Introductory Course in Neuroscience

Neuromorphic Engineering I

'Organizing principles' in neural and neuromorphic electronic systems

Part 1: Motivation, history, community
Part 2: Organizing principles
Part 3: The physiologist's friend chip

Tobi Delbruck
Inst. of Neuroinformatics
www.ini.unizh.ch/~tobi

Your neuroscience exam question will based on this lecture and the following reading

1. Mead, **Neuromorphic Electronic Systems**, Proc. IEEE, 1990
2. Delbruck & Liu, **A silicon visual system as a model animal**, Vision Research, 2004
3. Boahen, **Mimic the Nervous System with Neuromorphic Chips**, Scientific American, 2005

You can get these slides and papers via the ZNZ Neuroscience Course web page.

Part 1: Motivation for neuromorphic engineering, history, community

Natural computation

Flies acrobatically
Recognizes patterns
Navigates
Forages
Communicates

10^{-15} J/op

Digital silicon 10^{-7} to 10^{-11} J/op

10^8 to 10^4 times as efficient as digital silicon

The Cat is Out of the Bag: Cortical Simulations with 10^9 Neurons, 10^{13} Synapses

Rajagopal Ananthanarayanan[1], Steven K. Esser[1]
Horst D. Simon[2], and Dharmendra S. Modha[1]

[1]IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120
[2]Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720
{ananthr,sesser}@us.ibm.com, hdsimon@lbl.gov, dmodha@us.ibm.com

Response from Henry Markram, Blue Brain Project

Dear Bernie,

You told me you would string this guy up by the toes the last time Mohda made his stupid statement about simulating the mouse's brain.

I thought that having gone through Blue Brain so carefully, journalists would be able to recognize that what IBM reported is a scam - no where near a cat-scale brain simulation, but somehow they are totally deceived by these incredible statements.

I am absolutely shocked at this announcement. Not because it is any kind of technical feat, but because of the mass deception of the public.

1. These are point neurons (missing 99.999% of the brain; no branches; no detailed ion channels; the simplest possible equation you can imagine to simulate a neuron, totally trivial synapses; and using the STDP learning rule I discovered in this way is also is a joke).

2. All these kinds of simulations are trivial and have been around for decades - simply called artificial neural network (ANN) simulations. We even stooped to doing these kinds of simulations as bench mark tests 4 years ago with 10's of millions of such points before we bought the Blue Gene/L. If we (or anyone else) wanted to we could easily do this for a billion "points", but we would certainly not call it a cat-scale simulation. It

Computer vs. Brain

Anderson et al. 2003

Pentium 4

Cortex

1mm

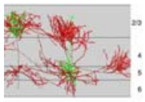

At the system level, brains are about 1 million times more power efficient than computers. Why?

Cost of elementary operation (turning on transistor or synapse) is about the same.
It's not some magic about physics.

Computer	Brain
Fast global clock	Self-timed
Bit-perfect deterministic logical state	Synapses are stochastic! Computation dances: digital→analog→digital
Memory distant to computation	Memory at computation
Fast high precision power hungry ADCs	Low resolution adaptive data-driven quantization
Devices frozen on fabrication	Constant adaptation and self-modification

Technology development has enabled the conventional approach to computation

1947
1 transistor

1997
10^9 transistors

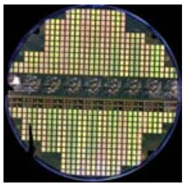

1. Moore's law: Number of transistors per chip doubles every 1.5 to 2 years
2. Cost/bit drops 29%/year
3. **True for last 45 years! Will continue at least another ~15y.**

How industry thinks of analog design

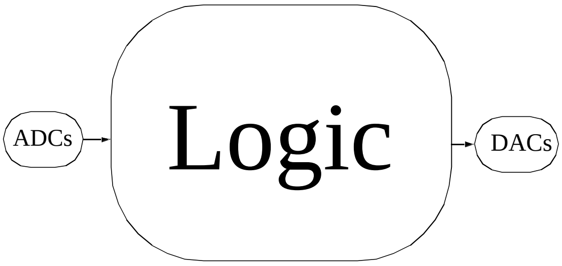

Synchronous logic

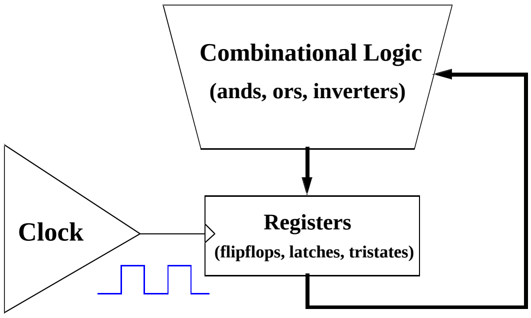

How to design logic (today)

```
architecture example_arch of example is
signal CountInternal: unsigned(7 downto 0);
attribute sync_set_reset of reset: signal is "true";
begin

        process(clock)
        begin
                if rising_edge(clock) then
                        if reset='1' then
                                CountInternal<=to_unsigned(0,8);
                        elsif CountInternal=to_unsigned(13,8) then
                                CountInternal<=to_unsigned(0,8);
                        else
                                CountInternal<=CountInternal+1;
                        end if;
                end if;
        end process;
        count<=CountInternal;
end example_arch;
```

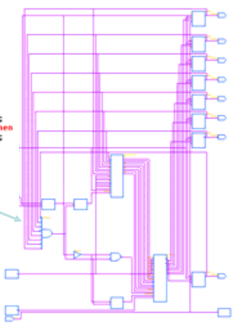

styles of processing

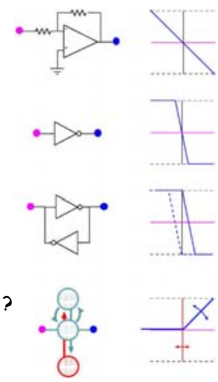

analog	constant **precise** analog gain by negative feedback
digital	**binary** signal restoration by high feedforward gain
digital	multi-stability by positive feedback & limits
cortex ?	variable analog gain, analog signal restoration, & multi-stability

Douglas 2001

Brief history of neuromorphic engineering

Physics of Computation Course

1982
1985

Carver Mead

Dick Feynman

John Hopfield

The World of Neuromorphic Engineering
2007

Shih-Chii Liu

Types of neuromorphic systems

- **Silicon retinas**—electronic models of retinas
- **Silicon cochleas**—electronic models of cochleas
- **Smart sensors** (e.g. tracking chips, motion sensors, presence sensors, auditory classification and localization sensors)
- **Networks of spiking neurons** – with self-modifying adaptive synapses
- **Central pattern generators** – for locomotion or rhythmic behavior
- **Models of specific systems:** e.g. bat sonar echolocation, lamprey spinal cord for swimming, lobster stomatogastric ganglion, electric fish lateral line
- **Multi-chip systems** that use the *address-event representation* (spikes) for inter-chip communication

---

Accomplishments of neuromorphic engineering

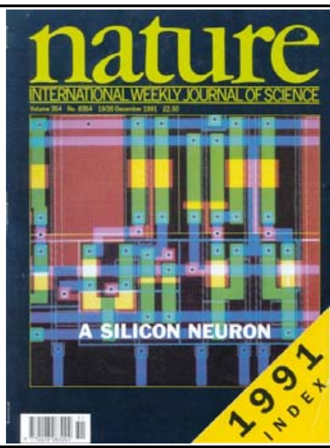

Mahowald, Douglas, 1991

Mahowald, Mead, 1991

Boahen, 2005

Companies that sprang from neuromorphic engineering

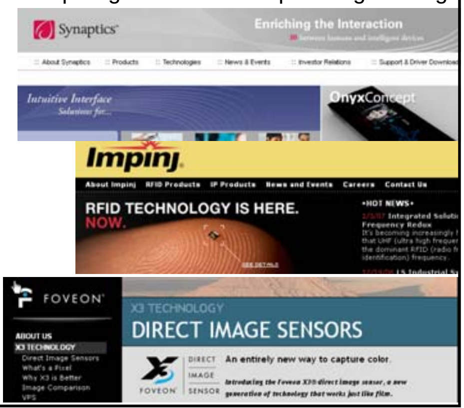

The Telluride Neuromorphic Engineering Workshop

- Focus is on
  - tutorials, hands-on workgroups
  - fostering the neuromorphic community
  - establishing long-lasting collaborations
- Running 12 years now, started by Rodney Douglas, Misha Mahowald, Terry Sejnowski, and Christof Koch.
- Funded by NSF & others, steadily at about $110k/yr.
- 60 people each year, about half invited and half applicants – **you can apply. Housing and part of travel is covered.**
- 3 weeks long each July, in the mountains in Colorado, USA.

Google "Telluride Neuromorphic" for more information

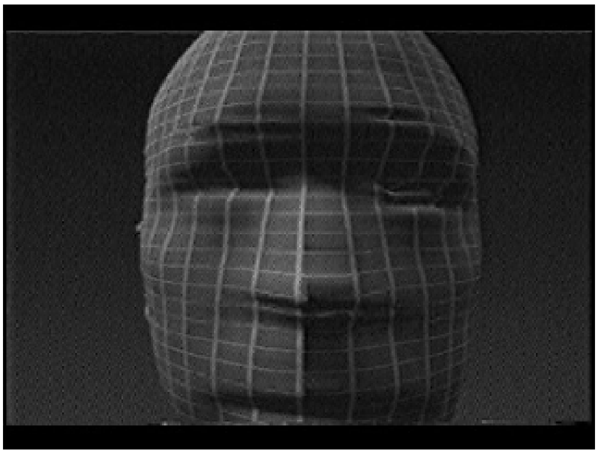

Part 2:
What are "organizing principles" as applied in neuromorphic engineering?

The fact that we can build devices that implement the same basic operations as those the nervous system uses leads to the inevitable conclusion that we should be able to build entire systems based on the organizing principles used by the nervous system.

Mead, 1990

Complementary devices, amplification

(Example #1)

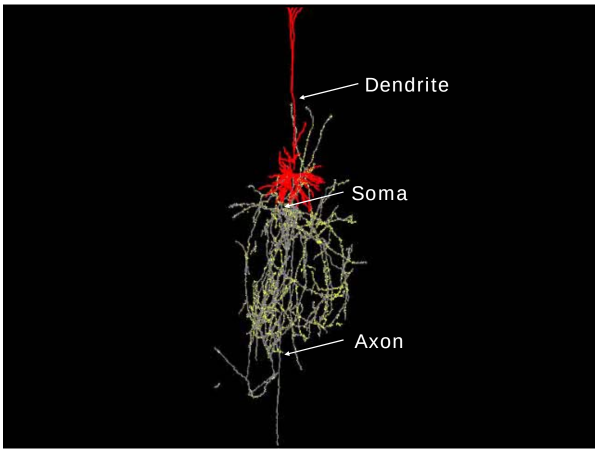

The membrane voltage is controlled by complementary channels

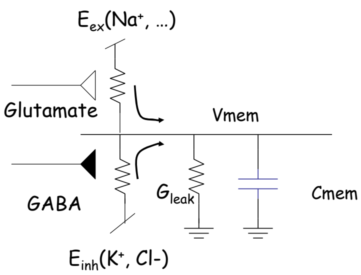

Almost no power is burned when both channels are off!

CMOS (complementary metal oxide semiconductor)

A CMOS inverting amplifier

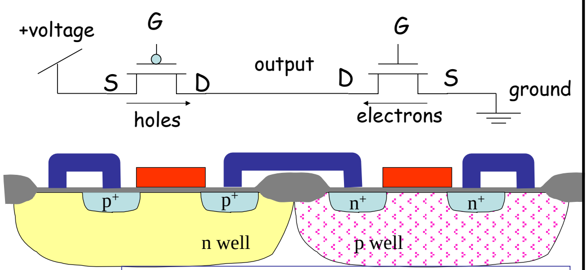

Organizing principle:
Use complementary devices to generate gain without burning excessive static power

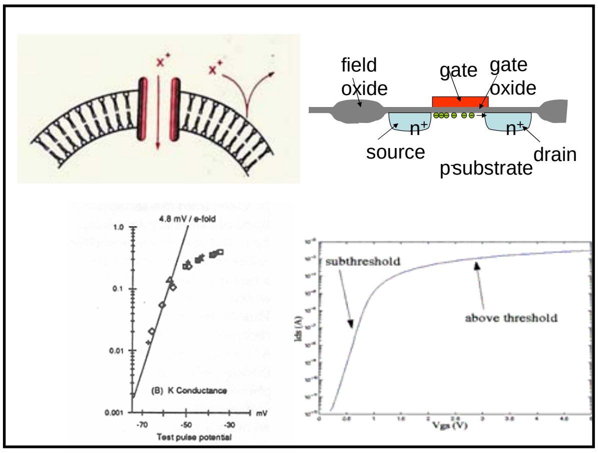

Mechanism of gain

Voltage sensitive channel conductance is exponential in membrane voltage	Transistor current is exponential function of gate voltage
Organizing principle: Use device physics (Boltzman distributions) to amplify	

Part 3:
Structure and function of the retina, as expressed in the "Physiologist's Friend Chip"

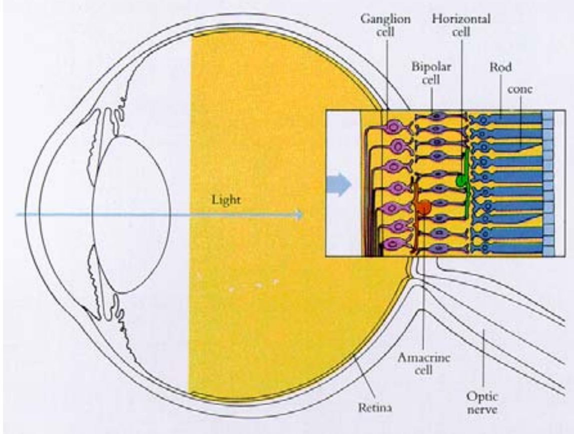

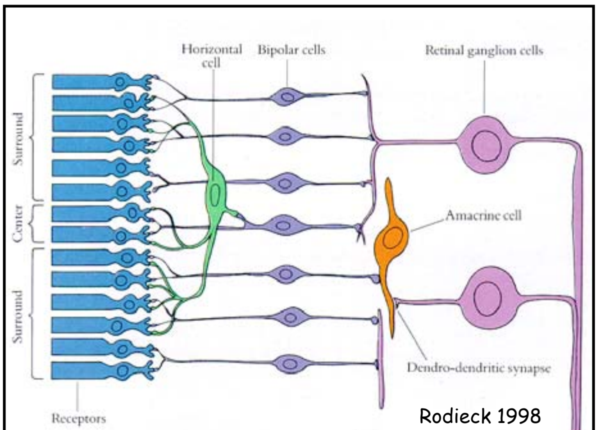

Rodieck 1998
The first steps in seeing

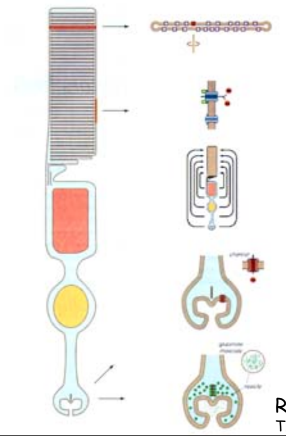

Rodieck 1998
The first steps in seeing

Biological photoreceptors adapt their operating point and gain

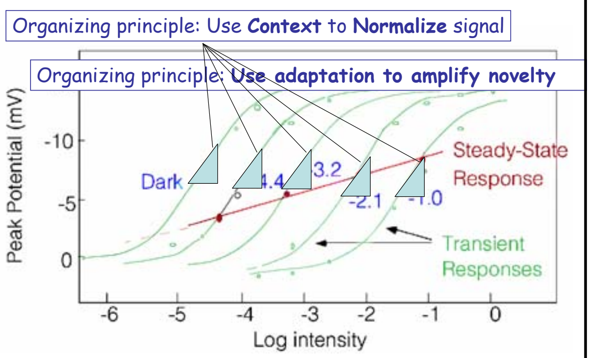

Norman & Perlman 1979

A silicon photoreceptor

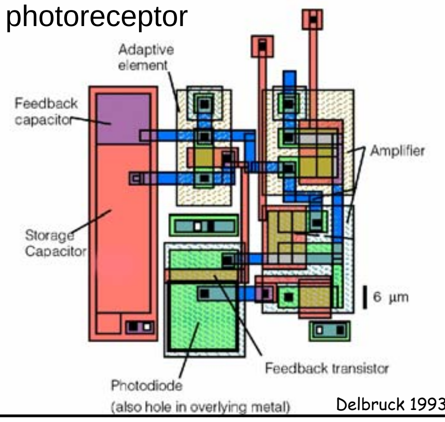

Delbruck 1993

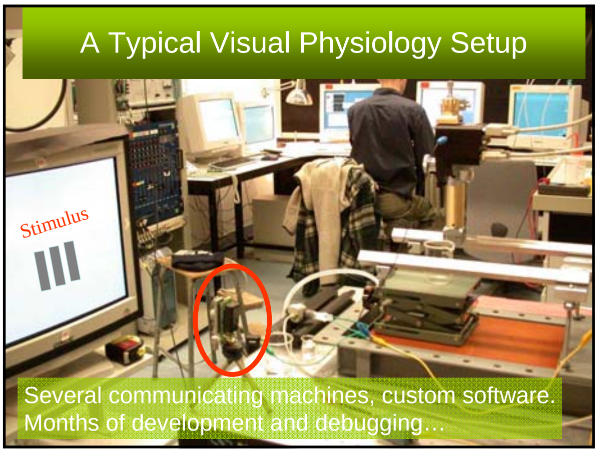

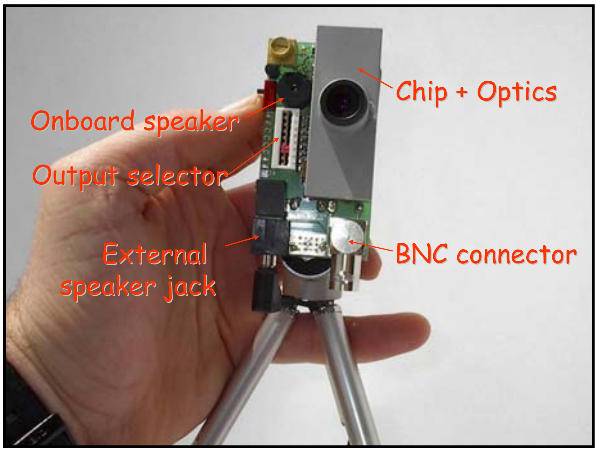

Physiologist's Friend Chip

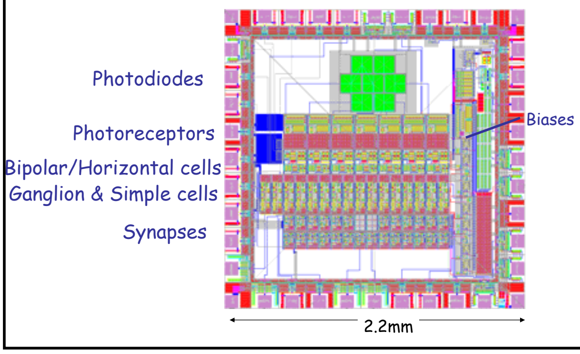

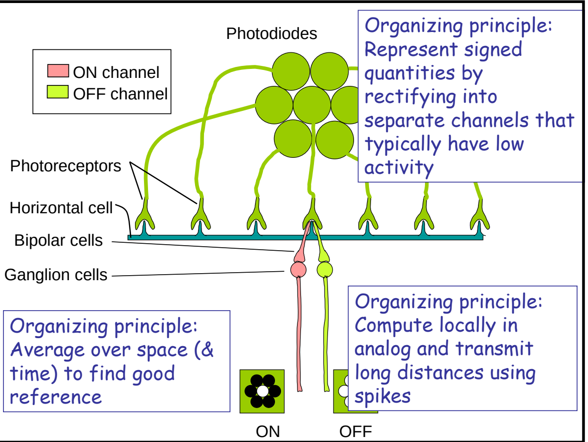

Review: Organizing principles

1. **Use device physics for computation**
   1. Sum currents onto nodes
   2. Use capacitance to integrate over time
   3. Use Boltzman physics to amplify
2. **Use complementary devices to amplify without burning excessive static power**
3. **Average over space (& time)** to find correct context and reduce noise
4. **Use Context to Normalize** signal
5. **Represent signed quantities by rectifying** into separate ON and OFF channels
6. Use **adaptation** to **amplify novelty**, not familiarity
7. Compute **locally in analog**, **communicate remotely using events**

IF WE ARE TO AVOID THE AI TRAP
WE HAD BETTER EVOLVE OUR SYSTEMS WITH REAL INPUT DATA
(BOTTOM UP)

THE AI TRAP:

1. ANNOUNCE INTENTION TO SOLVE AN OBVIOUSLY DIFFICULT PROBLEM
2. WORK LONG ENOUGH TO LEARN THAT IT IS MUCH MORE DIFFICULT THAN WAS INITIALLY SUPPOSED
3. FIND A TOY EXAMPLE THAT CONTAINS ONLY THE EASY PARTS OF THE PROBLEM
4. MAKE DEMO OF TOY EXAMPLE
5. DECLARE THE PROBLEM SOLVED WITHOUT REVEALING WHAT HAS BEEN LEARNED ABOUT THE HARD PARTS
6. GO TO STEP 1. OF A MORE DIFFICULT PROBLEM

Mead ca. 1990