9.523/6.861:
Aspects of a Computational Theory of Intelligence, 2018

Shimon Ullman + Tomaso Poggio

Xavier Boix + Andrew Francl + Yen-Ling Kuo

CENTER FOR
Brains
Minds+
Machines

9.523 overview

- Course description/logistic
- Motivations for this course: the greatest problem in science, CBMM, the MIT Quest
- A bit of history: Neuroscience and AI
- CBMM, the Visual Intelligence moonshot
- Module 1
  - Module 1, theory
  - Module 1, eccentricity
  - Module 1, invariance

GRADING:

- Final project 80%
- Class attendance **and participation** 20%

FINAL PROJECTS:
~5 pages, equivalent to ~2 weeks work. We + speakers will suggest some topics. You may also propose a topic that to be approved by us. Teams are encouraged (but no larger than 3 members).

CENTER FOR
Brains
Minds+
Machines

9.523/6.861:
Aspects of a Computational Theory of Intelligence

- This class is part of the interdisciplinary educational effort of the Center for Brains, Minds and Machines at MIT.

- This year, the lectures in the course are organized in order to present and discuss the 4 modules of the new CBMM project (2018-2023), which could be called "Visual Intelligence"

    - Visual Stream (module 1),
    - Memory and Executive Function (module 2),
    - The Cognitive Core (module 3),
    - Symbolic Compositional Models (module 4).

9.523/6.861:
Aspects of a Computational Theory of Intelligence

- Each talk will situate their module in the context of the other ones (we will share a couple of slides used in the first class). We asked each speaker

    - to outline the present research plan for the module for the next 5 years, stressing what is already cast in stone and what is still open
    - to provide the necessary background for the students
    - to discuss -- and stimulate students to think about -- possible interactions between modules.
    - to challenge students to contribute ideas
    - to propose 1-2 projects

- The class is an opportunity for developing plans of future CBMM research by faculty, TAs and students working together in the class and in the final projects.

9.523/6.861:
Aspects of a Computational Theory of Intelligence

- Of course the class is also an opportunity for future projects including master and PhD theses … if you like the combination of neuroscience and CS!

9.523 overview

- Course description/logistic
- Motivations for this course: the greatest problem in science, CBMM, the MIT Quest
- A bit of history: Neuroscience and AI
- CBMM, the Visual Intelligence moonshot
- Module 1
  - Module 1, theory
  - Module 1, eccentricity
  - Module 1, invariance

Motivation for 9.523: the Problem of Intelligence

Of course, intelligence is not solved,
not as as a scientific problem, not as an engineering problem.

Research is needed:

- for the sake of basic science
- for the engineering of tomorrow

Problem of Intelligence is not yet solved

Search Facebook

Building Jarvis

MARK ZUCKERBERG · MONDAY, DECEMBER 19, 2016

My personal challenge for 2016 was to build a simple AI to run my home -- like Jarvis in Iron Man. Within 5-10 years we'll have AI systems that are more accurate than people for each of our senses -- vision, hearing, touch, etc, as well as things like language. It's impressive how powerful the state of the art for these tools is becoming.

At the same time, we are still far off from understanding how learning works. Everything I did this year – natural language, face recognition, speech recognition and so on –are all variants of the same fundamental pattern recognition techniques. I spent about 100 hours building Jarvis this year, **but even if I spent 1,000 more hours, I probably wouldn't be able to build a system that could learn completely new skills on its own -- unless I made some fundamental breakthrough in the state of AI along the way**.

In a way, AI is both closer and farther off than we imagine. AI is closer to being able to do more powerful things than most people expect -- driving cars, curing diseases, ... Those will each have a great impact on the world, but **we're still figuring out what real intelligence is.**

CBMM's focus is the Science and the Engineering of Intelligence

We aim to make progress in understanding intelligence, that is in understanding how the brain makes the mind, how the brain works and how to build intelligent machines. We believe that the science of intelligence will enable better engineering of intelligence.

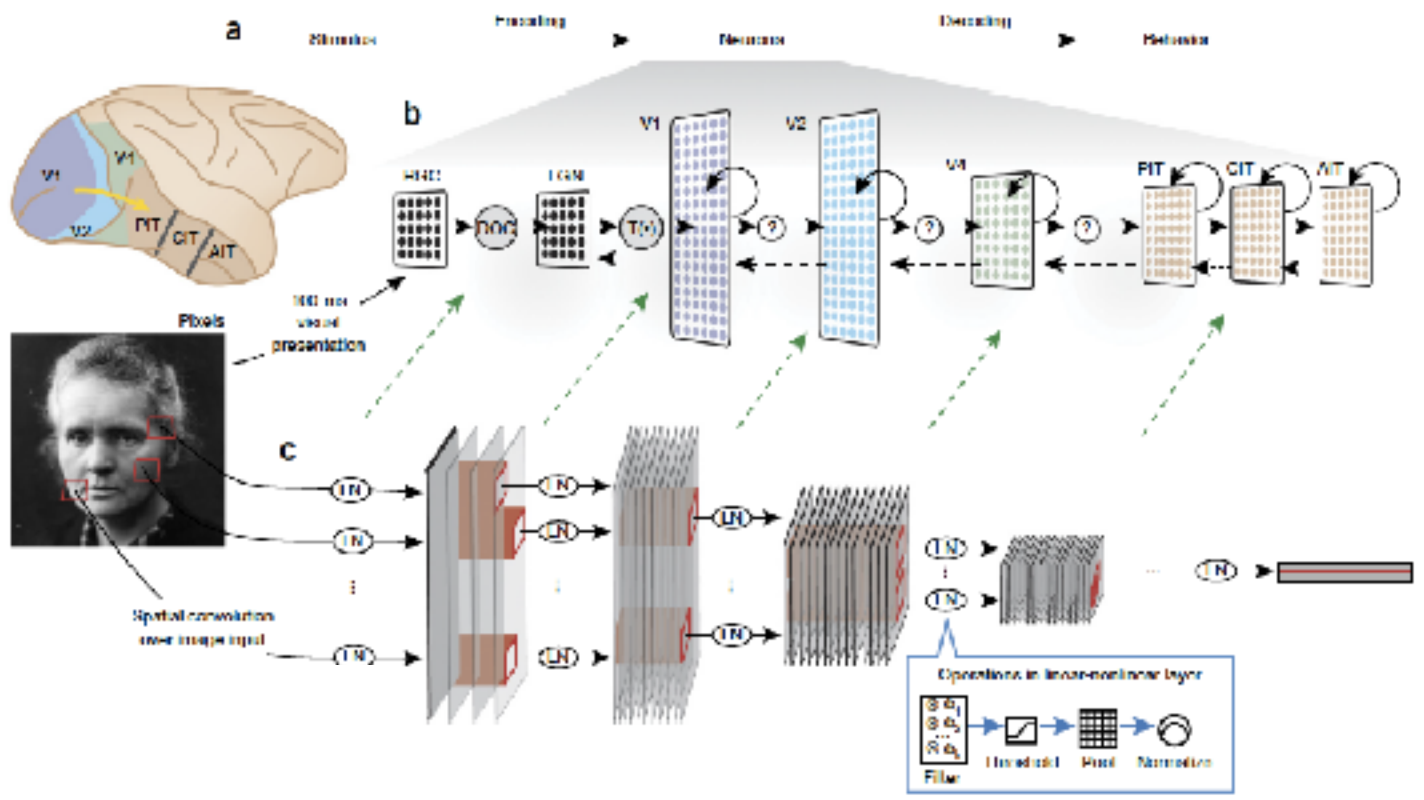

CENTER FOR Brains Minds+ Machines

CBMM approach to the problems of Intelligence

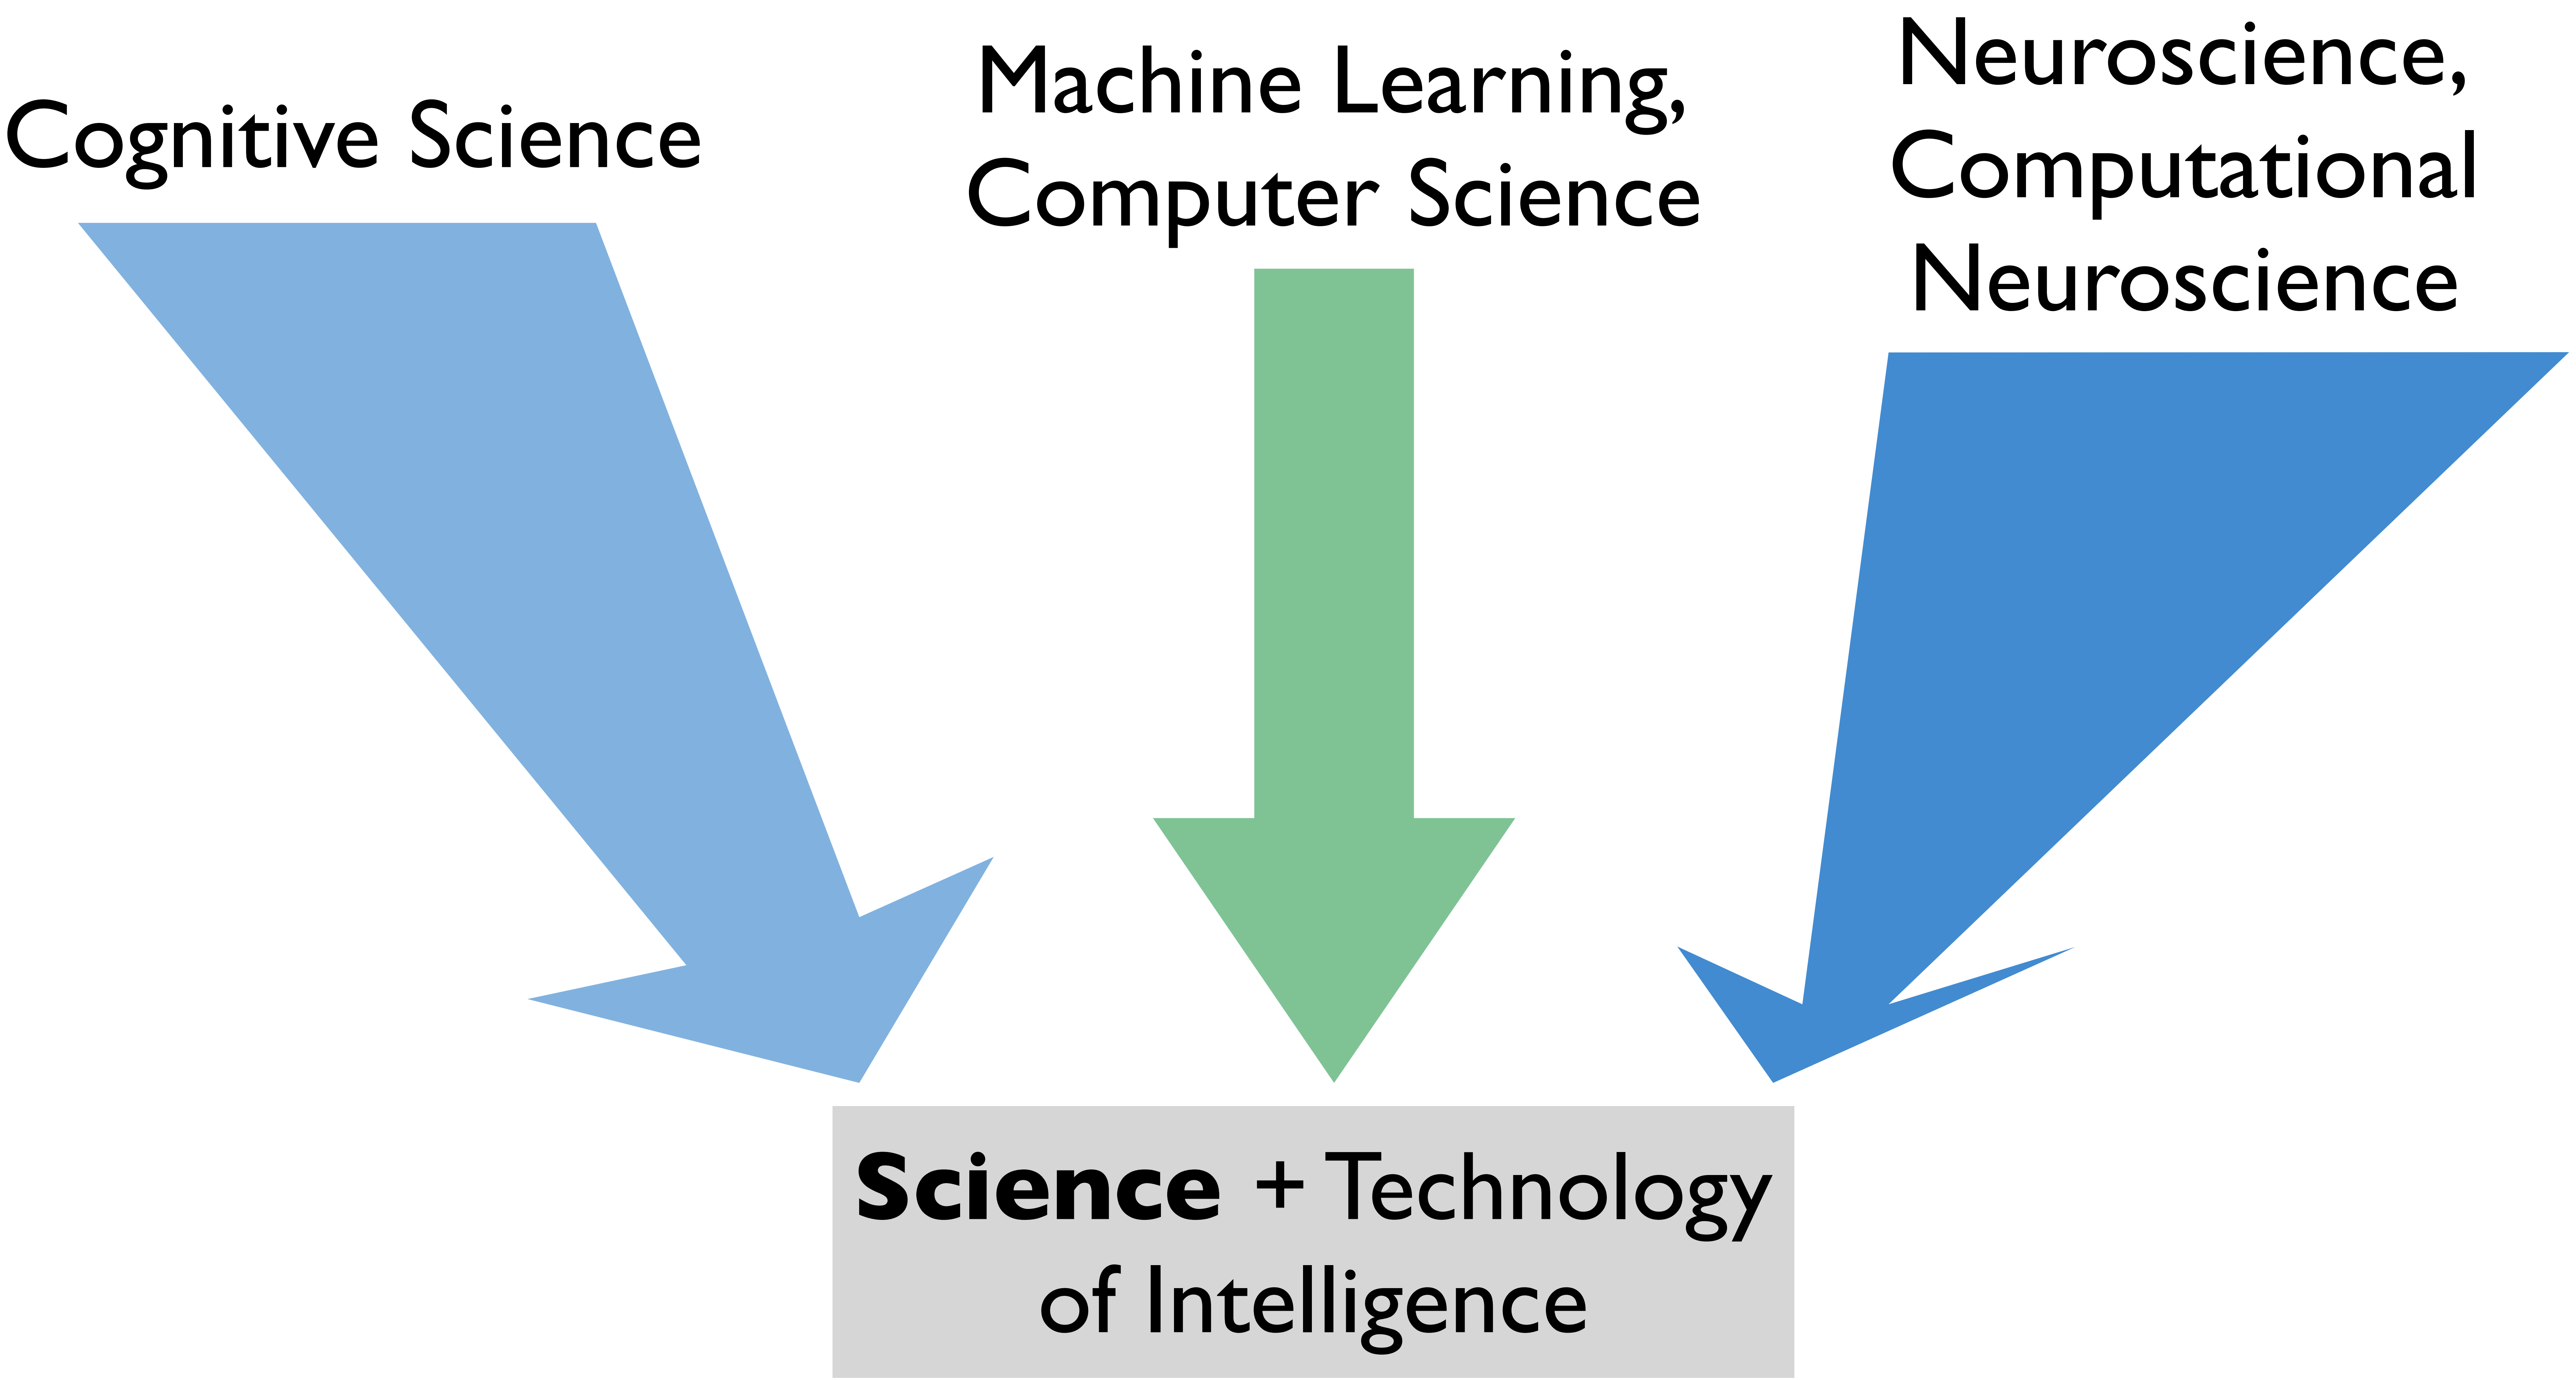

CENTER FOR Brains Minds+ Machines

Research, Education & Diversity Partners

MIT
Boyden, Desimone, DiCarlo, Kanwisher, Katz, McDermott, Poggio, Rosasco, Sassanfar, Saxe, Schulz, Tegmark, Tenenbaum, Ullman, Wilson, Winston

Harvard
Blum, Gershman, Kreiman, Livingstone, Nakayama, Sompolinsky, Spelke

Allen Institute
Koch

Howard U.
Chouika, Manaye, Rwebangira, Salmani

Hunter College
Chodorow, Epstein, Sakas, Zeigler

Johns Hopkins U.
Yuille

Queens College
Brumberg

Rockefeller U.
Freiwald

Stanford U.
Goodman

Universidad Central del Caribe (UCC)
Jorquera

University of Central Florida
McNair Program

UMass Boston
Blaser, Ciaramitaro, Pomplun, Shukla

UPR – Mayagüez
Santiago, Vega-Riveros

UPR– Río Piedras
Garcia-Arraras, Maldonado-Vlaar, Megret, Ordóñez, Ortiz-Zuazaga

Wellesley College
Hildreth, Wiest, Wilmer

Academic and Corporate Partners

IIT	A*star	Hebrew U.	MPI
Cingolani	Chuan Poh Lim	Weiss	Bülthoff

Genoa U.	Weizmann	City U. HK
Verri, Rosasco	Ullman	Smale

Google	IBM	Microsoft	Orcam	Siemens	Philips	GE	
DeepMind	Honda	Boston Dynamics	NVIDIA	Rethink Robotics	Schlumberger	Mobileye	Intel

CENTER FOR Brains Minds+ Machines

Summer Course at Woods Hole: Our flagship initiative led by G. Kreiman

Brains, Minds & Machines Summer Course

An intensive three-week course gives advanced students a “deep” introduction to the problem of intelligence

A community of scholars is being formed:

First reunion of alumni of summer school Aug. 26-27 in Woodshole, MA

CENTER FOR Brains Minds+ Machines

Forging connections between human and machine intelligence research, its applications, and its bearing on society.

The MIT Intelligence Quest will advance the science and engineering of both human and machine intelligence. Launched on February 1, 2018, this effort seeks to discover the foundations of human intelligence and drive the development of technological tools that can positively influence virtually every aspect of society.

The Institute's culture of collaboration will encourage life scientists, computer scientists, social scientists, and engineers to join forces to investigate the societal implications of their work as they pursue hard problems lying beyond the current horizon of intelligence research. By uniting diverse fields and capitalizing on what they can teach each other, we seek to answer the deepest questions about intelligence.

CENTER FOR Brains Minds+ Machines

Intelligence: The MIT Quest

CORE: Cutting-Edge Research on the Science + Engineering of Intelligence

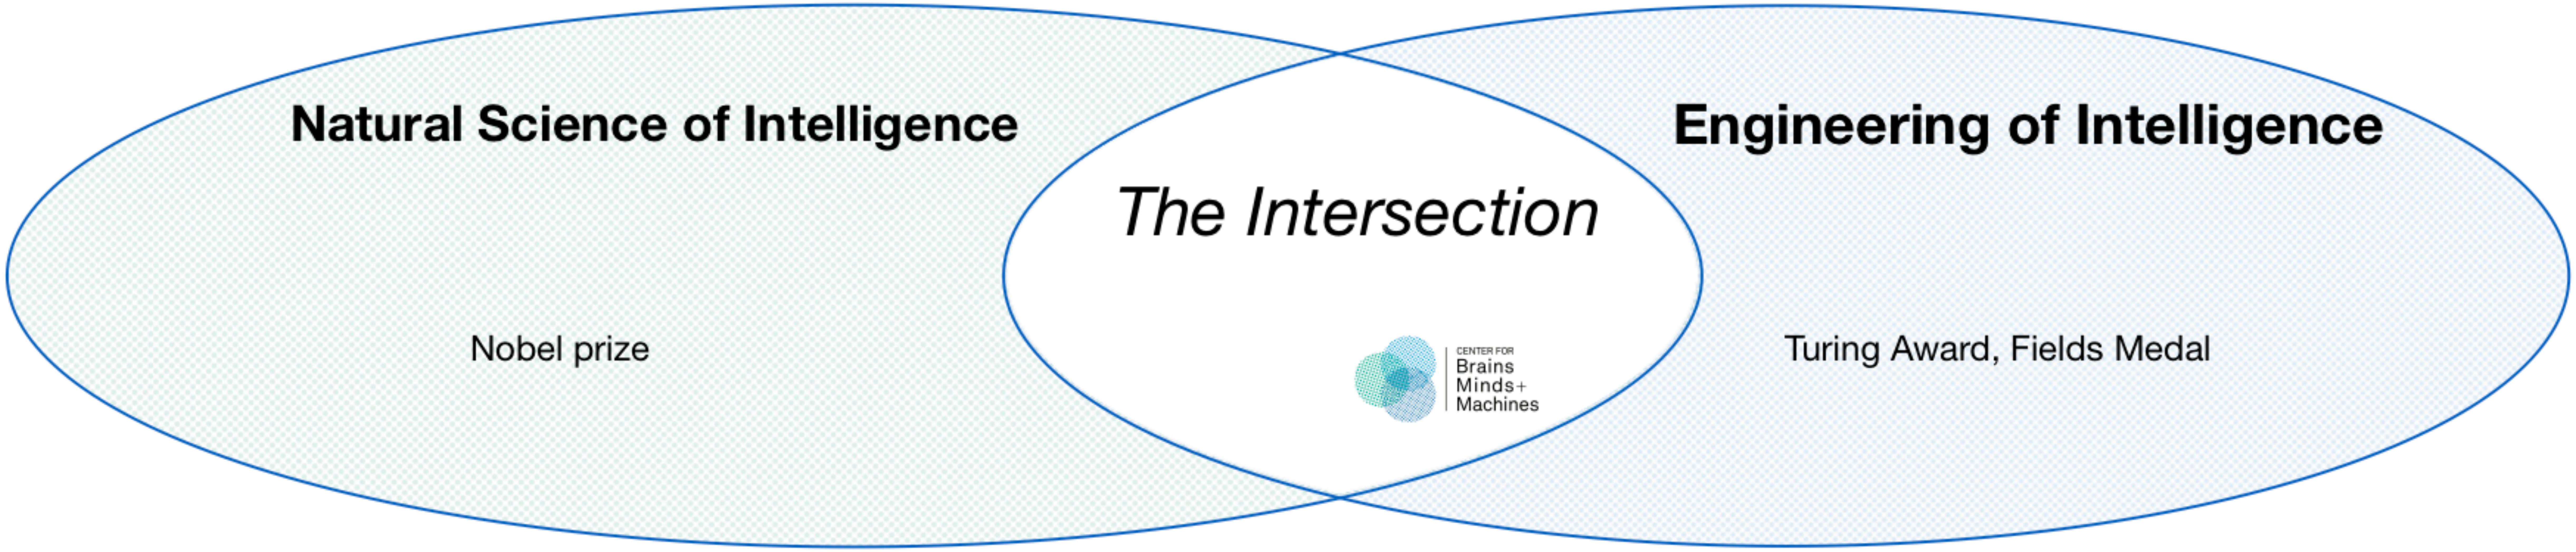

Logical for MIT…

Marvin Minsky

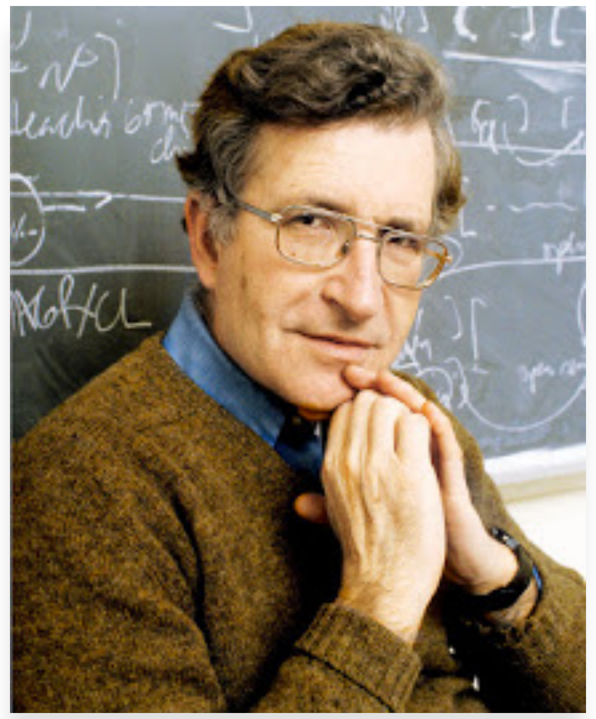

Noam Chomsky

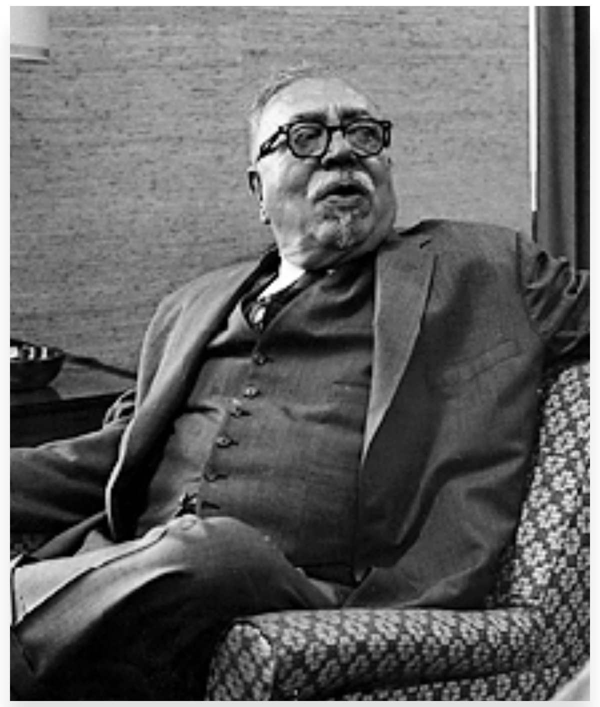

Norbert Wiener

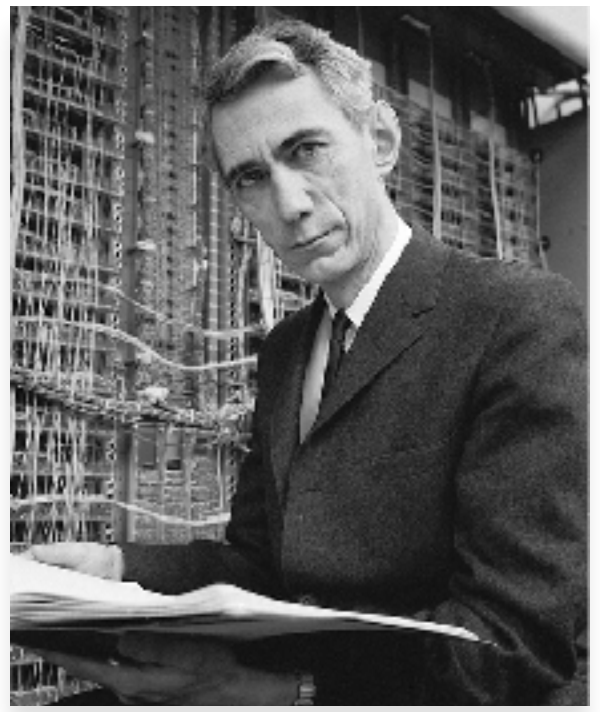

Claude Shannon

“The Golden Age” 1950 - 1970

MIT Intelligence Initiative

CENTER FOR
Brains
Minds+
Machines

Intelligence: The MIT Quest

2008

2012 - 2013

2018

9.523 overview

- Course description/logistic
- Motivations for this course: the greatest problem in science, CBMM, the MIT Quest
- Neuroscience and AI
- CBMM, the Visual Intelligence moonshot
- Module 1
    - Module 1, theory
    - Module 1, eccentricity
    - Module 1, invariance

How do we make significant progress?

Just a definition: I use the word *science* to mean *natural science*

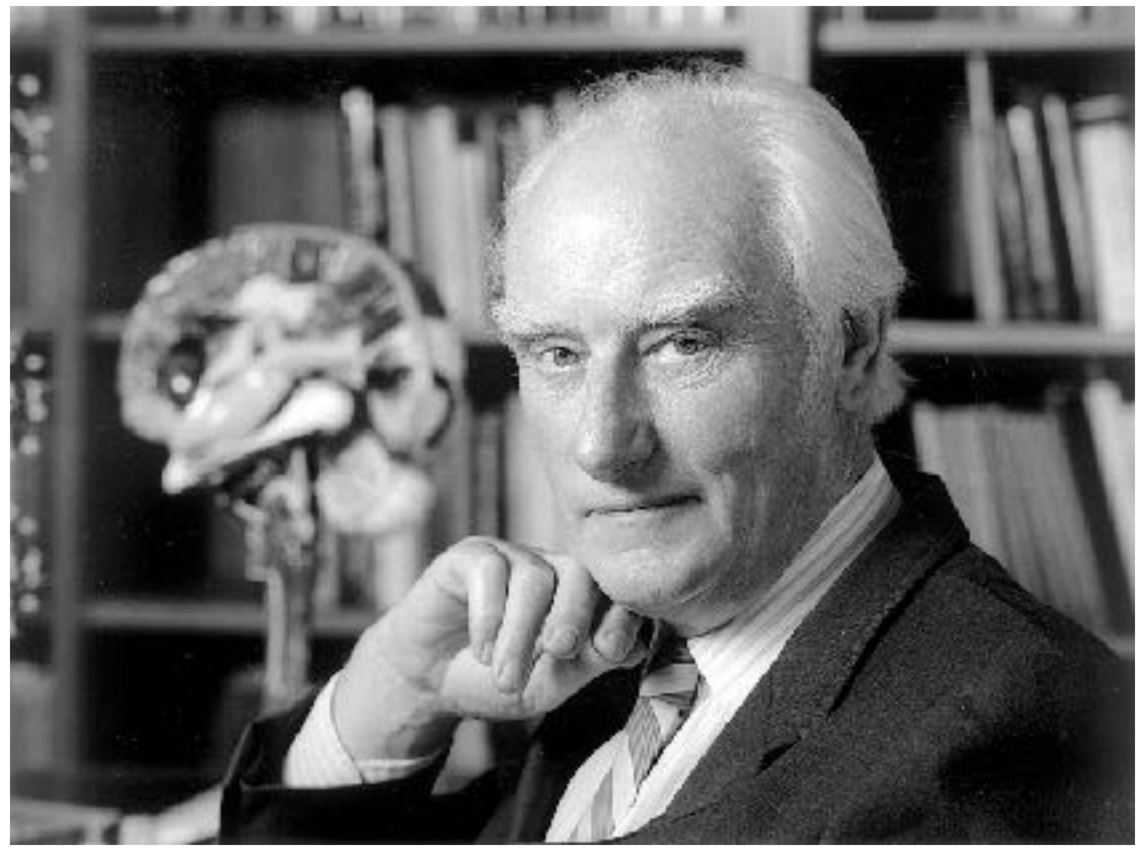

CBMM's focus is the Science and the Engineering of Intelligence

We aim to make progress in understanding intelligence, that is in understanding how the brain makes the mind, how the brain works and how to build intelligent machines. We believe that the science of intelligence will enable better engineering of intelligence.

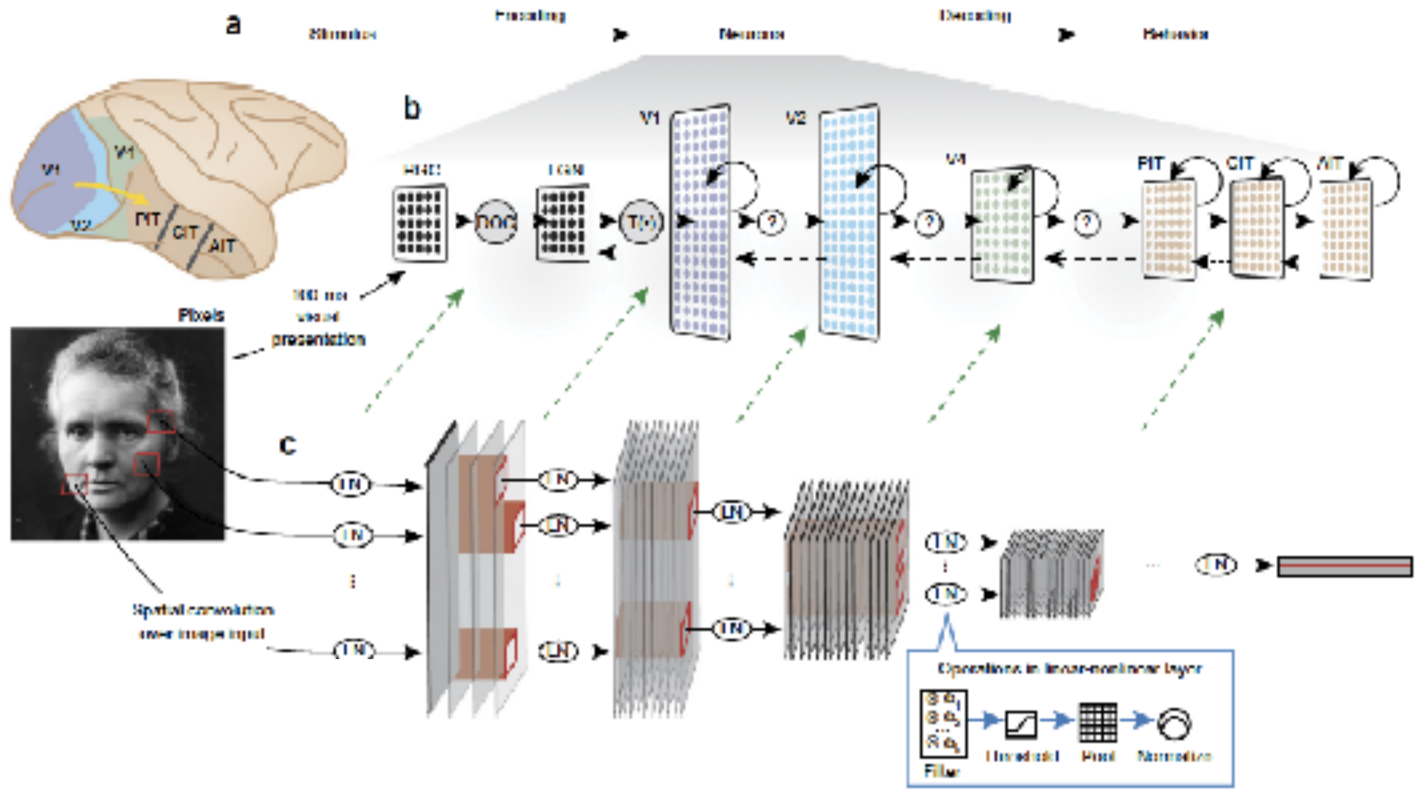

CENTER FOR Brains Minds+ Machines

Neuron
Review

Neuroscience-Inspired Artificial Intelligence

Demis Hassabis,[1,2,*] Dharshan Kumaran,[1,3] Christopher Summerfield,[1,4] and Matthew Botvinick[1,2]
[1]DeepMind, 5 New Street Square, London, UK
[2]Gatsby Computational Neuroscience Unit, 25 Howland Street, London, UK
[3]Institute of Cognitive Neuroscience, University College London, 17 Queen Square, London, UK
[4]Department of Experimental Psychology, University of Oxford, Oxford, UK
*Correspondence: dhcontact@google.com
http://dx.doi.org/10.1016/j.neuron.2017.06.011

The fields of neuroscience and artificial intelligence (AI) have a long and intertwined history. In more recent times, however, communication and collaboration between the two fields has become less commonplace. In this article, we argue that better understanding biological brains could play a vital role in building intelligent machines. We survey historical interactions between the AI and neuroscience fields and emphasize current advances in AI that have been inspired by the study of neural computation in humans and other animals. We conclude by highlighting shared themes that may be key for advancing future research in both fields.

The successful transfer of insights gained from neuroscience to the development of AI algorithms is critically dependent on the interaction between researchers working in both these fields, with insights often developing through a continual handing back and forth of ideas between fields. In the future, we

Two Main Recent Success Stories in AI

DL and RL come from neuroscience

R
L

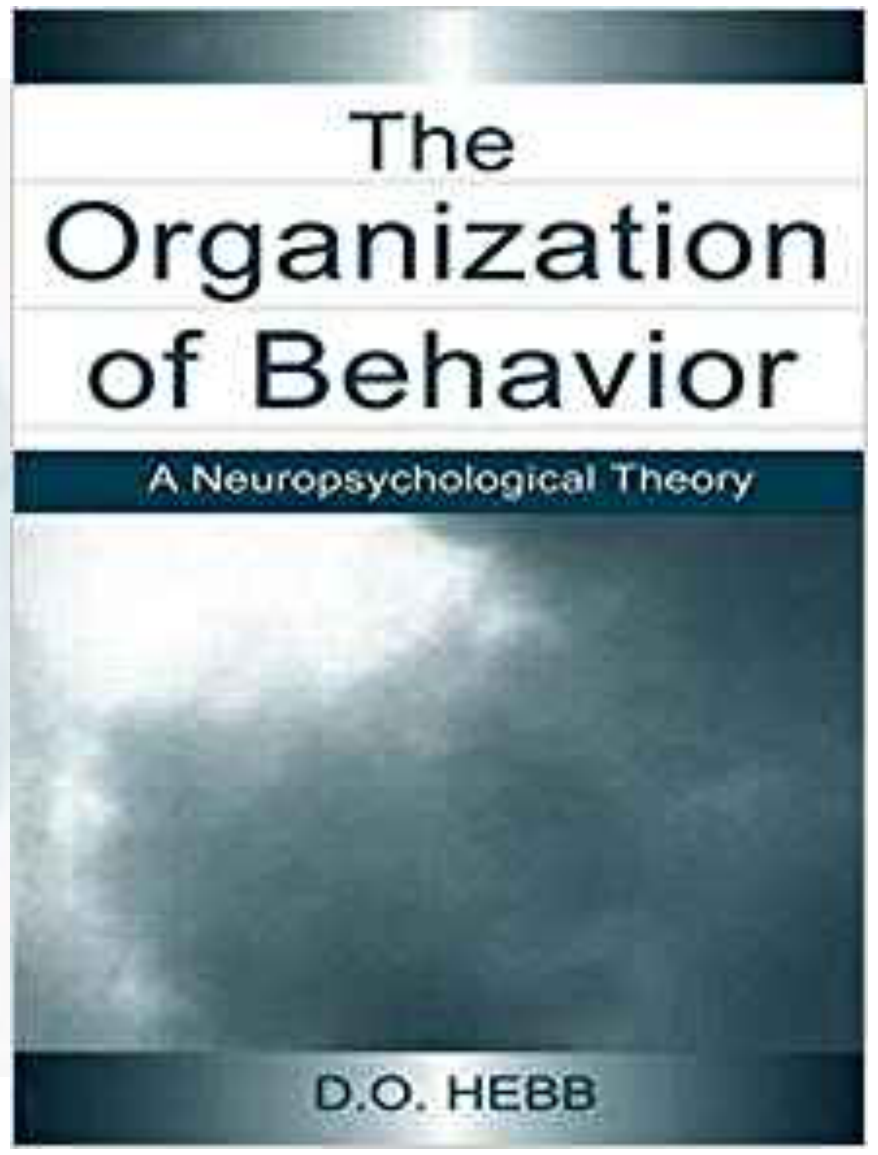

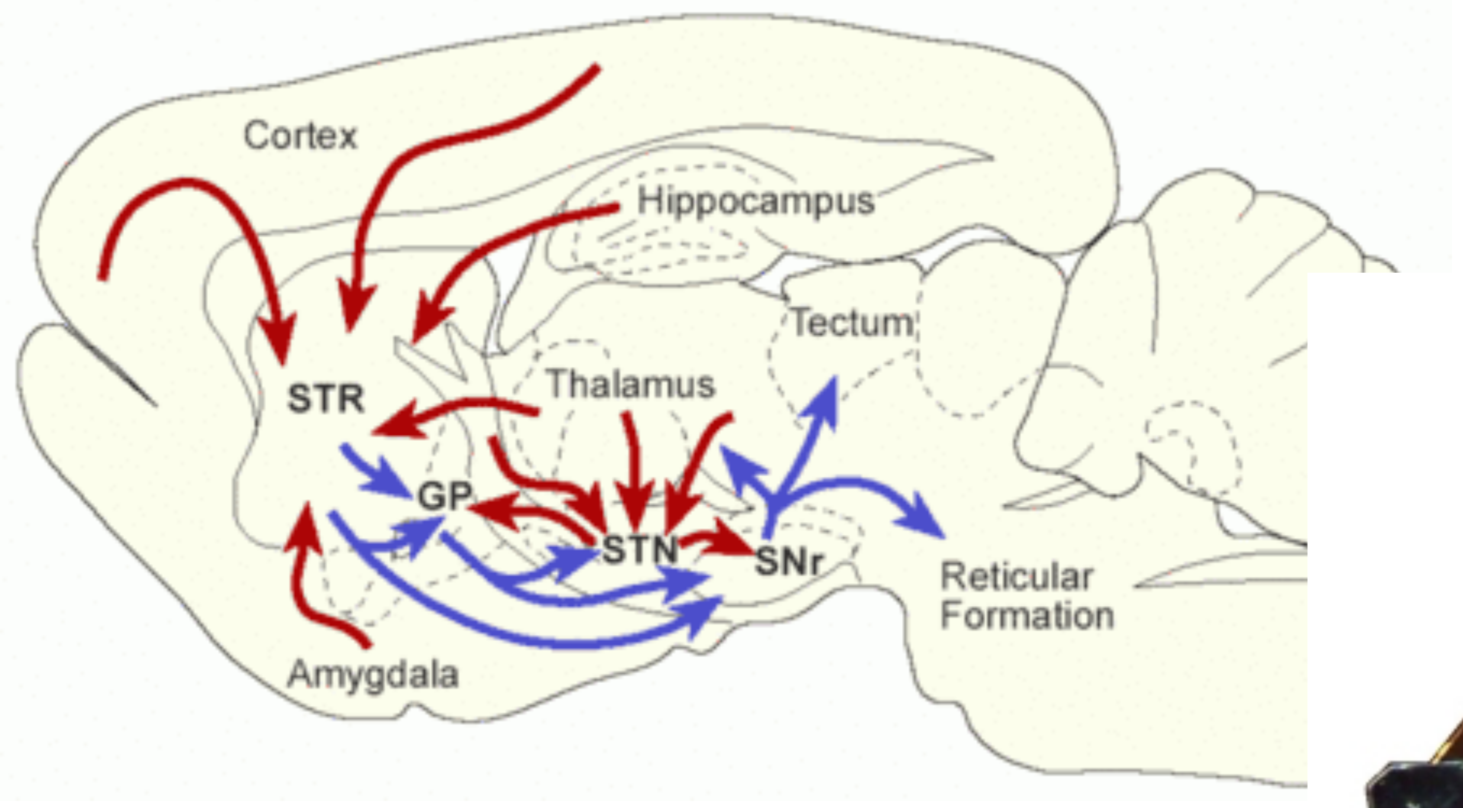

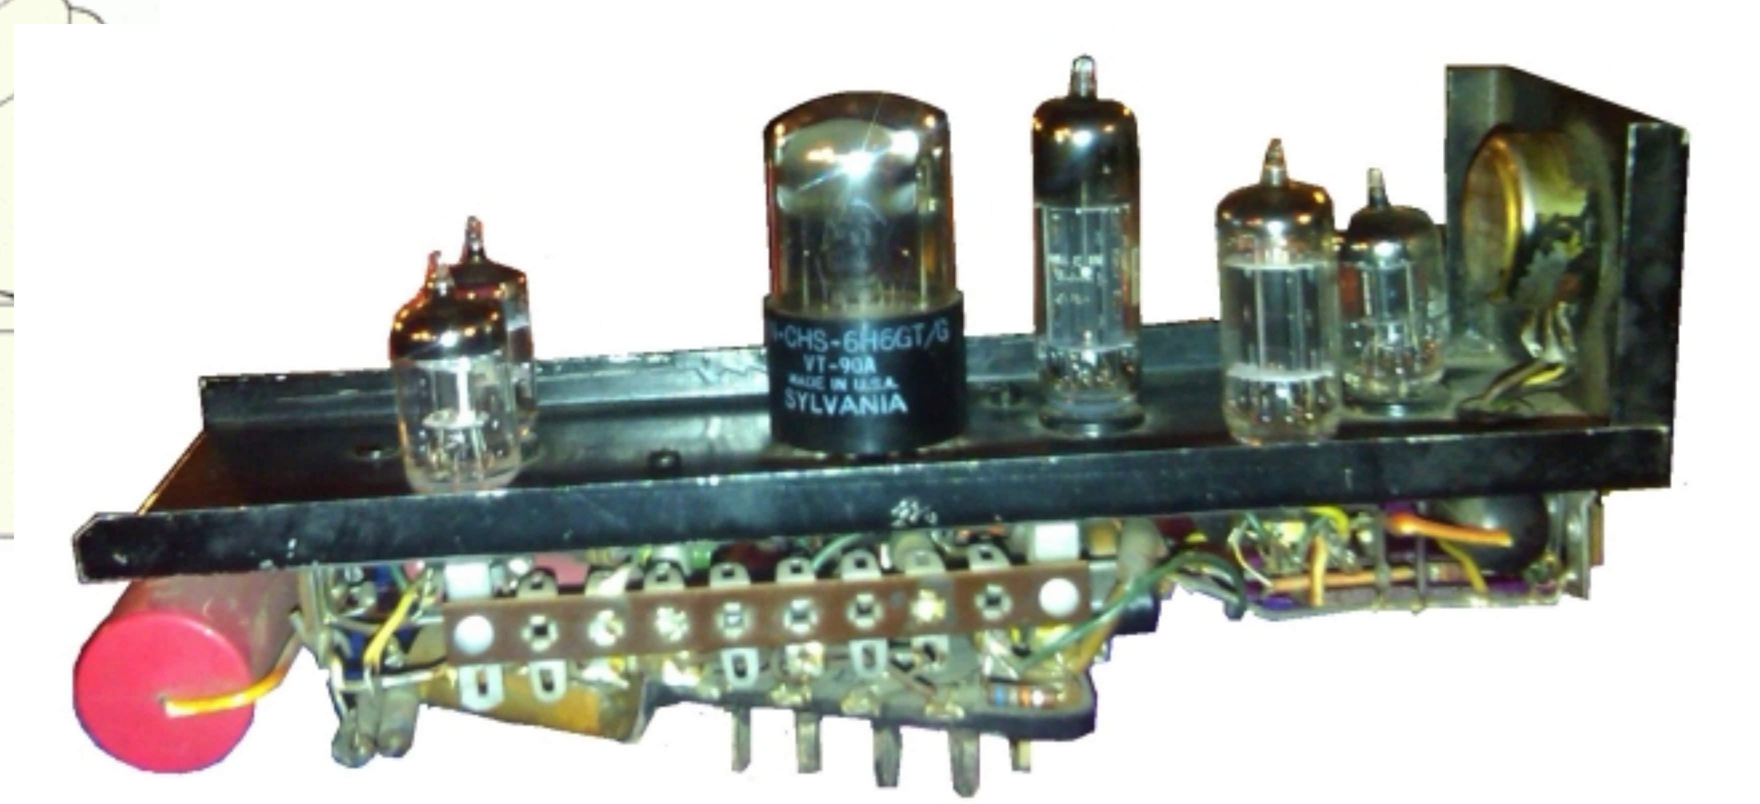

Minsky's SNARC

D
L

RECEPTIVE FIELDS AND FUNCTIONAL ARCHITECTURE IN TWO NONSTRIATE VISUAL AREAS (18 AND 19) OF THE CAT[1]

DAVID H. HUBEL AND TORSTEN N. WIESEL

Neurophysiology Laboratory, Department of Pharmacology, Harvard Medical School, Boston, Massachusetts

(Received for publication August 24, 1964)

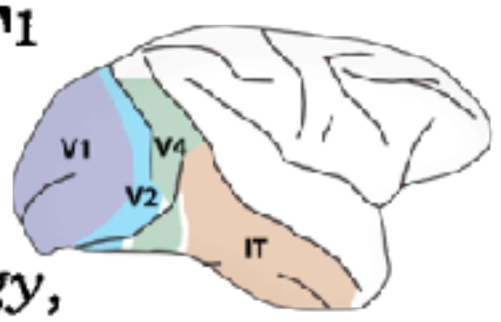

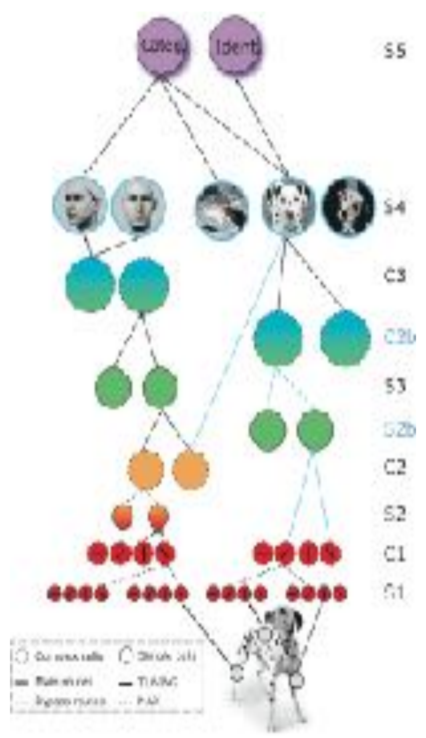

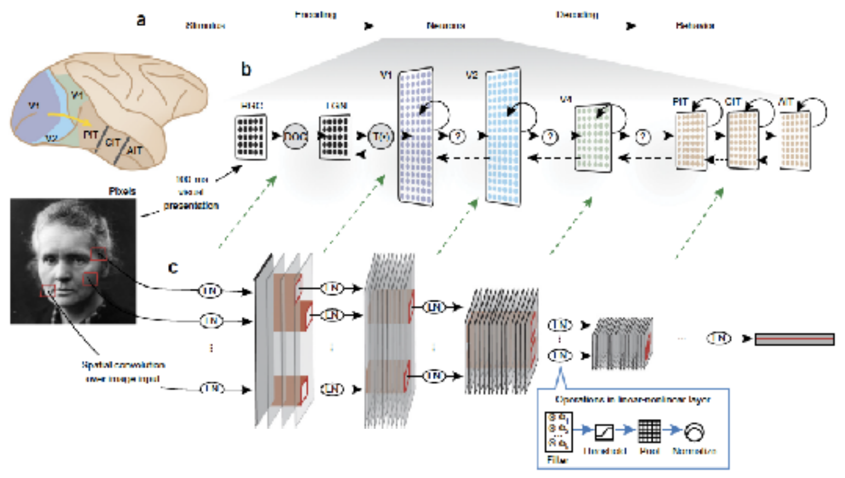

Background: State-of-the-art Machines (“Deep Learning”) Have Emerged From the Brain’s Visual Processing Architecture

Brains / Minds

(ventral visual stream)

What’s the engineering of the future?

Machines ...

State of the Art ResNets

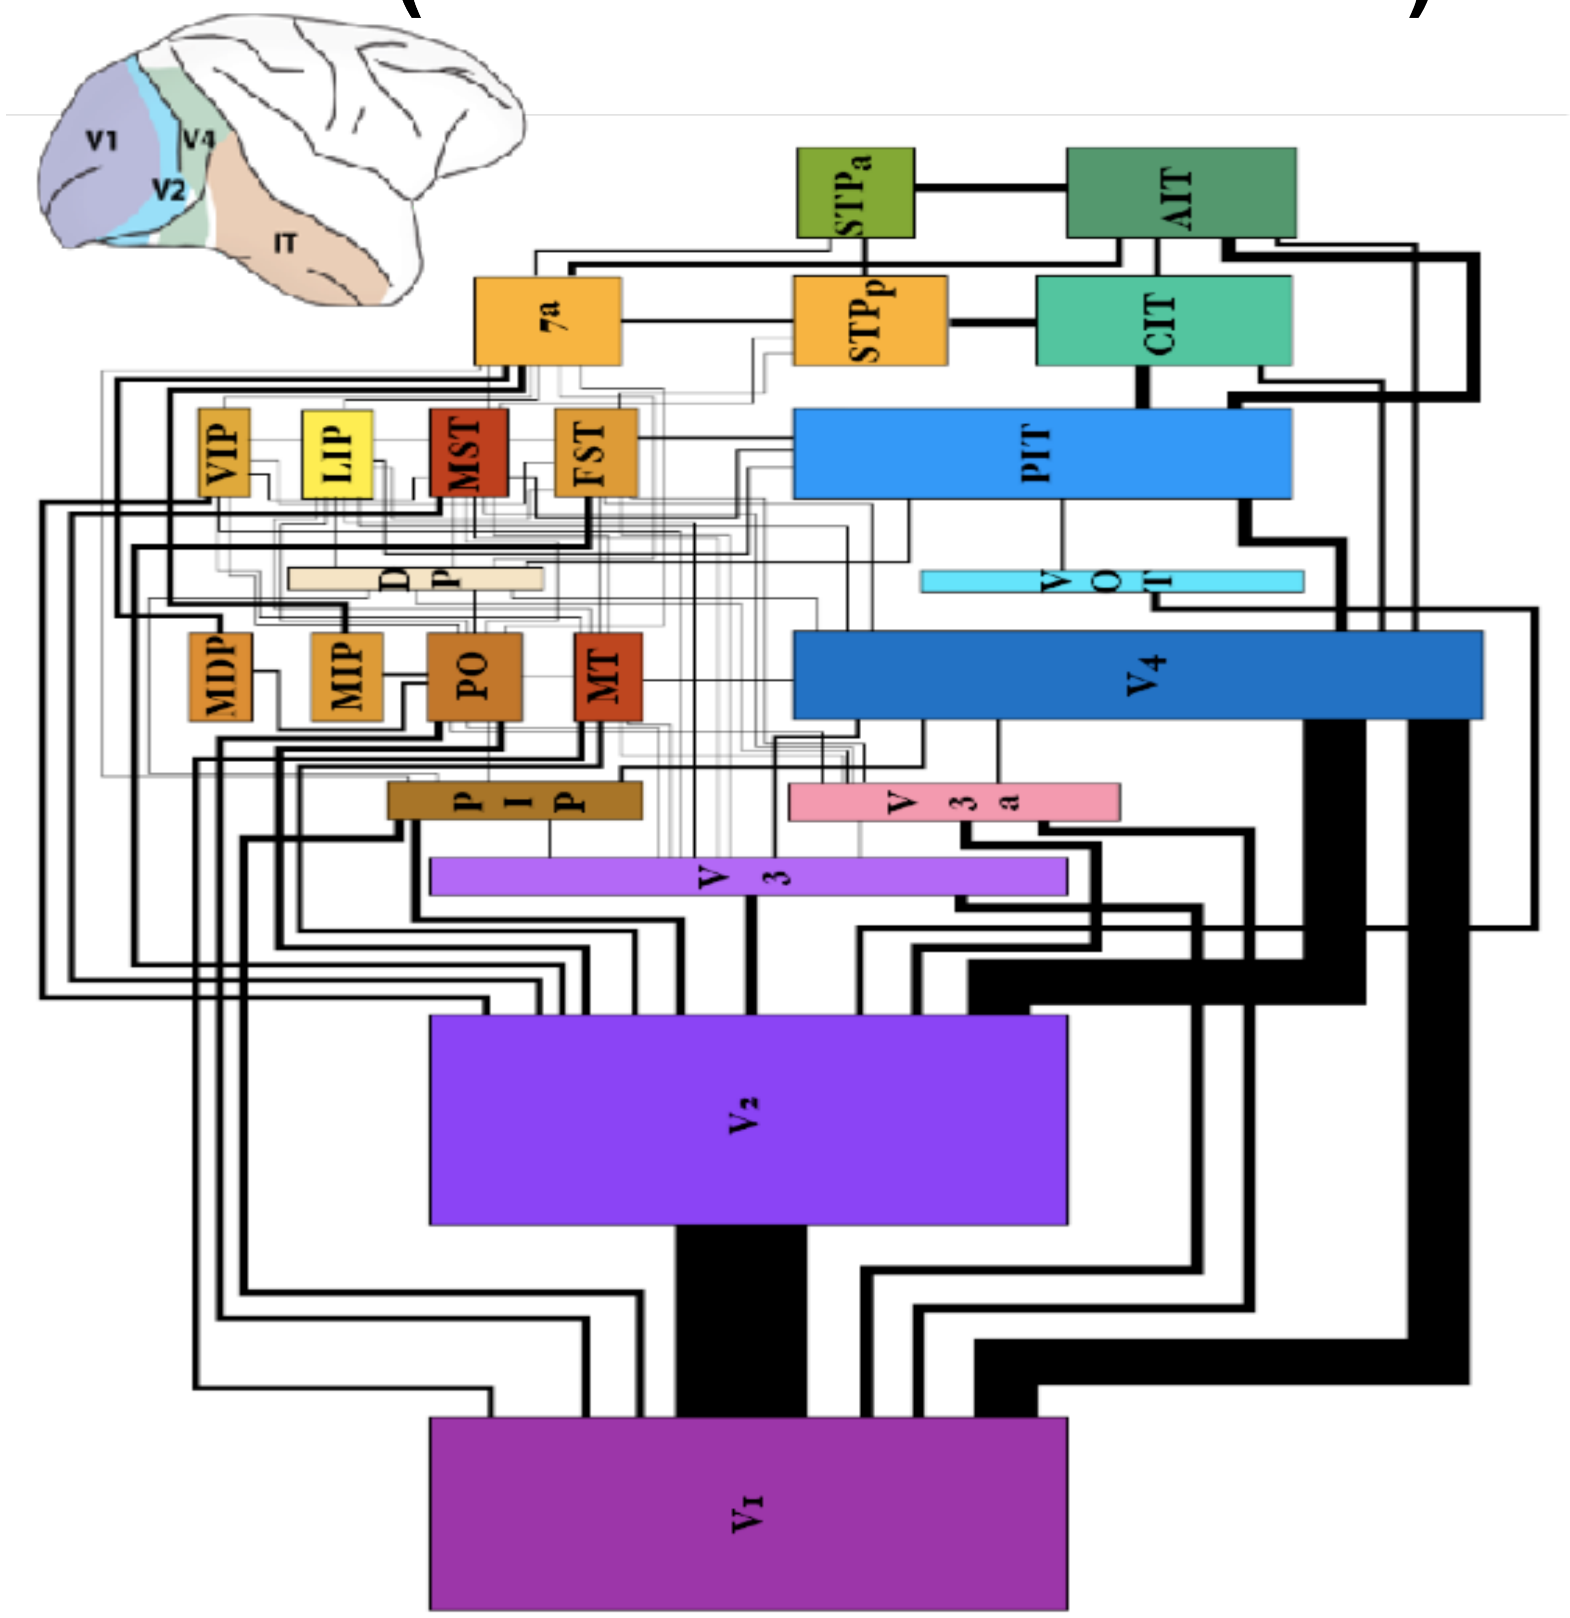

Desimone & Ungerleider 1989; vanEssen+Movshon

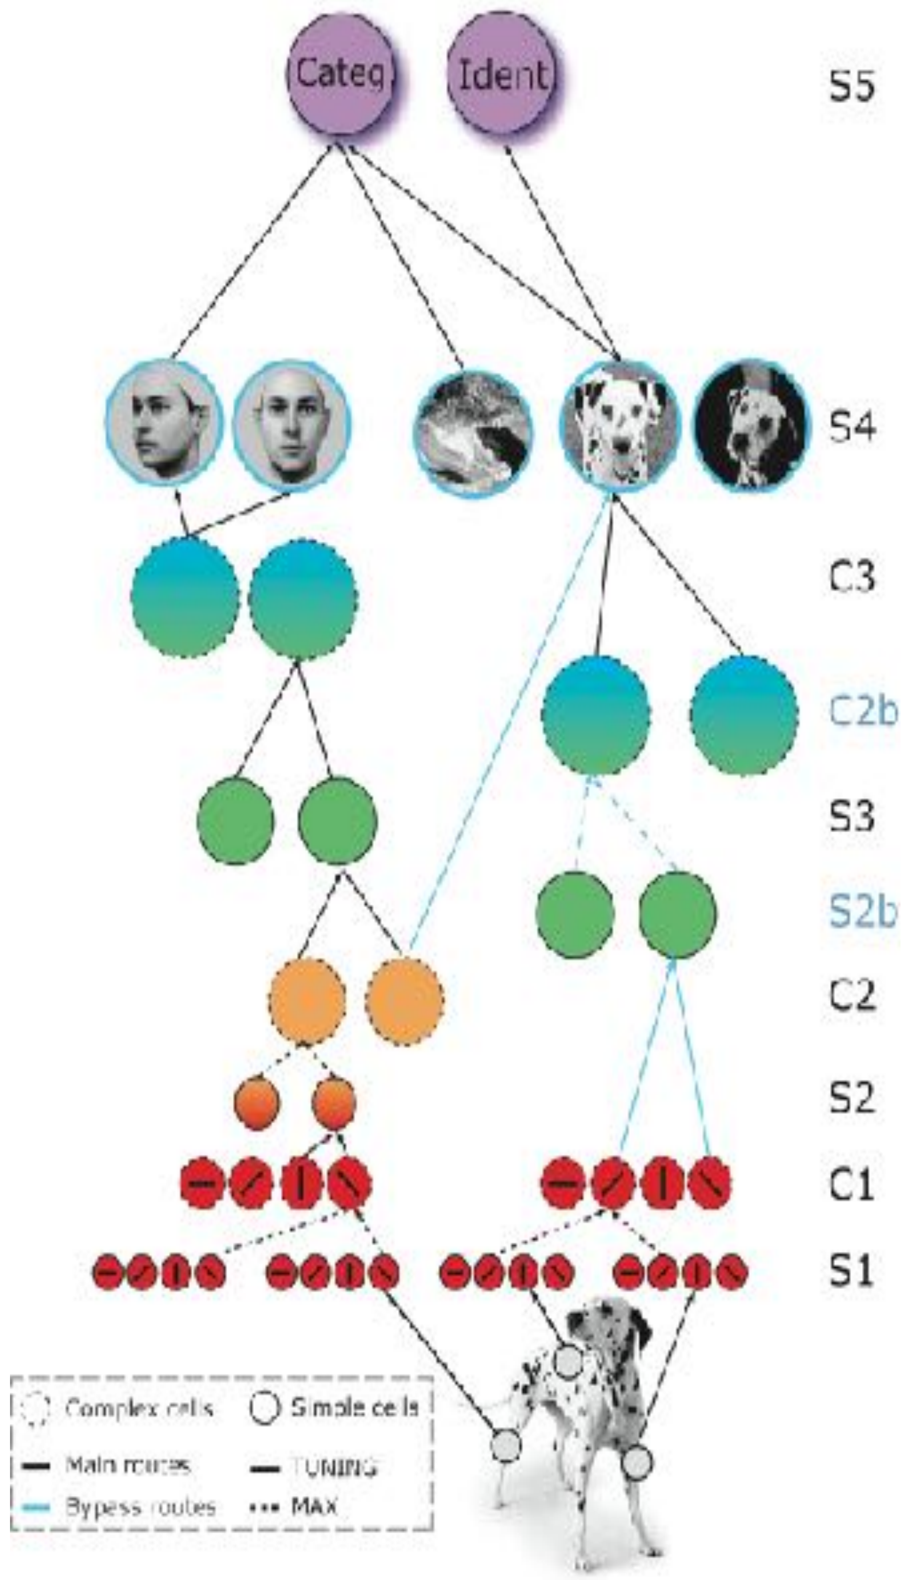

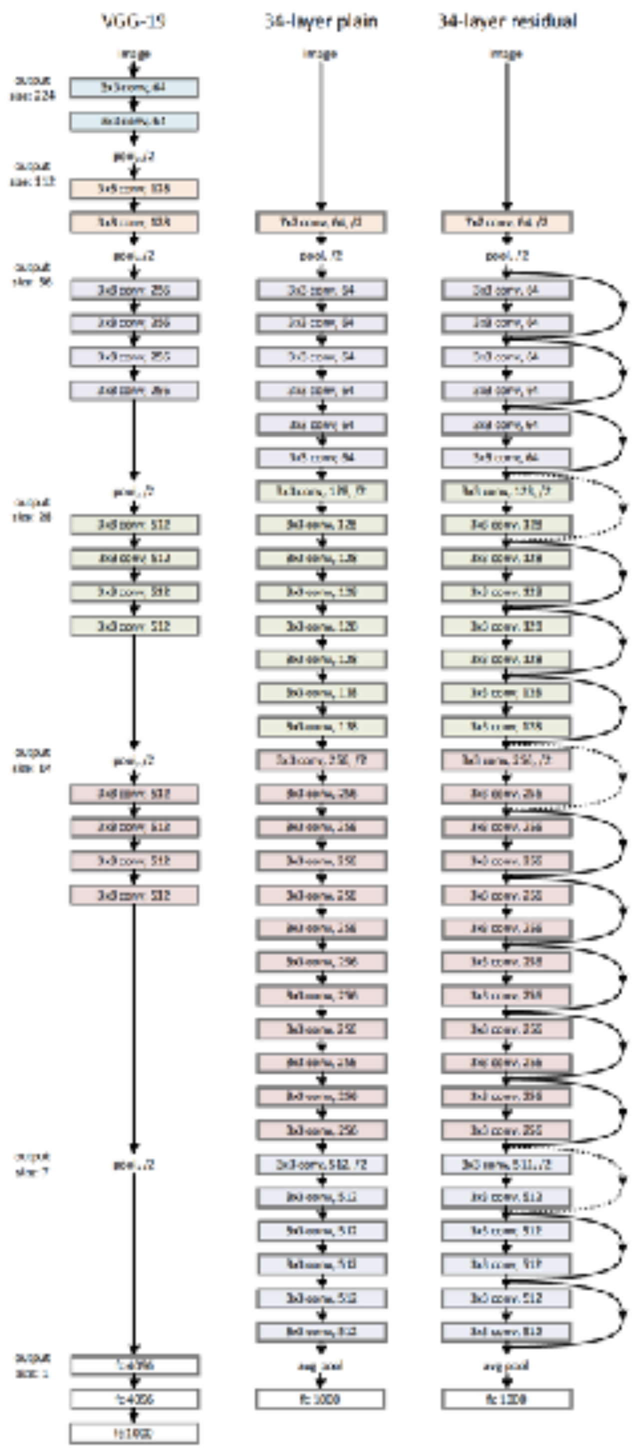

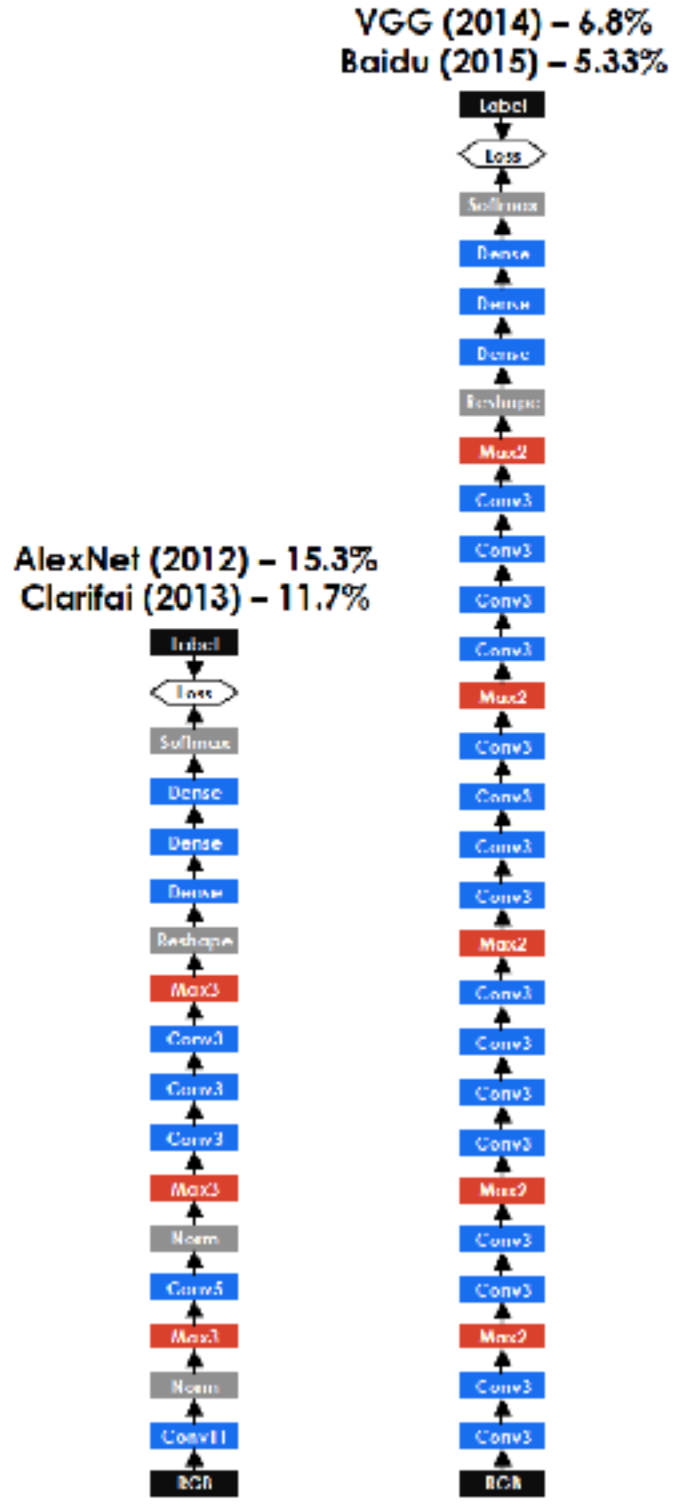

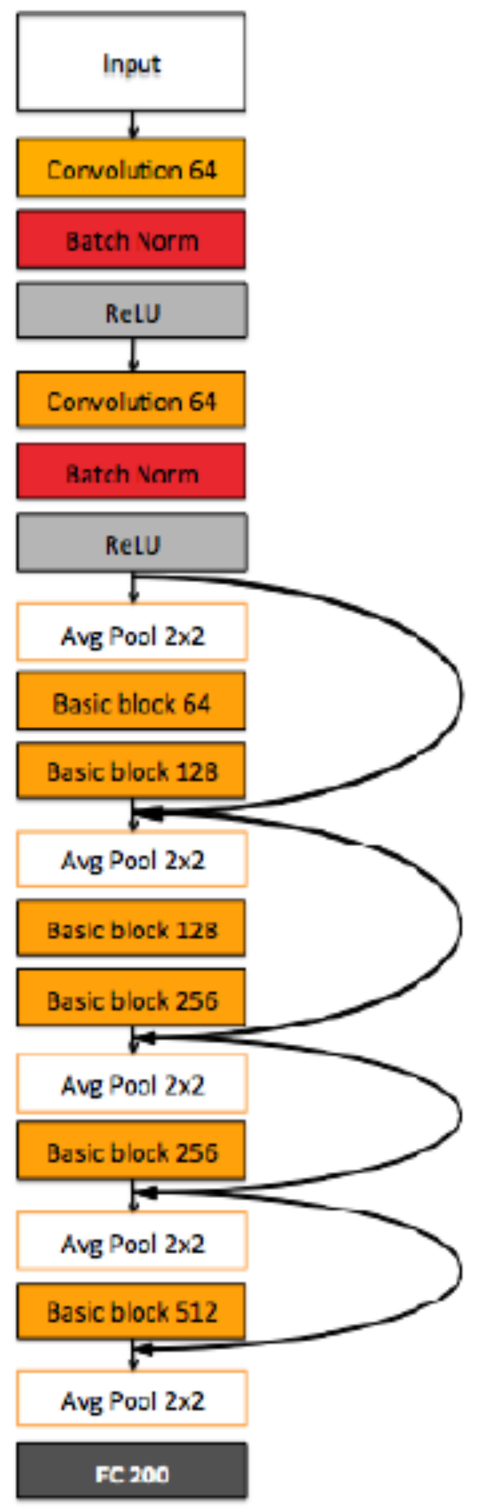

Convolutional networks

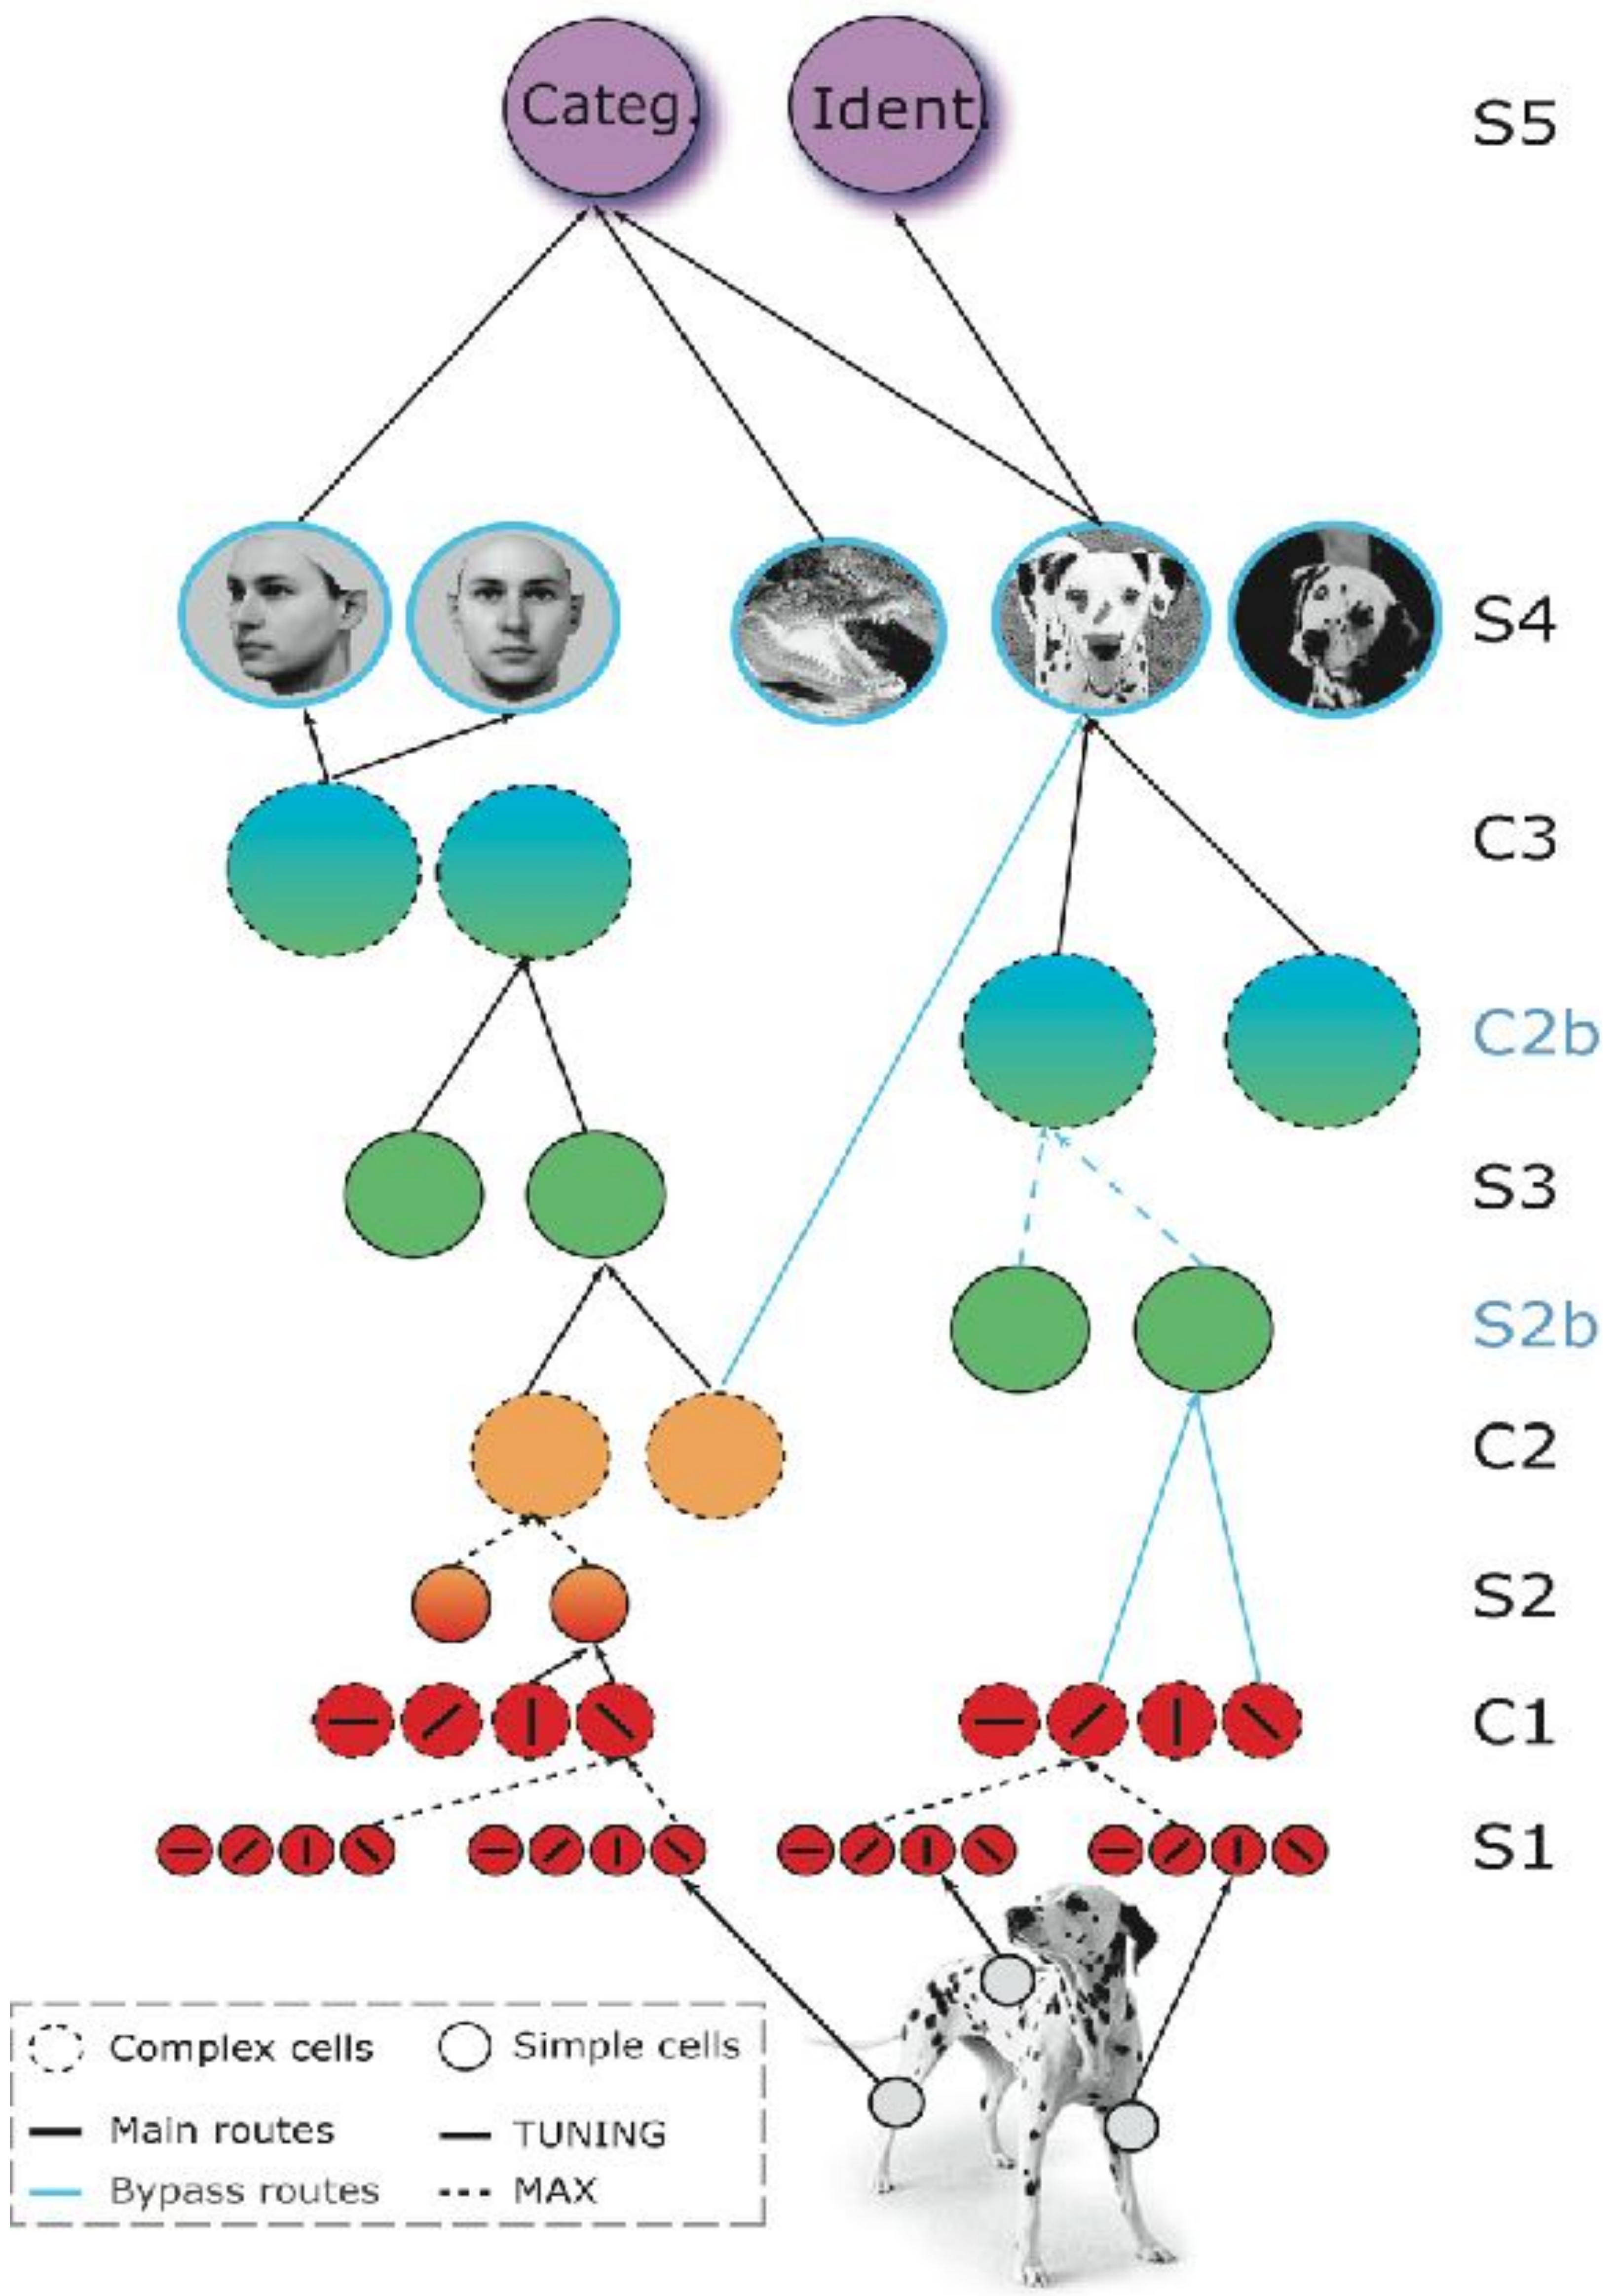

"Hubel-Wiesel" models include

Hubel & Wiesel, 1959: Fukushima, 1980, Wallis & Rolls, 1997; Mel, 1997; LeCun et al 1998; Riesenhuber & Poggio, 1999; Thorpe, 2002; Ullman et al., 2002; Wersing and Koerner, 2003; Serre et al., 2007; Freeman and Simoncelli, 2011….

Riesenhuber & Poggio 1999, 2000; Serre Kouh Cadieu Knoblich Kreiman & Poggio 2005; Serre Oliva Poggio 2007

9.523 overview

- Course description/logistic
- Motivations for this course: the greatest problem in science, CBMM, the MIT Quest
- Neuroscience and AI
- CBMM, the Visual Intelligence moonshot
- Module 1
  - Module 1, theory
  - Module 1, eccentricity
  - Module 1, invariance

Understanding the World

CBMM Organizational Chart (future)

- **Director** — *Tomaso Poggio*
  - **EAC**
  - **Managing Director** — *Kathleen Sullivan (MIT)*
    - **Administrative Assistant**
  - **Education Coordinator** — *Ellen Hildreth (WC)*
    - **Technology Director**
  - **Education Evaluation** — *Lizanne DeStefano (GT)*
  - **KT Coordinator** — *Boris Katz (MIT)*
  - **Diversity Coordinator** — *Mandana Sassanfar (MIT)*
  - **Deputy Director** — *Gabriel Kreiman (HU)*
  - **Associate Director & Trainee Coordinator** — *Matt Wilson (MIT)*
  - **Research Director** — *Kenneth Blum (HU)*
    - **Module 1: VISUAL STREAM** — *Tomaso Poggio, Shimon Ullman (MIT)*
    - **Module 2: BRAIN OS** — *Gabriel Kreiman (HU)*
    - **Module 3: COGNITIVE CORE** — *Nancy Kanwisher, Joshua Tenenbaum (MIT)*
    - **Module 4: TOWARDS SYMBOLS** — *Boris Katz, Shimon Ullman (MIT)*

CENTER FOR Brains Minds+ Machines

The CBMM Visual Intelligence moonshot

Shimon

Module One: Ventral Visual Stream

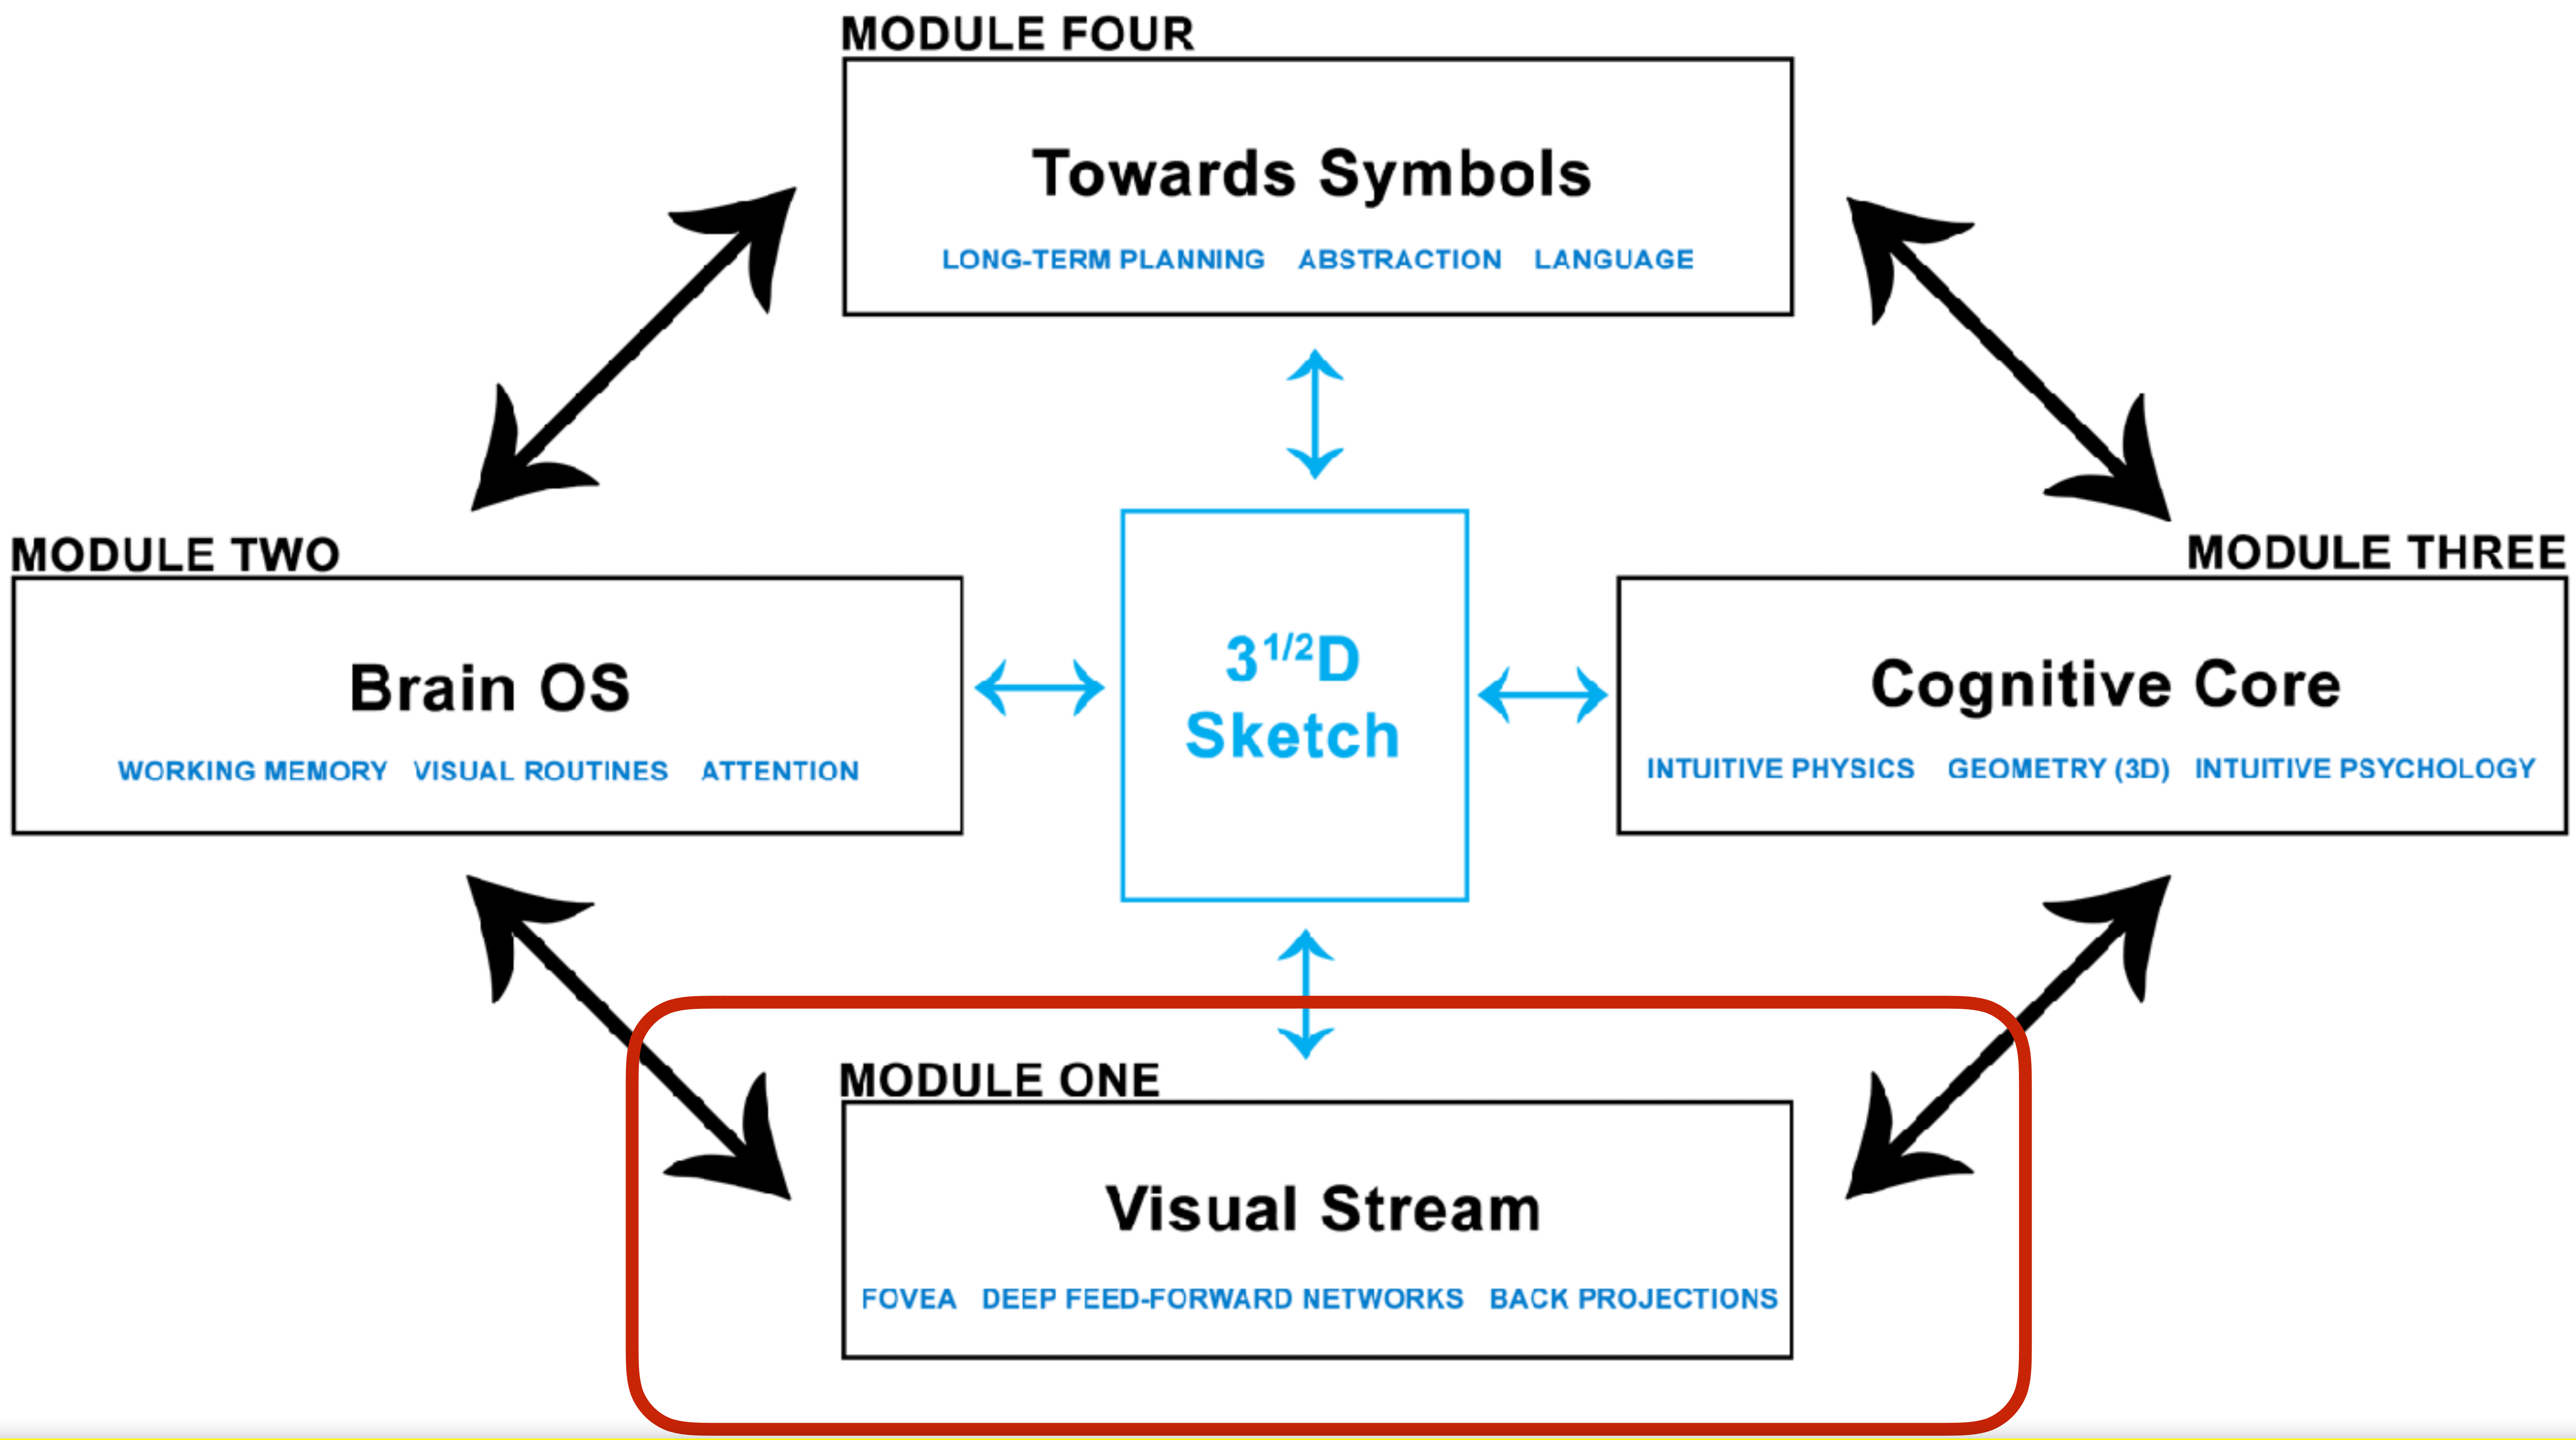

Foundation for other modules: Center of visual gaze, first 200 msec of visual processing

9.523 overview

- Course description/logistic
- **Motivations for this course: the greatest problem in science, CBMM, the MIT Quest**
- Neuroscience and AI
- CBMM, the Visual Intelligence moonshot
- Module 1
    - Module 1, theory
    - Module 1, eccentricity
    - Module 1, invariance

Importance of viewpoint invariance

Invariant Representations Lead to Lower Sample Complexity for a Supervised Classifier

Theorem *(translation case)*
Consider a space of images of dimensions $d \times d$ pixels which may appear in any position within a window of size $rd \times rd$ pixels. The usual image representation yields a sample complexity (of a linear classifier) of order $m = O(r^2 d^2)$;the oracle representation (invariant) yields (because of much smaller covering numbers) a sample complexity of order $m_{oracle} = O(d^2) = \frac{m_{image}}{r^2}$

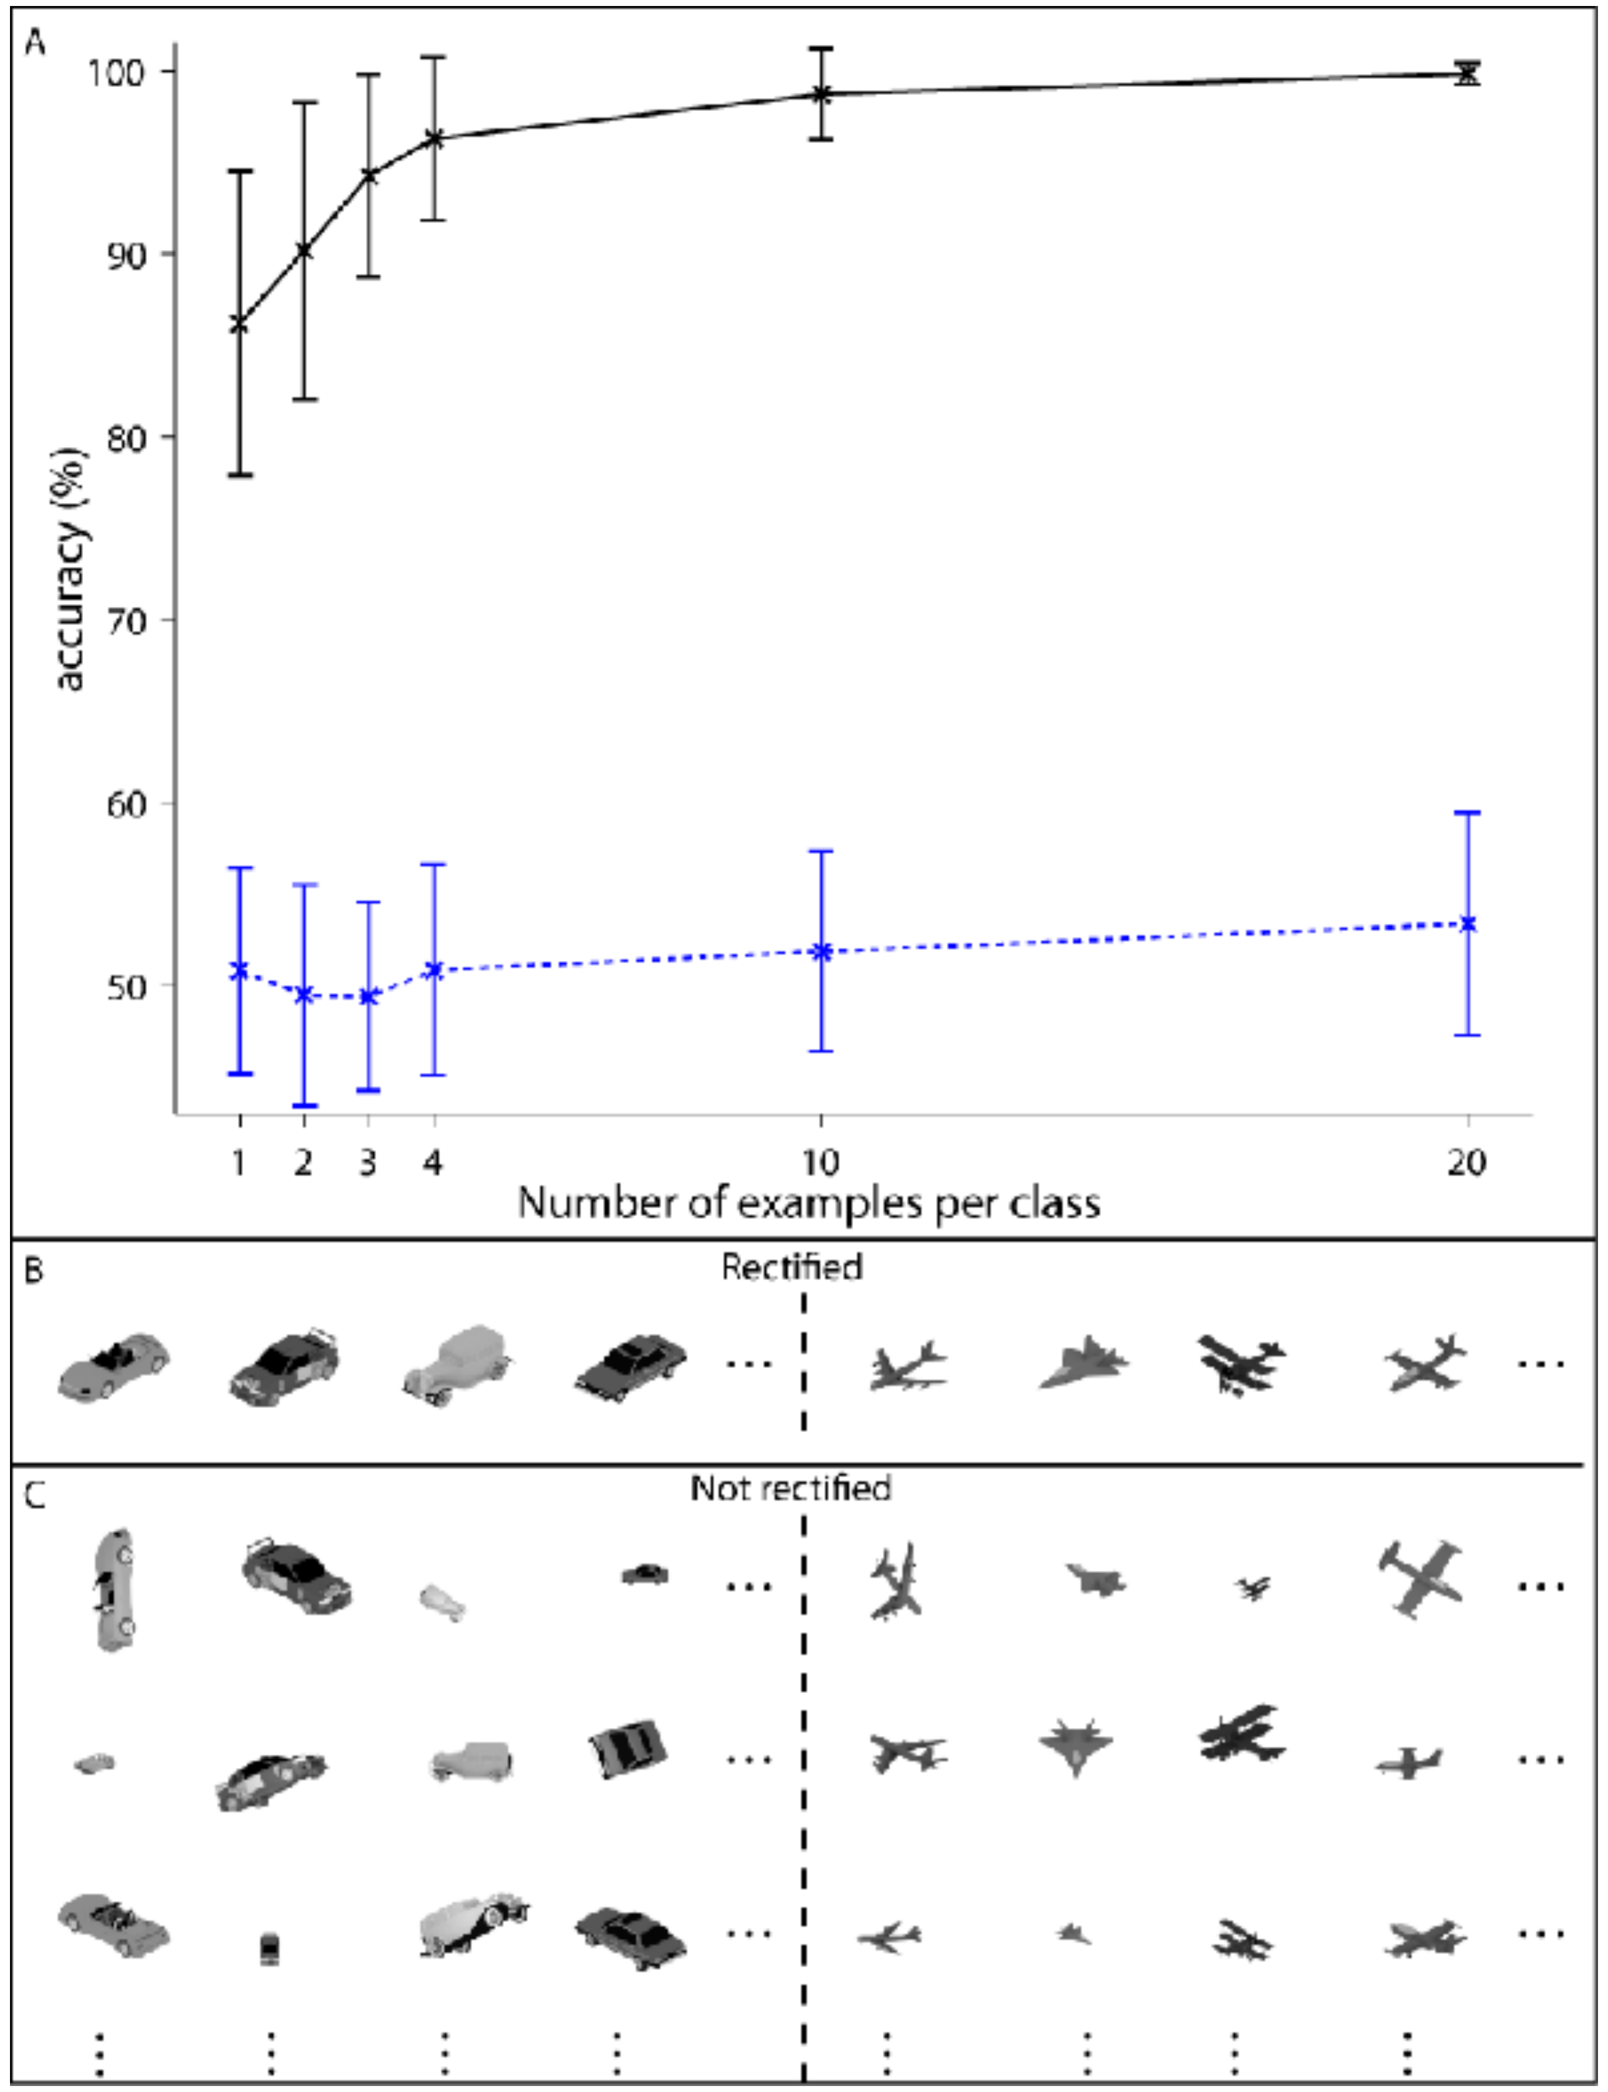

Puzzle of linear dependence of resolution on eccentricity

Note: we focus on the sampling layout of the retinal ganglion cells (RGCs) - the *outputs* of the retina.

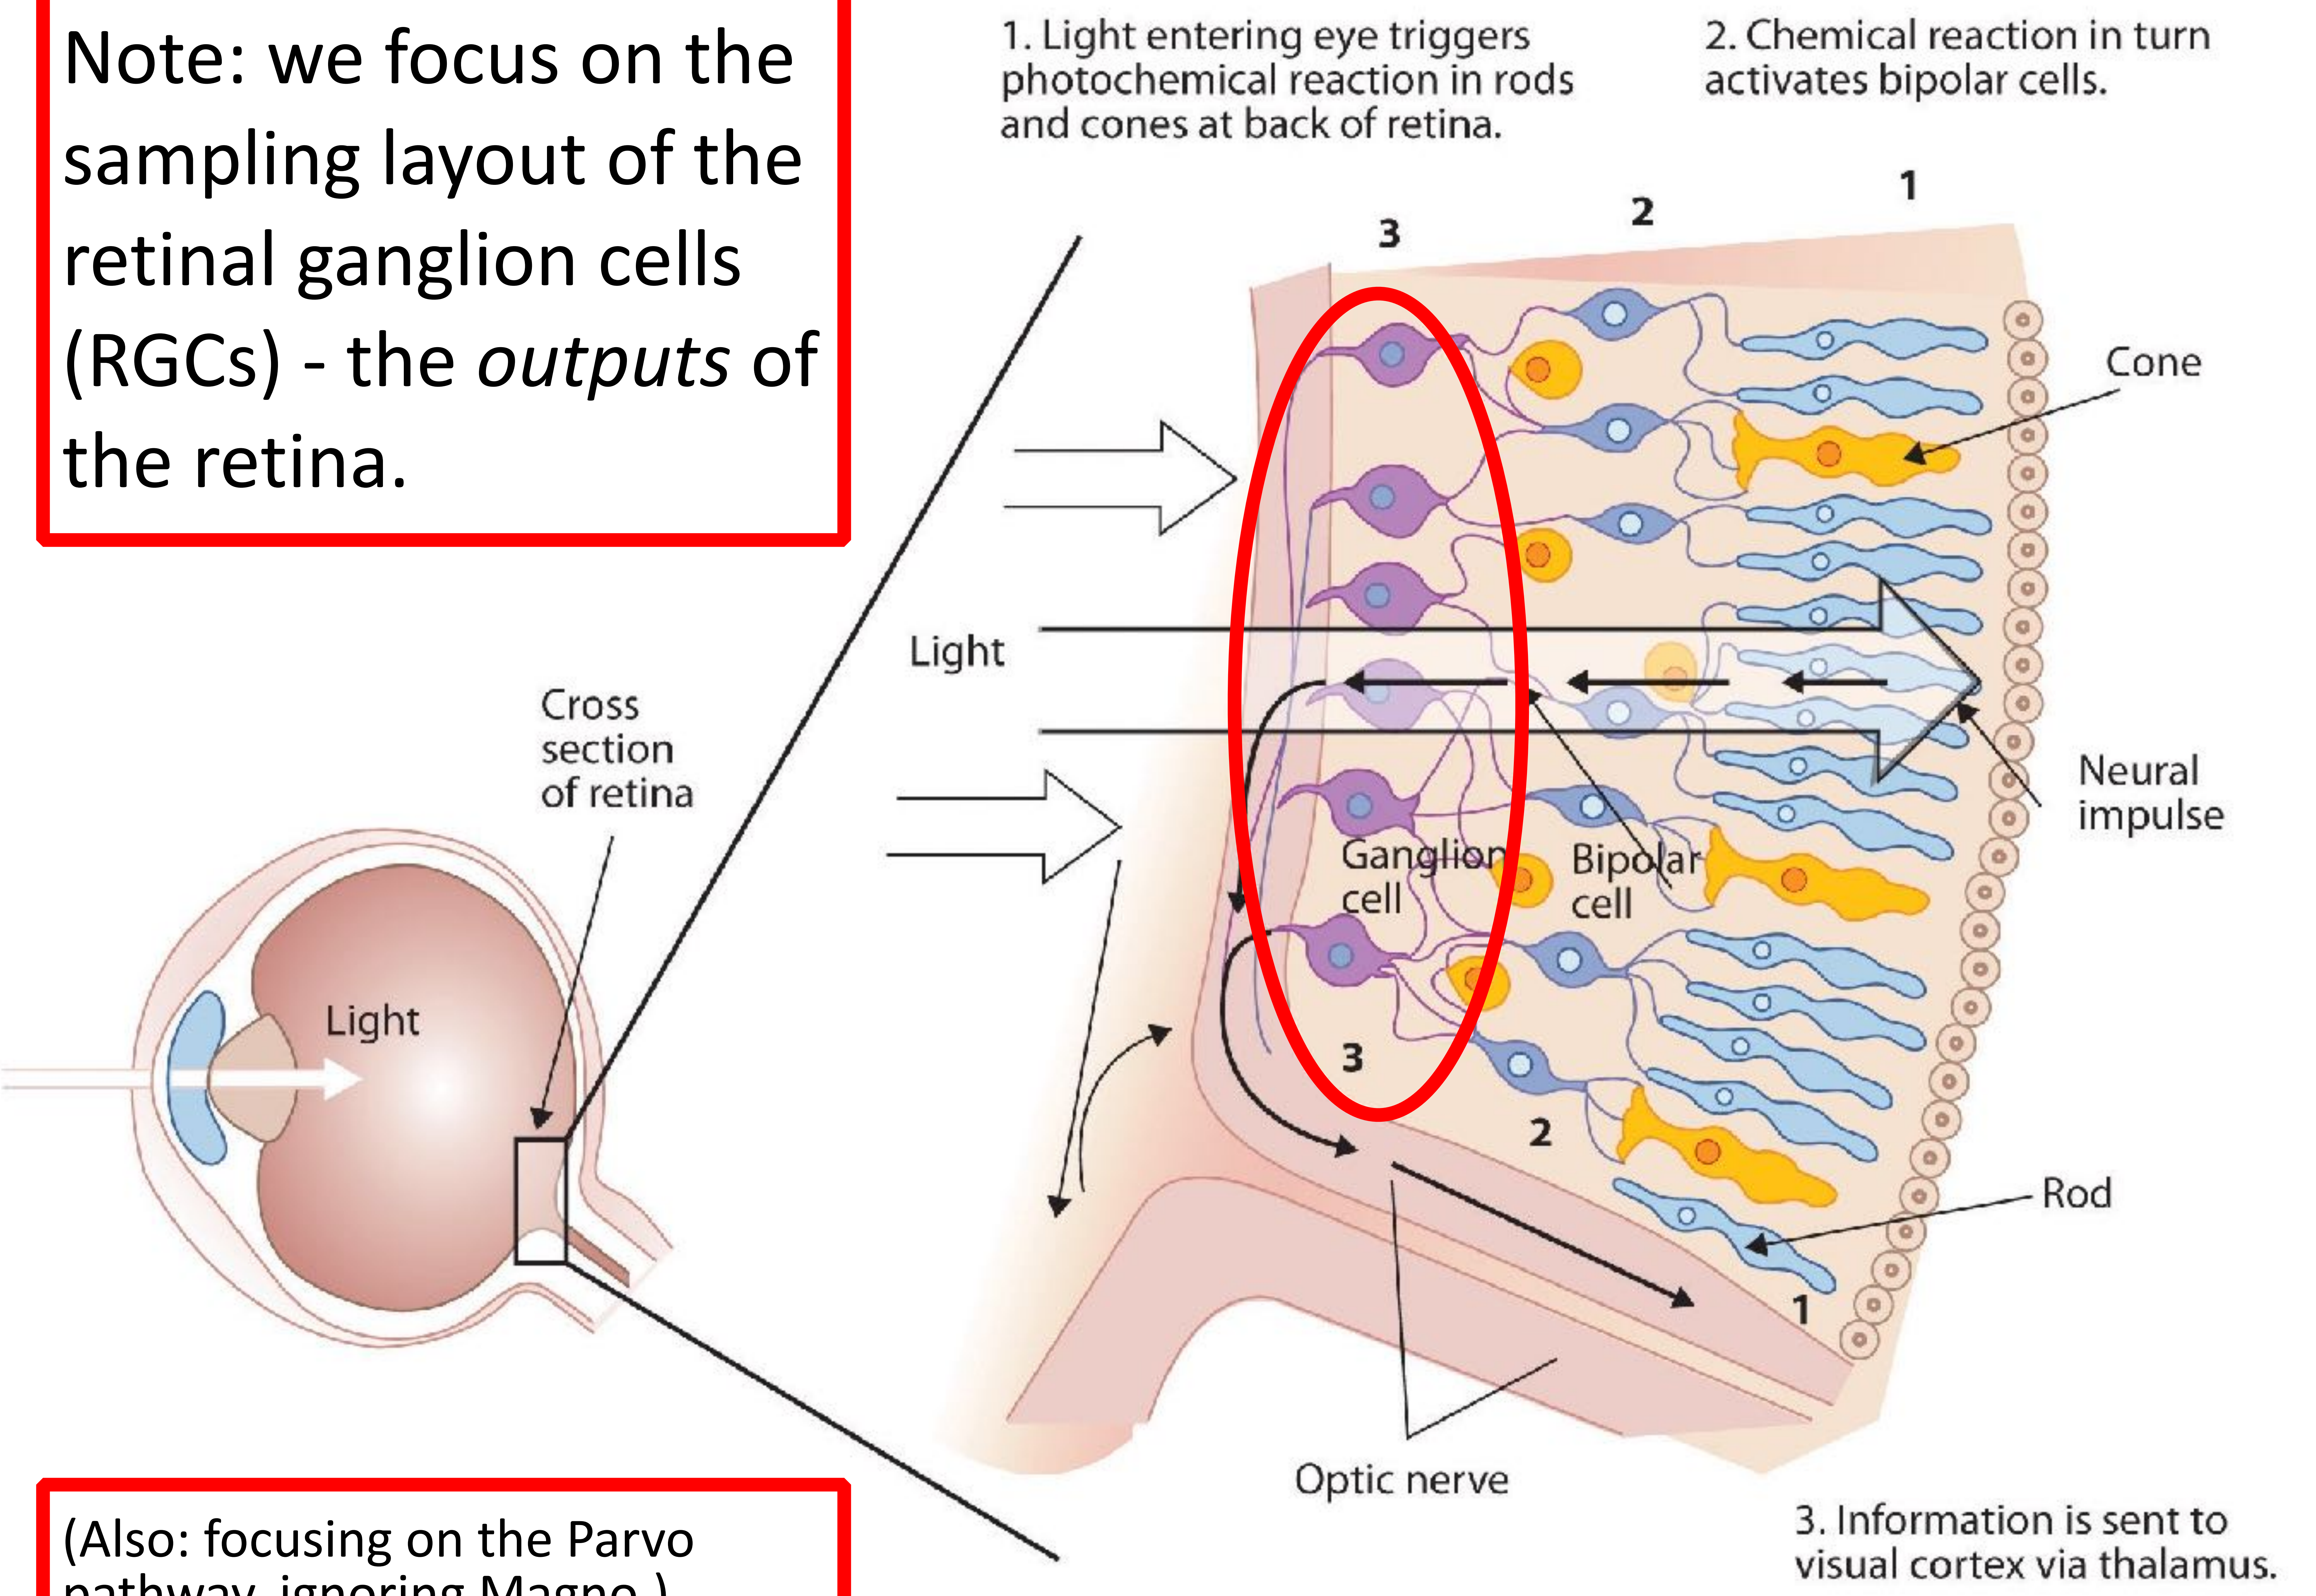

(Also: focusing on the Parvo pathway, ignoring Magno.)

Receptive field size vs. eccentricity - HW

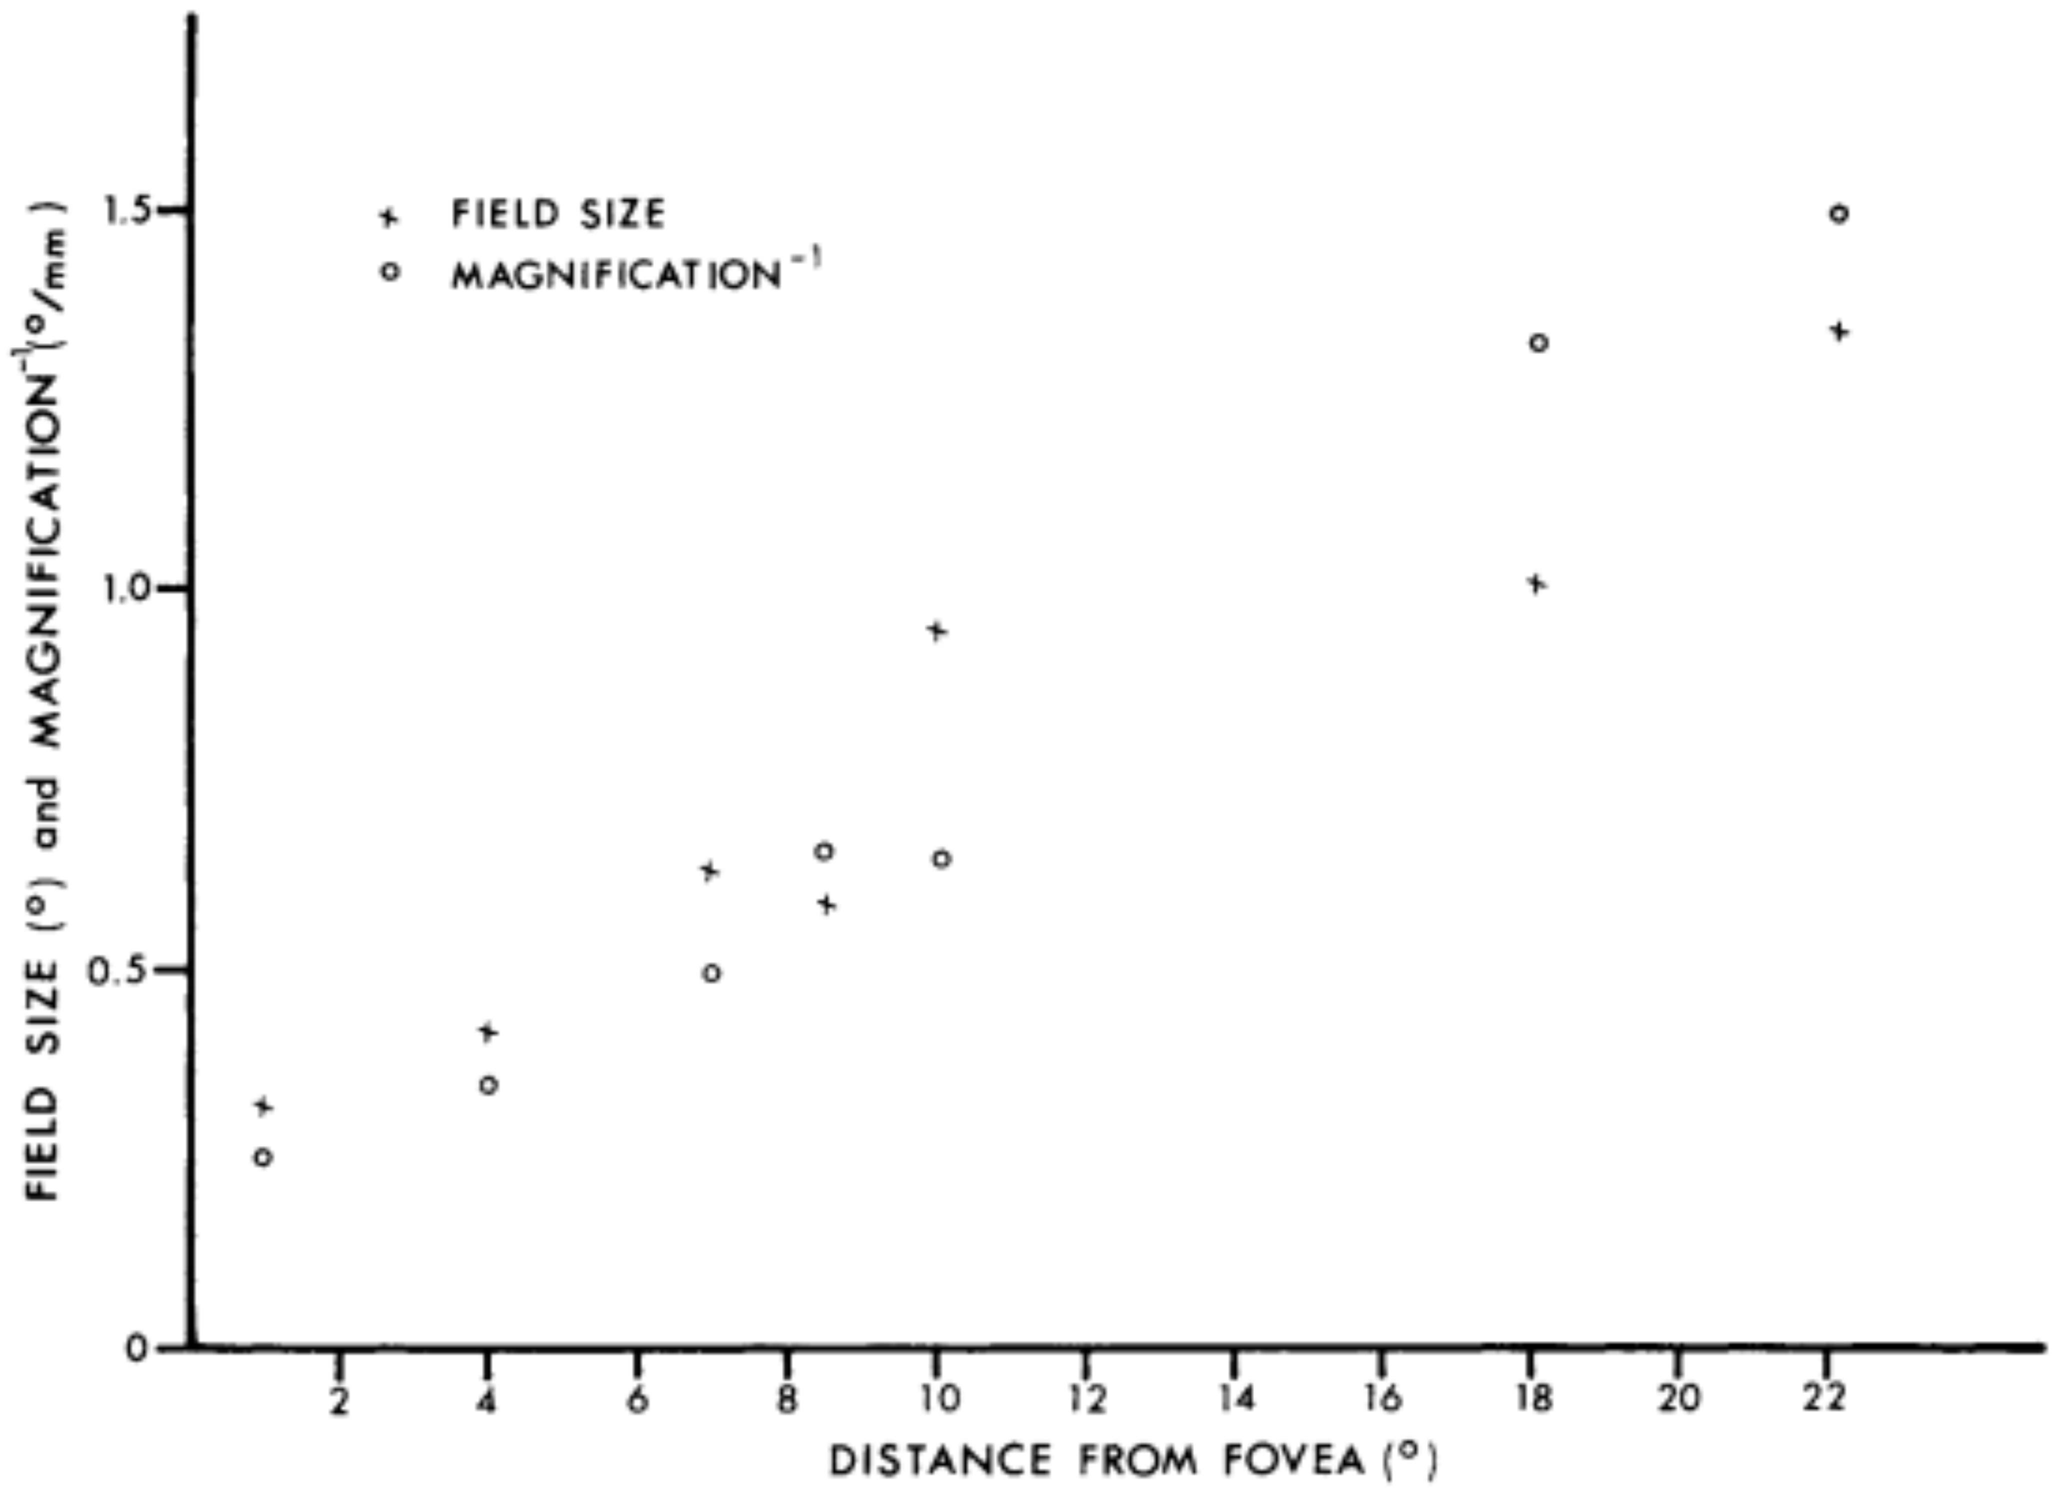

Fig. 6A Graph of average field size (crosses) and magnification^{-1} (open circles) against eccentricity, for five cortical locations. Points for 4°, 8°, 18° and 22° were from one monkey; for 1°, from a second. Field size was determined by averaging the fields at each eccentricity, estimating size from (length × width)$^{0.5}$.

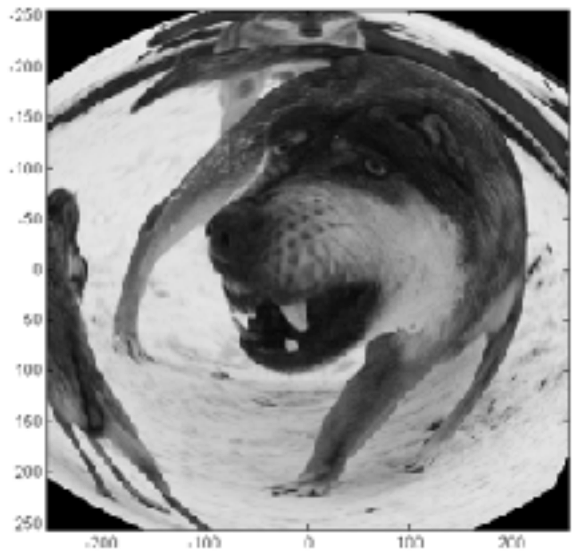

Hubel and Wiesel, 1971

Foundation for other modules ==> eye movements:
Center of visual gaze, first 200 msec of visual processing

E

N H

N F D

Z X T P

F T D Z U

H T P N F D

Z D T N U H P

H P X T Z F N

P D T N U H Z

Multiple scales?

Scatter of receptive field sizes in V1

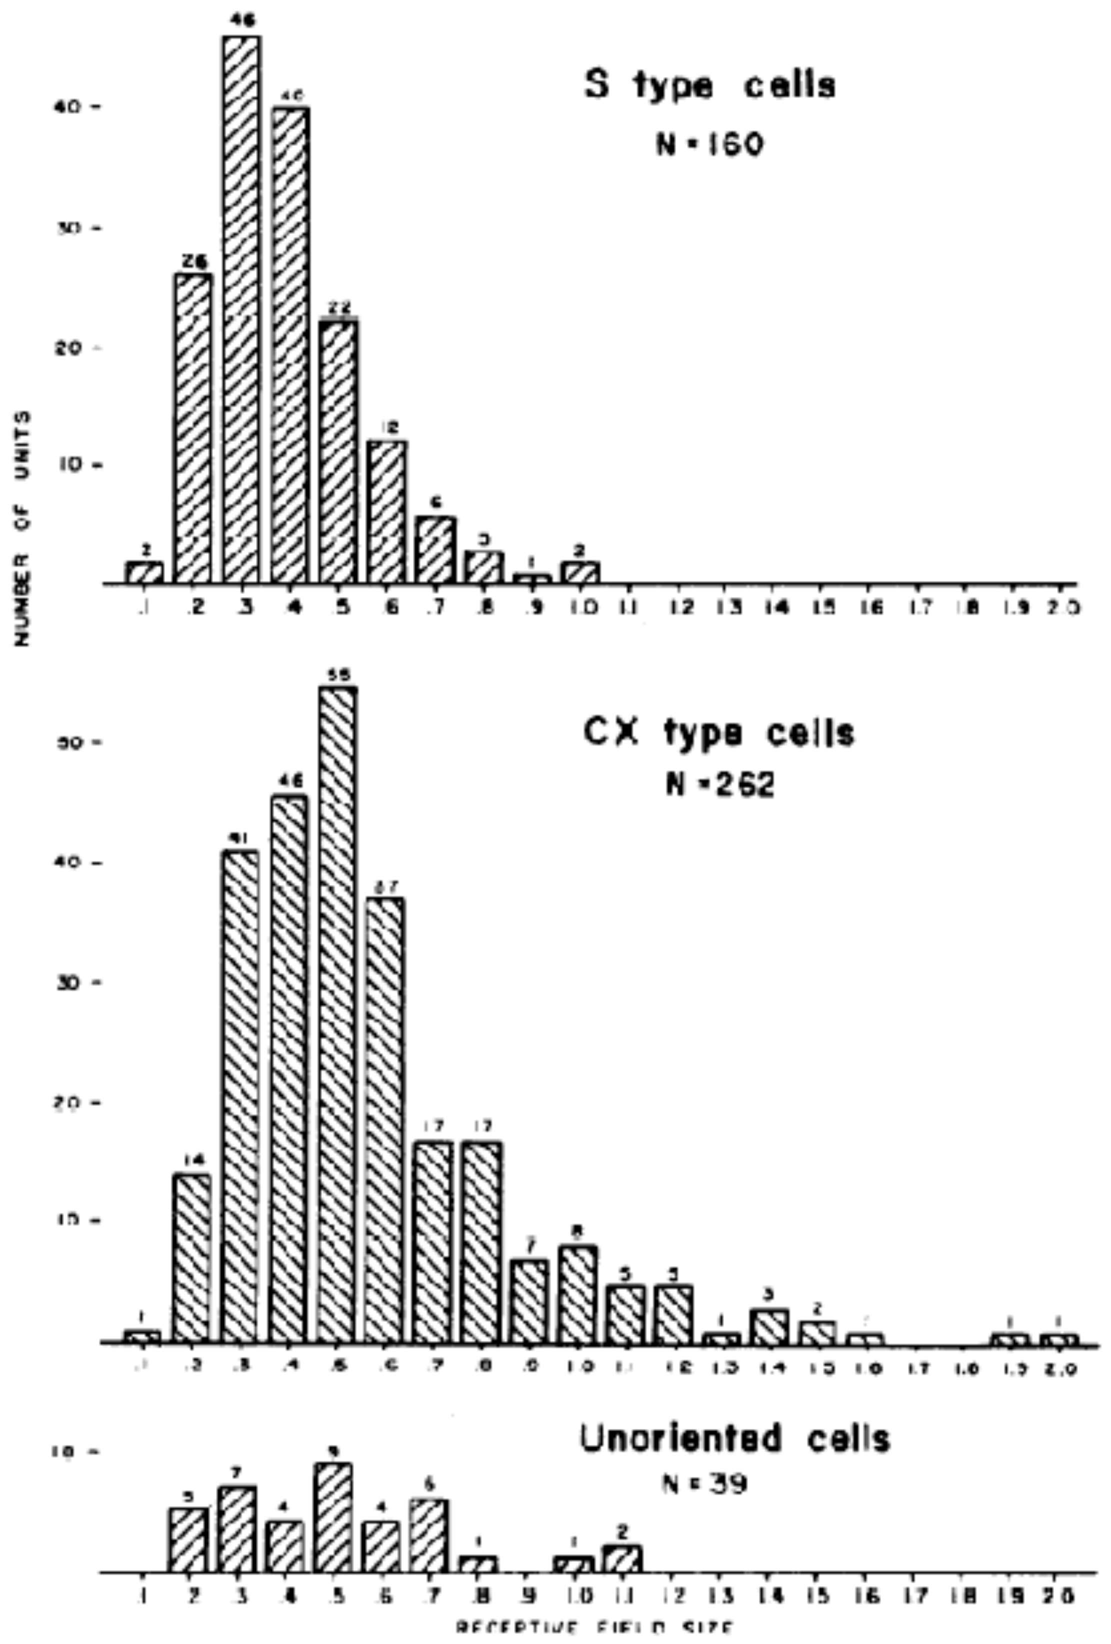

FIG. 18. Distribution of overall receptive-field width.

The data in Fig. 18, obtained from 589 cells, show a skewed distribution of sizes which vary over almost a 20-fold range. Only a relatively small fraction of this range is due to variation in retinal eccentricity. Even in a vertical penetration, where receptive fields scatter only over a small area, field size can vary over a large range, with the receptive fields of cells in layer 4 being small and those in layers 5 and 6 often being very large (10, 13).

Figure 18 also shows that CX-type cells, on the whole, have larger receptive fields than do S-type cells. This difference is statistically significant (F > 0.001), although there is clearly considerable overlap in these populations. This is due, in part,

Schiller, P., Finlay, B., Volman S. Quantitative Studies of Single Cells Properties in monkey striate cortex, 1976

Is invariance the computational reason for linear dependence of RF size on eccentricity??

Invariance window

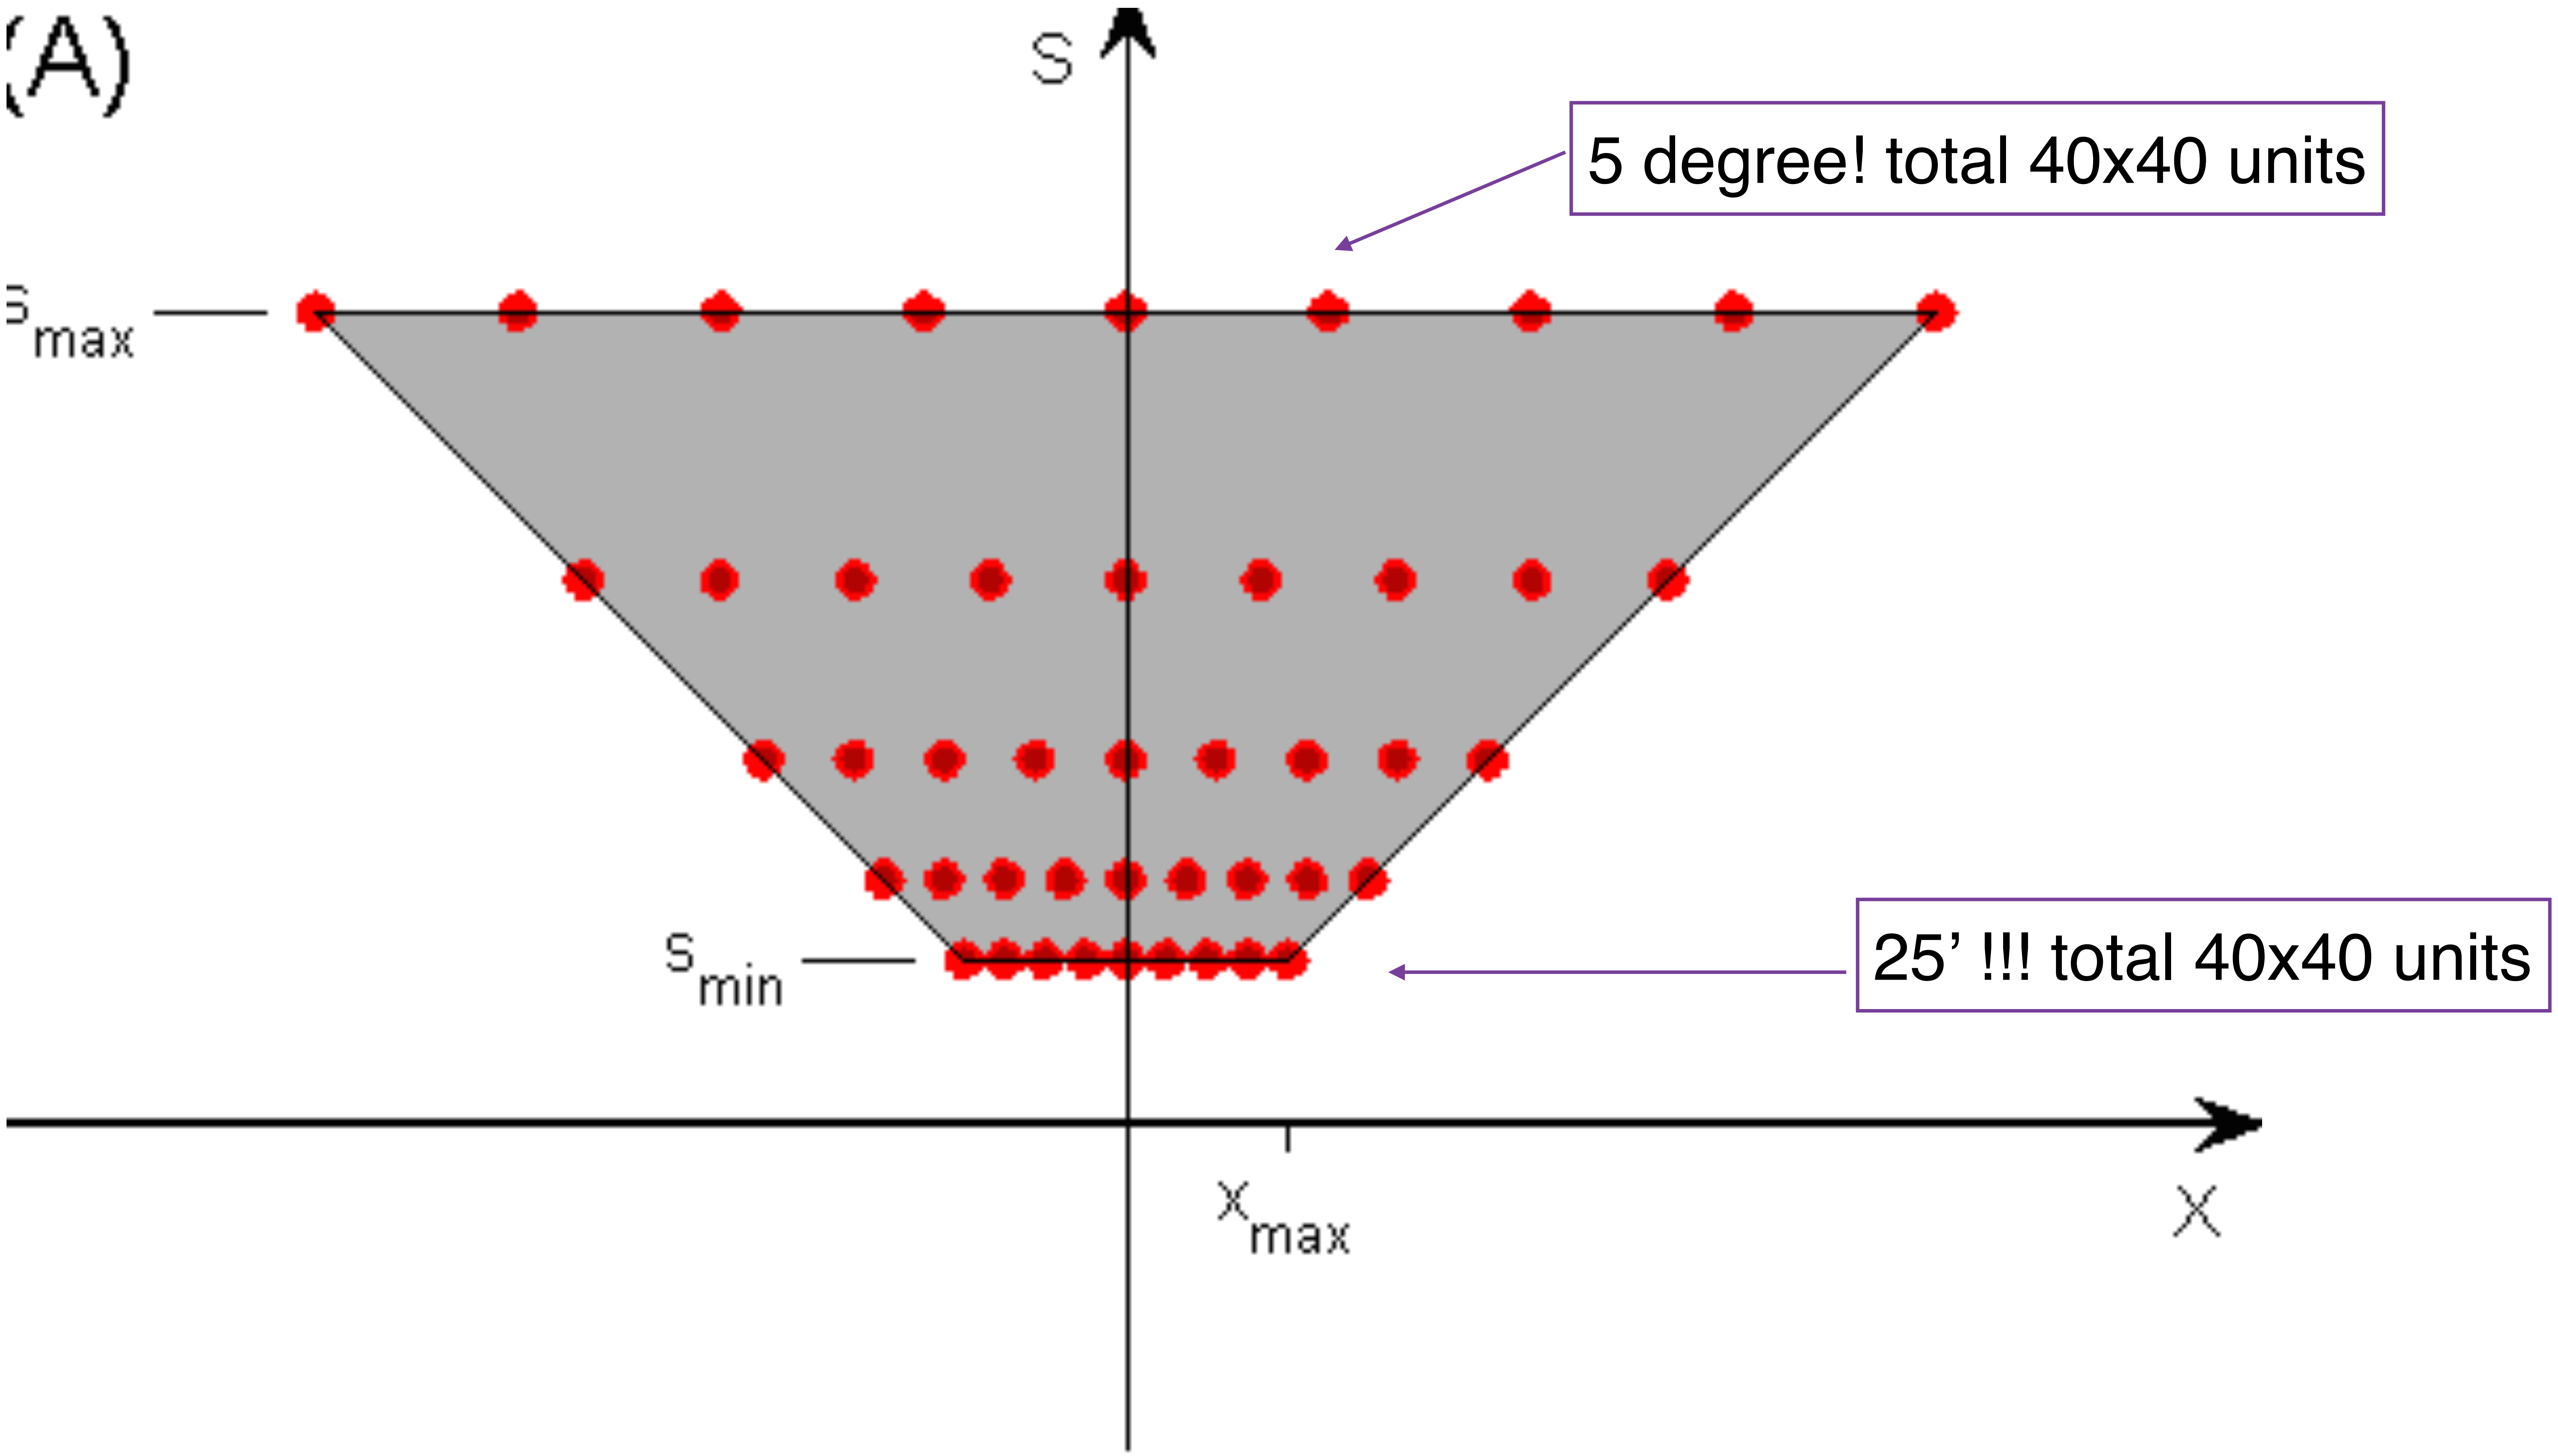

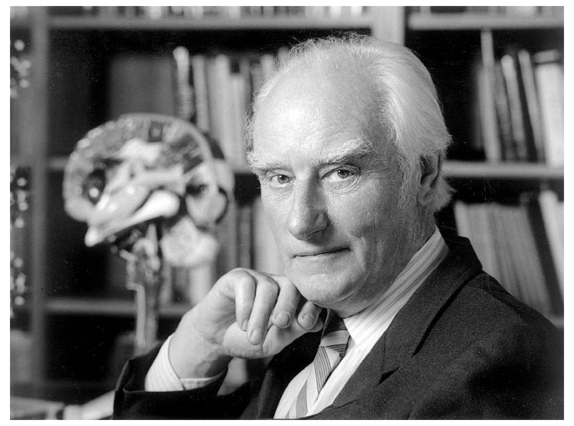

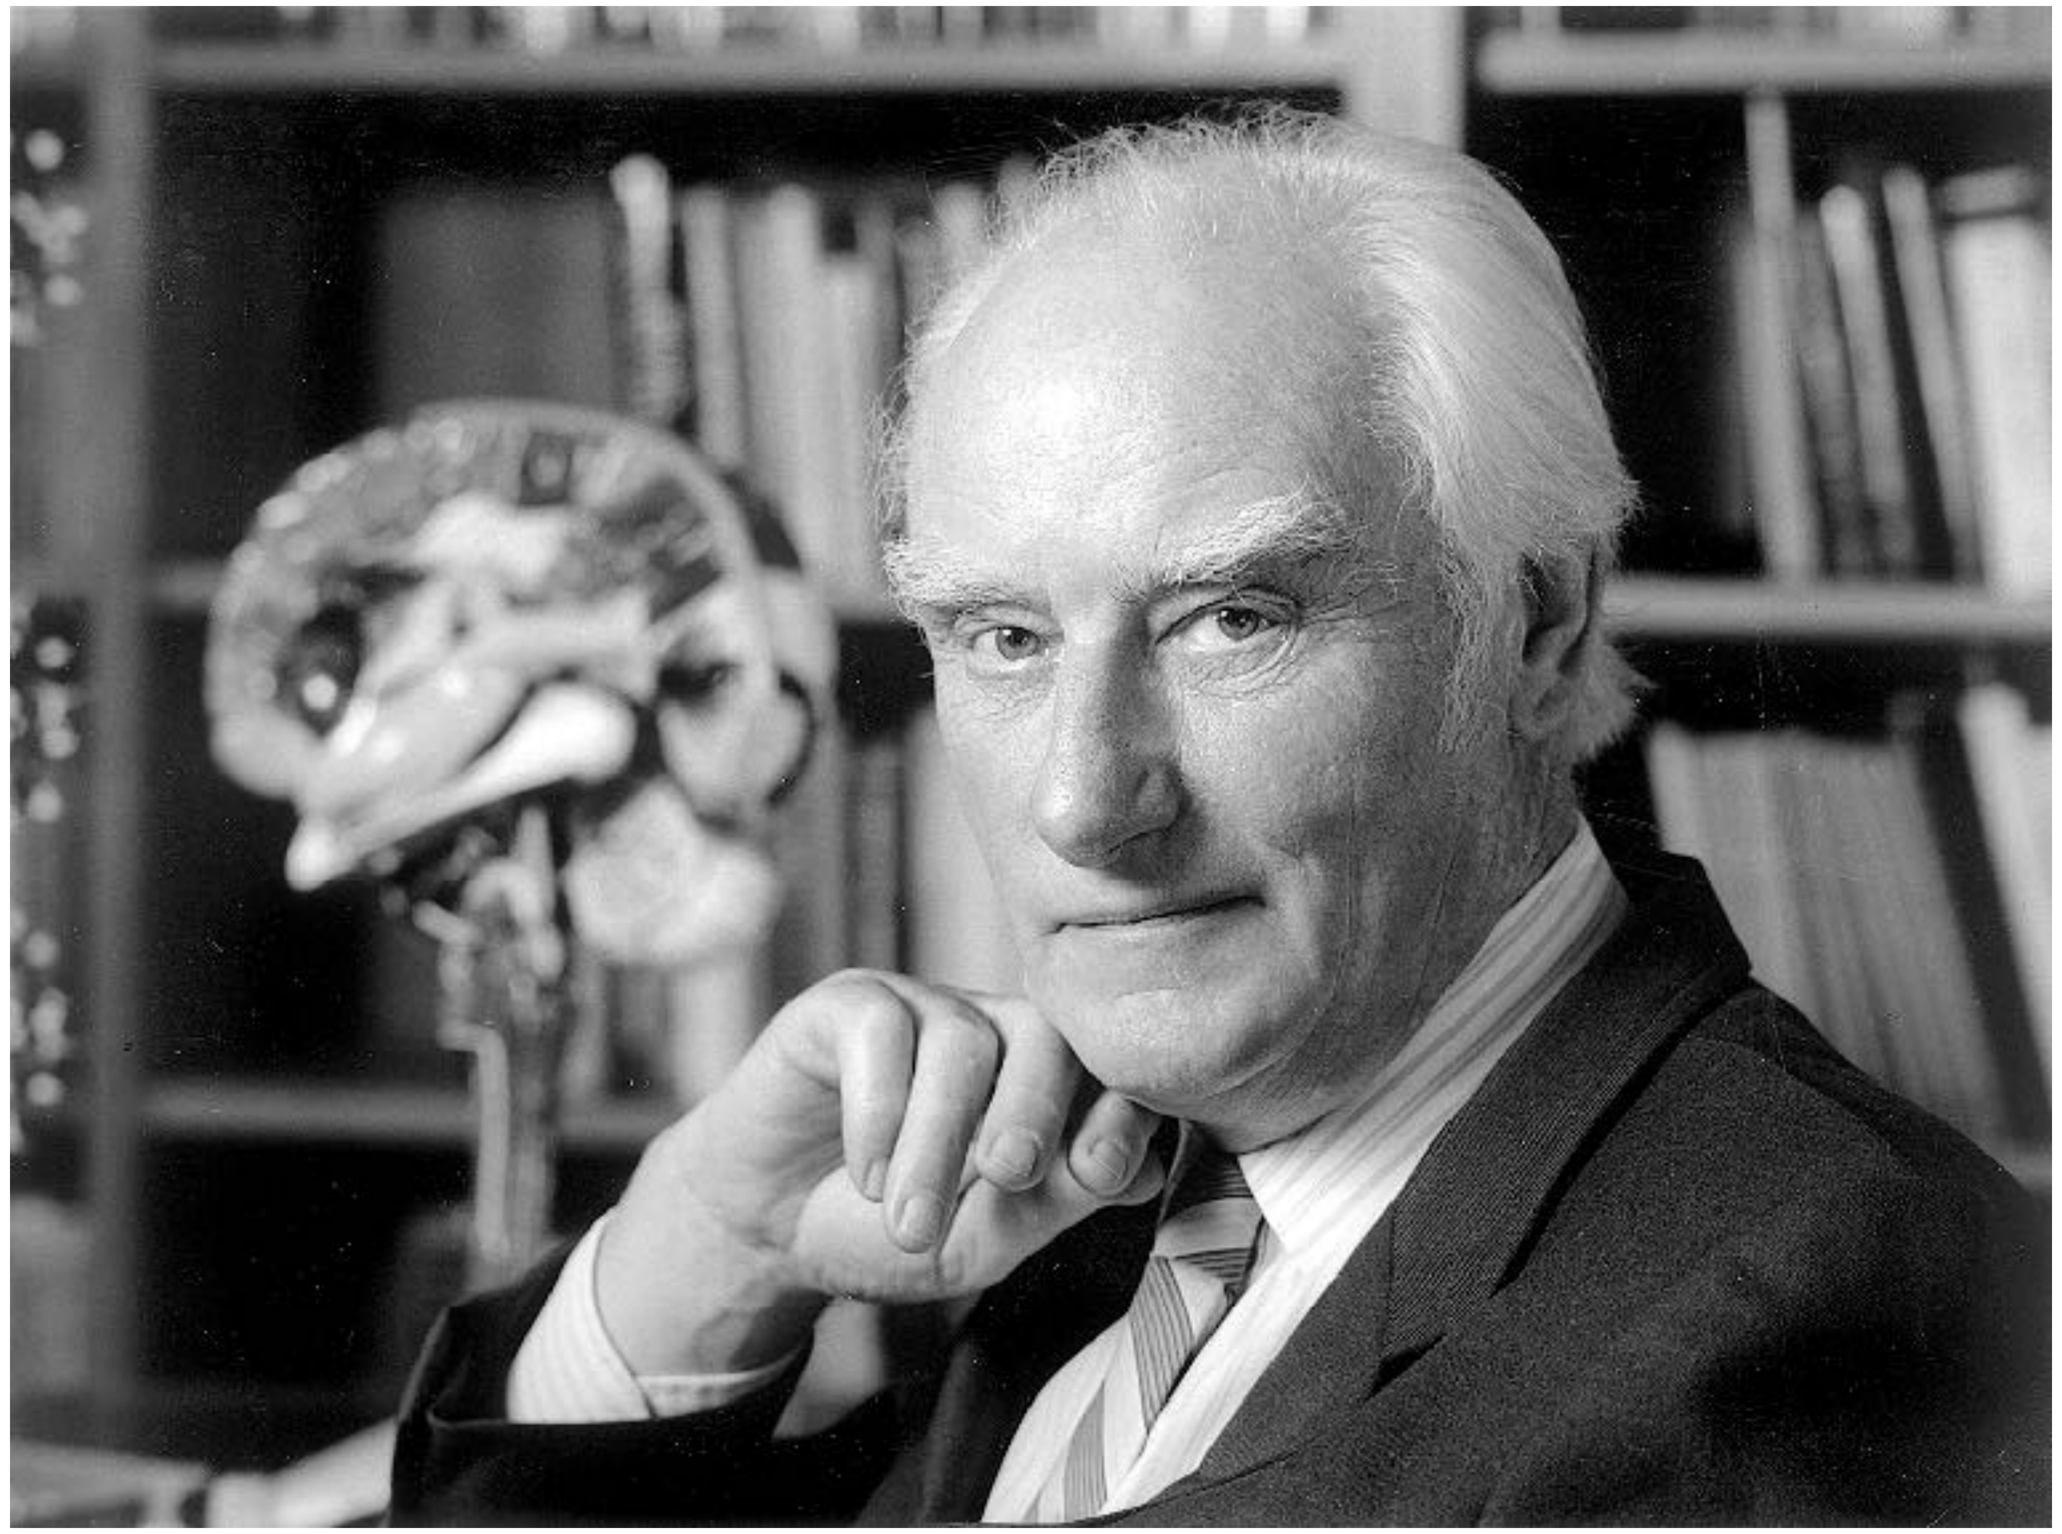

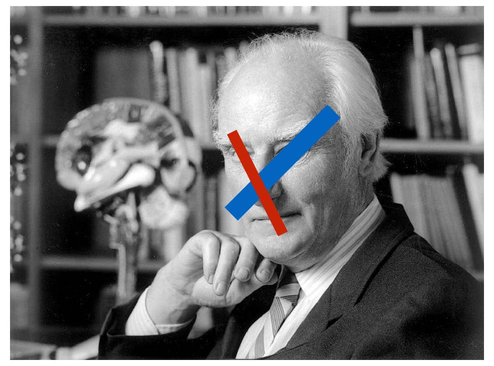

Eccentricity dependence of M in V1, V2, V4, IT

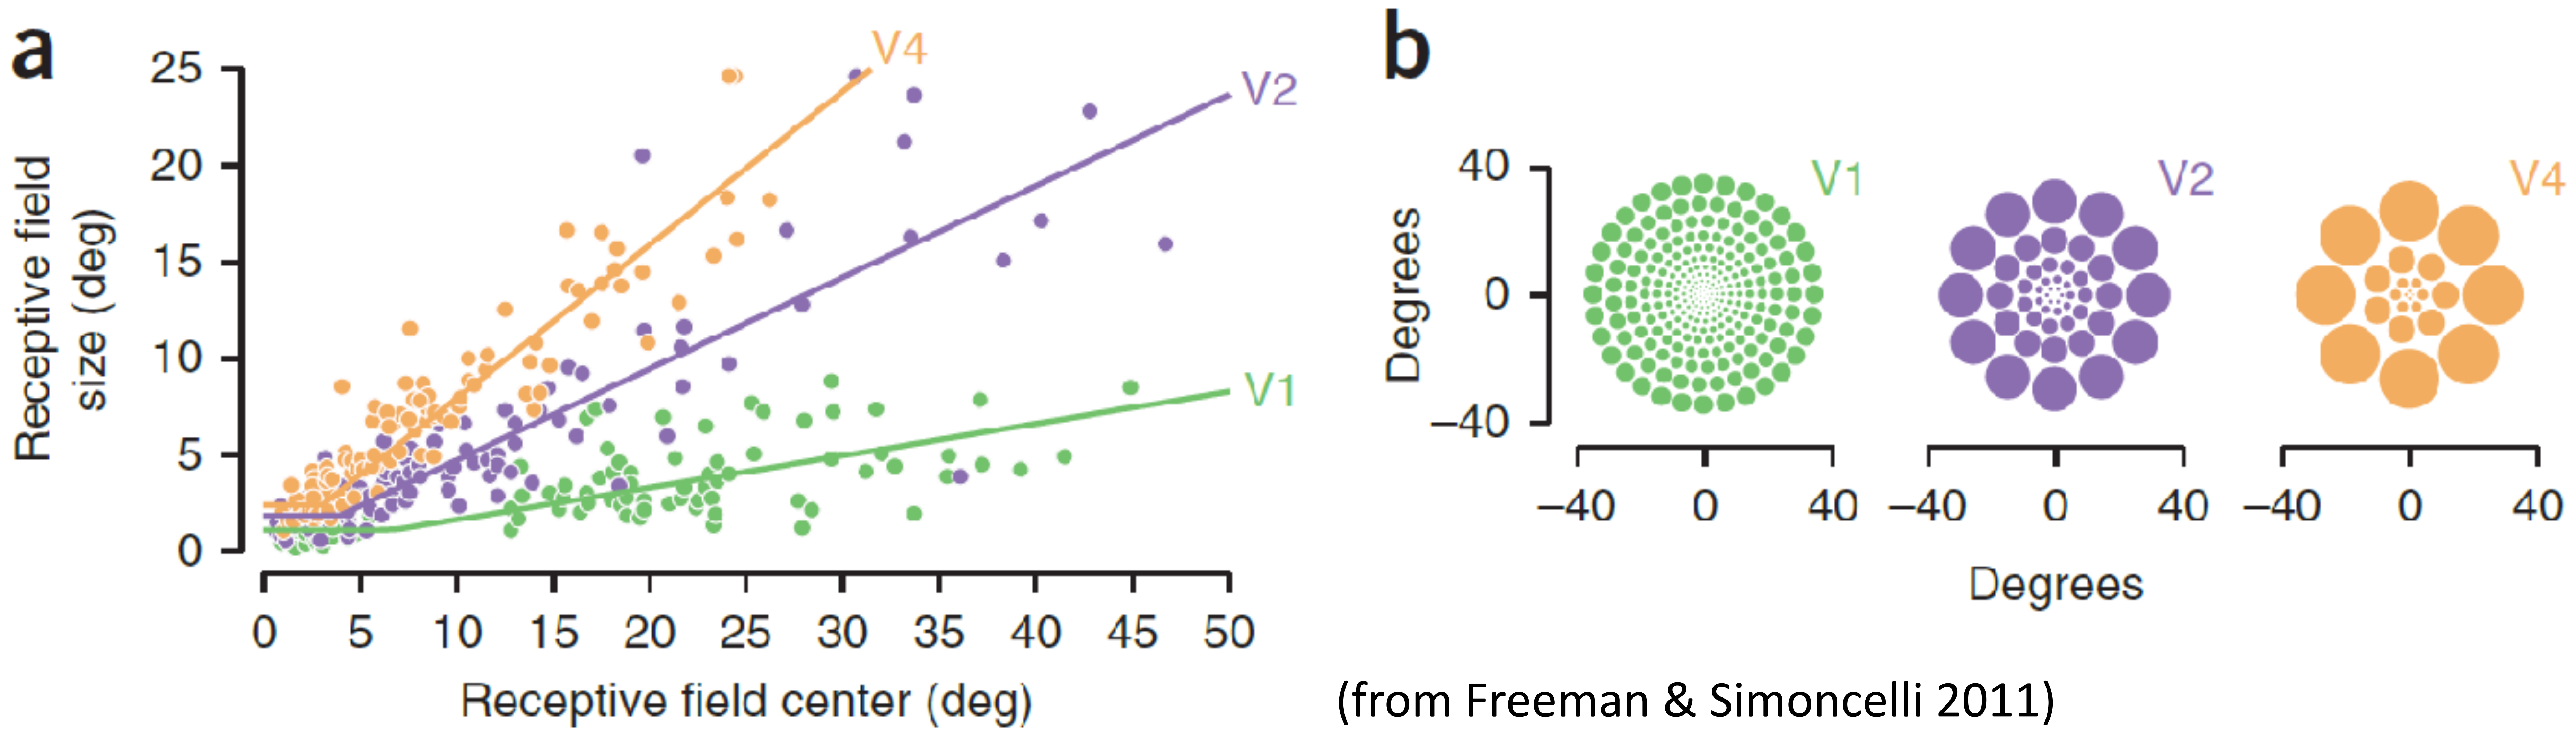

(from Freeman & Simoncelli 2011)

“Prediction” of Anstis observation

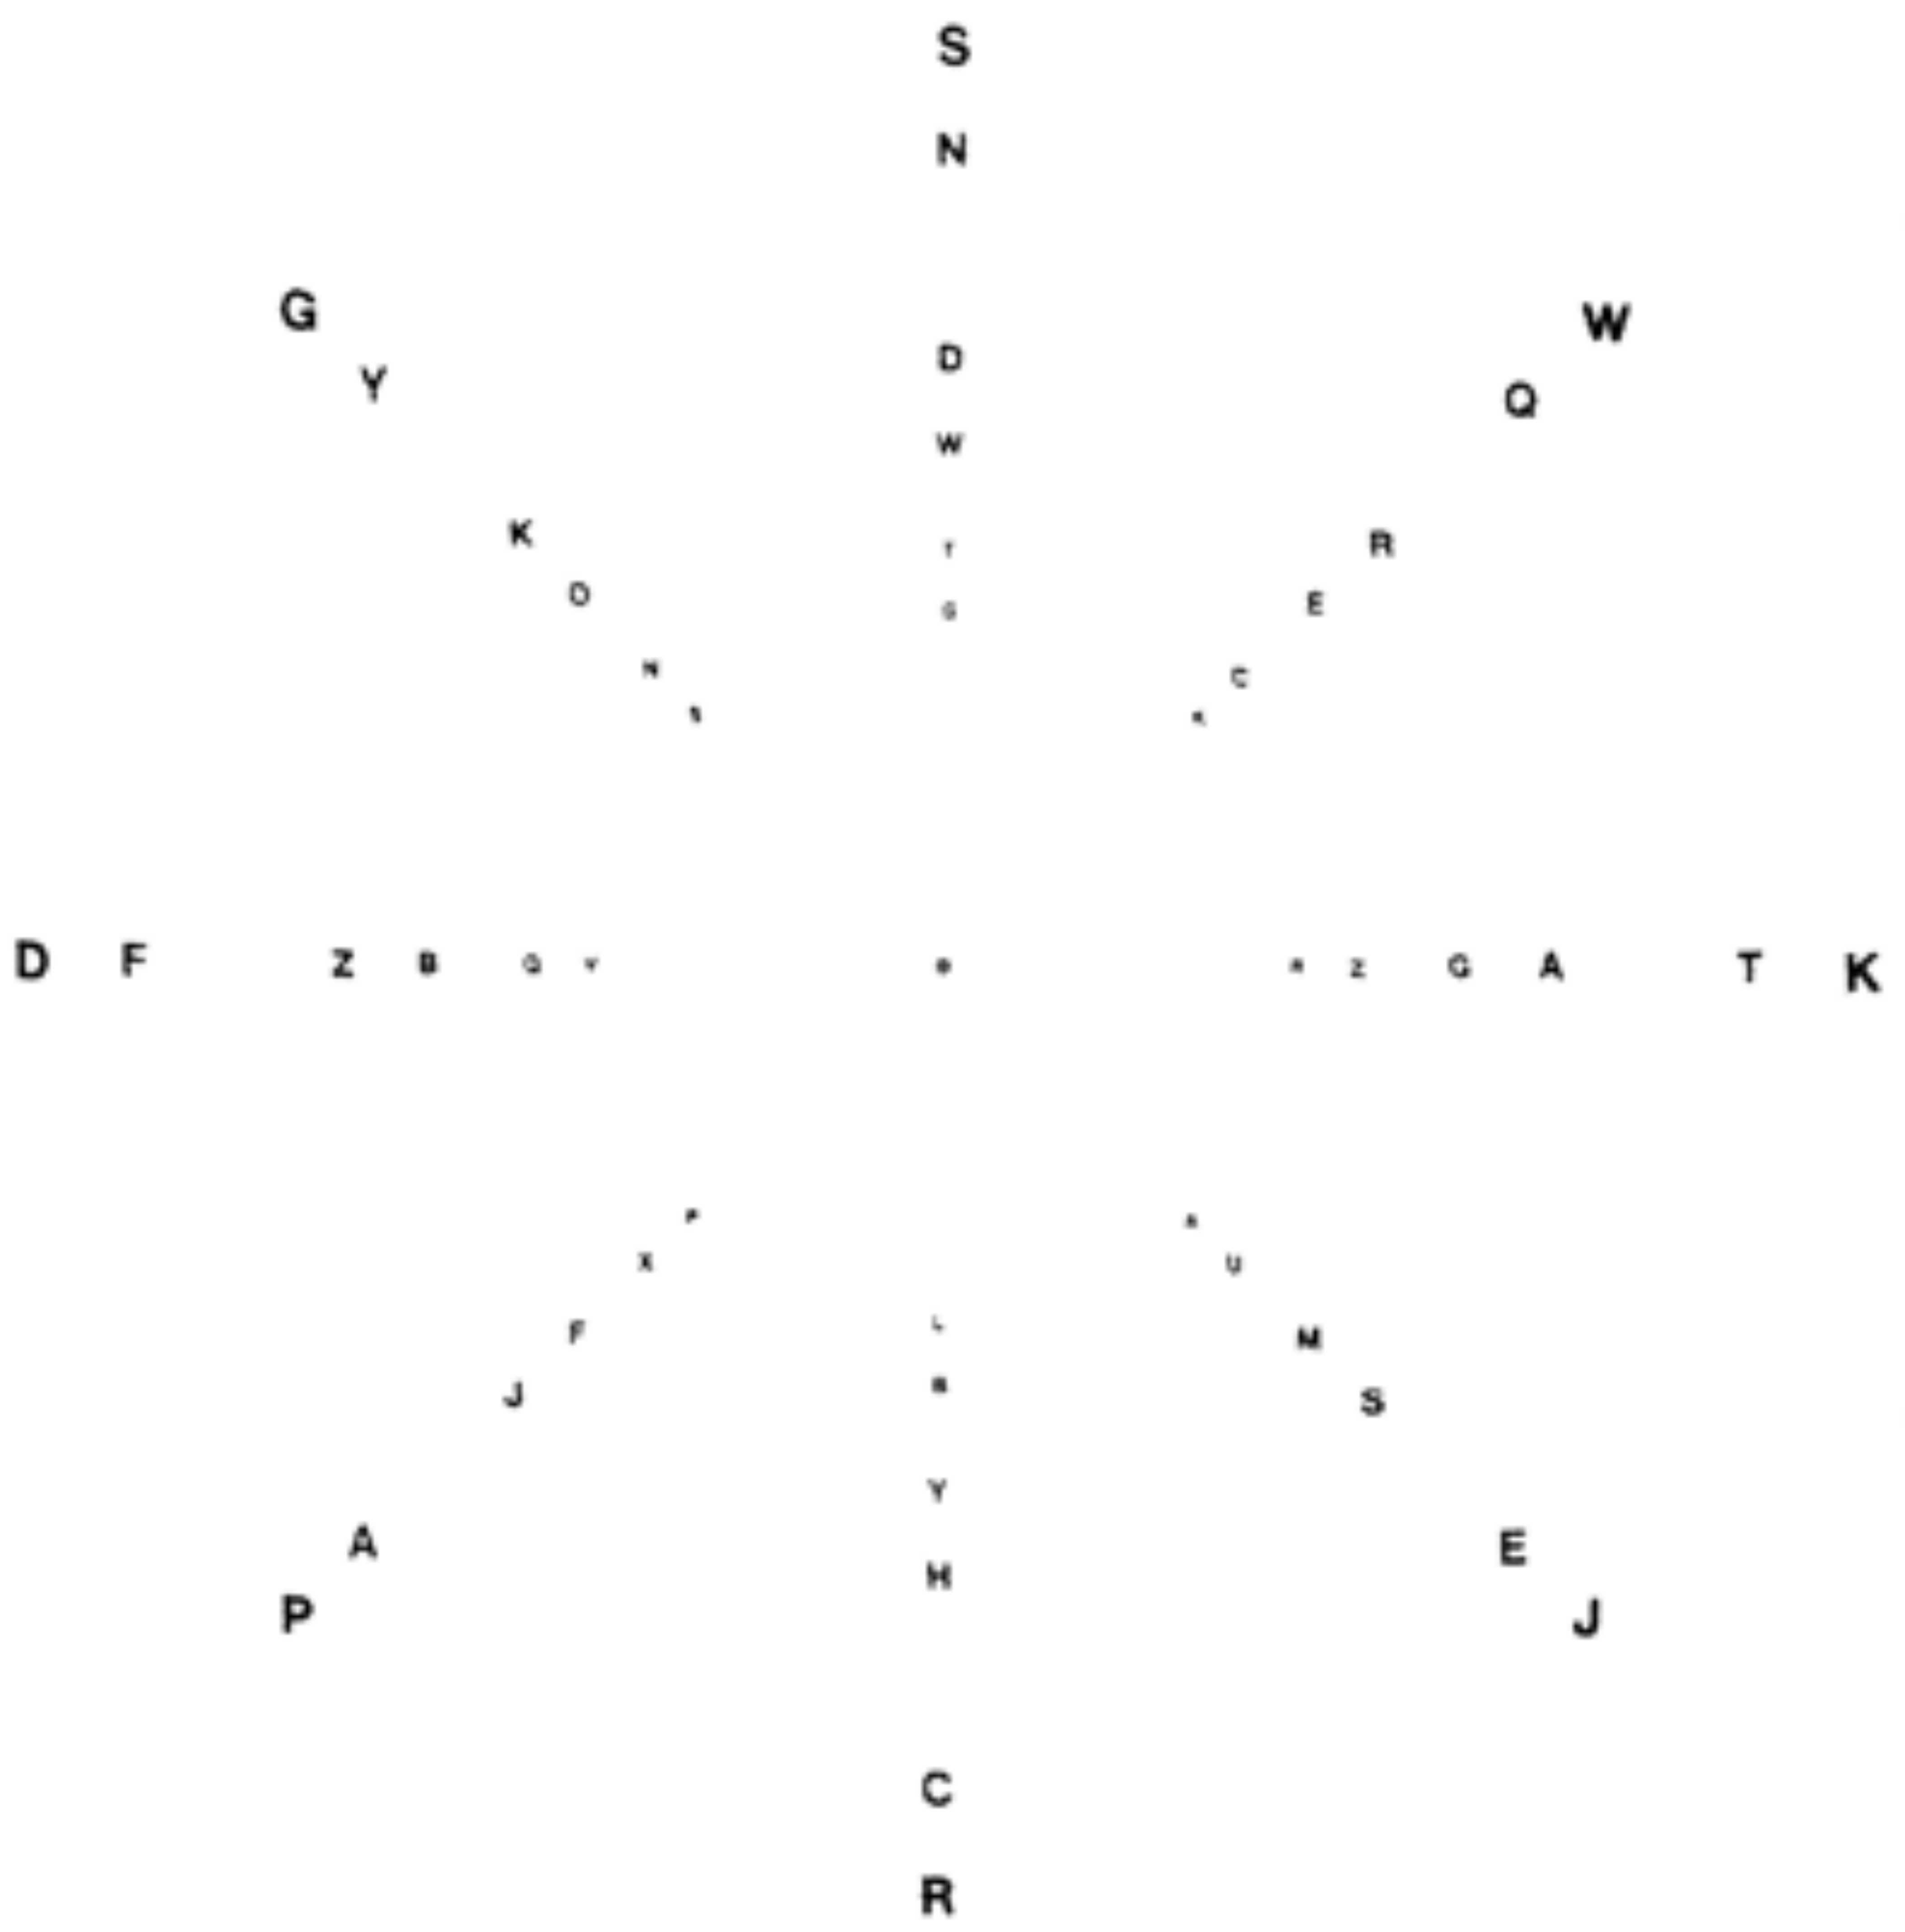

Anstis, 1974

Invariance for new images

Target	아 드 피 뤄 춘 선 머 르 타 예 간 방 우 시 켜
Distractor	마 므 티 뢔 훈 건 다 브 뎌 메 산 랑 은 지 려

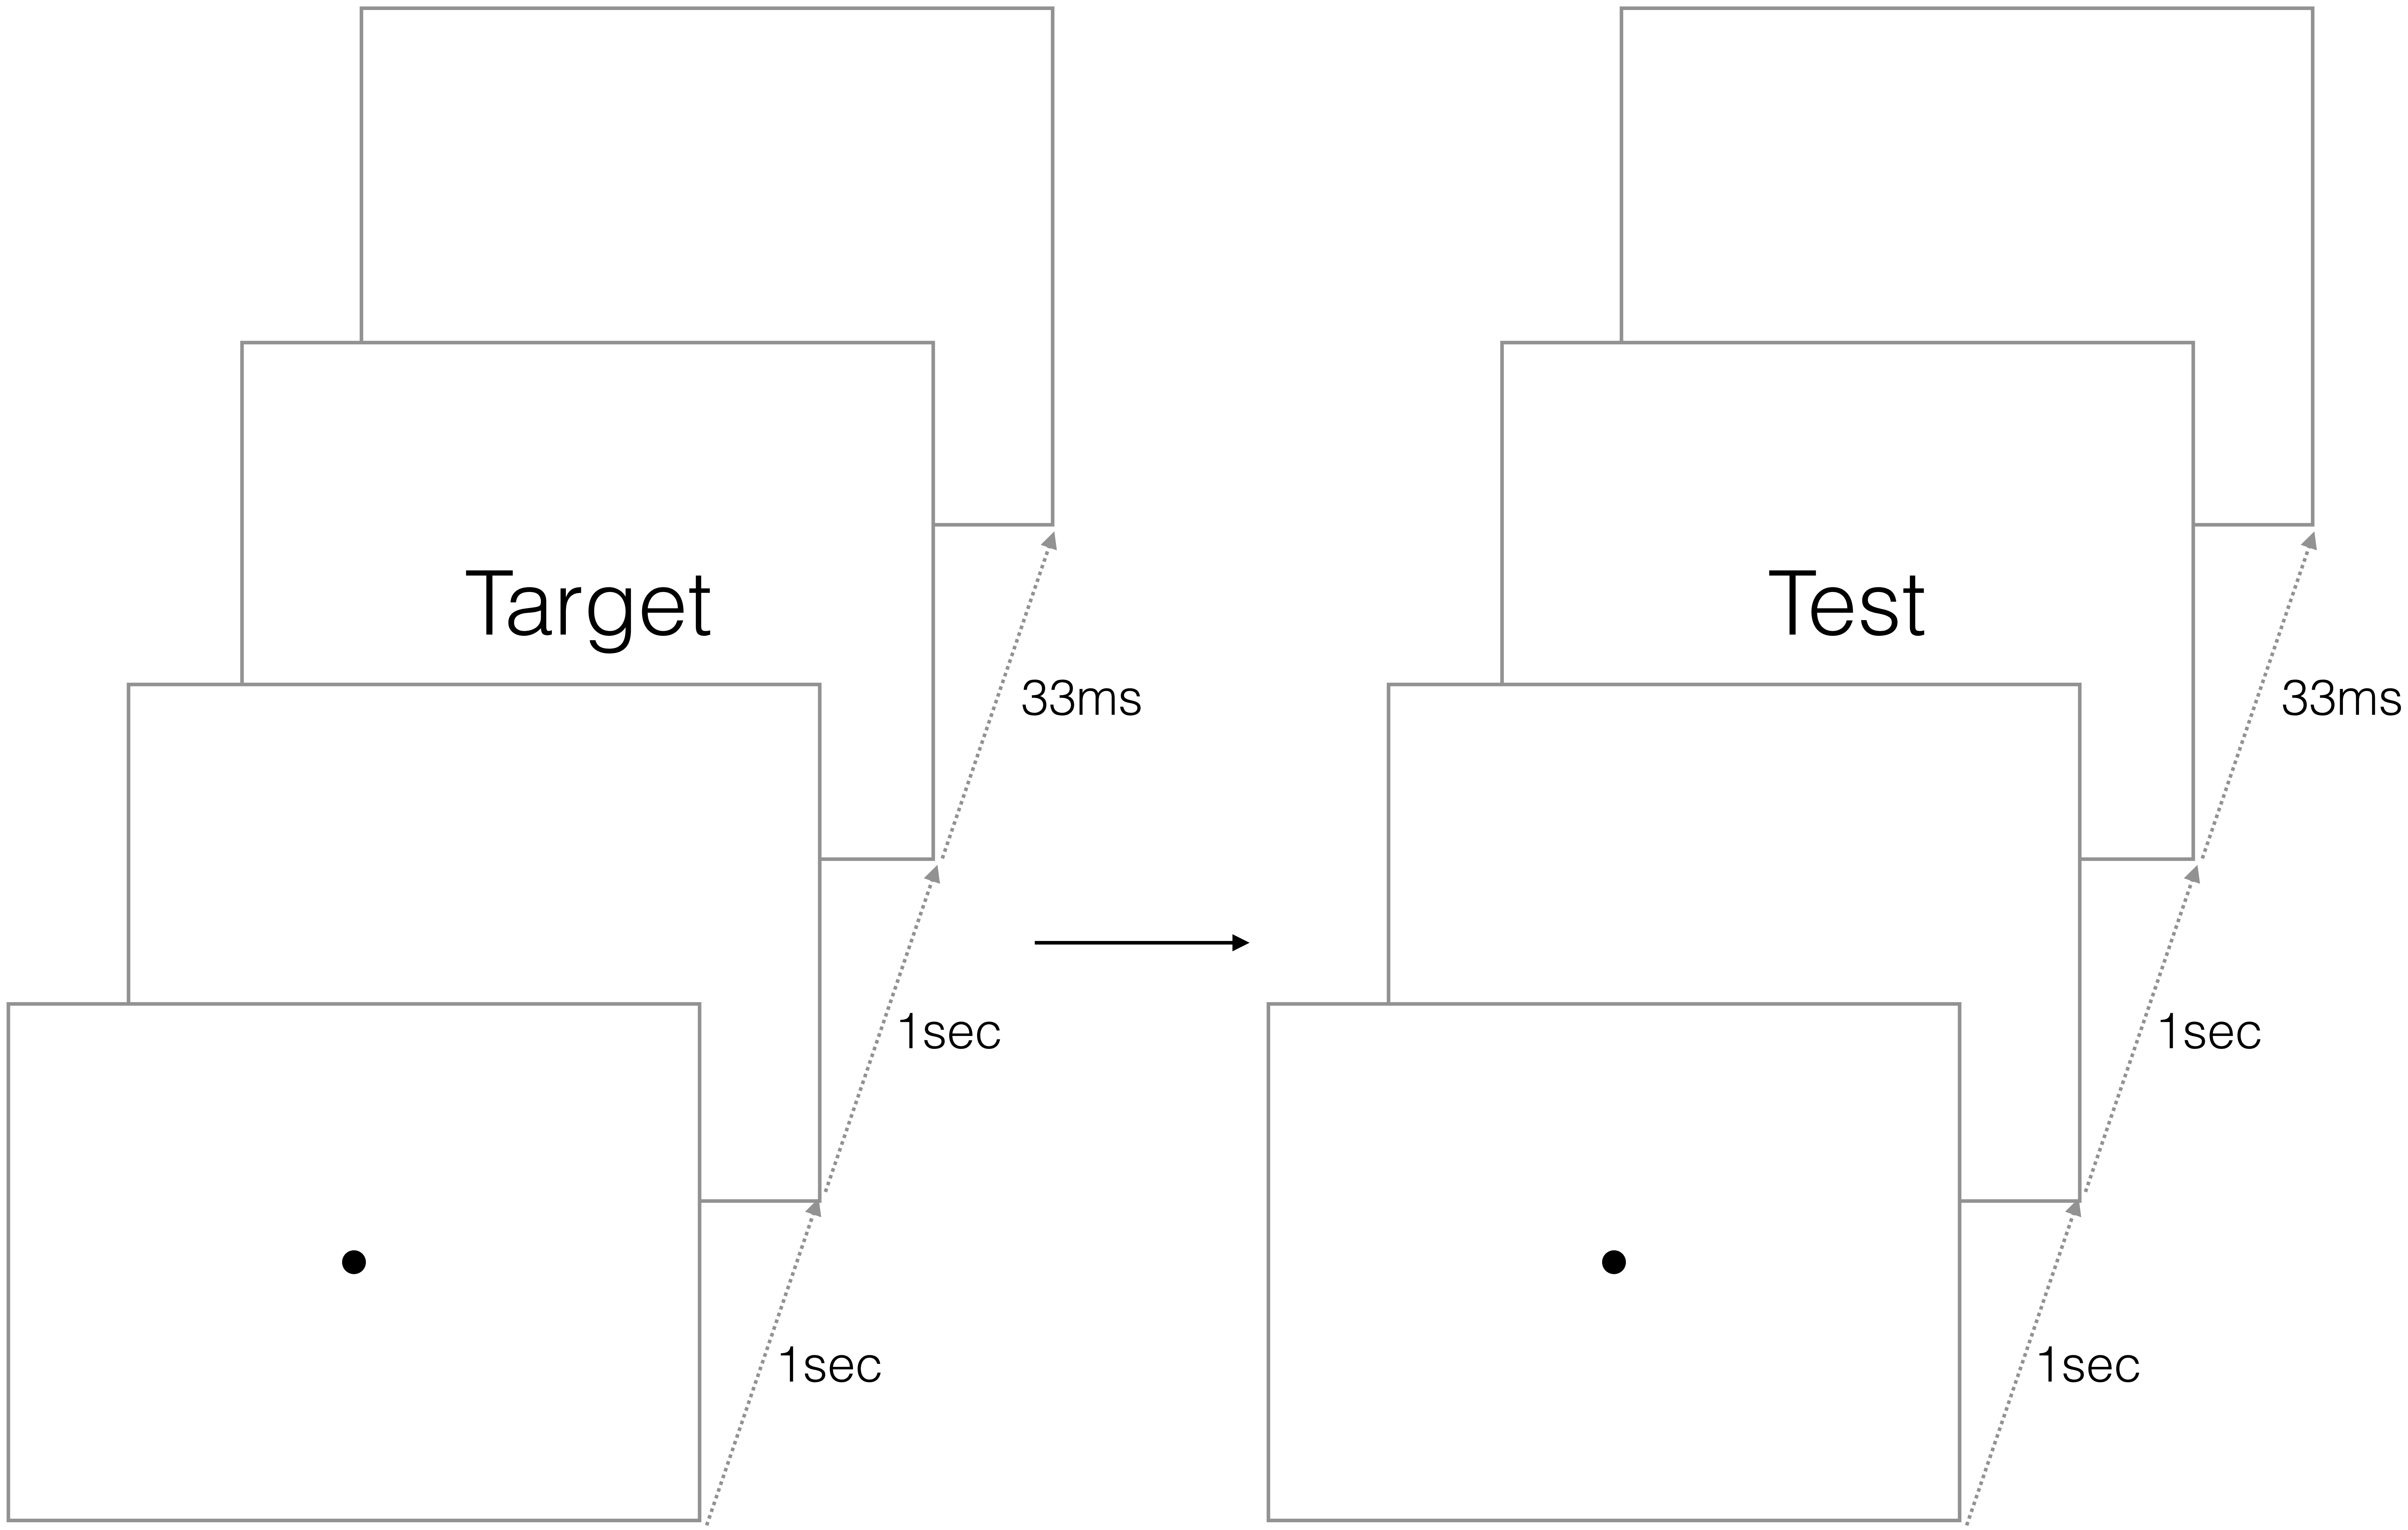

Yena Han, 2017

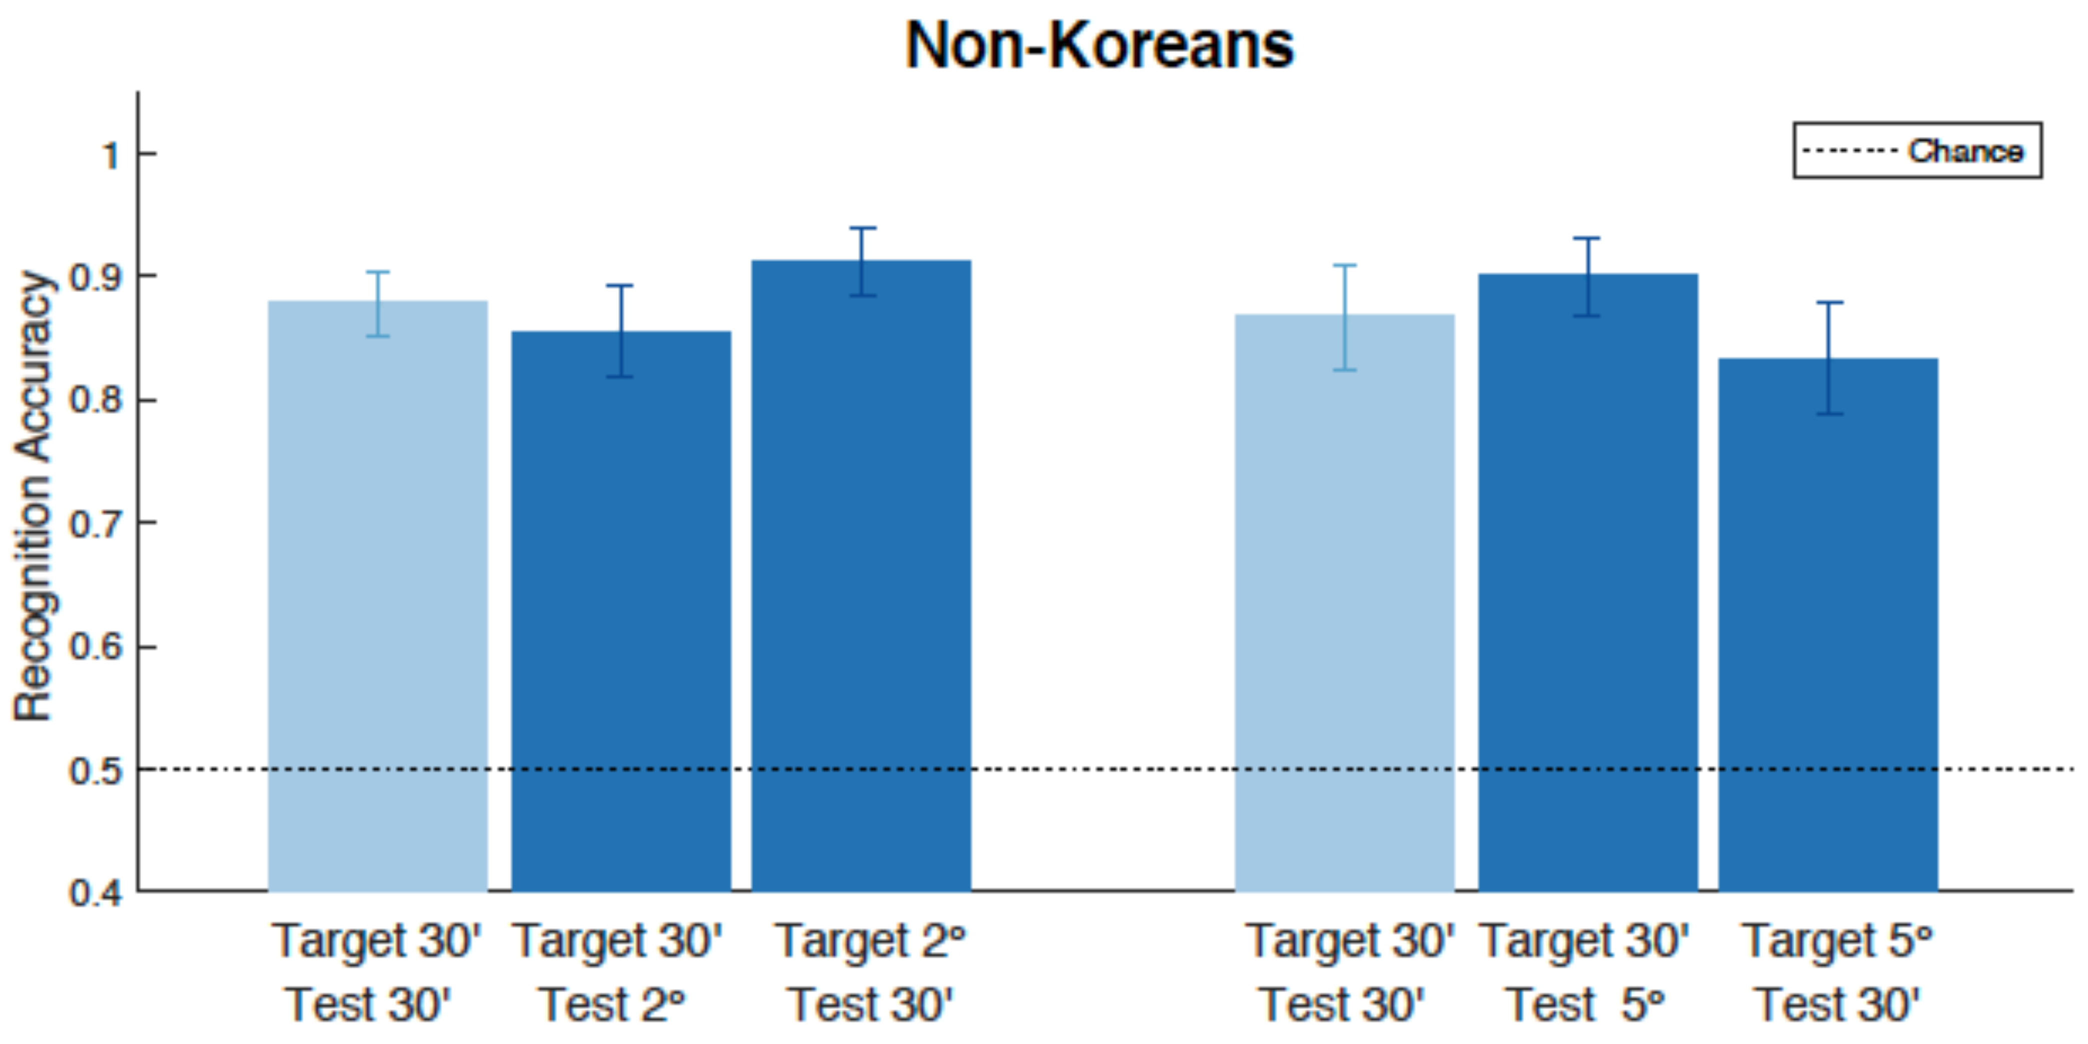

Non-Koreans
Chance
Recognition Accuracy
1
0.9
0.8
0.7
0.6
0.5
0.4
Target 30' Test 30'
Target 30' Test 2°
Target 2° Test 30'
Target 30' Test 30'
Target 30' Test 5°
Target 5° Test 30'

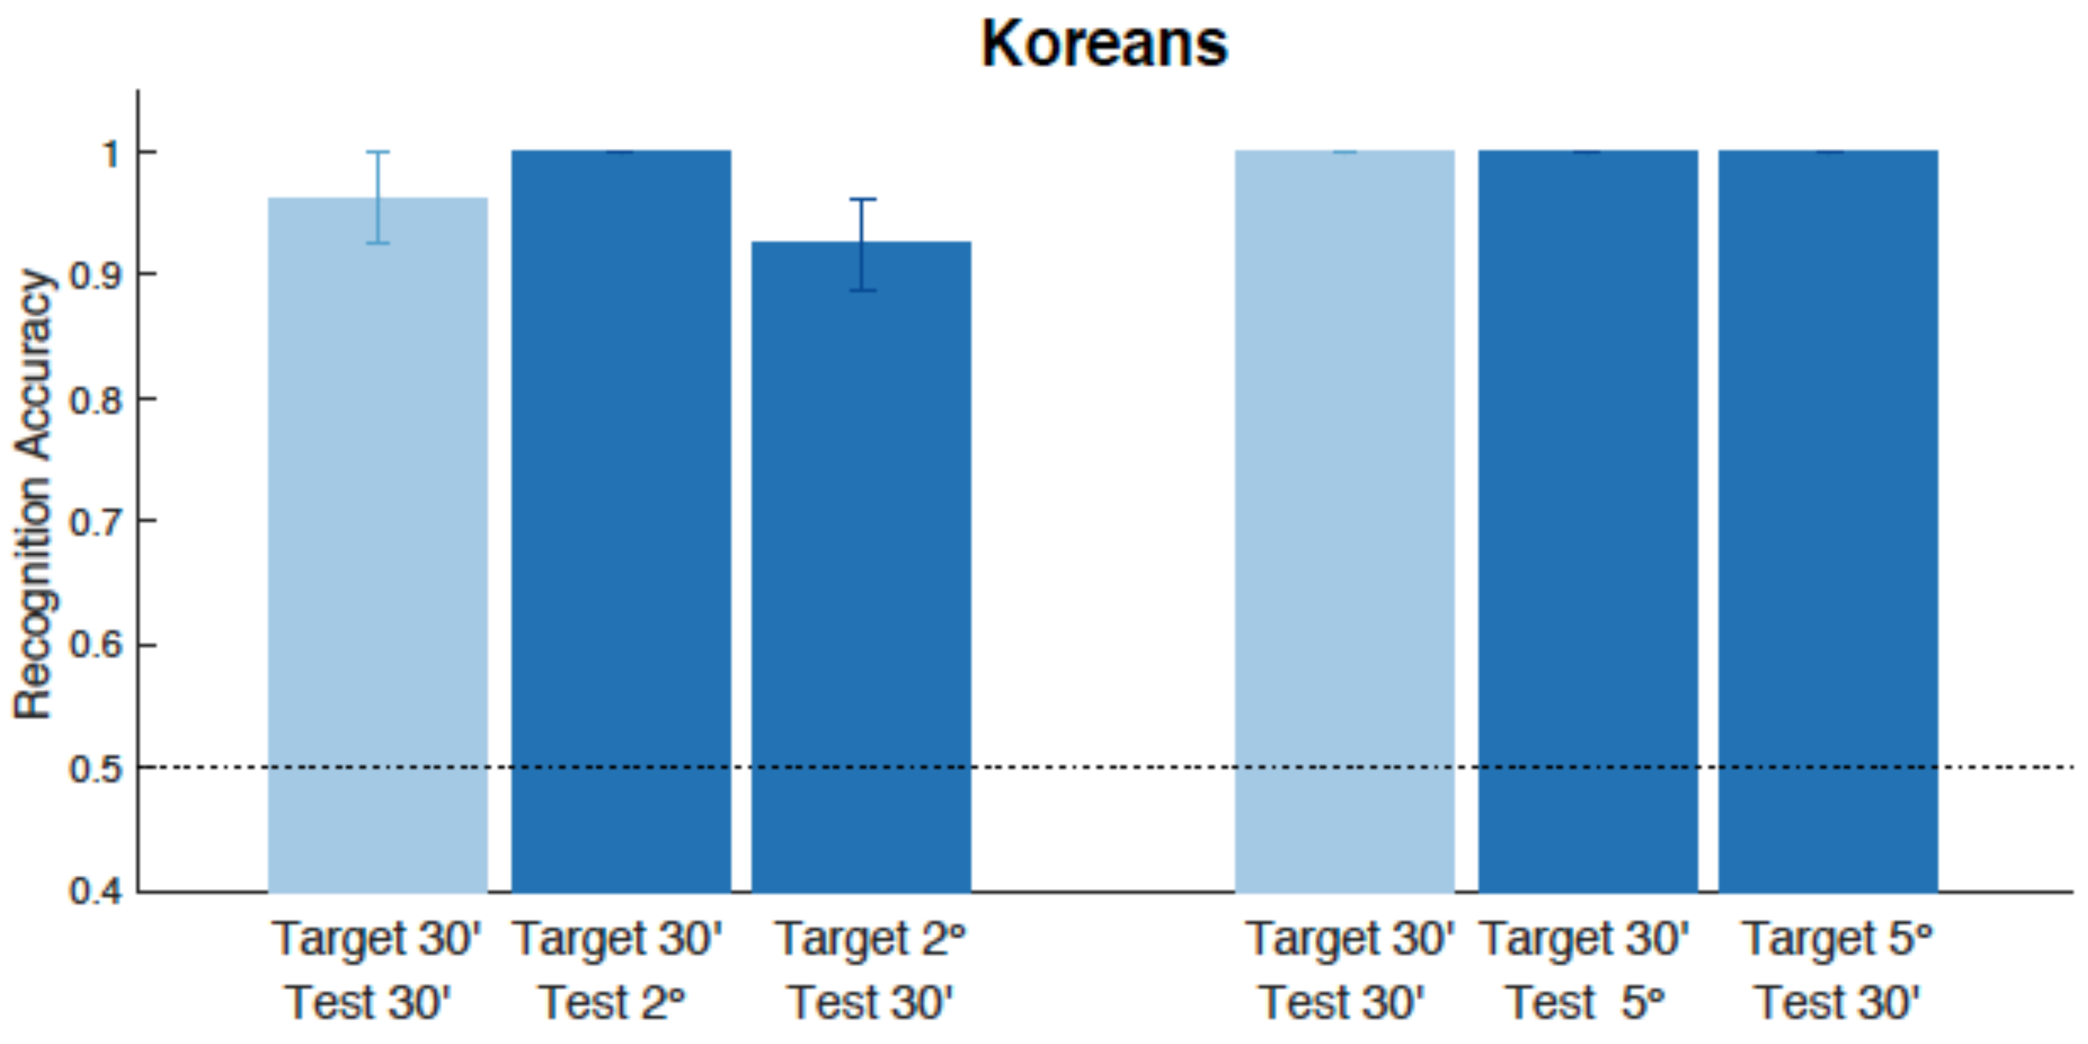

Koreans
Recognition Accuracy
1
0.9
0.8
0.7
0.6
0.5
0.4
Target 30' Test 30'
Target 30' Test 2°
Target 2° Test 30'
Target 30' Test 30'
Target 30' Test 5°
Target 5° Test 30'

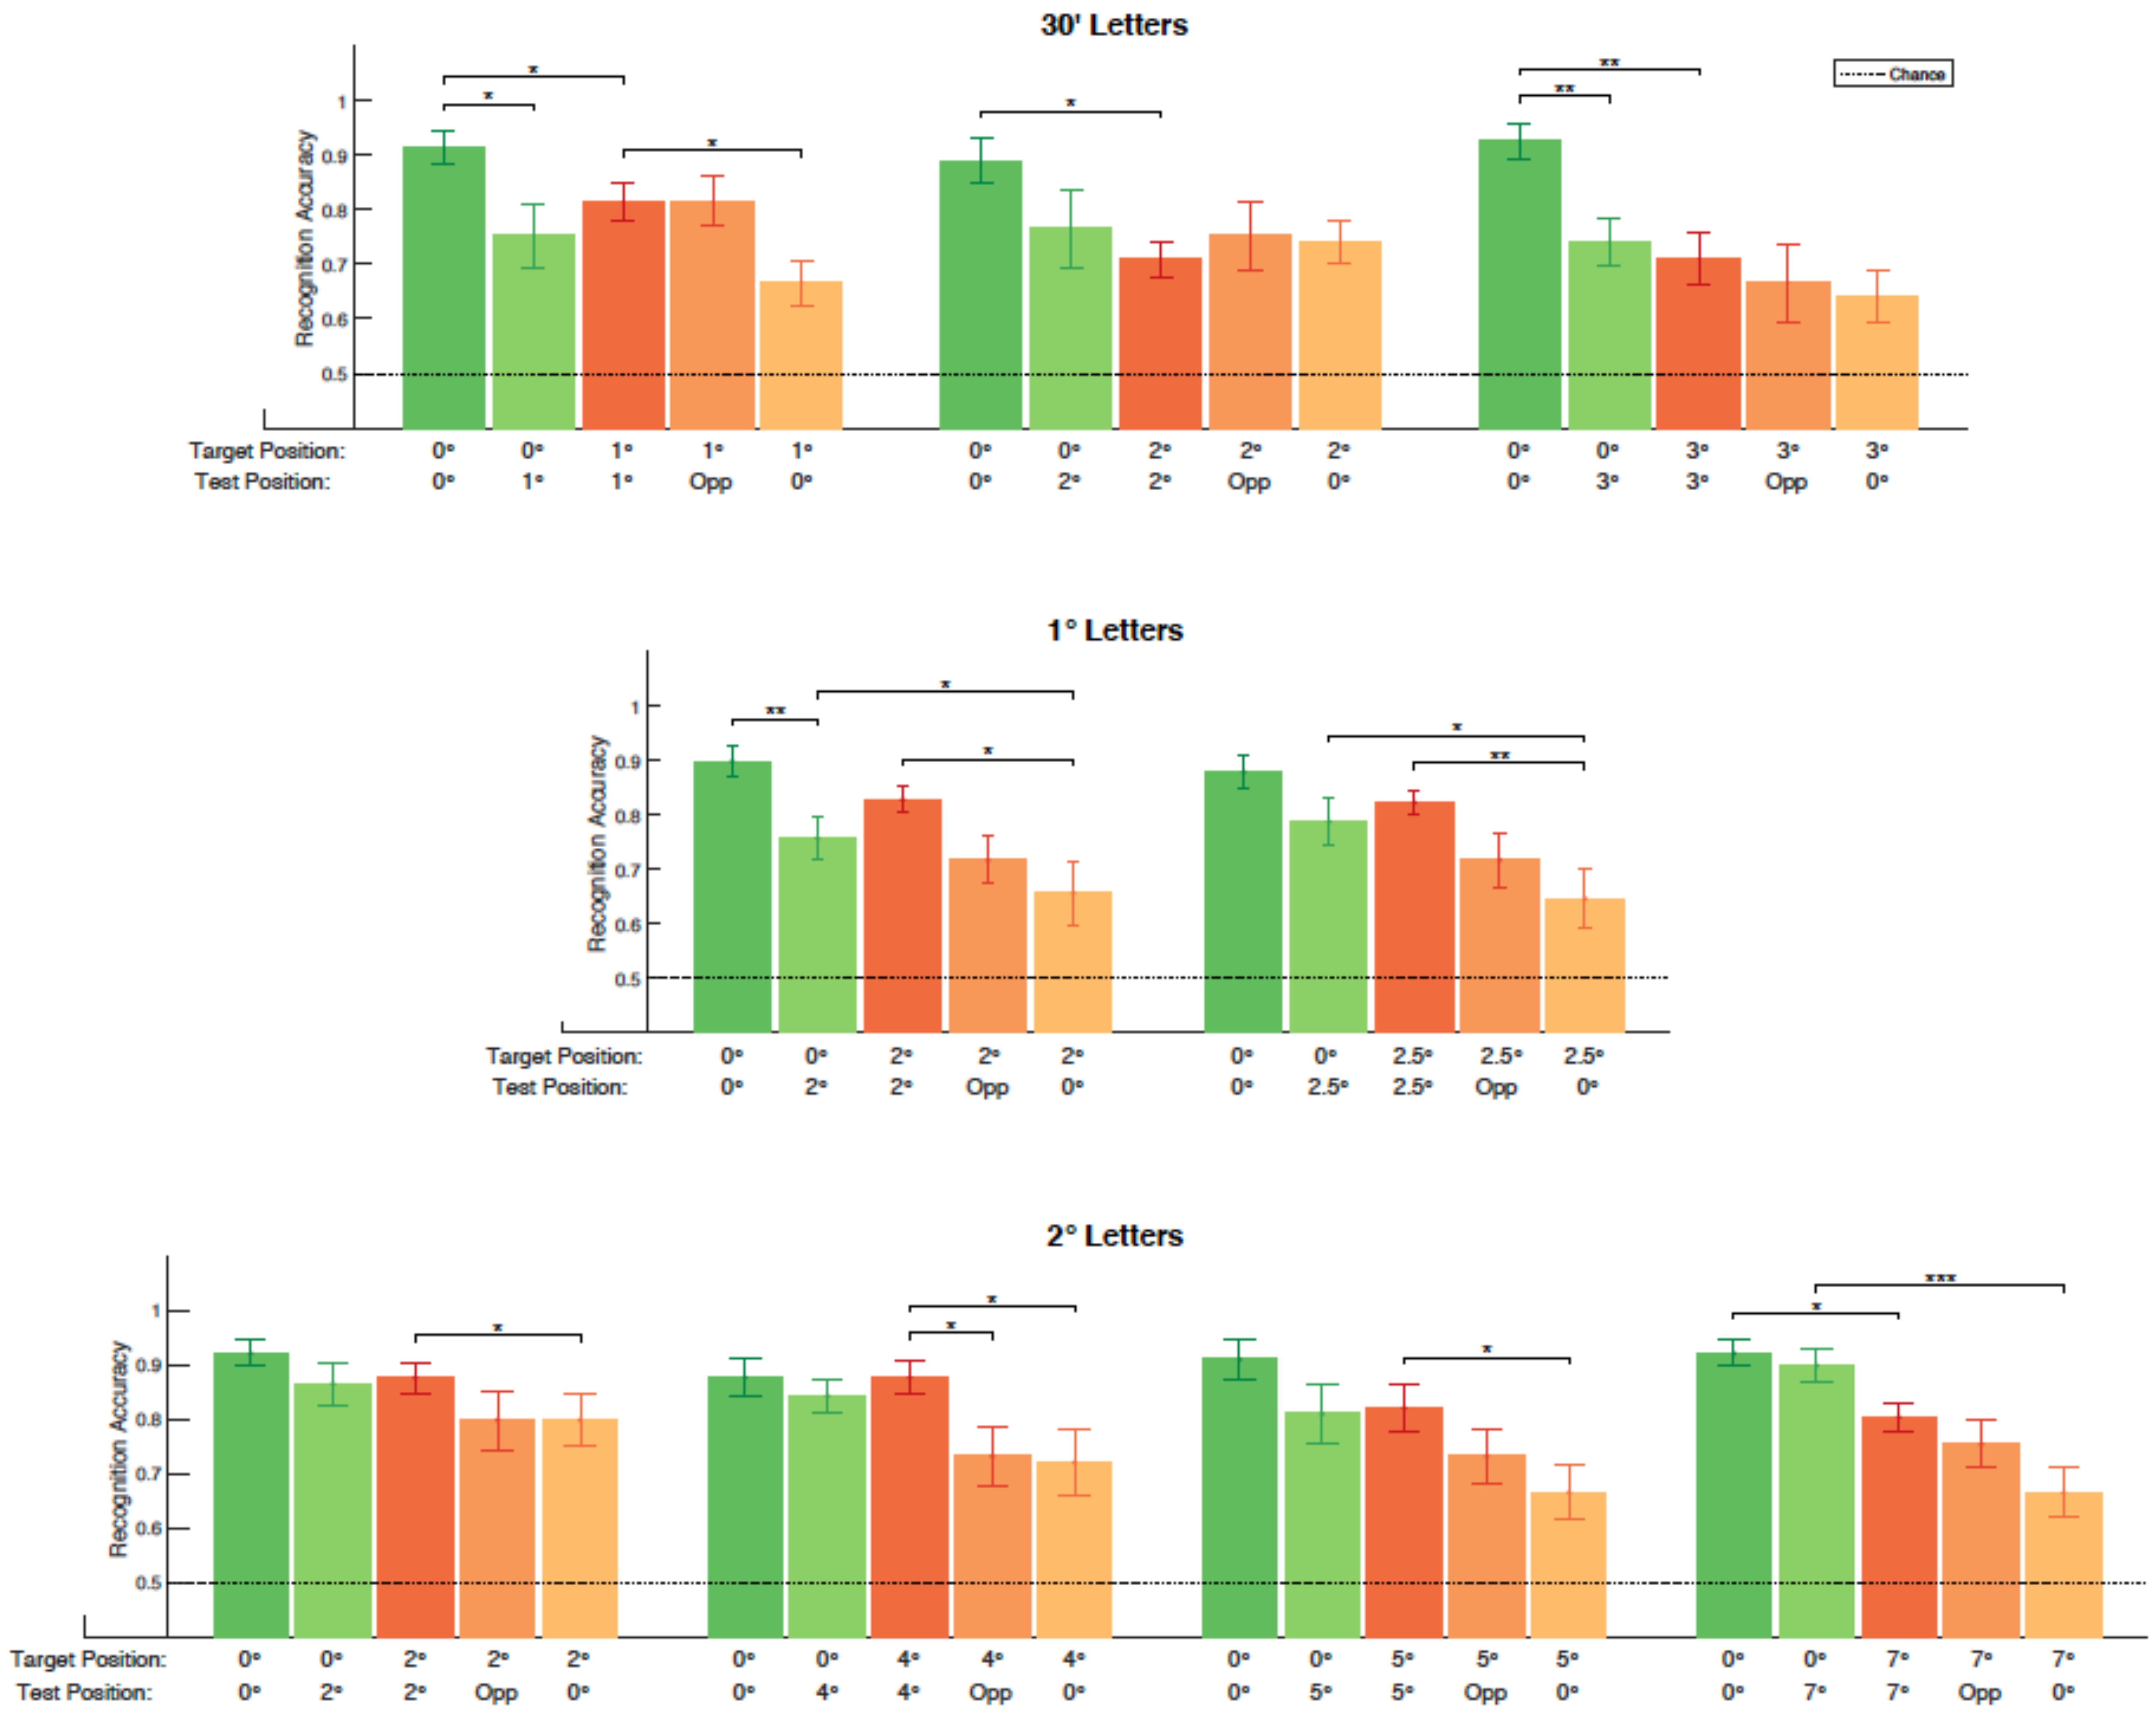
30' Letters
Chance
Recognition Accuracy
Target Position: 0° 0° 1° 1° 1° 0° 0° 2° 2° 2° 0° 0° 3° 3° 3°
Test Position: 0° 1° 1° Opp 0° 0° 2° 2° Opp 0° 0° 3° 3° Opp 0°
1° Letters
Recognition Accuracy
Target Position: 0° 0° 2° 2° 2° 0° 0° 2.5° 2.5° 2.5°
Test Position: 0° 2° 2° Opp 0° 0° 2.5° 2.5° Opp 0°
2° Letters
Recognition Accuracy
Target Position: 0° 0° 2° 2° 2° 0° 0° 4° 4° 4° 0° 0° 5° 5° 5° 0° 0° 7° 7° 7°
Test Position: 0° 2° 2° Opp 0° 0° 4° 4° Opp 0° 0° 5° 5° Opp 0° 0° 7° 7° Opp 0°

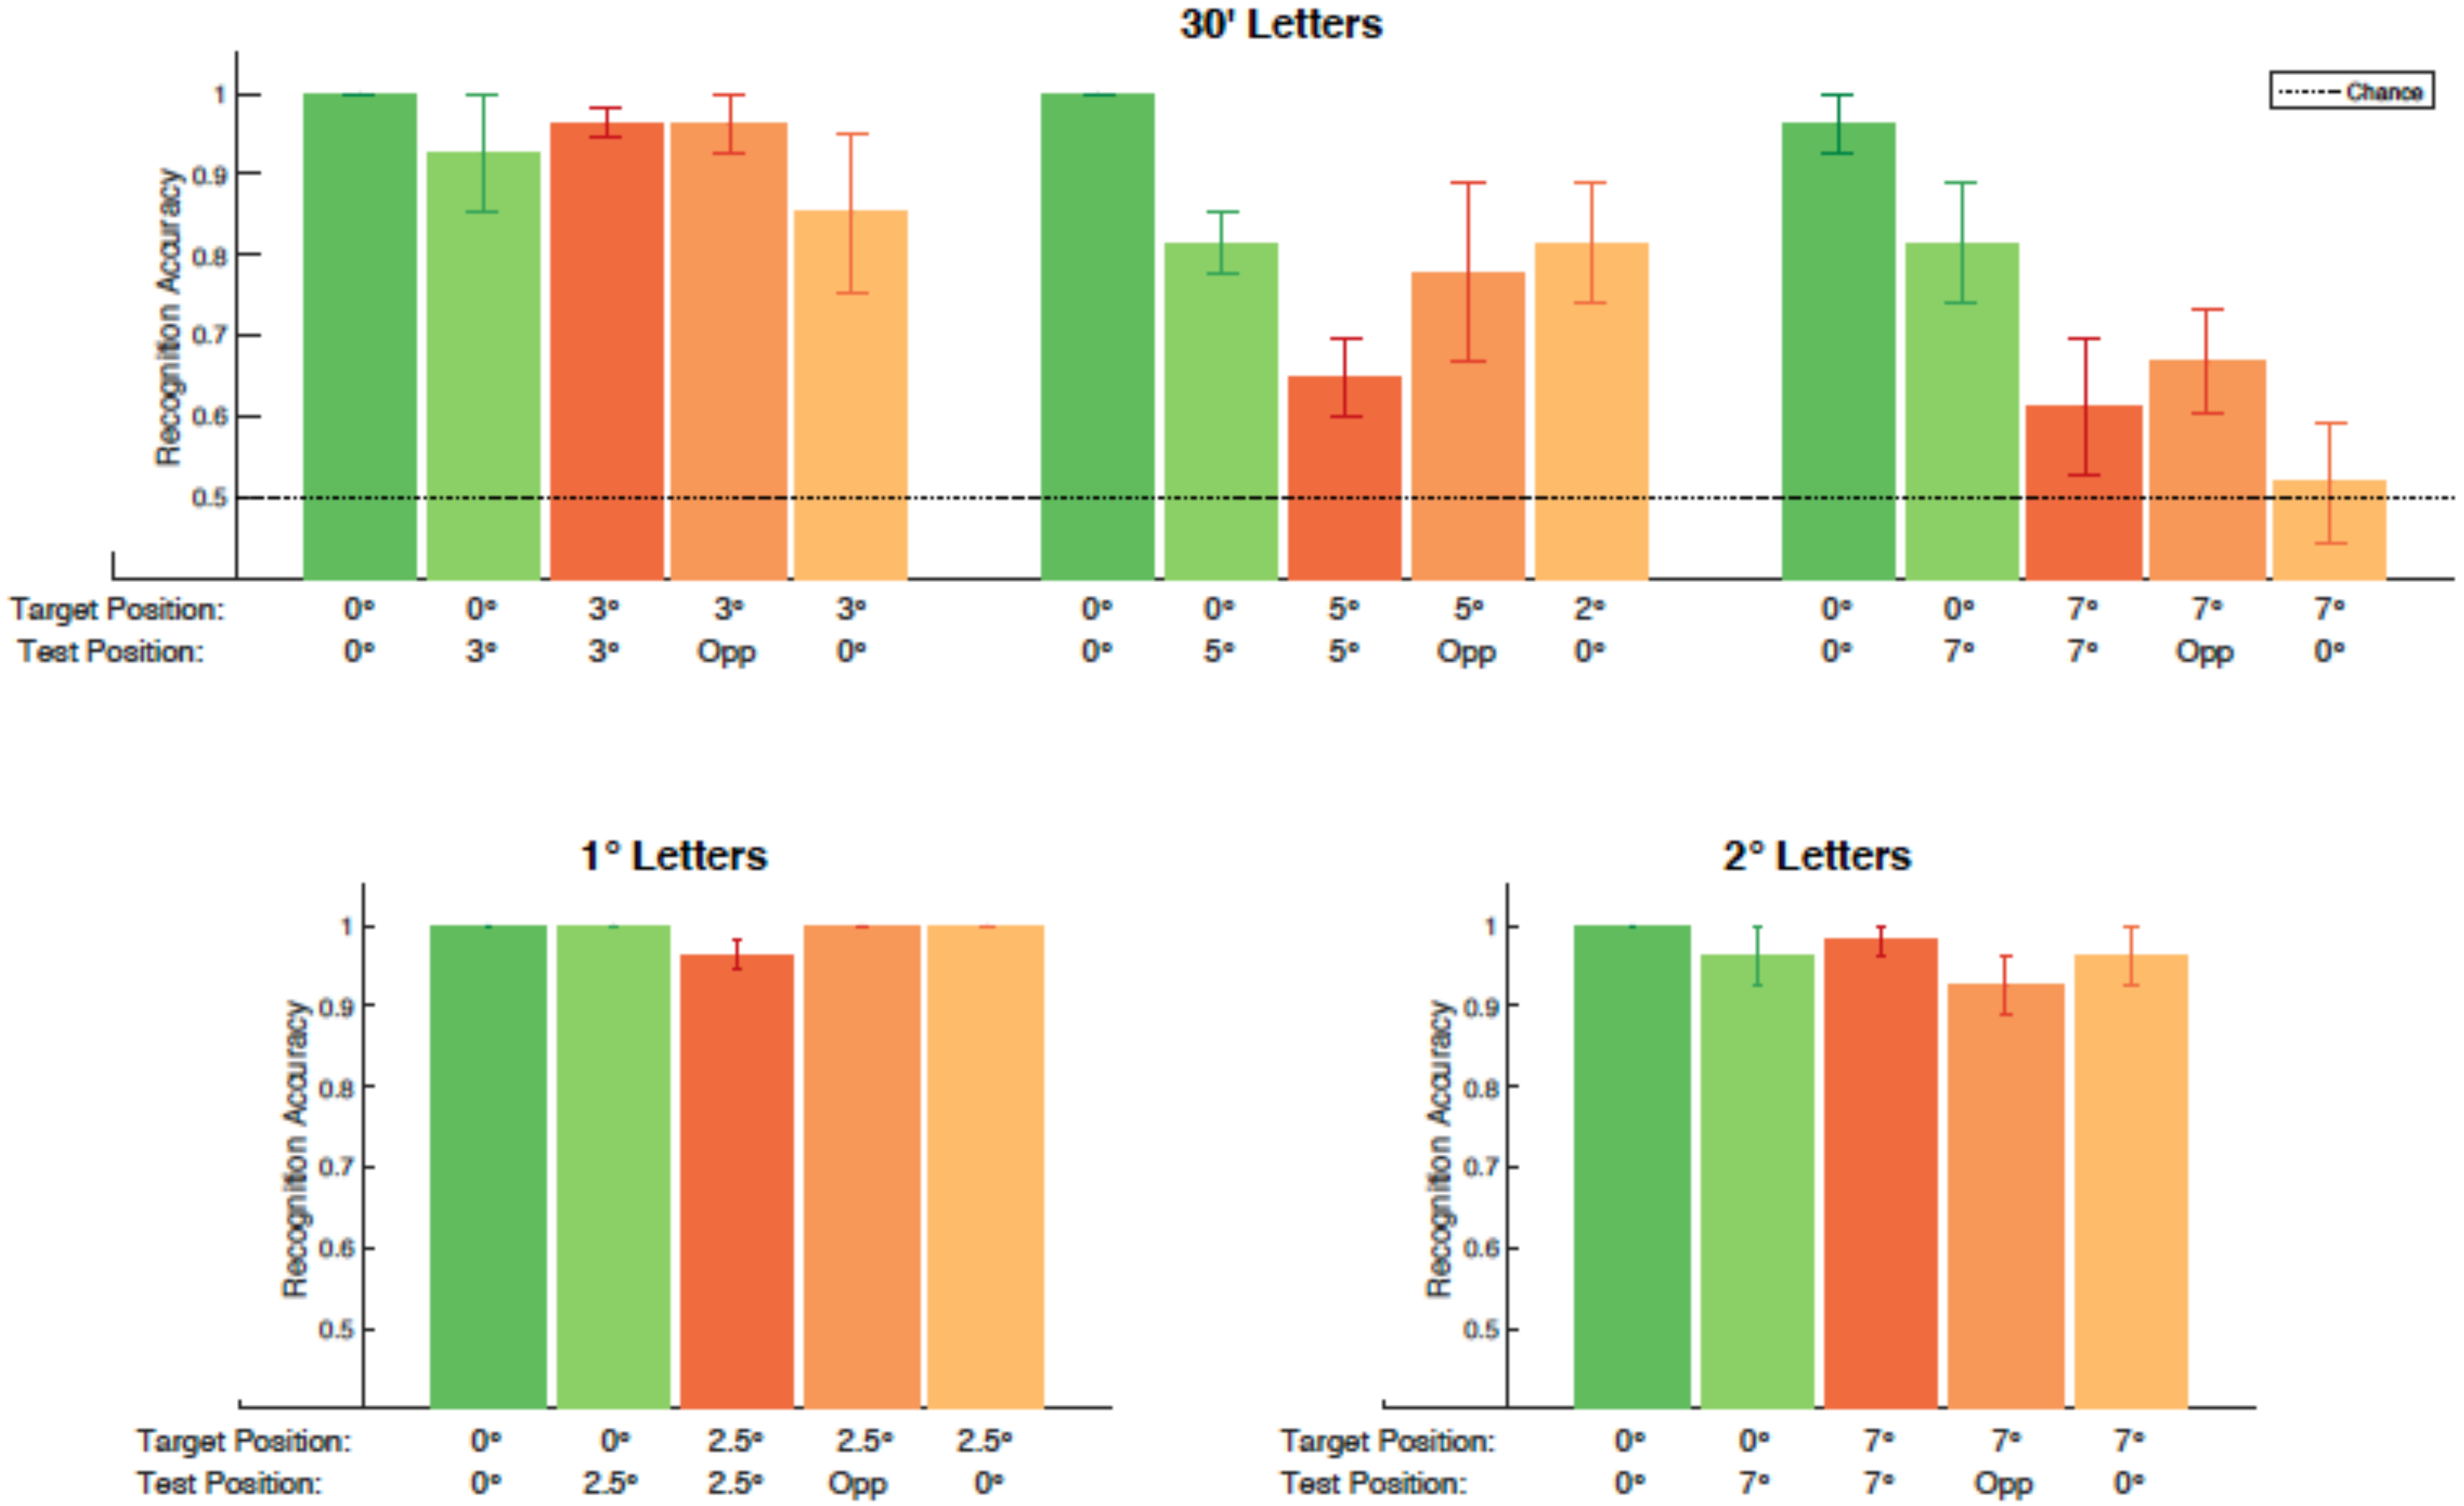
30' Letters
Chance
Recognition Accuracy
1
0.9
0.8
0.7
0.6
0.5
Target Position: 0° 0° 3° 3° 3° 0° 0° 5° 5° 2° 0° 0° 7° 7° 7°
Test Position: 0° 3° 3° Opp 0° 0° 5° 5° Opp 0° 0° 7° 7° Opp 0°
1° Letters
Recognition Accuracy
Target Position: 0° 0° 2.5° 2.5° 2.5°
Test Position: 0° 2.5° 2.5° Opp 0°
2° Letters
Recognition Accuracy
Target Position: 0° 0° 7° 7° 7°
Test Position: 0° 7° 7° Opp 0°

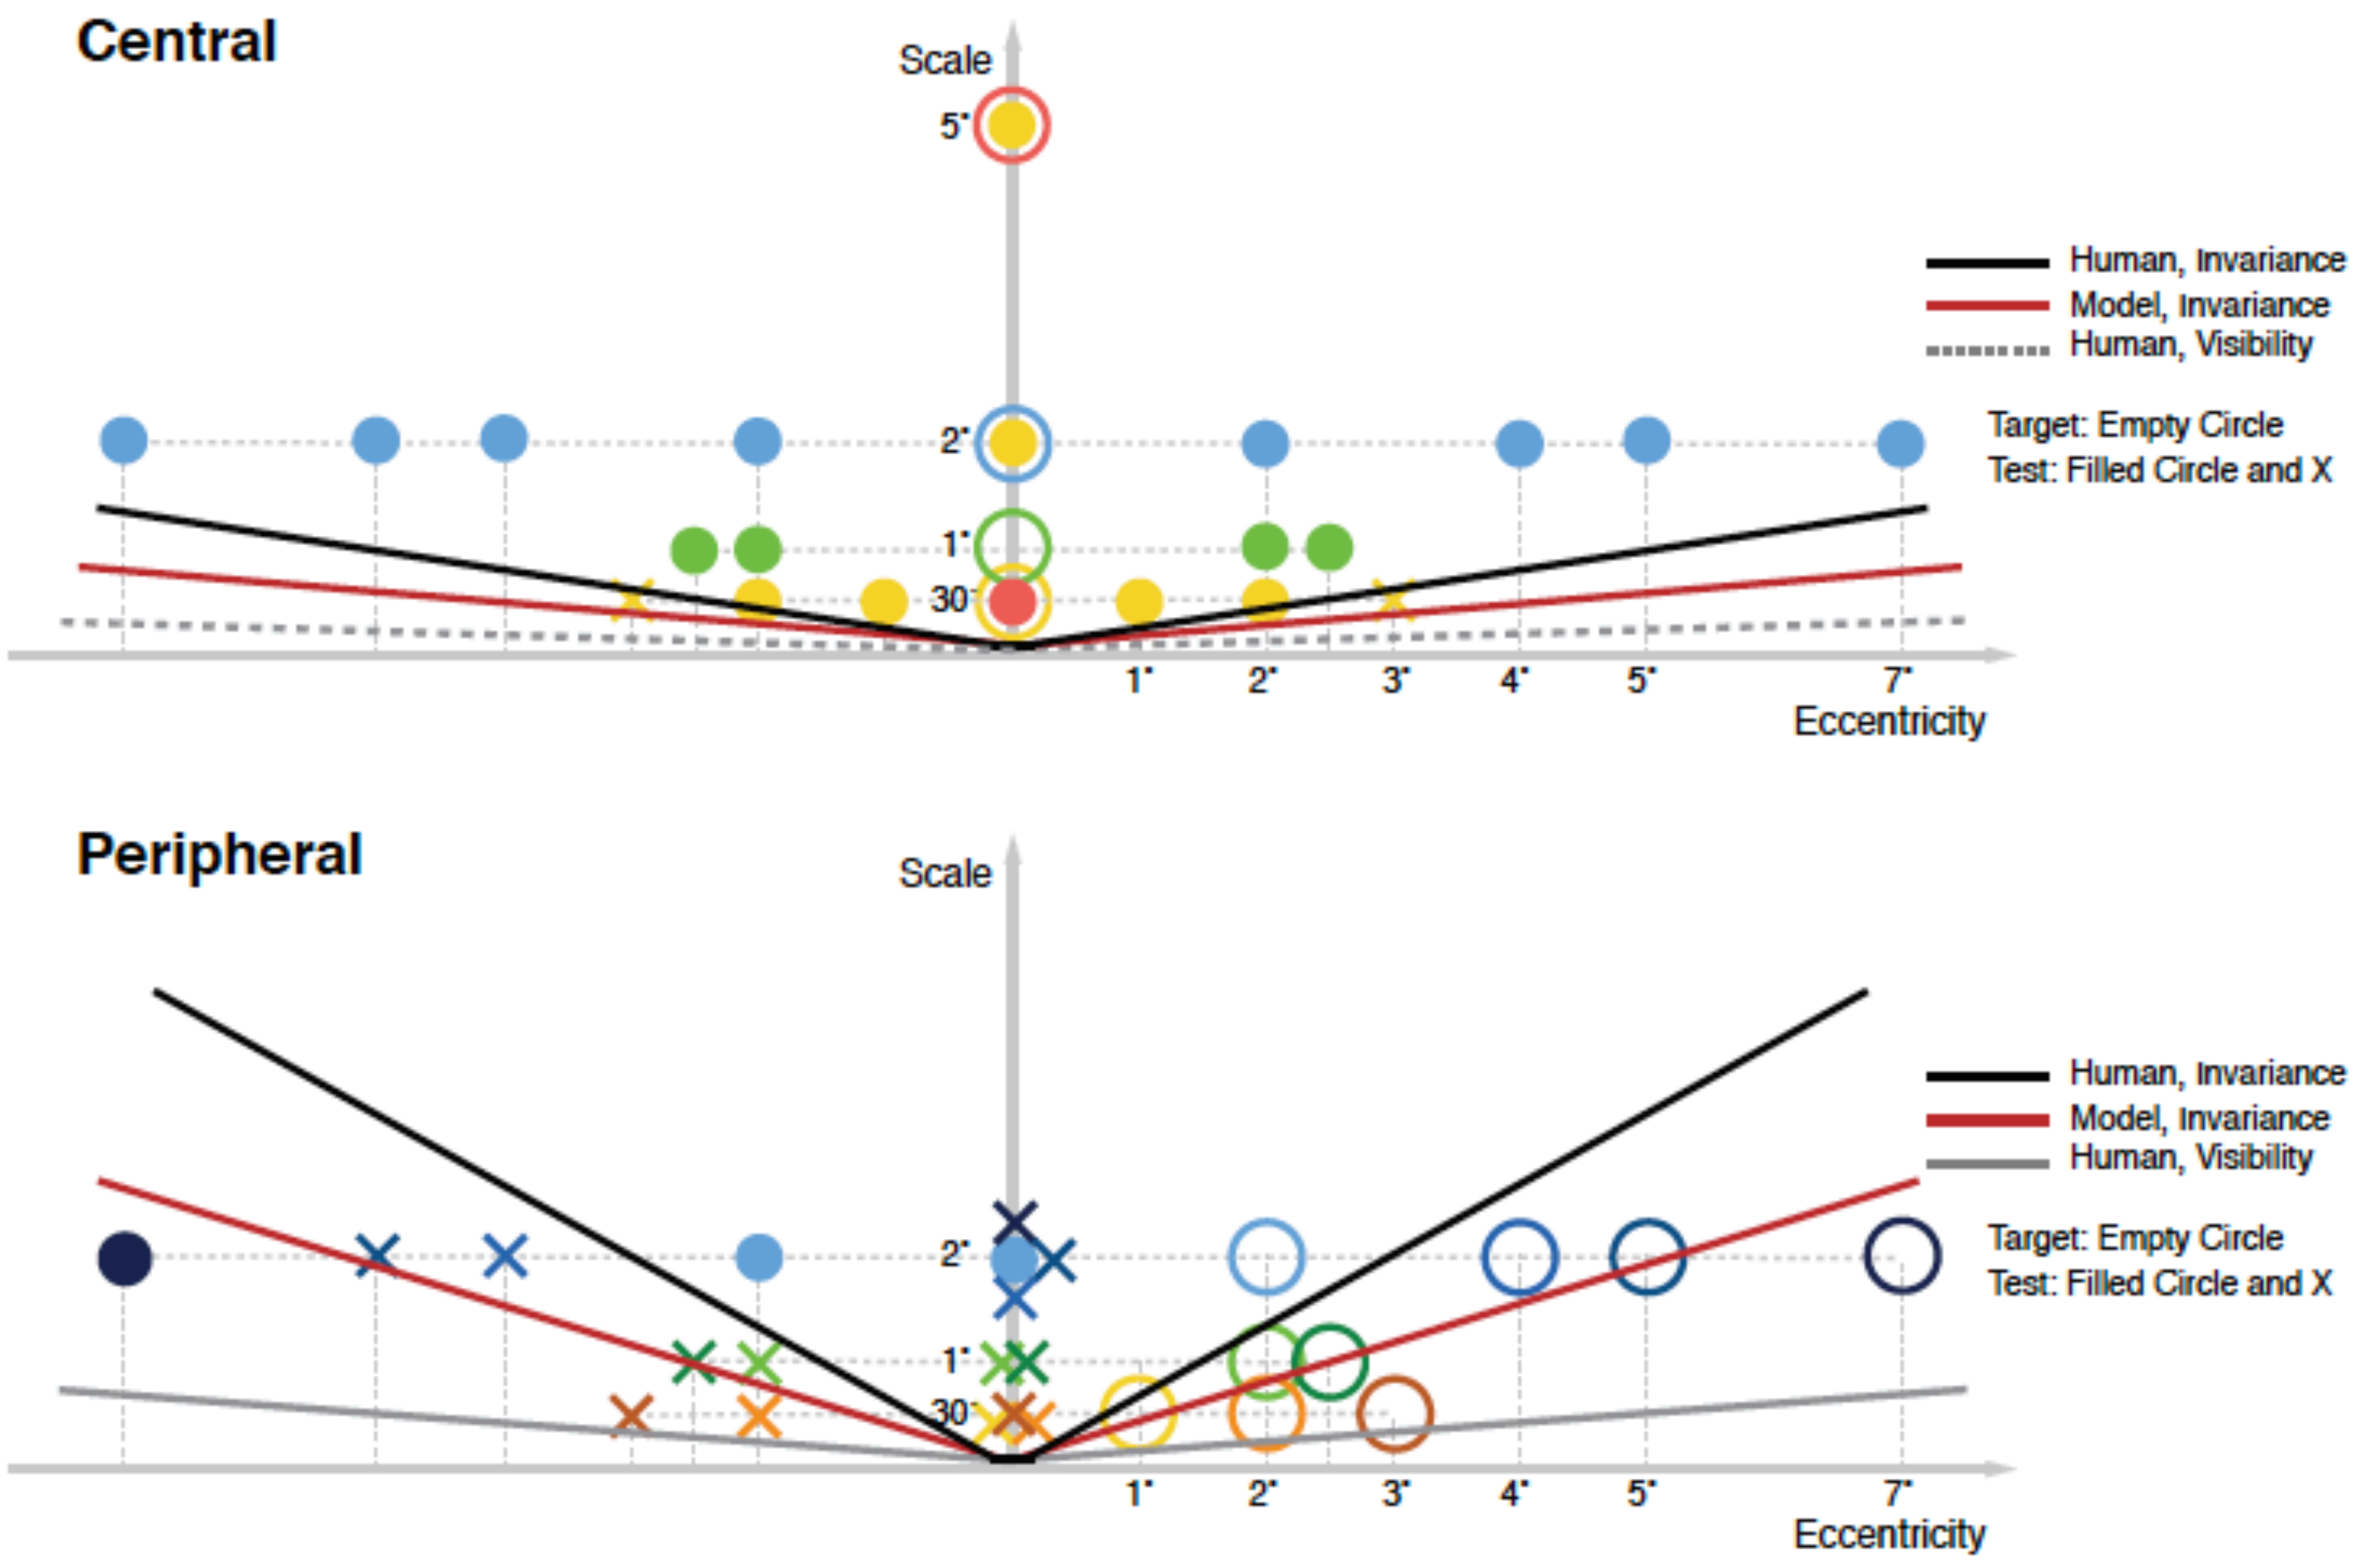

Central
Scale
5°
2°
1°
30'
1°
2°
3°
4°
5°
7°
Eccentricity
Human, Invariance
Model, Invariance
Human, Visibility
Target: Empty Circle
Test: Filled Circle and X
Peripheral
Scale
2°
1°
30'
1°
2°
3°
4°
5°
7°
Eccentricity
Human, Invariance
Model, Invariance
Human, Visibility
Target: Empty Circle
Test: Filled Circle and X

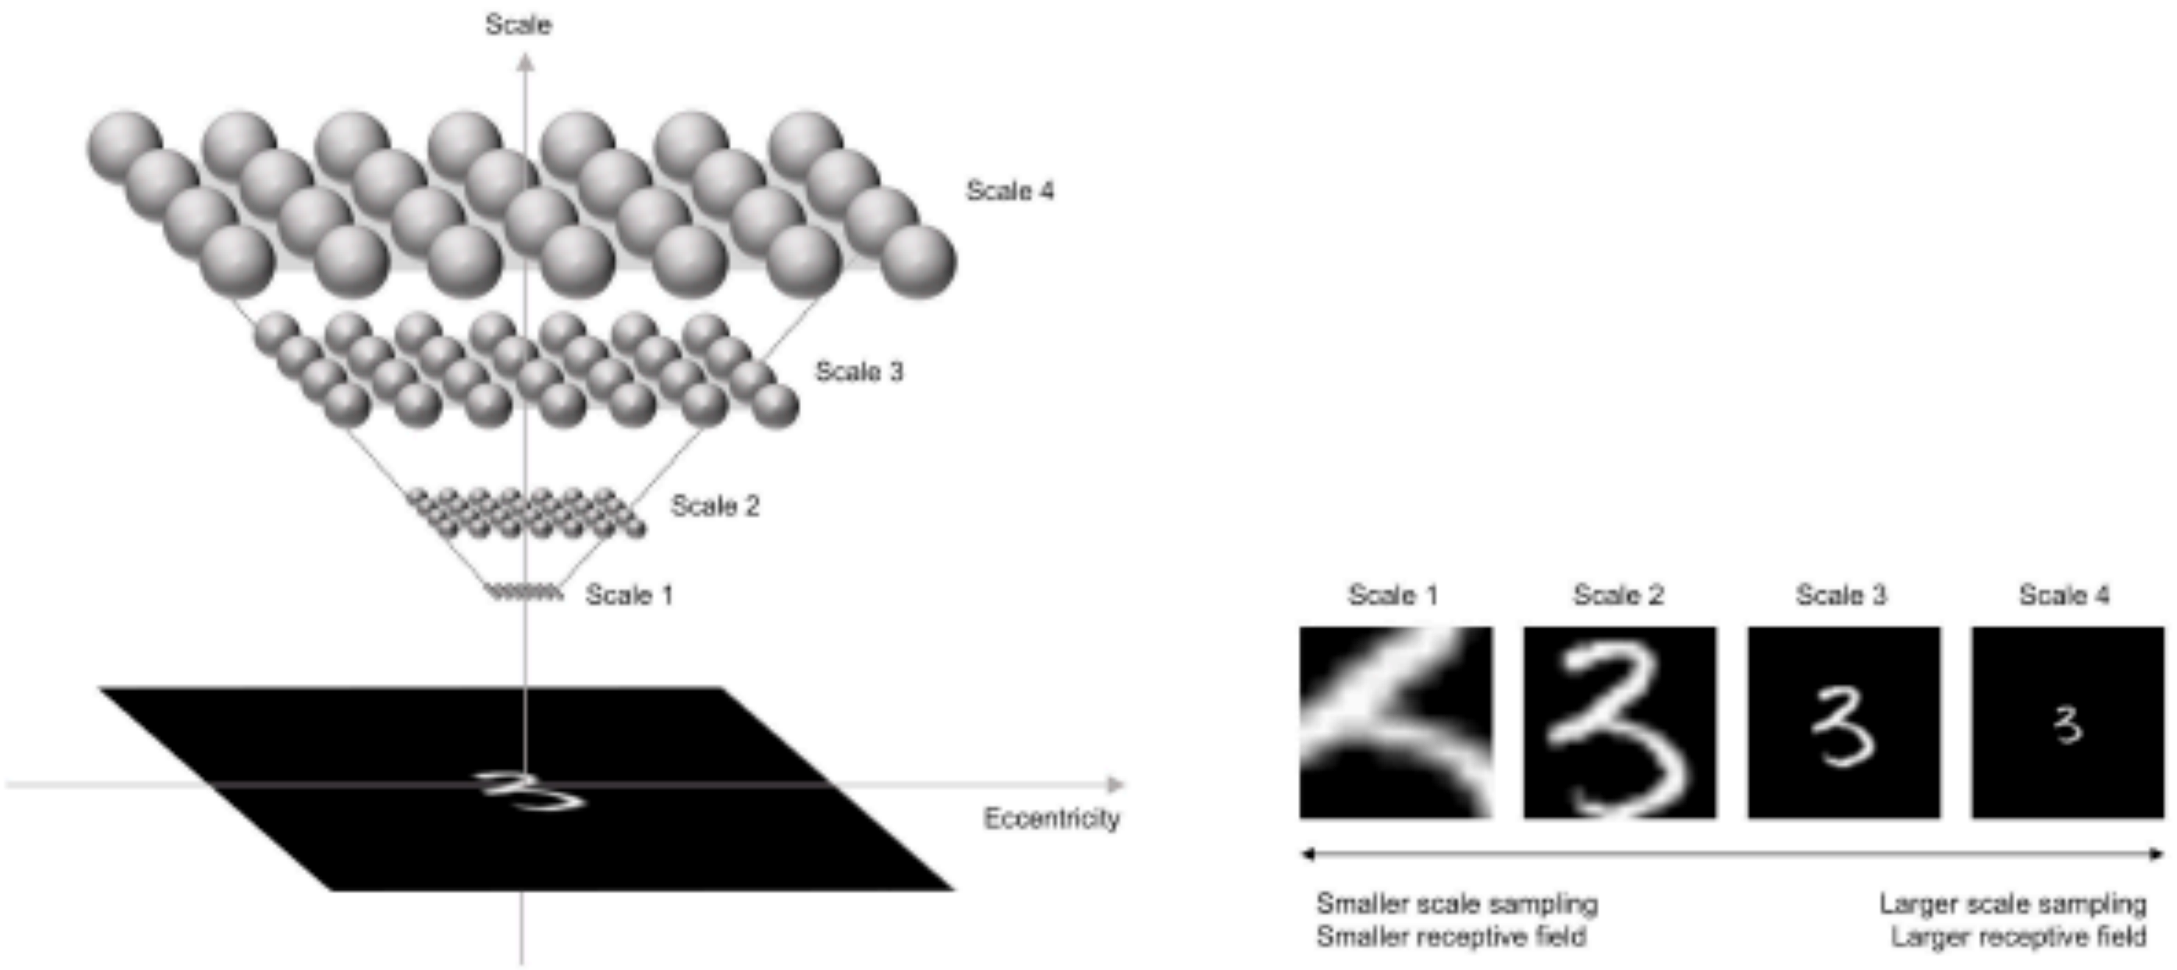
Scale
Scale 4
Scale 3
Scale 2
Scale 1
Eccentricity
Scale 1
Scale 2
Scale 3
Scale 4
Smaller scale sampling
Smaller receptive field
Larger scale sampling
Larger receptive field

(A) (B)

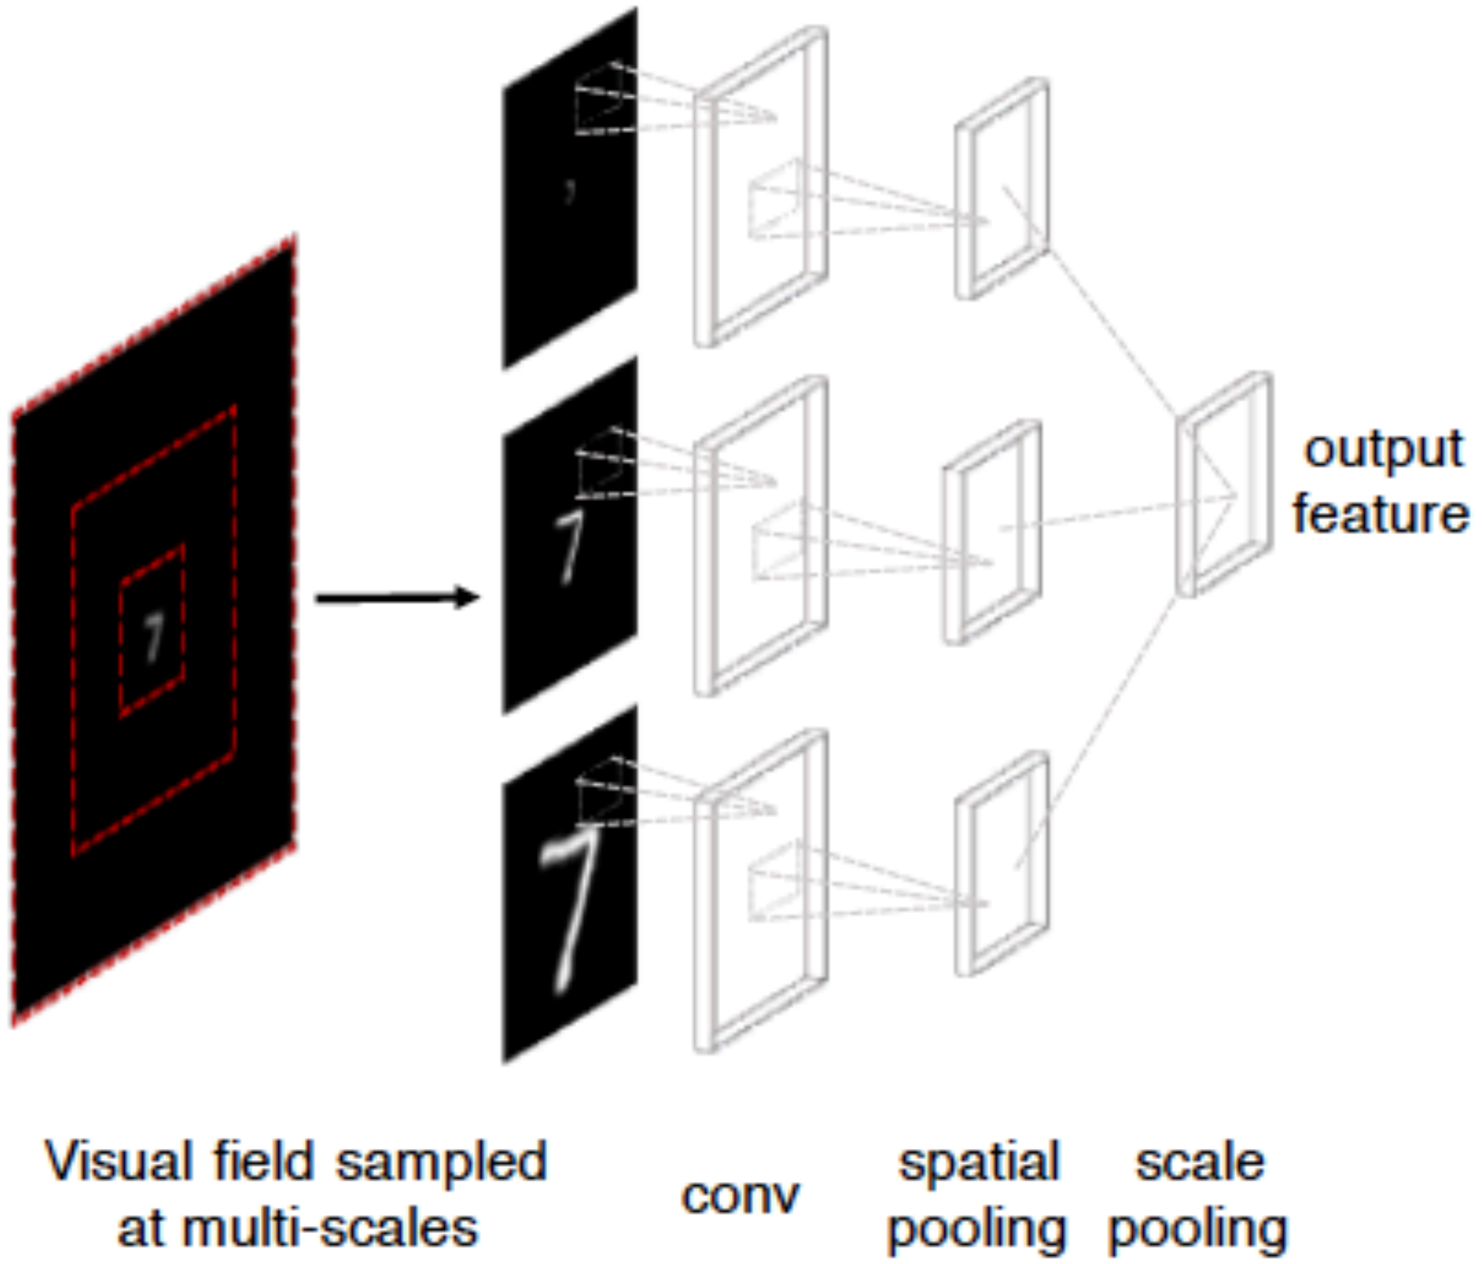
output
feature
Visual field sampled
at multi-scales
conv
spatial
pooling
scale
pooling

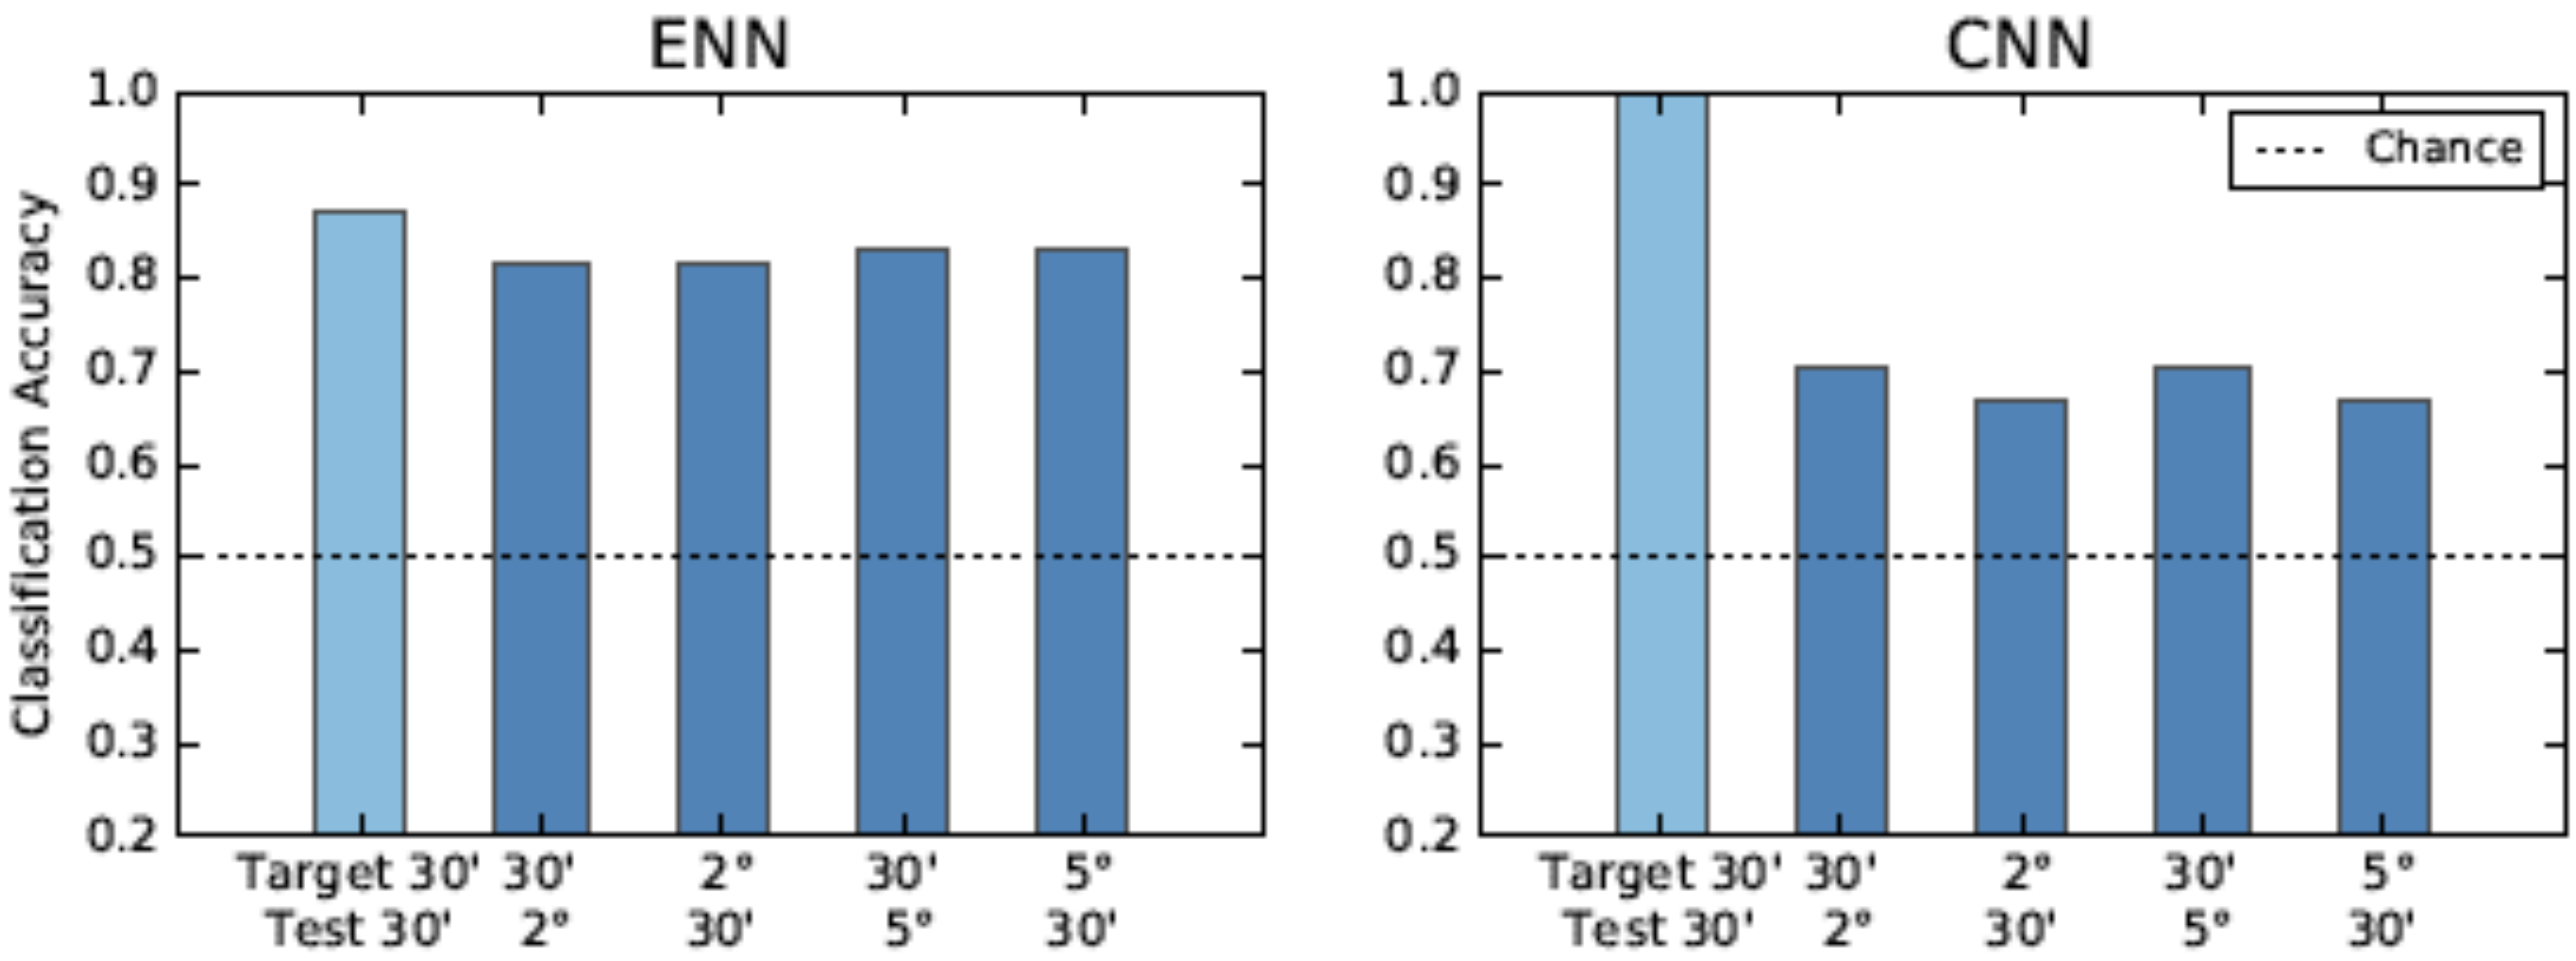
ENN
CNN
Chance
Classification Accuracy
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
Target 30'
Test 30'
30'
2°
2°
30'
30'
5°
5°
30'

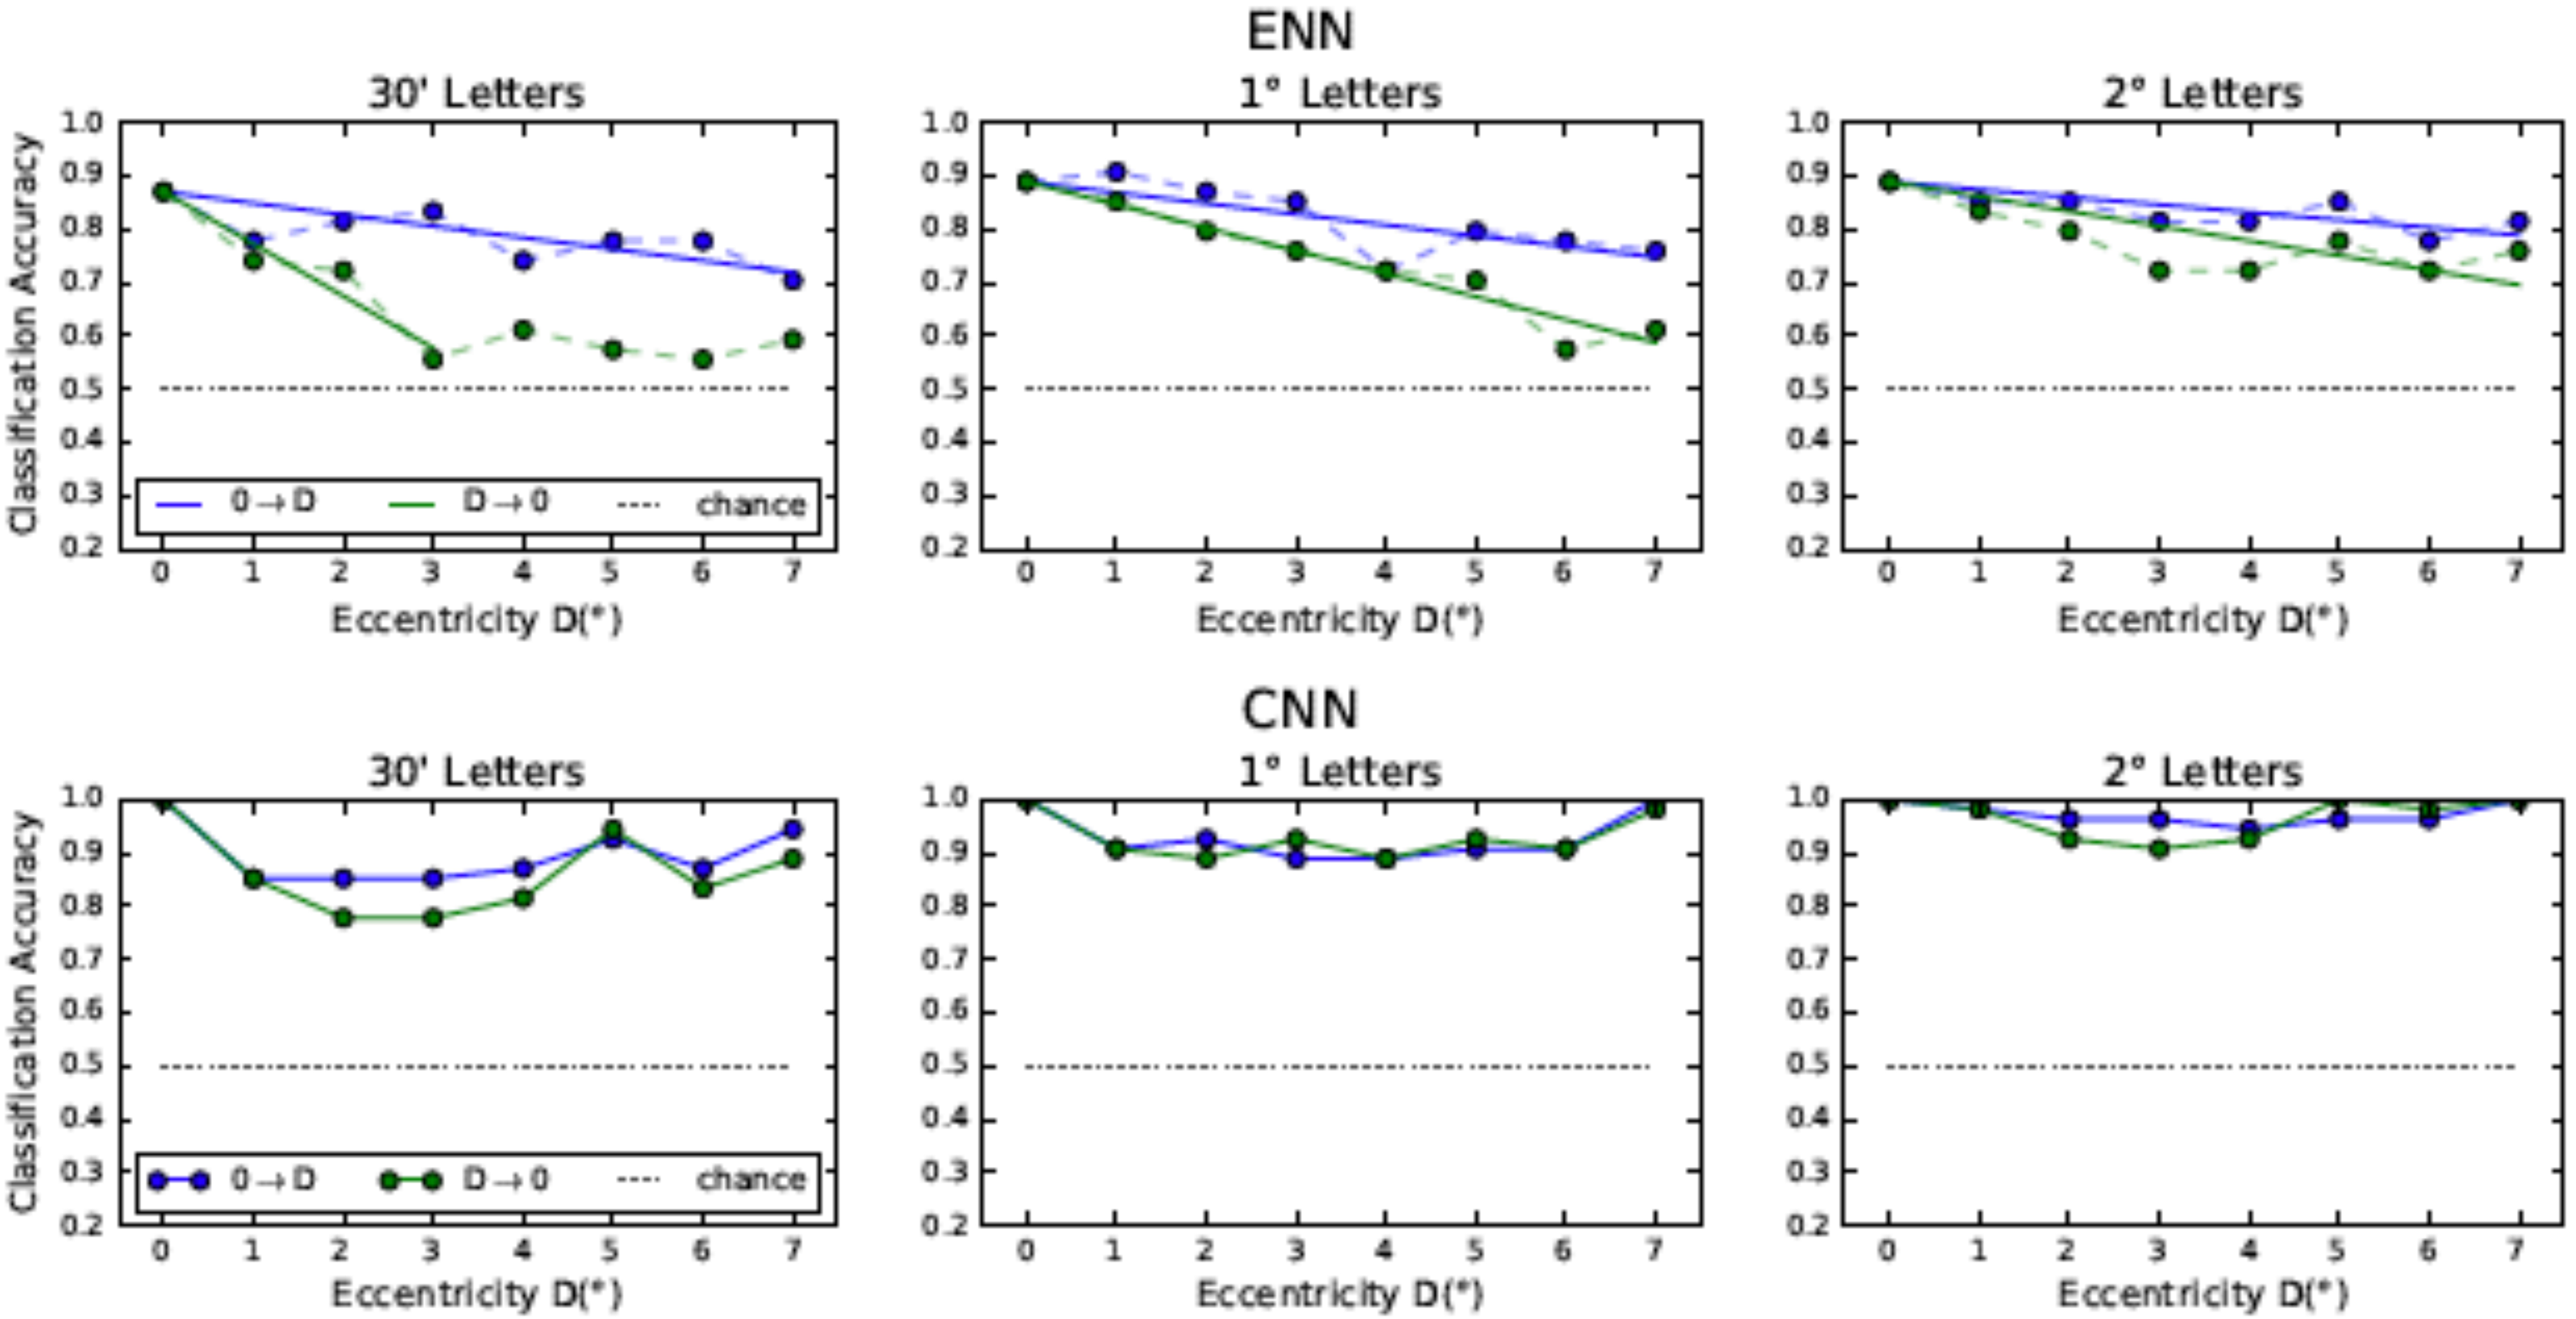
ENN
30' Letters
1° Letters
2° Letters
Classification Accuracy
Eccentricity D(°)
0 → D
D → 0
chance
CNN
30' Letters
1° Letters
2° Letters
Classification Accuracy
Eccentricity D(°)
0 → D
D → 0
chance


Module one determines computational strategy of vision (eye movements)
different from today’s CNNs

9.523 overview

- Course description/logistic
- **Motivations for this course: the greatest problem in science, CBMM, the MIT Quest**
- Neuroscience and AI
- CBMM, the Visual Intelligence moonshot
- Module 1
  - Module 1, theory
  - Module 1, eccentricity
  - Module 1, development/invariance

Prediction: class specific modules (like faces) in visual cortex because of *class-specific* invariances

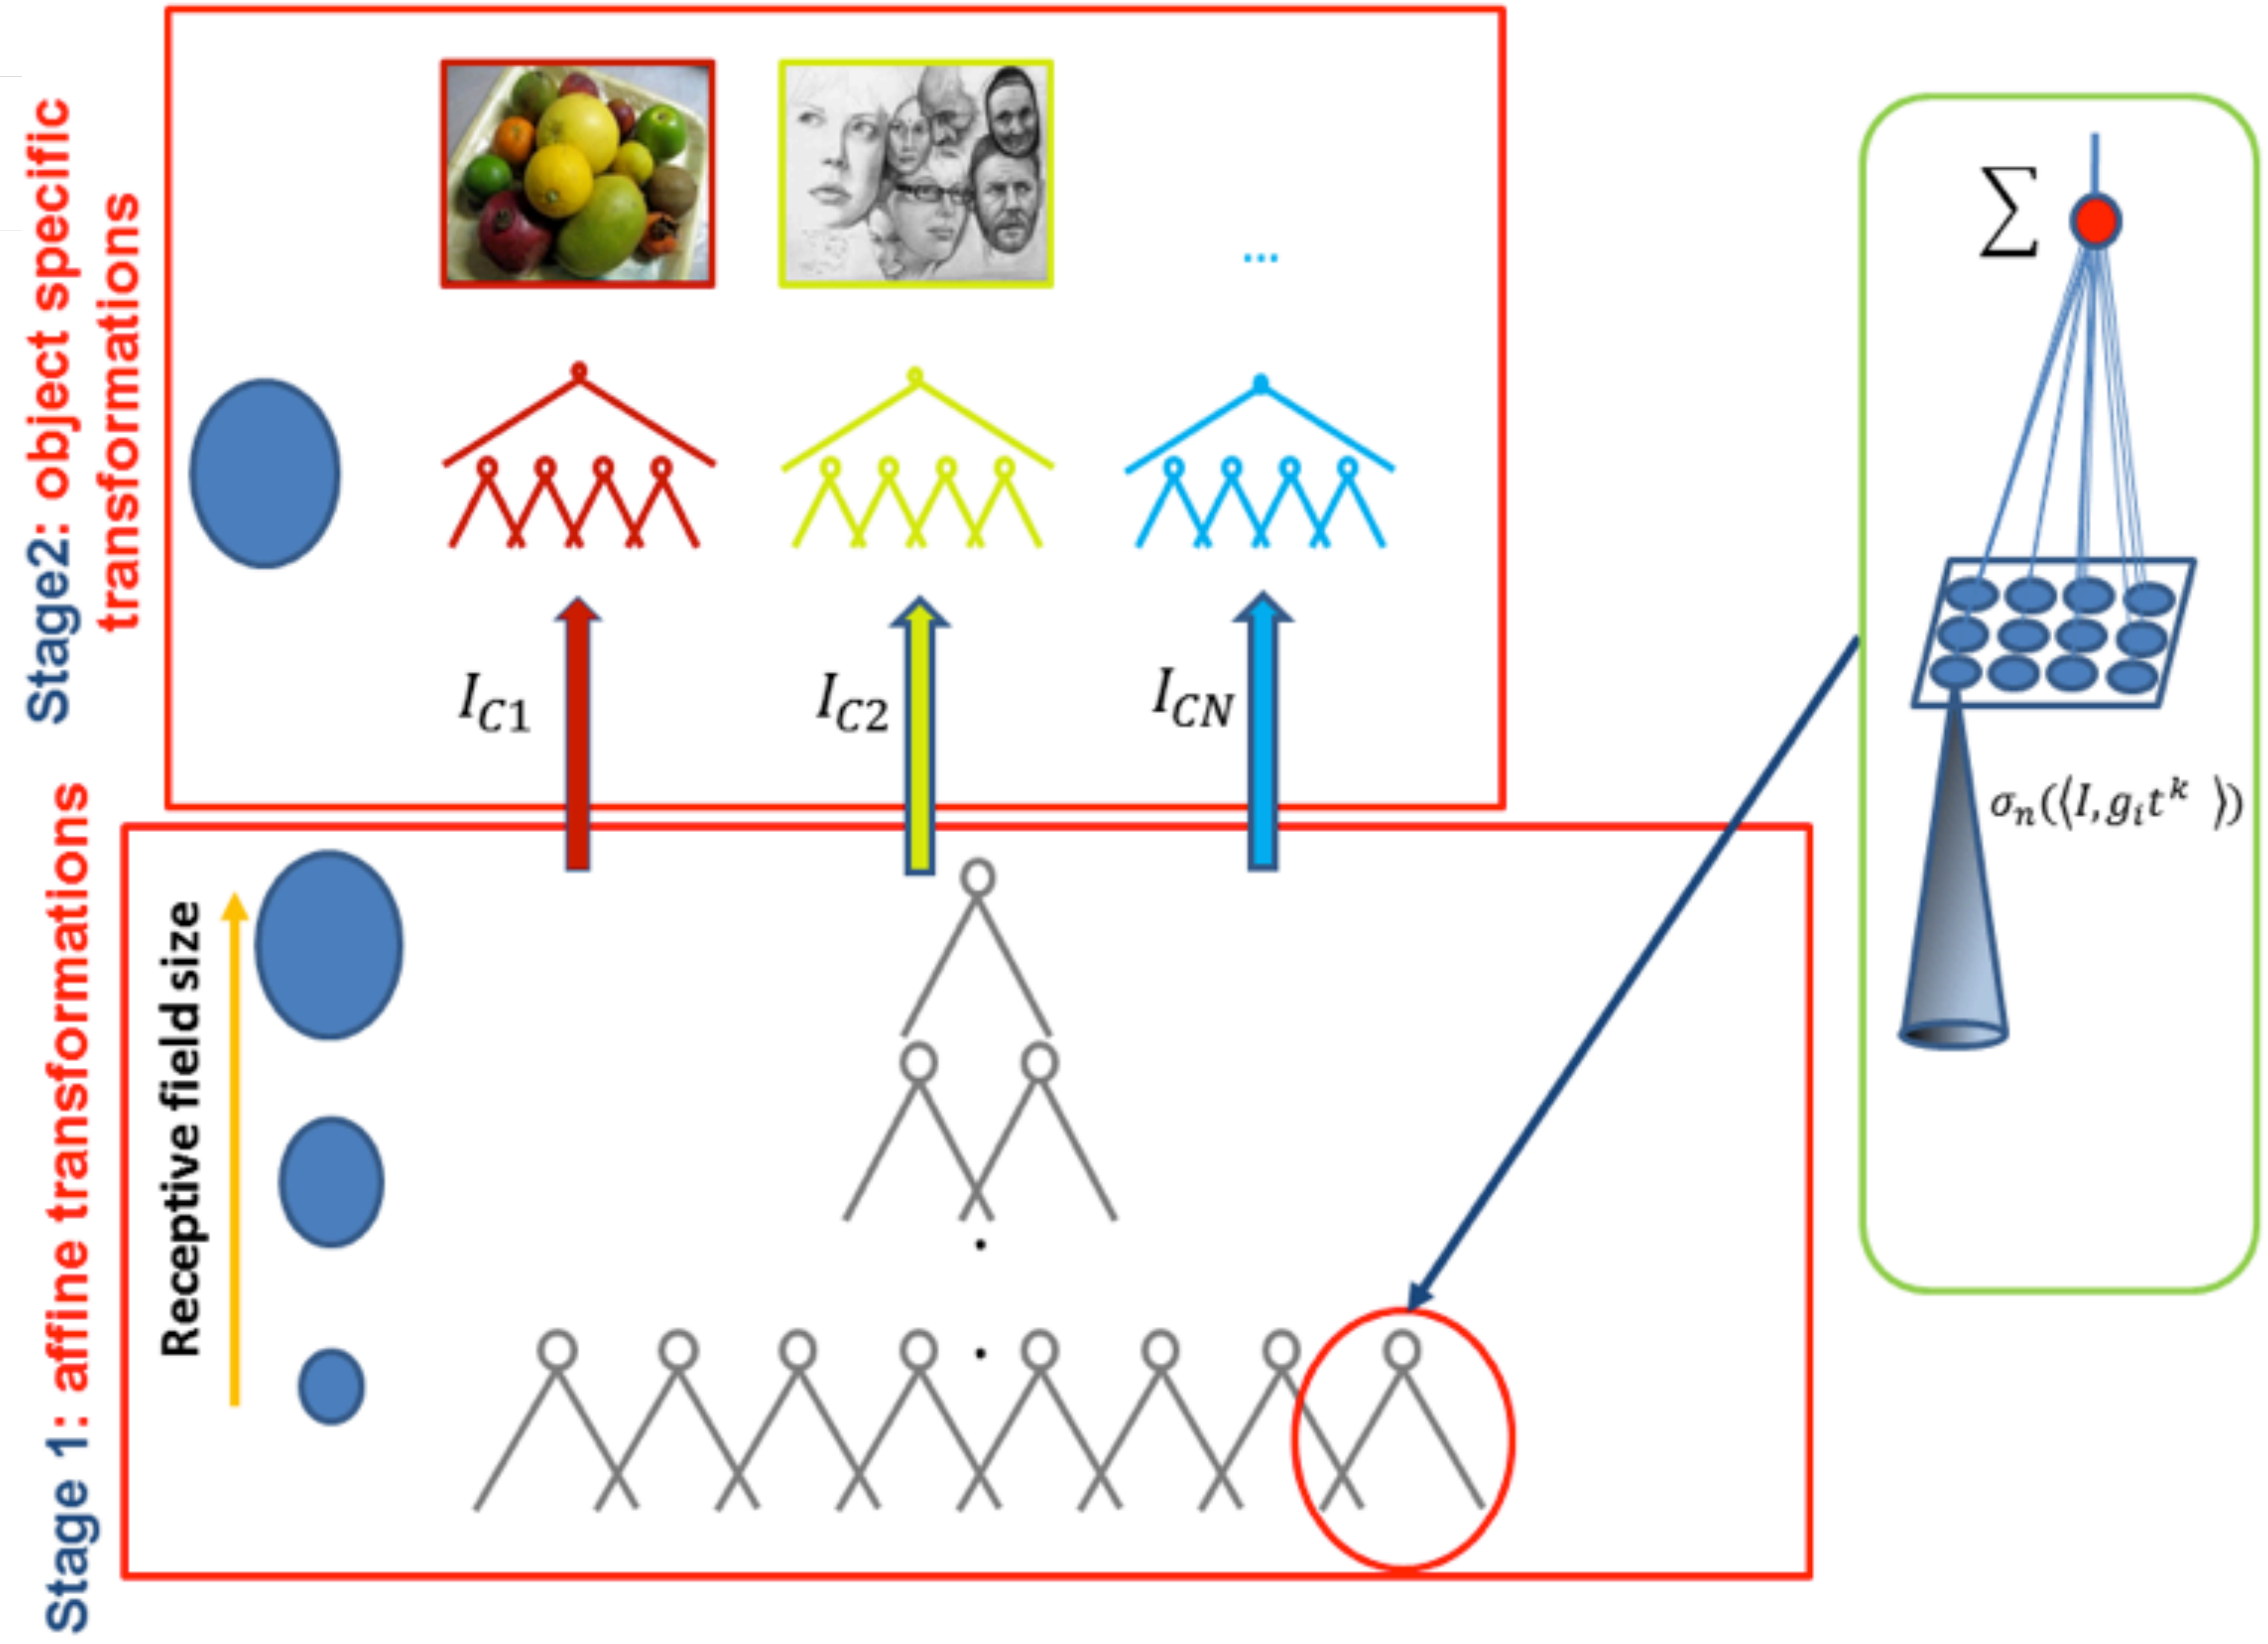

Leibo, Liao, Poggio, 2015

The Who Question: Face Recognition

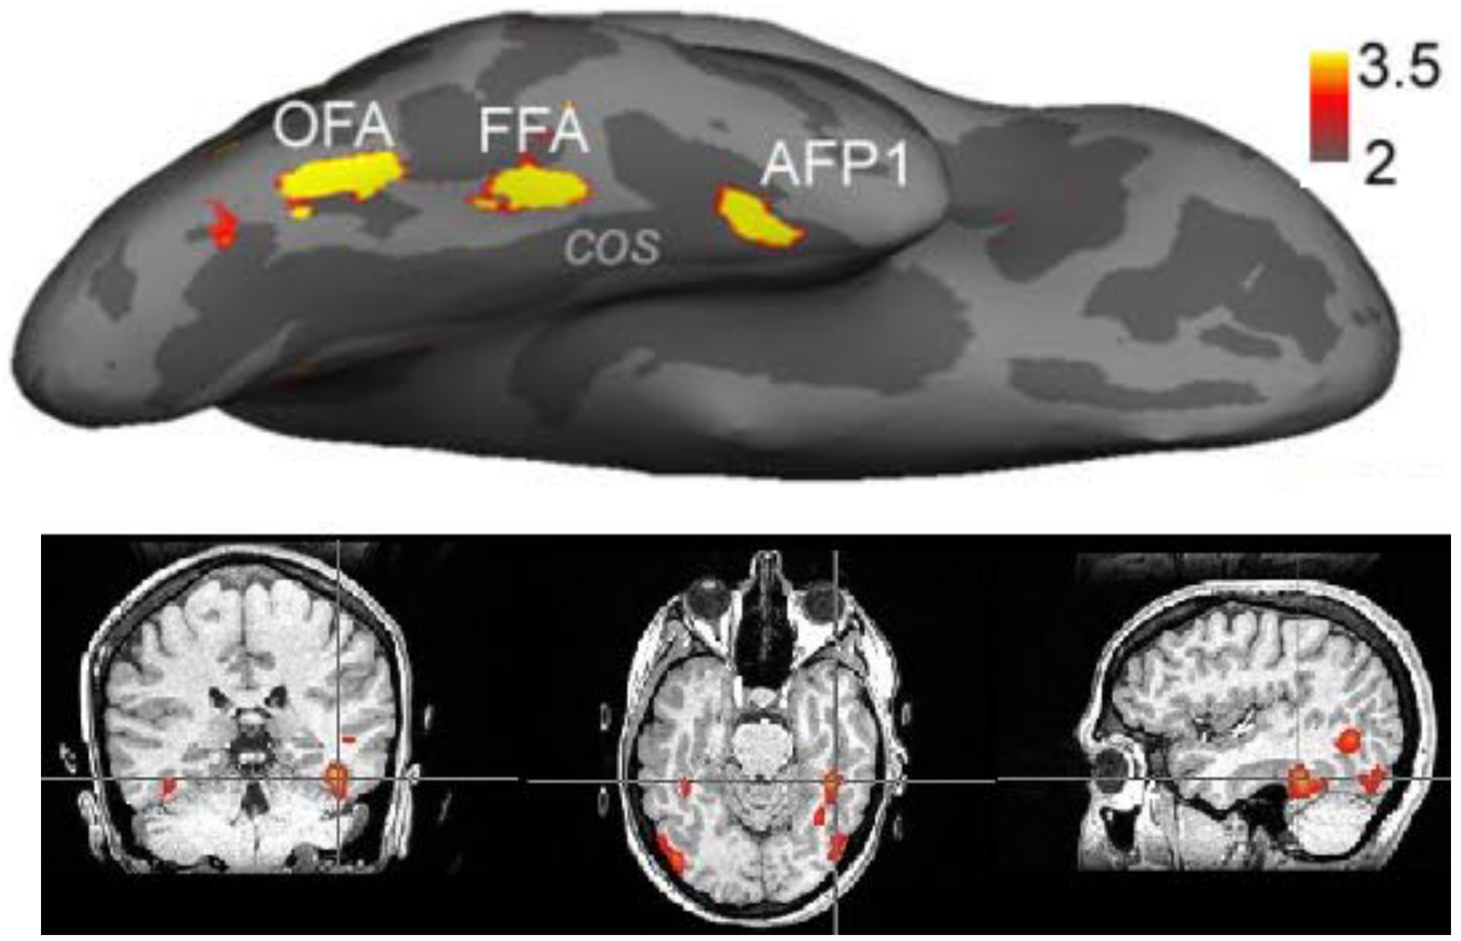

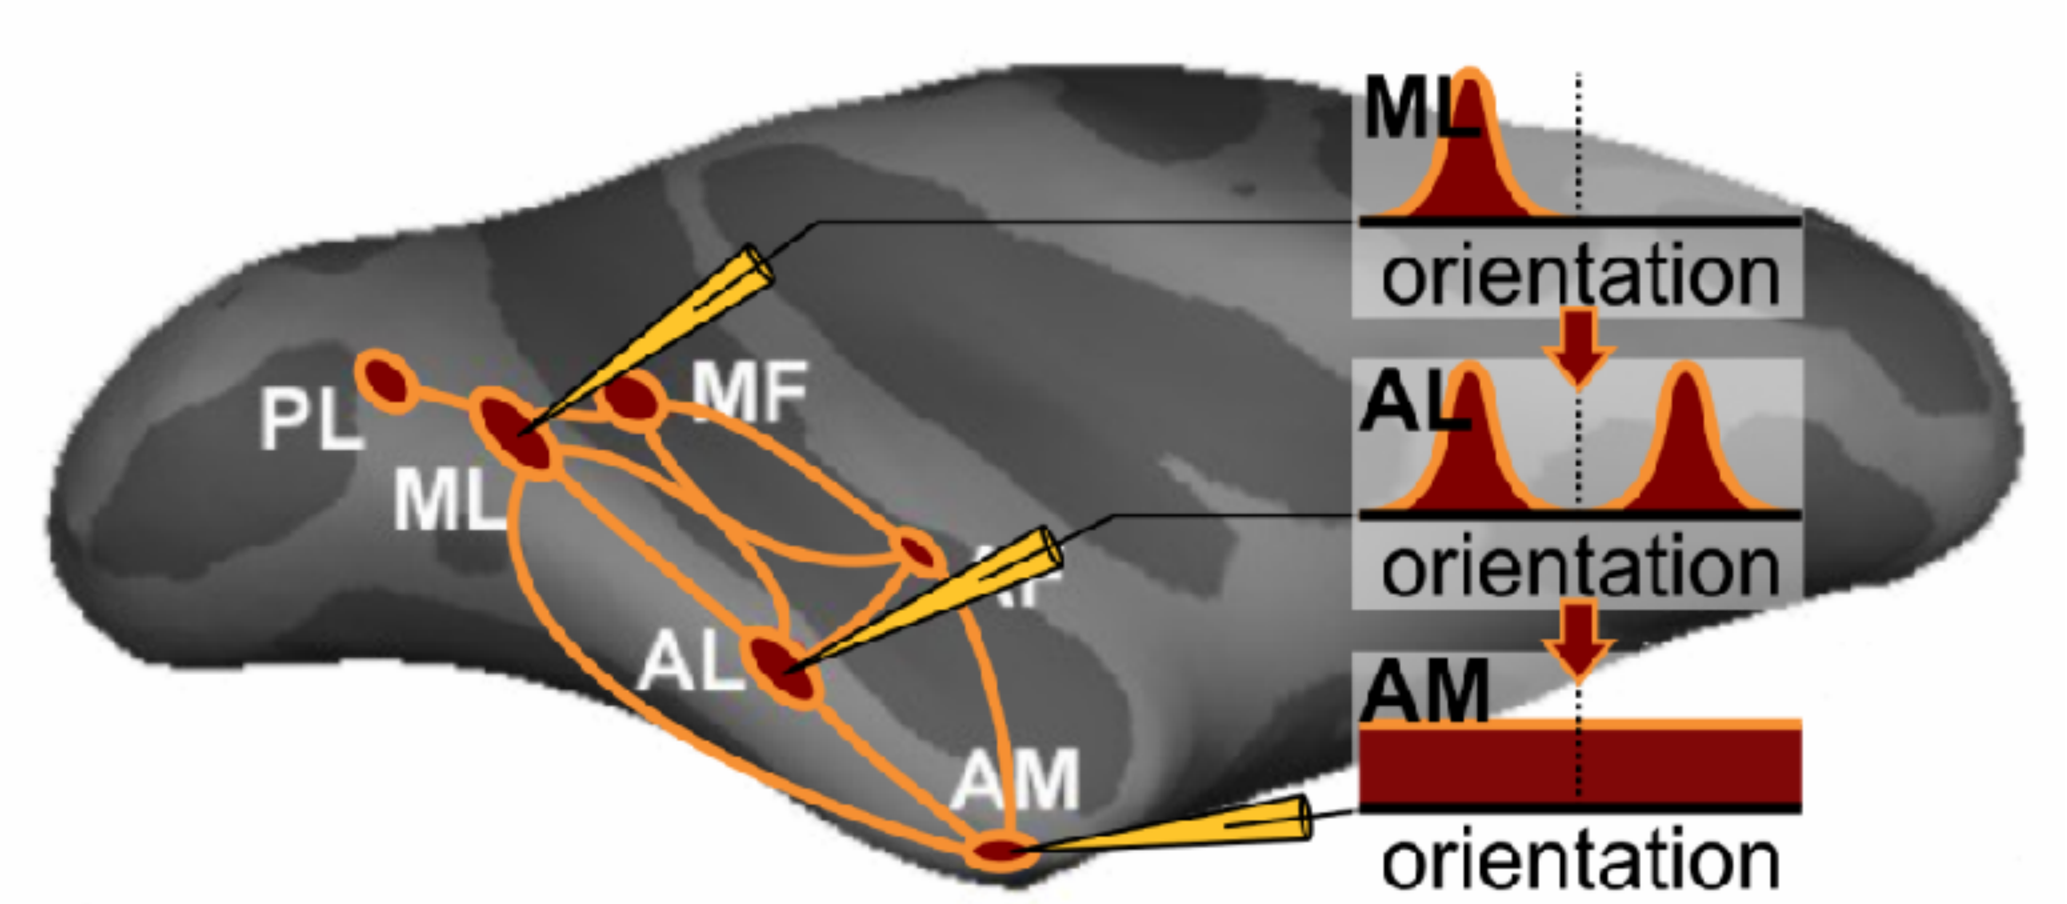

Winrich Freiwald and Doris Tsao

The Who Question: Face Recognition

Adults

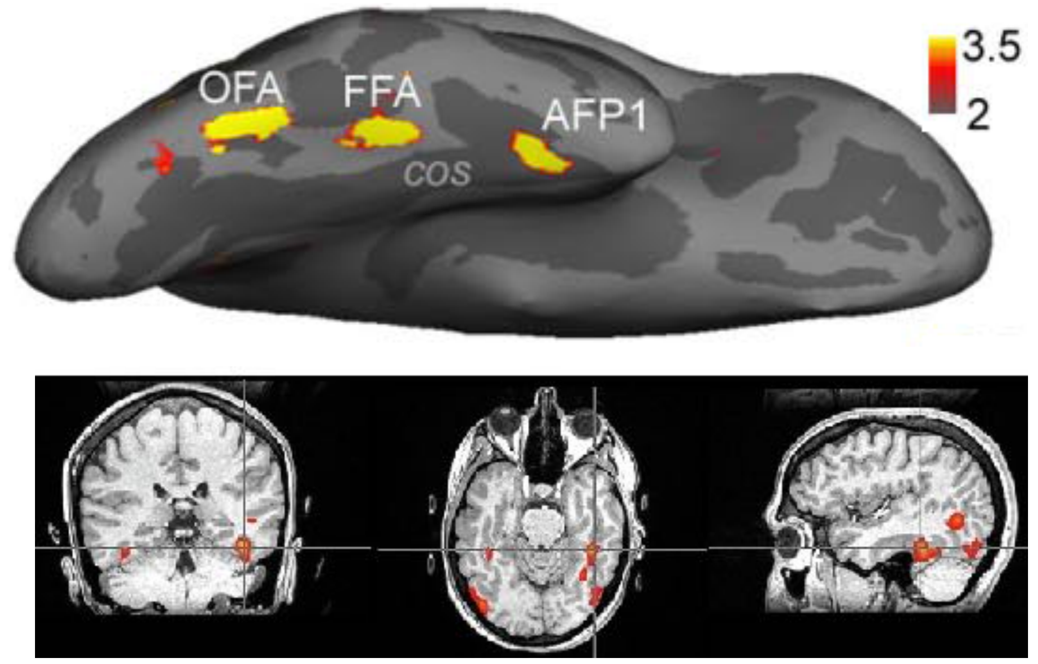

Kanwisher

Humans

Development

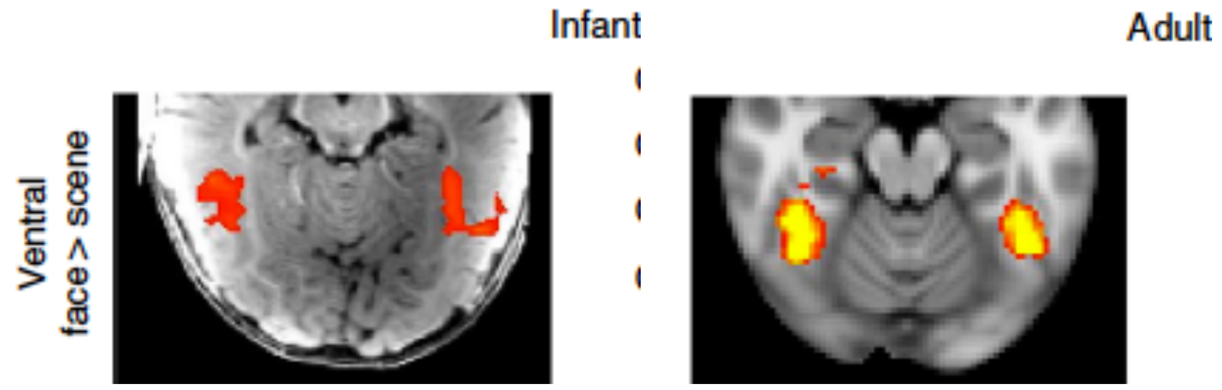

Deen, Richardson, Dilks, Takahashi, Keil, Wald,Kanwisher, Saxe, NN, 2017

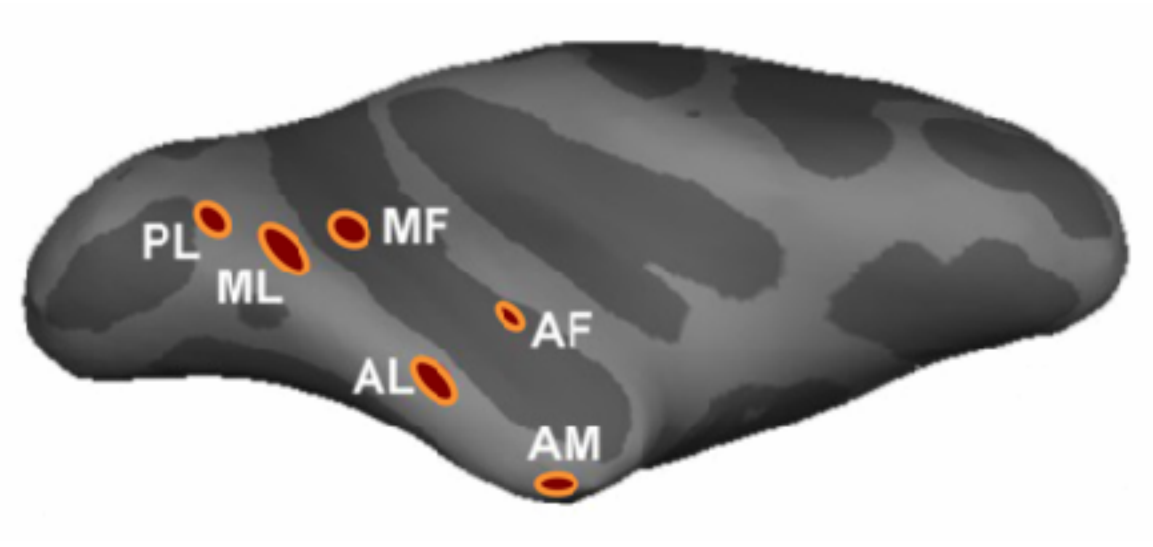

Winrich Freiwald and Doris Tsao

Macaque

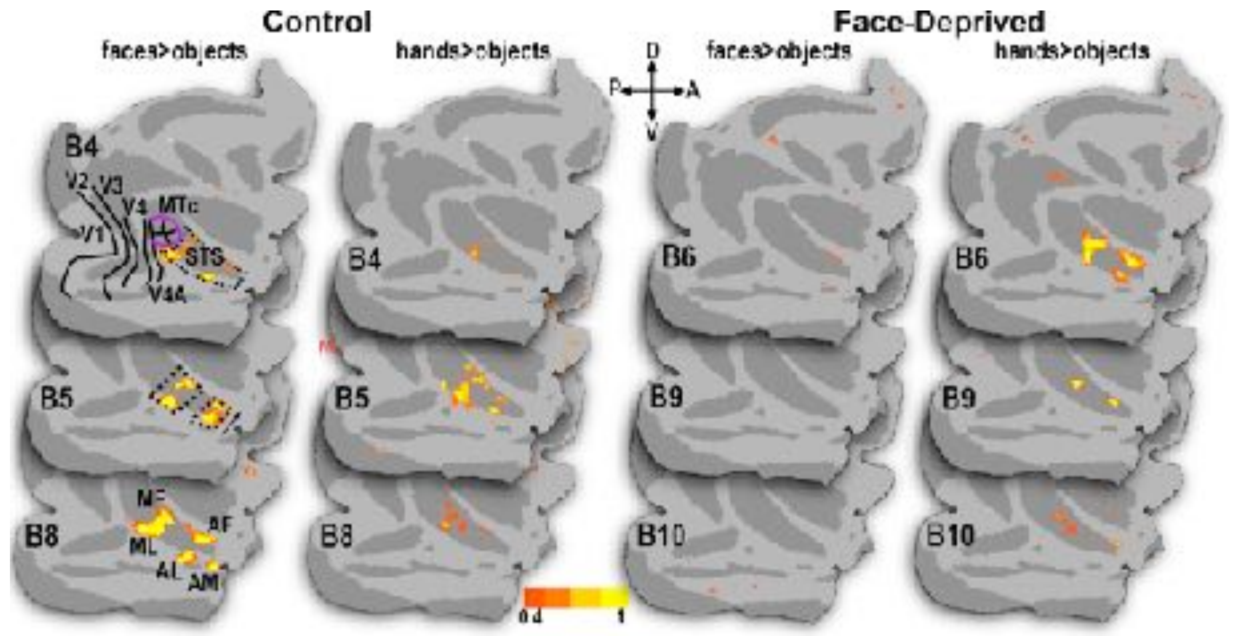

Marge Livingstone, 2017 unpublished

A HW-like module for tolerance to pose of a novel face

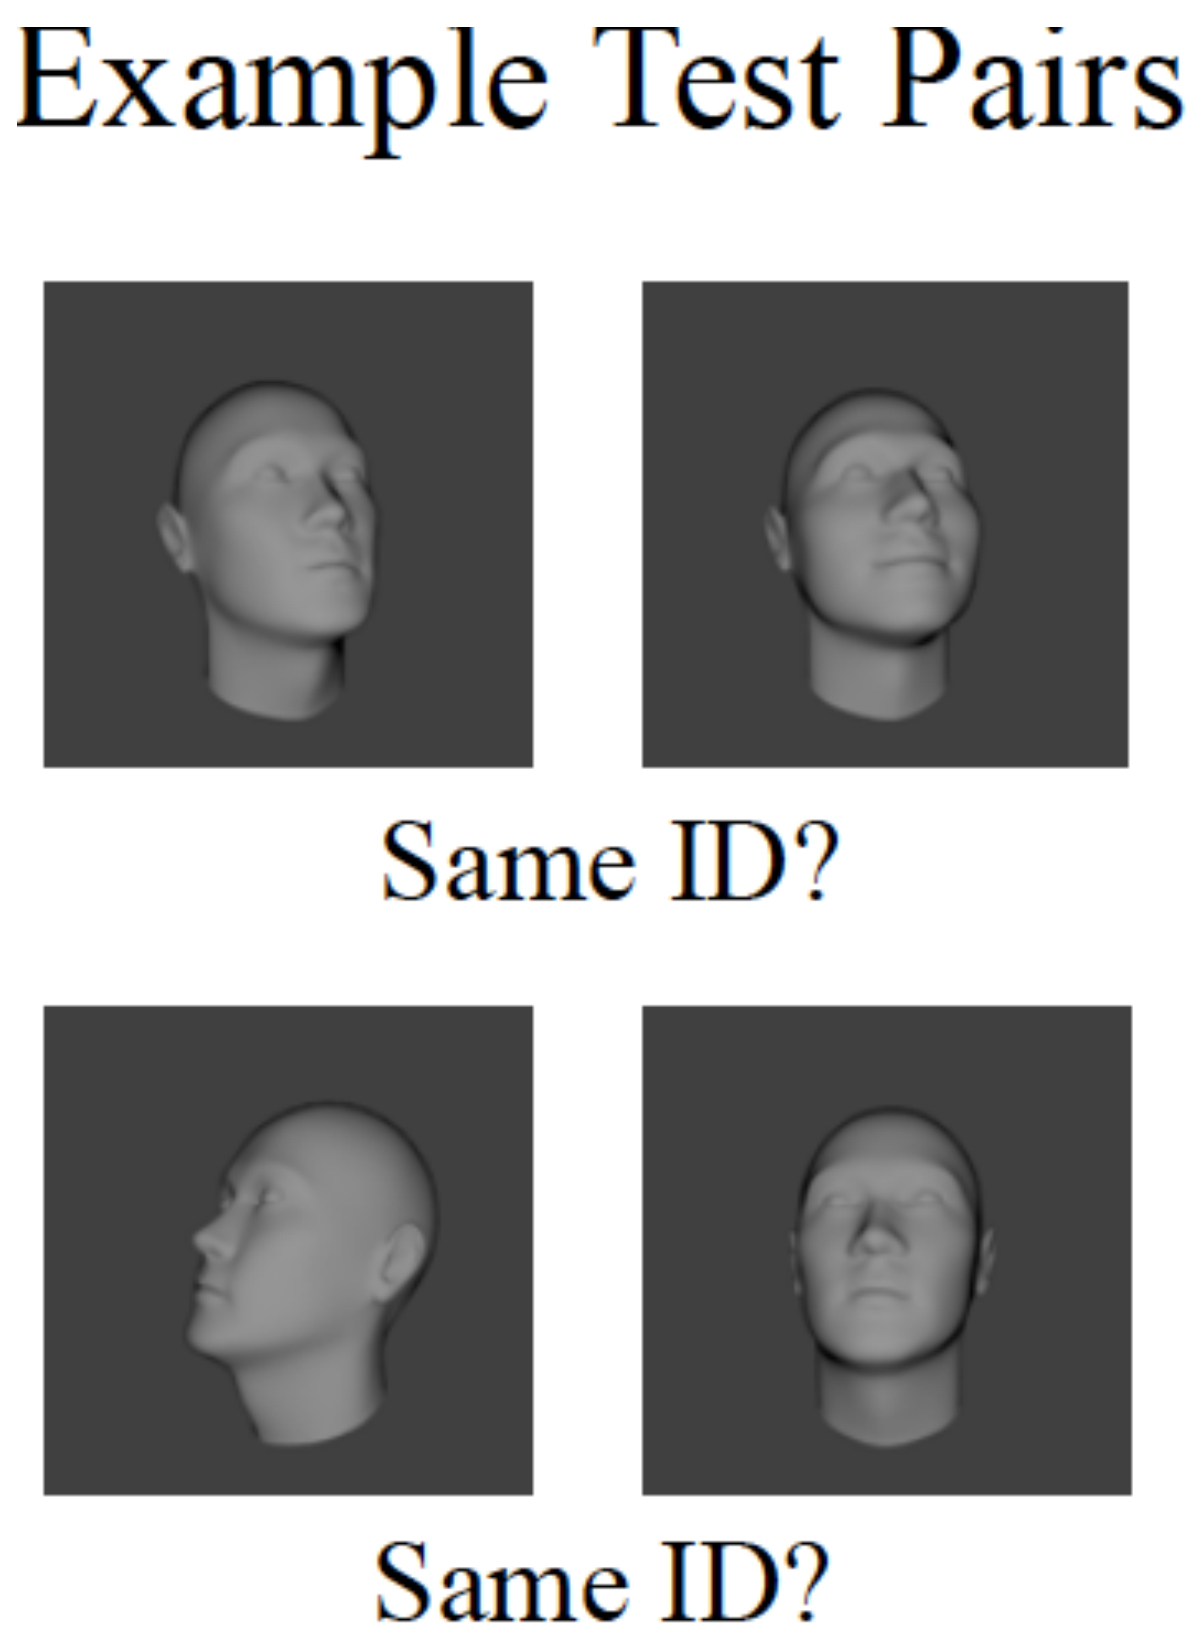

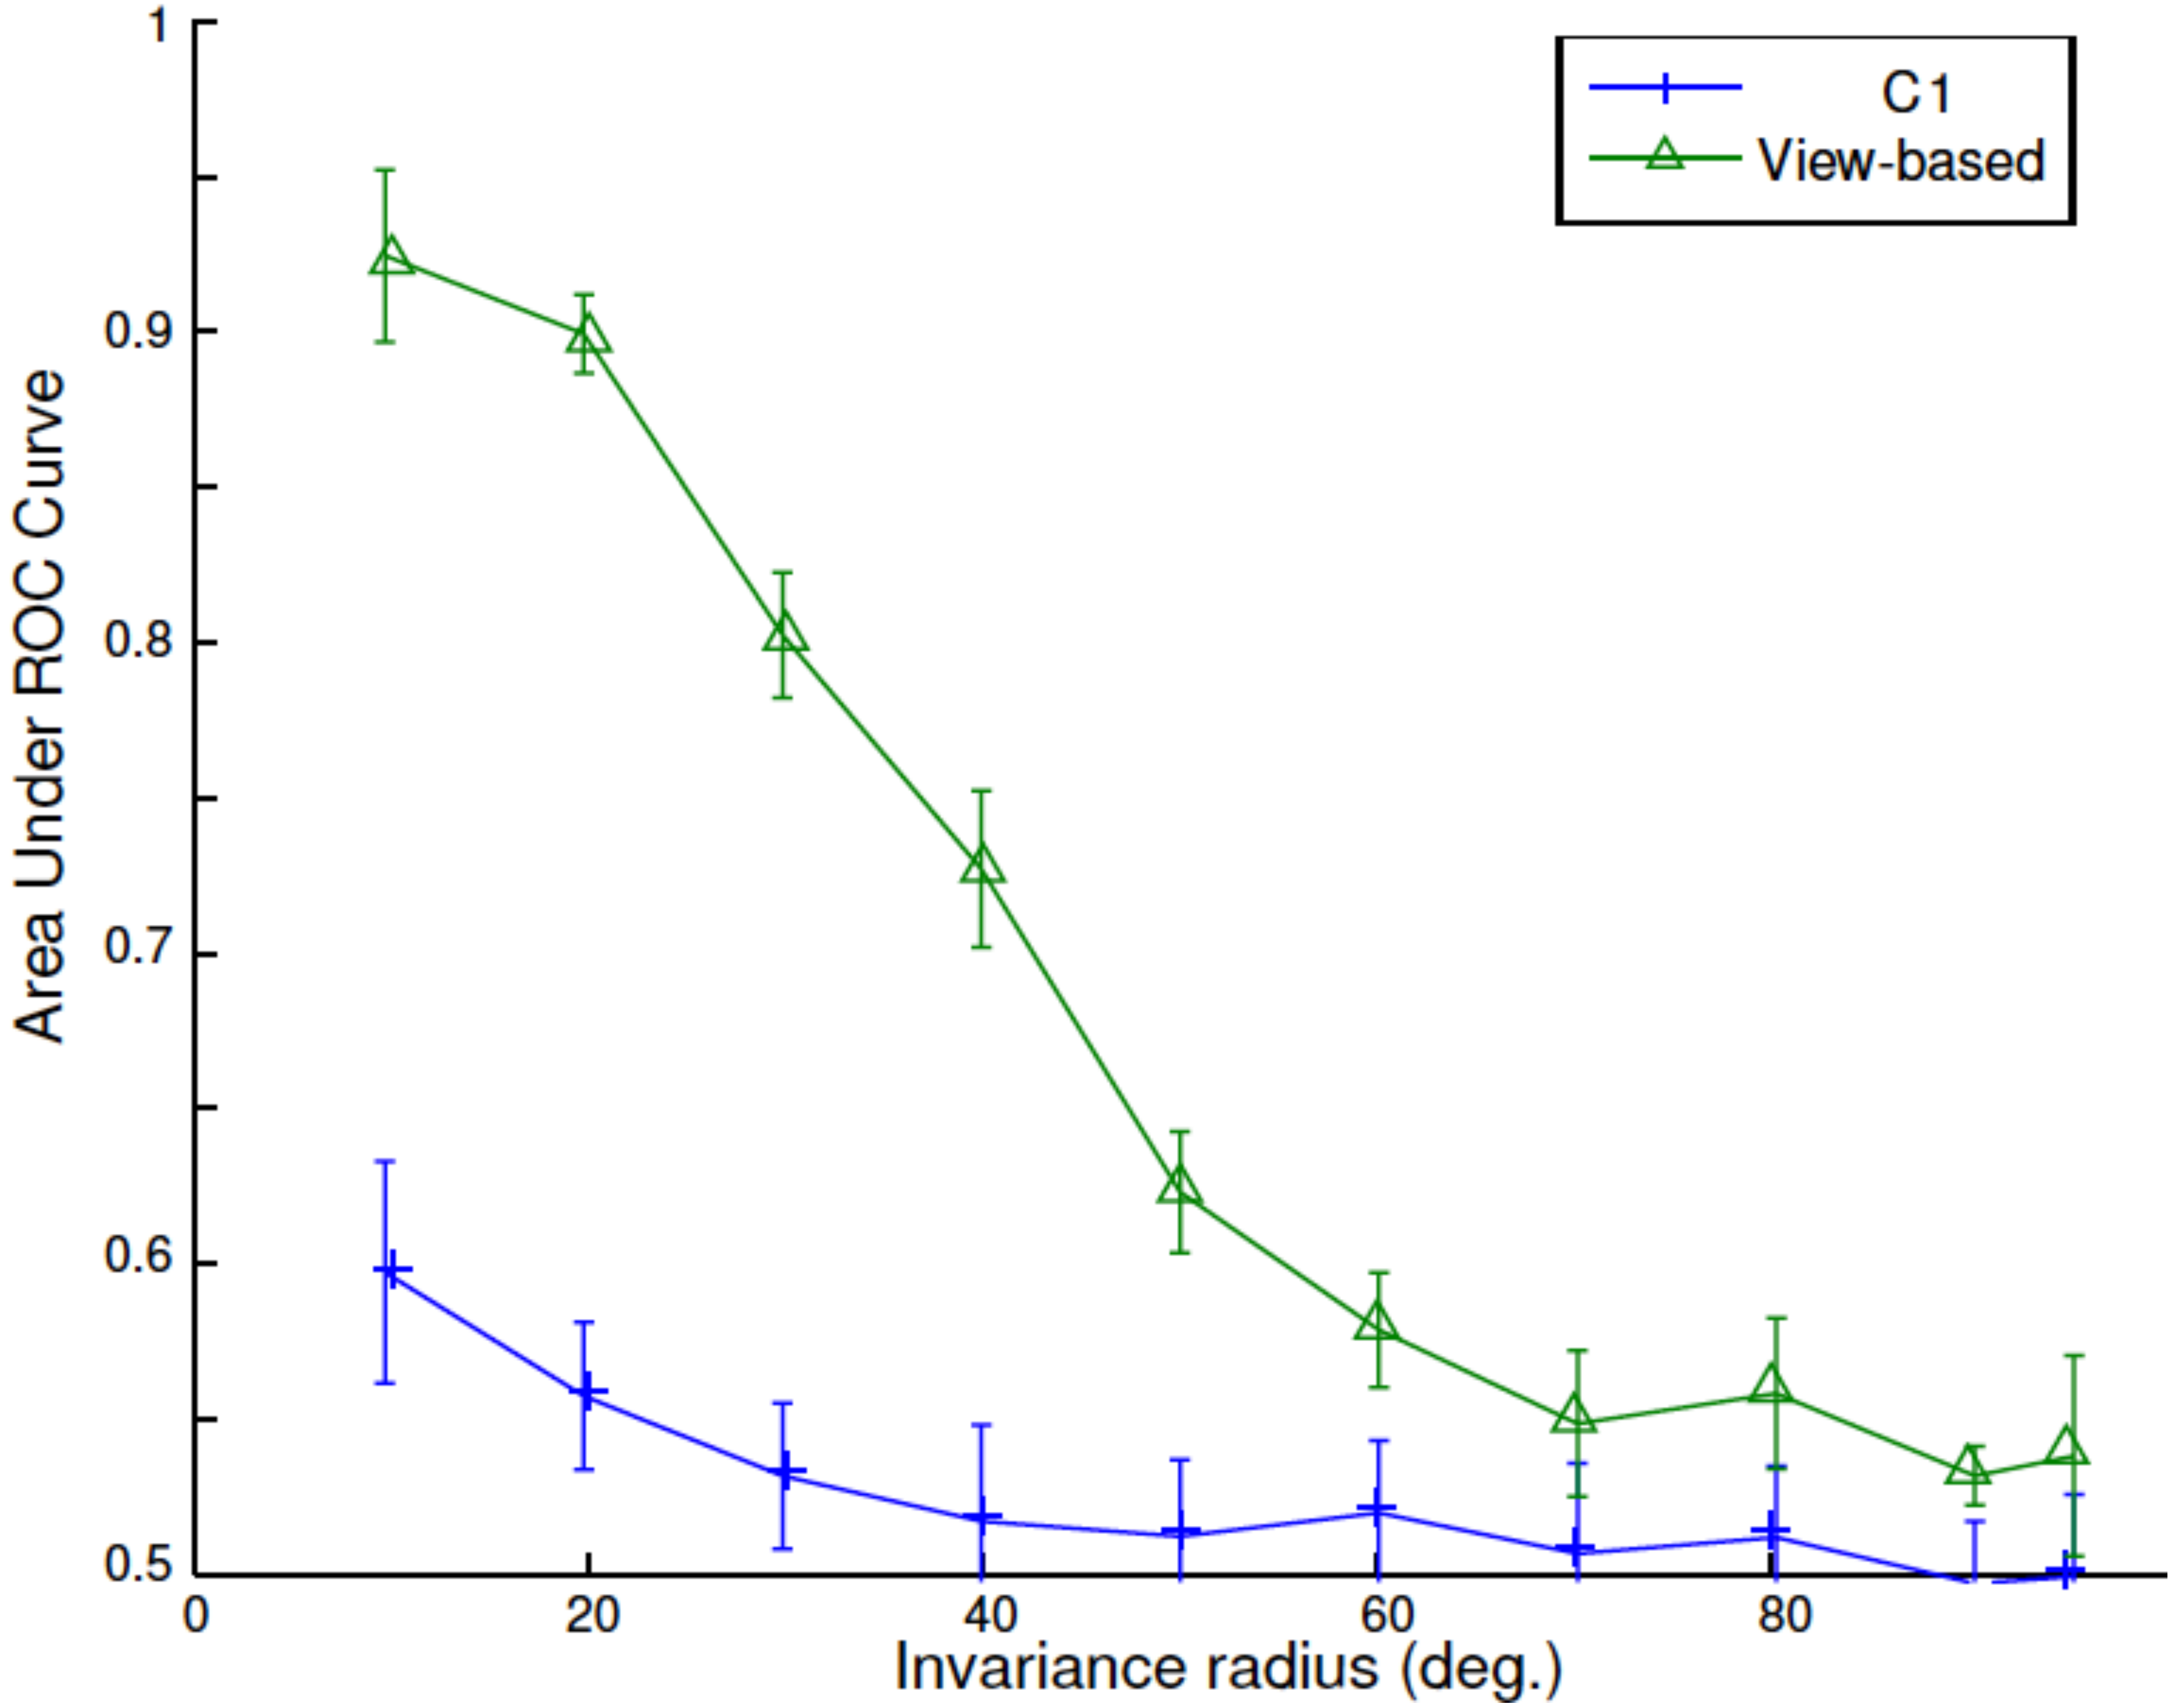

Leibo, Freiwald , Anselmi, Liao, Poggio, 2016

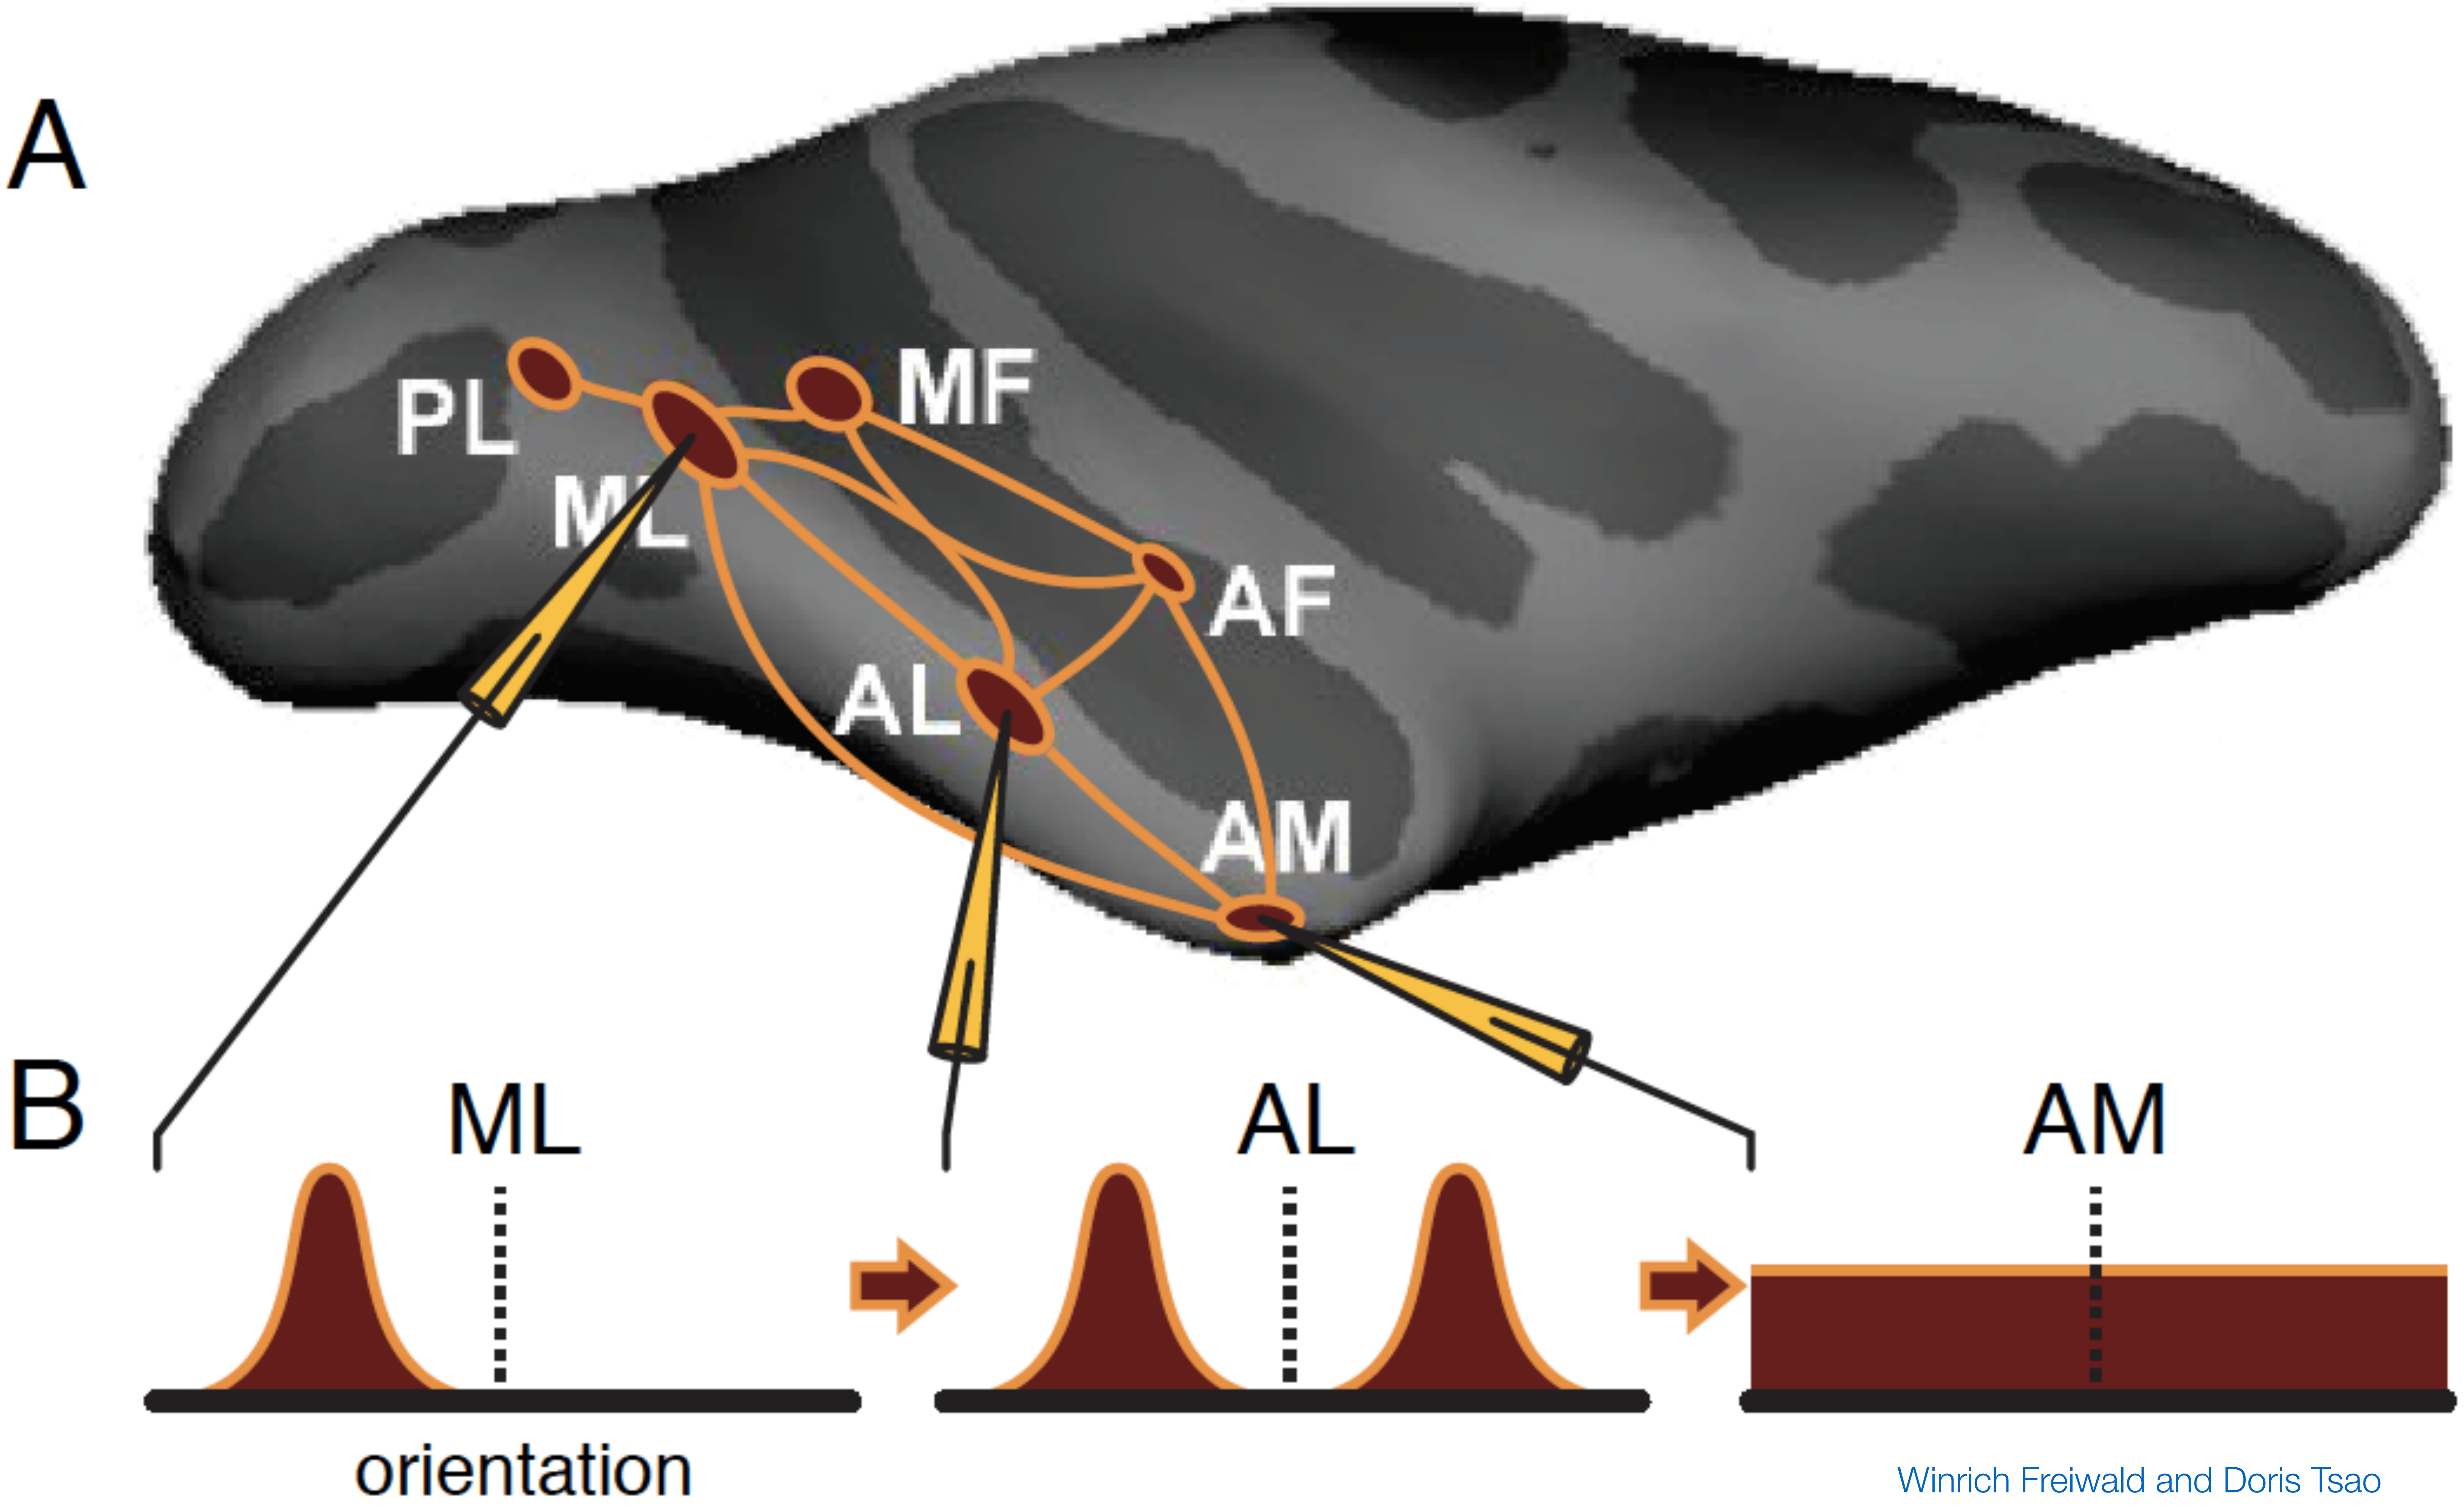

Winrich Freiwald and Doris Tsao

9.523 overview

- Course description/logistic
- **Motivations for this course: the greatest problem in science, CBMM, the MIT Quest**
- Neuroscience and AI
- CBMM, the Visual Intelligence moonshot
- Module 1
  - Module 1, theory
  - Module 1, eccentricity
  - Module 1, invariance

• Theory of Deep Local Hierarchical Networks

- When and why are deep networks better than shallow networks?

- What is the landscape of the empirical risk?

- How can deep learning generalize so well?

Computation in a neural net

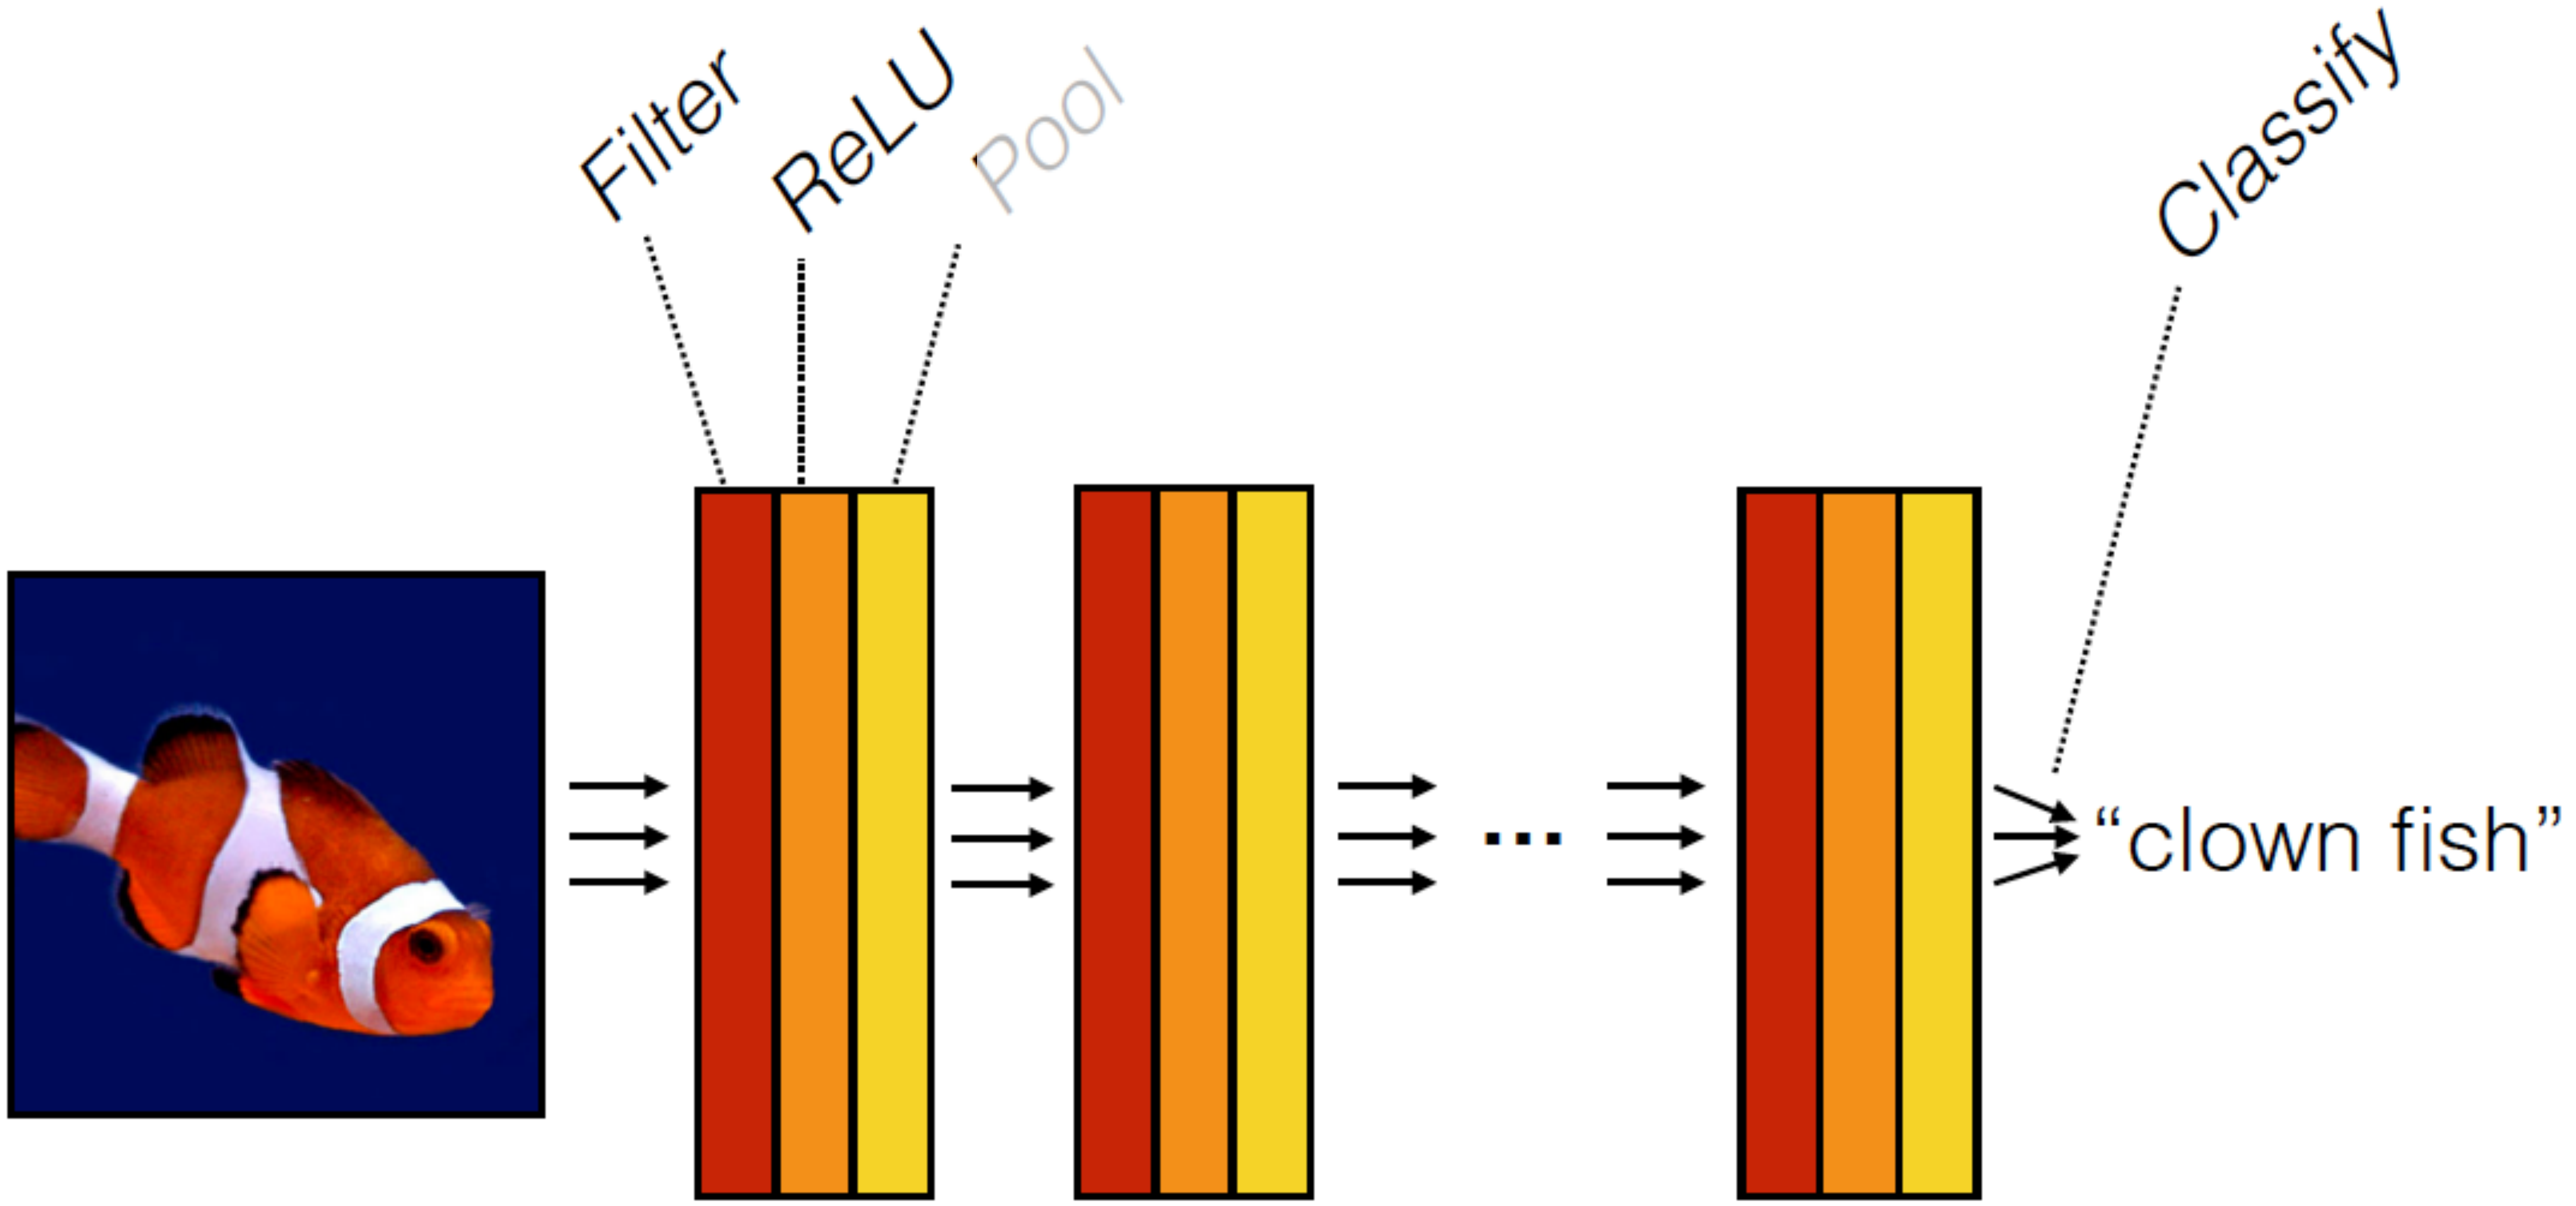

$$f(\mathbf{x}) = f_L(\dots f_2(f_1(\mathbf{x})))$$

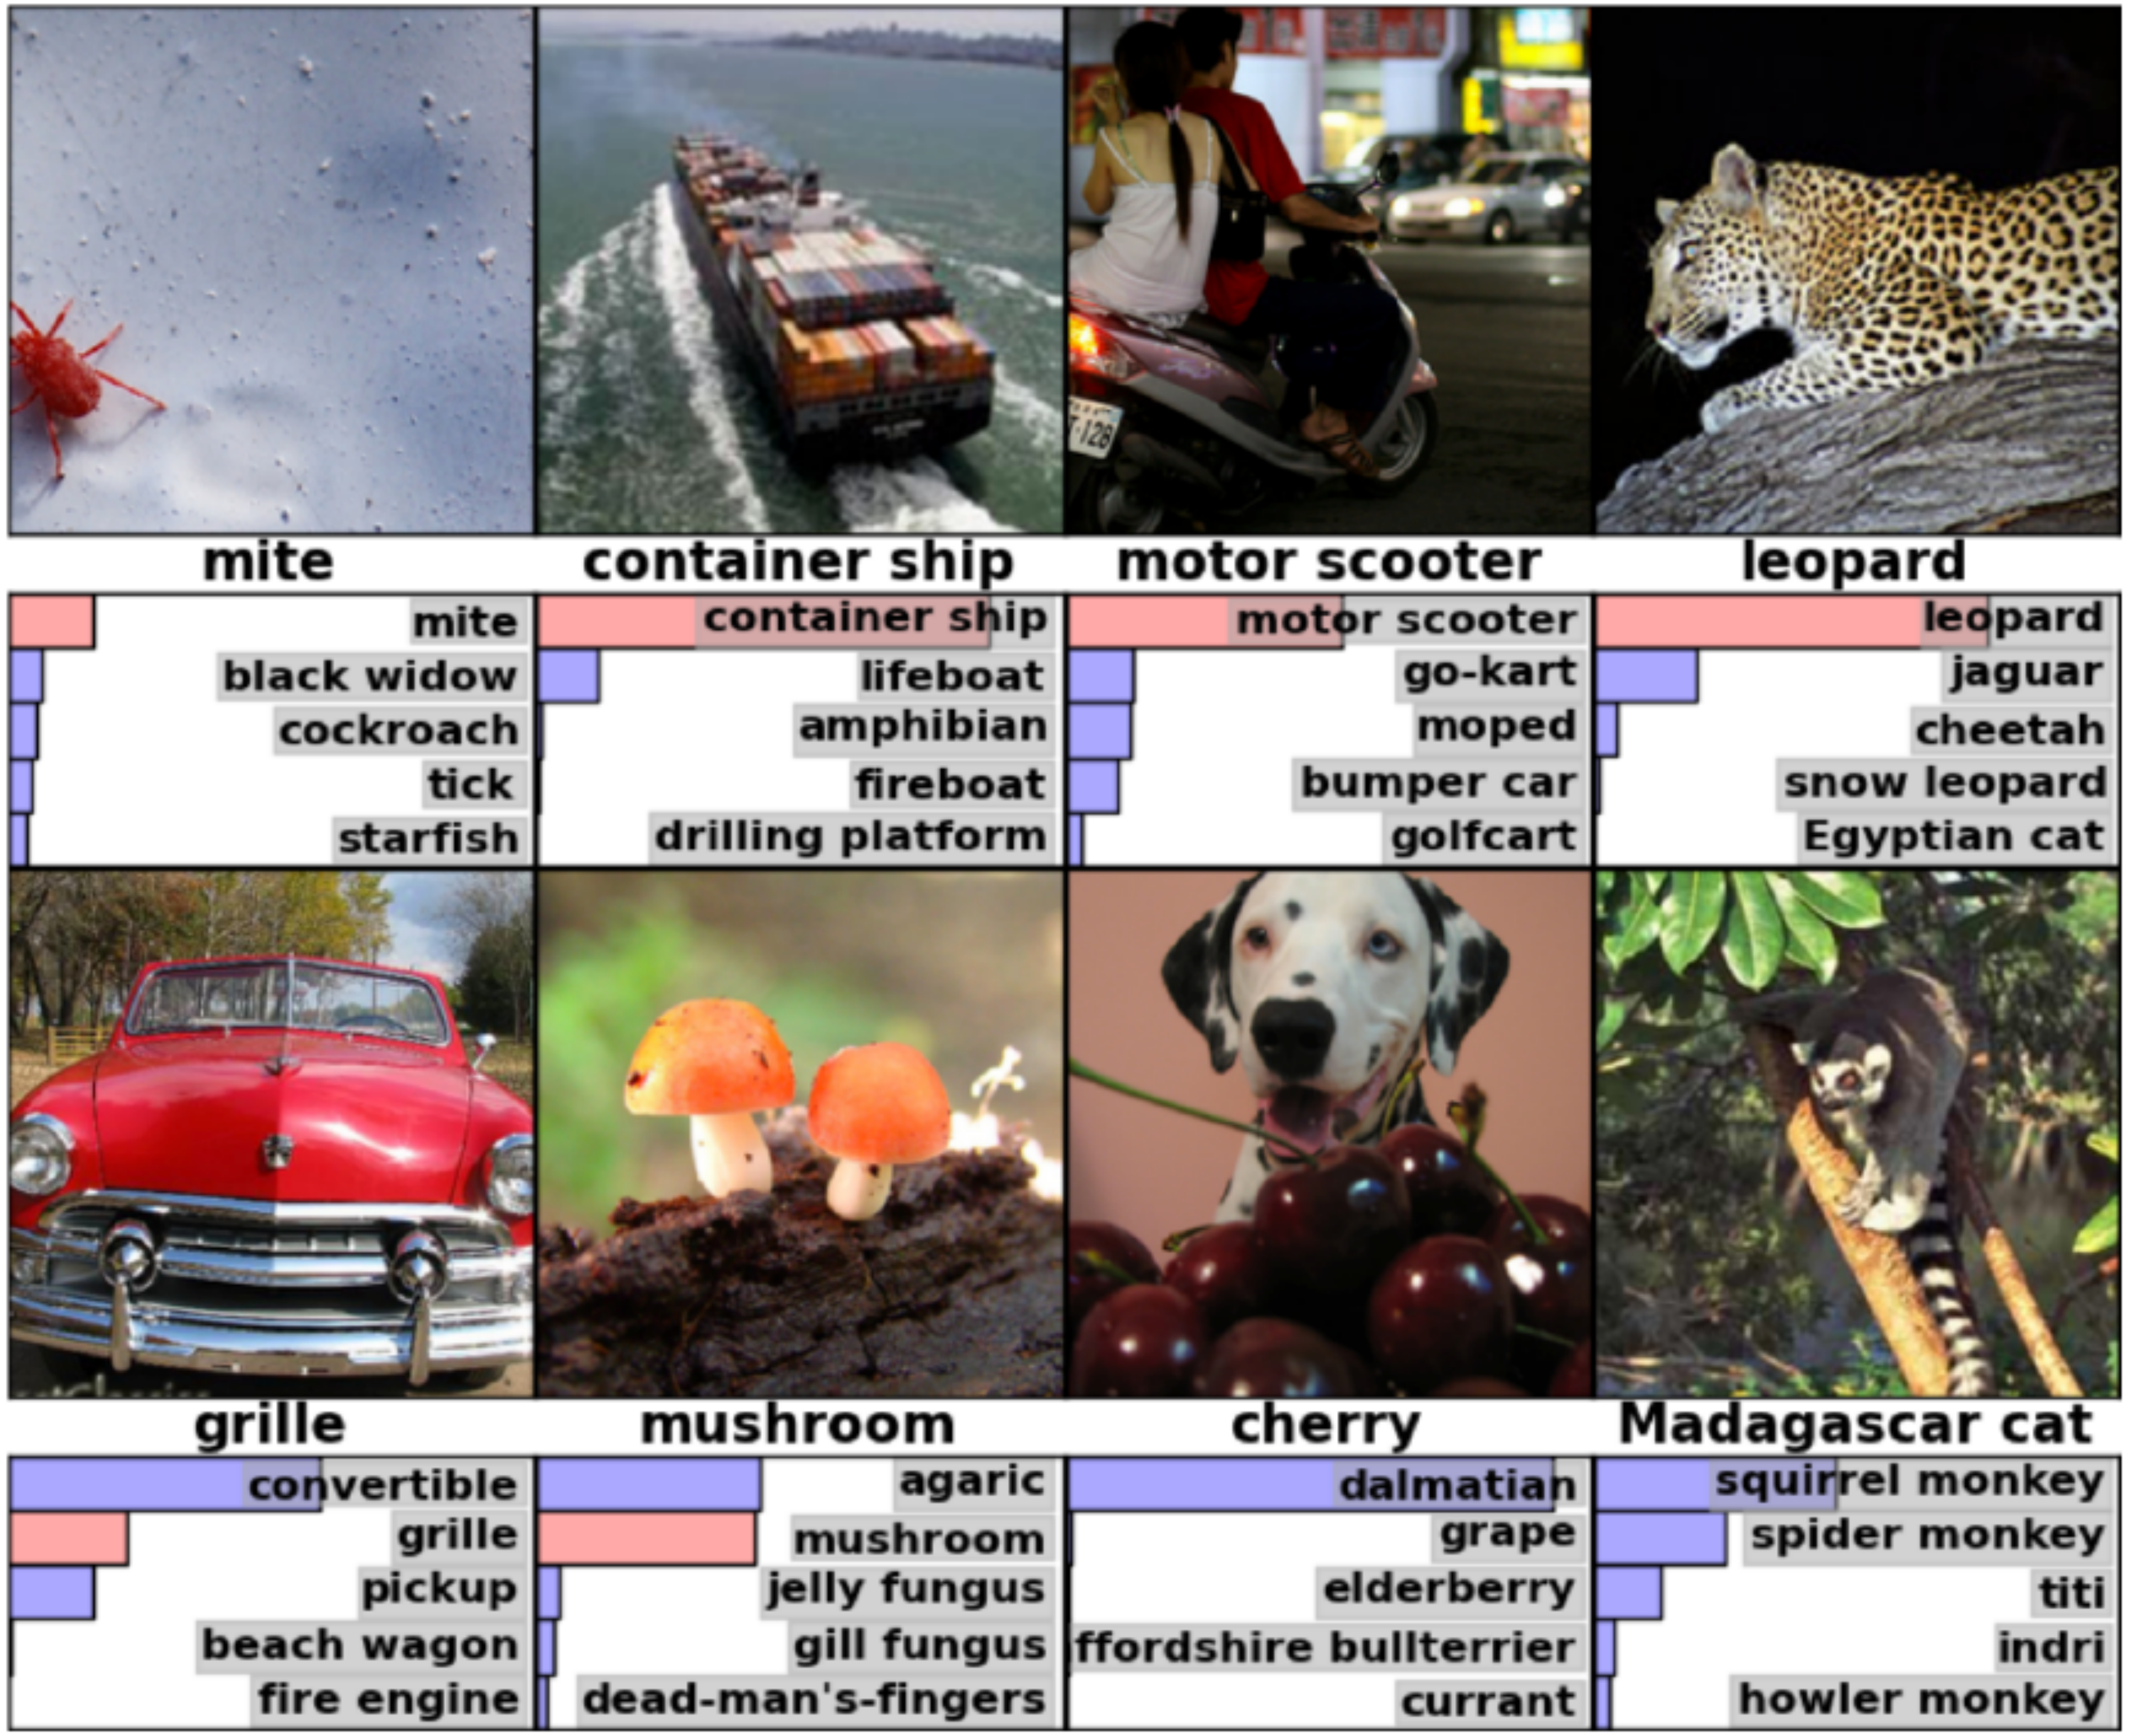

Krizhevsky et al. NIPS 2012

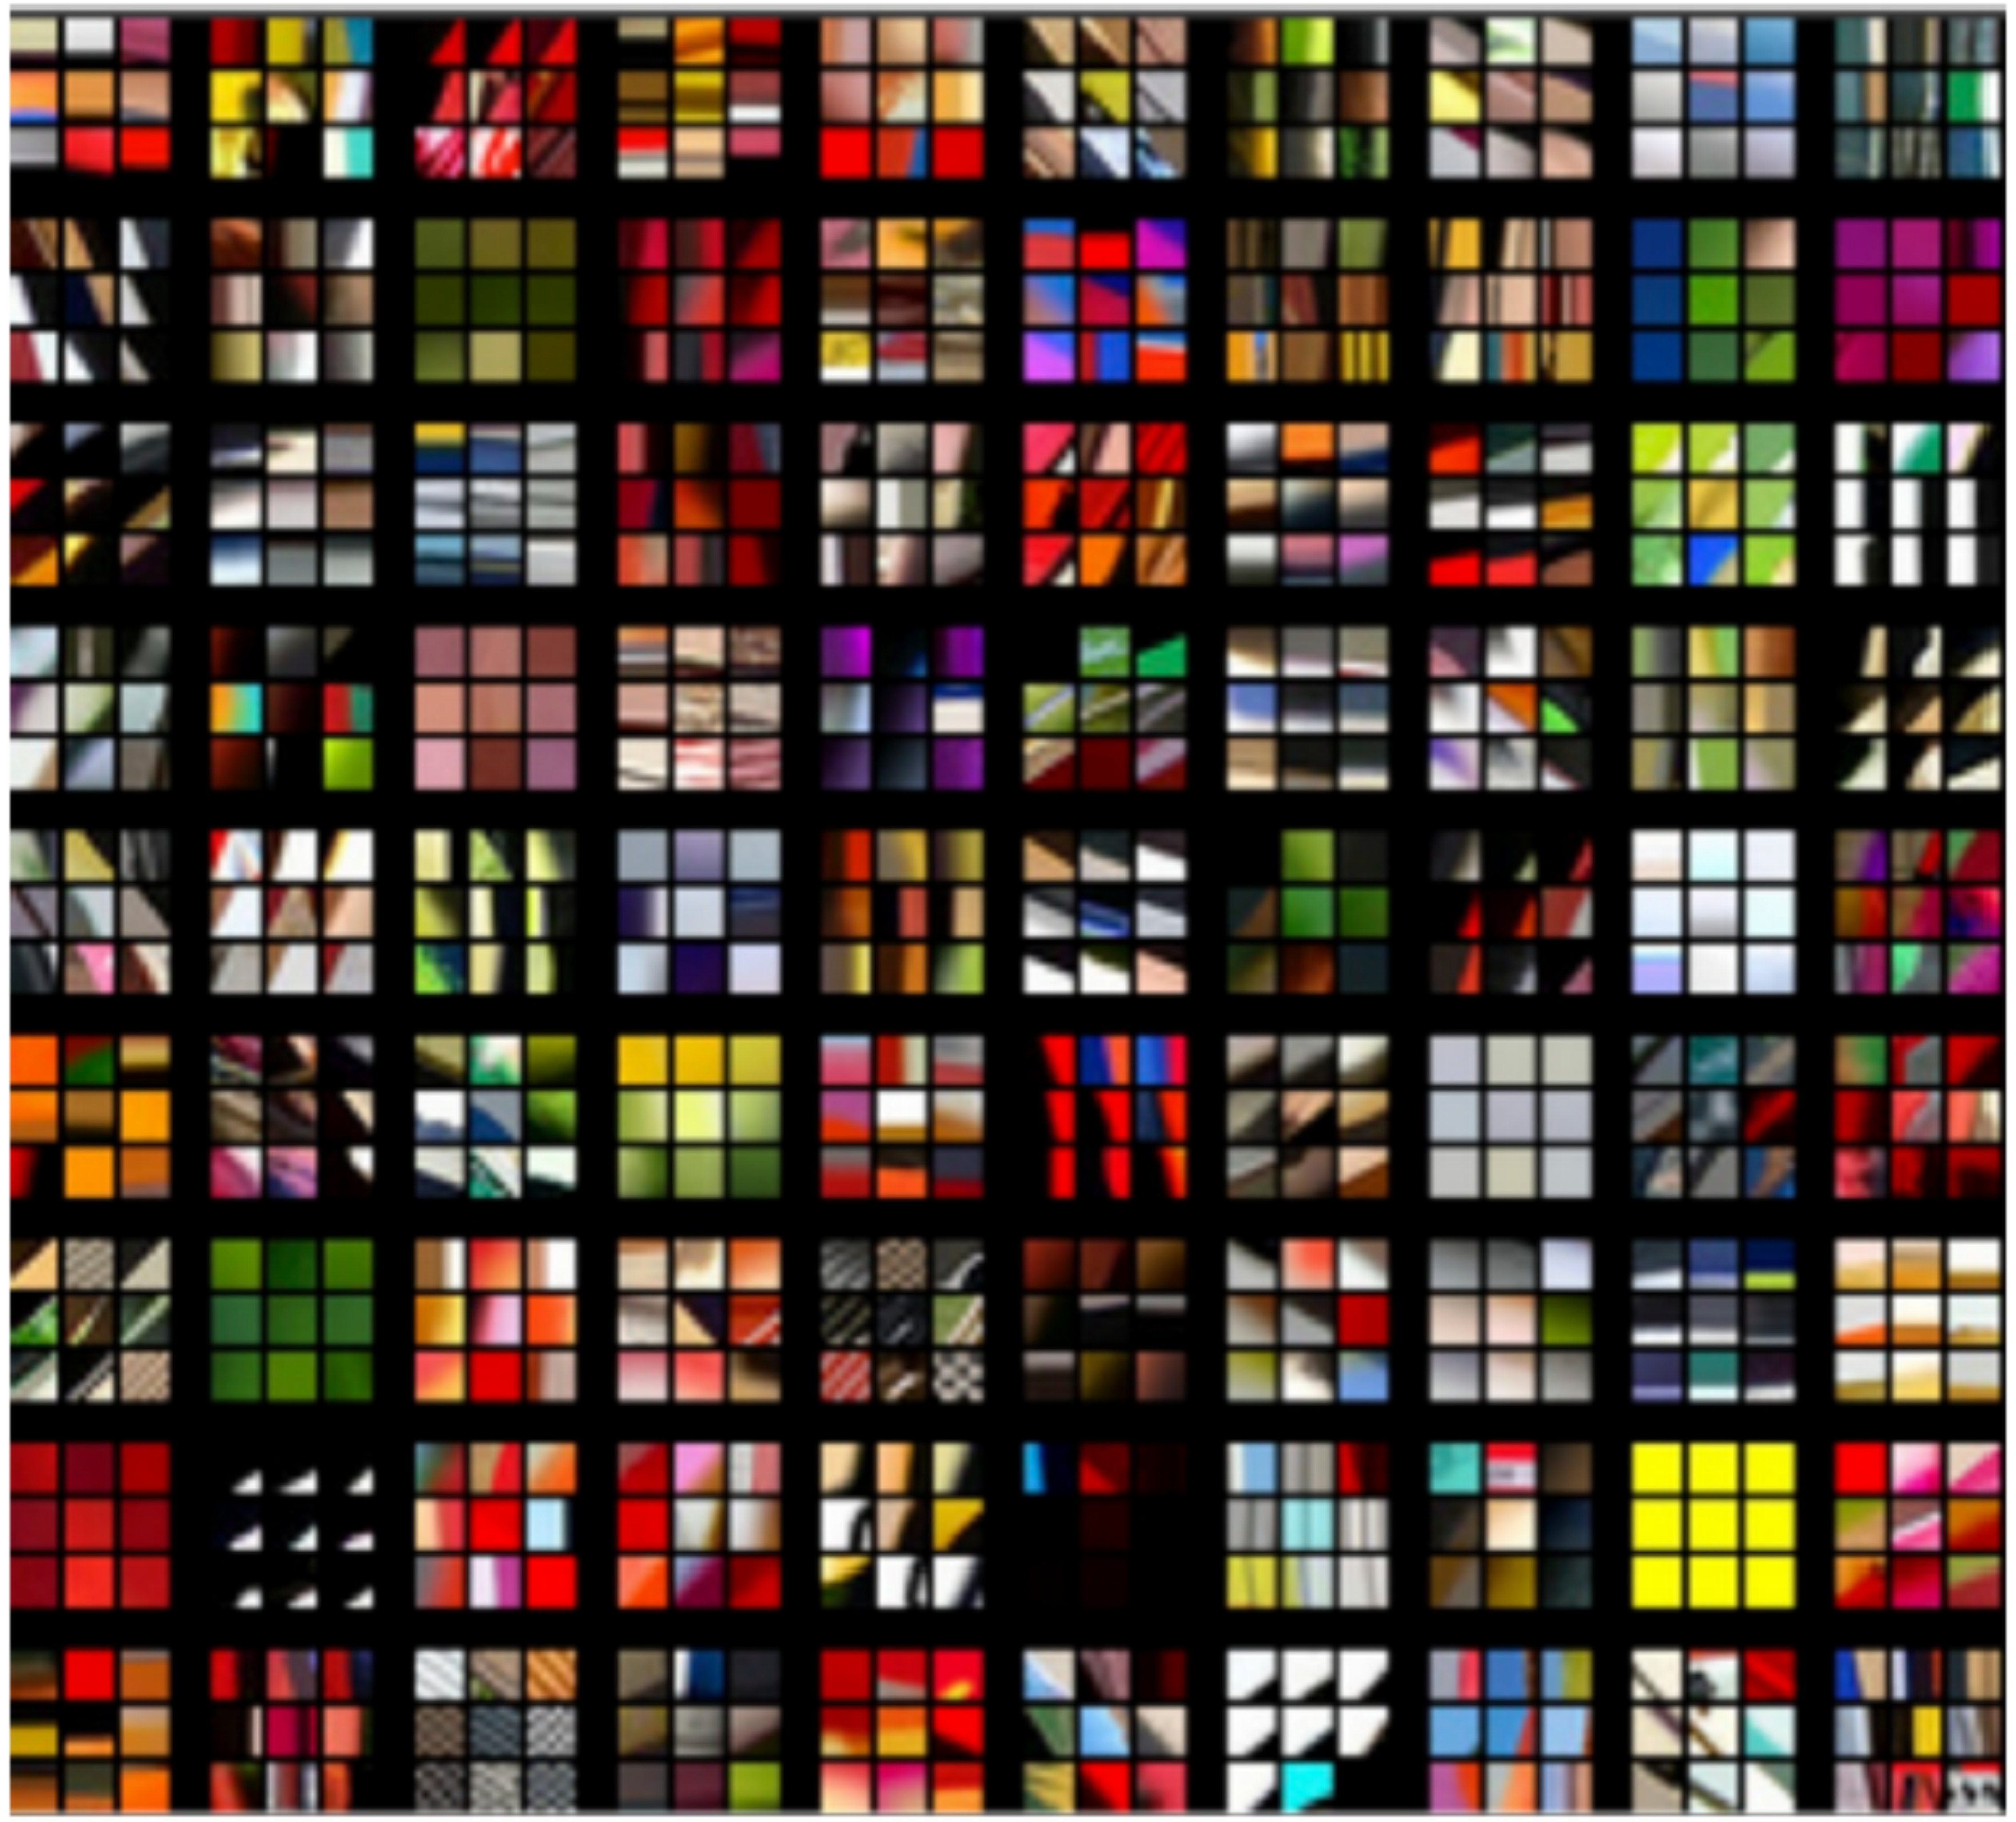

patches that strongly activate first layer filters

Zeiler & Fergus, ECCV 2014

Layer 2

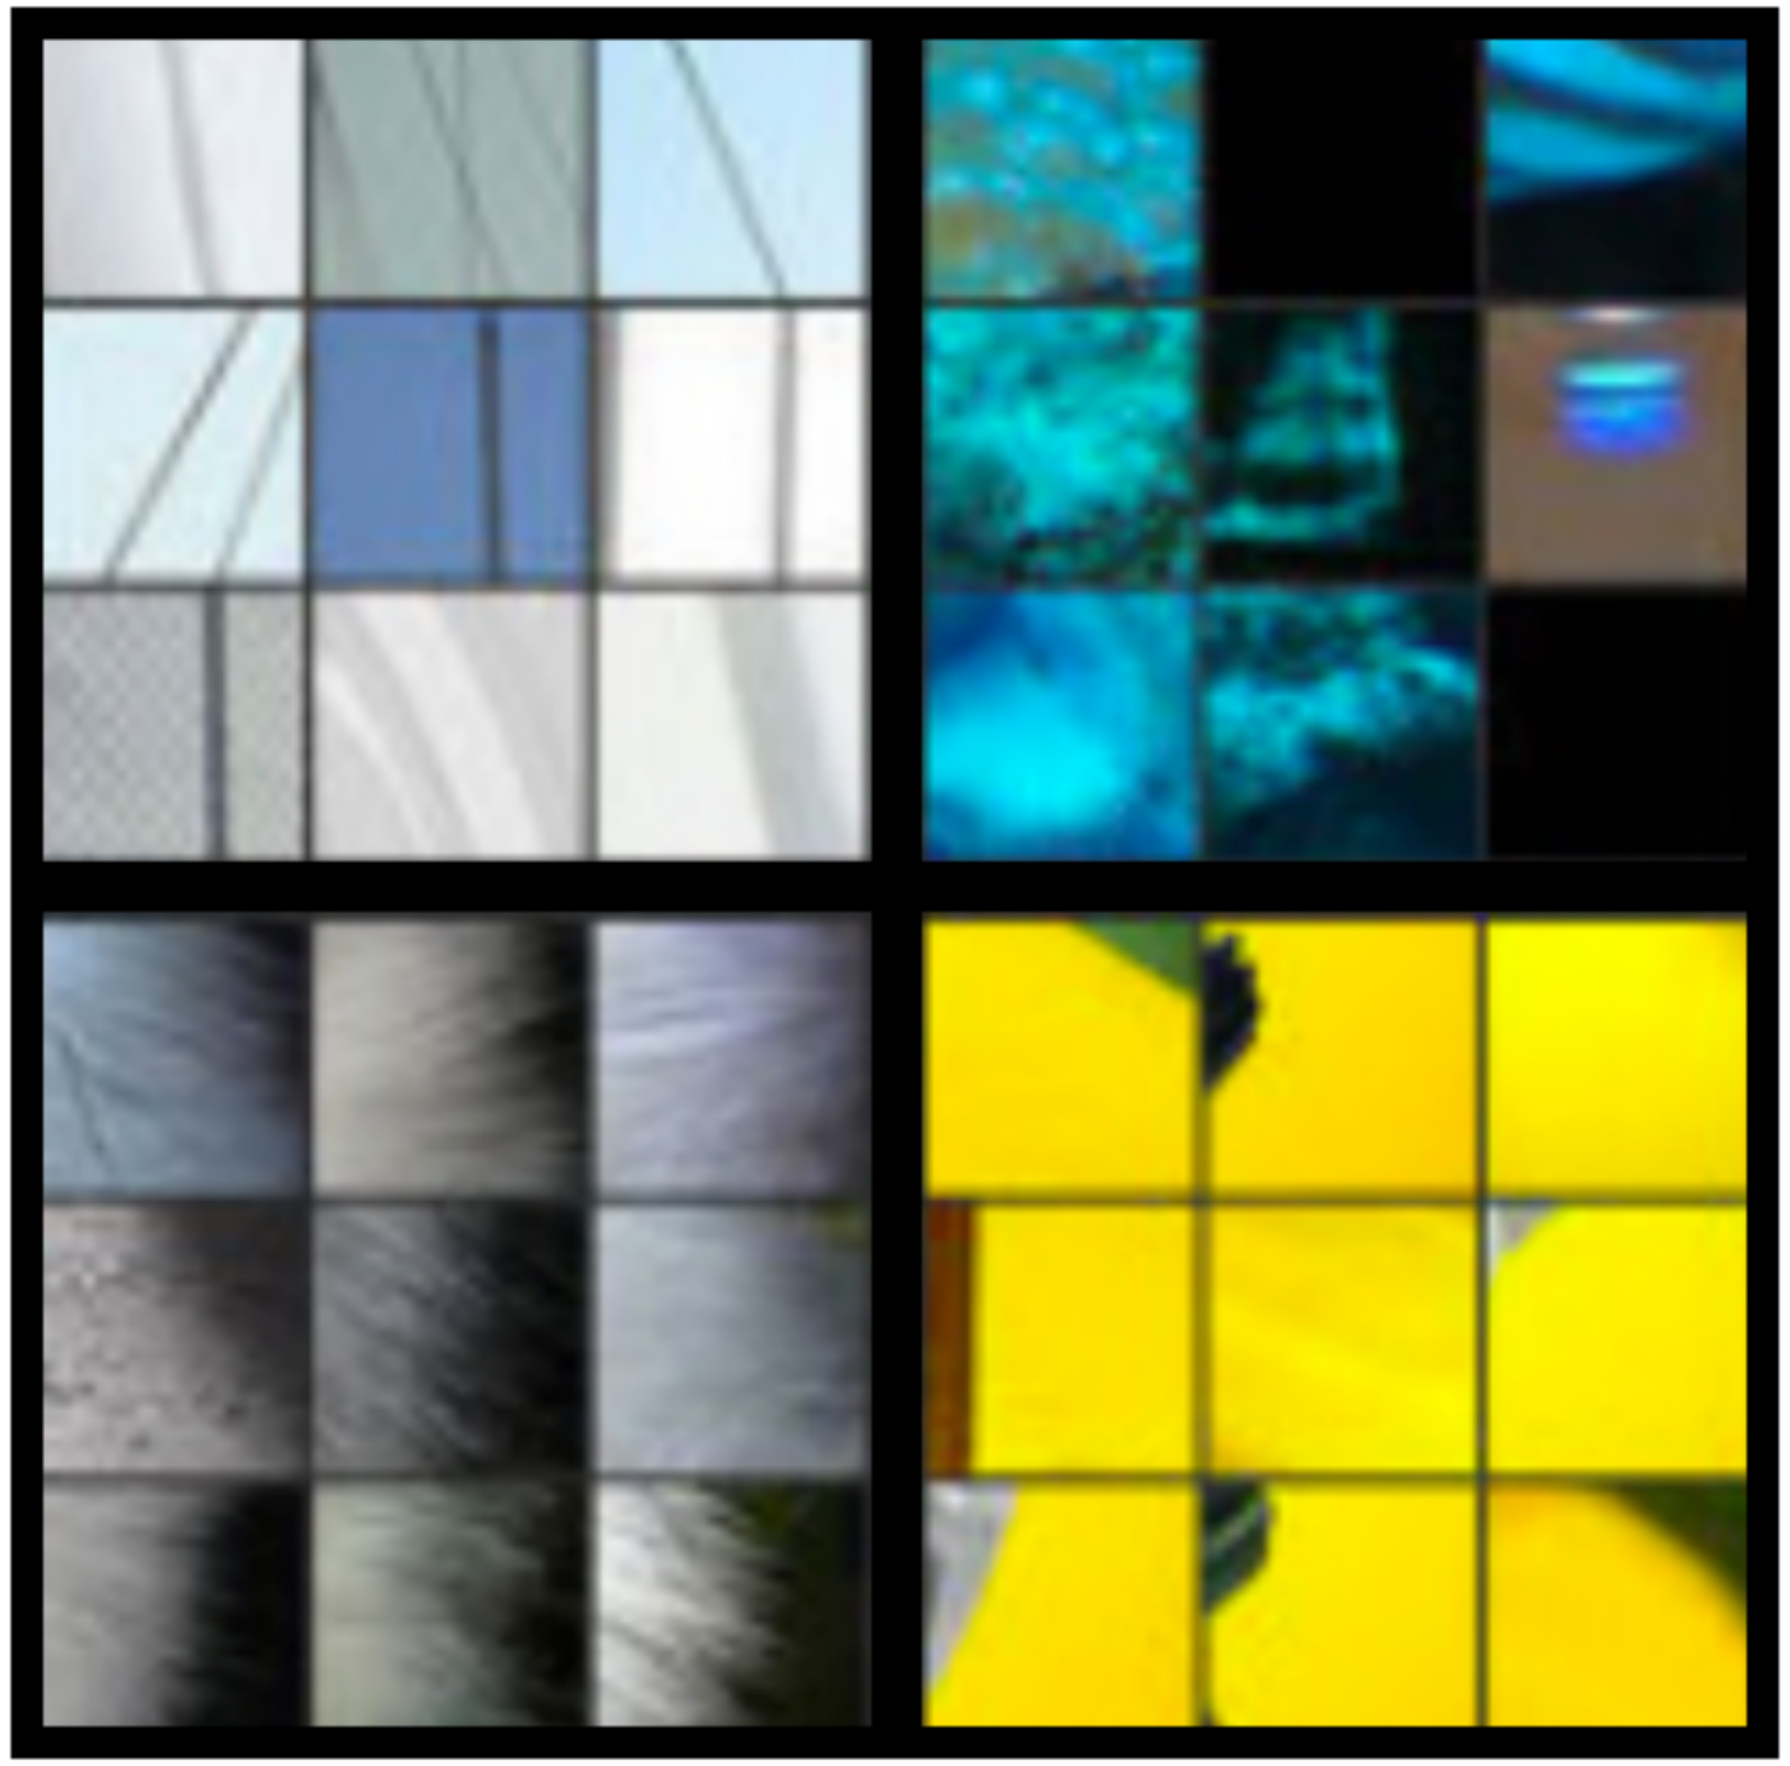

Layer 5

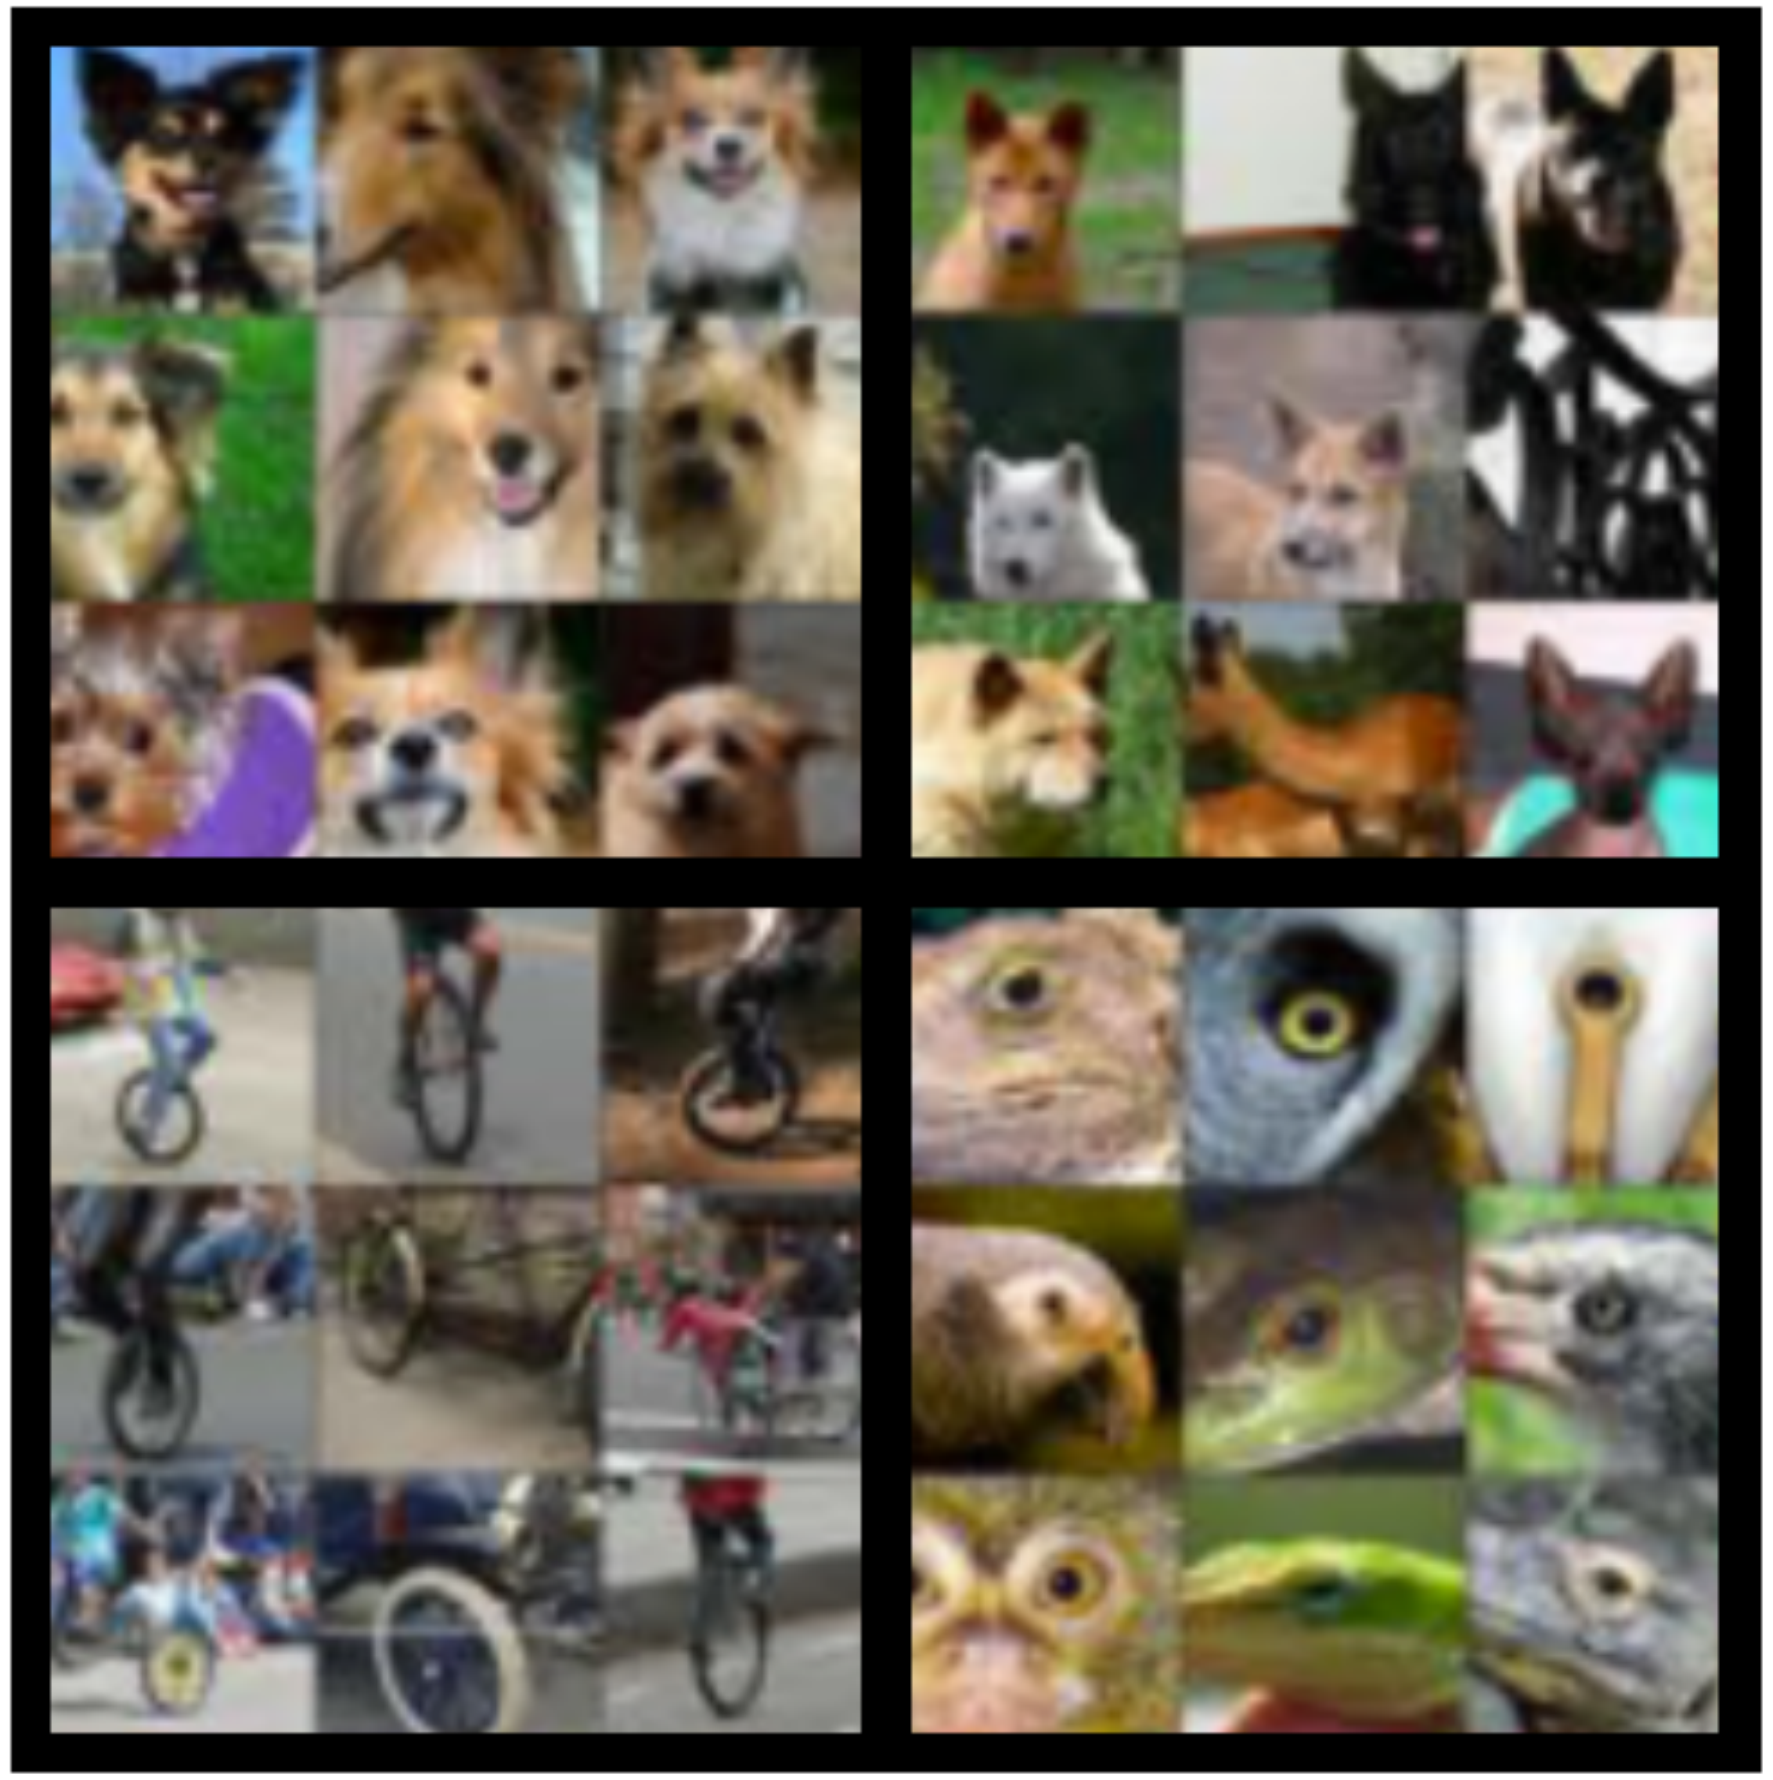

patches that strongly activate neurons on specified layer

Zeiler & Fergus, ECCV 2014

Using goal-driven deep learning models to understand sensory cortex

Daniel L K Yamins[1,2] & James J DiCarlo[1,2]

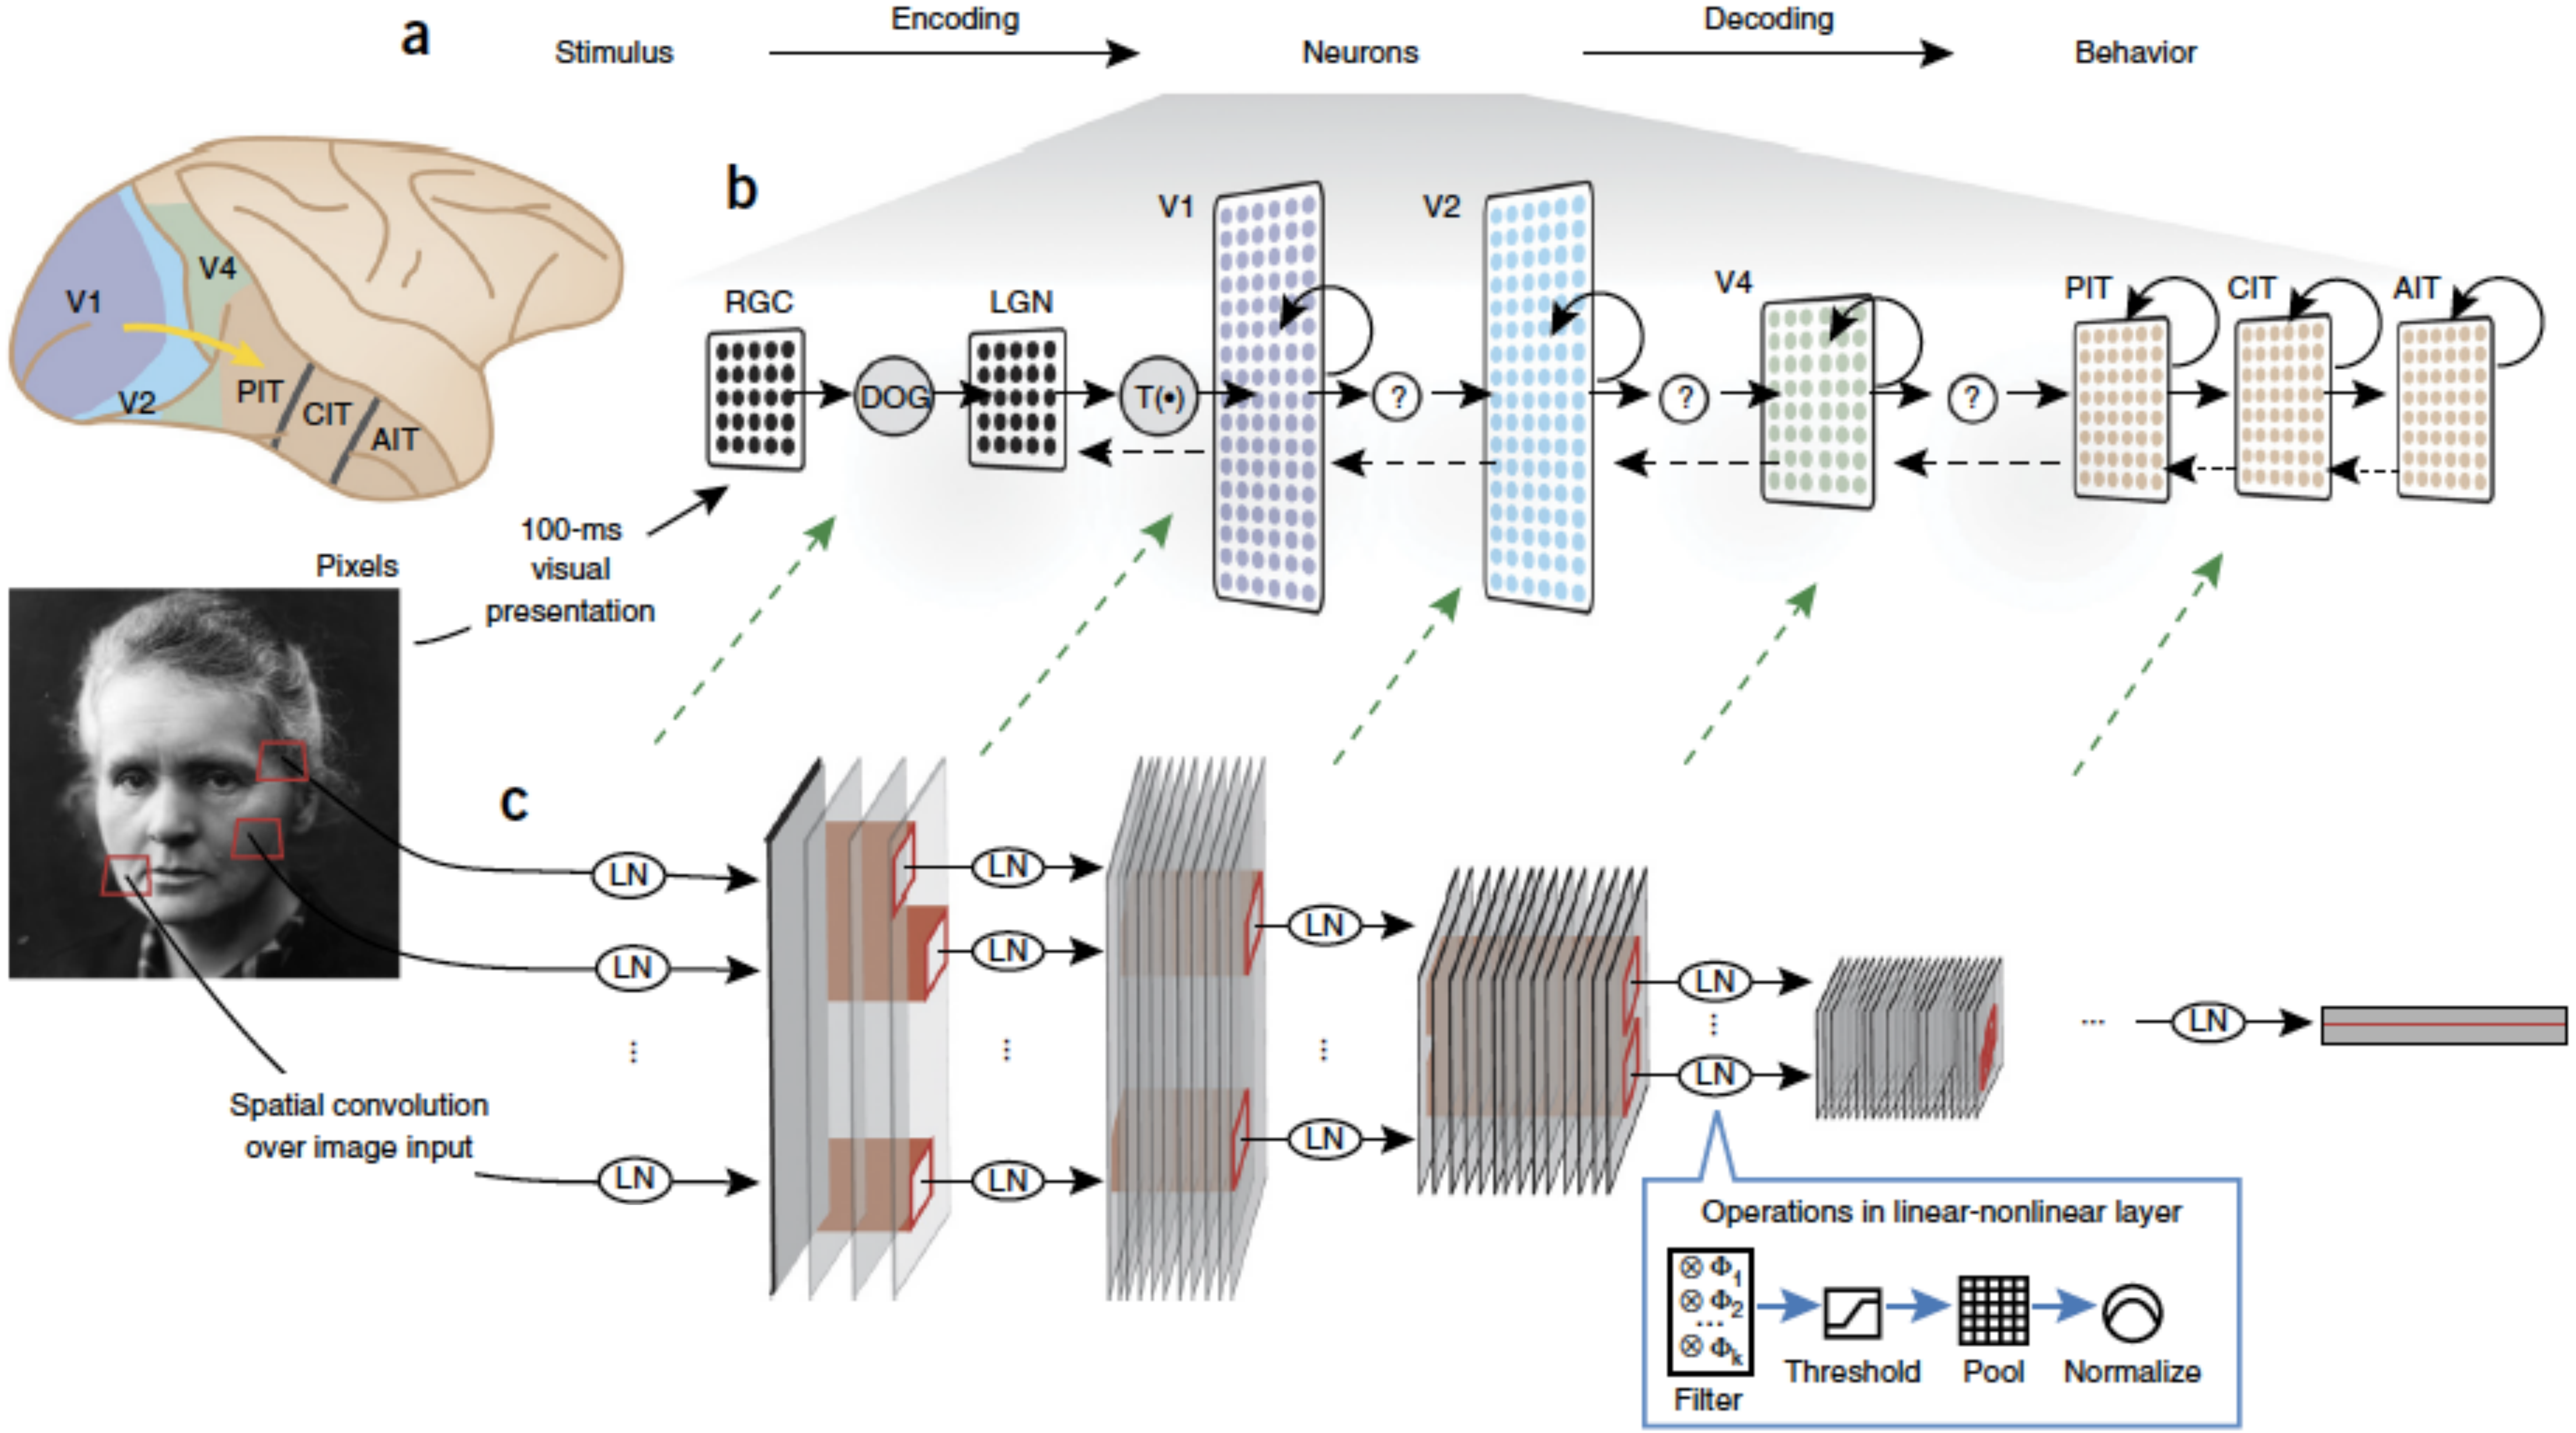

Computation in a neural net

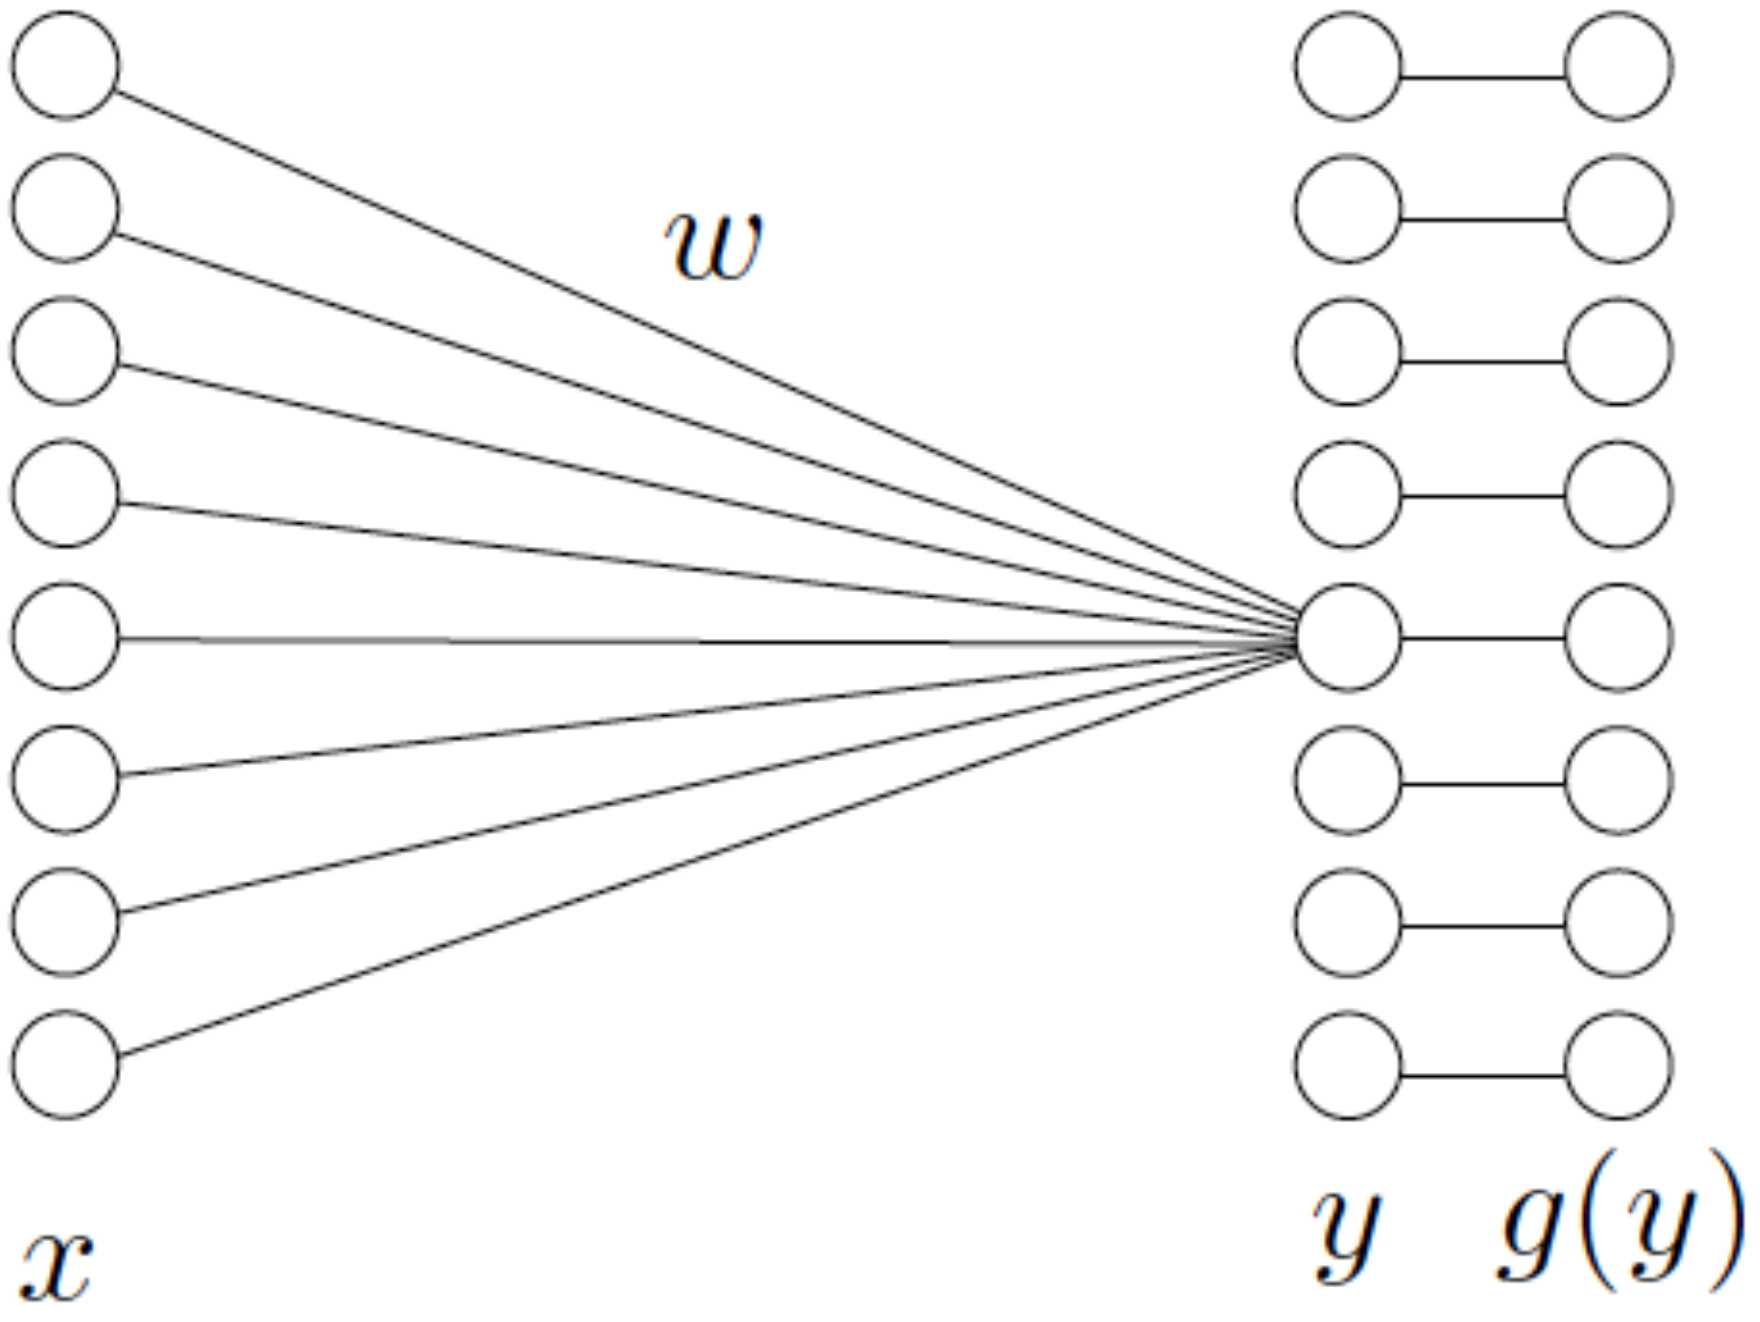

Rectified linear unit (ReLU)

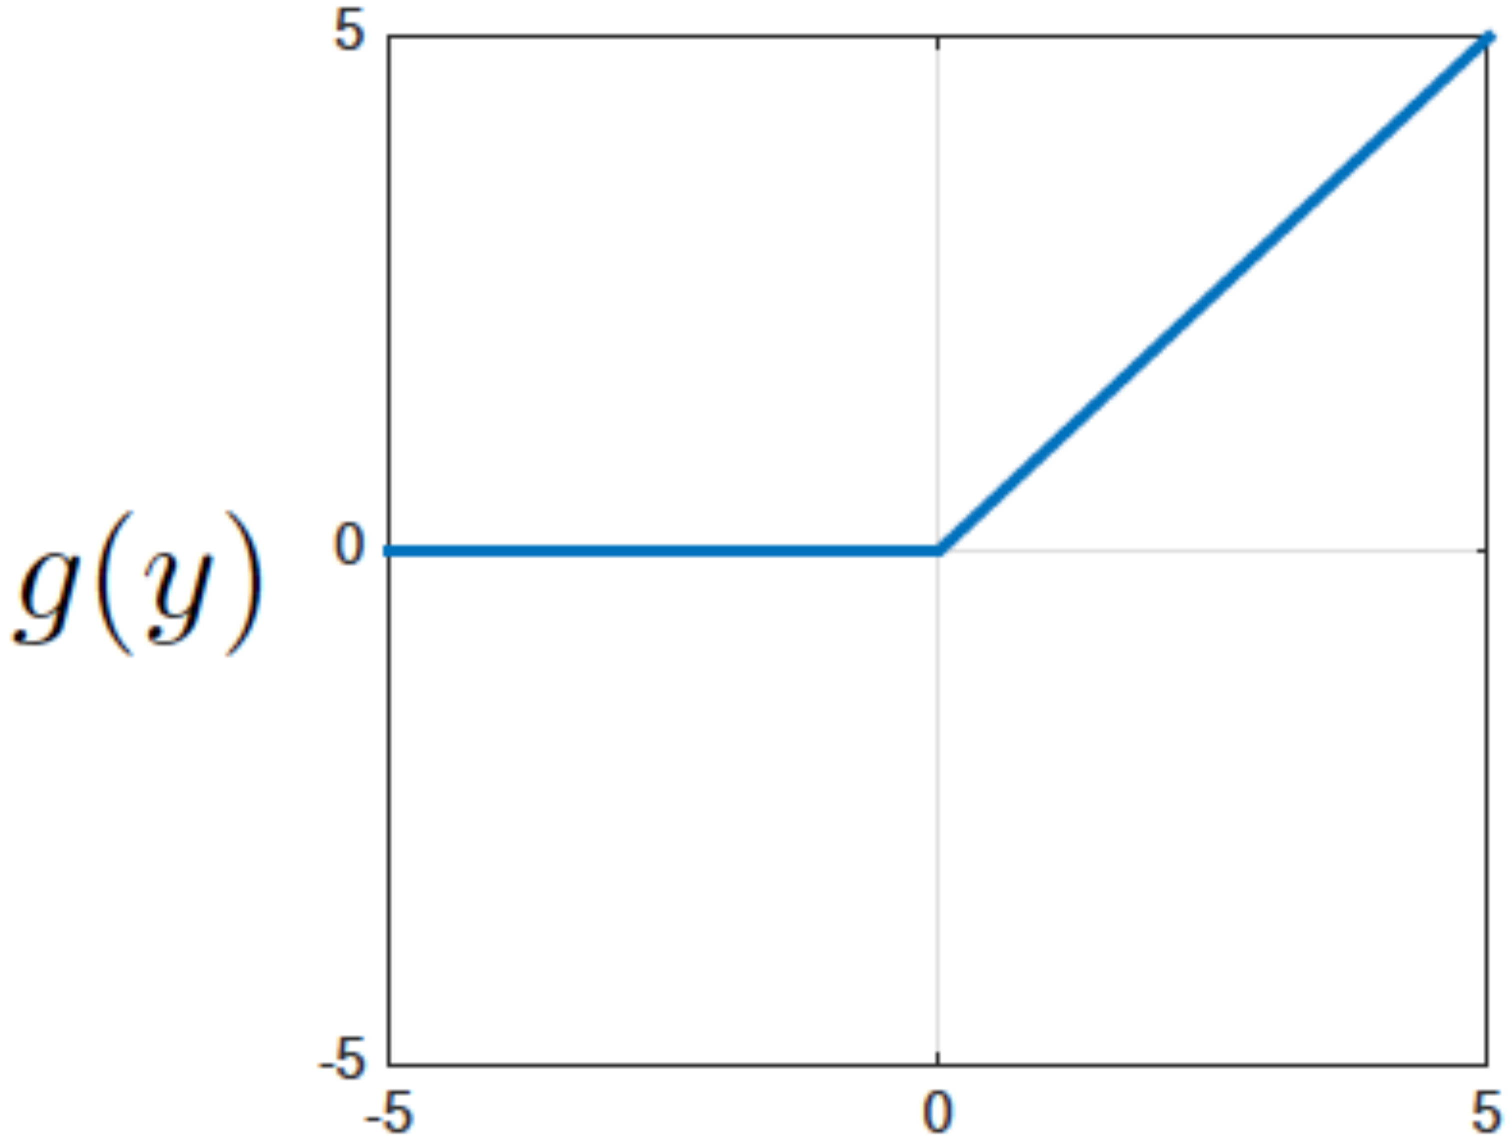

$$g(y) = \max(0, y)$$

Deep nets architecture and SGD training

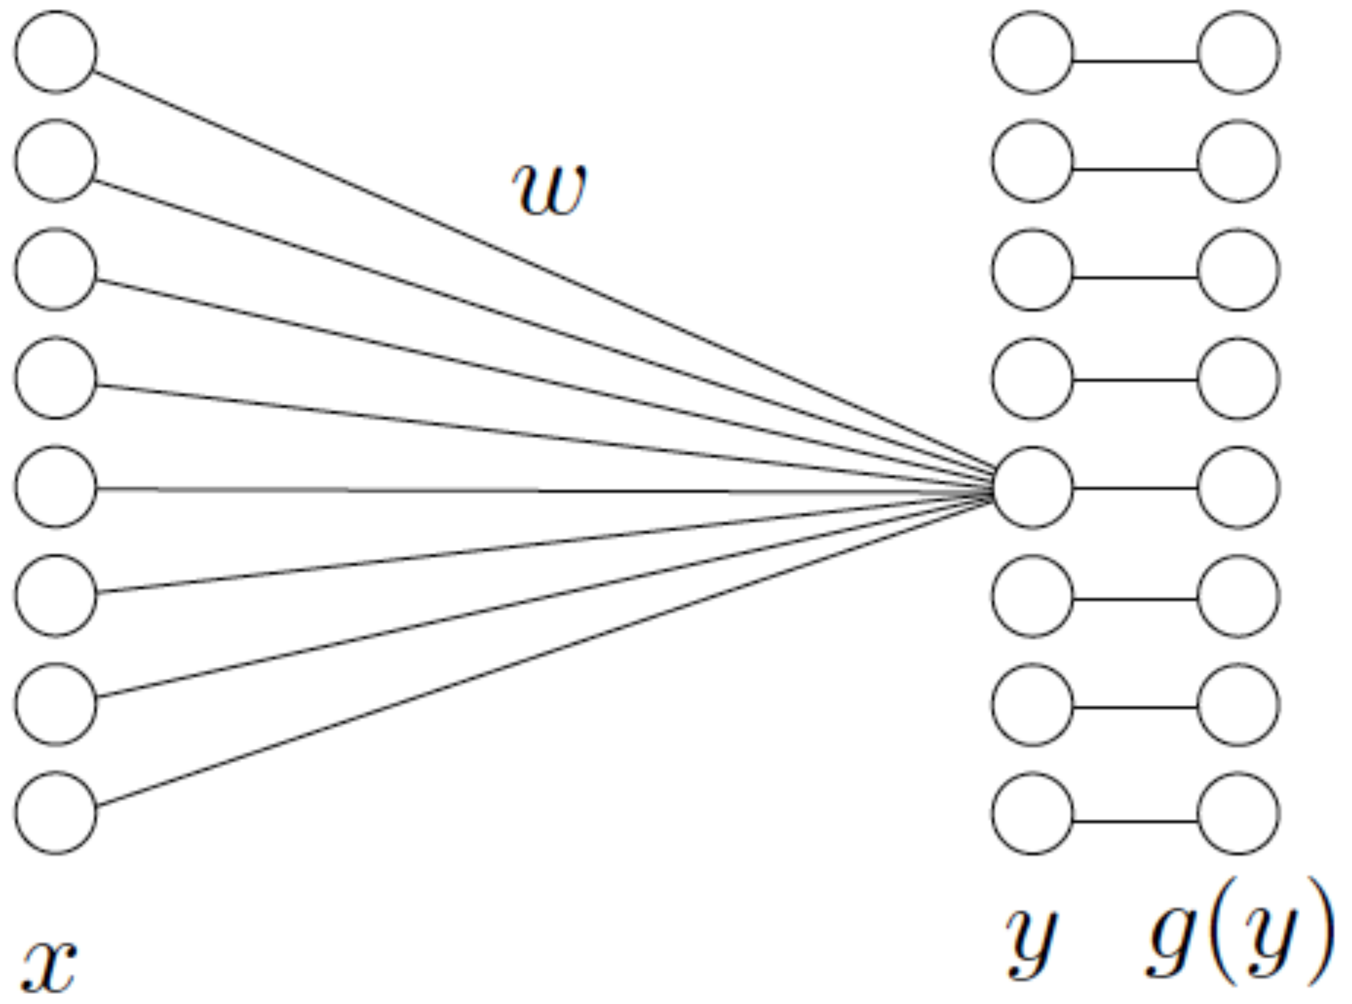

Rectified linear unit (ReLU)

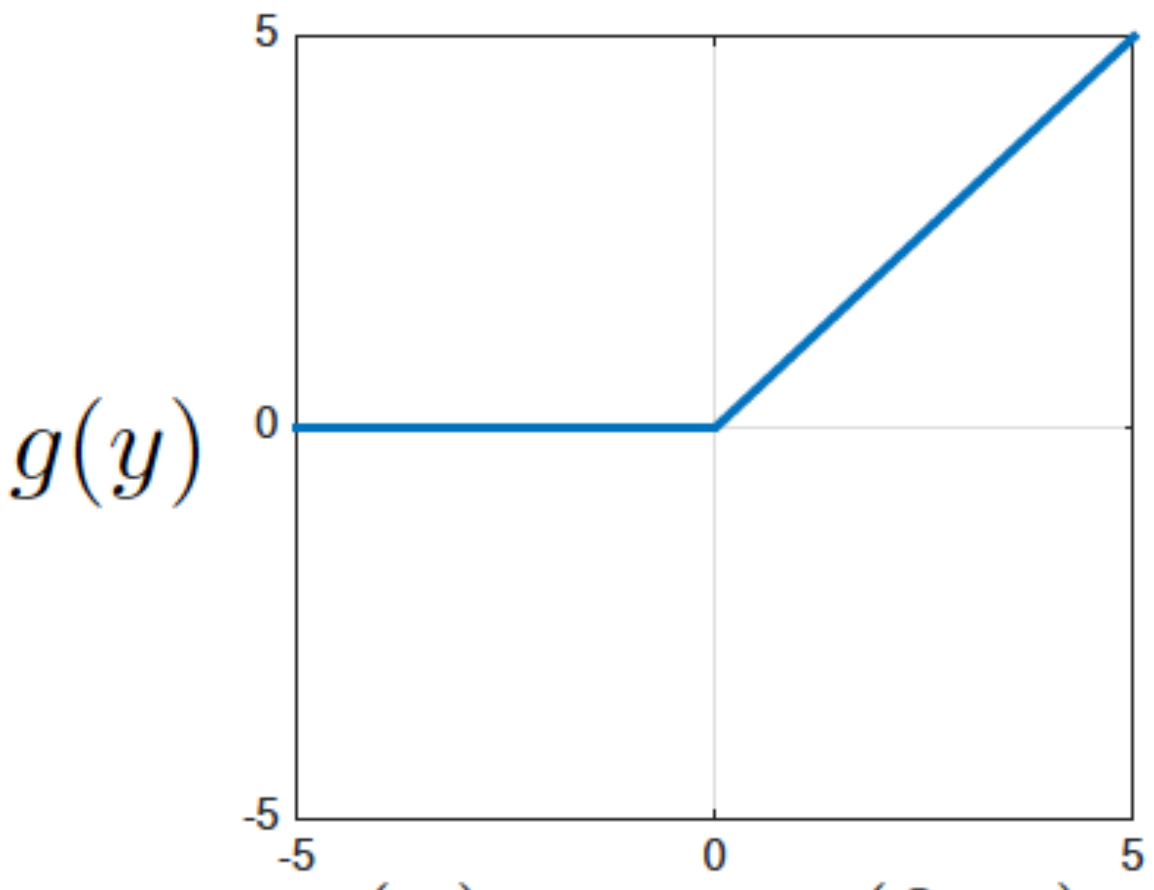

$$g(y) = \max(0, y)$$

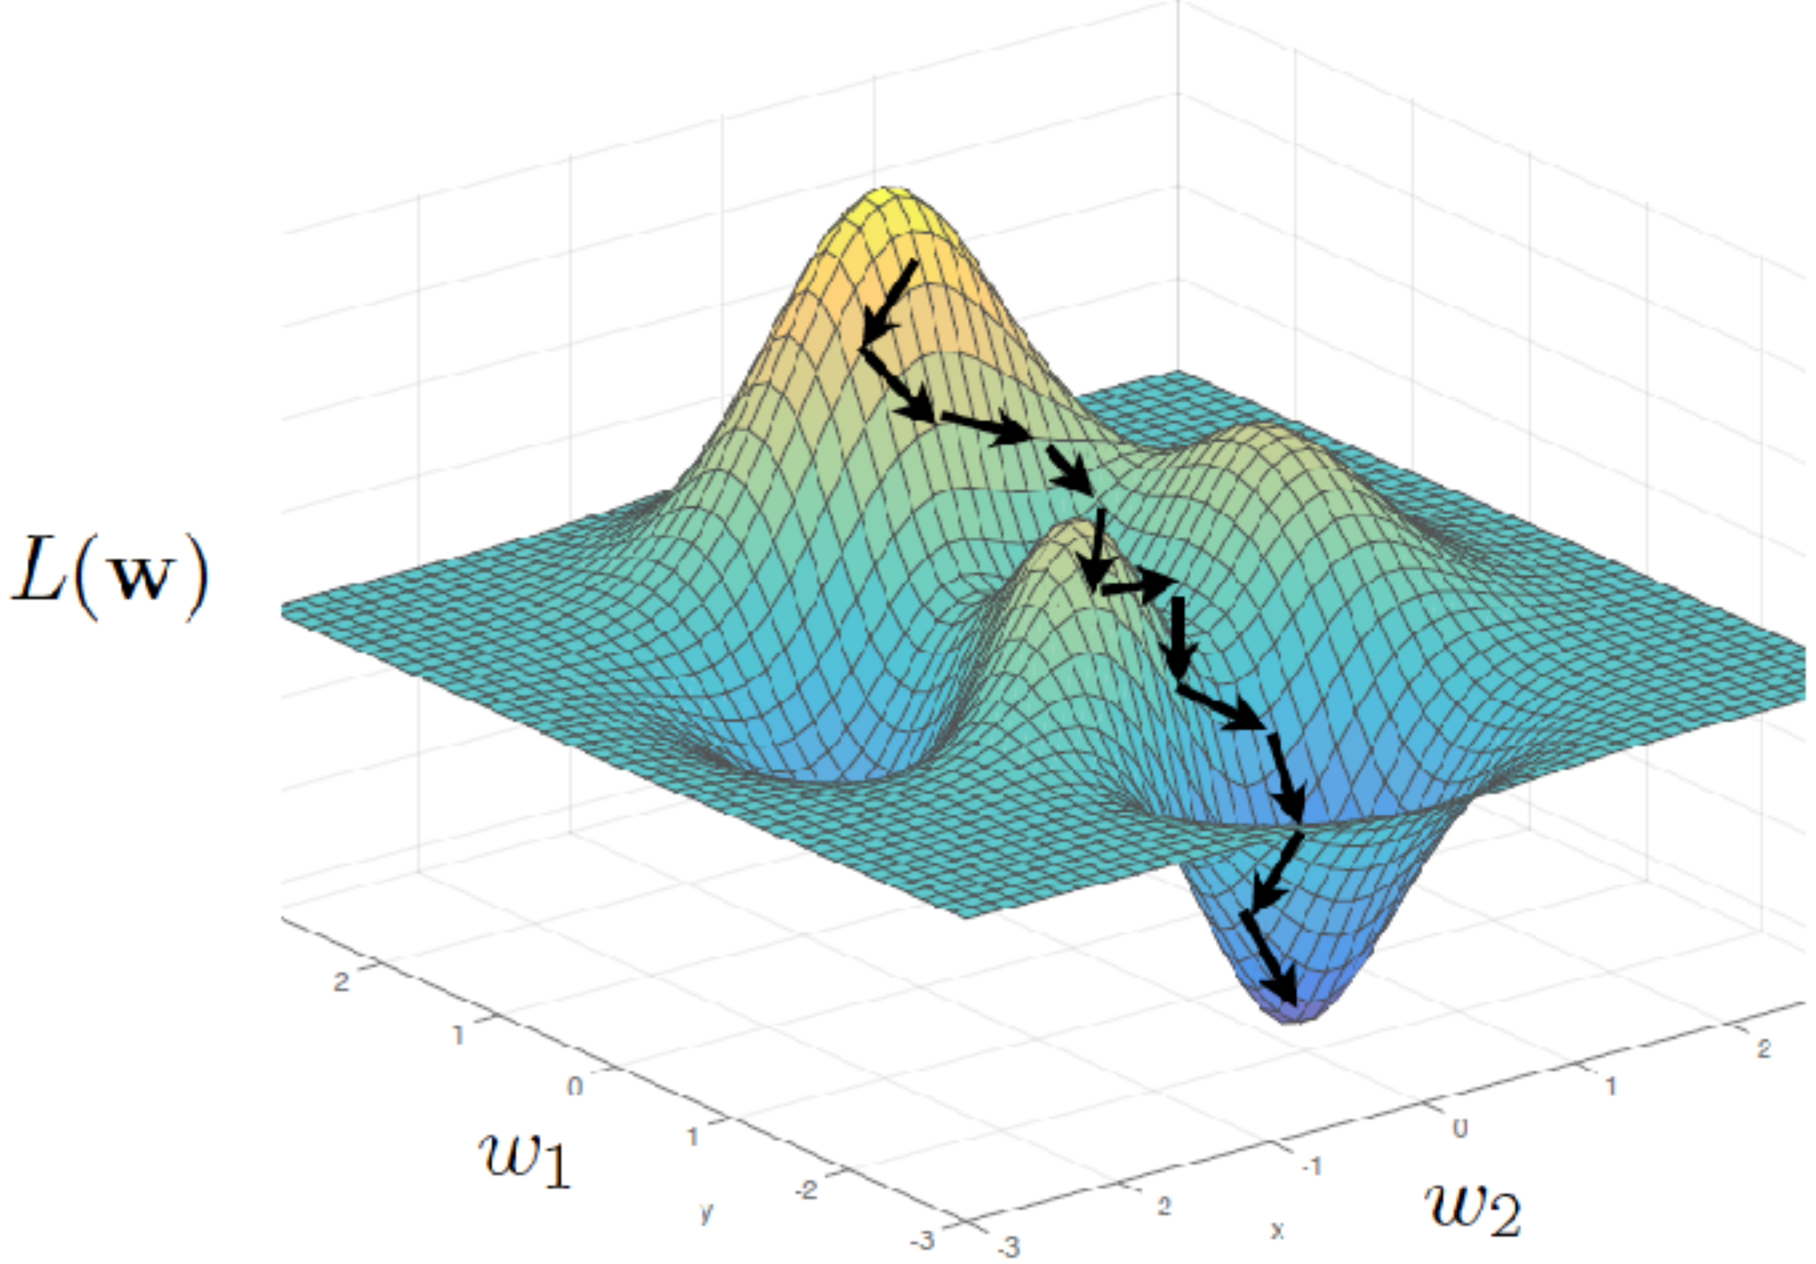

Gradient descent

$$\boxed{\underset{\mathbf{w}}{\operatorname{argmin}} \sum_i \ell(\mathbf{z}_i, f(\mathbf{x}_i; \mathbf{w})) = L(\mathbf{w})}$$

One iteration of gradient descent:

$$\mathbf{w}^{t+1} = \mathbf{w}^t - \eta_t \frac{\partial L(\mathbf{w^t})}{\partial \mathbf{w}}$$

learning rate

Computing the gradients

$$L(\mathbf{x}, \mathbf{w}) = f_L(\dots f_2(f_1(\mathbf{x}; w_1); w_2) \dots)$$

Chain rule

$$\frac{\partial L}{\partial \mathbf{w}_i} = \frac{\partial f_L}{\partial f_{L-1}} \frac{\partial f_{L-1}}{\partial f_{L-2}} \dots \frac{\partial f_i}{\partial \mathbf{w}_i}$$

This can be computed efficiently using **back-propagation**.

First run the net forward to see what kinds of errors it makes.

Then propagate errors back, using the chain rule, to tell each layer how to adjust its weights to reduce the error.

Theory I:
Why and when are deep networks better than shallow networks?

$$f(x_1,x_2,\ldots,x_8)=g_3(g_{21}(g_{11}(x_1,x_2),g_{12}(x_3,x_4))g_{22}(g_{11}(x_5,x_6),g_{12}(x_7,x_8)))$$

$$g(x)=\sum_{i=1}^{r} c_i\left|<w_i,x>+b_i\right|_+$$

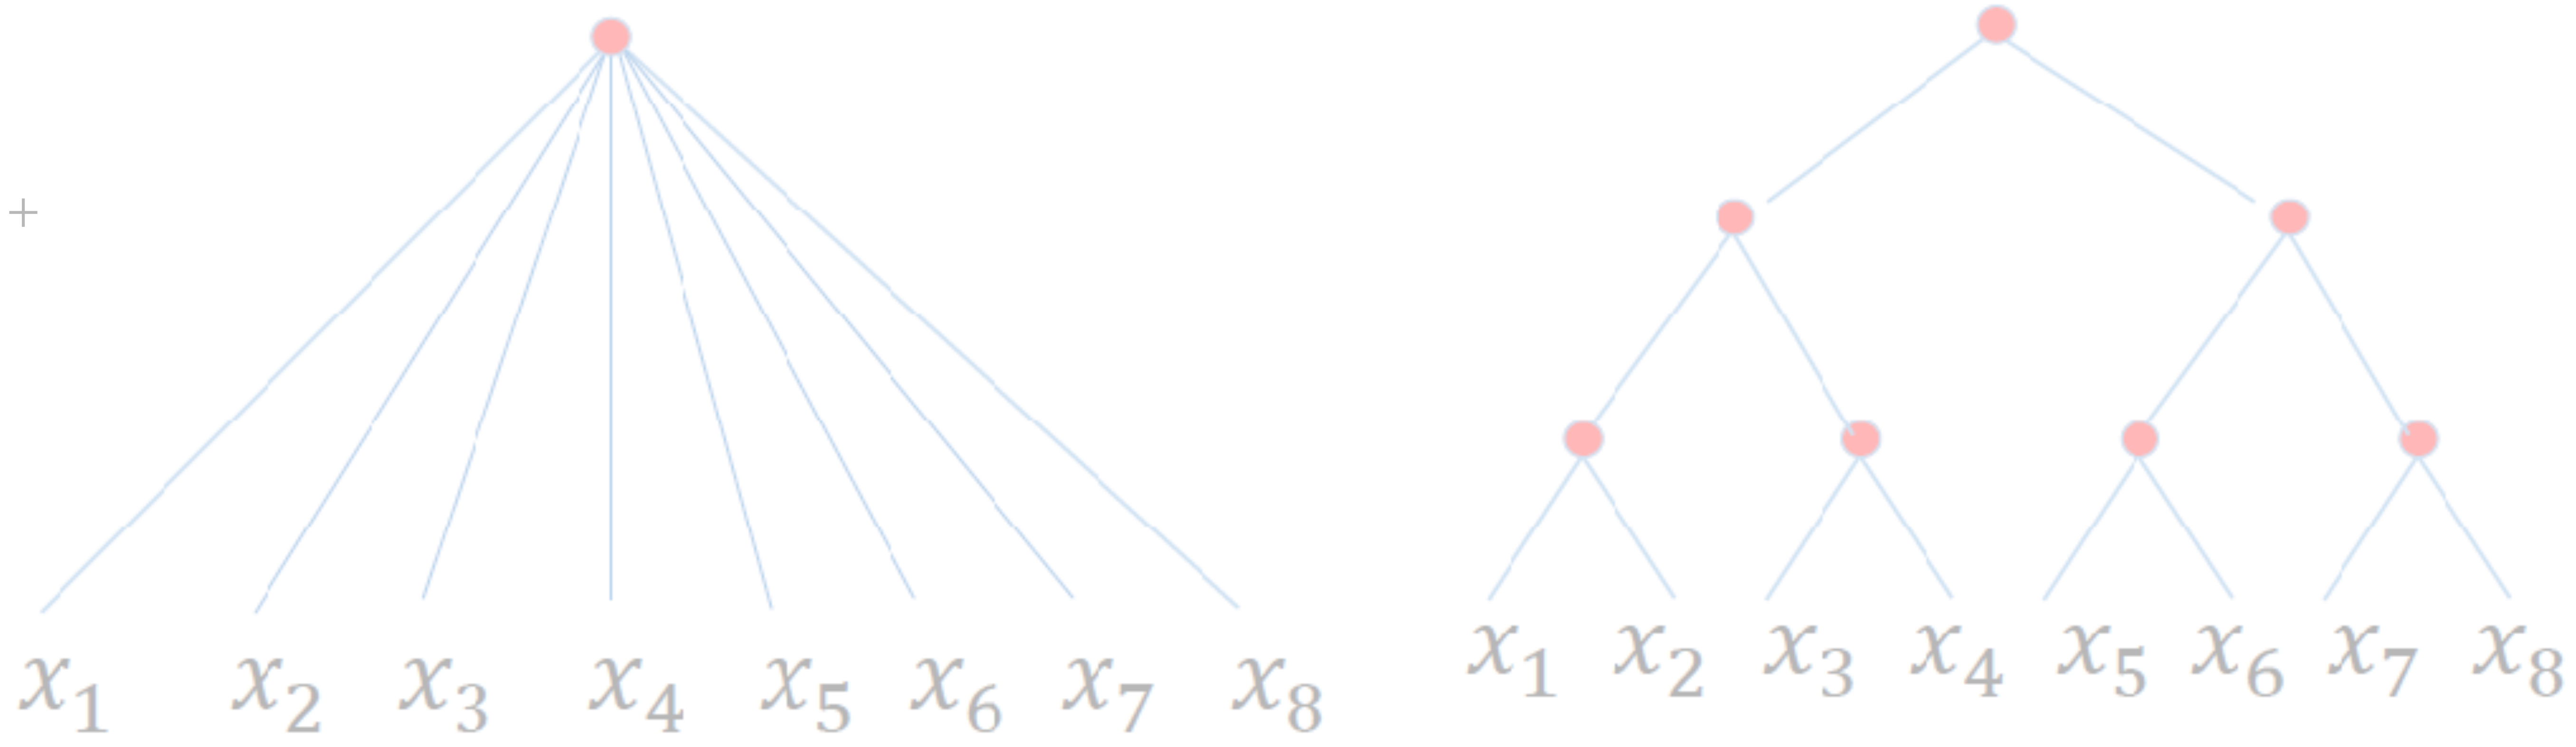

Theorem (informal statement)

Suppose that a function of d variables is compositional . Both shallow and deep network can approximate f equally well. The number of parameters of the shallow network depends exponentially on d as $O(\varepsilon^{-d})$with the dimension whereas for the deep network dance is dimension independent, i.e. $O(\varepsilon^{-2})$

Deep and shallow networks: universality

Theorem *Shallow, one-hidden layer networks with a nonlinear $\phi(x)$ which is not a polynomial are universal. Arbitrarily deep networks with a nonlinear $\phi(x)$ (including polynomials) are universal.*

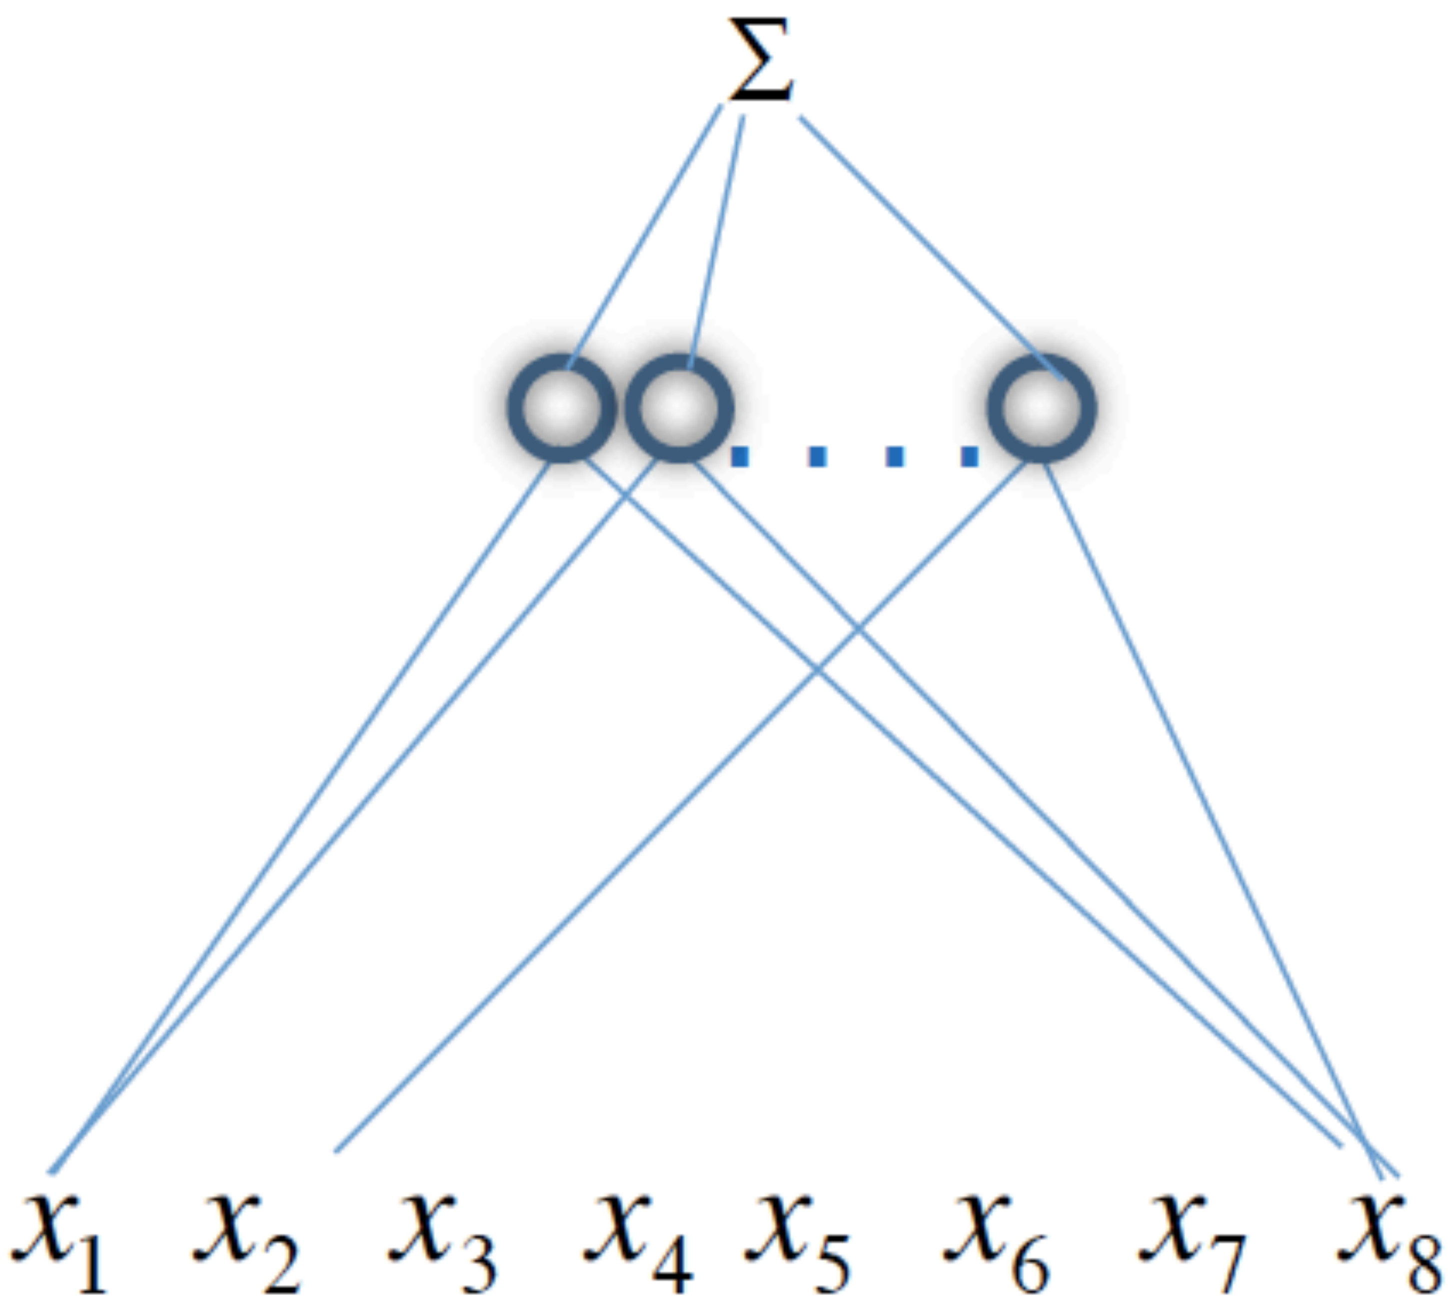

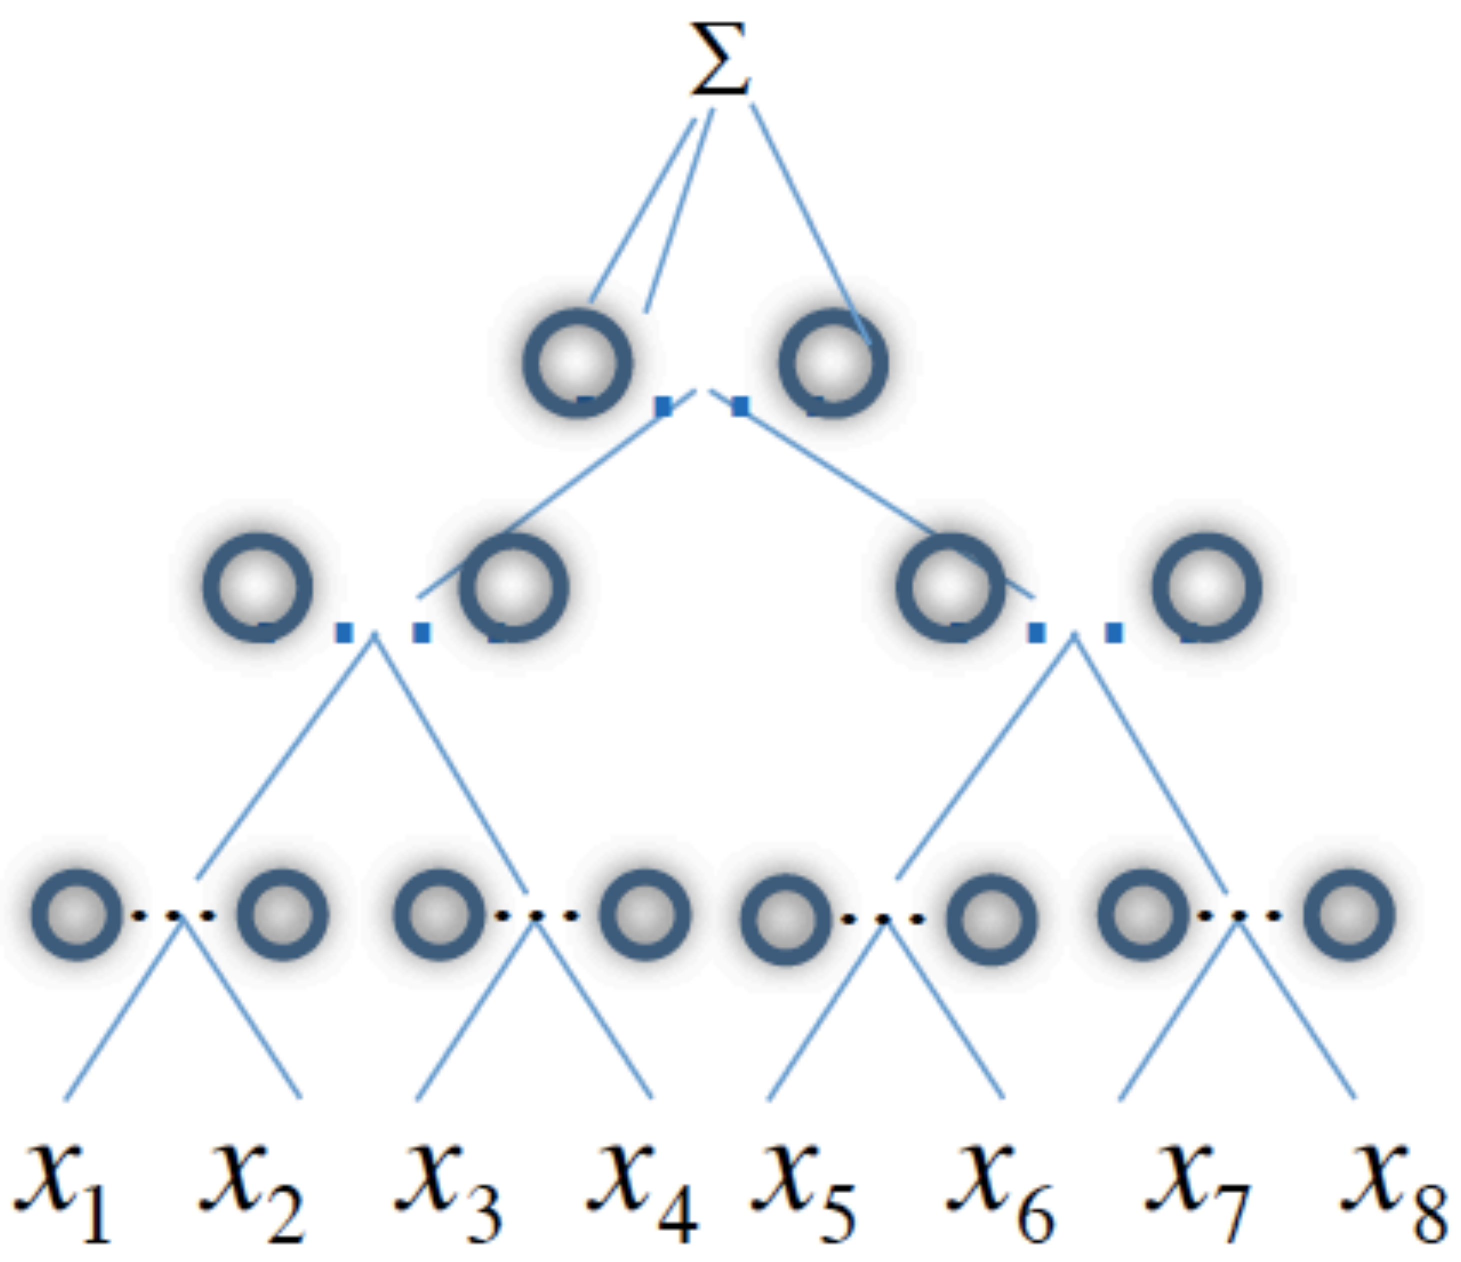

$$\phi(x) = \sum_{i=1}^{r} c_i \left| < w_i, x > + b_i \right|_+$$

Cybenko, Girosi, ….

Curse of dimensionality

$$y = f(x_1, x_2, \ldots, x_8)$$

Both shallow and deep network can approximate a function of d variables equally well. The number of parameters in both cases depends exponentially on d as $O(\varepsilon^{-d})$.

CENTER FOR Brains Minds+ Machines

Mhaskar, Poggio, Liao, 2016

When can the curse of dimensionality be avoided

Generic functions

$$f(x_1, x_2, \ldots, x_8)$$

Compositional functions

$$f(x_1, x_2, \ldots, x_8) = g_3(g_{21}(g_{11}(x_1, x_2), g_{12}(x_3, x_4))g_{22}(g_{11}(x_5, x_6), g_{12}(x_7, x_8)))$$

Mhaskar, Poggio, Liao, 2016

Microstructure of compositionality

target function

approximating
function/network

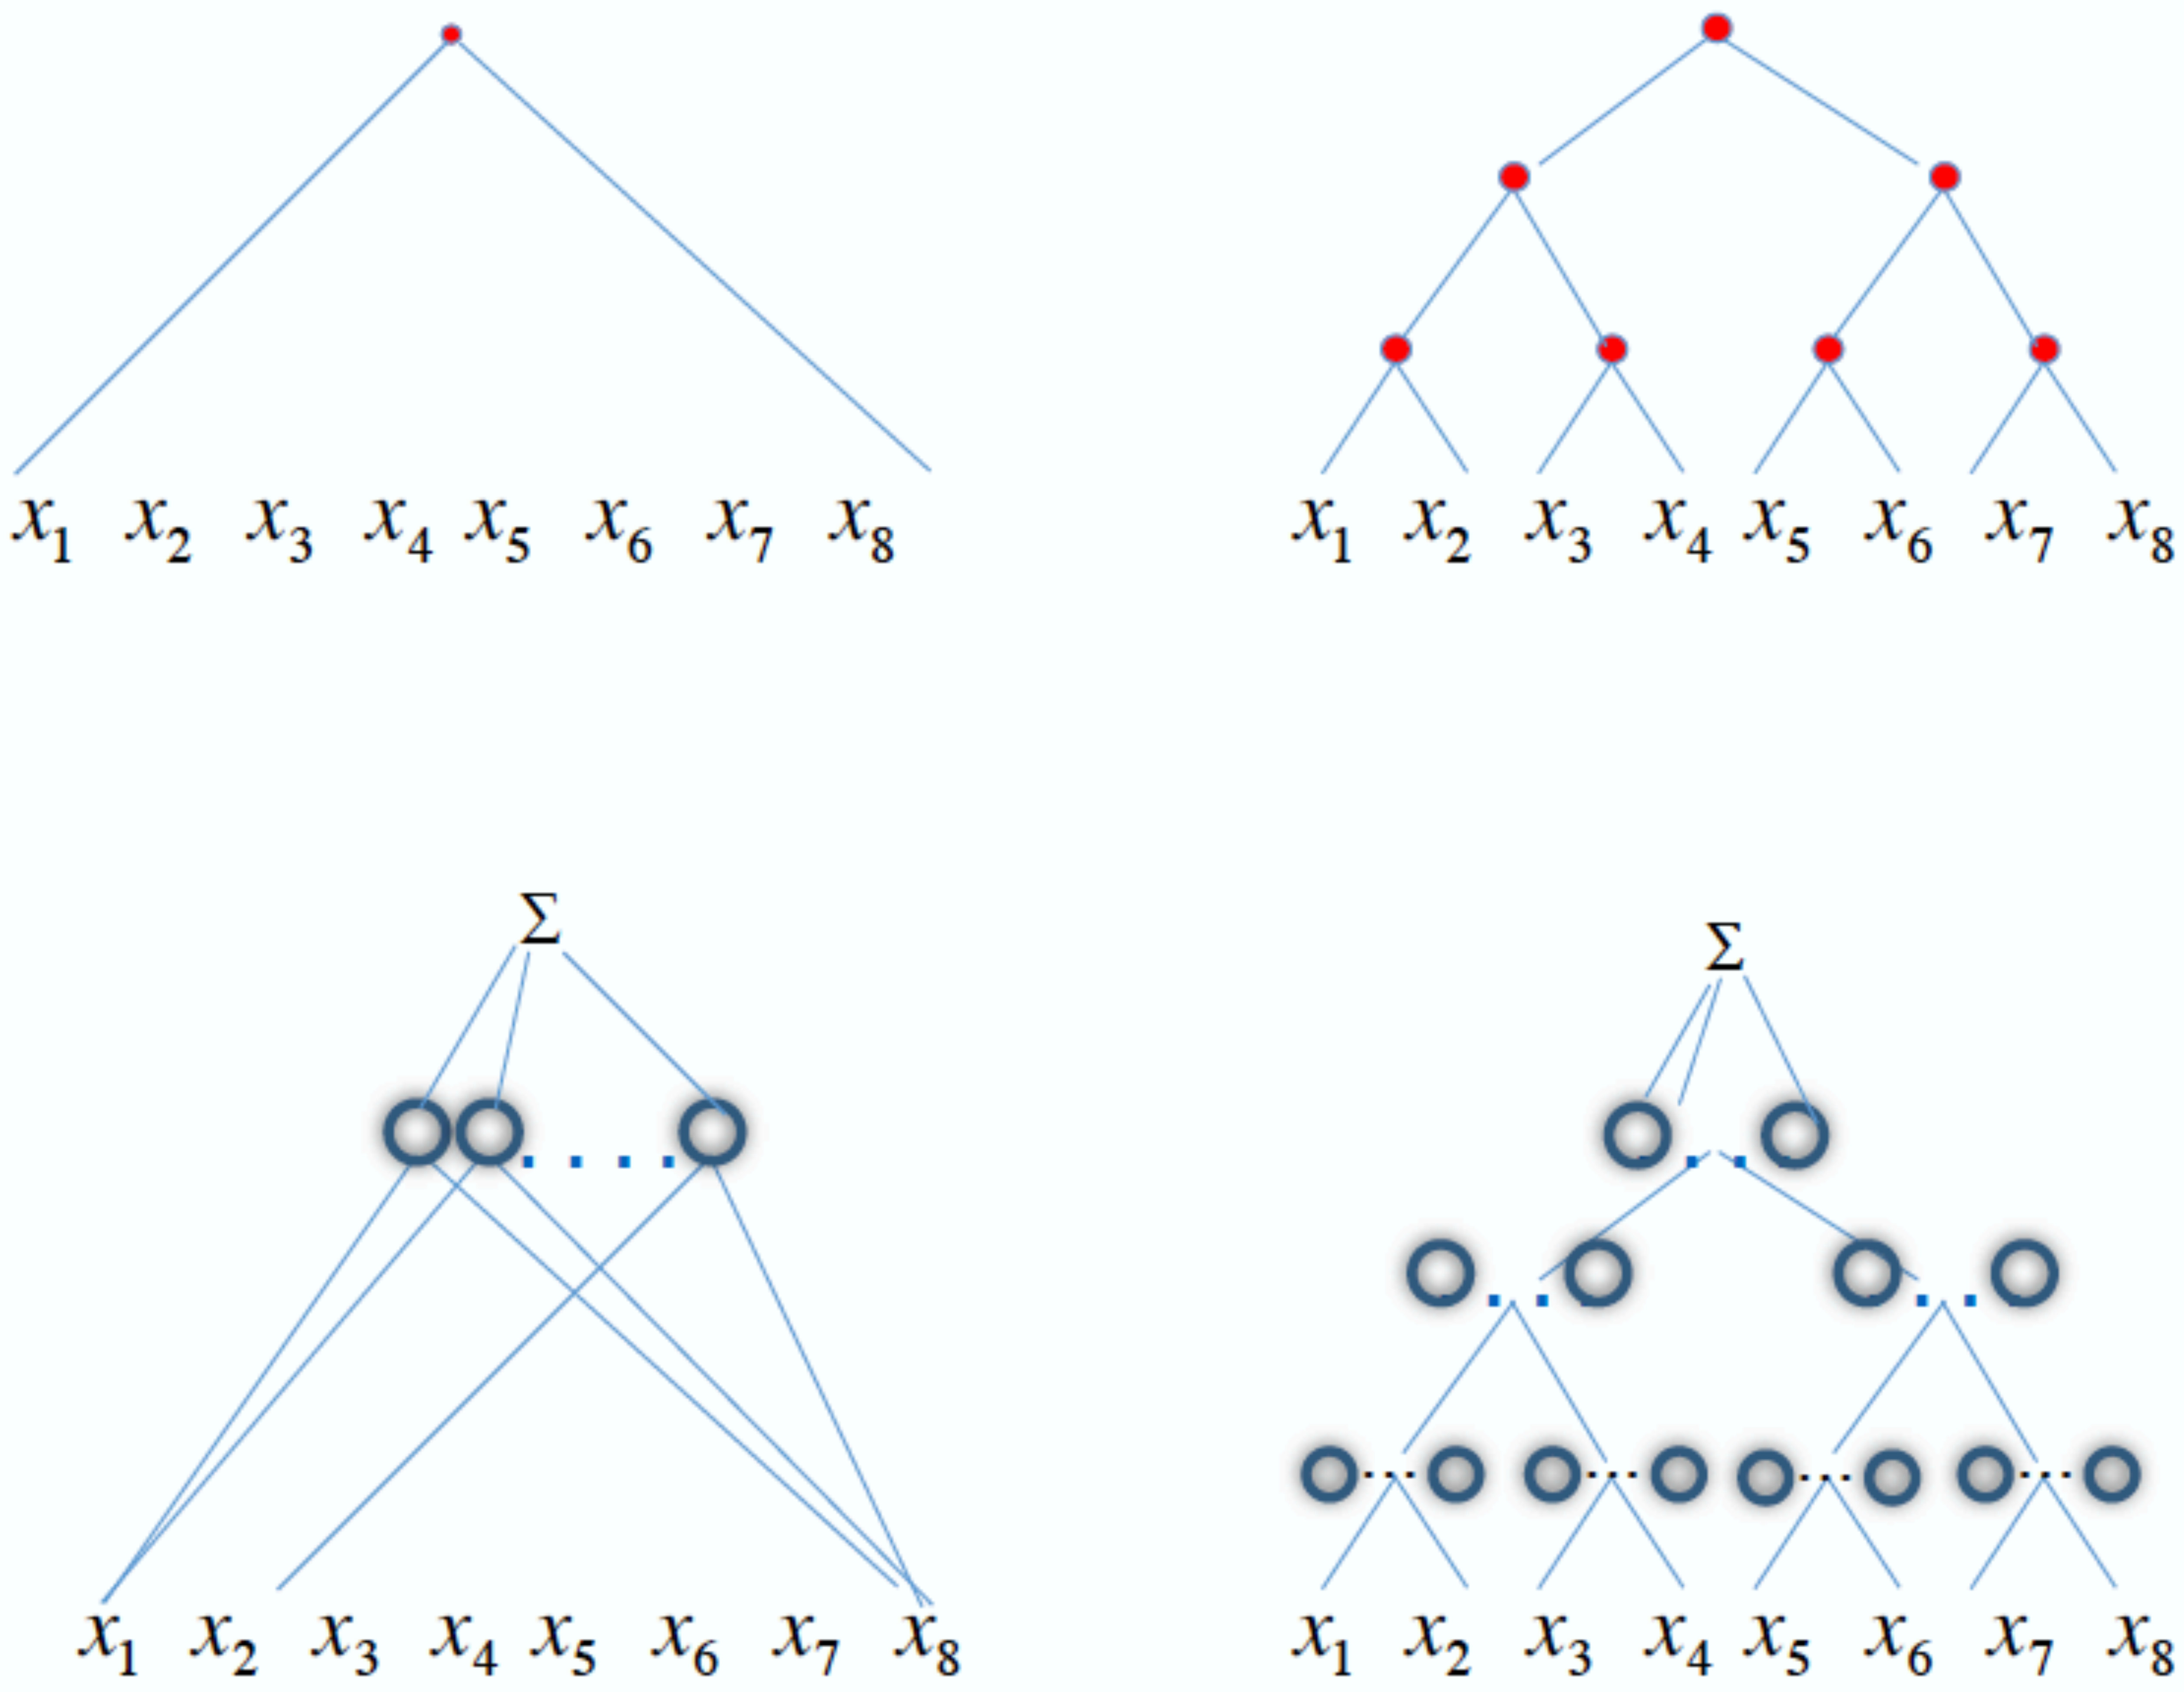

a *b*

Hierarchically local compositionality

$$f(x_1,x_2,\ldots,x_8)=g_3(g_{21}(g_{11}(x_1,x_2),g_{12}(x_3,x_4))g_{22}(g_{11}(x_5,x_6),g_{12}(x_7,x_8)))$$

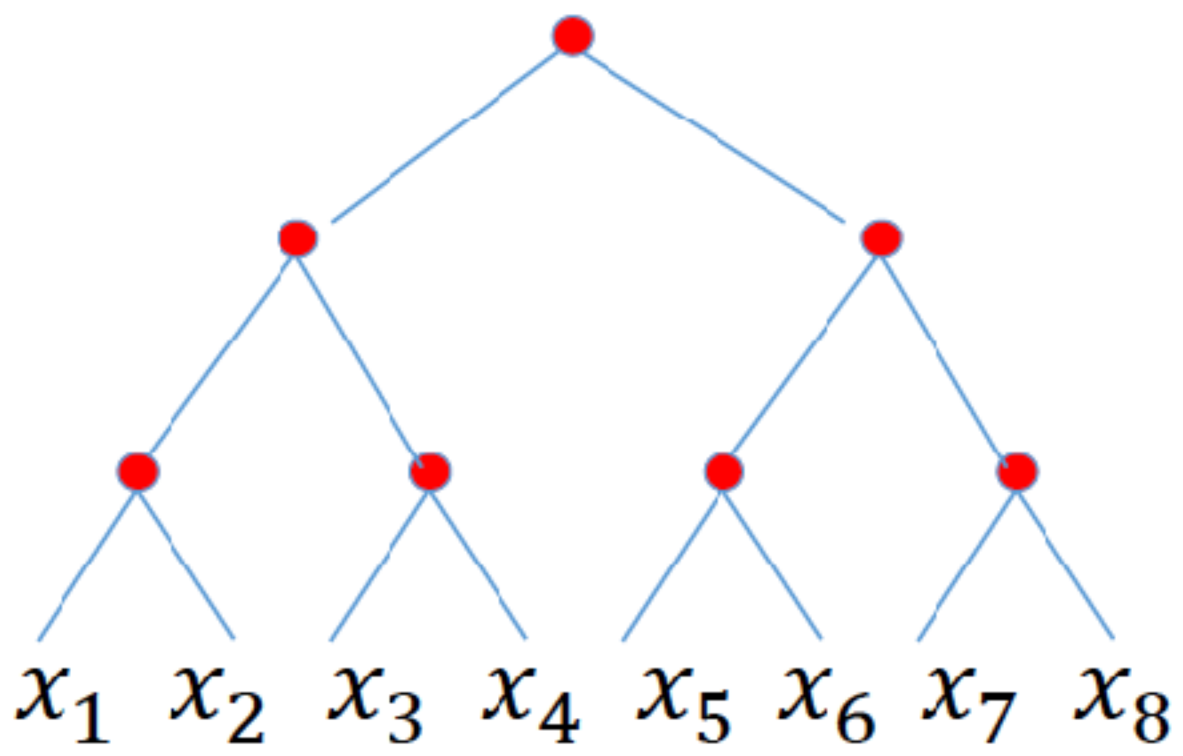

Theorem (informal statement)

Suppose that a function of d variables is hierarchically, locally, compositional . Both shallow and deep network can approximate f equally well. The number of parameters of the shallow network depends exponentially on d as $O(\varepsilon^{-d})$ with the dimension whereas for the deep network dance is $O(d\varepsilon^{-2})$

Mhaskar, Poggio, Liao, 2016

Locality of constituent functions is key **not** weight sharing: CIFAR

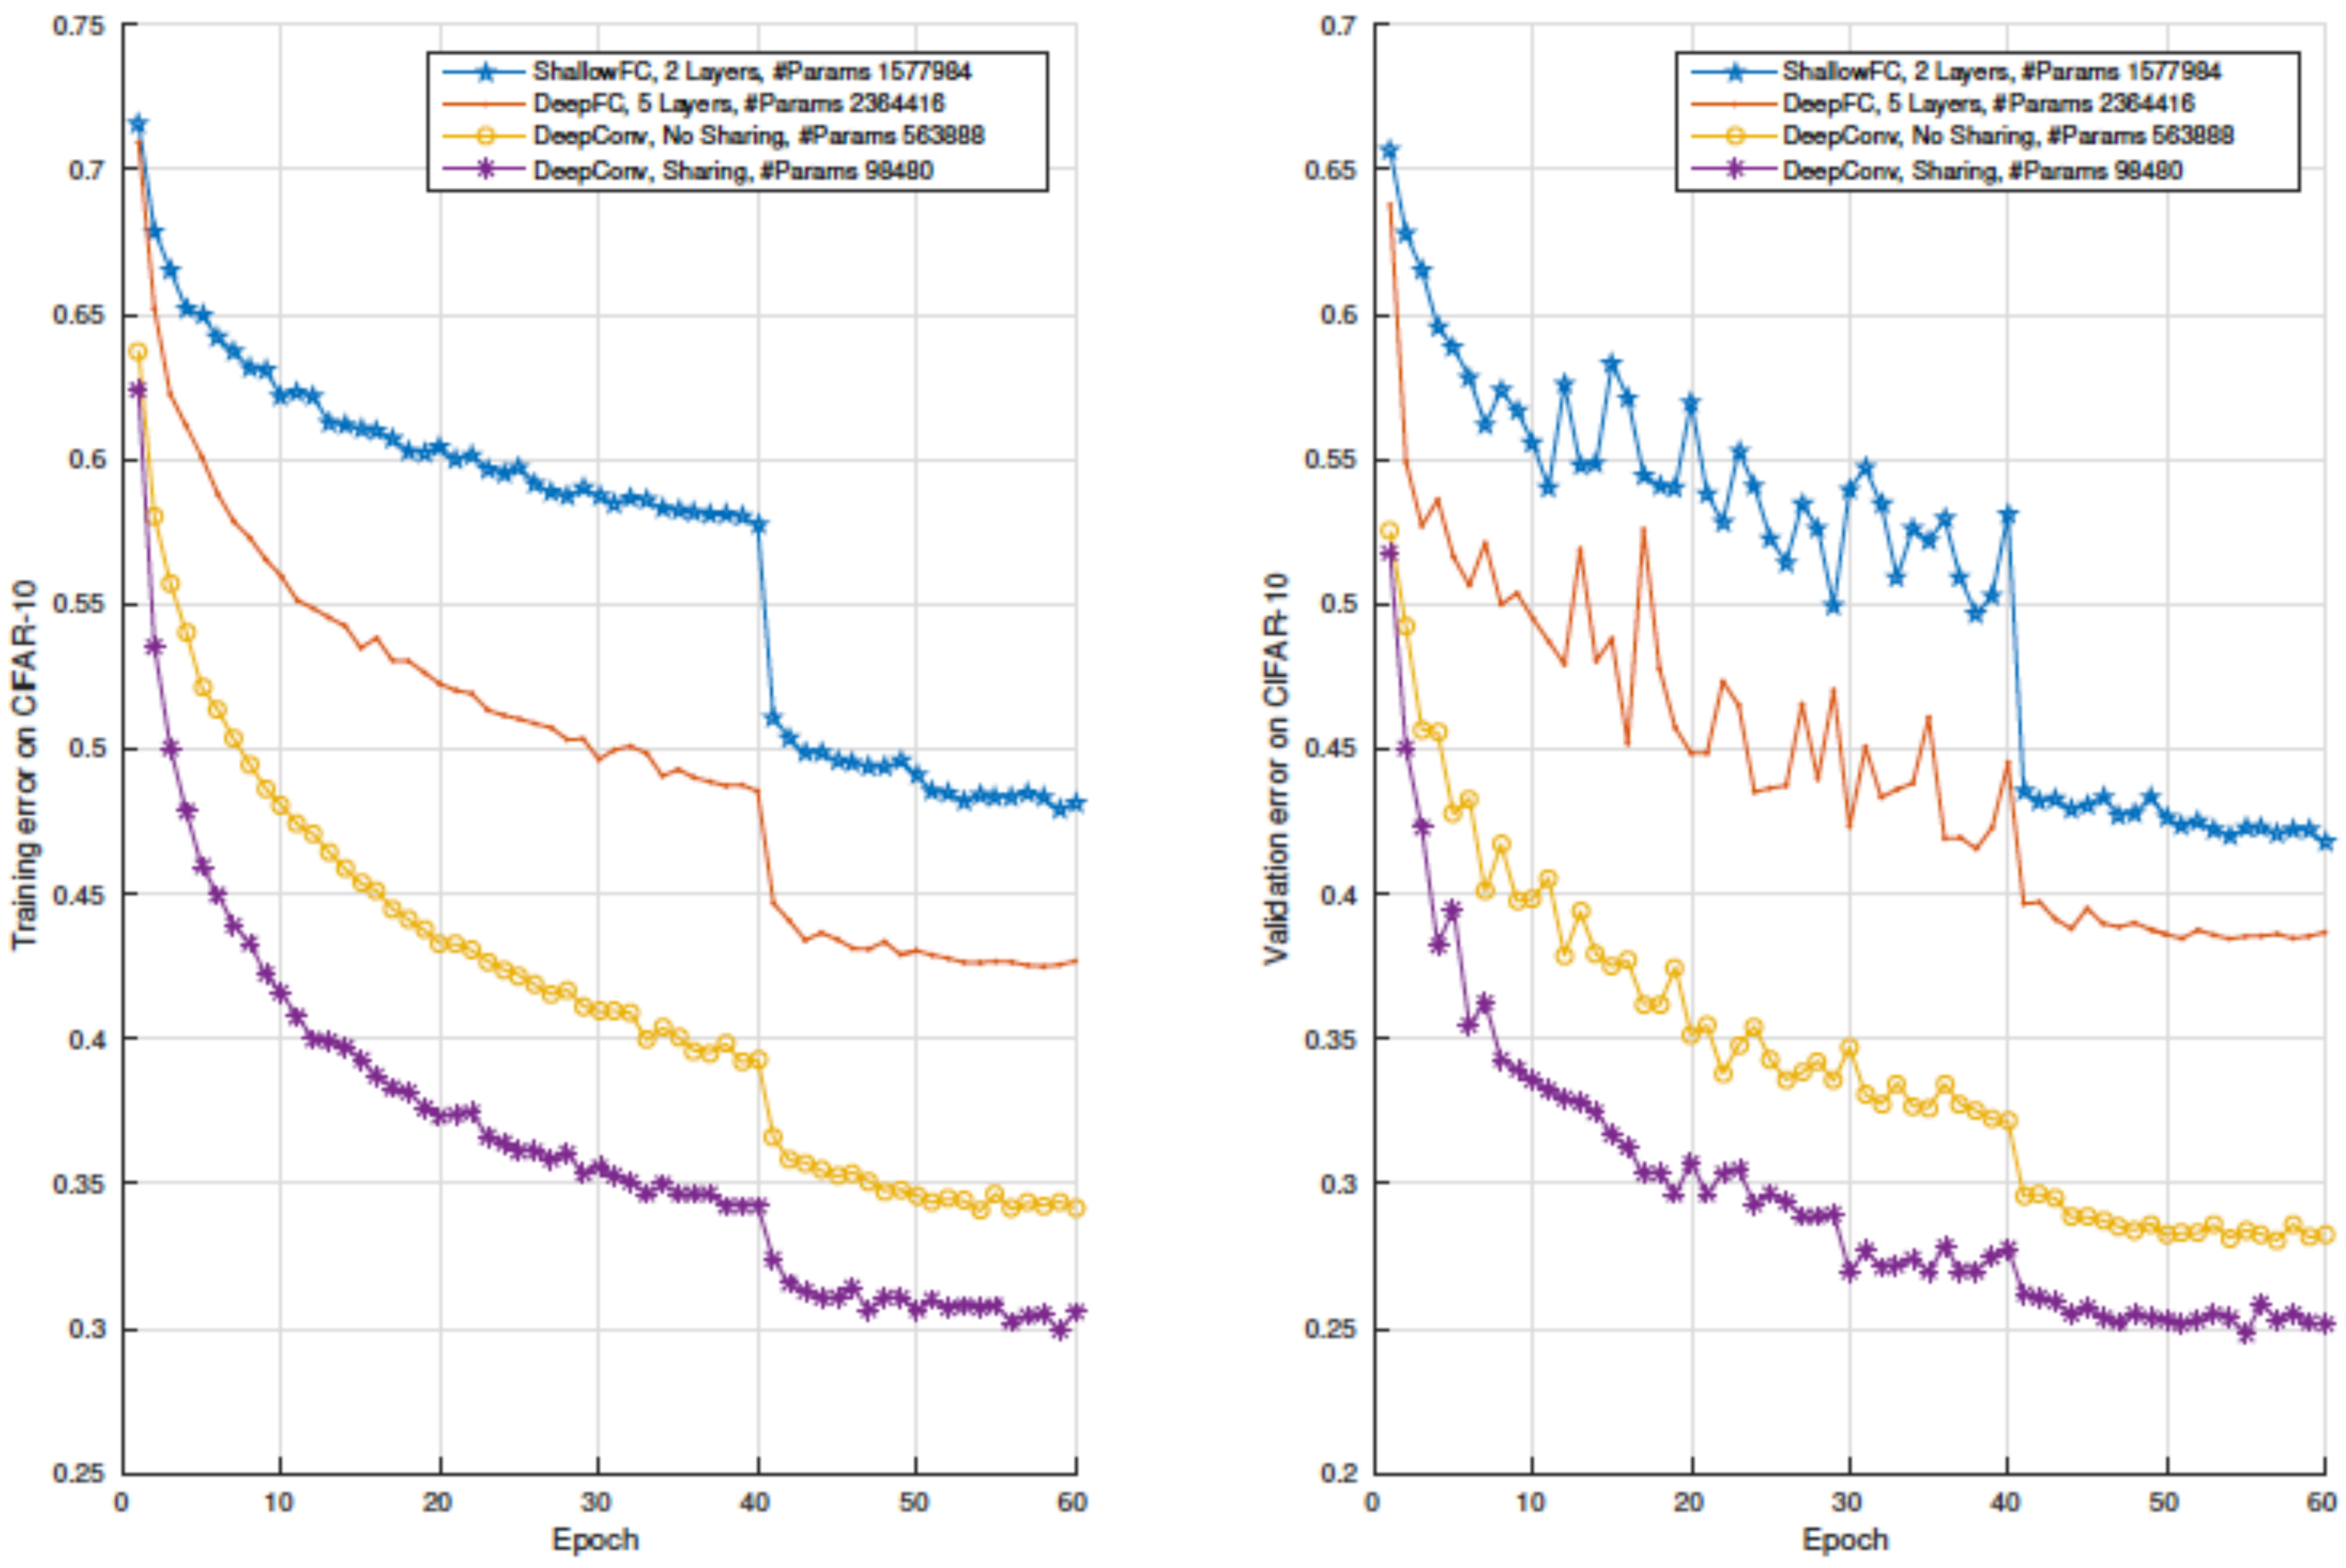

Open problem: why compositional functions are important for perception?

Which one of these reasons:
Physics?
Neuroscience? <===
Evolution?

Opportunity for theory projects!

Theory II:
What is the Landscape of the empirical risk?

Theorem (informal statement)

Replacing the RELUs with univariate polynomial approximation, Bezout theorem implies that the system of polynomial equations corresponding to zero empirical error has a very large number of degenerate solutions. The global zero-minimizers correspond to flat minima in many dimensions (generically unlike local minima). Thus SGD is biased towards finding global minima of the empirical risk.

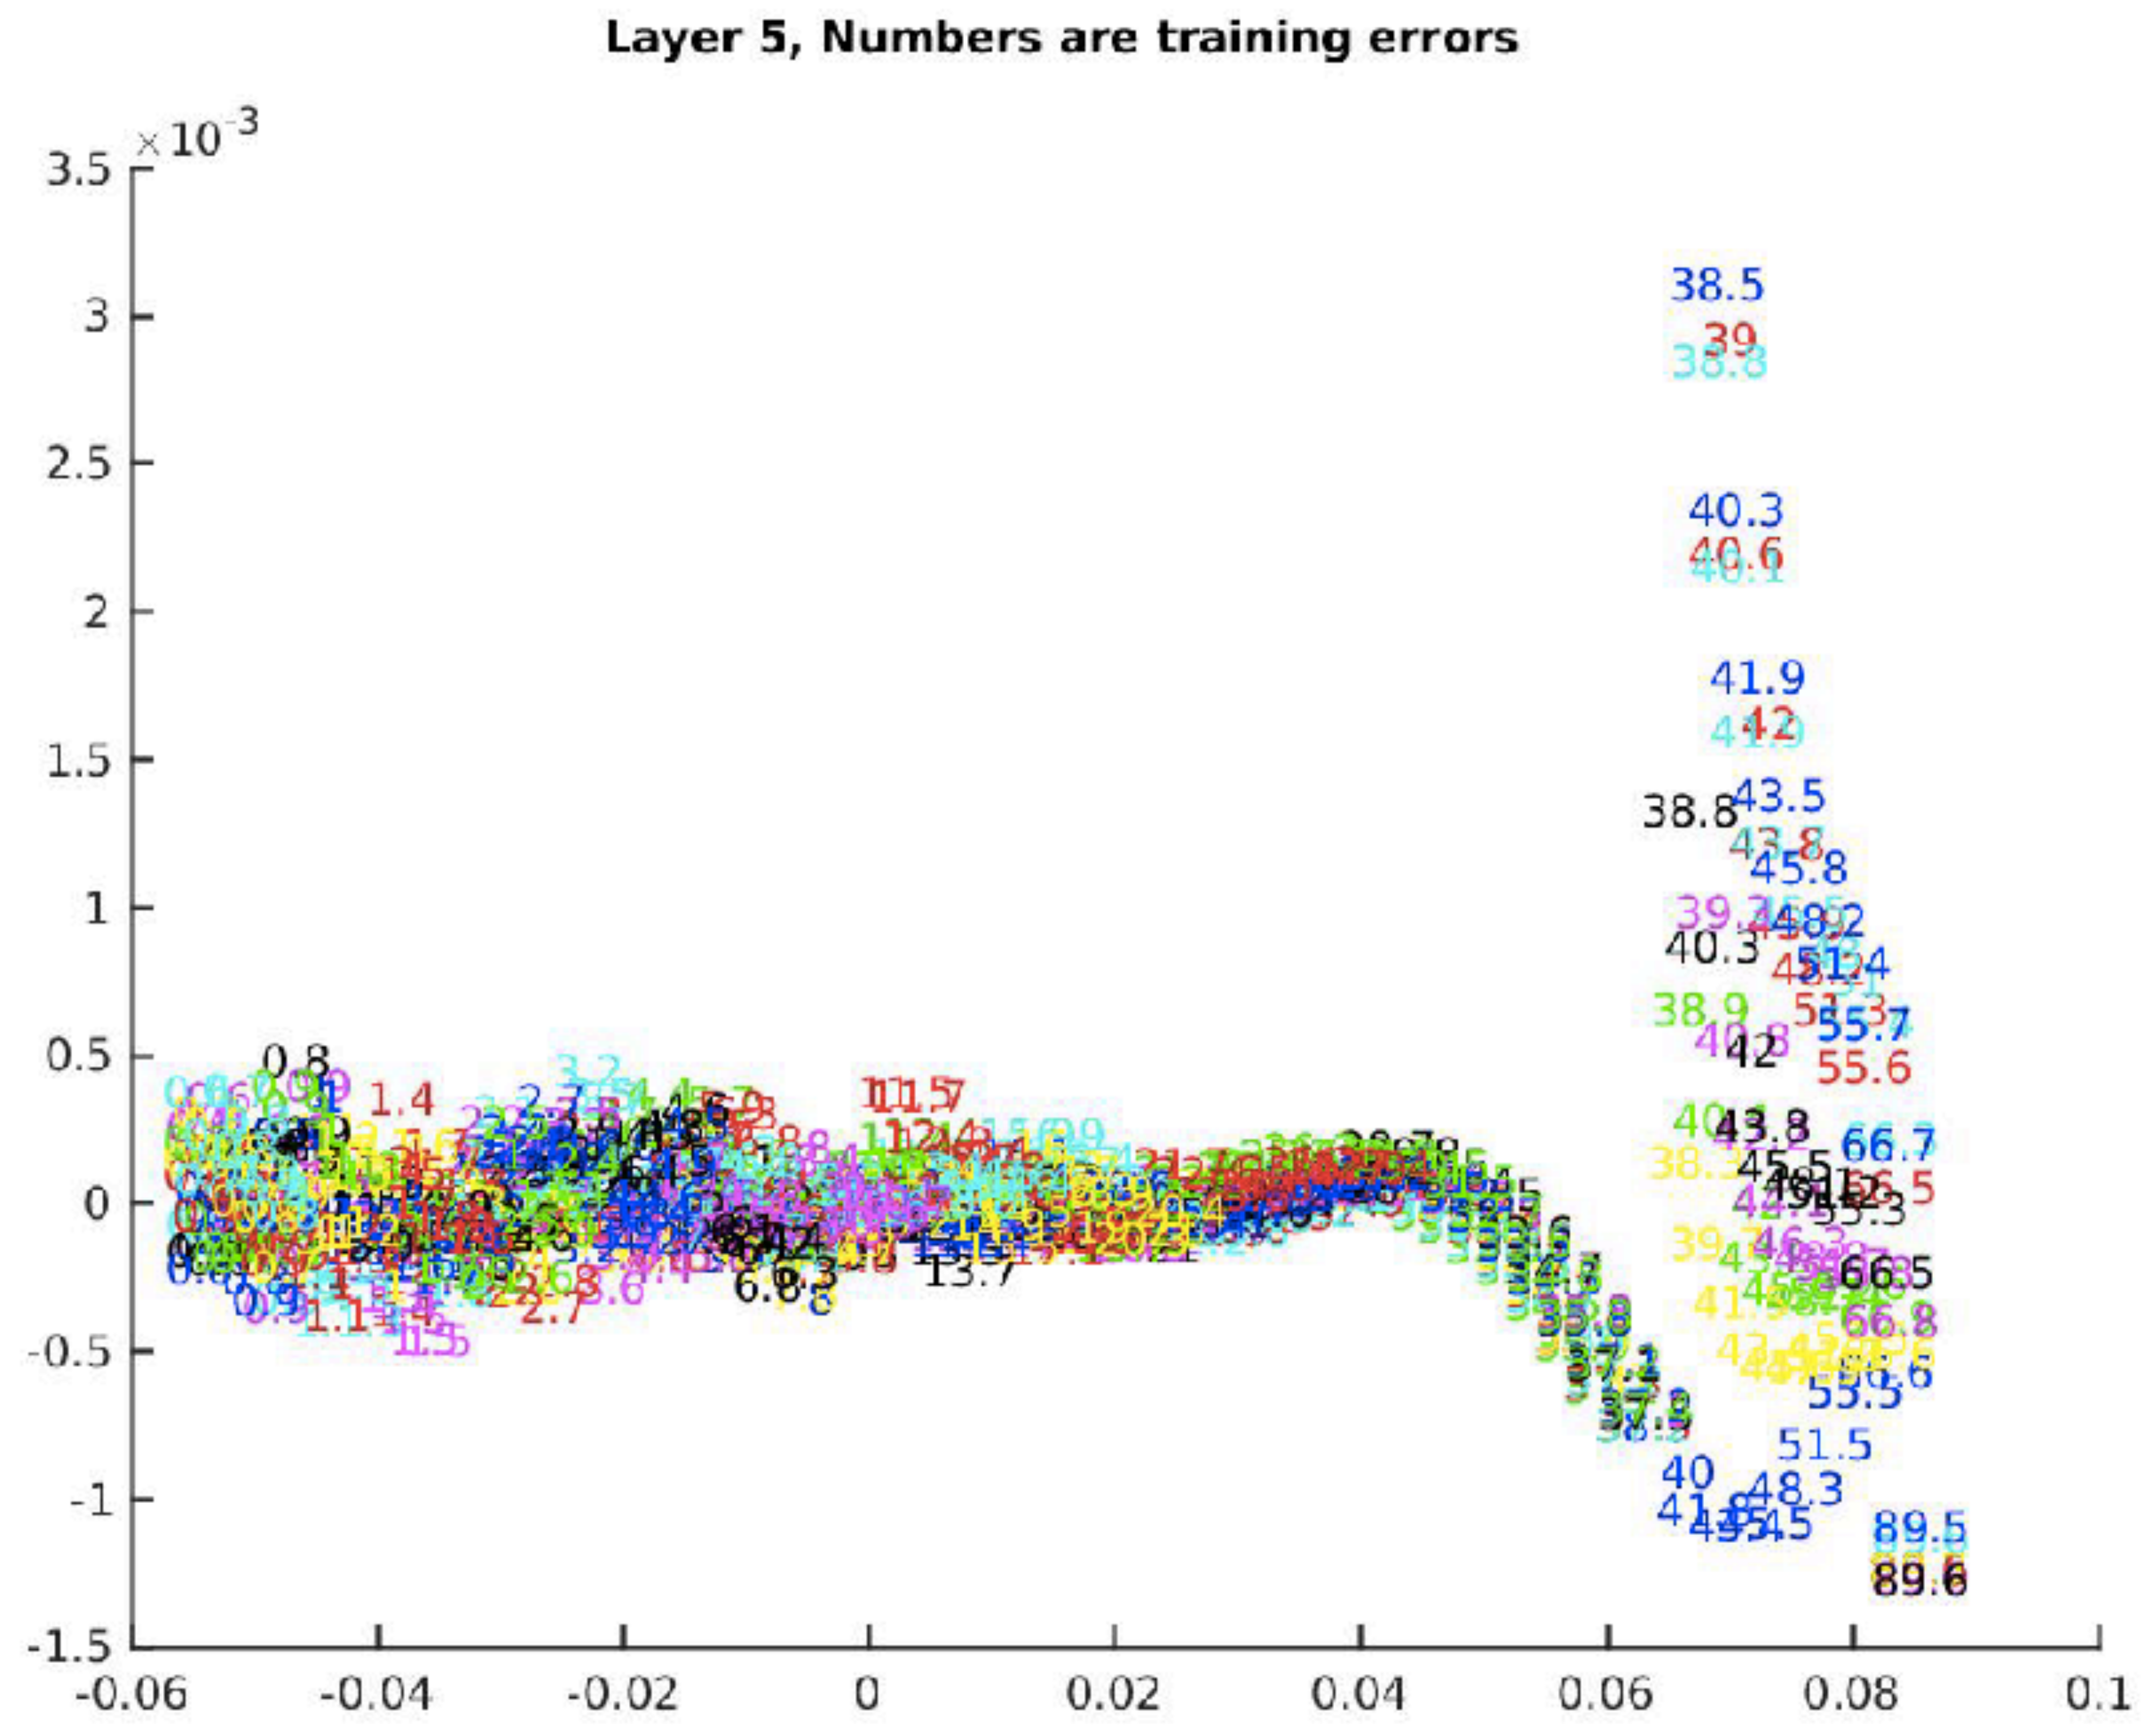

CENTER FOR Brains Minds+ Machines

Liao, Poggio, 2017

Theory III: How can underconstrained solutions generalize?

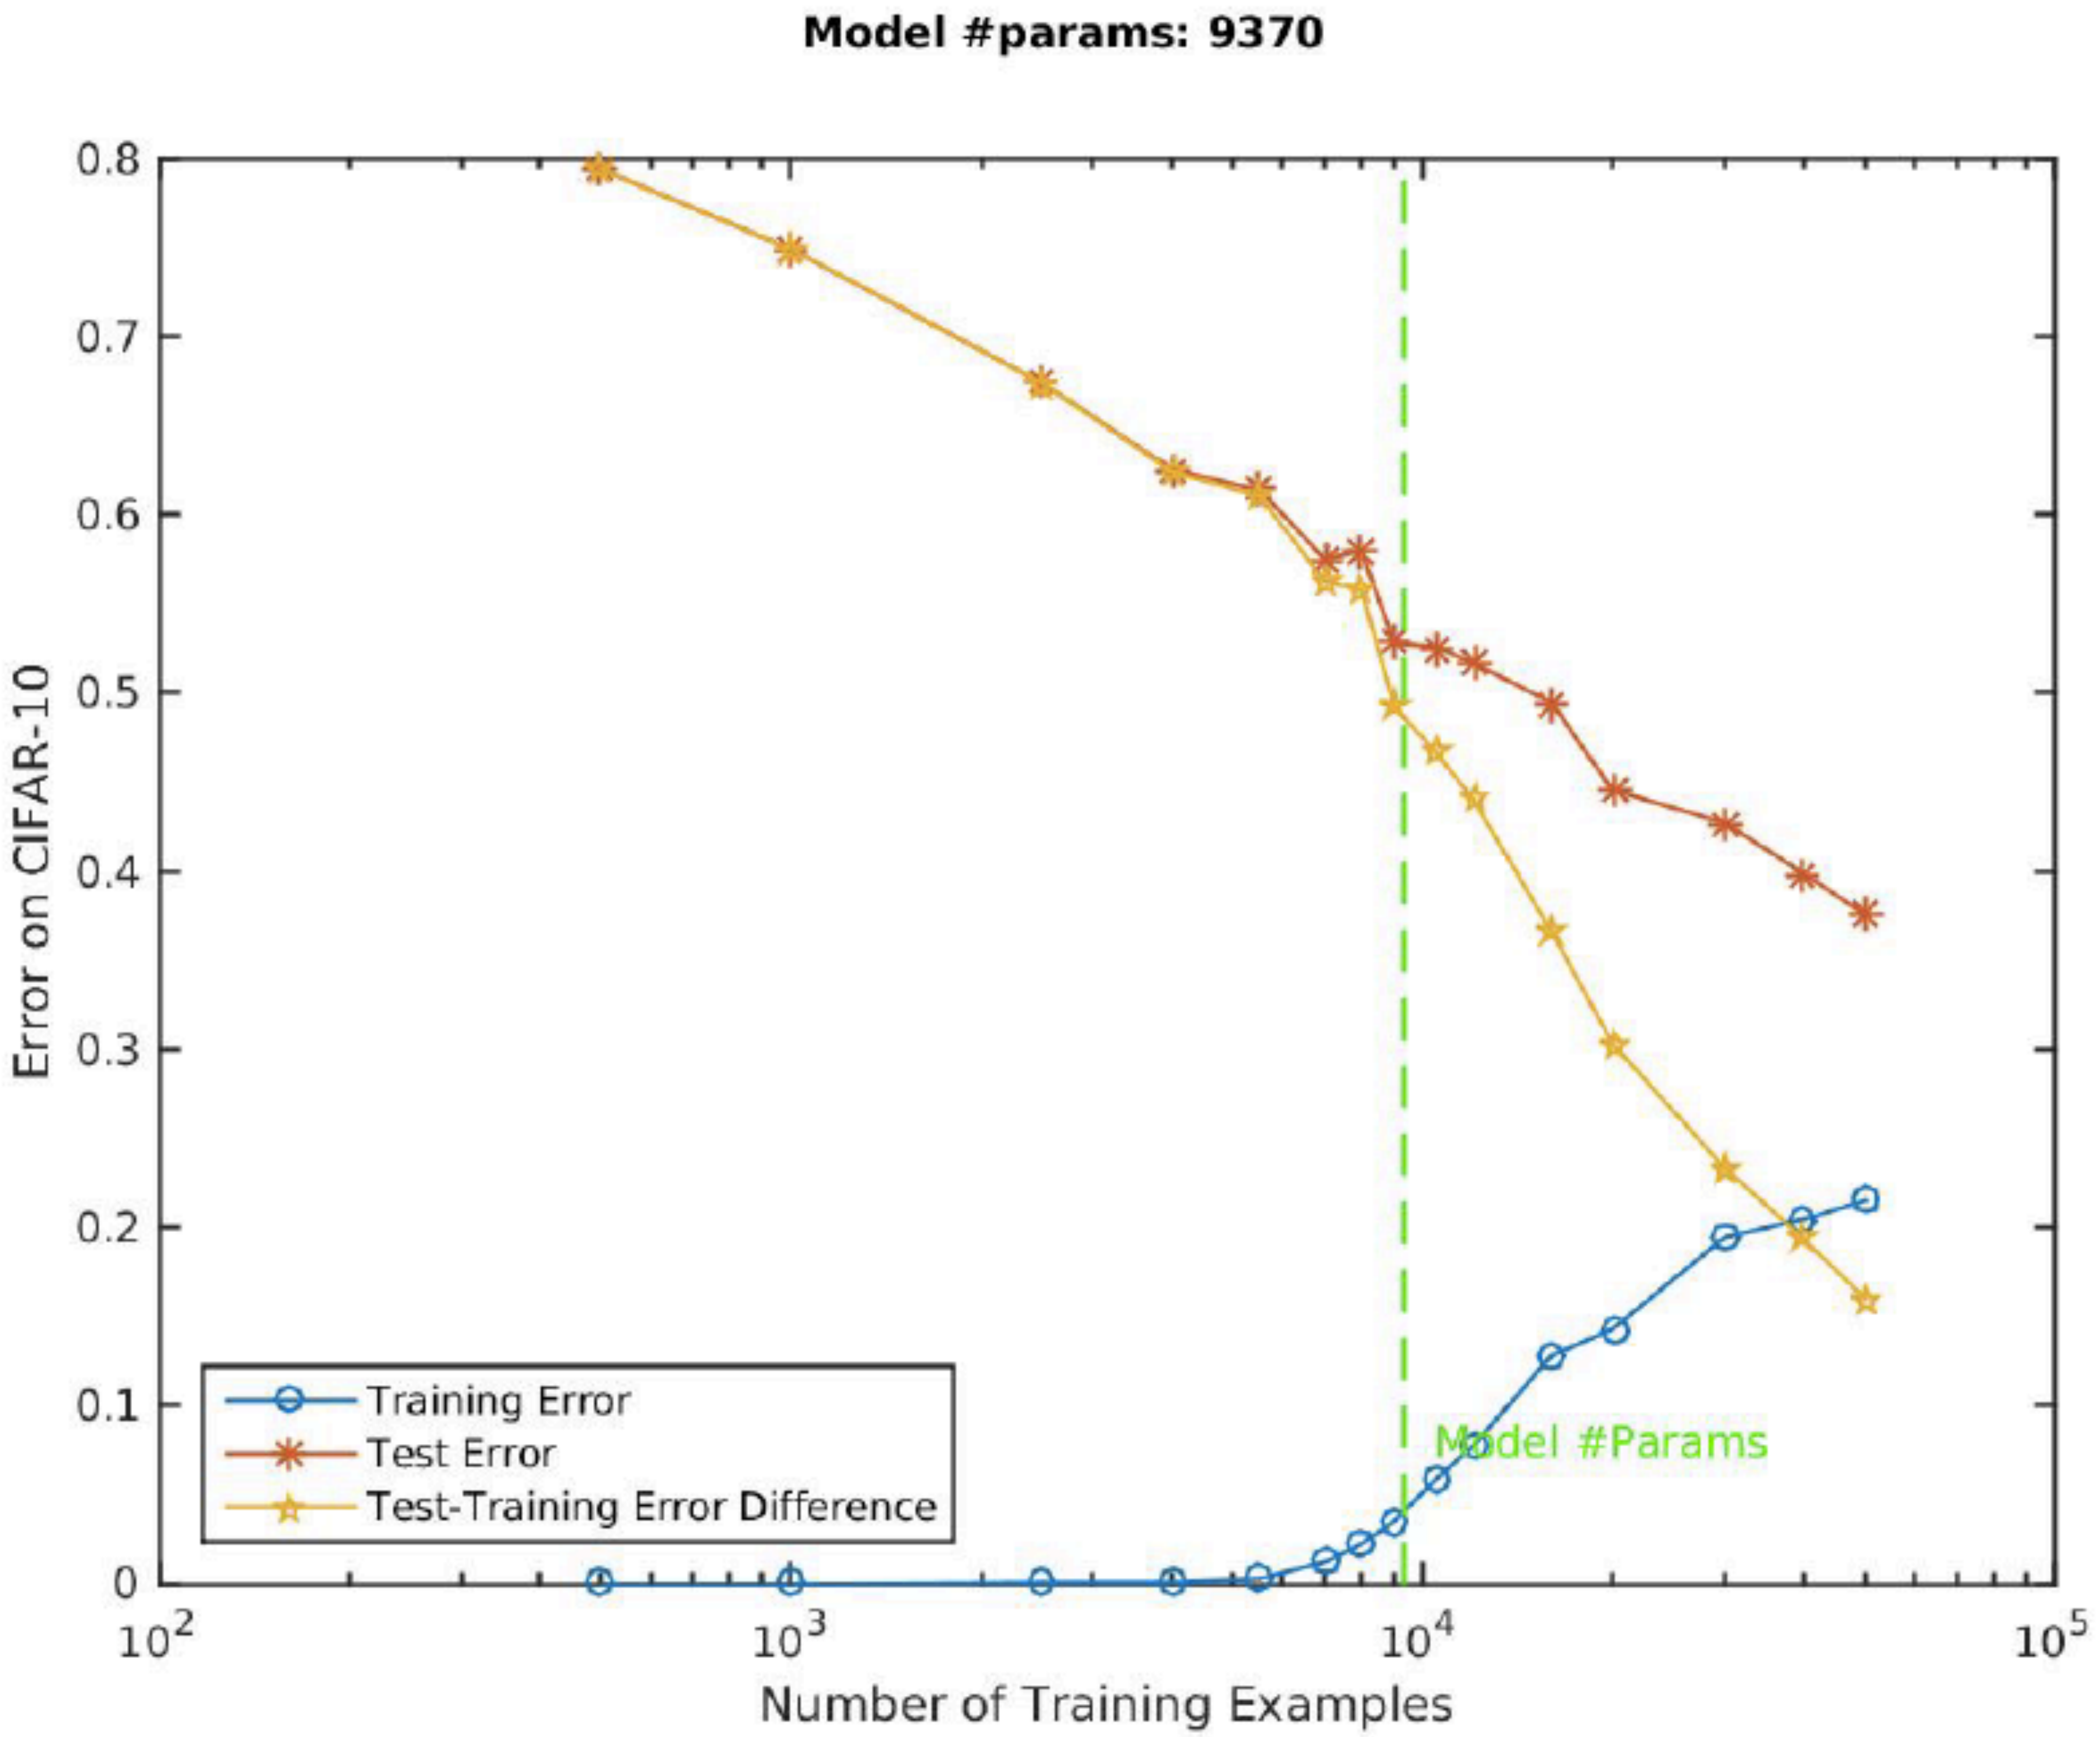

CENTER FOR Brains Minds+ Machines

Good generalization with less data than # weights

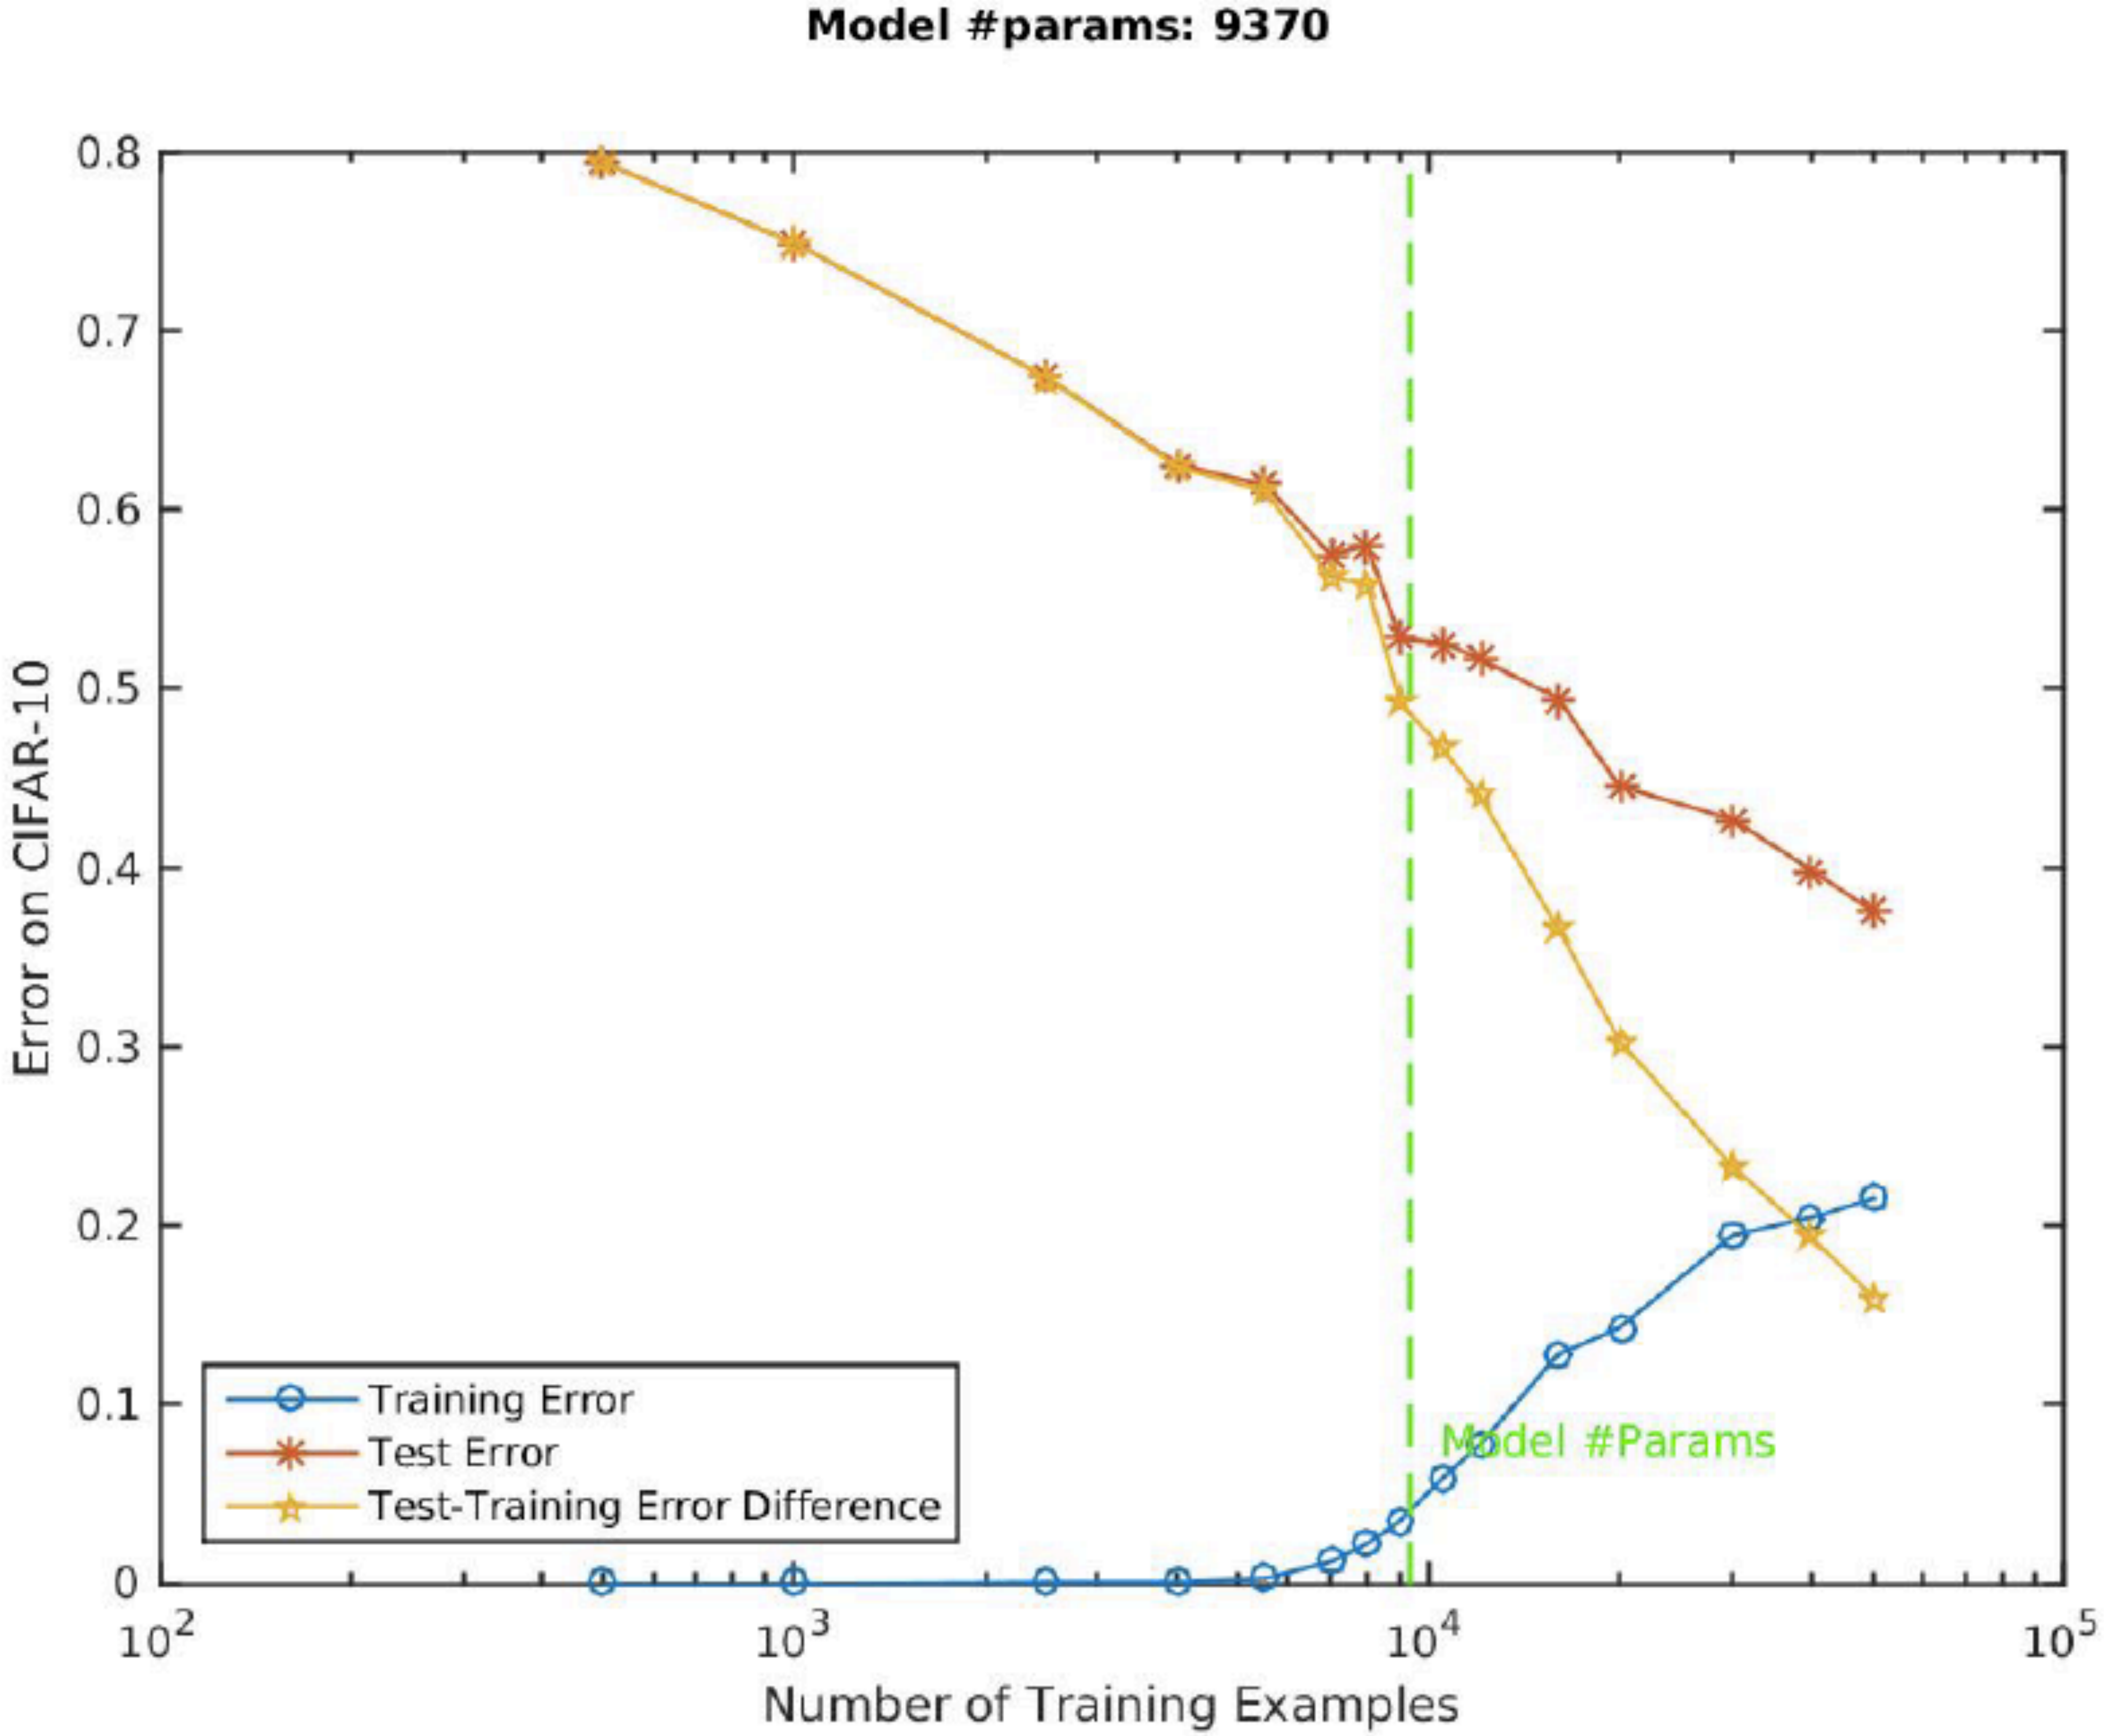

CENTER FOR
Brains
Minds+
Machines

Randomly labeled data

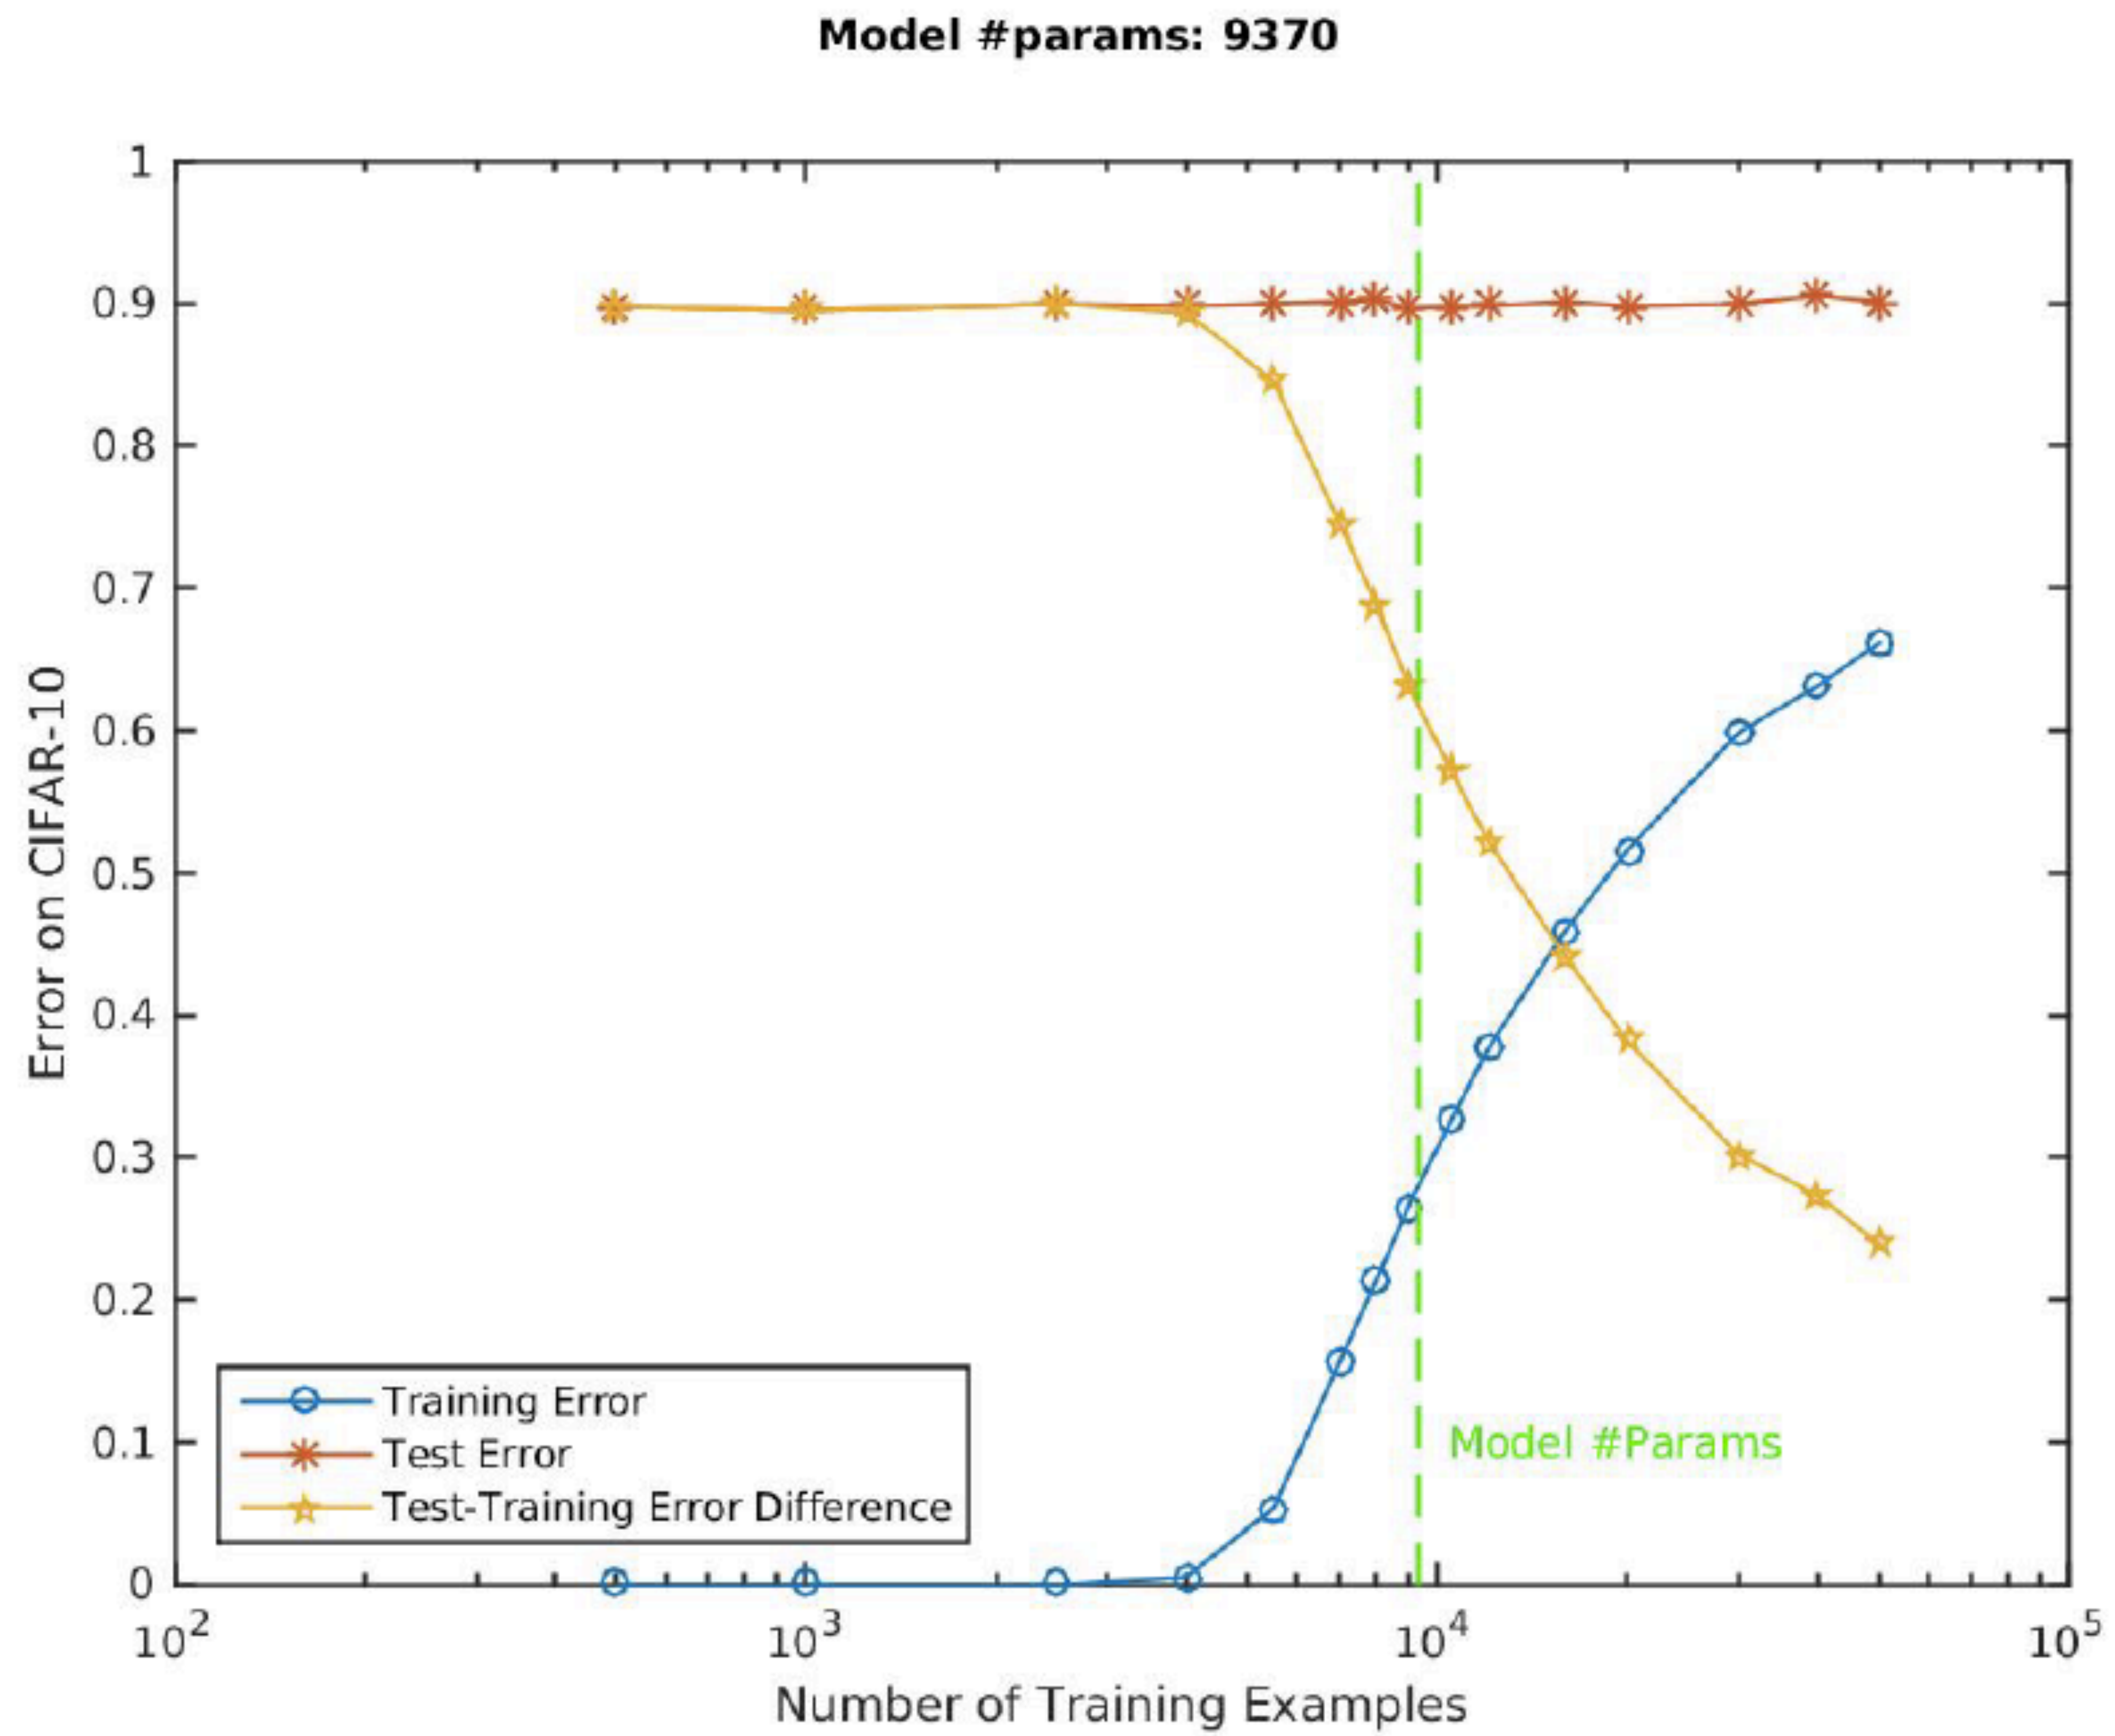

CENTER FOR
Brains
Minds+
Machine

Beyond today's DLNNs: several scientific questions…

Why do Deep Learning Networks work? ===>

In which cases will they fail?

Is it possible to improve them?

Is it possible to reduce the number of labeled examples?

CENTER FOR Brains Minds+ Machines

Module One: Ventral Visual Stream

What we aim to accomplish (years 5-10):

- A mathematical theory when & how deep networks work: will provide guidance on when networks will succeed (or fail) when applied to other sensory data streams.
- An understanding of the role of retinal sampling, and feedback and bypass connections in the success of ventral visual processing
- Novel ventral visual stream models: improved visual task performance & improved match/explanation of neurophysiology and behavior.
- First systematic data on the development of the primate ventral stream and its dependence on its visual experience "diet".


MODULE FOUR
Towards Symbols
LONG-TERM PLANNING ABSTRACTION LANGUAGE
B. Katz, A. Barbu, S. Ullman, J. Tenenbaum
G, Kreiman, M. Wilson, B. Desimone
J. Tenenbaum, N. Kanwisher, Spelke
MODULE TWO
Brain OS
WORKING MEMORY VISUAL ROUTINES ATTENTION
3 1/2 D
Sketch
MODULE THREE
Cognitive Core
INTUITIVE PHYSICS GEOMETRY (3D) INTUITIVE PSYCHOLOGY
Depository of vision routines… synthesizing routines as needed
Running routines…
MODULE ONE
Visual Stream
FOVEA DEEP FEED-FORWARD NETWORKS BACK PROJECTIONS

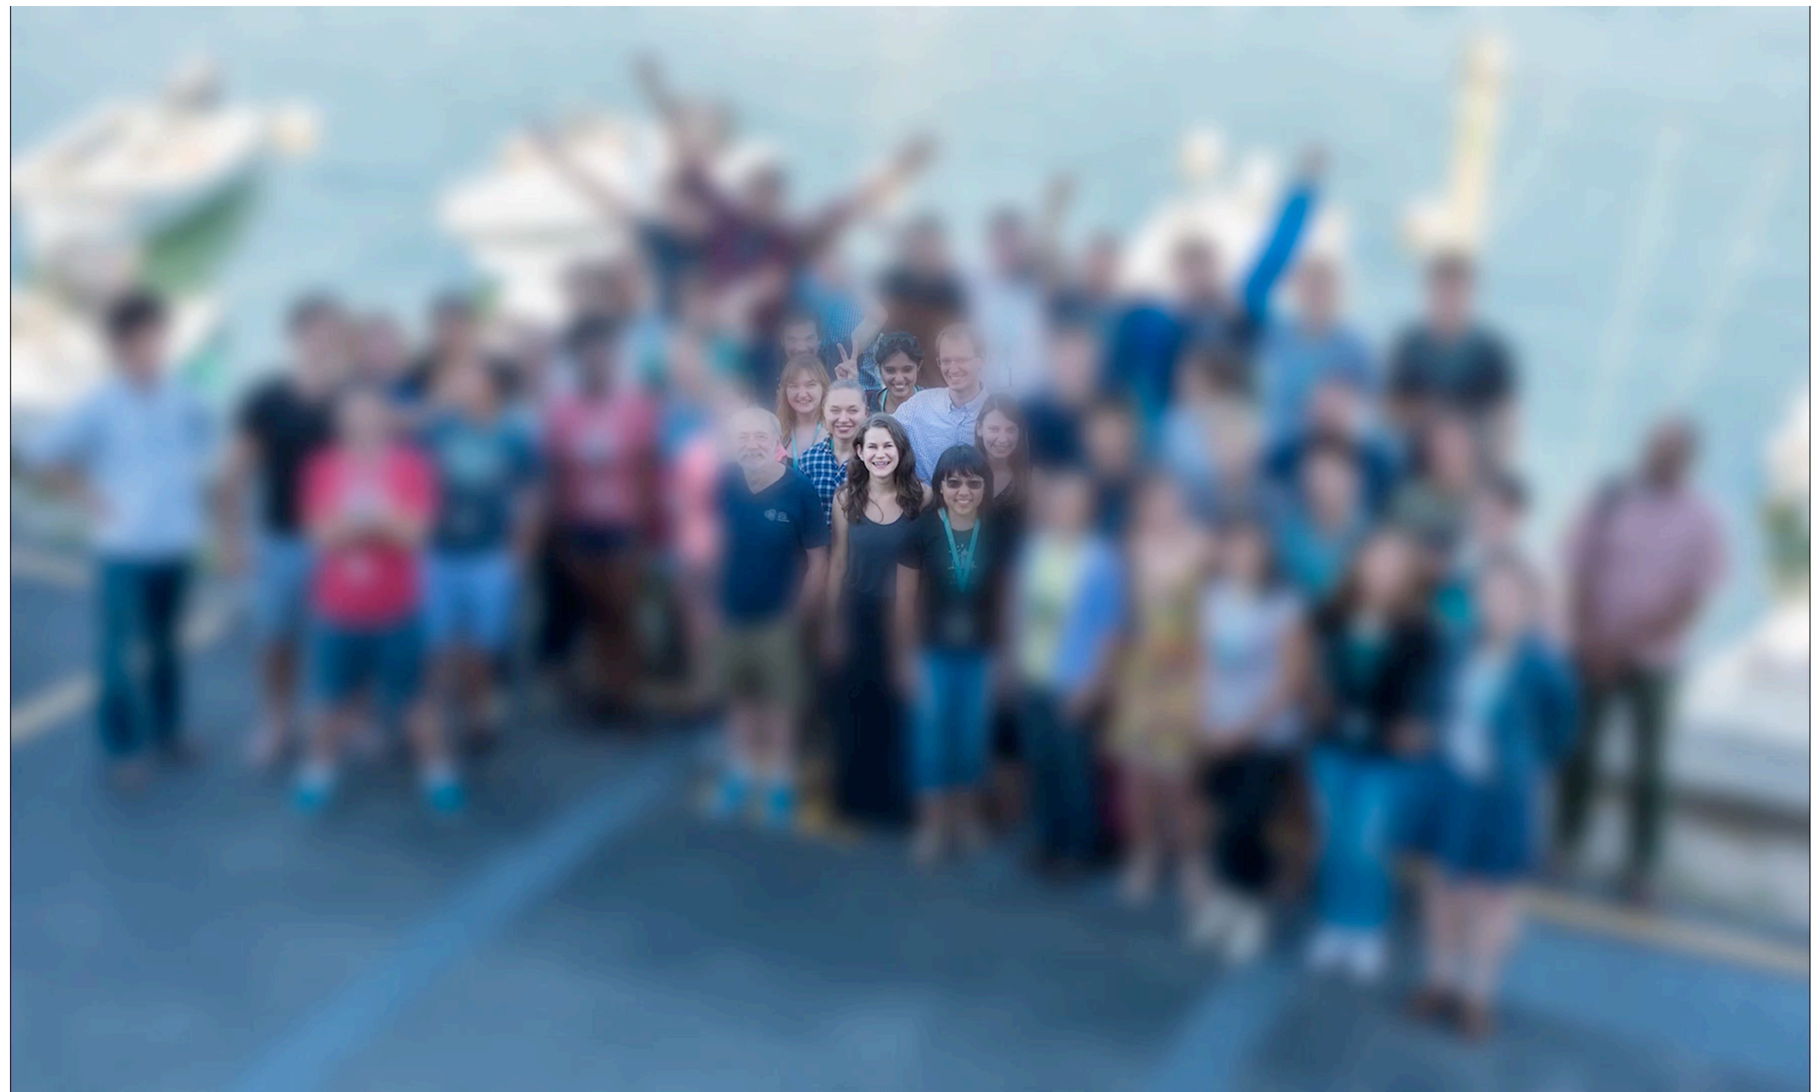

9.523 overview

- Module 1
  - Module 1, theory
  - Module 1, eccentricity
  - Module 1, invariance
- Projects

General projects/proposals

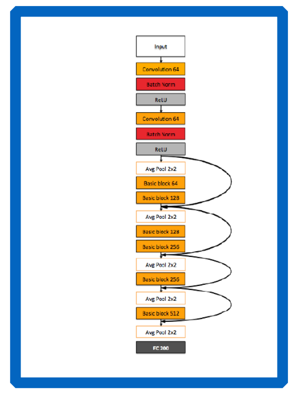

- new architectures/class of applications from basic DCN block (example GAN + RL/DL + …)

- new semisupervised training framework, avoiding labels: implicit labeling…predicting next “frame”:
  - coloring…

- new learning algorithms instead of SGD, that are biologically plausible:
  - trail and error, genetic algorithms…
  - stream and counter stream
  - …

- 3 1/2 D sketch specs

- Visual awareness experiment?