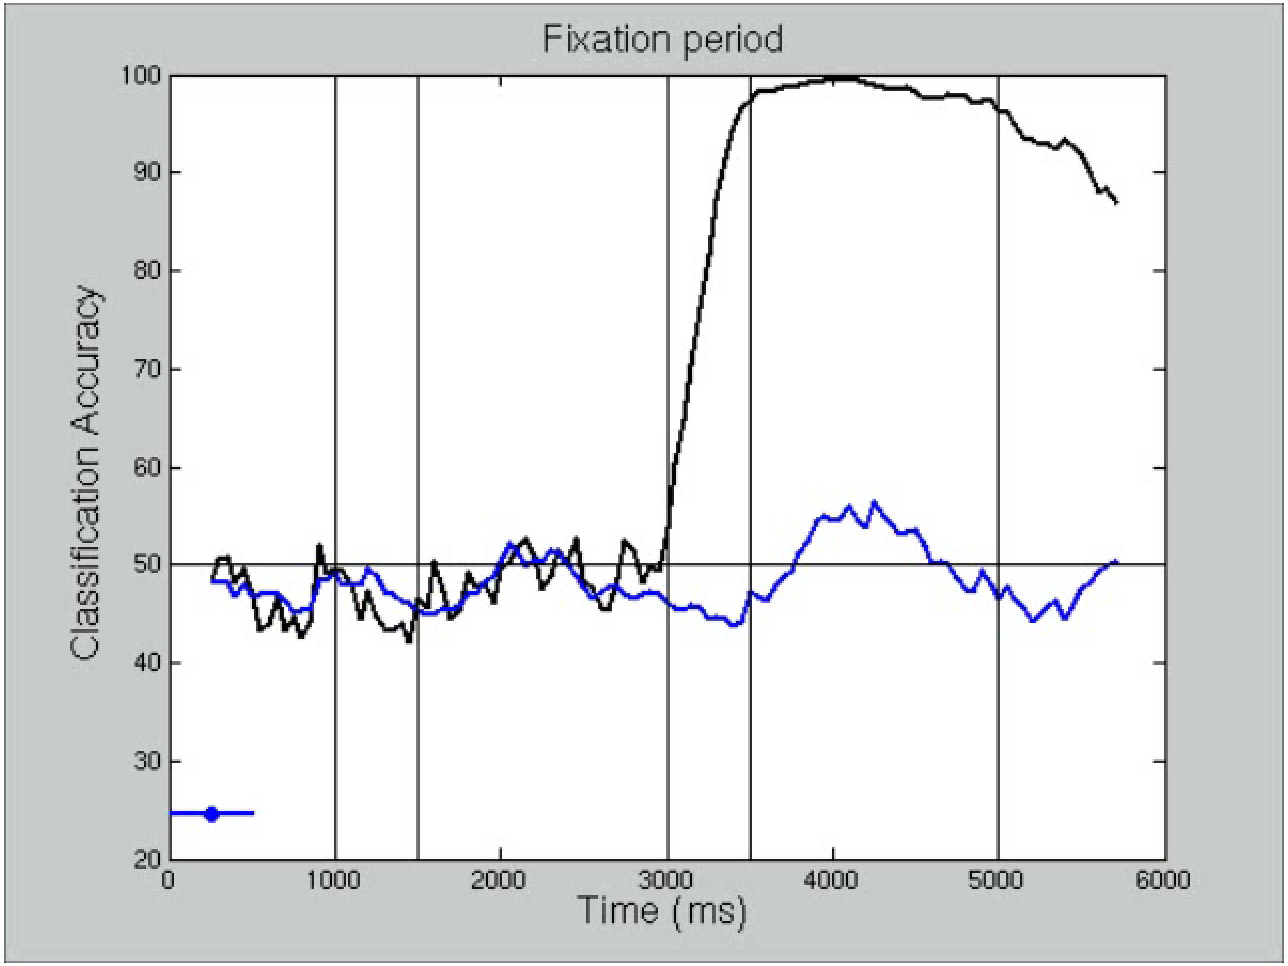

Feature task

Dynamic population coding

Passive fixation

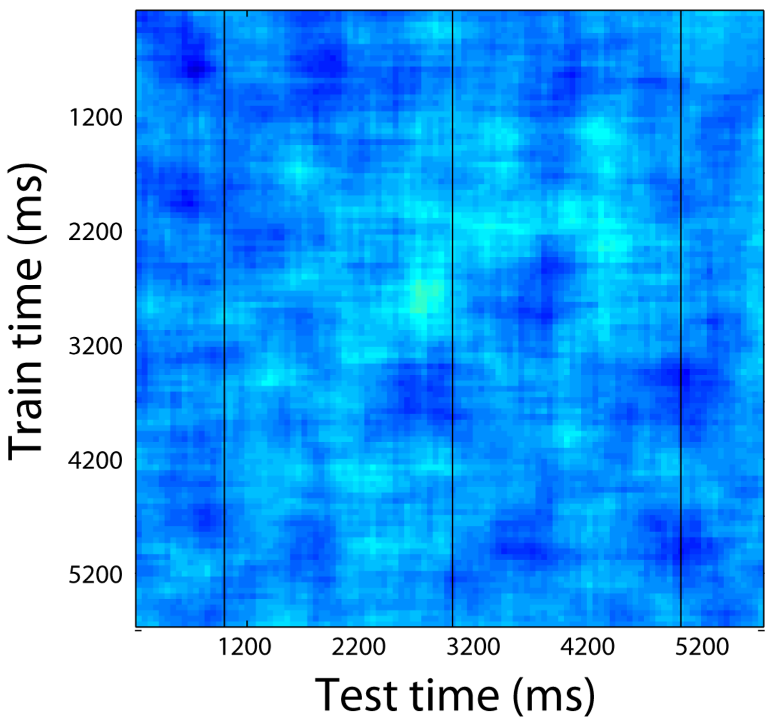

Feature task

Dynamic population coding

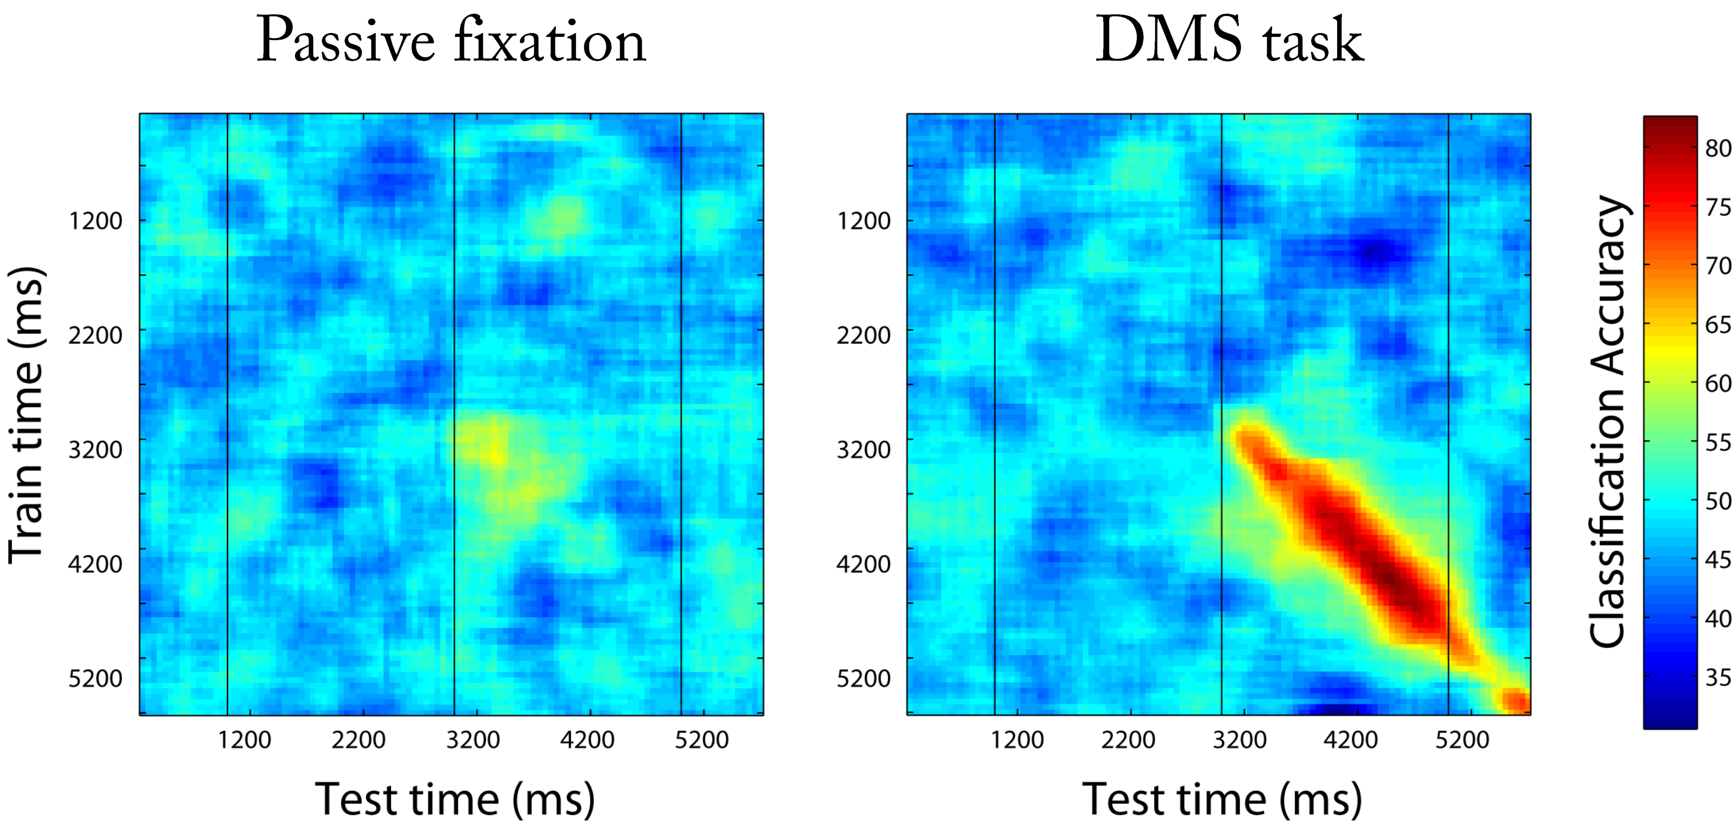

Spatial task

The dynamics can be seen in individual neurons

Feature task

Spatial task

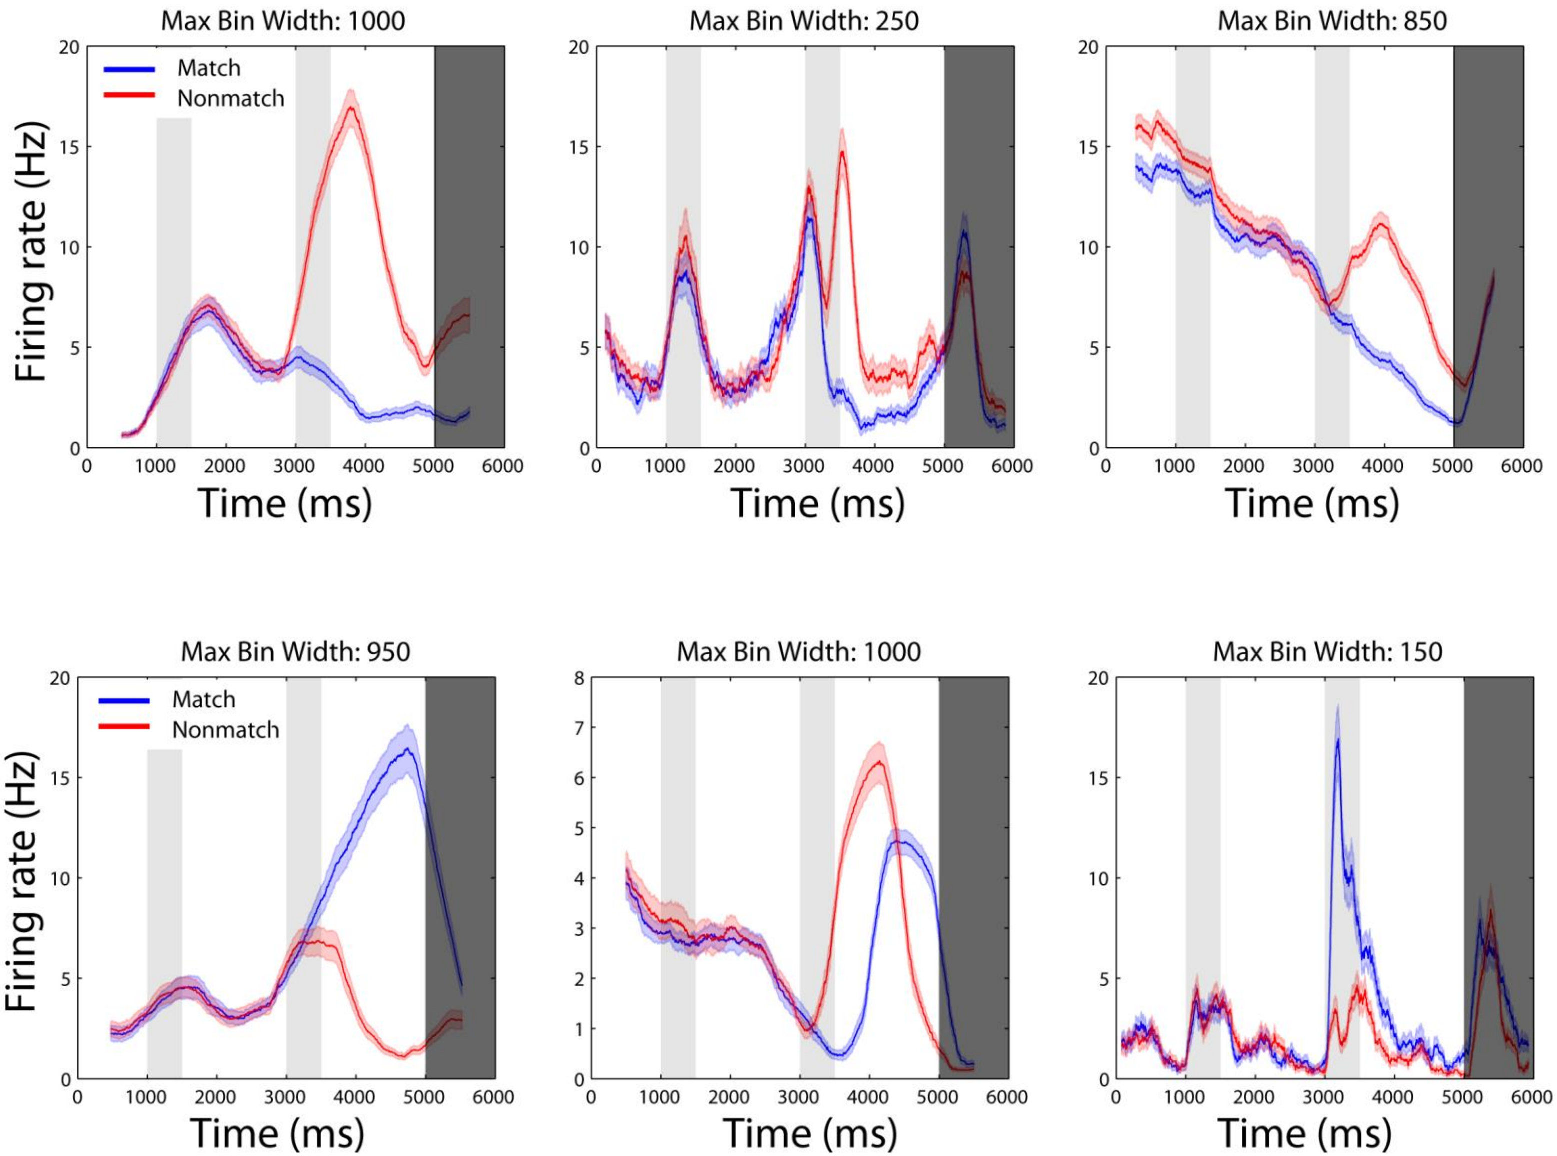

Do neurons become specialized to new information or do they multiplex information?

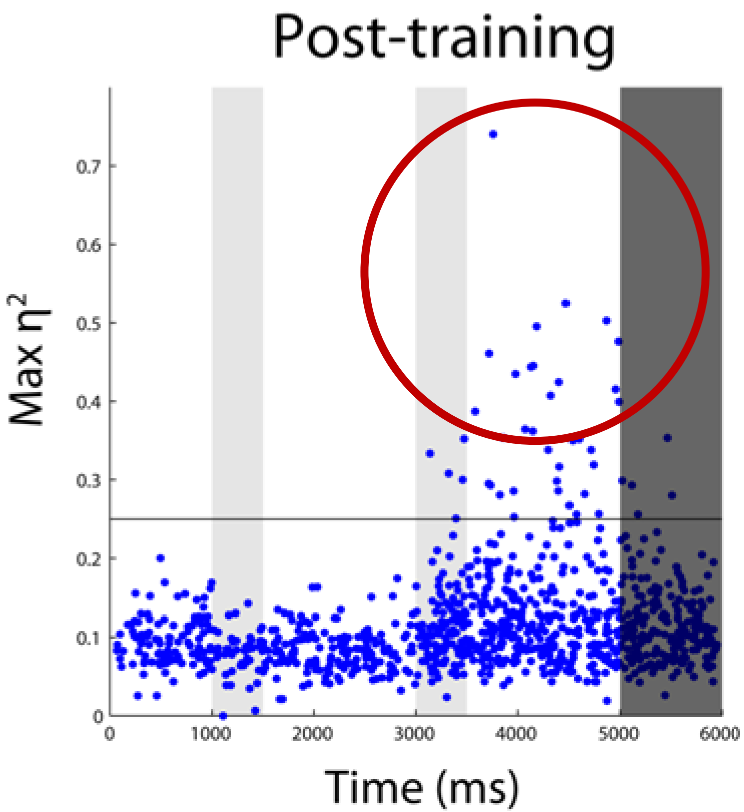

- Find highly selective match/nonmatch neurons (ANOVA, $p < 10^{-6.5}$)
- 50 neurons feature task
- 18 neurons spatial task
- Use highly selective match/nonmatch neurons to decode *visual information*

Highly selective match/nonmatch neurons also contain stimulus information

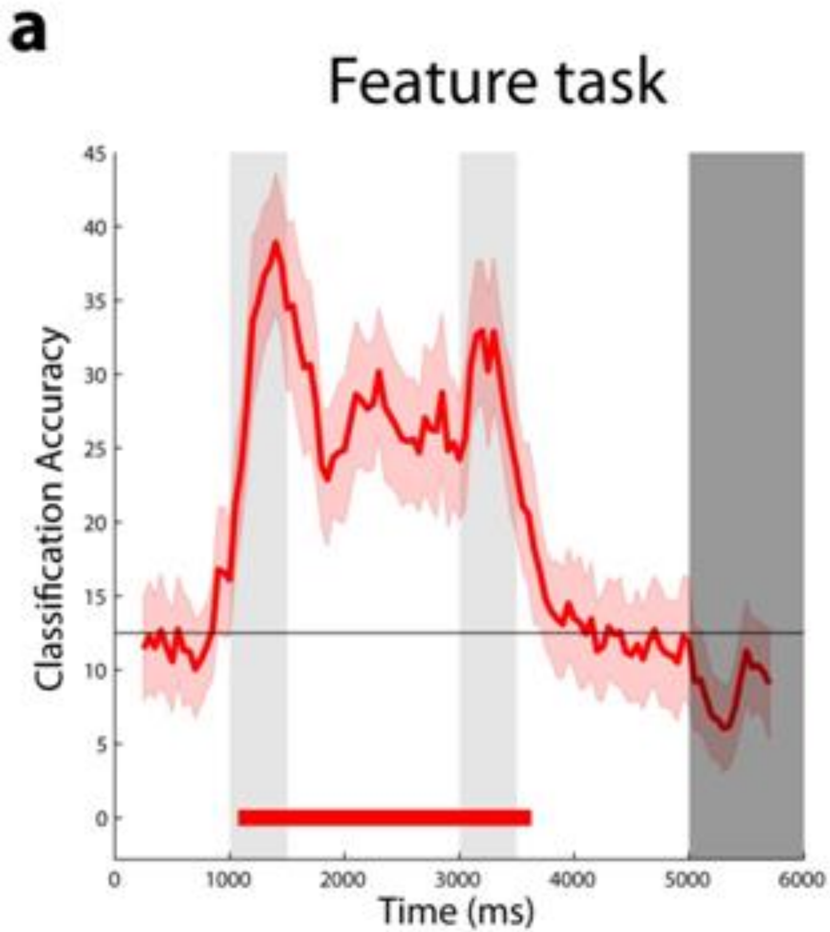

How is information coded in neural activity?

Answer:

Information is contained in a dynamic population code, where at each point in time, a small number of neurons contain all the information present in the larger population.

These neurons contain information about multiple variables.

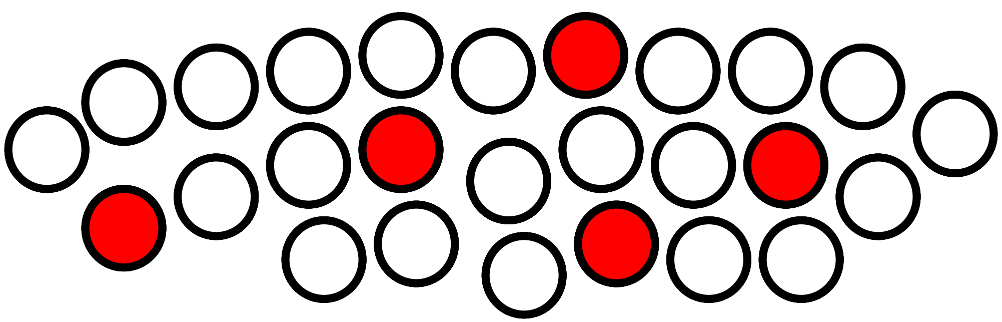

Are there regions differences within lateral PFC?

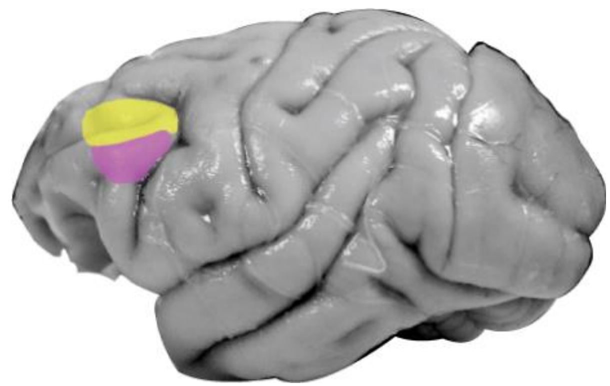

Domain specific theories

(Wilson et al., 1993 Science; O'Scalaidhe et al., Science 1997; Romanski et al., Nat Neuro, 2002)

Integrative theories

(Rao, et al., Science 1997; Rainer et al., PNAS 1998)

The debate continues

(O'Reilly, 2010 TINS, Wilson et al., 2010 TINS).

Dorsolateral vs. ventrolateral PFC visual information

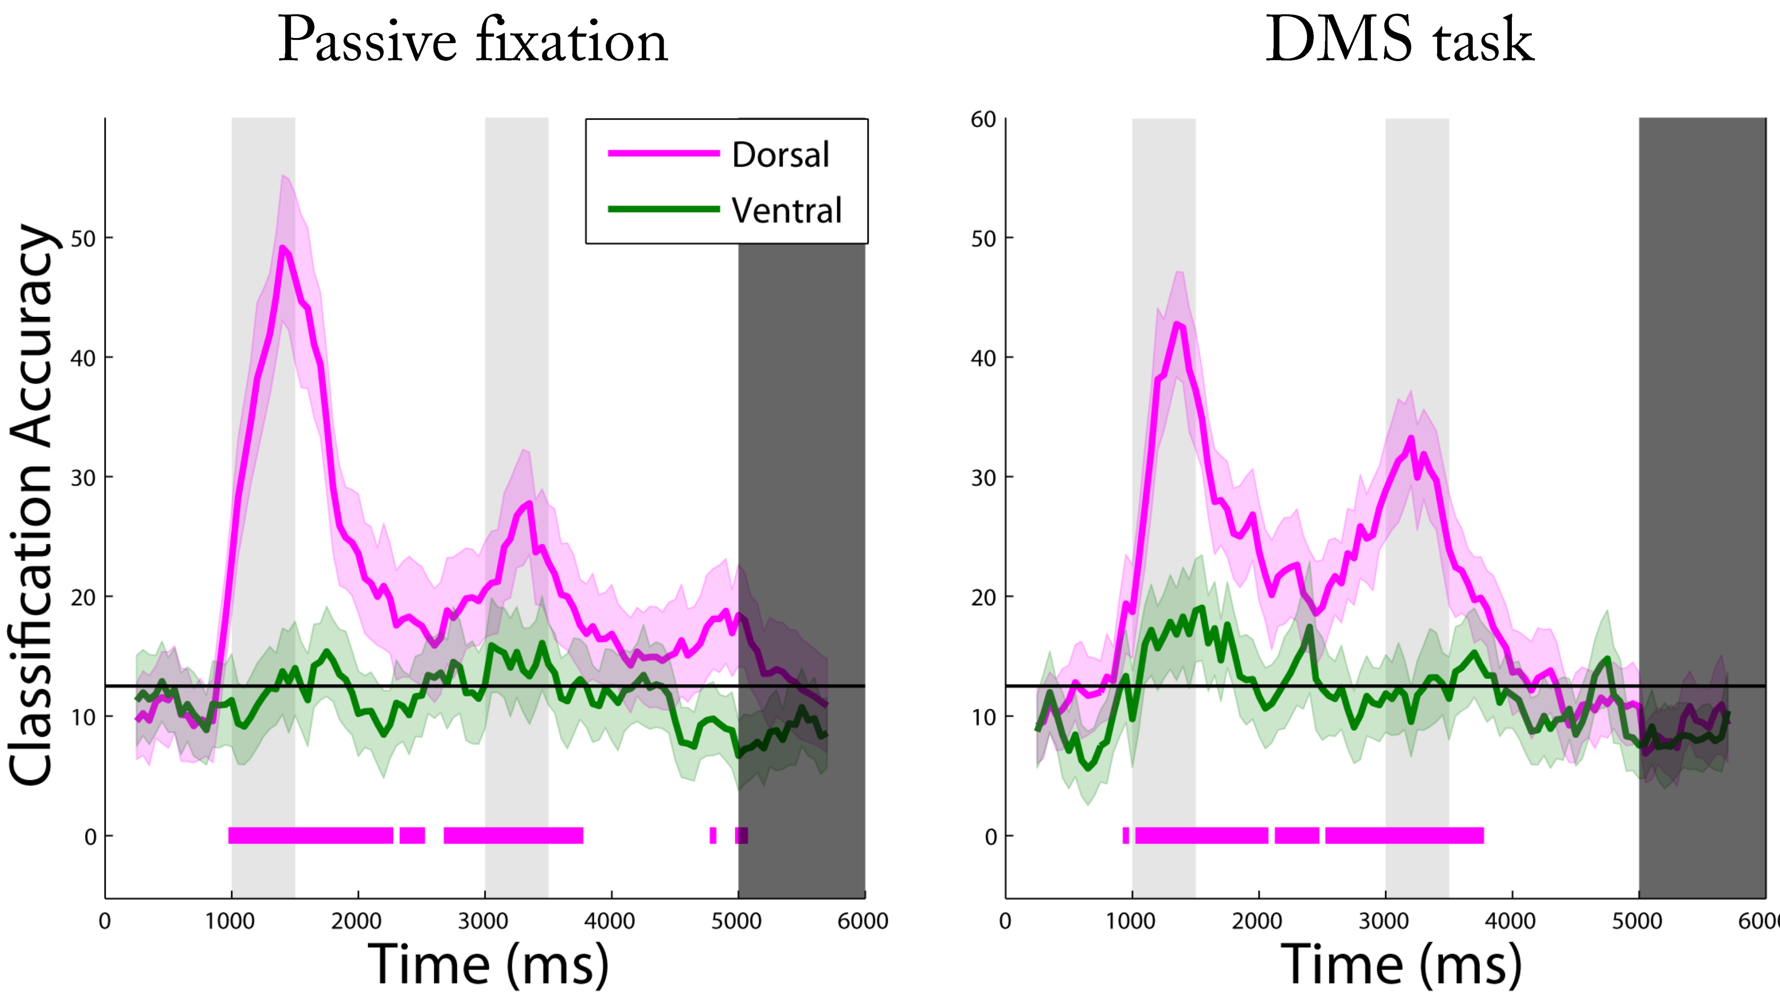

Feature task

Dorsolateral vs. ventrolateral PFC visual information

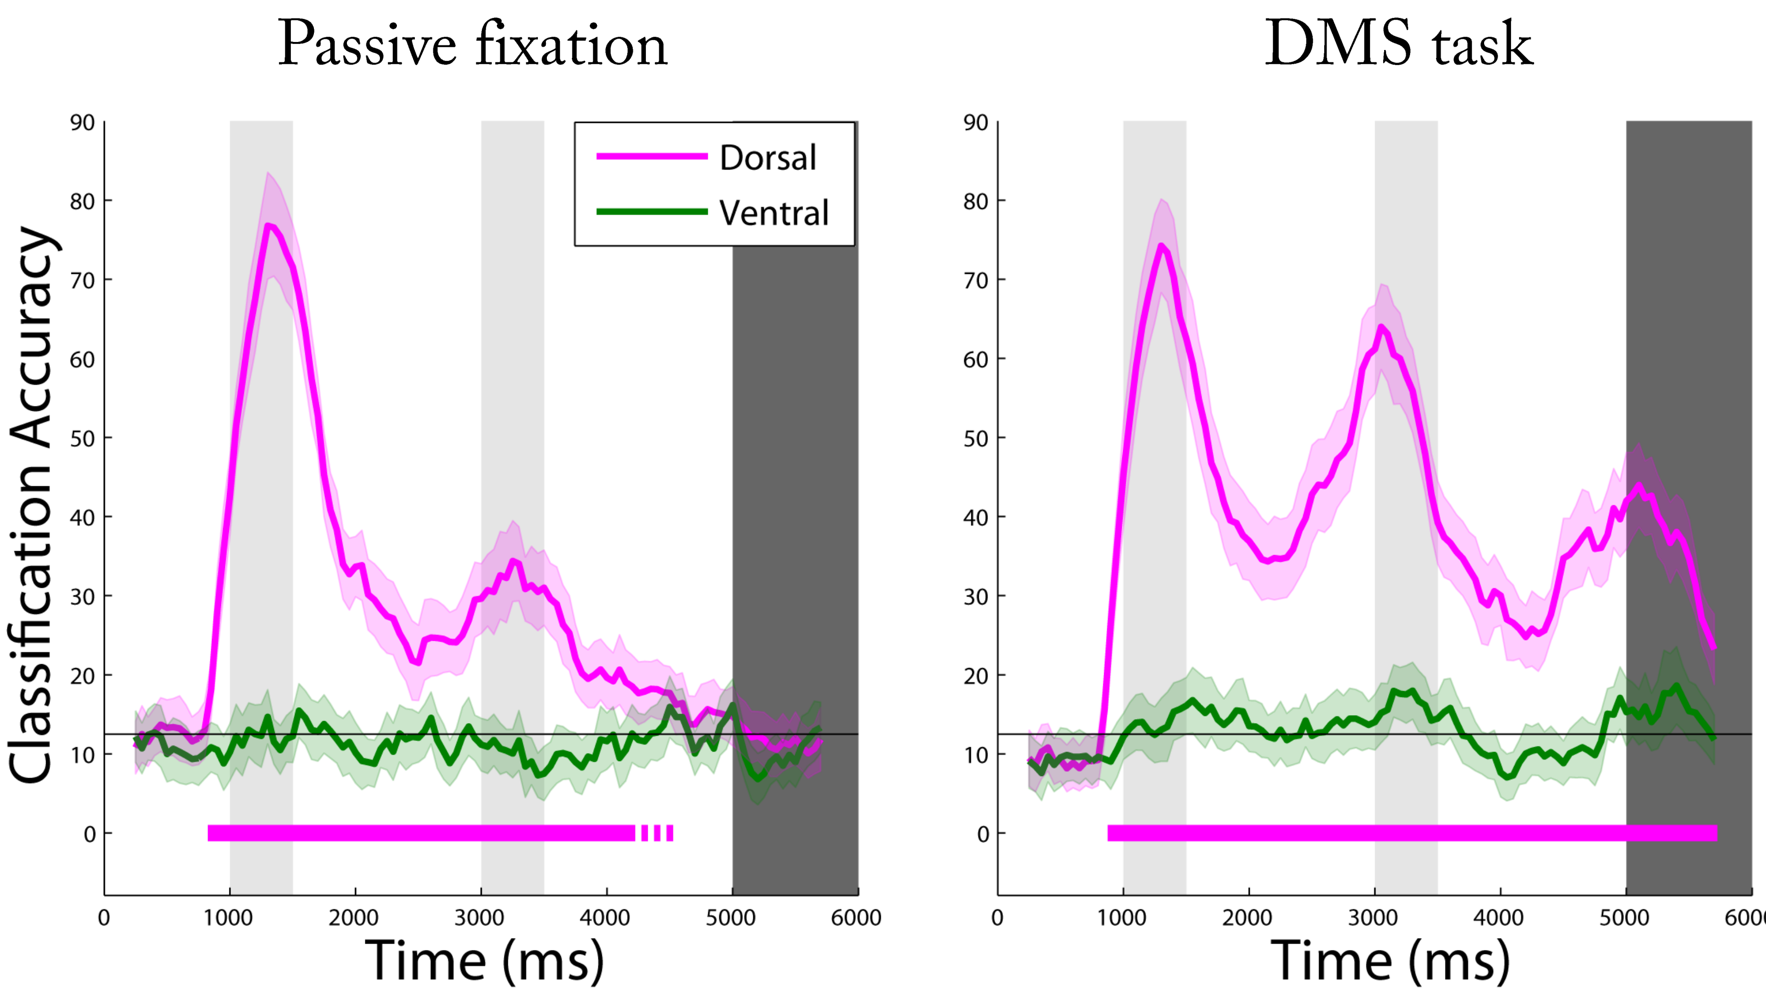

Spatial task

Dorsolateral vs. ventrolateral PFC match/nonmatch information

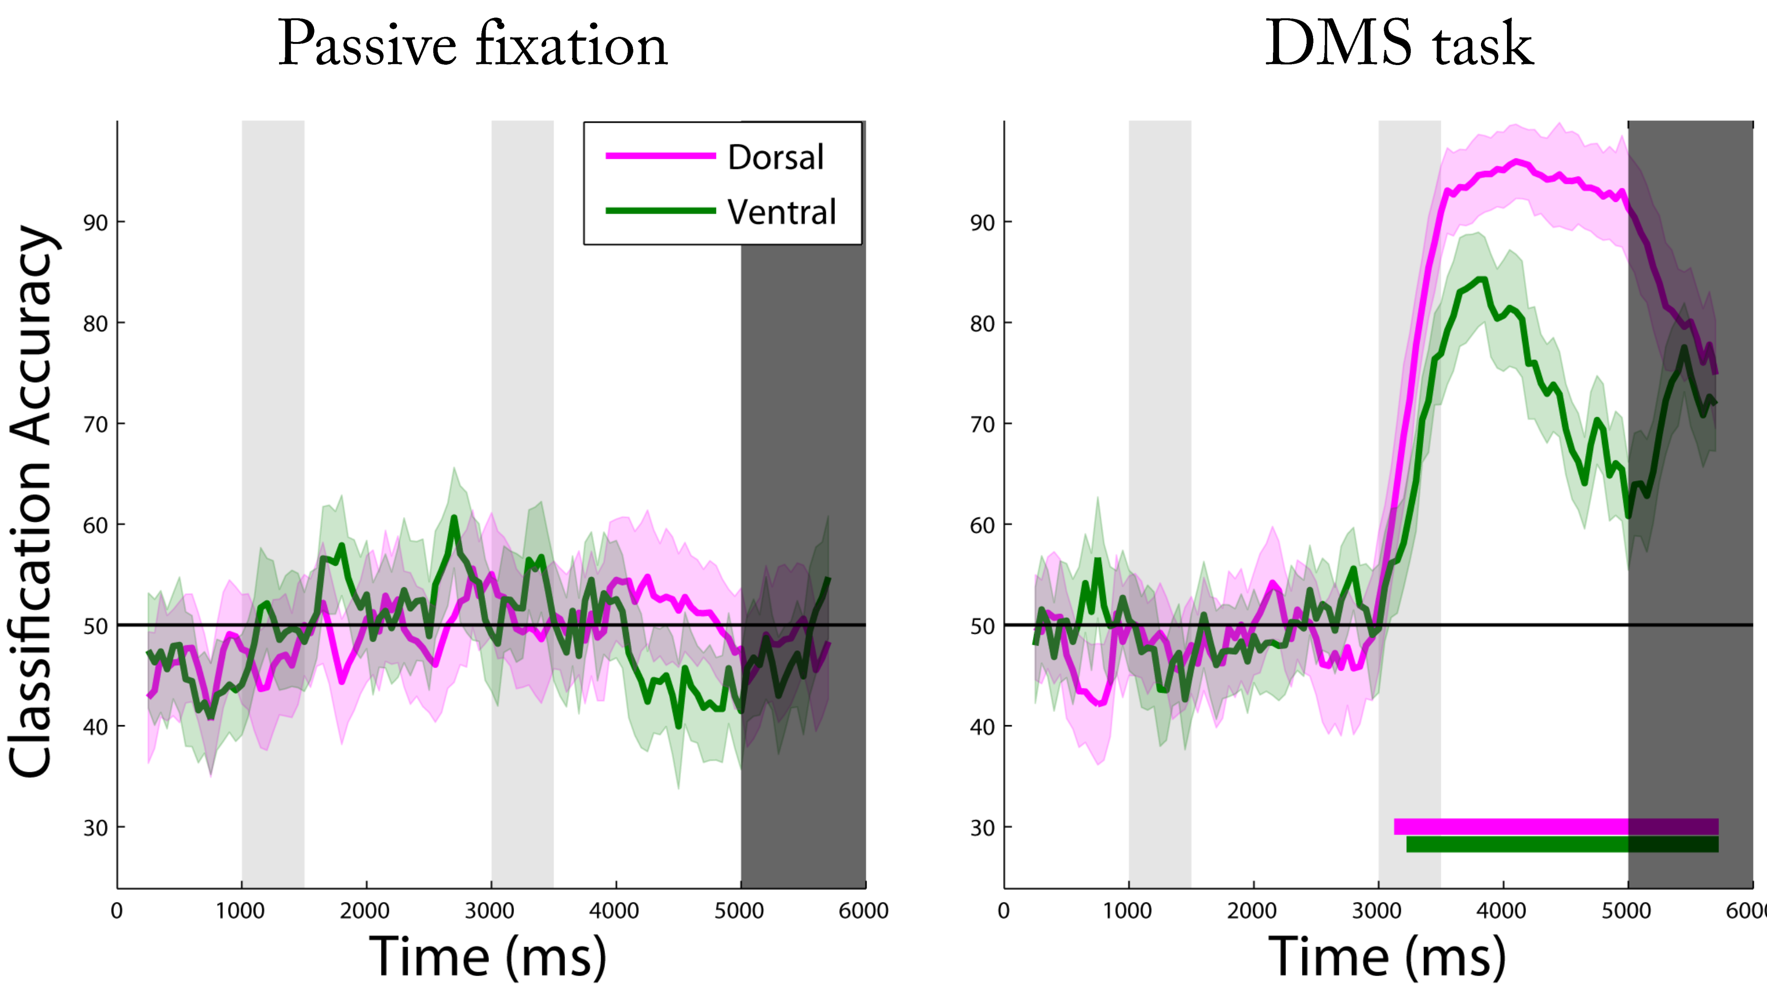

Feature task

Dorsolateral vs. ventrolateral PFC match/nonmatch information

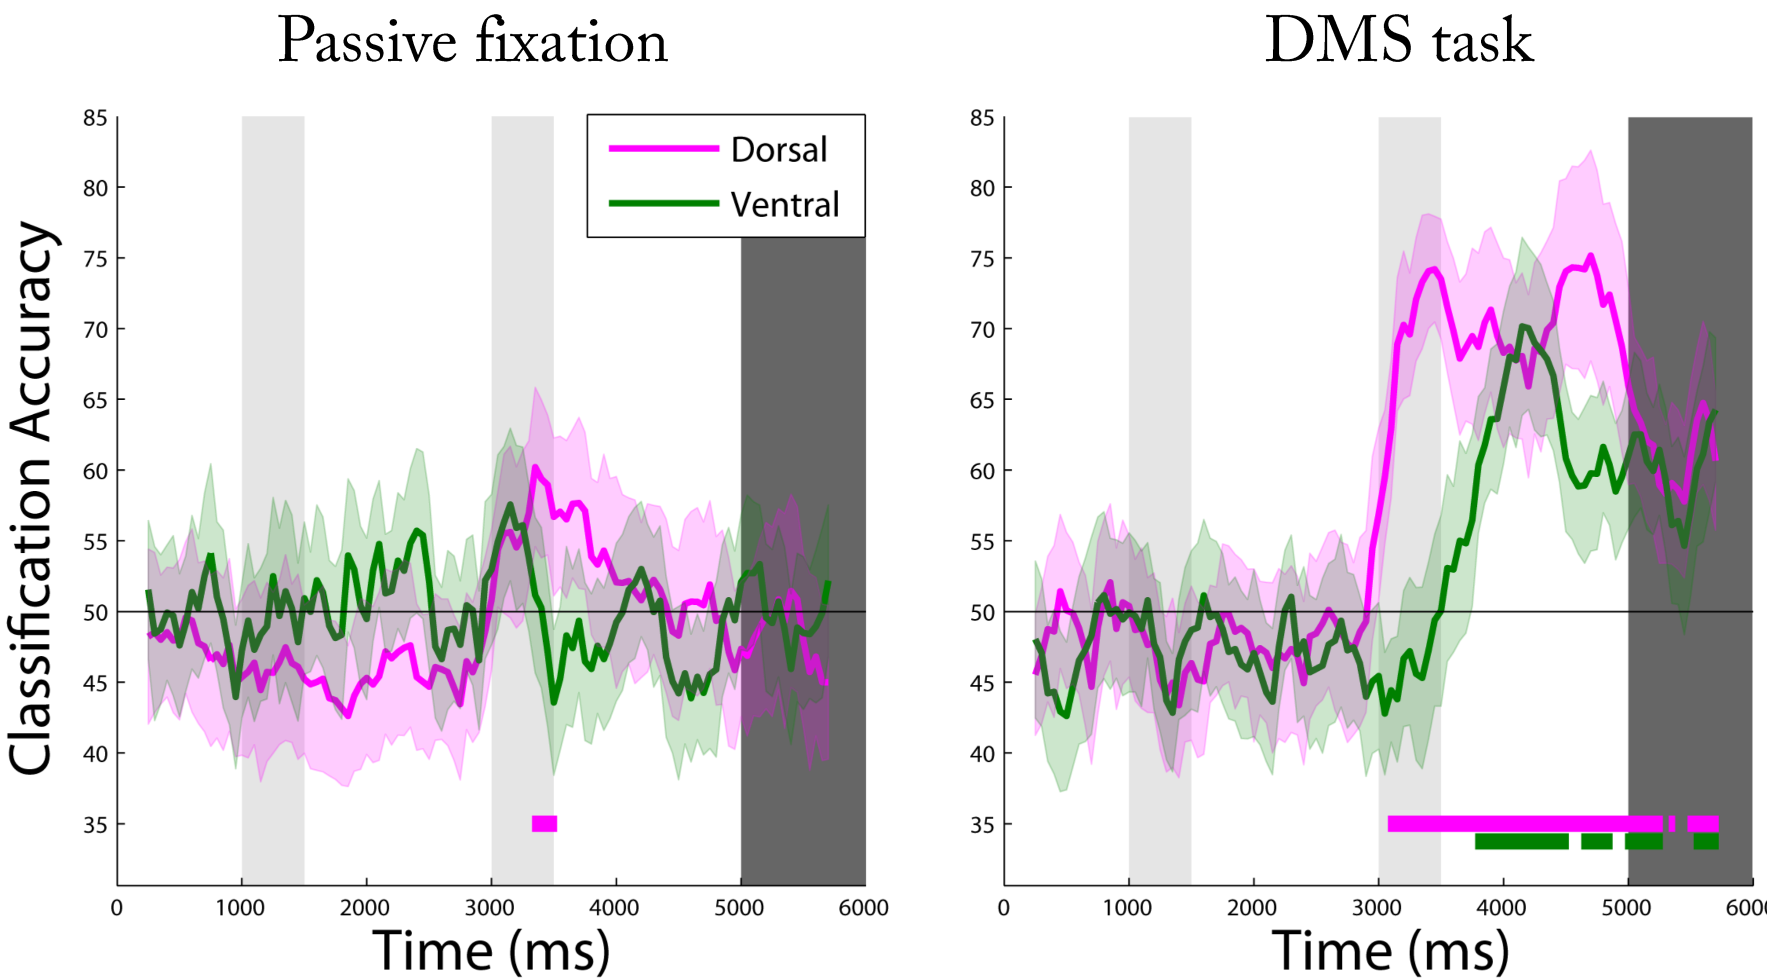

Spatial task

Summary

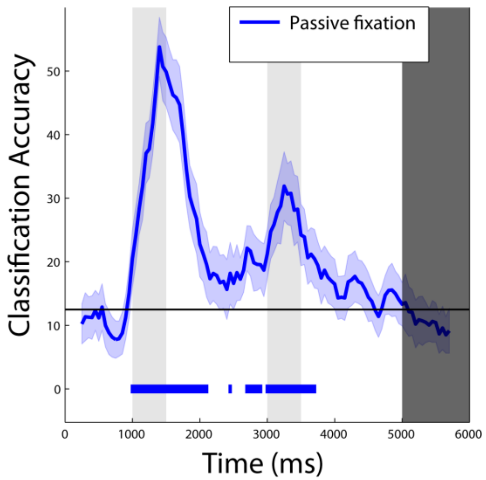

1. How does the information change after learning a new task?

Visual information is largely unchanged; there is a large increase in task relevant information

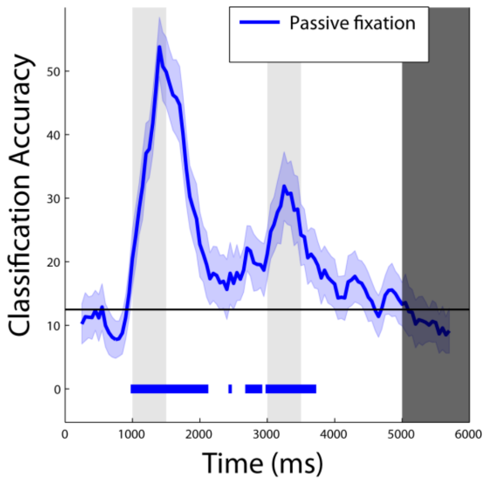

Visual information is largely unchanged; there is a large increase in task-relevant information

Information is contained in a compact neural population code

2. How is information coded in dynamic activity?

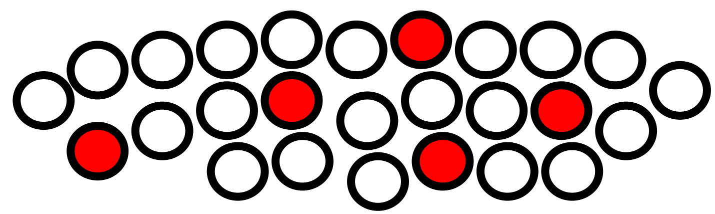

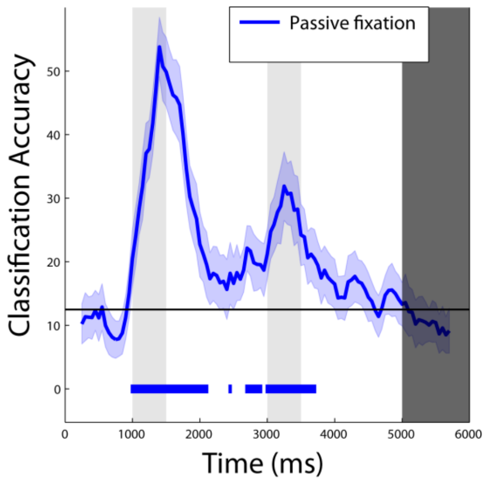

Visual information is largely unchanged; there is a large increase in task-relevant information

Information is contained in a compact dynamic population code

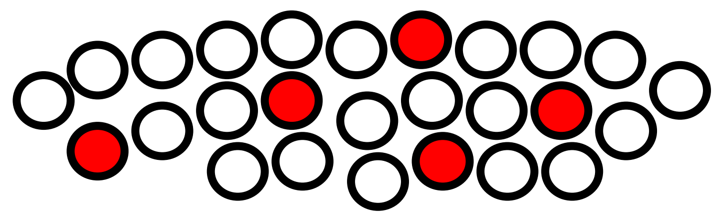

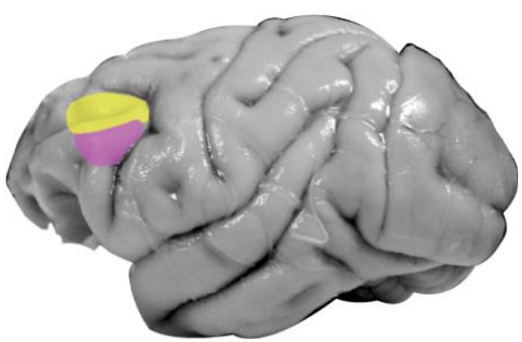

Visual information is primarily in dorsal PFC, while task relevant information is widespread

3. Are there regions differences within dorsolateral PFC?

Conclusions

Using machine learning is a powerful way to analyze neural data

Thanks to:

B. Jarosiewicz for providing some slides

ONLINE SUBJECT EVALUATIONS ARE NOW OPEN

http://web.mit.edu/subjectevaluation

- You have until Monday, Dec. 16 at 9 AM
- Please evaluate all subjects in your list
- Don’t forget your TAs
- Write comments

Your feedback is read and valued!

Also, class projects due December 11th
email them to: shimon.ullman@gmail.com