High Accuracy Decoding of Dynamical Motion from a Large Retinal Population

Olivier Marre, Vicente Botella-Soler, Kristina D. Simmons, Thierry Mora,
Gašper Tkačik, Michael J. Berry Ii

High Accuracy Decoding of Dynamical Motion from a Large Retinal Population

Olivier Marre[1,2]*, Vicente Botella-Soler[3], Kristina D. Simmons[4], Thierry Mora[5], Gašper Tkačik[3], Michael J. Berry II[1]

1 Department of Molecular Biology and Neuroscience Institute, Princeton University, Princeton, United States of America, 2 Institut de la Vision, INSERM UMRS 968, UPMC UM 80, CNRS UMR 7210, Paris, France, 3 Institute of Science and Technology Austria, Klosterneuburg, Austria, 4 Department of Psychology, University of Pennsylvania, Philadelphia, United States of America, 5 Laboratoire de Physique Statistique, École Normale Supérieure, CNRS and UPMC, Paris, France

* olivier.marre@inserm.fr

Abstract

Motion tracking is a challenge the visual system has to solve by reading out the retinal population. It is still unclear how the information from different neurons can be combined together to estimate the position of an object. Here we recorded a large population of ganglion cells in a dense patch of salamander and guinea pig retinas while displaying a bar moving diffusively. We show that the bar's position can be reconstructed from retinal activity with a precision in the hyperacuity regime using a linear decoder acting on 100+ cells. We then took advantage of this unprecedented precision to explore the spatial structure of the retina's population code. The classical view would have suggested that the firing rates of the cells form a moving hill of activity tracking the bar's position. Instead, we found that most ganglion cells in the salamander fired sparsely and idiosyncratically, so that their neural image did not track the bar. Furthermore, ganglion cell activity spanned an area much larger than predicted by their receptive fields, with cells coding for motion far in their surround. As a result, population redundancy was high, and we could find multiple, disjoint subsets of neurons that encoded the trajectory with high precision. This organization allows for diverse collections of ganglion cells to represent high-accuracy motion information in a form easily read out by downstream neural circuits.

Author Summary

It remains unclear how the brain is able to track the location of moving objects by reading the spike trains received from the retina. To address this question, we recorded a large population of ganglion cells, the retinal output, in a dense patch of salamander and guinea pig retinas while displaying a bar moving in complex motion. From previous studies, the naive expectation was that individual ganglion cells would spike when an object was moving on their receptive field center and that the entire population's activity would resemble a "hill" that continuously tracked the object's location. However, our analysis revealed that this picture did not hold. Instead, ganglion cells fired sparsely and coded for the bar

trajectory even when it was far from their receptive field center. Nevertheless, we showed that the bar's position could be reconstructed from retinal activity with an accuracy better than the spacing between photoreceptors, when using more than 100 cells. We also showed that the retinal code was highly redundant, over-representing the same information more than 6-fold. Yet, this unexpected representation allowed for precise object tracking using a simple decoder, as long as the temporal structure of the spike trains was accounted for.

Introduction

Our current understanding of how sensory neurons collectively encode information about the environment is limited. Being able to "read out" this information from neural ensemble activity is a major challenge in neuroscience. Discriminating among a discrete set of stimuli based on their sensory responses has been attempted in brain areas including the retina [1–3], sensory cortex [4, 5], and motor cortex [6]. Some studies have attempted the more difficult task of reconstructing a time varying signal from neural activity [7–9]. Being able to reconstruct a dynamical stimulus from the neural activity would greatly improve our understanding of the neural code, but there have been only a few attempts in the visual system [10–12].

The retina is an ideal circuit in which to attempt to decode a dynamical stimulus because recordings from a large, diverse, complete population of ganglion cells—the retinal output—have recently become possible [13, 14]. In contrast to cortical recordings, it is possible to drastically reduce the proportion of "hidden variables" in the network, i.e. unrecorded neurons that could carry relevant information. Furthermore, since the retina encodes all visual information available to the brain, the decoding performance of a complete retinal population can be rigorously compared to behavioral performance for an equivalent task [15, 16].

Motion tracking is of major ecological relevance. Amphibians can capture small moving prey, like flies, at a variety of speeds and distances from their body [17–19], implying that the retina must track such motion accurately. Furthermore, humans can use fixational eye movements to discriminate stimuli separated by roughly two cone photoreceptors, which is highly challenging without an accurate retinal representation of image movement [20].

Yet a reconstruction of the position of an object moving randomly from the activity of sensory neurons has not been achieved in the vertebrate visual system. The classical view suggests that ganglion cells will signal the position of the bar when it is at the peak of their receptive field, making this reconstruction easy by tracking the peak firing rate in the retinal map [21, 22]. However, the non-linear computations performed by the retinal network seem to make this picture more complex than intially thought: for example, a synchronous peak in the activity could signal sudden changes of speed [23] or a sudden reversal of motion [24]. So it is unclear how a complex trajectory can be decoded from the retinal output, and which neurons are most useful for that purpose.

In order to study how neural populations encode dynamic motion, we recorded from a large fraction of the ganglion cells in a patch of the retina while stimulating with a randomly moving dark bar. We show that the entire trajectory of the bar can be estimated by a linear decoder with a precision better than the average spacing between cones using groups of 100+ ganglion cells. Our analysis also revealed an unexpected structure to the population code. Instead of representing the object's location with a "hill" of neural activity, the population activity was sparse and broadly distributed in space. The retina thus used a distributed and redundant code,

in which several disjoint groups could be read-out to reconstruct the stimulus with high precision.

Results

Population recording and linear decoding

We used a large multi-electrode array with 252 electrodes to record the responses of ganglion cells in the salamander and guinea pig retinas (see Methods), while presenting a randomly moving bar (Fig 1A, 1B). The density of the electrodes allowed us to record a large fraction (at least 50%) of the ganglion cells in the retina patch covered by the array [14]. The neurons were recorded over a compact region and had highly overlapping receptive fields. Up to 189 cells were recorded simultaneously (see example spike rasters in Fig 1C) over several hours. The stimulus was a dark bar (width = 100 μm) on a grey background moving diffusively over the photoreceptor layer (with a diffusion constant $D = 62.5 mm^2 s^{-1}$). The trajectory was a random walk with a restoring force to keep the bar close to the array (Fig 1A, and Methods), spanning a region corresponding to roughly 10 degrees of visual angle.

One of our primary motivations in choosing this stimulus ensemble was to select patterns of motion that are more complex than constant velocity punctuated by sudden discontinuities. Many previous studies of motion processing in the retina have used objects moving at a constant velocity, sometimes punctuated by sudden changes of velocity [22–25]. While such
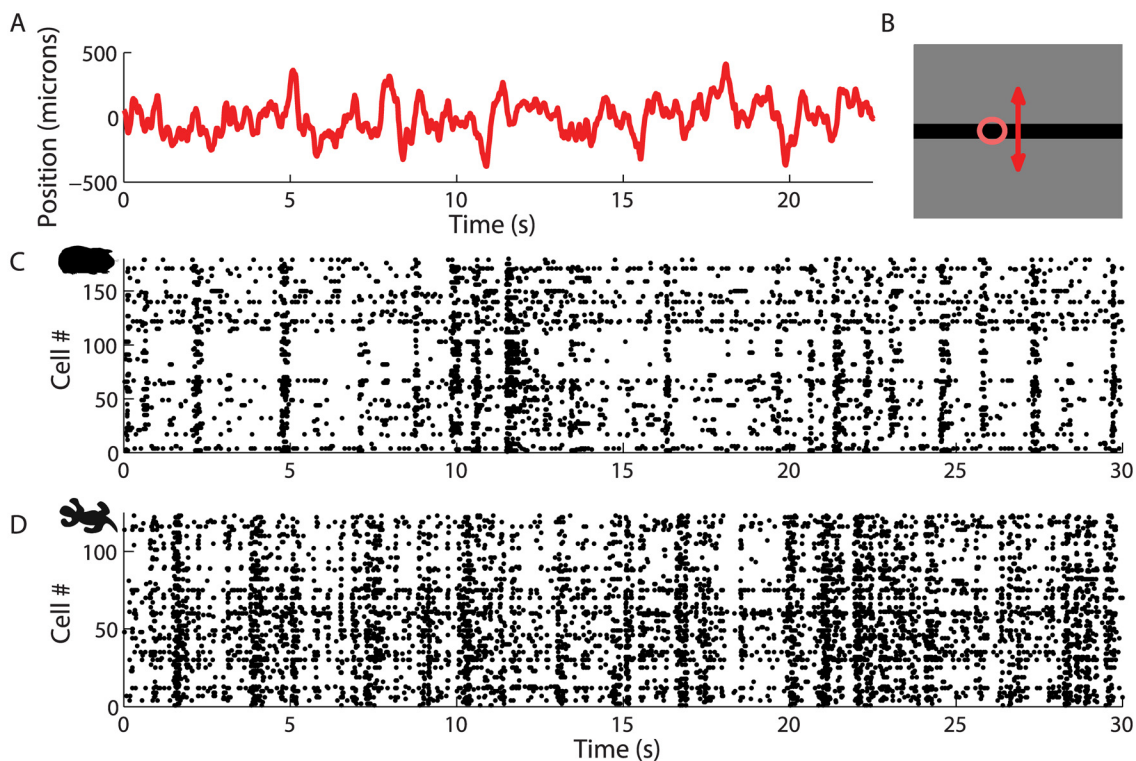
Fig 1. Ganglion cell spike trains during random bar motion. A: Position of the bar as a function of time. **B:** Example of one stimulus frame; motion is perpendicular to the bar (red arrow). Ellipse fitted to the spatial receptive field profile of one representative ganglion cell (pink). **C:** Spiking activity of 180 cells in the guinea-pig retina in response to a bar moving randomly with the trajectory shown in A. Each line corresponds to one cell, and the points represent spikes. The order of the cells along the y-axis is arbitrary. **D:** Spiking activity of 123 cells in the salamander retina responding to a bar moving randomly. Same convention as **C**.

doi:10.1371/journal.pcbi.1004304.g001

studies have contributed greatly to our understanding of retinal motion processing, they are not necessarily representative of the patterns of motion that an animal encounters in the natural environment. Obviously, not all objects in the world move at nearly constant velocity for long periods. In addition, we suspected that ganglion cell responses to complex motion might not be simply related to the manner in which they respond to constant velocity motion. Thus, we wanted a more general class of motion patterns. Building on the analogy to white noise stimuli, we selected a class of diffusive motion, which constitutes a broad ensemble of motion patterns without making highly specific choices about the set of trajectories.

Under these stimulus conditions, the activity of ganglion cells was sparse and strongly modulated (Fig 1C and 1D). Our goal was to decode the position of the bar at all times from this retinal activity. We used a linear decoder that took as an input the spike trains of all the ganglion cells and predicted the position of the bar as a function of time [26]. A temporal filter is associated with each cell and added to the prediction each time the cell fires an action potential (Fig 2A). The filter shapes were fitted on a training fraction of the data (2/3) to minimize the squared error between the predicted and true trajectories (see Methods for details). To test the decoder, the prediction was evaluated on the previously withheld test fraction of the data (1/3).

With 140 cells recorded in a single experiment in the guinea pig retina, the prediction was very precise and followed closely the position of the bar (Fig 2B). The normalized cross-correlation (see Methods) between the two traces was CC = 0.95 (corresponding to an error of $\sim$ 5μm, or 0.15 degree of visual angle) on the testing and the training dataset, which also indicated that the effects of over-fitting were minimal. Similar results were found in the salamander retina (with CC = 0.9 in the testing set, CC = 0.93 on the training dataset, Fig 2C). They were confirmed in n = 4 retinas for the salamander, and n = 2 for the guinea pig. Similar results were obtained when forcing the filters to be causal, by restricting their parameters to the range of -500 ms to 0 ms before a spike (Fig 2E; CC = 0.9). Thus, the activity of a large population of the ganglion cells contained enough information to reconstruct the position of an object moving randomly across the visual field with high precision. Furthermore, a simple readout mechanism—the linear decoder—was able to perform this reconstruction.

How large a population of cells is necessary to decode the position of the bar so precisely? To address this question, we selected random subsets of cells, performed linear decoding on them, and quantified their decoding performance by measuring the cross-correlation between the real and the predicted trajectories. The decoding performance as a function of the number of selected ganglion cells grew rapidly for few cells and then slowed down for ≳70 cells in both the guinea pig and the salamander retina (Fig 2D, 2E). Therefore, recording a large number of ganglion cells in the same area of the retina was essential to be able to track the trajectory of a moving object with high accuracy.

Performance estimation

While the high cross-correlation between the real and predicted bar trajectories demonstrates high performance, it is instructive to make this comparison along other dimensions. The decoding performance is expected to depend significantly on the frequency of bar motion for many reasons. Ganglion cells are unlikely to be able to follow rapid motion (high frequencies), because of their temporal integration time; similarly, ganglion cells are unlikely to be able to follow extremely slow motion (low frequencies), because of adaptation. Alternatively, many ganglion cells have highly transient responses that might emphasize moderately high frequencies. Finally, a large neural population might be able to extract information at higher and lower frequencies than individual ganglion cells.
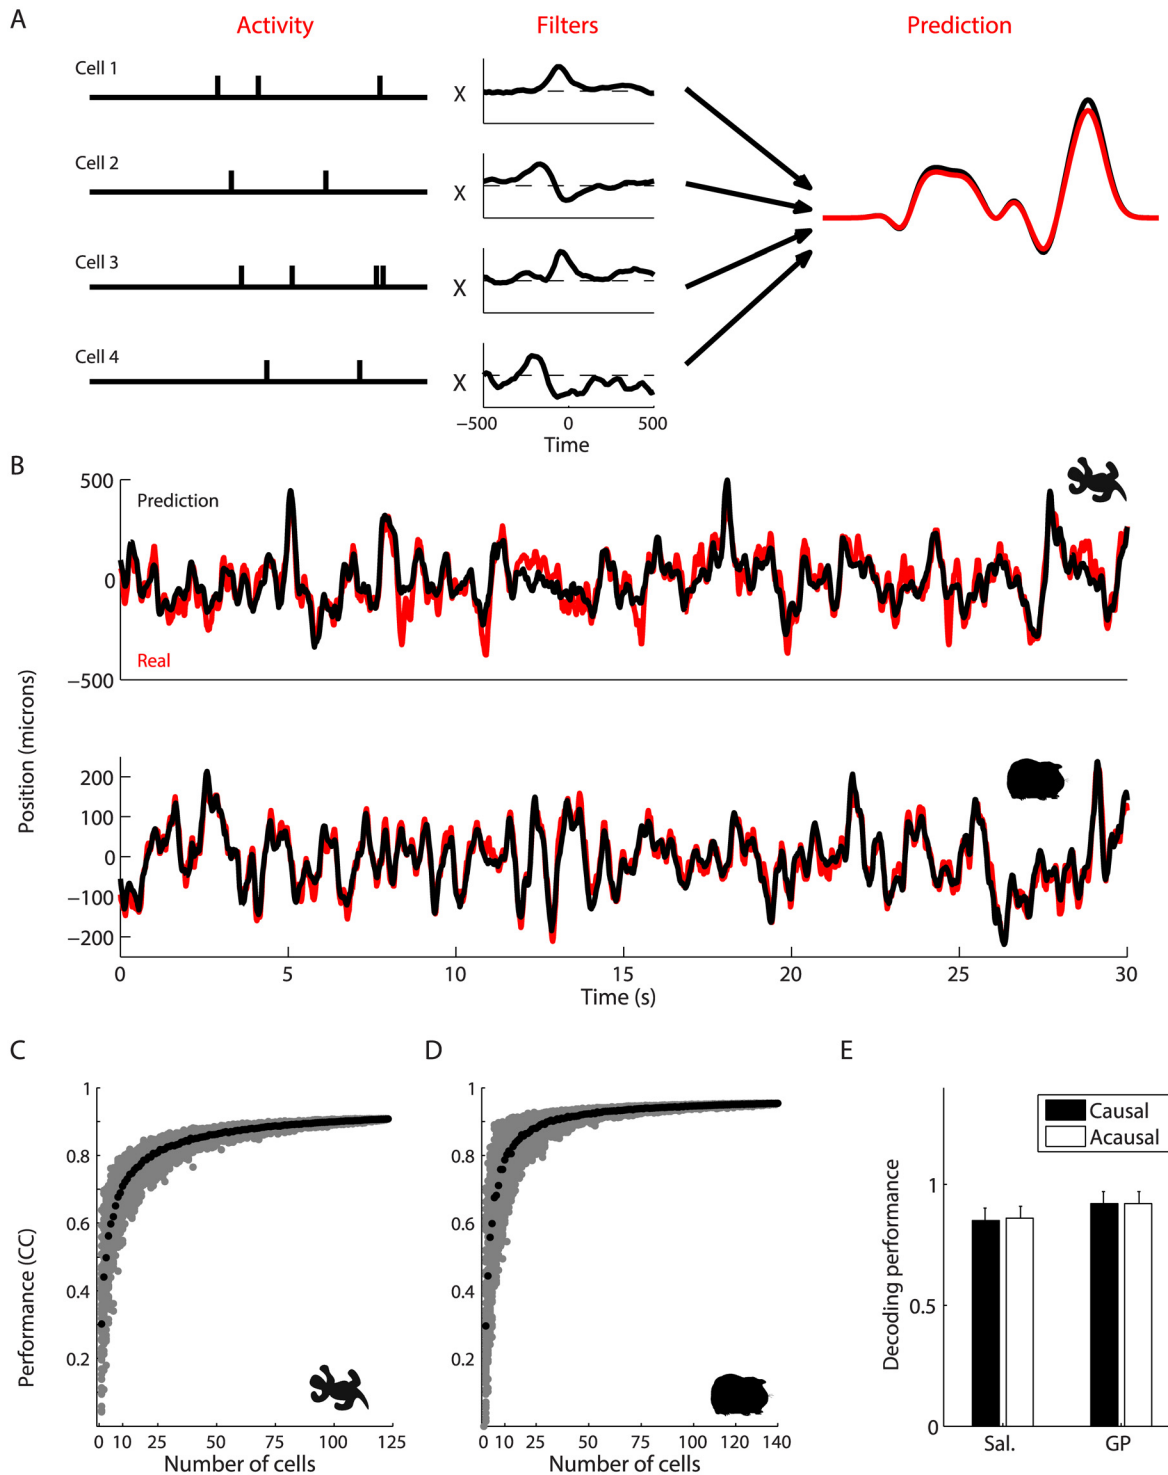
Fig 2. High-accuracy reconstruction of the bar's trajectory. A: Schematic of the linear decoding method, here for 4 cells. A temporal filter is associated with each cell. Each time the cell spikes, its filter is added to the ongoing reconstruction at the time of the spike. The filters are optimized on part of the data to have the lowest reconstruction error and then tested on the rest of the data. **B:** Top: prediction of the bar's position (black) from the activity of 123 cells in the salamander retina versus the real trajectory (red). Bottom: prediction of the bar's position (black) from the activity of 178 cells in the guinea pig retina versus the real trajectory (red). **C, D:** Decoding performance plotted against the number of cells in the salamander (C) and guinea pig (D). Gray points correspond to random subsets of cells, black to the average performance. **E:** Histogram of the average decoding performance across all experiments using either causal (black) or acausal (white) decoding filters; results shown for the entire recorded neural population in both the salamander (left) and guinea pig (right).

doi:10.1371/journal.pcbi.1004304.g002

Fig 3. The prediction error is in the hyperacuity range for salamander retina. A: Squared error as a function of frequency (red), compared to the power spectrum of the trajectory (blue). **B:** Root mean-squared error as a function 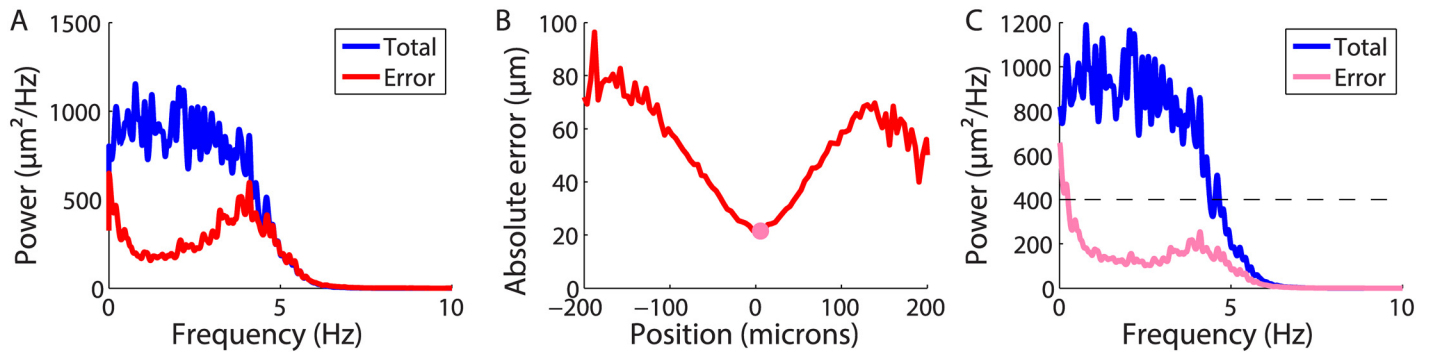of position (red) with the point of minimal error (pink circle). **C:** Error spectrum as a function of frequency (pink), for the position labeled in **B**. Power spectrum of the full trajectory (blue); error spectrum corresponding to the spacing between cone photoreceptors (dashed line).

The capacity of the retinal circuit to track the position of a moving object is also related to the geometry and signal-to-noise ratio of its sensors. We thus wanted to analyze more precisely the prediction errors and relate them, at least qualitatively, to the properties of the photoreceptor array. In some cases, an upper bound has been derived for the decoding performance that is related to the properties of the photoreceptor array (spacing between photoreceptors, transfer function of the photoreceptors, etc) [27].

To quantify the dependence of decoding performance on frequency, we generated a new bar trajectory: a white noise trajectory low-pass filtered at 5 Hz (see Methods). This stimulus ensemble is desirable because a large band of frequencies are represented equally (Fig 3A; blue). This allowed us to explore how the retina represents both slow and fast fluctuations in the bar's position. We recorded the responses of 158 neurons in the salamander retina to this moving bar, performed the same linear decoding, and estimated the power spectral density of the error (Fig 3A; see methods). The error spectrum was clearly frequency-dependent with a dip between 1 and 4 Hz.

Note that the error estimation over this large band illustrates the benefit of the white noise stimulus. If we had used the random walk described in the previous section, we could not have measured the error spectrum for the highest frequencies (above 3 Hz), because this trajectory does not have enough power in this band. Nevertheless, for the lower frequencies (below 2Hz), we found error spectra that were consistent with what was measured with this white noise stimulus. In the guinea pig, where only the random walk stimulus was employed, we also found an error spectrum that was consistent with what we found in the salamander for the same stimulus ensemble.

The bar's motion also spanned a wider field than the array, so we hypothesized that errors would be larger for eccentric positions. We estimated the mean absolute error as a function of the position of the bar itself (Fig 3B). As expected, the error was position-dependent with the minimal error found for the position that was at the middle of the bar's range of motion. Note that this dependence on position can have more than one explanation. It can come from an incomplete sampling of the side positions by the recorded cells, but also from the fact that, when minimizing the error, more weight is given to the center positions, since they are explored more often. The error spectrum estimated above thus reflected an average between heterogeneous error levels taken at different positions of the bar. We thus aimed at separating simultaneously the effects of position and signal frequency on the prediction error.

For that purpose, we designed an error measure that took these two aspects into account. For each frequency, we filtered the error signal (i.e. the difference between the traejctory and the prediction) with a band-pass filter centered on that frequency, and then computed the variance separately for each position (see Methods). For proper normalization, we required that a weighted average of this position-frequency spectrum over the positions gives back the power spectrum shown in (Fig 3A). At the middle of the bar's range of motion, the error spectrum had a soft dip between 1 and 4 Hz, a slightly broader range of frequencies than when averaging over all bar positions (Fig 3C). Over that range of frequencies, the error was significantly below the average spacing between cones in the salamander retina [28, 29], illustrated by the dashed line. To our knowledge, this is the first demonstration of high precision in decoding the time-varying trajectory of an object's motion in the vertebrate visual system.

Contribution from different cell types

We split our population of ganglion cells into different functional types to understand how they contributed to estimating the moving bar's trajectory. Here we focus on analyzing single cells. When using the linear decoder, there was no difference between the decoding performance of OFF and ON cells ($p \geq 0.75$, Mann Whitney U-test). A similar result was obtained for the guinea pig retina, with no significant difference in decoding performance between OFF and ON cells ($p \geq 0.7$, Mann Whitney U-test). Furthermore, when we compared the overall firing rate during random bar motion, we found that ON and OFF cells were similarly responsive (in the salamander, 1.12 ± 0.09 Hz (ON cells) and 1.96 ± 0.49 Hz (OFF cells), with no significant difference, ($p \geq 0.37$, Mann Whitney U-test); in the guinea pig, 1.15 ± 0.14 Hz (ON cells) and 1.18 ± 0.2 Hz (OFF cells), and $p \geq 0.9$).

We went on to sub-divide the salamander ganglion cell population into six functional types according to the temporal dynamics of the receptive field center mechanism, as in previous studies [14, 30]. Biphasic and monophasic OFF cells were brisk, firing at high rates with spikes clustered in bursts, while the other types (medium and slow OFF cells) were sluggish, having slower temporal dynamics, lower peak firing rates and longer refractory periods. However, we did not find any clear difference between the decoding performance for the different cell types ($p \geq 0.05$, Mann Whitney U-test, for all the pairwise comparisons), similar to the result for ON and OFF cells.

These results illustrate that the functional properties of ganglion cells cannot be solely described by their classical receptive fields. A striking aspect of this sub type analysis was that ON cell contained as much information about the stimulus than OFF cells. If an ON cell is monophasic, and if it integrates the visual input linearly like a LN model does, in principle, it should not respond to the bar leaving the receptive field, and these ON cells should carry less information about the stimulus than OFF cells. This is not what we observed. Our results can be understood as arising from the biphasic temporal response of some ON cells, and/or from the pooling of excitatory and inhibitory non-linear subunits, as found in Y-type cells [31–35]. In this latter case, individual ON-type subunits would have small enough receptive fields that they can be excited by the trailing edge of the dark bar, which causes an increase in the local light intensity when it moves. Thus, ON cells with nonlinear receptive field subunits can also encode local motion with high accuracy.

Decoding based on the neural image

Our naive expectation was that ganglion cells would fire when the bar moved in their receptive field center, such that the population had a peak of activity that travelled with the moving

object [25]. In this case, the moving object's position is represented well by the peak in the neural image. Following a sudden reversal of motion, there is an error in the location of the neural image [24]. However, the spatial location represented by the synchronous burst of firing following motion reversal was shifted in the new direction of motion, so that it begins to "correct" for the retina's mistaken representation of the object position. As a result, a decoder based on the population vector could achieve good results for an object moving at constant velocity punctuated by sudden reversals [22]. A related example comes from place cells in the hippocampus, where a simple population vector decoder can reconstruct the position of the animal with good resolution [7, 8].

Thus, we wanted to test if this previously used population vector decoding could be successful at predicting the bar's trajectory, and see if the retinal activity showed a moving hill that followed the bar's position. Compared to our previous method of linear decoding, it is interesting to note that population vector decoding corresponds to linear decoding with a temporal window of just a single time bin, and the weight of the filter being related to the value of the receptive field center.

We first constructed the "neural image" as in previous studies, where its peak corresponded well with the position of the moving bar [24, 25]. The neural image is the spatial pattern of firing in the ganglion cell population as a function of time. We calculated it by plotting the firing rate of the cells as a function of their receptive field position (see Methods). Despite the success of the neural image in tracking simpler object motion [22, 25], we found that it did a poor job with our complex trajectory (CC = 0.06, Fig 4A). Similar results were obtained on the 3 salamander retinas (average performance CC = 0.09 ± 0.04). The performance of the neural image decoder was also poor for the guinea pig retina (CC = 0.17, n = 1).

One of the reasons for this poor performance might be that the bar can move to regions where not all of the ganglion cells are recorded by our array. To take this into account, we divided the neural image at each position by the total number of spikes recorded at this position (see Methods). Another reason for the poor performance could be that each cell can respond with a different latency and have different temporal dynamics. So we tried to improve the neural image by taking into account the temporal extent of the receptive field. To this end, we added the temporal profile of the receptive field each time a neuron spiked. Together these two corrections significantly improved the correlation between the neural image and the real trajectory, but the final performance was CC = 0.25 for the salamander and CC = 0.28 for the guinea pig, far below the performance of the linear decoder.

One of the primary reasons for this poor performance was that neural activity was too sparse to continuously represent the moving object's location (Fig 4B). We quantified the sparseness of a cell as in [36], by measuring the amount of time where the firing rate was above 5% of its maximum value. On average, we obtained that it was above this threshold 11% of the time (n = 117 cells), which means that neurons remained inactive 89% of the time. The population activity was also sparse: it was above 10% of its maximum only 36% of the time. The neural image performance slightly improved for epochs where the global firing rate was higher, as illustrated by a small negative correlation between the decoding error and the global firing rate (CC = −0.13). However, the performance was still poor during periods of relatively high firing rate: CC = 0.27 for the salamander if we restrict ourselves to epochs where the total firing rate was at least half of the max firing rate). Thus, there was no clear moving hill of neural activity from which the position of the bar could be inferred. Instead, we observed sparse neural activity with no obvious spatial structure.
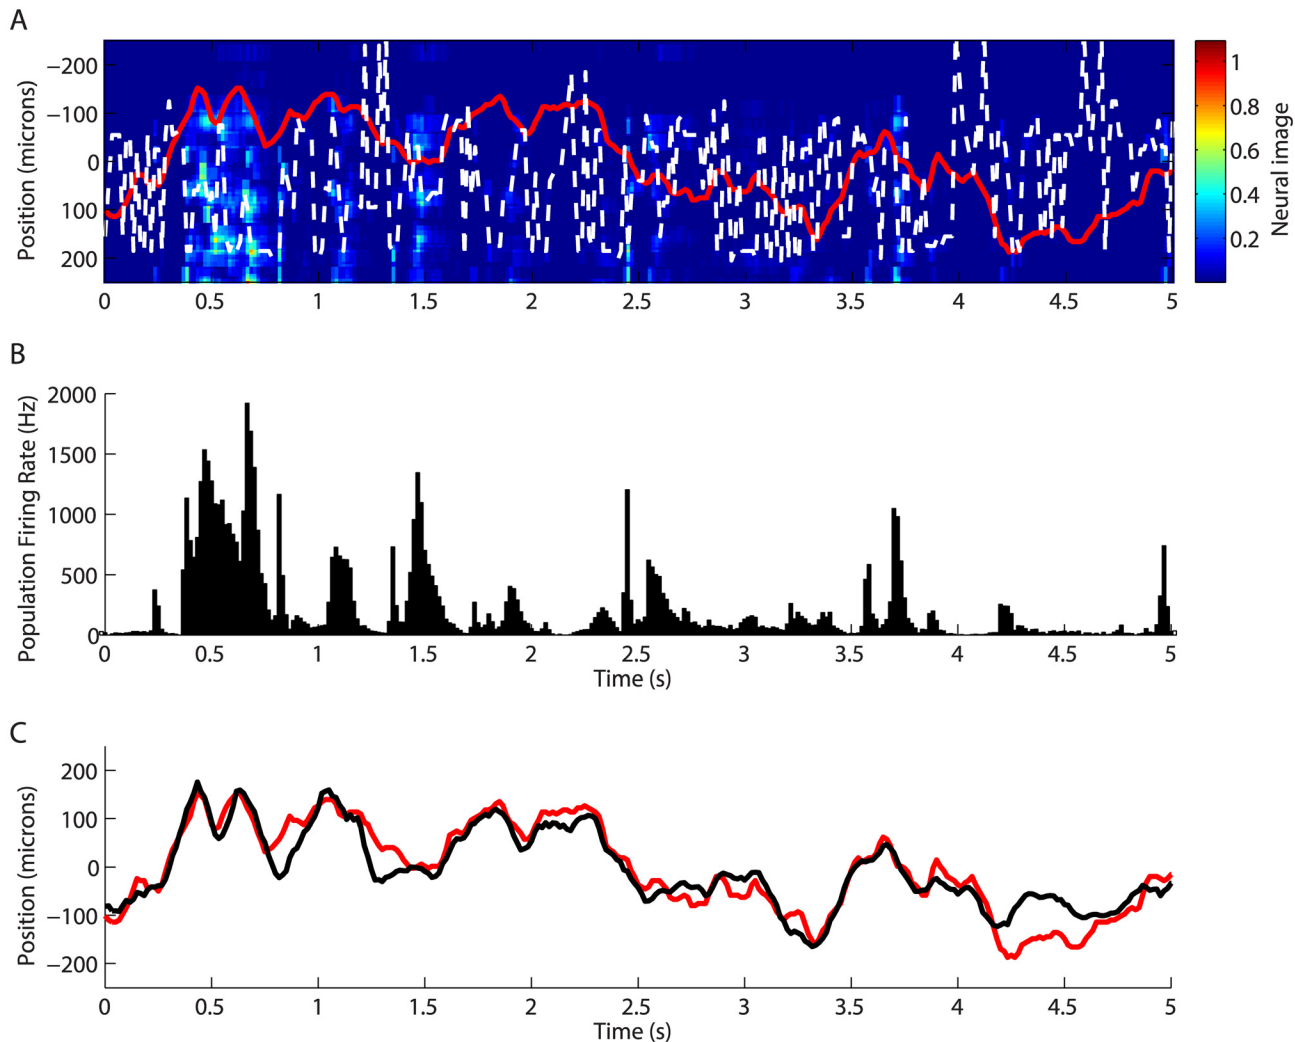
Fig 4. Decoding based on the neural image in the salamander retina. A: Neural image in response to the moving bar. Color plot: neural image of the ganglion cell's population activity at each point in time. White points: most likely position of the bar inferred from the peak in the neural image. Real trajectory in red. B: Population firing rate summed over all the cells as a function of time, for the same time window than A. C: Prediction of the bar's trajectory using the linear decoding (black); real trajectory (red).

doi:10.1371/journal.pcbi.1004304.g004

Coding for motion in the receptive field surround

From this lack of spatial structure we hypothesized that even the ganglion cells whose receptive fields were far from the bar carried information about the trajectory. To test this, we displayed the randomly moving bar stimulus in three different average locations, each separated by 430 microns and lasting 20 minutes (see Methods and Fig 5A). We then estimated the bar's trajectory for each stimulus ensemble one-at-a-time. The trajectory could be decoded for the three locations at nearly the same level of high performance (Fig 5B, 5C, 5D).

We then tried to decode the trajectory with each individual cell, thus re-learning the decoding filter for each of them. We plotted the individual decoding performance of different cells as a function of the distance between the cell's receptive field center coordinate and the mean position of the bar (Fig 5E). The decoding performance displayed a mild decrease as a function of
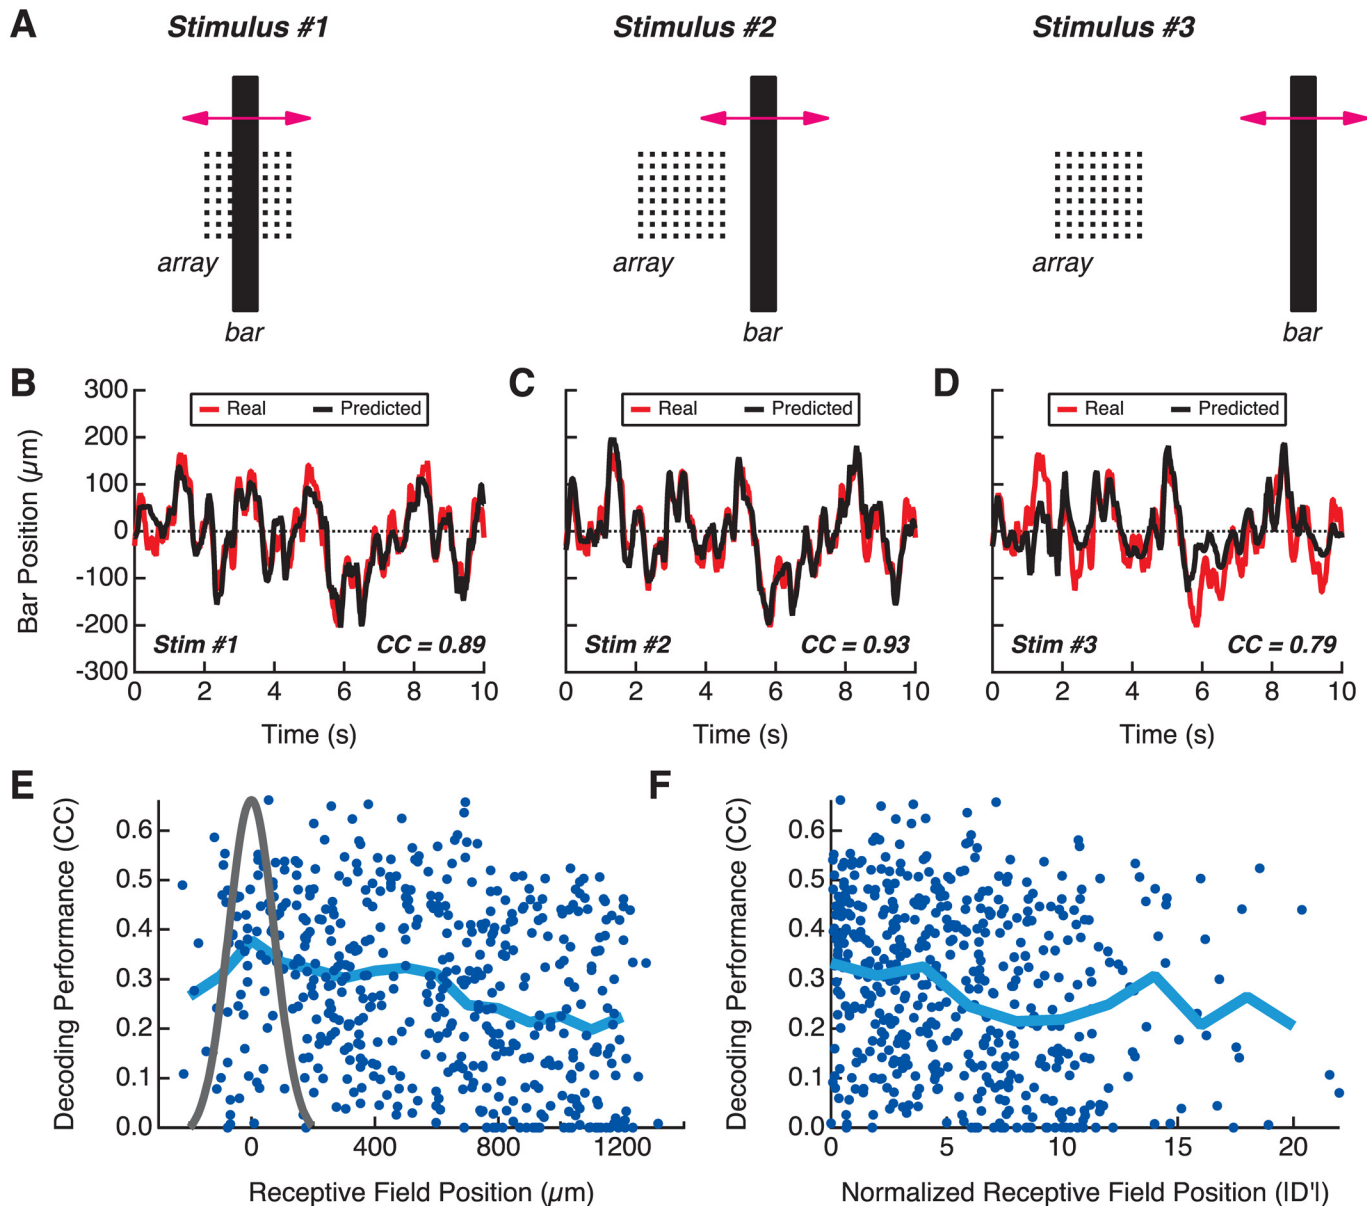
Fig 5. Coding for motion in the receptive field surround in the salamander retina. A: Schematic of the experiment: the bar is randomly moved at three different average locations relative to the array. **B, C, D:** Prediction of the bar's trajectory using the linear decoding (black); real trajectory (red), for the three average locations above. **E:** Performance of linear decoding (blue) for individual ganglion cells (dots) plotted as a function of the distance between their receptive field center coordinate and the average bar position; probability distribution of bar position (black). Blue line: average decoding performance as a function of the distance. **F:** Performance of linear decoding (blue) for individual ganglion cells (blue dots) plotted as a function of the normalized distance between their receptive field and the average bar position (see text). Blue line: average decoding performance as a function of the normalized distance.

doi:10.1371/journal.pcbi.1004304.g005

distance (on average, a CC loss of 0.11 per mm). This decrease was partly due to an increasing number of cells falling to $CC = 0$. Many cells still carried information about the bar at distances of more than 800 μm away from the bar, while the average receptive field radius is 115μm [30]. We then checked if these results could be explained by some ganglion cells having large receptive fields that would still overlap with the bar trajectory. We defined a normalized distance between the bar and each receptive field as the difference between their average position divided

by the receptive field width (see methods). When estimating this distance, the bar width is taken into account by convolving the receptive field with the bar. In these units, the point at which the strength of the surround exceeds that of the center (the zero crossing point) is roughly 1.5 [37]. However, if we want to know where motion is entirely in the surround, we need to take into account the range of bar positions. If we assume the bar rarely goes further than 3σ (where $\sigma = 73\mu m$), then this normalized distance is roughly $\frac{219\mu m}{115\mu m} \sim 1.9$. Therefore, normalized positions beyond ~ 3.5 show examples of ganglion cells where motion is entirely in the surround, many of which still show high fidelity motion tracking. These results show that the representation of the moving bar was distributed across a large population of cells spread widely in spatial location, far beyond what the extent of the receptive field center would predict. This also explains why the neural image failed to code for the bar position.

We can gain additional insight into why ganglion cells can precisely encode motion on their surround by inspecting the corresponding neural responses. When the bar moved over the receptive field center, ganglion cells fired in discrete firing events on a background of silence (Fig 6A, 6B). Similar spike train structure has been seen under many other stimulus conditions [14, 36, 38]. Motion in the surround also triggered responses with similar event-like structure (Fig 6C, 6D, 6E, 6F). Another way to explore the nature of the precise coding of an object's trajectory is to compute the spike-triggered speed average, which is analogous to the standard spike-triggered stimulus average [39]. This reveals that the speed was higher than average for a period ~ 200 ms before a spike (depending on the cell) and slightly below average for a period of almost a second before a spike (Fig 6G, 6H). Thus, ganglion cells tended to spike after a brief acceleration. Importantly, this spike-triggered speed average was nearly invariant to the overall location of the moving bar (Fig 6G, 6H; red, green, blue). Together, these results support the interpretation that the structure of the code for a moving object's location was substantially similar for motion in both the center and surround, with sudden accelerations triggering firing events in ganglion cells. Yet at the same time, ganglion cells could encode the position of the moving bar quite well. Even though the spike-triggered speed average was similar for the 3 positions, the decoding filters were very different (Fig 6I, 6J). The positive ranges on the y-axis correspond to the bar getting closer to the receptive field. So for the green and red curves, the cell tends to signal when the bar gets closer to its receptive field. In the far surround case, this might correspond to excursions of the bar into the closer surround. Thus, we found that ganglion cells could encode different features of the trajectory depending whether the bar is close or far from its receptive field.

A redundant and distributed code

Because of the distributed nature of the code, we wondered how much redundancy was present in the ganglion cell population. As seen above (Fig 2), after a rapid initial increase, the decoding performance essentially saturated as a function of the number of cells. This saturation suggests that the information encoded by a new cell was highly redundant with the rest of the population. To quantify the redundancy in a principled way, we estimated the information rate between the real and predicted trajectories for single cells and for groups of cells of different sizes. Note that this measure is a lower bound on the true mutual information between the spike trains of a group of cells and the stimulus [26, 40].

Information about the moving bar's trajectory encoded by subsets of ganglion cells increased as more cells were added (Fig 7A). But for the same number of cells, the performance varied substantially depending on the particular cells chosen in the subset. A natural hypothesis for this variability is that some cells carry more information about the stimulus than others. To
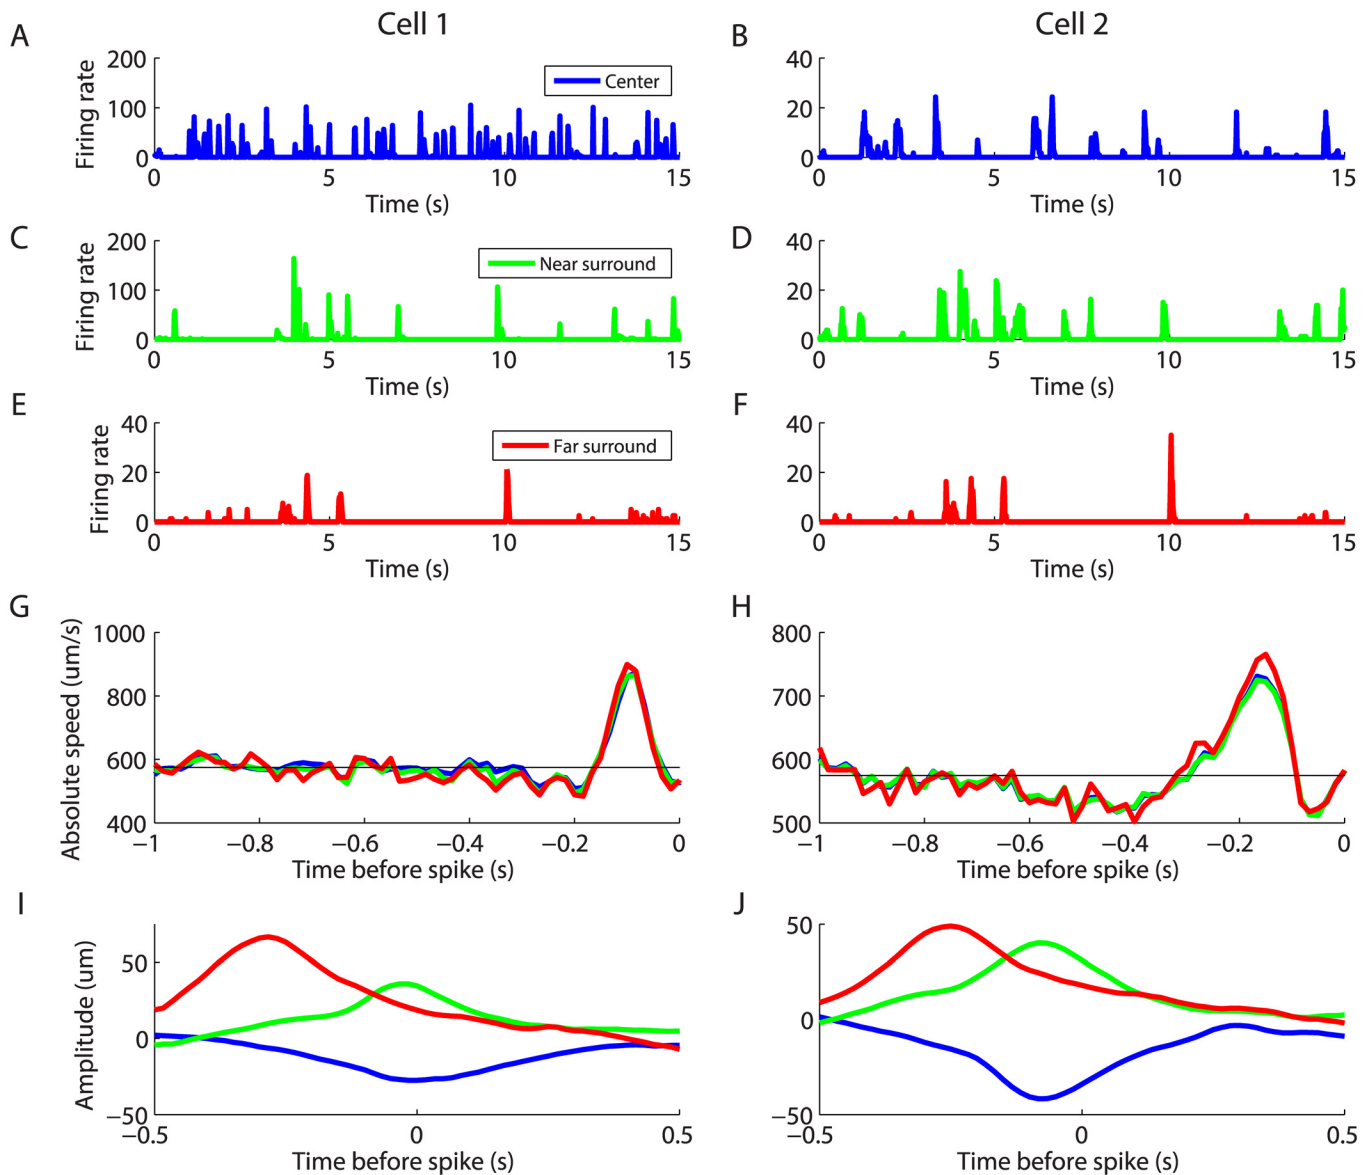
Fig 6. Responses to motion in the receptive field surround in salamander retina. Responses are shown for two example cells, cell 1 (first column) and cell 2 (second column). For each cell, PSTH of the response to the same trajectory is plotted when the trajectory is displayed over the receptive field center (blue, **A** and **B**); in the near surround (green, **C** and **D**); and in the far surround (red, **E** and **F**). **G, H:** average time course of the bar speed before a spike for different positions of the trajectory (color code matches that in A-F). **I, J:** Decoding filter of the same cells for different positions of the trajectory (same color code)

doi:10.1371/journal.pcbi.1004304.g006

explore this relationship, we compared the "total" information encoded by a group of ganglion cells (y-axis in Fig 7A and 7B) versus the sum of the individual informations encoded by each cell in the group (x-axis in Fig 7B). Plotted in these units, there was much less variability in the information encoded by different groups of ganglion cells. The same result was obtained in the guinea pig retina (Fig 7C, 7D).

In making this comparison, we picked subsets randomly. To better test the predictive power of the sum of individual informations, we looked for the subsets with the best and the worst
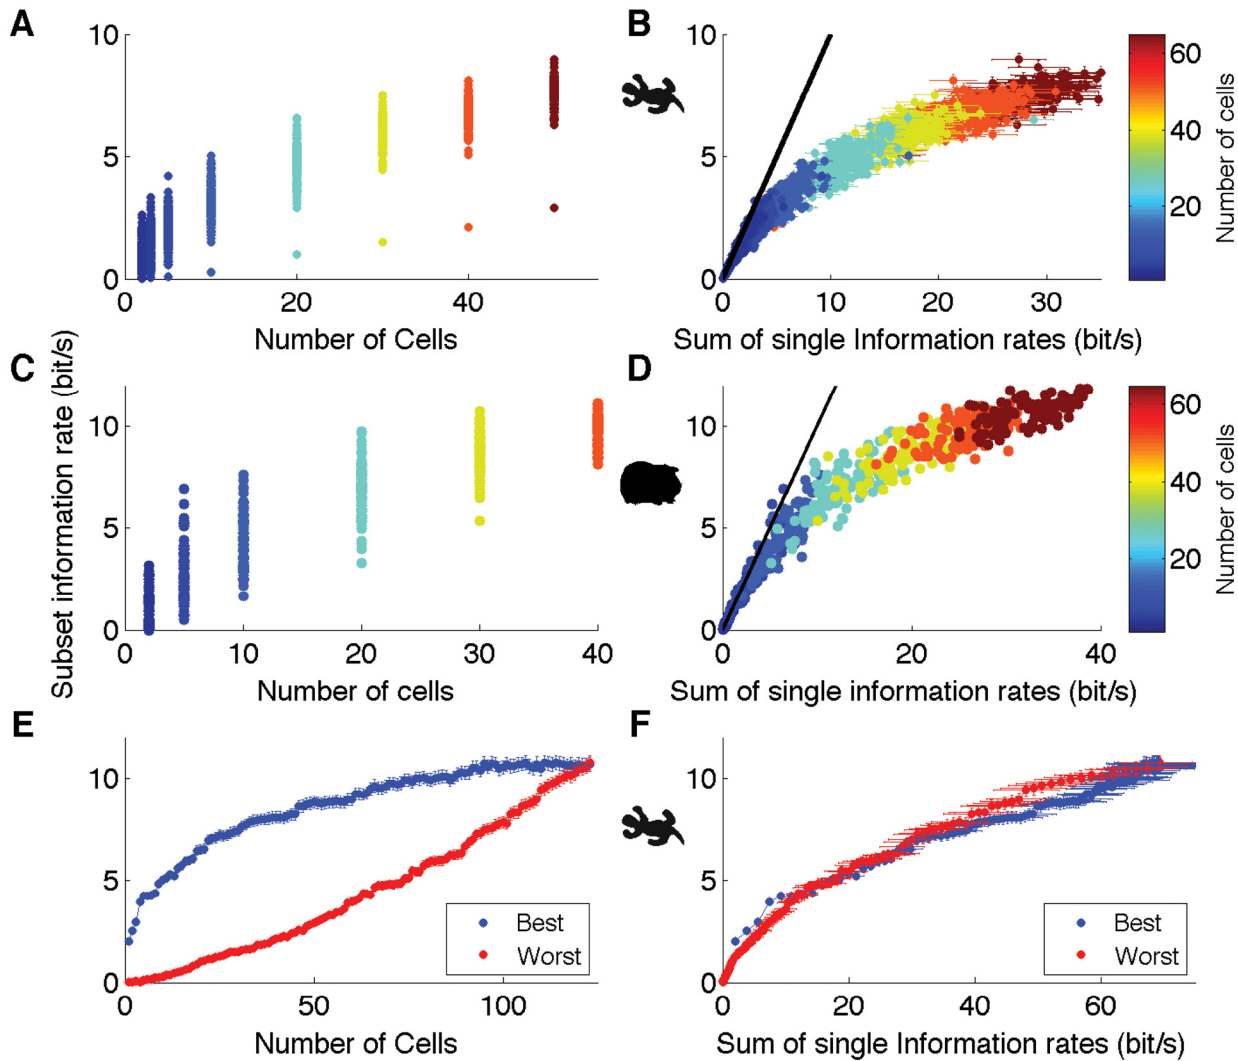
Fig 7. Redundancy of the retinal code. A, B: Information rate obtained when decoding with different subsets of ganglion cells (dots) plotted against the number of cells (A) or the sum of the individual information rate of each cell of the subset (B); color scale indicates the number of cells in a subset. **C, D:** Plots analogous to panels A, B but for the guinea pig retina. **E, F:** Information rate for the best (blue) and worst (red) subsets of salamander ganglion cells (see text) plotted against the number of cells (E) or the sum of the individual information rate of each cell of the subset (F).

doi:10.1371/journal.pcbi.1004304.g007

decoding performances for the same number of cells. There were too many combinations of cells having a given group size to allow for a comprehensive search over all possible subsets. Instead, we ordered the cells by their individual informations, and for each number of cells N, we estimated the total information for the N "best" cells and the N "worst" cells. While there was a large difference in the performance for the best and the worst subsets when plotted against the number of cells (Fig 7E), this difference was almost entirely compensated when plotting them against the sum of individual information (Fig 7F). For a group of ganglion cells having a similar sum of individual informations (i.e. a difference less than 0.5 bits/s), the average difference between worst and best subsets was 8% of the best. Part of the remaining difference is likely due to the error in the estimation of the individual informations (horizontal error bars) due to
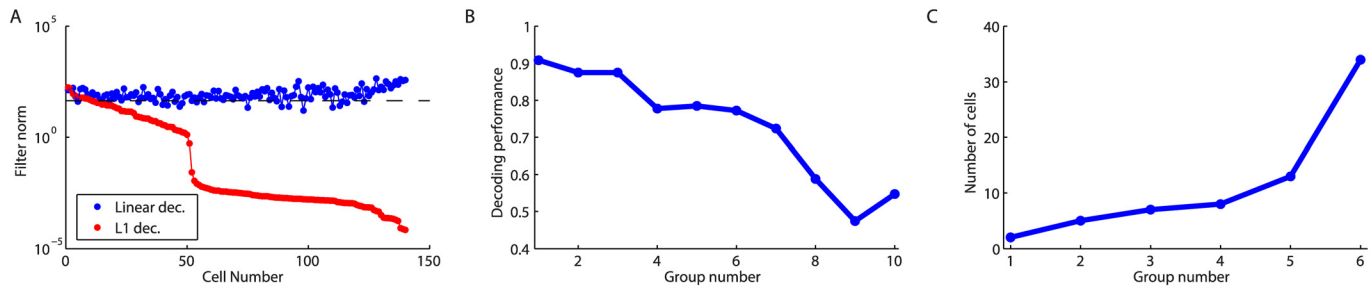
Fig 8. Regularized decoding. A: Comparison of the norm of the filters of each cell for the linear decoding (blue) against the regularized L1 decoding (red). The dashed line separate the 10 cells with the highest regularized norm. B: Decoding performance of groups of 10 cells sorted by their L1 norm, from the highest to the lowest. C: Number of cells in each subset that could reach a decoding performance of 0.85.

doi:10.1371/journal.pcbi.1004304.g008

the finite size of the data. In contrast, the "worst" cell groups had on average 55% less information than "best" groups when matched for number of cells.

These results together show that we could predict the decoding performance of a group of ganglion cells from the properties of individual cells. In both the guinea pig and the salamander, small groups (less than 10 cells) had a total information that was close to the sum of individual informations, so the different cells carried nearly independent information. But as the sum of individual informations increased, it became much larger than the total information encoded by the group. Thus, the redundancy among ganglion cells (defined as one minus total information divided by sum of individual informations) increased with the number of cells and became quite large as the performance saturated, with a 6.4-fold redundancy for the 123 cells in the salamander, and 6.6-fold for 140 cells in the guinea pig.

This suggests that multiple subsets of cells might be able to encode the bar trajectory. To find these disjoint subsets, we used L1 regularized decoding, a method that simultaneously minimizes the error in the prediction while setting to zero as many filter coefficients as possible (see Methods). Compared to the previous optimization where there were no constraints on the filter amplitude, here most filters were driven to zero (Fig 8A), but the prediction performance remained unchanged (CC = 0.93). The filters with the highest amplitude corresponded to the neurons that were the most useful at decoding the trajectory. We ordered them from the highest to the lowest filter amplitude, and split them into groups of 10 that we used to decode the bar position. The decoder was then re-trained without L1-regularization for each group of 10. As expected, the cells with higher filter amplitude were better at decoding the trajectory than those with lower amplitude (Fig 8B), which confirmed that the L1 method picked the most informative cells. Of particular interest, we also found that the performance remained above CC = 0.75 for the first seven groups, indicating many different groups of 10 neurons could decode the trajectory with high accuracy. We then looked for disjoint subsets that would give a prediction performance above CC = 0.85, which was 90% of the performance for the entire population. We first tried to decode using only the cell with the highest filter amplitude, and then kept adding cells to the subset, ordered from the highest to the lowest amplitude, until the decoding performance using the subset reached 0.85. We then restarted the same process after discarding the cells already used in the subset. By iterating this selection, we could find 6 disjoint subsets of increasing size that were all able to decode the trajectory of the bar (Fig 8C). All these subsets had a decoding performance between 0.85 and 0.87. These analyses together demonstrate that we could define 6 or 7 subsets with high decoding performance, similar to the estimated redundancy of information in the population.

So far we have estimated redundancy for random subsets, including ganglion cells from all the different types. Perhaps, the redundancy is significantly different for cells of the same functional type. We could define one cell type in the salamander retina, the biphasic OFF type cell (Fig 9B), whose receptive fields clearly tiled visual space (Fig 9A), as expected from previous studies [14, 30]. Using the same method as before, we took many subsets of cells belonging to the same cell type and plotted the sum of the individual information against the information carried by the subset. The values obtained in this redundancy plot were similar to the ones obtained with random subsets (Fig 9C). When averaging over subsets having the same sum of individual informations, we found that the information encoded by small subsets was essentially identical for ganglion cells of the same functional type as for random functional types. However, for large enough subsets, the information encoded by cells of the same functional type was significantly smaller than for random subsets (Fig 9D).
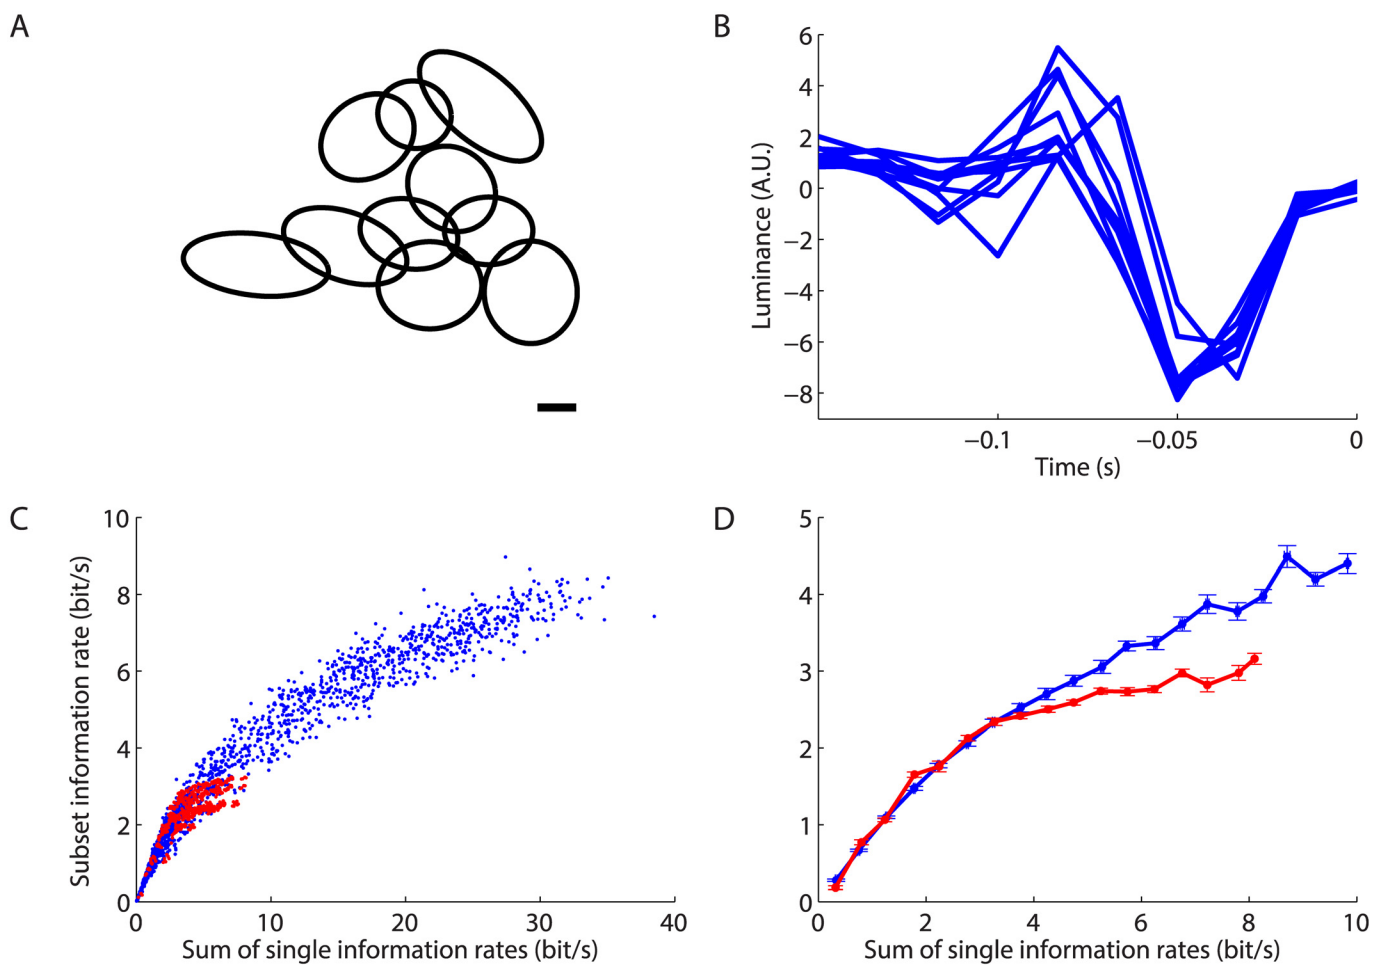

Fig 9. Redundancy of cells of the same type. A: Ellipses fitted to the spatial receptive field profile for the biphasic OFF cells. Scale bar: 100 microns. B: Temporal profiles of the fast biphasic OFF cell type. C: Plots analogous to Fig 7B. Information rate obtained when decoding with different subsets of ganglion cells (dots) plotted against the sum of the individual information rate of each cell of the subset; subsets are either picked randomly (blue), or from the biphasic OFF cells (red). D: Average information rate plotted against the sum of the individual information rate of each cell of the subset; subsets are either picked randomly (blue), or from the biphasic OFF cells (red).

doi:10.1371/journal.pcbi.1004304.g009

Discussion

We have shown that we could reconstruct the entire time-varying trajectory of an object in complex motion from the activity of a large population of sensory neurons in the early visual system. Using linear decoding, populations of 100+ retinal ganglion cells were able to estimate the location of the object during its motion with an accuracy better than the spacing between cone photoreceptors in the salamander (20 μm [28, 29]). Such hyperacuity performance has, to our knowledge, never before been demonstrated for a full time-varying reconstruction of a sensory stimulus. This result is consistent with an analysis of optimal performance of motion estimation that is possible by integrating signals over an array of photoreceptors [27].

The naive expectation was that many individual ganglion cells would fire spikes when an object was moving on their receptive field center and that the activity of the entire population would resemble a "hill" of higher activity that continuously tracked the location of the moving object. While the spatial map of the retina certainly localizes ganglion cell activity to some extent, our analysis revealed that neural activity was too sparse, and too distributed in space in a region extending into the receptive field surround, to closely track the object's location. Instead, the ganglion cells exhibited a highly distributed and redundant population code that allowed several distinct groups of cells to achieve super-resolution in tracking the location of a diffusively moving object.

Comparison to previous results

Linear decoding was first developed and applied to the H1 interneuron in the fly to estimate wide-field motion [10]. Linear decoding has also been applied to small populations of retinal ganglion cells to estimate the light intensity of a spatially uniform, white noise stimulus [11, 41], and to reconstruct a checkerboard stimulus for large populations of primate parasol cells [12]. However, in these previous studies, the LN model worked well for the stimuli employed. Some of the most notable quantitative successes of the LN model have been demonstrated with spatially uniform stimulation [42–44]. Pillow et al (2008) showed that the LN model performed very well for parasol cells under their stimulus ensemble, perhaps due to the fact that the checkers were large enough to cover a cell's entire receptive field center.

This is not the case for the kind of diffusive motion we used here, which includes pauses, starts, and reversals that trigger transient bursts of firing that cannot be explained by the LN model [24, 35]. One important factor in the failure of the LN model is the presence of nonlinear spatial subunits inside ganglion cell receptive field [34, 35, 39, 45, 46]. During spatially uniform stimulation, all nonlinear spatial subunits are activated together, allowing models that do not include this structure to approximate well the ganglion cell's firing rate. But during irregular motion, the subunits can be independently activated, such that a model without this structure performs poorly [35, 47].

Other studies have tried to estimate a visual stimulus from the activity of populations of ganglion cells. The speed of a moving bar could be estimated from the activity of parasol ganglion cells [1, 2]. However, this study was restricted to the case of motion at constant speed, so only one number was estimated: the speed of the moving bar. For the case of a moving texture that switched speeds every 500 ms, discrete classification of the speed was possible [23] (see also [48]). Similar classification studies have been performed in the visual cortex [4, 49]. Our study differs from these by reconstructing the full time-varying position of the bar, a task that is much more difficult than the classification of several stimuli or the estimation of a single parameter.

Structure of retina's population code

One of the most significant surprises in this study was the lack of precise spatial structure of the retina's population code for this diffusive motion stimulus. One key factor was the complexity of our motion stimulus. Most early studies of how ganglion cells respond to motion involved an object moving at constant velocity [22], a case in which the picture of the neural image seems to hold. In our stimulus ensemble, the object moved at varied speed, including discontinuities of motion and moments of high acceleration. Motion discontinuities like the sudden onset or reversal of motion trigger transient bursts of firing [24, 35], as does acceleration in general [23]. Finally, while our diffusive motion is more complex than the stimuli aforementioned, it is still far from the complexity of a natural movie. It would be interesting for future studies to investigate if the responses to natural movies show a precise spatial structure or not.

Another striking result was how well ganglion cells encoded diffusive motion of a dark bar in their surrounds, which partly explained the diffuse spatial structure of the code. Motion in the surround has long been known to be able to generate excitatory responses in ganglion cells [50]. Our analysis demonstrates that this activation is not just a global alert signal, but can be used to encode the precise trajectory of the bar. Detailed investigation has revealed that the response of ganglion cells to gratings drifting in their surround can be modeled as the pooling of excitatory and inhibitory non-linear subunits [51, 52]. In the salamander, the surround response to annulus of modulated luminance can also be modeled by pooling non-linear subunits [53]. Many other studies have demonstrated the presence of disinhibitory circuits within the receptive field center, involving transient amacrine cells turning off sustained inhibition from sustained amacrine cells [54–57]. These circuits may allow precise information about motion in the surround to be conveyed by the firing of ganglion cells.

Finally, we found that the decoding performance of ON and OFF-type ganglion cells was comparable. This is surprising if one thinks about how the moving bar interacts with the classical receptive field. In this view, the dark bar should cause excitation of OFF cells and inhibition of ON cells. Thus, ON cells should fire at lower rates and convey less visual information. Neither of these properties was observed in our data. However, we can understand these results if we instead think about ganglion cells as having nonlinear subunits in their receptive fields. Such subunits have smaller receptive fields than the entire ganglion cell, and therefore can be excited by the trailing edge of the moving dark bar, which is effectively an ON-type stimulus. Many ganglion cell types have been shown to possess spatial subunits, including brisk cells like the alpha cell in the guinea pig [32] and in the mouse [34] and the fast OFF cell in the salamander [33] as well as sluggish cells like the local edge detector [56, 58].

One might also think that our results apply primarily to retinal ganglion cells with transient response types and that sustained ganglion cells might be able to continuously follow a moving object. However, our recorded populations included sustained cells. More specifically, the fast OFF cells in the salamander, which were the focus of the original work on motion anticipation [25] as well as a recent study using a neural image decoder [22] have a sustained response to smooth motion. Yet these same cells exhibit sporadic, event-like firing for diffusive motion. Thus, the stimulus ensemble changes the qualitative nature of ganglion cell spike trains, and studies of how sustained cells respond to constantly drifting gratings do not predict how those cells would respond to a diffusive motion. And in fact, our recordings in the guinea pig also showed sparse, event-like firing across the ganglion cell population (Fig 1C) in contrast to the intuition derived from many studies of mammalian retina using drifting gratings.

But important caveats to our findings are needed: our stimulus ensemble—a diffusively moving bar—is artificial and only explores a subset of all possible motion patterns. It is

conceivable that more naturalistic stimuli could strongly activate dedicated subpopulations of ganglion cells. Furthermore, we could only define broad functional groups rather than "true" cell types with distinct anatomy. Perhaps true anatomical cell types would have more clearly differentiated function.

A flexible code

Clearly, the distributed and redundant neural code that we have observed is not how a human engineer would design a system to track motion. In most theoretical studies where a population of neurons has to code for the value of a parameter, the single cell selectivity is modeled by a tuning curve. As a result, the activity in response to the stimulus is local, and neurons code best for stimuli at the peak of their sensitivity, or at the side of their tuning curves, depending on the noise level [59]. Experimentally, this corresponds to the case of the neural image. While the peak of neural activity does track objects moving at constant velocity [22, 25], we have shown that, in our case of diffusive motion, the retinal code was too sparse and spatially diffuse for the peak of neural activity to track the position of the bar. This organization challenges the classical theory that the position of an object could be simply tracked by following the peak of firing rate in the retinal output.

Even though the retinal code did not have a straightforward spatial organization, it could still be read out with one of the simplest of all decoding mechanisms: the linear decoder. One of the implications of the success of the linear decoder is that it makes information easy to extract by subsequent neural circuits [60]. More generally, a fundamental function of sensory circuits might be to compute and actively maintain an "explicit" or linearly decodable representation of the most relevant features of the environment [61]. This appears to be the case for IT cortex: the identity of an object can be decoded linearly from IT neurons, but not from V1 neurons [62]. Our results thus show that in this sense, the retina has an "explicit" representation of the position of a moving object.

The high redundancy of the retinal code was both notable and somewhat surprising. This observation is consistent with a previous study finding extensive redundancy during stimulation with natural movie clips [38], so the high redundancy we observed is not only due to the structure of the stimulus we employed here. Redundancy results from the high spatial overlap of ganglion cell's receptive fields: in the salamander, the number of receptive field centers covering a point in visual space is roughly 60 [63]. It also results from the distributed nature of the code, namely the fact that cells of different functional type do not convey categorically different visual information.

Why might such a distributed and redundant code be a beneficial organization? One advantage is that neural circuits in the central brain would have access to high precision motion tracking information from several groups of ganglion cells. This is useful because there are multiple features in the visual scene beyond just the location of a moving object that the brain might want to follow. For instance, if such circuits sampled all $\sim$ 100 ganglion cells with receptive fields overlapping one point in visual space, they could extract high resolution spatial and chromatic information about object identity in addition to the object's trajectory. Alternatively, if $\sim$ 100 ganglion cells with distributed spatial locations were sampled, then information about the characteristics of a larger object could be extracted in addition to its motion trajectory. Thus, the distributed, redundant population code maintains maximum flexibility with respect to the purposes of downstream neural circuits.

Of course, our analysis does not establish whether the brain uses a linear decoder nor how it might learn effective decoding kernels. In particular, we have shown that the decoders that reached the best performance had to be learned on the moving bar data directly. This means

that the detailed parameters of the decoder that achieves high performance might depend on the properties of the stimulus ensemble.

By demonstrating the broad spatial selectivity of ganglion cells, our study shows that neurons with complex tuning curves or mixed selectivity [64] do not appear only when merging information from different modalities in associative areas of the brain, but already at the earliest stage of sensory processing. Interestingly, these features of the retinal code were only apparent when the stimulus dynamics were sufficiently rich. The redundancy of trajectory representation that we uncovered could be a trade-off between discrimination and generalisation [65]. Future studies will have to understand how a representation of motion that is invariant to the context can be extracted from the retinal activity.

Methods

Ethics statement

This study was performed in strict accordance with the recommendations in the Guide for the Care and Use of Laboratory Animals of the National Institutes of Health. The protocol was approved by the Institutional Animal Care and Use Committee (IACUC) of Princeton University (Protocol Number: 1828).

Recordings

Retinal tissue was obtained from larval tiger salamanders (Ambystoma tigrinum) of either sex and continuously perfused with oxygenated Ringer's medium at room temperature. For guinea pig experiments, the tissue was perfused with AMES solution and maintained at 37 °C. Ganglion cell spikes were recorded extracellularly from a multi-electrode array with 252 electrodes spaced 30 μm apart (custom fabrication by Innovative Micro Technologies, Santa Barbara, CA). Details of the recording and spike sorting methods are described elsewhere ([14, 63]. Experiments were performed in accordance with institutional animal care standards.

The main dataset presented in this paper is available for download on the Dryad repository (doi: 10.5061/dryad.1dp55).

Visual stimulation

The stimuli were displayed on a CRT screen with a 60 Hz refresh rate [38] and a background light of 12 mW/m^2. We focused the stimulation plane precisely on the photoreceptor plane during each experiment. The stimulus presented was a dark bar of 100% contrast moving randomly over a gray background. The length of the bar was 2300 μm and the width 100 μm. The trajectory was given by Brownian motion with a spring-like force to bring the bar back to the center:

$$\frac{dv}{dt} = -\frac{v}{\tau} + \sigma \Gamma(t) - \omega_0^2 x \tag{1}$$

where x is the position of the bar, $v = \frac{dx}{dt}$ is the velocity of the bar, $\sigma^2 = 0.05 m^2 s^{-3}$, $\Gamma(t)$ is a Gaussian white noise, $\omega_0 = 9.42$ Hz and $\tau = 50$ ms.

For the experiments where we tested the decoding performance as a function of frequency, the trajectory was sampled every stimulus frame (16.7 ms) from a Gaussian distribution of standard deviation 60 μm and low-pass filtered at a frequency of 5 Hz using a Butterworth filter. The cut-off was chosen to be at 5 Hz because salamander cones do not seem to convey much information beyond this frequency [11].

Spatio-temporal receptive fields were measured by reverse correlation to a flickering checkerboard composed of squares of 69 μm size that were randomly selected to be black or white at a rate of 30 Hz [30].

Since salamander eyes are $\sim$ 4 mm in diameter, one degree of visual angle should correspond to approximately 35 microns in the retina plane, assuming a focal length of 2 mm. So the size of the array corresponds to 12–13 degree of visual angle.

Linear decoding

We used a linear model that takes the spike trains as an input and gives a prediction about the position of the bar as an output:

$$p(t) = \sum_{i,j} K_i(t - t_{ij}) + C \tag{2}$$

where t_{ij} is the j-th spike of the neuron i, and $p(t)$ is the predicted position of the bar over time. The filters K_i and the constant C are found by minimization of $\langle (p(t) - x(t))^2 \rangle$ (least square minimization) [11], where $x(t)$ is the real position of the bar, and angular brackets denote averages over time. The filters extended 500 ms before and after the spike.

All the results shown here were cross-validated: we trained our model on 2/3 of the data, and tested it on the other 1/3. We used one hour of recordings, which corresponds to 40 minutes for the training set, and 20 minutes for the test set. If the performance on the testing was below 0 (e.g. for decoding using a single cell), we set it to 0, since a negative performance is an effect of overfitting.

To characterize decoding performance, we estimated the normalized cross-correlation (CC) between the prediction and the real trajectory. For two signals $x(t)$ and $y(t)$, it is classically defined as $CC = \frac{\langle (x(t) - m_x) \cdot (y(t) - m_y) \rangle}{\sigma_x \cdot \sigma_y}$, where m_x and σ_x are the mean and standard deviation of $x(t)$, respectively, and $\langle \rangle$ denotes an average over time.

Position-frequency analysis

To get an estimation of the error signal as a function of both the real position and the frequency, we filtered the error $e(t) = p(t) - x(t)$ by a bandpass filter that did not introduce a phase delay, i.e. a Morlet Wavelet. The wavelet was normalized so that the averaged squared value of the filtered signal matched the corresponding value in the total power spectrum. For each possible position, we picked the times where the bar was at this position, and took the average of the corresponding squared values in the filtered signal. Formally, if $e_f(t)$ is the filtered version of $e(t)$ for frequency f, then the error $e(f, p)$ for position p and frequency f was computed as:

$$\frac{1}{N_p} \sum_{i=1}^{N_p} e_f(t_i)^2 \tag{3}$$

where the t_i's are all the times such that $x(t_i) = p$, and N_p is the number of such points.

Note that the error spectra measured here are in microns^2 Hz^{-1}. This means that, if the trajectory was a sinusoid at a given frequency, the variance of the decoding error should be the value estimated in this error spectrum. The dashed line representing cone spacing in Fig 3 corresponds to the square of the average distance between cones in the salamander (20μm). If the error spectrum value at frequency f is below this dashed line, it means that a sinusoidal trajectory at this frequency should be decoded with a root-mean-squared error lower than 20 μm. For more complex stimuli, the error should be integrated over the entire frequency band of the stimulus before being compared to this line.

Neural image

To construct the neural image, we assigned a spatial position to each cell as the peak of its receptive field. The activity of the cells was binned in 16.6 ms bins (corresponding to the refresh rate of the stimulus). At each time bin, we counted the number of spikes emitted by the cells at each spatial location. The resulting matrix was then smoothed across spatial positions with a Gaussian smoothing kernel of width 21 μm. The "most likely" position was obtained by taking the peak location of the neural image at each time where there was at least one cell firing. We also tried to decode the position by estimating an average of the position weighted by the firing rate, but the performance was even worse (CC = 0.027).

We then improved this neural image analysis in two ways. First, to take into account the fact that there might be "holes" in the coverage, we normalized the neural image with the following method. We assumed that, if we were to record all the retinal ganglion cells, the total firing rate should be the same for each spatial position. If this is not the case, the neural image will be biased to favor some positions just because cells spike more often there. To remove this bias, we divided the neural image at each position by the total number of spikes recorded at this position. Mathematically, it means that if we call the neural image $A(x, t)$, where x is the spatial position and t the time, we divided $A(x, t)$ by $\sum_t A(x, t)$

Second, to take into account the fact that each cell will respond with a specific latency and time course, we tried to improve the neural image by taking into account the temporal extent of the receptive field. To this end, each time a neuron at position x spiked at time t, instead of simply incrementing the neural image at (x, t), we added the temporal profile of the receptive field to the window $(x, [t - \tau; t])$.

Information estimation

The information rate between the true and decoded bar trajectory was estimated from their mutual coherence, $\gamma(f)$, as $I = -\int_0^{f_{max}} df \, \log_2(1 - \gamma^2(f))$[40]. Debiased coherence was computed using Chronux [66] using trajectory windows of duration $256/60s$ sampled at 60 Hz. The frequency range of integration $[0, f_{max}]$ was determined as the contiguous range where the estimated coherence is above zero at 10^{-2} significance level. Alternatively we considered a larger frequency range by lowering the significance threshold to 10^{-1}, and performing explicit debiasing by estimating the information on multiple subsets of the data of various sizes, and extrapolating to infinite sample size [67]. The results agreed within error bars. To validate our method, we generated synthetic "real" and "decoded" traces with Gaussian statistics and power spectra that were similar to the real traces, and of similar length. In this case the true information is analytically computable. We compared this exact result to our estimators to assess their accuracy and verify that the chosen parameters (window length, determination of frequency range) were suitable. For information rates above 0.5 bit/s the estimated error in the information rate is between 5–10%, while for rates of ∼ 0.1 bit/s the error increases to about 20%.

The relation between the normalized cross-correlation and the mutual information estimated this way was examined. If there were no temporal correlation, the mutual information would be $I = -0.5 \log(1 - CC^2)$. The measured information was significantly different from the value predicted by this formula. This is unsurprising since the formula does not take into account the different frequencies. However, we could fit the information values with a linear interpolation of this formula, i.e. $I = -0.5 \, \alpha \log(1 - CC^2) + \beta$. With this formula, we could explain 98% of the variability in the information values. So an empirical relationship could be inferred between the two quantities. Note that α was larger than one, indicating a synergy between

frequencies. Redundancy in a subset of cells A was defined as $R = 1 - \frac{MI_A}{\sum_{i \in A} M_i}$, where MI_A is the mutual information for the subset, and M_i the mutual information for each cell.

Regularized decoding

In L1-regularized decoding the least squares minimization problem $\langle (p(t) - x(t))^2 \rangle$ is substituted by $\langle (p(t) - x(t))^2 \rangle + \lambda \Sigma_i \|K_i\|_1$ where $\lambda \geq 0$ is the regularization parameter. Unlike the simple least squares case, this regularized minimization cannot be reduced to a simple straightforward linear algebra problem and a numerical solution is required. For this purpose we have used the implementation of the interior-point method by Boyd et al [68]. The value of the regularization parameter λ is usually chosen to minimize the mean-squared error (MSE). However, our goal is to assign zero weight filters to as many cells as possible while keeping a good decoding performance. Therefore, we choose a larger-than-optimal regularization parameter by allowing a 5% decrease in performance (CC) and a 20% increase of the MSE. Also, we run a 5-fold cross-validation procedure on the training set to validate the choice of regularization parameter.

Normalized distance

For each cell i, we estimated its receptive field using a classical checkerboard stimulus (see above) and convolved it with the bar. From the resulting spatial distribution, we estimated the average position m_i and the receptive field width s_i as the standard deviation of this distribution. The normalized distance was then defined as $d_i = \frac{|m_b - m_i|}{s_i}$. Note that the receptive fields were measured with the same stimulus display. If there was a blur in the optical system that we used to display the stimulus, it would equally affect the receptive field estimation. So our distance measure should not be affected by a blur in the optical system.

Acknowledgments

Author Contributions

Conceived and designed the experiments: OM MJB. Performed the experiments: OM KDS. Analyzed the data: OM VBS TM GT MJB. Wrote the paper: OM TM GT MJB.

References

1. Frechette ES, Sher A, Grivich MI, Petrusca D, Litke AM, Chichilnisky EJ. (2005) Fidelity of the ensemble code for visual motion in primate retina. J Neurophysiol. Jul; 94(1):119–35. doi: 10.1152/jn.01175.2004 PMID: 15625091

2. Chichilnisky E.J. and Kalmar R.S. (2003). Temporal Resolution of Ensemble Visual Motion Signals in Primate Retina. J Neurosci., July 30, 23(17):6681–6689 PMID: 12890760

3. Schwartz G, Macke J, Amodei D, Tang H, Berry MJ 2nd. (2012) Low error discrimination using a correlated population code. J Neurophysiol. Aug; 108(4):1069–88. doi: 10.1152/jn.00564.2011 PMID: 22539825

4. Oram MW, Foldiak P, Perrett DI, Sengpiel F. (1998) The 'Ideal Homunculus': decoding neural population signals. Trends Neurosci. Jun; 21(6):259–65. doi: 10.1016/S0166-2236(97)01216-2 PMID: 9641539

5. Kay KN, Naselaris T, Prenger RJ, Gallant JL. Identifying natural images from human brain activity. (2008) Nature. Mar 20; 452(7185):352–5. doi: 10.1038/nature06713 PMID: 18322462

6. Georgopoulos AP, Schwartz AB, Kettner RE. (1986). Science. 1986 Sep 26; 233(4771):1416–9. doi: 10.1126/science.3749885 PMID: 3749885

7. Wilson MA, McNaughton BL. (1993) Dynamics of the hippocampal ensemble code for space. Science. Aug 20; 261(5124):1055–8. doi: 10.1126/science.8351520 PMID: 8351520

8. Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. J Neurophysiol. 1998 Feb; 79 (2):1017–44. PMID: 9463459

9. Pasley BN, David SV, Mesgarani N, Flinker A, Shamma SA, Crone NE, Knight RT, Chang EF (2012). Reconstructing speech from human auditory cortex. PLoS Biol. Jan; 10(1):e1001251. doi: 10.1371/journal.pbio.1001251 PMID: 22303281

10. Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D. (1991) Reading a neural code. Science. Jun 28; 252(5014):1854–7. doi: 10.1126/science.2063199 PMID: 2063199

11. Warland DK, Reinagel P, Meister M. Decoding visual information from a population of retinal ganglion cells (1997). J Neurophysiol. Nov; 78(5):2336–50. PMID: 9356386

12. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP (2008) Spatio-temporal correlations and visual signalling in a complete neural population. Nature 454: 995–9. doi: 10.1038/nature07140 PMID: 18650810

13. Field GD, Gauthier JL, Sher A, Greschner M, Machado TA, Jepson LH, Shlens J, Gunning DE, Mathieson K, Dabrowski W, Paninski L, Litke AM, Chichilnisky EJ. (2010). Functional connectivity in the retina at the resolution of photoreceptors. Nature. 2010 Oct 7; 467(7316):673–7. doi: 10.1038/nature09424 PMID: 20930838

14. Marre O., Amodei D., Deshmukh N., Sadeghi K., Soo F., Holy T.E., Berry M.J. II. Mapping a Complete Neural Population in the Retina. Journal of Neuroscience, 24 October 2012, 32(43):14859–14873. doi: 10.1523/JNEUROSCI.0723-12.2012 PMID: 23100409

15. Segev R, Schneidman E, Goodhouse J, Berry MJ 2nd (2007). Role of eye movements in the retinal code for a size discrimination task. J Neurophysiol. 2007 Sep; 98(3):1380–91. doi: 10.1152/jn.00395.2007 PMID: 17625063

16. Jacobs AL, Fridman G, Douglas RM, Alam NM, Latham PE, Prusky GT, Nirenberg S. (2009). Ruling out and ruling in neural codes. Proc Natl Acad Sci USA. Apr 7; 106(14):5936–41. doi: 10.1073/pnas.0900573106 PMID: 19297621

17. Deban SM, Wake DB, Roth G (1997). Salamander with a ballistic tongue. Nature Sept 4; 389; 27–28. doi: 10.1038/37898

18. Roth G. Visual Behavior in Salamanders (Springer, Berlin, 1987).

19. Llinas R. Frog Neurobiology: A Handbook. Berlin: Springer-Verlag, 1976.

20. Pitkow X, Sompolinsky H, Meister M. A neural computation for visual acuity in the presence of eye movements. (2007) PLoS Biol. Dec; 5(12):e331. doi: 10.1371/journal.pbio.0050331 PMID: 18162043

21. Rodieck RW, Stone J. Response of cat retinal ganglion cells to moving visual patterns. J Neurophysiol. 1965 Sep; 28(5):819–32. PMID: 5867881

22. Leonardo A, Meister M (2013). Nonlinear dynamics support a linear population code in a retinal target-tracking circuit. J Neurosci. Oct 23; 33(43):16971–82. doi: 10.1523/JNEUROSCI.2257-13.2013 PMID: 24155302

23. Thiel A, Greschner M, Eurich CW, Ammermuller J, Kretzberg J. (2007) Contribution of individual retinal ganglion cell responses to velocity and acceleration encoding. J Neurophysiol. Oct; 98(4):2285–96. doi: 10.1152/jn.01342.2006 PMID: 17596411

24. Schwartz G, Taylor S, Fisher C, Harris R, Berry MJ 2nd. (2007) Synchronized firing among retinal ganglion cells signals motion reversal. Neuron. Sep 20; 55(6):958–69. doi: 10.1016/j.neuron.2007.07.042 PMID: 17880898

25. Berry MJ 2nd, Brivanlou IH, Jordan TA, Meister M. Anticipation of moving stimuli by the retina. Nature. 1999 Mar 25; 398(6725):334–8. doi: 10.1038/18678 PMID: 10192333

26. F Rieke, D Warland, RR van Steveninck, W Bialek (1999). Spikes: exploring the neural code.

27. Ruderman DL and Bialek W (1992). Seeing beyond the Nyquist limit. Neural Comp 4, 682–690. doi: 10.1162/neco.1992.4.5.682

28. Sherry DM, Bui DD, Degrip WJ. (1998). Vis Neurosci. 1998 Nov–Dec; 15(6):1175–87. doi: 10.1017/S0952523898156201 PMID: 9839981

29. Zhang J., Yang Z. and Wu S.M. (2004) Immunocytochemical analysis of spatial organization of photoreceptors and amacrine and ganglion cells in the tiger salamander retina. Visual Neuroscience, 21, 157–166. doi: 10.1017/S0952523804042075 PMID: 15259567

30. Segev R, Puchalla J, Berry MJ 2nd. (2006) Functional organization of ganglion cells in the salamander retina. J Neurophysiol. Apr; 95(4):2277–92. doi: 10.1152/jn.00928.2005 PMID: 16306176

31. Shapley RM, Victor JD (1980). The effect of contrast on the non-linear response of the Y cell. J Physiol. May; 302:535–47. PMID: 7411468

32. Demb JB, Zaghloul K, Haarsma L, Sterling P (2001). Bipolar cells contribute to nonlinear spatial summation in the brisk-transient (Y) ganglion cell in mammalian retina. J Neurosci. Oct 1; 21(19):7447–54. PMID: 11567034

33. Gollisch T, Meister M. (2010). Eye smarter than scientists believed: neural computations in circuits of the retina. Neuron. Jan 28; 65(2):150–64. doi: 10.1016/j.neuron.2009.12.009 PMID: 20152123

34. Schwartz G.W., Okawa H., Dunn F.A., Morgan J.L., Kerschensteiner D., Wong R.O. and Rieke F. (2012). The spatial structure of a nonlinear receptive field. Nat Neurosci. Nov; 12(15):1572–80. doi: 10.1038/nn.3225

35. Chen E., Marre O., Fisher C., Schwartz G., Levy J., da Silveira R., and Berry M.J. II (2013). Alert Response to Motion Onset in the Retina. J Neurosci. Jan 2; 33(1):120–32. doi: 10.1523/JNEUROSCI.3749-12.2013 PMID: 23283327

36. Berry MJ, Warland DK, Meister M (1997). The structure and precision of retinal spike trains. Proc Natl Acad Sci USA. 1997 May 13; 94(10):5411–6. doi: 10.1073/pnas.94.10.5411 PMID: 9144251

37. Peichl L and Wassle H (1979). Size, scatter and coverage of ganglion cell receptive field centres in the cat retina. J Physiol 291: 117–141. doi: 10.1113/jphysiol.1979.sp012803 PMID: 480198

38. Puchalla JL, Schneidman E, Harris RA, Berry MJ. (2005) Redundancy in the population code of the retina. Neuron May 5; 46(3):493–504. doi: 10.1016/j.neuron.2005.03.026 PMID: 15882648

39. Olveczky BP, Baccus SA, Meister M (2003). Segregation of object and background motion in the retina. Nature. May 22; 423(6938):401–8. doi: 10.1038/nature01652 PMID: 12754524

40. Borst A, Theunissen FE. (1999). Information theory and neural coding. Nat Neurosci. Nov; 2(11):947–57. doi: 10.1038/14731 PMID: 10526332

41. Ala-Laurila P, Greschner M, Chichilnisky EJ, Rieke F. Cone photoreceptor contributions to noise and correlations in the retinal output. Nat Neurosci. 2011 Sep 18; 14(10):1309–16. doi: 10.1038/nn.2927 PMID: 21926983

42. Keat J, Reinagel P, Reid RC, Meister M. (2001). Predicting every spike: a model for the responses of visual neurons. Neuron. 2001 Jun; 30(3):803–17. doi: 10.1016/S0896-6273(01)00322-1 PMID: 11430813

43. Pillow JW, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ. (2005). Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. J Neurosci. Nov 23; 25(47):11003–13. doi: 10.1523/JNEUROSCI.3305-05.2005 PMID: 16306413

44. Fairhall AL, Burlingame CA, Narasimhan R, Harris RA, Puchalla JL, Berry MJ 2nd. (2006). Selectivity for multiple stimulus features in retinal ganglion cells. J Neurophysiol. Nov; 96(5):2724–38. doi: 10.1152/jn.00995.2005 PMID: 16914609

45. Hochstein S, Shapley RM. (1976). Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. J Physiol. Nov; 262(2):265–84. doi: 10.1113/jphysiol.1976.sp011595 PMID: 994040

46. Baccus SA, Olveczky BP, Manu M, Meister M. (2008). A retinal circuit that computes object motion. J Neurosci. Jul 2; 28(27):6807–17. doi: 10.1523/JNEUROSCI.4206-07.2008 PMID: 18596156

47. Chen, E. (2013). Motion processing in the retina. PhD thesis, Princeton University.

48. Greschner M, Bongard M, Rujan P, Ammermuller J. (2002) Retinal ganglion cell synchronization by fixational eye movements improves feature estimation. Nat Neurosci. Apr; 5(4):341–7. doi: 10.1038/nn821 PMID: 11914721

49. Radons G, Becker JD, Dulfer B, Kruger J. (1994). Analysis, classification, and coding of multielectrode spike trains with hidden Markov models. Biol Cybern. 1994; 71(4):359–73. doi: 10.1007/BF00239623 PMID: 7948227

50. McIlwain JT (1964). Receptive fields of optic tract axons and lateral geniculate cells: peripheral extent and barbiturate sensitivity. J Neurophysiol. Nov; 27:1154–73. PMID: 14223976

51. Passaglia CL, Enroth-Cugell C, Troy JB. (2001). Effects of remote stimulation on the mean firing rate of cat retinal ganglion cells. J Neurosci. Aug 1; 21(15):5794–803. PMID: 11466451

52. Passaglia CL, Freeman DK, Troy JB. (2009). Effects of remote stimulation on the modulated activity of cat retinal ganglion cells. J Neurosci. Feb 25; 29(8):2467–76. doi: 10.1523/JNEUROSCI.4110-08.2009 PMID: 19244521

53. Takeshita D, Gollisch T (2014). Nonlinear spatial integration in the receptive field surround of retinal ganglion cells. J Neurosci. May 28; 34(22):7548–61. doi: 10.1523/JNEUROSCI.0413-14.2014 PMID: 24872559

54. Geffen MN, de Vries SEJ, Meister M (2007) Retinal ganglion cells can rapidly change polarity from off to on. PLoS Biol 5(3): e65. doi: 10.1371/journal.pbio.0050065 PMID: 17341132

55. Venkataramani S, Taylor WR. (2010). Orientation selectivity in rabbit retinal ganglion cells is mediated by presynaptic inhibition. J Neurosci. 2010 Nov 17; 30(46):15664–76. doi: 10.1523/JNEUROSCI.2081-10.2010 PMID: 21084622

56. Russell TL, Werblin FS. (2010) Retinal synaptic pathways underlying the response of the rabbit local edge detector. J Neurophysiol. May; 103(5):2757–69. doi: 10.1152/jn.00987.2009 PMID: 20457864

57. Manu M, Baccus SA. (2012). Disinhibitory gating of retinal output by transmission from an amacrine cell. Proc Natl Acad Sci USA. Nov 8; 108(45):18447–52. doi: 10.1073/pnas.1107994108

58. van Wyk M, Taylor WR, Vaney DI (2006). Local edge detectors: a substrate for fine spatial vision at low temporal frequencies in rabbit retina. J Neurosci. Dec 20; 26(51):13250–63. doi: 10.1523/JNEUROSCI.1991-06.2006 PMID: 17182775

59. Butts DA, Goldman MS. (2006). Tuning curves, neuronal variability, and sensory coding. PLoS Biol. Apr; 4(4):e92. doi: 10.1371/journal.pbio.0040092 PMID: 16529529

60. Rigotti M, Ben Dayan Rubin D, Wang XJ, Fusi S. (2010). Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. Front Comput Neurosci. 2010 Oct 4; 4:24. doi: 10.3389/fncom.2010.00024 PMID: 21048899

61. DiCarlo JJ, Cox DD. (2007). Untangling invariant object recognition. Trends Cogn Sci. Aug; 11(8):333–41. doi: 10.1016/j.tics.2007.06.010 PMID: 17631409

62. Hung CP, Kreiman G, Poggio T, DiCarlo JJ. (2005). Fast readout of object identity from macaque inferior temporal cortex. Science. Nov 4; 310(5749):863–6. doi: 10.1126/science.1117593 PMID: 16272124

63. Segev R, Goodhouse J, Puchalla J, Berry MJ 2nd. (2004) Recording spikes from a large fraction of the ganglion cells in a retinal patch. Nat Neurosci. Oct; 7(10):1154–61. doi: 10.1038/nn1323 PMID: 15452581

64. Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S (2013). The importance of mixed selectivity in complex cognitive tasks. Nature. May 30; 497(7451):585–90. doi: 10.1038/nature12160 PMID: 23685452

65. Barak O, Rigotti M, Fusi S (2013). The sparseness of mixed selectivity neurons controls the generalization-discrimination trade-off. J Neurosci. Feb 27; 33(9):3844–56. doi: 10.1523/JNEUROSCI.2753-12.2013 PMID: 23447596

66. Observed Brain Dynamics, Partha Mitra and Hemant Bokil, Oxford University Press, New York, 2008.

67. Strong SP, Koberle R, de Ruyter van Steveninck RR, Bialek W (1998). Entropy and information in neural spike trains. Phys Rev Lett 80, 197–200. doi: 10.1103/PhysRevLett.80.197

68. Kim S., Koh K., Lustig M., Boyd S. and Gorinevsky D. An Interior-Point Method for Large-Scale l_1-Regularized Least Squares (2007). IEEE JOURNAL OF SELECTED TOPICS IN SIGNAL PROCESSING, VOL. 1, NO. 4, 606–617. doi: 10.1109/JSTSP.2007.910971