Multi-view Representation Learning with Applications to Functional Neuroimaging Data

Po-Hsuan (Cameron) Chen

Final Public Oral
June 20, 2017

PRINCETON UNIVERSITY

How does the human brain work?

functional Magnetic Resonance Imaging (fMRI)

Functional magnetic resonance imaging (fMRI) data

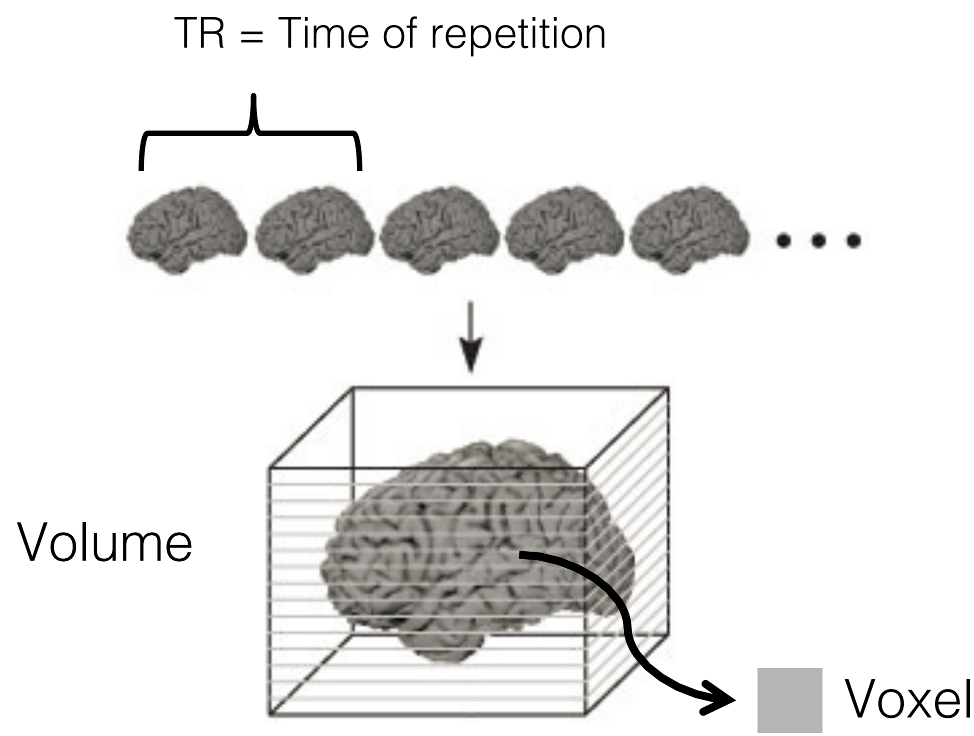

Data collection

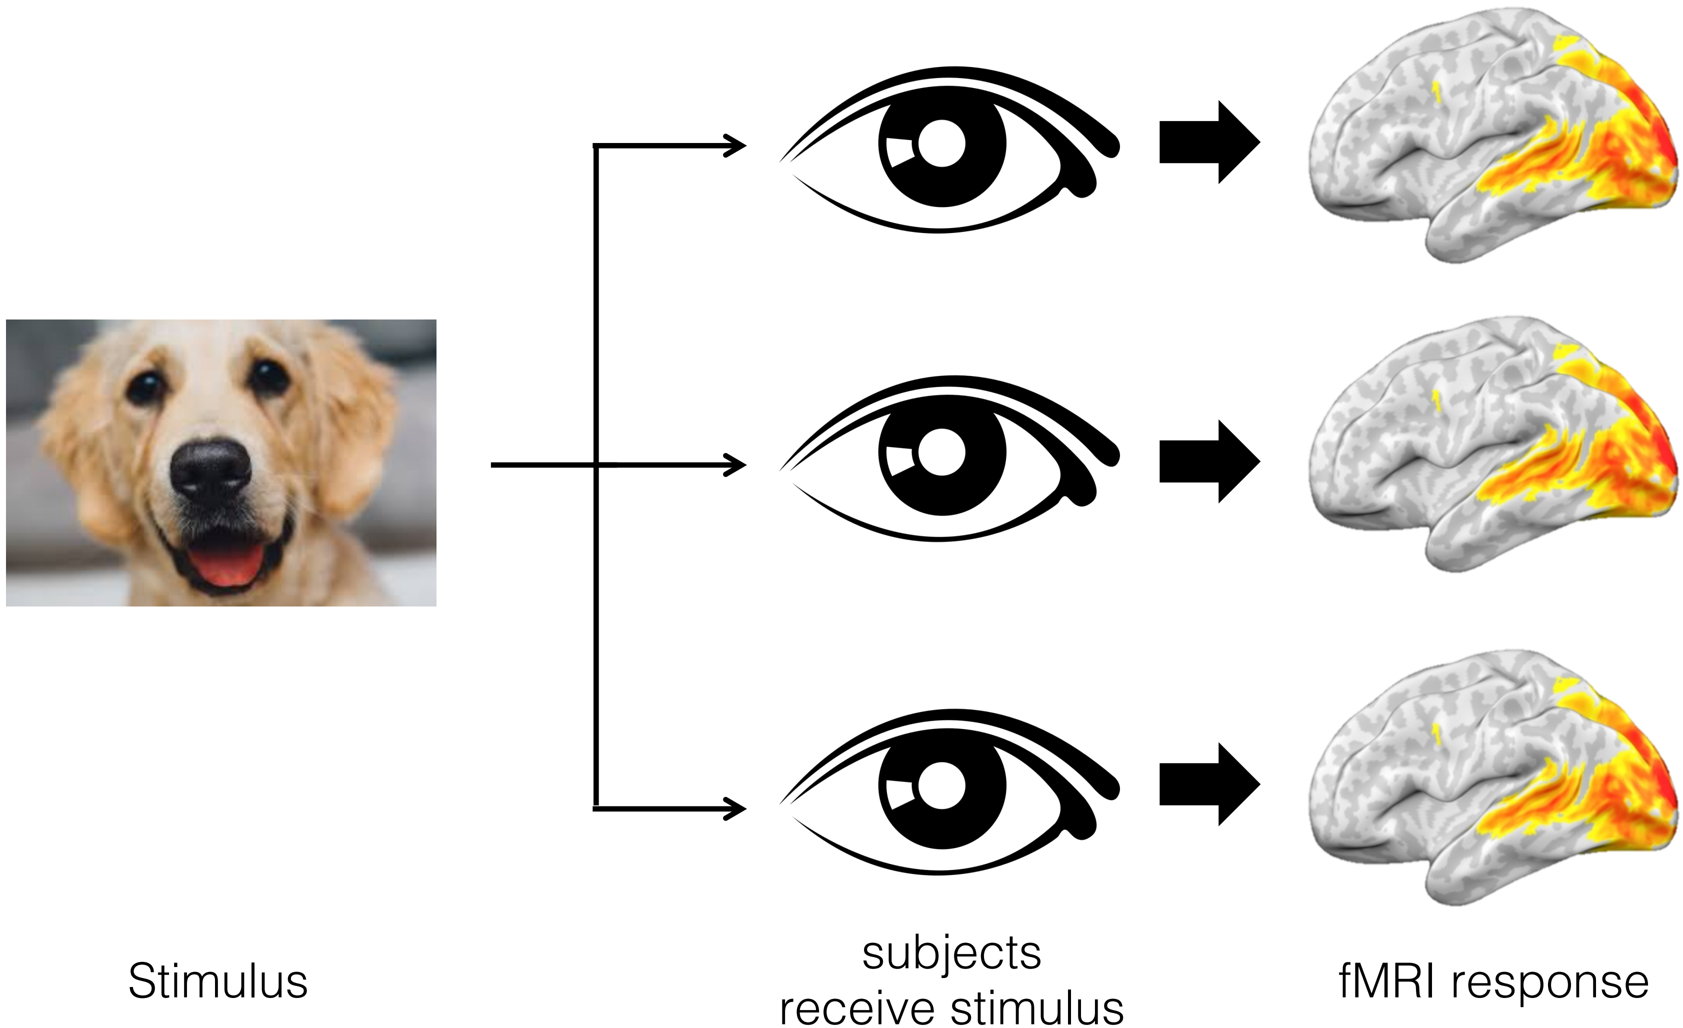

Three interesting problems

fMRI to Stimulus
(decoding)

Stimulus to fMRI
(encoding)

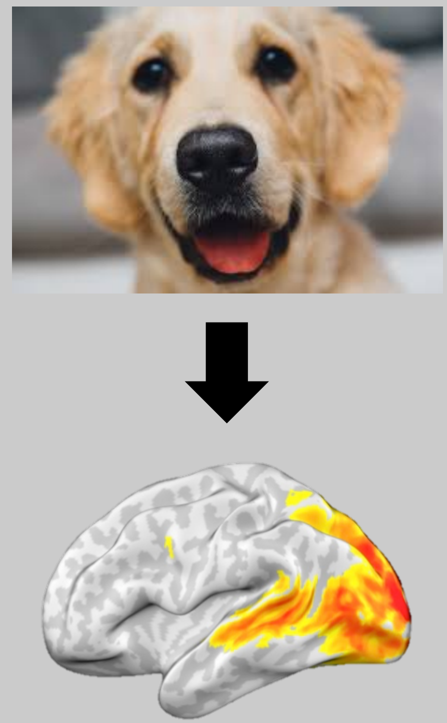

fMRI to fMRI

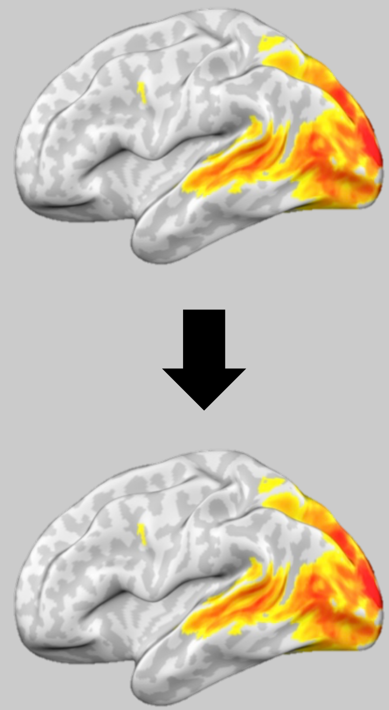

A coherent multi-view framework for all three problems

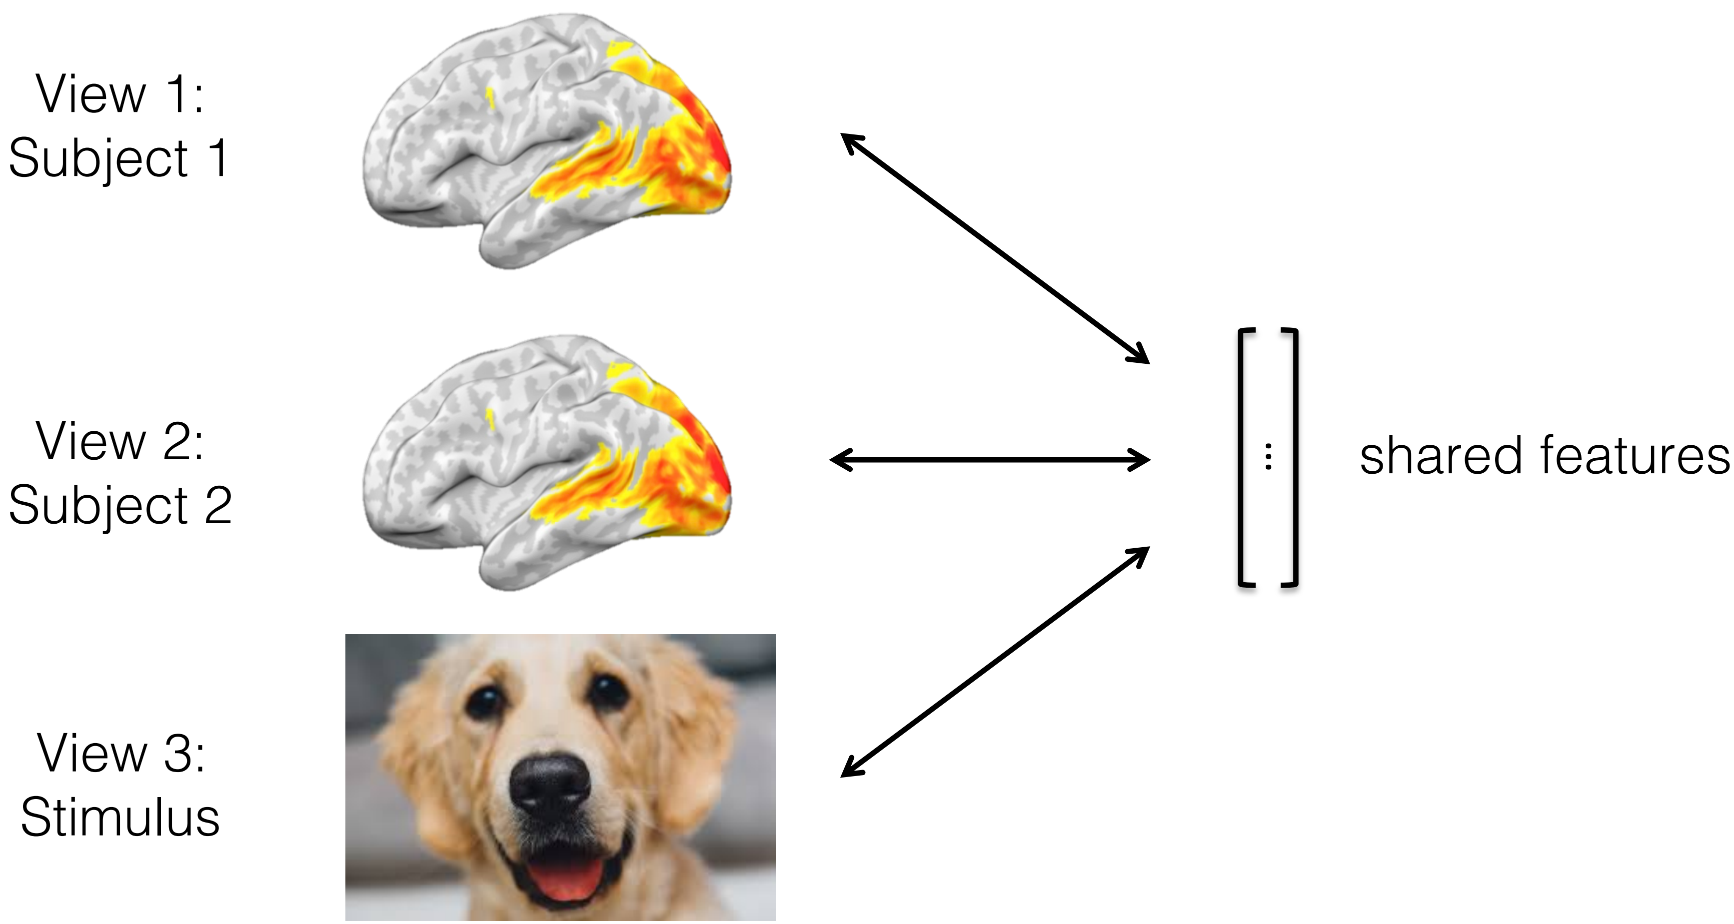

Outline

I. A Shared Response Model (SRM)

II. SRM on Neuroimaging Data

III. Discussions and Extensions of SRM

IV. Conclusion

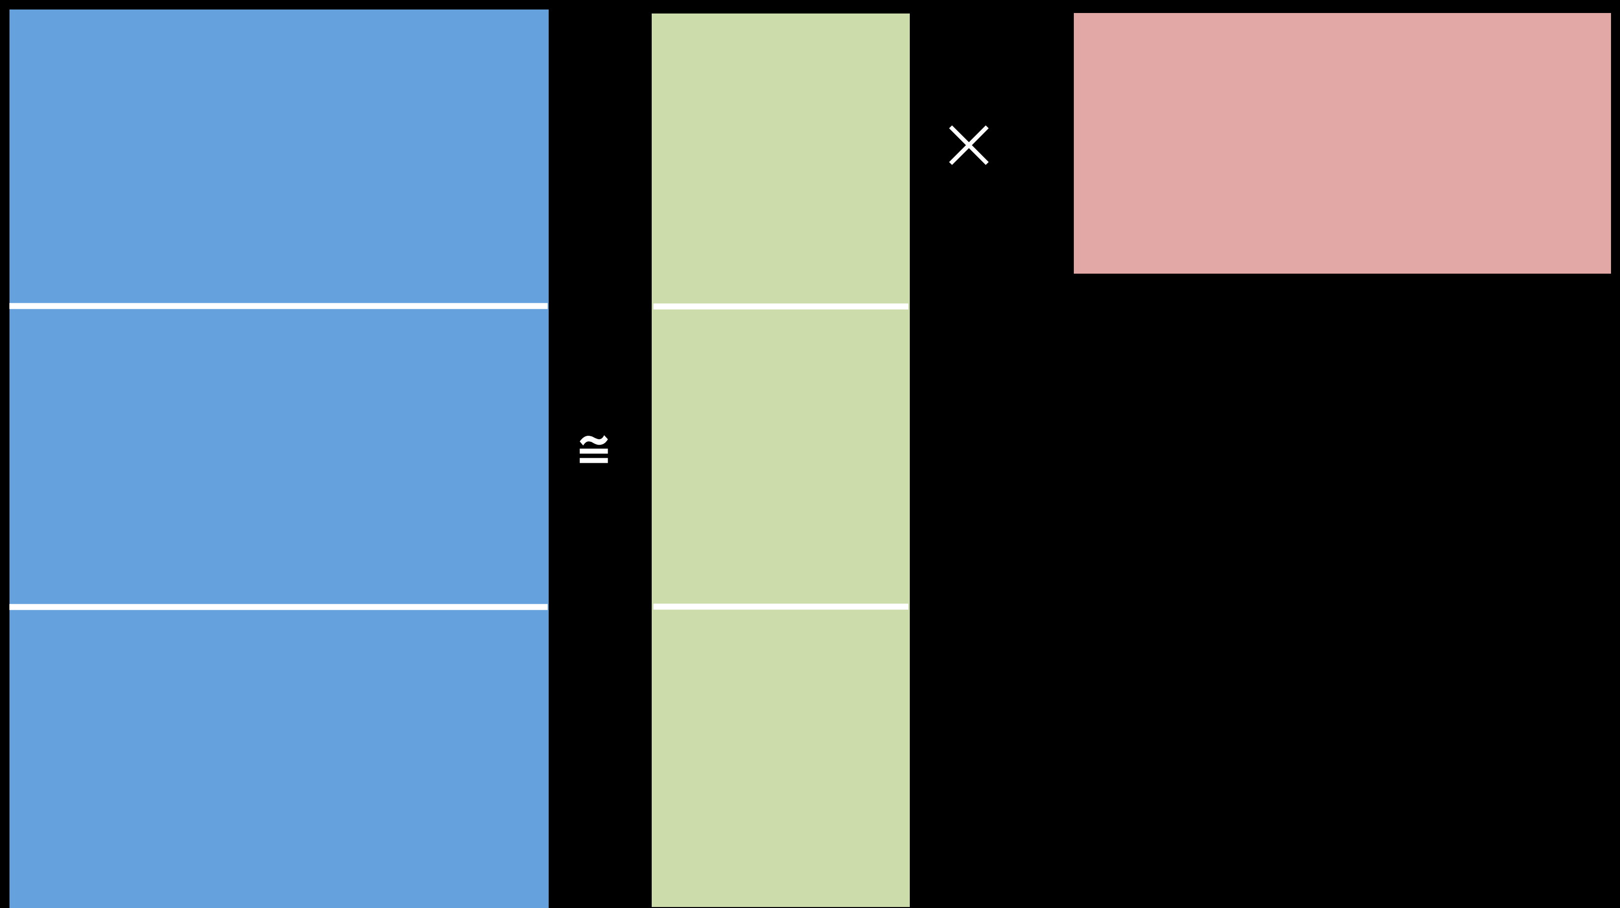

Part I: A Shared Response Model

From a multi-view perspective

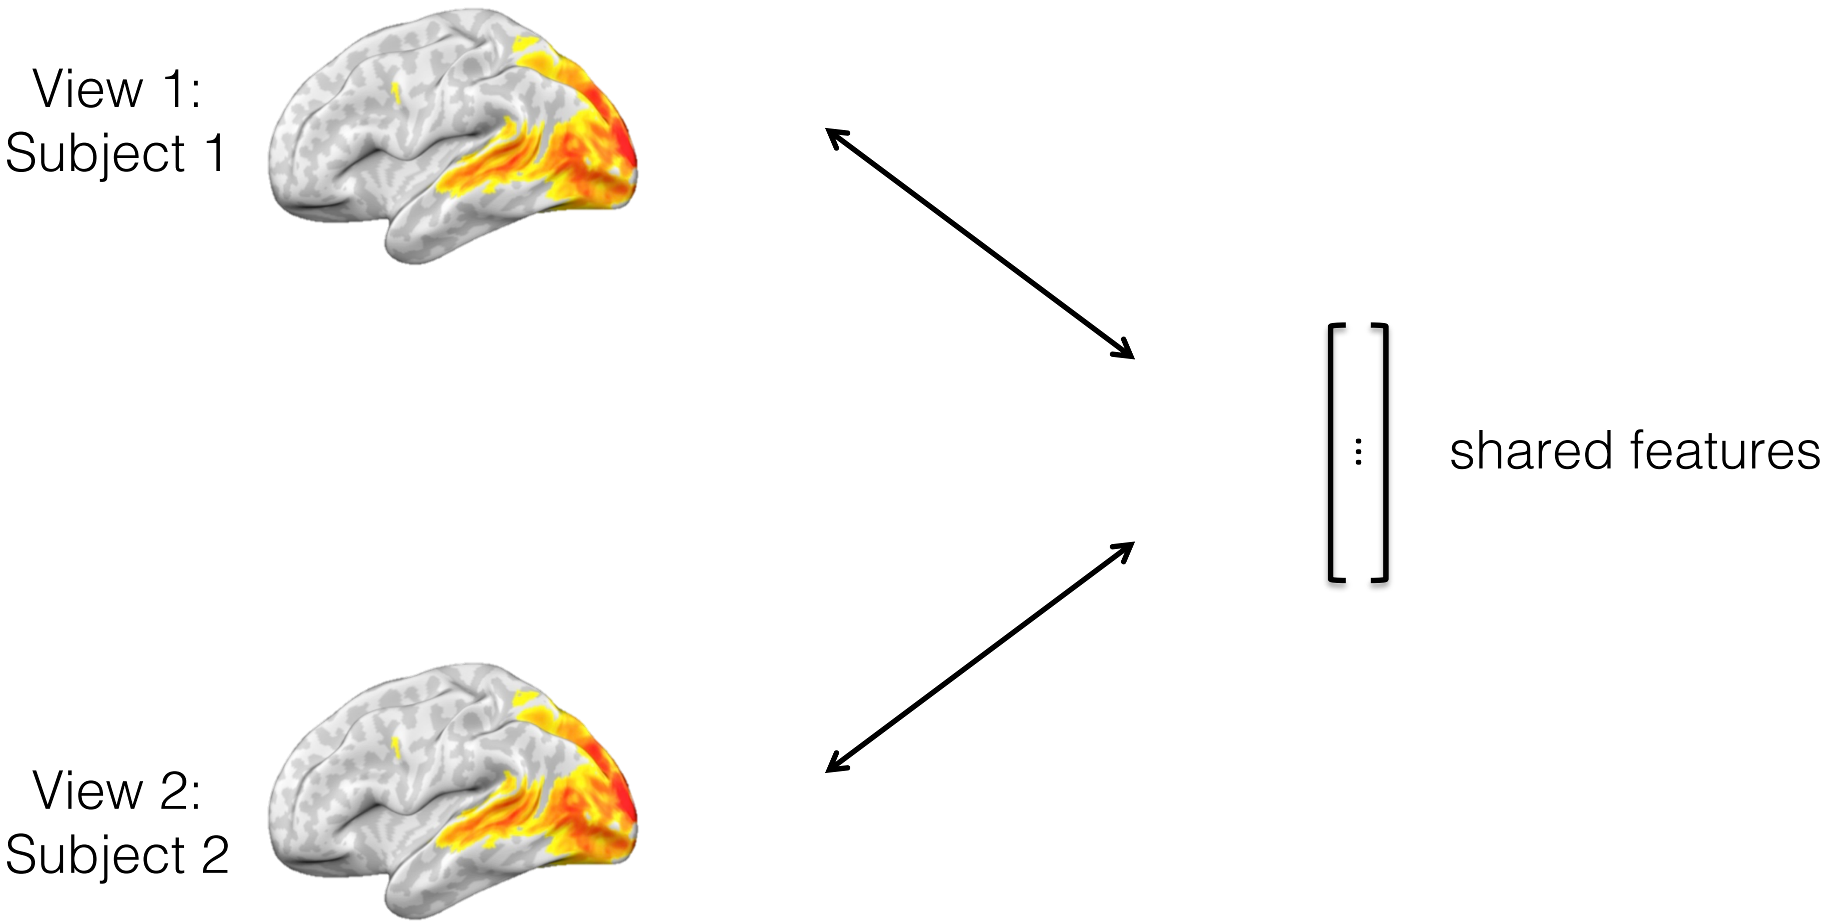

From a multi-view perspective

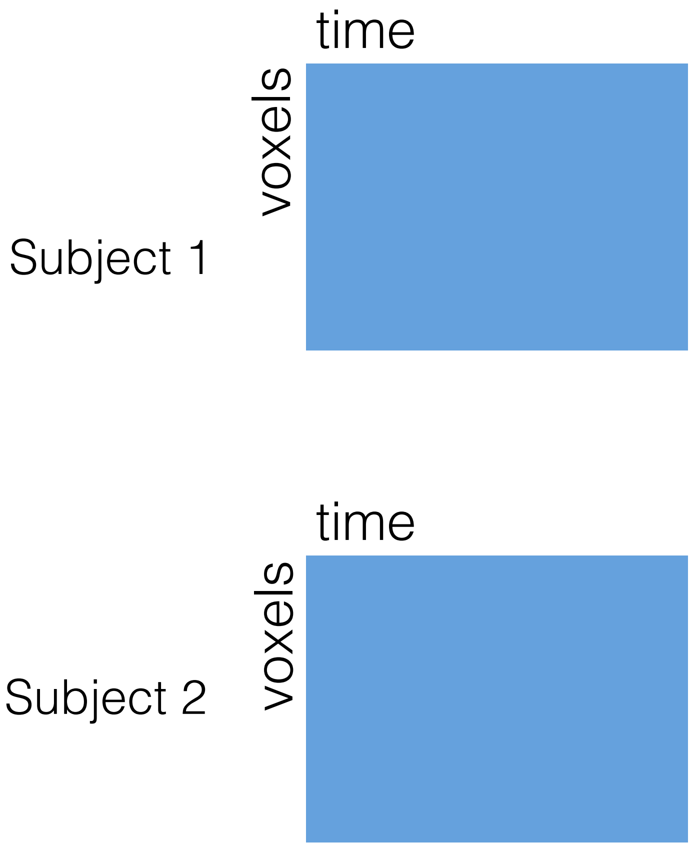

fMRI data matrix

From a multi-view perspective

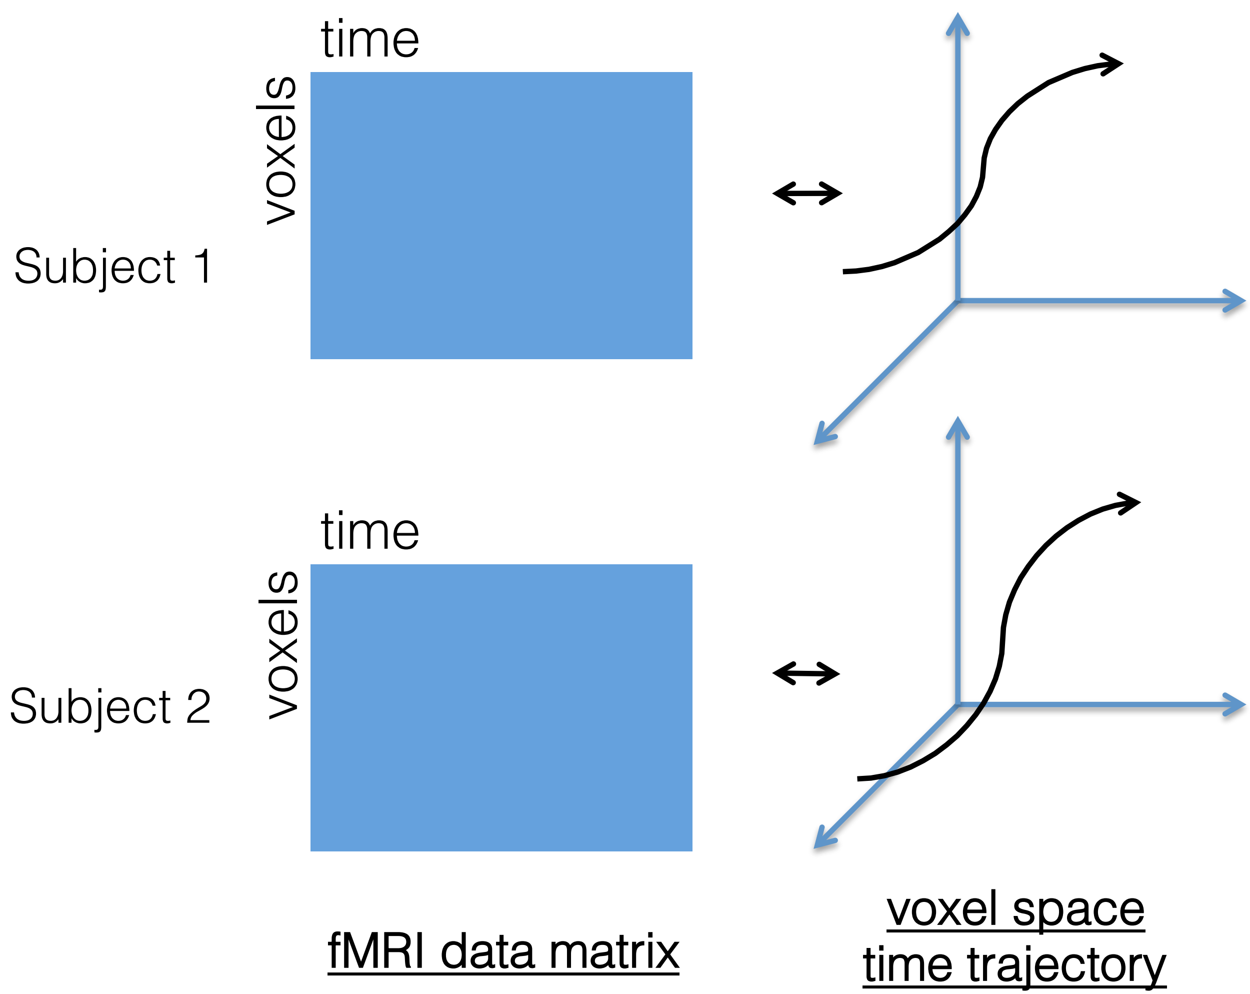

From a multi-view perspective

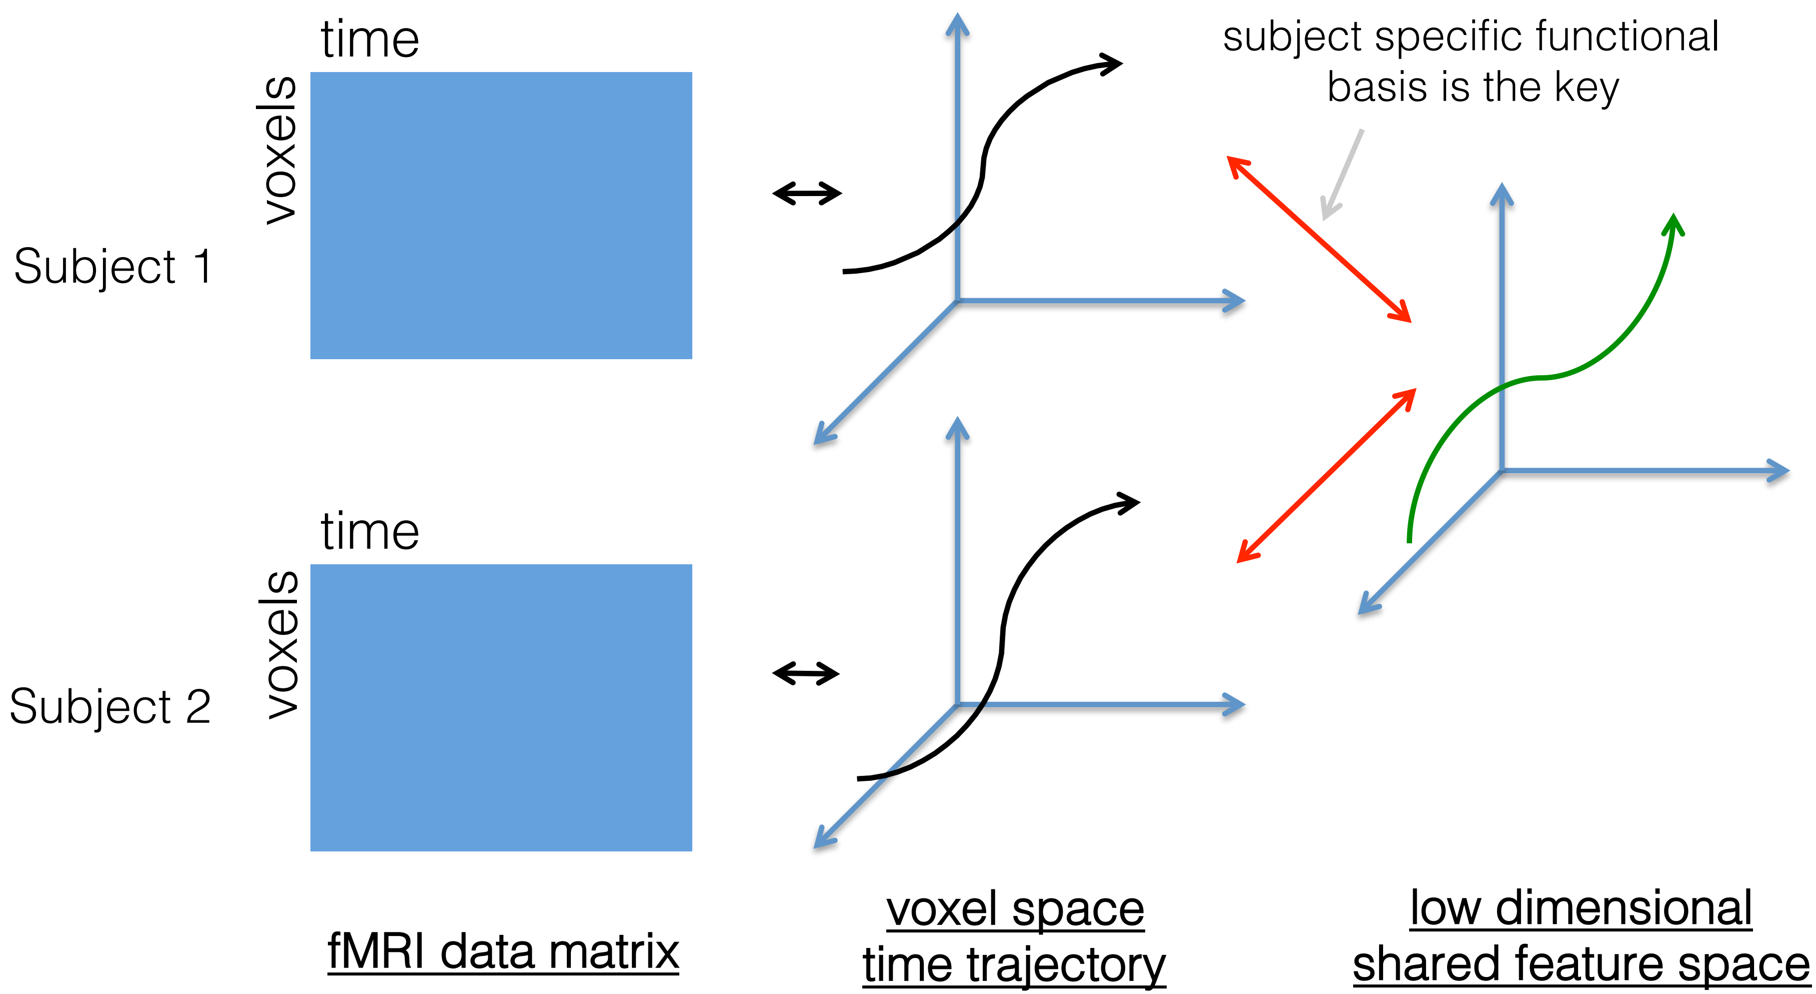

Data collected while subjects receiving stimulus

Temporally synchronized naturalistic stimuli

1. Sample a wide range of response from the subject
2. Use time as anchor for learning shared response

sherlock

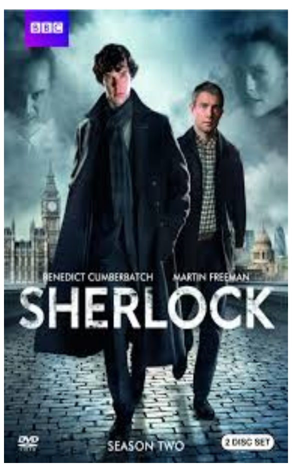

movie watching

raider

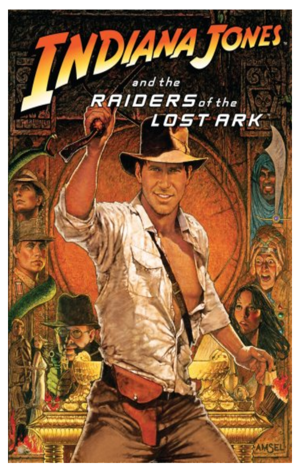

movie and image watching

forrest

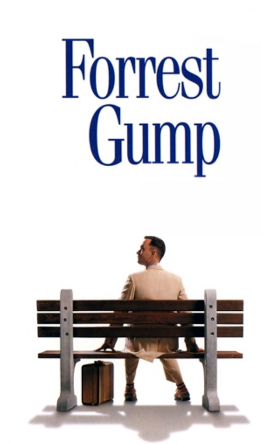

auditory film listening

audiobook

audio book listening

Factor Model

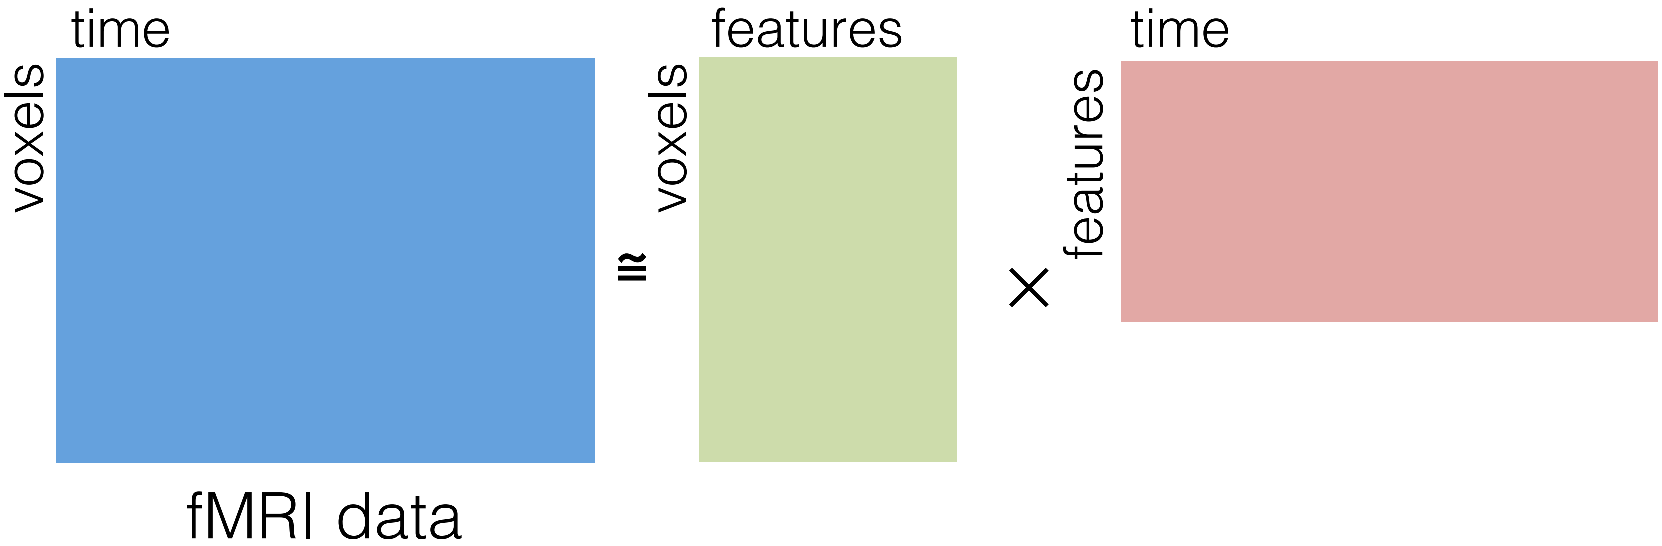

fMRI response as linear combination of functional basis

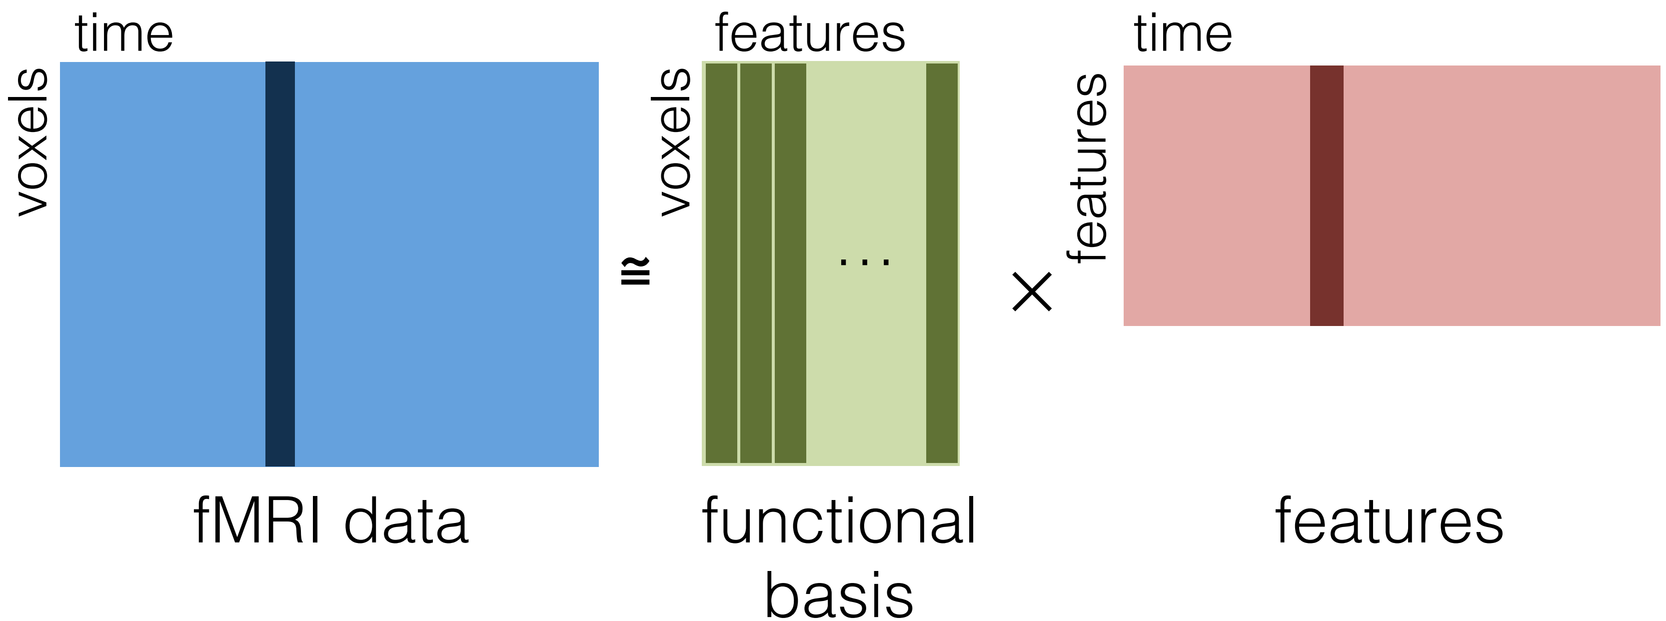

Learning what is shared across subjects

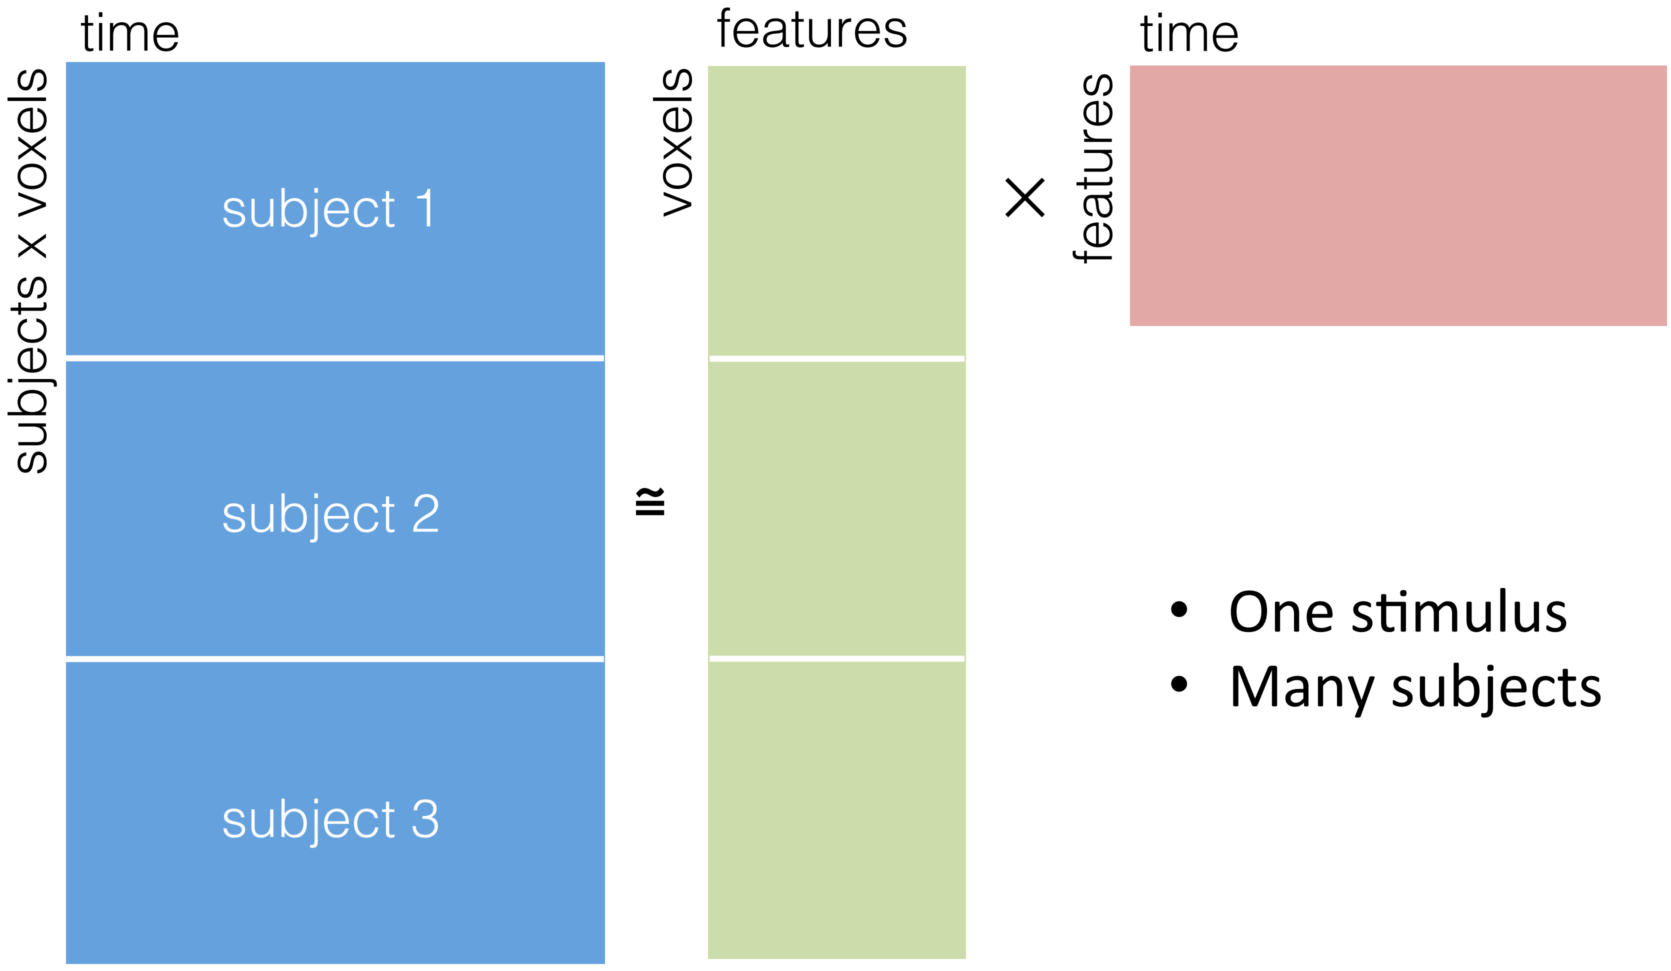

fMRI data

subject specific basis

shared features

- One stimulus
- Many subjects

fMRI data as linear combination of subject specific basis

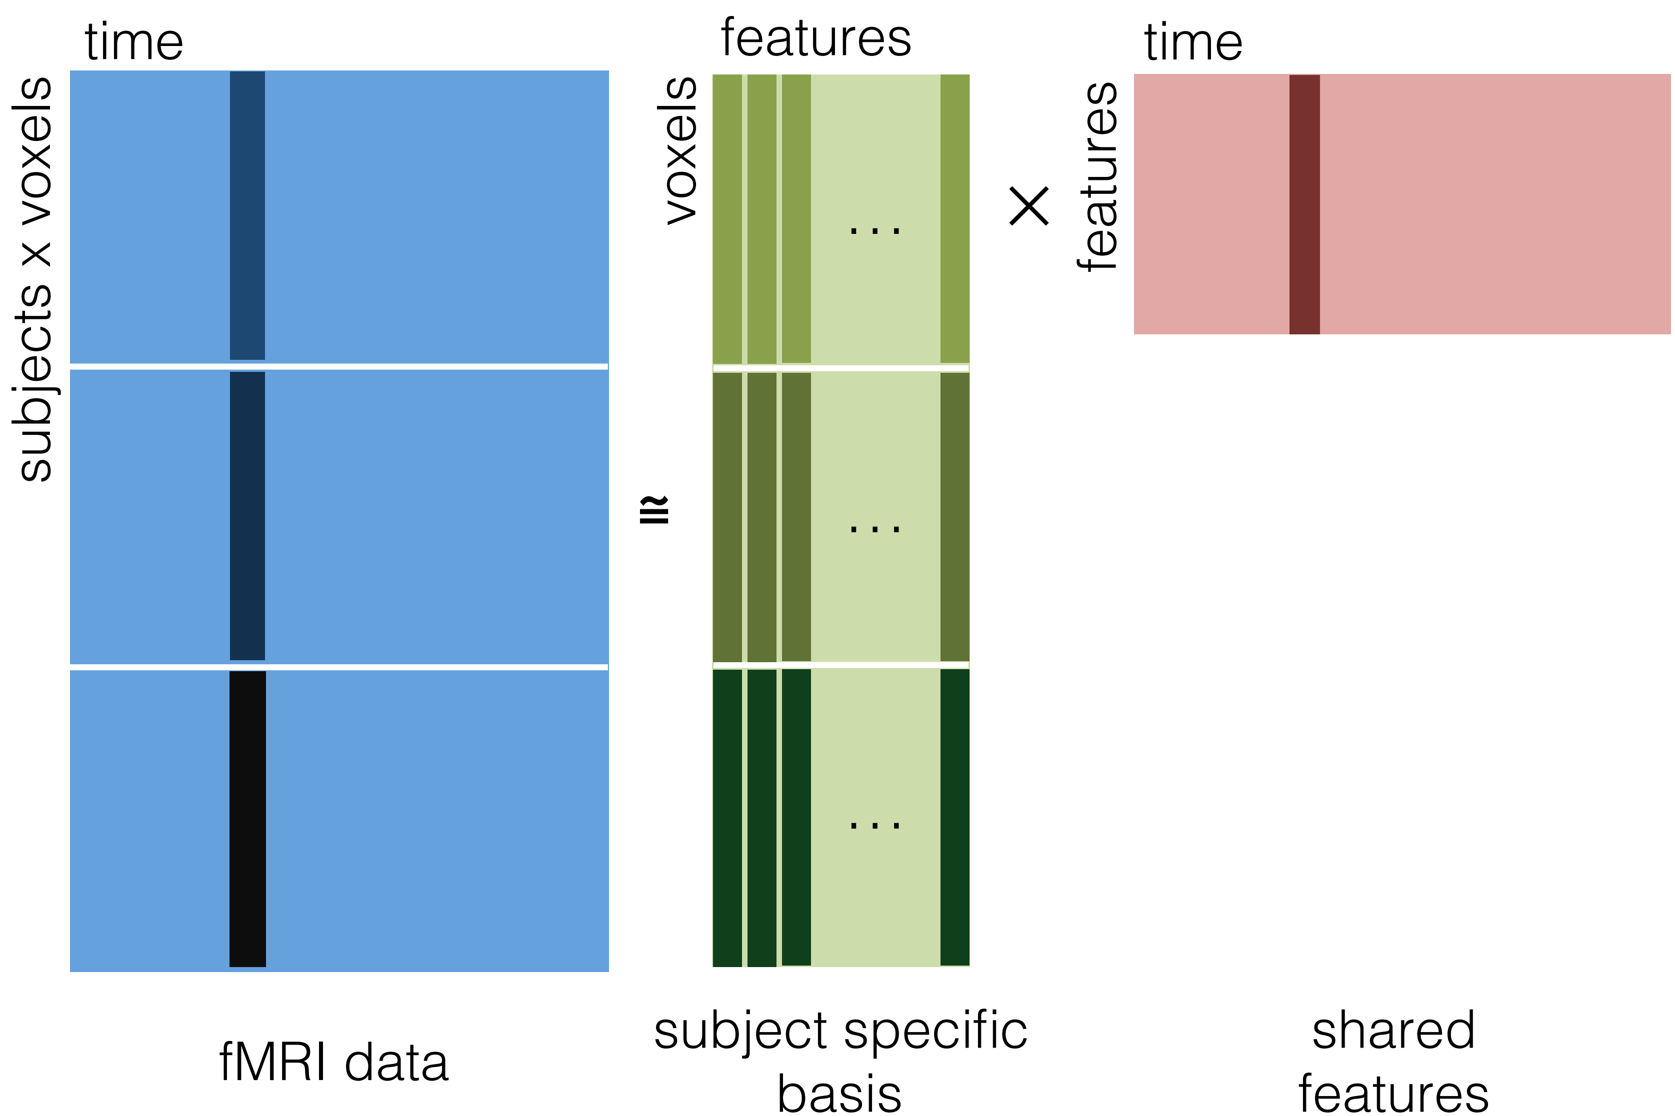

Shared Response Model in one figure

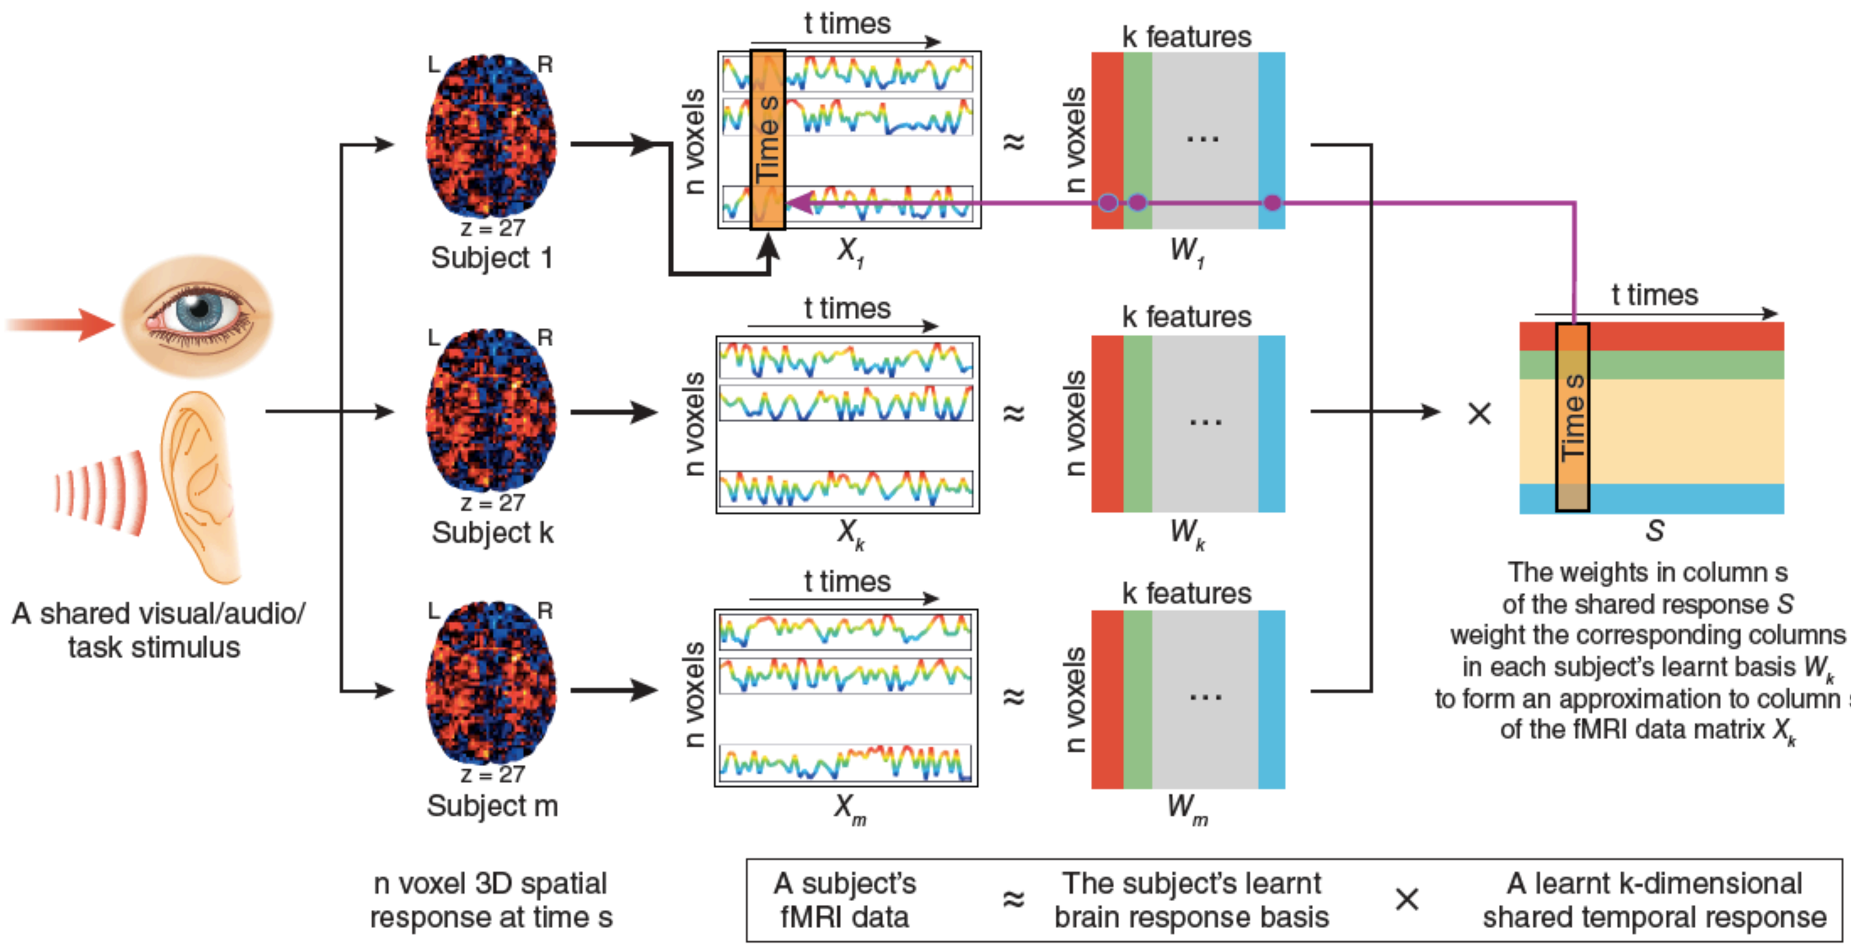

[J. Cohen et al., Nature Neuroscience, 2017]

A generative model

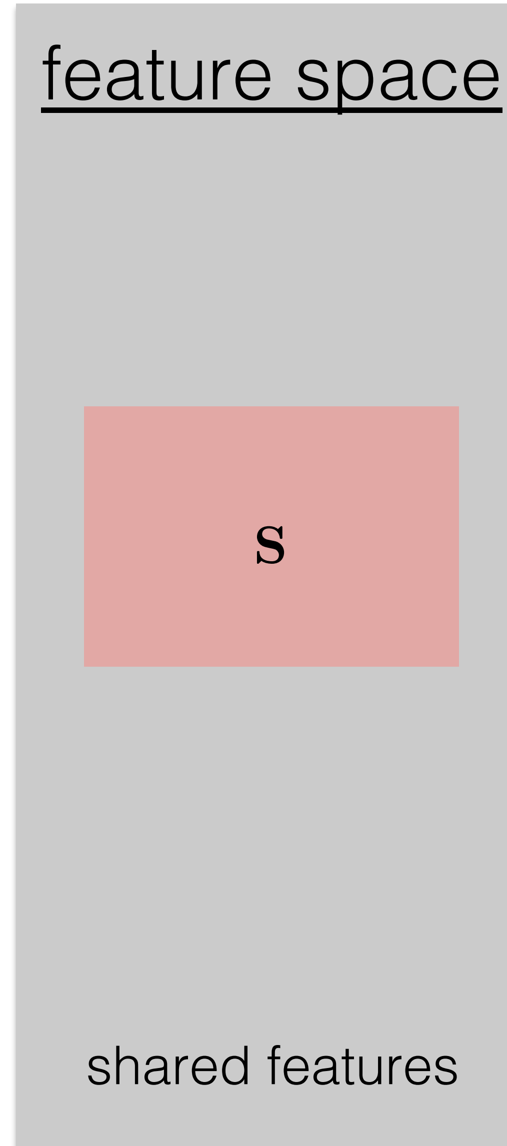

A generative model

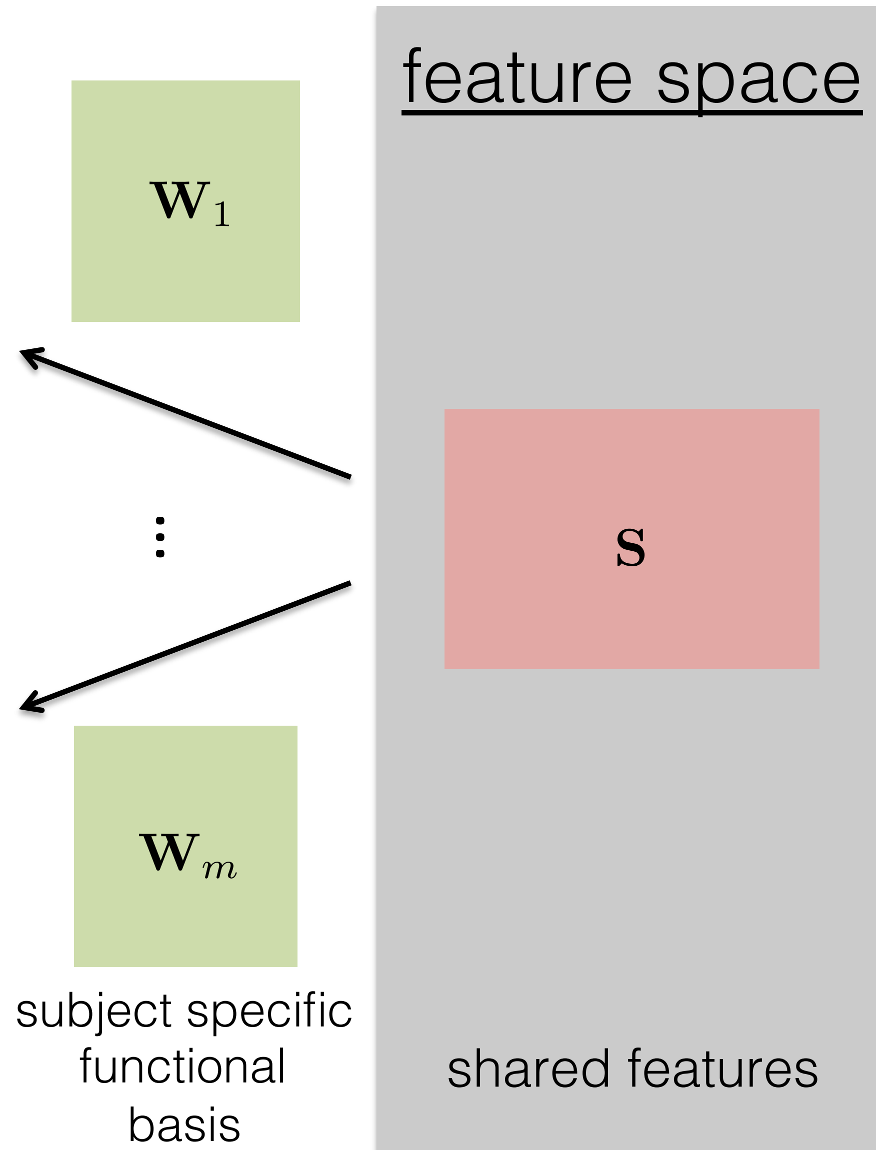

A generative model

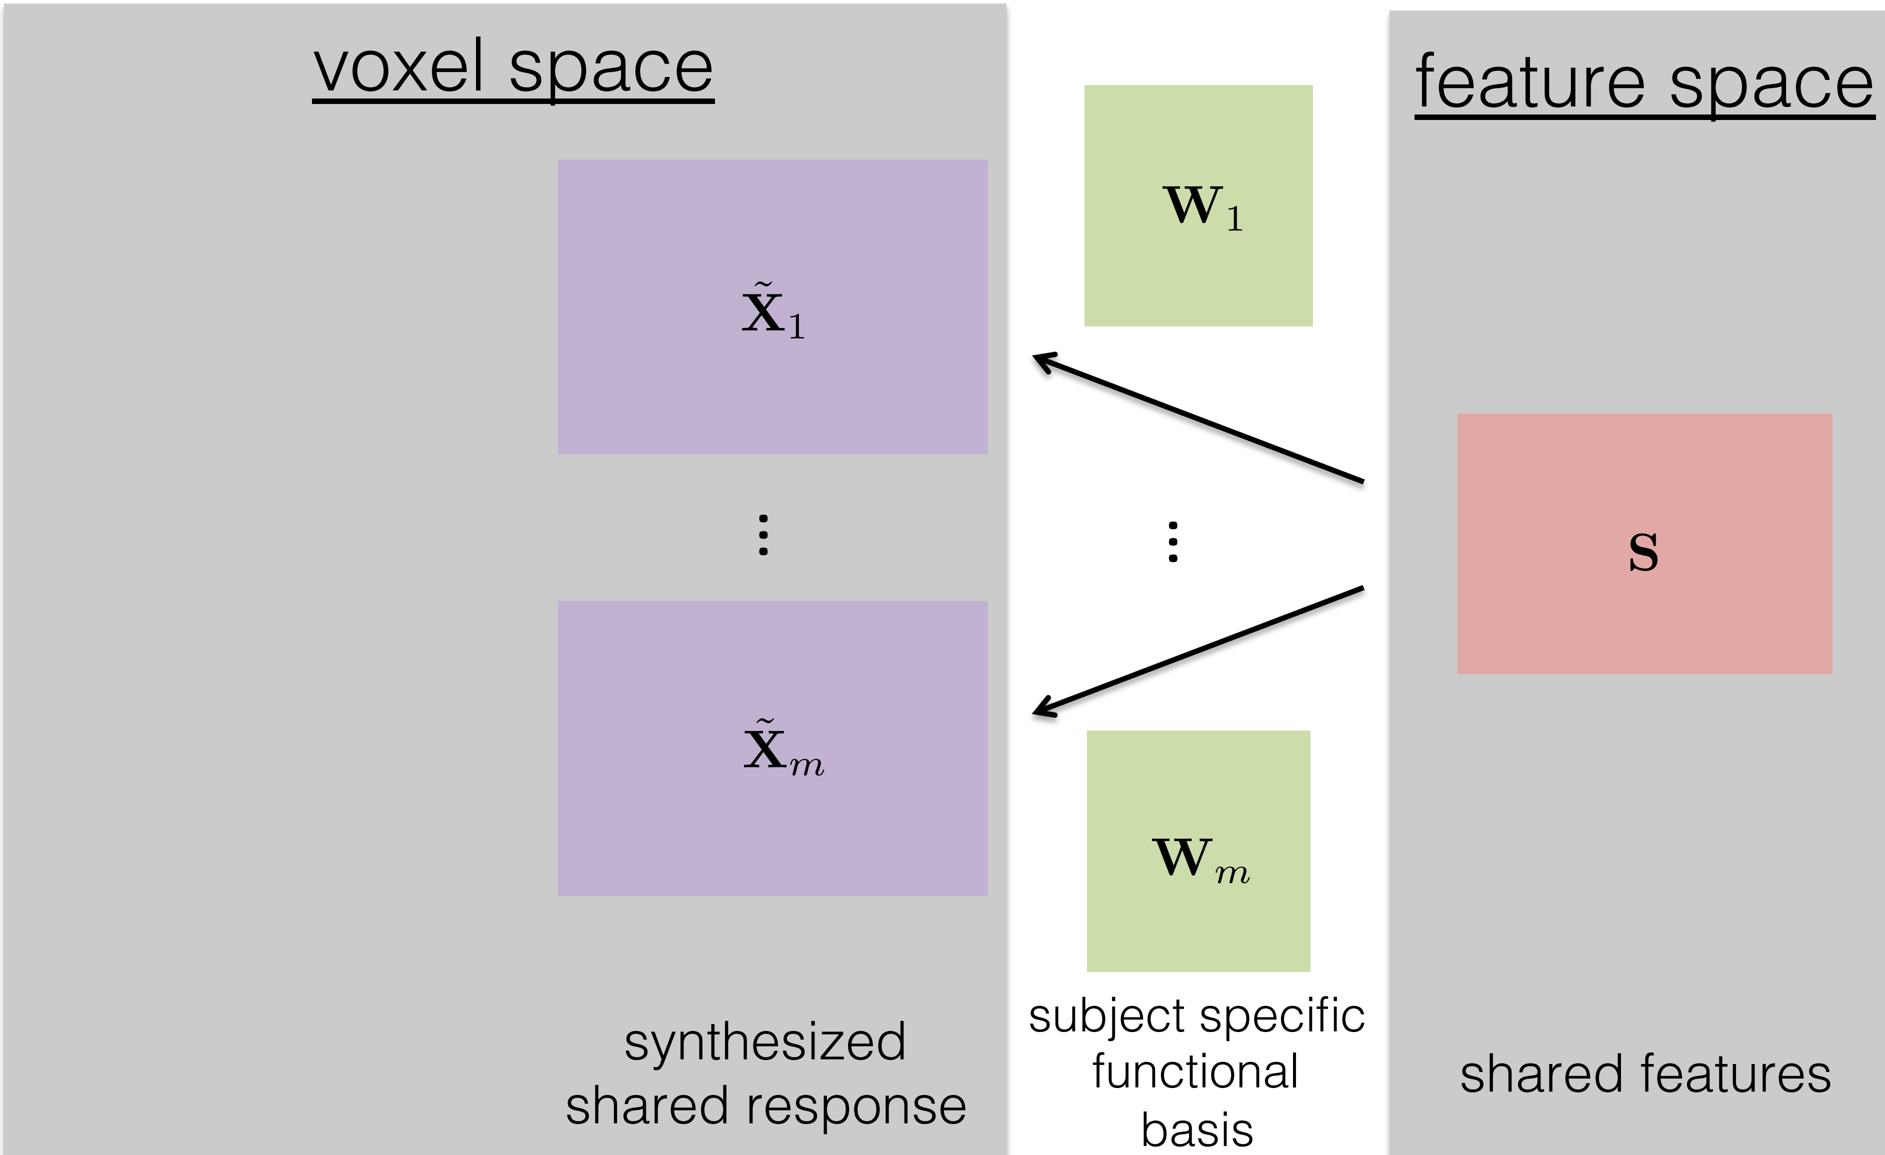

A generative model

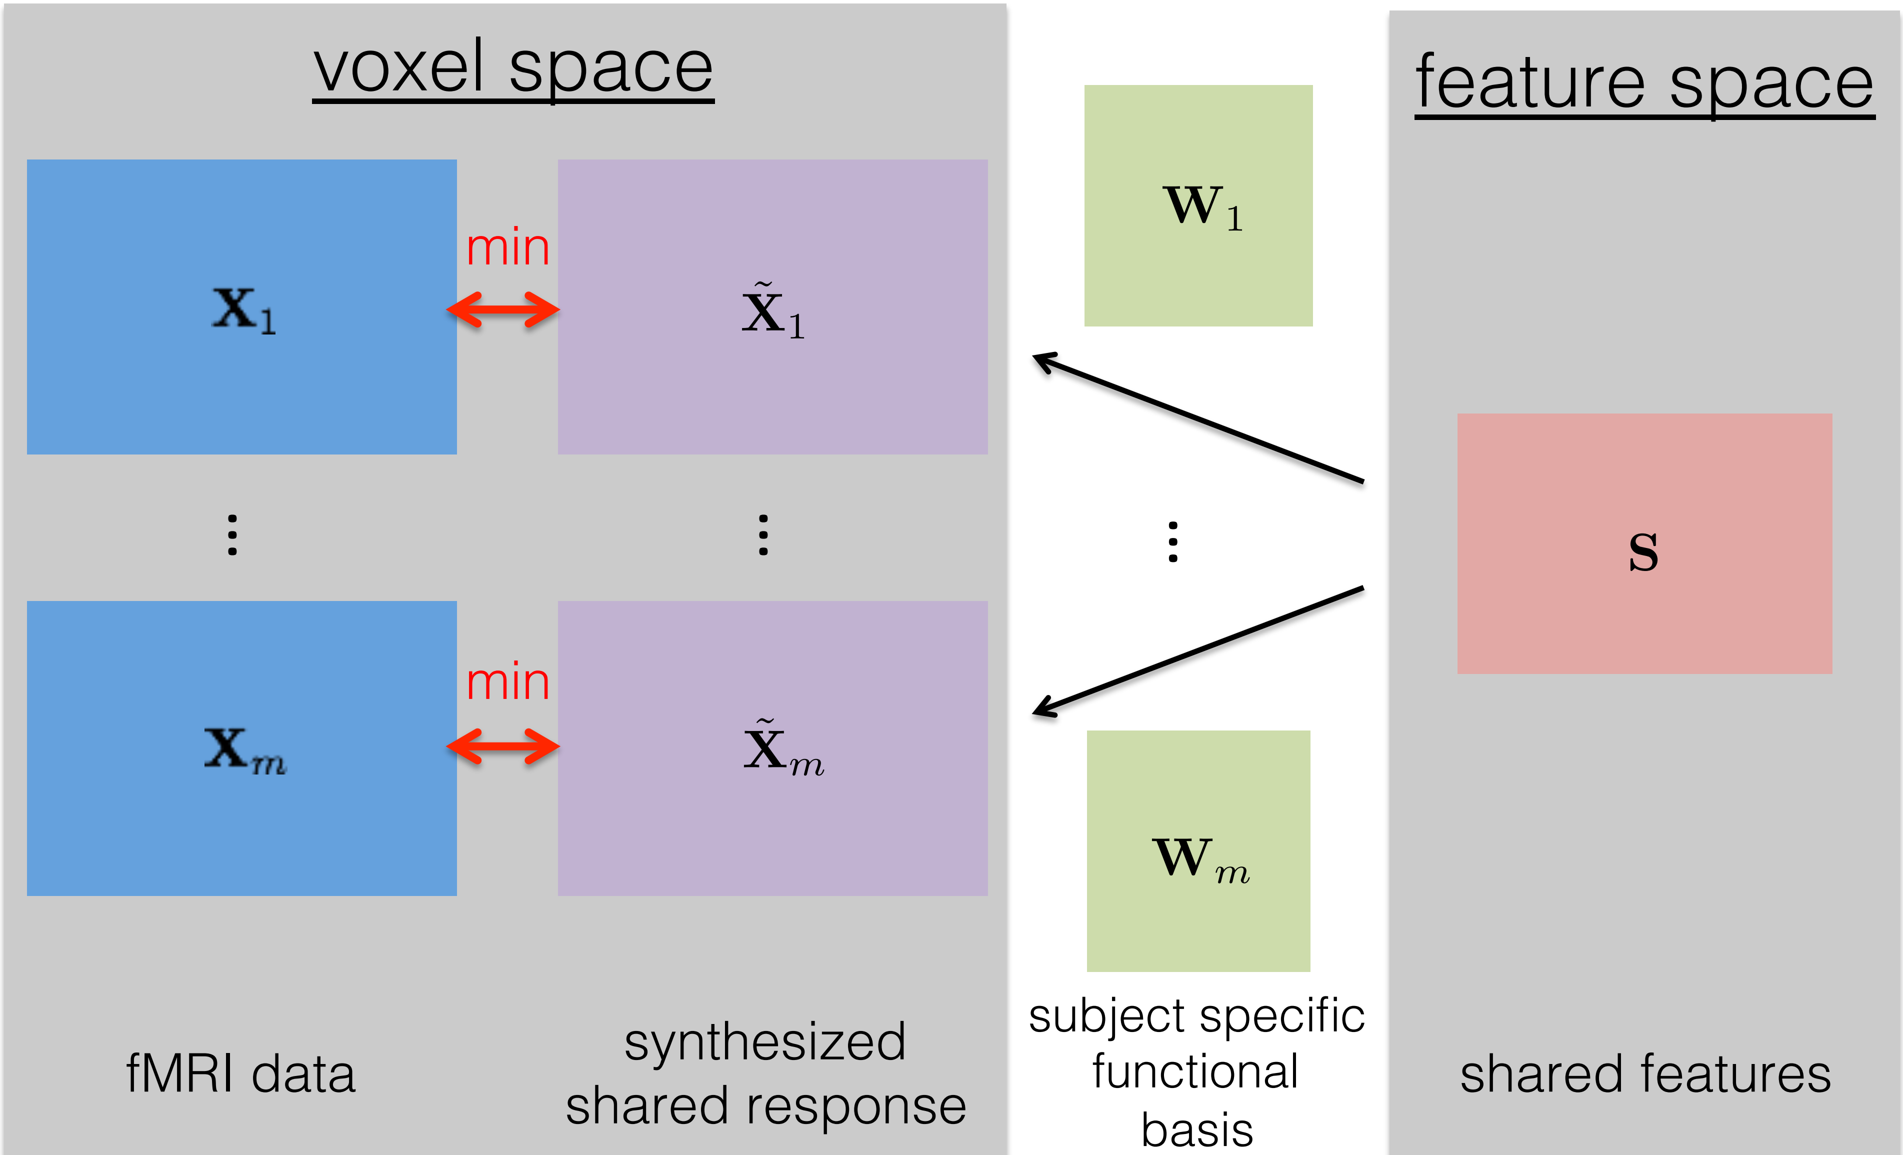

Shared Response Model (SRM) is a latent variable model

$$s_t \sim \mathcal{N}(0, \Sigma_s)$$
$$x_{it} | s_t \sim \mathcal{N}(W_i s_t + \mu_i, \rho_i^2 I)$$
$$W_i^T W_i = I$$

W_i not square

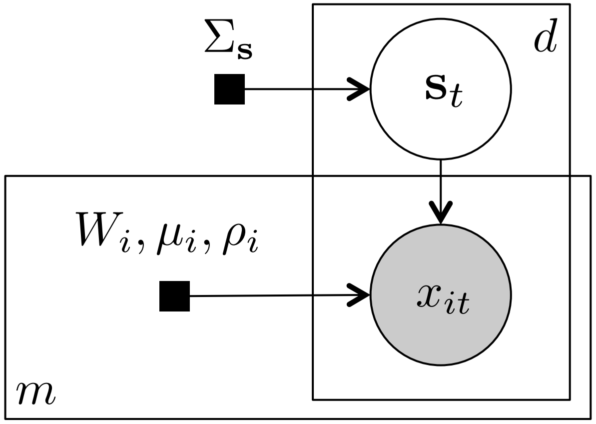

s_t shared elicited response at time t

W_i functional basis for subject i

x_{it} observations of subject i at time t

ρ_i^2 noise level for subject i's data

- Feature identification with dimensionality reduction
- Constrained EM algorithm

[P.-H. Chen et al., NIPS , 2015]

Constrained EM algorithm

E-step :

$$\begin{aligned}
\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t] &= (W\Sigma_s)^T(W\Sigma_s W^T + \Psi)^{-1}(x_t - \mu), \\
\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t\mathbf{s}_t^T] &= \mathbf{Var}_{\mathbf{s}|x}[\mathbf{s}_t] + \mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t]\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t]^T \\
&= \Sigma_s - \Sigma_s^T W^T(W\Sigma_s W^T + \Psi)^{-1} W\Sigma_s + \mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t]\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t]^T
\end{aligned}$$

M-step :

$$\begin{aligned}
&\mu_i^{\text{new}} = \tfrac{1}{d}\textstyle\sum_t x_{it}, \\
&W_i^{\text{new}} = A_i(A_i^T A_i)^{-1/2}, \quad A_i = \tfrac{1}{2}\big(\textstyle\sum_t (x_{it} - \mu_i^{\text{new}})\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t]^T\big), \\
&\rho_i^{2\,\text{new}} = \tfrac{1}{dv}\textstyle\sum_t \big(\|x_{it} - \mu_i^{\text{new}}\|^2 - 2(x_{it} - \mu_i^{\text{new}})^T W_i^{\text{new}}\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t] + \mathbf{tr}(\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t\mathbf{s}_t^T])\big), \\
&\Sigma_s^{\text{new}} = \tfrac{1}{d}\textstyle\sum_t (\mathbb{E}_{\mathbf{s}|x}[\mathbf{s}_t\mathbf{s}_t^T]).
\end{aligned}$$

- Learning W on Stiefel manifold

Shared features, subject specific functional basis, and subject specific response

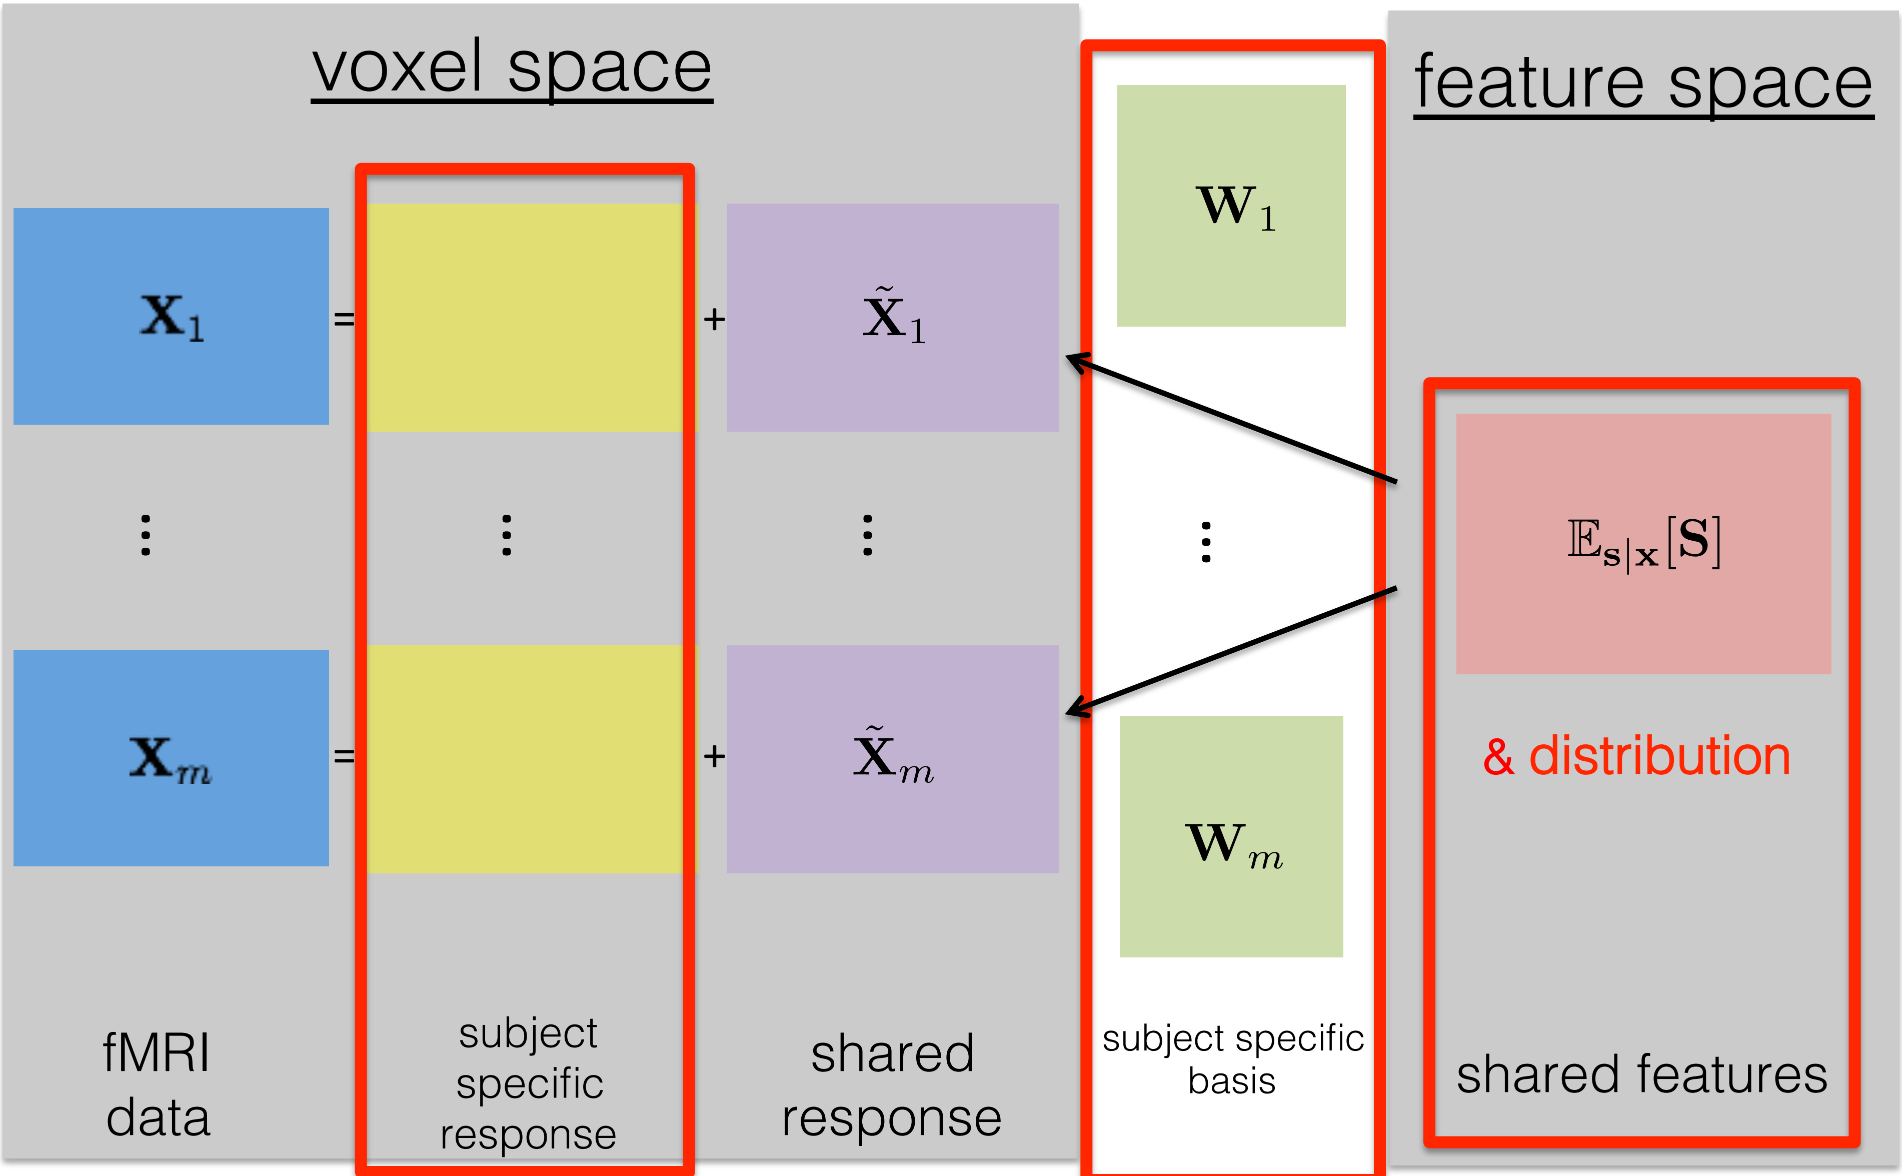

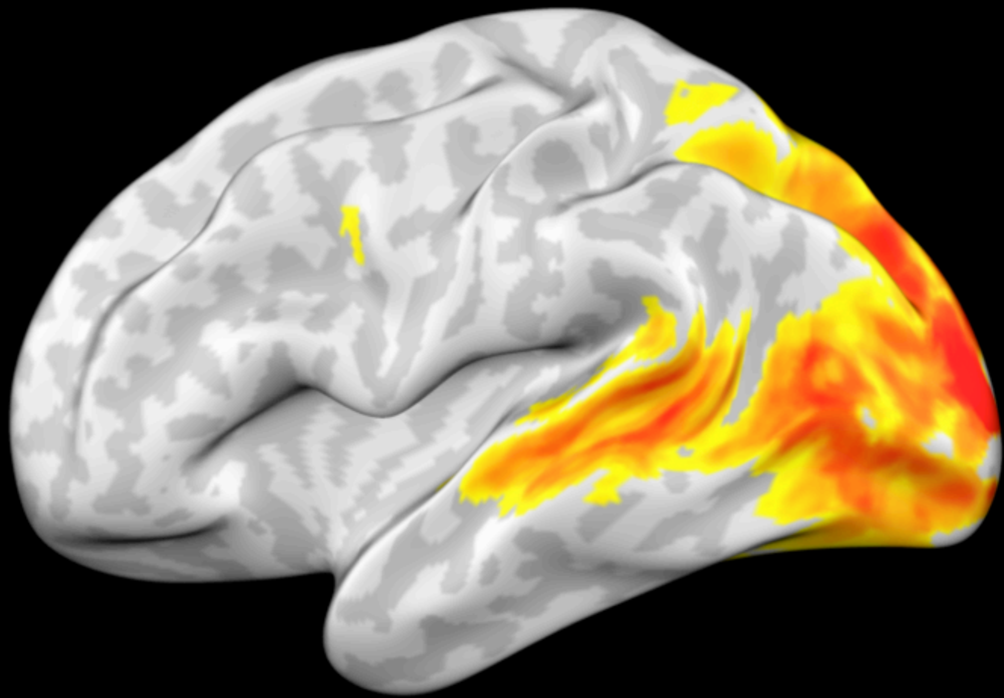

Part II: Shared Response Model on Neuroimaging Data

Evaluation with various datasets

- Different MRI machines
- Different institutes
- Different subjects
- Different preprocessing protocols
- Different brain regions
- Different data size

sherlock

raider

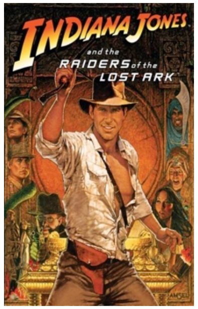

forrest

audiobook

SRM on fMRI

1. Generalize to new subject

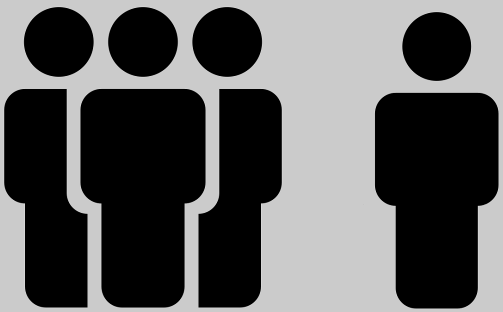

2. Generalize to new stimulus

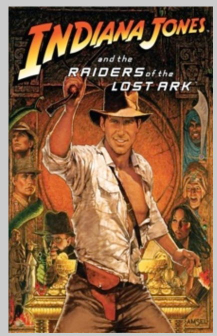

3. Decoupling shared and individual response

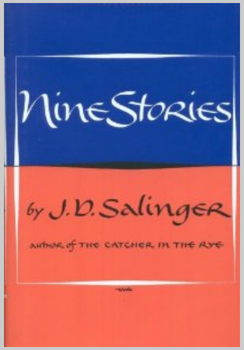

4. SRM with retinotopy

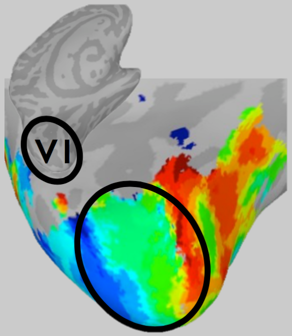

5. Searchlight SRM

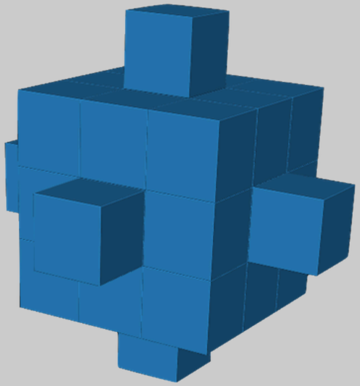

6. Bridging shared space and text semantic space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

1. Generalize to new subject
2. Generalize to new stimulus

1. Generalize to new subject

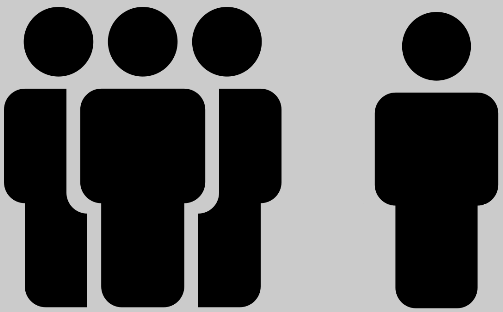

2. Generalize to new stimulus

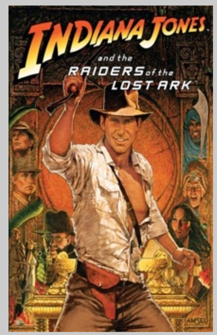

3. Decoupling shared and individual response

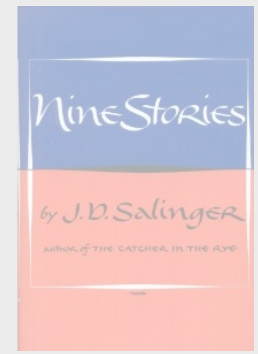

4. SRM with retinotopy

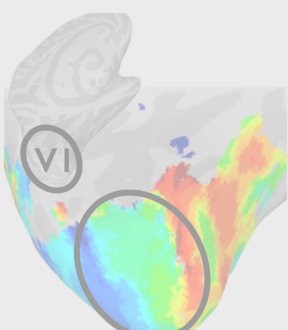

5. Searchlight SRM

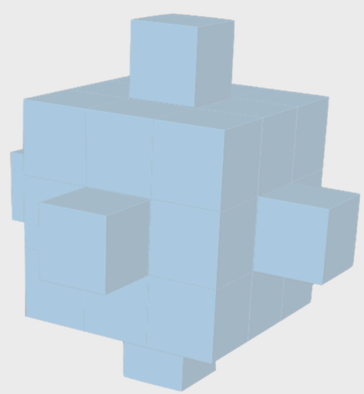

6. Bridging shared space and word embedding space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

Generalization to new subject with time segment matching

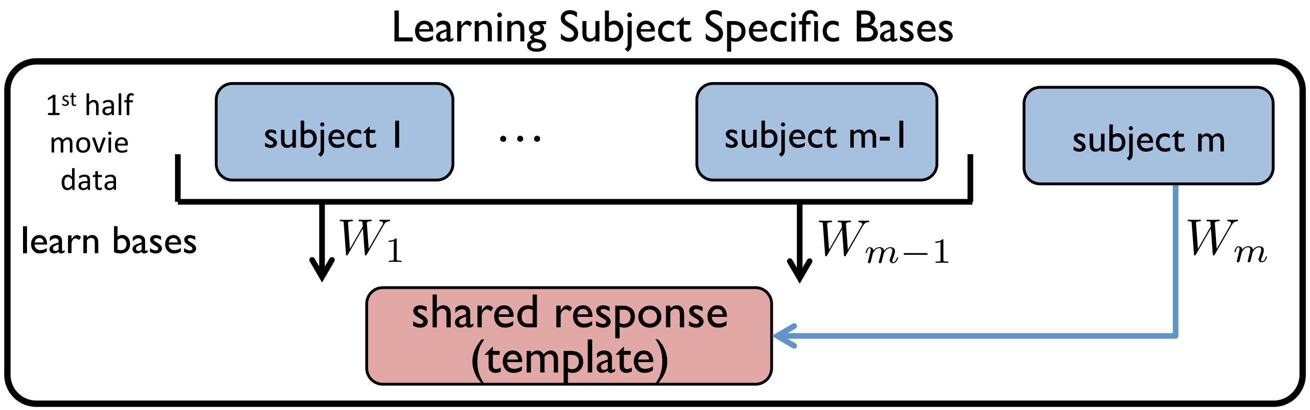

Datasets

sherlock

forrest

raider

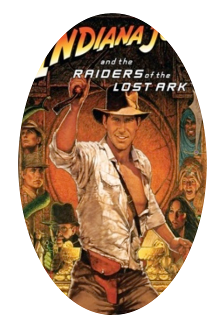

Generalization to new subject with time segment matching

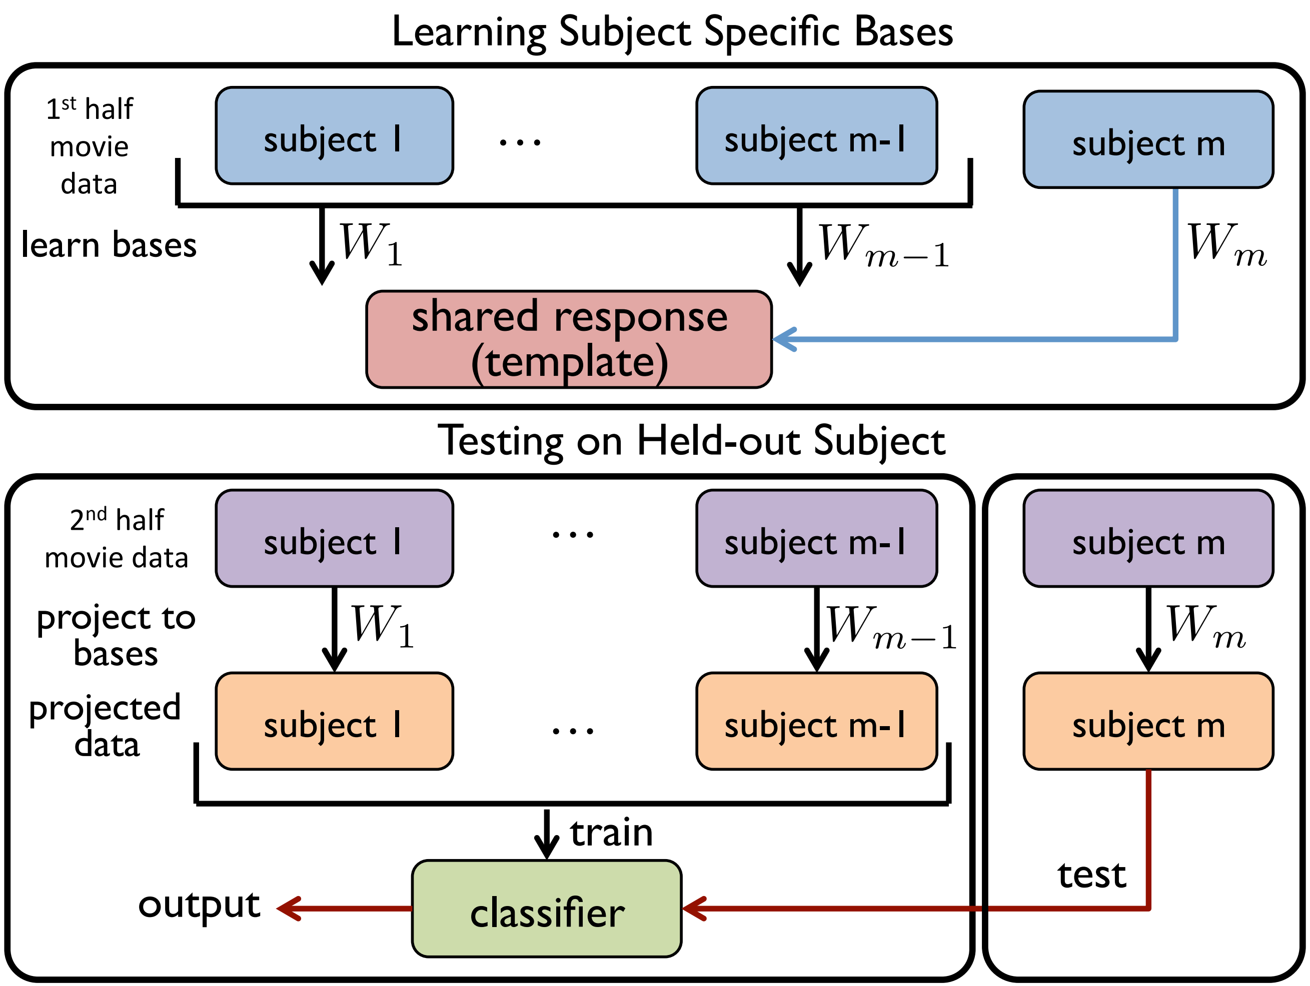

Datasets

sherlock

forrest

raider

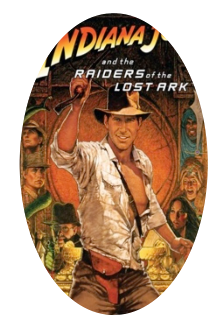

Generalization to new subject with time segment matching

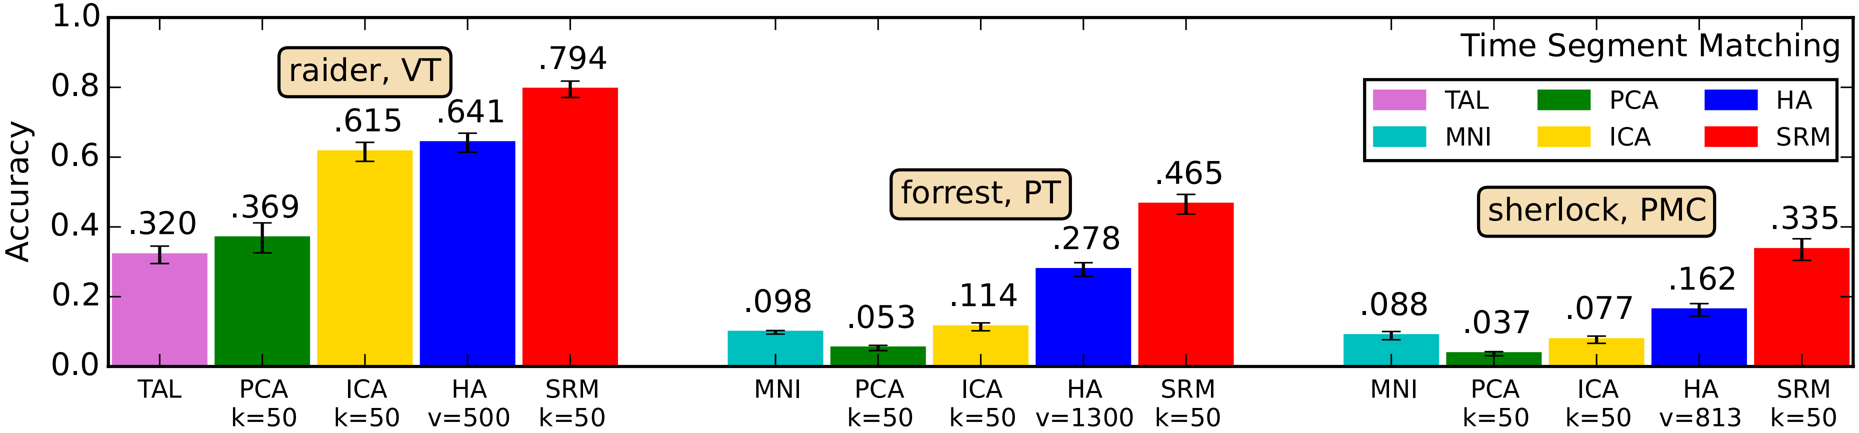

[P.-H. Chen et al., NIPS , 2015]

Generalization to new subject and distinct stimulus with image classification

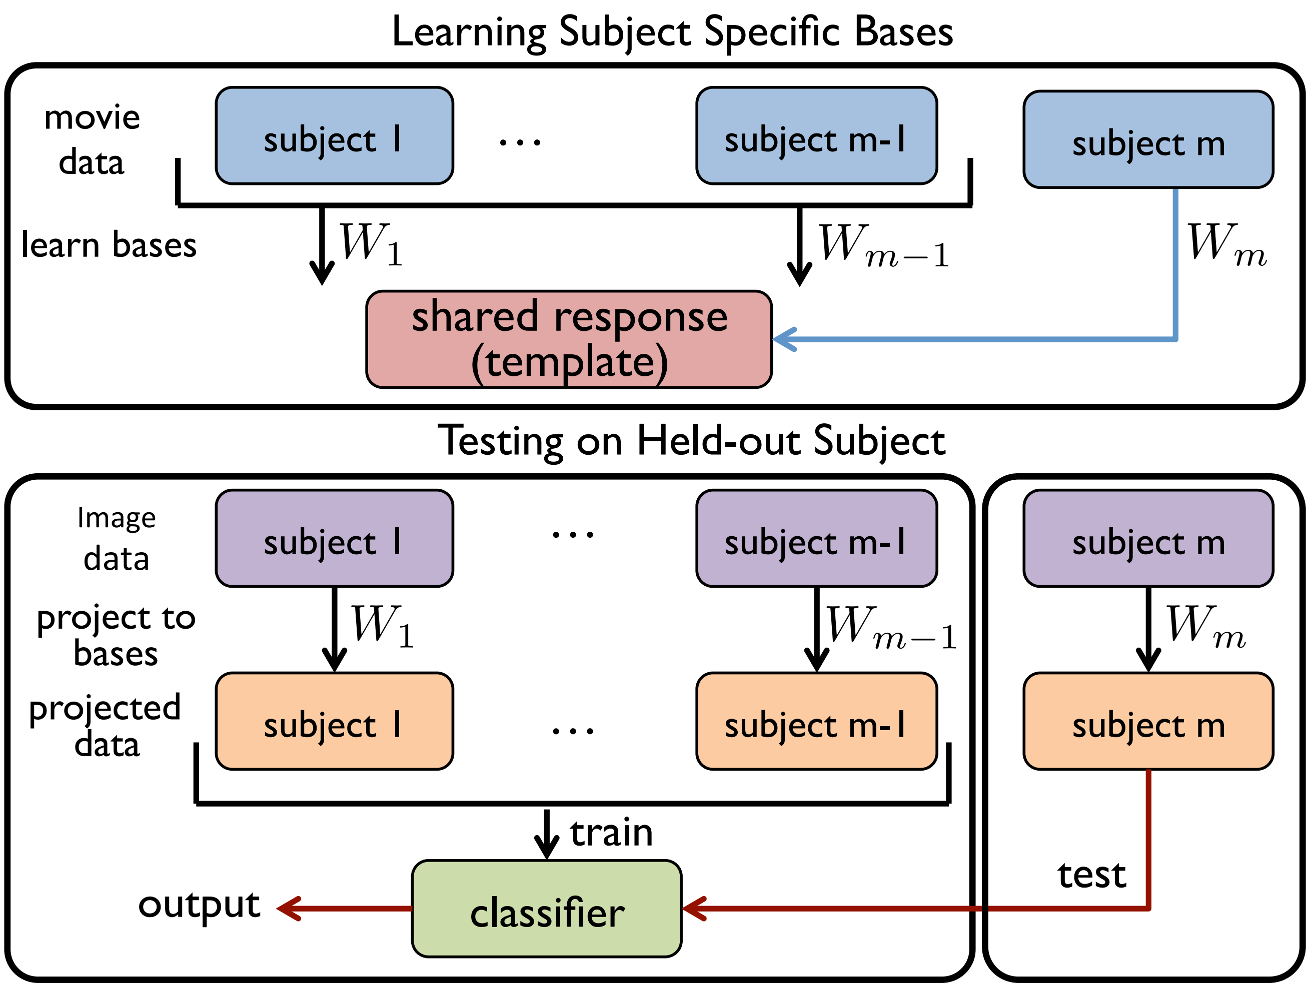

Dataset

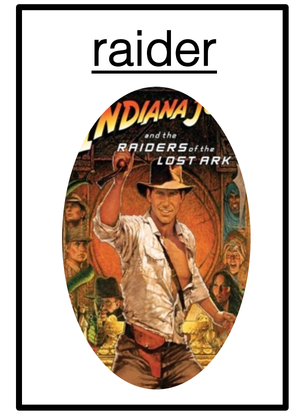

Generalization to new subject and distinct stimulus with image classification

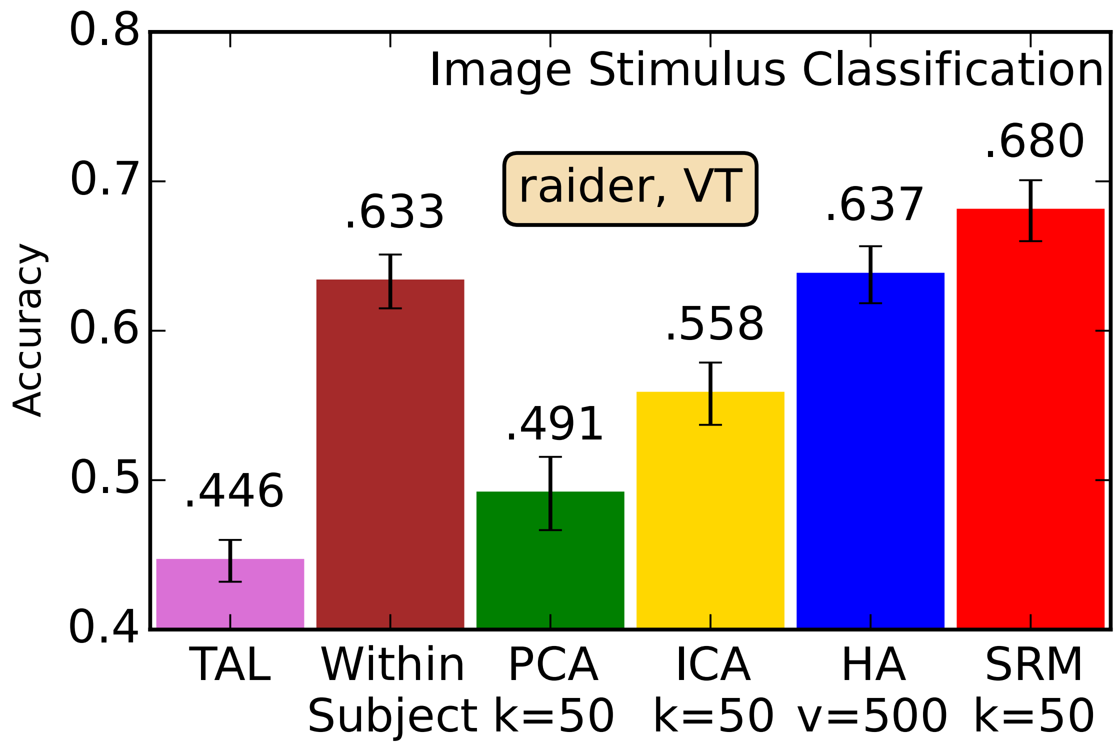

- Outperforms within-subject classification

[P.-H. Chen et al., NIPS , 2015]

3. Decoupling shared and individual response

1. Generalize to new subject

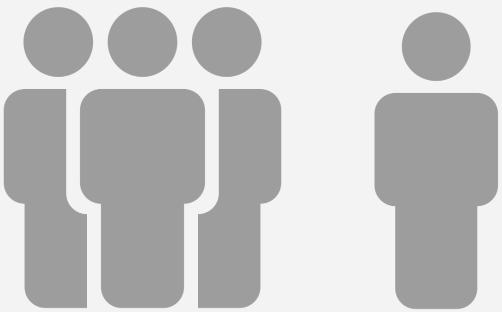

2. Generalize to new stimulus

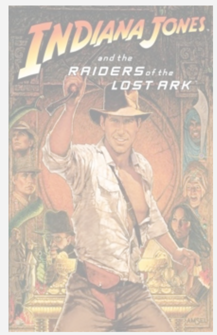

3. Decoupling shared and individual response

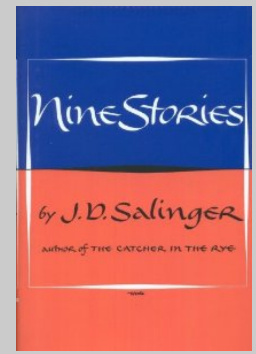

4. SRM with retinotopy

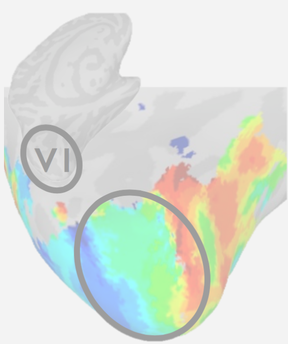

5. Searchlight SRM

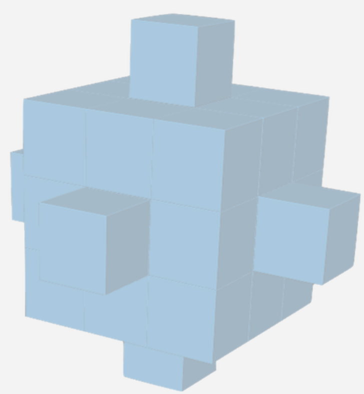

6. Bridging shared space and word embedding space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

Classifying mental states

- 40 subjects listening to narrated story
- Separate 40 subjects into 2 groups
- Two groups receive different prior contexts
- Leading to different interpretations of the story
- Predict prior context of a left-out subject

Dataset

audiobook

Classifying mental states

Group 1

shared by all | within group | individual

Group 2

shared by all | within group | individual

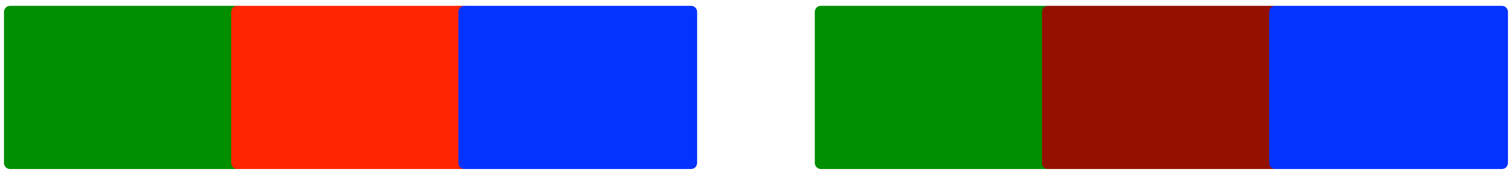

Group prediction accuracy with SRM

0.72±0.06

[P.-H. Chen et al., NIPS , 2015]

Classifying mental states

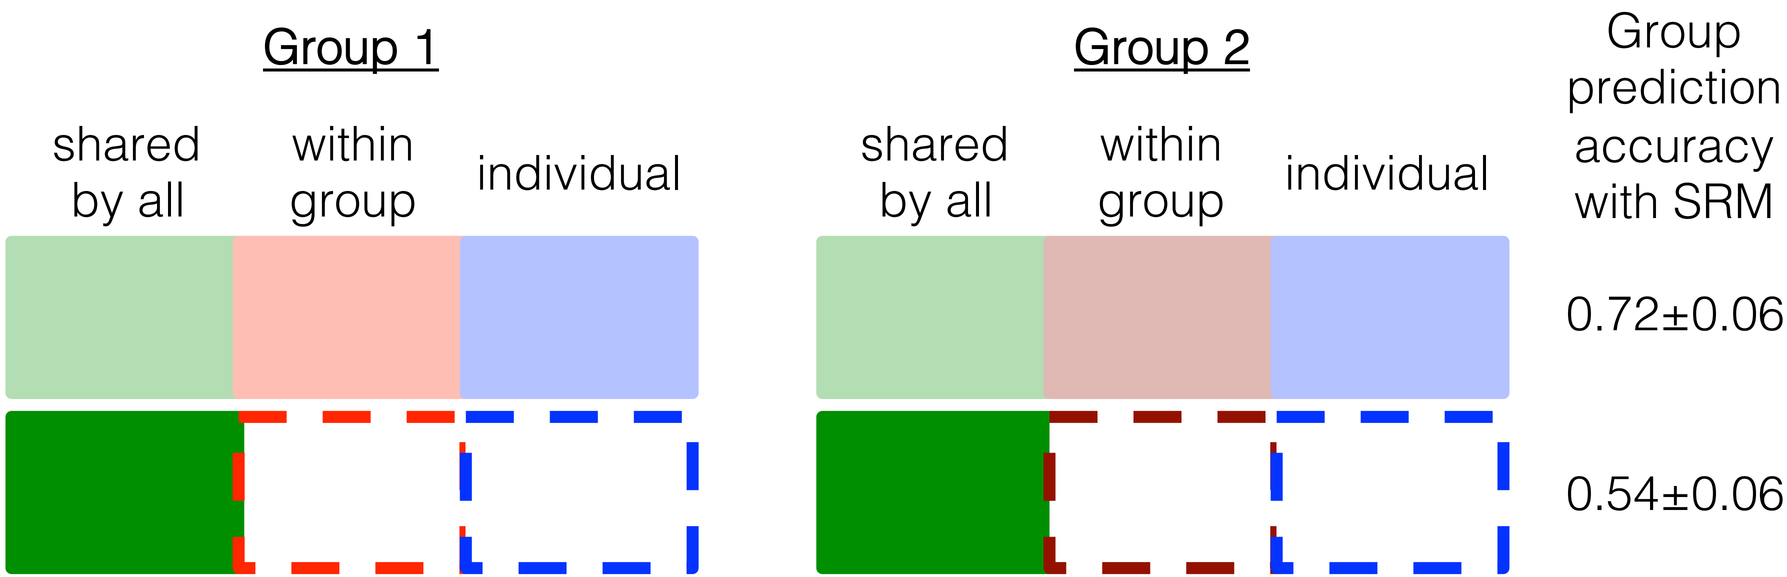

[P.-H. Chen et al., NIPS , 2015]

Classifying mental states

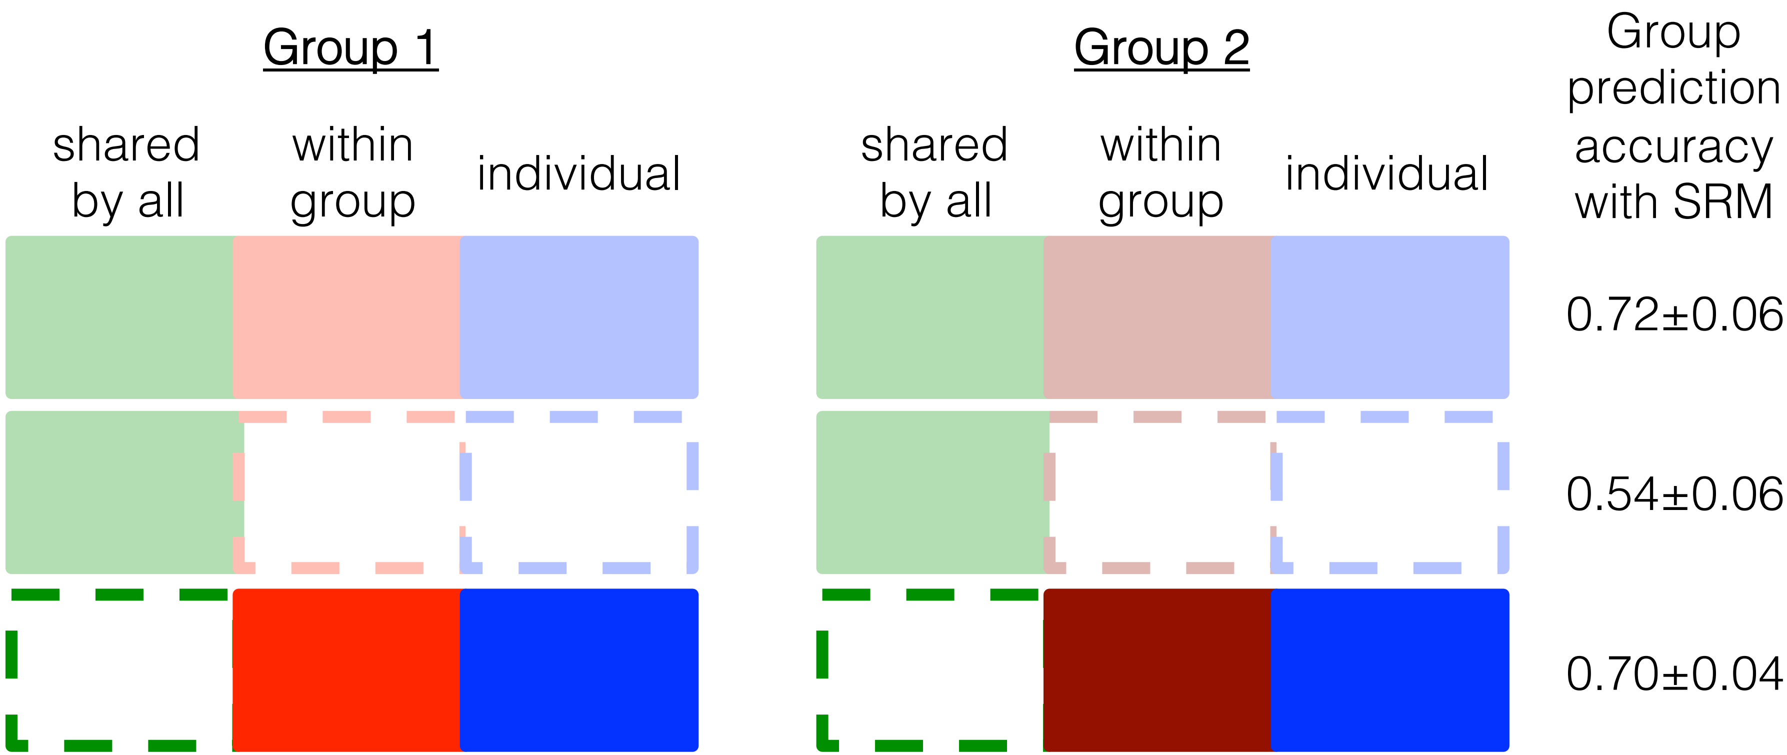

[P.-H. Chen et al., NIPS , 2015]

Classifying mental states

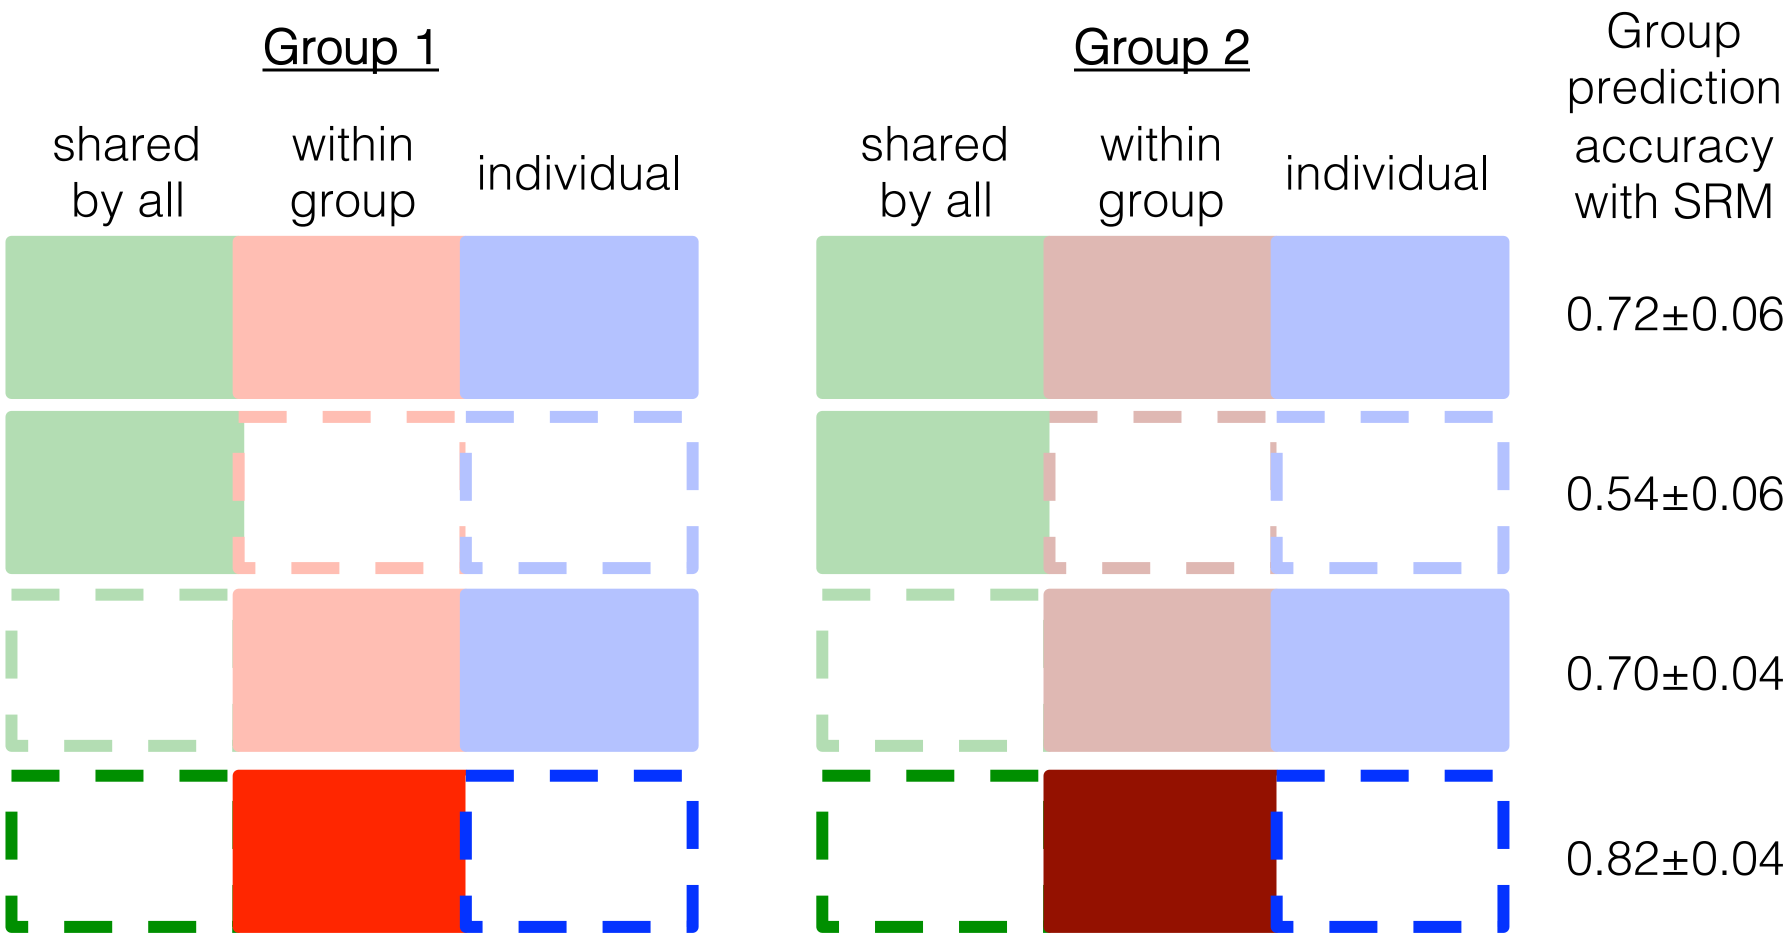

[P.-H. Chen et al., NIPS , 2015]

4. SRM with retinotopy

1. Generalize to new subject

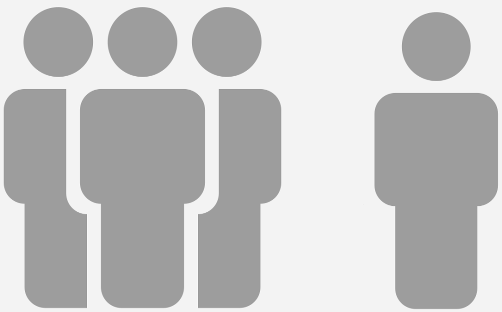

2. Generalize to new stimulus

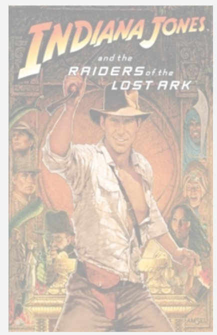

3. Decoupling shared and individual response

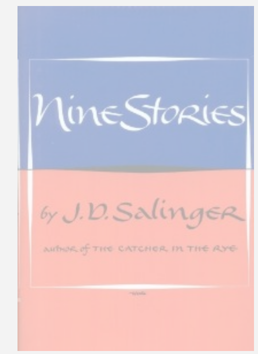

4. SRM with retinotopy

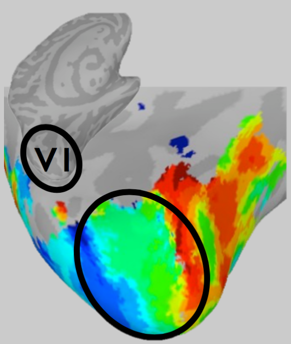

5. Searchlight SRM

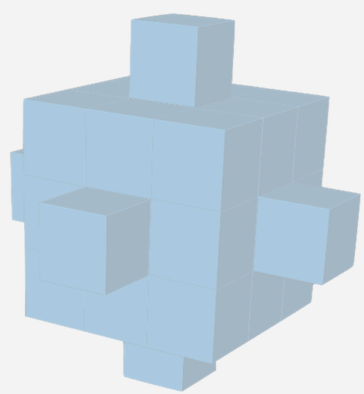

6. Bridging shared space and word embedding space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

Mapping Visual Field Maps: Retinotopy

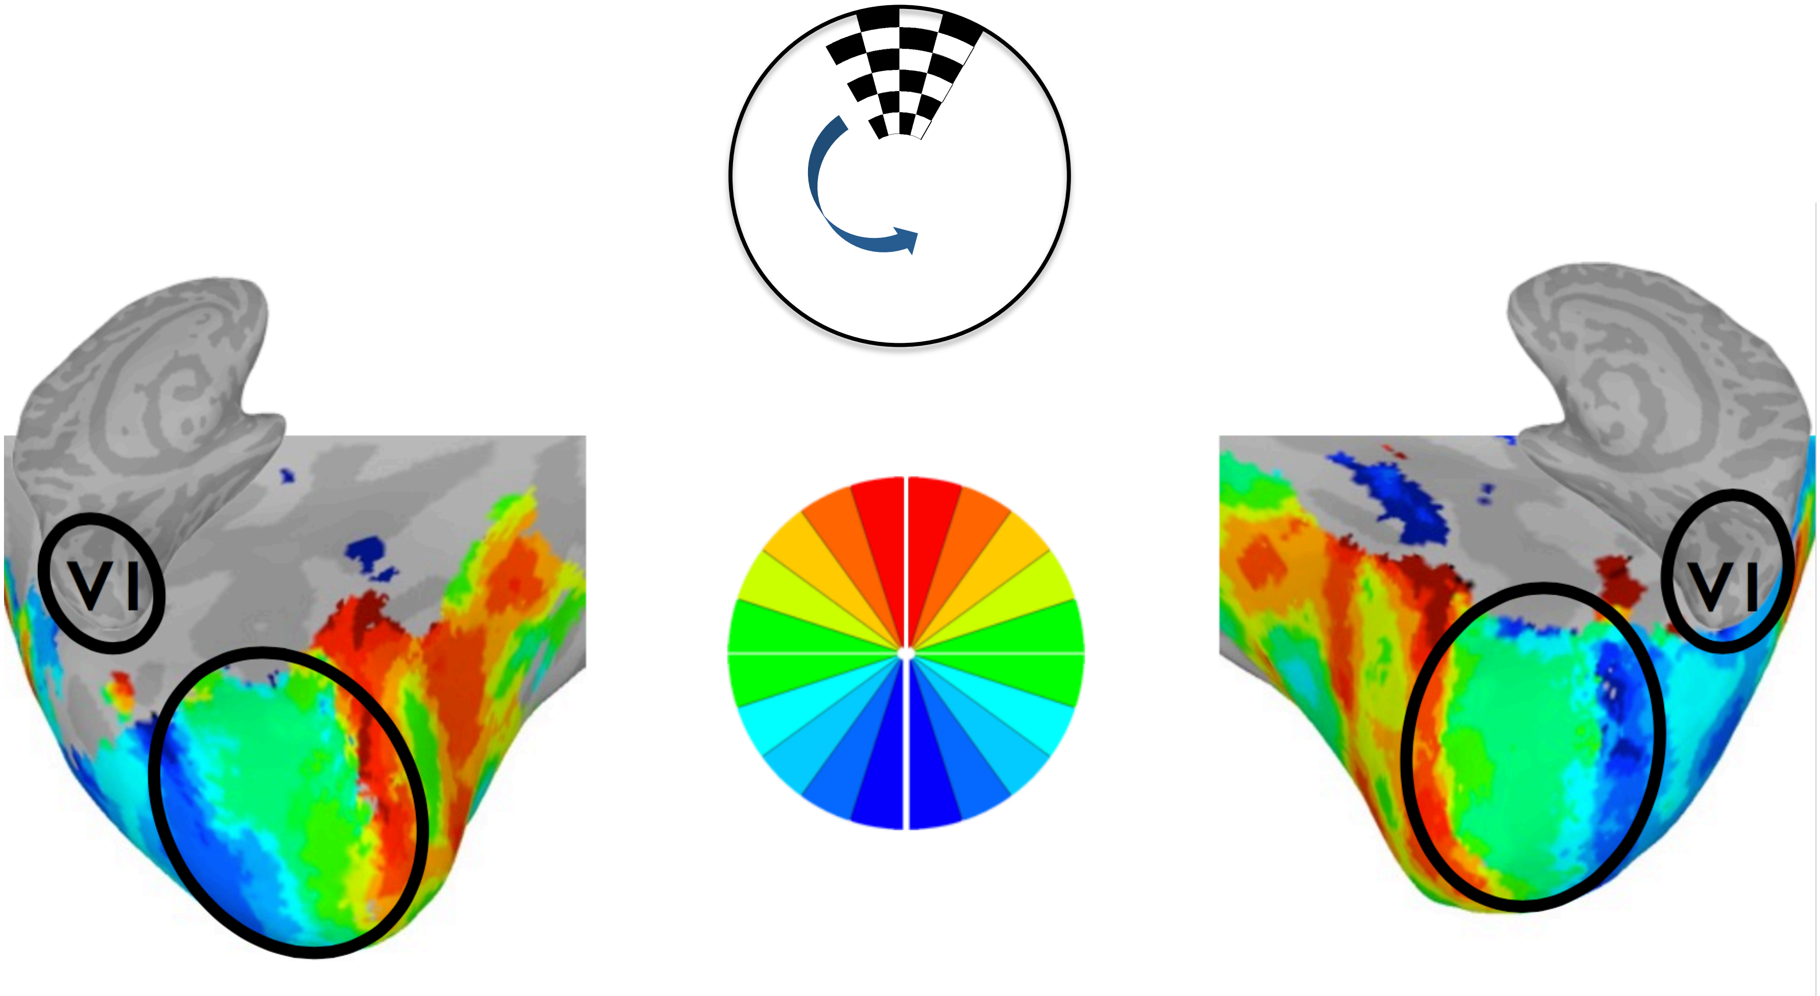

[Work by Michael J. Arcaro]

Original Phase Maps vs. SRM

Sanity check: (W_i*transformed_data_i)

Phase map comparison between original phase maps and phase maps derived from data reconstructed in same subject post hyperalign. NOTE: original data was not masked and includes more of cortex. Data threshold a $p < .0001$

Orig SRM

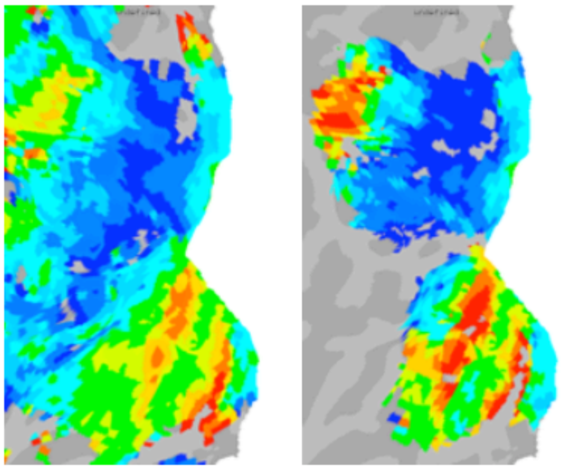

Orig SRM

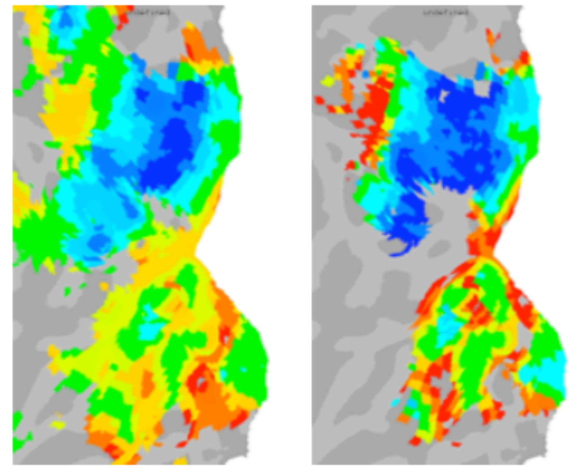

Orig SRM

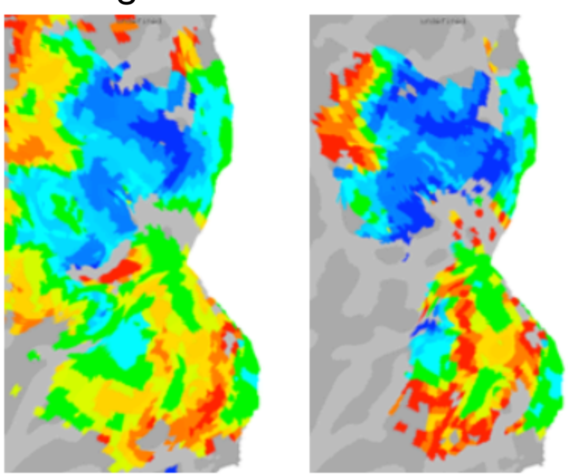

Orig SRM

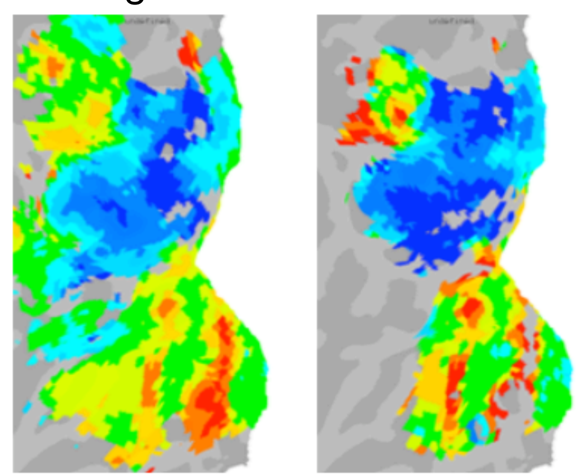

[Work by Michael J. Arcaro]

Transformation between subjects

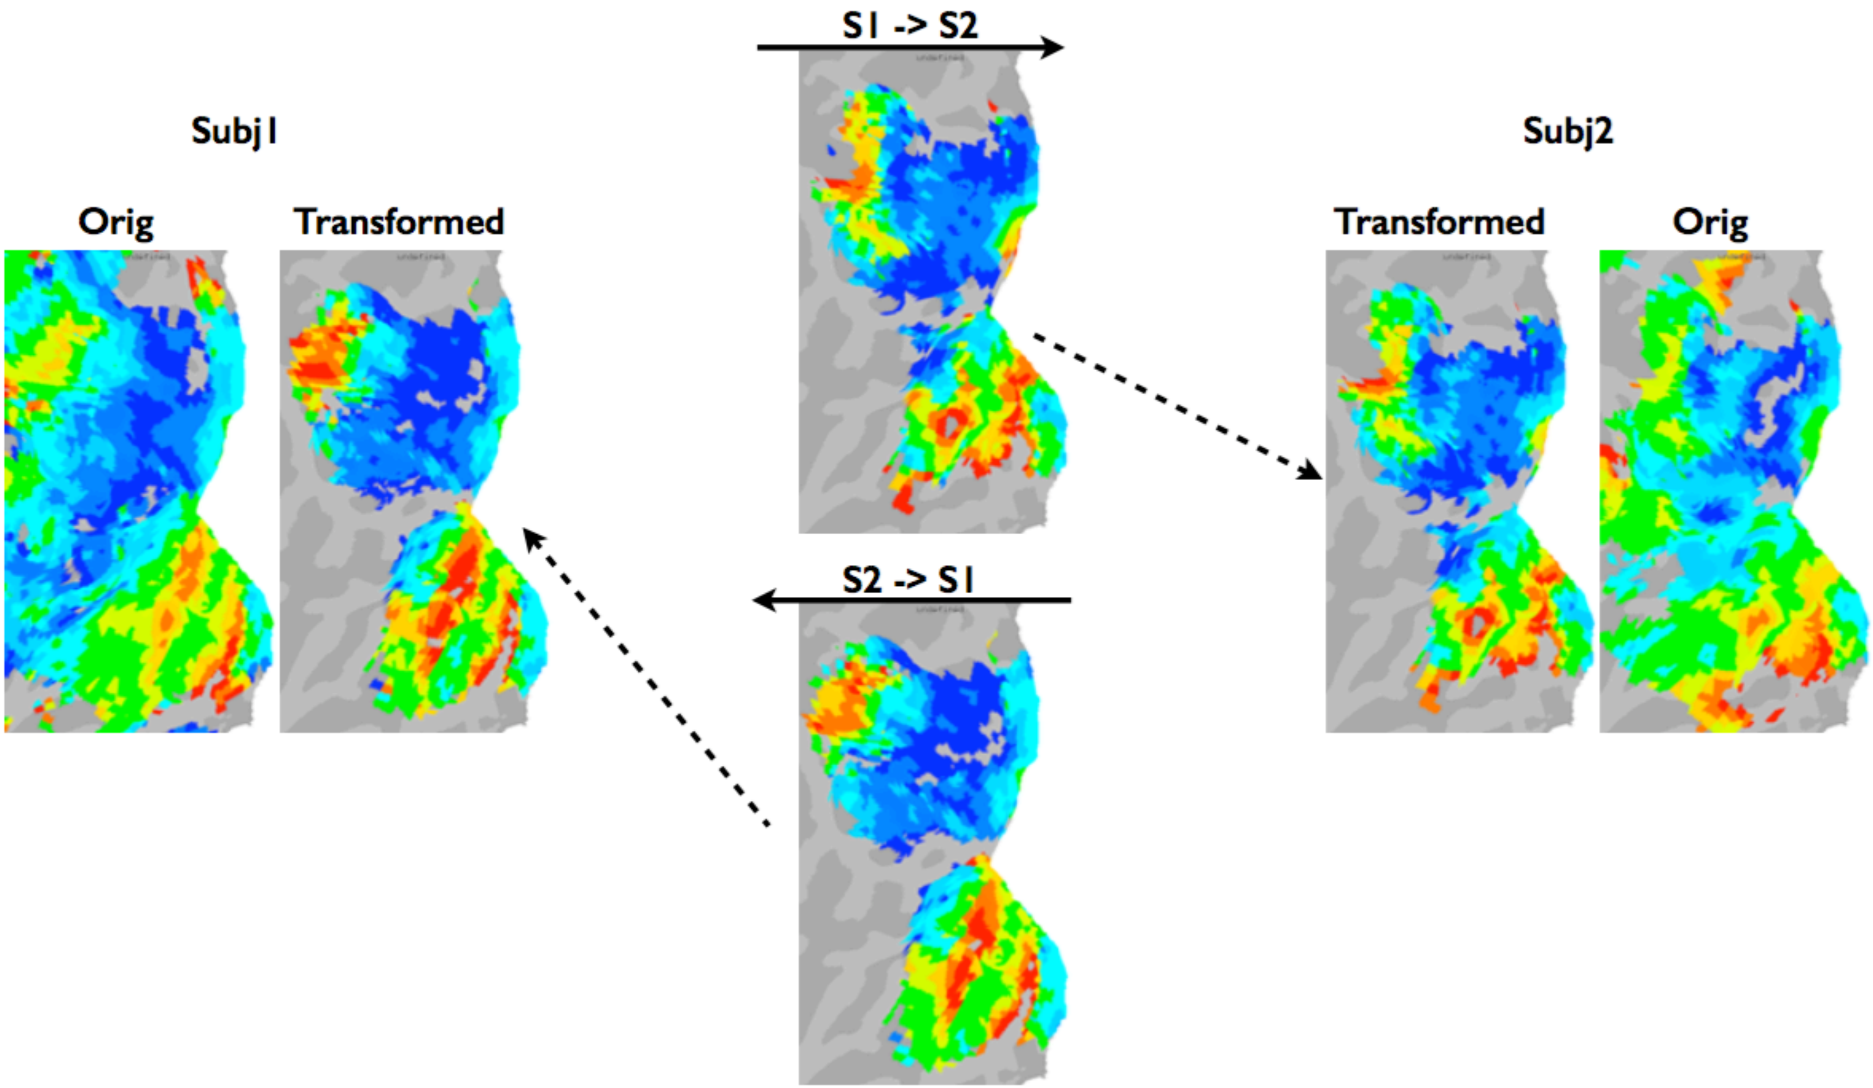

[Work by Michael J. Arcaro]

5. Searchlight SRM

1. Generalize to new subject

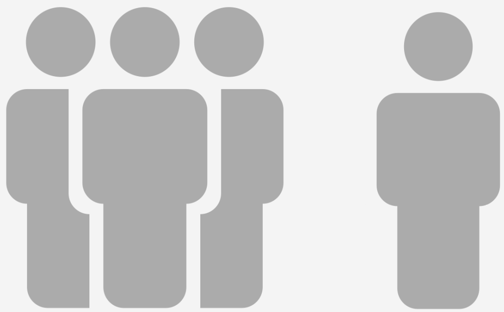

2. Generalize to new stimulus

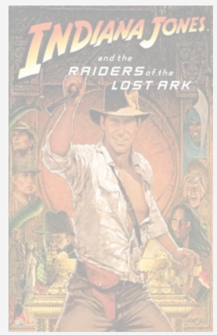

3. Decoupling shared and individual response

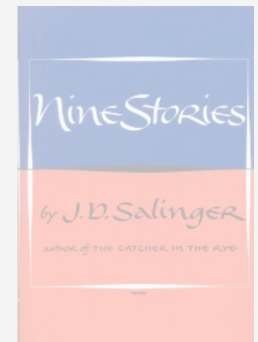

4. SRM with retinotopy

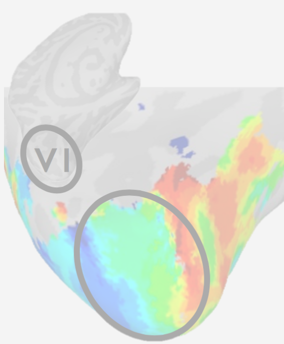

5. Searchlight SRM

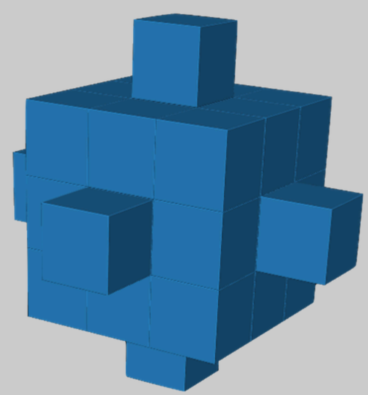

6. Bridging shared space and word embedding space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

Searchlight SRM

- localized analysis across the whole brain

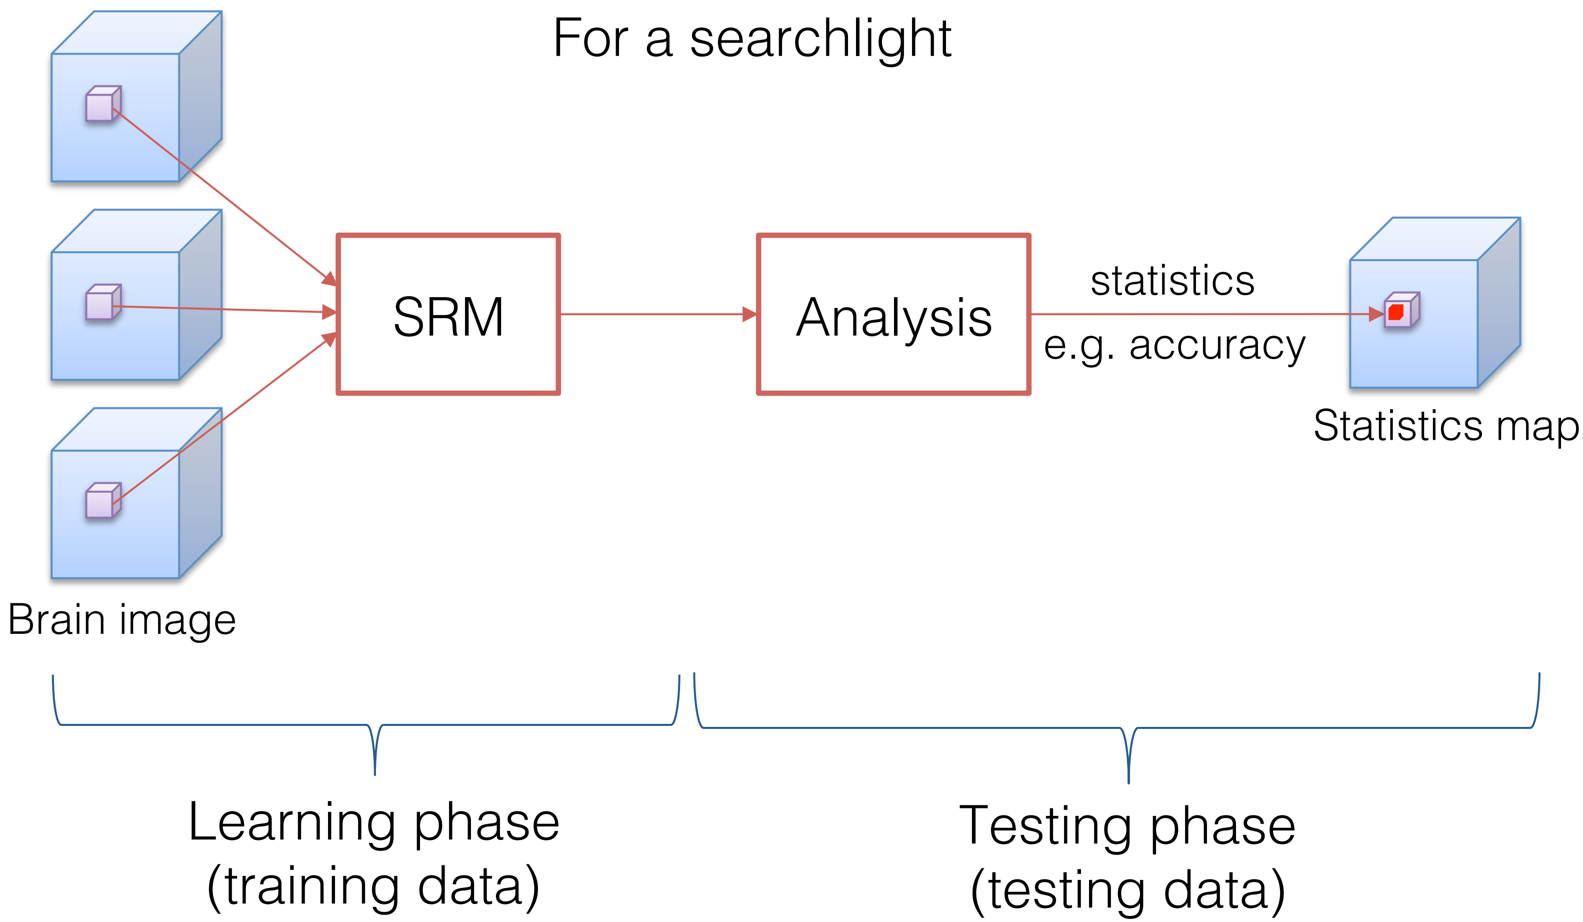

[Zhang and Chen et al., ArXiv , 2016]

Time segment matching with searchlight SRM

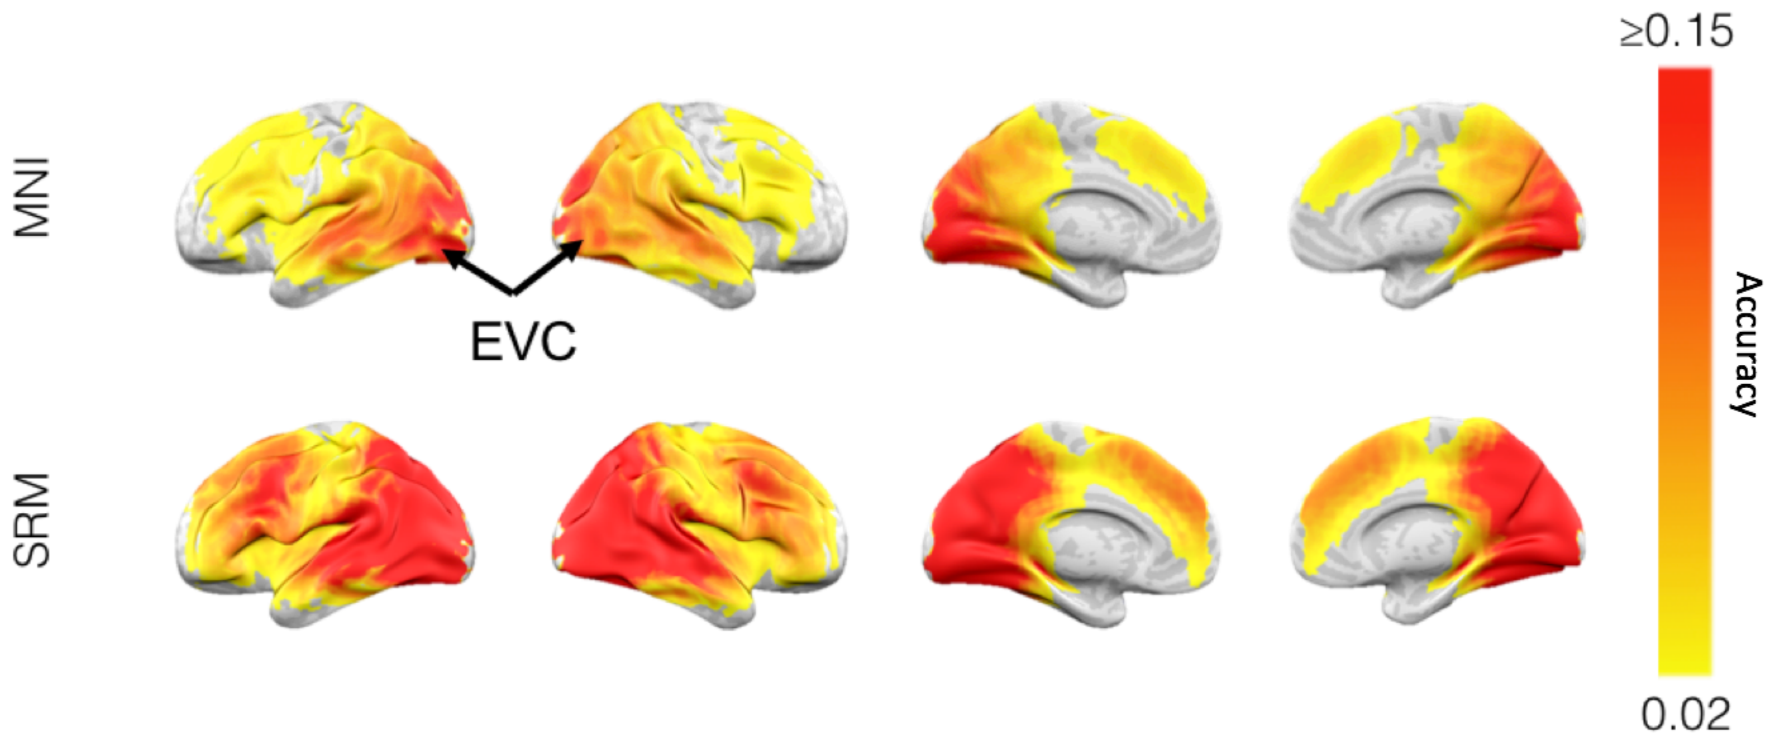

Accuracy map from time segment matching experiment (Sherlock)

[Zhang and Chen et al., ArXiv , 2016]

6. Bridging fMRI shared space and text semantic space

1. Generalize to new subject

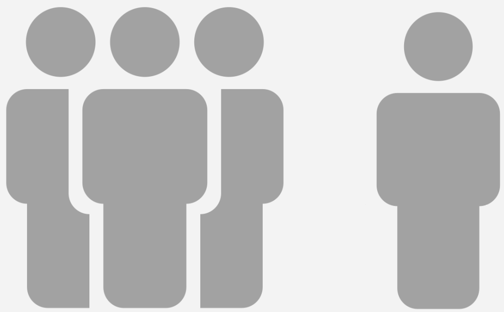

2. Generalize to new stimulus

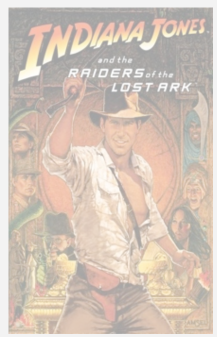

3. Decoupling shared and individual response

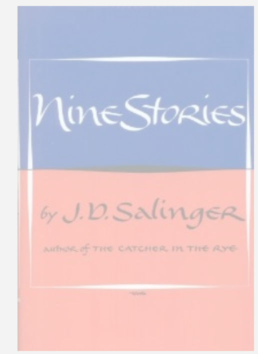

4. SRM with retinotopy

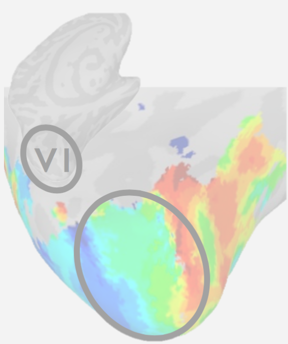

5. Searchlight SRM

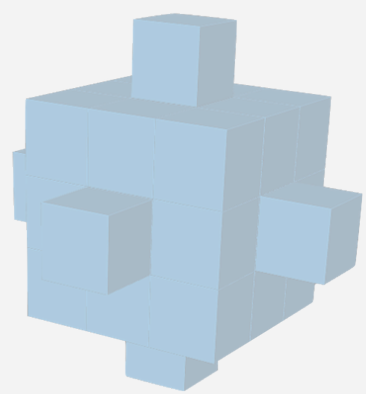

6. Bridging fMRI shared space and semantic space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

Bridging fMRI shared space and text semantic space

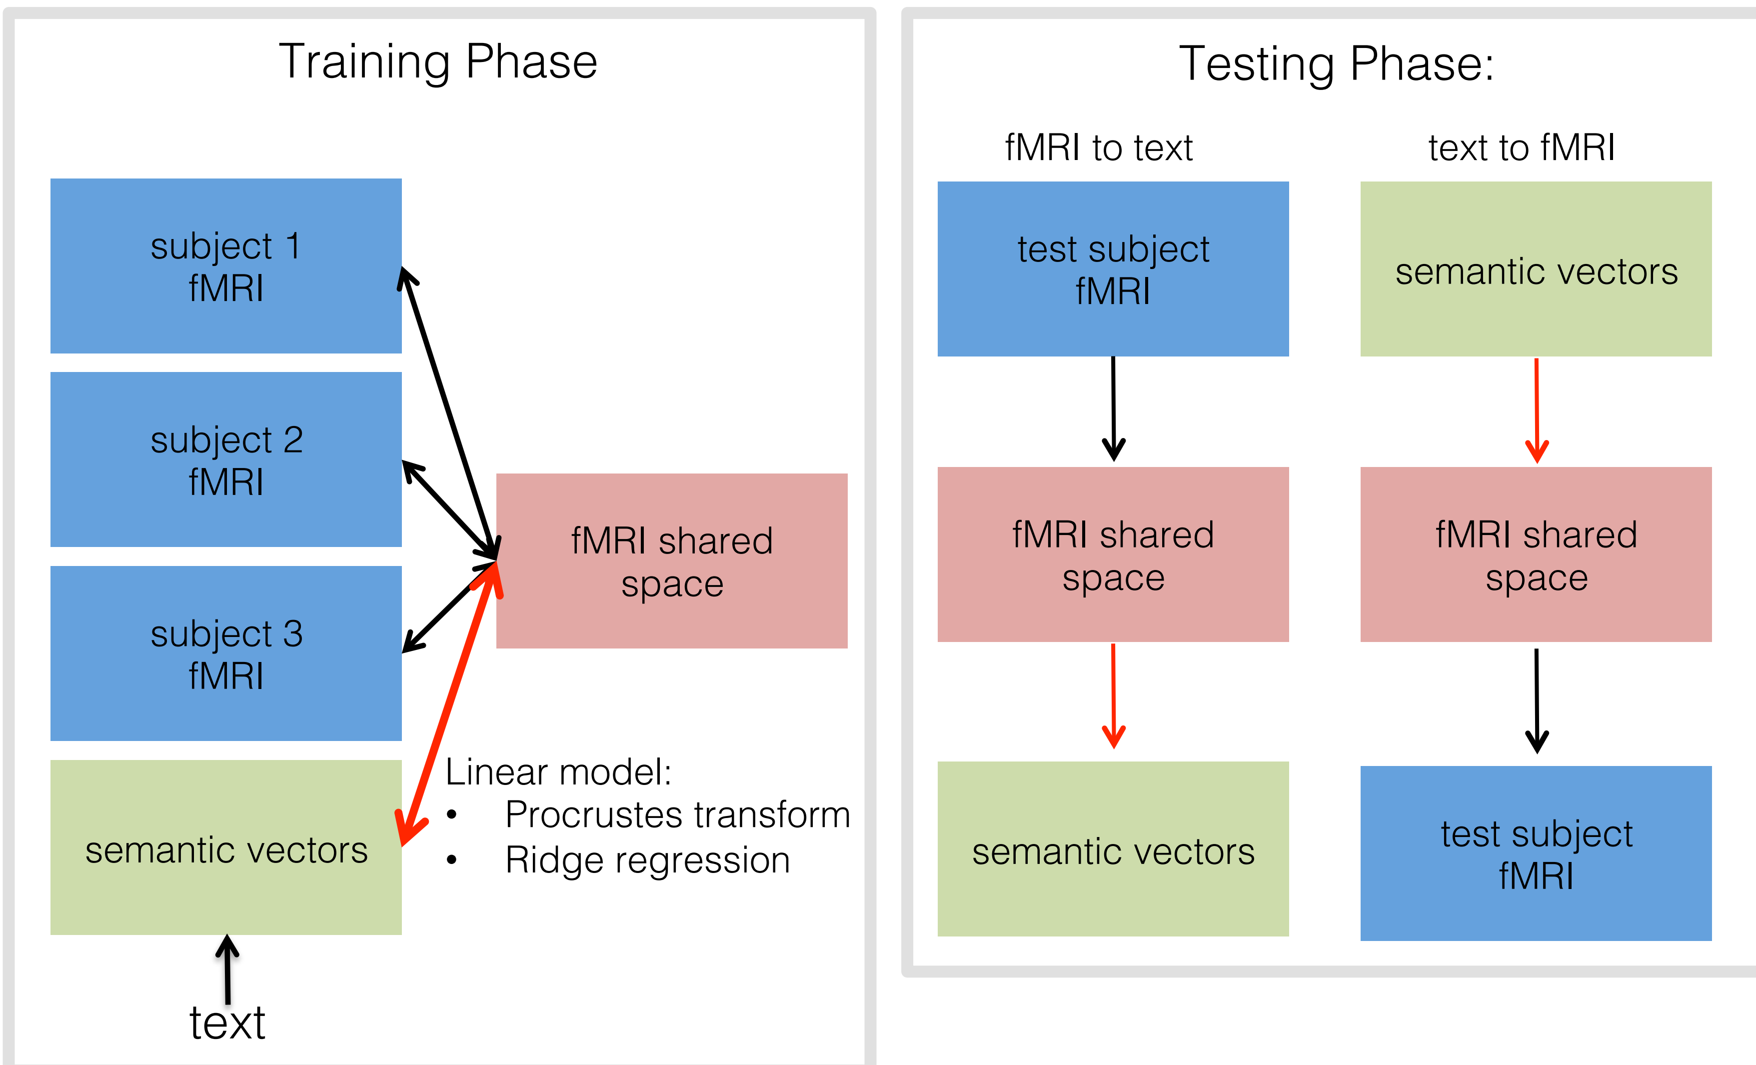

[Vodrahalli and Chen et al. 2016 arXiv]

Bridging shared space and word embedding space

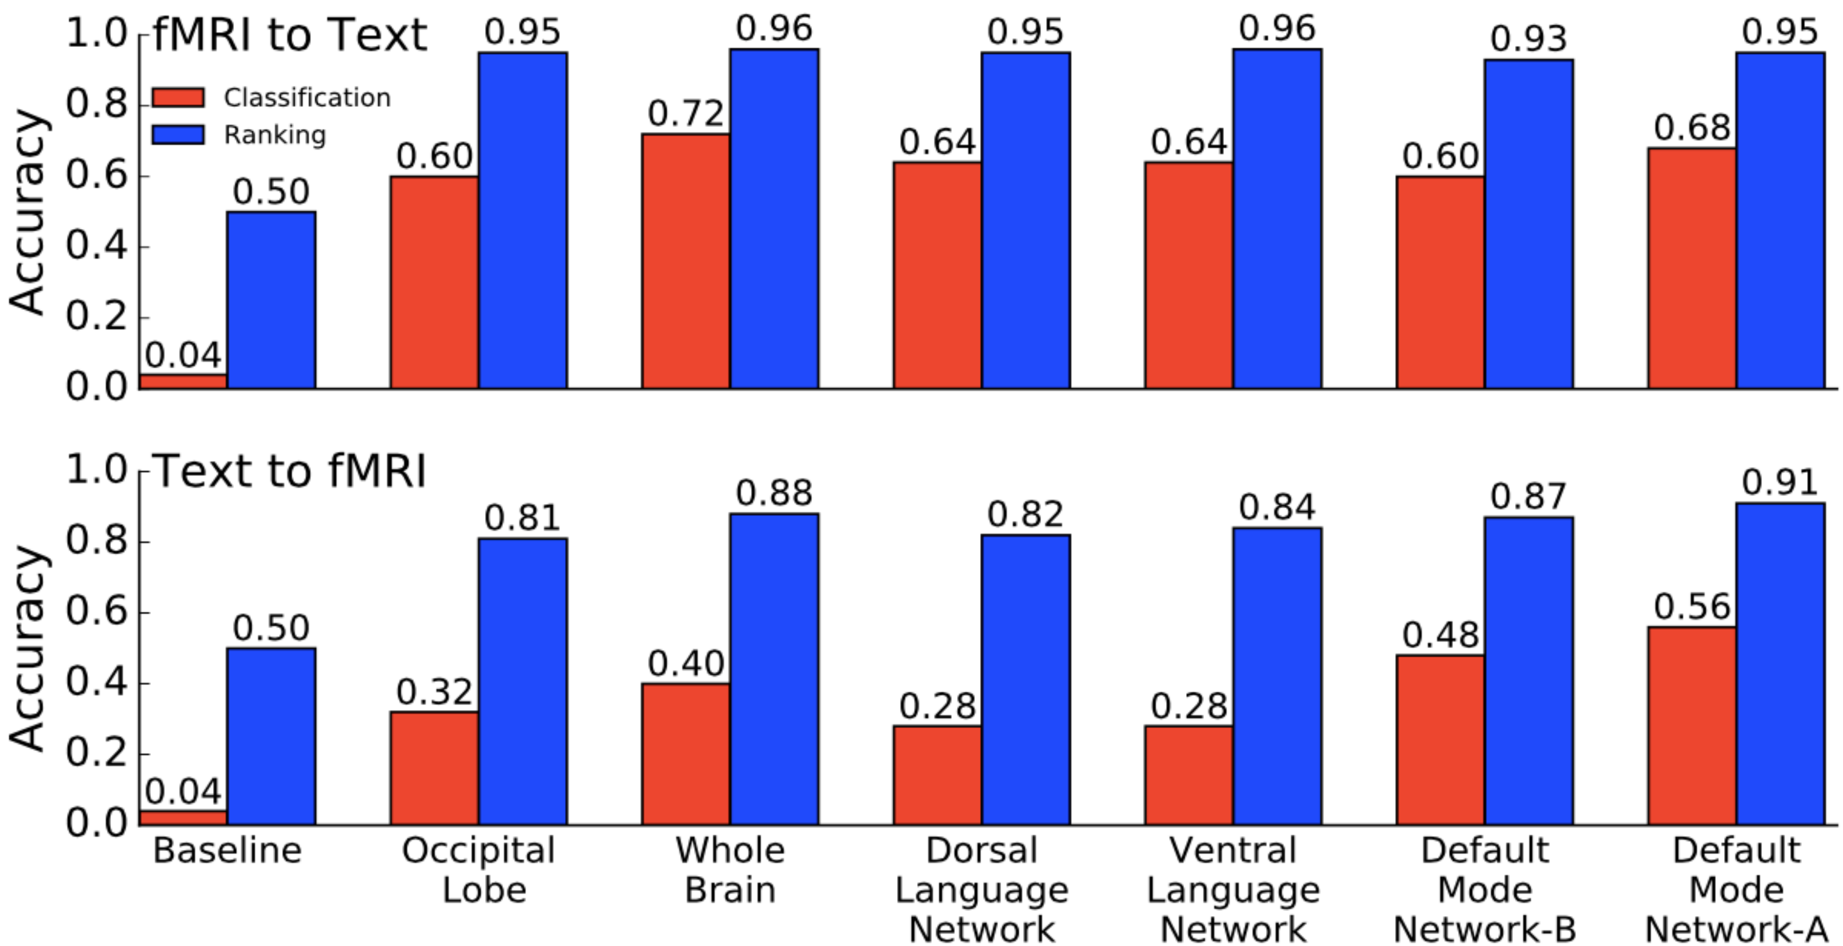

[Vodrahalli and Chen et al. 2016 arXiv]

Part III Discussions and Extensions of SRM

SRM on fMRI

1. Generalize to new subject

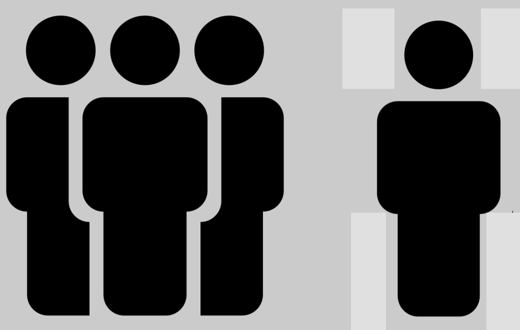

2. Generalize to new stimulus

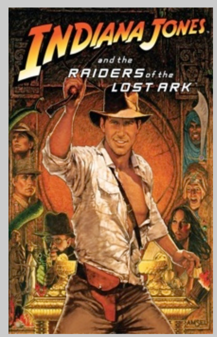

3. Decoupling shared and individual response

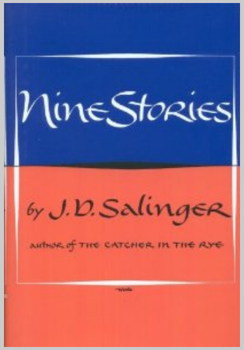

4. SRM with retinotopy

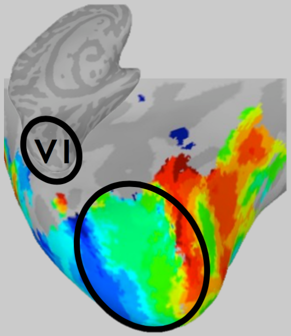

5. Searchlight SRM

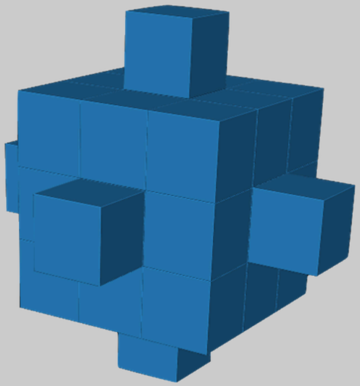

6. Bridging shared space and text semantic space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

How can SRM help?

What can SRM do?

- Multi-subject data driven de-noising
- Aggregation of multi-subject data
- Generalizable to new subject and new stimulus
- Outperform within subject classification
- Decoupling of shared and individual response

Can I use SRM on my data?

- Temporally synchronized stimuli
  - No problem!
- Non-temporally synchronized stimuli
  - Might also work with preprocessing!

When should you consider using SRM?

1. I want to figure out what's shared/not shared in my multi-view data (multi-subject, multi-modality, multi-region, fMRI + stimulus, etc)

2. I have multi-view dataset, I want better prediction accuracy!

A series of extensions of SRM

Semi-supervised SRM

$$\min_{\psi,\theta} (1-\alpha)\,\mathcal{L}_{Align}(\psi) + \alpha\mathcal{L}_{Sup}(\theta;\psi) + R(\theta)$$

[Turek et al. ICASSP 2017]

Independent factor SRM

Algorithm 1: Shared Response ICA (SR-ICA)

```
input  : Data matrices X_i, number of factors k, convergence threshold τ,
         max iteration N, number of subjects m
output: Subject-specific maps W_i and shared response S
W_i^0 ← initialization with random orthonormal columns ;
for n in 1 to N do
    S ← (1/m) Σ_{i=1}^m W_i^{n-1+} X_i          ▷ (·)^+ is pseudo-inverse;
    for i in 1 to m do
        W_i^n ← (E{X_i g(S)} − E{X_i g'(S)} W_i^{n-1+})^+;
        W_i^n ← W_i^n (W_i^{nT} W_i^n)^{-1/2};
    end
    converged ← True;
    for i in 1 to m do
        if max |W_i^{nT} W_i^{n-1} − I| ≥ τ then
            converged ← False;
        end
    end
    return W_i, S;
end
```

[Zhang et al. ArXiv 2016]

Scaling up to thousand subjects

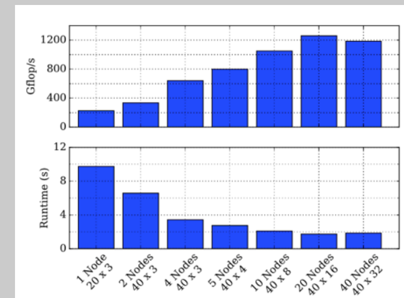

[Anderson et al. IEEE Bigdata 2016]

Kernelized SRM

$$\min \|\Phi_i - \Phi_i \tilde{A}_i \tilde{S}\|_F^2$$

$$\text{s.t.}\,\tilde{A}_i^T \mathbf{K}_i \tilde{A}_i = I_k.$$

Gaussian Process SRM

$$\begin{aligned}
\mathbf{s}_{ri} &\sim \mathcal{GP}(0, \mathbf{K}_{\mathbf{s}_i}(t,t')),\\
\mathbf{x}'_{mt}|\mathbf{s}_t &\sim \mathcal{N}(W_m \mathbf{s}_t + \mu_m, \rho_m^2 I),\\
\text{s.t. } W_m^T W_m &= I,\\
[\mathbf{s}_{r1} \ldots \mathbf{s}_{ri} \ldots \mathbf{s}_{rN}]^T &= [\mathbf{s}_1 \ldots \mathbf{s}_t \ldots \mathbf{s}_T],
\end{aligned}$$

And more!

Code ready to use!

https://github.com/IntelPNI/brainiak

- Simple setting, one line code to fit a model to your data

Open source software contribution

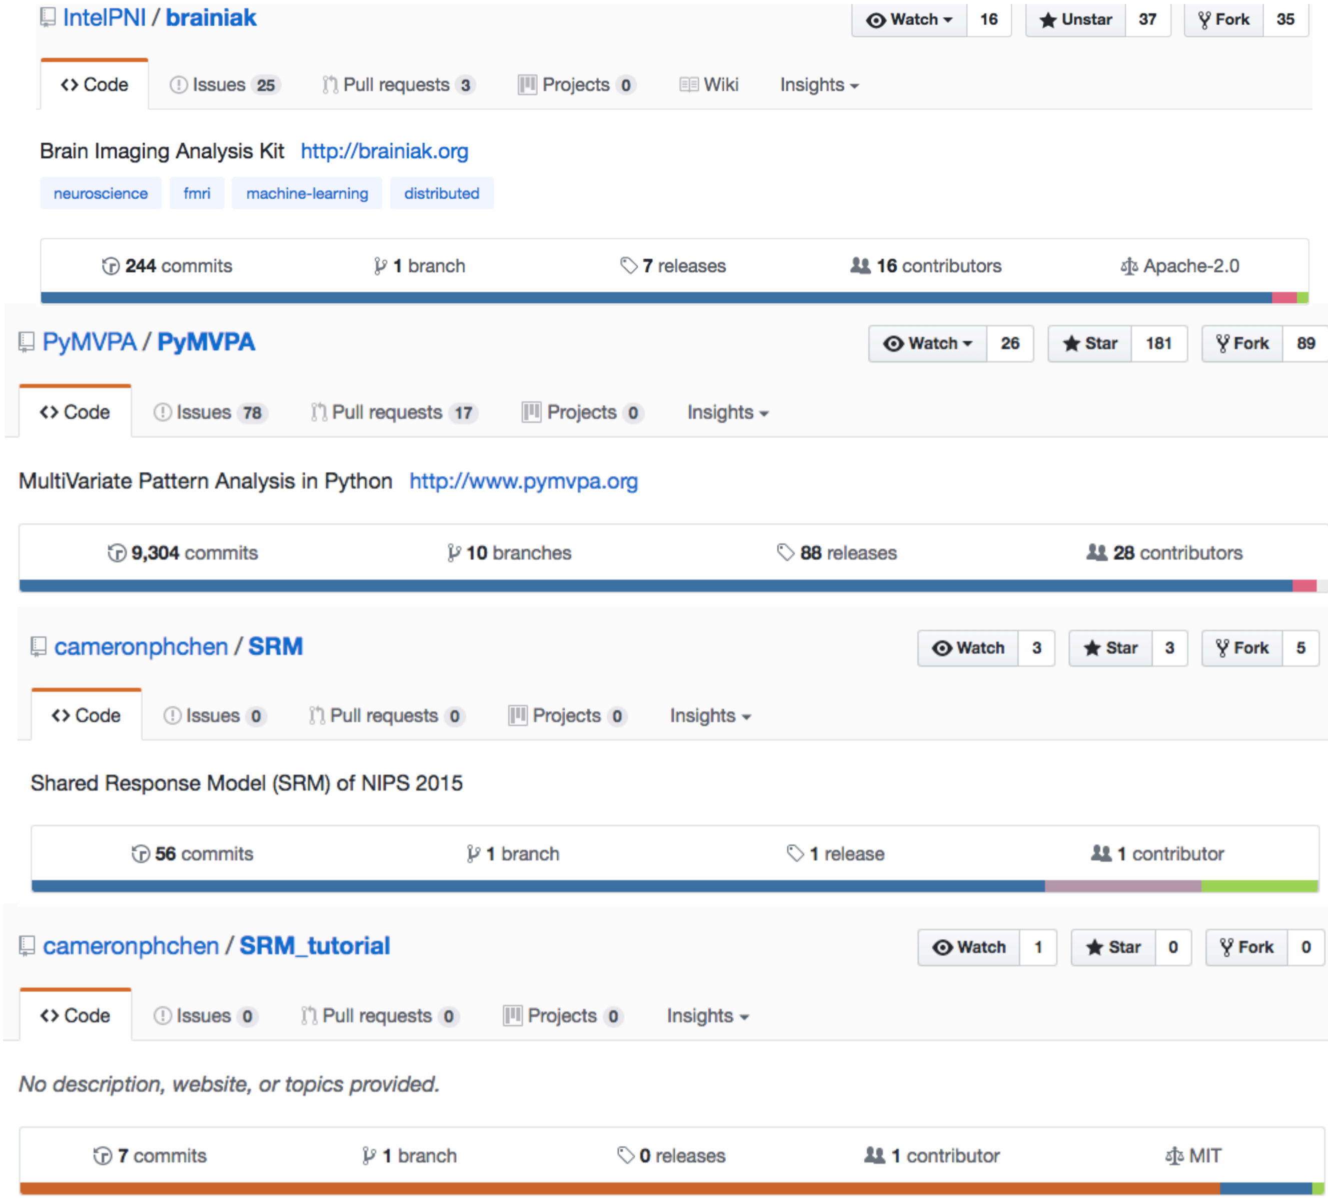

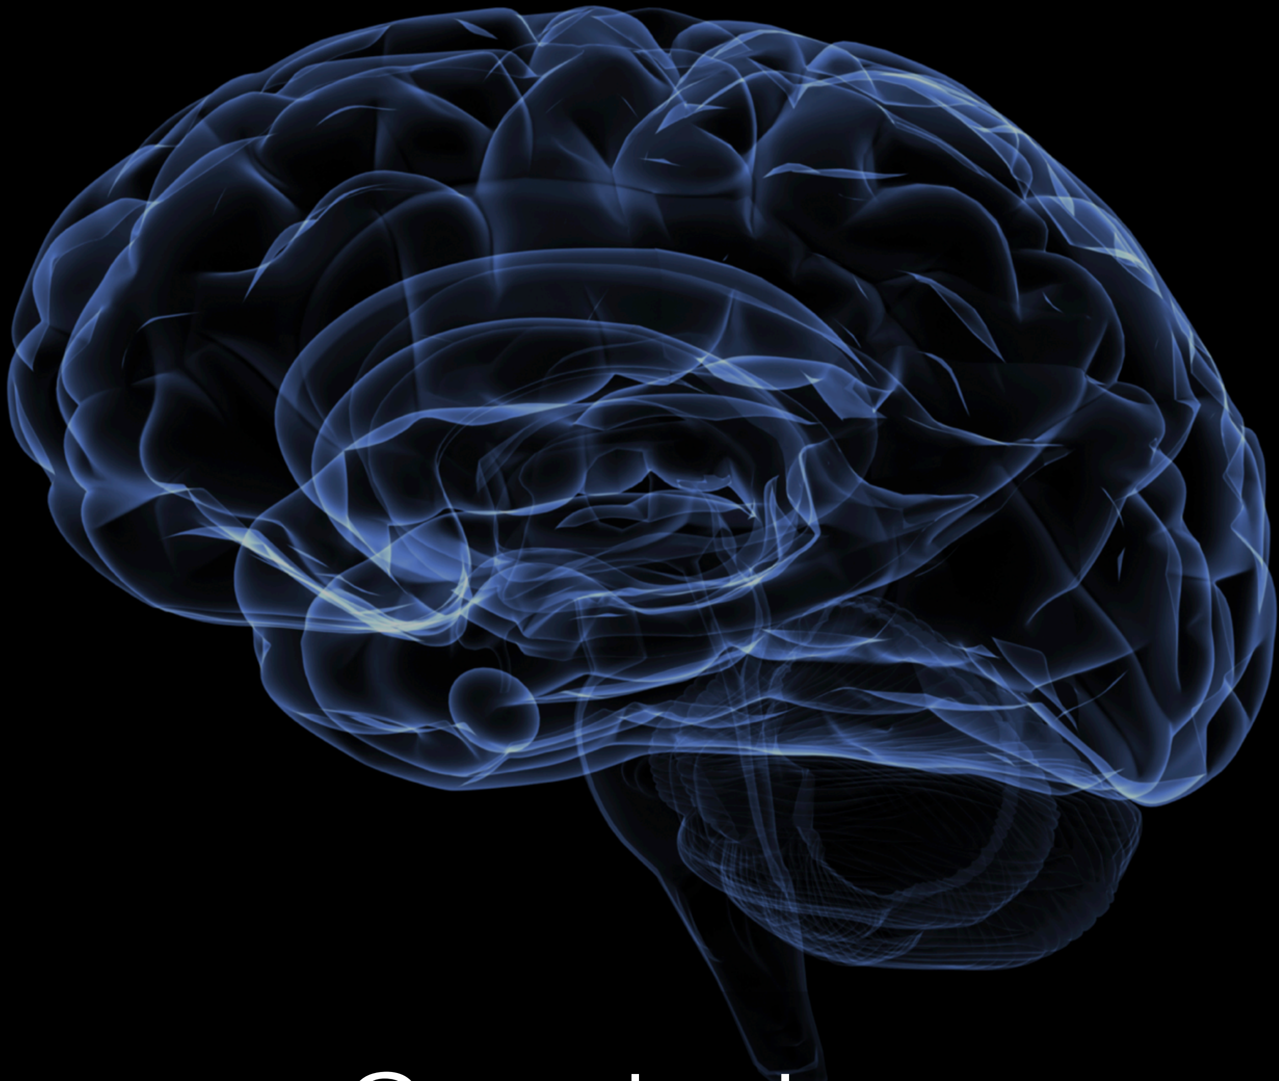

Conclusion

Conclusion

Proposed a multi-view learning framework

Conclusion

Proposed a multi-view learning framework

Developed SRM and many other models from the framework

Conclusion

Proposed a multi-view learning framework

Developed SRM and many other models from the framework

Demonstrated these models on real fMRI in various settings

How can these help us learn more about the brain ?

Increase statistical power from aggregated data

How can these help us learn more about the brain ?

Increase statistical power from aggregated data

Learn more about the distribution of information in the brain

How can these help us learn more about the brain ?

Increase statistical power from aggregated data

Learn more about the distribution of information in the brain

Open up new possibilities for analyzing neuroimaging data

The Spirit Carries On!

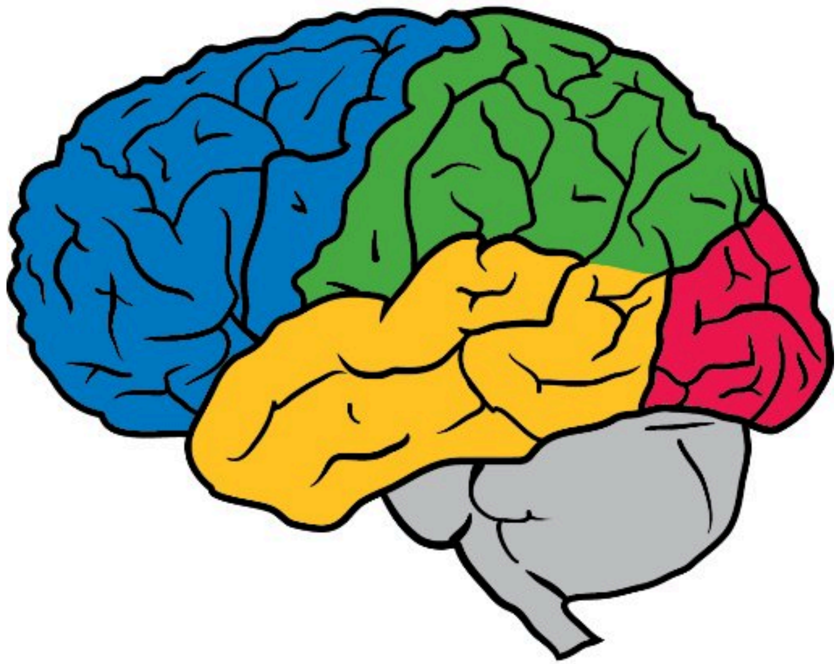

Google Brain

Research and Machine Intelligence

Machine Learning and Deep Learning on Healthcare and Medical Imaging

Acknowledgement

- Prof. Peter Ramadge
- Prof. Uri Hasson
- Prof. Ken Norman

Thesis Readers

- Prof. Janice Chen
- Prof. Yuxin Chen

General Committee

- Prof. David Blei
- Prof. Paul Cuff

Ramadge Lab

- Hejia Zhang
- Hossein Valavi
- Xu (Mia) Chen
- Yun Wang
- David Eis
- Hao Xu
- Pingmei Xu
- Alex Lorbert
- Selina Man

PNI

- Prof. Jonathan D Cohen
- Prof. Kenneth A Norman
- Prof. Nicholas Turk-Browne
- Prof. Jonathan Pillow
- Prof. Janice Chen
- Prof. Jeremy R. Manning
- J. Benjamin Hutchinson
- Christopher A. Baldassano
- Edwin C. Clayton
- Mingbo Cai
- Michael Shvartsman
- Yida Wang
- Yaara Yeshurun
- Michael J. Arcaro
- Kiran Vodrahalli
- Sebastian Musslick
- Anqi Wu

Intel

- Theodore L Willke
- Javier Turek
- Xia (Ivy) Zhu
- Mihai Capota

Haxby Lab

- Prof. James Haxby
- Prof. Michael Hanke
- Prof. Yaroslav O. Halchenko
- J. Swaroop Guntupalli

- ELE Staff
- PNI Staff
- PNI Help Desk

Funding

- Google PhD Fellowship
- Princeton University Fellowship in Natural Science and Engineering
- Taiwan Ministry of Education Study Abroad Scholarship
- Intel
- NSF

Friends and Family!


CAUTION
GAP

Thank you!!

Back up

A coherent multi-view framework for all three problems

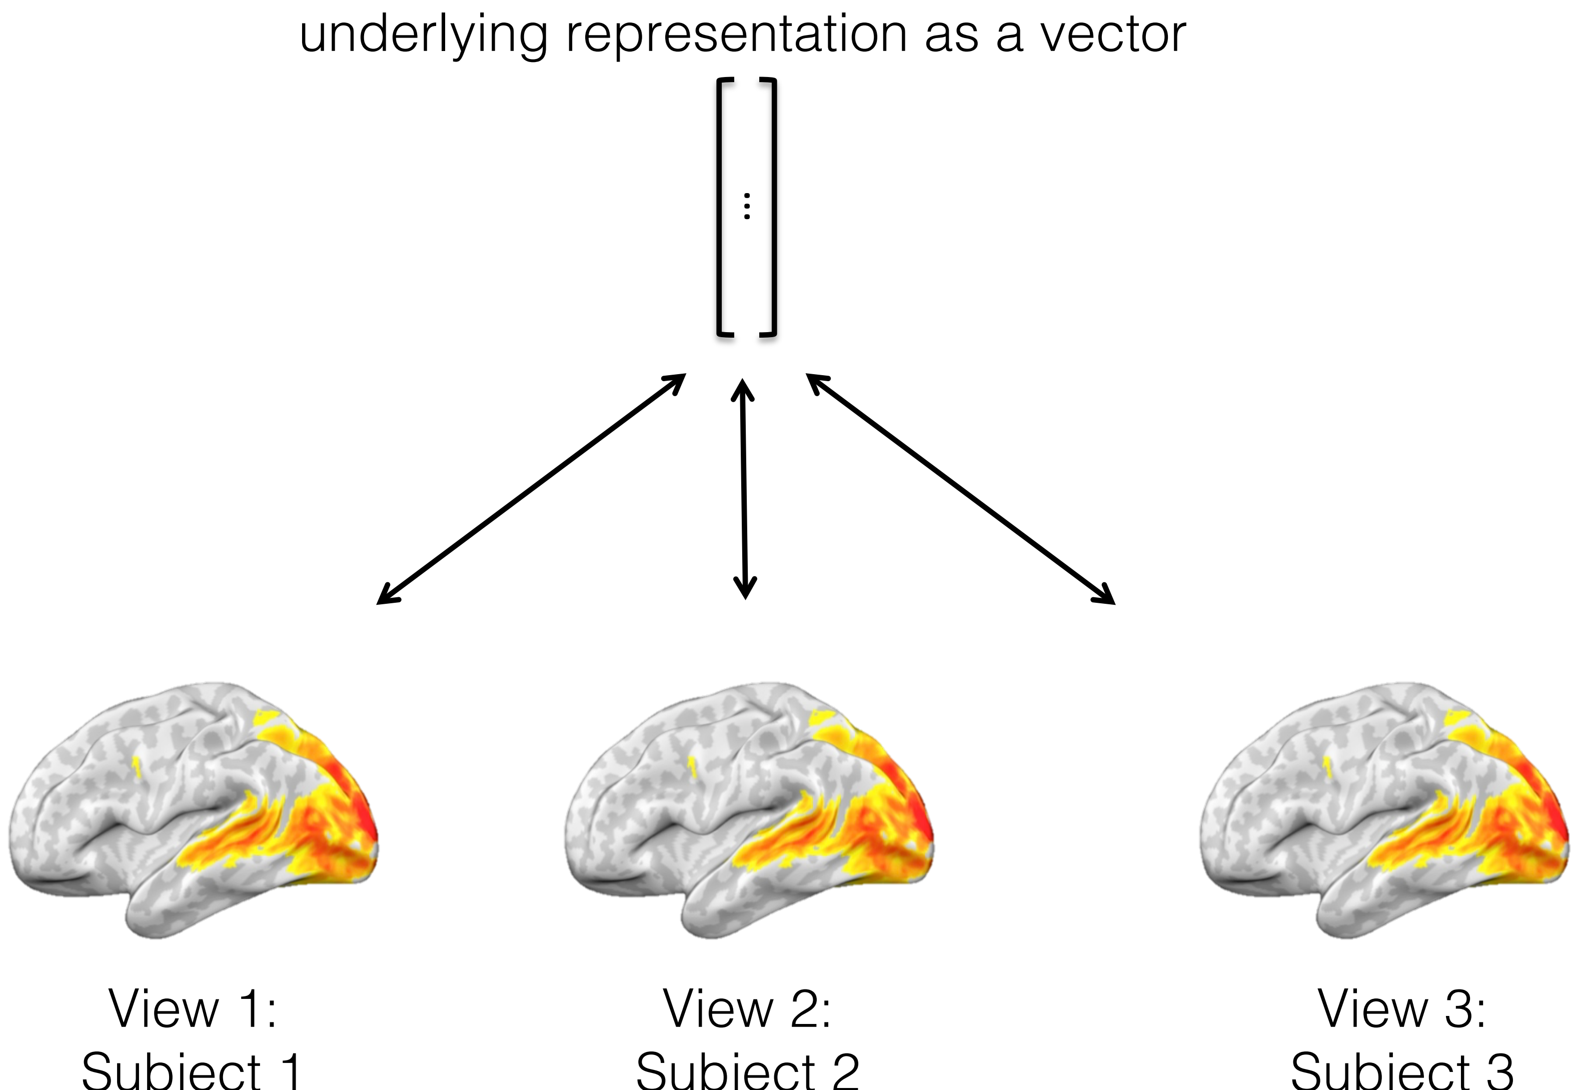

What is multi-view learning?

- Exist an unknown underlying representation, and each view is a realization of it
- Multi-view learning models estimate transformations between views and representation

Questions to think about before using SRM on fMRI data

What are the views?

What is the hypothesis that we are testing?

Which space are we analyzing in?

Multi-view in fMRI data can be of various forms

Multi-subject

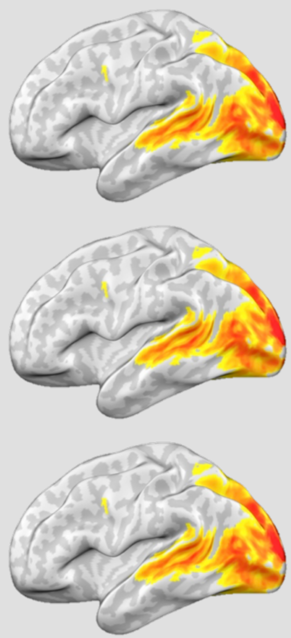

Stimulus+fMRI

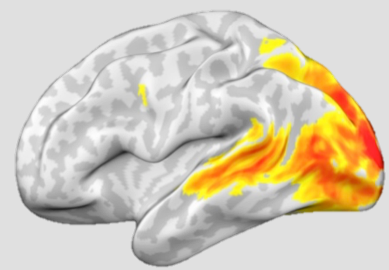

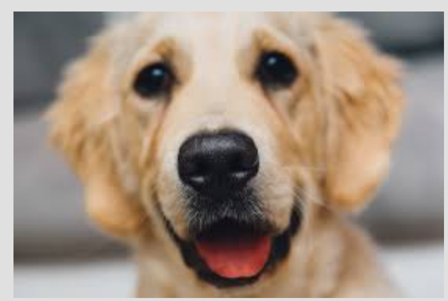

and more!

Conventional ML models disregard variability across views

Tailarach

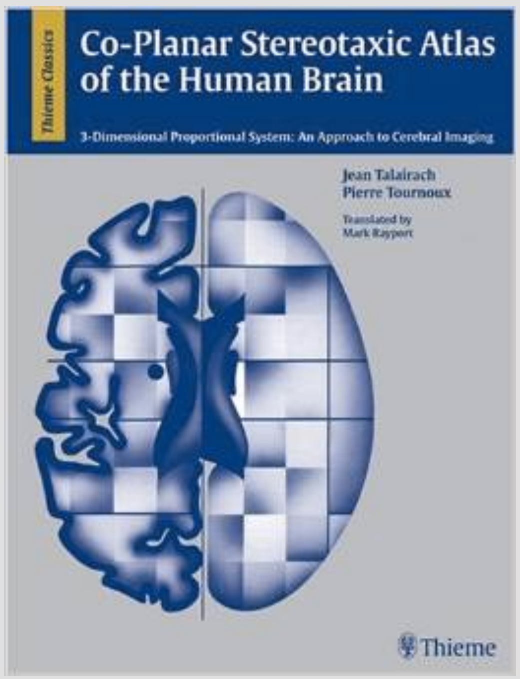

Cortical Surface

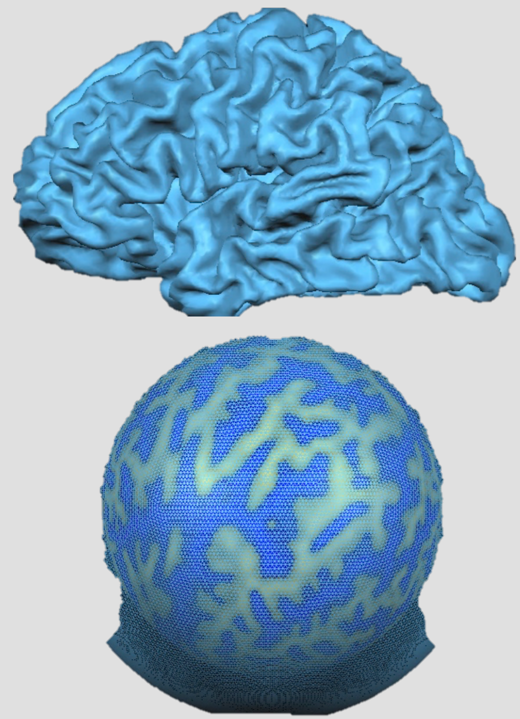

Challenges

Tailarach

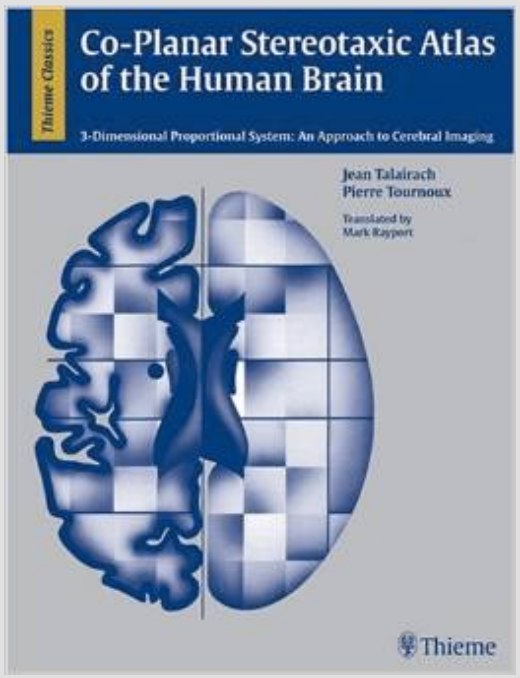

Cortical Surface

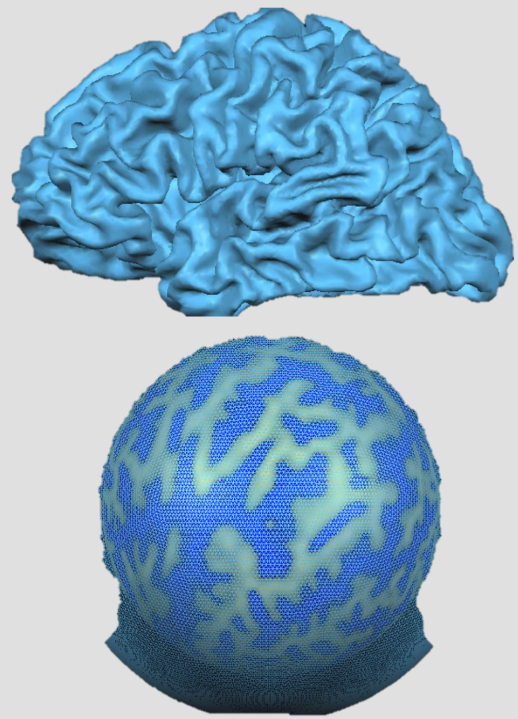

Conventional ML models disregard variability across views

Tailarach

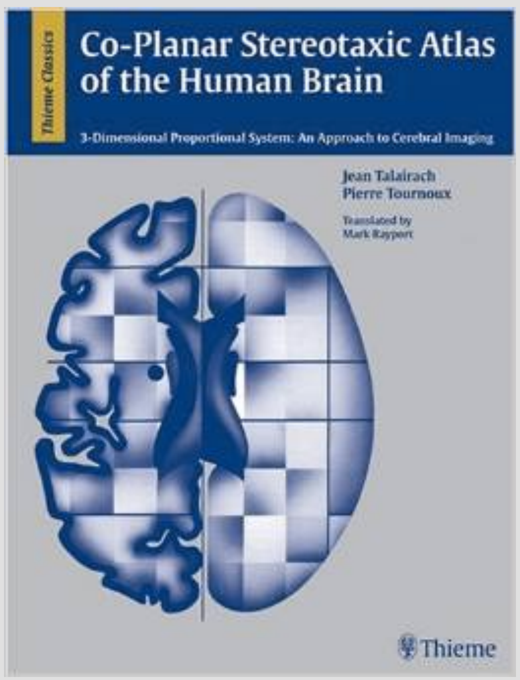

Cortical Surface

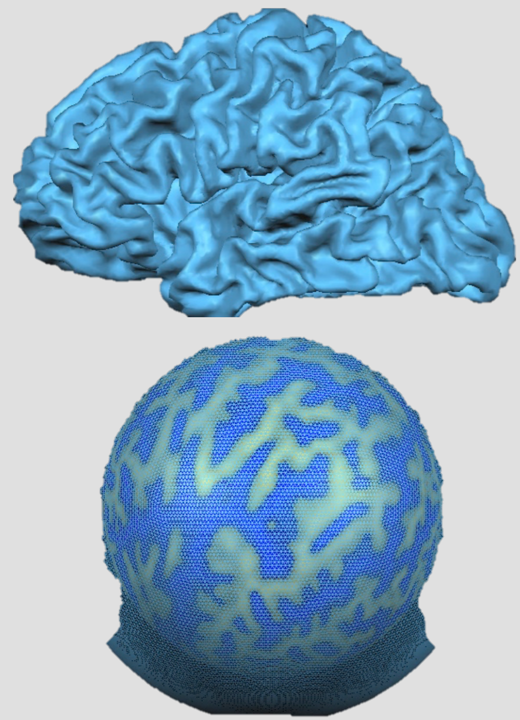

The Need for Multi-view Learning in Neuroimaging

Generalizing findings across subjects

Aggregating data for statistical power

Mapping data between views

What is multi-view learning?

underlying representation as a vector

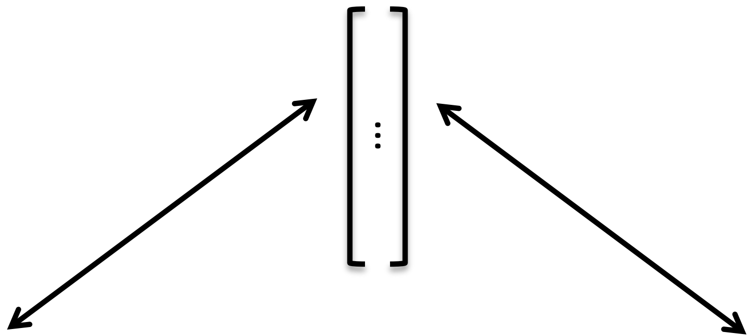

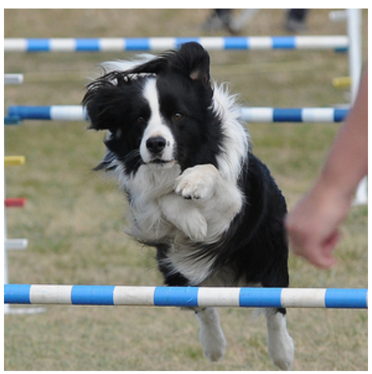

view1: image

"black and white dog
jumps over bar"

view2: text

Neuroimaging measures brain activity

Multi-view Representation Learning Examples

Image Caption Generation

"man in black shirt is playing guitar."

"construction worker in orange safety vest is working on road."

"girl in pink dress is jumping in air."

"black and white dog jumps over bar."

Multi-Language Translation

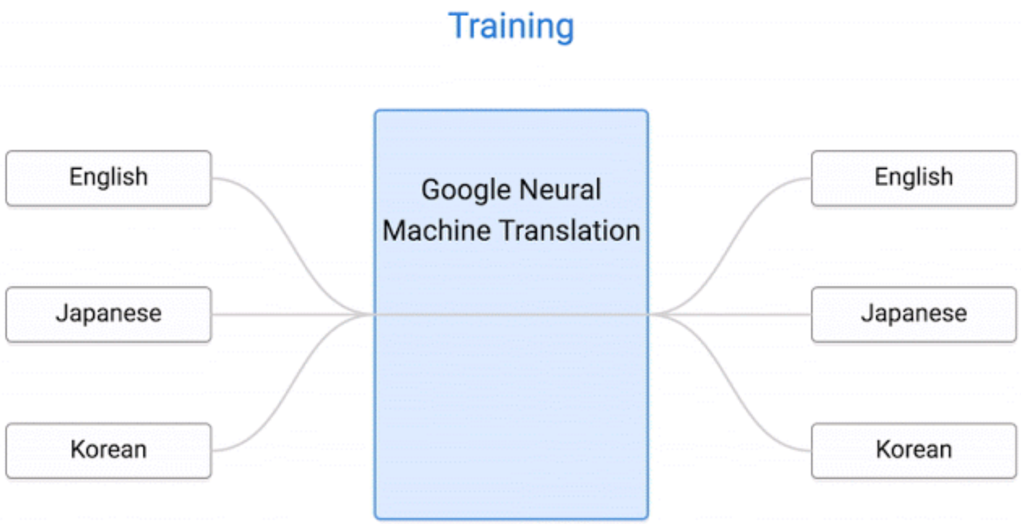

source: Karpathy and Li. "Deep visual-semantic alignments for generating image descriptions." CVPR 2015.
https://research.googleblog.com/2016/11/zero-shot-translation-with-googles.html

The Need for Multi-view Learning in Neuroimaging

Example: Multi-Language Translation as Multi-view Representation Learning

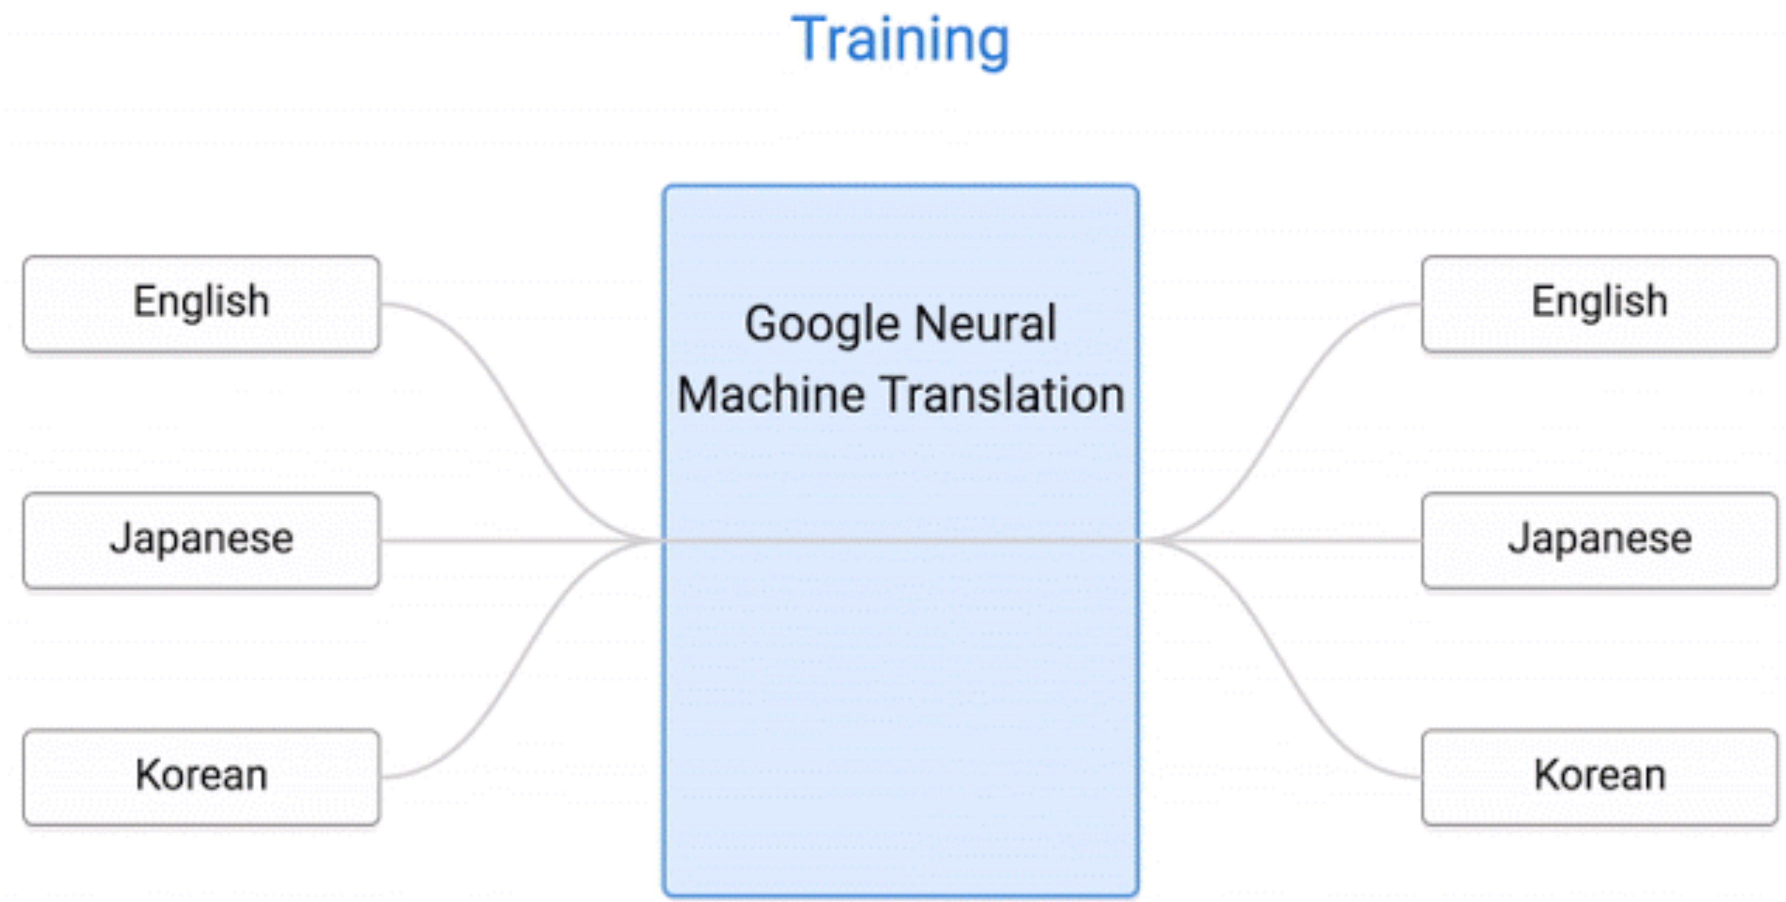

source : https://research.googleblog.com/2016/11/zero-shot-translation-with-googles.html

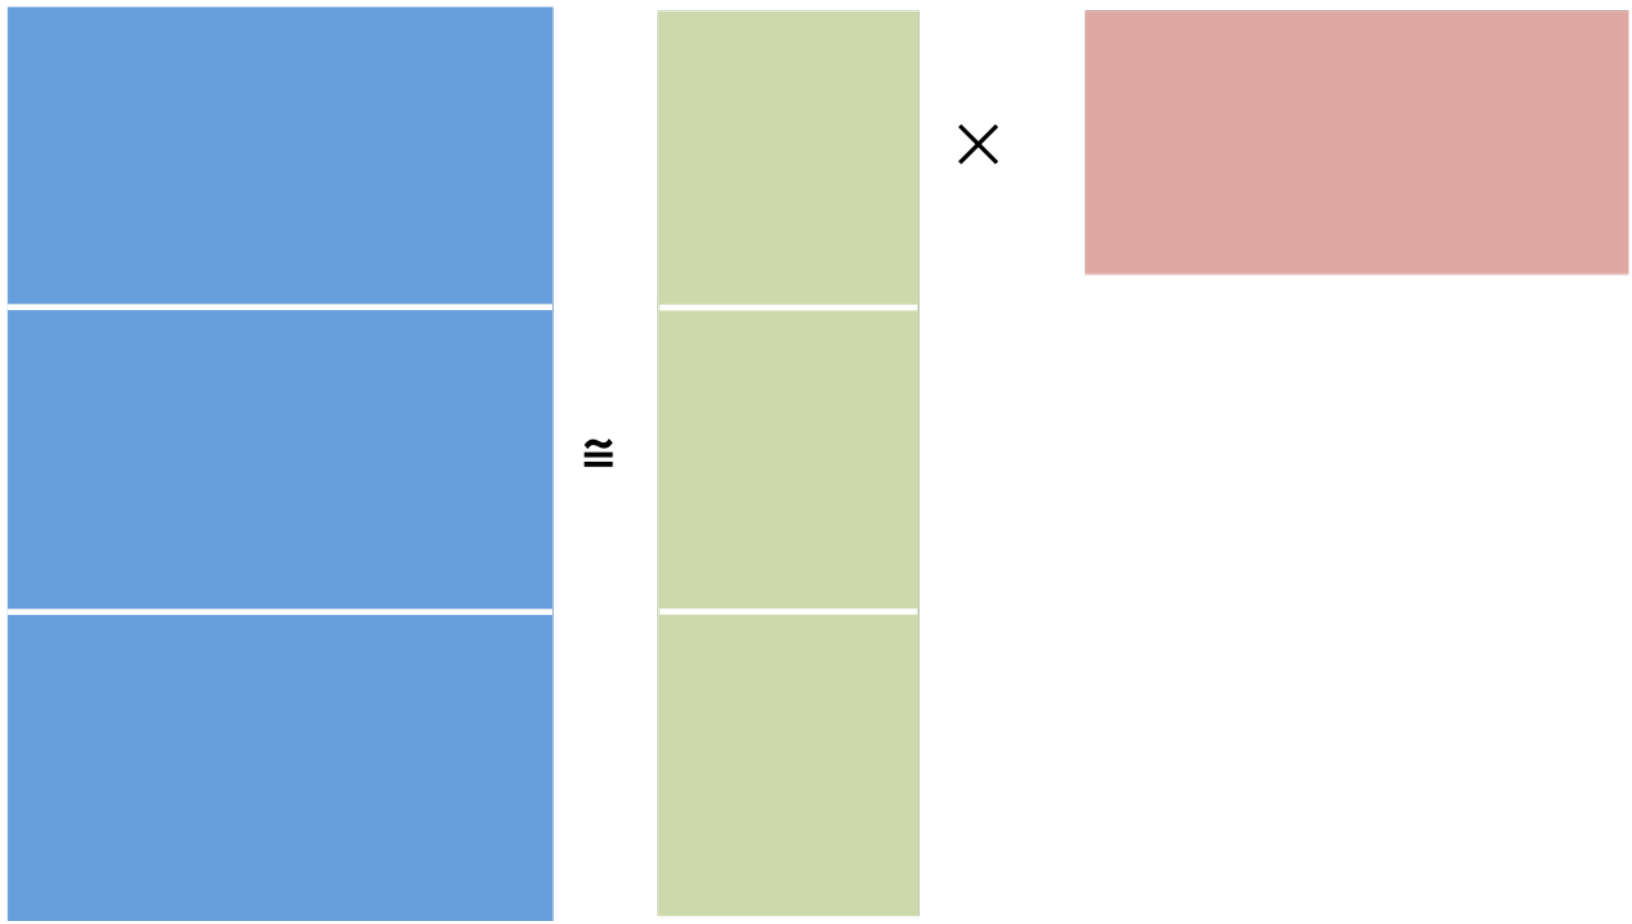

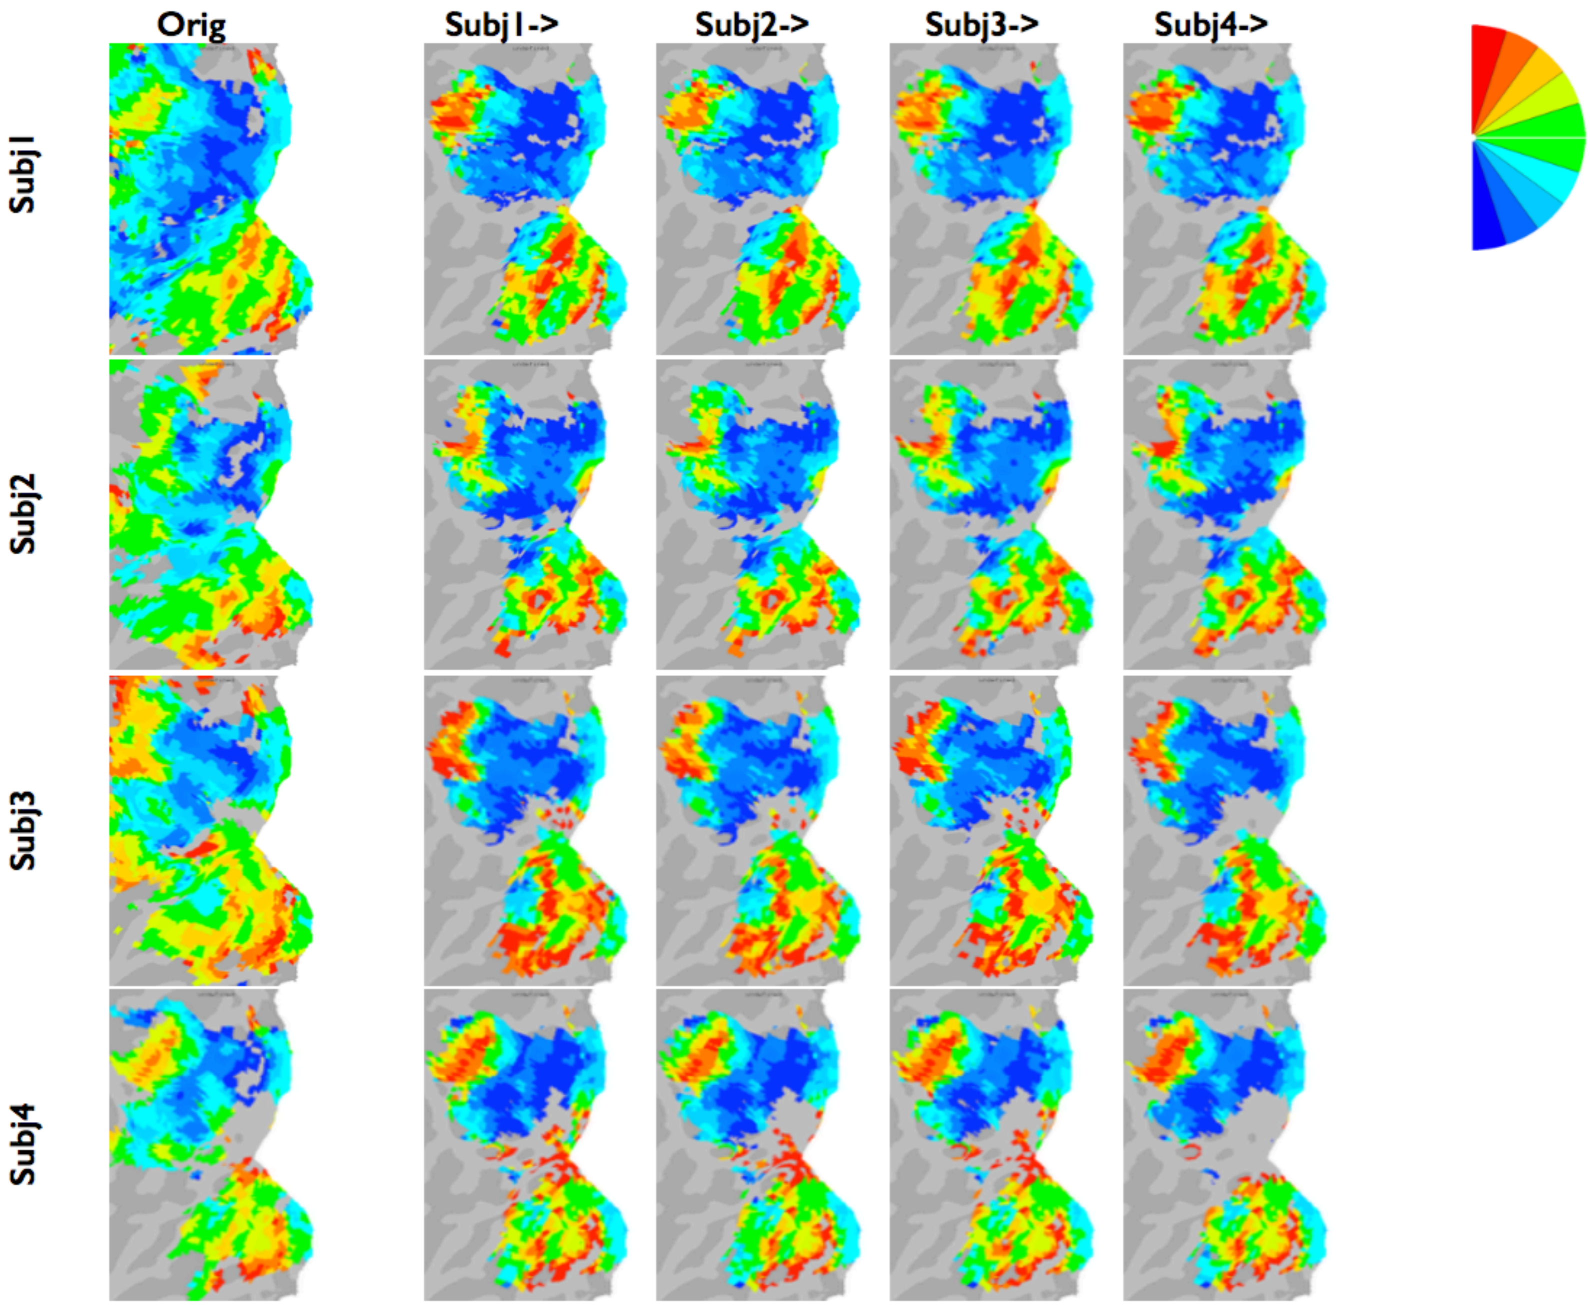

[Work by Michael J. Arcaro]

A generative model

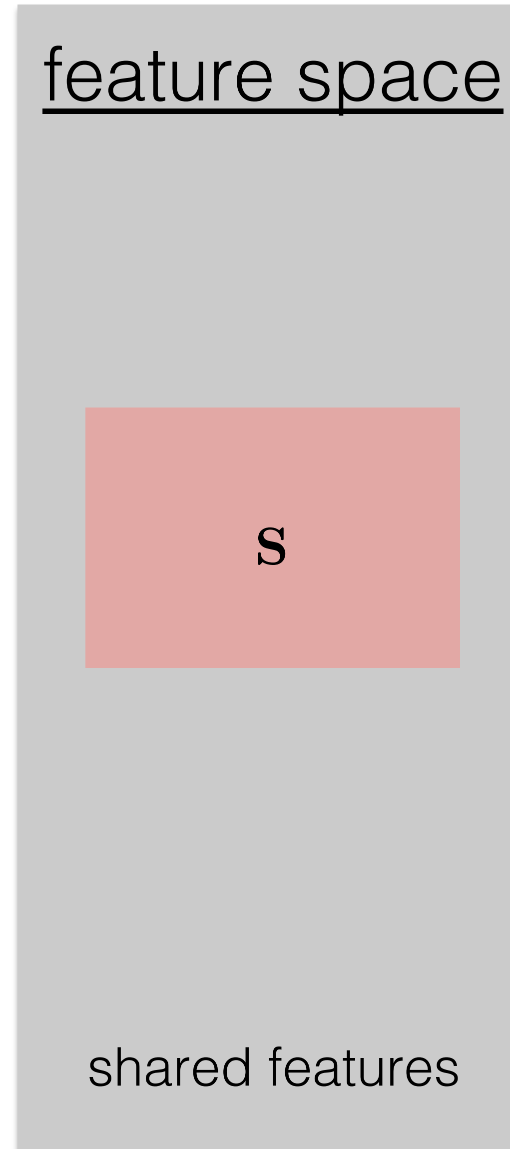

A generative model

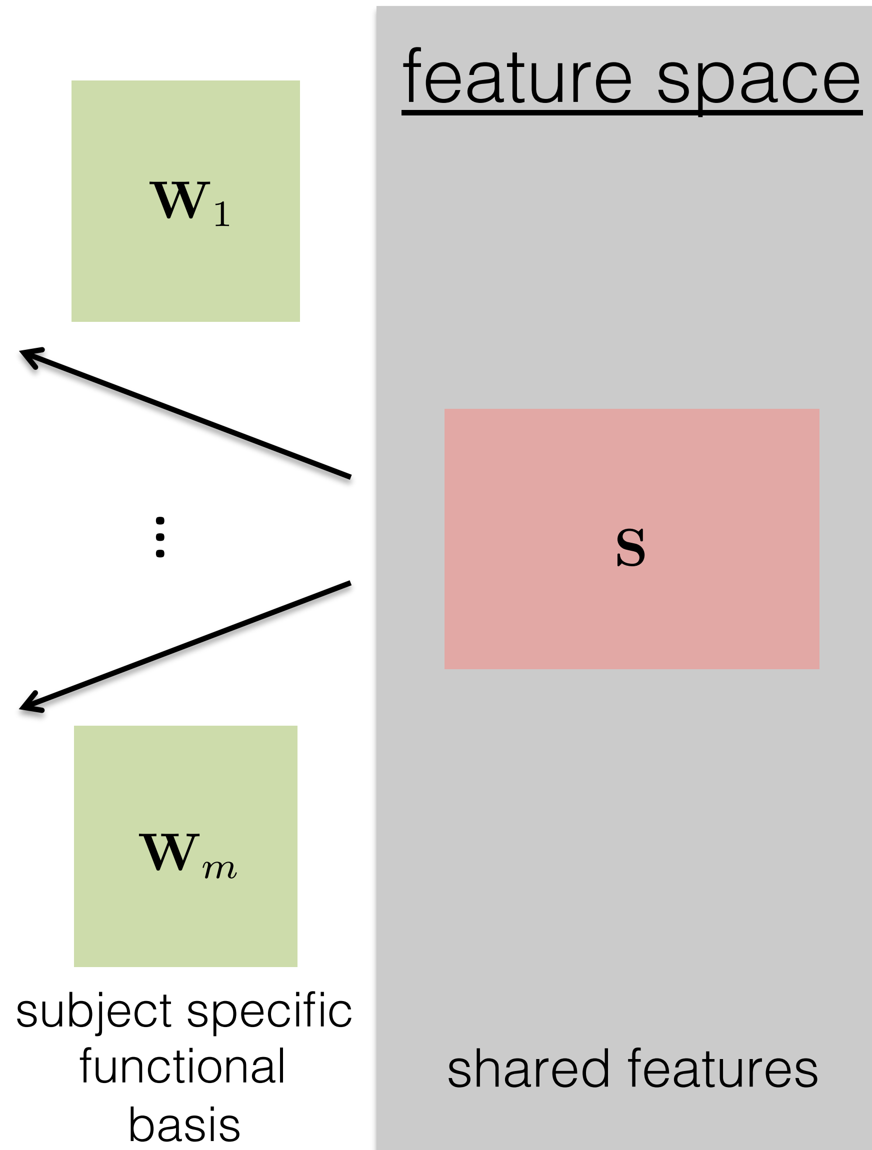

A generative model

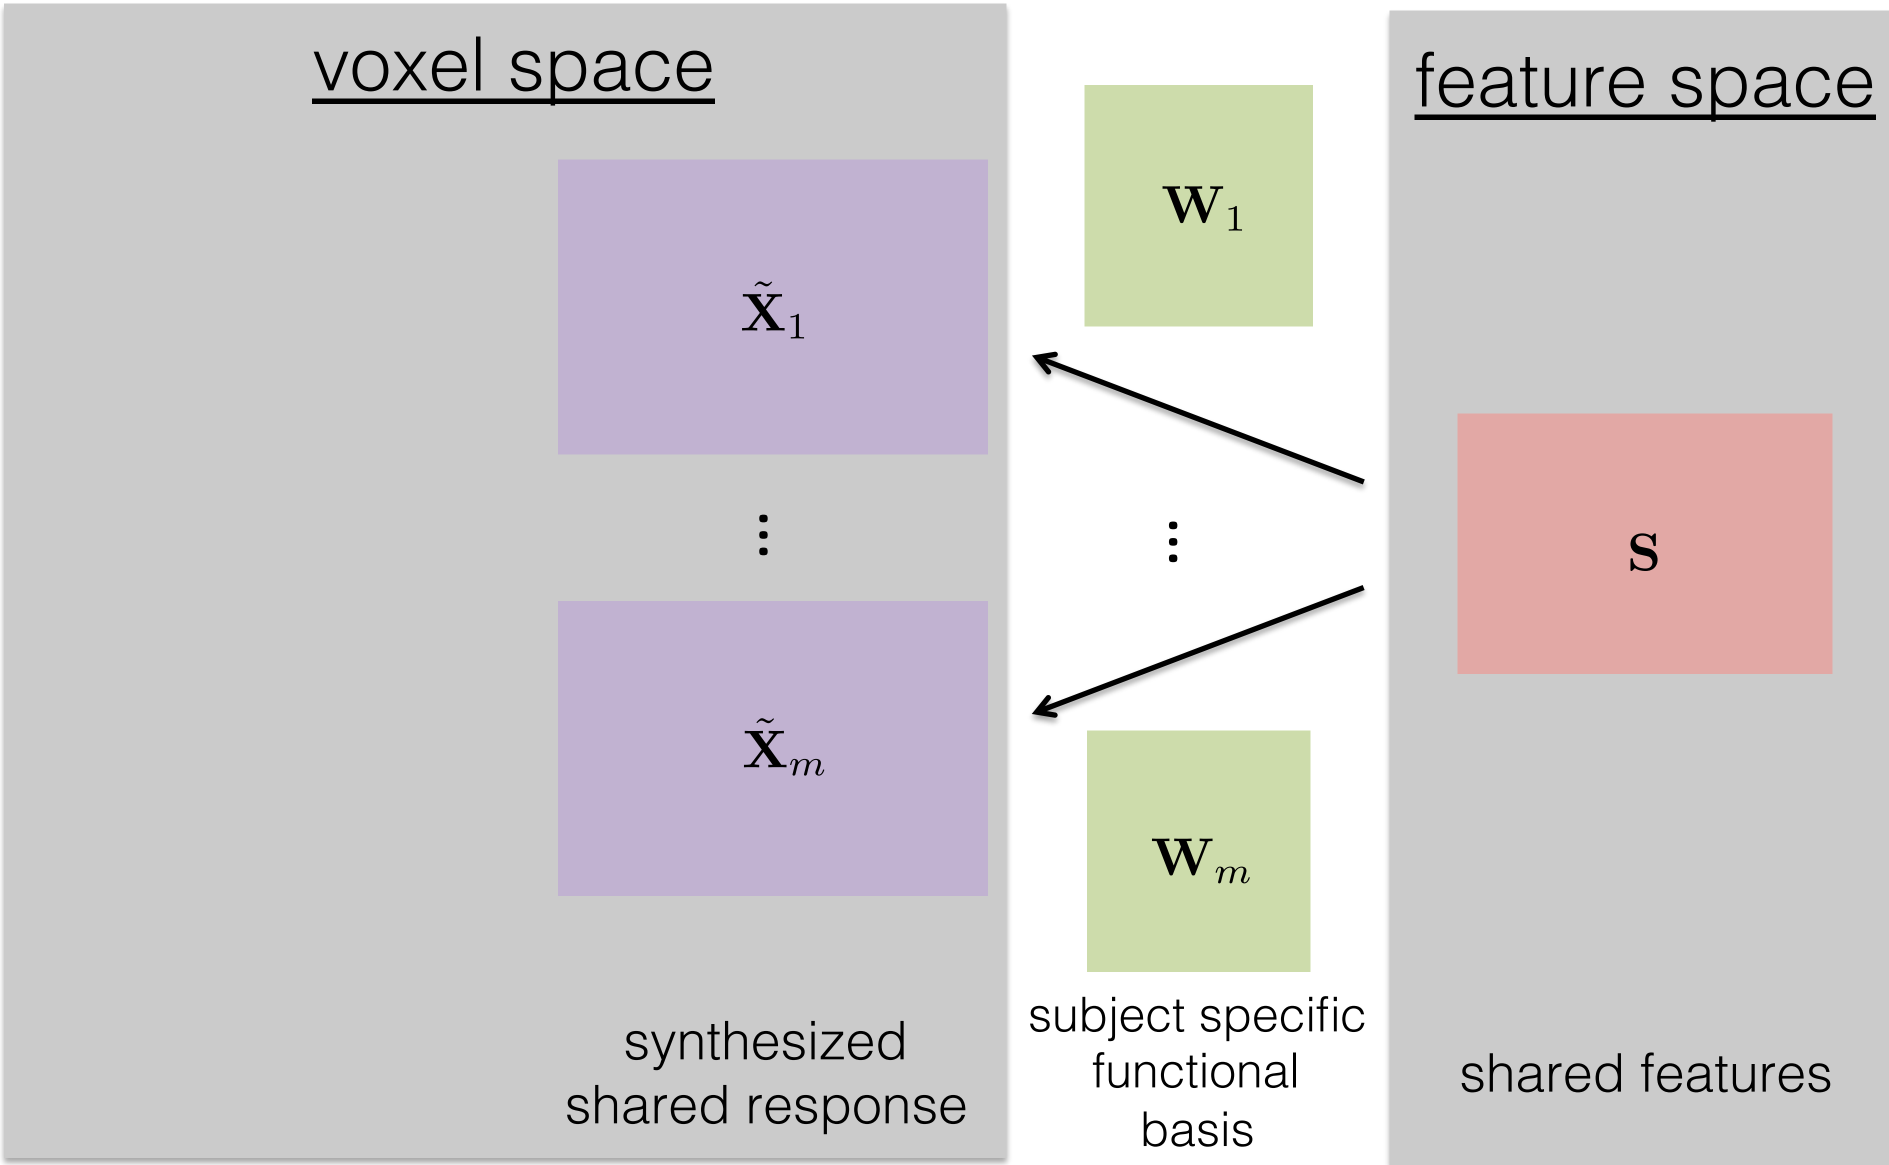

A generative model

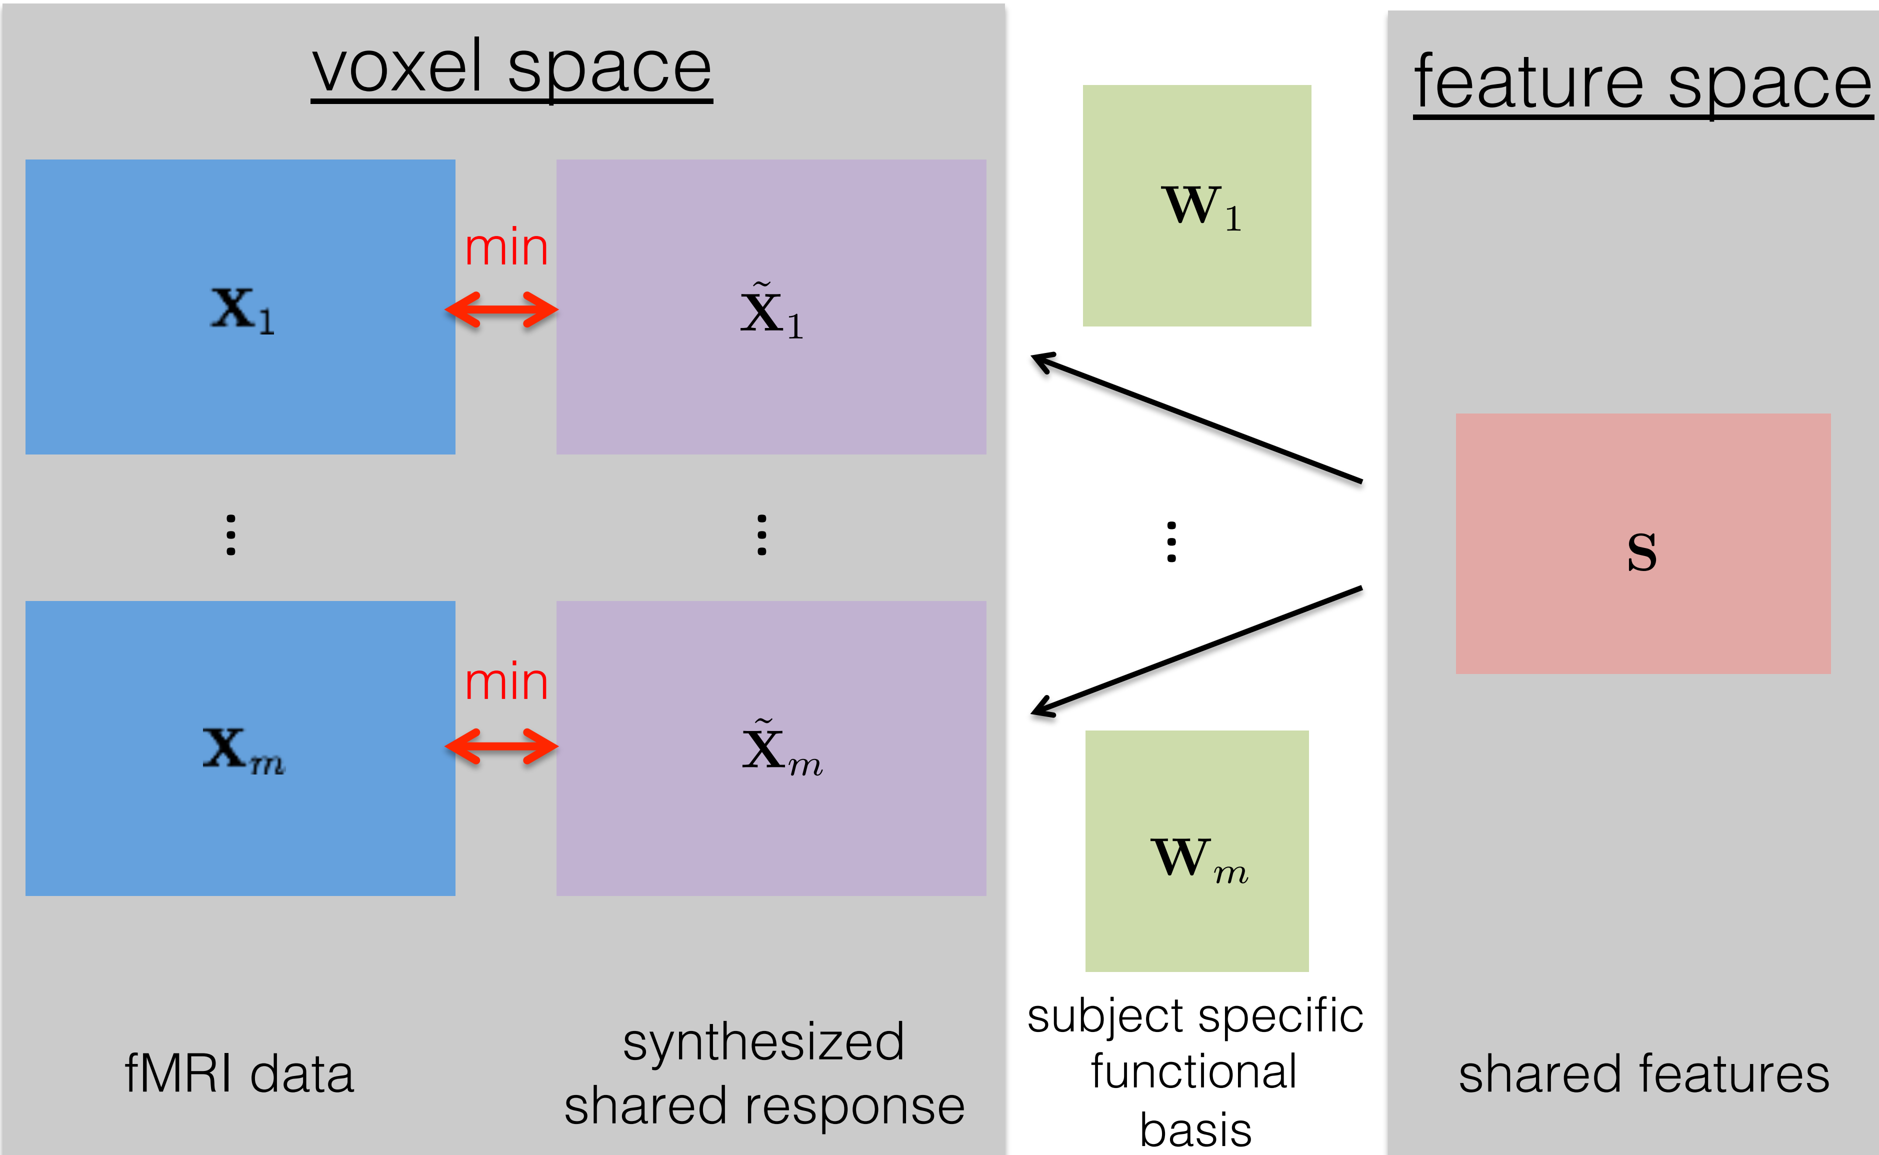

Given data from training views,

<u>Prediction:</u>
can we predict the brain response of a test subject?

<u>Classification:</u>
given brain response from a test subject, can we classify what’s the stimulus?

Differentiating between groups

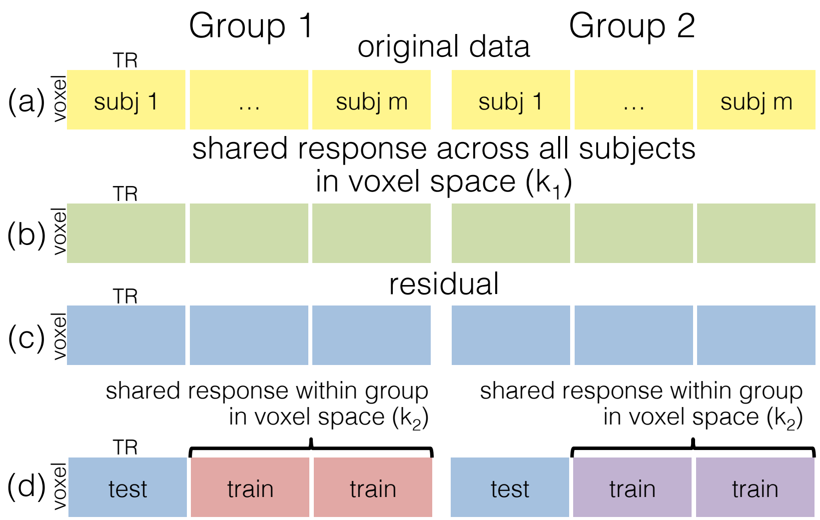

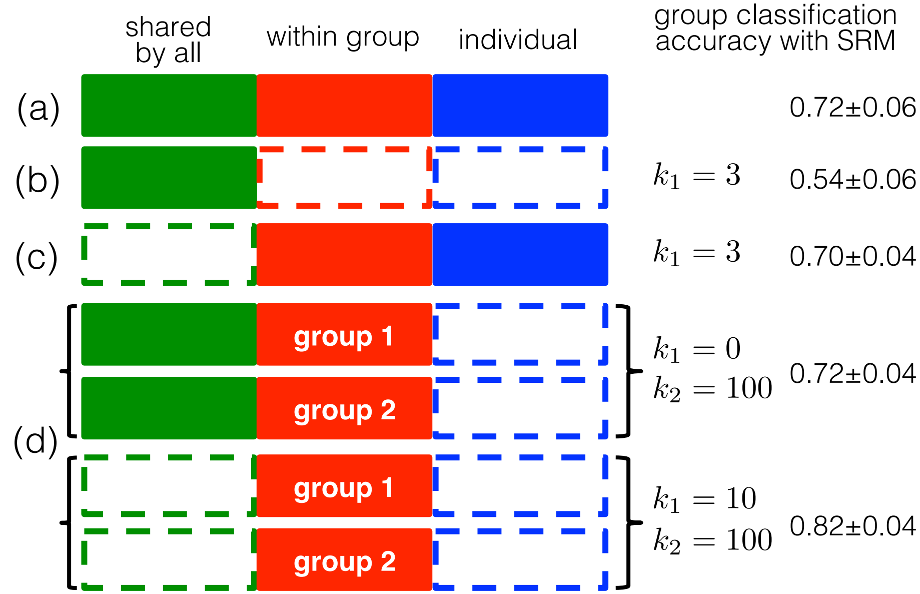

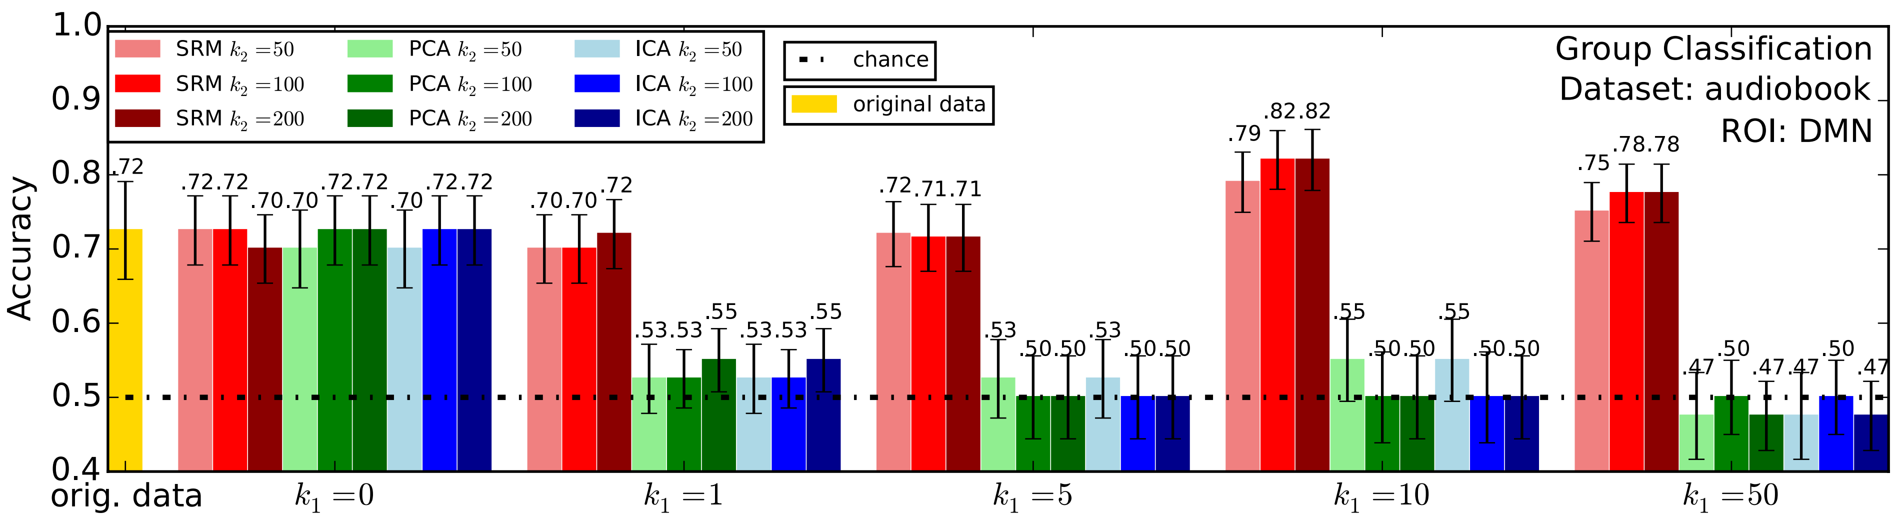

SRM with non-temporally synchronized dataset

SRM with non-temporally synchronized dataset

- Each observation is a noisy sample of the brain state

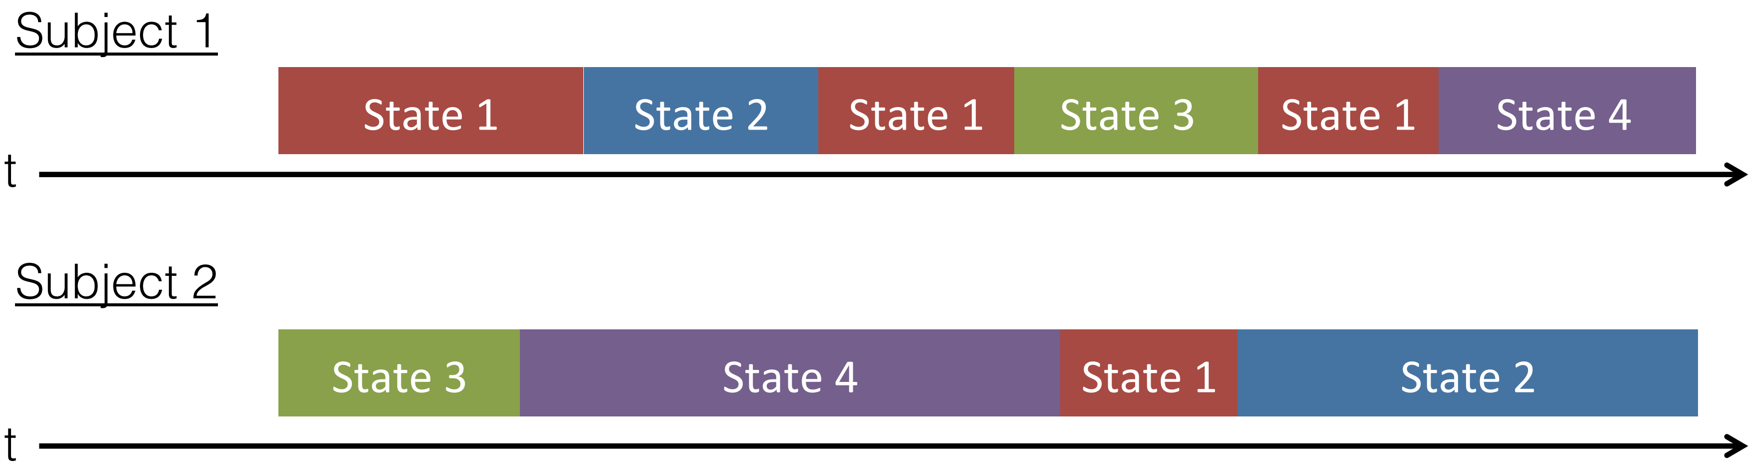

SRM with non-temporally synchronized dataset

Step 1: reordering

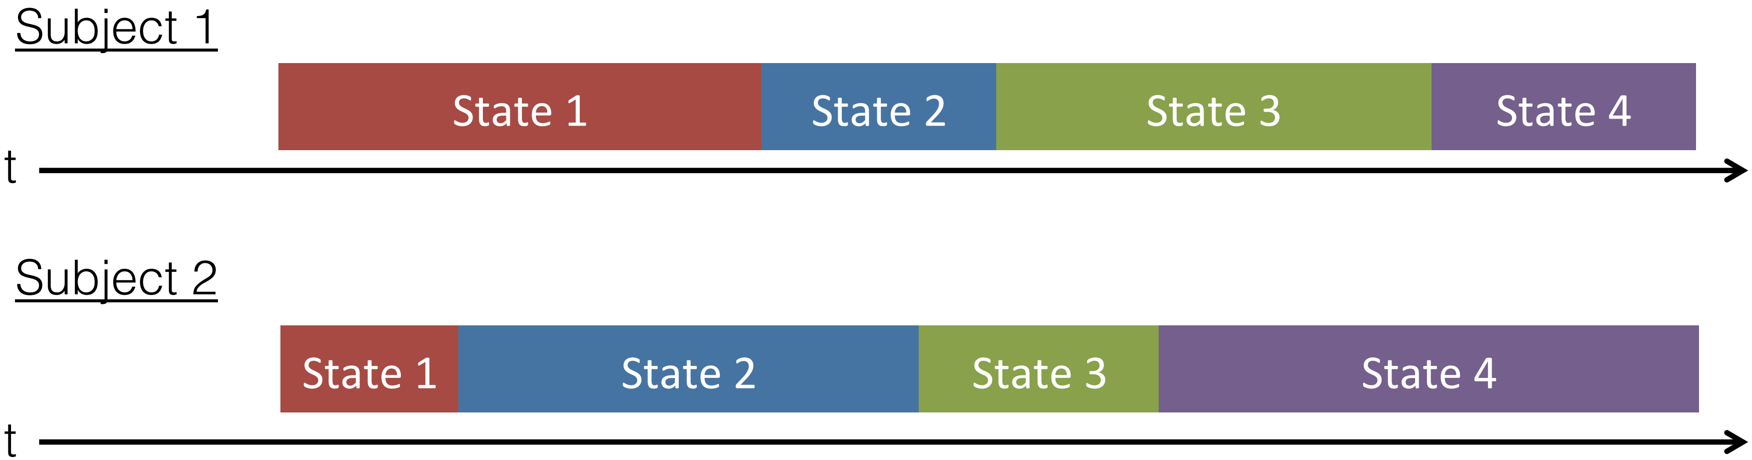

SRM with non-temporally synchronized dataset

Step 2: down sampling

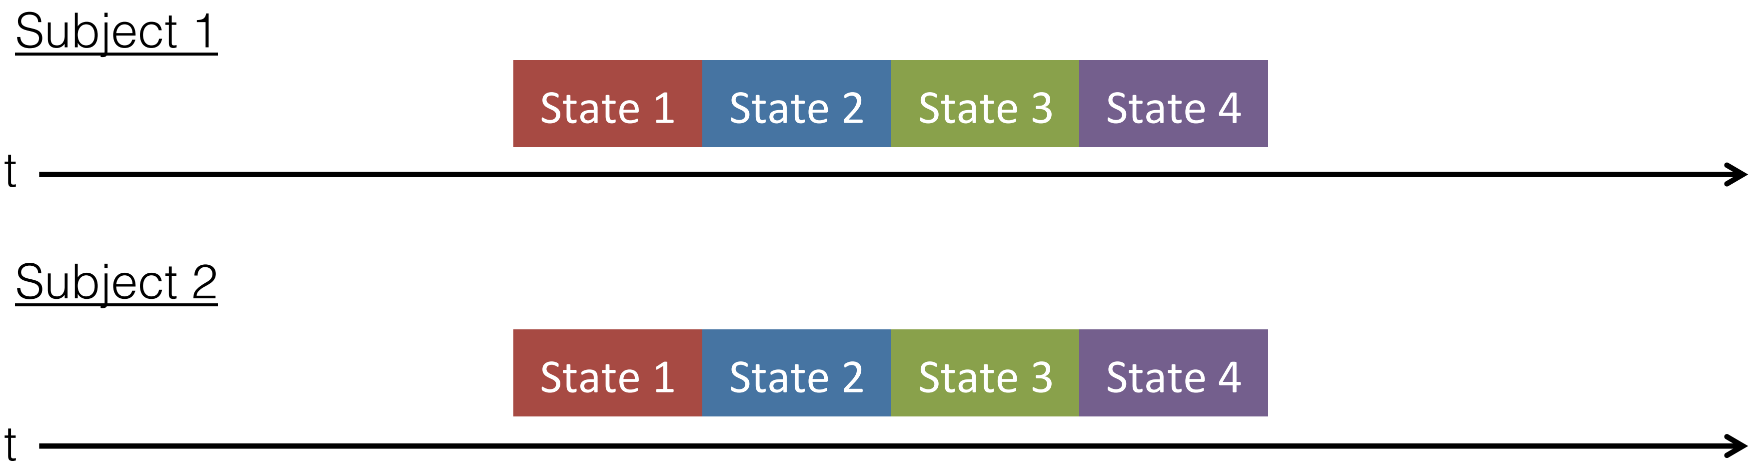

Step 3: fit SRM with preprocessed data

Quantifying dimensionality of shared response

Quantifying dimensionality of shared response

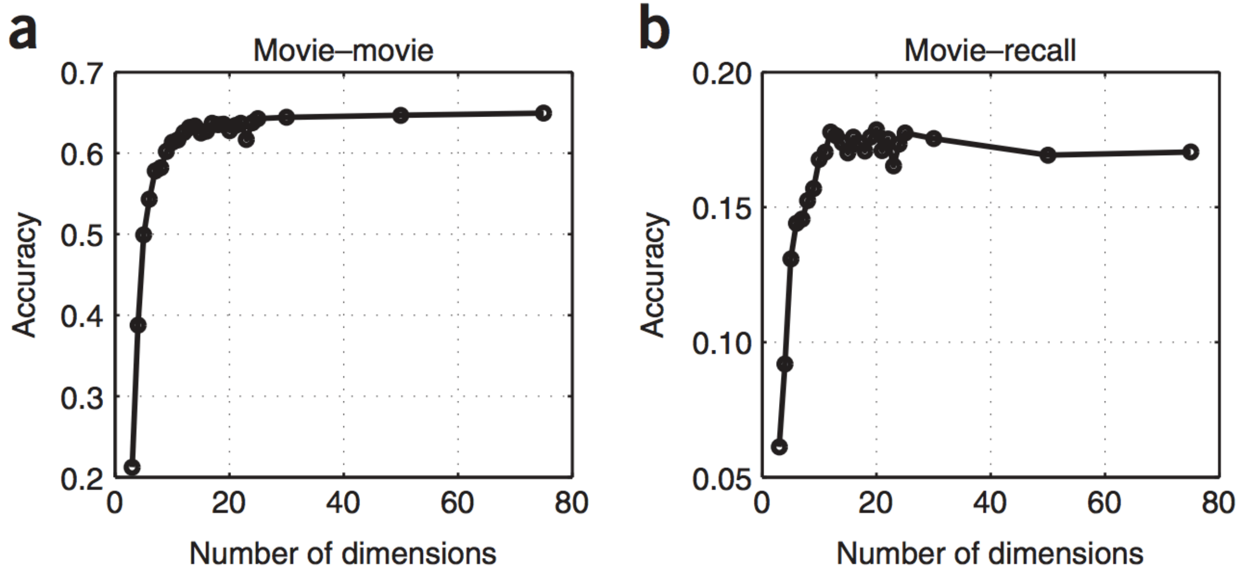

[J. Chen et al., Nat. Neur., 2017]

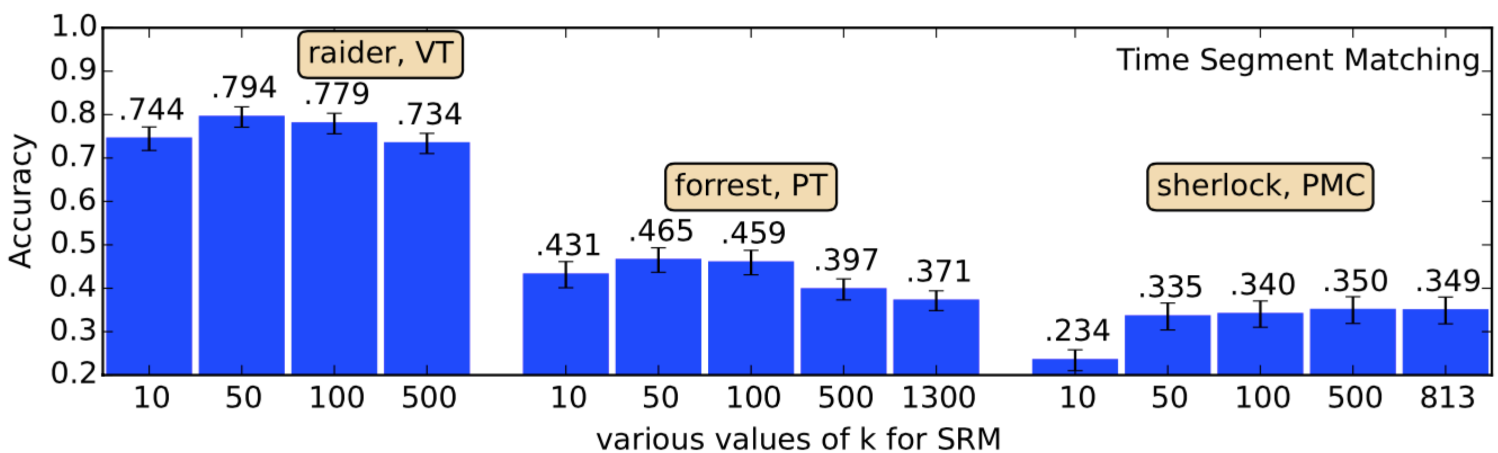

[P.-H. Chen et al. NIPS, 2015]

Quantifying dimensionality of shared response

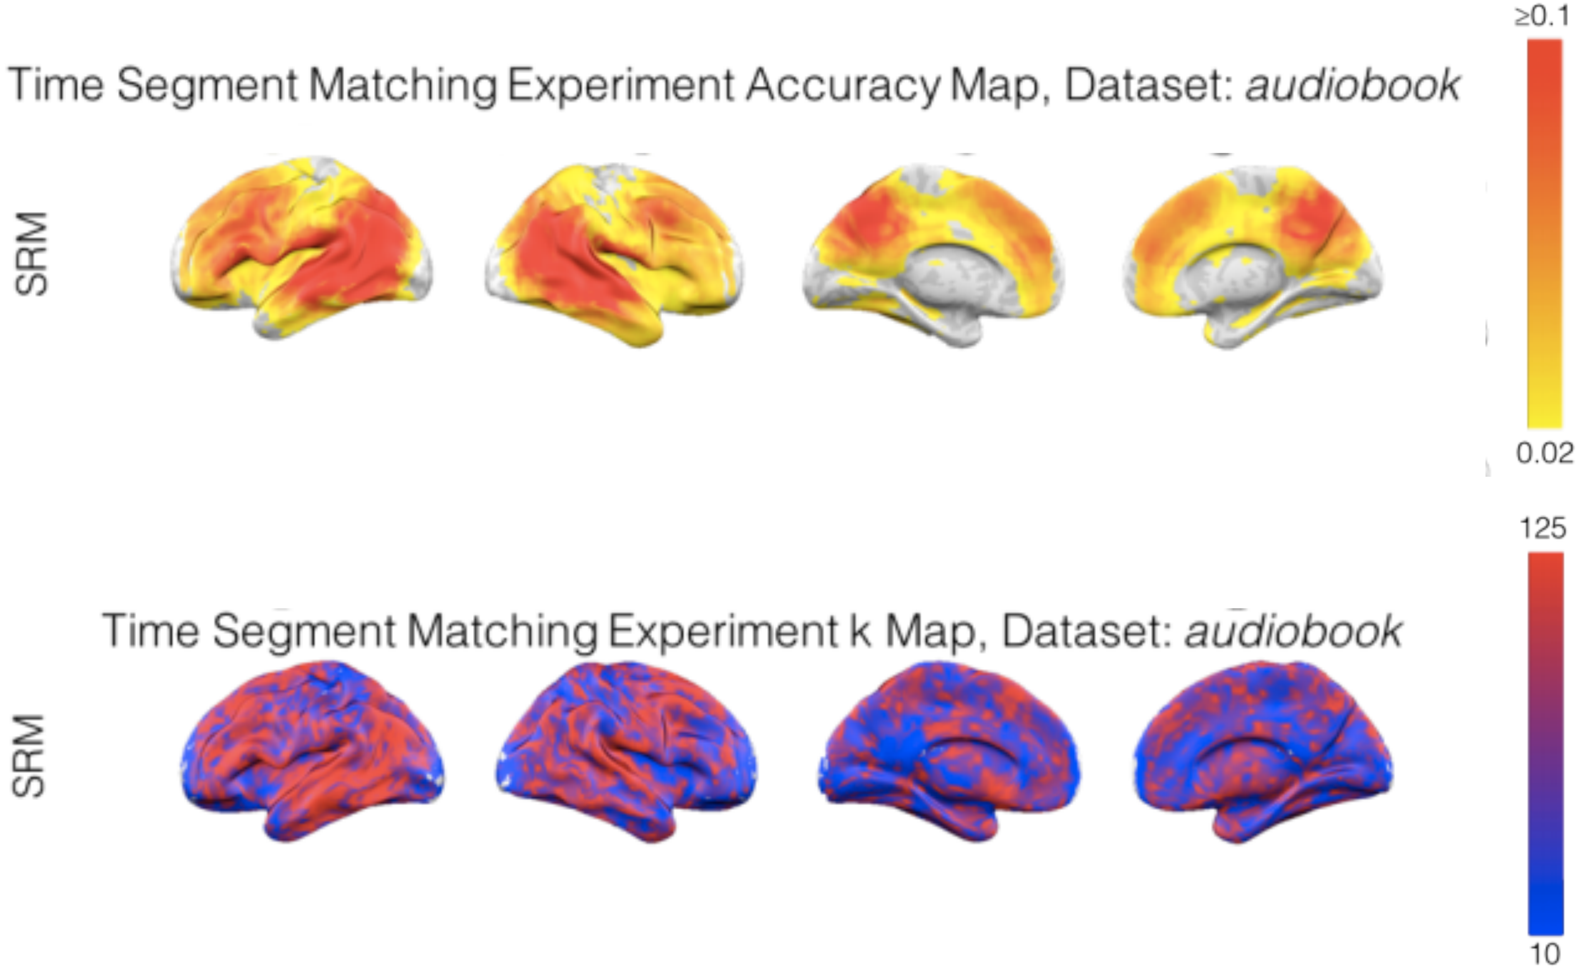

[H. Zhang et al. ArXiv, 2016]

Amount of data required to train SRM

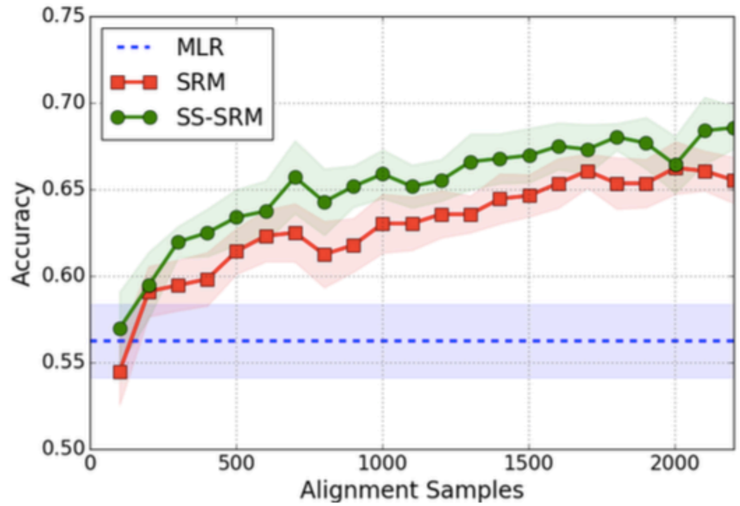

raider image category classification

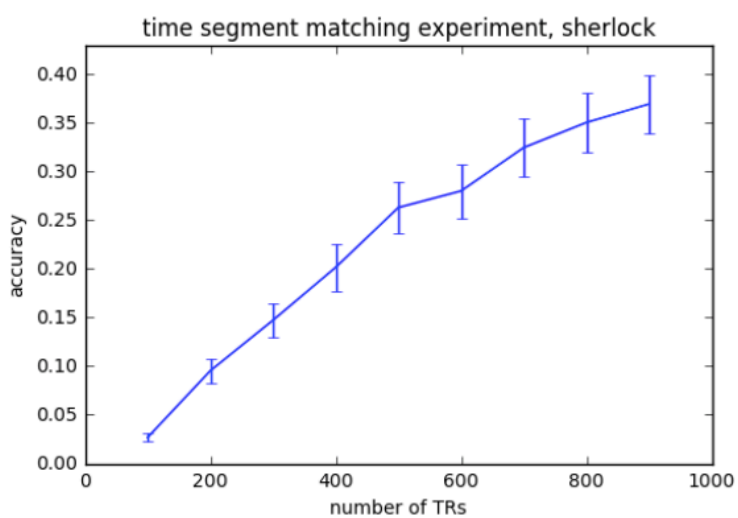

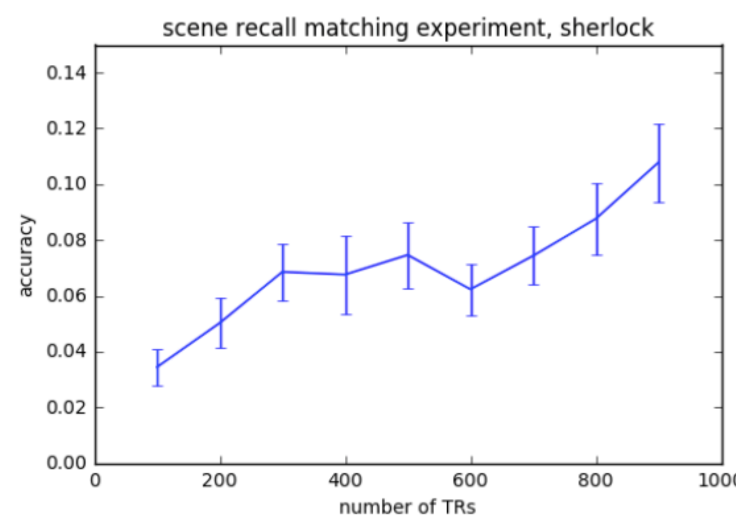

Amount of data required to train SRM

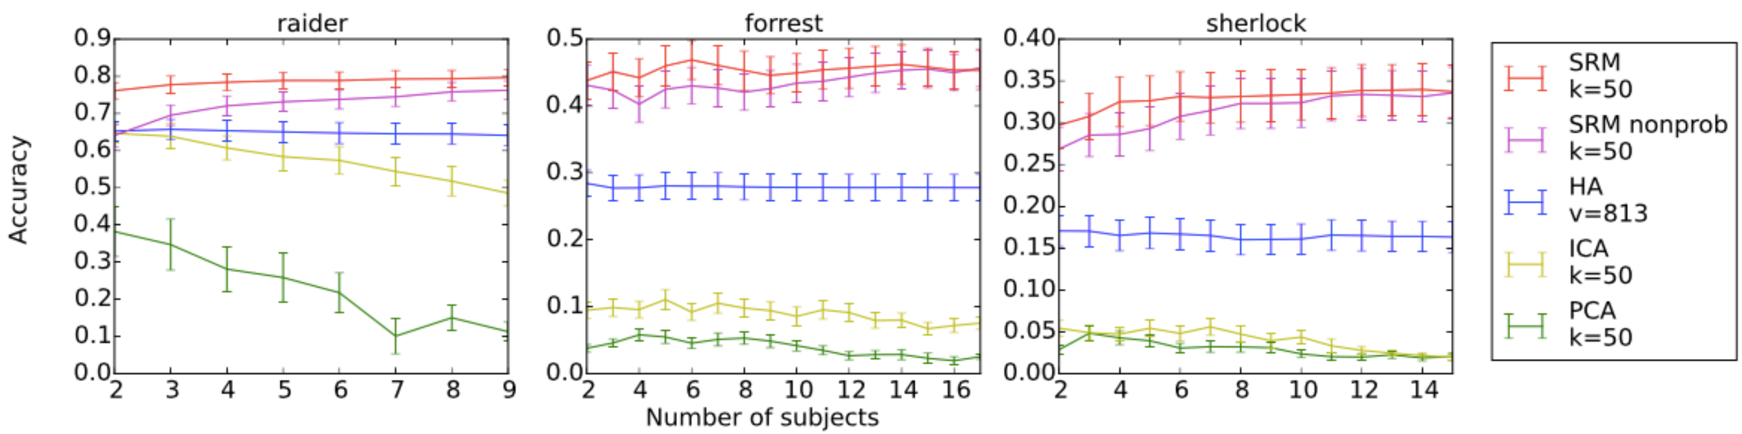

Figure 3.11: Effect of the number of subjects used in SRM training on the classification 18s time segments of a held out subject for three datasets and distinct ROIs. Error bars: ±1 stand. error.

A multi-subject convolutional autoencoder

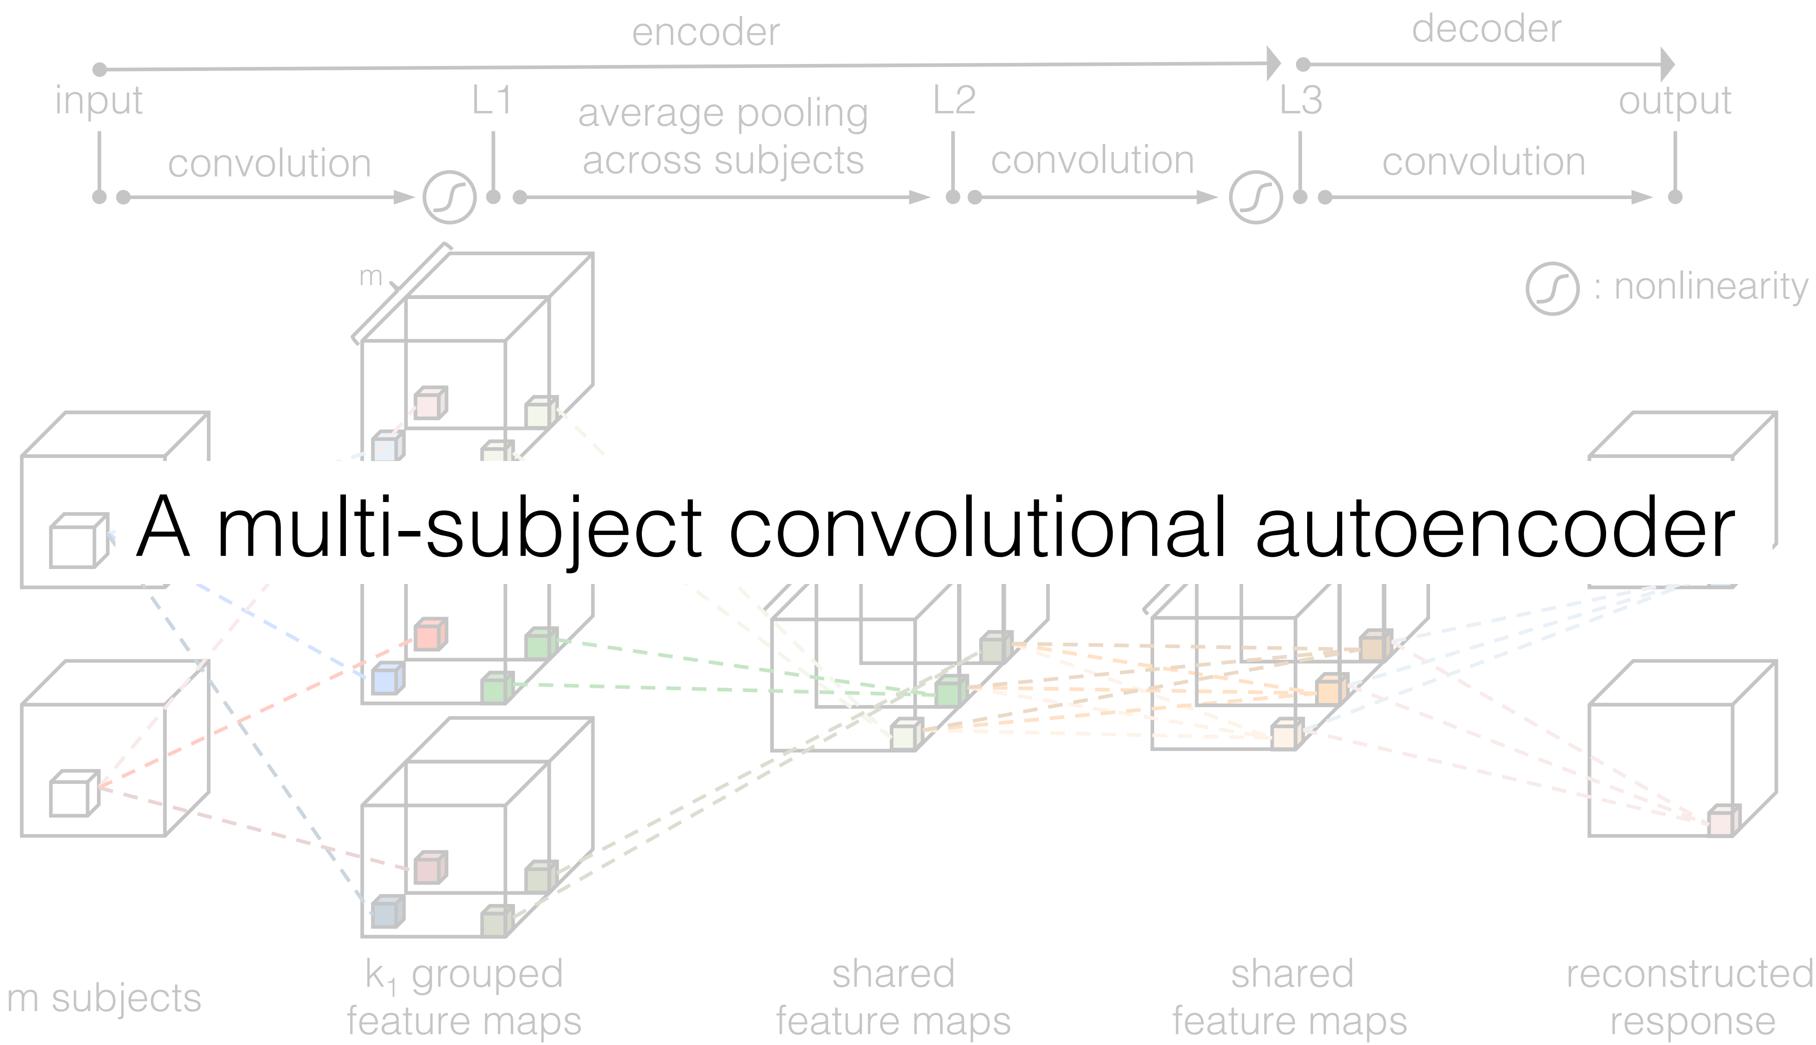

Dispersion of cross-subject mapping makes it hard to interpret the brain maps

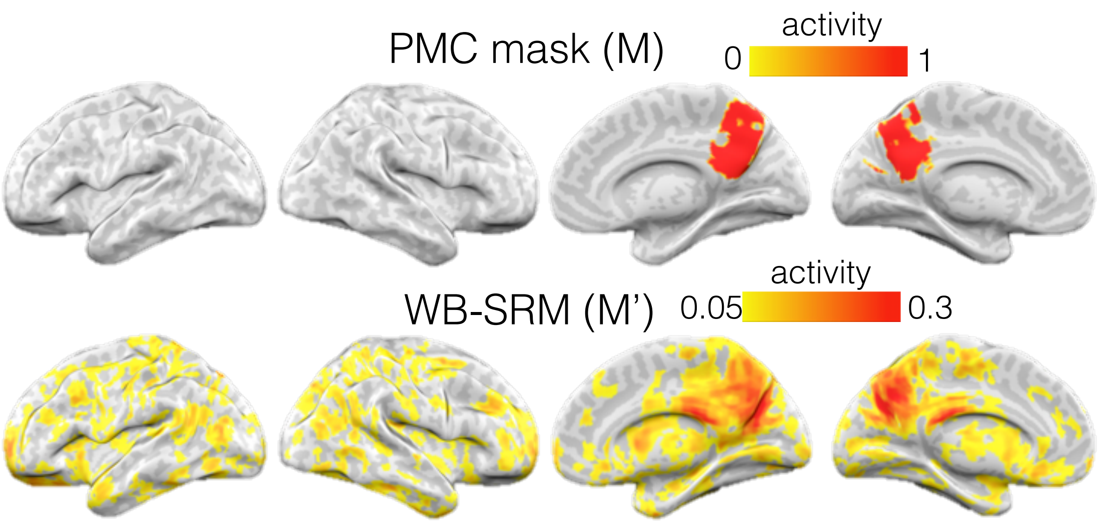

- Regularization
- Searchlight analysis

A multi-subject convolutional autoencoder (CAE)

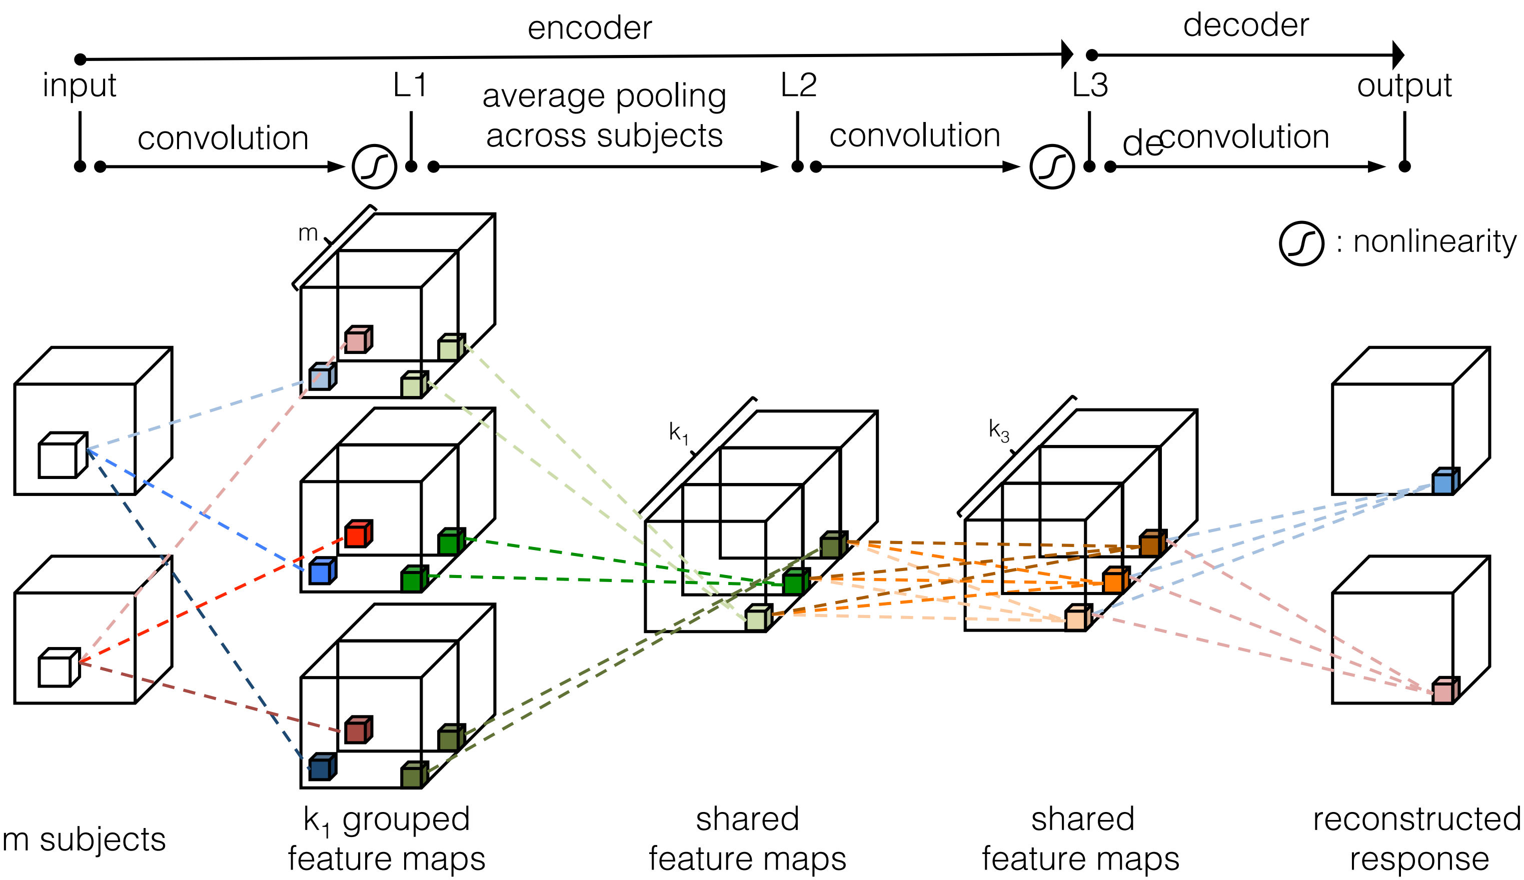

- Unified view of searchlight analysis and convolution operation
- Non-linear model for multi-subject fMRI data

Discovering information distribution in the brain

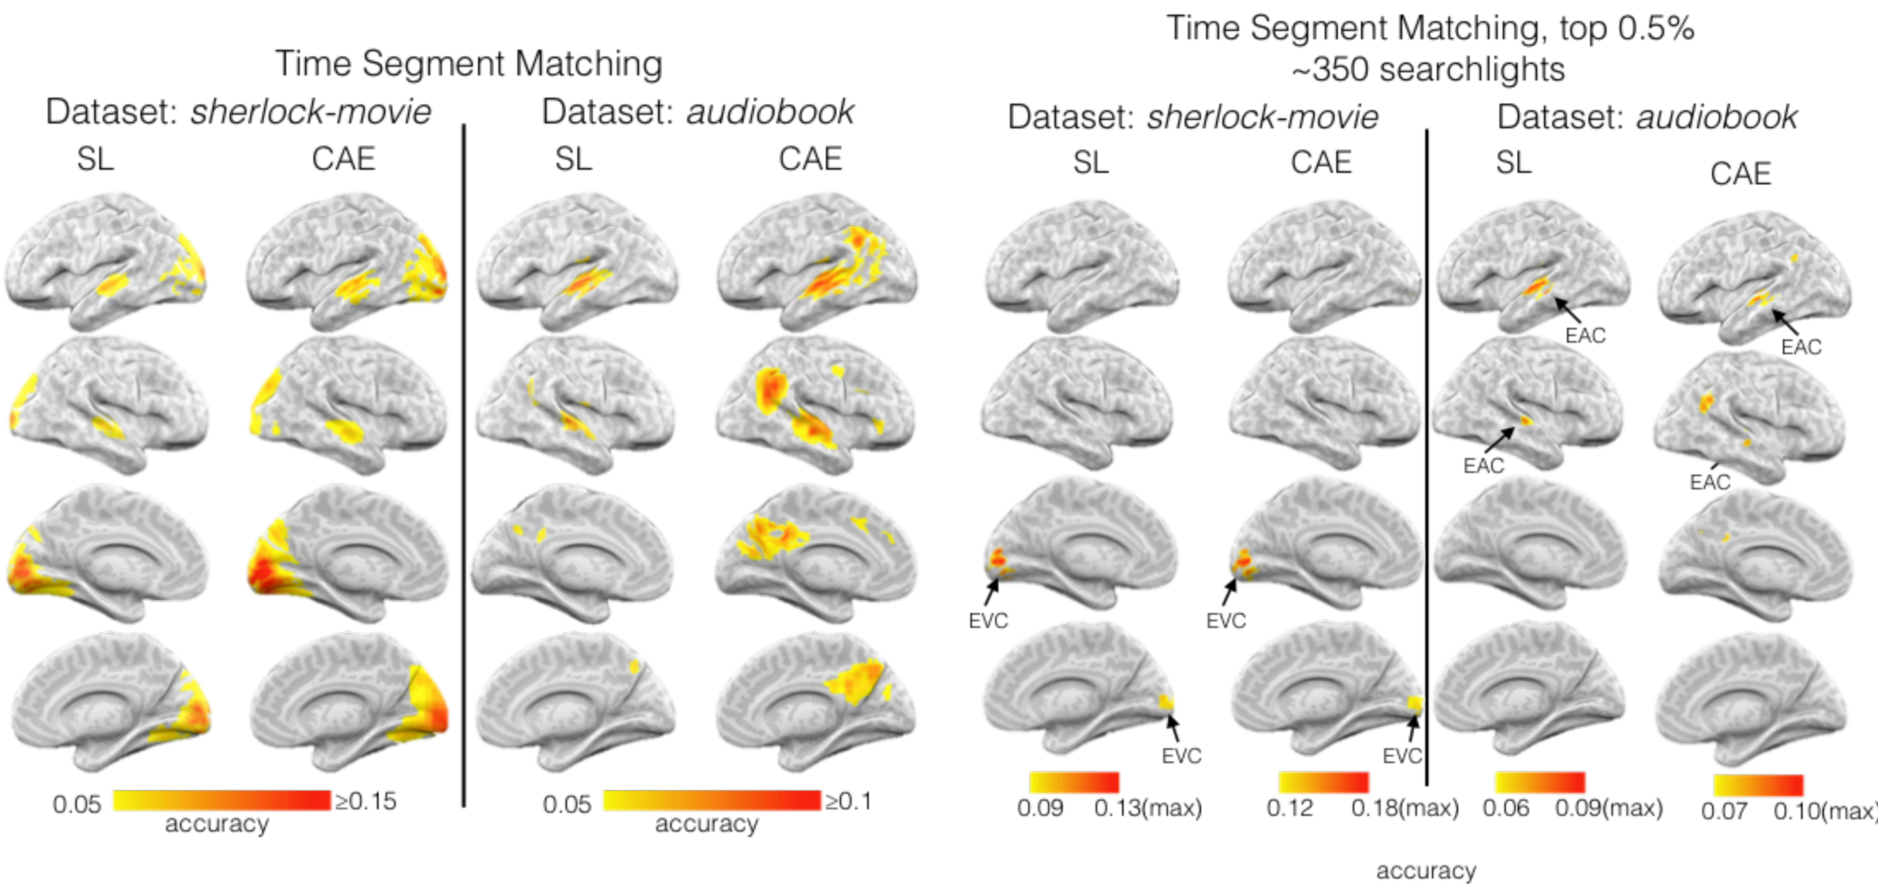

Demonstrating local information propagation with CAE

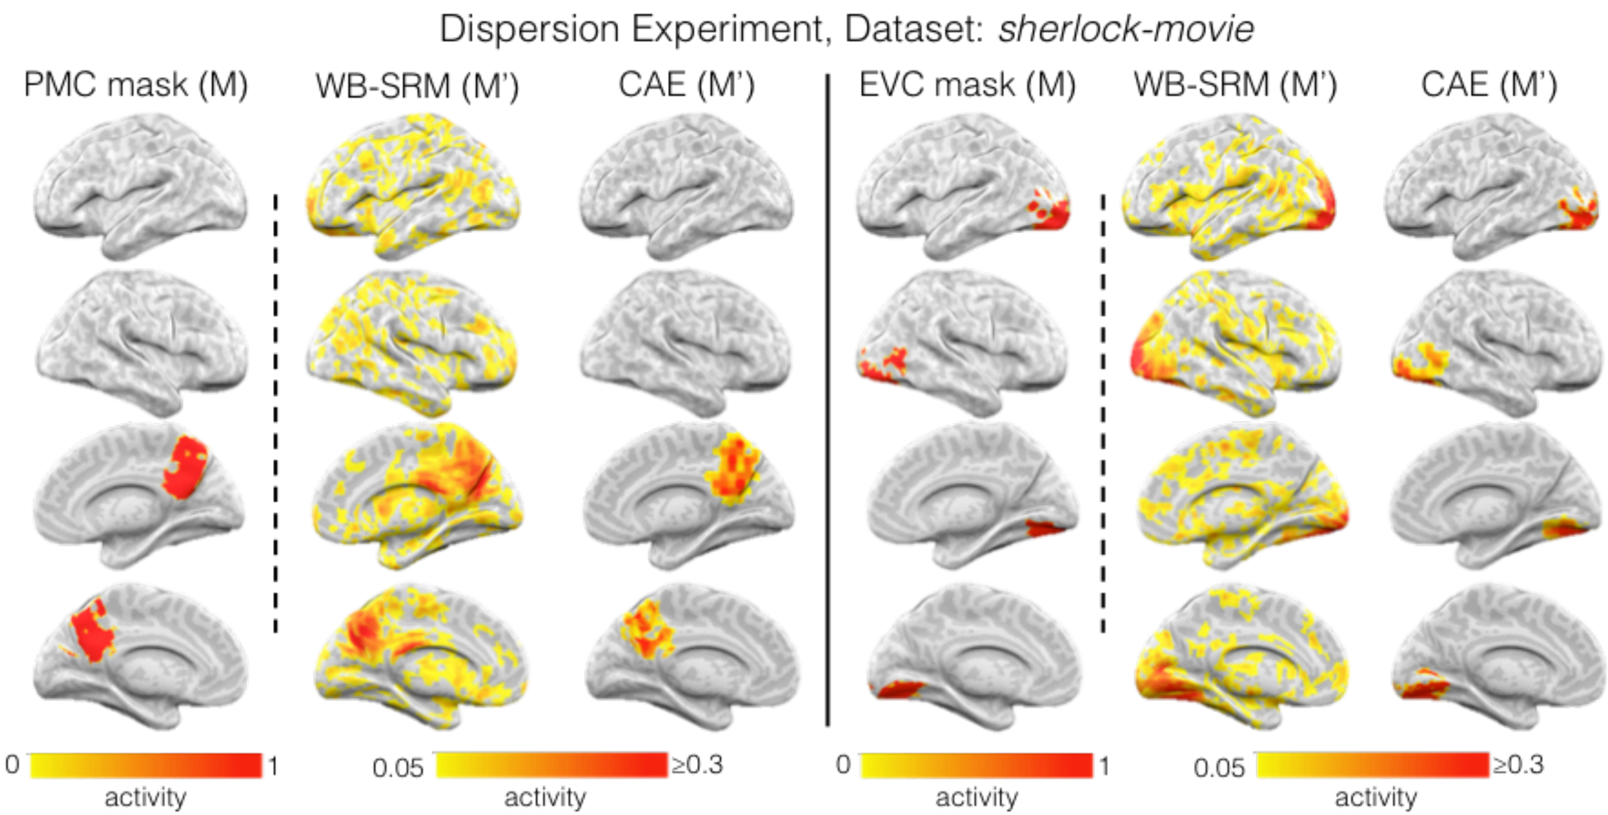

Transfer Learning on fMRI Datasets

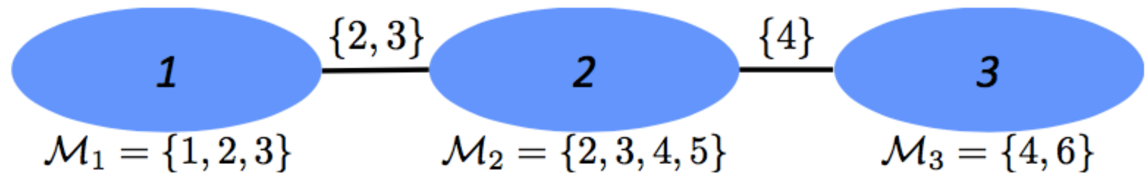

Figure 1: A simple dataset graph. Nodes represent datasets, edges indicate the presence of shared subjects, the edge labels indicate the set of indices of the shared subjects. $\mathcal{M}_d$ is the set of subject indices in dataset d.

Dataset	Type	Samples	Num. Subjs
greeneyes [23]	Audio	450 TRs	40
milky [24]	Audio	297 TRs	18
vodka [24]	Audio	297 TRs	18
schema [25]	Audio	937 TRs	31
sherlock [26]	Movie	1973 TRs	16
sherlock-recall [26]	Recall	34 scenes	16

Table 1: Information on fMRI datasets. Each TR is 1.5 seconds. Each scene is the averaged response when recalling the scene.

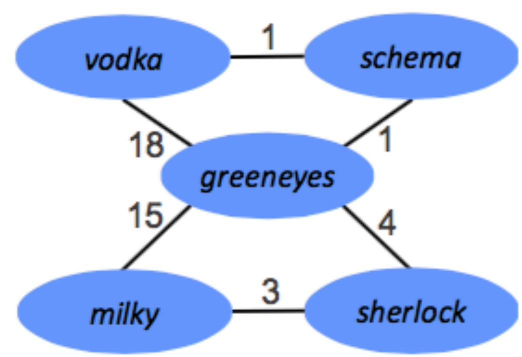

Figure 3: Structure of datasets as a graph. Num. shared subjects labeled on edges

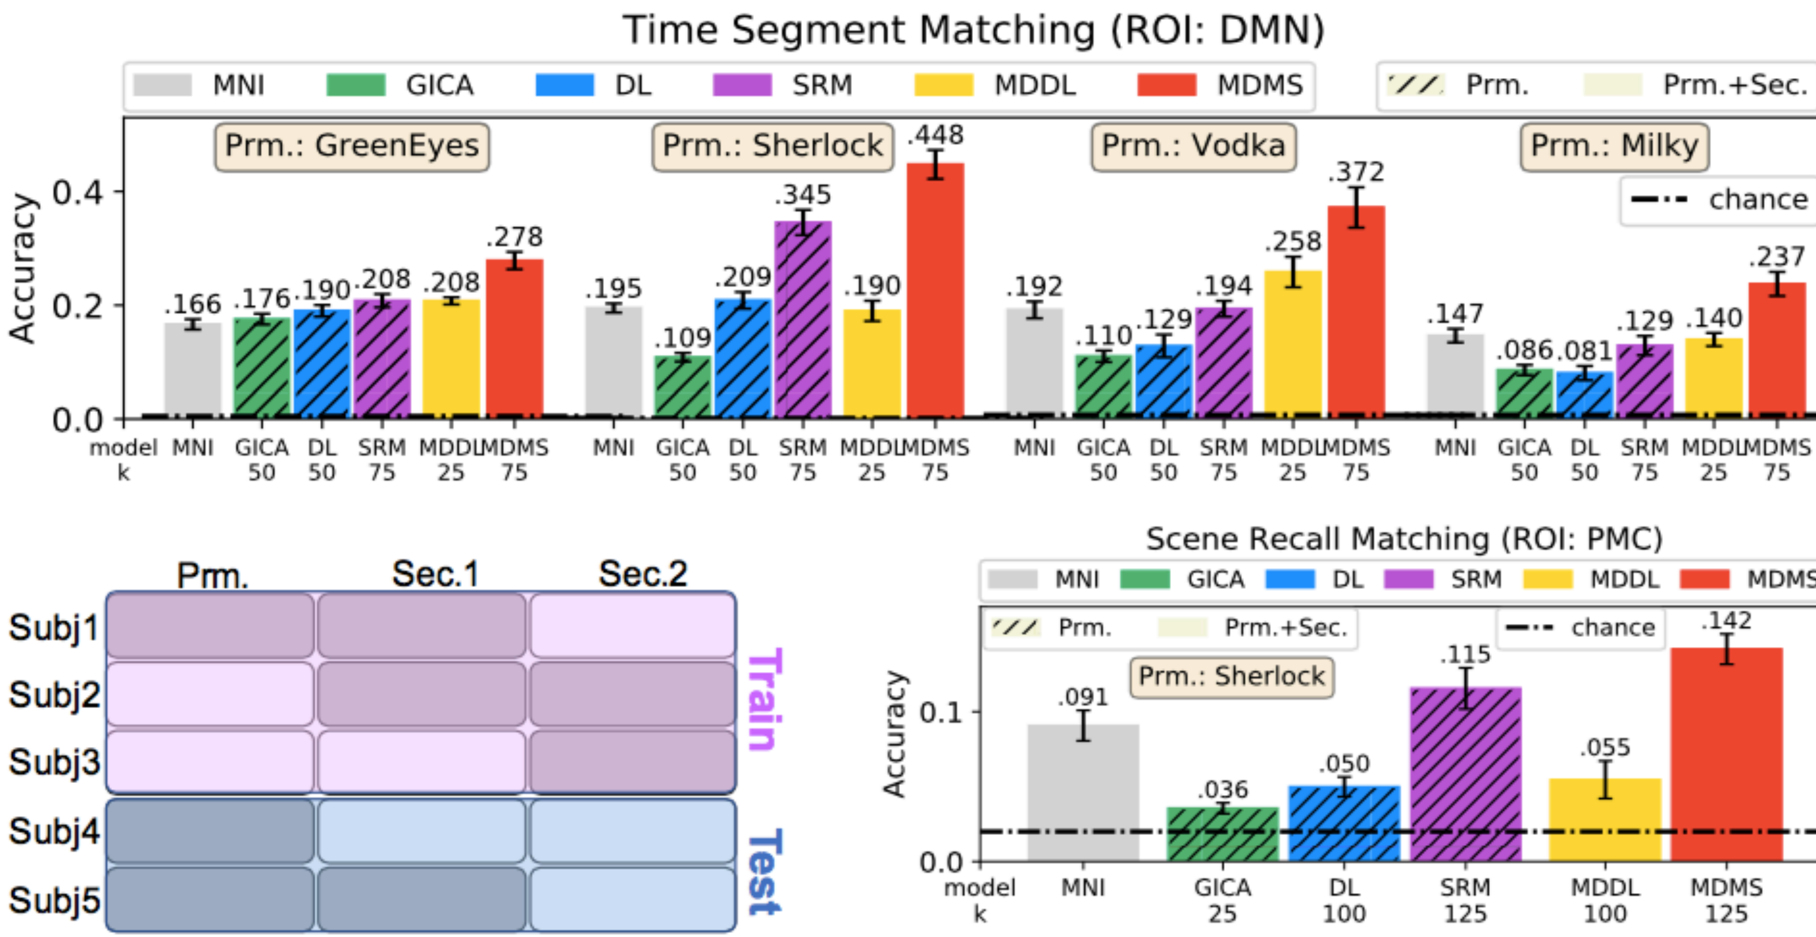

Figure 5: Experiment 1. **Top**: Results of time segment matching experiment on DMN (other ROIs in Supp. Mat.). Chance accuracy: *greeneyes*: 0.005; *sherlock*: 0.001; *vodka*: 0.008; *milky*: 0.008. k selected based on cross-validation. **Bottom (left)**: An example of random partition of training and testing subjects. Available observations are grey blocks, missing observations are white blocks. Testing subjects are completely left-out in all datasets. **Bottom (right)**: Scene recall matching on PMC. Each subject has data for 34 scenes on average, but there are 50 possible scenes (classes), so the chance accuracy is 0.02. k selected based on cross-validation.

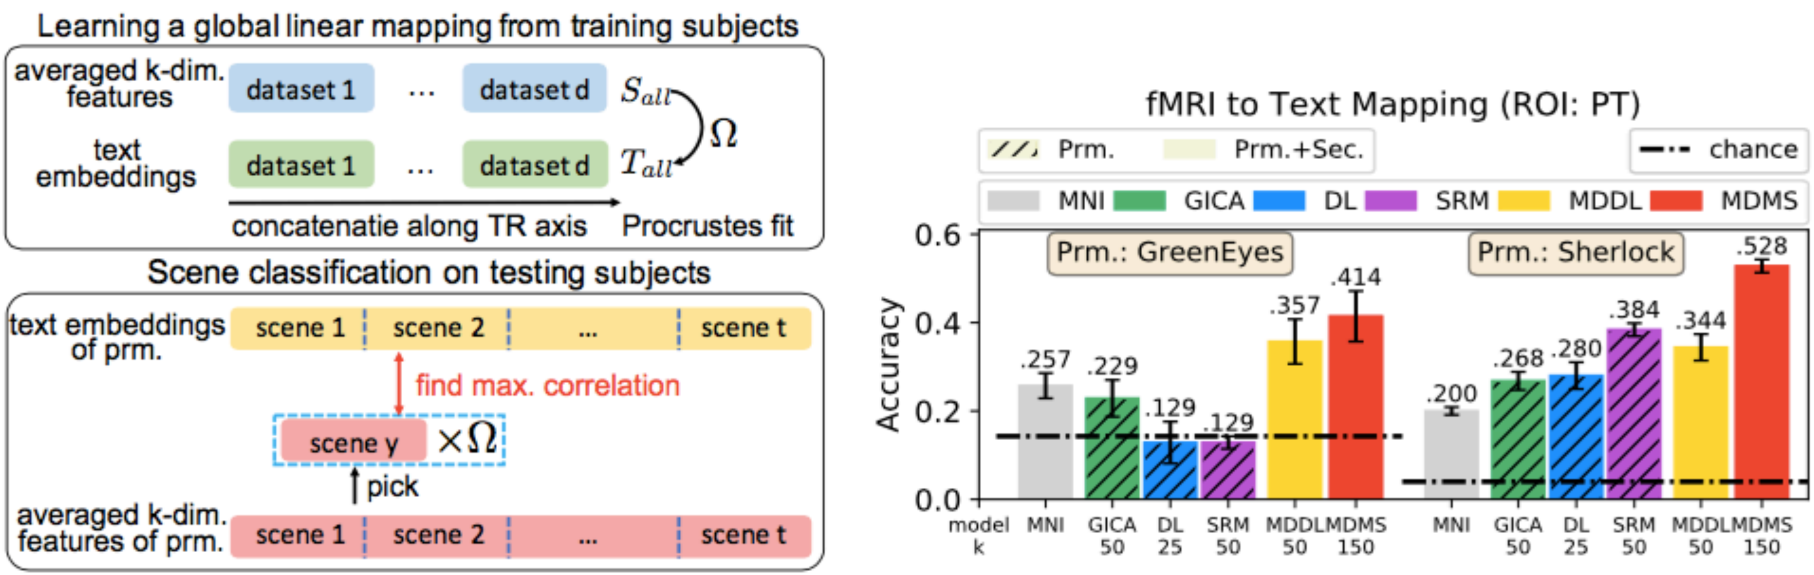

Figure 6: Experiment 2. **Left**: Learn linear transformation and perform scene classification. **Right**: fMRI data to text embedding transformation classification accuracy. Results on other ROIs in Supp. Mat. Chance accuracy: *greeneyes*:0.14; *sherlock*: 0.04. k selected based on cross-validation.

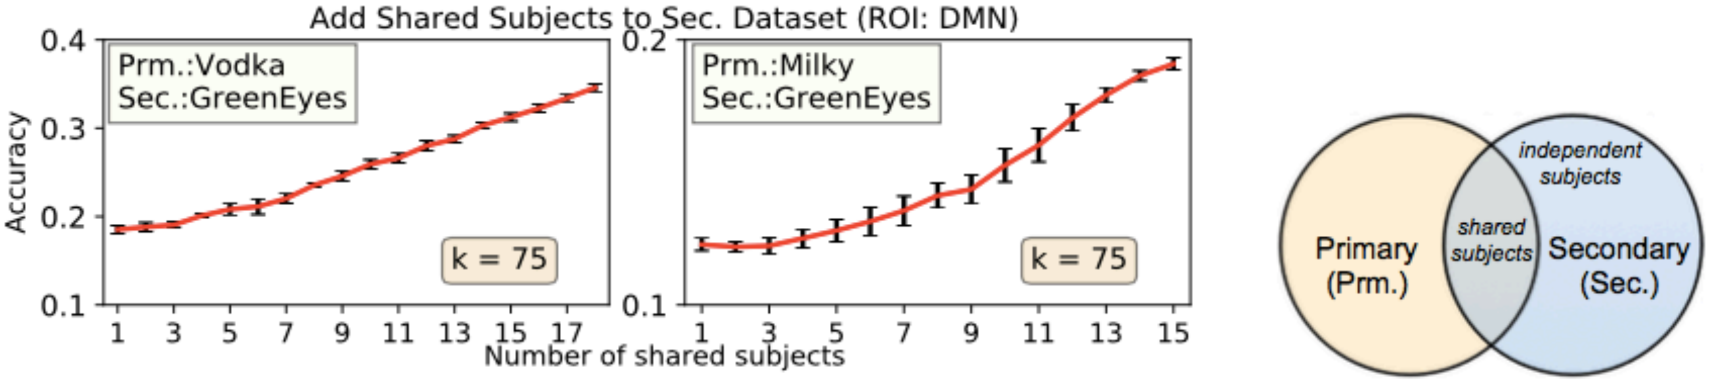

Figure 8: Experiment 4. **Left**: Time segment matching accuracy on prm. dataset when using all independent subjects and different number of shared subjects in sec. dataset. Chance accuracy and k same as experiment 1. Error bar computed across subjects. **Right**: Definition of shared subjects and independent subjects.

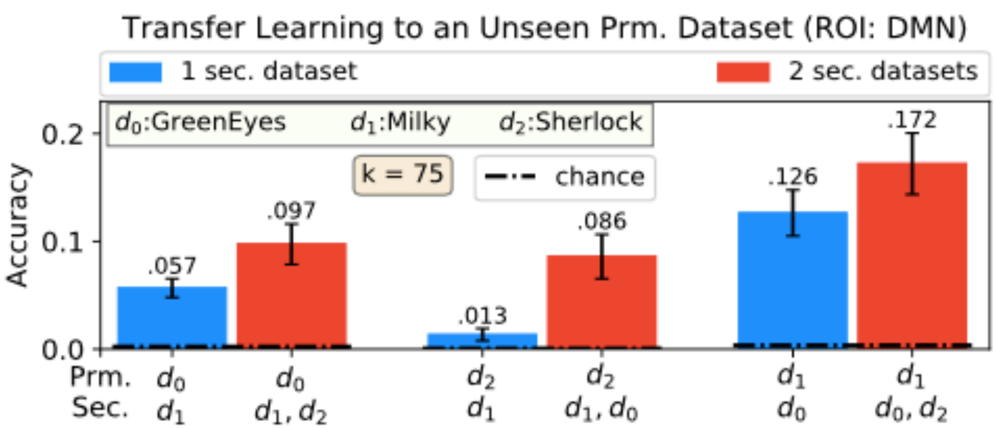

Figure 7: Experiment 3. Time segment matching accuracy on prm. dataset using subject specific basis learned from 1 or 2 secondary datasets. Results on other ROIs in Supp. Mat. Chance accuracy: *greeneyes*: 0.0025; *milky*: 0.004; *sherlock*" 0.0005. k same as experiment 1. Error bar computed across subjects.

Jupyter notebook examples

Need jupyter notebook and brainiak properly installed with python 3

1. git clone https://github.com/cameronphchen/SRM_tutorial.git
2. cd SRM_tutorial
3. chmod +x download-data.sh
4. ./download-data.sh
5. jupyter notebook

Code ready to use on your dataset

https://github.com/IntelPNI/brainiak

- Simple setting, one line command to fit SRM on your data
- Handles different numbers of voxels across subjects/views

SRM on fMRI

1. Generalize to new stimulus
2. Generalize to new subject
3. Decoupling shared and individual response
4. SRM with retinotopy
5. Searchlight SRM
6. Bridging shared space and word embedding space

SRM on fMRI

1. Generalize to new subject

2. Generalize to new stimulus

3. Decoupling shared and individual response

4. SRM with retinotopy

5. Searchlight SRM

6. Bridging shared space and word embedding space

A man, startling awake, sweating in his bed. A single bed in the dullest, plainest room. He sits up, calming himself, letting his breathing return to normal.

Why searchlights?

Structured Sparsity

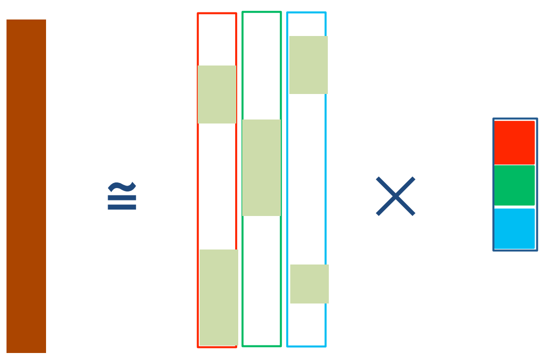

t-th column of X_i W_i t-th column of S

Searchlight

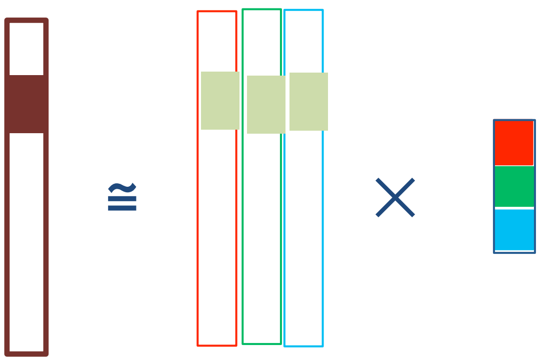

t-th column of X_i W_i t-th column of S

Semi-supervised SRM

Dataset	Experiment	MLR	SRM	SS-SRM
raider	Image category	56.25%	65.53%	68.57%
sherlock	Scene recall	4.28%	5.31%	6.12%

Table 1. Comparison of average accuracy for brain decoding experiments.

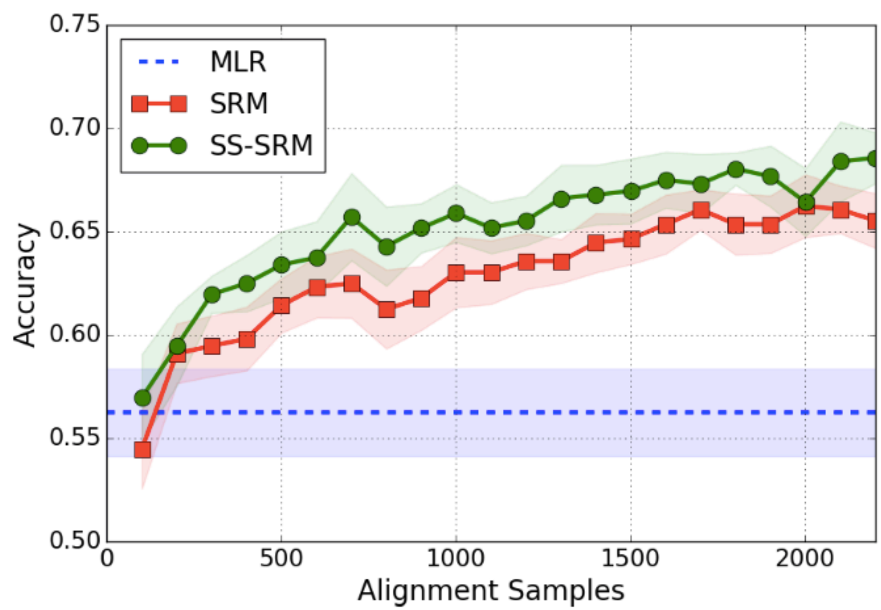

Fig. 1. Average accuracy as a function of the number of alignment samples.