EJN fens European Journal of Neuroscience

European Journal of Neuroscience, Vol. 34, pp. 766–779, 2011 doi:10.1111/j.1460-9568.2011.07805.x

NEUROSYSTEMS

Visual adaptation and novelty responses in the superior colliculus

Susan E. Boehnke,[1,*] David J. Berg,[2,*] Robert A. Marino,[1] Pierre F. Baldi,[3] Laurent Itti[2,4] and Douglas P. Munoz[1,5,6,7]
[1]Centre for Neuroscience Studies, Queen's University, Kingston, Ontario, Canada
[2]Neuroscience Graduate Program, University of Southern California, Los Angeles, CA, USA
[3]School of Information and Computer Sciences, University of California at Irvine, CA, USA
[4]Computer Science Department, University of Southern California, Los Angeles, CA, USA
[5]Department of Biomedical and Molecular Sciences, Queen's University, Kingston, ON, Canada
[6]Department of Psychology, Queen's University, Kingston, ON, Canada
[7]Department of Medicine, Queen's University, Kingston, ON

Keywords: Bayesian, dishabituation, habituation, repetition suppression, rhesus macaque

Abstract

The brain's ability to ignore repeating, often redundant, information while enhancing novel information processing is paramount to survival. When stimuli are repeatedly presented, the response of visually sensitive neurons decreases in magnitude, that is, neurons adapt or habituate, although the mechanism is not yet known. We monitored the activity of visual neurons in the superior colliculus (SC) of rhesus monkeys who actively fixated while repeated visual events were presented. We dissociated adaptation from habituation as mechanisms of the response decrement by using a Bayesian model of adaptation, and by employing a paradigm including rare trials that included an oddball stimulus that was either brighter or dimmer. If the mechanism is adaptation, response recovery should be seen only for the brighter stimulus; if the mechanism is habituation, response recovery ('dishabituation') should be seen for both the brighter and dimmer stimuli. We observed a reduction in the magnitude of the initial transient response and an increase in response onset latency with stimulus repetition for all visually responsive neurons in the SC. Response decrement was successfully captured by the adaptation model, which also predicted the effects of presentation rate and rare luminance changes. However, in a subset of neurons with sustained activity in response to visual stimuli, a novelty signal akin to dishabituation was observed late in the visual response profile for both brighter and dimmer stimuli, and was not captured by the model. This suggests that SC neurons integrate both rapidly discounted information about repeating stimuli and novelty information about oddball events, to support efficient selection in a cluttered dynamic world.

Introduction

Efficient selection of important events among temporal clutter requires the ignoring of repeating stimuli, thereby emphasizing novel and potentially important ones. This simple form of non-associative learning has been referred to as adaptation, habituation, and repetition suppression, depending on the era and field of study (Grill-Spector *et al.*, 2006; Krekelberg *et al.*, 2006; Clifford *et al.*, 2007; Kohn, 2007). From an information processing perspective, adaptation serves to adjust the operating point of a sensory system, to maximize the efficiency of sensory coding and increase differential sensitivity to novel events (Muller *et al.*, 1999; Dragoi, 2002; Dragoi *et al.*, 2002; David *et al.*, 2004; Dean *et al.*, 2005). This can be achieved through incremental updating over time of a Bayesian prior, which can then bias the processing of incoming sensory data (Itti & Baldi, 2005; Stocker & Simoncelli, 2006).

Correspondence: Susan E. Boehnke, as above.
E-mail: susan.boehnke@queensu.ca

*S.E.B. and D.J.B. contributed equally to this work.

Received 6 June 2011, accepted 16 June 2011

Electrophysiological evidence of response reduction with stimulus repetition has been observed throughout the visual system, from the retina (Smirnakis *et al.*, 1997; Brown & Masland, 2001; Hosoya *et al.*, 2005) and thalamus (Solomon *et al.*, 2004), to the visual cortex (Maffei *et al.*, 1973; Movshon & Lennie, 1979; Muller *et al.*, 1999; Motter, 2006) and frontal eye fields (Mayo & Sommer, 2008). These studies usually focus on perception; however, stimulus repetition effects also have profound, although less studied, consequences for visual orientation: the latency and magnitude of the visual response influences the timing of eye and head movements to foveate the stimulus (Dorris *et al.*, 2002; Fecteau *et al.*, 2004; Corneil *et al.*, 2008).

The ideal place to study visual repetition effects is the superior colliculus (SC) – the phylogenetically conserved hub of the visual orienting system (Ingle, 1975; Huerta & Harting, 1983; Dean *et al.*, 1989; Munoz *et al.*, 2000; May, 2006) – which is integrated with all other visual areas in the brain. Many visually responsive neurons in the SC also have movement responses that are time-locked to saccades (Mohler & Wurtz, 1976). The superficial layers of the SC receive visual input directly from the retina and from the early visual cortex,

whereas the intermediate layers receive more complex visual and cognitive input from various cortical areas, the basal ganglia and cerebellum (May, 2006). Therefore, early (e.g. retinal) and late (e.g. cortical) sources of visual adaptation can be compared directly by examining repetition effects across all SC visually responsive neurons located in different layers.

We explored how the magnitude and onset latency of SC visual responses changed with repetition. We modeled these changes by using a Bayesian approach to provide a quantitative definition of adaptation, which was then used to predict the consequences of changes to stimulus timing and intensity. To dissociate simple adaptation from higher-level learning processes (e.g. habituation), we compared responses to rarely presented brighter or dimmer oddball stimuli. Response decrement resulting from adaptation should follow the adaptation model and recover with a brighter but not a dimmer stimulus; however, response decrement resulting from habituation should recover (dishabituation) after any novel stimulus change (Sokolov, 1963; Baars, 1988).

Materials and methods

All procedures were approved by the Queen's University Animal Care Committee and were in full compliance with the Canadian Council on Animal Care guidelines on the care and use of laboratory animals. Experiments were performed with two male rhesus monkeys (*Macaca mulatta*) weighing between 9 and 12 kg. The surgical techniques used to prepare the animals for behavioral and physiological recordings have been described previously in detail (Marino *et al.*, 2008). Briefly, monkeys were implanted with a head post for head fixation, a recording chamber over the SC, and eye coils to measure eye position with the search coil technique. On the evening prior to surgery, the animal was placed under nil per os (water *ad libitum*), and prophylactic treatment with antibiotics was initiated [5.0 mg/kg enrofloxacin (Baytril)]. On the day of the surgery, anesthesia was induced with intramuscular ketamine (6.7 mg/kg). A catheter was placed intravenously to deliver fluids (lactated Ringer) at a rate of 10 mL/kg/h to a maximum of 60 mL/kg throughout the duration of the surgical procedure. Intramuscular glycopyrolate (0.013 mg/kg) was administered to control salivation and bronchial secretions, and to optimize heart rate. An initial dose was delivered at the start of surgery, followed by a second dose 4 h into the surgery. General anesthesia was maintained with gaseous isofluorene (2–2.5%) after an endotracheal tube had been inserted (under sedation induced by an intravenous bolus of propofol, 2.5 mg/kg). Heart rate, pulse, pulse oximetry saturation (S_pO_2), respiration rate, fluid levels, circulation and temperature were monitored throughout the surgical procedure. The analgesic buprenorphine (0.01–0.02 mg/kg intramuscular) was administered throughout the surgery and during recovery (8–12 h). The anti-inflammatory agent ketoprofen (2.0 mg/kg first dose, 1.0 mg/kg additional doses) was administered at the end of surgery (prior to arousal), on the day after the surgery, and every day thereafter (as required). Monkeys were given 2 weeks to recover prior to onset of behavioral training.

Monkeys were trained to perform a variety of oculomotor tasks for liquid reward. Real-time control of the experimental task and visual display was achieved with REX version 6.0 (National Eye Institute, NIH, Bethesda, MD, USA). Monkeys were seated in a primate chair 60 cm away from a CRT monitor (Mitsubishi XC2935C; 75-Hz refresh rate, 71.5 × 53.5 cm; usable field of view of 62° × 48°). Visual stimuli were presented within a darkened environment. Dark adaptation was prevented by dimly illuminating the monitor screen for 800 ms during the inter-trial interval. Physiological activity was monitored from 109 single neurons, using tungsten electrodes (Frederick Haer; 0.5–5.0 mΩ), with stimulus events and spike times being collected, and waveforms digitized, through the Plexon MAP system (Plexon Inc., Dallas, TX, USA). Further analysis was performed offline with custom MATLAB-based software (Mathworks, Natick, MA, USA).

Cell classification

When a neuron was first isolated, its visual receptive field was established by use of a simple fixation task in which white light stimuli (42.5 cd/m^2, 100 ms in duration, 0.25° diameter spot) were presented in pseudorandom order to 182 possible locations distributed across 60 (horizontal) × 50 (vertical) degrees of visual angle, the order of which was designed so that no two subsequent stimuli appeared within the typical response field of an SC neuron, in order to prevent adaptation effects. The centroid of the receptive field was then determined by use of a cubic spline function, and this location was used for all subsequent studies on the neuron. Because we were interested in adaptation of visual responses, we limited data collection to encountered cells that had a visual response.

We then collected further information to characterize the neuron relative to known SC cell types. First, we made careful measurements of microelectrode depth referred to the dorsal surface of the SC (as determined by the electrode depth that first elicited multiunit visual-only activity). Second, neural recordings taken during four interleaved saccade tasks [step, gap, memory-delay, and visual-delay; described in detail elsewhere (Munoz & Wurtz, 1993)] were used to classify visual and motor responses; critically, the visual-delay task dissociated visual and motor activity (Fig. 1A). In this task, the animal starts each trial by fixating a central fixation point. A target was then presented randomly in the center of the response field or at a location opposite to the vertical and horizontal meridian. To receive a reward, the animal had to maintain fixation until the fixation point disappeared (the delay period: 500–800 ms randomized) and then make a saccade to the target.

We refer to visual responses as 'transient' (a short visual burst) or 'sustained' (a visual burst followed by an extended period of low-frequency activity), as described previously (White *et al.*, 2009). This terminology is in line with descriptions of visual neurons in the geniculostriate pathway. We classified neurons as visual-transient (VT), visual-sustained (VS), visuomotor-transient (VmT), and visuomotor-sustained (VmS) (see Fig. 1B for single-unit examples), using two indices: a visual–motor index and a transient–sustained index. The visual–motor index was constructed with information from the saccade-aligned spike density function (Gaussian, σ = 5 ms) from the visual-delay task (Fig. 1A). The spike density function was first low-pass filtered by iterative convolution with a five-tap binomial kernel: 100 iterations in the forward direction and 100 in the backward direction. The result had no phase shift and approximated convolution with a Gaussian of σ = 14.12 ms. The timing and magnitude of the peaks and troughs of the waveform were then estimated by finding the zero crossings of the numerical gradient. A strong peak in activity from 25-ms pre-saccadic to 5-ms post-saccadic initiation was taken as evidence for a visuomotor neuron, and we quantified this feature with a simple probability measure:

$$P = \frac{1}{2}\left[1 + erf\left(\frac{\theta - T}{\delta\sqrt{2}}\right)\right] \quad (1)$$

where T and δ were the mean and standard deviation of non-perisaccadic peaks, θ was the value of the perisaccadic peak, and $erf(x)$ was the Gauss error function. Large probabilities indicated

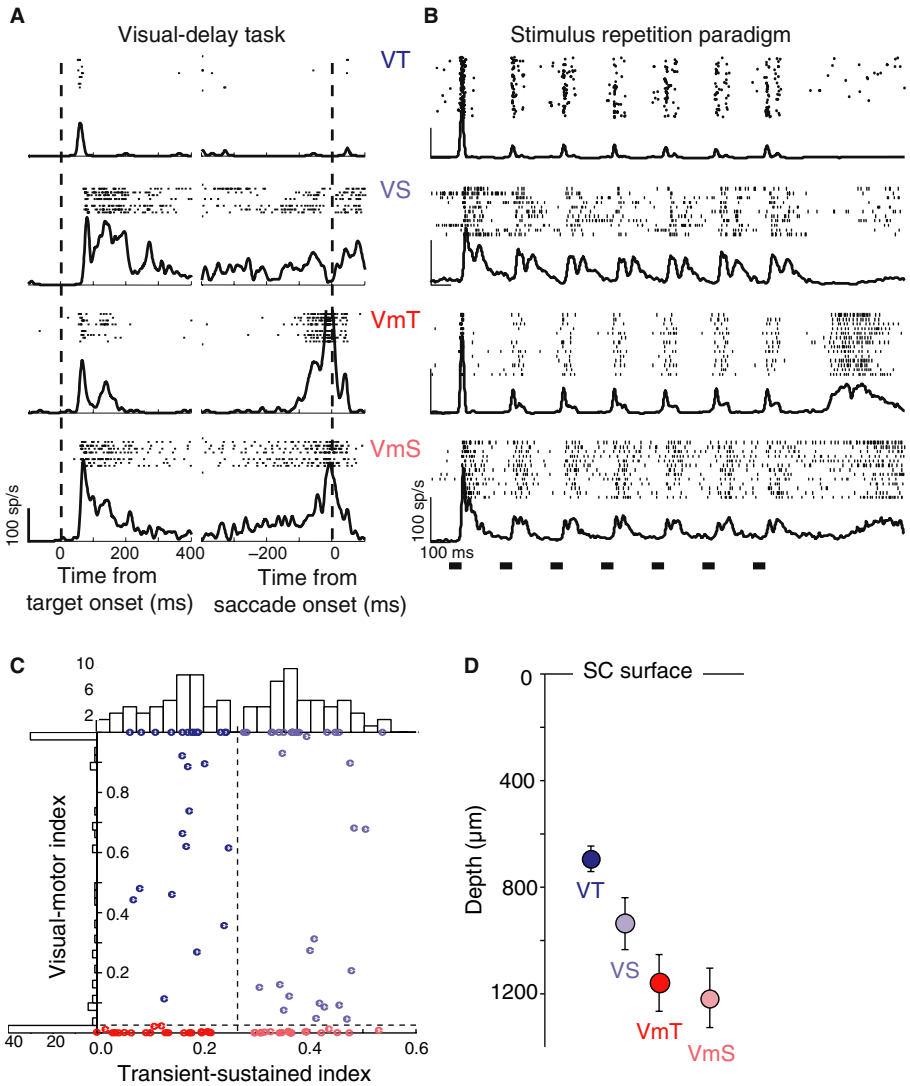

FIG. 1. (A) Raster plots and spike density waveforms (σ = 5 ms) recorded from representative VT, VS, VmT and VmS neurons for a delayed saccade task, which was used to facilitate neural classification. Data are aligned on target appearance (left column) and saccade onset (right column) in the delayed saccade task when the target appeared in the neuron's response field. (B) The response of the same single neurons to the seven stimuli in the standard repetition paradigm. The black bars across the bottom of the abscissa represent the stimulus timing. Spikes for individual trials are presented in raster format (only a subset of trials are shown for display purposes) and overlaid with a mean spike density function (σ = 5 ms). (C) Scatter plots and histograms of the metrics used to classify cells. The transient–sustained index is plotted against the visual–motor index for each cell (color indicates cell class), with smaller numbers indicating more motor and more transient responses, respectively, as measured from responses in the visual delay task shown above. The histograms show the number of cells with each parameter value, using a bin width of 0.025 units. The dashed lines show the cutoff values that were used to separate classes of neurons into the four categories. (D) The mean depth for each cell class. The cell classes had significantly unequal variances (Bartlett's test, T_3 = 15.50, P = 0.0014), and a Kruskal–Wallis test was therefore conducted to evaluate differences in depth among the four cell classes. Cell classes significantly differed in depth ($C^2_{3,108}$ = 20.10, P = 0.0002), and pairwise Wilcoxon rank-sum tests (Bonferroni-corrected) indicated that VT cells differed significantly from both VmT and VmS cells (z = 3.78, $P < 0.001$, and z = 3.79, $P < 0.001$, respectively).

the presence of motor activity, and if no peak was found, a probability of 0.0 was assigned. To confirm that the peak activity was related to a robust motor response and not to residual sustained visual activity or noise, the smallest trough was measured in a small window (±25 ms) around saccade onset. We computed the probability that activity at the trough (or at saccade initiation if no trough existed) was higher than the average pre-saccadic baseline activity (−900 ms to −50 ms pre-saccade) using Eqn 1, but where T and δ were the mean and standard deviation of the baseline, and θ was the value at the trough. Finally, the visual–motor index was computed as $1 - P_pP_t$, where P_p was the probability from the peak measurement, and P_t was the probability from the trough measurement. We considered cells with a visual–motor index < 0.025 to be visuomotor cells.

To compute the transient–sustained index, we aligned spike density functions to target appearance in the visual-delay task (Fig. 1A) and divided the post-stimulus visual response into early (transient) and later (sustained) components. Each time point in the first 400 ms of post-stimulus activity was compared with a baseline (700 ms pre-target) using Eqn 1, where T and δ were the mean and standard deviation of the baseline, and θ was the value of the time point. Intervals of post-stimulus activity where each point in the interval had a high probability ($P \geq 0.99$) were identified, and if the first region had raised activity for > 10 ms, its start and end (maximum of 400 ms) identified the early component; otherwise, the whole post-stimulus interval from 0 to 400 ms was taken as the early component. The later component was then identified as the remaining interval until 550 ms after stimulus appearance (minimal delay interval). The VT index was

then calculated as $S/(T + S)$, where S was the mean activity in the later (sustained) component, and T was the mean activity in the early (transient) component. The distribution of index values for each metric is shown in Fig. 1C. To divide the cells into transient and sustained classes, we chose a value of 0.2625, which was a natural division in the distribution of index values.

SC neurons have well-characterized responses ranging from purely visual to purely motor (Mohler & Wurtz, 1976; Mays & Sparks, 1980; Munoz & Wurtz, 1995; McPeek & Keller, 2002). Visual-only cells with transient visual responses and no saccade-related activity (VT) tended to be located more superficially than the other classes of visually responsive cells (Fig. 1D). Thus VT cells were typically found in the upper superficial gray layer (e.g. superficial layers of the SC), where retinal Y-type cells terminate directly and indirectly through the magnocellular lateral geniculate nucleus and V1 (May, 2006). Visual-only cells that had sustained visual responses (VS) typically paused during saccadic eye movements (Fig. 1A). Previously, we showed that VS neurons were sensitive to color signals, whereas VT cells were not (White *et al.*, 2009), suggesting parvocellular input. These features, along with a mean depth of about 900 μm (Fig. 1D), suggest that they were located in the lower superficial layers. This area receives visual input from higher occipital and parietal areas (Graham *et al.*, 1979; Tigges & Tigges, 1981). The VS neurons that we identified were probably the same as the 'visual-tonic' neurons described previously (McPeek & Keller, 2002; Li & Basso, 2008). Finally, visuomotor cells with transient or sustained activity – VmT or VmS – were easily characterized, because of their bursts of activity being time-locked to saccades and their bursts of visual activity being time-locked to stimulus onset (Fig. 1A). Our sample of visuomotor neurons was, by necessity, biased towards those with robust visual responses, and were always found more than 1 mm below the dorsal surface (Fig. 1D).

Behavioral task

Monkeys actively fixated a central fixation spot (grayscale circle of 0.25° diameter presented at 1.1 cd/m^2) while a series of seven light flashes (i.e. stimuli, 0.25° in diameter, 55 ms in duration; Figs 1B and 2A) were presented in the receptive field of the monitored neuron. In the main paradigm, these seven stimuli were separated by intervals 200 ms in duration [i.e. 255-ms interstimulus interval (ISI)]. Monkeys received a small liquid reward for maintaining fixation within a small computer-controlled window (1–3° square window) for the duration of each trial. If fixation was broken prior to the end of the trial, the trial was aborted, eliminated from further analysis, and recycled back into the trial sequence. In the main paradigm, 70% of the trials (control trials) consisted of seven equiluminant stimuli (1.1 cd/m^2), and 30% of the trials (oddball trials) were identical except that the fourth stimulus could be brighter (10%, 5 cd/m^2), dimmer (10%, 0.1 cd/m^2), or absent (10% of trials). These trial types were randomly interleaved. Trains of seven stimuli were chosen because they allowed for examination of responses before and after presentation of the oddball, and they constituted a comfortable trial duration for the monkey to maintain steady fixation. The ISI of 255 ms was chosen because maximal inhibition of return was observed in monkeys at a cue–target onset asynchrony of that interval (Fecteau *et al.*, 2004).

For a subset of neurons (19 tested and 17 analyzed, two removed for having no response to some stimuli), the ISI was varied systematically between 155, 255 and 455 ms in the control condition only, to investigate the effects of ISI on the repetition effect. Typically, the 255-ms ISI block was collected first, because that was part of the main paradigm used with the oddball trials. If neuronal isolation remained strong after that paradigm, additional files were obtained at other ISIs collected in pseudorandom order.

Neural analysis

Single neurons or pairs of single neurons were recorded from a single electrode, isolated online with the window discriminator in PLEXON, and verified and optimized offline with the PLEXON OFFLINE SORTER (Plexon Inc., Dallas, TX, USA). The timing of events in the trial sequence was then calculated automatically with custom MATLAB (MATLAB 6.5; Mathworks, Natick, MA, USA) software during offline analysis. We recorded from a total of 109 neurons in the control task with oddball trials. Of these, recordings from 98 neurons (60 from monkey Q; 38 from monkey Y) had sufficiently good spike isolation throughout recording, a mean visual response > 40 spikes/s, responses to all seven stimuli during the control trials, and at least six trials of each oddball condition. Typically, there were 10–20 repetitions of each oddball condition and 70–140 repetitions for the control condition. Two spike density functions were created for each trial of each condition by convolving the trains of action potentials with a Gaussian kernal ($\sigma = 5$) or by convolving with a combination of growth and decay functions that resemble a postsynaptic potential given by:

$$R(t) = (1 - e^{-t/\tau_g})(e^{-t/\tau_d}) \quad (2)$$

where R is the firing rate as a function of time t, τ_g is a time constant for the growth phase, and τ_d is a time constant for the decay phase. Time constants of 1 and 20 for the growth and decay phases, respectively, were chosen, following the practice of others (Thompson *et al.*, 1996). Spike density functions were aligned on the first stimulus onset, and activity from repeated trials were averaged to generate a mean spike density function (for both functions separately) for each neuron for each condition. The magnitude of the first peak response to each visual stimulus was calculated off the Gaussian spike density function, and the onset of that response was calculated from the spike density function created by rate function R, which provides a more accurate measurement of onset time. This was performed for the responses to each of the seven stimuli in every condition, for each neuron. To find the first peak in the visual response and its onset, a custom computer algorithm written in MATLAB looked for all peaks in the spike density function in the epoch from 50 ms after stimulus onset until the end of the ISI. It then marked the highest peak (usually the first one) on a visual display. Each first-peak calculation was manually examined and changed if it was incorrect (e.g. if a late noisy peak was incorrectly chosen as the first main peak by the algorithm). Once the peak had been determined, the onset latency of that visual response was calculated by an algorithm that looked backwards in time (maximum of 40 ms back) along the descending slope from the first peak, in order to find the point at which the response became significantly greater than the mean neural activity in an epoch spanning 25 ms before to 25 ms after the onset of the visual stimulus that generated that response. Again, each of these was examined on the visual display, and manually adjusted if necessary (e.g. if unusual noise levels or sustained activity from the preceding response unduly lengthened the response onset latency (ROL) calculation of the algorithm).

Receiver operating characteristic (ROC) analysis

The ROC was used to quantify the time course of differences between control and oddball conditions after the presentation of the fourth stimulus in the sequence. For each cell, we computed the area under the

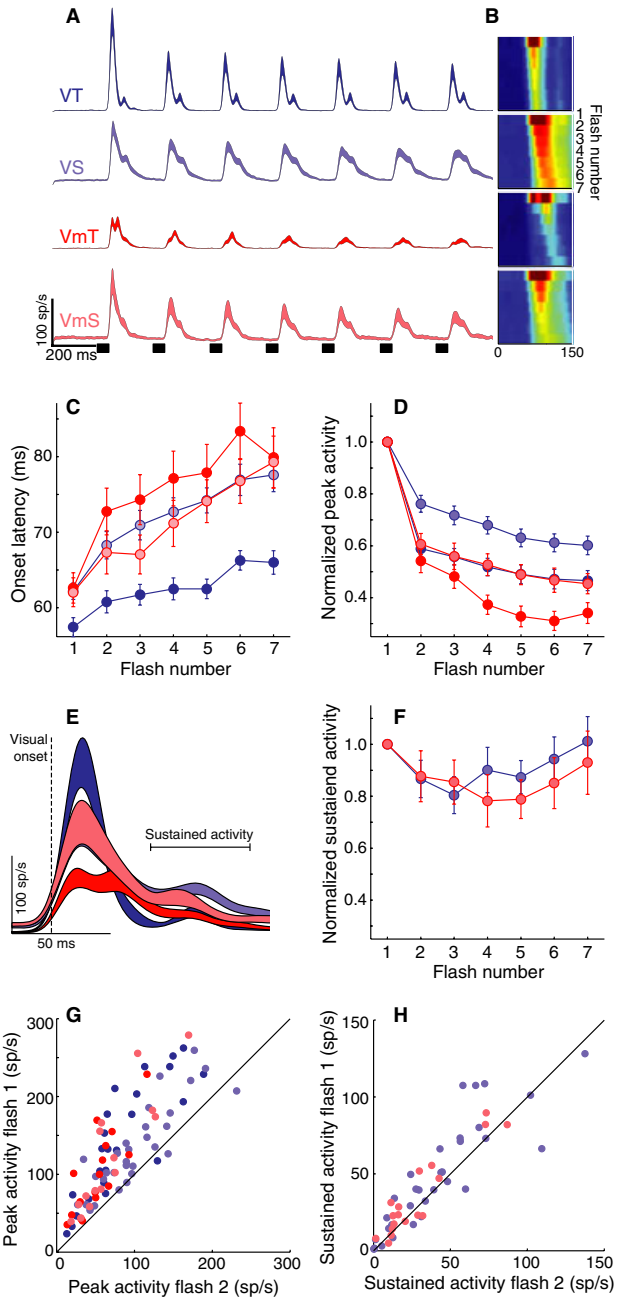

FIG. 2. (A) Spike density functions for VT (n = 32), VS (n = 32), VmT (n = 16) and VmS (n = 18) neurons in response to seven repeated stimuli (shown as small dark bars at the bottom of the trace) in the center of each neuron's response field. (B) Color coding of intensity of neural activity in response to the seven stimuli (time of the response to a given stimulus on the horizontal axis, response to each stimulus descending vertically, color-coded for normalized spike rate). Note the shift in onset latency with each stimulus repetition. (C) Changes in mean response onset latency across stimulus number for each neural type. (D) Changes in peak response magnitude across stimulus number for each neural type, normalized to the response on the first stimulus. (E) Population spike density waveforms in response to the first target stimulus, aligned on response onset to show the early (transient) and later (sustained) components of the visual response. (F) Normalized mean sustained activity (50–100 ms after onset of visual response) is plotted for the seven stimuli for the VS and VmS neuron populations. (G and H) Scatter plots showing the relationship between the responses to the first and second stimulus for the transient peak (G) and sustained portion (H) of the neural response. Standardized major axis regression analysis revealed that this relationship had a slope greater than unity for the peak activity (F-test; $F_{1,96}$ = 72.32, P < 0.01), but not for the sustained activity (F-test; $F_{1,48}$ = 0.99, P = 0.32).

ROC curve between the control condition and each oddball condition in a 50-ms sliding window centered on the time point of interest (gray box, bottom left of Fig. 6A). The window's left edge started at the onset of the visual response and advanced 1 ms in time (depicted by the solid line) until the right edge reached the onset of the next (fifth) stimulus. This resulted in one ROC area measurement for each cell and each time point on the interval shown (which started at half the window size, 25 ms after the onset of the visual response). To control for variation in timing of each cell's response to the fourth stimulus, it was necessary to first align each spike density function to the onset of visual activity in response to the fourth stimulus. As a result of this realignment, the waveforms from different neurons and conditions had slightly different lengths, as the time of the visual onset varied with cell type and condition, and consequently, the ROC analysis was performed over a different length of time for each cell and condition. As a result, fewer cells entered the ROC area calculation near the end of the analysis interval. The length of the analysis interval shown in Fig. 6 (below the VT spike density function) was the maximum interval that could be chosen that still contained more than 50% of the cells for all classes and conditions (all but VmT cells had more than 80% at the end of this interval). All cells in all conditions had at least 122 ms from the onset of the visual response to the onset of the next stimulus, and the median was 152 ms.

Implementation of the computational model

The Bayesian model of adaptation is summarized in Fig. 3A. The model is based on surprise theory using a Poisson-gamma model, which is described in detail elsewhere (Itti & Baldi, 2005; Baldi & Itti, 2010), but is summarized here with the differences in our implementation being noted. The model consists of two stages of Bayesian learning, which are identical except for their input sources (Fig. 3A), so for clarity the equation subscripts are omitted from the following discussion. We consider that each Bayesian learner receives one-dimensional Poisson-distributed spike trains (from the retina and visual cortex or from the previous stage of learning), represented internally as a family of models, $M(\lambda)$, which are all the possible one-dimensional Poisson distributions of firing rates ($\lambda > 0$). Each learner builds probability distributions (hypotheses or beliefs) $P(M(\lambda))$, concerning which of these models best represents the current state of the stimulus. As is typical in iterative Bayesian learning, the prior and posterior distributions are chosen from the same functional form (conjugate priors), so that the posterior distribution at one time step is used as the prior distribution for the next. When the data are Poisson-distributed ($D = \lambda$), the gamma probability density function is the conjugate prior, $P(M(\lambda))$:

$$P(M(\lambda)) = \gamma(\lambda; \alpha, \beta) = \frac{\beta^{\alpha}\lambda^{\alpha-1}e^{-\beta\alpha}}{\Gamma(\alpha)} \quad (3)$$

with shape $\alpha > 0$ and inverse scale $\beta > 0$, and where $\Gamma(\alpha)$ is the Euler gamma function of α. Given an input sample $D = \lambda$, the posterior distribution of beliefs over the possible input firing rates is also a gamma distribution characterized by:

$$\alpha' = \zeta\alpha + \lambda + \varepsilon \quad \text{and} \quad \beta' = \zeta\beta + 1 \quad (4)$$

where α' and β' are the shape and inverse scale of the posterior distribution, ζ is a temporal parameter (forgetting factor, $0 < \zeta < 1$), which determines the rate of learning, and ε is a constant representing

noise. The second stage of Bayesian learning takes as input the expected value of the first stage's posterior distribution:

$$E[P(M(\lambda)|D)] = E[\gamma(\lambda;\alpha,\beta)] = \frac{\alpha}{\beta} \quad (5)$$

The output of the system is then calculated from the final Bayesian learner as the Kullback–Leibler (KL) divergence (Kullback, 1959) between prior and posterior distributions over all possible firing rates, which summarizes the amount of learning or adaptation that just resulted from observing the data D:

$$\begin{aligned} &KL(P(M(\lambda)), P(M(\lambda)|D)) = KL(\gamma(\lambda;\alpha,\beta), \gamma(\lambda;\alpha',\beta')) \\ &= \alpha' \log\frac{\beta}{\beta'} + \log\frac{\Gamma(\alpha')}{\Gamma(\alpha)} + \beta'\frac{\alpha}{\beta} + (\alpha-\alpha')\Psi(\alpha) \end{aligned} \quad (6)$$

where $\Psi(\alpha)$ is the digamma function of α. This differs from the Itti and Baldi implementation (Itti & Baldi, 2005; Baldi & Itti, 2010), where the KL divergence is computed at each learning stage, and the system output is the product of the outputs at each stage. Figure 3B shows the time dynamics of the system for each stage in response to a control trial.

The Itti and Baldi implementation uses five learning stages, each having the same temporal parameter. In this experiment, we found that only two stages, but with the temporal parameter of each stage allowed to be different, adequately predicted the peak firing rates of the neurons. Additionally, in their implementation of Eqn 4, the temporal parameter is applied to the prior distribution's α and β parameters before computation of the Bayesian update. As a result, there is always a baseline output. We computed the update so that, if the posterior and prior distributions are the same, the output of the system is 0.

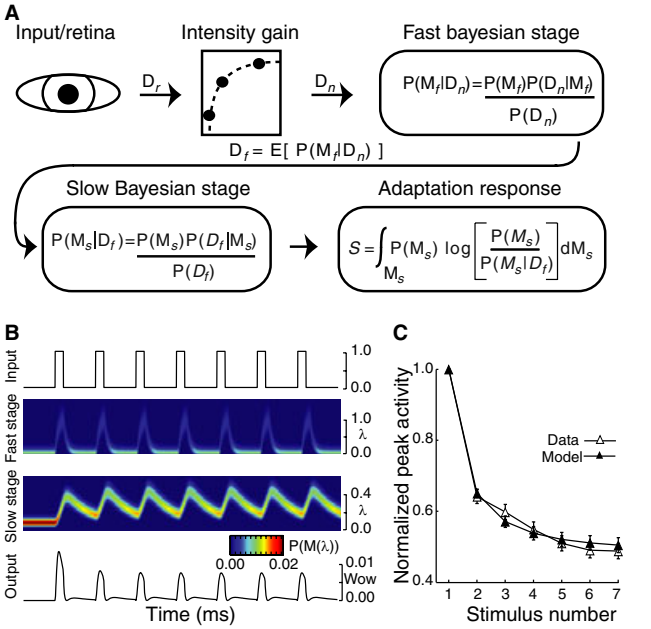

FIG. 3. A. Schematic of the Bayesian adaptation model. Light stimulating the retina was modeled as a square wave of unity amplitude (D_r; 1.1 and 0.9 for the brighter oddball conditions) and passed through a static gain function that was constant for all model neurons (see Materials and methods). Two stages of Bayesian learning supply the adaptation dynamics. In each stage (subscripts omitted), the learning process builds hypotheses or beliefs (probability distribution) over a class of internal models M that represent all possible values of its input. As new sensory data D_r are collected, the Bayes theorem provides the mechanics to turn a prior set of hypotheses $P(M)$ about which model best characterizes the input data into a posterior set of hypotheses $P(M|D)$, given the likelihood of the data $P(D|M)$ under the assumptions of model M. The fast Bayesian stage quickly adapts to the input, and passes the expectation of its posterior beliefs D_f as input to the second Bayesian stage. A posterior set of beliefs is computed in the same fashion as for the fast learner, but with a slower learning dynamic. The adaptation response is then calculated for every data observation as the KL divergence (Kullback, 1959) between the slow learner's prior and posterior hypotheses, signaling the amount of shift in the model's beliefs caused by each new observation. (B) Detailed view of the model dynamics across each stage during a control trial (see Materials and methods for a detailed description of the model). The top trace represents the input stimulus from control trials. The two central images show, for each Bayesian learner, the distribution of beliefs about which of the possible Poisson firing rates (y-axis) best characterizes the input over the course of a single trial (x-axis). Hotter colors indicate that, at a given point in time, there is a higher belief (probability) in a particular firing rate. The bottom panel shows the final output of the system. (C) Population mean and standard error of the model (filled symbols) and neural (open symbols) normalized peak responses to the seven stimuli in the control condition.

Model fitting

Model parameters were estimated for each neuron individually by fitting the peaks in the model's output to the seven peak magnitudes in each neuron's response profile from the 255-ms ISI condition. Each model neuron consisted of three parameters: two time constants (ζ, Eqn 4) that controlled the speed of learning in the two Bayesian learners, and the baseline noise parameter (ε, Eqn 4), which was globally set for all model neurons. The data were fitted initially with all three parameters free for each cell. After fitting, a probability density function of the baseline parameter was estimated by kernel density estimation with automatic bandwidth selection (Sheather & Jones, 1991), which was implemented in the R statistical package (www.r-project.org). A quadratic function was fitted (least squares method) to three points centered on the maximum of this curve, and the analytic maximum of the quadratic function was used as an estimate of the most likely value of the baseline parameter. Fitting was then performed again with the baseline parameter fixed for all cells to the most likely value, reducing the model to the two temporal parameters. The best parameters were determined by using the Nelder–Mead simplex method (Lagarias *et al.*, 1998) built into MATLAB to minimize the error of the following process. First the input signal (seven stimuli) was simulated as a square wave with unit amplitude, and the adaptation model's response was computed for a given parameter set. Parameters were encoded such that the second stage's learning rate was guaranteed to be slower than the first stage's. Model and cell responses were normalized by the response to the first stimulus. Normalization eliminated the need for scaling parameters in the model without affecting the morphology of the adaptation. The error was then computed as the median absolute difference between the model's peak response to each stimulus and the cell's peak responses (disregarding the first stimulus, which always had zero error). Several alternative error functions were explored, and median absolute error was chosen because it gave highly significant, and qualitatively the best, overall fits and predictions. Several other error functions also gave significant results. Additionally, a single-parameter model significantly fitted the data; however, the two-parameter model produced less total error for all conditions combined, and less median error for all but the 155-ms ISI condition. Qualitatively, the mean population responses of the neural data were in agreement with the mean population responses of the two-parameter model.

To account for the relationship between stimulus brightness and a cell's peak firing rate, the adaptation model used a gain function

(Fig. 3A). Because only three intensity levels were considered, this amounted to finding two gain factors to represent the 10% brighter and 10% dimmer stimuli. After finding of the best temporal parameters during control trials, a single set of gain parameters was chosen that minimized the error between all model neurons and all real neurons simultaneously, considering only the brighter and dimmer conditions. For this dataset, the gain factors were ~1.1 and ~0.9, respectively. Because the gain factors were very close to the 10% brighter and 10% dimmer inputs, the gain function could have been omitted with little loss of model fit quality.

Model evaluation

To evaluate the model fits without assumptions about the distribution of data or errors, several statistics were computed in the permutation (randomization) framework. Goodness-of-fit was assessed by using the median of the absolute error between all neuron and model responses, for all conditions, as a test statistic in a repeated measures permutation design (stimuli 2–7 for each condition). To assess whether cell and model responses came from the same underlying distribution, the permutation equivalent of a two-factor repeated measures ANOVA was performed. The final test (paired-error or reliability test) indicated whether, overall, the model was able to predict neuronal responses better than other cells from the same class (which might be thought of as an upper bound). First, for each condition separately, all pairwise combinations of neurons (restricted to within neuron class) were evaluated with the error function. This distribution of values represents the errors that occurred when each cell was used to predict other cell responses, and served as a summary of the variability (reliability) of the repetition effect within a class of neurons. Higher values indicated that neurons responded very differently from one another. By use of the permutation equivalent of a two-factor repeated measures ANOVA, this distribution was compared with the distribution of errors generated by model predictions. Figures 4D and 5C show the distribution of model errors for the ISI, and oddball manipulations, respectively. This test compared directly the quality of our model fits to the variability of the repetition effect. We reason that a well-performing model should be, on average, as, or more, consistent with the neurons' response than neurons of the same class are with each other. All permutation tests were carried out with the Monte Carlo method, with 30 000 iterations.

Results

Effects of stimulus repetition

Figure 2 illustrates the main effect of this study – the large response decrement that occurred with repeated stimulation (seven stimuli) of the receptive field of visually responsive neurons in the SC. Response decrement was observed for all four types of visual neuron classified: VT (n = 32), VS (n = 32), VmT (n = 16), and VmS (n = 18) (see Fig. 1B for examples of individual neuron responses). Following the appearance of the first stimulus, neurons of each cell type discharged a robust phasic response (Fig. 2A and E). The early transient part of this response was dramatically affected by repeated stimulation: the peak response magnitude decreased (Fig. 2A, B and D) and ROL increased (Fig. 2B and C). A mixed analysis of variance (ANOVA) with a between-subjects factor (four cell classes) and a within-subjects factor (seven stimuli) was conducted, and revealed significant main effects of cell class (peak, $F_{3,94}$ = 9.15, P < 0.01; ROL, $F_{3,94}$ = 8.8, P < 0.01) and stimulus repetition (peak, $F_{6,94}$ = 378.1, P < 0.01; ROL, $F_{6,94}$ = 89.9, P < 0.01), and an interaction (peak, $F_{18,564}$ = 5.8, P < 0.01; ROL, $F_{18,564}$ = 3.3, P < 0.01).

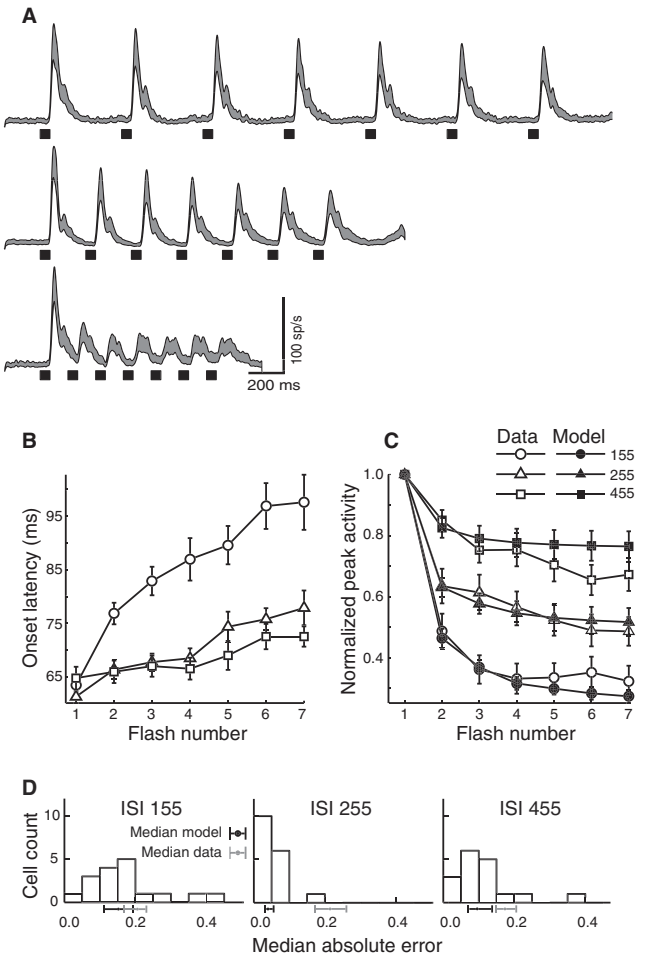

FIG. 4. Effect of changing ISI on the repetition effect. (A) Population spike density waveforms recorded from 17 visually responsive neurons in the SC in response to seven stimuli (55 ms) presented with ISIs of 155, 255 and 455 ms. As ISI increased, the repetition effect was reduced. At short ISIs, the response onset latency (B) was increased and the peak response magnitude (C) was decreased with stimulus repetition. Without changing the parameters used to generate the model fit to control trials (Fig. 3C), the model (closed symbols) predicted the neurons' peak response magnitude (open symbols) across the different rates of stimulus presentation (two-factor repeated measures permutation test, P = 0.89). (D) The paired-error test (see Materials and methods) indicated that, on average, each model was a better predictor of the peak activity of its corresponding real neuron than other neurons of the same class (P < 0.01). Histograms of median absolute errors between each neuron and its corresponding model for the three ISI conditions are shown. The black dot and line below each axis show the median error and 95% bootstrapped confidence interval (30 000 iterations) of model errors. The gray dot and line show the median error and 95% bootstrapped confidence interval from the distribution of pairwise errors between actual neurons (see Materials and methods). Note that, in all cases, the models' median error is less than the median error between neurons.

All cell types decreased their peak response magnitude with repetition (VT, $F_{6,186}$ = 143.06, P < 0.001; VS, $F_{6,186}$ = 71.69, P < 0.001; VmT, $F_{6,90}$ = 114, P < 0.001; VmS, $F_{6,102}$ = 71.5, P < 0.001), and the majority of the decrease occurred on the second stimulation. This was verified statistically: the ratio of peak magnitude between the first and second stimuli (mean = 0.36) was greater than the ratio of peak magnitude between the second and seventh stimuli (mean = 0.15) (t_{97} = 7.9, P < 0.001; paired t-test). This relationship

was confirmed for all cell types independently ($P = 0.04$ or less). ROL increased in a mostly linear fashion, with repetition for all cell types (VT, $F_{6,186} = 18.1$, $P < 0.001$; VS, $F_{6,186} = 42.5$, $P < 0.001$; VmT, $F_{6,90} = 15.9$, $P < 0.001$; VmS, $F_{6,102} = 17.5$, $P < 0.001$).

In summary, VT neurons, which are likely to be found in the most superficial retino-recipient SC layers (May, 2006), had strong adaptation (~50%) but the smallest ROL increase (~10 ms) of all cell types. VS neurons, which are most likely to be found in the lower superficial layers (Tigges & Tigges, 1981) showed the least adaptation (~35%) but a large increase in ROL with repetition (~15 ms). Finally, visuomotor neurons of the intermediate SC layers showed both strong adaptation (> 50%) and a large increase in ROL (> 15 ms), particularly the VmT cells. Indeed, some VmT cells (not described here) completely lost their visual response after only a few trials (Goldberg & Wurtz, 1972) and thus could not be studied in our paradigms.

To determine whether the later components of the visual response in cells with a significant sustained component (VS and VmS as defined by our cell classes) were also affected by repetition, we calculated the average activity from 50 to 100 ms after the response onset (Fig. 2E). The sustained activity was less affected by repetition (Fig. 2F and H) than the early transient component (Fig. 2D and G). An ANOVA on the sustained activity of VS and VmS neurons showed a far smaller main effect of repetition ($F_{6,48} = 6.75$, $P < 0.01$) than that seen in the transient component, and no main effect of cell class, or an interaction (F-values < 1).

Modeling the response decrement with a Bayesian framework

The effect of stimulus repetition on response magnitude was modeled by use of a simple Bayesian model of stimulus adaptation that monitored the temporal dynamics of streams of stimuli (see Fig 3A and B, and Materials and methods). The model relies on a recently developed Bayes-optimal theory of novelty, which has been shown to provide a quantitative account of adaptation in early visual neurons (Itti & Baldi, 2005; Baldi & Itti, 2010) This model provided a principled theoretical foundation for quantifying the effects of adaptation in terms of a hypothetical optimal Bayesian learner: stimuli that, over time, gave rise to no significant learning caused a rapid decrease in response (adaptation); in contrast, stimuli that caused a shift in the model's current estimates gave rise to significant learning and to vigorous model responses. Each neuron was modeled individually as three stages consisting of a static gain function and two Bayesian learners (Fig. 3A). Parameters were estimated by fitting each model neuron's peak responses to a real neuron's peak responses to all seven stimuli (see Materials and methods). The model was able to significantly fit the repetition effect on response magnitude (goodness-of-fit test, $P < 0.01$), and the population responses for model and real neurons overlapped (Fig. 3C).

The effect of stimulus presentation rate

We generated predictions for the repetition effect's dependence on the rate of stimulus presentation by altering the ISI of inputs to the population of model neurons. If the response decrement followed the adaptation model predictions, then decreasing the ISI to 100 ms would cause a stronger repetition effect, whereas increasing the ISI to 400 ms would allow recovery of the effect of previous stimulation. We tested these predictions in a subset of 17 neurons (three VT, seven VS, three VmT, and four VmS) by repeating control trials (seven identical stimuli of 55-ms duration) with these different ISIs (onset to onset times of 155, 255 and 455 ms). Figure 4A shows the combined spike density functions from this subpopulation. There was a clear effect of ISI on both peak response magnitude ($F_{2,32} = 29.14$, $P < 0.01$) and ROL ($F_{2,32} = 28.5$, $P < 0.01$). That is, the shorter the ISI, the more dramatic the repetition effect. This was confirmed by an interaction between ISI and repetition (peak, $F_{12,192} = 8.44$, $P < 0.01$; ROL, $F_{12,192} = 6.9$, $P < 0.01$). Reducing the ISI led to an increase in ROL (Fig. 4B), and a reduction in response magnitude (Fig. 4C). The main effect of repetition, as expected, was significant (peak, $F_{6,96} = 46.29$, $P < 0.01$; ROL, $F_{6,96} = 32.27$, $P < 0.01$). Remarkably, we found that our simple model was able to predict the pattern of response magnitudes observed in these other ISI conditions (Fig. 4C) without a change in parameters (goodness-of-fit test, $P < 0.01$) and was a better predictor of neural activity than other neurons (see Fig. 4D).

The effect of rare changes in stimulus luminance

We also modeled the effect of inserting rarely presented luminance oddball stimuli (brighter, dimmer, and absent) into the stimulus sequences. If the pattern of changes was caused purely by adaptation, the neural response to the oddball should follow the predictions of the adaptation model and recover somewhat with a brighter, but not a dimmer oddball stimulus. The response to the subsequent (5^{th}) normal luminance stimulus should show the opposite pattern – an increased response after a previous dimmer stimulus and a decreased response after a previous brighter stimulus. Alternatively, neurons could show response recovery with all of the rare stimuli, akin to 'dishabituation'. To test these predictions, we examined the visual response of SC neurons to sequences of seven stimuli where the fourth stimulus was of higher intensity (10% of trials), was of lower intensity (10% of trials), was absent (10% of trials), or had no change (70%). Figure 5A illustrates the population responses recorded from each type of neuron for the third, fourth and fifth stimuli. We found that the peak magnitude of the neural data conformed to the predictions of the adaptation model, using the same parameters as in the control condition (see Fig. 5B to contrast physiological data with model fits). The model fitted the data significantly (goodness-of-fit test, $P < 0.01$), and was a better predictor of neural activity than other neurons in the same class (Fig. 5C).

Figure 5D and E shows the normalized difference [(oddball − control)/(oddball + control)] between the oddball and control trials for ROL and the peak magnitude, respectively (all cell types were collapsed, because the changes in the early part of the visual response were qualitatively the same in all cell types). Significance was tested with a Bonferroni-corrected t-test (critical $t = 2.49$, $P < 0.05$). As expected, there was no difference between control and oddball trials on the third stimulus for any condition (all t-values < 2.49). Presentation of the brighter stimulus in the fourth position led to a larger-magnitude response ($t_{97} = 5.67$) at a shorter latency ($t_{97} = -6.36$), whereas presentation of the dimmer stimulus showed the opposite effect – a smaller peak response ($t_{97} = -3.7$) with a longer latency ($t_{97} = 8.79$). Furthermore, the changes in the latency and magnitude of the fourth response had predictable effects on the earliest part of the response to the fifth stimulus. If the fourth stimulus was brighter, the response to the fifth stimulus was reduced ($t_{97} = -6.22$) and arrived later in time ($t_{97} = 6.72$) than in the control condition. In contrast, when the fourth stimulus was dimmer, the response to the fifth stimulus was larger in magnitude ($t_{97} = 7.6$) but not significantly earlier in time ($t_{97} = -2.36$). In the absent fourth stimulus condition, the response to the fifth stimulus was much larger in magnitude ($t_{97} = 11.24$) and occurred earlier in time ($t_{97} = -6.02$). These

analyses indicate that the pattern of changes observed in the timing and magnitude of the early transient component of the response to oddball stimuli is consistent with predictions of Bayesian adaptation to stimulus intensity.

Sustained responses to novel events

The best-fitting adaptation models often did not produce any output during the inter-trial interval (sustained response), but a sustained response that showed some modulation with repeated stimuli was observed in sustained cell classes (Fig. 2E and F). To investigate whether the sustained activity could possibly reflect something other than simple adaptation, we analyzed the later part of the visual response to oddball trials (Fig. 6). First, we realigned the visual responses in control and oddball conditions to the onset of the response to the fourth stimulus (see traces on the left side of each panel in Fig. 6), correcting for the ROL difference between brighter and dimmer stimuli and between cells (Fig. 5D). We then performed an ROC analysis on the response, spanning from 25 to 120 ms after response onset (right side of each panel in Fig. 6), to determine when the responses to the control and oddball trials became significantly different (see Materials and methods). This analysis interval started earlier than that used in Fig. 2, to show the effect of the first visual volley in the ROC plots. For all cell types except VmT, immediately after the onset of the visual response there was a significant difference in the transient response in the brighter condition that became insignificant approximately 30 ms after ROL. That is, the peak transient activity faithfully reflected stimulus intensity – the brighter stimulus elicited the strongest response, and the dimmer stimulus the weakest. However, later in the response of the VS and VmS cells (bottom panels), there was a significant increase in activity for both the brighter and dimmer oddball conditions, possibly representing a dishabituation signal reflecting the novelty of the oddball stimuli. The time points after ROL when the brighter and dimmer stimulus responses diverged from control responses were 73 and 68 ms, respectively, for VS neurons, and 95 and 81 ms, respectively, for VmS neurons (Fig. 6; see vertical dotted lines and *P*-values plotted below the ROC area curves). To demonstrate how consistent this was across individual neurons, in Fig. 6E and F the mean sustained firing rate after the fourth stimulus for control trials is plotted against that for oddball trials for VS and VmS neurons, respectively. Points falling above the unity line show neurons whose rate was higher after the oddball stimulus than after the control stimulus (sustained epochs of 80–110 and 90–120 ms for VS and VmS neurons, respectively). In the inset graphs, we show the grand mean firing rate with standard error bars for the control and oddball stimuli. The rate for oddball stimuli was signficantly greater than for control for each comparison (paired *t*-tests, one-tailed; all *P*-values < 0.002) There was no change in the later portion of the visual response for VT and VmT neurons (Fig. 6A and B).

In sum, the oddball manipulation shows that the pattern of effects seen in the peak of the transient response was consistent with the adaptation to light intensity computed by our model; however, a

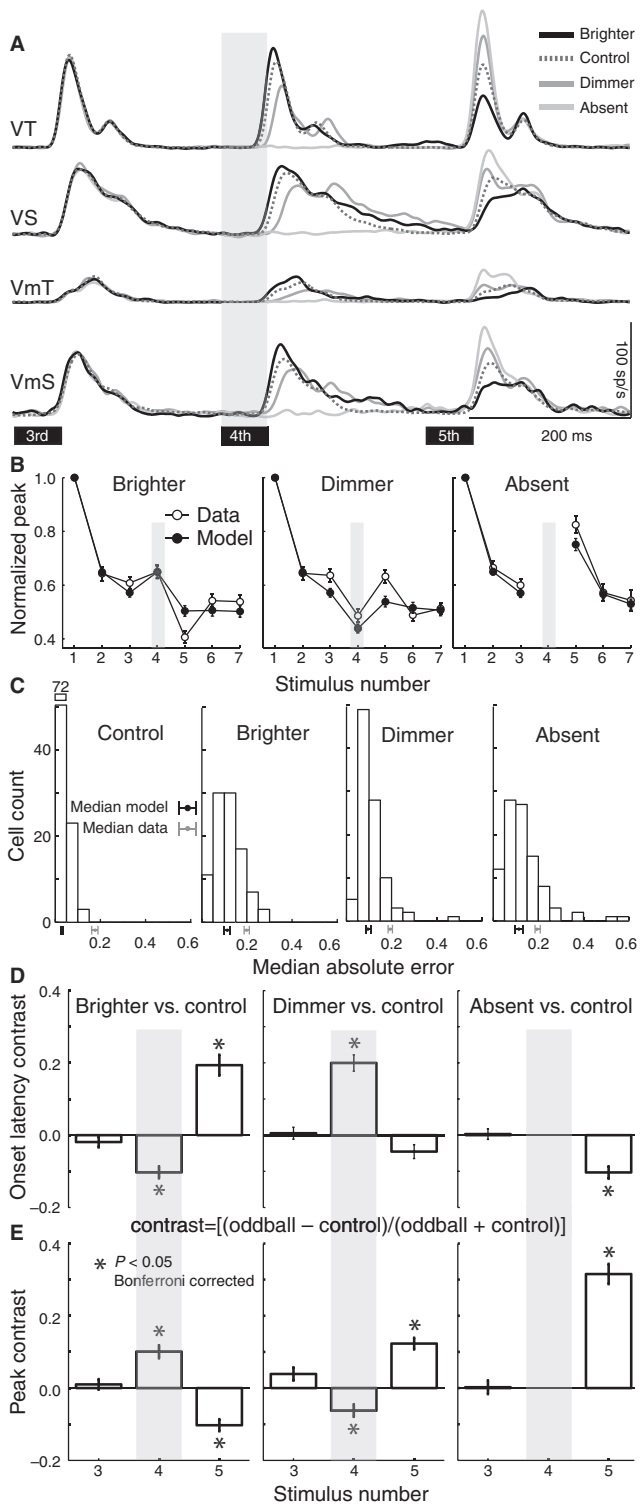

FIG. 5. Changes in the early transient part of visual responses to oddball stimuli presented in the fourth stimulus position in the sequence as depicted by the gray shaded bar across the different panels. (A) Population spike density waveforms showing the response to the third, fourth and fifth stimuli are plotted for the four classes of visually responsive neurons studied (same neurons as in Fig. 2). (B) Changes in the normalized peak response magnitude across stimulus number for model (closed symbols) and neural responses (open symbols) were not significantly different (two-factor repeated measures permutation test, $P = 0.61$). The paired-error test (see Materials and methods) indicated that, on average, each model was a better predictor of the peak activity of its corresponding real neuron than other neurons of the same class ($P < 0.01$). (C) Histograms of median absolute errors between each neuron and its corresponding model for the control and oddball conditions; conventions are the same as in Fig. 4D. (D, E) Mean normalized difference (i.e. contrast) in ROL (D) and mean normalized difference in peak response magnitude (E) between the control condition and each oddball condition calculated as [(control − oddball)/(control + oddball)]. Negative and positive values mean that oddball conditions had lower and higher values, respectively, than control conditions. Asterisk indicates a significant difference from 0 (Bonferroni corrected t-test, 1-tailed, $P < 0.05$).

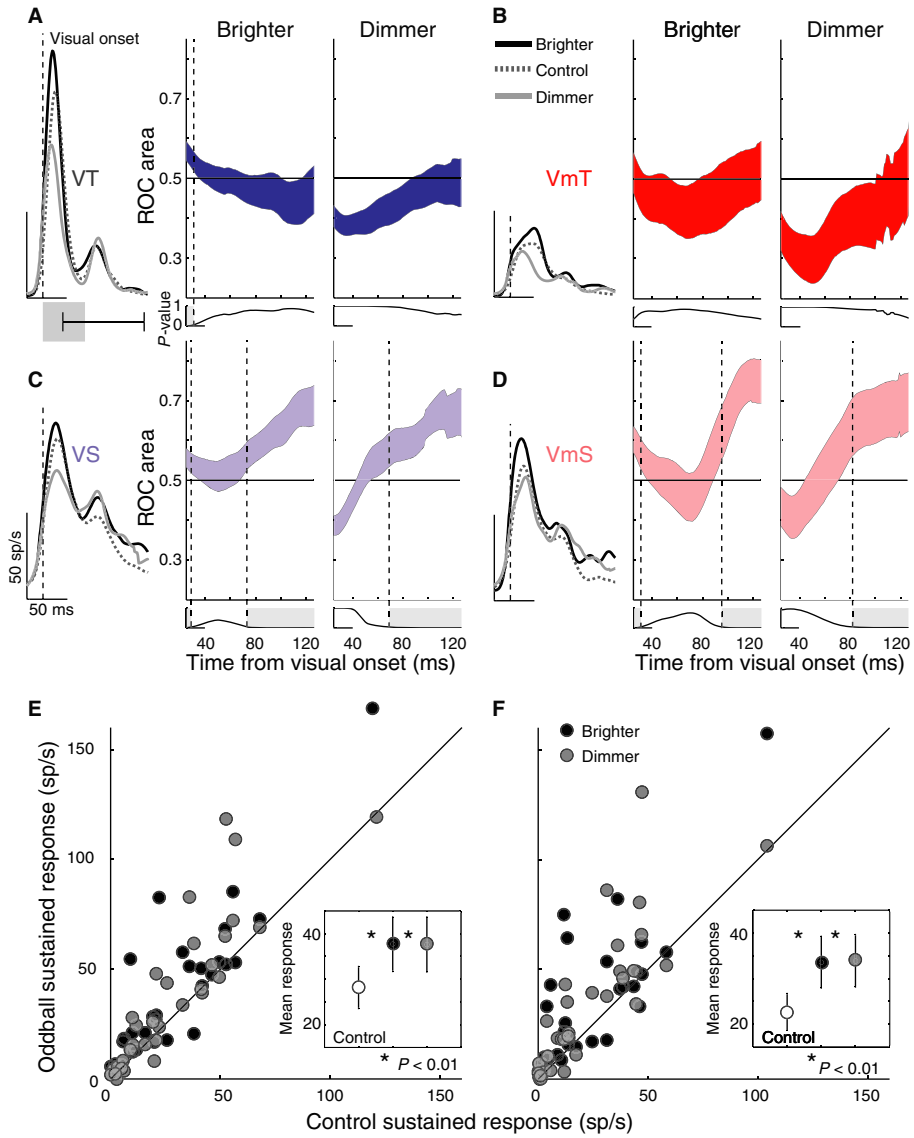

FIG. 6. Changes in the later sustained part of the visual response to oddball stimuli presented in the fourth stimulus position are shown for each neuron type: (A) VT; (B) VmT; (C) VS; (D) VmS. An ROC analysis was performed to determine at what points in time the later (sustained) part of the visual response became significantly different between control trials and either brighter or dimmer trials. The overlaid spike density functions show the average activity for the control, brighter and dimmer conditions aligned with the onset of the transient visual response (see Materials and methods) to the fourth stimulus. The filled colored regions represent the standard error of the ROC area across all cells of the same class, for the brighter and dimmer conditions separately. An ROC area of 0.5 or less indicates no difference between the control and oddball conditions at that particular time point, whereas values > 0.5 indicate the oddball condition had more activity on average. Each point was tested with a one-tailed *t*-test to determine whether the ROC area was significantly > 0.5. The *P*-value of this test is plotted below each ROC area plot, and the vertical dotted lines and light gray shaded regions indicate when the *P*-value crossed the significance threshold ($P < 0.05$). All ROC areas for all time points were normally distributed (Kolmogorov–Smirnov, $P < 0.05$). (E) Scatterplot of the mean sustained activity for the fourth stimulus in control trials vs. mean sustained activity in the brighter or dimmer oddball stimuli for VS neurons. The inset graph shows the population mean sustained activity with standard error bars for the control, brighter and dimmer stimuli. Asterisks show significant differences between oddball and control activity rates (paired *t*-test, one-tailed). (F) As described in E, but for VmS neurons.

significant dishabituation signal (enhanced responses to both brighter and dimmer stimuli) not seen in the model response was present in the later part of the visual response only in neurons with sustained activity.

Discussion

The timing and magnitude of the visual response of SC neurons underwent significant modification following stimulus repetition: the earliest part of the visual response decreased in magnitude and increased in latency with repetition (Fig. 2). The modulation of this early response with repetition was successfully modeled with our Bayesian adaptation model (Fig. 3), and predictions made about the effect of changing the rate of stimulus presentation (Fig. 4) and the intensity of rare stimuli (Fig. 5) were confirmed with neural data. The repetition effect was strongly dependent on the rate of stimulus presentation (Fig. 4), with the repetition effect increasing in magnitude as the interval between stimuli was reduced. For brighter or dimmer oddball stimuli, the main features of the repetition effect followed simple adaptation to light – larger, earlier responses to the brighter oddball stimulus, and smaller, later responses to the dimmer oddball stimulus, and the opposite pattern in response to the next (non-

oddball) stimulus. In contrast, the later, sustained component of the visual response was modulated much less by repetition, as observed previously in V4 (Motter, 2006), and was inconsistent with our Bayesian adaptation model. Finally, in response to either brighter or dimmer oddball stimuli, we observed an increase in response (e.g. a dishabituation) in this later sustained firing, suggestive of a 'novelty response'.

Comparison with other studies

Reductions in response magnitude with repetition have been previously observed in cortical areas, including V1 (Muller *et al.*, 1999), V4 (Motter, 2006) and frontal eye fields (Mayo & Sommer, 2008), and also in single neurons of the SC (Goldberg & Wurtz, 1972; Woods & Frost, 1977) and multiunit activity of the superficial layers of the SC (Mayo & Sommer, 2008). In the context of an attentional cueing task, a repetition effect has been described in the SC (Robinson & Kertzman, 1995; Dorris *et al.*, 2002; Bell *et al.*, 2004; Fecteau *et al.*, 2004) and lateral intraparietal area (Robinson *et al.*, 1995). The present study is the first to systematically explore the repetition effect using long stimulus sequences studied across different cell types and layers in the SC, the first to report the increase in response onset latency with repetition in the SC, and the first to explore the mechanism of this response decrement through modeling and experimental manipulations (oddball stimuli).

We observed a significant increase in ROL with repetition and with changes in stimulus intensity (oddball experiment). Modulation of ROL with intensity is consistent with previous reports on intensity modulations in the SC (Bell *et al.*, 2006; Li & Basso, 2008). Increases of ROL with stimulus repetition are evident in some results obtained from V4 (Motter, 2006; Hudson *et al.*, 2009), although they have not been explicitly described in the SC. There is one study in the SC which did not show an ROL increase with repetition (Mayo & Sommer, 2008), and there were interesting stimulus differences between their study and the present one that may explain why: the two stimuli in their sequence were shifted spatially in order to activate different retinal receptive fields but the same relatively large receptive fields in the frontal eye fields or superficial layers of the SC. Thus, their failure to see the ROL increases with repetition, while observing the decrease in response magnitude, may suggest that the ROL effect occurs very early in visual processing (e.g. retina, lateral geniculate nucleus of the thalamus, or input to V1), whereas the magnitude decrease occurs more centrally (e.g. V1 or beyond), although the anatomical locus of these effects remain to be explicitly tested. There are a few possible explanations for the ROL increase: there could be complete elimination of the earliest spikes of the response, owing to adaptation, which artificially shifts the ROL, or there could be reduced numbers of cells converging to provide the response, thus delaying generation of the first spikes by excitatory postsynaptic potentials. The increase in ROL is reminiscent of the increase in ROL as stimulus contrast is reduced (Bell *et al.*, 2006; Li & Basso, 2008), almost as though repetition was reducing the contrast of subsequent stimuli.

Mechanisms of adaptation

Grill-Spector *et al.* (2006) recently proposed three models for the mechanisms underlying adaptation. Adaptation may reflect a proportional reduction in firing rate in response to repetition (i.e. fatigue), a change in the tuning of neural responses for the repeated stimulus (i.e. sharpening), or a reduction in processing time for repeated stimuli (i.e. facilitation). On the basis of our data, the facilitation model can be discarded, because it predicts that the latency of the response (ROL) will be earlier with repetition, and we uniformly found the opposite. The sharpening model is possible, although it would predict that some neurons would have no response with repetition, and some (the best tuned for that stimulus) would show little response decrement. We found that, generally, SC neurons showed a graded reduction in response. Some form of the fatigue model is therefore most likely to account for the repetition effect observed on the early transient part of the visual response, and it is also the most closely related to our Bayesian model of adaptation. An important addition, however, is that we found a two-stage model with fast and slow dynamics to be necessary to best explain our neural data, a refinement indicating that at least 2 mechanisms (but possibly still within a single neuron) with different temporal sensitivities may be contributing to the adaptation effect. Note, however, that none of these models can yet account for the increase in ROL with repetition.

Alternatively, some portion of the response reduction may be affected locally in the SC by an increase in global inhibition from the basal ganglia. The intermediate layer of the SC projects the transient visual response monosynaptically to the substantia nigra compacta (Redgrave & Gurney, 2006), the response is then processed through the basal ganglia, and the substantia nigra pars reticulata (SNr) projects back to the intermediate layer of the SC (Hikosaka & Wurtz, 1983; Jiang *et al.*, 2003) to modulate neuronal firing via GABAergic synapses (Isa *et al.*, 1998; Kaneda *et al.*, 2008). A visual transient that is not accompanied by a response or reward (as in our simple fixation task) could result in increased SNr inhibition with each repetition (or less disinhibition), and thus reduced subsequent responses. The same mechanism could also account for our dishabituation effect following an intensity 'oddball' stimulus. VS and VmS neurons responded to oddball stimuli that were either brighter or dimmer with an increase in late sustained activity with a latency of approximately 140–160 ms after stimulus onset. A transient reduction of the inhibition from the SNr after an oddball stimulus is recognized by the basal ganglia as novel could account for the later increase in the sustained component of VS and VmS neuronal activity (i.e. disinhibition). This 'novelty signal' might then, in turn, be broadcast to the entire visual system from the SC (Boehnke & Munoz, 2008).

Implications for learning theory

In this study, we designed a paradigm to study simple learning phenomena in a behaving primate that have been studied previously in exquisite detail in *Aplysia* (Castellucci *et al.*, 1970; Carew *et al.*, 1971). Given the differences in the complexity of the organisms, it is not clear that terminology and concepts are easily transferable, but some discussion is at least warranted. The response decrement with repetition that we observed on the initial transient part of the visual response has been called 'habituation' in V4 (Motter, 2006) and 'adaptation' in the frontal eye fields (Mayo & Sommer, 2008). Given how that transient response changed with our stimulus intensity oddballs, we believe that this decrement in the transient component is best described as adaptation. We have described the increased sustained activity after the brighter or dimmer oddball stimuli as a 'dishabituation-like' or 'novelty' signal. It is also possible that this increase represents a phenomenon called sensitization (Marcus *et al.*, 1988; Hawkins *et al.*, 2006), which amplifies responses like the dishabituation process. Sensitization has been shown to be an

independent process from dishabituation because, at least in *Aplysia*, it develops at a different time (Rankin & Carew, 1988). Our experiment was not designed to differentiate these two processes, although sensitization usually requires a noxious stimulus, which we did not employ. We also did not objectively determine the discriminability of our stimuli, although the neurons clearly differentiated them. The use of brighter and dimmer stimuli as oddballs had the advantage of simplicity, and allowed for the dissociation of habituation from adaptation. However, as the stimuli were identical in shape, size, and color, there may have been a counteracting generalization process that prevented a larger recovery of response (dishabituation/sensitization) than might have been possible with a more distinctly different stimulus. These are questions for future studies. Importantly, this study represents an initial step in extending to primates the detailed understanding of these simple learning phenomenon achieved in simpler animals such as *Aplysia*, and the oculomotor system is an excellent candidate system with which to investigate these questions.

Implications for psychophysical studies

Our results are consistent with psychophysical findings on stimulus duration perception (Eagleman, 2008), where repeating stimuli are perceived as shorter in duration than an initial stimulus (Rose & Summers, 1995; Pariyadath & Eagleman, 2008) and any novel stimulus presented (Tse *et al.*, 2004; Pariyadath & Eagleman, 2007). In our sustained cell types, repetition reduced the size of visual responses, and novelty (brighter or dimmer oddball stimuli) caused an increased firing in the later sustained epoch. Thus, the first stimulus and any novel stimuli caused a larger sustained response than repeated stimuli, and may represent a neural correlate of the aforementioned perceptual findings. The timing of the novelty response also matches that of the N2 component of the human event-related potential in response to visual oddball stimuli (Folstein & Van Petten, 2008), a component that is thought to reflect detection of novelty or mismatch. We did not observe any response, early or late, when the fourth stimulus was absent (see Fig. 5A). A stimulus omission mismatch response in audition occurs only when the onset to onset time of the sequence of stimuli is < 150 ms (Yabe *et al.*, 1997), so perhaps it is not surprising that it was not observed with our longer ISI. Late event-related potential responses such as the P300 are observed with omitted visual stimuli (Tarkka & Stokic, 1998); however, the timing of such a response would coincide with the time when our neurons were responding to the fifth stimulus. The enhancement of the fifth stimulus response after a missing stimulus might, in part, reflect a P300 response, although it is difficult to determine this.

A previous visual event (attentional cue) also has implications for processing of a subsequent visual target for a manual or saccadic response: at separation intervals similar to those used here, the response to a subsequent target stimulus is slowed, a phenomenon referred to as ''inhibition of return'' or IOR (Klein, 2000; Fecteau & Munoz, 2006). We show that continued repetition of a visual stimulus (akin to having multiple cues) during fixation further reduces and delays the visual response. This presumably would lead to even slower reaction times and greater IOR. Indeed, it was recently shown that IOR for manual responses increased as the number of repeating cues is increased (Dukewich & Boehnke, 2008).

Information processing in the SC

Our results demonstrate that SC neurons' peak transient responses are consistent with a model of adaptation in which there is output of an information quantity related to the amount of learning caused by a new stimulus based on recent stimulation history. This is quite different from the most widely used quantitative definition of information (Shannon, 1948), where the information content of a piece of data, or a stimulus, is related to its probability (i.e. rare events are very informative). Although useful for the high-fidelity transmission of data, Shannon information does not quantify the subjective impact of stimuli on an observer – an important quantity when temporally changing signals are being processed.

Adaptation in the SC serves to rapidly decrease the early neural representation of repeating visual events at a particular spatial location (reducing the chance of reflexive orienting to that location), and to increase the representation of temporal outliers. Visual events that are not oriented upon first presentation, and are subsequently repeated, are not likely to contain immediately relevant information, and there is little to be learned. In this sense, adaptation acts as a simple and fast heuristic process to bias selection away from behaviorally irrelevant events in the absence of goal-directed orienting signals. Behaviorally relevant events may also manifest as more subtle changes in a stream of stimuli, and orienting to these novel events may require reinstatement of a previously adapted response. The slower dishabituation signal observed later in the response profile may serve as an additional heuristic process to support orienting, albeit delayed, to temporally adapted, yet novel, stimuli. Our data suggest that, by combining these heuristic processes, the primate orienting system achieves an efficient trade-off between fast selection of temporal outliers and slower detection of novel events.

Acknowledgements

We thank A. Lablans, D. Brien, M. Lewis and S. Hickman for outstanding technical support. This work was supported by research grants to D.P. Munoz and L. Itti from the Human Frontiers Science Program (RGP39/2005), National Science Foundation (Collaborative Research in Computational NeuroScience, BCS-0827764), and Defense Advanced Research Projects Agency (government contract no. HR0011-10-C-0034). Additional support was provided to L. Itti from the General Motors Corporation, and the Army Research Office (grant number W911NF-08-1-0360). Additional support was provided from Canadian Institutes of Health Research, no. CNS-90910, and the Canada Research Chair Program to D.P. Munoz, and an NIH Biomedical Informatics Training grant (LM-07443-01) to P.F. Baldi. The authors affirm that the views expressed herein are solely their own, and do not represent the views of the United States government or any agency thereof.

Abbreviations

ISI, interstimulus interval; KL, Kullback–Leibler; ROC, receiver operating characteristic; ROL, response onset latency; SC, superior colliculus; SNr, substantia nigra pars reticulata; VmS, visuomotor-sustained; VmT, visuomotor-transient; VS, visual-sustained; VT, visual-transient.

References

Baars, B.J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press, New York.

Baldi, P. & Itti, L. (2010) Of bits and wows: a bayesian theory of surprise with applications to attention. *Neural Netw.*, **23**, 649–666.

Bell, A.H., Fecteau, J.H. & Munoz, D.P. (2004) Using auditory and visual stimuli to investigate the behavioral and neuronal consequences of reflexive covert orienting. *J. Neurophysiol.*, **91**, 2172–2184.

Bell, A.H., Meredith, M.A., Van Opstal, A.J. & Munoz, D.P. (2006) Stimulus intensity modifies saccadic reaction time and visual response latency in the superior colliculus. *Exp. Brain Res.*, **174**, 53–59.

Boehnke, S.E. & Munoz, D.P. (2008) On the importance of the transient visual response in the superior colliculus. *Curr. Opin. Neurobiol.*, **18**, 544–551.

Brown, S.P. & Masland, R.H. (2001) Spatial scale and cellular substrate of contrast adaptation by retinal ganglion cells. *Nat. Neurosci.*, **4**, 44–51.
Carew, T.J., Castellucci, V.F. & Kandel, E.R. (1971) An analysis of dishabituation and sensitization of the gill-withdrawal reflex in aplysia. *Int. J. Neurosci.*, **2**, 79–98.
Castellucci, V., Pinsker, H., Kupfermann, I. & Kandel, E.R. (1970) Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in aplysia. *Science*, **167**, 1745–1748.
Clifford, C.W., Webster, M.A., Stanley, G.B., Stocker, A.A., Kohn, A., Sharpee, T.O. & Schwartz, O. (2007) Visual adaptation: neural, psychological and computational aspects. *Vision Res.*, **47**, 3125–3131.
Corneil, B.D., Munoz, D.P., Chapman, B.B., Admans, T. & Cushing, S.L. (2008) Neuromuscular consequences of reflexive covert orienting. *Nat. Neurosci.*, **11**, 13–15.
David, S.V., Vinje, W.E. & Gallant, J.L. (2004) Natural stimulus statistics alter the receptive field structure of v1 neurons. *J. Neurosci.*, **24**, 6991–7006.
Dean, P., Redgrave, P. & Westby, G.W. (1989) Event or emergency? Two response systems in the mammalian superior colliculus. *Trends Neurosci.*, **12**, 137–147.
Dean, I., Harper, N.S. & McAlpine, D. (2005) Neural population coding of sound level adapts to stimulus statistics. *Nat. Neurosci.*, **8**, 1684–1689.
Dorris, M.C., Klein, R.M., Everling, S. & Munoz, D.P. (2002) Contribution of the primate superior colliculus to inhibition of return. *J. Cogn. Neurosci.*, **14**, 1256–1263.
Dragoi, V. (2002) A feedforward model of suppressive and facilitatory habituation effects. *Biol. Cybern.*, **86**, 419–426.
Dragoi, V., Sharma, J., Miller, E.K. & Sur, M. (2002) Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. *Nat. Neurosci.*, **5**, 883–891.
Dukewich, K.R. & Boehnke, S.E. (2008) Cue repetition increases inhibition of return. *Neurosci. Lett.*, **448**, 231–235.
Eagleman, D.M. (2008) Human time perception and its illusions. *Curr. Opin. Neurobiol.*, **18**, 131–136.
Fecteau, J.H. & Munoz, D.P. (2006) Salience, relevance, and firing: a priority map for target selection. *Trends Cogn. Sci.*, **10**, 382–390.
Fecteau, J.H., Bell, A.H. & Munoz, D.P. (2004) Neural correlates of the automatic and goal-driven biases in orienting spatial attention. *J. Neurophysiol.*, **92**, 1728–1737.
Folstein, J.R. & Van Petten, C. (2008) Influence of cognitive control and mismatch on the N2 component of the ERP: a review. *Psychophysiology*, **45**, 152–170.
Goldberg, M.E. & Wurtz, R.H. (1972) Activity of superior colliculus in behaving monkey II. Effect of attention on neuronal responses. *J. Neurophysiol.*, **35**, 560–574.
Graham, J., Lin, C.S. & Kaas, J.H. (1979) Subcortical projections of six visual cortical areas in the owl monkey, *Aotus trivirgatus*. *J. Comp. Neurol.*, **187**, 557–580.
Grill-Spector, K., Henson, R. & Martin, A. (2006) Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn. Sci.*, **10**, 14–23.
Hawkins, R.D., Cohen, T.E. & Kandel, E.R. (2006) Dishabituation in aplysia can involve either reversal of habituation or superimposed sensitization. *Learn. Mem.*, **13**, 397–403.
Hikosaka, O. & Wurtz, R.H. (1983) Visual and oculomotor functions of monkey substantia nigra pars reticulata IV. Relation of substantia nigra to superior colliculus. *J. Neurophysiol.*, **49**, 1285–1301.
Hosoya, T., Baccus, S.A. & Meister, M. (2005) Dynamic predictive coding by the retina. *Nature*, **436**, 71–77.
Hudson, A.E., Schiff, N.D., Victor, J.D. & Purpura, K.P. (2009) Attentional modulation of adaptation in V4. *Eur. J. Neurosci.*, **30**, 151–171.
Huerta, M.F. & Harting, J.K. (1983) Sublamination within the superficial gray layer of the squirrel monkey: an analysis of the tectopulvinar projection using anterograde and retrograde transport methods. *Brain Res.*, **261**, 119–126.
Ingle, D. (1975) Sensorimotor function of the midbrain tectum II. Classes of visually guided behavior. *Neurosci. Res. Program Bull.*, **13**, 180–185.
Isa, T., Endo, T. & Saito, Y. (1998) The visuo-motor pathway in the local circuit of the rat superior colliculus. *J. Neurosci.*, **18**, 8496–8504.
Itti, L. & Baldi, P. (2005) A principled approach to detecting surprising events in video. *Proceedings 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, **1**, 631–637.
Jiang, H., Stein, B.E. & McHaffie, J.G. (2003) Opposing basal ganglia processes shape midbrain visuomotor activity bilaterally. *Nature*, **423**, 982–986.
Kaneda, K., Isa, K., Yanagawa, Y. & Isa, T. (2008) Nigral inhibition of GABAergic neurons in mouse superior colliculus. *J. Neurosci.*, **28**, 11071–11078.
Klein, R.M. (2000) Inhibition of return. *Trends Cogn. Sci.*, **4**, 138–147.
Kohn, A. (2007) Visual adaptation: physiology, mechanisms, and functional benefits. *J. Neurophysiol.*, **97**, 3155–3164.
Krekelberg, B., Boynton, G.M. & van Wezel, R.J. (2006) Adaptation: grom single cells to BOLD signals. *Trends Neurosci.*, **29**, 250–256.
Kullback, S. (1959). *Information Theory and Statistics*. Wiley, New York.
Lagarias, J.C., Reeds, J.A., Wright, M.H. & Wright, P.E. (1998) Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM J. Optimization*, **9**, 112–147.
Li, X. & Basso, M.A. (2008) Preparing to move increases the sensitivity of superior colliculus neurons. *J. Neurosci.*, **28**, 4561–4577.
Maffei, L., Fiorentini, A. & Bisti, S. (1973) Neural correlate of perceptual adaptation to gratings. *Science*, **182**, 1036–1038.
Marcus, E.A., Nolen, T.G., Rankin, C.H. & Carew, T.J. (1988) Behavioral dissociation of dishabituation, sensitization, and inhibition in aplysia. *Science*, **241**, 210–213.
Marino, R.A., Rodgers, C.K., Levy, R. & Munoz, D.P. (2008) Spatial relationships of visuomotor transformations in the superior colliculus map. *J. Neurophysiol.*, **100**, 2564–2576.
May, P.J. (2006) The mammalian superior colliculus: laminar structure and connections. *Prog. Brain Res.*, **151**, 321–378.
Mayo, J.P. & Sommer, M.A. (2008) Neuronal adaptation caused by sequential visual stimulation in the frontal eye field. *J. Neurophysiol.*, **100**, 1923–1935.
Mays, L.E. & Sparks, D.L. (1980) Dissociation of visual and saccade-related responses in superior colliculus neurons. *J. Neurophysiol.*, **43**, 207–232.
McPeek, R.M. & Keller, E.L. (2002) Saccade target selection in the superior colliculus during a visual search task. *J. Neurophysiol.*, **88**, 2019–2034.
Mohler, C.W. & Wurtz, R.H. (1976) Organization of monkey superior colliculus: intermediate layer cells discharging before eye movements. *J. Neurophysiol.*, **39**, 722–744.
Motter, B.C. (2006) Modulation of transient and sustained response components of V4 neurons by temporal crowding in flashed stimulus sequences. *J. Neurosci.*, **26**, 9683–9694.
Movshon, J.A. & Lennie, P. (1979) Pattern-selective adaptation in visual cortical neurones. *Nature*, **278**, 850–852.
Muller, J.R., Metha, A.B., Krauskopf, J. & Lennie, P. (1999) Rapid adaptation in visual cortex to the structure of images. *Science*, **285**, 1405–1408.
Munoz, D.P. & Wurtz, R.H. (1993) Fixation cells in monkey superior colliculus II. Reversible activation and deactivation. *J. Neurophysiol.*, **70**, 576–589.
Munoz, D.P. & Wurtz, R.H. (1995) Saccade-related activity in monkey superior colliculus I. Characteristics of burst and buildup cells. *J. Neurophysiol.*, **73**, 2313–2333.
Munoz, D.P., Dorris, M.C., Pare, M. & Everling, S. (2000) On your mark, get set: brainstem circuitry underlying saccadic initiation. *Can. J. Physiol. Pharmacol.*, **78**, 934–944.
Pariyadath, V. & Eagleman, D. (2007) The effect of predictability on subjective duration. *PLoS ONE*, **2**, e1264.
Pariyadath, V. & Eagleman, D.M. (2008) Brief subjective durations contract with repetition. *J. Vis.*, **8**, 11.1–11.6.
Rankin, C.H. & Carew, T.J. (1988) Dishabituation and sensitization emerge as separate processes during development in aplysia. *J. Neurosci.*, **8**, 197–211.
Redgrave, P. & Gurney, K. (2006) The short-latency dopamine signal: a role in discovering novel actions? *Nat. Rev. Neurosci.*, **7**, 967–975.
Robinson, D.L. & Kertzman, C. (1995) Covert orienting of attention in macaques III. Contributions of the superior colliculus. *J. Neurophysiol.*, **74**, 713–721.
Robinson, D.L., Bowman, E.M. & Kertzman, C. (1995) Covert orienting of attention in macaques II. Contributions of cortex. parietal. *J. Neurophysiol.*, **74**, 698–712.
Rose, D. & Summers, J. (1995) Duration illusions in a train of visual stimuli. *Perception*, **24**, 1177–1187.
Shannon, C.E. (1948) A mathematical theory of communication. *Bell System Tech. J.*, **27**, 379–423.
Sheather, S.J. & Jones, M.C. (1991) A reliable data-based bandwidth selection method for kernel density estimation. *J. R. Statist. Soc.*, **53**, 683–690.
Smirnakis, S.M., Berry, M.J., Warland, D.K., Bialek, W. & Meister, M. (1997) Adaptation of retinal processing to image contrast and spatial scale. *Nature*, **386**, 69–73.

Sokolov, E.N. (1963) Higher nervous functions; the orienting reflex. *Annu. Rev. Physiol.*, **25**, 545–580.
Solomon, S.G., Peirce, J.W., Dhruv, N.T. & Lennie, P. (2004) Profound contrast adaptation early in the visual pathway. *Neuron*, **42**, 155–162.
Stocker, A.A. & Simoncelli, E.P. (2006) Noise characteristics and prior expectations in human visual speed perception. *Nat. Neurosci.*, **9**, 578–585.
Tarkka, I.M. & Stokic, D.S. (1998) Source localization of P300 from oddball, single stimulus, and omitted-stimulus paradigms. *Brain Topogr.*, **11**, 141–151.
Thompson, K.G., Hanes, D.P., Bichot, N.P. & Schall, J.D. (1996) Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *J. Neurophysiol.*, **76**, 4040–4055.
Tigges, J. & Tigges, M. (1981) Distribution of retinofugal and corticofugal axon terminals in the superior colliculus of squirrel monkey. *Invest. Ophthalmol. Vis. Sci.*, **20**, 149–158.
Tse, P.U., Intriligator, J., Rivest, J. & Cavanagh, P. (2004) Attention and the subjective expansion of time. *Percept. Psychophys.*, **66**, 1171–1189.
White, B.J., Boehnke, S.E., Marino, R.A., Itti, L. & Munoz, D.P. (2009) Color-related signals in the primate superior colliculus. *J. Neurosci.*, **29**, 12159–12166.
Woods, E.J. & Frost, B.J. (1977) Adaptation and habituation characteristics of tectal neurons in the pigeon. *Exp. Brain Res.*, **27**, 347–354.
Yabe, H., Tervaniemi, M., Reinikainen, K. & Naatanen, R. (1997) Temporal window of integration revealed by MMN to sound omission. *Neuroreport*, **8**, 1971–1974.