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

Lecture 1: Motivations and Overview
2026.01.05

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

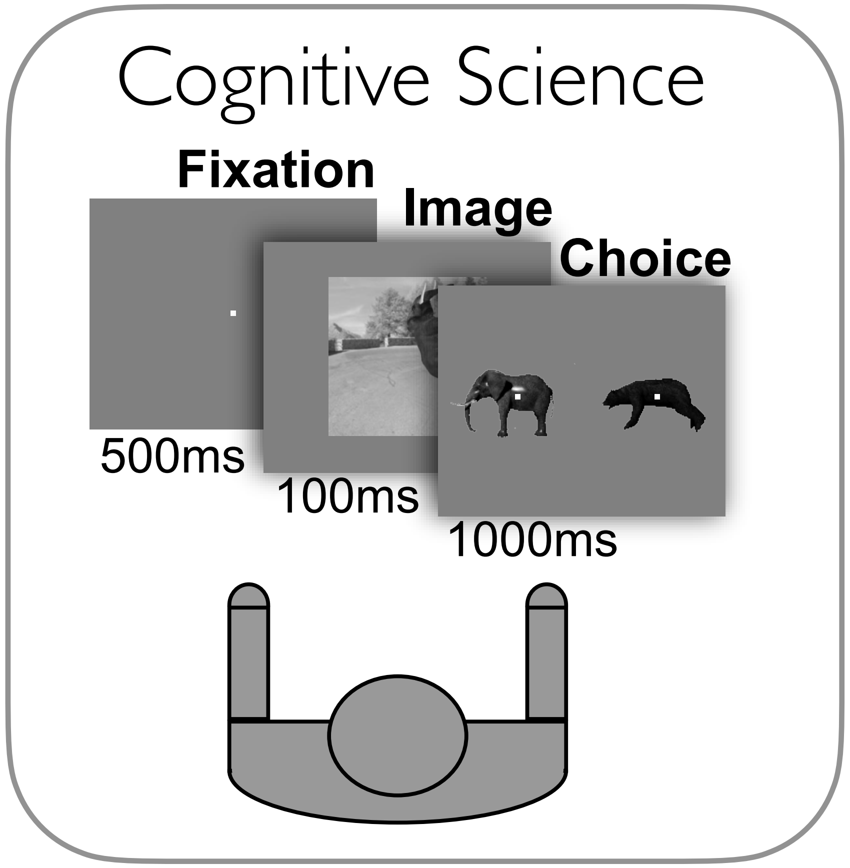

Benchmarking Humans and Machines at Scale

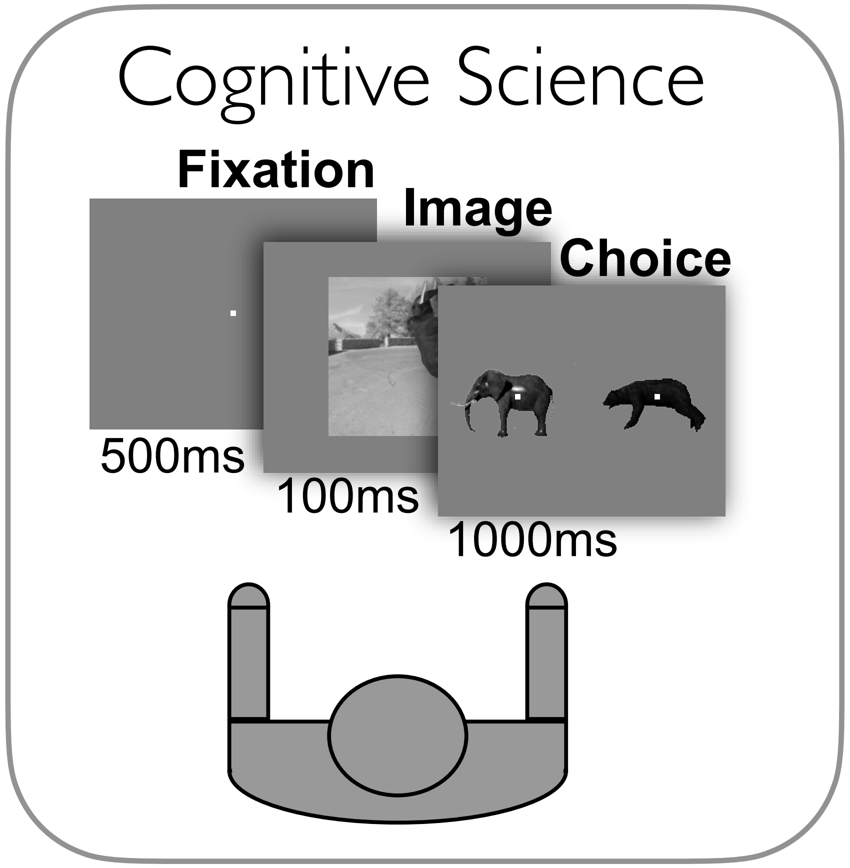

Benchmarking Humans and Machines at Scale

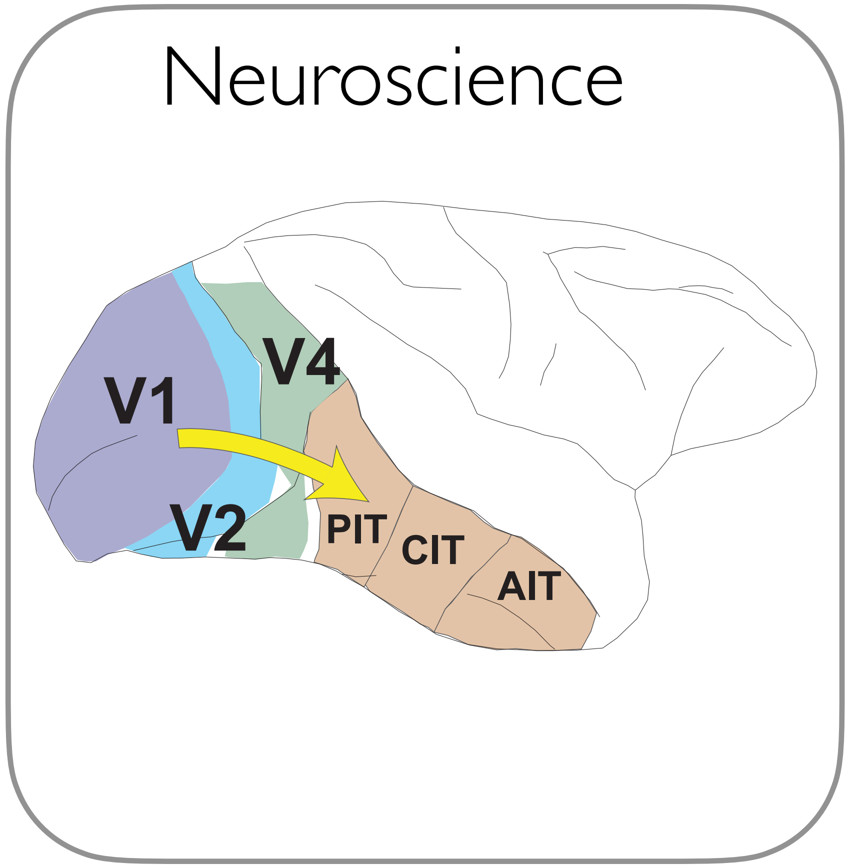

Making Predictive Models of Brain Data

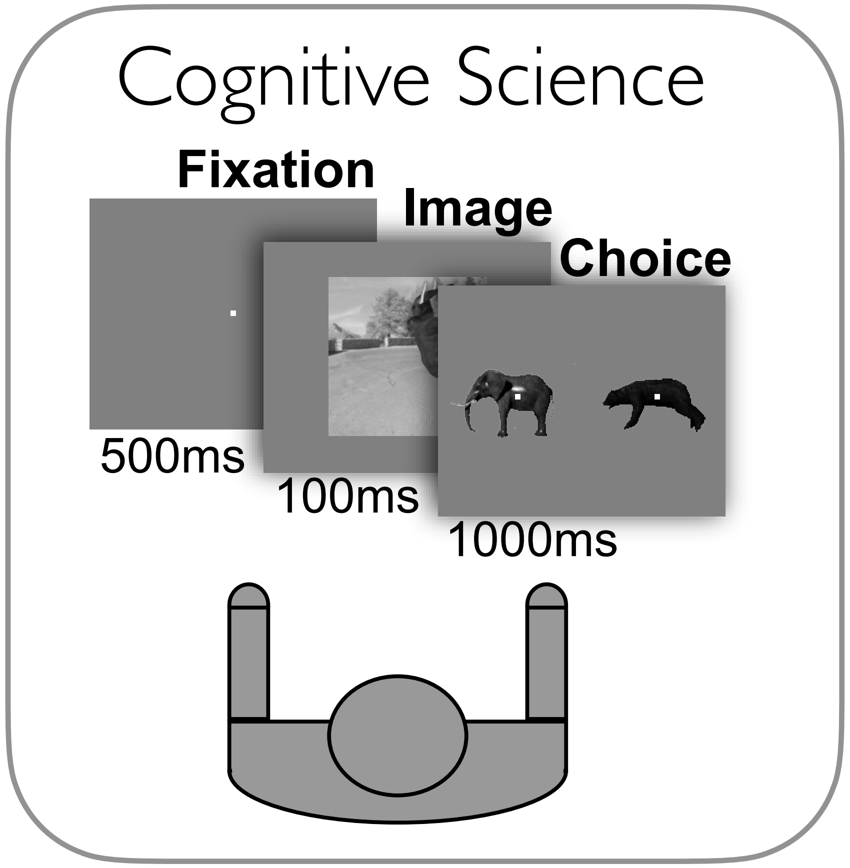

Benchmarking Humans and Machines at Scale

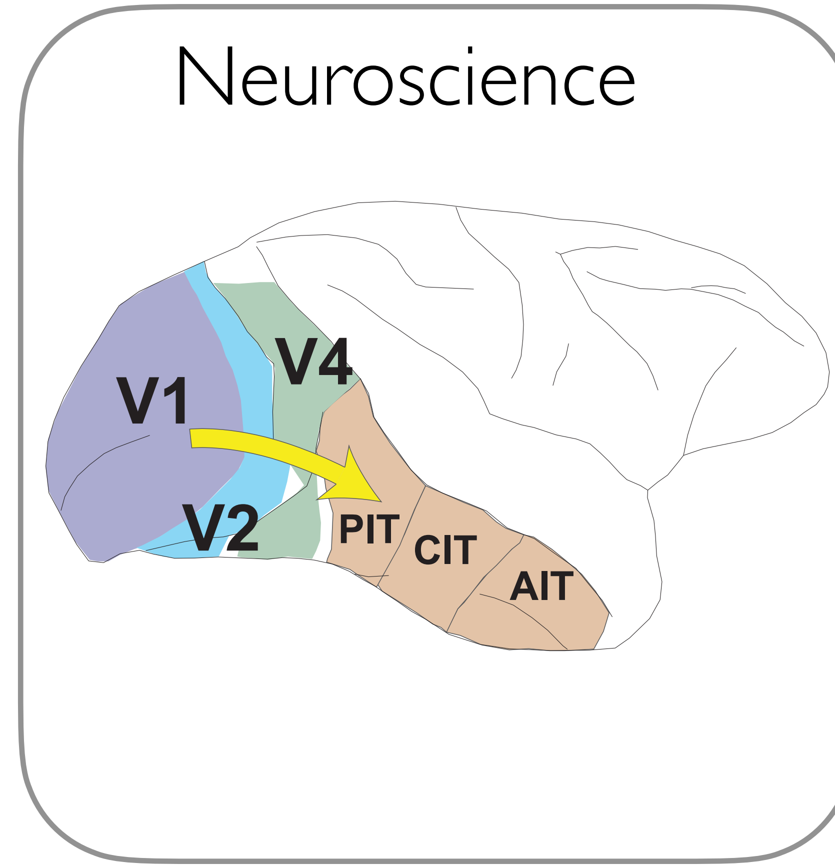

Making Predictive Models of Brain Data

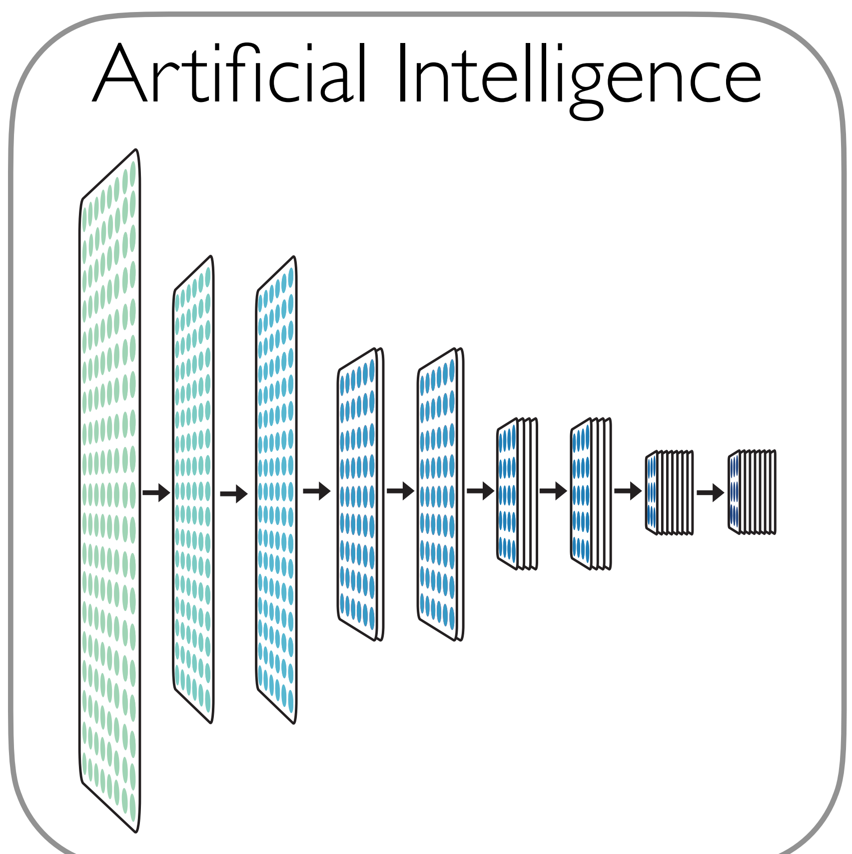

Building Neural Networks to Solve Cognitive Tasks

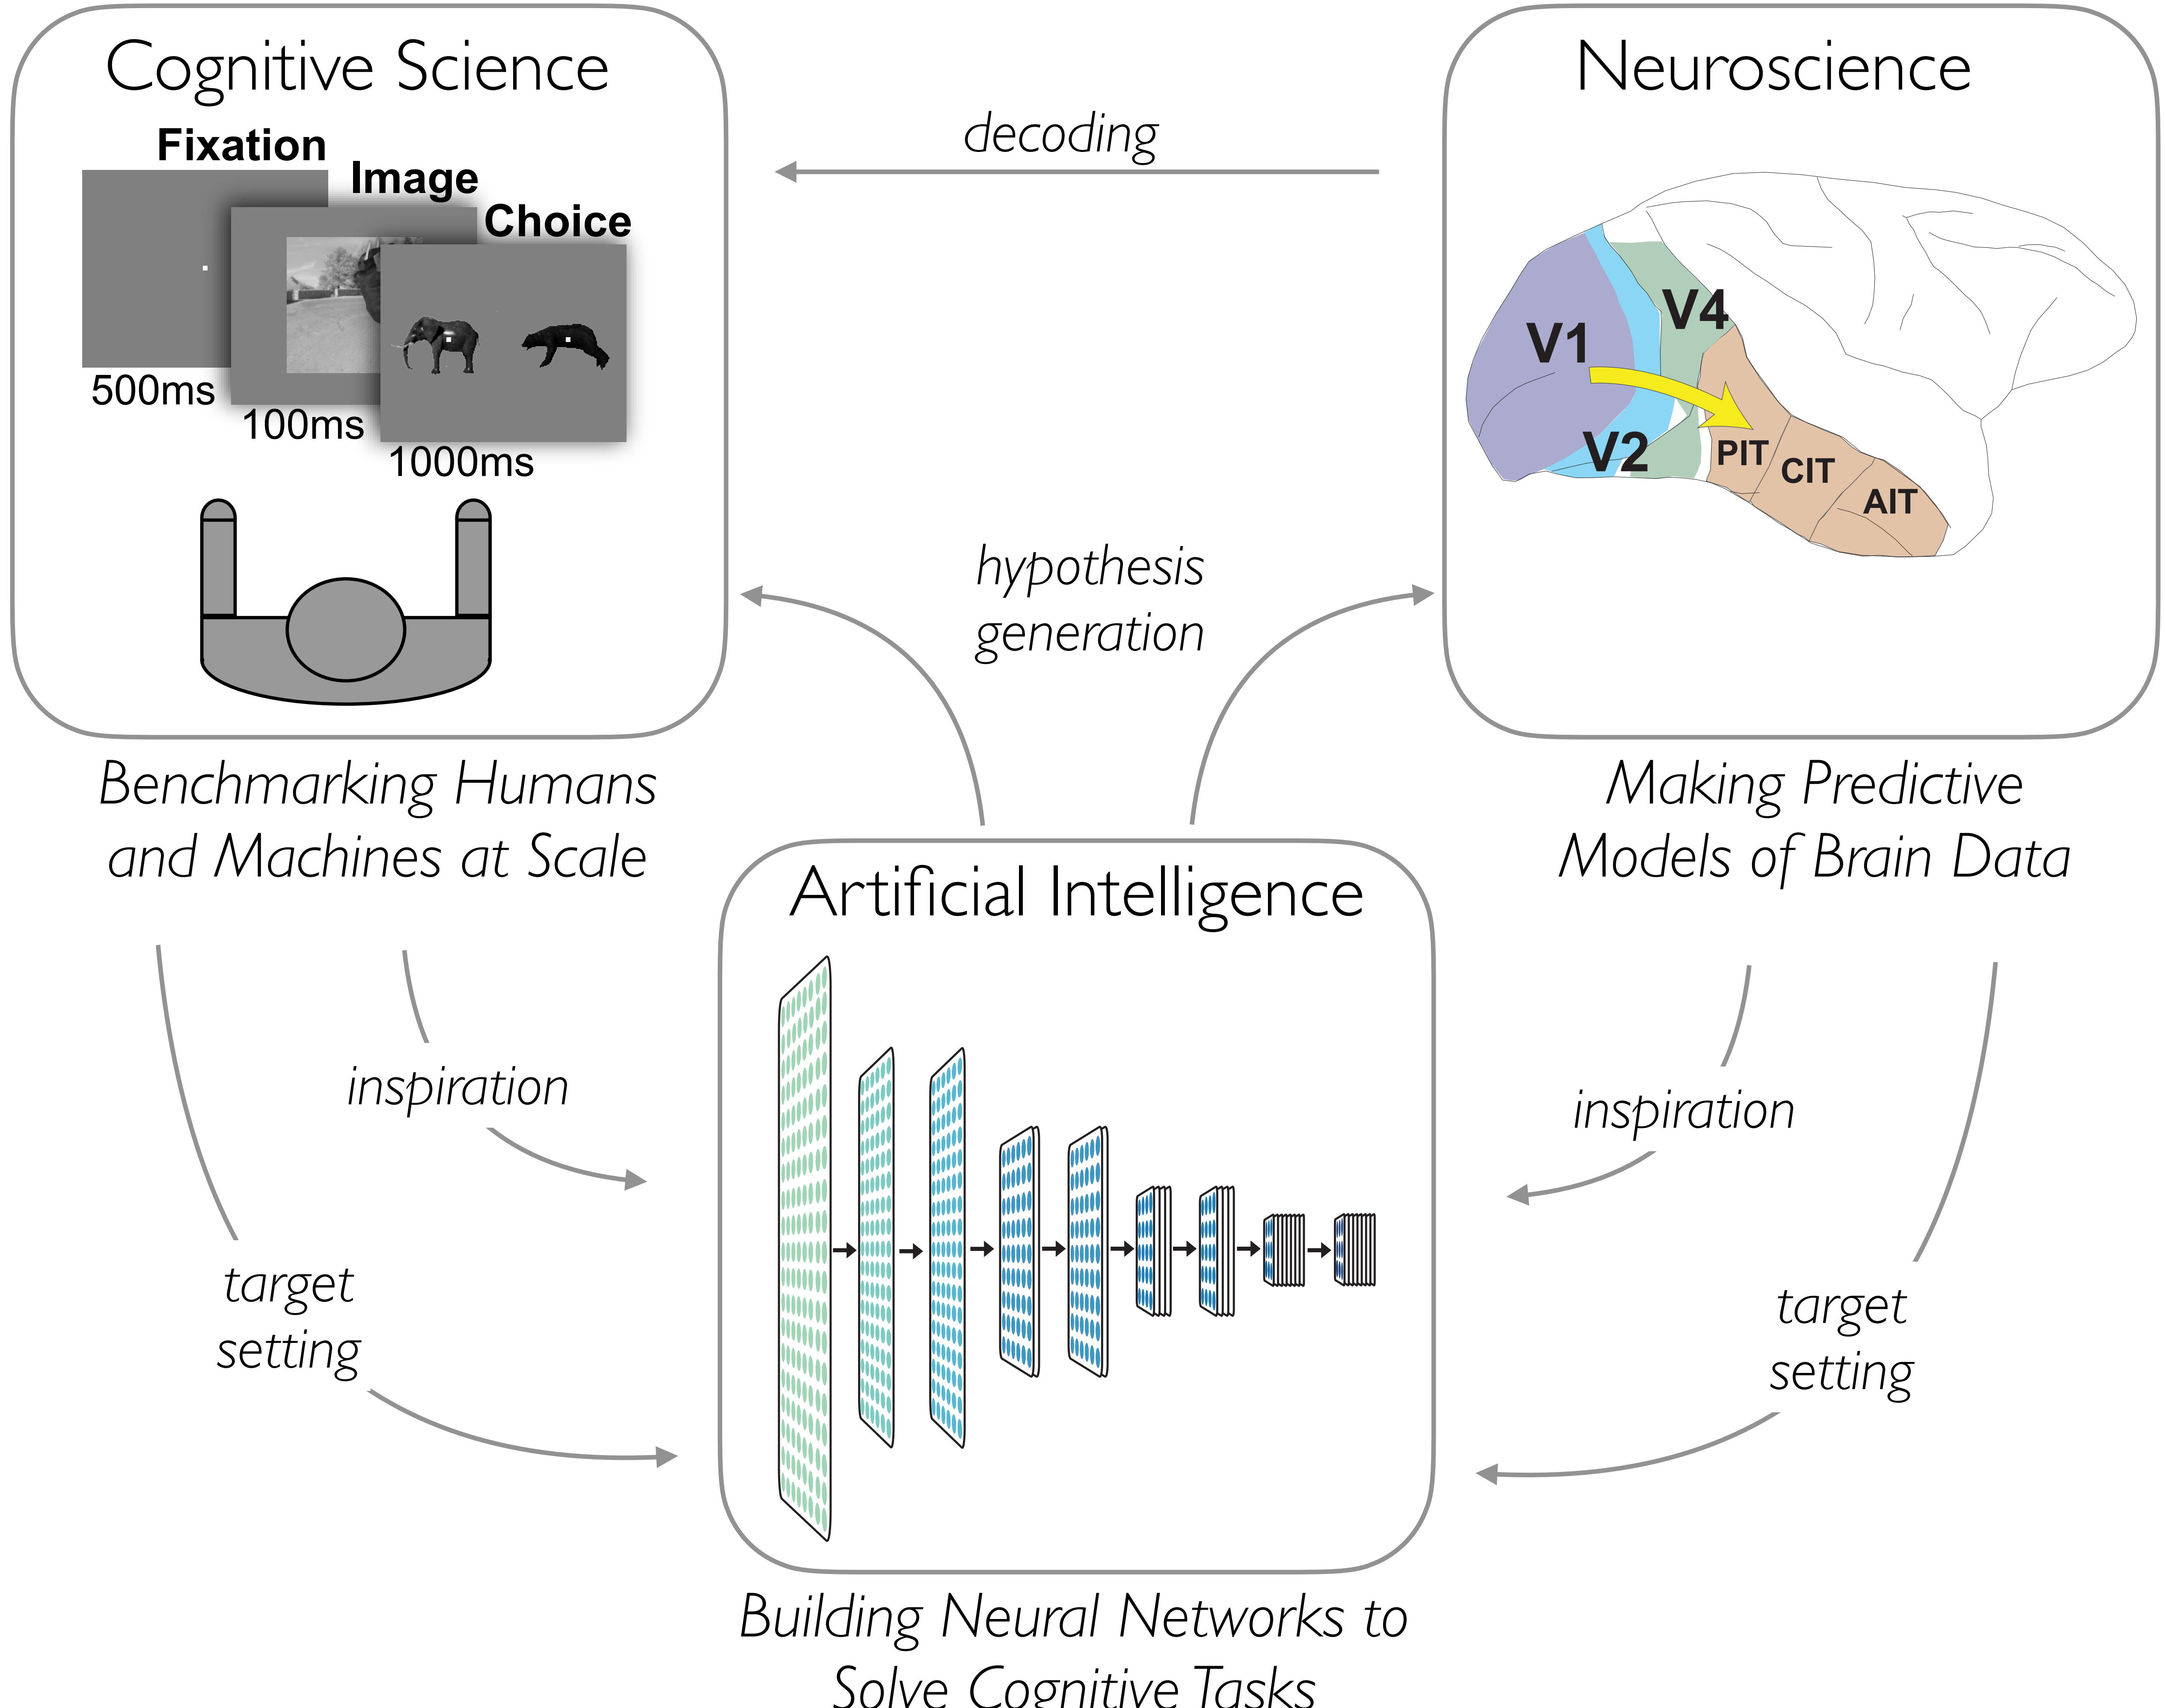
Cognitive Science
Fixation
Image
Choice
500ms
100ms
1000ms
Neuroscience
decoding
V1
V4
V2
PIT
CIT
AIT
hypothesis generation
Benchmarking Humans and Machines at Scale
Making Predictive Models of Brain Data
Artificial Intelligence
inspiration
inspiration
target setting
target setting
Building Neural Networks to Solve Cognitive Tasks

Question: What does it mean to understand how the brain works?

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

sensory

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

motor

sensory

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

motor

decision making

sensory

memory

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

motor

decision making

sensory

memory

learning

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

motor

decision making

sensory

memory

communication

learning

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

motor

decision making

sensory

Being an agent

memory

communication

learning

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors:

motor

decision making

sensory

Being an agent
Being with other agents

memory

communication

learning

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors.

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors.

Answer: To be able to easily build an artificial system that:

(1) behaves like the human at high resolution

Question: What does it mean to understand how the brain works?

Observation: Constellation of interesting behaviors.

Answer: To be able to easily build an artificial system that:

(1) behaves like the human at high resolution

(2) whose internal parts can be mapped to the parts of the brain at some chosen level of resolution

To be able to **easily build** an artificial system that:

(1) behaves like the human at high resolution

(2) whose internal parts can be mapped to the parts of the brain at some chosen level of resolution

easily build = simulate at low cost

To be able to easily build an **artificial system** that:

(1) behaves like the human at high resolution

(2) whose internal parts can be mapped to the parts of the brain at some chosen level of resolution

artificial system = neural network

To be able to easily build an artificial system that:

(1) **behaves like the human** at high resolution

(2) whose internal parts can be mapped to the parts of the brain at some chosen level of resolution

behaves like human = (a) has similar types of input sensors

(b) has similar types of output actuators

(c) generates similar input/output map

(d) develops and learns similar way

To be able to easily build an artificial system that:

(1) **behaves like the human** at high resolution

(2) whose internal parts can be mapped to the parts of the brain at some chosen level of resolution

develops and **learns** similar way

changes itself in response to environmental input

starts like a baby; changes itself independently of a environmental input

To be able to easily build an artificial system that:

(1) behaves like the human at high resolution

(2) whose internal parts can be mapped to the parts of the brain at some chosen level of resolution

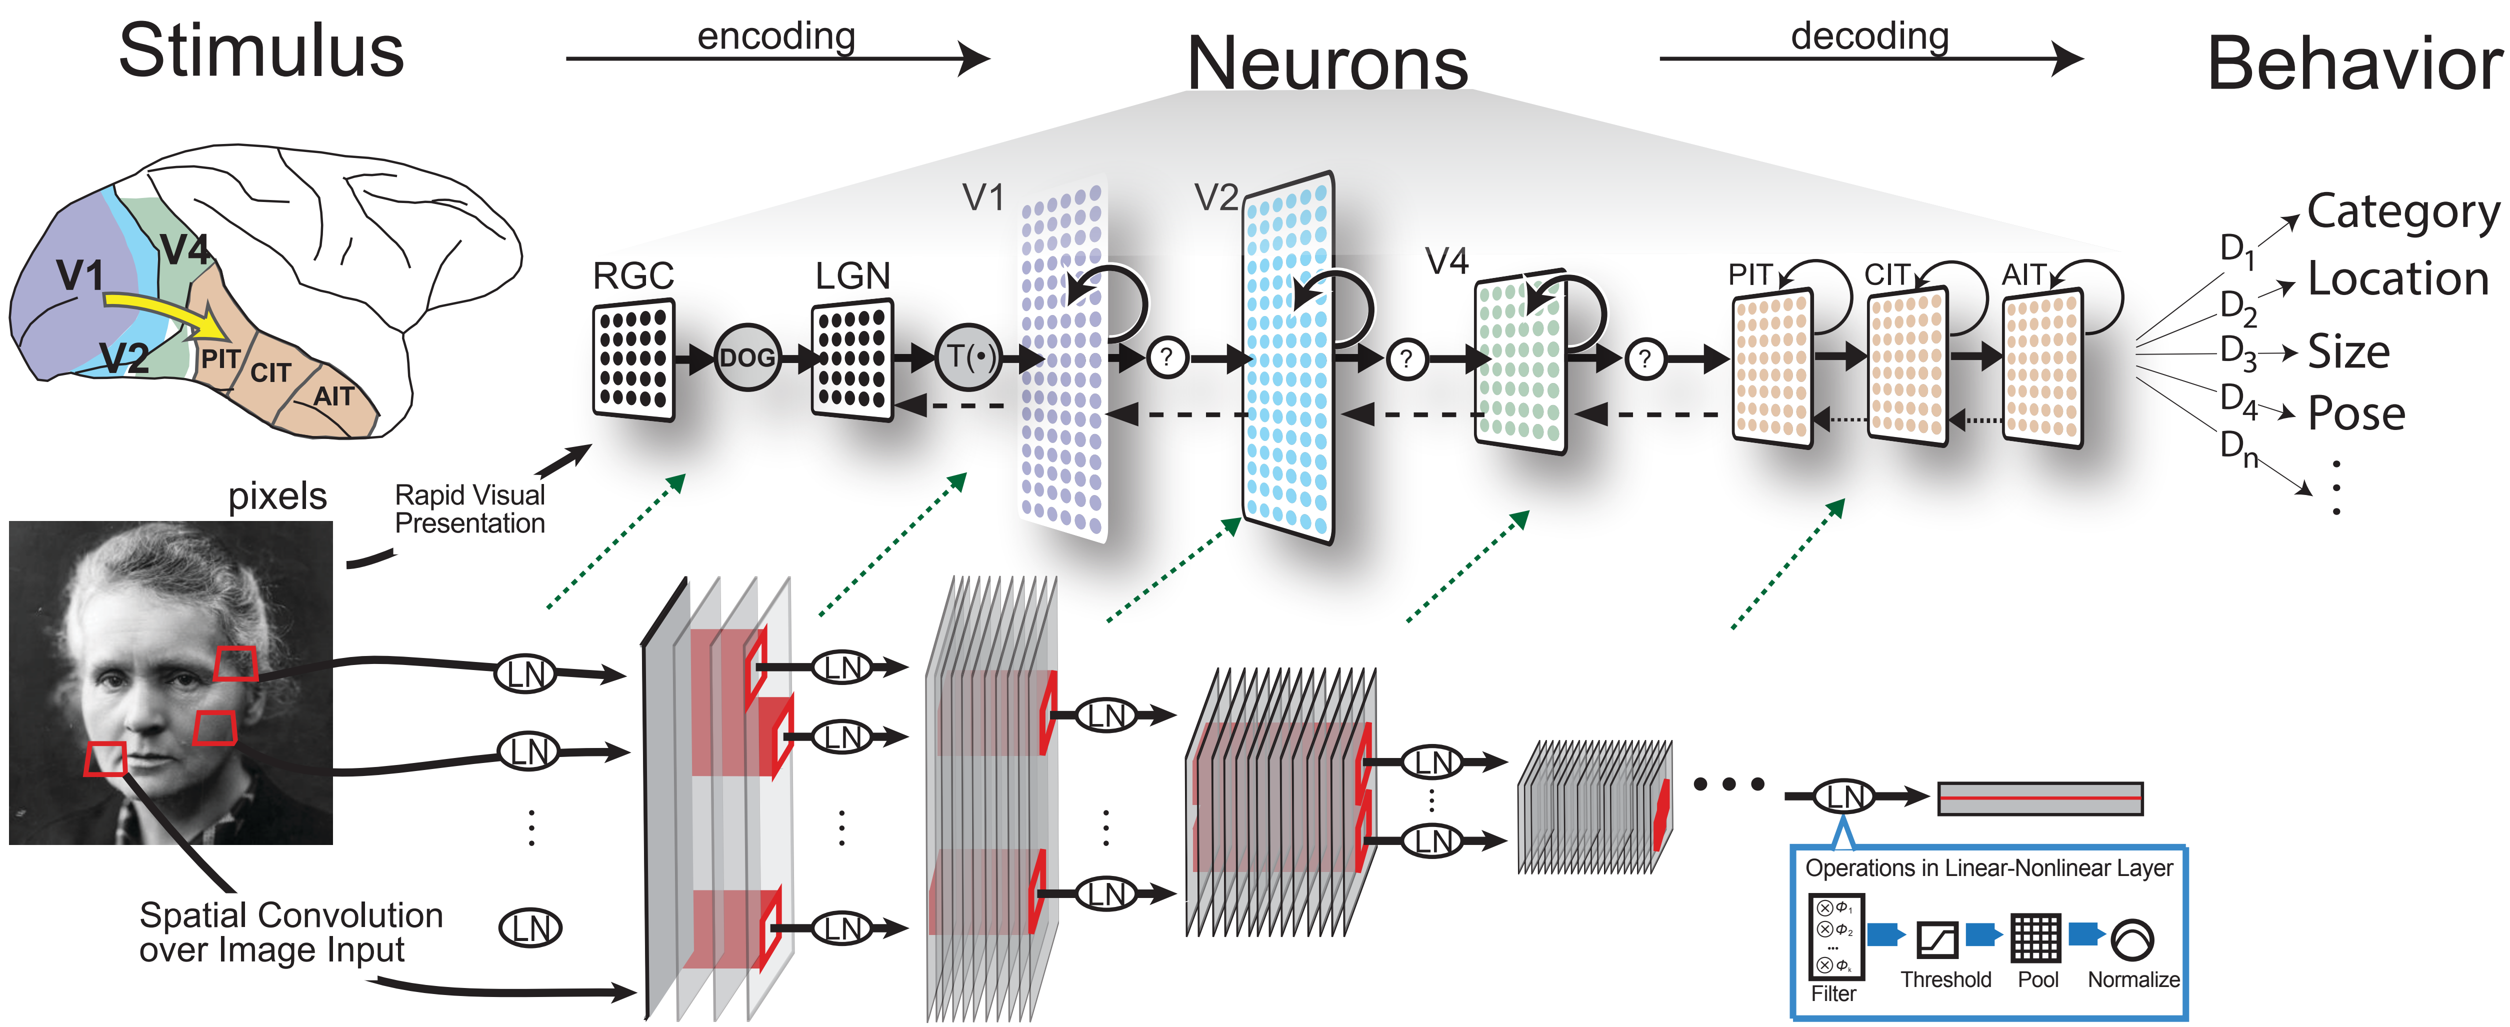

Many Different Computational Goals

- Sensory processing
  - visual, auditory, somatosensory processing (occipital, temporal)
  - navigation (hippocampus?)

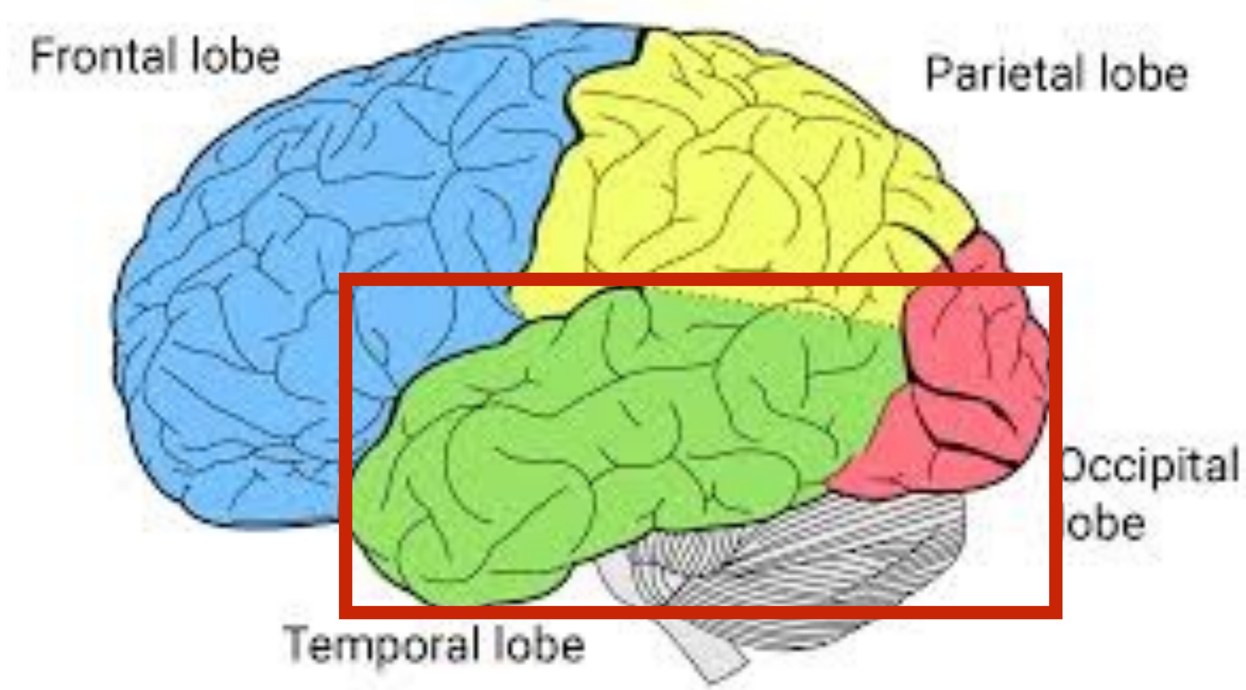

Many Different Computational Goals

- Sensory processing
  - visual, auditory, somatosensory processing (occipital, temporal)
  - navigation (hippocampus?)

- motor command production & execution (motor cortex)

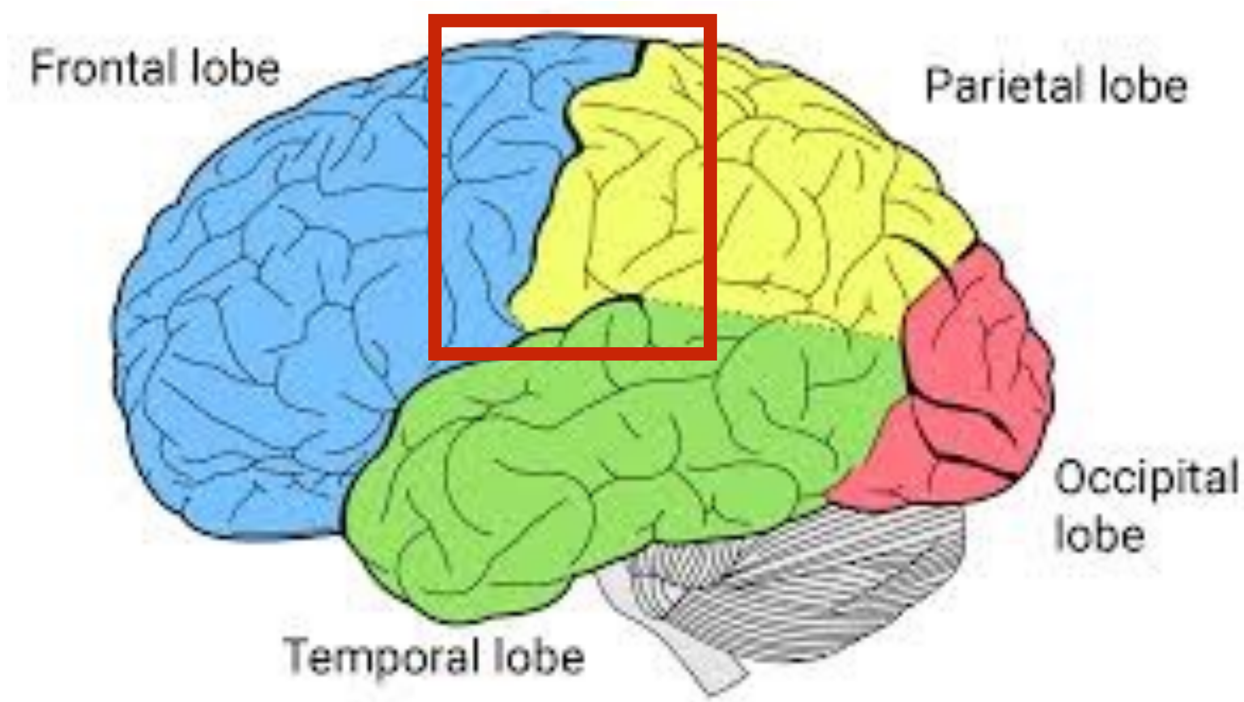

Many Different Computational Goals

- Sensory processing
  - visual, auditory, somatosensory recognition (occipital, temporal)
  - navigation (hippocampus?)

- motor command production & execution (motor cortex)

- memory, decision making and planning (hippocampus, prefrontal cortex)

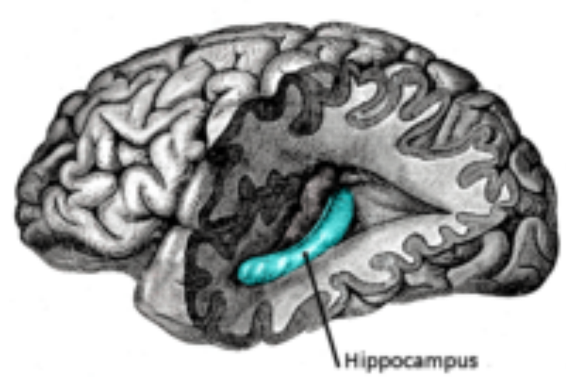

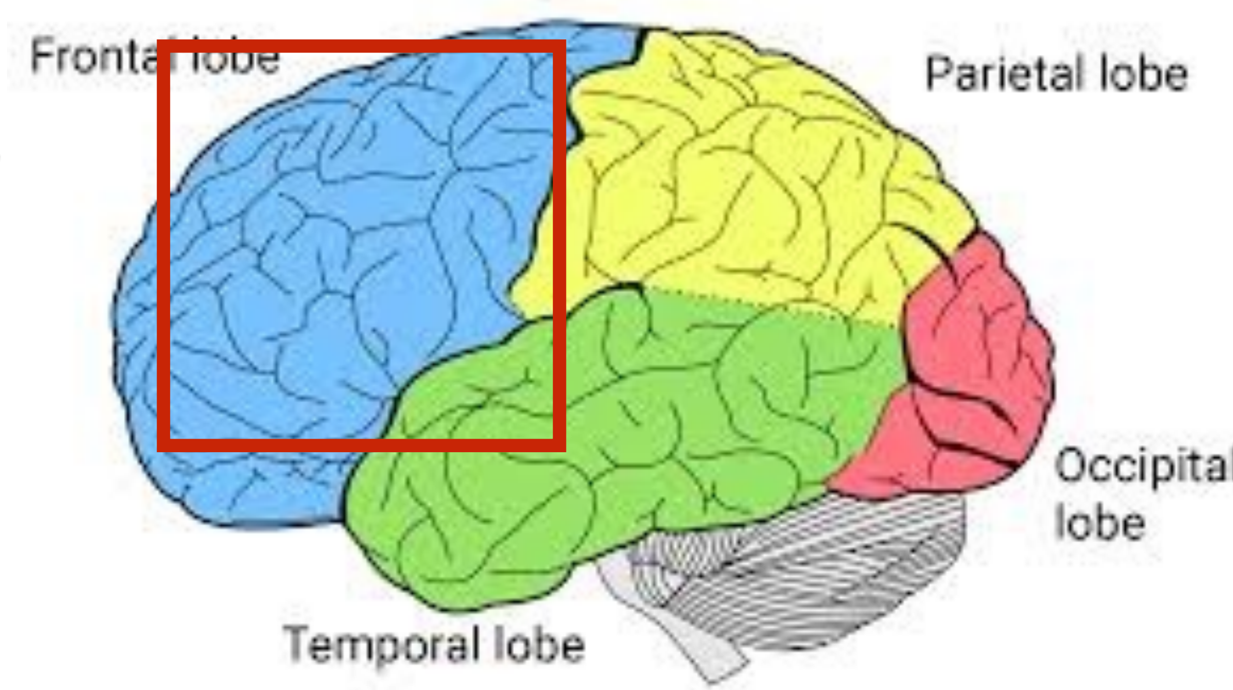

Many Different Computational Goals

- Sensory processing
  - visual, auditory, somatosensory recognition (occipital, temporal)
  - navigation (hippocampus?)
- motor command production & execution (motor cortex)
- memory, decision making and planning (hippocampus, prefrontal cortex)
- language

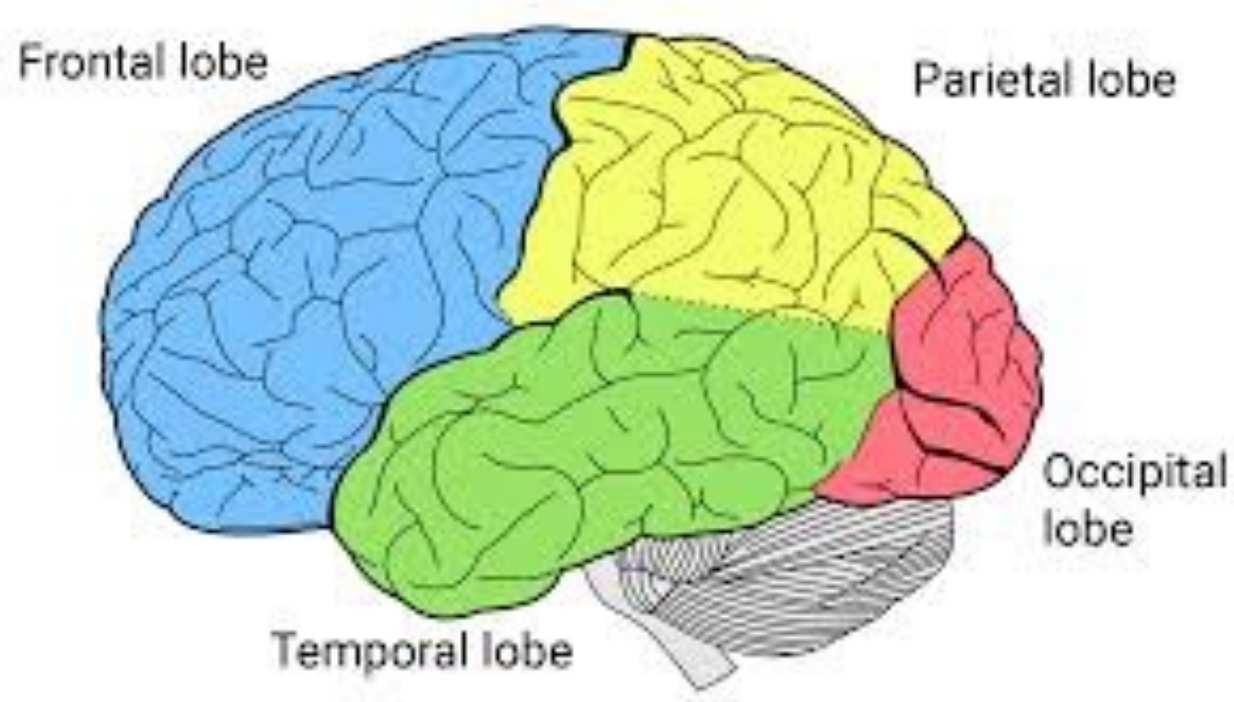

Many Different Computational Goals

- Sensory processing
  - visual, auditory, somatosensory recognition (occipital, temporal)
  - navigation (hippocampus?)
- motor command production & execution (motor cortex)
- memory, decision making and planning (hippocampus, prefrontal cortex)
- language
- emotions

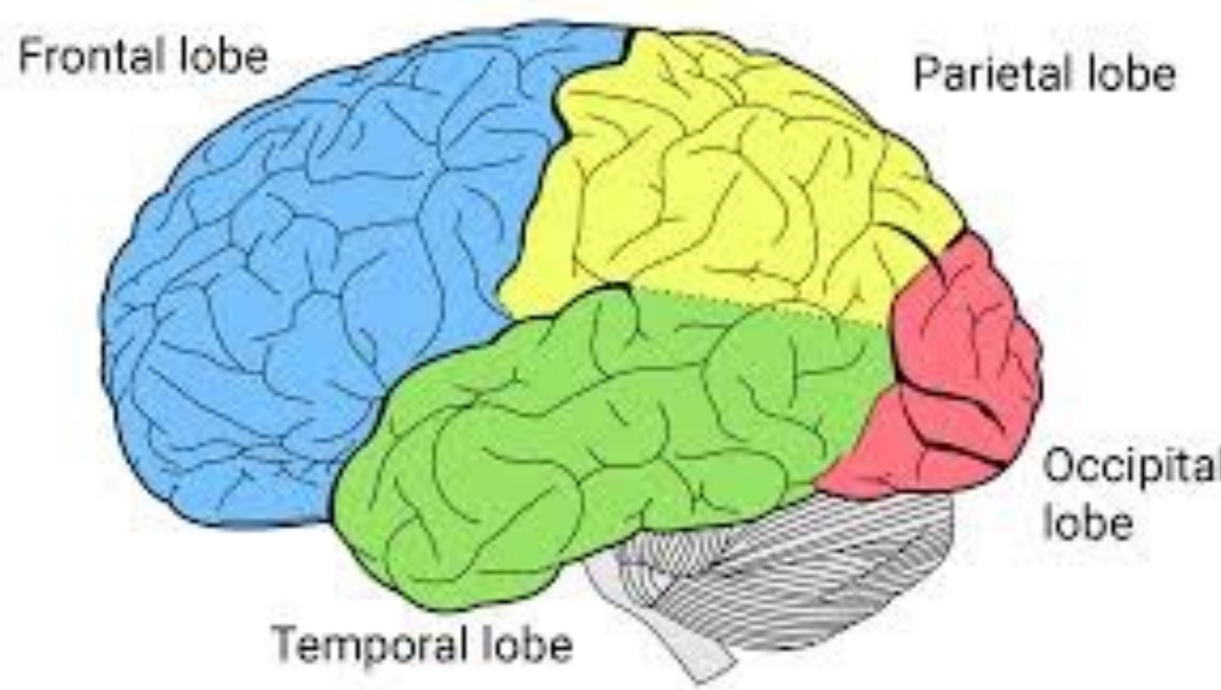

A Few Basic Ideas

The neuroscience origins of neural networks

The Neuron Doctrine

The "Strong" Neuron Doctrine:

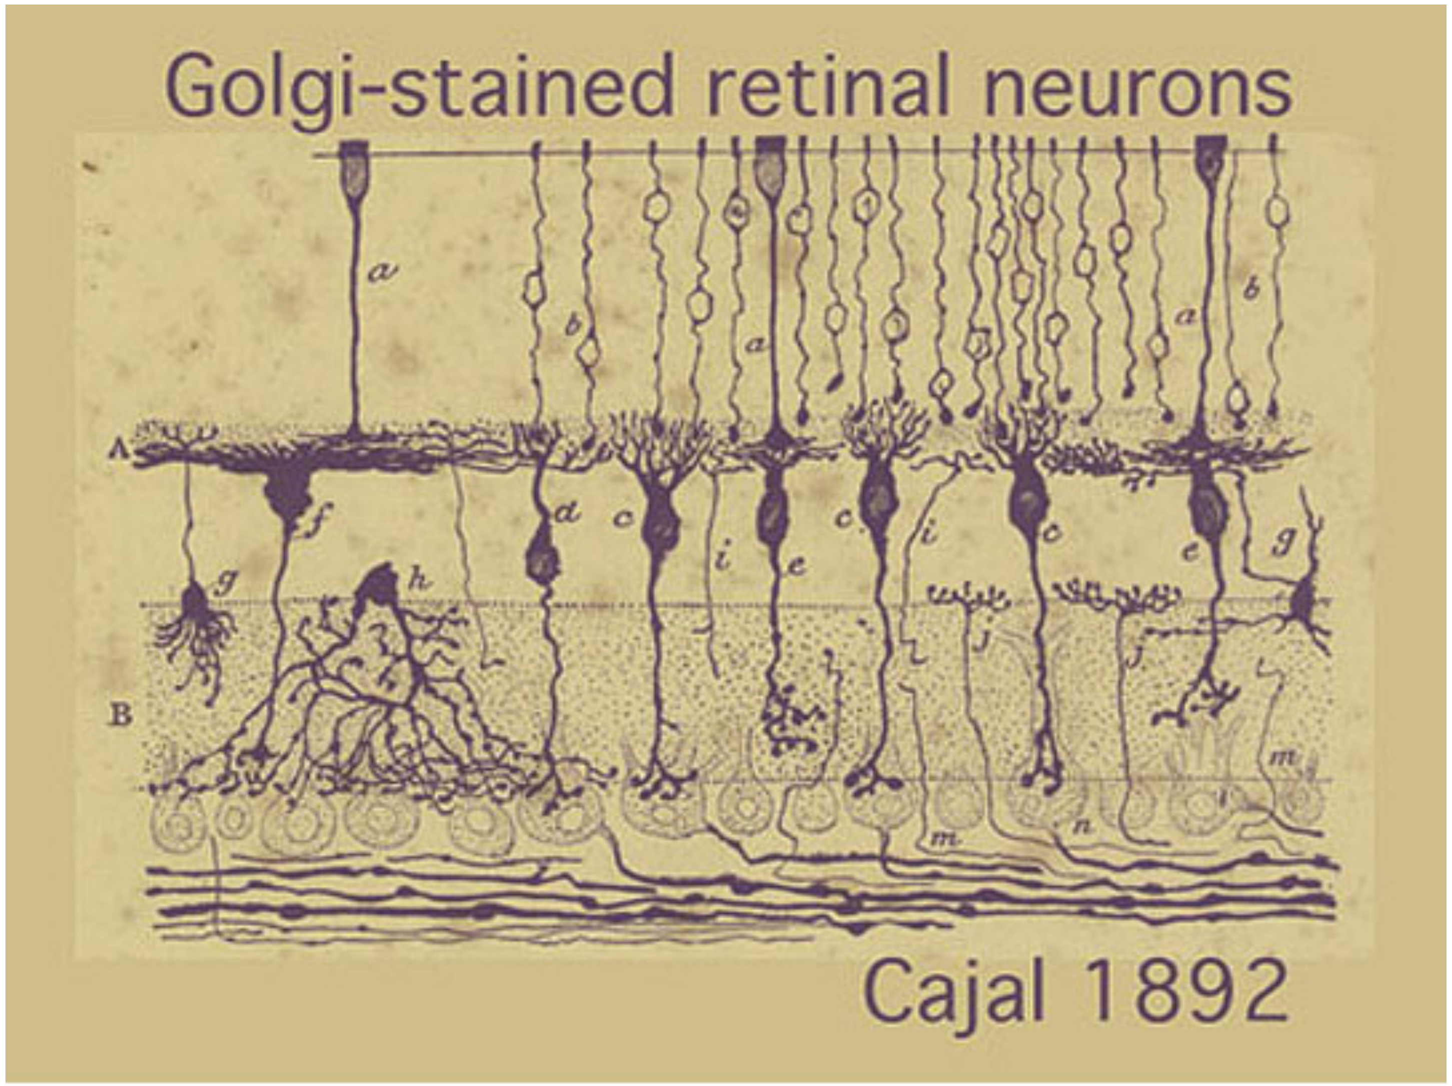

Fig. 2. A drawing done by Cajal to show some of the neurons of the retina in vertical section.

The Neuron Doctrine

The "Strong" Neuron Doctrine:

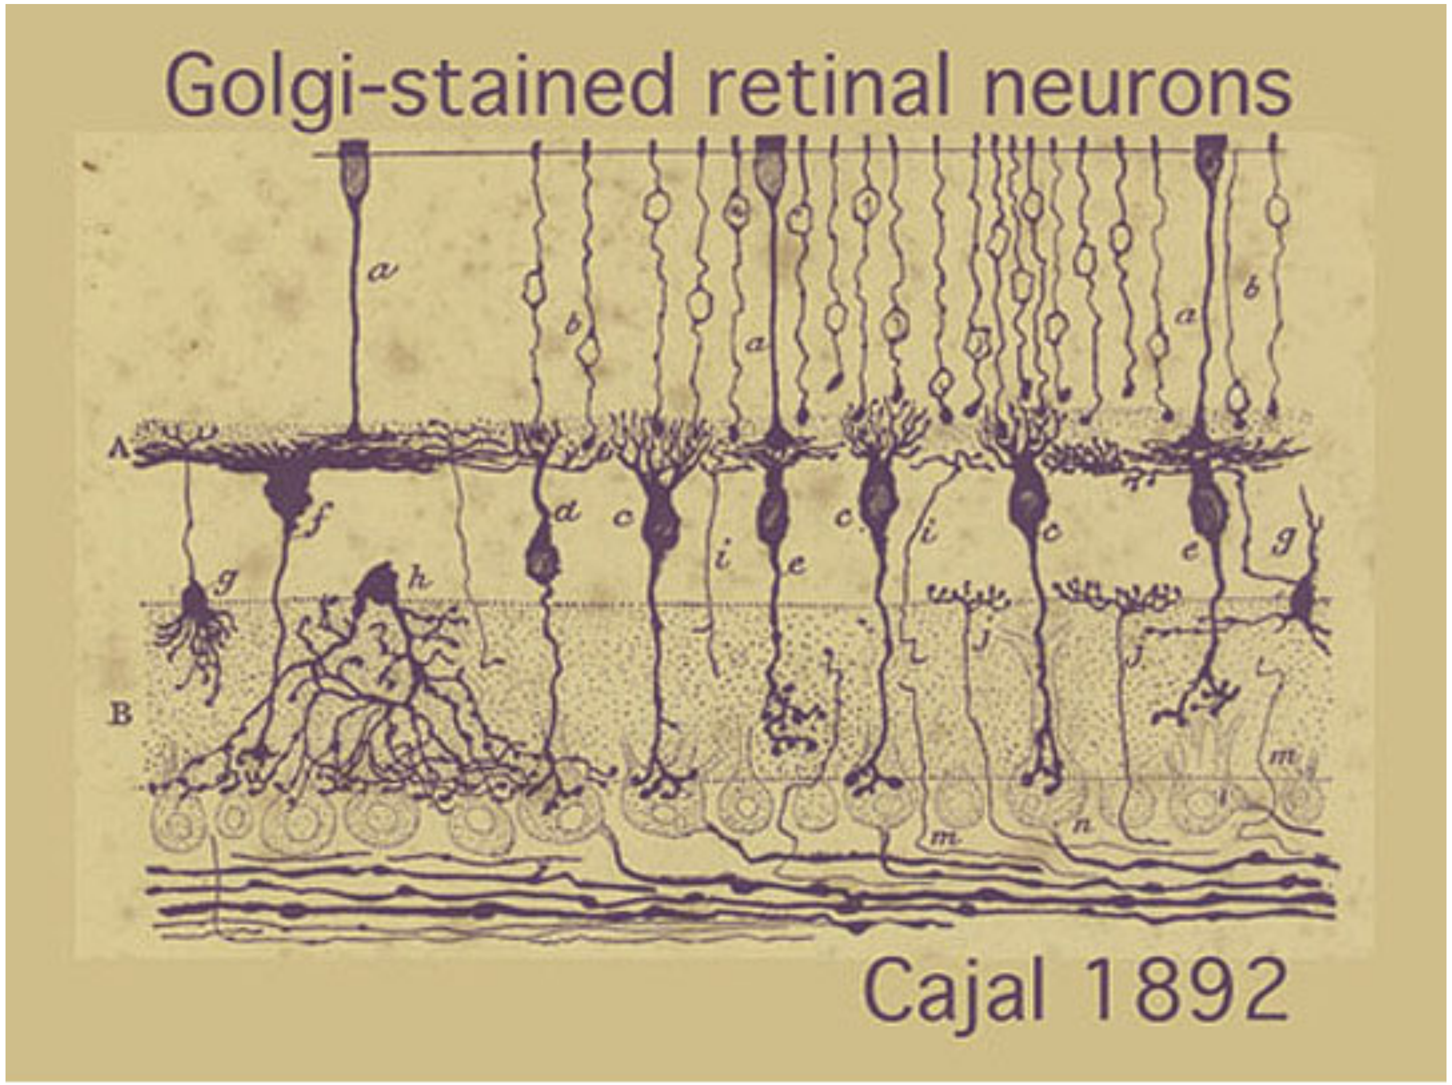

Fig. 2. A drawing done by Cajal to show some of the neurons of the retina in vertical section.

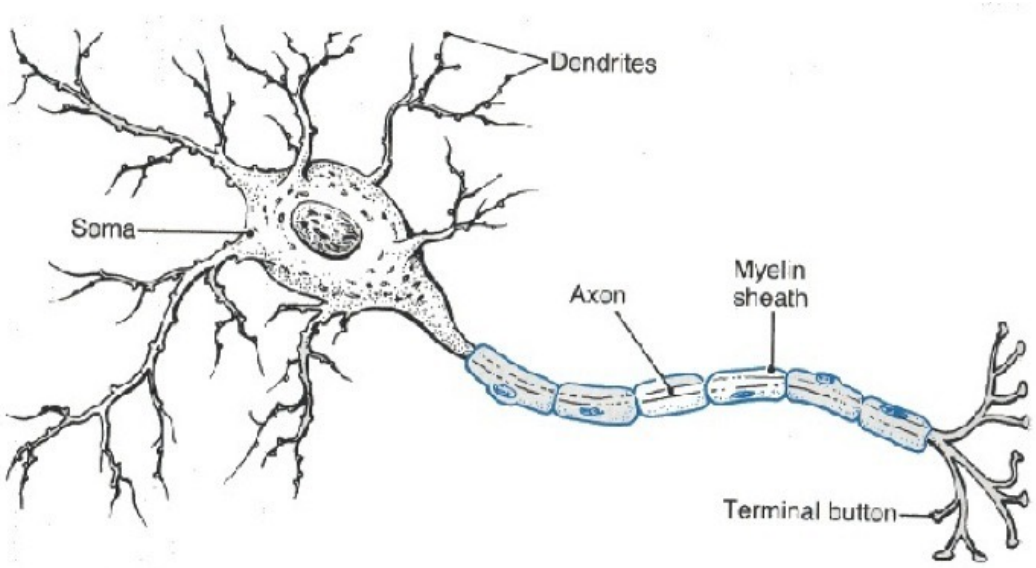

a neuron receives signals at its dendrites and cell body and transmits them, as action potentials, along the axon in one direction: away from the cell body

The "Strong" Neuron Doctrine:

i. *The nervous system is made up of discrete cells ("neurons"), connected by extracellular junctions (synapses) into a directed graph*

The "Strong" Neuron Doctrine:

i. *The nervous system is made up of discrete cells ("neurons"), connected by extracellular junctions (synapses) into a directed graph*

ii. *neurons are "excitable cells" that "fire" by a mechanism of electrochemical (de)polarization [Hodgkin-Huxel model]*

The Neuron Doctrine

The "Strong" Neuron Doctrine:

i. *The nervous system is made up of discrete cells ("neurons"), connected by extracellular junctions (synapses) into a directed graph*

ii. *neurons are "excitable cells" that "fire" by a mechanism of electrochemical (de)polarization [Hodgkin-Huxel model]*

iii. *the firing pattern of a neuron is a parameterized function that "integrates" the firing patterns of the neurons that synapse onto it*

The “Strong” Neuron Doctrine:

i. *The nervous system is made up of discrete cells (“neurons”), connected by extracellular junctions (synapses) into a directed graph*

ii. *neurons are “excitable cells” that “fire” by a mechanism of electrochemical (de)polarization [Hodgkin-Huxel model]*

iii. *the firing pattern of a neuron is a parameterized function that “integrates” the firing patterns of the neurons that synapse onto it*

iv. *the parameters of the function are plastic and therefore learnable*

Artificial Neural Networks (ANNs)

McCulloch and Pitts (1943)

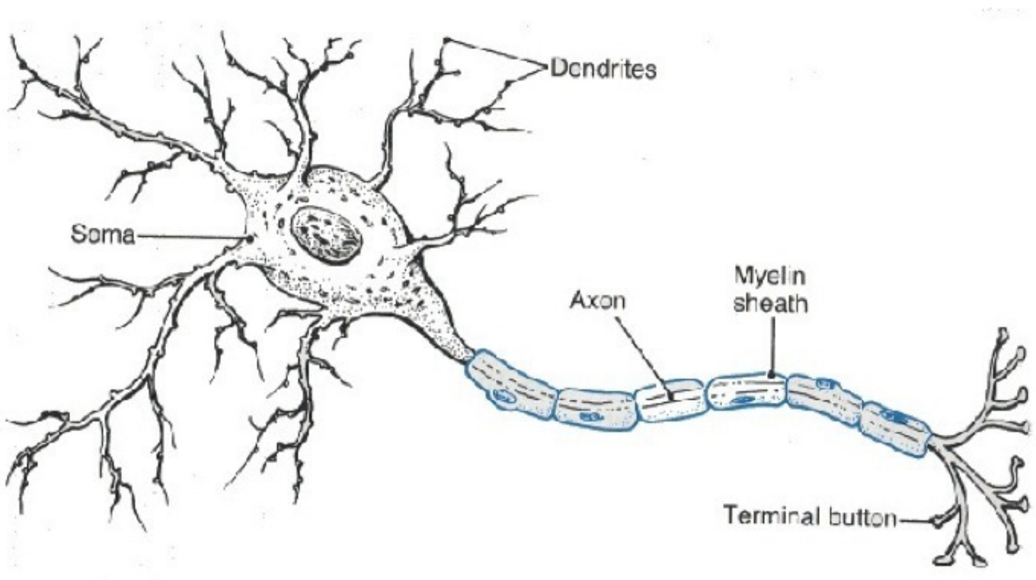

Artificial Neural Networks (ANNs)

McCulloch and Pitts (1943)

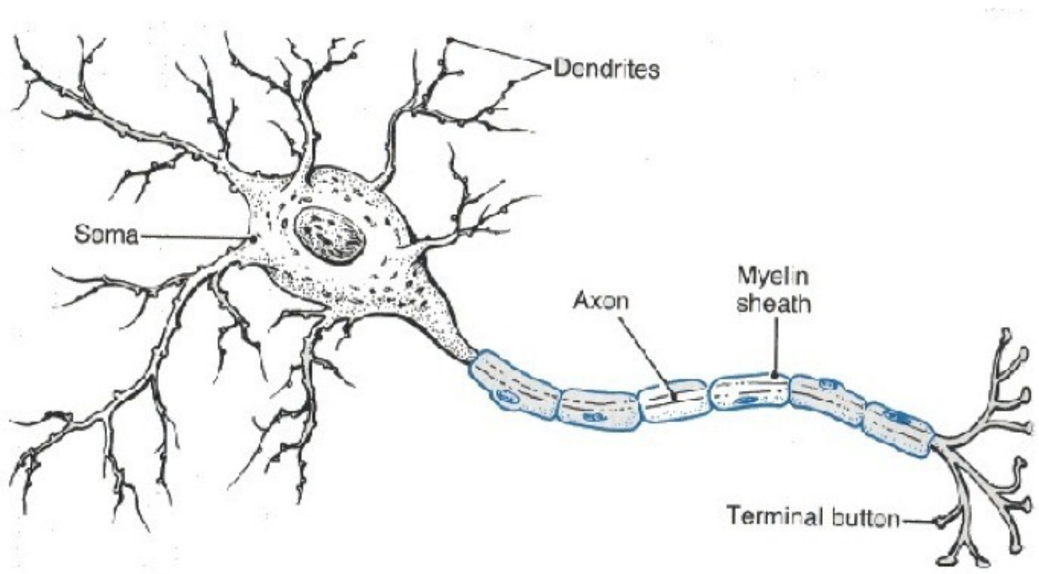

$$y_k = \phi\left(\sum_{j=0}^{m} w_{kj} x_j\right)$$

$$\phi : \mathbb{R} \longmapsto \mathbb{R}$$

some (nonlinear) activation function

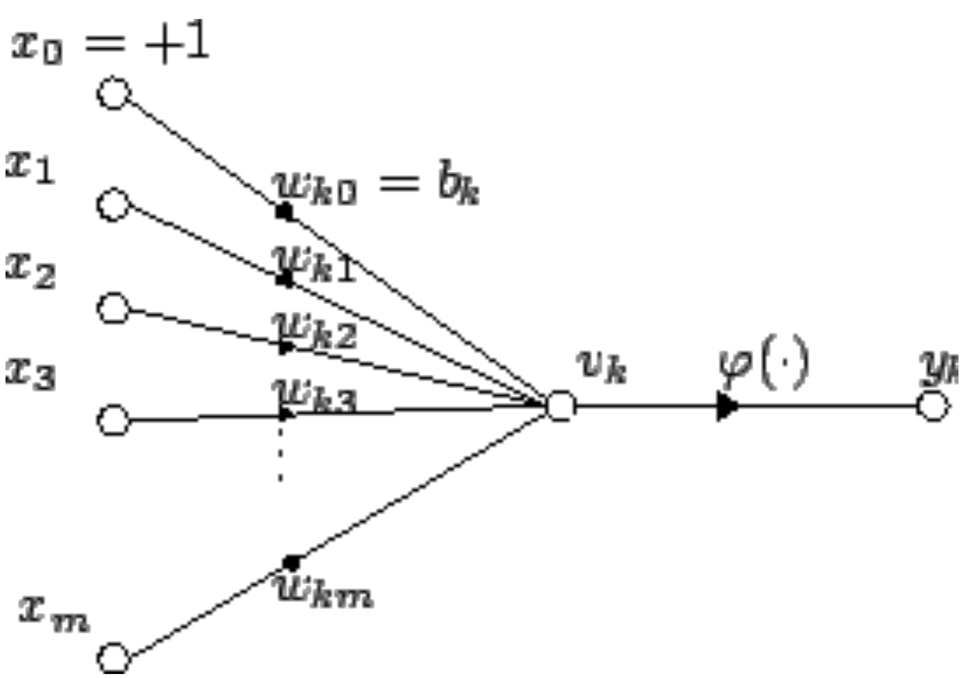

$$w_{kj} \in \mathbb{R}^{m+1}$$

"synaptic strengths"

Artificial Neural Networks (ANNs)

McCulloch and Pitts (1943)

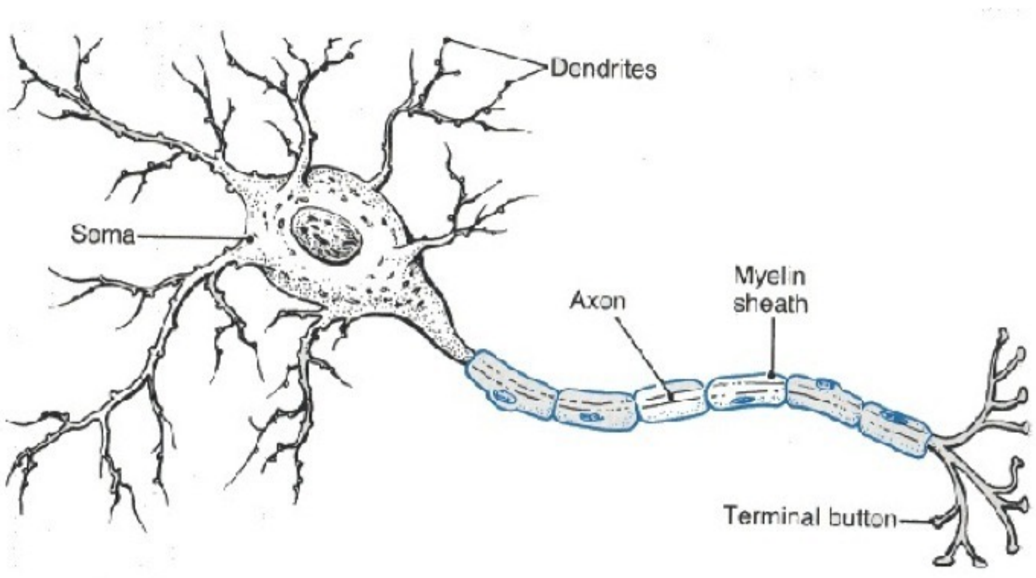

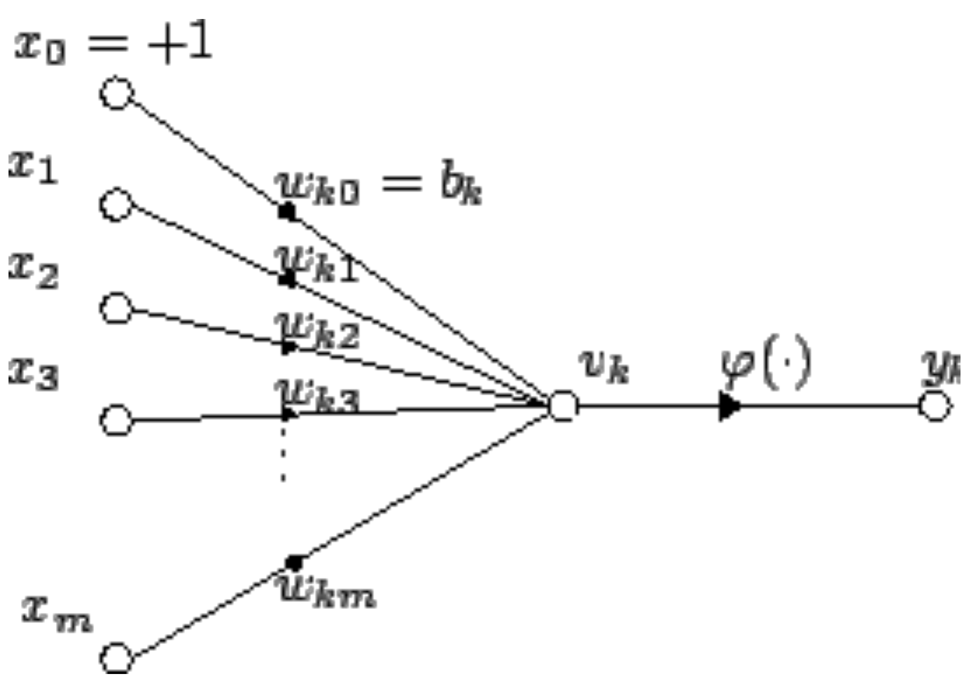

$$y_k = \phi\left(\sum_{j=0}^{m} w_{kj} x_j + b_j\right)$$

$$\phi : \mathbb{R} \longmapsto \mathbb{R}$$

some (nonlinear) activation function

$$w_{kj} \in \mathbb{R}^{m+1}$$

"synaptic strengths"

$$b_j \in \mathbb{R}$$

"biases"

Artificial Neural Network (ANN) Models of the Brain

Core (obvious) idea: Model brain systems with ANNs

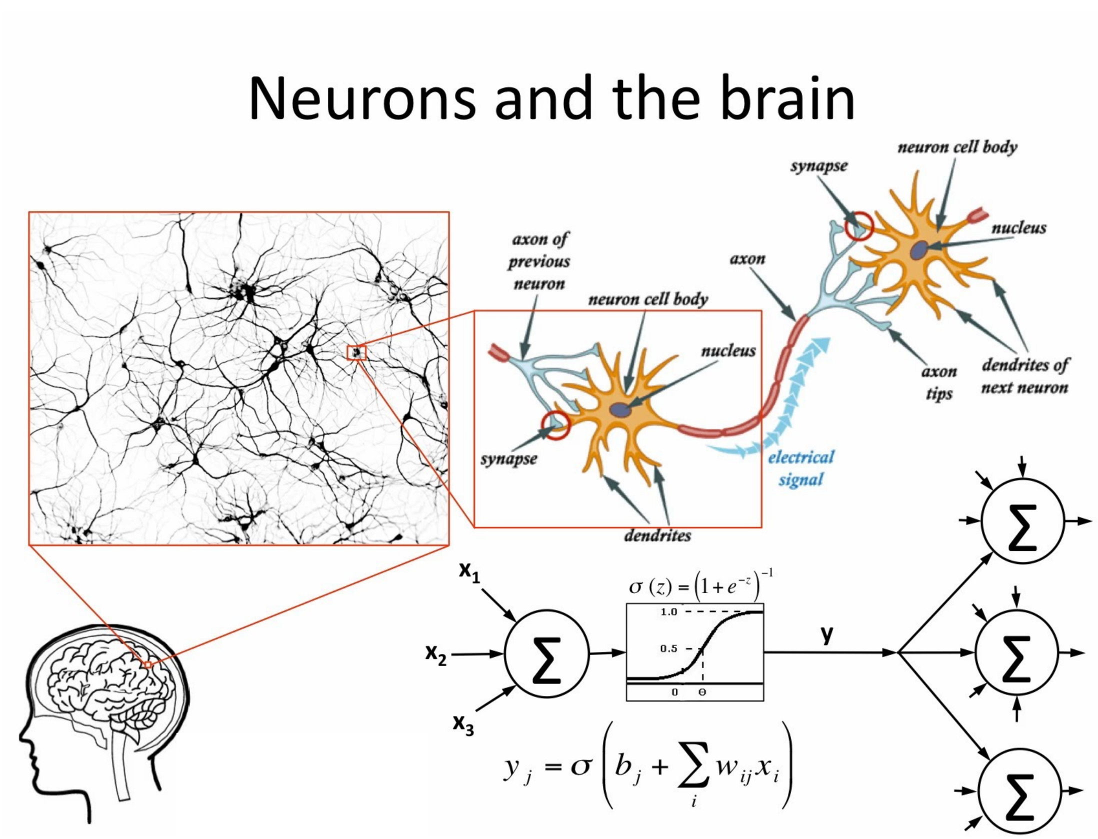

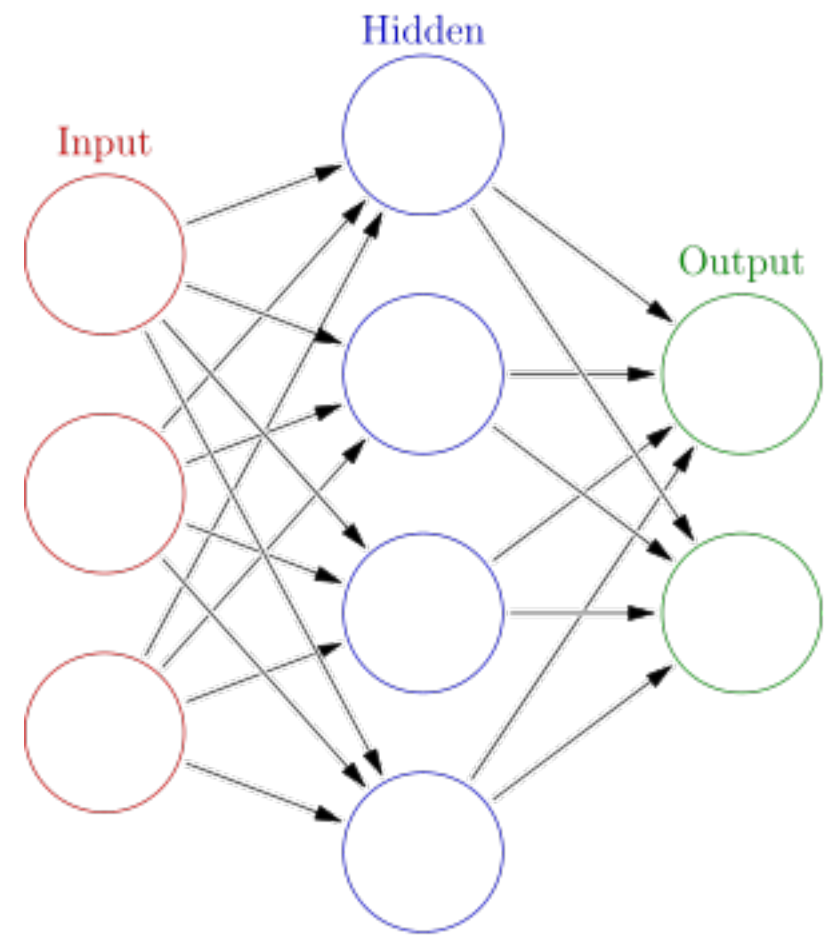

Artificial Neural Network (ANN) Models of the Brain

Core (obvious) idea: Model brain systems with ANNs

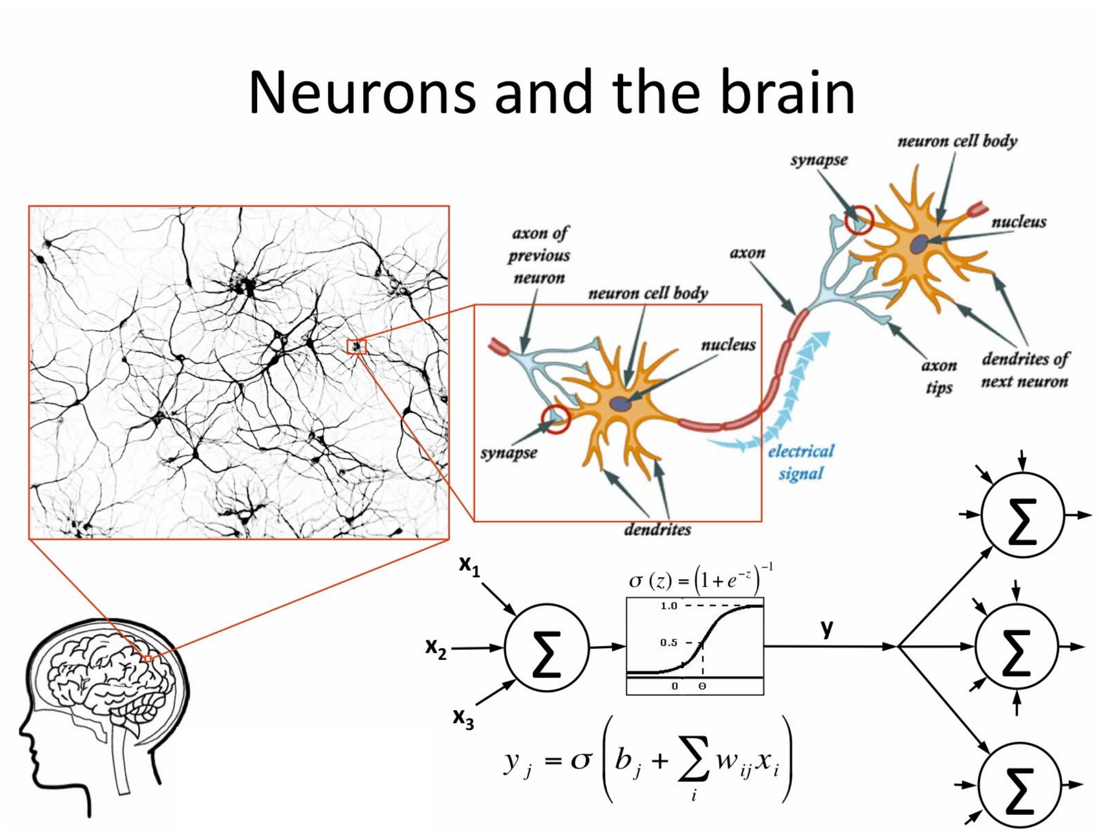

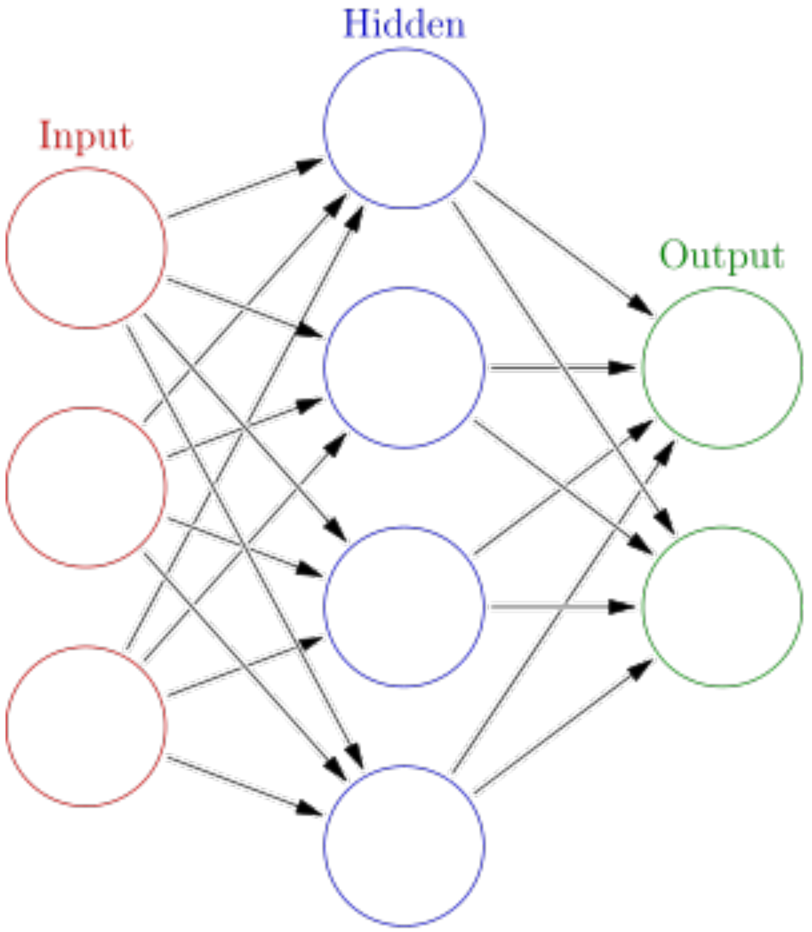

But how to find the correct parameters?*

*both continuous parameters like weights and discrete parameters of the architecture

Artificial Neural Network (ANN) Models of the Brain

Core (obvious) idea: Model brain systems with ANNs

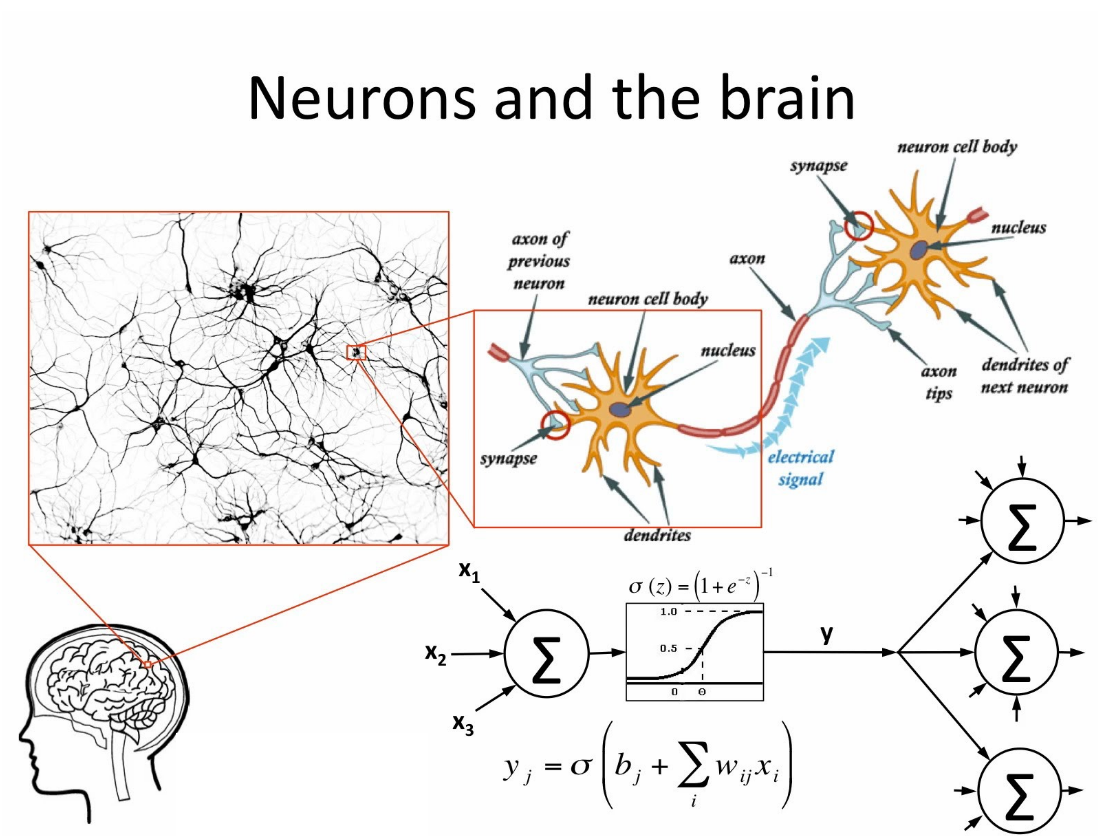

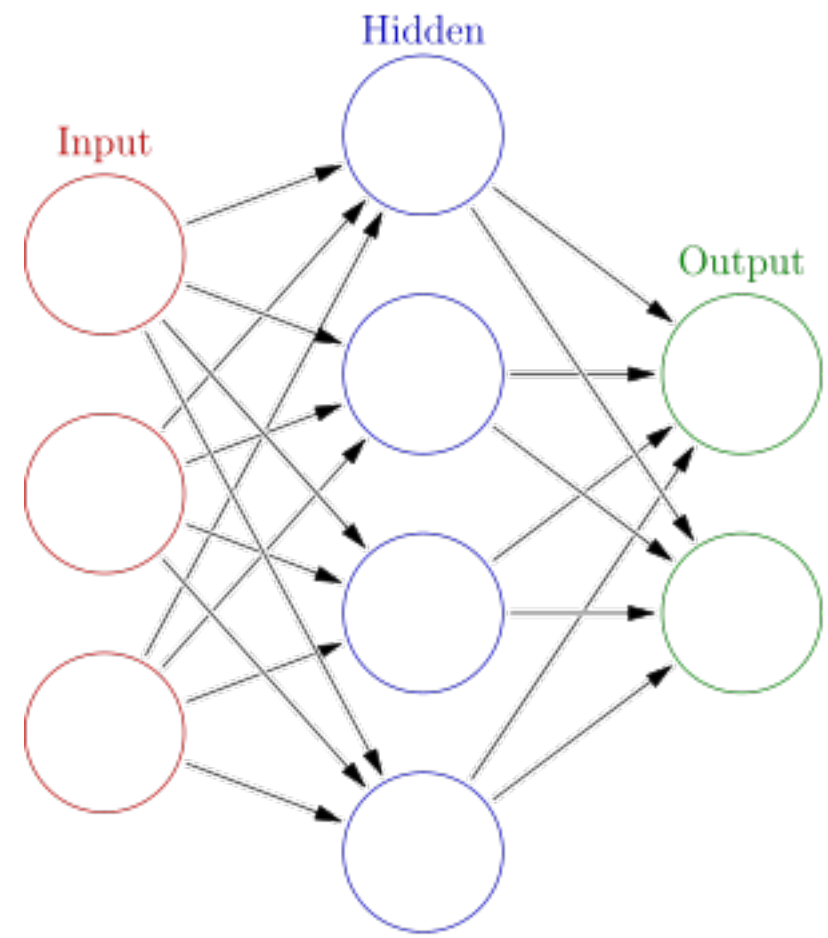

But how to find the correct parameters?*

*both continuous parameters like weights and discrete parameters of the architecture

How to measure model correctness? (and model "understanding"?)

*both continuous parameters like weights and discrete parameters of the architecture

Artificial Neural Network (ANN) Models of the Brain

Core (obvious) idea: Model brain systems with ANNs

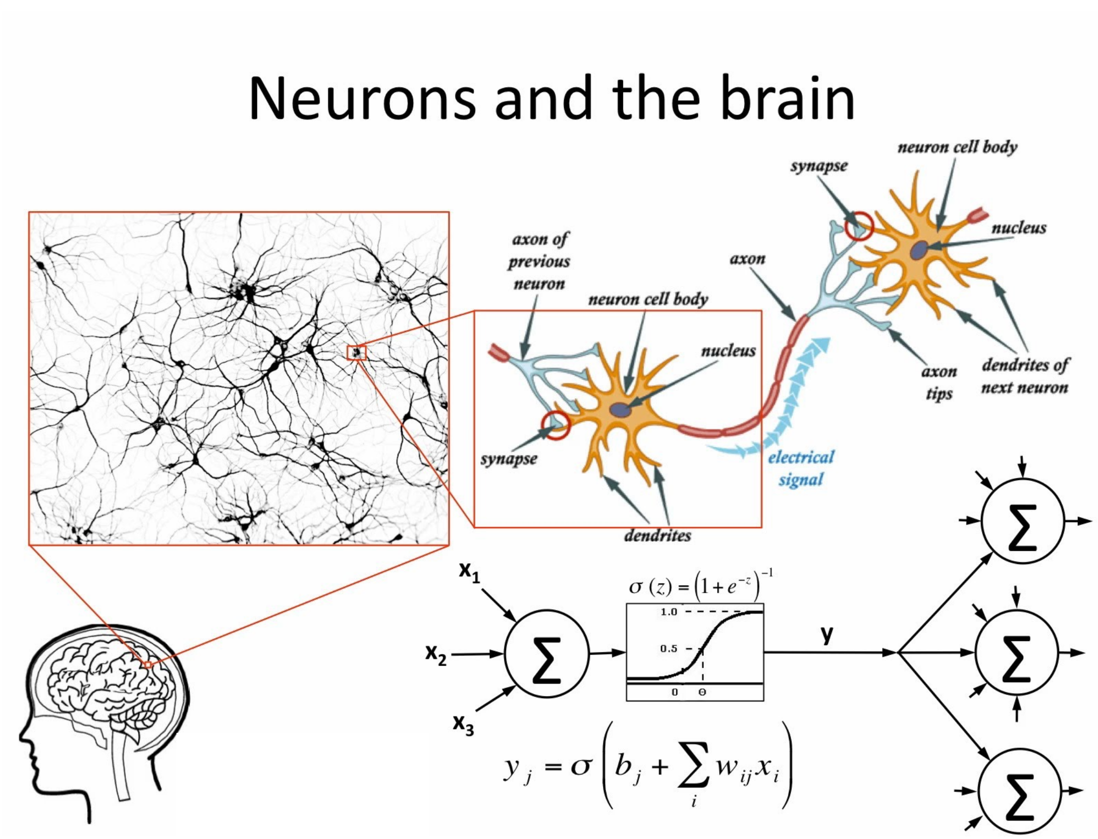

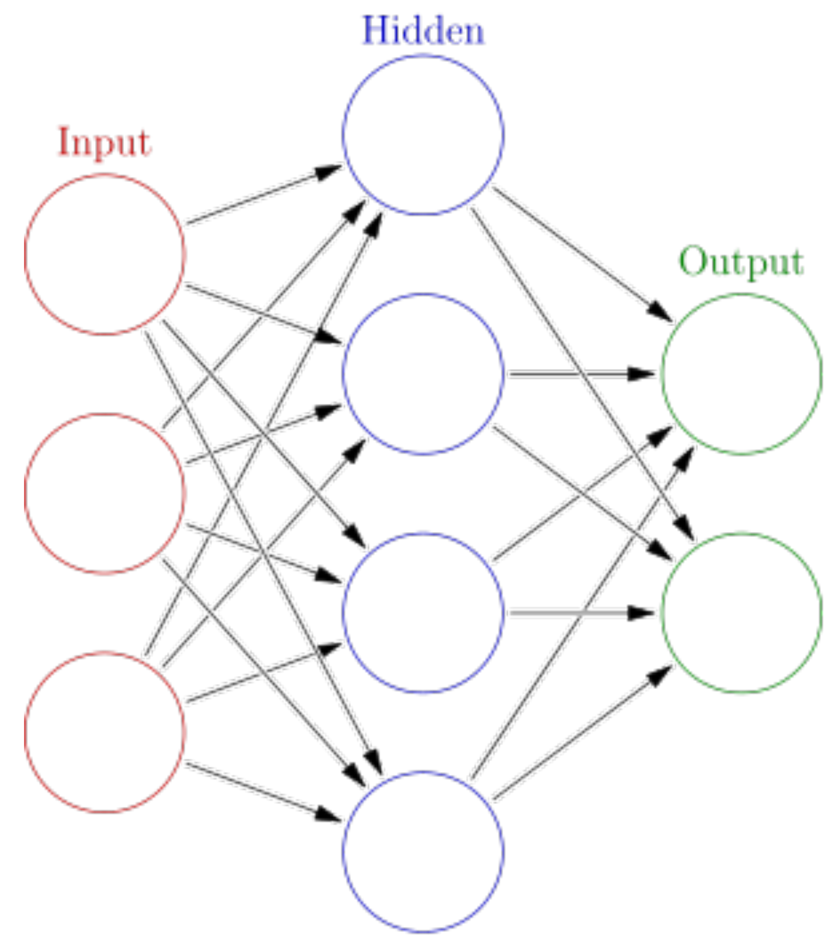

But how to find the correct parameters?*

*both continuous parameters like weights and discrete parameters of the architecture

How to measure model correctness? (and model “understanding”?)

This course is teach you how to do these things.

Artificial Neural Network (ANN) Models of the Brain

Core (obvious) idea: Model brain systems with ANNs

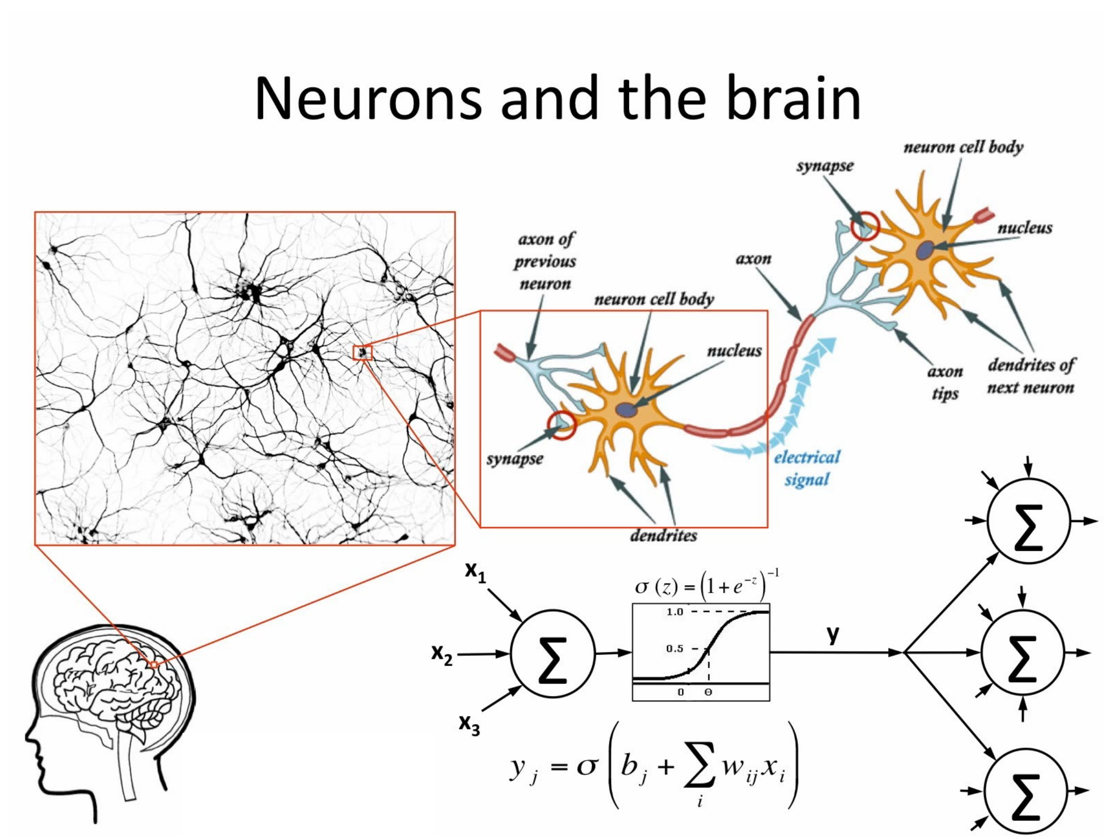

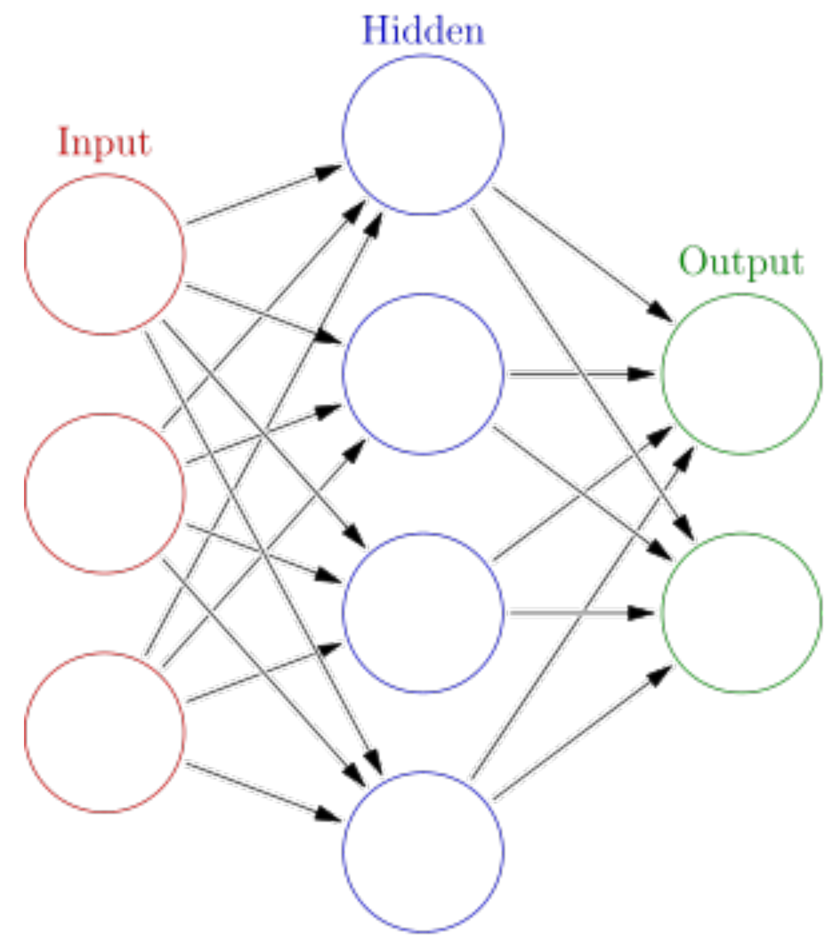

But how to find the correct parameters?

Case Study:

The Problem of Entity Extraction

Problem: Entity Extraction

Understanding complex, noisy data streams is a critical part of cognition.

Problem: Entity Extraction

Understanding complex, noisy data streams is a critical part of cognition.

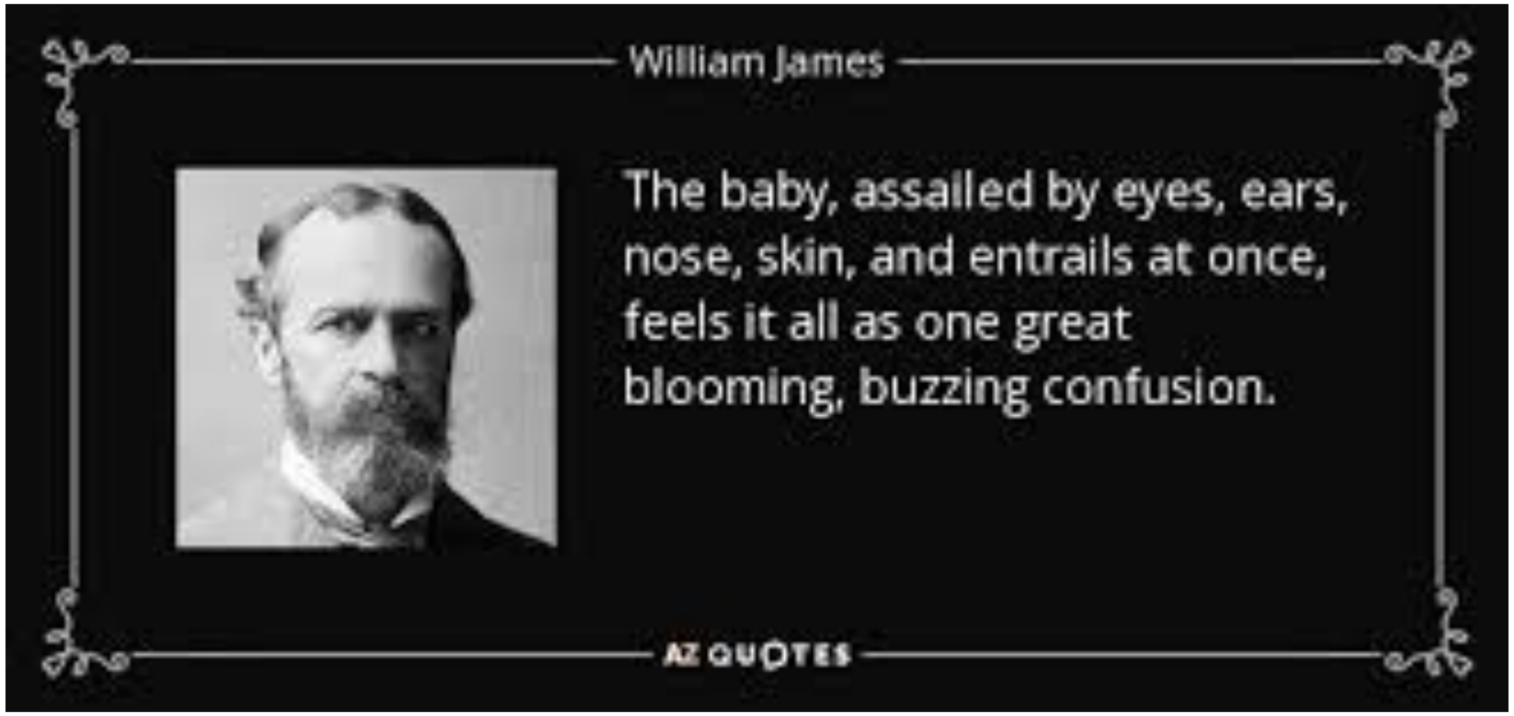

Without sophisticated parsing and entity extraction, the world would be "as one great blooming, buzzing confusion" (for babies or otherwise).

Problem: Entity Extraction

Understanding complex, noisy data streams is a critical part of cognition.

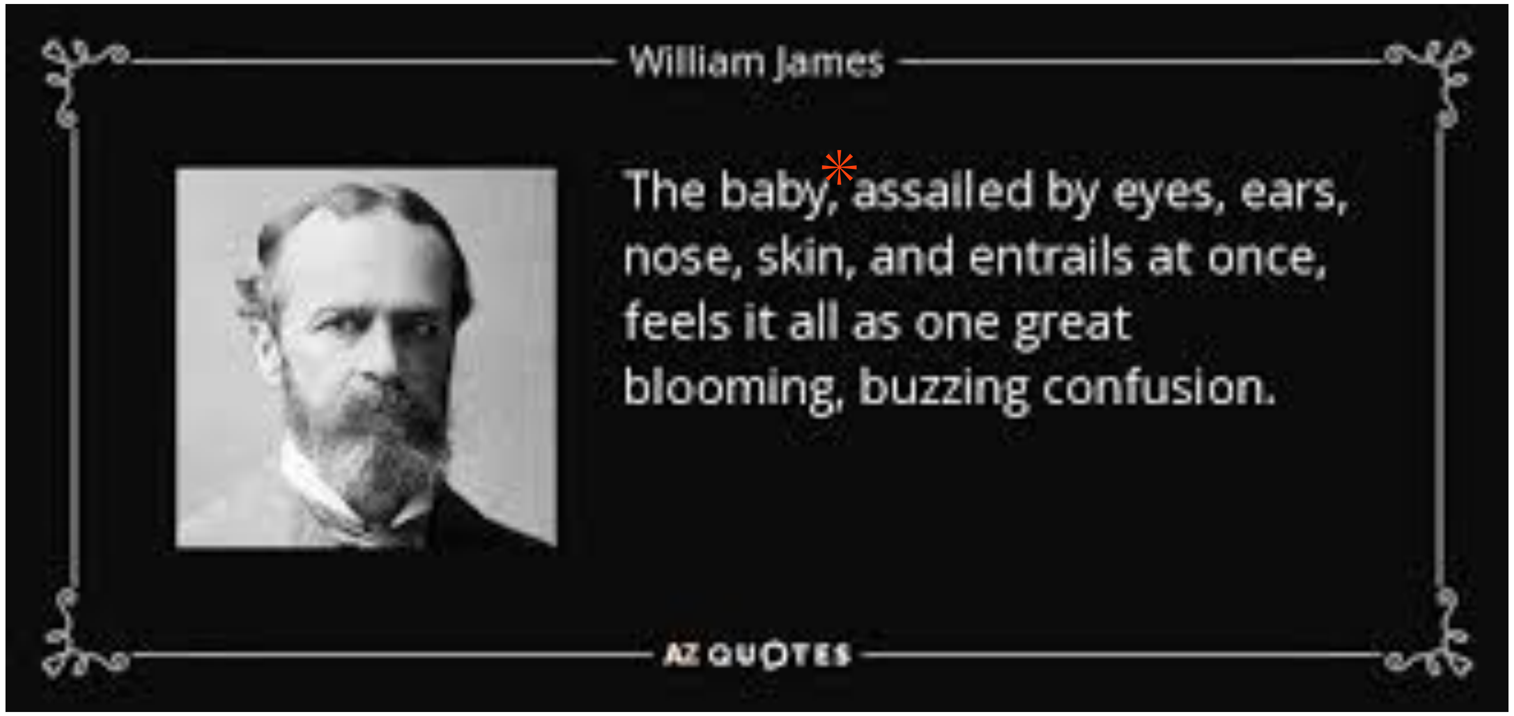

Without sophisticated parsing and entity extraction, the world would be "as one great blooming, buzzing confusion" (for babies or otherwise).

*actually not clearly true for babies . . .

Problem: Entity Extraction

Understanding complex, noisy data streams is a critical part of cognition.

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

Problem: Entity Extraction

View: position, size, pose, illumination

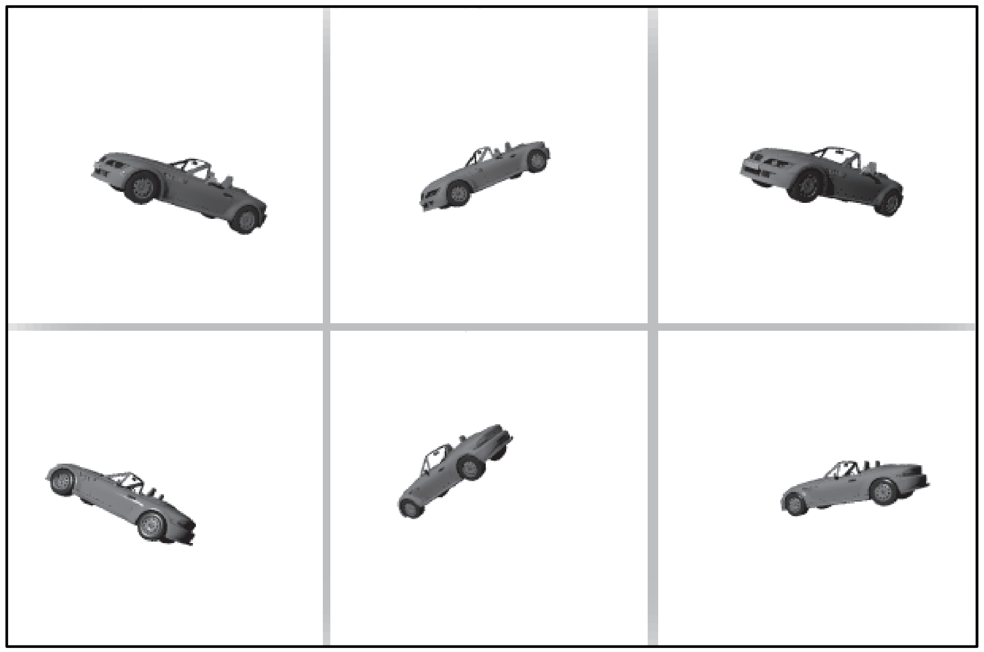

Background variation

Distortion & Noise

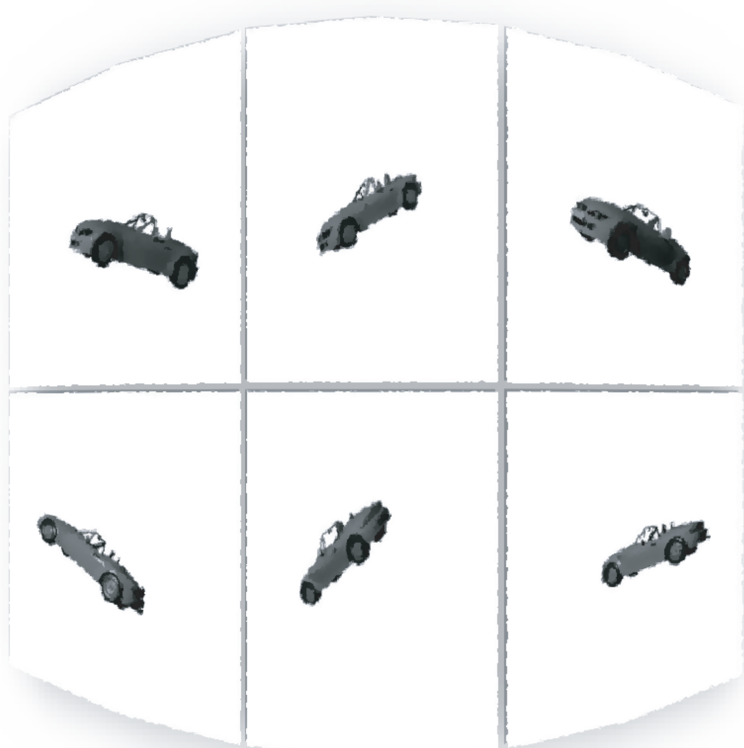

Geometric variation

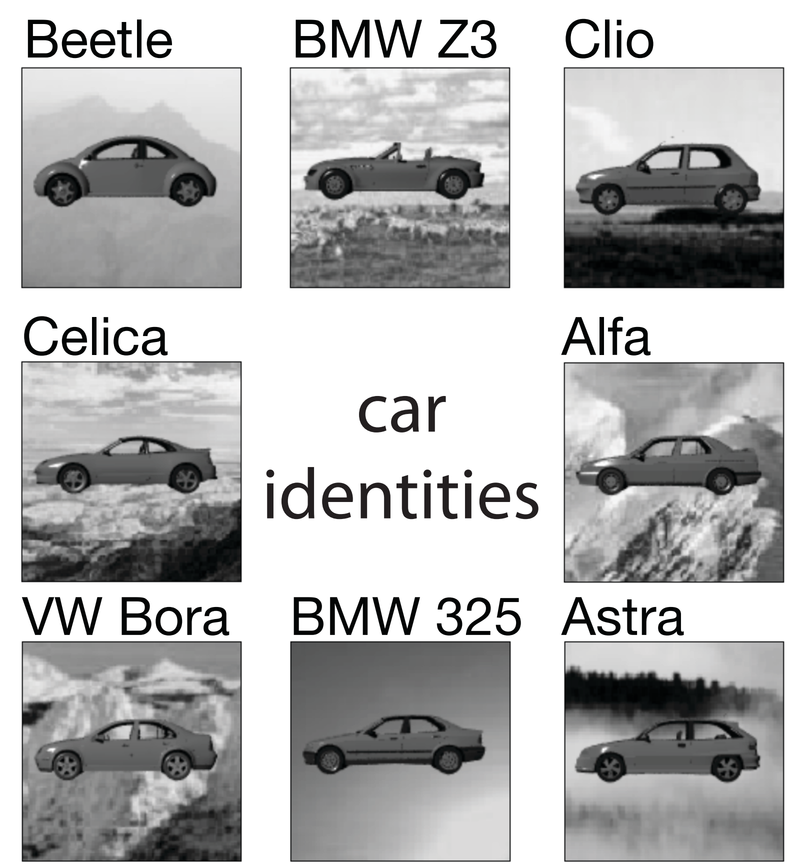

"Visual object recognition"

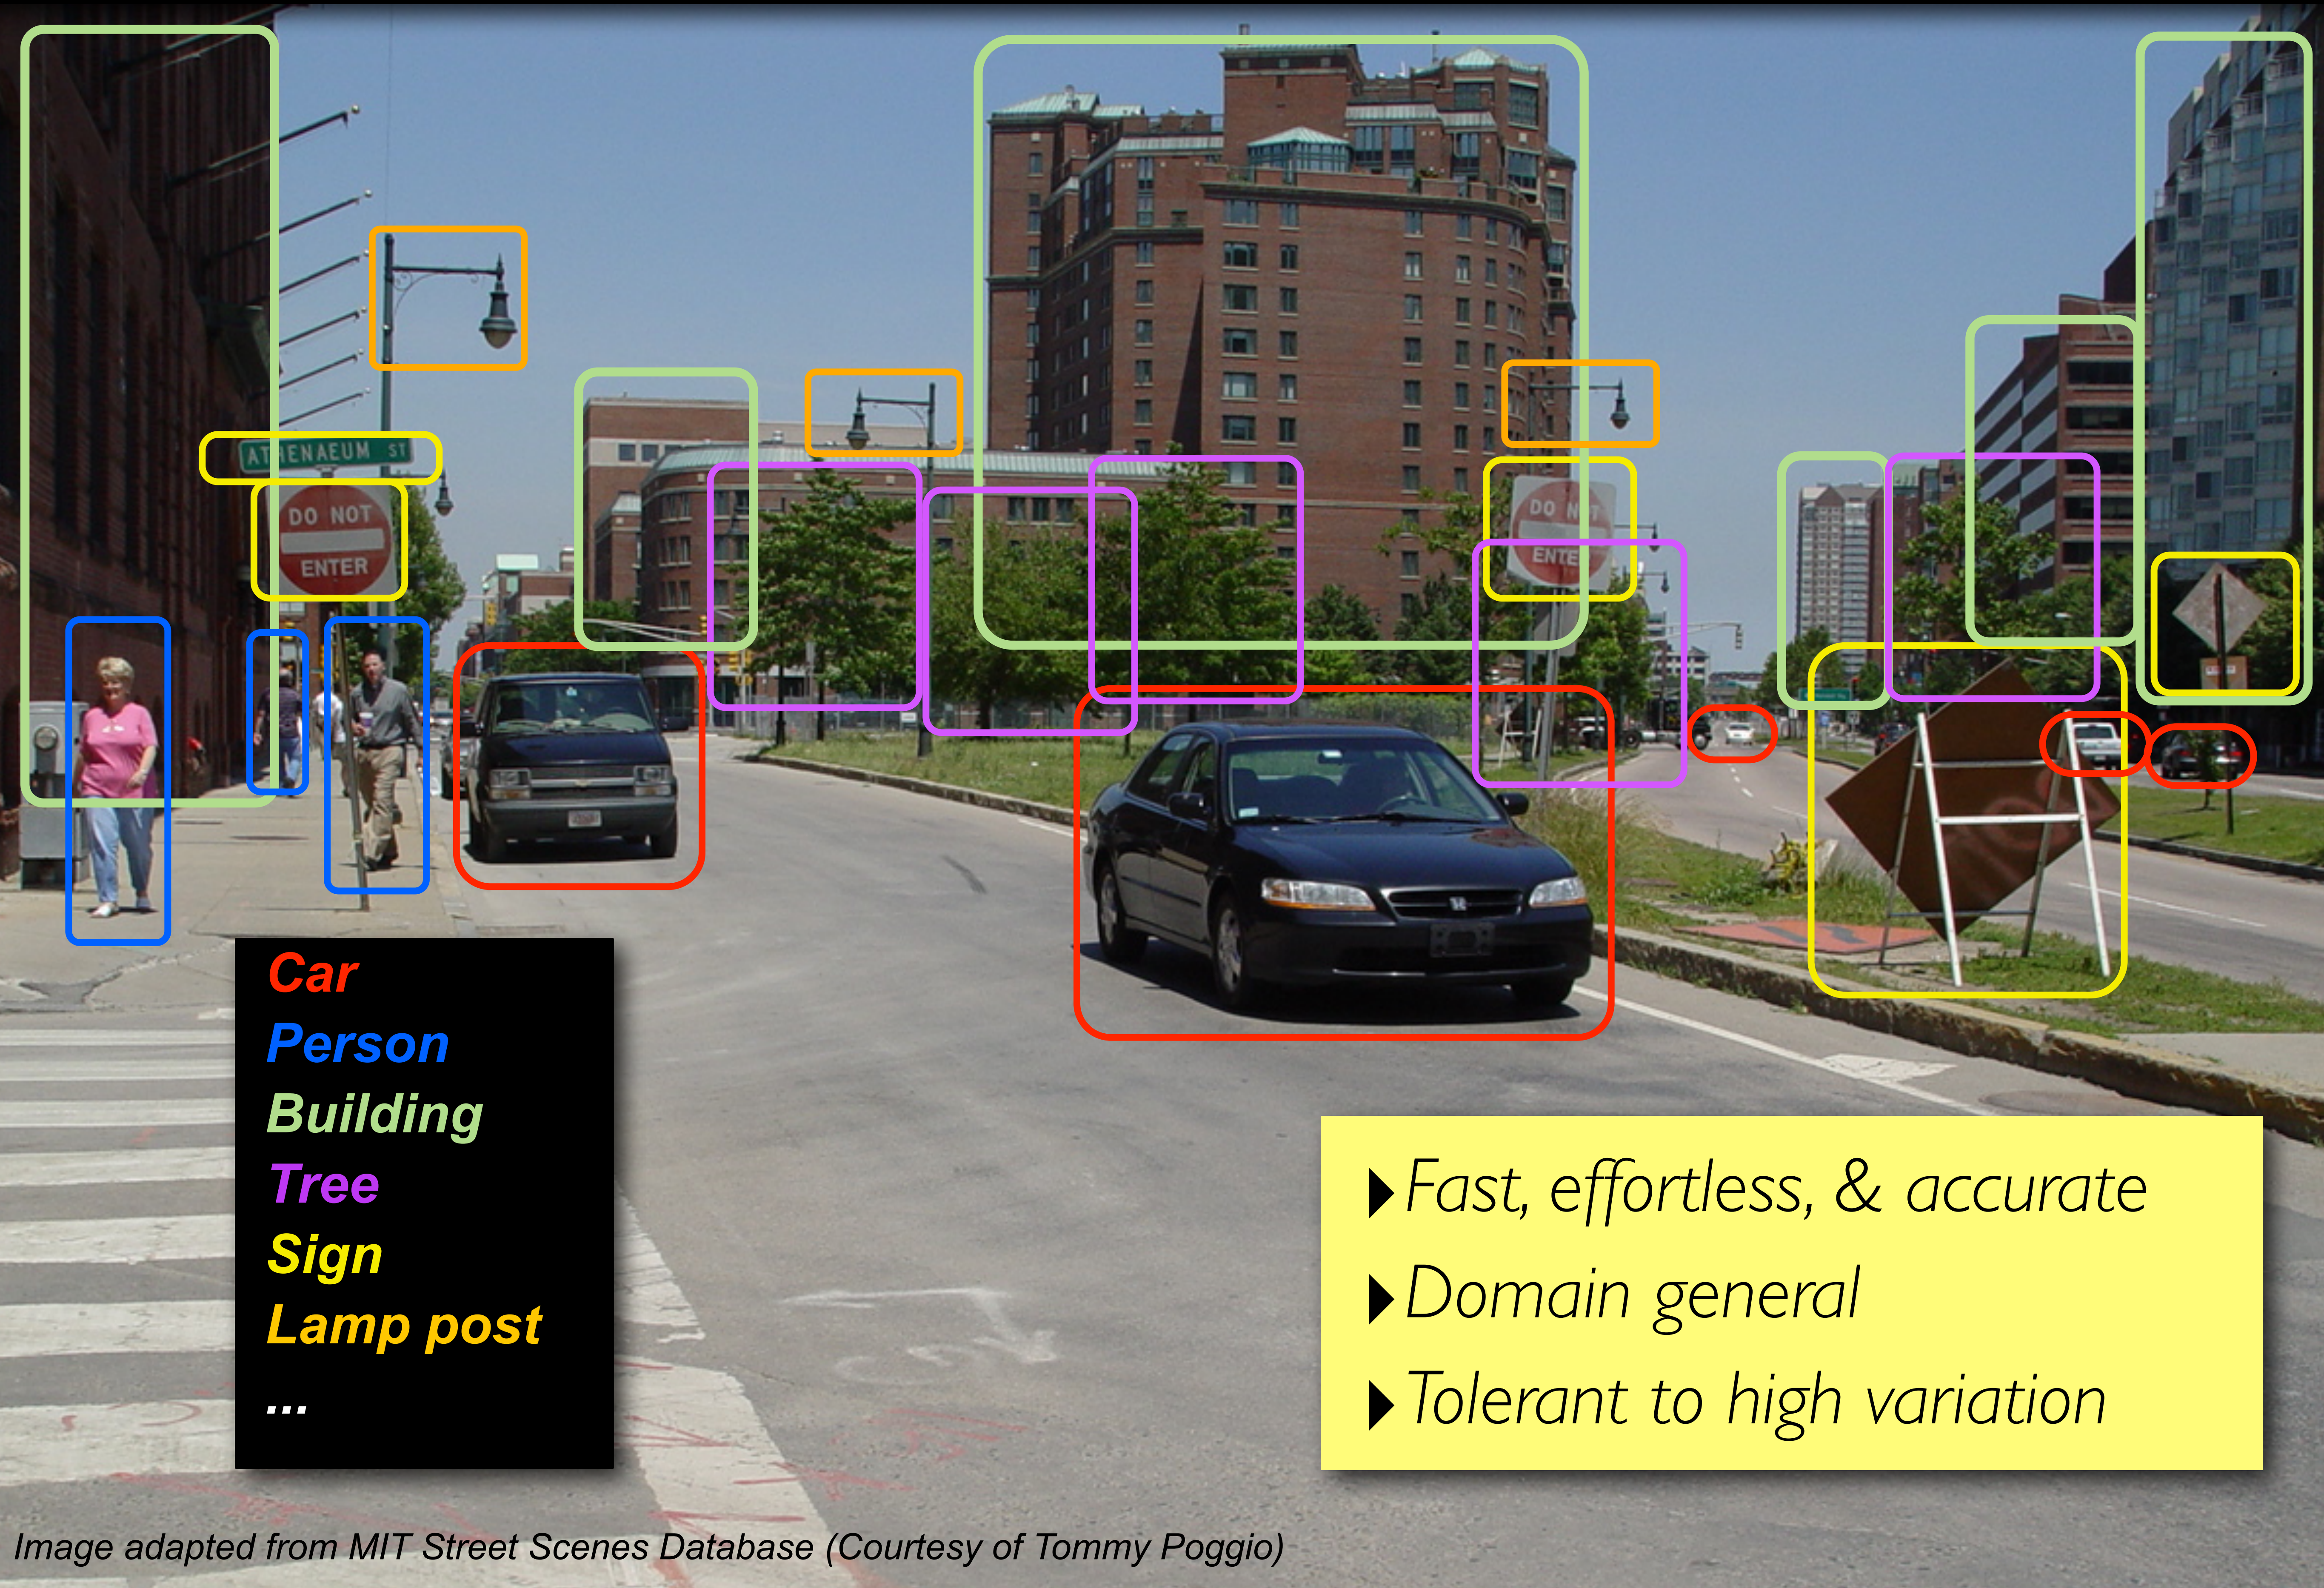

Car
Person
Building
Tree
Sign
Lamp post
...

- *Fast, effortless, & accurate*
- *Domain general*
- *Tolerant to high variation*

Image adapted from MIT Street Scenes Database (Courtesy of Tommy Poggio)

Problem: Entity Extraction

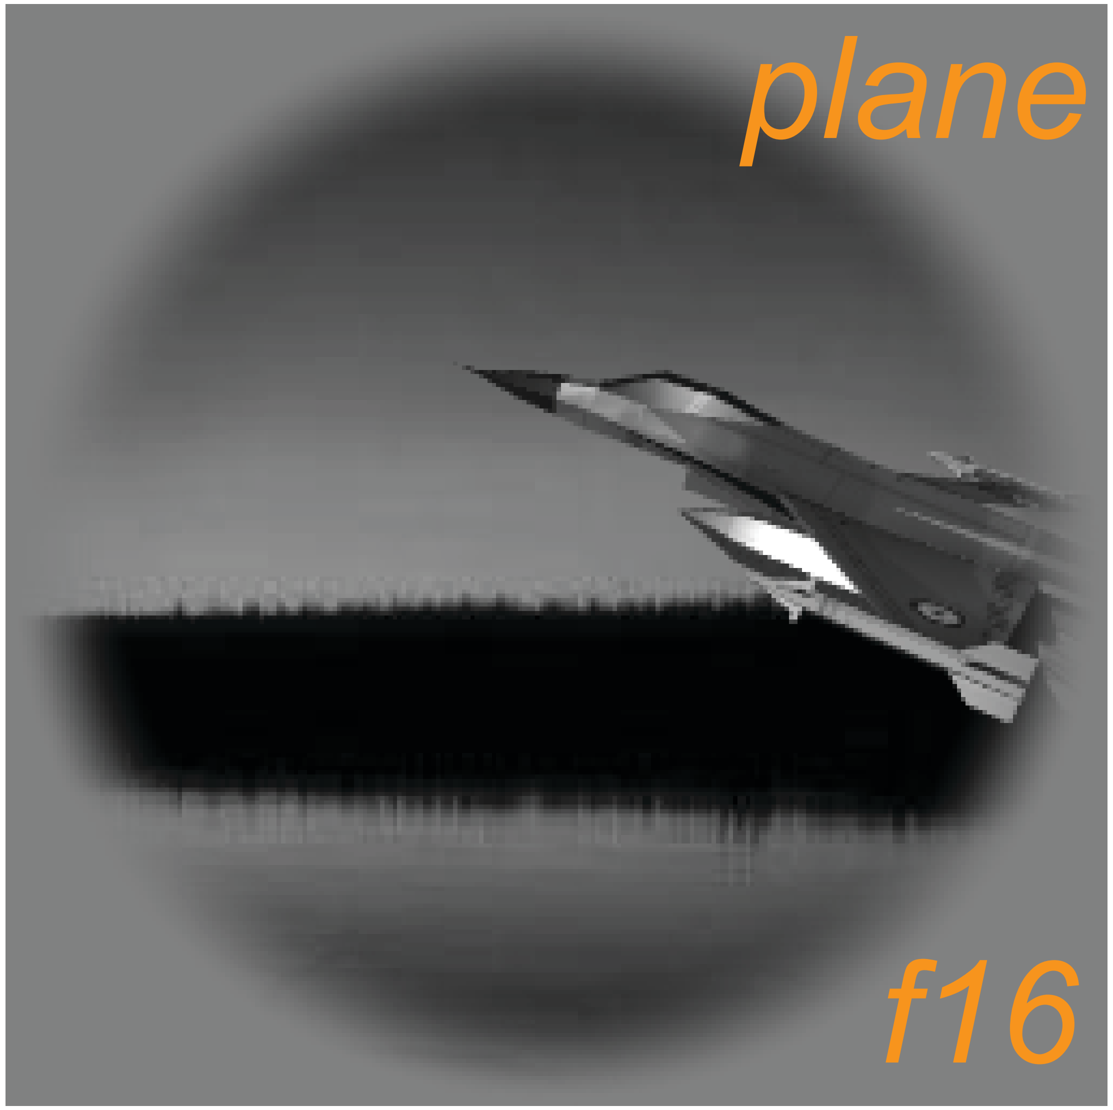

Category

Identity

from Hong et al (2016)

Problem: Entity Extraction

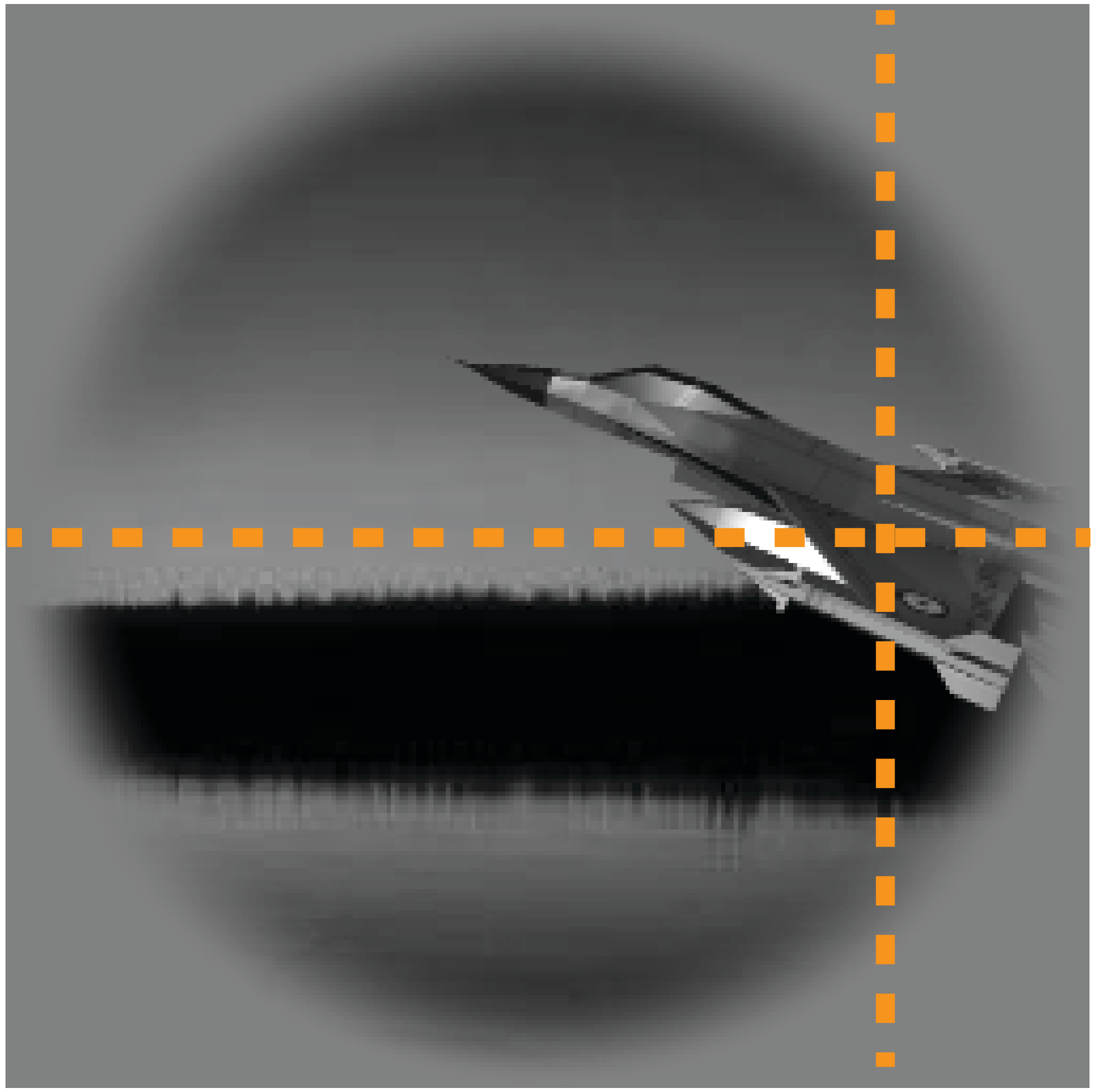

Position

from Hong et al (2016)

Problem: Entity Extraction

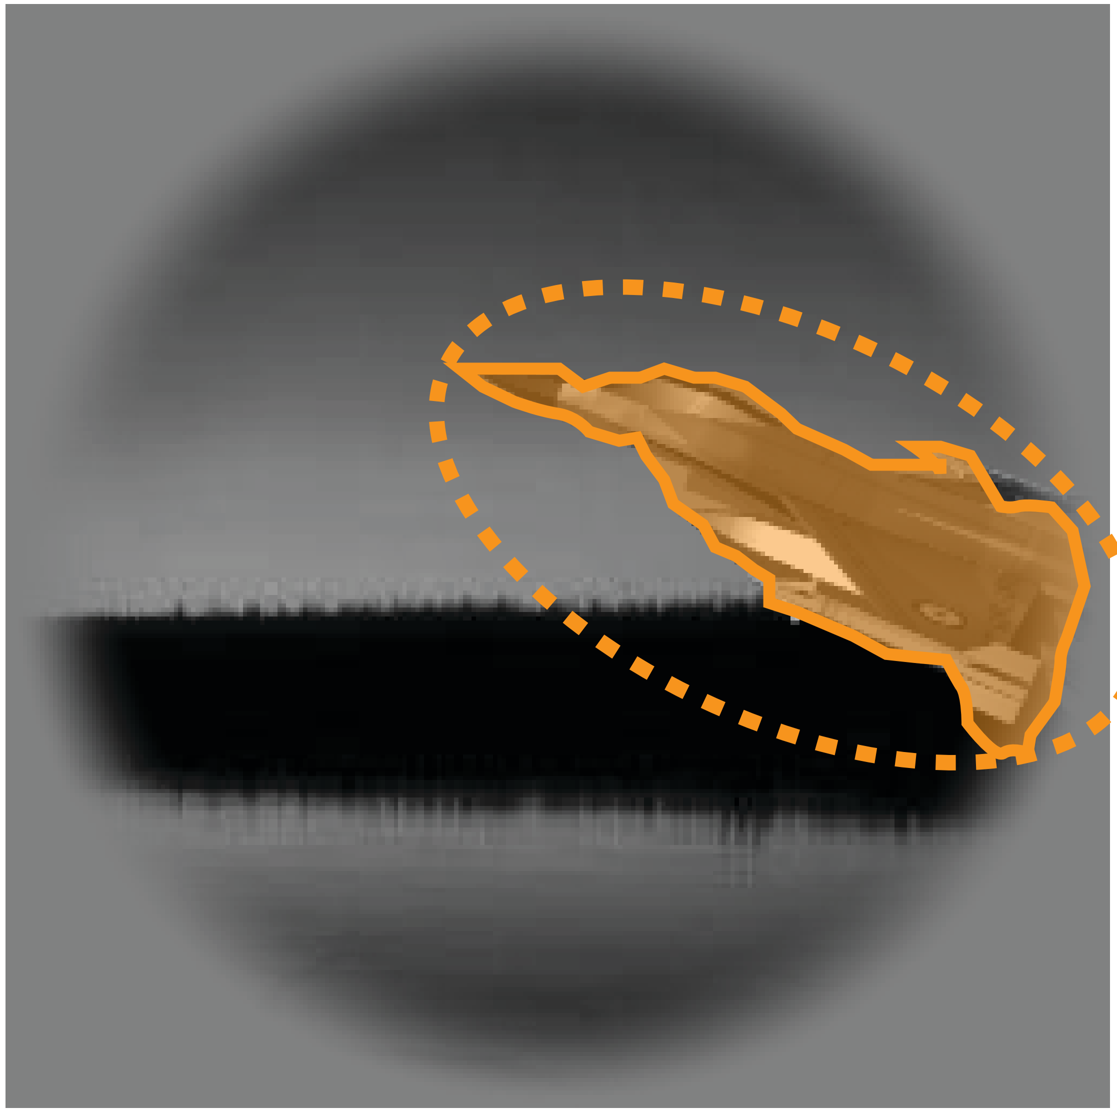

Size

from Hong et al (2016)

Problem: Entity Extraction

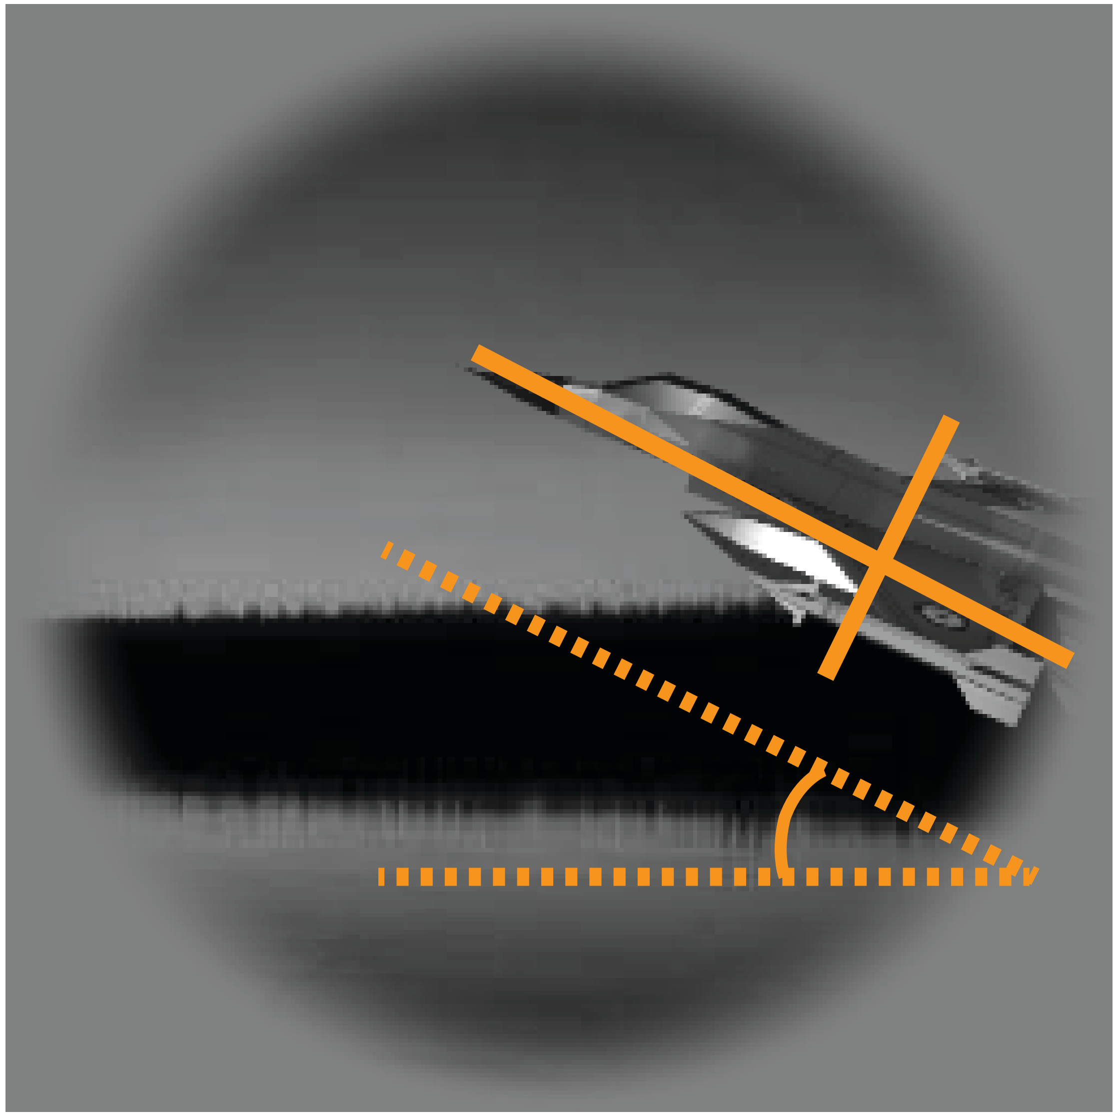

Aspect Ratio and Angle

from Hong et al (2016)

Problem: Entity Extraction

We can quickly assess the scene as a whole.

Category

Identity

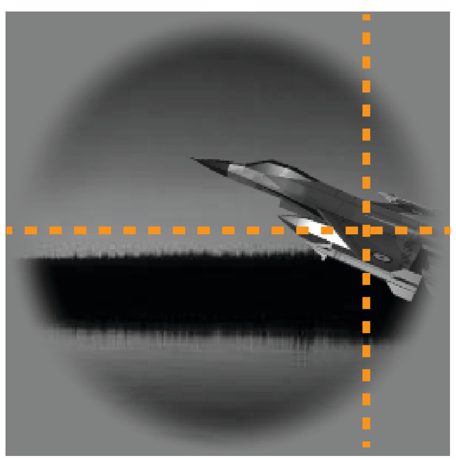

X and Y Axis Position

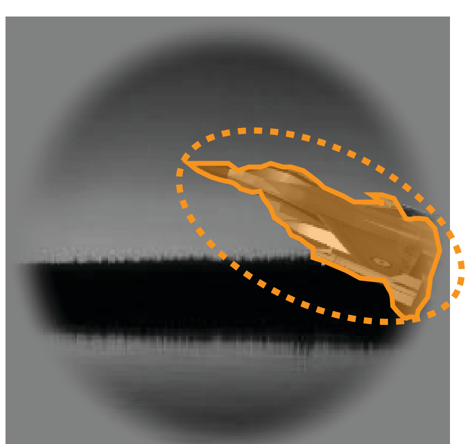

Perimeter

2-D Retinal Area

3-D Object Scale

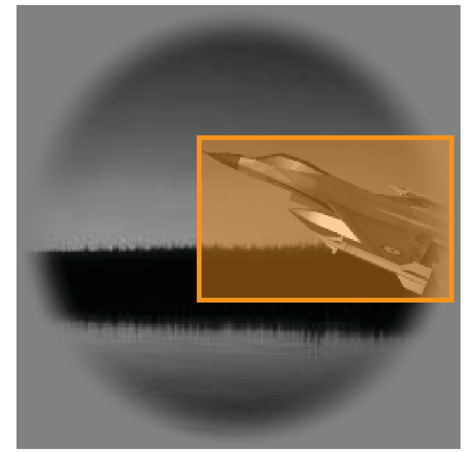

Bounding Box

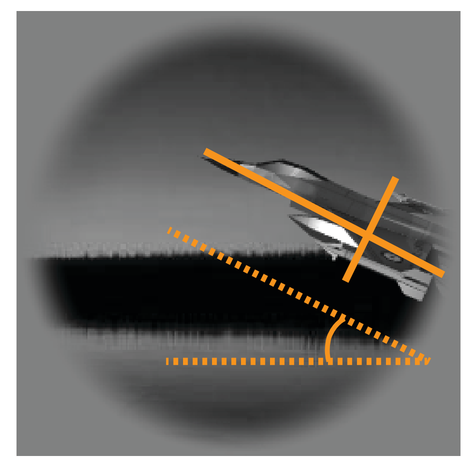

Aspect Ratio

Major Axis Length

Major Axis Angle

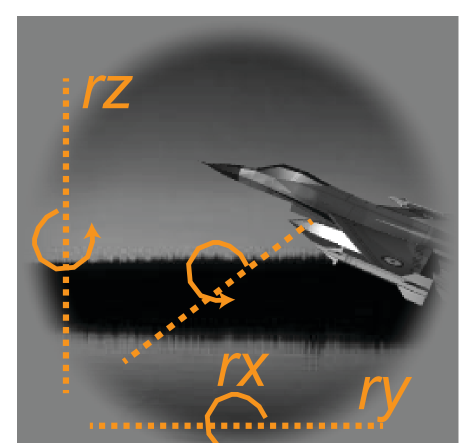

Pose in each axis

from Hong et al (2016)

Problem: Entity Extraction

Understanding complex, noisy data streams is a critical part of cognition.

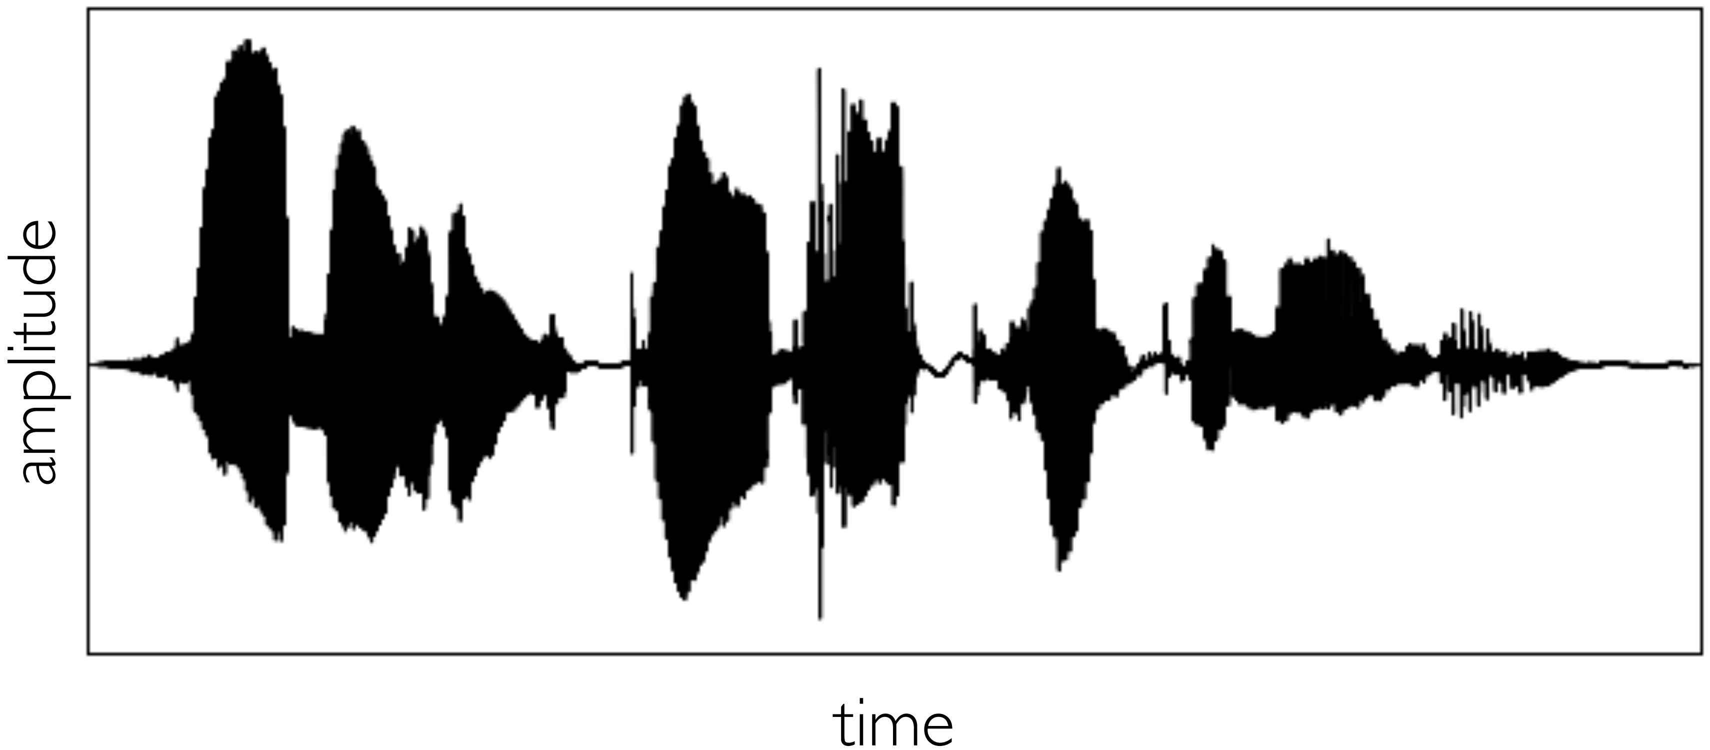

Problem: Entity Extraction

Understanding complex, noisy data streams is critical part of cognition.

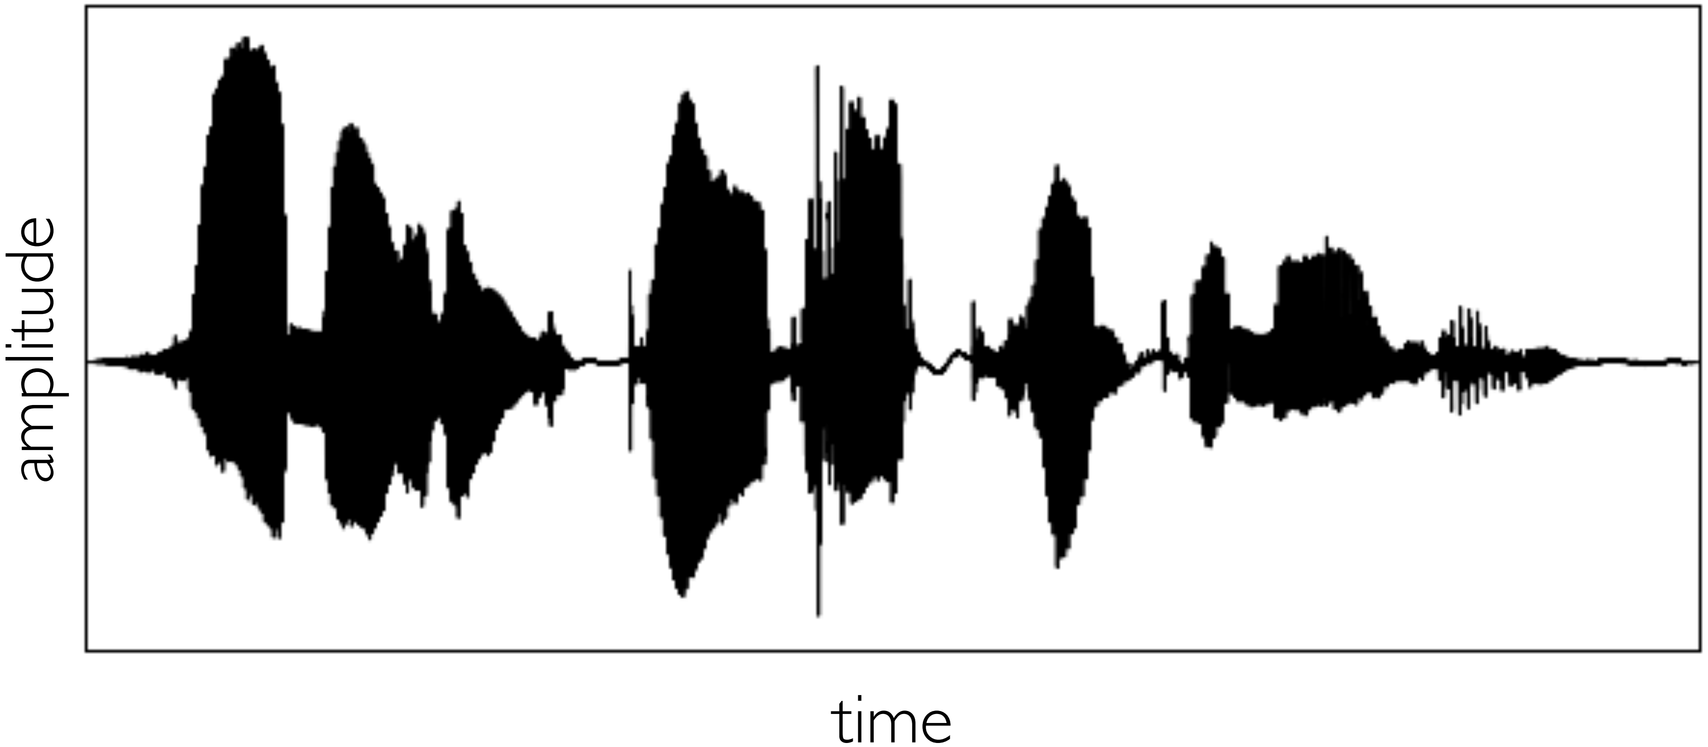

"Hannah is good at compromising."

Problem: Entity Extraction

Understanding complex, noisy data streams is critical part of cognition.

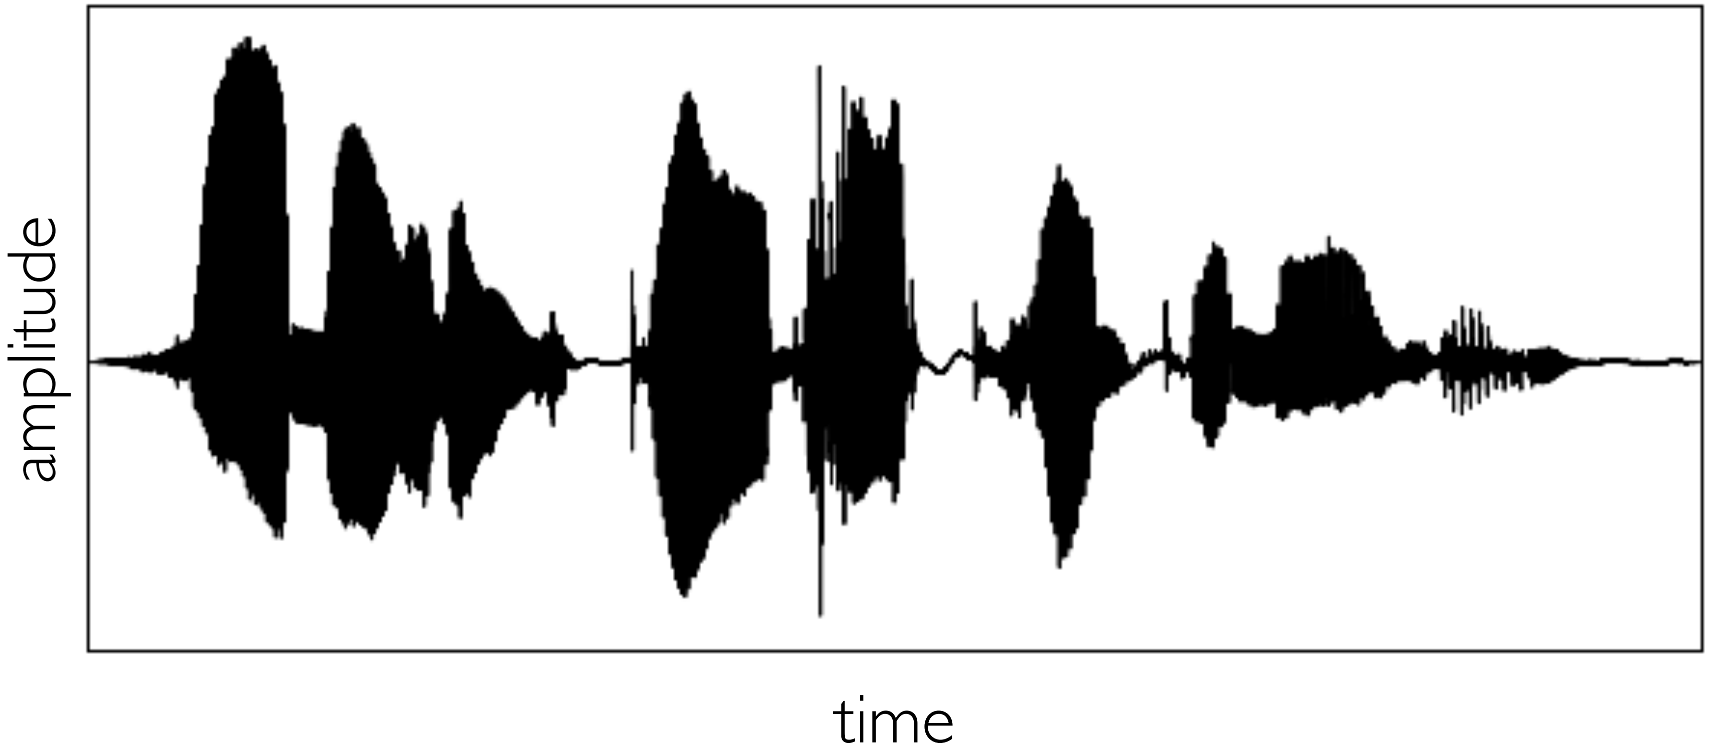

"Hannah is good at compromising."

variation sources: speaker identity

background noise

reverberation

...

“Explicit” vs. “Implicit” representations

A working definition of an “explicit” representation = a basis in which a problem is linearly separable

Explicit representation

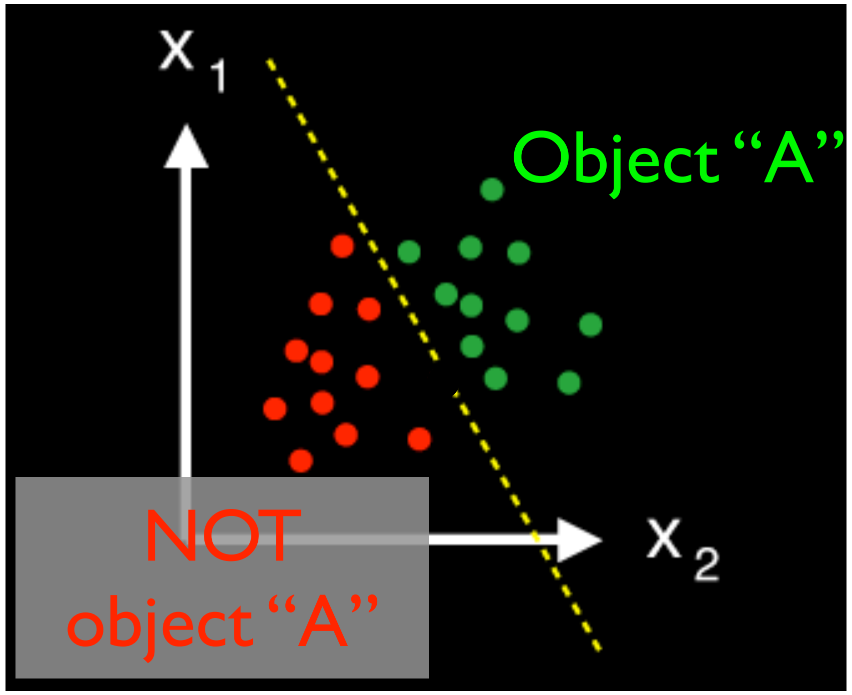

Linearly separable

Implicit representation

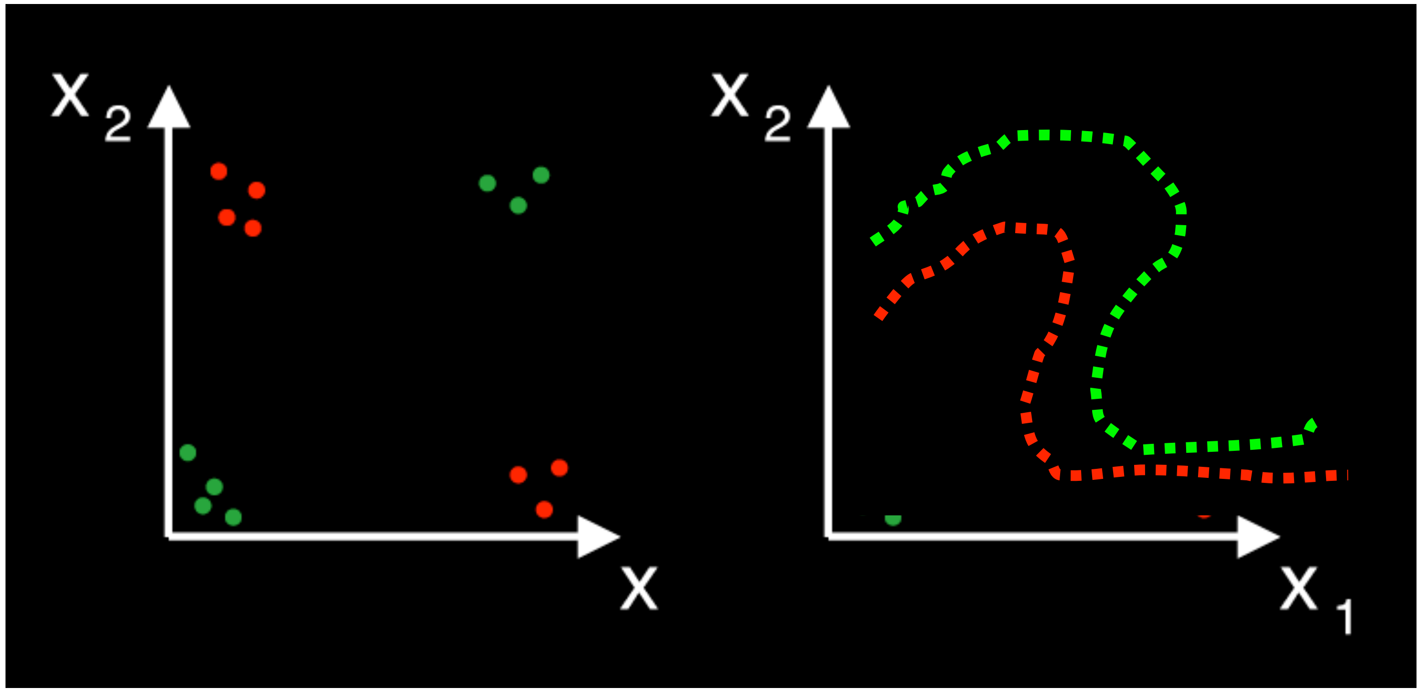

Not linearly separable

The same concept applies to higher dimensional spaces

The Computational Crux of the problem

You need SELECTIVITY for different objects

You need TOLERANCE to changes in the retinal image

Computationally easy (e.g. templates)	Computationally easy (e.g. simply integrate)	BOTH -- computationally difficult!

The Computational Crux of the problem

Any population representation

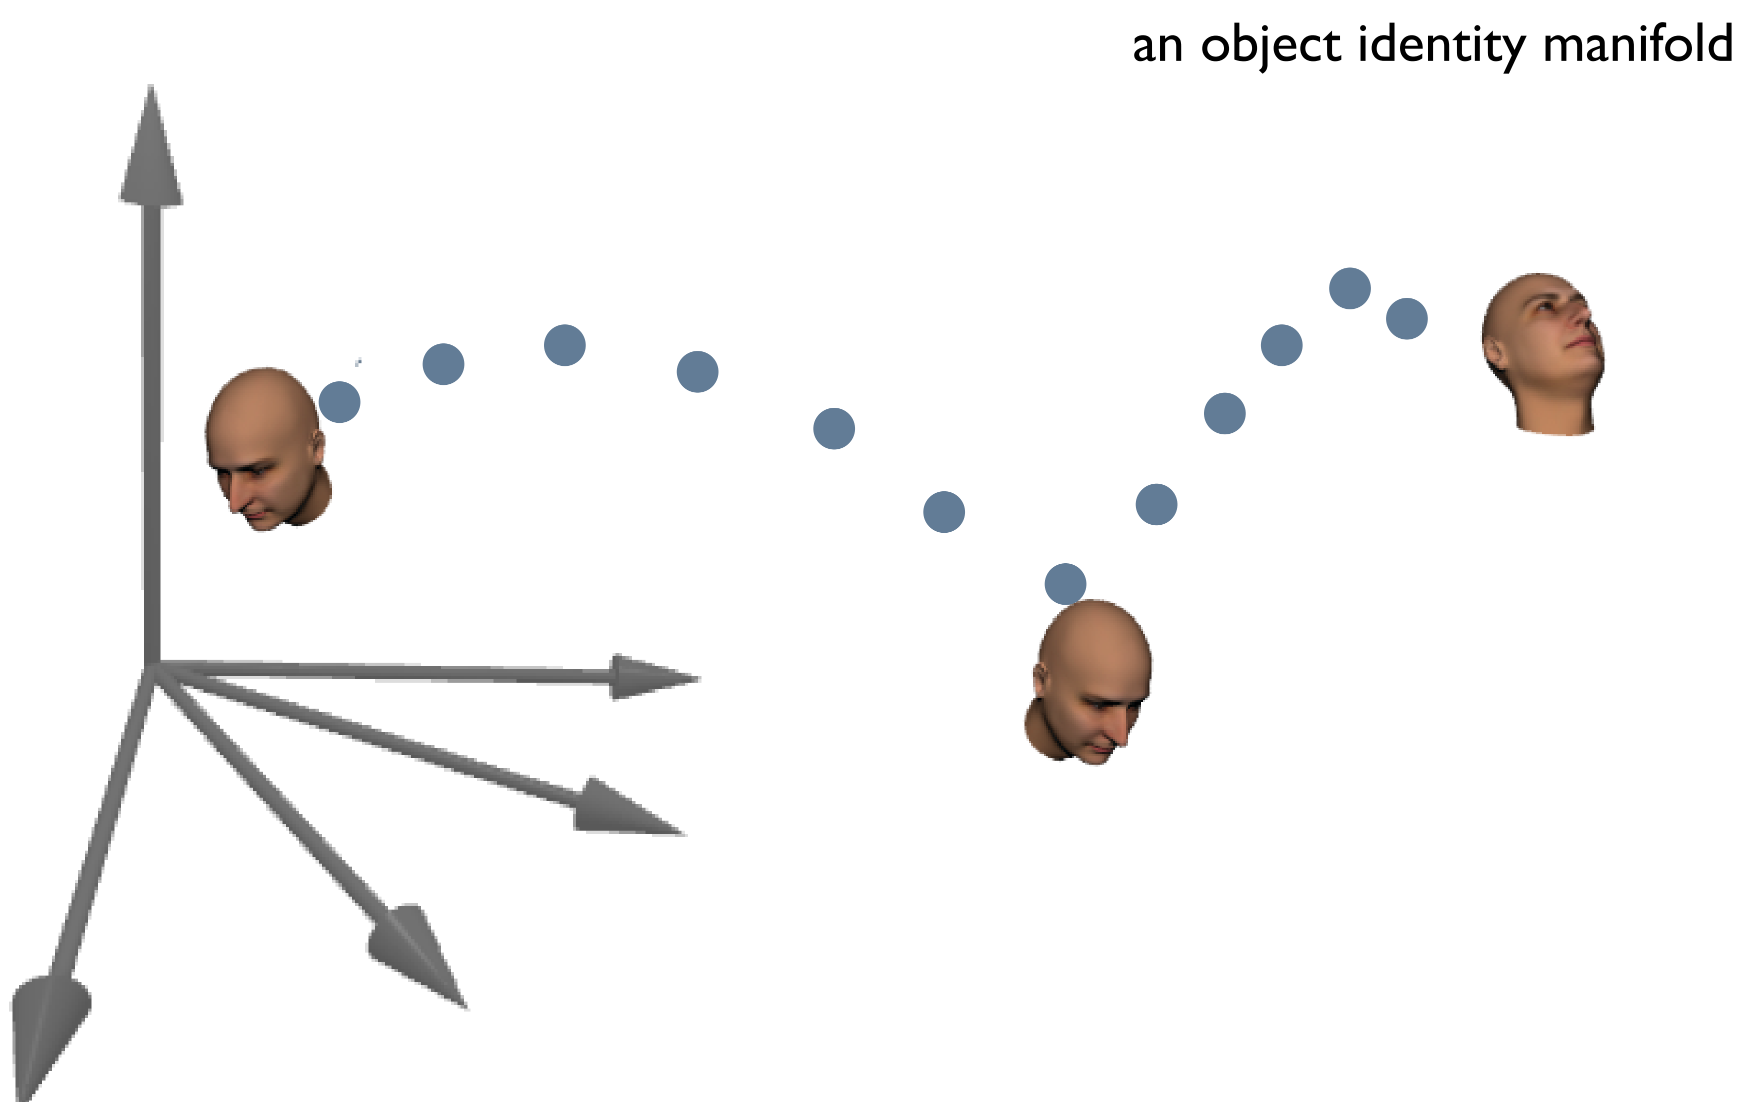

The Computational Crux of the problem

A "good" population representation

Explicit object representation

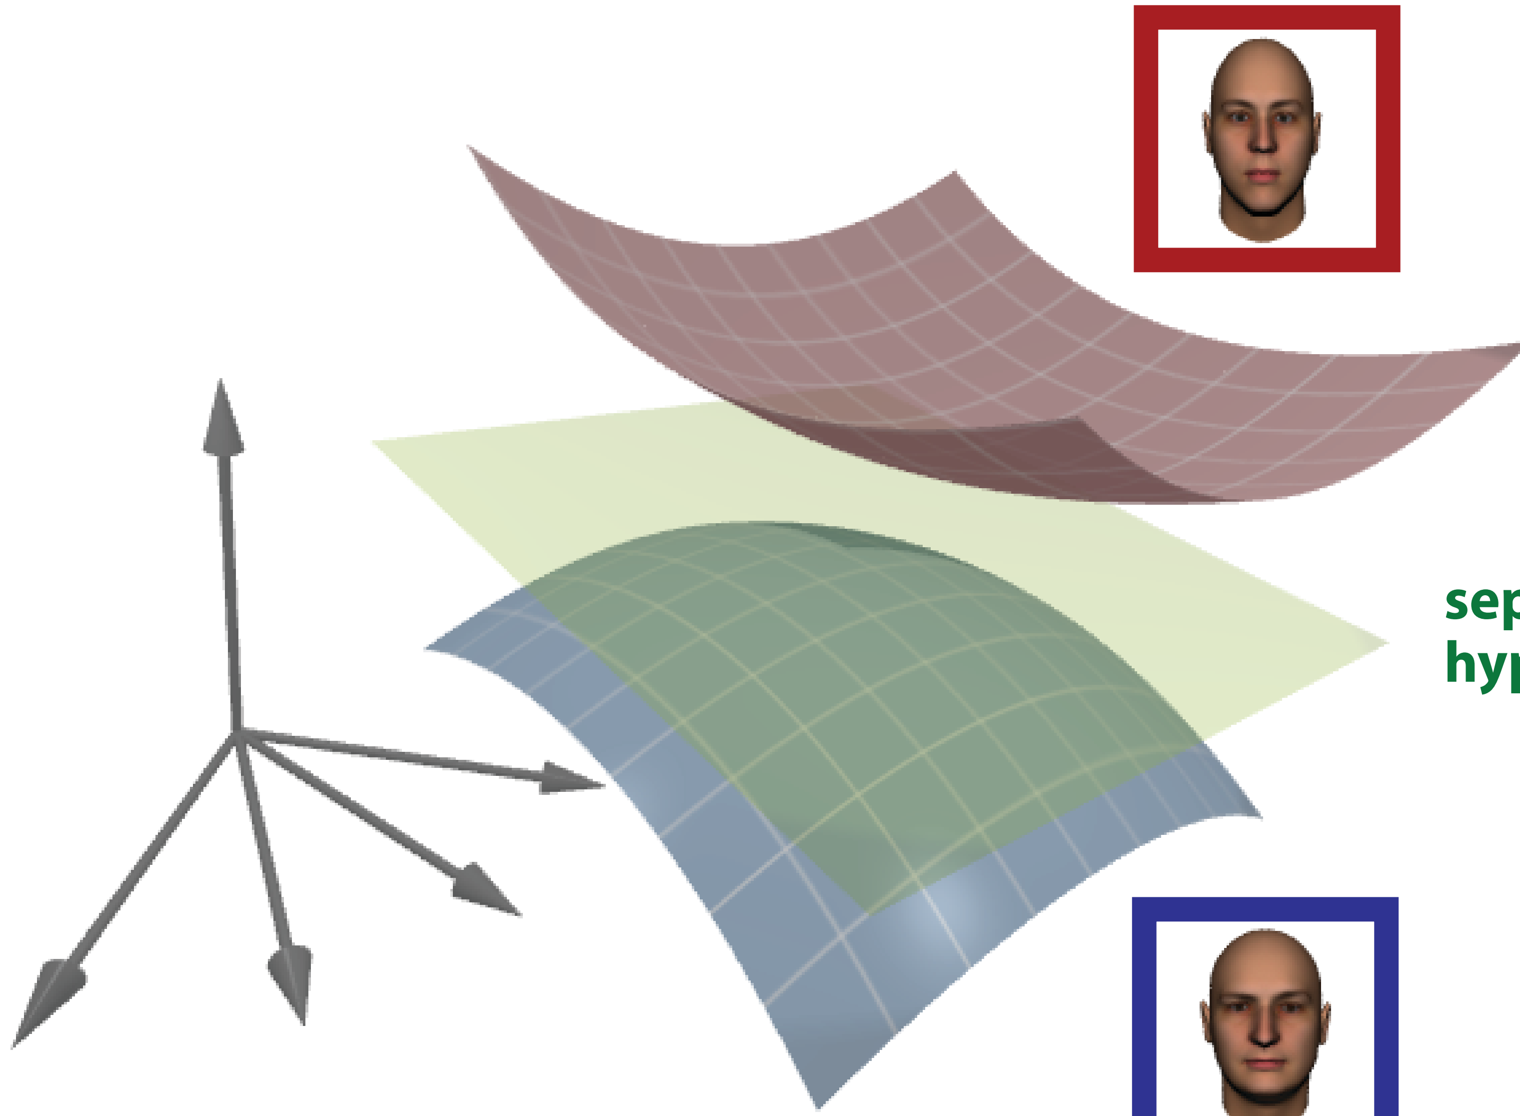

*DiCarlo and Cox, **TICS** (2007)*

The Computational Crux of the problem

A "bad" population representation

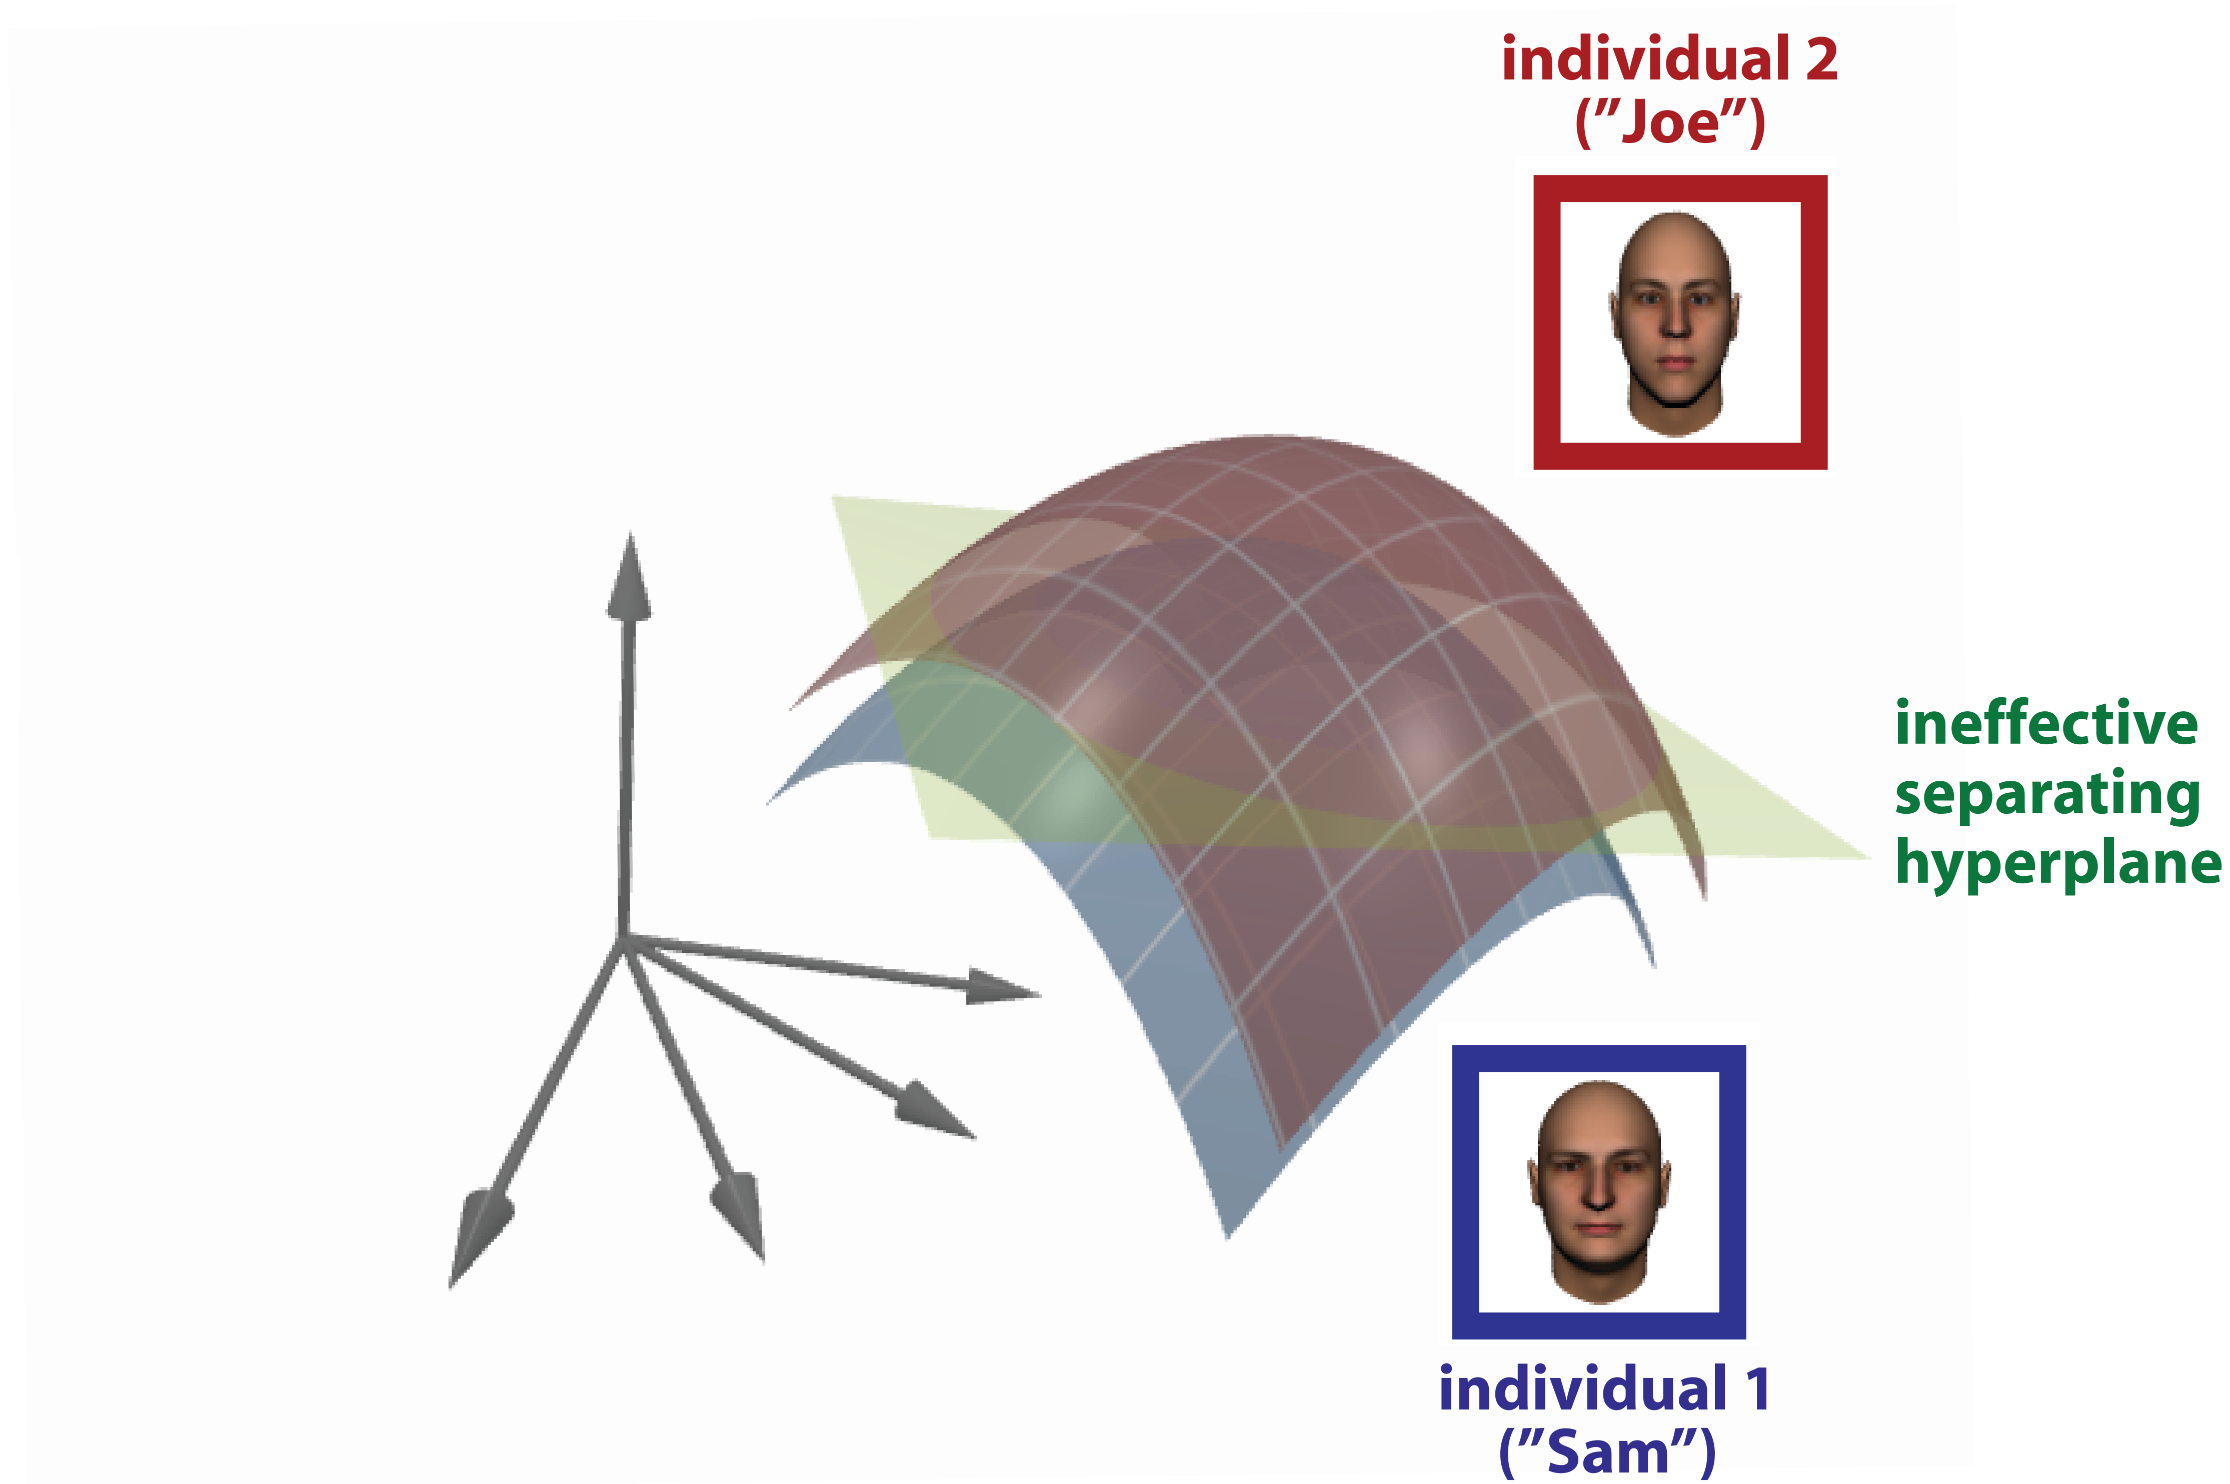

*DiCarlo and Cox, **TICS** (2007)*

The Tangling of Object Manifolds

Actual pixel representation

(~ retinal image representation)

Implicit object representation

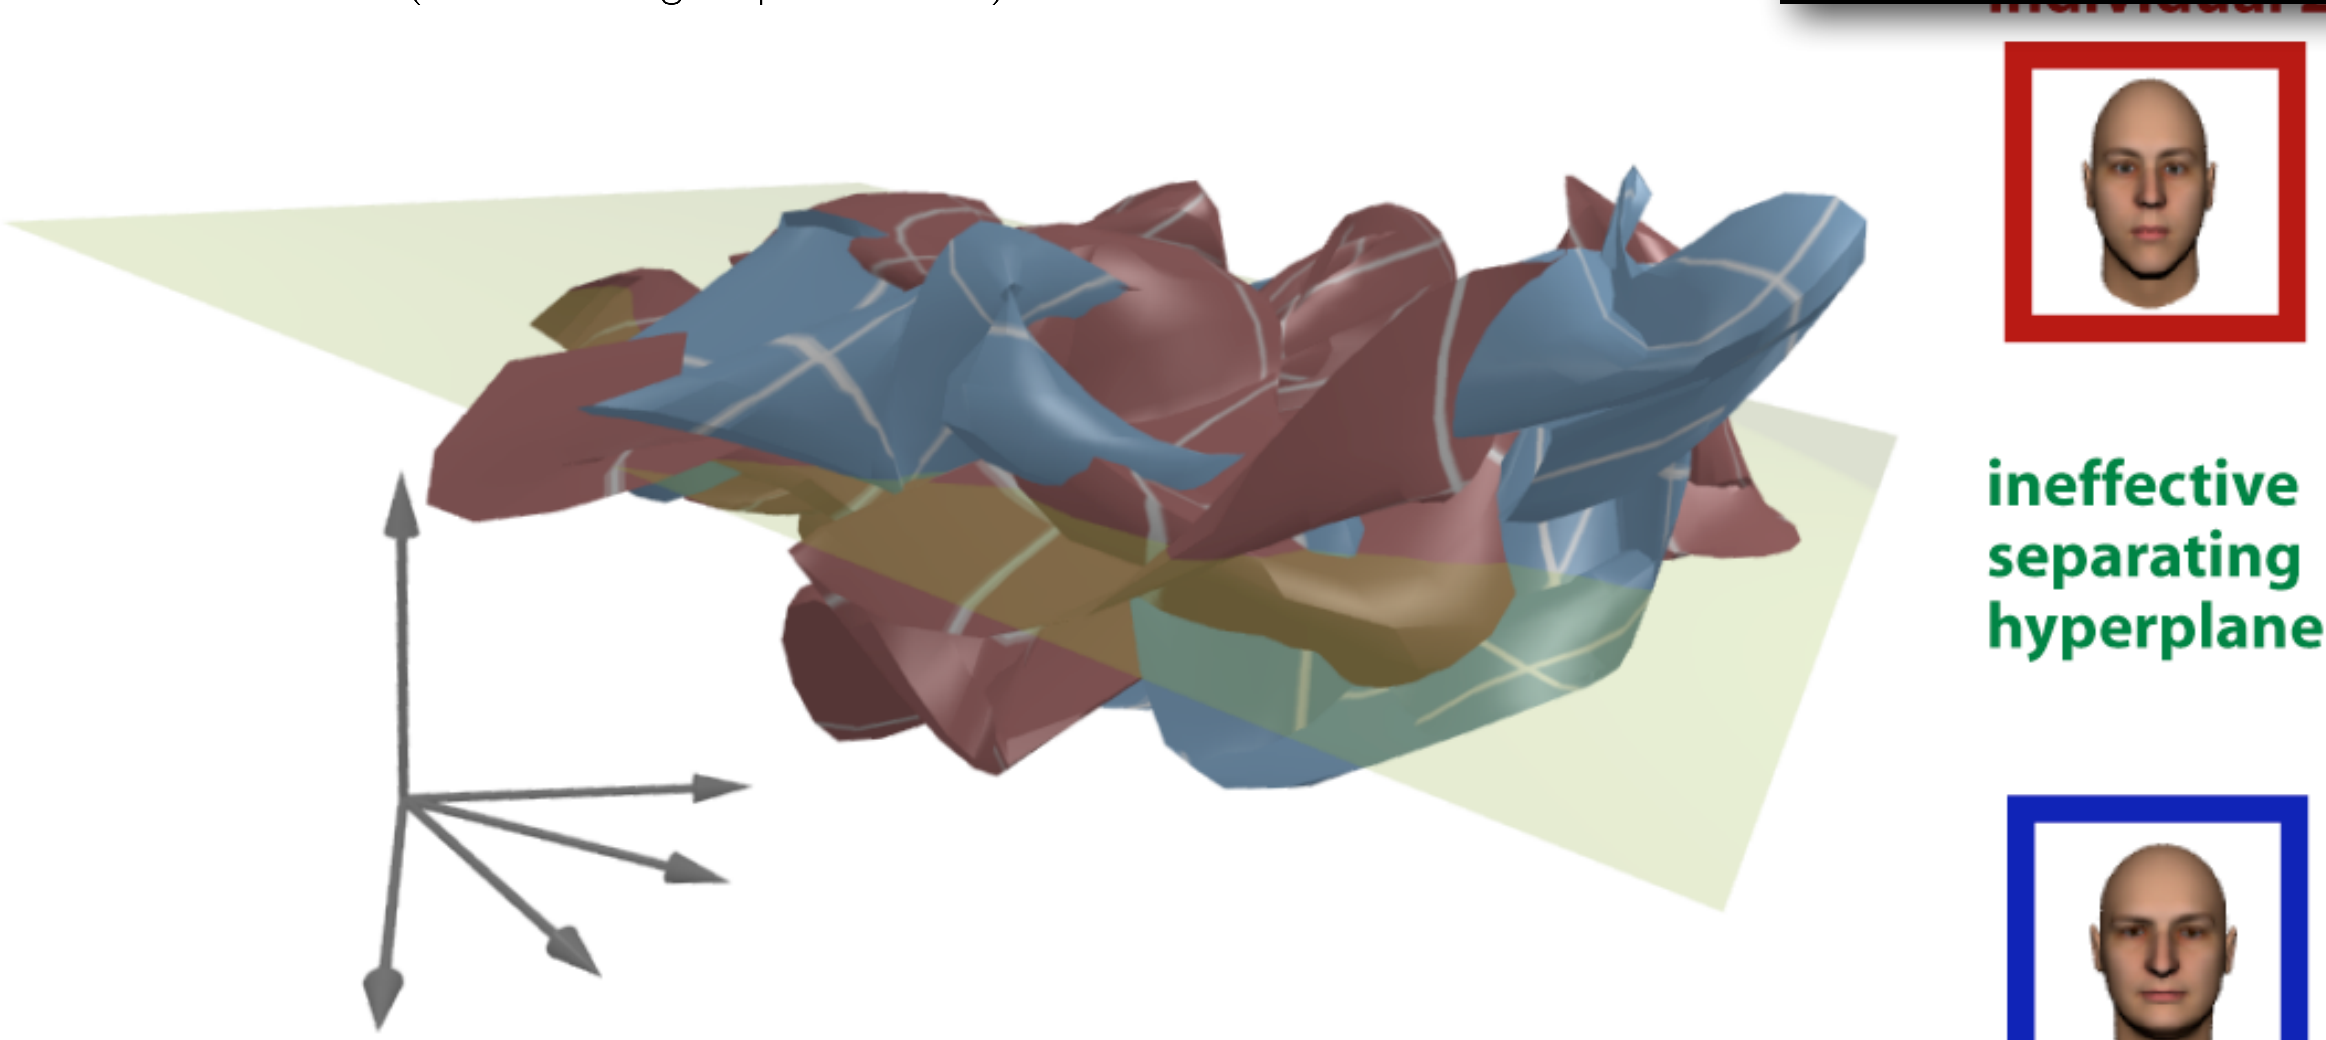

object manifolds are "tangled"

(Due to identity-preserving image variation.)

*DiCarlo and Cox, **TICS** (2007); Pinto, Cox, and DiCarlo, **PLoS Comp Bio** (2008)*

Problem: Entity Extraction

Axes of natural variation of natural
"physics" representation of world

e.g.

retinal photoreceptor voltage

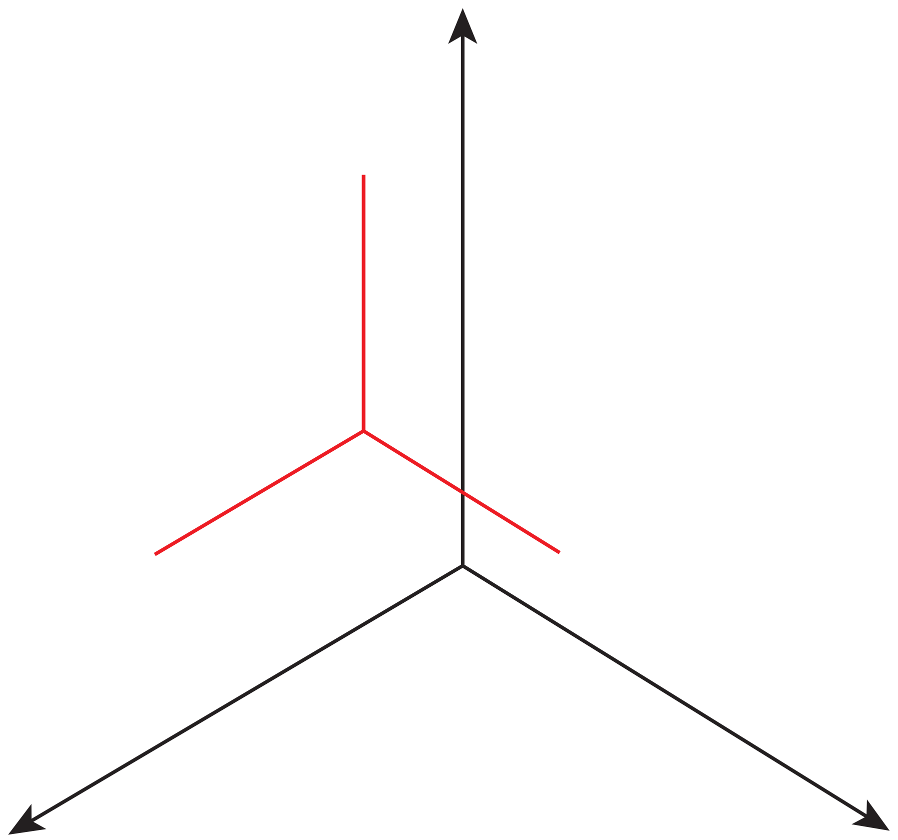

Problem: Entity Extraction

Axes of natural variation for
natural **behavioral** events

e.g.

deforming face moving in complex-
lighted environment

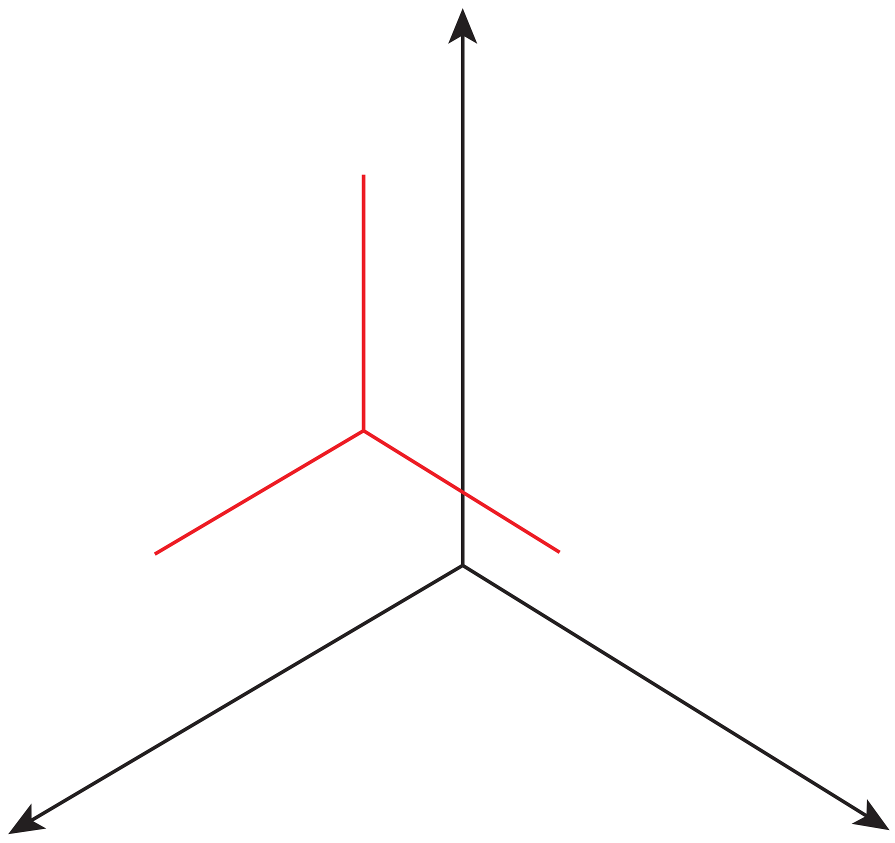

Problem: Entity Extraction

Axes of natural variation for natural **behavioral** events (e.g. deforming face moving in complex-lighted environment)

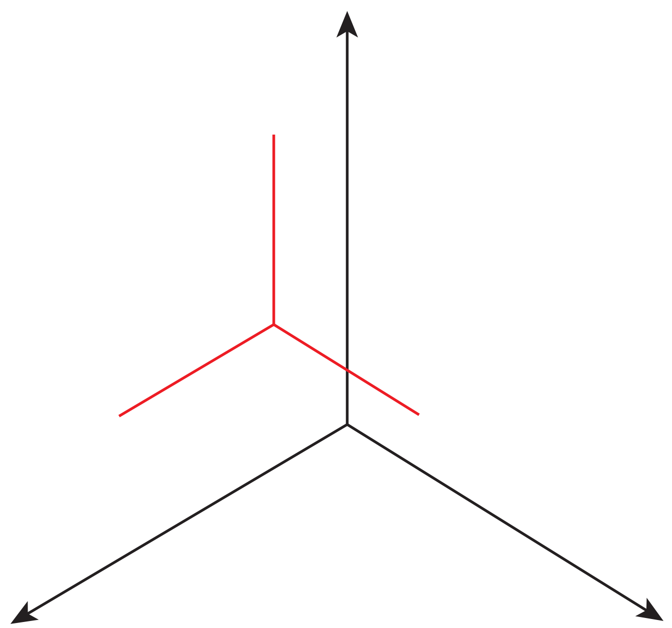

are misaligned with

Axes of natural variation of natural **“physics”** representation of world e.g. retinal photoreceptor voltage

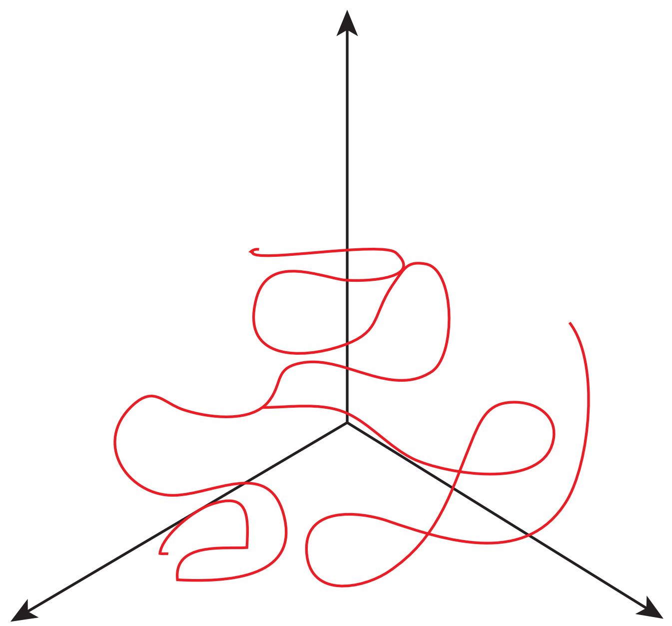

Problem: Entity Extraction

Axes of natural variation for
natural **behavioral** events
(e.g. deforming face moving in
complex-lighted environment)

are misaligned with

Axes of natural variation of natural
“physics” representation of world
e.g. retinal photoreceptor voltage

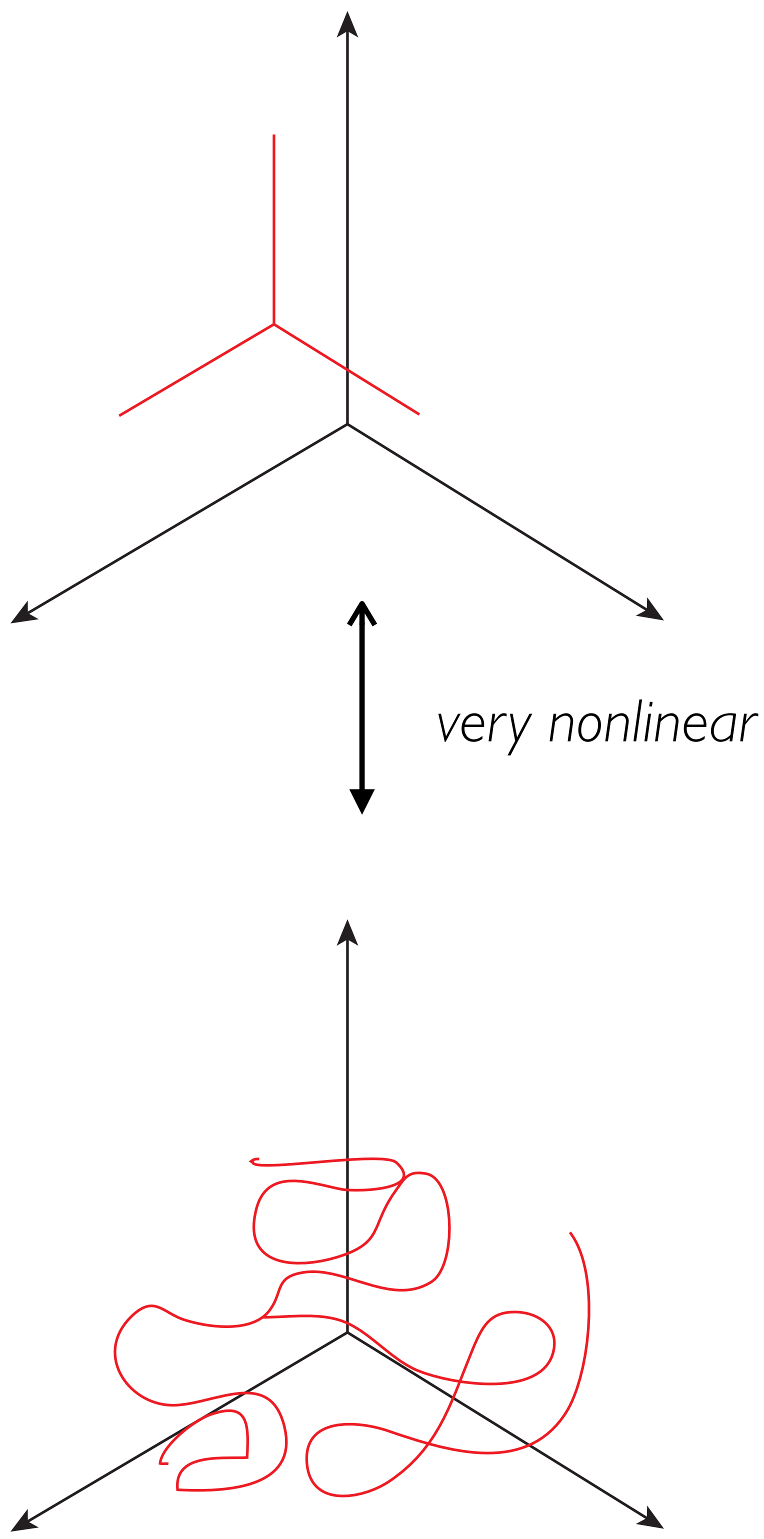

Problem: Entity Extraction

Why is the problem hard computationally?

1. Nonlinear misalignment between physical and behavioral dimensions

Problem: Entity Extraction

Why is the problem hard computationally?

1. Nonlinear misalignment between physical and behavioral dimensions

2. Needs to be done ***fast***, and thus, presumably, massively in parallel

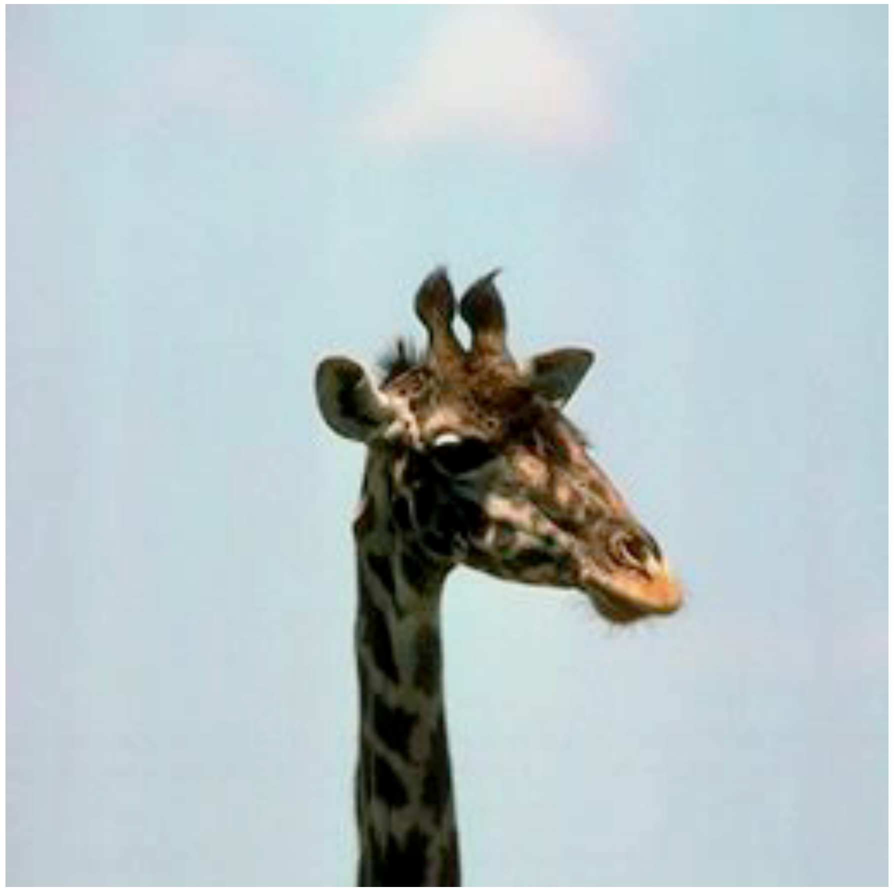

A Modern Approach

NeuroAI Pathways

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

Theo Dobzhansky

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

Theo Dobzhansky

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

Gordon Shepherd

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

Theo Dobzhansky

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

Gordon Shepherd

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

CS 375

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

Theo Dobzhansky

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

Gordon Shepherd

computational

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

CS 375

“Nothing in biology makes sense except in light of evolution”

Dobzhansky

Restated:

Behavior is highly constraining of the brain

“Nothing

Gordon Shepherd

<u>computational</u>

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

CS 375

“Nothing in biology makes sense except in light of evolution”

Dobzhansky

Restated:

Behavior is highly constraining of the brain,

as revealed by computational models

“Nothing

Gordon Shepherd

computational

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

CS 375

Heuristic of "Goal-Driven Modeling"

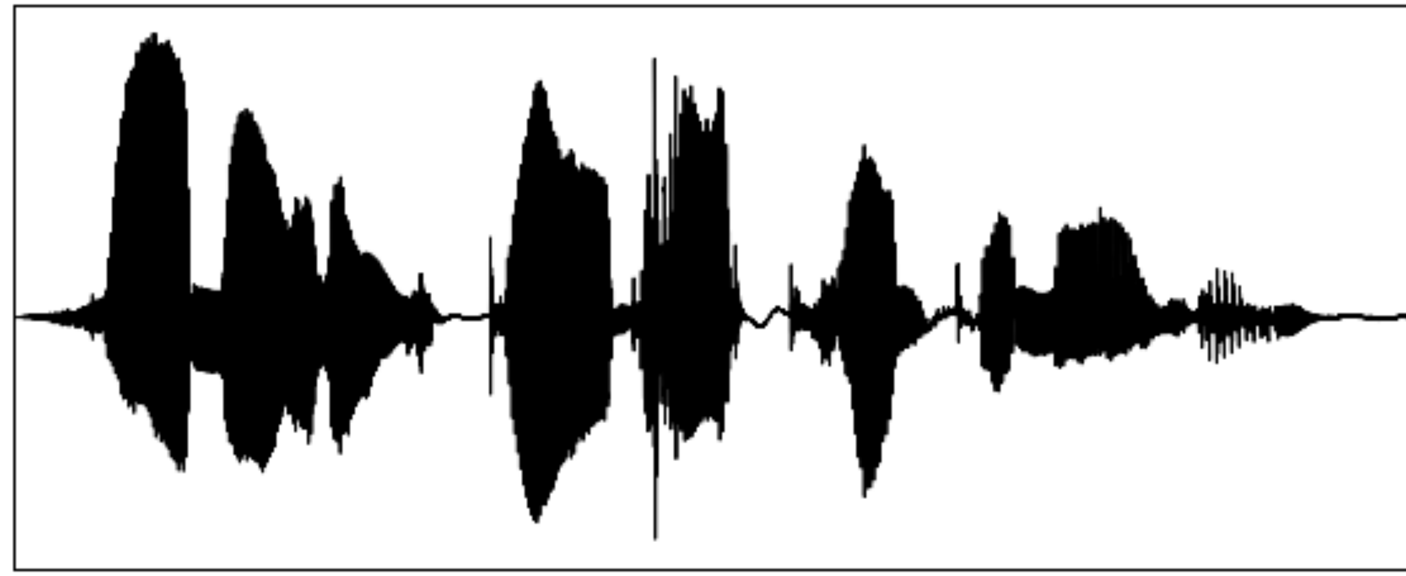

visual
cortex

auditory
cortex

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

"Hannah is good at compromising"

Heuristic of "Goal-Driven Modeling"

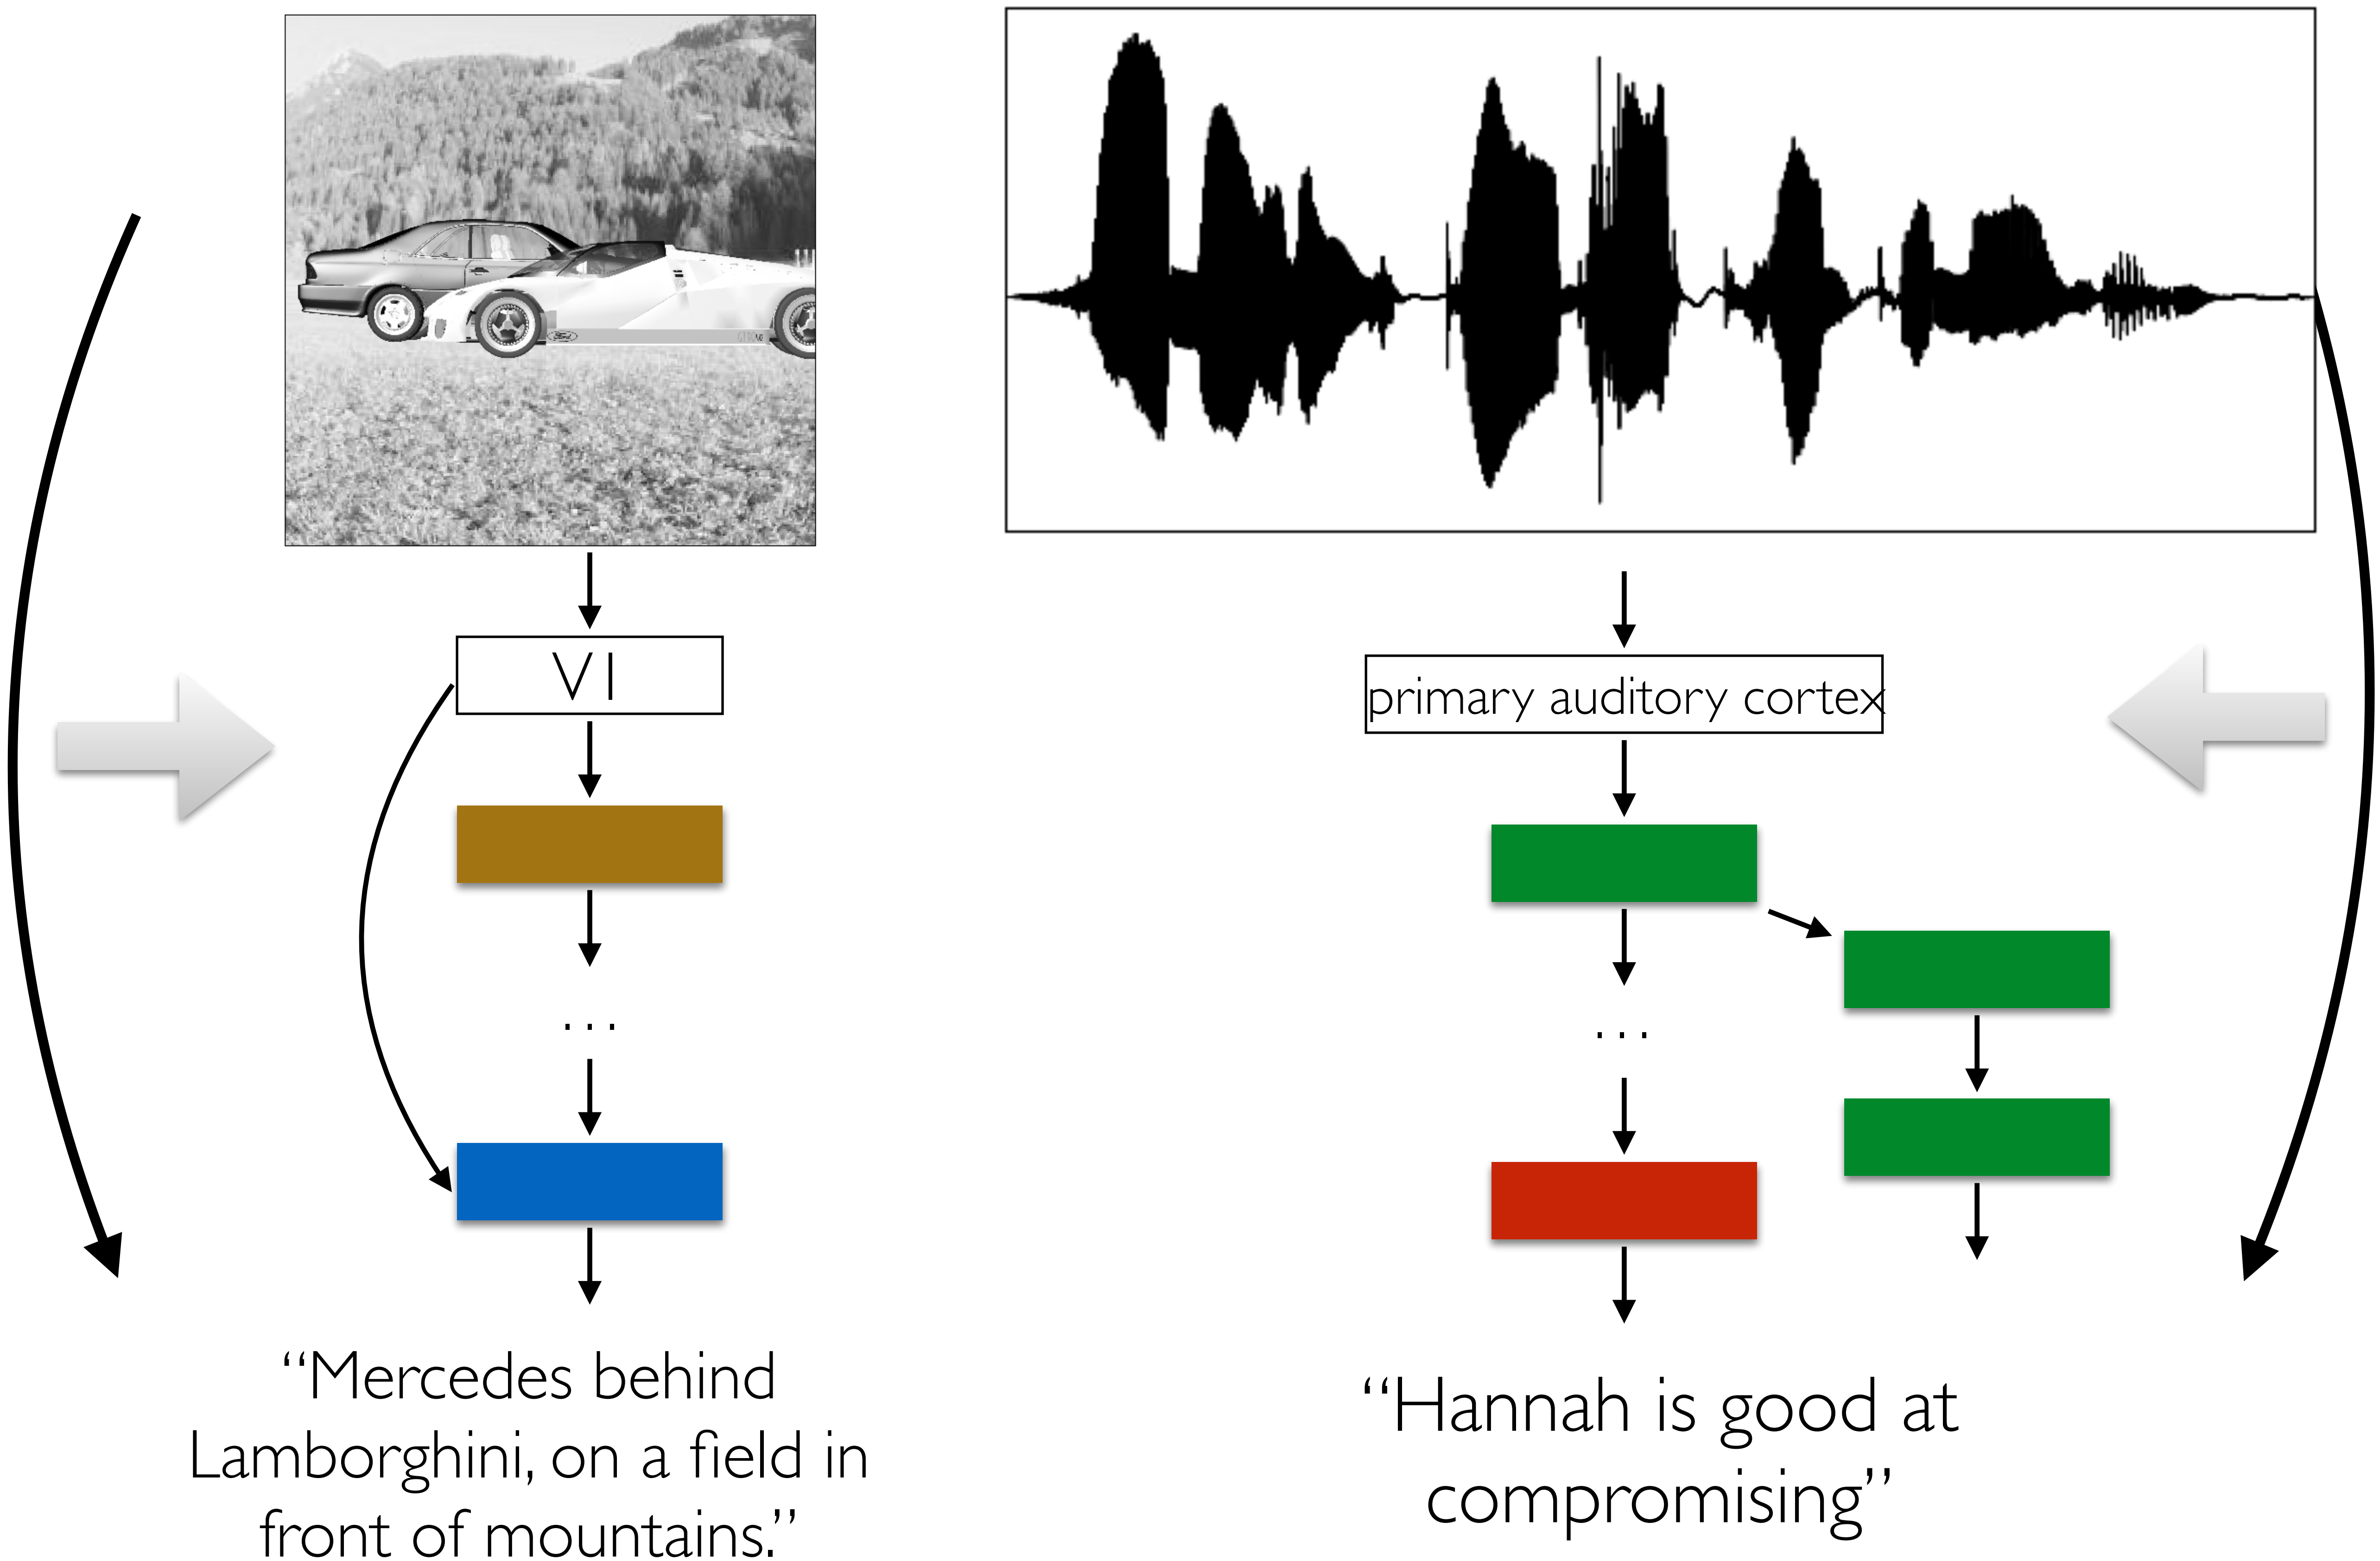

1. Formulate comprehensive model class (**CNNs**)

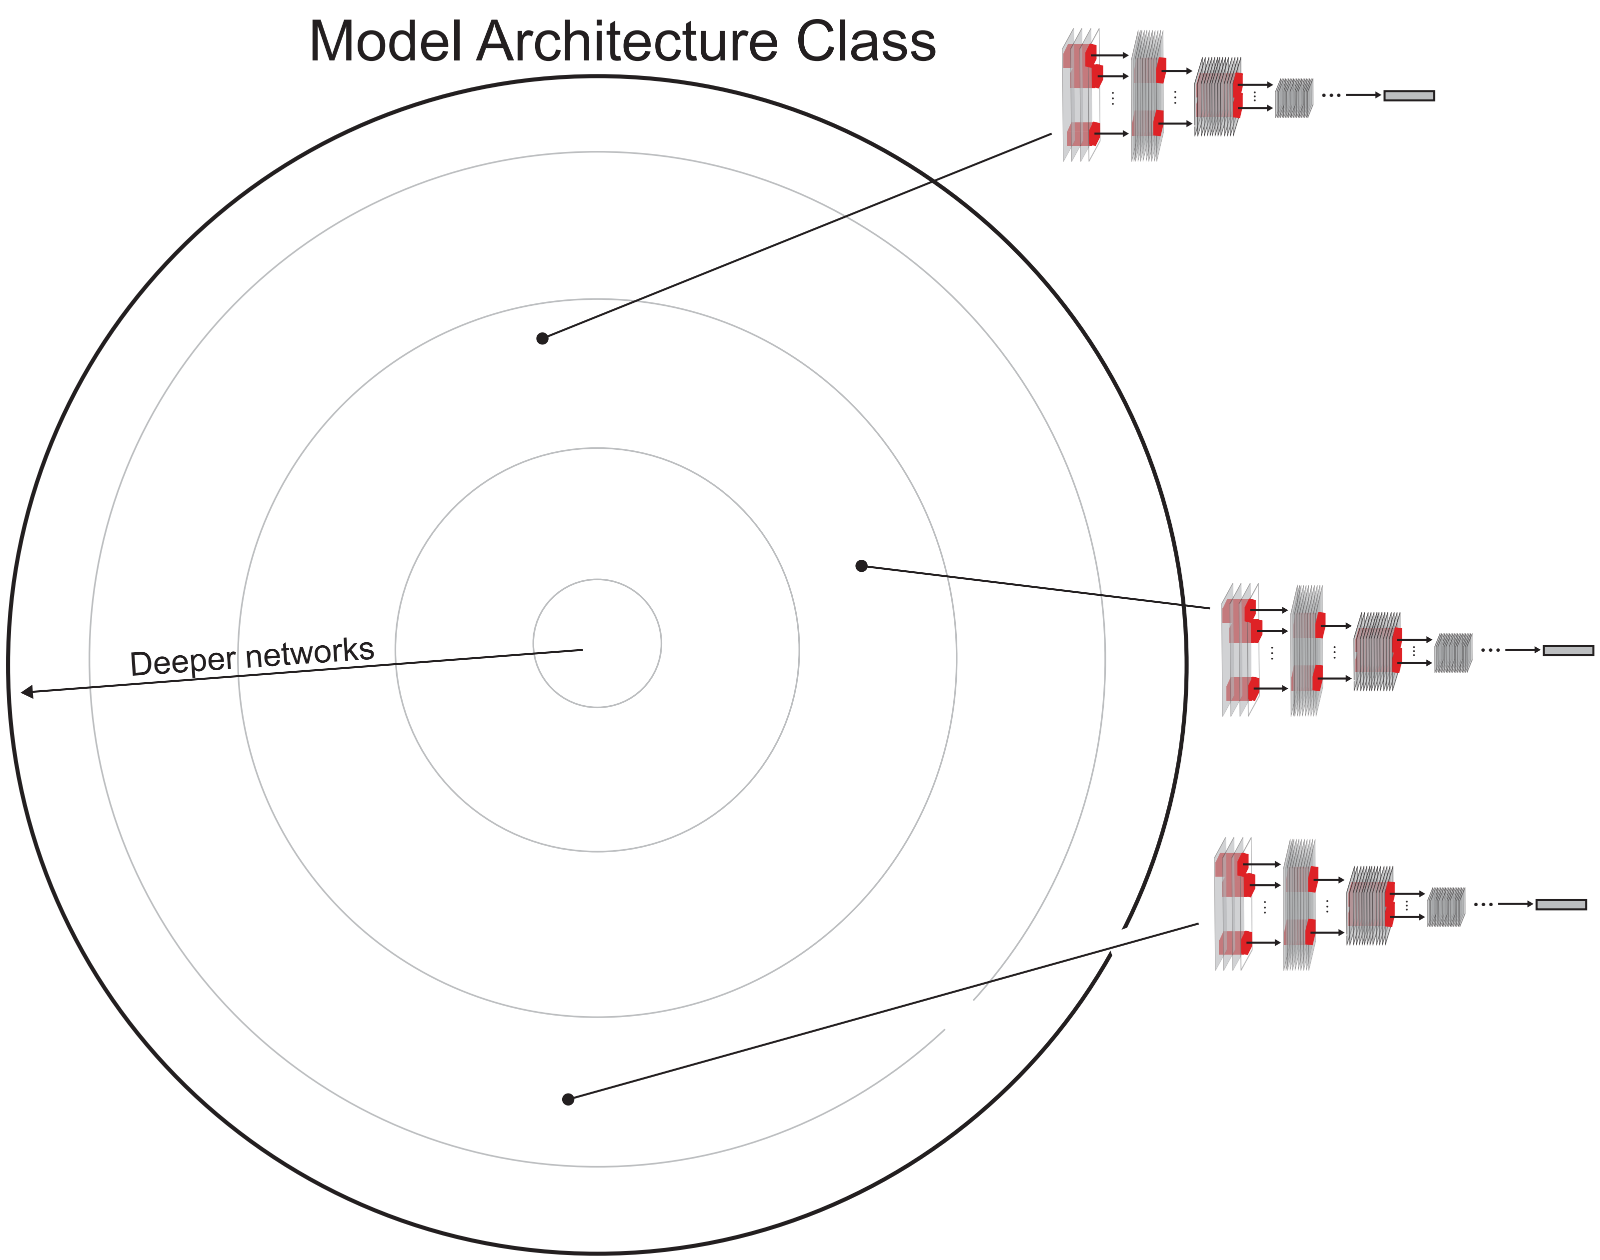

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

1. Formulate comprehensive model class (**CNNs**)

2. Choose challenging, ethologically-valid tasks (**categorization**)

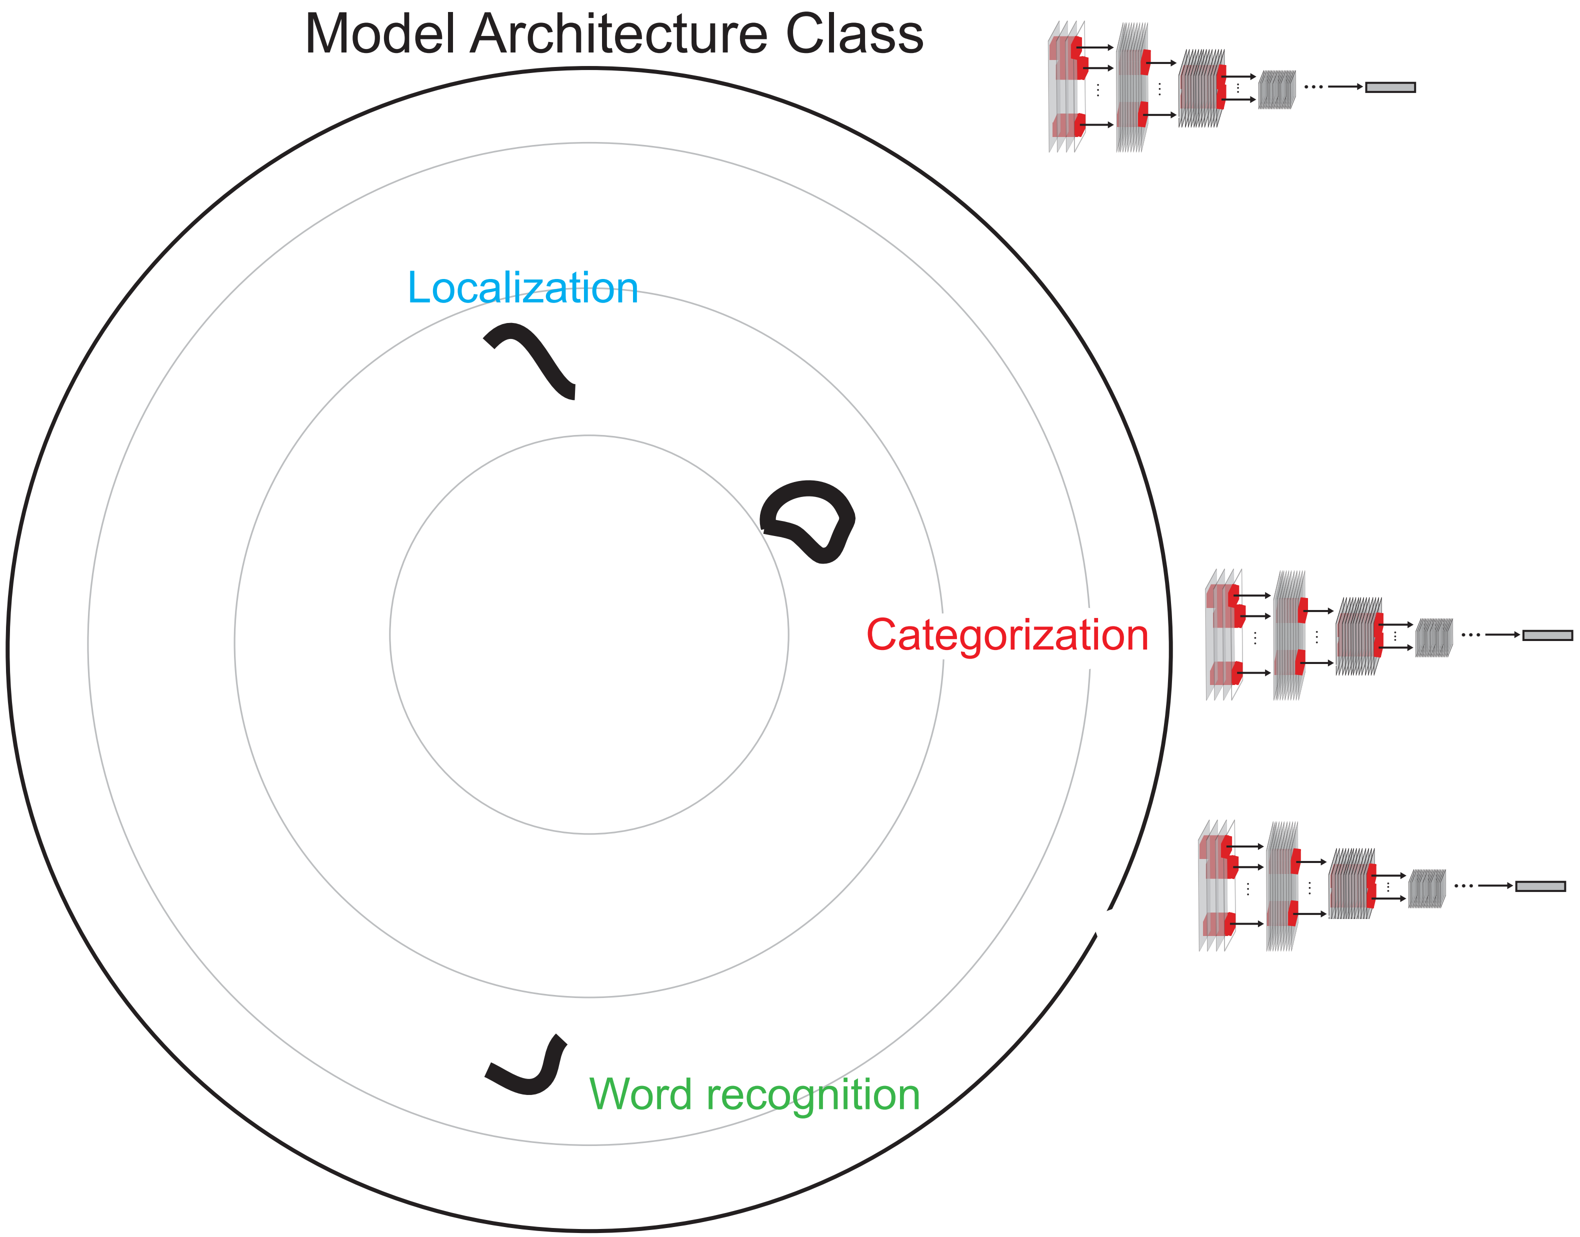

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

1. Formulate comprehensive model class (**CNNs**)

2. Choose challenging, ethologically-valid tasks (**categorization**)

3. Implement generic learning rules (**gradient descent**)

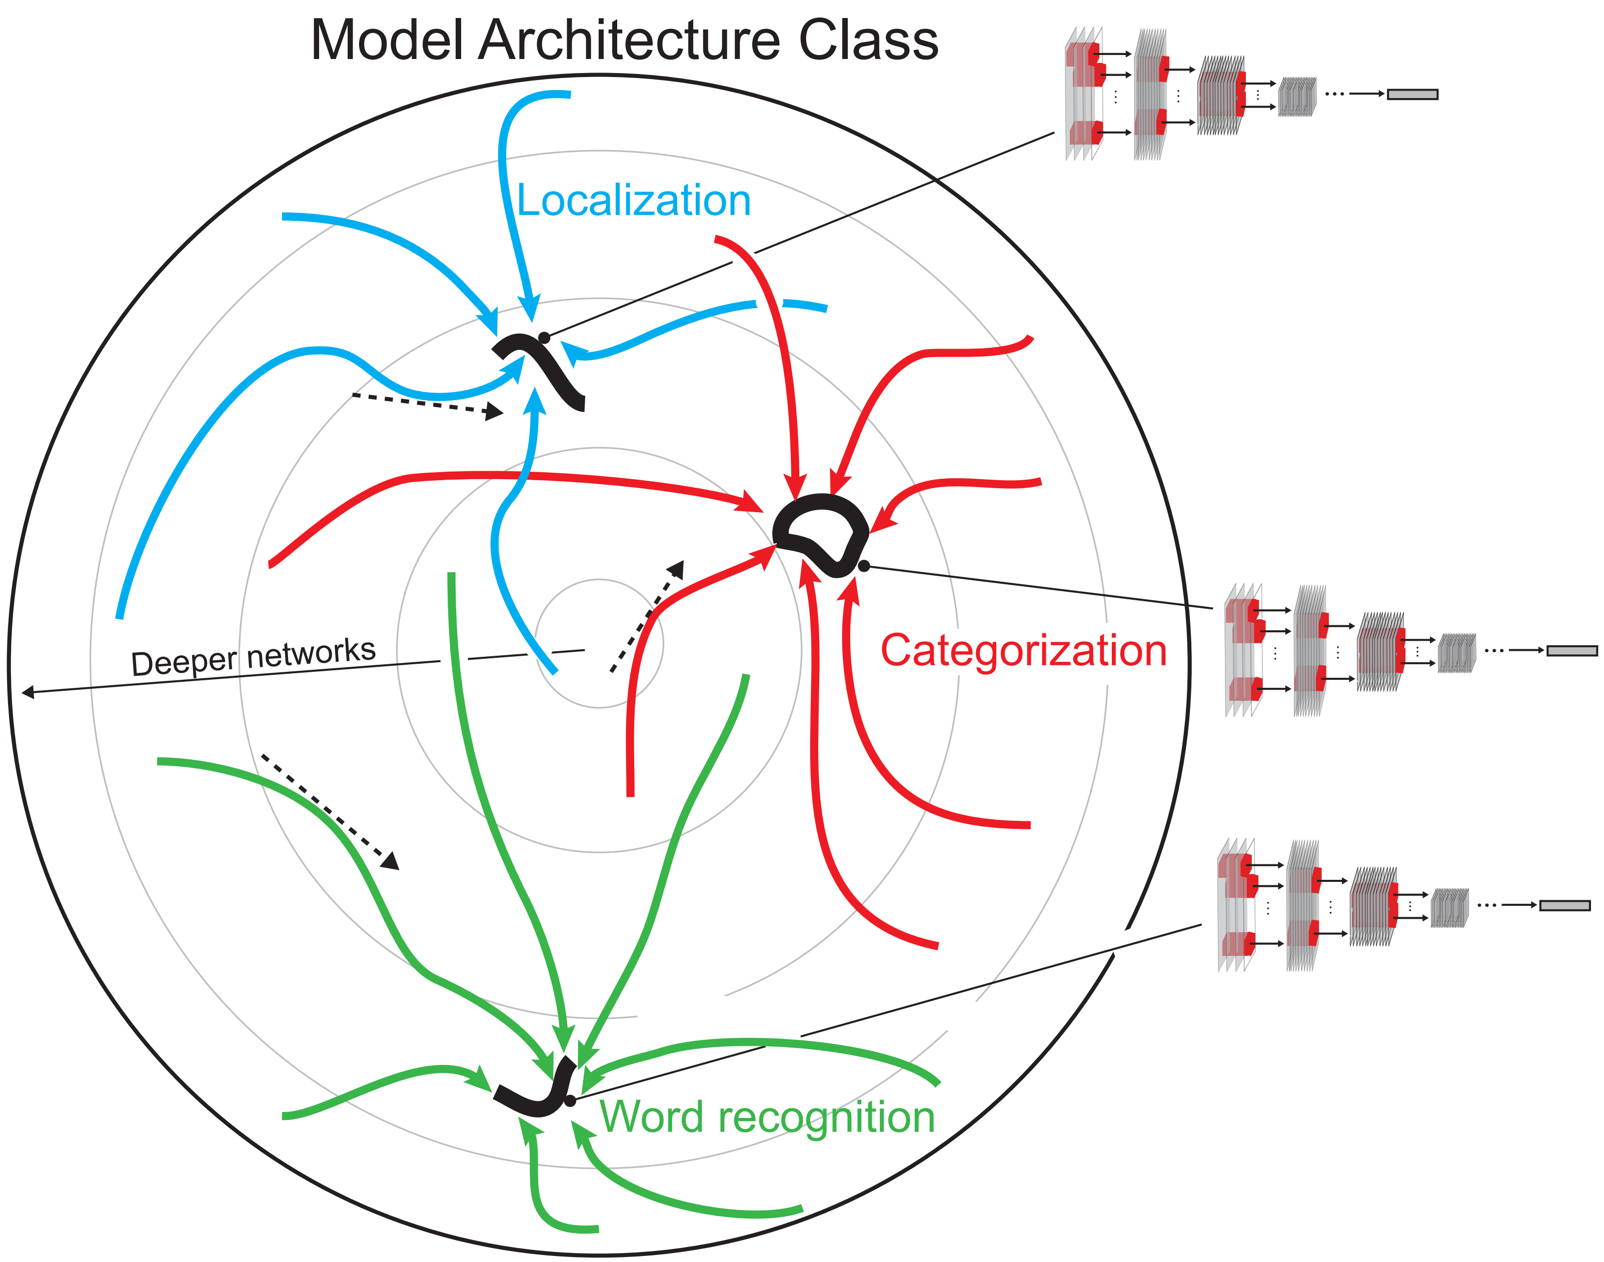

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

1. Formulate comprehensive model class (**CNNs**)

2. Choose challenging, ethologically-valid tasks (**categorization**)

3. Implement generic learning rules (**gradient descent**)

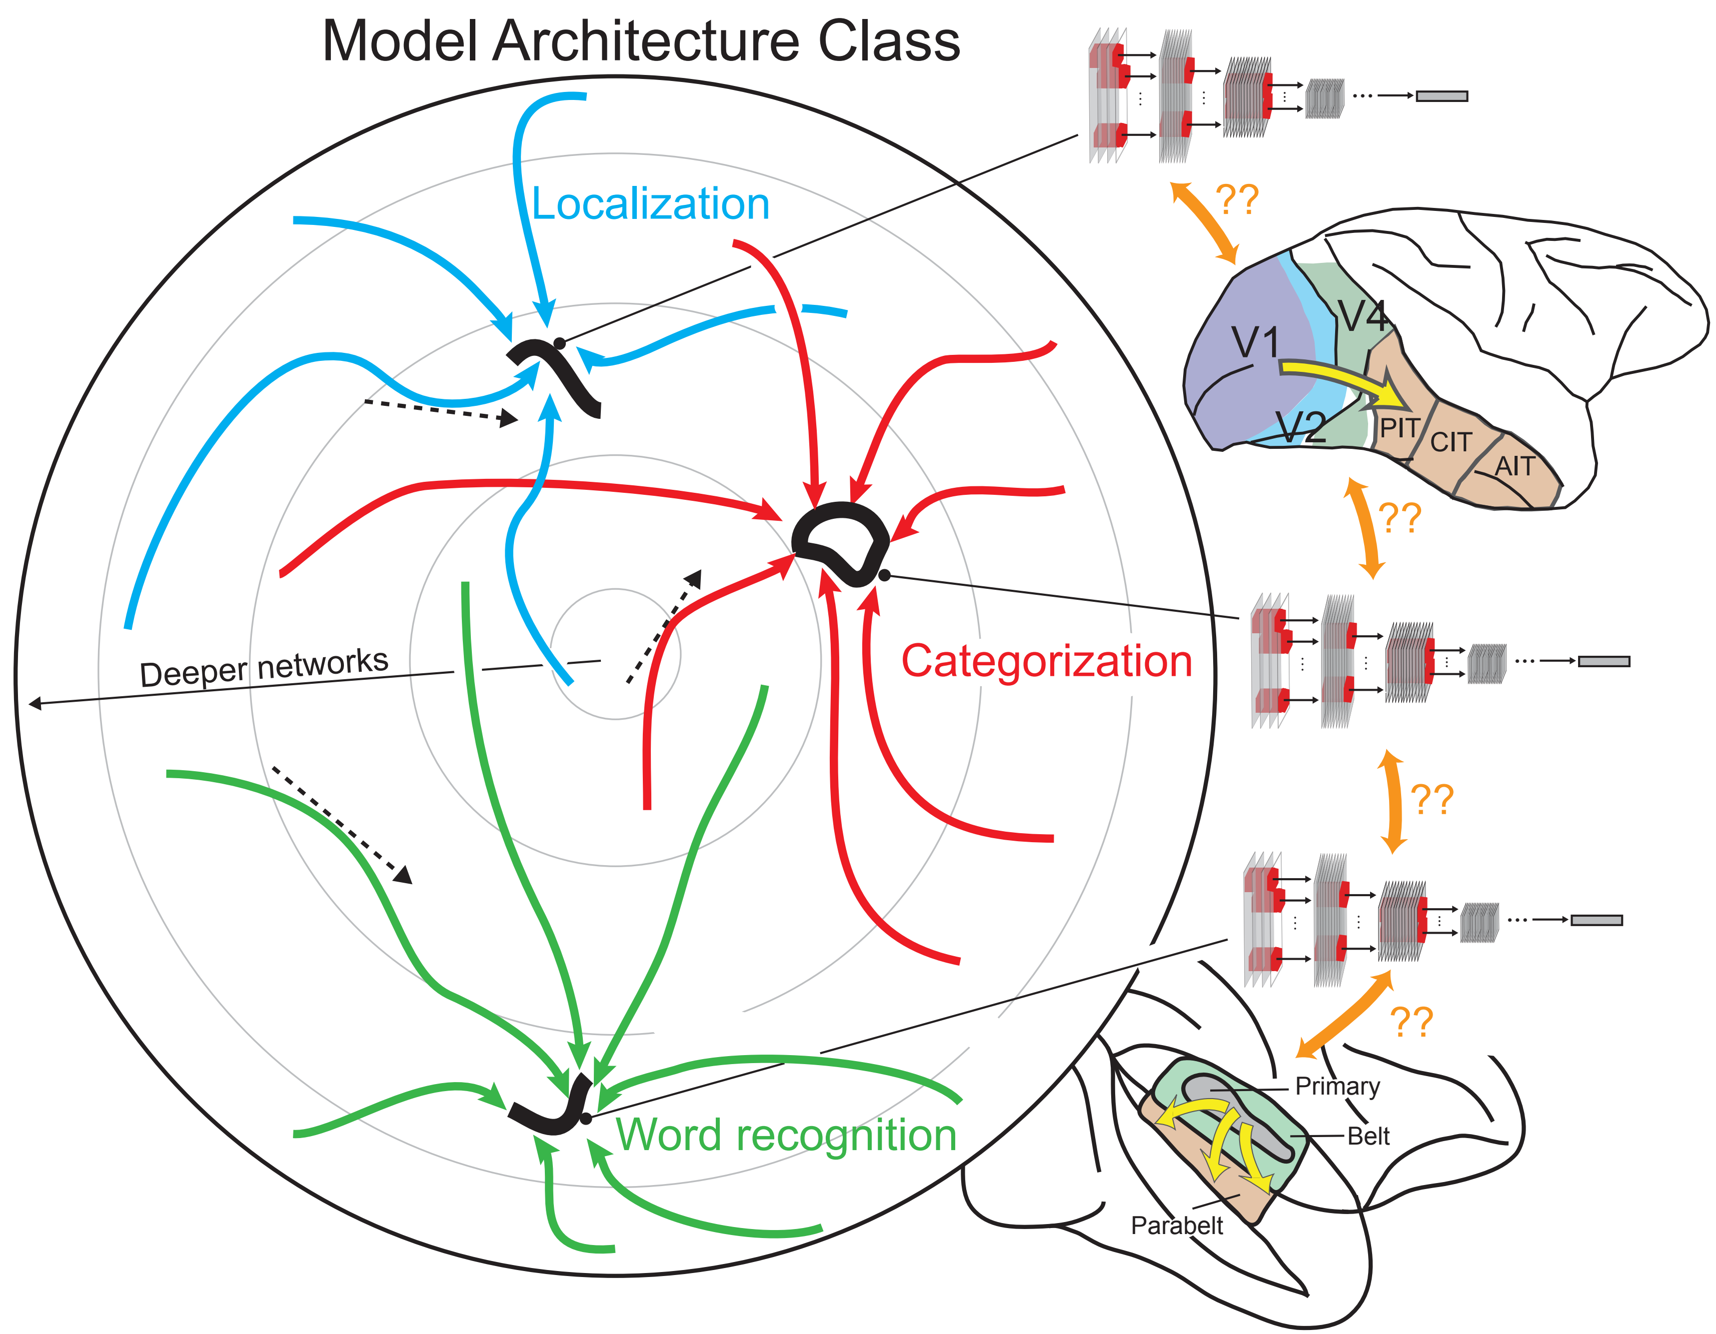

Yamins & DiCarlo.
Nat. Neuro. *(2016)*

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

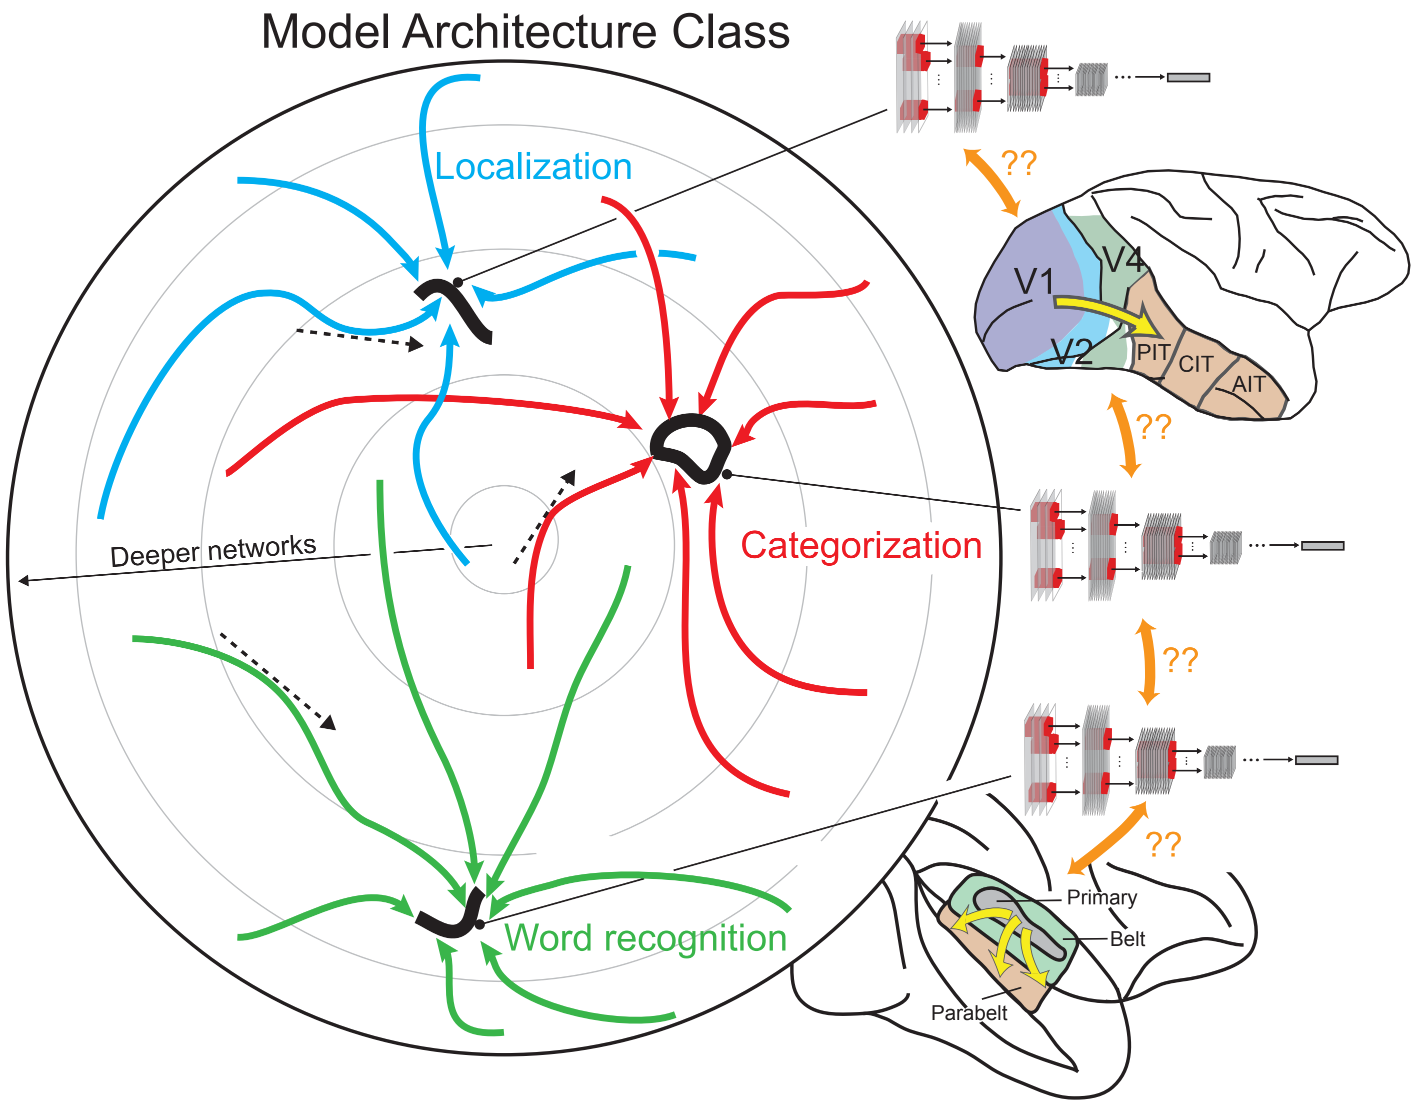

$$\underset{a \in \mathcal{A}}{\operatorname{argmin}}[L(p_a^*)]$$

where p^* is result of

$$\frac{dp_a}{dt} = -\lambda(t) \cdot \langle \nabla_{p_a} L(x) \rangle_{x \in \mathcal{D}}$$

A = *architecture class* ***L*** = *loss function* ***D*** = *dataset*

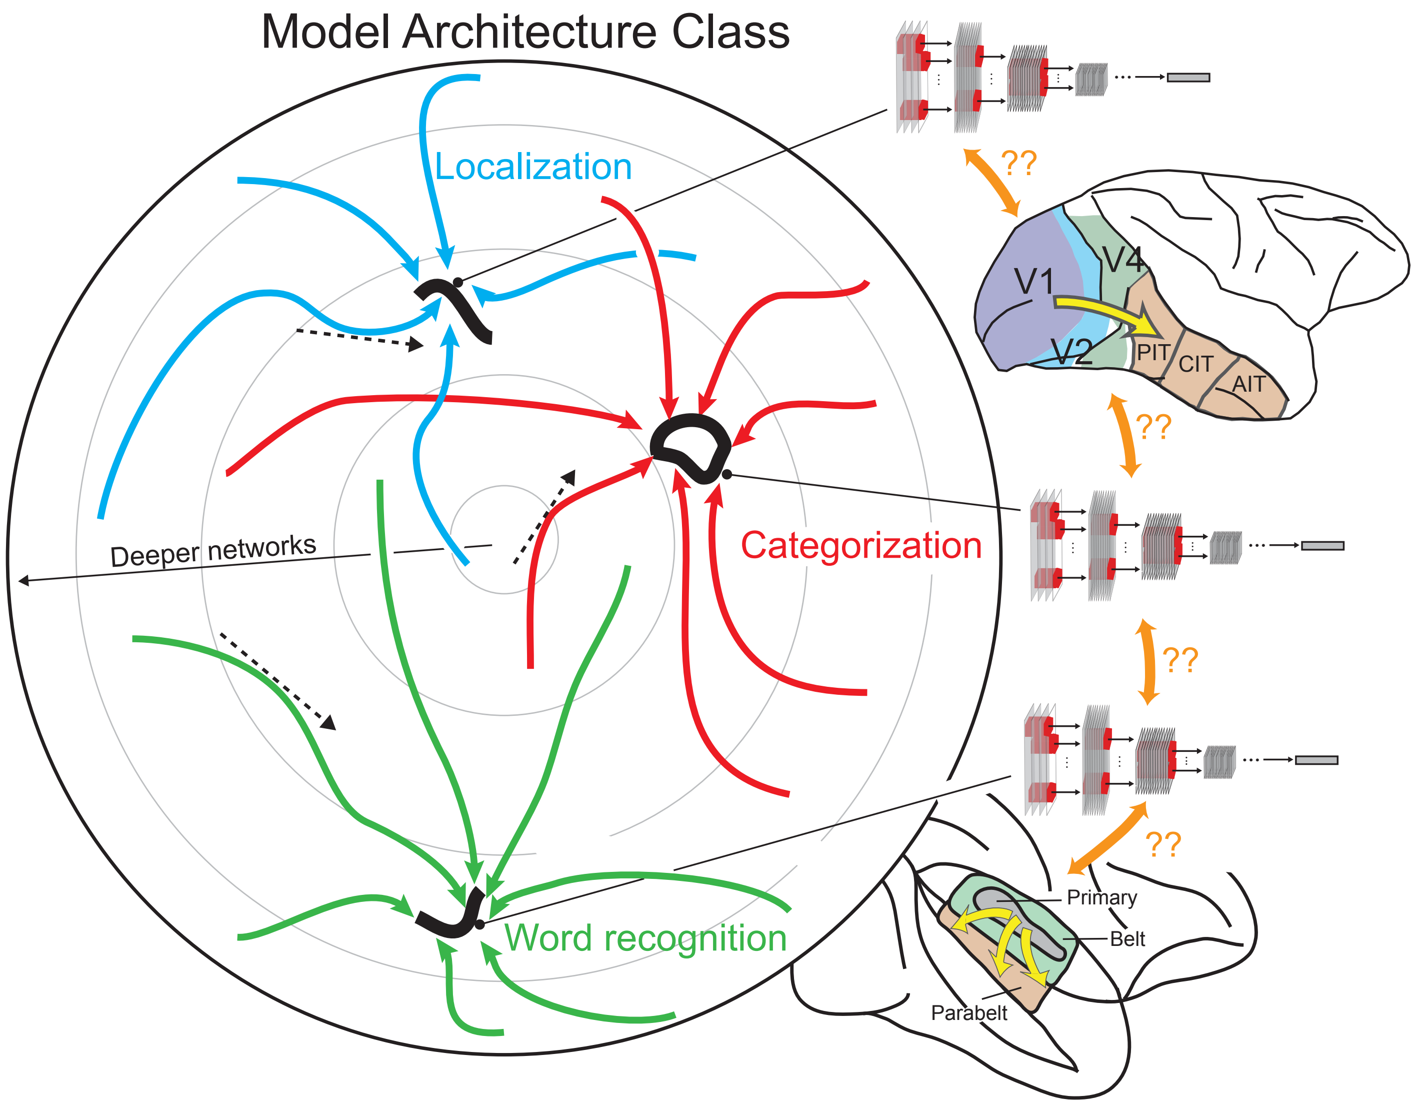

3.

$$\underset{a \in \mathcal{A}}{\operatorname{argmin}}[L(p_a^*)]$$

where p^* is result of

$$\frac{dp_a}{dt} = -\lambda(t) \cdot \langle \nabla_{p_a} L(x) \rangle_{x \in \mathcal{D}}$$

"learning rule"

1.

A = *architecture class*

2.

L = *loss function* ***D*** = *dataset*

"task"

Four Principles of Optimization-Based Modeling

1.

A = *architecture class*

2.

T = *task/objective*

3.

D = *dataset*

4.

L = *learning rule*

Four Principles of Optimization-Based Modeling

1.

A = *architecture class* ***= circuit neuroanatomy***

2.

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

3.

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

4.

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

Four Principles of Optimization-Based Modeling

1.

A = *architecture class* ***= circuit neuroanatomy***

solving

2.

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

situated in

3.

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

updating according to

4.

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

A = *architecture class*

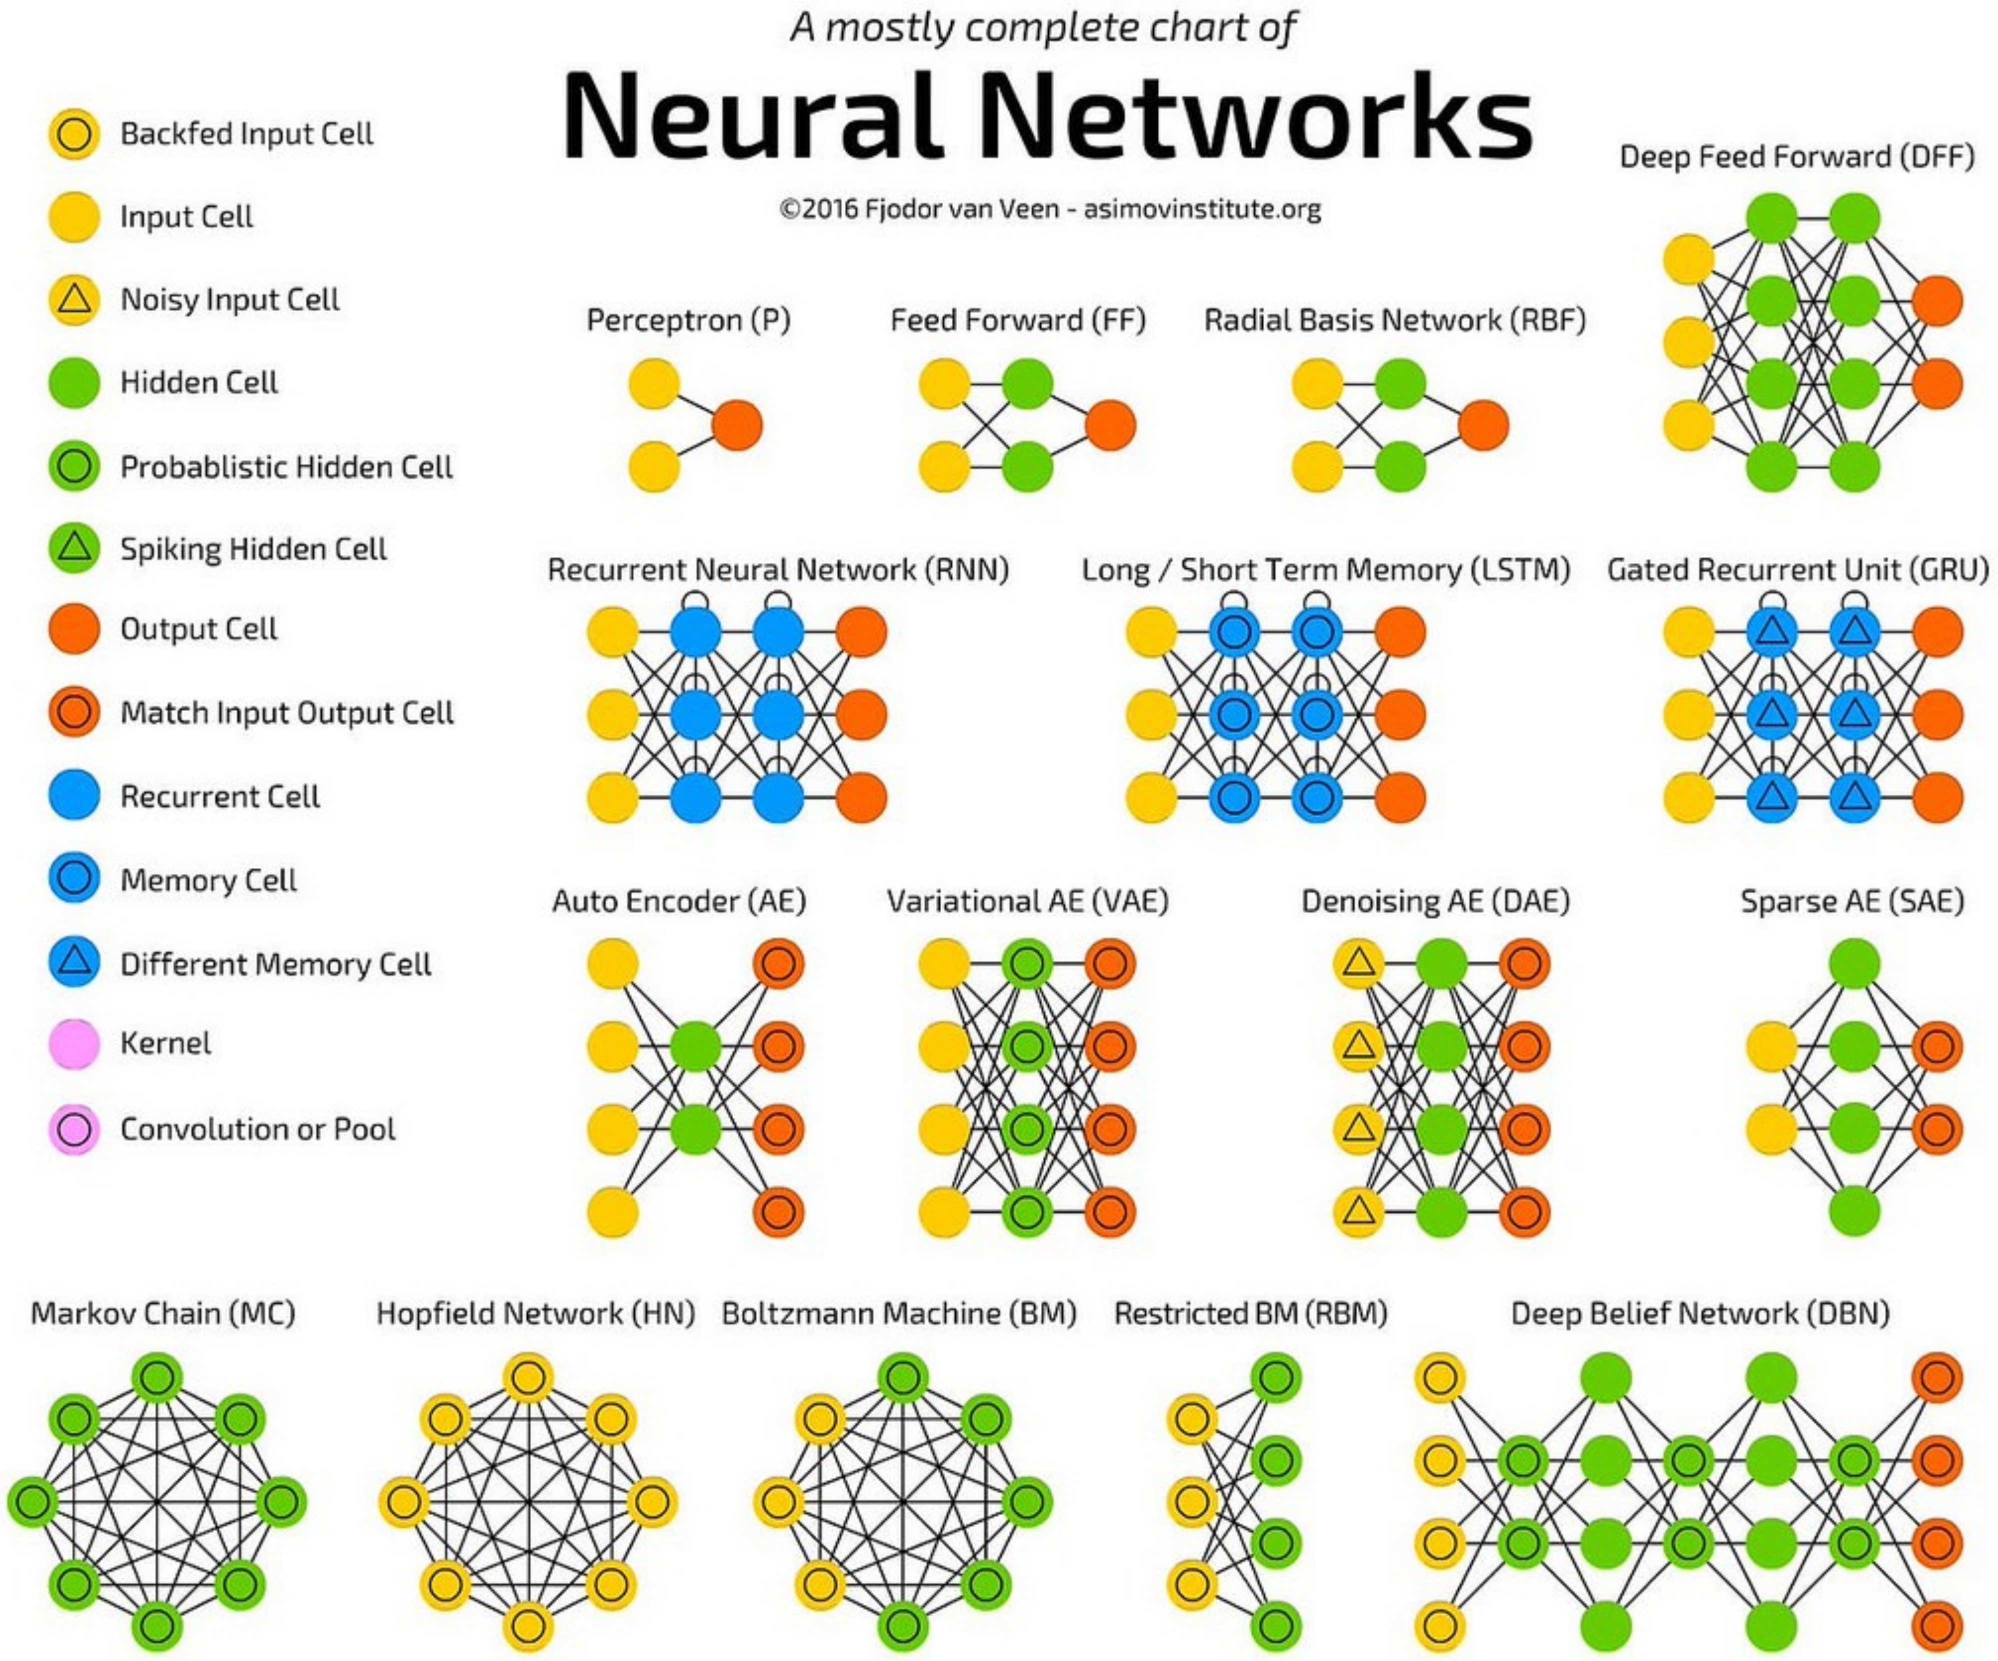

A *= architecture class*

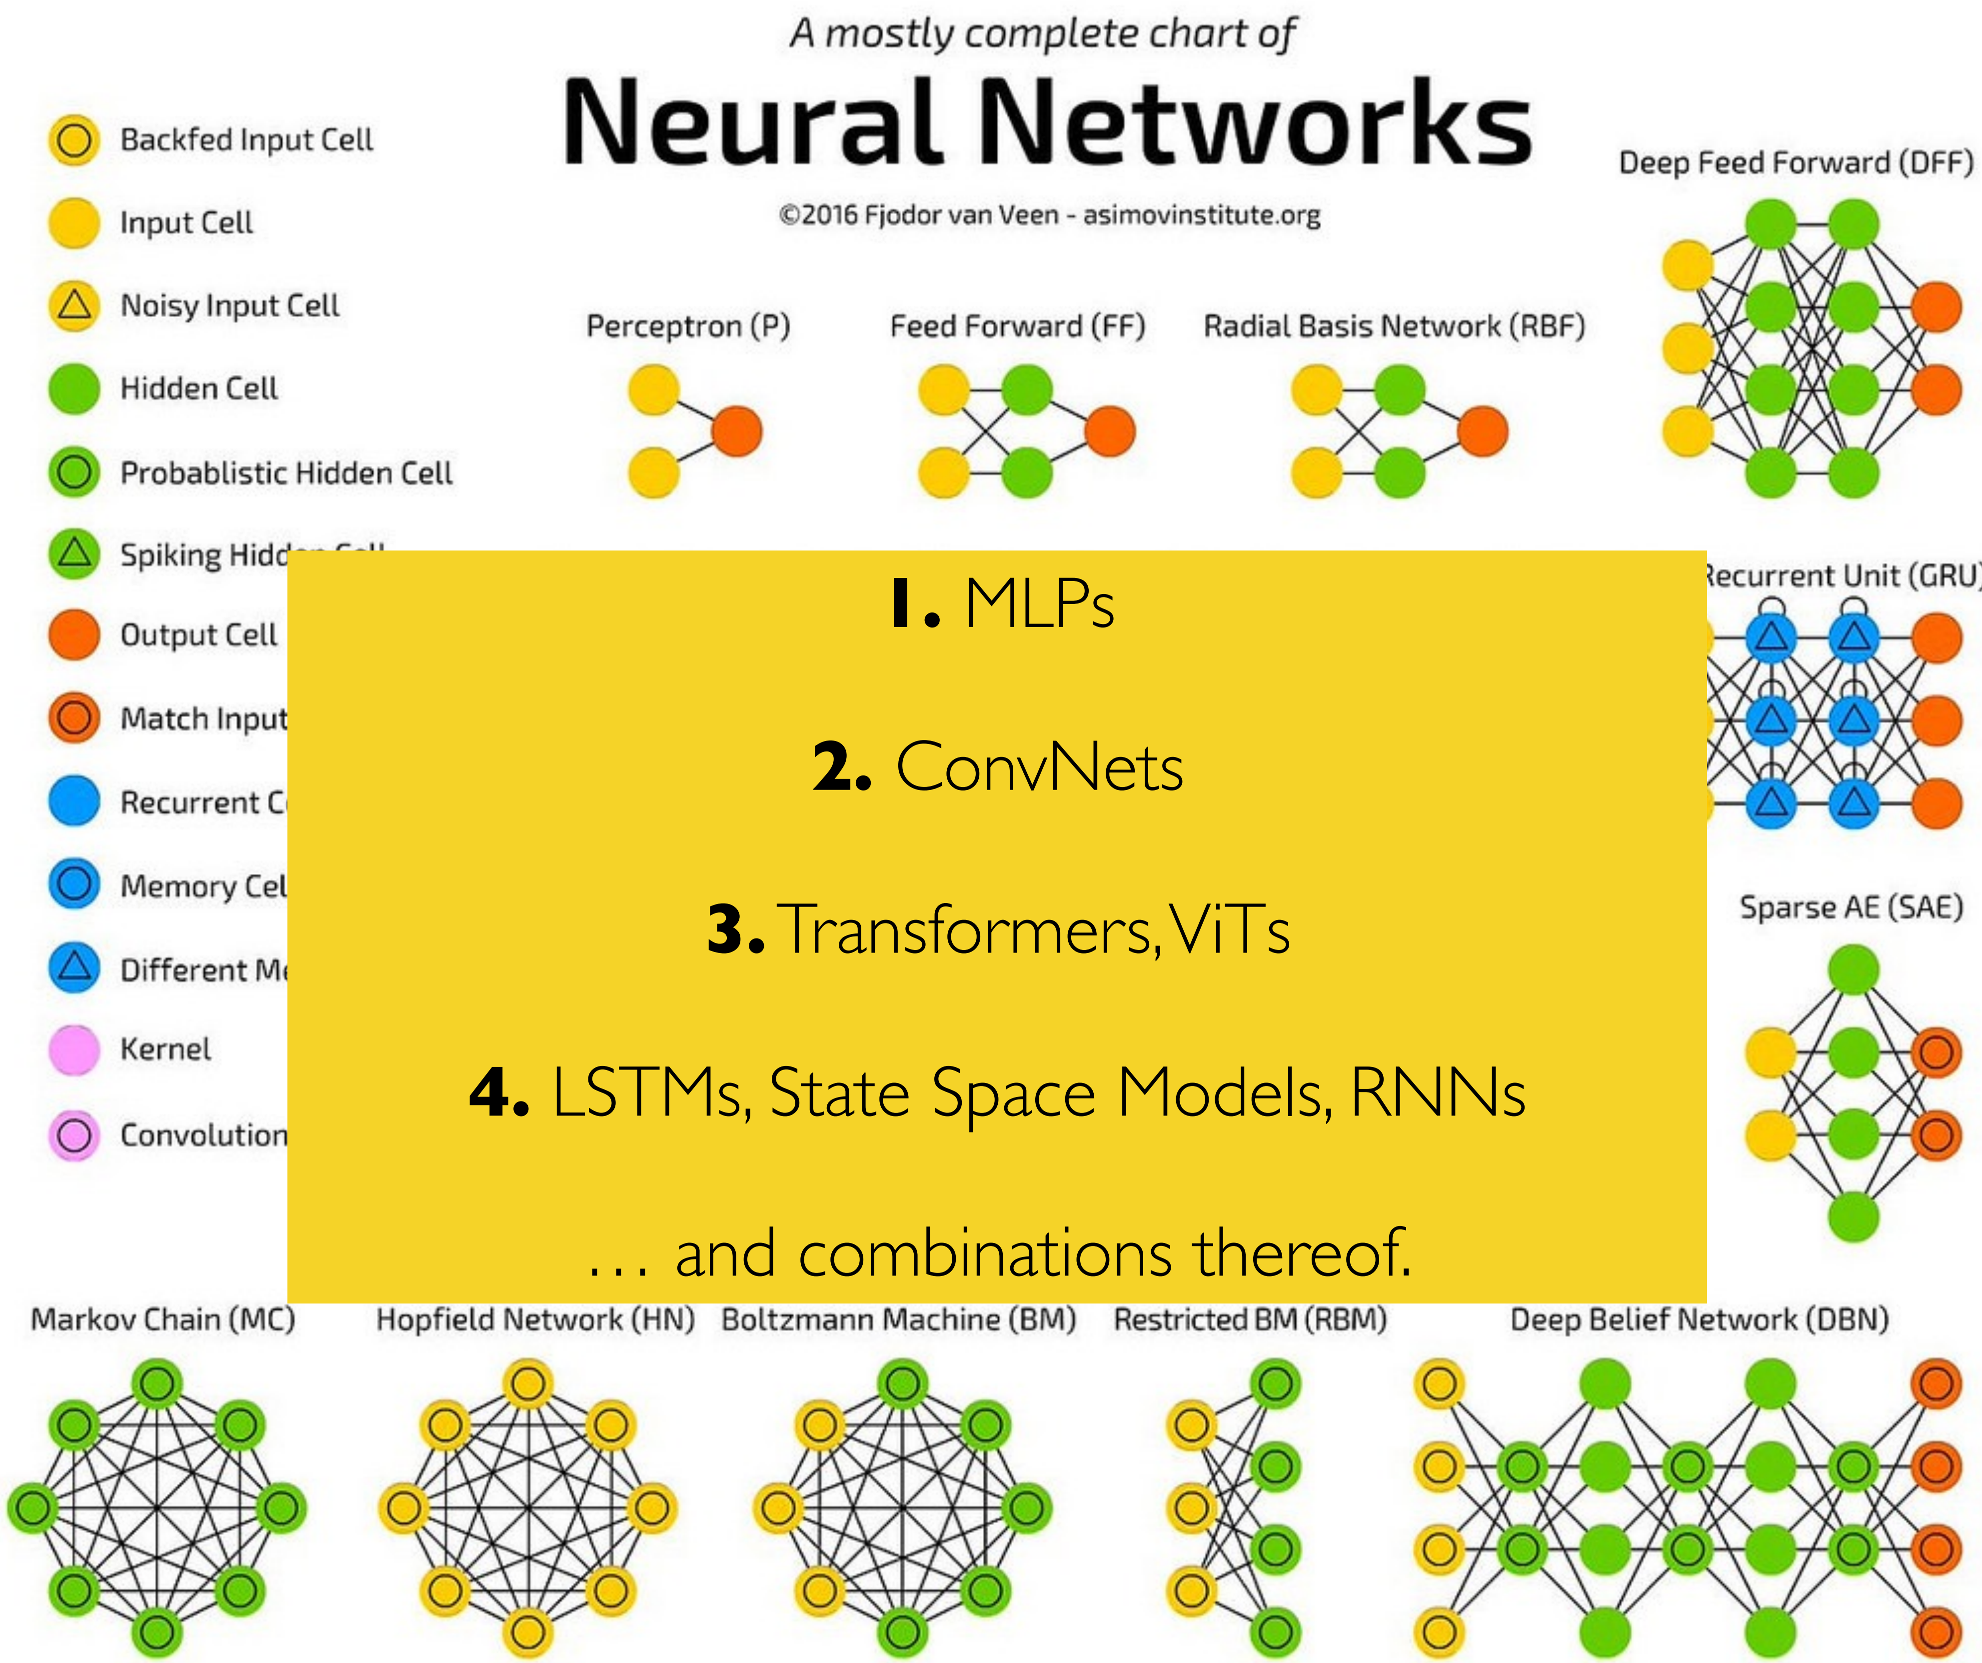

1. MLPs

2. ConvNets

3. Transformers, ViTs

4. LSTMs, State Space Models, RNNs

… and combinations thereof.

Task: L = *loss function* **D** = *dataset*

Surface Normals

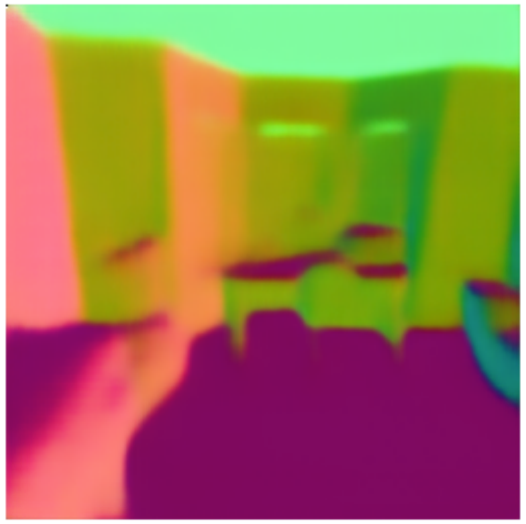

Image Reshading

2D Texture Edges

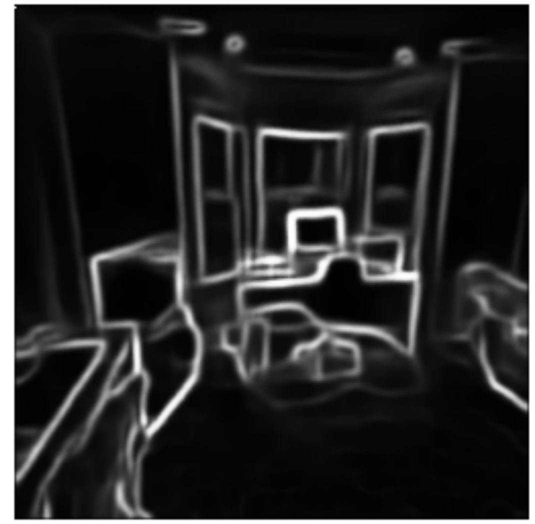

Vanishing Points

Unsupervised 2.5D Segm.

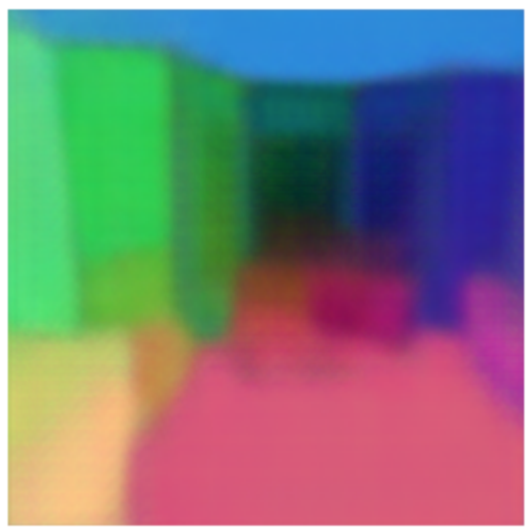

Room Layout

. . . and beyond

Scene Classification

Top 5 prediction:
home_office
office
television_room
computer_room
office_cubicles

3D Keypoints

3D Occlusion Edges

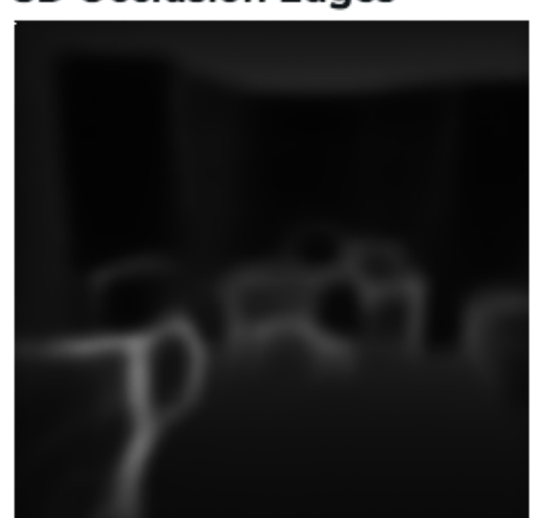

Neuroscience Methods

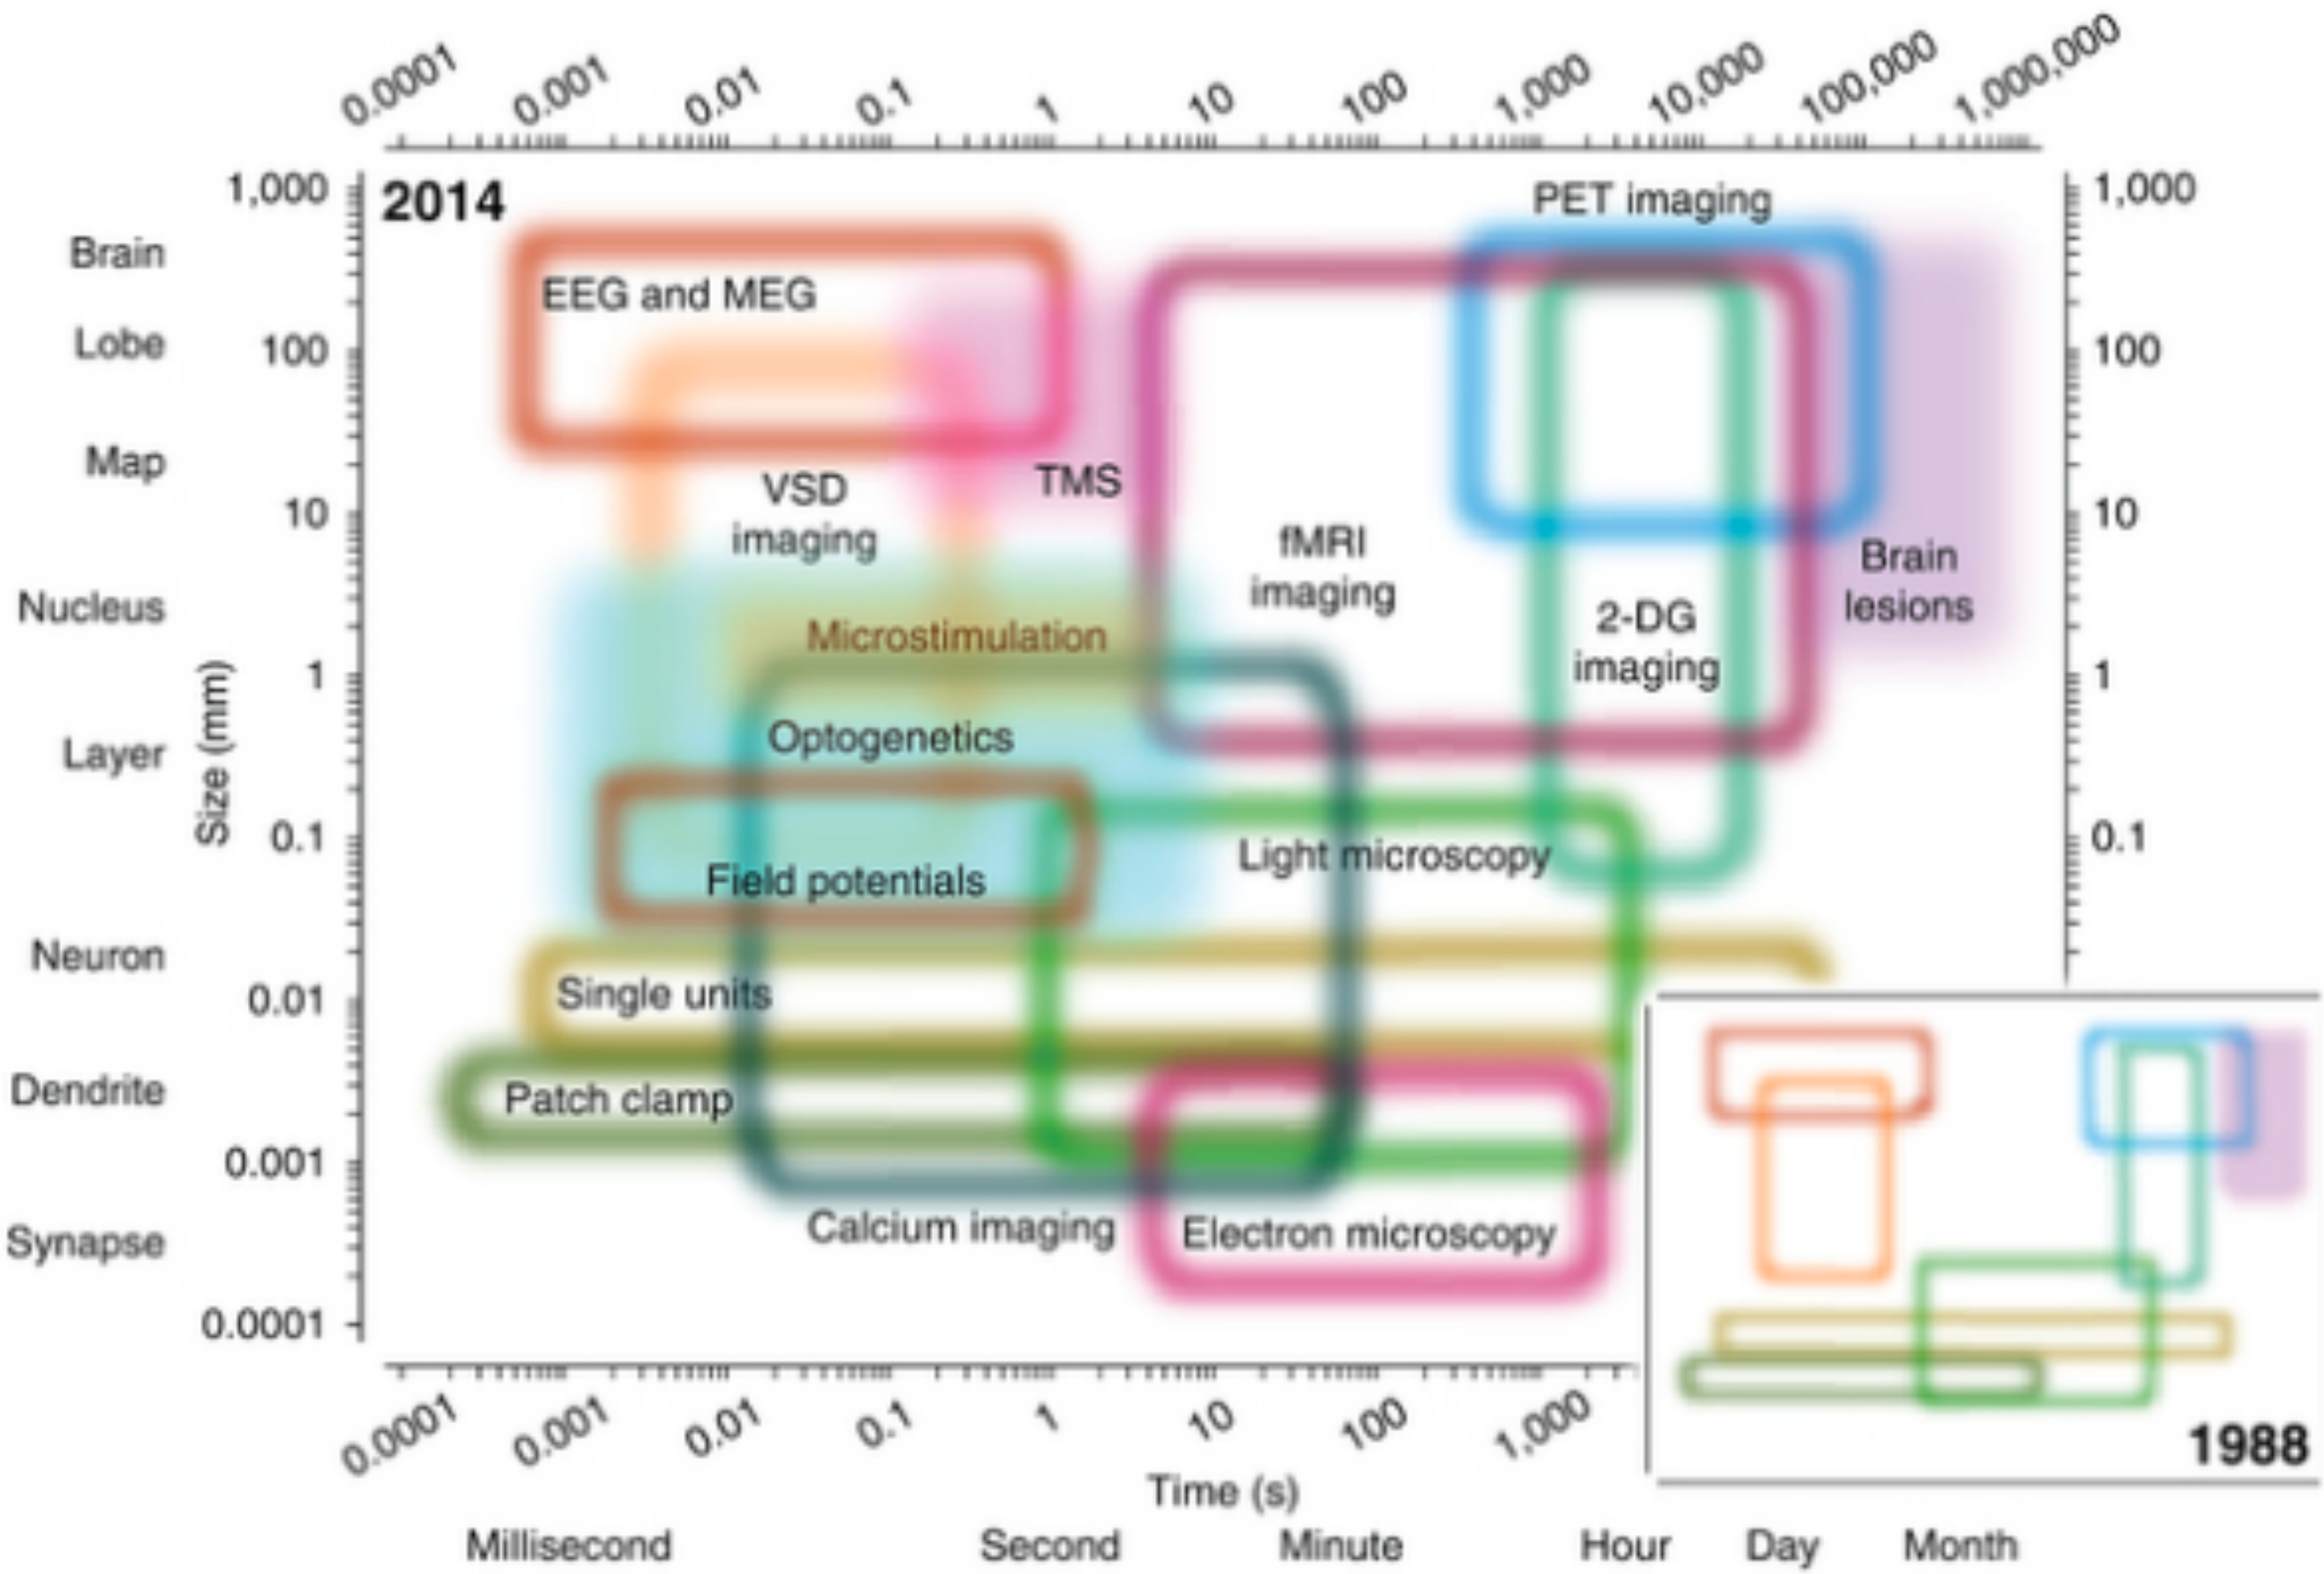

Neural Response Predictivity

Cognitive Task Performance

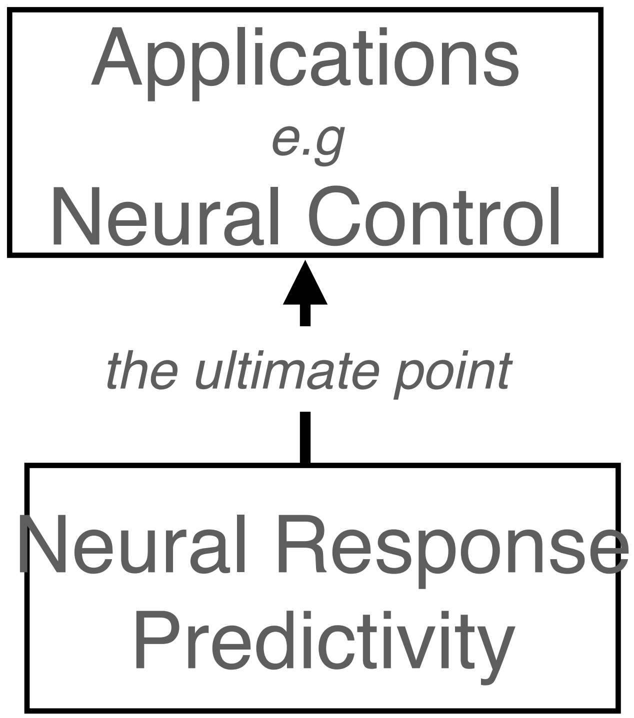
Applications
e.g
Neural Control
the ultimate point
Neural Response
Predictivity

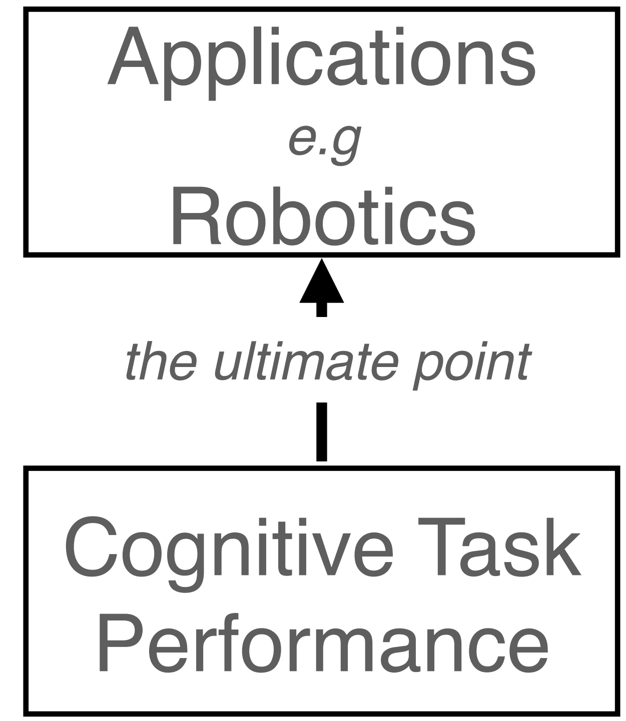
Applications
e.g
Robotics
the ultimate point
Cognitive Task
Performance

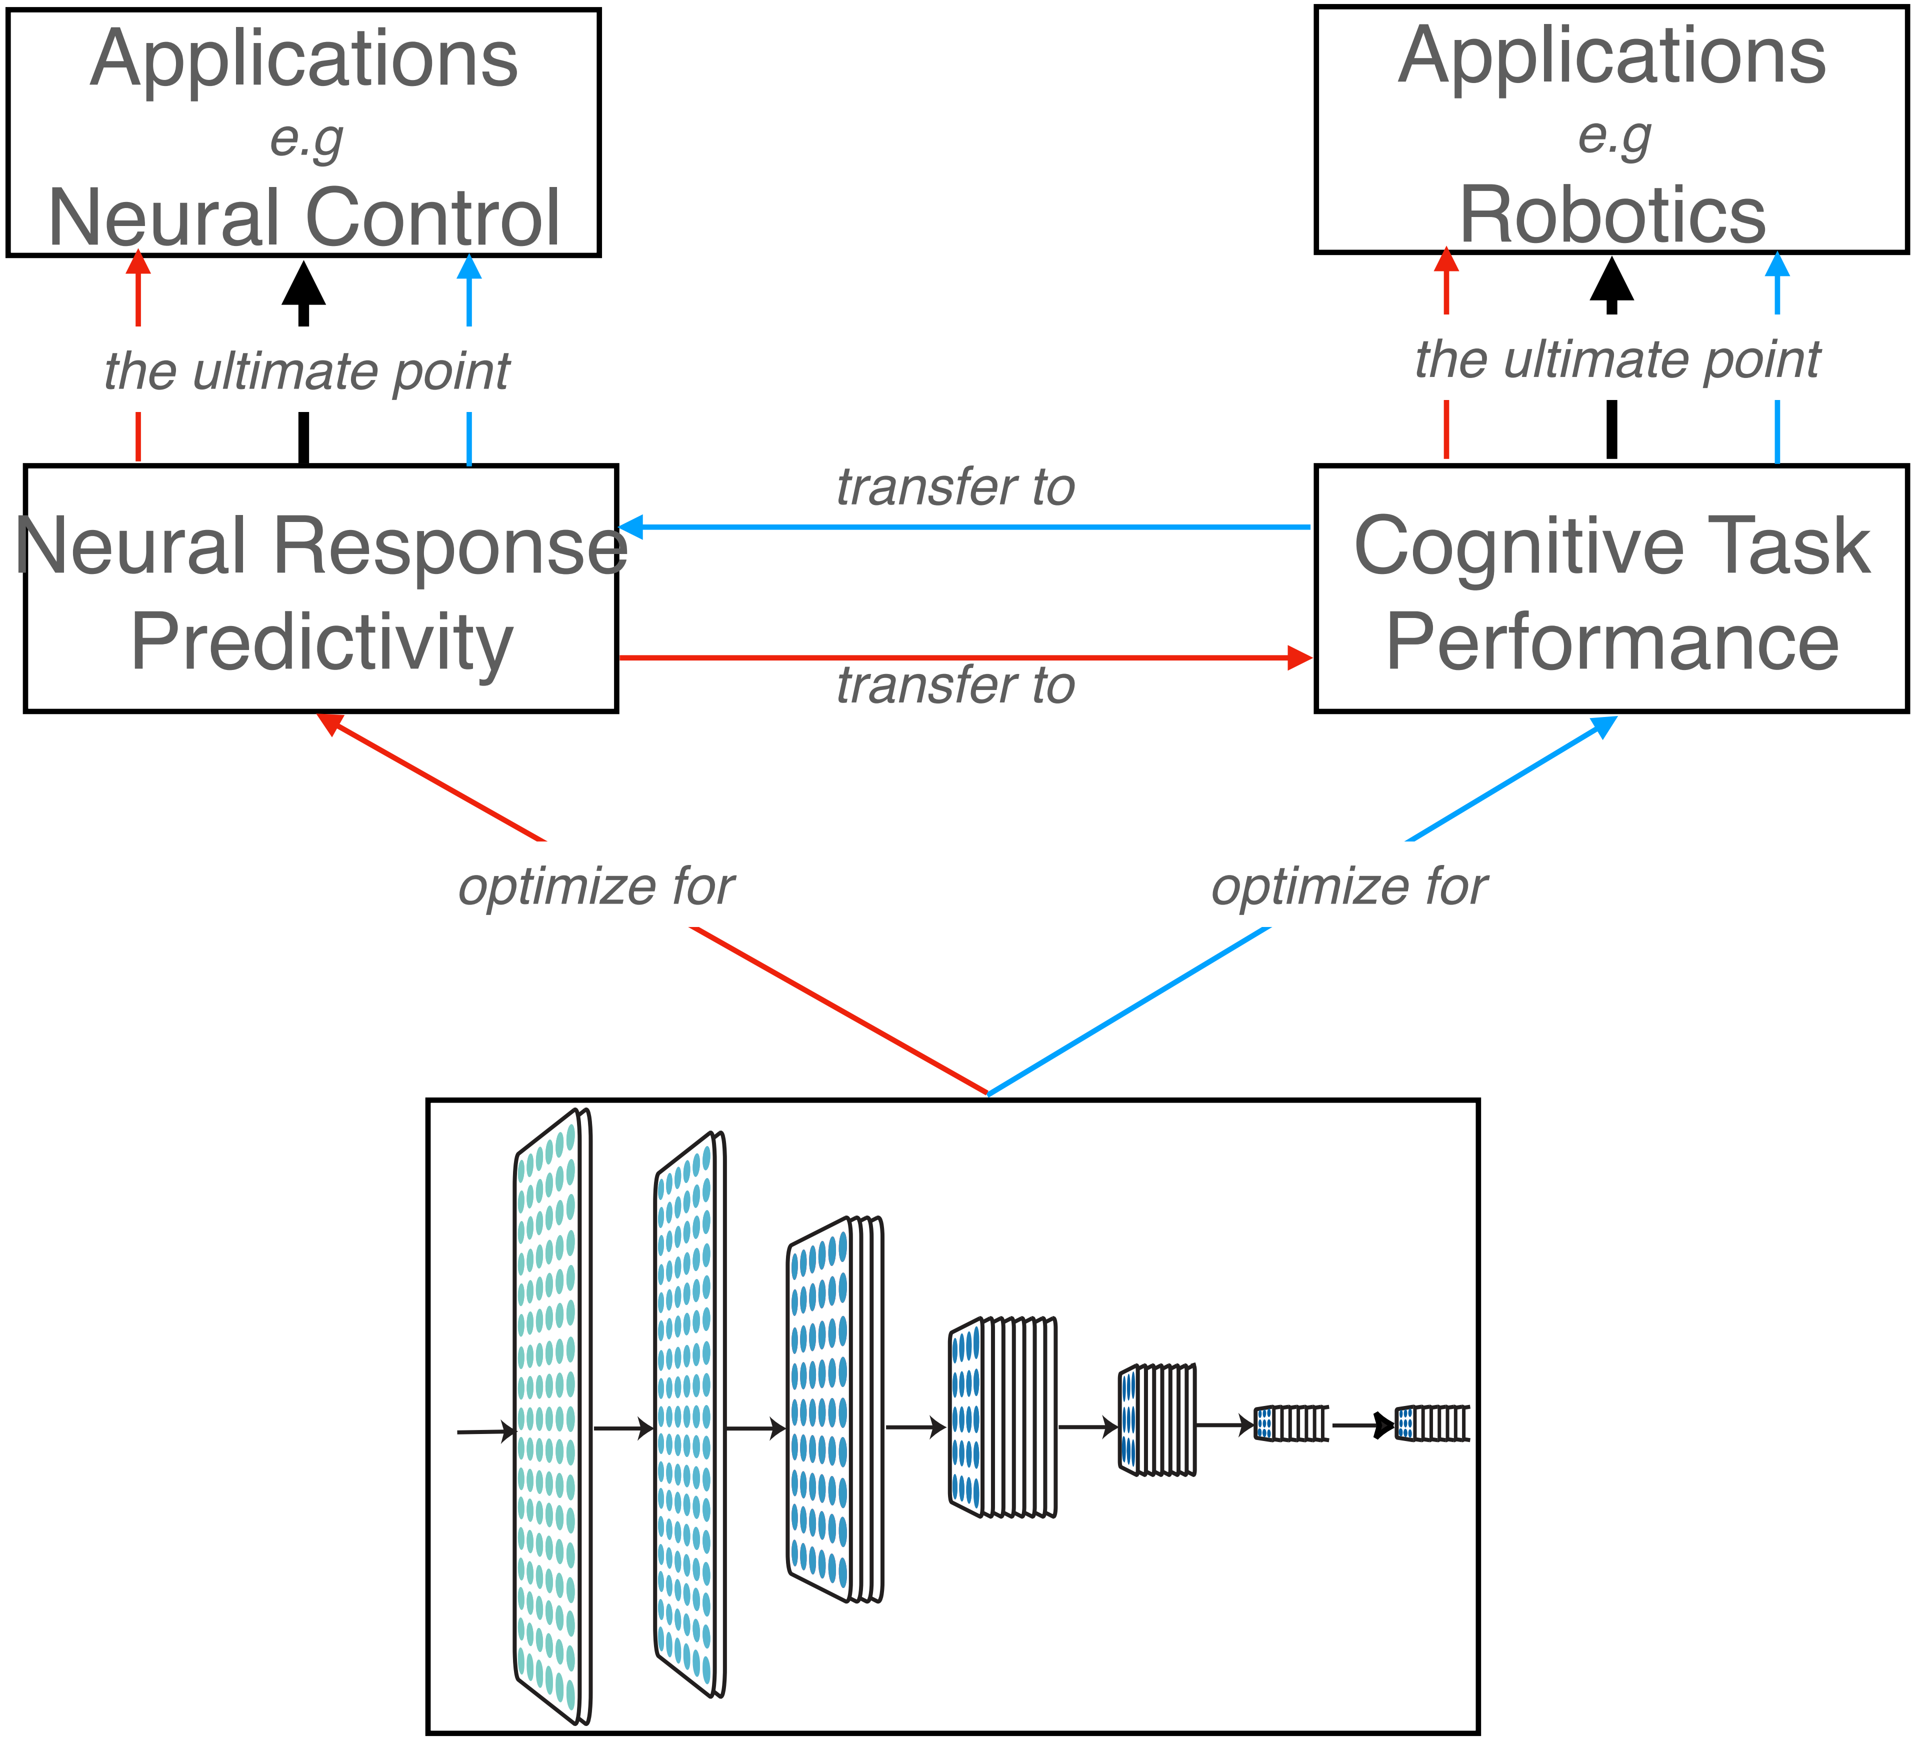
Applications
e.g
Neural Control
Applications
e.g
Robotics
the ultimate point
the ultimate point
Neural Response Predictivity
transfer to
transfer to
Cognitive Task Performance
optimize for
optimize for

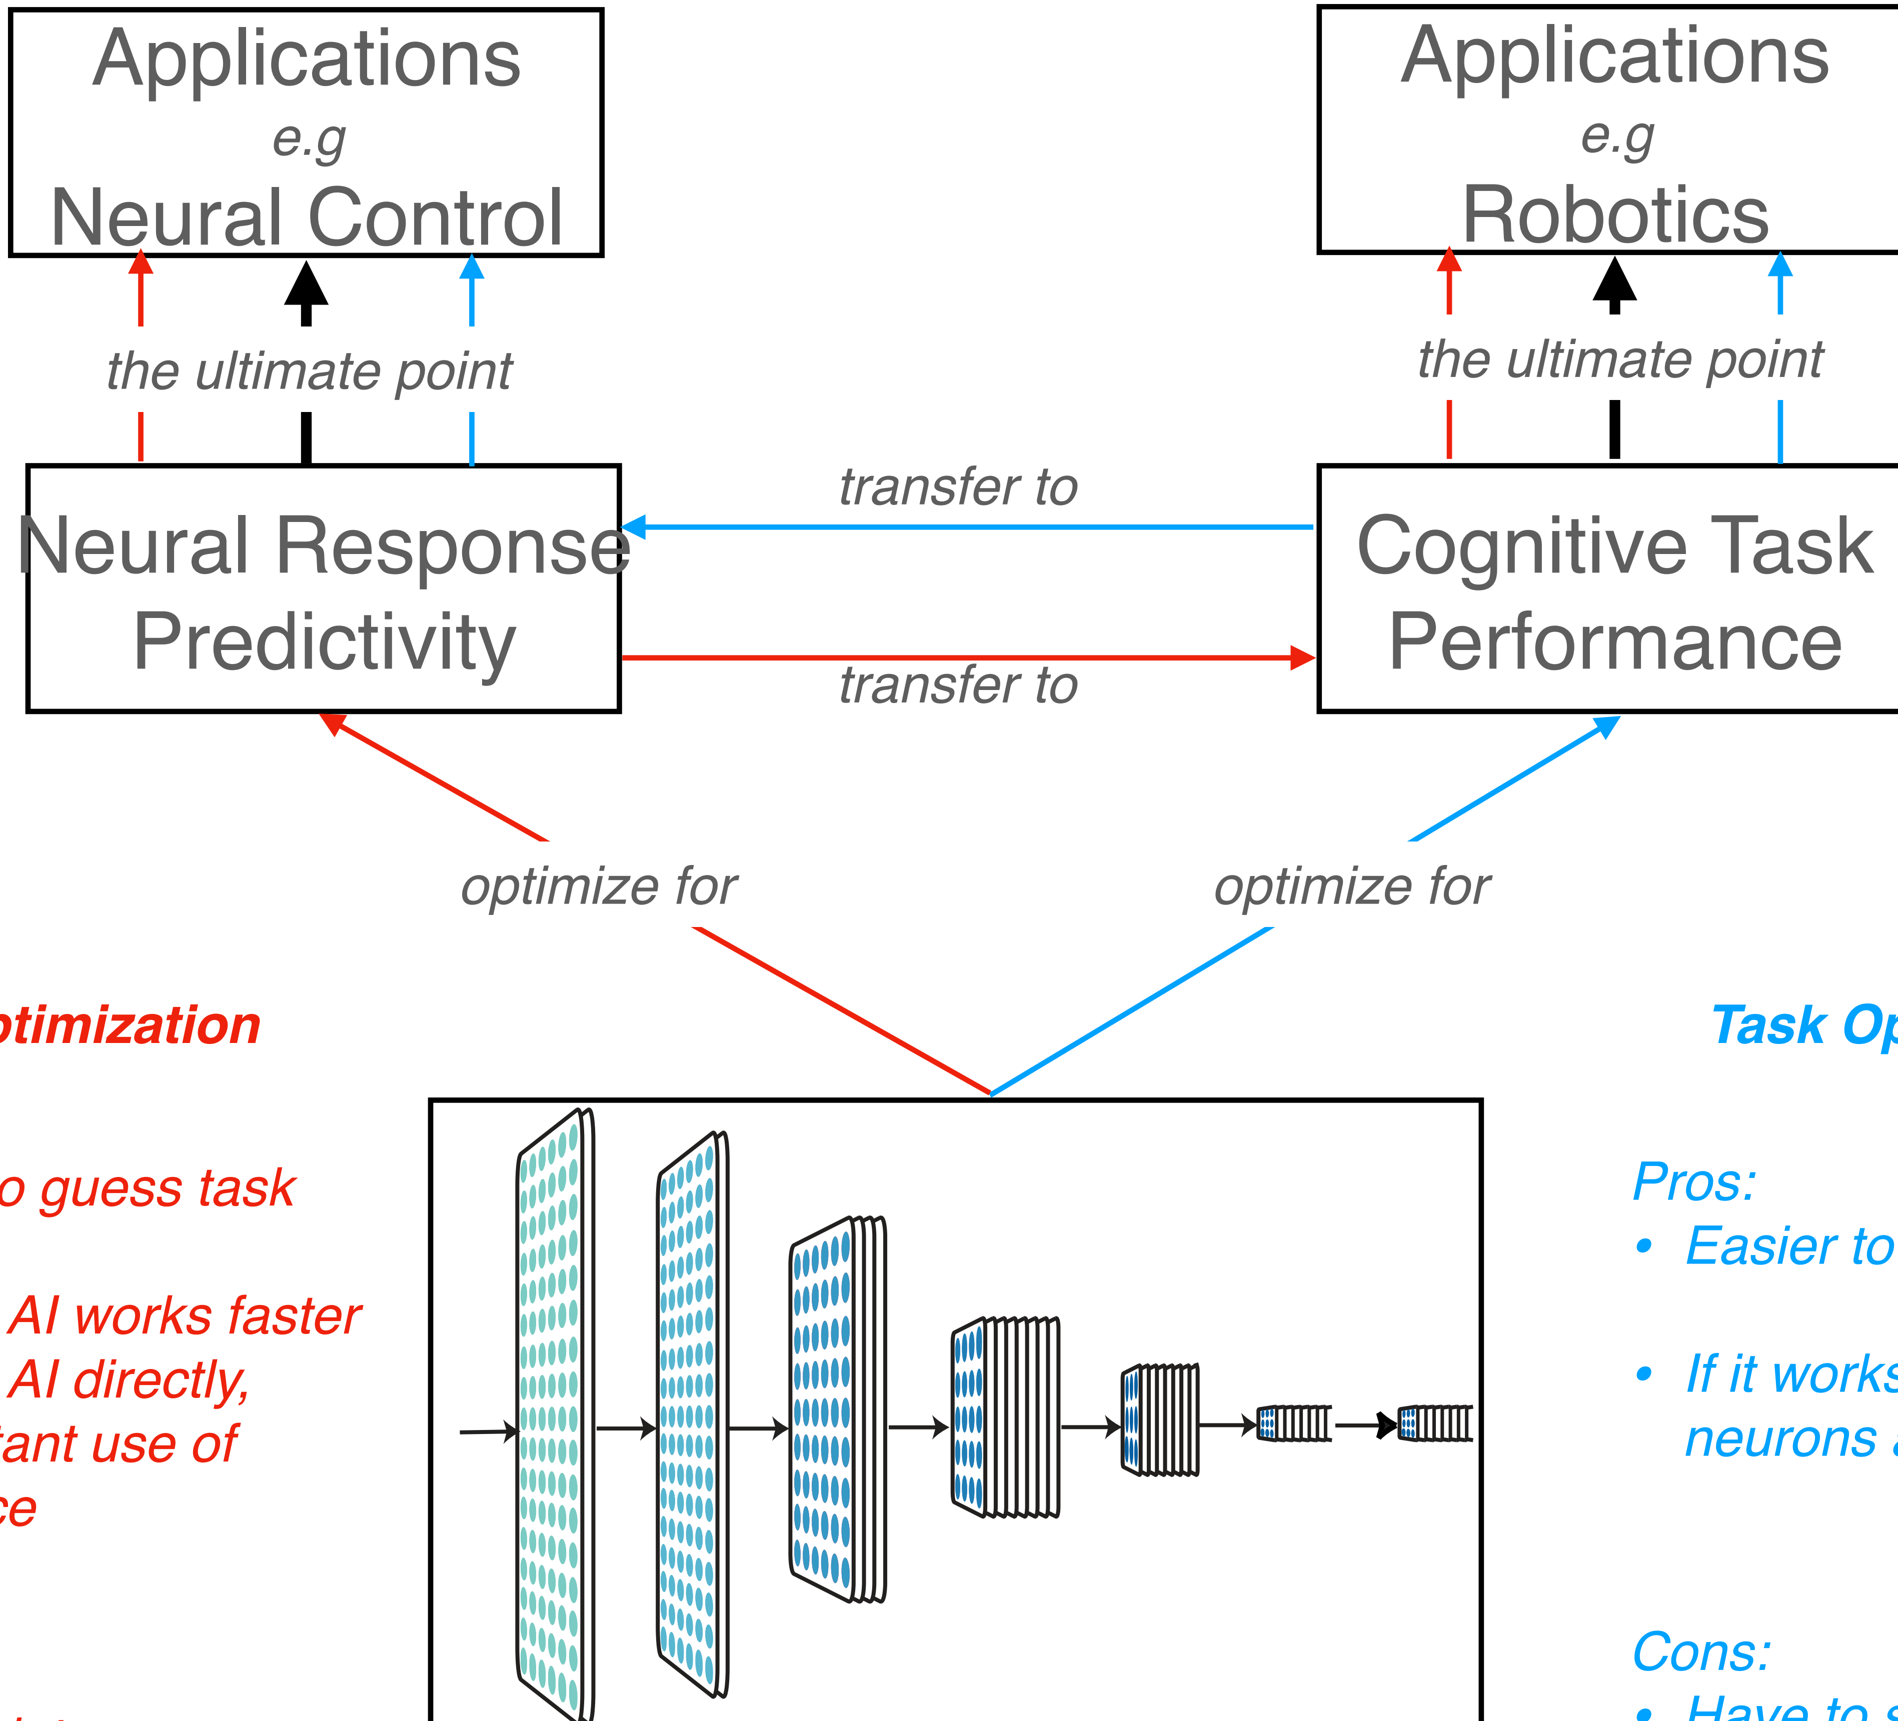

Neural Optimization

Pros:

- *Don't have to guess task*
- *If transfer to AI works faster than solving AI directly, really important use of Neuroscience*

Cons:

- *Hard to get data*
- *Doesn't explain "why" neurons are as they are*

Task Optimization

Pros:

- *Easier to get training data*
- *If it works, it explains "why" neurons are as they are*

Cons:

- *Have to shoot in the dark as to task choice*

Neuroscience Methods

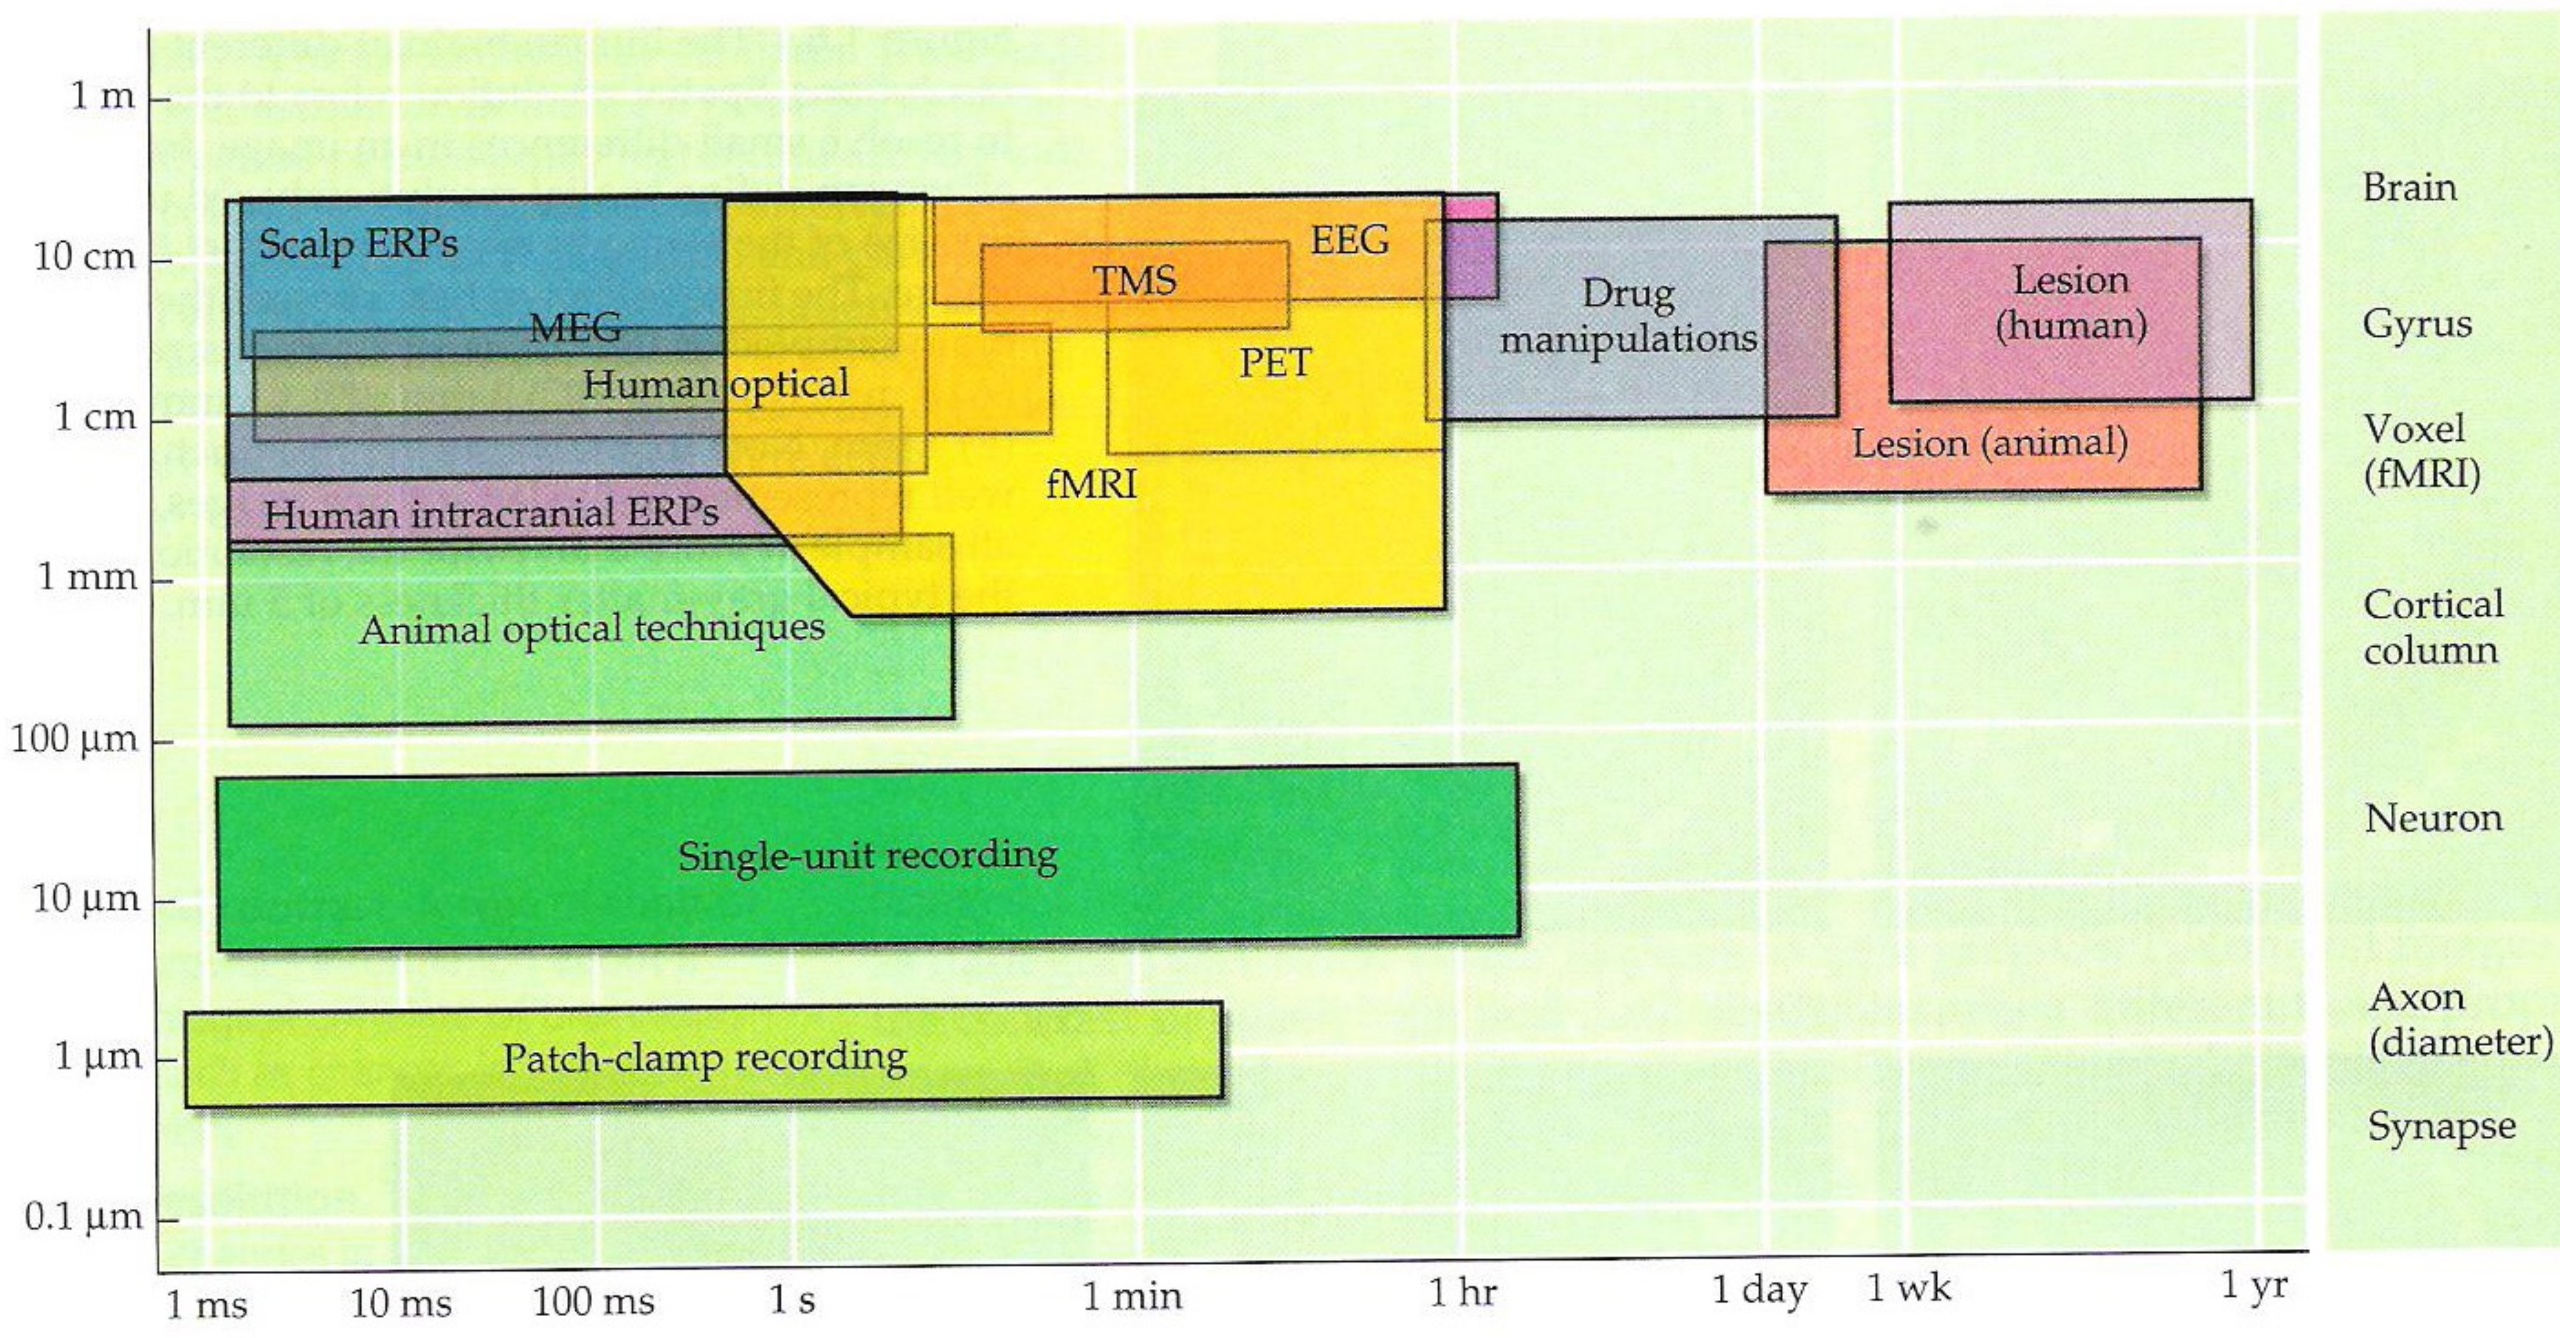

Neuroscience Methods

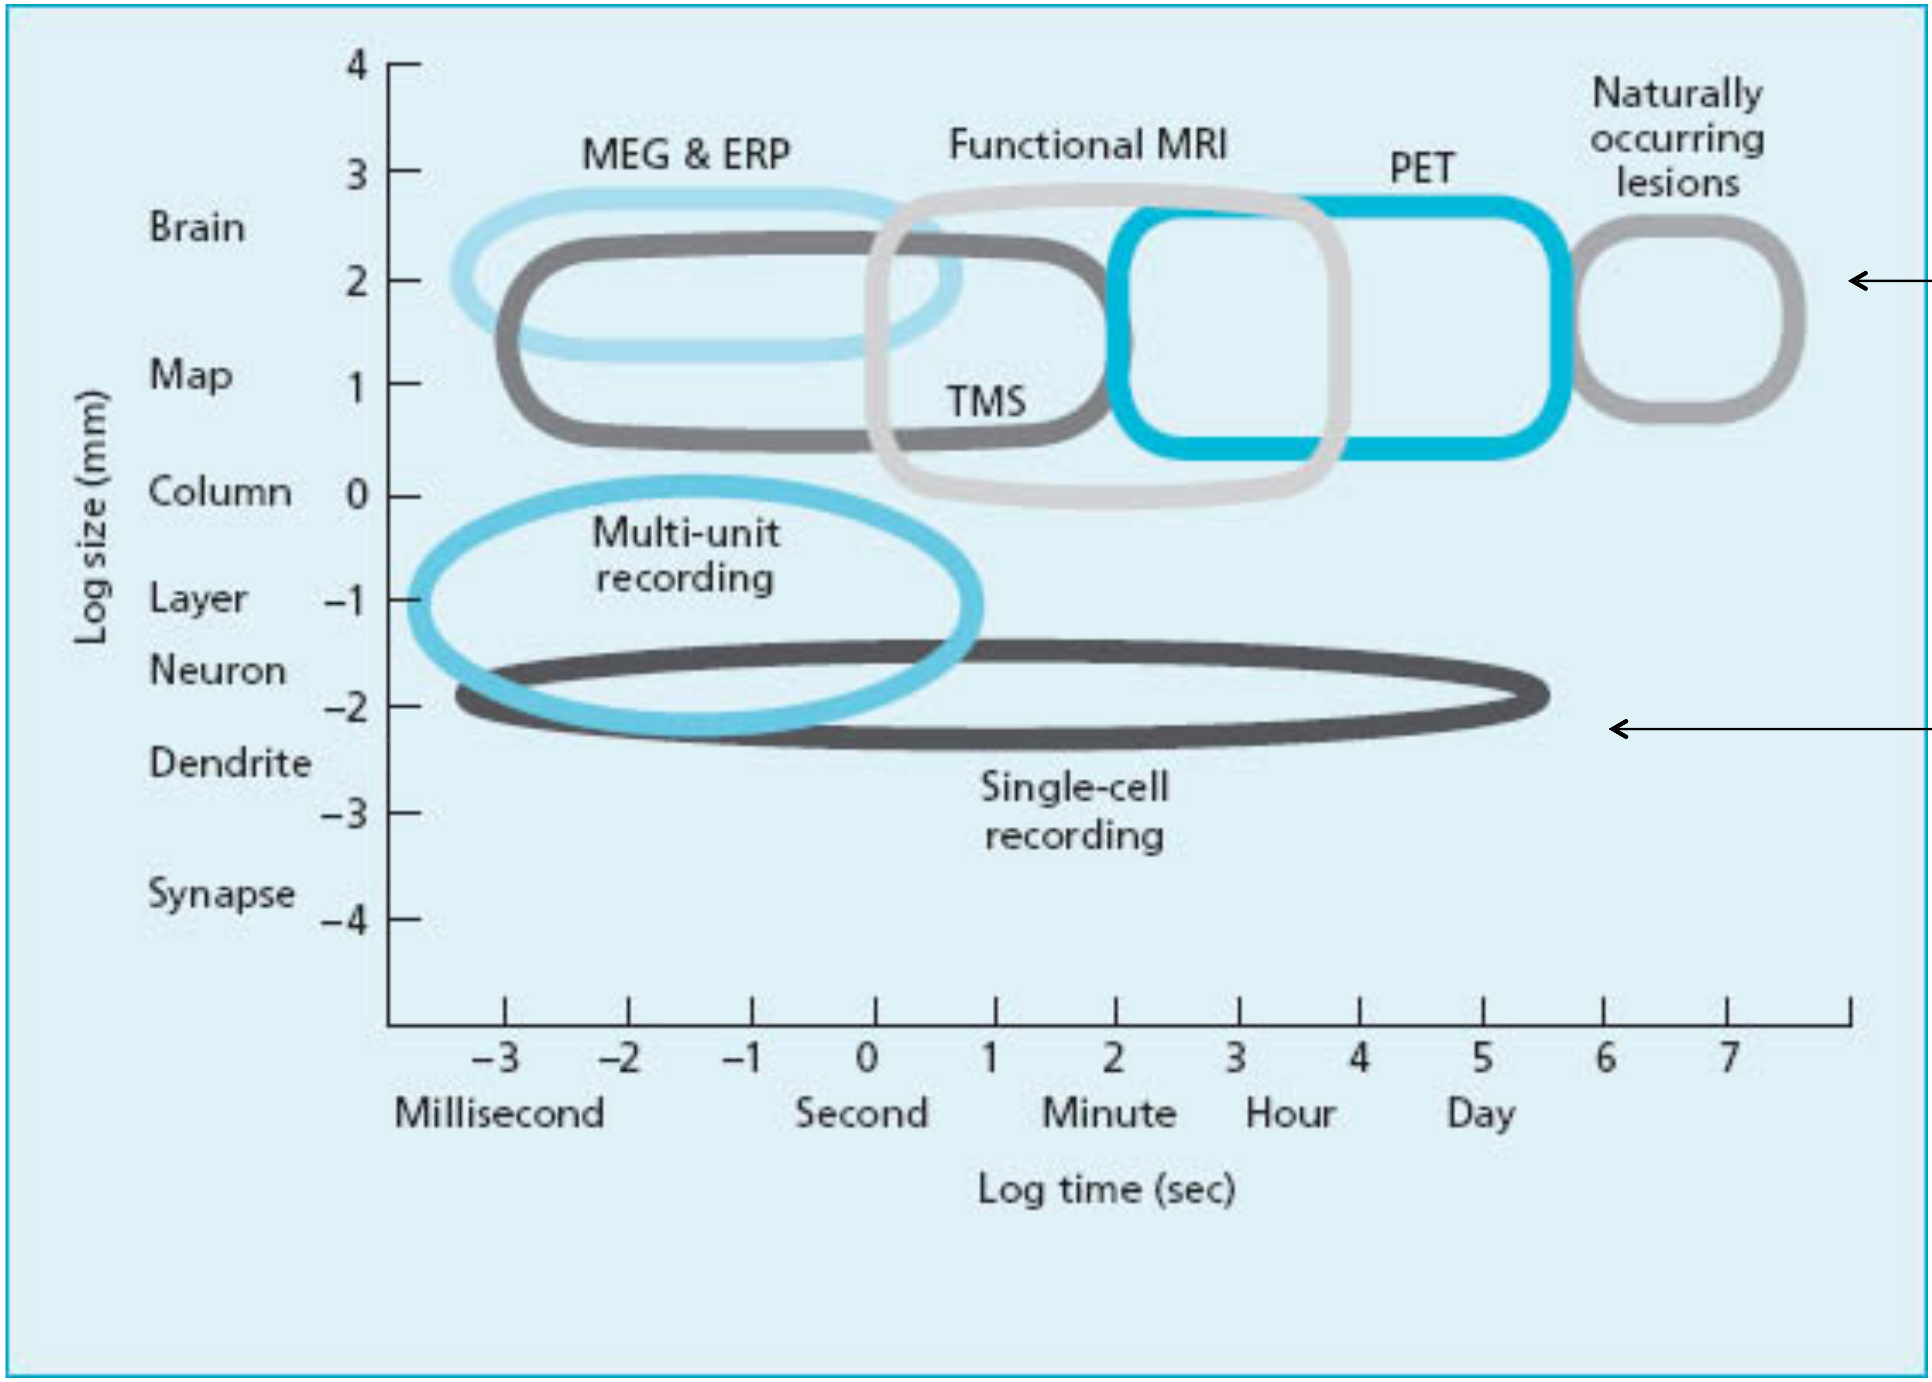

Methods available for studying awake behaving humans

can be used in awake behaving Macaques

Class Outline

Tour of Session Calendar

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Vision

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Perception

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Perception

Beyond Perception

Course Session Outline

Date	Session	
01/05	Introduction to NeuroAI	
01/07	Visual Systems Neuroscience Background	**Perception**
01/12	DNN Models of the Visual System	
01/14	Model-Brain Mapping Methods	
01/19	[NO CLASS-MLK DAY]	
01/21	Unsupervised Learning and the Brain	
01/26	Recurrent Model of Vision	
01/27		
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development	
02/02	Topography and Functional Organization	
02/04	**Andreas Tolias (Stanford):** The Enigma Project	
02/09	Auditory and Somatosensory Models	
02/11	Memory and the Hippocampus	**Memory**
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)	
02/18	Navigation and the MEC	
02/23	**Aran Nayebi (CMU):** Models of Agents	
02/25	The Motor System	
02/26		
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain	
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain	
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution	
03/11	Learning Rules in the Brain	
03/16	Project Presentations	
03/22		

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Perception

Memory

Motor

Course Session Outline

Date	Session	
01/05	Introduction to NeuroAI	
01/07	Visual Systems Neuroscience Background	**Perception**
01/12	DNN Models of the Visual System	
01/14	Model-Brain Mapping Methods	
01/19	[NO CLASS-MLK DAY]	
01/21	Unsupervised Learning and the Brain	
01/26	Recurrent Model of Vision	
01/27		
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development	
02/02	Topography and Functional Organization	
02/04	**Andreas Tolias (Stanford):** The Enigma Project	
02/09	Auditory and Somatosensory Models	
02/11	Memory and the Hippocampus	**Memory**
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)	
02/18	Navigation and the MEC	
02/23	**Aran Nayebi (CMU):** Models of Agents	
02/25	The Motor System	**Motor**
02/26		
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain	
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain	**Language**
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution	
03/11	Learning Rules in the Brain	
03/16	Project Presentations	
03/22		

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Perception

Memory

Cognition and Agency

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Perception

Memory

Cognition and Agency

Learning & Evolution

Course Session Outline

Date	Session
01/05	Introduction to NeuroAI
01/07	Visual Systems Neuroscience Background
01/12	DNN Models of the Visual System
01/14	Model-Brain Mapping Methods
01/19	[NO CLASS-MLK DAY]
01/21	Unsupervised Learning and the Brain
01/26	Recurrent Model of Vision
01/27	
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development
02/02	Topography and Functional Organization
02/04	**Andreas Tolias (Stanford):** The Enigma Project
02/09	Auditory and Somatosensory Models
02/11	Memory and the Hippocampus
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)
02/18	Navigation and the MEC
02/23	**Aran Nayebi (CMU):** Models of Agents
02/25	The Motor System
02/26	
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution
03/11	Learning Rules in the Brain
03/16	Project Presentations
03/22	

Neural Dynamics

Logistics

Date	Session	Assignments
01/05	Introduction to NeuroAI	Coding Assignment 1 Released
		(training and evaluating your own neural network)
01/07	Visual Systems Neuroscience Background	
01/12	DNN Models of the Visual System	
01/14	Model-Brain Mapping Methods	
01/19	[NO CLASS-MLK DAY]	
01/21	Unsupervised Learning and the Brain	
01/26	Recurrent Model of Vision	
01/27		Coding Assignment 1 Due
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development	Coding Assignment 2 Released
02/02	Topography and Functional Organization	
02/04	**Andreas Tolias (Stanford):** The Enigma Project	
02/09	Auditory and Somatosensory Models	
02/11	Memory and the Hippocampus	
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)	Coding Assignment 2 Due
02/18	Navigation and the MEC	
02/23	**Aran Nayebi (CMU):** Models of Agents	
02/25	The Motor System	
02/26		Final Project Validation Due
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain	
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain	
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution	
03/11	Learning Rules in the Brain	
03/16	Project Presentations	
03/22		Presentation Write-Ups Due

Date	Session	Assignments
01/05	Introduction to NeuroAI	Coding Assignment 1 Released (training and evaluating your own neural network)
01/07	Visual Systems Neuroscience Background	
01/12	DNN Models of the Visual System	
01/14	Model-Brain Mapping Methods	
01/19	[NO CLASS-MLK DAY]	
01/21	Unsupervised Learning and the Brain	
01/26	Recurrent Model of Vision	
01/27		Coding Assignment 1 Due
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development	Coding Assignment 2 Released
02/02	Topography and Functional Organization	
02/04	**Andreas Tolias (Stanford):** The Enigma Project	
02/09	Auditory and Somatosensory Models	
02/11	Memory and the Hippocampus	
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)	Coding Assignment 2 Due
02/18	Navigation and the MEC	
02/23	**Aran Nayebi (CMU):** Models of Agents	
02/25	The Motor System	
02/26		Final Project Validation Due
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain	
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain	
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution	
03/11	Learning Rules in the Brain	
03/16	Project Presentations	
03/22		Presentation Write-Ups Due

Date	Session	Assignments
01/05	Introduction to NeuroAI	Coding Assignment 1 Released
		(training and evaluating your own neural network)
01/07	Visual Systems Neuroscience Background	
01/12	DNN Models of the Visual System	
01/14	Model-Brain Mapping Methods	
01/19	[NO CLASS-MLK DAY]	
01/21	Unsupervised Learning and the Brain	
01/26	Recurrent Model of Vision	
01/27		Coding Assignment 1 Due
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development	Coding Assignment 2 Released
02/02	Topography and Functional Organization	
02/04	**Andreas Tolias (Stanford):** The Enigma Project	
02/09	Auditory and Somatosensory Models	
02/11	Memory and the Hippocampus	
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)	Coding Assignment 2 Due
02/18	Navigation and the MEC	
02/23	**Aran Nayebi (CMU):** Models of Agents	
02/25	The Motor System	
02/26		Final Project Validation Due
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain	
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain	
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution	
03/11	Learning Rules in the Brain	
03/16	Project Presentations	
03/22		Presentation Write-Ups Due

Date	Session	Assignments
01/05	Introduction to NeuroAI	Coding Assignment 1 Released (training and evaluating your own neural network)
01/07	Visual Systems Neuroscience Background	
01/12	DNN Models of the Visual System	
01/14	Model-Brain Mapping Methods	
01/19	[NO CLASS-MLK DAY]	
01/21	Unsupervised Learning and the Brain	
01/26	Recurrent Model of Vision	
01/27		Coding Assignment 1 Due
01/28	**Cliona O'Doherty (Stanford):** Modeling Infant Development	Coding Assignment 2 Released
02/02	Topography and Functional Organization	
02/04	**Andreas Tolias (Stanford):** The Enigma Project	
02/09	Auditory and Somatosensory Models	
02/11	Memory and the Hippocampus	
02/16	[NO CLASS-PRESIDENT'S DAY] (BBScore Evening Session)	Coding Assignment 2 Due
02/18	Navigation and the MEC	
02/23	**Aran Nayebi (CMU):** Models of Agents	
02/25	The Motor System	
02/26		Final Project Validation Due
03/02	**Scott Lindermann (Stanford):** Dynamical Systems in the Brain	
03/04	**Greta Tuckute (Harvard):** Language, LLMs, and the Brain	
03/09	**Tony Zador (Cold Spring Harbor):** Models of Brain Evolution	
03/11	Learning Rules in the Brain	
03/16	Project Presentations	
03/22		Presentation Write-Ups Due

CS375/Psych279 Homework 1: Training Your Own Neural Network

Overview

In this assignment, you will implement, train, and visualize the behavior of an AlexNet convolutional neural network (CNN) using PyTorch. You will visualize the kernels of the firs layer of the neural network and analyze their response patterns. You will replicate some basic findings of classical work by Hubel and Wiesel in silico by measuring orientation selectivinty of several artificial neurons early in the model. Specifically you will:

1. **Implement the AlexNet model** and understanding its architecture.
2. **Implement a training loop** capable of training the network on the ImageNet dataset.
3. **Measure kernel responses** for various spatial frequencies and orientations of sinusoidal grating stimuli.
4. **Visualize the learned kernels** in the first layer of your model.

CS375/Psych279 Homework 1: Training Your Own Neural Network

Overview

In this assignment, you will implement, train, and visualize the behavior of an AlexNet convolutional neural network (CNN) using PyTorch. You will visualize the kernels of the firs layer of the neural network and analyze their response patterns. You will replicate some basic findings of classical work by Hubel and Wiesel in silico by measuring orientation selectivinty of several artificial neurons early in the model. Specifically you will:

1. **Implement the AlexNet model** and understanding its architecture.
2. **Implement a training loop** capable of training the network on the ImageNet dataset.
3. **Measure kernel responses** for various spatial frequencies and orientations of sinusoidal grating stimuli.
4. **Visualize the learned kernels** in the first layer of your model.

Submission Instructions

1. **Code**
   - Submit your modified `train.py` file with all tasks completed. Include your name at the top of the assigment.
2. **Report**
   - Provide a PDF or Markdown report that includes:
     - A brief explanation of the code you implemented
     - An image of the accuracy, loss and circular variance plot, along with a description of the final accuracy values. Observer the trends in loss decrease, accuracy increase and kernel circular variance. Specifically remark on when during the training do the filters seem to get tuned for direciton selectivity.
     - An image of the kernel visualziation plot of the first layer along with a brief description of some qualitative properties of some of the filters.
     - Visualizations of 3 individual filters of your choice and their rotation and frequency selectivity plots. Pick filters that illustrate a clear bias and describe what they seem to be selective for.

CS375/Psych279 Homework 1: Training Your Own Neural Network

Overview

In this assignment, you will implement, train, and visualize the behavior of an AlexNet convolutional neural network (CNN) using PyTorch. You will visualize the kernels of the firs layer of the neural network and analyze their response patterns. You will replicate some basic findings of classical work by Hubel and Wiesel in silico by measuring orientation selectivinty of several artificial neurons early in the model. Specifically you will:

1. **Implement the AlexNet model** and understanding its architecture.
2. **Implement a training loop** capable of training the network on the ImageNet dataset.
3. **Measure kernel re**
4. **Visualize the lear**

Policy:

Feel free to use ChatGPT ; but please do your own work and and run your own training.

Submission Instr

1. **Code**
   - Submit your modified `train.py` file with all tasks completed. Include your name at the top of the assigment.
2. **Report**
   - Provide a PDF or Markdown report that includes:
     - A brief explanation of the code you implemented
     - An image of the accuracy, loss and circular variance plot, along with a description of the final accuracy values. Observer the trends in loss decrease, accuracy increase and kernel circular variance. Specifically remark on when during the training do the filters seem to get tuned for direciton selectivity.
     - An image of the kernel visualziation plot of the first layer along with a brief description of some qualitative properties of some of the filters.
     - Visualizations of 3 individual filters of your choice and their rotation and frequency selectivity plots. Pick filters that illustrate a clear bias and describe what they seem to be selective for.

The Final Project

https://github.com/neuroailab/bbscore_public/

README

BBScore: Brain-Behavior Scoring Framework

BBScore is a comprehensive framework for benchmarking deep learning models against neural (fMRI, ephys) and behavioral datasets. It handles the complex pipeline of loading model weights, preprocessing stimuli (images/videos), extracting feature activations, and scoring them against biological data.

📁 Project Structure

- `benchmarks/`: Definitions tying data and scoring together (e.g., `NSD`, `Algonauts`).
- `data/`: Scripts to download and preprocess datasets (Stimuli and Neural assemblies).
- `metrics/`: Mathematical implementations of scores (Ridge, RSA, PLS).
- `models/`: Wrappers for deep learning models (HuggingFace, TorchVision, Custom).
- `mongo_utils/`: Helpers for database injection (Advanced use).

Attendance & Participation Policy

1. You must attend class in person

2. You must be on time

3. You must participate!

If you can't make it to a given session, let us know more than 24 hours in advance. And, don't let it happen much!

You are expected to be present in person for the final presentation.

Logistics

Website: cs375.stanford.edu

Class Structure: Mixture of Lectures and Guest Lectures

Assignments: 2 coding assignments, final project proposal, final project presentation & writeup

Grading: participation (25%), coding assignments (40%), project presentation (15%) project write-up (20%)

Office Hours: Wednesdays 4:30-5:30, Wu Tsai Neuro Institute 2nd Floor lounge

Tools: http://cs375.stanford.edu/software-tools.html