

System Models of Brain-Like Vision & Language

Martin Schrimpf

research scientist, MIT

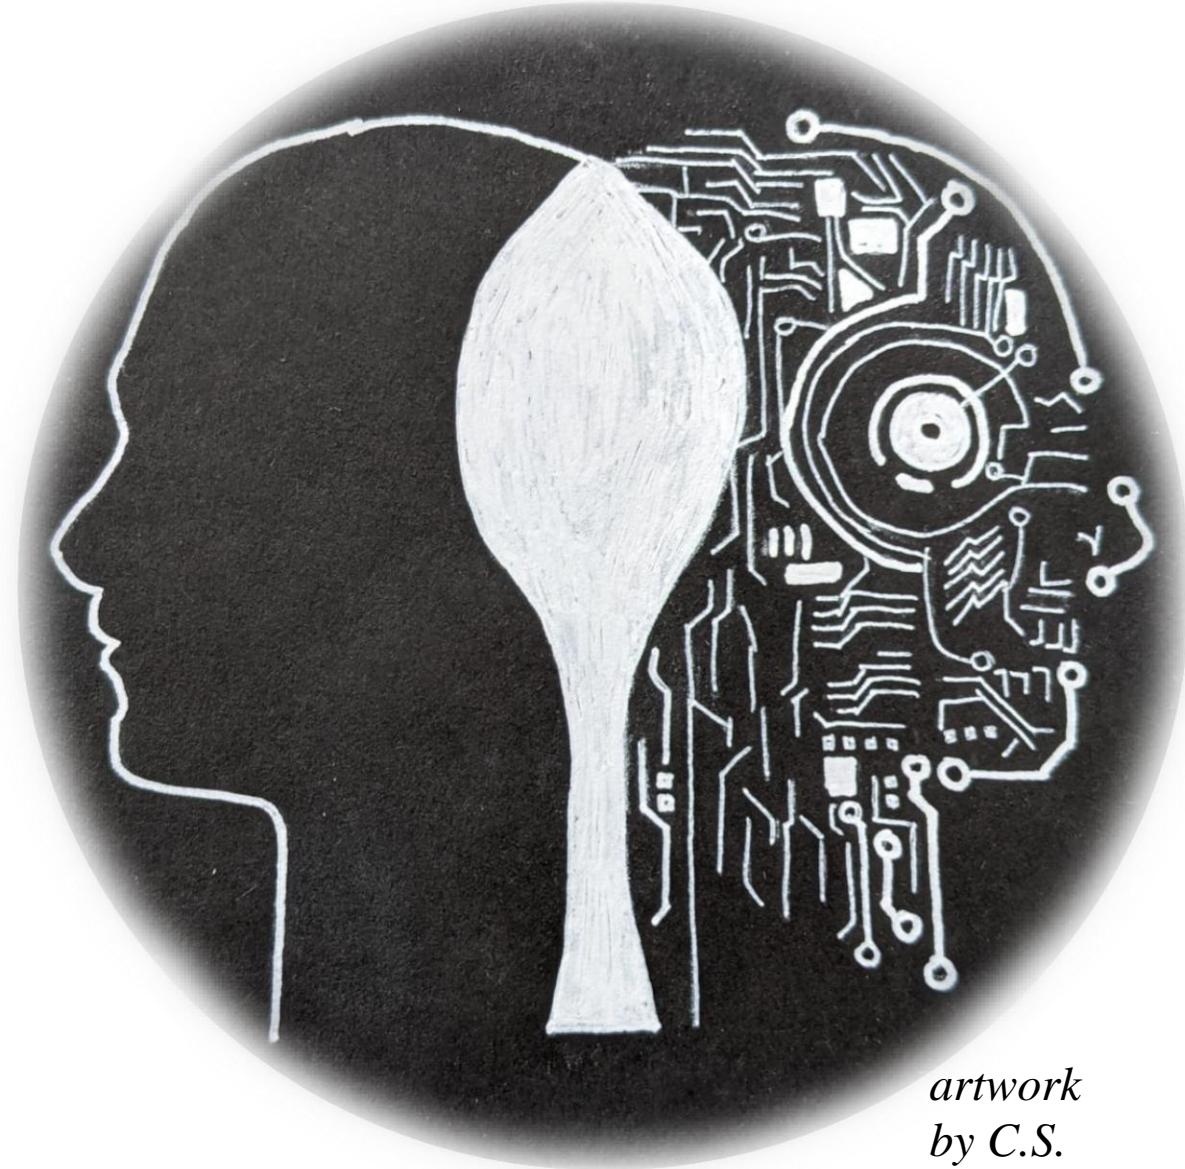
tenure-track assistant professor, EPFL



msch@mit.edu



@martin_schrimpf



*artwork
by C.S.*



Jonas Kubilius



Tiago Marques



Kohitij Kar



Franziska Geiger



Joel Dapello



Paul McGrath



Mike Ferguson



Jim DiCarlo



Greta Tuckute



Carina Kauf



Idan Blank



Corey Ziemba



Tony Movshon



Pouya Bashivan



Kailyn Schmidt



Jon Prescott-Roy



Ratan Murty



Ev Fedorenko



Chengxu Zhuang



Eshed Margalit



Rishi Rajalingham



Dan Yamins



Michael Lee



Anna Ivanova



Eghbal Hosseini



Robert Ajemian



Nancy Kanwisher



Josh Tenenbaum

Funding:

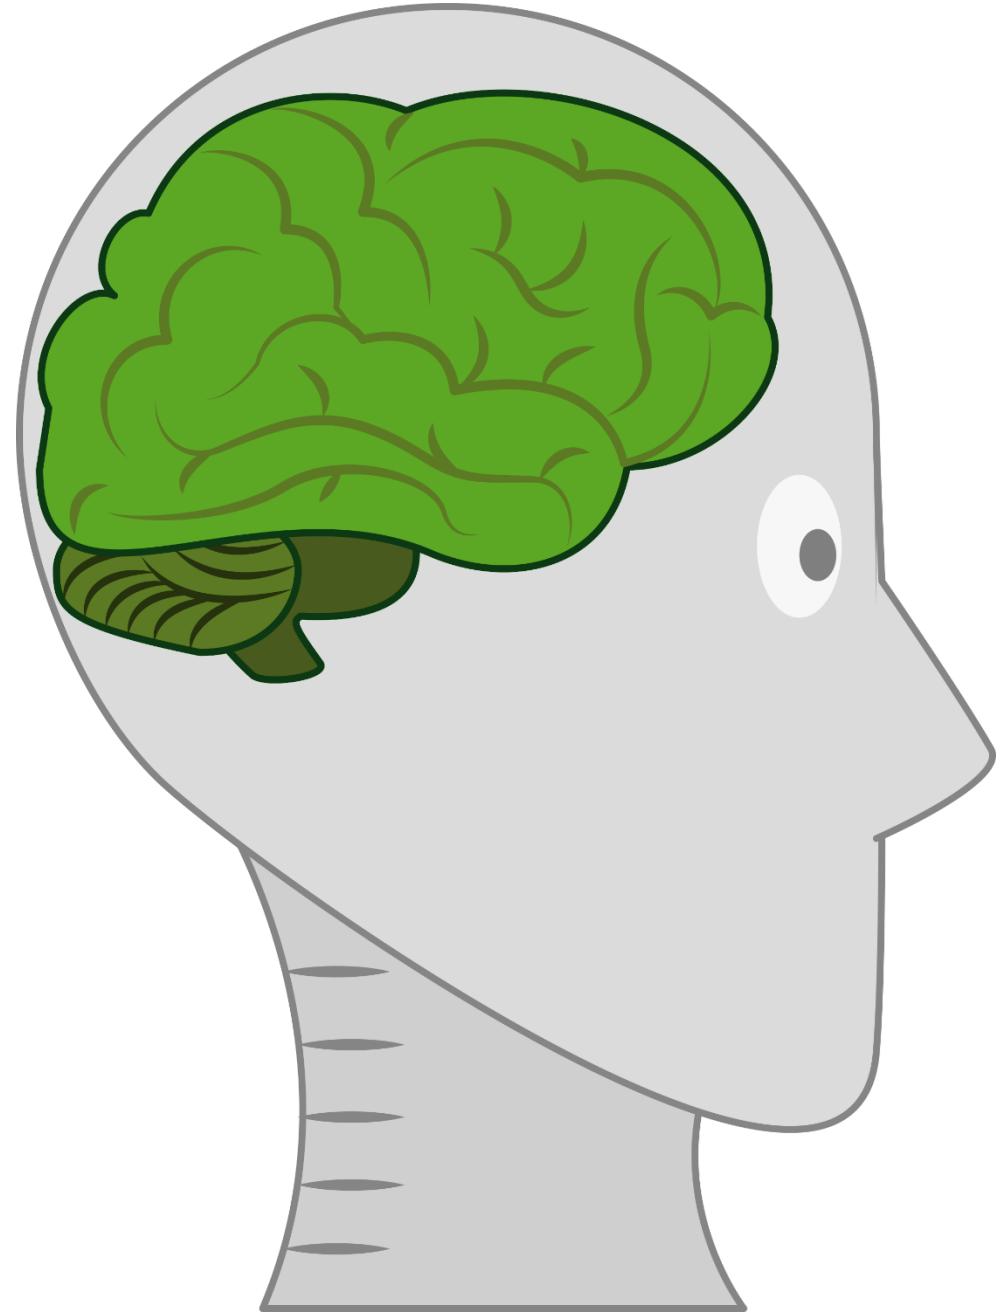


McGOVERN INSTITUTE
FOR BRAIN RESEARCH



Center for Brain-Inspired Computing



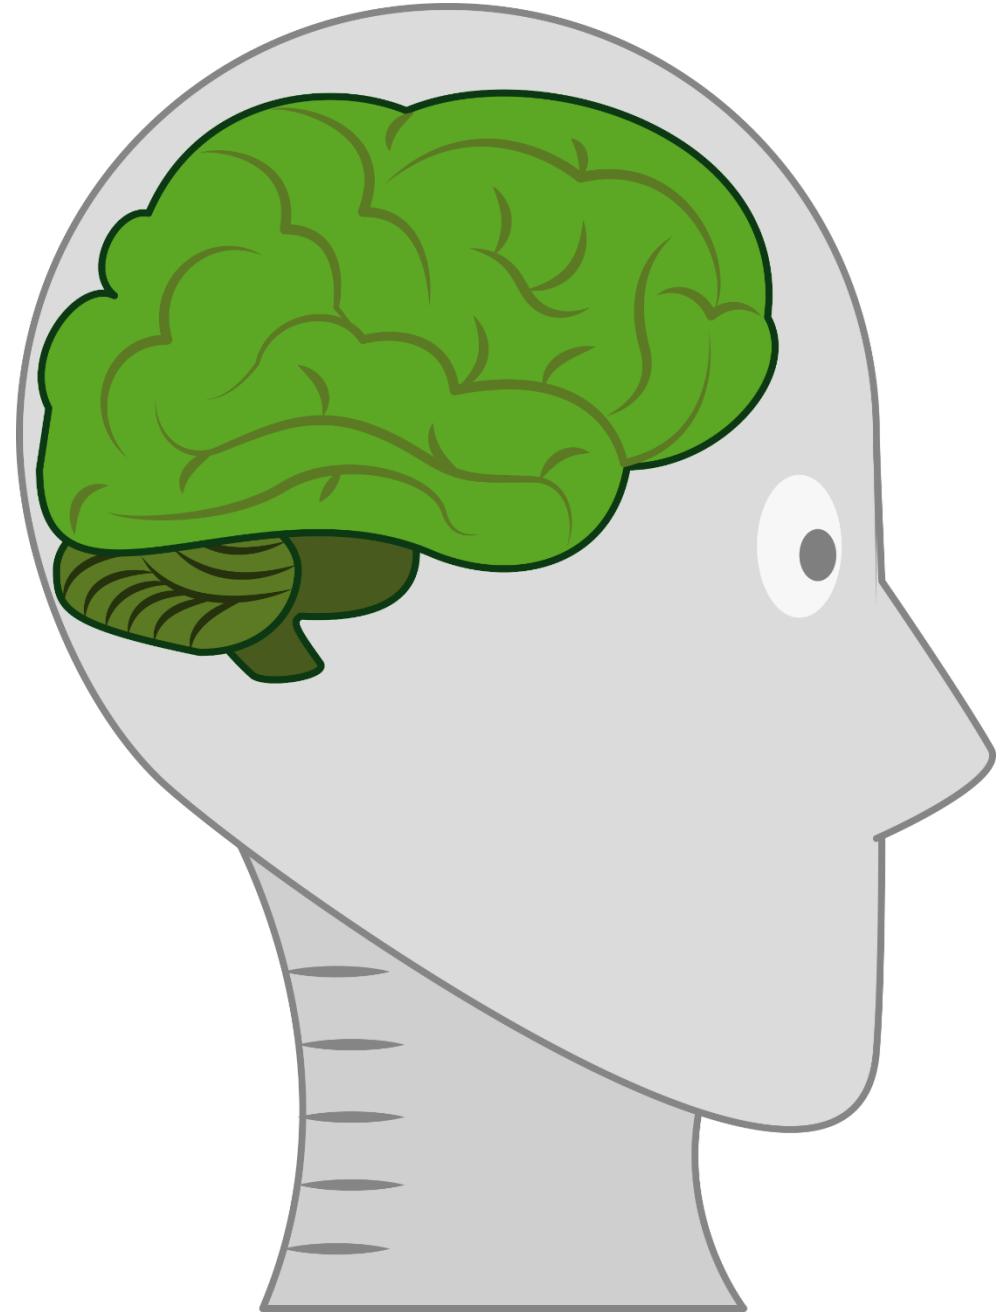


A core goal of our field:
model human intelligence and
underlying neural mechanisms

Computational
understanding
of natural intelligence

Next-generation
AI algorithms

Future clinical
applications



A core goal of our field:
model human intelligence and
underlying neural mechanisms

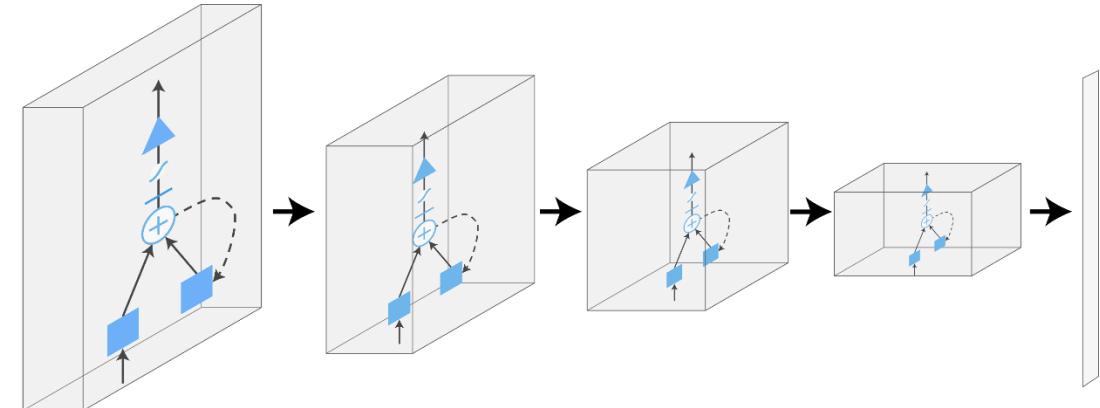
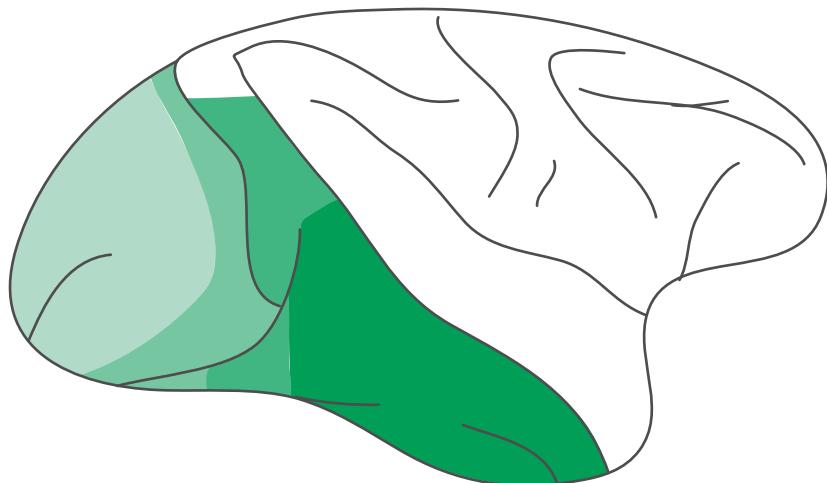
How can we even tell
we're making progress?

This talk: a proposal for integrative
model testing in vision & language

Modeling Primate Visual Intelligence

Evaluate model **alignment to all experiments** in the domain of visual intelligence.

The alternative of **evaluating isolated models on isolated data** is horribly inefficient

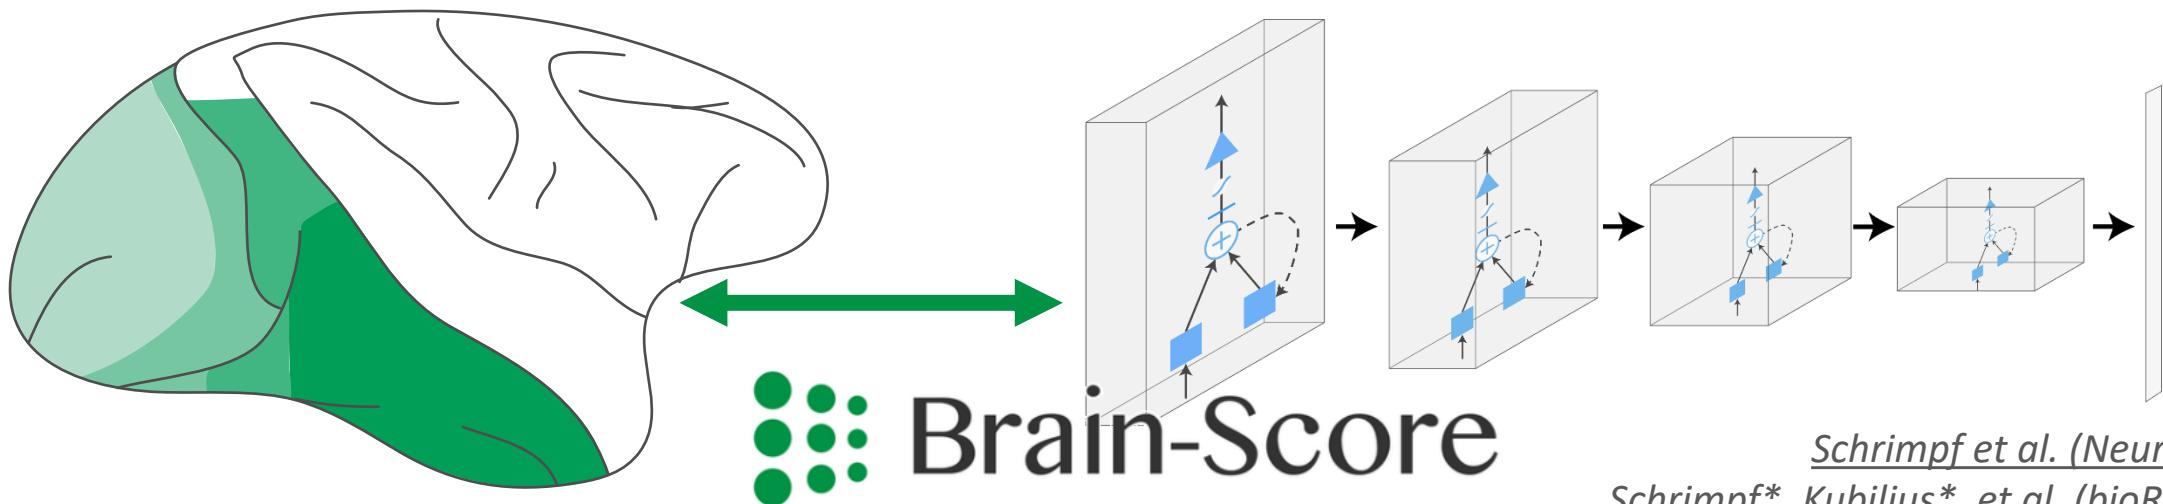


Schrimpf et al. (Neuron 2020)
Schrimpf, Kubilius*, et al. (bioRxiv 2018)*

Modeling Primate Visual Intelligence

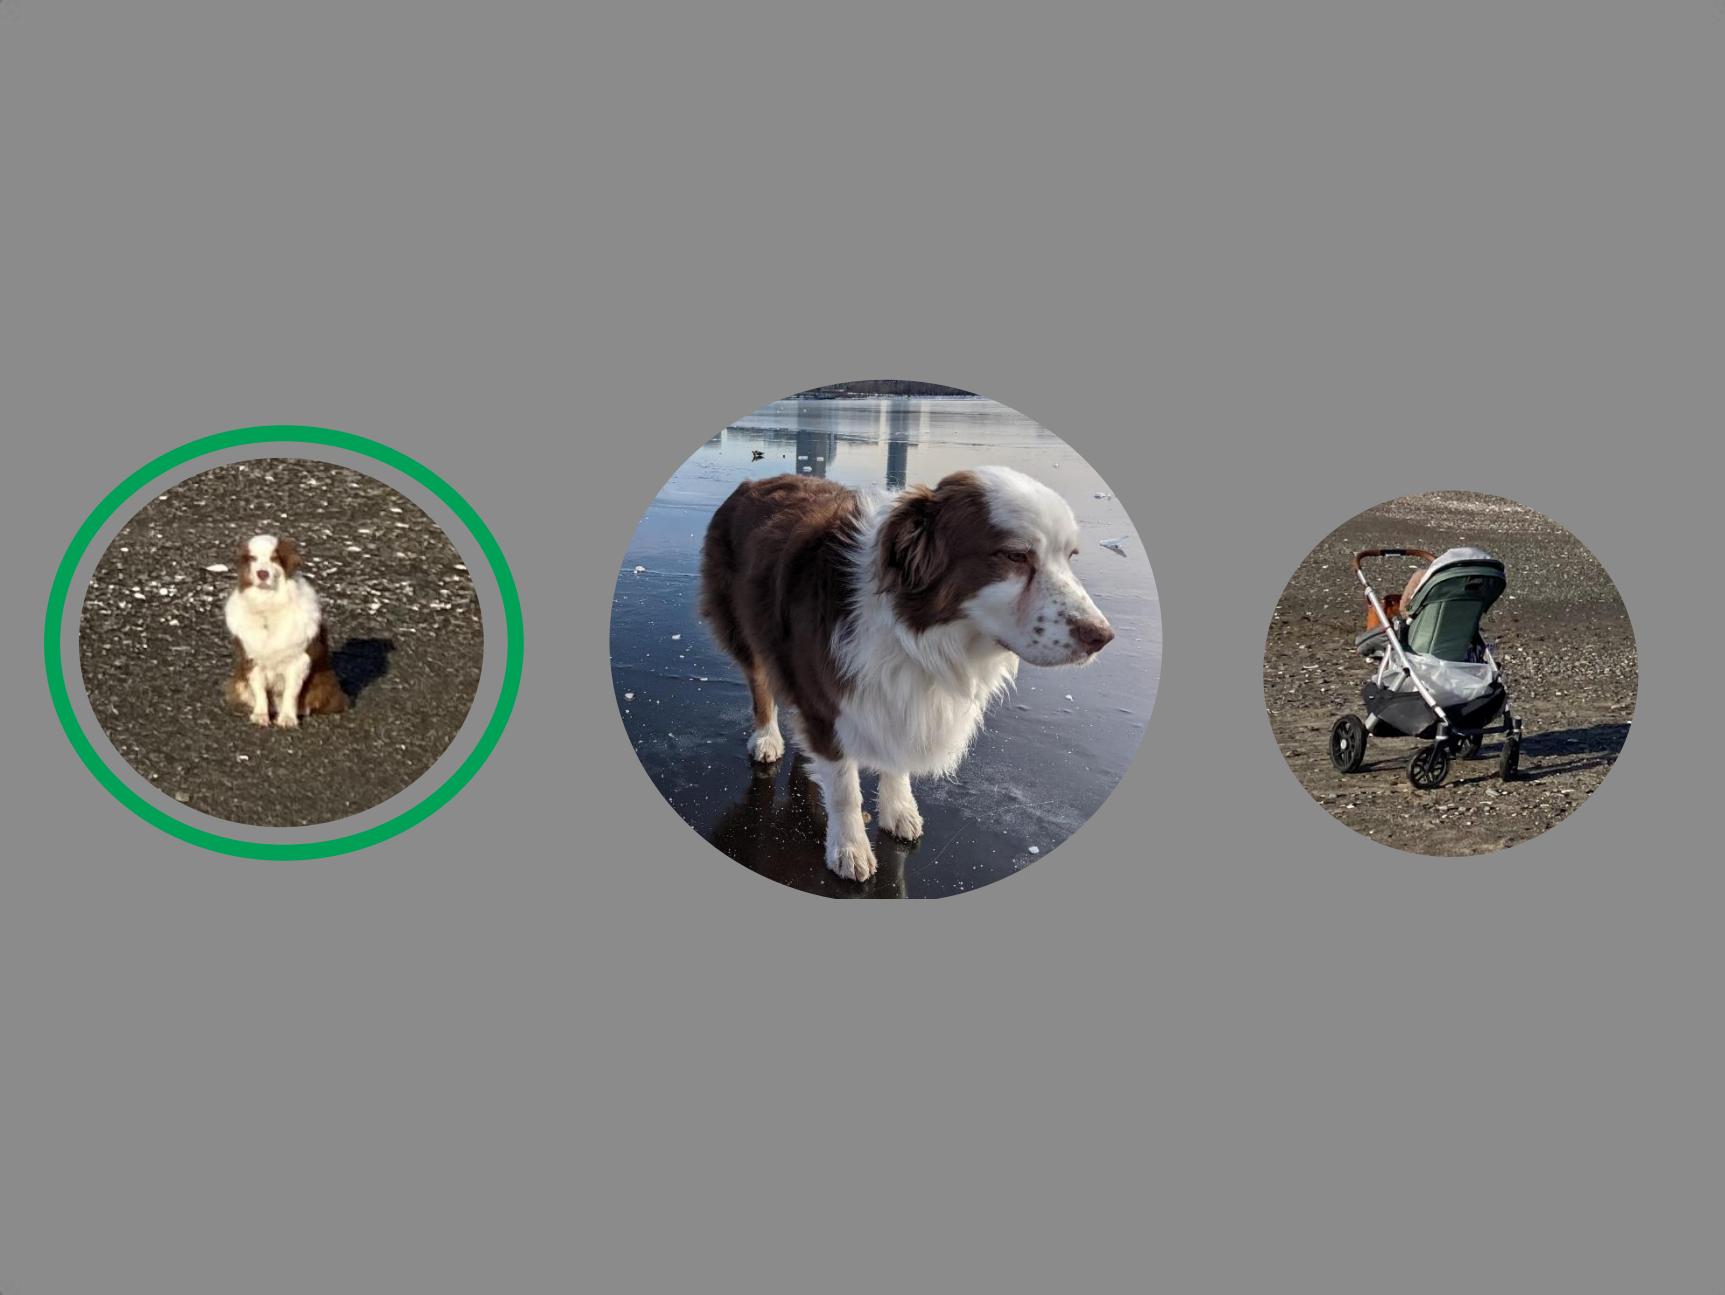
Evaluate model **alignment** to all experiments in the domain of visual intelligence.

- Test behavioral alignment by showing the **same images** to humans and models and comparing their outputs
- Test neural alignment at the level of **spike rates** since they linearly predict visual behaviors (e.g. *Majaj**, *Hong**, *et al.* 2015; *I Gusti Bagus et al.* 2022). Include **non-human primates** since their visual behaviors and visual cortex are highly similar to humans (e.g. *Kriegeskorte et al.* 2008; *Rajalingham et al.* 2018)



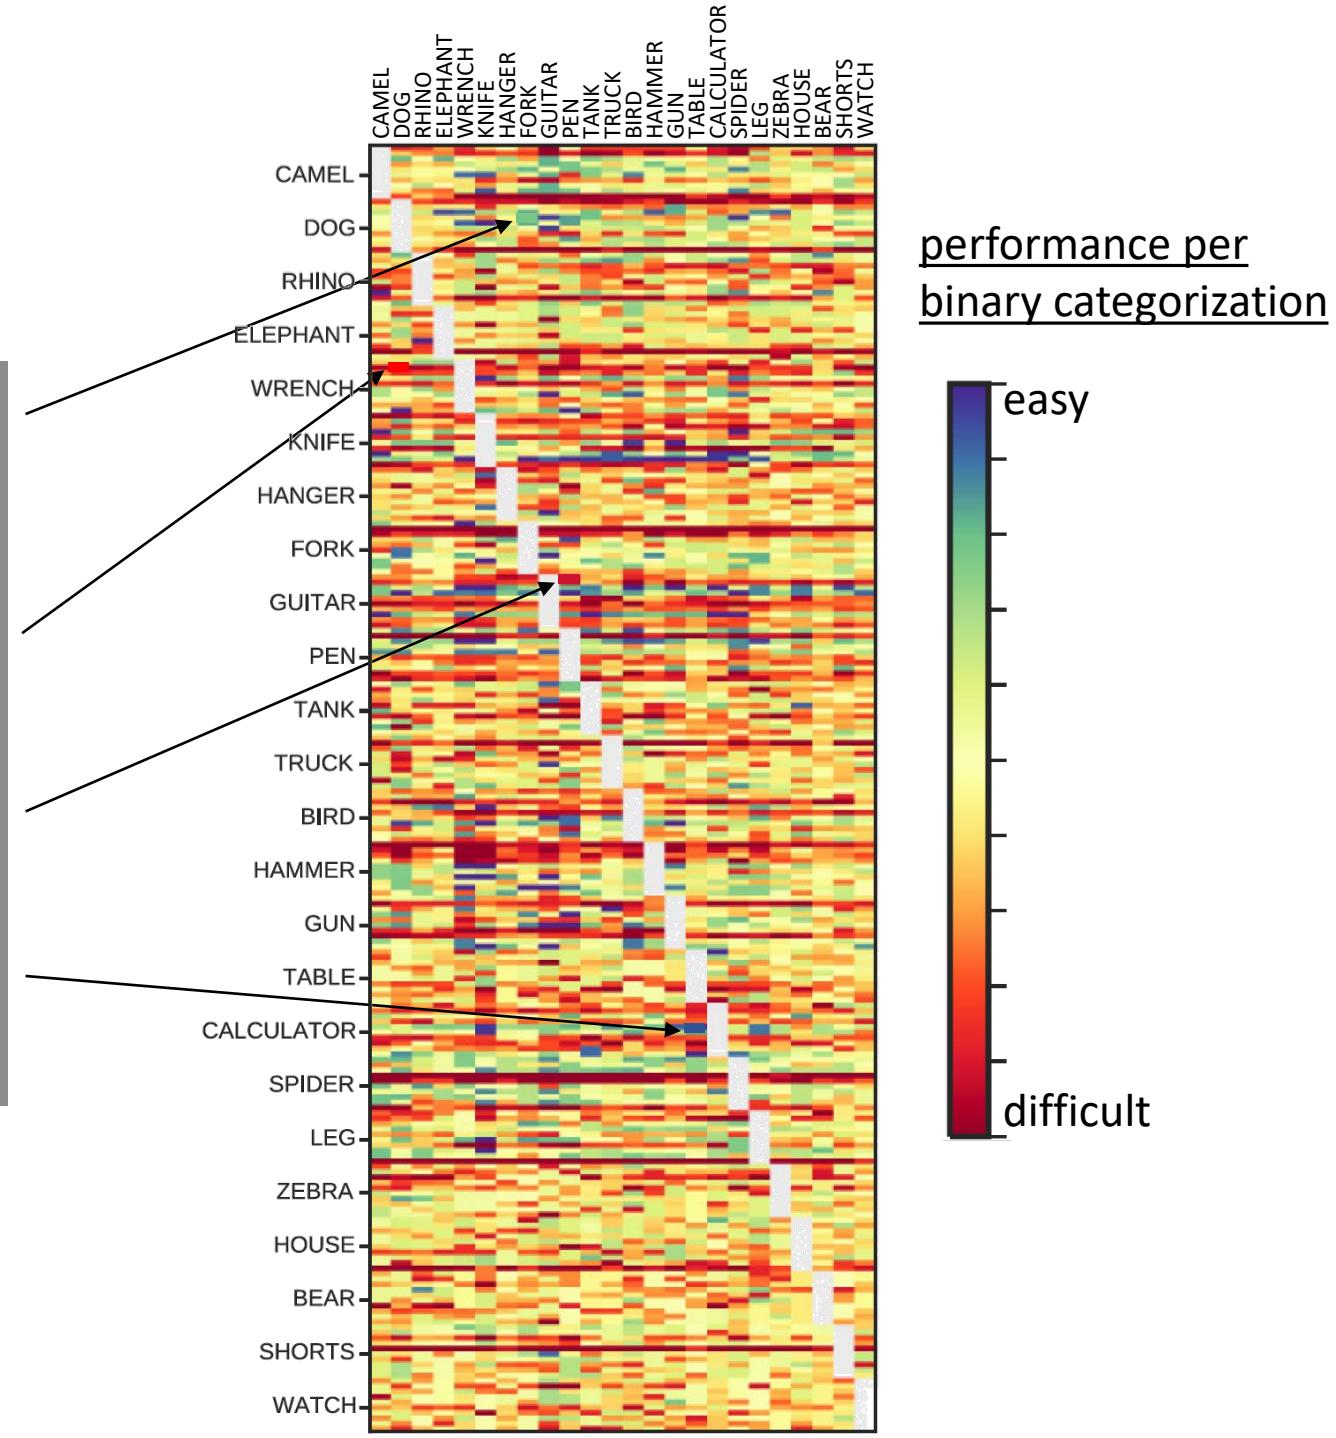
Schrimpf et al. (Neuron 2020)
Schrimpf, Kubilius*, et al. (bioRxiv 2018)*

Behavioral benchmark

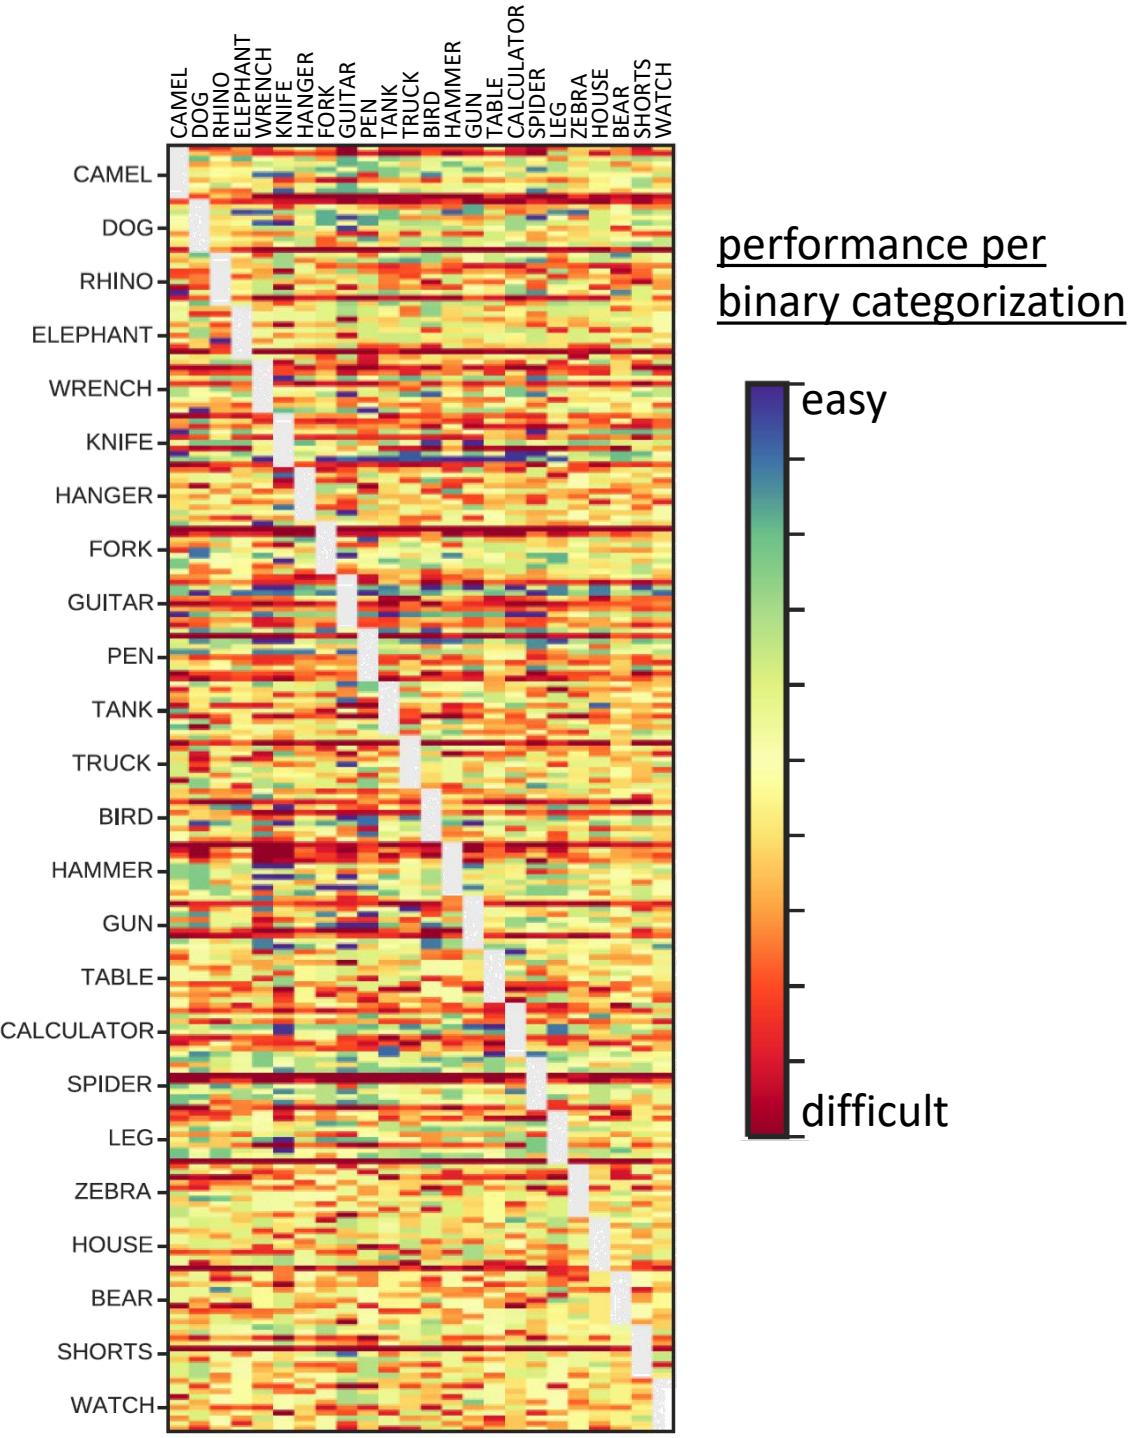


Behavioral benchmark

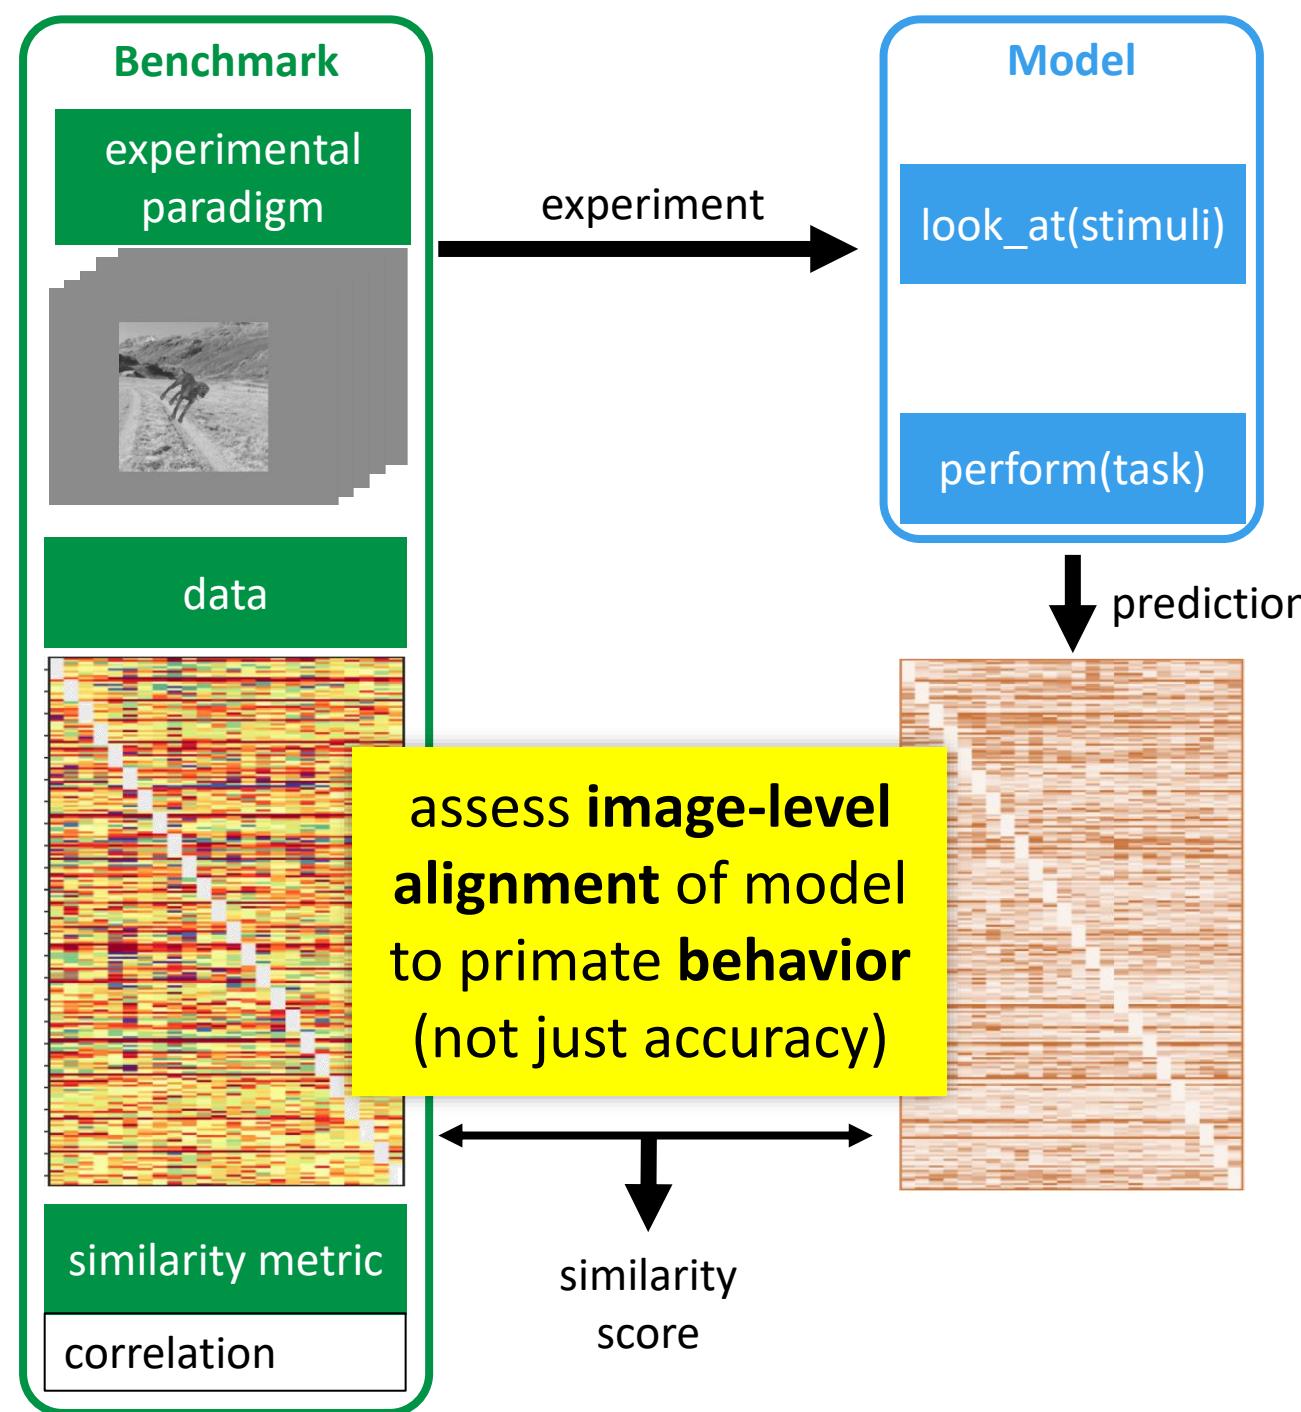
Behavioral benchmark



Behavioral benchmark

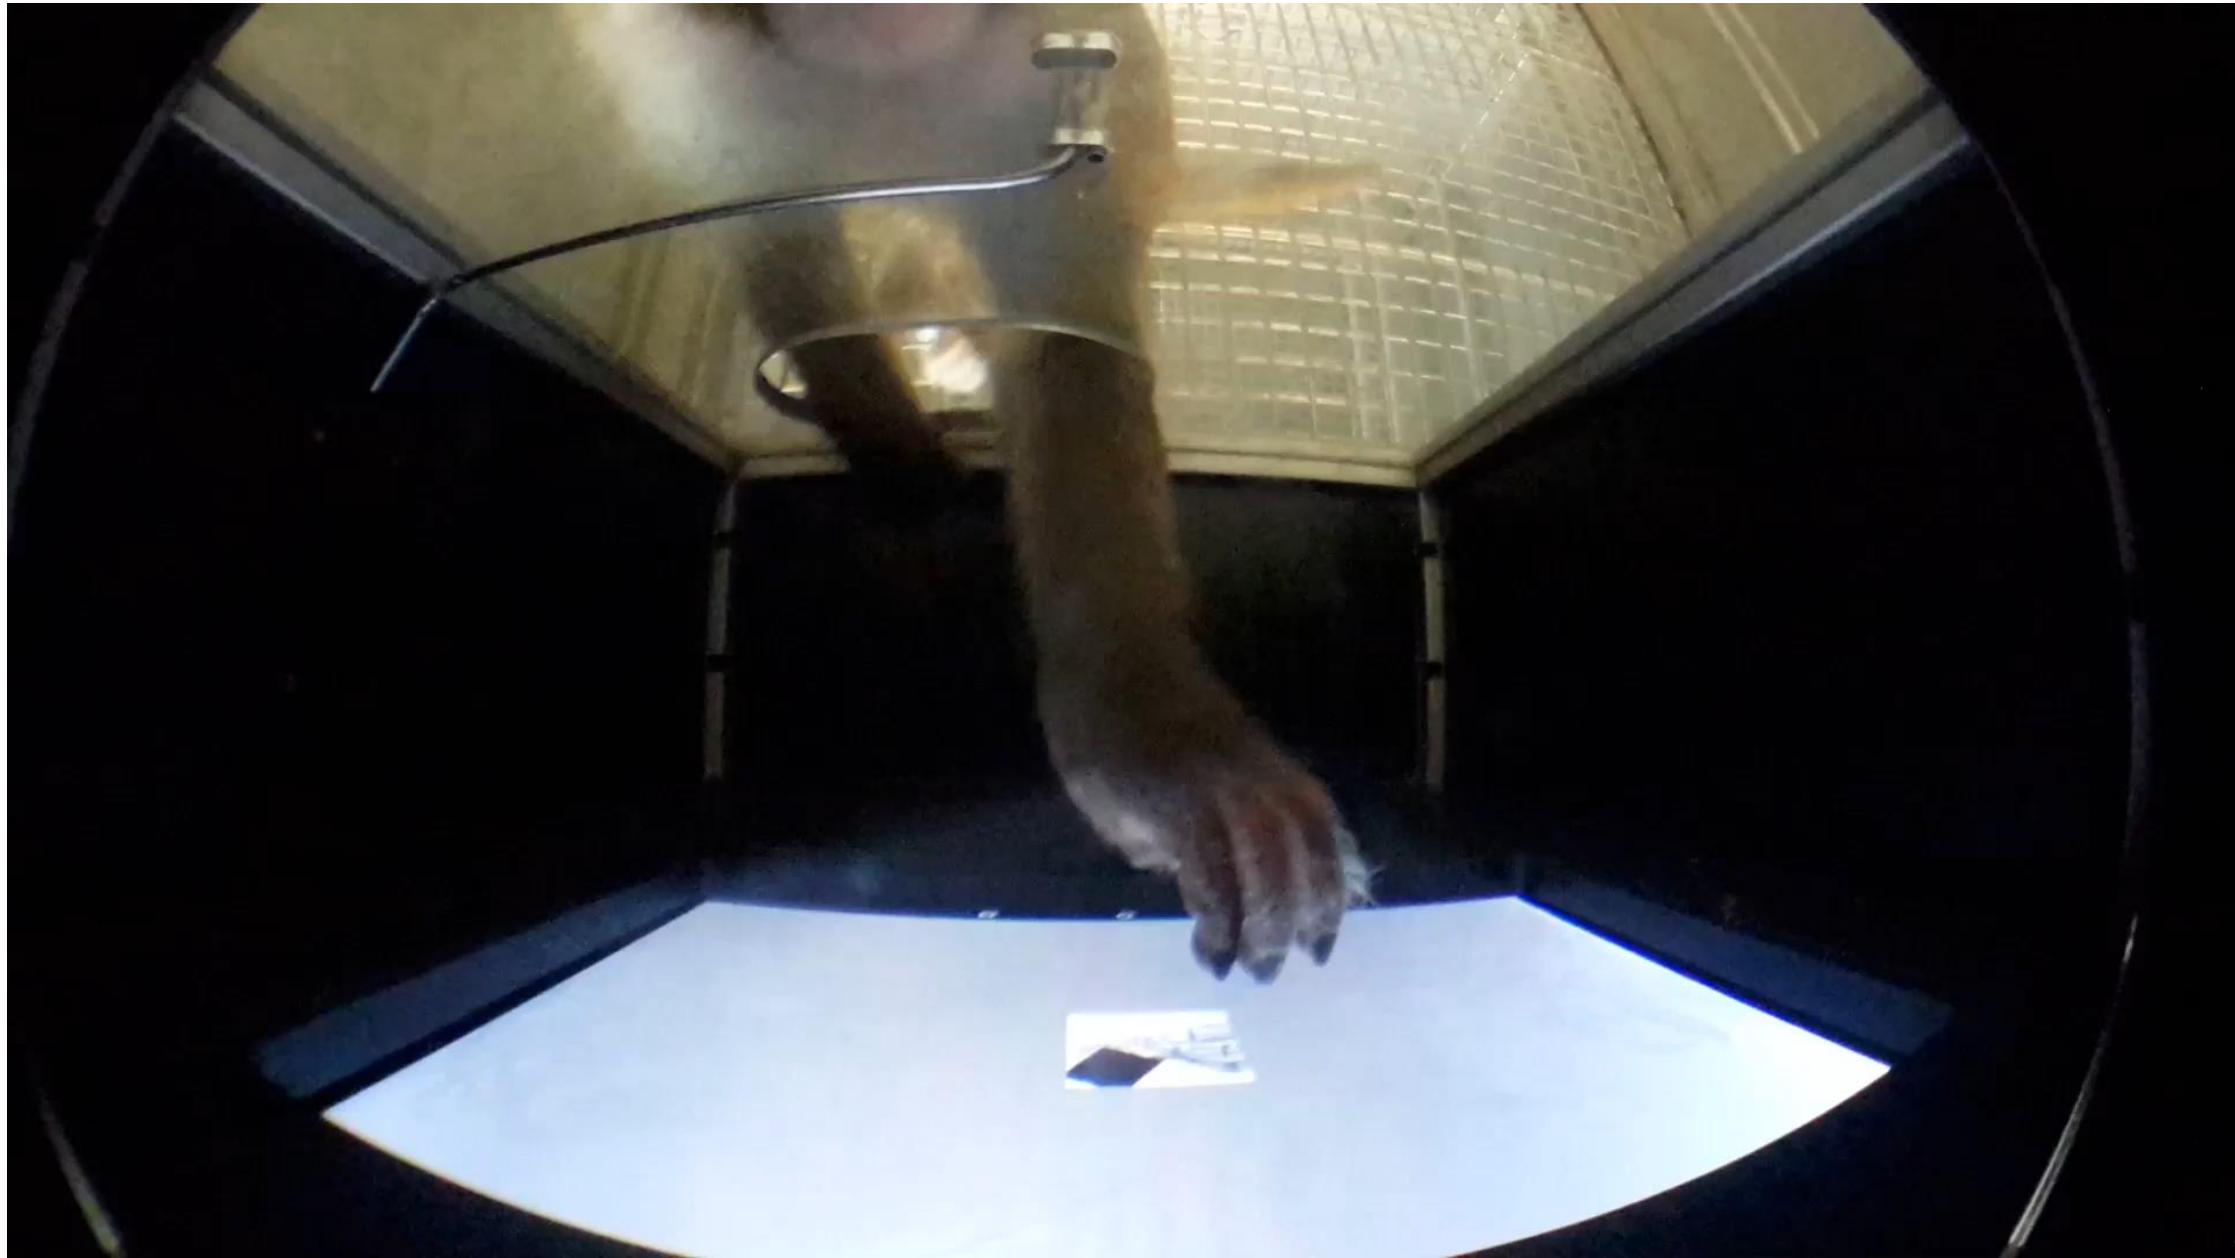


Behavioral benchmark



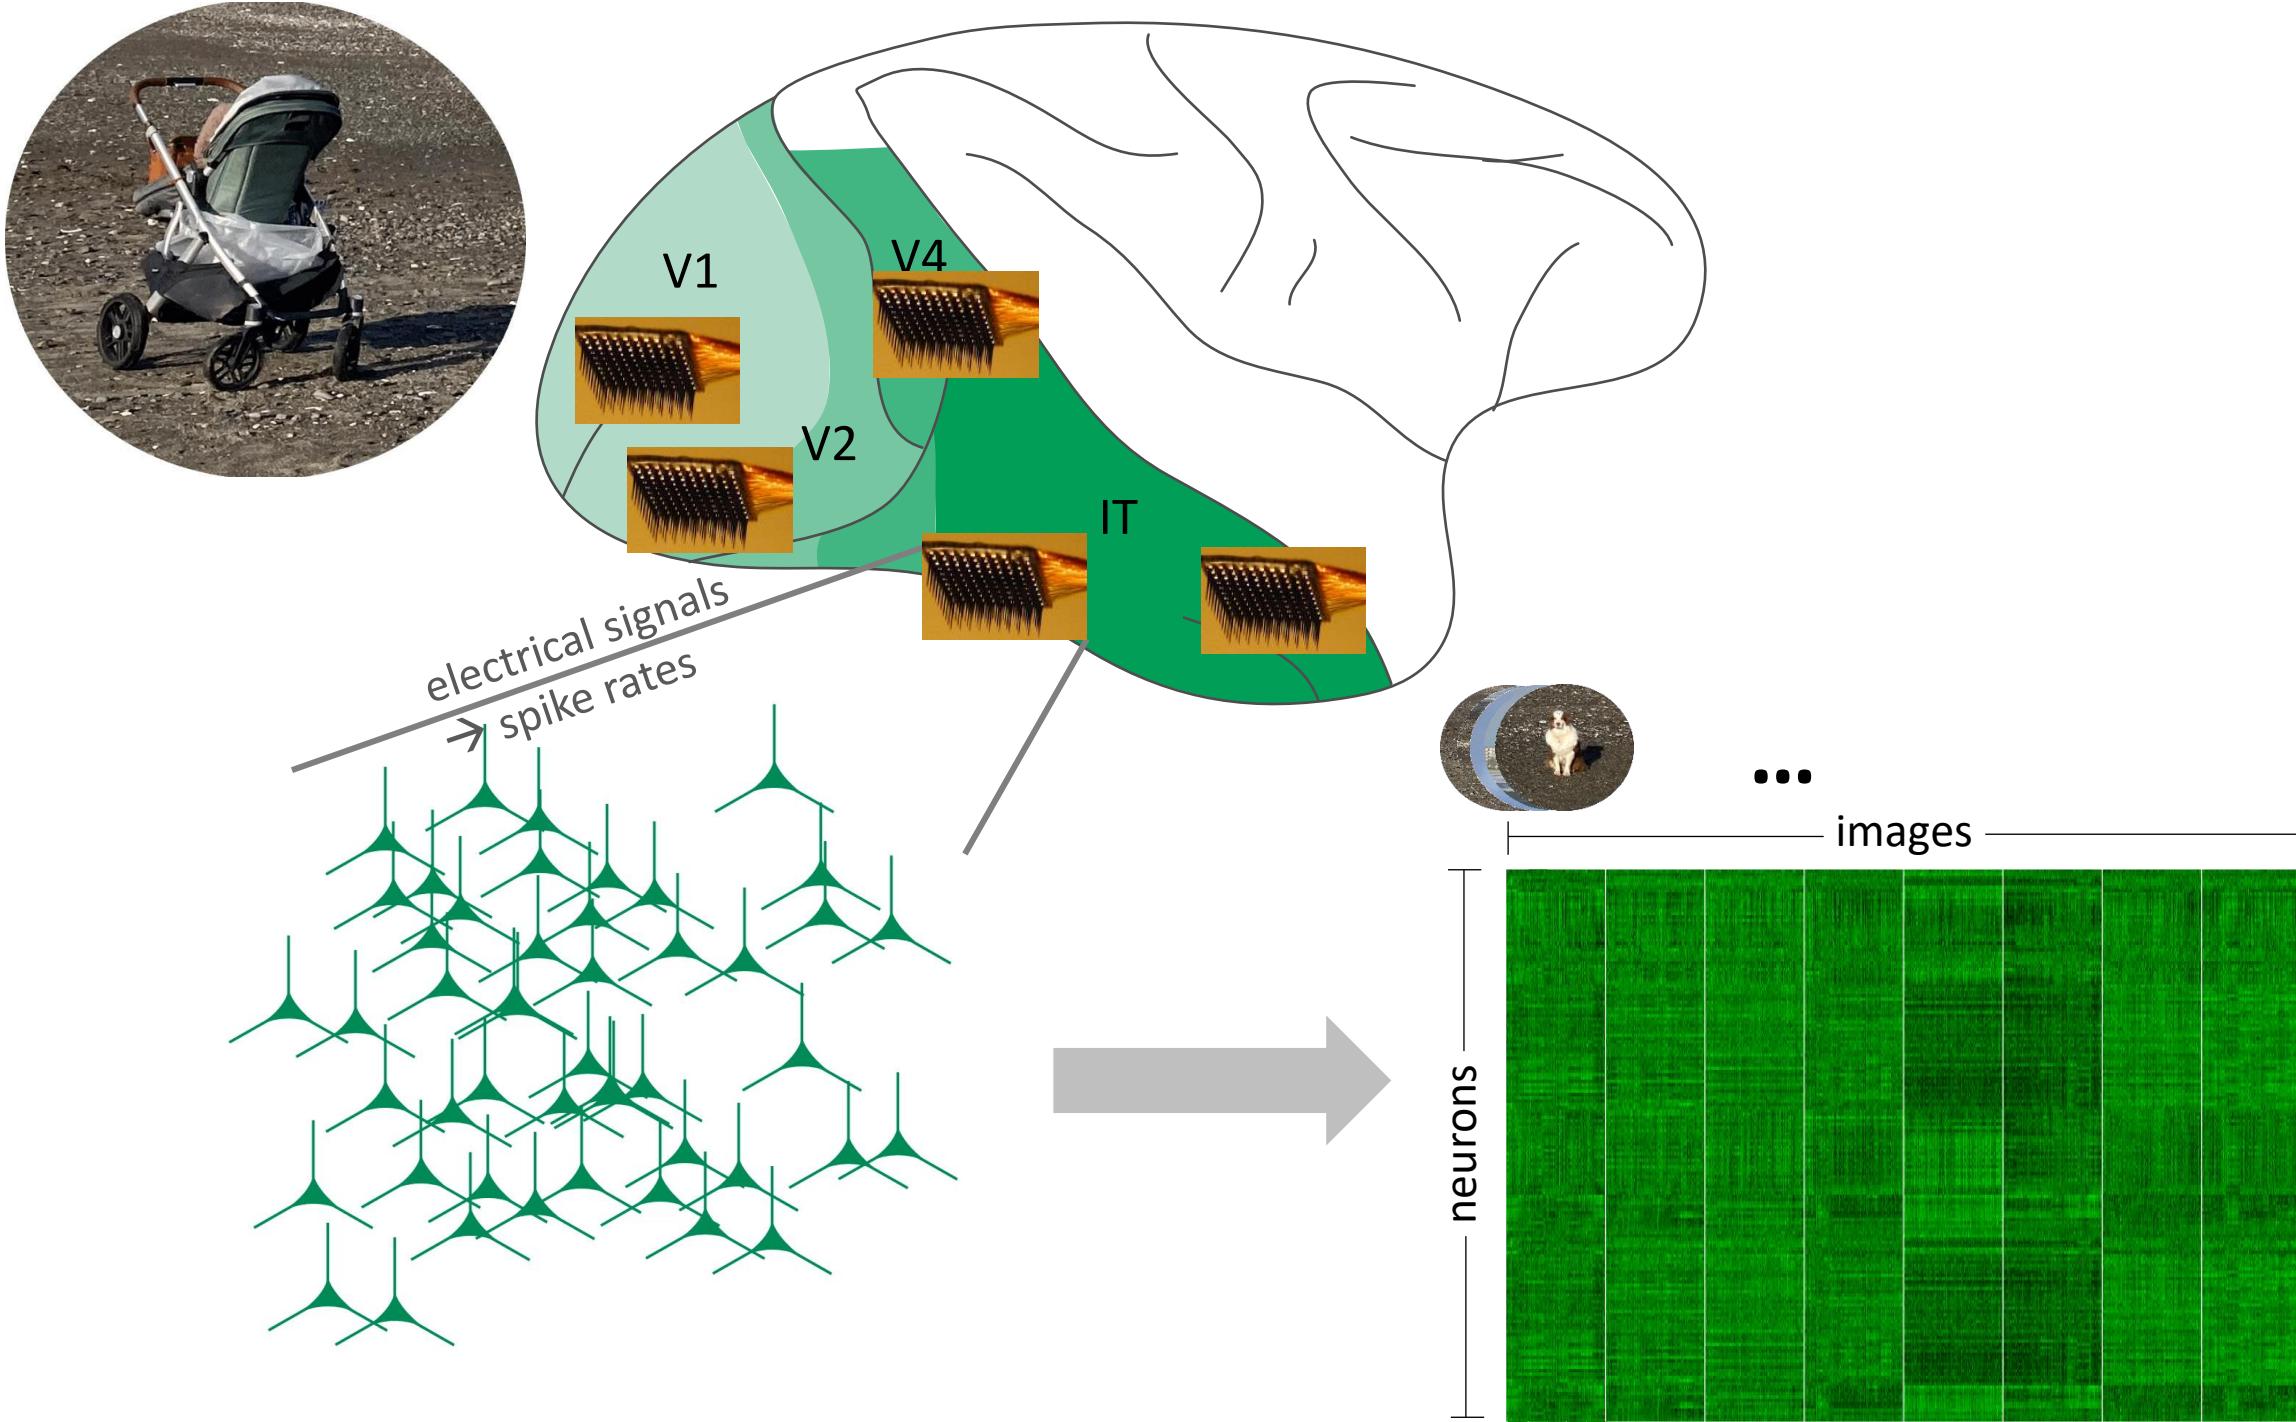
Neural benchmarks

Neural benchmarks

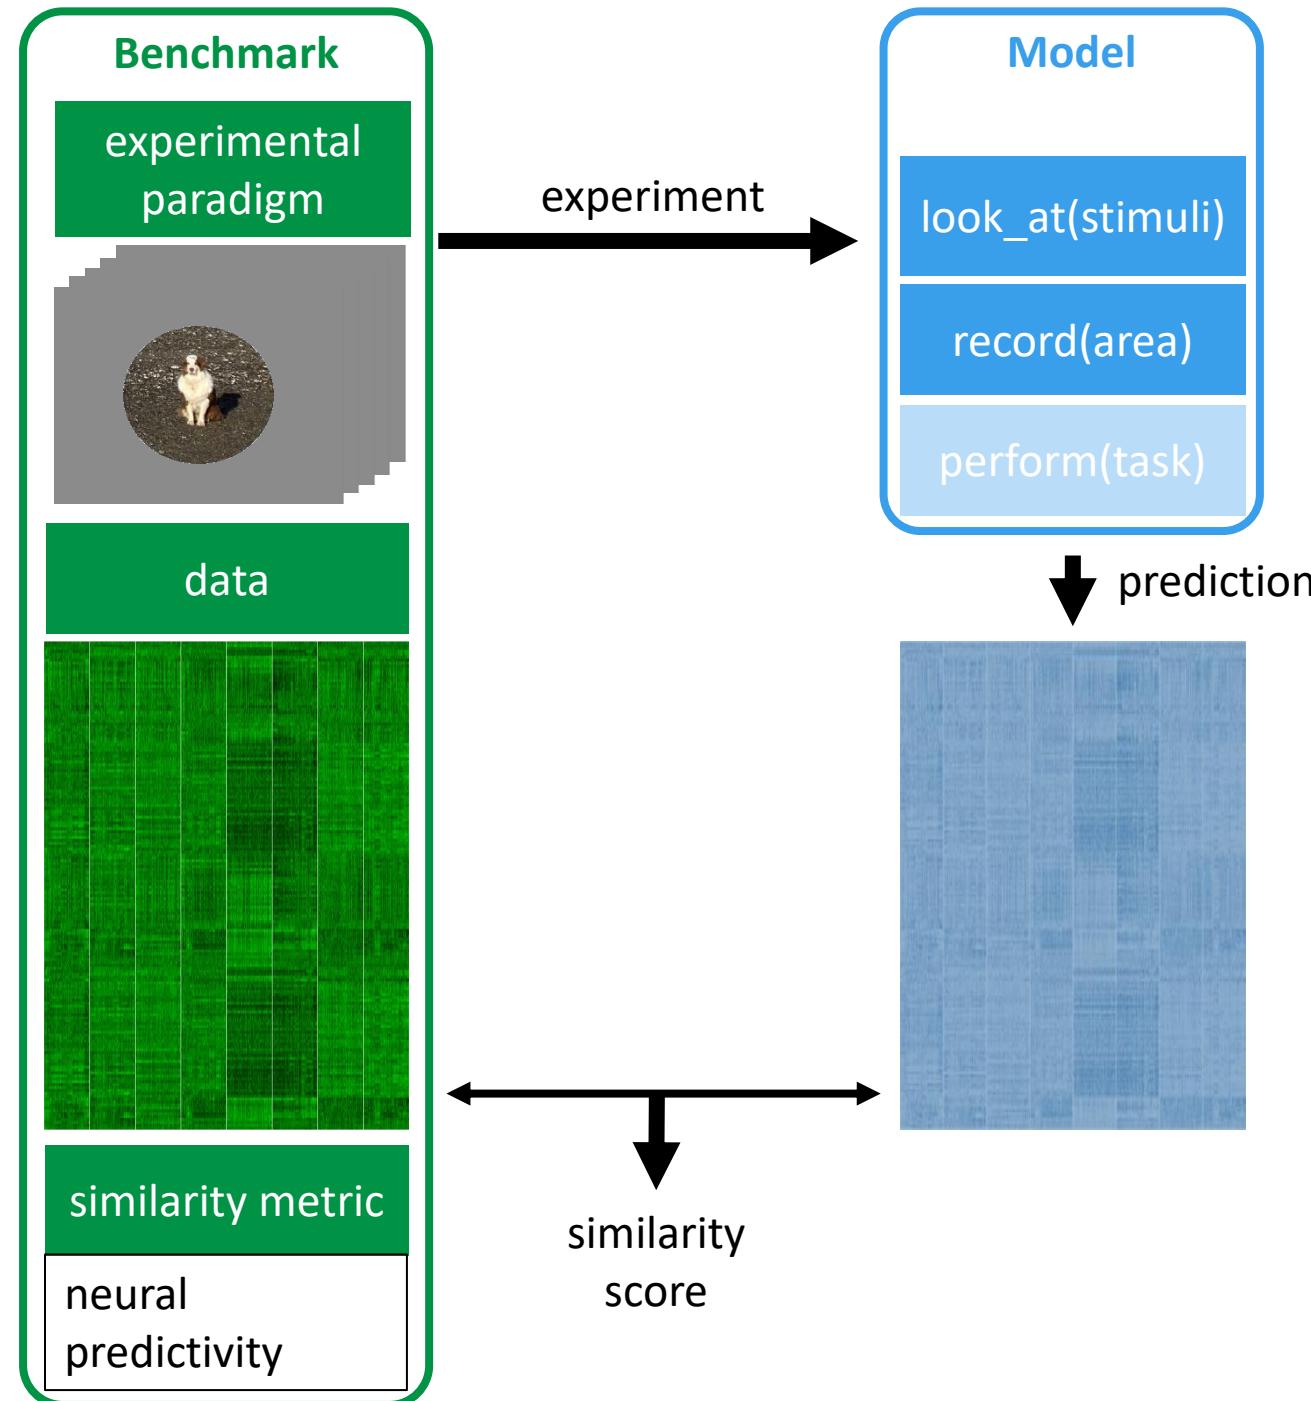


video courtesy of Kailyn Schmidt

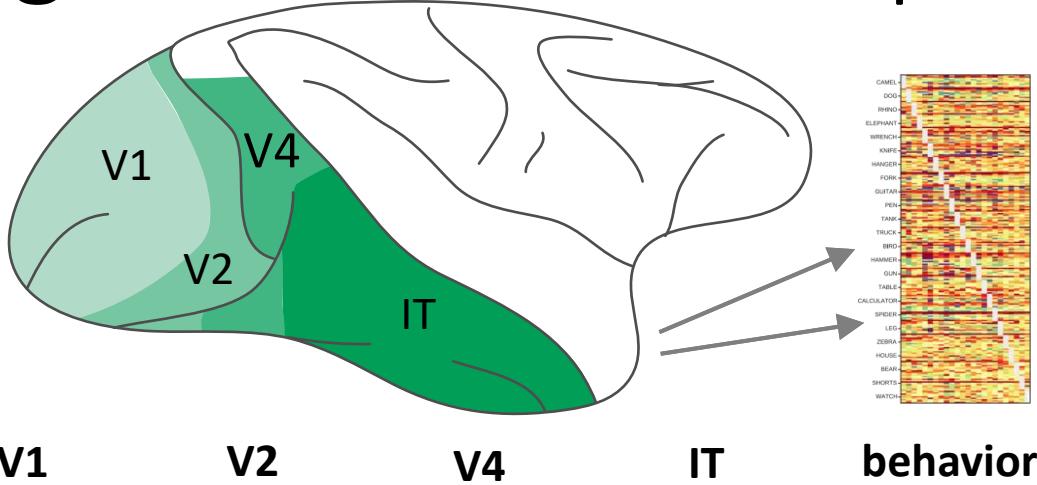
Neural benchmarks



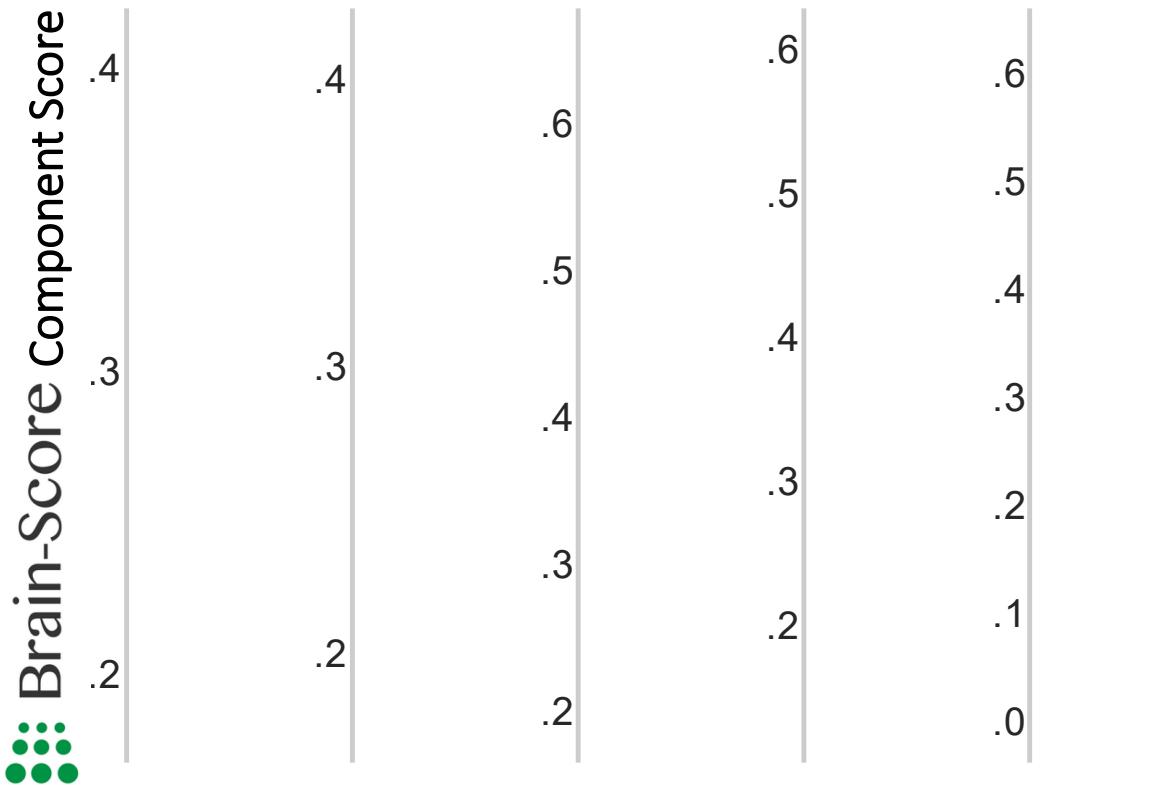
Neural benchmarks



Integrative model comparison on Brain-Score



V1 V2 V4 IT behavior



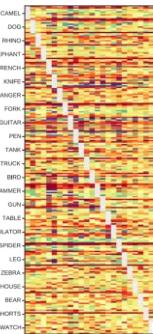
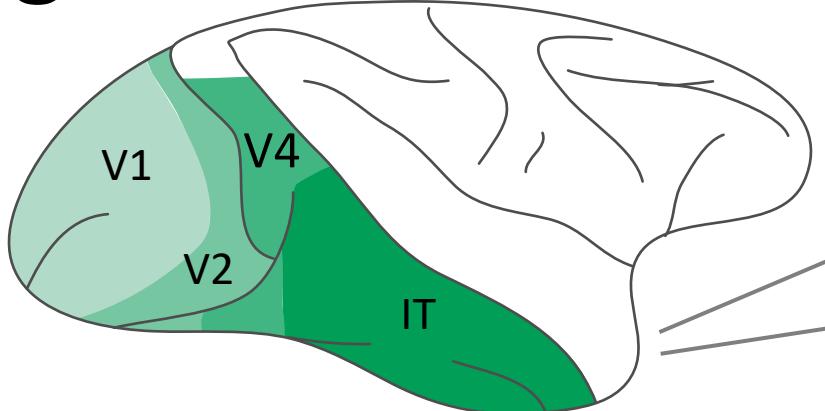
Schrimpf, Kubilius*, et al. (bioRxiv 2018)*

V1, V2 data: Freeman, Ziembra*, et al. (NatNeuro 2013)*

V4, IT data: Majaj, Hong*, et al. (JNeuro 2015)*

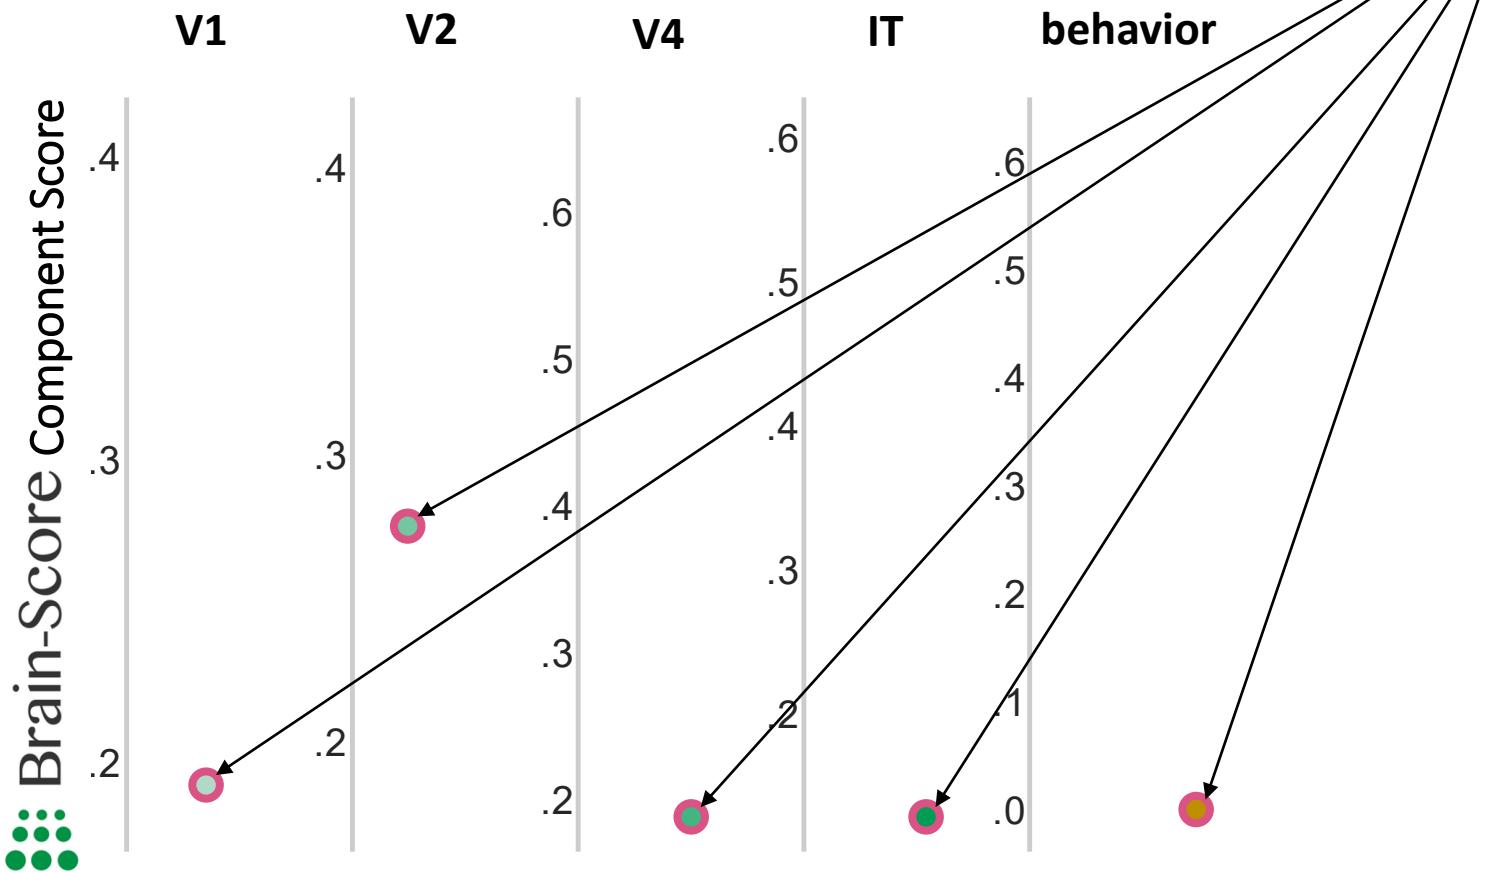
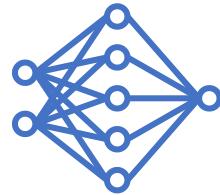
behavioral data: Rajalingham, Issa*, et al. (JNeuro 2018)*

Integrative model comparison on Brain-Score



Model candidates tested:

hmax *classic neuroscience model*



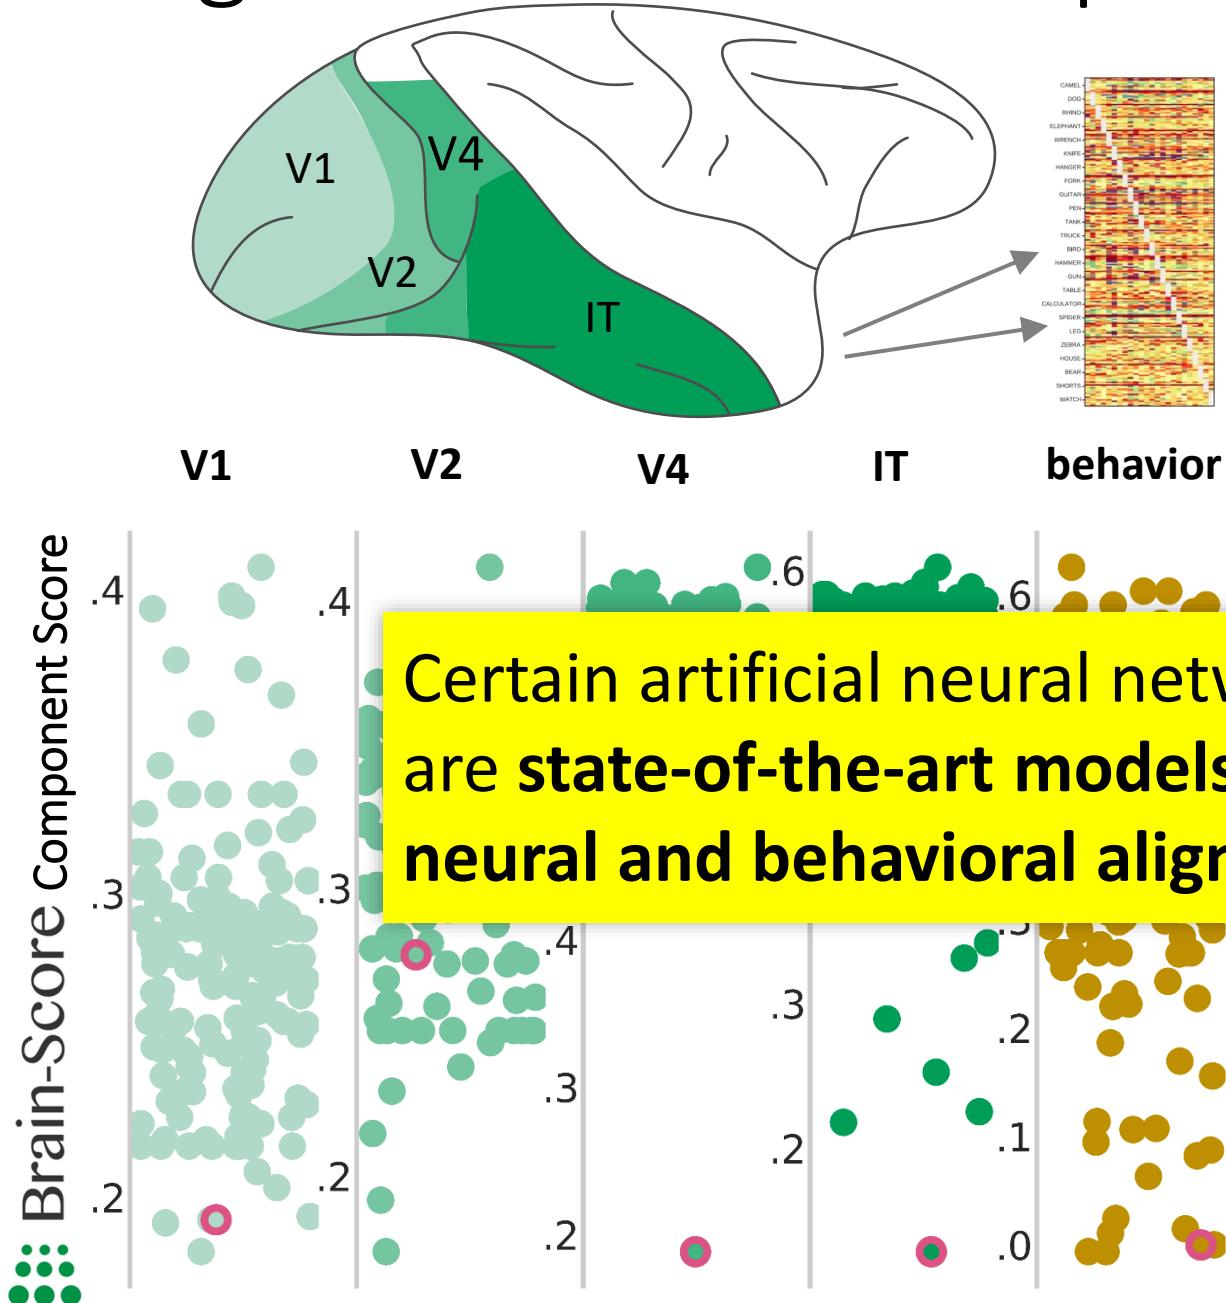
Schrimpf*, Kubilius*, et al. (bioRxiv 2018)

V1, V2 data: Freeman*, Ziembra*, et al. (NatNeuro 2013)

V4, IT data: Majaj*, Hong*, et al. (JNeuro 2015)

behavioral data: Rajalingham*, Issa*, et al. (JNeuro 2018)

Integrative model comparison on Brain-Score



Model candidates tested:

hmax
vgg-16
vgg-19
densenet-121
densenet-169
densenet-201
inception_resnet_v2
inception_v1
inception_v2
inception_v3
inception_v4
mobilenet_v1_0.25_128
mobilenet_v1_0.25_160
mobilenet_...
mobilenet_v2_1.3_224
mobilenet_v2_1.4_224

ML models



*translated into System Models:

- assign layers to regions
- assign pixels to visual degrees

Schrimpf*, Kubilius*, et al. (bioRxiv 2018)

V1, V2 data: Freeman*, Ziembra*, et al. (NatNeuro 2013)

V4, IT data: Majaj*, Hong*, et al. (JNeuro 2015)

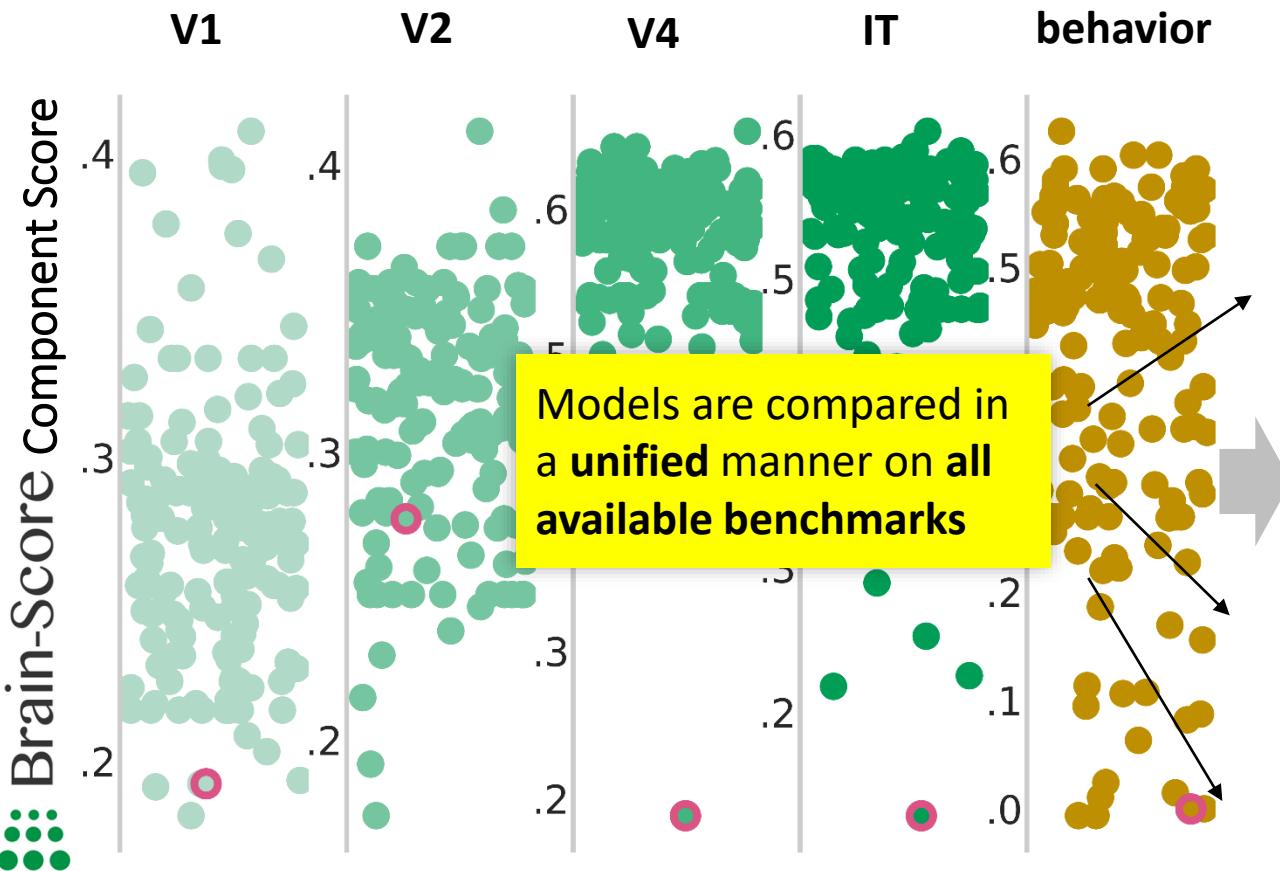
behavioral data: Rajalingham*, Issa*, et al. (JNeuro 2018)

Brain-Score: Integrative Benchmarking



Leaderboard About Compare Participate

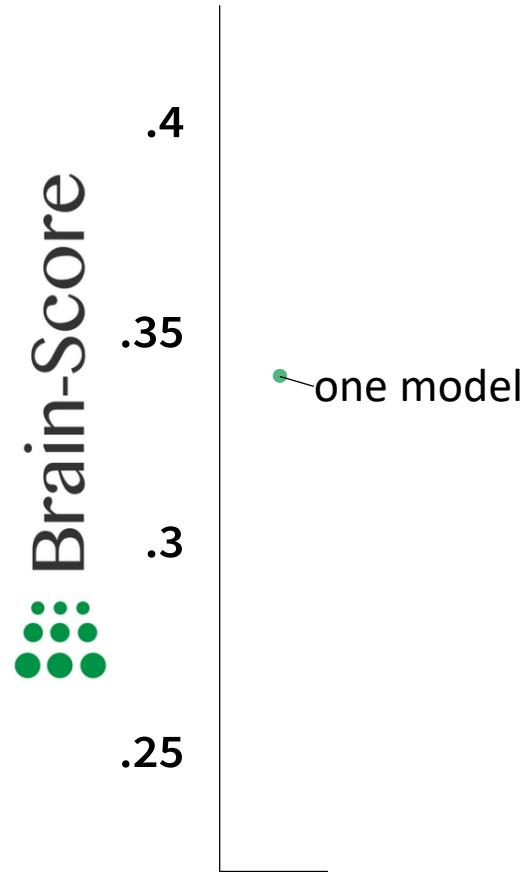
www.Brain-Score.org



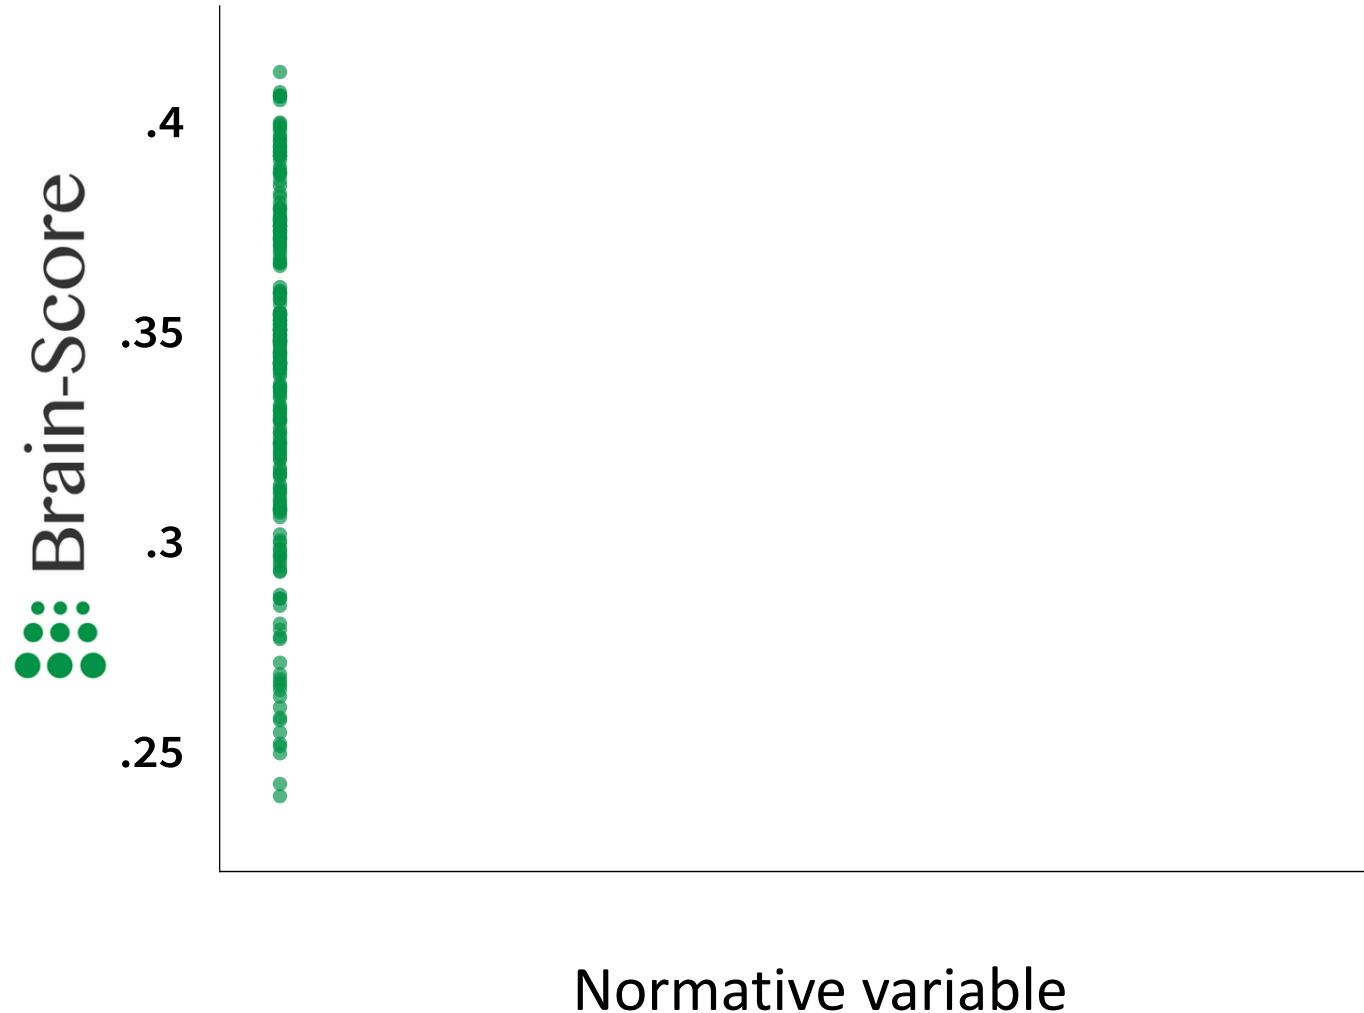
Rank

	Model submitted by
1	fnetb1_cutmixpatch_augmix_robust Alexander Riedel
2	resnext101_32x8d_wsl Martin Schrimpf
3	snet50_finetune_cutmix_e3_robust Alexander Riedel
4	effnetb1_272x240 Alexander Riedel
5	ustom_model_cv_18_dagger_408 William Berrios
6	resnet-152_v2 Brain-Score Team
7	voneresnet-50-non_stochastic Tiago Marques
8	pnasnet_large Brain-Score Team
9	resnet-152_v1 Brain-Score Team
10	AdvProp_efficientnet-b6 Joel Dapello
11	



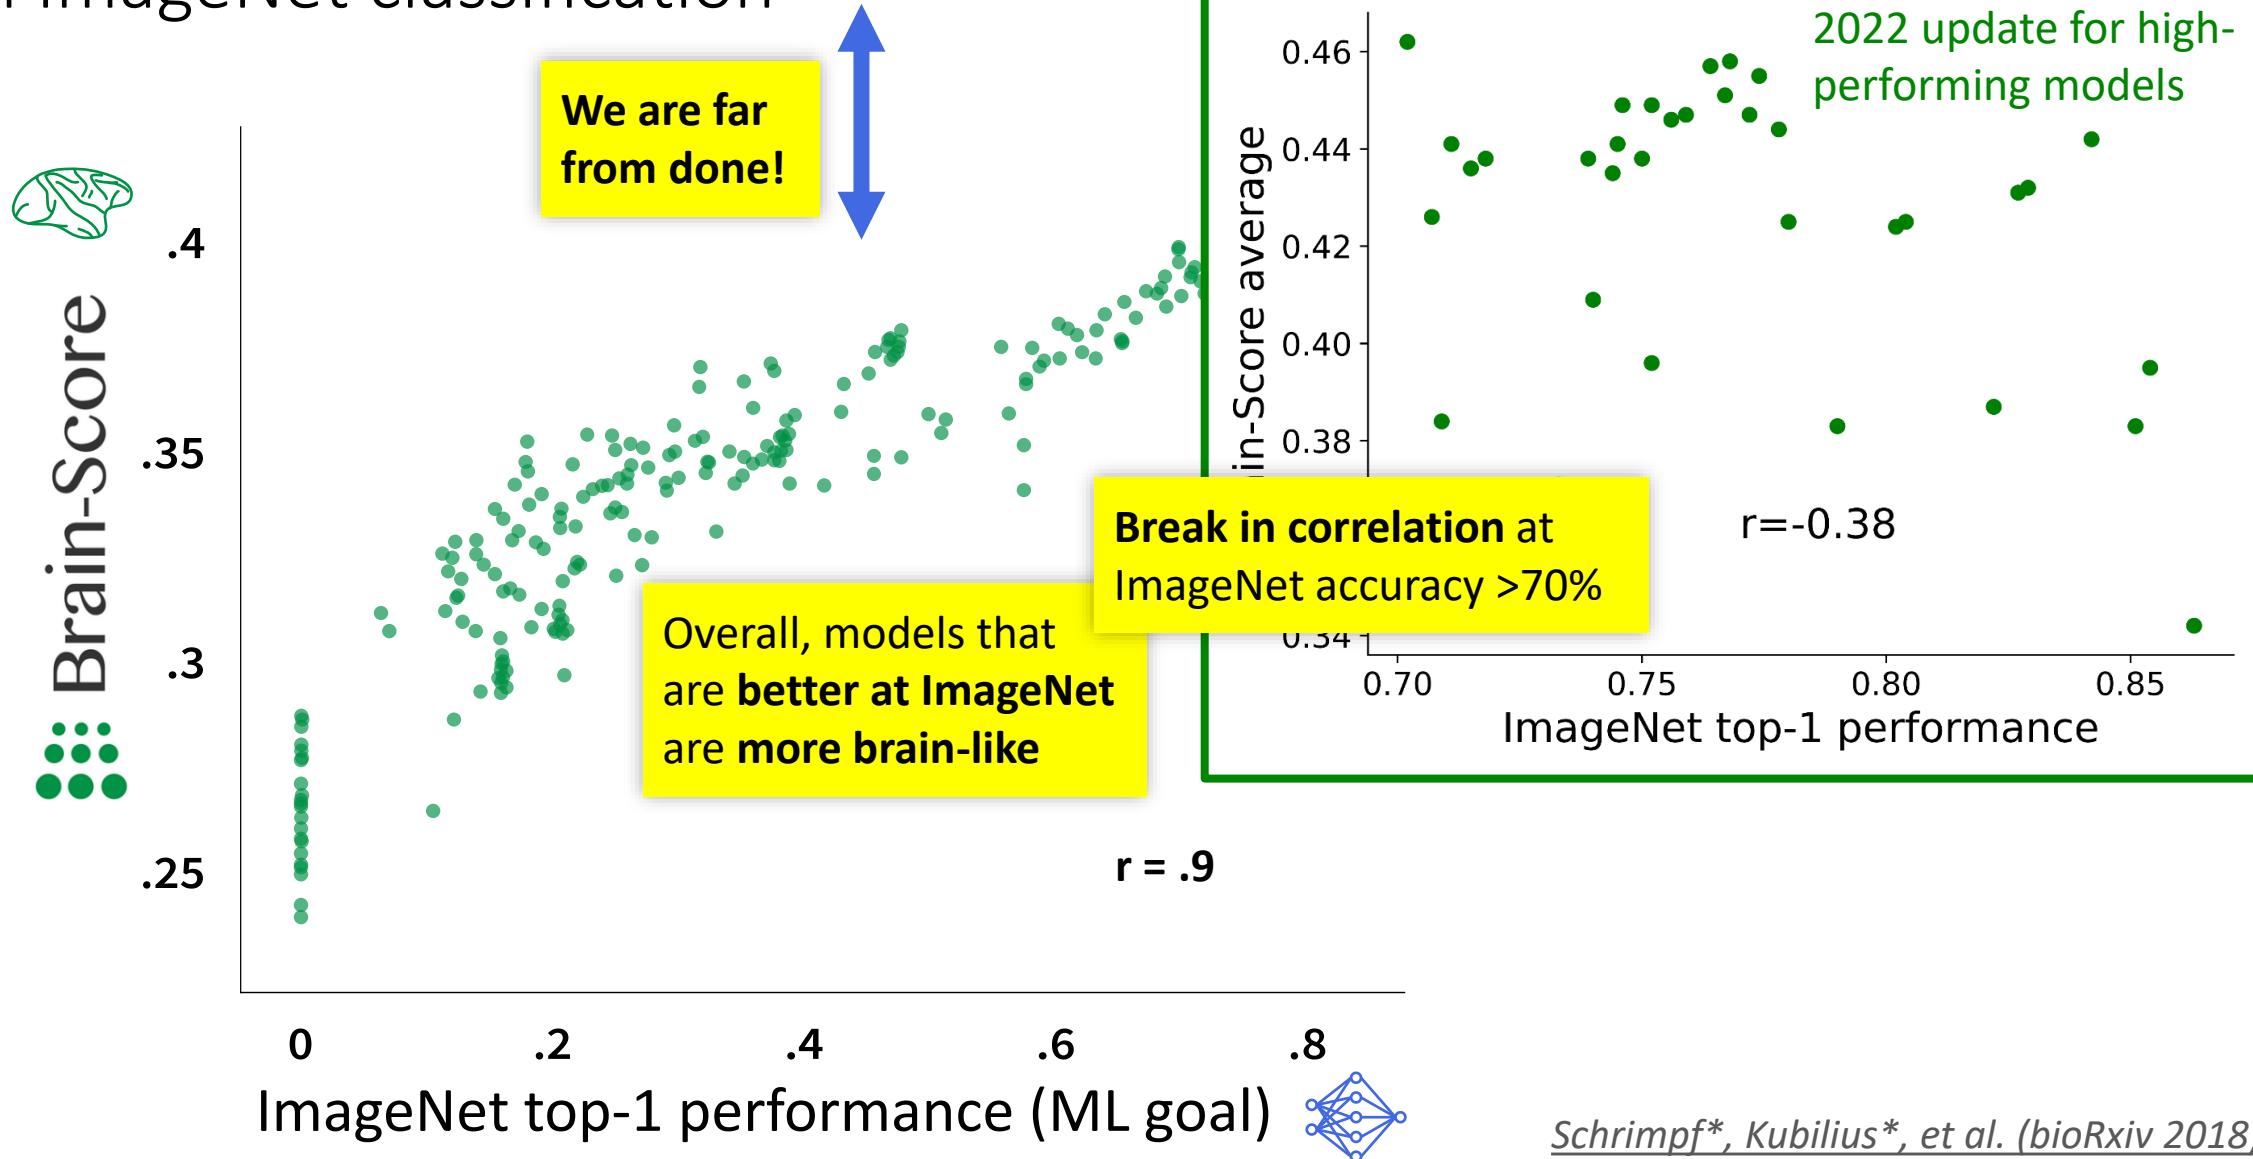


What explains the model differences?



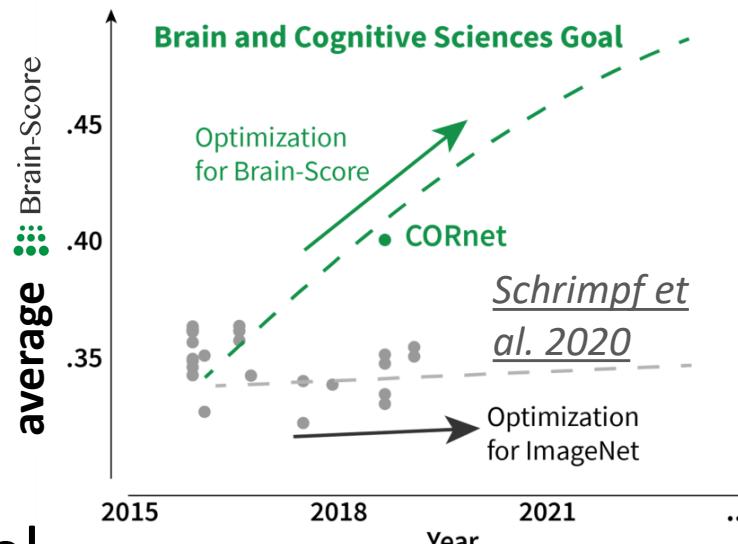
cf. Yamins, Hong*, et al. (PNAS 2014)*
Schrimpf, Kubilius*, et al. (bioRxiv 2018)*

Task performance correlates with Brain-Score on ImageNet classification



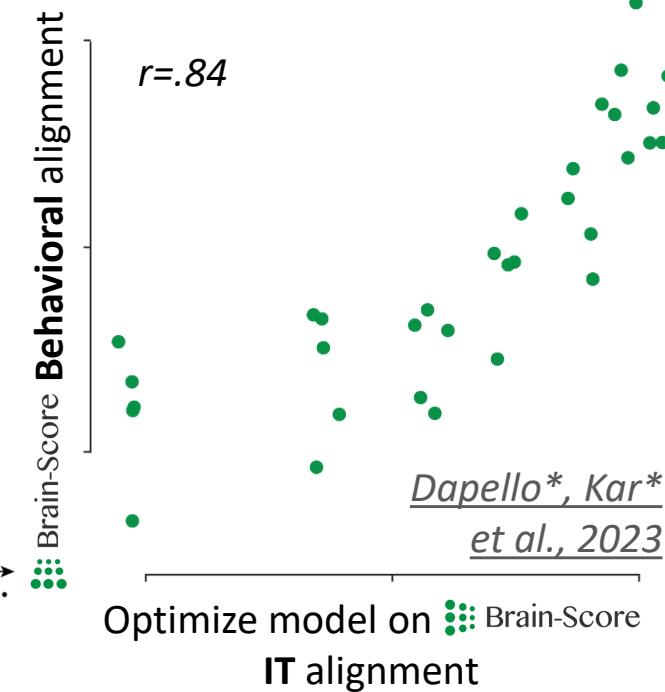
What Brain-Score enables

1 Track and guide progress of modeling primate vision



2 Relate brain benchmarks to one another.

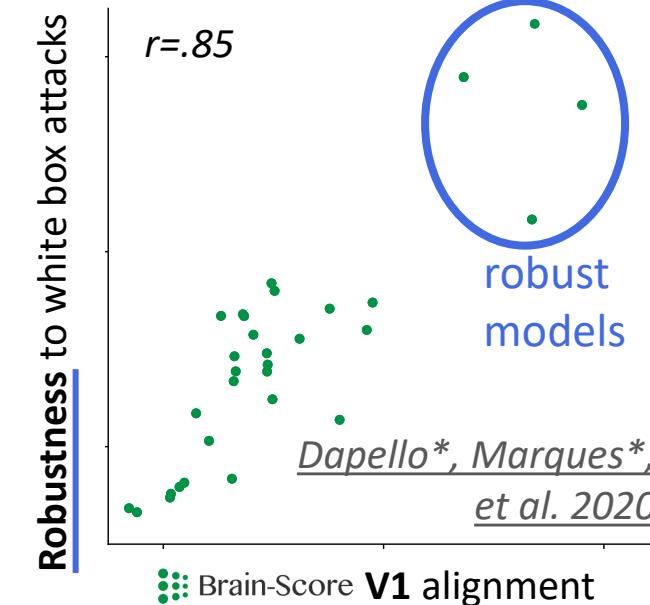
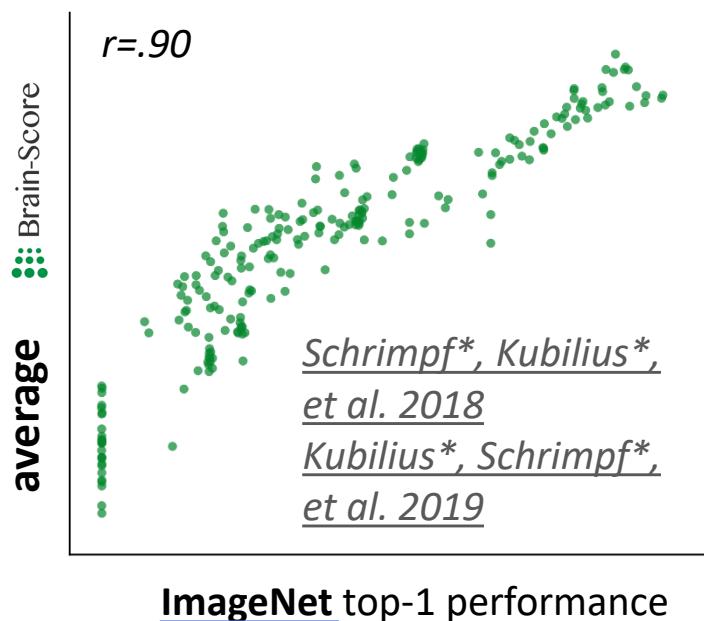
Improve IT component of model
→ improve behavior



3 Relate brain benchmarks to engineering desiderata

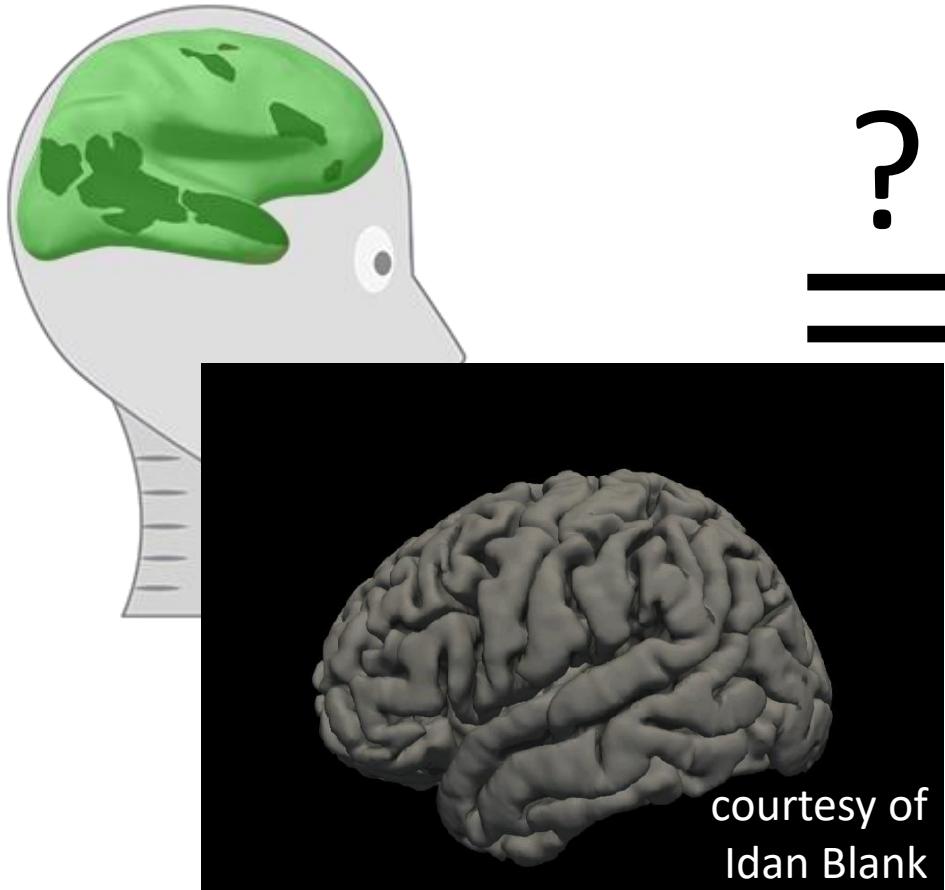
ImageNet → Brain-Score;

More V1-like
→ improved robustness

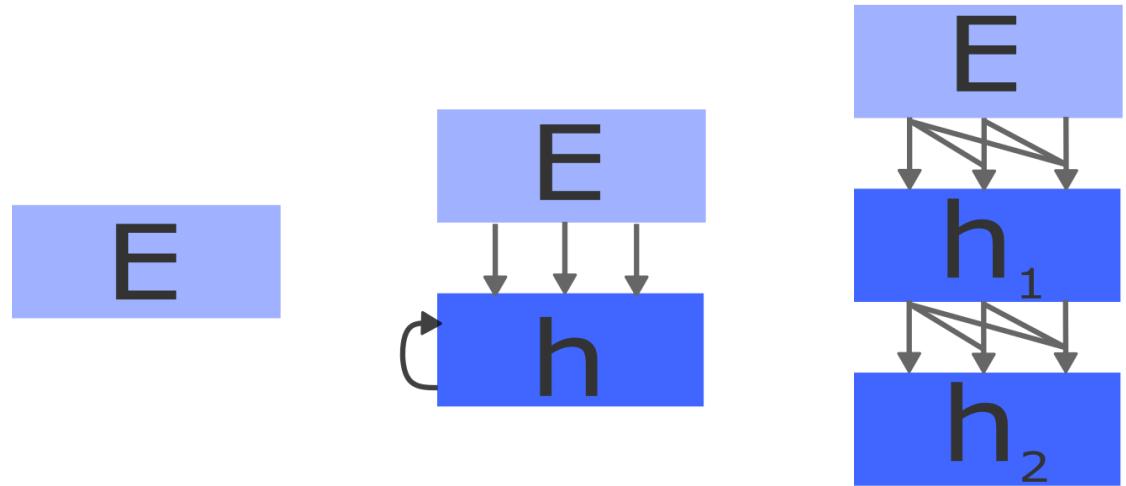


Integrative testing yields insights across domains of intelligence such as language

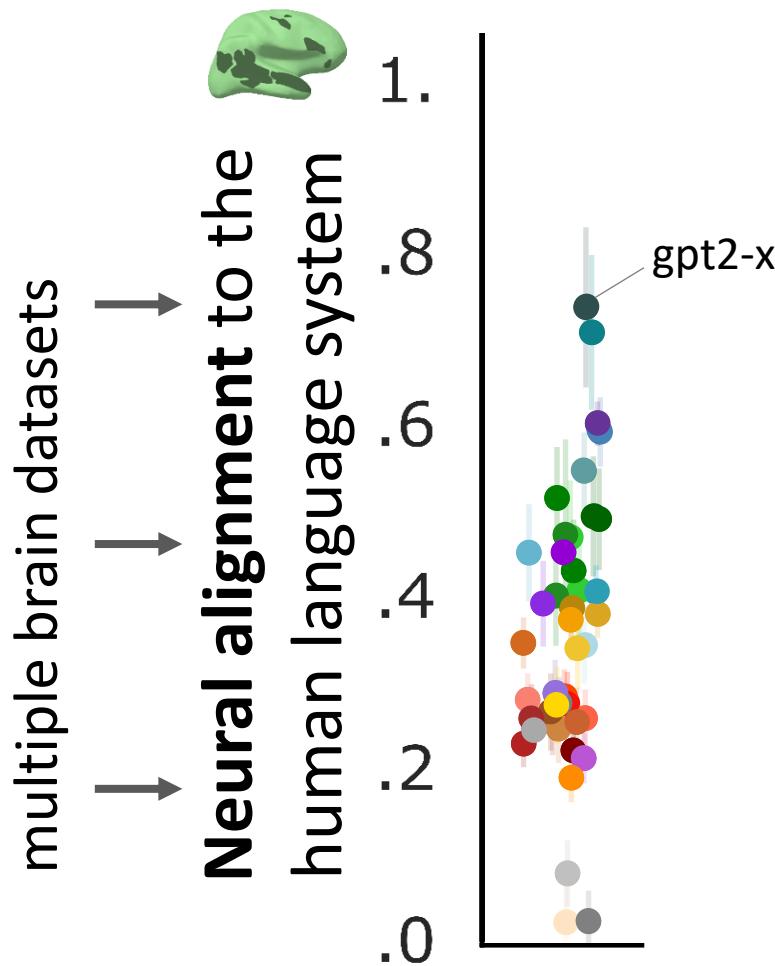
Humans



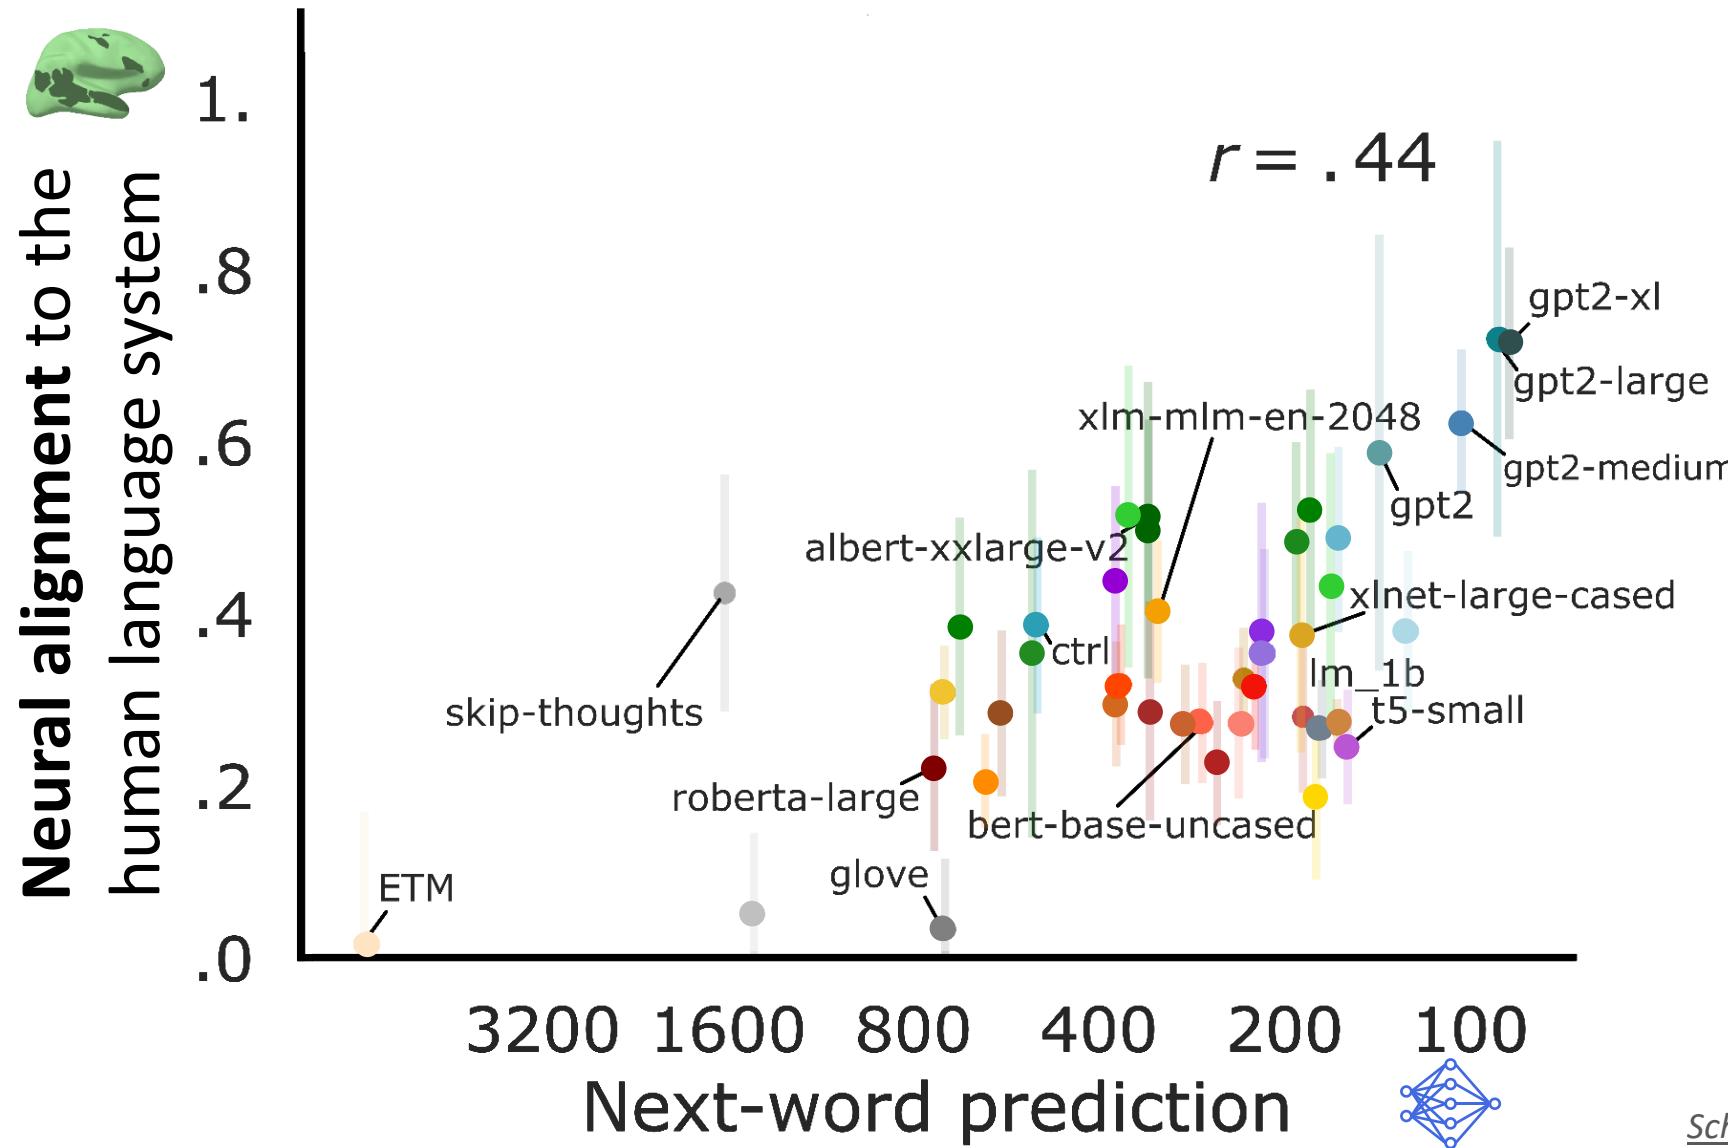
Models



Particular ML language models predict the human language system



The better models can predict the next word,
the more brain-like they are



Contributions



To meaningfully **make progress** in modeling human intelligence (behavior + underlying neural activity), we must *integrate experimental results*.

Brain-Score is an implementation of this approach, currently in the domains of vision and language.

- *Identify* most brain-like models
- Provide *empirical constraints* for developing new models
- Discover key *relationships* between neural function, behavior, and computation

