Lecture 10
Network models

jochen.braun@nat.uni-magdeburg.de

January 13, 2005

Goals:

1) Validity of firing rate models
2) Formalization of firing rate models
3) Example of feedforward network

Credits:

Dayan & Abbot, Chapter 6

1 The need to simplify

A typical network of cortical neurons involves at a minimum some millions of excitatory and inhibitory neurons (150,000 per mm^3), each with complex dendrites and axons, with many types of membrane channels for Na^+, K^+, Ca^{2+}, etc. and receiving input and emitting output through 8,000 synapses on average. Models of this scale are extremely expensive and unwieldy (many time-scales!). Somehow we need to capture the function with a smaller model involving fewer dynamic units.

To simplify, we can consider the firing rate of sub-populations rather than the spiking of neurons. For example, we can consider a cortical column as a single model node. However, we loose effects that depend on precise spike timing.

2 Firing-rate models

For one such model node, we define the input (pre-synaptic) firing rate vector $\vec{u}(t)$ and the output (post-synaptic) firing rate $v(t)$. The total synaptic current I_s mediates between input and output rates. When can we ignore spikes and use rates? To find out, look closely first at one particular synapse b, with post-synaptic current $K_s(t)$ and weight w_b, and then at all such synapses together. The total current produced by pre-synaptic spikes i

$$I_b(t) = w_b \sum_i K_s(t - t_i) = w_b \int_{-\infty}^{t} K_s(t - \tau)\, \rho_b(\tau)\, d\tau \qquad \rho_b(\tau) = \sum_i \delta(\tau - t)$$

$$I_s(t) = \sum_b I_b(t) = \sum_b w_b \int_{-\infty}^{t} K_s(t - \tau)\, \rho_b(\tau)\, d\tau$$

where we assume that all synapses are independent. We want, but don't yet dare, to replace spike train $\rho_b(t)$ by spike rate $u(t)$.

Upon reflection, we find that we may ignore "spike train variability" when it is smoothed out by (i) slow synapse potentials $K_s(t)$ or (ii) many uncorrelated variabilities. The latter is true only as long a neural firing remains uncorrelated!

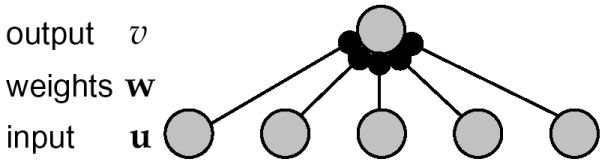

Figure 7.1: Feedforward inputs to a single neuron. Input rates $\mathbf{u}$ drive a neuron at an output rate v through synaptic weights given by the vector $\mathbf{w}$.

3 Time constants

The synaptic kernel $K_s(t)$ describes the current pulse at the postsynaptic soma caused by a presynaptic spike. It reflects the time-course of the synaptic conductance and the passive and active properties of dendritic cables. A common approximation is a current step followed by exponential decay with time-constant τ_s:

$$K_s(t) = \frac{1}{\tau_s} \exp\left(-\frac{t}{\tau_s}\right)$$

In this case the total synaptic current can be described by a differential equation:

$$I_s = \sum_b w_b \int_{-\infty}^{t} K_s(t - \tau)\, u_b(\tau)\, d\tau \qquad K_s(t) = \frac{1}{\tau_s} \exp\left(-\frac{t}{\tau_s}\right)$$

$$\tau_s \frac{dI_s}{dt} = -I_s + \sum_b w_b\, u_b$$

For an electrotonically compact dendrite, τ_s reflects the time-constant of the synaptic conductance, which may be as short as a few milliseconds (e.g. AMPA glutamate receptors). For a synapse on the distal part of a thin dendrite, it may be larger. Measured values are typically small (milliseconds).

Having described I_s in terms of input firing rates u, we need to complete the model by computing the ouput firing rate v. For constant input currents I_s, the steady-state firing rate is described by the so-called "activation function":

$$v = F(I_s)$$

$$F(I_s) = [I_s - \gamma]_+ \qquad \gamma = \text{threshold}$$

$$F(I_s) = v_{max} \frac{I_s^\alpha}{I_{thresh}^\alpha + I_s^\alpha}$$

In the general case, when the input current changes with time, we need to consider the somatic membrane capacitance and resistance, which limit how closely the firing rate v can follow the input current I_s. Roughly speaking, v will be a low-pass filtered version of I_s and can be modelled as

$$\tau_r \frac{dv}{dt} = -v + F(I_s)$$

Don't be deceived by the familiar shape of this equation! Here, v is a firing rate and NOT a membrane potential. Accordingly, the time-constant τ_r is NOT the membrane time-constant. Most network models use a value of τ_r that is considerably less than the membrane time-constant. More detailed simulations show that the effective value of τ_r depends on the firing-rate regime. Low-pass filtering by the membrane time-constant can be neglected if the neuron is always firing (i.e., always close to threshold) but must be taken into account otherwise.

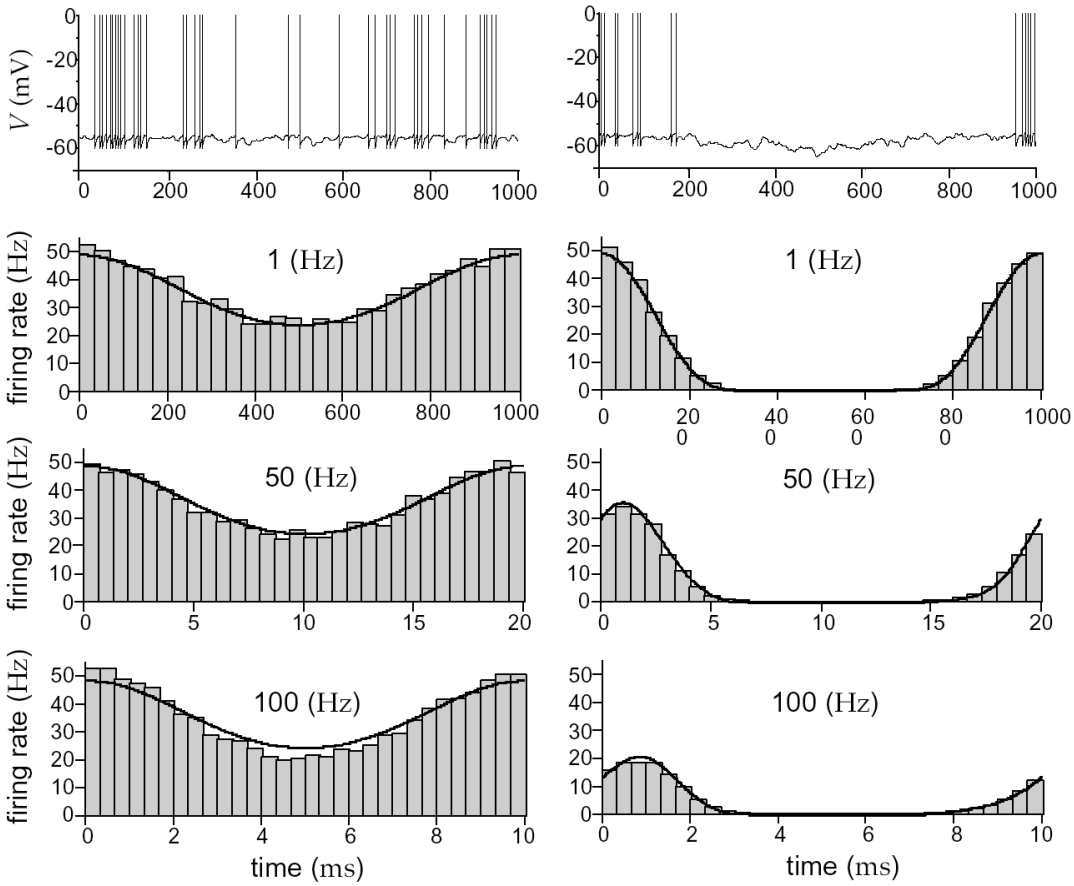

Figure 7.2: Firing rate of an integrate-and-fire neuron receiving balanced excitatory and inhibitory synaptic input and both constant and sinusoidally varying injected current. For the left panels, the constant component of the injected current was adjusted so the firing never stopped during the oscillation of the varying part of the injected current. For the right panel, the constant component was lowered so the firing stopped during part of the cycle. The upper panels show two representative voltage traces of the model cell. The histograms beneath these traces were obtained by binning spikes generated over multiple cycles. They show the firing rate as a function of the time during each cycle of the injected current oscillations. The different rows show 1, 50, and 100 Hz oscillation frequencies for the injected current. The solid curves show the fit of a firing-rate model that involves both instantaneous and low-pass filtered effects of the injected current. (Adapted from Chance *et al*, 2000.)

4 Compact firing-rate models

We now consider ways of combining our equations for I_s and v into a more compact model:

$$\tau_s \frac{dI_s}{dt} = -I_s + \sum_b w_b u_b = -I_s + \boldsymbol{w} \cdot \boldsymbol{u} \qquad\qquad \tau_r \frac{dv}{dt} = -v + F(I_s)$$

If $\tau_r >> \tau_s$, we can replace I_s by its equilibrium value

$$I_s = \sum_b w_b u_b = \boldsymbol{w} \cdot \boldsymbol{u} \qquad \Rightarrow \qquad \tau_r \frac{dv}{dt} = -v + F\left(\sum_b w_b u_b\right) = -v + F(\boldsymbol{w} \cdot \boldsymbol{u})$$

Alternatively, if $\tau_r << \tau_s$, we can replace v by its equilibrium value $F(I_s)$ and use

$$\tau_s \frac{dI_s}{dt} = -I_s + \sum_b w_b u_b = -I_s + \boldsymbol{w} \cdot \boldsymbol{u} \qquad \text{with} \qquad v = F(I_s)$$

In both cases, the steady-state firing rate v_∞ is given by

$$v_\infty = F\left(\sum_b w_b u_b\right) = F(\boldsymbol{w} \cdot \boldsymbol{u})$$

These approximations suffice to show the computational potential of network models. More accurate models would explicitly model individual spikes. However, provided that the spikes in a spiking model do not synchronize, the predictions of rate models are typically quite accurate.

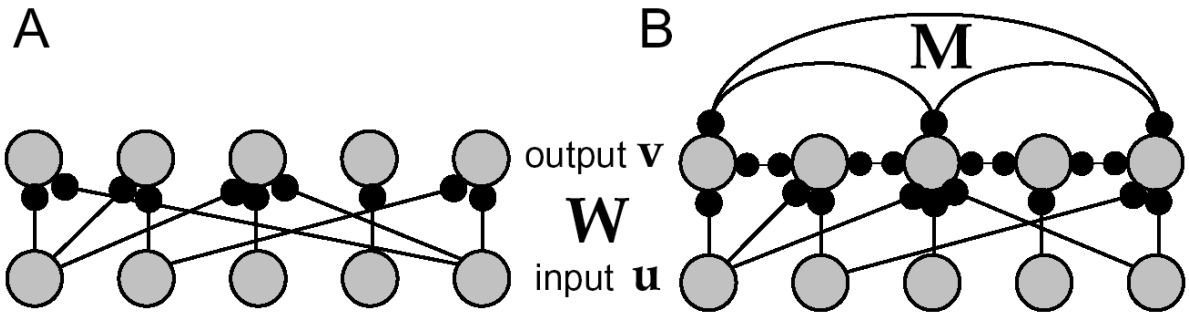

Figure 7.3: Feedforward and recurrent networks. A) A feedforward network with input rates **u**, output rates **v**, and a feedforward synaptic weight matrix **W**. B) A recurrent network with input rates **u**, output rates **v**, a feedforward synaptic weight matrix **W**, and a recurrent synaptic weight matrix **M**. Although we have drawn the connections between the output neurons as bidirectional, this does not necessarily imply connections of equal strength in both directions.

5 Feedforward and recurrent networks

In a feedforward network, N_u input units with rates $\boldsymbol{u}$ are driving N_v output units with rates $\boldsymbol{v}$, it is convenient to use vector notation and assemble the synaptic weights w_{oi} (i.e., from input unit i to output unit o, not the order of indices) into a matrix $\boldsymbol{W}$. The output rates are then determined by a system of equations as follows:

$$\tau_r \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{F}(\boldsymbol{W} \cdot \boldsymbol{u}) \qquad \text{or} \qquad \tau_r \frac{dv_o}{dt} = -v_o + F\left(\sum_i W_{oi}\, u_i\right)$$

In a recurrent network, there exist additional interconnections between output neurons, which are described by a synaptic matrix $\boldsymbol{M}$. Matrix element $M_{aa'}$gives the weight of the connection from output unit a' to output unit a (again note order of indices). The equations governing this recurrent network are:

$$\tau_r \frac{d\boldsymbol{v}}{dt} = -\boldsymbol{v} + \boldsymbol{F}(\boldsymbol{M} \cdot \boldsymbol{v} + \boldsymbol{W} \cdot \boldsymbol{u}) \qquad \text{or} \qquad \tau_r \frac{dv_o}{dt} = -v_o + F\left(\sum_{o'} M_{oo'}\, v_{o'} + \sum_i W_{oi}\, u_i\right)$$

In biologically realistic networks, the connectivity matrices $\boldsymbol{W}$ and $\boldsymbol{M}$ exhibits certain patterns, for example, that weights originating from one neuron must all have the same sign, because individual neurons are either excitatory or inhibitory ("Dale's law"). However, connections originating from one NODE (i.e., a population of excitatory and inhibitory neurons) can easily have different signs. Accordingly, a rate model does not have to be constrained by Dale's Law.

6 Continuous labelling

If a neural population encodes a continuous sensory or motor variable (e.g., a visual orientation, an auditory frequency, or a movement direction), it is easier to identify model neurons/nodes by their preferred value rather than by an integer index. In this case, the synaptic matrix is also expressed in this way. It is often true that the synaptic weight functions depend not on the absolute variable value but only on differences between such values so that the matrix becomes symmetric. In this case:

$$W_{oi} = W(\theta_o, \theta_i) = W(\theta_o - \theta_i) \qquad W_{oo'} = W(\theta_o, \theta'_o) = W(\theta_o - \theta'_o)$$

If the number of nodes in a network is large and the density of coverage ρ_θ is high, we can approximate the sums in our system equations by integrals:

$$\tau_r \frac{dv(\theta)}{dt} = -v(\theta) + F\left\{\rho_\theta \int \left[W(\theta,\theta')\, u(\theta') + M(\theta,\theta')\, v(\theta')\right] d\theta'\right\}$$

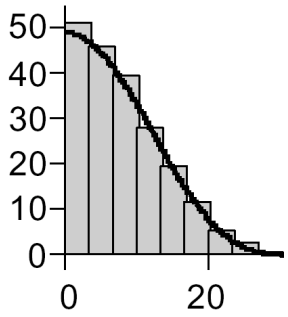

$$\sum_{i \text{ neurons}} a_i[Hz] \approx \rho \left[\frac{1}{\circ}\right] \int a(\theta)[Hz]\, d\theta[^\circ]$$

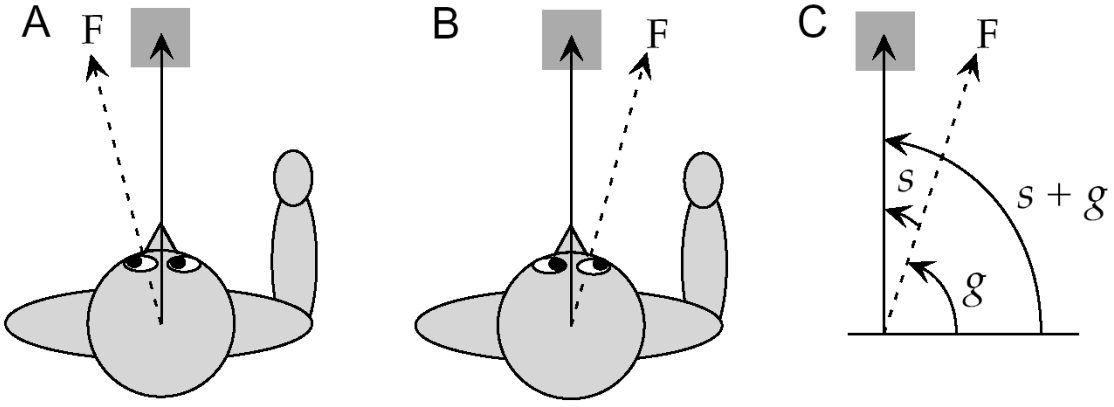

Figure 7.4: Coordinate transformations during a reaching task. A, B) The location of the target (the grey square) relative to the body is the same in A and B, and thus the movements required to reach toward it are identical. However, the image of the object falls on different parts of the retina in A and B due to a shift in the gaze direction produced by an eye rotation that shifts the fixation point F. C) The angles used in the analysis: s is the angle describing the location of the stimulus (the target) in retinal coordinates; g is the gaze direction angle, indicating the orientation of the eyes relative to the body. The direction of the target relative to the body is $s + g$.

7 Feedforward example: coordinate transform

Reaching movements necessitate coordinate transforms between retinal and body coordinates. For example, target direction in body coordinates is the sum of target direction in eye coordinates and gaze direction in body coordinates.

Visual neurons are sensitive to particular retinal locations; their receptive field is fixed in eye coordinates. In contrast, neurons in pre-motor cortex depend on the relationship between stimulus and body. With the head fixed and gaze direction is changed, the tuning curve of these motor neurons remains unchanged. However, when the gaze direction is fixed and the head is rotated, the tuning curve shifts by exactly the amount of head rotation. We conclude that the receptive fields of these pre-motor neurons are fixed in BODY coordinates (not EYE coordinates).

How can visual neurons with receptive fields in EYE coordinates drive pre-motor neurons with receptive fields in BODY coordinates?

A clue to the solution is the observation of "gaze-dependent gain-modulation" of visual neurons in area 7a. While the receptive field is fixed in EYE coordinates, the amplitude (or "gain") of the response is "modulated" by gaze position.

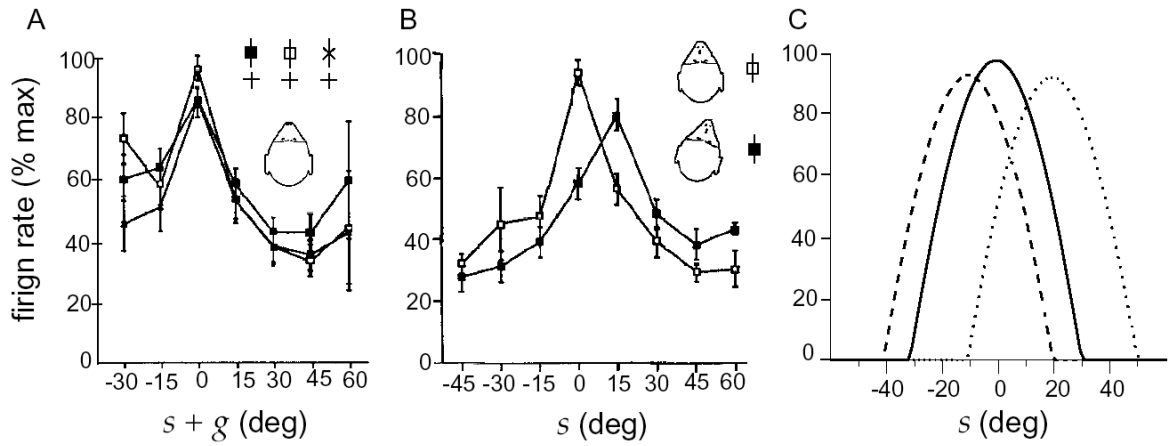

Figure 7.5: Tuning curves of a visually responsive neuron in the premotor cortex of a monkey. Incoming objects approaching at various angles provided the visual stimulation. A) When the monkey fixated on the three points denoted by the cross symbols, the response tuning curve did not shift with the eyes. In this panel, unlike B and C, the horizontal axis refers to the stimulus location in head-based, not retinal, coordinates ($s + g$, not s). B) Turning the monkey's head by 15° produced a 15° shift in the response tuning curve as a function of retinal location, indicating that this neuron encoded the stimulus direction in head-based coordinates. C) Model

Following Pouget and Sejnowski (1995), we construct an input layer of area 7a neurons whose responses depend on both stimulus position s (in retinal coordinates) and gaze position g (in body coordinates), as follows

$$u = f_u(s - \xi, g - \gamma) = \frac{\exp\left[\kappa\left(g - \gamma\right)\right]}{1 + \exp\left[\kappa\left(g - \gamma\right)\right]} \exp\left(-\frac{(s - \xi)^2}{2\sigma^2}\right)$$

where ξ and γ are the preferred stimulus position and the critical gaze position, respectively. These functions describe a Gaussian tuning for $s - \xi$ and a sigmoidal increase in response gain (or amplitude) for $g - \gamma$. These neurons are thus "stimulus-tuned" and "gaze-modulated".

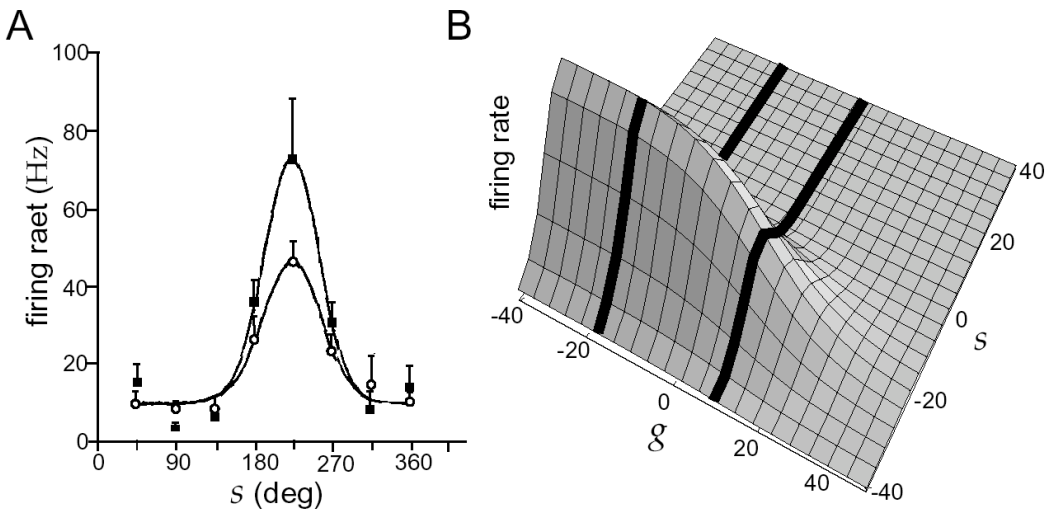

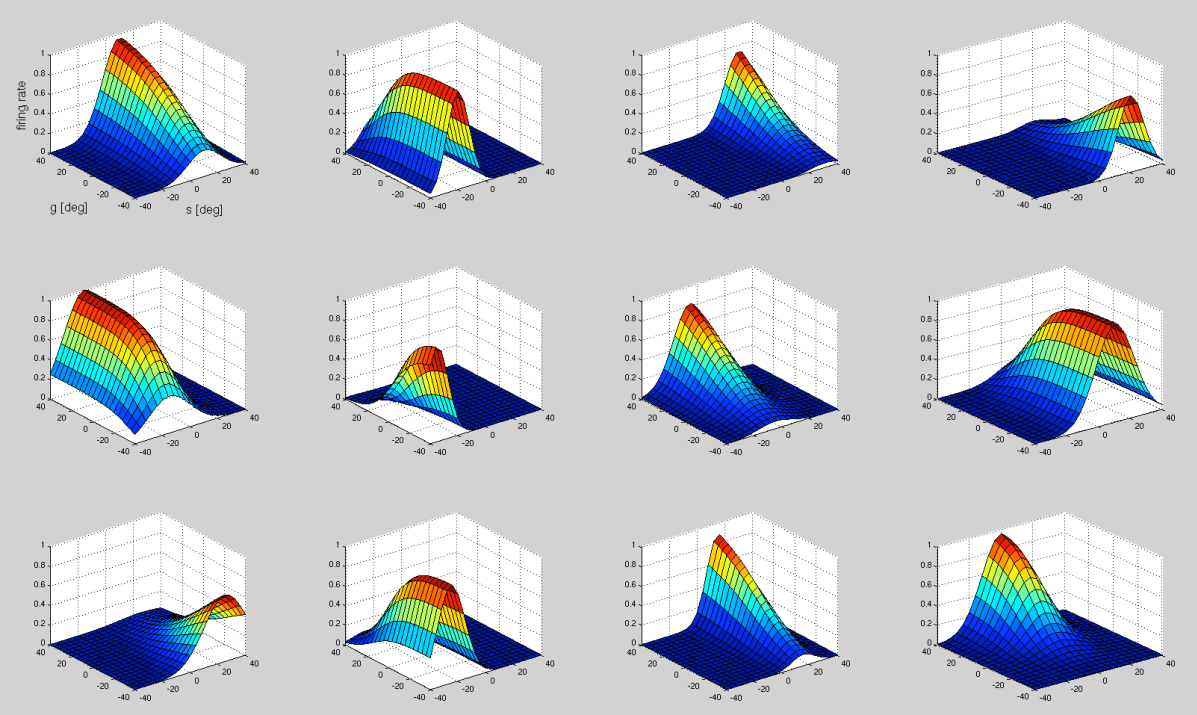

Our input layer u models visual area 7a and consists of a population of such gain-modulated units $u(\xi, \gamma)$, representing all possible combinations of ξ and γ. This population feeds via a connectivity matrix $w(\xi, \gamma)$ into an output unit v, which models a unit in pre-motor cortex. We neglect dynamic effects and consider only the steady-state response of the output unit, which is given by

$$v_\infty = F\left[\rho_\xi \, \rho_\gamma \int w(\xi, \gamma) \, f_u(s - \xi, g - \gamma) \, d\xi \, d\gamma\right]$$

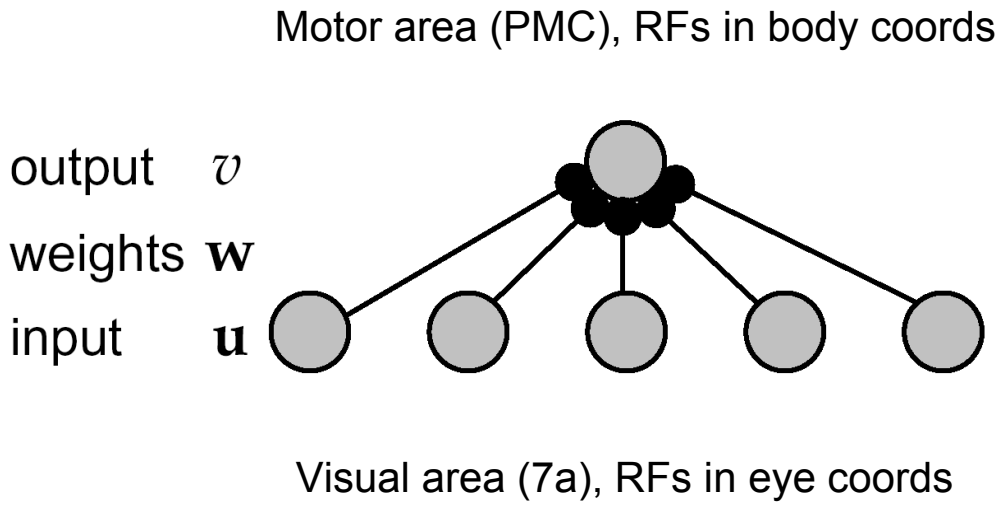

Is it possible to choose a connectivity matrix $w(\xi, \gamma)$ such that the output response is a function of $s + g$, in other words, to respond to stimulus location in BODY coordinates? To see that the answer is 'yes', we change the integration variables as follows:

$$\xi \rightarrow \xi - g$$
$$\gamma \rightarrow \gamma + g$$

$$v_\infty = F\left[\rho_\xi \, \rho_\gamma \int w(\xi - g, \gamma + g) \, f_u(s - \xi + g, -\gamma) \, d\xi \, d\gamma\right]$$

We can now see that v_∞ is a function of $s+g$ provided that $w(\xi-g, \gamma+g) = w(\xi, \gamma)$ (i.e., provided that he g-dependencies of w cancel). This, in turn, is the case if w is a function of $\xi+\gamma$, so that $w(\xi, \gamma) = w(\xi+\gamma)$.

$$w(\xi, \gamma) = w(\xi + \gamma)$$

$$v_\infty(s + g) = F\left[\rho_\xi \, \rho_\gamma \int w(\xi + \gamma) \, f_u(s + g - \xi, -\gamma) \, d\xi \, d\gamma\right]$$

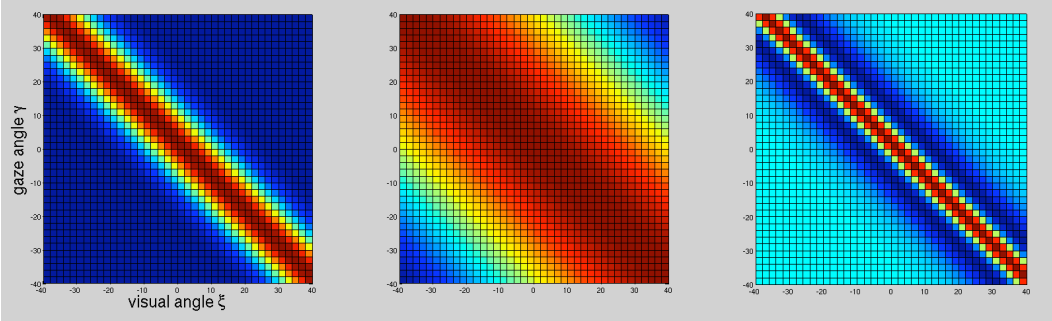

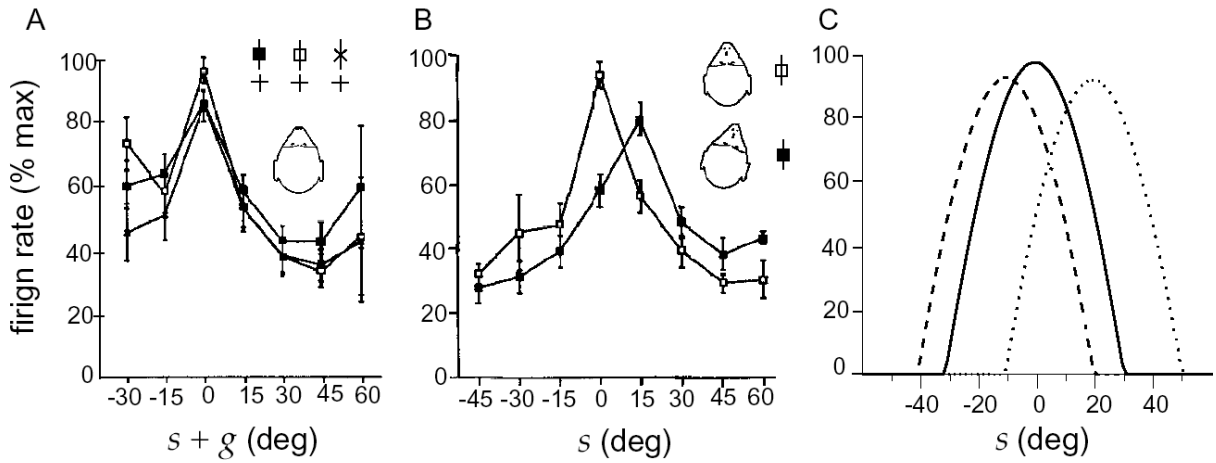

C) Model
tuning curves based on equation 7.16 shift their retinal tuning to remain constant in body-based coordinates. The solid, heavy dashed, and light dashed curves refer to $g = 0^\circ$, 10°, and -20° respectively. The small changes in amplitude arise from the limited range of preferred retinal location and gaze angles in the model. (A,B adapted from Graziano, Hu & Gross, 1997; C adapted from Salinas and Abbott, 1995.)

If synaptic weights are given by $w(\xi + \gamma$, the resulting tuning curve shifts as a function of gaze direction, but would remain constant if plotted as a function of $s + g$. Note that all three peaks occur at $s + g = 0$.

Next lecture:

recurrent networks