

Working towards human-level intelligence...

Dileep George

Vicarious

Intelligence:

Ability to model the world and to act on it

Intelligence:

Ability to model the world and to act on it

It is possible to act on the world without modeling it

Old brain

sponges, jellyfish, flatworm

fish

amphibians

mammaliaforms

non-human primates

modern humans

-600  
million years

-450

-300

-150

0

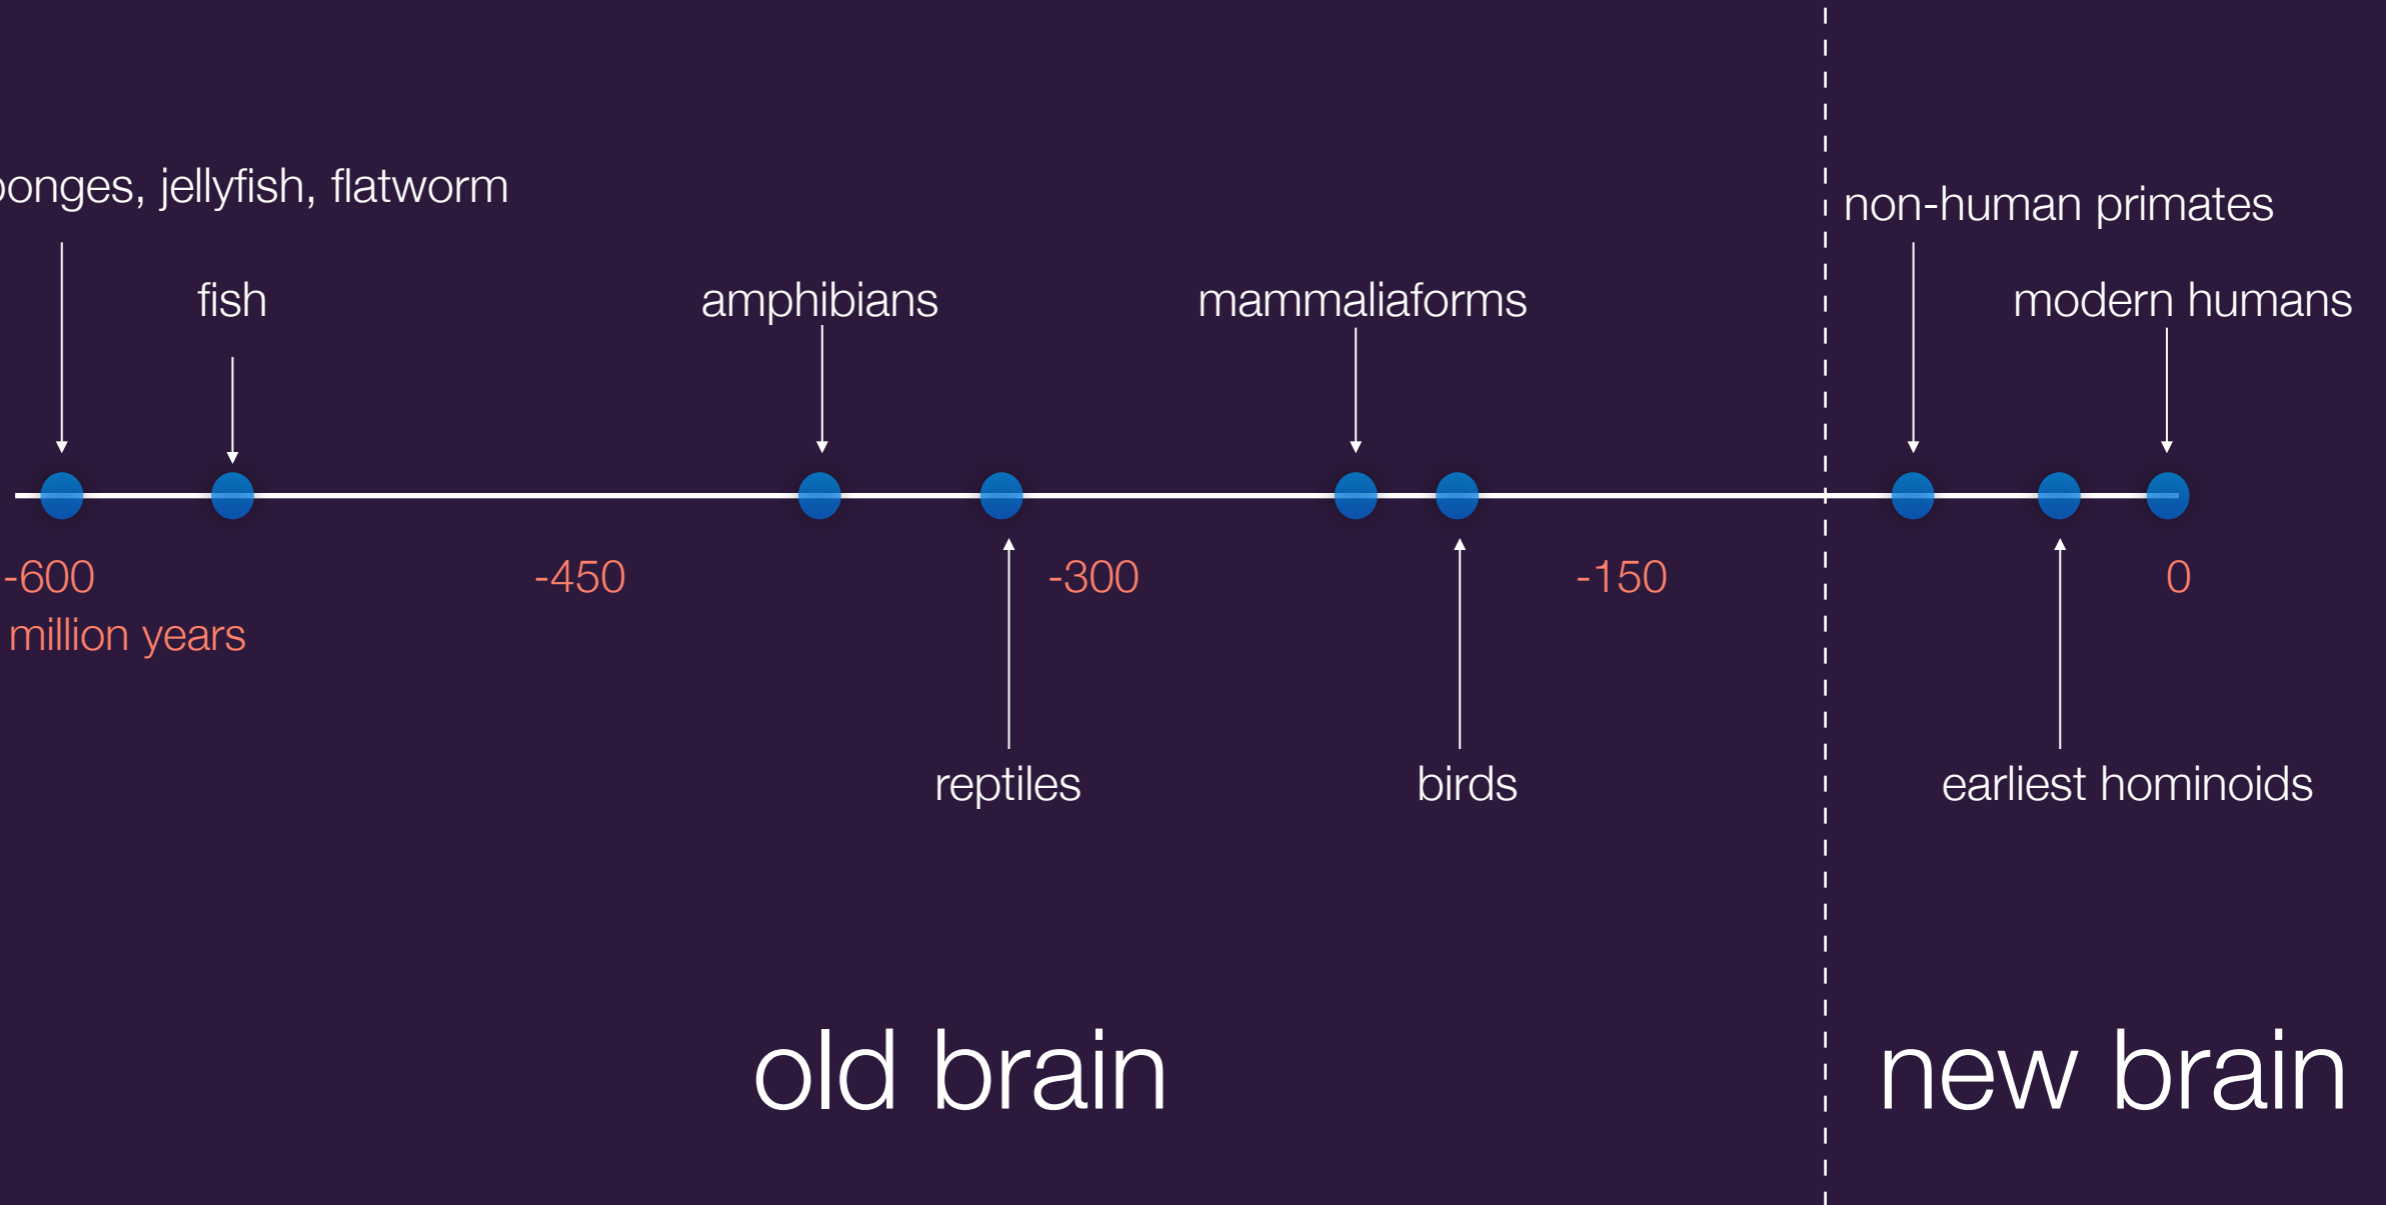
reptiles

birds

earliest hominoids

old brain

new brain



Dinosaurs appeared

Dinosaurs became extinct

Homo sapiens appeared

---

Jan 1st

Sep 21st

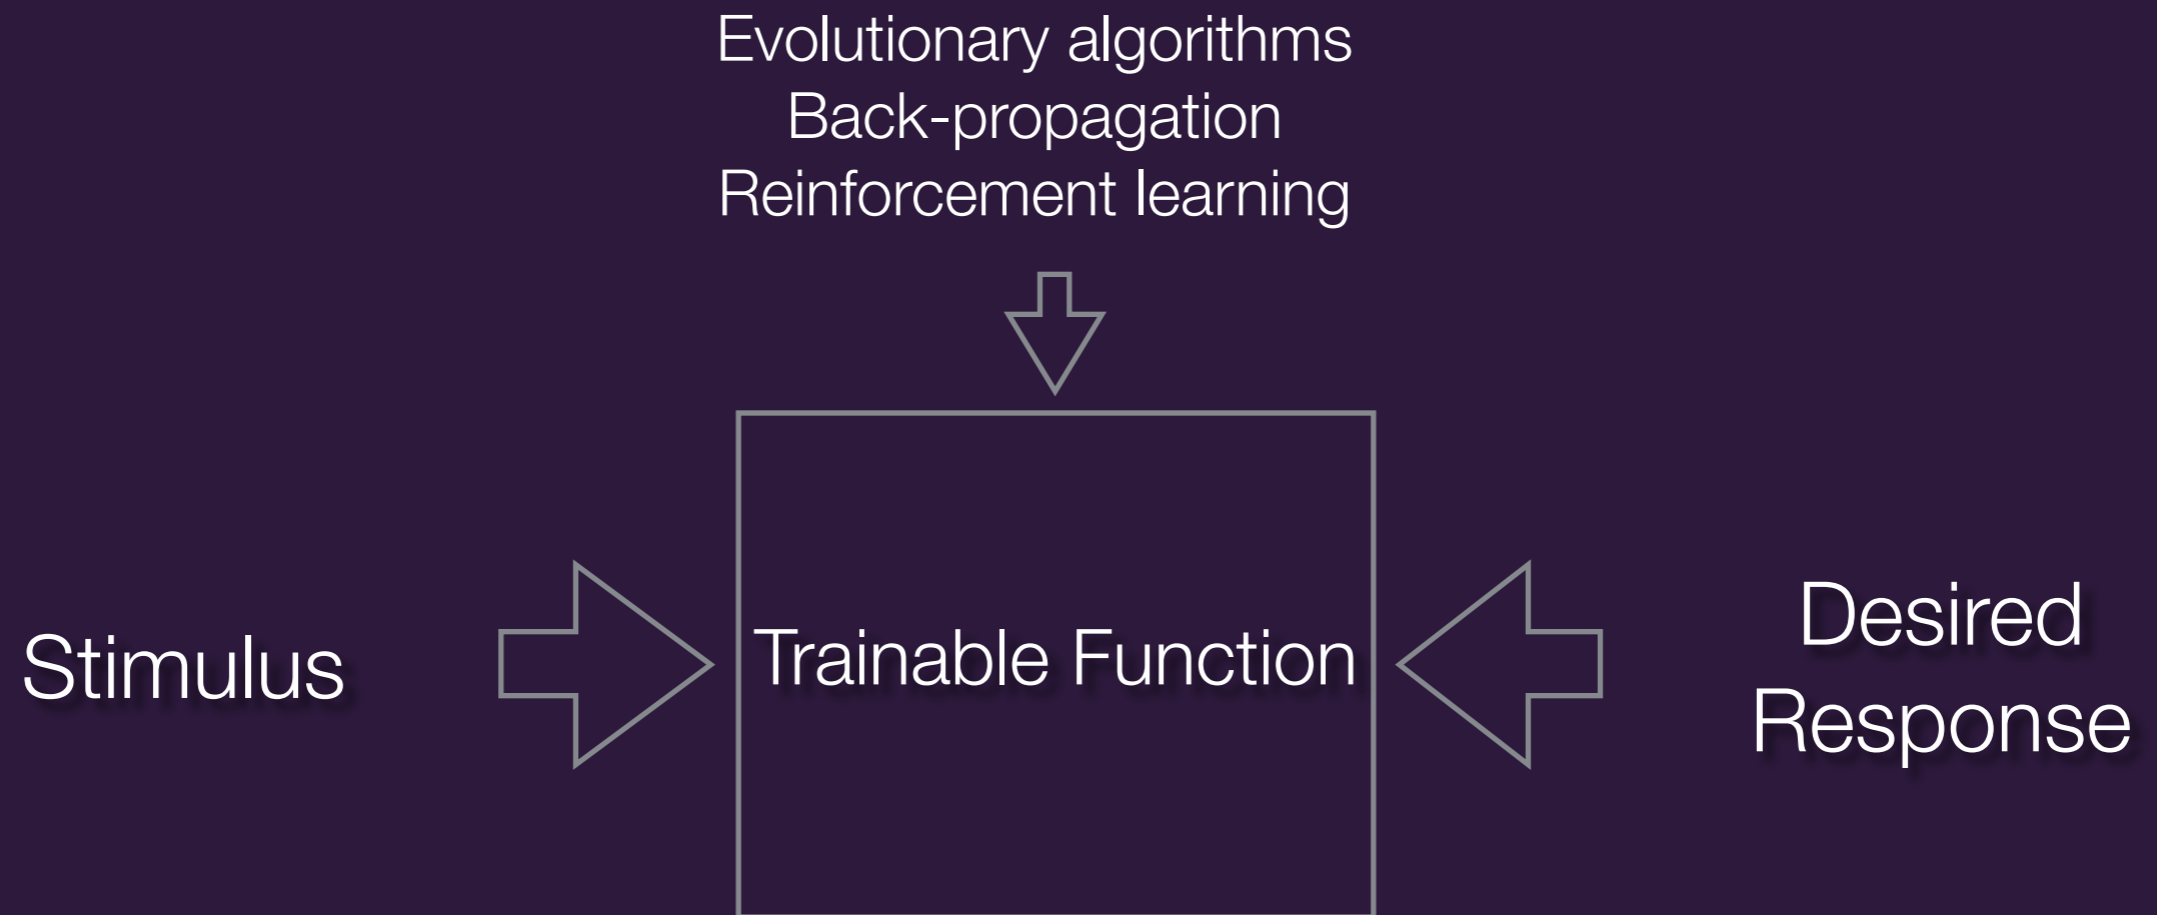
Dec 31st

Old brain was very successful. Dinosaurs had walnut-sized brains

# Old-brain style (model-free) learning

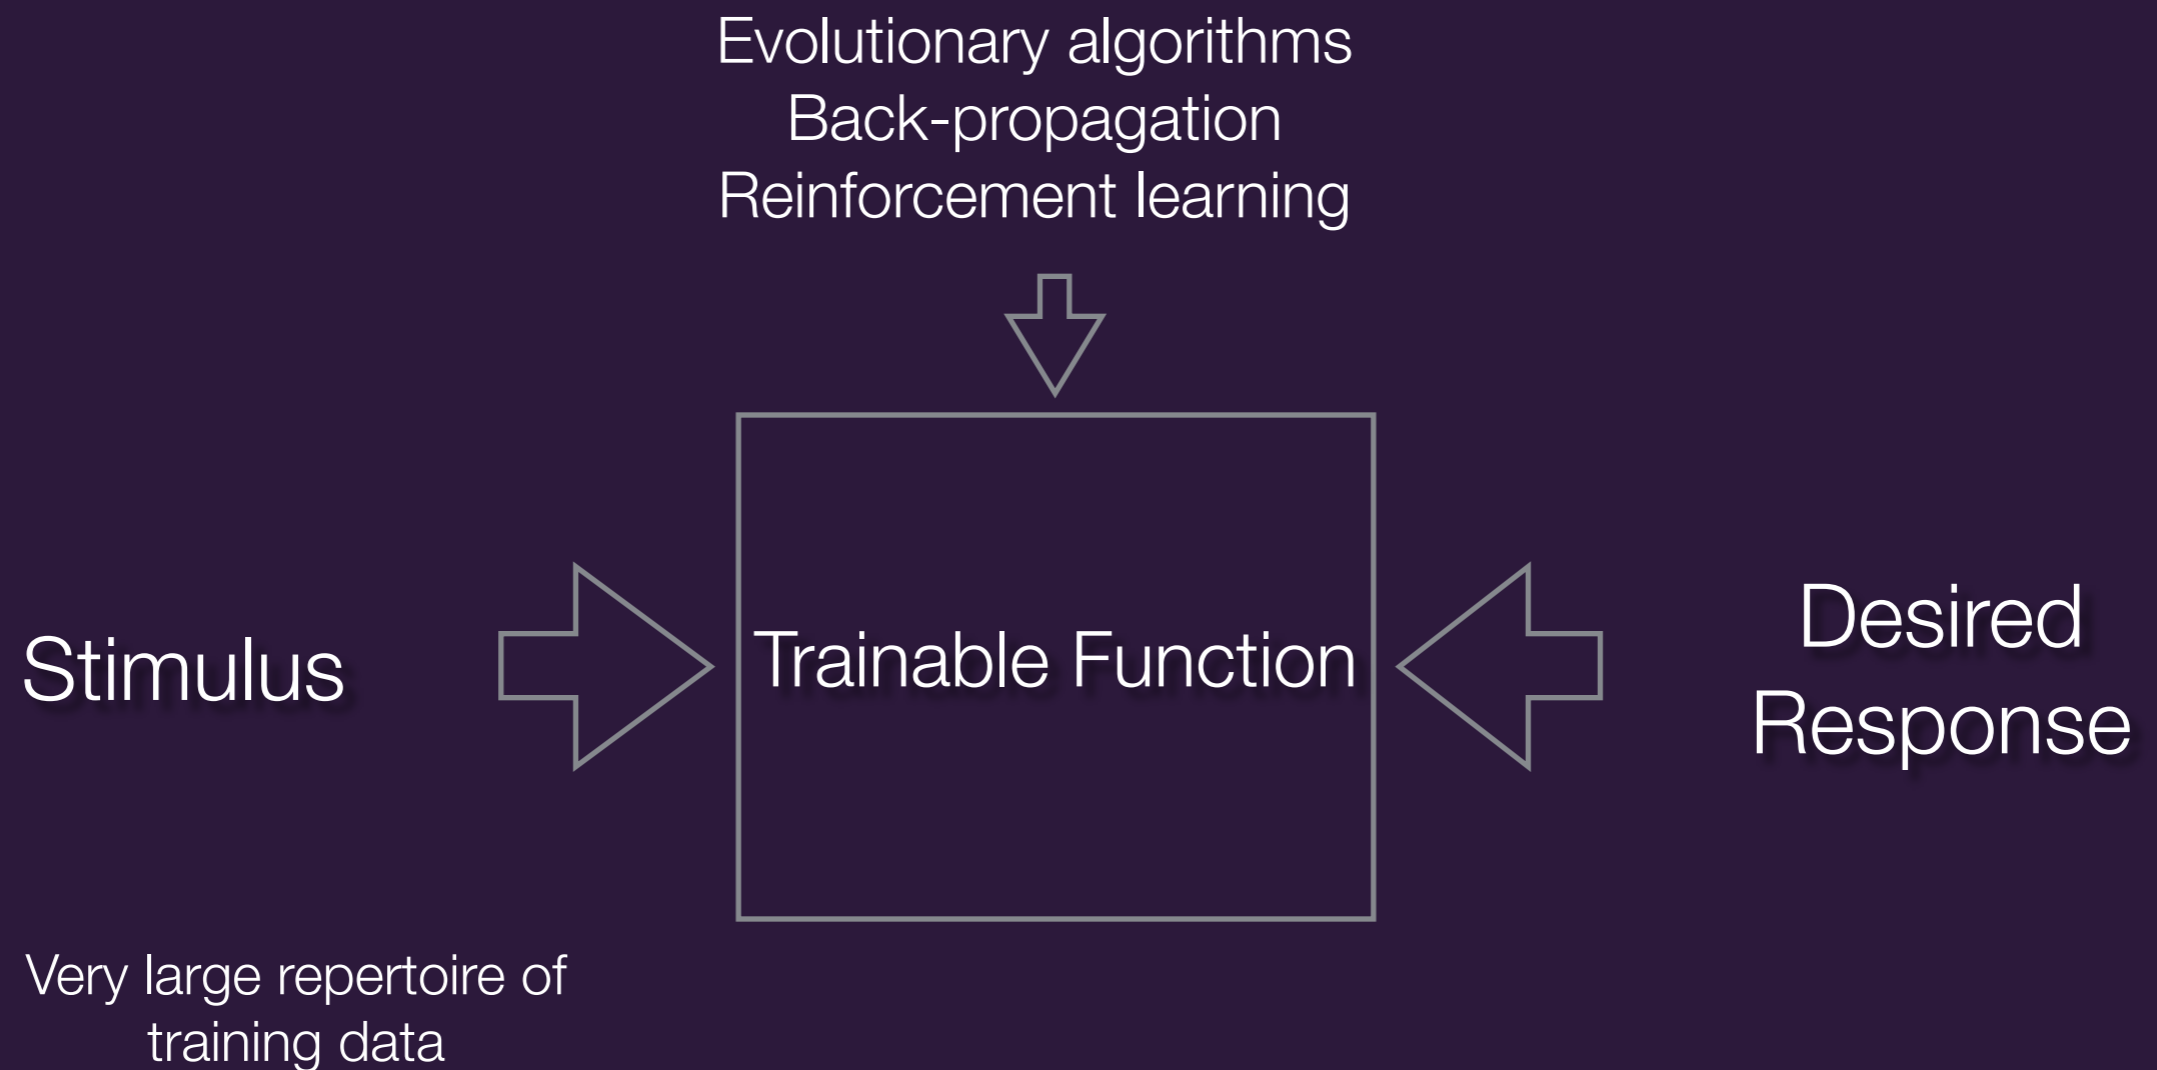


# Old-brain style (model-free) learning





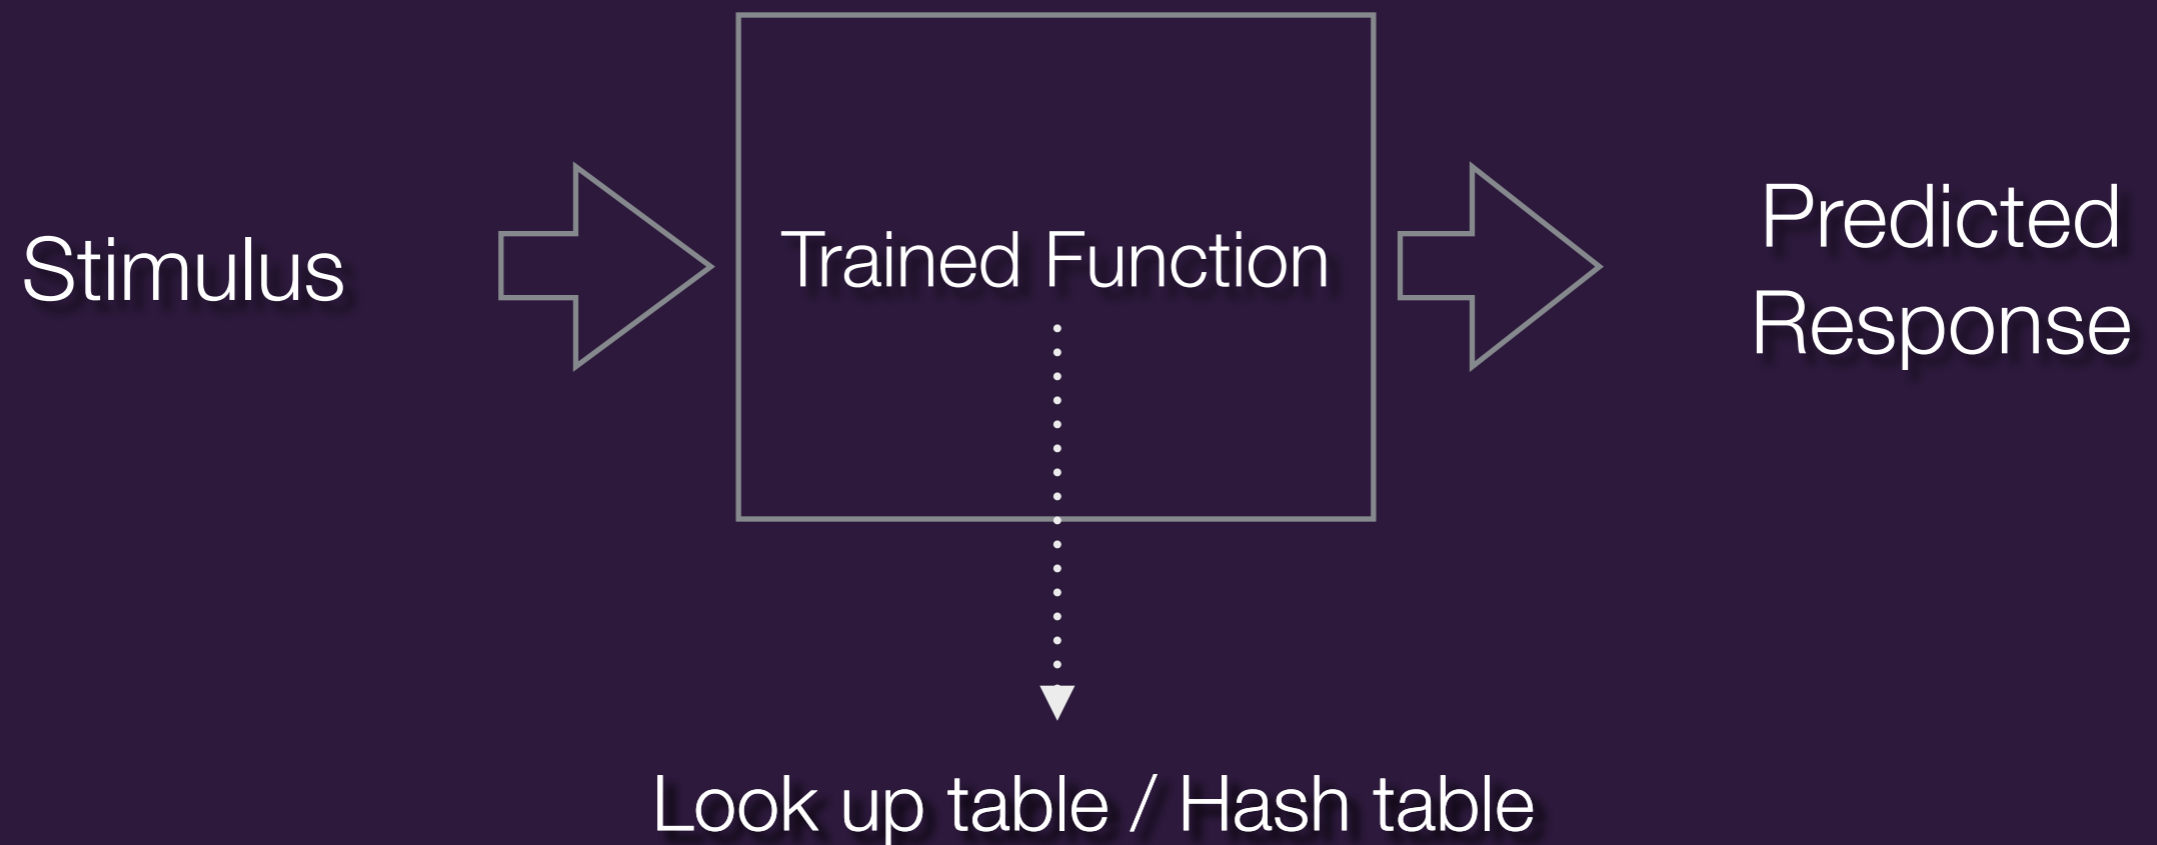
# Old-brain style (model-free) learning



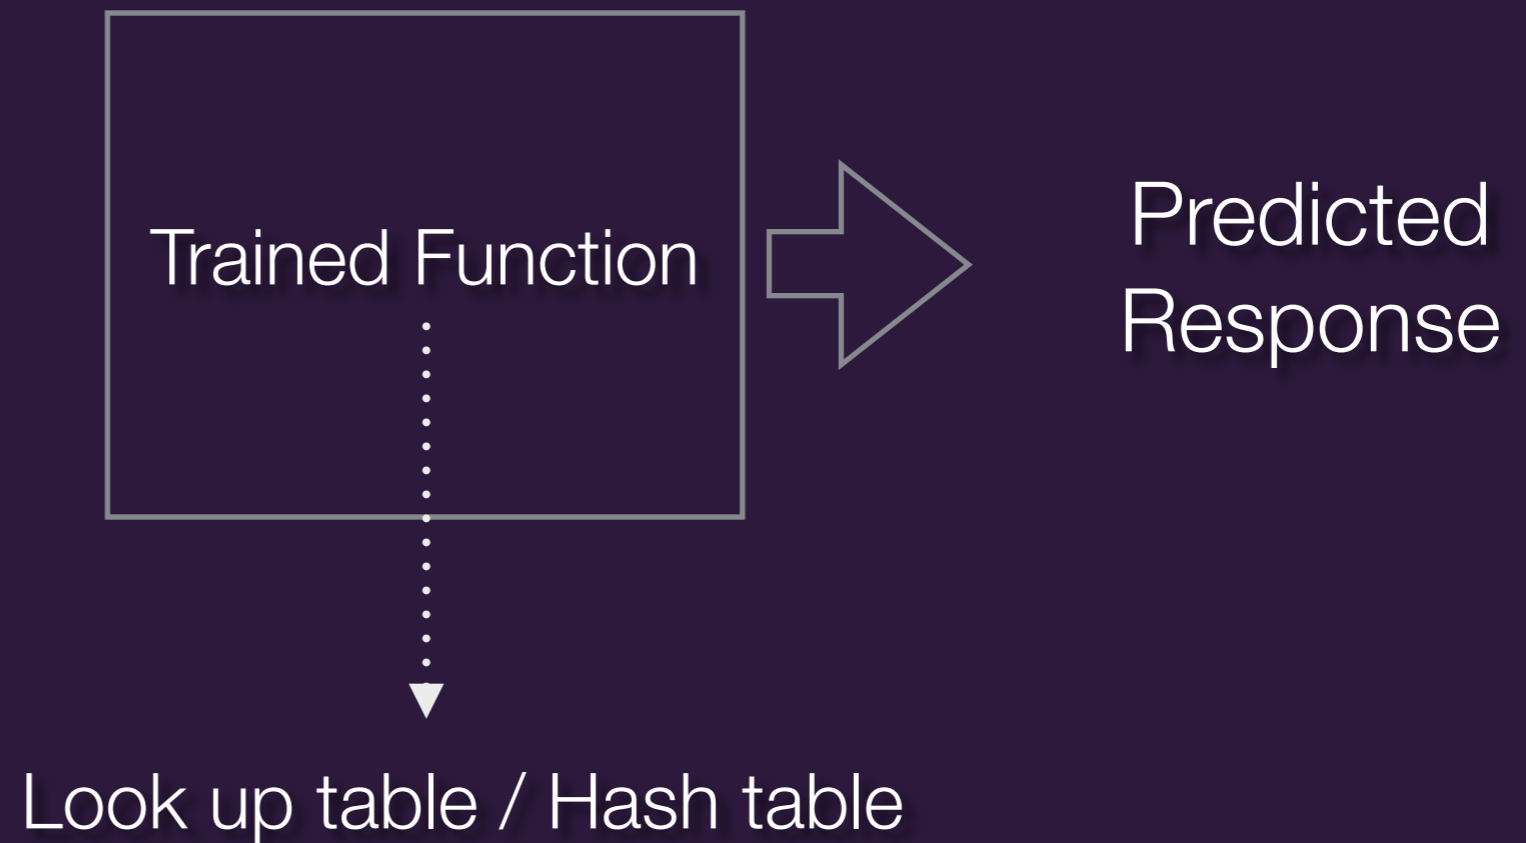
# Old-brain style (model-free) learning



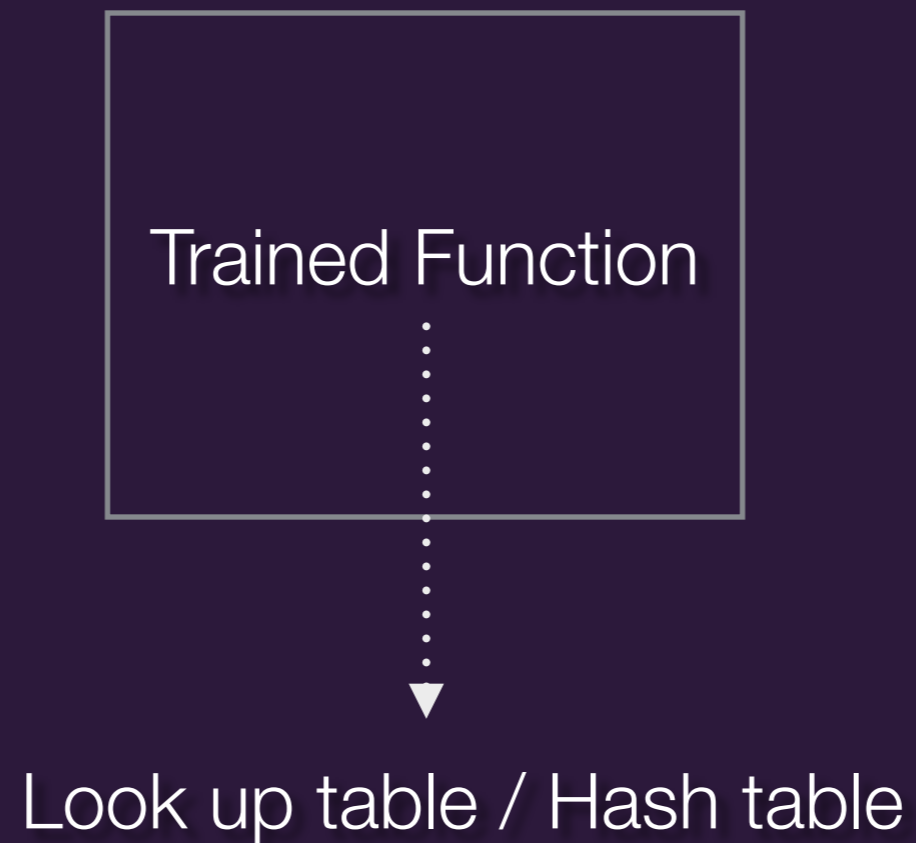
# Old-brain style (model-free) learning



# Old-brain style (model-free) learning



# Old-brain style (model-free) learning



# Old-brain style (model-free) learning



Black box with limited capability for thinking, imagination, creativity or planning

# Many modern methods are still like the old brain

- Extremely large number of training examples
- Inscrutable, black box classifiers
- Not generative. Models lack explanatory power

# Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images

Anh Nguyen  
University of Wyoming  
anguyen8@uwyo.edu

Jason Yosinski  
Cornell University  
yosinski@cs.cornell.edu

Jeff Clune  
University of Wyoming  
jeffclune@uwyo.edu

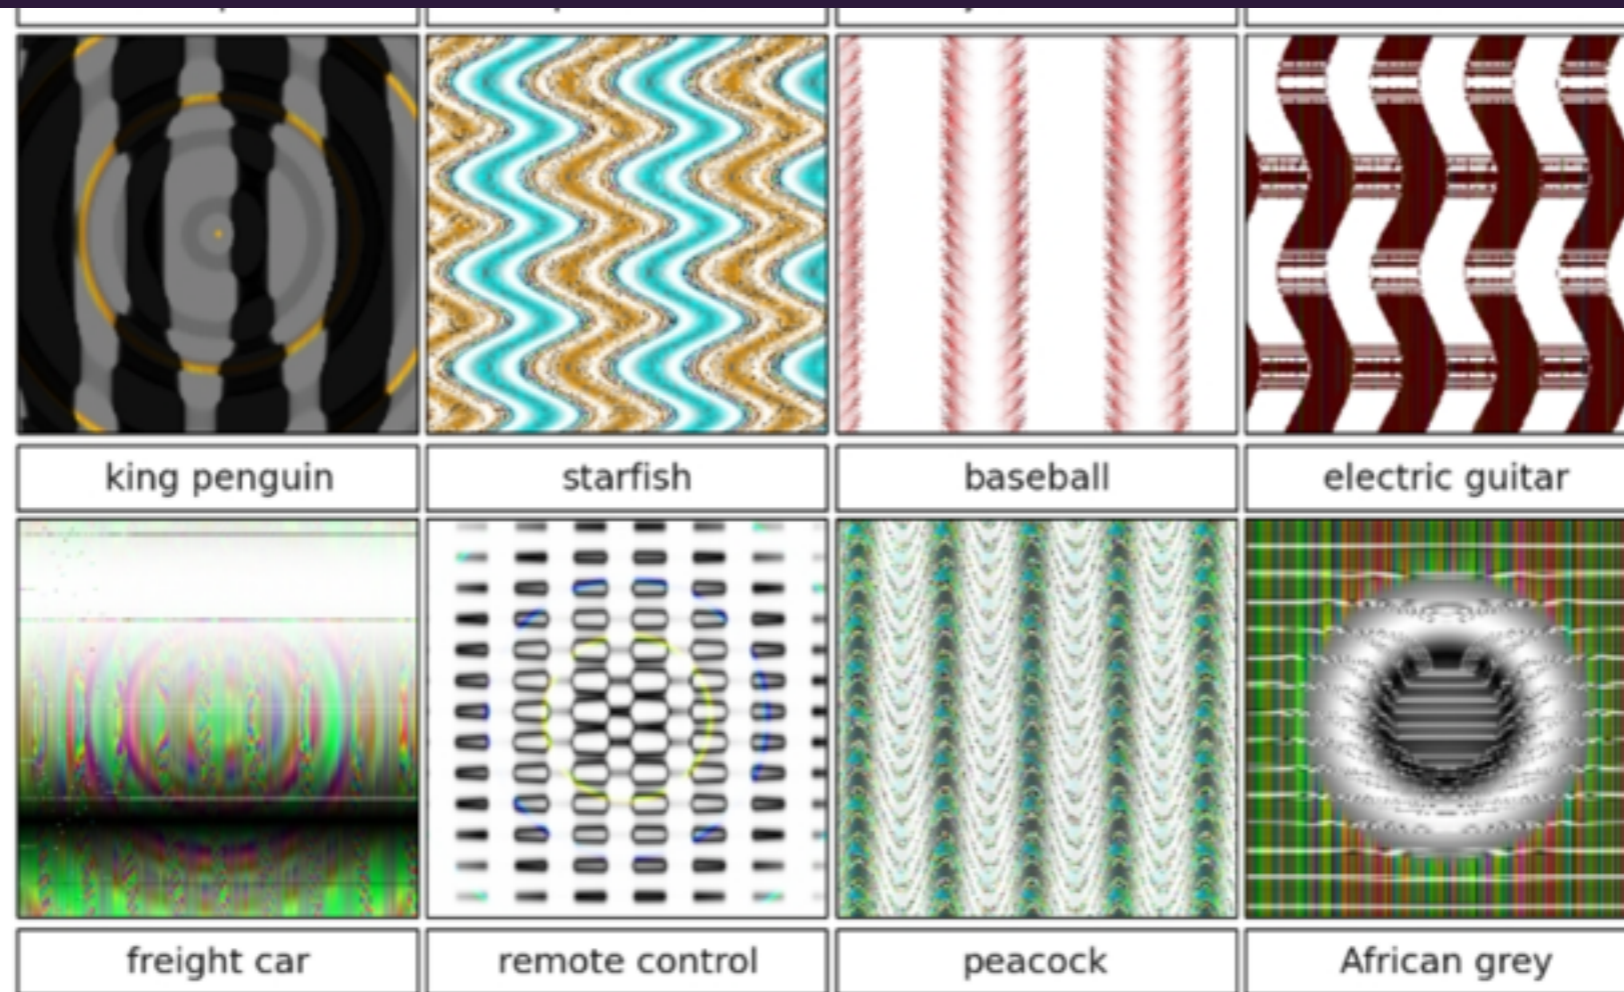


Figure 1. Fooling images that are unrecognizable to humans






starfish




Perhaps, the brain can tell us more about how to  
build truly intelligent machines

Neocortex




Old brain




Assumptions



Learning efficiency



Generality



By discovering neocortex, nature stumbled upon the 'magic middle' of learning architectures

This 'magical architecture' is :


General enough to solve multiple problems like vision, audition, somatosensory etc.

But specific enough to learn efficiently.


(This also implies that the same set of architectural assumptions work well for multiple domains)

Your brain is not very good at recognizing QR codes









Complex cells (Feature pooling)

Simple Cells (Feature detection)

# Hubel & Wiesel

**RECEPTIVE FIELDS, BINOCULAR INTERACTION  
AND FUNCTIONAL ARCHITECTURE IN  
THE CAT'S VISUAL CORTEX**


**BY D. H. HUBEL AND T. N. WIESEL**

*From the Neurophysiology Laboratory, Department of Pharmacology  
Harvard Medical School, Boston, Massachusetts, U.S.A.*


*(Received 31 July 1961)*

Neuroscience has had many more advances since 1961.  
Can't we use that?

How?



How can we decipher  
information-processing principles  
from the brain?




Neocortex




Physics of World's Data




Computational  
Framework







Neocortex  
Source of assumptions/  
constraints






Neocortex  
Source of assumptions/  
constraints




Physics of World's Data  
To find correspondence  
with neocortex properties

Neocortex  
Source of assumptions/  
constraints




Physics of World's Data  
To find correspondence  
with neocortex properties

Computational Framework  
Understand why neocortex  
does what it does to design  
algorithms





# Observed hierarchy in the cortex



Hierarchical  
structure of data


Efficiency and re-use.

# Spikes




?

?



# Blue brain project





Machine learning, statistics

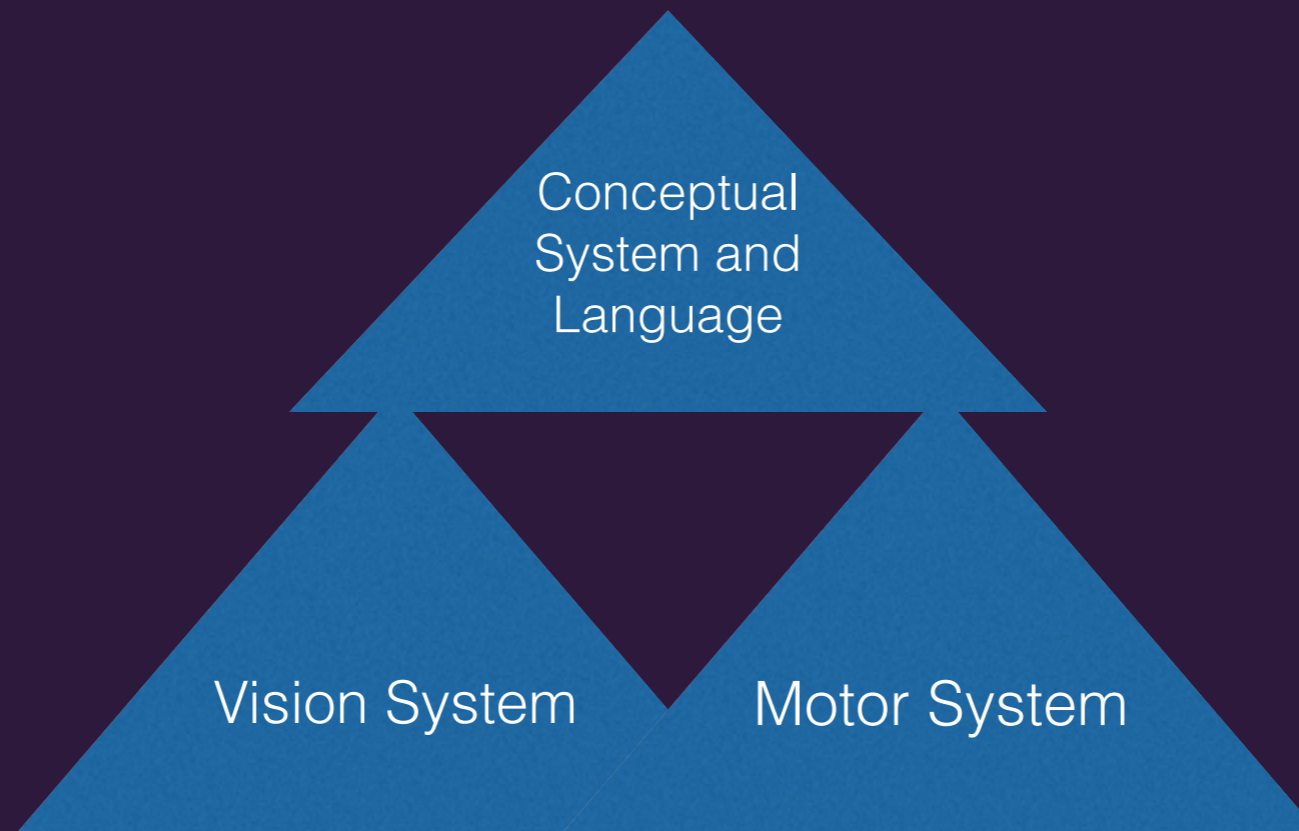


# Brain corporation



# New-brain research questions

- Hierarchy + Feedback + Temporal Learning & Inference
- Scene understanding
- Sensori-motor integration
- Concept learning
- Language understanding



Conceptual  
System and  
Language

Vision System

Motor System

How we solved CAPTCHAs

# We stuck to new-brain methods:

- Small amounts of training data
  - Trained on clean examples
  - Trained only on positive examples
- Unsupervised training
- ....

# Few training examples



**Interpolation** of millions of training images

VS

**Extrapolation** from few examples

# Solves All Variations



Total training examples < 1000 !




Our algorithm produces very detailed segmentation





even when the contours are occluded

**a**



# ICDAR

short

hungry

BEER

Campus

Summer's

animal

hot





# Scene parsing

input

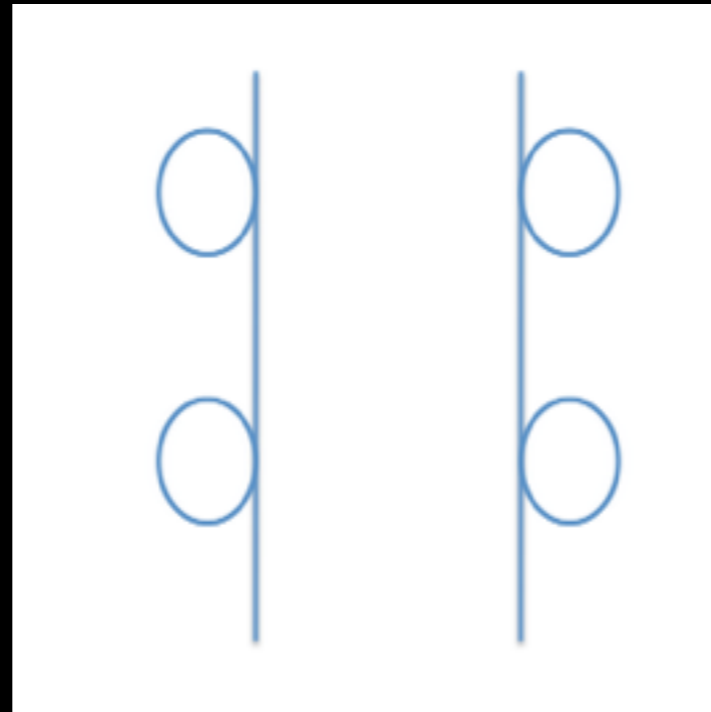
1st (closest)

2nd


3rd



Systems with imagination..




What is this picture?




It is the picture of a bear climbing a tree. Can you see the bear?



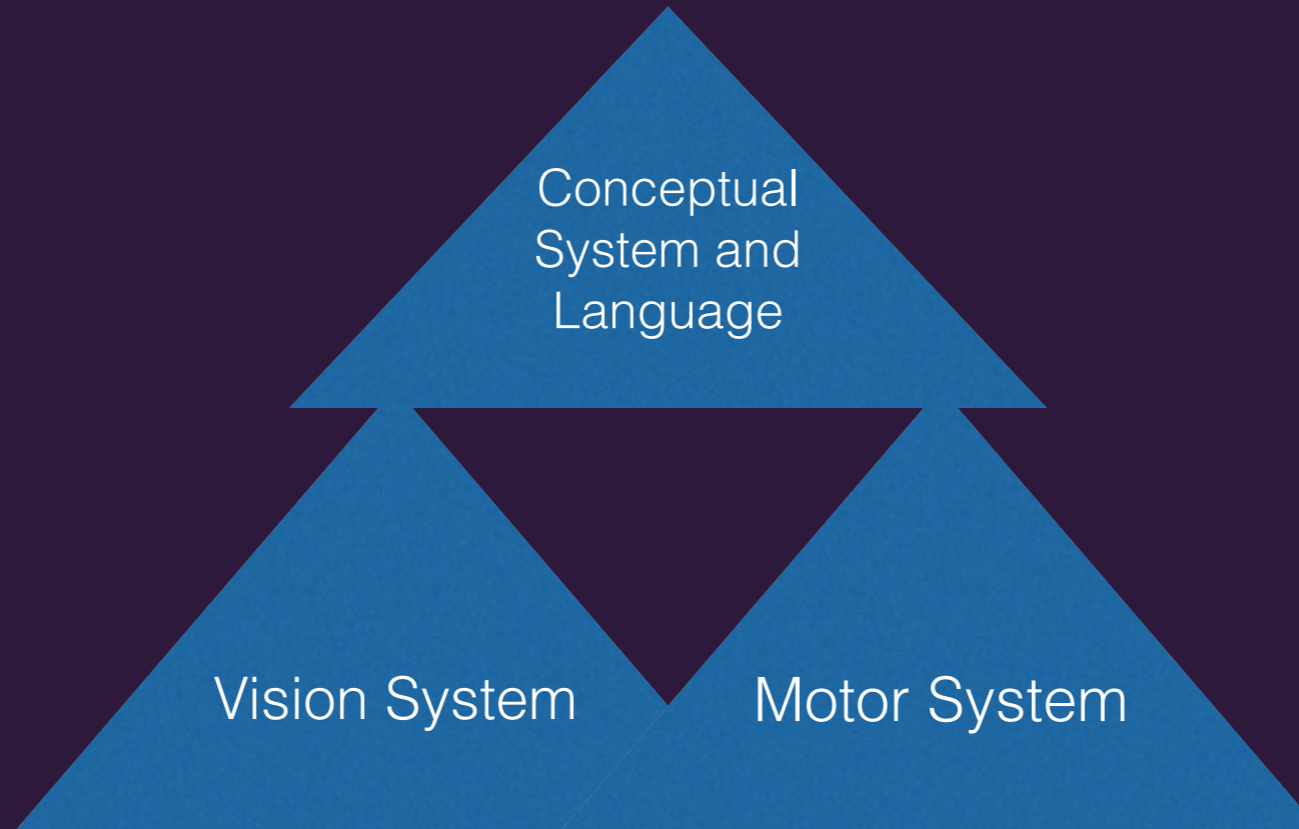


climbing\_bear



Imaginations from our system!


Why is imagination important for AGI?




Conceptual  
System and  
Language

Vision System


Motor System




Wind  
Tunnel



Kite



Glider



Wind  
Tunnel


Major breakthroughs in A.I will need some  
wind-tunnel style research

When?

Will we solve the fundamental research problems in N years?

$N \leq 5$	No way
$5 < N \leq 10$	Small possibility
$10 < N \leq 20$	$> 50\%$

## Increase in brain size during human evolution





About intelligence explosion...

Self-limiting forces on A.I that will prevent an uncontrolled growth



- Data limit
- Intervention limit
- ...

- Data limit
  - Super-human A.I will encounter limits of available data
  - Limit on recursive improvement
    - Going over the same data again and again does not increase the information in the data. (Data Processing Inequality)
  - Natural dynamics of data generation

Suppose super-human A.I was created before we had any knowledge about how gravity worked and how earth moved around the sun.

How long will it take for the A.I to discover that it takes the earth 365 days to go around the sun?

The dynamics of the world imposes limits on how fast data can be acquired.











vicarious