Critique of Pure Learning

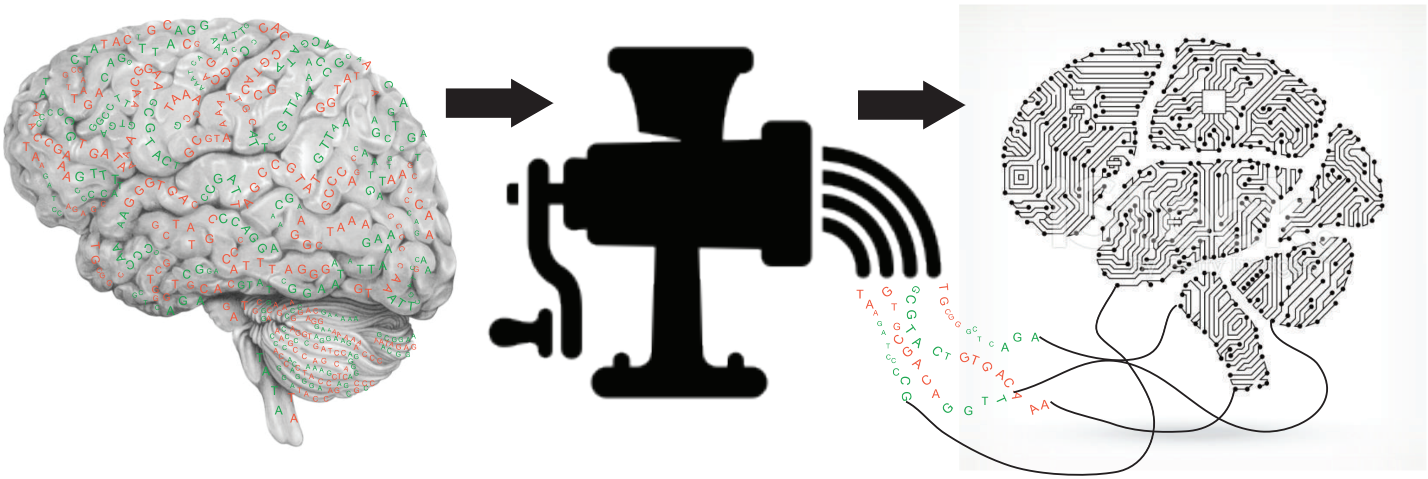

Anthony Zador
Cold Spring Harbor Laboratory

Goal: Build machines with general intelligence

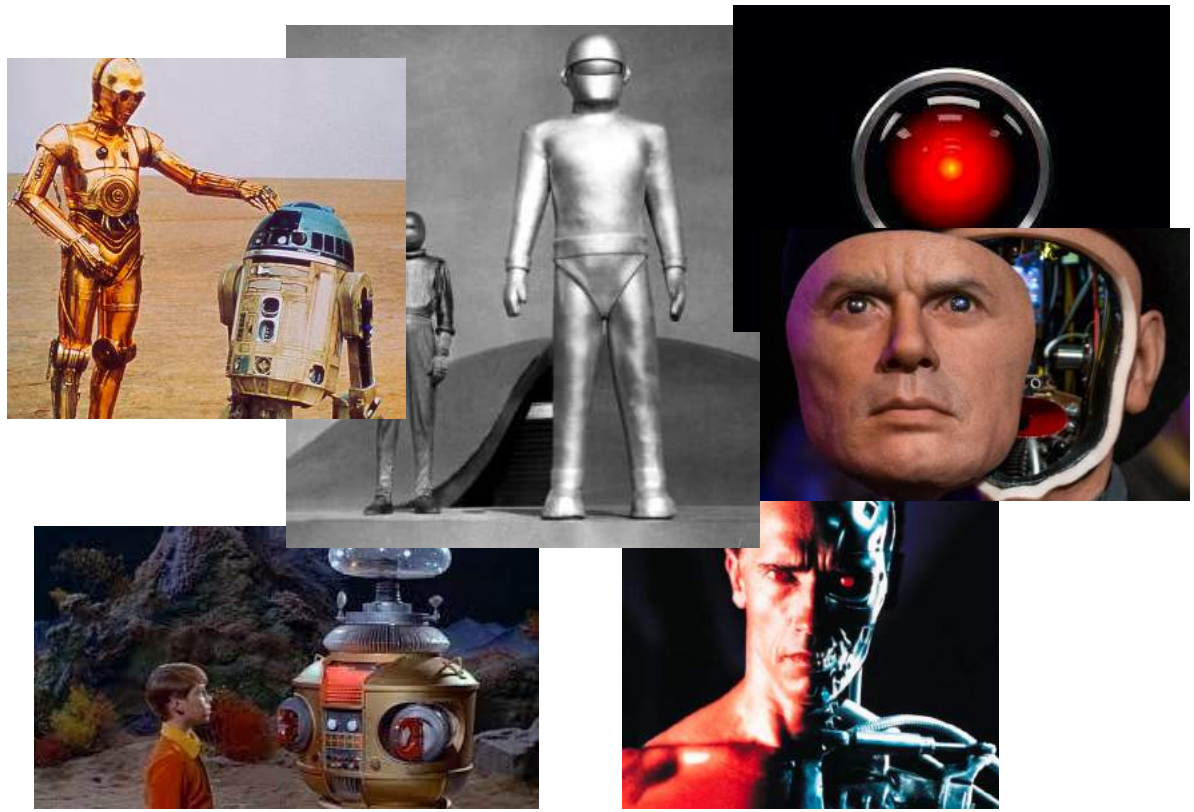

Why aren't we there yet??!?
And how do we get there?

Traditional GOFAI

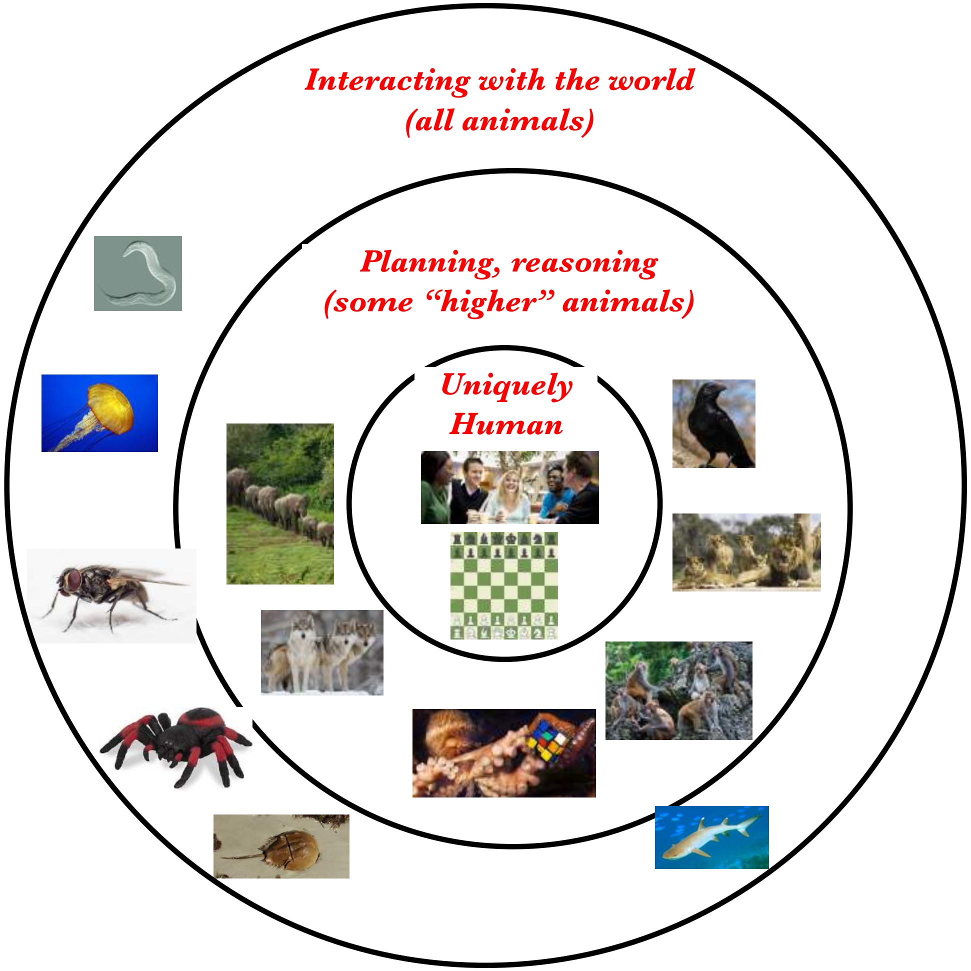

Sapiocentric AI

Moravec's Paradox

"Encoded in the large, highly evolved sensory and motor portions of the human brain is a billion years of experience about the nature of the world and how to survive in it. The deliberate process we call reasoning is, I believe, the thinnest veneer of human thought, effective only because it is supported by this much older and much more powerful, though usually unconscious, sensorimotor knowledge. We are all prodigious Olympians in perceptual and motor areas, so good that we make the difficult look easy. Abstract thought, though, is a new trick, perhaps less than 100 thousand years old. We have not yet mastered it. It is not all that intrinsically difficult; it just seems so when we do it."

--- AI Pioneer Hans Moravec (1988)

<u>*What animals do is hard—cognition and reasoning is the easy part*</u>

Copernican revolution

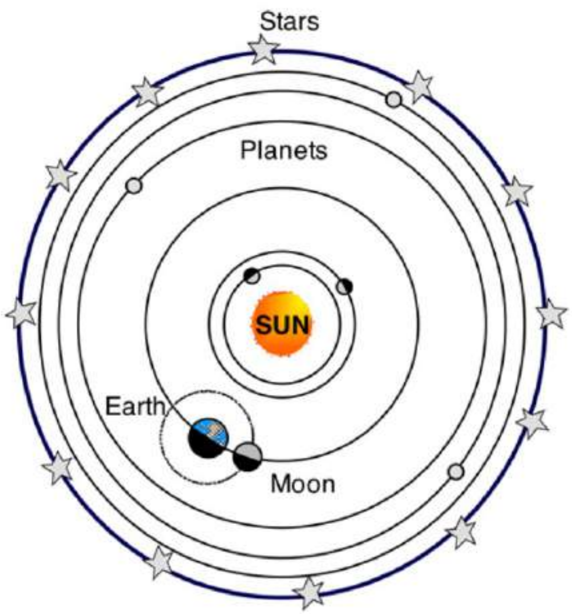

Heliocentric Theory

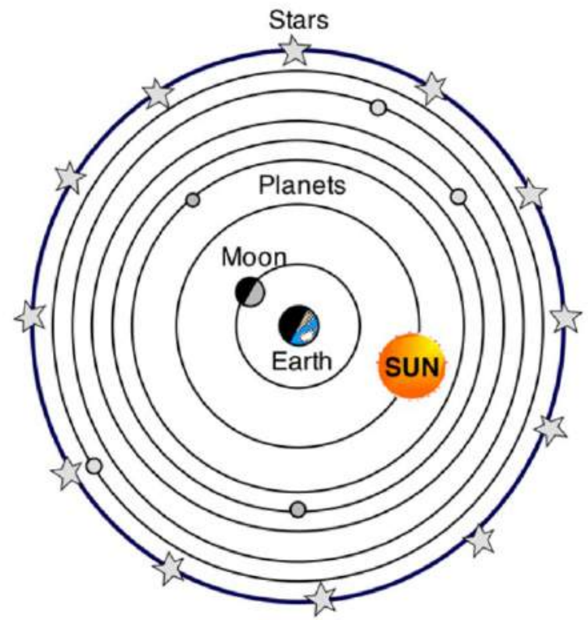

Geocentric Theory

Moravecian revolution

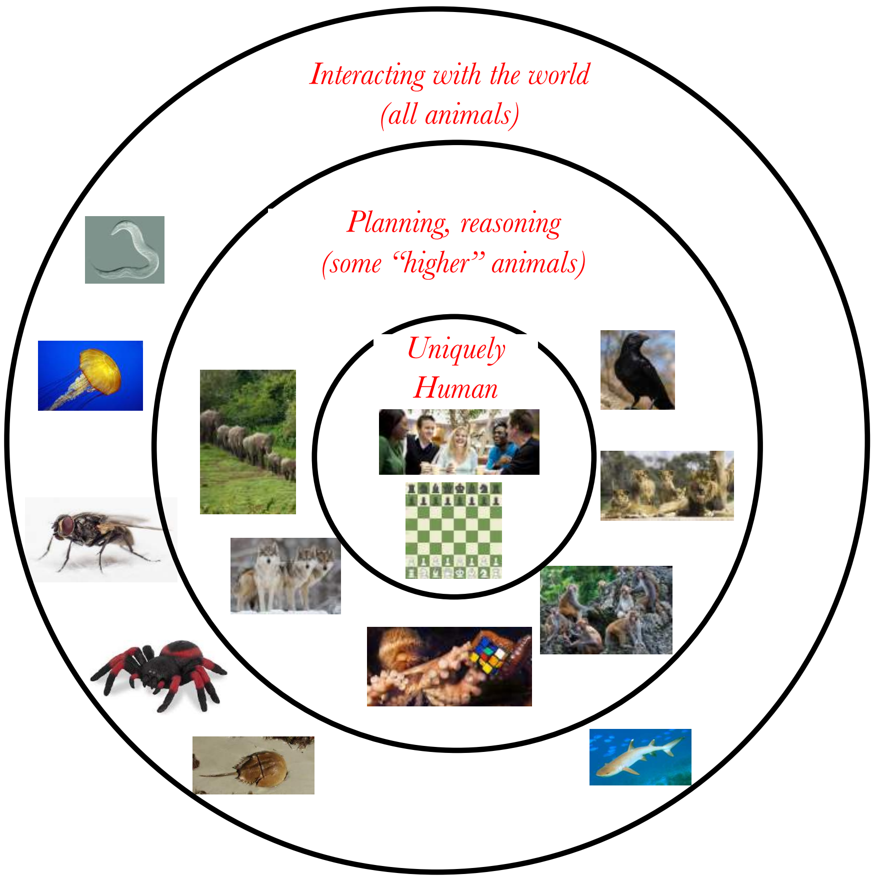

Sapiocentric AI

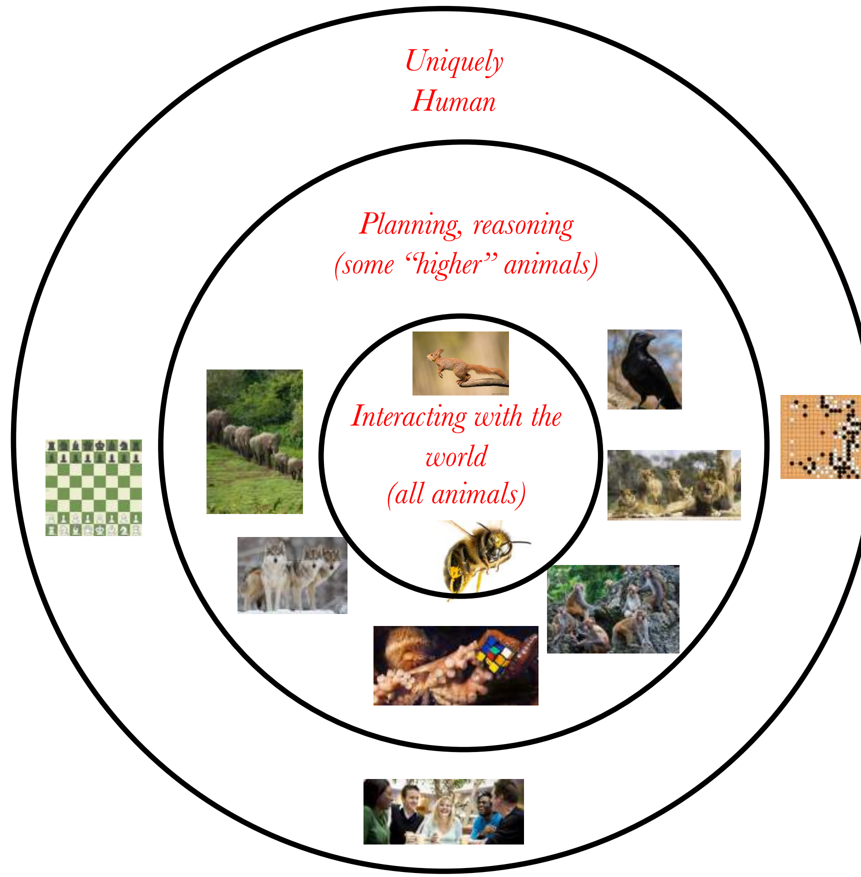

Zoocentric AI

Our goal: Achieve mouse-level AI

Neuro

AI

Mouse-level AI would provide a solid foundation

NeuroAI Scholars Program at CSHL

- For people with PhDs in AI; not a traditional postdoc
- Focus on applying insights from neuro to AI
- Scholars are imbedded for 2 years in CSHL neuro labs
- Scholars are independent, expected to forge new paths

https://www.cshl.edu/research/neuroscience/neuroai/

Outline

AI relies mainly on *learning*; animals rely on *innate structure*

→ *Animal learning is no better than machine learning.*

Innate structure: Genome → wiring diagram

Wiring diagram is compressed through "genomic bottleneck."

→ *Genomic bottleneck algorithm*

Outline

AI relies mainly on *learning*; animals rely on *innate structure*

→ *Animal learning is no better than machine learning.*

Innate structure: Genome → wiring diagram

Wiring diagram is compressed through "genomic bottleneck."

→ *Genomic bottleneck algorithm*

AI relies on learning

Backpropagation

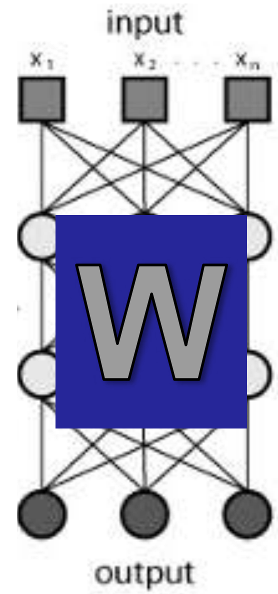

Many labeled pictures

Many synaptic weights

Networks start tabula rasa—blank slates.
So the need a lot of labeled examples to learn all those parameters!

Humans do not rely mostly on supervised learning

Labeling objects every second for a year would only be 3×10^7 training examples

How about unsupervised learning?

Reinforcement learning

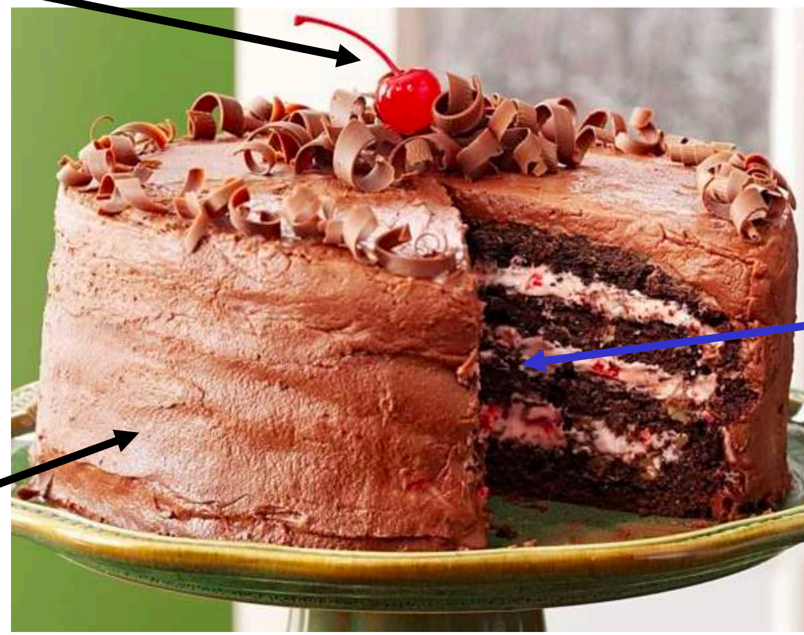

Unsupervised learning

Supervised learning

"If intelligence is a cake, the bulk of the cake is unsupervised learning, the icing on the cake is supervised learning, and the cherry on the cake is reinforcement learning." ---Yann Lecun (2016)

Even unsupervised learning isn't enough

< 10^7 seconds

< 10^2 seconds

< 10 seconds

A lot of behavior is largely innate

Termite mound

Bird nest

Beaver dam

Beaver damming is not learned from parents

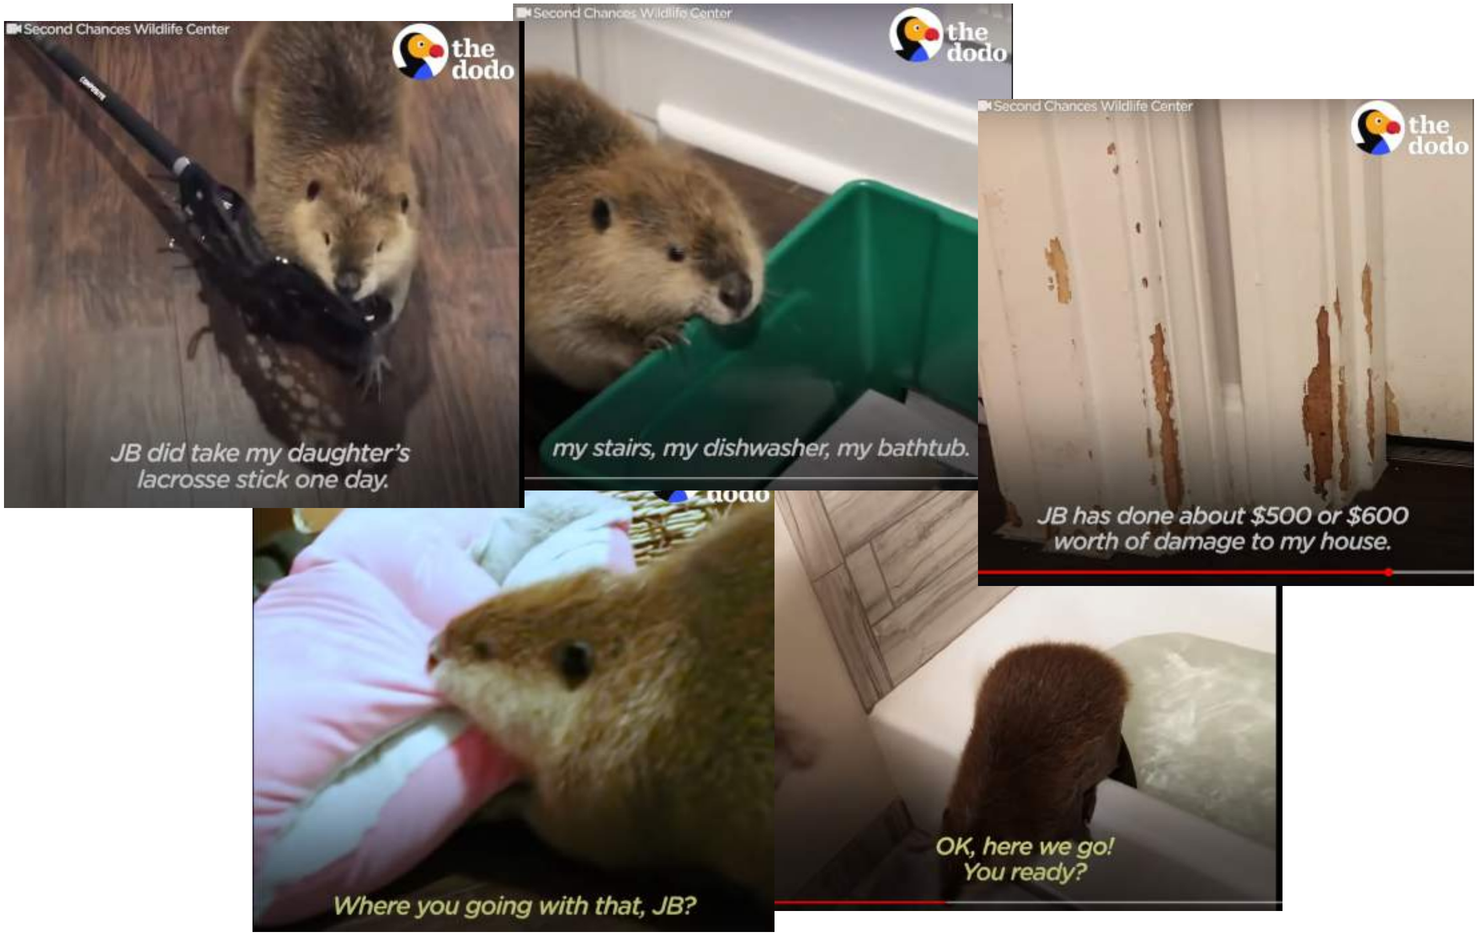

https://www.youtube.com/watch?v=DggHeuhpFvg

"Justin Beaver" the rescue beaver

Burrow building is not learned from parents

Burrowing in peromyscus

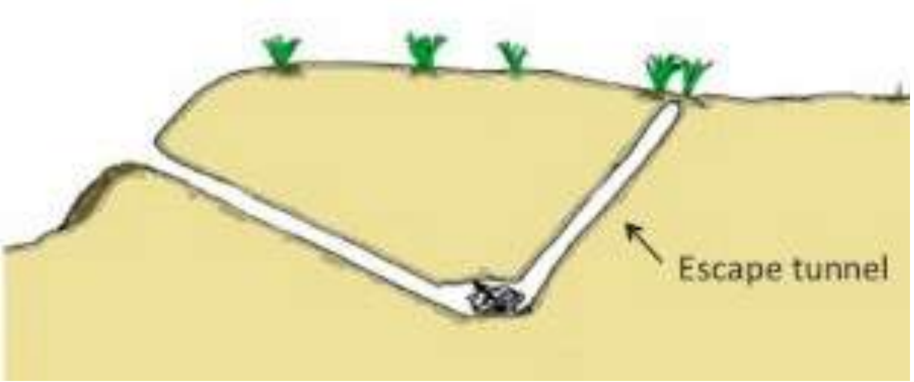

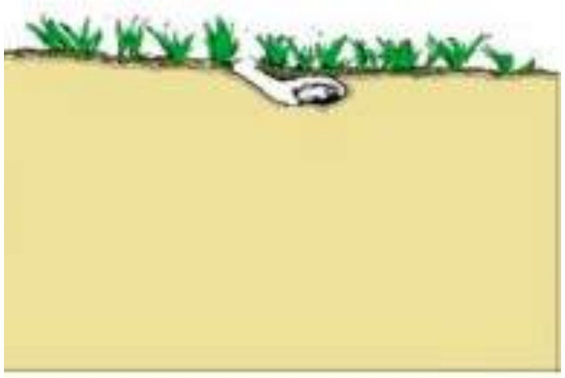

Old field mice

Deer mice

Cross-fostering → genetic, not learned

Hopi Hoekstra

Innate structure provides an evolutionary advantage

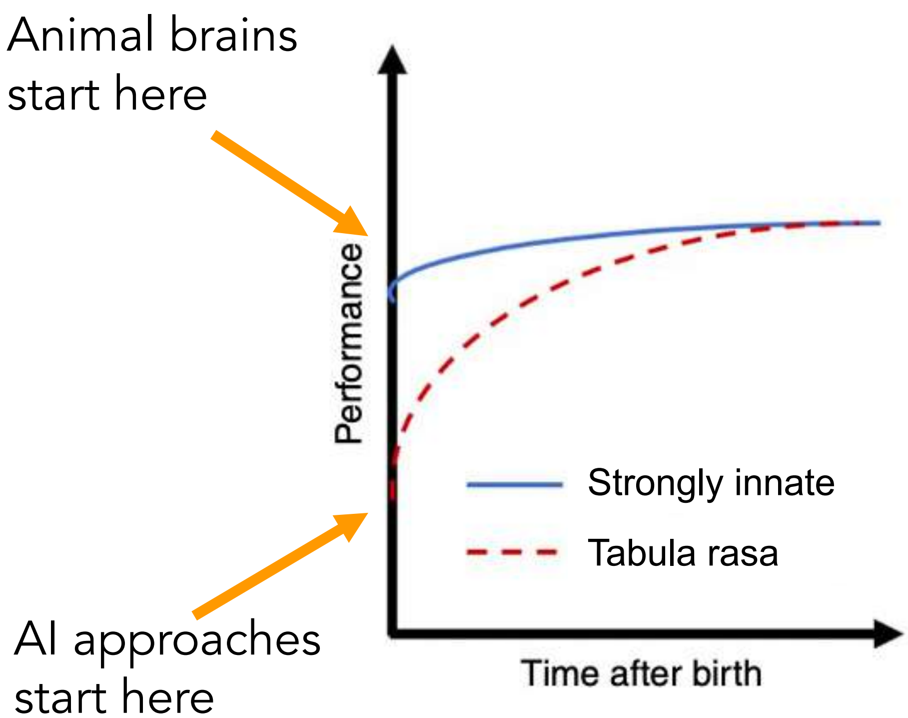

Evolution will typically maximize fitness at birth by incorporating innate structure

Innate priors speed learning

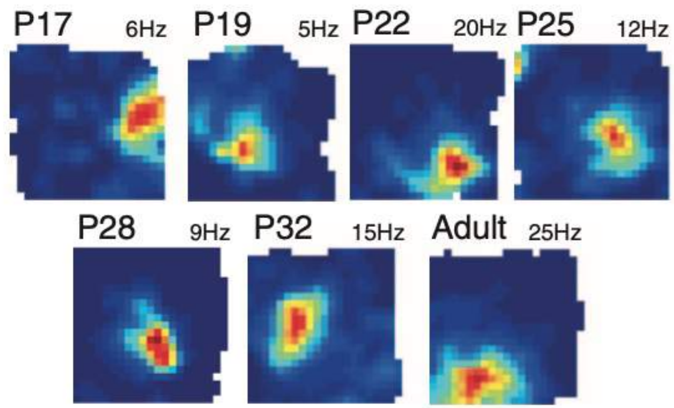

"a rudimentary map of space is already present when 2 1⁄2-week-old rat pups explore an open environment outside the nest for the first time."

The ability to acquire the map is already present

Development of the Spatial Representation System in the Rat

Rosamund F. Langston,[1*†] James A. Ainge,[1,2*] Jonathan J. Couey,[1] Cathrin B. Canto,[1] Tale L. Bjerknes,[1] Menno P. Witter,[1] Edvard I. Moser,[1‡] May-Britt Moser[1]

18 JUNE 2010 VOL 328 **SCIENCE** www.sciencemag.org

Innate priors speed learning: Imprinting

Konrad Lorenz
(Nobel '73)

Innate priors speed learning

Fusiform face area (FFA) in humans specialized for faces

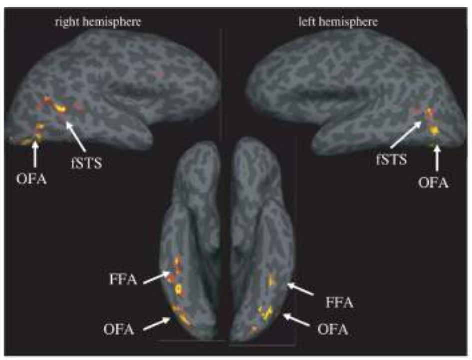

Kanwisher & Yoval 2008
Saygin et al, 2016
Murty et al, 2020

Innate priors enable learning

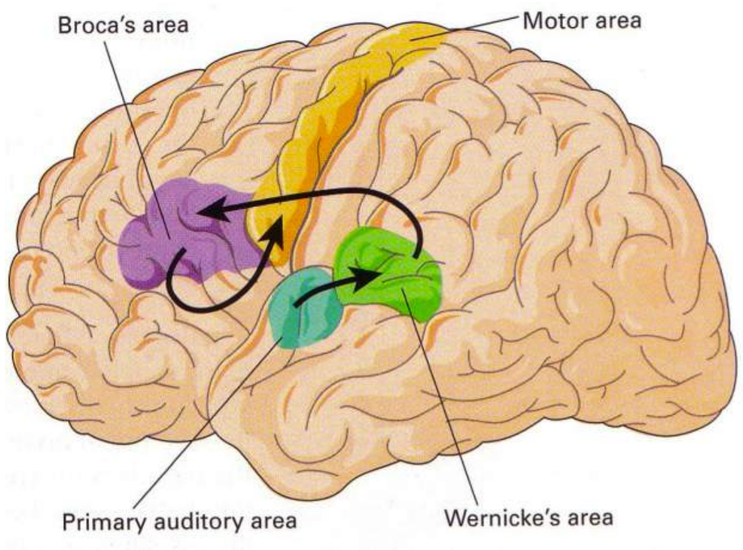

Humans have stereotyped language areas

Koko

- 1000 signs
- No syntax

Animals are about as good at learning as we are

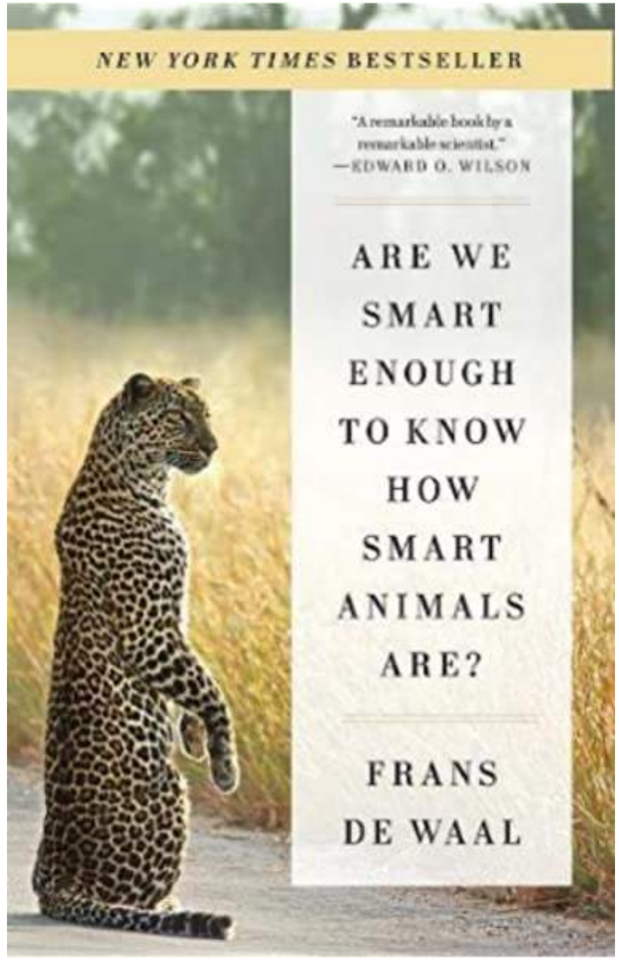

We just have better priors, acquired through cultural transmission

Animals aren't even that good at learning

In the lab, we do not attempt to train animals "end-to-end" as in RL

Training requires careful shaping, adapting the task to what animals do innately

The brain comes mostly prepackaged

The brain comes mostly prepackaged

Summary so far

Artificial intelligence relies on "learning"

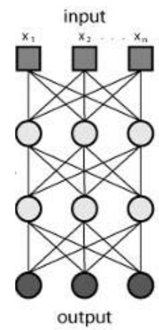

Natural intelligence relies on innate strategies

Where do innate behaviors come from?

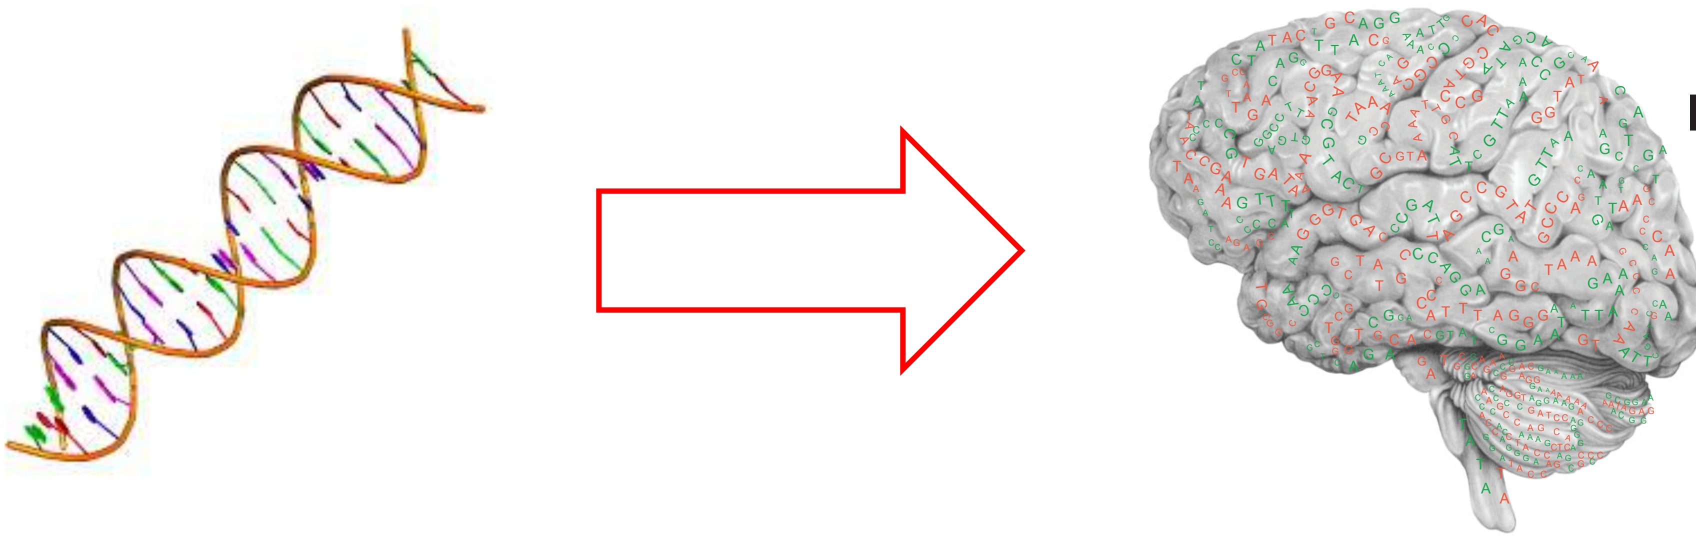

<u>"Innate" behaviors are encoded in the genome</u>

Genome encodes (almost) everything inherited by individual

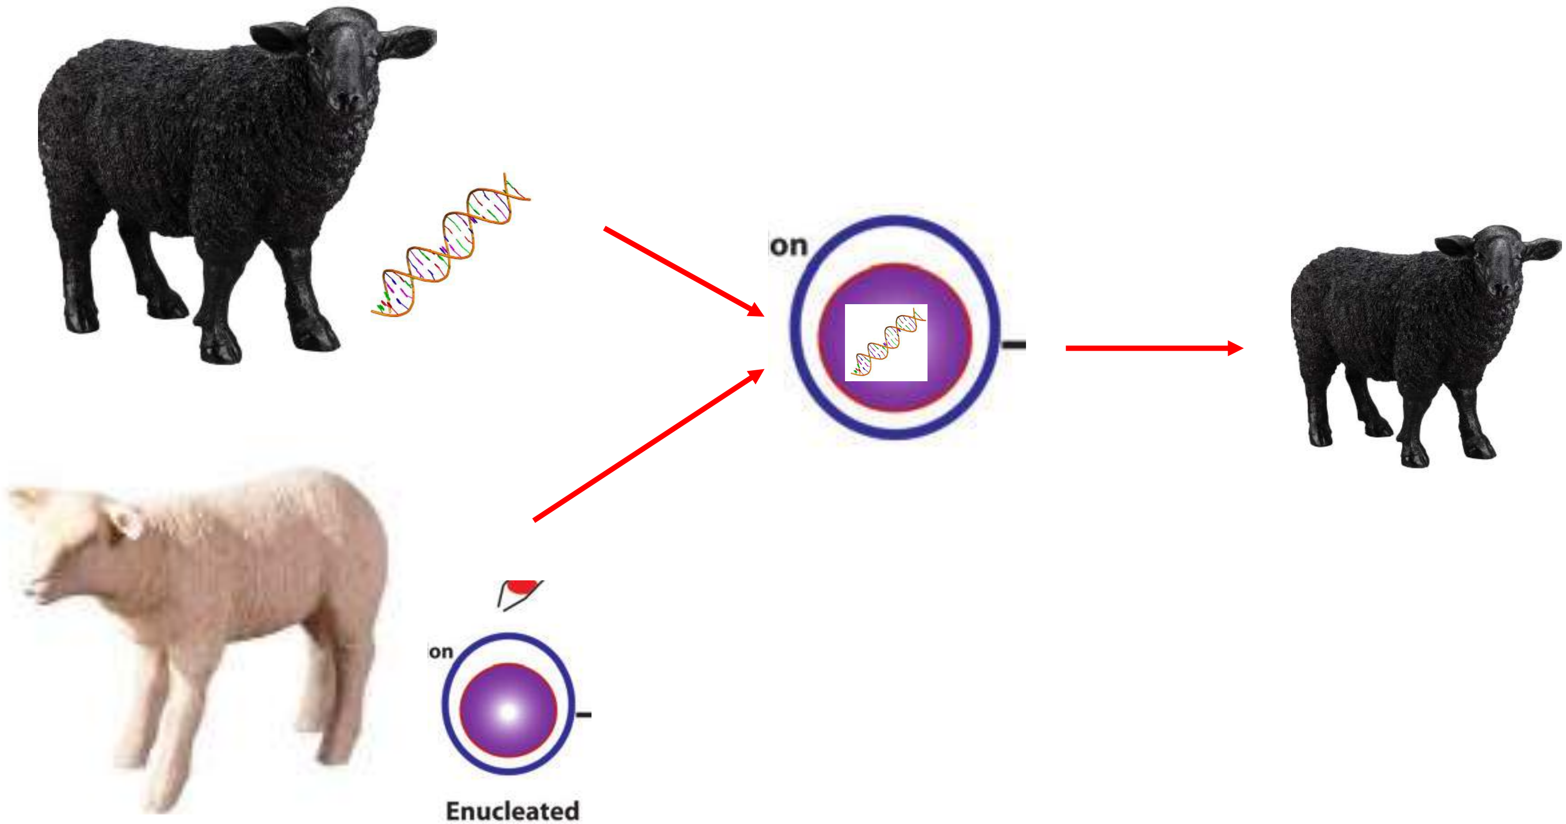

Dolly the sheep

What does the genome encode?

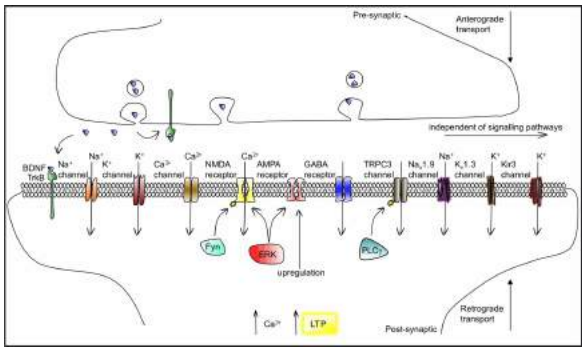

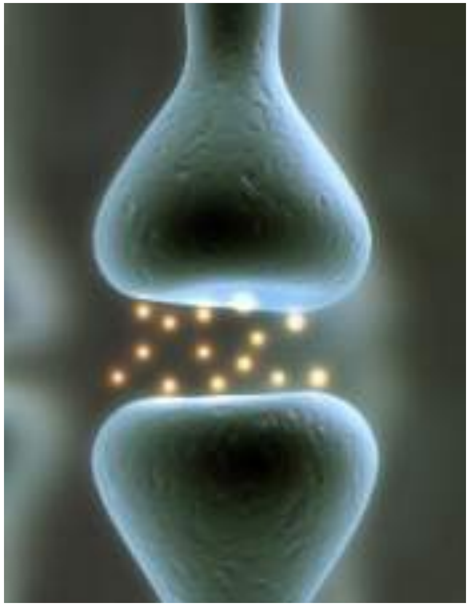

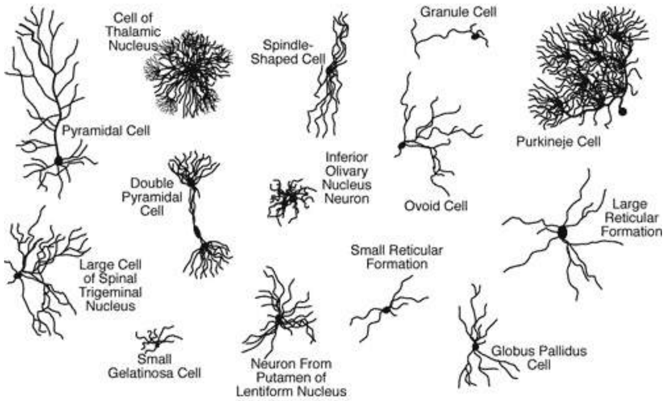

Molecules Synapses Neurons

We know a lot about the neuronal parts list

Genome encodes the initial wiring diagram

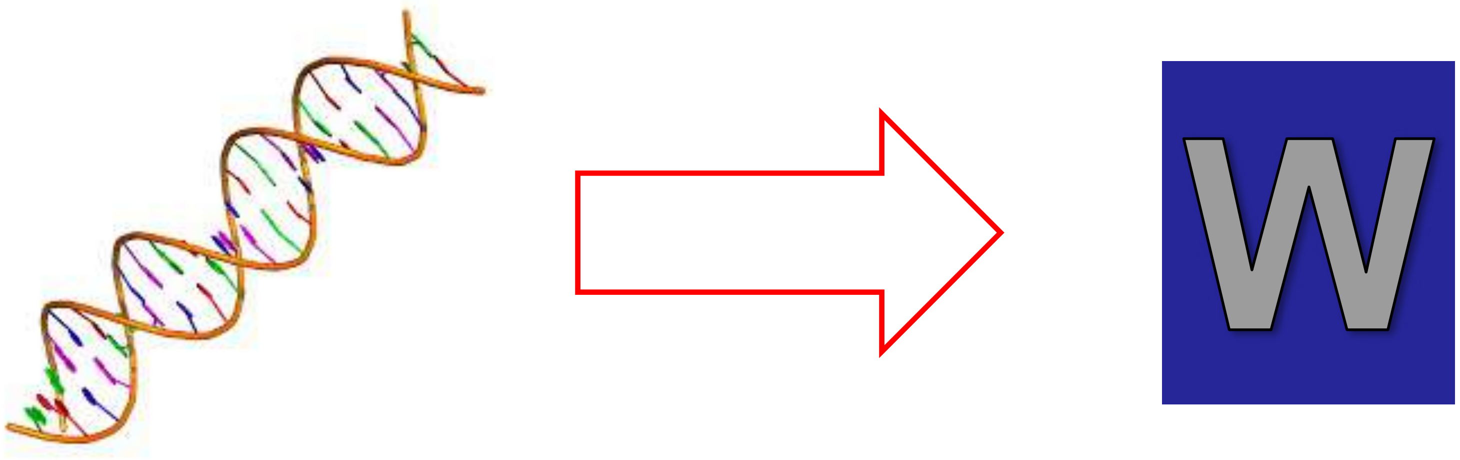

How many parameters does it take to specify a simple brain?

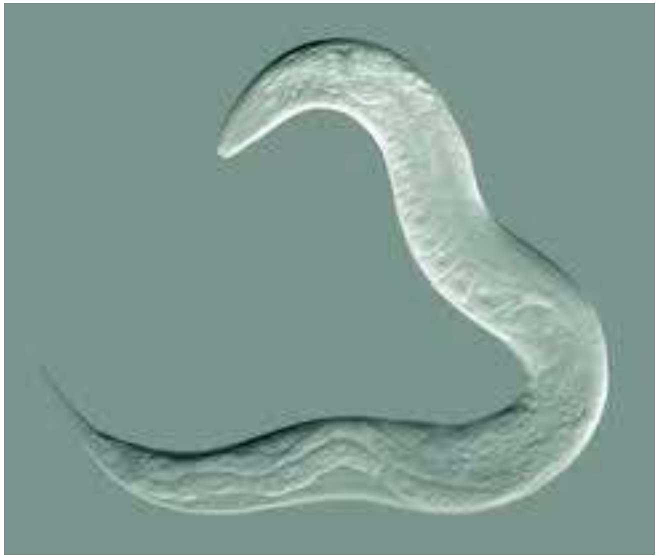

C elegans: 302 neurons, 7000 synapses

Is that a lot or a little?

How many parameters does it take to specify a simple brain?

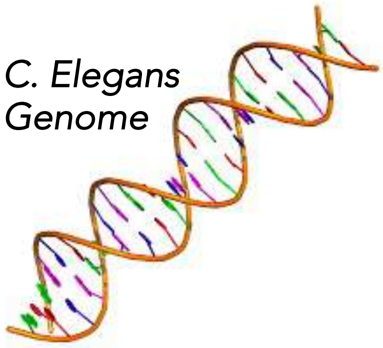

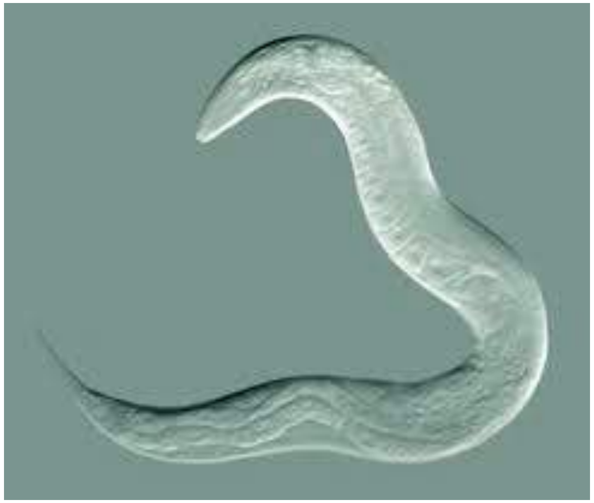

~100M bits

7000 synapses

Every connection could (in principle) be explicitly specified in the genome

How many parameters does it take to specify a complex brain?

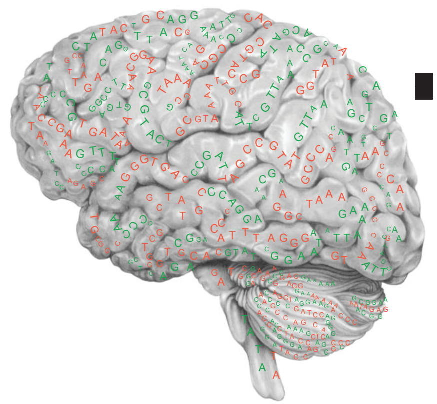

Human: $\sim 10^{11}$ neurons, $\sim 10^{14}$ synapses

<u>Naïve calculation:</u>

Potentially $(10^{14})^2$ connections, so 10^{28} parameters (or bits)

How many parameters does it take to specify a complex brain?

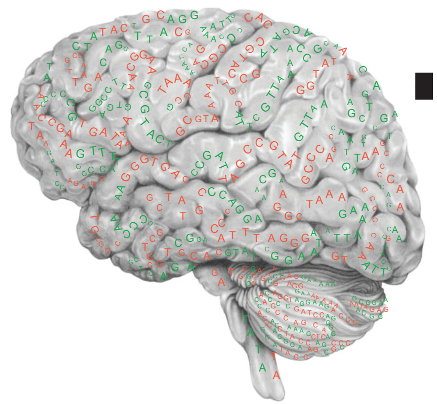

Human: ~10^{11} neurons, ~10^{14} synapses

<u>*Less naïve calculation, exploiting sparsity:*</u>

➔ *~10^{15} bits/brain*

How many parameters does it take to specify a complex brain?

Human Genome

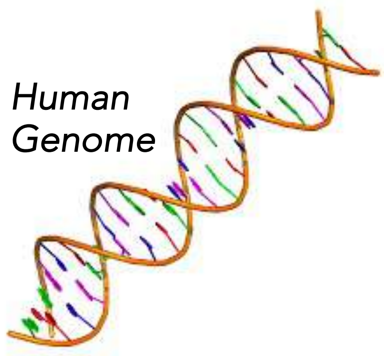

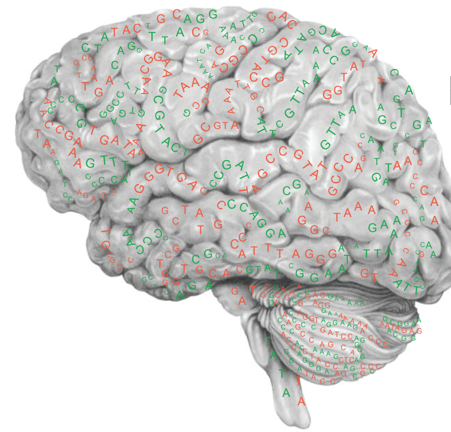

~10^9 bits

~10^{15} bits

1 GB ← ?? missing factor of 10^6 ?? → 1 PB

"Genomic bottleneck"

1 GB ← ?? missing factor of 10^6 ?? → 1 PB

MYSTERY???

The genome specifies rules for wiring a brain

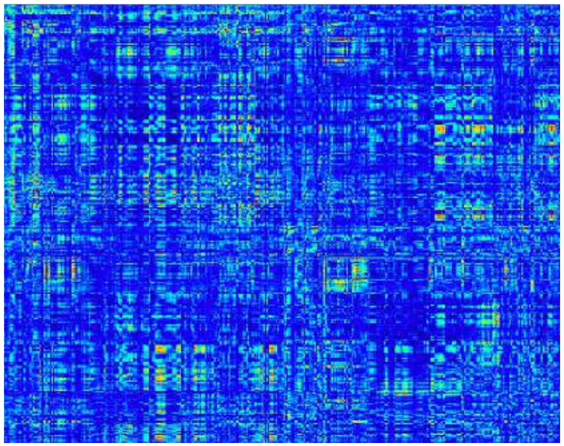

<u>Naïve</u>

302^2 parameters
(arbitrary connection matrix)

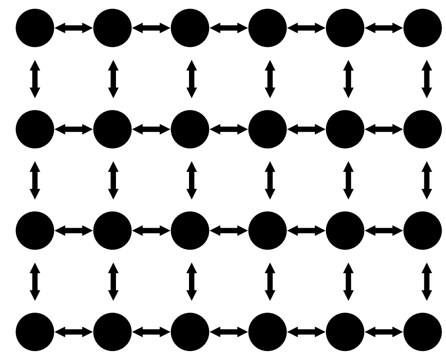

<u>Structured/sparse</u>

(many fewer parameters)

The actual complexity of the wiring diagram of the brain (at birth) is much less than the apparent complexity

Summary Part 2

Artificial intelligence relies on “learning” from examples

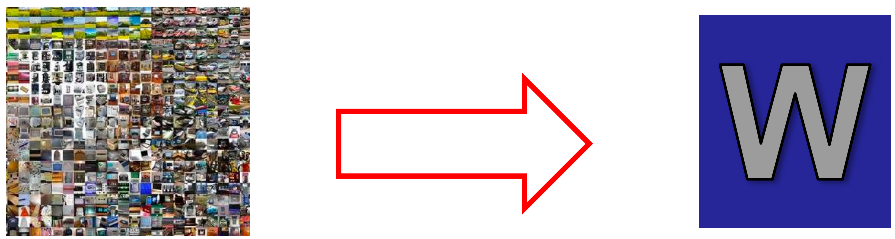

Natural intelligence relies on genomic compression of W

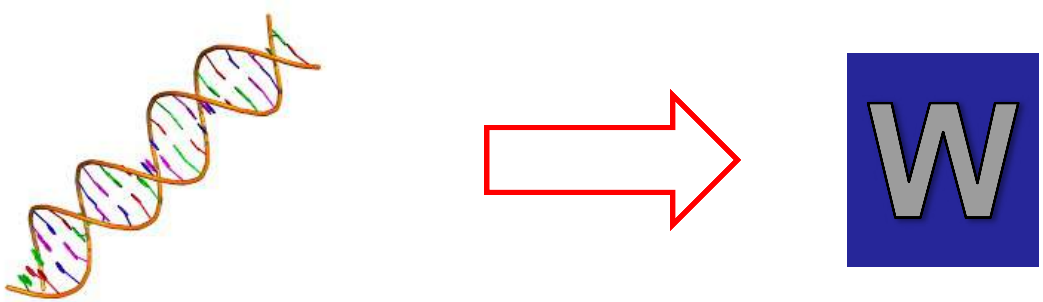

Outline

AI relies mainly on *learning*; animals rely on *innate structure*

→ *Animal learning is no better than machine learning.*

Innate structure: Genome → wiring diagram

Wiring diagram is compressed through "genomic bottleneck."

→ *Genomic bottleneck algorithm*

Genomic bottleneck algorithm

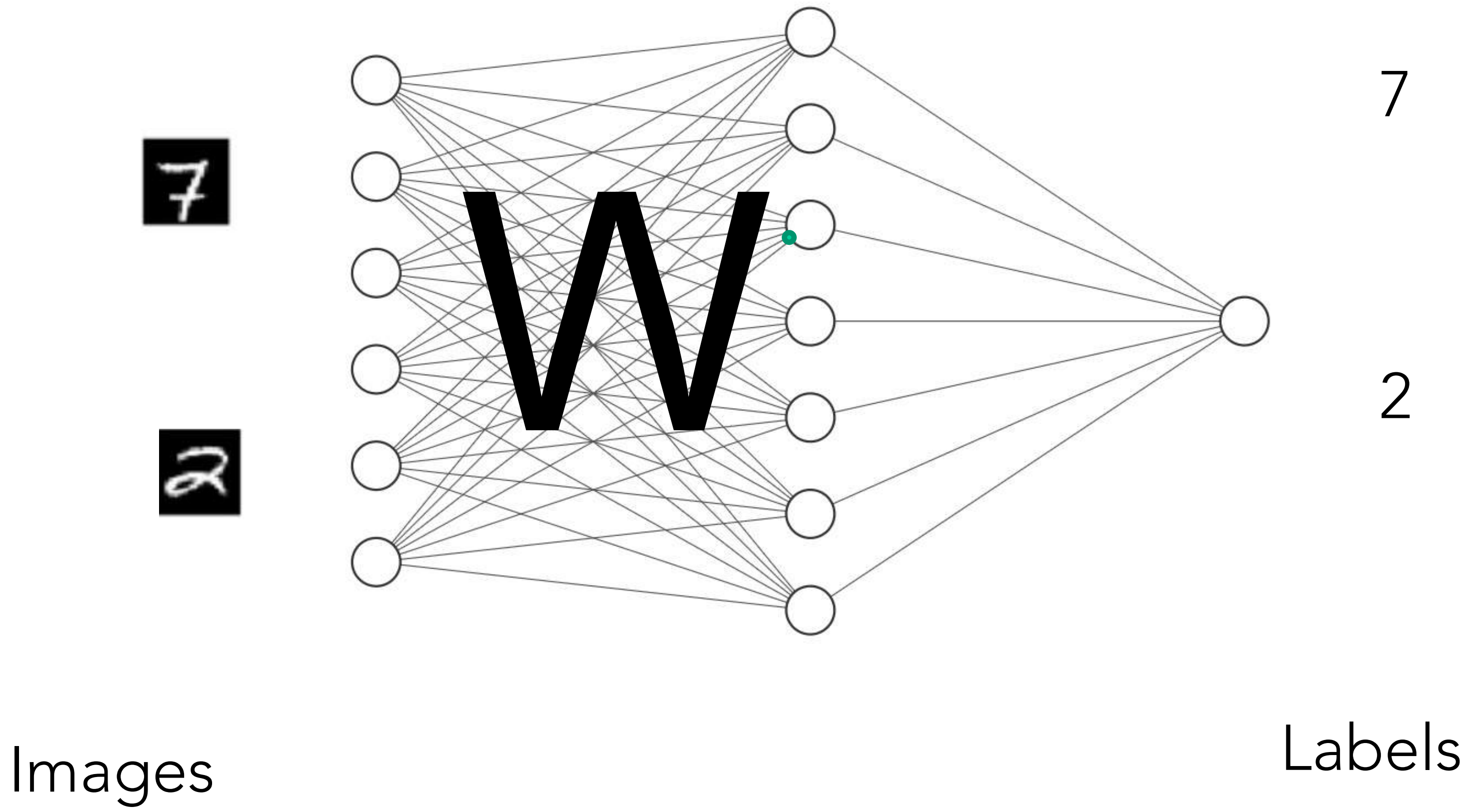

All the information that we learn is imbedded in the weight matrix ***W***

Genomic bottleneck algorithm

Our goal is to compress the weight matrix W
while still maintaining performance

Compression will be performed via a "genome"

Genomic bottleneck algorithm

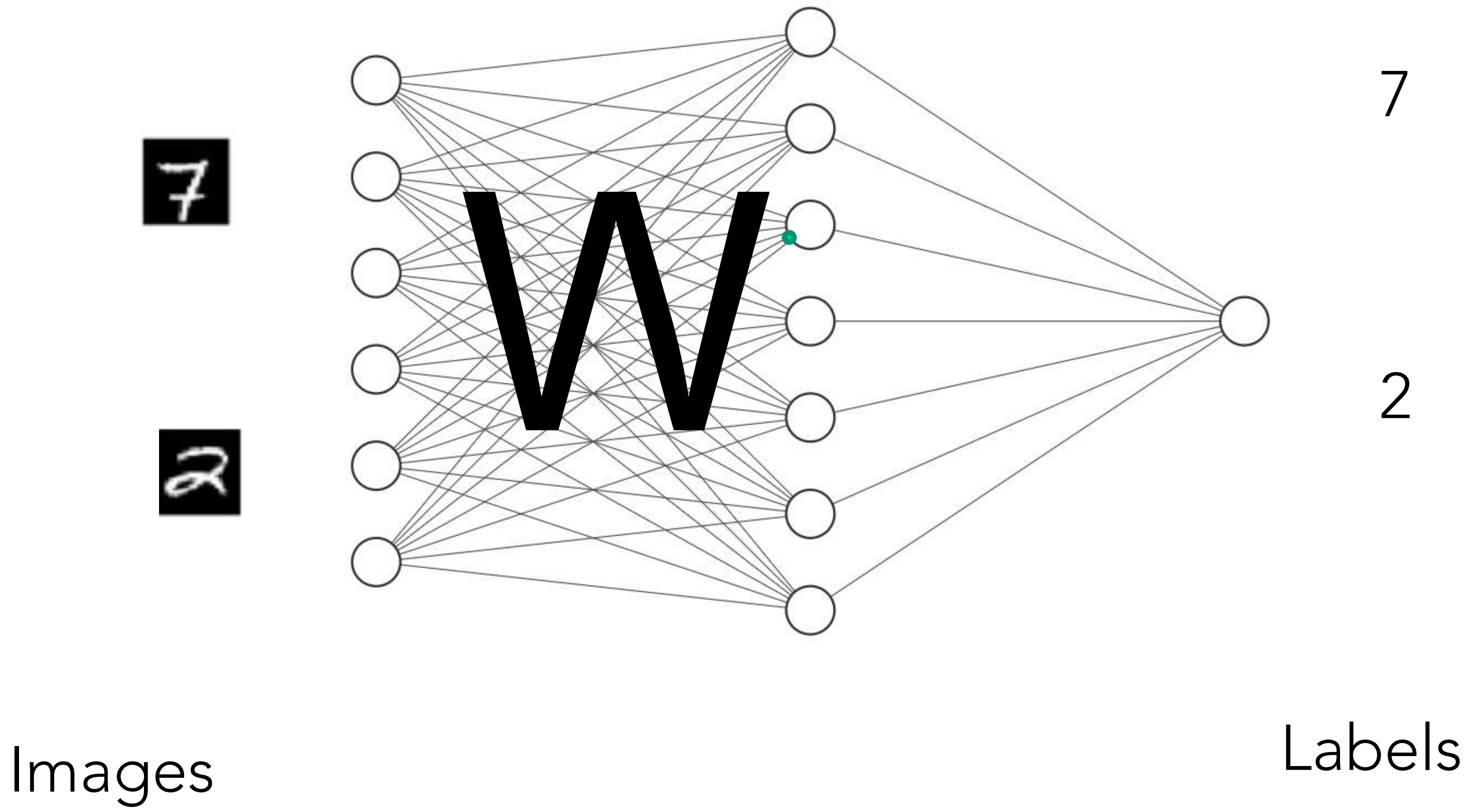

All the information that we learn is imbedded in the weight matrix ***W***

Genomic bottleneck algorithm

$$W \xrightarrow{\text{Compress}} G \xrightarrow[\textit{Development}]{\text{Uncompress}} W^*$$

Weight matrix
that solves a problem

Compressed
Representation
"Genome"

Approximate weight
matrix

Joint work with Alex Koulakov

Genomic bottleneck algorithm - implementation

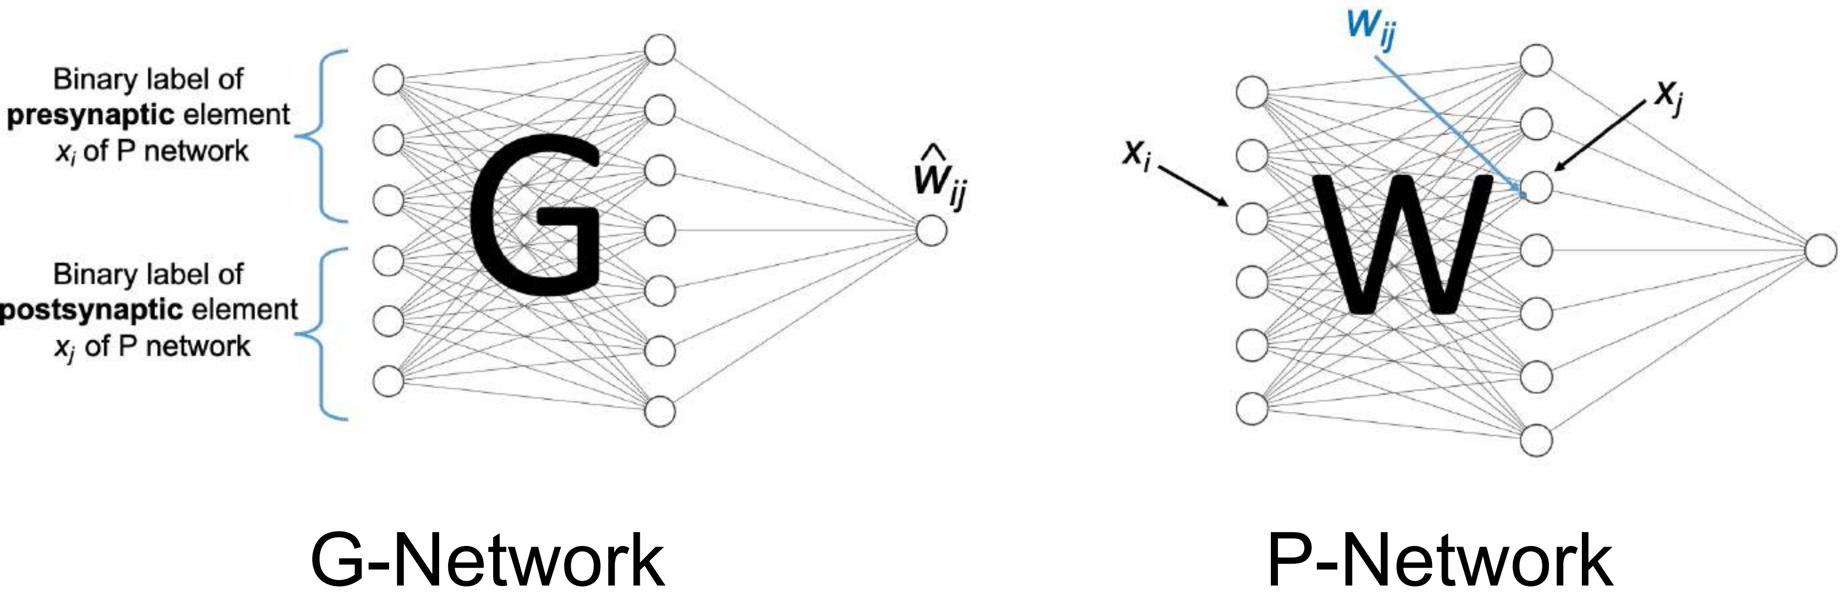

Compression occurs via the "genomic network **G**"

The input to **G** is a pair of neurons

Each neuron x_i is identified by unique binary strings. E.g. x_i = { 0 1 0 1 0 0 1}

Thus if there are *N* input neurons *1..N*, x_i is represented by a *log N* bit string

Joint work with Alex Koulakov

Genomic bottleneck algorithm - implementation

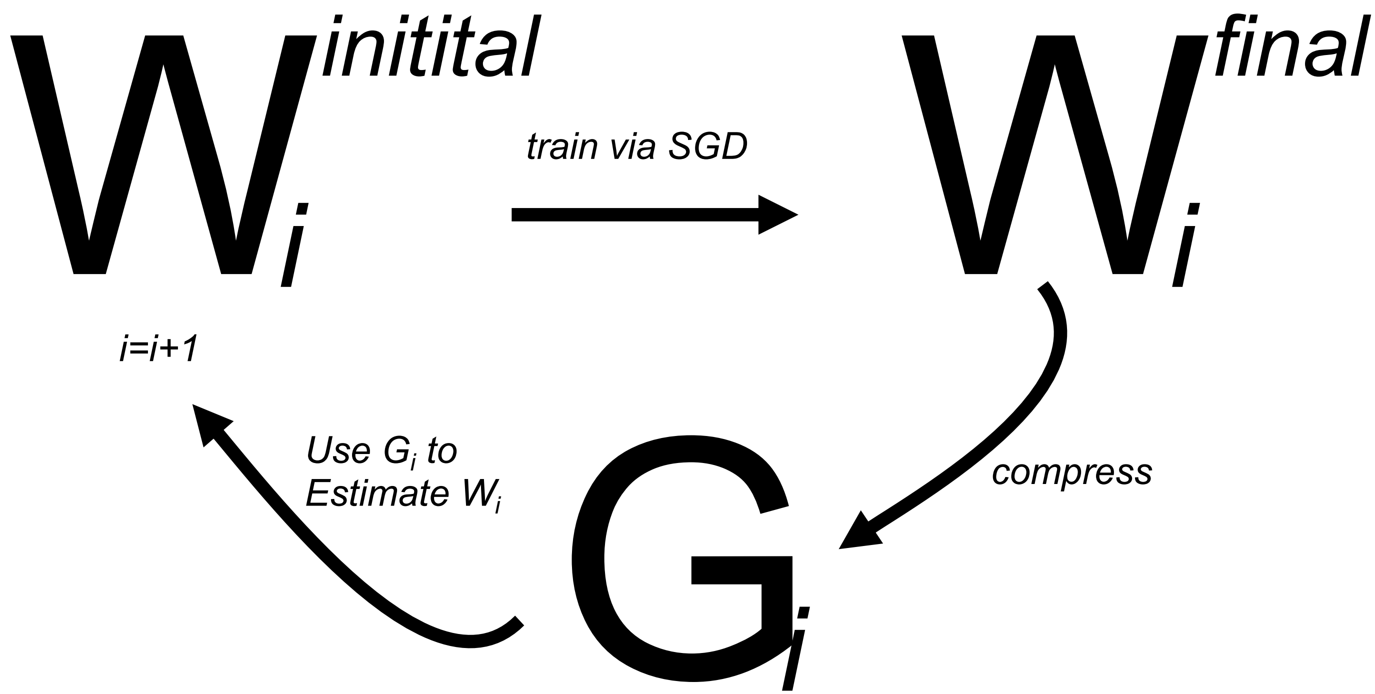

Joint work with Alex Koulakov

Genomic bottleneck compresses MNIST

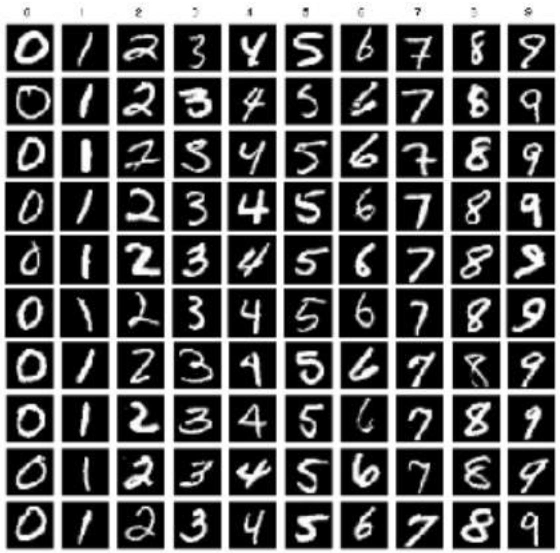

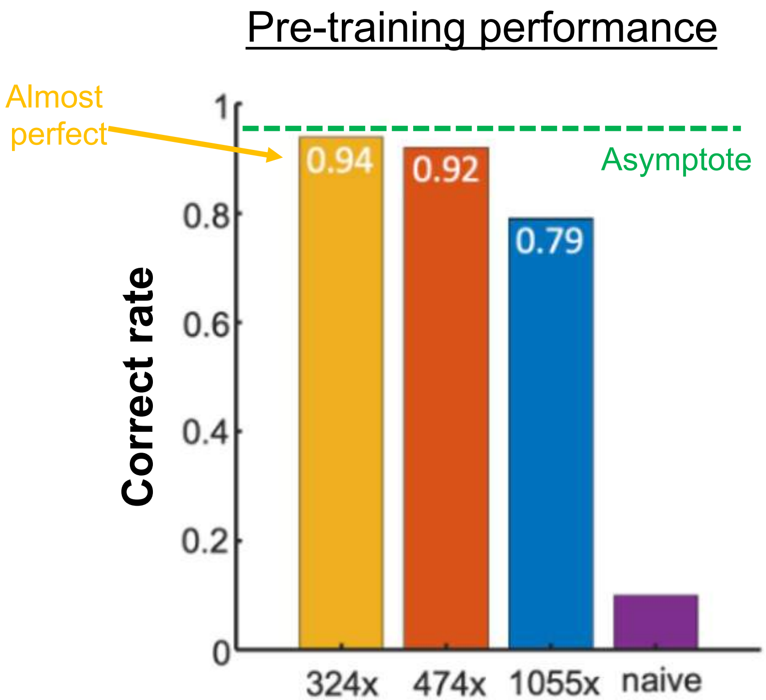

Joint work with Alex Koulakov

Genomic bottleneck compresses CIFAR10

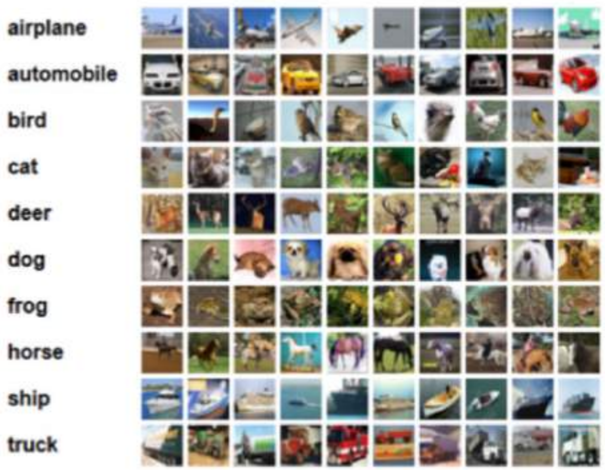

Pre-training performance

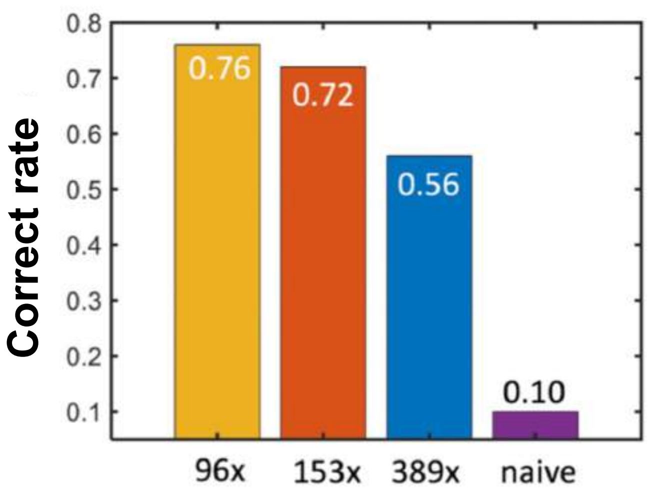

Joint work with Alex Koulakov

Compressed CIFAR10 transfers to SVHN

SVHN

Transfer from CIFAR10 to SVHN

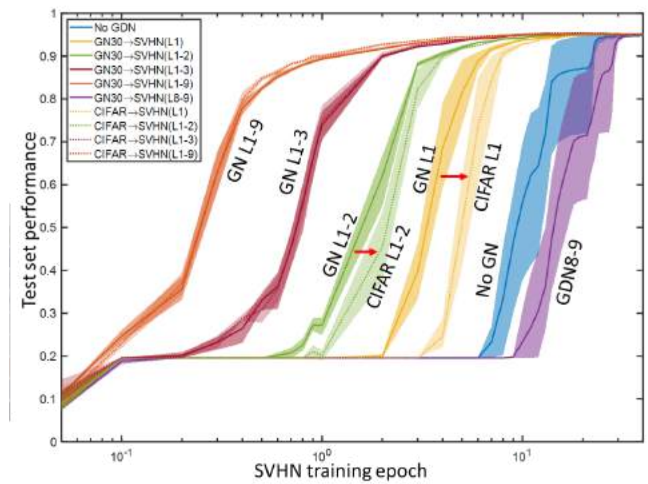

Joint work with Alex Koulakov

Dataset - Beam Rider

- The player's ***objective is to clear the Shield's 99 sectors of alien craft while piloting the Beamrider ship***.
- In this environment, the ***observation is an RGB image of the screen,*** which is an array of shape ***(210, 160, 3)***.
- **Total number of actions - 9**

Genomic bottleneck in Reinforcement Learning

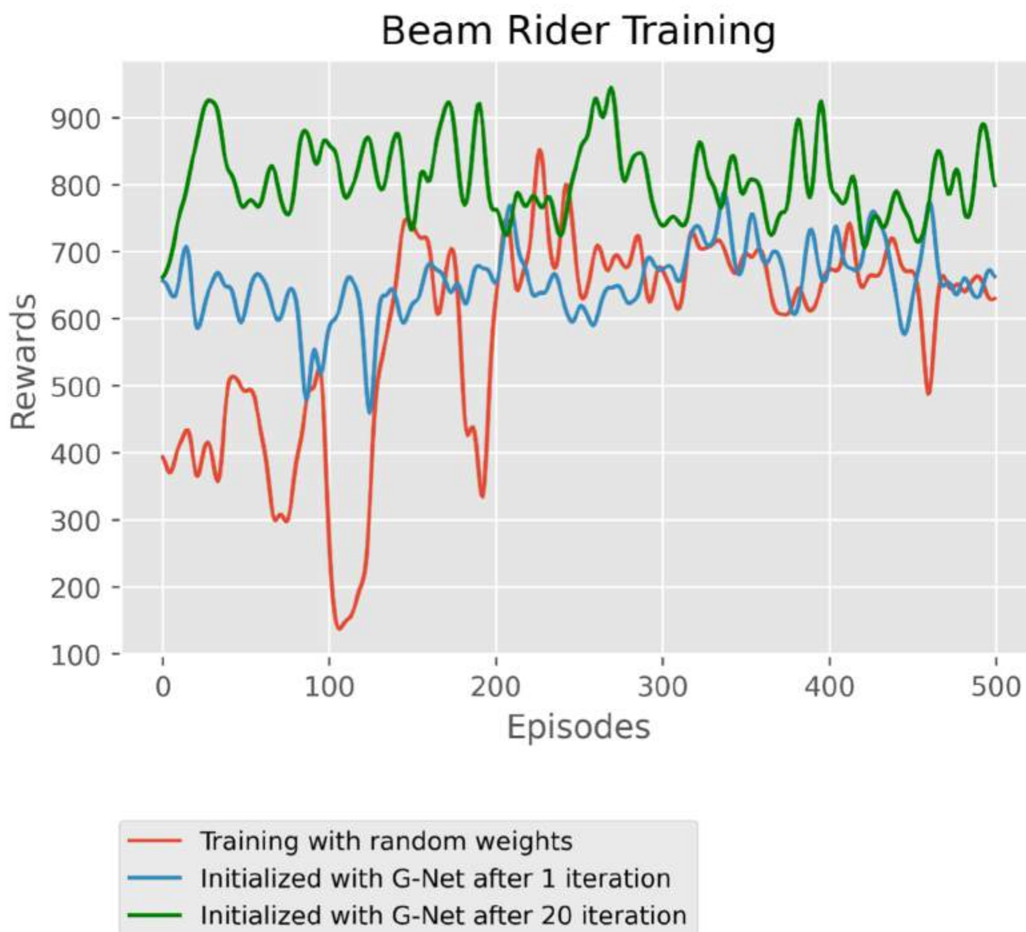

- Number of parameters in the P-net: ***3,295,915***
- Number of parameters in the G-net: ***9,994***
- Compression ratio is ***330x***

Joint work with Alex Koulakov/Divyansha

Genomic bottleneck in Reinforcement Learning

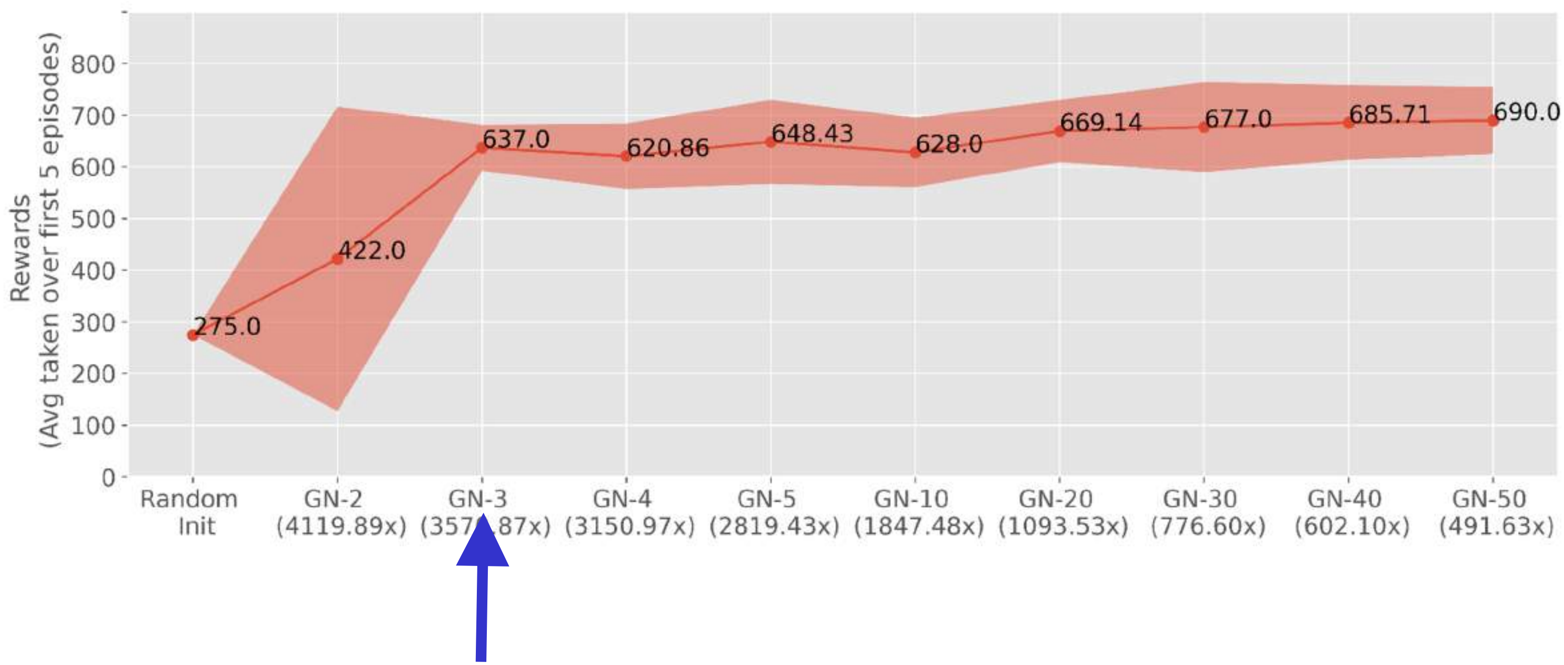

Good performance up to 3600-fold compression

Average Rewards for the first 5 episodes for the P nets generated from different G nets (Results averaged over 8 runs)

Joint work with Alex Koulakov/Divyansha

Next up: Open-AI Gym HalfCheetah

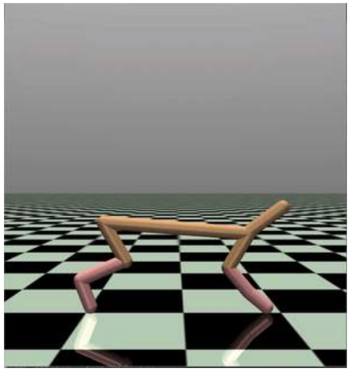

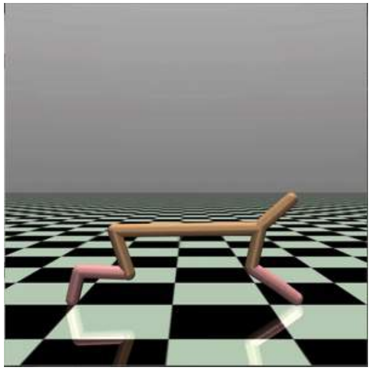

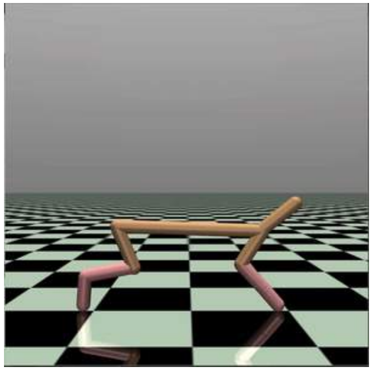

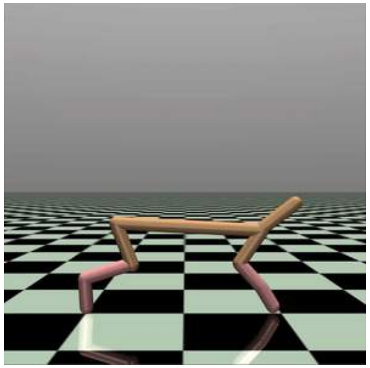

Joint work with Alex Koulakov/Divyansha

Genomic bottleneck algorithm-summary

Can compress MNIST & CIFAR x100-1000-fold

Complex compressed networks show good transfer learning

Promising results with reinforcement learning problems

We Need an "Embodied Turing Test"

Traditional Turing test

Embodied Turing test

Why are humans so successful?

~~Learning!?~~

Although humans *are* very good at learning,
we may not be much better than eg other apes.

I do not believe this is the key to our success

Humans have been around for >100,00 years, and almost died out 70,000 years ago! We only started succeeding very recently. Why?

Human success derives from cultural transmission

Oral transmission

Written transmission

Cultural transmission breaks the genomic bottleneck

Summary

AI relies mainly on *learning;* animals rely on *innate structure*

→ *Animal learning is no better than machine learning.*

Innate structure: Genome → wiring diagram

Wiring diagram is compressed through "genomic bottleneck."

→ *Genomic bottleneck algorithm*

Can compress MNIST & CIFAR x100-1000-fold

Complex compressed networks show good transfer learning

NeuroAI Scholars Program at CSHL

- For people with PhDs in AI; not a traditional postdoc
- Focus on applying insights from neuro to AI
- Scholars are imbedded for 2 years in CSHL neuro labs
- Scholars are independent, expected to forge new paths

https://www.cshl.edu/research/neuroscience/neuroai/

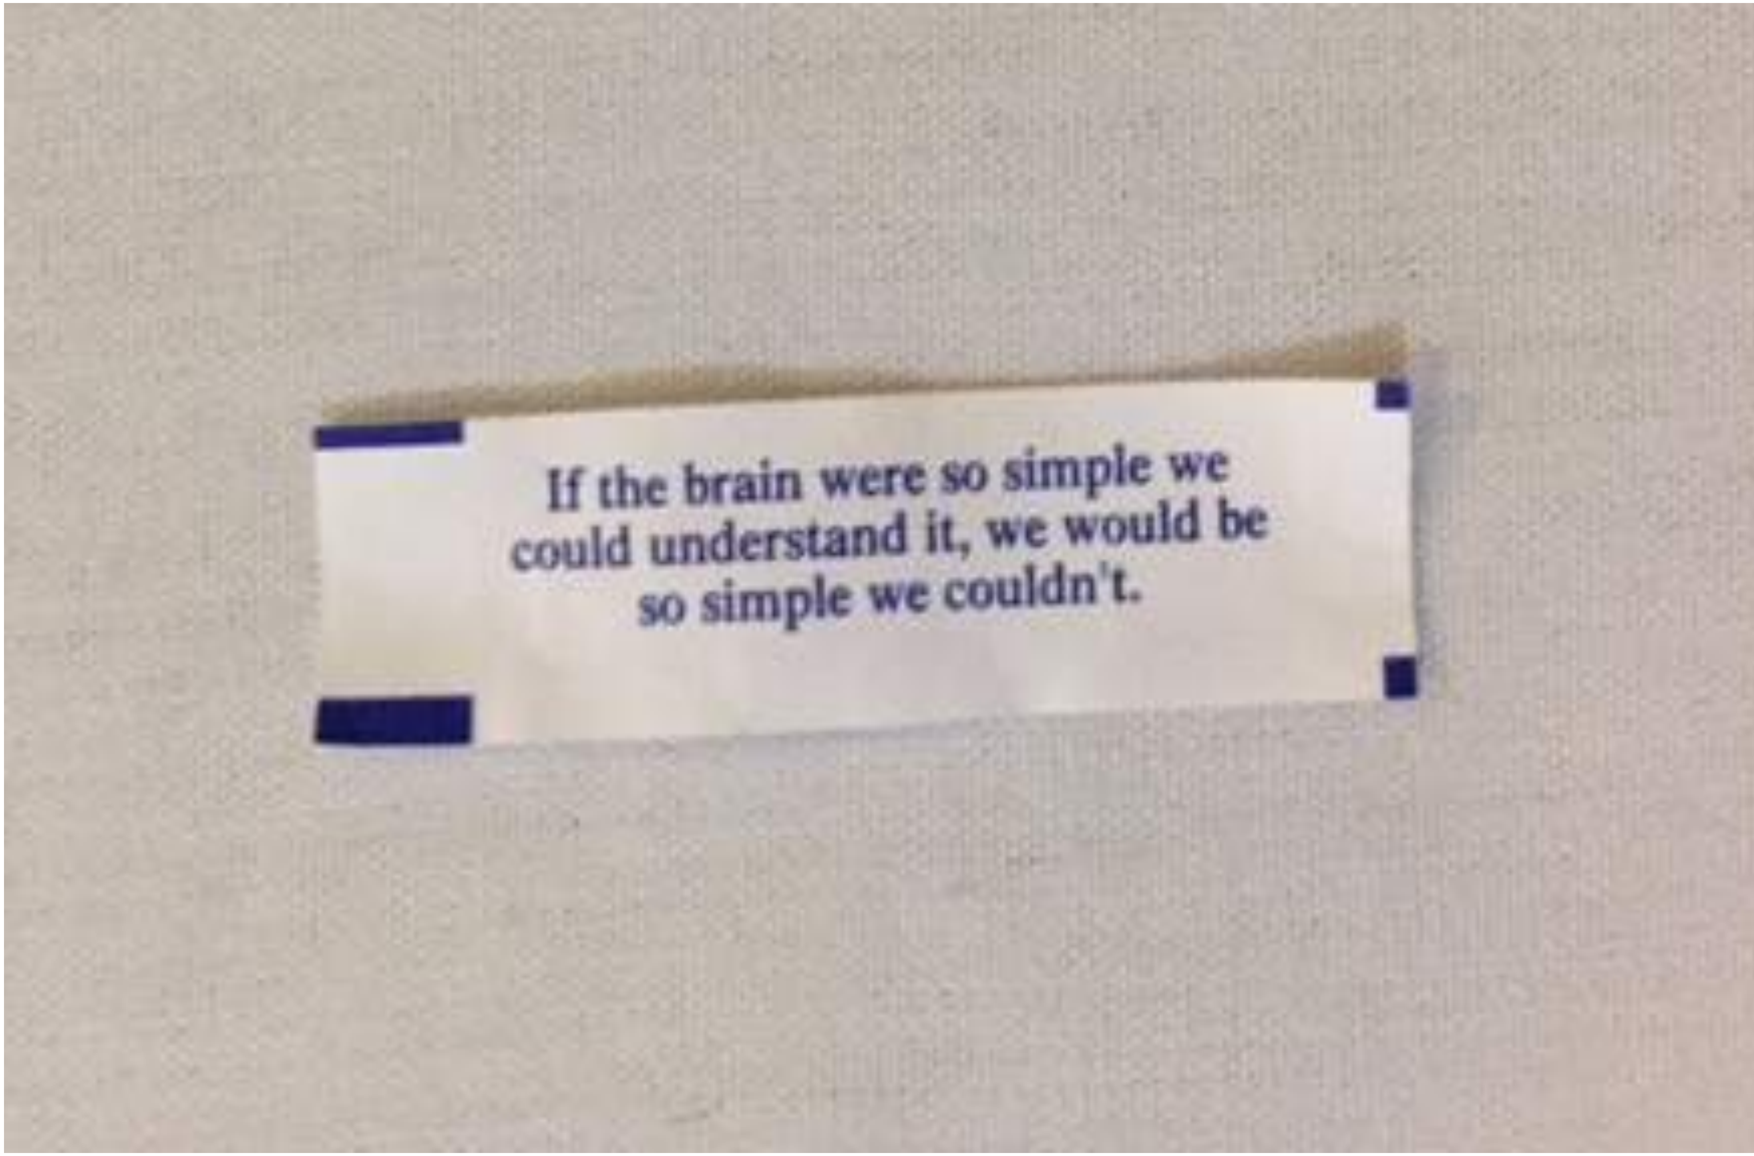
If the brain were so simple we
could understand it, we would be
so simple we couldn't.