电子科技大学

University of Electronic Science and Technology of China

IEEE Spectrum Seminar: 2
The Brain of a New Machine

By M. Versace & B. Chandler
Given by Fei You

游飞 Fei You, feiyou@uestc.edu.cn

Outline

- Preface
- The Brain and Neurons
- The Progress of AI
- The Difference - Brain vs. Computer
- Memristors -- the missing link
- To Build a Brain
- Future MoNETA
- Conclusion

Preface

- Stop us if you've heard this one before:
  - In the near future, we'll be able to build machines that learn, reason, and even emote their way to solving problems, the way people do.

The Singularity

- 不断失败的预言
  - 1960年代，晶体管制造成功，人们预测20年后机器就会比人类更加聪明
  - 50年后的今天，除了实现一定程度的自动化，还没有达到该预言
  - 还有没有希望？
  - 科学家都失望了
    - real artificial intelligence can't be done on traditional hardware, with its rigid adherence to Boolean logic and vast separation between memory and processing.

The Hope

- MoNETA
  - – Modular Neural Exploring Traveling Agent
  - – The brain on a chip
  - – Designed by the department of cognitive and neural systems, Boston University
    - It will **perceive** its surroundings, decide which information is useful, **integrate** that information into the emerging structure of its reality, and in some applications, **formulate** plans that will ensure its survival.

The Brain and Neurons

- The Human Brain
  - The human brain is the center of the human nervous system.

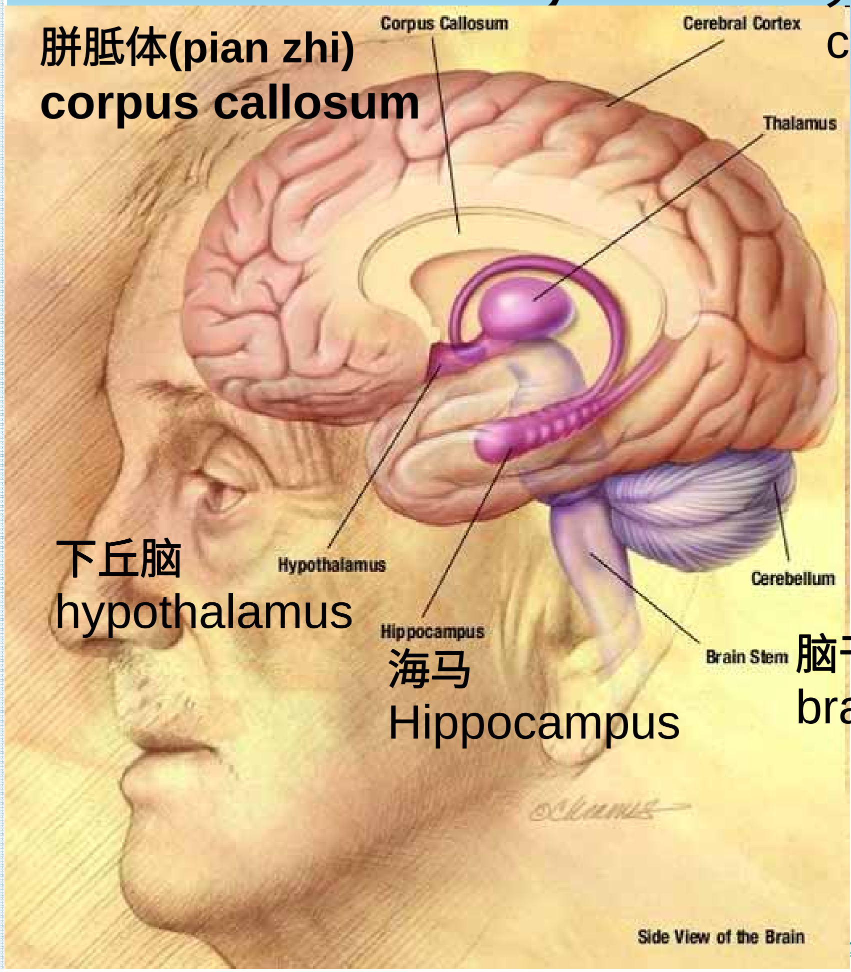

胼胝体(pian zhi)
corpus callosum

大脑皮层
cerebral cortex

丘脑
thalamencephalon

下丘脑
hypothalamus

海马
Hippocampus

小脑
cerebellum

脑干
brainstem

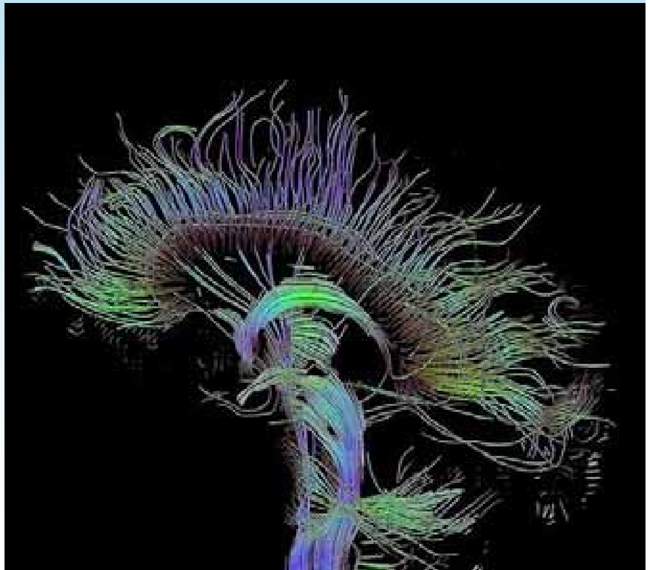

Diffusion Tensor Imaging (DTI) measurement of a human brain
扩散张量成像

The Brain and Neurons

- Neurons

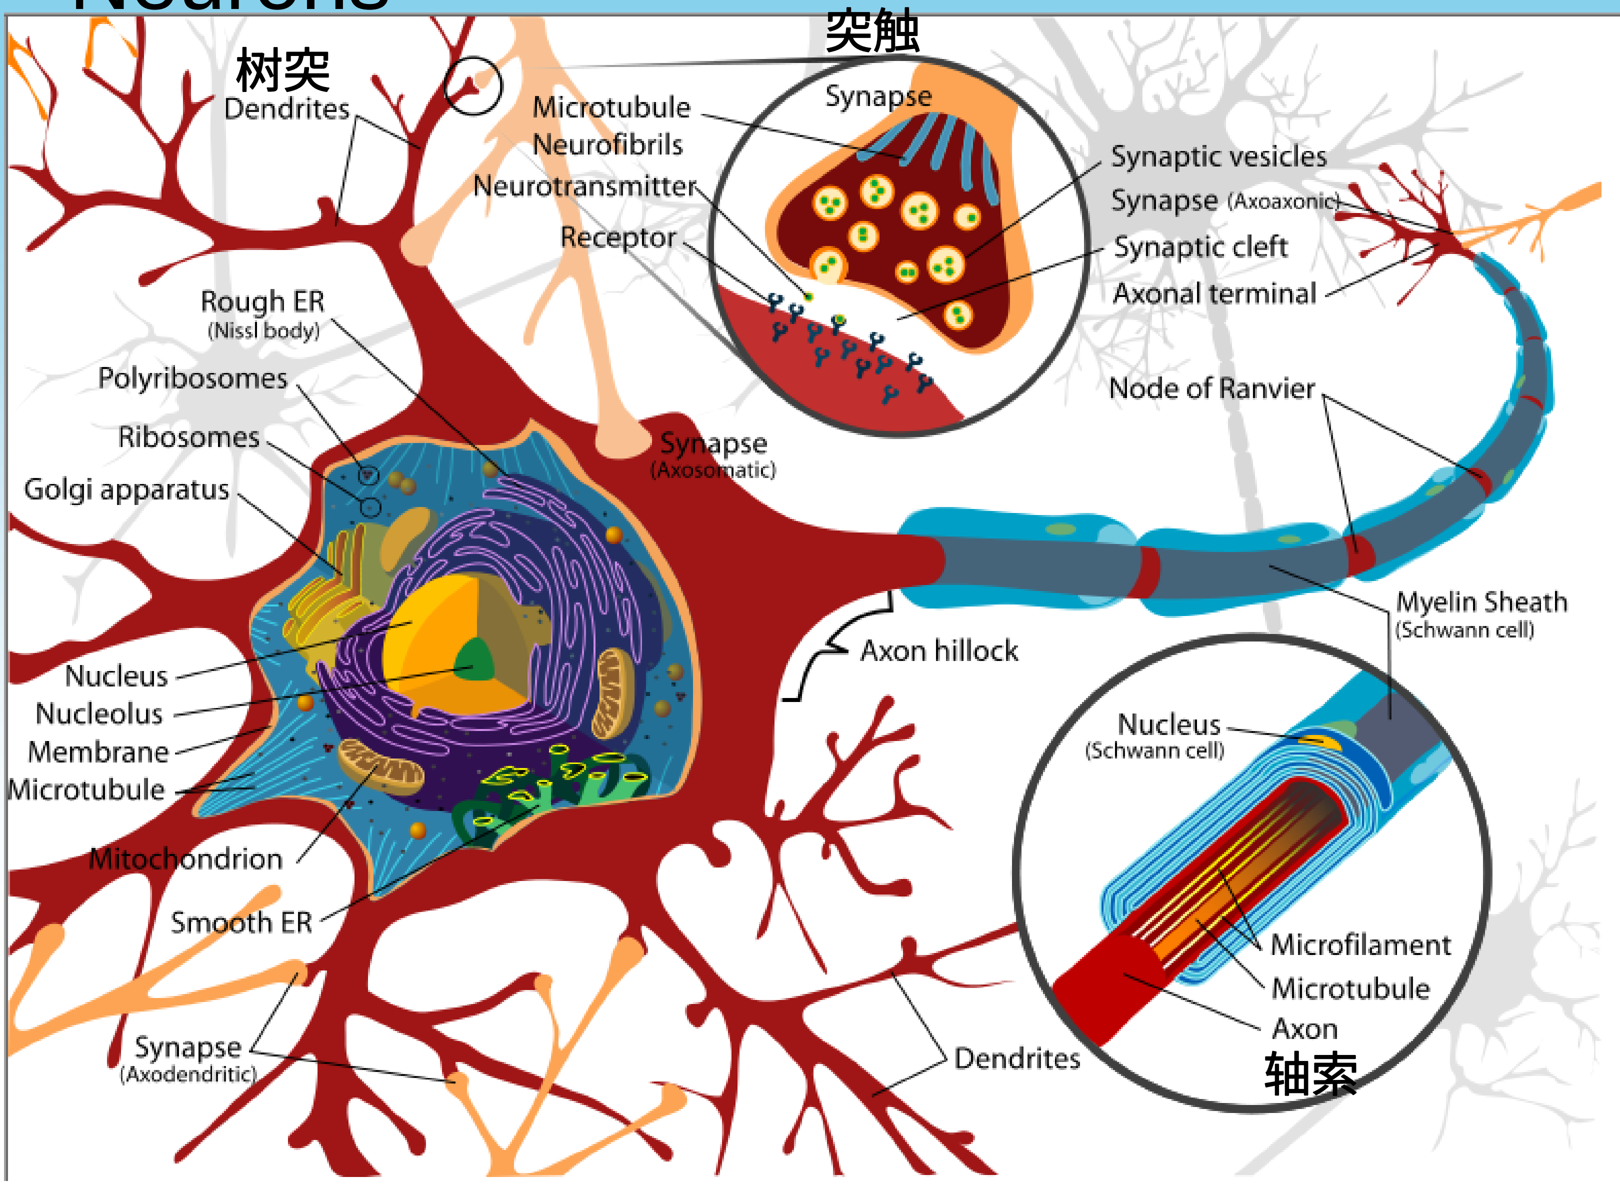

The Brain and Neurons

- Information Processing of Neurons
  - Neurons talk to each other by way of dendrites and axons.
    - in-boxes (dendrites)
    - out-boxes (axons)
  - Transmitting electrical impulses from one neuron to another
  - Most of the processing performed in the junctions , called synapse, between neurons.

The Brain and Neurons

- Dendrites
  - Greek, “Tree”
  - The branched projections of a neuron that act to conduct the electrochemical stimulation received from other neural cells to the cell body
  - It integrates the input from other neurons -- Sodium, calcium, and potassium channels are all implicated in contributing to input modulation -- a weak input from a distal synapse can be amplified --to send action potentials

The Brain and Neurons

- Dendrites
  - Coding

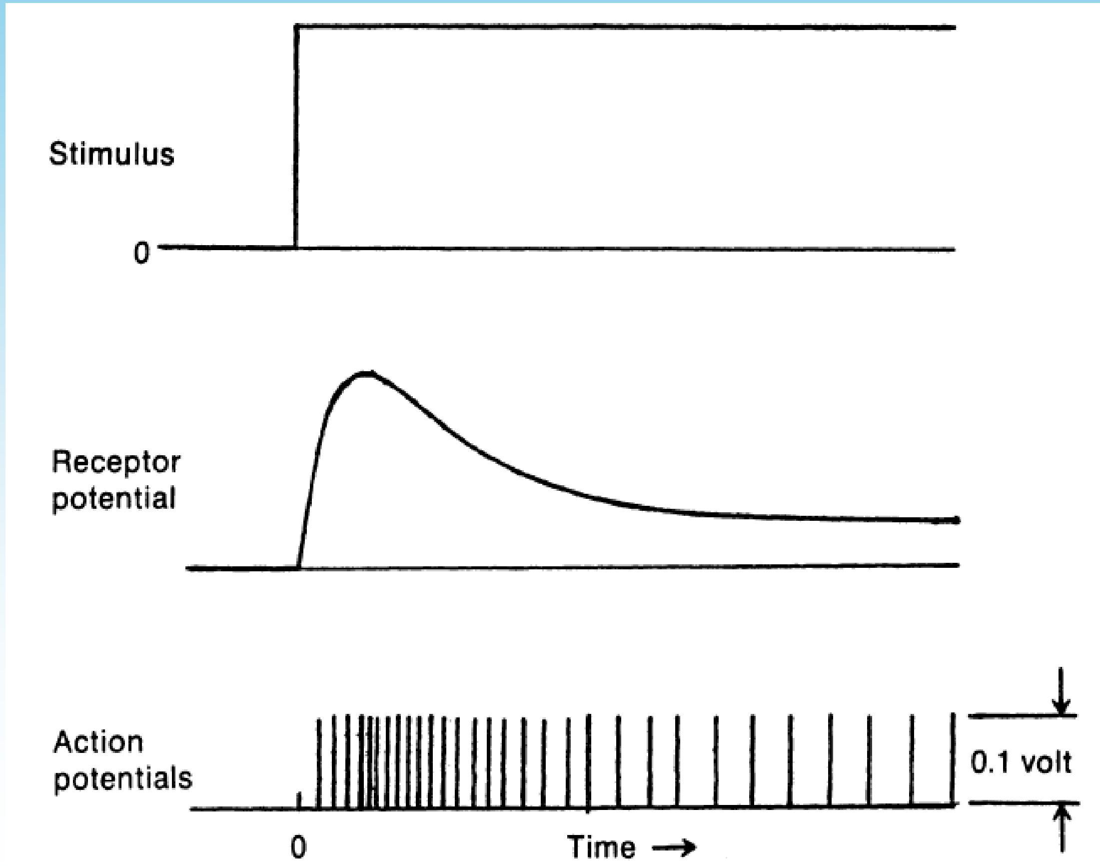

The Brain and Neurons

- Axons
  - An axon is a long, slender projection of a nerve cell, or neuron, that conducts electrical impulses away from the neuron's cell body

Axon of 9 day old mouse
with growth cone visible.

The Brain and Neurons

- Synapse
  - A junction that permits a neuron to pass an electrical or chemical signal to another cell (neural or otherwise)
  - Types
    - Chemical synapse
      - Transmitting and Receiving chemicals, called a neurotransmitter
    - Electrical synapse
      - Connected by channels that are capable of passing electrical current, causing voltage changes

The Brain and Neurons

- Synapse

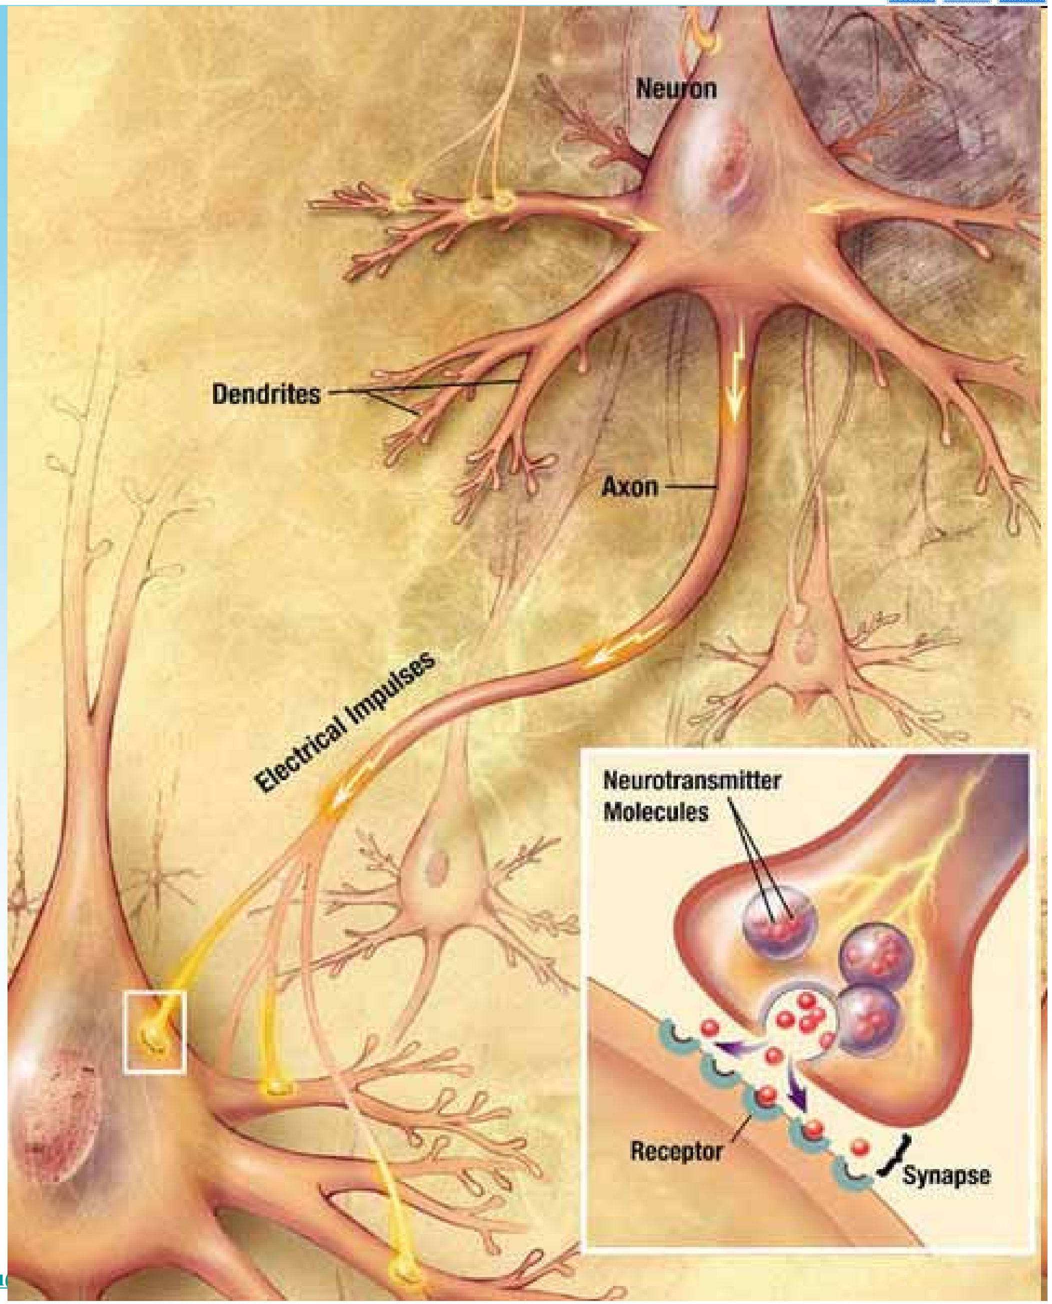

The Progress of AI

- **Artificial Intelligence**
  - **What is Intelligence?**
    - Deduction, reasoning, problem solving
    - Knowledge representation
    - Planning
    - Learning
    - Natural language processing
    - Motion and manipulation
    - Perception
    - Social intelligence
    - Creativity
    - General intelligence

The Progress of AI

- Latest Update
  - **“Watson” from IBM @ Jeopardy**
    - Natural language processing

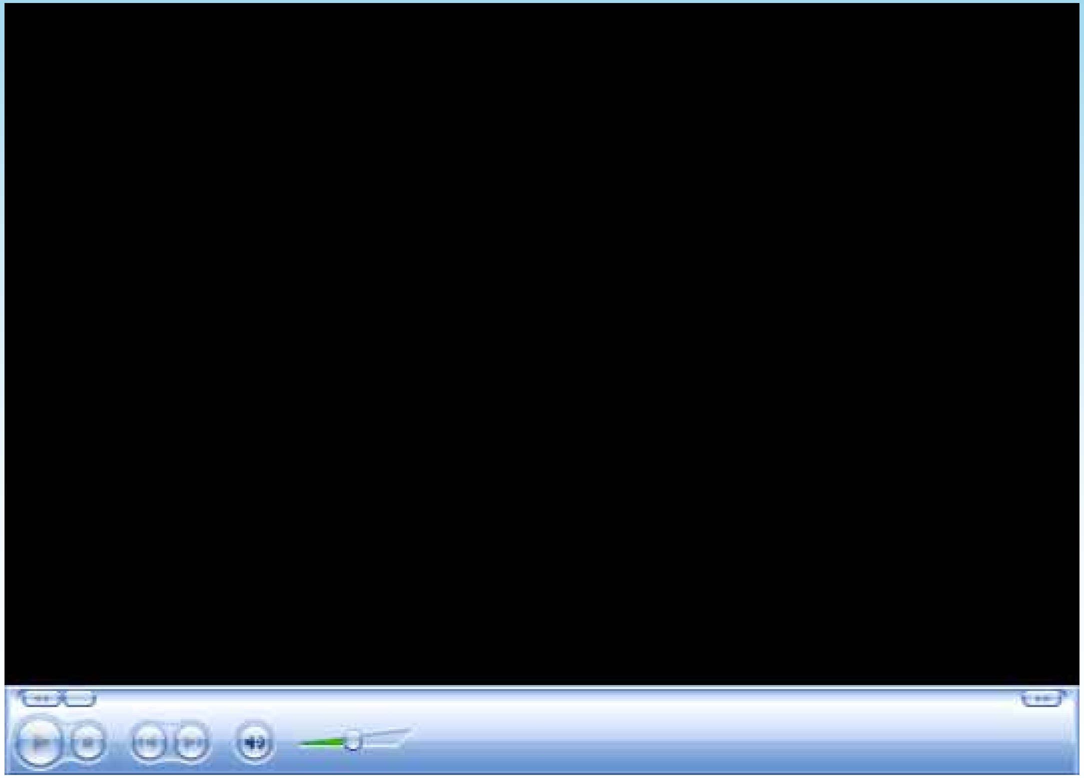

The Progress of AI

- Latest Update
  - **ASIMO from Honda**
    - **A**dvanced **S**tep in **I**nnovative **MO**bility

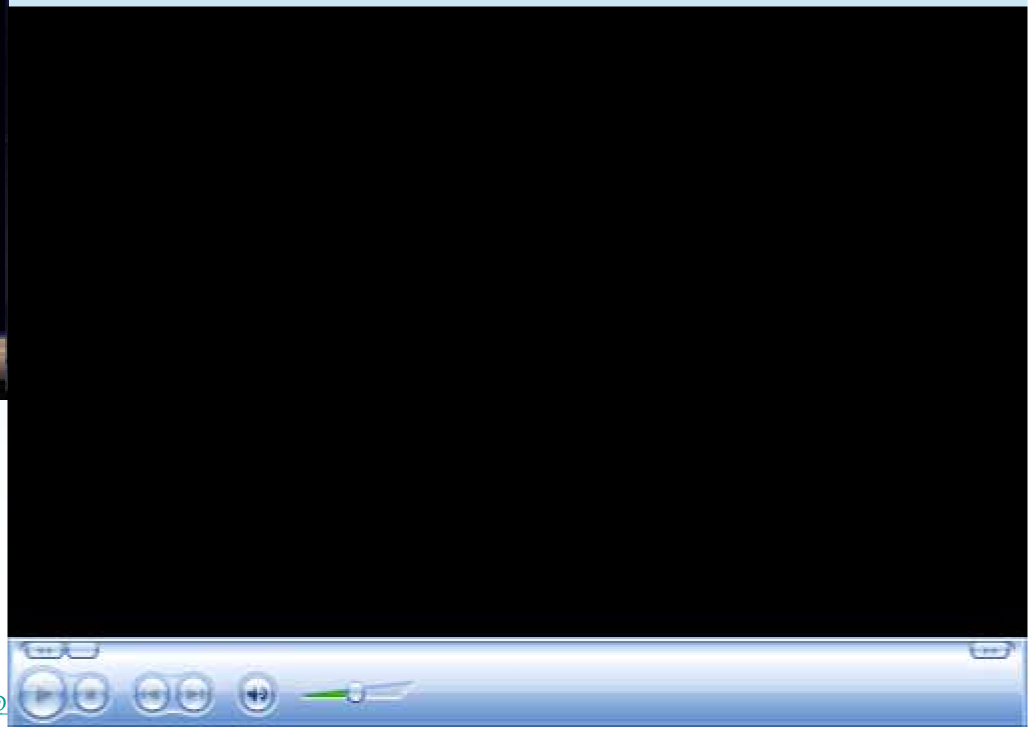

The Progress of AI

- **Competitions and prizes in artificial intelligence**
  - 1 General machine intelligence, $100,000
  - 2 Conversational behavior
  - 3 Pilotless aircraft, $30,000
  - 4 Driverless cars, $2,000,000
  - 5 Data-mining and prediction, $1,000,000
  - 6 Robot soccer
  - 7 Logic, reasoning and knowledge representation
  - 8 Games, $10,000

The Progress of AI

- **Competitions and prizes in artificial intelligence**
  - **International Aerial Robotics Competition**

The Progress of AI

- **Competitions and prizes in artificial intelligence**
  - **DARPA Grand Challenge**


The Progress of AI

- **Competitions and prizes in artificial intelligence**
  - The RoboCup and FIRA

DARPA Grand Challenge

- The authors’ work is supported and inspired by DARPA
  - The Defense Advanced Research Projects Agency (DARPA) is an agency of the United States Department of Defense responsible for the development of new technology for use by the military.

Formed	1958
Headquarters	Arlington, Virginia
Employees	240
Annual budget	$3.2 billion

DARPA Grand Challenge

- The Grand Challenge is to create autonomous vehicles that could drive themselves without human intervention.
  - Location: California's Mojave desert
  - Distance: 212 km
  - Prize: $2 million

DARPA Grand Challenge

- Winner in 2005
  - Stanley, Stanford's Volkswagen Touareg (大众途锐)
    - Led by Associate Professor Sebastian Thrun, director of the Stanford Artificial Intelligence Lab
    - Stanley is currently located at the Smithsonian National Museum of American History

DARPA Grand Challenge

- DARPA Urban Challenge in 2007
  - Location: George Air Force Base
  - Distance: 96 kilometres
- Winner
  - Tartan Racing, a collaborative effort by Carnegie Mellon University and General Motors Corporation
- Second
  - The second place finisher earning the $1 million prize was the Stanford Racing Team with their entry "Junior", a 2006 Volkswagen Passat.

DARPA Grand Challenge

- Winner in 2007

DARPA Grand Challenge

- DARPA Urban Challenge
  - Challenge
    - busy city streets
    - navigate basic traffic conditions according to California law
    - merging
    - passing
    - parking
    - negotiating intersections

DARPA Grand Challenge

- DARPA Urban Challenge
  - Results
    - 大多数参赛车装备了最先进的传感器、定位系统，多达14个的Blade Server(模块化设计的服务器)
    - 然而碰到面包大小的障碍物就直接歇菜
    - 数小时后，将近一半的车队都退出了比赛，例如因为停车场超速，与其他车共用单行道时擦挂

IBM HS20 blade server. Two bays for 2.5" (6.4 cm) SCSI hard drives appear in the upper left area of the image.

The difference

- The difference between MoNETA and the state-of-the-art AI
  - State-of-the-art AI
    - Even the most helpful AI must be programmed explicitly to carry out its one specific task.
  - MoNETA
    - A general-purpose intelligence that can be set loose on any problem
    - One that can adapt to a new environment without having to be retrained constantly
    - One that can tease the single significant morsel out of a gluttonous banquet of information

The difference

- The difference between MoNETA and the winners of DARPA Challenge
  - MoNETA-enabled military scout vehicle
    - It will be able to go into a mission with partially known objectives that change suddenly.
    - It will be able to negotiate unfamiliar terrain,
    - Recognize a pattern that indicates hostile activity
    - make a new plan
    - hightail it out of the hostile area
    - If the road is blocked, it will be able to make a spur-of-the-moment decision and go off-road to get home.

Brain vs. Computer

Brain

Computer

Brain vs. Computer

- The Humble rat can beat any state-of-the-art AI
  - a hungry rat will explore creatively for food
  - follow familiar, memorized safe routes
  - integrate signals from different senses
  - recognize never-seen dangerous objects such as a mousetrap and will avoid them
  - disengage its current plan and switch to its next priority
  - 2 gram brain vs. AI with a vehicle of hardware
  - tiny-bulb level power consumption

Brain vs. Computer

- Computer

Characteristic

1) *the memory and processor are separated by a data channel, or a bus, physically separated — a significant physical distance separates the areas where the data is stored from the areas where it is manipulated.*
2) *That channel's fixed capacity means that only limited amounts of data can be "checked out" and worked on at any given instant..*
3) *To minimize the amount of traffic flowing on the fixed-capacity bus, the registers with a cache memory are inserted.*

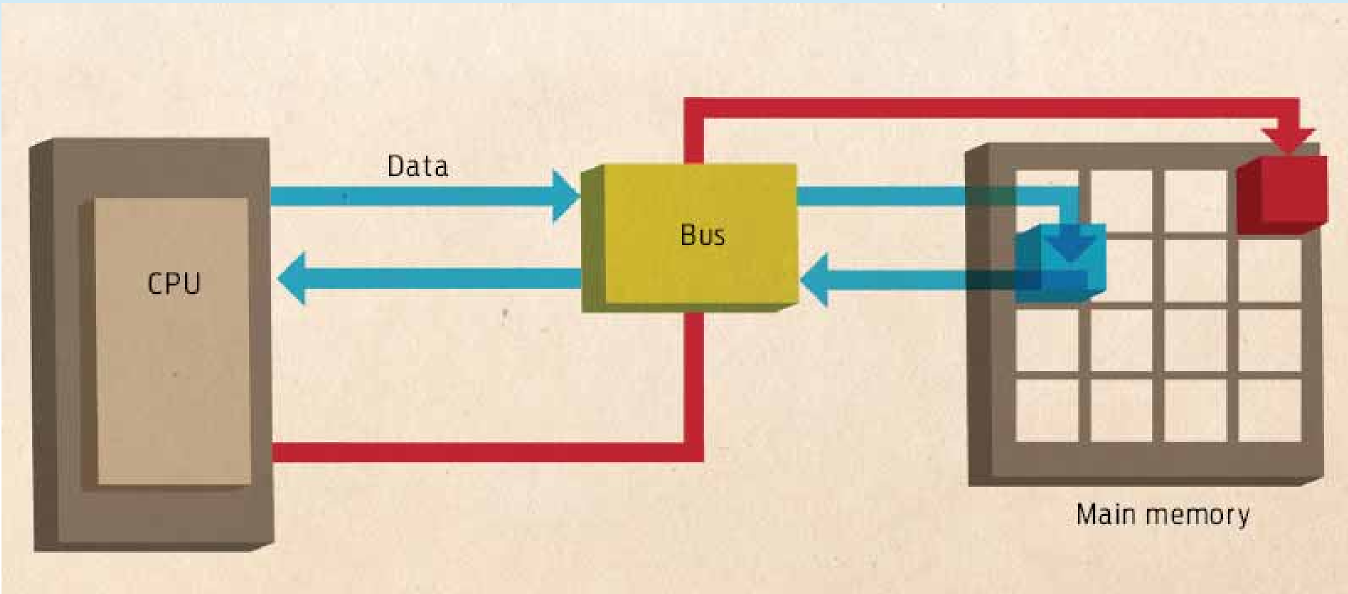

Brain vs. Computer

- Brain

Characteristic

1) *all processing has already taken place during the information transfer*
2) *There is no need for the brain to take information out of one neuron, spend time processing it, and then return it to a different set of neurons.*
3) *in the mammalian brain, storage and processing happen at the same time and in the same place.*

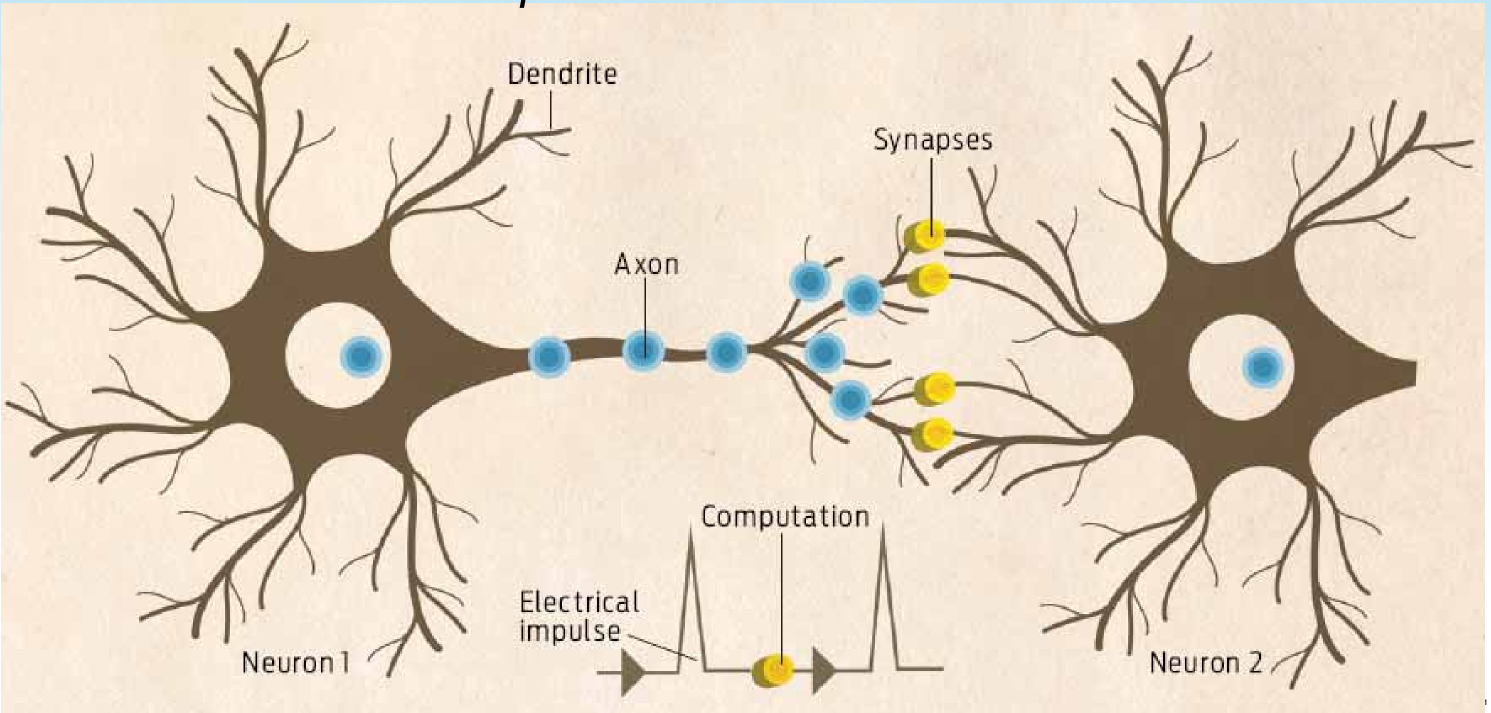

Brain vs. Computer

- Computational neuroscience based AI
  - the study of brain function in terms of the information processing properties of the structures that make up the nervous system.
- The major difference is about the Architecture
  - Computational neuroscience has focused largely on building software that can simulate or replicate a mammal's brain in the classic von Neumann computer architecture.

Brain vs. Computer

- Brain Simulation
  - Step 1: Synapse
    - Store the synapse's state
    - Change the state requires data bus throughput
    - Time consumption on data R/W
  - Step 2: Neuron
    - Multiply that synapse sequence by up to 8000 synapses (a single rat neuron)
  - Step 3: Brain
    - Multiply the neuron by billions
  - Results: an entire millisecond of brain activity

Brain vs. Computer

- Power Consumption
  - 20-watt light bulb vs. a dedicated power plant
- Operation Voltage
  - 100 mV vs. 1 V
- Replicating the brain structure described above is not totally impossible with today's silicon technology, but it would be fantastically inefficient.

Brain vs. Computer

- The missing key component of the brain on a chip
  - A component can process the information while transferring it.
- Memristor is the missing link

Memristors -- the missing link

- The missing link

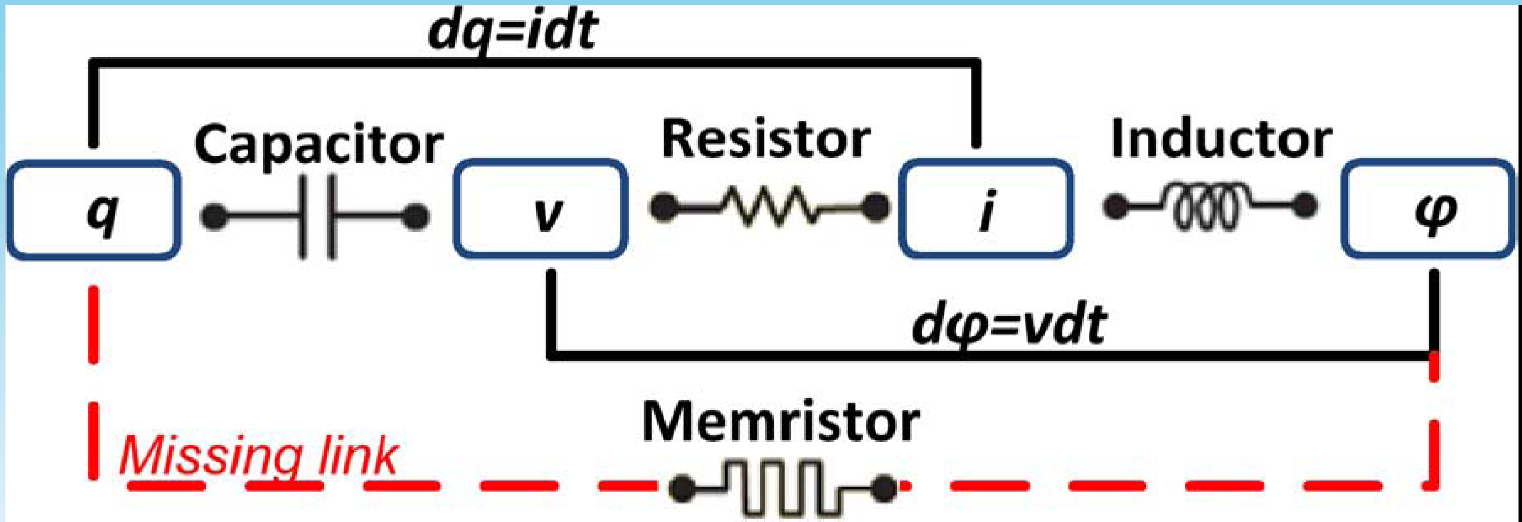

Fig. 1. Four fundamental circuit elements: Resistance ($dv = R \cdot di$), capacitance ($dq = C \cdot dv$), inductance ($d\varphi = L \cdot di$), and memristance ($d\varphi = M \cdot dq$) which is the missing link that Chua argued.

Memristor theory was formulated and named by Leon Chua, UC Berkeley, in a 1971 paper

Memristors -- the missing link

- Memristor

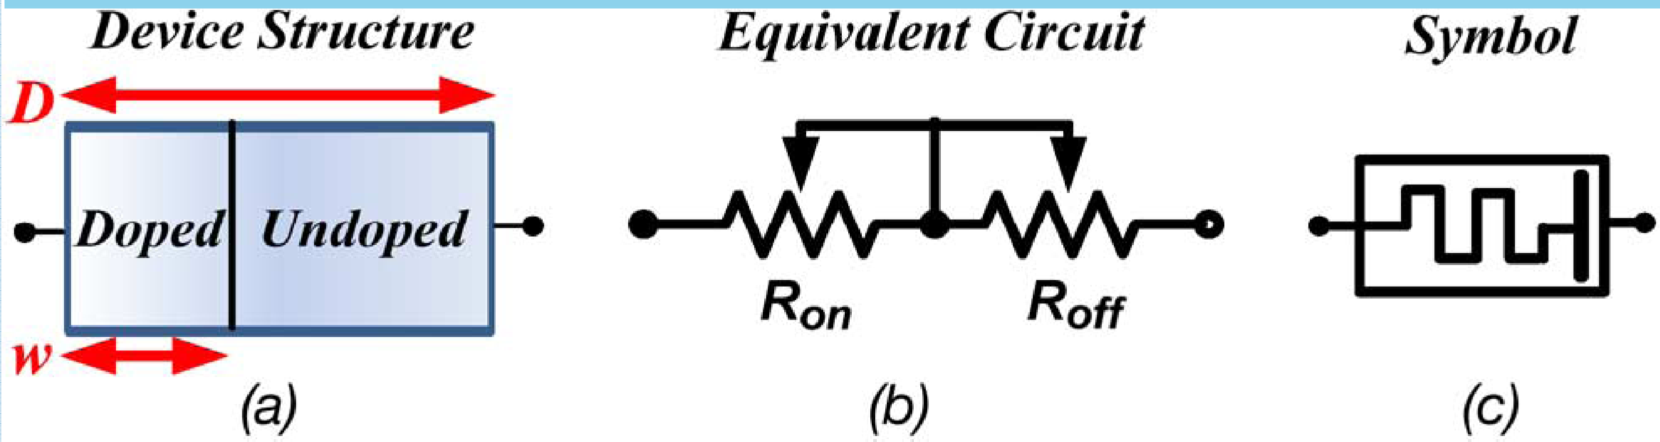

$$\left.\begin{array}{l} M(q) = d\varphi/dq. \\ d\varphi = v\,dt \\ dq = i\,dt \end{array}\right\} \quad v = M(q)\cdot i \longrightarrow v = R(w)\cdot i.$$

The memristance M(q) is equal to voltage over current which is also known as the resistance in the linear case. Therefore, memristance has the same unit (ohm) as resistance, and is logically a charge-controlled resistance.

Memristors -- the missing link

- Memristor
  - A memristor would behave like a resistor with a conductance that changed as a function of its internal state and the voltage applied.
  - A memristor could remember how much current had gone through it, it could work as an essentially nonvolatile memory.
  - Whatever its past state, or resistance, it freezes that state until another voltage is applied to change it. Maintaining that state requires no power.
  - Their biggest potential all along has been as a realistic analogue to synapses in brains.

Memristors -- the missing link

- Birth
  - In 2008, a team at HP Labs announced the development of a switching memristor based on a thin film of titanium dioxide.

R. Stanley Williams

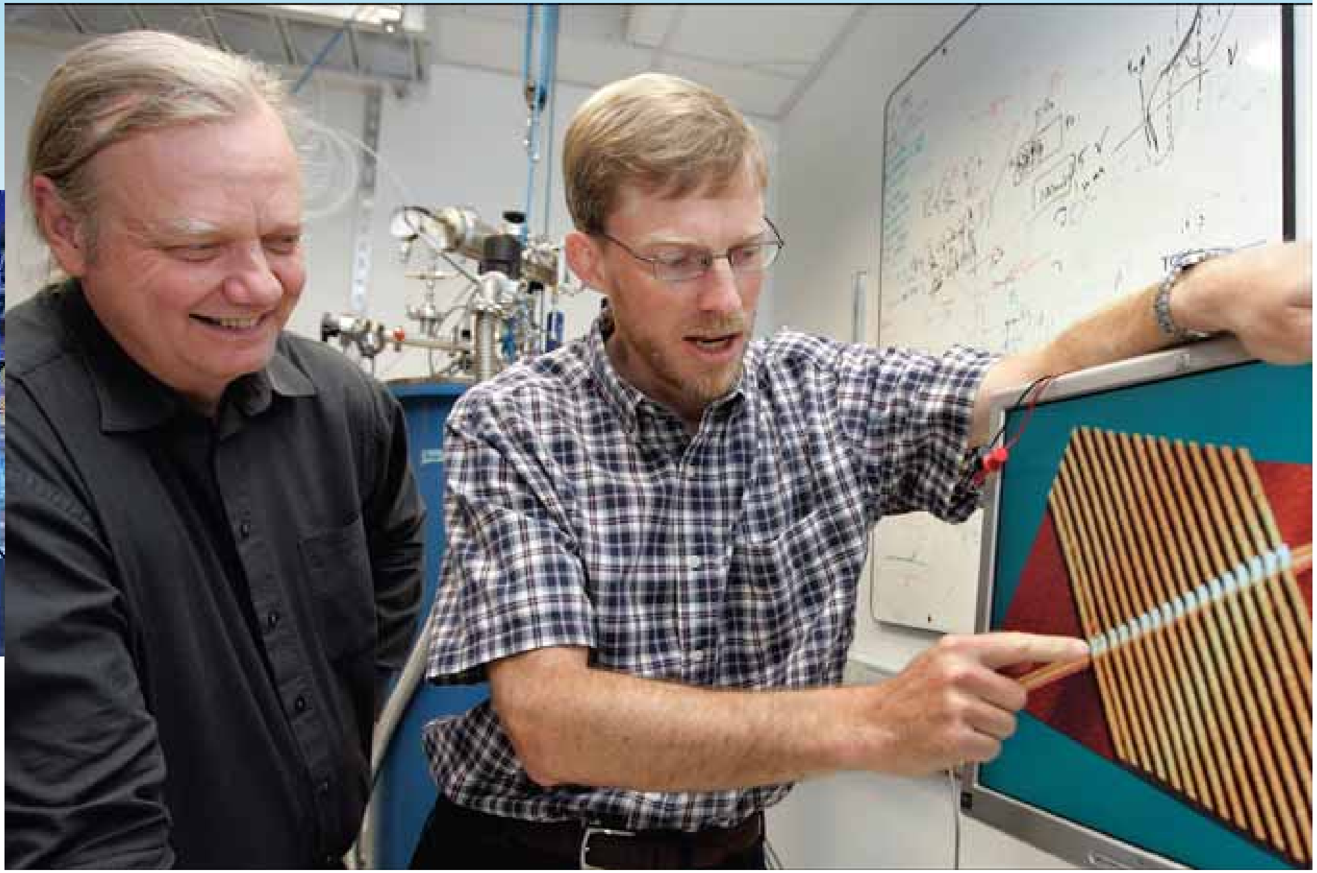

Memristors -- the missing link

- CROSSBAR ARCHITECTURE

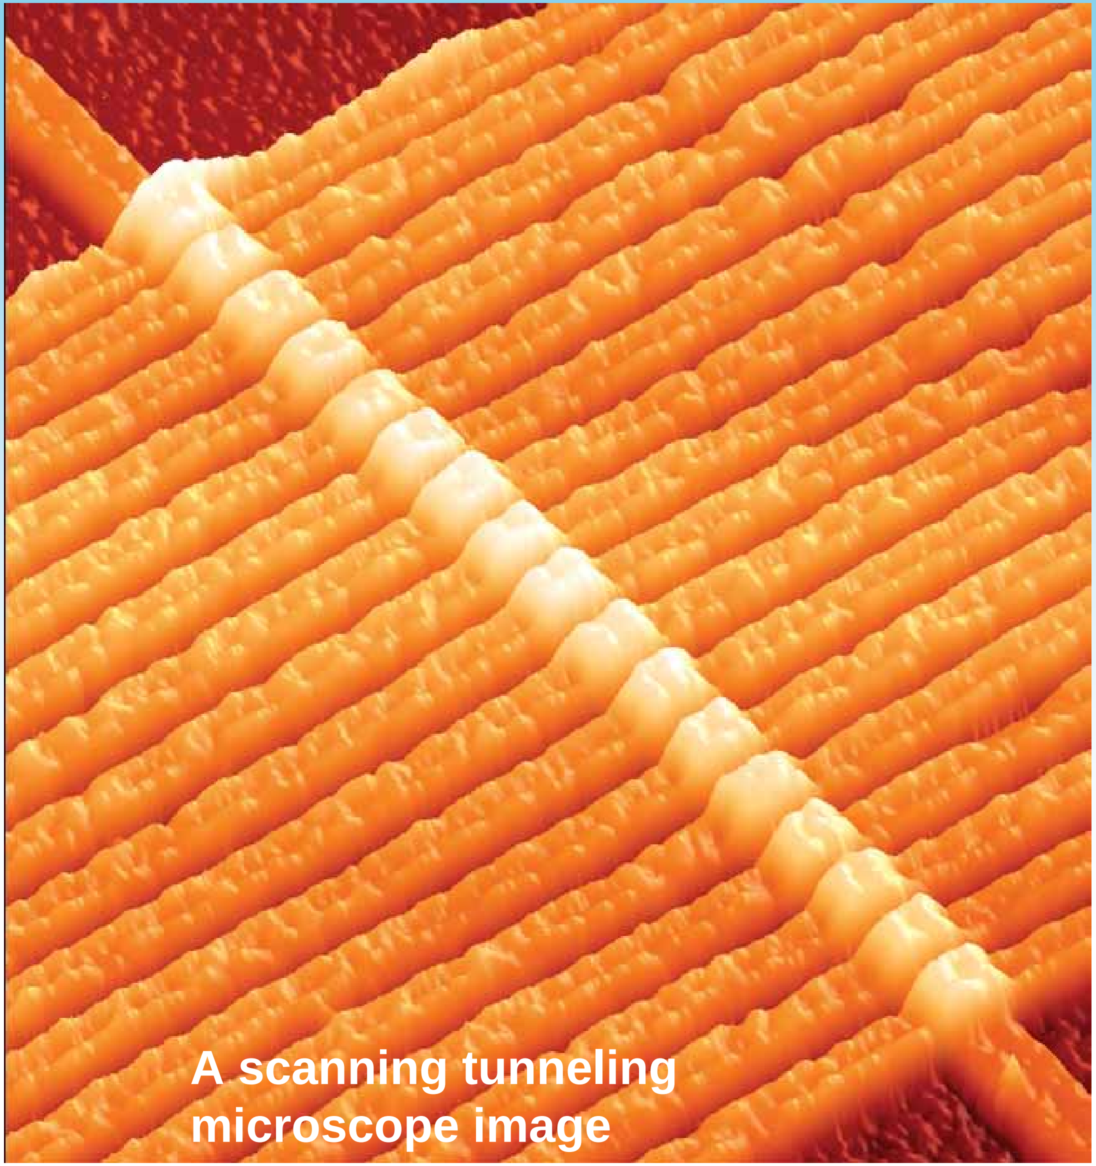

A scanning tunneling microscope image

Memristors -- the missing link

- How Memristance Works

The crossbar architecture is a fully connected mesh of perpendicular wires. Any two crossing wires are connected by a switch. To close the switch, a positive voltage is applied across the two wires to be connected. To open the switch, the voltage is reversed.

A switch is a 40-nanometer cube of titanium Dioxide (TiO_2) in two layers: The lower TiO_2 layer has a perfect 2:1 oxygen-to-titanium ratio, making it an insulator. By contrast, the upper TiO_2 layer is missing 0.5 percent of its oxygen (TiO_{2-x}), so x is about 0.05. The vacancies make the TiO_{2-x} material metallic and conductive.

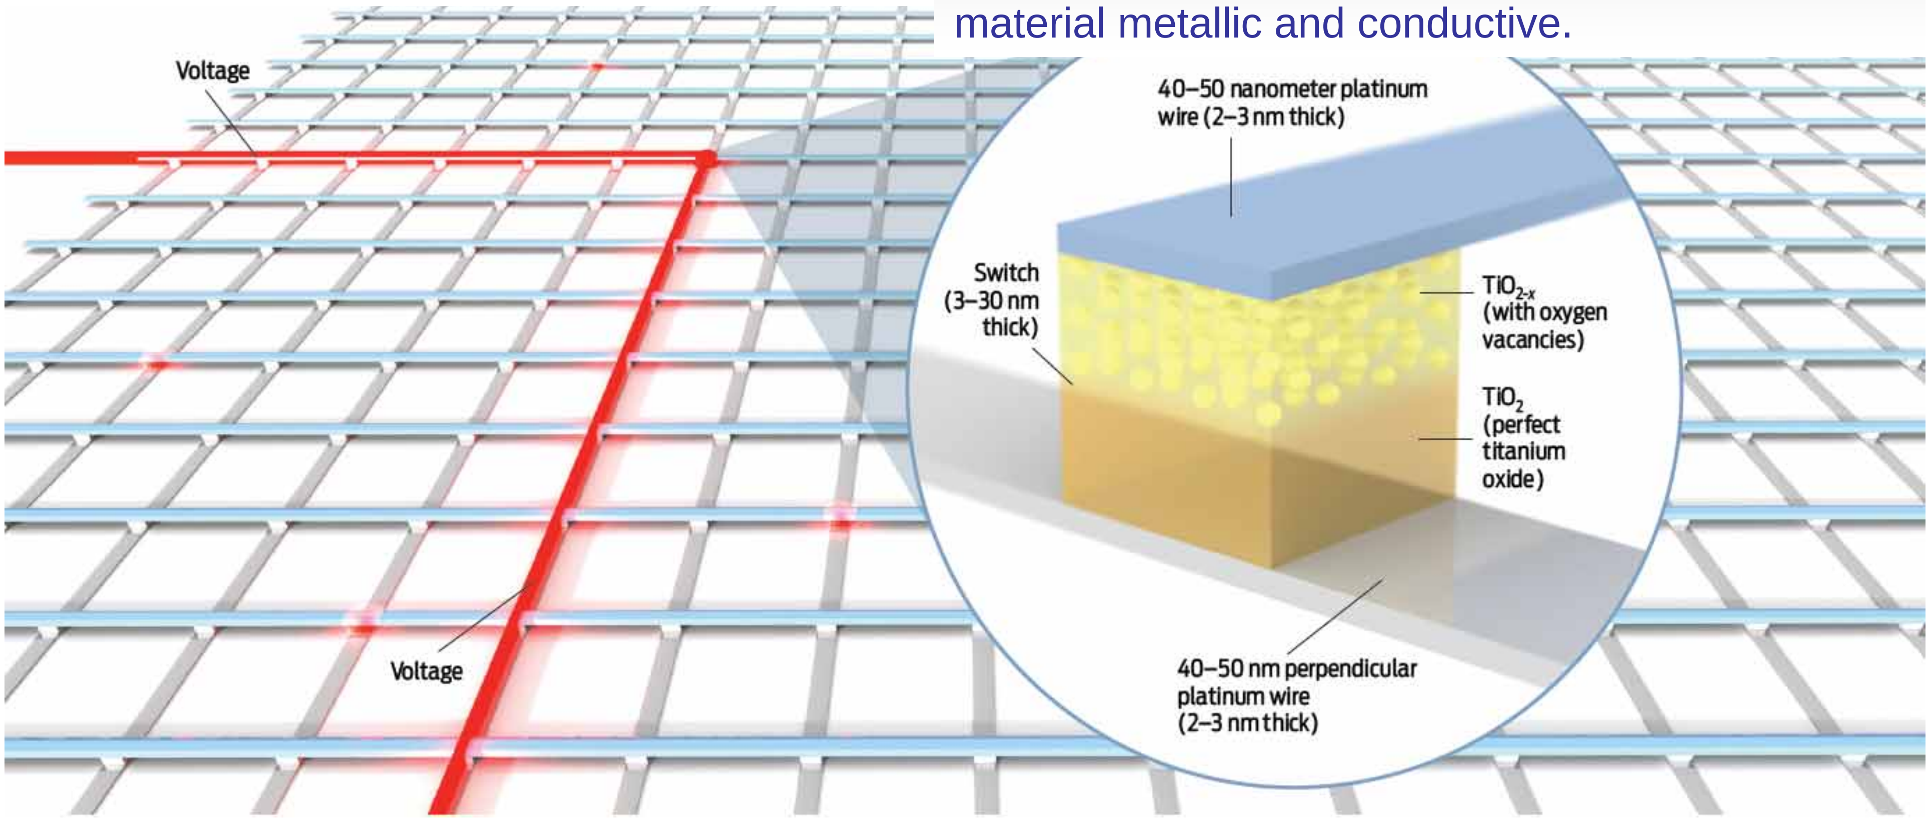

Memristors -- the missing link

- How Memristance Works

APPLIED MEMRISTANCE: The oxygen deficiencies in the TiO_{2-x} manifest as “bubbles” of oxygen vacancies scattered throughout the upper layer. A positive voltage on the switch repels the (positive) oxygen deficiencies in the metallic upper TiO_{2-x} layer, sending them into the insulating TiO_2 layer below. That causes the boundary between the two materials to move down, increasing the percentage of conducting TiO_{2-x} and thus the conductivity of the entire switch. The more positive voltage is applied, the more conductive the cube becomes.

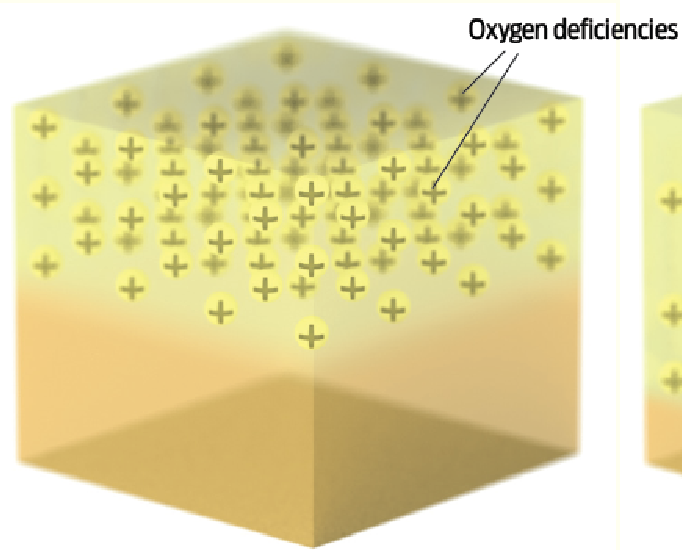

A negative voltage on the switch attracts the positively charged oxygen bubbles, pulling them out of the TiO_2. The amount of insulating, resistive TiO_2 increases, thereby making the switch as a whole resistive. The more negative voltage is applied, the less conductive the cube becomes.

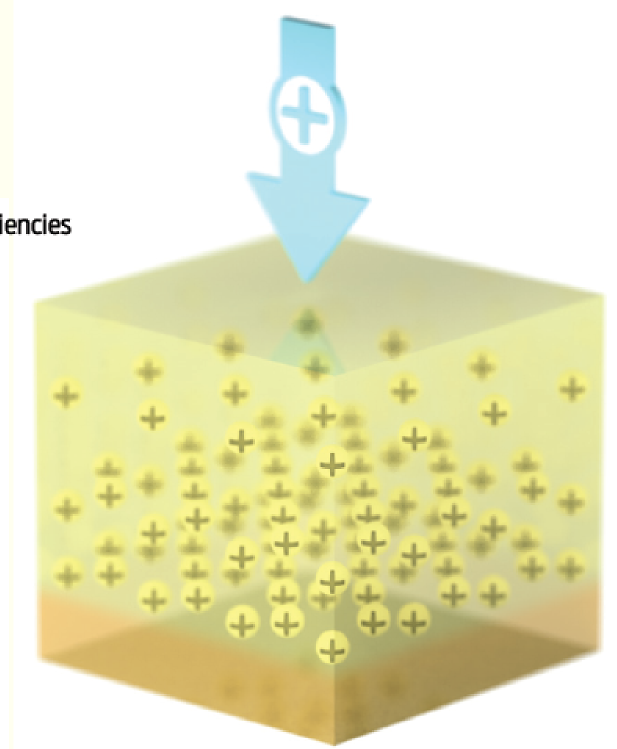

What makes this switch special—memristive—is that when the voltage is turned off, positive or negative, the oxygen bubbles do not migrate. They stay where they are, which means that the boundary between the two titanium dioxide layers is frozen. That is how the memristor “remembers” how much voltage was last applied.

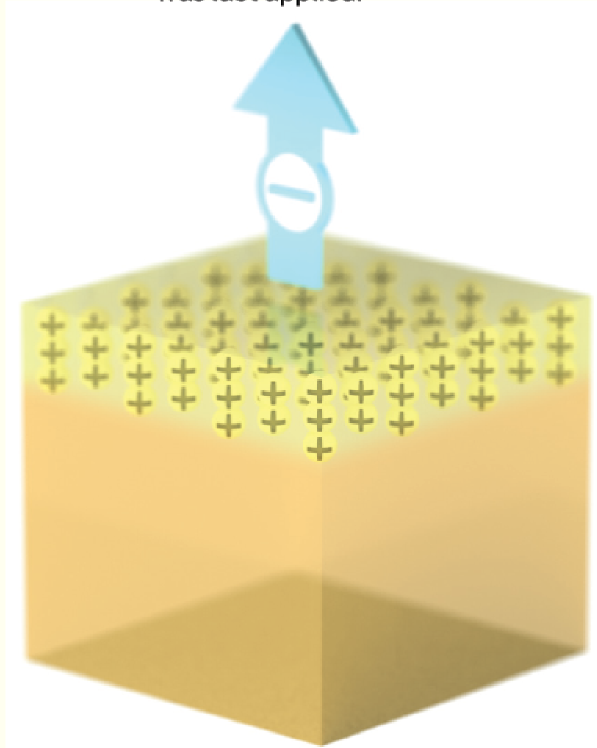

Memristors -- the missing link

- Advantages
  - Dense, Cheap, and Tiny
  - Memristors bring data close to computation, the way biological systems do.
  - They use very little power to store that information, just as the brain does.
  - Simultaneously move and manipulate data
  - Drastically cutting power consumption and space

Memristors -- the missing link

- Disadvantages
  - They also have a high failure rate at present
    - Must bear an intriguing resemblance to the brain’s synapses.
      - Brains gracefully degrade their performance as synapses are lost, without sudden system failure.

To Build a Brain

- Change your architecture to merge memory and computation
- Each silicon core is directly connected to its own immediately accessible megacache made up of millions of memristors, meaning that every single core has its own private massive bank of memory.

To Build a Brain

- Software – the brain models that will populate the hardware
  - biological algorithms will create this entity: MoNETA.
  - Org: Center of Excellence for Learning in Education, Science, and Technology (CELEST), Boston University
  - Computational modelers, neuroscientists, psychologists, and engineers collaborate with researchers from Harvard, MIT, Brandeis, and BU.

To Build a Brain

- Operating System
  - Cog Ex Machina. Cog
  - built by HP principal investigator Greg Snider
  - To allow the brain models and the neuromorphic hardware to interact
    - Neuromorphic computation means computation that can be divided up between hardware that processes like the body of a neuron and hardware that processes the way dendrites and axons do.

To Build a Brain

- Operating System

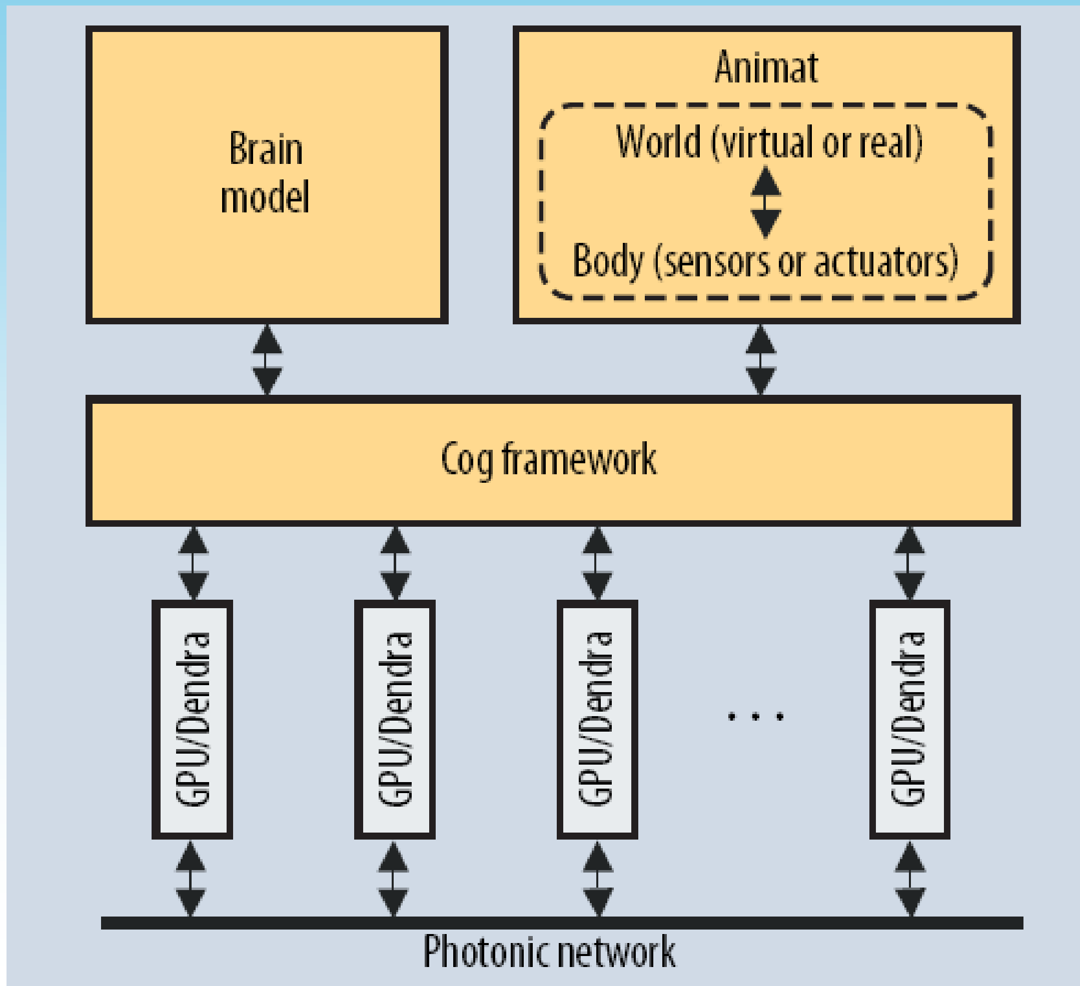

To Build a Brain

- Learning Scheme

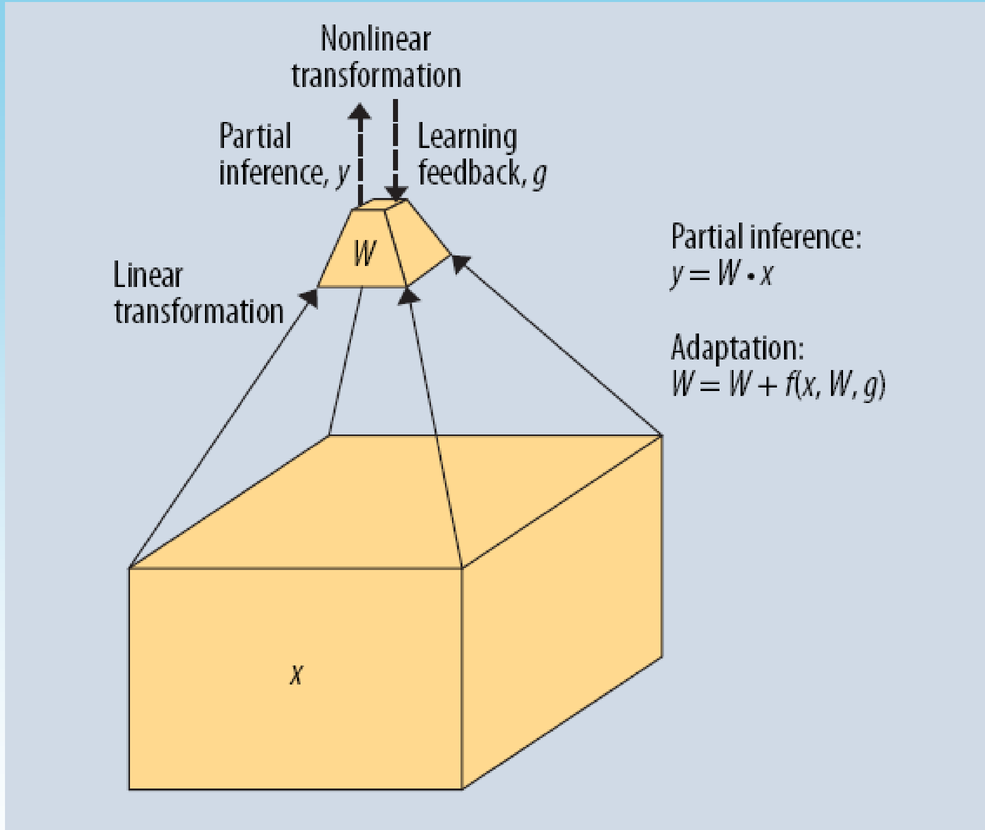

Figure 3. How learning occurs. Linear and nonlinear transformations cooperate to implement the learning that corresponds to long-term memory. Cog combines a linear transformation's current state, *W*, with the input, *x*, to produce *y*, a partial inference. It then combines *g*, *x*, and *W* to implement learning. The learning function, *f*, determines the type of learning that occurs.

To Build a Brain

- Learning Scheme

(a) Standard photographic compression

(b) Cog recaptures the details

To Build a Brain

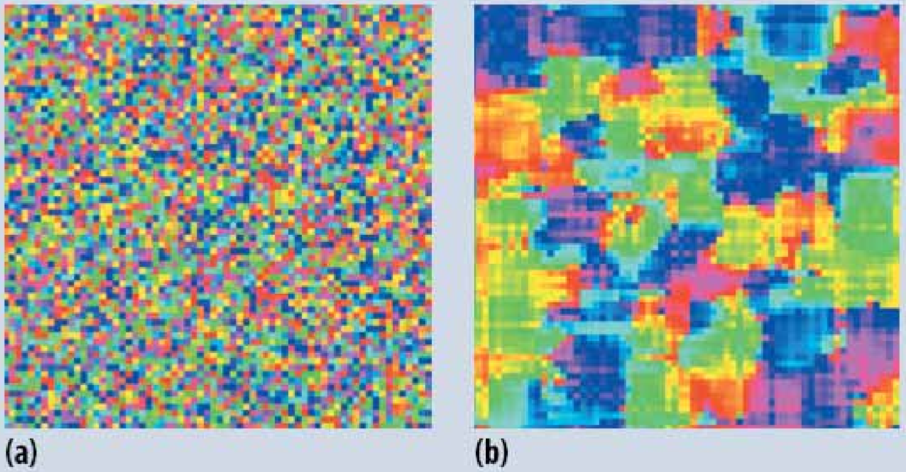

Figure 7. Topographic map of orientation selectivity. (a) Before learning, there are no identifiable clusters of orientation selectivity. (b) Clusters emerge after learning, with different colors denoting particular orientations.

- Learning Scheme

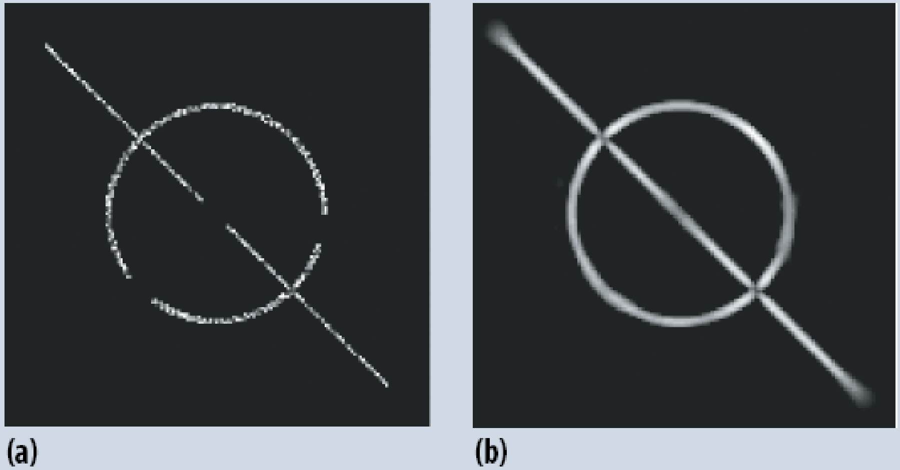

Figure 9. Simple boundary completion. (a) A noisy figure with a broken line and broken circle becomes (b) smooth and filled in using a boundary completion algorithm implemented with Cog.

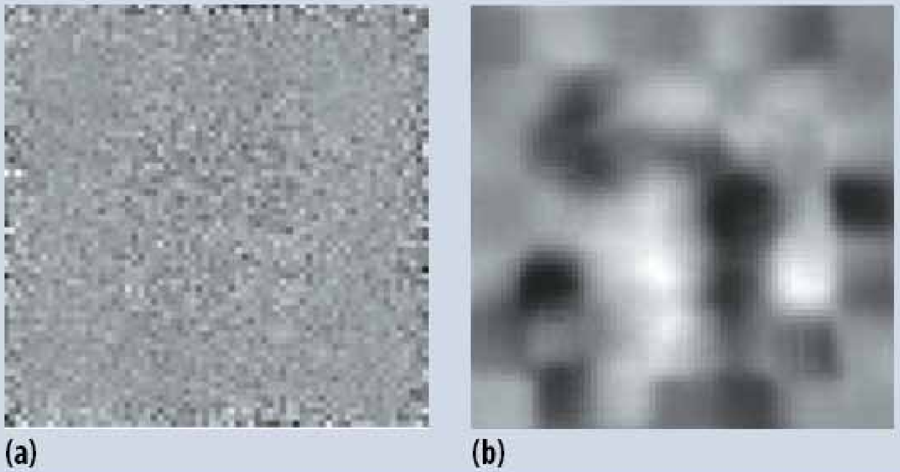

Figure 8. Topographic map of ocular dominance. (a) Before learning, there are no identifiable columns of ocular dominance. (b) Columns emerge after learning, with white denoting left eye, black denoting right eye, and gray denoting intermediate.

To Build a Brain

- Hardware Cores
  - “Neuron-type” CPU
    - flexible, letting it handle any operation you throw at it.
    - resemble those of the neuron
    - sucks up a lot of power,
    - make up only a small percentage of the system

To Build a Brain

- Hardware Cores
  - "Dendritic" core
    - optimized for only a specific kind of computation
    - the complicated linear algebra operations that approximate what happens inside a dendrite.
    - Approximate the distributed computation that dendrites carry out.
    - store extraordinary amounts of state information in their massive memristor-based memory banks
    - Memristors, finally, will act as the synapses that mediate the information transfer between the dendrites and axons of different neurons.

游飞 Fei You, feiyou@uestc.edu.cn http://staff.uestc.edu.cn/youfei

Future MoNETA

- MoNETA will be a general-purpose mammaliantype intelligence, an artificial, generic creature known as an animat.
- The key distinguishing feature: it won’t have to be explicitly programmed.
- as adaptable and efficient as a mammal’s brain.
- it will learn dynamically.

Future MoNETA

- Truly general-purpose intelligence can emerge only when everything happens all at once:
  - all perception (including auditory and visual inputs, or the brain areas responsible for the generation of fine finger movements), emotion, actions, and reactions combine and interact to guide behavior.
- Practical limitations:
  - lack of a unified theory of the brain

Future MoNETA

- MoNETA should learn throughout its lifetime without needing constant reprogramming or needing to be told a priori what is good for it, and what is bad.
- The animat will learn about objects in its environment, navigate to reach its goals, and avoid dangers

Future MoNETA

- Test
  - We will test our animat in a classic trial called the Morris water navigation task.
    - neuroscientists teach a rat to swim through a water maze
    - synchronize vision, touch, spatial navigation, emotions, intentions, planning, and motor commands.
    - If we can train the animat to negotiate this maze, we'll be confident that we have taken an important first step toward simulating a mammalian intelligence.

Future MoNETA

- Test
  - Survival of the fittest
    - working with thousands of candidate animats at once, all with slight variations in their brain architectures.
    - cull the best ones from the bunch and keep tweaking them
  - Real world
    - Once the memristor-based neural processing chip is manufactured, we will build it into robotic platforms that venture into the real world.

Future MoNETA

- Goals
  - to make machines that behave as if they are intelligent, emotionally biased, and motivated, *without* the constraint that they are actually aware of these feelings, thoughts, and motivations.

Future MoNETA

- Prediction
  - Neuromorphic chips will eventually come in as many flavors as there are brain designs in nature: fruit fly, earthworm, rat, and human. All our chips will have brains.

Conclusion

- 小器件也有大智慧