Lecture 4: Temporal representation in spiking neurons (cont.)

4.1 Using and analyzing temporal representation

4.1.1 Decoding signals in LIF neurons

Following the procedure outlined above for a typical pair of LIF neurons results in a decoder as shown (see slides). Note:

- The filter was found for bandlimited (0-30Hz) Gaussian white noise on 4 s of data.
- The filter is not symmetrical in time but skewed in the positive direction
- The filter has a negative time component (non-causal).
- The envelope of the decoder extends just over 5 ms or so.

Using this decoder as described above results in successful representation of signals in the ensemble (see slides).

Looking at the power spectrum of the original signal compared to the LIF neuron signal gives some hints about what is happening in the encoding/decoding process:

- The high frequency high-power components of the spiking neuron signals are removed by the filter (low pass filter)
- Conversely, high frequency is introduced by the encoding process.
- Overall, the original signal bandwidth and power is well-preserved.

Two outstanding issues:

1. Correlation times (rate and timing codes)
    (a) Correlation time is the maximum time window for which the autocorrelation function is above half its maximum value. This means that the future values of the signal are (somewhat) predictable out to that window (i.e. the signal past that window is statistically independent of the current value of the signal). So generally, a wider bandwidth means a shorter correlation time.

    $$Autocorrelation(\tau) = \int_{-\infty}^{\infty} f(t)f(t+\tau)\,dt$$

    Note: Fourier transform of the autocorrelation function is the power spectrum. (Matlab: x=fft(s); Power=fftshift(x.*conj(x)); or x=fft(xcorr(s)); Power=fftshift(sqrt(real(x).^2+imag(x

   (b) The same decoder can be used for high and low correlation times. In fact, these make the signal look more like a rate and timing code respectively. So, depending on the characteristics of the signal, the code will seem to change (i.e. signals with high freq will demand a timing code).
   (c) However, if the entire ensemble (not just one signal) tends to have a specific kind of correlation time, then (not surprisingly) an optimal decoder for that ensemble will do a better job than the optimal decoder for another ensemble that this ensemble is a subset of (see slides).
   (d) In any case, this result means that these methods are effective over regimes typically called ‘rate’ and/or ‘timing’ codes. This is very useful, and shows that the distinction can be ignored (more or less).

2. Neural plausibility: Post-synaptic currents (PSCs)
   (a) PSCs are ellicited upon the reception of neurotransmitters by dendrites. They can be mathematically described quite compactly (see slide).
   (b) PSCs look surprisingly like optimal filters (although they are strictly causal).
   (c) Decoding with PSCs instead of optimal filters results in only a small reduction in information transmission, but a huge gain in neural plausibility (see slide).
   (d) Reductions in coding accuracy using PSCs can be made up for by including more neurons, as we learned from the section on population coding.
   (e) Nevertheless, examing optimal filtering is useful because it provides:
      i. a means of comparing information processing characteristics of models to real neurons
      ii. bounds for non-optimal decoding (i.e. PSC decoding)
      iii. similarity between the optimal linear decoder and the PSC decoder suggests linear decoding is a good characterization of neural representation
      iv. optimal filter is more amenable to certain analyses which may (or may not) inform filtering in the non-optimal case.
   (f) Given these considerations, all future models use PSC filtering instead of optimal filtering.

4.1.2 Information transmission in real neurons

A brief summary of the information characteristics observed for real neurons:

- measures are surprisingly consistent across many different neural systems
- cricket cercal system (measures wind velocity), between about 150 and 300 b/s.[1] (between 1.1 and 3 bits per spike)

---

[1] Miller et al. calculate the rate to be about 40 bits per second, but they used a 100 ms binned rate code to calculate information.

- bullfrog sacculus (senses ground-borne vibrations) rates of about 3 bits per spike
- motion-selective H1 neurons in the blowfly; about 3 bits per spike
- salamander retina, rate of about 3.4 bits/spike
- primate visual area V5 (aka MT, for sensing motion), information transmission rates of 1–2 bits per spike
- highest transmission rates: bullfrog auditory neurons, rates as high as 7.8 bits per spike (for natural stimuli, broadband stimuli had transmission rates of about 1.4 bits/spike)
- Sum: natural sensory systems are generally able to encode stimuli with between about 1 and 3 bits of information per spike (Rieke, et al, 'Spikes' has a good review)

Notice:

- these are impressively high transmission rates that approach the optimal possible coding efficiencies
- in frog sacculus, cricket cercal system, bullfrog auditory system, and electric fish electrosensory system, the codes are between 20 and 60% efficient
- efficiency significantly increases when stimuli are restricted to be more like naturally occurring stimuli of interest to the organism
- the estimation of information transmission rates using this method places a *lower* bound on the amount of information transmitted by the code, efficiencies are likely even higher.

4.1.3 Information transmission in model neurons

When looking at the LIF model (or any other deterministic model), we should realize that technically, there is no noise in the system, so we can't compute a signal to noise ratio, and thus we can't determine information transmission in these neurons. However, there is a source of error that acts analogously to noise: our linear estimate. The linear estimate is imperfect, so there will always be some deviation from the value being represented. We can use this to determine an information transfer rate in these neurons. Specifically, we know that information in a channel is (Hartley-Shannon law):

$$I = \frac{1}{2}\log_2(1 + SNR)$$

We can find the SNR by comparing the variance explained by our decoding, $\left\langle \hat{x}^2 \right\rangle_x$ (i.e., assuming zero mean of the estimate), to the error that remains unexplained by this decoding. Recall that our decoder is

$$h = \frac{\langle xR \rangle_x}{\langle R^2 \rangle_x}$$

where $R$is the encoded signal. Recall that:

$$\begin{aligned} hR &= \hat{x} \\ h^2R^2 &= \hat{x}^2 \\ \left\langle \hat{x}^2 \right\rangle &= \left\langle h^2R^2 \right\rangle_x \end{aligned}$$

so we have our explained variance:

$$\begin{aligned} \left\langle \hat{x}^2 \right\rangle_x &= \left\langle h^2R^2 \right\rangle_x \\ &= \left[\frac{\langle xR \rangle_x}{\langle R^2 \rangle_x} \right]^2 \left\langle R^2 \right\rangle_x \\ &= \frac{\langle xR \rangle_x^2}{\langle R^2 \rangle_x} \end{aligned}$$

and our unexplained variance:[2]

$$\begin{aligned} MSE &= \left\langle x^2 \right\rangle_x - \left\langle \hat{x}^2 \right\rangle_x \\ &= \left\langle x^2 \right\rangle_x - \frac{\langle xR \rangle_x^2}{\langle R^2 \rangle_x} \end{aligned}$$

Which gives a signal-to-noise of:

$$\begin{aligned} SNR &= \frac{\left\langle \hat{x}^2 \right\rangle_x}{MSE} \\ &= \frac{\langle xR \rangle_x^2}{\langle x^2 \rangle_x \langle R^2 \rangle_x - \langle xR \rangle_x^2} \end{aligned}$$

Plugging this back into the information expression earlier gives two equivalent expressions, one with respect to the input

$$I = \frac{1}{2} \log_2 \left[\frac{\left\langle x^2 \right\rangle_x}{\left\langle x^2 \right\rangle_x - \frac{\langle x \cdot R(x) \rangle_x^2}{\langle R^2(x) \rangle_x}} \right]$$

and the other with respect to the response

$$I = \frac{1}{2} \log_2 \left[\frac{\left\langle R^2(x) \right\rangle_x}{\langle R^2(x) \rangle_x - \frac{\langle x \cdot R(x) \rangle_x^2}{\langle x^2 \rangle_x}} \right].$$

This expression has been derived assuming that we are working in the frequency domain (hence $\hat{x} = hR$), so to get an information rate (i.e., bits/s) we need to sum the information per frequency channel and convert it to time

$$InfoRate = \frac{1}{2} \frac{\Delta\omega}{2\pi} \sum_n I_n.$$

---

[2] We know this first expression is true because $\left\langle [x - \hat{x}]^2 \right\rangle_{x,\eta} = \left\langle x^2 \right\rangle - 2h \langle xR \rangle_x + \left\langle h^2R^2 \right\rangle_x$, which, subbing in for $h = \frac{\langle xR \rangle_x^2}{\langle R^2 \rangle_x}$ gives this result.

(Recall that the sampling time step $\Delta t = \frac{2\pi}{\Delta\omega}$ and the Nyquist limit states that the max frequency is twice that step).

Note:

- Now we have a means of calculating information transmission rates (and bits per spike) to allow us to compare model neurons to real neurons.

4.2 More realistic neural models

4.2.1 Adapting LIF neuron

Let's begin to look at trying to model other important properties of neural action potential generation. Simplest is to add adaptation to the LIF neuron.

- The most common cortical neurons are called 'regular spiking'
- Regular spiking cells have adaptation, that is, their spike rate slows down as they continue to spike (see slide)
- This happens because of a slow hyperpolarizing K current in these cells
- The effects of this current are captured by adding a variable resistor, R_{adapt}, to the model (see slide)
- Equations:

$$\begin{aligned} \frac{dV}{dt} &= -\frac{1}{\tau^{RC}}\left(V\left(1+\frac{R}{R_{adapt}}\right)-J_M R\right) \qquad (4.1)\\ \frac{dR_{adapt}}{dt} &= \frac{R_{adapt}}{\tau_{adapt}} \end{aligned}$$

- The channels that account for this conductance stay closed until an action potential is fired, and then they slowly open (with time constant τ_{adapt}), essentially lowering the reset voltage of the cell, making it more difficult for the next spike to be fired.
- When firing stops, the channels start to close, again raising R_{adapt}.
- This simple addition is very good at approximating full conductance models (see slide)
- Also, it significantly improves information transmission efficiency with out loss of accuracy (see slide)

4.2.2 Brief overview of neural firing

Recall that LIF doesn't model the action potential at all, just pre-threshold voltage changes due to input. It 'sticks on' a spike when the threshold is crossed. To understand the action potential dynamics, we have to look at the biophysical processes.

Resting state (see slide):

- Resting potential is about -70mV
- This is the state at which all of the ions are in equilibrium
- Forces moving ions: electrostatic gradient, concentration gradient, sodium-potassium pump (K in, Na out)
- Na+ and Cl- are concentrated outside the cell, K+ and Anions (A-) are more concentrated inside
- The membrane is permeable to everything but A- (though more to K+ than Na+)
- The Na–K pump ensures that the cell doesn't deplete K and fill up with Na (if this happened, the resting potential would shift).
- If the potential is moved away from the resting state, these passive ion movements cause the potential to move back to the resting state. This is where the leak potential in the LIF model comes from.

Action potential generation:

- The potential is driven away from the resting state by input from dendrites (depolarizing or hyperpolarizing)
- when the neuron is depolarized to it's threshold, voltage-gated channels for Na open. These allow Na to rush in, depolarizing the cell.
- shortly thereafter, voltage-gated K channels open (and the Na channels slowly start to close) allowing K to rush out and rapidly repolarize the cell
- there are other channels in cells that open on different time scales, allowing for slight changes in this typical firing of an action potential (e.g. hyperpolarizing adapting currents, see below).
- The first quantitative model of this was the Hodgkin-Huxley model (giant squid axon). This is a 'class II' neuron (Hopf bifurcation, jump in firing rates).
- (See slides) HH model is a 4D nonlinear system. The dynamics of the parameters governing activation (m, n) and inactivation (h) have eqns of the same form, but are all different functions of membrane potential, V. These functions are empirically determined (i.e. fit to data from squid axon; see slide). The dynamics are that these variables approach some asymptotic value (e.g., $N(V)$) with some time constant (e.g., $\tau_n(V)$).

4.2.3 θ-neuron

- Notice that during neuronal firing, there is a membrane potential that depends on the dynamics of ion channels that open and close depending on the membrane potential. That is, there is a dynamic system of coupled parameters. For this reason, dynamic systems theory is often used to understand neuronal dynamics.
- 'class I' neurons have a zero minimum spiking frequency and increase monotonically (like the LIF) - these are the most common in cortex.
- To get this extra behavior, we need to add another fast (though slower than Na) K current (the 'A-current') (see slide).
- The circuit now includes the Na current, the K current, the A-current, and the leak current but has no delta function generator. This is because the dynamics of Na and K take care of the spike generation.
- The A-current has a near-rest threshold (for both activiation and inactivation), which means that weak, sustained input slowly inactivates the channels until a spike is fired, which rapidly activates the channels, bringing the cell back to rest, where slow inactivation begins again (hence very slow firing rates are possible).
- In DST, the addition of this current changes the bifurcation from a Hopf (in Hodgkin-Huxley/class II neurons) to a saddle-node bifurcation (in class I neurons).
- The θ-neuron models this kind of bifurcation (see slide)
  - In (a) and (c) the upper diagram shows the location of critical points on the invariant circle, the middle graph shows the behavior of θ, and the lower graph is the trace of $(1 - \cos\theta)$ showing the spikes. (a) Excitable regime with $\beta + \sigma = -0.3$. The stable state is the node on the right. The single spike is evoked by a pulse stimulus (marked by the triangle) that forces θ past the saddle fixed point on the right. (b) Meeting of the saddle and node points with $\beta + \sigma = 0$. The trajectory has an infinite period. (c) Oscillatory regime where the stable state is now a limit cycle with $\beta + \sigma = 0.3$. Periodic behavior of the phase variable and spikes in $(1 - \cos\theta)$ are present. (d) Phase evolution and its analog to membrane voltage states. Note that the spike occupies a small region near π. A strong enough stimulus will push θ past the threshold and into the excited region. Here the regenerative dynamics that summarize active conductances carry the phase through the spike.
- This model is canonical (i.e. it represents the class of models with a similar bifurcation).
- It also has good/similar decoding compared to LIF neurons. This suggests LIF neurons are reasonably good appriximations to these dynamics (without actually having the dynamics explicitly).
- LIF neurons run 100 times faster.

4.2.4 Wilson neuron

- This is the most realistic of the neurons we have seen so far, even though it is a ‘reduced model’; i.e., there are direct mappings between neurophysiology and model parameters.

- it’s a class I, adapting neuron with spike dynamics; captures spike height changes, spike shape, and after-hyperpolarizations (overshoot of the resting value after a spike).

- To understand the reduction, start with a Hodgkin-Huxley (HH) neuron model. This is the first (and most famous) quantitative neural model, and is based on the giant squid axon (it’s class II (hopf bifurcation)).

- Rinzel simplifications:

  1. Na activation is very fast ($M(V)$ is really small), so we can eliminate the dynamics and allow $m = M(V)$ (see slide)
  2. Na inactivation is about equal and opposite to K activation, so we can let $h = 1 - n$ (see slide), again removing a differential equation (we combine h and n into one variable, R)

- This is a 2D model of a class II neuron. We need the A-current

- Rose and Hindmarsh simplification: A-current inclusion can be approximated by making the dynamics for R a quadratic (see slide). So we have a 2D model of a class I neuron.

- We can include adaptation by putting in another variable with slow dynamics that acts like R_{adapt} in the adapting LIF (see slide).

- The result is the Wilson model with adaptation. This compares favourable over a wide range of stimulation to real cortical data (see slide).

- Again, decoding this gives good info transfer (2 bits/spike).

- BUT: Takes 600 times longer to run than the LIF

4.2.5 Summary

- We have progressed from simple, ‘phenomenological’ models through to more complete models that include adaptation, spike dynamics, and ion channel dynamics. This gives a good overview of available models in computational neuroscience.

- The main class we haven’t discussed are compartmental models. These are much the same, but model spatial distribution of the neurons (we have considered ‘point’ neurons only), by using Rall’s cable equations (see here: http://diwww.epfl.ch/~gerstner/SPNM/node17.htm

- As for information transmission, all of the models have info rates between 1-3 bits/spike, comparable to real neurons (see slide).
- Adaptation seems to help improve efficiency without adversely affecting accuracy.
- Keep in mind we are always using Gaussian white noise here. Real neurons often seemed tuned to natural signals (i.e., with more specific spectra). To model this, we'd have to know how to tune model parameters to increase efficiency.
- In conclusion, we'll use LIF neurons because they are very computationally efficient and have reasonable info transmission efficiency and accuracy.