15-883: Computational Models of Neural Systems

Lecture 1.1:

Brains and Computation

David S. Touretzky
Computer Science Department
Carnegie Mellon University

Why Do Modeling?

- Models help to organize and concisely express our thoughts about the system being modeled.

- Good models make testable predictions, which can help guide experiments.

- Sometimes a computational model must be implemented in a computer simulation in order to explore and fully understand its behavior.
  - Surprising behavior may lead to new theories.

Models of the Nervous System

- Hydraulic network (Descartes): nerves = hoses that carry fluid to drive the muscles

- Clockwork: systematic and representational

- Telephone switchboard: communication

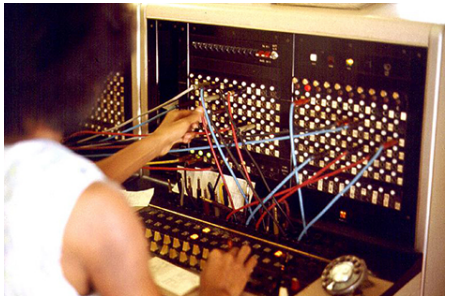

- Digital computer (“electronic brain”): computational

Metaphors can serve as informal theories.

Help to frame the discussion.

But limited in predictive power.

Do Brains Compute?

Most scholars believe the answer is “yes”.

Brains are meat computers!

Some consider this conclusion demeaning.

Computers are machines. I am not a machine!

Some try to find reasons the answer could be “no”.

Example: if unpredictable quantum effects played a crucial role in what brains do, then the result would not be describable as a computable function.

How Can A Physical System Perform "Computation" ?

It's a subjective judgment. What to look for:

1) Its physical states correspond to the representations of some abstract computational system.

2) Transitions between its states can be explained in terms of operations on those representations.

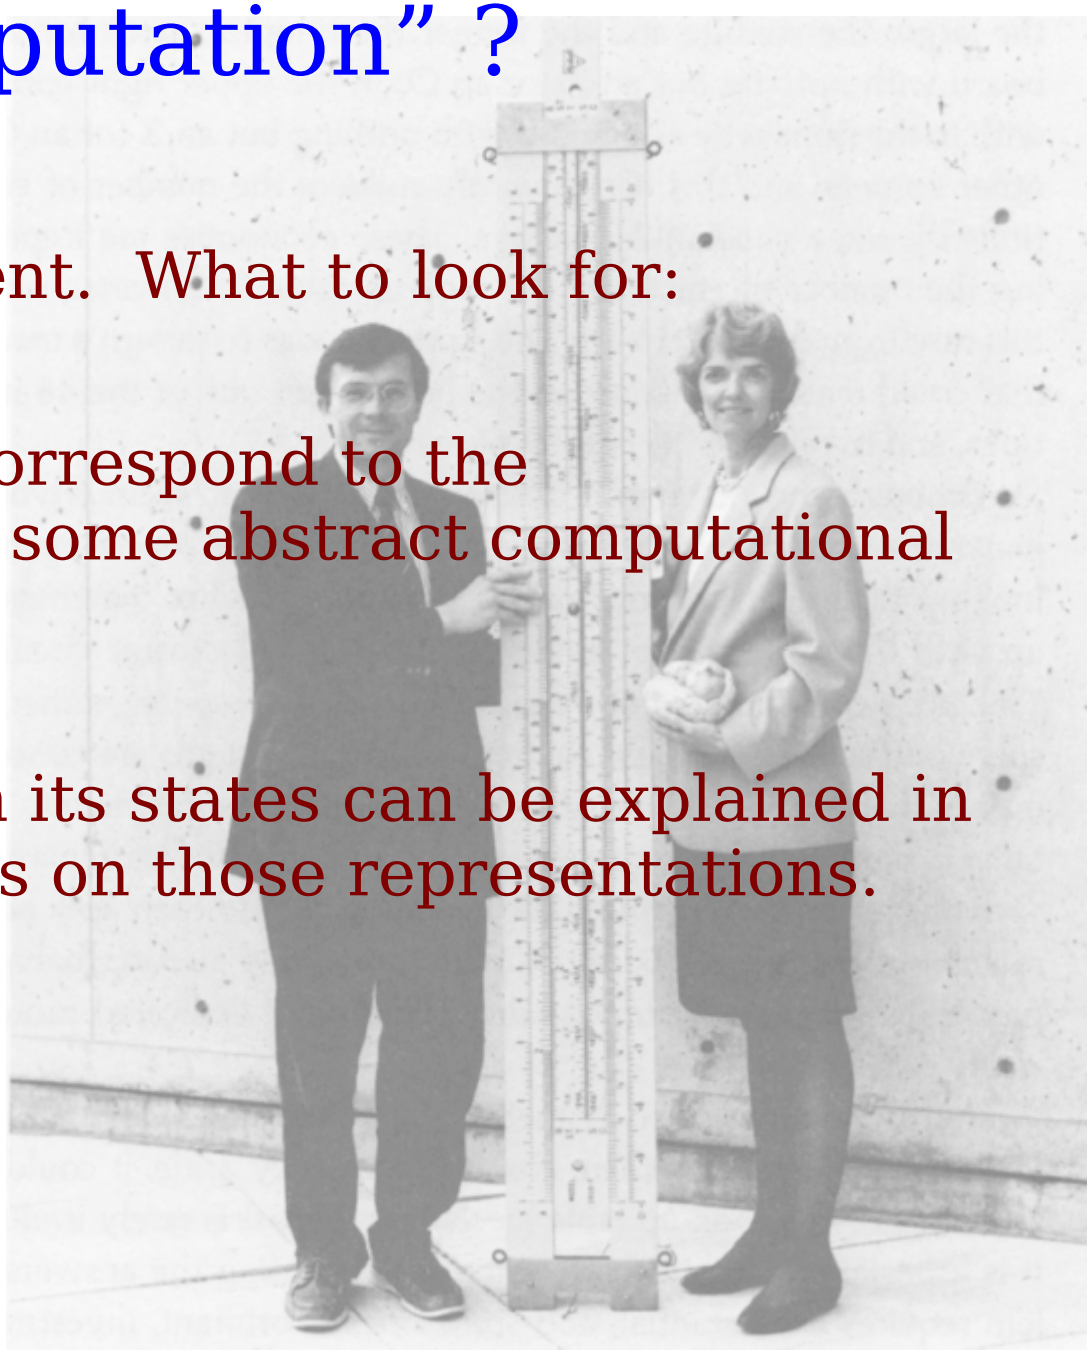

Terry Sejnowski and Patricia Churchland, authors of *The Computational Brain*

Physical Computation: The Slide Rule

- Abstract function being computed: multiplication
  - Input: a pair of numbers
  - Output: a number

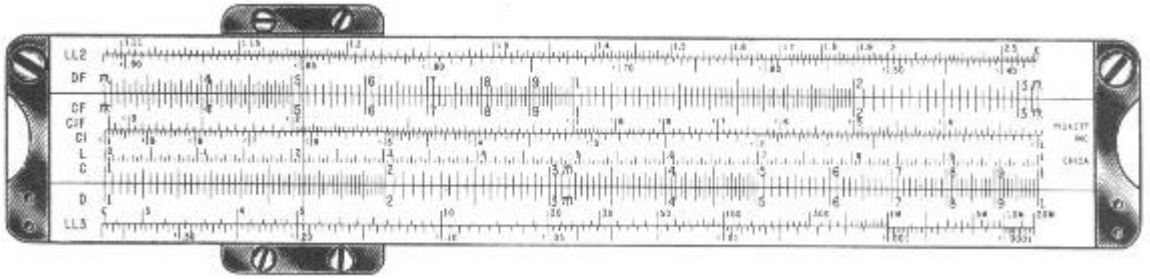

- Physical Realization:
  - First input = point on surface of the (fixed) D scale
  - Second input = point on surface of the (sliding) C scale
  - Output = point on surface of the (fixed) D scale

Slide Rule Computation: Multiply 2.05 by 3

- Move the sliding C scale so that the digit “1” is at 2.05 on the D scale.

C
D

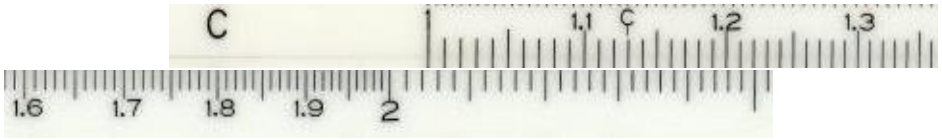

- Slide the cursor so that the red index is over the 3 on the C scale. Read the result 6.15 on the D scale.

C
D

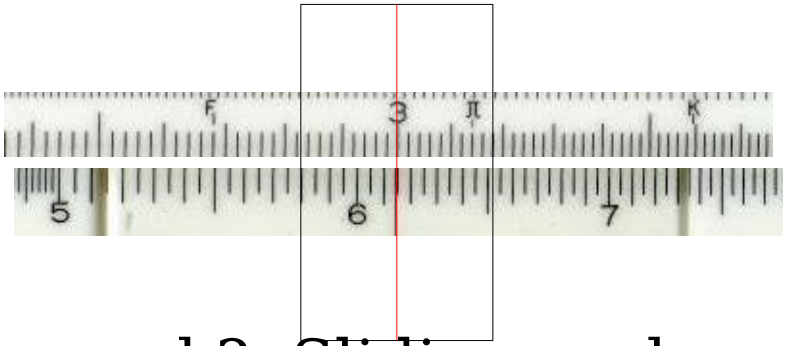

- Why does this work? Sliding scale and cursor = addition. Multiplication comes from adding logs.

Tinkerytoy Tic-Tac-Toe Computer

Designed by Danny Hillis at MIT.

See Scientific American article for details.

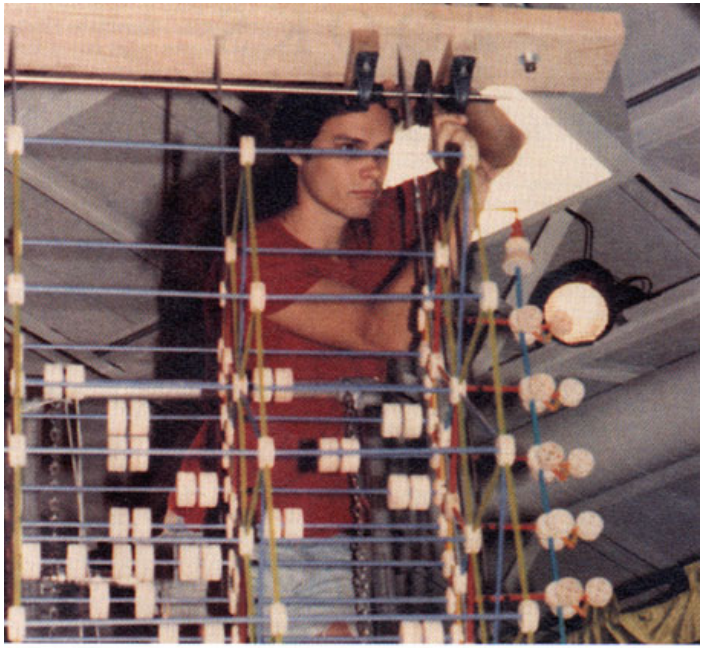

Edward Hardebeck helps to assemble the Tinkertoy computer

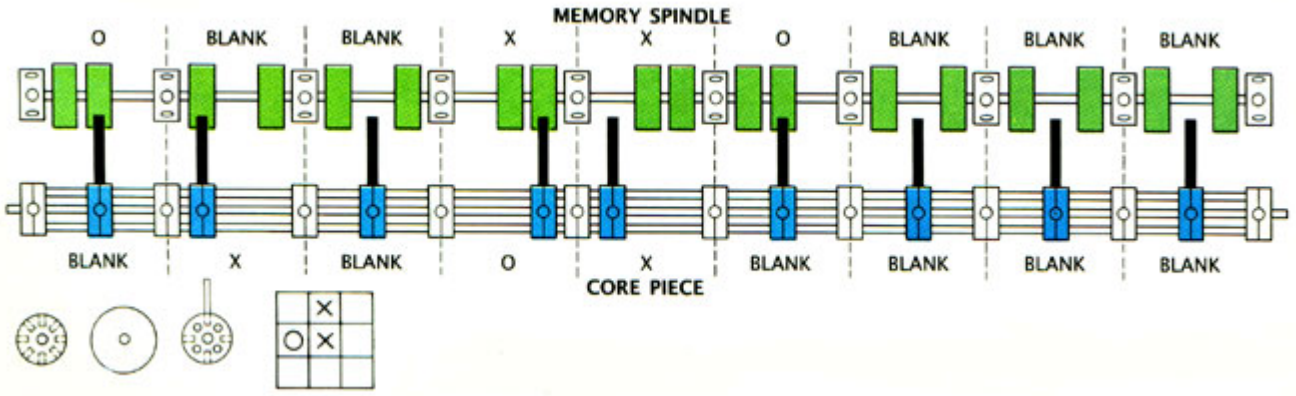

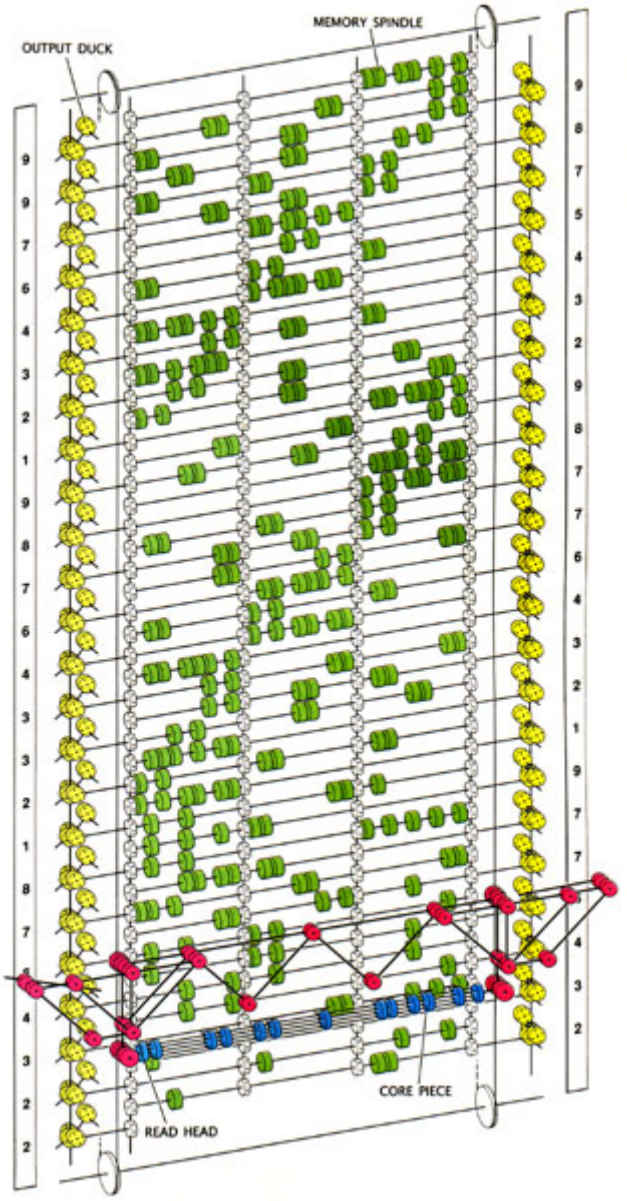

The Tinkertoy computer: ready for a game of tic-tac-toe

How Big Are Meat Computers? Some Numbers

	Neurons	Synapses
Humans	**86 billion**	**150 trillion** (1.5×10^{14})
Rats	200 million	450 billion (4.5×10^{11})
1 mm^3 of cortex	10^5	10^9

A cortical neuron averages 4.12×10^3 synapses (cat or monkey.)

Computational Resources

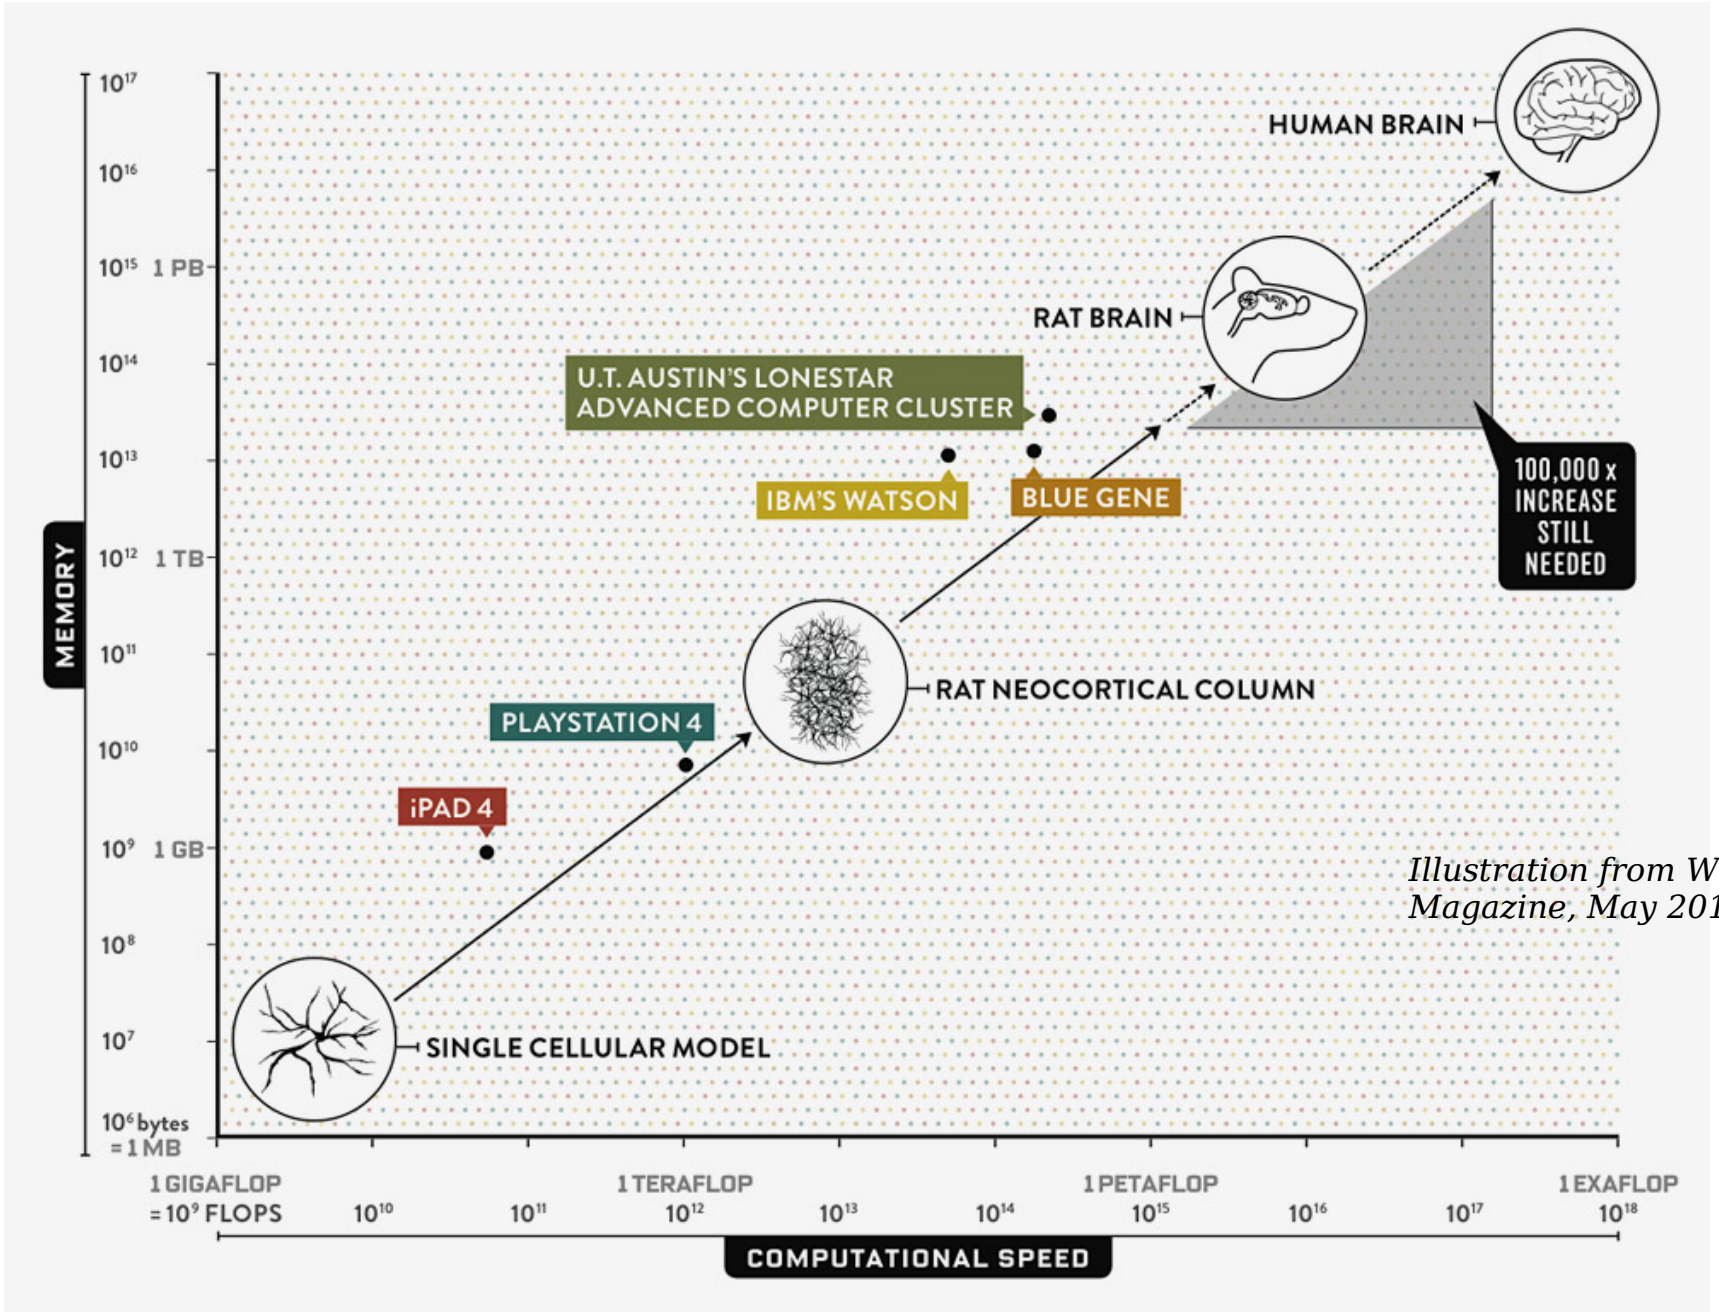

Illustration from Wired Magazine, May 2013.

“Building A Brain”

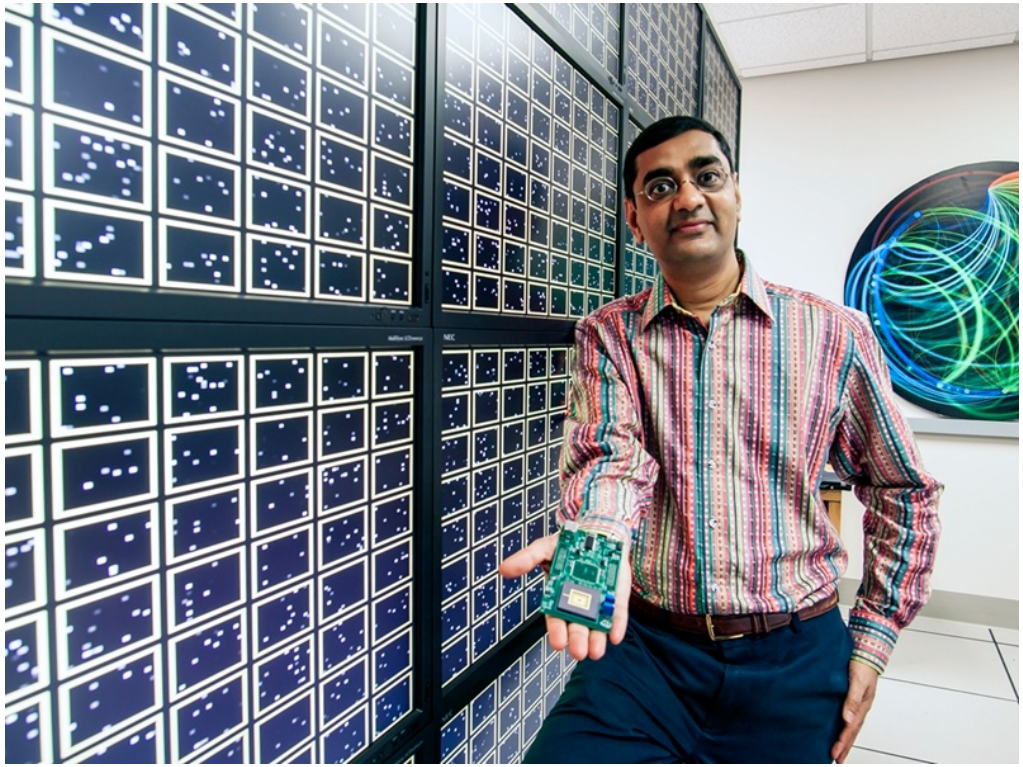

IBM's Dharmendra Modha

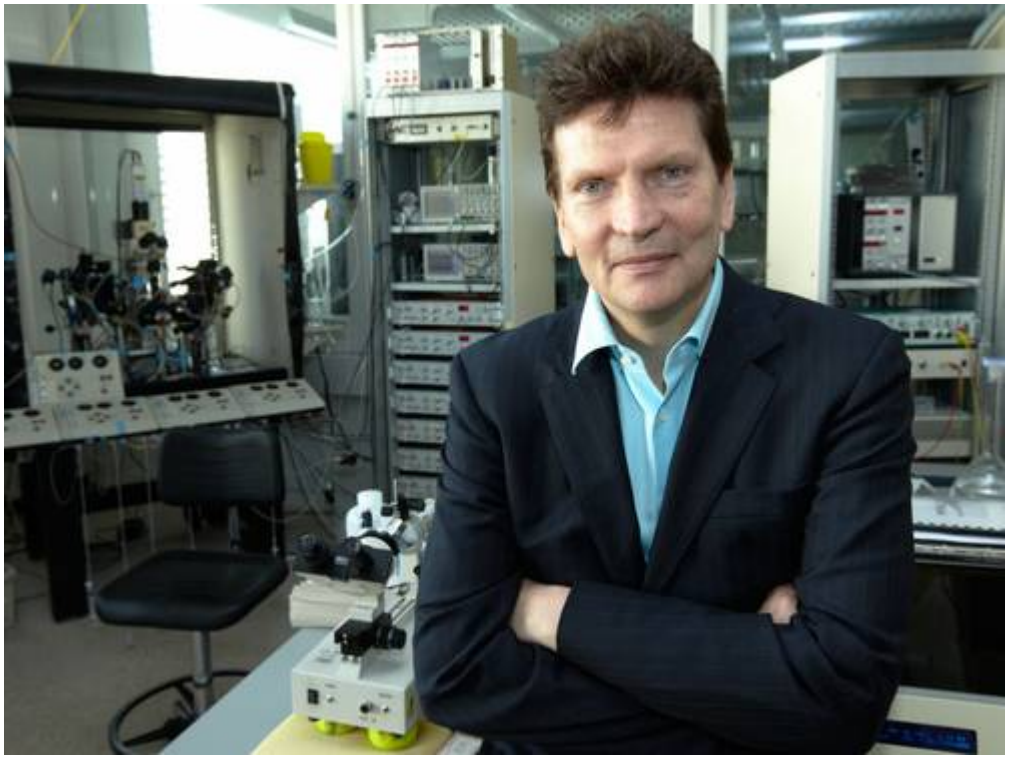

EPFL's Henry Markram

Do We Have All the Math We Need to Understand the Brain?

- Probably not yet.
- People have tried all kinds of things:
  - Chaos theory
  - Dynamical systems theory
  - Particle filters
  - Artificial neural networks (many flavors)
  - Quantum mechanics
- We can explain simple neural reflexes, but not memory or cognition.
- Current theories will probably turn out to be as wrong as Aristotelian physics.

Which Rock Hits the Ground First?

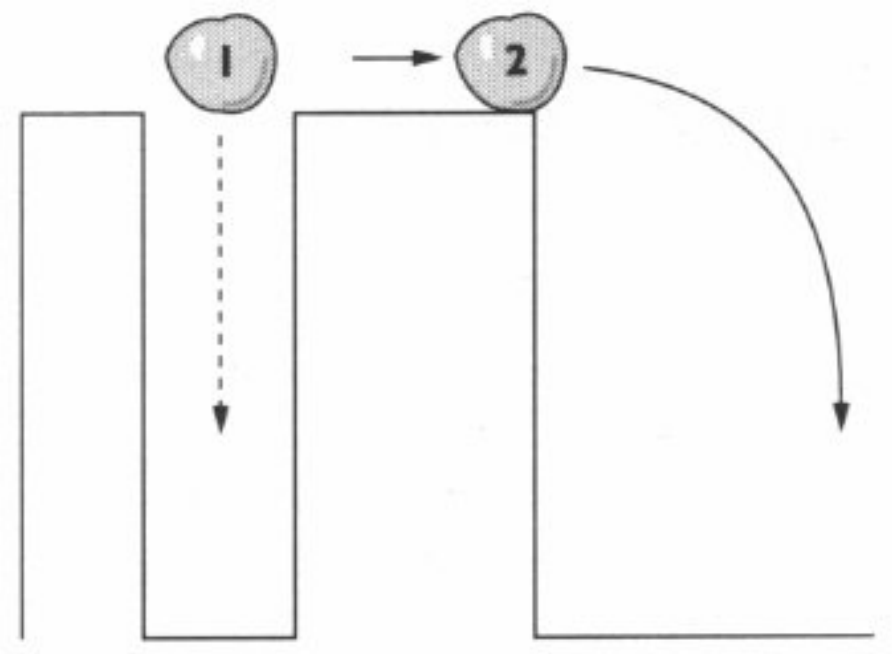

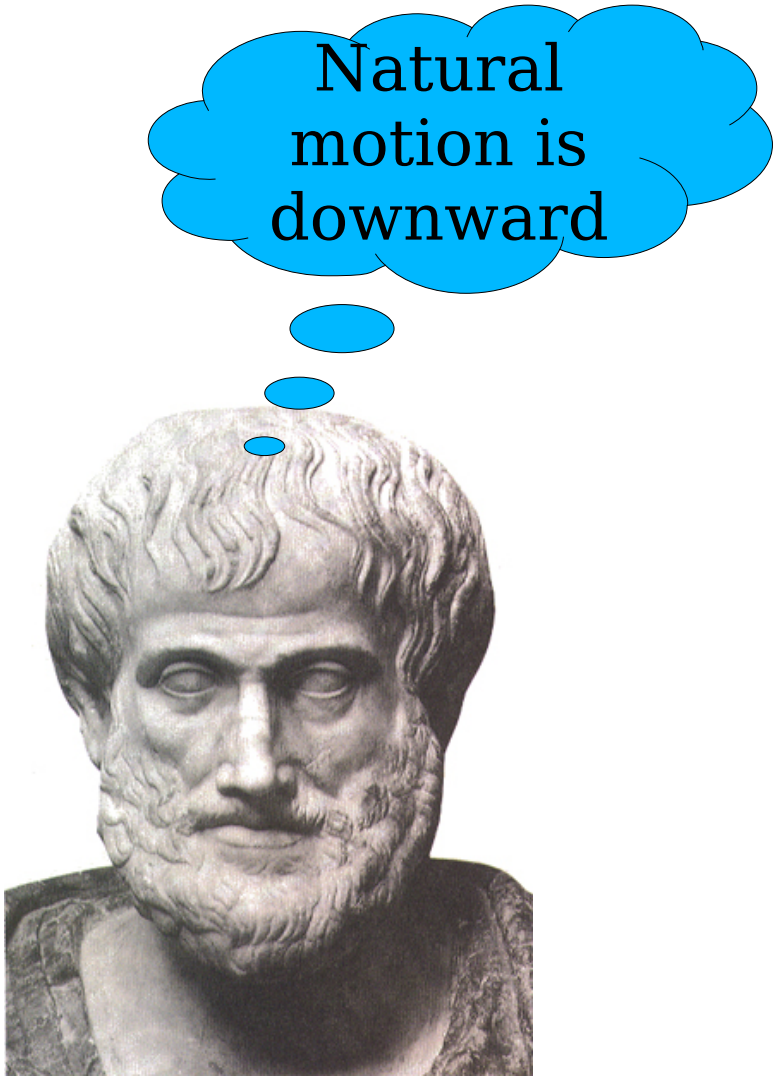

Aristotle (384-322 BCE)

Aristotelian Motion

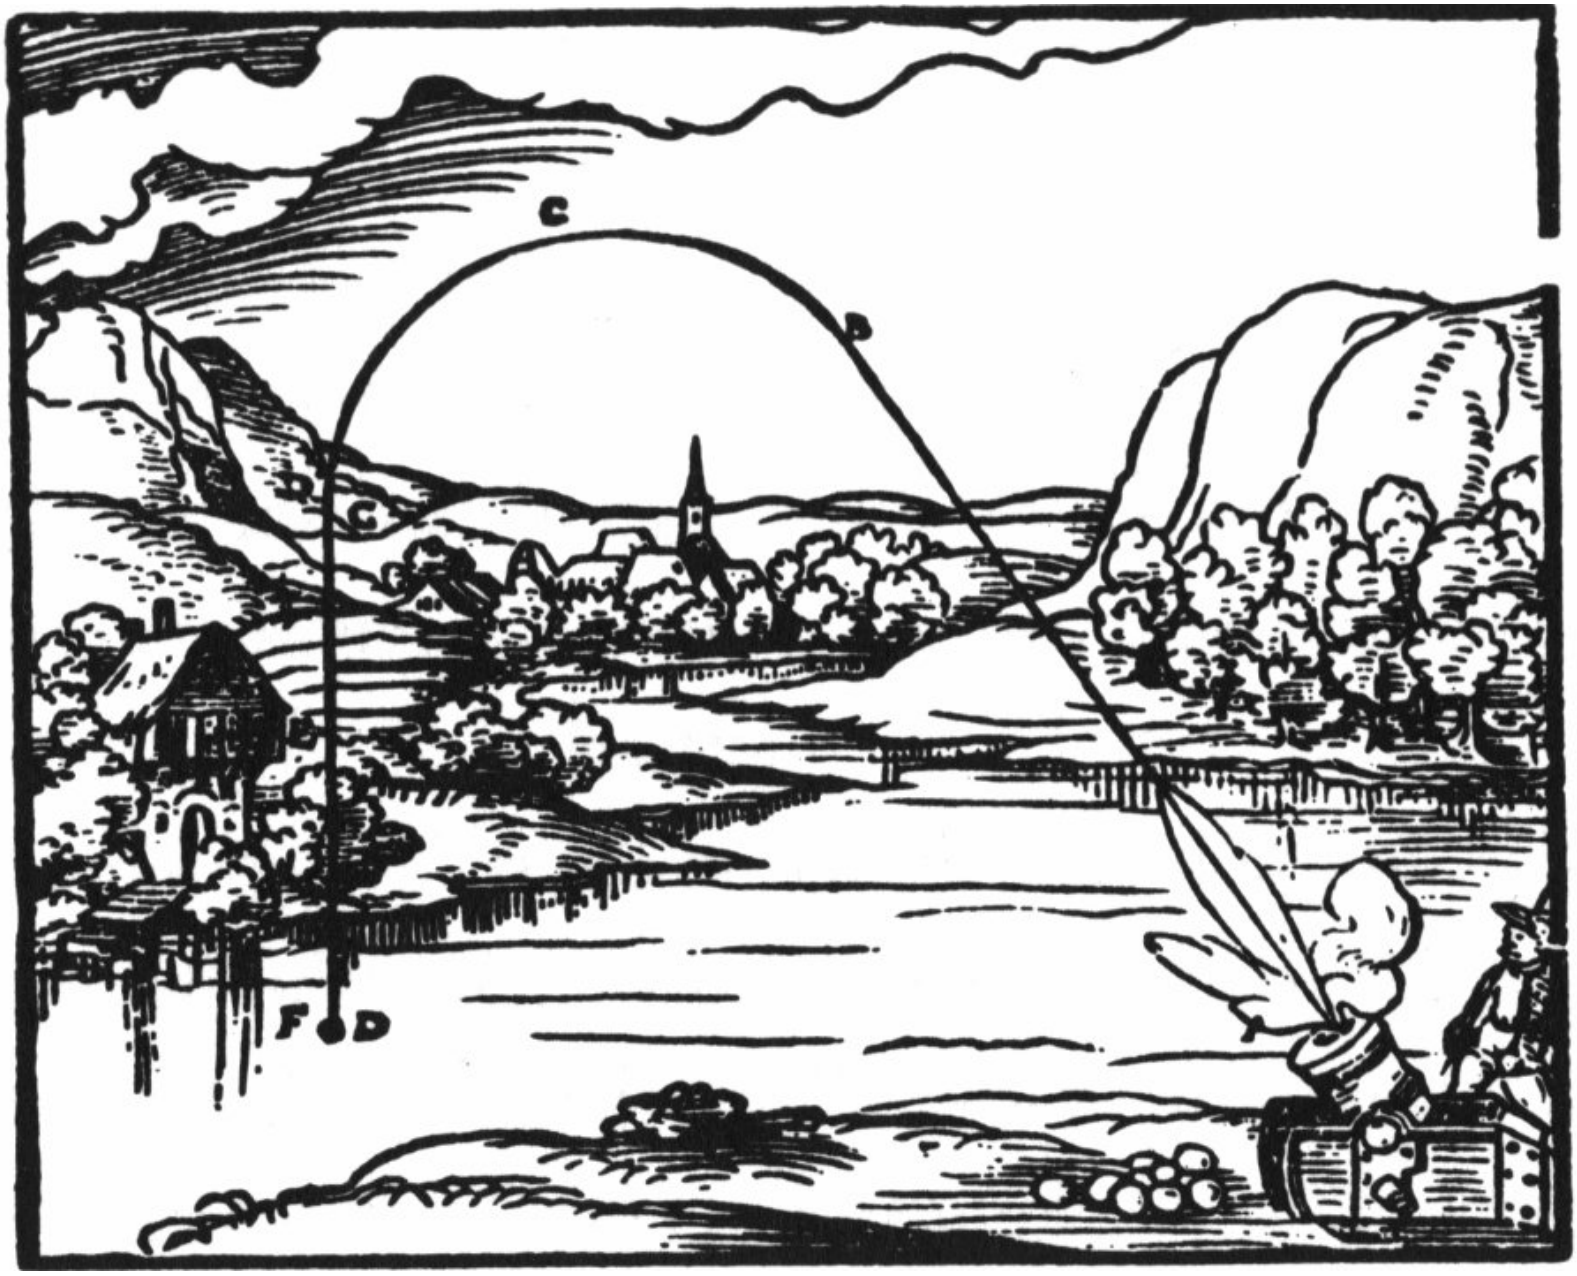

Galileo: Motion is Parabolic and Independent of Mass

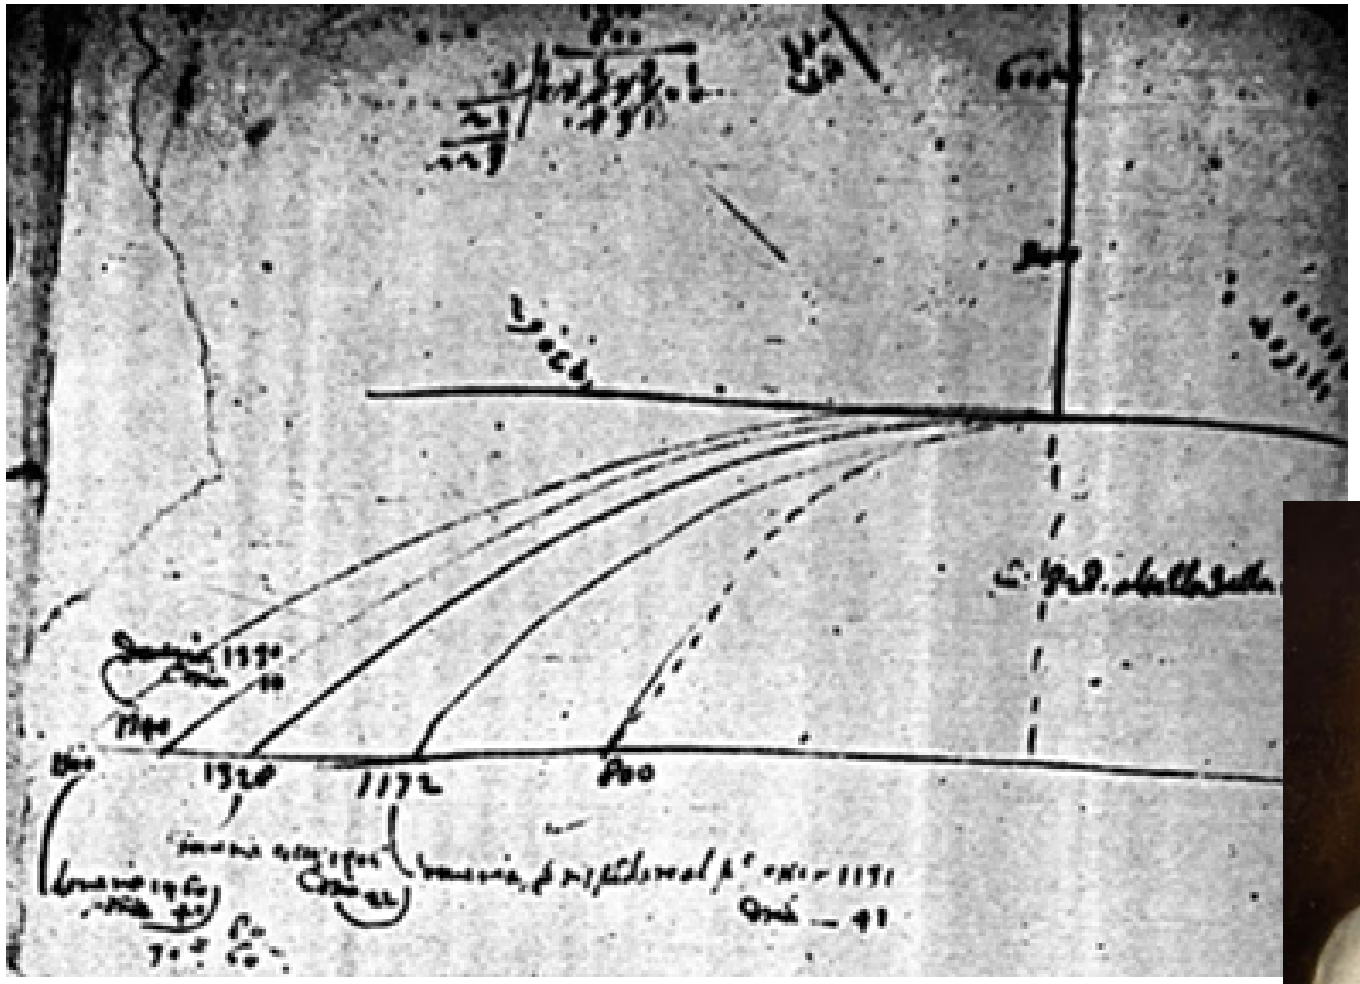

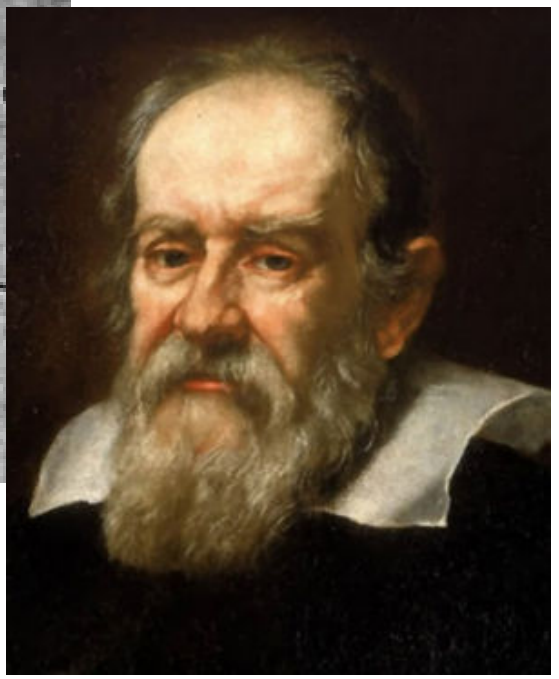

Galileo Galilei (1564-1642)

Why a Parabola? Need Calculus

$$a(t) = -9.8\, m/s^2$$

$$v(t) = \int a(t)\,dt = -9.8\,t + v_0$$

$$h(t) = \int v(t)\,dt = -9.8\,t^2/2 + v_0 t + h_0$$

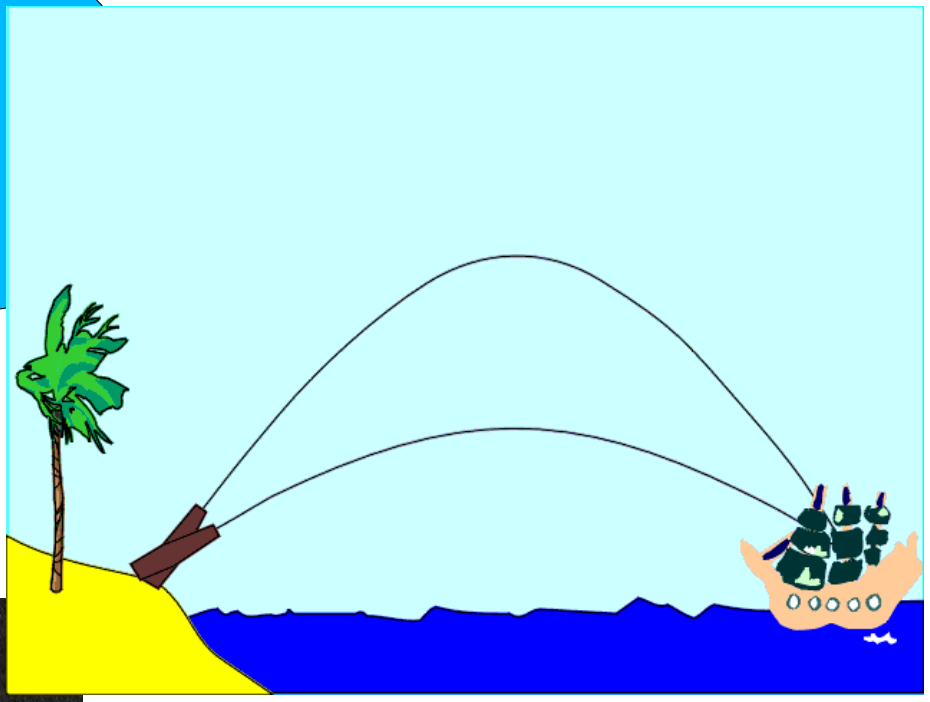

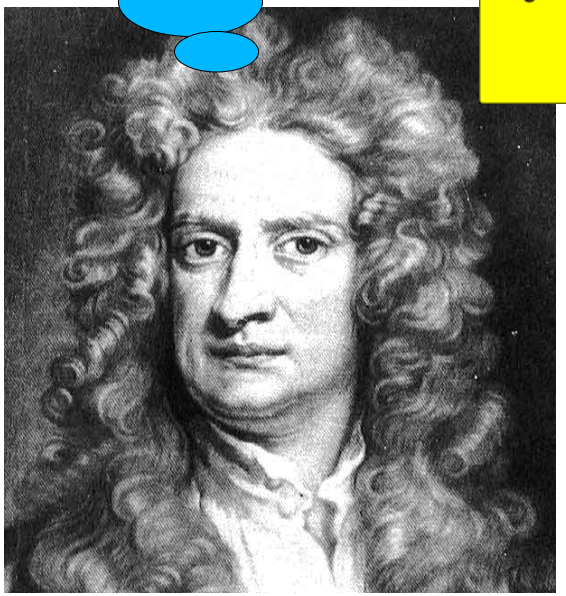

Isaac Newton (1643-1727)

Relativistic Motion: Curved Spacetime

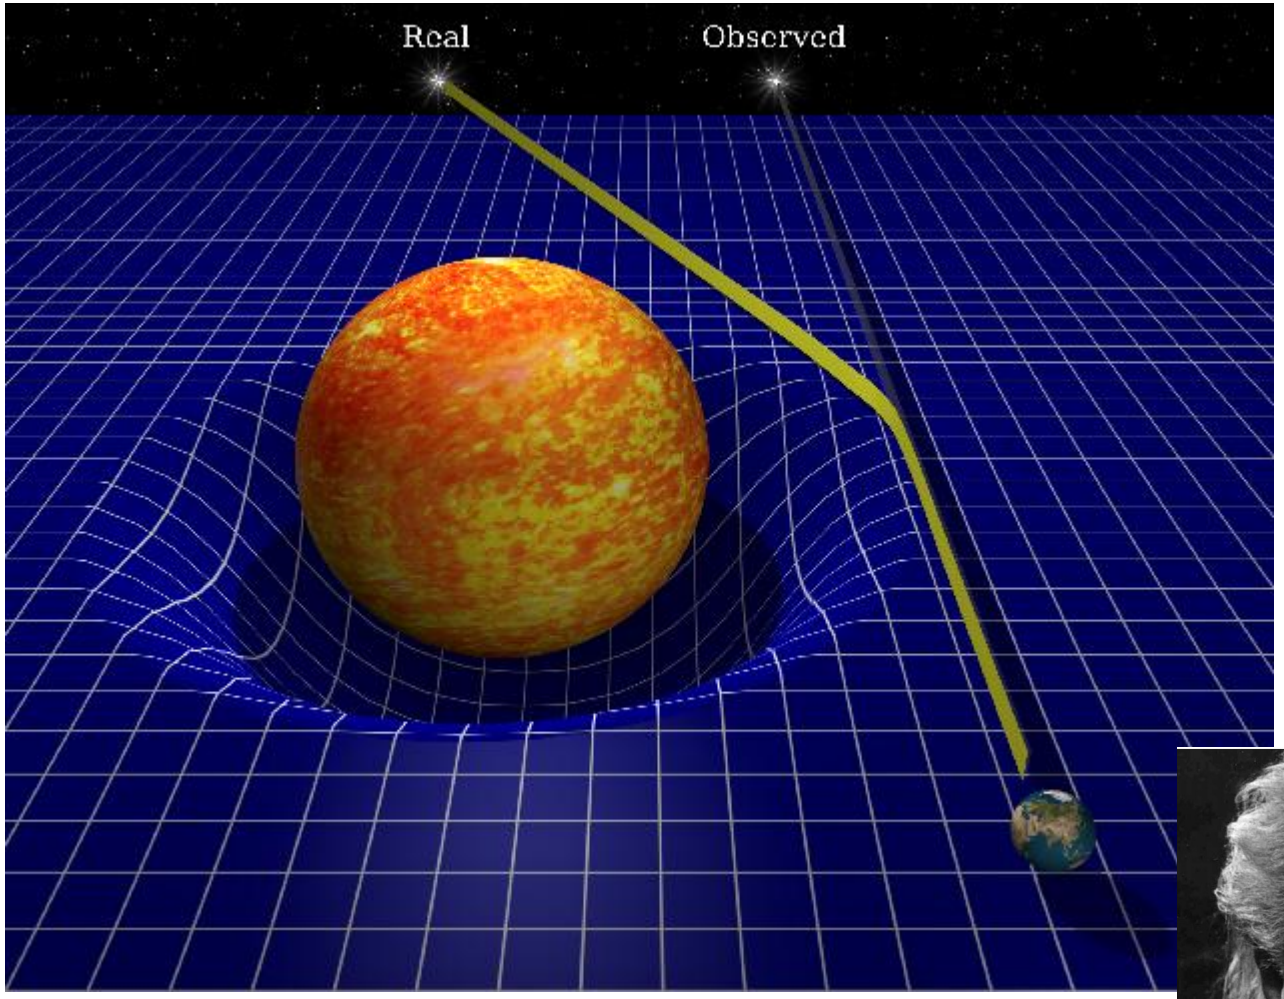

For this theory you need tensor calculus.

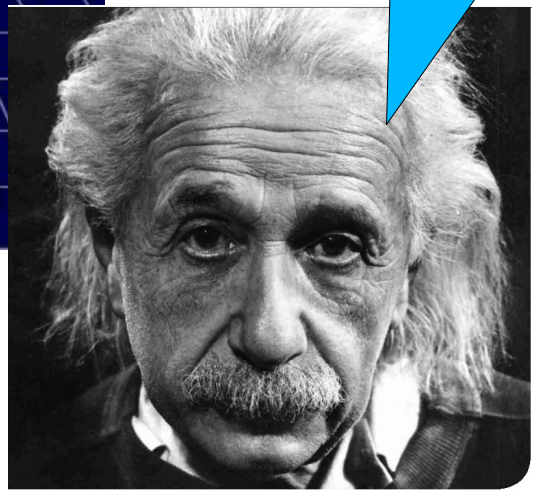

Albert Einstein (1879-1955)

Computational Processes Posited in the Brain

- Table lookup / associative memory.
- Competitive learning; self-organizing maps.
- Principal components analysis.
- Gradient descent error minimization learning.
- Temporal difference learning.
- Dynamical systems (attractor networks, parallel constraint satisfaction).

This course will explore these models and how they apply to various brain structures: hippocampus, basal ganglia, cerebellum, cortex, etc.

Some Representative Successes (1)

Dopamine cells in the midbrain fire in response to rewards, but also in response to neutral stimuli that have become associated with rewards. But they can also stop firing with further training, or pause when a reward is missed.

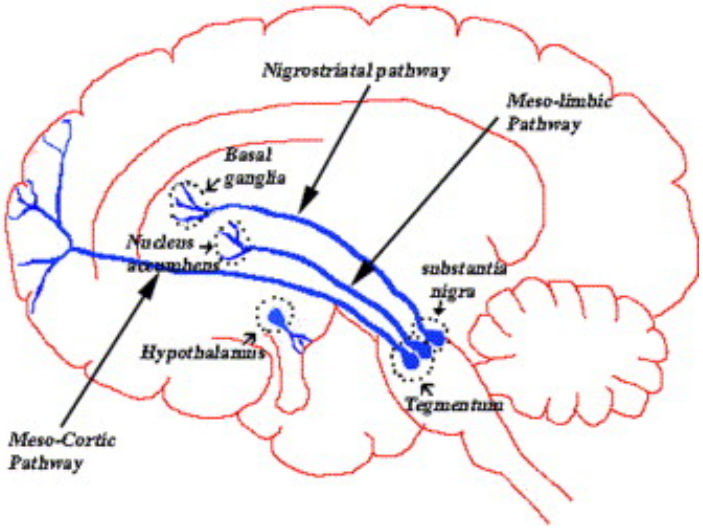

Why should they do that?

Temporal difference learning, a type of *reinforcement learning*, neatly explains much of the data.

Some Representative Successes (2)

Most cells in **primary visual cortex** (V1) get input from both eyes but have a dominant eye that they respond more to. Radioactive tracing shows zebra-like “ocular dominance” stripes.

How does this structure emerge?

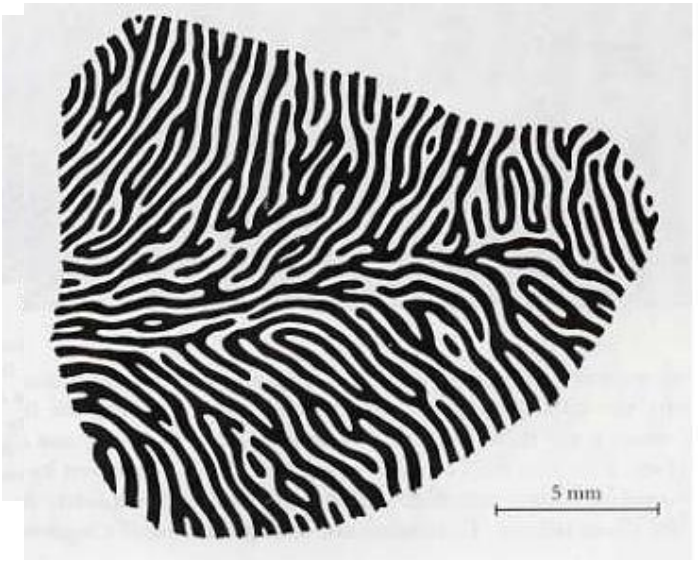

Competitive learning algorithms, a type of *unsupervised learning*, can account for the formation of ocular dominance and orientation selectivity in V1.

Science vs. Engineering

- Science: figure out how nature works.
  - Good models are as simple as possible.
  - Models should reflect reality.
  - Models should be falsifiable (make predictions).
- Engineering: figure out how to make useful stuff.
  - “Good” means performs a task faster/cheaper/more reliably.
  - Making a system more “like the brain” doesn't in itself make it better.
- Holy grail for CS/AI people: use insights from neuroscience to solve engineering problems in perception, control, inference, etc.
  - Hard, because we don't know how brains work yet.

Want to Build a Brain?
Some Disappointing News:

- We're still in the early days of neural computation.
- Our theories of brain function are vague and wrong.

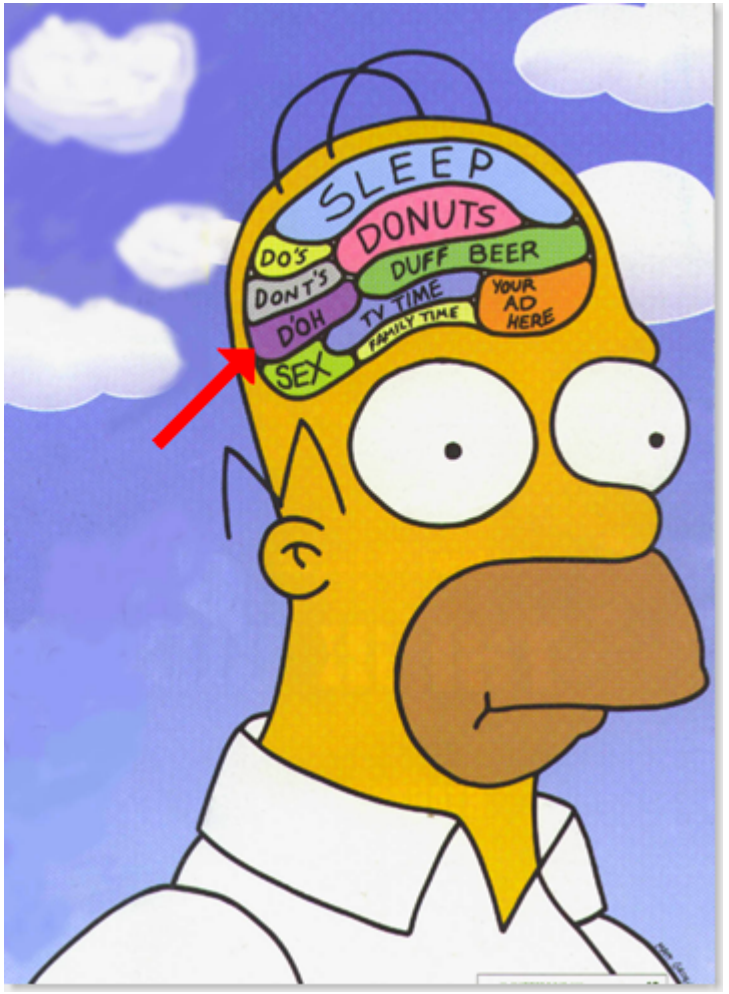

The Misunderstood Brain

- We know **a lot** about what makes neurons fire.
- We know **a good deal** about wiring patterns.
- We know **only a little** about how information is represented in neural tissue.
  - Where are the “noun phrase” cells in the brain?
- We know **almost nothing** about how information is processed.
- This course explores what we do know. There is progress every month.
- It's an exciting time to be a computational neuroscientist.

Brain Anatomy

- How do we uncover the anatomy of the brain?
  - Anatomical dissection + visual inspection
  - Microscopy: look for different cell types; staining
  - Tract tracing using radioactive labels, viruses, diffusion tensor imaging
  - Mapping gene expression

- The Allen Brain Institute's mouse brain atlas lists 737 distinct brain areas. What about humans?

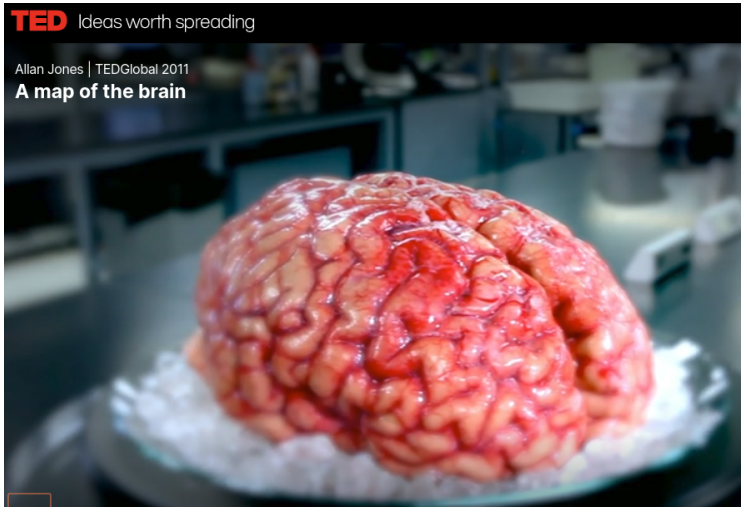

TED Talk Video:
A Map of the Brain.

What's So Special About the Human Brain?

- In general, brain size correlates with body size.
- Humans are smarter than rats and monkeys because we have bigger brains.
- Who has the biggest brain?
  - Dog: 0.4 kg
  - Human: 1.5 kg
  - Elephant: 5 kg
  - Killer whale: 5.4 – 6.8 kg
  - Sperm whale: 8 kg
- Why are humans the smartest?
  - We have the most cortical neurons.
  - Elephants have more neurons (257B vs. 86B) but most of them are in the cerebellum.

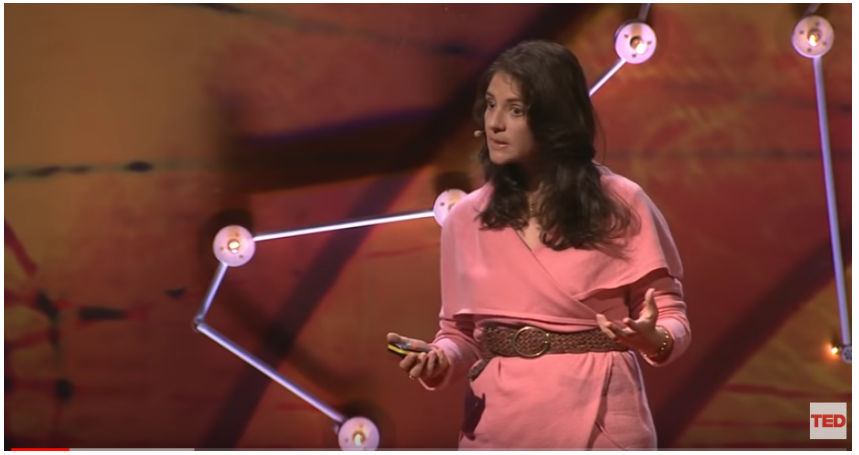

Suzana Herculano-Houzel measured the number of neurons in the brains of humans and other animals.