Rapid Neural Coding in the Retina with Relative Spike Latencies

Tim Gollisch* and Markus Meister†

Natural vision is a highly dynamic process. Frequent body, head, and eye movements constantly bring new images onto the retina for brief periods, challenging our understanding of the neural code for vision. We report that certain retinal ganglion cells encode the spatial structure of a briefly presented image in the relative timing of their first spikes. This code is found to be largely invariant to stimulus contrast and robust to noisy fluctuations in response latencies. Mechanistically, the observed response characteristics result from different kinetics in two retinal pathways ("ON" and "OFF") that converge onto ganglion cells. This mechanism allows the retina to rapidly and reliably transmit new spatial information with the very first spikes emitted by a neural population.

During natural vision, our gaze remains fixed for a mere fraction of a second. Sudden movements of the eye, called saccades, partition visual processing into short episodes (*1*, *2*). Each saccade exchanges the image that falls onto the retina; the new visual stimulus is then encoded into neural activity to be transmitted to the brain. Our visual system can analyze and classify a new complex scene in less than 150 ms (*3*), but the nature of the neural code that underlies this rapid visual processing has been elusive. Neurons in the vertebrate retina fire with remarkable temporal precision (*4*, *5*), so single spikes can, in principle, carry substantial information about visual stimuli. In order to assess how the retina transmits new visual information after a saccade, we investigated the responses of retinal ganglion cells to flashed visual images.

Spike trains were recorded simultaneously from many ganglion cells in the isolated salamander retina. The stimulus was a uniform gray field followed by appearance of a square-wave grating. Eight different shifted versions of the grating were used in a pseudo-random sequence. A ganglion cell typically responded to the appearance of the grating with a short burst of spikes (Fig. 1), and the vast majority of cells responded to most or even all of the stimuli. We characterized each burst by two numbers: the latency of the first spike after stimulus onset and the total spike count in the burst. For certain cell types, in particular fast and biphasic OFF cells (fig. S2), the spike count was very similar for all stimuli (Fig. 1B). By contrast, the spike latency for these cells varied across stimuli by as much as 40 ms (Fig. 1C). For repeats of the same stimulus, this latency was very reproducible, with a standard deviation of only 3 to 5 ms.

We calculated how much information the spike latency or the spike count conveys about which grating had been presented. Perfect iden-

tification of the stimulus among eight possibilities amounts to a maximum of 3 bits. The spike latency of a ganglion cell transmitted up to 2 bits of information on a single trial. The spike count provided considerably less information for the majority of all recorded cells (Fig. 1D). Subsequent brain regions may thus learn more about the stimulus from noting the time of the first spike after stimulus onset than by waiting for all spikes and noting the average firing rate.

In several sensory systems, shorter spike latencies result from stronger stimulation (*6–9*). This does not account for the present dependence of latency on spatial pattern. Stronger stimuli often generate higher spike counts, and indeed, gratings of higher contrast produced both more spikes and shorter latencies (fig. S3). By contrast, we observed a pronounced spatial tuning of the spike latency even when there were virtually no variations in spike count (Fig. 1); in some cases, shorter latencies even occurred in combination with fewer spikes (fig. S3).

Downstream brain centers can interpret the latency of a single neuron only if the onset time of the stimulus is known (*10*). If the new retinal image was initiated by an eye movement, then the brain does know the onset time, but it is unclear whether this motor information gets distributed to visual centers. We therefore asked what information can be extracted from visual signals alone by comparing latencies from neurons in the population (Fig. 2). For many pairs of ganglion cells, the difference between first spike times was strongly tuned with respect to the presented stimuli (Fig. 2C). In fact, the information contained in the latency difference reached values higher than 2 bits (Fig. 2D)—more than that from any single-cell absolute latency. One reason was the particular robustness of the latency difference to retinal noise. Each cell's latency underwent some trial-to-trial variation, but these fluctuations were often positively correlated in cells recorded simultaneously; when cell 1 fired earlier than usual, cell 2 tended to do the same (Fig. 2B). As a result, the latency difference (Fig. 2C) fluctuated less than expected from the noise in individual cell latencies (Fig. 2A). To assess the relevance of this compensation, we destroyed the noise correlations artificially by pairing the response of

cell 1 with the response of cell 2 on the subsequent trial; this led to a substantial information loss of up to ~20% (Fig. 2D).

Stimuli of greater strength tend to produce shorter spike latencies in the sensory response. If the latencies of different neurons in the population are affected in similar fashion, downstream circuits might use the difference in spike latencies to extract stimulus quality independent of stimulus strength (*11*). We therefore presented the flashed gratings at different contrast levels. As expected, individual latencies increased at lower contrast. However, the shape of the latency tuning curve was well preserved at each contrast level (Fig. 2, E and F). Furthermore, the contrast-dependent shifts of the latency tuning curves were similar for different cells. As a result, the latency difference between two neurons was almost perfectly invariant to changes in contrast (Fig. 2G). In fact, a downstream decoder could recover almost all the spatial information without knowing anything about the contrast of the stimulus (Fig. 2H).

How can the observed latency code be explained in terms of neural mechanisms? We start by considering a standard framework for visual responses (*12*, *13*) and exploring its prediction for first-spike latencies. In this picture, the stimulus is first passed through a linear filter that summarizes retinal integration of the image over space and time (Fig. 3A). The "activation" signal emerging from the filter can be interpreted as the membrane potential of the ganglion cell. When this signal crosses a preset threshold, the model neuron fires a spike. For each ganglion cell, we measured the spatiotemporal filter in a separate reverse-correlation experiment (fig. S1) (*14*), whereas the threshold remained as a single free parameter. In using this model to process grating stimuli, one quickly finds that it cannot account for the observed responses. Because the stimulus is integrated linearly, a certain grating may elicit strong excitation, but then its sign-reversed counterpart will elicit inhibition and produce no spikes at all, counter to what was observed experimentally (Fig. 3, B and C).

Thus, one is forced to include nonlinear processing steps. Ganglion cells draw their excitatory input from bipolar cells (Fig. 3D). These have comparatively small receptive fields [<100 μm; (*15*)] and respond to light in essentially linear fashion (*16*), but transmission to retinal ganglion cells may involve a degree of rectification (*17–19*). Furthermore, bipolar cells come in two major types: ON bipolars are excited by an increase in light intensity and OFF bipolars by a decrease. Individual ganglion cells can receive inputs from both types (*8*, *20–24*). To include this structure of the inner retina into the model, we replaced the single spatiotemporal filter by a set of parallel filters that mimicked spatially local ON and OFF bipolar cells (Fig. 3E). Transmission from bipolar to ganglion cells was approximated by a half-wave rectifying function (Fig. 3F). Under these conditions, all stimuli led to

---

Department of Molecular and Cellular Biology and Center for Brain Science, Harvard University, 16 Divinity Avenue, Cambridge, MA 02138, USA.

*Present address: Max Planck Institute of Neurobiology, Am Klopferspitz 18, 82152 Martinsried, Germany.
†To whom correspondence should be addressed. E-mail: meister@fas.harvard.edu

excitatory activation of the ganglion cell (Fig. 3G). Moreover, the predicted timing of the first spike agreed remarkably well with the observed latency tuning curve (Fig. 3H) and outperformed alternative circuit schemes (fig. S6) for virtually all cells with substantial latency tuning (fig. S7).

Note that all the elements of this model are well-known components of retinal circuitry. Closer inspection reveals how the circuit accomplishes latency tuning. First, the rectifying synapses ensure that every stimulus excites the ganglion cell: any image change within the receptive field will activate some set of bipolar cells that transmit their excitation to the ganglion cell (18). The lasting and synchronous activation of both ON and OFF pathways by flashed gratings emphasizes their nonlinear summation and, thus, the need for separate filters in the model. By contrast, earlier studies of ganglion cells (13) were based on spatially homogeneous stimuli that only transiently activate one pathway at a time and, therefore, allow for simpler models (14). Second, the measured ON filters have slower kinetics than the OFF filters (Fig. 3E and fig. S5), such that ON stimuli affect ganglion cell spiking ~30 ms later than OFF stimuli. This is consistent with prior observations (23) and probably results from a signal transduction delay at the synapse between photoreceptors and ON bipolars (25, 26). Thus, the proportion of light and dark stimulation within the receptive field determines the relative contribution of the ON and OFF pathways and modulates the time of the first threshold crossing.

This hypothesis for latency coding relies intimately on the convergence of parallel neuronal pathways with intrinsic kinetic differences. If this picture is correct, removal of one of the pathways should lead to a breakdown of latency tuning. We therefore exposed the retina to 2-amino-4-phosphono-butyrate, a metabotropic glutamate receptor agonist that blocks neural transmission to ON bipolar cells (26). The results were as predicted: Fast OFF ganglion cells ceased responding to about half of the stimuli (fig. S8), consistent with a loss of all the ON filters of the model as shown in Fig. 3F.

Although grating stimuli are convenient for systematic investigations, they do not capture the complex statistics of natural scenes. We thus briefly flashed a photographic image onto the retina (Fig. 4A). Across repeated presentations, the image was shifted to many different locations. In this way, spikes from a single ganglion cell could be used to simulate a population of identical neurons with different receptive field locations. Responses to the natural image resembled those to the gratings. For fast OFF cells, almost all image presentations elicited spike bursts that varied in latency by about 40 ms (Fig. 4B), and this latency was systematically related to the stimulus. Indeed, by simply plotting the differential spike latencies as a gray-scale code, we obtained a rather faithful neural representation of the raw visual image (Fig. 4C). This demonstrates the high quality of the latency information. Subsequent brain regions could use this for local image computations; for example, a neuron that detects spike coincidence among multiple ganglion cells would be selective for contour lines or edges in the image.

The corresponding neural image created from spike counts (Fig. 4D) is more blurred and noisy, and the highest values are observed near edges in the stimulus. In the flat regions, the center-surround antagonism of ganglion cell receptive fields (fig. S1) reduces firing activity. But because the effect of the receptive-field surround is delayed relative to the center (20), it does not affect the first spike in a burst. In fact, latency and spike count may serve to encode complementary stimulus features, which could support a rapid scene analysis with subsequent refinement (27). Furthermore, in natural vision, the ongoing fixational eye movements after a saccade may well affect the spike count throughout fixation, but should have negligible effects on the timing of the first spike.
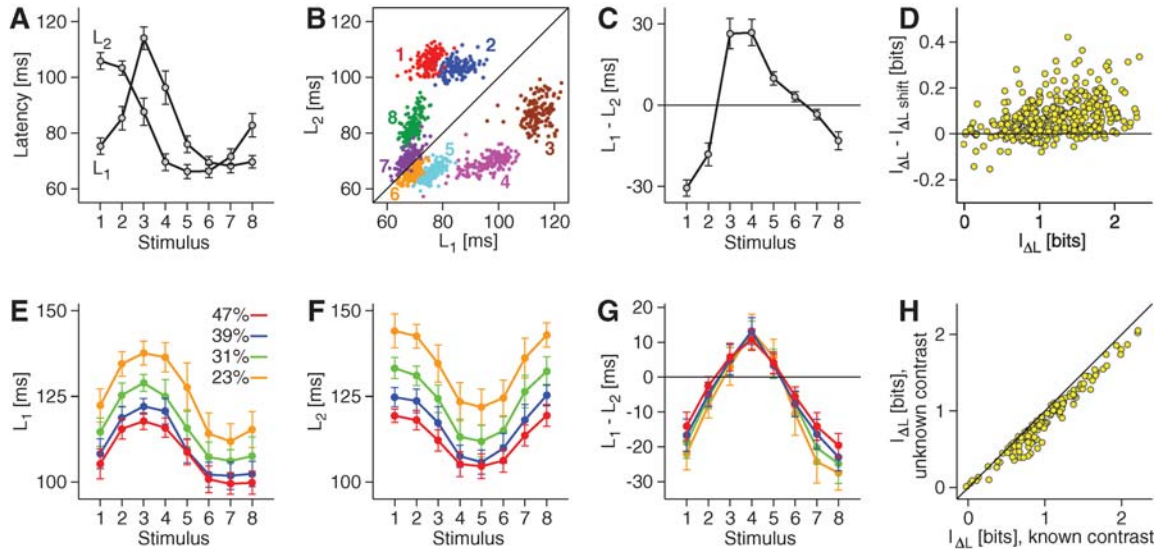
Altogether, our results suggest that a population code based on differential spike latencies can be a powerful mechanism to rapidly transmit a new visual scene. Rapid saccades are ubiquitous in animal vision (1). In salamanders, they result from turns of the head and constitute a vital part of the approach to a prey (2). During a saccade, many ganglion cells are strongly suppressed (19, 28), such that the first spike after a saccade is easily recognized. The differential latency of these
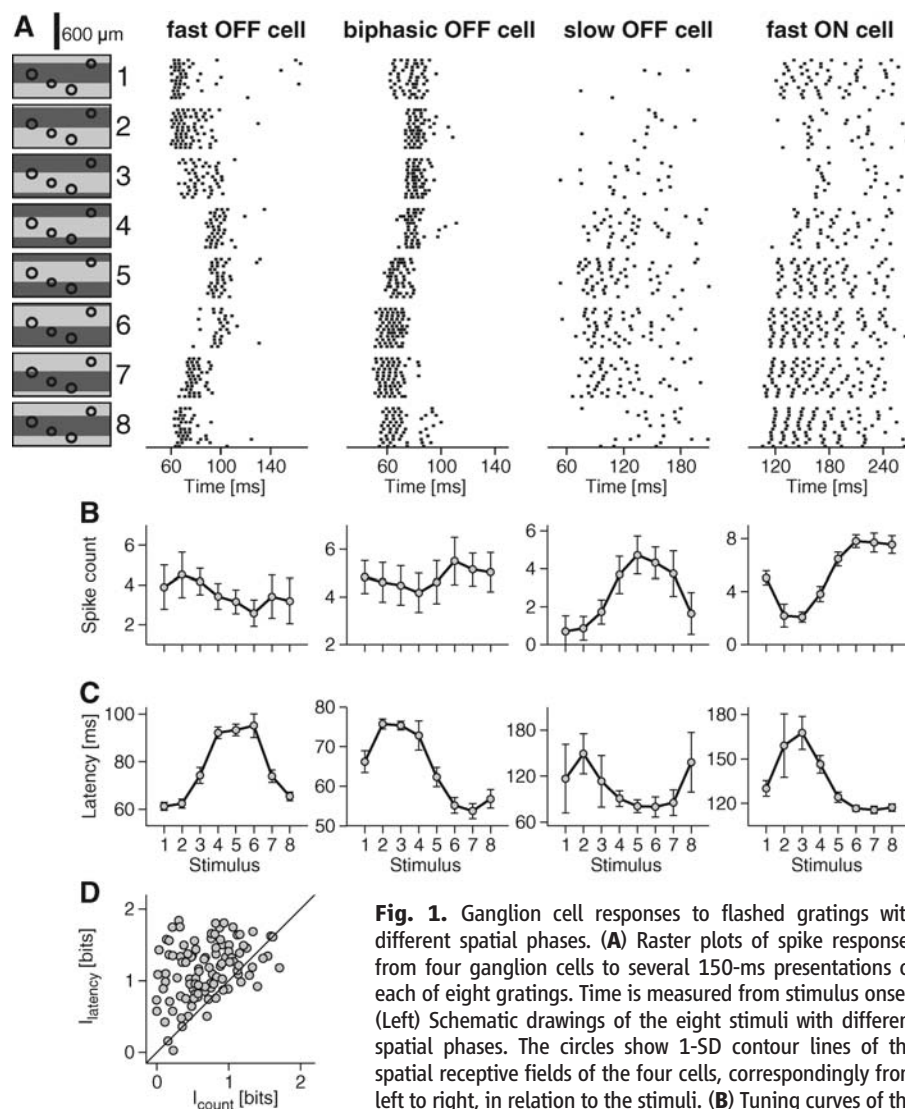
Fig. 1. Ganglion cell responses to flashed gratings with different spatial phases. (**A**) Raster plots of spike responses from four ganglion cells to several 150-ms presentations of each of eight gratings. Time is measured from stimulus onset. (Left) Schematic drawings of the eight stimuli with different spatial phases. The circles show 1-SD contour lines of the spatial receptive fields of the four cells, correspondingly from left to right, in relation to the stimuli. (**B**) Tuning curves of the elicited spike count. Here and in subsequent figures, all error bars show the standard deviation across trials with the same stimulus. (**C**) Tuning curves of the first-spike latency. "Fast OFF" and "biphasic OFF" cells typically showed strong tuning in the latency and only mild tuning in spike count; despite their names, these cell types receive input from both ON and OFF pathways (19). "Slow OFF" and "ON" cells, on the other hand, displayed good tuning in the spike count and often did not respond with spikes to all stimuli. The relatively long latencies are typical for cold-blooded animals. (**D**) Information about the stimulus identity contained in the spike count and in the latency, respectively, for all recorded cells. For a subdivision of the data by ganglion cell type, see fig. S2.
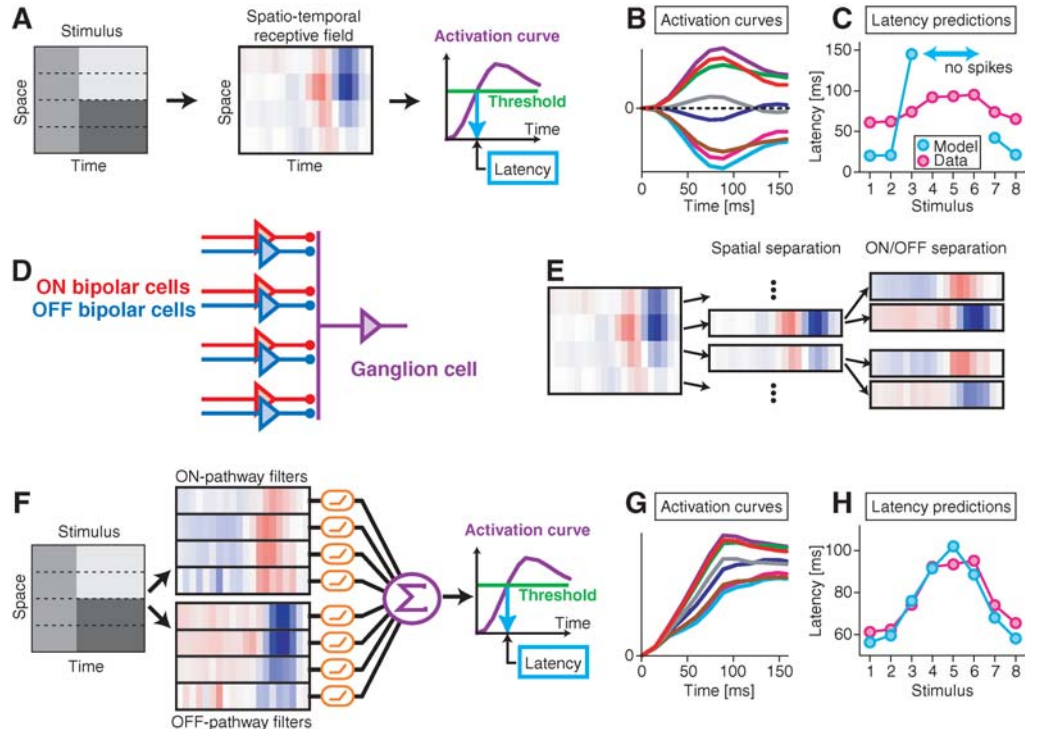
Fig. 2. (**A** to **D**) Encoding by relative latencies of pairs of ganglion cells. (A) Latency tuning curves for two simultaneously recorded fast OFF cells. (B) Scatter plot of latencies (L_1, L_2) for the two cells. The diagonally elongated distributions of the data show that L_1 and L_2 were positively correlated across trials with the same stimulus (*14*). (C) Tuning curve of the latency difference $L_1 - L_2$. (D) Information theoretical analysis of latency differences from simultaneously recorded cell pairs. The information $I_{\Delta L}$ about the stimulus contained in the latency difference $\Delta L = L_1 - L_2$ is plotted against the information loss $I_{\Delta L} - I_{\Delta L\ shift}$ that occurred when L_2 was shifted by one trial with respect to L_1. (**E** to **H**) Contrast-invariant encoding by pairs of ganglion cells. (E) Latency tuning curves for a fast OFF cell whose responses were recorded for flashed gratings at different Michelson contrast levels. (F) Latency tuning curves of a second, simultaneously recorded fast OFF cell. (G) Tuning curves of the latency difference for the two neurons. (H) Information about the stimulus pattern, carried by latency differences whether the contrast level is known or not. All cell pairs that were recorded at the four different contrast levels were analyzed. The data points near the diagonal show that little information is lost by ignoring the stimulus contrast.

Fig. 3. Modeling the response latencies of retinal ganglion cells. (**A**) Standard framework for modeling ganglion cell responses. The stimulus (left) is gray illumination followed by a grating. This is convolved with a spatiotemporal filter (middle) representing the ganglion cell's receptive field (fig. S1). When the resulting activation curve exceeds a preset threshold, the first spike is fired (right). (**B**) Activation curves computed for each of the eight grating stimuli, by using the measured spatiotemporal filter for a sample fast OFF ganglion cell (first cell in Fig. 1). (**C**) Predicted and measured dependence of the latency on the stimulus. The threshold is the only free parameter of the model and was optimized from a χ^2 fit to the measured latency tuning curve. Several stimuli did not lead to positive activation and thus did not predict spikes. (**D**) Retinal interneuron pathways that motivate a revised model. Each small subregion of the receptive field activates both ON and OFF bipolar cells. The ganglion cell pools inputs across subregions and from both bipolar types. (**E**) Separation of the spatiotemporal filter into spatial subunits and subsequently into ON and OFF pathway contributions. [See (*14*) and fig. S5 for measurement of these contributions.] (**F**) Multi-pathway model of the response: Each subregion of the stimulus is passed through an ON filter and an OFF filter. The filter output is half-wave rectified and then pooled with all other outputs to yield the activation curve. (**G**) Activation curves computed for each of the eight grating stimuli, using the model in (F). Note that each stimulus produced excitation. (**H**) Measured latency tuning curves and predictions of the model in (F), after optimization of the threshold.

spikes encodes fine spatial detail (Figs. 1 and 4), yet it is almost entirely invariant to the overall stimulus contrast level (Fig. 2). Furthermore, it is robust to retinal noise (Fig. 2), and it provides information about the pattern in the shortest possible time, namely, with the very first spikes.

Many vertebrates have specific ganglion cells, often multiple types, that combine inputs from both ON and OFF pathways (*20–23*), as evidenced by their ON/OFF response characteristics or by their dendritic morphology that connects to both ON and OFF bipolar cells. These neurons

are candidate carriers of a latency code (Fig. 3). Synapses in the early visual pathway are very efficient, such that short spike bursts are reliably transmitted from retina to cortex (*29*). Moreover, certain neurons in visual cortex are exquisitely sensitive to the coincidence of spikes on their
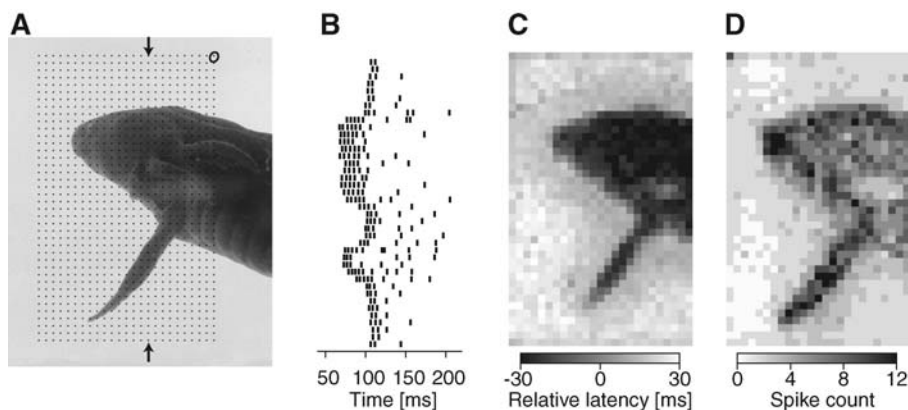
Fig. 4. Responses of a fast OFF ganglion cell to a flashed natural image. (For results from other cell types, see fig. S9.) (**A**) Photograph of a swimming salamander larva projected on the retina. The ellipse in the upper right corner shows a sample 1-SD outline of a ganglion cell receptive field. In each of 1000 presentations, the image was shifted slightly, and the grid of dots marks the resulting centers of the receptive field. Presentations were separated by gray illumination at the mean intensity of the photograph. The image onset produced luminance changes at most locations. (**B**) Spike trains of the ganglion cell for receptive-field locations along the column marked by the arrows in (A). (**C**) Gray-scale plot of the differential spike latency on single-trial presentations at the locations marked with dots in (A). The reference latency was chosen as the average value at all locations (10). (**D**) Corresponding gray-scale plot of the spike counts.

afferents (30), which is one possible readout mechanism for a latency code. Cortical neurons themselves carry substantial sensory information in their response latencies (6, 7, 31). Thus, it is conceivable that early aspects of sensory processing operate on the basis of the classification of spike latency patterns.

References and Notes

1. M. F. Land, *J. Comp. Physiol. A* **185**, 341 (1999).
2. C. Werner, W. Himstedt, *Zool. Jahrb. Abt. Allg. Zool. Physiol. Tiere* **89**, 359 (1985).
3. S. Thorpe, D. Fize, C. Marlot, *Nature* **381**, 520 (1996).
4. M. Meister, M. J. Berry, *Neuron* **22**, 435 (1999).
5. V. J. Uzzell, E. J. Chichilnisky, *J. Neurophysiol.* **92**, 780 (2004).
6. T. J. Gawne, T. W. Kjaer, B. J. Richmond, *J. Neurophysiol.* **76**, 1356 (1996).
7. D. S. Reich, F. Mechler, J. D. Victor, *J. Neurophysiol.* **85**, 1039 (2001).
8. M. Greschner, A. Thiel, J. Kretzberg, J. Ammermüller, *J. Neurophysiol.* **96**, 2845 (2006).
9. N. B. Sawtell, A. Williams, C. C. Bell, *Curr. Opin. Neurobiol.* **15**, 437 (2005).
10. S. M. Chase, E. D. Young, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 5175 (2007).
11. J. J. Hopfield, *Nature* **376**, 33 (1995).
12. E. J. Chichilnisky, *Network* **12**, 199 (2001).
13. J. Keat, P. Reinagel, R. C. Reid, M. Meister, *Neuron* **30**, 803 (2001).
14. Materials and methods are available as supporting material on *Science* Online.
15. W. A. Hare, W. G. Owen, *J. Neurophysiol.* **76**, 2005 (1996).
16. S. A. Baccus, M. Meister, *Neuron* **36**, 909 (2002).
17. J. D. Victor, R. M. Shapley, *J. Gen. Physiol.* **74**, 671 (1979).
18. J. B. Demb, K. Zaghloul, L. Haarsma, P. Sterling, *J. Neurosci.* **21**, 7447 (2001).
19. M. N. Geffen, S. E. de Vries, M. Meister, *PLoS Biol.* **5**, e65 (2007).
20. F. S. Werblin, J. E. Dowling, *J. Neurophysiol.* **32**, 339 (1969).
21. F. M. de Monasterio, *J. Neurophysiol.* **41**, 1435 (1978).
22. F. R. Amthor, E. S. Takahashi, C. W. Oyster, *J. Comp. Neurol.* **280**, 97 (1989).
23. D. A. Burkhardt, P. K. Fahey, M. Sikora, *Vis. Neurosci.* **15**, 219 (1998).
24. J. L. Coombs, D. Van Der List, L. M. Chalupa, *J. Comp. Neurol.* **503**, 803 (2007).
25. J. F. Ashmore, D. R. Copenhagen, *Nature* **288**, 84 (1980).
26. X. L. Yang, *Prog. Neurobiol.* **73**, 127 (2004).
27. M. D. Menz, R. D. Freeman, *Nat. Neurosci.* **6**, 59 (2003).
28. B. Roska, F. Werblin, *Nat. Neurosci.* **6**, 600 (2003).
29. P. Kara, R. C. Reid, *J. Neurosci.* **23**, 8547 (2003).
30. W. M. Usrey, J. M. Alonso, R. C. Reid, *J. Neurosci.* **20**, 5461 (2000).
31. S. Panzeri, R. S. Petersen, S. R. Schultz, M. Lebedev, M. E. Diamond, *Neuron* **29**, 769 (2001).
32. We thank F. Engert and members of the Meister laboratory for advice. This work was supported by grants from the National Eye Institute (M.M.) and the Human Frontier Science Program Organization (T.G.).

Predicting Human Interactive Learning by Regret-Driven Neural Networks

Davide Marchiori[1] and Massimo Warglien[2]*

Much of human learning in a social context has an interactive nature: What an individual learns is affected by what other individuals are learning at the same time. Games represent a widely accepted paradigm for representing interactive decision-making. We explored the potential value of neural networks for modeling and predicting human interactive learning in repeated games. We found that even very simple learning networks, driven by regret-based feedback, accurately predict observed human behavior in different experiments on 21 games with unique equilibria in mixed strategies. Introducing regret in the feedback dramatically improved the performance of the neural network. We show that regret-based models provide better predictions of learning than established economic models.

The surge of interest in the neural bases of economic behavior (1–3) prompts the question of how well neural networks can model human interactive decision-making (4). This question implies two issues: the choice of the network architecture and the selection of input information to the network that has to be both economically and neurophysiologically motivated.

Interactive learning differs from individual learning in that, given *n* agents, each agent adapts to behaviors that are modified by the concurrent learning of the other *n*–1 agents. It has an obvious relevance in economic contexts, but (more generally) much of human learning that occurs in social contexts has an interactive nature. Experimental game theory has provided a large set of laboratory data on human interactive learning in repeated games (5), often contradicting the predictions of standard game theory. The need for models of interactive learning in games arises from the difficulties of ordinary game-solution concepts to explain both the trajectories and the long-run stationary state of experimentally observed human behavior in repeated games. Games with unique equilibria in mixed strategies are an especially interesting case, because Nash equilibrium not only fails to approximate behavior in early rounds but also is often a poor predictor of the stable behavior emerging in the long run.

Until now, two main modeling strategies have been used with some success in trying to fit and predict how humans learn in repeated games in a laboratory setting. One modeling strategy extends a classical paradigm of learning theory (i.e., rein-

[1]Interdepartmental Center for Research Training in Economics and Management (CIFREM), University of Trento, Italy. [2]Advanced School of Economics and Department of Business Economics and Management, Ca' Foscari University, Venezia, Italy.

*To whom correspondence should be addressed. E-mail: warglien@unive.it