CS375 / Psych 249:
Large-Scale Neural Network Models for Neuroscience

Lecture 4: Deep CNNs and the Ventral Visual Stream — Part 2

2025.01.15

Daniel Yamins
Stanford Neurosciences Institute
Stanford Artificial Intelligence Laboratory
Departments of Psychology and Computer Science
Stanford University

Layer-area correspondence

Deep convolutional models improve predictions of macaque V1 responses to natural images

Santiago A Cadena, George H Denfield, Edgar Y Walker, Leon A Gatys, Andreas S Tolias, Matthias Bethge, Alexander S Ecker

doi: https://doi.org/ 10.1101/201764

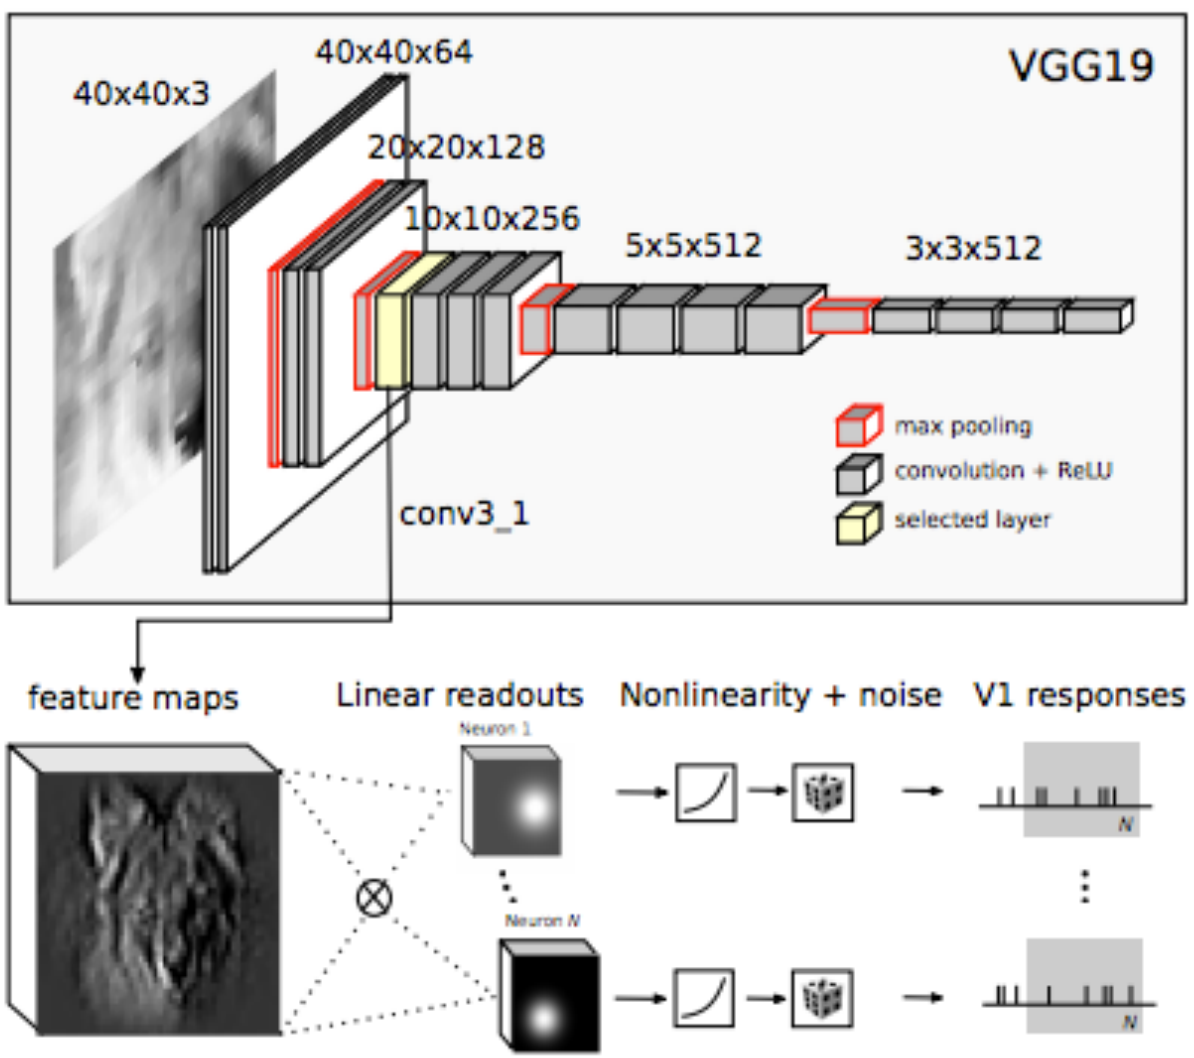

Layer-area correspondence

Deep convolutional models improve predictions of macaque V1 responses to natural images

Santiago A Cadena, George H Denfield, Edgar Y Walker, Leon A Gatys, Andreas S Tolias,
Matthias Bethge, Alexander S Ecker

doi: https://doi.org/ 10.1101/201764

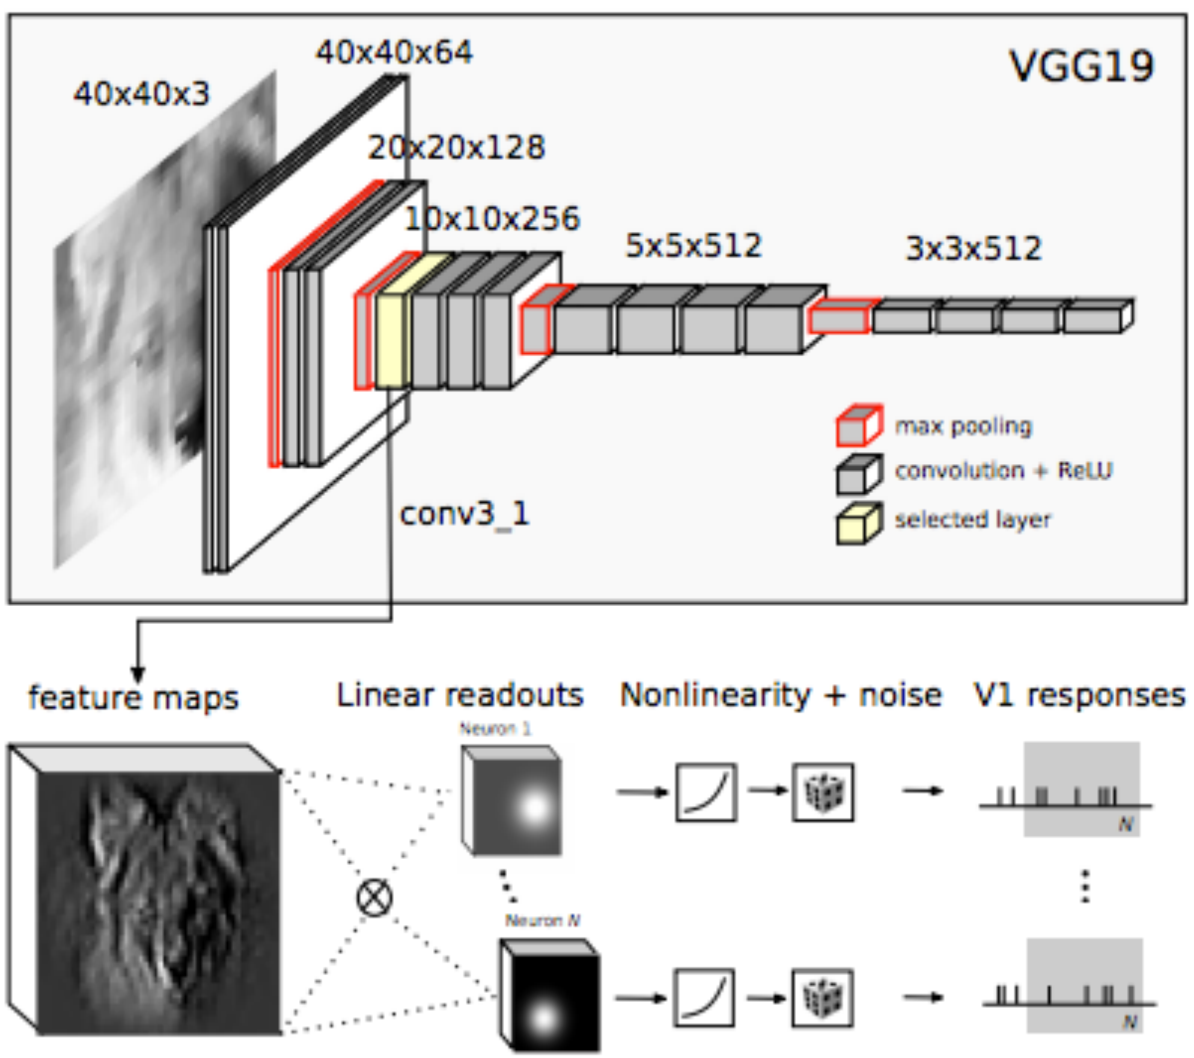

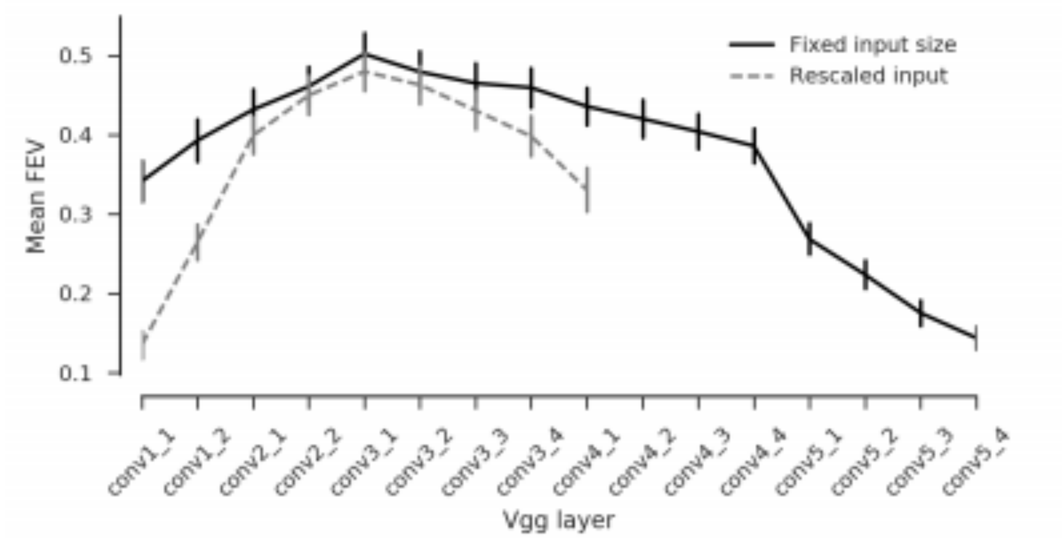

Layer-area correspondence

Deep convolutional models improve predictions of macaque V1 responses to natural images

Santiago A Cadena, George H Denfield, Edgar Y Walker, Leon A Gatys, Andreas S Tolias,
Matthias Bethge, Alexander S Ecker

doi: https://doi.org/ 10.1101/201764

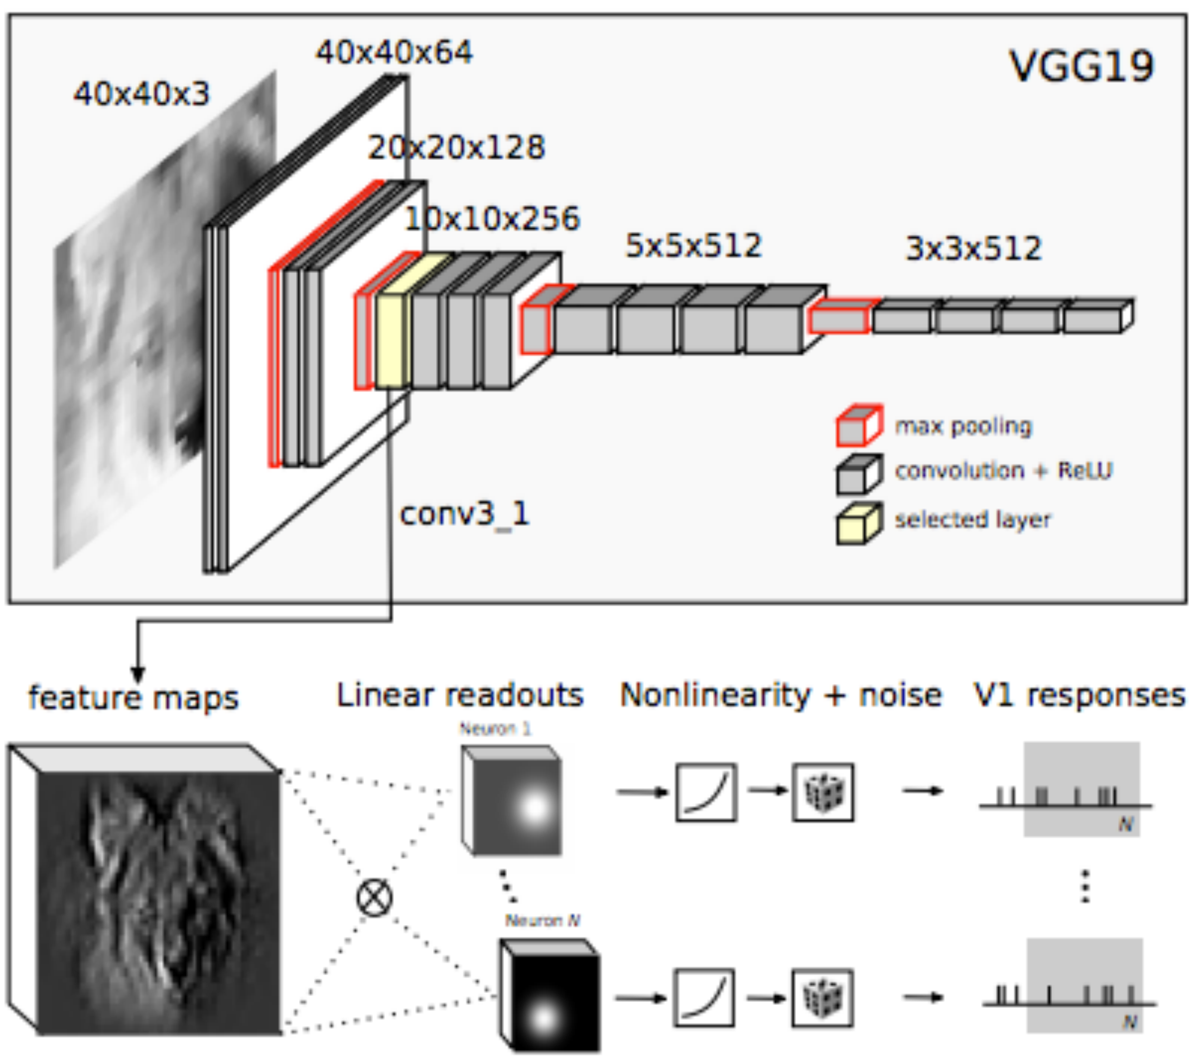

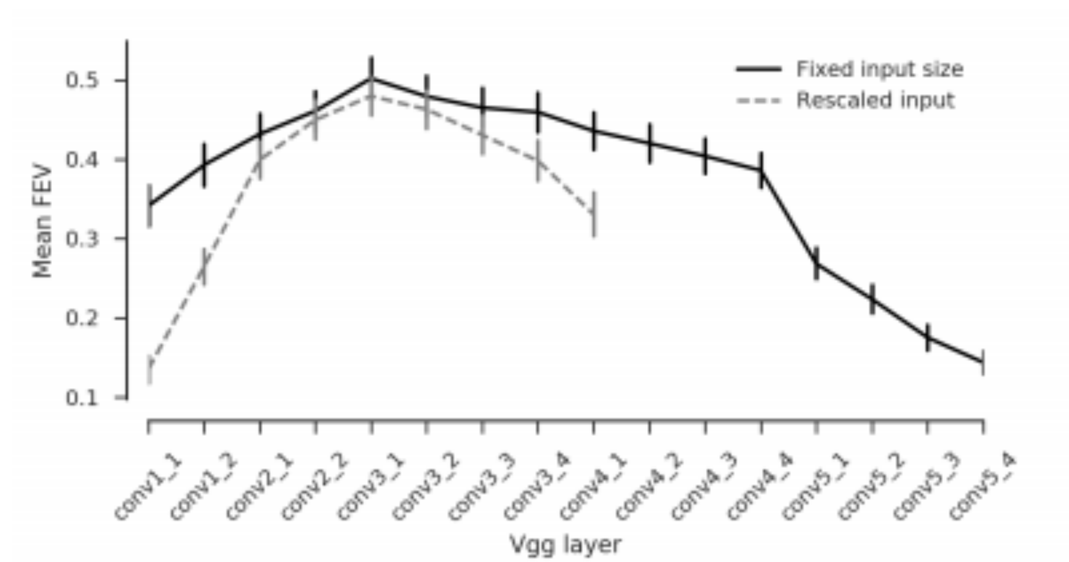

finer resolution

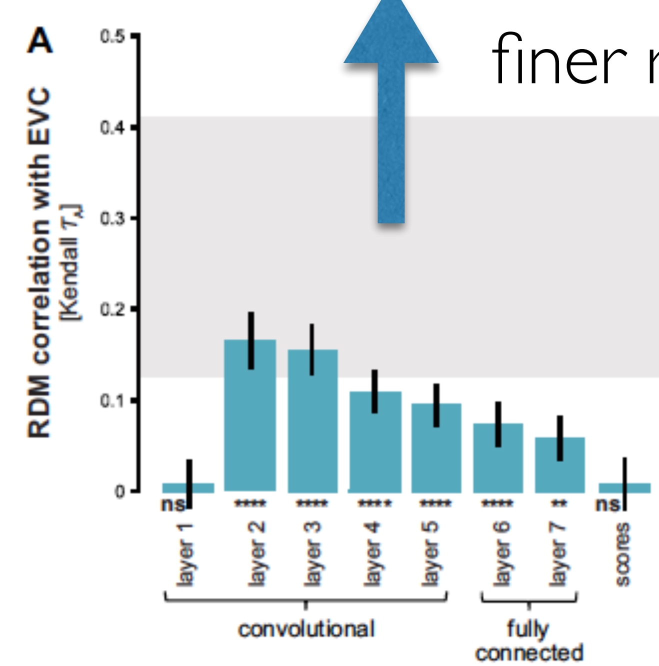

Layer-area correspondence

Deep convolutional models improve predictions of macaque V1 responses to natural images

Santiago A Cadena, George H Denfield, Edgar Y Walker, Leon A Gatys, Andreas S Tolias, Matthias Bethge, Alexander S Ecker

doi: https://doi.org/ 10.1101/201764

50% explained variance vs

- 17% for Linear-Nonlinear-Poisson (with gabor filters)
- 39% for Berkeley Wavelet Transform

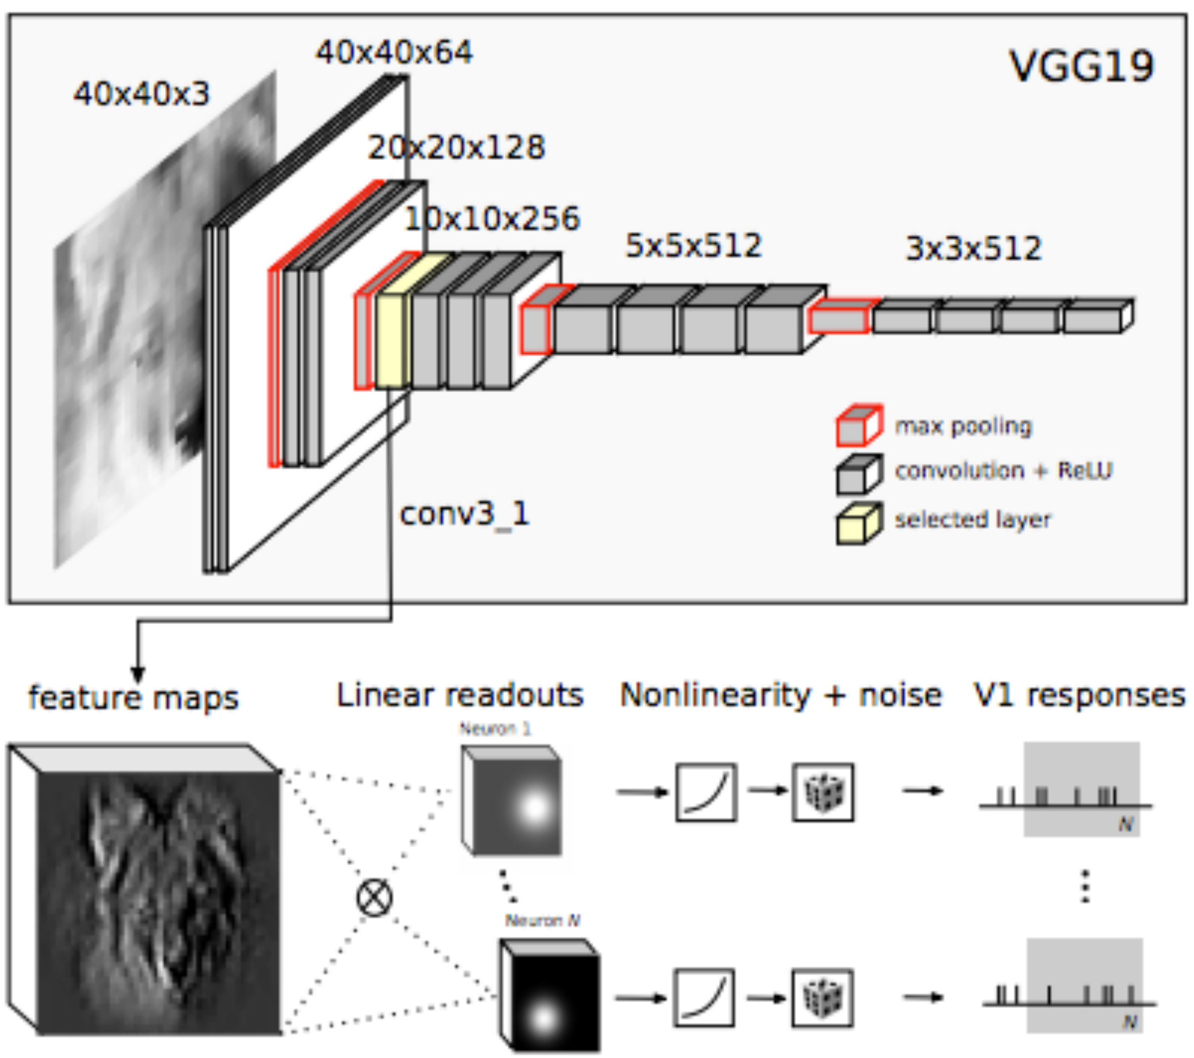

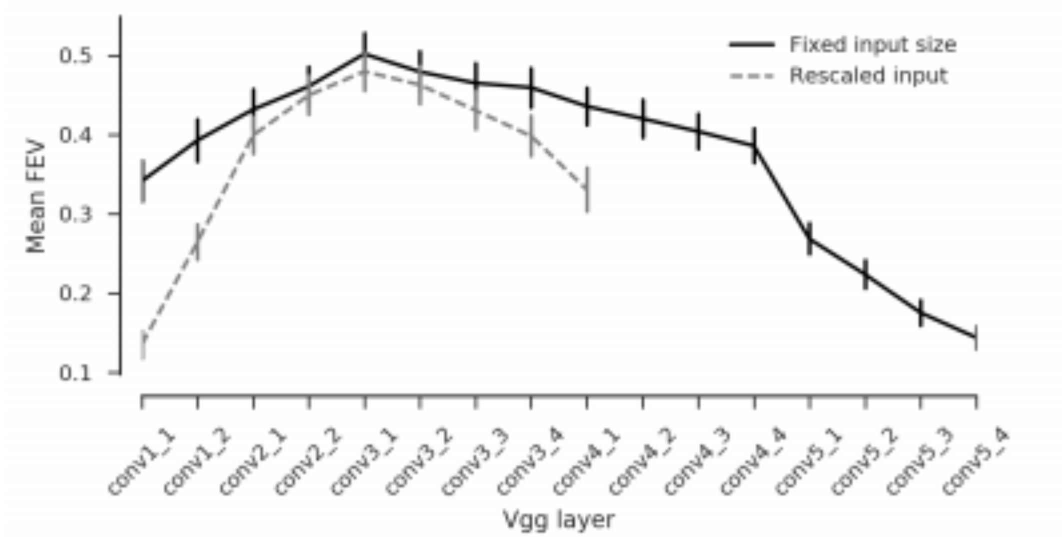

finer resolution

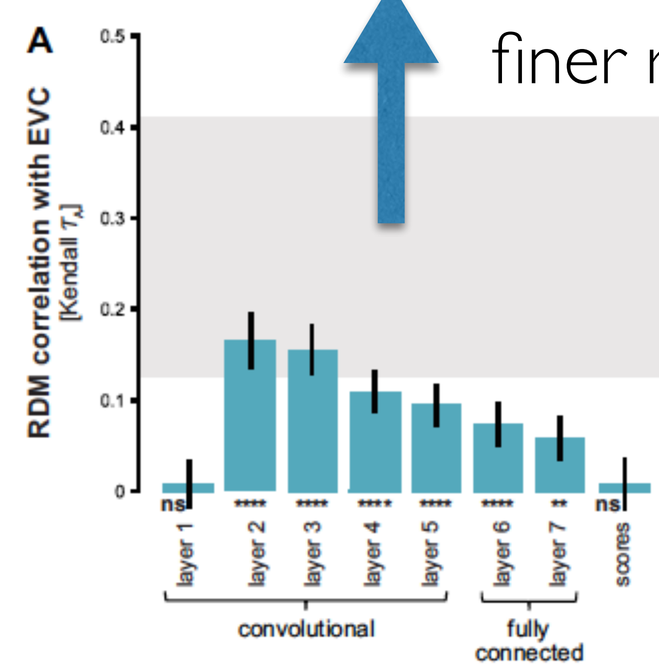

Layer-area correspondence

Deep convolutional models improve predictions of macaque V1 responses to natural images

Santiago A Cadena, George H Denfield, Edgar Y Walker, Leon A Gatys, Andreas S Tolias, Matthias Bethge, Alexander S Ecker

doi: https://doi.org/ 10.1101/201764

50% explained variance vs

- 17% for Linear-Nonlinear-Poisson (with gabor filters)
- 39% for Berkeley Wavelet Transform

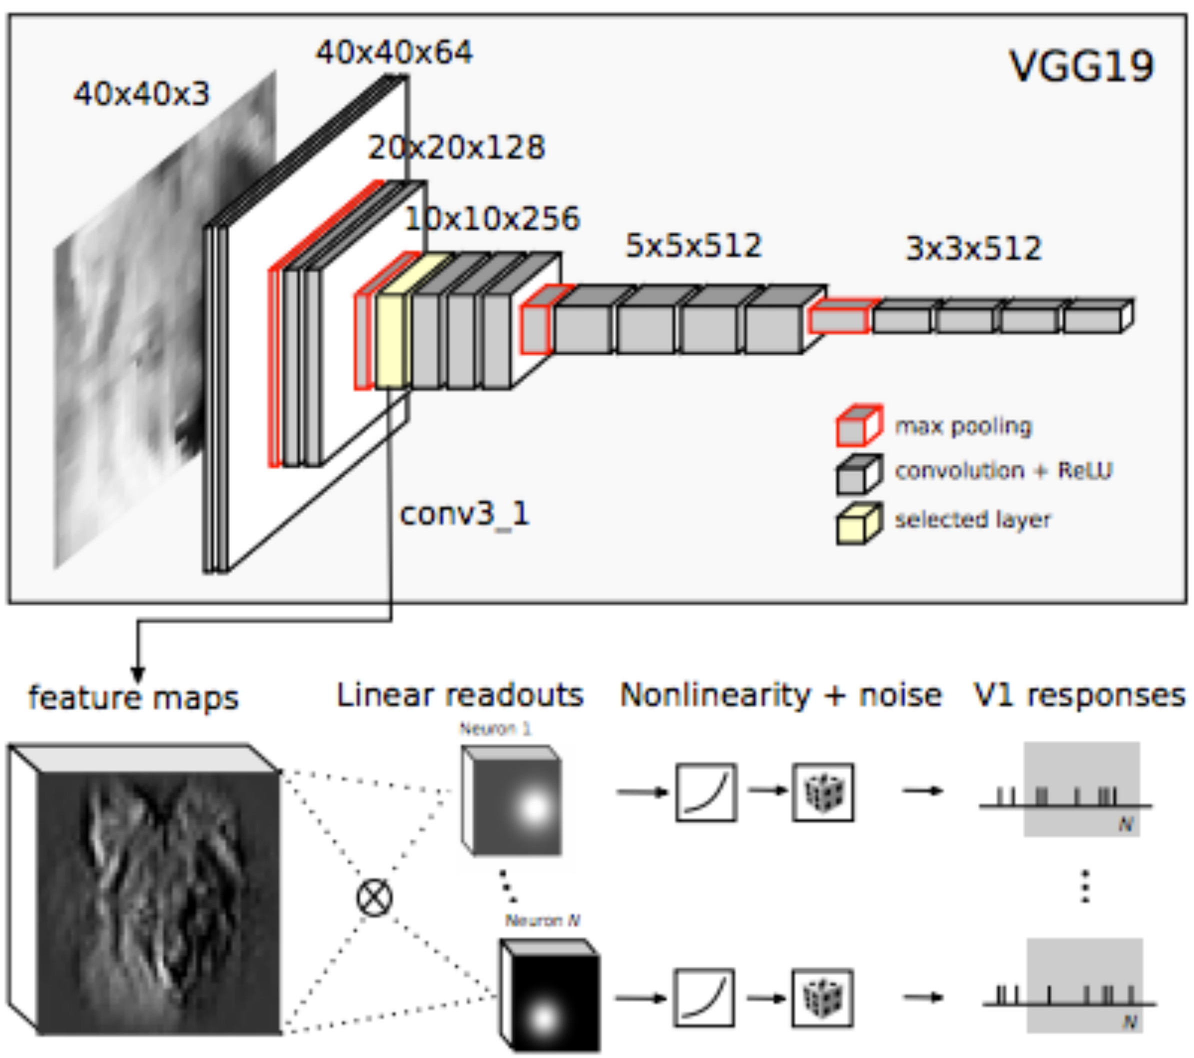

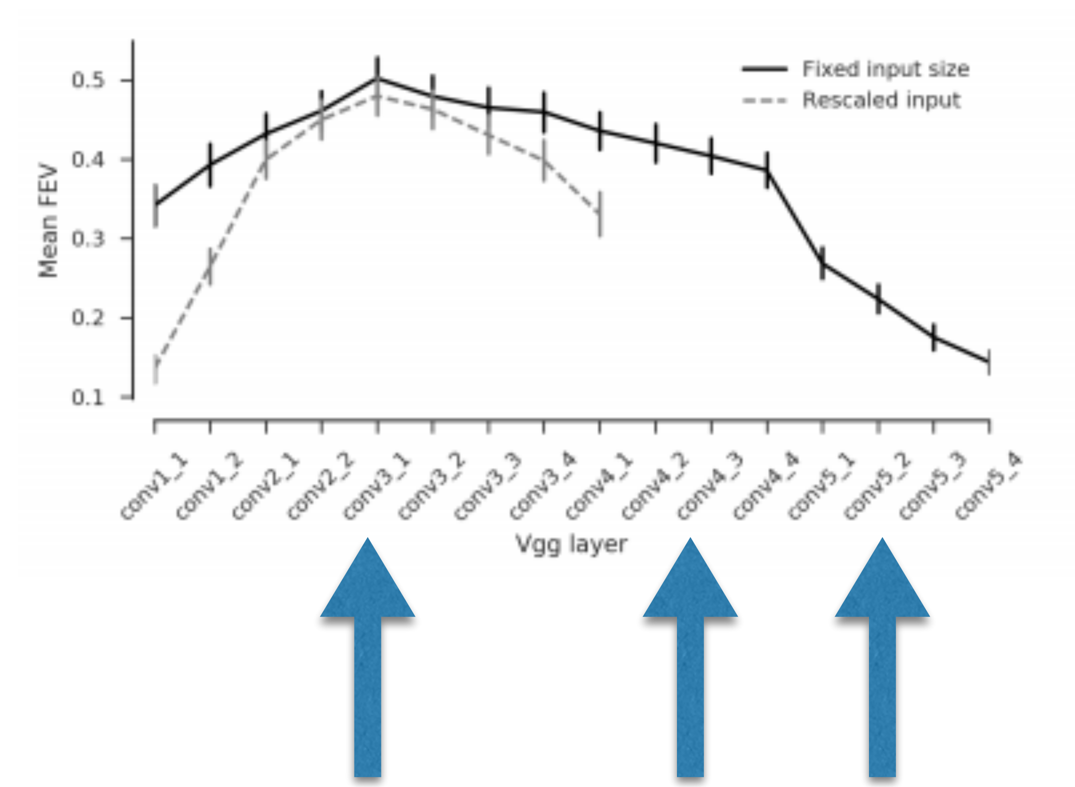

Peak V1

Peak V4 (unpublished)

Peak IT (unpublished)

Layer-area correspondence

Deep convolutional models improve predictions of macaque V1 responses to natural images

Santiago A Cadena, George H Denfield, Edgar Y Walker, Leon A Gatys, Andreas S Tolias, Matthias Bethge, Alexander S Ecker

doi: https://doi.org/ 10.1101/201764

50% explained variance vs

- 17% for Linear-Nonlinear-Poisson (with gabor filters)
- 39% for Berkeley Wavelet Transform

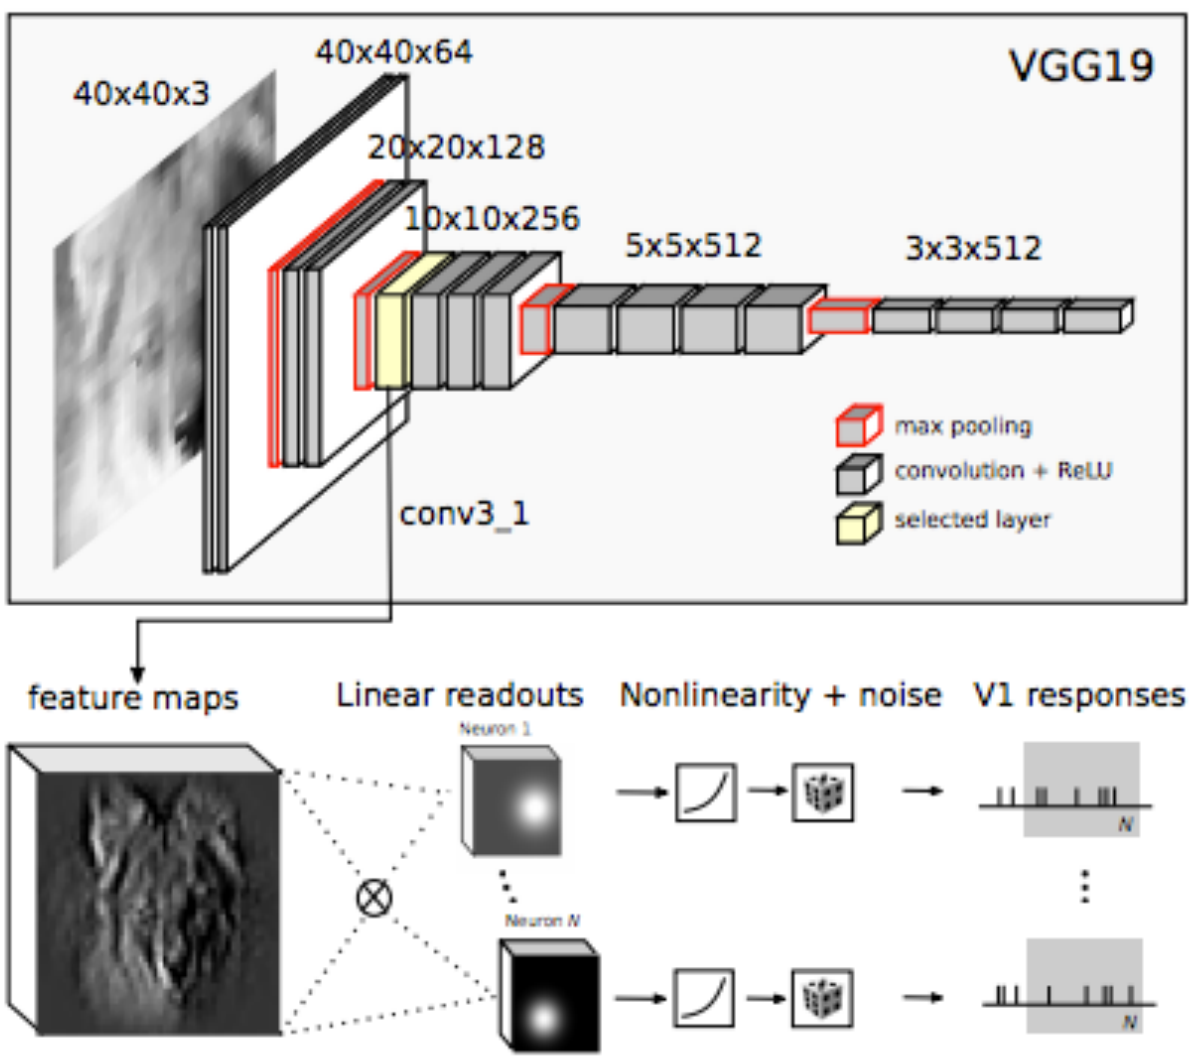

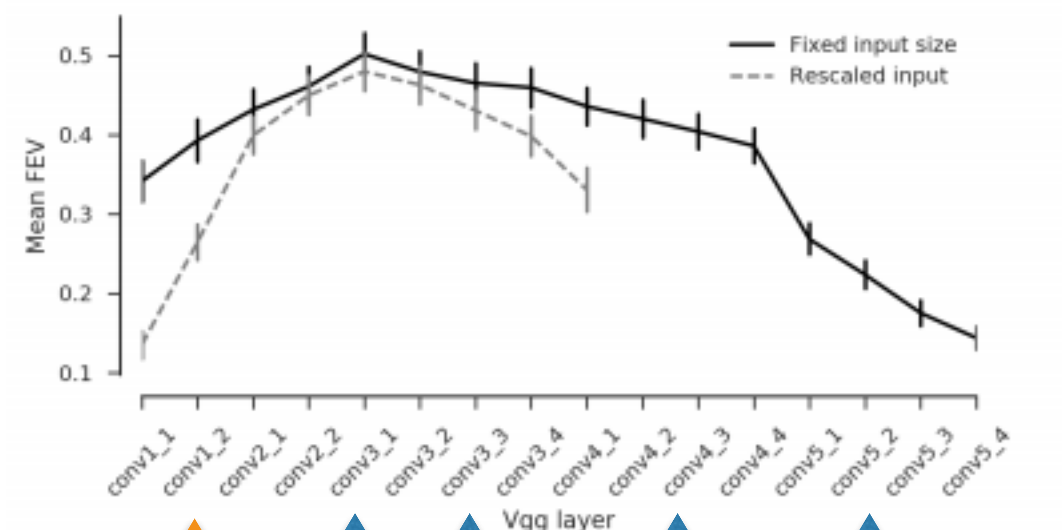

Peak V1

Peak V4

Peak IT

Peak V2

subcortical??

Layer-area correspondence

Deep Learning Models of the Retinal Response to Natural Scenes

Lane T. McIntosh*[1], Niru Maheswaranathan*[1], Aran Nayebi[1],
Surya Ganguli[2,3], Stephen A. Baccus[3]
[1]Neurosciences PhD Program, [2]Department of Applied Physics, [3]Neurobiology Department
Stanford University
{lmcintosh, nirum, anayebi, sganguli, baccus}@stanford.edu

Layer-area correspondence

Deep Learning Models of the Retinal Response to Natural Scenes

Lane T. McIntosh[*1], **Niru Maheswaranathan**[*1], **Aran Nayebi**[1],
Surya Ganguli[2,3], **Stephen A. Baccus**[3]
[1]Neurosciences PhD Program, [2]Department of Applied Physics, [3]Neurobiology Department
Stanford University
{lmcintosh, nirum, anayebi, sganguli, baccus}@stanford.edu

Three-layer CNN best fits retinal ganglion cell response patterns to natural images.

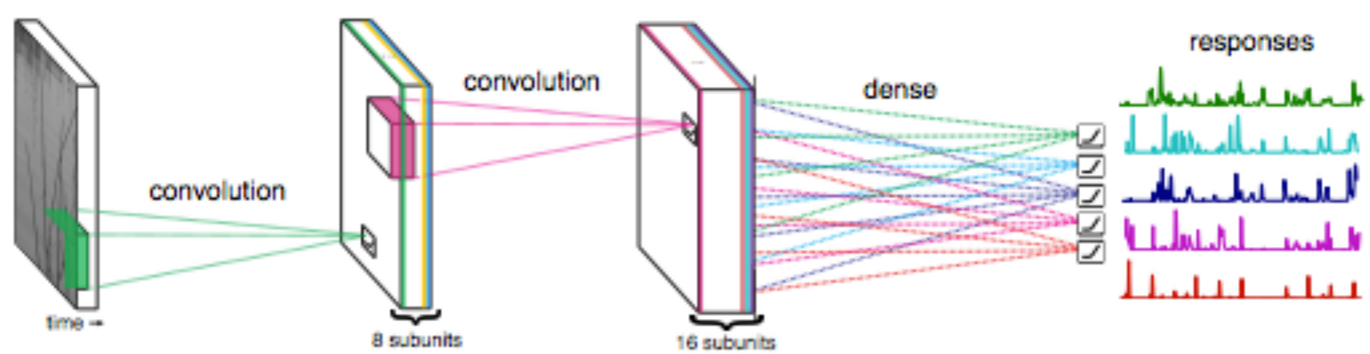

Layer-area correspondence

Better models of the ventral visual stream:

- V4 at 6th convolutional layer
- pIT at 7th convolutional layer
- cIT/aIT at layers 8-10, depending on neurons position on A/P axis

Dan Bear

Jonas Kubilius

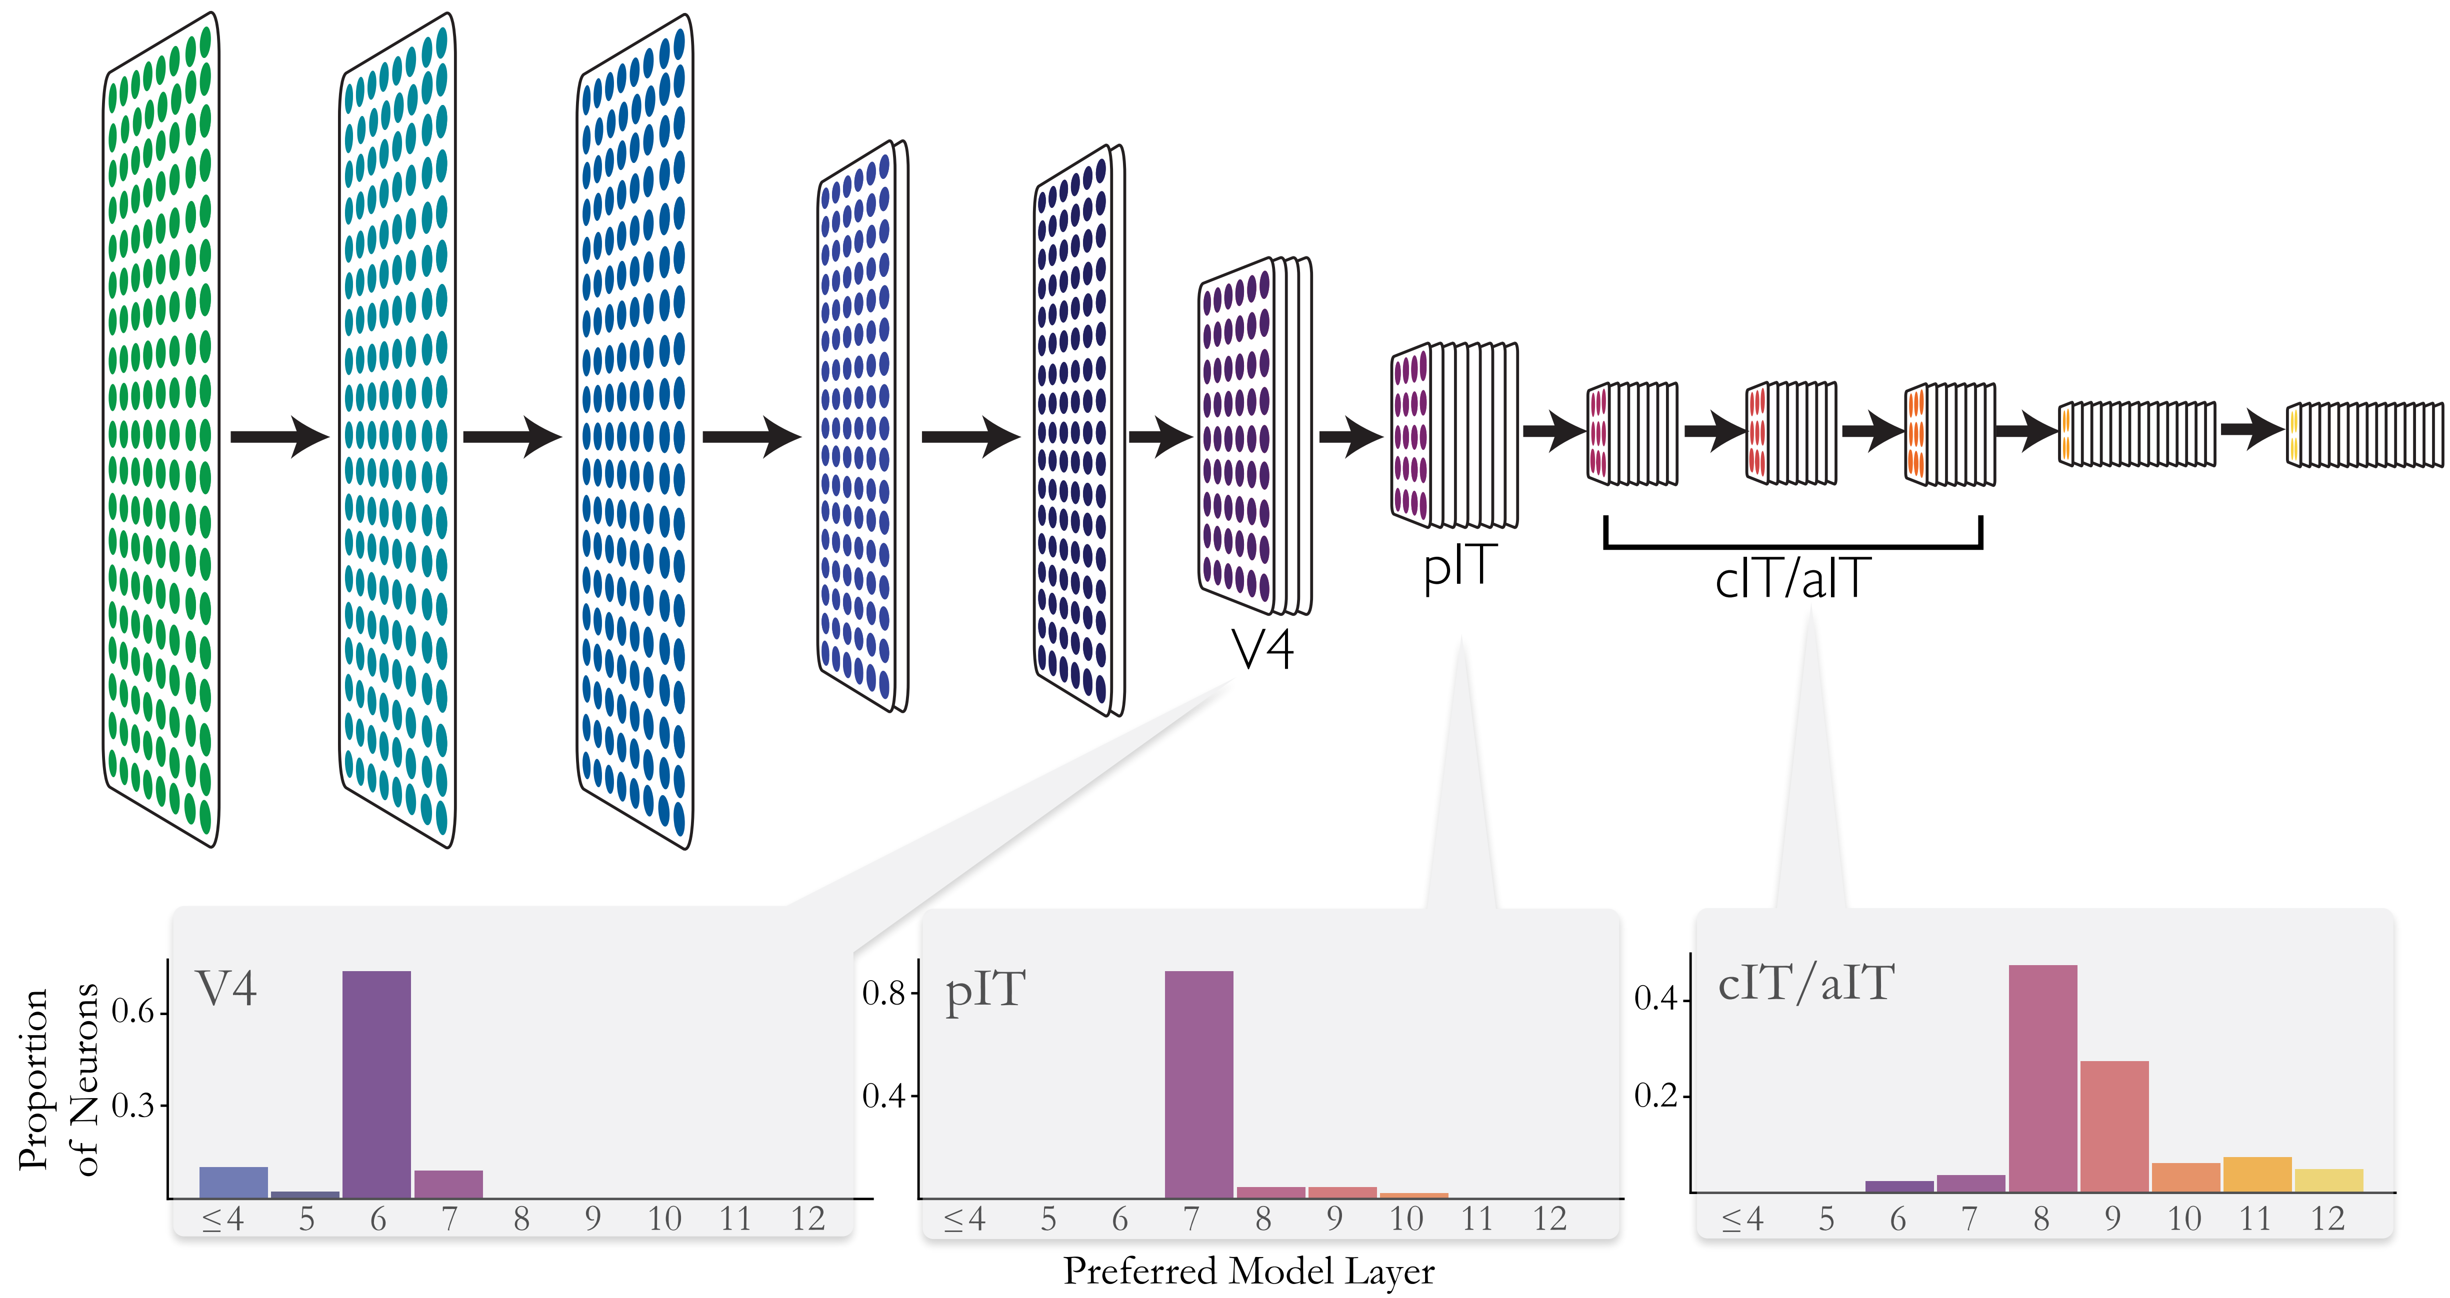

Post-AlexNet Developments

(1) Residual Connections and ResNets

(2) Vision Transformers

Post-AlexNet Developments

(1) Residual Connections and ResNets

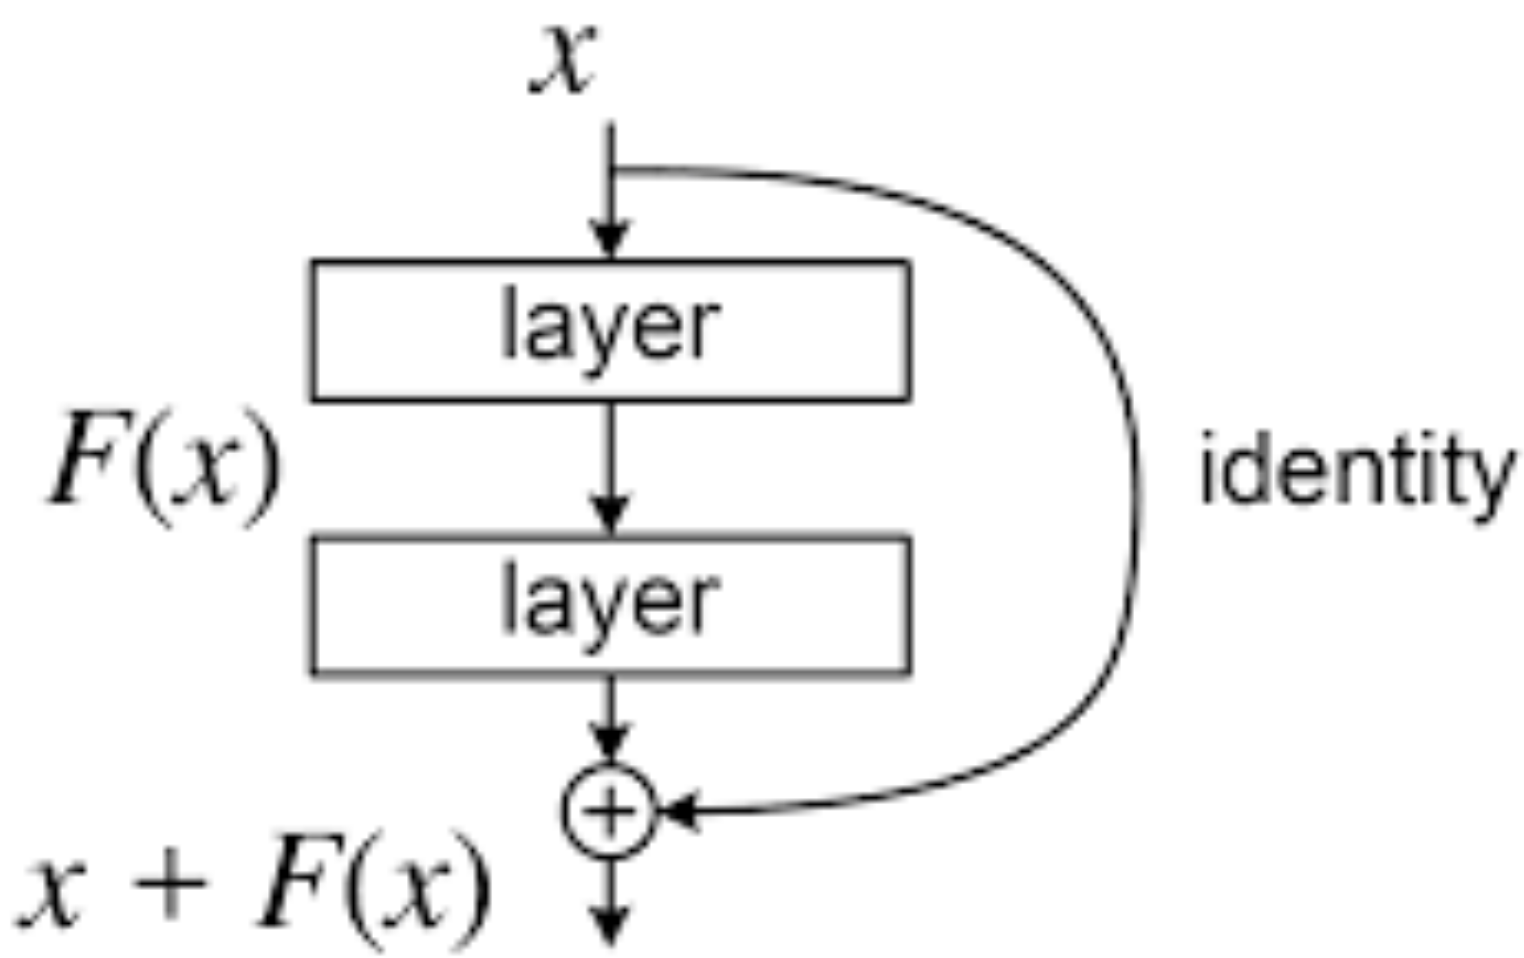

Residual connection stabilizes gradient backflow.

Post-AlexNet Developments

(1) Residual Connections and ResNets

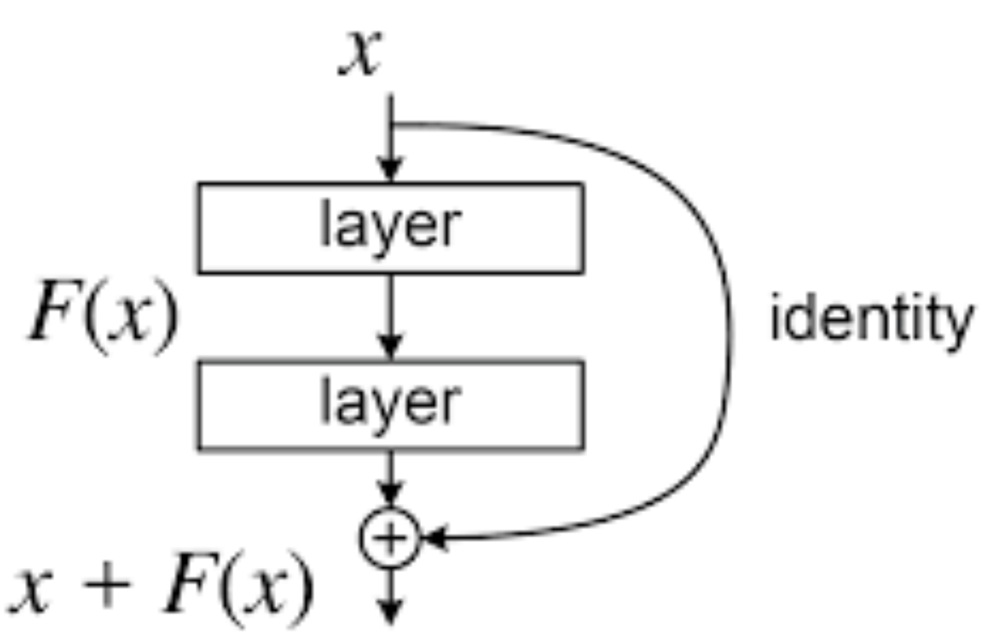

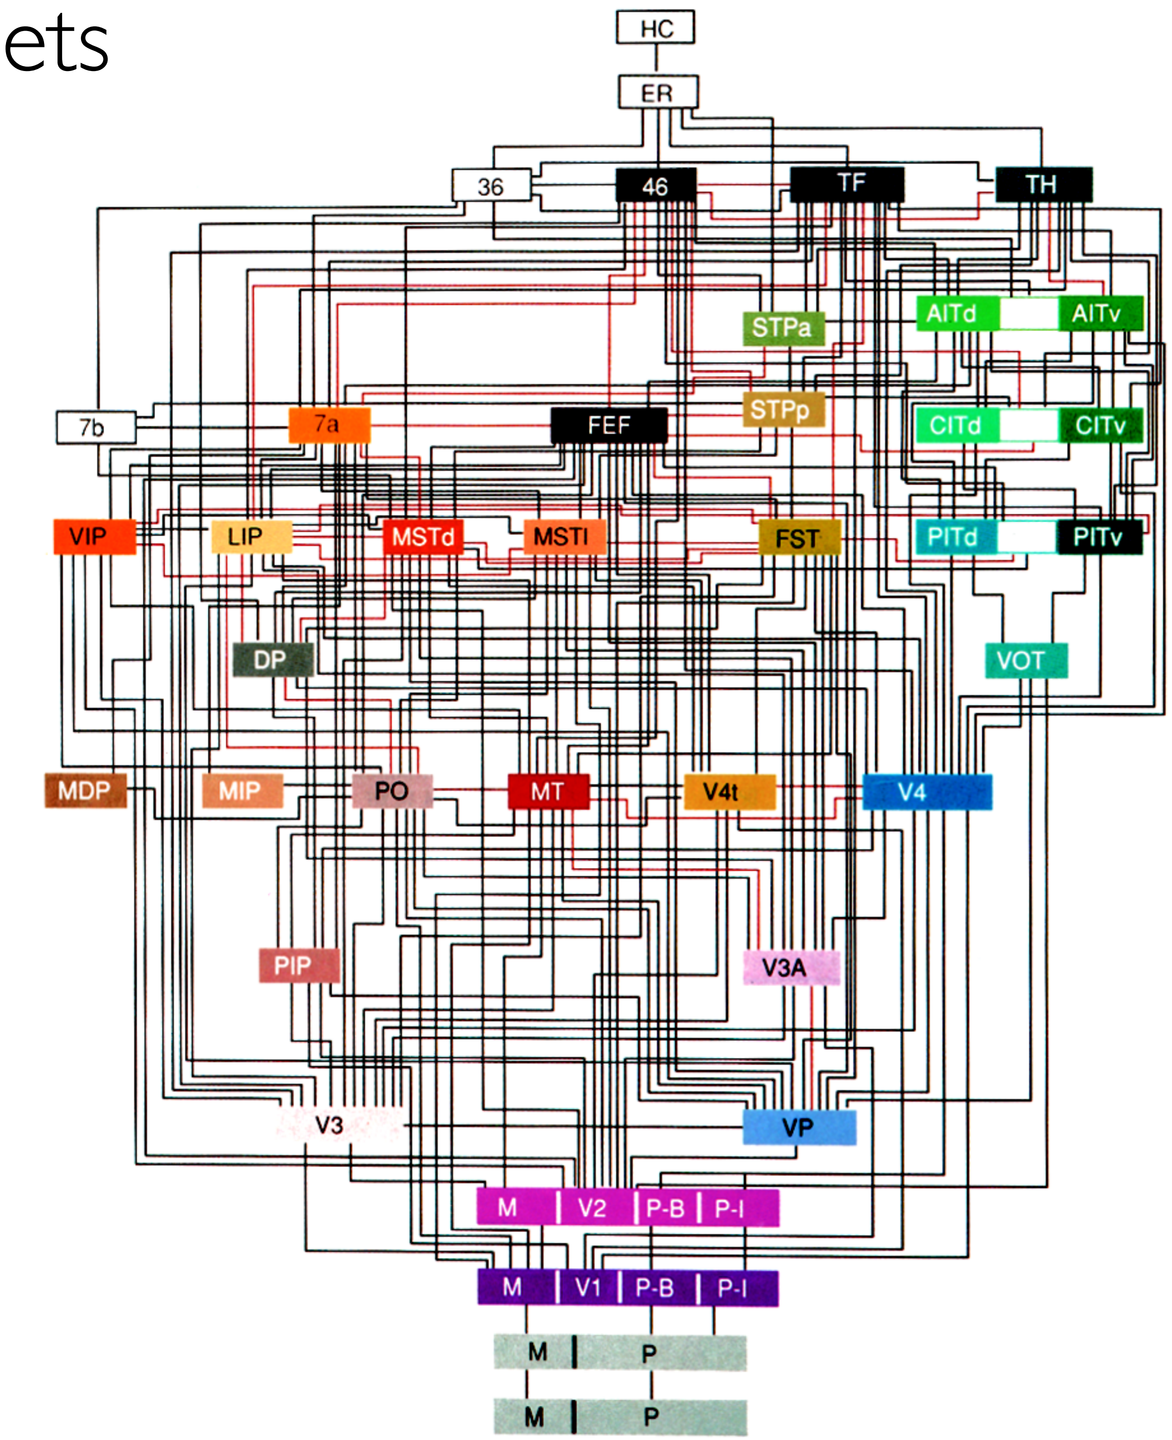

Residual connection stabilizes gradient backflow.

Lots of skip connections present in actual brain.

Post-AlexNet Developments

(2) Vision Transformers

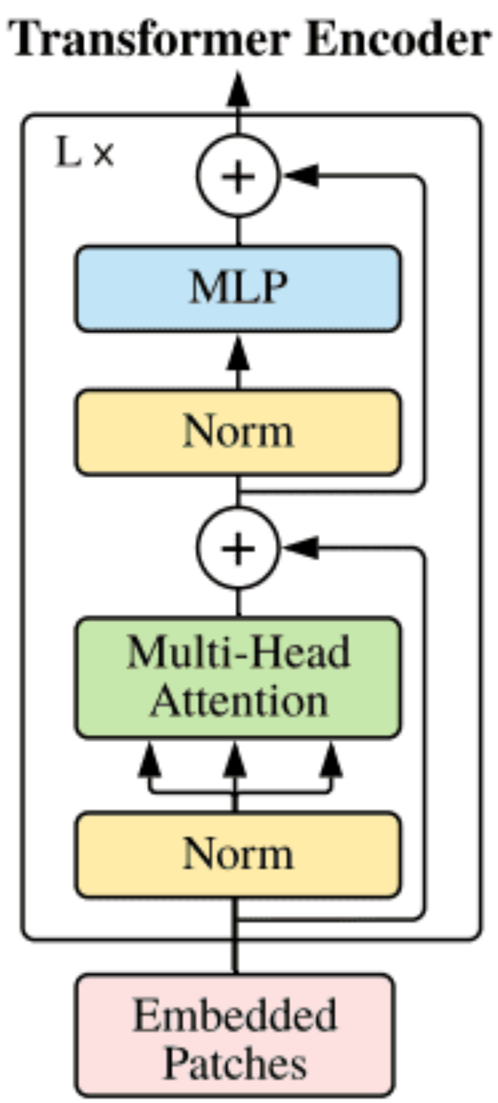

Post-AlexNet Developments

(2) Vision Transformers

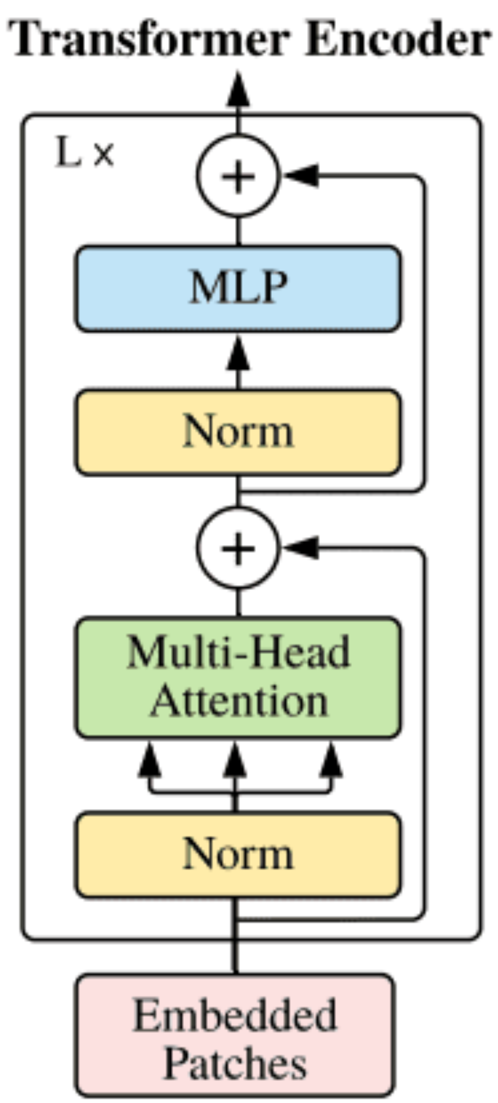

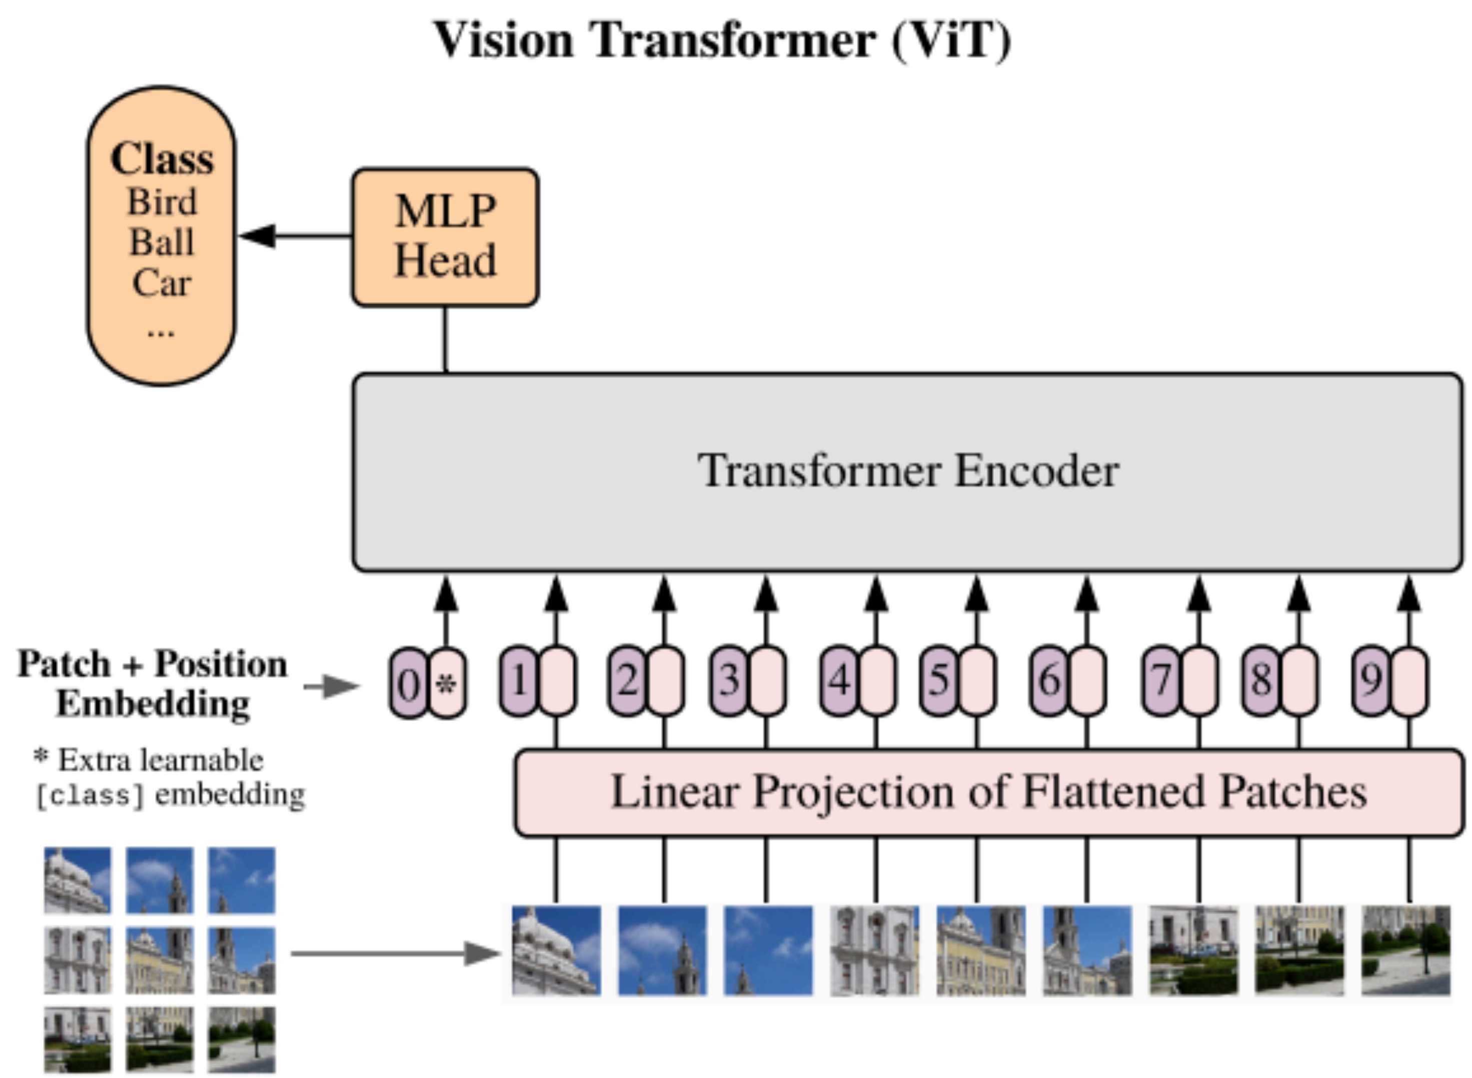

Post-AlexNet Developments

(2) Vision Transformers

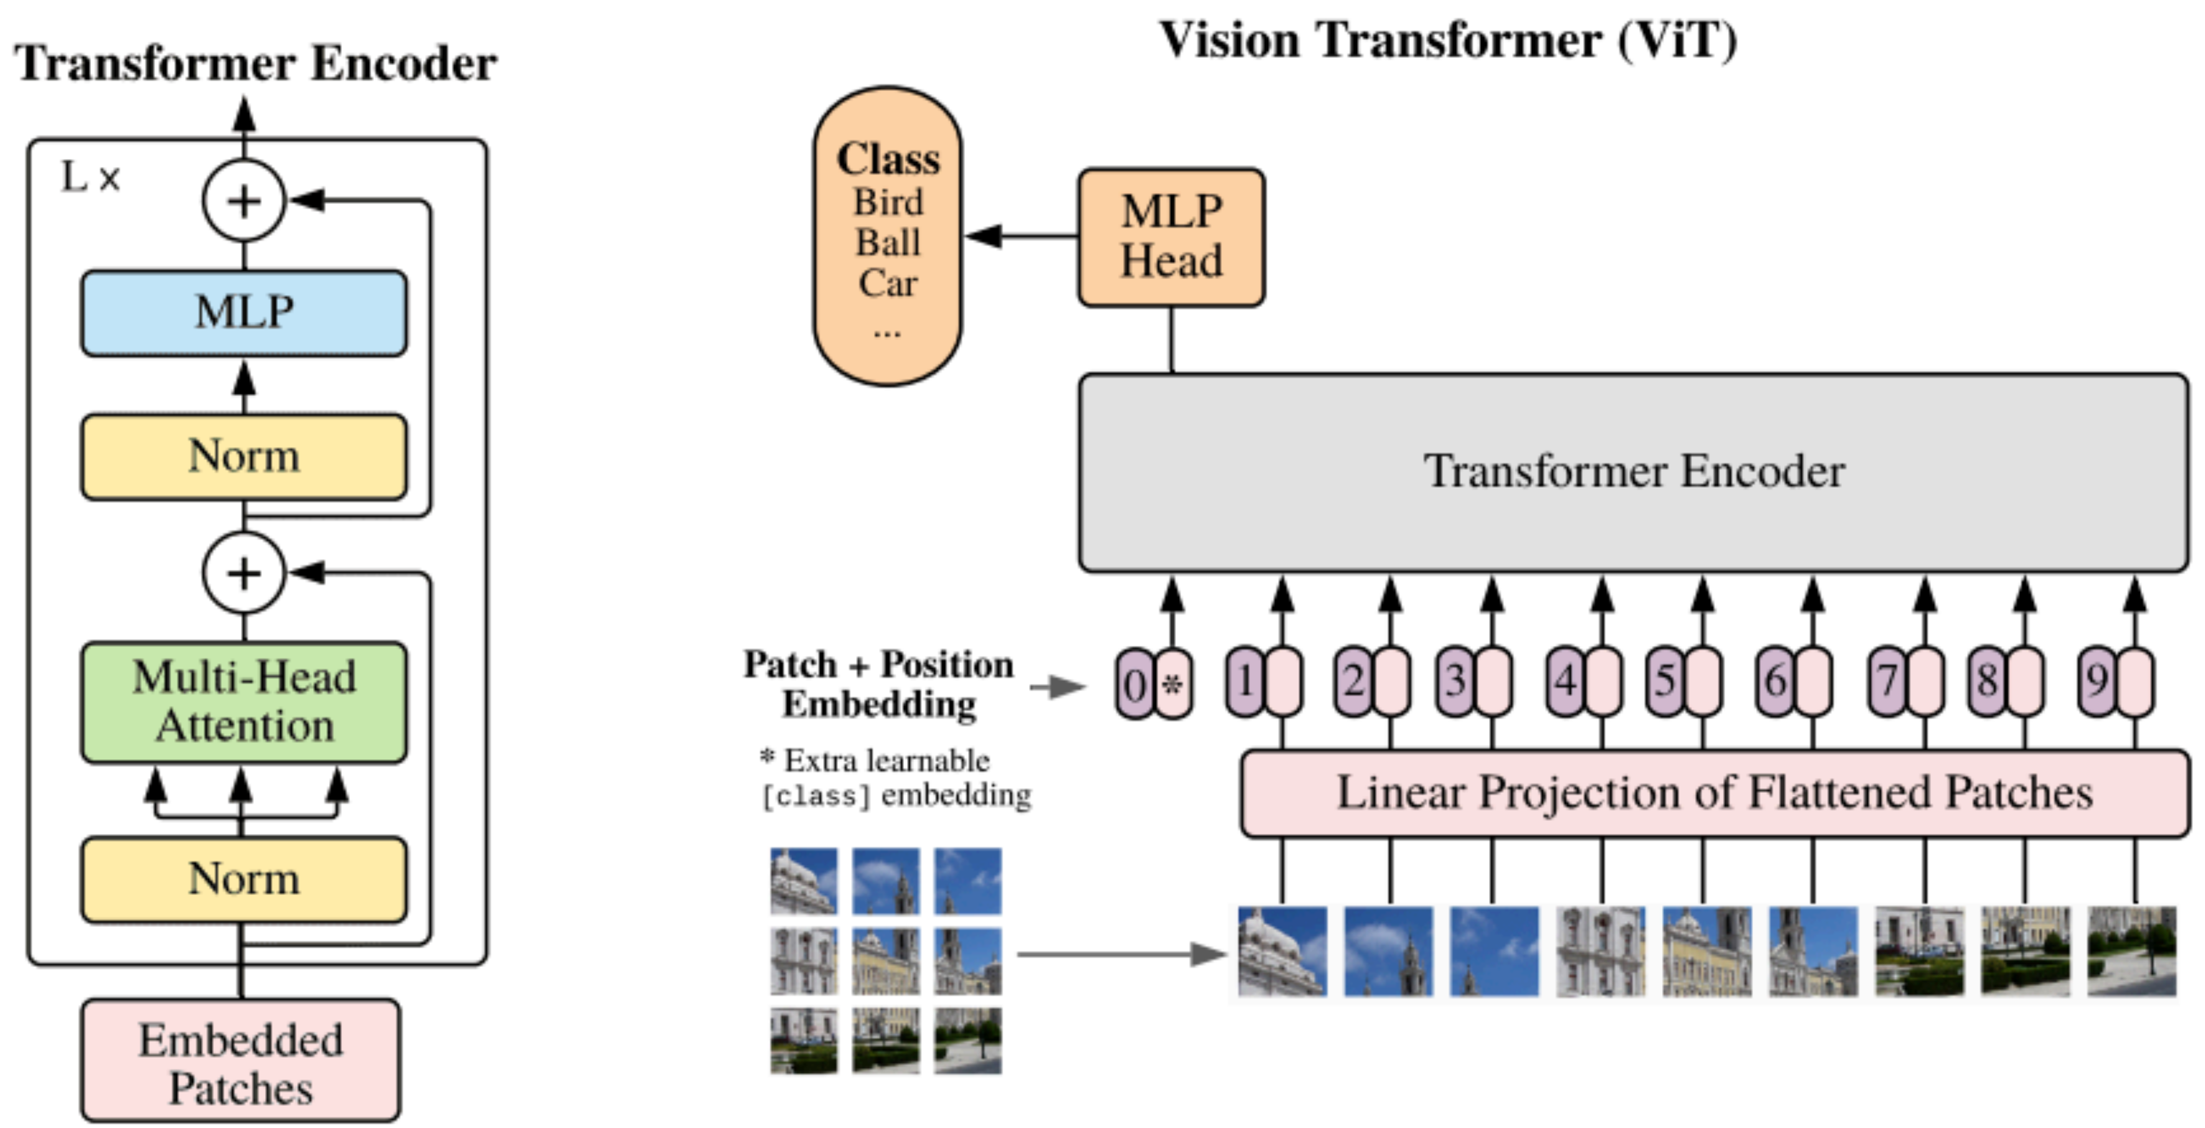

NB: still hierarchical, still with residual connections, potential locality from patches . . .

Post-AlexNet Developments

(2) Vision Transformers

Looking at receptive field analysis of ViTs vs ResNet:

Post-AlexNet Developments

(2) Vision Transformers

Looking at receptive field analysis of ViTs vs ResNet:

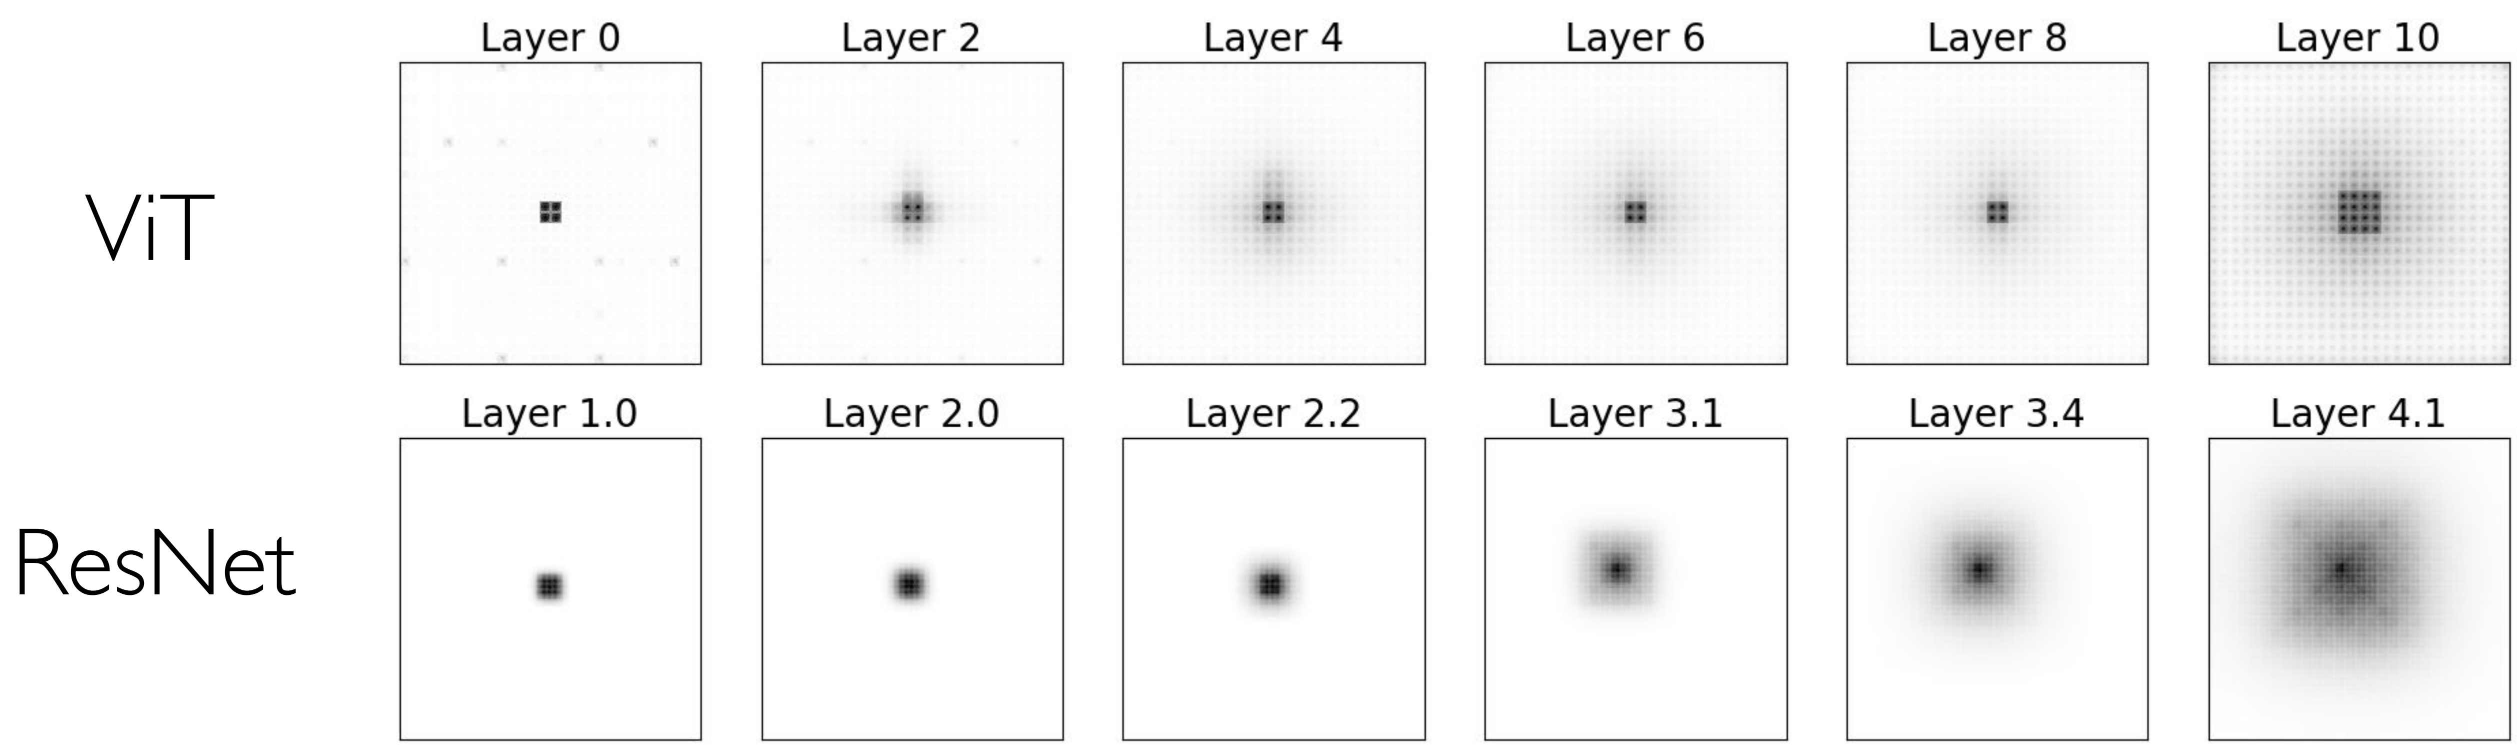

Post-AlexNet Developments

(2) Vision Transformers

Looking at receptive field analysis of ViTs vs ResNet:

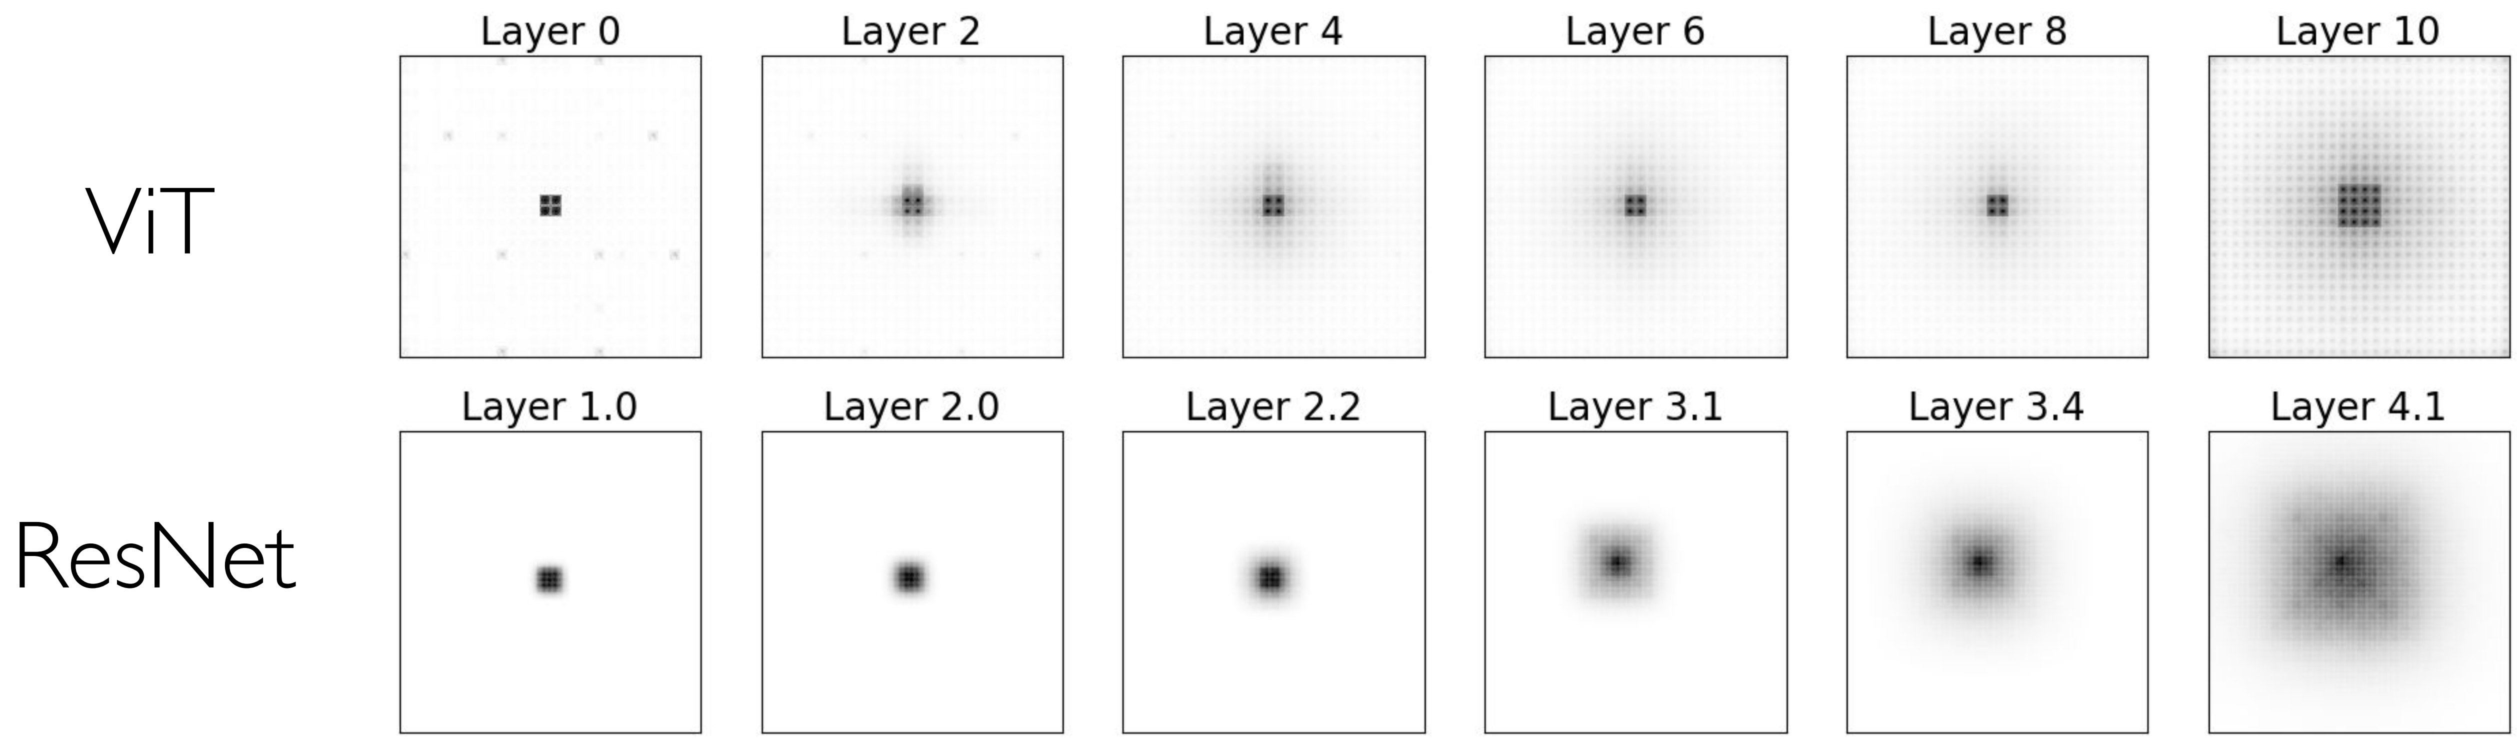

. . . we see learned ViT is <u>mostly local</u>, with increasing receptive field sizes.

Post-AlexNet Developments

(2) Vision Transformers

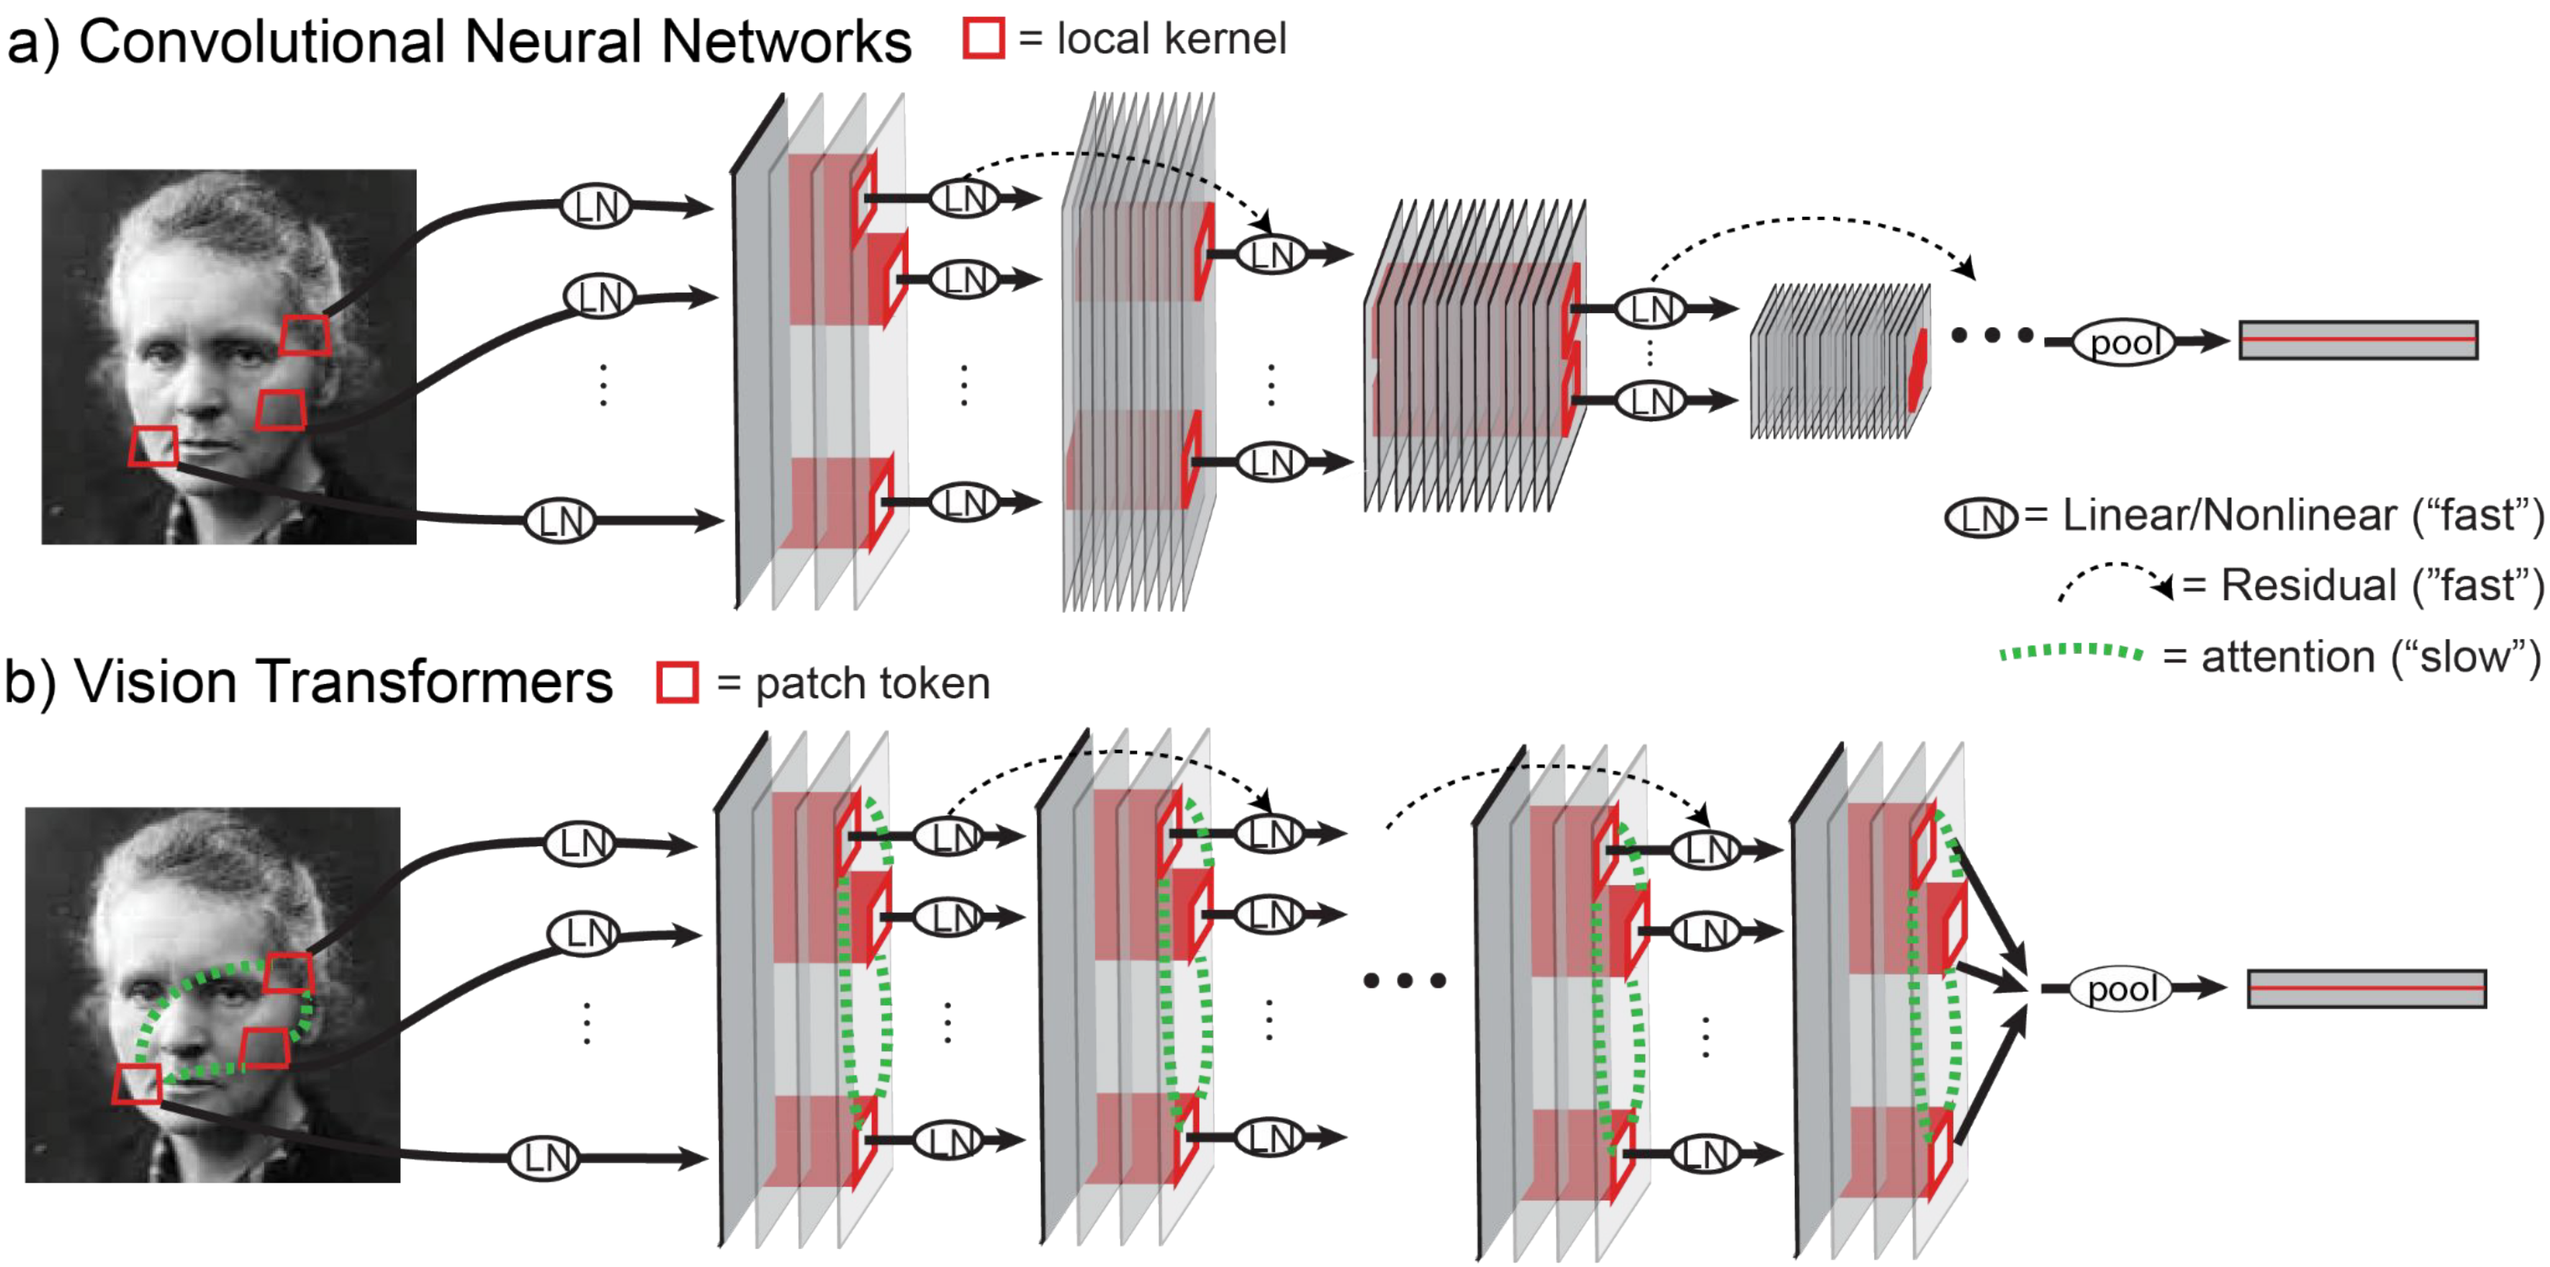

<u>ViT is a bit like a CNN with sparse global connections.</u>

Principles of Visual Architecture

(1) Hierarchical (2) Mostly local (3) Rectification-like nonlinearity

(4) Some residual connections (5) Normalization

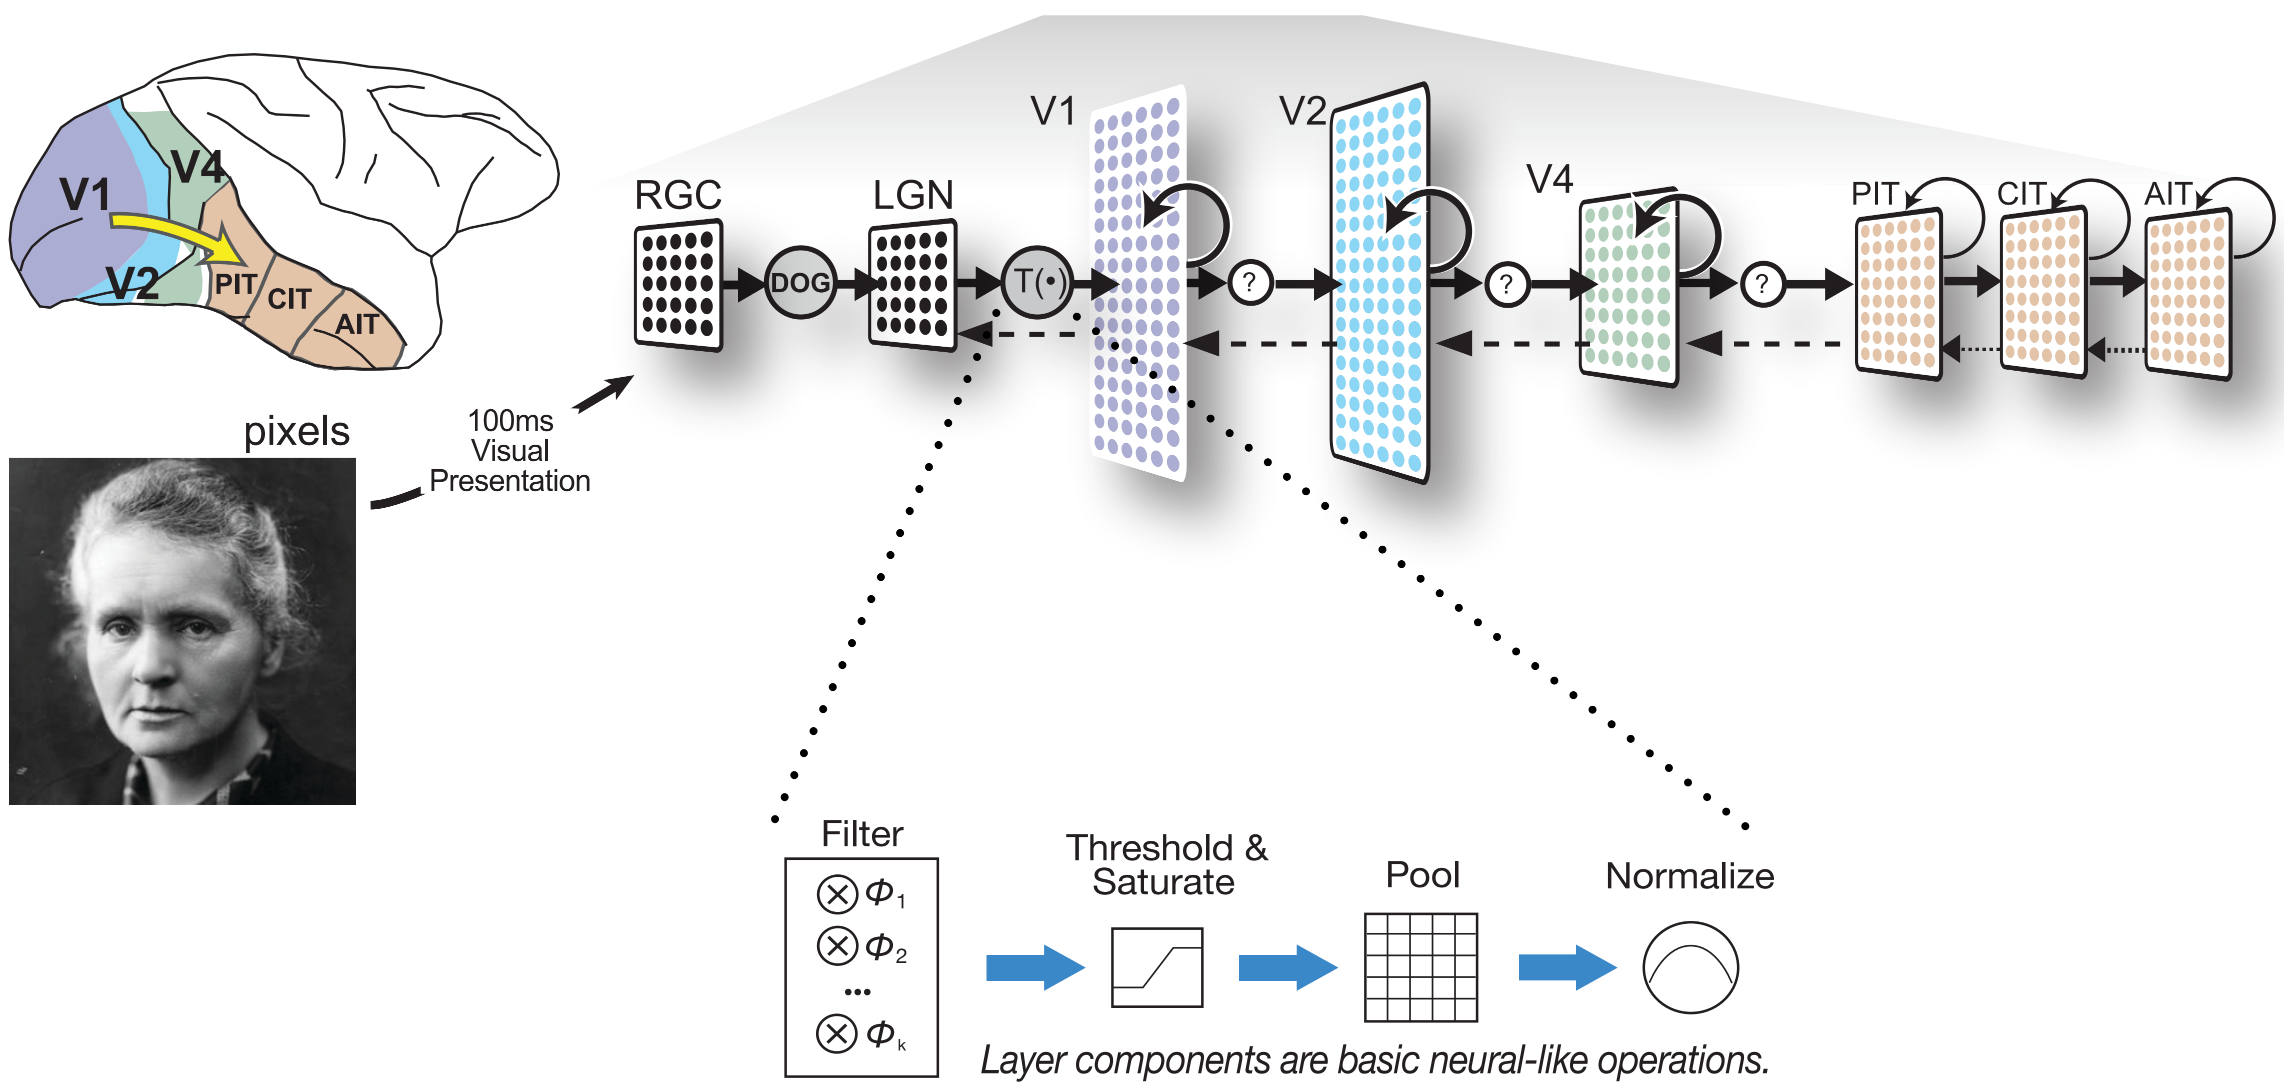

Behavioral "Top-Down" constraints

Complement standard "from below" approach …

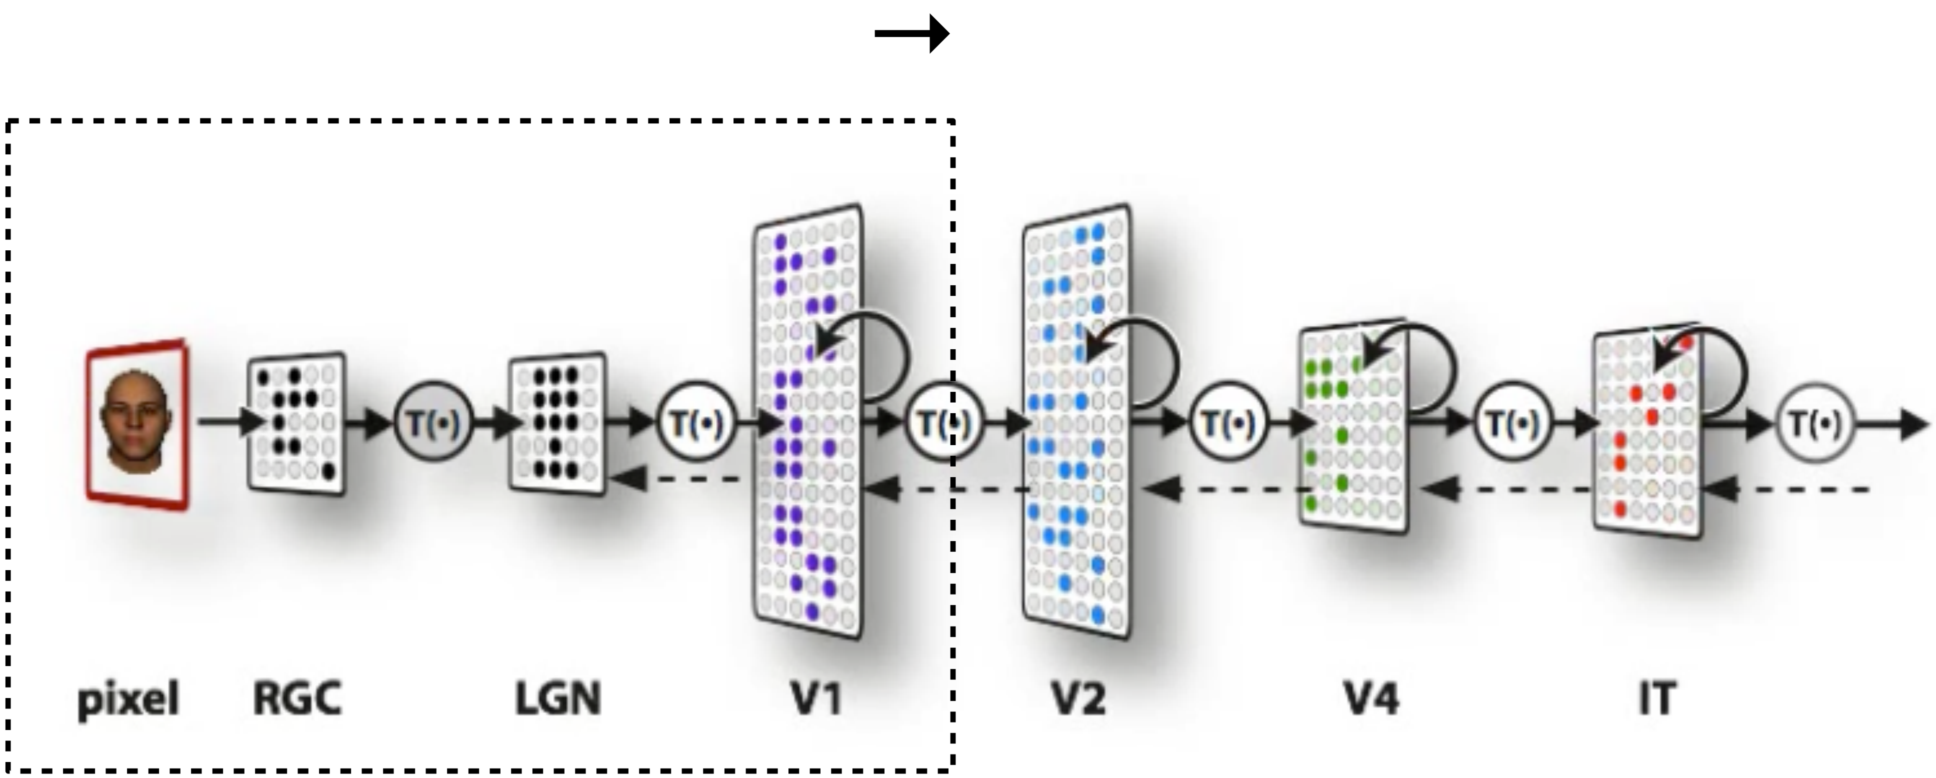

Behavioral "Top-Down" constraints

Complement standard "from below" approach …

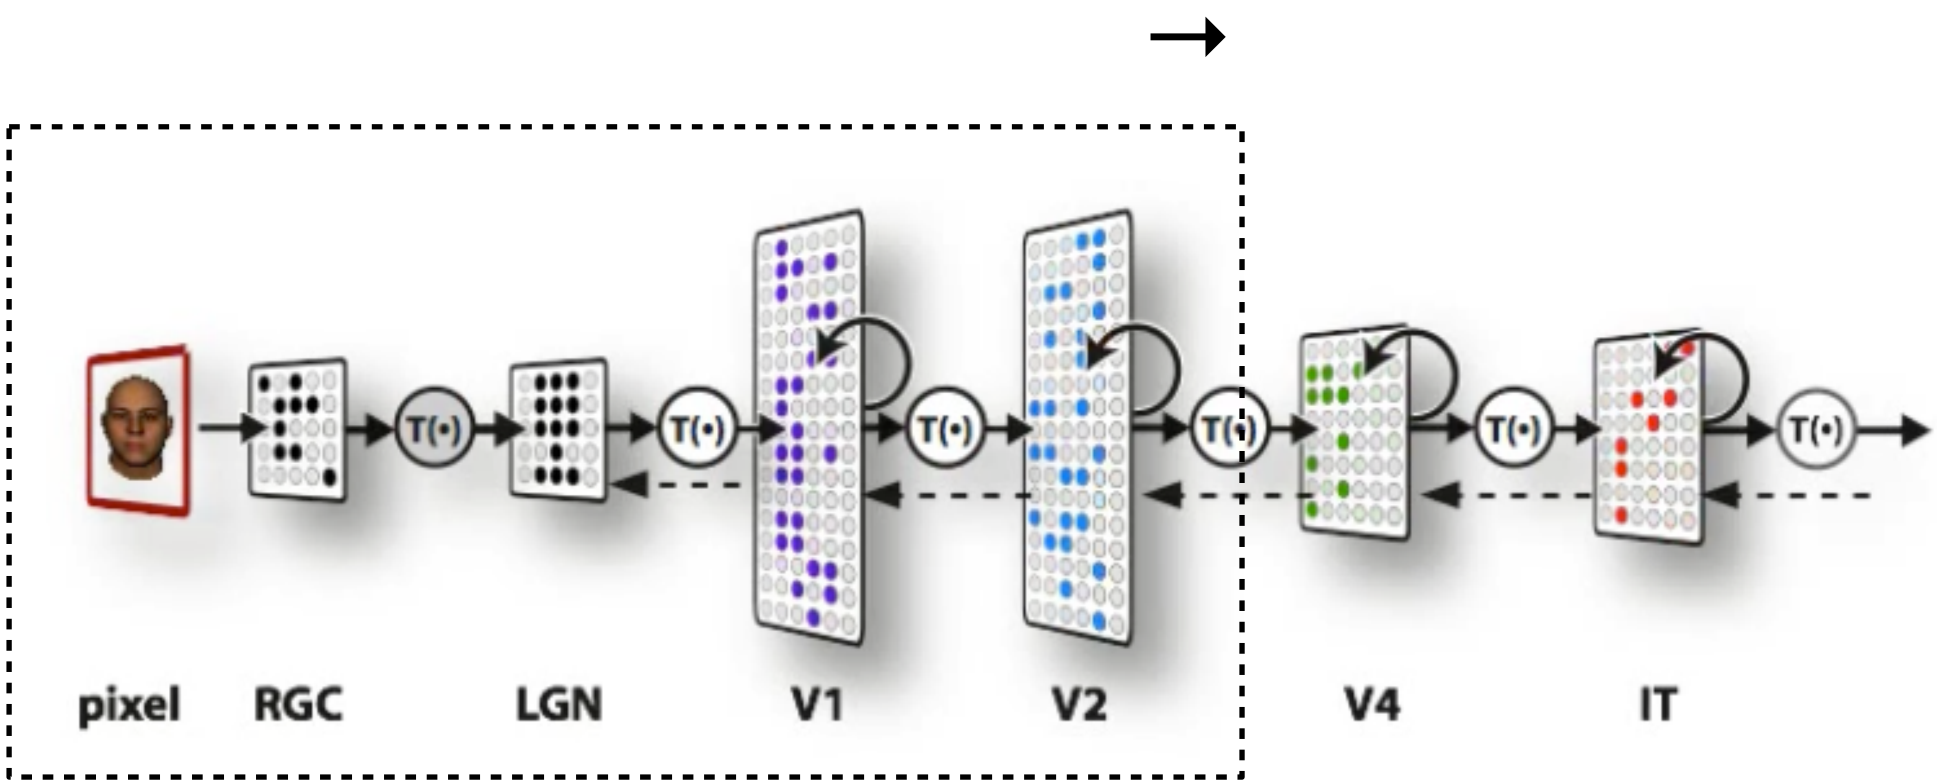

Behavioral "Top-Down" constraints

Complement standard "from below" approach …

→

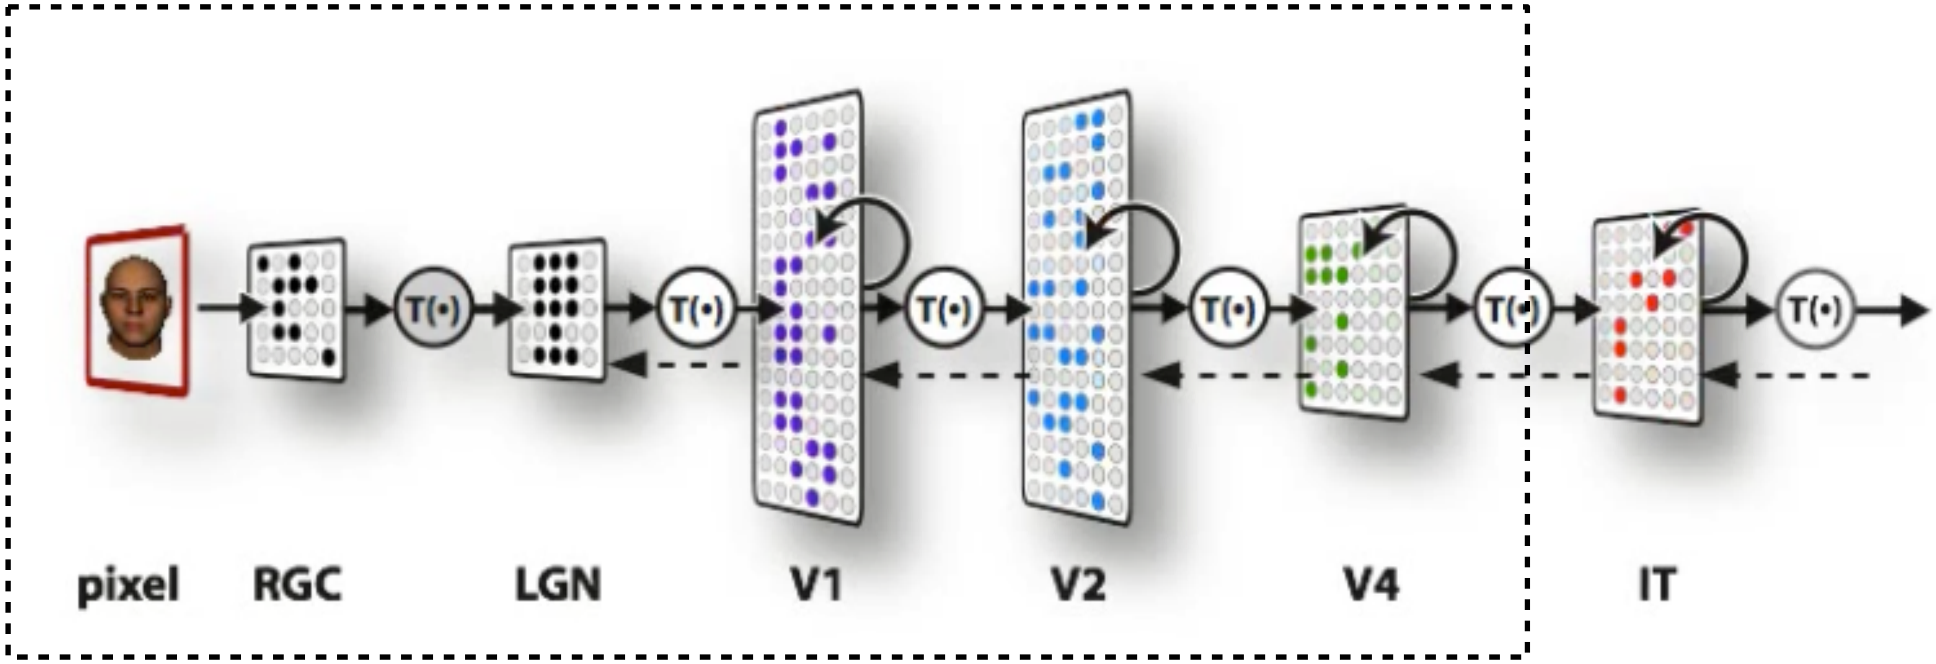

Behavioral "Top-Down" constraints

Complement standard "from below" approach …

→

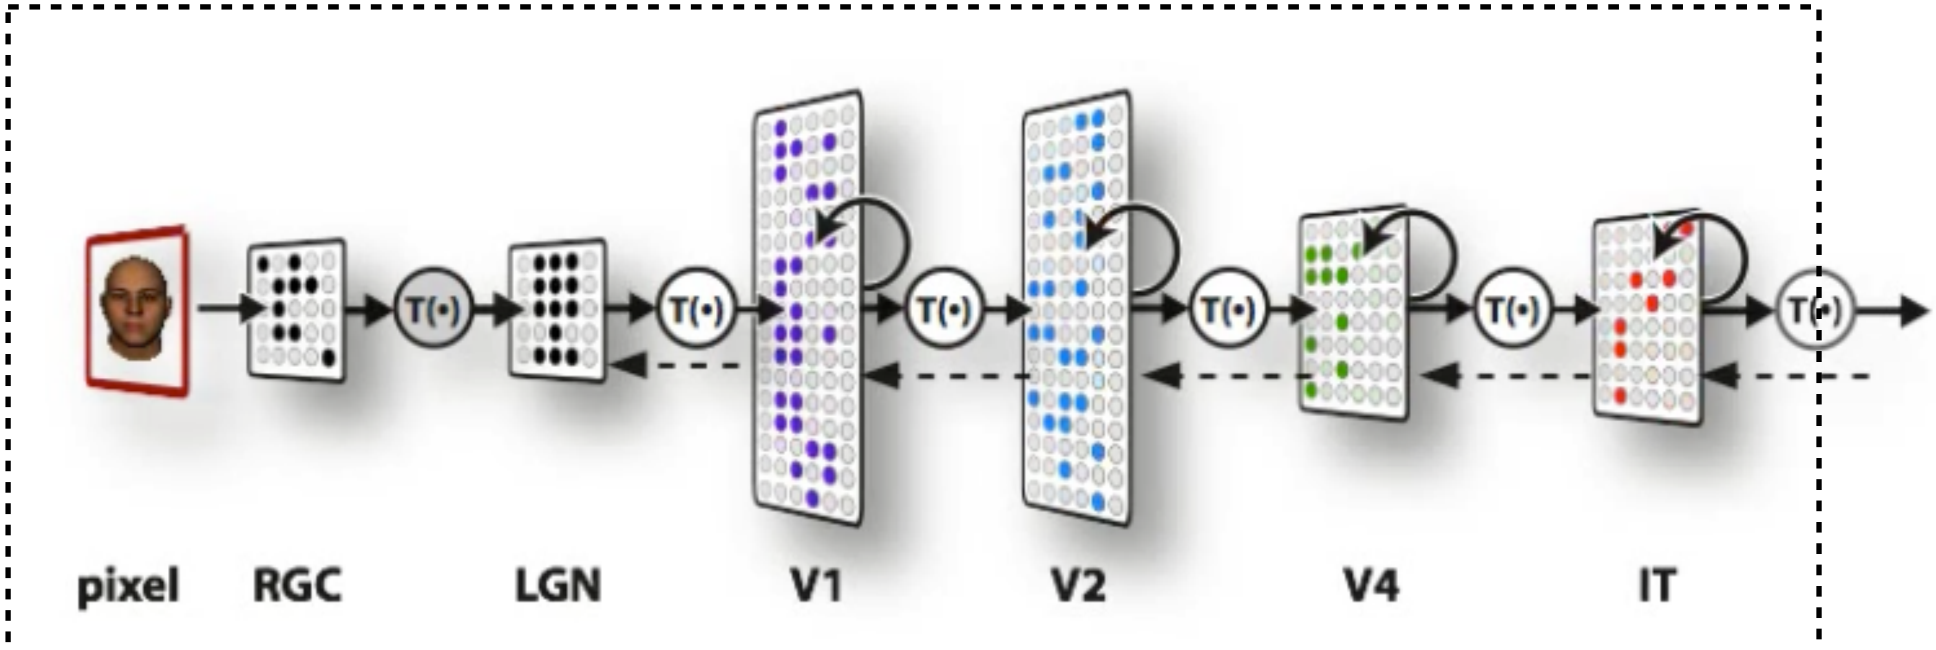

Behavioral "Top-Down" constraints

Complement standard "from below" approach … with behavioral constraints

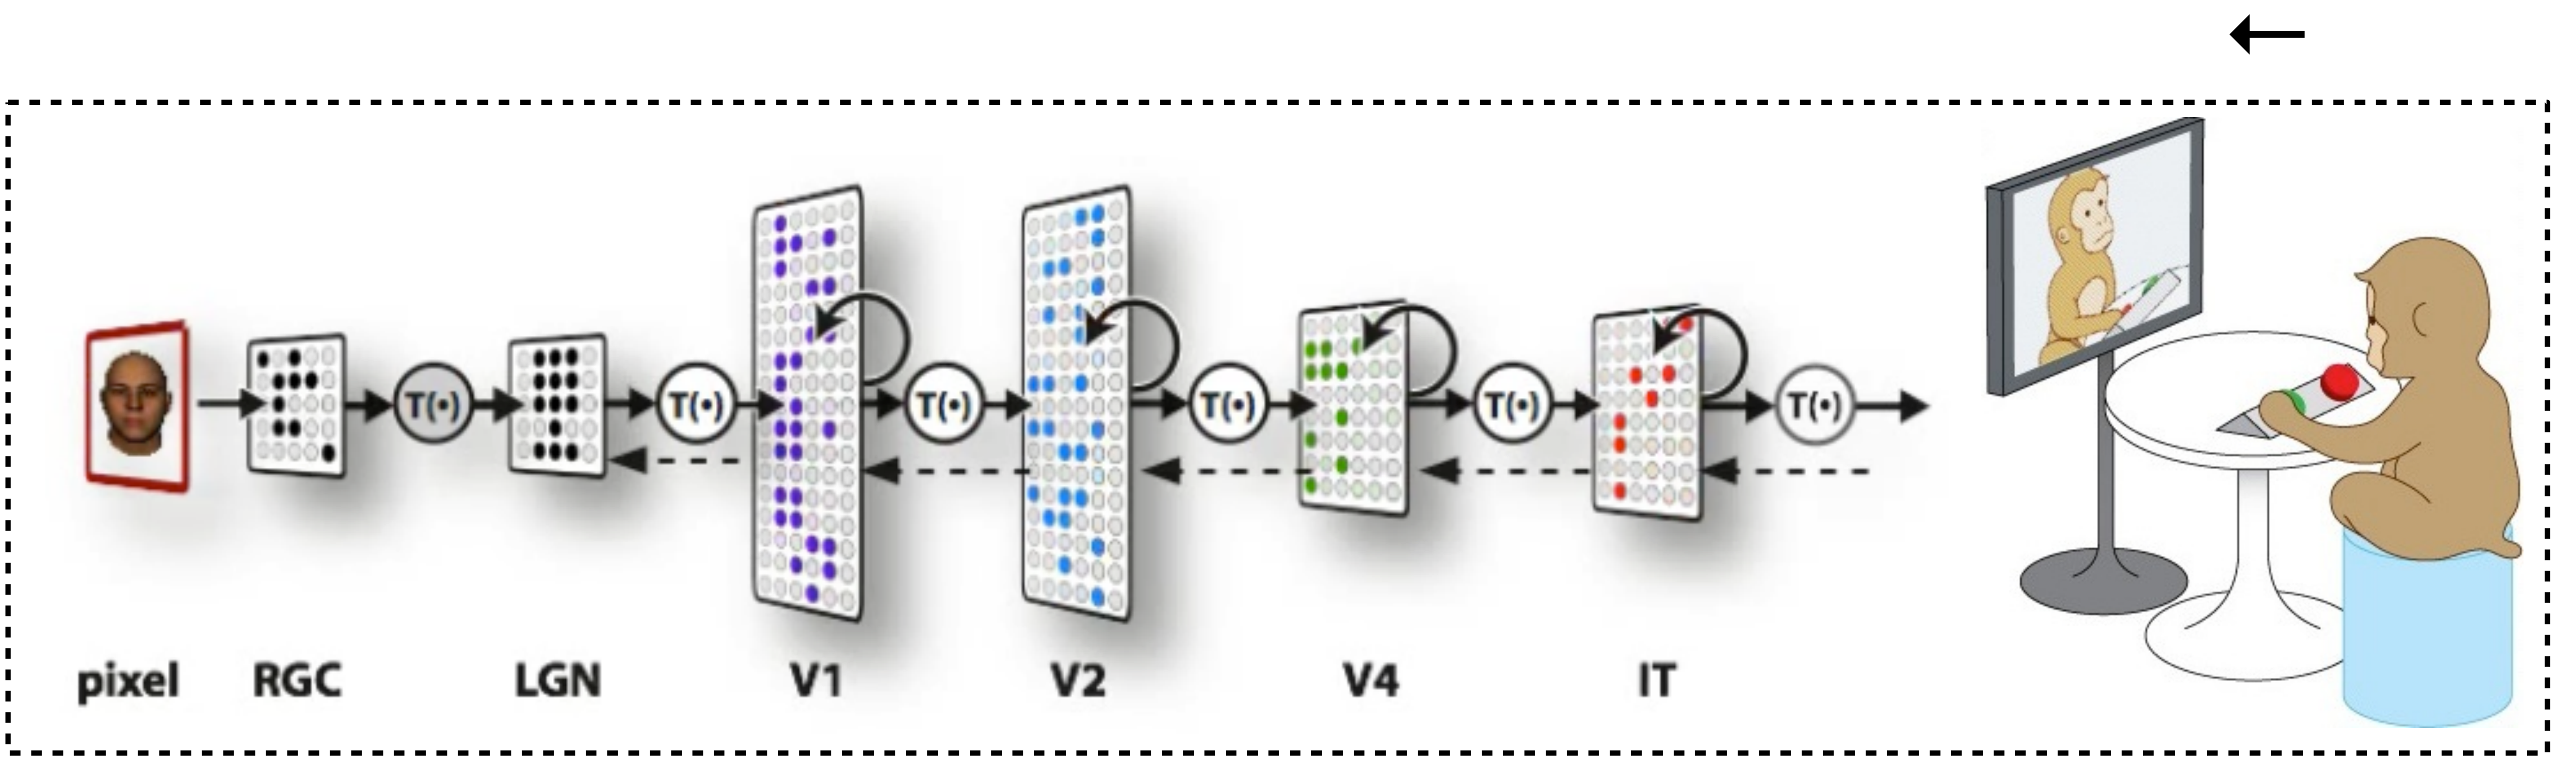

Behavioral "Top-Down" constraints

Complement standard "from below" approach … with behavioral constraints

←

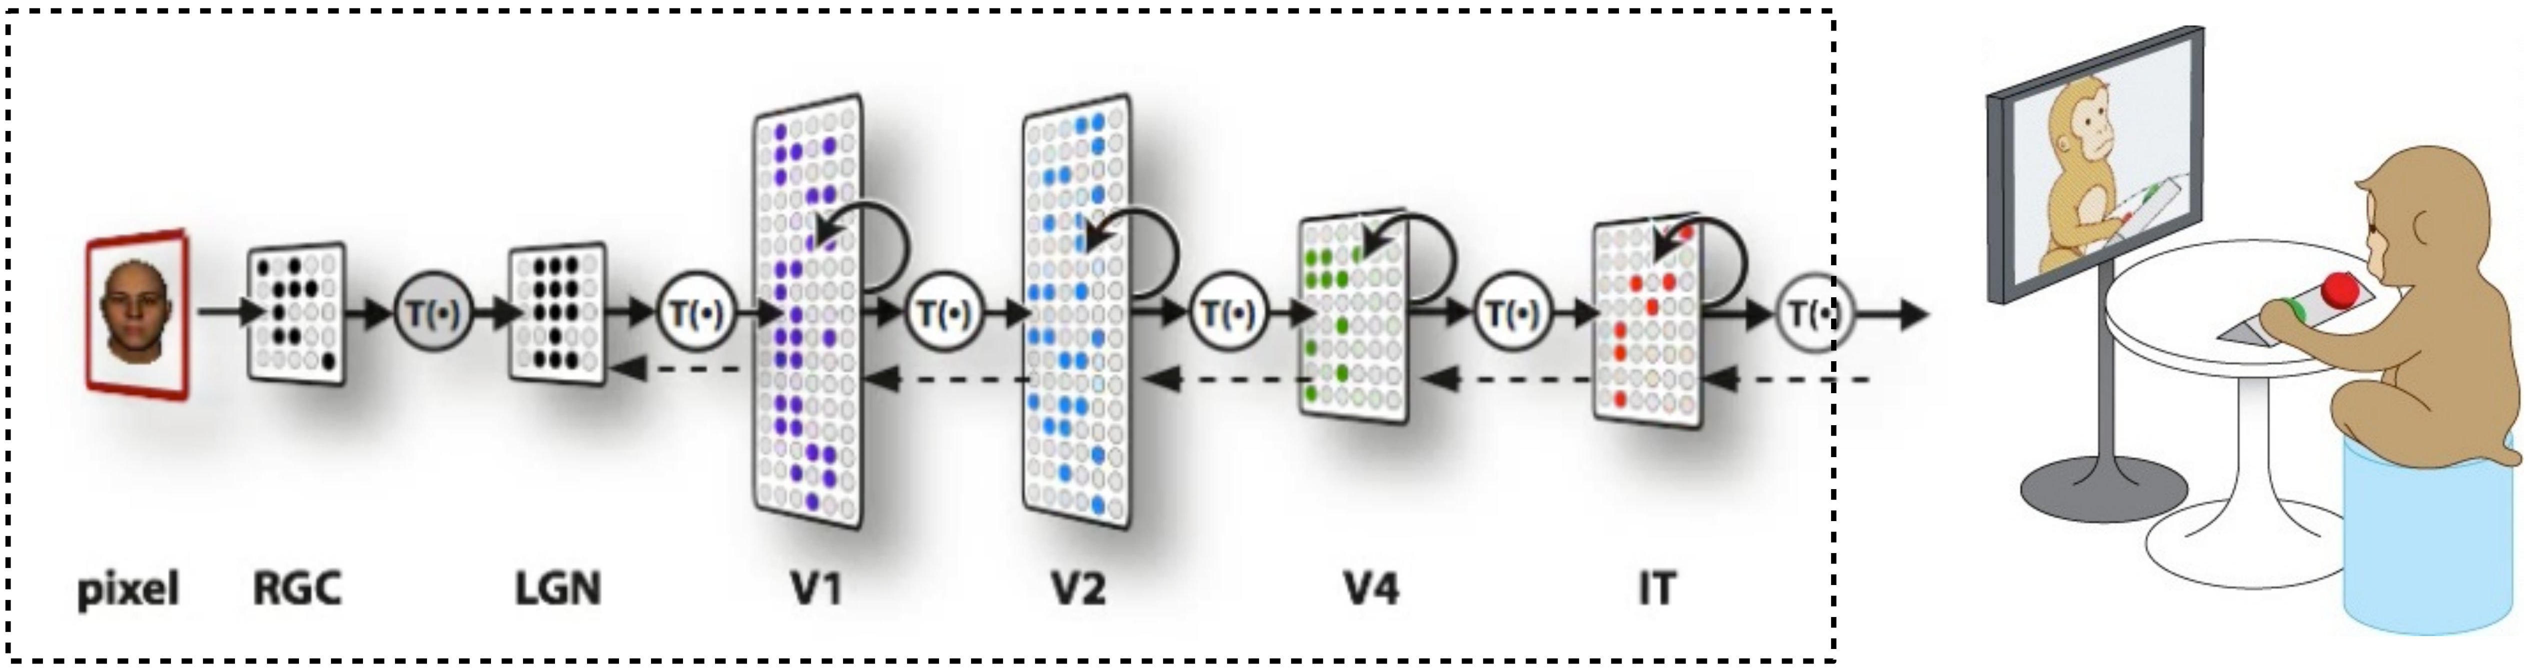

Behavioral "Top-Down" constraints

Complement standard "from below" approach … with behavioral constraints

←

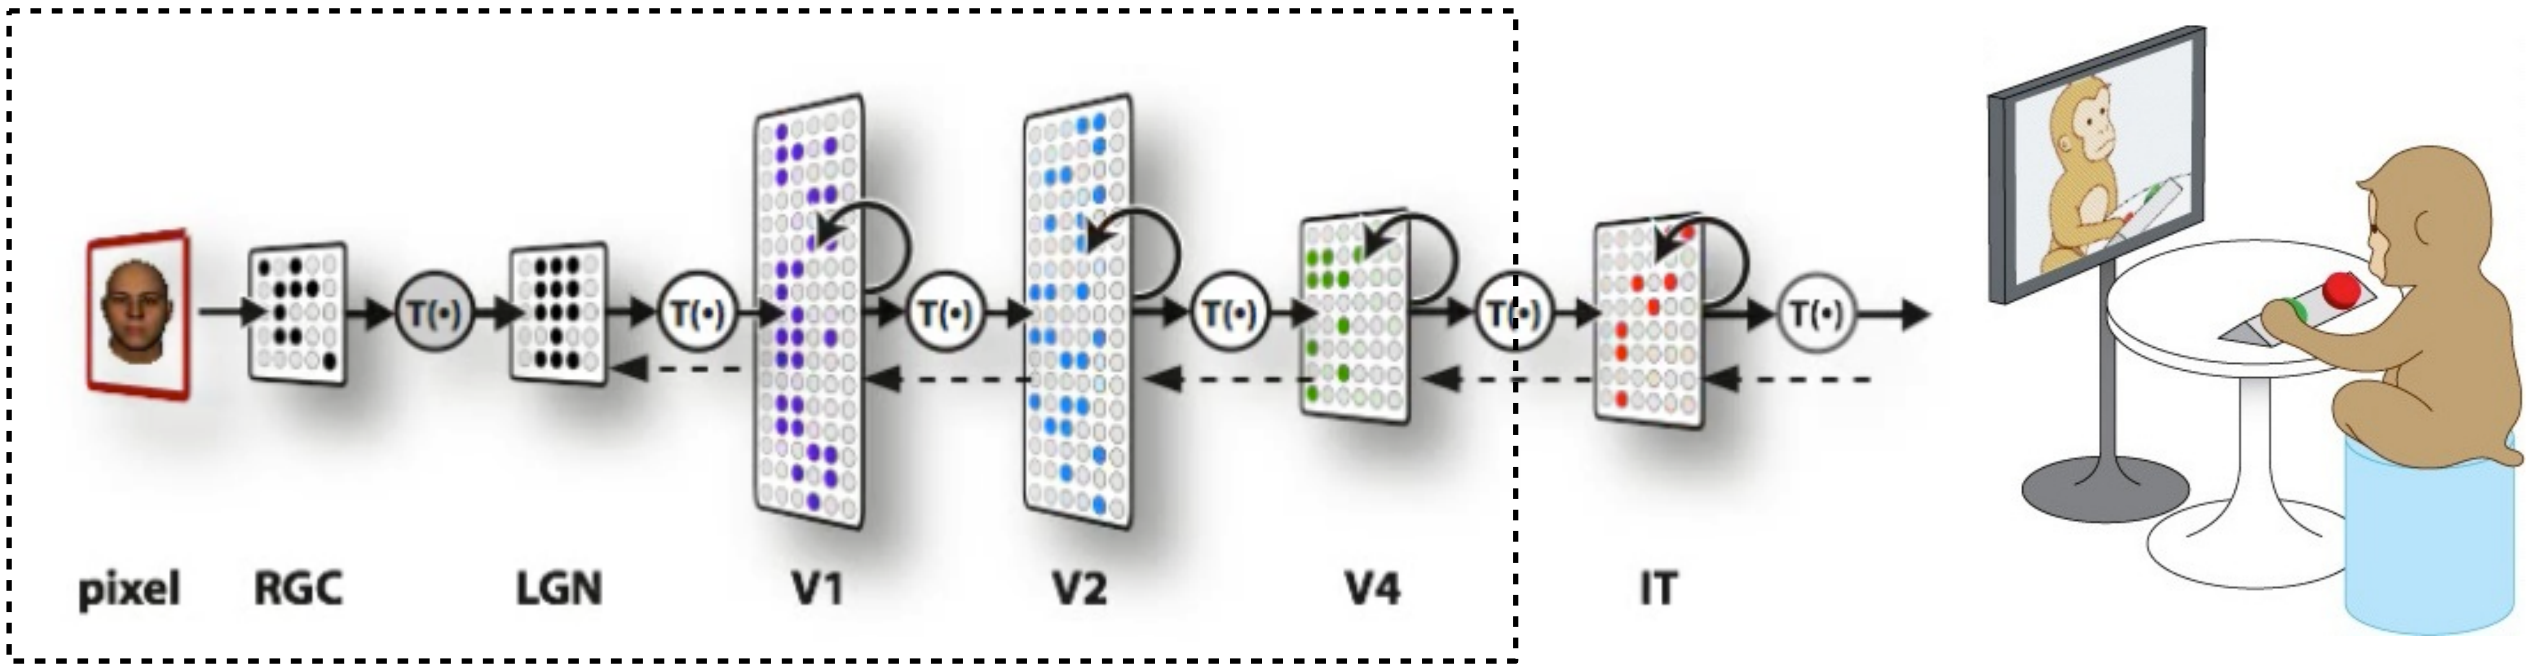

Behavioral "Top-Down" constraints

Complement standard "from below" approach … with behavioral constraints

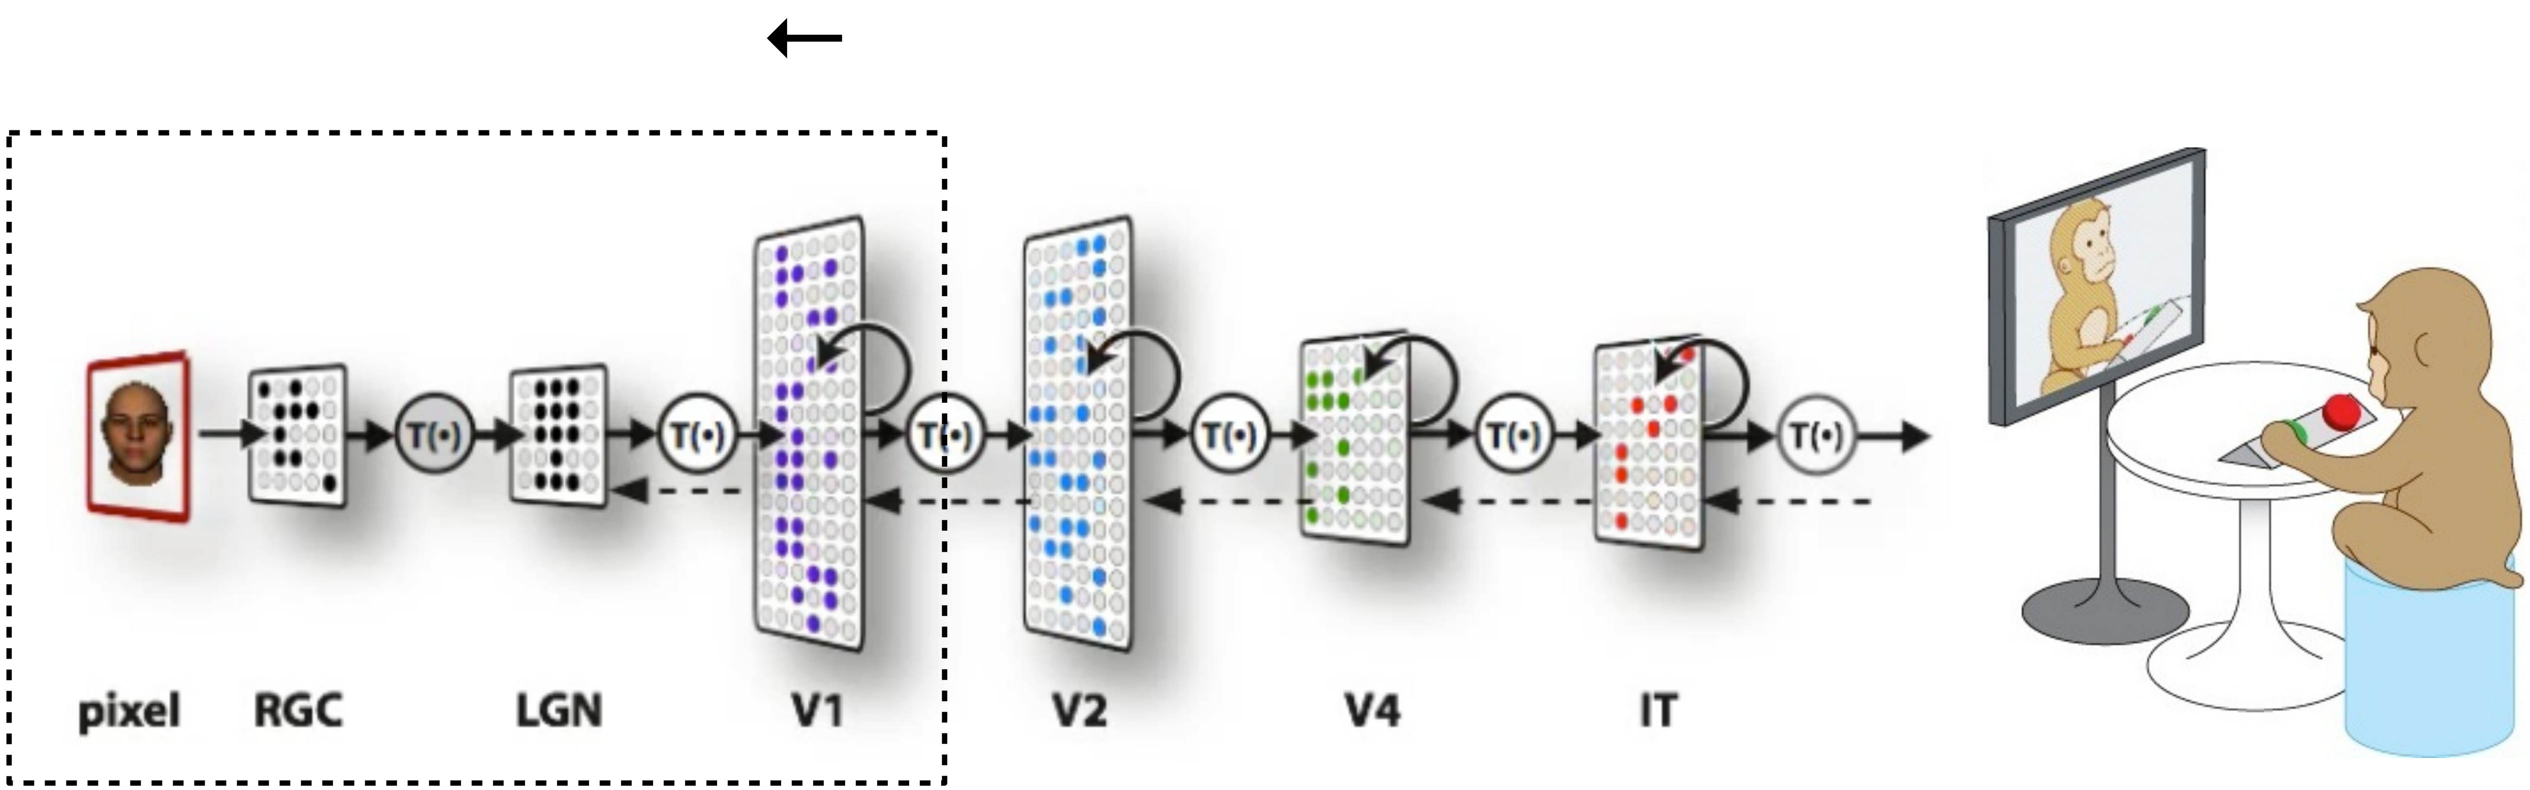

Beyond categorization

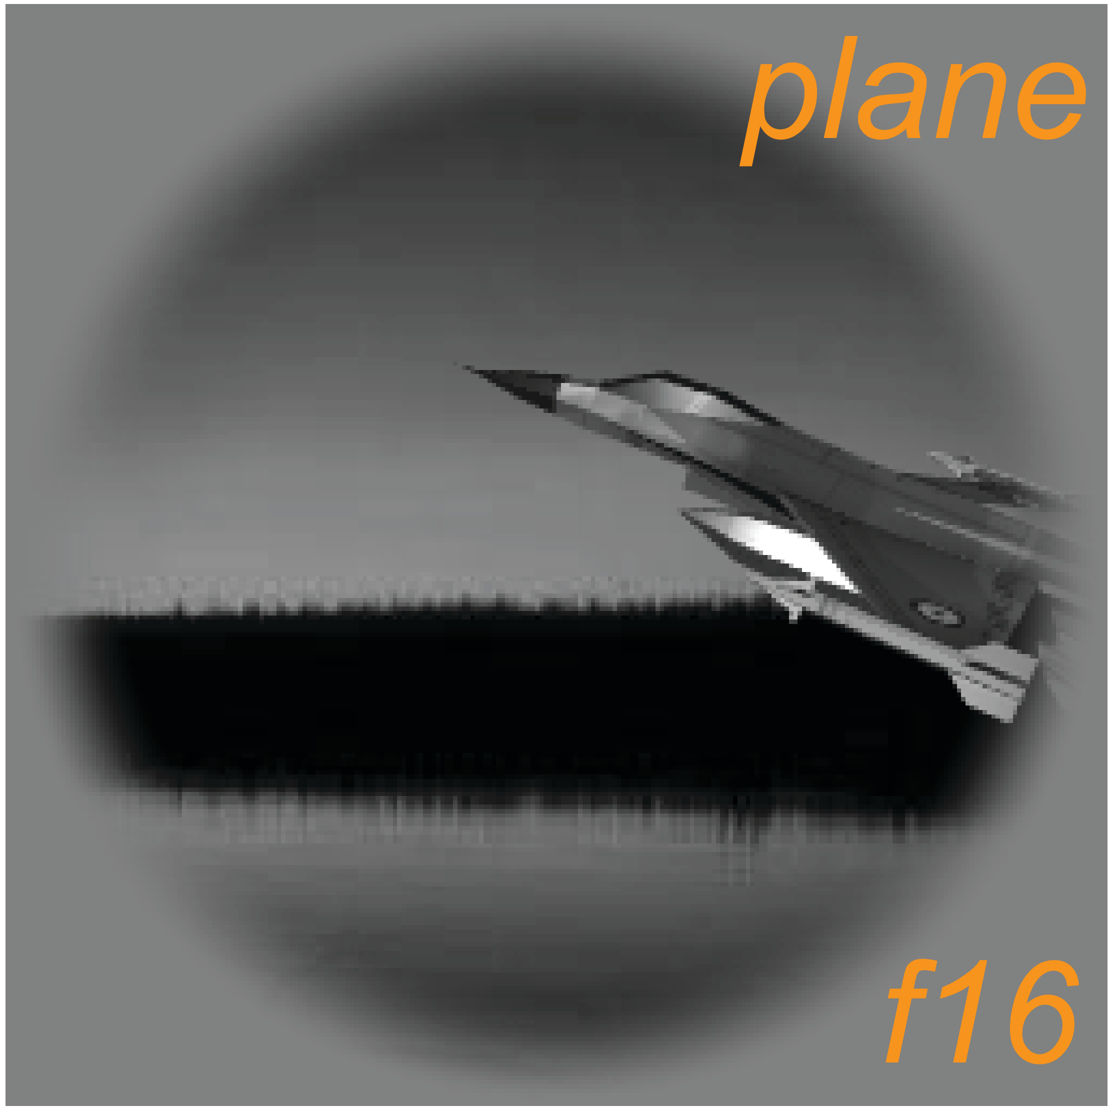

Category

Identity

Beyond categorization

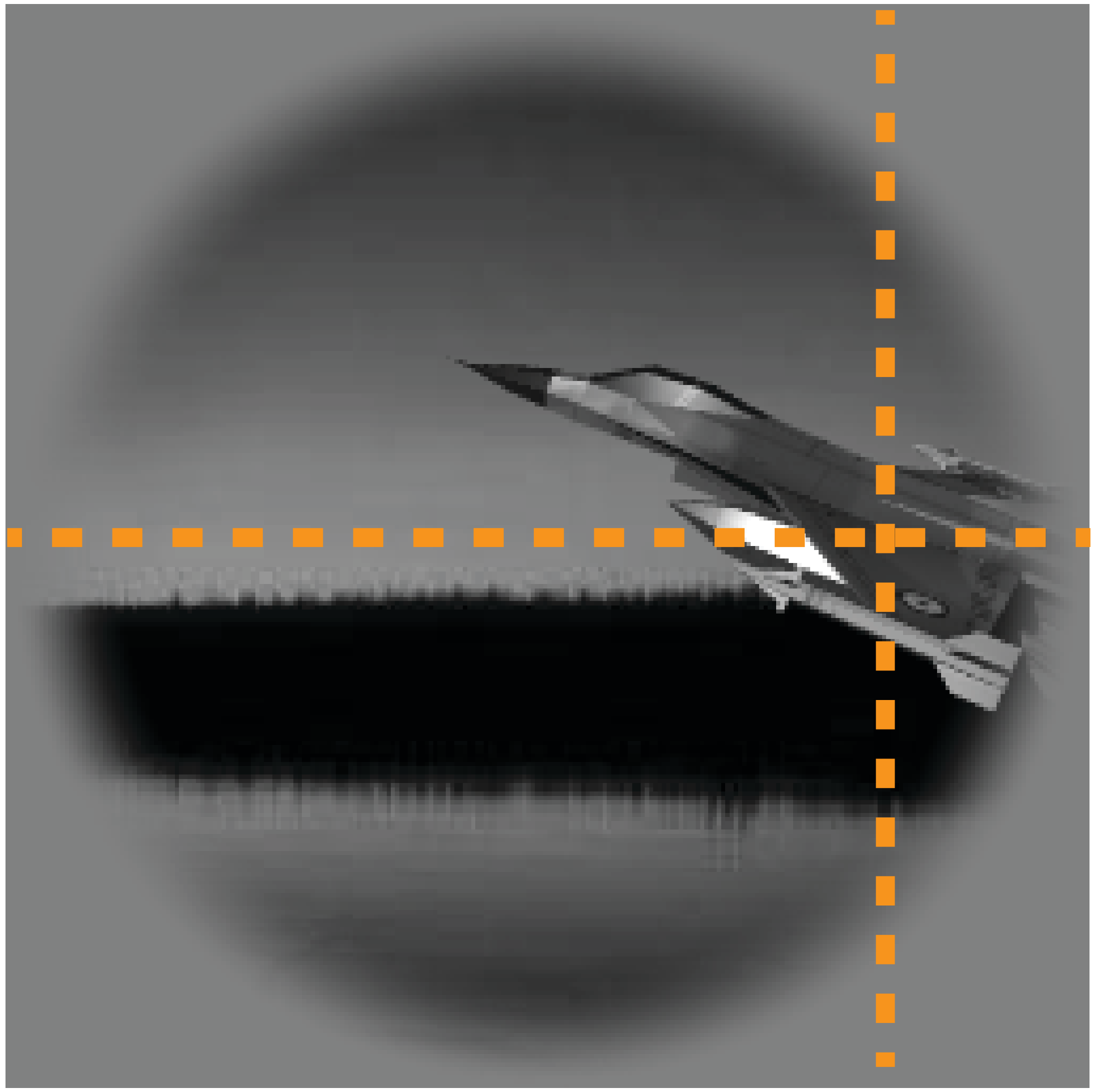

Position

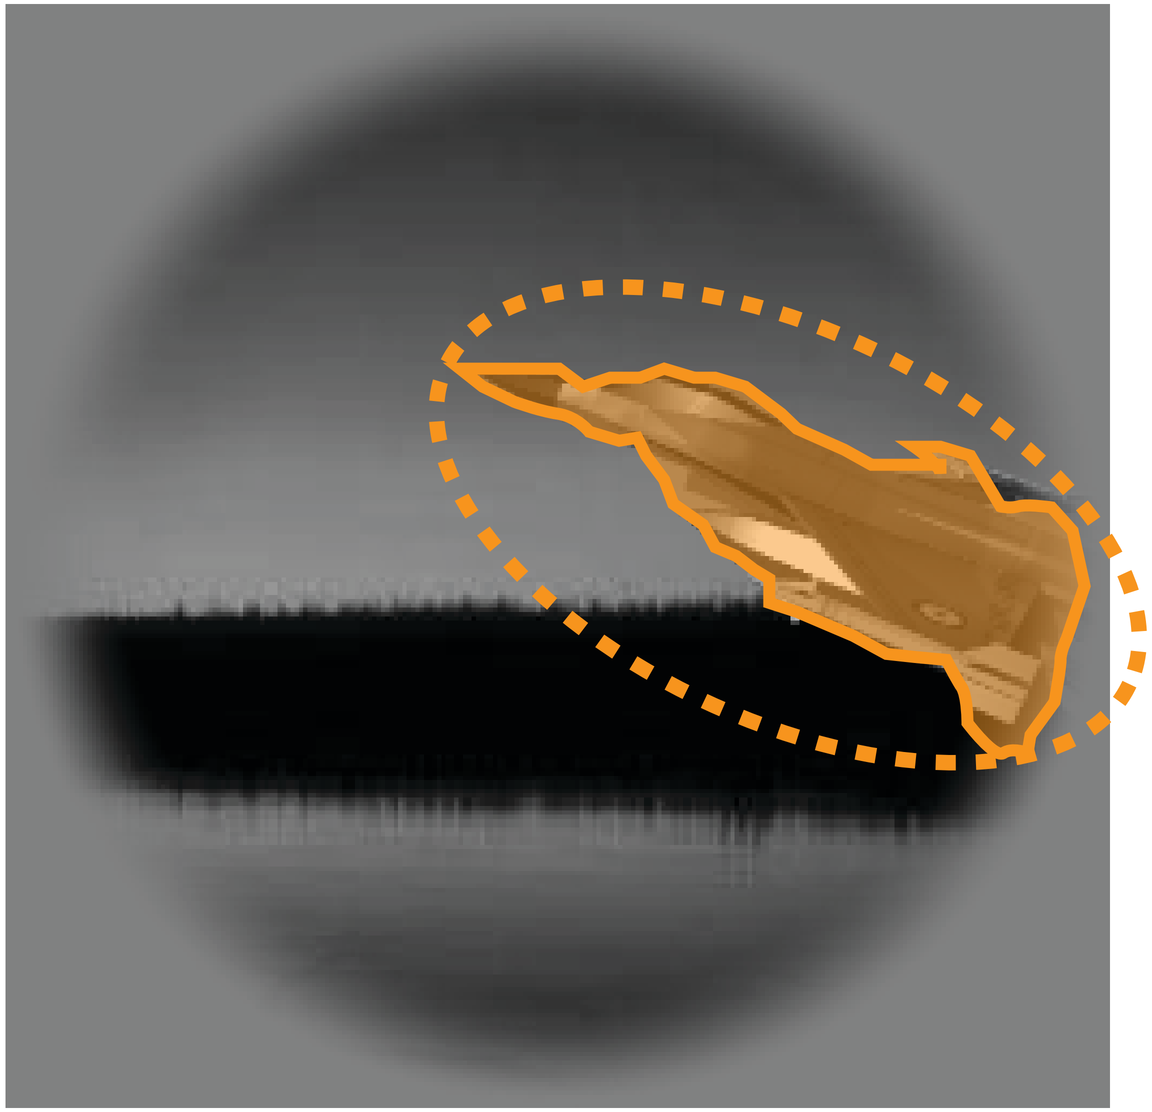

Size

Beyond categorization

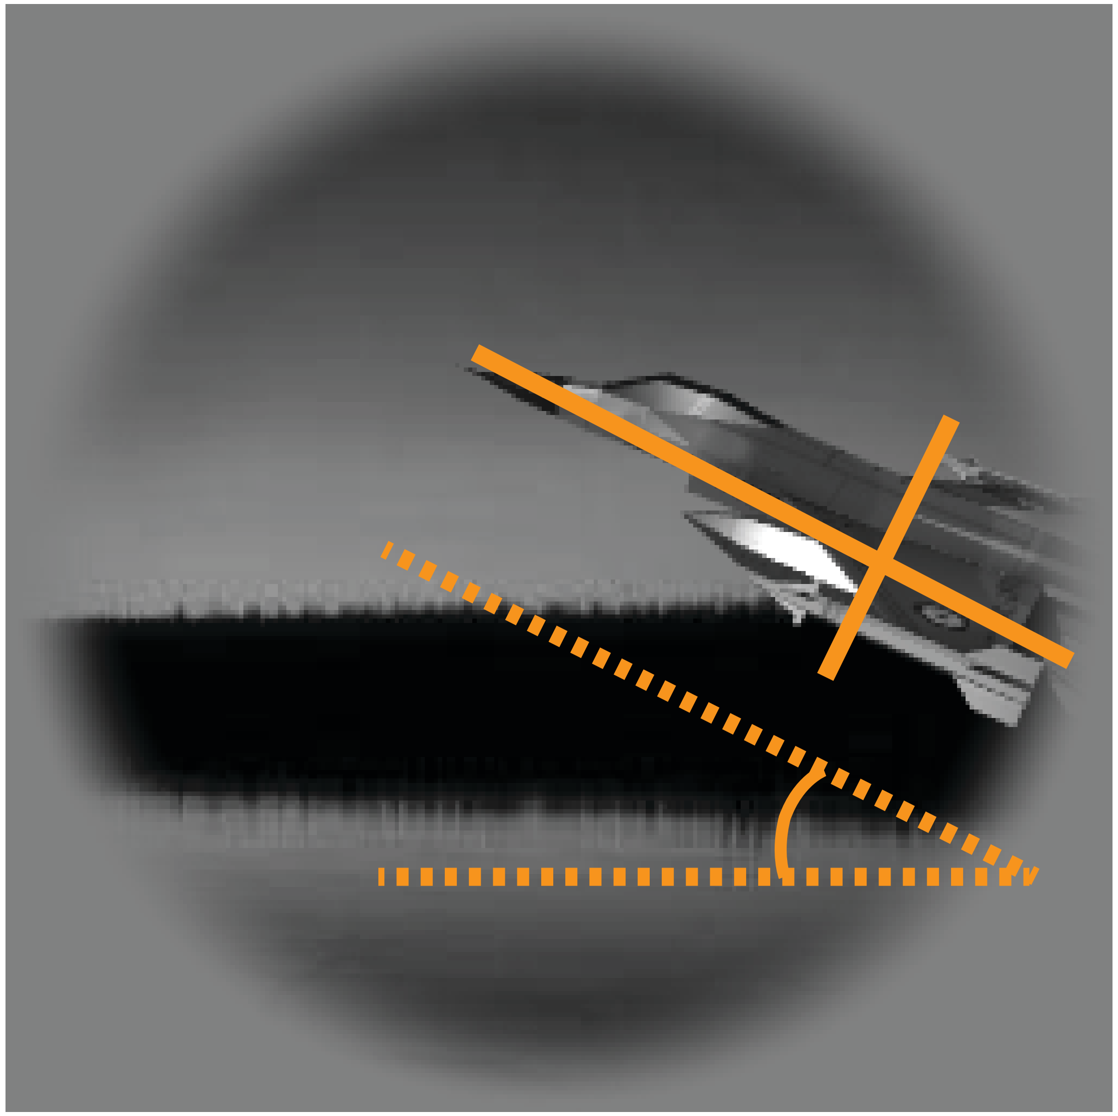

Aspect Ratio and Angle

Beyond categorization

We can quickly assess the scene as a whole.

Category

Identity

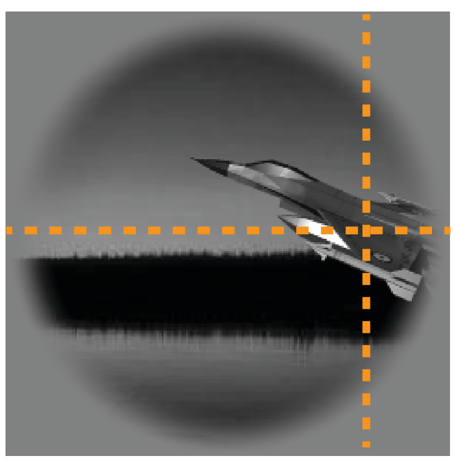

X and Y Axis Position

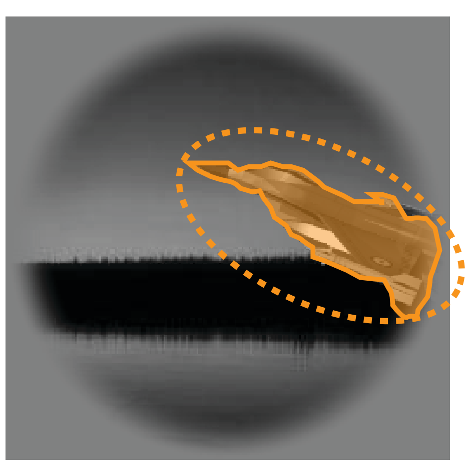

Perimeter

2-D Retinal Area

3-D Object Scale

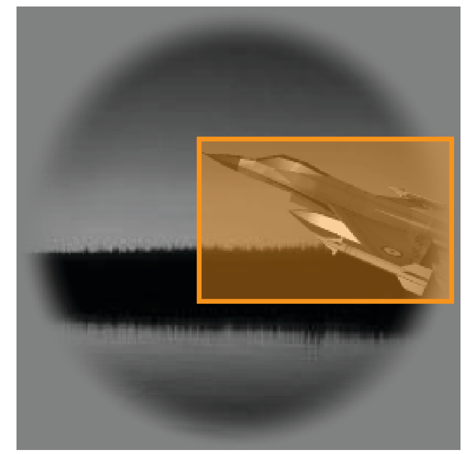

Bounding Box

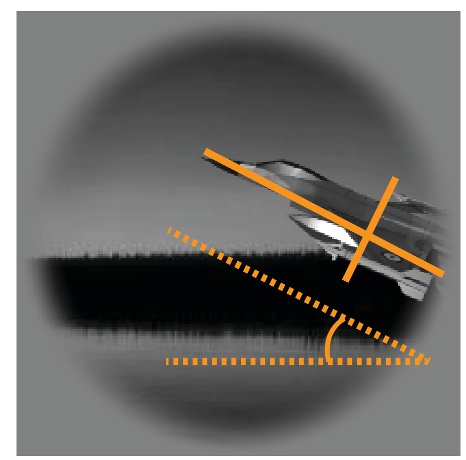

Aspect Ratio

Major Axis Length

Major Axis Angle

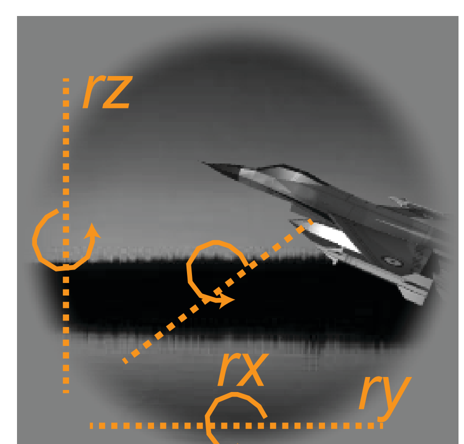

Pose in each axis

Where and how are all these properties coded neurally?

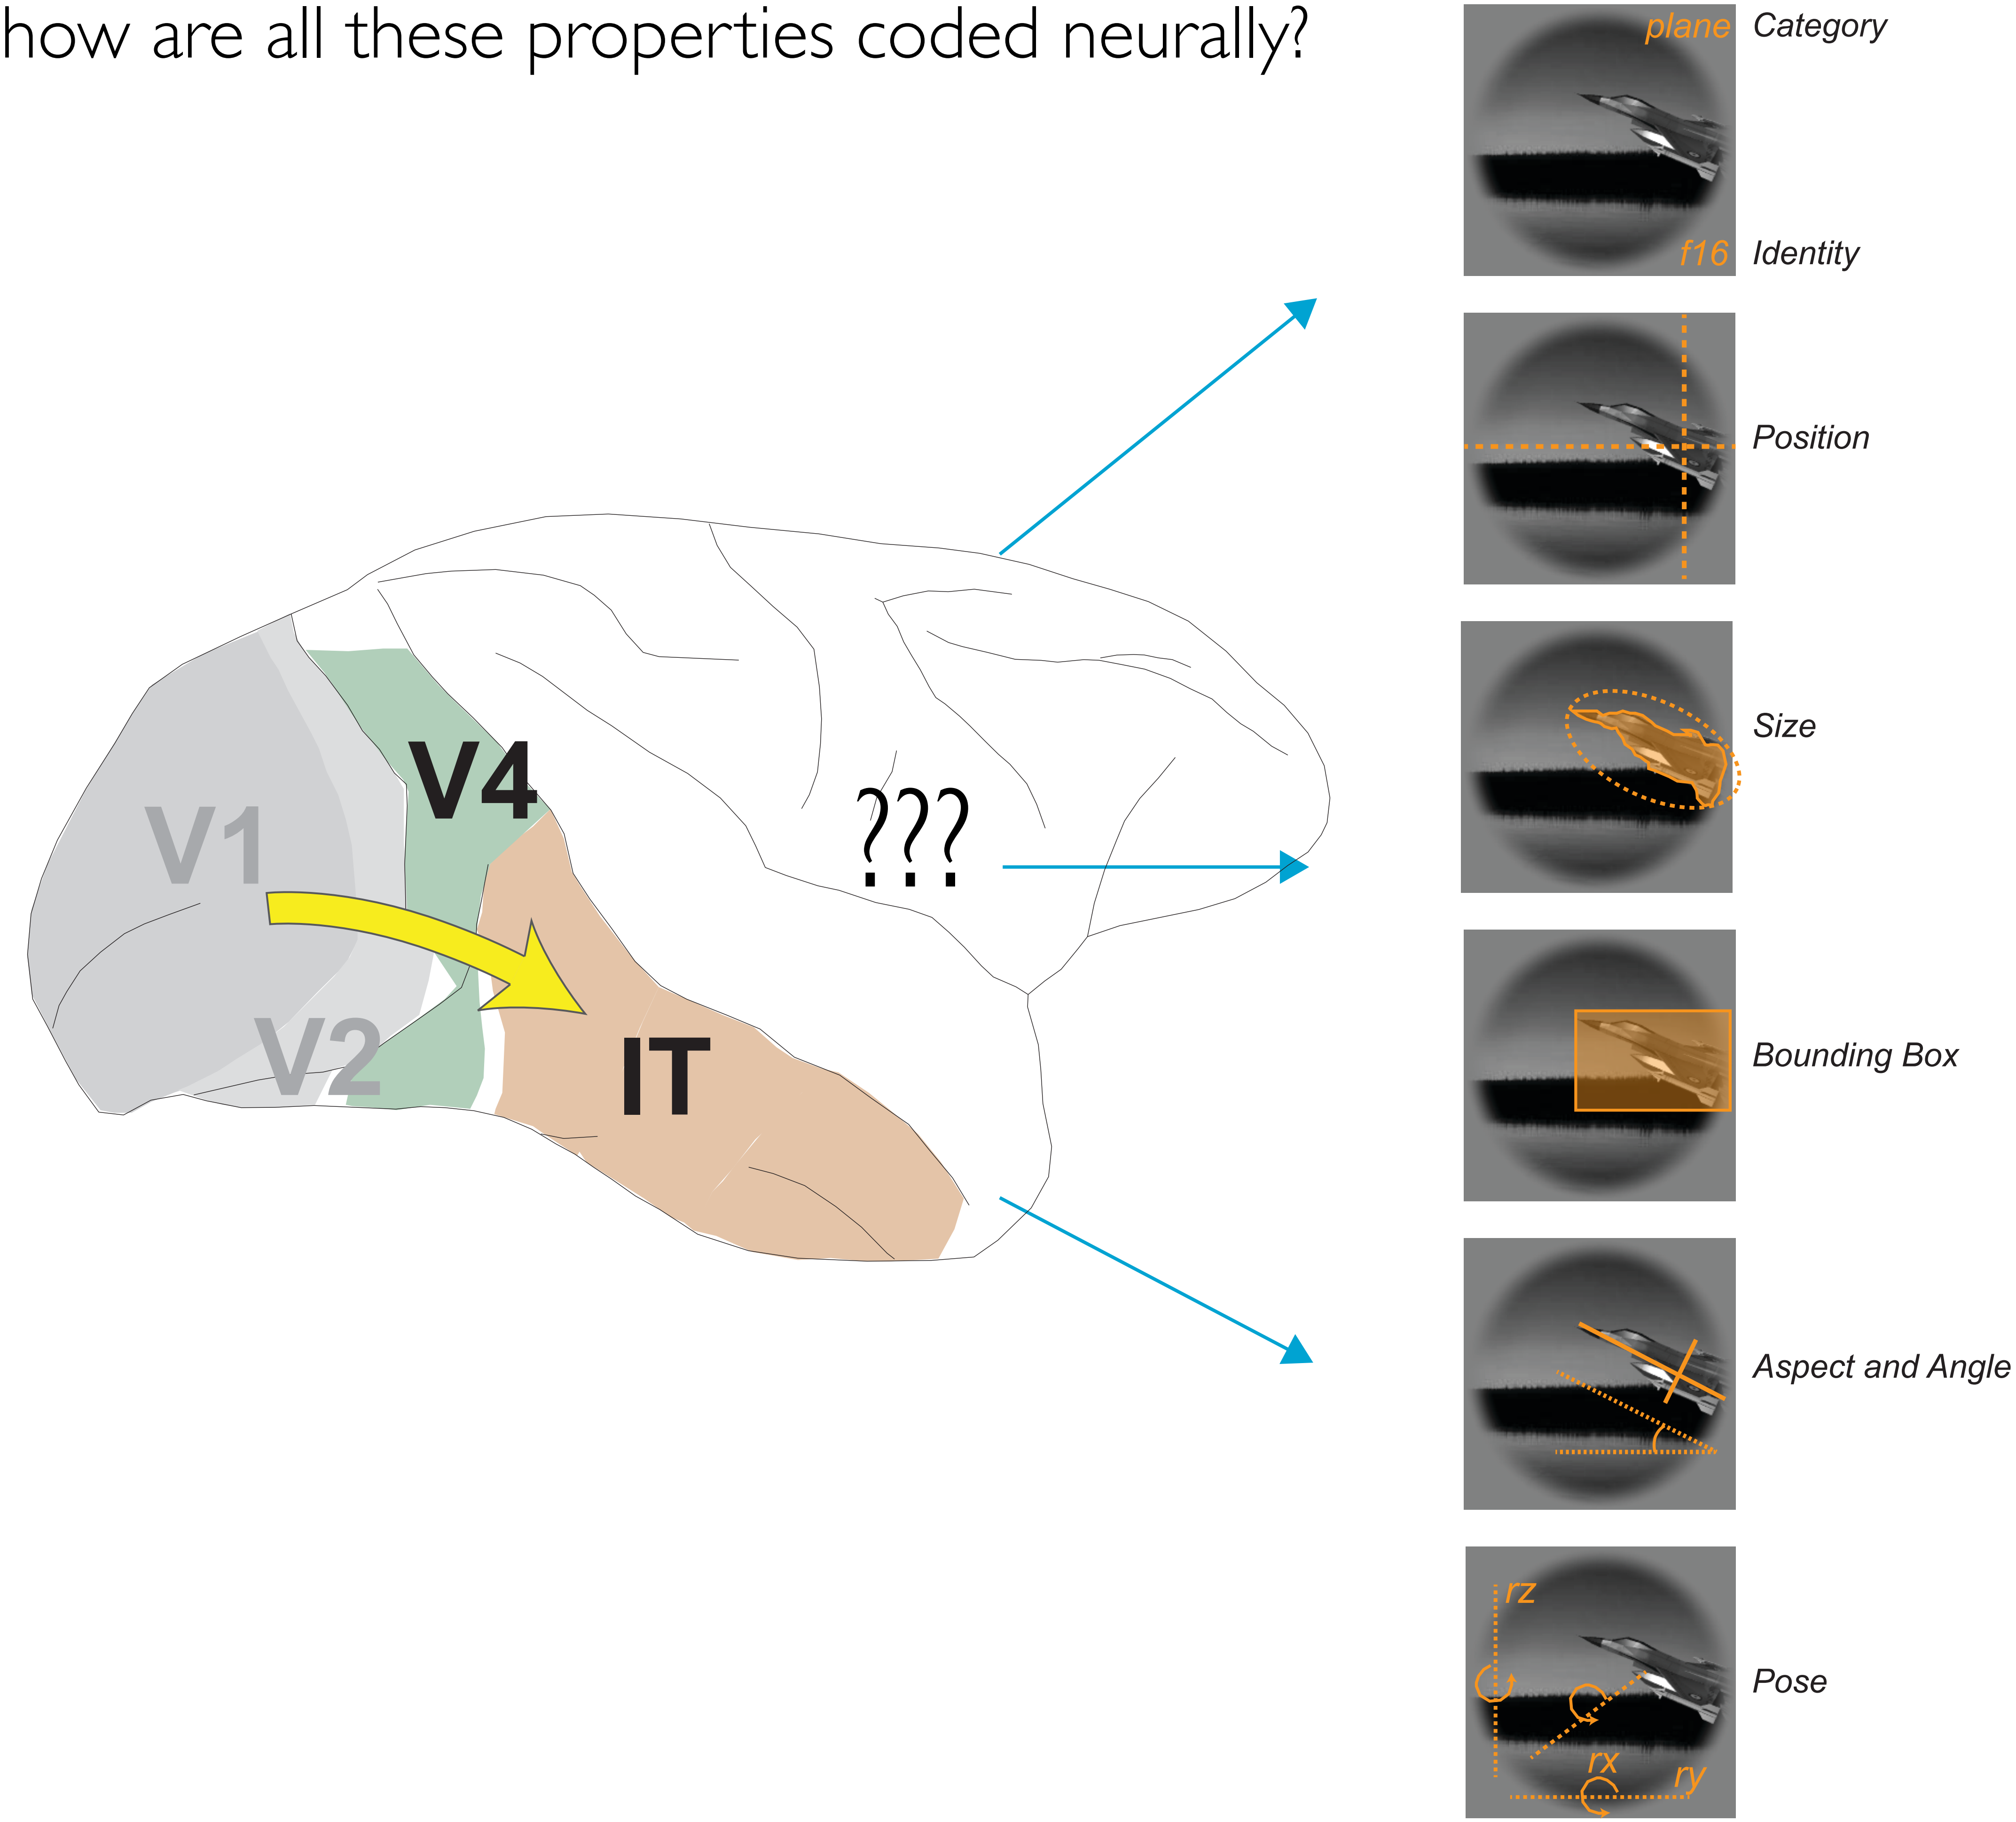

Beyond categorization

“Standard word model” predicts: **not at the top of the ventral stream**.

Aggregation over identity-preserving transformations, e.g. translation.

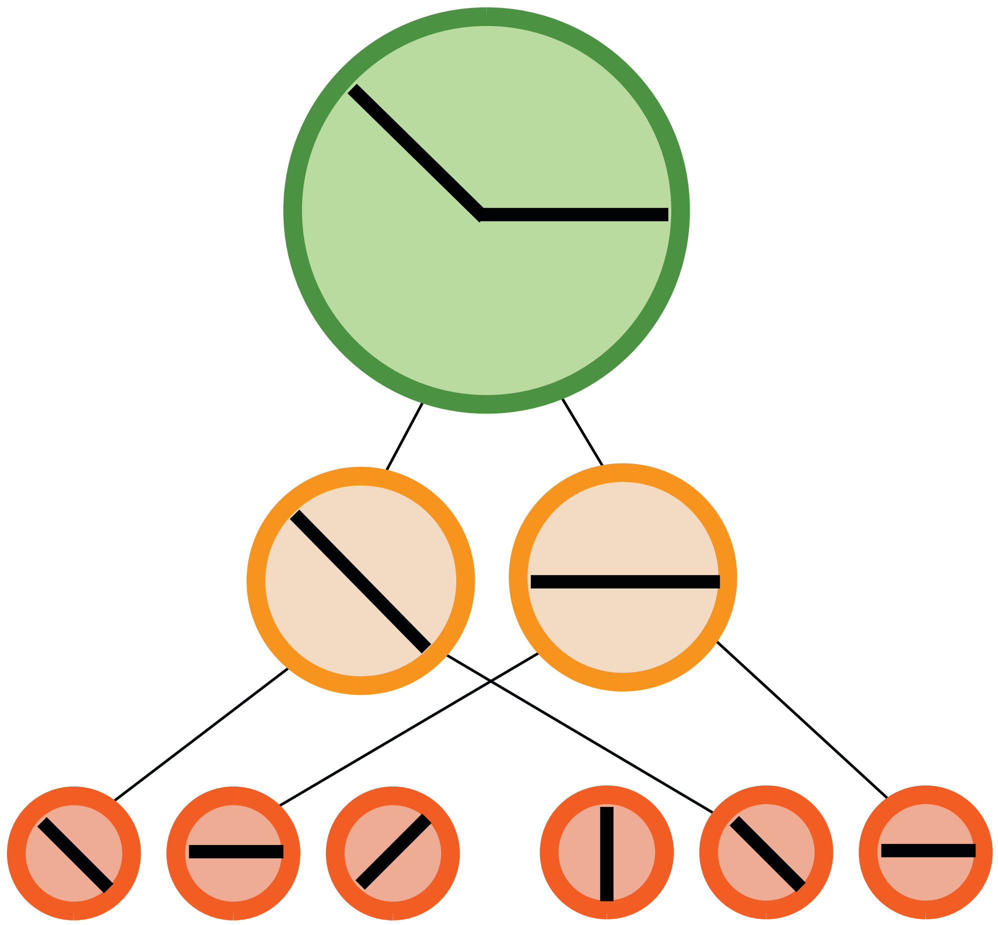

Beyond categorization

"Standard word model" predicts: **not at the top of the ventral stream**.

Aggregation over identity-preserving transformations, e.g. translation.

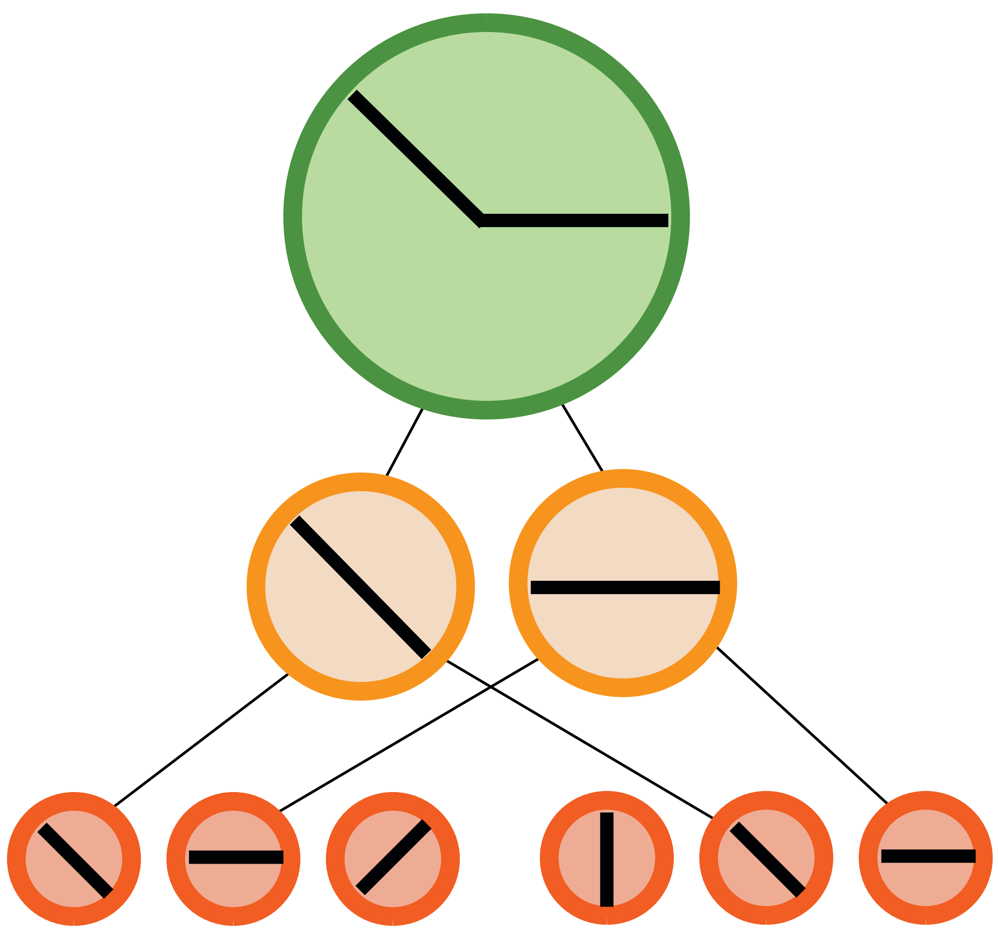

Receptive Field Size ↑

Category Invariance ↑

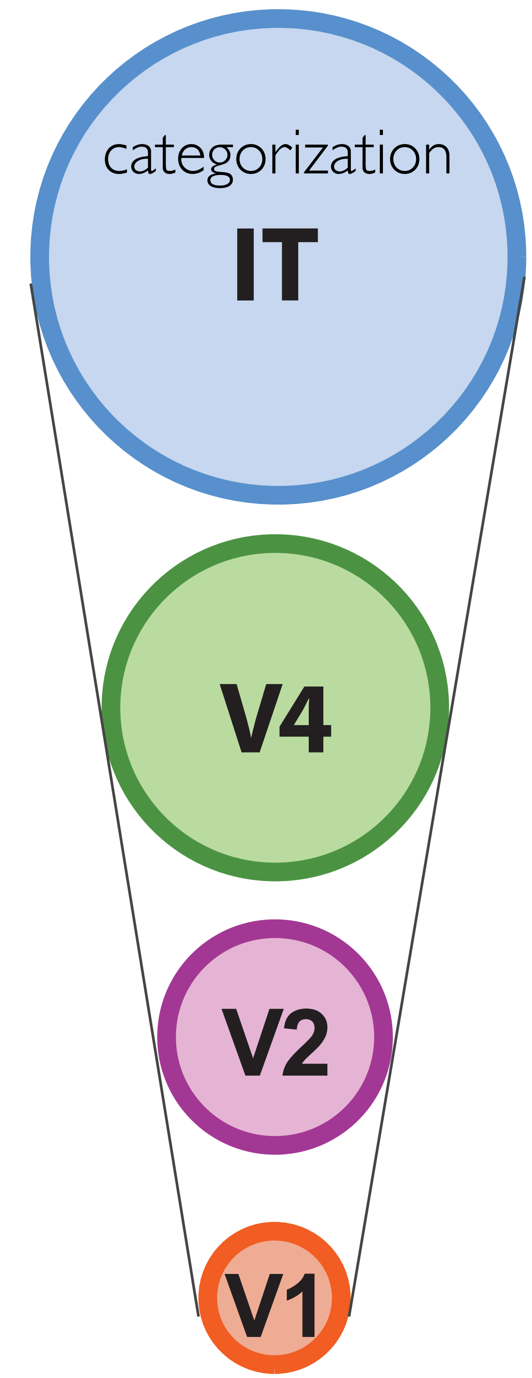

"Standard word model" predicts: **not at the top of the ventral stream**.

Aggregation over identity-preserving transformations, e.g. translation.

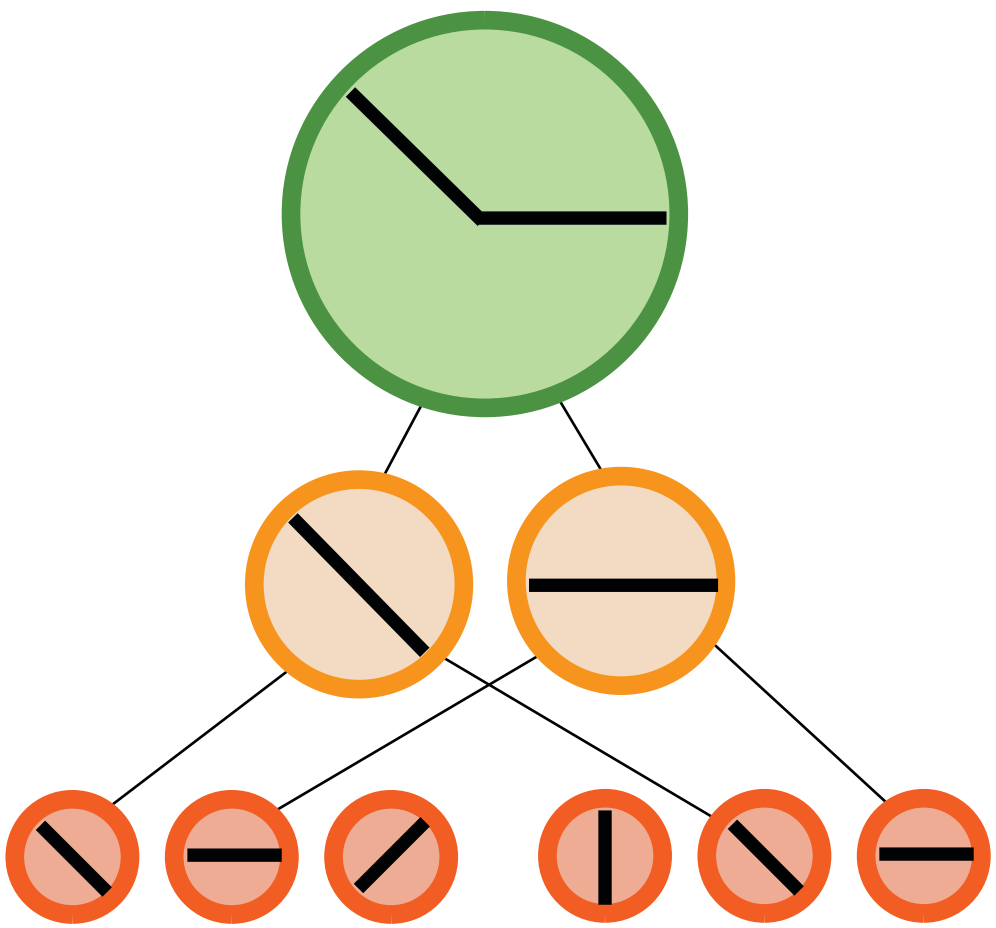

Receptive Field Size ↑

Category Invariance ↑

(e.g.) Position Sensitivity ↓

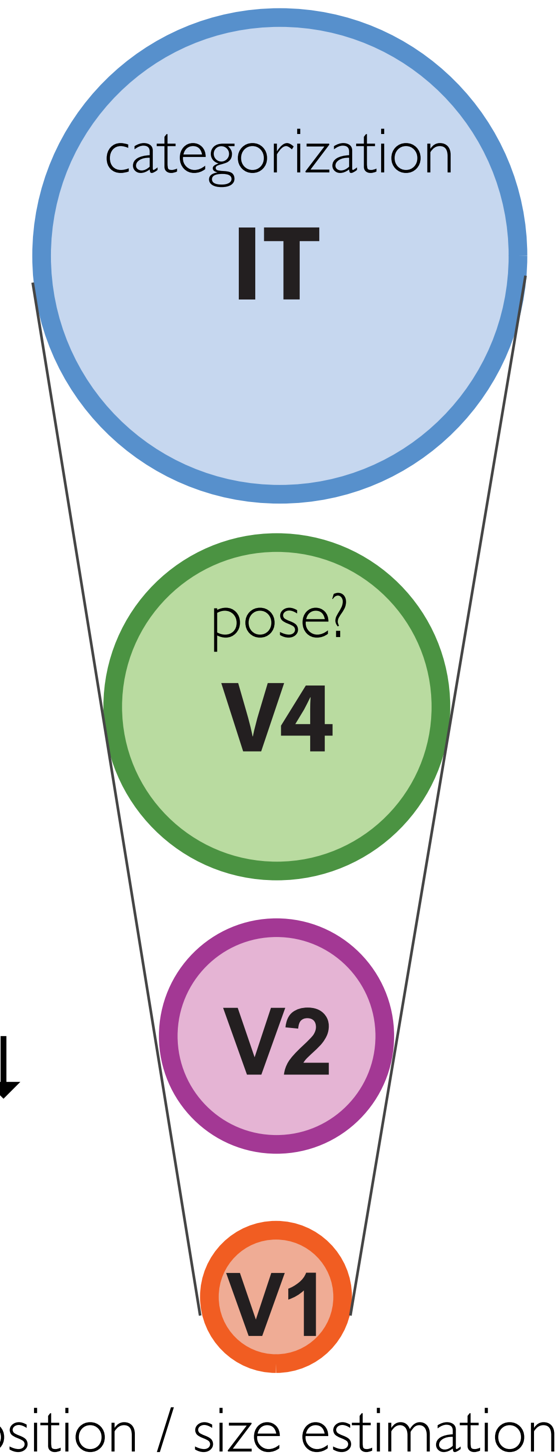

Where and how are all these properties coded neurally?

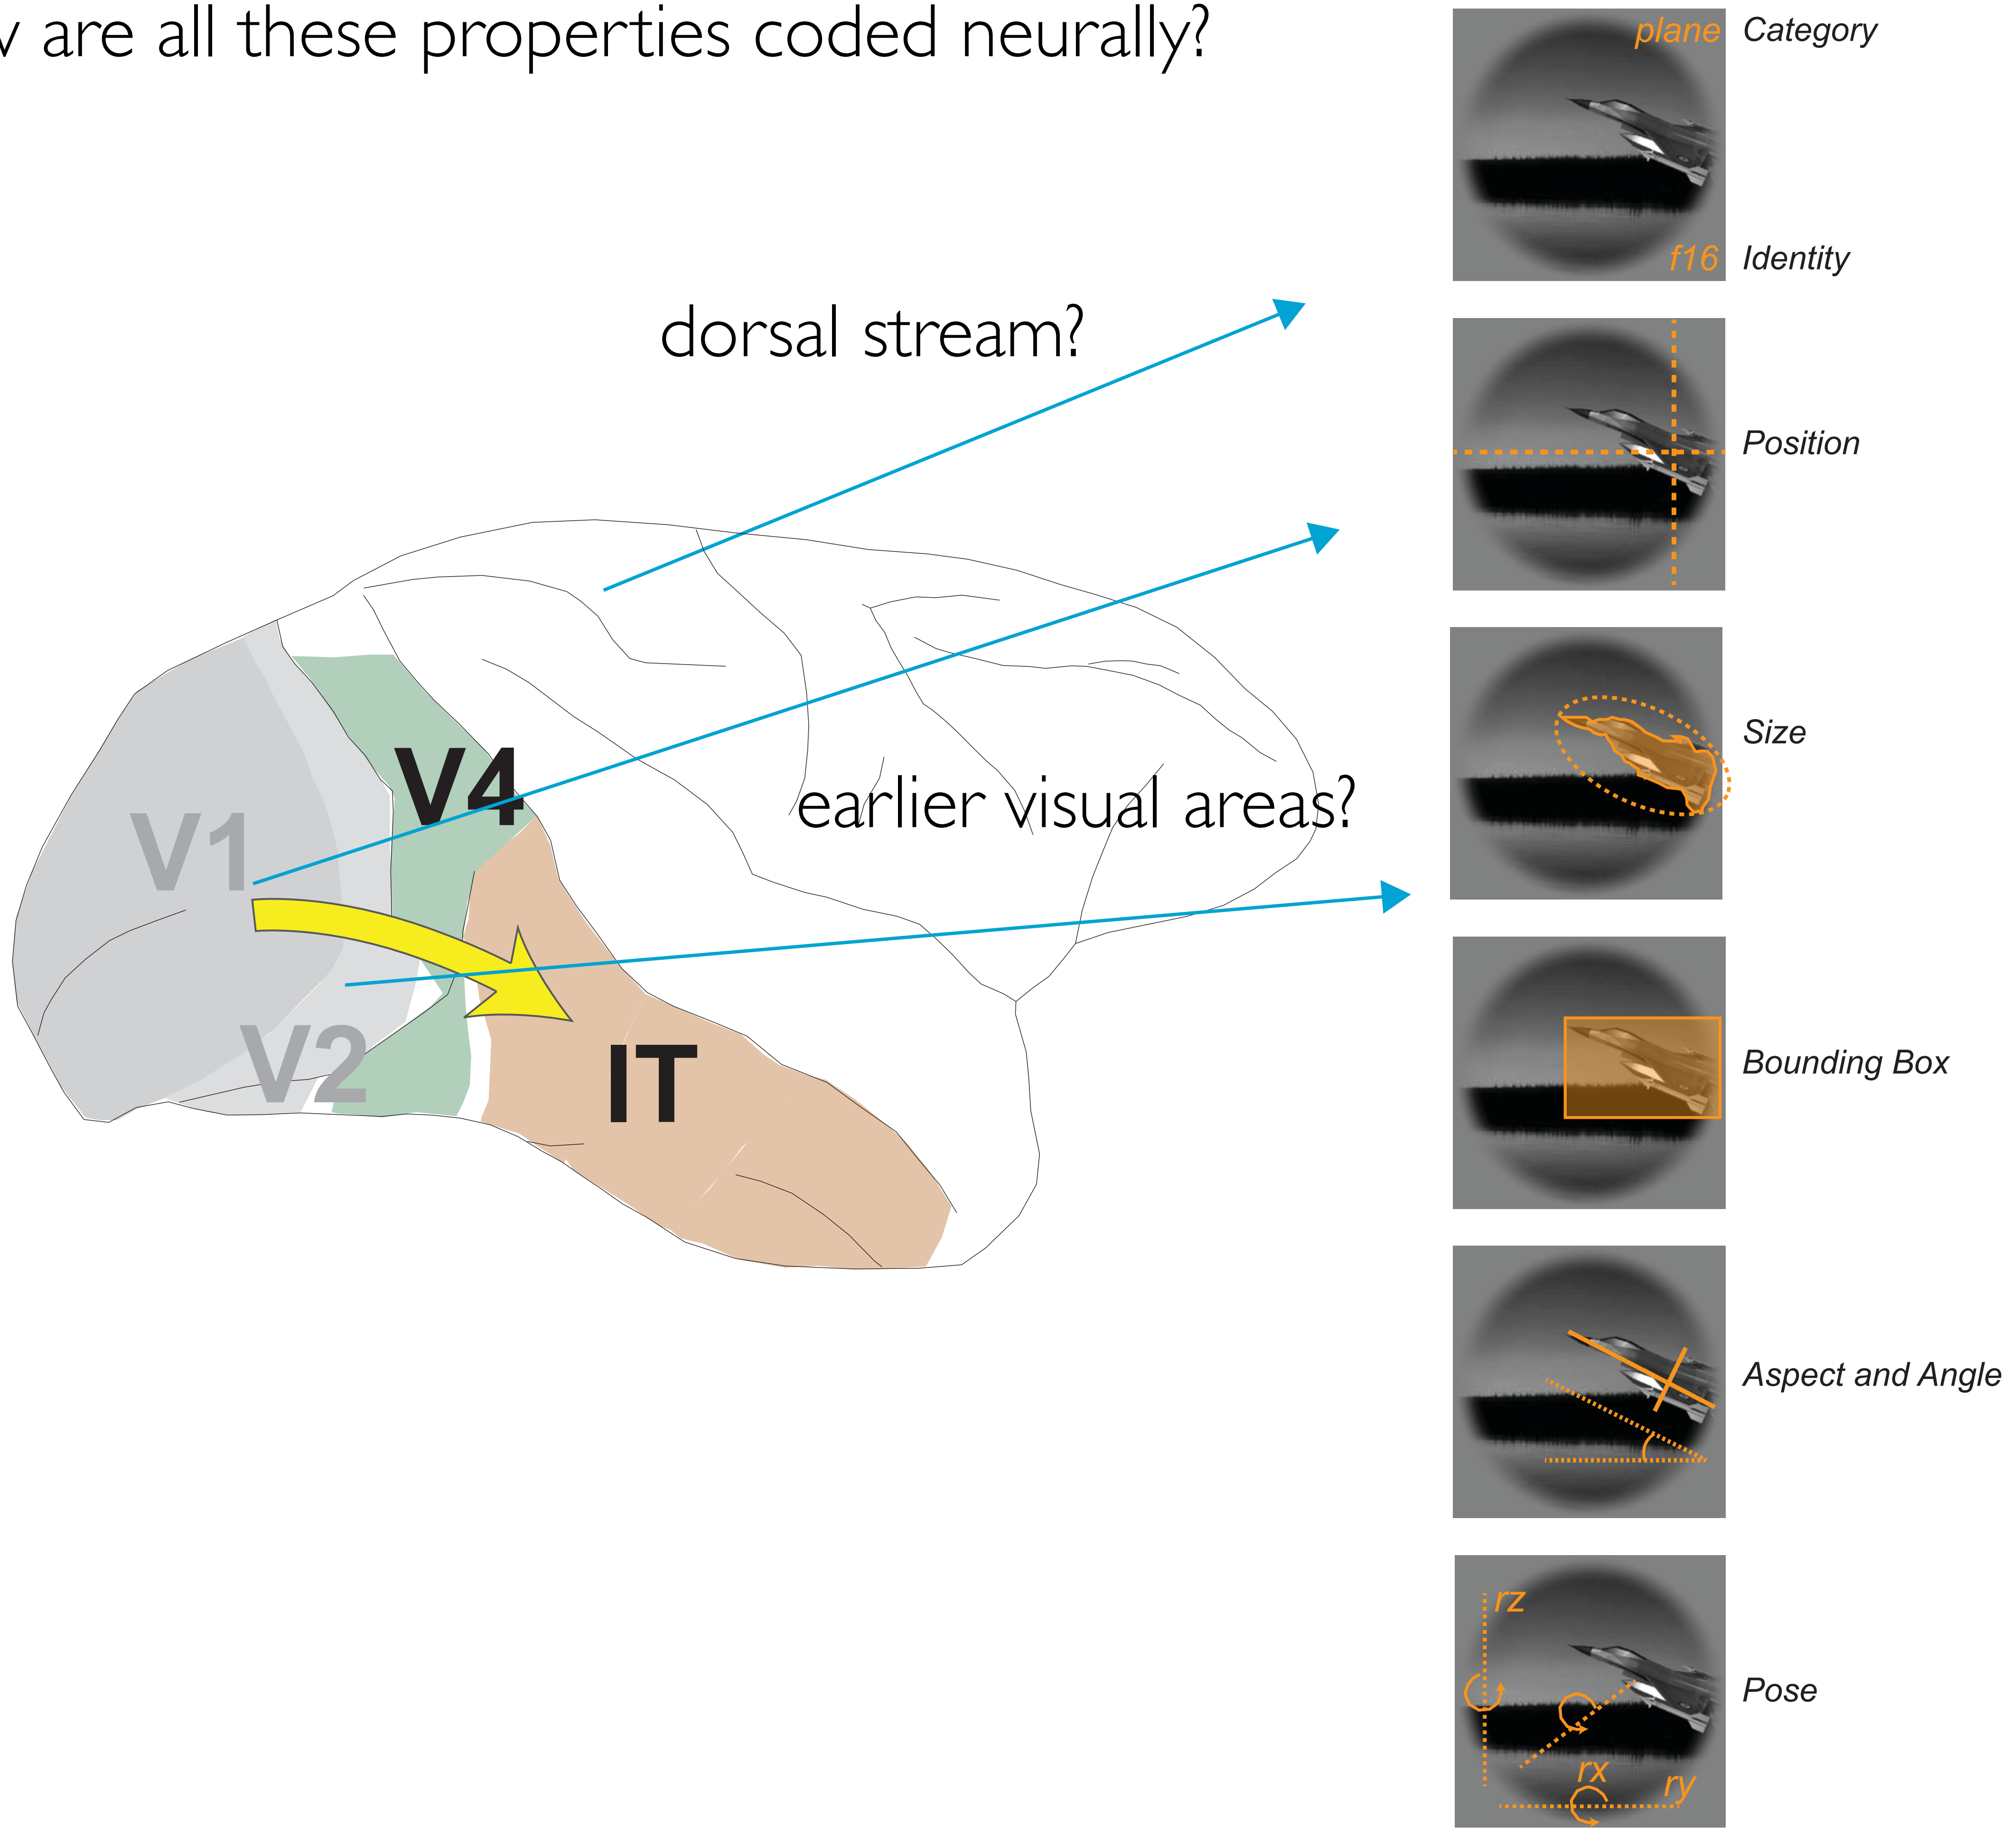

Somewhat newish ideas about IT?

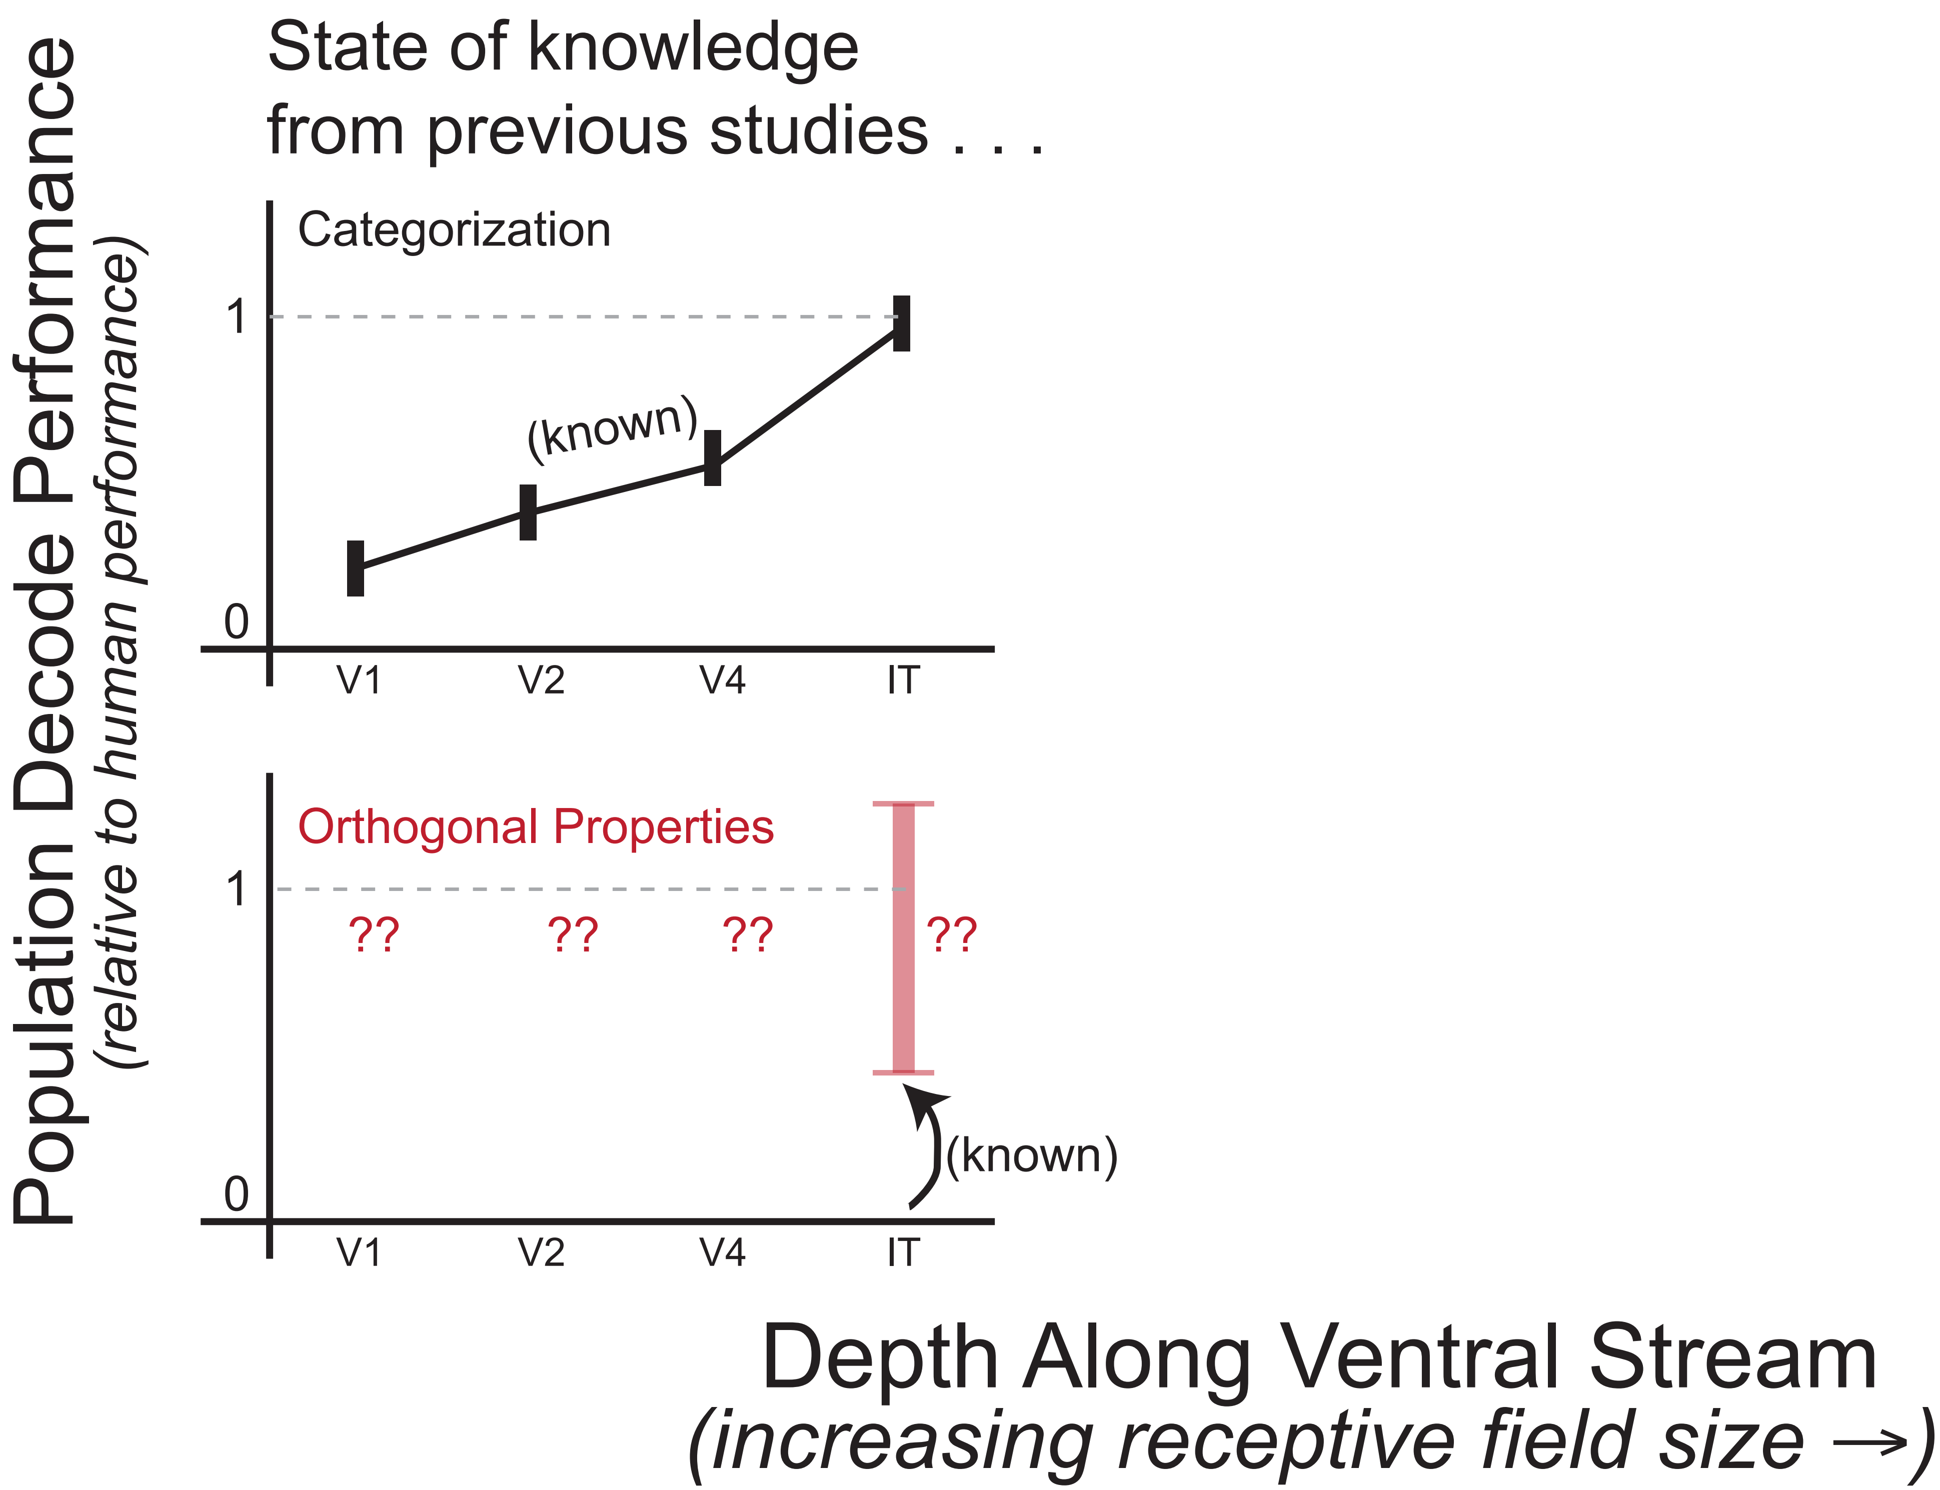

Somewhat newish ideas about IT?

State of knowledge from previous studies . . .

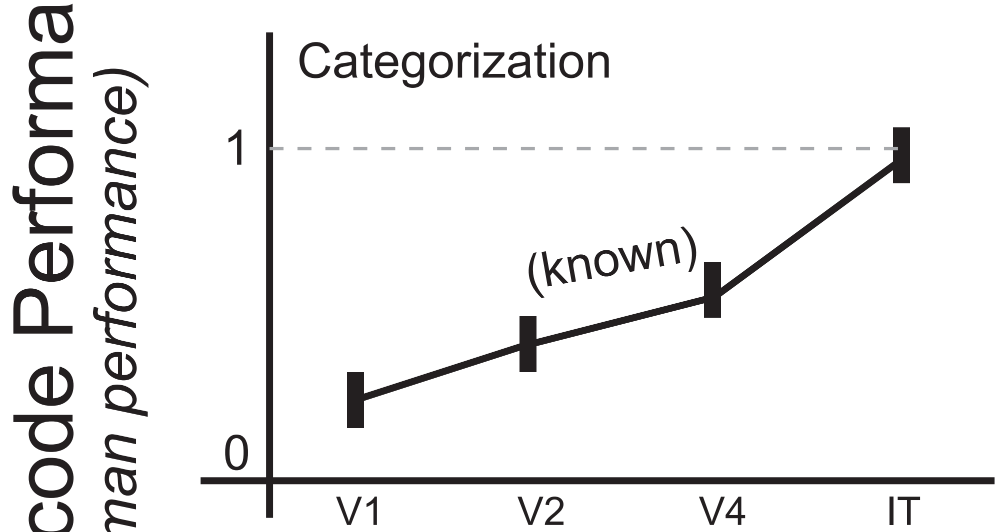

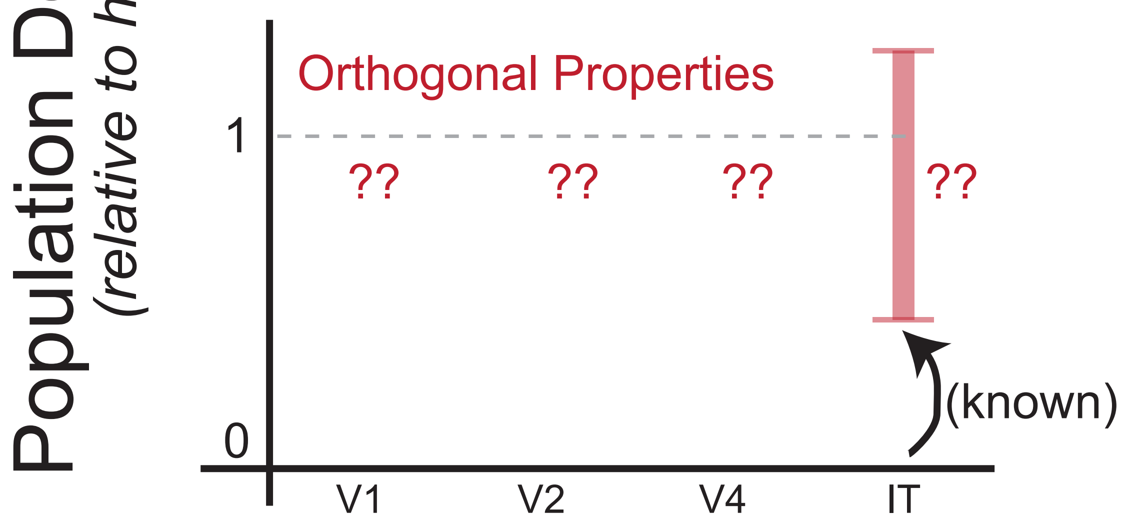

Multiple hypotheses consistent with the existing data . . .

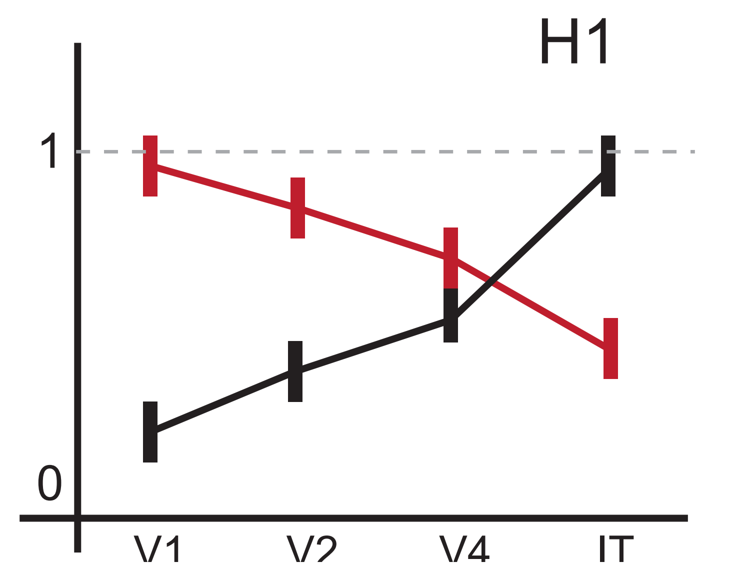

H1: Tolerance / sensitivity tradeoff?

Depth Along Ventral Stream

(increasing receptive field size →)

Somewhat newish ideas about IT?

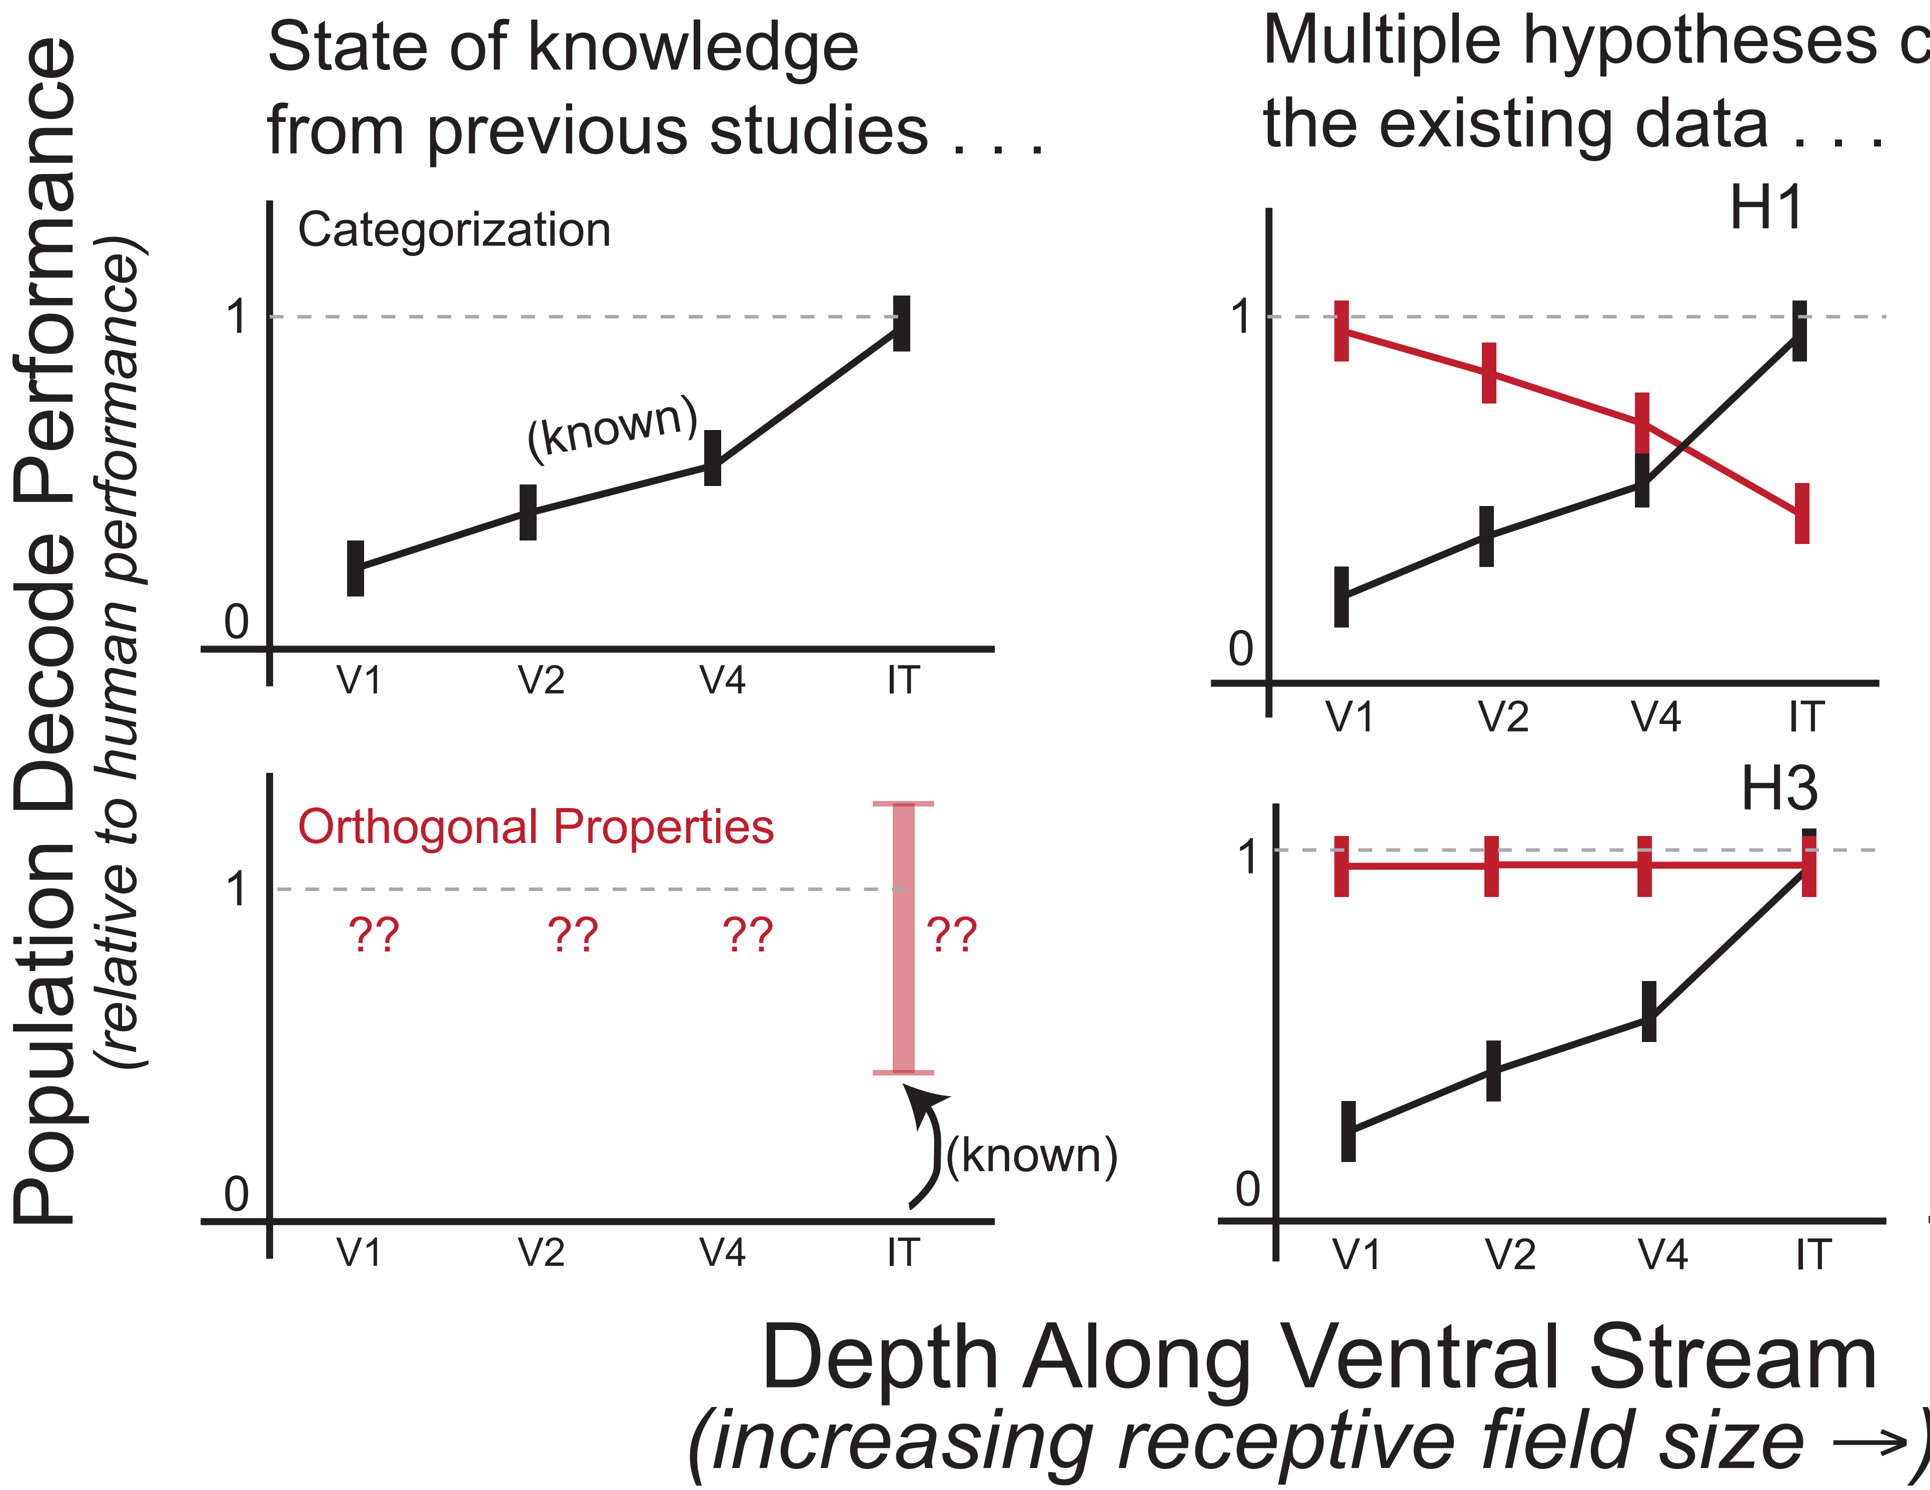

H3: Information preservation?

Beyond categorization

Unexpected observation:

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

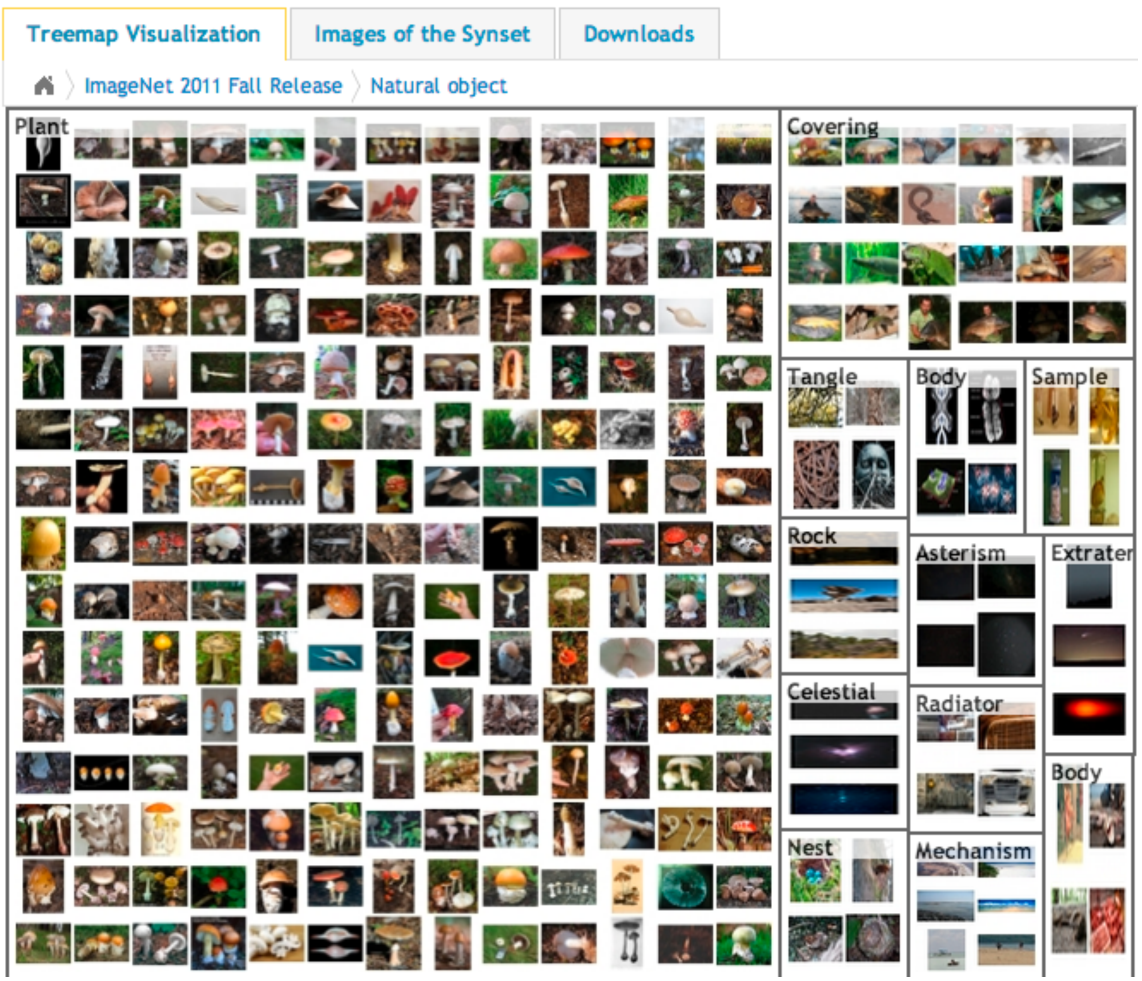

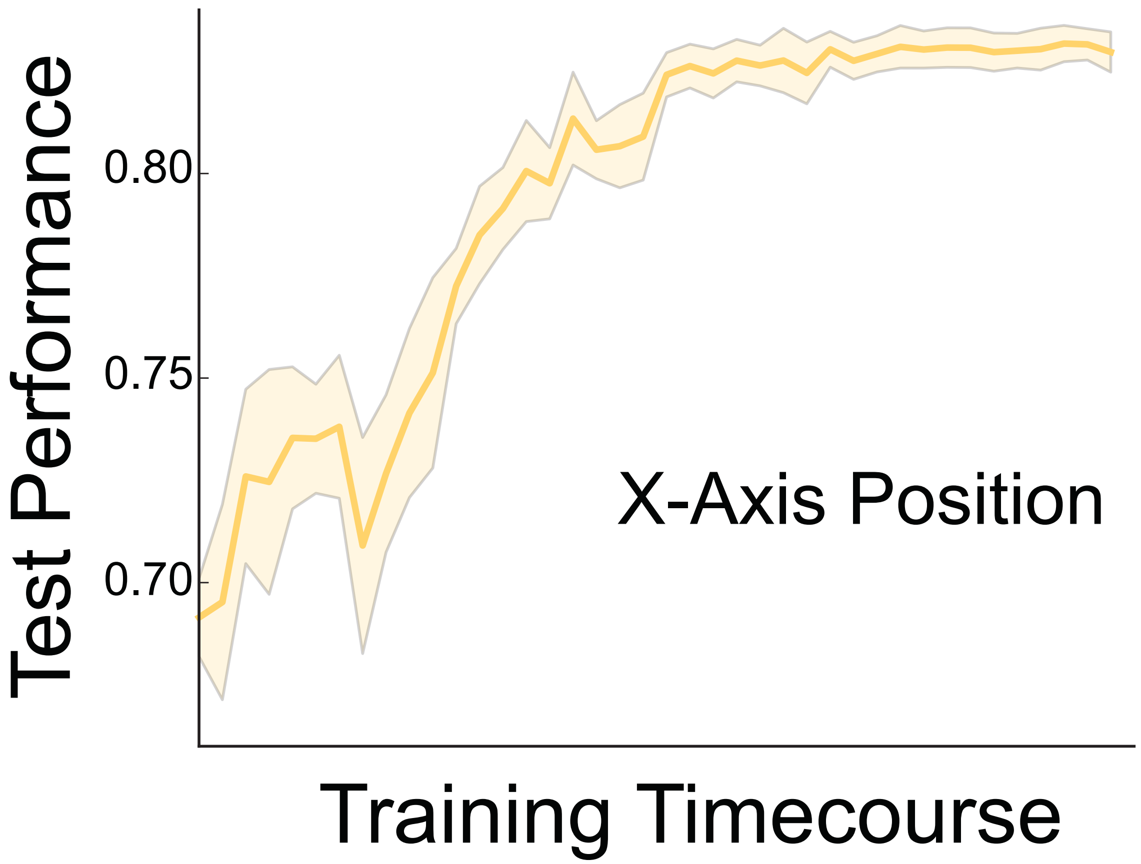

Training on
categorization task

Increased performance on
position estimation task.

even though the goal was to become INVARIANT to position

Beyond categorization

Category optimization → improved performance on non-categorical tasks.

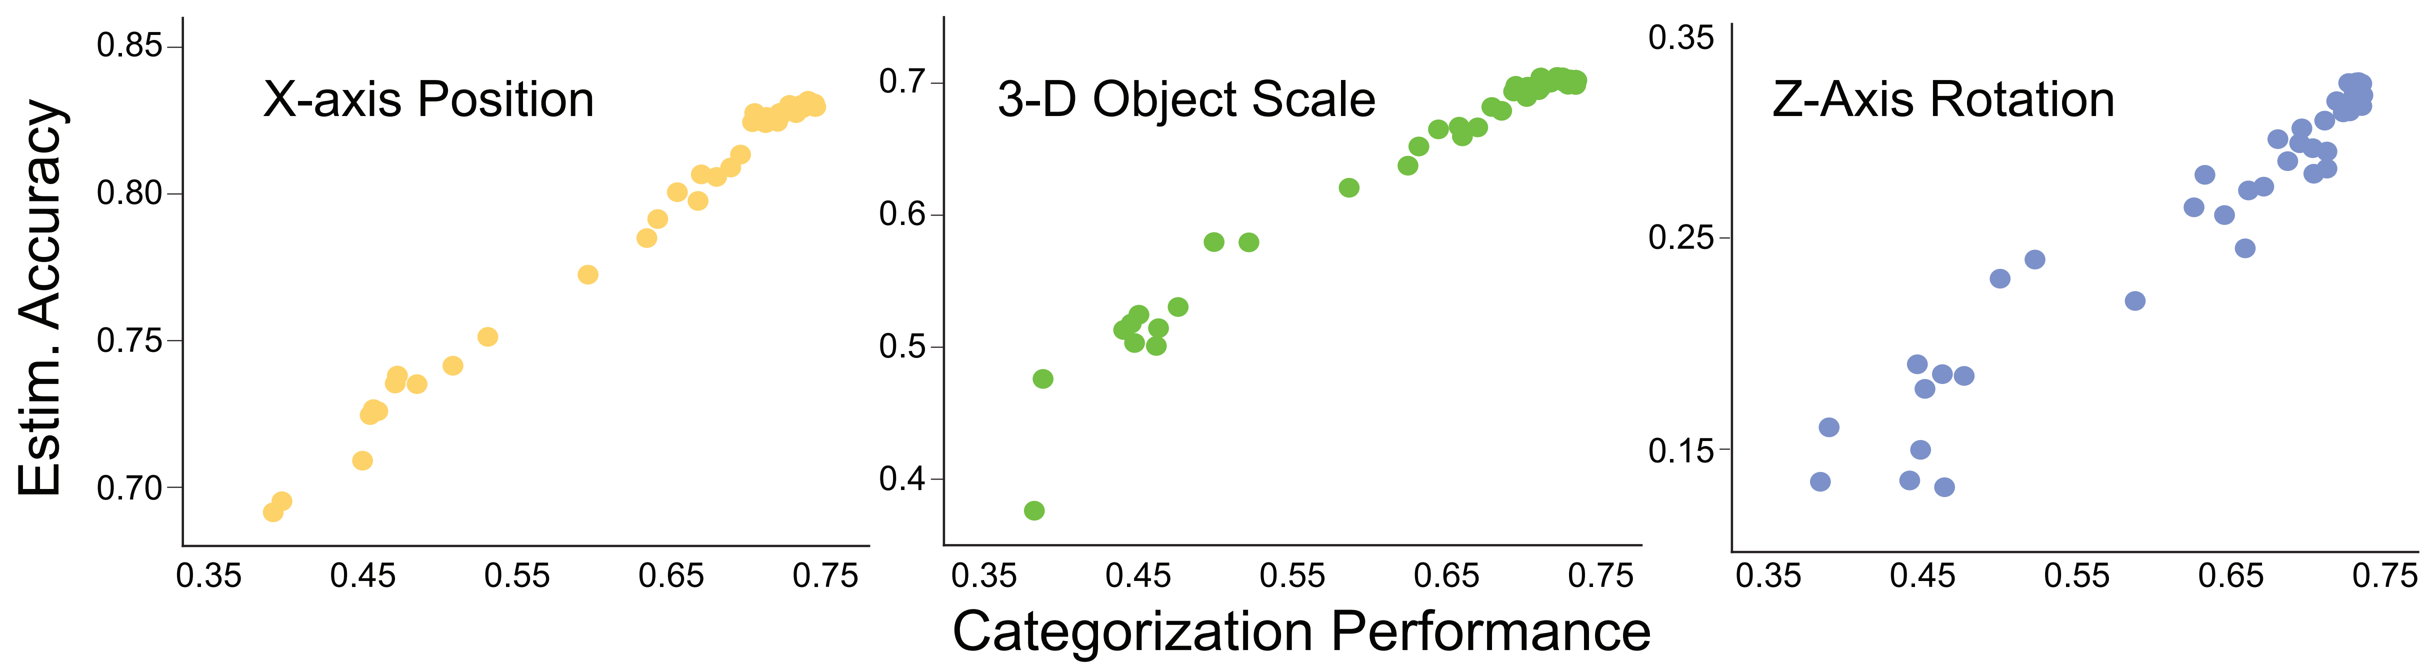

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

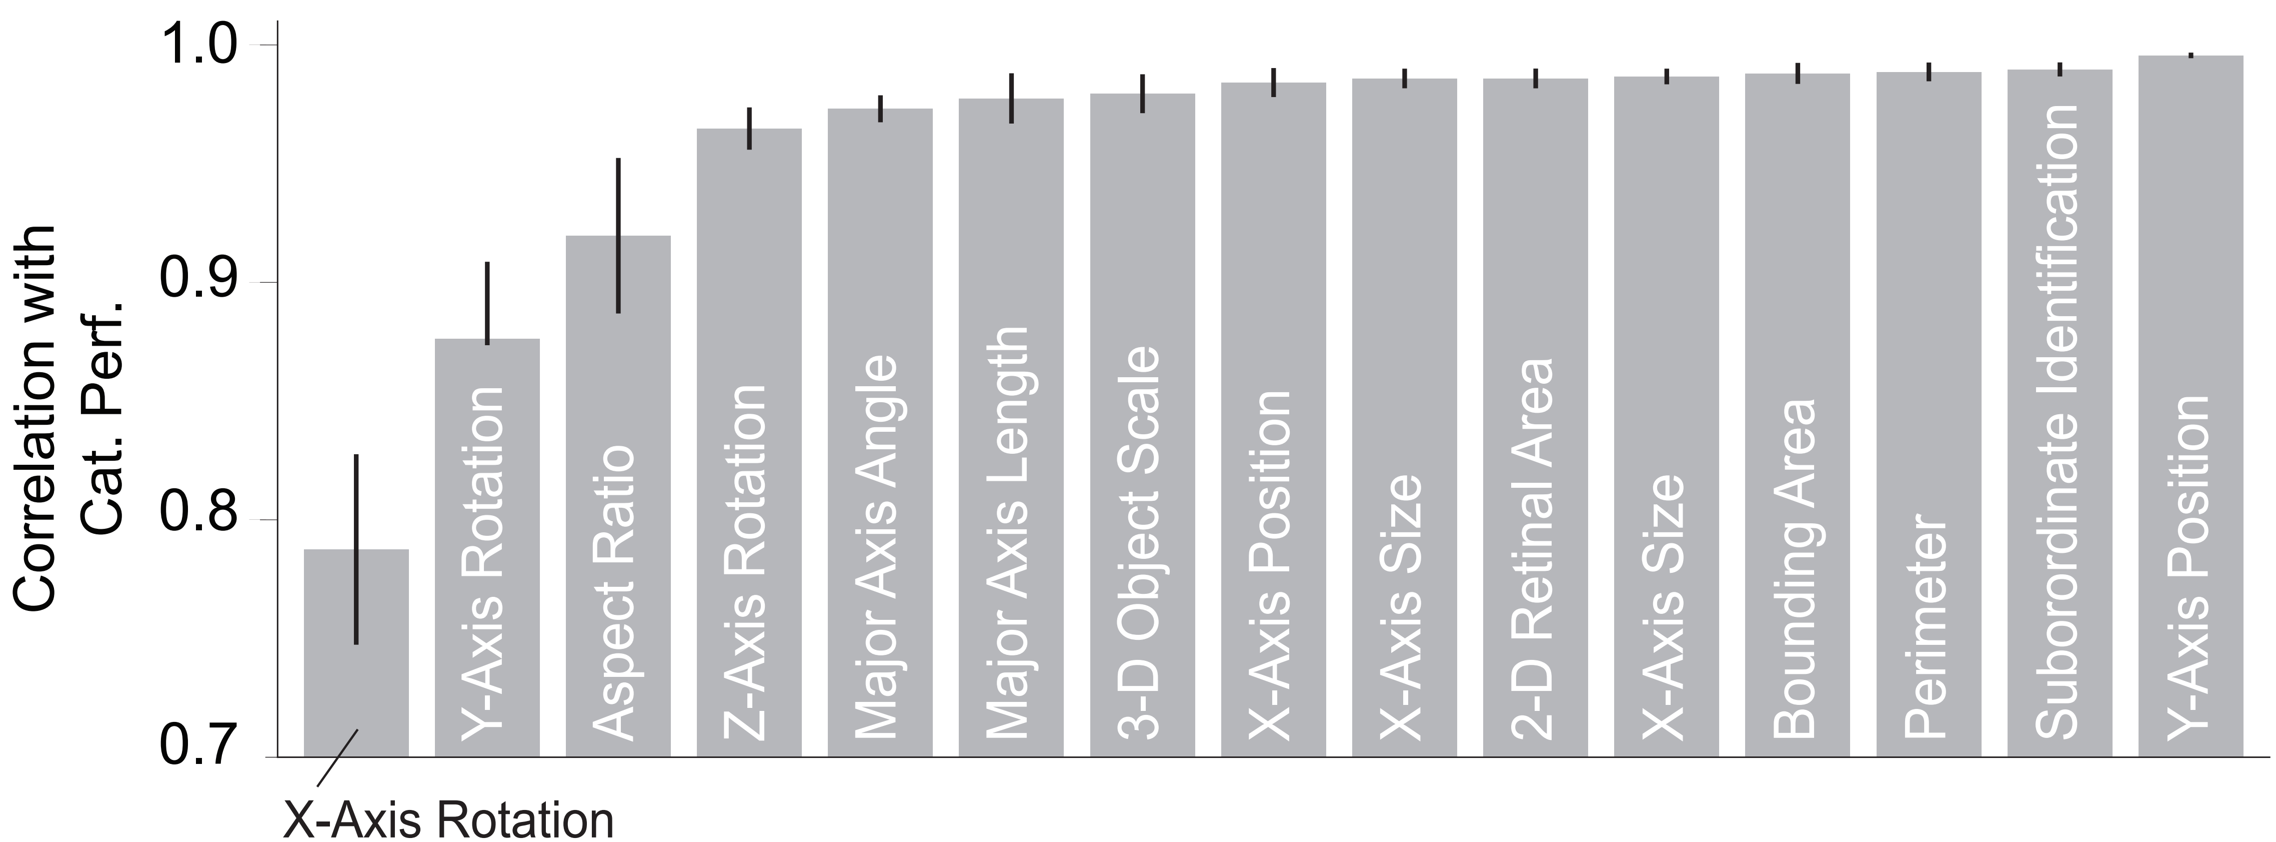

Beyond categorization

Unexpected observation #2:

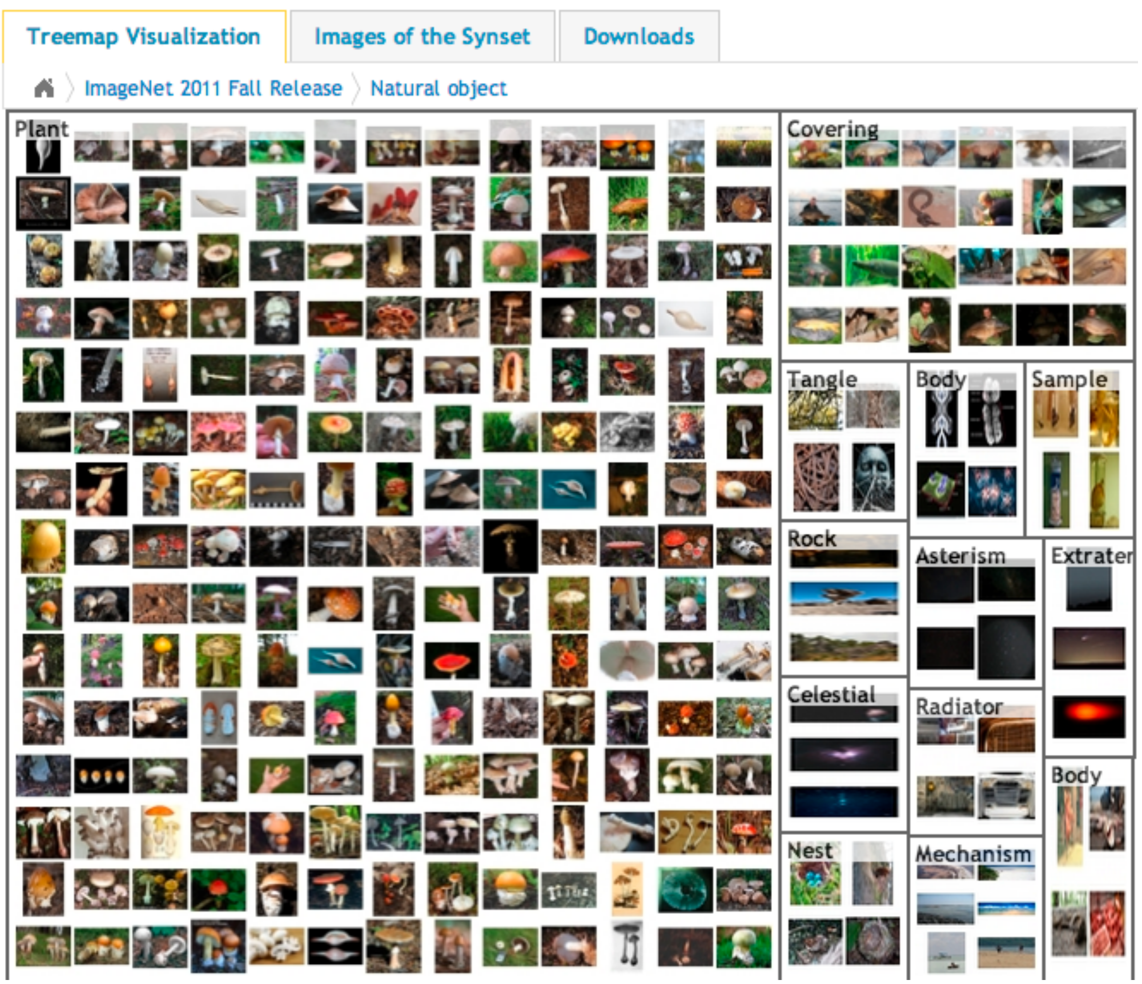

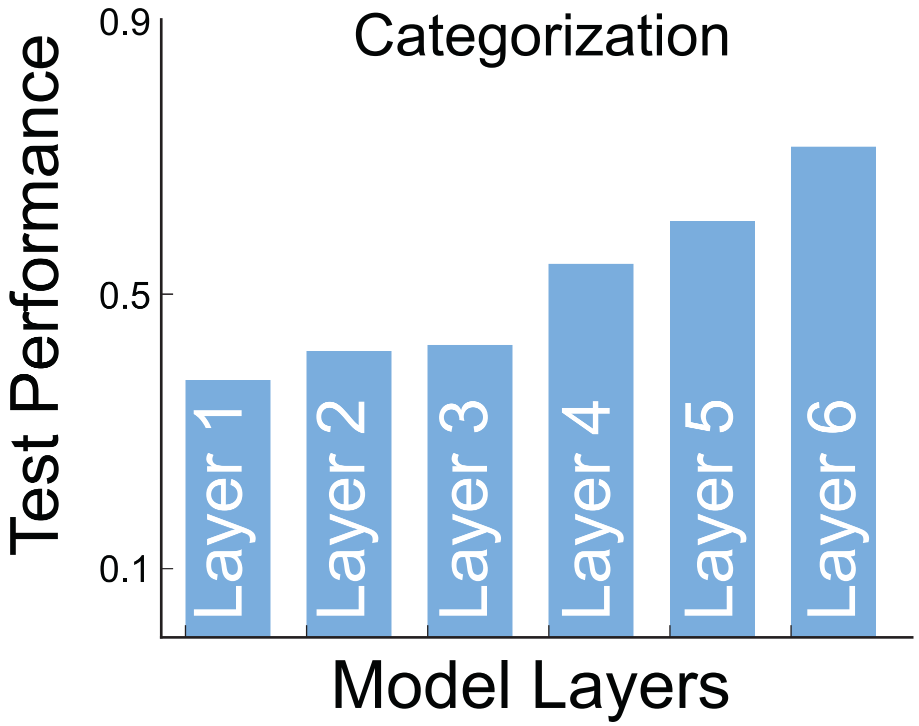

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

Beyond categorization

Unexpected observation #2:

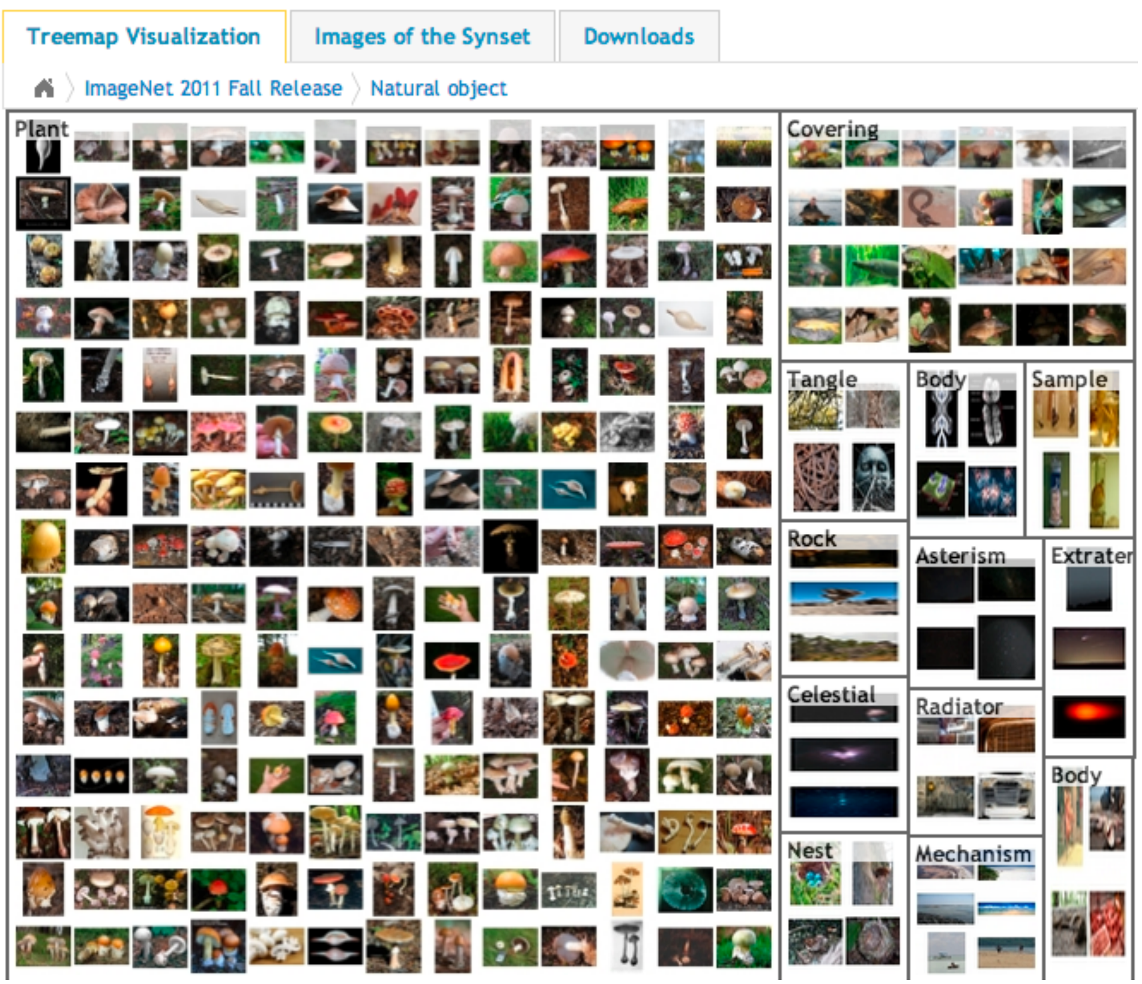

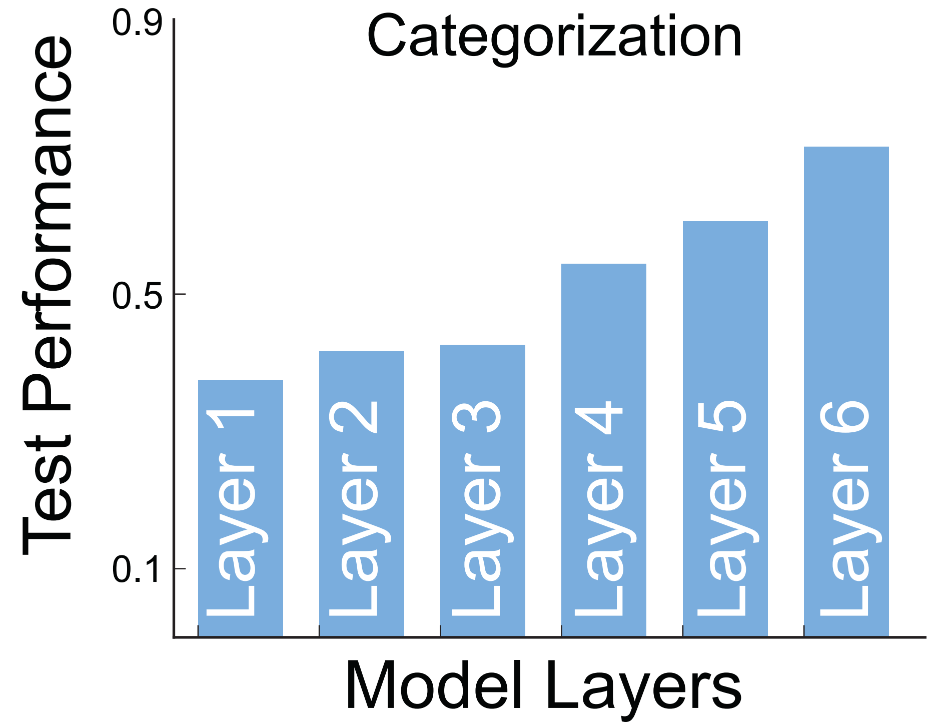

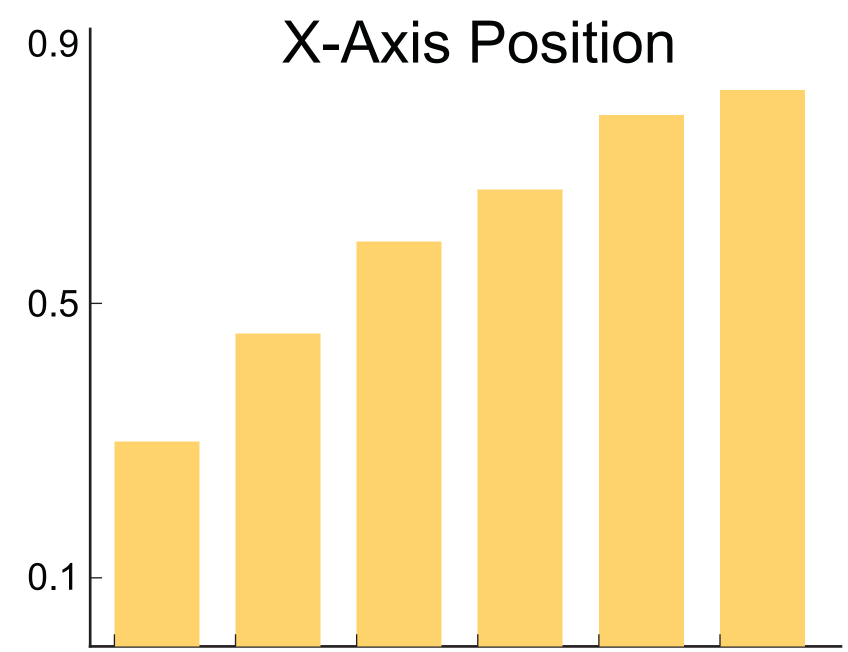

Increased performance on
position estimation task
at each model layer.

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

Beyond categorization

For all tasks of visual interest we could measure in our test dataset:

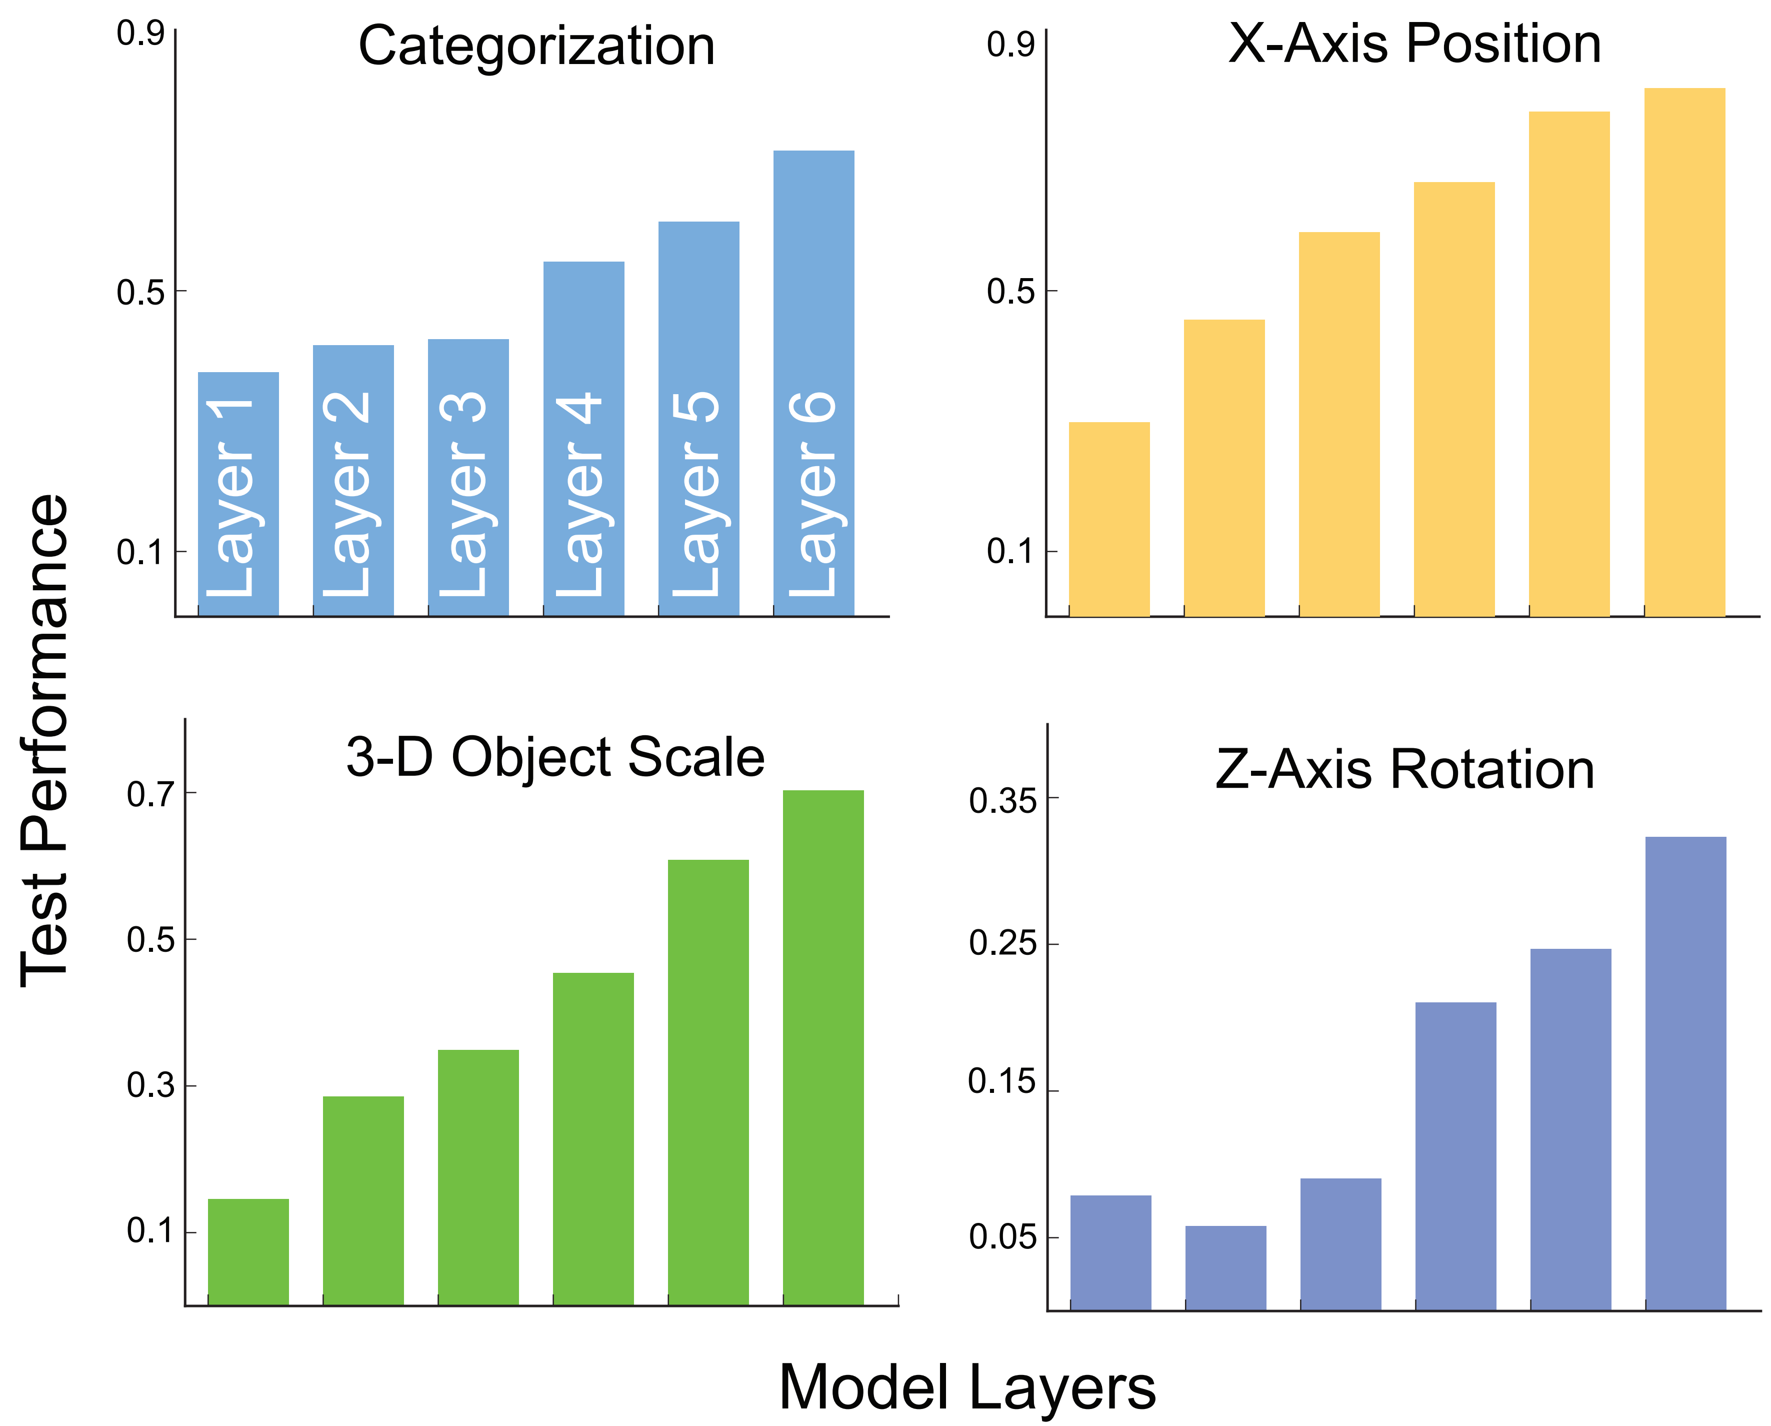

Performance on non-categorical tasks increases at each layer.

Beyond categorization

What do the data say?

Population Decoding

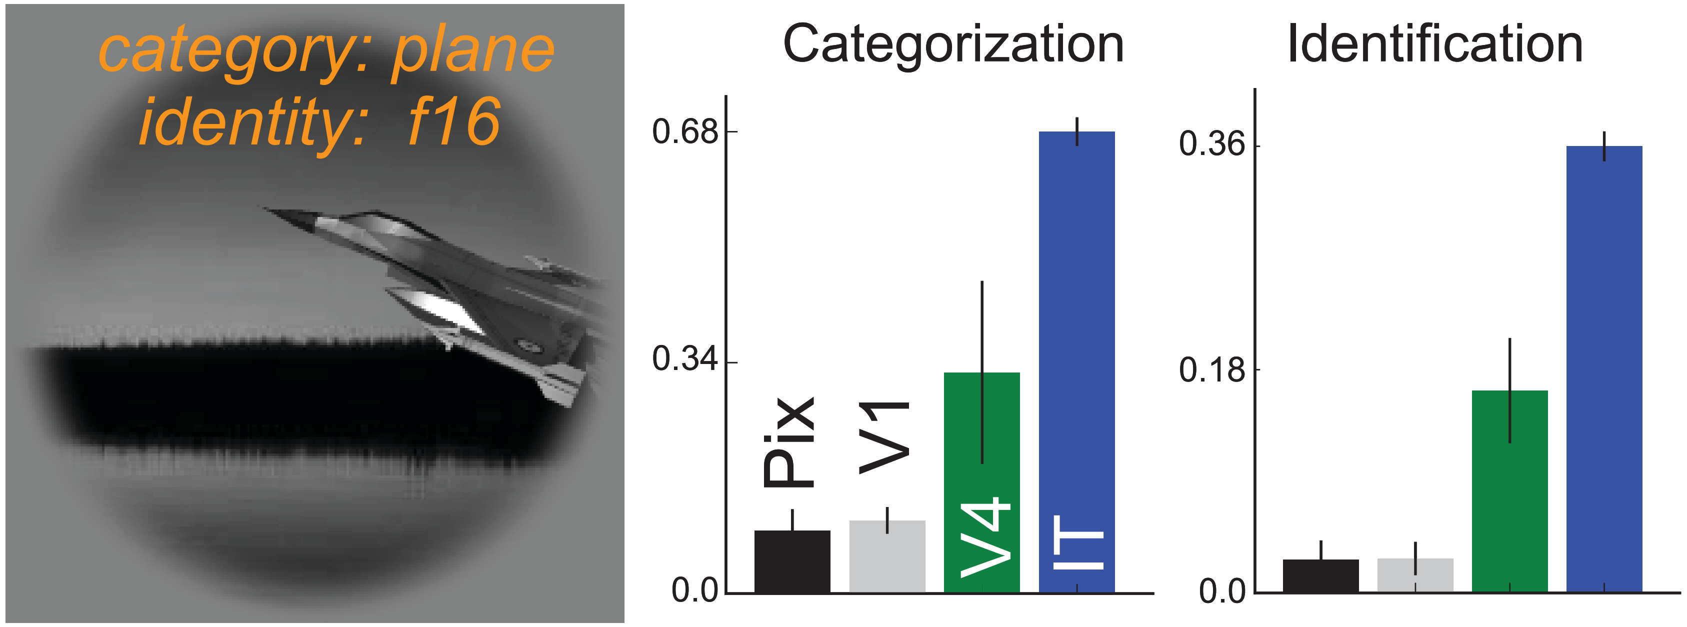

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

IT cortex

V4 cortex

V1-like model

pixel control

Population Decoding

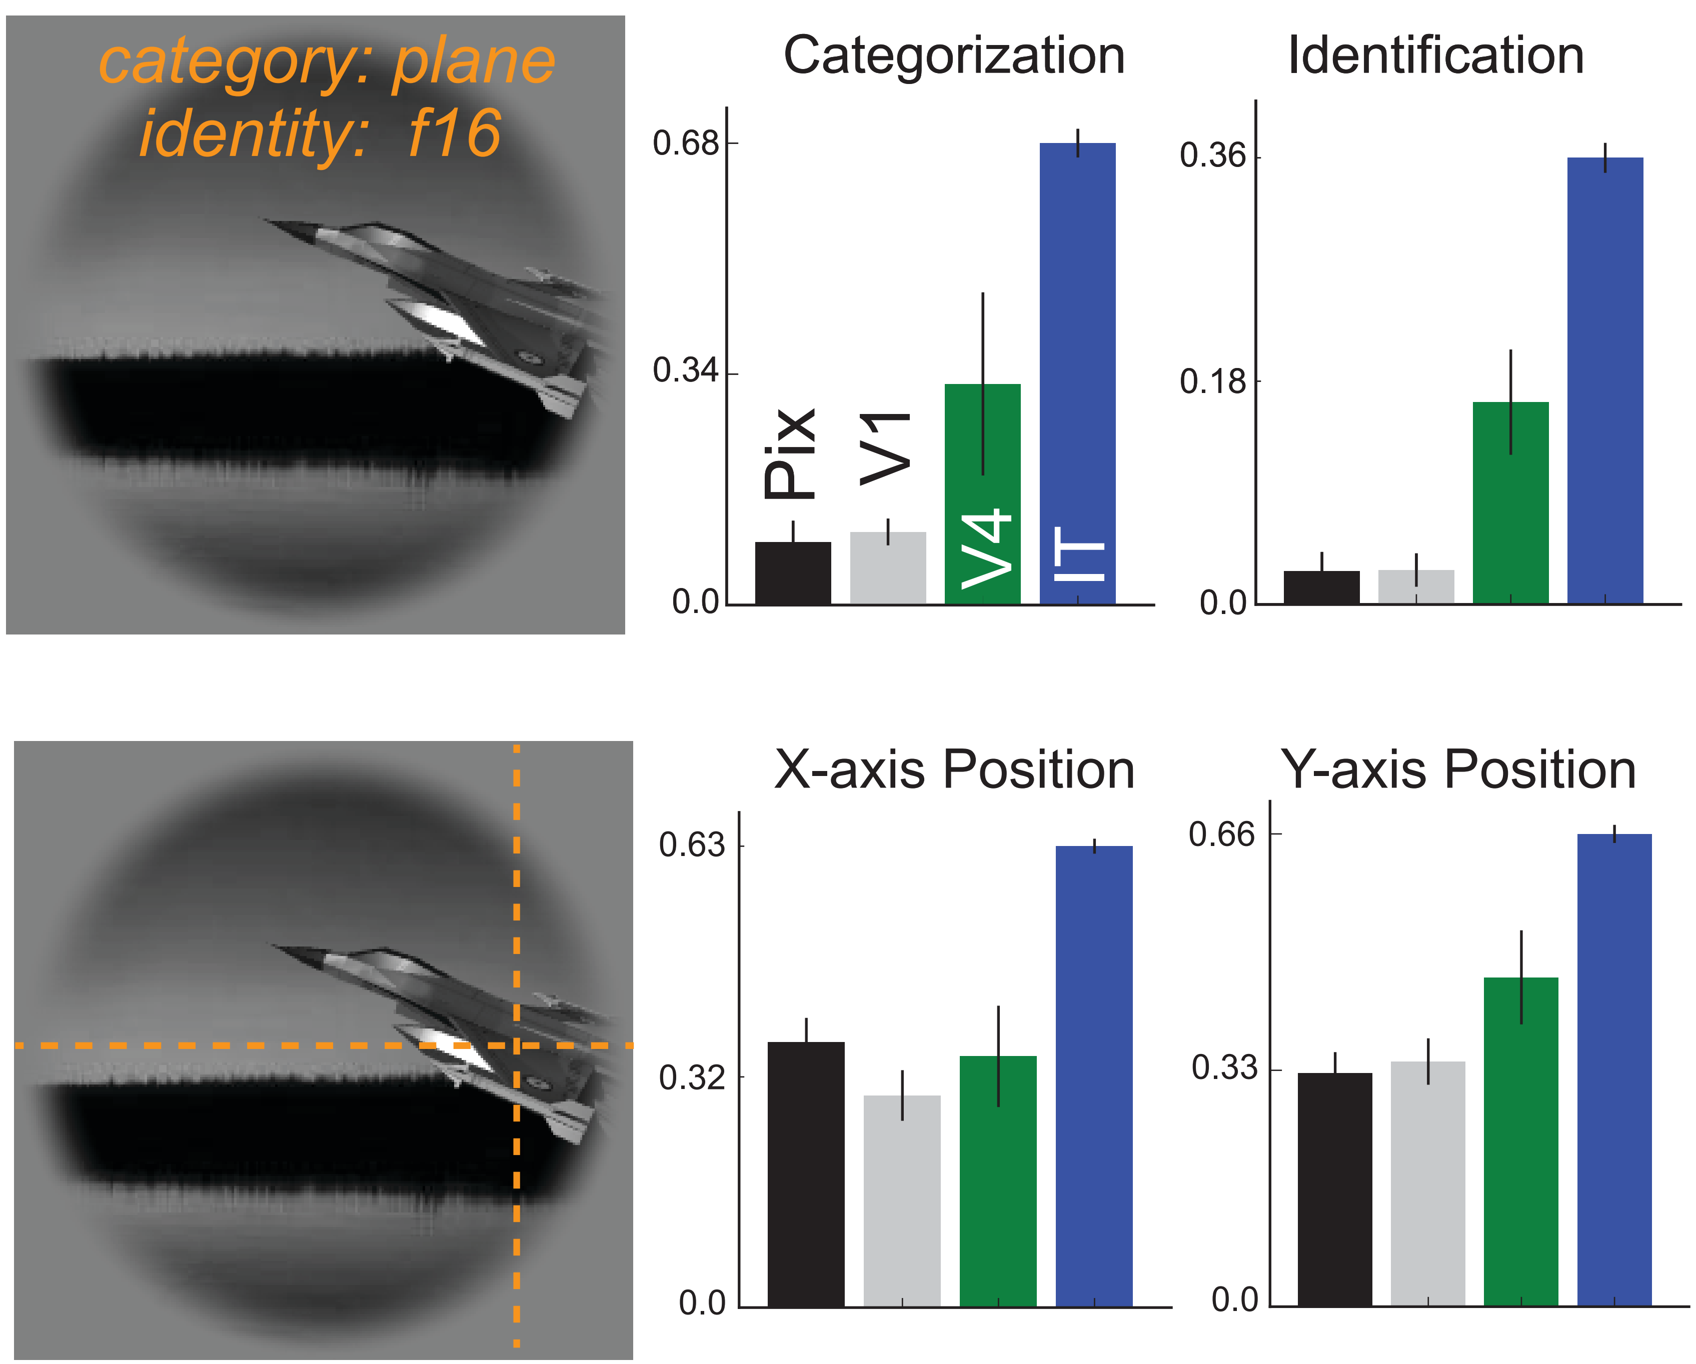

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

IT cortex

V4 cortex

V1-like model

pixel control

Population Decoding

IT > V4, V1 for all tasks **V4 > V1** for most tasks

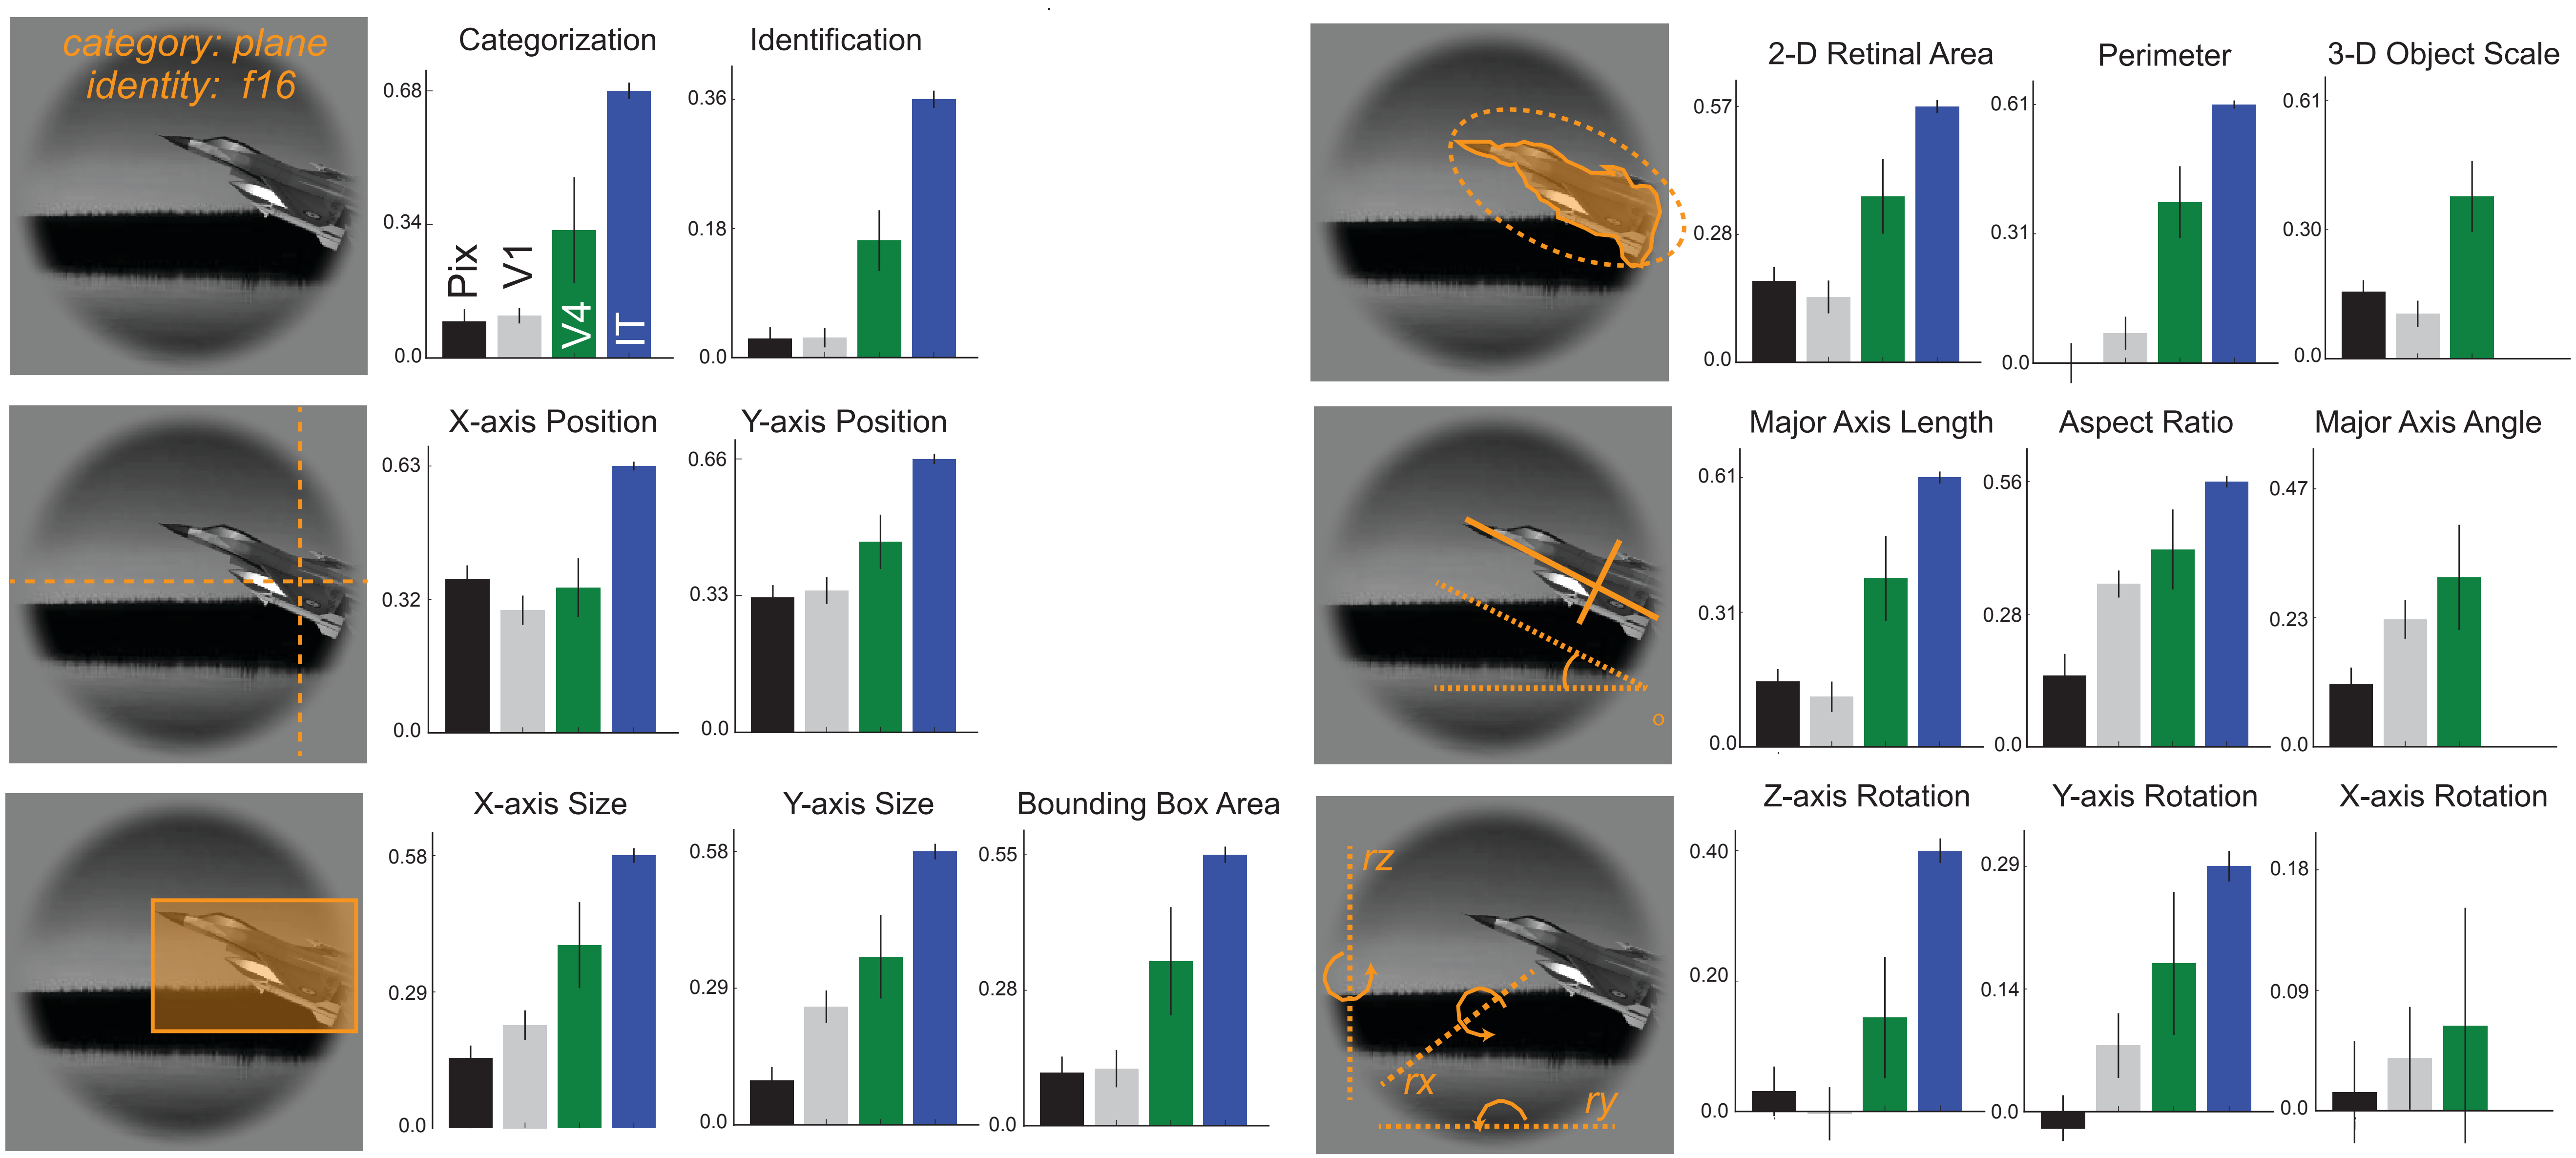

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

IT cortex

V4 cortex

V1-like model

pixel control

Single Site Responses

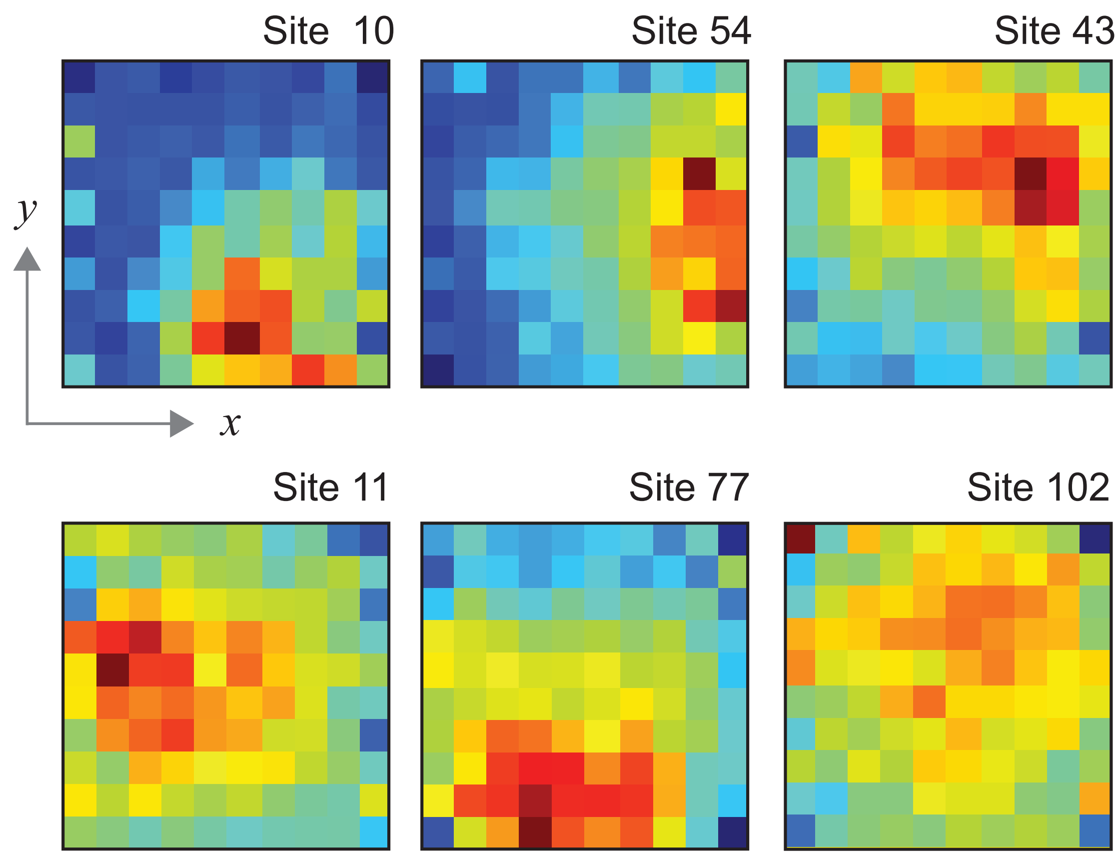

Best single position-encoding sites.

heat map value at x, y =
response averaged over all images where object center is in position x, y

Single Site Responses

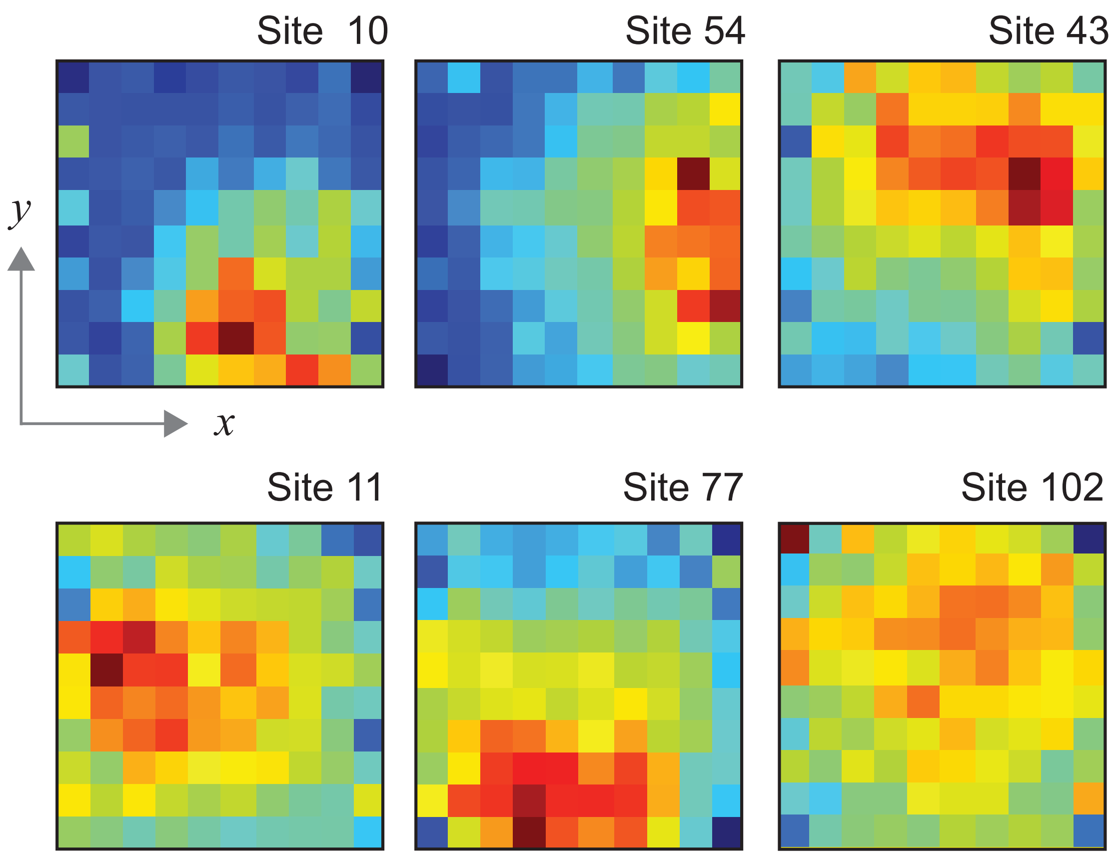

Best single position-encoding sites.

heat map value at x, y =
response averaged over all images where object center is in position x, y

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

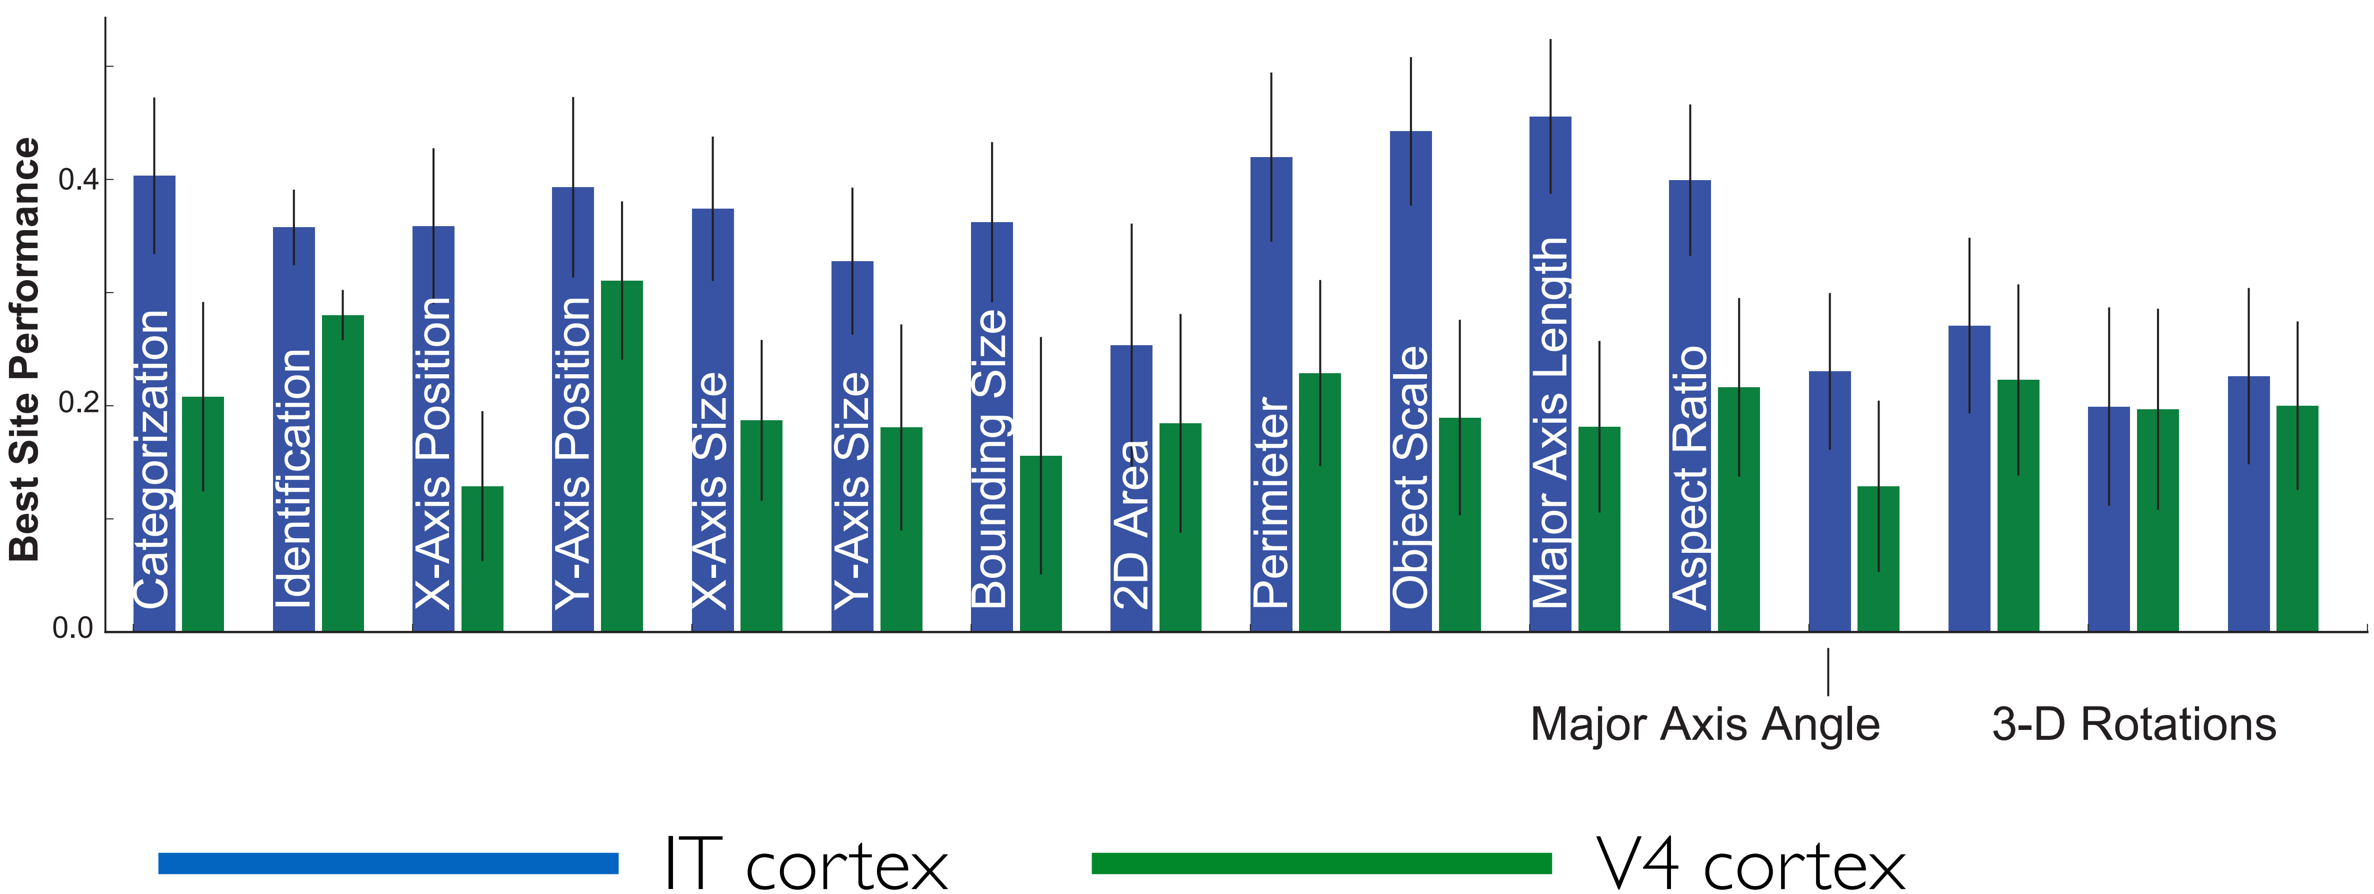

Population Decoding

"Standard" receptive field-mapping stimuli w/ position and orientation variation:

X-position

Y-position

Orientation

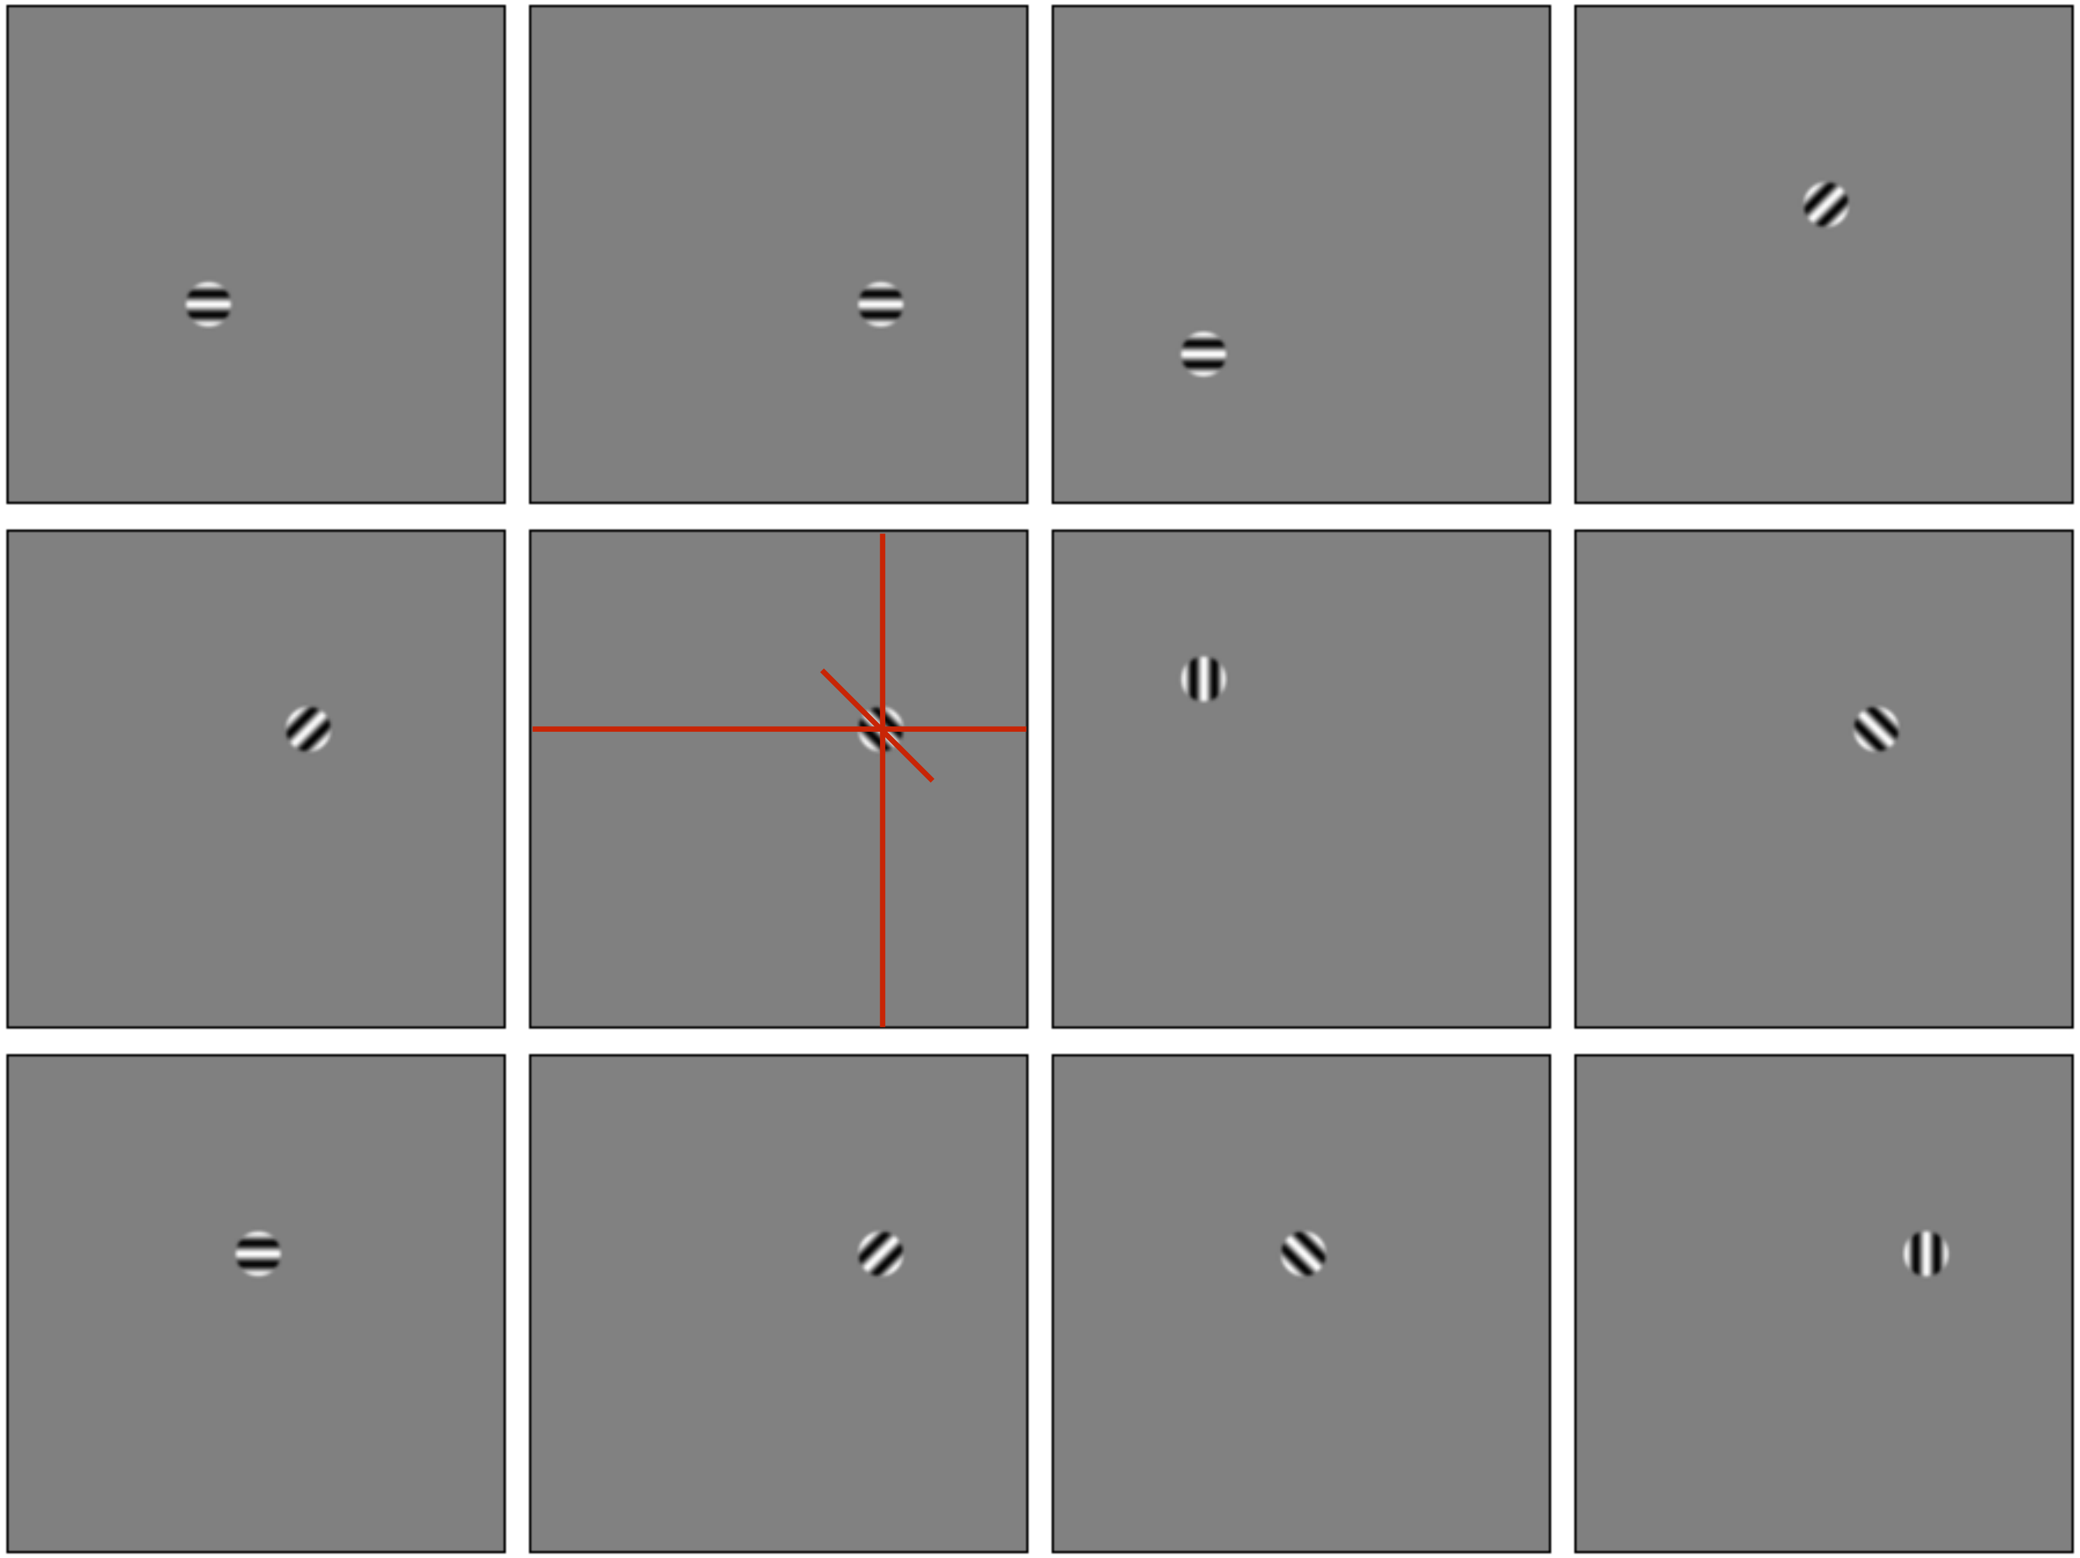

V1 > V4, IT for "standard" tasks

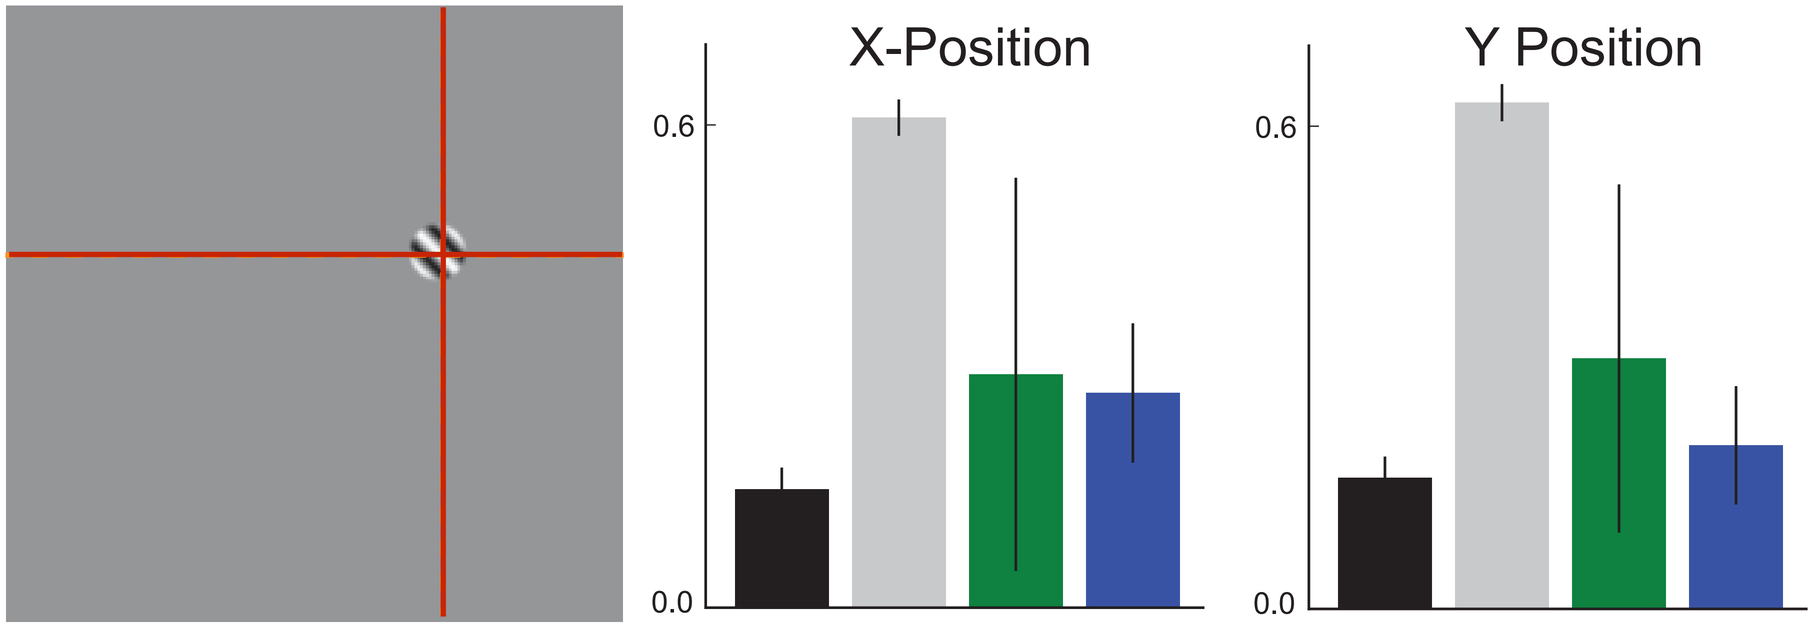

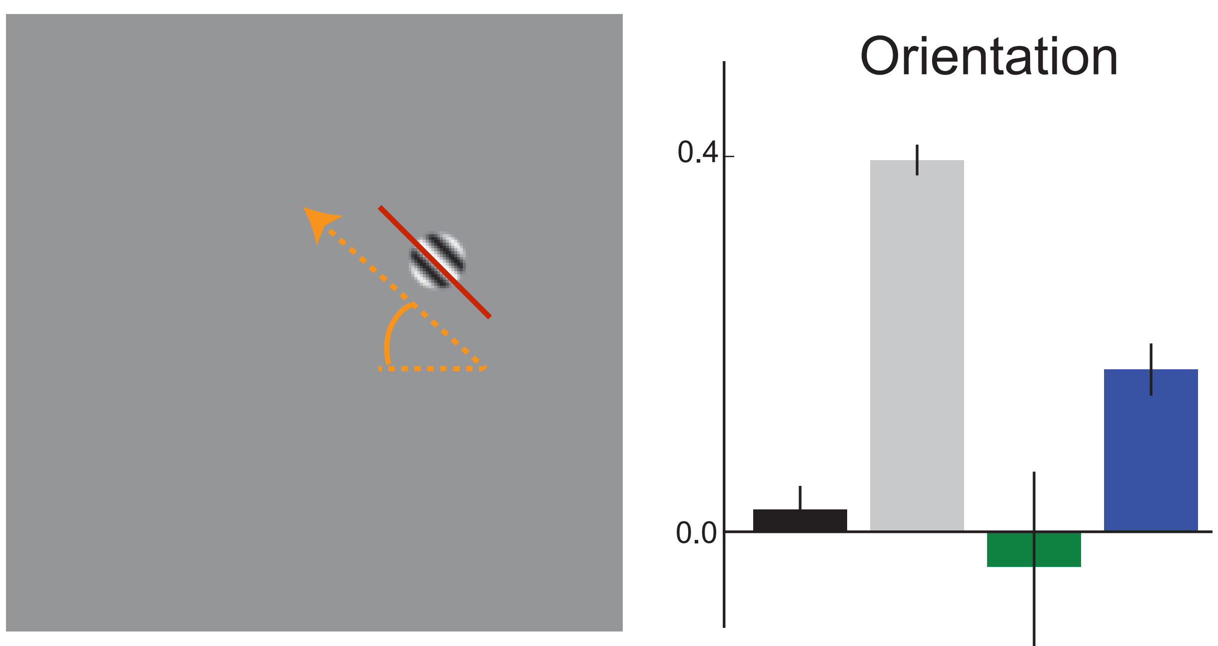

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

IT cortex

V4 cortex

V1-like model

pixel control

Human Psychophysical Measurements

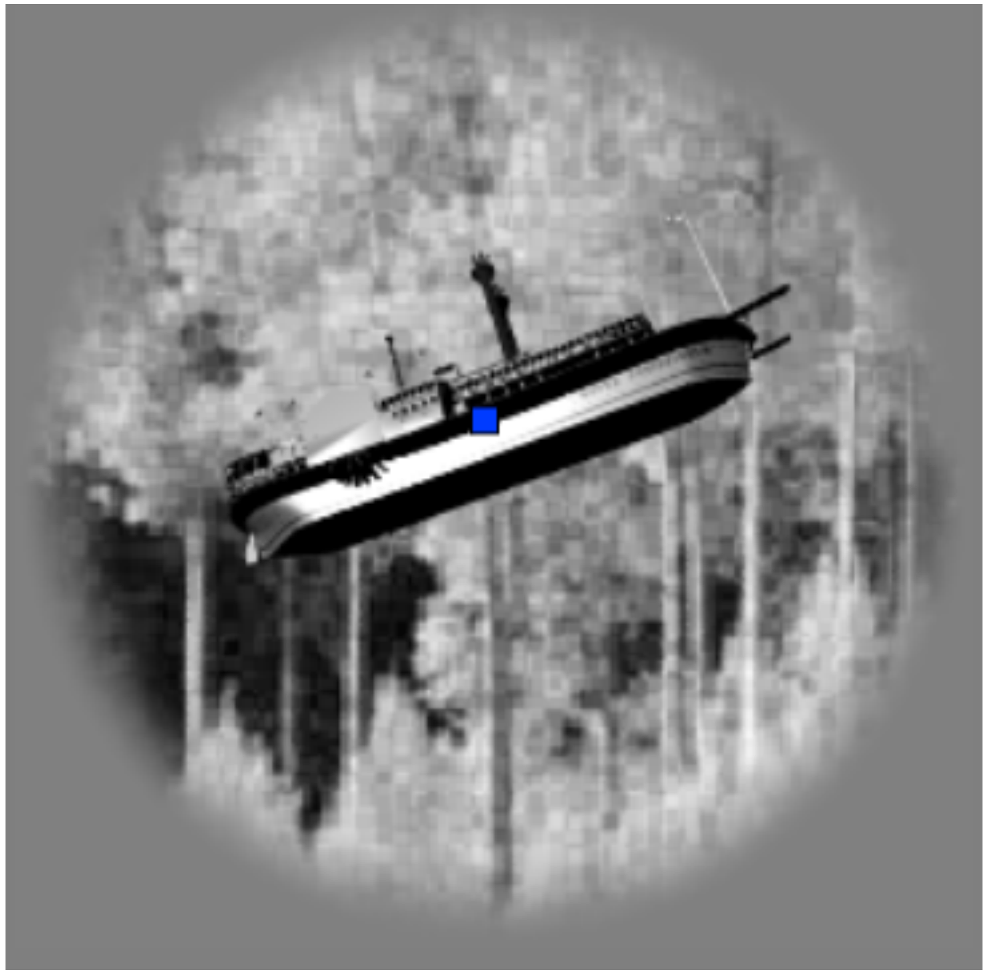

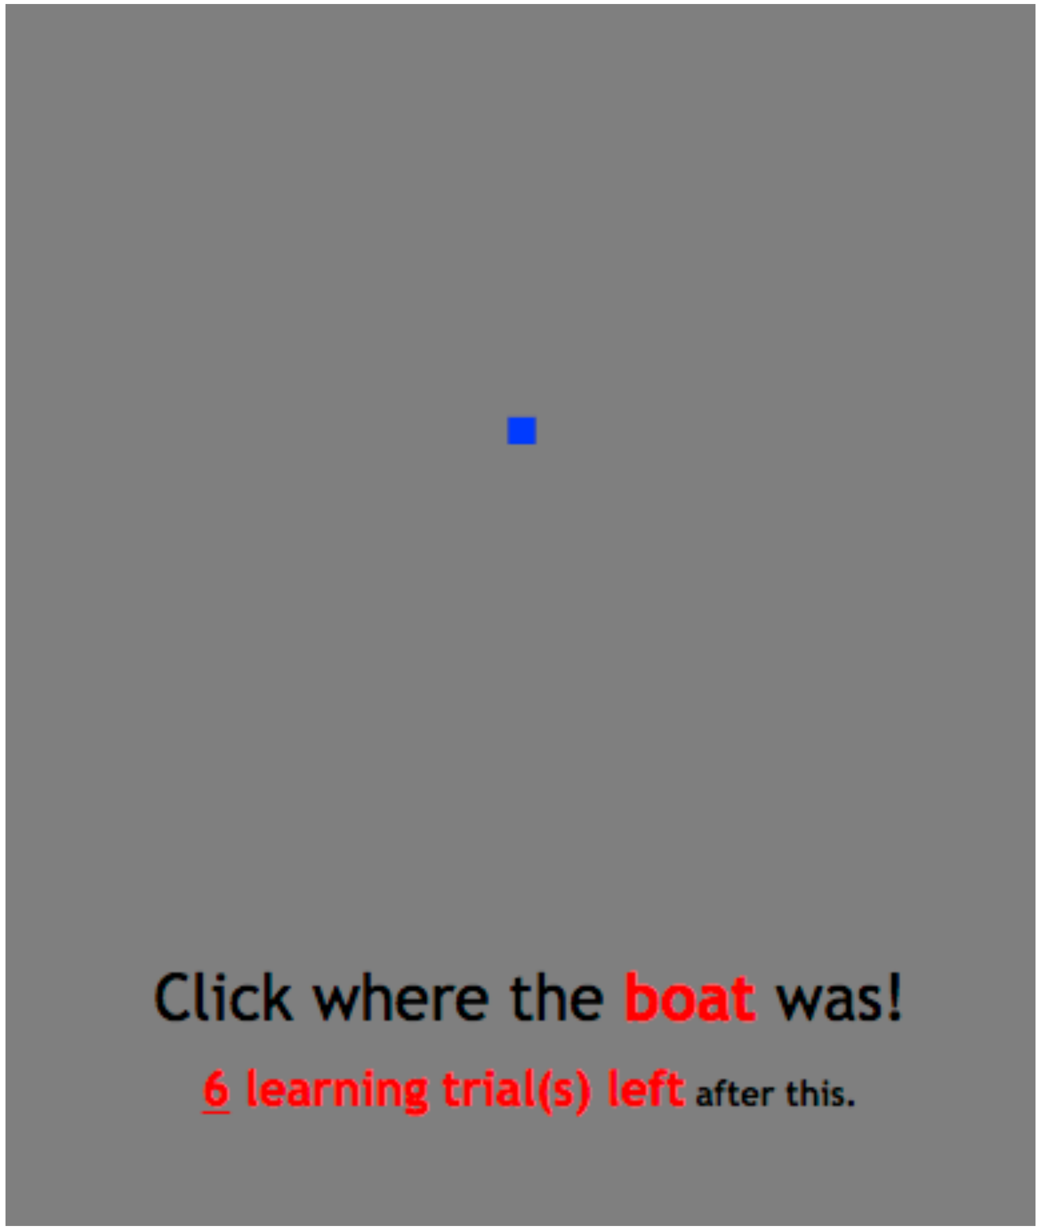

Monkey Neurons vs Humans

performance ~ k * log(N)

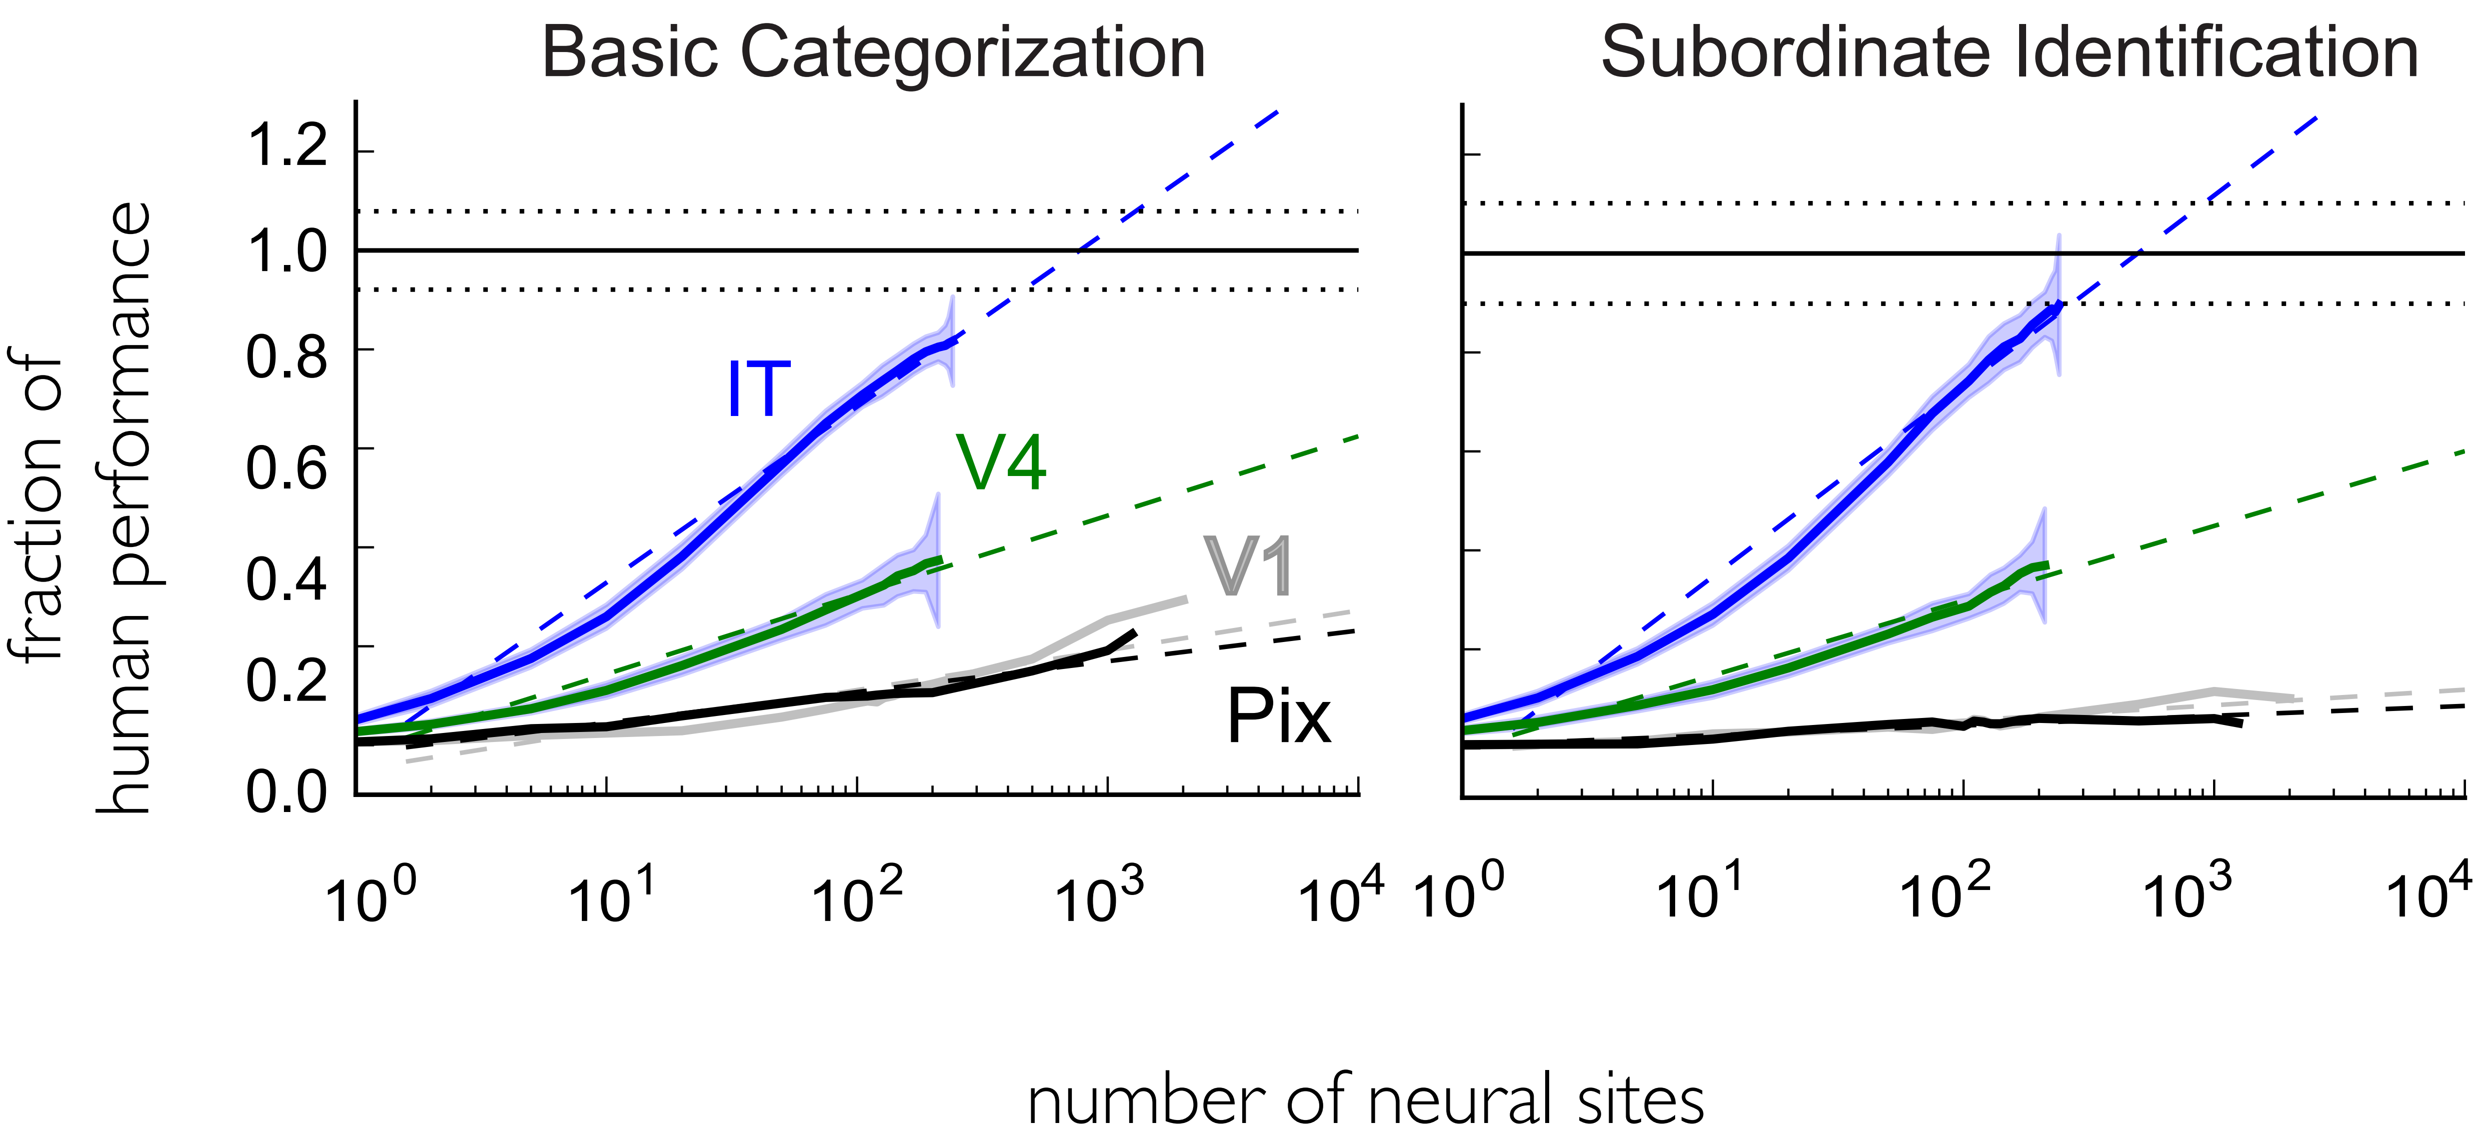

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

Monkey Neurons vs Humans

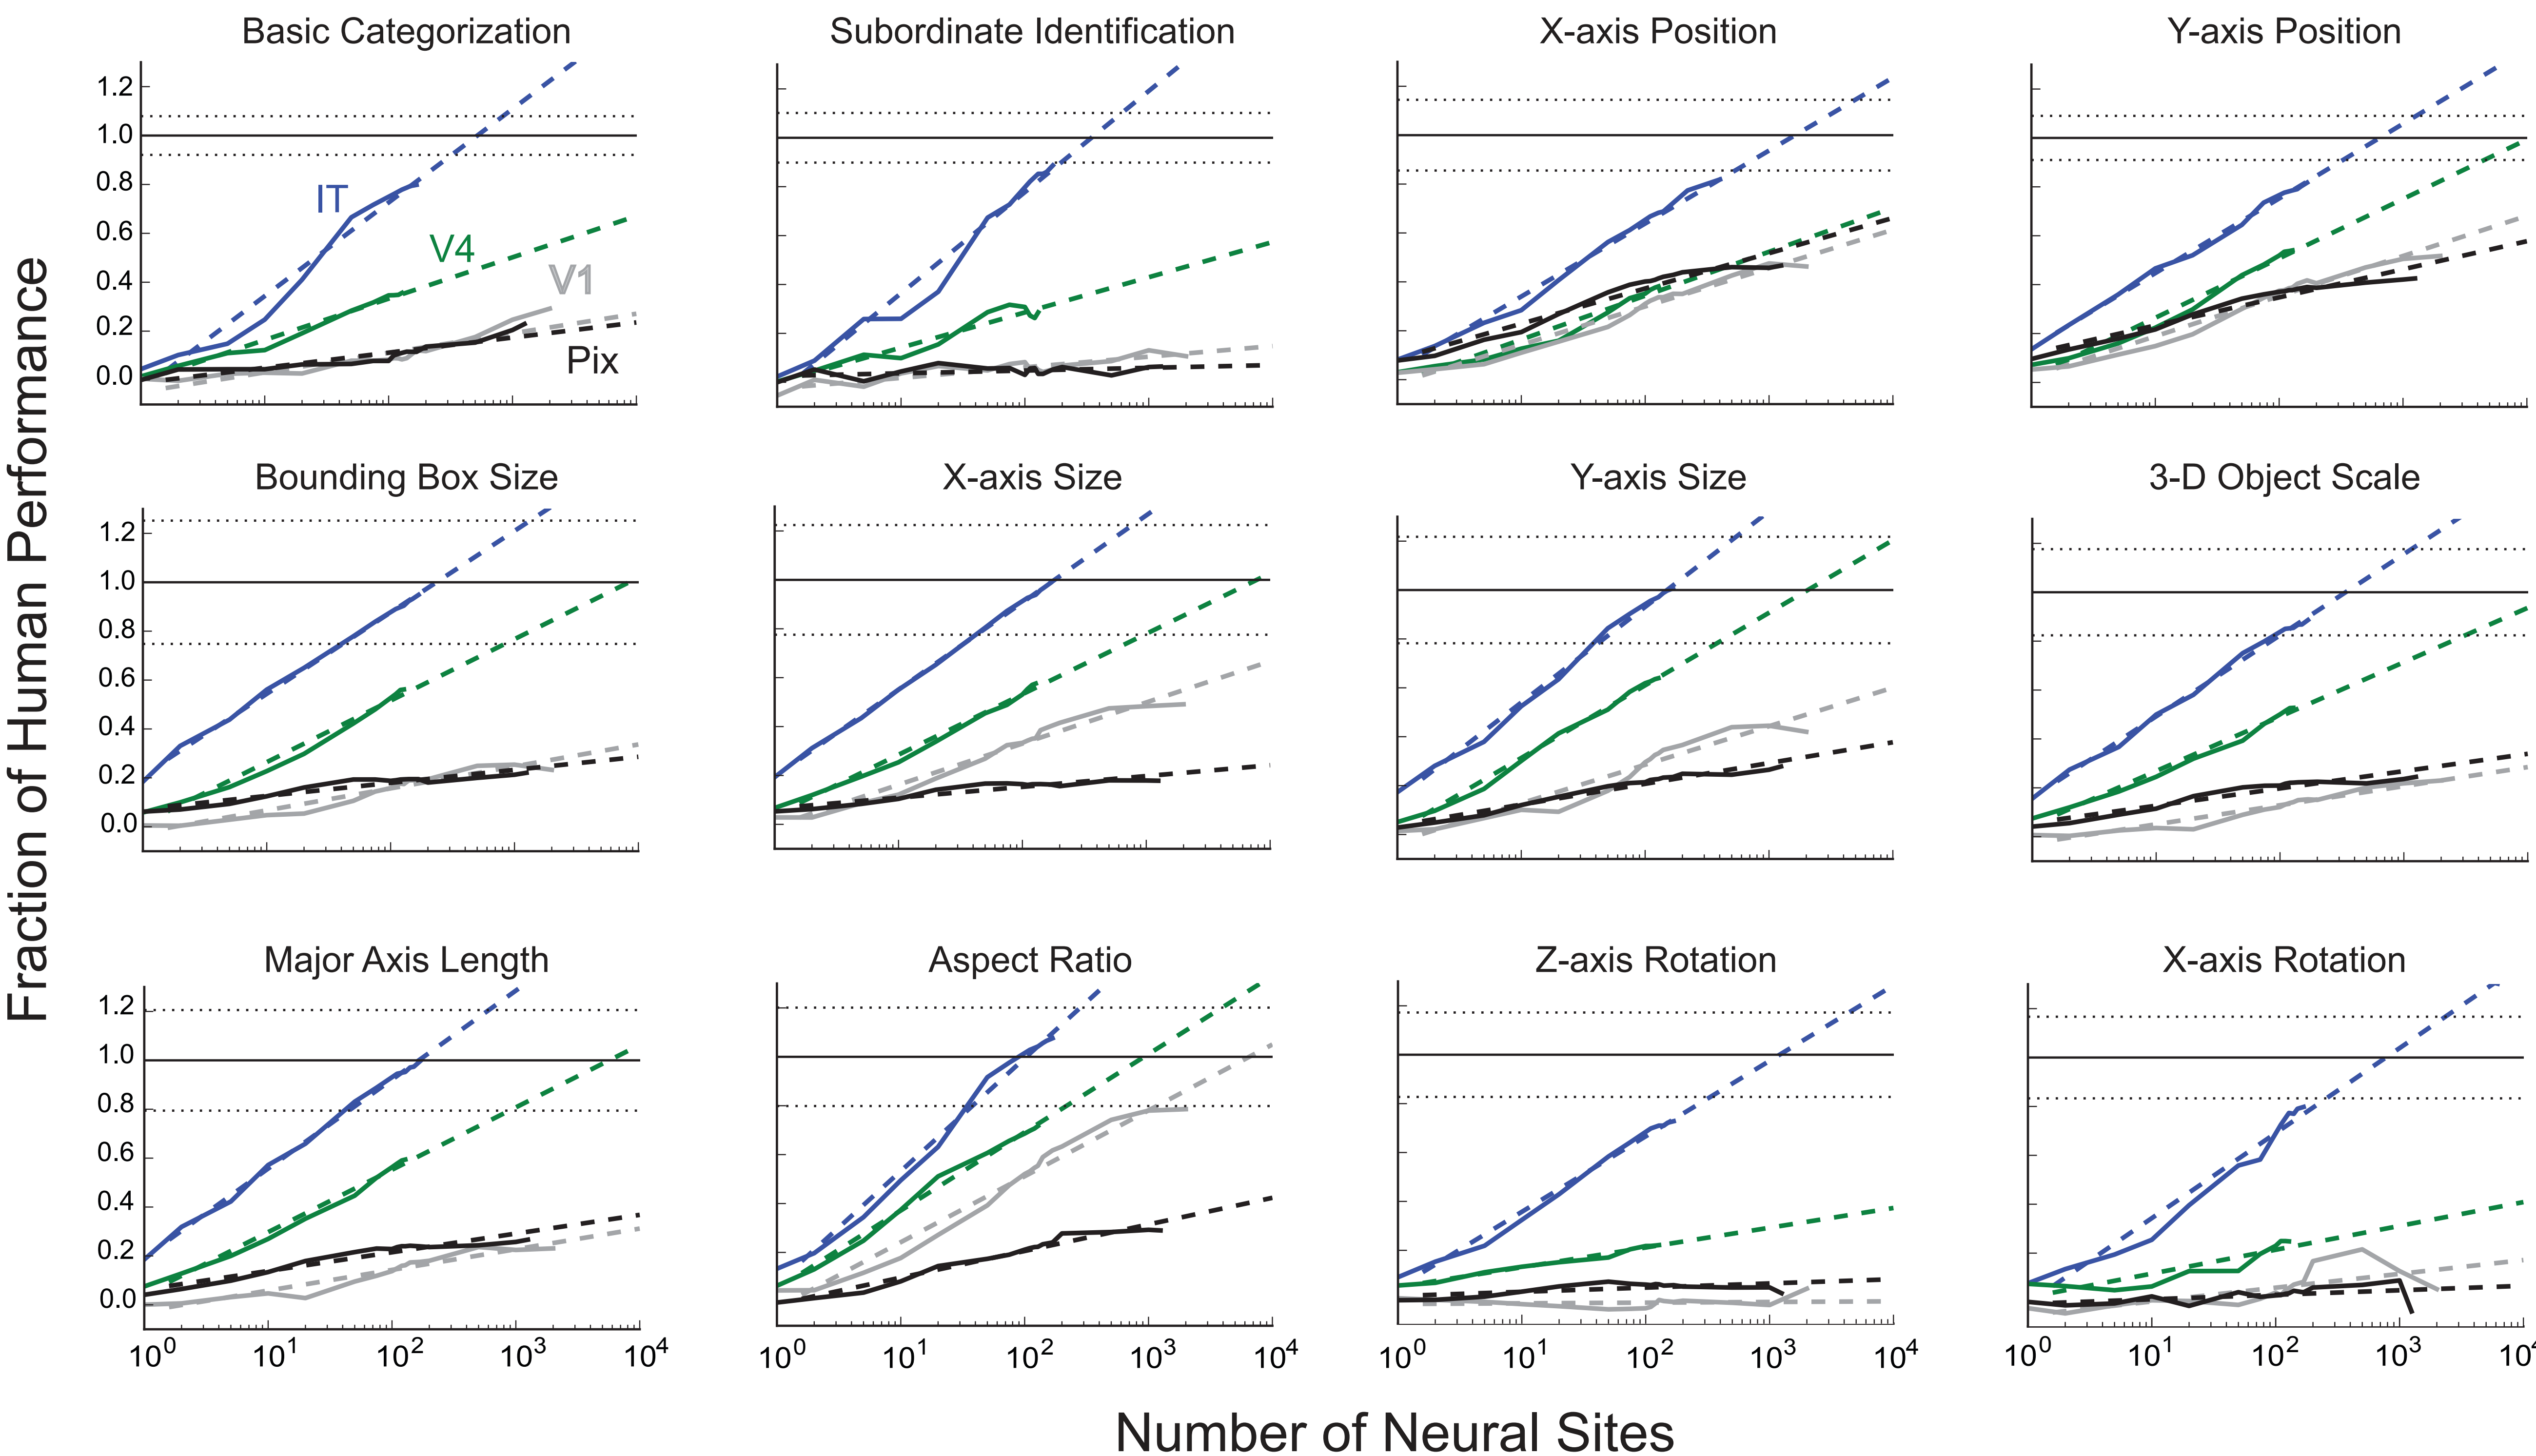

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

Monkey Neurons vs Humans

	IT	V4	V1	Pix
Basic Categorization	773 ± 185	2.2×10^6	—	—
Subordinate Identification	496 ± 93	4.4×10^6	—	—
X-axis Position	1414 ± 403	5.2×10^5	3.0×10^7	—
Y-axis Position	918 ± 309	2.5×10^4	8.7×10^6	—
Bounding Box Size	322 ± 90	1.7×10^4	—	—
X-axis Size	256 ± 87	9.8×10^3	3.4×10^7	—
Y-axis Size	237 ± 87	3.8×10^3	9.5×10^6	—
3-D Object Scale	401 ± 90	3.2×10^4	—	—
Major Axis Length	201 ± 70	1.1×10^4	—	—
Aspect Ratio	163 ± 61	951 ± 59	6.5×10^3	—
Major Axis Angle	804 ± 136	3.2×10^6	—	—
Z-axis Rotation	1932 ± 1061	—	—	—
Y-axis Rotation	369 ± 115	2.8×10^5	—	—
X-axis Rotation	1570 ± 530	—	—	—

— = more than 10 billion sites required

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

Mean over tasks, human-parity for IT is at ~**700** multi-unit trial-averaged sites.

Monkey Neurons vs Humans

	IT	V4	V1	Pix
Basic Categorization	773 ± 185	2.2×10^6	—	—
Subordinate Identification	496 ± 93	4.4×10^6	—	—
X-axis Position	1414 ± 403	5.2×10^5	3.0×10^7	—
Y-axis Position	918 ± 309	2.5×10^4	8.7×10^6	—
Bounding Box Size	322 ± 90	1.7×10^4	—	—
X-axis Size	256 ± 87	9.8×10^3	3.4×10^7	—
Y-axis Size	237 ± 87	3.8×10^3	9.5×10^6	—
3-D Object Scale	401 ± 90	3.2×10^4	—	—
Major Axis Length	201 ± 70	1.1×10^4	—	—
Aspect Ratio	163 ± 61	951 ± 59	6.5×10^3	—
Major Axis Angle	804 ± 136	3.2×10^6	—	—
Z-axis Rotation	1932 ± 1061	—	—	—
Y-axis Rotation	369 ± 115	2.8×10^5	—	—
X-axis Rotation	1570 ± 530	—	—	—

— = more than 10 billion sites required

Hong, Yamins*, Majaj & DiCarlo.* ***Nat. Neuro.*** *(2016)*

Mean over tasks, human-parity for IT is at ~**350000** single-unit single-trial neurons.

Somewhat newish ideas about IT?

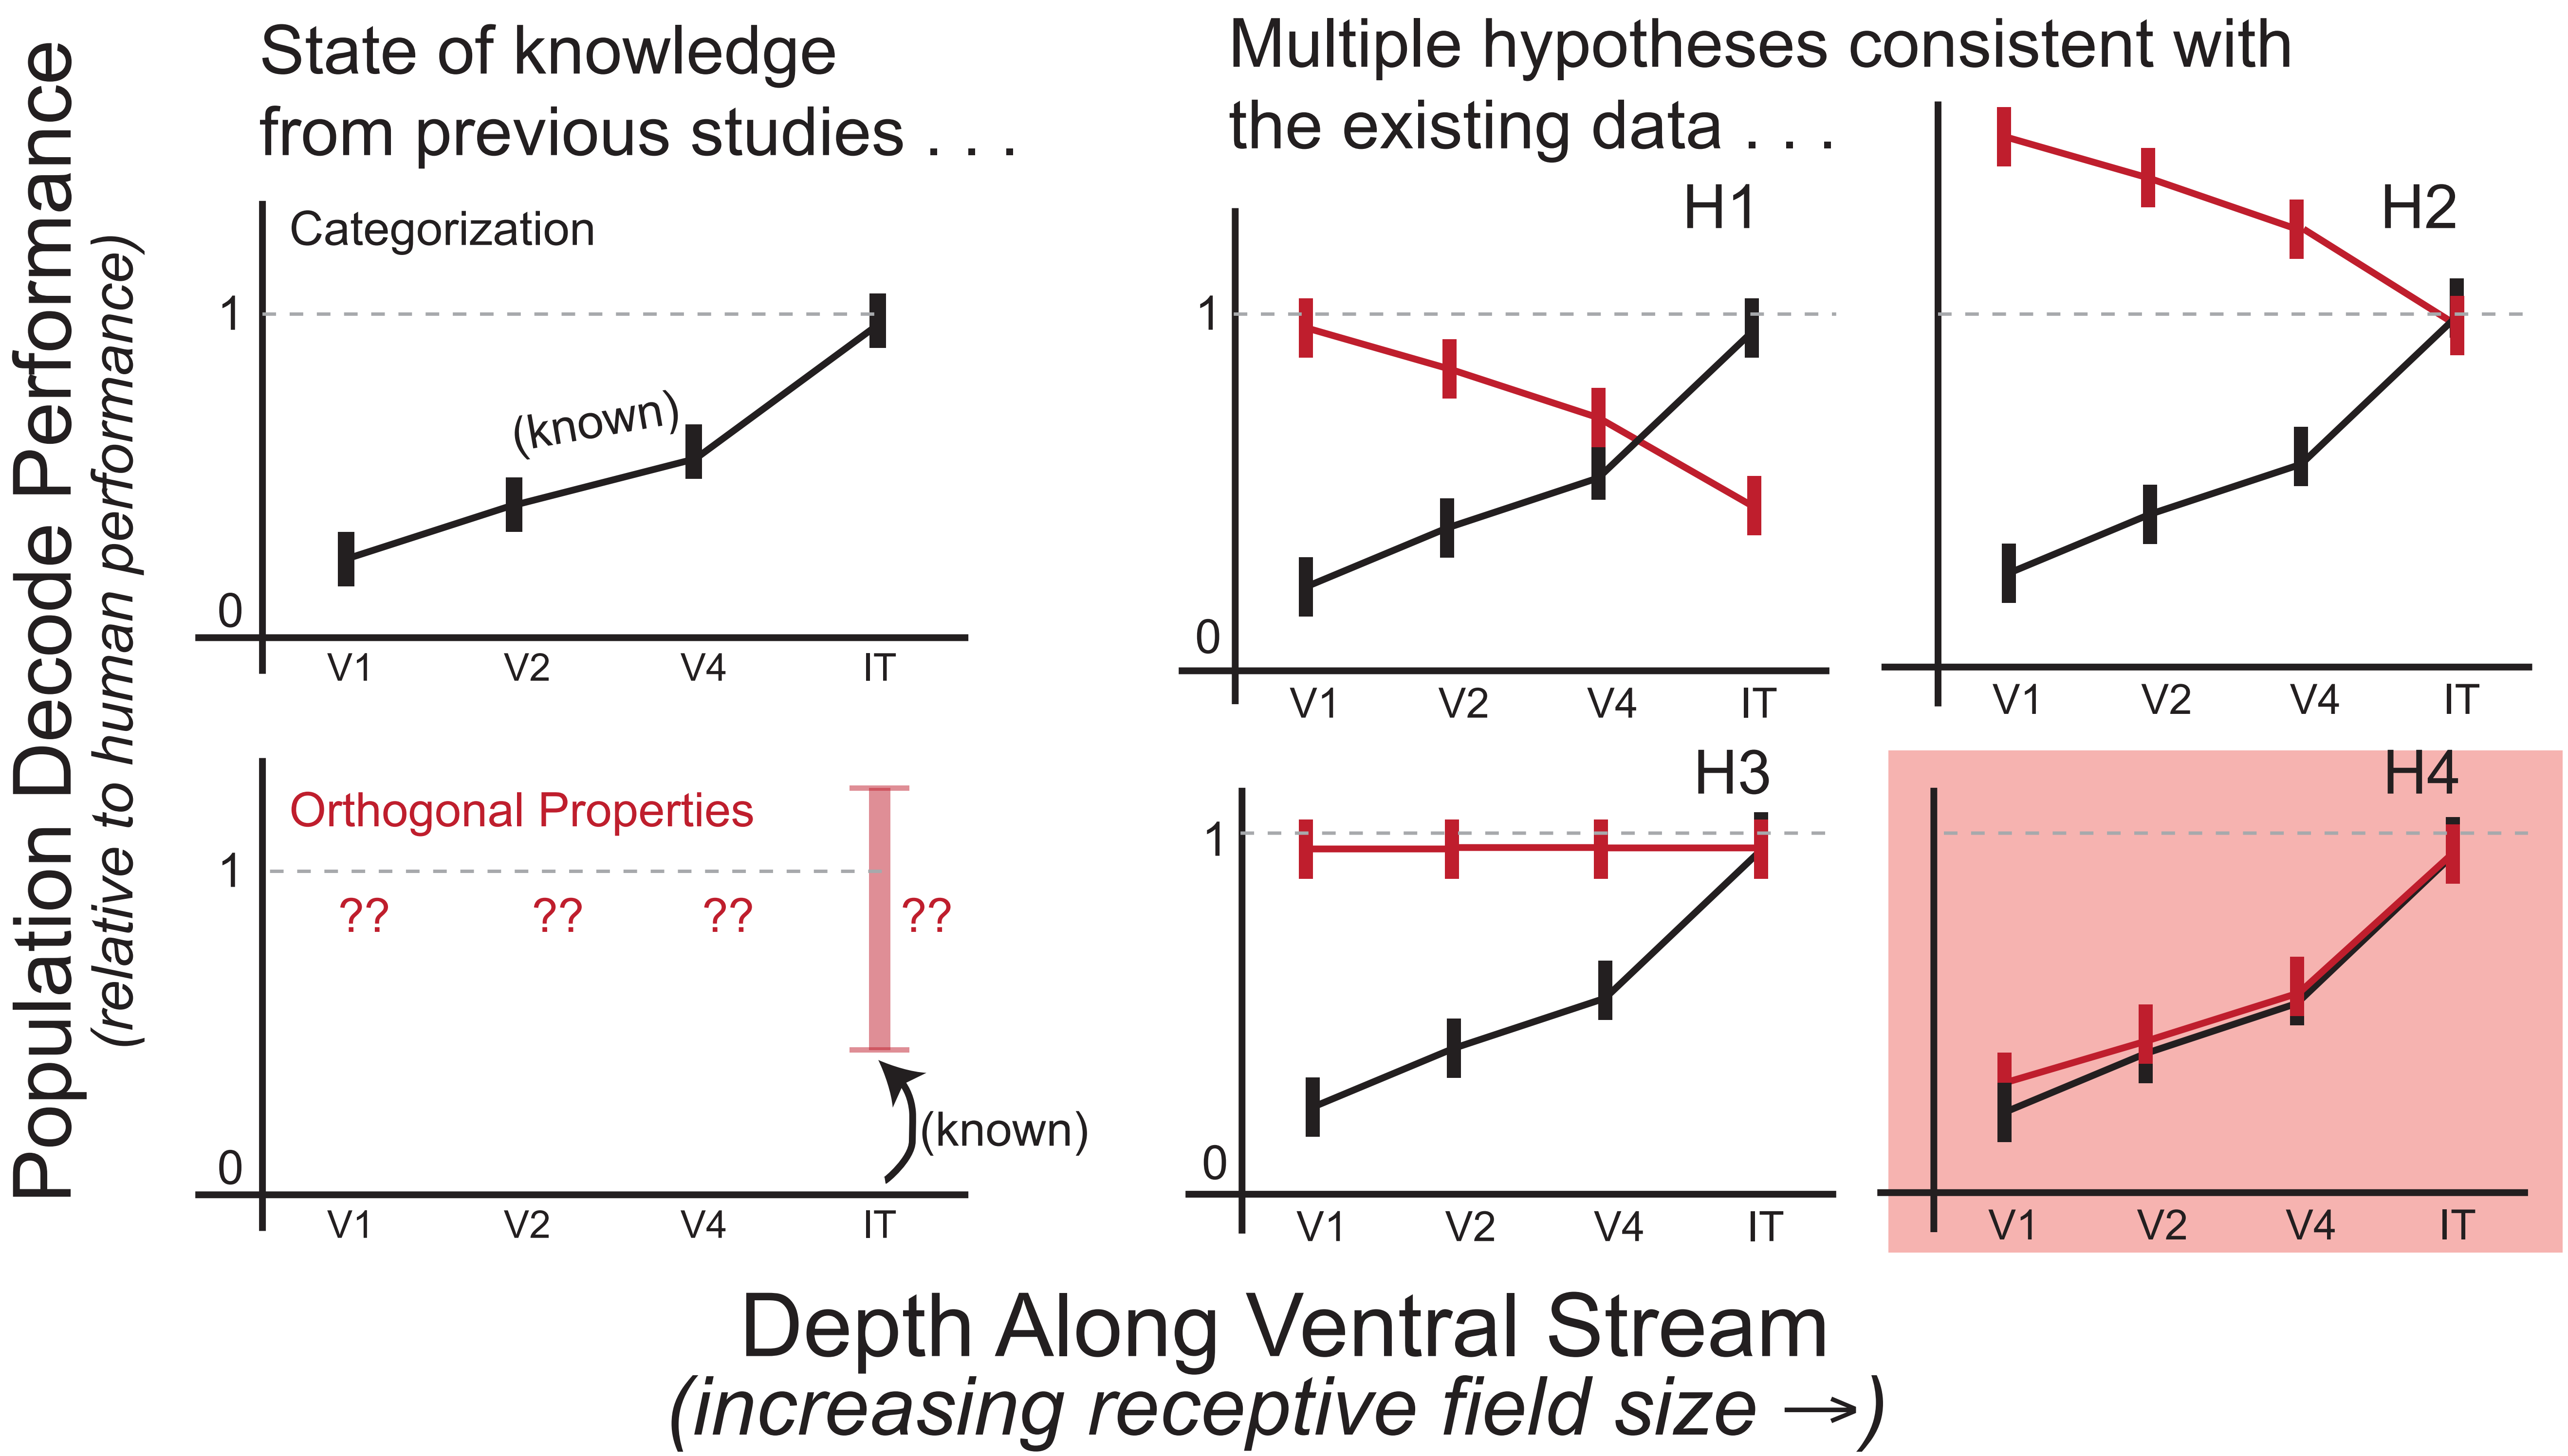

H4: Simultaneous build-up of encoding

Somewhat newish ideas about IT?

1. IT is *NOT* invariant. Strict generalization of simple-to-complex cells: **no**.

2. "Lower-level" properties are not that low-level — at least, with complex objects and backgrounds.

3. Categorization and non-categorical properties "go together" — *not* just that "not all (e.g.) position information is lost" (MacEvoy 2013, DiCarlo 2003)

Provides support to a hypothesis for what IT does:

"Inverting the generative model of the scene"

But what type of understanding is this?

But what type of understanding is this?

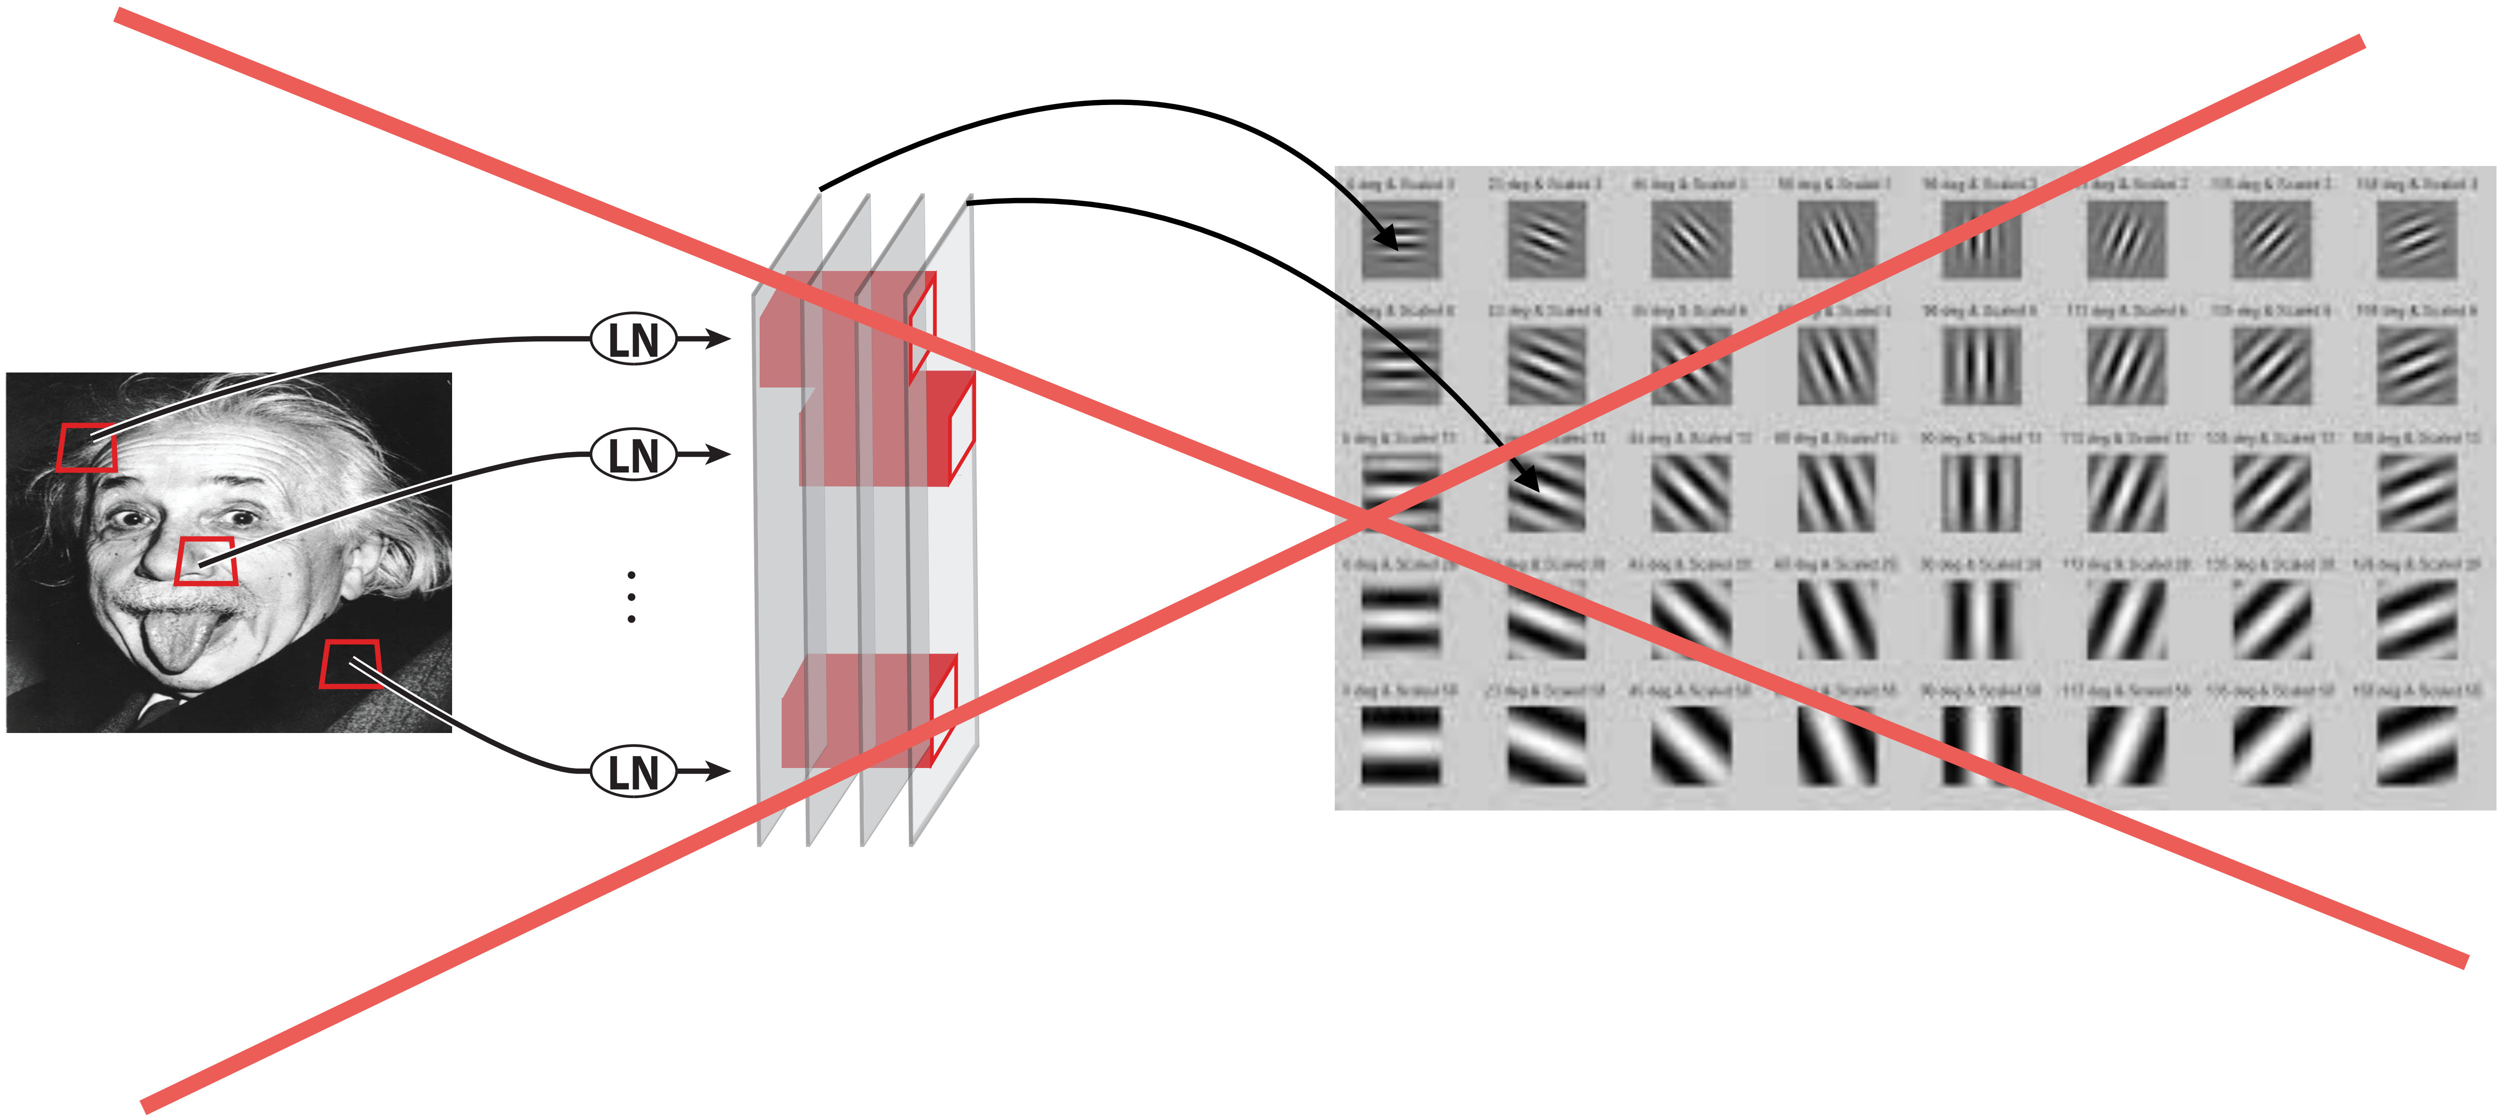

not saying this type of understanding is impossible …

Principle of "Goal-Driven Modeling"

> Formulate comprehensive model class (**CNNs**)

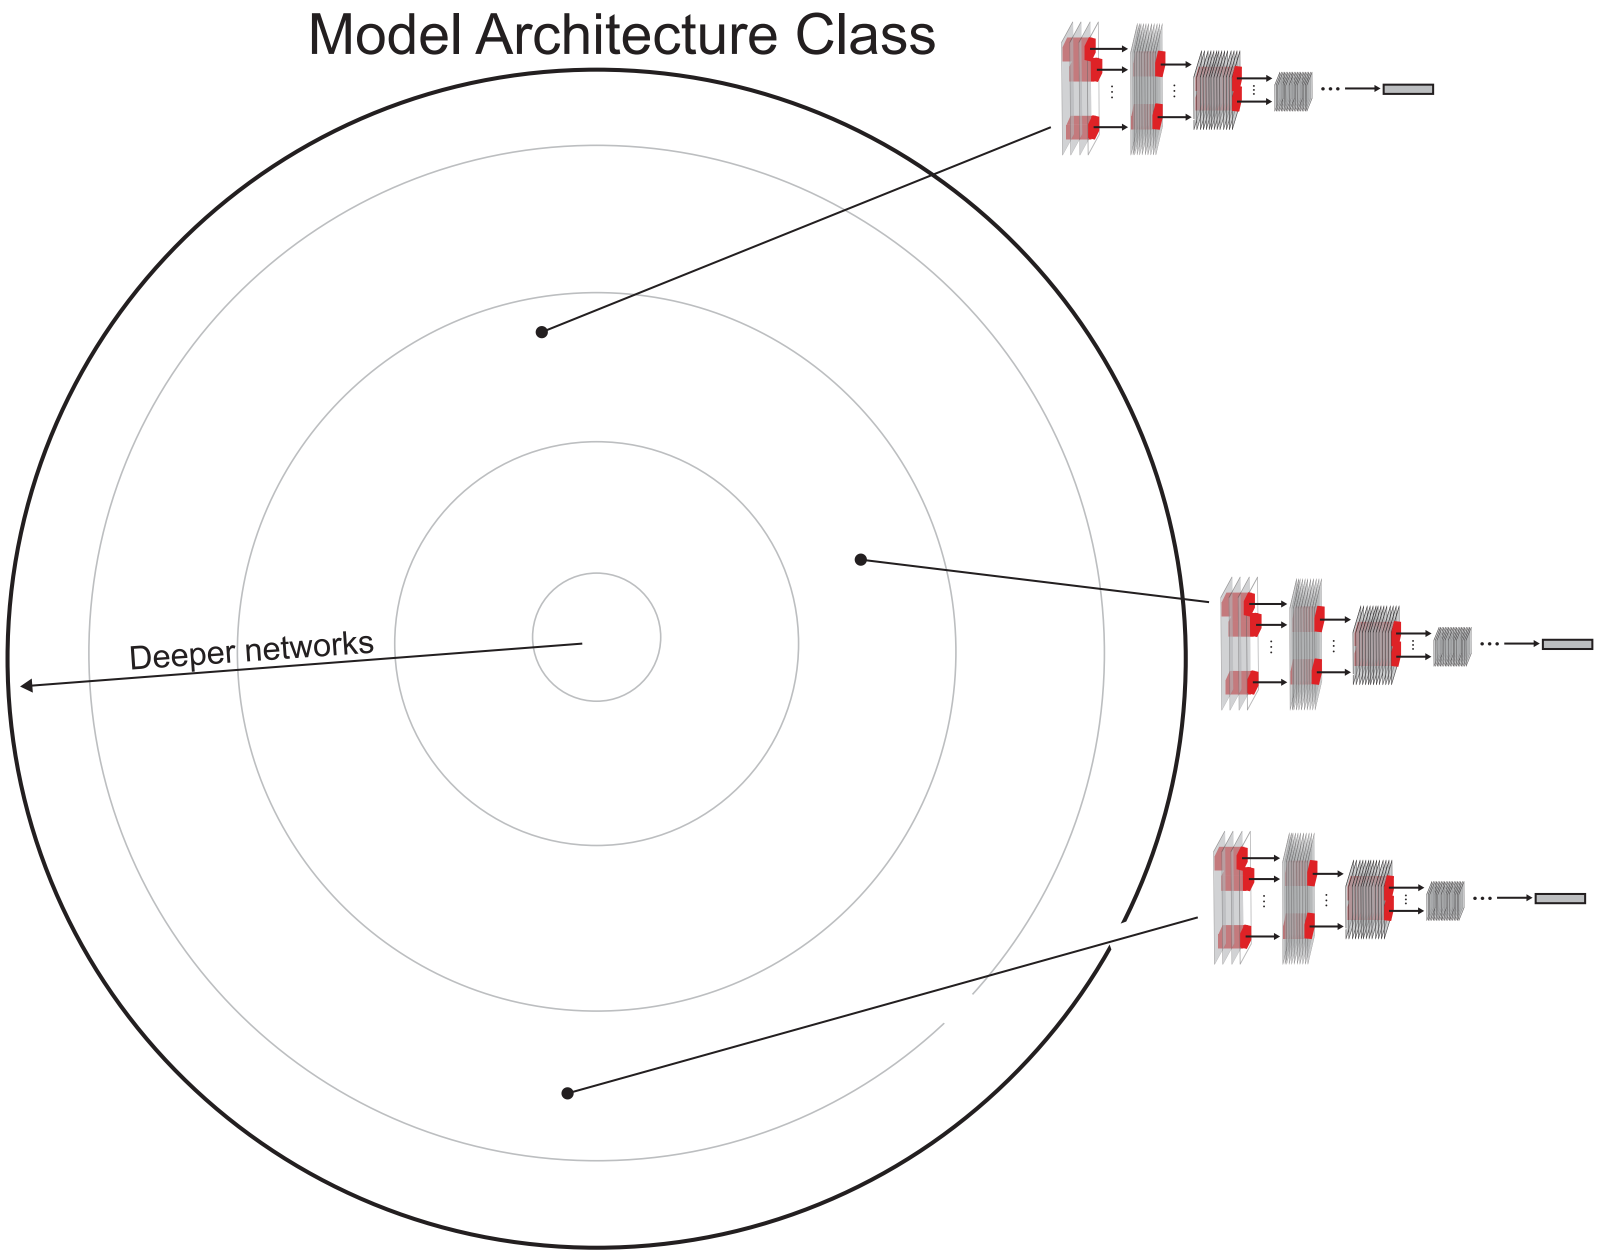

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

> Formulate comprehensive model class (**CNNs**)

> Choose challenging, ethologically-valid tasks (**categorization**)

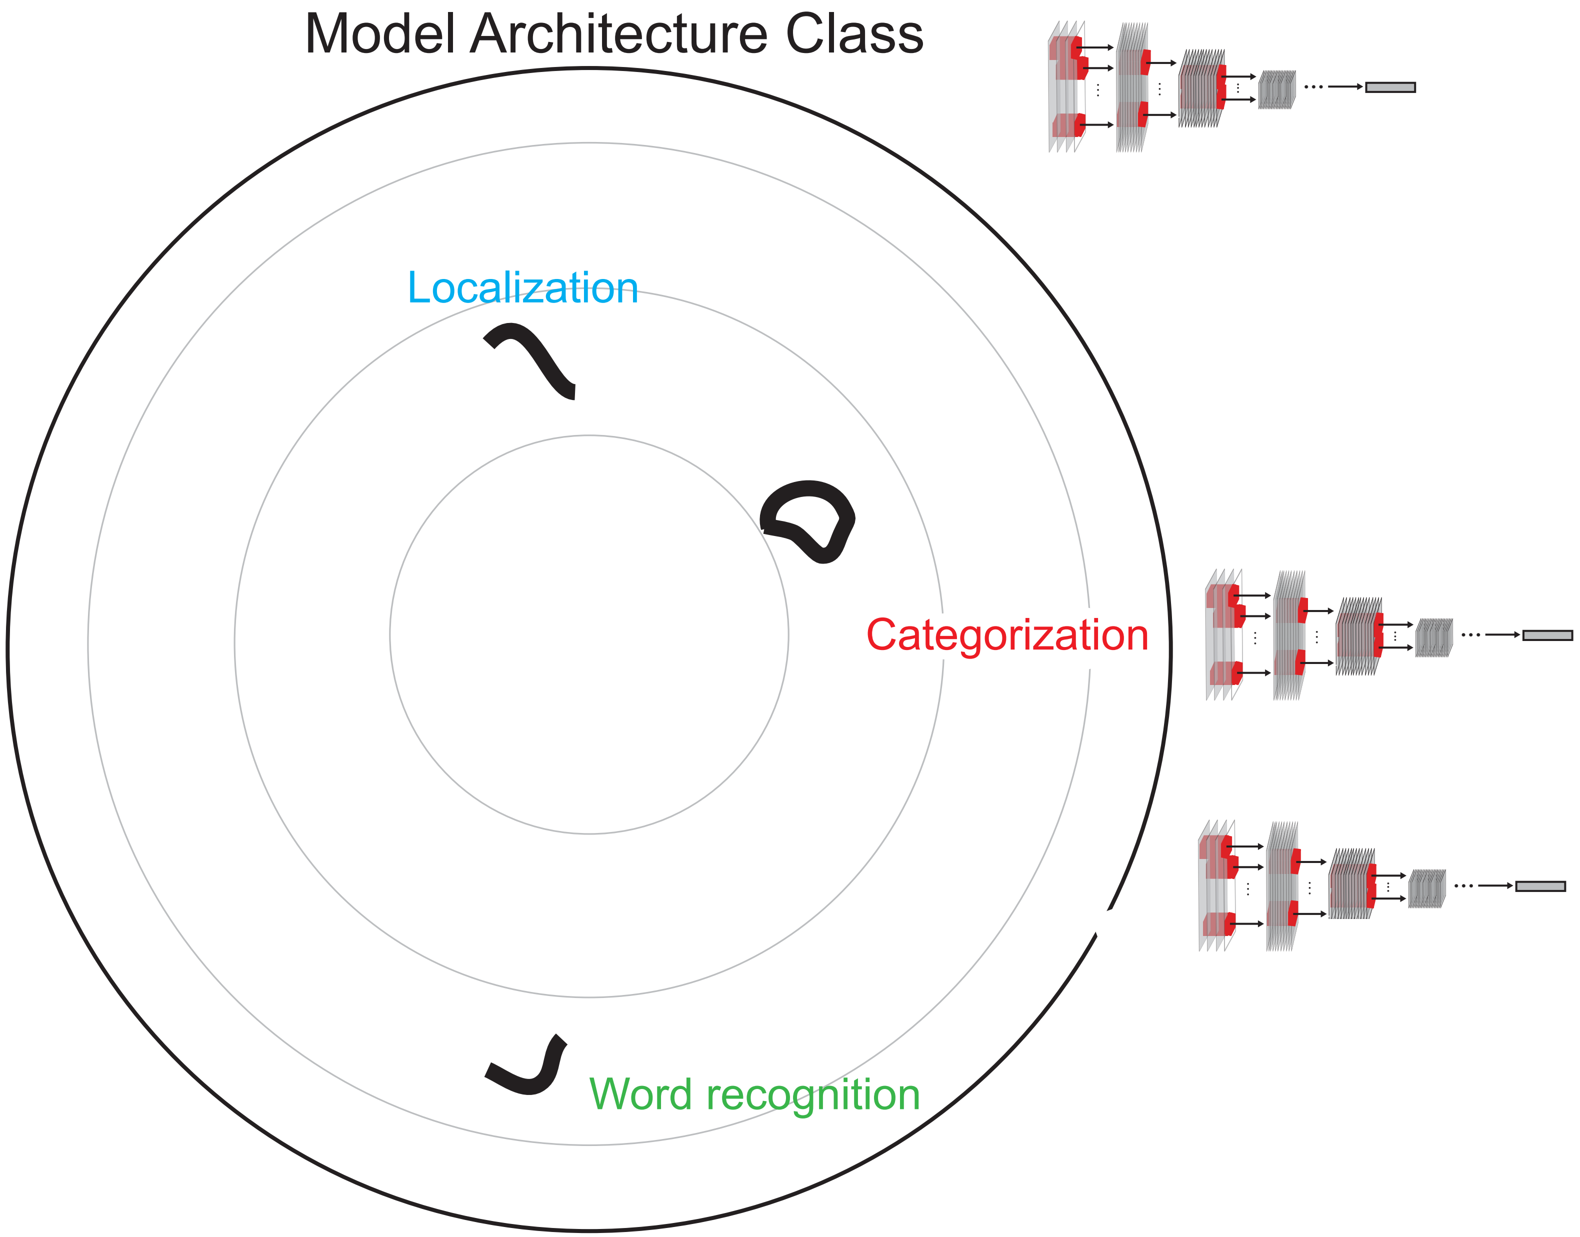

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

> Formulate comprehensive model class (**CNNs**)

> Choose challenging, ethologically-valid tasks (**categorization**)

> Implement generic learning rules (**gradient descent**)

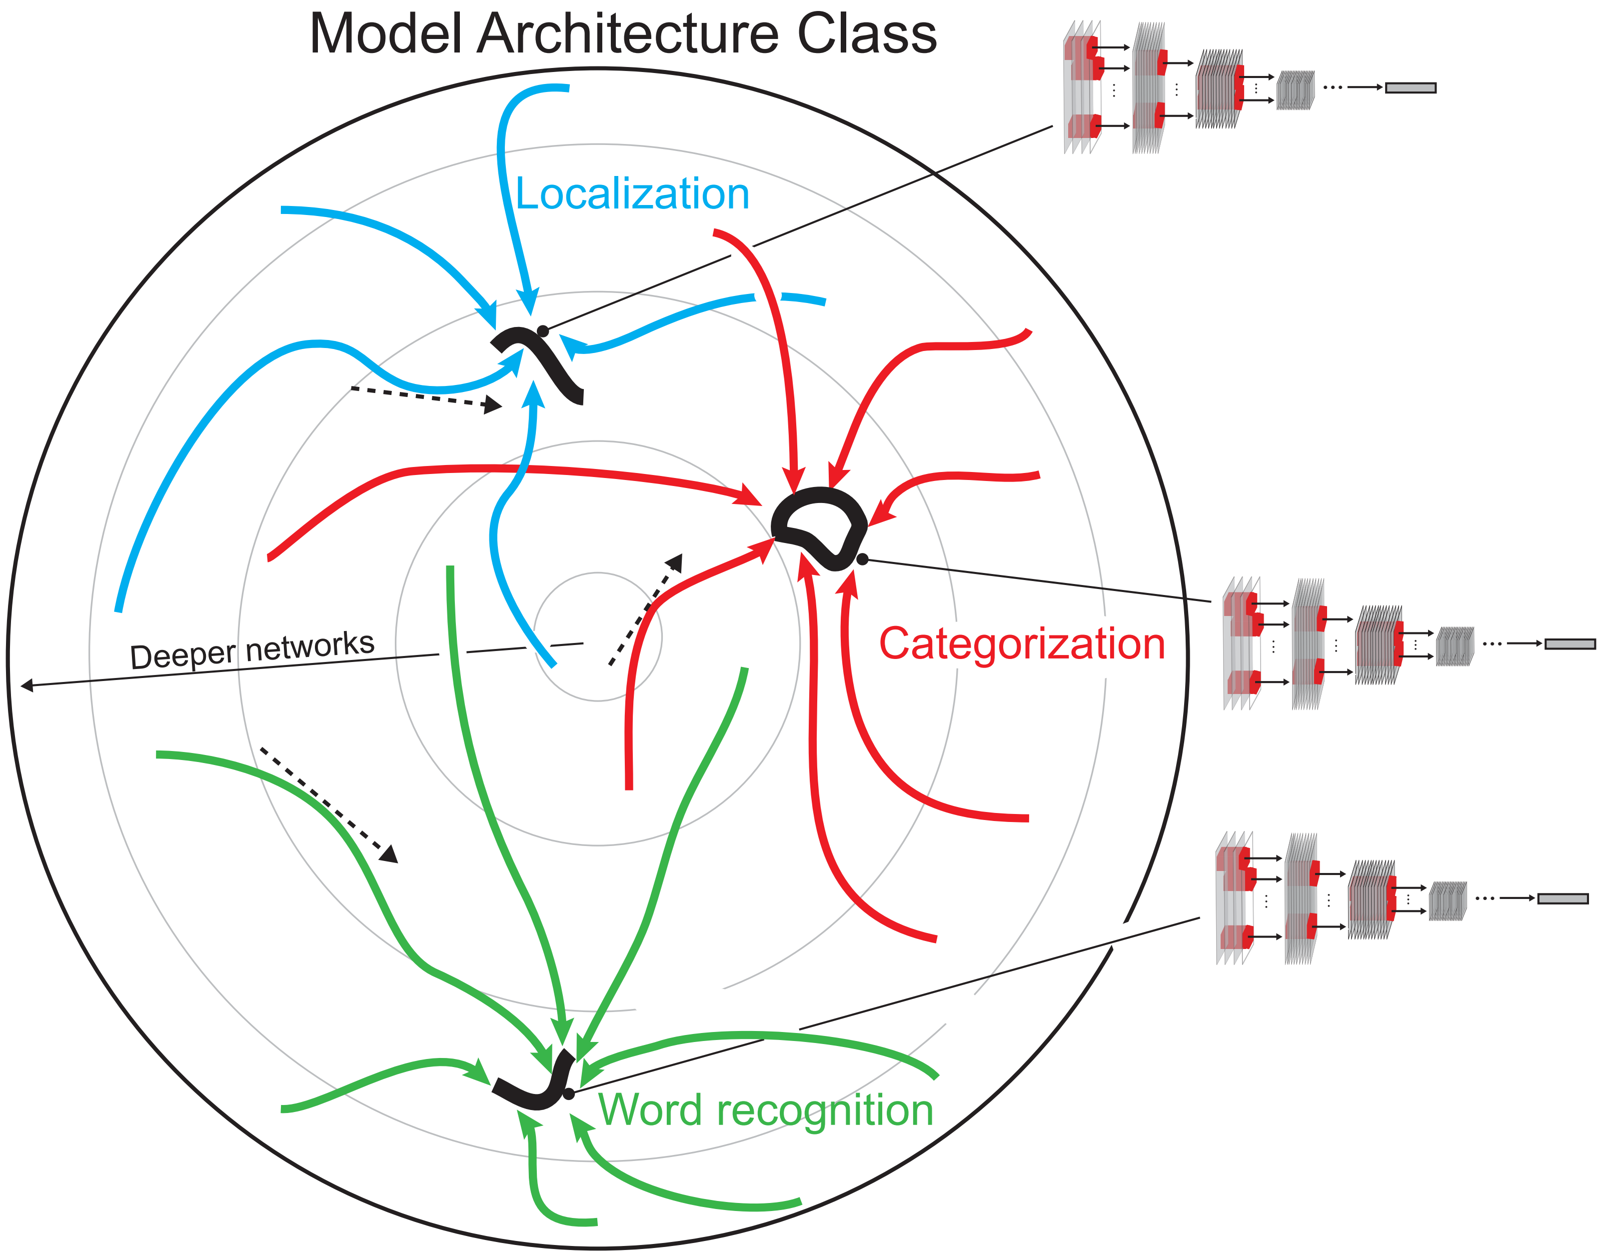

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

> Formulate comprehensive model class (**CNNs**)

> Choose challenging, ethologically-valid tasks (**categorization**)

> Implement generic learning rules (**gradient descent**)

> Map to brain data. (**ventral stream**)

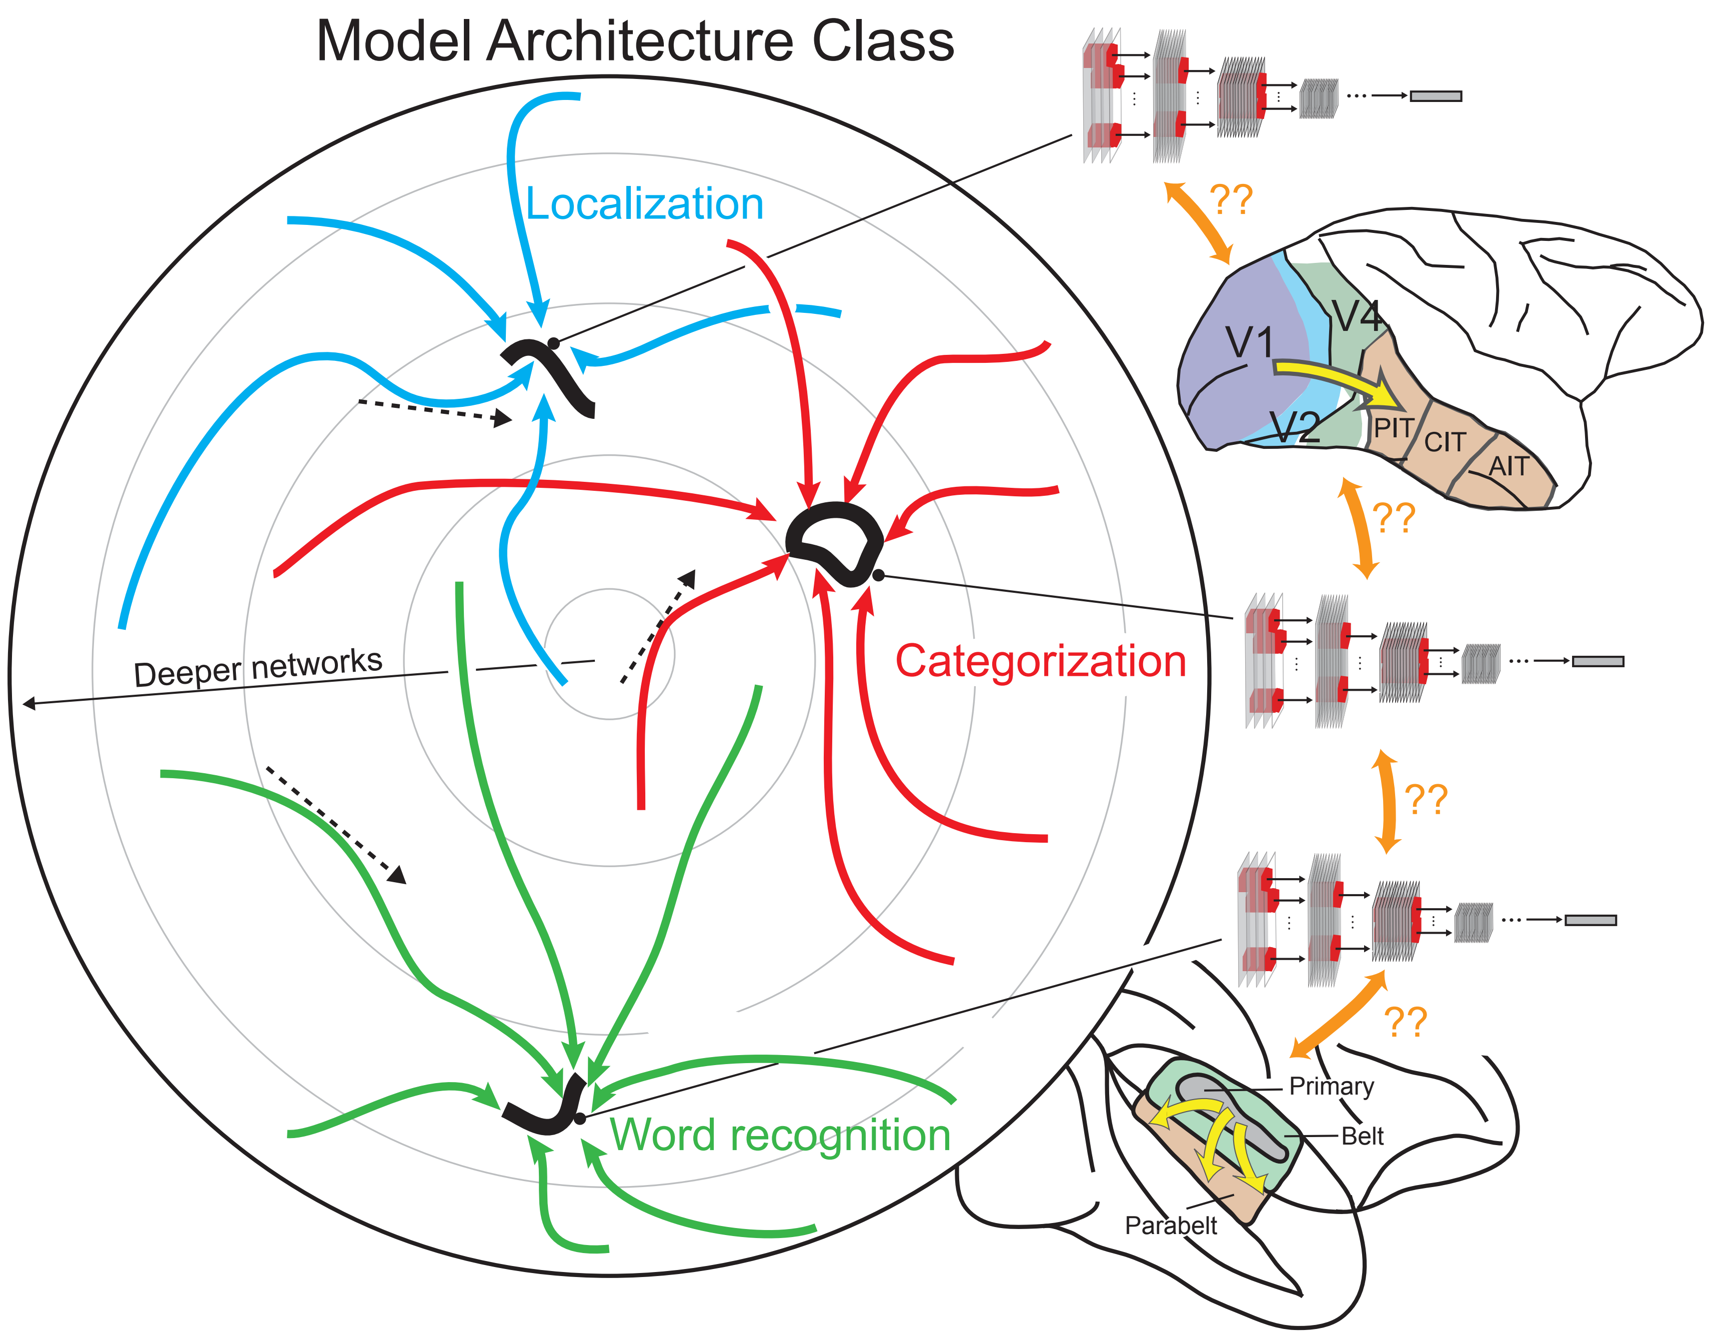

Yamins & DiCarlo.
***Nat. Neuro.** (2016)*

Four Principles of Goal-Driven Modeling

1.

A = *architecture class*

2.

T = *task/objective*

3.

D = *dataset*

4.

L = *learning rule*

Four Principles of Goal-Driven Modeling

1.

A = *architecture class*

2.

T = *task/objective*

3.

D = *dataset*

4.

L = *learning rule*

Best proxies thus far for ventral stream:

A = *ConvNets of reasonable depth*

T = *multi-way object categorization*

D = *ImageNet images*

L = *evolutionary architecture search +*
filter learning through gradient descent

Four Principles of Goal-Driven Modeling

1.

A = *architecture class* ***= circuit neuro-anatomy***

2.

T = *task/objective* ***= ecological niche***

3.

D = *dataset* ***= environment***

4.

L = *learning rule* ***= natural selection + synaptic plasticity***

Best proxies thus far for ventral stream:

A = *ConvNets of reasonable depth*

T = *multi-way object categorization*

D = *ImageNet images*

L = *evolutionary architecture search + filter learning through gradient descent*

Four Principles of Goal-Driven Modeling

1.

A = *architecture class* **= *circuit neuro-anatomy***

solving

2.

T = *task/objective* **= *ecological niche***

situated in

3.

D = *dataset* **= *environment***

updating according to

4.

L = *learning rule* **= *natural selection* + *synaptic plasticity***

Best proxies thus far for ventral stream:

A = *ConvNets of reasonable depth*

T = *multi-way object categorization*

D = *ImageNet images*

L = *evolutionary architecture search + filter learning through gradient descent*

"Nothing in biology makes sense except in light of evolution"

Theo Dobzhansky

"Nothing in biology makes sense except in light of evolution"

Theo Dobzhansky

"Nothing in neuroscience makes sense except in light of behavior"

Gordon Shepherd

"Nothing in biology makes sense except in light of evolution"

Theo Dobzhansky

"Nothing in neuroscience makes sense except in light of behavior"

Gordon Shepherd

<u>computational</u>

Nothing in neuroscience makes sense except in light of ***optimization***.

Stanford Neuro AI Lab

"Nothing in biology makes sense except in light of evolution"

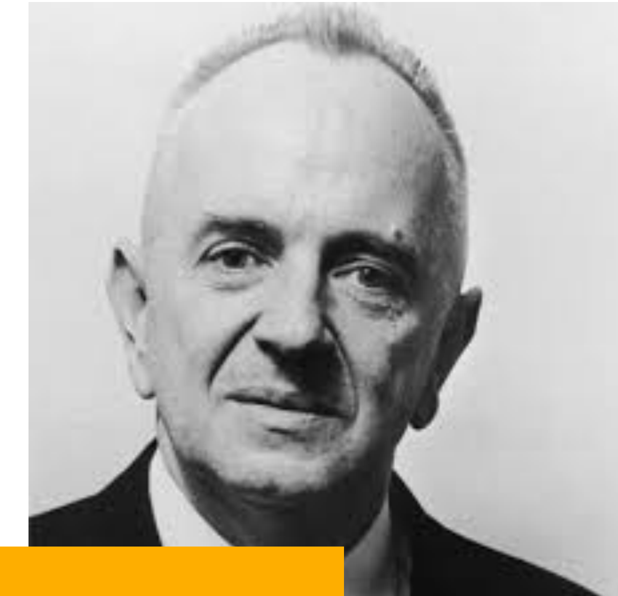

Dobzhansky

Restated:

Behavior is highly constraining of the brain.

"Nothing

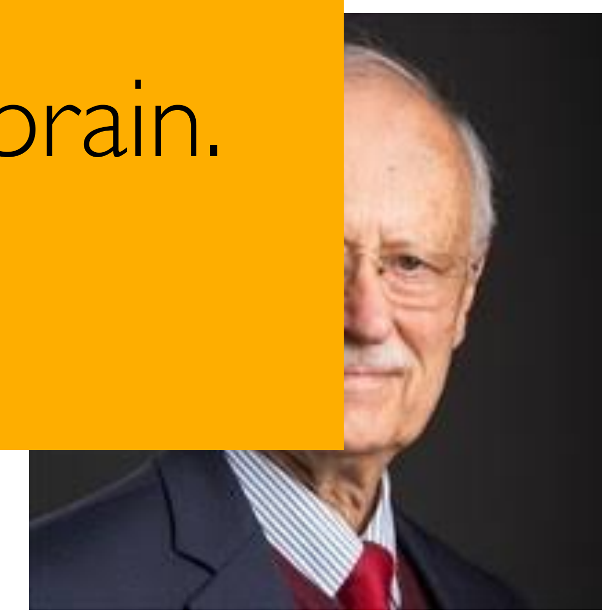

Gordon Shepherd

Nothing in computational neuroscience makes sense except in light of ***optimization***.

Stanford Neuro AI Lab

~~Principle~~ of "Goal-Driven Modeling"

Heuristic of "Goal-Driven Modeling"

~~Principle~~ of "Goal-Driven Modeling"

Heuristic of "Goal-Driven Modeling"

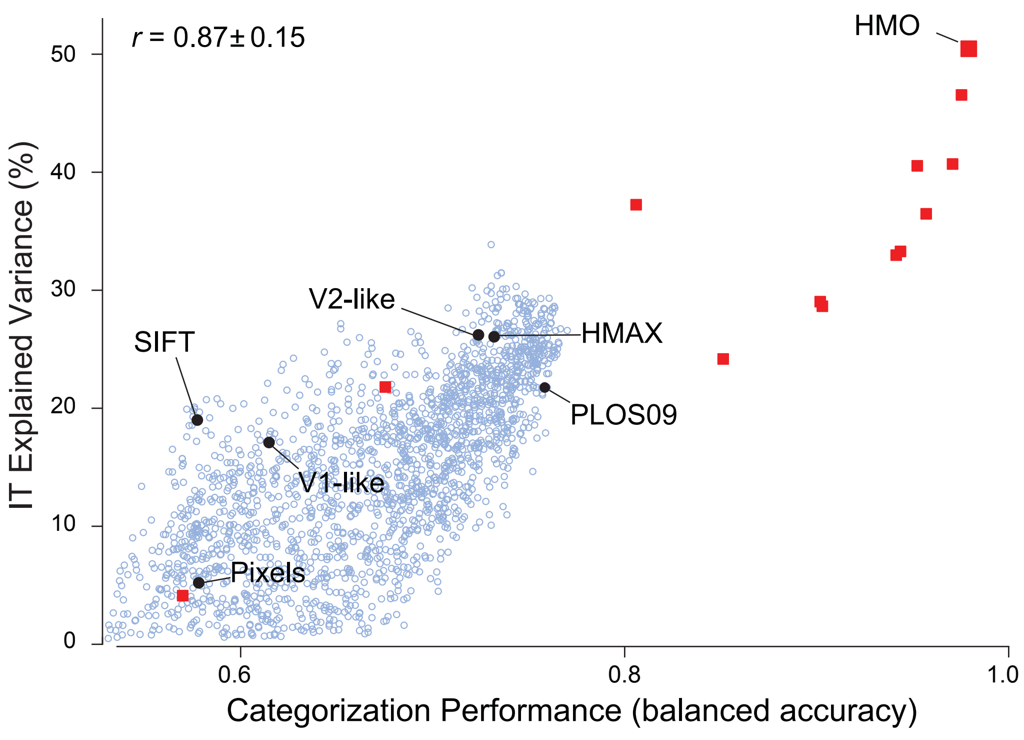

res-net?

… after all at some point, for any given task, you'll probably "go over the hump" … perhaps when you exceed human performance or overfit on that task

Can we go beyond vision?

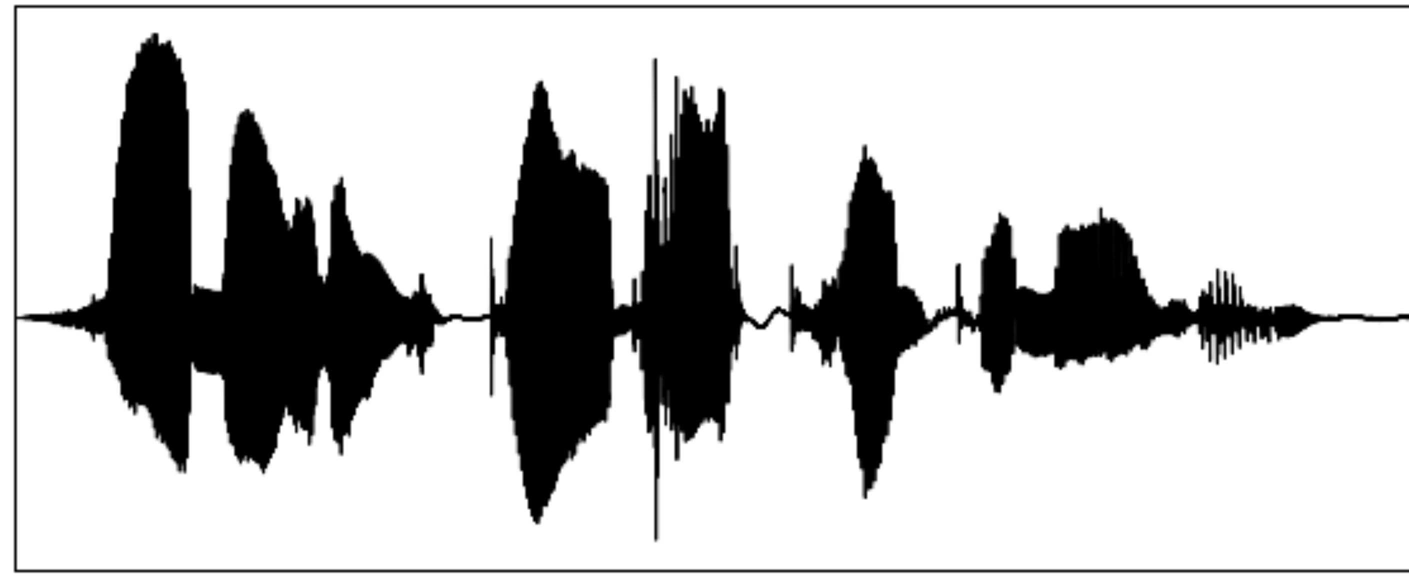

visual cortex

auditory cortex

"Mercedes behind Lamborghini, on a field in front of mountains."

"Hannah is good at compromising"

Can we go beyond vision?

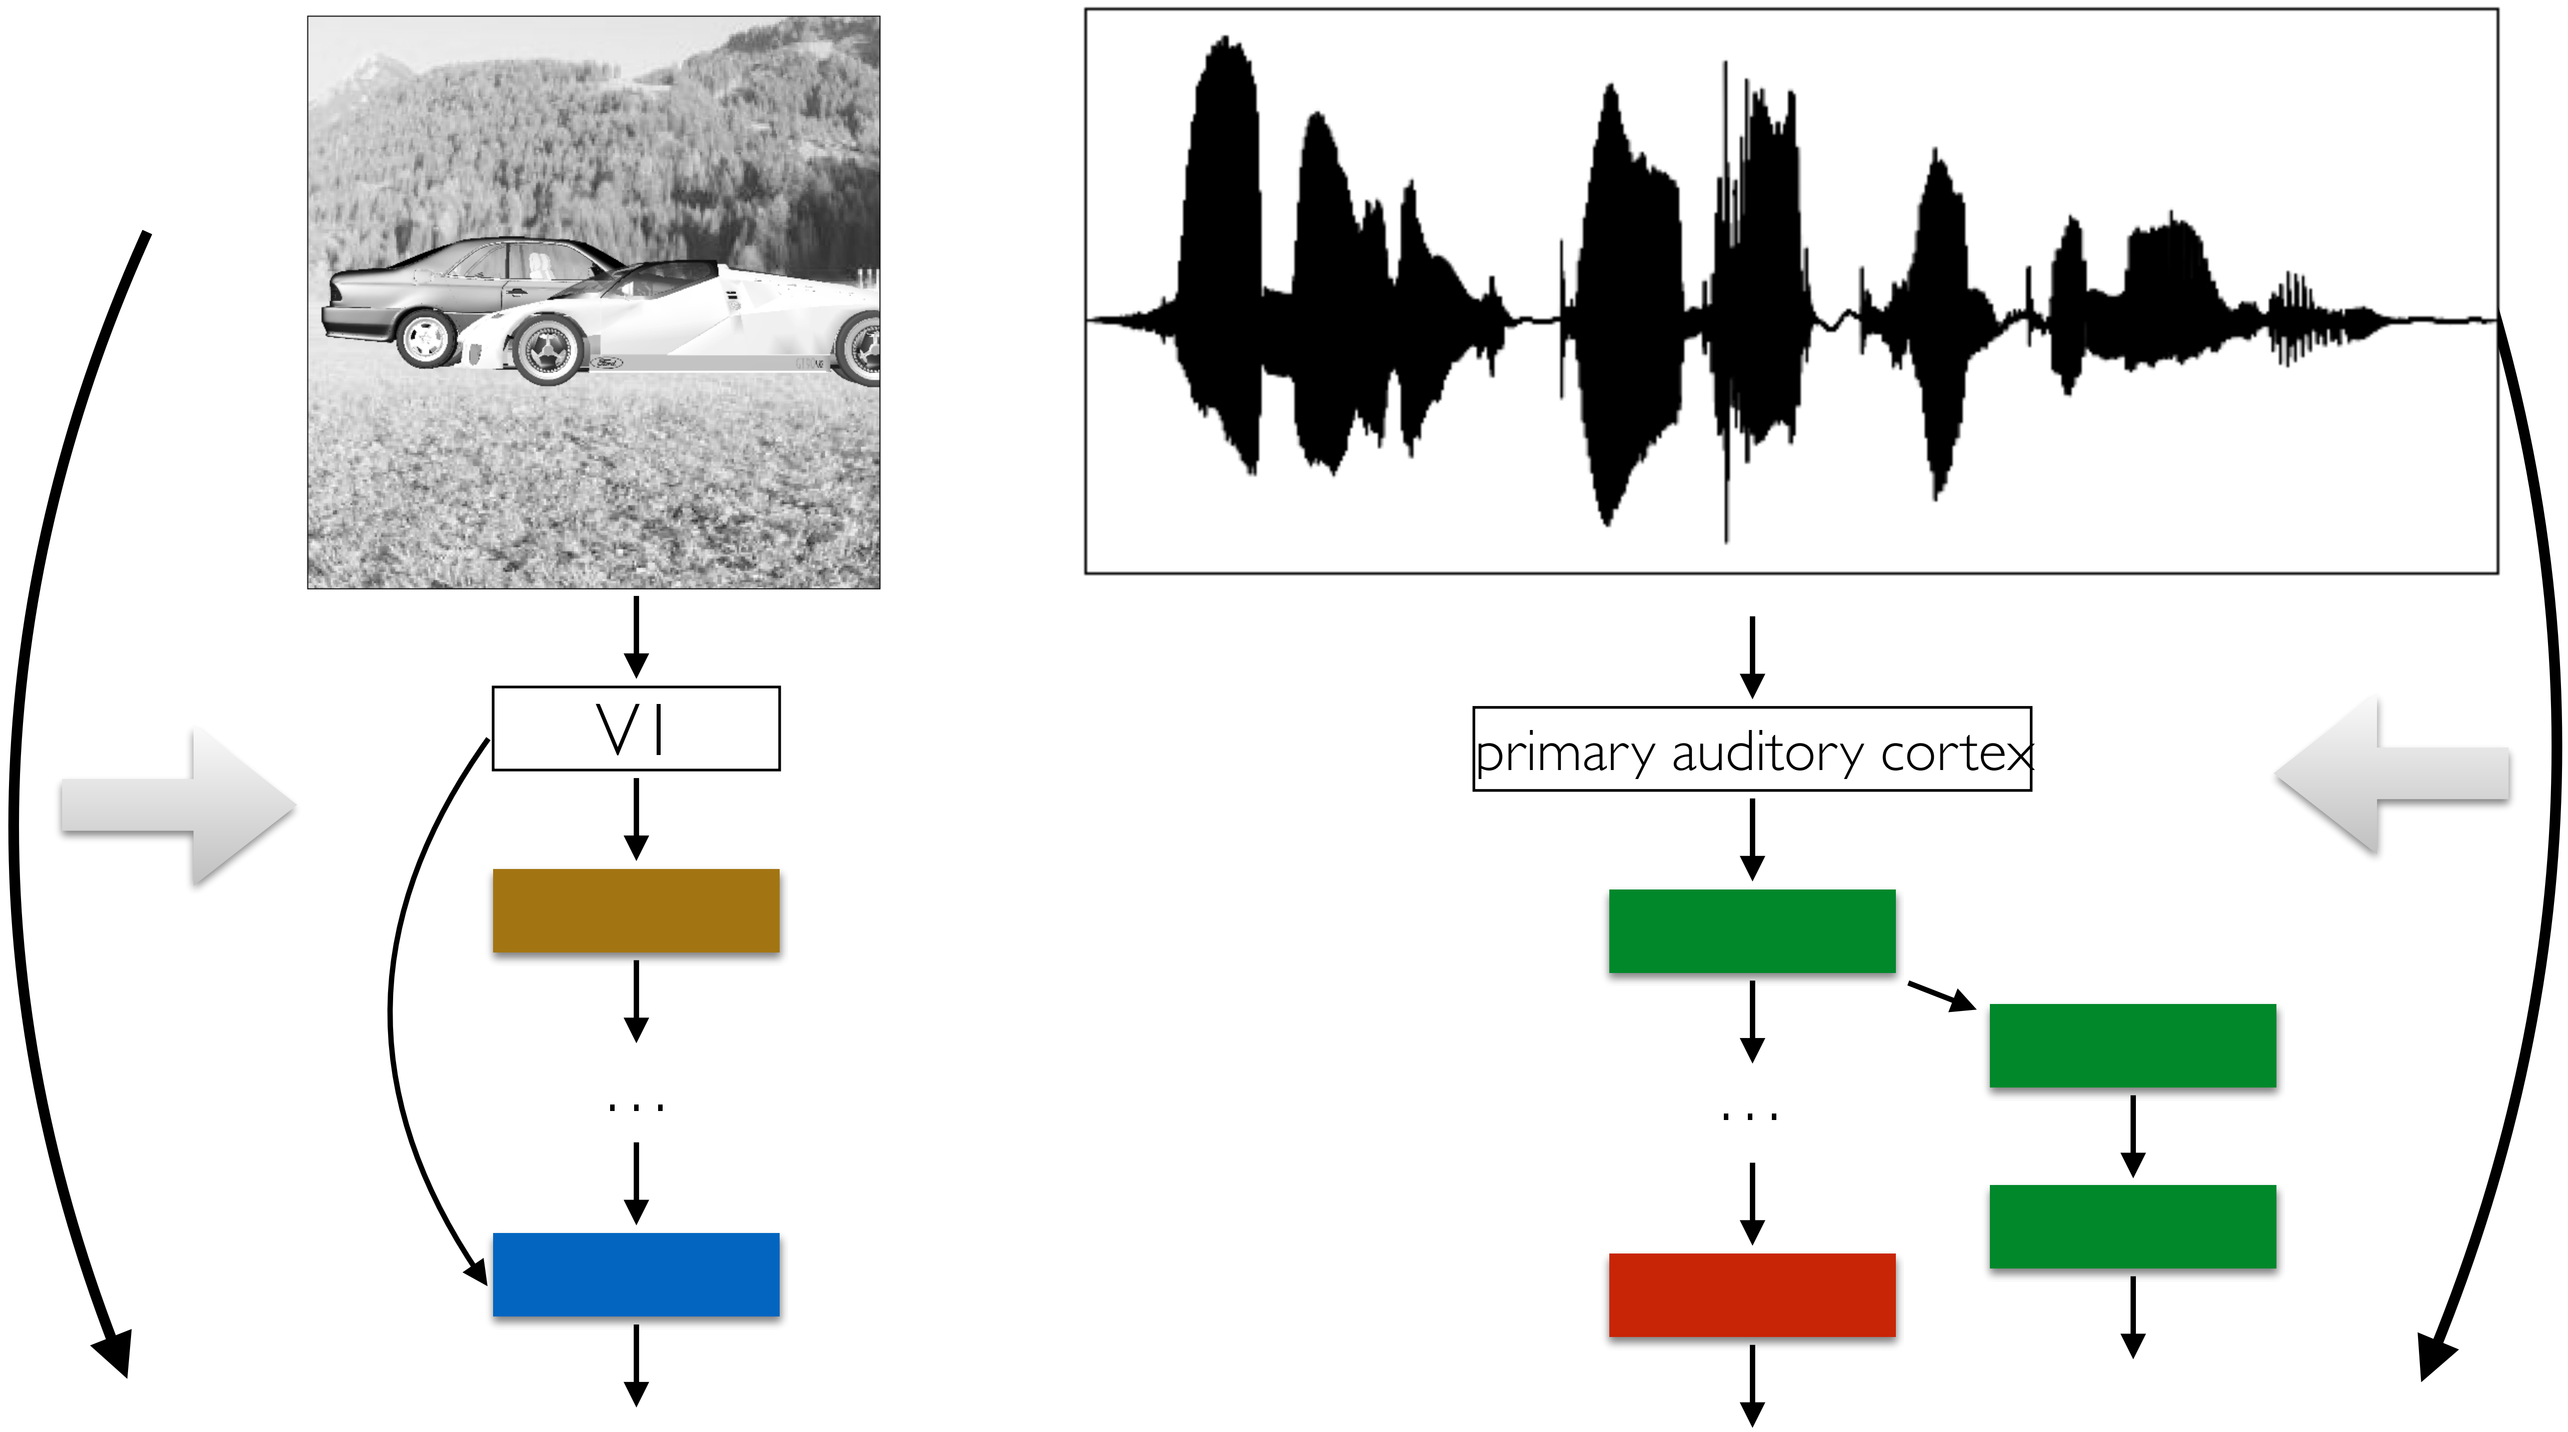

"Mercedes behind Lamborghini, on a field in front of mountains."

"Hannah is good at compromising"

> Formulate comprehensive model class (**RNNs**)

> Choose challenging, ethologically-valid tasks (**task switching/ memory**)

> Implement generic learning rules (**??**)

> Map to brain data. (**Parietal cortex, PFC**)

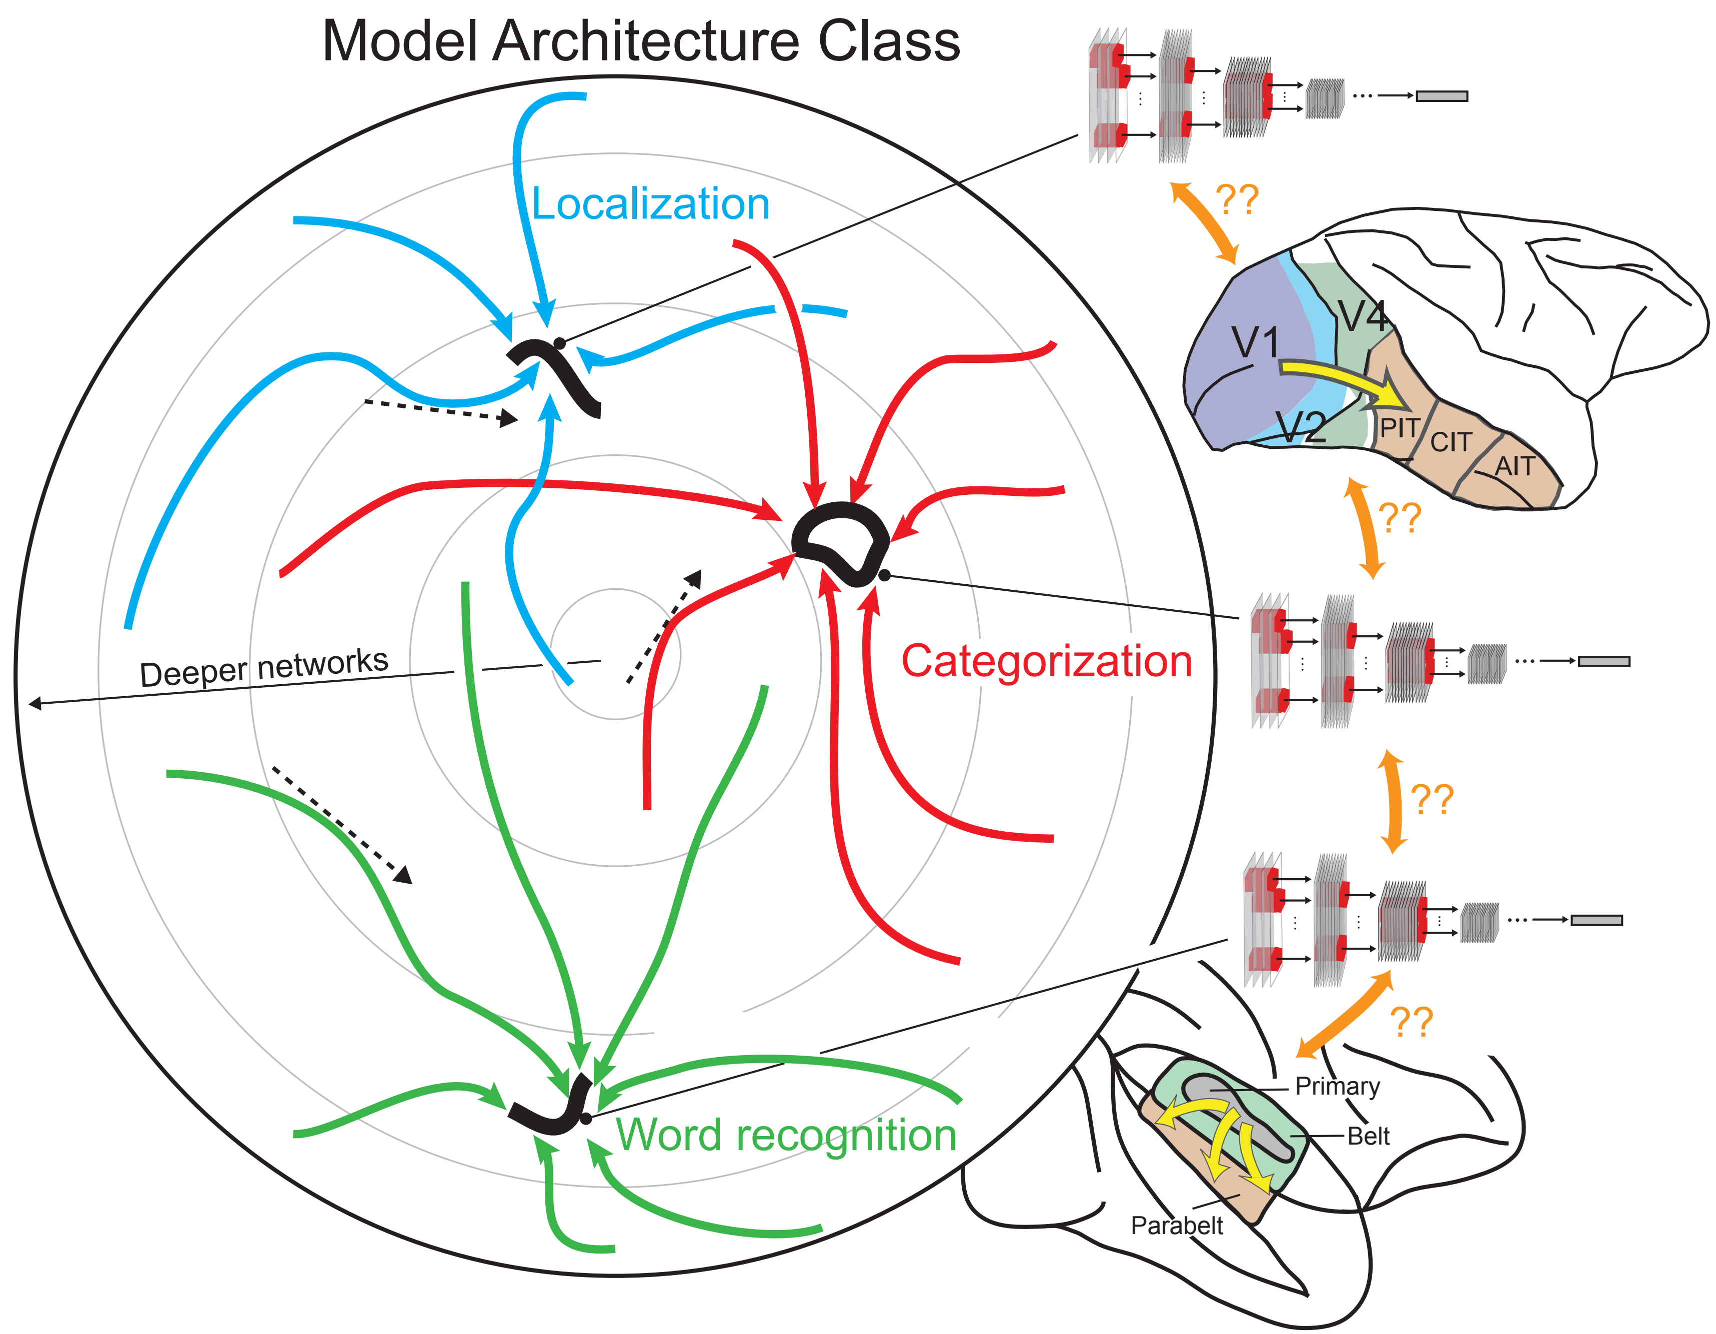

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

Big Problems in Each Area

***bad** = obviously deeply wrong as model of the brain or behavior

1. ✗ ***bad***

***A** = architecture class*

e.g. **CNNs**

2.

***T** = task/objective*

e.g. **Object Categorization**

3.

***D** = dataset*

e.g. **ImageNet**

4.

***L** = learning rule*

e.g. **Arch. Srch. + Grad. Desc.**

PROBLEM

Big Problems in Each Area

***bad** = obviously deeply wrong as model of the brain or behavior

1. ✗ ***bad***

A = *architecture class*

e.g. **CNNs**

2.

T = *task/objective*

e.g. **Object Categorization**

3.

D = *dataset*

e.g. **ImageNet**

4.

L = *learning rule*

e.g. **Arch. Srch. + Grad. Desc.**

PROBLEM

RECURRENCE and FEEDBACK!!?

Big Problems in Each Area

***bad** = obviously deeply wrong as model of the brain or behavior

PROBLEM

1. ✗ ***bad***

A = *architecture class*

e.g. **CNNs**

RECURRENCE and FEEDBACK!!?

2. ✗ ***bad***

T = *task/objective*

e.g. **Object Categorization**

TOO MUCH LABELLED DATA REQUIRED!!?

3.

D = *dataset*

e.g. **ImageNet**

4.

L = *learning rule*

e.g. **Arch. Srch. + Grad. Desc.**

Big Problems in Each Area

***bad** = obviously deeply wrong as model of the brain or behavior

1. ❌***bad***

A = *architecture class*

e.g. **CNNs**

2. ❌***bad***

T = *task/objective*

e.g. **Object Categorization**

3. ❌***bad***

D = *dataset*

e.g. **ImageNet**

4.

L = *learning rule*

e.g. **Arch. Srch.** + **Grad. Desc.**

PROBLEM

RECURRENCE and FEEDBACK!!?

TOO MUCH LABELLED DATA REQUIRED!!?

REAL NOISY VIDEO DATASTREAMS vs STEREOTYPED CLEAN STILL IMAGES

Big Problems in Each Area

***bad** = obviously deeply wrong as model of the brain or behavior

1. **bad**

A = *architecture class*

e.g. **CNNs**

2. **bad**

T = *task/objective*

e.g. **Object Categorization**

3. **bad**

D = *dataset*

e.g. **ImageNet**

4. **bad**

L = *learning rule*

e.g. **Arch. Srch.** + **Grad. Desc.**

PROBLEM

RECURRENCE and FEEDBACK!!?

TOO MUCH LABELLED DATA REQUIRED!!?

REAL NOISY VIDEO DATASTREAMS vs STEREOTYPED CLEAN STILL IMAGES

BACKPROP AND ITS DISCONTENTS

So far, we've done the basic idea

Date	Session
01/06	Introduction to NeuroAI
01/08	Visual Systems Neuroscience Background
01/13	DNN Models of the Visual System I
01/15	DNN Models of the Visual System II
01/20	[NO CLASS-MLK DAY]
01/22	Recurrent Models in Vision and Beyond
01/27	**Guest Lecture — Meenakshi Khosla (USCD):** *Mapping Neural Networks to the Brain*
01/28	
01/29	Unsupervised Learning and the Brain
02/03	**Guest Lecture — Arash Afraz (NIH):** *Model-Driven Brain Perturbation*
02/05	Auditory and Somatosensory Models
02/10	**Guest Lecture — Rhodri Cusack (Trinity):** *Models of Development and Learning*
02/11	
02/12	**Guest Lecture — Josh McDermott (MIT):** *Leveraging Models of Auditory Cortex*
02/17	[NO CLASS-PRESIDENT'S DAY]
02/19	Learning Rules in the Brain
02/24	Models of the Motor System
02/25	
02/26	**Guest Lecture — Scott Linderman (Stanford):** *Dynamical Systems Models in Neuroscience*
03/03	**Guest Lecture — Greta Tuckute (MIT):** *The Human Language Network & LLMs*
03/05	The Hippocampus: Memory and Spatial Navigation
03/10	Topographic Models: A Unified Theory of the Brain
03/12	**Guest Lecture — Robert Hawkins (Stanford):** *Cognitive Modeling*

Basic idea

Next we'll fix some of the problems . . .

Date	Session
01/06	Introduction to NeuroAI
01/08	Visual Systems Neuroscience Background
01/13	DNN Models of the Visual System I
01/15	DNN Models of the Visual System II
01/20	[NO CLASS-MLK DAY]
01/22	Recurrent Models in Vision and Beyond
01/27	**Guest Lecture — Meenakshi Khosla (USCD):** *Mapping Neural Networks to the Brain*
01/28	
01/29	Unsupervised Learning and the Brain
02/03	**Guest Lecture — Arash Afraz (NIH):** *Model-Driven Brain Perturbation*
02/05	Auditory and Somatosensory Models
02/10	**Guest Lecture — Rhodri Cusack (Trinity):** *Models of Development and Learning*
02/11	
02/12	**Guest Lecture — Josh McDermott (MIT):** *Leveraging Models of Auditory Cortex*
02/17	[NO CLASS-PRESIDENT'S DAY]
02/19	Learning Rules in the Brain
02/24	Models of the Motor System
02/25	
02/26	**Guest Lecture — Scott Linderman (Stanford):** *Dynamical Systems Models in Neuroscience*
03/03	**Guest Lecture — Greta Tuckute (MIT):** *The Human Language Network & LLMs*
03/05	The Hippocampus: Memory and Spatial Navigation
03/10	Topographic Models: A Unified Theory of the Brain
03/12	**Guest Lecture — Robert Hawkins (Stanford):** *Cognitive Modeling*

Basic idea

Fixing problems

. . . and then go beyond vision.

Date	Session
01/06	Introduction to NeuroAI
01/08	Visual Systems Neuroscience Background
01/13	DNN Models of the Visual System I
01/15	DNN Models of the Visual System II
01/20	[NO CLASS-MLK DAY]
01/22	Recurrent Models in Vision and Beyond
01/27	**Guest Lecture — Meenakshi Khosla (USCD):** *Mapping Neural Networks to the Brain*
01/28	
01/29	Unsupervised Learning and the Brain
02/03	**Guest Lecture — Arash Afraz (NIH):** *Model-Driven Brain Perturbation*
02/05	Auditory and Somatosensory Models
02/10	**Guest Lecture — Rhodri Cusack (Trinity):** *Models of Development and Learning*
02/11	
02/12	**Guest Lecture — Josh McDermott (MIT):** *Leveraging Models of Auditory Cortex*
02/17	[NO CLASS-PRESIDENT'S DAY]
02/19	Learning Rules in the Brain
02/24	Models of the Motor System
02/25	
02/26	**Guest Lecture — Scott Linderman (Stanford):** *Dynamical Systems Models in Neuroscience*
03/03	**Guest Lecture — Greta Tuckute (MIT):** *The Human Language Network & LLMs*
03/05	The Hippocampus: Memory and Spatial Navigation
03/10	Topographic Models: A Unified Theory of the Brain
03/12	**Guest Lecture — Robert Hawkins (Stanford):** *Cognitive Modeling*

Basic idea

Fixing problems

Beyond Vision