Primer on Neural Integrators and Neural Oscillators

David J. Heeger

New York University
2018

Neural integrator

Let's start with a simple leaky neural integrator:

$$\tau \frac{dv}{dt} = -v + \lambda x + (1-\lambda) y \tag{1}$$

$$y = v\,. \tag{2}$$

The value v is the membrane potential, x is the input, y is the output firing rate (and also the recurrent drive), τ is the intrinsic membrane time-constant, and λ is (for now) a constant ($0 < \lambda < 1$). The values of v, x, and y are time-varying, i.e., $v(t)$, but I'm leaving that out as well to make the notation simpler. **Eq. 1** is a first-order, linear differential equation. It's a simplified version of the membrane equation in which I'm assuming that the resting potential is 0 and that the synaptic inputs x and y act as current injections (rather than conductance changes). **Eq. 2** assumes that the output firing rate y is equal to the underlying membrane potential v. We can of course add a proportionality constant to convert from mV to spikes/sec, but I'm leaving that out to keep it simple.

This simple form in **Eq. 1** is problematic because it allows negative firing rates (the value of y could be positive or negative depending on the input). To fix that, we use a complementary pair of neurons (analogous to ON- and OFF-center retinal ganglion cells) that receive complementary copies of the input, x and $-x$, and in which the firing rates are a halfwave-rectified copy of the underlying membrane potential fluctuations (optionally again with scale factor to convert from mV to spikes/sec):

$$\tau \frac{dv}{dt} = -v + \lambda x + (1-\lambda)\left(y^{+} - y^{-}\right), \tag{3}$$

where

$$\begin{aligned} y^{+} &= \lfloor v \rfloor = \max(v,0) \\ y^{-} &= \lfloor -v \rfloor = \max(-v,0) \\ y &= y^{+} + y^{-} \\ v &= y^{+} - y^{-} \end{aligned}\,. \tag{4}$$

A leaky neural integrator is a shift-invariant, linear system. In signal processing terms, it acts as a recursive low-pass filter with an exponential impulse response function. **Fig. 1** plots the responses y to a step input z, for each of several values of the modulator λ. The value of the modulator λ changes the effective time constant. We can see this by rewriting either **Eq. 1** (or equivalently **Eq. 3**) and simplifying:

$$\tau \frac{dv}{dt} = \lambda\left(x - v\right) \tag{5}$$

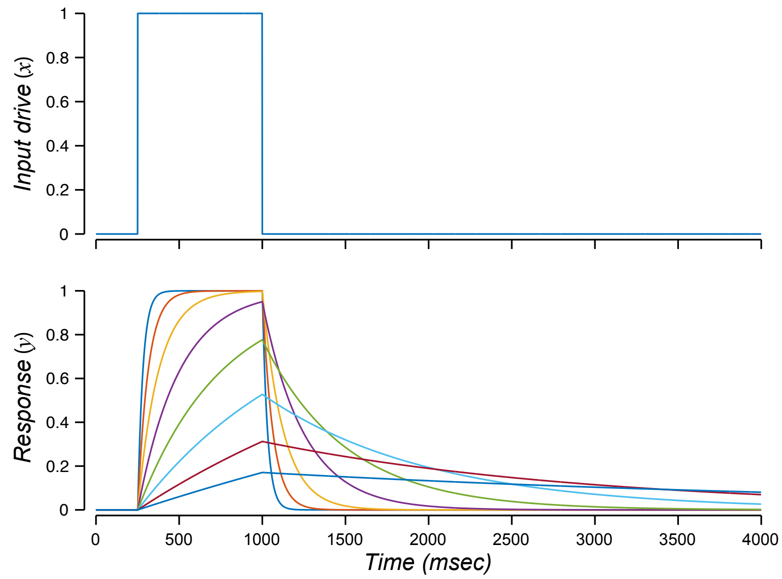

Figure 1. Leaky neural integrator. Top panel, time-course of input. Bottom-panel, response time-courses for various values of the modulator λ.

i.e.,

$$\tau' \frac{dv}{dt} = -v + x \text{ such that } \tau' = \frac{\tau}{\lambda} . \tag{6}$$

The value of τ is the intrinsic time-constant and τ' is the effective time-constant. If the input drive x is constant over time, then the responses y exhibit an exponential time course with steady state $y = x$, and time constant τ'_x. If λ is large (close to 1) then the responses follow the input such that the response increases rapidly when the input is turned on, and the responses decrease rapidly when the input is turned off. If λ is small (close to 0) then the responses exhibit sustained activity after the input is turned off.

Sustained activity

We can get the best of both worlds by allowing the value of the modulator λ to vary over time (**Fig. 2**). In this example, the value of λ was equal to 1 until just before the input was about to be turn off (i.e., for $t < 1000$) and λ was equal to 0 after the input turned off. The responses followed the input for $t < 1000$ because λ was large (=1, corresponding to a short effective time

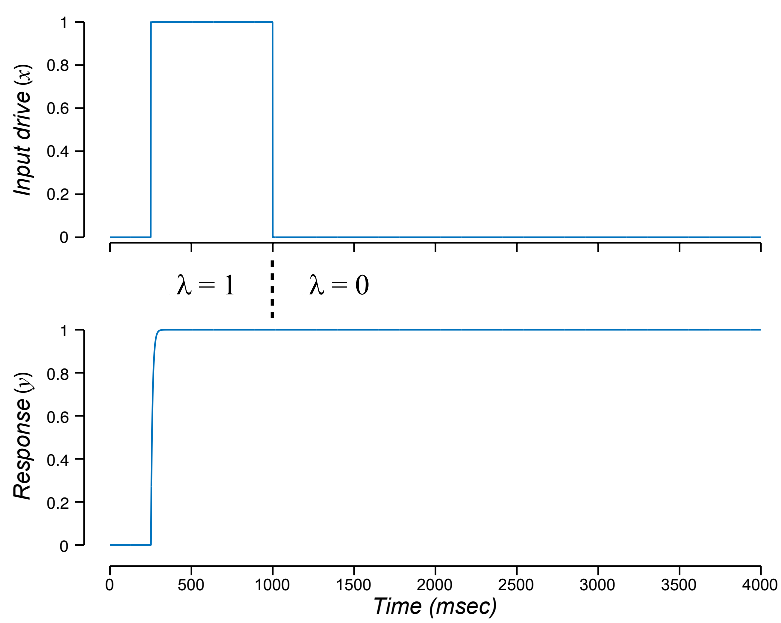

Figure 2. Sustained activity. Top panel, time-course of input drive. Bottom-panel, response time-course. $\lambda = 1$ for $t < 1000$ and $\lambda = 1$ for $t \geq 1000$.

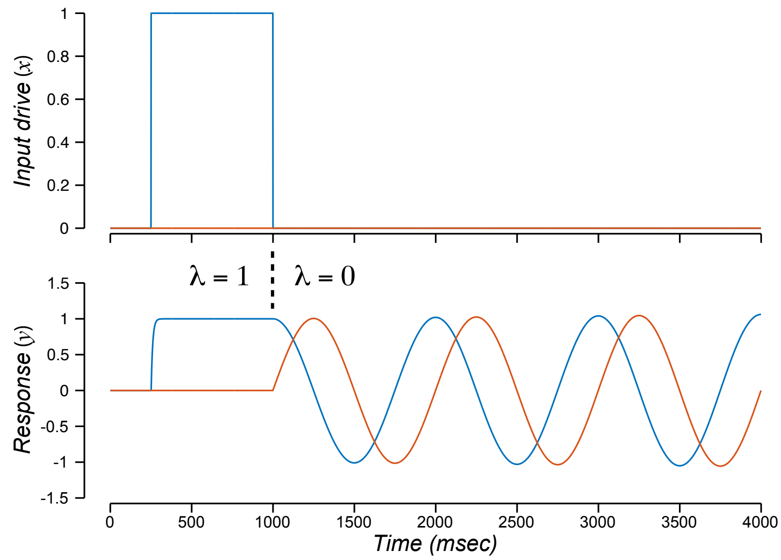

Figure 3. Neural oscillator. Top panel, time-course of input drive. Colors, input drive to each of the two neurons. Bottom-panel, response time-course. Colors, responses of each of the two neurons. $\lambda = 1$ for $t < 1000$ and $\lambda = 0$ for $t \geq 1000$.

constant). The value of λ then switched to be small (=0, corresponding to a long effective time constant) before the input turned off, so the output responses exhibited sustained activity (**Fig. 2**).

Neural oscillator

A generalization of this includes a network with more than 1 neuron such that the output response of each neuron depends on a weighted sum of all the other neurons:

$$\tau \frac{d\mathbf{v}}{dt} = -\mathbf{v} + \lambda \mathbf{x} + (1-\lambda)\hat{\mathbf{y}} \tag{7}$$

$$\hat{\mathbf{y}} = \mathbf{W}\mathbf{y} \,. \tag{8}$$

We use boldface lowercase letters to represent vectors and boldface uppercase to denote matrices. The variables ($\mathbf{v}$, $\mathbf{y}$, $\hat{\mathbf{y}}$, $\mathbf{x}$) are each functions of time. The time-varying output responses are represented by a vector $\mathbf{y} = (y_1, y_2, \ldots, y_j, \ldots, y_N)$ where the subscript j indexes different neurons in the network. The time-varying inputs are represented by another vector $\mathbf{x} = (x_1, x_2, \ldots, x_j, \ldots, x_M)$. The recurrent drive $\hat{y}_j$ to each neuron is a weighted sum of the outputs, and the weights are given by the recurrent weight matrix $\mathbf{W}$. We can use the same trick as above to ensure non-negative firing rates. A simple example corresponds to when the recurrent weight matrix $\mathbf{W}$ is the identity matrix. For that special case, you have a collection of independent neural integrators, each with their own input, but all with the same effective time constant.

A neural oscillator (**Fig. 3**) corresponds to the special case with two neurons that are mutually interconnected, with a particular form for the recurrent weight matrix:

$$\mathbf{W}\mathbf{y} = \begin{pmatrix} 1 & -2\pi\omega\tau \\ 2\pi\omega\tau & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} , \tag{9}$$

In this case there are two neurons y_1 and y_2 that are mutually interconnected. The value of ω determines the oscillation frequency. Intuitively, the responses oscillate because the recurrent weight matrix approximates a rotation matrix for small angles of rotation, i.e., $\cos(\theta) \approx 1$ and $\sin(\theta) \approx \theta$ when θ is small. If, at one instant in time, $\mathbf{y} = (1 \ 0)^t$, then an instant later the responses will have changed akin to rotating slightly around the unit circle.

Figure 4. Stability of sustained activity. Responses when $\lambda = 0$ for 3 different values of the recurrent weight. Yellow, responses grow without bound when recurrent weight is greater than 0. Blue, responses decay to zero when recurrent weight is less than 0. Orange, responses are constant over time when recurrent weight is equal than 0.

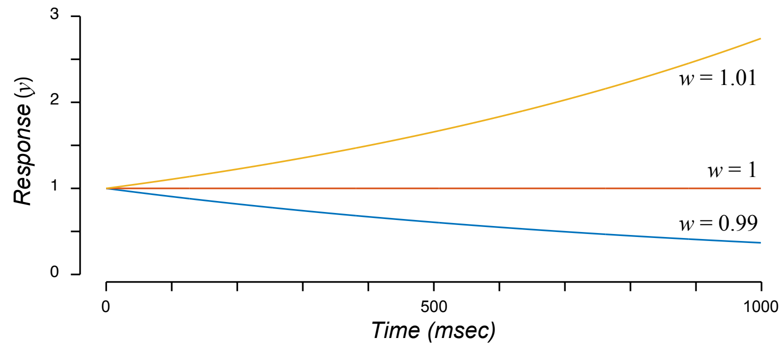

Sustained activity and stable oscillations

The dynamics of the responses depend on the recurrent weight matrix $\mathbf{W}$ (**Fig. 4**). This is particular important when $\lambda = 0$ (corresponding to the sustained activity in **Fig. 2** and the stable oscillating activity in **Fig. 3**). If the recurrent weights are too small then the responses decay over time. If the weights are too big then the responses grow without bound. A simple is example is given by a network with 3 neurons and a diagonal recurrent weight matrix:

$$\mathbf{W} = \begin{pmatrix} 1.01 & 0 & 0 \\ 0 & 1.0 & 0 \\ 0 & 0 & 0.99 \end{pmatrix} . \tag{10}$$

For this case, **Eqs. 7-8** simplify:

$$\tau \frac{dv_j}{dt} = -v_j + \lambda x_j + (1-\lambda)\hat{y}_j \ , \tag{11}$$

$$\hat{y}_j = w_j v_j , \tag{12}$$

where the recurrent weights w_j are the elements along the diagonal of the recurrent weight matrix $\mathbf{W}$. When $\lambda = 0$, this equation simplifies further:

$$\tau \frac{dv_j}{dt} = -v_j + w_j v_j \ . \tag{13}$$

For $w_j = 1$, the responses are constant over time (the derivative in **Eq. 13** is 0). For $w_j > 1$, the response grow over time. And for $w_j < 1$, the response decay over time.

In general, for an arbitrary recurrent weight matrix, the dynamics of the responses depend on the eigenvalues and eigenvectors of the recurrent weight matrix $\mathbf{W}$. When the eigenvectors and eigenvalues of the recurrent weight matrix are composed of complex values, the responses exhibit oscillations. For example, the recurrent weight matrix in the neural oscillator example (**Fig. 3**) is an anti-symmetric, 2x2 matrix (**Eq. 9**), with complex-valued eigenvalues and eigenvectors. The real-parts of the eigenvalues determine stability. In this case, the real parts of the eigenvalues are equal to 1 (the weight matrix was in fact scaled so that the eigenvalues had real parts that were equal to 1). The corresponding eigenvectors define an orthonormal coordinate system (or basis) for the responses. The responses during the period of stable oscillations (when $\lambda = 0$) are determined entirely by the projection of the initial values (the responses just before the input was turned off) onto the eigenvectors. Eigenvectors with corresponding eigenvalues equal to 1 are sustained. Those with eigenvalues less than 1 decay to zero (smaller eigenvalues decay more quickly). Those with eigenvalues greater than 1 grow without bound (which is why the weight matrix was scaled so that the largest eigenvalues = 1). The imaginary part of the eigenvalues of the recurrent weight matrix (in this example equal to $2\pi\omega\tau$) determine the oscillation frequency ω.