CS-E577005 : Computational Theories of the Brain

Prof. Stéphane Deny
Department of Neuroscience and Biomedial Engineering
Department of Computer Science
Aalto University

Lecture 1 : Crash course about the brain

Presentations

- What is your name?
- What are you currently studying (what program are you in)?
- What would you like to learn about the brain?

Your Brain vs a MacBook Pro (chip M1)

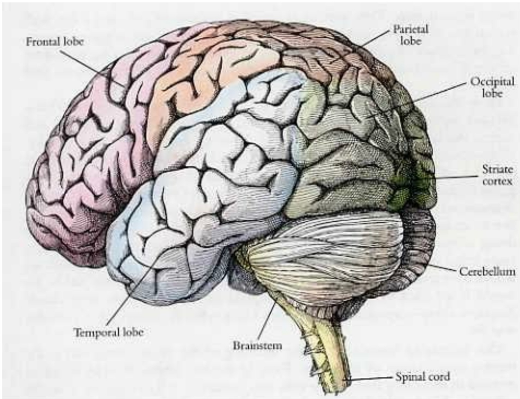

- Number of neurons:
- Number of operations per second:
- Average consumption:

- Number of tansistors:
- Number of operations per second:
- Average consumption:

Your Brain vs a MacBook Pro (chip M1)

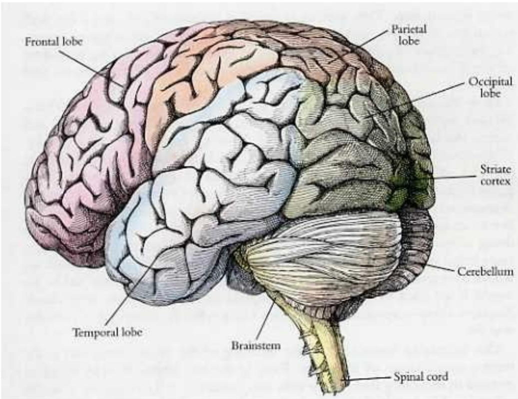

- Number of neurons: 86B
- Number of operations per second: ?
- Average consumption: 20W

- Number of tansistors: 34B
- Number of operations per second: 3.2GHz
- Average consumption: 30-50W

At the origins of neuroscience: Anatomical description of the brain

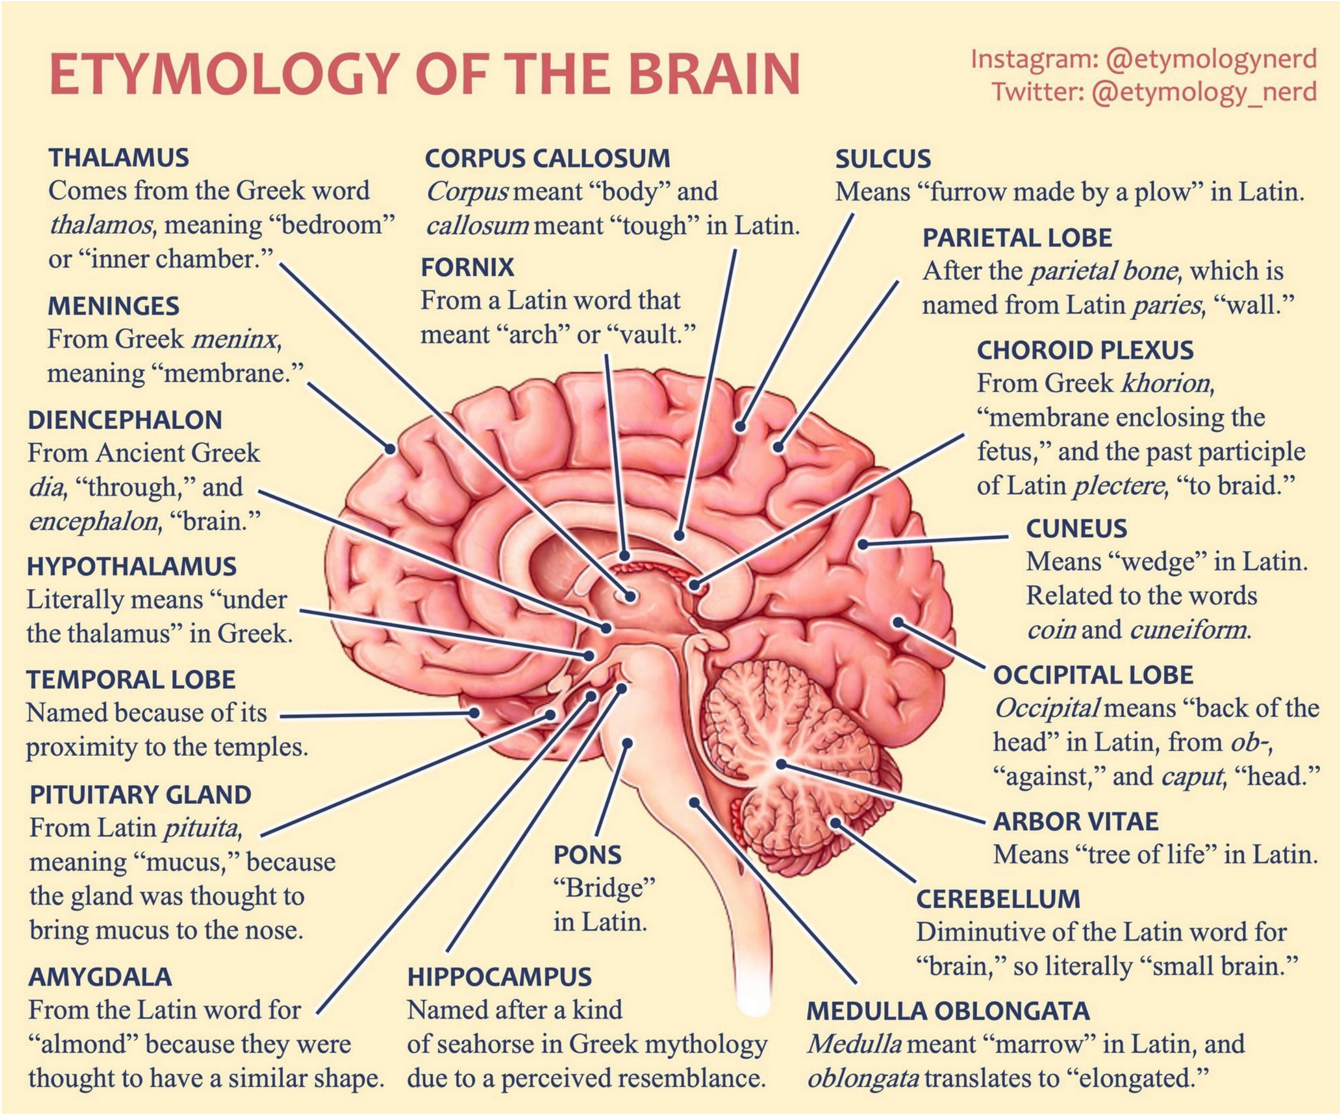

Species and brain size

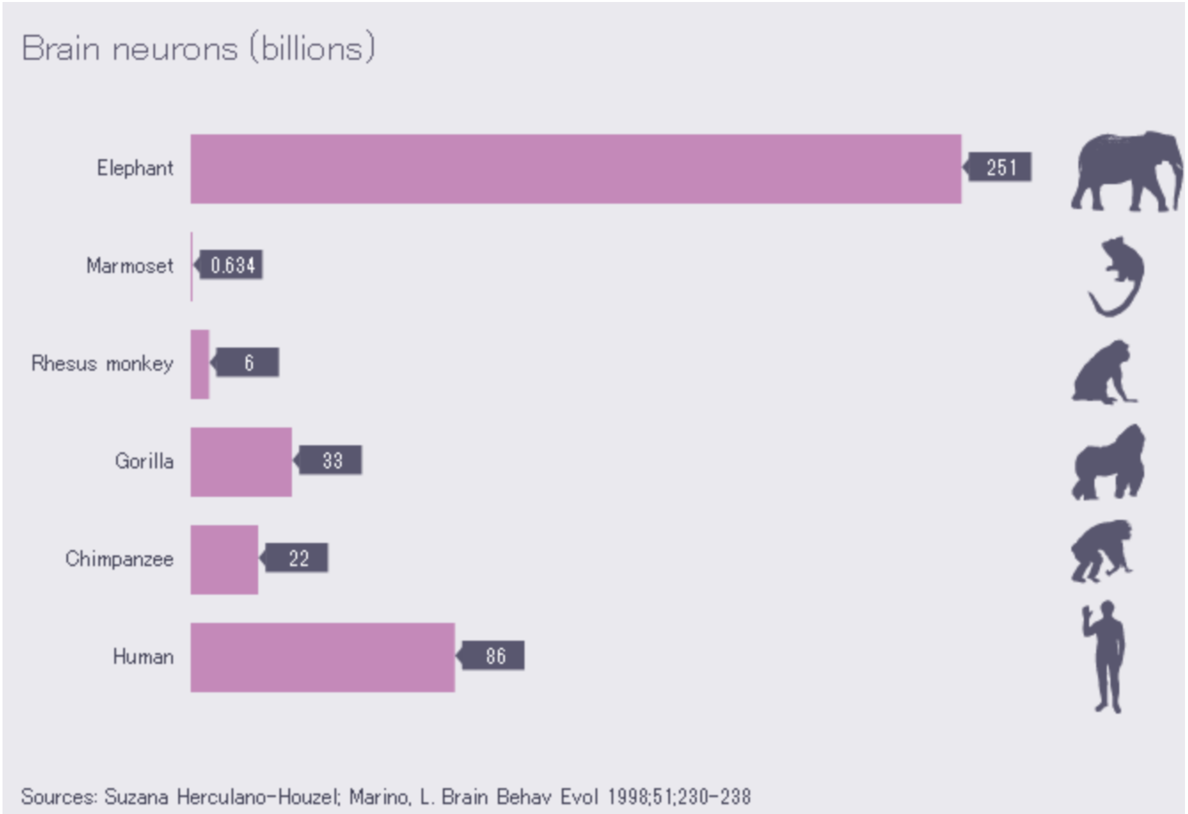

source: Wikipedia

Species and cortex size

Brain

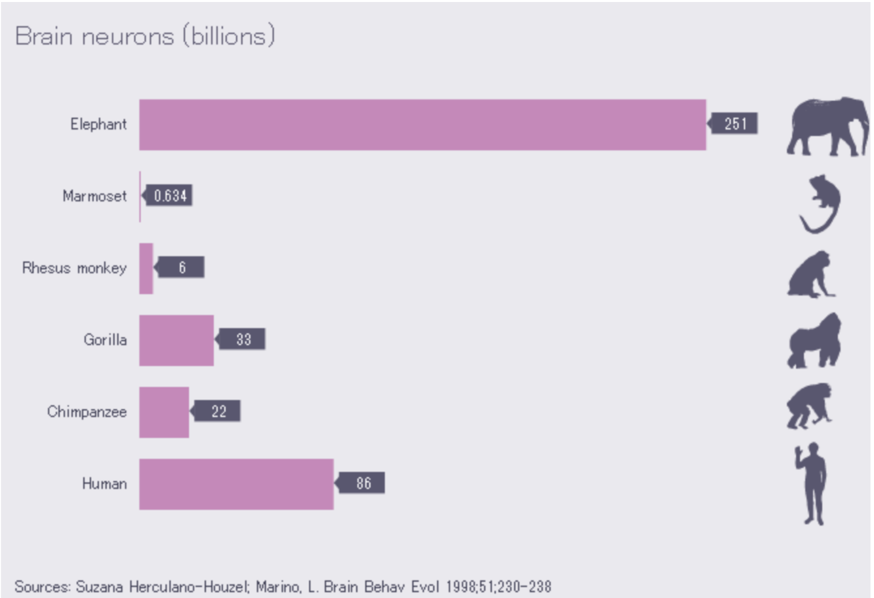

Cortex

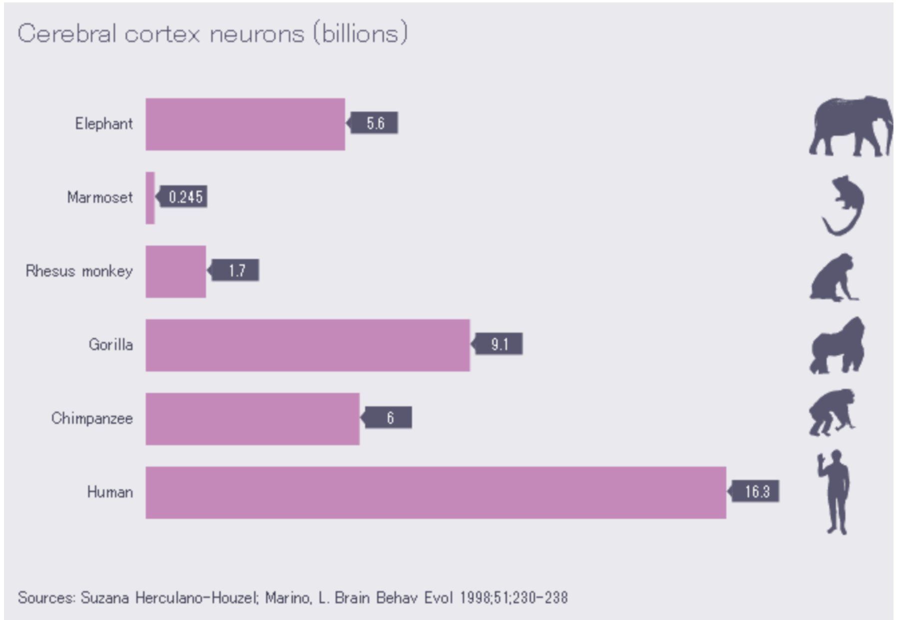

What is a cortex?

Cortex

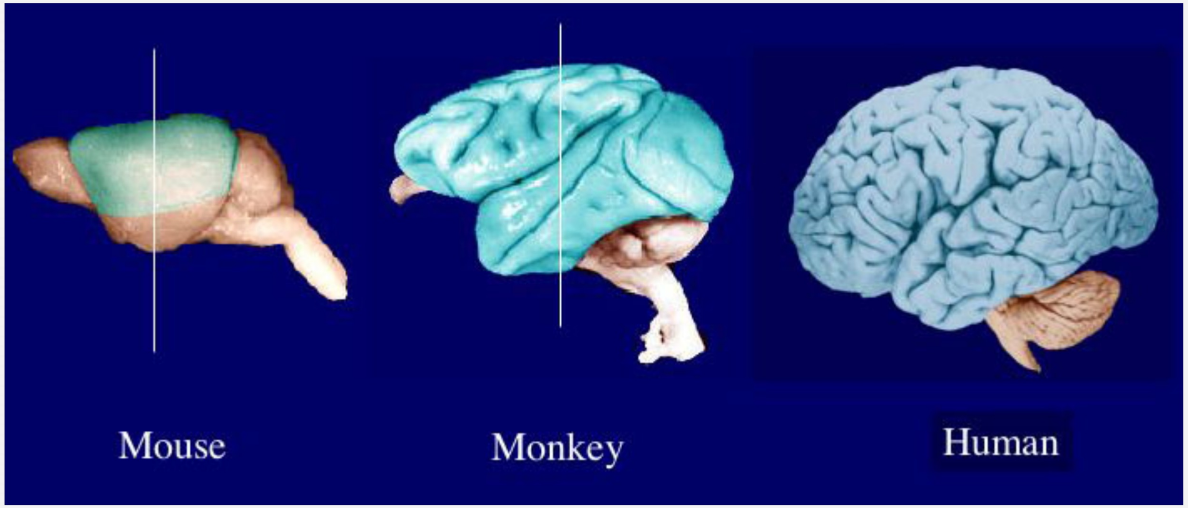

Fig. 1-2 Comparison of neocortex among mouse, monkey and human.

The neocortical surfaces are colored blue. (Adapted from Comparative Mammalian Brain Collections)

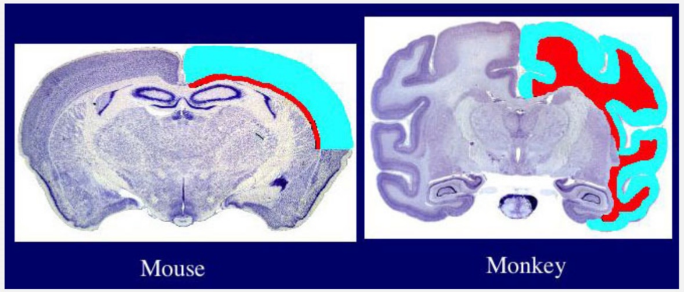

Fig. 1-3 Comparison of neocortex between mouse and monkey.

Neocortex is shown by blue. The white matter, which consists of nerve fibers that travels from and to the cortex is shown by red. (The Nissl photos are from Comparative Mammalian Brain Collections)

Cortex is a thin layered structure surrounding mammalian brains. It is the hallmark of mammalian brains and not present in birds or in reptiles.

Brain regions and their function

- Different brain regions specialize in different functions
- The functions of brain regions have historically been observed by studying the deficits of people with lesions to these regions
- Example: sensory areas, memory areas, motor areas, spatial navigation areas, emotion areas, decision areas
- They are now a multitude of tools to study the brain regions and their role:
  - invasive: electrodes, optogenetic microscopy
  - non-invasive: fMRI, MEG, ultrasound imaging

Exemple: ablation of the hippocampus and amnesia

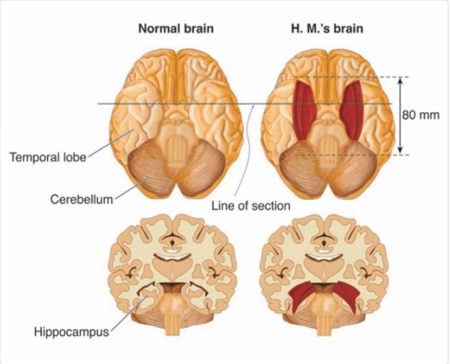

Bilateral resection of the anterior temporal lobe in patient HM.

Biography [edit]

Henry Molaison was born on February 26, 1926, in Manchester, Connecticut, and experienced intractable epilepsy that has sometimes been attributed to a bicycle accident at the age of seven.[note 1] He had minor or partial seizures for many years, and then major or tonic-clonic seizures following his 16th birthday. He worked for a time on an assembly line but, by the age of 27, he had become so incapacitated by his seizures, despite high doses of anticonvulsant medication, that he could not work nor lead a normal life.

In 1953, Molaison was referred to William Beecher Scoville, a neurosurgeon at Hartford Hospital.[3] Scoville localized his epilepsy to the left and right medial temporal lobes (MTLs) and suggested their surgical resection. On September 1, 1953, Scoville removed Molaison's medial temporal lobes on both hemispheres including the hippocampi and most of the amygdalae and entorhinal cortex, the major sensory input to the hippocampi.[14] His hippocampi appeared entirely nonfunctional because the remaining 2 cm of hippocampal tissue appeared to have atrophied and some of his anterolateral temporal cortex was also destroyed.

After the surgery, which was partially successful in controlling his seizures, Molaison developed severe anterograde amnesia: although his working memory and procedural memory were intact, he could not commit new events to his explicit memory. According to some scientists, he was impaired in his ability to form new semantic knowledge.[15]

Researchers argue over the extent of this impairment. He also had moderate retrograde amnesia, and could not remember most events in the one- to two-year period before surgery, nor some events up to 11 years before, meaning that his amnesia was temporally graded.

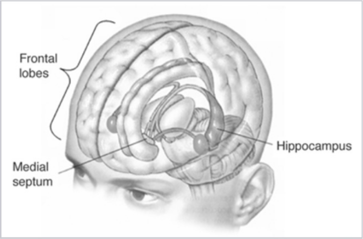

Most of Molaison's two hippocampi were removed bilaterally.

His case was first reported by Scoville and Brenda Milner in 1957, who referred to him by "H.M."[14] His full name was not revealed to the wider public until after his death.[11] While researchers had told him of the significance of his condition and of his renown within the world of neurological research, he was unable to internalize such facts as memories.[11]

Near the end of his life, Molaison regularly filled in crossword puzzles.[16] He was able to fill in answers to clues that referred to pre-1953 knowledge. For post-1953 information he was able to modify old memories with new information. For instance, he could add a memory about Jonas Salk by modifying his memory of polio.[2]

Example: The motor and somato-sensory homonculus

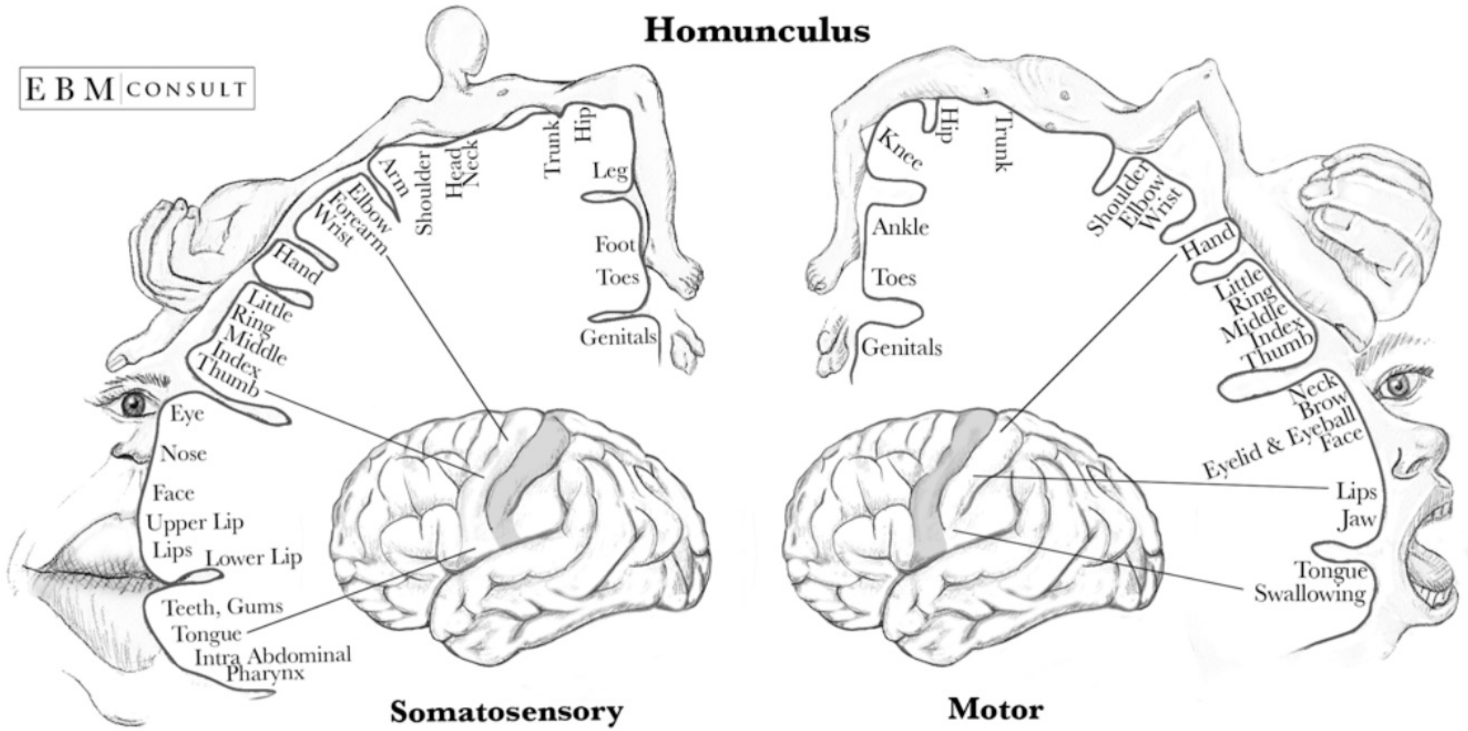

Brain regions and their function

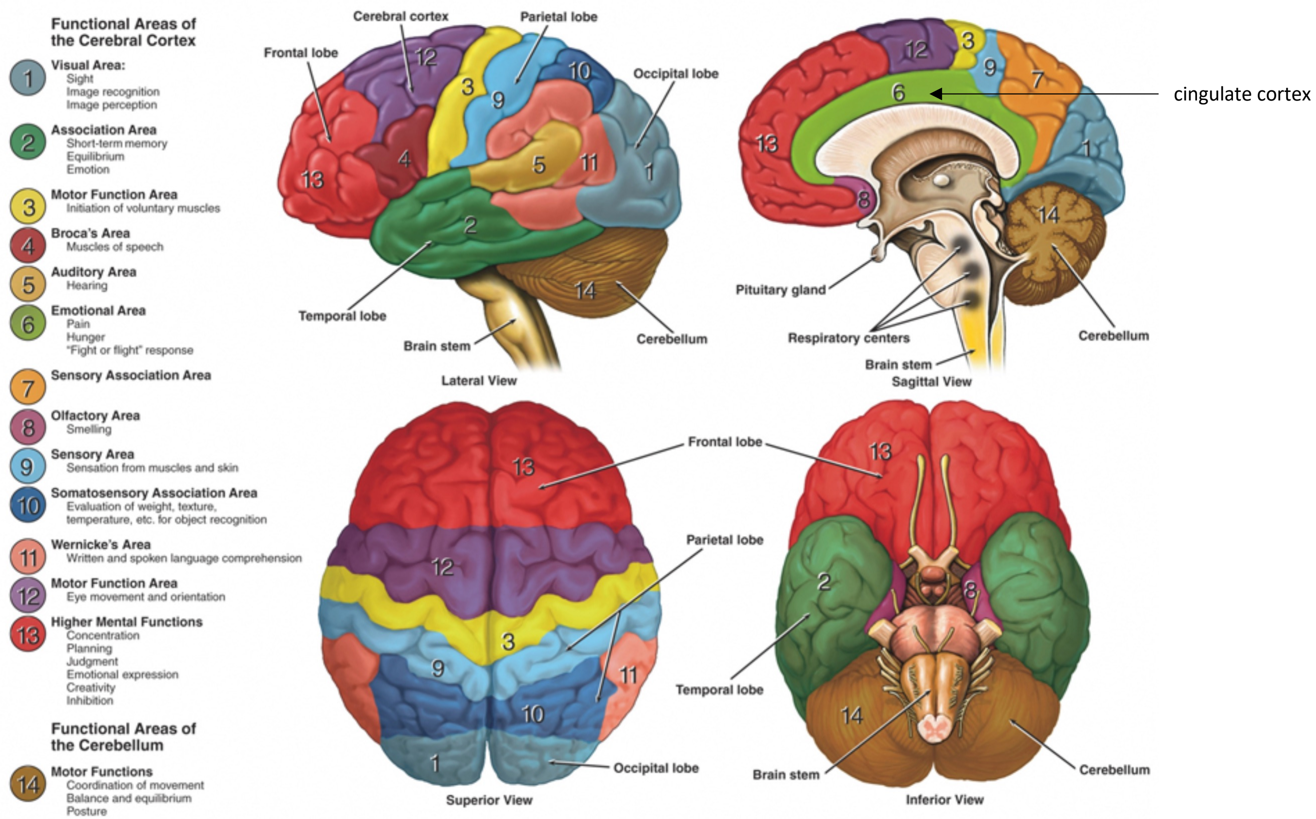

Brain tissue under the microscope

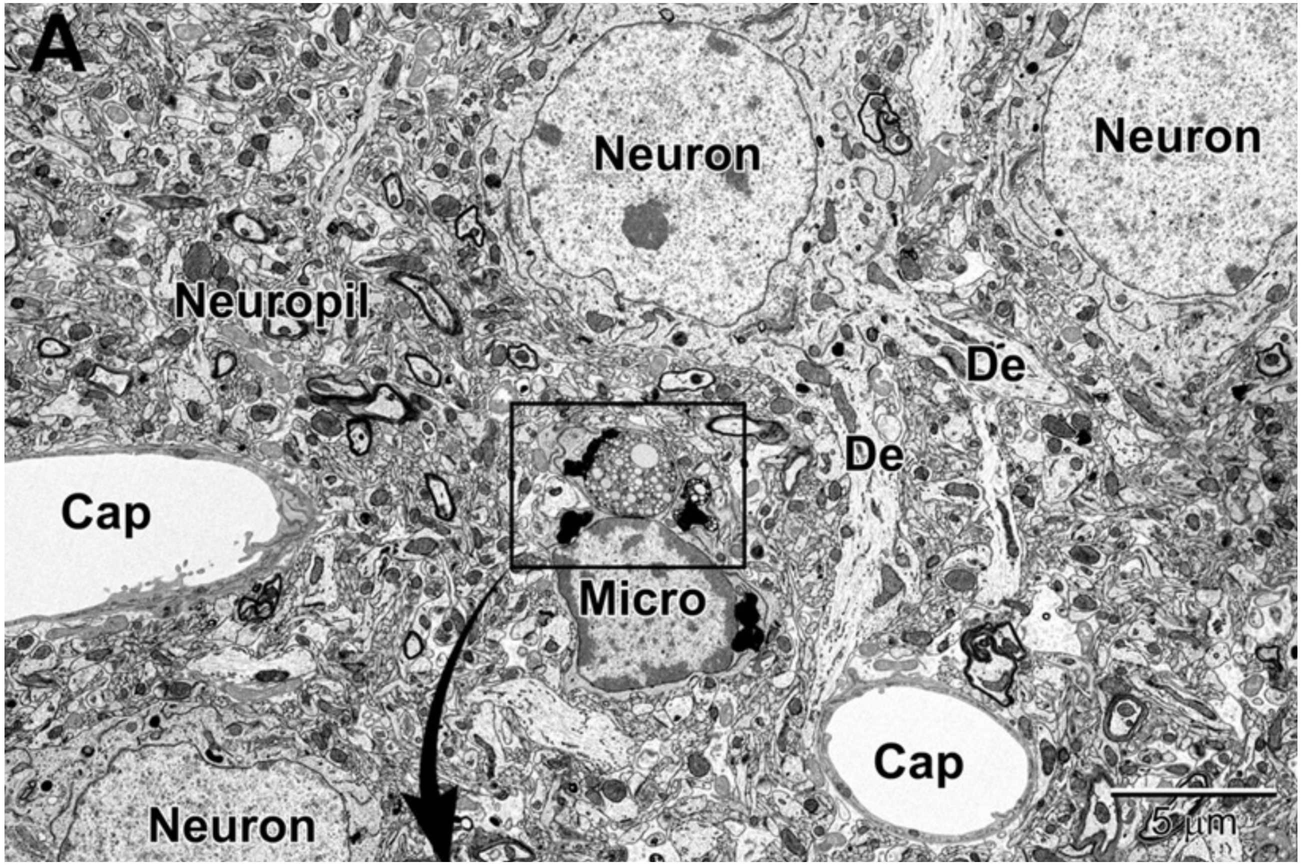

Cortex (mouse)

Microglia are strategically situated between neurons and capillaries (Cap), and function as the resident immune cell and phagocyte required for maintaining brain health throughout life.

source: https://www.frontiersin.org/articles/10.3389/fcell.2021.629503/full

Brain tissue reconstructed

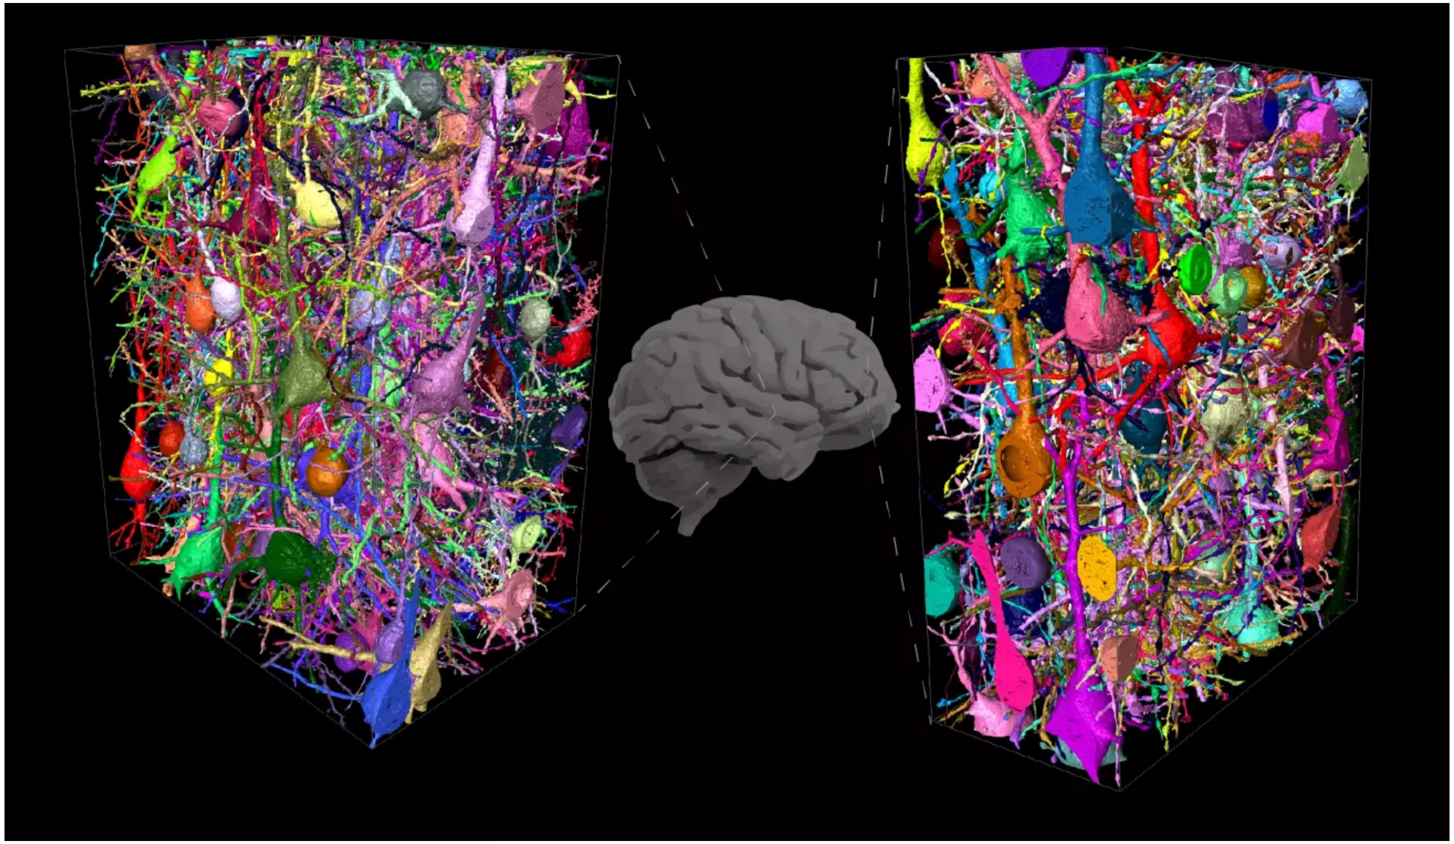

source: https://brain.mpg.de/547207/silence-for-thought?c=83711

Neurons

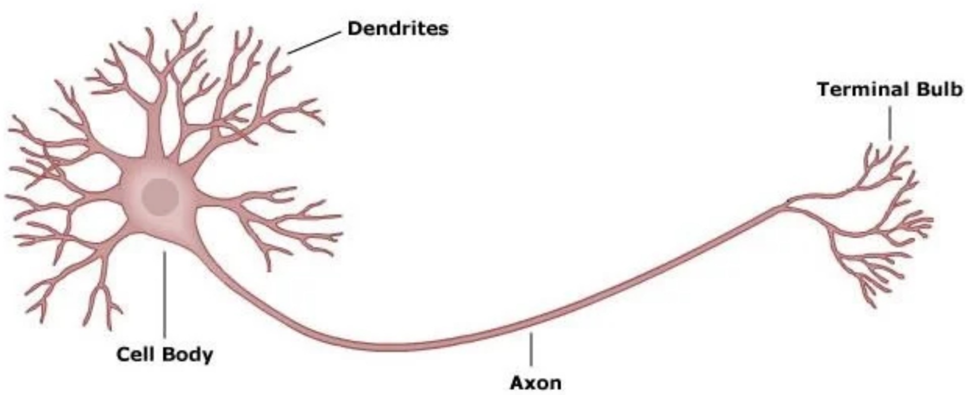
Dendrites
Terminal Bulb
Cell Body
Axon

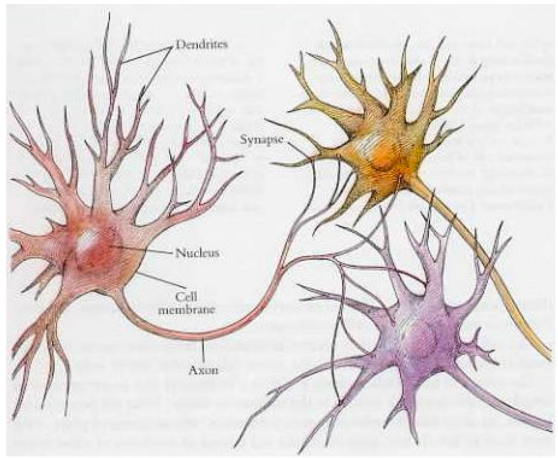
Dendrites
Synapse
Nucleus
Cell
membrane
Axon

Action potential ("spike")

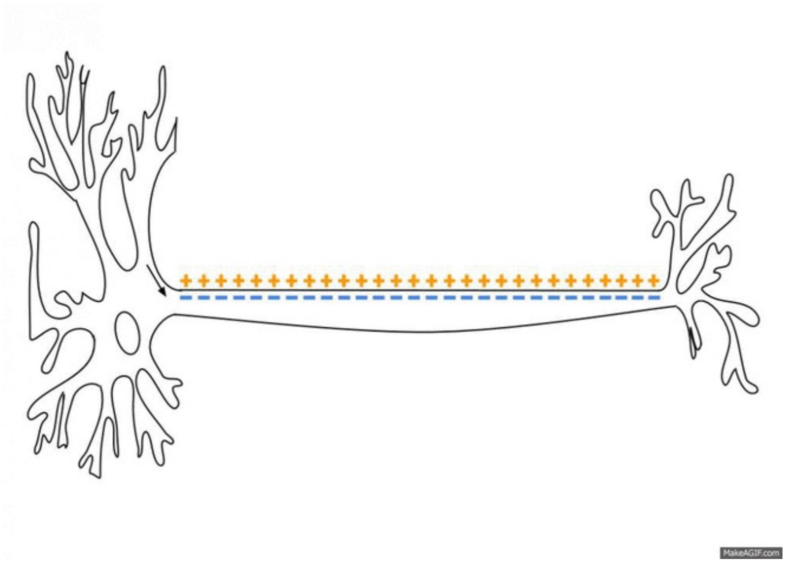

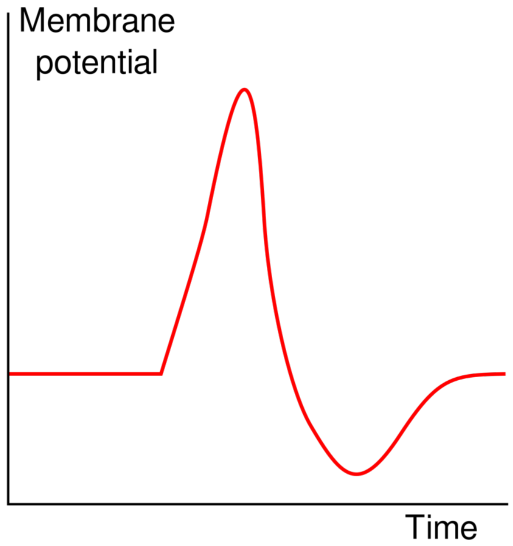

https://oertx.highered.texas.gov/courseware/lesson/1790/student-old/?task=2
https://courses.lumenlearning.com/suny-ap1/chapter/the-action-potential/

Synapse

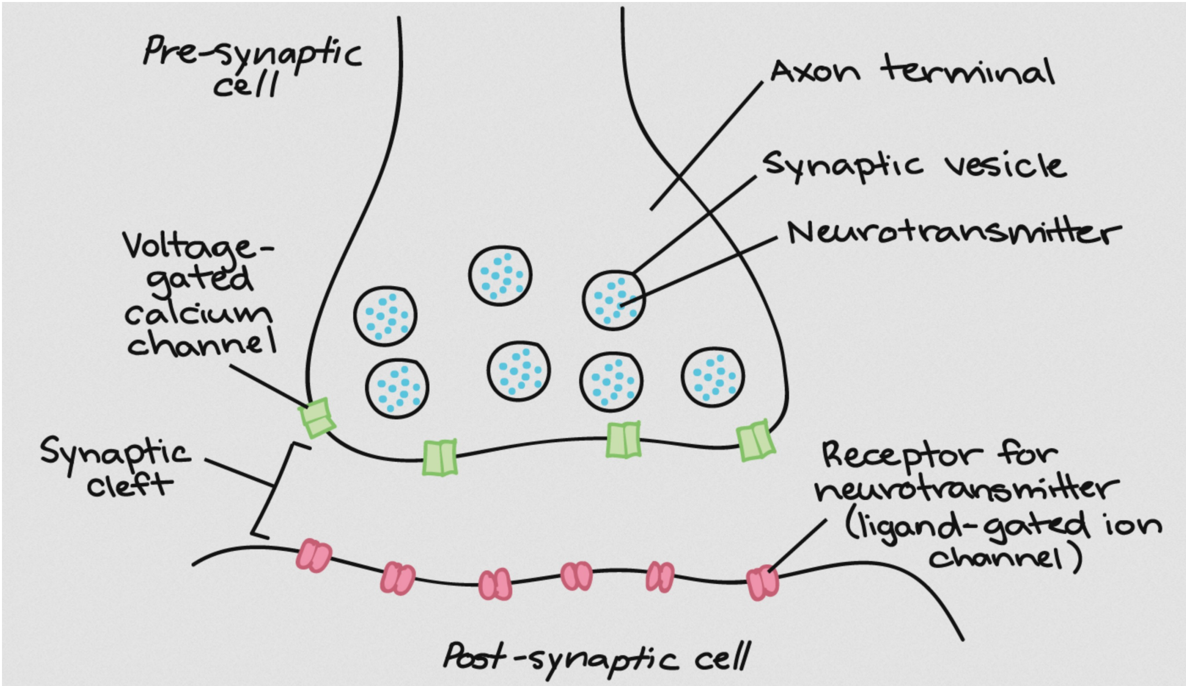

Synapses can be inhibitory or excitatory

The leaky integrate-and-fire model of a neuron

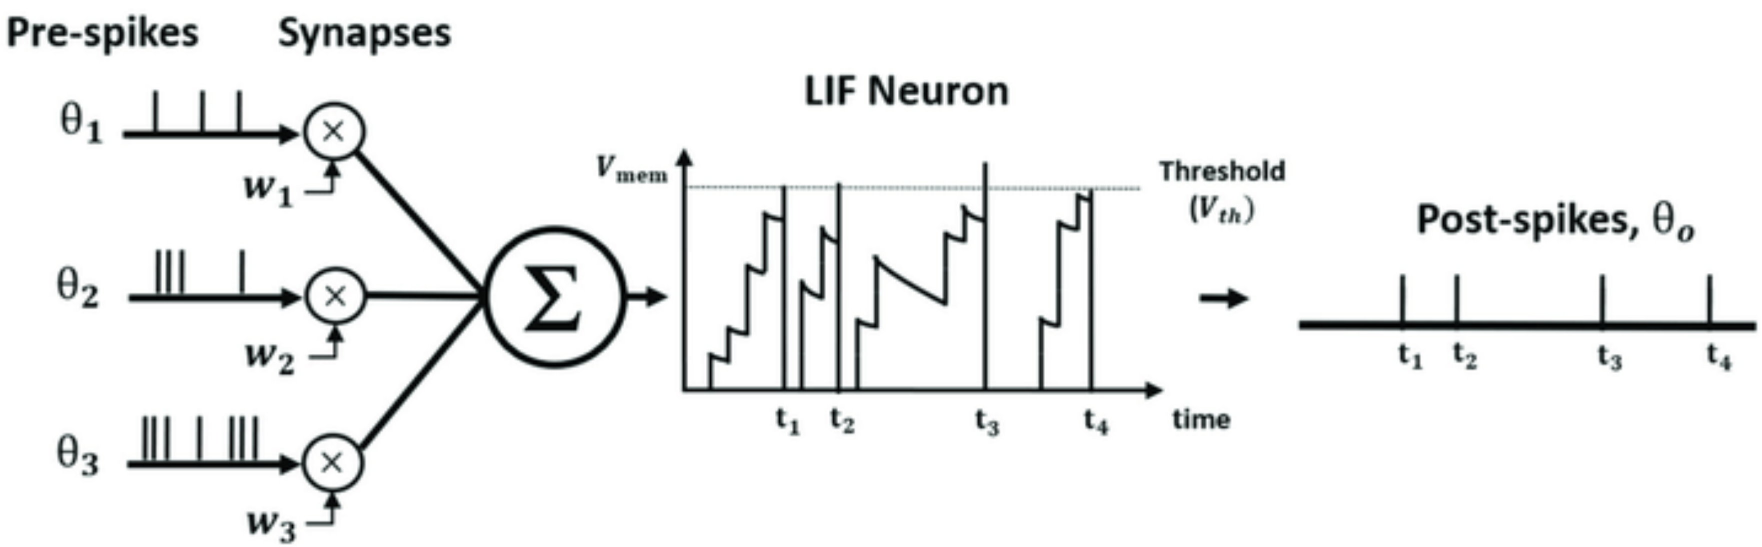

Synaptic plasticity

- **Hebb's rule (1949):**

  "Neurons than fire together wire together"

- **Spike-timing-dependent plasticity (Markram, Tsodyks 2000):**

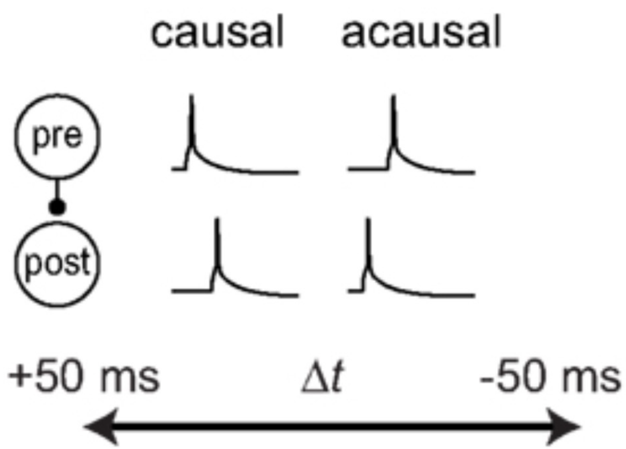

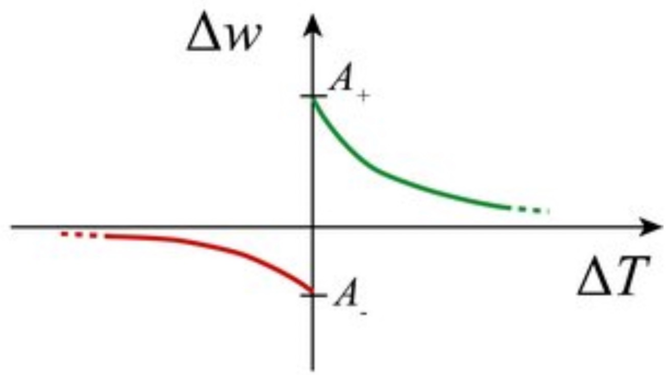

Neuron types

Ramon y Cajal (1852-1934)

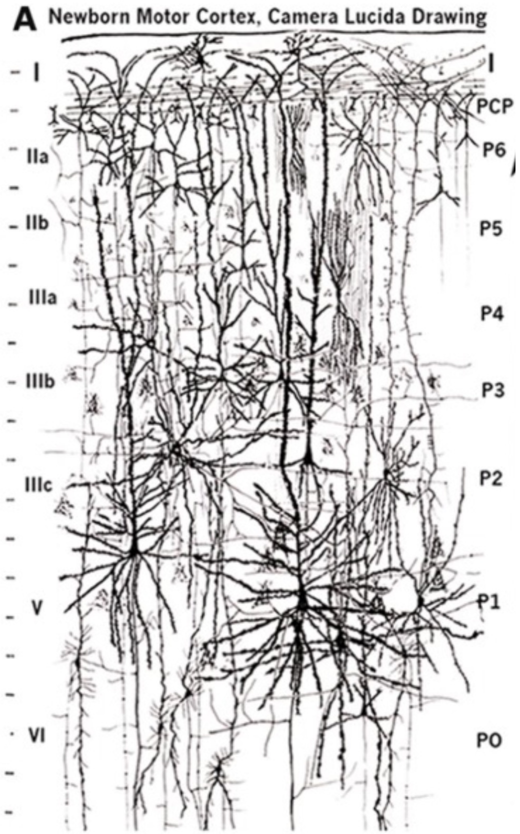

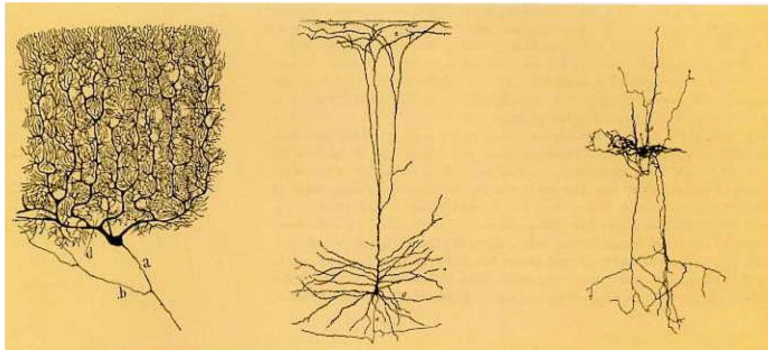

Left: The cerebellar Purkinje cell, shown in a drawing by Santiago Ramon y Cajal, presents an extreme in neuronal specialization. The dense dendritic arborization is not bushlike in shape, but is flat, in the plane of the paper, like a cedar frond. Through the holelike spaces in this arborization pass millions of tiny axons, which run like telegraph wires perpendicular to the plane of the paper. The Purkinje cell's axon gives off a few initial branches close to the cell body and then descends to cell clusters deep in the cerebellum some centimeters away, where it breaks up into numerous terminal branches. At life size, the total height of the cell (cell body plus dendrites) is about 1 millimeter. *Middle*: Ramon y Cajal made this drawing of a pyramidal cell in the cerebral cortex stained The cell body is the small black blob. *Right*: This drawing by Jennifer Lund shows a cortical cell that would be classed as "stellate". The dark blob in the center is the cell body. Both axons (fine) and dendrites (coarse) branch and extend up and down for distance of 1 millimeter.

Neuron types

Principles of connectivity among morphologically defined cell types in adult neocortex

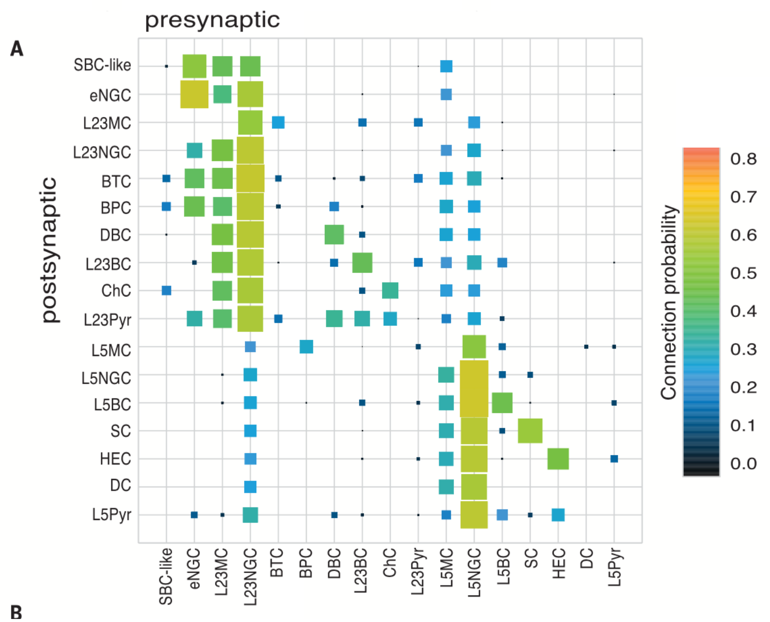

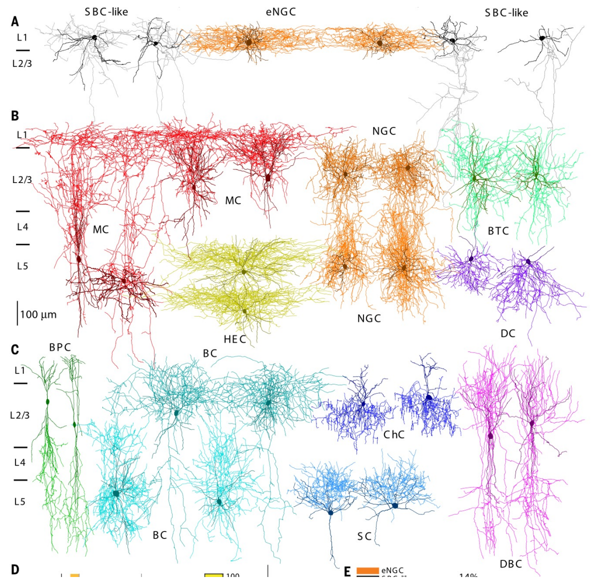

Receptive fields: ON and OFF cells in the retina

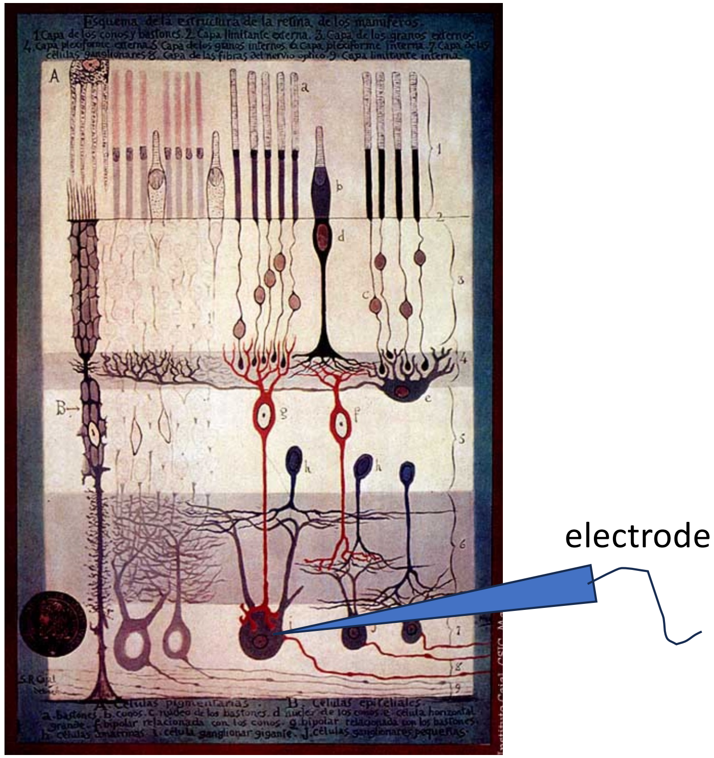

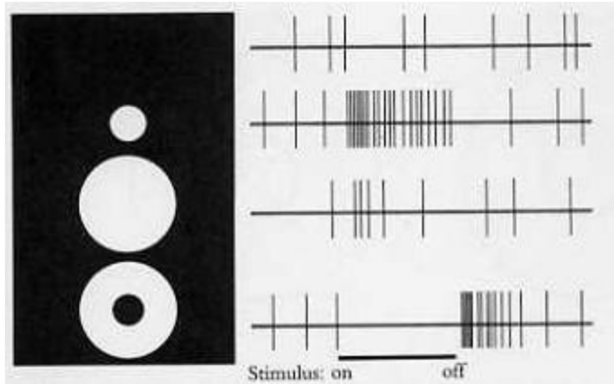

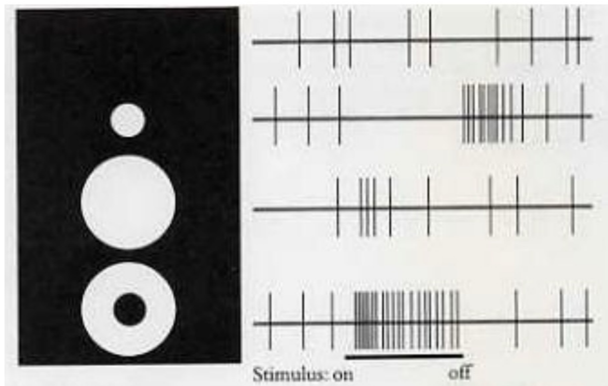

Stephen Kuffler (1953)

Receptive fields: orientation selective cells in V1

Simple cell

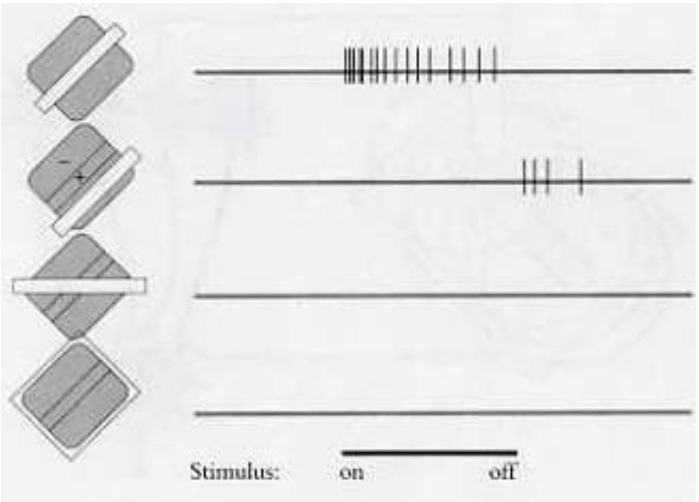

Complex cell

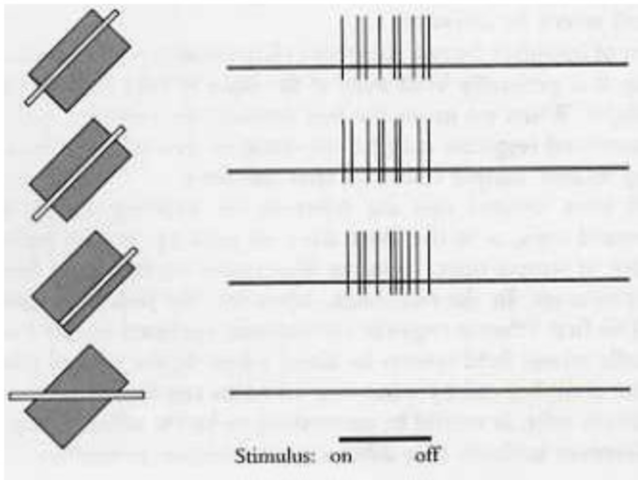

David Hubel and Torsten Wiesel, 1962

Development of the visual system depends on environment

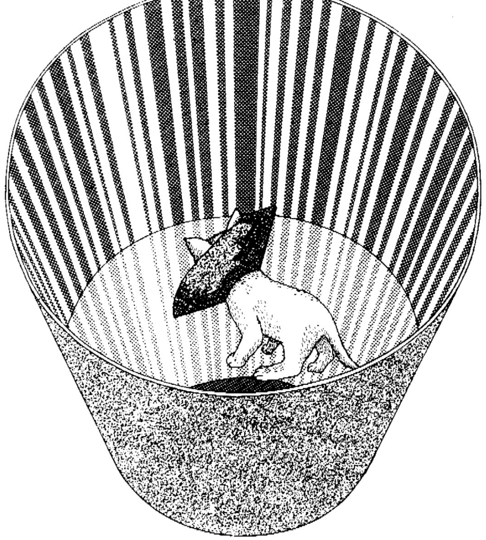

Fig. 1. The visual display consisted of an upright plastic tube, about 2 m high, with an internal diameter of 46 cm. The kitten, wearing a black ruff to mask its body from its eyes, stood on a glass plate supported in the middle of the cylinder. The stripes on the walls were illuminated from above by a spotlight. The luminance of the dark bars was about 10 cd. m^{-2} and of the bright stripes about 130 cd. m^{-2}: they were of several different widths. For this diagram the top cover and the spotlight have been removed from the tube.

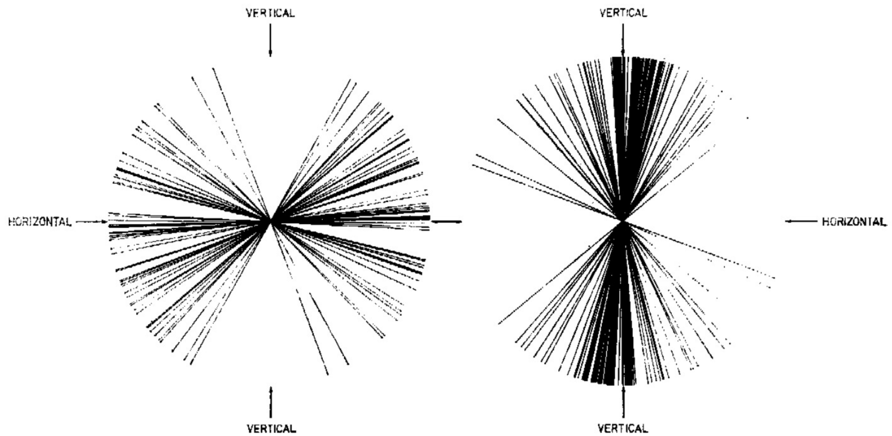

Fig. 2. These polar histograms show the distributions of optimal orientations for fifty-two neurones from a horizontally experienced cat on the left, and seventy-two from a vertically experienced cat on the right. The slight torsion of the eyes, caused by the relaxant drug, was assessed by photographing the pupils before and after anaesthesia and paralysis. A correction has been applied for torsion, so the polar plots are properly orientated for the cats' visual fields. Each line shows the optimal orientation for a single neurone. For each binocular cell the line is drawn at the mean of the estimates of optimal orientation in the two eyes. No units have been disregarded except for one with a concentric receptive field and hence no orientational selectivity.

Blakemore and Cooper 1970

Place cells

Fig. 1: In 2005 distinct firing patterns, associated with location and navigation, were traced to the entorhinal cortex (shown here in pink), a lobe which inputs into the hippocampus (blue)

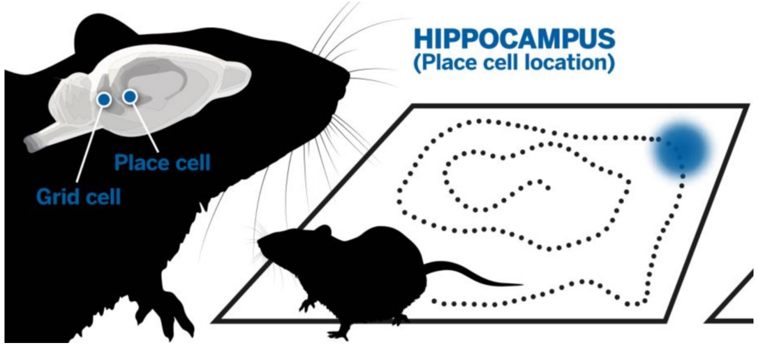

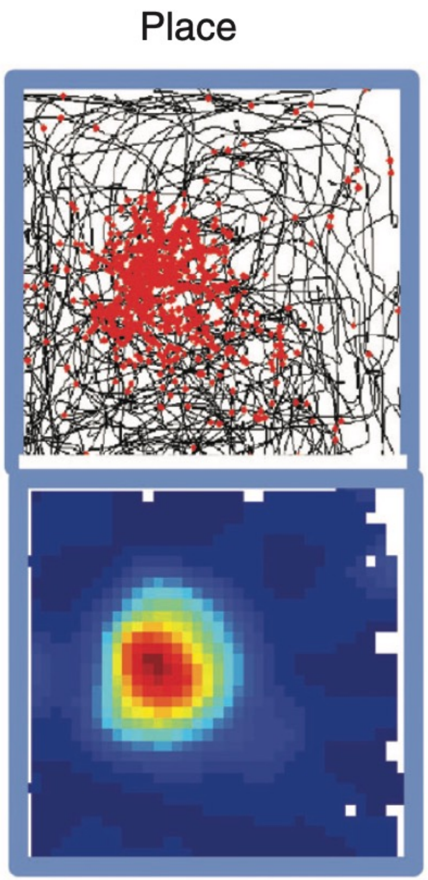

John O'Keefe 1971

Grid cells

Fig. 1: In 2005 distinct firing patterns, associated with location and navigation, were traced to the entorhinal cortex (shown here in pink), a lobe which inputs into the hippocampus (blue)

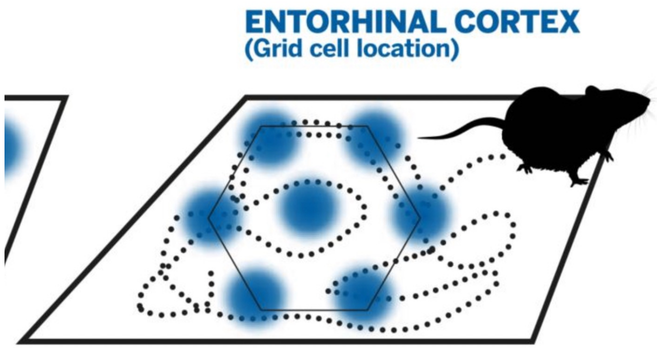

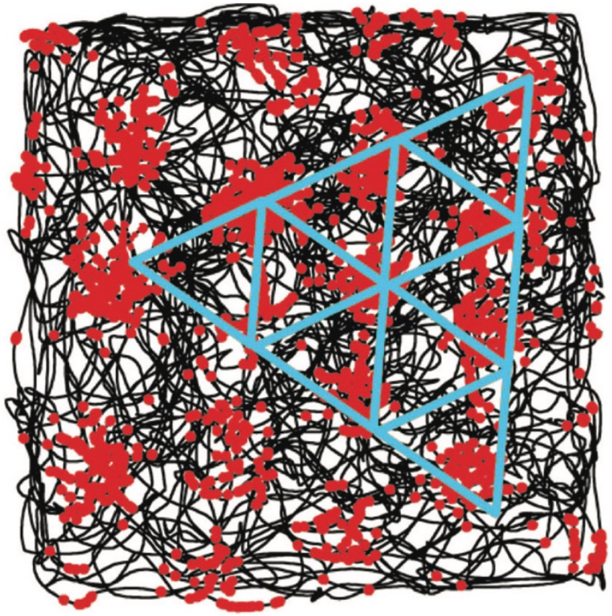

Edvard & May-Britt Moser, 2005

What is a computational theory of the brain?

?

?

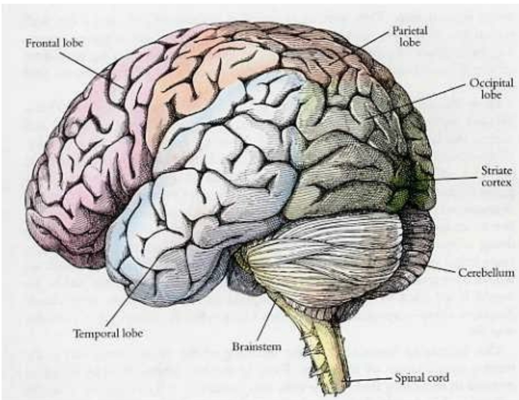

?

What is a computational theory of the brain?

- The brain performs some functions:
  examples: seeing, acting, deciding, navigating, memorizing, producing speech

- Computational theories of the brain aim at understanding how the brain performs these functions:
  - algorithm used?
  - implementation on neural substrate?
  - with given constraints (e.g., number of neurons, energy consumptions)
  - also: what really are the functions that the brain is performing??

After the break

- Ray and Andrea will present the topics that will be covered in class.
- For this Friday end of day, you will need to choose 1-3 topics, and fill in you preferences on MyCourse (“Preferred Topics”).

- T1_PSYCHOPHYSICS
  - ADDITIONAL_MATERIAL
  - ST1_JIM_BINDING_MODEL
  - ST2_FIT_FEATURE_INTEGRATION_THEORY
  - ST3_ART_ADAPTIVE_RESONANCE_THEORY
- T2_EFFICIENT_CODING
  - ST1_RETINAL_CELL_MODEL
  - ST2_VISUAL_CORTEX_MODEL
- T3_DEEP_LEARNING_MODELS
  - ST1_VISUAL_SYSTEM
  - ST2_MOTOR_AND_NAVIGATION
- T4_MEMORY_HOPFIELD
- T5_ATTRACTOR_NETWORKS
- T6_CONTROL_THEORY
- T7_CHAOS

Homework

- Choose 1 to 3 topics/subtopics (see slides and "Reading Material" on MyCourse)
- Fill your preference on MyCourse ("Preferred Topics"), no later than this Friday end of day. Check that you will be available on the week where the project will be presented to the class.
- If there is only 1 topic you like, you may choose this 1 only topic on MyCourse. We will try to accommodate you. However, if we cannot, you may be assigned to a topic you did not choose at all.
- If you would like to drop out, drop out now. Choosing a topic = committing to present

Cortex

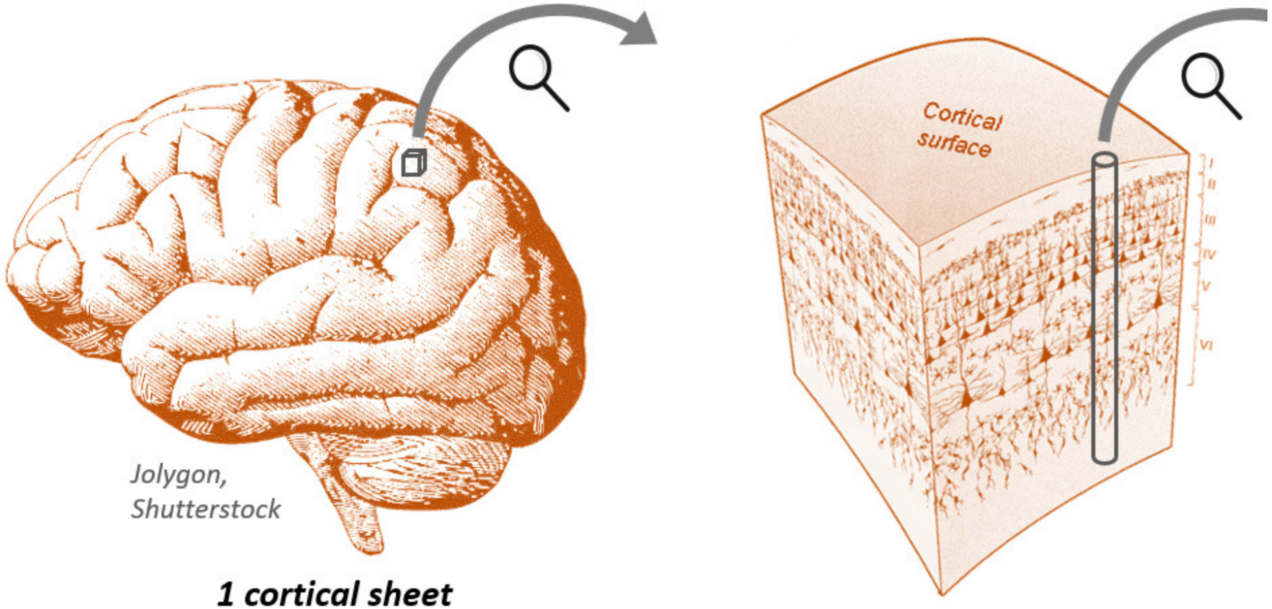

2.5 mm thick

1 cortical sheet
2 million macrocolumns
200 million minicolumns
20 billion neurons

1 macrocolumn
100 minicolumns
10.000 neurons

Synaptic plasticity

Spike-timing-dependent plasticity

2000

LETTER Communicated by Laurence Abbott

An Algorithm for Modifying Neurotransmitter Release Probability Based on Pre- and Postsynaptic Spike Timing

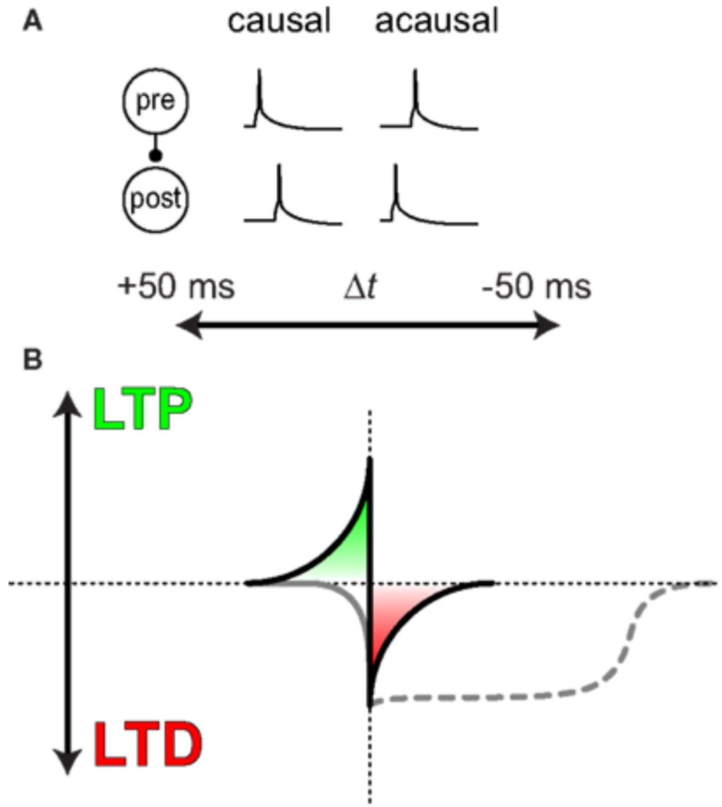

Hebb's rule (1949).

$$w_{ji}^{k+1} = w_{ji}^{k} + \Delta w_{ji}^{k}, \qquad (1)$$

where

$$\Delta w_{ji}^{k} = Ca_i^k X_j^k \qquad (2)$$

- Hebb's original learning rule (2) referred exclusively to excitatory synapses, and has the unfortunate property that it can only increase synaptic weights, thus washing out the distinctive performance of different neurons in a network, as the connections drive into saturation..

Basic facts about brain: cells types

Ramon y Cajal

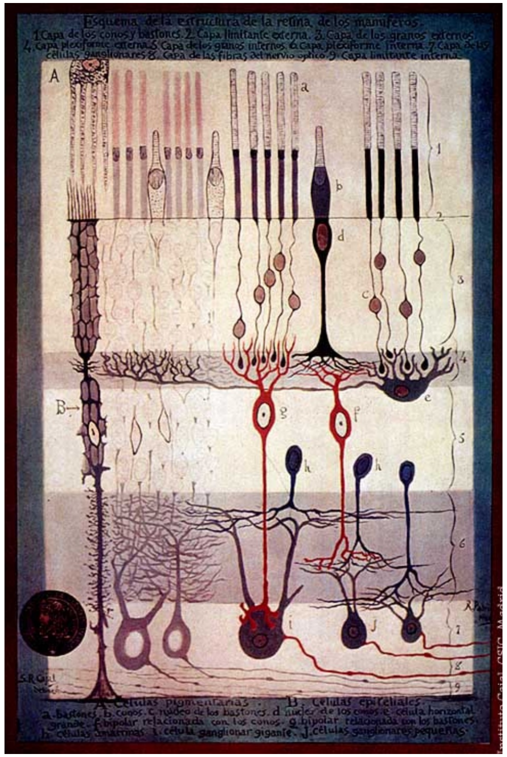

Introduction to the visual system

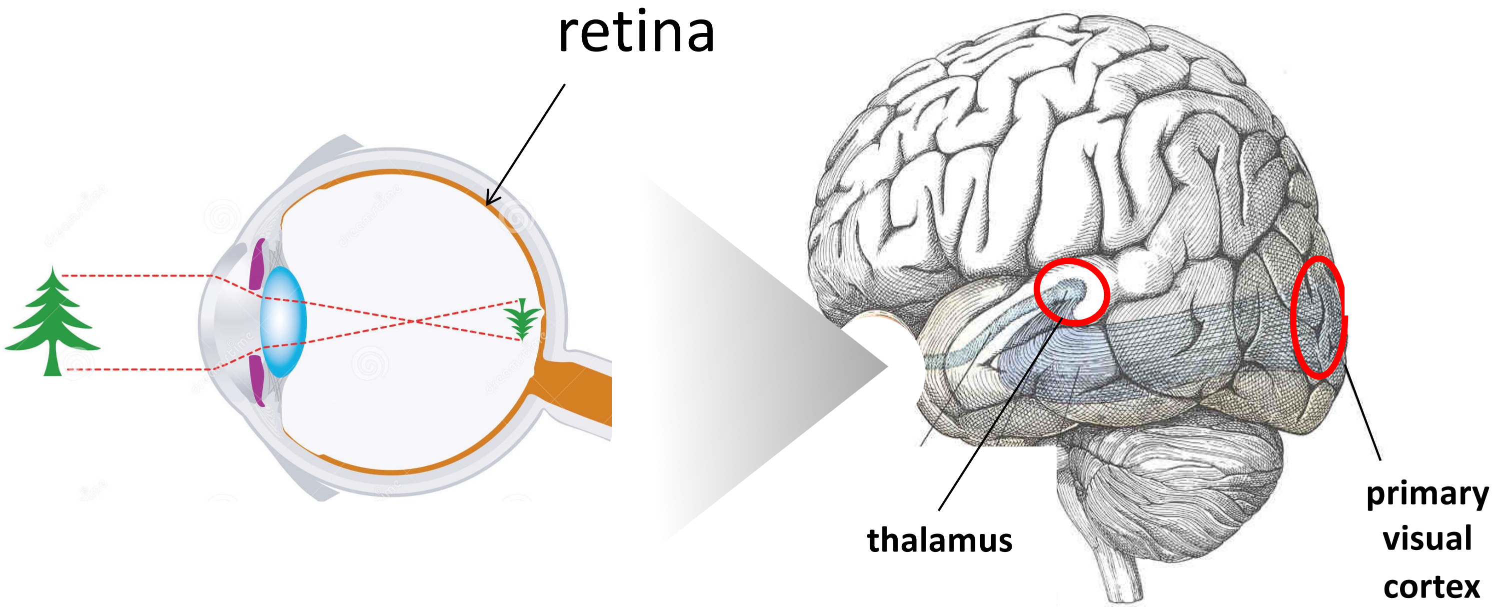

Group formation mechanism – 8 groups

- For next week, each student will have selected 1 to 3 topics/subtopic, on MyCourse.
- We will assign each student to a topic/subtopic based on theses preferences
- In case of conflict (too many students interested by the same topic), we will try to find an arangement or toss a coin.
- The groups are formed around the topics chosen. Most groups will have 2 students. No group can have more than 3 students.

Dates and rooms

Thu 25.04.2024 14:15 - 16:00, R001/M205

Thu 02.05.2024 14:15 - 16:00, R001/M205

(9.5. Ascension Day - no teaching)

Thu 16.05.2024 14:15 - 16:00, R001/M205

Thu 23.05.2024 14:15 - 16:00, R001/M205

Thu 30.05.2024 14:15 - 16:00, R001/M205

Thu 06.06.2024 14:15 - 16:00, R001/M205 (on exam week)

SUPPL SLIDES

Basic facts about brain: synaptic plasticity

LETTER Communicated by Laurence Abbott

An Algorithm for Modifying Neurotransmitter Release Probability Based on Pre- and Postsynaptic Spike Timing

2000

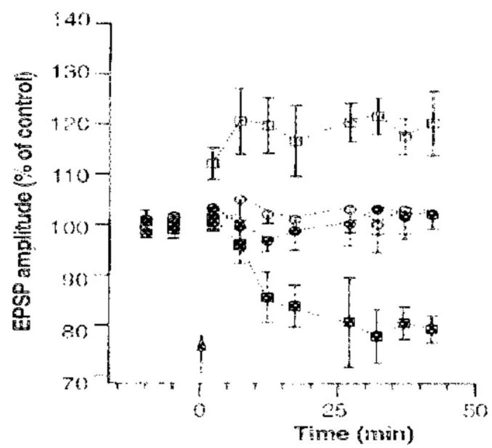

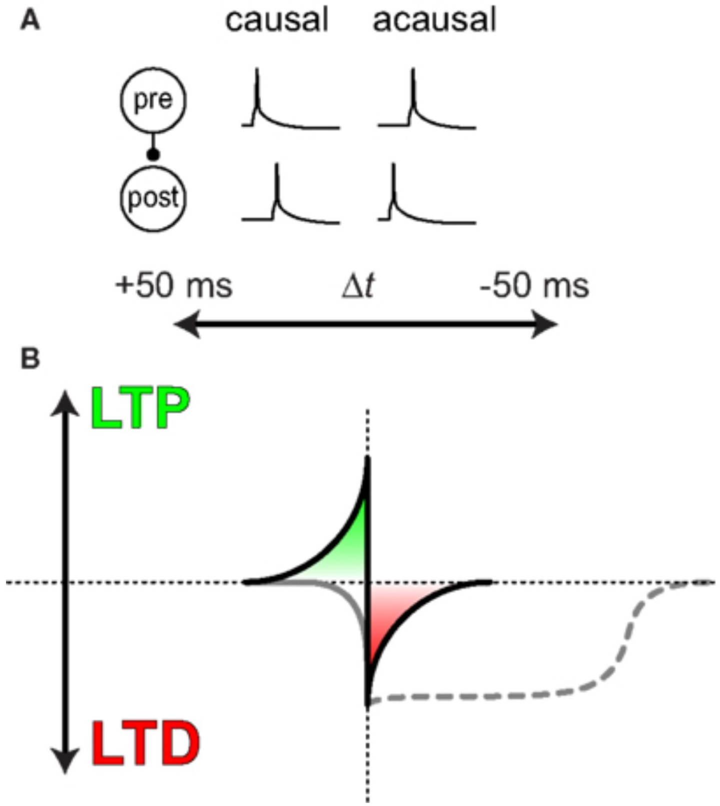

Spike-timing-dependent plasticity

Visual system: ventral and dorsal stream

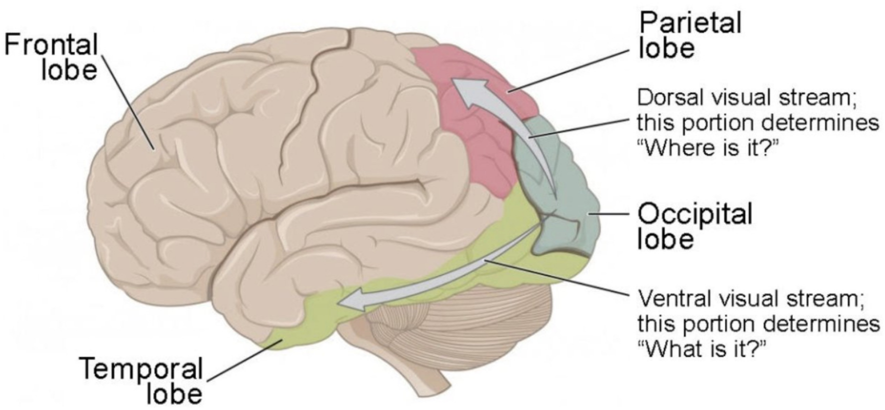

Dorsal and ventral streams

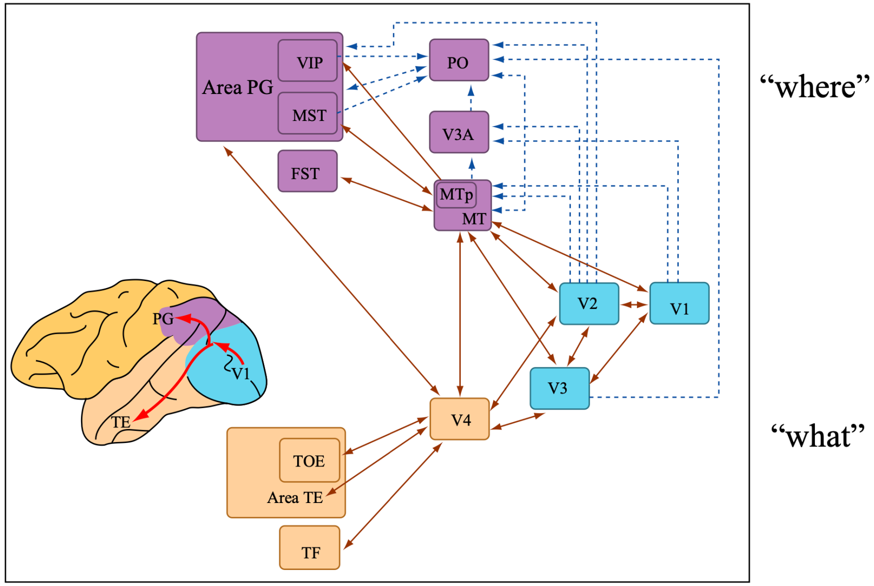

Figure by MIT OpenCourseWare.

DOI: 10.1093/brain/awg248 Advanced Access publication September 23, 2003 *Brain* (2003), **126,** 2463–2475

Ventral occipital lesions impair object recognition but not object-directed grasping: an fMRI study

Thomas W. James,[1] Jody Culham,[2] G. Keith Humphrey,[2] A. David Milner[3] and Melvyn A. Goodale[2]

[1]*Vanderbilt Vision Research Center, Vanderbilt University, Nashville, USA,* [2]*CIHR Group on Action and Perception, University of Western Ontario, London, Canada and* [3]*Wolfson Research Institute, University of Durham, Durham, UK*

Correspondence to: Thomas W. James, Psychology Department, 111 21st Street S., Nashville, TN 37203, USA Email: tom.james@vanderbilt.edu

Summary

D.F., a patient with severe visual form agnosia, has been the subject of extensive research during the past decade. The fact that she could process visual input accurately for the purposes of guiding action despite being unable to perform visual discriminations on the same visual input inspired a novel interpretation of the functions of the two main cortical visual pathways or 'streams'. Within this theoretical context, the authors proposed that D.F. had suffered severe bilateral damage to her occipitotemporal visual system (the 'ventral stream'), while retaining the use of her occipitoparietal visual system (the 'dorsal stream'). The present paper reports a direct test of this idea, which was initially derived from purely behavioural data, before the advent of modern functional neuroimaging. We used functional MRI to examine activation in her ventral and dorsal streams during object recognition and object-directed grasping tasks. We found that D.F. showed no difference in activation when presented with line drawings of common objects compared with scrambled line drawings in the lateral occipital cortex (LO) of the ventral stream, an area that responded differentially to these stimuli in healthy individuals. Moreover, high-resolution anatomical MRI showed that her lesion corresponded bilaterally with the location of LO in healthy participants. The lack of activation with line drawings in D.F. mirrors her poor performance in identifying the objects depicted in the drawings. With coloured and greyscale pictures, stimuli that she can identify more often, D.F. did show some ventral-stream activation. These activations were, however, more widely distributed than those seen in control participants and did not include LO. In contrast to the absent or abnormal activation observed during these perceptual tasks, D.F. showed robust activation in the expected dorsal stream regions during object grasping, despite considerable atrophy in some regions of the parietal lobes. In particular, an area in the anterior intraparietal sulcus was activated more for grasping an object than for just reaching to that object, for both D.F. and controls. In conclusion, we have been able to confirm directly that D.F.'s visual form agnosia is associated with extensive damage to the ventral stream, and that her spared visuomotor skills are associated with visual processing in the dorsal stream.

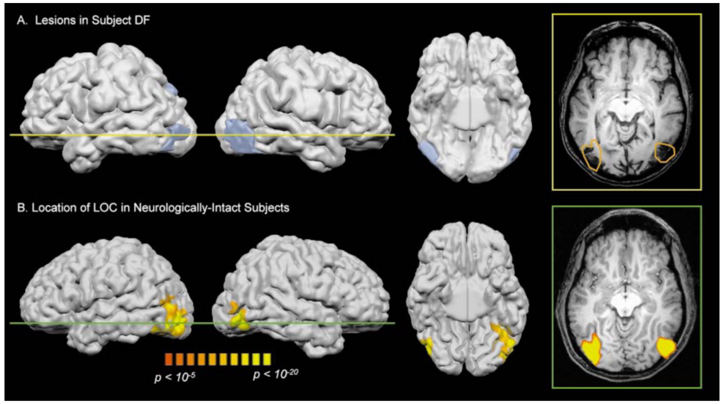

Fig. 4 Ventral stream lesions in D.F. shown in comparison with the expected location of the lateral occipital complex in healthy participants. (**A**) D.F.'s brain has been rendered at the pial surface (outer grey matter boundary). Lesions were traced on slices that indicated tissue damage and rendered on the pial surface in pale blue. Lateral views of the left and right hemispheres are shown, as is a ventral view of the underside of the brain. The rightmost image shows a slice through the lesions ($z = -8$). (**B**) The expected location of LOC based on group data from seven neurologically intact participants (Culham, 2003) is shown on one individual's pial surface and on a slice through the $z = -8$ plane. The activation in the slice is outlined in orange in panel **A** for comparison with the lesions in D.F.'s brain.

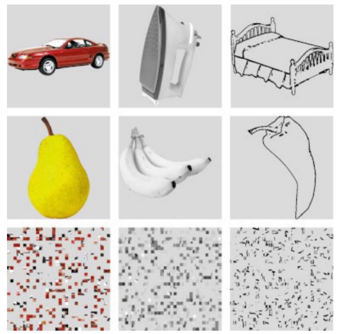

Fig. 1 Examples of intact and scrambled images of common objects in colour, greyscale and line drawing formats.

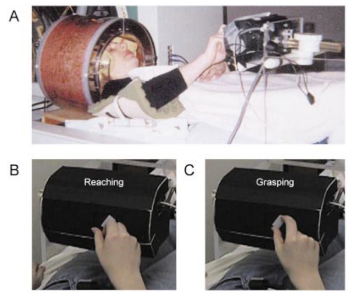

Fig. 2 Experimental setup for reaching and grasping experiment. (**A**) The participant lay supine in the magnet with the torso and head tilted to view illuminated shapes directly on an adjustable rotating drum, the 'grasparatus'. The participant's right arm was restrained to allow movement of the wrist and limited movement of the elbow, but no movement at the shoulder, which could be translated into head motion. (**B**) During reaching trials, participants extended their arm to touch the target object, an elongated translucent rectangle, with the knuckles. (**C**) During grasping trials, the participant used a precision grip with the index finger and thumb positioned along the long axis of the object.

fellows studying visual neuroscience at the University of Western Ontario. The age of the healthy participants ranged between 25 and 32 years. The ethical review boards of both the University of Western Ontario and the Robarts Research Institute approved protocols for the procedures, and the

Can a motion-blind patient reach for moving objects?

Thomas Schenk,[1] Norbert Mai,[1] Jochen Ditterich[1] and Josef Zihl[2]
[1]Neurologische Klinik, Klinikum Grosshadern, Ludwig-Maximilians-Universität München, Marchioninistr. 23, 81366 Germany
[2]Psychologisches Institut, Ludwig-Maximilians-Universität München, Germany

Keywords: akinetopsia, dorsal stream, human, interceptive action, V5/MT

Abstract

It has been claimed that the visual brain is organized in two separate processing streams for spatial vision: one for perception and one for action. To determine whether motion vision is also divided into vision for action and for perception we examined the interceptive behaviour of the motion-blind patient LM. The task for LM and three age-matched control subjects was to reach-and-grasp for an object that moved away. Three experiments were conducted to examine the effects on perfomance of target speed (Expt 1), observation time (Expt 2) and visual feedback (Expt 3). As LM is only able to reach for objects which move at 0.5 m/s or less, her performance is inferior to that of controls who can reach for objects moving at 1.0 m/s, but it is better than would be expected from her performance in psychophysical experiments on her motion vision. Kinematic analysis of LM's reaching movements showed that she adapted the speed of her moving hand to the speed of the target but only when full vision was available. In contrast to normal subjects, LM required long observation times and vision of her moving hand to produce successful reaching responses. Thus, the impairment of both perception and action in LM suggests that the motion area MT/V5 is located at an early stage of the extrastriate hierarchy and provides input to both the perception and the action processing streams.

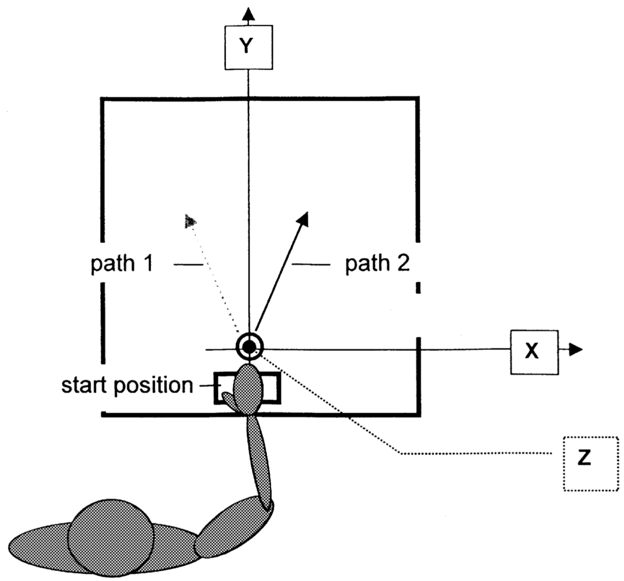

FIG. 1. Experimental setup as seen from above. The target moves either along path 1 or path 2 away from the subject. At the beginning of the trial the subject's hand rests on the hand button (starting position). The coordinate system used for the calculations described in the Appendix has the start position of the target object as its origin.

Ventral stream

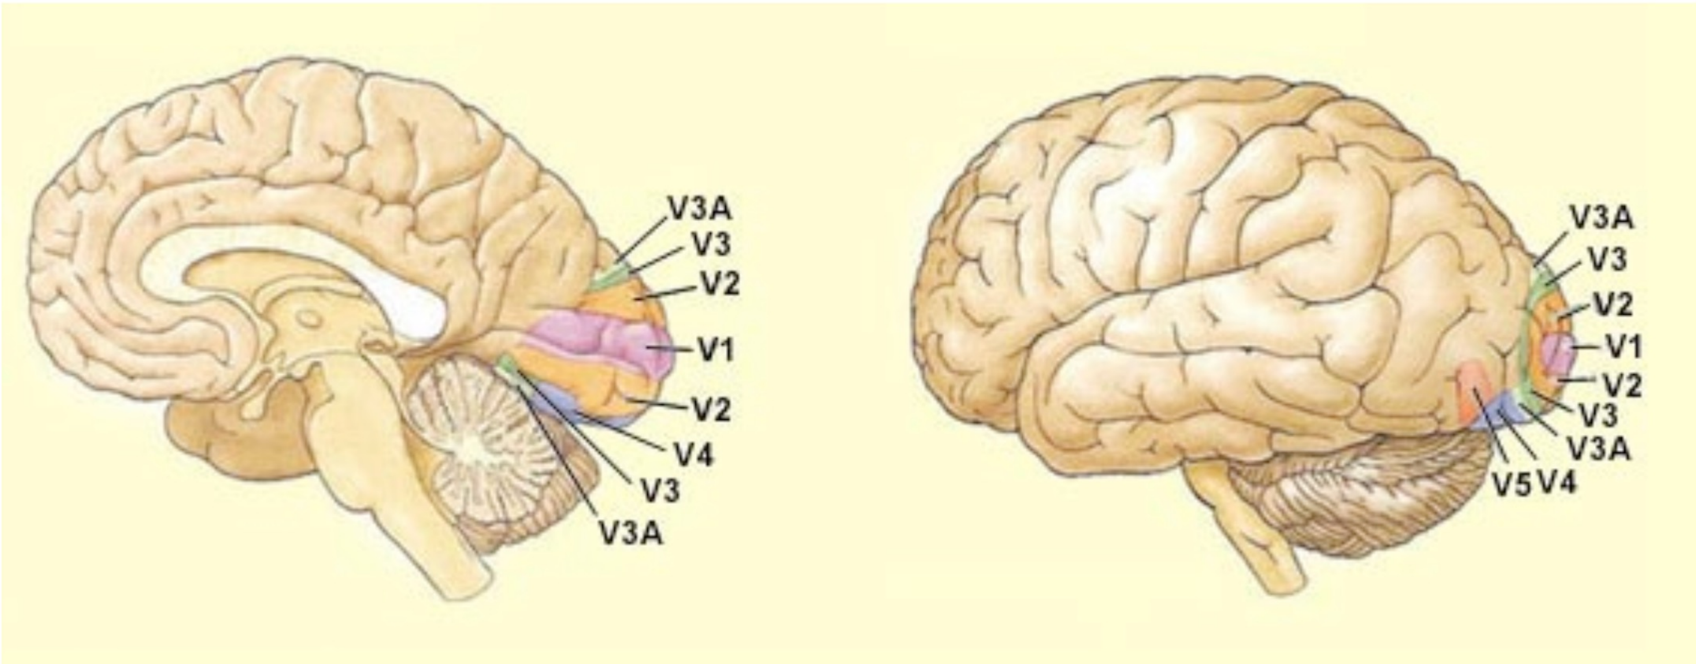

V2

A functional and perceptual signature of the second visual area in primates

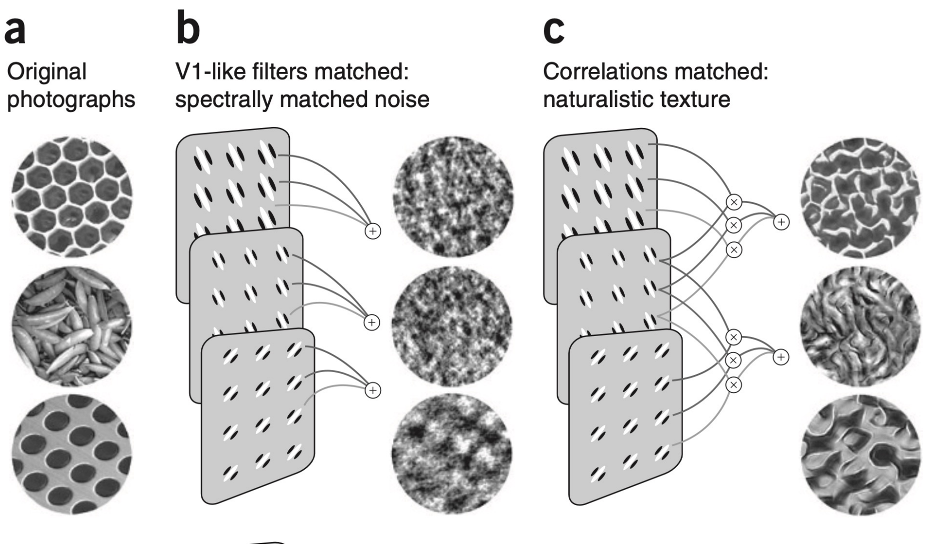

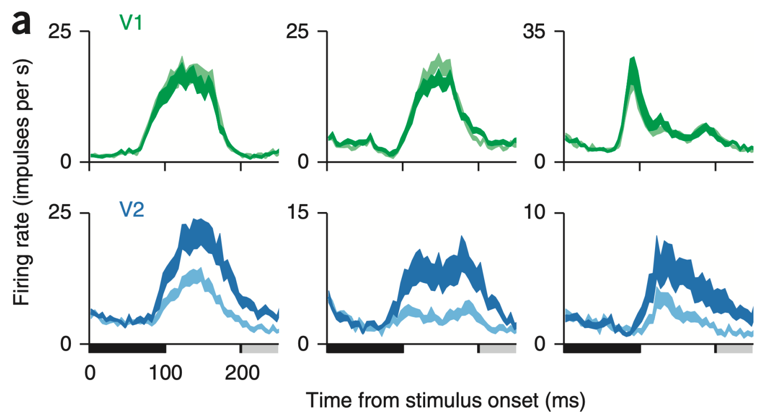

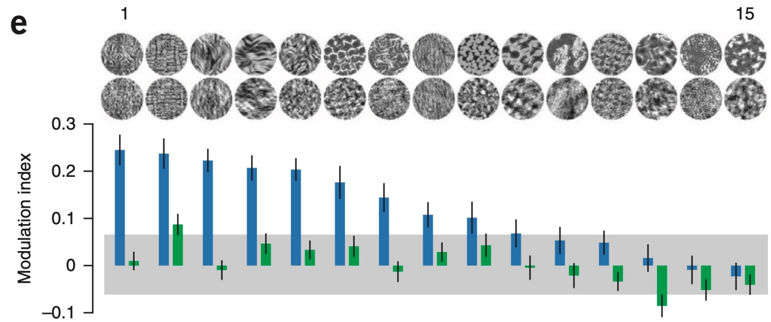

V4

Selectivity for Polar, Hyperbolic, and Cartesian Gratings in Macaque Visual Cortex

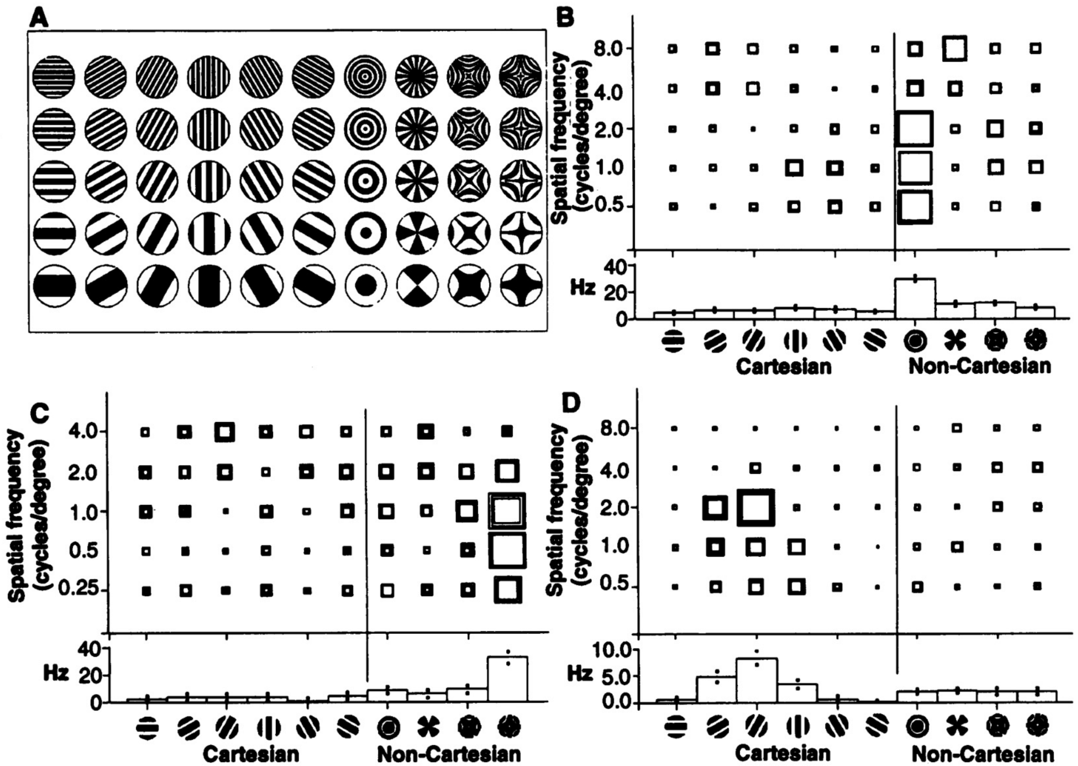

A Cortical Area Selective for Visual Processing of the Human Body

Paul E. Downing,[1]* Yuhong Jiang,[2] Miles Shuman,[2] Nancy Kanwisher[2,3]

Despite extensive evidence for regions of human visual cortex that respond selectively to faces, few studies have considered the cortical representation of the appearance of the rest of the human body. We present a series of functional magnetic resonance imaging (fMRI) studies revealing substantial evidence for a distinct cortical region in humans that responds selectively to images of the human body, as compared with a wide range of control stimuli. This region was found in the lateral occipitotemporal cortex in all subjects tested and apparently reflects a specialized neural system for the visual perception of the human body.

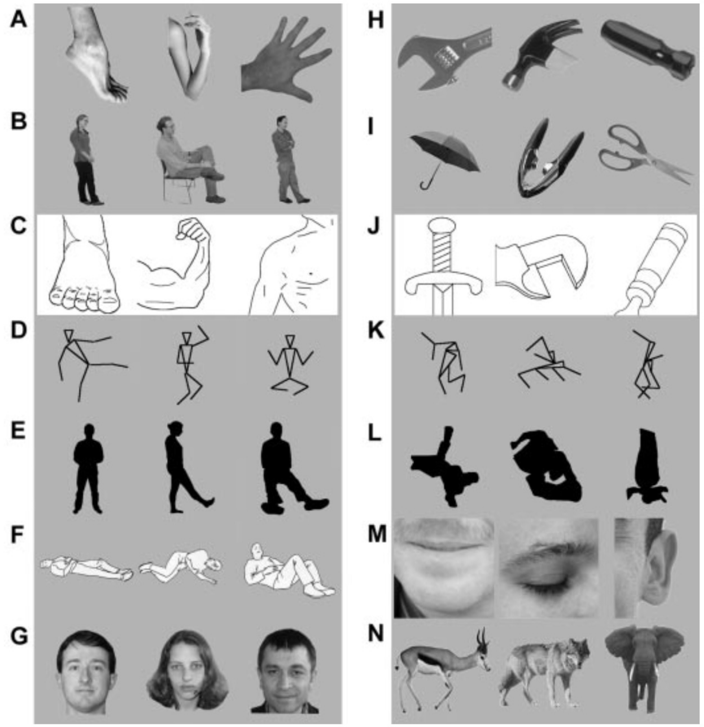

Fig. 2. Stimulus examples. The EBA response was high to human body parts (**A**) and whole human bodies (**B**) whether presented as photographs, line drawings (**C**), stick figures (**D**), or silhouettes (**E**), and was not attenuated to images that depict little implied motion (**F**). The low response to whole faces (**G**) was the single exception found to the preference for human bodies. In contrast, the EBA response was significantly lower to object parts (**H**) and whole articulated objects (**I**), whether represented as photographs or line drawings (**J**), as well as to scrambled control versions of stick figures (**K**) and silhouettes (**L**). The responses to face parts (**M**) and to mammals (**N**) were intermediate.

Functional specificity in the human brain: A window into the functional architecture of the mind

Nancy Kanwisher[1]

McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA 02139

This contribution is part of the special series of Inaugural Articles by members of the National Academy of Sciences elected in 2005.

Contributed by Nancy Kanwisher, April 16, 2010 (sent for review February 22, 2010)

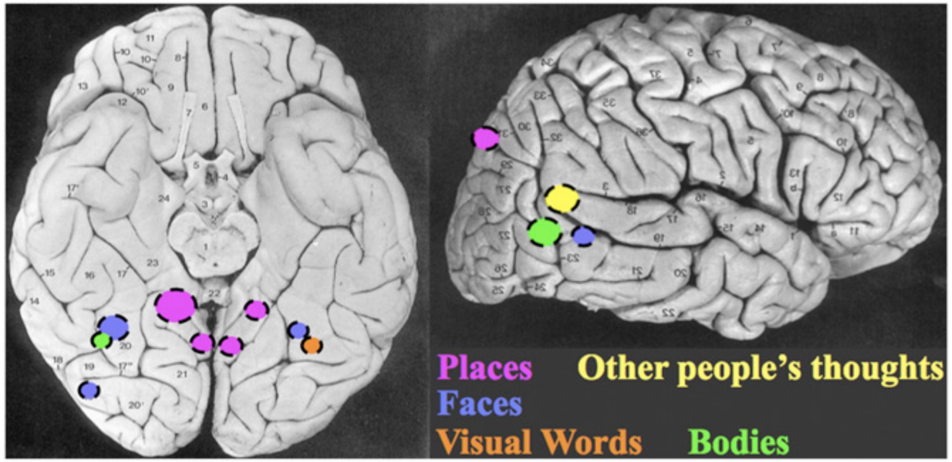

Fig. 1. This schematic diagram indicates the approximate size and location of regions in the human brain that are engaged specifically during perception of faces (blue), places (pink), bodies (green), and visually presented words (orange), as well as a region that is selectively engaged when thinking about another person's thoughts (yellow). Each of these regions can be found in a short functional scan in essentially all normal subjects.

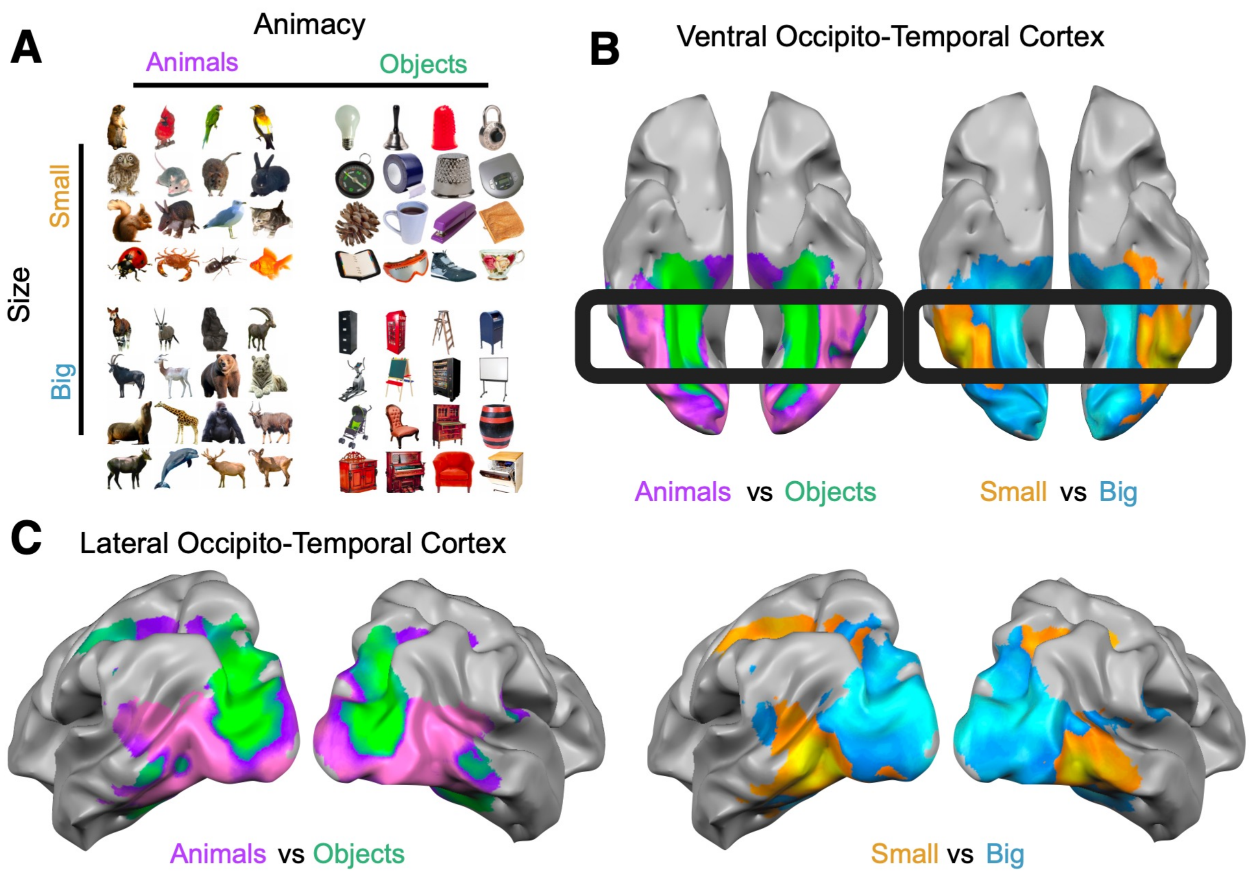

Tripartite Organization of the Ventral Stream by Animacy and Object Size

Talia Konkle[1,2] **and Alfonso Caramazza**[1,2]
[1]Department of Psychology, Harvard University, Cambridge, Massachusetts 02138, and [2]Center for Mind/Brain Sciences (CiMeC), University of Trento, 38068 Rovereto (TN), Italy

The first deep convolutional network: Neocognitron

Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position

1980

Kunihiko Fukushima
NHK Broadcasting Science Research Laboratories, Kinuta, Setagaya, Tokyo, Japan

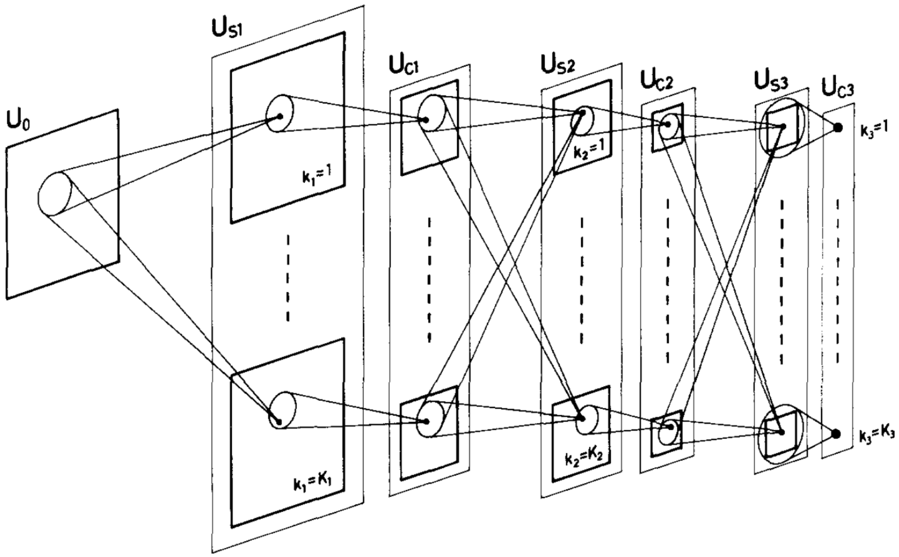

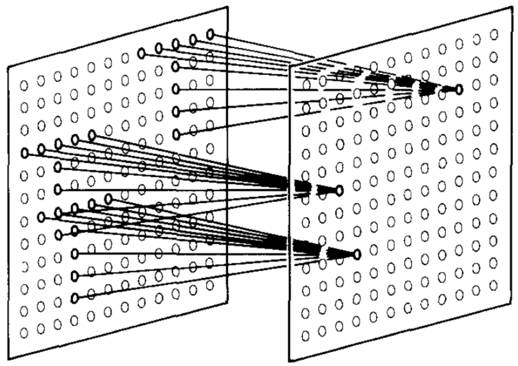

Fig. 3. Illustration showing the input interconnections to the cells within a single cell-plane

The first deep convolutional network: Neocognitron

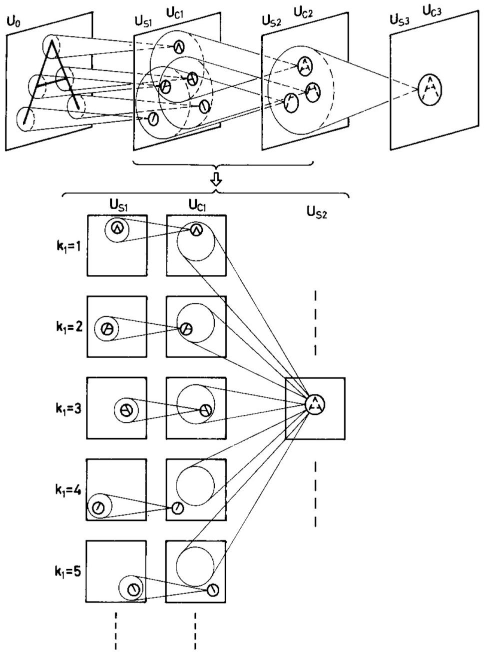

Fig. 5. An example of the interconnections between cells and the response of the cells after completion of self-organization

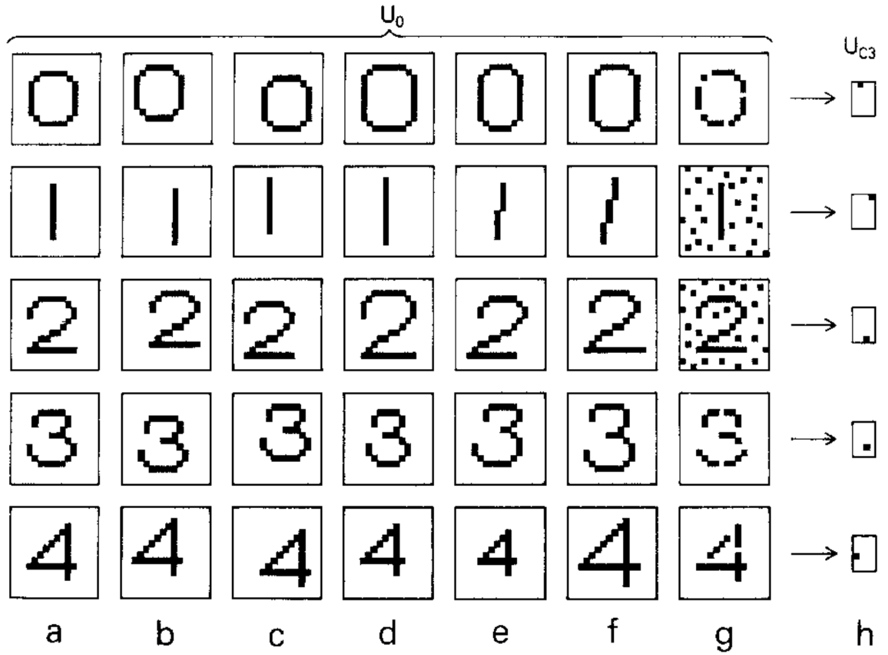

Fig. 6. Some examples of distorted stimulus patterns which the neocognitron has correctly recognized, and the response of the final layer of the network

Backpropagation

Handwritten Digit Recognition with a Back-Propagation Network

1990

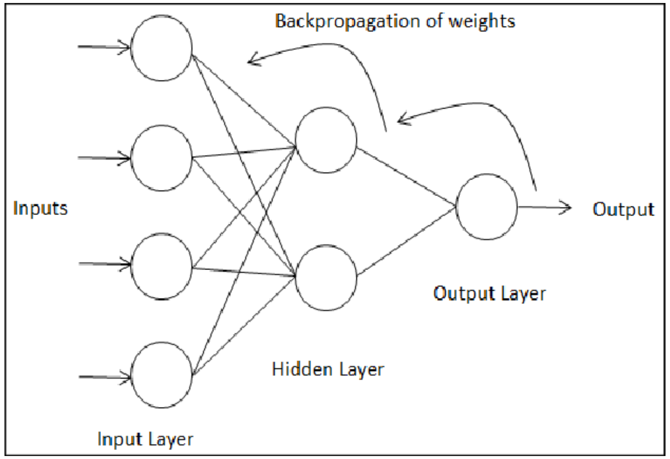

The network output is given by:

$$\widehat{y} = net(x)$$

where net() is composed of layers defined by:

$$a_i^l = ReLU\left[\sum_j w_{ij} a_j^{l-1}\right]$$

We define the loss function:

$$\mathcal{L} = ||\widehat{y} - y||_2$$

and the backpropagation update rule:

$$w_{ij}^{t+1} = w_{ij}^t - \epsilon * \frac{d\mathcal{L}}{dw_{ij}^t}$$

Success of a ConvNet on MNIST

Handwritten Digit Recognition with a Back-Propagation Network

Figure 2: Examples of normalized digits from the testing set.

Success of a ConvNet on ImageNet

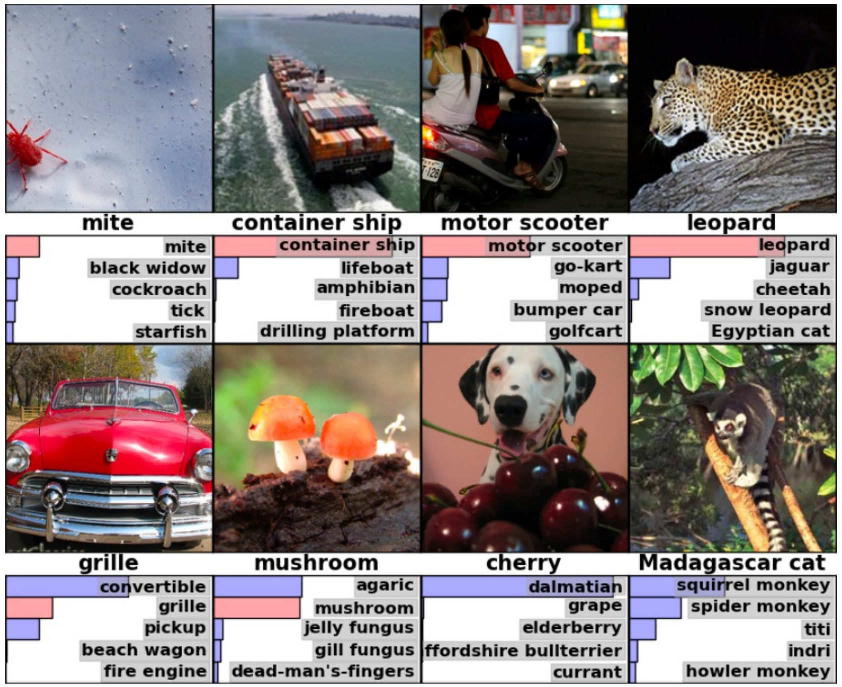

ImageNet Classification with Deep Convolutional Neural Networks

2012

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

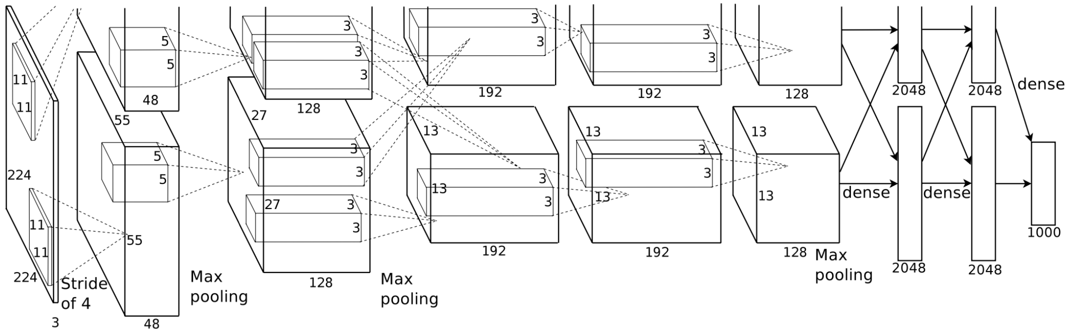

Figure 2: An illustration of the architecture of our CNN, explicitly showing the delineation of responsibilities between the two GPUs. One GPU runs the layer-parts at the top of the figure while the other runs the layer-parts at the bottom. The GPUs communicate only at certain layers. The network's input is 150,528-dimensional, and the number of neurons in the network's remaining layers is given by 253,440–186,624–64,896–64,896–43,264–4096–4096–1000.

The success of deep learning for image recognition

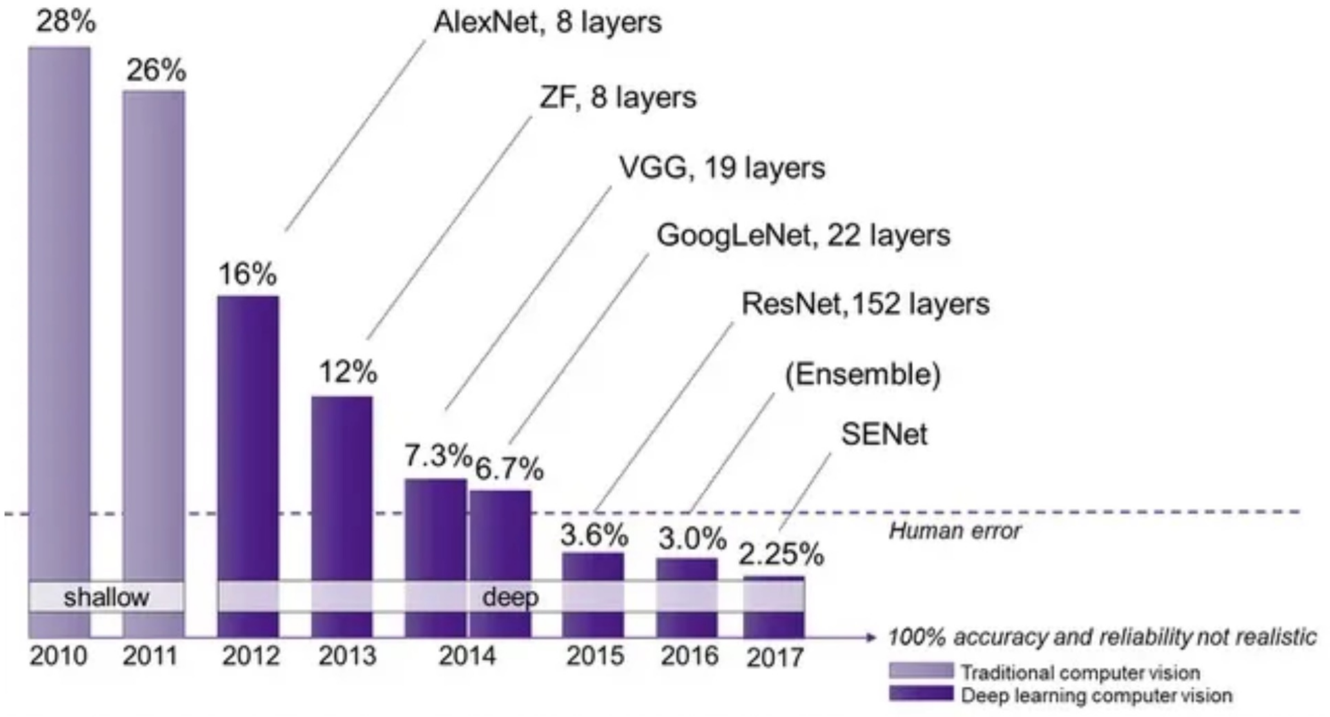

Comparison between the representations learned by a deep network and the brain

Performance-optimized hierarchical models predict neural responses in higher visual cortex

Daniel L. K. Yamins[a,1], Ha Hong[a,b,1], Charles F. Cadieu[a], Ethan A. Solomon[a], Darren Seibert[a], and James J. DiCarlo[a,2]

2014

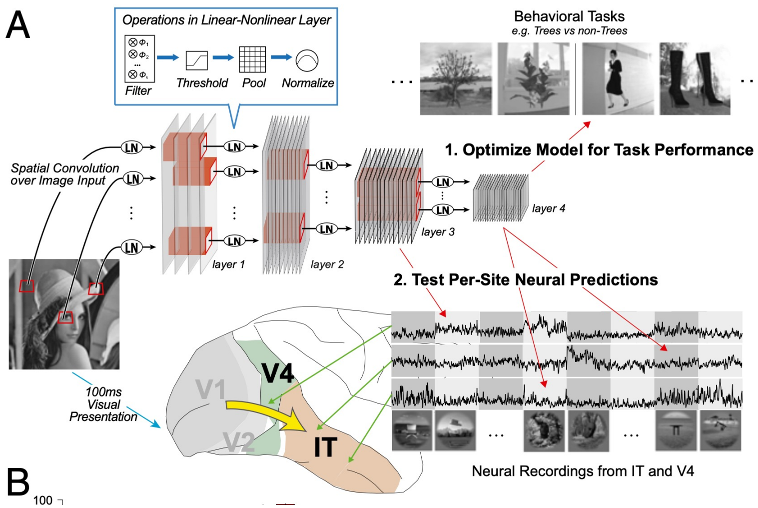

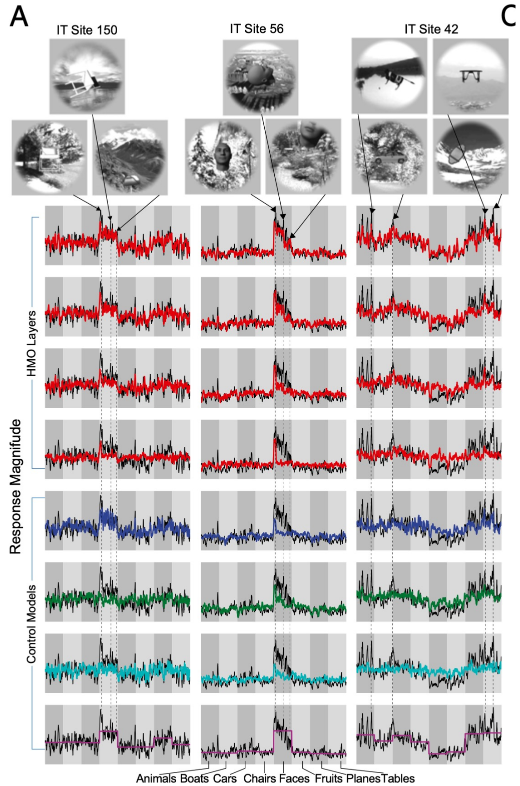

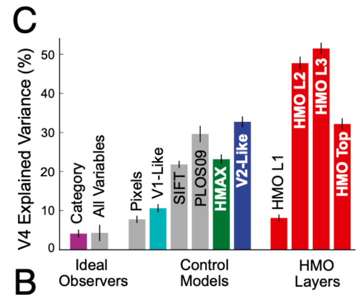

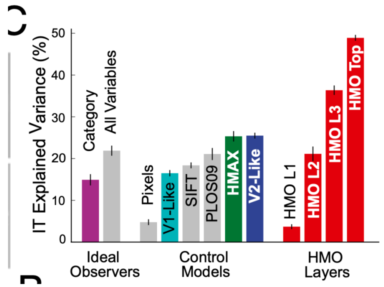

Correspondance between receptive fields in early layers

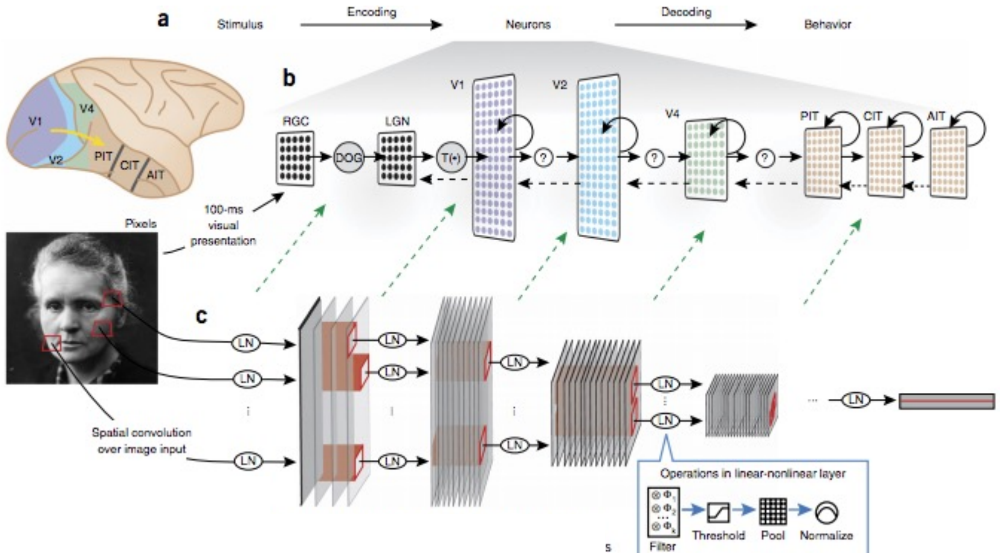

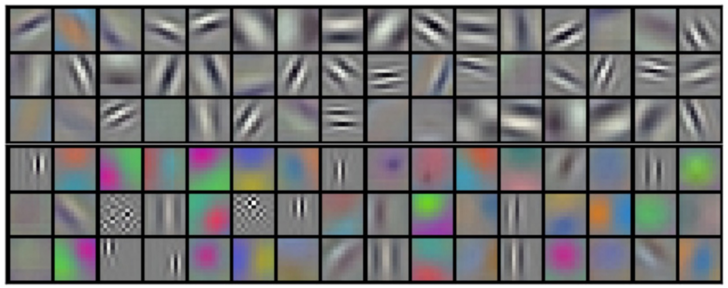

CNNs learn representations similar to those in cortex

Yamins & DiCarlo, 2016

But they skip past the retinal center-surround stage

Correspondance between receptive fields in early layers

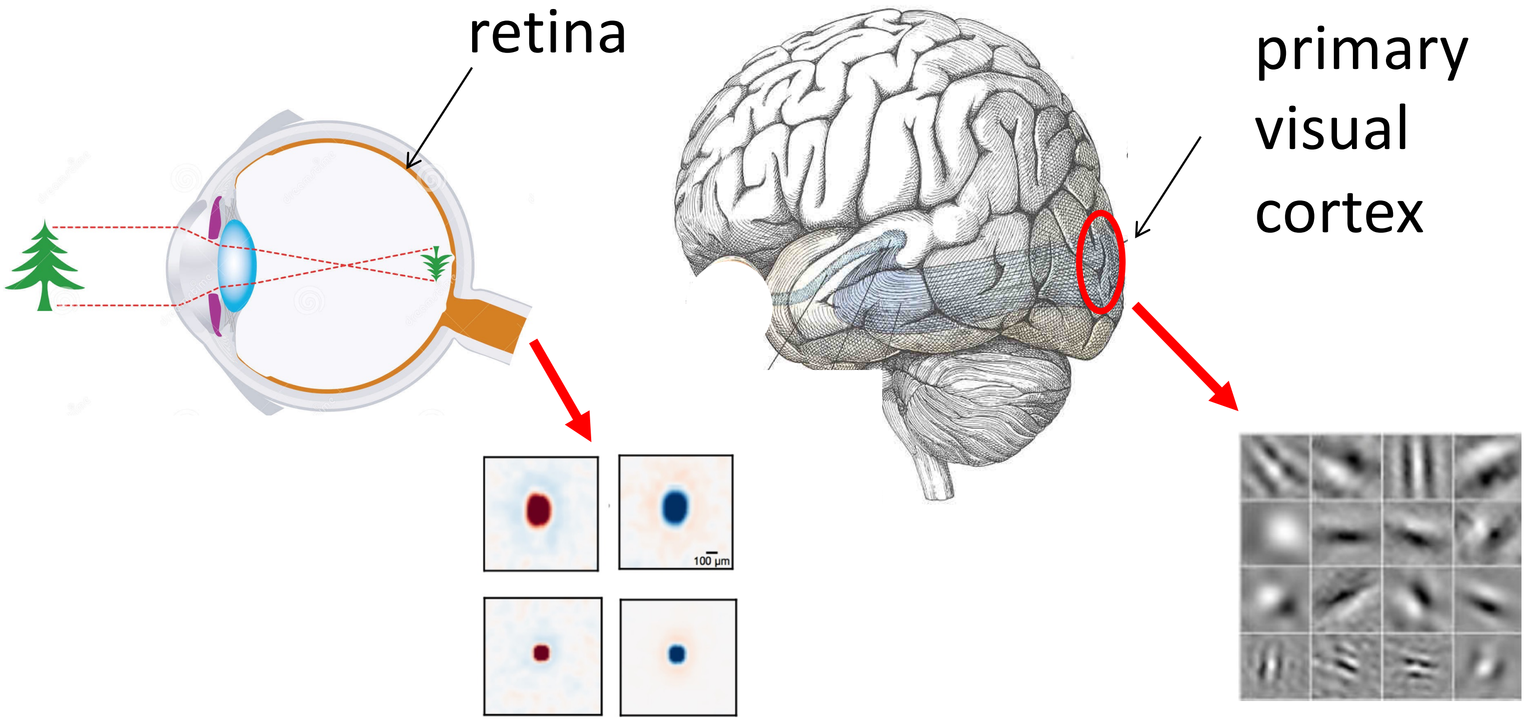

Kuffler, 1953

Hubel & Wiesel, 1960

Receptive fields are concentric in the retina and oriented in cortex

Towards a unified model from retina to cortex

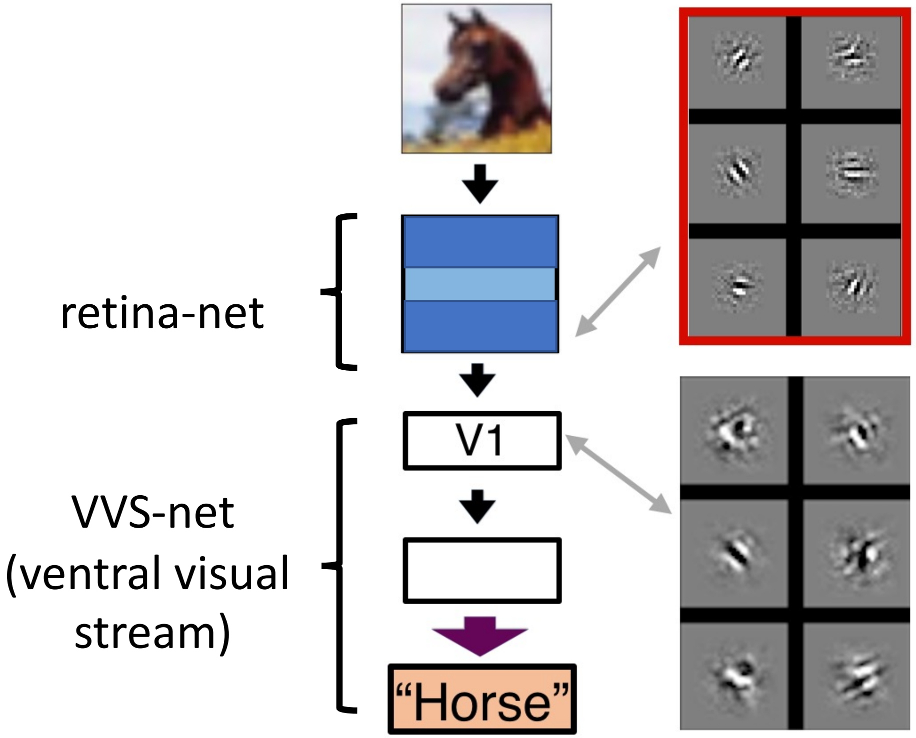

Mismatch with the retina

Lindsey, Ocko*, Ganguli, Deny, ICLR 2019*

Anatomical constraints on the visual system

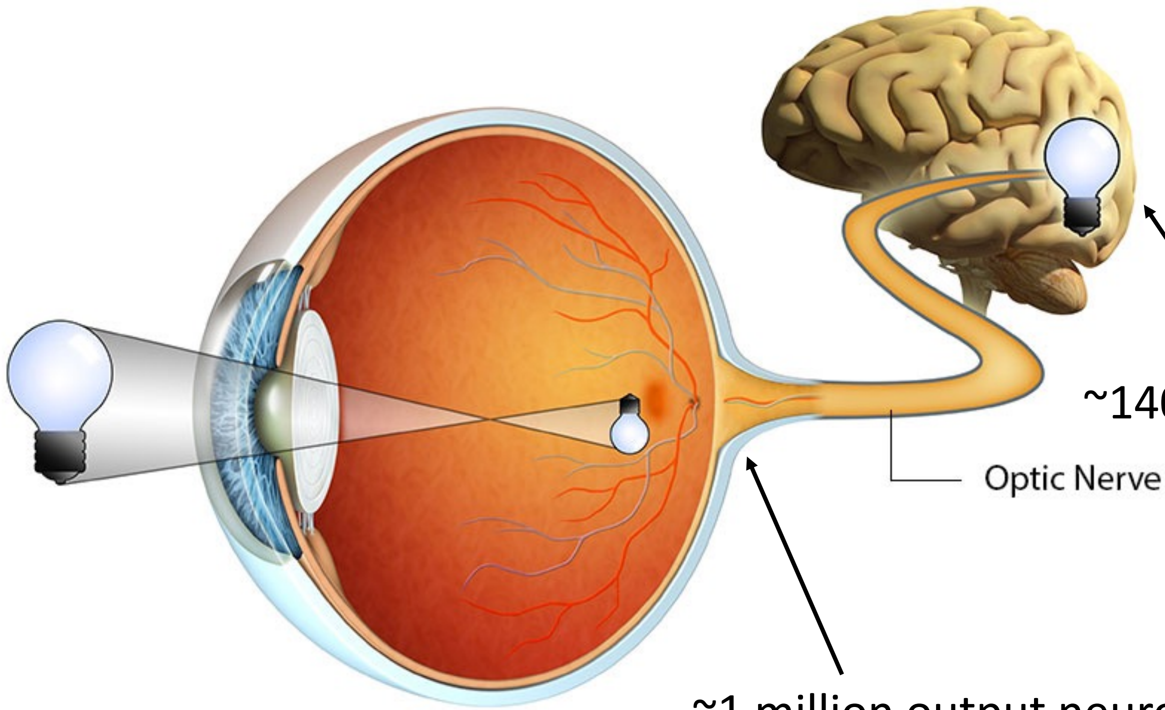

~140 million V1 neurons

~1 million output neurons
(retinal ganglion cells)

Lindsey, Ocko*, Ganguli, Deny, ICLR 2019*

An anatomically constrained deep convolutional model of the visual system

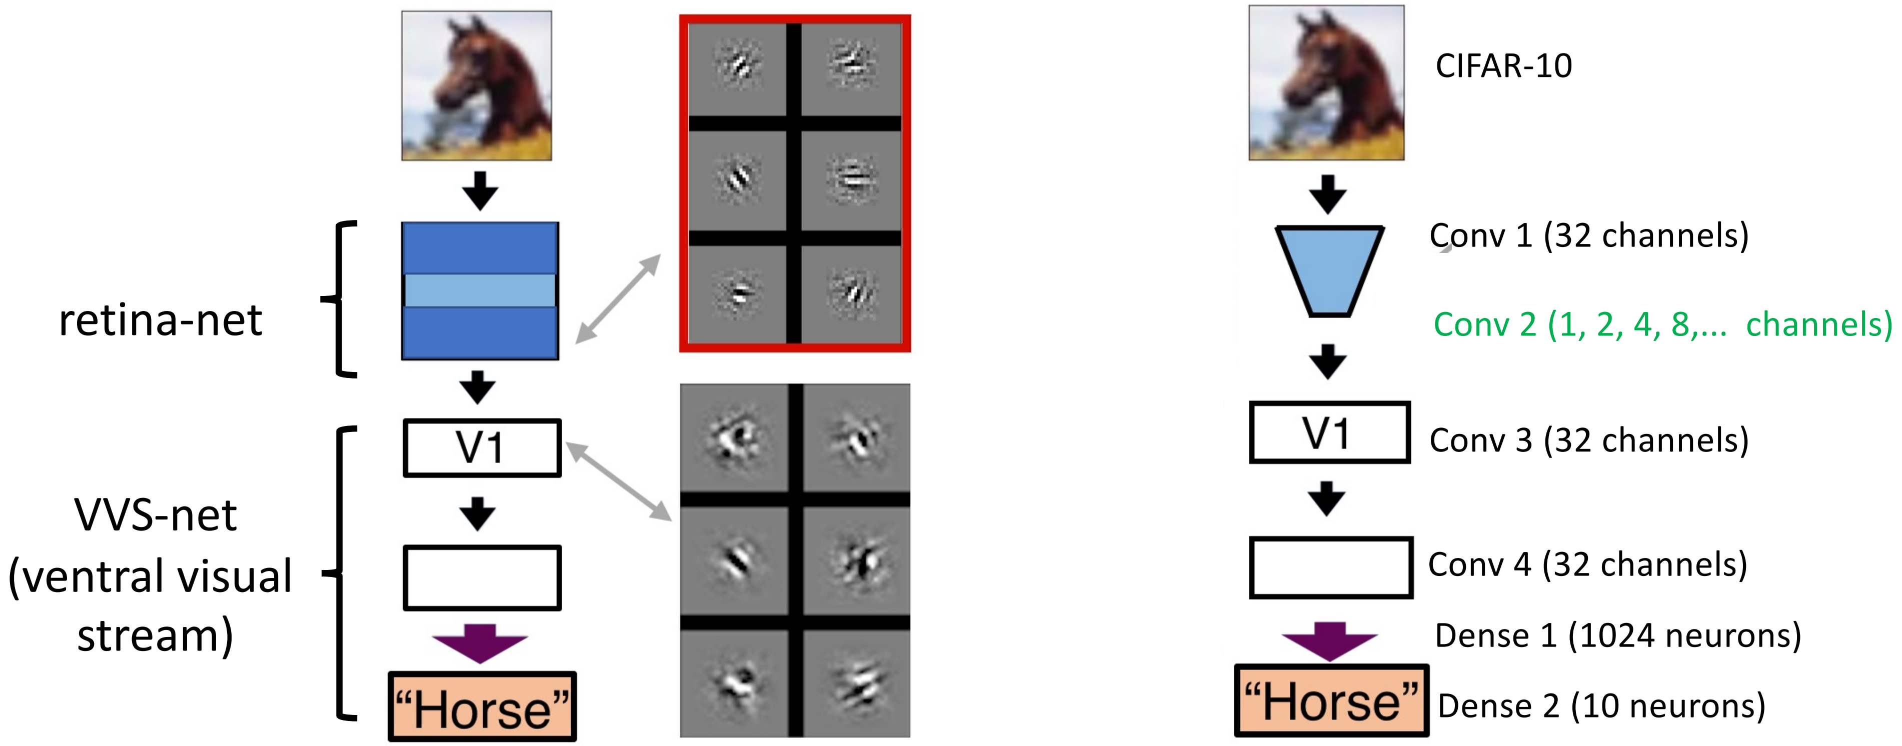

Lindsey, Ocko*, Ganguli, Deny, ICLR 2019*

An anatomically constrained deep convolutional model of the visual system

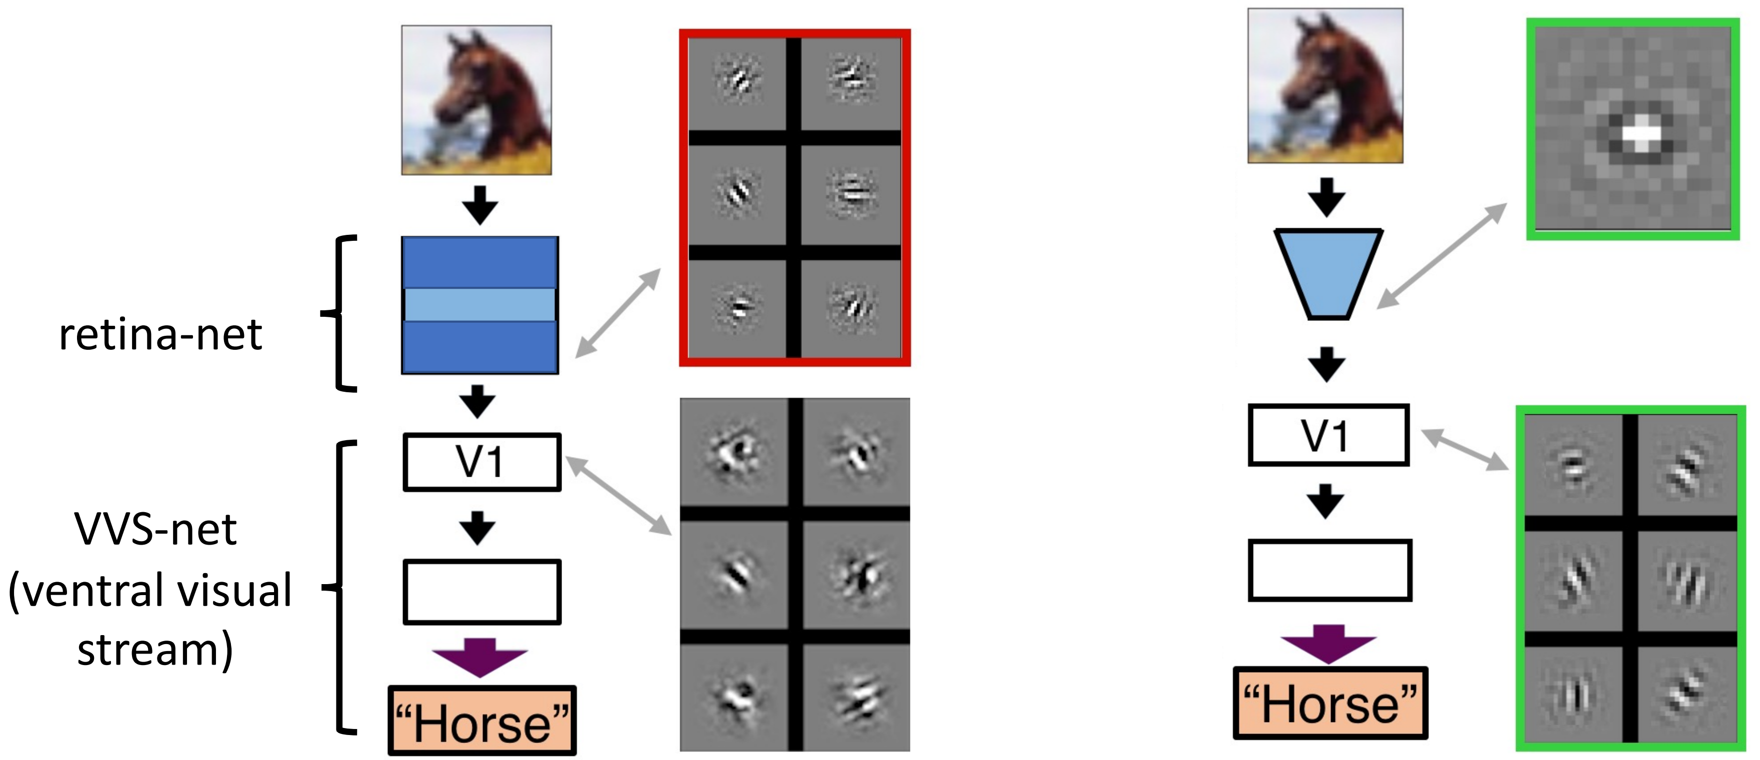

Lindsey, Ocko*, Ganguli, Deny, ICLR 2019*

Role of Recurrent connections in the brain

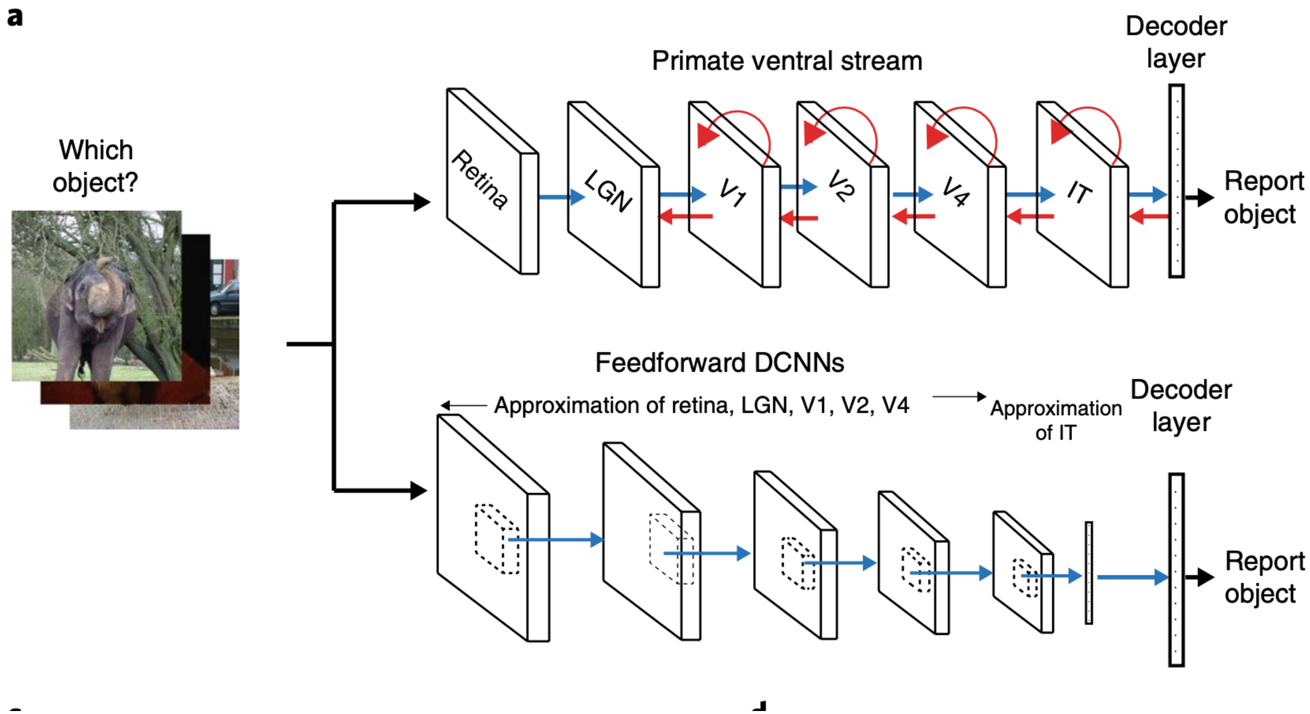

Kar et al 2019

Role of recurrent connections in the brain

nature neuroscience
ARTICLES
https://doi.org/10.1038/s41593-019-0392-5

Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior

Kohitij Kar[1,2]*, Jonas Kubilius[1,3], Kailyn Schmidt[1], Elias B. Issa[1,4] and James J. DiCarlo[1,2]

Non-recurrent deep convolutional neural networks (CNNs) are currently the best at modeling core object recognition, a behavior that is supported by the densely recurrent primate ventral stream, culminating in the inferior temporal (IT) cortex. If recurrence is critical to this behavior, then primates should outperform feedforward-only deep CNNs for images that require additional recurrent processing beyond the feedforward IT response. Here we first used behavioral methods to discover hundreds of these 'challenge' images. Second, using large-scale electrophysiology, we observed that behaviorally sufficient object identity solutions emerged ~30 ms later in the IT cortex for challenge images compared with primate performance-matched 'control' images. Third, these behaviorally critical late-phase IT response patterns were poorly predicted by feedforward deep CNN activations. Notably, very-deep CNNs and shallower recurrent CNNs better predicted these late IT responses, suggesting that there is a functional equivalence between additional nonlinear transformations and recurrence. Beyond arguing that recurrent circuits are critical for rapid object identification, our results provide strong constraints for future recurrent model development.

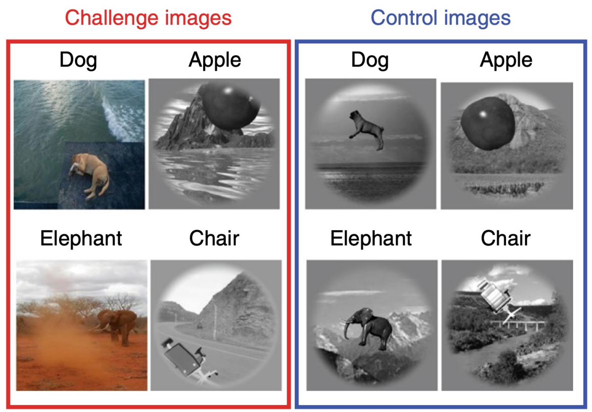

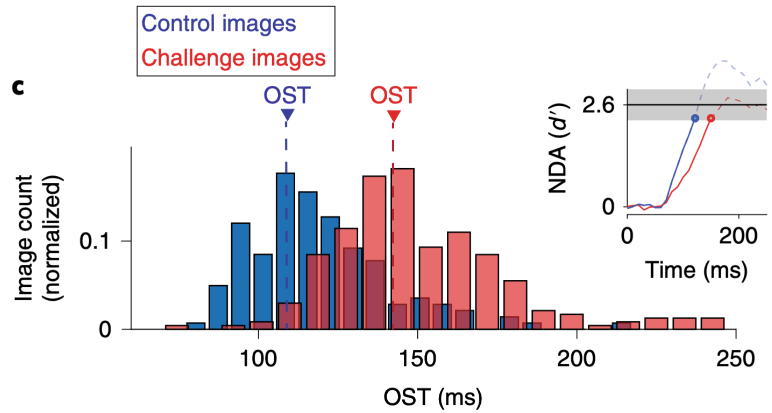

Interplay between memory and vision

Article

When the ventral visual stream is not enough: A deep learning account of medial temporal lobe involvement in perception

Tyler Bonnen,[1,4,*] Daniel L.K. Yamins,[1,2,3] and Anthony D. Wagner[1,3]
[1]Department of Psychology, Stanford University, Stanford, CA, USA
[2]Department of Computer Science, Stanford University, Stanford, CA, USA
[3]Wu Tsai Neurosciences Institute, Stanford University, Stanford, CA, USA
[4]Lead contact
*Correspondence: bonnen@stanford.edu
https://doi.org/10.1016/j.neuron.2021.06.018

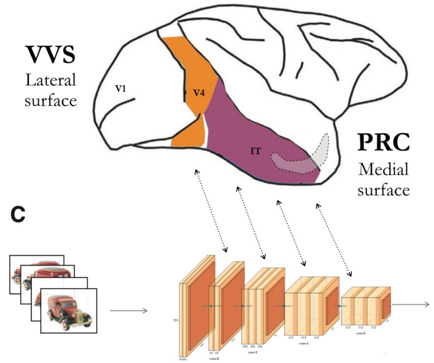

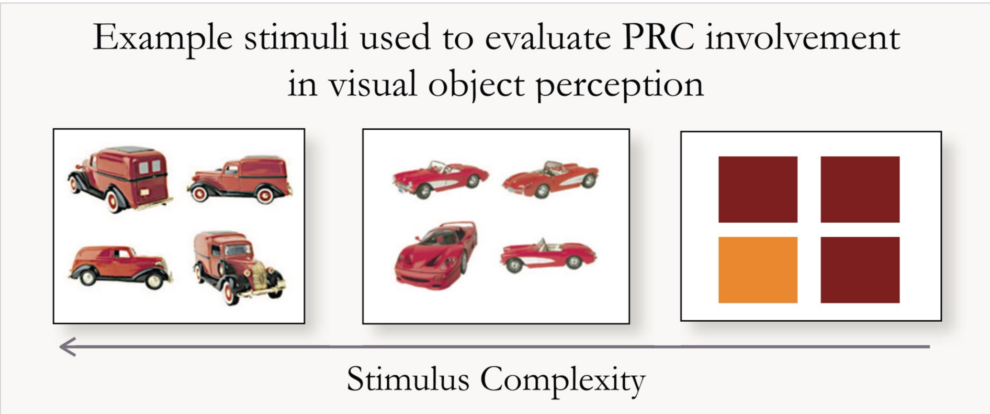

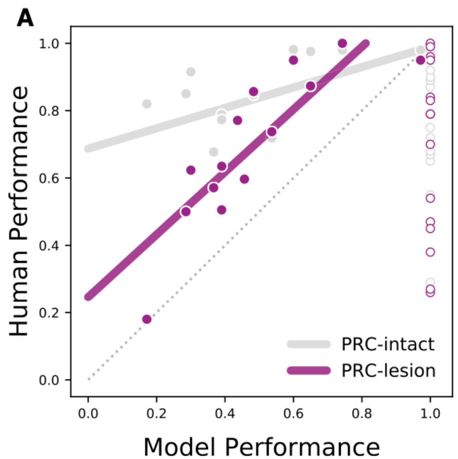

Interplay between memory and vision

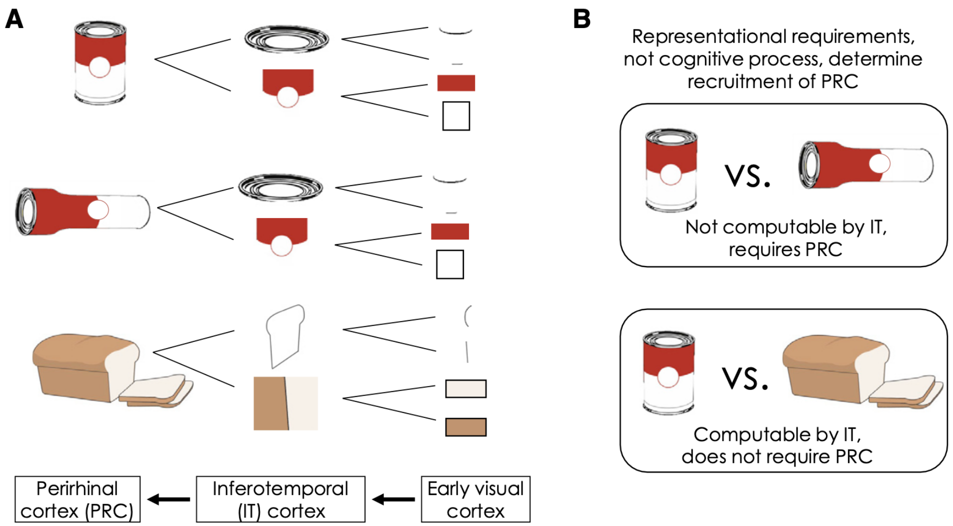

Figure 1. Representational-hierarchical account of cognition
(A) Throughout the VVS and extending into the PRC, the features of an object are represented with increasing complexity. The PRC is at the apex, containing complex representations of objects.
(B) According to this theory, the same cognitive *process* (in this example, perceiving whether two objects are the same or different) is solved by different regions in the hierarchy. Thus, it is *representational content*, rather than cognitive process, that determines the division of labor throughout the pathway. Illustrations by Alexander Jacob.

Next to understand : Mental rotation in the brain!

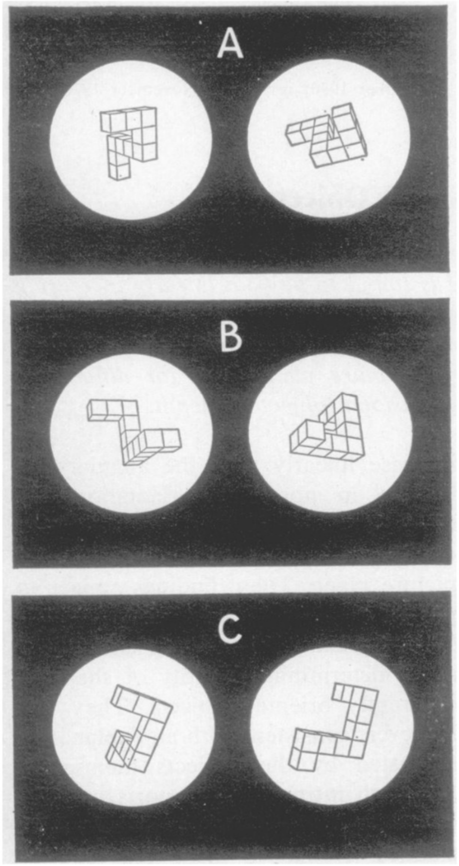

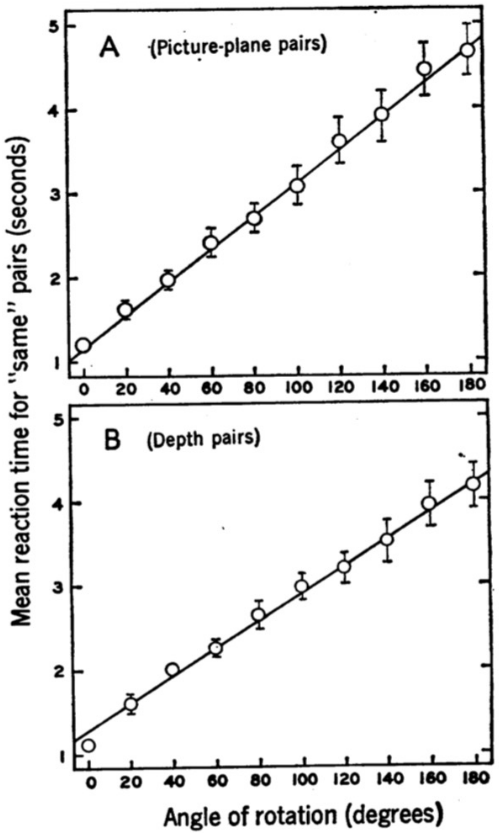

Kids *learn* how to do it

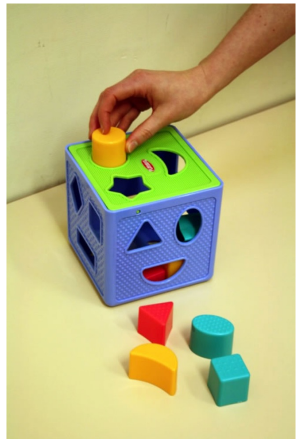

Mental Rotation of Three-Dimensional Objects

Shepard & Metzler, Science, 1971

Mental rotation in the brain

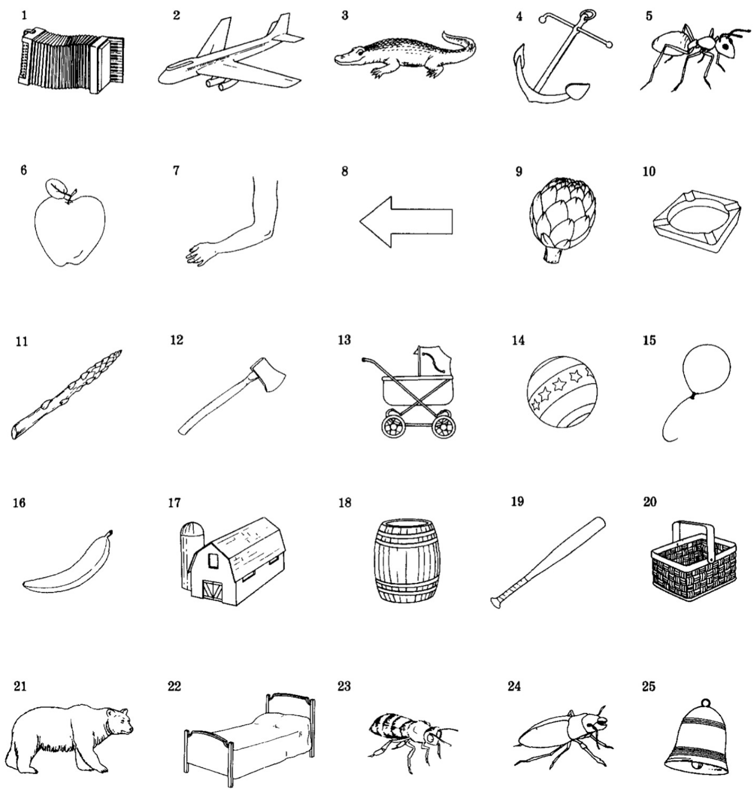

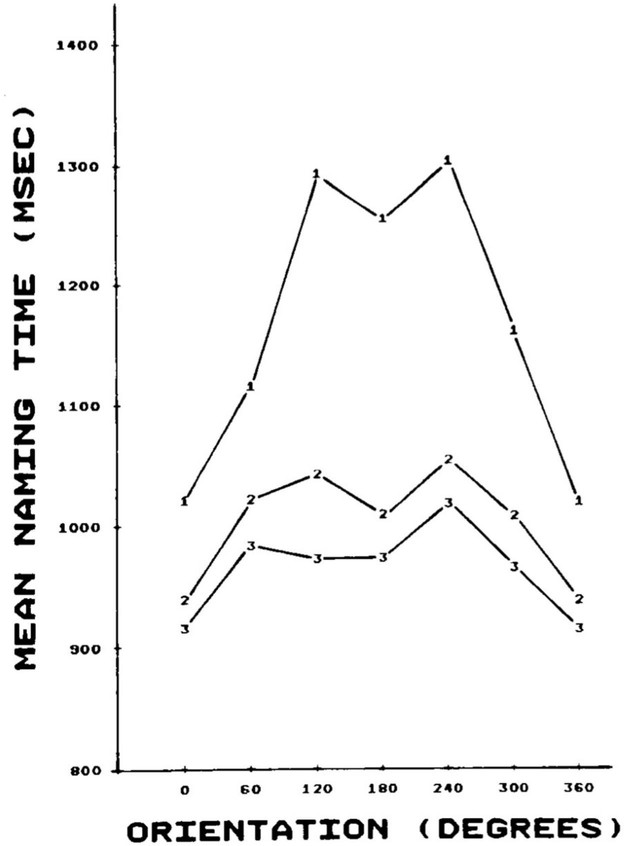

Jolicoeur 1985

Supplementary slides

What about the dorsal stream?

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

Shahab Bakhtiari
Mila & McGill University
bakhtias@mila.quebec

Patrick Mineault
patrick.mineault@gmail.com

Tim Lillicrap
DeepMind
timothylillicrap@google.com

Christopher C. Pack
McGill University
christopher.pack@mcgill.ca

Blake A. Richards
CIFAR, Mila & McGill University
blake.richards@mila.quebec

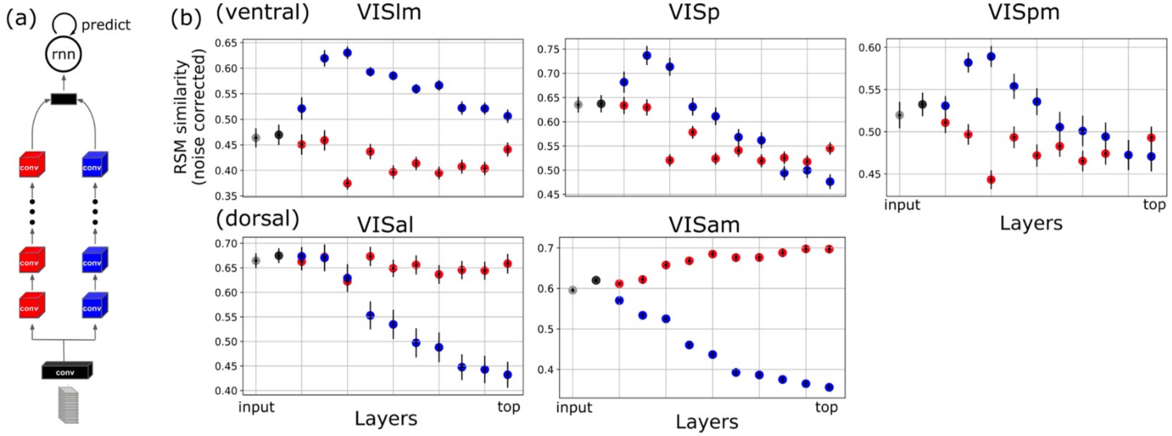

Figure 3: Representational Similarity Analysis between all the visual areas and the ANN trained with CPC. (a) The schematic of the ANN architecture with two pathways (ResNet-2p) used as the backbone of CPC. (b) Representational similarity between all the layers of the ANN with two pathways (trained with CPC) and the ventral (top: VISlm, VISp, VISpm) and the dorsal (bottom: VISal, VISam) areas

Mental rotation in the brain

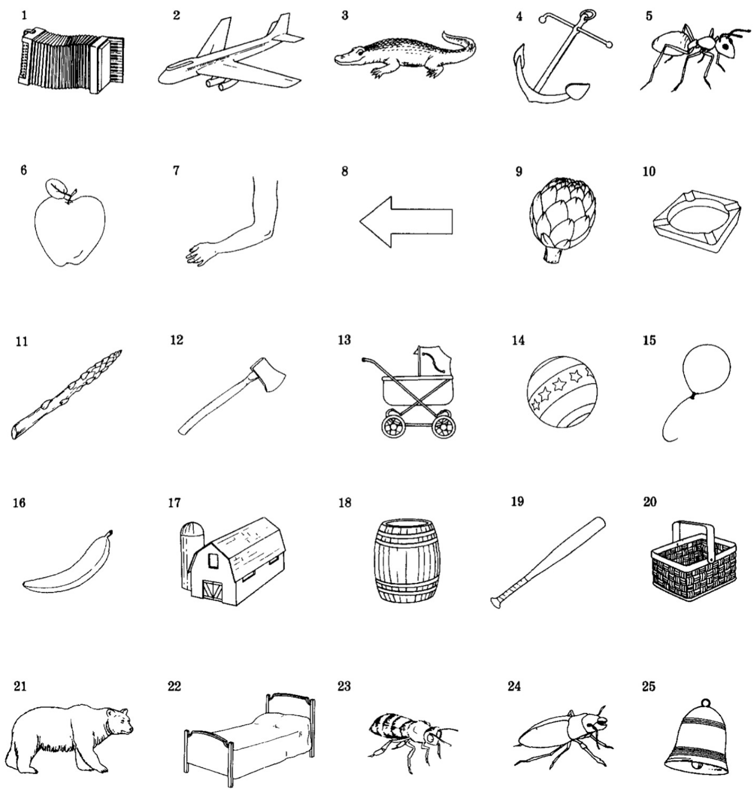

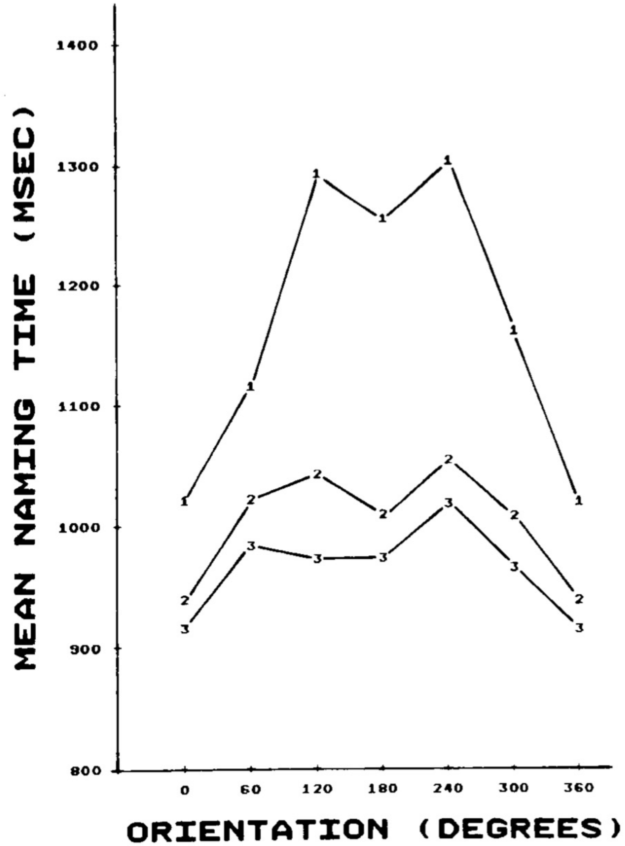

Jolicoeur 1985

Recurrent connections in the brain might play a key role in mental rotation and other reasoning tasks

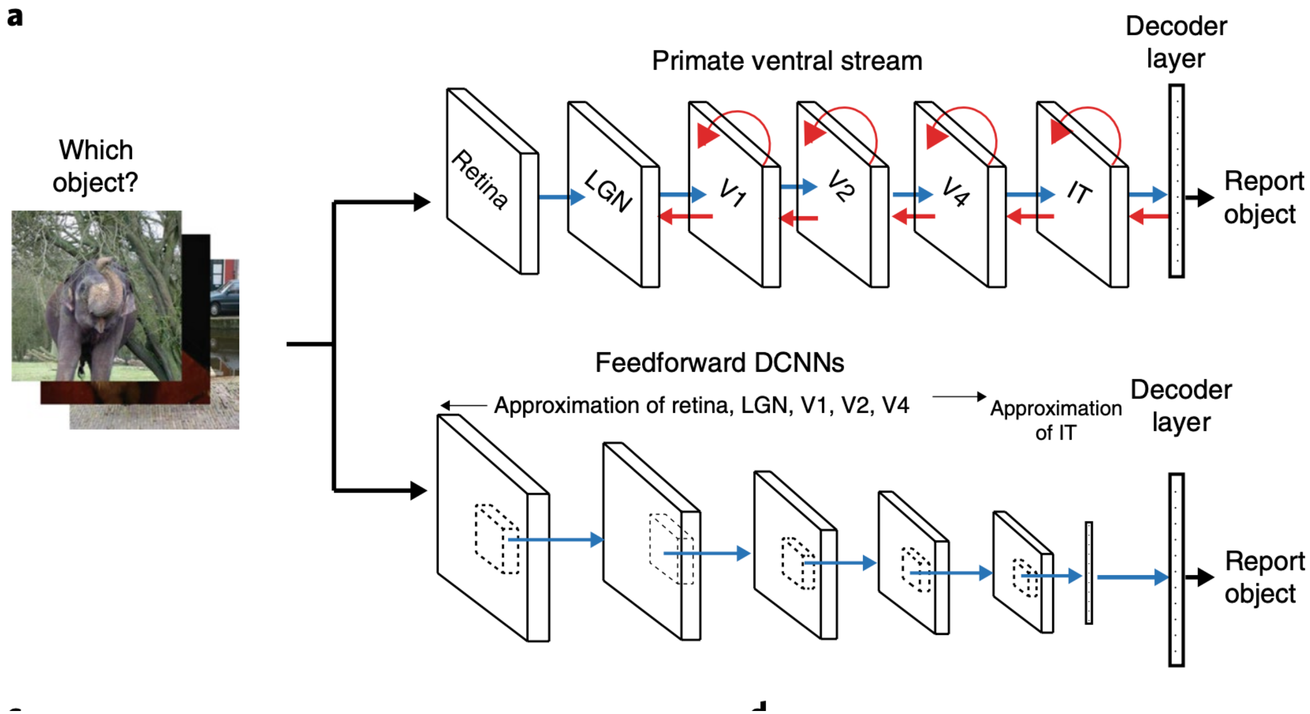

Kar et al 2019

Deep learning model of mental rotation

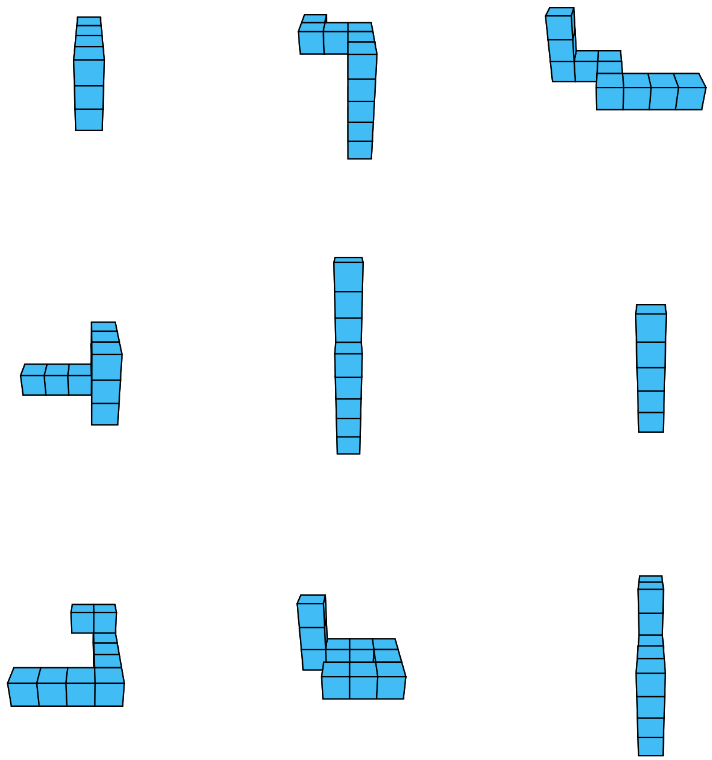

- includes 338 objects
- each object consists of 10 cubes
- each object features three right-angled elbows and a possible loop in just a few cases

Aleksandr Krylov

Deep learning model of mental rotation

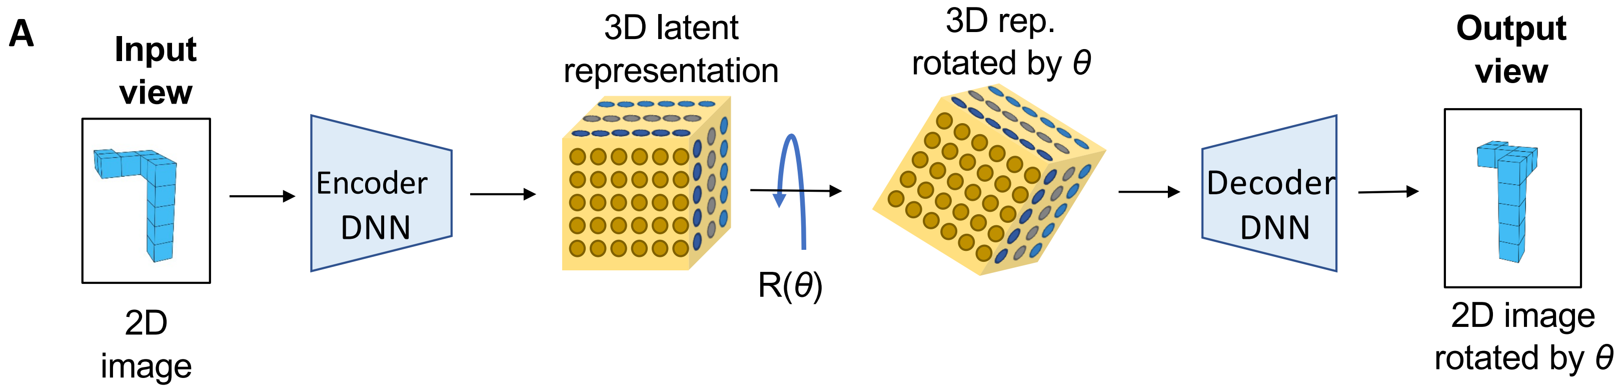

Aleksandr Krylov

Deep learning model of mental rotation

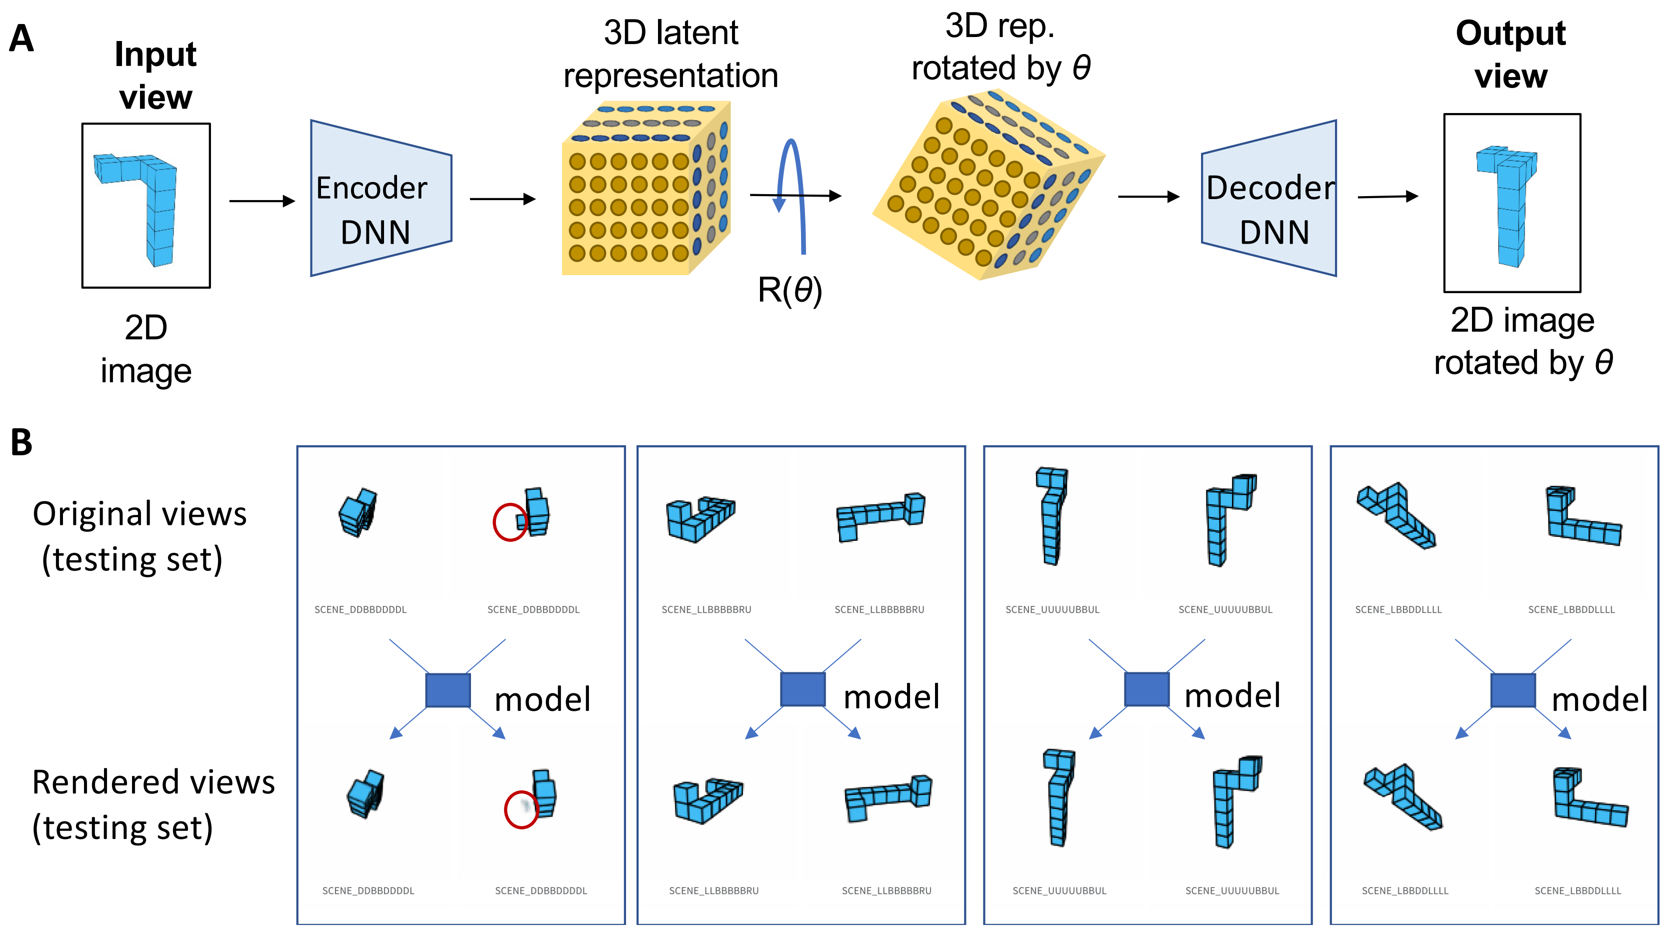

Minh Dinh

Deep learning model of mental rotation

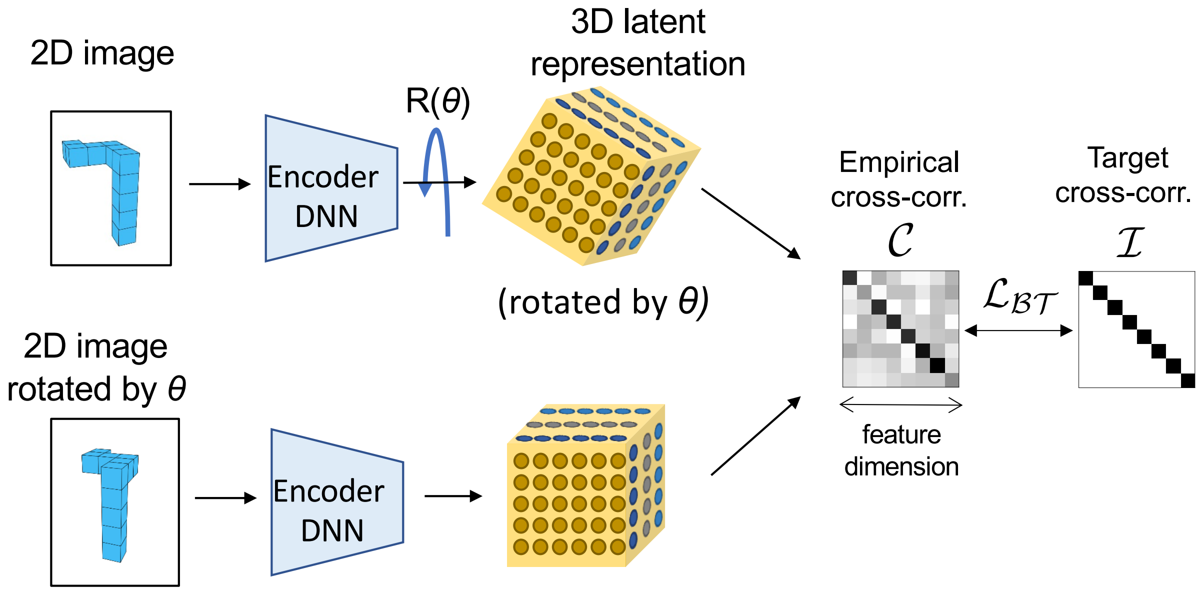

Minh Dinh

Deep learning model of mental rotation

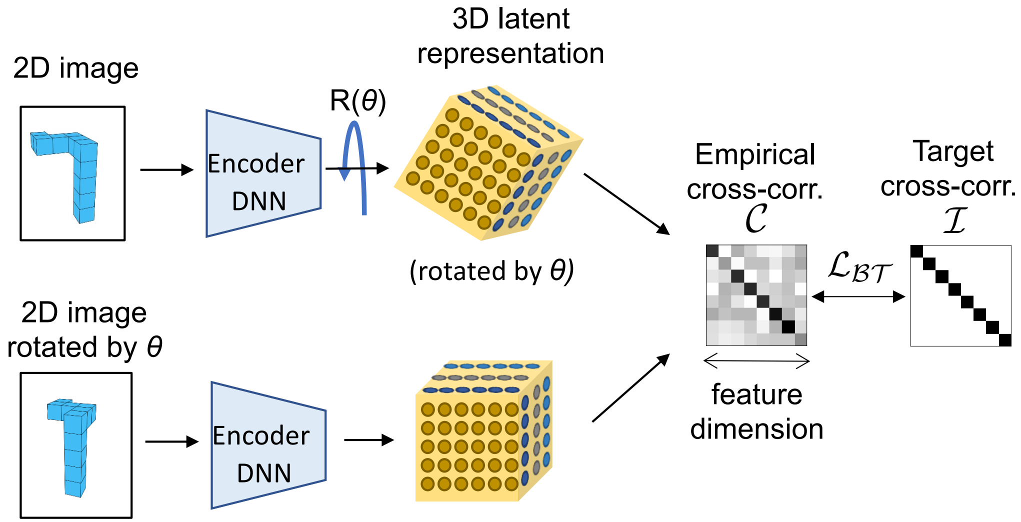

- $\|Y - Y'\|$: embedding of the same object should be *very* **close**
- $\|Yi - Yj\|$: embedding of the same object should be *very* **far**
- $\|Y - \underline{Y}\|$: embedding of the wrong rotation should be **far**
- $\|Y - y\|$ **:** embedding of the wrong representation should be **far**

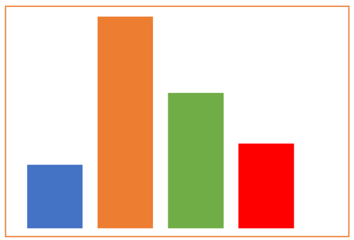

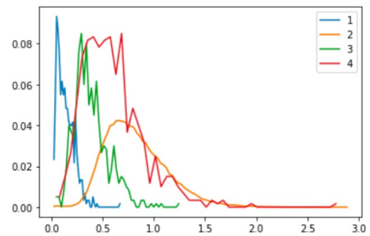

Minh Dinh

Deep learning model of mental rotation

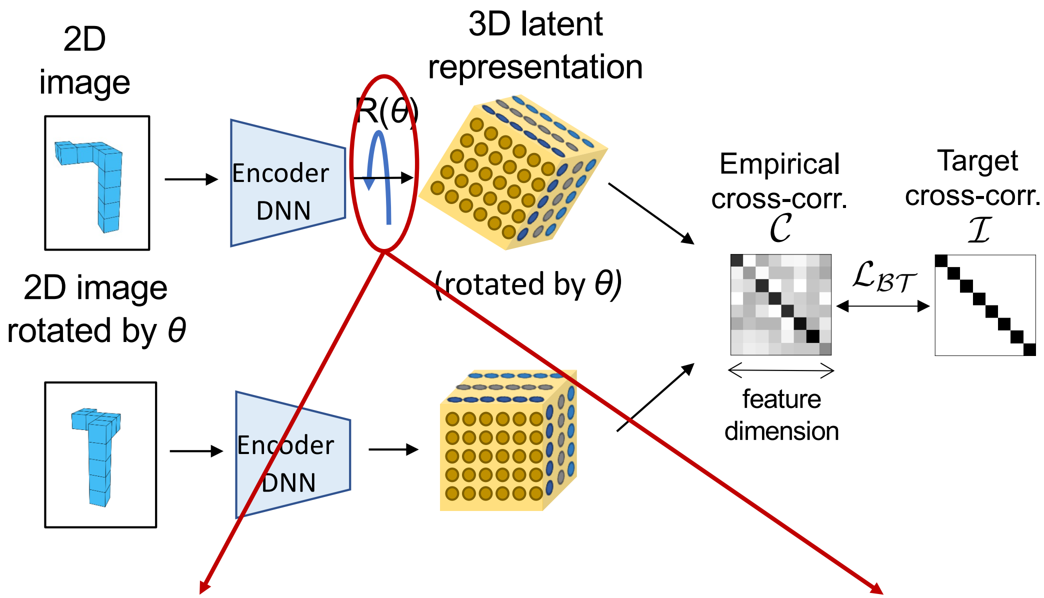

- $\|Y - Y'\|$: embedding of the same object should be *very* **close**
- $\|Yi - Yj\|$: embedding of the same object should be *very* **far**
- $\|Y - \underline{Y}\|$: embedding of the wrong rotation should be **far**
- $\|Y - y\|$: embedding of the wrong representation should be **far**

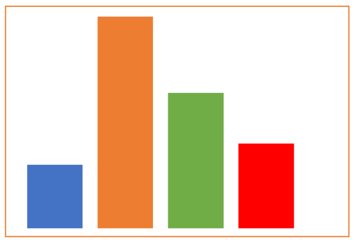

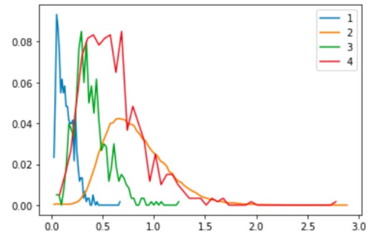

hyp1: progressive rotation with elementary steps until the upright pose is found

hyp2: trial-and-error rotation steps until the upright pose is found

Riina Pöllänen

Invariant representations in the brain

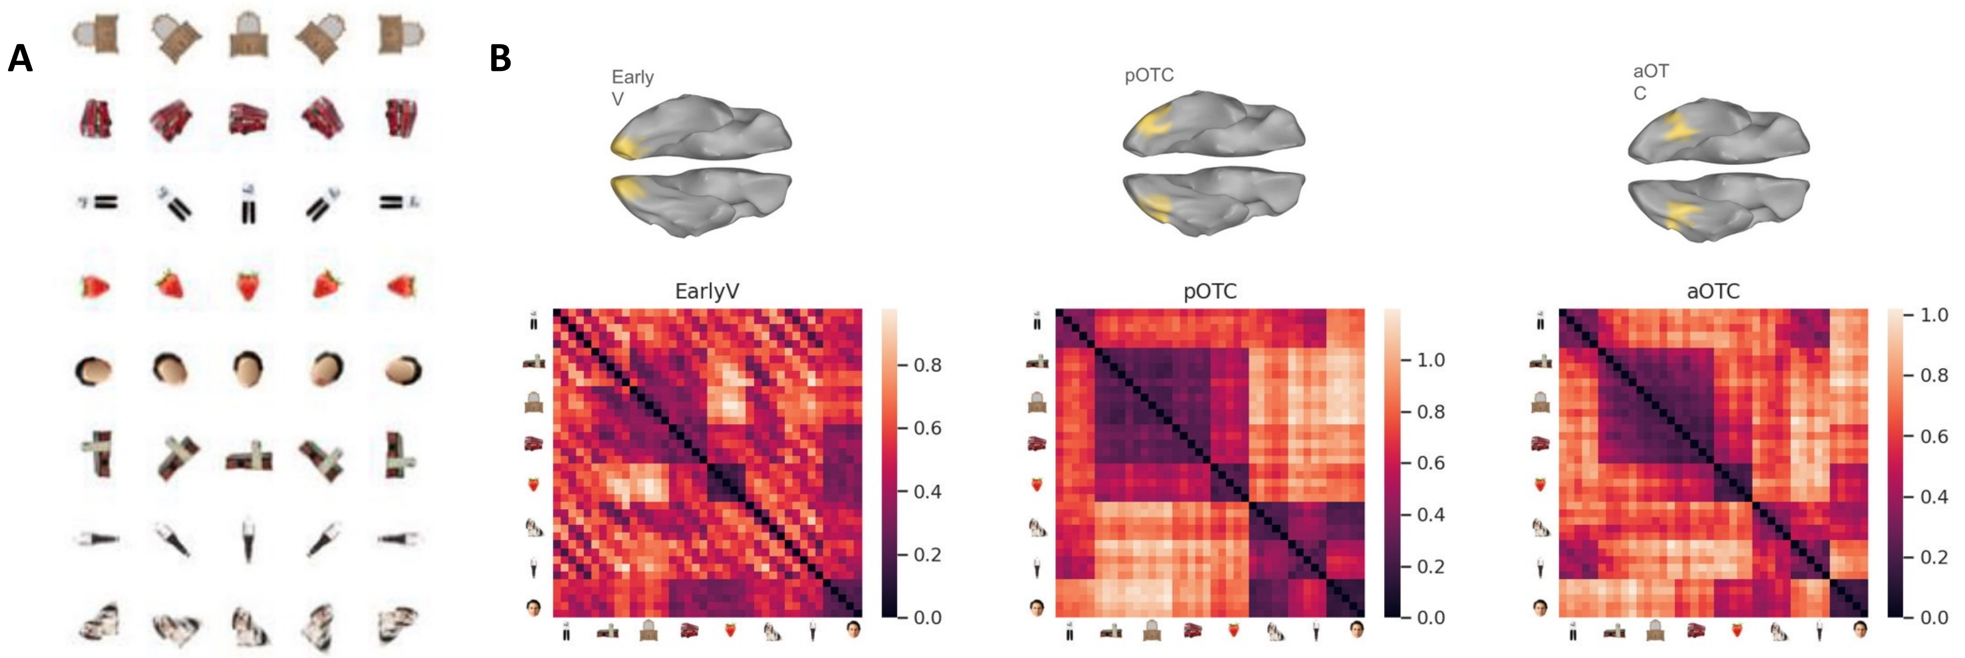

Matias Koponen

Metzler exp. re-analysis in favor of trial-and-error hypothesis

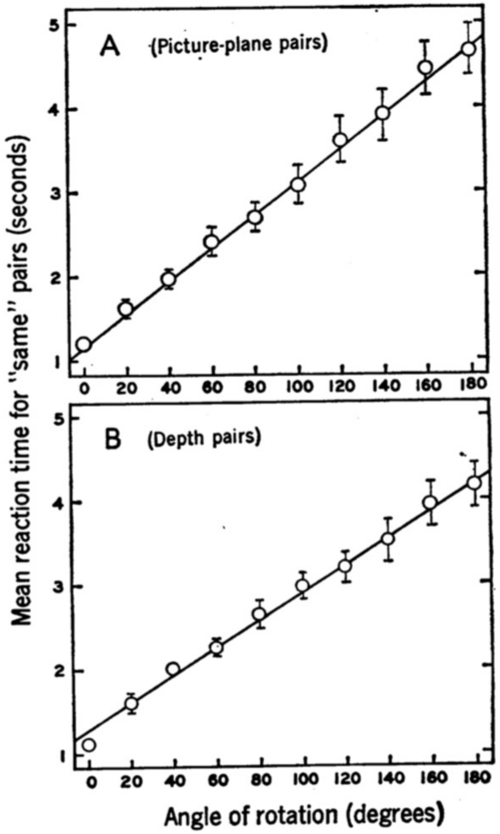

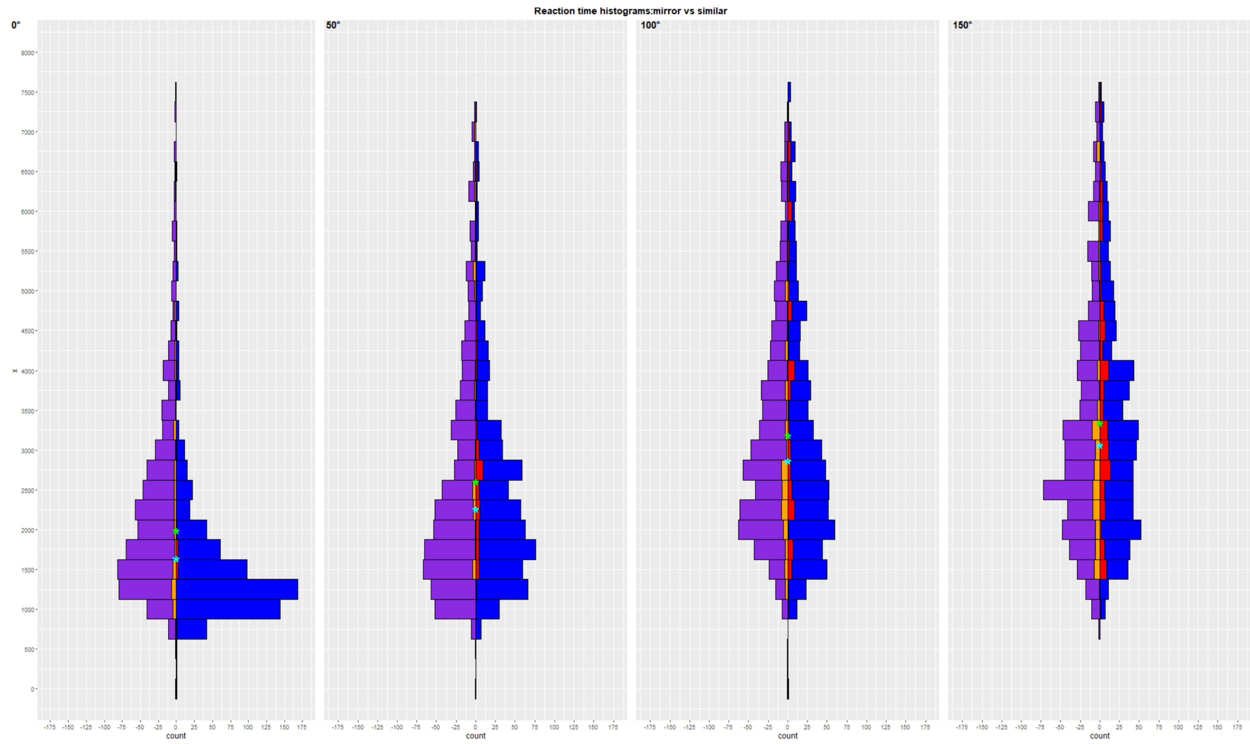

Experiements

- MEG to see how the invariance to rotated objects unfolds in time
- Does it unfold progressively with angle of rotation (hyp1)?
- Or does it suddenly jump to the correct invariant representation?

The visual system organization

- ventral vs dorsal stream
- what do they do?
- how were they historically discovered?

Crash-course about the brain: neurons

- They are cells containing dendrites and an axon
- They communicate via synapses
- synapses can be modified via experience: Hebb learning rule (STDP)
- They send precisely timed signals called ‘action potentials’ or spikes
- They can be approximated by a sum or their input and spike when a thershold is reached (leaky-integrate and fire model)

The visual system: organization of the ventral stream

- retina, thalamus, V1

- simple cells and complex cells in V1

- hypercolumns for orientation, occular dominance, color,

- different patches corresponding to different

The brain in numbers

- 86 Billion neurons (80% in the cerebellum, 18% in the cortex)
- 100 to 10,000 connections on average per neuron
- 20 Watts consumption (i.e., a light bulb)

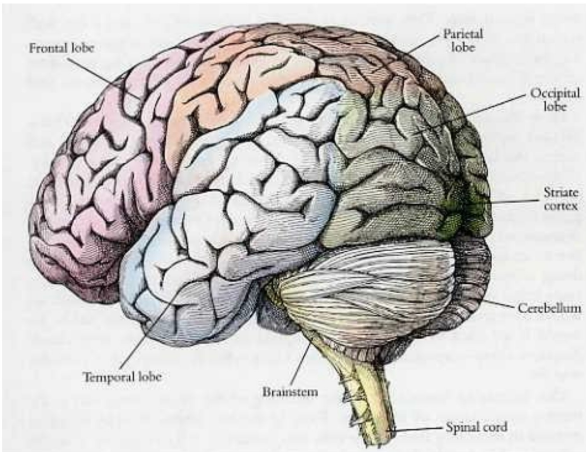

This view of a human brain seen from the left and slightly behind shows the cerebral cortex and cerebellum. A small part of the brainstem can be seen just in front of the cerebellum.