



Figure 7. Buildup Rate of LIP Neurons

T_{in} : neurons tuned to the direction of the stimulus. T_{out} : neurons tuned to a direction 180° away from the stimulus direction. T_{90} : neuron tuned to a direction 90° away from the stimulus direction. Blue: two-choice experiment. Red: four-choice experiment. (A) LIP data from Churchland et al., 2008. (B) Model.

optimal action selection (as we do in the SCb layer), but not optimal evidence integration (which is why we do not use it in the LIP layer). Moreover, attractor dynamics can provide an optimal solution for action selection, but, importantly, only for a limited family of distributions, one of which is Poisson-like. This is a critical point, as it emphasizes the strong link between the response distribution and optimal inference.

Our framework is sufficiently powerful that it can be extended in several directions, including incorporating prior information, dealing with time-varying stimuli, and taking into account nontrivial reward functions when selecting actions. This last extension is critical. We have shown how the evidence accumulation and the response selection can be optimized in neural circuits, but we have not shown how to optimize reward rates. Optimizing reward rate is a complex problem that depends crucially on the cost function and the stopping process (Kiani et al., 2008). This lies beyond the scope of the present paper, but it is an important issue, which we intend to explore in future studies. It remains to be seen if it can be incorporated in the PPC framework. We believe that a promising idea is to explore whether LIP encodes the expected reward as a function of saccade direction and amplitude. Recent experimental data suggest that LIP might indeed represent either expected reward for all actions or the probability that an action will maximize reward (Platt and Glimcher, 1999; Sugrue et al., 2004). Either way, our framework should be applicable, since these quantities are similar to probability distributions over stimulus values.

EXPERIMENTAL PROCEDURES

Network Simulations with LNP Neurons

The MT layer contained 100 stochastic spiking neurons with bell-shaped tuning curves to direction of motion. At each time step, the probability of a spike in neuron i was determined according to

$$p(r_i^{MT}(t_n) = 1) = [\delta t(c\Delta d \exp(K_{MT}(\cos(s_0 - s_i) - 1)) + d_{null}c + r_{spont} + n_i)]^+$$

where $r_i^{MT}(t_n)$ is the response of neuron i within the interval $[t_n - \delta t, t_n]$, s_0 is the direction of motion of the random dots, s_i is the preferred direction of neuron i , c is the percentage of dots moving in direction s_0 (the coherence level), Δd is the difference in drive between the preferred and null directions ($d_{pref} - d_{null}$), d_{null} is the drive in the null direction, r_{spont} is the spontaneous firing rate, n_i is a random variable used to induce correlations, and $[\cdot]^+$ is the threshold-linear operator: $[x]^+ = \max(0, x)$. The parameters of the model were: $d_{pref} = 0.4$, $d_{null} = -0.2$, $r_{spont} = 20$, $K_{MT} = 4$ (as reported in MT [Britten et al., 1993]), $\delta t = 1$ ms. Note that in the equation above, the coherence c cannot be factored out of the equation. As a result, the spike statistics in MT are not exactly in the Pois-

son-like family as defined in Equation 4, because the kernel $h(\cdot)$ depends on both s and c .

The noise n_i consisted of independent Gaussian noise convolved with a circular Gaussian kernel,

$$n_i = \sum_j A_{\eta_j} \exp(K_{\eta_j}(\cos(s_i - s_j) - 1))\eta_j,$$

with all η_j drawn independently from a zero-mean Gaussian distribution and K_{η_j} and A_{η_j} set to 2 and 10^{-5} , respectively. These were chosen so that the average correlation coefficients in spike trains were approximately 0.2 between neurons whose preferred directions differed by less than 90° , and approximately 0 for neurons whose preferred directions differed by more than 90° . These values are close to the ones that have been reported in vivo (Zohary et al., 1994).

In the LIP layers, we used 100 Linear-Nonlinear-Poisson (LNP) neurons. In the linear step, the membrane potential proxy of neuron i , denoted $m_i(t_n)$, is obtained from

$$m_i(t_{n+1}) = \left(1 - \frac{\delta t}{\tau}\right)m_i(t_n) + \frac{1}{\tau} \left(\sum_j W_{ij}^{ff} r_j^{MT}(t_n) + \sum_j W_{ij}^{rec} r_j^{LIP}(t_n) \right) + u(t_n) \quad (9)$$

where W_{ij}^{ff} and W_{ij}^{rec} are the matrices for the feedforward and recurrent weights, respectively, and $u(t)$ is an urgency signal (see below). The time constant, τ , was set to 1 s. The linear step is followed by a nonlinear one in which the membrane potential proxy, $m_i(t_n)$, is used to determine the probability that neuron i emits a spike between times t_n and $t_n + \delta t$,

$$p(r_i^{LIP}(t_n) = 1) = [m_i(t_n)]^+ \quad (10)$$

We used translation-invariant weights for both the feedforward and lateral connections (W_{ij}^{ff} and W_{ij}^{rec}),

$$W_{ij} = W(s_i - s_j) = a \exp(K(\cos(s_i - s_j) - 1)) + b.$$

For the feedforward weights, W_{ij}^{ff} , we used $a = 0.25$, $K = 5$ and $b = 0$, and for the lateral weights, W_{ij}^{rec} , we used $a = 0.35$, $K = 10$, and $b = -0.11$.

In the experiment of Shadlen and Newsome (2001), each trial starts with the appearance of M visual targets, where M is the number of choices. This triggers a response in the subset of LIP neurons whose receptive fields overlap with the visual targets. To model this activity, we initialize the firing rate of the neurons in the LIP layer according to

$$p(r_i^{LIP}(t_1)) = \frac{2}{(M/2 + 0.5)} \sum_{m=1}^M p_0 \exp(K_0(\cos(s_m - s_i) - 1))$$

where M is the number of possible directions for the moving dots and $\{s_1, \dots, s_M\}$ are the positions of the targets corresponding to the M choices. The parameters were set to $p_0 = 0.042$ (corresponding to a firing rate of 42 spikes/s for $\delta t = 1$ msec) and $K_0 = 4$. Given the width of these tuning curves (determined by the parameter K_0), the resulting population activity is almost perfectly flat for $M \geq 8$. Accordingly, we used $M = 8$ when simulating ‘‘continuous’’ decision making. This visually induced activity served as a starting point for the accumulation of evidence.