CSE 40171:
Artificial Intelligence

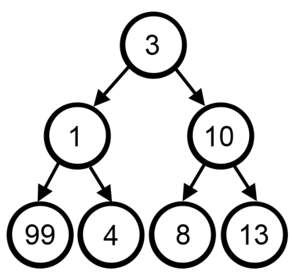

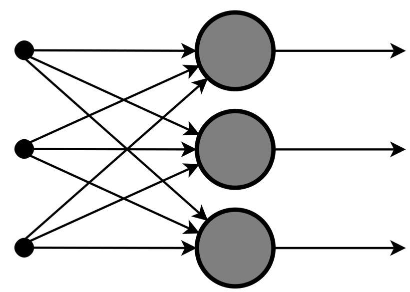

Artificial Neural Networks with Functional Fidelity:
Models of Neural Network Dynamics

Homework #7 has been released
It is due at 11:59PM on 12/2

Project Updates are Due **tonight** at 11:59PM

(See Course Website for Instructions)

Molecular Mechanisms

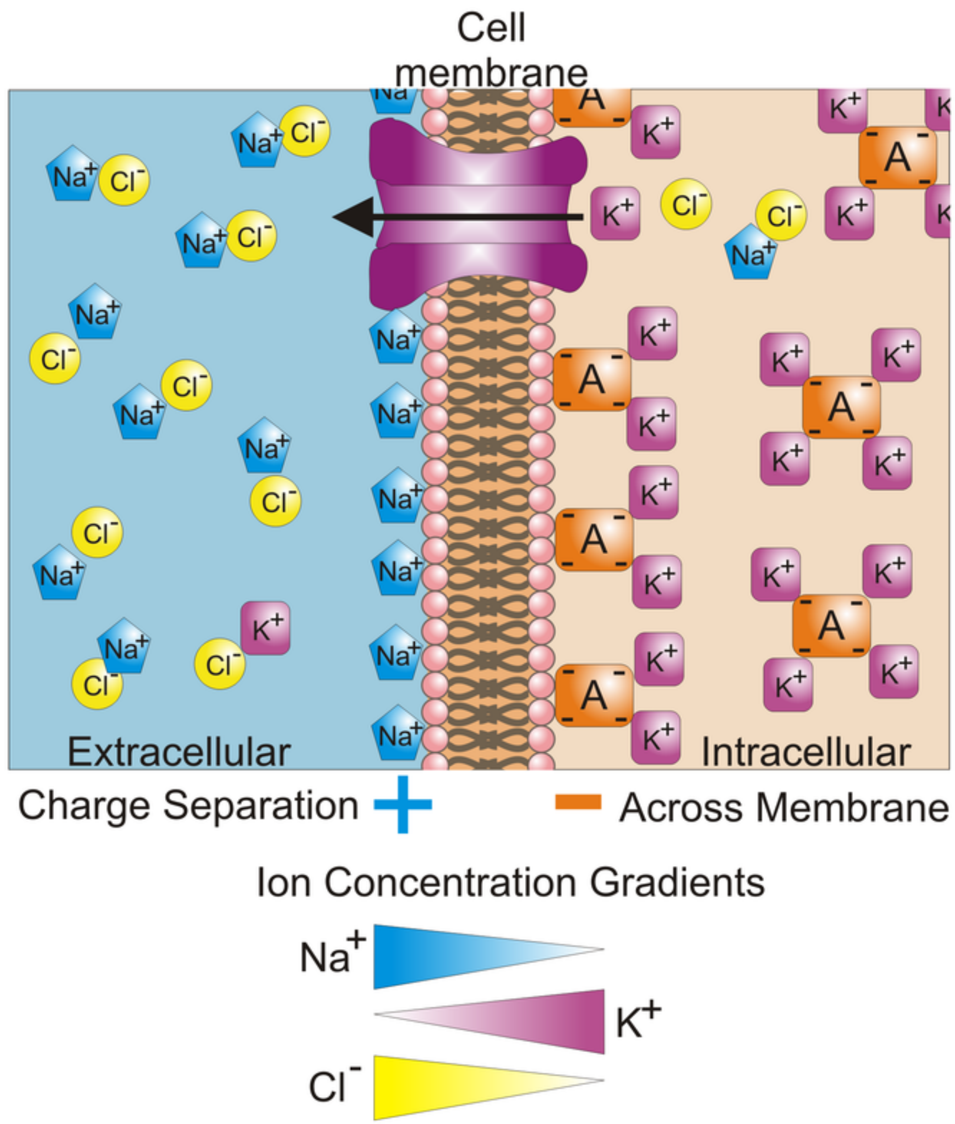

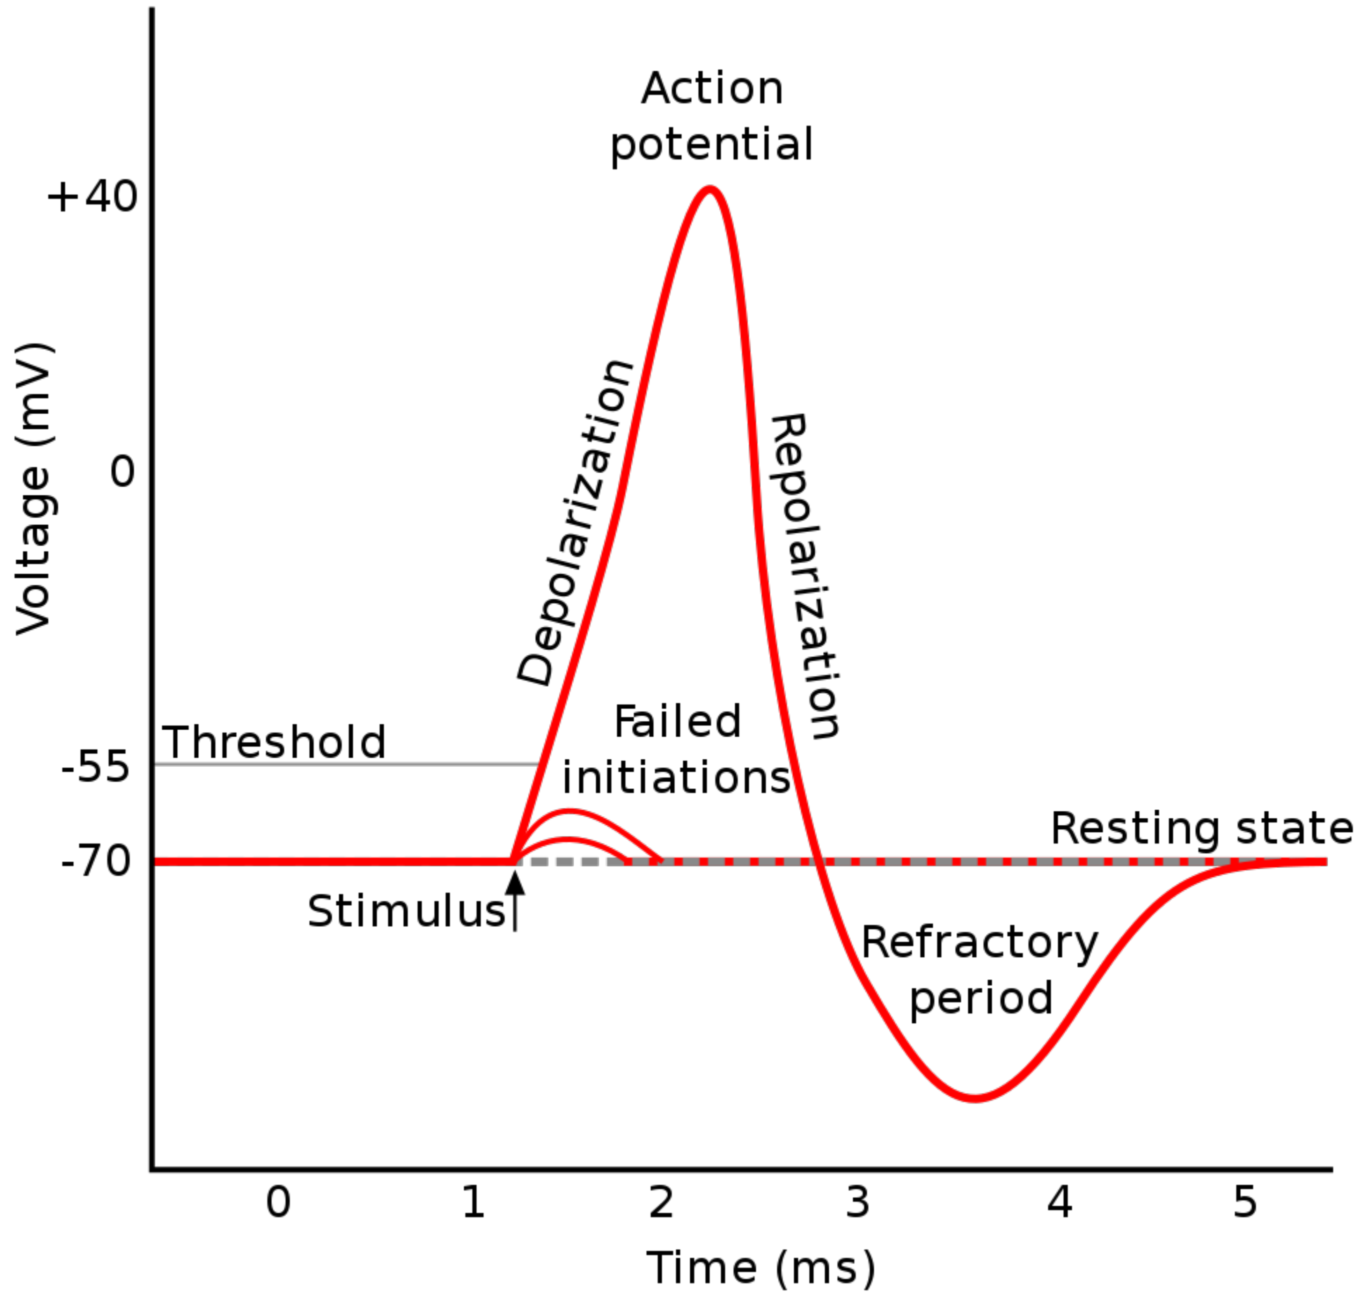
Action potential
+40
0
-55
-70
Voltage (mV)
Depolarization
Repolarization
Threshold
Failed initiations
Resting state
Stimulus
Refractory period
0
1
2
3
4
5
Time (ms)

Spike Trains

Action potentials convey information through their timing:

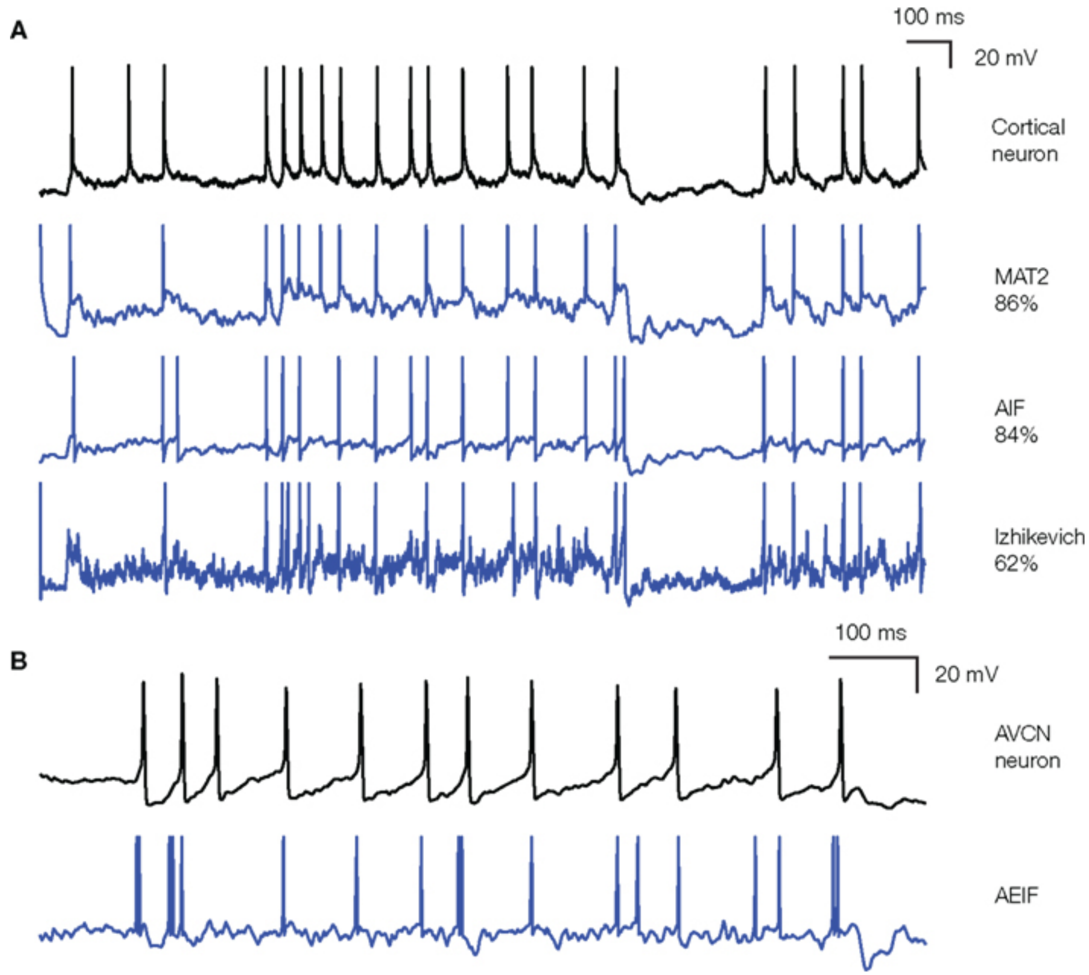

Rossant et al. Front. Neurosci. 2011

Firing Rates Approximated by Different Procedures

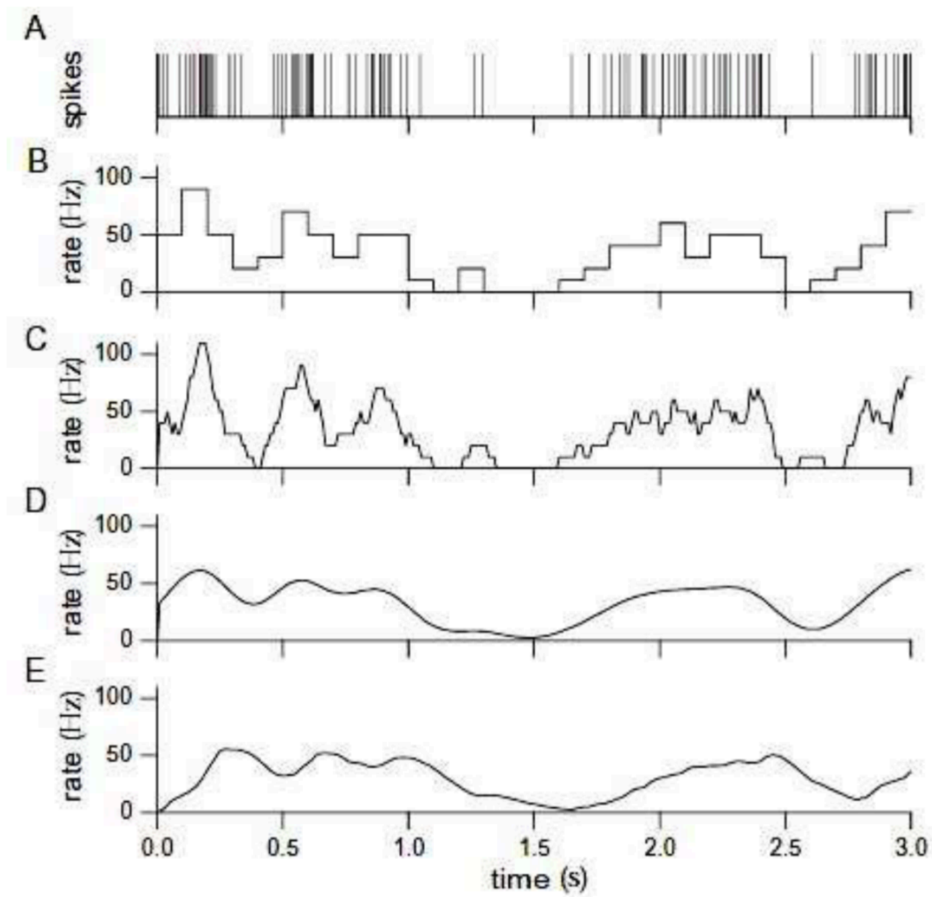

Spike train

Binwidth = 100ms

Smoothed (sliding window)

Smoothed (Gaussian window)

Dayan and Abbott 2001

Reminder of Tuning Curves

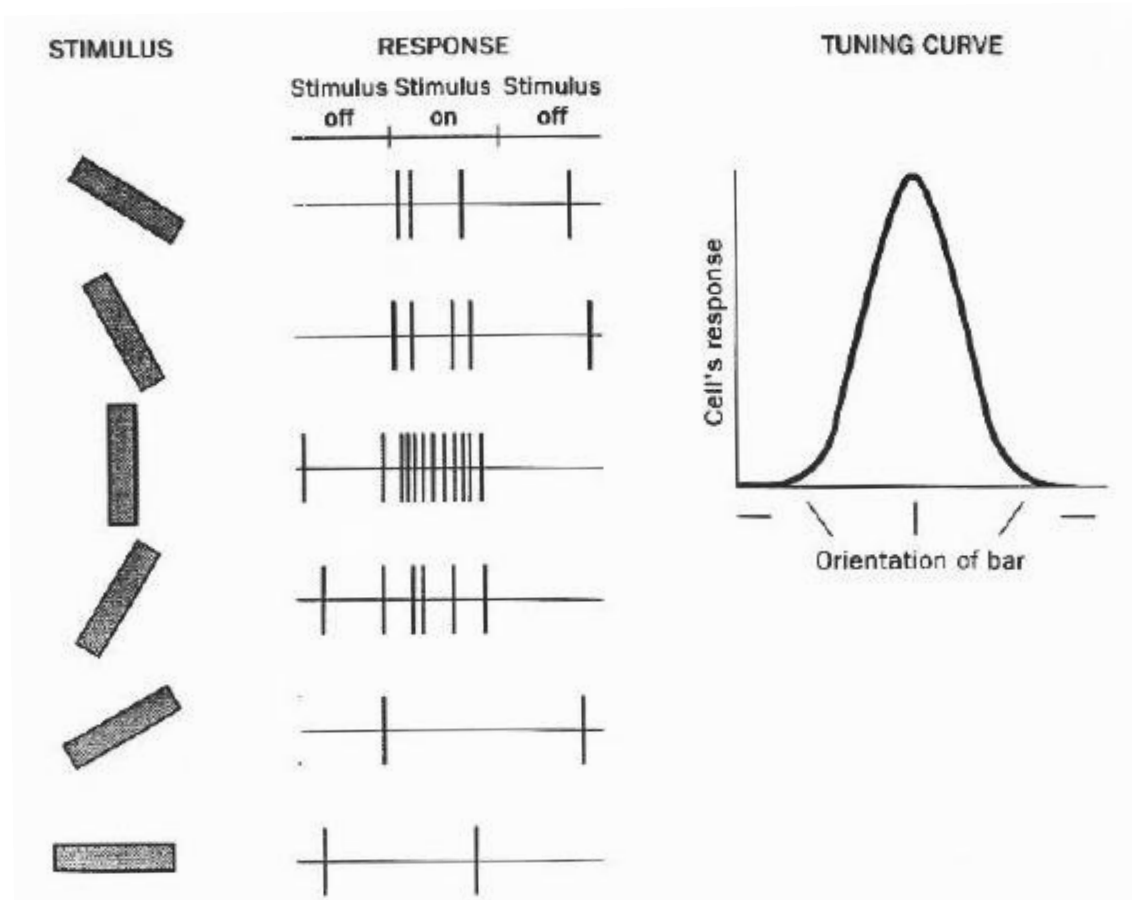

FIGURE 4.8 Response of a single cortical cell to bars presented at various orientations.

Hubel and Wiesel 1968

Modeling a Tuning Curve

Gaussian Tuning Curve:

orientation evoking max. response

max. avg. response rate

$$f(s) = r_{\mathrm{max}} \exp\left(-\frac{1}{2}\left(\frac{s - s_{\mathrm{max}}}{\sigma_f}\right)^2\right)$$

orientation angle of light bar

determines width of tuning curve

Modeling a Tuning Curve

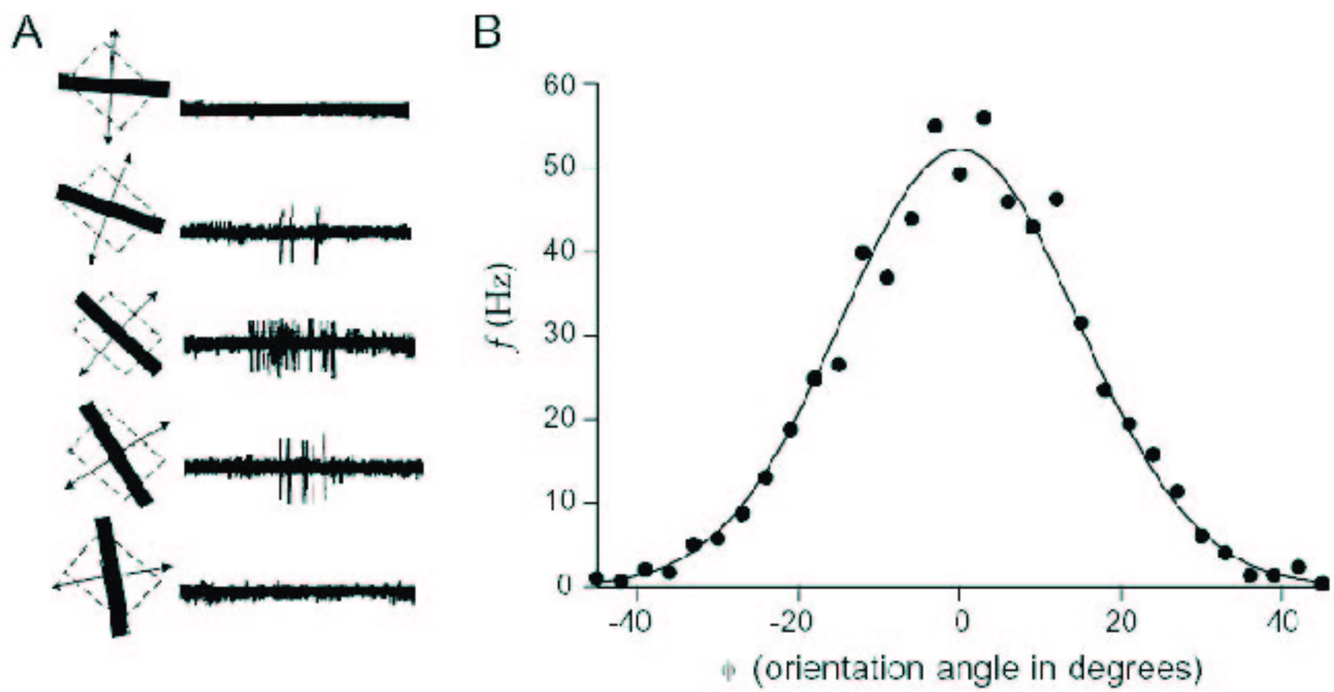

Dayan and Abbott 2001 (original: Wandell 1995)

Note: Spike Count Variability

- Tuning curves allows us to predict the average firing rate
- They **do not** describe how the spike-count firing rate r varies about its mean value from trial to trial
    - likely that single-trial responses can only be modeled probabilistically

Describing the stimulus

Neurons responding to stimuli must encode parameters that can vary over a large dynamic range.

Weber's law: how different two stimuli have to be to be reliably discriminated. The just noticeable difference Δs is proportional to the magnitude of the stimulus s, such that $\Delta s \,/\, s$ is constant.

Fechner's law: noticeable differences set the scale for perceived stimulus intensities. Integrating Weber's law, the perceived intensity of a stimulus of absolute intensity s varies as log s.

Adapting to the stimulus

Sensory systems make numerous adaptations to adjust to the average level of stimulus intensity.

Model this by describing responses to fluctuations about a mean stimulus level.

$s(t)$ is defined so that its time average over the duration of the trial is 0:

$$\int_0^T dt\ s(t)/T = 0$$

Stimuli and Time Averages

Scenario 1: 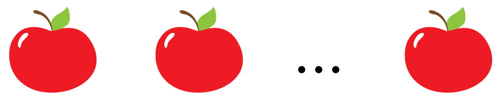

t_1 t_2 t_n

Scenario 2: 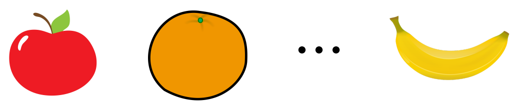

t_1 t_2 t_n

Spike-Triggered Average:

$$C(\tau) = \left\langle \frac{1}{n} \sum_{i=1}^{n} s(t_i - \tau) \right\rangle \approx \frac{1}{\langle n \rangle} \left\langle \sum_{i=1}^{n} s(t_i - \tau) \right\rangle$$

time interval

Computing the spike-triggered average stimulus

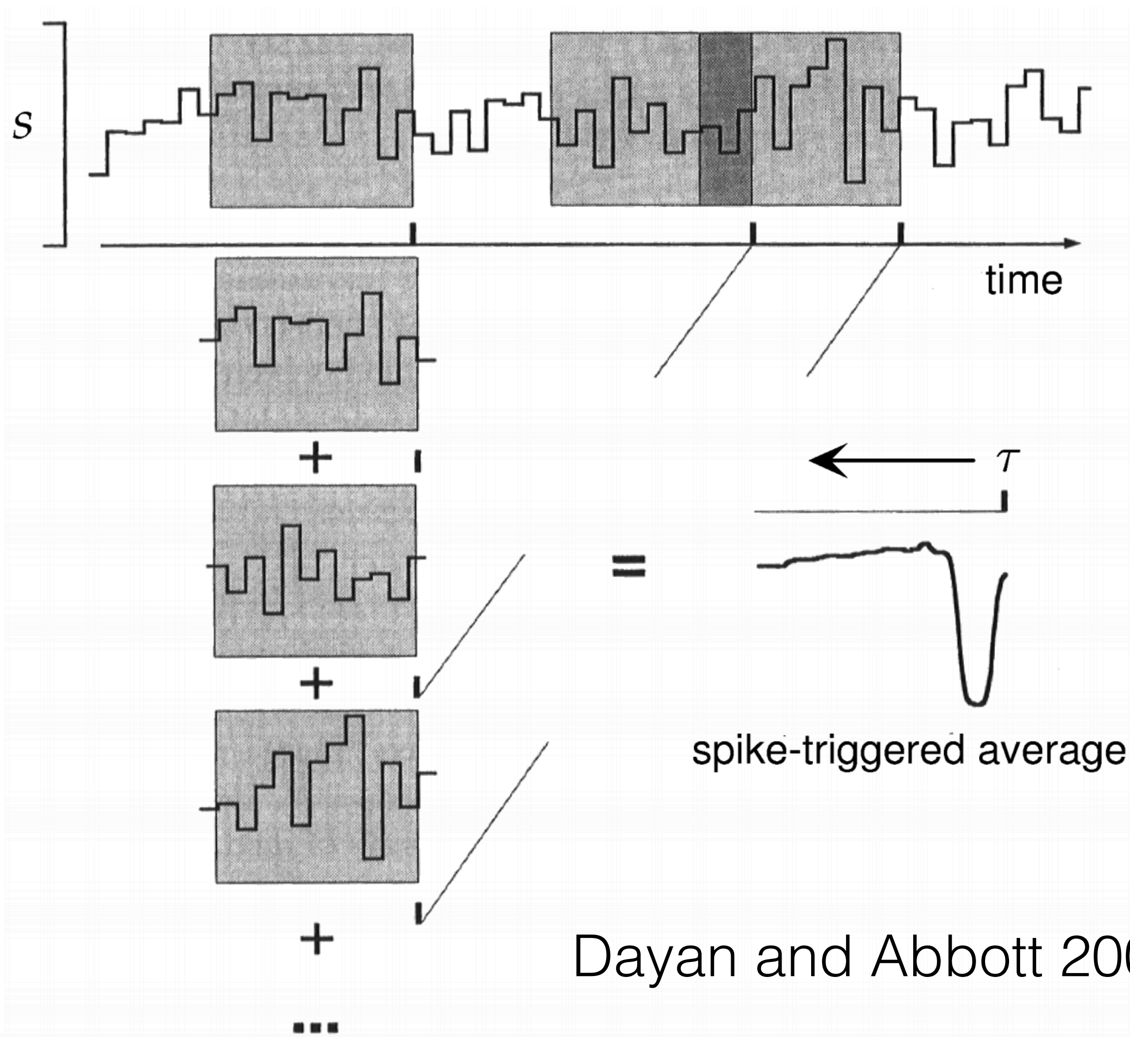

Dayan and Abbott 2001

What does incorporating these dynamics into an artificial neural network provide us with?

Spiking Artificial Neural Networks

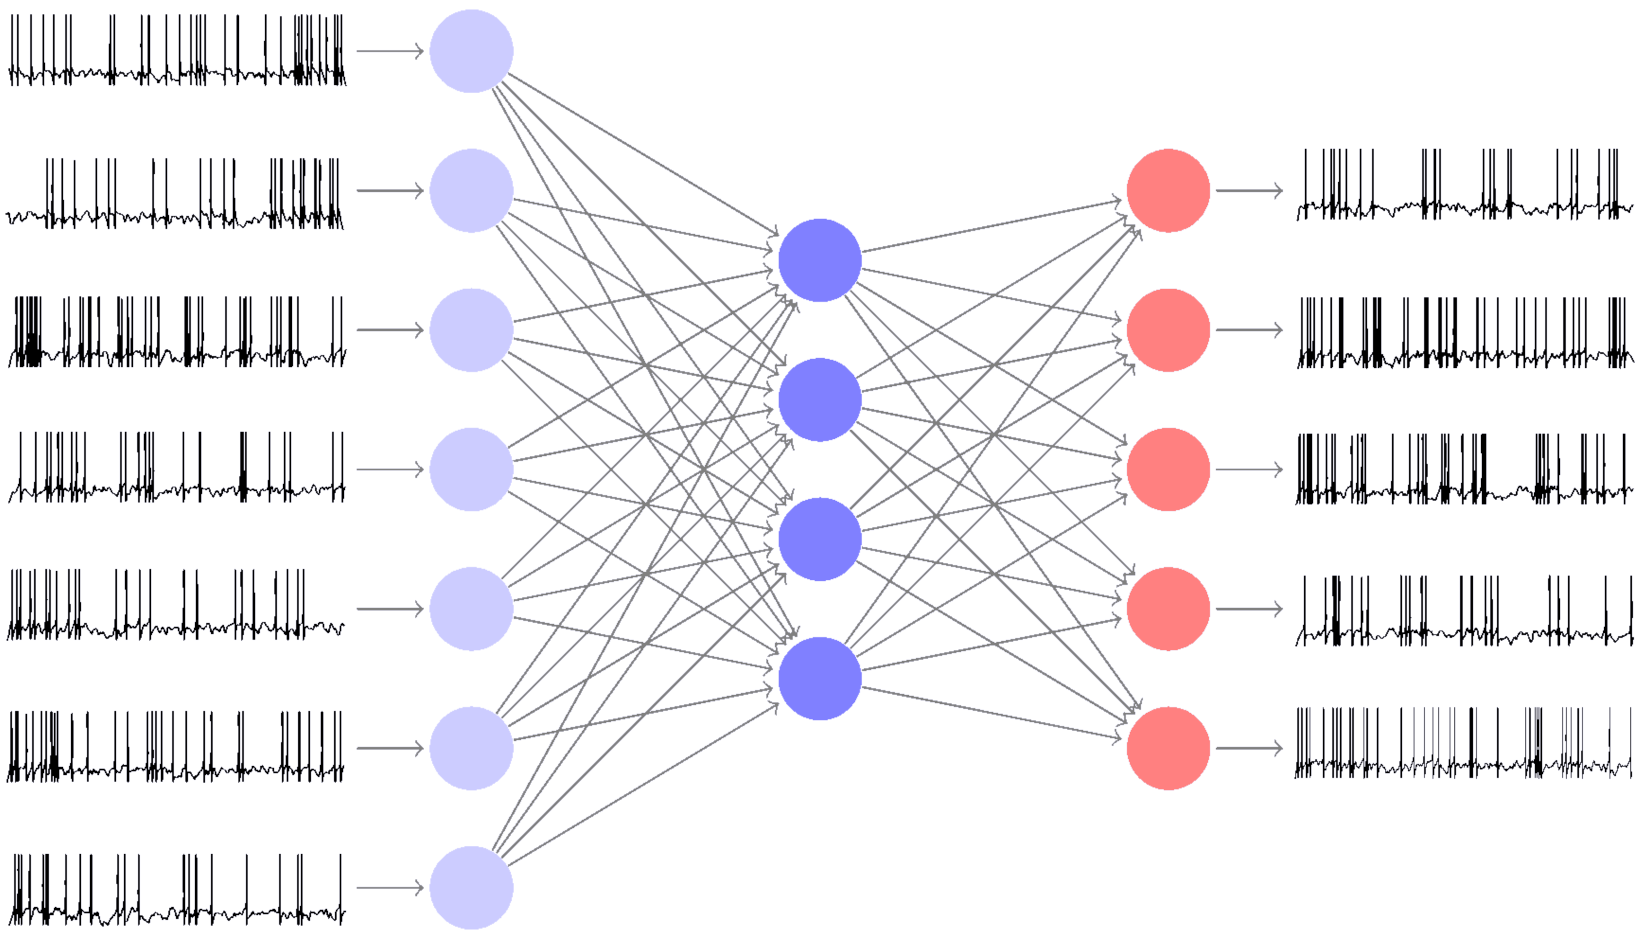

Image Credit: https://fzenke.net/index.php/2017/02/19/learning-in-multi-layer-spiking-neural-networks/

Advantage of Spiking Artificial Neural Nets

- Neuroscientists believe that information is encoded and decoded by a spike train
    - Do neurons communicate with a rate or temporal code?
- Temporal coding suggests that a single spiking neuron could replace hundreds of hidden units in a conventional artificial neural network

The Neural Code

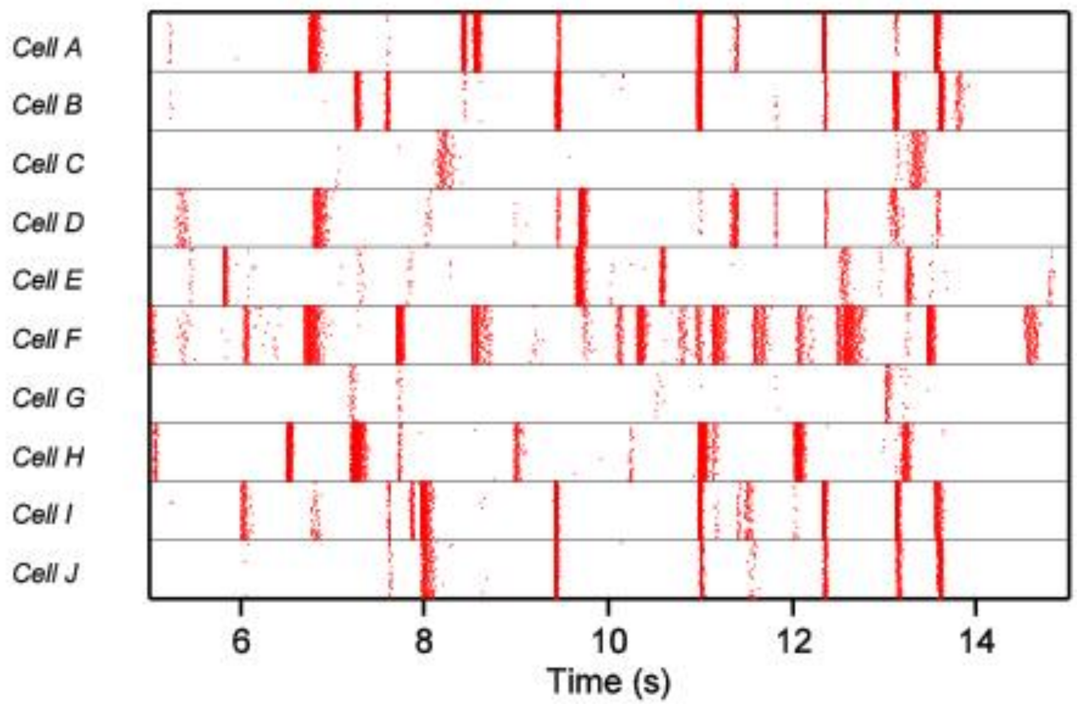

Image Credit:Alan Litke

Spiking neural networks consider temporal information

- Not all neurons are activated in every iteration of propagation
- A neuron is activated when its “membrane potential” reaches a threshold
- After activation, a signal is produced that is sent to connected neurons, raising or lowering their membrane potential

Unit in a Spiking Artificial Neural Network

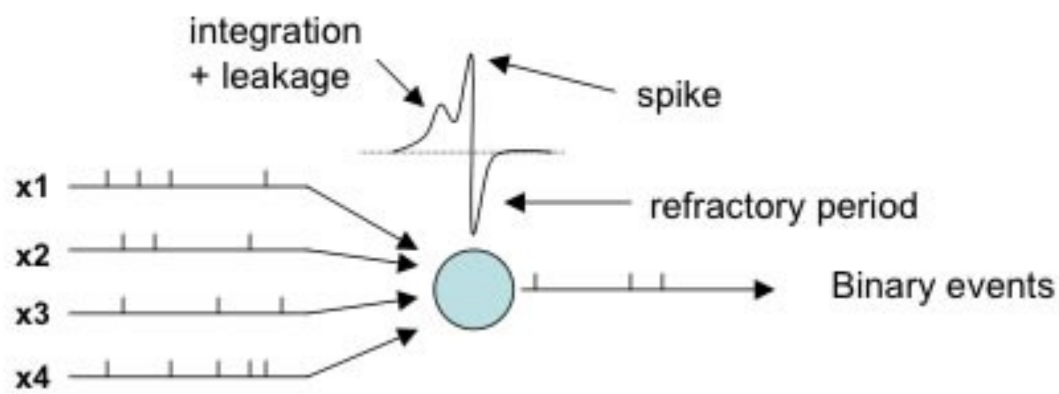

Image Credit: https://lis2.epfl.ch/CompletedResearchProjects/EvolutionOfAdaptiveSpikingCircuits/

Neuromorphic Architectures

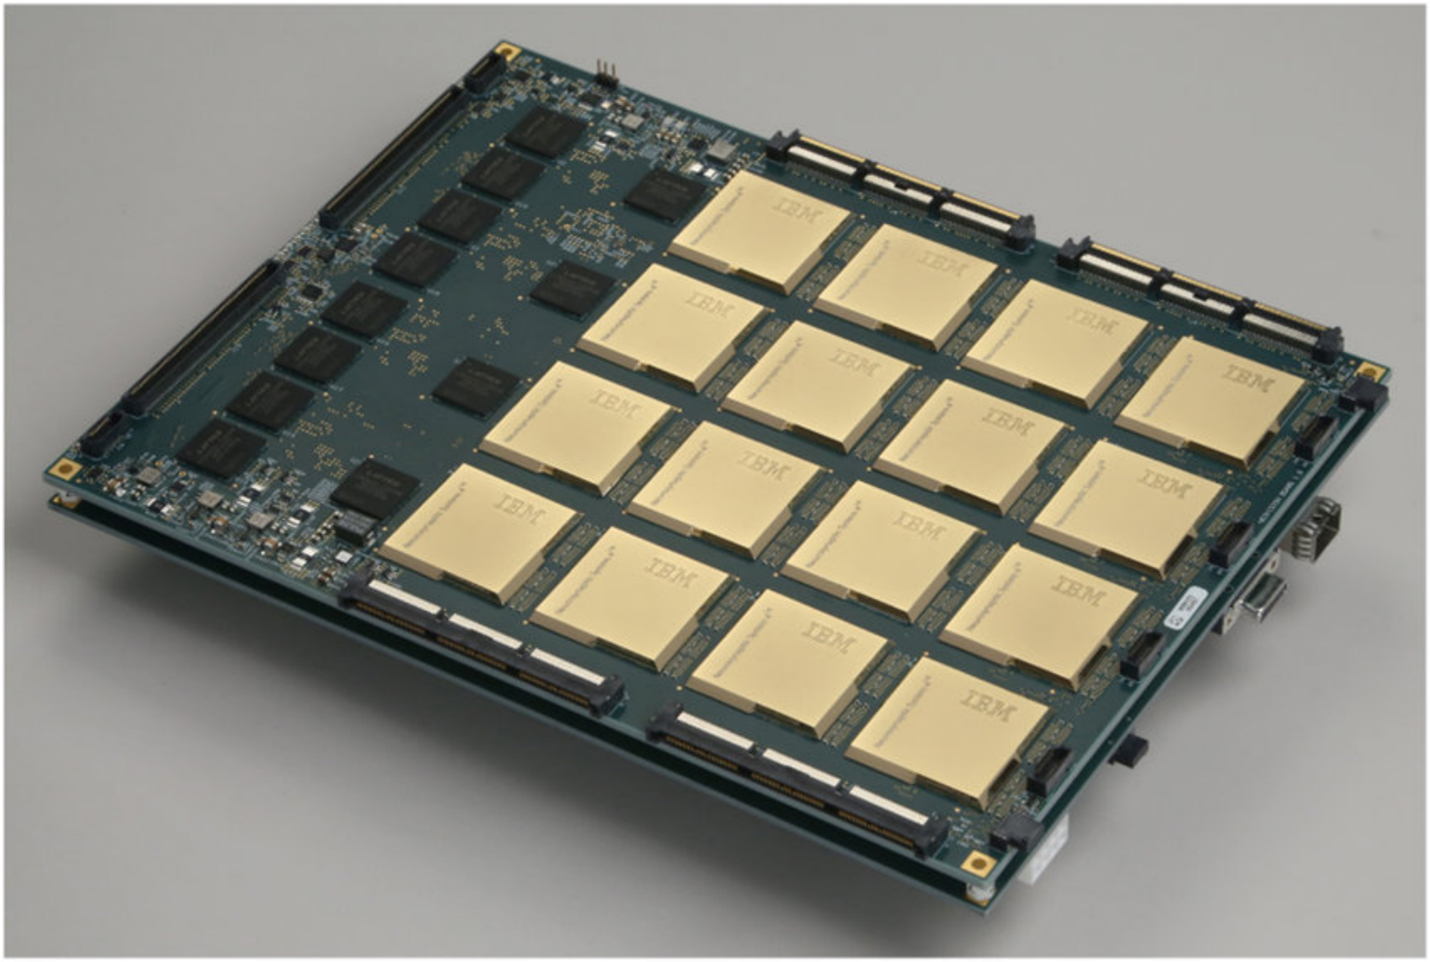

Image Credit: IBM Corporation

Key advantages of neuromorphic hardware

- Energy efficiency
- Execution speed
- Tolerance to local failures
- Ability to learn

Neurogrid (Stanford)

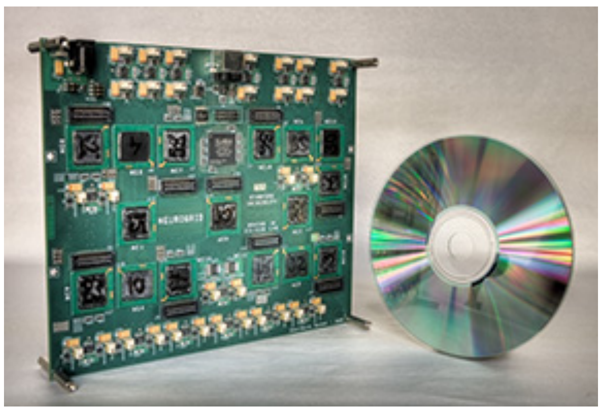

Image Credit: Brains in Silicon Group, Stanford University

- Analog computation to emulate ion channel activity
- Digital communication for structured connectivity patterns
- Simulates 1 million neurons and 6 billion synapses
- Consumes less than 2 watts of power

BrainScaleS (Human Brain Project)

- 200,000 neurons and 40 million synapses per system
- 20 such systems in the first version of the system
- Simulation of plasticity models
- Runs 10,000x faster than real time

Image Credit: https://flagship.kip.uni-heidelberg.de/public/BrainScaleS/

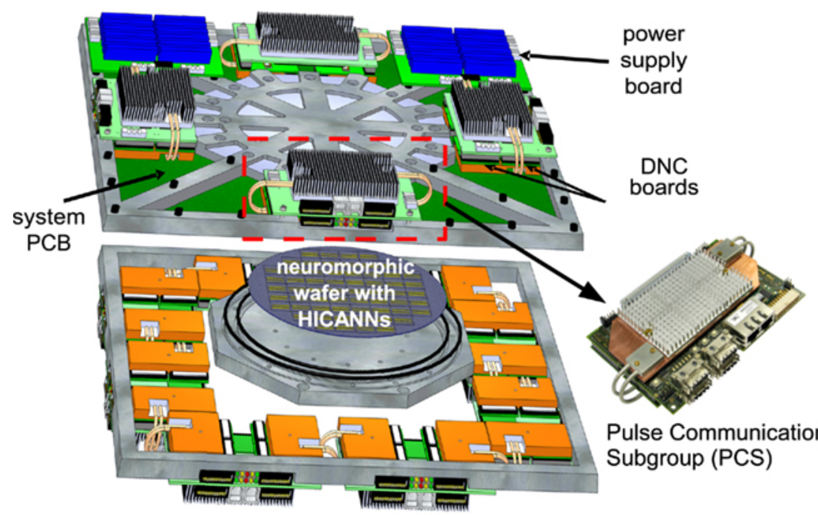

Image Credit: Schemmel et al. ISCAS 2010

SpiNNaker (Human Brain Project)

- Custom chips based on ARM
- Each chip has 18 cores and 128M of shared local RAM
- Over 1,000,000 cores available
- Based on numerical models of neuron dynamics

Image Credit: University of Manchester

TrueNorth (IBM)

- 4,096 cores, each with 256 programmable neurons (~1,000,000 neurons)
- ~268M programmable synapses
- 5.4B transistors, but only consumes 70 milliwatts of power
- Typical CPU: 1.4B transistors and 35+ watts of power
- Designed for pattern recognition

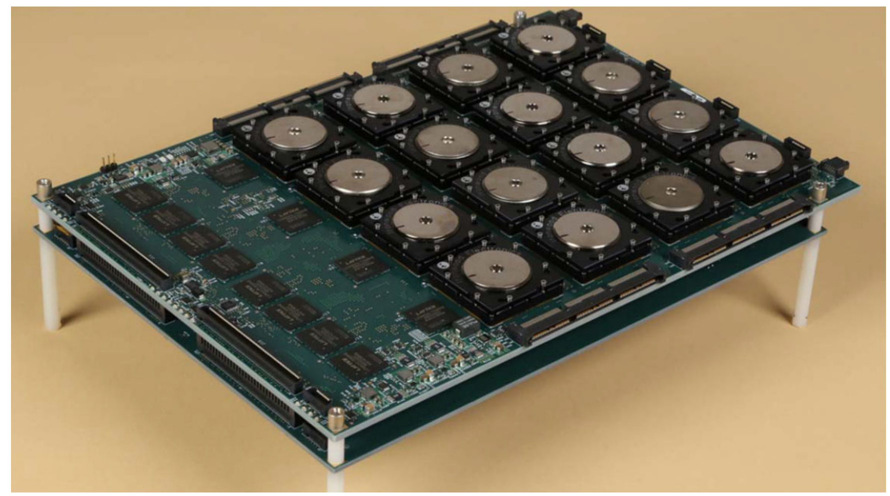

Second Order Effects: What is the model of computation?