The functional specialization of visual cortex emerges from training parallel pathways with self-supervised predictive learning

neuro
Institut-Hôpital neurologique de Montréal
Montreal Neurological Institute-Hospital

Mila

McGill

LiNC LAB

Blake Richards
CIFAR

January 12th, 2022
Dynamical Principles of Biological and Artificial Neural Networks

“What” versus “where” pathways in the brain

Dorsal pathway:
Where are things, and
where are they moving?

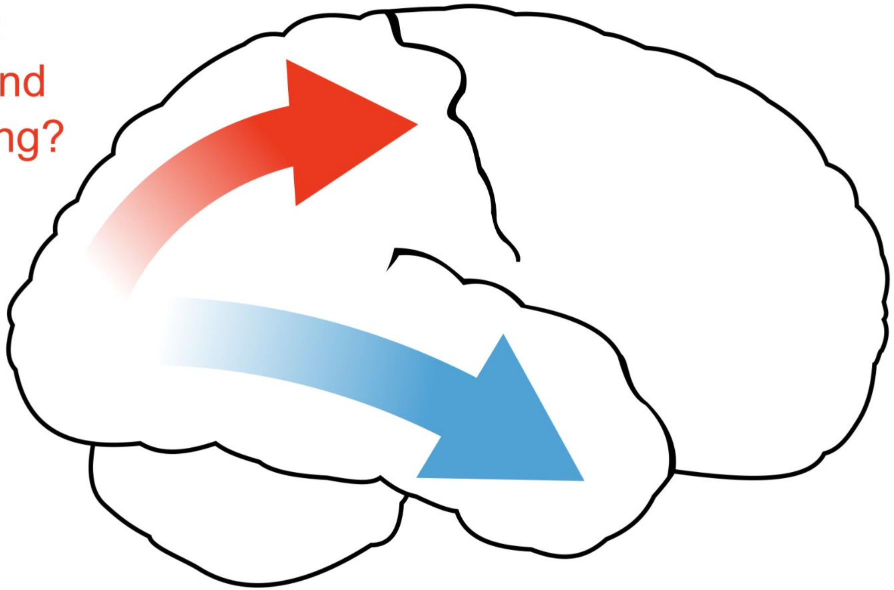

Ventral pathway:
What things are there, and
what do they mean?

Why does the brain have these specialized pathways?

Visual prediction requires two different, competing forms of invariance

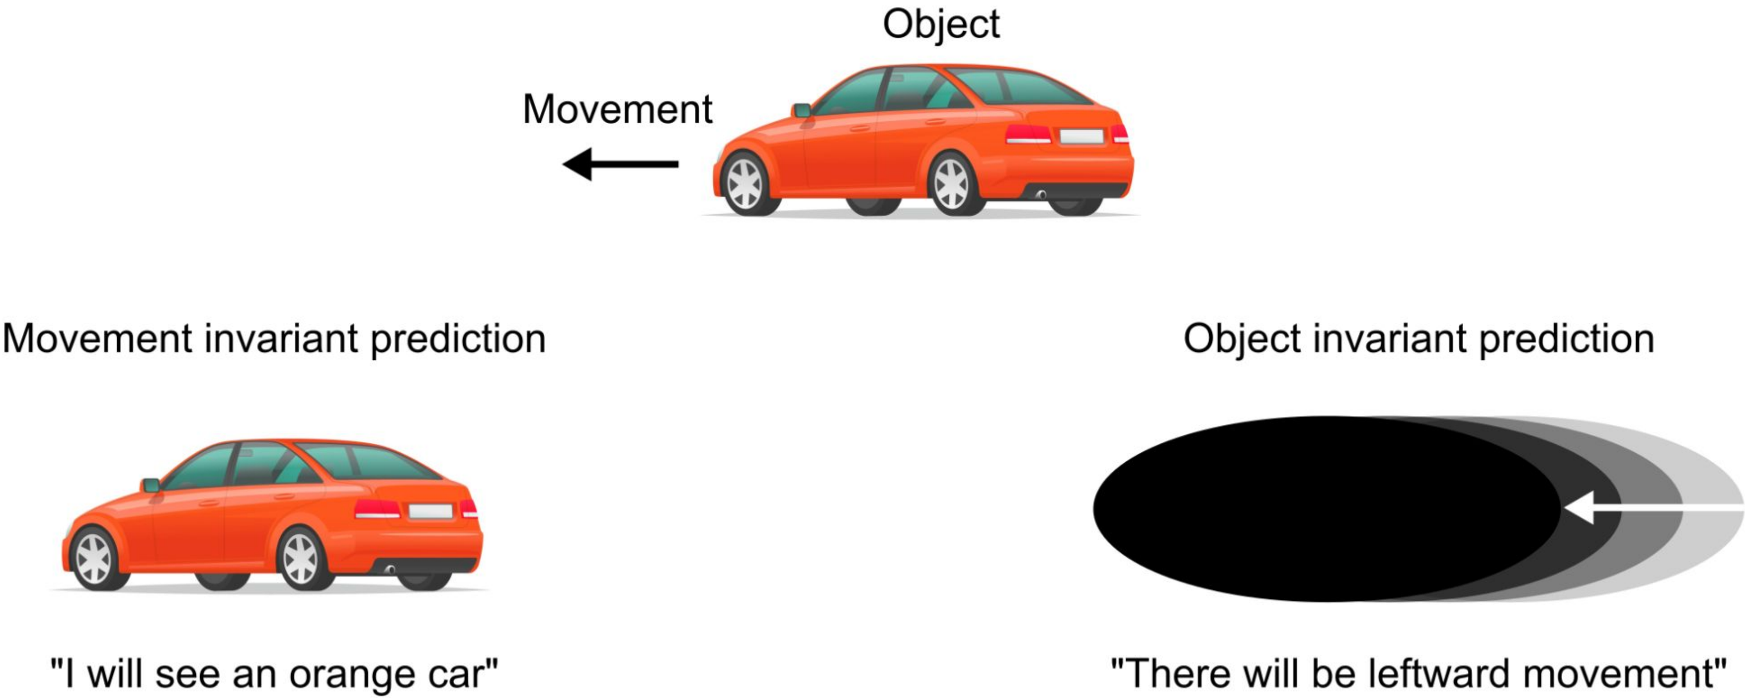

Hypothesis:

The "what" versus "where" specialization in the emerges from optimization for visual predictions

Approach:

Train ANNs with a self-supervised predictive loss, do they develop representations similar to those in "what" and "where" pathways in the brain?

Preview:
Self-supervised learning (with some anatomical segregation), but not supervised learning, induces "what" and "where" specialized pathways

Paper:
S Bakhtiari, PJ Mineault, TP Lillicrap, CC Pack and BA Richards
Neural Information Processing Systems 2021 (spotlight)

Neural data: Allen Brain Observatory

de Vries et al. Nature Neuro, 2019:
“Natural movies elicited responses from the most neurons”

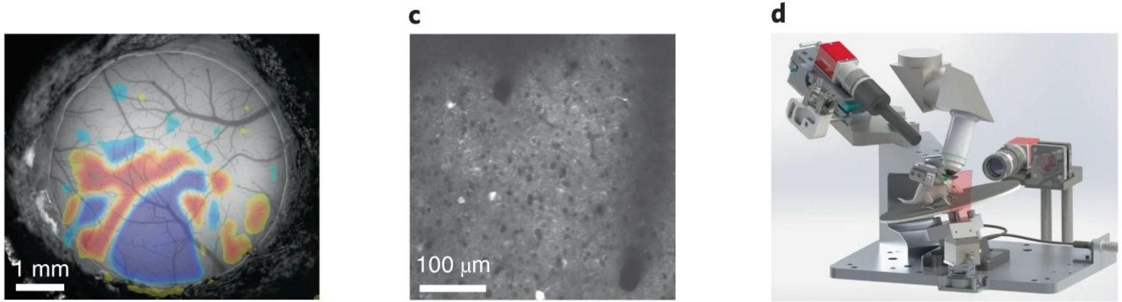

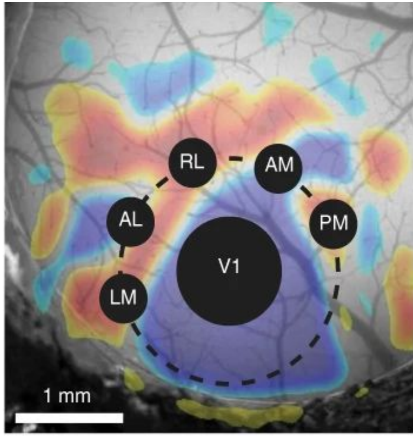

Comparison tool: Representation Similarity Analysis

r_i : response of neuron *i*

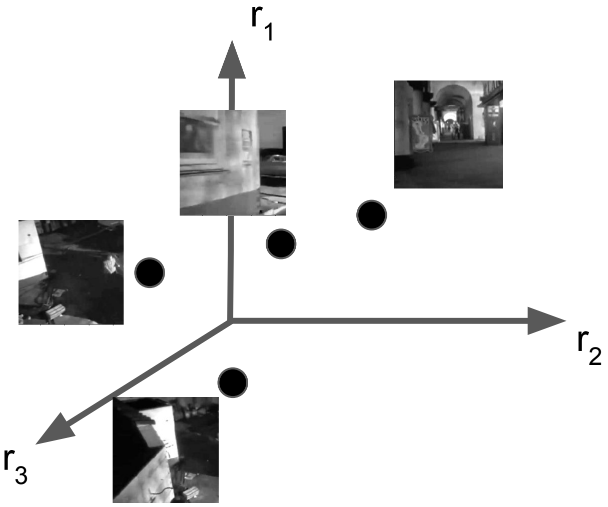

Comparison tool: Representation Similarity Analysis

r_i : response of neuron i

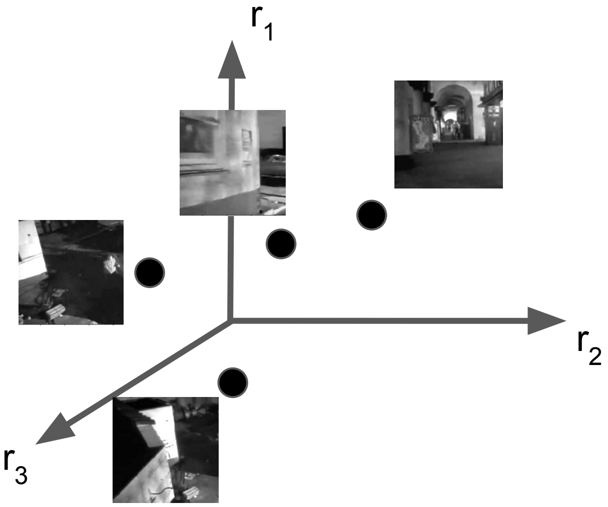

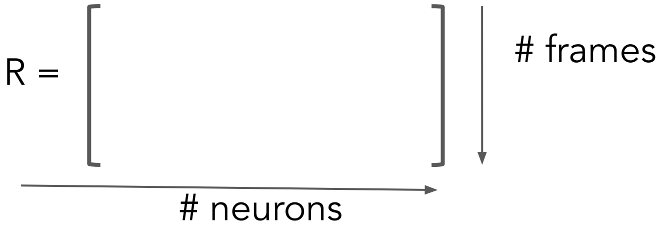

Comparison tool: Representation Similarity Analysis

r_i : response of neuron i

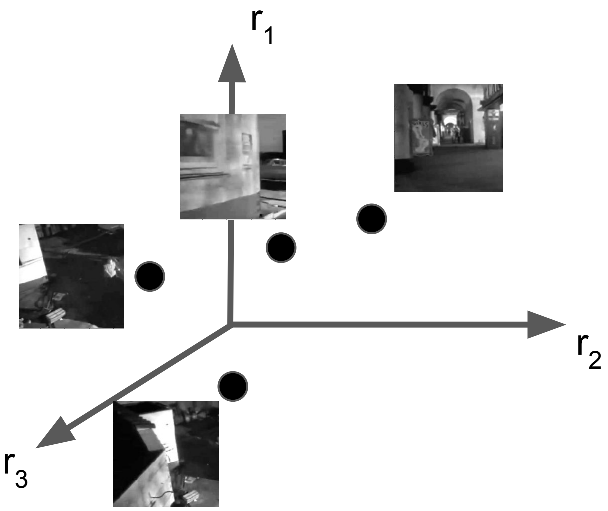

$$R = \begin{bmatrix} \quad \end{bmatrix}$$

frames

neurons

e.g. RR^T

$$RSM = \begin{bmatrix} \quad \end{bmatrix}$$

frames

frames

“Ventral” versus “Dorsal” pathways in mouse cortex

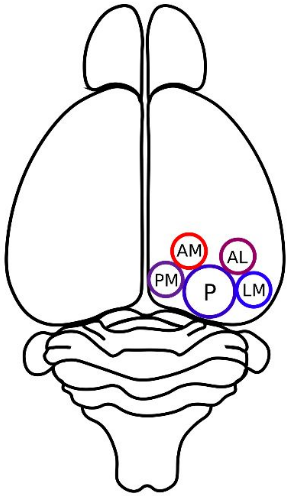

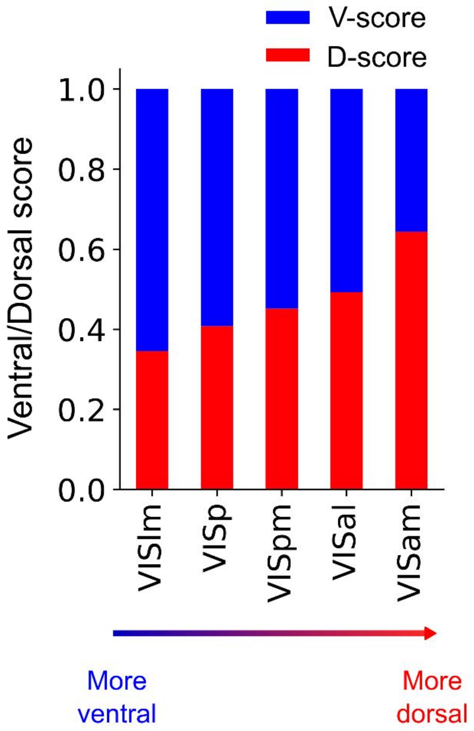

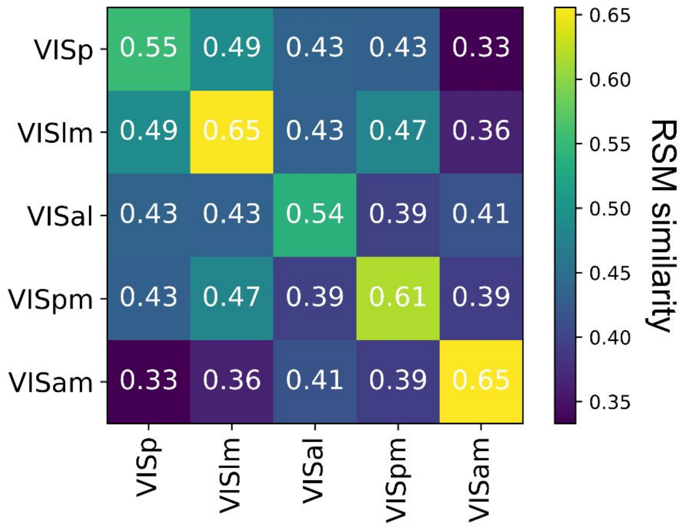

Loss: Contrastive Predictive Coding (CPC)

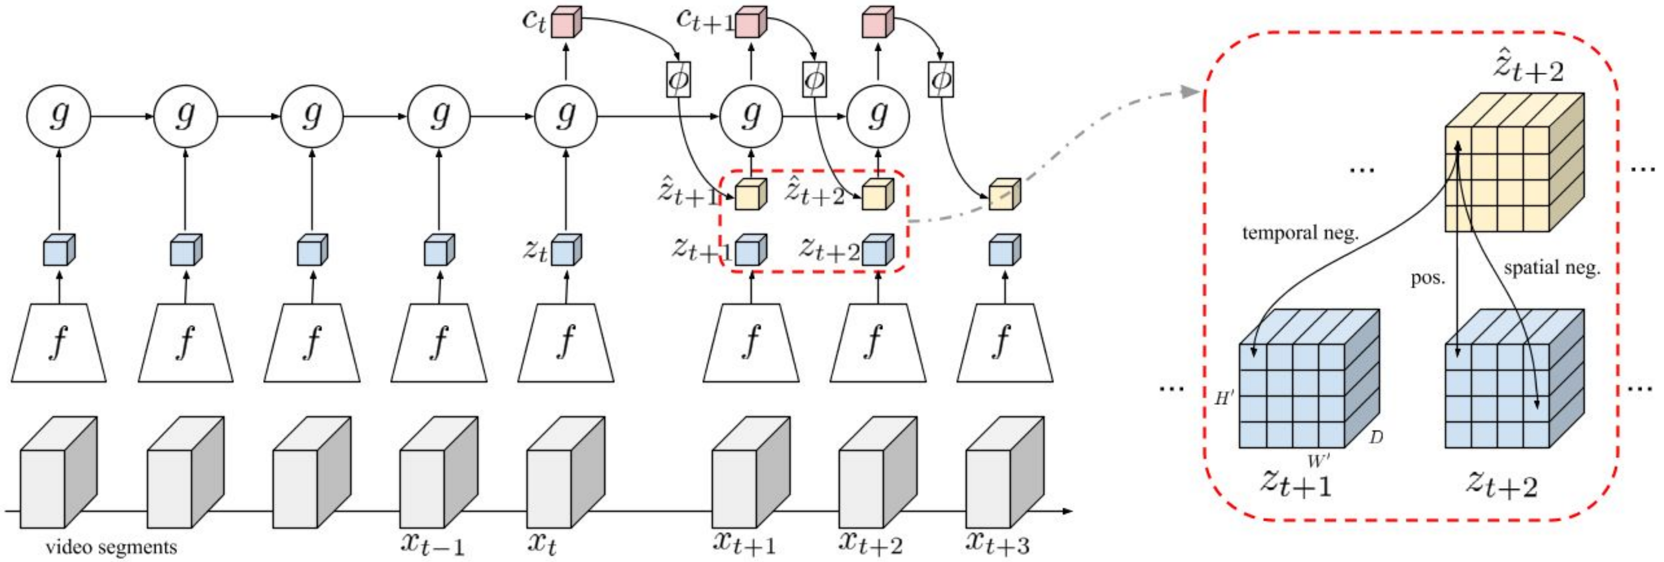

$$\mathcal{L} = -\sum_{i,k}\left[\log \frac{\exp(\hat{z}_{i,k}^{\top} \cdot z_{i,k})}{\sum_{j,m} \exp(\hat{z}_{i,k}^{\top} \cdot z_{j,m})}\right]$$

Han et al. (2019), *ICCVW*

Training data: UCF101

Training a single pathway induces ventral-like representations

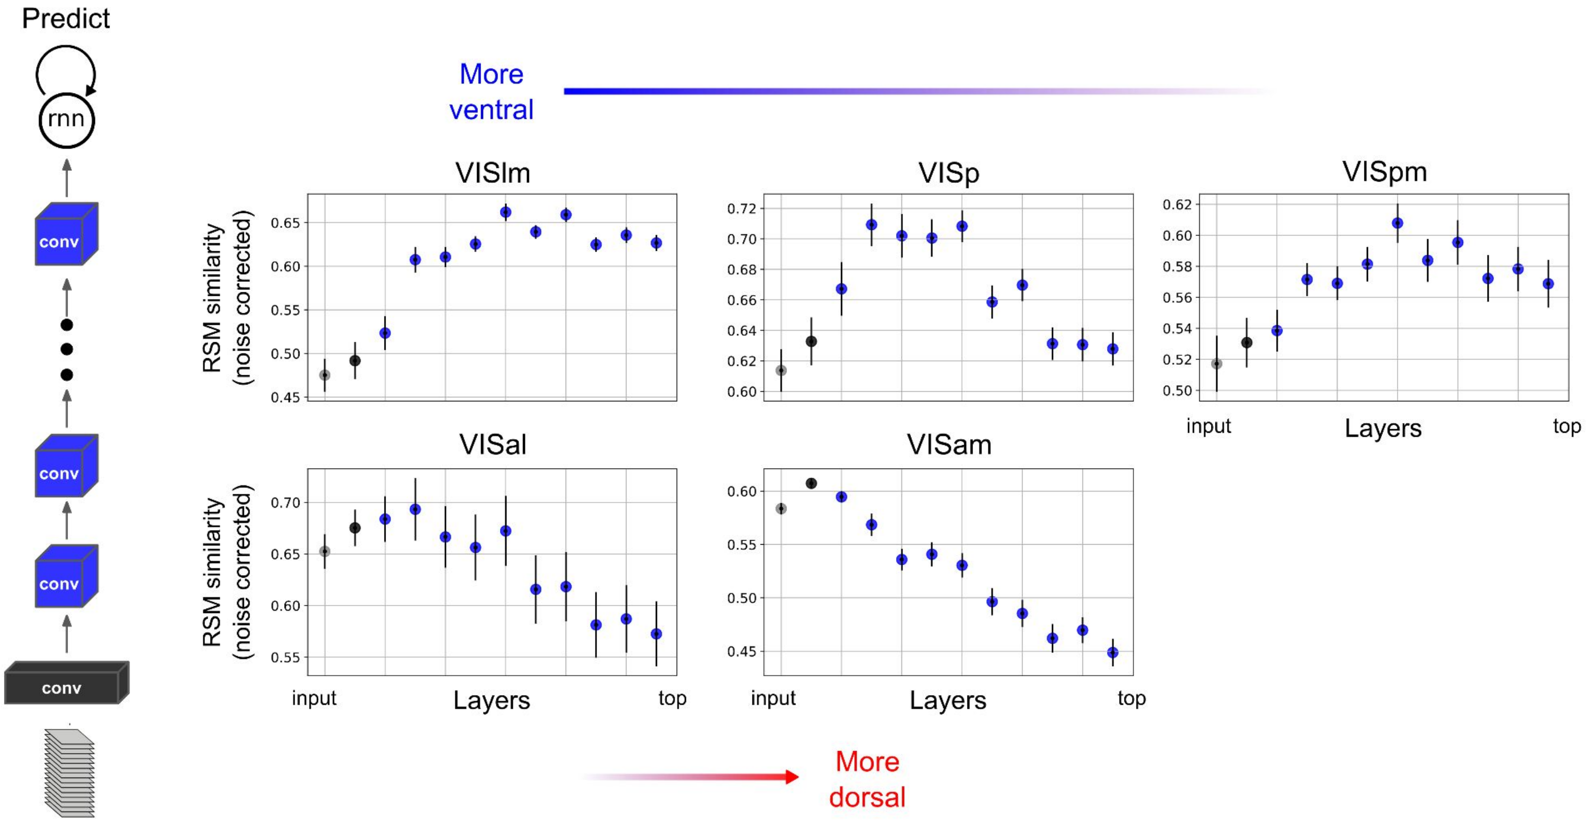

Training a single pathway induces ventral-like representations

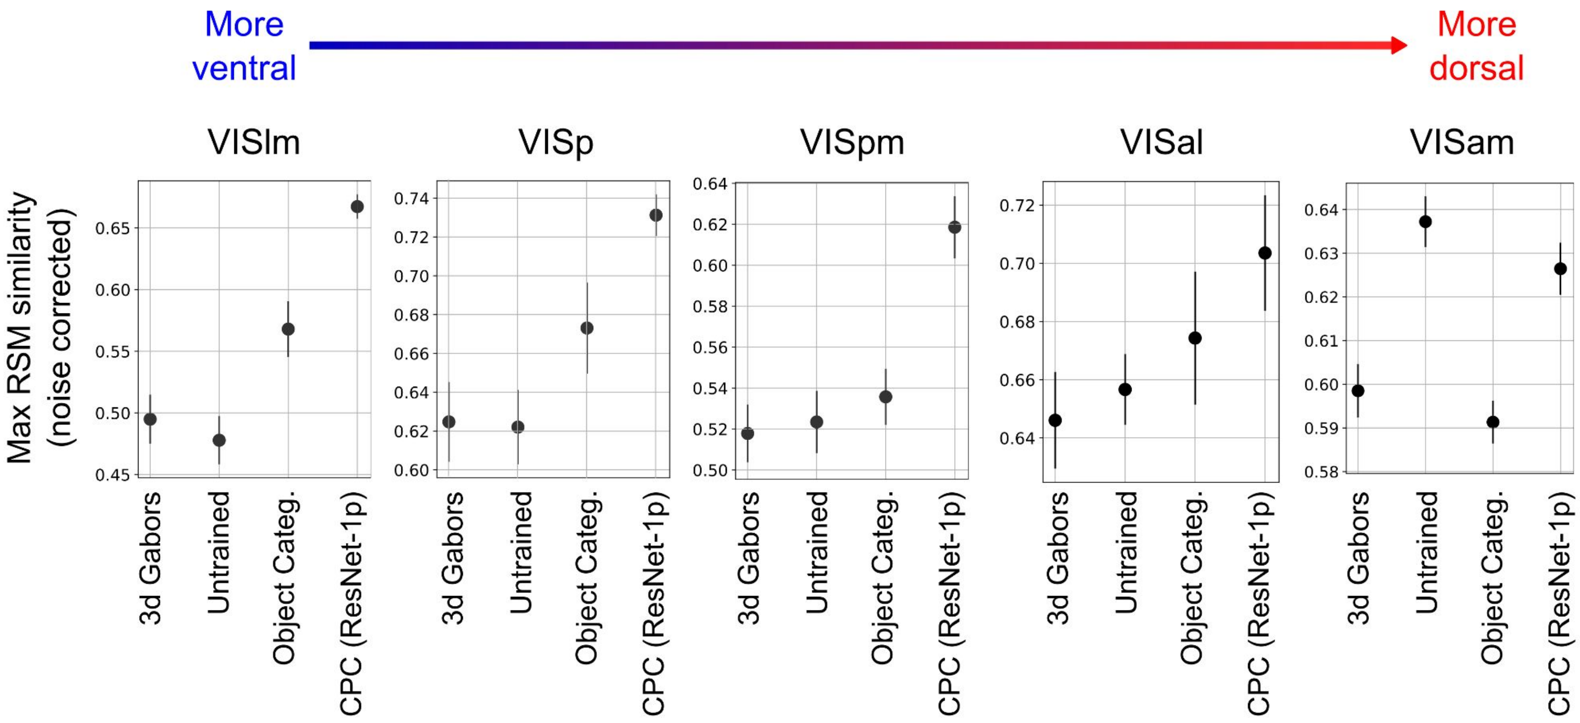

Two pathways leads to both ventral- and dorsal-like representations

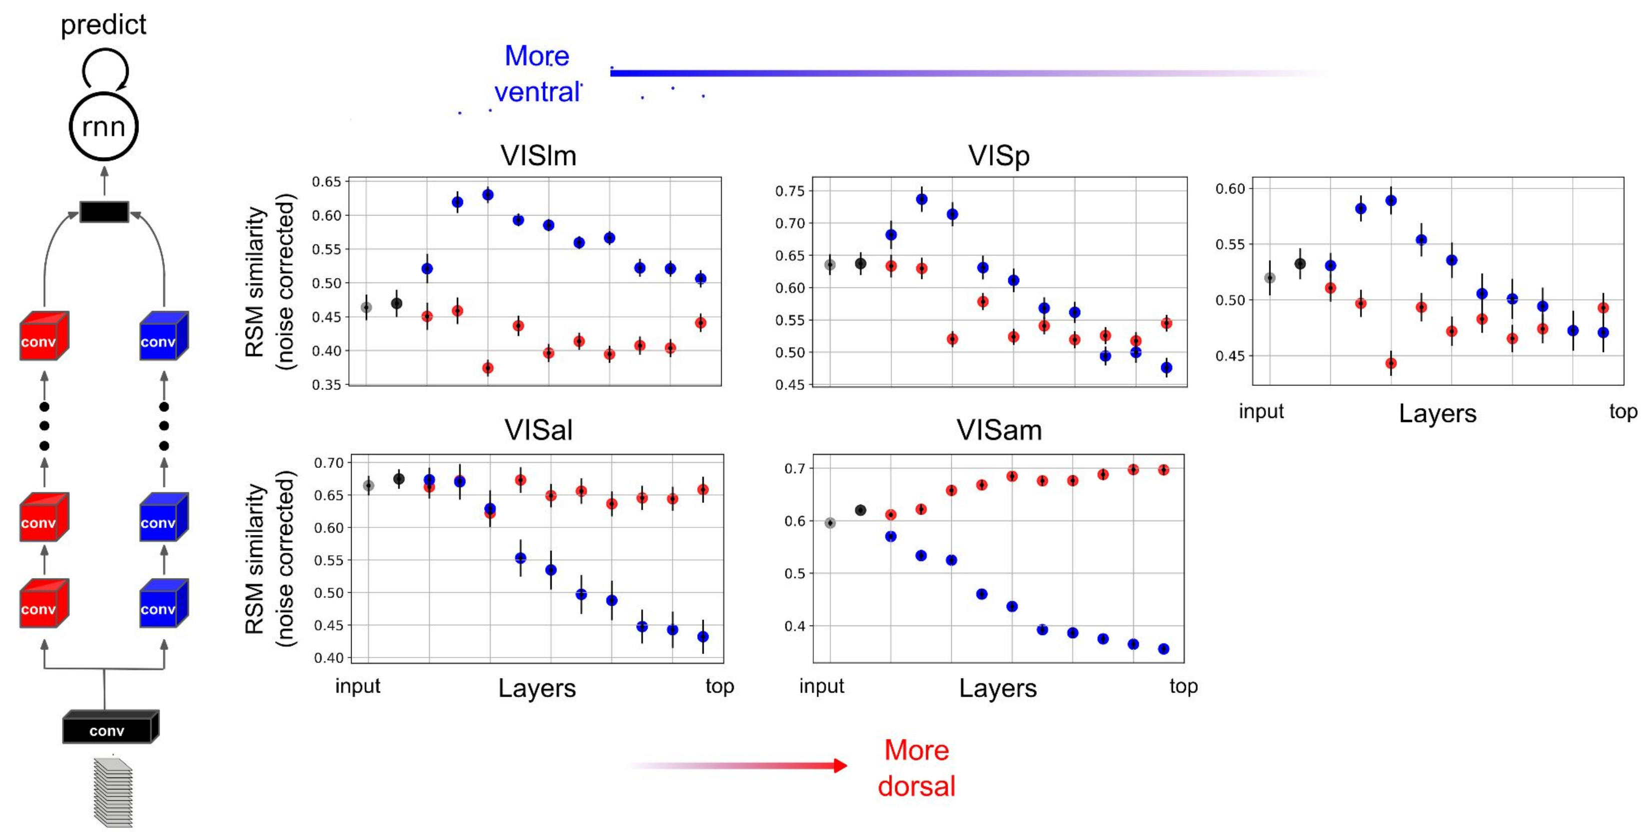

Two pathways leads to both ventral- and dorsal-like representations

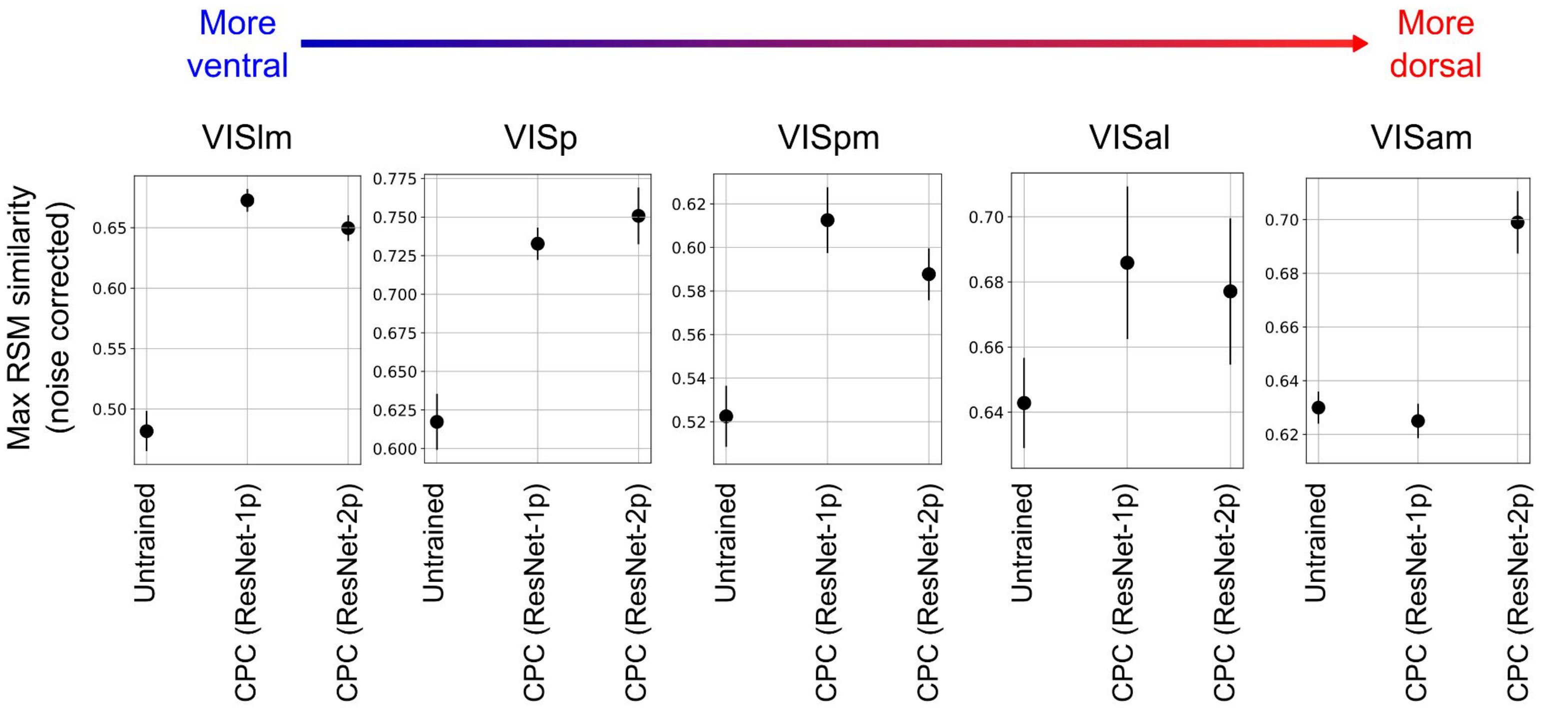

Functional specialization in ventral- and dorsal-like pathways

Object Categorization (CIFAR-10)

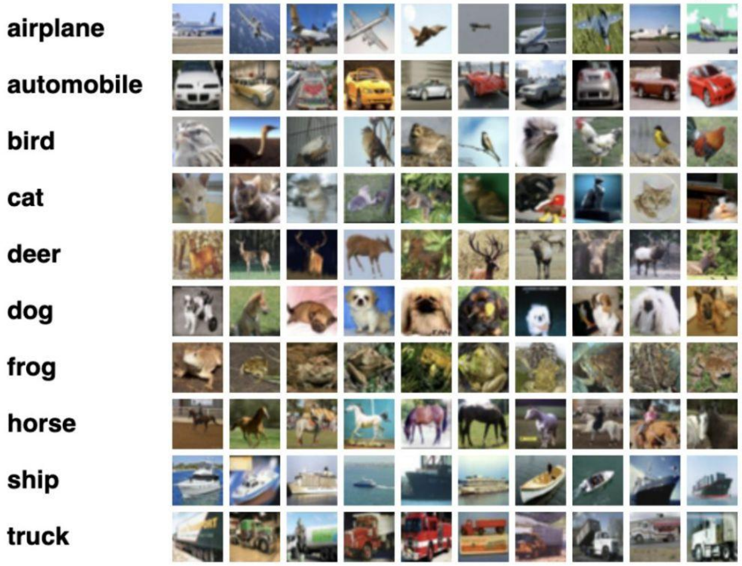

Motion Discrimination (4 directions)

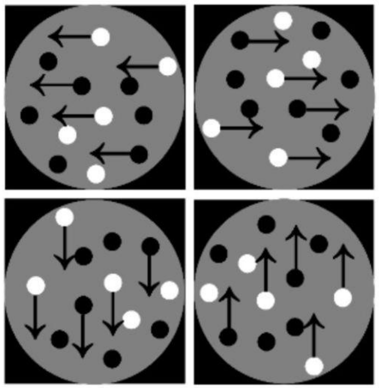

Functional specialization in ventral- and dorsal-like pathways

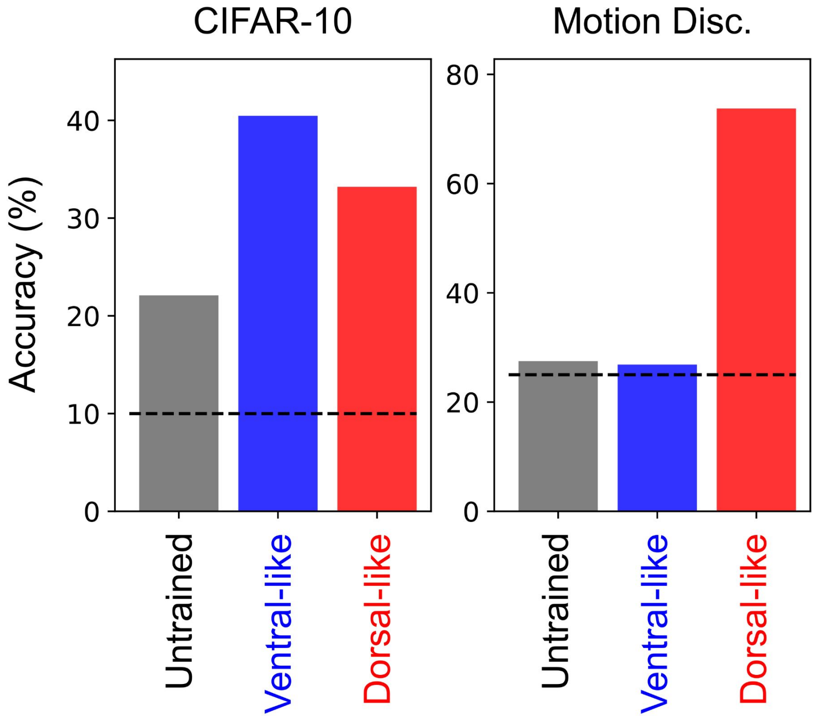

Two pathways allows for good prediction with fewer parameters

	Top-3 accuracy	Num. parameters
ResNet-1p	94.64 (0.68)	435k
ResNet-2p	93.472 (0.83)	285k

Two pathways allows for good prediction with fewer parameters

	Top-3 accuracy	Num. parameters
ResNet-1p	94.64 (0.68)	435k
ResNet-2p	93.472 (0.83)	285k

The single pathway model has some dorsal-like units, but not many

	% dorsal-like units
ResNet-1p	38.96 (1.30)
ResNet-2p (D-path)	60.09 (3.73)
ResNet-2p (V-path)	17.00 (1.05)

Supervised training to categorize actions does not induce specialization

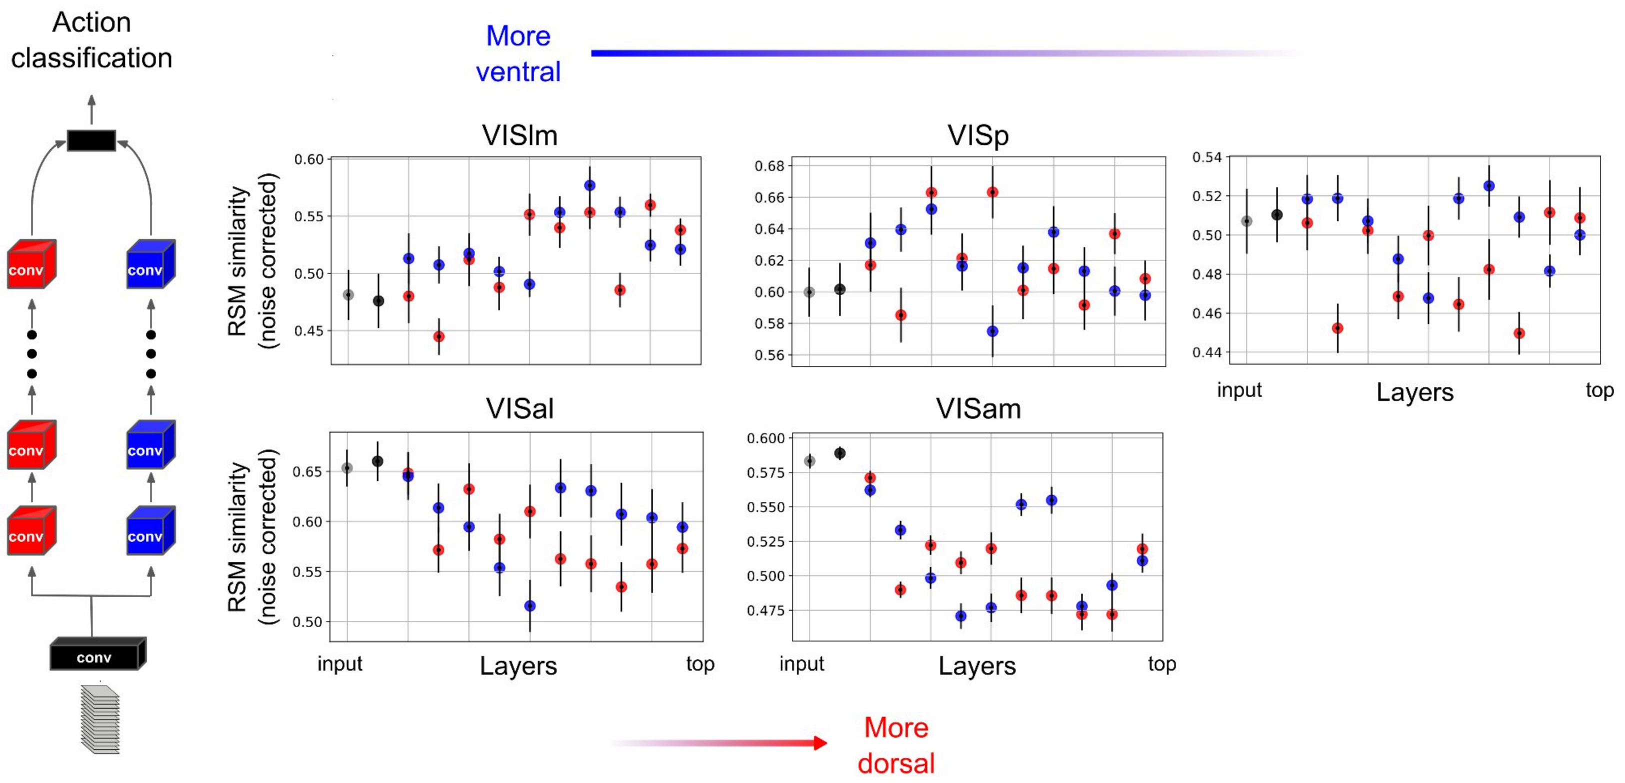

Supervised training to categorize actions does not induce specialization

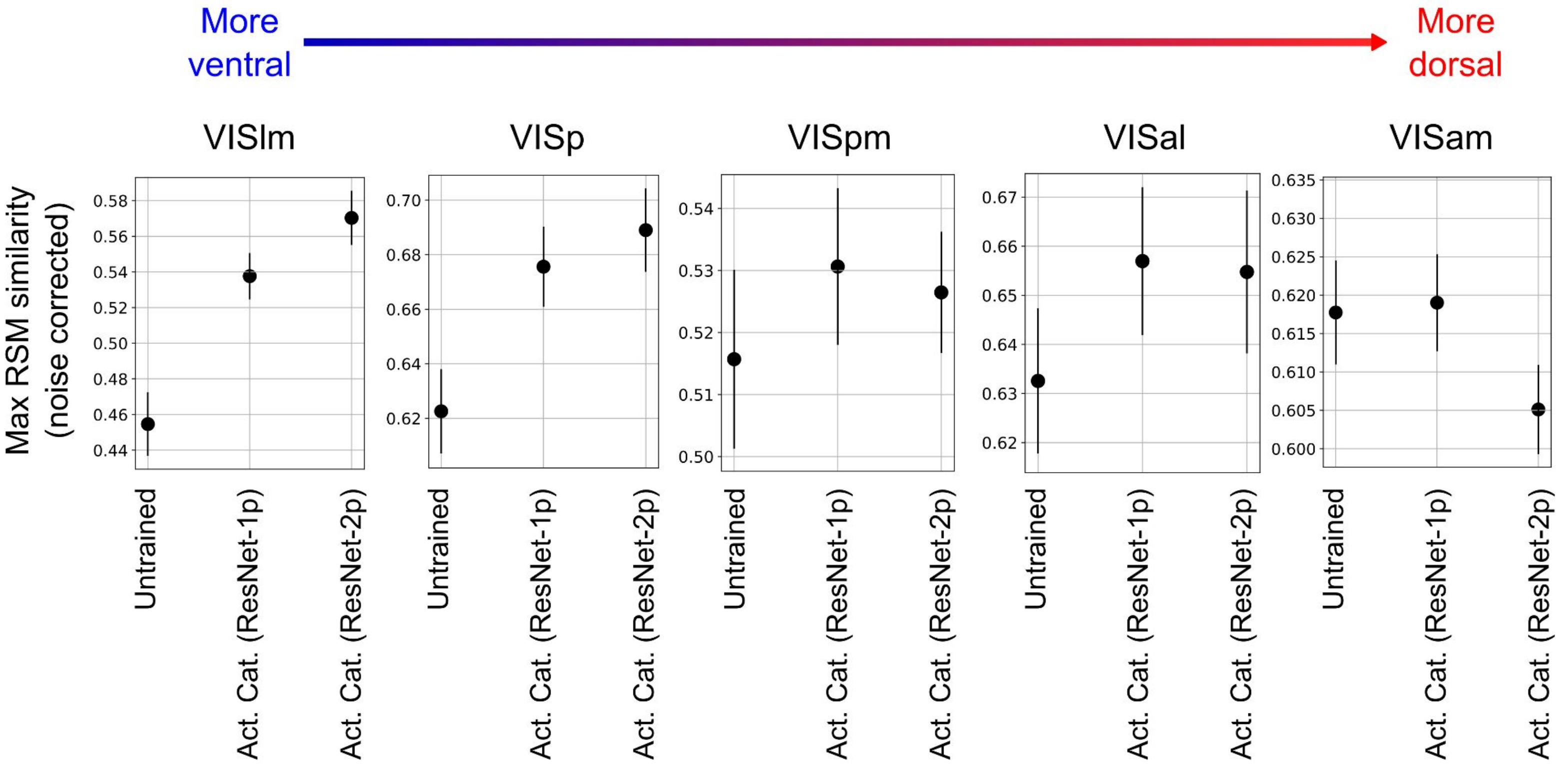

Conclusion

These results demonstrate that optimizing to predict upcoming visual inputs is <u>sufficient</u> to induce "what" versus "where" specialization, two *notes*:

(1) We can remain agnostic as to whether the optimization was evolution or learning in early life

(2) We cannot rule out the possibility that there are other losses (including supervised ones) that can also work

These results demonstrate that optimizing to predict upcoming visual inputs is <u>sufficient</u> to induce “what” versus “where” specialization, two *questions*:

(1) Can we find inductive biases to ensure that specific pathways take on specific functions?

(2) Could we do a better job matching neural representations by adding some additional losses?

A final thought:
parallel pathways with self-supervised learning may be a good strategy for control in AI systems

Control needs a "where" pathway

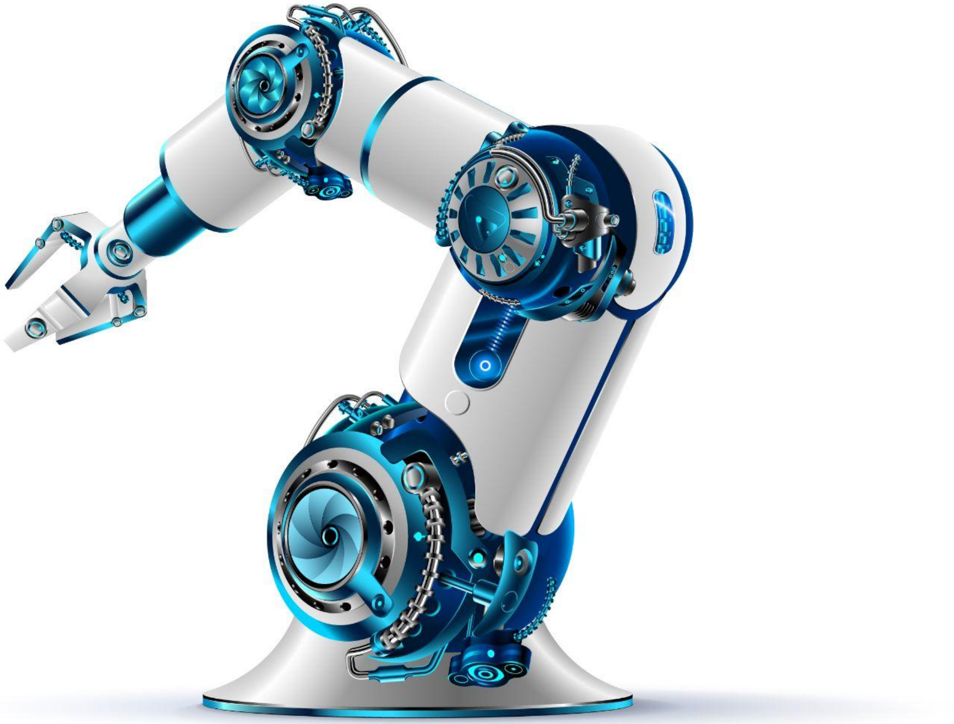

Thanks to:

Shahab Bakhtiari
Timothy Lillicrap
Patrick Mineault
Christopher Pack
The Allen Brain Observatory

neuro
Institut-Hôpital neurologique de Montréal
Montreal Neurological Institute-Hospital

LiNC LAB

McGill

Funding:

HEALTHY BRAINS FOR HEALTHY LIVES

NSERC CRSNG

Mila

CIFAR

HIBALL
HELMHOLTZ International BigBrain Analytics & Learning Laboratory

Two pathways leads to both ventral- and dorsal-like representations

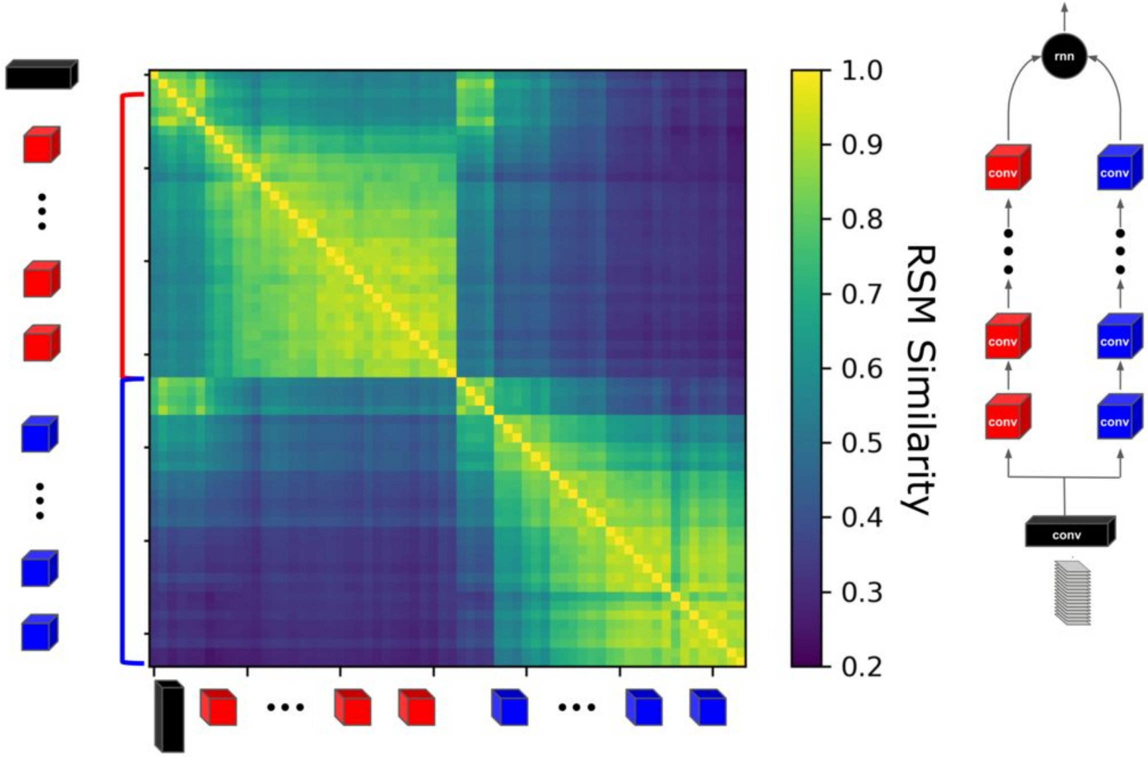

Two pathways leads to both ventral- and dorsal-like representations

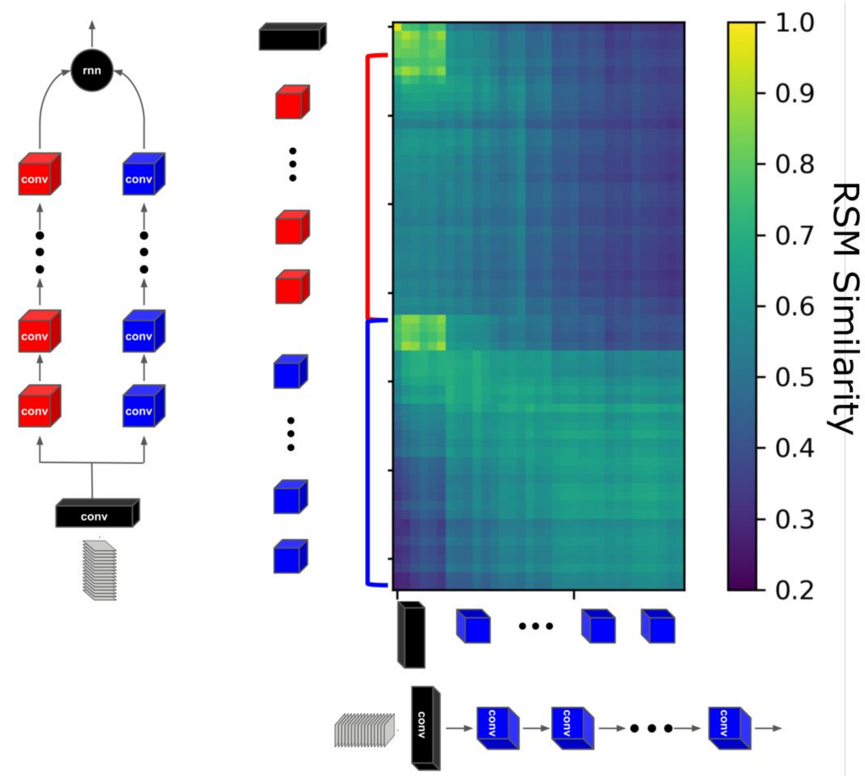

Functional specialization in ventral- and dorsal-like pathways

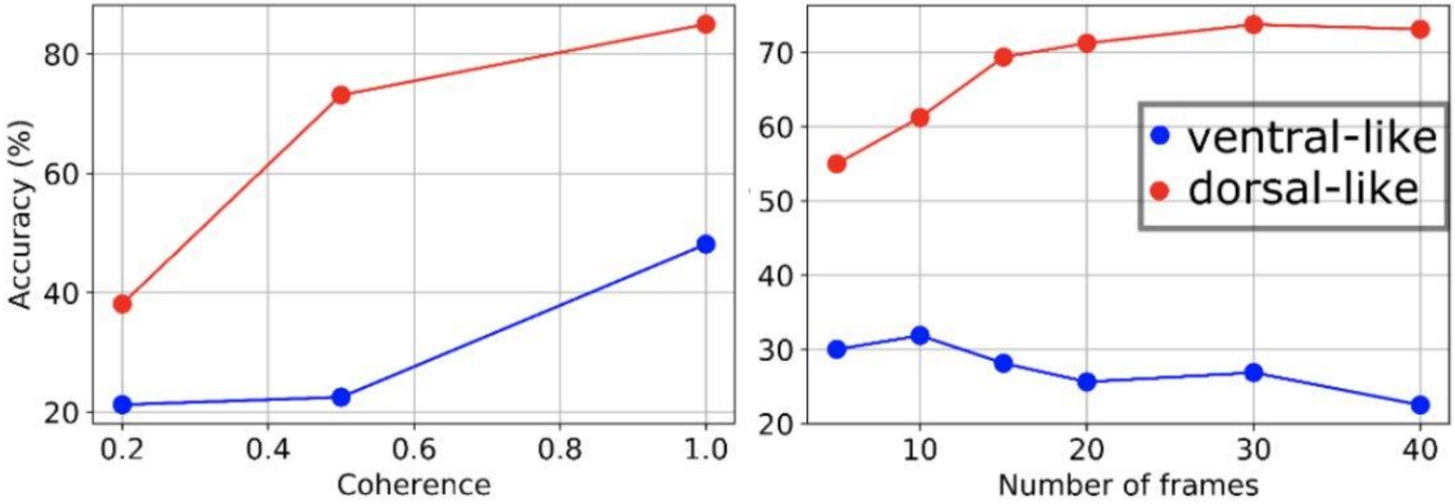

Supervised training does not induce dorsal-like representations

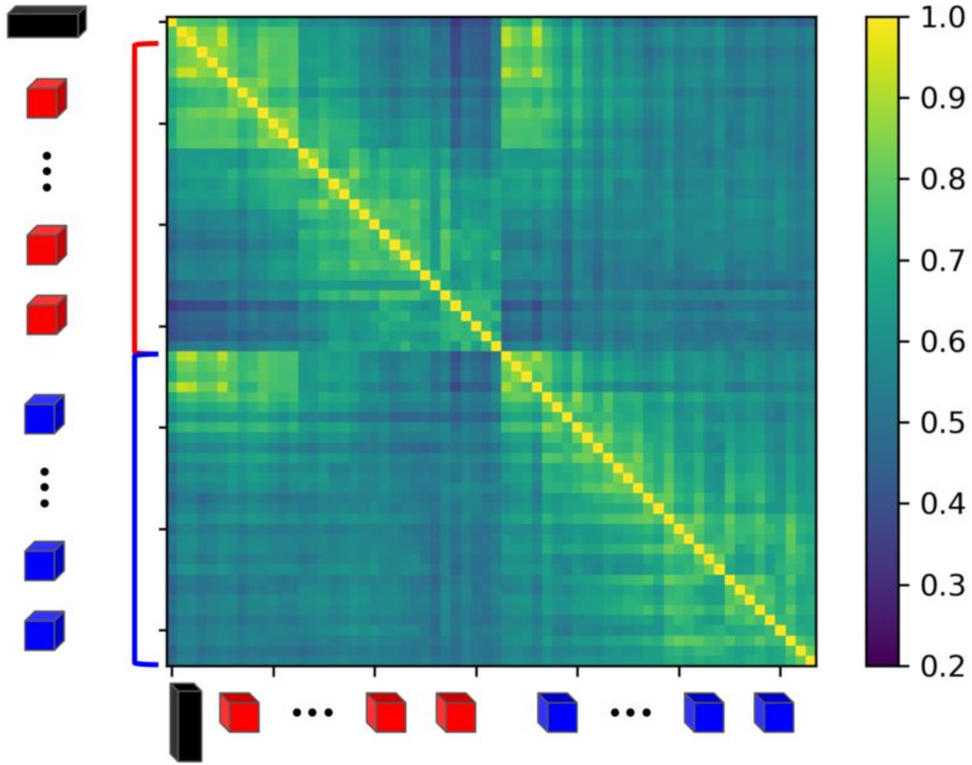

Ventral- and dorsal-like representations compete during learning, but two pathways mitigate the competition

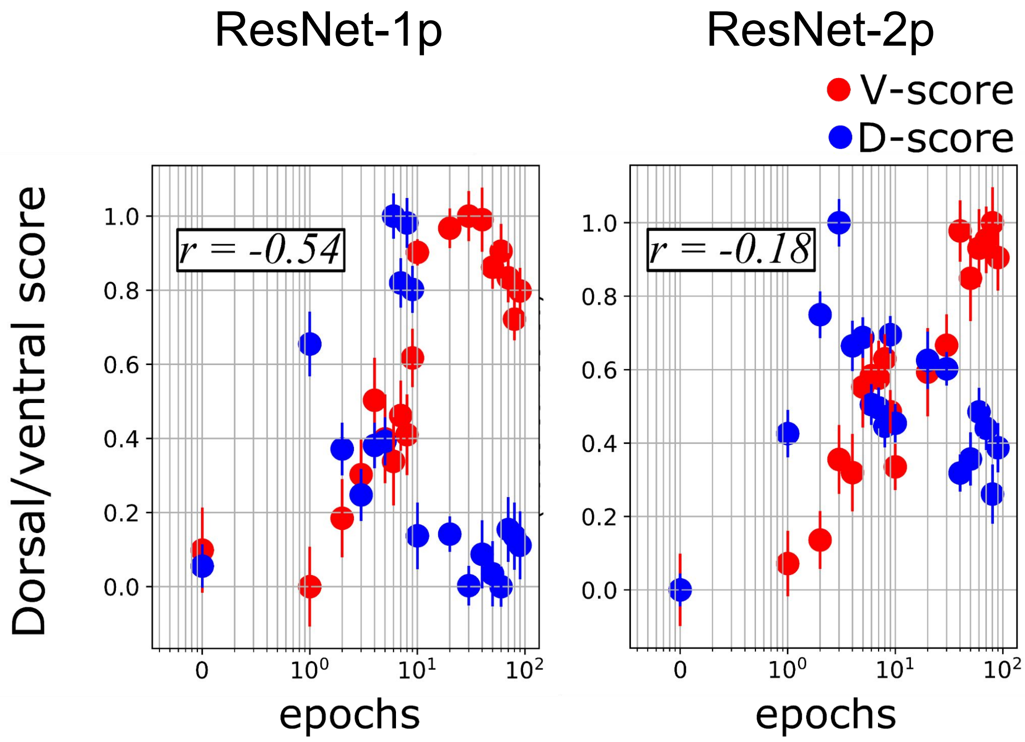

Our results match those of another recent study, showing that self-supervised learning induces better matches to mouse visual cortex than supervised learning

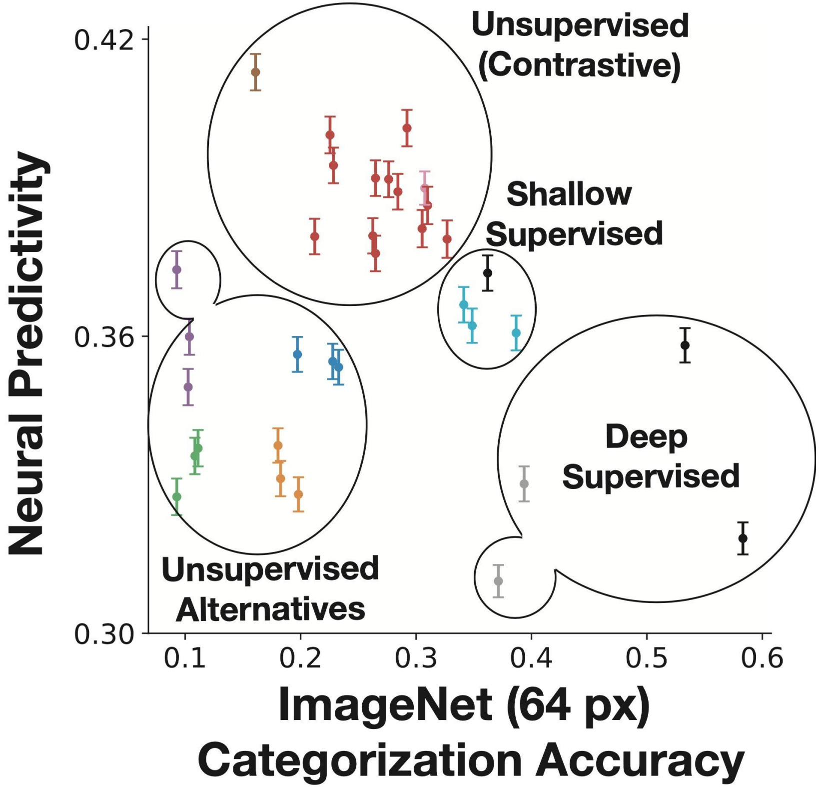

Nayebi et al. (2021), biorXiv: https://doi.org/10.1101/2021.06.16.448730

There are important differences between rodent and primate visual cortex, different loss functions may be required

Mouse (30 g) Tree shrew (200 g) Mouse lemur (60 g) Macaque (10 kg)

Luongo et al. (2021), biorXiv:
https://doi.org/10.1101/2021.07.04.451059

We have another paper coming out demonstrating good fit to primate dorsal stream using self-motion estimation

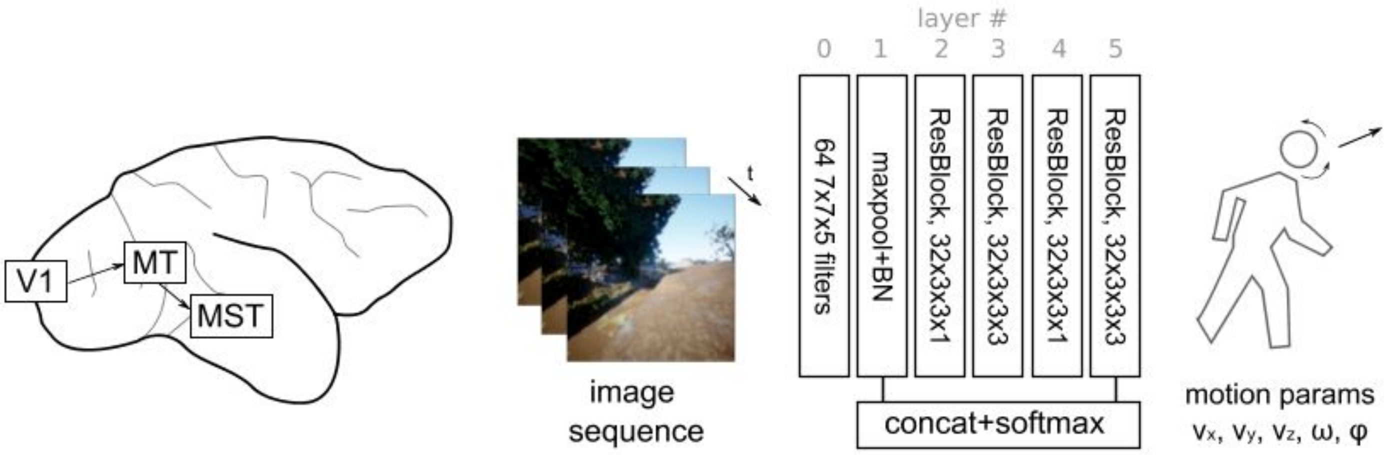

Minault et al. (2021), *in review*

We have another paper coming out demonstrating good fit to primate dorsal stream using self-motion estimation