Neuron, Vol. 40, 823–833, November 13, 2003, Copyright ©2003 by Cell Press

The Information Content of Receptive Fields

Thomas L. Adelman,[1,*] William Bialek,[2] and Robert M. Olberg[3]
[1]Department of Molecular Biology
[2]Department of Physics
Princeton University
Princeton, New Jersey 08544
[3]Department of Biological Sciences
Union College
Schenectady, New York 12308

Summary

The nervous system must observe a complex world and produce appropriate, sometimes complex, behavioral responses. In contrast to this complexity, neural responses are often characterized through very simple descriptions such as receptive fields or tuning curves. Do these characterizations adequately reflect the true dimensionality reduction that takes place in the nervous system, or are they merely convenient oversimplifications? Here we address this question for the target-selective descending neurons (TSDNs) of the dragonfly. Using extracellular multielectrode recordings of a population of TSDNs, we quantify the completeness of the receptive field description of these cells and conclude that the information in independent instantaneous position and velocity receptive fields accounts for 70%–90% of the total information in single spikes. Thus, we demonstrate that this simple receptive field model is close to a complete description of the features in the stimulus that evoke TSDN response.

Introduction

Stimuli in the natural world are extremely complex, and a complete record of the sensory stimuli impinging on the nervous system would be immense. Yet the neurons that sense, encode, and interpret these complex inputs are generally characterized using extremely simple and low-dimensional descriptions. This raises the question of whether these characterizations are merely utilitarian simplifications, or whether they correctly and completely describe a real dimensionality reduction that takes place in the nervous system. Here we address this question for the target-selective descending neurons (TSDNs) of dragonflies.

A commonly used low-dimensional description, and one of the most productive concepts in neuroscience, has been the idea of a receptive field for visual neurons. In the simplest and most qualitative view, the receptive field outlines that region of the visual space to which a neuron is responsive. More quantitatively, we might distinguish subregions of the receptive field that are excitatory or inhibitory ("on" or "off"), so that the probability of the neuron generating an action potential is increased or decreased by light in the respective subregion. The classical examples of these ideas are the (on)center-(off)surround structure of the receptive fields in retinal ganglion cells of frogs (Barlow, 1953) and cats (Kuffler, 1953), the similar lateral inhibitory structure in *Limulus* retina (Hartline and Ratliff, 1958), and the oriented on and off subregions of simple cells in primary visual cortex (Hubel and Wiesel, 1962). The concept of a receptive field has been generalized to encompass "fields" that are defined over different spaces and modalities, such as the spectrotemporal receptive fields of auditory neurons (Theunissen et al., 2000), wind direction receptive fields of the cockroach cercal system (Kolton and Camhi, 1995), somatosensory ("barrel") cortex receptive fields of the rat whisker system (Armstrong-James and Fox, 1987), and the place fields of the hippocampal neurons (O'Keefe and Dostrovsky, 1971).

In the classical work, receptive fields were mapped using simple stimuli, and it is natural to ask whether the description derived in these simple experiments generalizes to more natural stimulus conditions. Certainly in the visual system it is known that spatially restricted stimuli presented in a larger context can generate different responses than they would generate in isolation, even if the context is defined by image elements that lie completely outside of the classical receptive field (for an early overview of these phenomena, see Allman et al., 1985). In order to test the applicability of the receptive field concept in more natural stimulus contexts, we need to translate the schematic receptive field that one often draws as a summary of the experiments into a more quantitative model. Further, one has to define a measure of success or "goodness of fit" that allows us to test this model in the richer stimulus context.

In the general case, sensory stimuli are complex functions of time. Neurons respond to these stimuli with action potentials or spikes. A model of the neuron provides some way of predicting the spike train from the stimulus, and a receptive field model tells us that rather than looking, for example, at the full history of the visual movie leading up to a spike, the probability of a spike will be determined only by some limited features of the movie such as the position of a spot or the weighted average of light intensity over a small region. Such models are very useful, but raise an obvious question: given a candidate model, how do we test whether it is correct? One idea is to compare the predicted versus actual probability of a spike at each instant of time; the model yields the predicted probability directly from the movie, and the actual spiking probability of the neuron can be determined by repeating the movie many times as in a conventional poststimulus time histogram. But once the predicted and actual probabilities, or equivalently the time-dependent spike rates $r(t)$, have been determined, there is still no obvious way to compare them. For example, it might be that calculating the mean square difference between the predicted and actual spike rates would give a reasonable measure of their similarity. But it is not clear that this would capture our intuition about goodness of fit: should we, for example, weight more

*Correspondence: tadelman@princeton.edu

heavily differences at times where the measured response is most reproducible, or perhaps where the response is maximal, or most likely to generate responses from downstream neurons? Even if we agree on a metric for comparing different models of the neural response, it is not clear that this metric would provide an interpretable absolute measure of the quality of the model—a score of 0.8 out of 1.0 seems better than 0.5, but is 0.8 good in absolute terms?

Shannon's information theory (Shannon, 1948) provides us with a model-independent way of asking how well a receptive field model captures the structure of the neural response. Specifically, we can think of the neuron as establishing some correlation between sensory inputs and spike train output; the mutual information between input and output provides the unique measure of this correlation, which is consistent with several intuitively plausible conditions. When we narrow our view of the input stimulus from its full structure down to what is "visible" through the window provided by the receptive field, we necessarily lose some of this information; stated more positively, a successful receptive field model would have the property that the stimulus as seen through the receptive field captures most if not all of the mutual information between the stimulus and the spike train. In this approach, the absolute quality of the receptive field model is measured by the information loss, which has an immediate interpretation in terms of the number of different sensory stimuli that can be reliably distinguished. Our goal in this paper is to show how this general principle can be translated into practical tools for assessing the success of the receptive field picture in a reasonably natural stimulus context. Related recent work includes Agüera y Arcas et al. (2003) and Sharpee et al. (2003).

Here we apply these ideas to the target-selective descending neurons (TSDNs) of the dragonfly visual system. Dragonflies make frequent aerial pursuits of flying targets both for prey capture and for interactions with other dragonflies (Corbet, 1999). The cells most clearly involved in controlling the pursuit task are the TSDNs—a set of 16 bilaterally symmetric visual interneurons in the neck connective, descending from the brain and synapsing within the thoracic ganglia. The TSDNs have several properties that make them good candidates for being the chief guiding neurons linking vision to target pursuit behavior (Olberg, 1986; Frye and Olberg, 1995): they respond only to small spots moving in the visual field; they have large axon diameters, and therefore high-speed spike conduction as required for fast pursuits; and they do not synapse directly onto muscles, but spikes induced by intracellular stimulation of the TSDNs do result in wing motion. TSDN responses depend on both the position and velocity of the target: a cell responds only if the target is in a particular range of positions (the position receptive field) and simultaneously has a velocity from a particular range of vector velocities (the velocity receptive field). If either of these conditions is not met, the cell does not fire. The position receptive fields for all of the TSDNs are large, often covering an area of the visual field that is $45^\circ \times 45^\circ$ or greater.

For the dragonfly TSDNs, the simplest receptive field model is just a map of spike probability (i.e., spike rate) as a function of target position and (vector) velocity at one instant of time, each in two dimensions. In principle one could imagine sensitivity to particular combinations of position and velocity (e.g., circumferential motion), but the simplest hypothesis is that the probability of generating an action potential is a product of terms reflecting independent position and velocity sensitivity. This is an enormous simplification, since the cells could instead signal complex features of the dynamic target trajectory—accelerations, turns, etc.—and since these are high-level neurons involved in a complicated task, it would not be unreasonable to expect this complexity to be reflected in the neural response. Thus, as for visual neurons in vertebrate retina or mammalian cortex, the receptive field model of dragonfly TSDNs states that spiking is determined by (and hence presumably encodes for the animal) a very limited set of features extracted from a rich dynamic world of visual inputs.

The set of TSDNs provides an interesting test case for studies of a neural code because behaviorally relevant stimuli are encoded by a population of modest size but clearly not just by individual neurons. Much current thinking about coding in populations of neurons assumes the validity of a receptive field or tuning curve model for individual cells, but direct tests of such models in natural stimulus contexts remain to be done. We have succeeded in recording simultaneously from most of the population of TSDNs using a new flexible electrode technology, taking advantage both of the dragonfly anatomy and of novel fabrication methods. Here we focus on the analysis of these data one neuron at a time to test the validity of the receptive field model for each cell. This preparation also provides us with an ideal model system in which to study more complex aspects of the neural code; subsequent work will consider population coding, the significance of single spikes for target-trajectory discrimination, and the effect of complexity and correlations in the stimulus on the neural response times.

Results

We recorded the spiking activity of multiple TSDNs while the dragonfly watched a movie of a small, dark target moving on a light background (Figure 1). The target moved along 31 random, continuous trajectories, all drawn from the same probability distribution. One of these was repeated 30 times and the 30 remaining were presented only once. The goal of this paper is to propose a simple receptive field model that characterizes the response of these cells and, equally important, to test whether this model is correct. We do this by analyzing the responses to the 30 repeated and 30 nonrepeated trajectories in slightly different ways.

The 30 repeated trajectories are well suited for studying the precision of the neural response. A short example of the response of one the TSDNs is shown in Figure 1E. Here the variability in the spike times between identical stimuli gives a sense of the timing precision that can be conveyed by these cells. The 30 nonrepeated trajectories are better suited for quantifying the model because they sample a larger fraction of the stimulus space. It is reasonable to expect that these two quantifications of the neural response should be related. For example, if the timing and number of spikes were precise and

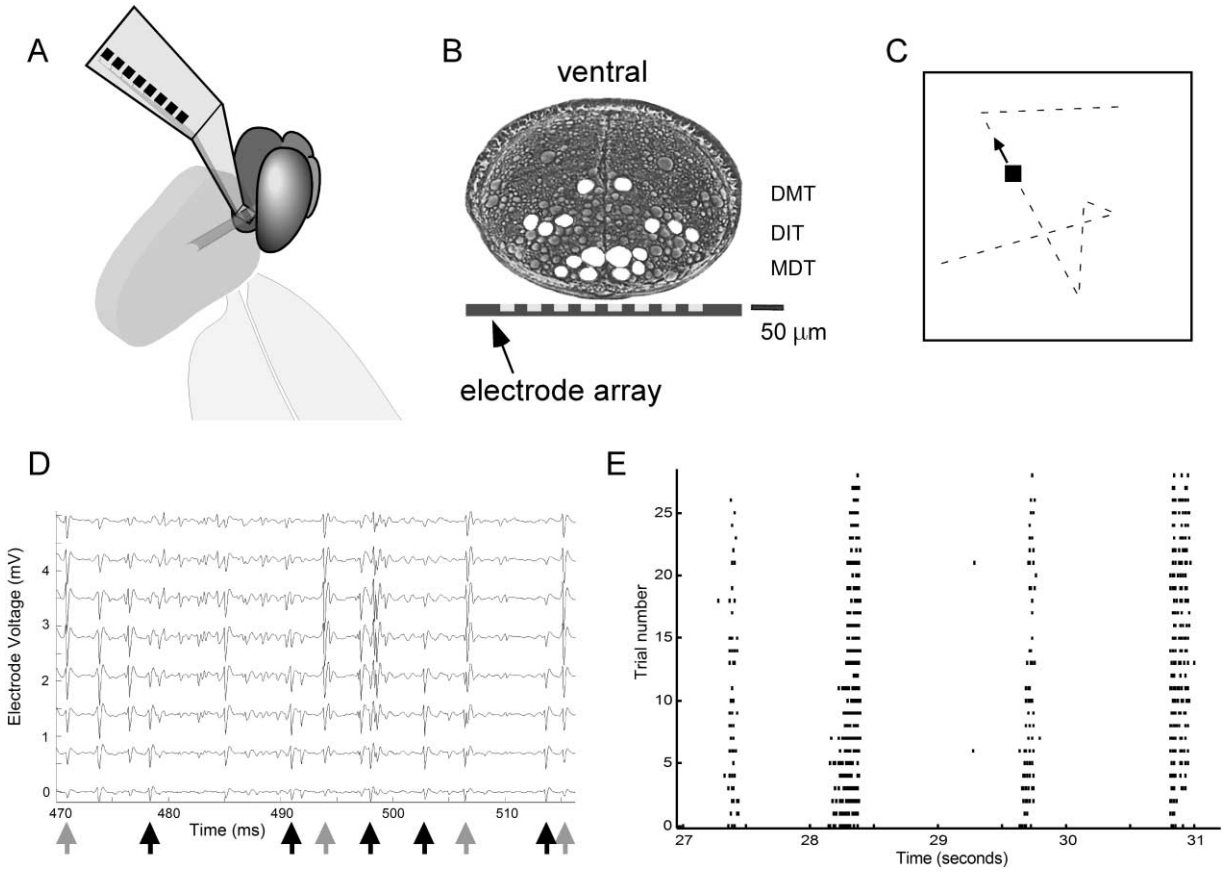

Figure 1. Dissection and Recording

(A) The dragonfly is positioned ventral side up, and a dissection is made in the neck to allow access to the ventral nerve cord. An eight-channel electrode array is positioned under the nerve cord (on the dorsal side).

(B) A schematic of the cross-section of the ventral nerve cord with the electrode array positioned dorsal to the nerve cord. During recording, the nerve cord is gently pressed against the array's eight electrodes (shown as white rectangles) using a small platform pressing from the ventral side (not shown). The TSDNs are marked in white, and the three tracks are identified to the left of the nerve cord. In this experiment, we only recorded from the DIT and MDT tracks.

(C) A schematic of the stimulus. A dark spot moves on a light background with a trajectory consisting of connected 0.25 s segments within which the spot moves at constant velocity (see Experimental Procedures). Here the dashed line indicates a short piece of a trajectory and the arrow indicates the direction of motion of the spot. (Only the spot is visible to the dragonfly; the arrow and dashed line are guides to the reader.)

(D) A short section of the recording from the electrode array. A single spike from one TSDN will generally be seen on multiple electrodes, and different patterns of spike weights can be used to distinguish the different TSDNs. Here, for example, repeated and distinct weight patterns can be readily distinguished by visual inspection. Some readers may be interested to compare that, for example, the spikes at {471, 494, 507, 515} (gray arrows) are clearly different from the spikes at {478, 491, 498, 503, 514} (black arrows).

(E) A section of the raster plot of the response to the repeated stimuli for cell A. The stimulus is repeated 30 times, and the spike times are indicated in the figure. A single trial lasts 30 s, a 4 s section of which is shown.

consistent across repeated trials, then a good descriptive model of the neuron should also exhibit this precision; in particular, the receptive fields should have correspondingly sharp, well-defined features. Here we use ideas from information theory to make this intuition precise. In the first section of the Results, we examine the information content of the receptive field model, and in the second section we look at the model-independent measure of information derived from the 30 repeated trajectories. We then compare these results to test how much of the total information is captured by the model.

Information in Position and Velocity Receptive Fields

The most general description of the target motion is a two-dimensional movie, and a full description of the stimulus history preceding each spike would be quite complex. Instead of using this complex history, we hope to find a low-dimensional description of the stimuli that lead to a spike. The model we develop and quantify is a severe simplification: that the responses of the TSDNs are driven by instantaneous snapshots of the target position and velocity.

On observing a single spike at a particular time, we know more about the position and velocity of the target than we did before observing the spike. Specifically, before observing a spike, we can only base our estimate of target position on the prior distribution (Figure 2), but after we see a spike, we can localize the target more precisely because we know that it is in the receptive field of the cell. Figure 3A shows position receptive fields

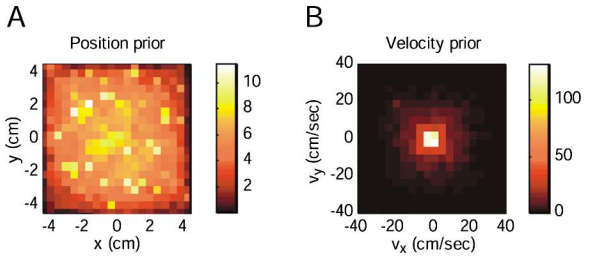

Figure 2. Sampling of the Stimulus Space by the Moving Target: The Position and Velocity prior Distributions

(A) Sampling of display screen position. The color scale value indicates the amount of time spent in each bin in seconds. The orientation shown is as viewed by the dragonfly (positive vertical corresponds to dorsal, and positive horizontal corresponds to the dragonfly's right). Positions are indicated in display coordinates as displacement from the center of the screen. The horizontal and vertical edges each correspond to a viewing angle of 74°.

(B) Sampling of velocity space. The orientation is the same as in (A) and the axes correspond to the velocity in display coordinates. The maximum velocity (40 cm/s) corresponds to approximately 330°/s (see Experimental Procedures).

of two cells calculated using the 30 nonrepeated trajectories (see Experimental Procedures). The receptive fields are notably different from the prior distribution, and it is this reduction in the uncertainty of the target position on observation of a single spike that is quantified by the mutual information between target position and single spikes, $I(\vec{x}; \text{spike})$. Similarly, $I(\vec{v}; \text{spike})$ gives the mutual information between target velocity and single spikes. Each information value is calculated directly from the appropriate receptive field and prior distribution through Equation 3 in the Experimental Procedures.

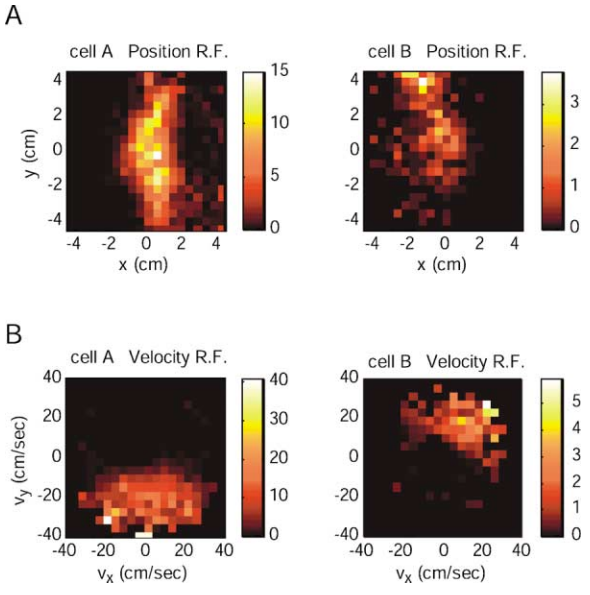

Figure 3. Position and Velocity Receptive Fields for Two of the Cells Studied Here

Each receptive field (A, position; B, velocity) is calculated at the time offset where its information peaks (Figure 5). The axes here are as described in Figure 2. The color bar indicates the spike rate in spikes/s.

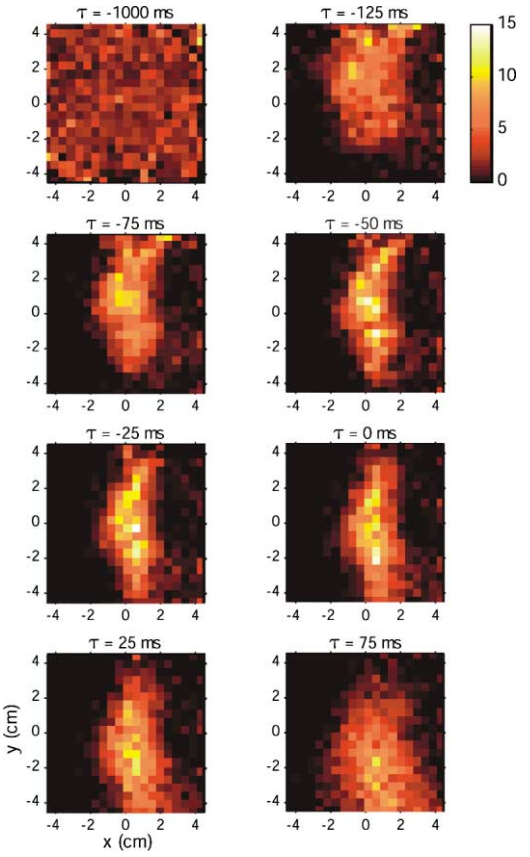

Figure 4. Time Evolution of the Position Receptive Field of Cell A

Each figure is a receptive field calculated using a stimulus position taken at a time that is shifted from the spike time by the indicated time offset τ. The time offsets span symmetric time steps from −25 ms (where the information peaks for this cell; Figure 5). Details of the calculation are discussed in the text. Note that the receptive field seems to move from the top to the bottom of the screen as the latency shortens. This evolving structure is caused by the velocity receptive field (Figure 3): to produce a spike, each individual trajectory must pass through the receptive field while also moving from top to bottom. The color bar indicates the spike rate in spikes/s.

Due to neural delays, a spike might provide more information about target position prior to the spike than at the time of the spike. Therefore, in constructing the receptive field, we may not want to use the position of the target at the time of the spike, but prefer to use the position at a time prior to the spike. To understand the receptive field, it is therefore necessary to determine which delay is best. To determine the appropriate delay, we calculate a set of position and velocity receptive fields using a range of delays (Figure 4 shows this for the position receptive field of cell A) and calculate the information in each, creating a plot of information value versus delay, as shown in Figure 5 for cells A and B. We are primarily interested in the delay that results in the most information. For all cells and for both position and velocity, these curves peak near zero delay (generally ±50 ms) and go to zero far from the spike time—since a spike carries no information about the stimuli of the distant past or future. The peak of these curves gives the time of the stimulus for which a spike carries the

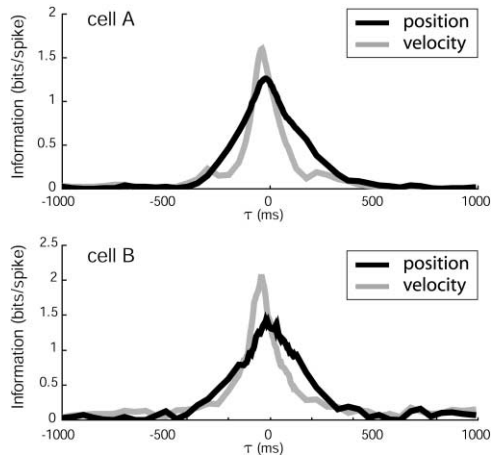

Figure 5. Information in Position and Velocity Receptive Fields versus Time Offset

A negative time offset corresponds to the stimulus occurring prior to the spike. Correlations in the stimulus broaden and add extra structure to these curves. For example, the side peaks in velocity are caused by a false receptive field that will occur from the velocities immediately preceding and following the velocity at the main peak time.

most information. The delay times of the peaks for all of the cells are summarized in Figure 6A.

In Figure 6B we report the information in the receptive fields calculated with the time offset corresponding to the largest information value for each of the cells. Here it can be seen that, in all of the cells, spikes contain more information about target velocity than about position. This is reasonable in a pursuit task and is consistent with dragonfly behavior. Rather than taking a direct path toward the target (where information about velocity is not required), dragonflies move along intercepting trajectories (Olberg et al., 2000), and such trajectories require information about both position and velocity.

The information in the independent receptive field model is just the sum of these two information values, $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$, since the statement of independence is equivalent to saying that the information values simply add. Like the assumption that the position and velocity receptive fields are the interesting things to calculate, this assumption of independence is also part of the model and will be implicitly tested in the next section. This sum is shown in Figure 6C for each cell (together with the total information per spike, which will be discussed in the next section).

In addition to the information values, which are the primary goals of this analysis, the delay times in Figure 6A are of interest on their own. It is striking that all of the velocity delays are quite similar (45 $\pm$ 5 ms), which may indicate that this value is intrinsic to the dragonfly's calculation of velocity. The position delays are not so consistent. These delays are generally shorter than for velocity and possibly indicate that the dragonfly responds to target position more quickly than target velocity, although other explanations are possible. For example, the short delays could instead be due to the dragonfly extrapolating the target position from some time in the past (e.g., -40 ms) to some time in the more recent past (e.g., -20 ms). Two cells have positive values for the position delays, which indicates that the information in the position receptive field peaks *after*

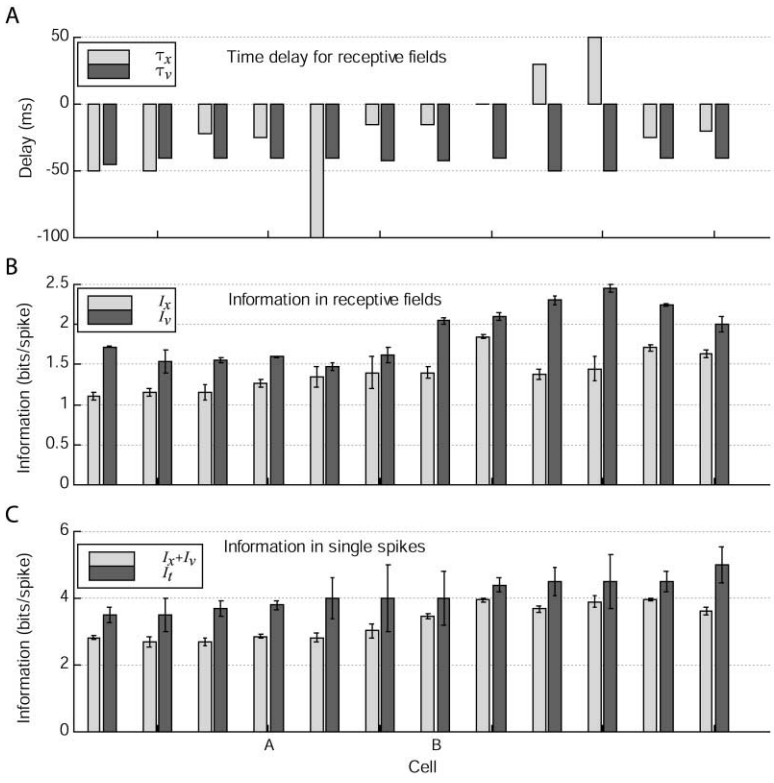

Figure 6. Histograms Comparing Several Quantities for Each of the Cells Studied Here

The two example cells used throughout this text are indicated as cell A and B, and the other cells are not labeled.

(A) The offset time that gives the maximal information when used in calculation of the position (τ_x) and velocity (τ_v) receptive fields (see text and Figure 5). Negative (positive) times indicate that the information peaks at a time before (after) the spike.

(B) Information values in the position ($I(\vec{x}; \text{spike})$, or I_x) and velocity ($I(\vec{v}; \text{spike})$, or I_v) receptive fields calculated at the time delays in (A).

(C) Comparing the information in the receptive field model ($I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$, or $I_x + I_v$) to the stimulus-independent total information per spike (I_t).

the spike has occurred. This most likely results from an extrapolation of target position that is performed by the dragonfly. This interpretation is consistent with the observation that the cells with the largest values of information about velocity, and therefore most able to make the best extrapolation, are the ones that exhibit the extrapolation most clearly. This result is not surprising for a pursuit task: for example, in catching a ball, one reaches or runs to where the ball will be, not to where it is at the time the reach is initiated.

The comparison and addition of the information values used in this section needs to be slightly qualified due to correlations between position and velocity in the stimulus. These correlations are unavoidable in this experiment due to constraints imposed by using continuous trajectories in a finite display size. At issue is that the correlations limit the interpretation of $I(\vec{x}; \text{spike})$ and $I(\vec{v}; \text{spike})$ as independent quantities, as is implicitly assumed when we compare or add these terms. The effect of these correlations can be treated as a small error in the values of $I(\vec{x}; \text{spike})$ and $I(\vec{v}; \text{spike})$, and in the Experimental Procedures we show that the error to use in comparing $I(\vec{x}; \text{spike})$ and $I(\vec{v}; \text{spike})$ is less than 0.1 bits for all cells, and that the error to use in adding $I(\vec{x}; \text{spike})$ and $I(\vec{v}; \text{spike})$ is less than 0.25 bits. In all of the cells, $I(\vec{x}; \text{spike})$ and $I(\vec{v}; \text{spike})$ are much larger than these errors.

Total Information per Spike

In the previous section, by considering the spike-conditional stimulus ensemble (de Ruyter van Steveninck and Bialek, 1988), we determined how much information the spikes provide about particular stimulus features, in our case the position and velocity of a target at a particular time. However, we can regard the spikes purely as an output, without reference to the stimulus, and directly calculate the amount of information they contain (Brenner et al., 2000).

In general, the information in a neural response is determined by its entropy, reduced by the entropy of the noise in the response, which limits its capacity to transmit meaningful information. Since we are focusing on the arrival time of single spikes, the total entropy is easy to compute and the noise entropy can be estimated by repeating the same stimulus many times to determine the variability of the response. This quantifies the sense that the response to repeated stimuli indicates the precision of a cell's response to features in the stimulus and can be easily calculated using the responses of the TSDNs to the 30 repeated trials. Crucially, neither the computation of the total entropy nor the noise entropy depends on any reference to features in the stimulus. We refer to the difference of these entropies as the total information, or I_t. Thus, I_t gives the single-spike information without assumptions about the features in the stimulus that influence production of a spike (such as the receptive fields used in the previous section).

The information about a low-dimensional feature of the stimulus must always be less than the total information (see Appendix). In our case, I_t specifies the total information carried by single spikes and $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$ refers to the information that single spikes have about particular aspects of the stimulus, so it must always be the case that

$$I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike}) \leq I_t. \quad (1)$$

The important point of this equation is that it allows us to compare the total information (I_t) to the amount of information captured by a low-dimensional model ($I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$). If the model is very good, then we expect these values to be very close, but if the model is poor, we will find that spikes have significantly less information about the corresponding features of the stimulus.

This comparison is the primary goal of this paper, and it is made directly in Figure 6C, where I_t and $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$ are plotted together for each cell. From this we conclude that the independent position and velocity receptive field model provides a good description of the single-spike response of these cells. Specifically, for all cells, the model accounts for most of the information in single spikes and falls short by only 0.5 to 1.5 bits, depending on the cell.

The formal procedures for calculating $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$ and I_t are superficially quite different, and it may not seem at first that these two quantities should be directly compared. But this comparison is valid, as is shown in the Appendix. To gain an intuition of this comparison, Figure 7 shows a simulation of a cell with a simplified, one-dimensional receptive field. The figure shows a target trajectory and the simulated response of a cell to this target. There are two spike-conditional projections of motion, each of which give response distributions from which we can calculate an information: from the projection onto the time axis we determine the total information per spike, and the projection onto the position axis gives the information in the receptive field. In this sense, these information measures are just different projections of the same quantity and can be quantitatively compared. Furthermore, if a cell were sensitive to more features of the stimulus trajectories, then the projections of the rate onto the position axis would necessarily be less structured than its projection onto the time axis, and hence information would be lost in the construction of the receptive field model. More generally, no projection can result in a richer distribution of rates than observed along the time axis, and hence it is this description that sets the limit on the information conveyed by single spikes.

Discussion

The primary purpose of this paper is to test whether the independent position and velocity receptive fields provide a valid and complete description of single spikes in the dragonfly TSDNs. We can answer this question quantitatively by comparing the information in the receptive field model ($I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$) to the total, model-independent information (I_t). If all of the total information can be accounted for by the receptive field description, then we know that the description completely captures the stimulus features to which the cell is responding. If not, we will be able to tell by how much the model falls short. For all cells studied here, we find that $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike}) \approx I_t$, but that $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$ generally falls short by 0.5 to 1.5 bits. That is, for all of the measured TSDNs, position and velocity

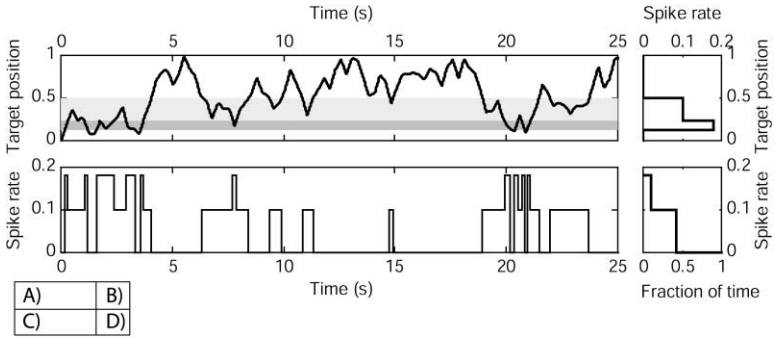

Figure 7. Schematic Illustrating the Comparison of Total Information to Information in the Receptive Field for a One-Dimensional Trajectory

We have predetermined the receptive field of this simulated cell and calculated the spike rate completely from this receptive field. As shown in this schematic, the probability distributions from which the information values are calculated are different projections of the spike rate, and therefore it is meaningful to compare them.

(A) The stimulus trajectory. The two bands indicate the positions that are active in the receptive field.

(B) The receptive field of the simulated cell. This is used to determine the information in the model (Equation 3).

(C and D) The simulated average response rate (C, the poststimulus time histogram). In an experiment, this would be determined by repeating a stimulus over multiple trails as in Figure 1E, but in this simulation it is determined directly from the receptive field. This is used to determine the total information (Equation 8). This calculation is invariant to shuffling and rescaling of the time axis, and this is done in (D), which shows the rates from (C) reorganized into descending order with a normalized abscissa. Because the rate distribution in (B) is nearly identical to (D), the information values calculated from these distributions will also be very close.

Of course, in an actual experiment, the distributions are found in a slightly different order: first we present the stimulus to the animal (A), then we determine the average spike rate (C), and from this we can determine the receptive field (B) and the rate distribution (D). Furthermore, in the experiments presented in this paper where the receptive field space is very large, we calculate the receptive fields by presentations of many trajectories, rather than the repeated presentation of the same trajectory. Finally, the trajectory here has been chosen so that its prior on the receptive field space is uniform. This is a simplification that allows us to compare by inspection the target position and spike rate distributions. Without this, as is the case in the experiments presented here, we must normalize by the prior.

receptive fields account for most (70%–90%) of the information in single spikes.

It is of some surprise that such a simple receptive field description captures so much of the information for all of these cells. The dragonfly watches a complicated time history of a stimulus trajectory and extracts features from this that are useful in catching a target. This is a difficult task, requiring a high level of speed and precision, and there is no a priori reason for believing that the important features of the trajectory should be summarizable in terms of independent receptive fields.

That the total information is slightly larger than the information in the independent receptive field model indicates that there are additional features of the stimulus that influence spiking that are not described by the model, but as this difference is small, we know that these features are not very significant in determining the single-spike response. Even though small, it is worth speculating about features of the trajectories that are not captured by the independent receptive field model. A likely candidate is that the TSDNs might be responding to features of the joint position and velocity space $(\vec{x}, \vec{v})$ that cannot be described in the independent position and velocity spaces. For example, selectivity for target trajectories that are curved, or where velocity preferences otherwise vary with position, would require a description in the joint space. There are other features of the target trajectory to which the cells could be responding: for example, linear acceleration or curvature. The possibilities listed to this point are still within the framework that the TSDNs are essentially responding to a stimulus event that occurred at a snapshot in time. Instead, the dragonfly could be extracting much more complex features from the short- or long-term history of the stimulus: for example, that the target visits positions in a particular sequence or that the stimulus affects the TSDNs with several latencies. We emphasize that I_t is completely independent of such assumptions about latencies and reflects the total information regardless of whether a single latency or multiple nonuniform latencies are active in the system. This is an important point: in most neural coding models, one of the most severe dimensionality reductions is that the neuron only responds to some finite history of the stimulus, and a proper test for completeness of a model must address whether the correct amount of stimulus history is accounted for. Therefore, although more complex history dependence or latencies could be relevant to interpreting the response, these features can have no more than 0.5 to 1.5 bits/spike. Given all of these possibilities, many of which could be useful in catching a flying target, it is surprising that the response of these high-level cells can be so simply described.

On the other hand, the simplicity of this model should not be overstated. Because it contains both position and velocity information, there is an implied dynamic. Examples of this can be seen in the evolving receptive field structure in Figure 4 and also in that several of the cells seem to extrapolate or predict target position (Results). As another example, we note that this model is at least comparable in complexity to a common model for the receptive fields of retinal ganglion cells that consist of a two-dimensional receptive field followed by temporal filter (Rodieck, 1965), i.e., two independent filters, as with the dragonfly. In analogy to the ganglion cell temporal filter, it is worth noting that the velocity receptive fields of the TSDNs force a temporal component to the response, in that all of the cells spike only when the target moves quickly through the receptive field; for this reason, most of the TSDNs spike in short, well-defined bursts, and only to the transient motion of the target through the position receptive fields. Finally, we note from Figure 6C that the TSDNs carry approximately 4 bits per spike, which is large compared to other central neurons (Rieke et al., 1997; Buračas and Albright,

1999). We can gain an intuition for this by imagining that all 4 bits referred to position, in which case 3 independent spikes would give enough information to localize a target to one part in 2^{12} or, for example, to reduce an attack angle uncertainty of $90^\circ \times 90^\circ$ to $1.5^\circ \times 1.5^\circ$. It does not seem unreasonable that such cells could provide the speed and precision that is required to catch flying prey.

In this work, we have both introduced a theoretical framework and applied this framework to the analysis of particular experiments on the dragonfly TSDNs. We emphasize that the framework is very general, and that even within the dragonfly system we can use the same information theoretic tools to ask questions that lead toward a more complete understanding of neural coding under natural stimulus conditions. At the same time, we also note that all of the analysis presented here should be interpreted within the context of the particular stimulus ensemble that we used; as with all characterizations, we can only draw conclusions within the stimulus ensemble used in our study. Therefore, it would be interesting to repeat the approach used here but with different stimuli. For example, how applicable is the receptive field model when there are multiple targets present, or when the target moves against a complex background? How does the receptive field model deal with changes in features such as target size, contrast, or shape? All of these variations are feasible in a natural context. One particular point of interest is that the velocity receptive fields have latencies of approximately 40 ms. While we have used stimulus trajectories that are primarily smooth on this timescale, one can imagine flight paths for the dragonfly and prey that have multiple evasive actions and turns on this timescale. It would be interesting to provide stimuli with all timescales shorter than 40 ms and see how well the receptive field model continues to work. It could be, for example, that velocity receptive fields do not change, but continue to refer to the stimulus at a 40 ms delay; or it could be that they broaden, since it may take less time to determine a less informative quantity; or it could be that the position receptive fields become more informative as the velocity receptive fields become less informative. In addition to changing features in the stimulus, we could extend our analysis beyond single spikes to look at patterns of spikes, either across time or across the population of cells. Possibly, such patterns could stand for new features of the trajectory. All of these issues have a precise information theoretic formulation exactly parallel to the analysis introduced here and will be addressed in subsequent work.

We began the study of the TSDNs without a priori knowledge of the type of information that would be represented in the nervous system. Because responses to most features in the stimulus space would likely produce position and/or velocity receptive fields, we needed some indication of whether these receptive fields are appropriate descriptions of the responses of the TSDNs. We have found that single spikes generally contain slightly more information about velocity than position, and that together these account for most, but not all, of the total information. By comparing the information in the receptive field model to the total information in the spikes of the individual cells, we have been able to quantify to what extent the receptive fields capture the cells' selectivity and have found that 70% to 90% of the information provided by individual spikes can be accounted for by the simple receptive field model.

Experimental Procedures

Our approach in this study was to record concurrently from multiple TSDNs in *Aeshna canadensis* and to determine the responses of the TSDNs to a moving target. Concurrent TSDN activity was recorded using an electrode array.

Mounting and Dissection

The mounting and dissection of the dragonfly were similar to the procedures outlined by Olberg (1986), and here we only briefly summarize. The TSDNs project from the brain to the thoracic ganglia through the ventral nerve cord. The dragonfly was mounted ventral side up to a rigid metal bar, and a small patch of cuticle in the neck was removed to allow access to the fused cervical connectives by the electrode array, as shown in Figures 1A and 1B.

Recording

An array of extracellular electrodes was placed beneath the neck connective to record spike activity of the TSDNs as shown in Figure 1. As the TSDNs are situated dorsally in the ventral nerve cord, the array was placed against the dorsal surface of the cord. To accomplish this, we manufactured an electrode array with a special 3D geometry.

The array device consists of eight evenly spaced, side-by-side wires encased in an insulating and flexible polyimide film ($\sim$11 μm thick). A small window in the polyimide at one end of each wire allowed electrical contact with the connective, and a larger window at the other end allowed connection to the amplifiers. This flexible array was glued to a thin rigid support that was shaped to reach over the thorax of the dragonfly, down into the neck incision, and under the connective to rest against its dorsal side. The array geometry is shown in Figure 1 with an example set of recordings from the electrodes.

The electrodes were placed across the width of the connective. The electrodes were spaced 65 μm center-to-center, 10 μm wide along the array axis, and 50 μm long. The voltages from the electrodes were amplified using A-M Systems Model 1700 amplifiers (gain = 1 K, low-frequency cut-off = 300 Hz, high-frequency cut-off = 10 kHz), and the complete voltage versus time record for each channel was stored on a computer hard drive after being digitized to 12 bit precision at 40 kHz per channel (multiplexed) using a National Instruments PCI-MIO-16E-1 board.

Spike Sorting

Spikes thus recorded from the TSDNs ranged between 0.1 and 1.0 mV. Spike times were identified as the time of the lowest value below a threshold on all of the eight channels. The value of the minimum on each of the eight channels was determined, and other features of the spike shape were ignored in the subsequent analysis. The minima define an eight-dimensional vector for each spike, and examination of two-dimensional projections of this vector space show many distinct clusters. An agglomeration algorithm (Gordon, 1999) was used to identify cluster membership.

To verify the above spike-sorting techniques, in multiple preparations we have made simultaneous intracellular recordings of two TSDN cells concurrent with the extracellular recordings (although not for the specific animal used for this paper). In all cases we have found that clearly identifiable clusters corresponded to spikes from individual cells. On the other hand, not all spikes from an individual cell would lie within the cell's cluster. We suggest that the predominant cause for this is temporal overlapping of spike waveforms, and here we make no attempt to sort these overlap events. The consequence of these methods was that in any given recording and for spikes above a threshold value appropriate for the TSDNs, we found many well-defined individual clusters within a very low-density background of spikes that fall within no cluster.

Visual Stimuli

The visual stimulus consists of a single small dark spot moving against a light visual field (Olberg, 1986). The motion of the spot was computer controlled to give pseudo-random trajectories as illustrated in Figure 1C, which have the advantage that the TSDNs do not habituate nearly as quickly as they do to simpler trajectories. The trajectories were constructed as a series of 240 straight line segments, each segment having a 0.25 s duration, connected end-to-end with no pauses between the segments, giving a full trajectory of 1 min in duration. Each segment had a speed randomly chosen from a flat distribution with a 40 cm/s cut-off and a randomly chosen direction uniform over 0 to 2π. If the entire segment did not lie within the boundary of the screen, both a new speed and new direction were chosen until the segment was contained entirely within the screen. Thirty-one 1 min trajectories were presented to the animal with half-minute rests between each presentation, during which the target was kept stationary at the starting point of the subsequent trajectory. One of the 31 trajectories was selected to be shown repeatedly, interleaved between each of the remaining 30 random trajectories. A complete experiment took 1.5 hr.

The coverage of stimulus space by the random trajectories is shown in Figure 2. Note that in the (V_x, V_y) representation, the velocity distribution was strongly peaked near $V_x = V_y = 0$; this is because the stimuli were constructed from a flat distribution of speeds (speed $= \sqrt{V_x^2 + V_y^2}$), and each concentric annulus of constant speed had approximately the same number of segments in the stimulus, leaving the larger annuli with the same number of visits spread over a larger area.

The stimulus was displayed on a Tektronix 608 oscilloscope, rastered in the vertical direction at 500 Hz with 20 lines per cm. The 9×9 cm screen was placed 6 cm from the dragonfly, thus filling a $74^\circ \times 74^\circ$ field of view. The oscilloscope trace was dimmed where necessary (using in-house electronics) to create a square spot 0.42 cm on a side, seen as 4° by the dragonfly. The brightness of the screen was 35.4 μW/cm^2/sr, and the dimmed spot had a brightness of 2.3 μW/cm^2/sr. The screen was placed in front of the dragonfly so it could be seen by the fovea and adjusted to include most of the receptive fields of the TSDNs.

The construction of the target trajectories was chosen as a compromise between a stimulus that solved various practical considerations and one that the animal might encounter in nature. In the natural environment, dragonflies pursue targets along interception trajectories. While on an ideal interception course, the target will be at a stationary position from the perspective of the dragonfly (Olberg et al., 2000), and from this perspective the position and velocity of a moving target gives information about the correction that is needed to return to an interception course. We think of our stimulus as a series of segments corresponding to such slip trajectories. We have further simplified these segments by keeping the velocity constant within each segment and making them all have the same duration, chosen to be substantially longer than the dragonfly's response delay times. Many other ideas for the stimulus construction, such as random flicker or use of a spot that flashes at random locations on the screen, were ruled out because they did not produce responses from the TSDNs.

Throughout this paper we represent the stimulus in the coordinates of the display screen. This is chosen over representation of dragonfly viewing angles. In the viewing angle coordinate system, target trajectories that are straight lines on the display screen will appear as arcs, thus making it difficult for the reader to determine whether or not an arc originated from a straight or curved trajectory. For the display used here, the horizontal and vertical extrema of the screen are at ± 4.5 cm ($\pm 37^\circ$).

Receptive Fields

In this paper we quantify a receptive field as the probability of the neuron spiking as a function of a stimulus parameter, P (spike|stimulus), where the stimulus parameter is either the vector position, $P(\text{spike}|\vec{x})$, or vector velocity, $P(\text{spike}|\vec{v})$, of the target. Practically, we evaluate the receptive field by binning the stimulus parameter and computing a histogram of those bins that are traversed at some fixed time relative to the spike, and we vary the time between the spike and our measurement of the receptive field to determine the optimal time delay. All receptive fields are calculated using the 30 nonrepeated trials. Example position and velocity receptive fields of two cells are shown in Figure 3. Alternatively, we could have quantified the receptive field as the conditional distribution of the stimulus given a spike, P(stimulus|spike). These descriptions are related through Bayes' rule,

$$P(\text{spike}|\text{stimulus}) = P(\text{stimulus}|\text{spike})P(\text{spike})/P(\text{stimulus}), \quad (2)$$

where P(spike) is the overall spike probability, or the firing rate, and P(stimulus) is the prior probability distribution of the stimulus presented in the experiment.

The information in the receptive field is a measure of the reduction in uncertainty of the target position or velocity upon observation of a single spike. Without knowledge of the spike arrival times, the target's position and velocity are known to be distributed according to the prior distributions, $P(\vec{x})$ and $P(\vec{v})$ (Figure 2). On occurrence of a spike from a particular cell at a given time, the target's position and velocity will be described by the probability distributions $P(\vec{x}|\text{spike})$ and $P(\vec{v}|\text{spike})$ derived from the receptive fields of the cell. If either of these distributions is more concentrated than the prior distribution, the uncertainty in the position and/or velocity has been reduced. The information content of a receptive field quantifies this reduction in the uncertainty, and, for position, is determined from (see Appendix)

$$I(\vec{x};\,\text{spike}) = \int d^2x P(\vec{x}|\text{spike})\log_2\left[\frac{P(\vec{x}|\text{spike})}{P(\vec{x})}\right], \quad (3)$$

and the corresponding expression for velocity defines $I(\vec{v};\,\text{spike})$.

Accounting for Correlations in the Stimulus

Correlations between position and velocity in the prior stimulus distribution can remove our ability to distinguish between position and velocity sensitivity in the dragonfly's response. To check the influence of correlations in the stimulus, we compare the sum of the information values of the position and velocity receptive fields to the information in the joint distribution, using the identity

$$I(\vec{x};\,\text{spike}) + I(\vec{v};\,\text{spike}) = I(\vec{x},\vec{v};\,\text{spike}) - I(\vec{x},\vec{v}|\text{spike}) + I(\vec{x};\vec{v}), \quad (4)$$

where $I(\vec{x};\vec{v})$ is the mutual information between position and velocity in the entire stimulus ensemble, and $I(\vec{x};\vec{v}|\text{spike})$ is the mutual information between position and velocity for the subset of stimuli that leads to a spike. This equation relates information in the independent model (left side) to terms that account for correlations in $\vec{x}$ and $\vec{v}$ in both the stimulus and the dragonfly's response (right side). $I(\vec{x},\vec{v};\,\text{spike})$ is of interest here since it accounts for the response of the dragonfly to the joint $(\vec{x},\vec{v})$ stimulus space. Unfortunately, we cannot measure $I(\vec{x},\vec{v};\,\text{spike})$, as the joint $(\vec{x},\vec{v})$ space is too large to sample thoroughly, and for this reason we have proposed the independent model that we study in this paper. Since information is always positive or zero, if $I(\vec{x};\vec{v}) = 0$, we would have $I(\vec{x};\,\text{spike}) + I(\vec{v};\,\text{spike}) \leq I(\vec{x},\vec{v};\,\text{spike})$, which is reasonable, since the independent model should contain less information than the full joint response. But if $I(\vec{x};\vec{v}) \neq 0$, then we cannot make this comparison to $I(\vec{x},\vec{v};\,\text{spike})$, and in general the addition of $I(\vec{x};\,\text{spike})$ and $I(\vec{v};\,\text{spike})$ will not be meaningful. Therefore, for $I(\vec{x};\,\text{spike}) + I(\vec{v};\,\text{spike})$ to be useful, $I(\vec{x};\vec{v})$ must be small. For the stimulus used here, we calculate $I(\vec{x};\vec{v}) < 0.25$ bits, significantly smaller than $I(\vec{x};\,\text{spike})$ or $I(\vec{v};\,\text{spike})$ and not much larger than the error bars for most of these quantities.

Since $I(\vec{x};\vec{v})$ is small, we can treat this as an error term to the more interesting but approximate relation

$$I(\vec{x};\,\text{spike}) + I(\vec{v};\,\text{spike}) \approx I(\vec{x},\vec{v};\,\text{spike}) - I(\vec{x},\vec{v}|\text{spike}). \quad (5)$$

Here, all of the terms refer to the TSDN responses to different features of the stimulus without influence of the correlations in the stimulus. It is also emphasized that $I(\vec{x};\vec{v}|\text{spike})$ describes the amount of information spikes contain about correlated features in the stimulus (*not* intrinsic correlations in the stimulus). Furthermore, to within the small value of $I(\vec{x};\vec{v})$, we have the approximate bound

$$I(\vec{x};\,\text{spike}) + I(\vec{v};\,\text{spike}) \lesssim I(\vec{x},\vec{v};\,\text{spike}). \quad (6)$$

Here we have shown that to within the small value of $I(\vec{x}; \vec{v})$, it is meaningful to treat $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$ as the information in the independent receptive field model. Whether this is a good description of the response of the dragonfly (i.e., not only is $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike}) \leq I(\vec{x}, \vec{v}; \text{spike})$, but $I(\vec{x}, \vec{v}; \text{spike})$ is fairly close to $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$) is addressed by comparing $I(\vec{x}; \text{spike}) + I(\vec{v}; \text{spike})$ to the total information, which is done in the Results.

A related question is whether correlations between position and velocity in the stimulus might cause errors in our interpretation of the individual receptive fields. For example, consider a stimulus ensemble with built-in correlations such that the target always has a particular velocity, $\vec{v}_0$, when at a particular position, $\vec{x}_0$. In this case, a cell that had $\vec{x}_0$ prominent in its position receptive field will necessarily have $\vec{v}_0$ appear in its velocity receptive field even if the cell had no intrinsic preference for $\vec{v}_0$. To test the significance of this effect for the actual stimulus ensemble used in these experiments, we can estimate the influence of the position velocity correlations by calculating the information in a "velocity correlation receptive field" calculated purely from the actual position receptive field. That is, from each cell's position receptive field, we calculate the receptive field in velocity that we would measure solely induced by the correlations in the stimulus. This is given by

$$P_{\text{corr}}(\text{spike}|\vec{v}) = \int d^2x P(\text{spike}|\vec{x}) P(\vec{x}|\vec{v}). \tag{7}$$

Calculating this value for all cells, we find that the information content of the correlation induced receptive field is less than 0.1 bits. This indicates that correlations between position and velocity do not produce significant correlations between the position and velocity receptive fields calculated in Equation 3, and that these correlations can be ignored in the current analysis.

Total Information

The stimulus-independent information, which we refer to as the total information, is calculated using the following equation (see also Brenner et al., 2000, and Appendix Equation 16)

$$I_t = \frac{1}{T} \int_0^T dt \, \frac{r(t)}{\bar{r}} \log_2 \left[\frac{r(t)}{\bar{r}} \right] \text{ bits/spike}, \tag{8}$$

where I_t is the average information carried by the arrival time of a single spike, $r(t)$ is the average instantaneous rate calculated from the response to the 30 repeated trajectories, ($\bar{r}$) is the average value of $r(t)$, and T is the length of the trial.

Finite Size Effects in Information Calculations

Our analysis of the data involves the computation of information from various probability distributions. These distributions refer both to the arrival times of spikes (see Equation 8) and to the position and velocity vectors describing the stimulus (see Equation 3). Here we estimate the distributions by uniformly discretizing these continuous variables and simply counting the number of times each bin is occupied. With finite data sets, all information theoretic quantities will depend on the number of samples or trials in the experiment and on the bin size of our discretization. Here we take a conservative approach to these problems. The dependence of information on sample size N is examined explicitly (as in Strong et al., 1998; Golomb et al., 1997), and for all reported results we see the expected behavior

$$I(N) = I_{\text{inf}} + A/N; \tag{9}$$

I_{inf} is then our estimate. The success of this correction can be seen in Figure 5, where I_{inf} is close to zero at times far from the occurrence of the spike. Without this correction, all of the information values in this figure would be about 0.5 bits larger, and in particular, at times far from the spike time where nonzero information would be physically unreasonable.

This estimate of information is studied as a function of the discretization scale, and we look for a plateau—overly coarse discretization leads to information loss, while overly fine discretization leads to artifactually large estimates because of undersampling. For these experiments, we find that discretizing the target's position and velocity to 12 × 12 bins and time to 12 ms lies within the plateau for all cells. Errors could arise from statistical noise in the sample, from failure to use a stimulus that is long enough to sample the distribution of stimuli ergodically, or from difficulties in extracting I_{inf}; all of these have been discussed elsewhere (Brenner et al., 2000). We have examined these errors for the present data and find that the dominant source of error is in the expected plateau in information versus discretization scale not completely leveling off, and we report these errors conservatively. It is possible that smaller error bars could be derived from a more sophisticated estimation procedure.

Appendix

In the interest of completeness, we derive here all of the equations needed for the analysis presented in the text. To a large extent our discussion follows that of Brenner et al. (2000).

Imagine that we observe neural responses in a large time window $0 < t < T$. During this time, a complete description of the stimulus is given by the entire trajectory $\tilde{x}(t)$. If we are told the arrival time t_0 of a single spike, then on average this provides an information about the stimulus that is given by

$$I(\{\tilde{x}(t)\}; t_0) = \int \{d\tilde{x}(t)\} \int_0^T dt_0 P[\{\tilde{x}(t)\}, t_0] \log_2 \left[\frac{P[\{\tilde{x}(t)\}, t_0]}{P[\{\tilde{x}(t)\}] P(t_0)} \right] \text{ bits}, \tag{10}$$

where we write $\{\tilde{x}(t)\}$ to remind us that, in principle, the occurrence of a single spike can provide information about *any* aspect of the trajectory; to compute the full information we need to integrate over the space of functions $\{\tilde{x}(t)\}$ as opposed to the space of positions $\tilde{x}(t)$ at one particular time.

We can decompose Equation 10 in several ways. We begin by writing

$$P[\{\tilde{x}(t)\}, t_0] = P[t_0|\{\tilde{x}(t)\}] P[\{\tilde{x}(t)\}]. \tag{11}$$

But the probability $P[t_0|\{\tilde{x}(t)\}]$ of a spike at t_0, given that we know the full stimulus $\{\tilde{x}(t)\}$, must be proportional to the time-dependent firing rate $r(t_0)$ that we would measure in a conventional peristimulus time histogram as in Figure 1E. Because we are considering the probability distribution for exactly one spike, we must obey the normalization conditions

$$\int_0^T dt_0 P[\{\tilde{x}(t)\}, t_0] = \int_0^T dt_0 P(t_0) = 1, \tag{12}$$

which means that

$$P[\{\tilde{x}(t)\}, t_0] = \frac{1}{T} \frac{r(t_0)}{\bar{r}} \tag{13}$$

$$P(t_0) = \frac{1}{T}, \tag{14}$$

where $\bar{r}$ is the average spike rate over the whole window $0 < t_0 < T$. Thus, we can substitute into Equation 10 to obtain

$$I(\{\tilde{x}(t)\}; t_0) = \int \{d\tilde{x}(t)\} P[\{\tilde{x}(t)\}] \frac{1}{T} \int_0^T dt_0 \frac{r(t_0)}{\bar{r}} \log_2 \left[\frac{r(t_0)}{\bar{r}} \right], \tag{15}$$

where $P(t_0) = \bar{r}$ is the average spike rate over the whole window $0 < t_0 < T$. Now we note that if the window T is sufficiently large, the integral over this time will be independent of the precise choice of trajectory, since averaging over time already is equivalent to averaging over an ensemble of signals (ergodicity). Thus, we can drop the integral over trajectories and obtain more simply

$$I(\{\tilde{x}(t)\}; t_0) = \frac{1}{T} \int_0^T dt_0 \frac{r(t_0)}{\bar{r}} \log_2 \left[\frac{r(t_0)}{\bar{r}} \right]. \tag{16}$$

This is Equation 8 of the text and Equation 2.5 of Brenner et al. (2000). We emphasize that this provides an estimate of the information conveyed by the arrival time of a single spike without any assumptions regarding which features of the stimulus are being encoded. Note also that this is an exact formula for the information carried by a single spike, even if successive spikes do not carry independent information.

As an alternative, we can decompose Equation 10 by writing

$$P[\{\tilde{x}(t)\}, t_0] = P[\{\tilde{x}(t)\}|t_0]P(t_0). \quad (17)$$

Substituting, we find

$$I(\{\tilde{x}(t)\}; t_0) = \int_0^T dt_0 P(t_0) \int \{d\tilde{x}(t)\} P[\{\tilde{x}(t)\}|t_0] \log_2 \left(\frac{P[\{\tilde{x}(t)\}|t_0]}{P[\{\tilde{x}(t)\}]} \right). \quad (18)$$

A crucial point is that we expect the arrival of a spike at t_0 to tell us something about the trajectory in the neighborhood of the time t_0, but that what we are being told is invariant to this time. More precisely, if we compute the integral

$$\int \{d\tilde{x}(t)\} P[\{\tilde{x}(t)\}|t_0] \log_2 \left(\frac{P[\{\tilde{x}(t)\}|t_0]}{P[\{\tilde{x}(t)\}]} \right),$$

we will find that for different values of the spike time t_0, different portions of the trajectory $\tilde{x}(t)$ are most important in the integrand, but the final integral is independent of t_0. Thus, we can write

$$I(\{\tilde{x}(t)\}; t_0) = \int \{d\tilde{x}(t)\} P[\{\tilde{x}(t)\}|t_0] \log_2 \left(\frac{P[\{\tilde{x}(t)\}|t_0]}{P[\{\tilde{x}(t)\}]} \right), \quad (19)$$

which is Equation A.11 of Brenner et al. (2000) and Equation 2 of de Ruyter van Steveninck and Bialek (1988).

In general, Equation 19 is not so useful, since it expresses the information carried by single spikes as an integral over the full space of trajectories $\{\tilde{x}(t)\}$. This is why the mathematically equivalent integral over time-dependent rates, Equation 16, is so important in providing a practical model-independent measure of information transmission. On the other hand, if we suspect that the neuron is selective for only some features $\mathbf{f} \equiv \{f_1, f_2, \ldots, f_K\}$ of the trajectory, then there is an analogous equation for the information that the spike arrival time provides about these features,

$$I(\mathbf{f}; t_0) = \int d^K f \, P(\mathbf{f}|t_0) \log_2 \left[\frac{P(\mathbf{f}|t_0)}{P(\mathbf{f})} \right]. \quad (20)$$

With $\mathbf{f}$ identified as the position $\tilde{x}(t_0 - \tau)$ at some time τ relative to the spike, this corresponds to Equation 3 in the text.

If the features $\mathbf{f}$ are computed from the trajectory $\{\tilde{x}(t)\}$, then in this computation—which represents a projection from the (very) high-dimensional space of trajectories onto the limited K-dimensional space of features—we can only lose information or at best preserve the information that is there. Thus, we have

$$I(\mathbf{f}; t_0) \leq I(\{\tilde{x}(t)\}; t_0), \quad (21)$$

with equality *only* if the neuron's response is determined entirely by the features $\mathbf{f}$ and by no other aspects of the stimulus. This inequality allows us to ask in a rigorous way whether a limited set of features (target position and velocity for the neurons discussed here) capture the information that the neuron provides about the sensory input.

Acknowledgments

We are grateful for the help and advice of Blaise Agüera y Arcas, Naama Brenner, Adrienne Fairhall, Thomas Oliver, and Robert de Ruyter van Steveninck. We would like to thank Doug Shire and the Cornell NanoScale Science & Technology Facility for manufacture of the flexible electrode arrays. T.L.A. would like to thank Ron Hoy for support during the initial stages of this project. This work was primarily supported by the NEC Research Institute; portions of this manuscript were written while T.L.A. was working with David Tank at Princeton University.

Received: May 2, 2003
Revised: August 4, 2003
Accepted: September 23, 2003
Published: November 12, 2003

References

Agüera y Arcas, B., Fairhall, A.L., and Bialek, W. (2003). Computation in a single neuron: Hodgkin and Huxley revisited. Neural Comput. *15*, 1715–1749.

Allman, J., Meizen, F., and McGuiness, E. (1985). Stimulus specific response from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. Annu. Rev. Neurosci. *8*, 407–430.

Armstrong-James, M., and Fox, K. (1987). Spatiotemporal convergence and divergence in the rat S1 "barrel" cortex. J. Comp. Neurol. *263*, 265–281.

Barlow, H.B. (1953). Summation and inhibition in the frog's retina. J. Physiol. *119*, 69–88.

Brenner, N., Strong, S.P., Koberle, R., Bialek, W., and de Ruyter van Steveninck, R.R. (2000). Synergy in a neural code. Neural Comput. *12*, 1531–1552.

Buračas, G.T., and Albright, T.D. (1999). Gauging sensory representations in the brain. Trends Neurosci. *22*, 303–309.

Corbet, P.S. (1999). Dragonflies: Behavior and Ecology of *Odonata* (Ithaca, NY: Comstock).

de Ruyter van Steveninck, R.R., and Bialek, W. (1988). Real-time performance of a movement sensitive neuron in the blowfly visual system. Proc. R. Soc. Lond. B *234*, 379–414.

Frye, M.A., and Olberg, R.M. (1995). Visual receptive-field properties of feature detecting neurons in the dragonfly. J. Comp. Physiol. [A] *177*, 569–576.

Golomb, D., Hertz, J., Panzeri, S., Treves, A., and Richmond, B. (1997). How well can we estimate the information carried in neuronal responses from limited samples? Neural Comput. *9*, 649–665.

Gordon, A.D. (1999). Classification (Boca Raton, FL: Chapman & Hall/CRC).

Hartline, H.K., and Ratliff, F. (1958). Spatial summation of inhibitory influences in the eye of Limulus and the mutual interaction of receptor units. J. Gen. Physiol. *41*, 1049–1066.

Hubel, D.N., and Wiesel, T.N. (1962). Receptive fields, binocular interactions and functional architecture in the cat's visual cortex. J. Physiol. *160*, 106–154.

Kolton, L., and Camhi, J.M. (1995). Cartesian representation of stimulus direction: parallel processing by two sets of giant interneurons in the cockroach. J. Comp. Physiol. [A] *176*, 691–702.

Kuffler, S.W. (1953). Discharge patterns and functional organisation of mammalian retina. J. Neurophysiol. *16*, 37–68.

O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. Brain Res. *34*, 171–175.

Olberg, R.M. (1986). Identified target-selective visual interneurons descending from the dragonfly brain. J. Comp. Physiol. [A] *159*, 827–840.

Olberg, R.M., Worthington, A.H., and Venator, K.R. (2000). Prey pursuit and interception in dragonflies. J. Comp. Physiol. [A] *186*, 155–162.

Rieke, F., Warland, D., de Ruyter van Steveninck, R.R., and Bialek, W. (1997). Spikes: Exploring the Neural Code (Cambridge, MA: MIT Press).

Rodieck, R.W. (1965). Quantitative analysis of cat retinal ganglion cell response to visual stimuli. Vision Res. *5*, 583–601.

Shannon, C.E. (1948). A mathematical theory of communication. Bell System Tech. J. *27*, 379–423 and 623–656.

Sharpee, T., Rust, N.C., and Bialek, W. (2003). Maximally informative dimensions: analyzing neural responses to natural signals. To appear in Advances in Neural Information Processing *15*, S. Becker, S. Thrun, and K. Obermayer, eds. (Cambridge, MA: MIT Press), in press.

Strong, S.P., Koberle, R., de Ruyter van Steveninck, R.R., and Bialek, W. (1998). Entropy and Information in Neural Spike Trains. Phys. Rev. Lett. *80*, 197–200.

Theunissen, F., Sen, K., and Doupe, A. (2000). Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. J. Neurosci. *20*, 2315–2331.