

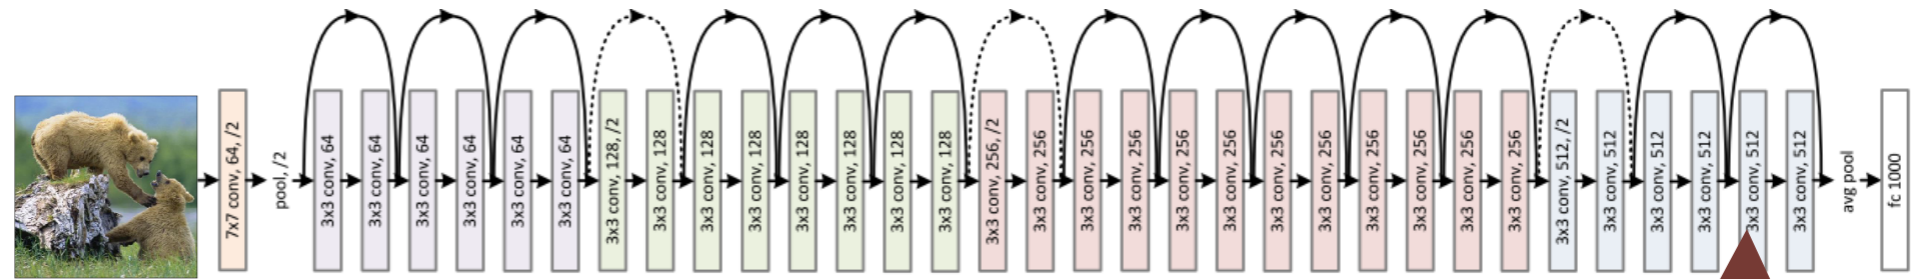
Factorized visual representations in the primate visual system and deep neural networks

Jack Lindsey
Columbia University



Work conducted with Elias Issa

Deep neural network models predict neural activity in visual brain regions

Task-Optimized Model

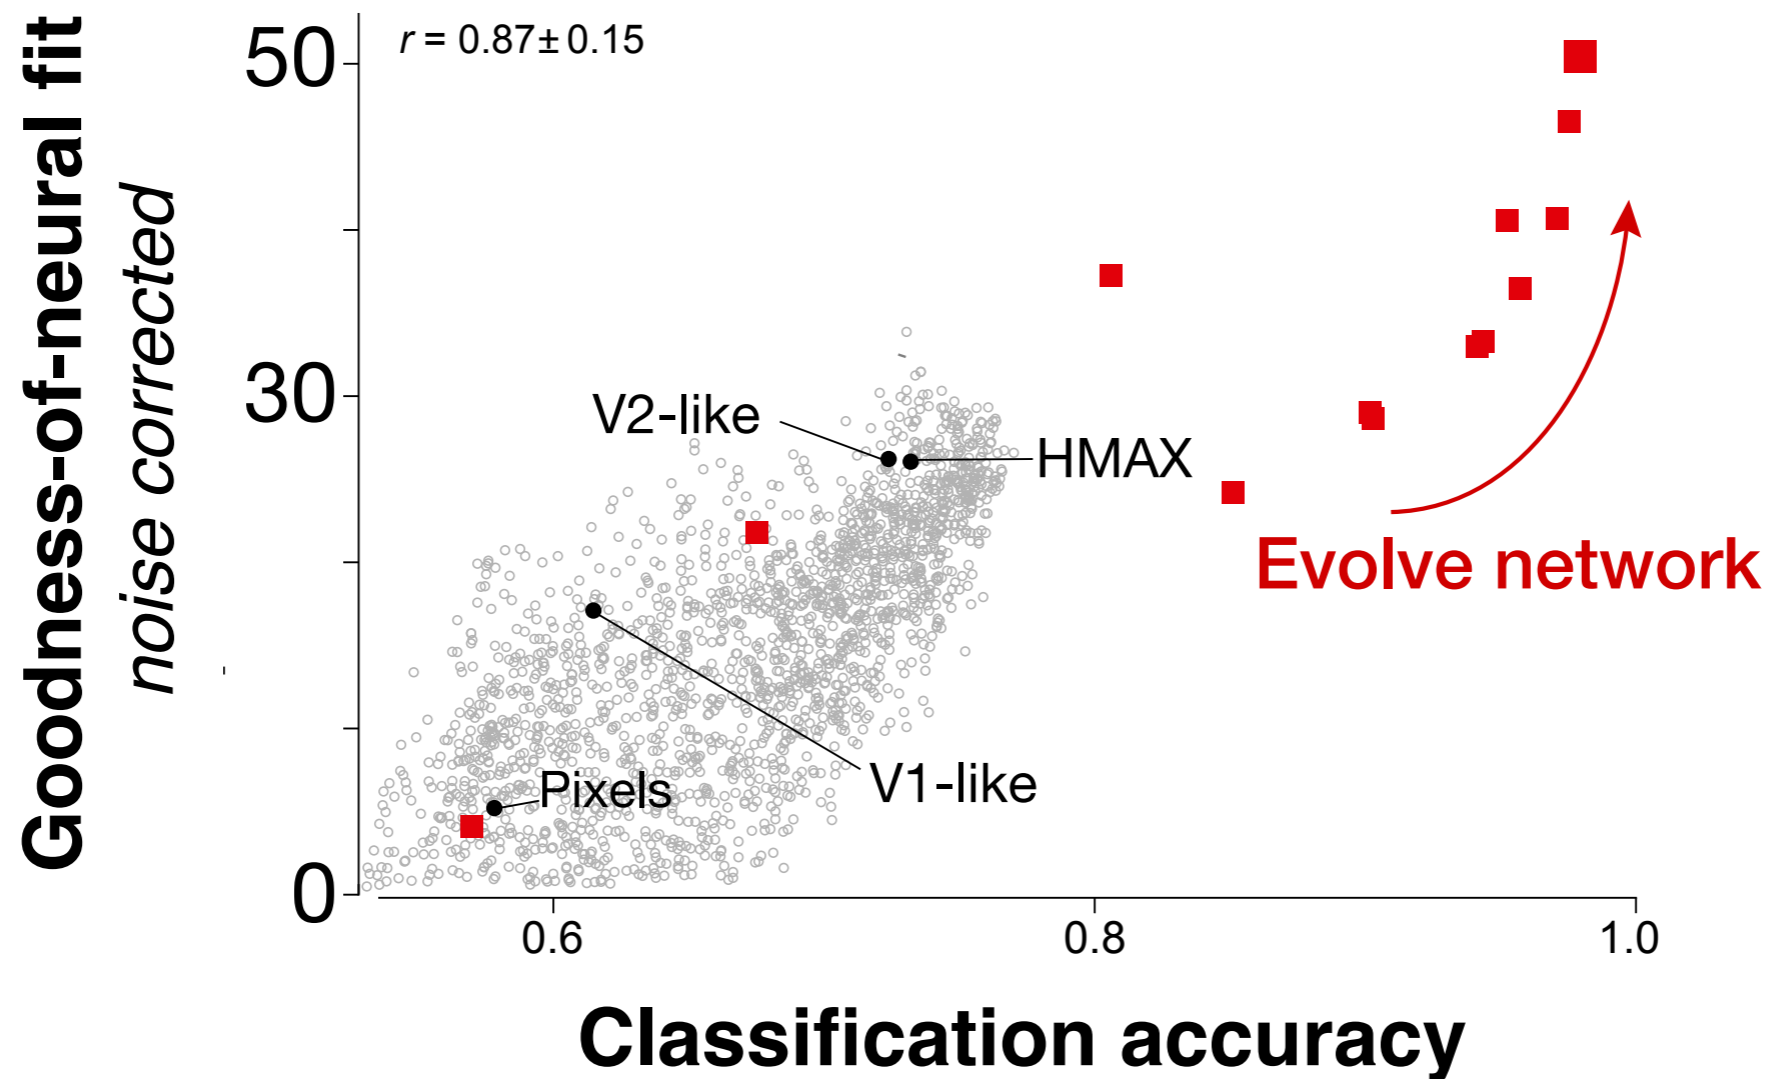


Visual Cortex Data

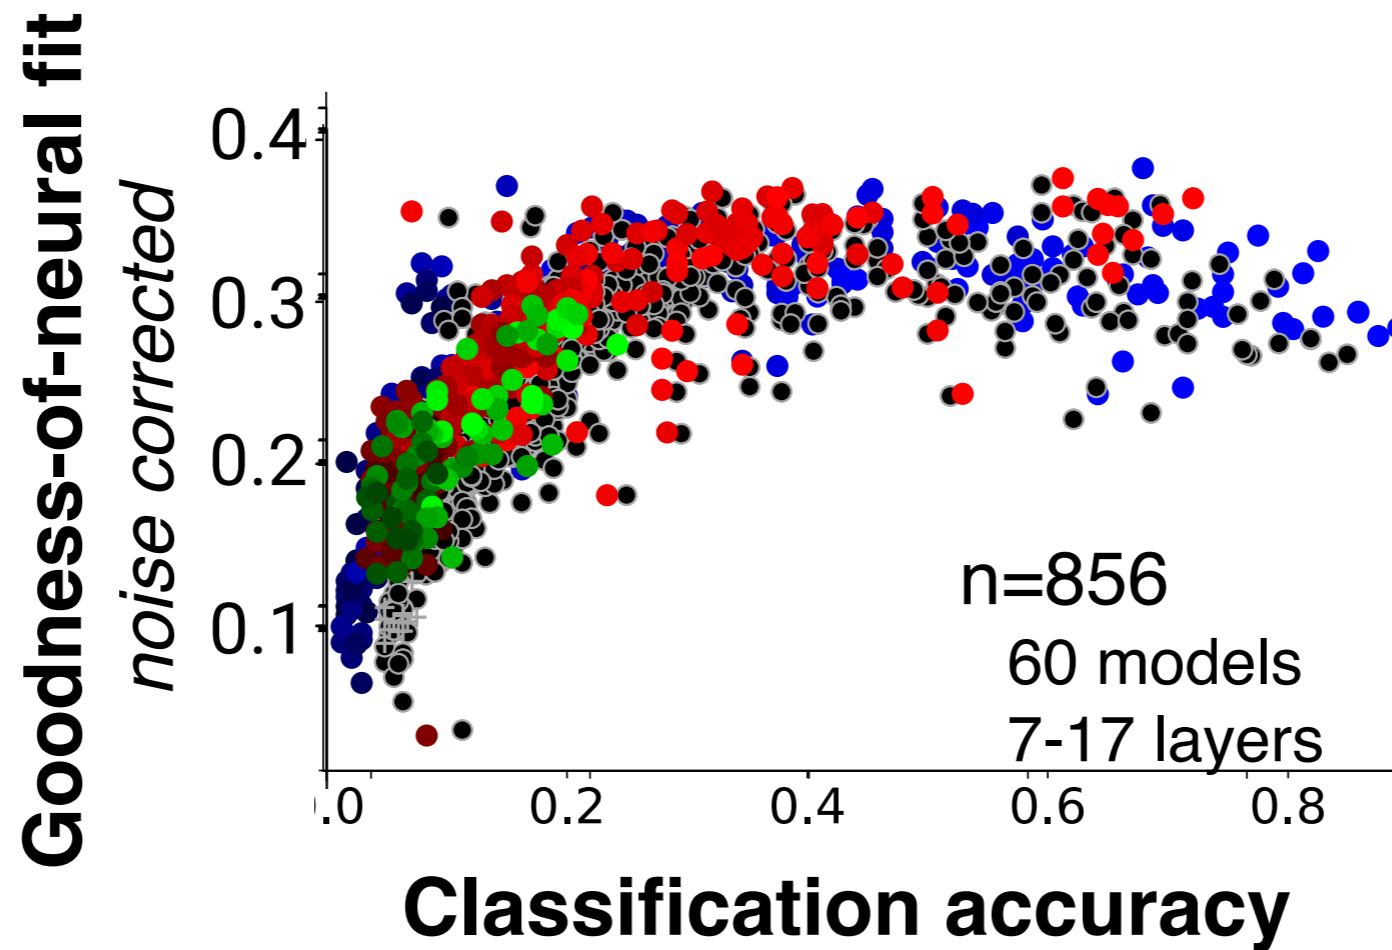


Compare high-level representations

How well does a model's image classification performance predict its fit to neural data?

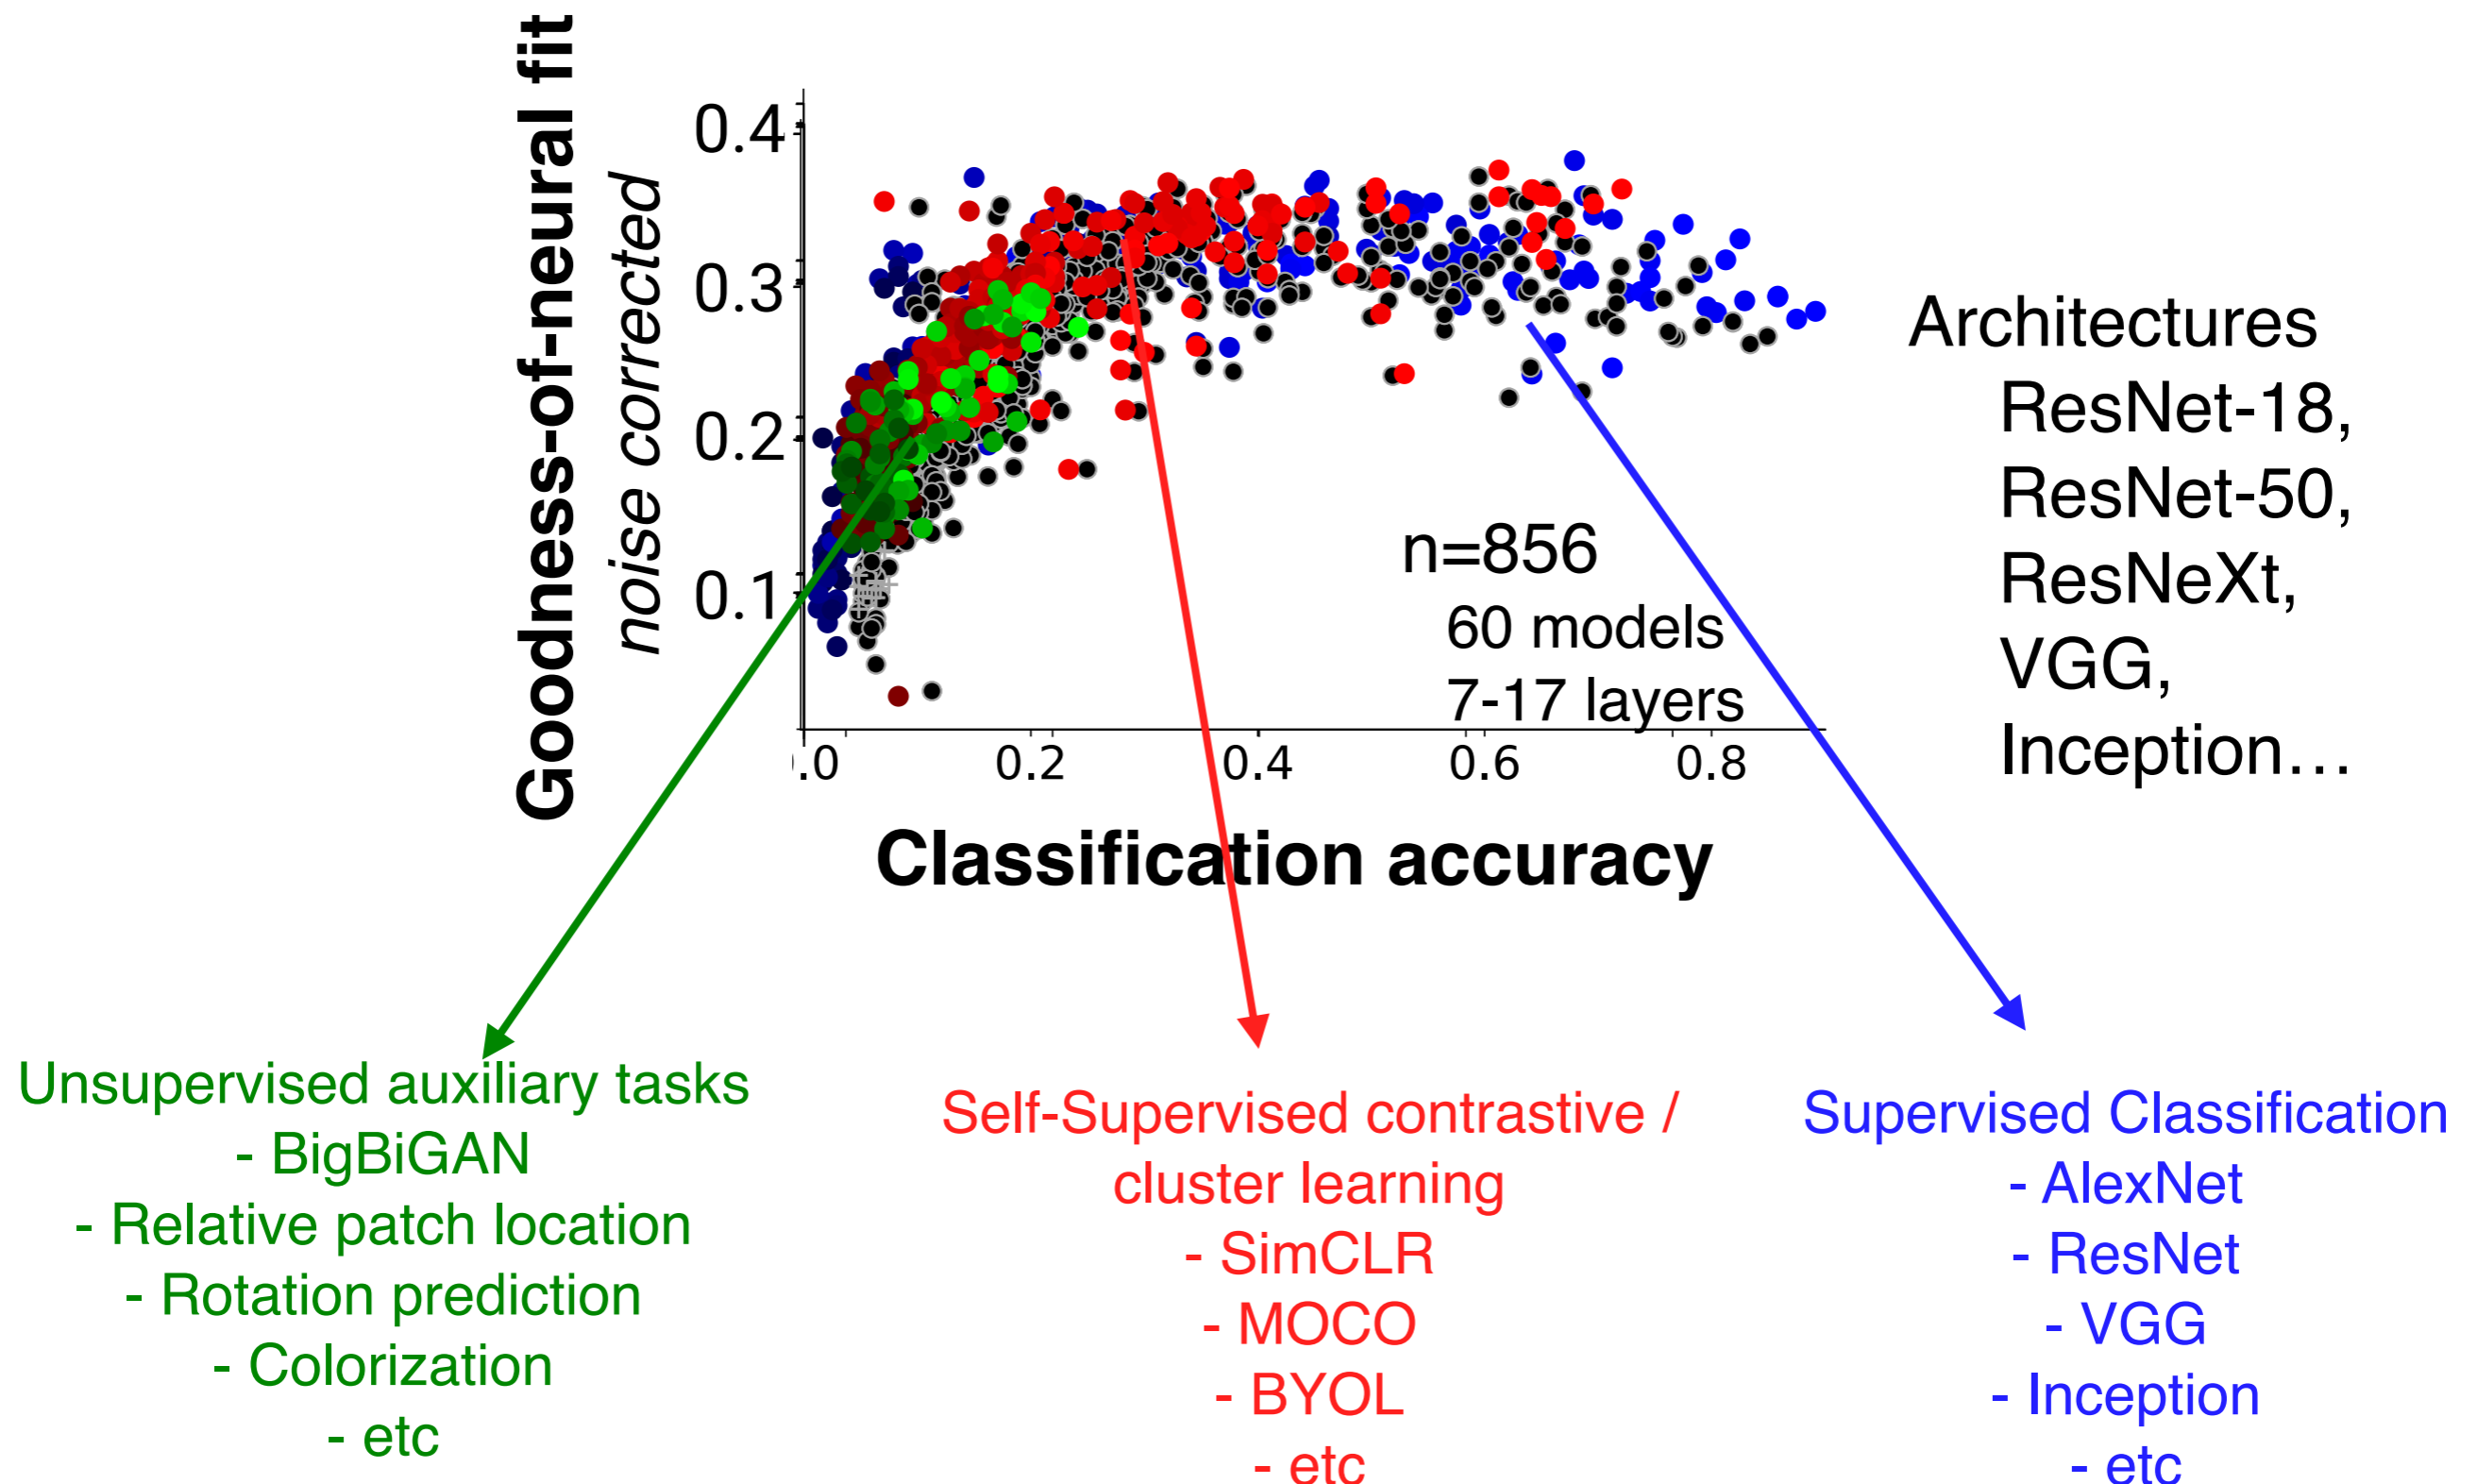


How well does a model's image classification performance predict its fit to neural data?



Architectures
ResNet-18,
ResNet-50,
ResNeXt,
VGG,
Inception...

How well does a model's image classification performance predict its fit to neural data?

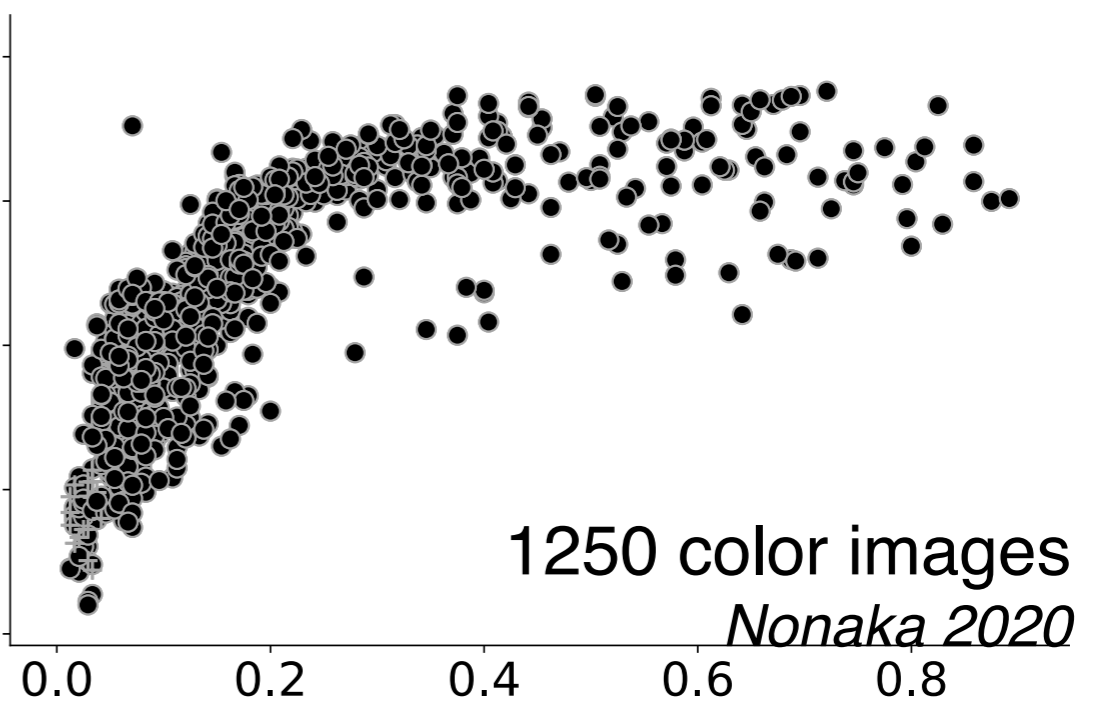
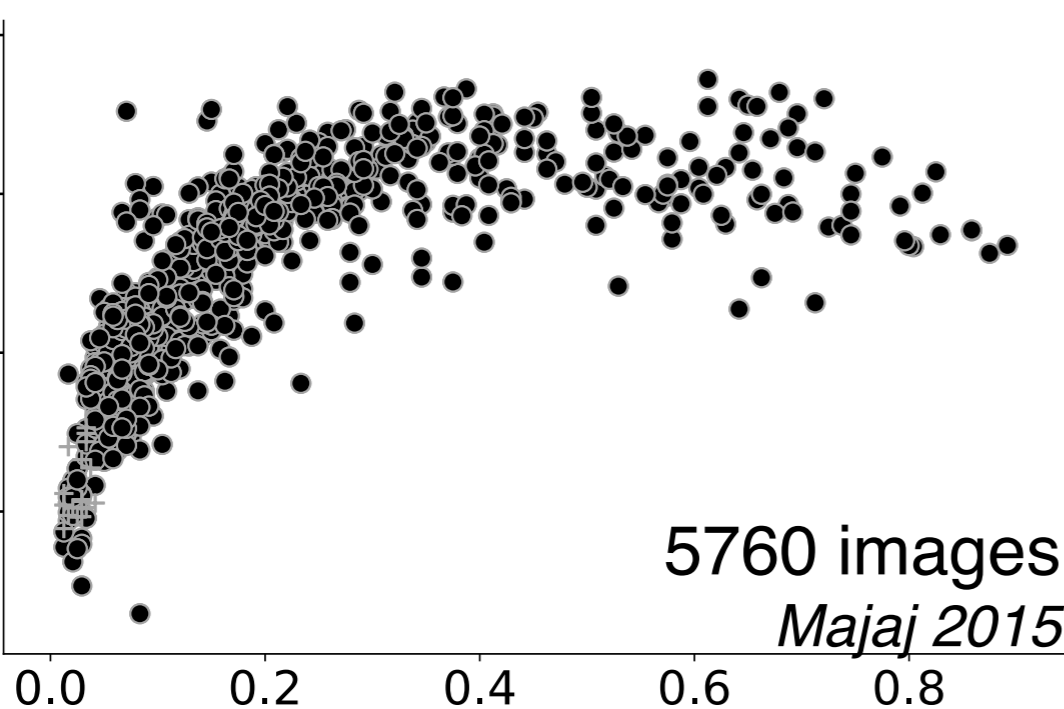

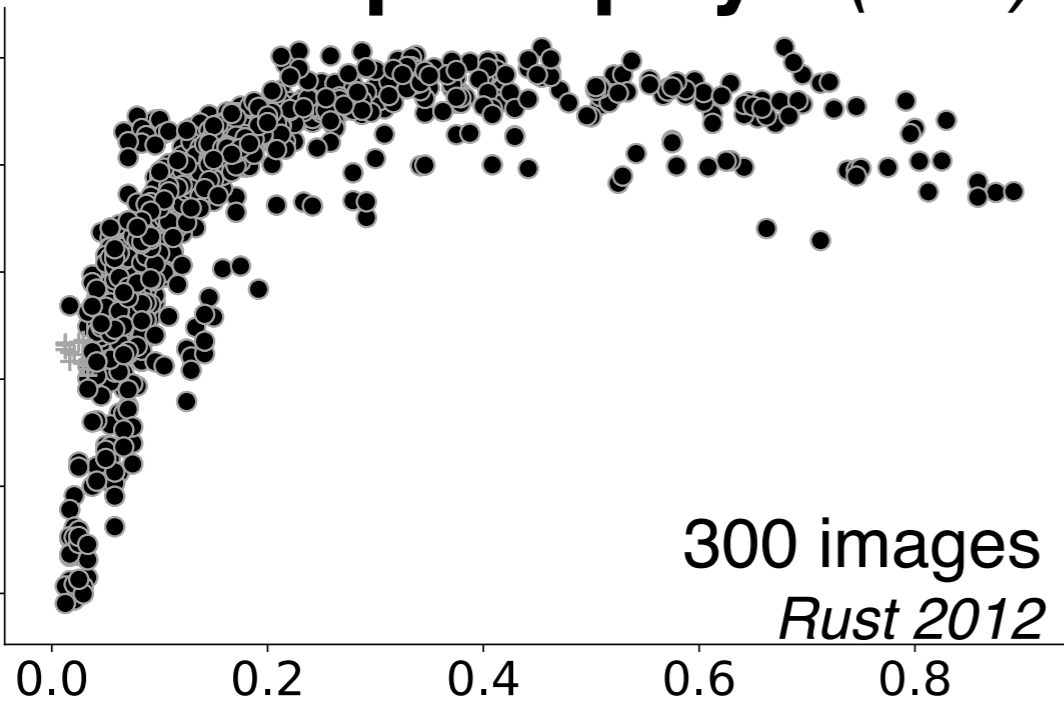


Trend holds across multiple datasets

Macaque Ephys (*ITC*)

Human fMRI (*HVC*)

Goodness-of-neural fit
noise corrected




Classification accuracy


Moving beyond classification

What additional properties of a model *make some better than others* at predicting activity in the primate visual system?

Conceptual gap between models & visual perception

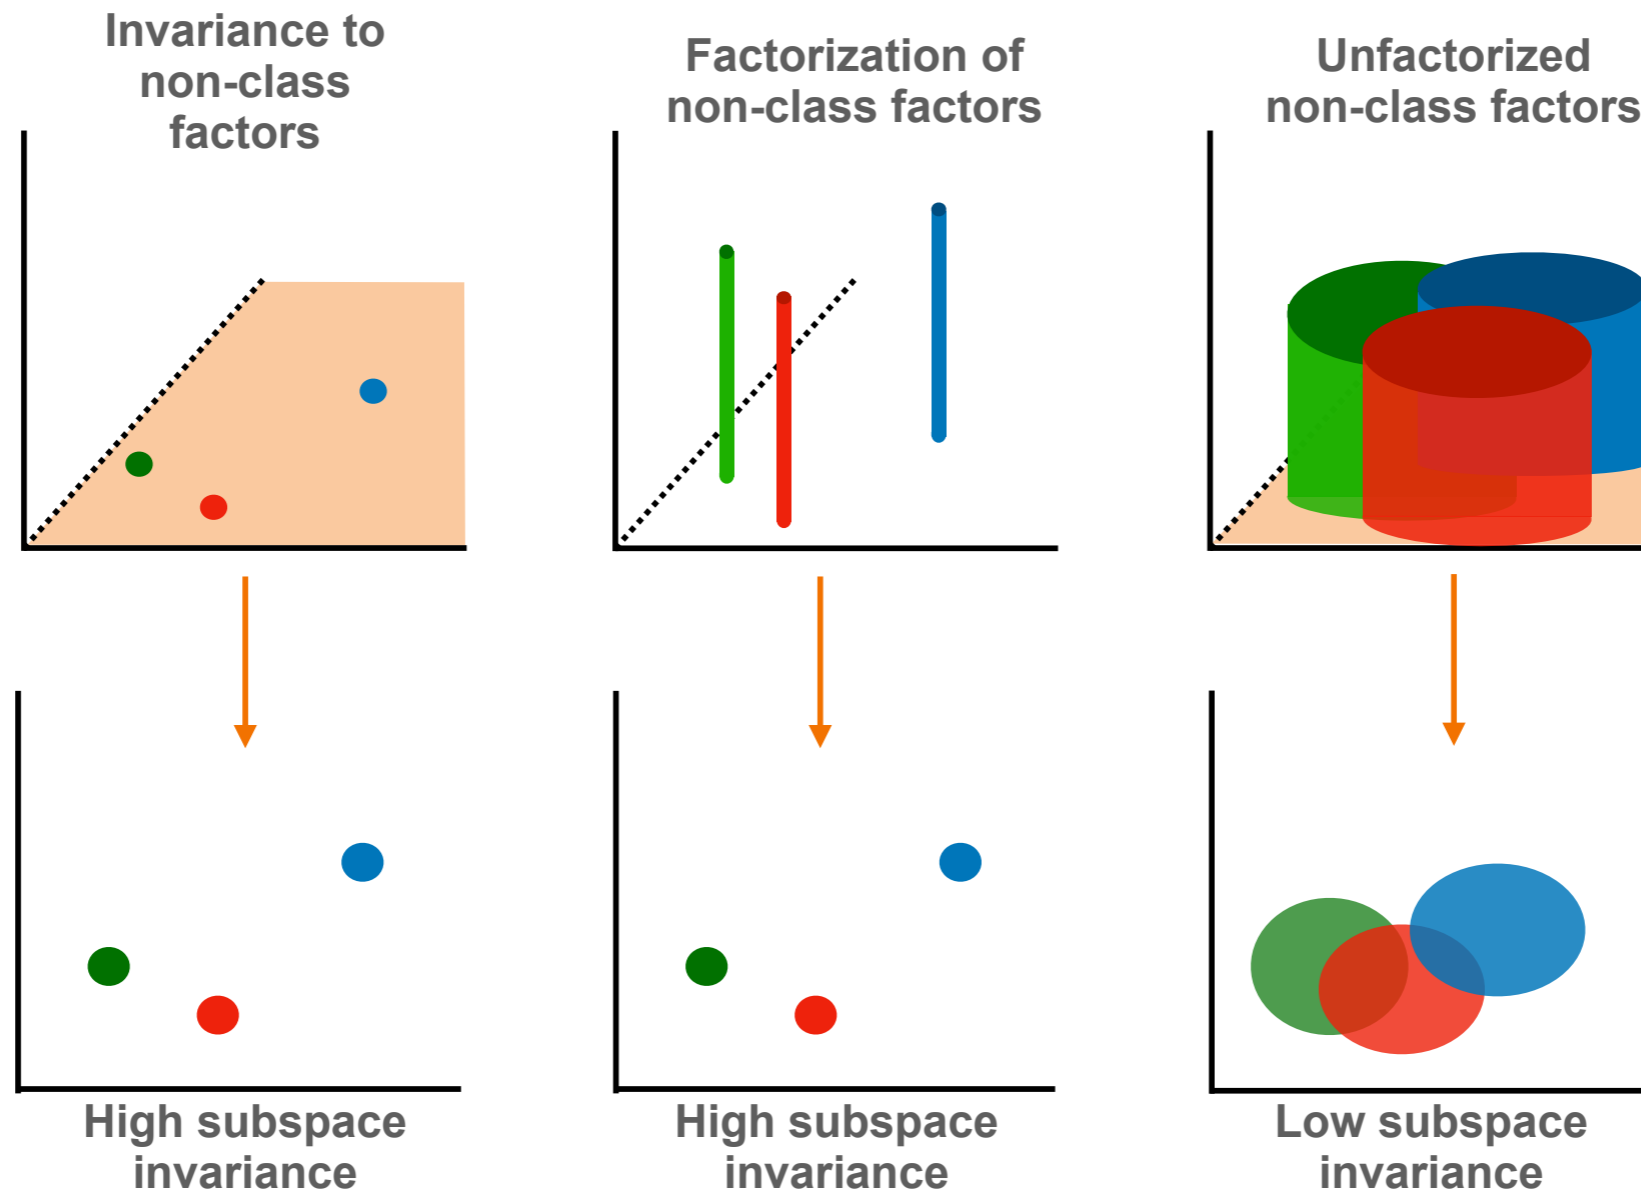


Most *models optimize for classification of object identity, irrespective of pose & scene composition.*



The brain supports *a rich understanding of many scene properties.*

Candidate property: factorization



Factorization measures the extent to which different parameters drive overlapping vs. orthogonal variance in neural activity space

Metrics

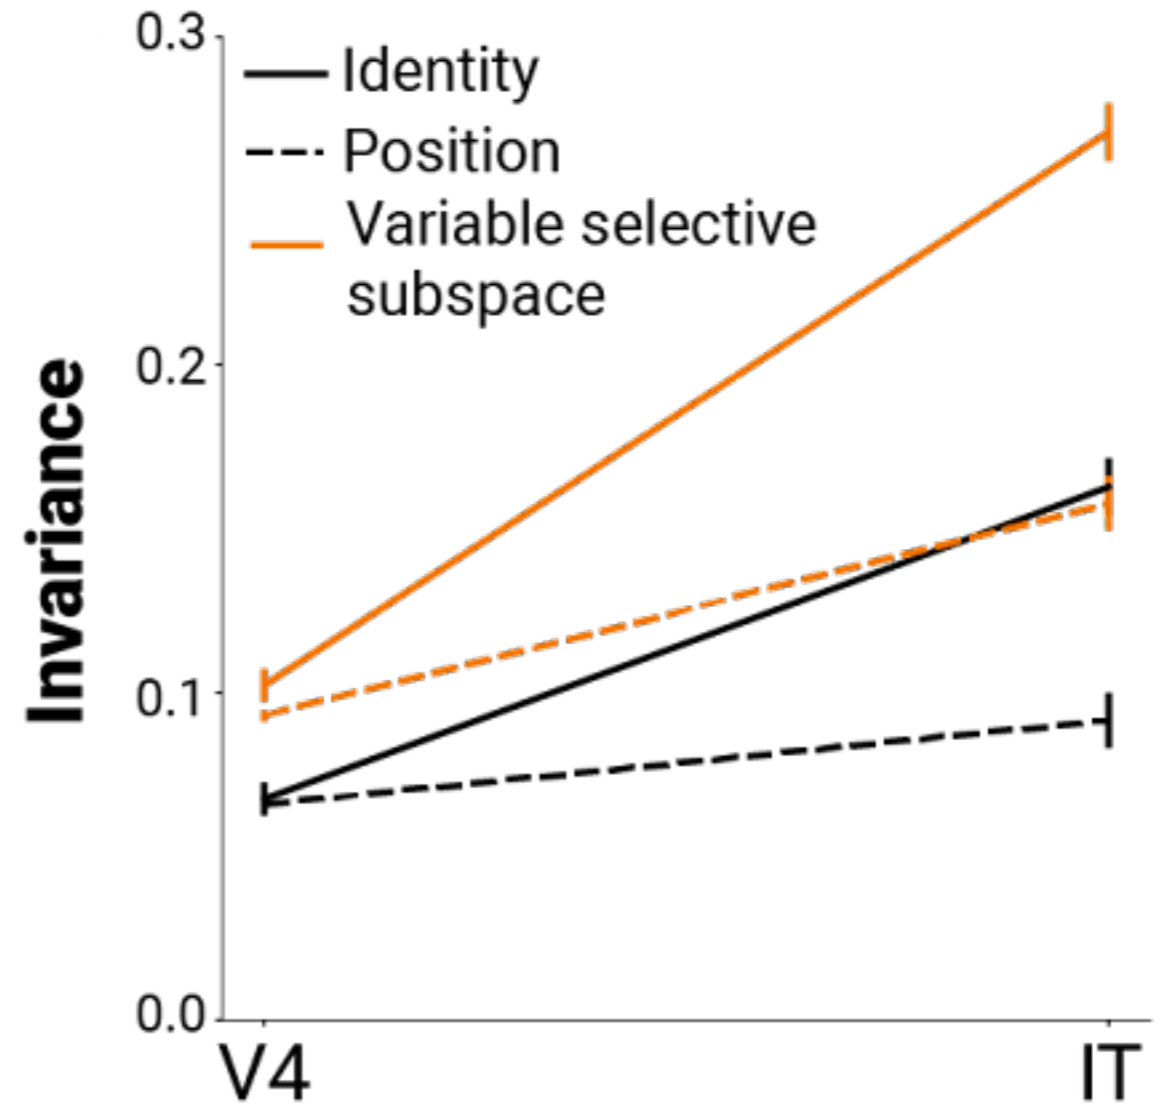
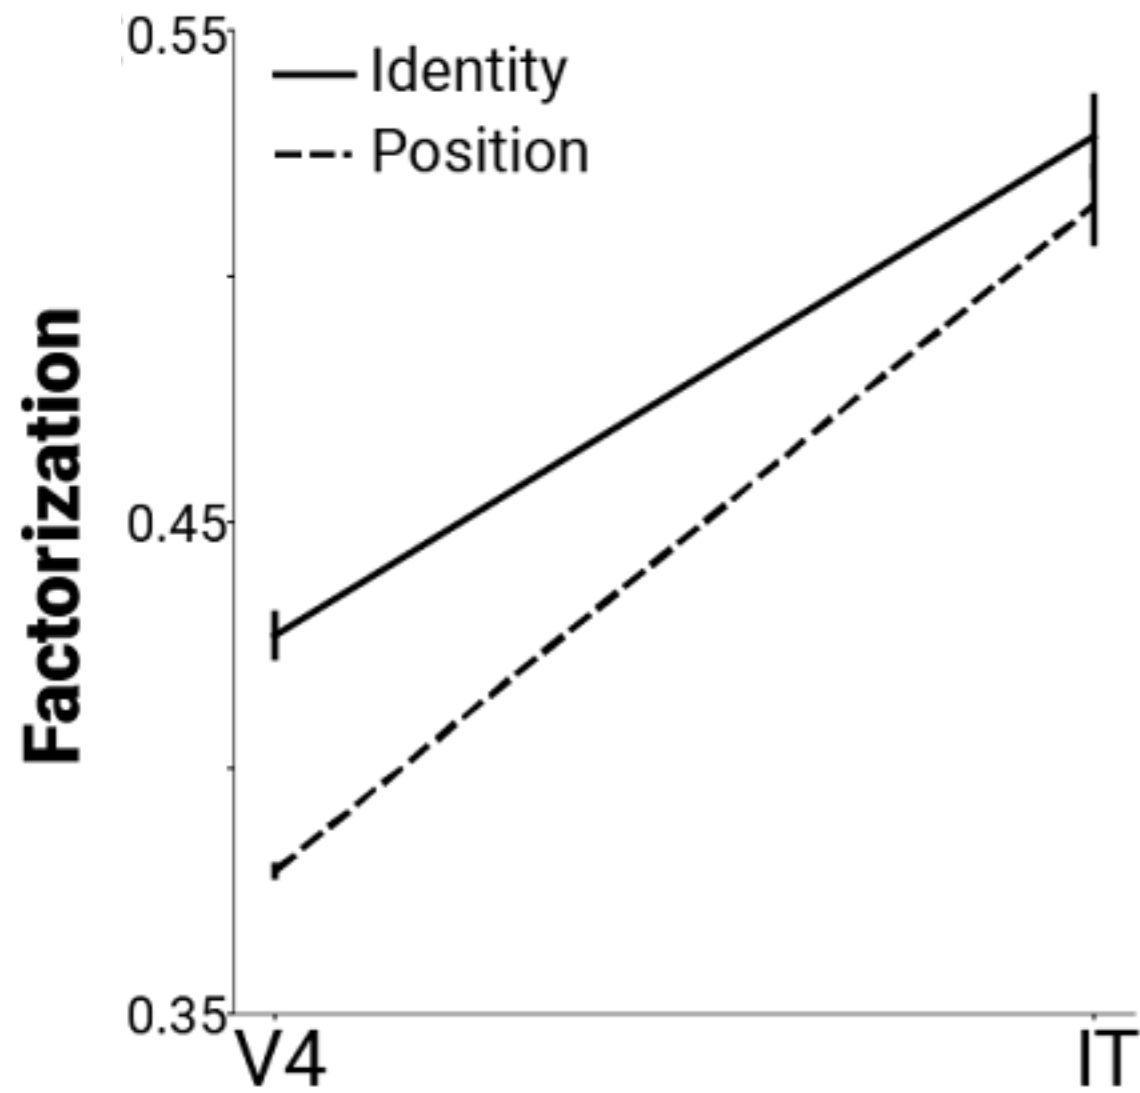
$$\text{Invariance: } 1 - \frac{\text{var}(P)}{\text{var}(\text{all})}$$

$$\text{Factorization: } 1 - \frac{\text{var}(P \mid \text{subspace}[Q])}{\text{var}(P)}$$

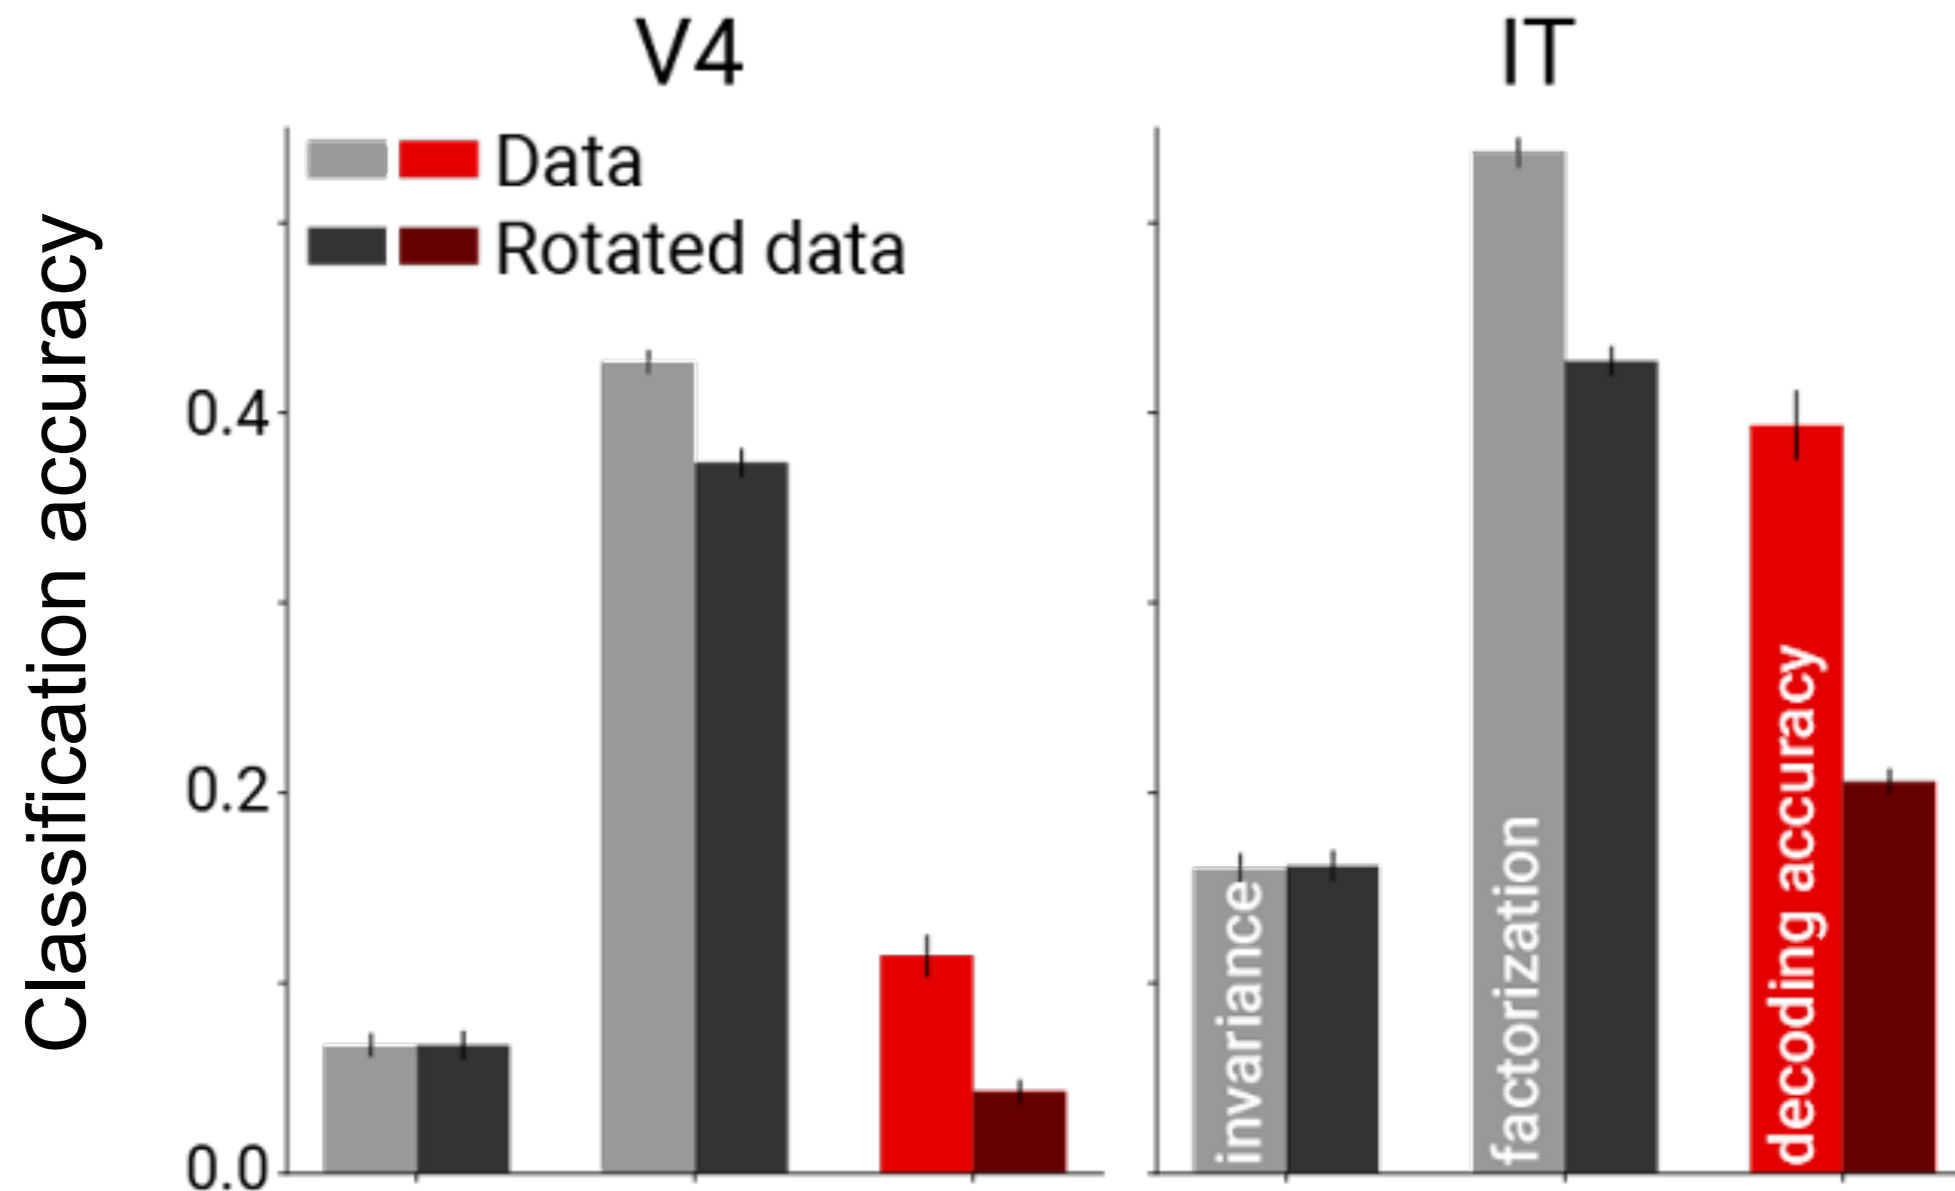
Notation

- $\text{var}(P)$ = “variance induced by parameter P (controlling all other parameters)”
- $\text{subspace}(P)$ = “top principal components of variance induced by parameter P ”
- $\text{var}(P \mid S)$ = “var(P) when projected onto S ”

Factorization increases across the visual stream



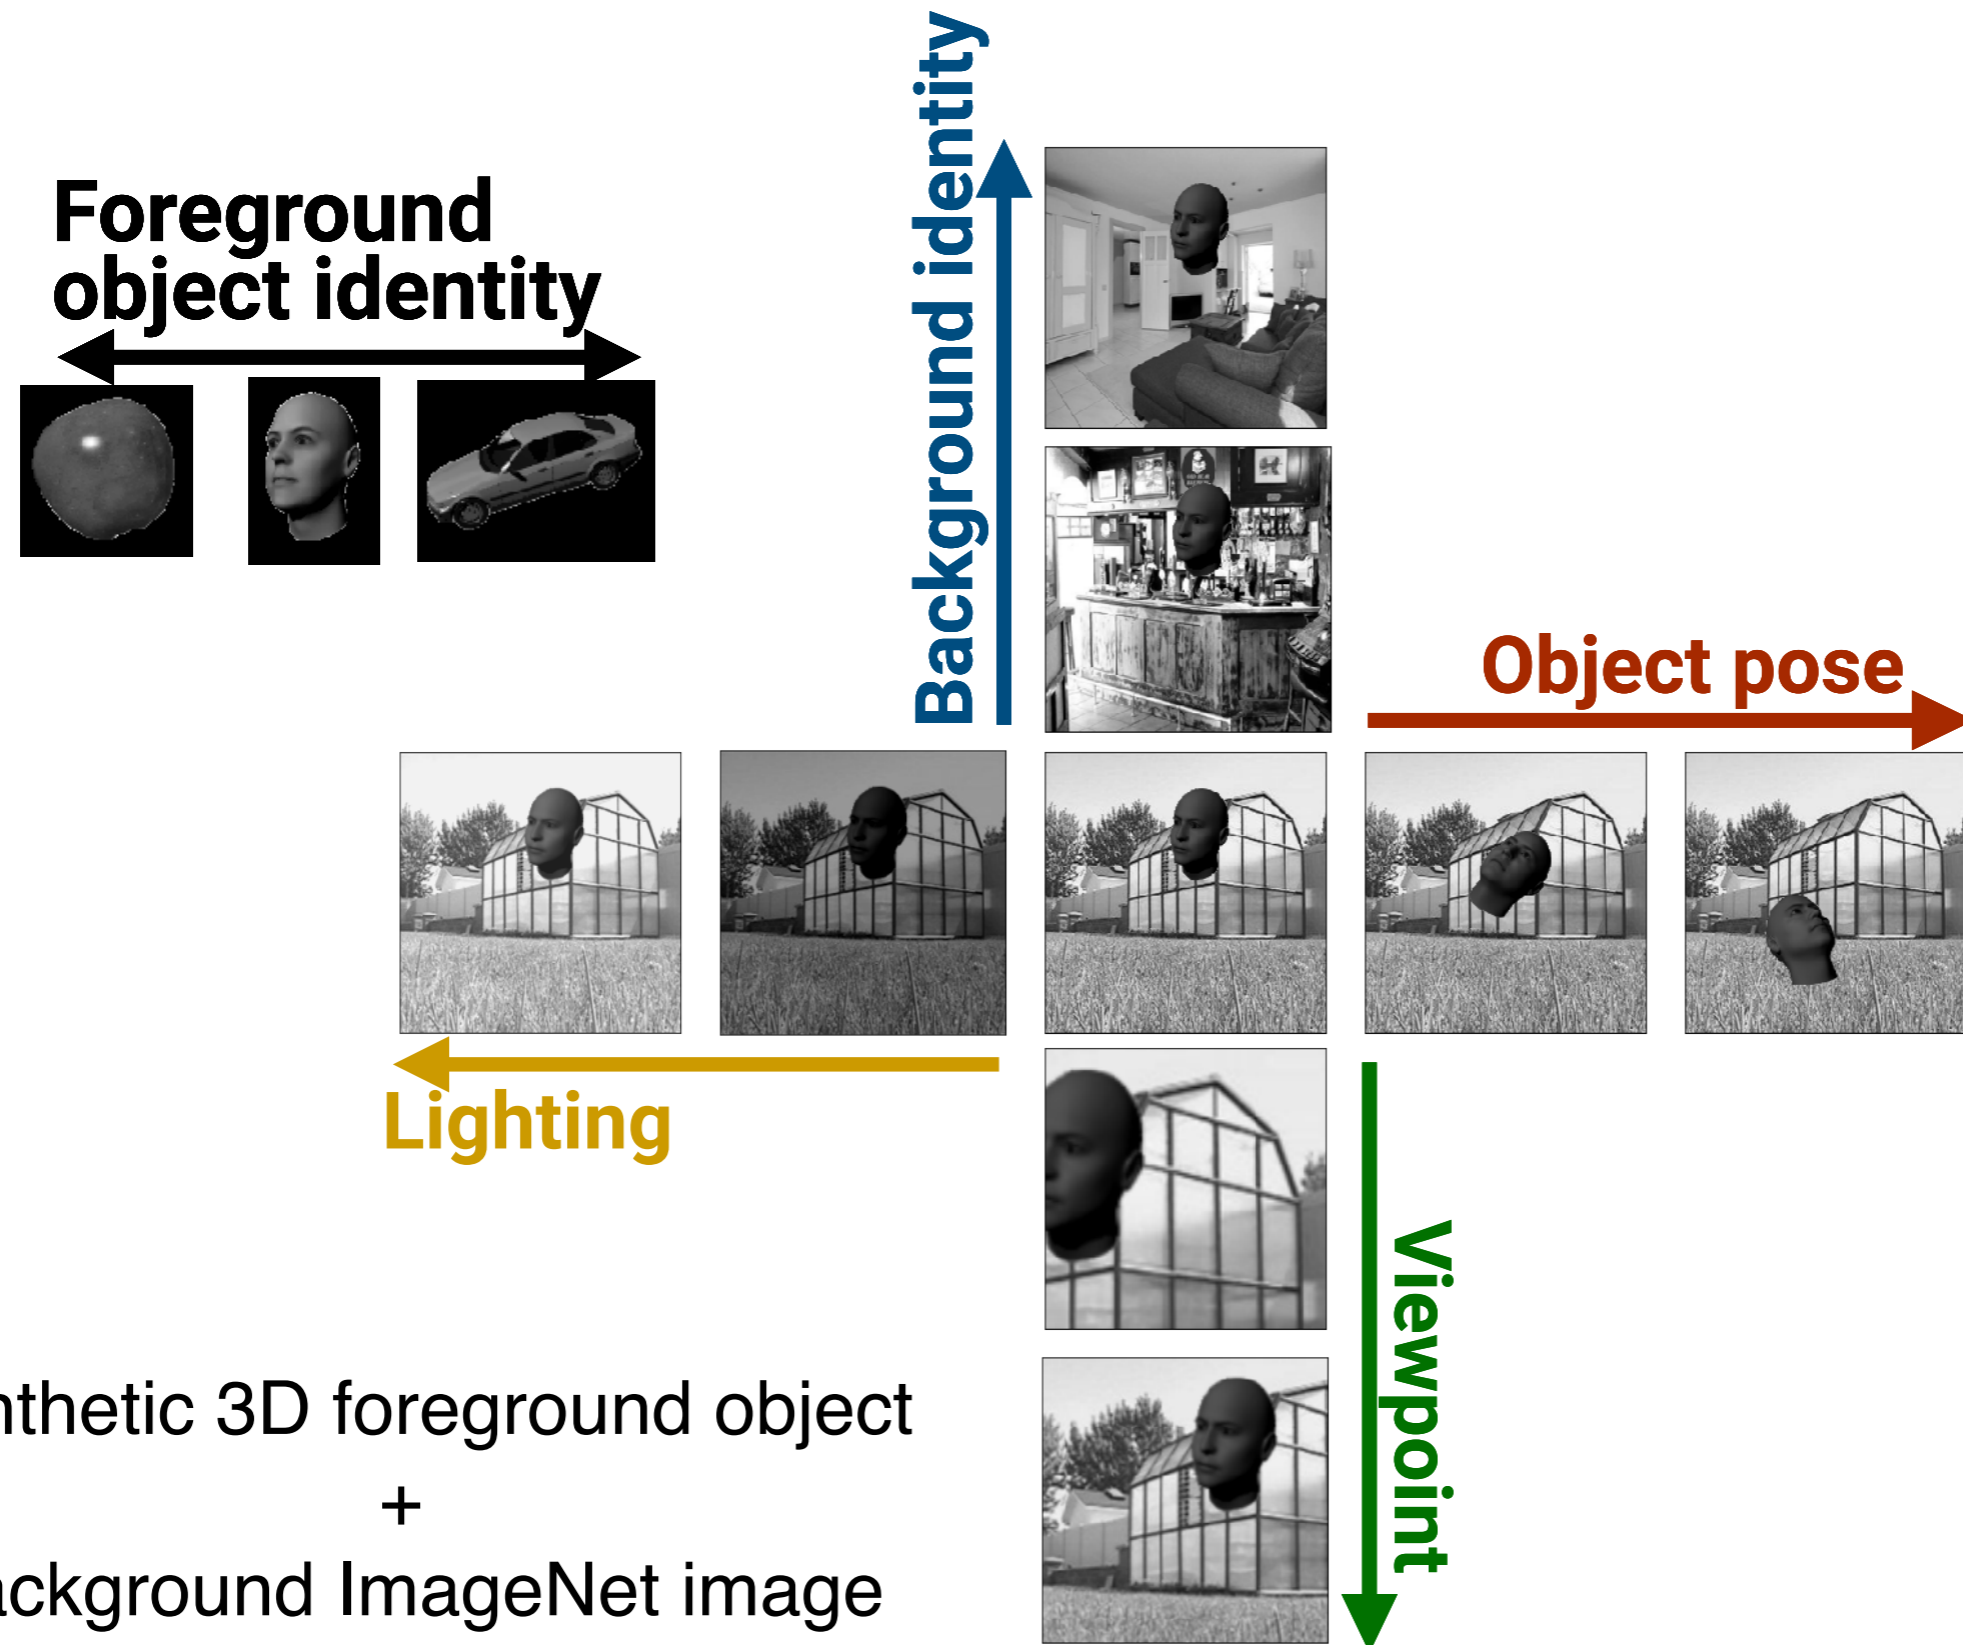
Factorization is important for object decoding



- Synthetic data with reduced factorization but otherwise identical statistics is much worse at supporting object decoding

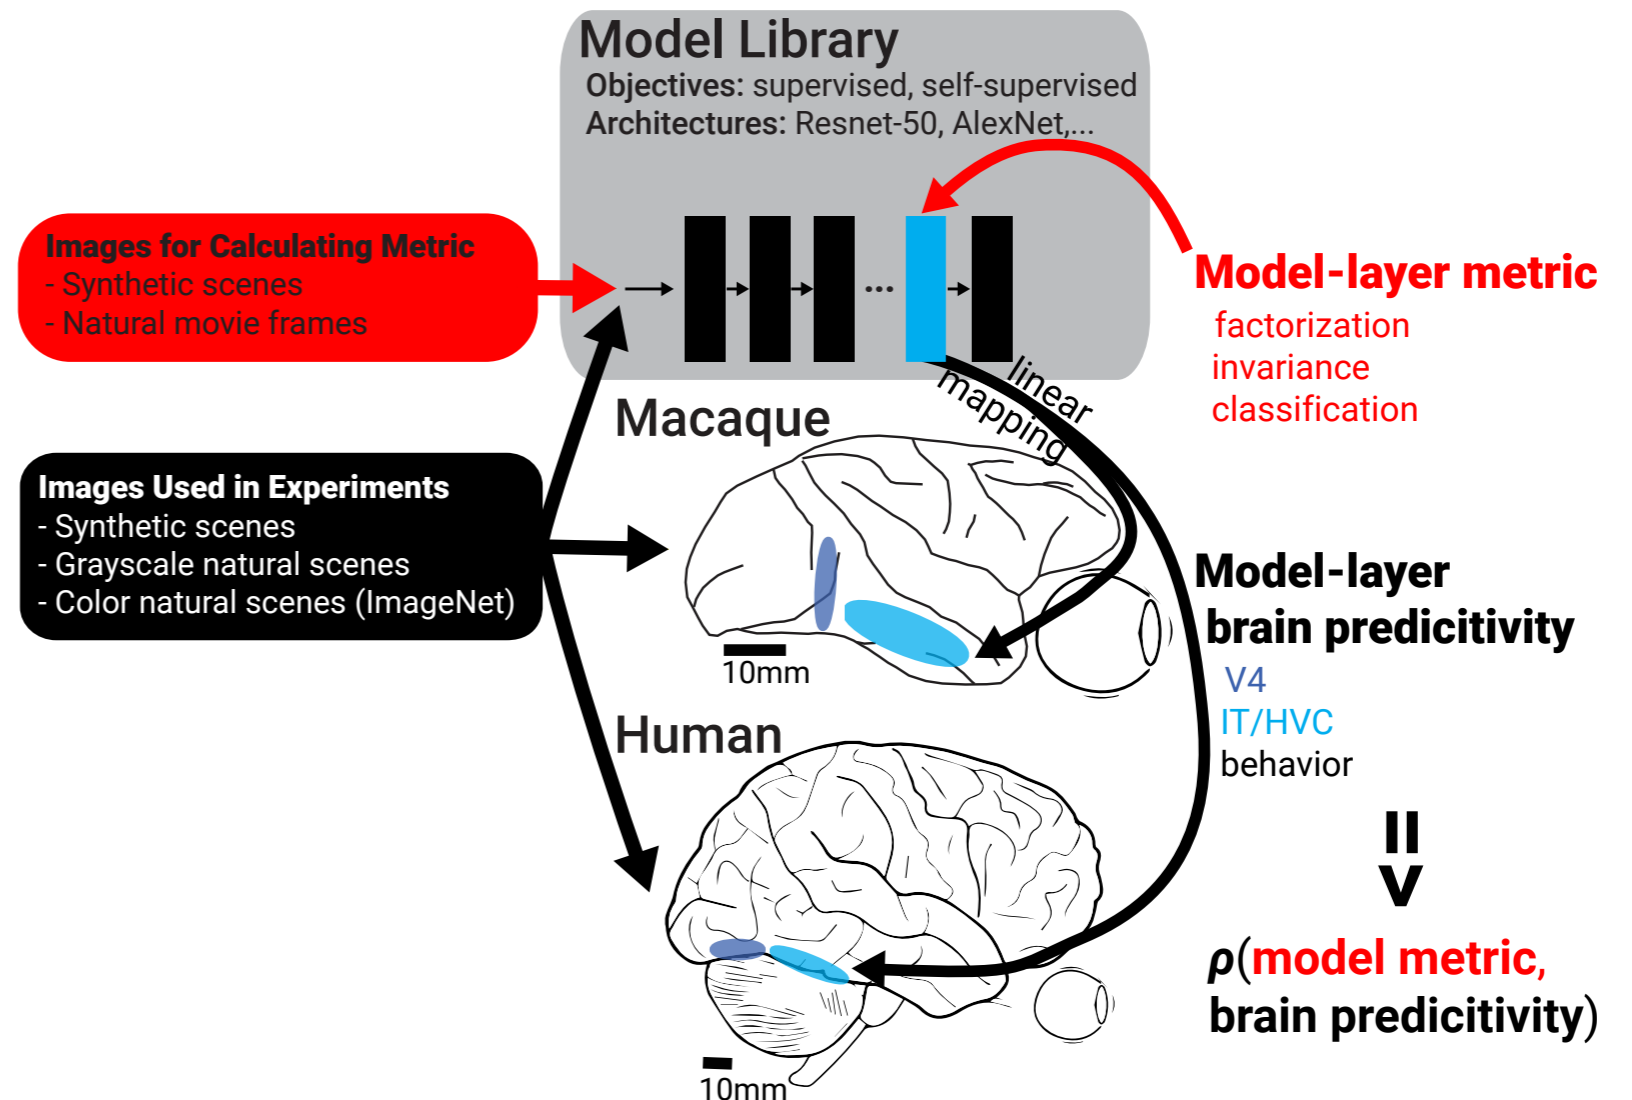
- Can't do much more than this with existing neural data
- No large-scale datasets of visual responses to parametrically varied scenes
- Solution: large-scale analysis of deep neural network models

A synthetic stimulus set to independently control scene variables



Synthetic 3D foreground object
+
Background ImageNet image

Strategy: what model properties are associated with match to neural data?



1) Compute scene encoding properties of model


- Factorization + invariance metrics
- Involves >1k PCA runs over >100k dimensions

2) Assess how well models match neural data


- Fit linear regression from model neurons to data neurons / voxels / behavior in response to the same images
- Involves >1k ridge regressions over >100k dimensions

Example: viewpoint factorization


Ephys 1




Ephys 2



fMRI 1



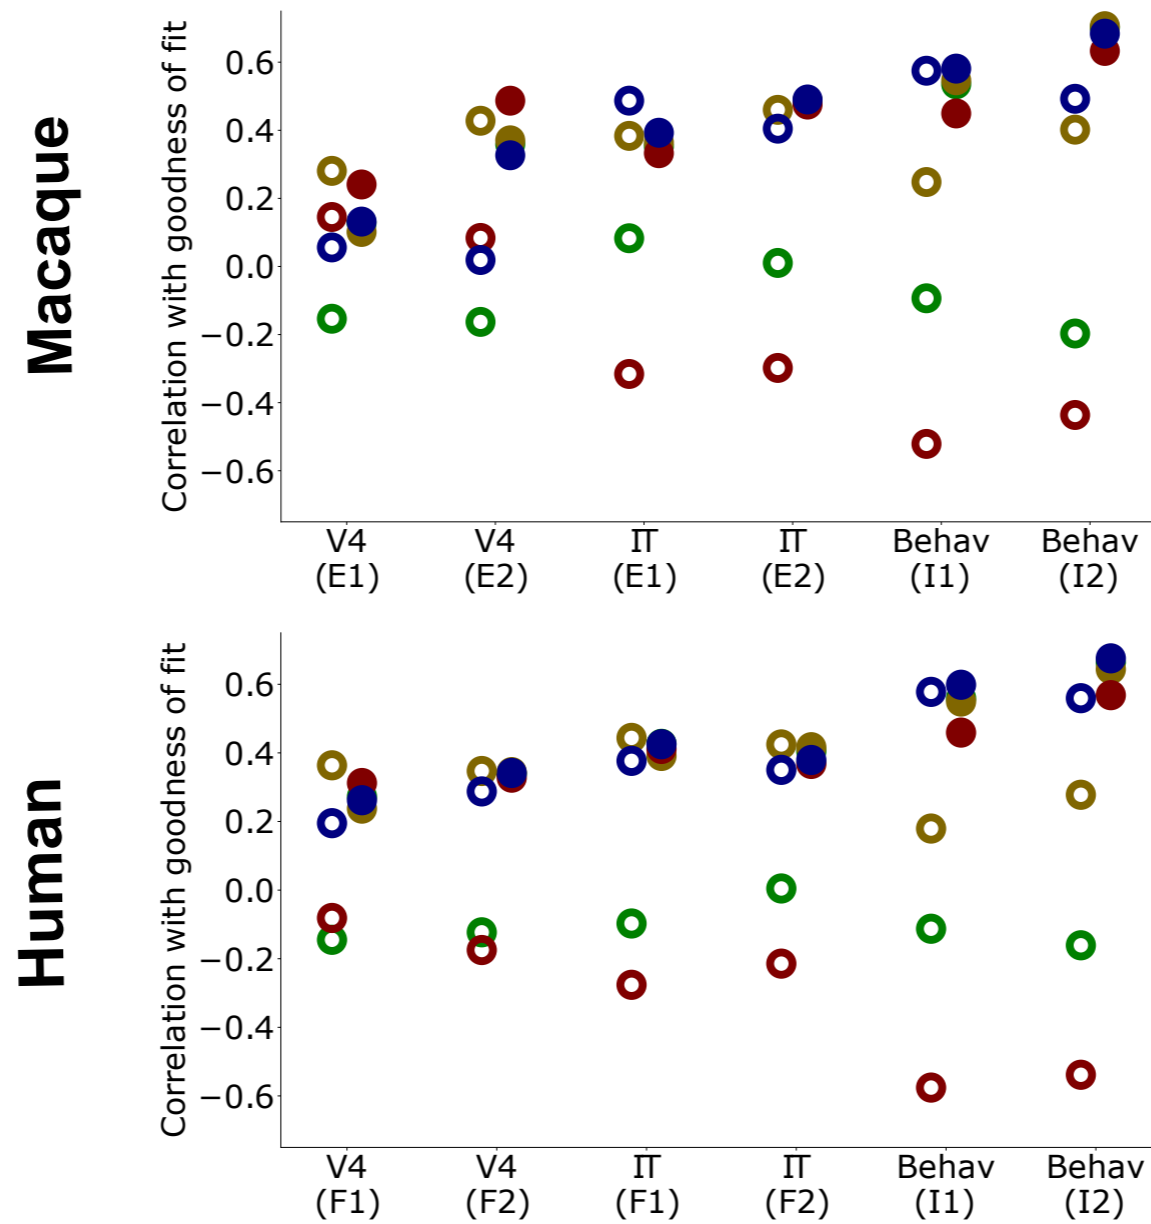
fMRI 2



- Shown here: last layer of each network
- Take correlation values, average across datasets + layers
- Repeat for other metrics


More brain-like models tend to factorize scene variables

● Factorization ○ Invariance
Viewpoint Object pose Background Lighting



Factorization computed from natural videos

Viewpoint



Object motion


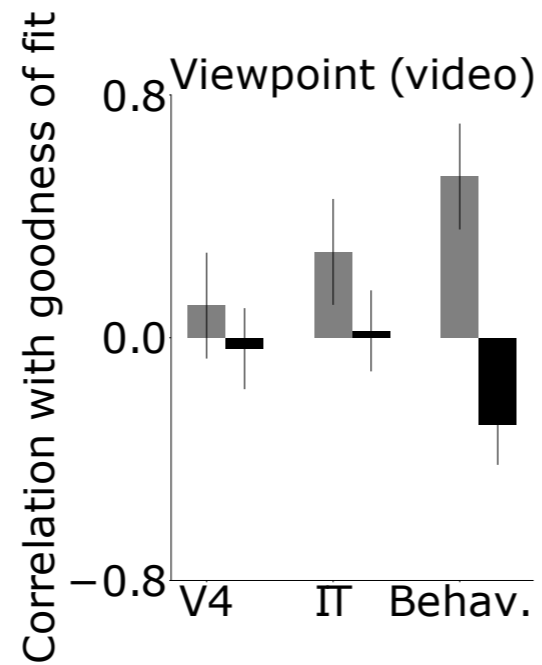


Macaque


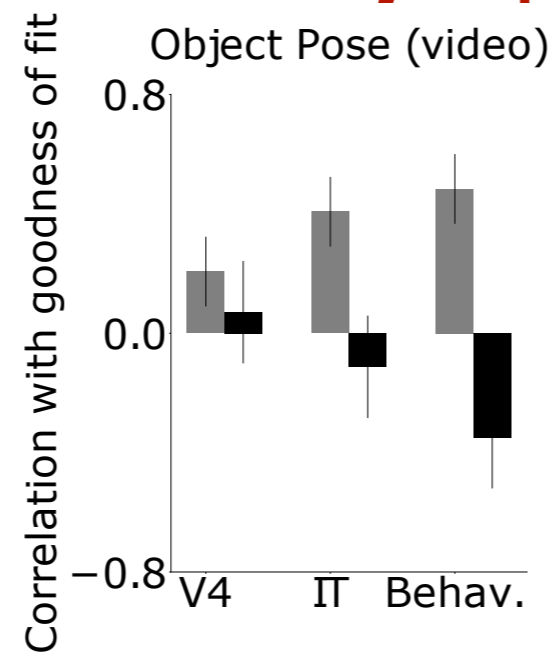
Human

Viewpoint (video)

Factorization
Invariance




Object pose (video)




Factorization complementary to classification performance as an indicator of brain-like models


V4



IT



Behavior



Classification performance

Viewpoint fact. (+ class)

Background fact. (+ class)

Object pose fact. (+ class)

Lighting fact. (+ class)

Follow-up directions

- Auxiliary objective functions to directly incentivize factorization
 - Shown to improve out-of-distribution model performance (Ying et al., 2023)
- Unsupervised discovery of parameters being factorized by a given representation
- What is the biological substrate of parameter-driven subspaces?

Thank you!