Behavioral/Systems/Cognitive

Climbing Neuronal Activity as an Event-Based Cortical Representation of Time

Jan Reutimann,[1] **Volodya Yakovlev,**[2] **Stefano Fusi,**[1] **and Walter Senn**[1]

[1]Institute of Physiology, University of Bern, 3012 Bern, Switzerland, and [2]Department of Neurobiology, Institute of Life Science, Hebrew University, Jerusalem, 91904, Israel

The brain has the ability to represent the passage of time between two behaviorally relevant events. Recordings from different areas in the cortex of monkeys suggest the existence of neurons representing time by increasing (climbing) activity, which is triggered by a first event and peaks at the expected time of a second event, e.g., a visual stimulus or a reward. When the typical interval between the two events is changed, the slope of the climbing activity adapts to the new timing. We present a model in which the climbing activity results from slow firing rate adaptation in inhibitory neurons. Hebbian synaptic modifications allow for learning the new time interval by changing the degree of firing rate adaptation. This event-based representation of time is consistent with Weber's law in interval timing, according to which the error in estimating a time interval is proportional to the interval length.

Key words: delay activity; neuronal clock; firing rate adaptation; Hebbian synaptic plasticity; learning cellular timers; Weber law in interval timing

Introduction

In many areas of the brain neural activity triggered by sensory input can persist for several seconds after the stimulus is removed. Such persistent activity is thought to underlie working memory and has been found in prefrontal cortex (Fuster, 1973; Funahashi et al., 1989; Quintana and Fuster, 1992; Miller et al., 1996; Rainer et al., 1999), inferotemporal cortex (Fuster, 1982; Yakovlev et al., 1998), parietal cortex (Quintana and Fuster, 1992; Constantinidis and Steinmetz, 1996; Chafee and Goldman-Rakic, 1998), and perirhinal cortex (Erickson and Desimone, 1999). Persistent activity can be more than just a static snapshot of sensory input from the past; often the delay activity is observed to increase or decrease monotonically over time (Kojima and Goldman-Rakic, 1982; Sakai and Miyashita, 1991; Quintana and Fuster, 1992; Miller et al., 1996; Chafee and Goldman-Rakic, 1998). Climbing activity, i.e., delay activity that steadily increases between two subsequent stimuli, has been linked with the anticipation of various events, like a motor response (Constantinidis and Steinmetz, 1996), the end of the delay interval (Romo et al., 1999; Reutimann et al., 2001), or the identity of the sample or match stimulus in delayed matching-to-sample (DMS) tasks [retrospective vs prospective coding (Rainer et al., 1999; Mongillo et al., 2003)]. Increasing delay activity also can be associated with reward expectation, for instance in prefrontal cortex (Wa-

tanabe, 1996), striatum (Kawagoe et al., 1998; Hassani et al., 2001), and thalamus (for review, see Schultz, 2000; Komura et al., 2001).

Recent experiments have addressed the question of how a change in the duration of the delay period is reflected in the time-varying delay activity (Kojima and Goldman-Rakic, 1982; Komura et al., 2001; Brody et al., 2003b). In these experiments the build-up of activity is stretched in time, rather than shifted. The stretching causes the slope of the activity profile to decrease with the length of the delay period. This is in agreement with the scaling property of interval timing found in psychophysical studies on humans (Rakitin et al., 1998) and recently confirmed by *in vivo* experiments with monkeys (Leon and Shadlen, 2003). When the duration of a time interval is estimated, the error distribution scales linearly with the length of the interval. Although psychophysics furnished many important results about the phenomenology of interval timing, the cellular mechanisms by which time is represented and handled in the brain are only beginning to be discovered (Buonomano and Karmakar, 2002; Durstewitz, 2003; Gallistel, 2003).

We present a model of how climbing activity can form an adaptable, neuronal representation of the time passing between two events. The model suggests a cellular implementation of a cortical clock that is very general and applicable to the different experimental data cited above where increasing or decreasing delay activity is observed. However, to confine ourselves to a specific example, we put the model in the context of delay activities measured in inferotemporal (IT) cortex of macaque monkeys performing a delayed matching-to-sample task (Yakovlev et al., 1998). Although this experiment originally was designed for a different study, for two cells the length of the inter-stimulus interval (ISI) has been changed during the course of the experiment. Both of these cells show an ISI-dependent delay activity

that increases from the sample to the test stimulus, and that may reflect the expectation time of the test stimulus (see also Reutimann et al., 2001). Our goal is not to present data that give a definitive answer to the question of how time is represented in IT during the DMS task but to study quantitatively the conditions for generating and modifying such a possible representation. Because the behavior of the two recorded cells is particularly evident and consistent with other observations (Kojima and Goldman-Rakic, 1982; Komura et al., 2001; Brody et al., 2003b), it may represent a general paradigm for the neural encoding of time.

The typical ISIs are of the order of seconds, both in our experimental data and in the literature. This fact presents a serious modeling problem; it is not clear how neurons and synapses operating on a millisecond time scale can encode information about time intervals on the order of seconds. Moreover, how can such an internal representation of time be adapted to reflect changes in the stimulus interval duration? To answer these questions, we consider a model of interacting populations that are subject to neuronal adaptation and synaptic plasticity. In this model the delay activity that decreases across the ISI results from slow firing rate adaptation, as observed in intracellular recordings *in vitro*; if a neuron is injected with strong currents, it will respond with spike trains of decreasing frequency and eventually will stop firing (Rauch et al., 2003). Climbing activity is attributable to the firing rate adaptation in inhibitory interneurons. Hebbian synaptic plasticity may tune the climbing activity to the recent ISIs and may underlie the learning of time intervals.

Materials and Methods

Neuron and network simulations. We consider different populations of leaky integrate-and-fire neurons (see Fig. 2): sustained delay activity (sDA) neurons, excitatory (Ex) neurons showing climbing activity, inhibitory (Inh) neurons showing decreasing activity, and spontaneously firing background (BG) neurons projecting onto the Ex and Inh populations (not depicted in Fig. 2). In the limit of many presynaptic neurons and small asynchronous synaptic potentials, the total synaptic current driving a postsynaptic neuron becomes Gaussian-distributed, with mean μ and variance σ^2. The leaky integrate-and-fire neuron converts this current into a firing rate (in units of spikes/sec $\equiv$ Hz) according to Amit and Brunel (1997a,b) and Rauch et al. (2003):

$$\nu = F(\mu, \sigma) = \left(\tau_{rp} + \tau_m \int_{\frac{V_{reset}-\mu}{\sigma}}^{\frac{V_\theta-\mu}{\sigma}} \sqrt{\pi} \cdot e^{-2}(1 + erf(z))dz \right)^{-1},$$

(1)

where for notational convenience the currents and voltages are expressed in the same units (mV). The formula gives the neuronal firing rate in the steady state as a function of the spike emission threshold V_θ, the reset potential V_{reset}, the refractory period τ_{rp}, and the membrane time constant τ_m.

The mean and variance of the input current from population x onto a single neuron of population y (with x standing for Ex, Inh, sDA, and BG and y standing for Ex and Inh) are given by:

$$\mu_{y,x} = c N_x \nu_x J_{y,x} \tau_m , \sigma_{y,x}^2 = \mu_{y,x}J_{y,x} .$$

(2)

Here N_x is the number of neurons in population x, c (the "connectivity") is the probability that a neuron of population x is connected to a neuron of population y, ν_x is the firing rate of a neuron in x, and $J_{y,x}$ is the synaptic strength of a connection from x to y. The time constant τ_m in the expression for the mean current $\mu_{y,x}$ imposes the correct units (because the above formula for the firing rate considers the individual postsynaptic currents as delta functions, the only (positive) time constant entering in

Table 1. Stimulation parameters

Parameter	Value	Description and unit
ν_{sDA}^{rest}	2	Spontaneous activity of sDA neurons (spikes/sec)
ν_{sDA}^{da}	30	Delay activity (constant during ISI) (spikes/sec)
ν_{sDA}^{vr}	70	Visual response of sDA neurons (spikes/sec)
ν_{BG}^{rest}	11	Activity of background neurons (spikes/sec)
ν_{BG}^{vr}	14	Visual response of background neurons (spikes/sec)
c	0.1	Connectivity (probability of synaptic contact)
N_{sDA}	1000	Number of sustained delay activity neurons
N_{Ex}	1000	Number of excitatory ("expectation") neurons
N_{Inh}	1000	Number of inhibitory neurons
N_{BG}	10000	Number of background neurons
$J_{Ex,BG}$	0.082	Synaptic strength BG $\rightarrow$ Ex (mV)
$J_{Ex,sDA}$	0.05	Synaptic strength sDA $\rightarrow$ Ex (mV)
$J_{Ex,Ex}$	0.01	Synaptic strength Ex $\rightarrow$ Ex (mV)
$J_{Ex,Inh}$	−0.1	Synaptic strength Inh $\rightarrow$ Ex (mV)
$J_{Inh,BG}$	0.08	Synaptic strength BG $\rightarrow$ Inh (mV)
$J_{Inh,sDA}$	0.6	Synaptic strength sDA $\rightarrow$ Inh (mV)
$J_{Inh,Inh}$	−0.2	Synaptic strength Inh $\rightarrow$ Inh (mV)
τ_{rec}	200	Recovery-from-adaptation time constant (msec)
τ_{net}	10	Network time constant (msec)
μ_0^a	22	Adaptation threshold (mV)
Q^a	100	Adaptation constant (sec mV)
r	0.001	Learning rate (mV)
Θ_{pre}	10	Presynaptic learning threshold (spikes/sec)
Θ_{post}	12	Postsynaptic LTP/LTD threshold (spikes/sec)
τ_m	20	Membrane time constant, all neurons (msec)
V_Θ	20	Spike emission threshold, all neurons (mV)
V_{reset}^{Ex}	15	Reset potential Ex neurons (mV)
V_{reset}^{Inh}	0	Reset potential Inh neurons (mV)
τ_{rp}^{Ex}	5	Refractory period Ex neurons (msec)
τ_{rp}^{Inh}	20	Refractory period Inh neurons (msec)

the model is τ_m). The total input current onto population y is obtained by summing over all presynaptic populations,

$$\mu_y = \sum_x \mu_{y,x} \text{ and } \sigma_y^2 = \sum_x \sigma_{y,x}^2 .$$

The expected mean firing rate of the population $y = $ Ex is given by $\nu_y = F(\mu_y, \sigma_y)$. For the inhibitory population $y = $ Inh, the firing rate is $\nu_y = F(\mu_y - \mu^a, \sigma_y)$, where μ^a is an adaptation current as defined below. The evolution of the population firing rate of Ex and Inh follows:

$$\frac{d\nu}{dt} = \frac{F(\mu, \sigma) - \nu}{\tau_{net}},$$

(3)

with τ_{net} being the network time constant and μ the total input current. The neurons in the sDA and BG population are assumed to fire with some fixed (Poisson) rates at rest and in response to the visual stimulus. Parameter values for the full network simulations underlying Figures 3, 5, and 6 (with the indicated changes) are given in Table 1.

Firing rate adaptation. The behavior of a neuron, as observed in the *in vitro* experiments (see Fig. 3a,b), suggests that it is able to respond with a sustained activity up to some critical input current (adaptation threshold), above which the neuron gradually loses its firing ability, to the point at which it completely fails to emit action potentials. The gradual decline depends on how much the input exceeds the critical value. The neuron recovers quickly when the input current falls below the adaptation threshold.

To account for these observations, we model an input-dependent adaptation current μ^a, which reduces the synaptic input current, μ_y, to give a mean $\mu = \mu_y - \mu^a$ of the total current. The dynamics of μ^a depends on the synaptic input current. If μ_y is below some adaptation threshold, $\mu_y < \mu_0^a$, the adaptation current decays with a time constant τ_{rec} toward zero. If the synaptic input current is above this threshold, $\mu_y \geq \mu_0^a$, the adaptation current builds up proportional to the excess $\mu_y - \mu^a$, with an adaptation time constant τ_{adapt}. This adaptation time constant is mono-

tonically decreasing with μ_y, $\tau_{adapt} = \tau_{rec} + Q^a/[\mu_y - \mu_0^a]$, where Q^a is some adaptation constant. The stronger the synaptic input current, the faster the adaptation current builds up, until it eventually cancels the input current (i.e., until $\mu^a = \mu_y$). In summary, we have:

$$\frac{d\mu^a}{dt} = \begin{cases} \dfrac{-\mu^a}{\tau_{rec}}\mu_y - \mu^a & \mu_y < \mu_0^a \\ \dfrac{Q^a}{\tau_{rec} + \dfrac{Q^a}{\mu_y - \mu_0^a}} & \mu_y \geq \mu_0^a \end{cases}. \qquad (4)$$

For simplicity we assume that only inhibitory cells undergo slow adaptation (y = Inh). This is justified if inhibitory cells receive relatively strong inputs and therefore show stronger firing rate adaptation. Slow firing rate adaptation also can be more prominent in inhibitory than in excitatory cells in response to the same input current. When cortical neurons are stimulated *in vitro* for several seconds, late adaptation acting on a time scale of 1–4 sec is found both in inhibitory and excitatory neurons. Preliminary data from our lab show that this late adaptation in fact is more pronounced for inhibitory neurons, although they can sustain higher firing rates. Whereas in adolescent rats 34 of 44 recorded layer 5 pyramidal cells exhibited a firing rate decay of 0.5–1 Hz/sec with an initial rate of ~30 spikes/sec (Rauch et al., 2003), five of five inhibitory cells exhibited an adaptation of 1–4 Hz/sec starting from an initial rate of 60 spikes/sec (personal communication by G. La Camera, A. Rauch, and D. Thurbon). Other works show that inhibitory low-threshold spiking neurons slowly adapt their firing rate, for instance by 39% of the initial rate within 600 msec (Gibson et al., 1999) or by 38% in 1 sec for 37% of the sampled inhibitory neurons (Gupta et al., 2000). The parameter values chosen for the firing rate adaptation of the model cell (Table 1) produce a behavior that is consistent with these data. The phenomenon of slow frequency adaptation might be attributable to slow sodium channel inactivation (Fleidervish et al., 1996; Powers et al., 1999), to calcium-activated potassium channels, or to a change in the electrogenic sodium–potassium pump (Sanchez-Vives et al., 2000).

Synaptic plasticity. Modifications of the synaptic strength from the sDA to the Inh population occur only if the presynaptic firing rate exceeds a threshold θ_{pre}. The change of the synaptic strength is proportional to the excess $\nu_{pre} - \theta_{pre}$, with a proportionality constant r ("learning rate"). The sign of the change depends on whether the postsynaptic firing rate exceeds a threshold θ_{post} or not. If $\nu_{post} > \theta_{post}$, the synapse undergoes long-term potentiation (LTP); if $\nu_{post} < \theta_{post}$, it undergoes long-term depression (LTD). Hence,

$$\frac{dJ}{dt} = r \cdot \lfloor \nu_{pre} - \Theta_{pre} \rfloor \cdot sign\,(\nu_{post} - \Theta_{post}), \qquad (5)$$

where $\lfloor x \rfloor = 0$ for $x < 0$ and $\lfloor x \rfloor = x$ else. Parameter values are given in Table 1. Note that the postsynaptic activity contributes only to the sign to the synaptic modification (although this is not reflected directly in the sketch of Fig. 4d). More realistic rules that further differentiate the strength of the postsynaptic current will similarly ensure the convergence toward a balanced LTP/LTD state. For instance, a synaptic modification that is a sigmoidal function of $\lfloor \nu_{pre} - \theta_{pre} \rfloor \cdot (\nu_{post} - \theta_{post})$ will stabilize itself even if θ_{pre} would be above the presynaptic frequency of the sustained delay activity, and learning would take place only during the time of the stimulus presentation. The tuning mechanism would still work because the degree of adaptation at the end of the delay interval is reflected in the strength of the Inh response to the second stimulus (compare Fig. 4d). The saturation of LTP at high values of $\lfloor \nu_{pre} - \theta_{pre} \rfloor \cdot (\nu_{post} - \theta_{post})$ helps to limit the positive feedback loop that might, during the presentation of the first stimulus, counteract the stabilizing negative feedback caused by adaptation during the second stimulus.

Experimental procedure and data analysis. When the monkey signaled to be ready to perform a trial of the DMS task, a sample stimulus was presented on a computer monitor (0.5 sec); after a fixed delay interval (usually 5 or 8 sec) a test stimulus (0.5 sec) appeared (see Fig. 1a). The test stimulus matched the sample stimulus in one-half of the trials. After a variable post-test stimulus interval of 0.5–1.5 sec, the monkey had to shift a lever to the left (match) or to the right (non-match) and was rewarded

with fruit juice if performing correctly (>90%). For further details, see Yakovlev et al. (1998).

To obtain the activity plots in Figure 1c, we calculated mean firing rates across bins of 250 msec and averaged over the indicated number of trials (n; see legend of Fig. 1). For neurons showing stimulus-selective delay activity we averaged over trials with the same pair of sample and test stimuli (see Fig. 1c1). For neurons showing unselective delay activity, different sample stimuli were pooled together (see Fig. 1c2,c3). Bars in Figure 1d represent mean firing rates within 0.5 sec bins, again averaged over five consecutive trials. To extract the slope of the activity increase, we fit the first 5 sec of the delay activity linearly (compare Fig. 5a, right). The slopes of these linear fits were averaged over five consecutive trials and plotted as a function of the trial number (see Fig. 1e).

Testing Weber's law. To test the scalar variability of interval timing (see Fig. 6), we performed reduced and full network simulations. In the reduced simulations (see Fig. 6b) the noise in the Ex population firing rate was modeled by a stochastic process with a correlation length of 2 msec, a linearly increasing mean, and a variance proportional to the mean firing rate $\nu_{Ex}(t)$ (with a proportionality constant equal to one over the number of Ex neurons). The slope was tuned such that the first passage times for crossing the 40 spikes/sec threshold were centered on the onset of the test stimulus (i.e., 5 and 8 sec after the sample stimulus, respectively). The readout threshold was distributed normally at ~40 spikes/sec. In the full network simulations we considered 250 neurons in the Inh, Ex, and sDA population and 2500 neurons in the BG population (sample traces of the Ex population firing rate are shown in Fig. 6a). The parameter values were the same as in the corresponding meanfield simulations of these networks. The synaptic strengths to the Inh neurons were adjusted to obtain peak activities at 5 and 8 sec, respectively.

Results

Expectation effect in delayed matching-to-sample tasks

When an animal has learned to perform a temporal prediction task repeatedly, it likely has formed some internal representation of the time passing toward the predicted event. Cells in IT may be able to represent such a passage of time between two events, as suggested by the behavior of two paradigmatic neurons in IT recorded in an experiment studying selective delay activity (Yakovlev et al., 1998). Many IT neurons during this experiment showed a gradually increasing (climbing) activity within the inter-stimulus interval of a delayed matching-to-sample task that started after the sample stimulus and peaked at the time of the test stimulus (Fig. 1a,c2). In general, neurons with three different behaviors during the ISI were observed (Fig. 1c): (1) stimulus-selective neurons with nearly constant, sustained delay activity, (2) neurons with climbing (increasing) activity, and (3) neurons with decreasing activity during the delay interval. The increasing/decreasing activity neurons tended to be unselective, i.e., they were each responding with the same temporal profile and with a similar cell-specific strength to all 30 of the different visual stimuli.[a] For two cells showing increasing activity, the length of the delay interval was changed from 5 to 8 sec and back again after several trials (Fig. 1d1–d5). These cells showed a steadily climbing activity, starting with low activity at the beginning of the delay interval and peaking at the time when the test stimulus was presented. In the first few trials after the interval had been changed to 8 sec, the activity still increased during 5 sec but then stayed constant afterward. This biphasic dynamics disappeared several trials later when the activity increase became slower but lasted now for 8 sec, the length of the delay interval. When the monkey

---

[a]The mean delay activity was not significantly selective according to a statistical test based on a one-way ANOVA. The probabilities that for the two cells of Figure 1e the mean delay frequencies were the same for all of the stimuli and that the measured values were different from stimulus to stimulus simply because of statistical fluctuations were $p = 0.63$ and $p = 0.23$, respectively. Usually these probabilities were $<10^{-4}$ for clearly selective cells (Yakovlev et al., 1998).
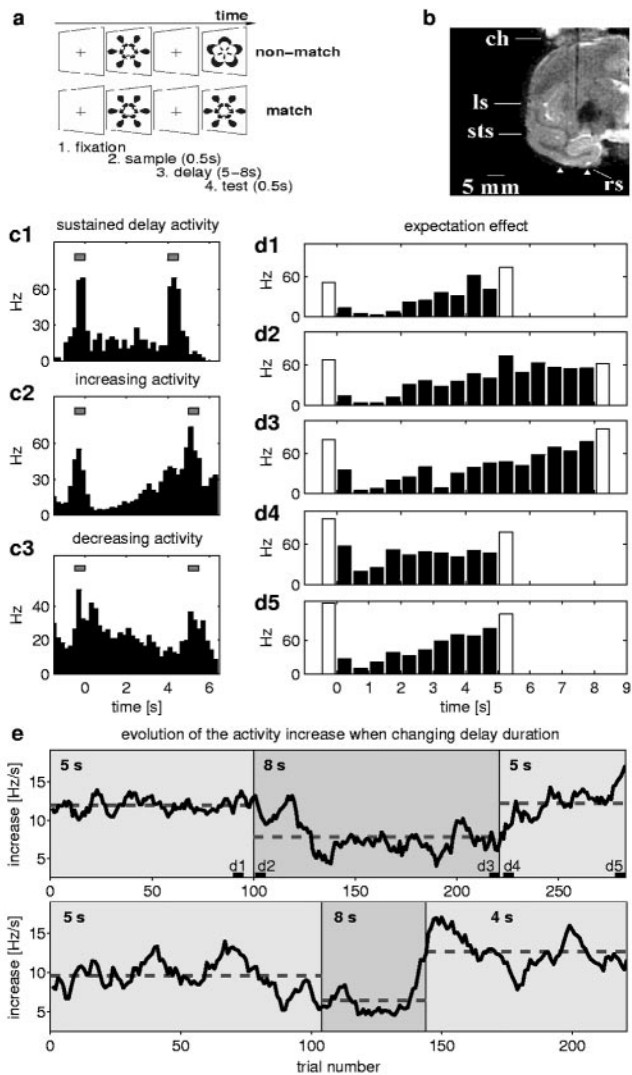
Figure 1. Expectation effect in IT: climbing activity always peaks at the expected time of the test stimulus. *a*, In the DMS task, a sample stimulus is followed by a test stimulus after 5 or 8 sec, with the test stimulus being either a match (bottom) or a non-match (top). *b*, A coronal MRI image of the right hemisphere of the brain of one of the studied monkeys (from Yakovlev et al., 1998). The tip of the electrode was in the rhinal sulcus (rs) in the area between the two arrows. *c*, Example behaviors of three classes of IT cells (average of *n* trials). *c1*, Stimulus-selective sustained delay activity (*n* = 2). *c2*, Unselective, increasing activity (*n* = 150). *c3*, Unselective, decreasing activity (*n* = 24). Horizontal bars indicate the time of the sample and test stimulus presentations. *d*, Effect of changing the delay interval from 5 to 8 sec and back for the cell shown in *c2*. The slope of activity increase is adjusted to the delay interval (empty bars, stimulus presentation; dark bars, delay interval). *d1*, Last trials with a 5 sec interval: increasing activity throughout the interval. *d2*, First trials with a 8 sec interval: increasing activity for the first 5 sec and constant activity thereafter. The cell still "expects" the test stimulus to appear after 5 sec. *d3*, Last trials with a 8 sec interval: increasing activity throughout the interval. *d4*, The shallow slope remains after changing the interval back to 5 sec. *d5*, Last trials with the 5 sec interval: again increasing activity, peaking at the test stimulus onset. *e*, Slope of activity increase obtained from a linear fit of the firing rate across the corresponding delay interval, plotted for two cells (top and bottom panels). Dashed lines represent average slopes over trials with the same ISI length (5, 8, or 4 sec). Trials shown in *d1* – *d5* are marked by horizontal bars in the top panel.

again was confronted with a 5 sec interval, the familiar situation with a 5-sec-long climbing activity again could be observed after a transient of a few trials. Plotting the slope of the activity profile against the trial number shows that the slope adaptation maximally requires some tens of trials (see Fig. 5*e*). These properties of

the adapting climbing activity may be connected to the monkey's expectation of an important event (the test stimulus) to occur at some time in the near future.

Slow firing rate adaptation determines the dynamics of delay activity

The three classes of neurons with sustained, increasing, and decreasing activity were interconnected mutually in a model of IT (Fig. 2). We assume that stimulus-selective sustained delay activity (sDA) populations have formed, showing constant delay activity throughout the ISI. Each of these sDA populations is selective to a specific, familiar visual stimulus. The existence of such populations has been postulated previously, and cells have been found in IT that behave in this classical "working memory"-like way (Fuster, 1982; Miyashita and Chang, 1988). Stimulus-selective sDA populations are assumed to drive the inhibitory (Inh) and the excitatory (Ex) populations, which show decreasing and increasing activity during the delay interval, respectively. Because these two populations receive input from all of the stimulus-selective populations, they are themselves unselective to the stimulus. The decay of the activity of the Inh neurons is caused by a slow firing rate adaptation operating on a time scale of seconds. If a stimulus is presented, the corresponding sDA population will be highly activated. Without adaptation, mutually connected populations quickly would reach persistent delay activity. Because adaptation gradually reduces the Inh firing rate, the activity of the excitatory cells Ex gradually increases because of disinhibition.

We implemented slow firing rate adaptation in the inhibitory neurons, which depends on the strength of the input current. This type of neural adaptation was observed in different classes of neocortical cells of the rat *in vitro* (Rauch et al., 2003), and it can be even more pronounced in inhibitory neurons (see Materials and Methods). Because our model requires only that the slow firing rate adaptation is stronger in the inhibitory neurons, we neglect the firing adaptation in the Ex neurons. When a prolonged noisy current of a few seconds is injected into such cells, the mean firing rate tends to decrease linearly with time (Fig. 3*a*). Strong current injection above a critical value leads a faster decline in firing rate, starting at a higher initial rate (Fig. 3*b*) (see also Fleidervish et al., 1996). Below some threshold current no adaptation is observed. The phenomenology of this firing rate adaptation also is captured by the model: the stronger the total synaptic current in the Inh neurons, the steeper the firing rate adaptation (Fig. 3*c*). Because firing rate adaptation is a function of the total synaptic current, the postsynaptic firing rate can drop to zero with constant drive, both in the data and the model. Note that this would not be possible if adaptation would originate in the postsynaptic firing rate; a drop in the firing rate then would reduce the adaptation also, and the firing rate would saturate at a positive value. Current-based adaptation is also the reason that the firing rate curves can cross, such that a stronger neuronal drive eventually will lead to a smaller firing rate (Fig. 3*c*). Although in the data these crossings are distributed on a somewhat longer time scale (Fig. 3*a*, the crossing between the dotted and dashed curve is at 25 sec), we tuned the model adaptation such that they can arise within the first 10 sec. The precise crossing times within this interval, however, are not essential [because in the network simulations the total drive of the Inh neurons varies only little (see Fig. 5*a*, filled triangles), the effective crossings always lie near each other, as indicated by the open triangles in Fig. 5*a*]. A full simulation of the interacting populations sDA,
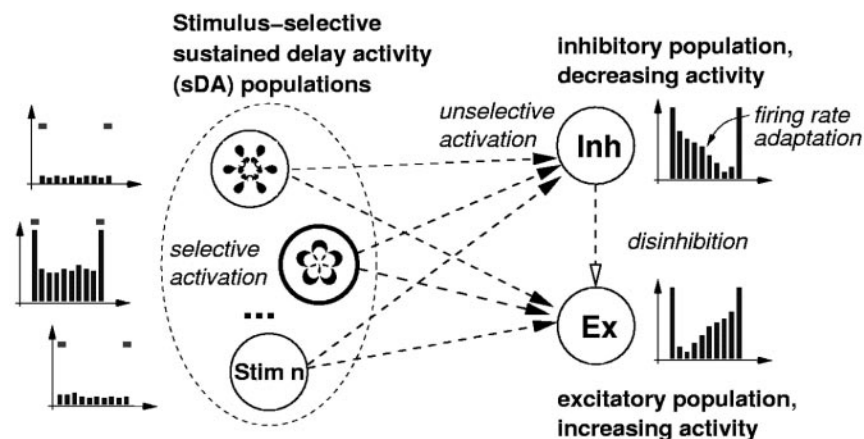
Figure 2. A model network of IT generating climbing activity. Populations of stimulus-selective neurons, displaying sustained delay activity (sDA populations), project onto unselective populations of inhibitory (Inh) and excitatory (Ex) neurons, displaying decreasing and increasing activity, respectively. With the presentation of a sample stimulus (first horizontal bar in left activity plots), the corresponding selective sDA population (marked with a thick circle) responds strongly, and it remains active during the subsequent delay interval. The other sDA populations remain spontaneously active (top and bottom left activity plots) and contribute to the background activity. Because of the firing rate adaptation, the activity of the Inh neurons decreases (top right activity plot), which in turn causes climbing activity in the Ex neurons via disinhibition (bottom right activity plot). The strong responses at the end of the delay period are caused by the test stimulus presentation (second horizontal bar in left activity plots).
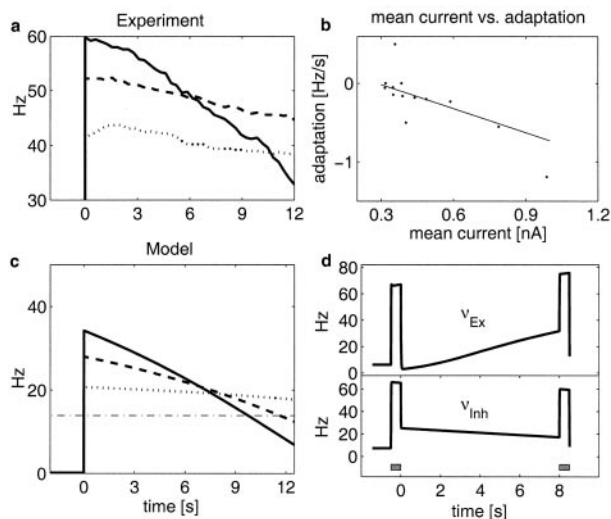

Figure 3. Firing rate adaptation produces decreasing activity. *a*, Experimental evidence for slow adaptation in neocortical cells of rat somatosensory cortex. Shown are mean firing rates (using a sliding window of 2 sec) of one cell in response to noisy currents of at least 10 sec duration, with a mean of 1.0, 0.8, and 0.6 nA, and each with an SD of 0.32 nA (from top to bottom). The firing rate adaptation is approximately linear in time and is steeper for strong input currents. *b*, Firing rate adaptation (in Hz/sec, slope of the activity time course) versus mean input current measured in the experiment: the stronger the input, the more adaptation (pooled data of 5 cells, linear fit). *c*, Response of the Inh population in the model to a 30 spikes/sec firing rate of the selected sDA population, turned on at 0 sec, with different synaptic strengths. Crossings of the firing rate curves can arise because, at any value of the mean firing rate, adaptation is based on the total synaptic input current: the stronger the input current, the more adaptation. The model adaptation is tuned such that the different crossings arise somewhere within the first 10 sec, although this feature is not always shared by the experimental data. The dashed/dotted line indicates a possible LTP/LTD modification threshold θ_{post} (which should be below the crossings; see Results). *d*, Full network simulation of a single trial, including visual stimuli. Firing rate adaptation in the Inh neurons causes disinhibition of the Ex neurons and results in decreasing and increasing activities, as seen in the data (Fig. 1*c3*, *c2*, respectively).

Inh, and Ex, with an additional background population, qualitatively reproduces the increasing and decreasing activity found in the experiments (Fig. 3*d*, as compared with the data shown in Fig. 1*c2*, *c3*). A weak firing rate adaptation in the Inh neurons is enough to

cause a pronounced firing rate increase in the Ex neurons during the delay period.

Changing the degree of adaptation via Hebbian learning

Because the synaptic drive to the Inh population determines the degree of firing rate adaptation, it can be modified by synaptic plasticity. We implemented a simple Hebbian rule that modifies the synaptic strength of the excitatory input from sustained delay activity cells onto inhibitory cells (Fig. 4*a*). According to this rule synapses are upregulated if the pre- and postsynaptic firing rate are both high and are downregulated if only the presynaptic firing rate is high but the postsynaptic is low (Fig. 4*b*). Hence, whether a synapse undergoes LTP or LTD depends on whether the postsynaptic activity is above or below some activity threshold (θ_{post}). No synaptic modification occurs if the presynaptic firing rate is below a critical level (θ_{pre}). In the simplest form of the Hebbian plasticity in which the postsynaptic firing rate determines only the sign of the synaptic modification, this critical firing rate, θ_{pre}, should be below the firing rate of the presynaptic delay activity such that there is also some learning during the delay interval (Fig. 4*c*), although for more realistic rules this restriction is not necessary (compare with Materials and Methods).

A change in the length of the delay interval is manifested directly in a change of the LTP/LTD ratio (Fig. 4*d*). By a self-regulating process the firing rate adaptation is adjusted such that the Inh neurons activity eventually ceases at the end of the new delay interval. If the delay interval is stretched, the postsynaptic activity will stay for longer below the θ_{post} threshold, and LTD dominates LTP (Fig. 4*d*, top left). As a consequence, the total synaptic current decreases; this causes less firing rate adaptation and therefore a shallower decay of the Inh activity. After a few trials, when the delay activity just vanishes at the end of the new interval, LTP again balances LTD (top right panel). The same self-regulatory process correctly adjusts the firing rate if the delay interval is shortened (bottom panels). In this case the LTD regimen shrinks, which causes LTP to dominate. The increased synaptic current causes the firing rate adaptation to become stronger. As a result, LTD again increases, and a new LTP/LTD balance is obtained. This balance is reached approximately when the activity of the Inh neuron vanishes at the end of the new interval and hence when the Ex neurons peak at the onset of the test stimulus.

Simulations of the interacting populations validate this self-tuning property. The 120 consecutive trials of the DMS experiment were run, and the inter-stimulus interval was changed from 5 to 8 sec and back. The activity increase in the Ex population is the observable that can be compared directly with the experiments. As in the data, after the interval is stretched, the activity first saturates at the end of the old interval until it again culminates just before the second stimulus sets in (Fig. 5*a1*–*a5*, right, as compared with Fig. 1*d*). When the interval is shortened again, the increasing delay activity is cut initially at an intermediate level until it again saturates toward the onset of the test stimulus. When the slope of the Ex population activity is extracted at each
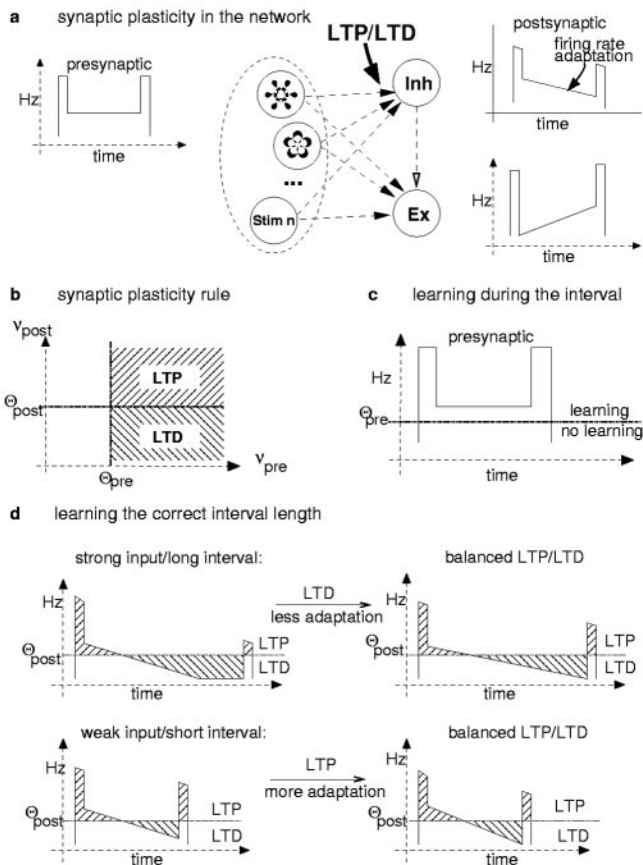
Figure 4. Hebbian plasticity adjusts the firing rate adaptation to fit the ISI length. *a*, The full model considers Hebbian synaptic plasticity from the stimulus-selective sDA population(s) to the inhibitory population Inh. *b*, Synaptic plasticity rule. If the presynaptic activity v_{pre} exceeds a threshold θ_{pre}, the synaptic strength is upregulated or downregulated (LTP/LTD), depending on whether the postsynaptic activity v_{post} is above or below θ_{post}. *c*, During the stimulation and the delay period, the activity of the presynaptic cells exceeds θ_{pre}, and synaptic modifications are triggered. *d*, Top panels, If the delay interval is too long, such as after a change in the protocol, or if the input to the Inh population is too strong, then the postsynaptic frequency stays for longer in the LTD regimen than in the LTP regimen, resulting in a net depression of the synaptic strengths. Weaker synapses imply less firing rate adaptation, and this leads to a shift toward the balanced case, in which case LTD and LTP cancel each other. Bottom panels, If, on the other hand, the input is too weak or the interval too short, then a net LTP will strengthen the synapses, causing more firing rate adaptation in the next trials. A balanced situation is reached again in which the rate of adaptation is adjusted to the interval duration.

Figure 5. The slope of the climbing activity adapts to the interval length. The simulations qualitatively reproduce the experimental data shown in Figure 1. *a*, Activity of the Inh population (left) and Ex population (right) during five of the 100 trials. The network is initialized with a synaptic strength onto the Inh population ($J_{Inh,DA}$) such that the climbing activity extends over 5 sec (*a1*). After the ISI is stretched to 8 sec, the activity still reaches a steady state after 5 sec (*a2*) but then adapts to span the full interval linearly (*a3*). The network recalibrates after switching back to 5 sec (*a4, a5*). The different initial responses of the Inh neurons (filled triangles) are approximately equalized at the beginning of the delay interval (open triangles). Dashed lines (right) represent the linear fits of the climbing activity during the 0.5– 4.5 sec interval. *b*, Evolution of the slope extracted from the fit in *a* while the delay interval is switched from 5 to 8 to 5 sec. Vertical bars indicate the trials shown in *a*. *c*, Evolution of the synaptic strength $J_{Inh,DA}$ that underlies the change of the slope.

trial during the interval stretching and squeezing protocol, a similar time course is observed as in the experiment (Fig. 5*b* as compared with Fig. 1*e*). In agreement with the data, the slope always adapts after a few presentations to the imposed ISI. The change in the slope is reflected in the corresponding change of the synaptic strength from the sustained delay activity population to the inhibitory population (Fig. 5*c*). The synaptic plasticity in these connections faithfully traces the changes in the inter-stimulus intervals with a delay of a few trials.

The self-regulation of the firing rate arises by a negative feedback that keeps the final activity at the end of the delay interval constant. To understand this self-regulation, we first note that the Hebbian plasticity by itself would lead to a positive feedback; an increased synaptic strength leads to a stronger postsynaptic response, and this further potentiates the synapses. Firing rate adaptation, however, can convert this positive feedback loop into a negative feedback loop because a stronger drive eventually leads to a smaller postsynaptic response. Such an inversion is possible
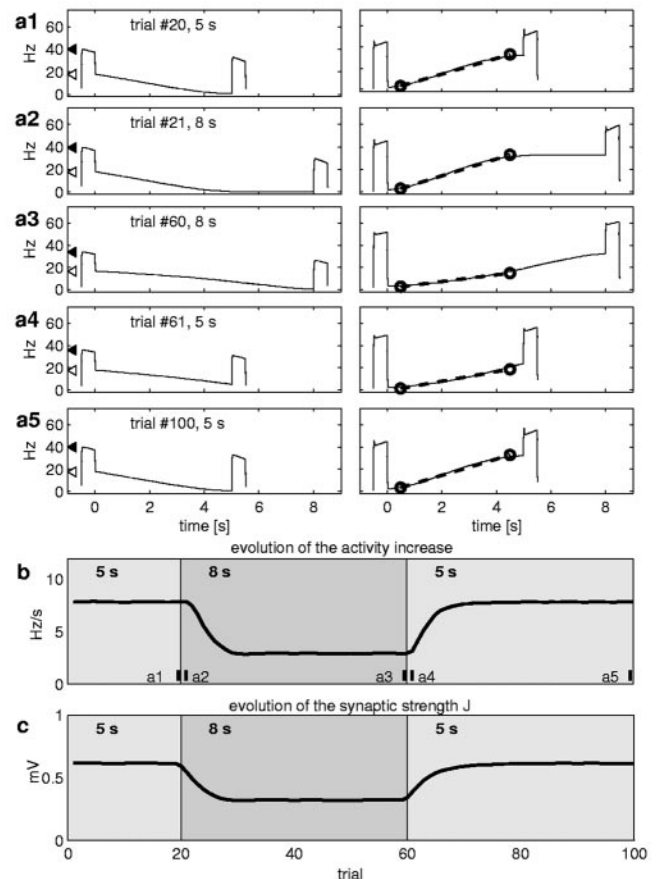
because the current-dependent adaptation leads to crossings in the firing rate curves: the stronger the input current, the earlier the postsynaptic firing rate drops to zero (Fig. 3*a,c*) and the earlier the LTD regimen is reached (below the dashed/dotted line in Fig. 3*c*, indicating a possible LTP/LTD threshold θ_{post}). Hence, a stronger drive shortens the initial LTP interval and lengthens the subsequent LTD interval, leading to an overall dominance of LTD and therefore to a dominance of the negative feedback loop.

Because of the negative feedback loop, the system always finds a steady state in which LTP and LTD are balanced across the whole duration of the stimulus presentations and the delay interval. Because adaptation leads to a linear decay of the firing rate in time, this balance is reached approximately if one-half of the time the postsynaptic activity is above and one-half of the time it is below the LTP/LTD modification threshold θ_{post}, independently of the interval length. The firing rate at the end of the delay interval remains approximately the same, because this is true for the firing rate at the beginning of the delay interval. In fact, although different synaptic strengths lead to different initial re-
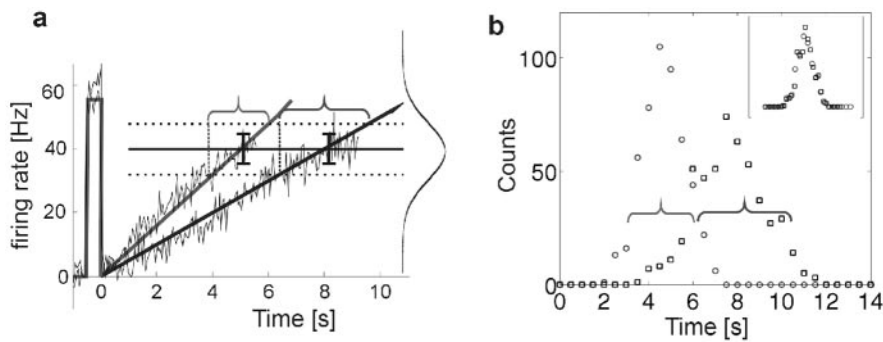
Figure 6. Weber law of interval timing: the variability of the expectation time is proportional to the length of the ISI. *a*, Downstream neurons can read out the expected time of the test stimulus from the Ex population activity by a simple threshold mechanism (thick horizontal line). Such readout neurons have to cope with the variability of the Ex population firing rate (noisy traces) and the variability of their own activation threshold (dotted horizontal lines, assumed to be Gaussian, distributed around the mean readout threshold, as indicated by the density function on the right). The firing rate variability around the mean threshold is the same for an ISI of 5 and 8 sec (two vertical error bars) but maps differently onto the variability of the corresponding readout time of the two ISIs (left and right brackets, respectively). *b*, Histograms of the threshold crossing times of 500 trials based on the reduced simulations, with an ISI of 5 sec (circles) and 8 sec (squares). The width of the histograms represents the variability of the readout times (same brackets as in *a*). In agreement with Weber's law, the variability is proportional to the ISI length, as confirmed by the superposition of the centered and normalized histograms (inset, stretching factors 1/5 and 1/8, respectively).

sponses (Fig. 5*a*, filled arrows), the response at the beginning of the delay activity is approximately equalized (Fig. 5*a*, open arrows). The apparent equalization is caused by the adaptation during the stimulus presentation, followed by the downscaling of the activity caused by the reduction of the input when the stimulus is turned off. Note that, compared with Figure 3*c*, the initial response in Figure 5*a* varies much less and that after the first 0.5 sec of the stimulus presentation the synaptic drive is reduced by more than one-half.

Timing variability is proportional to the interval length
The activity in the Ex population is well suited to represent the time between the sample and test stimulus. A simple threshold mechanism is enough to read out the expected time of the test event from the Ex activity. The readout can be realized by cells downstream of the Ex population that start firing as soon as this population crosses a certain activity threshold, say 40 spikes/sec (Fig. 6*a*).

Because the slope of the delay activity faithfully reflects the ISI length, the threshold crossings of the Ex firing rates will give accurate estimates of the changing delay intervals. Although the mean of the readout time will be centered at a new ISI after a few trials, however, some readout variability will remain. This readout variability has two sources: (1) the variability in the instantaneous spike rate of the Ex population because of the finite number of Ex neurons and (2) the stochasticity of the readout threshold itself (Fig. 6*a*). Reduced simulations considering these two noise sources show that the readout variability is always proportional to the ISI length (see Materials and Methods). In fact, the scaled histograms of the readout times for different interval lengths perfectly overlap (Fig. 6*b*). This scalar variability of the readout times also is confirmed when the stochastic simulations of the Ex population were replaced by full network simulations with integrate-and-fire neurons (sample traces in Fig. 6*a*).

To understand the linear scaling property, we first consider the variability of the readout threshold only. If the Ex population firing rate were deterministic, this variability would be mapped directly onto the variability of the readout time via linear activity ramp, as shown in Figure 6*a*. By this linear mapping the readout variability becomes indirectly proportional to the slope of the activity ramp and directly proportional to the ISI length. The

same argument also applies to the variability of the Ex population firing rate. The stochasticity of the firing rate affects the variability of the readout time in the same way as if the variability of the readout threshold would be increased by a fixed amount. Note that the firing rate variability of the Ex population is always the same when the mean firing rate reaches the mean readout threshold (40 spikes/sec). Because this fixed firing rate variability is mapped linearly onto the readout time, the readout time variability remains proportional to the ISI length.

Discussion
We have presented a general model of interacting neuronal populations that generate an event-based representation of time by slowly increasing activity. The model reproduces the expectation effect observed in IT cells during a DMS task. These cells show climbing activity during the delay interval that peaks at the time when the test stimulus is expected to occur. In agreement with the data, the slope of the climbing activity adapts within a few trials to a new ISI length.

Increasing and decreasing activities as cellular timers
It has been postulated that time-varying delay activity is a good candidate for a cortical representation of time (Hopfield and Brody, 2000; Brody et al., 2003b). Our model could be applied easily to such stimulus-dependent increasing or decreasing delay activity as it is observed in these and other works (Kojima and Goldman-Rakic, 1982; Miller et al., 1996; Romo et al., 1999; Komura et al., 2001) (for review, see Brody et al., 2003a). The model, in principle, is also compatible with a recent study that directly addresses the neuronal representation of time during active interval estimation (Leon and Shadlen, 2003). When an animal has learned to perform some task in time, it likely has formed an internal representation of the time passing toward the behaviorally relevant event, independently of whether timing is essential to perform well (as in an interval estimation task) or whether it is just useful additional information to improve the performance (as in a delayed matching-to-sample task).

Firing rate adaptation causing climbing activity
The characteristic of our model is that climbing activity in the excitatory neurons is caused by firing rate adaptation in inhibitory neurons. To serve as a basis for a self-regulation of the activity slope, however, adaptation needs to show an additional feature: the stronger the input current, the earlier the firing rate must drop below a critical level. Such an inversion is present in some of the cited data, and it is captured by the model of a current-dependent firing rate adaptation (compare Fig. 3 and Materials and Methods). It remains unclear, however, whether this behavior is shared by the specific type of interneurons involved in the generation of the delay activity. An additional requirement of the model is that the activity curves for sustained input currents approximately cross within the first 10 sec (Fig. 3*c*), although it is not important where the precise crossings are within this time span. Slowly increasing inhibition of the Inh neurons may represent a possible alternative.

Modification of the slope via synaptic plasticity

The inter-stimulus interval is represented faithfully by the climbing activity if the steady-state firing rates are reached at the time when the second stimulus appears. Because the slope of the activity decrease depends on the total postsynaptic current, it can be controlled by Hebbian plasticity. Our simplified learning rule is based on mean firing rates, but it is consistent with recent findings on spike-timing-dependent plasticity (Sjöström et al., 2001). LTP is triggered if both the pre- and postsynaptic firing rates are high, whereas LTD is triggered if only the presynaptic, but not the postsynaptic, firing rate is high. Such a rule can be implemented easily in terms of spike-driven synaptic dynamics (Fusi et al., 2000; Senn, 2002). The interesting point is that it is possible to adjust the timing in the range of seconds by a rule that is purely local in time and that cannot sense the time span between the two visual stimuli.

The suggested mechanism for the slope adaptation is simple and robust. The synaptic strength onto the inhibitory neurons is modified such that the activity across the delay interval always becomes balanced around the LTP/LTD threshold (θ_{post}; Fig. 4d). In this balanced state the initial and final firing rates of the climbing activity are independent of the length of the inter-stimulus interval. Although the exact value of θ_{post} is not crucial, an additional slow synaptic process, as realized in the Bienenstock–Cooper–Munro rule (Bienenstock et al., 1982) or in a nonlinear form of spike-timing-dependent plasticity (Senn et al., 2001), is required to bring θ_{post} into the regimen of the adapted firing rates of the inhibitory neurons (Fig. 3c).

Adapting the slope of the delay activity via Hebbian plasticity, however, may limit the speed by which this slope can change. In the present data set the slope, and therefore the hypothetical synaptic strength from the sDA to the Inh neurons, adapts to the new time interval within a few trials. This relatively slow adaptation might be ascribed to the fact that the switch between the two intervals was performed only twice during the whole experiment. If it were repeated several times, the monkey could learn the two delay intervals and, in principle, realize after a single trial that the time interval has changed. There are preliminary data in a slightly different experimental context (Brody et al., 2003a) suggesting that this might happen. In this case, different and mutually inhibiting Inh populations might be learned, one coding for a short and one coding for a long time interval. After the segregation into these two Inh subpopulations, a small bias triggered by a change of the ISI may shift the activity quickly from one to the other subpopulation. It would be interesting to study the conditions for such a symmetry breaking in the sDA-to-Inh connections and in the Inh population.

Adaptive climbing activity and the Weber law of interval timing

Psychological studies refer to the scalar variability, or the Weber law, of interval timing (Staddon et al., 1999): individuals, when estimating a time interval (without counting), tend to make errors that are distributed around the exact interval duration with a width proportional to the interval length. If the error distribution is scaled proportionally to the interval length to be estimated, these scaled distributions overlap almost perfectly (Gibbon, 1977; Rakitin et al., 1998), evidence for a relative, scaled representation of time intervals (for a theoretical analysis of the scaling law, see Gibbon, 1992). This scaling property of interval timing might be a corollary of the event-based representation of time via climbing activity; the relative times between the triggering event (sample stimulus) and the test event (test stimulus) are mapped

linearly onto the firing rates of the Ex neurons. As a consequence, the variability of the readout times scales with the length of the interval (Fig. 6), as observed in the psychophysical experiments.

Synaptic facilitation fails to reproduce the experimental data

There might be alternative models explaining climbing activity; in a previous version we considered synaptic facilitation in the feed-forward and recurrent connections onto the Ex population as a source of the slowly increasing activity (Reutimann et al., 2001). One consequence of that model is that the visual response of the Ex neurons to the second (test) stimulus should be stronger in the match than in the non-match case. This is because the feed-forward connections from the stimulus-selective population would facilitate during the sustained delay activity and thus fire vigorously if the same visual stimulus were presented again. However, additional inspection of the total 500 trials showed that the magnitude of the second response is independent of a match or non-match. This makes it unlikely that climbing activity is caused by synaptic facilitation. In turn, if the reason for the increased/decreased activity is in the postsynaptic neuron, as it is the case for firing rate adaptation, no difference between match and non-match is seen. The indifference arises because the postsynaptic neuron itself is not stimulus-selective, although the different projections onto that neuron are (Fig. 4a).

Experimental prediction: stimulus-dependent expectation

One prediction is that the slope of the decreasing activity in inhibitory neurons should adapt similarly to a change in the ISI length as the slope of the increasing activity in the excitatory neurons does. Another prediction of our model is that each sample stimulus can be associated with its own expectation time for the test stimulus. Because the learning rule affects only the synapses that are selectively active during the delay period (Fig. 4a), the adjustment of the slope to a specific delay interval becomes itself stimulus-selective. Our model predicts that changing the delay interval for one sample stimulus will not affect the expectation time that follows a different sample stimulus. For example, one may present for 20 trials sample stimulus A, followed by a delay interval of 5 sec, and stimulus B, followed by a delay interval of 8 sec. After these trials the very same cell would show climbing activity that would peak after 5 sec in response to stimulus A and after 8 sec in response to stimulus B. This example would show that different event-based representations of time may coexist in a single cell.

References

Amit D, Brunel N (1997a) Dynamics of a recurrent network of spiking neurons before and following learning. Network 8:373–404.

Amit D, Brunel N (1997b) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb Cortex 7:237–252.

Bienenstock E, Cooper L, Munro P (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. J Neurosci 2:32–48.

Brody C, Romo R, Kepecs A (2003a) Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. Curr Opin Neurobiol 13:204–211.

Brody C, Hernandez A, Zainos A, Lemus L, Romo R (2003b) Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. Cereb Cortex 13:1196–1207.

Buonomano D, Karmakar U (2002) How do we tell time? Neuroscientist 8:42–51.

Chafee M, Goldman-Rakic P (1998) Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. J Neurophysiol 79:2919–2940.

Constantinidis C, Steinmetz M (1996) Neural activity in posterior parietal

area 7a during the delay periods of a spatial memory task. J Neurophysiol 76:1352–1355.

Durstewitz D (2003) Self-organizing neural integrator predicts interval time through climbing activity. J Neurosci 23:5342–5353.

Erickson C, Desimone R (1999) Responses of macaque perirhinal neurons during and after visual stimulus association learning. J Neurosci 19:10404–10416.

Fleidervish I, Friedmann A, Gutnick M (1996) Slow inactivation of Na$^+$ current and slow cumulative spike adaptation in mouse and guinea-pig neocortical neurones in slices. J Physiol (Lond) 493:83–97.

Funahashi S, Bruce C, Goldman-Rakic P (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J Neurophysiol 61:331–349.

Fusi S, Annunziato M, Badoni D, Salamon A, Amit DJ (2000) Spike-driven synaptic plasticity: theory, simulation, VLSI implementation. Neural Comput 12:2227–2258.

Fuster J (1973) Unit activity in prefrontal cortex during delayed-response performance: neural correlates of transient memory. J Neurophysiol 36:61–78.

Fuster J (1982) Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. J Neurosci 2:361–375.

Gallistel C (2003) Time has come. Neuron 38:149–150.

Gibbon J (1977) Scalar expectancy theory and Weber's law in animal timing. Psychol Rev 84:279–325.

Gibbon J (1992) Ubiquity of scalar timing with a Poisson clock. J Math Psychol 35:283–293.

Gibson J, Beierlein M, Connors B (1999) Two networks of electrically coupled inhibitory neurons in neocortex. Nature 402:75–79.

Gupta A, Wang Y, Markram H (2000) Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. Science 287:273–278.

Hassani O, Cromwell H, Schultz W (2001) Influence of expectation of different rewards on behavior-related neuronal activity in the striatum. J Neurophysiol 85:2477–2489.

Hopfield J, Brody C (2000) What is a moment? Cortical sensory integration over brief intervals. Proc Natl Acad Sci USA 97:13919–13924.

Kawagoe R, Takikawa Y, Hikosaka O (1998) Expectation of reward modulates cognitive signals in the basal ganglia. Nat Neurosci 1:411–416.

Kojima S, Goldman-Rakic P (1982) Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. Brain Res 248:43–49.

Komura Y, Tamura R, Uwano T, Nishijo H, Kaga K, Ono T (2001) Retrospective and prospective coding for predicted reward in the sensory thalamus. Nature 412:546–549.

Leon M, Shadlen M (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. Neuron 38:317–327.

Miller E, Erickson C, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J Neurosci 16:5154–5167.

Miyashita Y, Chang H (1988) Neuronal correlate of pictorial short-term memory in the primate temporal cortex. Nature 331:68–70.

Mongillo G, Amit DJ, Brunel N (2003) Retrospective and prospective persistent activity induced by Hebbian learning in a recurrent cortical network. Eur J Neurosci 18:2011–2024.

Powers R, Sawczuk A, Musick J, Binder M (1999) Multiple mechanisms of spike-frequency adaptation in motoneurones. J Physiol (Lond) 93:101–114.

Quintana J, Fuster J (1992) Mnemonic and predictive functions of cortical neurons in a memory task. NeuroReport 3:721–724.

Rainer G, Rao SC, Miller E (1999) Prospective coding for objects in primate prefrontal cortex. J Neurosci 19:5493–5505.

Rakitin C, Gibbon J, Penney T, Malapani C (1998) Scalar expectancy theory and peak-interval timing in humans. J Exp Psychol Anim Behav Process 24:15–33.

Rauch A, La Camera G, Lüscher H-R, Senn W, Fusi S (2003) Neocortical pyramidal cells respond as integrate-and-fire neurons to *in vivo*-like input currents. J Neurophysiol 90:1598–1612.

Reutimann J, Fusi S, Senn W, Yakovlev V, Zohary E (2001) A model of expectation effects in inferior temporal cortex. Neurocomputing 38–40:1533–1540.

Romo R, Brody C, Hérnandez A, Lemus L (1999) Neural correlates of parametric working memory in the prefrontal cortex. Nature 399:470–473.

Sakai K, Miyashita Y (1991) Neural organization for the long-term memory of paired associates. Nature 354:152–155.

Sanchez-Vives M, Nowak L, McCormick D (2000) Cellular mechanisms of long-lasting adaptation in visual cortical neurons *in vitro*. J Neurosci 20:4286–4299.

Schultz W (2000) Multiple reward signals in the brain. Nat Rev Neurosci 1:199–207.

Senn W (2002) Beyond spike-timing: the role of nonlinear plasticity and unreliable synapses. Biol Cybern 87:344–355.

Senn W, Markram H, Tsodyks M (2001) An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. Neural Comput 13:35–68.

Sjöström P, Turrigiano G, Nelson S (2001) Rate, timing and cooperativity jointly determine cortical synaptic plasticity. Neuron 32:1149–1164.

Staddon J, Higa J, Chelaru I (1999) Time, trace, memory. J Exp Anal Behav 71:293–301.

Watanabe M (1996) Reward expectancy in primate prefrontal neurons. Nature 382:629–632.

Yakovlev V, Fusi S, Berman E, Zohary E (1998) Inter-trial neuronal activity in inferotemporal cortex: a putative vehicle to generate long-term associations. Nat Neurosci 1:310–317.