NBio 136 - Computational Neuroscience - Course Information

January:

22 Introduction, Firing Rates and Tuning Curves (Ch 1)
25 Introduction to MatLab, Poisson Spiking (Ch 1)
29 Poisson Spiking, Spike-Triggered Averages (Ch 1)

Encoding - Decoding

February:

1 Reverse Correlation Methods (Ch 2)
5 Modeling Visual Responses (Ch 2)
8 Population Coding (Ch 3)
12 Population Decoding (Ch 3)
15 Discrimination from Neural Responses (Ch 3)
19 No Class
22 No Class
26 To Be Announced

March:

1 Information Theory (Ch 4)
5 Information Maximization (Ch 4)

Biophysics

8 Electrical Properties of Neurons (Ch 5)
12 Integrate-and-Fire Models (Ch 5)
15 The Hodgkin-Huxley Model (Ch 5)
19 Channels and Diffusion of Ions (Ch 6)
22 Synapses (Ch 6)
26 No Class
27 Special Wednesday Class – Multicompartment Neurons and Cables (Ch 6)
29 No Class

April:

2 No Class
5 No Class
9 Network Models (Ch 7)
12 Network Models (Ch 7)
16 Network Models (Ch 7)
19 Coordinate Transformations - Parietal Cortex (Ch 7)

Adaptation & Learning

23 Synaptic Plasticity (Ch 8)
26 Hebbian Learning (Ch 8)
30 Memory Models (Ch 8)

May:

3 Models of Development (Ch 8)
7 Classical Conditioning (Ch 9)

Office Hours:

Larry Abbott - Volen 317 (enter through 316) - abbott@brandeis.edu - 62876
Tues 1-2, Fri 1-3

Patrick Drew – drew@brandeis.edu – 63138

Stephen Van Hooser – vanhoosr@brandeis.edu - 62319

Course Web Site: http://www.bio.brandeis.edu/classes/nbio136/

Tues. 22 Jan

Stimulus - Response Mapping

- APs (1ms, .1 Volt) are common encoding in cortex, Thalamus (LGN), RGCs
- In e.g. photoreceptors, graded response occurs (analog encoding of stimulus)

Why the analog → discrete conversion?

- Can't propagate analog signals w/o local connections (< 1mm)
  e.g. gap junctions

- Since all spikes are identical, spike timing is the only information there is
  - Means time is used to show changes in input, **AND** time is part of the code itself.
- Code is also stochastic, given stimulus will rarely produce identical output
  - To paper over this, condense the train down to 'firing rate'

Firing Rate(s)

- Spike Count Rate: $\frac{\text{n. spikes}}{\text{t. of stim}}$

- Time-dependent FR: Bin-up rasters and count avg. number of spikes in each
  Once you have a value for each bin, you have a continuous signal again, and one that accounts for the randomness in individual traces

Probabilistic relationship b/t signal + spiks:

Narrow bin Δt centered at time t

Prob. of spike: $P[\text{spike}] = r(t) \cdot \Delta t$!

Models:

$S(t) \rightarrow R(t)$ from stim. to firing rate

$R(t) \rightarrow$ spikes Probabilistic train generation (e.g. Poisson process)
- Divide time into bins, get prob. of a spike in that bin based on $R(t)$, then roll the dice

Average Firing Rate: $\langle r \rangle$

- Just avg. of entire $R(t)$ function: 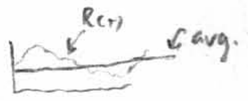

- To measure selectivity of cell, plot $\langle r \rangle$ vs. Stim. Strength/Value
  this gives a tuning curve

Tuning Curves:

{ S → constant for given trial
{ Resp → Avg. rate given S

1) Sigmoidal f

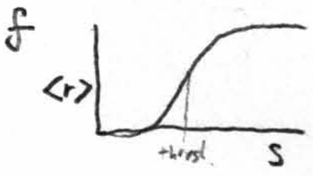

2) Gaussian f (like θ-tuning)

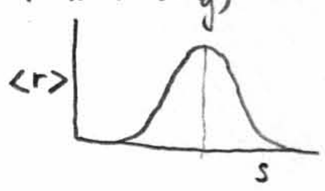

"S" doesn't necessarily mean external stimulus (e.g. in motor cortex, "S" is really command to start moving/continue/etc.

3) Cosine

Non-discrete Instantaneous Firing Rate

- Rather than chopping rasters into bins and making a histogram, use a window function as a filter, then slide window along
  (bin case is equivalent, but with no two windows overlapping)
- Sliding window minimizes aliasing
- Also, not restricted to square bin. Could use, e.g., gaussian window
- If window is centered on t it's actually using spikes from the future to generate the rate
  - Instead can use window that only projects into past
  - Only constraint on diff. windows is that area under curve = Δt

Tues 29 Jan

Continuous vs. discrete

$S(t)$: A continuous function

S_i -or- $\{S_1, S_2, S_3, \ldots S_N\}$: Discrete value representation of $S(t)$

where $S_i = S(i \cdot \Delta t)$

Integration & Derivatives

dS/dt : Change of stimulus over time (continuous)

$S_{i+1} \rightarrow S_i + x$: Update rule to get value for next timestep

---

$\int_0^T S(t)\,dt$: Total stimulus over range of time

$\Delta t \sum_{i=1}^{N} S_i$: Sum of all discrete stimulus values

Representing spike trains

- Can make a δ that's 1 at Δt's with spikes, 0 in other Δt bins
- In continuous notation, as $\Delta t \rightarrow 0$, bins must be infinitely tall as well as thin so area under spike curve remains constant across Δt sizes. a.k.a a δ rule
- Since Δt value is essentially arbitrary, the important quantities are those that can be extracted away from Δt

  e.g. $P(t) = r(t)\,\Delta t$, any two Δt's give diff. results, $r(t)$ is consistent

Stimulus → Response → Spike Train Mapping

- Is Poisson $P(t) = r(t) \cdot \Delta t$ actually a good way to generate trains?
  - Problem is it doesn't consider previous spiking history, in particular, there is a lower probability of many spikes in a row (ignores refractory period behavior)
- 'Patched' Poisson Model
  - Drop $r(t)$ to zero (artificially) when you spike, then let it decay back to its ballistic $r(t)$

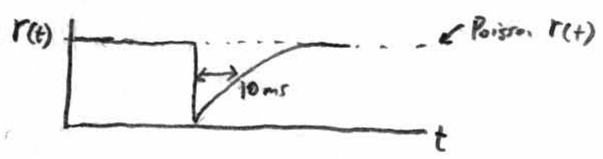

- Exponential Recovery
  - At every timestep, the difference b/t r and r_0 decreases by the same proportion (e.g. 20% recovery)
  - The 'x' value is always dependent on Δt, so to make an invariant measure, let:

    $x = \exp\left(-\frac{\Delta t}{\tau}\right)$ where $\tau \approx 10$ ms

    This way the curve is always the same, regardless of the Δt you've chosen

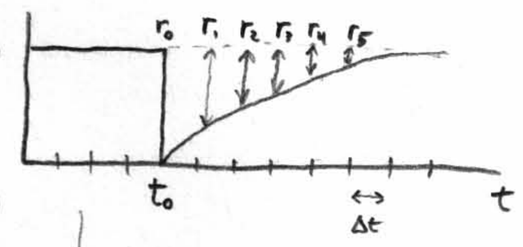

$(r_0 - r_1) = x r_0$ (where $0 \leq x \leq 1$)

$(r_0 - r_2) = x(r_0 - r_1)$

$\Downarrow$

$$r_{i+1} = r_0 + e^{-\frac{\Delta t}{\tau}} \cdot (r_i - r_0)$$!

Nature of the Neural Code

- Since all spikes are identical, it's all about the timing
- DEBATE: Timing relative to what?
  1) Timing relative to stimulus (rate code)
     - Spikes are a direct representation of stimulus signal
  2) Timing relative to other spikes (ISI code)
     - e.g. in the same way that the '1' in "12" really means 10, the spikes in a sequence each represent different things and change the meaning of subsequent spikes → slow spike = upswing, fast = downswing for example
  3) Timing relative to other neurons (spike timing)
     - Importance lies in interval b/t cell A's spike and cell B's
- But really it seems like ISI coding either doesn't exist, or happens in very weak ways (in Larry's opinion)
  - Typically 95% of information conveyed by cells comes from rate code and 5% from ISI coding
  - Harder to write off #3 since multi-unit recording is still young

---

Stimulus → Rate Mapping

- Looking at a stimulus and output spikes, how do you decide what feature the neuron is responding to?
  - First, average across trials to deal with stochasticity of response
  - Then for each spike in the trace, look at stimulus in a recent-history window. If you average the windows you get:

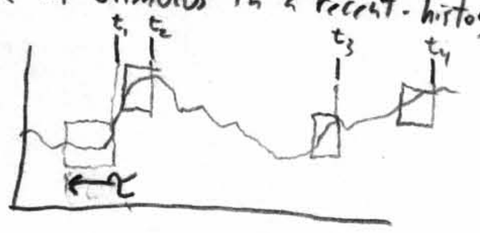

Spike Triggered Average

$$C(\tau) = \frac{1}{N} \sum_{i=1}^{N} S(t_i - \tau)$$

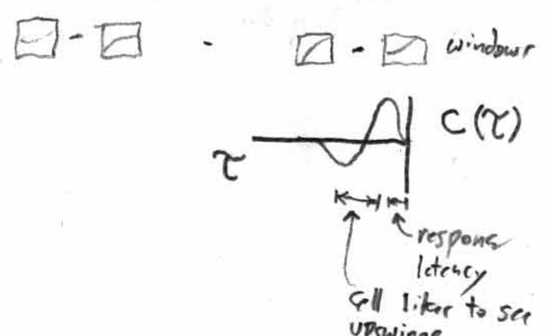

- Problem with STA is if you give it a stimulus that doesn't provide what the cell is looking for (i.e. you can bias your interpretation through stimulus selection)
- Solution is to use a stim. that has all possible Patterns w/ equal probability: White Noise (each S_i is just a random value)
- Problem w/ W.N. is getting larger, more structured patterns b/c if variance is too small signal will just average out to meaninglessness

Fri 1 Feb

Spike-Triggered Averages (a.k.a. Reverse

- Grab stimulus window just before spike, then average all the windows collected
- To deal with dc of avg. stimulus, you usually rig it so it divides its time evenly above and below a zero value

$$C(\tau) = \frac{1}{N}\sum_{i=1}^{N} S(t_i - \tau)$$ $t_1, t_2, \ldots$ = spike times

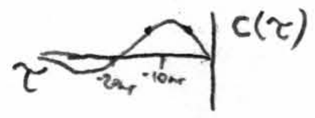

- NOTE: if there's actually a value on the positive side of the plot, it means there's some sort of correlated pattern to S (thus the white noise to decorrelate stim)

- What to do if you have r(t) instead of individual spike times? Use probability of spikes at any given time bin

$$C(\tau) = \sum_{bins} Prob(bin)\cdot S(t_{bin} - \tau)$$

$$= \sum_{bins} r(t_{bin})\,\Delta t \cdot S(t_{bin} - \tau)$$

as bins → 0, this turns into an integral

$$= \frac{1}{n}\int_0^T dt\; r(t)\,S(t-\tau)$$

↑ avg. number of spikes in 0–T interval

! $$= \frac{1}{\langle r\rangle T}\int_0^T dt\; r(t)\,S(t-\tau)$$

$$\langle r\rangle = \frac{1}{T}\int_0^T dt\; r(t)$$

Correlation Function (of stim + response)

! $$Q_{rs}(\tau) = \frac{1}{T}\int_0^T dt\; r(t)\,s(t+\tau)$$

- If you get a significantly pos you know they're related

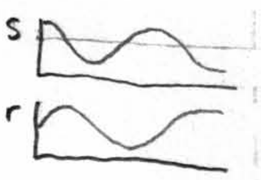

The response is correlated to stim but shifted in time. By looking at $Q_{rs}(\tau)$ for diff. τ's we can shift the curves to see if they align.

So really,

$$C(\tau) = \frac{Q_{rs}(-\tau)}{\langle r\rangle}$$

White Noise

- With reverse correlation it's easy to stack the deck by only showing a particular type of stimulus.

e.g. H1 Neuron of Fly visual system

S: velocity of drum surrounding fly (averaged to 0 over trial)

- You need to be careful to provide all different velocities (otherwise 5°/s cell might look like a 3°/s if you never showed it anything faster than 3°/s)

- Creating white noise

S_{bin} = rand() w/ 2 conditions:

1) avg (S) = 0
2) variance = $\frac{(x - avg.)^2}{\#\text{ of samples}} = \frac{\sigma_s^2}{\Delta t}$

As you choose a smaller Δt, you need to scale your variance so you can directly compare your variance with that of others w/ different Δt's

What's an STA good for?

- Allows you to predict r(t) from S(t)
- Take C(τ) and superimpose it on S over the window size and see how good the template match is,

$$r(t) = \int_0^\infty dt\; D(\tau)\,S(t-\tau)$$

S(t−τ) ← stim going into past

D(τ) ← Weighting function for how important (and what sign) S is at various τ offsets

↑ Looking into stim's past

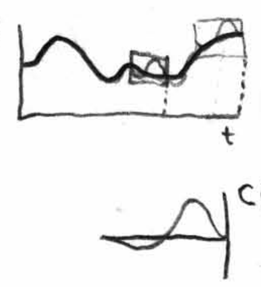

$$D(\tau) = \frac{\langle r\rangle}{\sigma_s^2}\,C(\tau)$$

- Receptive field is proportional to C(τ) but scaled to go from stimulus units, to firing rate units
- σ_s^2 has units of (stimulus)^2 · time

Why are C(τ) and D(τ) equivalent

- Difference of S(t−τ) − C(τ) should tell how good a match stim and RF are,

$$\int d\tau\,(S(t-\tau) - C(\tau))^2 = \int d\tau\; S^2(t-\tau) - 2\int d\tau\,(S(t-\tau)C(\tau) + \int dt\,(C(\tau)\ldots$$

What is mapping b/t size of $\int D(\tau)\cdot S(t-\tau)$?

$$L(t) = \int_0^\infty D(\tau)\,S(t-\tau)$$

So L(t) is linear function of template match

$$r(t) = F(L(t))$$

Where F is a 'static nonlinearity' mapping from template convolution to an output

So then, with Prob = r(t)·Δt, we can simulate this theoretical neuron

The Static Nonlinearity

- F() exists to deal w/ a) impossibility of negative r(t)
  b) impossibility of ∞ r(t)

a) is easily solved w/ rectification $F = [L]_+ = \begin{cases} L & \text{if } L \geq 0 \\ 0 & \text{if } L < 0 \end{cases}$

b) is solved w/ saturation

sigmoid smoothes out both of these

- So think of L(t) as input current to neuron and F() as biophysics of cell to translate current to firing rate

Vision

- Since images are 2D stimuli, S(t) is no longer a single value. We're now addressing locations S(x,y,t)
- Also means STA is function of 3 variables C(x,y,τ), also D(x,y,τ)

Tues 5 Feb

Modeling Process

1) $r = F(L) = [r_0 + \alpha L(t)]_+$ (basal firing rate) where $L()$ is the linear mapping of input signal...
2) $L(t) = \int_0^\infty d\tau\ D(\tau)\ S(t-\tau)$ (Any stimulus) and $D(\tau)$ is the temporal RF
3) $D(\tau) = \frac{\langle r \rangle}{\sigma_s^2} C(\tau)$ ← spike-triggered avg. w/ white noise at variance $\frac{\sigma_x^2}{\Delta t}$

- To model a cell, get the STA from real cell, then plot your computed L against the output firing rates, The resulting function is your F() static non-linearity

{Dot Product: Take every component of two vectors, multiply them and sum the results}

Visual System

- $r = F(L)$ same as before, but L is spatial as well as temporal:
- $L = \int d\tau \int dx \int dy\ \overbrace{D(x,y,\tau)}^{RF}\ \overbrace{S(x,y,t-\tau)}^{image}$
- $D = \frac{\langle r \rangle}{\sigma_s^2} C(x,y,\tau)$
- Spike triggered avg. will show where D() is positive (white) and negative (black) or doesn't care (grey) at a given τ
- Dorsal pathway (parietal) is 'where' pathway for targeting objects for motor acts, saccades, etc.
- Ventral pathway (temporal) is 'what' pathway for recognising particular objects
- Thalamus Processing fairly unimportant — just passes on RGC activity patterns (except during sleep)
- RGC RFs flip over time: maximal response when stim shifts from dark center + light surround to light center + dark surround (for an on-center cell)
  - Low latency on center (wants to see light at ≈ 20 ms), slightly longer latency on surround (40 ms?).
  - At $\tau \approx 100$ ms, prefers negative of its pattern
  - DoG shows it likes to see light-dark change in terms of space. Reversal of RF shows it also wants to see change over time.
- V1 cells lose the circular symmetry: want to see elongated contrast regions this gives you edge detection & orientation selectivity.
  - Modeled using gabor function (oscillatory base w/i gaussian envelope)

$$D(x,y) = \exp\left(-\frac{x^2}{\sigma_x^2} - \frac{y^2}{\sigma_y^2}\right)\cos(kx - \phi)$$

(σ's define fatness of RF in x & y axes; cos: spatial frequency; ϕ: spatial Phase)

(Symmetrical w/ zero ϕ)

$k = 2\pi/\text{wavelength}$

  - But this is just the spatial RF. Next time, its temporal aspects...

Tues 12 Feb

- Stimulus → Response (Encoding) produces spike-rate based representation
- Decoding problem is how to interpret state of the world from representation
  - To prevent problem of underestimating brain's schema (because you use a dumb decoding schema) find mathematically optimal model to place proper limits on task
- Easiest approach; use just two stimuli and measure discrimination
  - Can you reconcile neural responses and behavioral responses

Cricket Jumping Exmpl

- Cricket judges angle of incoming wind recorded from $\approx 10^3$ hairs
- Hairs are hinged at various orientations and axons project to small ganglion w/ only four neurons

- Vector Interlude
  - 2D vector can be broken down to component sizes along x- and y-axes
  - Decoding problem would be recreating vector from component values. Create unit vectors $\hat{x}$ & $\hat{y}$ and multiply by the coefficients

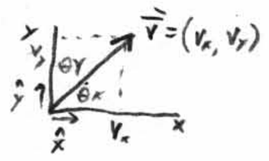

  - RESTRICTIONS: axes must be orthogonal
    must be able to compute cos(
    $\therefore V_x = |\vec{V}| \cos(\theta_x)$
    $V_y = |\vec{V}| \cos(\theta_y)$

- The cricket does this sort of vector decomposition, establishing a (tipped) x & y coordinate plane
  - Since it can't represent negative rates, it's really using four axes
  - Each of the 4 ganglion cells fires proportionally to vector projected onto its axis

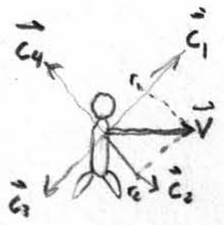

$r_1 \propto \cos(\theta_1)$ r
$r_2 \propto \cos(\theta_2)$
⋮

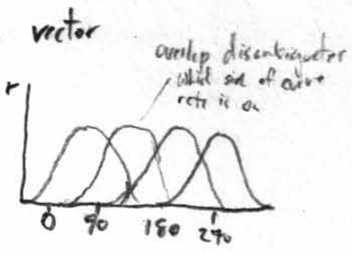

  - Cricket response peaks are all 90° apart, so it really has set up a coordinate scheme
  - Now that we understand the encoding, we know how to reconstruct $\vec{V}$ from the rates r_1, r_2, r_3, r_4:

$$\vec{V} \propto r_1\vec{C}_1 + r_2\vec{C}_2 + r_3\vec{C}_3 + r_4\vec{C}_4$$

  - Since there's some noise, there's variation in the r_n measurements, so the decoded vector may be ±6°

Monkey Reaching

- Monkey cells also do cosines, but in 3D. Rather than using 6 cosines (w/ no negative numbers) only 3 are needed relative to some baseline firing rate to define negative lower-bound

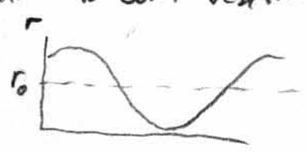

$$\vec{V}_{real} = \sum_c (r_i - r_0)\vec{C}_i$$

↑ This is a way to do the decoding, but it's a very partial sum since we're not seeing all the cells. So the fact that they get larger estimation errors doesn't necessarily mean the system can't do any better than the recorded error
  - Using optimal method shows that monkey really can do well, even w/ a small # of cells

Probability

r → response s → stimulus

P[r] = (Prob. of response r occurring)

P[r | s] = (Prob. of getting resp. r to stimulus s)
- Can be measured by distribution of responses to given stimulus exposure
- This entirely describes the encoding scheme

P[s | r] = (prob. of a stimulus having occurred to elicit response r)
- Both a measure of your best estimate, AND how confident you are about it
- This is the decoding scheme

P[s] = (prob. of stimulus s appearing)
- Knowledge of likelihood of certain events can be used as well:

P[r, s] = (prob of getting response r, & having s appear)

$= P[r|s] \cdot P[s]$ ← These expressions must be equal, So we can derive:

$= \underbrace{P[s|r]}_{\text{Prob of getting both together}} \cdot \underbrace{P[r]}_{\text{Prob. of one appearing}}$

$$P[s|r] = \frac{P[r|s] \cdot P[s]}{P[r]}$$ Baye's Theorem

- Both P[r|s] and P[r] can easily be extracted from the data, but P[s] is trickier. Thus the interest in finding the statistics of natural images

- Baye's Theorem informs optimal decoding methods. E.g.,
  - Sest – most likely
  - Utility – rate of error

Fri 15 Feb

- Math -

- An integral is just a continuous dot product

$$\int_0^\infty d\tau\, D(\tau)\cdot S(t-\tau)$$

$$= \Delta\tau\, D' \cdot S(i:-1:\text{length}(D)+1)$$

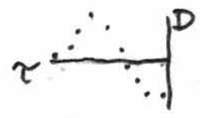

- Bayesian Statistics

P[r|s]: Prob of response r to stimulus s

P[s|r]: Prob that stimulus s caused response r

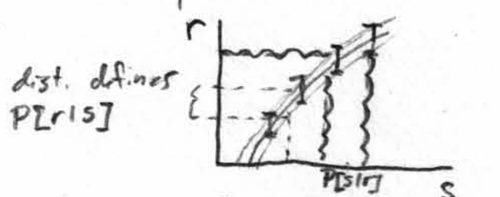

Resp. will vary over for same stimulus.

Bayes Theorem:

$$P[s|r] = \frac{P[r|s]\cdot P[s]}{P[r]}$$

1) MAP decoding

- If you know P[s|r], choose s to maximize this quantity; the most likely stimulus.

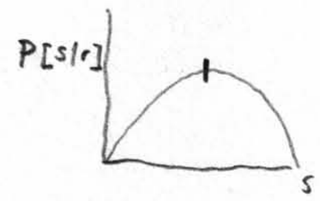

2) Max Likelihood

Choose s to maximize P[r|s]

3) Bayesian Methods

- Cost function measures how costly it is to screw up, e.g. try to minimize $\langle (S_{estimate} - S)^2 \rangle$

Max likelihood is most generally applicable.

- From a set of responses you generate a set of stimulus estimates and compare to the true stimulus (S)

$r_1, r_2, r_3 \rightarrow S_{est1}, S_{est2}, S_{est3}$

- Calculate variance of estimates as compared to true value
  → average the squared difference between S_{est} and S

$$\sigma^2_{est}(S) = \langle (S_{est} - S)^2 \rangle$$

Unbiased Estimator means $\langle S_{est} \rangle = S$ (true value)

Theorem:

If you have an Unbias Estimator, then it can't possibly do better than the Max Likelihood method:

$$\sigma^2_{est} \geq \sigma^2_{ML} = \frac{1}{I_F(S)}$$

If Fisher Information is large, you can more accurately estimate the maximum likelihood. This sets an upper bound on how well you can do.

- $I_F(S)$ is computed based on P[r|s]

- Most information is given when slope is largest b/c changing S has biggest effect on r

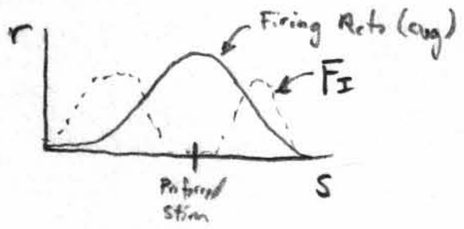

- Data showing it's easier to estimate angles of bars of diff. line lengths. Psychophysical data shows peoples' estimates v. closely approximate σ_{ML}

Discrimination

- Only two possible values for S: +, −
- MT is tuned to large-scale motions (i.e. perceives optic flow)
- Stimulus is % coherence in random moving dot field
  100% means all dots drift in same direction
  50% means half drift upward, others drift randomly
- Neurons prefer a direction and dislike the opposite direction
- So, based on r → infer S, and also have animal report S
- Task gets hard at ≈10% coherence. But once success rate gets close to chance, it stops putting in the effort and starts guessing.
- Next, compare discrim. behavior to recording from individual neurons
  - Cell firing rates for + stims are significantly higher than for − when coherence ≥ 10%
  - As coherence goes down, the + and − rate distributions converge

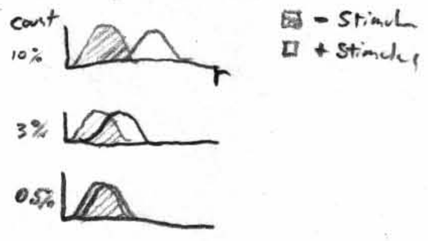

  - Agreement b/t single neurons & discrim. is very good (only on avg. though, not trial to trial)
- So, how do you do optimal decoding based on histograms?
  - They are ▨: P[r|S−]
    □: P[r|S+]
  - One approach is to choose a threshold: r ≥ thresh → +
    r < thresh → −
- Prob we'll give the right answer is $P[r \geq thresh | +] : \int_{thresh}^{\infty} dr\, P[r|+]$

True Stim	P[correct]		P[incorrect]	
+	P[r ≥ thresh | +]	β	P[r < thresh | +]	1−β
−	P[r < thresh | −]	1−α	P[r ≥ thresh | −]	α

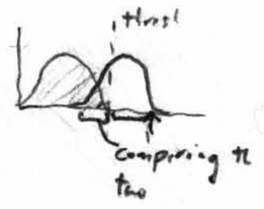

β = hit rate
α = false alarms

- Both α and β depend on where you put the threshold
  - If it's too high β will be small, but α will be 0; all correct
  - If it's too low, you'll get all + cases, but miss − cases

ROC (Receiver Operating Characteristic) Curve compares α + β given thresh (z)

- Area under curves corresponds to animal's performance under conditions (or, can be no better than area)

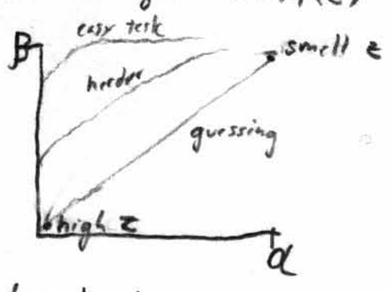

Performance of ideal observer of single cell, actually doesn't do worse than animal, which has thousands more cells to pool from.

- Direction-Selectivity lies in columns, so you could stimulate cells and improve or hurt performance based on appropriateness of column for stimulus.
  - After you stop stimulating, animal guesses in other direction b/c they're worried about having too many [UP]'s in a row.

tues 26 feb

Biophysics of spikes

- Function is mostly accounted for in our firing-rate models since most important signals are spike based anyway. } uses extracellular data
- Mechanism isn't addressed though. Simple + Complex cells are functionally v. different, but there's only a smoothly varying distribution of cell types. So where does the difference come from? } uses intracellular data

Two aspects of Cell Modelling

1) Morphology - cable theory. How much do distant dendritic signals affect the cell membrane potential

2) Electrical Properties - Conductance-based models. What currents go where? How are spikes generated?

Conductance Models

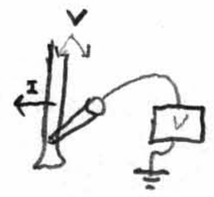

- Voltage is relative to extracellular: V is the potential across the membrane.
- Simple, single-compartment model imagines spherical cell w/ only a single voltage at all points on membrane
  - This is unrealistic, of course; somatic + dendritic potentials are different both in scale and time scale (faster + larger response to nearby stimulation)
- Voltages com from ions (since there are so many acids in cell giving off H^+, lots of negatively charged particles left over)
- The abundance of negative charges causes repulsion so the - ions distribute along membrane (it's as far apart as they can get).
- Negative charges on inside attract + ions on outside
  - This is a capacitor, The charges can't actually get through the insulator

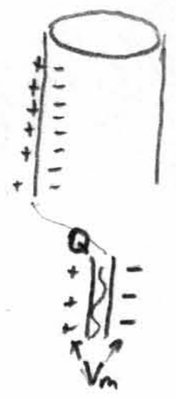

- If we could count the charges, Q, we can find the Voltage b/c of the relationship: $Q = CV$
  (C: capacitance (a function of the insulator))
- Biggest factor in determining C is size of membrane. More membrane area, A, means more space for charge to build up upon:

$$C = c_m A$$

(A: cell's surface area; c_m: Capacitance per unit area (typically 10 nF/mm²))

- Typical cell has $A = .01$ to $.1$ mm²

  $C = .1$ to 1.0 nF

- Farads are the constant that take you from coulombs of charge to volts: 1 coulomb = 1.0 F · 1 Volt

  e.g. Typical cell has .1nF * -70mV, From this we get -7×10^{-12} coulombs (or 10^9 singly charged ions)

- What we care about though isn't charge, but <u>current</u>, but current is just charge moving (current is derivative of charge: $\frac{\text{coulombs}}{\text{sec}}$ = amp's) $I = \frac{dQ}{dt}$
- So change $Q = CV$ to its time derivative:

$$I = C\frac{dV}{dt}$$

- We can use this relationship to guess how much current is necessary to cause an arbitrary voltage deflection, e.g.,

$$\left(\frac{1\,mV}{1\,ms}\right) 0.1\,nF = 0.1\,nA$$ } an epsp-scale event

- Or, in the case of an action potential is much larger:

$$\frac{100\,mV}{1\,ms} \cdot 0.1\,nF = 10\,nA$$ } ap scale event

(Total realistic range)

Typical Range Summary

$c_m = 10\,nF/mm^2$

$C = 0.1 - 1\,nF$

$V \approx 1 - 100\,mV$ (actual range $-90\,mV$ to $+50\,mV$)

$I \approx 0.1 - 10\,nA$

Where does range of Voltages come from?

- Ion channels allow cell to modify its own conductance
- So why is it in this specific range of tens-of-millivolts?
  - Temperature

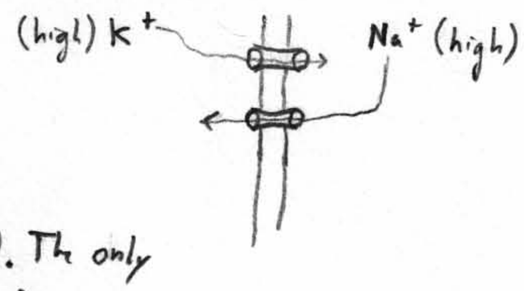

- Sodium will rush down its strong voltage gating
- Potassium doesn't want to leave (flow up electrical potential grating). The only reason they move out is because of their thermal energy (see Boltzmann constant)
  - So how much energy does it take to get an ion out of the cell? It's on order of $q \cdot |V|$ and $q|V| \approx k_B \cdot temp$
    (q: charge of ion)
  - So $V_T = \frac{k_B T}{q}$
    (V_T: temperature-dependent voltage)
  - But Boltzmann constant is coefficient of a single charge. More commonly you see V_T derived from <u>moles</u> of charge:

$$V_T = \frac{k_B T}{q} = \frac{RT}{F} \approx 25\,mV$$

- Since this is all the kick you can get, we need to work with range of voltages on same order as 25mV so we don't either get swamped by its noise, or not get significantly pushed by it (if we're using voltages that are too big).

Reversal Potentials

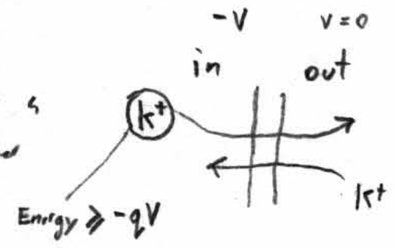

- When a K^+ channel is opened, there's a probability that K^+ will leave and another that one will leave

$$P(\text{into cell}) = [K^+ \text{ outside}]$$

$$P(\text{out of cell}) = [K^+ \text{ inside}] \cdot e^{-qV/k_B \cdot T}$$

($e^{-qV/k_B \cdot T}$: Probability that ion has enough thermal energy to get across V gradient (Prob that Energy $\geq -qV$))

- (In general, probability of something getting enough energy from the environment to reach some threshold is $e^{-threshold/k_B Temp}$)
- The point where K^+ will equilibrate is the Nernst potential when nothing will move:

$$[outside] = [inside] \cdot e^{qV/k_B T}$$

fri 1 March

Integrating Differential Equations

- $C\frac{dV}{dt} = I$ current charges up cell, voltage change is proportional to capacitance

- $C_m\frac{dV}{dt} = \frac{I}{A}$ same, but divided by area

$C = C_m \cdot A$ (unit capacitance and surface area)

$I = C\frac{dV}{dt}$

$I_{electrode}$ → Inward is positive

$I_{membrane}$ → Inward is negative

- Current is membrane current + electrode current

$$\frac{I}{A} = \frac{I_m + I_e}{A}$$

- E is equilibrium potential — voltage at which inward + outward current through the channel cancel one another
  - High external $[Na^+]$ pushes Na^+ into cell via concentration gradient. Also, the cell's negative charge attracts the ions. To negate inward push, raise the potential: $E_{Na} = +50mV$
  - High internal $[K^+]$ pushes them out. To keep them in, lower potential:

$E_K = -80mV$

$E_{Cl^-} = -60mV$

$E_{Ca^{++}} = 150mV$

- Increasing a given ion's conductance will pull V_m toward the ion's equilibrium potential

- We know current goes through zero at $V_m = E_{ion}$. So presume it's a straight line

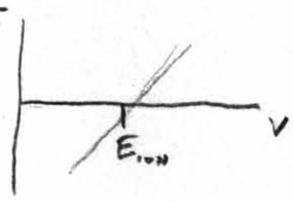

- So $I_{channel} = (V - E_{ion})$, scaled by conductance (the goodness of the channel & the number of channels)

  eg. $I_{Na} = -\overline{G_{Na}}(V - E_{Na})$

  $I_K = -\overline{G_K}(V - E_K)$

- Since these $\overline{G_{ion}}$ include cell's area. Divide it out to get the conductance per unit area:

$$\frac{I_m}{A} = -g_{Na}(V - E_{Na}) - g_K(V - E_K) - \dots$$

So: $c_M\frac{dV}{dt} = -g_{Na}(V - E_{Na}) - g_K(V - E_K) - \dots + \frac{I_e}{A}$ ← only electrode current has to account for area

- But where does g_{ion} come from?

  Depend on voltage, time, internal calcium, external neuro-modulators, ligands...

For first model, consider resting state and don't perturb it too much (as if g's are constant)

- Thus you only have one conductance g_{Leak} and one reversal pot. E_{leak}

$C_m\frac{dV}{dt} = -g_L(V - E_L)$ $\quad g_L = g_{Na} + g_K + g_{Cl} \dots$

$E_L = \frac{g_{Na}E_{Na} + g_K E_K}{g_L}$

- Also consider current injection:

$$C_m\frac{dV}{dt} = -g_{leak}(V - E_{leak}) + \frac{I_e}{A}$$

How to get these values for the model?

1) don't inject any current and wait for voltage to settle

so, $0 = -g_L(V - E_L) + 0$

Then E_L = resting potential

2) To get g_L, inject some current, and find new E_L

$g_L(V - E_L) = \frac{I_e}{A}$

$V = E_L + \frac{I_e}{g_L \cdot A}$

$g_L \cdot A = G_L$ (the total cell conductance) $= \frac{1}{R_{membrane}}$ $\quad R_m \approx 10 - 100\ M\Omega$

Imagine injecting 0.1 nA and $R_m = 1.0 \times 10^6$

then $I_e R_m = 10^{-3}V = 10\ mV$ via Ohm's Law, $V = IR$

$C_m\frac{dV}{dt} = -g_L(V - E_L) + \frac{I_e}{A}$

$\frac{c_m}{g_L}\frac{dV}{dt} = E_L - V + \frac{I_e}{g_L A}$

↑ this is membrane time constant $\tau_m = R_m C$

(*) $$\tau_m\frac{dV}{dt} = E_L - V + R_m I_e$$

↑ membrane time constant, ↑ resting Potential, ↑ Input Resistance, ↑ Injected Current

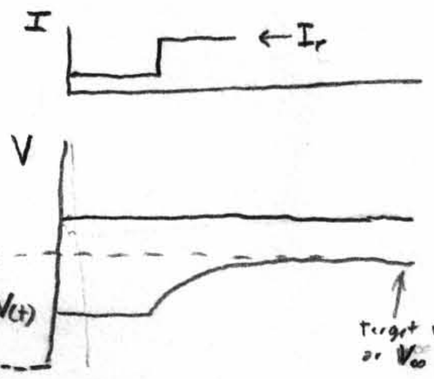

- So we know when V is going when we inject some amount of current. What are the dynamics of getting there?

- Exponential decay again. For each time step, distance is divided by some constant factor.

- At given time, distance is $V(t) - V_\infty$

- To get distance at next time step, multiply by the factor:

$K \cdot (V(t) - V_\infty) = V(t+1) - V_\infty$

where k is diff in time over time constant:

$e^{-\frac{\Delta t}{\tau_m}}(V(t) - V_\infty) = V(t+\Delta t) - V_\infty$

In normal form:

$$V(t+\Delta t) = V_\infty + e^{-\frac{\Delta t}{\tau_m}}(V(t) - V_\infty)$$

Other formulations of same equation:

$V(t_2) = V_\infty + e^{-\left(\frac{t_2 - t_1}{\tau_m}\right)}(V(t_1) - V_\infty)$ $\quad t = t_1,\ t + \Delta t = t_2$

$V(t) = V_\infty + (V(0) - V_\infty)\cdot e^{-\frac{t}{\tau_m}}$ $\quad t_1 = 0,\ t_2 = t$

- To prove equation satisfies (*), take its time derivative:

$\frac{\partial V(t)}{\partial t} = \frac{-1}{\tau_m}(V_0 - V_\infty)e^{-\frac{t}{\tau_m}}$

$\tau_m\frac{dV(t)}{dt} = -(V(t) - V_\infty) \dots$

[eureka]

- Returning to (*), we want to vary $I_e \to I_e(t)$, so we'll numerically solve the resulting equation since the integrals are too evil

  - So pick a small Δt of current and approximate I_e to be constant at that time.

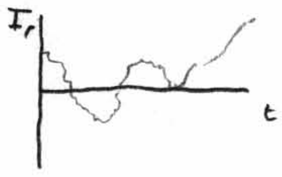

  - Then plug in the constant I_e

$$V(t+\Delta t) = E_L + R_m I_e(t) + \left(V(t) - E_L - R_m I_e(t)\right)\cdot e^{\frac{\Delta t}{\tau_m}}$$

- Integrate & Fire Model adds wrinkle of Action Potentials
  - Follow previous model, then if V reaches $V_{threshold}$ steps in a spike-and-return-to-reset pattern and return to model
    - Works since AP duration is so short
  - Since we can compute difference in time between V_{rest} getting to V_{thresh}, we can now determine the firing rate.

$V_{thresh} \approx -50mV$

$V_{rest} \approx -80$ to $-60mV$

Tues 5 March

Integrate & Fire Model:

- $\tau_m \frac{dV}{dt} = V_{rest} - V + RI$
- If $V \gg V_{thr}$, spike + reset

[Means $V \to V_{rest} + RI$ and does so exponentially with a rate set by τ_m]

Parameters

$\tau_m = 20ms$
$V_{rest} = -65mV$
$R = 10-100\ M\Omega$
$V_{th} = -50mV$
$V_{reset} = V_{rest}$

- So far model doesn't take a synaptic input, only injected current
- We can analytically determine the firing rate from this equation

$V_{steady\ state} = V_{rest} + RI$ if $V_{ss} < V_{th}$

$I_{threshold} = \frac{V_{th} - V_{rest}}{R}$

(amount of current needed to charge membrane up to threshold)

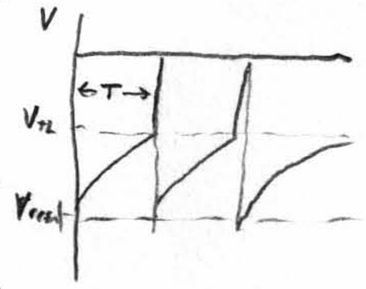

- If current is constant, time from rest of one spike to initiation of the next is also constant:

The time to go from V_{reset} to V_{th}

$V(t) = V_{rest} + RI + (V(0) - V_{rest} - IR)e^{-\frac{t}{\tau_m}}$

$V_{th} = V_{rest} + RI + (V_{reset} - V_{rest} - IR)e^{-\frac{T}{\tau_m}}$ ← T tells us time from reset to threshold

$\frac{V_{th} - V_{rest} - RI}{V_{reset} - V_{rest} - RI} = e^{-\frac{T}{\tau_m}} \Rightarrow \ln\left(\frac{V_{th} - V_{rest} - RI}{V_{reset} - V_{rest} - RI}\right) = \frac{-T}{\tau_m}$

$$T = \tau_m \cdot \ln\left(\frac{V_{reset} - V_{rest} - RI}{V_{th} - V_{reset} - RI}\right)$$

then rate is $\frac{1}{T}$

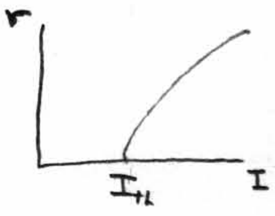

- Note that: if no spikes occur, you get ∞

Synaptic Inputs

- The I in the canonical equation is unrealistic bc it only accounts for injected current
- With synaptic input, current comes from changing ion channel conductances
  - To gauge synapse's activity, we need to know synapse's reversal Potential and conductance:

$I_s = g_s(V - E_s)$

E_s { ex-glut: $E_s \approx 0mV$ (K^+ & Na^+); in-$GABA_a$ $E_s \approx -80mV$; in-$GABA_b$ $E_s \approx -60mV$

- To model g_s dynamics, instantaneously move $g_s \Rightarrow g_s + g_{unitary}$ at spike time. Then exponentially decays to zero:

$\tau_s \frac{dg_s}{dt} = -g_s$

$g_s(t + \Delta t) = g_s(t) e^{-\frac{\Delta t}{\tau_s}}$ τ_s { AMPA: $\tau_s \approx 2ms$; NMDA: $\tau_s \approx 100ms$; $GABA_a$: $\tau_s \approx 5ms$; $GABA_b$: $\tau_s \approx 80ms$

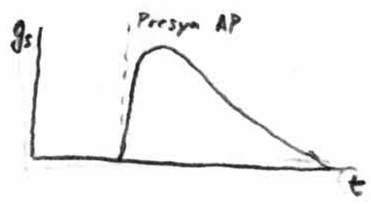

Unified Equation:

$$\tau_m \frac{dV}{dt} = V_{rest} - V - g_s R_m (V - E_s)$$

↑ Need to include membrane capacitance since we divided the left side by C_m

Improving Integrate & Fire

Two crude approximations:

1) APs are phony (but since they occur so quickly, this isn't such a problem)
2) assume constant conductances in sub-threshold regime

Problems:

- Lack of saturation with increasing current (ignores limits on firing rate)
  - One solution is to hold V_m at reset to mimic a refractory period
- Lack of Spike Rate Adaptation
  - Firing rate will lessen over time in response to constant current
  - In real cells there are Ca^{++}-gated K^+ channels. As intracellular $[Ca^{++}]$ increases (due to spikes), the outward K^+ current mutes synaptic input
    - So add to the tail of our equation: $-g_{SRA} \cdot r_m (V - E_K)$

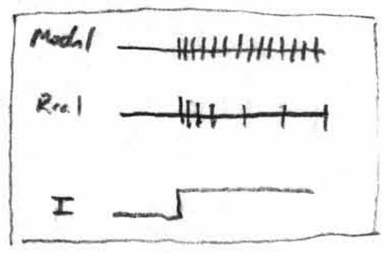

with dynamics: $g_{SRA} \to g_{SRA} + \Delta$ {on spike}

$\tau_{SRA} \frac{dg_{SRA}}{dt} = -g_{SRA}$ {to model decay between spikes}

[talk yesterday said that realistically, g_{SRA} never saturates under natural conditions]

Off Topic:

- Integrating exponential decay:
  - For equations of form

$\tau \frac{dg}{dt} = -g$

$\tau\, dg = -g \cdot dt$

$\tau \frac{dg}{g} = -dt$ } Separation of variables

$\tau \int_{g(0)}^{g(t)} \frac{dg}{g} = -\int_0^t dt$

$\tau \ln\left(\frac{g(t)}{g(0)}\right) = -t$

Hodgkin Huxley Model

$C_m \frac{dV}{dt} = I_m + I_e$ | $I_m = -\left(g_{Na}(V - E_{Na}) + g_K(V - E_K) + g_{leak}(V - E_{leak})\right)$; I_e = constant

(Question is dynamics of g_{Na} and g_K)

$C_m \frac{dV}{dt} = -g_{Na}(V - E_{Na}) - g_k(V - E_K) + g_L(V - E_L) + I_e/area$

Note: this is still of form $\tau_V \frac{dV}{dt} = A - V$

$= g_{Na}E_{Na} + g_K E_K + g_L E_L + I_e/area - (g_{Na} + g_K + g_L)V$

$\frac{C_m}{g_{Na} + g_e + g_L} \frac{dV}{dt} = \frac{g_{Na}E_{Na} + g_K E_K + g_L E_L + I_e/area}{g_{Na} + g_K + g_L} - V$

fri 8 March

Hodgkin Huxley Model

$$C_m \frac{dV}{dt} = -g_L(V-E_L) - g_{Na}(V-E_{Na}) - g_K(V-E_K) + I_e/A$$

- The conductances are the interesting part since they change in response to V_m (i.e. the channels are gated, unlike leak)
- $g_K = \bar{g}_K P$
  - $\bar{g}_K$: Maximal conductance per unit area (if every channel is open)
  - P: The fraction of channels open at any given time -or- prob. of one being open
- Now P has all the channel dynamics in it. Can be deduced from channel structure: 4 identical subunits, all of which must undergo a conformational change for the channel to be 'open'

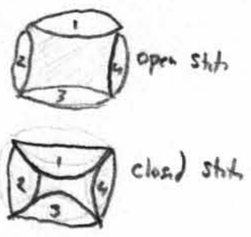

- Assumption that all subunit states are independent
  - n = Probability that a given subunit is in the 'open' state
  - $\therefore n^4$ = Probability that all subunits are 'open,' and thus the channel is open
- Now problem is to describe n. First, what are the rates of transition between state (as a function of voltage)

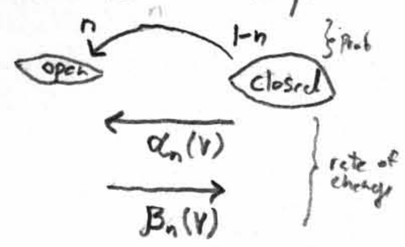

$$\frac{dn}{dt} = \alpha_n(V)(1-n) - \beta_n(V)\,n$$

- $\alpha_n(V)$: rate constant
- $(1-n)$: Prob. of finding a closed subunit
- $\beta_n(V)$: rate constant
- n: Prob. of finding an open subunit

$$\tau_n(V)\frac{dn}{dt} = n_\infty(V) - n$$

$$\tau_n(V) = \frac{1}{\alpha_n(V)+\beta_n(V)}$$

$$n_\infty(V) = \frac{\alpha_n(V)}{\alpha_n(V)+\beta_n(V)}$$

n will go to n_∞ with a rate determined by $\tau_n(V)$:

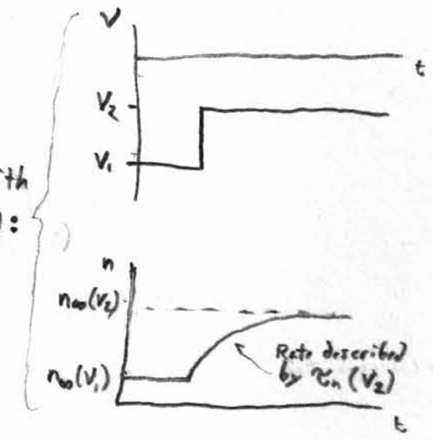

- So to finish, what are α_n and β_n (or equivalently $\tau_n(V)$ and $n_\infty(V)$):
  - We know α_n increases as V increases, and that the rate is pretty slow (on molecular terms)
    - likely reason for the slowness is an 'energy barrier'
    - When cell is depolarized, there is still a small hump the subunits have to get over to move to the open state
    - As V_m gets more + more depolarized, the energy barrier gets smaller (as the hump disappears)
      - Thus less energy must be borrowed from the environment to make the transition
  - There are well established rules for borrowing energy as a function of environment's temperature:

$$\text{Prob}(\Delta E) = e^{-\frac{\Delta E}{kT}}$$

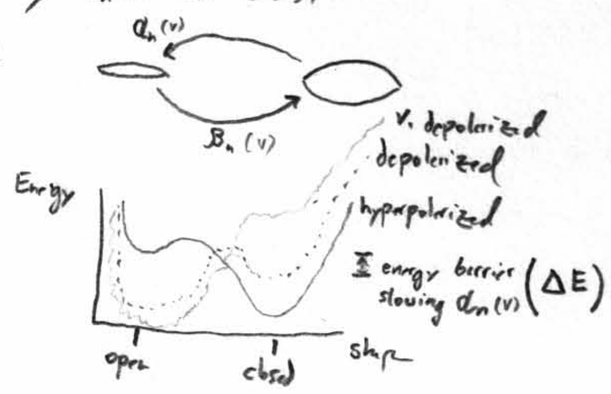

So,

$$\alpha_n(V) \propto e^{-\frac{\Delta E}{kT}} = e^{-\frac{A+BV}{k_B T}}$$

| When higher voltages make it easier to borrow ΔE (?)

- Thus α and β are just exponentials with different signs, but this is based on the simplifying assumptions of only 2 states per subunit, etc.
- Accordingly most units are closed at hyperpolarized state, most open at depolarized, with some transition in between
- At edges, energy barrier is small, so transitions are fast. In middle it's larger so slower transitions

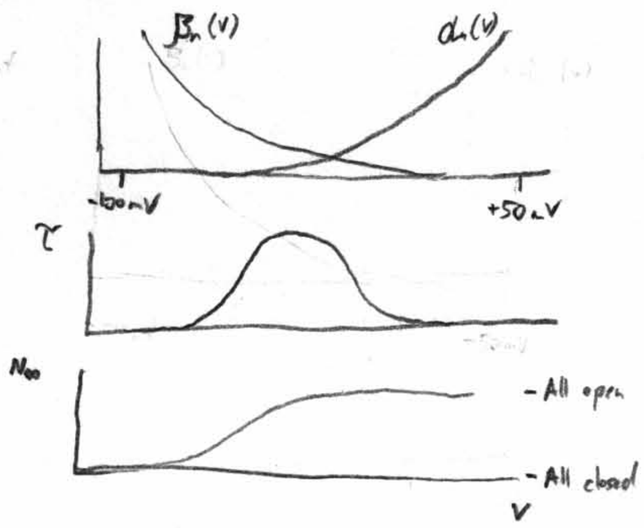

Sodium Channel

- Applied same thinking to Na channel, but Na structure has 4 subunits, where they guessed there were 3 (the subunits aren't quite independent)

$$\bar{g}_{Na} P = \bar{g}_{Na} m^3$$

- But Na acts like there are 2 doors deciding whether the channel is opened:

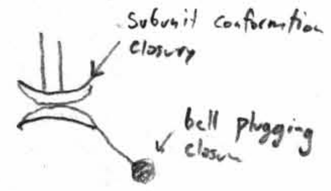

$$g_{Na} = \bar{g}_{Na} m^3 h$$

$$\tau_m(V)\frac{dm}{dt} = m_\infty(V) - m$$

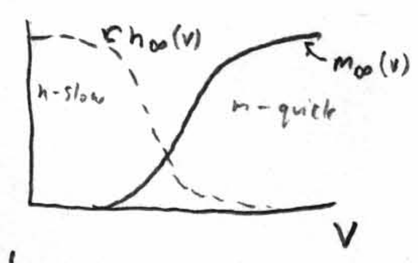

- Since h + m kick in inversely, the only way there's ever any current is that m is much more rapid in its activation than h's inactivation
- So Na is a transient channel:

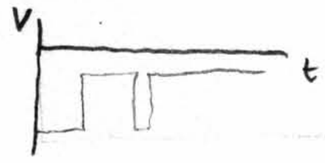

$$Na \sim m^3 h$$

- $m\uparrow$: activation
- $m\downarrow$: deactivation
- $h\uparrow$: inactivation
- $h\downarrow$: disinactivation

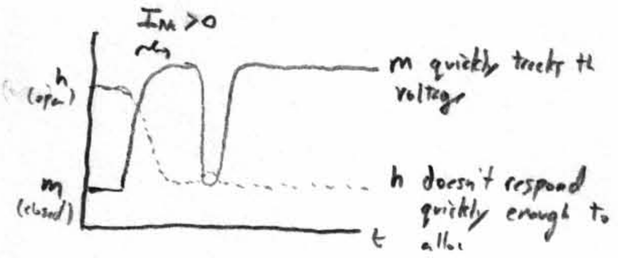

- Current only passes when both m and h are at high values (their open state)
- Hyperpolarizing the cell increases h, so a subsequent voltage spike will cause more current than would have been w/ less prior depolarization:

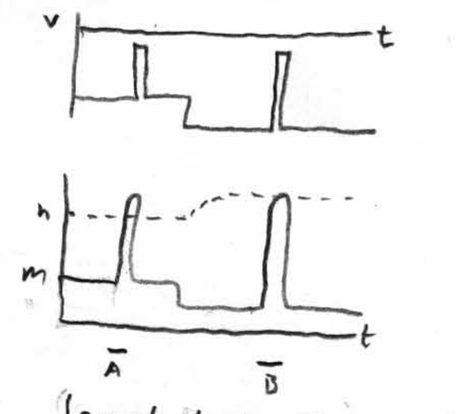

(current at time B is larger)

So final equation is:

$$C_m \frac{dV}{dt} = -g_L(V-E_L) - \bar{g}_{Na} m^3 h (V-E_{Na}) - \bar{g}_K n^4 (V-E_K) + I_e/A_{area}$$

$$\tau_n(V)\frac{dn}{dt} = n_\infty(V) - n$$

$$\tau_n(V) = \frac{1}{\alpha_n(V)+\beta_n(V)}$$

$$n_\infty(V) = \frac{\alpha_n(V)}{\alpha_n(V)+\beta_n(V)}$$

tues 12 march

Hodgkin Huxley Continued

$$C_m \frac{\partial V}{\partial t} = -g_L(V-E_L) - \bar{g}_{Na} m^3 h (V-E_{Na}) - \bar{g}_K n^4 (V-E_K) - \frac{I_e}{Area}$$

Behavior

- V_{rest} is ≈ -65 mV (slightly off from E_L b/c of small K^+/Na^+ currents)
- At threshold, there is a sharp Na current which quickly inactivates
  Recovery (deinactivation) happens gradually but only once V_m repolarizes
  - Remember; current is entirely due to the h value changing more slowly than m
- Potassium has a gradually increasing outward current, followed by a tail where there is still some small current shortly after V_m falls

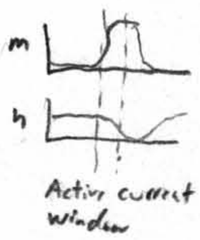

Currents

- These two active conductances do a good job characterizing the squid giant axon, but there are many other currents in more complicated neurons
- Currents follow stereotyped form:

$$I_{conductance} = \bar{g} M^p \cdot H^q \cdot (V-E) \qquad \tau_m \frac{\partial M}{\partial t} = M_\infty - M$$

M: activation variable; H: inactivation variable; (V−E): Driving force

$$\tau_n \frac{\partial H}{\partial t} = H_\infty - H$$

1) Persistent Currents ((activation only) – q is always zero

2) Transient Currents have activation & inactivation

Two types of transitions from not-firing to firing

- Type I

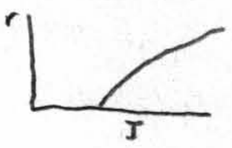

- Type II

Useful only for on-off signaling, like fight/flight in squid giant axon

3) Anomalous currents are inactivation only (e.g. H, anomalous rectifier)

4) Non-voltage gated currents

e.g. Ca^{++}-dependent currents when $M_\infty(V) \rightarrow M_\infty(V, [Ca^{++}])$

Channels

- All these currents are really based on how we think the channel works
- Remember K^+ current presumes all four subunits must change to 'open' state for channel to be open

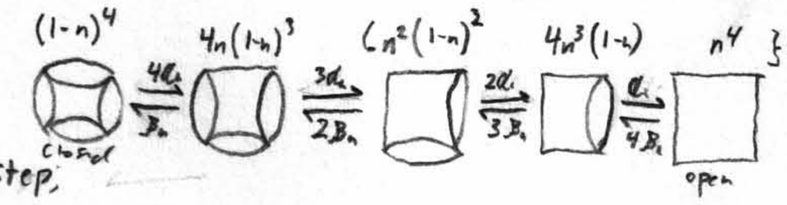

- Procedure: At each time step,
  1) Examine state
  2) If "closed" open w/ probability $\alpha_n \Delta t$
  3) If "open" close w/ probability $\beta_n \Delta t$
- You can measure these probabilities by poking onto membrane and measuring current in response to voltage

Sodium Current ($P = m^3 h$)

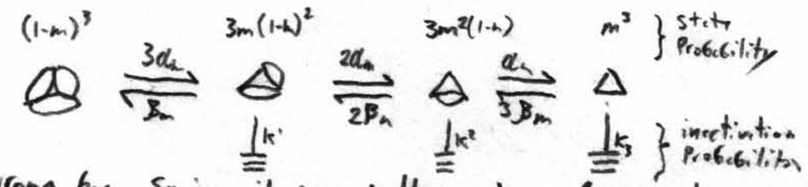

- H+H got h-inactivation wrong by saying it was voltage-dependent. Not possible b/c inactivating ball is on inside of cell, ∴ no voltage
- α and β should be exponential based on stochastic model

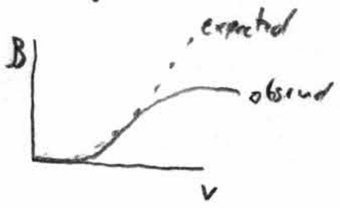

- Saturation of observed β shows h is not voltage dependent, but <u>state-dependent</u>

- The inactivation probability is just a constant for any given activation state, so $\beta_h = k_1 3m(1-m^2) + k_2 3m^2(1-m) + k_3 m^3$
  - H+H's mistake is b/c $m \rightarrow m_\infty$ <u>very</u> quickly, so if you substitute $m_\infty(V)$ for our state-dependent m, we are suddenly back to voltage-dependence
  - Since m_∞ function is identical to the 'β_h' that they observed (w/ saturation), perhaps they should have seen it

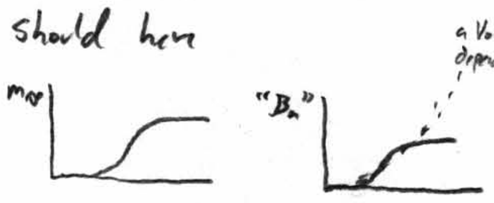

Cell Structure & Compartmental Models

- Long axons will not have same voltage down the length. Instead there is a rapid fall off as you move away from the current source

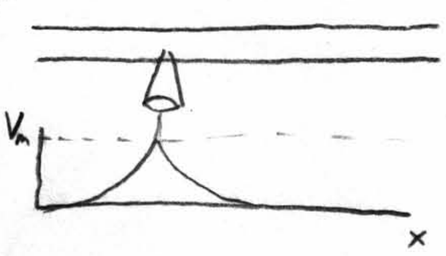

- By ohm's law, we know that difference in voltage between two points is:

$$V_1 - V_2 = I_L R_L$$

I_L: Longitudinal current; R_L: Longitudinal resistance

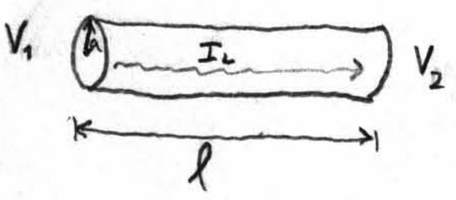

$$R_L = \frac{r_L l}{\pi a^2}$$, where r_L is a universal constant: $r_L \approx 1-3\ k\Omega/mm$

- For example, if $a = 2\mu m$ and $l = 100\mu m$, $R_L \approx 8 M\Omega$; letting you determine whether there is a significant voltage drop between points.
- So where does the decaying charge go?
  1) Charges in capacitance, like when dendritic current reaches the much larger volume at the soma
  2) Across the membrane in the cell itself (i.e., leak)

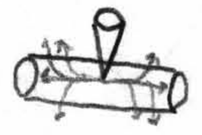

  - This leak can be calculated based on the surface area:

$$G_m = g_m A \qquad r_m = 1 M\Omega/mm^2$$

$$R_m = r_m/A = \frac{r_m}{2\pi a^2 \cdot l}$$

- Interaction b/t longitudinal resistance and membrane resistance encapsulated in λ, Electrotonic length constant:

$$\lambda = \sqrt{\frac{r_m a}{2 r_L}}$$

- So you know that the signal will go λ (≈ 1mm) before it seriously begins to degrade

fri 15 march

The Cable Equation

- Injected current charges voltage which decays exponentially with distance
- Length Constant is when longitudinal resistance and membrane (leak) resistance are equal
- To double λ, you need to make cable 4x as big (gets expensive quickly)
- Can use cable equation to gauge compartment size (make them small enough so you don't need to worry about decay within one.

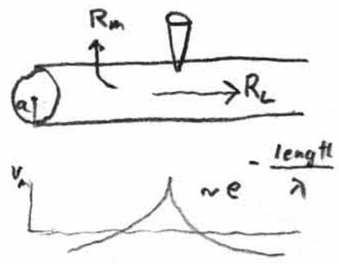

$r_m = \Omega \cdot mm^2$

$R_m = \frac{r_m}{2\pi a \lambda}$ circumference

$r_L = \Omega \cdot mm$

$R_L = \frac{r_L \lambda}{\pi a^2}$ area

$\lambda = \sqrt{\frac{r_m a}{2 r_L}} \approx 1m \quad \tau = r_m C_m$

- We can String compartments together and model each as an H.H. neuron, but each with an additional current from neighboring compartments: I_L

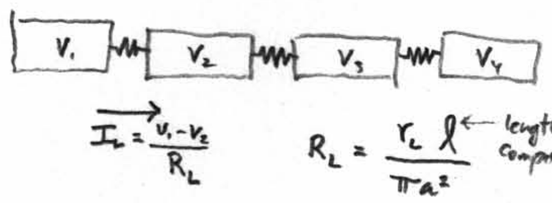

- Now we just need to know R_L, which we can get from the longitudinal resistance over the distance between compartments
  - NOTE: don't need to include R_m, because it is already accounted for within H.H. calculation for each compartment
- Action potentials go faster as axon size increases (b/c you can raise the voltage of neighboring compartments more easily w/ lower R_L)
- AP speed is fully a function of length constant (how far ahead depolarization can spread) and capacitance (how quickly those depolarizations charge up): $V \approx \frac{\lambda}{\tau_m} \approx \sqrt{a}$ (for unmyelinated)
  - So one way to speed up is to boost the radius.
  - More economical approach is insulation (lowering r_m)
- Myelination makes the axon model more like a series of compartments (b/c membrane currents are cut off in between nodes of Ranvier)

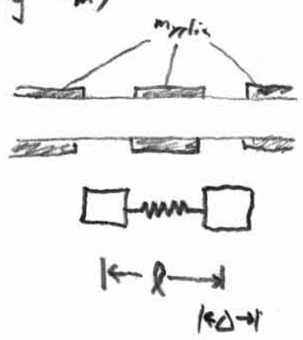

$V = \frac{\ell}{\tau_{compartment}}$ where $\tau_c = C_m \cdot R_L$

$C_m = c_m \cdot 2\pi a \cdot \Delta$

$R_L = \frac{r_L \ell}{\pi a^2}$

$\therefore \tau_c = \frac{c_m r_L \cdot 2\Delta\ell}{a}$

$V = \frac{\ell}{\tau_c} = \frac{a}{2 c_m r_L \Delta} \propto a$

So we now only need to double the area to double the speed (not quadruple it as w/ unmyelinated)

- There is an optimal ratio of radius to myelination width, which is precisely what nature has provided us
- Multi-compartment models allow us to simulate entire cell, w/ compartments chosen to be biggest chunks that could reasonably have a common voltage
  - e.g., Soma would all be one compartment, dendrite + axon would have bunches in series
- Still big debate on what sorts of conductances + processes there should be in dendrites

Synapses

- In H.H. model, post-synaptic side accounted for as just another current
- Presynaptic side more complex b/c of failures (since a Presynaptic AP doesn't guarantee a postsynaptic current)
  - Each presynaptic terminal only transmits with a Release of ≈ 0.1 on average
- However, P_{rel} varies over time which acts as a switch to silence the cell

Two Processes which affect P_{rel}: Depression + Facilitation

- Under low-activity conditions, P_{rel} is near close to 1.0
- As you increase activity P_{rel} will increase or decrease, then decay back to baseline when activity dies down
- No activity: $\tau_p \cdot \frac{dP_{rel}}{dt} = P_o - P_{rel}$

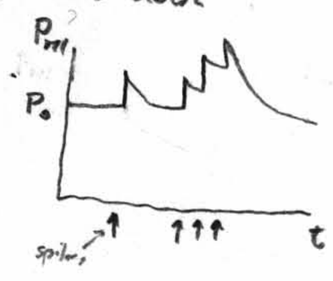

- When presyn. spike:
  - Facil: $P_{rel} \rightarrow P_{rel} + f_F(1 - P_{rel})$

So when a sequence of spikes come in, the probability will increase; could be used for burst-detection, transmitting last spikes faithfully, but ignores earlier ones

- Depression: $P_{rel} \rightarrow P_{rel} - f_D \cdot P_{rel}$

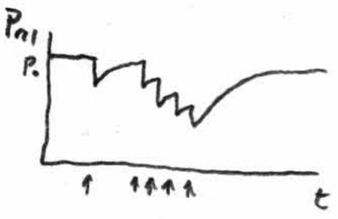

This one ignores extended spike trains but will react to onset of spikes after a period of inactivity

- When input is at a constant rate, P_{rel} will reach equilibrium based on frequency: $r - \langle P_{rel} \rangle \sim \frac{1}{r}$
  - So prob of release is inverse of incoming rate, meaning its own output is barely different under 50 Hz stimulation than under 100 Hz stimulus
  - When rate changes, the before and after output are about the same. But in the transition, the output jumps: $\frac{1}{r} \cdot (r + \Delta r)$ since new input rate is divided by older, smaller depression factor

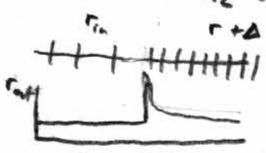

- So neuron is reporting ratio of rate increase, but without caring about absolute values
- Since they're only reporting unpredicted portions of input signal, this is effectively decorrelating the stimulus.

fri 22 march

Information Theory

1) How interesting is a spike train? (Entropy)
2) How much is the ST telling us? About what? (Mutual Information)

Quantifying "unexpectedness"

- A given event is interesting if a) it occurred, and b) it was low-probability:
  ↑ Surprise ⇔ P[r]↓

But how to quantify this relationship?

r ≈ response
s ≈ stimulus
P[r] = Prob. of getting response r

- If there are two independent spike trains, the probability of seeing the event in both is $P[r_1, r_2] = P[r_1]\,P[r_2]$ — the product
  - The 'surprise' of both happening is the sum of the two surprises:

$$\text{Surp}(P[r_1, r_2]) = \text{Surp}(P[r_1]\,P[r_2]) = \text{Surp}(P[r_1]) + \text{Surp}(P[r_2])$$

i.e., this is a logarithm, b/c $\log x \cdot y = \log x + \log y$
But since log is always increasing, it doesn't satisfy the Surp = 1/P[r] property. So use the negative log:

$$\text{Surprise}(r) = -\log_2 P[r]$$

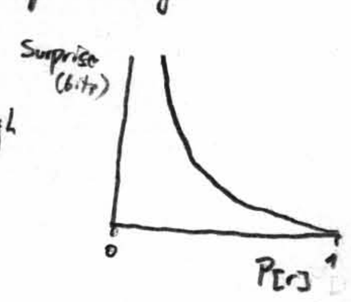

  - low prob. events very surprising, high prob. not interesting.
  - Use log base 2 so we can use bits

- Entropy looks over the entire spike train, then weights the events by the frequency — it takes the observed prob. into account:

$$\text{Entropy} = \sum_r P[r]\,\text{Surprise}[r] = \langle \text{Surprise} \rangle$$

$$H = -\sum_r P[r] \log_2 P[r]$$

e.g. Prob of getting single spike is 1.0,
Prob. of getting something else is 0:

$$H = -\Sigma\, 0 \log_2(0) - 1 \log_2(1.0)$$
$$0 \cdot 1 \quad - 1 \cdot 0$$

H = 0, so boring things are boring

($x = P[\text{getting 1 spike}]$)

$$H = -(1-x)\log_2(1-x) - x \log_2(x)$$

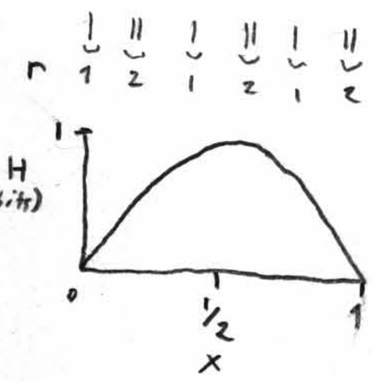

- Boring when entirely one, or entirely two. Individual events encode the maximum amount of information when 'letters' are equally common
  - Not surprising that with 2-letter alphabet you can potentially encode 1 bit

- Do real neurons try to maximize their potential information?
  r ≈ firing rate, $0 \le r \le r_{max}$
  To maximize entropy, $P[r] = \text{constant} = \frac{1}{r_{max}}$
  i.e., you should use all firing rates equally

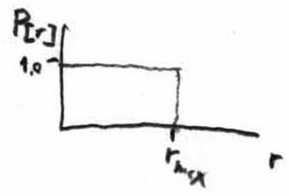

- Imagine the firing rate is a sigmoidal function of stimulus
  - Measure prob. of different stimuli in real world (e.g. contrasts), then find relation: $P[s] = P[r]\frac{df}{ds}$ b/c

$$P[s]\Delta s = P[r] \cdot (f(s+\Delta s) - f(s))$$

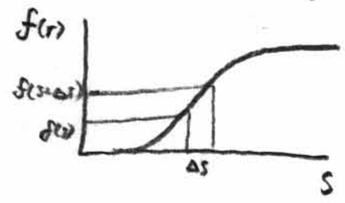

- So $P[s] = \frac{1}{r_{max}}\frac{df}{ds} \rightarrow f(s) = r_{max}\int_{s_{min}}^{s} ds'\,P[s']$
- Neurons will adjust their firing rate curves to keep stimuli from world within their dynamic range

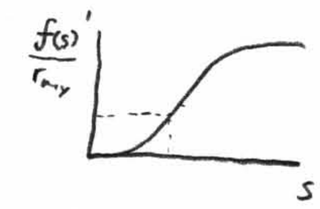

- What maximizes entropy?
  - If constraint is $0 \le r \le r_{max}$, P[r] = constant
  - If constraint is $\langle r \rangle = r_{avg}$, e.g., to constrain energy usage, P[r] = exponential

Information

- Entropy just says how interesting a spike train is
- Mutual information accounts for whether the variability is related to the stimulus
  - It's the entropy, with the variability that's not due to the stimulus subtracted out:

$$I_{mutual} = H - H_{noise}$$

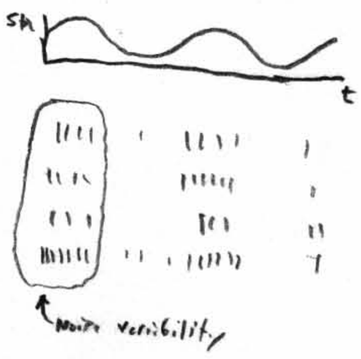

$$H = -\sum_r P[r] \log_2 P[r]$$

$$H_s = -\sum_r P[r|s] \log_2 P[r|s]$$

$$H_{noise} = \sum_s P[s]\,H_s$$

$$I_m = -\sum_r P[r] \log_2 P[r] - \left(-\sum_{s,r} P[s]P[r|s] \log_2 P[r|s]\right)$$

$$P[r] = \sum_s P[s]\,P[r|s]$$

$$I_m = \sum_{s,r} P[s]\,P[r|s] \log_2\left(\frac{P[r|s]}{P[r]}\right)$$

Since $P[r,s] = P[s]\,P[r|s]$
$= P[r]\,P[s|r]$, we can substitute...

- Mutual information = $H_{response} - noise$
  $= H_{stimulus} - noise$ (neuron can only be as interesting as signal it's encoding)

In nervous system, $I_m \approx 1\text{–}3$ bits/spike

wed 27 March

1. If there's a fixed range of firing rates, the max possible entropy comes from distributing P[rates] evenly
- Optimal distribution for a tuning curve: $r = f(s)$ is based on prob. of stim in environment $f(s) = \int^{s} ds'\, P[s]$
  - Deemphasizes parts that are commonly (?) seen in order to distribute firing rate probability

2. If you constrain the avg firing rate, P[r] is a decaying exponential to maximize entropy

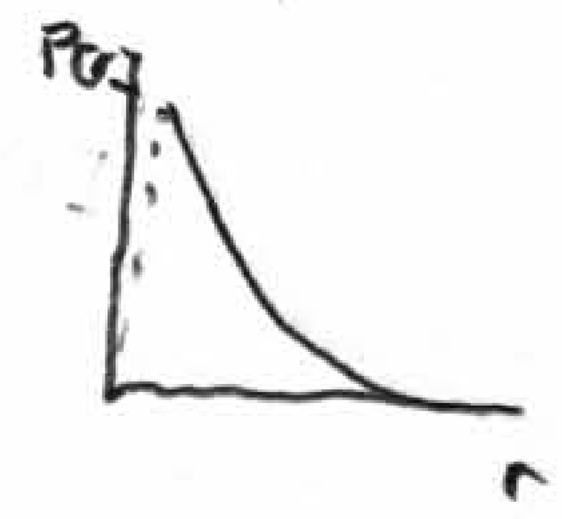

For more complex setups, rates are influenced by multiple cells, so apply the info. theory
- We want to maximize info over the entire population
  1) Maximize individual cells (histogram equalization)
  2) Minimize redundancy (independence) ⇄ A trade-off
  3) Deal with noise (averaging)

- Independence
  - We want for cell 1's firing rate to not tell you anything about cell 2's or cell n's:

$$P[r_1, r_2, \ldots r_n] = P[r_1]\,P[r_2] \ldots P[r_n]$$

  - 'Factorizing' is the way you get zero redundancy. You break up the n-variable function into a product of one-variable functions, and it's the same function

- Remembering simple cells / LGN,

$$r = \int dx\,dy\, D_s(x,y)\, s(x,y)$$
$$r(t) = \int_0^{\infty} d\tau\, D_t(\tau)\, s(t-\tau)$$

- So can we predict what these receptive fields should look like to maximize independence?
- Because of position invariance, we do calculations in Fourier space
- To maximize population, we want the power to be the same across the frequency space (histogram equalization)

whitening filter:
- Looking at stats of neural activity, peak is invariably at 5Hz. To maximize info, though, we want a flat histogram, so we spread our dynamic range over the range of stimuli

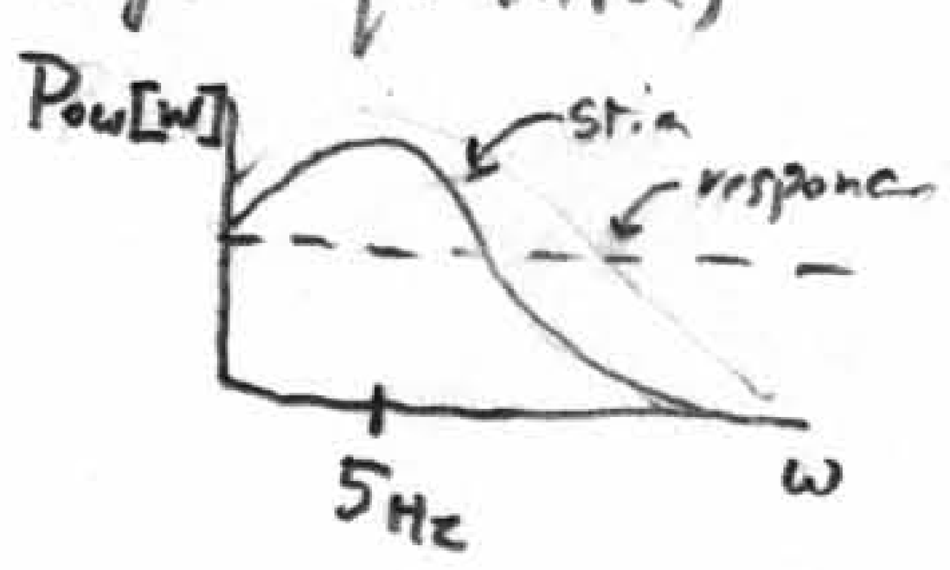

  - Since you're boosting parts of the stimulus that don't have much power, you're suddenly receptive to noise in that boosted region

Matched Filter
- Attempt to get rid of noise which would get spuriously boosted by the whitening filter

  - To be successful, filter should be derived from signal statistics

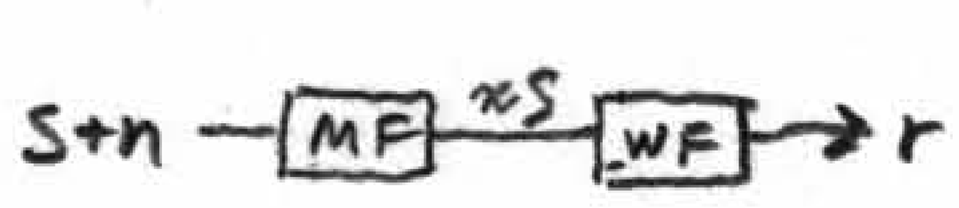

  - Under low stim conditions, dealing with noise is paramount, so the MF dominates, whereas when there's a lot of power in the stimulus, the noise is nominal, so whitening filter dominates

- Difficult to go from set of stimuli to set of filters that will factorize
  - Closest approximation is to use decorrelation:

$$\langle (r_1 - \langle r_1 \rangle)(r_2 - \langle r_2 \rangle) \rangle = 0$$

  - Idea that dreaming checks system, where noise from factorized images generates realistic images lower down

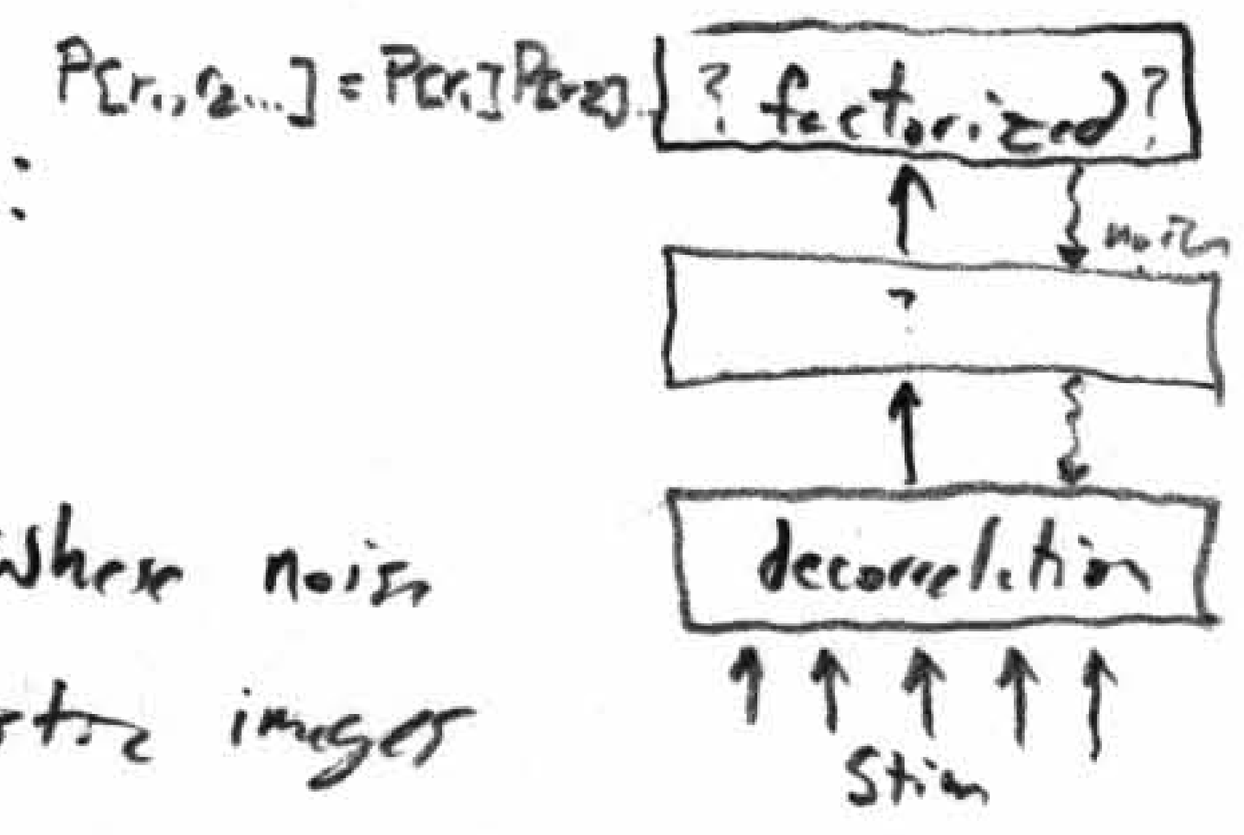

Information From Spike Trains

- $I_{mut} = H[r] - H[r|s]$

$$= -\sum_r P[r] \log_2 P[r] + \sum_{r,s} P[s]\, P[r|s] \log_2 P[r|s]$$

(under $P[r]$: ?; under $P[s]$: known; under $P[r|s]$: ?) — subtract out part that's just due to the cell, not the stimulus

- To figure out P[r], move a sliding bin over spike train
- To get P[r|s], you need to repeat the stimulus and compare the responses to each presentation

- Results of Info. Theory analysis
  1) I_m varies with firing rate (more info with high spike rates)
  2) $I_m/\text{spike} \approx 1\text{-}3$ bits (individual spikes more meaningful at low rates)
     - Bigger bins makes you lose patterns: ||| | is equivalent to ||||
     - So instead, you can break into sub-bins: 11100001 (1ms, 5-10ms)
     - Then rather than looking for P[4 spikes], (coarse resolution) we look for P[11100001] (fine resolution)
       - Ideally, we want fine → 0 (letter specificity), coarse → ∞ (word size)
  3) To get all possible information out of spike train, individual timing (fine) is very important. But the word size is only 1 spike.
     - For 90-95% of time, combinations of spikes are unimportant. Two reasons: 1) Clumps of spikes are correlated anyway, so you don't get much more info. from their relationship 2) Refractory period influences spacing

∴ So it's a simple code with a huge mess of single-character 'words'

tues 9 april

Networks → Plasticity

- Basic network questions are how to treat neurons and how to treat synapses
  - Cells could use H.H. or integrate + fire
  - Synapses can consider conductances, short term dynamics (e.g. depression)
- To stay simple (and fast)
  - Cells are just firing rates (spiking probabilities)
  - Synapses are current injectors
- Stochastic neuron model good in that the bio. data always varies by some amount, so simulating every spike may be overkill for comparison purposes.
- Also, the units in real networks receive noise signals from sources not controlled by experimentor, so there's yet more variability ∴ firing rate probabl, okay

Static Case

- If we know F/I transfer function, all we need is the input current (from synapses) and we can calculate the cell's rate in response
- Rather than treating synaptic inputs as trains of individual spikes, we just take their average over some window (e.g. if input is spiking 1/3 of the time, change it to a constant value with a value of 1/3)

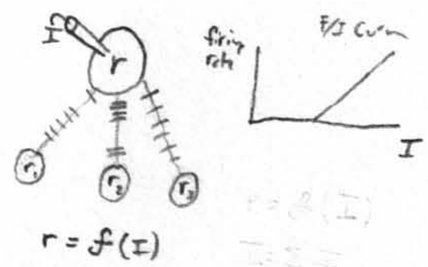

$r = f(I)$

$I = \sum_i I_{max} r_i \tau_s$

↑ Decay rate of this synaptic current

- $I_{max} \cdot \tau_s$ is area under current curve → syn. weight

- Synaptic weight denoted w_{ij} ← Presynaptic, ← Postsynaptic
- Total synaptic input for given unit: $I_i = \sum_{j \neq i} w_{ij} \cdot r_j$

(e.g. $I_1 = w_{12} r_2 + w_{13} r_3 + w_{14} r_4$)

- Then, given total current we use the F/I curve to calculate firing rate:

$$r_i = f(I_i)$$

Dynamic Case

We also need to consider external inputs (e.g. from LGN to our cortical network), so just add another input current which is a function of the stimulus (presuming stim doesn't vary over time):

$$I_i = I_i^{ext}(Stim) + \sum_j w_{ij} r_j$$

- Since model assumes reactions of cells are instantaneous, it's not trivial to make stim. a function of time
- So add a time constant for change in unit's input currents:

current form

$$\tau_s \frac{dI_i}{dt} = -I_i + \left(I_i^{ext} + \sum_j w_{ij} r_j\right) \Rightarrow r_i = f(I_i)$$

- Other idea is that since firing rate is really a function of voltage, not current, we should model:

voltage form

$$\tau_m \frac{dV_i}{dt} = V_{rest} - V_i + R_m \left(I_i^{ext} + \sum_j w_{ij} r_j\right) \Rightarrow r_i = F(V_i)$$

- Difference is whether relevant time constant is synaptic τ (fast, $\approx 2ms$) or membrane τ (much slower, $\approx 20ms$)
  - But neurons can follow inputs much faster than the 20ms τ_m can resolve.
  - Voltage form is useful for determining <u>if</u> you fire (i.e. in subthreshold realm) while current form useful for determining <u>how much</u> you fire (at-threshold realm)

fri 12 april

Firing Rate Models (current or voltage?)

- Steady State

$$r_i = f(I_i)$$

$$I_i = I_i^{ext} + \sum_j w_{ij} r_j$$

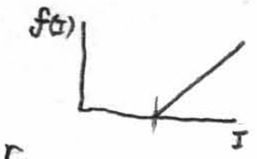

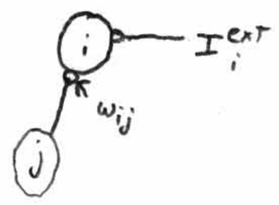

- Dynamic
  - Neuron doesn't react instantaneously, so r_i equation now invalid

$$\tau_r \frac{dr_i}{dt} = -r_i + f(I_i)$$

change from: 

to: 

  - Question becomes whether τ_r is based on slow, membrane time-constant or fast, synaptic time constant

In Summary:

$$\boxed{\tau_r \frac{dr_i}{dt} = -r_i + f\left(I_i^{ext} + \sum_j w_{ij} r_j\right)}$$

- Considering the f function
  - f will usually be a rectified linear function, but if we temporarily do away with the rectification, we can analytically examine the resulting linear model
  - So, assuming a linear f with a threshold at zero, the equation becomes

look at single-unit case first:

$$\tau_r \frac{dr_i}{dt} = -r_i + I_i^{ext} + \sum_j w_{ij} r_j$$

$$\tau_r \frac{dr}{dt} = -r + I^{ext} + \omega r \quad (\omega \text{ = autapse})$$

$$= -(1-\omega) r + I^{ext}$$

$$\frac{\tau_r}{1-\omega} \cdot \frac{dr}{dt} = -r + \frac{I^{ext}}{1-\omega} \qquad \left(\text{So } r \to \frac{I^{ext}}{1-\omega} \text{ with } \tau_{eff} = \frac{\tau_r}{1-\omega}\right)$$

Two main effects seen in equation

1) Speed: In inhibitory case ($\omega<0$), response is lower, but τ_{eff} shrinks and response is faster.

2) Amplification: if $(0<\omega<1)$, response increases, amplifying I^{ext}, but speed also decreases

i.e. Excitation reinforces itself, so it is slow to come down again

(Idea is that cortical cells operate in regime 2: slow but powerful)

3) if $(\omega>1)$ model becomes unstable, equation no longer decays to $\frac{I^{ext}}{1-\omega}$ but explodes instead. (epilepsy?)

4) if $(\omega=1)$, equation becomes $\tau \frac{dr}{dt} = I^{ext}$. So r just becomes the integral of I^{ext}.

- So instead of I affecting the value of r, it affects its slope
- Is also a memory device, since it tracks, e.g. the number of input pulses we've seen

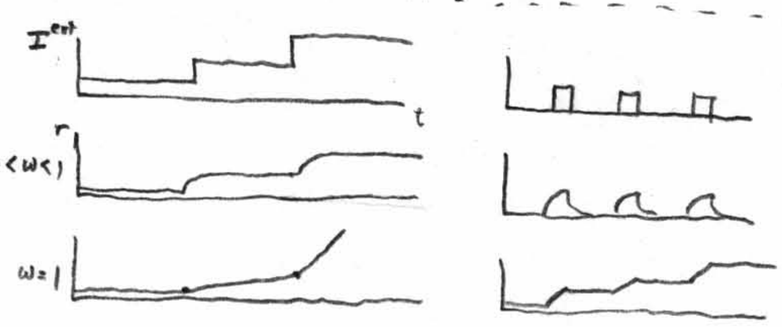

Integrator Neurons

- Commonly seen in oculomotor pathway
  - Saccade commands are sent in brief pulses
  - Integrator circuit 'counts' motor signals and 'remembers' where the eye is currently pointing.
- More general memory systems too
  - Cell accepts initial input, then is cut-off from it and maintains itself in a loop.

Returning to the Real Model

- Trick is to turn multi-unit model into something that looks like single-unit answer is to 'diagonalize w':

$$r_i = \sum_{\mu=1}^{N} C_\mu^{(t)} V_i^\mu$$

($C_\mu^{(t)}$: coordinates; V_i^μ: axes)

$$\begin{matrix} & j \to \\ i \downarrow & \end{matrix} \begin{pmatrix} & \\ & w \\ & \end{pmatrix} \begin{pmatrix} \\ r \\ \\ \end{pmatrix}$$

- Condition on V is that:

$$\sum_{j=1} w_{ij} V_j^\mu = \lambda_\mu V_i^\mu$$

(V_i^μ: eigenvectors(?); λ_μ: eigenvalue)

- Plugging in:

$$\tau_r \sum_\mu V_i^\mu \frac{dC_\mu}{dt} = -\sum_\mu C_\mu V_i^\mu + \sum_\mu \lambda_\mu C^\mu V_i^\mu + I_i^{ext}$$

$$\tau_r \frac{dC_\mu}{dt} = -C_\mu + \lambda_\mu C_\mu + \sum_i V_i^\mu I_i$$

- Just like original equation, but matrix now just a number... uh...
- Similar to 1-unit ω λ in ω's place

Looking at some cases:

2) amplification is seen on particular connections, so your amplification can be more selective to some particular aspect of the input pattern

- Removing our simplifications

1) ω is symmetric
   - Doesn't make any sense (e.g. what if an E is connected to an I?)
   - If we ditch the symmetry requirement, we get a new behavior type: oscillations

2) f is linear
   - A non-linear model allows you to control the $\omega>1$ case where before you would get exponential explosion.

Wilson-Cowan Model

- Two pools of neurons: Excitatory and Inhibitory. All units in pool share r

$$\tau_E \frac{dr_E}{dt} = -r_E + \left[w_{EE} r_E + w_{EI} r_I + I^{ext}\right]_+$$

$$\tau_I \frac{dr_I}{dt} = -r_I + \left[w_{II} r_I + w_{IE} r_E + I^{ext}\right]_+$$

- If $(\lambda_\mu - 1) > 0$, linear model would blow up, but the rectification at zero keeps it in range

if $\lambda \in \mathbb{R}$, exponential behavior

$\lambda \notin \mathbb{R}$, oscillatory behavior

tues 16 april

Networks

$$\tau_r \frac{dr_i}{dt} = -r_i + f\left(I_i^{ext} + \sum_{j=1}^{N} w_{ij} r_j\right)$$

- Two complications
  1) Huge weight matrix
  3) Strange f nonlinearity
- Weight matrix (self-coupling in N=1 case) effects:
  - If excitatory - amplifies and slows time constant (b/c it's not responding just to the input signal)
  - When w is a matrix, there is selective amplification of certain inputs

Simplify again to single-unit case

$$\tau_r \frac{dr}{dt} = -r + f\left(I^{total}(r)\right)$$

- To find a stable state, set derivative to 0, i.e. at r in:

$$r = f\left(I^{Total}(r)\right)$$

- Assume f is a typical sigmoid
- When $f() > r$, firing rate will flow upwards, if $f() < r$ will flow downwards
- With low I^{ext}, only stable fixed point is at zero, but if current can push rate past threshold[2], cell will reach the upper fixed point (which is now stable)
  - We now have a flip flop memory where the neuron will switch between bistable states based on firing rate vs. some threshold

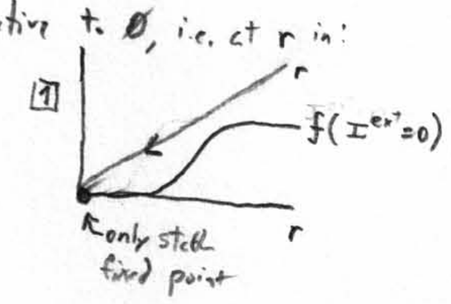

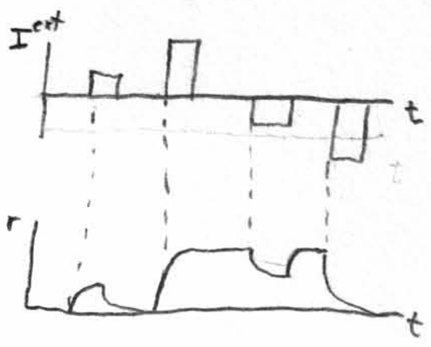

We can use selective amplification to model orientation selectivity in SCs

- H+W idea was that orientation selectivity comes from the input patterns from the LGN
- Problem w/ idea is the anatomy: Vast majority of input cells come not from LGN, but elsewhere in cortex.

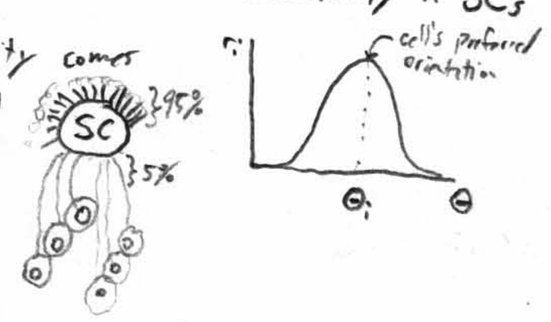

- Rather than labeling cells simply by index, use preferred theta: $r(\theta_i)$

likewise weights: $w(\theta_i, \theta_j)$

likewise current: $I^{ext}(\theta_i, \Theta)$ ← cell, ← stimulus

Underlying assumption is that all cells in same orientation selectivity are the same

- Assumption for weights is that values are just a function of the diff. in orientation between cells: $w(\theta_i - \theta_j)$

$I^{ext}(\theta_i - \Theta)$

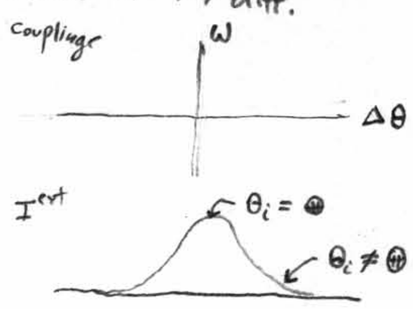

- One problem w/ H+W model is 'iceberg effect' as contrast ↓, orientation selectivity narrows.

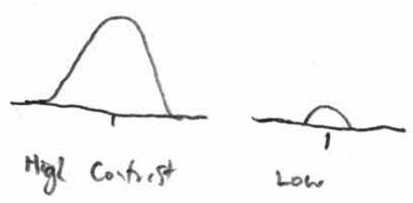

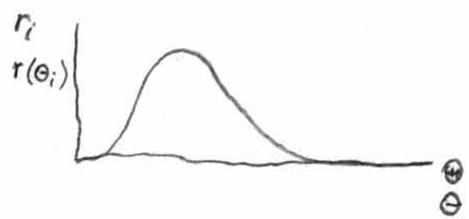

Alternative Model

- Set up weights to tightly connect units w/ similar preferred θ, w/ inhibition from those far from its preference
- Advantage is that selectivity width is constant since its controlled by the zero-point in the weight distribution
- Model's assumptions are that:
  1) There is amplification in cortex
  2) The target of amplification is orientation tuning
- To test model, knock out the amplifying connections
  - Result is lower activity, but orientation selectivity is still there
  - So amplification is occurring, but it's not responsible for selectivity

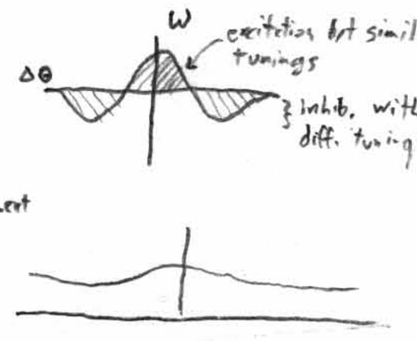

Complex Cells

- Orientation selective, but spatial phase invariant w drifting gratings
- Frequency doubled response to counterphase gratings ⇒
- Argument for model is that amplification is acting on the spatially invariant part
  - Set w_{ij} to be independent of spatial phase (i.e., rather than amplifying inputs w/ similar θ, we amplify signals with different phase - selectivity) ↑?
- If cells are tightly coupled, cells behave as complex cells

  If cells are weakly coupled, you get simple cells
  - So presumably there is a spectrum of cells in cortex, and data show that there is a bimodal histogram w/ seperate simple + complex populations (but spectrum idea would expect a single peak...)
  - Argument that two-populations may just be an artifact: If the input → firing rate conversion is sharp, its possible for cells w/ moderate couplings to be pushed toward looking like their couplings are more extreme than they are.

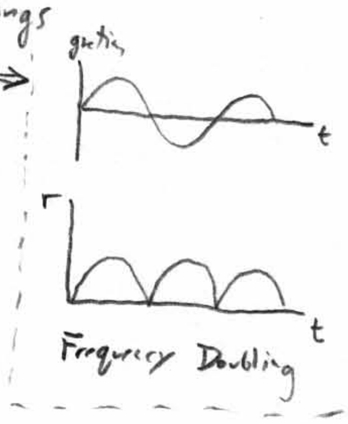

fri 19 april

Network Dynamics

1) Amplification – increase selectivity e.g., Simple Cells ← complex recog.
   decrease selectivity e.g., Complex Cells ← invariance

2) Slowing Down – integration (i.e., memory). Since neurons are driving themselves they don't notice instantly that input changed
   - Short Term Memory as the maintenance of a pattern of activity within a network
   - Long Term Memory stored in synaptic weights

Autoassociation Memory

- The beginning of a bit of info causes the full memory to bootstrap itself into activity. (Seen in, e.g., object recognition)

$$\tau_r \frac{dr}{dt} = -r_i + f\left(\underbrace{I^{ext}}_{\uparrow} + \sum w_{ij} r_j\right)$$

The initial input that gives rise to the full memory, it sets the partial-memory activation pattern
If $r_i(t=0) \approx M_i$, $r_i(t=later) = M_i$

- 'Domains of attraction', when you set the initial condition, the system will drift to the nearest metal

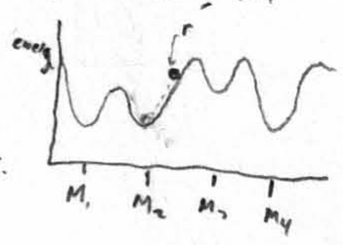

- To focus/broaden certain memories, we shape the surface to make wells steeper, wider, etc.
- Agrees w/ tip-of-the-tongue phenomenon; if you are thinking of many things, you're adding noise to the system. So even though you know you're in the right neighborhood, you can't find an initial point that will get to the attractor.

Modeling Memory

- We set initial state $r_i(0) \approx M_i^1$, where input is: $\sum^n w_{ij} r_j$
- So our update rule is:

$$r_i(t+\Delta t) = \text{sign}\left[\sum_{j=1}^{N} w_{ij} r_j(t)\right]$$

- So how do you design ω to make the memory state move to proper fixed point given initial condition and then to stay there (i.e., make sure the wells in ω are in proper places)?

neurons $i = 1, 2, 3, \ldots n$
memories $a = 1, 2, 3, \ldots p$
$r_i = M_i^a$ ← which memory
↖ which neuron
$r_i(0) \approx M_i^1 : r_1 \to M_1$
$r_i = \begin{cases} 1 & \text{active} \\ -1 & \text{inactive} \end{cases}$ $m_i^a = \begin{cases} 1 & \\ -1 & \end{cases}$

- Simplest answer is the Hebb Rule:
  Neurons that are active together get excitatory connections; others inhibitory.

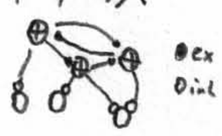

- Mathematically, construct ω from outer product:

$$w_{ij} = M_i^1 M_j^1$$

$$\text{sgn}\left[\sum^N M_i^1 M_j^1 M_j^1\right] = \text{sign}\left[M_i^1 \sum^N (M_j^1)^2\right] = \underbrace{M_i^1 M_j^1}$$

Since we're just looking at the sign, magnitude doesn't matter

$$w_{ij} = M_i^1 M_j^1 + M_i^2 M_j^2 + M_i^3 M_j^3 + \ldots + M_i^N M_j^N$$ } i.e. strong connections for active, weak for inactive

- To probe memory trace for M^1

$$\text{sign}\left[\sum w_{ij} M_j^1\right] = \text{sign}\left[\sum \left(M_i^1 M_j^1 + \sum_{a=2}^{p} M_i^a M_j^a\right) M_j^1\right] \quad \pm\sqrt{N} \text{ (b/c 1 or -1)}$$

$$= \text{sign}\left[\underbrace{M_i^1 \sum (M_j^1)^2}_{\text{The memory we're correcting}} + \underbrace{\sum_{2}^{p} M_i^a \overbrace{\sum^{N} (M_j^a M_j^1)}}_{\text{other memories interfering with memory 1}}\right]$$

$$= \text{sign}\left[N M_i^1 \pm \sqrt{N} \sum_{a=2}^{N} M_i^a\right] \leftarrow \text{good memory } N \text{ times, bad memory only } \sqrt{N} \text{ times}$$

$$= \text{sign}\left[N M_i^1 \pm \sqrt{n}\sqrt{p-1}\right] \leftarrow$$

- So $\sqrt{n}\sqrt{p-1} < N$, or else we've exceeded capacity

- So we need $\sqrt{n}\sqrt{p-1} < N$ or else we've exceeded the network's capacity obscuring the stored memory:

$$\sqrt{N}\sqrt{p-1} < N$$
$$p - 1 < N$$
$$p = 0.14N$$

} Hopfield Networks

- So if $N = 10^6$, you can store 140,000 bits in the memory

tues 30 april

Function Approximation

- Simpl feedforward network where

$$r = \sum_{i=1}^{N} w_i f_i(s)$$

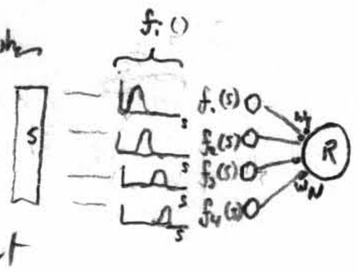

- Delta rule provides weight change based on approximating desired output
- Each input's tuning curve is displaced relative to the others. So "learning" is just strengthening the bumps that line up w/ the input/target relationship
- Delta rule: (supervised learning) every time you see a stimulus, change w by a rule

$$\Delta w = \sum_{i=1}^{N} (h(s) - r) f_i(s)$$

($h(s) - r$: error; $f_i(s)$: activity of given input)

- As error gets smaller, change magnitude also gets smaller
- Neurons w/ high firing rates have their weights boosted more
- Bacterial Chemotaxis strategy (Reward Learning)
  - Delta rule is given the error, and whether you improved in last trial
  - Do random walk w/ conditional turns
    - If target signal just increased, keep going forward
    - else, pick new direction randomly in param. space

- Unsupervised Learning
  - Hebb rule & Principl Component Analysis
    - Since you're condensing signals from hundreds of inputs into a singl number – the output – it makes sense to align value with axis of greatest variability.
      - e.g. you maximize information by reporting based on the Projection of the points onto the P.C. axis

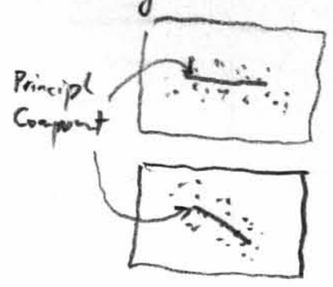

  - Hebb rule causes weights to drift toward reporting the PC axis
  - Modified hebb rules can also find second- third- & n-most interesting axes
  - In a nutshell:

$$\tau_w \frac{dw_i}{dt} = u_i v \quad \Rightarrow \quad \frac{1}{2}\tau_w \frac{d|\vec{w}|^2}{dt} = v^2$$

(v: output signal; u_i: input signal)

- Problem w/ Hebb rule is that over time $|\vec{w}| \to \infty$
- Also, there needs to be some communication between neurons w/ similar input populations. Otherwise all cells will gravitate toward some representation (the most interesting one)
  - So there must be some sort of competition to force neurons to show some diversity
- One quick fix is to hold $|\vec{w}|$ constant. This prevents both the runaway strengthening, and forces competition since increases in one weight must be paired w/ reductions in others

$$\tau_w \frac{dw_i}{dt} = u_i v - \alpha v^2 w_i \quad \Rightarrow \quad \frac{1}{2}\tau_w \frac{d|\vec{w}|^2}{dt} = v^2 - \alpha v^2 |\vec{w}|^2$$

$$= v^2(1 - \alpha|\vec{w}|^2)$$

- Other plasticity question is how changes in one synapse affect other synapses onto same cell
  - Tetanus LTP induction in hippocampal CA1 boosts strength of synapses where the input was coactive w/ the postsynaptic cell. Inactive synapses remain unchanged

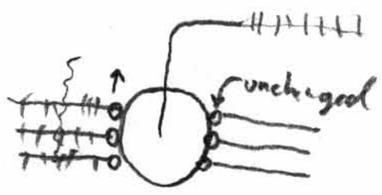

  - In LTD case, lack of postsynaptic activity paired w/ presynaptic spikes leads to drop in strength

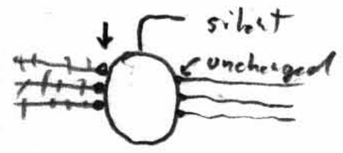

fri 3 may

LTP & LTD

1) Hebb - If pre 'causes' post fire → potentiate
2) Neurons that fire together wire together
   → Modified to just be pre/post correlation

Initial definition meant 1+2 firing together

- Lost in typical LTP/LTD protocol is fact that tetanus/LFS never happen in healthy CNS, so the learning + memory angle fades

Spike-Timing Dependent Plasticity

- Pre- before post- firing strengthens synapse. Reverse pairing occurs
- Cool because it only requires 50 pairings or so — reasonable behavior naturally — and returns to idea of causality
- The P_{ca} curve actually varies from system to system.

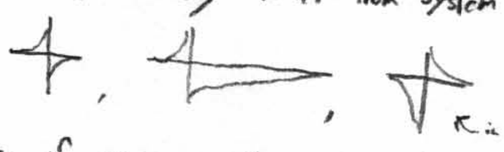

- In case of presyn spikes that are randomly spaced (uncorrelated) the synapse will generally depress (b/c LTD curve has larger area)
- Experimentally, we know about STDP w/ single synapses, but for afferent bundles, we have only the tetanus-type LTP example
  → Cliff's model takes an IaF neuron w/ 1000+ synapses and applies STDP rules
- General fear in Hebbian-type modification is runaway strengthening of weights (thus the need for global constraints)
  - Here synapses are just required to a) be >0, and b) be $<MAX$
  - When all inputs are poisson generators there is a segregation bt strong + weak synapses w/ few in the middle
    - Happens bc there is competition where strong synapses fight to control the postsynaptic spikes
    - Initially strong synapses influence spiking which makes it more likely for future pre-post times to be well correlated - leading to continued strength
  - Partitioning rule is also homeostatic: as input rates increase, the proportion of strong synapses drops, keeping postsyn rates at in reasonable range.

tues 7 may

STDP

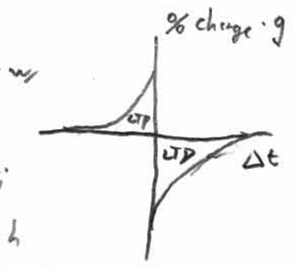

- Extrapolating from single synapse properties to cells w/ thousands of inputs
- One purpose is homeostasis: won't to stay w/i dynamic range rather than saturating at high input
- Also nice b/c just boosting an input's rate doesn't strengthen the synapse: instead it's correlation that decides
- Competitive pressure: only the best correlated inputs will strengthen: if a new set of inputs w/ higher correlation than current 'top' group, it will bump the old off the top.

Less-Toy Example

- There's a ≈ 10ms lag period b/t inputs from low latency visual afferents' response to stimulus arriving and the slower afferents' signal
  - So even though we receive input at time t, it's not until t+10 that we actually get enough input to spike
- Applying STDP to a range of variably latent inputs, strengthens the inputs w/ the small lag and lets go of the more delayed ones.

Whisker System of Rat

- Primary sensory cortex ('barrel cortex') has whisker topic map
- Each cell has primary whisker it responds to, and has much less response to neighbors.
- If you clip the primary whisker, you knock out its main input (then can test all response by wiggling stump)
  - The decreased input weakens the involved afferents
  - Next, intracortical inputs from neighboring barrel strengthen, Meaning our cell is now firing in a manner correlated to this new whisker
    - Now conditions are set for the direct thalamic inputs for neighboring whisker to strengthen
- Since the thalamic pathway is lower latency, the STDP competition sets in against intracortical Pathway → low latency wins
  - Thus intracortical inputs are acting as a 'teacher' adapting the cell to a new input

Place Cells

- If the rat walks from one place field to the next, STDP predicts that the synapses from the first to the second will be strengthened (pre before post) and the other directional synapse weakened
  - After the connection is strengthened enough, the place-field of the second cell should extend to encompass the first cell's field, so the center of mass drifts forward toward the earlier place field.
- As a result, place fields along (and in the neighborhood of) the repeated path will be elongated such that the bump of population activity Centered at the current location will spill over in the direction of the repeated path.
  - So the hippocampus isn't so much signaling where the rat <u>is</u>, but instead where it <u>will be</u> if it heads back toward the path.